



Learning, Words and Actions: Experimental Evidence on Coordination-Improving Information

Nicolas Jacquemet, Adam Zylbersztein

► To cite this version:

Nicolas Jacquemet, Adam Zylbersztein. Learning, Words and Actions: Experimental Evidence on Coordination-Improving Information. B.E. Journal of Theoretical Economics, 2013, 13 (1), pp.215-247. 10.1515/bejte-2012-0018 . halshs-00845123

HAL Id: halshs-00845123

<https://shs.hal.science/halshs-00845123>

Submitted on 16 Jul 2013

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Learning, words and actions: experimental evidence on coordination-improving information*

Nicolas Jacquemet[†]

Adam Zylbersztein[†]

July 2013

Abstract

We experimentally study an asymmetric coordination game with two Nash equilibria: one is Pareto-efficient, the other is Pareto-inefficient and involves a weakly dominated strategy. We assess whether information about the interaction partner helps eliminate the imperfect equilibrium. Our treatments involve three information-enhancing mechanisms: repetition and two kinds of individual signals: messages from partner or observation of his past choices. Repetition-based learning increases the frequencies of the most efficient outcome and the most costly strategic mismatch. Moreover, it is superseded by individual signals. Like previous empirical studies, we find that signals provide a screening of partners' intentions that reduces the frequency of coordination failures. Unlike these studies, we find that the transmission of information between partners, either via messages or observation, does not suffice to significantly increase the overall efficiency of outcomes. This happens mostly because information does not restrain the choice of the dominated action by senders.

Keywords: Coordination Game, Communication, Cheap-talk, Observation.

JEL Classification: C72, D83.

*This paper is a revised version of *CES Working Paper* n°2010-64. We thank Ibrahim Ahamada, Juergen Bracht, Timothy Cason, Boğaçhan Çelen, Gary Charness, Nick Feltovich, Pierre Fleckinger, Guillaume Fréchette, Nobuyuki Hanaki, Frédéric Koessler, Michal Krawczyk, Stéphane Luchini, Andreas Ortmann, Drazen Prelec, Stéphane Robin, Jean-Marc Tallon, and Erik Wengström for inspiring discussions; participants in numerous conferences, workshops and seminars for insightful comments; Maxim Frolov for his technical assistance in running the experiments, as well as financial support from University Paris 1 and the Paris School of Economics. Nicolas Jacquemet acknowledges the support of the Institut Universitaire de France. Adam Zylbersztein is grateful to the State of Sao Paulo Research Foundation (FAPESP), the *Collège des Ecoles Doctorales de l'Université Paris 1 Panthéon-Sorbonne*, the Alliance Program and the Columbia University Economics Department for their support.

[†]Paris School of Economics and University Paris I Panthéon-Sorbonne. Centre d'Economie de la Sorbonne, 106 Bd. de l'Hopital, 75013 Paris, France. Nicolas.Jacquemet@univ-paris1.fr; adam.zylbersztein@malix.univ-paris1.fr.

Table 1: The experimental game

Player A	Player B	
	l	r
L	(9.75;3)	(9.75;3)
R	(3;4.75)	(10;5)

1 Introduction

In coordination games, people may fail to synchronize their actions because of strategic uncertainty – the risk associated with not knowing how other players will play the game. In such contexts, which give rise to multiple equilibria, theoretical refinements are characterized by assumptions on players’ beliefs about other players’ behavior. For instance, the idea that players rely on the own-payoff-maximizing behavior of others is enough in theory to rule out equilibria supported by non-credible strategies that may undermine the value of interaction. In his 1981 paper, Rosenthal conjectured that an imperfect equilibrium may nowadays happen depending on: *(i)* the stakes of the game and *(ii)* what is known about the interaction partner’s behavior. Existing experimental evidence has explored the first part of the conjecture, and confirmed the puzzle: the lower the surplus from relying on other’s ability to maximize own payoff, the more likely the failure to coordinate on efficient outcomes. The question of what kind of institutions can help overcome coordination failures for a *given* payoff structure, *i.e.* imposed by the incentives of the economic situation of interest, remains largely open.

This paper fills this gap by exploring the second dimension of the conjecture – the importance of the availability of information about the interaction partner. To better illustrate the challenge, Table 1 presents the normal-form version of a game originally proposed by Selten (1975) and extensively discussed by Rosenthal (1981). In the original sequential version of the game, player A moves first and chooses between L and R . If R is chosen, player B can maximize both players’ payoffs by choosing r , or undermine them both by choosing l . The only subgame perfect equilibrium is (R, r) which leads to the Pareto efficient payoff (10,5). But this outcome cannot be reached unless the likelihood that player B chooses l is perceived to be high enough by player A. The normal-form game thus has two pure-strategy Nash equilibria : (R, r) and (L, l) , since decision r is only weakly dominant for player B. The decision to play L involves less strategic uncertainty for player A – for the payoff structure presented in Table 1, for instance, the secure option L dominates the expected payoff of reliance (through choosing R) for probabilities of action r as low as 0.036.¹

¹This payoff structure coincides with treatment 1 in the seminal contribution by Beard and Beil (1994), who were the first to study subjects’ behavior under varying stakes in a sequential move game. They found that 66%

Our experimental treatments change what player A knows about the interaction partner, while holding constant the payoff structure of the game. The first kind of information we introduce is a simple repetition of the game. Pairs are rematched according to a perfect stranger procedure, so that player As are able to acquire better knowledge about the general population of player Bs. Two further treatments implement specific information about the current interaction partner. One treatment allows player Bs to provide player As with “soft” information about their intended behavior through cheap-talk communication: in this communication treatment, player Bs send closed-form messages to player As prior to the decision-making stage of the game. In the observation treatment, we turn to “harder” information in the form of actual decisions taken by player Bs: before making their own decisions, player As are informed about all the decisions their current interaction partners made in the past. To our best knowledge, we are the first to study this kind of mechanism in an *asymmetric* coordination problem, in which agents’ interests are aligned, but the stakes and the strategic risks they face differ. An advantage of our design as compared with symmetric games is that players’ functions in the information-transmission process (that is, being either the sender or the receiver) are a natural consequence of their strategic position in the game, rather than an outcome of a random draw.

Our results are twofold. First, we empirically demonstrate that coordination failures in this game have a different structure from that suggested by previous experimental studies. All previous experiments elicit player Bs’ actions conditional on player As being reliant. Given that player As seldom rely on their partners, and that the proportion of own-payoff-maximizers among those relied upon is generally high, these studies conclude that player As often fail to rely on player Bs’ own-payoff-maximization. This does indeed happen in the cases where player A interacts with a perfect own-payoff-maximizer and nonetheless chooses L (what we call a *type I error*, following the idea that player A’s decision provides a test of player B’s preferences). Yet, conditioning the observability of player Bs’ choices on player As’ prior reliance makes it impossible to judge whether action L is virtually justified or not. The only coordination failure this design might capture is *type II error*, in which player A relies on a player B who then uses the weakly dominated strategy. In order to fully account for the occurrence of coordination failures in the game, one thus needs to observe not only player As’ degree of non-reliance, but also player Bs’ behavior unconditional on what player As do. Our experiment implements the normal form of the game, so that decisions are elicited from both players in each interaction. Thanks to this design, we show that the empirical

(from 20% to even 80% across treatments, with an average of 54.5%) of player As prefer the mistrustful choice L , while the preference to maximize own gains is almost universal (97.8%) among subjects in the role of player B who are trusted by their partners. These results were confirmed on Japanese subjects in Beard, Beil, and Mataga (2001). Goeree and Holt (2001) applied the strategy method to the decision of player Bs. When asked what they would do if player A chose R , the odds that player Bs would choose r vary from 53% to 100% depending on the relative payoffs of each action. The rate of secure choices from subjects in the role of player As varies from 16% to 80%.

puzzle raised by Rosenthal’s conjecture is more sophisticated than previously thought. Although type I errors are more widespread than type II errors, non-reliance is often justified. This occurs because unreliability (through decision l) is widespread among player Bs, and insensitive to the treatment variables. In this sense, our study provides genuine evidence that a relatively weak degree of reliance from player As may be rational given the behavior of the population of player Bs. Coordination failures due to the strategic uncertainty faced by player As still prevail, however, since a lot of opportunities to reach the efficient outcome are missed. This is what enhanced information seeks to resolve.

Second, we observe that repetition-based learning stimulates reliance which, in turn, increases the frequency of the most efficient outcome. However, this improvement comes at a price in terms of coordination failures: since player As learn about the general population of player Bs rather than their actual interaction partner, more attempts to rely on the payoff-maximizing behavior fail. Any kind of subject-specific information, through cheap-talk or observation, appears to be a substitute for repetition-based learning. Both kinds of positive signals (through either an announcement to play the weakly dominant strategy or a reputation to have done so frequently in the past) provide an accurate screening of player Bs’ intentions and induce a rise in player As’ reliance. If signals are negative, on the other hand, communication appears to be more informative than observation: a non-reassuring message allows for better strategic coordination than an imperfect reputation.

Many experimental studies of communication and observation in coordination games report that such institutions help to overcome coordination failures and foster efficient cooperation.² In sharp contrast to those results, we report that although communication and observation restrict coordination failures, they do not significantly improve the overall efficiency of outcomes, which happens mostly due to the persistence of player Bs’ inefficient behavior. Developing an institutional environment that not only helps player As predict partners’ intentions, but also strongly disciplines player Bs’ actions, is an important challenge for future research.

