Overcoming Coordination Failure in a Critical Mass Game: Strategic Motives and Action Disclosure[☆]

Aidas Masiliūnas

Aix-Marseille University, CNRS, EHESS, Centrale Marseille and AMSE, Centre de la Vieille Charité, 2 rue de la Charité, 13236 Marseille, France.

July 26, 2017

Abstract

We study whether coordination failure is more often overcome if players can disclose their actions at a lower cost. In an experiment subjects first choose their action and then choose whether to disclose this action to other group members, and disclosure costs are varied between treatments. We find that no group overcomes coordination failure when action disclosure costs are high, but half of the groups do so when the costs are low. Simulations with a belief learning model can predict which groups will overcome coordination failure, but only if it is assumed that players are either farsighted, risk-seeking or pro-social. To distinguish between these explanations we collected additional data on individual preferences and the degree of farsightedness. We find that in the low cost treatment players classified as more farsighted more often deviate from an inefficient convention and disclose this action, while the effect of risk and social preferences is not significant.

Keywords: lock-in, coordination failure, learning, strategic teaching, farsightedness, collective action, critical mass, response time
JEL classification: C72, C92, D83

[☆]I am indebted to my PhD supervisors Friederike Mengel and J. Philipp Reiss for their guidance and support. I would also like to thank Jordi Brandts, David Cooper, Lu Dong, Kyle Hyndman, Ernesto Reuben, Arno Riedl, Martin Sefton, Roberto Weber, colleagues at Maastricht University and participants at seminars and conferences in Maastricht (7th M-BEES, BEElab meeting, AiO seminar, PhD Colloquium, GAMES 2016), Utrecht (Experimental Social Science on Social Dilemmas), Prague (ESA European meeting 2014), Karlsruhe (Doctoral Research Seminar), Marseille (Aix-Marseille School of Economics), Toulouse (IMEBESS 2015, NAG 2016), Nottingham (NIBS 2015 Workshop), Montreal (11th World Congress of the Econometric Society), Malta (SEET 2016 workshop), Bonn (Ratio seminar), Naples (Labsi workshop), Cergy (7th ASFEE conference), Vilnius (5th Conference on Economic Research), Wageningen (SABE/IAREP Conference), Zurich (ETH) and Munich (TU Munich) for comments and suggestions. Financial support from the European Union (grant PIEF 2009-235973) and GSBE at Maastricht University is gratefully acknowledged. This paper was completed as a part of my PhD dissertation at Maastricht University.

Email address: `aidas.masiliunas@gmail.com` (Aidas Masiliūnas)

1. Introduction

Governments in Tunisia and Egypt remained in power for more than two decades until they were swiftly overthrown by a wave of protests. Success of the protest movements has often been attributed to information and communication technologies (ICT) that spread at the onset of the protest movements and were used by many protesters,[1] but is has also been argued that the role of ICT has been exaggerated and it may have had adverse effects.[2] A better understanding of the role played by ICT in the democratization process would have important implications for public policy, yet scientific evidence is lacking and the few studies that do address this question use empirical country-level data, making the causal relationship hard to establish (see Kedzie, 1997, Best and Wade, 2009, and a review in Meier, 2011). One way to establish causality is to run an experiment, and we do so by varying one aspect of information technology – lower costs to disseminate information about events – in a controlled laboratory environment.

In the experiment we use a critical mass game that has been used to model the dilemma faced by citizens in oppressed countries (Angeletos et al., 2007; Edmond, 2013). In a critical mass game participants choose between a safe action (call it "stay out") with a fixed payoff, and a risky action (call it "revolt") that pays a large payoff if sufficiently many group members choose this action. Three outcomes are possible in this game: efficient coordination, when everyone is revolting, coordination failure, when everyone is staying out, and miscoordination, when both actions are played by some participants.[3] Coordination failure may occur even though efficient coordination is preferred by all citizens.

Once coordination failure occurs, it is difficult to overturn because of two factors: the history of inefficient coordination and imperfect observability of the actions taken by others. Revolts often start suddenly, after years with no signs of protests (Kuran, 1989). This history of no revolts establishes an inefficient convention, as defined by Lewis (1969) and Young (1996): staying out is customary (everyone has been staying out), expected (everyone expects others to stay out) and self-enforcing (staying out is optimal as long as others do so). In the experiment an inefficient convention was established by starting the game with a high participation threshold, which in similar games leads to coordination on the inefficient equilibrium (e.g. Brandts and Cooper, 2006a). The second factor that inhibits efficient coordination is imperfect observability of revolt levels. Once protests start, it is

[1] For example, almost half of the protesters in Egypt said that they have taken and shared pictures or videos of the protests (Tufekci and Wilson, 2012) and social networking services in both Egypt and Tunisia were primarily used to raise awareness inside the country about the civil movements (Mourtada and Salem, 2011).

[2] It has been stated that social media may be a result rather than a cause of discontent with the regime (Shirky, 2011), online activism may crowd out activism in the streets (Morozov, 2012) and protests movements are not hampered by attempts to shut down communication channels (Hassanpour, 2014). For a broader overview of the debate see Kalathil and Boas (2003), Morozov (2012) or Lynch (2011).

[3] The term "coordination failure" may be used to denote failure to coordinate on an equilibrium, but we use terminology that is consistent with the literature, such as Van Huyck et al. (1991) and Devetag and Ortmann (2007).

difficult to obtain accurate information about the participation levels: geographical spread inhibits direct observation, and lack of objective mass media inhibits indirect observation. When mobilization attempts are not observed, informational cascades that boost participation levels fail to start (Lohmann, 1994). However, observability can be improved if players have access to information technology, which facilitates action disclosure. In the experiment players had an option to disclose their action by paying a certain cost, and only the disclosed actions were observed in the feedback stage. Disclosure costs were varied across treatments: disclosure was cheap in the LOW treatment but expensive in the $HIGH$ treatment. Differences in costs represent heterogeneity in the access to information technology: if such technology is available, it is inexpensive to inform a large number of people about the chosen action through social media, but the cost increases if information technology is not available.

We study the cost of information disclosure because there is evidence that ICT helps mobilization in large part because easier information dissemination alleviates the collective action problem. Protest movements in Russia following electoral fraud in 2011 were facilitated by the spread of online social networks which lowered the costs of collective action: fractionalization between different social networks decreased the frequency of protests while the spread of social media reduced neither the share of pro-government vote or nor the approval of the government (Enikolopov et al., 2016). During economic downturns protests are more common in the areas with high mobile phone coverage because increased mobile phone usage increases responsiveness both to the economic conditions and to the number of other people protesting (Manacorda and Tesei, 2016). Governments that seek to maintain the status quo seem to understand the potential consequences of ICT as some governments limit information dissipation by blocking communication channels (Dainotti et al., 2014) while others censor certain content, especially posts that encourage social mobilization (King et al., 2013).

ICT could alleviate the collective action problem either by making it easier to coordinate when and where to protest (termed “tactical coordination” by Little, 2016) or by increasing the number of people who protest (“political coordination”). Tactical coordination is easier if participants can use social media such as Facebook *before* taking political action, and political coordination is improved if participants use Twitter or Youtube to disseminate information about turnout levels *after* taking political action. In experiments, the role of ICT in improving tactical coordination could be modeled through pre-play communication, while improvements in political coordination could be modeled through post-play disclosure. Previous literature has shown that two-way pre-play communication can improve coordination in two-player Pareto-ranked coordination games (Cooper et al., 1992) and in nine-player minimum and median games (Blume and Ortmann, 2007), one-way free-form communication from the leader helps overcome coordination failure (Brandts et al., 2014) while restricted communication helps prevent coordination failure (Sahin et al., 2015), although when the efficiency gains are small the effect is small and not persistent, and there is no effect in overcoming coordination failure (Dong et al., 2017). In this paper we study the role of ICT in decreasing disclosure costs, which has not been investigated before.

The first goal of our study is to determine whether lower disclosure costs help overcome coordination failure. We find that they do: half of the groups overcome coordination failure in *LOW*, while no group does so in *HIGH*. This treatment difference cannot be explained by differences in histories because all groups in both treatments experience coordination failure when the participation threshold is high. The treatment difference is small even immediately following the decrease in the threshold, but the gap grows over time as some groups in *LOW* start to overcome coordination failure. All groups that overcome coordination failure sustain efficient coordination even when the participation threshold is increased back to its original level, leading to a persistent gap between two treatments. We show that the treatment difference can be explained by a learning model and initial treatment differences in terms of disclosure rates.

Our second goal is to understand the motivation behind the decision to deviate from the inefficient convention. These deviations could be myopically optimal given beliefs and preferences, or they could be used strategically, to influence the future choices of other group members. Although the strategic teaching hypothesis has been postulated as a potential reason for deviations (Brandts and Cooper, 2006a, Devetag, 2003, Friedman, 1996), no attempt has been made to test the existence of such strategic motives. In our experiment the identification of strategic motives is possible because of two novel features: costly action disclosure and the elicitation of farsightedness. Disclosure decisions provide additional evidence about the motives because players who revolt solely for myopic reasons would not pay to disclose their actions, while players who revolt for strategic reasons would disclose if the cost was lower than the expected increase in future earnings. To independently measure the ability to plan ahead, we designed a separate task which was implemented at the end of the experiment together with additional tests to elicit beliefs, risk attitudes and pro-social preferences.

We find evidence that deviations from an inefficient convention are motivated by strategic considerations. Many players are willing to pay to disclose their actions and disclosures are systematic: revolts are disclosed much more often than stay outs and disclosures are most frequent at the start of each block. We also find that in *LOW* subjects classified as more farsighted tend to initiate revolts more often.

1.1. Related Literature

Experiments on overcoming coordination failure in coordination games have been conducted by Brandts and Cooper (2006a) and Hamman et al. (2007), although their focus on coordination failures in organizations has lead to different design choices.[4] Both studies

[4]The game used in this study does not have the weak-link structure, thus it is not necessary for all participants to cooperate for the efficient equilibrium to be implemented (although used the weak-link structure in the first block of the game to achieve convergence to the inefficient equilibrium); after the initial block in which groups coordinate on the inefficient equilibrium, literature on coordination failure in organizations changes the bonus that is paid to group members if the efficient equilibrium is achieved, while we model a decrease in the strength of the regime by lowering the participation threshold; participants in our experiment choose one of two actions, leading to two Nash equilibria, while the strategy space and the

find that whether a group overcomes coordination failure can be predicted by the outcomes in the first round, and some players (called "leaders" by Brandts and Cooper, 2006a) are more likely to initiate the transition than others ("laggards"). Our study extends these findings by showing that whether a person is a leader or a laggard partly depends on their farsightedness level.

Other studies have focused on the role of feedback in coordination games. Brandts and Cooper (2006b) show that groups receiving full feedback about the choices of other participants overcome coordination failure more often than groups receiving information only about the minimum effort level. Devetag (2003) finds that full feedback improves efficient coordination in a game that is similar to ours, but without the initial period of inefficient coordination. On the other hand, Van Huyck et al. (1990) note that informing players about the distribution of actions does not improve coordination in a minimum effort game. In our study the level of feedback is determined endogenously, allowing us to investigate not only the effects, but also the causes of action disclosure. Also, heterogeneity in received feedback allows us to investigate the relationship between observed actions and choices more precisely and to determine how efficient coordination is affected by small increases in observed actions.

Although the strategic teaching hypothesis has not been tested in N-person coordination games, it was tested in two player games with two Pareto-ranked Nash equilibria (Hyndman et al., 2009), three non Pareto-ranked equilibria (Terracol and Vaksmann, 2009) and a unique equilibrium (Hyndman et al., 2012). All three studies find evidence for strategic teaching by showing that (efficient) coordination is less likely when: (i) the cost of strategic teaching is increased (Hyndman et al., 2009), (ii) players are rematched in each round or (iii) information about the payoffs of other players is not displayed (Hyndman et al., 2012). We contribute to the strategic teaching literature in several ways. First, we investigate whether strategic teaching is possible in games with more than two players. For strategic teaching to be profitable in two player coordination games the teacher must be able to affect the beliefs of the other party about one's future choices. In N-person games strategic teaching by one player may be not enough, even if all other players are receptive learners. To find strategic teaching profitable, players must expect that others are using strategic teaching too, and each player would prefer others to perform strategic teaching rather than performing it themselves, leading to a collective action problem and potentially lowering the willingness to use strategic teaching. Second, we test the strategic teaching hypothesis in a different way and using a smaller treatment manipulation, making it unlikely that the treatment difference could be explained by a theory that is not based on players acting strategically.[5] Costly action disclosure allows us to identify teaching at-

number of equilibria in the minimum effort game are larger.

[5]Findings made in two player games, while consistent with strategic teaching, might also be explained by other factors: larger costs of strategic teaching would lower the likelihood that a player will choose an efficient action even for players who are myopic but choose stochastically; rematching players every round could increase strategic uncertainty and make learning more difficult; social preferences might be switched off when players are not informed about payoffs of other participants.

tempts more precisely, and we show that personal characteristics, such as farsightedness, can partly explain why some players use strategic teaching while others do not.

Several studies have investigated the role of observability in bank run games, which have two Pareto-ranked equilibria: in the efficient equilibrium all players keep money deposited and in the inefficient one deposits are withdrawn (Diamond and Dybvig, 1983). Garratt and Keister (2009) use a setup related to ours: decisions are made simultaneously and efficient coordination is easy in the first stage, establishing the play of the efficient equilibrium, but more difficult in the second stage when randomly chosen depositors are forced to withdraw. The study finds that bank runs are more common when depositors have three opportunities to withdraw and observe previous withdrawals. Kiss et al. (2014) manipulate observability in a sequential bank run, and show that first depositors withdraw less frequently when they know that their action will be observed by the second or third depositor. This finding could be explained by strategic teaching, especially as second depositors react to feedback: they withdraw more often after observing a withdrawal and less often after observing no withdrawal. Kinateder et al. (2015) show theoretically and experimentally that bank run rates are lower when there is an option to signal the decision to keep the money deposited. Costly signaling is used by a sizable number of participants, especially at the start of the sequence, a finding that is consistent with the strategic teaching hypothesis. Such signals are also effective as they decrease the probability of a subsequent withdrawal.

More generally, the critical mass game resembles other repeated coordination games with strategic complementarities, complete information and more than two players, such as entry games (Heinemann et al., 2004; Duffy and Ochs, 2012), order-statistic games (Van Huyck et al., 1990, 1991, 2007; Kogan et al., 2011), games with network externalities (Ruffle et al., 2015; Mak and Zwick, 2010) and step-level public goods games (Rapoport and Eshed-Levy, 1989; Offerman et al., 1996; Sonnemans et al., 1998). For similar one-shot games, see Heinemann et al. (2009) or Keser et al. (2012). Global games of regime change use a similar payoff function, but assume incomplete information about the participation threshold (Angeletos et al., 2007, tested in an experiment by Shurchkov, 2013).

There are several theoretical studies that investigate the role of ICT and strategic motives in protest movements. Barbera and Jackson (2017) assume a continuum of players who differ in terms of how they feel about the regime. The study shows that revolts are more likely to succeed if they are preceded by a commonly observed demonstration: the most dissatisfied types participate in the demonstration, affecting the beliefs of the moderate types who subsequently participate in a revolt. However, such a mechanism does not work if demonstrations are costless, e.g. if information is transmitted via polls or social media. Little (2016) assumes that social media makes it easier to express one's opinion about the regime, and finds that it does not necessarily lead to higher protest levels because negative opinions can be shared as often as positive ones, and subjects discount the size of grievances if they know that complaining is easy. Kiss et al. (2016) use a sequential game in which players have a preference either for the status quo or for the regime change. The study shows that when the revolution is possible and types are private information, there is a unique equilibrium in which revolution is successful when the actions of individual players are known (social media), but there may be multiple equilibria, including some in which revolution fails, when only the total revolt levels are known (mass media). Lohmann (1994, 1993, 2000) develops a signaling theory of mass action in which citizens with private information about regime strength take part in mass political action to induce others to join. De Mesquita (2010) uses a coordination game with uncertainty about the level of anti-government sentiment. The study finds that vanguards undertake some costly action (e.g. acts of terror) to make the anti-government support seem larger and increase future revolt levels. Edmond (2013) uses a model in which citizens are not perfectly informed about the strength of the regime and thus the regime has incentives to provide biased information. Although the setting which we analyze is different from any of these theoretical models,[6] our experiment reinforces the view that participation in political action could be used strategically and that ICT might increase the likelihood of revolts.

2. Experimental Design

At the start of the experiment participants were matched into groups of six and played a variant of a critical mass game (Heinemann et al., 2009; Devetag, 2003) for 33 rounds. In each round players went through four steps: they (i) chose an action, (ii) chose whether to disclose this action, (iii) reported beliefs about the actions chosen by other group members and (iv) received feedback about the actions of those group members who chose to disclose.

2.1. Action Choice

In each round participants chose an action, R (for *"risky"* or *"revolt"*) or S (for *"safe"* or *"stay out"*).[7] The payoff of S was fixed while the payoff of R depended on whether the total number of players who chose R (denoted by $\#R$) exceeded the participation threshold (denoted by θ):

$$\pi(R) = \begin{cases} 100 \text{ ECU} & \text{if } \#R \geqslant \theta \\ 5 \text{ ECU} & \text{if } \#R < \theta \end{cases}$$
$$\pi(S) = \quad 60 \text{ ECU}$$

This payoff function with $2 \leqslant \theta \leqslant 6$ leads to two Pareto-ranked stage game Nash

[6] Most theoretical studies include heterogeneity across the players either in terms of beliefs about the regime strength or in terms of preferences towards regime change. A strategic player thus can affect the actions of other players by altering their beliefs about the unknown variable, a process that is made possible by ICT. In our paper the strength of the regime is public knowledge, and all players favour regime change over status quo. Alternatively, our game could be interpreted as having two types of players, but computer makes a decision to stay out on behalf of those who favor status quo. Decreased difficulty to coordinate in the second block could be due to some players now favoring regime change instead of the status quo.

[7] In the experiment we used neutral language, labeling actions as "A" and "B". Screenshots of the decision screen are reproduced in Appendix K.

equilibria in pure strategies:[8] in the efficient equilibrium everyone chooses R and in the inefficient one everyone chooses S.

Our main objective is to study how an efficiency-improving transition could take place following coordination failure. To establish coordination failure we divided the game into three blocks and set the participation threshold equal to 6 in the first one, so that 100 ECU is received only if all six players in a group choose R. With this parameter choice, a critical mass game reduces to a binary minimum effort game, in which large groups typically converge to the inefficient equilibrium (Van Huyck et al., 1990). To test whether coordination failure can be overcome we made efficient coordination easier in the second block by reducing the participation threshold to 5. In the third block the threshold was increased back to 6 to test if the history of efficient coordination improves coordination in a more difficult environment.

The transfer of precedent from one block to the other was facilitated by keeping all elements of the game constant, except for the participation threshold, and by framing the game in terms of several blocks in a game rather than as separate games. At the start of each block players were informed about the duration of the current block (13 rounds in block 1, 15 rounds in block 2, 5 rounds in block 3) and about the value of the participation threshold, but no information was provided about the parameter values in future blocks.

2.2. Action Disclosure, Belief Elicitation and Feedback

After choosing an action, R or S, players had an option to disclose their action to other group members. The cost of not disclosing an action was 1 ECU, while the cost of disclosing depended on the treatment to which a player was assigned. The cost parameters were chosen to generate sufficient difference according to the strategic teaching model (for details, see Appendix C):

- In *LOW* the disclosure cost was 2 ECU.
- In *HIGH* the cost was 80 ECU.

