

Signaling without common prior: An experiment*

Michalis Drouvelis[†]
University of York

Wieland Müller[‡]
*CentER, TILEC,
Tilburg University*

Alex Possajennikov[§]
University of Nottingham

April 21, 2009

Abstract

The common prior assumption is pervasive in game-theoretic models with incomplete information. This paper investigates experimentally the importance of inducing a common prior in a two-person signaling game. For a specific probability distribution of the sender's type, the long-run behavior without an induced common prior is shown to be different from the behavior when a common prior is induced, while for other distributions behavior is similar under both regimes. We also present a learning model that allows players to learn about the other players' strategies and the prior distribution of the sender's type. We show that this learning model accurately accounts for all main features of the data.

Keywords: common prior, signaling, experiment, learning

JEL classification Codes: C72, C92, D83

*For helpful comments we thank Robin Cubitt, Henrik Orzen and seminar participants at the 2007 CREED-CeDEX workshop and the 3rd Spain Italy Netherlands meeting on Game theory. Furthermore, we thank the audience of a CeDEX seminar at the University of Nottingham for participating in a simplified pilot experiment. The second author acknowledges financial help from the Netherlands Organisation for Scientific Research (NWO) through a VIDI grant. The third author acknowledges financial support for the experiment from the School of Economics of the University of Nottingham.

[†]Department of Economics and Related Studies, University of York, Heslington, York YO10 5DD, United Kingdom, Email: *md539@york.ac.uk*.

[‡]Department of Economics, Tilburg University, P.O.Box 90153, 5000LE Tilburg, The Netherlands, Email: *w.mueller@uvt.nl*.

[§]School of Economics, University of Nottingham, University Park, Nottingham NG7 2RD, United Kingdom, Email: *alex.possajennikov@nottingham.ac.uk*.

1 Introduction

Game theory usually assumes that all components of a game such as the number of players, strategy sets, payoffs, and the prior beliefs of the players are commonly known. Arguably, the assumption of commonly known game components is an exception rather than the rule in many real-world applications of game theory. This raises the question of what happens if this assumption is violated. If players are uncertain about any specific component of a game, one usually assumes that players at least agree on the set of possible realizations of this component and that a chance move at the beginning of the game assigns an element out of this set of possible components to each player (Harsanyi, 1967). But here, still, it is usually assumed that the probability distribution of this chance move is common knowledge among players. However, what happens if either players do not agree on the set of possible realizations or if no information about the probability distribution of the chance move is available? In this paper we concentrate on the latter question. That is, on the case where the so-called common prior is initially not induced among players.

Surely, the common prior assumption plays an important role in game theory in general and in information economics in particular. Although it is possible to build models without this assumption, it would be difficult, if not impossible, to solve them.¹ Almost all applications we are aware of make this assumption. However, it is not clear whether the assumption is really satisfied in the situations that these applications are trying to describe. For instance, in labour market relationships, it is hard to imagine that managers are aware of the (exact) probability distribution of skilled and unskilled job candidates and it seems likely that often their employment decisions are based on previous learning experience. The same holds for the case of auctions: information about how the valuations of other bidders are distributed will often be unknown or is likely to be very imprecise.

This paper investigates experimentally what happens when a common prior belief about the underlying probability distribution of types of a player is not induced. This is done using a specific example of a two-person signaling game that has two separating equilibria.² The main question the paper tries to answer is whether behavior is different in situations when the common prior belief is induced and when it is not induced. To answer this question, we report the results

¹See Dekel et al. (2004) for a thorough exposition of the issues that arise when players do not only learn about opponents' strategies but also about nature's moves in Bayesian games.

²Note that although prior beliefs do not matter in a separating equilibrium, they matter along the adjustment path in a repeated play of the game. It is the latter what is being analyzed in this paper.

of six treatments that differ in the underlying prior distribution of the sender's type and whether this prior distribution is commonly known or not. In our game, the sender has two possible types and the probability p of the sender being of type 1 is either $1/4$, $1/2$, or $3/4$.

We chose a signaling setup because it allows for a simple situation that reflects the potential importance of having (or learning) correct beliefs. The sequential nature of the signaling game reduces strategic uncertainty about actions and thus allows focusing on consequences of what players know (or learn) about types of other players rather than about their strategies. Signaling applications are common in economic modeling, and thus it is potentially important to extend the analysis to situations where a common prior belief is not established at the start of repeated and random encounters. Various signaling games have been subjected to earlier experimental tests (e.g. Brandts and Holt, 1992, 1993; Banks et al., 1994, Cooper et al., 1997). Thus, there is already a certain amount of knowledge about behavior in such situations that can be built upon.

We conjectured that knowing or not knowing the prior distribution (characterized by p) may change the long-run outcome of this game for one value of p whereas the outcome would not be changed for the other values. This conjecture is based on a "shortcut" version of a learning model suggested by Brandts and Holt (1993, 1996) where players learn the opponent's strategy over time. Indeed, when the prior is known, this learning model predicts that play converges to one equilibrium of our signaling game when p is small ($1/4$) but converges to the other equilibrium when p takes on higher values ($1/2$ or $3/4$). If the prior is not known it is reasonable to assume that subjects start with a uniform prior about the sender's type. In this case we demonstrate that, under specific circumstances, in an adjusted learning model (that also allows learning about the probability of the sender's type) play converges to the equilibrium predicted when the prior p is equal to $1/2$ and known, independent of the true underlying prior. Hence, knowing or not knowing the prior is predicted to lead to different long-run outcomes only when $p = 1/4$.

Our experimental results support this hypothesis. Indeed, we find that in treatments with the prior $p = 1/2$ or $p = 3/4$, behavior appears to be the same independent of whether or not the prior is known, with play clearly converging to the equilibrium predicted by the learning process mentioned above. But when the prior of $p = 1/4$ is not known, we find that although play does not converge to a pure equilibrium (as predicted by the naive learning model), there is a clear difference in comparison to the case of a known prior of $p = 1/4$.

Earlier experimental work by Güth and Ivanova-Stenzel (2003) has shown that knowing or not knowing the distribution of types in specific asymmetric auctions did not lead to significant dif-

ferences in behavior. While this result is remarkable, in their study these authors confine themselves to only reporting the aforementioned result. In our study we try to go beyond the work presented by these authors. First, as mentioned above, we use a simple shortcut version of a naive learning model to derive hypotheses about the question when knowing or not knowing the prior may make a difference in a signaling game. Second, we present results confirming the hypotheses about differences in adjustment patterns for different values of the prior, thereby providing an example where not knowing the prior distribution leads to differences in long-run behavior. Third, we propose an adaptation of the naive learning model of Brandts and Holt (1993, 1996) that allows players to also learn about the prior distribution in case it is not initially induced. More importantly, we show that this learning model tracks the evolution of outcomes in the various treatments quite accurately. In particular, using one set of parameters, we show that the learning model predicts no differences in behavior when the prior is $p = 1/2$ or $p = 3/4$ independent of whether the prior is commonly known or not, but does predict clear differences when the prior is $p = 1/4$.

Despite the importance of learning in games whose components are not commonly known, up to now surprisingly little work has been done in this area. Next to the paper cited above, there is one more related experimental study by Oechssler and Schipper (2003). These authors analyze the question whether subjects can learn the payoff structure of their rival in simple matrix games. Oechssler and Schipper are able to construct the games that subjects appear to be playing (the so-called subjective games) and show that these games often differ from the real games, although subjects come close to playing an equilibrium of the subjective game.³

2 The Signaling Game

Our experiment is based on the following signaling game.

		Receiver				Receiver	
		Type t_1		Type t_2		Type t_2	
			a_1	a_2		a_1	a_2
Sender	m_1	15, 10	80, 80	Sender	m_1	80, 80	15, 30
	m_2	25, 10	50, 50		m_2	50, 50	25, 30

Nature (N) first selects the type of the sender and the sender observes his type: with probability p the type is t_1 and with probability $1 - p$ the type is t_2 . Each type of the sender can then send either

³See also the studies by Camerer et al. (1993) and Costa-Gomes et al. (2001) that try to uncover subjects' reasoning process by using lookup patterns in simple simultaneous-move or sequential-move games.