2 Analysis of the game

In the one-shot simultaneous move game presented in Figure 1, one player has a weakly dominated strategy: for player B, playing l is never payoff improving. Because the dominance is only weak, all outcomes are rationalizable: for each action, there is some belief about opponent’s behavior to which this action is a best reply. Only two of them are pure-strategy Nash equilibria: (R, r) , which is payoff dominant; and (L, l) , which arises because L is the appropriate action against a player B choosing l . In fact, L is a risk dominant action for player A; and the strategic risk associated with playing the dominant action is quite extreme: while payoffs slightly increase should the most efficient outcome be reached when R is played, they strongly decrease compared with

²See, *e.g.*, Cooper, DeJong, Forsythe, and Ross (1992), Crawford (1998), Charness (2000), Duffy and Feltovich (2002), Duffy and Feltovich (2006), Blume and Ortmann (2007).

the safe choice L in the case of a coordination failure (R, l) . As stressed above, existing lab implementations of this game observed a high rate of non-reliance (increasing in the difference in payoff between the two equilibria) and little unreliability conditional on being relied on (decreasing in the difference in payoff between the two actions).

2.1 Learning

In the baseline treatment, we allow both players to experiment with these features of the game through trial and error across several repetitions of the one-shot game (see Section 3 for a detailed description of the design).³ Because the payoffs are asymmetric, the room for learning is very different for both players. On player Bs' side, one might think that the use of weakly dominated strategies mainly comes from inexperienced players. If that is true, we should observe a decrease in the share of decisions l as players gain some experience. For player As, by contrast, the choice between L and R depends on the anticipated behavior of their current interaction partners. As player As meet different player Bs from one play to the next, the baseline allows them to learn the actual distribution of behaviors in the population of player Bs.

2.2 Words

In the communication treatment, we introduce specific information about the current interaction partner through one-way, closed-form communication. The effect of communication on the outcomes of the game depends on the link between messages and future actions. To that regard, the theoretical properties of cheap-talk from player B to player A in this game are an illustration of the Aumann (1990) critique of Farrell (1988). Using the Farrell and Rabin (1996) nomenclature, the message "I will choose r " is self-committing because if it is enough to convince player A to play R , it subsequently creates the incentives for B to play in accordance with the message. But the message is not self-signalling: player B always wants player A to choose R , whether his true intention is to play r or not. A rational player A should thus anticipate that the message announcing action r is only meant to lead him to choose R and does not tell him anything about the true intentions of player B. As a result, communication should not matter in this game.

Ellingsen and Östling (2010) study cheap-talk in such contexts in terms of non-rationalizable beliefs. They formalize the idea originally introduced by Crawford (1998) that communication provides reassurance to the receiver about the degree of rationality of the sender. Ellingsen and Östling (2010) relax the assumption that players believe with probability 1 that their opponents are rational. The maintained assumption is that players have a weak (lexicographic) preference for being honest – *i.e.* players are truthful whenever they are indifferent between messages. In

³Although the original Rosenthal conjecture concerns a one-shot game, Beard and Beil note in their original paper (pp. 261-262) that it seems equally valid for a repeated game, and that learning through experience may affect people's behavior independently of payoff-related factors.

that case, all players signal their true intentions. In our game, this implies that the message “I will choose r ” can be used by player Bs to signal their ability to disregard weakly dominated strategies, hence reassuring player As about the actual risk induced by seeking for the payoff dominant outcome.

In the experimental literature devoted to the effect of cheap-talk communication, many papers focus on the Stag Hunt game, which shares some important features with the Rosenthal game studied here: there is a conflict between efficiency and strategic risk, and (in the most standard version of the game) one-way communication is self-committing but not self-signalling. In this context, one-way communication between players is generally found to foster efficiency to a very large extent – as compared to the outcomes without communication.⁴ Based on the theoretical properties discussed above, as well as this previous evidence, we thus expect cheap-talk to improve the outcomes of the game, through a reassurance effect that increases the completeness of information between players.

2.3 Actions

Our third experimental condition follows recent experimental literature contrasting the coordination properties of communication with the performance of observation of current interaction partner’s past actions. The most important difference with the information coming from cheap-talk communication is that past actions are not cheap: the signal is costly because it comes from actual decisions associated with variations in payoffs. Duffy and Feltovich (2002) introduce both cheap-talk and observation of partner’s most recent action in three 2 x 2 games: Prisoner’s Dilemma, Stag Hunt and Chicken.⁵ Stag Hunt is the closest to our setting. In the version of the Stag Hunt game they implement, the safe action leads to the same payoff regardless of opponent’s choice, so that messages are highly credible (*i.e.* both self-signalling and self-committing). This is not the case in the Chicken game, where messages are self-committing but not self-signalling (like in our game), but the nature of strategic interaction is different from the one we study here, because there is no Pareto dominant outcome. Duffy and Feltovich (2002) show that both treatments lead to an increase in frequency of the Nash equilibria and improve the efficiency of outcomes whatever the structure of the game. Cheap-talk appears to be more effective than observation in the Stag Hunt game – “words speak louder than actions”. Observation brings about better results in the

⁴See Crawford (1998) for an earlier survey of the theoretical and experimental literature on cheap-talk, Ellingsen and Östling (2010) for a detailed survey of the evidence in Stag Hunt games, and Cooper, DeJong, Forsythe, and Ross (1992) and Charness (2000) for a comprehensive experimental study.

⁵There are still few papers that implement this kind of feedback information. Bracht and Feltovich (2009) apply these two treatments in the Gift-Exchange Game. The results show a striking contrast between treatments: while observation is effective in reinforcing cooperation, the effect of communication visibly lags behind. See also Çelen, Kariv, and Schotter (2010) for further discussion of the effects of communication and observation on strategic behavior in the lab.

Chicken game – “actions speak louder than words”. The game we study is different from both of them: as in Stag Hunt, one equilibrium is payoff dominant, but as in Chicken, communication is not self-signalling. How observation and communication will perform to improve the outcomes in our game is thus still an open question. Our observation treatment moreover increases the amount of information compared with previous studies: we provide players with the full history of past decisions, rather than only the previous action. The design of this treatment therefore enables us to identify the weight of each historical component in the overall perception of reputation.

3 Experimental design

For the sake of replication, our core game is based on the original experiment of Beard and Beil (1994). Among the various payoff combinations they use, we chose treatment 1, presented in Table 1, due to its several attractive features: *(i)* as in the original setting, it does not lead to any conflict of interests between partners; *(ii)* the rate of player As’ unreliant choices related to this treatment is remarkable: 65.7%, and *(iii)* this is the only treatment where deviations from the dominant strategy by player B’s were observed (in 17% of all cases where player A made a reliant decision R).⁶ The Baseline condition implements a simple repetition of the one-shot game over an undefined number of rounds. We study the role of information in the game through two information boosting devices applied prior to the decision-making stage of every interaction: cheap-talk messages from Bs to As and historical information on how each player B acted in all the previous periods.

3.1 Baseline treatment

Our focus on the effect of information in the game led us to introduce several modifications to the original experiment: *(i)* the one-shot game is played repeatedly by experimental subjects, and *(ii)* we implement the normal form of the game rather than the genuine sequential form.

The experiment involves 10 rounds, each consisting of the core game presented in Table 1. Roles are fixed, so that each subject takes 10 decisions as either player A or player B. The experiment was designed so as to remain as close as possible to a one-shot game. First, the pairs are rematched after each round using a round-robin perfect stranger design (each session involves 20 subjects). Second, although the number of repetitions is pre-determined, we do not reveal it to subjects in

⁶Given the payoff structure of this game, one hypothesis is that social preferences explain player Bs’ behavior. In fact, even if the payoff dominant issue is reached, player B always earns less than player A: it could then be that player Bs take into consideration relative payoffs rather than their own earnings (see, *e.g.*, Fehr and Schmidt (1999)). We have tested this hypothesis through companion experiments, reported in Jacquemet and Zylbersztein (2011), in which the baseline treatment is compared with a treatment that equalizes payoffs between players in the Pareto-Nash equilibrium. We unambiguously reject the hypothesis that aversion to inequality is enough to account for player Bs’ striking behavior.

order to avoid end-game effects – in the experimental instructions we only inform them that the game contains several rounds. Last, we associate take-home earnings from the experiment with only one round out of ten. For that purpose, one round is randomly drawn at the end of the experiment (the same for all subjects).

To ensure the homogeneity of rounds despite repetition, we also modify the sequentiality of the game originally introduced by Rosenthal (1981). As pointed out by Binmore, McCarthy, Ponti, Samuelson, and Shaked (2002), the repetition of one-shot multi-stage games may induce some unwarranted heterogeneity and selection bias in observed behavior, because players are induced to distinguish between rounds based on the decisions made in earlier stages of the game. Unlike the original Beard and Beil (1994) experiment, we therefore ask both player A and player B to make a decision in each period. To make it as close as possible to the original sequential game, we describe the decision phase to subjects as follows: player A is first asked to choose between L and R , then player B chooses between l and r . Payoffs depend only on player A’s decision if L is chosen, or on both players’s decision otherwise.

To sum up, our Baseline treatment implements a repeated version of the one-shot game originally studied by Beard and Beil (1994). We use a perfect stranger design, the normal form of the game, an unknown termination rule and a one-round compensation rule to avoid subjects calculating the expected value of the entire game. This should induce players to maximize their utility in each repetition of the one-shot game. As a first step towards assessing the role of information in the Rosenthal puzzle, intra-comparisons in the Baseline treatment hence allow us to assess the robustness of the results to repetition. We further increase the amount of information in two subsequent experimental treatments.