After the action disclosure stage we elicited beliefs using a procedure adapted from Heinemann et al. (2009): players were asked to report the probability that a randomly selected other group member will choose R in the current round. The belief elicitation task was incentivised and we used a binarized scoring rule that elicits truthful beliefs even if subjects are not risk neutral (Hossain and Okui, 2013). In coordination games stated beliefs may not be truthful because of hedging (see Armantier and Treich, 2013), therefore we reduced the incentives to hedge by either paying for the main part of the experiment or by playing a lottery in which the probability to receive the higher prize was determined by performance in a randomly selected belief elicitation task. Details about the belief elicitation task are presented in Appendix F.

[8] There also are equilibria in mixed strategies. In a symmetric mixed strategy equilibrium all players choose R with probability $I^{-1}_{\frac{11}{19}}(\theta-1, n-\theta+1)$, where I^{-1} is an inverse of the regularised incomplete beta function. In this equilibrium R is chosen with probability 0.72 if $\theta = 5$, and with probability 0.90 if $\theta = 6$.

At the end of each round players were informed only about the actions taken by group members who disclosed their actions. In particular, players were informed about:

- The total number of group members who chose R and disclosed ("observed R")
- The total number of group members who chose S and disclosed ("observed S")
- The total number of group members who did not disclose their actions ("unobserved")

In addition, players were informed about their round income and whether at least θ players chose R. A history box displaying own actions, stated beliefs and observed feedback in previous rounds was always visible on the computer screen. Feedback about the accuracy of reported beliefs was provided only at the end of the experiment to suppress information about the choices of group members who did not disclose their actions.

2.3. Procedure

Experiments were conducted in the BEElab at Maastricht University in February, 2014, using z-Tree (Fischbacher, 2007) and ORSEE (Greiner, 2015). 72 participants took part in the experiment, 36 in each treatment. The average duration of the experiment was 100 minutes and average earnings were 18.50 euros. Earnings in all parts of the experiment were denoted in ECU and exchanged into euros using the conversion rate of 250 ECU = 1 EUR. Negative round income was possible in *HIGH*, therefore in both treatments players started with an initial balance of 400 ECU.

After reading the instructions but before starting the experiment subjects answered several questions about the payment scheme and calculated earnings in hypothetical situations to ensure that instructions were well understood (questions are reproduced in figures K.9 and K.10 in Appendix K). The questionnaire was computerised and subjects could not continue until all the questions have been correctly answered.

2.4. Theoretical Predictions

In each stage game players chose an action (R or S) and whether to disclose it (D) or not (ND), thus each player had four strategies: (R, D), (R, ND), (S, D) and (S, ND). Disclosure is costly, therefore (R, ND) dominates (R, D) and (S, ND) dominates (S, D) so no action should be disclosed if a stage game was played once or if all players were myopic. However, actions might be disclosed on the path of a subgame perfect Nash equilibrium (SPNE) in the repeated game. For example, consider a strategy that prescribes R if all players disclosed their actions in all previous rounds and prescribes S otherwise. There is a SPNE in which all players use such strategies, and on the equilibrium path all players would be choosing R and disclosing their actions in all but the last round as long as disclosure costs are sufficiently small. In a similar way, there are SPNE in which players do not coordinate on a common action if, for example, all players use strategies that prescribe R only following such history of miscoordination.

Nash equilibrium is useful in making predictions about stage-game outcomes that may be reached in the long run, but the multiplicity of equilibria leads to vacuous predictions about the dynamic process in a repeated game. An alternative approach is to assume that players are boundedly rational and learn from experience. In particular, we will make predictions about the path of choices using a weighted fictitious play model (Cheung and Friedman, 1997), which typically fits data better than reinforcement learning in coordination games with Pareto-ranked equilibria (Battalio et al., 1997) and which can be extended with sophistication (Hyndman et al., 2009, Masiliunas, 2016). Weighted fictitious play assumes that players form beliefs based on observed history and choose an action that maximises immediate utility conditional on those beliefs. The model predicts that players who observed others choosing a particular action will expect this action to be chosen more often, and will therefore be more likely to choose this action themselves. We formulate this first prediction the following way:

Prediction M1. *Players who observe action R chosen more often are more likely to choose R and to believe that others will choose R.*

However, players in learning models assume that others are using stationary strategies, ignoring the possibility that others may be learning. These models therefore overlook strategic teaching, which may explain deviations from an inefficient convention. To test whether behaviour is driven by such strategic motives we extend the learning model by adding sophistication, and we compare the predictions of this model to the predictions of a standard learning model. The experiment was designed in a way such that the two learning models would make different predictions, allowing us to determine: (i) whether choice dynamics can be accurately approximated by belief learning and (ii) whether the explanatory power of a learning model is increased by adding sophistication.

We add sophistication to the standard learning model by assuming that some players are sophisticated and others are myopic. Myopic players act as predicted by the standard belief learning model (see prediction M1). Sophisticated players anticipate the learning process of myopic players and take into account not only the differences in immediate expected payoffs, but also the differences in future payoff flows generated by each action. A solution concept for N-person critical mass games with sophisticated and myopic players is presented in Masiliunas (2016); here we will only use the features of the model that are useful in making predictions about the treatment difference.[9]

Sophisticated players may deviate from an inefficient convention to increase future payoffs from efficient coordination. Furthermore, sophisticated players would be willing to pay to make their deviations observable because they value a deviation for its informational content rather than for the immediate payoff it generates. A decision to not deviate should not be disclosed because it makes myopic players less likely to choose R in the future, which can only decrease the payoffs of sophisticated players.

Prediction S1. *Action S is never disclosed, R is disclosed more often than S.*

[9]For alterantive ways to model sophistication see Milgrom and Roberts (1991), Kalai and Lehrer (1993) or Camerer et al. (2002)

The model with sophistication also predicts differences in disclosure rates between the two treatments. Action R should be disclosed more often in *LOW* than in *HIGH* because lower disclosure costs reduce the cost of strategic teaching. Since action S should never be disclosed, its disclosure rates should not depend on disclosure costs.

Prediction S2. *R is more often disclosed in LOW than in HIGH.*

In addition to the immediate effect of action disclosure, lower disclosure costs are predicted to help overcome coordination failure. This prediction follows from predictions M1, S1 and S2. Notice that S1 and S2 predict that only R will be disclosed, and more so in *LOW*, therefore M1 predicts that over time the learning process should lead to higher R levels in *LOW* than in *HIGH*. Higher R levels result in an increased frequency of instances when θ or more players choose R, in which case coordination failure would be overcome.

Prediction S3. *R is more frequently chosen and coordination failure is more often overcome in LOW than in HIGH.*

Finally, players with a longer planning horizon would be more likely to deviate from an inefficient convention because they put more weight on future benefits of efficient coordination.

Prediction S4. *The tendency to choose R and disclose it is higher among farsighted subjects.*

It is important to note that predictions S1, S2, S3 and S4 would not hold if all players were myopic, because no myopic player would be willing to pay to disclose an action.[10] Predictions S1, S2, S3 and S4 are therefore not only interesting on their own, but they also help to test whether players are motivated by strategic considerations.

3. Results

We address two main questions in this section. First, we want to know if lower action disclosure costs help overcome coordination failure. Second, we want to see if the transition process can be explained by a learning model, either with sophistication or without it.

3.1. Is Coordination Failure More Often Overcome with Low Action Disclosure Costs?

We compare the treatments in terms of two measures: *R levels*, defined as the total number of players in a group who choose R, and *transition frequency*, defined as the fraction

[10] Action disclosure could be myopically optimal only if players were willing to throw away money to reduce advantageous inequality ($\beta > 1$ in the model of Fehr and Schmidt, 1999). We tested the existence of such preferences in an additional task, where players could reduce their earnings without affecting the earnings of the other person. We found that such preferences were rare: 94% of participants chose not to decrease their own earings.

of groups that coordinate on R at the end of block 2.[11] The former metric captures the willingness to initiate a transition while the latter measures the success of these efforts.

3.1.1. Differences in R Levels

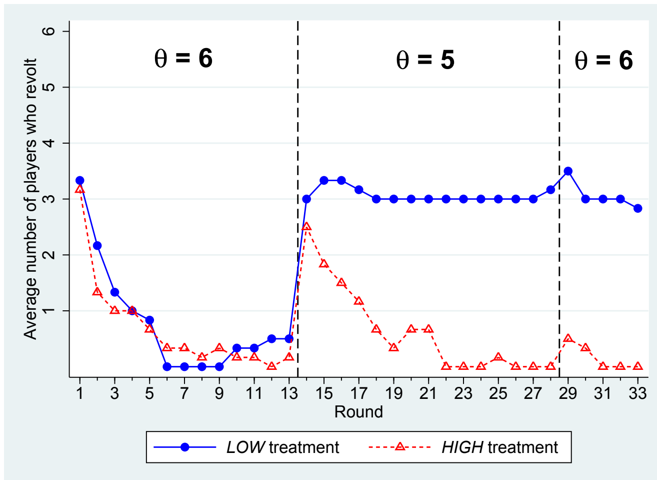

Figure 1: Comparison of R levels, by treatment. Vertical dashed lines separate the three blocks with different values of θ.

In block 1 there is no significant treatment difference in terms of R levels (figure 1). In both treatments about half of the players choose R in round 1, but afterwards R levels decrease until almost nobody is choosing R at the end of block 1. In the last five rounds of block 1 the fraction of players who choose R is 6% in *LOW* and 3% in *HIGH* while average beliefs about R choice are respectively 14% and 13%. A high participation threshold therefore leads to coordination failure regardless of disclosure costs.

The effect of disclosure costs becomes apparent when the participation threshold is lowered in block 2. The treatment difference in round 14 is small and insignificant (Mann-Whitney U test $p > 0.4810$) as a decreased participation threshold leads to a similar jump in both treatments. However, the gap between two treatments soon widens and eventually almost all players in *HIGH* stop choosing R while average R levels hardly change in *LOW*.

Result 1. *R levels in LOW are significantly higher than in HIGH in blocks 2 and 3, but there is no significant difference in block 1. In block 2 the treatment difference increases over time.*

Result 1 can be formally supported by regressing R levels on the treatment variable, number of rounds remaining until the end of the block and the interaction of these two variables. Table 1 shows that the main effect of the treatment variable is significantly

[11] We would not use data from groups that did not achieve coordination failure in the first block, as such groups would have prevented rather than overcome coordination failure. But in our experiments all groups experienced initial coordination failure.

Table 1: Random effects GLS regression using group level data (12 groups). Dependent variable: number of players in a group who choose R. Standard errors clustered on the group level.

	Block 1	Block 2	Block 3
LOW treatment	0.060	3.437**	2.900**
	(0.24)	(2.44)	(2.20)
Rounds remaining	0.172***	0.155***	0.133***
	(3.81)	(2.90)	(2.76)
Rounds remaining * *LOW* treatment	0.009	-0.143**	0.000
	(0.15)	(-2.28)	(0.00)
Constant	-0.353***	-0.450***	-0.100**
	(-2.79)	(-2.94)	(-2.28)
Observations	156	180	60

t statistics in parentheses

* p<0.10, ** p<0.05, *** p<0.01

different from zero in blocks 2 and 3, but not in block 1.[12] We use the estimated coefficients to predict the treatment difference in each round (table B.10 in Appendix B) and find no significant treatment difference in block 1 and at the start of block 2, but a significant difference at the end of block 2 and in block 3. The observation that the gap between two treatments grows over time in block 2 is also evident from the negative coefficient of the interaction term in the second column of table 1.

3.1.2. Differences in Transition Frequency

A large difference between two treatments is also observed if the comparison is performed using transition frequency instead of R levels. Convergence patterns on a group level are illustrated in figure 2. All groups converge to an inefficient equilibrium in the first block, and no group is in equilibrium in round 14, following a change in the participation threshold. However, whether a group moves to the efficient equilibrium in block 2 depends on the treatment: in *HIGH* all groups return to the inefficient equilibrium, while in *LOW* half of the groups move to the efficient one. The flat line of average R levels in *LOW* shown in figure 1 therefore does not tell the entire story: while no group is in equilibrium in round 14, all converge by round 18. All the groups that overcome coordination failure in block 2 stay in the efficient equilibrium in block 3, when the participation threshold is increased to its original value. A temporary decrease in the participation threshold can therefore have a lasting positive effect on coordination.

To test whether the treatment difference in terms of transition frequency is statistically significant, we compare the fractions of group members who choose R in the last round

[12]Results do not change if we exclude the interaction term: the treatment variable is significant at a 5% level in blocks 2 and 3 and not significant in block 1.

of each block. We find that there is a significant treatment difference in the last round of block 2 (Mann-Whitney U test $p = 0.0209$) and block 3 ($p = 0.0578$), but not in block 1 ($p = 0.9020$). Test results are very similar if instead of one last round we average over the last 2 or 3 rounds of a block.

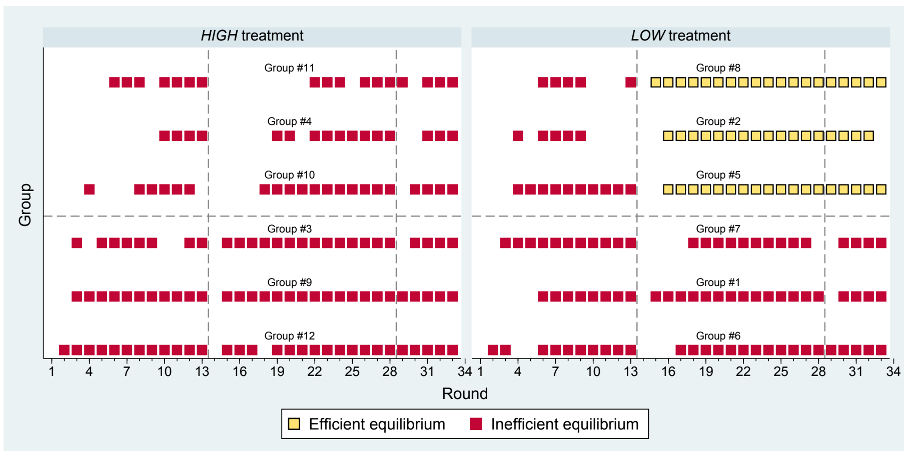

Figure 2: Convergence by group. A group is said to be in an efficient equilibrium if all group members choose R and in an inefficient equilibrium if all group members choose S. Empty space denotes rounds in which a group is not in equilibrium. Vertical dashed lines separate blocks with different participation thresholds. Groups are ordered by the total number of R choices in the entire experiment. Horizontal dashed lines show the median split, groups with higher R levels being at the top.

Result 2. *Coordination failure is overcome more often in LOW than in HIGH.*

Overall, we find that lower action disclosure costs increase both the R levels and the frequency of transitions to the efficient equilibrium. Since action disclosure costs should have no effect if all players were myopic, these findings seem to indicate that some players act strategically, a hypothesis that will be investigated in the next subsection.

3.2. Is Adaptation Explained by Belief Learning Models?

Treatment differences in terms of R levels and transition frequency are mainly driven by the adaptation process: figure 3 shows that groups which transition to the efficient equilibrium in *LOW* (top three groups in the right panel of figure 2) and groups which are most likely to choose R in *HIGH* (top three groups in the left panel of figure 2) exhibit similar patterns of choices in block 1 and round 14, but all groups in *LOW* overcome coordination failure while no group does so in *HIGH*. This subsection tests whether the adaptation process can be explained by a learning model with or without sophistication: we first test the predictions made by each model, then fit each model to the data and compare the goodness of fit and finally we simulate the path of choices using the estimated parameter values to test which model can most accurately predict whether a group will overcome coordination failure.

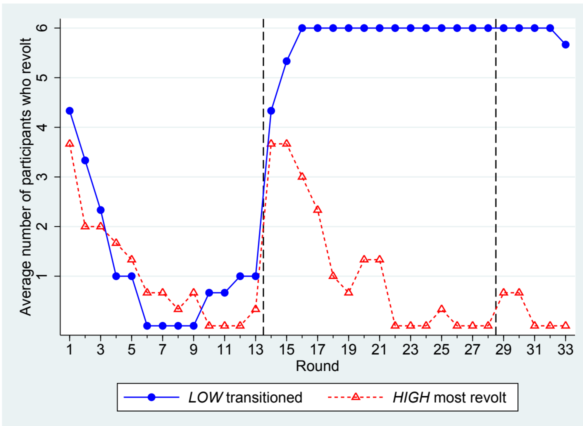

Figure 3: R levels for groups in *LOW* that transition to an efficient equilibrium and groups in *HIGH* that choose R most (three uppermost groups in figure 2 for each treatment).

3.2.1. Action Disclosure

A learning model with sophistication predicts that R will be disclosed more often than S and more often in *LOW* than in *HIGH* (predictions S1 and S2). Figure 4 shows that these predictions are mostly confirmed by the data.[13] In the first round of block 2 all players in *LOW* who choose R also disclose their actions, but only 40% do so in *HIGH*, a difference that persists until all groups converge to an equilibrium in round 18. Actions are more often disclosed at the start of a block than at the end; although this decrease is not the main interest of this paper, we suspect that the rate of disclosure decreases either because the lower number of remaining rounds decreases incentives for strategic teaching or because strategic uncertainty decreases as players converge to an equilibrium, reducing the need for disclosure.

The difference between panel (a) and panel (b) of figure 4 indicates that S is disclosed much less often than R in both treatments. Still, some players do disclose S in *LOW*, typically in rounds in which groups are not in equilibrium.

Result 3. *R is disclosed more frequently than S. R is disclosed more frequently in LOW than in HIGH.*

Result 3 is supported by a probit regression of a disclosure decision on action choice and treatment dummies (see table B.12 in Appendix B). The treatment effect is highly

[13]Occasional fluctuations and missing data in later rounds (panel (a) of figure 4) are caused by a small number of players choosing R in *HIGH*.

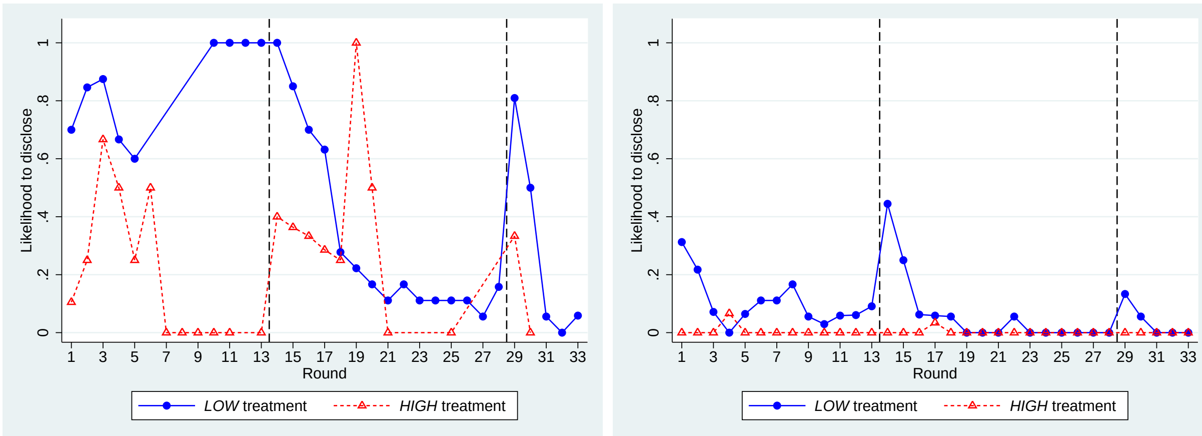

(a) Disclosure conditional on choosing R

(b) Disclosure conditional on choosing S

Figure 4: Frequency of disclosure in *LOW* and *HIGH* treatments.

significant if all data is taken into account or if we look only at disclosure conditional on choosing R. Action R is disclosed significantly more often than S, and the rate of disclosure decreases over time.

3.2.2. Belief Formation

The prediction that observing R increases both the beliefs about R choice and the likelihood to choose R (prediction M1) can be tested on a group level and on an individual level. On a group level, groups with higher initial observed R levels should overcome coordination failure more often. Table 2 provides evidence in favour of this prediction: all groups in which at least 4 players choose and disclose R in round 14 overcome coordination failure, while all other groups do not. This difference cannot be explained by path dependence alone, as groups in which many players choose R but do not disclose their actions do not overcome coordination failure (e.g. groups 4 and 11). Overall, the likelihood of a transition depends on the level of *observed* R choices, and this feedback is lacking when disclosure costs are high.