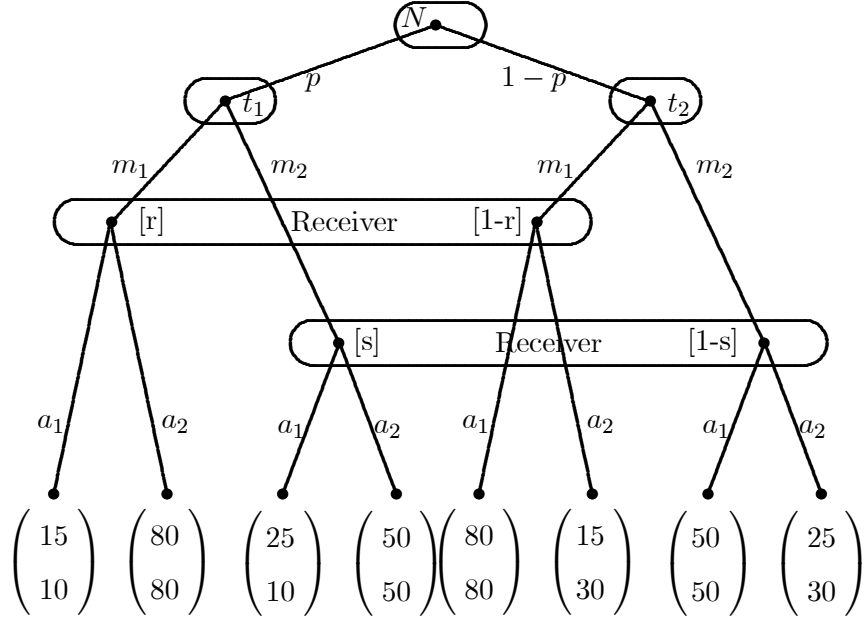


Figure 1: The signaling game used in the experiment

message m_1 or message m_2 . The receiver observes the message but not the type of the sender, and decides between actions a_1 and a_2 . The extensive form of the game is shown in Figure 1. The notation for the beliefs of the receiver is indicated in brackets $[\cdot]$ in her information sets: the receiver's posterior belief that the sender is of type t_1 is r after observing message m_1 and s after observing m_2 . The payoffs of the sender (receiver) at terminal nodes are the upper (lower) entries in the payoff vectors.

In the experiment we implemented the priors $p \in \{1/4, 1/2, 3/4\}$. When p is commonly known, weak perfect Bayesian equilibria have the following properties. First, for any value of p the game has two pure separating equilibria: $[(m_1, m_2), (a_1, a_2), r = 1, s = 0]$ and $[(m_2, m_1), (a_2, a_1), r = 0, s = 1]$. (Here, the first entry in the sender's strategy is the message chosen as type t_1 , and the second entry is the message chosen as t_2 . Analogously, the first entry in the receiver's strategy is the action chosen after observing message m_1 , and the second entry is the action chosen after message m_2 .) Second, the game has no pooling equilibria. Third, for each of the values of p mentioned above, the game has a hybrid equilibrium.⁴

When p is not commonly known, Bayesian equilibrium analysis is not possible. We use instead self-confirming equilibrium (Dekel et al., 2004) that allows players to have incorrect beliefs

⁴The game with $p = 1/4$ has the hybrid equilibrium $[(m_1, \frac{7}{15}m_1 + \frac{8}{15}m_2), (\frac{7}{13}a_1 + \frac{6}{13}a_2, a_1), r = \frac{5}{12}, s = 0]$, the game with $p = 1/2$ has the hybrid equilibrium $[(\frac{5}{7}m_1 + \frac{2}{7}m_2, m_1), (\frac{6}{13}a_1 + \frac{7}{13}a_2, a_2), r = \frac{5}{12}, s = 1]$, and the game with $p = 3/4$ has the hybrid equilibrium $[(\frac{5}{21}m_1 + \frac{16}{21}m_2, m_1), (\frac{6}{13}a_1 + \frac{7}{13}a_2, a_2), r = \frac{5}{9}, s = 1]$.

about the distribution of types and/or about the strategies of the other players. However, since in our experiment players receive feedback on realizations of types and on actions, self-confirming equilibria coincide with Nash equilibria.

Since static equilibrium concepts do not address the question of where beliefs come from, we look also at a dynamic learning model. Brandts and Holt (1993, 1996) suggest that learning in a repeated signaling game like the one in Figure 1 could be based on a naive reasoning process. We describe a version of such a learning process that combines “naive” initial beliefs with best reply behavior. The signaling game in Figure 1 was constructed in such a way that for a small probability p of type 1 the learning process *could* lead to a different outcome depending on whether the prior is commonly known or not; and that the outcome of the learning process *could* be the same independent of the knowledge of the prior for larger values of p .

Suppose the sender initially believes that the receiver will choose any of his two actions with equal probability independent of the message sent. Then both types of the sender will initially send message m_1 since, given the initial belief about the receiver’s reaction, the expected payoff of sending m_1 (47.5) is higher than the expected payoff of sending m_2 (37.5) for both types of the sender.

When the prior is $p = 1/4$ and common knowledge, the best response of the receiver after observing any of the two messages is to choose action a_1 . (The receiver’s expected payoff after observing m_1 is 62.5 for a_1 and 42.5 for a_2 , and the receiver’s expected payoff after observing m_2 is 40 for a_1 and 35 for a_2). So, initial play would be $[(m_1, m_1), (a_1, a_1)]$. Given the receiver’s initial reaction, the best response for type t_1 of the sender would be m_2 while the best response for type t_2 remains unchanged. Given this strategy (m_2, m_1) of the sender, the receiver’s best response will still be a_1 after observing message m_1 , but a_2 after observing message m_2 . These two strategies together— (m_2, m_1) for the sender and (a_1, a_2) for the receiver—constitute equilibrium behavior.

When $p = 1/2$, the best response of the receiver against the initial strategy of the sender mentioned above is action a_2 after m_1 (expected payoffs are 45 for a_1 and 55 for a_2) and a_2 after m_2 (expected payoffs are 30 for a_1 and 40 for a_2). So, in this case initial play would be $[(m_1, m_1), (a_2, a_2)]$. Given the receiver’s initial reaction, the best response for type t_1 of the sender remains unchanged while the best response of type t_2 of the sender would be m_2 (change in payoff from 15 to 25). Given this strategy (m_1, m_2) of the sender, the receiver’s best response will be a_2 after observing message m_1 , and a_1 after observing message m_2 . The two strategies together— (m_1, m_2) for the sender and (a_2, a_1) for the receiver—constitute, again, equilibrium behavior. When the commonly known prior is $p = 3/4$, the dynamic is the same as in case of $p = 1/2$. We summarize

	Prior commonly known		Prior not commonly known	
	Sender	Receiver	Sender	Receiver
$p = 1/4$	(m_2, m_1)	(a_1, a_2)	(m_1, m_2)	(a_2, a_1)
$p = 1/2$	(m_1, m_2)	(a_2, a_1)	(m_1, m_2)	(a_2, a_1)
$p = 3/4$	(m_1, m_2)	(a_2, a_1)	(m_1, m_2)	(a_2, a_1)

Note: The first entry in the sender’s strategy is the message chosen as type t_1 and the second entry is the message chosen as t_2 . The first entry in the receiver’s strategy is the action chosen after observing message m_1 and the second entry is the action chosen after m_2 .

Table 1: Predictions regarding long-run behavior based on the naive learning process.

these predictions in the second and third columns of Table 1.

When the prior probability p is not commonly known, it is perhaps reasonable to assume that subjective beliefs may be (close to) $p = 1/2$. Then initial play should be as described above when $p = 1/2$. Further adjustment depends on the speed of belief revision about the prior⁵ and about the opponent’s strategy. If players revise beliefs about p faster than beliefs about the opponent’s strategy, the receiver may switch to (a_1, a_1) and play may be attracted to the equilibrium $[(m_2, m_1), (a_1, a_2)]$. On the other hand, if beliefs about the opponent’s strategy are revised faster than beliefs about the probability p , players may continue playing as if $p = 1/2$ for long enough to get attracted to the equilibrium $[(m_1, m_2), (a_2, a_1)]$. Hence, it is possible that play in the above game will converge to the equilibrium $[(m_1, m_2), (a_2, a_1)]$ when the probability p is not commonly known. This is indicated in the last two columns in Table 1. Together with the predictions when p is commonly known, this implies that we expect long-run behavior to be quite similar when $p = 1/2$ or $p = 3/4$ independently of whether the probability is commonly known or not. On the other hand, we expect that long-run behavior will be different depending on the prior probability being commonly known or not when $p = 1/4$.

We want to emphasize that we view the above analysis as a “shortcut” version of a more elaborate learning model that we will introduce in Section 5.⁶ Of course, we do not expect subjects

⁵Similarly to the learning about strategies, with respect to learning about the unknown prior we assume that, starting with the flat prior of $p = 1/2$, players update this prior in each round given the feedback about the selected type.

⁶In this formal learning model we will spell out explicitly how the updating of beliefs about strategies and the

to converge to one of the equilibria within the first few periods. After all, subjects do not have to indicate an entire strategy, but only react at one information set at a time. Also, so far we assumed that subjects always choose a best reply given their beliefs without committing any decision errors. This assumption will be relaxed in Section 5 below. Here, we only wanted to show that a reasonable adjustment theory gives rise to the hypothesis that for some values of p it might make a difference whether or not players know the prior.