3.2 Experimental treatments

Pre-play communication. This treatment allows player Bs to provide information about their intended play to player As. In every round, prior to the decision-making phase, player B has to send a message to player A. In the experimental implementation of cheap-talk messages, several trade-offs must be solved. As argued by, *e.g.*, Farrell and Rabin (1996), cheap-talk ought to be meaningful, *i.e.* to have a precise meaning. Messages of the “I will do ...”-type, which might be considered slightly oversimplified, are nonetheless highly meaningful. Voluntary free-form communication, by contrast, improves the informational content of communication but always gives the sender an opportunity to send an empty message or a message that is either meaningless or imprecise – which is hard to interpret for both the receiver and the experimenter.⁷ Given that our

⁷Experimental results from Charness and Dufwenberg (2006, 2008) substantiate the idea that impersonal messages, which have been prefabricated by the experimenter, work effectively in coordination games, whereas in trust games a more customized free-form communication seems to be needed. Similarly, Bochet, Page, and Putterman (2006) find that free-form communication yields higher efficiency in a VCM game than numerical messages.

primary goal is to boost the amount of information, we want to encourage players to communicate in a precise and clear manner. We hence implement a fixed-form communication and limit the set of possible messages to three options only, out of which two contain precise information, while the third is empty. Before any decision takes place in the round, player B is asked to choose one out of the three following messages:

☐ I will choose r

 ☐ I will choose l

 ☐ I will choose either l or r

by clicking on the relevant field on his computer screen. This message is then displayed on player A's computer screen. Once player A has confirmed reception of the message, the round moves to the decision phase. It is highlighted in the written instructions that messages are not binding (decisions from player Bs can be anything following any of the messages) and do not affect experimental earnings.

Observation of historical information. In the third condition, we allow subjects in the role of player As to inspect all the decisions made by their current interaction partner in all previous rounds. In every round, before the decision-making phase, player B is asked to wait while player A is provided with the history of choices made by player B. Following, *e.g.*, Bolton, Katok, and Ockenfels (2004) we make available the full history of past decisions rather than only the last one (see, *e.g.*, Bracht and Feltovich (2009), Duffy and Feltovich (2002)). In each round, player As thus receive a list of all the decisions made so far by their current interaction partner. Since pairs are rematched before each round, this information is updated and extended accordingly. Once player A has confirmed that he is aware of player B's history, the decision-making phase starts.

3.3 Experimental procedures

For each treatment, we ran three sessions involving 20 subjects each. Upon arrival, participants are randomly assigned to their computers and asked to fill in a small personal questionnaire containing basic questions about their age, gender, education, etc. The written instructions are then read aloud. Players are informed that they will play some (unrevealed) number of rounds of the same game, each round with a different partner, and that their own role will not change during the experiment. Before starting, subjects are asked to fill in a quiz assessing their understanding of the game they are about to play. Once the quiz and all remaining questions are answered, the experiment begins. Prior to the first round, players are randomly assigned to their roles – either A or B. They are then anonymously and randomly matched to a partner and asked for their choice, R or L for player As, and r or l for player Bs. At the end of each round, each subject is informed solely about his own payoff. Once all pairs have completed a round of the game, subjects are informed whether a new round will start. In this case, pairs are rematched according to a perfect stranger round-robin matching procedure (in which any pair meets only once in the session). At

the end of the experiment, one round is randomly drawn and each player receives the amount in euros corresponding to his gains in that round, plus a show-up fee of 5 euros.

All sessions were conducted in the laboratory of University Paris 1 Panthéon-Sorbonne (LEEP) between June 2009 and March 2010. Subjects were recruited via the LEEP database from among individuals who had successfully completed the registration process on the Laboratory’s website.⁸ The experiment involved a total group of 180 subjects, 90 males and 90 females.⁹ 86% of them were students, of whom 85 were likely to have some knowledge of game theory due to their field of study.¹⁰ 36% had never taken part in any economic experiment in LEEP before. Participants’ average age is roughly 24. No subject participated in more than one experimental session. Each session lasted about 45 minutes, with an average payoff of 12 euros (including a 5 euro show-up fee).

3.4 Statistical procedure for mean comparisons

The experimental design raises the issue of two kinds of correlation in the data. First, since players make a sequence of decisions, each subject’s choices might be serially correlated. Second, interaction partners change after each round of the experiment, which might result in an inter-subject correlation. To account for this structure of the data, we perform statistical tests for the comparisons of means through parametric regressions that assume clustered standard errors at the session level. This specification is asymptotically robust to any misspecification of the OLS residuals (Williams, 2000; Wooldridge, 2003). We also apply a delete-one jackknife correction in order to account for a potential small sample bias. Note that observations in the first round are still independent within and between sessions. We therefore use two-sided Fisher’s exact test for mean comparisons in round 1.

3.4.1 Standard errors estimation

The data are split into clusters (at the session level) and we denote i each observation in each cluster s , with $i=\{1, \dots, N_s\}$ and $s=\{1, \dots, S\}$ so that the total number of observations is $N = \sum_{s=1}^S N_s$. We perform statistical tests for differences between means through linear probability models of the form:

$$y_{is} = \sum_{k=0}^K \beta_k x_{is,k} + \epsilon_{is}$$

in which y_{is} is a dummy dependent variable, $X_{is} = \{1, x_{is1}, \dots, x_{isK}\}$ is the set of explanatory variables including the intercept, $\{\beta_0, \dots, \beta_K\}$ is the set of unknown parameters, and ϵ_{is} is the

⁸The recruitment uses ORSEE (Greiner, 2004); the experiment is computerized through software developed under REGATE (Zeiliger, 2000).

⁹This 50-50 spread of genders is purely incidental.

¹⁰Disciplines such as economics, engineering, management, political science, psychology, mathematics applied to social science, mathematics, computer science, sociology.

error term. We consider regressions on dummy variables reflecting changes in the environment (for instance, experimental treatments). Because the endogenous variable is itself binary, we also have that: $E(y|X) = Pr(y = 1|X)$. In this specification, the parameters thus reflect the mean change in the probability of the outcome induced by the change in the environment. In the case of one explanatory variable, $y_{is} = \beta_0 + \beta_1 I_{is} + \epsilon_{is}$, for instance: $E(y_{is}|I_{is} = 1) - E(y_{is}|I_{is} = 0) = Pr(y_{is} = 1|I_{is} = 1) - Pr(y_{is} = 1|I_{is} = 0) = \beta_1$, so that the parameter measures the mean variation in the probability of y . Two-sided t -tests on each treatment-related parameter thus provide significance levels of the differences in means with respect to baseline condition.

To compute the standard errors, we allow for dependences inside clusters as well as unspecified heteroscedasticity across observations,¹¹ *i.e.* we assume that any two error terms i and j are independent between clusters, $Cov(\epsilon_{ig}, \epsilon_{jh}) = 0 \forall g \neq h$, but allow for any type of dependence within a cluster, $Cov(\epsilon_{ig}, \epsilon_{jg}) = \sigma_{ijg}^2 \forall i, j, g$. To that end, we correct the estimated covariance matrix at the cluster level using the following procedure, in which the model is written at the cluster level, $Y_s = X_s \beta + \epsilon_s$, where Y_s and ϵ_s are $[N_s \times 1]$ vectors, X_s is a $[N_s \times (K+1)]$ matrix, β is a $[(K+1) \times 1]$ vector:

1. Using the parameters estimated on pooled data, $\hat{\beta}_{OLS} = (X'X)^{-1}(X'Y)$, we calculate the vector of error terms in each cluster:

$$\hat{\epsilon}_s = Y_s - X_s \hat{\beta}_{OLS}$$

2. We then estimate the cluster robust covariance matrix (CRCME):

$$\hat{V}_{CRCME} = (X'X)^{-1} \left(\sum_{s=1}^S X_s' \hat{\epsilon}_s \hat{\epsilon}_s' X_s \right) (X'X)^{-1} \quad (1)$$

3.4.2 Correction for small sample bias in standard errors

The procedure described above provides a consistent estimator of the covariance matrix which can typically be biased in small samples. What is more, the bias is generally found to be negative, so that significance tests reject the null hypothesis too often. A first way to deal with this issue is to correct for the degrees of freedom by substituting $\tilde{\epsilon}_s = \sqrt{C_{df}} \hat{\epsilon}_s$, with $C_{df} = \frac{S(N-1)}{(S-1)(N-K)}$, in (1) – a procedure known in the literature as HC1. Bell and McCaffrey (2002) and Cameron, Gelbach, and Miller (2008) propose a more accurate correction, called HC3, which estimates the residuals as $\tilde{\epsilon}_s = \sqrt{\frac{S-1}{S}} [I_{N_s} - H_{ss}]^{-1} \hat{\epsilon}_s$, where I_{N_s} is a $[N_s \times N_s]$ identity matrix, and $H_{ss} = X_s (X_s' X_s)^{-1} X_s'$.

¹¹Heteroscedasticity is due to the linear probability specification. Even if the data generating process was *i.i.d* (*i.e.* $V(u_{is}) = \sigma^2$, and $E(u_{is}u_{jt}) = 0 \forall i \neq j$ and $\forall t$) the model entails that: $V(y|X) = Pr(y = 1|X)[1 - Pr(y = 1|X)] = X\beta(1 - X\beta)$.