Result 4. *Transitions to the efficient equilibrium occur in groups that observed high R levels in round 14.*

On an individual level, prediction M1 states that players who observe more R choices should report higher beliefs. To test this prediction we regress stated beliefs on the number of observed R and observed S in the previous two rounds. A correct specification must take into account the fact that our units of observation are not independent: stated beliefs might be correlated with the person's stated beliefs in other rounds as well as with the stated beliefs of other group members in the current round. Thus the most appropriate specification should include an individual random effect nested within a group random effect (three-level nested random effects model). Additionally, we estimate a standard two-

Table 2: The number of players in each group who chose R ("R levels") and the number of players who chose R and disclosed it in round 14 ("Disclosed R"). A transition is said to occur if all group members chose R in the last round of block 2.

Treatment	Group ID	R levels	Disclosed R	Transition?
HIGH	3	2	0	No
HIGH	4	4	3	No
HIGH	9	1	0	No
HIGH	10	3	1	No
HIGH	11	4	2	No
HIGH	12	1	0	No
LOW	1	1	1	No
LOW	6	1	1	No
LOW	7	3	3	No
LOW	2	4	4	Yes
LOW	5	4	4	Yes
LOW	8	5	5	Yes

level model with only a group-specific random component. As an additional robustness check we include a model with only one round lags.

Table 3 shows that observed R levels have a significant positive effect on beliefs about R choice in the subsequent round. The actual R level is also significant, but the significance level and the economic size are much lower than that of the observed R level. Observing more players choosing S is not significant, suggesting that players who do not disclose their actions are perceived in a similar way to those who choose S.

Models in table 3 treat observations from all rounds equally, therefore results are influenced by the correlation between choices and beliefs after play has converged to an equilibrium, potentially overstating the effect of belief learning. We therefore replicate the estimation using data only from rounds in which play has not yet converged, where convergence is said to occur if in the previous round all group members chose the same action, either R or S. In this subset the effect of observed feedback on beliefs is lower, but still highly significant (see table B.13 in Appendix B).

3.2.3. Estimation of a Weighted Fictitious Play Model

After establishing that beliefs are related to observed R levels, we can specify the adaptation process by fitting a belief learning model. We use a weighted fictitious play model (Cheung and Friedman, 1997), modified to accommodate three features of the game. First, we extended the original model specified for two player games to N-person games by assuming that the same probability of choosing R, denoted $b_i(t) \in [0, 1]$, is assigned to every other group member.[14] We make this homogeneity assumption because aggregate feedback does not allow players to distinguish between other anonymous group members

[14] This specification of fictitious play assumes that players expect each group member to use a stationary mixed strategy. In games with two strategies prior beliefs about the probability to choose R follow a beta distribution and are updated using Bayes rule.

Table 3: Three-level nested random effects model and two-level random effects GLS model (12 clusters). Dependent variable: stated probability that a randomly chosen group member chose R (from 0 to 100). Standard errors are robust to correlation within the observations of each group.

	Random effects GLS		Nested random effects	
	Two-period lags	One-period lags	Two-period lags	One-period lags
Observed R in (t-1)	11.20***	11.58***	10.47***	10.83***
	(12.90)	(9.19)	(9.36)	(8.47)
Observed R in (t-2)	2.440**		1.829	
	(2.33)		(1.57)	
Observed S in (t-1)	-2.940**	-2.041	-3.639***	-3.191**
	(-2.50)	(-1.37)	(-2.98)	(-2.00)
Observed S in (t-2)	0.916		-0.0777	
	(1.01)		(-0.07)	
Actual R level in (t-1)	0.362	3.682***	1.128	3.940***
	(0.31)	(3.84)	(0.89)	(4.50)
Actual R level in (t-2)	1.680**		1.836**	
	(1.98)		(2.10)	
Treatment: 1 = *LOW*	-1.029	0.0160	0.189	1.116
0 = *HIGH*	(-0.49)	(0.01)	(0.09)	(0.54)
Constant	13.91***	13.91***	13.82***	13.96***
	(6.14)	(6.77)	(6.30)	(6.86)
Observations	1944	2160	1944	2160

t statistics in parentheses
* p<0.10, ** p<0.05, *** p<0.01

(for a broader discussion and examples of use, see Rapoport, 1985, or Rapoport and Eshed-Levy, 1989). Second, the option to not disclose one's action results in incomplete feedback, thus additional assumptions need to be made about how unobserved actions are perceived. If beliefs were formed using Bayes rule, the interpretation of undisclosed actions would depend on beliefs about the probability that R is disclosed, which is unknown to the players. Instead of imposing Bayesian updating and correct beliefs we use a different approach and add a new parameter (ρ) which specifies how an unobserved action is perceived. In particular, if the number of other group members who did not disclose in round t is denoted by $\#U(t)$, the perceived R level would be calculated by $\#R(t)+\rho\times\#U(t)$, with $\rho \in [0,1]$.[15] If $\rho = 1$, unobserved actions are perceived as R. If $\rho = 0$, unobserved actions are perceived as S. The third feature is the change in the participation threshold from one block to the next, which might mean that experience in one block is not directly applicable to the next block. We therefore allow players to discount experience from a previous block by multiplying the observations from the first block by β_1 and the observations from the second block by β_2. Putting everything together, we use the following formula to determine the beliefs of player i about the likelihood that that any other group member will choose

[15]Equivalently, we could allow players to form three separate beliefs $b_i^R(t)$, $b_i^S(t)$ and $b_i^U(t)$ about the observed number of R, S and U in a standard way, and calculate beliefs about the perceived number of R as a convex combination of $b_i^R(t)$ and $b_i^U(t)$.

Table 4: Calculation of the weights assigned to previous experience.

	Round	1	2	3	...	14	15	...	29	30	...
(1)	Fraction choosing R	0.6	0.2	0.4	...	0.6	0.8	...	0.4	0.6	...
(2)	Discount factor	γ^{t-2}	γ^{t-3}	γ^{t-4}	...	γ^{t-15}	γ^{t-16}	...	γ^{t-30}	γ^{t-31}	...
(3)	Forgetting factor	$\beta_1\beta_2$	$\beta_1\beta_2$	$\beta_1\beta_2$	...	β_2	β_2	...	1	1	...

R in round t, for all $t > 1$:

$$b_i(t) = \frac{\sum_{u=1}^{t-1} \gamma^{u-1} \frac{\#R(t-u)+\rho\#U(t-u)}{n-1} D(t-u)}{\sum_{u=1}^{t-1} \gamma^{u-1} D(t-u)} \tag{1}$$

$$\text{where } D(\tau) = \begin{cases} \beta_1 \times \beta_2 & \text{if } \tau \in [1,13] \\ \beta_2 & \text{if } \tau \in [14,28] \\ 1 & \text{if } \tau \in [29,33] \end{cases}$$

Table 4 illustrates how beliefs are formed. Row (1) shows a hypothetical sequence of the observed fraction of other players choosing action R, including the correction for unobserved actions, while rows (2) and (3) show weights assigned to each observation. In each round beliefs are calculated by dividing the dot product of vectors (1), (2) and (3) by the dot product of vectors (2) and (3). The discount factor assigns lower weight to the history that occurred further back in time, because $\gamma \in [0,1]$. The forgetting factor reduces the weight assigned to observations in previous blocks. For example, when a player is in the third block (round 29 and above), observations from rounds 1–13 receive a weight of $\beta_1\beta_2$, observations from rounds 14–28 receive a weight of β_2 and observations for rounds above 29 receive a weight of 1. When a player is in the second block, all previous observations are multiplied by β_2, which cancels out, therefore observations from the first block receive a weight of β_1 and observations from the second block receive a weight of 1. For players in the first block all observations are multiplied by $\beta_1\beta_2$, which cancels out, leaving a weight of 1.

The model was fit by finding parameters γ, ρ, β_1 and β_2 that minimize the mean squared deviation between $b_i(t)$ and the stated belief. As in all of the estimations that will follow, we did not use data from the first period because there is no previous history to predict beliefs. The estimated coefficient values are: $\hat{\gamma} = 0.43$, $\hat{\rho} = 0.15$, $\hat{\beta}_1 = 0.28$ and $\hat{\beta}_2 = 0.99$. The low estimated value of ρ shows that most players react to an unobserved action as if it was S rather than R. Adding the β_1 parameter improves the fit substantially, and the low estimate shows that history in the first block has only a weak effect on choices in block 2. The estimated value of β_2, however, is close to 1, indicating that history in block 2 has a strong effect on beliefs in round 3 and no transitions occur at the start of block 3. The high estimated value of β_2 and the low value of β_1 therefore indicate an asymmetric effect of changes in participation threshold: a decreased threshold allows some groups to move from the inefficient equilibrium to the efficient one, but a subsequent increase does not lead to any transitions. A temporary decrease in the participation threshold can therefore move some groups from an inefficient to the efficient equilibrium.

We evaluate the fit of a belief learning model by comparing stated beliefs with fictitious play beliefs, calculated using the estimated population level parameters and the actual history of observed feedback. Figures A.5 and A.6 (in Appendix A) show that predicted beliefs track stated beliefs rather well even on an individual level. Largest discrepancies arise for players whose stated beliefs are not sensitive to observed history. Overall, both the regressions and the estimation of a belief learning model indicate that stated beliefs are affected by the observed history, as is predicted by a belief learning model.

Result 5. *Players who observed higher R levels in previous rounds expect that others will be more likely to choose R.*

3.2.4. Response to Stated Beliefs

The previous subsection has shown that the evolution of beliefs can be predicted by a weighted fictitious play model using feedback observed in the experiment. However, only 60% of players best-respond to their stated beliefs at the start of block 1 and block 2, almost entirely because of choosing R when S would be optimal (see model 1 in figure A.7). Here we will investigate whether utility functions of a different functional form could explain choices better. In all specification we assume that decisions are based on expected utilities:

$$EU(S) = u(60)$$

$$EU(R) = u(100) \times Pr(\#R \geqslant \theta) + u(5) \times (1 - Pr(\#R \geqslant \theta))$$

where $u(\cdot)$ is the utility function, $EU(S)$ and $EU(R)$ are the expected utilities of actions S and R and $Pr(\#R \geqslant \theta)$ is the subjective probability assigned to an event that the threshold will be exceeded. To calculate this probability we use elicited probabilistic beliefs $b_i(t)$ and construct a probability distribution over the number of group members who choose R. In these calculations we assume that beliefs are homogeneous (Rapoport, 1985; Rapoport and Eshed-Levy, 1989)[16] and that players expect the actions of other group members to be independently drawn from a binomial distribution with the stated probability. The probability that the threshold will be exceeded is calculated as follows:

$$Pr(\#R \geqslant \theta) = \sum_{k=\theta-1}^{n-1} (b_i(t))^k (1 - b_i(t))^{n-1-k} \binom{n-1}{k}$$

How expected utilities are mapped into choices is determined by a choice rule. We compare the predictions of the following choice rules:

1. **Deterministic choice** (DET). The action with a higher expected utility is always chosen, utility is equal to the monetary payoff.

[16] If separate beliefs were formed, the scoring rule is designed to elicit the expected value of all probabilistic beliefs.

Table 5: Parameter values, estimated by minimizing mean squared deviation (MSD).

Parameter	*DET*	*LOG*	*RISK*	*SOC*	*FAR*	*ALL*
$\hat{\lambda}$	–	11.41	0.02	34.60	12.93	2.82
$\hat{r}$	–	–	-1.85	–	–	-0.43
$\hat{\alpha}$	–	–	–	0.65	–	0
$\hat{h}$	–	–	–	–	6.75	3.60
MSD	242	197.88	181.31	183.83	181.31	180.25

2. **Logistic choice** (*LOG*). Utility is equal to the monetary payoff, but each action can be chosen with positive probability, determined by a logistic choice rule:

$$Pr(R) = \frac{e^{\lambda EU(R)}}{e^{\lambda EU(R)} + e^{\lambda EU(S)}}, \text{ with } \lambda \in [0, \infty)$$

3. **Risk preferences** (*RISK*). In addition to a logistic choice rule, this specification includes a utility function with constant relative risk aversion:

$$u(\pi_i) = \frac{\pi_i^{1-r}}{1-r}$$

4. **Pro-social preferences** (*SOC*). In addition to their own payoff, subjects also care about the average earnings of other group members, weighed by parameter $\alpha \in [0, 1]$:

$$u(\pi_i, \pi_{-i}) = (1-\alpha)\pi_i + \alpha \sum_{j \in N \setminus \{i\}} \frac{\pi_j}{n-1}$$

5. **Farsightedness** (*FAR*). This choice rule assumes that players take into account the payoffs from h_i upcoming rounds, where $h_i \in [1, \infty)$ is the length of the planning horizon.[17] We assume that farsighted players expect all other group members to use the *DET* choice rule, with beliefs formed using weighted fictitious play. The ability to anticipate the reaction of others allows players to choose a best-response path to the actions of other players in every future round. The expected utility of an action is calculated by adding the expected utilities from all rounds along this best-response path. Additional details of the *FAR* model are explained in Appendix G.
6. **All preferences and farsightedness** (*ALL*). This specification combines logistic choice with farsightedness, risk and pro-social preferences.

We estimate the parameters of the five choice rules with free parameters by minimizing the mean squared deviation (MSD), defined following Camerer and Ho (1999) and Nyarko and Schotter (2002):

[17] Players know the number of rounds in a given block, thus they cannot take into account the earnings from rounds that will not be played. Therefore the actual number of rounds taken into account is $\min\{h_i, T-t+1\}$, where T is the number of rounds in a block and t is the current round.

$$MSD = \sum_{i=1}^{72} \sum_{t=2}^{33} \left(x_i^t - \hat{Pr}(R)_i^t \right)^2$$

where x_i^t is equal to 1 if player i chose R in round t and equals 0 otherwise and $\hat{Pr}(R)_i^t$ is the probability that player i will choose R in round t, estimated by a model.

Table 5 shows the estimated parameter values and the goodness of fit for each choice rule. The model that combines all preferences and farsightedness fits data best, although specifications with only risk preferences or farsightedness fit almost equally well with a lower number of parameters. Figure A.7 in Appendix A compares experimental data with the predictions. All choice rules fit data well at the end of each block, but models DET and LOG under-predict R choice following coordination failure in LOW at the start of block 2 and all models under-predict R choice in LOW at the start of block 3.

The finding that a model which combines all preferences fits data best is not surprising as models are not penalized for having more free parameters. To account for the different numbers of parameters we calculate the Akaike information criterion (AIC) and Bayesian information criterion (BIC). The calculation of these information criteria require optimization by log-likelihood maximization instead of MSD minimization. Results of log-likelihood estimation, BIC and AIC values are reported in Appendix E.

So far we have been evaluating the in-sample fit of the belief learning model with various choice rules. Next, we want to test whether these models can explain the convergence patterns and heterogeneity in success rates between groups. In particular, we found that groups with higher observed R levels in round 14 overcome coordination failure more often (columns 1-3 in table 6). To test if learning models can explain these findings, we calculate beliefs in round 14 using estimated population-level fictitious play parameters and observed feedback in rounds 1-13, averaged across all players. Then we use each choice rule with weighted fictitious play to simulate the path of beliefs and choices in rounds 15-28 for each 6-person group, without using any empirical belief or choice data from these rounds. We run 1000 simulations for each model and count the frequency of simulations in which coordination failure is overcome, which is said to occur if in round 28 all group members hold beliefs under which R is a deterministic best-response.

Table 6 compares the results of simulations with experimental data. A model would predict data perfectly if 100% of simulations with groups #8, #2 and #5 would converge to an efficient equilibrium and 0% would do so for all other groups. No model does as good, but most predict higher success rates for groups that actually overcome coordination failure in the experiment. In particular, $RISK$ and ALL stand out, with a predicted gap of at least about 60 percentage points between the two clusters of groups. SOC and FAR do somewhat less well as SOC under-predicts the success of groups that do transition to the efficient equilibrium, while FAR over-predicts the success of groups that do not transition. DET and LOG predict that almost no group will overcome coordination failure, regardless of the choices in round 14.

Results of the simulations show that the correlation between final outcomes and initial choices that is observed in our experiment can be well explained by belief learning models

Table 6: Frequency of transitions to the efficient equilibrium in simulations, compared to the experiment. In simulations the percentage indicates the fraction of iterations in which R is the deterministic best-response to beliefs in round 28. In the experiment a transition occurs if all group members choose action R in round 28. Groups are sorted by observed R levels in round 14.

	Experiment		Simulations					
Group ID	(R $\wedge$ D) in 14	Transition?	*DET*	*LOG*	*RISK*	*SOC*	*FAR*	*ALL*
8	5	Yes	0.0%	0.1%	100.0%	99.4%	91.3%	99.6%
2	4	Yes	0.0%	0.0%	81.3%	58.3%	88.3%	83.1%
5	4	Yes	0.0%	0.0%	84.9%	60.8%	87.6%	86.8%
4	3	No	0.0%	0.0%	26.8%	11.5%	48.7%	29.8%
7	3	No	0.0%	0.0%	17.3%	7.1%	37.6%	20.5%
11	2	No	0.0%	0.0%	0.9%	0.3%	2.8%	1.2%
1	1	No	0.0%	0.0%	0.3%	0.2%	0.2%	0.2%
6	1	No	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
10	1	No	0.0%	0.0%	0.4%	0.0%	0.3%	0.2%
3	0	No	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
9	0	No	0.0%	0.0%	0.3%	0.1%	0.1%	0.2%
12	0	No	0.0%	0.0%	0.0%	0.0%	0.1%	0.0%

with risk preferences, social preferences or farsightedness. This is, of course, not the only possible explanation, and we review other possibilities in Appendix D. In short, we look at possibilities that players in the two treatments have different personal characteristics, experience different histories in block 1, or that the adaptation process is driven by reinforcement learning or reciprocity. Each of these explanations could potentially explain the treatment difference, but we show that none of them can fully explain experimental data.

3.3. Individual Heterogeneity in LOW

In *HIGH* all groups converge to an inefficient equilibrium, but in *LOW* there is heterogeneity, which is explored in this section. In particular, we have shown that whether a group in *LOW* overcomes coordination failure depends on the number of group members who choose R and disclose it in round 14 (result 4). We will therefore test whether R is more often chosen by players who are more risk seeking, pro-social or farsighted, as predicted by the theory (see Appendix H). Before showing the results, we will explain the tasks that we used in the lab to measure each of these characteristics. Each task was incentivised and these earnings were added to the earnings from the main part of the experiment.

3.3.1. Elicitation of Risk and Social Preferences

Risk preferences were elicited using a multiple price list (Holt and Laury, 2002), and the number of safe choices was used as a measure of risk aversion.[18] Note that objective risk was not present in the experiment and uncertainty about payoffs originated only from

[18] 5 subjects switched more than once, therefore results are slightly different if risk aversion is measured using the first switching round. In such case the coefficient of farsightedness turns significant at a 10% level in model 3 in table 9.

strategic uncertainty. But if players form probabilistic beliefs, a utility function with a lower Arrow-Pratt measure of risk aversion would make R myopically optimal for a wider range of beliefs (see Appendix H.3 for a proof).