The game with $p = 1/2$ is included in our experiment because this value is a natural first guess when p is not known. We do not expect that the behavior would depend on whether $p = 1/2$ is commonly known or not; treatments with $p = 1/2$ rather serve as a benchmark. We use other values of p on both sides of $1/2$ for completeness. If play in the game in Figure 1 is following the “naive” dynamic described above, the game with value $p = 1/4$ can be an example where knowing the prior distribution matters in the long-run, while the game with value $p = 3/4$ is another example where it does not matter.

The payoffs in the game were primarily chosen in such a way that for some values of p the learning process could lead to a different outcome depending on whether or not p is commonly known, and that for other values of p the outcome is the same independently of the knowledge of the prior.

In an effort to weaken their possible confounding influence on adjustment patterns, we also had an eye on various behavioral effects when designing the game. These were the similarity of payoffs of the two players and of the two sender types (to avoid relative payoff effects), the similarity of payoffs in the two long-run outcomes (to lessen the focality effect of the efficient outcome), small conflict between expected payoff maximization and maximin (to diminish risk aversion and ambiguity aversion effects), and the relative steepness of the best response (to give enough incentive to find it). Further desirable features we considered were that all messages are sent during the adjustment period (to give subjects experience in all information sets of the game), and that the dynamic is relatively fast (so that the convergence to a long-run outcome can occur in a reasonable period of time).

prior works and what we mean when we speak, as above, about the speed of belief revision.

3 Experimental design and hypotheses

Experimental design: Our experiments are based on a 3×2 factorial design. It includes treatments with three different values of the prior, namely $p \in \{1/4, 1/2, 3/4\}$ and it includes treatments with and without a commonly known prior. A treatment is denoted $K-p$ if the prior probability $p \in \{1/4, 1/2, 3/4\}$ of type 1 of the sender is commonly known and $N-p$ if the prior is not known. Hence, the six treatments are referred to as $K-1/4$, $K-1/2$, $K-3/4$, $N-1/4$, $N-1/2$, and $N-3/4$.

In the instructions, senders were referred to as “A-participants” and receivers as “B-participants”. Subjects were informed that A-participants could be of two types (“Type 1” and “Type 2”) and that at the beginning of each round, a random draw would determine the type of an A-participant. For the treatments with unknown prior, the instructions contained the following sentences: “The random draw is such that with a $X\%$ chance the A-participant will be of Type 1, and with a $(100 - X)\%$ chance of Type 2. You receive no information about the value of X , except that X is constant over all rounds of the experiment.” In the treatments with a commonly known prior, the latter sentence was omitted and X and $100 - X$ were explicitly given by either 75 and 25, 50 and 50, or 25 and 75. In all treatments subjects were informed that after the random draw the A-participant would be informed about his type while the B-participant would not be informed about the type of the A-participant.

The sender’s messages and the receiver’s actions were generally referred to as “decisions” and labelled “ C ” and “ D ” (for messages) and “ E ” and “ F ” (for actions). Each session consisted of 40 rounds to give subjects room for learning. Payoffs were given to senders in the form of two tables corresponding to the two types. Receivers saw payoff tables corresponding to different messages as this makes it clearer what they know and what they do not know. See Appendix A for further details of the instructions.

Given that our research topic concerns beliefs, we could have asked questions with regard to which beliefs players actually entertain during the game and whether the true distribution of types is eventually learnt. However, in this study we have chosen to limit ourselves to the analysis of behavior. The game appeared to us relatively complex, and properly eliciting beliefs may have created excessive demands on the cognitive effort of subjects, perhaps making behavior unnecessarily noisy.

The experiment was computerized using the *z-Tree* software (Fischbacher, 2007) and was conducted in March 2007 at the Centre for Decision Research and Experimental Economics (CeDEx)

Treatment	Value of the Prior	Prior	No. of sessions	No. of independent matching groups	No. of subjects
$K-1/4$	$1/4$	known	3	6	$6 \times 8 = 48$
$N-1/4$	$1/4$	not known	3	6	$6 \times 8 = 48$
$K-1/2$	$1/2$	known	2	4	$4 \times 8 = 32$
$N-1/2$	$1/2$	not known	2	4	$4 \times 8 = 32$
$K-3/4$	$3/4$	known	2	4	$4 \times 8 = 32$
$N-3/4$	$3/4$	not known	2	4	$4 \times 8 = 32$

Table 2: Overview of the experimental design

laboratory at the School of Economics of the University of Nottingham, United Kingdom. The recruitment of the subjects from the pool of University of Nottingham students registered with CeDEx for experiments was done via the *ORSEE* software (Greiner, 2004). Subjects were students from various fields of studies including economics.

There were three sessions each for treatments $K-1/4$ and $N-1/4$ (since this is the comparison we were mostly interested in), and two sessions for each of the other four treatments. Each session consisted of 16 subjects who were divided into two matching groups of eight subjects each (four senders and four receivers). The design is summarized in Table 2.

We randomly rematched subjects in every period within each matching group in order to create an environment as close as possible to a single-period interaction between subjects. The same matching protocol was used in all individual matching groups. Player roles were randomly assigned at the beginning of the experiment and were then kept fixed for the entire experiment. The assignment of types to senders in a matching group was random in each of the 40 rounds of the experiment but was done prior to the experiments for each value of p . This assignment of types to senders was then used for all matching groups in treatments with this value of p .

After each period, each subject observed the type of the sender, the message sent by the sender, the action taken by the receiver, and the realized payoffs in his/her own match. This is needed to be able to update beliefs. Some parts of this information can be inferred from other parts but we decided to give all of this information explicitly to facilitate learning. We decided not to provide the whole history of play up to the current period because our predictions are based on a short-memory best-reply adjustment process. Subjects were allowed to take notes (though almost nobody did).

Sessions lasted approximately 90 minutes. Subjects were paid according to accumulated earnings over all rounds. With the points converted to money at the rate of £0.05 for 10 points, the average payments were £11.73 (\$22.56) per subject.

Hypotheses: We want to compare behavior when the prior probability is unknown with behavior when this probability is commonly known. Based on the simple learning process described above, we expect long-run behavior to be quite similar when the prior probability is $p = 1/2$ or $p = 3/4$ independently of whether or not the value of p is commonly known. On the other hand, when the probability is $p = 1/4$ we expect the long-run behavior to be different depending on whether the value of p is commonly known or not (see Table 1). Based on the “naive” reasoning assumed in the learning process, we should also see short-run behavior in treatments with an unknown value of p to be similar to behavior when $p = 1/2$ is known.

4 Results

4.1 A First Look

We first look at the time series of senders’ and receivers’ strategies, averaged across sessions and subjects of the same treatment. Figures 2, 3, and 4 show data for the treatments with prior $p = 1/4$, $p = 1/2$, and $p = 3/4$. The data are grouped in blocks of 5 periods. The figures show the relative frequency of decisions made. More precisely, the two top panels in each figure show the relative frequency of message- m_1 choices for each of the two types of the sender, whereas the two bottom panels in each figure show the relative frequency of action- a_1 choices after each of the two possible messages.

Consider first Figure 2 which shows behavior in the treatments with prior $p = 1/4$. Concentrate first on behavior in treatment $K-1/4$ where the prior is known. In this case we see clear convergence towards the equilibrium $[(m_2, m_1), (a_1, a_2)]$ as predicted in Table 1. That is, senders of type 2 quickly learn to send message m_1 while receivers learn to respond to it with action a_1 . Senders of type 1 learn to send message m_2 , after which receivers react mostly with a_2 . Next, we observe that there is a clear gap between the graphs representing behavior in treatments $K-1/4$ and $N-1/4$. Furthermore, the graphs representing behavior in the two treatments, by and large, run in parallel (with the exception of the early behavior of receivers after observing message m_2). Finally, while behavior in treatment $K-1/4$ converges to a pure equilibrium, this does not seem to