Table 2: Summary of experimental evidence on Rosenthal’s game

Experiment	Payoff			Observed outcomes				Nb.	
	(L)	(R, r)	(R, l)	L	(R, r)	(R, l)	(r R)	r	obs.
Beard and Beil (1994)-Tr.1	(9.75; 3)	(10; 5)	(3; 4.75)	66%	29%	6%	83%	—	35
Beard et al. (2001)-Tr.1	(1450; 450)	(1500; 750)	(450; 700)	79%	18%	3%	86%	—	34
Goeree and Holt (2001)-Tr.2	(80; 50)	(90; 70)	(20; 68)	52%	36%	12%	75%	—	25
Goeree and Holt (2001)-Tr.3	(400; 250)	(450; 350)	(100; 348)	80%	16%	4%	80%	—	25
Baseline, round 1	(9.75; 3)	(10; 5)	(3; 4.75)	77%	23%	0%	100%	80%	30
Baseline, rounds 2-10	(9.75; 3)	(10; 5)	(3; 4.75)	48%	43%	9%	84%	81%	270
Baseline, overall	(9.75; 3)	(10; 5)	(3; 4.75)	51%	41%	8%	84%	81%	300

Note. The monetary payoffs displayed in the first three columns are in USD in Beard & Beil (1994), in cents of USD in Goeree & Holt (2001), in yens in Beard et al. (2001) and in euros in our treatments.

For an OLS regression, this corrected variance-covariance matrix amounts to implementing a delete-one jack-knife procedure:

$$\tilde{V}_{jackknife} = \frac{S-1}{S} \sum_{s=1}^S \left(\tilde{\beta}_{-s} - \hat{\beta} \right) \left(\tilde{\beta}_{-s} - \hat{\beta} \right)' \quad (2)$$

where $\tilde{\beta}_{-s}$ is the vector of coefficients estimated after leaving out the s th cluster.¹²

4 Results

The last three rows of Table 2 provide a summary of observed behavior in our baseline treatment along with results from previous experimental studies using the same game (top part of the Table). Our results are generally in line with what has been observed in other studies, despite the differences in the design described in Section 3: the average rate of non-reliance is 51%, very close to the one observed in Goeree and Holt (2001) who apply the strategy method to the sequential game. Unsurprisingly, the one-shot sequential games of Beard and Beil (1994); Beard, Beil, and Mataga (2001) are much better replicated by the first round of our baseline treatment than by the overall rate produced by the repetition of the game. Moreover, in all studies the likelihood of observing action r is alike and equal to roughly 80%. Once all repetitions of our baseline treatment are pooled, the outcomes are consistent with all the previous studies in terms of both

¹²All p -values presented in the section below are associated with statistics computed according to this HC3 procedure. We also ran robustness checks by implementing the HC1 correction, which generally leads to lower estimated standard errors. Our choice is thus conservative as regards our ability to find significant differences in behavior. Based on a correction closely related to the HC3 procedure, Angrist and Lavy (2009) find an inflation of the cluster-robust standard errors by 10% up to 50%.

Table 3: Overall effects of information treatments

	Decisions		Outcomes		Errors		
	Reliant A	Reliable B	Cooperation	Coordination	Type I	Type II	
	(R)	(r)	(R, r)	(R, r) $\cup$ (L, l)	(L, r)	(R, l)	
Baseline	1	0.233	0.800	0.233	0.433	0.567	0.000
	2-4	0.456	0.844	0.411	0.522	0.433	0.044
	5-7	0.589	0.767	0.489	0.622	0.278	0.100
	8-10	0.511	0.811	0.400	0.478	0.411	0.111
	Average	0.490	0.807	0.413	0.530	0.393	0.077
Communication	1	0.500	0.800	0.433	0.567	0.367	0.067
	2-4	0.467	0.811	0.422	0.567	0.389	0.044
	5-7	0.678	0.778	0.589	0.722	0.189	0.089
	8-10	0.667	0.811	0.600	0.722	0.211	0.067
	Average	0.593	0.800	0.527	0.660	0.273	0.067
Observation	1	0.167	0.767	0.033	0.133	0.733	0.133
	2-4	0.544	0.833	0.511	0.644	0.322	0.033
	5-7	0.578	0.800	0.511	0.644	0.289	0.067
	8-10	0.589	0.844	0.544	0.656	0.300	0.044
	Average	0.530	0.820	0.473	0.597	0.347	0.057

Note. For each treatment, we separate data according to the stage of the game – round 1, rounds 2-4, rounds 5-7, rounds 8-10, and the overall results (consecutive rows). In each row, we provide empirical frequencies of decision R by player A, decision r by player B, cooperation on the most efficient Nash equilibrium (R, r), coordination on the existing Nash equilibria (R, r) or (L, l), Type I error (L, r) and Type II error (R, l) (columns 1-6, respectively).

efficient coordination – outcome (R, r) – and strategic mismatches resulting in outcome (R, l). In what follows, we discuss whether and how information helps to improve efficiency and overcome coordination failures.

4.1 Aggregate treatment effects

Table 3 summarizes aggregate behavior elicited in each of the three treatments. For each treatment, we separate data into five categories related to the stage of the game – the initial round, rounds 2-4, rounds 5-7, rounds 8-10, and finally the overall results (consecutive rows). The first two columns of the table summarize unconditional average behavior of player As and Bs. The right-hand side describes the resulting outcomes: positive ones (efficiency and coordination) in the two middle columns, and failures in the last two. Thanks to our design, we are able to observe two sources of coordination failure: beyond the outcome arising when player A mistakenly relies on player B, which we classify as type II errors, we also identify type I errors – *i.e.* cases where player A should have relied on player B, since player B would have proved reliable in this case – resulting in outcome (L, r).

The likelihood of each outcome depends on both players' behavior. As shown in the second

Table 4: Statistical support to Table 3

	Reliant A Pr(R)		Reliable B Pr(r)		Cooperation Pr(R, r)		Coordination Pr($(R, r) \cup (L, l)$)		Type I errors Pr(L, r)		Type II errors Pr(R, l)	
	coef	p-val.	coef	p-val.	coef	p-val.	coef	p-val.	coef	p-val.	coef	p-val.
Model 1												
Intercept	0.233	0.000	0.800	0.000	0.233	0.000	0.433	0.003	0.567	0.001	0.000	0.410
BT_rounds2-4	0.222	0.000	0.044	0.820	0.178	0.001	0.089	0.620	-0.133	0.484	0.044	0.009
BT_rounds5-7	0.356	0.000	-0.033	0.717	0.256	0.000	0.189	0.173	-0.289	0.032	0.100	0.002
BT_rounds8-10	0.278	0.001	0.011	0.890	0.167	0.022	0.044	0.764	-0.156	0.266	0.111	0.043
CT	0.267	0.178	0.000	1.000	0.200	0.202	0.133	0.491	-0.200	0.385	0.067	0.122
CT_rounds2-4	-0.033	0.836	0.011	0.880	-0.011	0.941	0.000	1.000	0.022	0.898	-0.022	0.122
CT_rounds5-7	0.178	0.521	-0.022	0.838	0.156	0.563	0.156	0.528	-0.178	0.476	0.022	0.122
CT_rounds8-10	0.167	0.555	0.011	0.816	0.167	0.553	0.156	0.531	-0.156	0.535	0.000	1.000
OT	-0.067	0.256	-0.033	0.816	-0.200	0.006	-0.300	0.025	0.167	0.164	0.133	0.122
OT_rounds2-4	0.378	0.077	0.067	0.290	0.478	0.001	0.511	0.000	-0.411	0.007	-0.100	0.332
OT_rounds5-7	0.411	0.013	0.033	0.412	0.478	0.001	0.511	0.000	-0.444	0.005	-0.067	0.172
OT_rounds8-10	0.422	0.056	0.078	0.308	0.511	0.001	0.522	0.000	-0.433	0.006	-0.089	0.365
Model 2												
Intercept	0.490	0.000	0.807	0.000	0.413	0.000	0.530	0.000	0.393	0.000	0.077	0.000
CT	0.103	0.110	-0.007	0.941	0.113	0.211	0.130	0.018	-0.120	0.010	-0.010	0.779
OT	0.040	0.744	0.013	0.882	0.060	0.659	0.067	0.403	-0.047	0.507	-0.020	0.202
Model 3												
Intercept	0.233	0.000	0.800	0.000	0.233	0.000	0.433	0.003	0.567	0.001	0.000	0.658
BT_rounds2-10	0.285	0.000	0.007	0.952	0.200	0.000	0.107	0.458	-0.193	0.203	0.085	0.000
CT	0.267	0.178	0.000	1.000	0.200	0.202	0.133	0.491	-0.200	0.385	0.067	0.122
CT_rounds2-10	0.104	0.656	0.000	1.000	0.104	0.648	0.104	0.626	-0.104	0.637	0.000	1.000
OT	-0.067	0.256	-0.033	0.816	-0.200	0.006	-0.300	0.025	0.167	0.164	0.133	0.122
OT_rounds2-10	0.404	0.037	0.059	0.318	0.489	0.001	0.515	0.000	-0.430	0.005	-0.085	0.289

Note. Columns summarize the results of session-clustered (9 clusters in total, 100 observations per cluster, standard errors corrected with delete-one jackknife) OLS regressions of treatment-related variables on player A's decision R , player B's decision r , and outcomes: cooperative (R, r), coordinated (R, r) $\cup$ (L, l), Type I errors (L, r) and Type II errors (R, l) (columns 1-6, respectively). The intercept represents the reference frequency in round 1 of the Baseline treatment; dummies CT and OT correspond to the change in the intercept due to the Communication and observation treatments. Remaining coefficients are interpreted as an absolute change in the frequency of dependant variables due to achieving certain stages of the game. Prefix BT stands for the baseline treatment, CT and OT for the Communication and observation treatments, respectively. p -values come from two-tailed t -tests for nullity of coefficients.

column of Table 3, the behavior of player Bs is fairly stable regardless of time and experimental treatments.¹³ As a result, any difference between treatments and between rounds is very unlikely to be driven by changes in player Bs' behavior. If anything, the treatment effects of information occur because of changes in the way player As perceive player Bs, rather than through differences between populations of player Bs.