Social preferences were elicited using a Social Value Orientation (SVO) Slider Measure (Murphy et al., 2011), following a z-Tree implementation by Crosetto et al. (2012). In this measure subjects complete six allocation tasks[19] between themselves and some other participant, who was randomly selected from a different group to prevent meeting former group members. A graphical representation of all tasks is shown in figure I.8, Appendix I. At the end of the experiment one task was chosen for payment and subjects received either the amount that they sent or the amount that was sent to them by the other participant. Answers to each task were used to construct a continuous, uni-dimensional scale of SVO, known as the SVO angle, using the following formula:

$$\text{SVO angle} = \arctan\left(\frac{\sum_{t\in\{1,..,6\}}(P_t^O - 50)}{\sum_{t\in\{1,..,6\}}(P_t^S - 50)}\right)$$

where P_t^S is the amount allocated to oneself in task t and P_t^O is the amount allocated to the other person. SVO angle ranges from -16.26° for perfectly competitive individuals to 61.39° for perfectly altruistic individuals.

3.3.2. Elicitation of Farsightedness

We define farsightedness by the length of the planning horizon: myopic players take into account only the earnings in the current round, while farsighted players anticipate future rounds. The task used to measure farsightedness was therefore selected to identify whether players anticipate how their decisions made in the current round will affect outcomes in future rounds. Importantly, we chose to measure farsightedness in an individual choice task, in which players had to anticipate their own future actions.

The task was adapted from Bone et al. (2009) and presented as a decision tree with two decision nodes and two chance nodes (see figures K.14 - K.16 in Appendix K). In the second decision node one move dominates the other, making it trivial to choose correctly (97% of players chose the dominant move, both in Bone et al., 2009, and in our experiment). In the first node one decision also dominates the other, but only if subjects anticipate their own choices in the second node. Payoffs were set up such that players who ignore the payoffs that will not be reached would choose one move (we call it "correct move" as it would maximise payoffs for 97% of the players) while those who treat the second decision node as a chance node, expecting all final payoffs to be equally likely, would choose a different move.

To classify subjects more precisely we designed three tasks with different difficulty

[19]We added a seventh task, completed after the six original ones. This task tested whether subjects were willing to give up their earnings to reduce the level of advantageous inequality (see footnote 10). This additional task was not used to calculate the SVO angle.

levels,[20] which increased the frequency of correct moves from 21% in the hard task to 42% in the intermediate task and 65% in the easy task. Each subject completed all three tasks in the order of decreasing difficulty and one task was randomly chosen for payment. 88% of subjects chose the correct move in an easier task if they had chosen the correct move in a more difficult task, conforming with the Guttman scale (Guttman, 1950). To more accurately classify subjects who did not conform with the Guttman scale, the farsightedness score was calculated the following way:

- Score of 0 for players who failed in the easy task
- Score of 1 for players who solved the easy task but failed in the intermediate task
- Score of 2 for players who solved the easy and intermediate tasks but failed in the hard task
- Score of 3 for players who solved all three tasks

Instead of the length of the planning horizon, we could have measured other aspects of farsightedness. For example, farsighted players would need to be able to anticipate the reactions of other players, which could be measured in a two-player version of the game used by Bone et al. (2009), by a sequential game such as a centipede game (McKelvey and Palfrey, 1992) or a race game (Gneezy et al., 2010), or even a static dominance-solvable game (Johnson et al., 2002). However, in games with strategic uncertainty choices may differ from theoretical predictions for reasons other than the lack of farsightedness: for example, subjects may be pro-social or doubt the rationality of others. Correlation between choices in these games and behavior in the main part of the experiment could therefore be influenced by aspects other than farsightedness, such as the ability to form accurate beliefs. In an individual choice task decisions should not be affected by distributional concerns or beliefs about rationality, so we are more confident that it is the ability to plan ahead that is being measured.

Additionally, we evaluated farsightedness by the process data, in particular by the response times in the two decision nodes. Recently there has been an increasing use of procedural variables such as data of eye movements (Arieli et al., 2011; Polonio et al., 2015) or patterns of information search (Costa-Gomes et al., 2001; Johnson et al., 2002) to

[20] We manipulated the task difficulty by changing payoff differences. Two types of payoff differences are present in each task: the actual difference in expected payoffs in favour of the correct move if the second node decision was taken into account ("farsighted payoff difference"), and the difference in the average of all payoffs in favour of the incorrect move if all payoffs were treated as being equally likely ("myopic payoff difference"). The "hard task" was identical to one of the tasks used by Bone et al. (2009), with a farsighted payoff difference of 125 ECU and the myopic payoff difference of 138 ECU. In the "intermediate task" the correct choice was made more attractive by increasing the farsighted payoff difference to 175 ECU and at the same time reducing the myopic payoff difference to 125 ECU. In the "easy task" the distribution of payoffs was the same for both actions, reducing the myopic payoff difference to 0, but the payoffs were arranged in a way that increased the farsighted payoff difference to 300 ECU.

supplement choice data and improve the precision of the classification procedure. Response times in particular have been widely used for both classification and prediction (Rubinstein, 2007; Piovesan and Wengström, 2009, review in Spiliopoulos and Ortmann, 2016) and there is evidence that response times are correlated with strategic behavior (Rubinstein, 2016).

We hypothesize that response times reflect the intention to plan ahead: players who plan ahead need to make decisions in all possible second nodes before making the first node decision, therefore they should spend more time in the first node and less in the second one, while players who do not plan ahead would make a quick choice in the first node and spend more time in the second node. The second measure of farsightedness was therefore calculated as a ratio of time spent in the first node relative to the second node. The farsightedness measure based on response times has several advantages over the measure based on correct answers. The time ratio is a continuous variable, therefore subjects can be classified more accurately. The number of correct answers may also be confounded by other traits, such as intelligence, while response times capture the willingness to think ahead, regardless of the success of such efforts. It is reassuring to know that both measures are closely related: participants who solve more tasks correctly on average spend more time in the first decision node (table 7). This finding not only validates the use of the farsightedness measure, but provides additional support that response times can be used to identify strategic behavior (Rubinstein, 2016). As the next section will show, players who spend more time in the first node also more often initiate transitions to the efficient equilibrium in the main part of the experiment.

Table 7: Comparison of two farsightedness measures: farsightedness score is based on the number of correct answers, time ratio is calculated as the ratio of time spent in node 1 relative to time spent in node 2.

Farsightedness score	Average time ratio	N
0	1.20	25
1	1.35	21
2	1.91	18
3	2.50	8
Total	1.57	72

3.3.3. Results: Personal Characteristics and R Choice in LOW

We identify the types of players who are more likely to choose R by comparing the characteristics of subjects who choose R and then running a probit regression with four specifications. In the first two specifications the dependent variable is set to 1 for players who chose R and disclosed it in round 14 and is set to 0 otherwise. We favour this specification because observed R levels in round 14 accurately predict transition success, but it suffers from a problem of independence, as observations in round 14 are affected by the shared history in block 1. To account for the different histories we add the total number of observed R and S in block 1 as control variables. Furthermore, we replicate the analysis in models 3 and 4 using choices from round 1, which are independent. The difference between model 1 and 2 and between model 3 and 4 lies in the measurement of farsightedness: *farsighted score* is based on the number of correctly solved individual choice tasks, while *time ratio* measures time spent in the first decision node relative to time spent in the second decision node.

Table 8: Average scores of subjects who choose R and S in rounds 1 and 14, *LOW* treatment. Statistical significance of the difference between the participants who choose S and those who choose R is measured using a Mann-Whitney U test.

	Round 1		Round 14	
	S	R	S	R
Farsightedness: score	0.75	1.25	0.78	1.28
Farsightedness: time ratio	1.53	1.69	1.34*	1.89*
Risk aversion	6.38	5.60	6.06	5.83
Pro-social attitudes	20.75	24.52	23.34	22.35
Beliefs	30.75***	66.80***	35.44**	52.67**
Average income	63.54	67.99	62.68	69.35
N	16.00	20.00	18.00	18.00

* p<0.10, ** p<0.05, *** p<0.01

Both the average scores in table 8 and the probit regressions in table 9 show that players who are more farsighted and who report higher beliefs are more likely to choose R in rounds 1 and 14. Table 8 shows a large difference in the magnitude of farsightedness between players who choose R and S, although Mann-Whitney U test fails to reject the null hypothesis when the farsightedness score is used. Table 9 reports the marginal effects at the means of covariates for the four probit models. Farsightedness is positively related to the tendency to choose R and disclose, and the significance level depends on the dependent variable and on the method used to measure farsightedness (p-values range from 0.002 to 0.131). In addition to statistical significance, we find a large economic effect: the predicted likelihood to choose R and disclose in round 14 is 30% for subjects with a farsightedness score of 0 and increases to 81% for those with a score of 3, holding other factors constant at their mean (Wald test shows that the difference is statistically significant, p-value = 0.0324). The difference between subjects with a score of 0 and those with scores 1 or 2 is of expected sign, but not significant at the 5% level. The ratio of decision times predicts choices better than the number of correct answers, as is evident from higher significance levels and a higher pseudo-R^2. The coefficient of stated beliefs is statistically significant in all specifications except for model 1.

Result 6. *Participants in LOW who are classified as more farsighted are significantly more likely to choose R and disclose this action in rounds 1 and 14.*

Table 9: Probit models, dependent variable equals one if R is chosen and disclosed. Controls for gender and age. Standard errors are heteroskedasticity-robust.

	Round 14		Round 1	
	(1)	(2)	(3)	(4)
Farsighted score	0.153*		0.147	
	(1.67)		(1.62)	
Time ratio		0.289***		0.249***
		(3.08)		(2.86)
Risk aversion	-0.00779	-0.0233	-0.0835	-0.0838
	(-0.12)	(-0.34)	(-1.24)	(-1.13)
Pro-social attitudes	0.00279	0.00334	0.0106	0.0104
	(0.46)	(0.55)	(1.28)	(1.29)
Beliefs	0.00673	0.0117**	0.0104***	0.0120***
	(1.46)	(2.36)	(3.38)	(3.85)
Observed R	0.0448	0.0437		
	(1.40)	(1.40)		
Observed S	-0.0373	-0.0567		
	(-0.90)	(-1.34)		
Constant				
N	36	36	36	36
pseudo R^2	0.174	0.264	0.300	0.355

Marginal effects; t statistics in parentheses

* p<0.10, ** p<0.05, *** p<0.01

4. Discussion: Role of Strategic Motives

One of the goals of this paper is to investigate whether deviations from an inefficient convention could be explained by strategic motives. This section will briefly review our main findings, which are mostly consistent with the strategic hypothesis.

We find that many players are willing to pay to make their actions observable, and those who choose R are much more likely to do so (result 3). Strategic motives seem to be the most likely explanation for these findings because action disclosure provides no monetary benefits in the current round. If disclosures were motivated by non-strategic motives, the frequency of disclosure should not decrease over time and it would be hard to explain why less than 20% of players disclose their actions at the end of block 2 when it costs 2 ECU while 40% disclose at the start of block 2 when it costs 80 ECU (see figure 4). Some disclosures might be a result of mistakes, but a large gap between disclosures of R and S and the concentration of disclosures at the start of each block indicate a systematic pattern rather than random errors. We also find that R levels and disclosures of R are higher and transitions to the efficient equilibrium are more frequent in *LOW* (result 1 and result 2), just as predicted by the learning model with sophistication. Evidence for belief learning (result 5) and a strong correlation between observed R levels in round 14 and revolt success (result 4) provide further support that choosing R following coordination on S could be a part of a payoff-maximizing strategy if future earnings are taken into account. In fact, players in *LOW* who choose R and disclose it in round 14 on average earn 69 ECU per round, compared to 63 ECU earned by other players. In *HIGH*, however, strategic teaching is not profitable: those who choose R and disclose on average make only 42 ECU, compared to 54 ECU made by others. Finally, we find that in the low cost treatment players with a longer planning horizon are more likely to choose R and disclose this action, while the effect of social and risk preferences is not significant (result 6).

However, even as players seem to understand the benefits of having their actions observed, there are indications that the ability to make inter-temporal trade-offs is limited. First, higher action disclosure costs discourage action disclosure, but have only a small direct effect on R levels in round 14. A sophisticated player would be expected to compare the payoffs of choosing R and disclosing with the payoffs of choosing S and not disclosing, and choose R if the former is larger than the latter. Consequently, higher costs should reduce not only the tendency to disclose, but also the tendency to choose R. Yet we find that in round 14 higher costs discourage only action disclosure, but not the tendency to choose R. Thus low disclosure costs help overcome coordination failure mostly through higher observed R levels that alter beliefs and therefore choices in subsequent rounds. The second issue is that at the start of the second block 25% of players in *HIGH* choose R and do not disclose their actions. It may seem that sophisticated players should either choose R and disclose, or choose S and not disclose. One explanation for choosing R and not disclosing could be that such choices are driven by non-strategic motives, such as beliefs or preferences. But sophisticated players could also choose R without disclosing even if R was not myopically optimal, as long as a sufficiently large probability is assigned to the event that exactly $\theta - 1$ other group members will choose R. In such an event one's decision would be pivotal, and a choice of R would raise R levels above the threshold, informing other players about the true R levels via their payoffs.

5. Concluding Remarks

This paper addresses the social dilemma of inefficient conventions: a unilateral deviation from an inefficient convention is costly, even though a collective deviation would lead to a Pareto improvement. We are interested in how this improvement can be achieved: what features of the environment facilitate the diffusion of efficient conventions and what motivates individuals to deviate? We address these questions using a laboratory experiment that tests whether deviations from an inefficient convention are explained by strategic motives and whether such deviations occur more often when players can disclose their actions at a lower cost. We find support for the hypothesis that deviations are motivated by strategic reasons and we also find that inefficient conventions are overcome only if action disclosure costs are low.

The problem that motivated the design of our experiment was the failure to take political action, and our findings suggest several ways of how this problem could be solved. Availability of information technology could make it easier to inform others about one's action both directly, through social media, and indirectly, through the information shared by other parties. The physical location of events might matter too, as a large central space should make it easier to gauge the level of support, while also allowing the information to be shared more rapidly between the participants. However, before implementing any such interventions more research has to be done to verify the external validity of our findings and to understand the potential negative effects. For example, advances in information technology could crowd out physical protests or make it easier for the regime to use propaganda, censorship or surveillance (Morozov, 2012).

Our findings could also be applied to other settings with persistent inefficient conventions. An inefficient technology could become the standard if it exhibits steep returns to scale (Arthur, 1989): examples of such technological lock-in include QWERTY keyboard (David, 1985), inferior information technology (Shapiro and Varian, 1999) or light water nuclear reactors (Cowan, 1990). As older technologies are constantly being superseded by new innovations, explicit regulation of standards becomes too costly; instead, it would be desirable to design a decentralised system that would facilitate the diffusion of the most efficient technology. Designing the environment in the right way could also help overcome coordination failures that occur when groups adopt inefficient social customs (Akerlof, 1980) or when societies end up with inefficient economic and political institutions (North, 1990; Acemoglu, 2006).

Appendix A. Figures

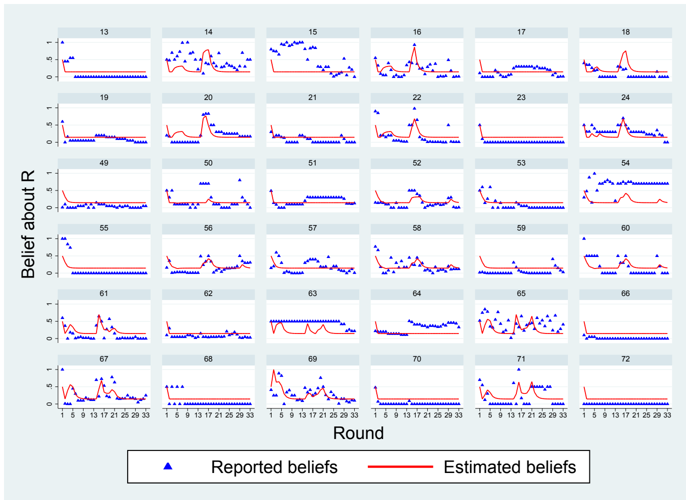

Figure A.5: Comparison of stated beliefs and estimated beliefs using a weighted fictitious play model with parameters that minimise MSD. Graphs by subject, *HIGH* treatment.

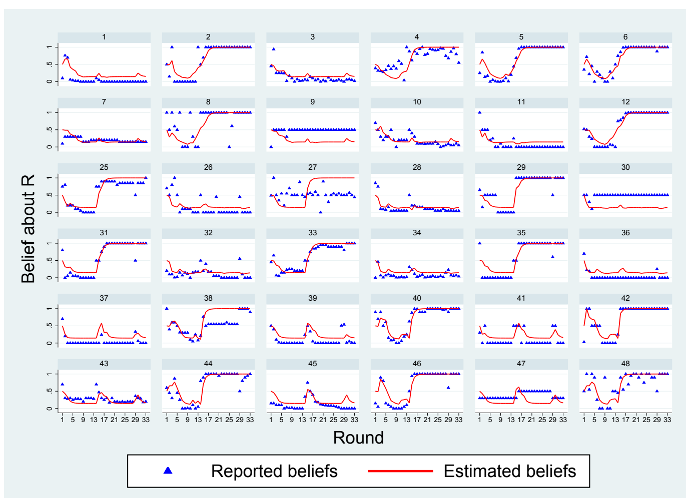

Figure A.6: Comparison of stated beliefs and estimated beliefs using a weighted fictitious play model with parameters that minimise MSD. Graphs by subject, *LOW* treatment.

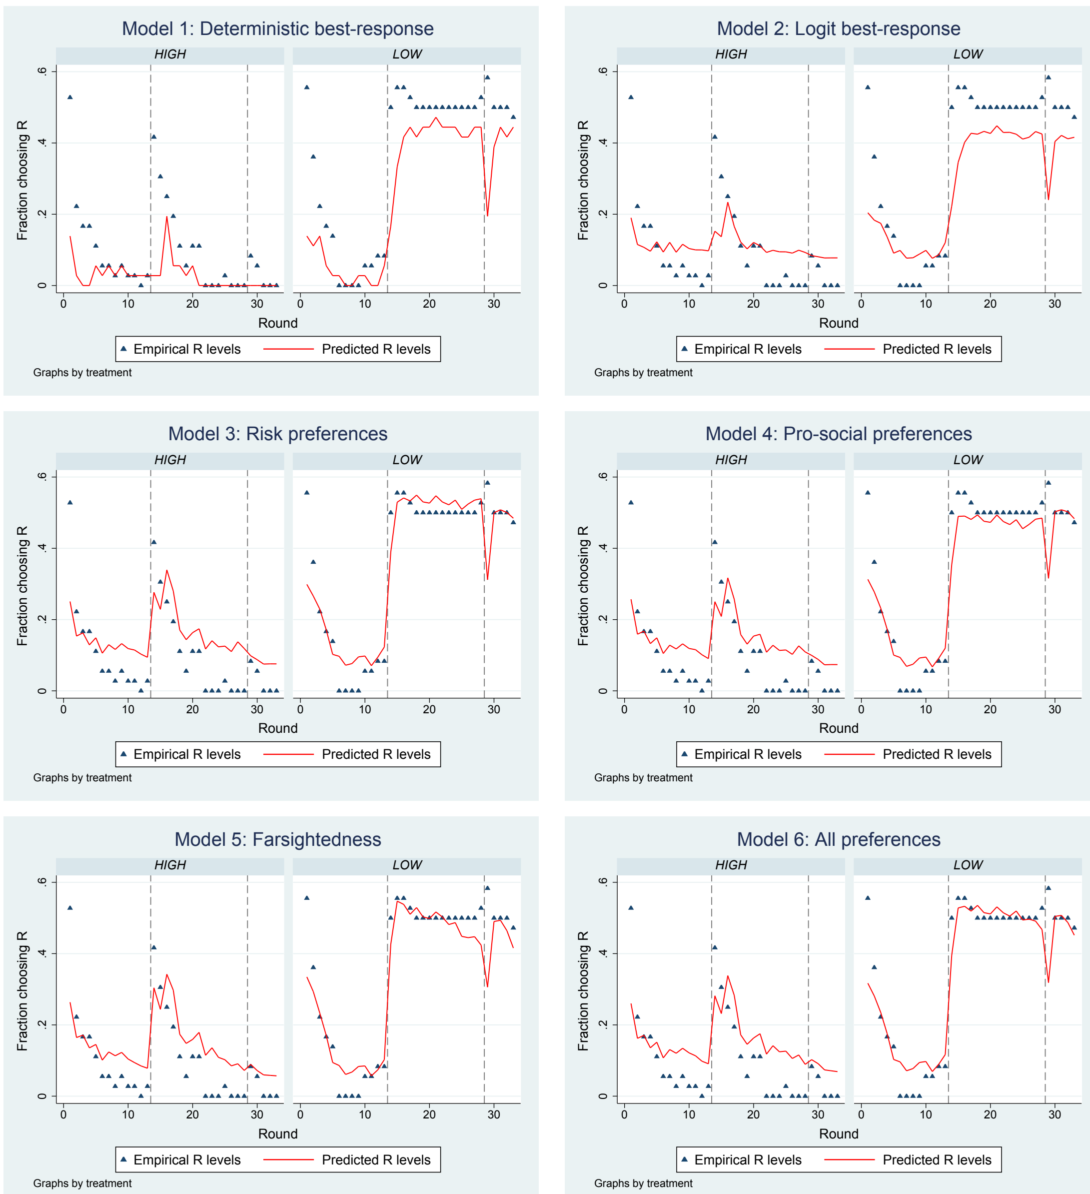

Figure A.7: Comparison of empirical R levels and R levels predicted by each of the six choice rules.