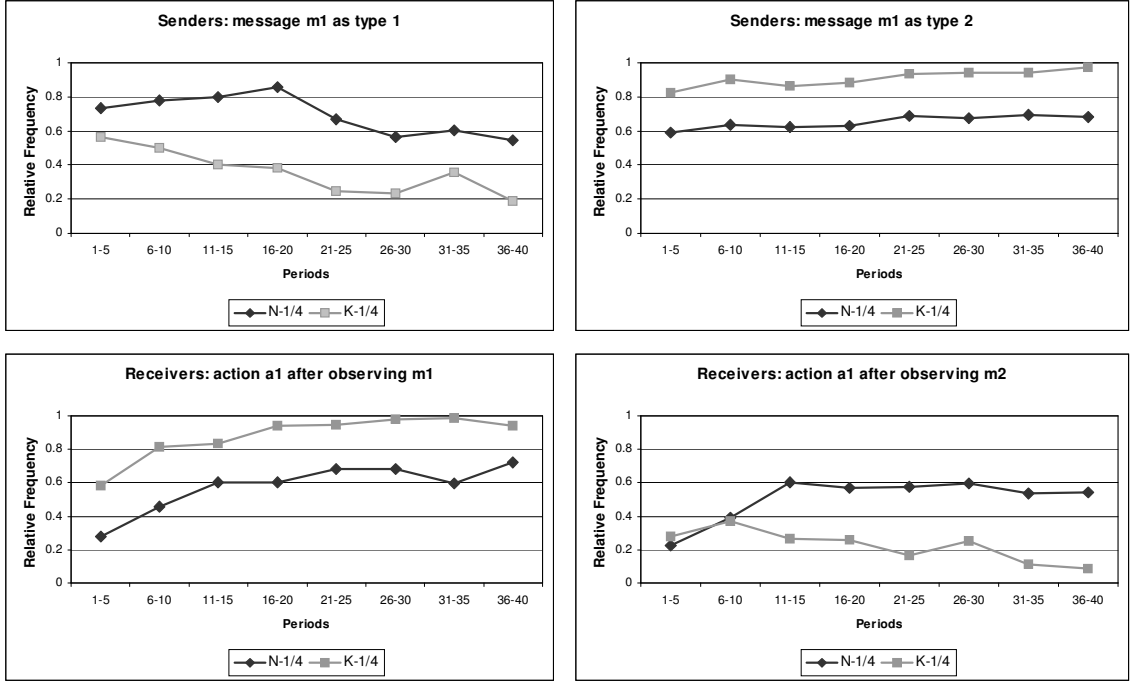


Figure 2: Time series of average strategies in treatments $N-1/4$ and $K-1/4$.

be the case in treatment $N-1/4$.⁷

Let us now consider behavior in the treatments with prior $p = 1/2$ and $p = 3/4$ as shown in Figures 3 and 4. The results in these two cases are very similar and clear-cut. Inspecting the graphs in each of the four panels of these two figures reveals that there is not much difference between behavior in treatments $N-x$ and $K-x$ for $x \in \{1/2, 3/4\}$. Also comparing the panels between the two figures, there does not seem to be much difference across all four treatments. In particular, when the sender is of type 1, almost exclusively m_1 is played and this message is answered by action a_2 . When the sender is of type 2, play converges to the senders choosing message m_2 which the receivers answer by mostly choosing action a_1 . Thus, we conclude that in all treatments with $p = 1/2$ or $p = 3/4$ play converges to the equilibrium $[(m_1, m_2), (a_2, a_1)]$ which is what we predicted in Section 2 (see Table 1).

Summarizing, Figures 2-4 suggest the following. First, in treatments with prior $p = 1/2$ or

⁷Recall from footnote 4, that the game with $p = 1/4$ has the hybrid equilibrium $[(m_1, \frac{7}{15}m_1 + \frac{8}{15}m_2), (\frac{7}{13}a_1 + \frac{6}{13}a_2, a_1), r = \frac{5}{12}, s = 0]$ and the game with $p = 1/2$ has equilibrium $[(\frac{5}{7}m_1 + \frac{2}{7}m_2, m_1), (\frac{6}{13}a_1 + \frac{7}{13}a_2, a_2), r = \frac{5}{12}, s = 1]$. In these equilibria either type 1 or type 2 of the sender chooses one of the messages with certainty and the receiver chooses one of the actions with certainty after observing message m_2 . Figure 2 clearly suggests that this is not what play converges to over time in treatment $N-1/4$.

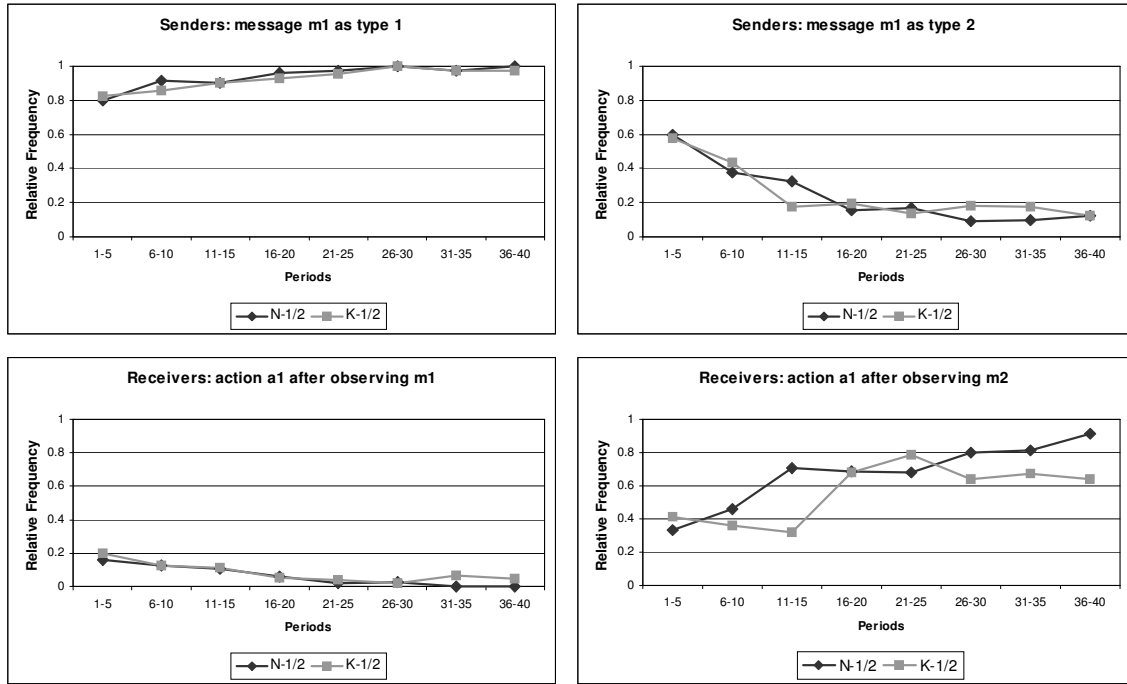


Figure 3: Time series of strategies in treatments $N-1/2$ and $K-1/2$.

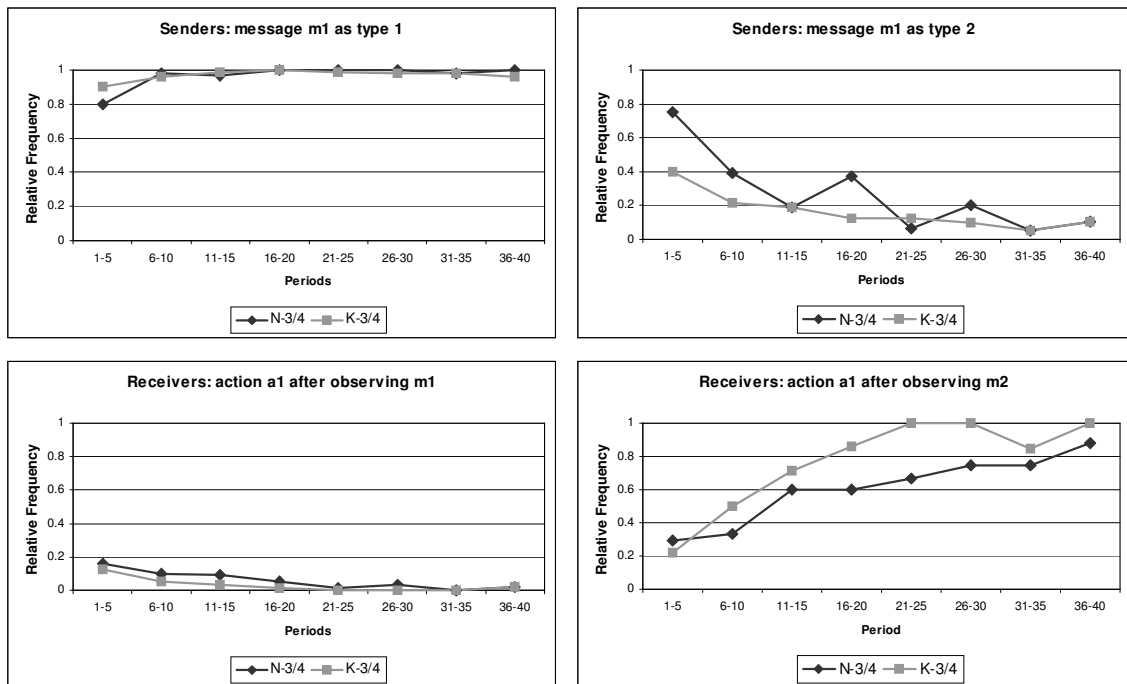


Figure 4: Time series of strategies in treatments $N-3/4$ and $K-3/4$.