Table 4 provides the results of several specifications of the linear probability model described in Section 3.4. Model 1 summarizes the changes due to repetition-based learning. Models 2 and 3 provide statistical tests of the effects of treatments. To sum up, the repetition of the one-

¹³Fisher's exact test does not reject the null hypothesis that player Bs' decisions in round 1 come from the same distribution in all treatments ($p=1.000$). Model 3 in Table 4 suggests that the average proportion of decisions r in rounds 2-10 does not significantly differ from the initial round BT in either treatment: $p=0.952$; CT: $p=1.000$, OT: $p=0.318$. Finally, on the basis of Model 1 in Table 4 we also test a joint hypothesis that the means in all treatments are statistically different in rounds 2-4 through $H_0: (BT_rounds2-4 = CT + CT_rounds2-4) \cap (BT_rounds2-4 = OT + OT_rounds2-4)$. No difference arises either in this early stage ($p=0.925$), or in rounds 5-7 ($p=0.932$) or rounds 8-10 ($p=0.917$).

shot game improves efficiency in the baseline, but at the same time increases the odds of type II errors. Specific information (through communication or observation) reduces the likelihood of coordination failures, but does not improve the efficiency of outcomes. We comment below on the main driving forces behind these results. We then turn to the way specific information about the interaction partner is taken into account by player As.

We first focus on repetition-based learning by comparing outcomes across rounds within the baseline treatment. As shown in Table 3, the rate of reliant decisions from player As more than doubles between round 1 and the subsequent occurrences of the game. Given the stability of player Bs' actual decisions, this suggests that over time player As update their beliefs about the population of player Bs. The main effect of this change in player As' behavior is an important improvement in the share of the efficient outcome, from 23% of first round outcomes to 43% of subsequent repetitions (the difference is significant at the 5% level according to Model 1). This comes at a price in terms of coordination failure: while the risk of type I errors falls (significantly only for rounds 5-7), type II errors become more likely (the rise is significant at the 5% level for all triplets of periods).

In the communication treatment, all outcomes become much less sensitive to the repetition of the game. As compared to the baseline situation, cheap-talk strongly increases the reliance rate in the first round, from 23% in the baseline treatment to half of decisions in the communication treatment ($p=0.060$ using Fisher's exact test), and only slightly in further repetitions of the game, from 52% to 60% ($p=0.269$ after testing $H_0 : BT_rounds2 - 10 = CT + CT_rounds2 - 10$ in Model 3). The overall increase in the share of efficient outcomes compared with the baseline is not significant ($p=0.211$, see Model 2). Note, however, that the proportion of efficient outcomes in the communication treatment in the first round is already very close to the one attained due to repetition in the baseline condition. The main effect of cheap-talk communication is an improvement in coordination which is significant at the 5% level (see Model 2). While the likelihoods of both type I and type II errors decrease, only the former is statistically significant, *i.e.* player As are less likely to mistakenly choose the secure option.

In contrast to communication, the observation treatment does not provide specific information to player As at the beginning of the game – the signal becomes available in round 2. The treatment appears to be anticipated by player As: the rate of reliance is lower in the first round compared with the baseline treatment, which hinders both cooperation ($p=0.052$ using Fisher's exact test) and coordination ($p = 0.020$). Once information becomes available – in rounds 2-10 – outcomes improve compared with the first round of the baseline treatment to reach levels similar to those observed after several repetitions.¹⁴ The only exception is that type II errors are significantly less frequent than in rounds 2-10 of the baseline ($p = 0.015$), so that player As' risk of relying in

¹⁴Based on Model 3, we assess the effect of observation against baseline through tests of $H_0 : BT_rounds2 - 10 = OT + OT_rounds2 - 10$. The differences are insignificant as regards reliance ($p = 0.709$), cooperation ($p = 0.544$) coordination ($p = 0.237$) and type I errors ($p = 0.421$)

vain on a partner falls. Compared with the communication treatment, the proportion of actions R in the observation treatment falls drastically in the first round ($p = 0.013$ using Fisher's exact test), which boosts the rate of type I errors ($p=0.009$) and reduces coordination and cooperation (both $p \leq 0.001$). In subsequent rounds, however, both experimental conditions provide similar outcomes.¹⁵

4.2 Informational content of signals

Table 5 reorganizes data according to the flow of information. As a benchmark, the first column of the table summarizes the outcomes observed in the baseline treatment when all rounds are pooled. For the communication treatment, the observations are conditioned on the message received by player A – "*I will choose r* ", "*I will choose l* ", "*I will choose either l or r* ". For the observation treatment, we use the reputation of each player B to separate the population into two groups: highly reliable ones and others. For that purpose, we construct a reputation index for each player B equal to the rate of decisions r amongst all decisions made prior to the current round. We classify player Bs in each round by comparing their reputation to the cut-off probability of decision r (0.964) that makes a risk-neutral player A indifferent between choosing L and R . Hence, prior to entering an interaction, each player B may either have a perfect reputation (denoted B_P) or an imperfect one (B_{IP}).¹⁶ The last column provides observations from the first round, in which no information is available.

For each treatment involving specific information about partner, this classification can thus organize data according to three kinds of informational content delivered to player As: a positive signal – a reassuring message $m(r)$ in the communication treatment and a perfect reputation B_P in the observation treatment, a negative signal – one of the non-reassuring messages $m(l)$, $m(l/r)$, and an imperfect reputation B_{IP} , or lack of information – like in the first round of observation treatment. Table 6 summarizes statistical tests for differences of proportions conditional on source and content of information, based on parametric regressions of outcomes on signal-related dummy variables discussed above.

In the communication treatment, the frequencies of empty messages and messages announcing the weakly dominated action are both equal to around 12%. 90% of player Bs who announce they will select r actually do so, while any other message makes the likelihood of choosing r fall

¹⁵We use Model 1 to test the joint hypothesis that in every triplet of rounds – 2-4, 5-7 and 8-10 – a given outcome is equally frequent in both treatments, that is $H_0 : (CT + CT_rounds2 - 4 = OT + OT_rounds2 - 4) \cap (CT + CT_rounds5 - 7 = OT + OT_rounds5 - 7) \cap (CT + CT_rounds8 - 10 = OT + OT_rounds8 - 10)$. We find $p = 0.688$ for reliance, $p = 0.669$ for cooperation, $p = 0.360$ for coordination, $p = 0.531$ for type I error, and $p = 0.949$ for type II error.

¹⁶Note, this way of separating player Bs implies that the first group comprises only those players that constantly played r before the current interaction. As a result, any player B with a perfect record who chooses l once in the game drops out from this category permanently, and becomes B_{IP} ever after.

Table 5: Informational content of signals

	Baseline	Communication			Observation		
		$m(r)$	$m(l)$	$m(l/r)$	B_P	B_{IP}	Unknown
Frequency within treatment	100%	75.7%	12.0%	12.3%	49.3%	40.7%	10%
Reliant A (R)	49.0%	72.2%	16.7%	21.6%	77.7%	32.0%	16.7%
Reliable B (r)	80.7%	90.3%	38.9%	56.8%	93.9%	68.9%	76.7%
Cooperation (R, r)	41.3%	65.6%	8.3%	16.2%	75.7%	23.8%	3.3%
Coordination ($((L, l) \cup (R, r))$)	53.0%	68.7%	61.1%	54.1%	79.7%	46.7%	13.3%
Type I error (L, r)	39.3%	24.7%	30.6%	40.5%	18.2%	45.1%	73.3%
Type II error (R, l)	7.7%	6.6%	8.3%	5.4%	2.0%	8.2%	13.3%
Nb of observations	300	227	36	37	148	122	30

Note. For each treatment in column, the rows provide the proportion of observed decisions (first two rows) and outcomes (last four rows). The first column pools all observations from the baseline. In the middle columns, data from all rounds of the communication game are split according to the message received by player A: "I will choose r ", denoted $m(r)$, "I will choose l ", $m(l)$, and "I will choose either l or r ", $m(l/r)$. For the observation game (right-hand side of the table), observations are classified according to the reputation of player B: in rounds 2-10, the reputation is perfect, and denoted B_P , if all previous decisions are r ; and imperfect otherwise, denoted B_{IP} . Reputation is unknown in round 1.

by 42 percentage points ($p=0.045$ after testing $H_0 : CT_ReassMess = CT_NonReassMess$) – only 39% (57%) of those sending message "I will choose l " ("I will choose either l or r ") selects r . Similarly, reputation is not a perfect predictor of reliability. Nonetheless, 94% of player Bs entering an interaction with a full record of weakly dominant actions continue to behave this way, while only 69% of the subjects with an imperfect reputation select r ($p < 0.001$ after testing $H_0 : OT_PerfRep = OT_ImPerfRep$). Thus, in both treatments positive signals provide an accurate screening of player Bs' intentions. The rate of reliable partners is much higher among those delivering such a signal – through either the message announcing a play r or a perfect reputation – than among others. Compared with the baseline treatment, the reliability rate among player Bs announcing decision r in the communication treatment is slightly improved with respect to the baseline ($p=0.249$), and significantly decreases among player Bs sending one of the two other messages ($p=0.002$). The reliability rate among players with a perfect reputation in the observation treatment is also improved with respect to the baseline, this time significantly ($p = 0.067$), while an imperfect reputation is only slightly detrimental in this respect ($p=0.242$).