Appendix B. Tables

Table B.10: Estimated treatment effect by round, calculated using the estimates of a multilevel GLS random effects regression and group level data. Standard errors clustered on group level.

Round	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
Coefficient	0.170	0.161	0.152	0.143	0.134	0.125	0.115	0.106	0.097	0.088	0.079	0.070	0.060	1.429	1.573	1.716
p-value	0.782	0.773	0.761	0.748	0.732	0.713	0.691	0.668	0.652	0.655	0.689	0.747	0.807	0.228	0.183	0.146
17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33
1.860	2.003	2.146	2.290	2.433	2.577	2.720	2.864	3.007	3.151	3.294	3.438	2.900	2.900	2.900	2.900	2.900
0.116	0.092	0.073	0.058	0.047	0.038	0.031	0.026	0.022	0.019	0.016	0.015	0.017	0.019	0.021	0.024	0.028

Table B.11: Time trends: estimated p-values for the coefficient of the number of rounds left.

Block	*HIGH* treatment	*LOW* treatment
1	0.0001	0.0000
2	0.0037	0.7351
3	0.0059	0.0410

Table B.12: Random-effects panel data probit regression (12 clusters). Dependent variable is equal to 1 if the action is disclosed.

	All data	Decisions after choosing R
Treatment = *LOW*	1.128***	1.227***
	(5.25)	(3.33)
Action = R	1.742***	
	(15.48)	
Round	-0.0517***	-0.0554***
	(-9.55)	(-7.63)
Constant	-2.011***	-0.192
	(-10.65)	(-0.70)
Observations	2376	545

t statistics in parentheses

* p<0.10, ** p<0.05, *** p<0.01

Table B.13: Three-level nested random effects model and two-level random effects GLS model (12 clusters). Dependent variable: stated probability that a randomly chosen group member chose R (from 0 to 100). Standard errors are robust to correlation within the observations of each group. Observations from rounds in which groups are in equilibrium are excluded (a group is in equilibrium if no player observed two different actions being chosen in the previous round)

	Random effects GLS		Nested random effects	
	Two-period lags	One-period lags	Two-period lags	One-period lags
Observed R in (t-1)	6.723***	6.815***	7.074***	7.051***
	(4.96)	(3.97)	(6.13)	(3.87)
Observed R in (t-2)	0.772		1.447	
	(0.36)		(0.66)	
Observed S in (t-1)	-3.950*	-3.019	-4.591**	-3.577*
	(-1.87)	(-1.36)	(-2.40)	(-1.91)
Observed S in (t-2)	1.110		1.083	
	(0.36)		(0.36)	
Actual R level in (t-1)	5.120**	6.680***	4.767***	6.508***
	(2.40)	(5.67)	(3.37)	(6.52)
Actual R level in (t-2)	1.857		1.231	
	(0.64)		(0.48)	
Treatment: 1 = *LOW*	-6.596*	-1.813	-6.148*	-1.633
0 = *HIGH*	(-1.66)	(-0.44)	(-1.89)	(-0.40)
Constant	17.51***	16.03***	17.51***	16.09***
	(3.03)	(3.22)	(3.44)	(3.37)
Observations	192	270	192	270

t statistics in parentheses

* p<0.10, ** p<0.05, *** p<0.01

Appendix C. Choice of the Game Parameters

Duration of blocks. The amount of time that participants can spend in the lab limits the number of rounds that can be implemented. We wanted each block to contain enough rounds to observe convergence, and assigned more rounds to the more important blocks of the experiment. The main question of the study is whether coordination failure is overcome in the second block, therefore the second block was the longest, with 15 rounds. In the first block we also wanted to observe convergence in actions and in beliefs, and a pilot study with 10 rounds in the first block suggested that beliefs are still being adapted at the end of the block. Therefore the number of rounds in the first block was increased to 13. Additionally, we did not want the number of rounds in the first block to be a multiple of 5, to be able to run comparable treatments in which players would not be informed about the block length. The third block is relatively less important than the first two, and we also anticipated that players will have understood the game well enough to justify a shorter duration, therefore it was set to 5 rounds. Results show that the duration was sufficiently long to observe convergence in each of the three blocks, as most groups are in equilibrium by the end of each block (see figure 2).

Disclosure costs. Disclosure costs were set to generate sufficient variation in the behavior of farsighted players. We assume that farsighted players compare the sum of payoffs with no strategic teaching (then the group remains in the inefficient equilibrium and payoffs are equal to 60 ECU in each of the 15 periods) to the payoffs if strategic teaching is used (strategic players choose and disclose R in the first t periods and receive a payoff of 5 ECU because of miscoordination, but efficient coordination is implemented in all periods after t, generating a payoff of 100 ECU). In addition, a cost of 1 ECU is paid when actions are not disclosed, and a cost of c ECU is paid when actions are disclosed.

Overall, the relationship between the disclosure cost (c) and the maximum number of periods that a player is willing to teach (t) must make players indifferent between using strategic teaching and not using it:

$$t(5-c)+(15-t)(100-1)=15(60-1)$$

It is straightforward to calculate that strategic players are willing to engage in strategic teaching if they expect a transition to the efficient equilibrium in $\frac{600}{c+94}$ rounds or less. The relationship between disclosure costs and the discrete values of the maximum teaching length is shown in table C.14. Higher disclosure costs mean that a transition to the efficient equilibrium must occur earlier to justify strategic teaching. In *LOW* the disclosure cost was set to 2 ECU, with which strategic players would be willing to teach for at most 6 periods. To ensure sufficient differences between treatments, the disclosure cost in *HIGH* was set to 80 ECU, which implies the maximum teaching length of 3 periods.

Table C.14: The maximum number rounds that strategic players would be willing to teach, for a range of disclosure cost values.

Disclosure cost (in ECU)	0-6	6-26	26-56	56-106	106-206	206-506
Max. rounds of strategic teaching	6	5	4	3	2	1

Appendix D. Alternative Explanations for the Treatment Difference

Treatment	Farsightedness score	Risk preference	SVO angle	Female	Age
High	1.22	5.92	16.85	0.58	21.06
Low	1.03	5.94	22.85	0.47	20.97
Range	0–3	2–10	-7.82 – 53.37	0–1	18–29

Table D.15: Descriptive statistics and their range by treatment

First, participants in the two treatments could differ in terms of their personal characteristics, which would lead to different responses to beliefs. To evaluate this possibility we compare the levels of farsightedness, risk and social preferences and demographic data, which were elicited at the end of the experiment (see section 3.3.1 for details). Table D.15 shows that most characteristics are similar, except for the SVO angle that measures pro-social preferences: subjects in *LOW* are more pro-social and the difference is marginally significant (Mann-Whitney U test p-value = 0.0615). If more pro-social participants were assigned to LOW, this difference could potentially explain the treatment effect because pro-social players should be more likely to choose R (see Appendix H.4). However, differences in social preferences should increase R levels in all rounds, while we find that the treatment difference is very small at the start of block 2, but grows over time. Alternatively, choices in the SVO task could have been influenced by outcomes in the main part of the experiment. Social preference elicitation task was conducted at the end of the experiment and subjects from groups that coordinated on the efficient equilibrium might have acted in a more pro-social way, for motives such as reciprocity. Note that differences in social preferences cannot be explained by direct reciprocity, as subjects were never matched with their group members from the main part of the experiment, but they might be explained by generalised reciprocity (see Stanca, 2009, for evidence of generalised reciprocity).

Second, the process that drives the treatment difference could be reinforcement learning instead of belief learning. If this was the case, separate attractions could be formed for each of the four strategies in a stage game: 'choose R and disclose it' (R, D), 'choose R and not disclose' (R, ND), 'choose S and disclose' (S, D) and 'choose S and not disclose' (S, ND). If at the start of block 2 there are no treatment differences between the frequencies with which each strategy is chosen, payoffs to players who chose (R, D) will be lower in *HIGH* than in *LOW*. Under reinforcement learning the attraction of (R, D) would be reduced more strongly in *HIGH* than in *LOW*. If few players choose (S, D), this would reduce R levels in *HIGH* more than in *LOW*. However, there is no evidence that large losses discourage subsequent R choice in *HIGH* at the start of block 2: players who choose (R, D) in round 14 make -75 ECU in that round, yet 83% of them continue choosing R in round 15; in comparison, those who choose (R, ND) earn 4 ECU, but only 22% of them choose R in round 15.

Another process that could explain the correlation between observed feedback and choices without the need to form beliefs is reciprocity: players who observe many group members choosing R could receive higher payoffs and they would be willing to reciprocate by choosing R more often in subsequent rounds. However, this explanation seems unlikely for several reasons. First, subjects are always informed about their payoffs, so observed R levels provide no additional information about one's earnings. Thus reciprocity should be induced by the actual R levels rather than the observed R levels. Second, R levels would affect earnings only if the participation threshold has been exceeded, and therefore there should be no difference in responses to higher observed R levels as long as the threshold has not been exceeded. Instead, we find that the effect of observed R levels is continuous: observing one additional group member choosing R has a positive effect, regardless of the R levels.[21] Third, reciprocity should be driven by subjects who chose R in the previous round because the payoffs of those who chose S are always fixed. In the probit regression we include a variable capturing the interaction between choosing R in the previous round and the observed R levels and find that the estimated coefficient is small and not significant (p = 0.973), suggesting that those who chose R and those who chose S are similarly affected by the observed history.

[21] p-values for the different numbers of observed R are: 0.008 (1 vs 0), 0.150 (2 vs 1), 0.068 (3 vs 2), 0.825 (4 vs 3), 0.000 (5 vs 4)

Appendix E. Model fitting by log-likelihood maximization

In this section we replicate the model fitting exercise from section 3.2.4, but we search for parameter values that maximize log-likelihood rather than minimize mean squared deviation. The upside of log-likelihood estimates is the possibility to compare the goodness of fit using the Akaike information criterion (AIC) and Bayesian information criterion (BIC), which penalize models for the number of parameters.[22]

The log-likelihood function takes the following form:

$$\log \mathcal{L}(\lambda, r, \alpha, h) = \sum_{i=1}^{72} \sum_{t=2}^{33} \log \left(1 - \sqrt{(x_i^t - \hat{Pr}(R)_i^t)^2} \right)$$

where x_i^t is equal to 1 if player i chose R in round t and equals 0 otherwise and $\hat{Pr}(R)_i^t$ is the probability that player i will choose R in round t, estimated by a model.

Both MSD and log-likelihood favor models with accurate predictions, but log-likelihood is less favorable to models with bold predictions because the function $\log(1 - \sqrt{\cdot})$ grows to minus infinity as the argument approaches 1. For this reason log-likelihood cannot be calculated for the *DET* model.

Table E.16: Parameter values, estimated by maximizing log-likelihood.

Parameter	*DET*	*LOG*	*RISK*	*SOC*	*FAR*	*ALL*
$\hat{\lambda}$	–	12.17	0.59	18.79	12.28	0.59
$\hat{r}$	–	–	-0.86	–	–	-0.86
$\hat{\alpha}$	–	–	–	0.35	–	0
$\hat{h}$	–	–	–	–	2.59	1
LL	–	-697.42	-651.15	-675.31	-669.38	-651.15
AIC	–	1396	1306	1355	1343	1310
BIC	–	1403	1318	1366	1354	1333

Table E.16 shows the estimated parameter values for each choice rule. The model with risk preferences does best regardless of the measure: AIC, BIC and even log-likelihood. The addition of pro-social preferences and farsightedness in the *ALL* model does not improve log-likelihood at all, and consequently has higher AIC and BIC values.

Table E.17 shows the results of simulations with parameters from table E.16. A comparison with the results of simulations performed using parameters estimated by minimizing MSD (table 6) shows that models with parameters estimated by MSD minimization simulate choices that are much closer to experimental data, compared to models estimated

[22] AIC $= -2\log(\mathcal{L}) + 2k$, BIC $= -2\log(\mathcal{L}) + k\log(N)$, where k is the model degrees of freedom and N is the number of observations.

Table E.17: Frequency of transitions to the efficient equilibrium in simulations, compared to the experiment. Parameters for the simulations were estimated by log-likelihood maximization. In simulations the percentage indicates the fraction of iterations in which R is the deterministic best-response to beliefs in round 28. In the experiment a transition occurs if all group members choose action R in round 28. Groups are sorted by observed R levels in round 14.

	Experiment		Simulations					
Group ID	(R $\wedge$ D) in 14	Transition?	*DET*	*LOG*	*RISK*	*SOC*	*FAR*	*ALL*
8	5	Yes	0.0%	0.1%	74.9%	18.6%	50.3%	74.9%
2	4	Yes	0.0%	0.0%	8.0%	1.4%	5.0%	8.0%
5	4	Yes	0.0%	0.0%	7.3%	1.0%	5.1%	7.3%
4	3	No	0.0%	0.0%	0.7%	0.0%	0.3%	0.7%
7	3	No	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
11	2	No	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
1	1	No	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
6	1	No	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
10	1	No	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
3	0	No	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
9	0	No	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
12	0	No	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%

by log-likelihood maximization. All models fail to predict transitions to the efficient equilibrium for all groups except for group #8, and even for this group only *RISK* and *ALL* models predict transitions sufficiently often.

Appendix F. Belief Elicitation Task

The belief elicitation task was designed to elicit precise and truthful probabilistic beliefs about the actions of other players. Screenshots of the decision screen are reproduced in figures K.12 and K.13. Instructions are reproduced in Appendix J.

To get a precise measure of beliefs we allowed subjects to choose any number between 0% to 100%, but elicitation was performed in two stages. The first stage displayed only the outcomes for all multiples of 5 between 0 and 100, and the second stage displayed the outcomes for the 11 strategies closest to the one chosen in the first stage.

To avoid hedging within one round, subjects were paid either their accumulated earnings from 33 rounds of the main part of the experiment, or they were paid for the belief elicitation tasks. In particular, at the end of the experiment a random draw was performed, so that with 80% probability a subject was paid for the main part of the experiment and with 20% probability for the belief elicitation task (Blanco et al., 2010, used a similar procedure to avoid hedging). The chances to be paid for the belief elicitation task were deliberately kept low to further reduce the temptation to hedge. We used a random-lottery incentive scheme by choosing one round for payment at the end of the experiment, a procedure that has been shown to avoid the income effect that appears if earnings are added up (Lee, 2008).

The second issue that may hinder the elicitation of truthful beliefs is risk preference: for example, risk-averse subjects may avoid reporting beliefs that are too extreme to lower the risk of receiving a low payment. It has been shown that deterministic payment schemes, such as the quadratic scoring rule (McKelvey and Page, 1990; Nyarko and Schotter, 2002), do not elicit truthful beliefs if subjects are not risk neutral (Schlag and van der Weele, 2013). Therefore we used a binarized scoring rule, which should elicit truthful beliefs irrespective of risk attitudes (Schlag and van der Weele, 2013). Experiments have also shown that a binarized scoring rule is less influenced by risk attitudes and is closer to the objective probabilities compared to the quadratic scoring rule (Hossain and Okui, 2013). The binarized scoring rule was implemented by paying subjects in lottery tickets instead of monetary earnings. If the belief elicitation task was chosen for payment, subjects participated in a lottery that yielded either 4000 ECU or 1000 ECU. Better performance in the belief elicitation task therefore did not affect the size of the earnings, but increased the chances to receive the higher reward.

Appendix G. Choice Rule for Models with Farsightedness

In this appendix we will explain how we calculate the utility for a farsighted player, used in models *FAR* and *ALL*. We assume that a farsighted player i is maximizing the sum of earnings over the time period h_i, which we will refer to as a "planning horizon". We treat the planning horizon as a personal characteristic, to be estimated from the data. Of course, the planning horizon cannot be longer than the number of periods that are remaining (which was known in the experiment), thus players add up the earnings from $\min\{h_i, T-t+1\}$ periods, where T is the length of the game and t is the current period.

Just as in other choice rules, farsighted players hold a probabilistic belief about the actions of other players in the current period.

A player with a planning horizon equal to 1 would therefore myopically respond to this belief. Players with a longer planning horizon must also form beliefs about the choices that others will take in future periods. We assume that farsighted players expect that all others are myopic: they form beliefs according to weighted fictitious play and choose an immediate best-reply, under standard preferences (as in model *DET*). Therefore, myopic player beliefs in round $t+1$ are solely determined by observed history in round t, beliefs in round $t+2$ are determined by observed history in rounds t and $t+1$, and so on.

The precise way how these paths are calculated is shown in tables G.18 and G.19. Table G.18 shows the path of choices when $\theta = 6$. In this case action R is the best-response only for players who hold a belief of at least 0.9, therefore in period $t+1$ action R would be chosen only by those who observed all other players choosing R in period t. Consequently, the inefficient convention is overcome only if player i and all myopic players choose R in period t. If the sophisticated player chooses R and 4 myopic players choose S, there will be one myopic player who will have observed all others choosing R, and who would subsequently choose R in period $t+1$. But this player would not observe anyone choosing R in round $t+1$ and hence all players would choose S from period $t+2$ onwards. In all other cases all myopic players would be choosing S from period $t+1$ onwards.

Table G.19 shows the path of choices when $\theta = 5$. With this threshold value action R is the myopic best-response if a player holds a belief of at least 0.72, thus in period $t+1$ action R would be chosen only by those who observed at least 4 out of 5 others choosing R. If in period t player i chooses R, and so do 4 or 5 myopic players, all players end up playing R in period $t+1$, as well as in all subsequent periods. The same happens if player i chooses S and all myopic players choose R. But if the number of players who choose R in period t is 4 or less, the group always converges to an inefficient equilibrium, because only the myopic players who chose S in period t observe 4 others choosing R. This results in 1 or 2 myopic players choosing R, and would lead all myopic players to choose R in period $t+2$, regardless of the value of γ. In all other cases, when the total number of players choosing R in period t is 3 or less, all myopic players choose S in period $t+1$ and all subsequent periods.

Note that in general we would need to make assumptions about the way that myopic players are learning, in particular about the value of γ, to calculate the path of choices of myopic players. In this case, however, the results will hold for all values of γ because the path of choices will be solely determined by the history in round t. For example, in the case of $\theta = 6$ and the history $\{R, R, R, R, R, S\}$ in round t and history $\{S, S, S, S, S, R\}$ in round $t + 1$, all myopic players would choose S in round $t + 2$, even if $\gamma = 1$, that is if rounds t and $t + 1$ were weighed equally, because the belief would never exceed 0.9.