$p = 3/4$ behavior appears to be the same independent of whether the prior is known or not, with play clearly converging to the equilibrium predicted by the learning process illustrated in Section 2. Second, in the treatment with known prior $p = 1/4$, play clearly converges to the equilibrium predicted by the naive learning process. However, when the prior $p = 1/4$ is not known, play does not converge to a pure equilibrium and there is a clear difference in comparison to the case of a known prior of $p = 1/4$. Hence, this first graphical inspection of the data supports our hypotheses stated at the end of Section 2.

Another observation, common to all treatments, is that the time series are rather constant in the late periods, while there are discernible trends in the early ones. We will consider the relatively stable behavior in the second half of the treatments as the basis for the statistical tests for long-run behavior.

4.2 A formal analysis

Let us now check the observation from the graphs more formally. We first ask for each of the priors $p \in \{1/4, 1/2, 3/4\}$ separately, whether there is a difference in behavior when the prior is known and not known.

We use the Wilcoxon-Mann-Whitney non-parametric test (Siegel and Castellan, 1988, Ch. 6) using matching groups as independent observations. This gives 6 observations for the treatments $N-1/4$ and $K-1/4$ and 4 observations for each of the other treatments. We use periods 21-40 for the main tests, checking the robustness of the results for other choices of periods.

To test whether, conditional on type realization, senders behave differently when the prior is known or not, we calculate the proportions of times senders chose message m_1 within each matching group. Similarly, to test whether, conditional on the message observed, receivers behave differently when the prior is known or not, we calculate the proportion of times receivers chose action a_1 .

The results are presented in Table 3. The tests presented in the table reject the one-sided null hypotheses of a difference between the strategy proportions in treatments $N-1/4$ and $K-1/4$ at the 5% significance level. Therefore there is a difference between the treatments in the direction predicted by the theory in Section 2. On the other hand, the two-sided null hypotheses of no difference between treatments $N-1/2$ and $K-1/2$, and between treatments $N-3/4$ and $K-3/4$ are not rejected. The results are similar when tests are done on the data from all periods, or on the data from the last block (periods 36-40) only.

Since the power of the test is low when there are only four observations per treatment, we

Proportions of strategies for Periods 21-40 and comparison tests				
	Senders		Receivers	
	$m_1 t_1$	$m_1 t_2$	$a_1 m_1$	$a_1 m_2$
$N-1/4$ vs $K-1/4$	0.59 vs 0.26 (0.032)**	0.68 vs 0.95 (0.023)**	0.67 vs 0.96 (0.018)**	0.56 vs 0.15 (0.019)**
$N-1/2$ vs $K-1/2$	0.99 vs 0.98 (0.495)	0.12 vs 0.16 (0.237)	0.01 vs 0.04 (0.281)	0.80 vs 0.68 (0.149)
$N-3/4$ vs $K-3/4$	1.00 vs 0.98 (0.850)	0.11 vs 0.10 (0.653)	0.02 vs 0.00 (0.166)	0.78 vs 0.96 (0.237)

Note: p -values in parentheses. $N-1/4$ vs $K-1/4$: 6 observations for each treatment,

$H_0 : Prop_{N-1/4} \leq Prop_{K-1/4}$ for $m_1|t_1$ and $a_1|m_2$; $H_0 : Prop_{N-1/4} \geq Prop_{K-1/4}$ for $m_1|t_2$ and $a_1|m_1$.

$N-1/2$ vs $K-1/2$ and $N-3/4$ vs $K-3/4$: 4 observations for each treatment, for all tests

$H_0 : Prop_{N-x} = Prop_{K-x}$, where $x = 1/2, 3/4$. ** $p < 0.05$.

Table 3: Results of tests regarding differences in behavior across treatments in Periods 21-40.

also ran subject-specific random-effect probit regressions to compare strategies across treatments. These regressions confirm the difference between treatments $N-1/4$ and $K-1/4$ and also do not find a significant difference between treatments with known and unknown prior when $p \neq 1/4$ (except in $N-3/4$ vs $K-3/4$ in the answer of receivers to the less frequent message m_2). The regressions also corroborate that the behavior in the second half of periods is stable as there are no significant time trends.

Result 1 *There is a significant difference in behavior between treatments with known and unknown prior when $p = 1/4$ in the directions predicted by Table 1 in Section 2. There are no significant differences in behavior between treatments with known and unknown prior when $p \neq 1/4$.*

The predictions of our simple learning model in Section 2 (as well as the more elaborate version presented in the next section) are based on the assumption that subjects in treatments with an unknown prior entertain the initial belief that each of the sender's type is equally likely. A comparison of behavior between treatments $N-p$ and $K-1/2$ in the early rounds (rounds 1-5) will show whether or not this assumption is justified. Note, though, that testing for possible differences in *behavior* only provides an indirect test of the hypothesis that the subjective prior maintained in

Proportions of strategies in Periods 1-5 and comparison tests				
	Senders		Receivers	
	$m_1 \mid t_1$	$m_1 \mid t_2$	$a_1 \mid m_1$	$a_1 \mid m_2$
$N-1/4$ vs $K-1/2$	0.73 vs 0.83 (0.590)	0.59 vs 0.58 (0.745)	0.28 vs 0.20 (0.398)	0.22 vs 0.42 (0.392)
$N-1/2$ vs $K-1/2$	0.80 vs 0.83 (0.505)	0.60 vs 0.58 (0.769)	0.16 vs 0.20 (0.885)	0.33 vs 0.42 (1.000)
$N-3/4$ vs $K-1/2$	0.80 vs 0.83 (0.772)	0.75 vs 0.58 (0.186)	0.16 vs 0.20 (0.773)	0.29 vs 0.42 (1.000)

Note: p -values in parentheses. $H_0 : Prop_{N-x} = Prop_{K-1/2}$, where $x = 1/4, 1/2, 3/4$.

Table 4: Results of tests regarding differences in behavior in Periods 1-5

the initial periods of the treatments with an unknown prior is the same as the prior in treatment $K-1/2$. However, since we did not explicitly elicit beliefs about the prior, this indirect method is the only way to shed light on the issue of what subjects' initial beliefs are when the prior is not known.

The tests on the behavior in the first five periods of the experiments compare, separately for senders and receivers, proportions of actions used in the treatments with an unknown prior with those used in the treatment with known prior $p = 1/2$. The results of the Wilcoxon-Mann-Whitney tests are presented in Table 4. The test results show that we cannot reject the hypothesis that behavior in the early periods of the treatments with unknown prior is the same as behavior in treatment $K-1/2$. This in turn implies that we cannot reject the hypothesis that subjects in the treatments with an unknown prior entertain beliefs about the prior which are similar to the ones in treatment $K-1/2$. The results are also corroborated by the subject-specific random-effects probit regressions.

Result 2 *There is no significant difference in the behavior in early periods between treatments with an unknown prior and the treatment with a known prior $p = 1/2$.*

5 Simulation of the learning model

In this section we present the results of simulations which are based on an extension of a learning model first proposed in Brandts and Holt (1993, 1996).⁸ In the model, players learn about the strategies of the other players; the extension allows the players to learn also about the prior distribution in the case when it is not initially induced.

The formal description of the learning model given below is tedious, but the main idea is simple. Players initially have a flat prior on the other player's strategies and choose a noisy best response given their belief about those strategies. In later rounds they update their beliefs in view of experience and continue choosing noisy myopic best responses. In treatments in which players do not know the probability of chance move, they initially have a flat prior on it and update this prior in each round given the feedback about the selected type. In treatments with a known probability of nature's move, players just learn over time about the other player's strategy. In treatments with an unknown probability of nature's move, they also learn about this probability. In what follows, we provide a formal description of the learning model.

Let us first concentrate on receivers. Let $A^S(t_i, m_j, t)$ be the "belief propensity" that type t_i sends message m_j in period t , and let $A^S(t_i, t) = A^S(t_i, m_1, t) + A^S(t_i, m_2, t)$. The belief at time t that type t_i sends message m_j is then $\Pr(m_j|t_i)_t = A^S(t_i, m_j, t)/A^S(t_i, t)$. Let $A^S(t_i, 1)$ be a given parameter⁹ of the model (we assume $A^S(t_1, 1) = A^S(t_2, 1)$) and let $A^S(t_i, m_j, 1) = A^S(t_i, 1)/2$. Note that this implies $\Pr(m_j|t_i)_1 = 1/2$ such that receivers believe in period 1 that each type of the sender chooses each of the possible messages with equal probability. Updating proceeds as follows: suppose that a receiver observes that type t_i of the sender chose message m_j at time t , then $A^S(t_i, m_j, t+1) = A^S(t_i, m_j, t) + 1$, $A^S(t_i, m_{-j}, t+1) = A^S(t_i, m_{-j}, t)$, and $A^S(t_{-i}, m_j, t+1) = A^S(t_{-i}, m_j, t)$ for $j = 1, 2$. That is, only the belief propensity of type t_i who sent message m_j in period t is increased by 1, whereas all other belief propensities (including the ones of the other type of the sender) remain unchanged.