Furthermore, player As appear to account for this information by relying more on player Bs delivering a positive signal: from 49% in the baseline, the reliance rate rises to 72% against a reassuring message and 78% against perfect reputation. As a result, any positive signal induces a significant increase in the rate of cooperation and coordination, along with a fall in type I errors (based on the corresponding regression models in Table 6, all comparisons with the baseline are significant at the 5% level, with exception of type I error in the observation treatment, for which $p = 0.062$). At the same time, both types of negative signals substantially decrease the likelihood of achieving the most efficient outcome (R, r) (both $p < 0.050$).

Table 6: Statistical support to Table 5

	Reliant A Pr(R)		Reliable B Pr(r)		Cooperation Pr(R, r)		Coordination Pr($(R, r) \cup (L, l)$)		Type I errors Pr(L, r)		Type II errors Pr(R, l)	
	coef	p-val.	coef	p-val.	coef	p-val.	coef	p-val.	coef	p-val.	coef	p-val.
Intercept	0.490	0.000	0.807	0.000	0.413	0.000	0.530	0.000	0.393	0.000	0.077	0.000
CT_ReassMess	0.232	0.009	0.096	0.249	0.243	0.019	0.157	0.038	-0.147	0.039	-0.011	0.813
CT_NonReassMess	-0.298	0.004	-0.327	0.002	-0.290	0.006	0.045	0.362	-0.037	0.521	-0.008	0.608
OT_Round1	-0.323	0.001	-0.040	0.682	-0.380	0.001	-0.397	0.000	0.340	0.000	0.057	0.486
OT_PerfRep	0.287	0.049	0.133	0.067	0.343	0.033	0.267	0.036	-0.211	0.062	-0.056	0.008
OT_ImPerfRep	-0.170	0.025	-0.118	0.242	-0.176	0.048	-0.063	0.112	0.057	0.159	0.005	0.840

Note. In each column, we present the results of session-clustered (9 clusters in total, 100 observations per cluster) OLS regressions of treatment-related variables on player A's decision R , player B's decision r , cooperative outcome (R, r) , coordinated outcome $(R, r) \cup (L, l)$, Type I error (L, r) and Type II error (R, l) (columns 1-6, respectively). The intercept represents the frequency in the baseline treatment, the coefficient of each explanatory variable is interpreted as a change induced by a given signal: reassuring (CT_ReassMess) and non-reassuring (CT_NonReassMess) message in the Communication Treatment; perfect (OT_PerfRep) and imperfect (OT_ImPerfRep) reputation in the Observation Treatment; OT_Round1 controls for the effect of the first round of the Observation Treatment, where no signals were available yet.

When comparing communication with observation, we find that the performance of positive signals is similar in terms of both cooperation and coordination, although we observe quantitatively higher rates following a perfect reputation rather than a reassuring message – 76% for coordination on the most efficient equilibrium following the former, against 67% following the latter; and 80% for coordination on either of the two Nash equilibria following a perfect reputation, against 69% following a reassuring message.¹⁷

Interesting differences arise, by contrast, as regards the effect of negative signals. Player Bs with an imperfect reputation are substantially more reliable than those sending a non-reassuring message – the difference in the likelihood of decision r is about 21 percentage points and statistically significant ($p=0.045$).¹⁸ Although player As take this into account (the analogous difference between the average likelihoods of reliance amounts to 13 percentage points, with $p=0.093$), we nonetheless find that an imperfect reputation deteriorates coordination as compared to the baseline, while non-reassuring messages improve the odds that player A accurately accommodates partner's true intentions. Comparing the two treatments, the difference in coordination amounts to 11 percentage points ($p = 0.056$).

4.3 The effects of long-term observation

In Table 7, we further investigate how subjects use the environment allowing for a long-term observation. We break down player Bs' reputation in the observation treatment (restricted to rounds 3-10) according to history (H) – the very last observation – and pre-history (PH) – the proportion of decisions r among all the remaining observations (preceding the last round). For

¹⁷None of these differences are significant, though: testing $H_0 : CT_ReassMess = OT_PerfRep$ gives $p=0.475$ for cooperation, $p=0.383$ for coordination, $p=0.572$ and $p=0.326$ for type I and type II errors, respectively.

¹⁸All comparisons made in this paragraph are based on tests of $H_0 : CT_NonReassMess = OT_ImPerfRep$ in the regressions in Table 6.

Table 7: Marginal effect of components of past history in the observation treatment

	N	Pr(R)		Pr(r)		Pr(R, r)		Pr($(R, r) \cup (L, l)$)	
		coef	p-val.	coef	p-val.	coef	p-val.	coef	p-val.
$H(r) \times PH(1)$	125	0.808	0.009	0.960	0.000	0.784	0.012	0.800	0.012
$H(r) \times PH [0.7; 0.9[$	29	-0.291	0.152	-0.132	0.336	-0.336	0.205	-0.248	0.251
$H(r) \times PH [0.5; 0.7[$	28	-0.487	0.063	-0.174	0.139	-0.498	0.052	-0.336	0.038
$H(r) \times PH [0; 0.5[$	17	-0.455	0.015	-0.489	0.068	-0.666	0.029	-0.388	0.091
$H(l) \times PH(1)$	9	-0.586	0.134	-0.182	0.106	-0.673	0.043	-0.578	0.069
$H(l) \times PH [0.7; 0.9[$	8	-0.433	0.060	-0.335	0.208	-0.534	0.148	-0.300	0.249
$H(l) \times PH [0.5; 0.7[$	8	-0.558	0.158	-0.210	0.408	-0.534	0.170	-0.300	0.079
$H(l) \times PH [0; 0.5[$	16	-0.683	0.016	-0.335	0.015	-0.722	0.018	-0.425	0.164

Note. The Table summarizes the results – estimated coefficients and corresponding p -values – of a session-clustered (3 clusters from the Observation Treatment, 80 observations per cluster, standard errors corrected with a delete-one jackknife) OLS regressions of dummies representing player A’s decision R , player B’s decision r , cooperative outcome (R, r) and coordinative outcome $(R, r) \cup (L, l)$ (columns 3-6, respectively) on a set of reputation-related variables. Only observations from rounds 3-10 from the observation treatment are included. For instance, variable $H(r)PH[0.5; 0.7[$ stands for a reputation in which the decision from the most recent round (history) is r , and the proportion of decisions r in the remaining rounds (pre-history) is greater or equal to 0.5 and less than 0.7.

each decision observed recently – either l or r – we consider the marginal effect of four intervals of proportions of decisions r in the remaining rounds: $[0; 0.5[$, $[0.5; 0.7[$, $[0.7; 0.9[$, 1 . We find that both components of reputation, historical and pre-historical, are taken into account by player As, and that consistent signals in which both elements point to the same behavior facilitate coordination to a greater extent than more ambiguous ones in which the two components diverge.

When the previous decision is r , reliance decreases as decisions r become less frequent in earlier stages (the coefficients of $H(r) \times PH \in [0; 0.5[$ and $H(r) \times PH \in [0.5; 0.7[$ are both negative with $p = 0.015$ and $p = 0.063$, respectively). At the other extreme, decision l in the previous round has a negative but statistically insignificant effect on player A’s behavior when reputation is perfect in earlier stages (the p -value for the coefficient of $H(l) \times PH(1)$ is 0.134). A bad pre-historical reputation (lower than 0.5), in turn, induces player As to rely on the most recent decision they observe ($p = 0.025$).¹⁹

As regards cooperation and coordination, reputations $H(r) \times PH [0; 0.5[$ and $H(r) \times PH [0.5; 0.7[$ lead to a substantial decrease in the likelihoods of both outcomes compared with a perfect reputation $H(r) \times PH(1)$ (cooperation: $p=0.029$, $p=0.052$; coordination: $p=0.091$, $p=0.038$, respectively). For a perfect pre-history, we observe a substantial positive effect of observing the most recent decision r instead of l , on the probability of achieving both (R, r) ($p=0.043$) and $(R, r) \cup (L, l)$ ($p=0.069$). Furthermore, reputations $H(l) \times PH [0; 0.5[$ and $H(l) \times PH [0.5; 0.7[$, where both pre-history and history are imperfect, provide a better screening of player Bs’ true intentions,

¹⁹These results are obtained through an additional t -test for equality of coefficients. For instance, in the latter case we test $H_0 : H(r) \times PH [0; 0.5[= H(l) \times PH [0; 0.5[$ in regression in column 3.