Since the farsighted player i anticipates the choices of other players, he will choose the action for period $t+1$ and all later periods in a way that maximizes the sum of payoffs from the game. Information about the actions presented in tables G.18 and G.19 can be used to compute the earnings for all players: round payoffs are equal to 60 when S is chosen, 100 when R is chosen and the threshold has been exceeded and 5 when R is chosen and the threshold has not been exceeded. Choosing R in round t becomes more attractive to player i as he puts higher probability on the event that 5 others will choose R in the case $\theta = 6$ or on the event that 4 others will choose R in the case $\theta = 5$, as in these two cases i's decision will determine to which equilibrium the play will converge.

Table G.18: Each table represents a choice path for every possible action profile in period t if $\theta = 6$.

	Actions						
Period	Player i	Other players					$\#R \geqslant \theta$?
t	R	R	R	R	R	R	Yes
$\geqslant t+1$	R	R	R	R	R	R	Yes

t	R	R	R	R	R	S	No
$t+1$	S	S	S	S	S	R	No
$\geqslant t+2$	S	S	S	S	S	S	No

t	R	R	R	R	S	S	No
$\geqslant t+1$	S	S	S	S	S	S	No

t	R	R	R	S	S	S	No
$\geqslant t+1$	S	S	S	S	S	S	No

t	R	R	S	S	S	S	No
$\geqslant t+1$	S	S	S	S	S	S	No

t	R	S	S	S	S	S	No
$\geqslant t+1$	S	S	S	S	S	S	No

	Actions						
Period	Player i	Other players					$\#R \geqslant \theta$?
t	S	R	R	R	R	R	No
$\geqslant t+1$	S	S	S	S	S	S	No

t	S	R	R	R	R	S	No
$\geqslant t+1$	S	S	S	S	S	S	No

t	S	R	R	R	S	S	No
$\geqslant t+1$	S	S	S	S	S	S	No

t	S	R	R	S	S	S	No
$\geqslant t+1$	S	S	S	S	S	S	No

t	S	R	S	S	S	S	No
$\geqslant t+1$	S	S	S	S	S	S	No

t	S	S	S	S	S	S	No
$\geqslant t+1$	S	S	S	S	S	S	No

Table G.19: Each table represents a choice path for every possible action profile in period t if $\theta = 5$.

	Actions						
Period	Player i	Other players					$\#R \geqslant \theta$?
t	R	R	R	R	R	R	Yes
$\geqslant t+1$	R	R	R	R	R	R	Yes

t	R	R	R	R	R	S	Yes
$\geqslant t+1$	R	R	R	R	R	R	Yes

t	R	R	R	R	S	S	No
$t+1$	S	S	S	S	R	R	No
$\geqslant t+2$	S	S	S	S	S	S	No

t	R	R	R	S	S	S	No
$\geqslant t+1$	S	S	S	S	S	S	No

t	R	R	S	S	S	S	No
$\geqslant t+1$	S	S	S	S	S	S	No

t	R	S	S	S	S	S	No
$\geqslant t+1$	S	S	S	S	S	S	No

	Actions						
Period	Player i	Other players					$\#R \geqslant \theta$?
t	S	R	R	R	R	R	Yes
$\geqslant t+1$	R	R	R	R	R	R	Yes

t	S	R	R	R	R	S	No
$t+1$	S	S	S	S	S	R	No
$\geqslant t+2$	S	S	S	S	S	S	No

t	S	R	R	R	S	S	No
$\geqslant t+1$	S	S	S	S	S	S	No

t	S	R	R	S	S	S	No
$\geqslant t+1$	S	S	S	S	S	S	No

t	S	R	S	S	S	S	No
$\geqslant t+1$	S	S	S	S	S	S	No

t	S	S	S	S	S	S	No
$\geqslant t+1$	S	S	S	S	S	S	No

Appendix H. Risk and Social Preferences

Appendix H.1. Preliminaries

Assume that player $i \in N$ chooses action $a_i \in \{1, 0\}$ where $a_i = 1$ if R is chosen and $a_i = 0$ if S is chosen. Define the action profile of all players except for i by a vector $a_{-i} = \times_{j \in N \setminus \{i\}} a_j$. The total number of players other than i who chose action R can be calculated using a dot product: $\#R = a_{-i} \cdot a_{-i}$. The vector of profits for all players can be expressed as a function of the action taken by player i and of the total number of other players who choose R:

$$\Pi(a_i, \#R) = \begin{cases} \{100, 100 \cdot \mathbb{1}^{\#R}, 60 \cdot \mathbb{1}^{n-\#R-1}\} & \text{if } \#R + a_i \geqslant \theta \text{ and } a_i = 1 \\ \{60\ \ , 100 \cdot \mathbb{1}^{\#R}, 60 \cdot \mathbb{1}^{n-\#R-1}\} & \text{if } \#R + a_i \geqslant \theta \text{ and } a_i = 0 \\ \{5\ \ \ \ , 5\ \ \ \ \cdot \mathbb{1}^{\#R}, 60 \cdot \mathbb{1}^{n-\#R-1}\} & \text{if } \#R + a_i < \theta \text{ and } a_i = 1 \\ \{60\ \ , 5\ \ \ \ \cdot \mathbb{1}^{\#R}, 60 \cdot \mathbb{1}^{n-\#R-1}\} & \text{if } \#R + a_i < \theta \text{ and } a_i = 0 \end{cases}$$

where $\mathbb{1}^x$ is a vector of ones of length x. The first entry in the profit vector specifies the profit to player i, the second entry specifies the profits to $\#R$ players who chose R and the third entry specifies the profits for the $n - \#R - 1$ players who chose S. Define the payoff to player i as $\Pi_i(a_i, \#R)$ and the payoff vector of other players as $\Pi_{-i}(a_i, \#R)$.

Define subjective beliefs about the number of other group members who will choose R by a probability mass function $\sigma(\#R) : \{0, 1, \ldots, n-1\} \to [0, 1]$.

The expected utility of action a_i, conditional on subjective beliefs σ, is:

$$EU(a_i, \sigma) = \sum_{b \in \{0, \ldots, n-1\}} \sigma(b) u(\Pi(a_i, b))$$

Action R maximizes expected utility if a function that measures the difference between expected utilities is positive:

$$d(\sigma, u) = EU(R, \sigma) - EU(S, \sigma) = \sum_{b \in \{0, \ldots, n-1\}} \sigma(b)(u(R, b) - u(S, b)) \geqslant 0 \tag{H.1}$$

Appendix H.2. Risk Neutral, Self-interested Players

Self interested players would take into account only their own payoffs, specified by the first entry of a payoff vector $\Pi(a_i, \#R)$:

$$\Pi_i(a_i, \#R) = \begin{cases} 100 & \text{if } a_i = 1 \text{ and } \#R \geqslant \theta - 1 \\ 5 & \text{if } a_i = 1 \text{ and } \#R < \theta - 1 \\ 60 & \text{if } a_i = 0 \end{cases} \tag{H.2}$$

For risk neutral players utility is equal to payoffs: $u(\Pi_i(a_i, \#R)) = \Pi_i(a_i, \#R)$.

Substitute (H.2) into (H.1) and define $Pr(\#R \geqslant \theta - 1) = \sum_{b=\theta-1}^{n} \sigma(b)$ to get:

$$
\begin{aligned}
d(\sigma, u) \geqslant 0 \quad &\Leftrightarrow \\
Pr(\#R \geqslant \theta - 1) * 40 + (1 - Pr(\#R \geqslant \theta - 1)) * (-55) \geqslant 0 \quad &\Leftrightarrow \\
Pr(\#R \geqslant \theta - 1) \geqslant \frac{11}{19} &
\end{aligned}
\tag{H.3}
$$

A risk neutral player will choose R if and only if he assigns a probability of at least $\frac{11}{19}$ to the event that the participation threshold will be exceeded.

Appendix H.3. Risk Attitudes

Now suppose that players are not necessarily risk neutral. Define $v(x)$ satisfying $v'(x) > 0$ as the utility of receiving x, and use (H.2) to calculate utility:

$$
u(a_i, \#R) = \begin{cases} v(100) & \text{if } a_i = 1 \text{ and } \#R \geqslant \theta - 1 \\ v(5) & \text{if } a_i = 1 \text{ and } \#R < \theta - 1 \\ v(60) & \text{if } a_i = 0 \end{cases} \tag{H.4}
$$

Note that if $\sigma(\cdot)$ is degenerate, there is no strategic uncertainty and thus the strict monotonicity of $v(\cdot)$ ensures that the shape of the utility function will not affect the optimal choice. If $\sigma(\cdot)$ is not degenerate and $Pr(\#R > \theta - 1) \in (0, 1)$, R will be the optimal action if:

$$
\begin{gathered}
d(\sigma, u) \geqslant 0 \Leftrightarrow \\
Pr(\#R \geqslant \theta - 1) * (v(100) - v(60)) + (1 - Pr(\#R \geqslant \theta - 1)) * (v(5) - v(60)) \geqslant 0 \Leftrightarrow
\end{gathered}
$$

$$
\frac{Pr(\#R \geqslant \theta - 1)}{1 - Pr(\#R \geqslant \theta - 1)} \frac{v(100) - v(60)}{v(60) - v(5)} \geqslant 0 \tag{H.5}
$$

Consider two utility functions, $v_1(x)$ and $v_2(x)$ such that v_1 exhibits higher risk aversion $\left(\frac{-v_1''(x)}{v_1'(x)} > \frac{-v_2''(x)}{v_2'(x)}\right)$. Then $\frac{v_1(100)-v_1(60)}{v_1(60)-v_1(5)} < \frac{v_2(100)-v_2(60)}{v_2(60)-v_2(5)}$ (Pratt, 1964) and therefore $d(\sigma, v_1) < d(\sigma, v_2)$.

We conclude that players with lower risk aversion coefficients are weakly more likely to choose R, holding other factors constant.

Appendix H.4. Social Preferences

Suppose that utility depends not only on own earnings π_i, but also on the earnings of all other group members. In particular, assume that the utility function is separable and can be written the following way:

$$
u_i(a_i, a_{-i}) = \Pi_i(a_i, \#R) + \alpha v(\Pi_{-i}(a_i, \#R))
$$

where $v(\cdot)$ captures the effect of the payoffs received by other group members. We assume that $\nabla_{\mathbb{1}}(v(\pi_1, \pi_2, \ldots, \pi_{n-1})) \geqslant 0$, that is $v(\Pi_{-i}(a_i, \#R))$ does not decrease if there is no other player whose payoff decreased. A selfish player would be characterised by $\alpha = 0$, and would take into account only his own payoff. Players with $\alpha > 0$ would be pro-social and those with $\alpha < 0$ would be anti-social. If $\alpha > 0$, altruism could be modeled by setting $v(\cdot) = \sum(\cdot)$, maximin preferences could be modeled by setting $v(\cdot) = \min(\cdot)$. R is optimal if the following holds:

$$
\begin{aligned}
& d(\sigma, u) = \\
\sum_{b \in \{0,\ldots,n-1\}} \sigma(b)(\Pi_i(R,b) + \alpha v[\Pi_{-i}(R,b)]) - \sum_{b \in \{0,\ldots,n-1\}} & \sigma(b)(\Pi_i(S,b) + \alpha v[\Pi_{-i}(S,b)]) = \\
\sum_{b \in \{0,\ldots,n-1\}} \sigma(b)(\Pi_i(R,b) - \Pi_i(S,b)) + \alpha \sum_{b \in \{0,\ldots,n-1\}} & \sigma(b)(v[\Pi_{-i}(R,b)] - v[\Pi_{-i}(S,b)]) \geqslant 0
\end{aligned}
$$

The first term is not affected by social preferences because it includes only the earnings of player i. Notice that $\Pi_{-i}(R,b) \geqslant \Pi_{-i}(S,b)$, $\forall b \in \{1, \ldots, n-1\}$, that is payoffs received by other group members are at least as high when i chooses R as when i chooses S. Also, $\Pi_{-i}(R, \theta - 1) \geqslant \Pi_{-i}(S, \theta - 1)$, because if a person is pivotal, choosing R would lead to the threshold being exceeded and would therefore provide higher payoffs for those group members who chose R, compared to the case if S was chosen. Thus if a positive weight is given to the event that $\theta - 1$ others will choose R, as we have assumed, the second term would be increasing in α. The value of $d(\sigma, u)$ would be higher for players who are more pro-social, defined by higher values of α, thus for such players R would be optimal under a broader set of parameter values.

Note that the positive effect of pro-social preferences on R choice holds only if players make deterministic choices. If choices are stochastic and depend on the payoff difference, increased pro-sociality may decrease the payoff difference between R and S, making R less likely to be chosen.

Appendix I. SVO Slider Measure

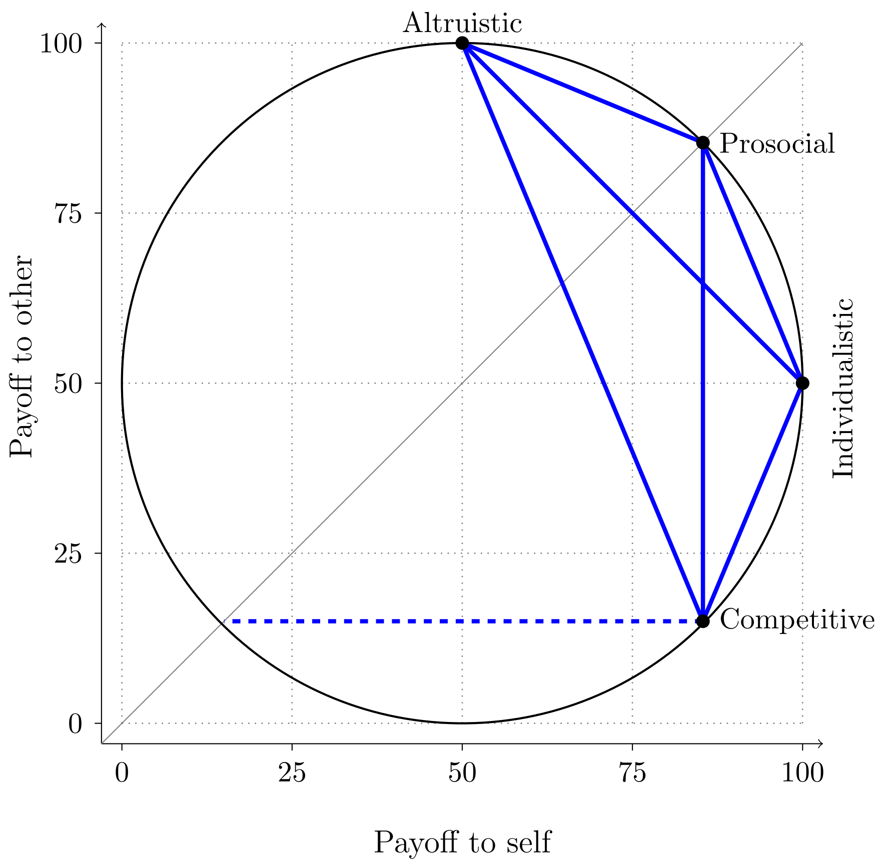

Figure I.8: Six allocation tasks of the Slider measure (solid lines) and the additional task (dashed line). Each line represents one allocation task, in which subjects choose one of the 9 allocations on the line.

Appendix J. Instructions

Instructions were identical for treatments HIGH and LOW, except for the disclosure costs, which are highlighted in red below.

INSTRUCTIONS

Welcome to the experiment. Please read the instructions carefully and take as much time as you need. The instructions are identical for all the participants with whom you will interact during this experiment.

If you have any questions please raise your hand. One of the experimenters will come to you and answer your questions. From now on communication with other participants is not allowed. If you do not conform to these rules you will be excluded from the experiment with no payment. Please also switch off your mobile phone at this moment.

In this experiment you can earn some money. How much you earn depends on your decisions and the decisions of the other participants. During the experiment we will refer to ECU (Experimental Currency Unit) instead of Euro. The total amount of ECU that you will have earned during the experiment will be converted into Euro at the end of the experiment and paid to you in cash confidentially. In this experiment the conversion rate that will be used to convert your ECU earnings into your Euro cash payment is: 250 ECU = 1 euro.

Overview

This experiment consists of multiple rounds. At the start of the experiment you will be randomly matched with 5 other participants. The group will remain constant throughout the experiment so you will interact with the same 5 other participants in all rounds. All other participants will face the same decision task as you, and you will not know the identity of each other. Every round consists of two tasks, a Decision Task and a Guess Task. You will do both tasks in every round, but you will be paid either for Decision Tasks or for the Guess Tasks. At the end of the experiment it will be randomly determined whether your payment is based on the Decision Tasks or on Guess Tasks. There is a chance of 4 out of 5 that your payment will be based on the Decision Tasks and a chance of 1 out of 5 that your payment will be based on the Guess Tasks.

Next we will describe what you have to do in each of the tasks.

Decision Task

- Your initial balance at the start of round 1 is 400 ECU.
- In each round you will have to choose either action A or action B.
- If you choose action B, your round income will be 60 ECU.
- If you choose action A, your round income depends on decisions of other participants in your group. If x or more others (not including yourself) choose A, your round income will be 100 ECU. If less than x others choose A, your round income will be 5 ECU. For example, if x=2, your round income will be 5 ECU if 0 or 1 other participant chooses A and your round income will be 100 ECU if 2,3,4 or 5 other participants choose A.

- The rounds will be grouped into a certain number of blocks. At the start of a block you will be informed about the number of rounds that the block contains and the value of *x* in that block. The value of *x* will stay the same in every round of the block, and you will be reminded about the value of *x* in every round.

- Everyone will make decisions at the same time, so you will not know the decisions of other participants before making your own decision.

Table 1 summarizes your choices and possible round income in the Decision Task.

Your choice	Your income if *x* or more other participants choose A	Your income if less than *x* other participants choose A
A	100	5
B	60	60

Table 1: Round income for the Decision Task (in ECU)

Making your action public or hidden

- After choosing an action in the Decision Task you will have to choose if you want to make this action public or if you want to make it hidden.

- If you make your action public, all other participants in your group will be informed about the action (A or B) that you chose in the Decision Task, but only after the Guess Task has been completed. If you make your action hidden, other participants will not know your action in the Decision Task.

- Making your action public will cost you 80 ECU [2 ECU] . Making your action hidden will cost you 1 ECU. This amount will be subtracted from your round income. If your round income is not sufficient to cover these costs, the difference will be subtracted from the initial balance or from your income in other rounds.

- Making your action public or hidden will affect only the information that others see, but it will not affect their income. Likewise, your income will not be affected by whether other participants in your group made their actions public or hidden.

Guess Task

- In the Guess Task you will be asked to guess what action, A or B, some other participant in your group chose in the Decision Task. We ask you to report your Guess on a scale from 0 to 100, where 0 means that you are sure the other participant chose B and 100 means that you are sure that the other participant chose A. The participant whose action you have to guess (call him "other participant") will be randomly chosen by the computer. All 5 other participants in your group have equal chances to be chosen.

- You will not earn any ECU from the Guess Tasks, but you will earn some lottery tickets. If your payment is based on the Guess Tasks, at the end of the experiment you will play a lottery in which you will receive either 4000 ECU or 1000 ECU. The probability to receive 4000 ECU will depend on the number of tickets that you received in one of the Guess Tasks. Which Guess Task is used will be determined randomly by the computer and every Guess Task has an equal chance of being selected.

- The exact way how the number of lottery tickets depends on your Guess and on the action of the other participant in the Decision Task is shown in Table 2. If you guess 0, you will receive 4000 ECU for sure if the other participant chose B and will receive 1000 ECU for sure if the other participant chose A. If you guess 100, you will receive 4000 ECU for sure if the other participant chose A and will receive 1000 ECU for sure if the other participant chose B. All other Guesses give a positive probability to receive either sum. Notice that as you increase your Guess, you receive more tickets if the other participant chose A but receive less tickets if the other participant chose B.