Next turn to senders. Let $A^R(m_j, a_k, t)$ be the belief propensity that message m_j is answered by action a_k in period t , and let $A^R(m_j, t) = A^R(m_j, a_1, t) + A^R(m_j, a_2, t)$. The belief that message

⁸Brandts and Holt (1992, 1993, 1996), Cooper et al. (1997), and Anderson and Camerer (2000) discuss different adjustment theories ("naive" dynamic, fictitious play, EWA) for experimental signaling games. However, all of these theories are based on belief revision in view of experience and on best response, thus their long-run predictions are likely to be the similar, although the speed of adjustment may differ.

⁹ $A^S(t_i, 1)$ corresponds to the parameter α in Brandts and Holt (1996).

m_j is answered by action a_k at time t is then $\Pr(a_k|m_j)_t = A^R(m_j, a_k, t)/A^R(m_j, t)$. Let $A^R(m_j, 1)$ be given (we will assume that $A^R(m_1, 1) = A^R(m_2, 1) = A^S(t_i, 1)$), and let $A^R(m_j, a_k, 1) = A^R(m_j, 1)/2$. Note that this implies $\Pr(a_k|m_j)_1 = 1/2$ so that senders believe in period 1 that a receiver will choose each of his two actions with probability $1/2$ in response to either message sent by the sender. Updating proceeds as follows: suppose a sender observes that message m_j is answered by action a_k at time t . Then $A^R(m_j, a_k, t+1) = A^R(m_j, a_k, t) + 1$, $A^R(m_j, a_{-k}, t+1) = A^R(m_j, a_{-k}, t)$, and $A^R(m_{-j}, a_k, t+1) = A^R(m_{-j}, a_k, t)$ for $k = 1, 2$. That is, only the belief propensity of the action that was chosen in response to the actual message sent in period t is increased by 1, whereas all other belief propensities (including the ones of the message not sent) remain unchanged.

We now turn to the description of how agents learn about the chance move selecting the sender's type. Let $A^T(t_i, t)$ be the belief propensity that Type t_i is the outcome of the chance move at time t , and let $A^T(t) = A^T(t_1, t) + A^T(t_2, t)$. The belief that type t_i is selected in period t is then $\Pr(t_i)_t = A^T(t_i, t)/A^T(t)$. Let $A^T(1)$ be given. If the probability p of type 1 is known, then $A^T(t_1, t) = p \cdot A^T(t)$ and $A^T(t_2, t) = (1 - p) \cdot A^T(t)$. This implies $\Pr(t_1) = pA^T(t)/A^T(t) = p$, and $A^T(t_i, t)$ is not updated. If the probability of type 1 is not known, then $A^T(t_1, t) = A^T(1)/2$. Note that this implies $\Pr(t_i)_1 = 1/2$ such that agents start out with a flat prior assigning equal probability to either type having been selected. Updating proceeds as follows: when a player observes that type t_i was realized at time t , then $A^T(t_i, t+1) = A^T(t_i, t) + 1$, $A^T(t_{-i}, t+1) = A^T(t_{-i}, t)$.

Given the beliefs, players calculate expected payoffs. For senders, $E^S[m_j|t_i] = S(t_i, m_j, a_1) \cdot \Pr(a_1|m_j) + S(t_i, m_j, a_2) \cdot \Pr(a_2|m_j)$ where $S(t_i, m_j, a_k)$ is the payoff of the sender when his type is t_i , the message sent is m_j , and action chosen by the receiver is a_k . For receivers, $E^R[a_k|m_j] = R(t_1, m_j, a_k) \cdot \Pr(t_1|m_j) + R(t_2, m_j, a_k) \cdot \Pr(t_2|m_j)$, where $R(t_i, m_j, a_k)$ is the payoff of the receiver when the sender's type is t_i , the message sent is m_j , action chosen by the receiver is a_k , and

$$\Pr(t_i|m_j) = \frac{\Pr(m_j|t_i) \Pr(t_i)}{\Pr(m_j|t_i) \Pr(t_i) + \Pr(m_j|t_{-i}) \Pr(t_{-i})}$$

is the posterior probability (using Bayes' rule) that the sender is of type t_i given that message m_j was sent.

We assume that players make (logistic) decision errors when choosing their actions. More precisely,

$$\begin{aligned} \Pr(m_j|t_i) &= \frac{\exp(\lambda \cdot E^S[m_j|t_i])}{\exp(\lambda \cdot E^S[m_j|t_i]) + \exp(\lambda \cdot E^S[m_{-j}|t_i])}, \\ \Pr(a_k|m_j) &= \frac{\exp(\lambda \cdot E^R[a_k|m_j])}{\exp(\lambda \cdot E^R[a_k|m_j]) + \exp(\lambda \cdot E^R[a_{-k}|m_j])}, \end{aligned}$$

where $\Pr(m_j|t_i)$ is a sender's probability to play message m_j when he is of type t_i and $\Pr(a_k|m_j)$ is a receiver's probability to play action a_k after observing message m_j . Here, λ is an error parameter that represents the strength with which agents choose a best response. If $\lambda \rightarrow 0$, agents choose their actions randomly, while if $\lambda \rightarrow \infty$, agents choose the action with the highest expected payoff with certainty.

The parameters of our learning model are the strength of the initial propensities on strategies $A^S(t_i, 1) = A^R(m_j, 1) = A$, the strength of the initial propensities on types $A^T(1)$ (in treatments with an unknown prior), and the strength of the best response λ .

Mimicking the procedures used in our experiments, the simulation model was run over 40 periods with eight simulated agents (four agents in each of the two player roles), using the same matching protocol as in the actual experimental sessions. Furthermore, the same outcomes of the chance moves selecting senders' types as in the experimental sessions were used in the simulations. To smoothen the effects of randomness, we use the average outcomes of 100 runs of the simulation model for each treatment and compare the simulated outcomes with the actual data. In each period, there are eight proportions of observed actual outcomes $p_{obs}(m_j, a_k|t_i)$ and eight proportions of simulated outcomes $p_{sim}(m_j, a_k|t_i)$. As a goodness-of-fit measure, for each treatment we use the metric $\sum_{\tau=1}^{40} (p_{obs}(m_j, a_k|t_i) - p_{sim}(m_j, a_k|t_i))^2$. We add up the metrics for all treatments and look for the set of parameters A , $A^T(1)$, and λ that give the best fit to the observed data.

To find the best-fitting parameters, we did a grid search. After a few initial trials to locate the likely values of parameters, the grid search was concentrated on the interval $[0.5, 5]$ with step size 0.5 for A , the interval $[11, 20]$ with step size 1 for $A^T(1)$, and the interval $[0.05, 0.25]$ with step size 0.05 for λ . In Figures 5-7 we present the results of simulations that use the set of parameters $A = 2$, $A^T(1) = 15$, and $\lambda = 0.15$ that give the best overall fit (that is, the best fit over all six treatments). These Figures also present the observed behavior and are organized in the same way as Figures 2-4 above.

Let us first consider the results of the simulations for the treatments with prior $p = 1/4$ which are shown in Figure 5. First of all, recall from our analysis in Section 4 that the main feature of the observed data is that there is a clear gap between shares of senders' and receivers' choices depending on whether or not the prior $p = 1/4$ is known. Figure 5 shows that the learning model accounts for this observation as it predicts the same differences in choice behavior in all subpanels of Figure 5. Next to this main observation, a few more features of the comparison of observed and simulated data are as follows. First, with regard to senders' behavior we observe that the learning