Table 8: Probit regressions on player As' reliance in the three experimental games

	Baseline			Communication			Observation		
	coef	p-val.	ME	coef	p-val.	ME	coef	p-val.	ME
Intercept	-3.56	0.03	—	-0.97	0.59	—	-1.77	0.12	—
Male_A	0.71	0.16	0.28	0.43	0.16	0.16	0.53	0.11	0.20*
1 _{Round=3}	0.35	0.42	0.14	0.00	0.99	0.00	0.37	0.39	0.14
1 _{Round=4}	-0.06	0.90	-0.02	-0.04	0.93	-0.01	0.28	0.58	0.10
1 _{Round=5}	0.07	0.88	0.03	0.17	0.68	0.06	0.00	0.99	0.00
1 _{Round=6}	-0.24	0.66	-0.10	0.07	0.88	0.03	-0.42	0.41	-0.16
1 _{Round=7}	-0.75	0.21	-0.27	-0.68	0.18	-0.26	-0.01	0.98	-0.01
1 _{Round=8}	-1.12	0.08	-0.38**	-0.24	0.65	-0.09	-0.17	0.79	-0.07
1 _{Round=9}	-1.12	0.11	-0.38**	-1.03	0.08	-0.39*	-0.69	0.30	-0.27
1 _{Round=10}	-2.07	0.01	-0.54***	-0.97	0.14	-0.37	-1.08	0.13	-0.41*
1 _{Trusted Before}	0.33	0.60	0.13	0.15	0.74	0.06	-0.10	0.83	-0.04
. $\times$ Round $\times$ Ratio_Rr	0.28	0.02	0.11**	0.23	0.01	0.08***	0.24	0.01	0.09***
Population_B	3.20	0.08	1.28*	-1.05	0.55	-1.28	-0.47	0.93	-0.05
Message_r	—	—	—	1.75	0.00	0.62***	—	—	—
Message_l	—	—	—	-0.30	0.46	-0.12	—	—	—
Reputation $\in$ (0.5; 0.6]	—	—	—	—	—	—	-0.71	0.42	-0.28
Reputation $\in$ (0.6; 0.7]	—	—	—	—	—	—	-0.03	0.95	-0.01
Reputation $\in$ (0.7; 0.8]	—	—	—	—	—	—	0.19	0.66	0.07
Reputation $\in$ (0.8; 0.9]	—	—	—	—	—	—	0.92	0.03	0.29**
Perfect_Reputation	—	—	—	—	—	—	1.92	0.00	0.66***
N	270			270			270		
Log L	-110.68			-117.73			-120.77		

Legend. *, **, *** indicate statistical significance at the 10%, 5%, 1% level.

Note. Panel probit regressions with individual random effects. The dependent variable is a dummy indicating whether player A makes a reliant decision R , observations come from rounds 2-10. Marginal effects (ME) are calculated under the assumption that the individual factors u_i are equal to 0. Male_A indicates whether player A is a male. The variables 1_{Round=3}, ..., 1_{Round=10} are round fixed effects – round 2 is omitted. 1_{Trusted Before} is a dummy variable switched to one once player A has chosen R in the past. Population_B is computed in each round (except for round 1) for every experimental session as the proportion of decisions r among all past decisions made by the population of player Bs. For the communication treatment (middle part), the flow of information is included through the content of received message ("I will choose r " or "I will choose l "). For the observation treatment (right-hand side) the reputation of the current interaction partner is accounted for through dummies for each 10% range (past reliability less than or equal to 50% is the reference).

which increases the likelihood of achieving a coordinated outcome, than a reputation $H(l) \times PH(1)$, where only the most recent decision is imperfect ($p=0.059$, $p=0.079$, respectively).²⁰

4.4 The determinants of reliance

We test the marginal significance of the effects highlighted above by estimating treatment-specific models of the probability of observing a reliant decision from player As. We include individual random effects to account for individual heterogeneity and condition observed behavior on time-dummies, observed past behavior and both group and individual information about player Bs. This conditioning controls for player As' common experience due to interactions with the same

²⁰These results are obtained through additional tests for equality of coefficients in the regression in column 6: $H_0 : H(l) \times PH(1) = H(l) \times PH[0.5; 0.7]$ and $H_0 : H(l) \times PH(1) = H(l) \times PH[0.7; 0.9]$, respectively.

player Bs over time. Due to the introduction of lagged variables, all models are estimated on data from rounds 2-10. The results from probit regressions are presented in Table 8.

In the baseline treatment, the round dummies indicate small variations over time. This effect of time reflects repetition-based learning about player Bs' behavior. The variable *Population_B* is constructed as the proportion of decisions r among all decisions made in the entire population of player Bs in the earlier rounds of the experimental session – this measures global variations from one round to the next. The effect of this variable thus measures how the true behavior of the population of player Bs is accommodated for in the experiment – although this information is never available as such to player As in this treatment. Both the coefficient and the marginal effect are significant. In the baseline treatment, repetition-based learning induces a positive correlation between reliance and the overall rate of reliable partners in the population of player Bs. Interestingly, this effect vanishes in the two information-enhancing treatments. Once any information flow becomes available (through either cheap-talk or observation) trustworthiness becomes driven only by the specific signal.

In the communication treatment, the information appears to rely solely on the reassuring message (“I will choose r ”) which substantially improves the odds of taking action R . Message “I will choose l ” does not significantly change the probability of decision R , as compared to an uninformative message. For the observation treatment, we discretize the reputation observed by player A into 6 intervals, and estimate separately the effect of perfect reputation. As compared to the reference category of a bad reputation (past reliability not exceeding 50%), only very high levels of reputation (over 80% of decisions r in the past) are able to induce a significant increase in reliance. Partners' perfect reputation induces the highest increase in the use of a reliant strategy, corresponding to a 66 percentage point increase in the likelihood of a reliance.

In all games, player As who receive positive signals (through the repetition of game in the baseline treatment, reassuring messages in the communication treatment and the quality of reputation in the observation treatment) are significantly more willing to rely on partners. One virtue of this change in behavior is that it improves efficiency through reinforcement-based learning. The dummy variable $1_{Trusted\ Before}$ indicates whether player A has already relied on his partner in the past. This variable *per se* is a very poor predictor of the likelihood of current reliance. We capture reinforcement through interacting this variable with time (variable *Round*) and a measure of success, *Ratio_Rr*, constructed as the proportion of outcomes (R, r) among all decisions R . In all regressions, estimated coefficients indicate that decision R becomes more likely over time when the historical attempts to rely on partners are more fruitful. This reinforcement-based learning is significant and has comparable marginal effects in all treatments.

Table 9: Average payoffs from the use of dominant strategies

	Reliability rate			Average payoffs					
	$Pr(B_P)$	$Pr(r B)$		$E(Payoff(B) r)$			$E(Payoff(A) R)$		
		B_P	B_{IP}	B_P	B_{IP}	B	B_P	B_{IP}	B
Baseline	55.6%	92.0%	66.7%	4.07€	4.08€	4.07€	9.56€	7.93€	8.85€
Communication	53.0%	92.3%	66.1%	4.42€	4.21€	4.34€	9.78€	8.41€	9.23€
— Following m(r)	62.1%	97.6%	77.9%	4.48€	4.50€	4.49€	9.85€	8.63€	9.39€
Observation	54.8%	93.9%	68.9%	4.61€	3.69€	4.26€	9.82€	8.21€	9.41€

Note. In each column, outcomes are conditioned on interacting with player Bs (B), who are then separated into two subgroups, depending on whether they enter the interaction with a perfect reputation (B_P) or not (B_{IP}). Due to such conditioning, data come from rounds 2-10. For each treatment in rows, the left-hand side reports the reliability rate in the main population player Bs as well as in each of the subgroups. In the right hand-side, the middle column provides the average payoff earned by player Bs who choose the payoff maximizing decision r . The last column provides the average payoff earned by player As who rely on their partner by choosing R .

4.5 Empirical efficiency of reliance

Table 9 summarizes the average payoffs earned by and against player Bs with different reputations when the decisions of the payoff maximizing strategy, r for player B and R for player A, are actually played.²¹ Compared with the baseline treatment, the introduction of either communication or observation does not greatly change the average welfare of trustful player As, while substantially benefiting player Bs. The welfare of reliant player As increases, though, when reliance is conditioned on either of the positive signals – conveyed through a reassuring message or an observed perfect reputation. Conditional on receiving a positive signal, observation is more welfare-improving than communication for both players. Keeping with a message announcing decision r brings on average 4.49€ to reliable player Bs, while observed perfect reputation yields on average 4.61€; for reliant player As, the average payoff increases from 9.39€ to 9.82€, respectively.

From player As' perspective, reliance is empirically dominated in all three experimental conditions. Recall that player A can always guarantee 9.75€ in the game by choosing L , in which case player B experiences a 3€ payoff. Compared with this threshold, reliance never dominates the secure choice: the average earnings are 8.85€ in the baseline, 9.23€ in the communication treatment and 9.41€ in the observation treatments. These average payoffs closely reflect a common lack of reliability in the population of player Bs. In all treatments, player Bs' past behavior is strongly related to their current intentions, in such a way that the emergence of two distinct types of player Bs – already reported in the context of observation treatment – extends to other experimental conditions. This is accounted for in the table by distinguishing between highly reliable B_P players (who always played r in the past and whose likelihood of choosing r is 92%-94%) and less reliable B_{IP} players (who deviated from playing r in the past, and whose probability of choosing

²¹Due to this conditioning on historical behavior, we use observations from rounds 2-10.

r is 66%-69%).²²

Although, in the communication treatment, player As cannot observe player Bs' history and directly distinguish between the two types, a reassuring message conveys indirect (and imperfect) information about partner's type. A strong majority (62%) of subjects sending a reassuring message are highly reliable B_P players who then choose r 97.6% of times and provide an average gain of 9.85€ to a reliant partner (as compared to 77.9% and 8.63€, respectively, observed among B_{IP} players who send this signal). Consequently, the presence of weakly reliable player Bs severely reduces the average gain in welfare from relying on a partner sending a reassuring message, which ultimately makes non-reliance empirically more beneficial.