Your Guess	Number of lottery tickets that you will receive if other participant chose A in the Decision Task	Number of lottery tickets that you will receive if other participant chose B in the Decision Task
0	0.00	100.00
5	9.75	99.75
10	19.00	99.00
15	27.75	97.75
20	36.00	96.00
25	43.75	93.75
30	51.00	91.00
35	57.75	87.75
40	64.00	84.00
45	69.75	79.75
50	75.00	75.00
55	79.75	69.75
60	84.00	64.00
65	87.75	57.75
70	91.00	51.00
75	93.75	43.75
80	96.00	36.00
85	97.75	27.75
90	99.00	19.00
95	99.75	9.75
100	100.00	0.00

Table 2 Number of lottery tickets received in the Guess Task

- Notice that Table 2 lists the payoffs only for the multiples of five, but you are free to choose any number between 0 and 100. To make it easier to choose, we will first ask you to choose a multiple of five, and then allow you pick a more precise Guess by showing earnings for 11 numbers closest to the one you picked. For example, if in the first table you choose number 50, you will be given numbers 45-55 to choose from in the second table.

- Your choice in the Guess Task only affects your payoff. It does not affect the payoffs of other participants.

What you will see at the end of a round

After both the Decision Task and the Guess Task have been completed, you will see the following information on the screen:

- Your choice in the Decision Task
- The number of other participants in your group who chose action A and made it public
- The number of other participants in your group who chose action B and made it public
- The number of other participants in your group who made their action hidden
- Your payoff from the Decision Task

You will not know the choices of participants who made their choices hidden.

Information about your choices and choices of other participants in previous rounds will always be available on the right side of your computer screen.

You will not receive any feedback about the Guess Task. The choice of a randomly selected participant and the number of lottery tickets that you earned in every round will be displayed at the end of the experiment.

How your cash earnings will be determined

At the end of the experiment it will be randomly determined whether you are paid for the Decision Tasks or for the Guess Tasks. There is a chance of 4 out of 5 that you will be paid for Decision Tasks and there is a chance of 1 out of 5 that you will be paid for Guess Tasks. Note that if you are paid for

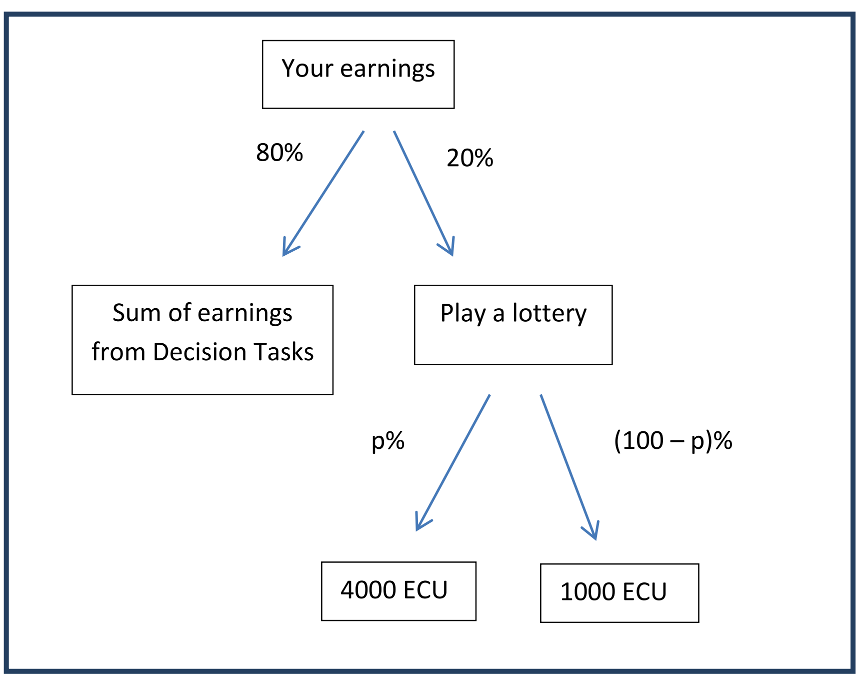

Figure 1. How your earnings from Part 1 will be determined. p is the number of lottery tickets that you got from one randomly chosen Guess Task.

Decision Tasks, your choices in the Guess Task will not influence your final earnings. If you are paid for Guess tasks, you will not receive the earnings that you made in the Decision tasks.

If your payment is based on Decisions Tasks, your final earnings will be determined by adding your round income from all Decision Tasks to the initial balance of 400 ECU.

If your payment is based on Guess Tasks, your total earnings will be either 4000 ECU or 1000 ECU. Which sum you will receive will be determined by spinning a virtual Lottery Wheel. The Lottery Wheel will have two sectors, Red and Blue, and will randomly stop in one of them. If the Lottery Wheel stops in the Red sector, you will receive 4000 ECU; if it stops in the Blue sector, you will receive 1000 ECU. The size of the Red sectors will be determined by the number of lottery tickets that you received in one Guess Task. For example, if you earned *p* lottery tickets, *p*% of the lottery wheel will be colored Red and the remaining (100 - *p*)% will be colored blue. Therefore the probability that you will receive 4000 ECU will be *p*% and the probability that you will receive 1000 ECU will be (100-p)% . Which Guess Task is used to determine the size of the Red sector will be determined randomly by the computer and all Guess Tasks have equal probability to be selected. The spin of the virtual Lottery Wheel will be performed on your computer screen at the end of the experiment.

How it will be determined if you are paid for the Decision Tasks or for the Guess Tasks

At the end of the experiment we will ask you to choose a number between 1 and 5. Afterwards we will ask one participant in this room to randomly draw one card out of 5 cards, numbered from 1 to 5. The draw will be public, so everyone will be able to observe what number has been drawn. If the drawn number is the same as the number that you entered, your payment will be based on the Guess Tasks. If the drawn number is different from the one that you have entered, your earnings will be based on the Decision Tasks.

What you will see at the end of the experiment

At the end of the experiment you will be informed about the choices of all participants in all rounds and your earnings from all Decision Tasks and all Choice Tasks. You will see if you are paid for the Decision Tasks or for the Guess Tasks. If you are paid for the Guess Tasks, you will see the result of a virtual Lottery Wheel spin.

After you have made decisions in all rounds, but before seeing the final feedback screen you will have to do some additional tasks in which you will have a chance to earn more ECU. Instructions for these additional tasks will be shown on your computer screen. After these tasks you will see your total earnings from the experiment. We will also ask you to complete a short questionnaire. After completing the questionnaire, please stay seated until we ask you to come to receive your earnings from the experiment. Your earnings will be paid in cash and in private.

If you have any further questions, please raise your hand now.

In order to ensure that everybody has understood the instructions, we will ask you to answer a few questions. Please click a button on your computer screen to start answering these questions.

Appendix K. Screenshots

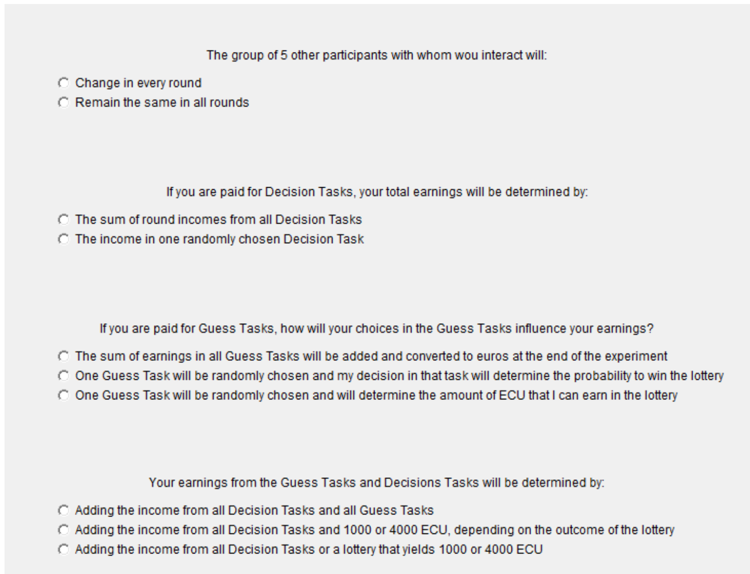

The group of 5 other participants with whom wou interact will:

- Change in every round
- Remain the same in all rounds

If you are paid for Decision Tasks, your total earnings will be determined by:

- The sum of round incomes from all Decision Tasks
- The income in one randomly chosen Decision Task

If you are paid for Guess Tasks, how will your choices in the Guess Tasks influence your earnings?

- The sum of earnings in all Guess Tasks will be added and converted to euros at the end of the experiment
- One Guess Task will be randomly chosen and my decision in that task will determine the probability to win the lottery
- One Guess Task will be randomly chosen and will determine the amount of ECU that I can earn in the lottery

Your earnings from the Guess Tasks and Decisions Tasks will be determined by:

- Adding the income from all Decision Tasks and all Guess Tasks
- Adding the income from all Decision Tasks and 1000 or 4000 ECU, depending on the outcome of the lottery
- Adding the income from all Decision Tasks or a lottery that yields 1000 or 4000 ECU

Figure K.9: First part of the questionnaire that had to be answered before subjects could start the experiment.

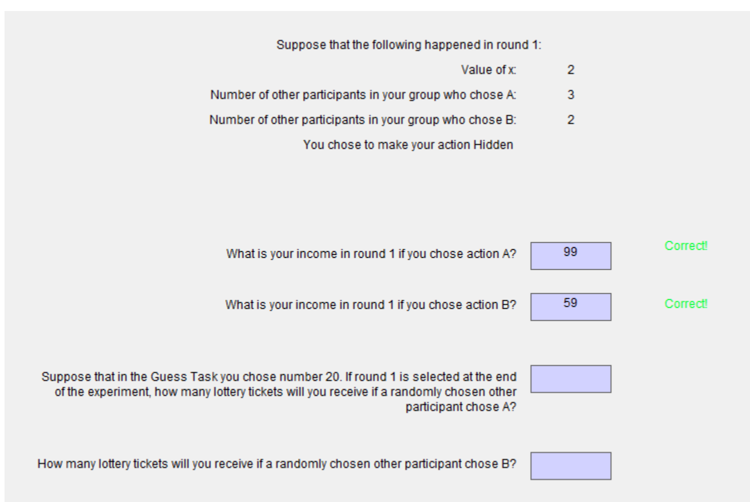

Suppose that the following happened in round 1:

Value of x: 2

Number of other participants in your group who chose A: 3

Number of other participants in your group who chose B: 2

You chose to make your action Hidden

What is your income in round 1 if you chose action A? 99 Correct!

What is your income in round 1 if you chose action B? 59 Correct!

Suppose that in the Guess Task you chose number 20. If round 1 is selected at the end of the experiment, how many lottery tickets will you receive if a randomly chosen other participant chose A?

How many lottery tickets will you receive if a randomly chosen other participant chose B?

Figure K.10: Second part of the questionnaire that had to be answered before subjects could start the experiment. Values for the first question were generated randomly.

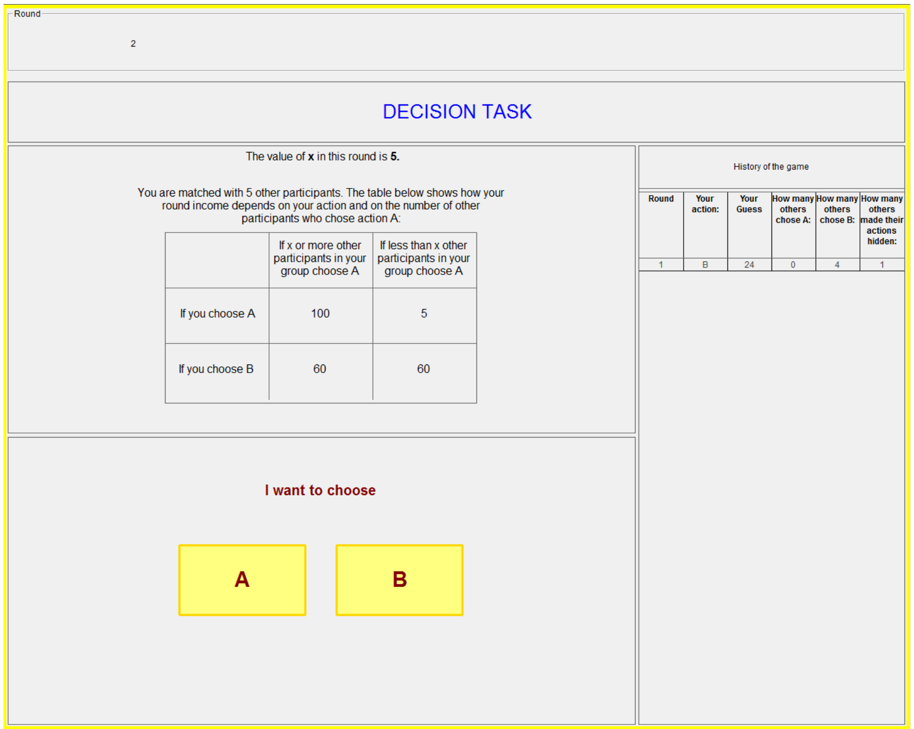

Figure K.11: Action choice stage

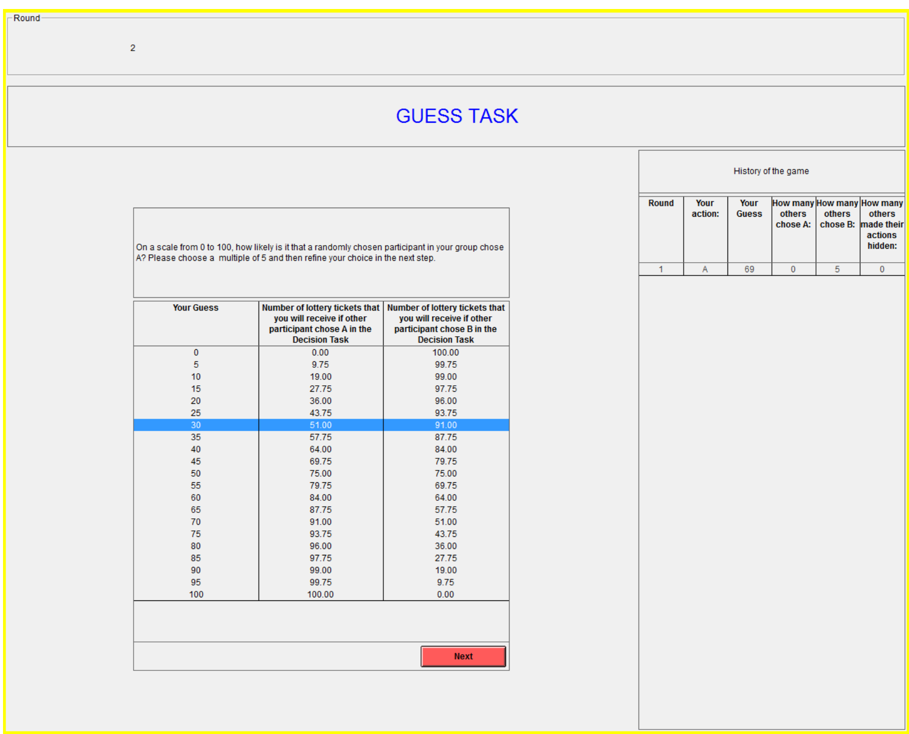

Figure K.12: First part of the belief elicitation task

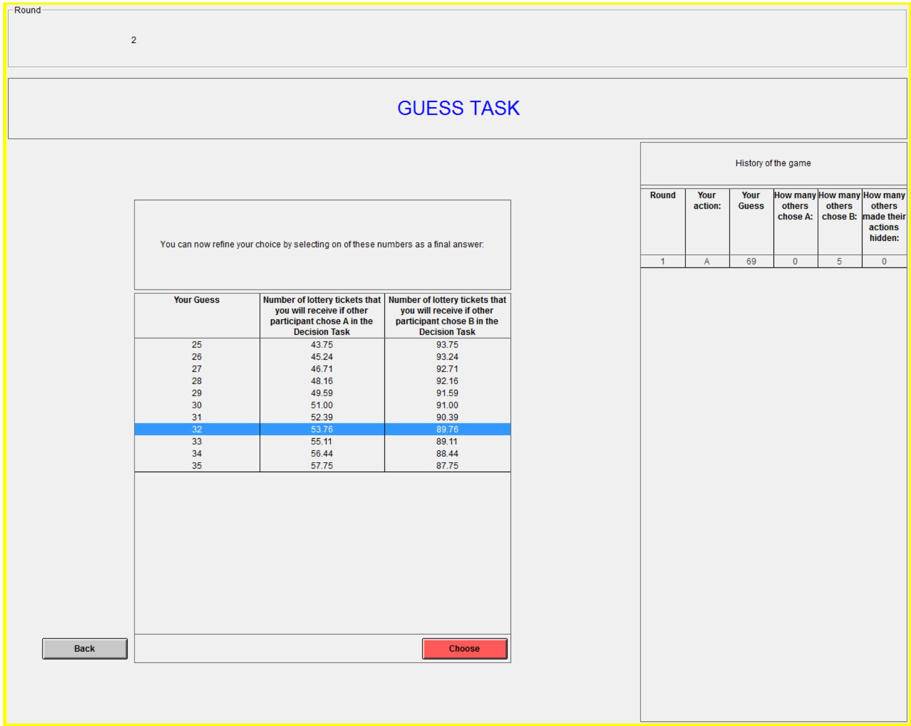

Figure K.13: Second part of the belief elicitation task

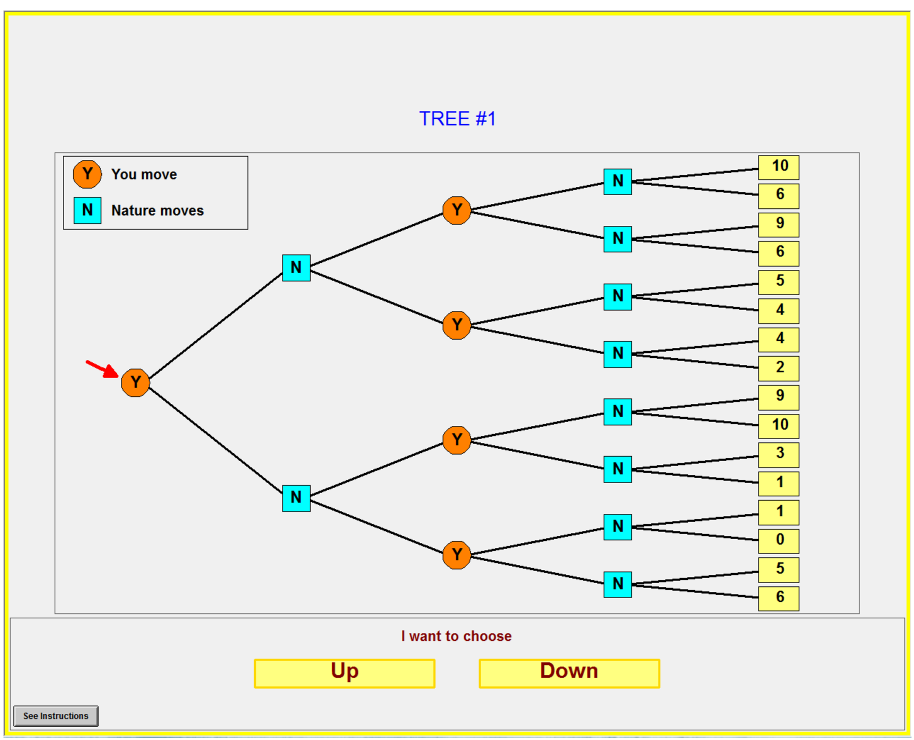

Figure K.14: Decision tree 1 from the task measuring farsightedness.