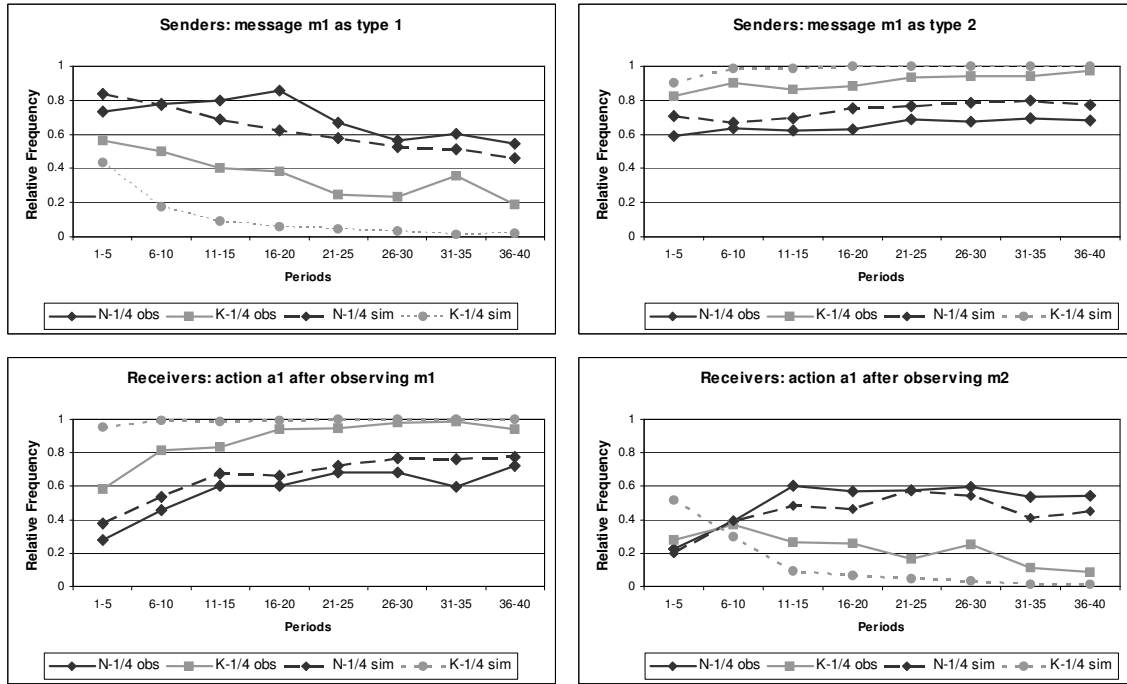


Figure 5: Observed and simulated time series of average strategies in treatments $N-1/4$ and $K-1/4$.

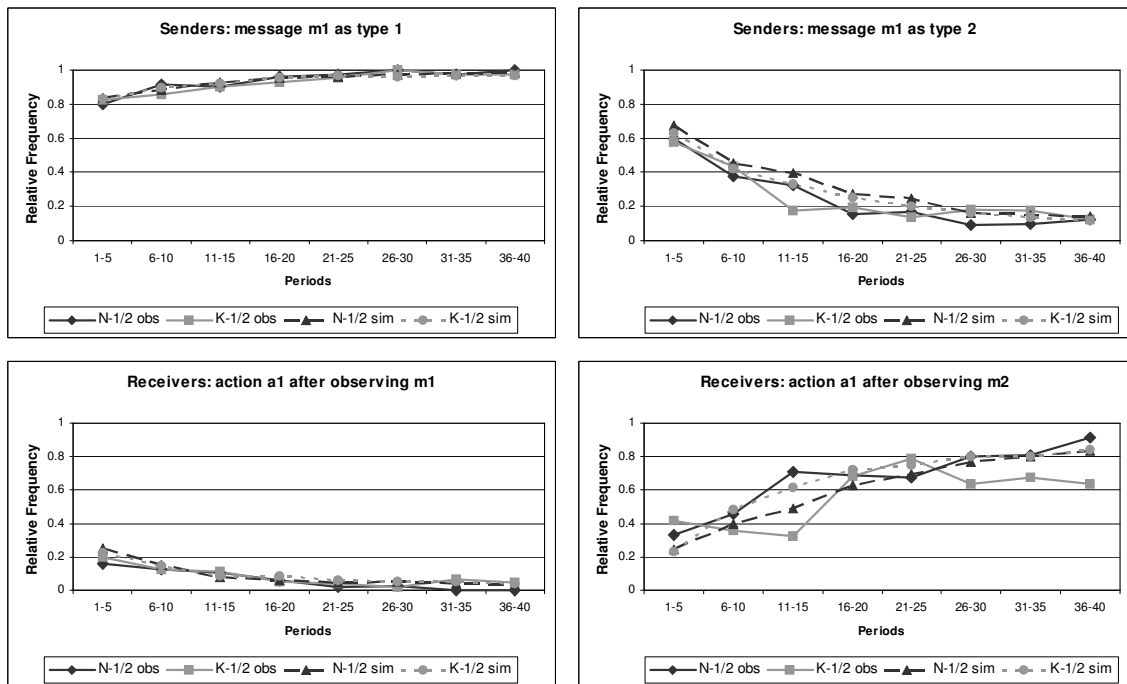


Figure 6: Observed and simulated time series of average strategies in treatments $N-1/2$ and $K-1/2$.

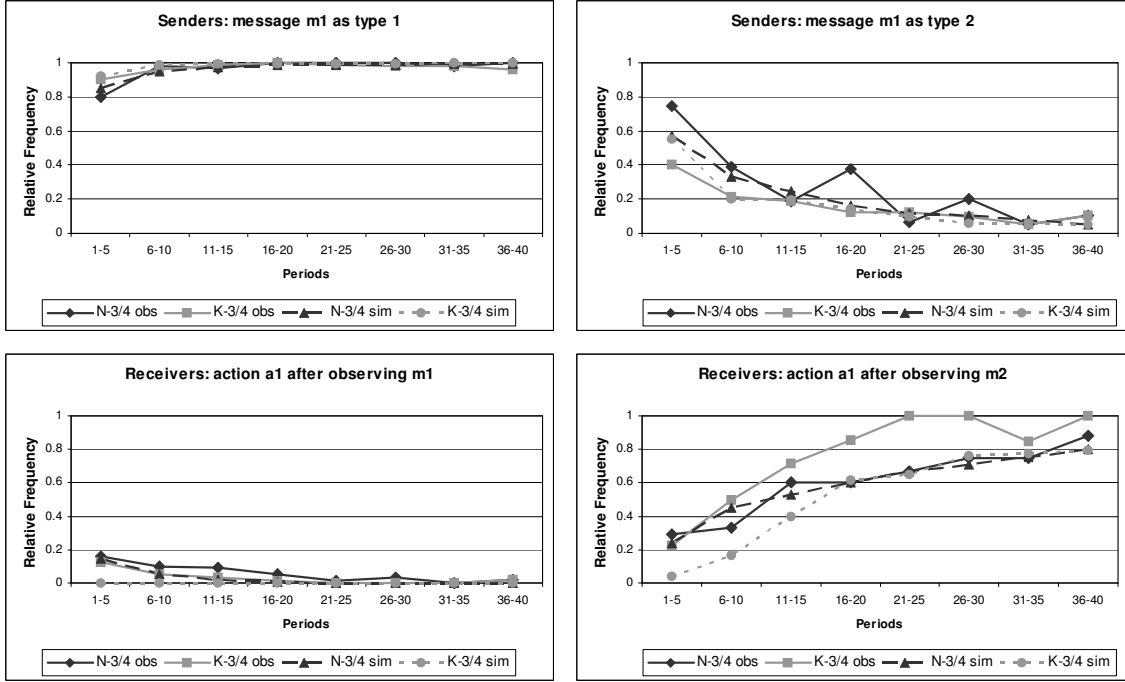


Figure 7: Observed and simulated time series of average strategies in treatments $N-3/4$ and $K-3/4$.

model slightly, but consistently, underestimates the choice of message m_1 for type t_1 of the sender and overestimates it for type t_2 in both treatments $N-1/4$ and $K-1/4$. In particular, the simulated senders of type t_1 learn much quicker not to choose message m_1 when the prior $p = 1/4$ is known (upper left panel in Figure 5). Second, with regard to receivers' behavior we observe that the learning model slightly, but again consistently, overestimates the choice of action a_1 after observing message m_1 and underestimates it after message m_2 in both treatments. In particular, simulated receivers after observing message m_1 learn much quicker to choose action a_1 when the prior $p = 1/4$ is known (lower left panel in Figure 5).

Let us now turn to Figures 6 and 7 that show the results of the simulations in the treatments with prior $p = 1/2$ and $p = 3/4$, respectively. In both of these treatments the fit is very good for the sender when the type is t_1 and for the receiver conditional on observing message m_1 . In fact, the fit in these cases is so good that the lines representing observed and simulated data can hardly be distinguished. Furthermore, whereas the fit for type t_2 senders is also quite good in both treatments, the learning model only has some problems tracking the behavior of receivers conditional on observing message m_2 in treatment $K-3/4$, as the predicted share is consistently lower than the observed share of message m_2 .

Summarizing the results of the simulations, it seems fair to say that they account for the primary patterns found in the data. Strategies of simulated players follow a similar evolution as in the actual experimental sessions. Note that our learning model is quite parsimonious as it has only three parameters. Furthermore, we did not need to make parameters dependent on treatments to have the model approximately replicate the patterns of strategy adjustment in all six treatments.