Observation, in contrast, allows player As to identify both types of player Bs, to reward a good reputation and to punish a bad one. This results in a strong inter-type gap in player Bs' average gains (4.61€ for B_P subjects vs. 3.69€ for B_{IP} subjects). Although a large difference in player As' average payoffs from relying on both types of player Bs (9.82€ for B_P subjects vs. 8.21€ for B_{IP} subjects) justifies such discriminatory behavior, player As do not appear to be conservative enough, since the average gain from action R (9.41 euros) remains lower than the value of the safe option.

5 Conclusion

We implement in the laboratory a 2 x 2 game with two Nash-equilibria, one of which is imperfect and arises if one player is reluctant to rely on the rationality of the current interaction partner. In his discussion of this game, Rosenthal (1981) conjectured that the imperfect equilibrium may arise depending on: the stakes of the game, and what is known about the interaction partner. Accumulated evidence from experiments confirms the first part of the conjecture (Beard and Beil, 1994; Beard, Beil, and Mataga, 2001; Goeree and Holt, 2001). To assess the role of information about the interaction partner, we replicate one of the payoff structures used by Beard and Beil (1994) under three information enhancing conditions. The game is repeated in the baseline so as to allow players to learn about the population of their partners. Two further treatments provide specific information about the current interaction partner through either one-way cheap-talk communication or observation of partners' historical decisions.

Observed behavior supports the claim that information helps to overcome coordination failures. Repetition of the game and the transmission of specific information increase the odds of reaching the most efficient outcome. Because repetition conveys information about the general population of partners, rather than the one actually met, this comes at a price in terms of action mismatches. The provision of subject-specific information appears as a substitute to repetition-

²²The differences in reliability between both types of player Bs, shown in the second column of Table 9, are statistically significant according to our parametric test for equality of proportions, with p -values equal to 0.052, 0.002, 0.002, 0.050 for subsequent rows.

based learning. The main difference between the two information-transmission mechanisms lies in the effect of negative signals: non-reassuring messages improve the chances of coordinating on a Nash equilibrium as compared to an imperfect reputation. This confirms previous findings that information-transmission mechanisms affect both the frequency and the form of coordination failures (Parkhurst, Shogren, and Bastian, 2004). Since the two information-transmission mechanisms have rather different screening properties in terms of the future actions of the sender, a natural direction for future research is to exploit their potential complementarities by combining them (see (Duffy and Feltovich, 2006) and (Çelen, Kariv, and Schotter, 2010) for applications to different games).

Our results thus confirm Rosenthal’s conjecture that the extent of information influences strategic interactions in this game. The experiment also raises interesting questions about individual behavior in the game. An important feature of our design is that it elicits decisions from both players in each round of the game. This allows us to assess the empirical efficiency of reliance induced by the actual likelihood of meeting a reliable partner. We find that the unreliability rate is high enough to justify unreliance in most situations. The behavior of our subjects is even insensitive to any of our experimental treatments. It means neither repetition-based learning, nor forward-looking information through messages nor even backward-looking information from observation of past decisions exert any discipline on the choice of weakly dominated actions. Identifying the dimensions of the environment that give rise to such behavior is next on our research agenda.

References

- ANGRIST, J. D., AND V. LAVY (2009): “The Effect of High School Achievement Awards: Evidence from a Randomized Trial,” *American Economic Review*, 99(4), 1389–1414.
- AUMANN, R. (1990): “Nash Equilibria Are Not Self-Enforcing,” in *Economic Decision-Making: Games, Econometrics, and Optimization: Contributions in Honor of Jacques H. Dreze*, ed. by J. J. Gabszewic, J.-F. Richard, and L. A. Wolsey, pp. 201–206. North-Holland, Amsterdam.
- BEARD, T. R., AND J. BEIL, RICHARD O. (1994): “Do People Rely on the Self-Interested Maximization of Others? An Experimental Test,” *Management Science*, 40(2), 252–262.
- BEARD, T. R., R. O. J. BEIL, AND Y. MATAGA (2001): “Reliant behavior in the United States and Japan,” *Economic Inquiry*, 39(2), 270–279.
- BELL, R. M., AND D. F. MCCAFFREY (2002): “Bias Reduction in Standard Errors for Linear Regression with Multi-Stage Samples,” *Survey Methodology*, 28(2), 169–179.
- BINMORE, K., J. MCCARTHY, G. PONTI, L. SAMUELSON, AND A. SHAKED (2002): “A Backward Induction Experiment,” *Journal of Economic Theory*, 104(1), 48–88.
- BLUME, A., AND A. ORTMANN (2007): “The effects of costless pre-play communication: Experimental evidence from games with Pareto-ranked equilibria,” *Journal of Economic Theory*, 132(1), 274–290.

- BOCHET, O., T. PAGE, AND L. PUTTERMAN (2006): "Communication and punishment in voluntary contribution experiments," *Journal of Economic Behavior & Organization*, 60(1), 11–26.
- BOLTON, G. E., E. KATOK, AND A. OCKENFELS (2004): "How Effective Are Electronic Reputation Mechanisms? An Experimental Investigation," *Management Science*, 50(11), 1587–1602.
- BRACHT, J., AND N. FELTOVICH (2009): "Whatever you say, your reputation precedes you: Observation and cheap talk in the trust game," *Journal of Public Economics*, 93(9-10), 1036–1044.
- CAMERON, A. C., J. B. GELBACH, AND D. L. MILLER (2008): "Bootstrap-Based Improvements for Inference with Clustered Errors," *Review of Economics and Statistics*, 90(3), 414–427.
- ÇELEN, B., S. KARIV, AND A. SCHOTTER (2010): "An Experimental Test of Advice and Social Learning," *Management Science*, 56(10), 1687–1701.
- CHARNESS, G. (2000): "Self-Serving Cheap Talk: A Test of Aumann's Conjecture," *Games and Economic Behavior*, 33(2), 177–194.
- CHARNESS, G., AND M. DUFWENBERG (2006): "Promises and Partnership," *Econometrica*, 74(6), 1579–1601.
- CHARNESS, G. B., AND M. DUFWENBERG (2008): "Broken Promises: An Experiment," *UCSB Working Paper*, 10(08).
- COOPER, R., D. V. DEJONG, R. FORSYTHE, AND T. W. ROSS (1992): "Communication in Coordination Games," *Quarterly Journal of Economics*, 107(2), 739–771.
- CRAWFORD, V. (1998): "A Survey of Experiments on Communication via Cheap Talk," *Journal of Economic Theory*, 78(2), 286–298.
- DUFFY, J., AND N. FELTOVICH (2002): "Do Actions Speak Louder Than Words? An Experimental Comparison of Observation and Cheap Talk," *Games and Economic Behavior*, 39(1), 1–27.
- (2006): "Words, Deeds, and Lies: Strategic Behaviour in Games with Multiple Signals," *Review of Economic Studies*, 73(3), 669–688.
- ELLINGSEN, T., AND R. ÖSTLING (2010): "When Does Communication Improve Coordination?," *American Economic Review*, 100(4), 1695–1724.
- FARRELL, J. (1988): "Communication, Coordination and Nash Equilibrium," *Economics Letters*, 27, 209–214.
- FARRELL, J., AND M. RABIN (1996): "Cheap Talk," *Journal of Economic Perspectives*, 10(3), 103–118.
- FEHR, E., AND K. M. SCHMIDT (1999): "A Theory of Fairness, Competition, and Cooperation," *Quarterly Journal of Economics*, 114(3), 817–868.
- GOEREE, J. K., AND C. A. HOLT (2001): "Ten Little Treasures of Game Theory and Ten Intuitive Contradictions," *American Economic Review*, 91(5), 1402–1422.

- GREINER, B. (2004): "An Online Recruitment System for Economic Experiments.," in *Forschung und wissenschaftliches Rechnen 2003. GWDG Bericht 63*, ed. by K. Kremer, and V. Macho, pp. 79–93. Ges. für Wiss. Datenverarbeitung, Göttingen.
- JACQUEMET, N., AND A. ZYLBERSZTEJN (2011): "What drives failure to maximize payoffs in the lab? A test of the inequality aversion hypothesis?," *CES Working Paper*, 2011(36).
- PARKHURST, G. M., J. F. SHOGREN, AND C. BASTIAN (2004): "Repetition, Communication, and Coordination Failure," *Experimental Economics*, 7(2), 141–152.
- ROSENTHAL, R. W. (1981): "Games of perfect information, predatory pricing and the chain-store paradox," *Journal of Economic Theory*, 25(1), 92–100.
- SELTEN, R. (1975): "Reexamination of the perfectness concept for equilibrium points in extensive games," *International Journal of Game Theory*, 4(1), 25–55.
- WILLIAMS, R. (2000): "A Note on Robust Variance Estimation for Cluster-Correlated Data," *Biometrics*, 56(2), 645–646.
- WOOLDRIDGE, J. M. (2003): "Cluster-Sample Methods in Applied Econometrics," *American Economic Review*, 93(2), 133–138.
- ZEILIGER, R. (2000): "A presentation of Regate, Internet based Software for Experimental Economics," <http://www.gate.cnrs.fr/~zeiliger/regate/RegateIntro.ppt>, *GATE*.