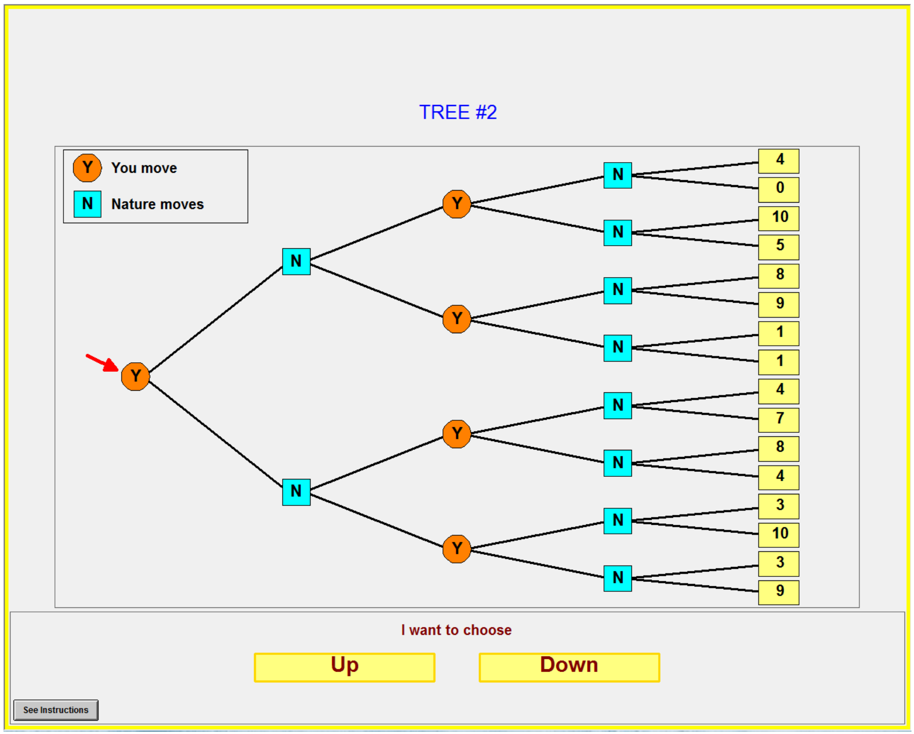

Figure K.15: Decision tree 2 from the task measuring farsightedness.

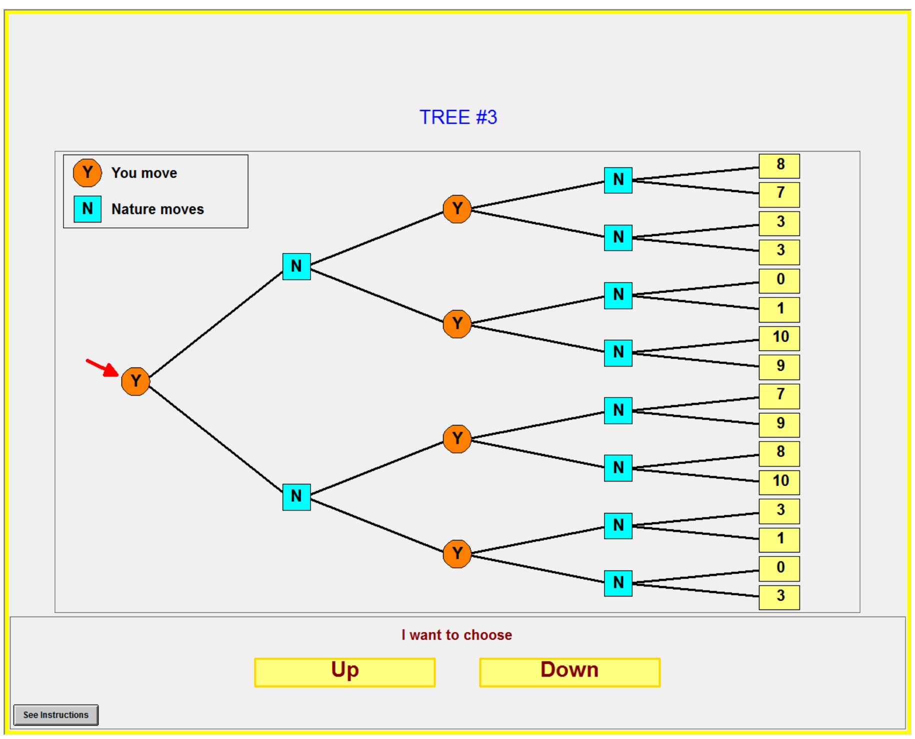

Figure K.16: Decision tree 3 from the task measuring farsightedness.

References

Acemoglu, D. (2006). A simple model of inefficient institutions. *The Scandinavian Journal of Economics*, 108(4):515–546.

Akerlof, G. A. (1980). A theory of social custom, of which unemployment may be one consequence. *The Quarterly Journal of Economics*, 94(4):749–775.

Angeletos, G.-M., Hellwig, C., and Pavan, A. (2007). Dynamic global games of regime change: Learning, multiplicity, and the timing of attacks. *Econometrica*, 75(3):711–756.

Arieli, A., Ben-Ami, Y., and Rubinstein, A. (2011). Tracking decision makers under uncertainty. *American Economic Journal: Microeconomics*, 3(4):68–76.

Armantier, O. and Treich, N. (2013). Eliciting beliefs: Proper scoring rules, incentives, stakes and hedging. *European Economic Review*, 62:17–40.

Arthur, W. B. (1989). Competing technologies, increasing returns, and lock-in by historical events. *The Economic Journal*, 99(394):116–131.

Barbera, S. and Jackson, M. O. (2017). A model of protests, revolution, and information. Working paper, available at SSRN: https://ssrn.com/abstract=2732864.

Battalio, R. C., Samuelson, L., and Van Huyck, J. B. (1997). Risk dominance, payoff dominance and probabilistic choice learning. Working Paper, Department of Economics, Texas A&M University.

Best, M. L. and Wade, K. W. (2009). The internet and democracy global catalyst or democratic dud? *Bulletin of science, technology & society*, 29(4):255–271.

Blanco, M., Engelmann, D., Koch, A. K., and Normann, H.-T. (2010). Belief elicitation in experiments: is there a hedging problem? *Experimental Economics*, 13(4):412–438.

Blume, A. and Ortmann, A. (2007). The effects of costless pre-play communication: Experimental evidence from games with pareto-ranked equilibria. *Journal of Economic Theory*, 132(1):274–290.

Bone, J., Hey, J. D., and Suckling, J. (2009). Do people plan? *Experimental Economics*, 12(1):12–25.

Brandts, J. and Cooper, D. J. (2006a). A change would do you good. an experimental study on how to overcome coordination failure in organizations. *The American Economic Review*, 96(3):669–693.

Brandts, J. and Cooper, D. J. (2006b). Observability and overcoming coordination failure in organizations: An experimental study. *Experimental Economics*, 9(4):407–423.

Brandts, J., Cooper, D. J., and Weber, R. A. (2014). Legitimacy, communication, and leadership in the turnaround game. *Management Science*, 61(11):2627–2645.

Camerer, C. and Ho, T.-H. (1999). Experience-weighted attraction learning in normal form games. *Econometrica*, 67(4):827–874.

Camerer, C. F., Ho, T.-H., and Chong, J.-K. (2002). Sophisticated experience-weighted attraction learning and strategic teaching in repeated games. *Journal of Economic Theory*, 104(1):137–188.

Cheung, Y.-W. and Friedman, D. (1997). Individual learning in normal form games: Some laboratory results. *Games and Economic Behavior*, 19(1):46–76.

Cooper, R., DeJong, D. V., Forsythe, R., and Ross, T. W. (1992). Communication in coordination games. *The Quarterly Journal of Economics*, 107(2):739–771.

Costa-Gomes, M., Crawford, V. P., and Broseta, B. (2001). Cognition and behavior in normal-form games: An experimental study. *Econometrica*, 69(5):1193–1235.

Cowan, R. (1990). Nuclear power reactors: a study in technological lock-in. *The journal of economic history*, 50(03):541–567.

Crosetto, P., Weisel, O., and Winter, F. (2012). A flexible z-Tree implementation of the Social Value Orientation Slider Measure (Murphy et al. 2011): Manual. *Jena Economic Research Papers*, 2012 – 062.

Dainotti, A., Squarcella, C., Aben, E., Claffy, K. C., Chiesa, M., Russo, M., and Pescapé, A. (2014). Analysis of country-wide internet outages caused by censorship. *IEEE/ACM Transactions on Networking (TON)*, 22(6):1964–1977.

David, P. A. (1985). Clio and the economics of qwerty. *The American economic review*, 75(2):332–337.

De Mesquita, E. B. (2010). Regime change and revolutionary entrepreneurs. *American Political Science Review*, 104(03):446–466.

Devetag, G. (2003). Coordination and information in critical mass games: an experimental study. *Experimental Economics*, 6(1):53–73.

Devetag, G. and Ortmann, A. (2007). When and why? a critical survey on coordination failure in the laboratory. *Experimental Economics*, 10(3):331–344.

Diamond, D. W. and Dybvig, P. H. (1983). Bank runs, deposit insurance, and liquidity. *Journal of political economy*, 91(3):401–419.

Dong, L., Montero, M., and Possajennikov, A. (2017). Communication, leadership and coordination failure. Working paper, available at SSRN: https://ssrn.com/abstract=2655861.

Duffy, J. and Ochs, J. (2012). Equilibrium selection in static and dynamic entry games. *Games and Economic Behavior*, 76(1):97–116.

Edmond, C. (2013). Information manipulation, coordination, and regime change. *The Review of Economic Studies*, 80(4):1422–1458.

Enikolopov, R., Makarin, A., and Petrova, M. (2016). Social media and protest participation: Evidence from Russia. CEPR Discussion Paper No. DP11278.

Fehr, E. and Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *The quarterly journal of economics*, 114(3):817–868.

Fischbacher, U. (2007). z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2):171–178.

Friedman, D. (1996). Equilibrium in evolutionary games: Some experimental results. *The Economic Journal*, 106(434):1–25.

Garratt, R. and Keister, T. (2009). Bank runs as coordination failures: An experimental study. *Journal of Economic Behavior & Organization*, 71(2):300–317.

Gneezy, U., Rustichini, A., and Vostroknutov, A. (2010). Experience and insight in the race game. *Journal of Economic Behavior & Organization*, 75(2):144–155.

Greiner, B. (2015). Subject pool recruitment procedures: organizing experiments with ORSEE. *Journal of the Economic Science Association*, 1(1):114–125.

Guttman, L. (1950). The basis for scalogram analysis. In Stouffer, S. A., Suchman, E. A., Lazarsfeld, P. F., Star, S. A., and Clausen, J. A., editors, *Measurement and prediction. [Studies in social psychology in World War II. Vol.4.]*. Princeton University Press, Princeton, NJ, US.

Hamman, J., Rick, S., and Weber, R. A. (2007). Solving coordination failure with all-or-none group-level incentives. *Experimental Economics*, 10(3):285–303.

Hassanpour, N. (2014). Media disruption and revolutionary unrest: Evidence from mubarak's quasi-experiment. *Political Communication*, 31(1):1–24.

Heinemann, F., Nagel, R., and Ockenfels, P. (2004). The theory of global games on test: experimental analysis of coordination games with public and private information. *Econometrica*, 72(5):1583–1599.

Heinemann, F., Nagel, R., and Ockenfels, P. (2009). Measuring strategic uncertainty in coordination games. *The Review of Economic Studies*, 76(1):181–221.

Holt, C. A. and Laury, S. K. (2002). Risk aversion and incentive effects. *American economic review*, 92(5):1644–1655.

Hossain, T. and Okui, R. (2013). The binarized scoring rule. *The Review of Economic Studies*, 80(3):984–1001.

Hyndman, K., Ozbay, E. Y., Schotter, A., and Ehrblatt, W. Z. (2012). Convergence: an experimental study of teaching and learning in repeated games. *Journal of the European Economic Association*, 10(3):573–604.

Hyndman, K., Terracol, A., and Vaksmann, J. (2009). Learning and sophistication in coordination games. *Experimental Economics*, 12(4):450–472.

Johnson, E. J., Camerer, C., Sen, S., and Rymon, T. (2002). Detecting failures of backward induction: Monitoring information search in sequential bargaining. *Journal of Economic Theory*, 104(1):16–47.

Kalai, E. and Lehrer, E. (1993). Rational learning leads to Nash equilibrium. *Econometrica*, 61(5):1019–1045.

Kalathil, S. and Boas, T. C. (2003). *Open networks, closed regimes: The impact of the Internet on authoritarian rule*. Carnegie Endowment.

Kedzie, C. (1997). Communication and democracy: Coincident revolutions and the emergent dictators. *RAND Dissertation*.

Keser, C., Suleymanova, I., and Wey, C. (2012). Technology adoption in markets with network effects: Theory and experimental evidence. *Information Economics and Policy*, 24(3-4):262–276.

Kinateder, M., Kiss, H. J., and Pintér, Á. (2015). Would depositors like to show others that they do not withdraw? theory and experiment. Working paper, available at SSRN: https://ssrn.com/abstract=2680152.

King, G., Pan, J., and Roberts, M. E. (2013). How censorship in china allows government criticism but silences collective expression. *American Political Science Review*, 107(02):326–343.

Kiss, H. J., Rodriguez-Lara, I., and Rosa-García, A. (2014). Do social networks prevent or promote bank runs? *Journal of Economic Behavior & Organization*, 101:87–99.

Kiss, H. J., Rodriguez-Lara, I., and Rosa-Garcia, A. (2016). Overthrowing the dictator: A game-theoretic approach to revolutions and media. Working paper, available at SSRN: https://ssrn.com/abstract=2298662.

Kogan, S., Kwasnica, A. M., and Weber, R. A. (2011). Coordination in the presence of asset markets. *American Economic Review*, 101(2):927–947.

Kuran, T. (1989). Sparks and prairie fires: A theory of unanticipated political revolution. *Public choice*, 61(1):41–74.

Lee, J. (2008). The effect of the background risk in a simple chance improving decision model. *Journal of Risk and Uncertainty*, 36(1):19–41.

Lewis, D. (1969). *Convention: A philosophical study*. John Wiley & Sons.

Little, A. T. (2016). Communication technology and protest. *The Journal of Politics*, 78(1):152–166.

Lohmann, S. (1993). A signaling model of informative and manipulative political action. *The American Political Science Review*, 87(2):319.

Lohmann, S. (1994). The dynamics of informational cascades: The monday demonstrations in Leipzig, East Germany, 1989-91. *World Politics*, 47(01):42–101.

Lohmann, S. (2000). Collective action cascades: An informational rationale for the power in numbers. *Journal of Economic Surveys*, 14(5):655–684.

Lynch, M. (2011). After Egypt: The limits and promise of online challenges to the authoritarian Arab state. *Perspectives on Politics*, 9(02):301–310.

Mak, V. and Zwick, R. (2010). Investment decisions and coordination problems in a market with network externalities: An experimental study. *Journal of Economic Behavior & Organization*, 76(3):759–773.

Manacorda, M. and Tesei, A. (2016). Liberation technology: mobile phones and political mobilization in Africa. CEPR Discussion Paper No. DP11278.

Masiliunas, A. (2016). Inefficient lock-in with sophisticated and myopic players. AMSE working paper WP 2016 - Nr 15.

McKelvey, R. D. and Page, T. (1990). Public and private information: An experimental study of information pooling. *Econometrica*, 58(6):1321–1339.

McKelvey, R. D. and Palfrey, T. R. (1992). An experimental study of the centipede game. *Econometrica*, 60(4):803–836.

Meier, P. (2011). *Do liberation technologies change the balance of power between repressive states and civil society?* PhD thesis, The Fletcher School of Law and Diplomacy.

Milgrom, P. and Roberts, J. (1991). Adaptive and sophisticated learning in normal form games. *Games and Economic Behavior*, 3(1):82–100.

Morozov, E. (2012). *The Net Delusion: The Dark Side of Internet Freedom*. PublicAffairs.

Mourtada, R. and Salem, F. (2011). Civil movements: The impact of Facebook and Twitter. *Arab Social Media Report*, 1(2):1–30.

Murphy, R. O., Ackermann, K. A., and Handgraaf, M. J. (2011). Measuring Social Value Orientation. *Judgment and Decision Making*, 6(8):771–781.

North, D. C. (1990). *Institutions, institutional change and economic performance*. Cambridge university press.

Nyarko, Y. and Schotter, A. (2002). An experimental study of belief learning using elicited beliefs. *Econometrica*, 70(3):971–1005.

Offerman, T., Sonnemans, J., and Schram, A. (1996). Value orientations, expectations and voluntary contributions in public goods. *Economic Journal*, 106(437):817–845.

Piovesan, M. and Wengström, E. (2009). Fast or fair? A study of response times. *Economics Letters*, 105(2):193–196.

Polonio, L., Di Guida, S., and Coricelli, G. (2015). Strategic sophistication and attention in games: an eye-tracking study. *Games and Economic Behavior*, 94:80–96.

Rapoport, A. (1985). Provision of public goods and the mcs experimental paradigm. *American Political Science Review*, 79(01):148–155.

Rapoport, A. and Eshed-Levy, D. (1989). Provision of step-level public goods: Effects of greed and fear of being gypped. *Organizational Behavior and Human Decision Processes*, 44(3):325–344.

Rubinstein, A. (2007). Instinctive and cognitive reasoning: A study of response times. *The Economic Journal*, 117(523):1243–1259.

Rubinstein, A. (2016). A typology of players: Between instinctive and contemplative. *The Quarterly Journal of Economics*, 131(2):859–890.

Ruffle, B. J., Weiss, A., and Etziony, A. (2015). The role of critical mass in establishing a successful network market: An experimental investigation. *Journal of Behavioral and Experimental Economics*, 58:101–110.

Sahin, S. G., Eckel, C., and Komai, M. (2015). An experimental study of leadership institutions in collective action games. *Journal of the Economic Science Association*, 1(1):100–113.

Schlag, K. H. and van der Weele, J. J. (2013). Eliciting probabilities, means, medians, variances and covariances without assuming risk neutrality. *Theoretical Economics Letters*, 3(1):38–42.

Shapiro, C. and Varian, H. R. (1999). *Information rules: a strategic guide to the network economy*. Harvard Business School Press, Boston, MA.

Shirky, C. (2011). The political power of social media: Technology, the public sphere, and political change. *Foreign affairs*, 90(01):28–41.

Shurchkov, O. (2013). Coordination and learning in dynamic global games: experimental evidence. *Experimental Economics*, 16(3):313–334.

Sonnemans, J., Schram, A., and Offerman, T. (1998). Public good provision and public bad prevention: The effect of framing. *Journal of Economic Behavior & Organization*, 34(1):143–161.

Spiliopoulos, L. and Ortmann, A. (2016). The BCD of response time analysis in experimental economics. Working paper, available at SSRN: https://ssrn.com/abstract=2401325.

Stanca, L. (2009). Measuring indirect reciprocity: Whose back do we scratch? *Journal of Economic Psychology*, 30(2):190–202.

Terracol, A. and Vaksmann, J. (2009). Dumbing down rational players: Learning and teaching in an experimental game. *Journal of Economic Behavior & Organization*, 70(1):54–71.

Tufekci, Z. and Wilson, C. (2012). Social media and the decision to participate in political protest: Observations from Tahrir square. *Journal of Communication*, 62(2):363–379.

Van Huyck, J. B., Battalio, R. C., and Beil, R. O. (1990). Tacit coordination games, strategic uncertainty, and coordination failure. *The American Economic Review*, 80(1):234–248.

Van Huyck, J. B., Battalio, R. C., and Beil, R. O. (1991). Strategic uncertainty, equilibrium selection, and coordination failure in average opinion games. *The Quarterly Journal of Economics*, 106(3):885–910.

Van Huyck, J. B., Battalio, R. C., and Rankin, F. W. (2007). Evidence on learning in coordination games. *Experimental Economics*, 10(3):205–220.

Young, H. P. (1996). The economics of convention. *The Journal of Economic Perspectives*, 10(2):105–122.