Note that the best-fitting values for parameters A and $A^T(1)$ are different. Since both these belief propensities are updated by 1 with each observation, the lower value for A means that the players adjust their beliefs about the opponent's strategy faster than their beliefs about the probabilities of the chance move. As discussed in the preliminary description of the learning model in Section 2, this is precisely what is needed for the model to produce the differences we observe.

6 Summary and conclusions

The paper studies experimentally the effect of not inducing a common prior belief in incomplete-information situations using a specific two-person signaling game. We implemented six versions of this signaling game that differ in the underlying prior distribution of the sender's type ($p \in \{1/4, 1/2, 3/4\}$) and whether or not the prior distribution is commonly known. Feedback in the experiments was such that players had the opportunity to learn over time about the strategy of the other players and the prior distribution of the sender's type. Furthermore, the game was designed so that, according to a plausible adjustment theory, early play in the games with the lowest value of p may direct behavior towards different long-run outcomes depending on p being commonly known or not, while initial information about p should be inconsequential with regard to long-run behavior in the games with higher values of p .

The main result of our paper is that these predictions are indeed borne out by the data. More precisely, whereas we find significant differences in behavior in the two games with $p = 1/4$, we cannot reject the hypothesis that there is no difference between corresponding pairs of games for the other values of p . The second main result is that the learning model we present (in which players learn both about the other player's strategy and the prior distribution of the sender's type in case it is not known) quite accurately tracks the main features of the adjustment process in our observed data.

We view our results as an existence proof that there are games (or, more speculatively, classes of games) with incomplete information for which inducing or not inducing a common prior

may make a difference with respect to agents' long-run behavior. This should be of some interest in the light of the results by Güth and Ivanova-Stenzel (2003) who showed that commonly known beliefs are inessential for bidding behavior in their specific asymmetric first-price and second price independent-value auctions.

Of course, the results presented here and elsewhere are so far not more than isolated findings and many more research efforts are necessary and desirable to get a fuller picture of the effect of inducing or not inducing a common prior. We see various avenues for fruitful future research. First, it would be interesting to consider different sender-receiver games and to test the ability of the learning model (or its short-cut version) presented in this paper to predict adjustment patterns in these games depending on whether or not a common prior is induced. Also, more experimental tests of other auction formats (especially those with common values) are warranted as auctions are an important application of game theory. Second, a more systematic inquiry into the relationship between the feedback given to players and the adjustment path in games with and without a common prior would be called for. After all, in this paper we concentrated on the arguably most simple case where subjects were informed about nature's move in their individual encounters at the end of each period. What happens in case this information is not given to subjects or in case information about the realization of the players' type can only indirectly or only ambiguously be derived from own payoffs? Third, in this paper we decided to abstain from properly eliciting beliefs due to the assumed complexity of the situation from a subject's perspective. But, certainly, it would be interesting to extend the analysis to directly learning about what beliefs subjects hold. Finally, other important classes of games with incomplete information and (not) commonly known priors should be subjected to thorough experimental tests, such as market games with incomplete information about e.g. costs or demand.

A Experiment Instructions

Instructions

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

In this experiment you can earn varying amounts of money, depending on which decisions you and other participants make. The experiment will consist of 40 rounds, in each of which you can earn Points. Your payoff at the end of the experiment is equal to the sum of your own payoffs from all rounds. For every 10 Points you will be paid 5p.

Description of the experiment

In each round of the experiment, two participants are randomly matched and interact with each other: one participant will be called “A-participant” and the other “B-participant.” The A-Participant can be of two types that we call “Type 1” and “Type 2.” At the beginning of each round, a random draw determines the type of the A-participant. The random draw is such that with an $X\%$ chance the A-participant will be of Type 1, and with a $(100 - X)\%$ chance of Type 2. You receive no information about the value of X , except that X is constant over all rounds of the experiment. After the random draw, the A-participant will be informed about his/her type. However, the B-participant will not be informed about the type of the A-participant. Knowing his/her type, the A-participant has to decide between options “C” and “D”. Then, the B-participant will be informed about which option was chosen by the A-participant. Knowing the option chosen by the A-participant, but not knowing his/her type, the B-participant will now have to choose between options “E” and “F”. After that, the payoffs of the two participants are determined according to the tables overleaf.

[Instead of the sentences with X , it was “The random draw is such that with a 25% chance the A-participant will be of Type 1, and with a 75% chance of Type 2.” in the treatment with known prior distribution $p = 1/4$, and the corresponding probabilities for the other treatments with known prior distribution.]

Payoffs

Generally, 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 the table on the left applies and for the A-participant of Type 2 the table on the right applies.

Payoff table for Type 1 of the A-participant:				Payoff table for Type 2 of the A-participant:			
		Decision of the B-participant				Decision of the B-participant	
		E	F			E	F
Decision of the A-participant	C	15, 10	80, 80	Decision of the A-participant	C	80, 80	15, 30
	D	25, 10	50, 50		D	50, 50	25, 30

The B-participant's payoffs

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

Payoff table for the B-participant if A-participant chose option "C":				Payoff table for the B-participant if A-participant chose option "D":			
		Decision of the B-participant				Decision of the B-participant	
		E	F			E	F
Type of the A-participant	1	15, 10	80, 80	Type of the A-participant	1	25, 10	50, 50
	2	80, 80	15, 30		2	50, 50	25, 30

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 will be of Type 1 and with a $(100 - X)\%$ chance of Type 2. You receive no information about the value of X , except that X is constant over all rounds of the experiment. [The last two sentences were "With a 25% chance the A-participant will be of Type 1, and with a 75% chance of Type 2." in the treatment with known prior distribution $p = 1/4$, and corresponding probabilities for the other treatments with known distribution.]
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" or "F".
5. Payoffs result as described above.

Number of rounds, role assignment and matching

The experiment consists of 40 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. (The computer program makes sure that half of the participants will be assigned the role of an A-participant and the other half the role of a B-participant.) You will then keep the same role during the entire experiment.

In each round the computer will randomly match participants in pairs of two (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 to whom you will be matched with in the current round.

References

- [1] Anderson, C.M., Camerer, C.F. (2000) "Experience-weighted attraction learning in sender-receiver signaling games", *Economic Theory* 16, 689-718.
- [2] Banks, J., Camerer, C., Porter, D. (1994) "An experimental analysis of Nash refinements in signaling games", *Games and Economic Behavior* 6, 1-31.
- [3] Brandts, J., Holt, C.A. (1992) "An experimental test of equilibrium dominance in signaling games", *American Economic Review* 82, 1350-1365.
- [4] Brandts, J., Holt, C.A. (1993) "Adjustment patterns and equilibrium selection in experimental signaling games", *International Journal of Game Theory* 22, 279-302.
- [5] Brandts, J., Holt, C.A. (1996) "Naive Bayesian learning and adjustment to equilibrium in signaling games," Working Paper, University of Virginia.
- [6] Camerer, C., Johnson, E., Rymon, T., Sen, S. (1993): "Cognition and framing in sequential bargaining for gains and losses", in: Binmore, K., Kirman, A., Tani, P. (eds.), *Frontiers of Game Theory*, MIT Press, Cambridge, MA.
- [7] Costa-Gomes, M., Crawford, V.P., Broseta, B. (2001) "Cognition and behavior in normal-form games: an experimental study", *Econometrica* 69, 1193-1235.
- [8] Cooper, D.J., Garvin, S., Kagel, J.H. (1997) "Signalling and adaptive learning in an entry limit pricing game", *RAND Journal of Economics* 28, 662-683.
- [9] Dekel, E., Fudenberg, D., Levine, D.K. (2004) "Learning to play Bayesian games", *Games and Economic Behavior* 46(2), 282-303.
- [10] Fischbacher, U. (2007) "z-Tree: Zurich toolbox for ready-made economic experiments", *Experimental Economics*, 10(2), 171-178.
- [11] Greiner, B. (2004) "An online recruitment system for economic experiments", in Kremer, K., Macho, V. (eds.) *Forschung und wissenschaftliches Rechnen 2003*, GWDG Bericht 63, Göttingen: Ges. für Wiss. Datenverarbeitung, 79-93.
- [12] Güth, W., Ivanova-Stenzel, R. (2003) "Asymmetric auction experiment with(out) commonly known beliefs", *Economics Letters* 80, 195-199.

- [13] Harsanyi, J.C. (1967) "Games with incomplete information played by 'Bayesian' players, part 1: The basic model, part II: Bayesian equilibrium points, part III: The basic probability distribution of the game", *Management Science* 14, 159–182, 320–334, 486–502.
- [14] Oechssler, J., Schipper, B. (2003) "Can you guess the game you're playing?", *Games and Economic Behavior* 43, 137-152.
- [15] Siegel, S., Castellan, N.J., Jr. (1988) *Nonparametric Statistics for the Behavioral Sciences*, 2nd edition, McGraw-Hill, New York.