Journal of Economic Theory **90**, 161–172 (2000)
doi:10.1006/jeth.1999.2573, available online at http://www.idealibrary.com on IDEAL®

An Experimental Investigation of Optimal Learning in Coordination Games[1]

Andreas Blume

Department of Economics, University of Iowa, Iowa City, Iowa 52242
andreas-blume@uiowa.edu

and

Uri Gneezy

Faculty of Industrial Engineering and Management, Technion, Haifa 32000, Israel
gneezy@econ.haifa.ac.il

Received May 27, 1999; revised July 13, 1999

This paper presents an experimental investigation of optimal learning in repeated coordination games. We find evidence for such learning when we limit both the cognitive demands on players and the information available to them. We also find that uniqueness of the optimal strategy is no guarantee that it will be used. Optimal learning can be impeded by both irrelevant information and the complexity of the coordination task. *Journal of Economic Literature* Classification Numbers: C72, C92. © 2000 Academic Press

Key Words: coordination; learning; language; common knowledge.

INTRODUCTION

This paper presents an experimental investigation of optimal learning in repeated coordination games that lack an *a priori* common-knowledge description. Coordination problems that lack a complete common-knowledge description are characteristic of situations without access to natural language such as tacit collusion, and of situations in which natural

[1] This paper was conceived while both authors were at the CentER for Economic Research at Tilburg University in The Netherlands. We thank CentER for its hospitality and financial support through the CentERpanel. A. Blume acknowledges support through a grant from the National Science Foundation (SBR-9808947). We are grateful to Ido Erev and Drew Fudenberg for comments.

0022-0531/00 $35.00
Copyright © 2000 by Academic Press
All rights of reproduction in any form reserved.

language is not entirely adequate, such as an organization facing novel challenges or parties trying to write a contract in a complex environment.[2]

We consider a simple, stylized setting with two players. Initially the players do not have a common language to distinguish either their roles or their actions. When actions are completely symmetric, first-round coordination is entirely a chance event; there are no *focal points* that would permit *a priori* coordination. The history of play, however, may desymmetrize the game, and players may follow rules that make use of such asymmetries. If they do use such rules, they may learn to commonly distinguish some of the actions. If their rule is optimal, i.e., they make the best possible use of the arising asymmetries, then we say that they learn optimally.

In general, there may be multiple optimal learning rules, which lead to a further, higher-order, coordination problem. Therefore we focus on the case where there is a unique optimal learning rule. Thus, if players are sufficiently sophisticated and attribute sufficient sophistication to others, there is at least the possibility that the unique optimal learning rule becomes focal.

Optimal learning in coordination games without a common-knowledge description was first investigated by Crawford and Haller [9].[3] Blume [3] considers optimal learning with partial languages that do not make distinctions among individual objects but have structure that permits fast learning.[4] To the best of our knowledge, a rigorous experimental investigation is still lacking. The present paper is the first step in such an investigation.

The games we consider are extremely simple two-player pure coordination games that are repeated twice. The presentation of the games ensures that in the first round of play players lack a common description of the game and thus are unable to guarantee first-round coordination. However, the games are designed in such a way that corresponding to any first-round outcome, there is a unique optimal way to play in the second round, making use of the common distinctions among actions that are introduced by first-round play. Thus, we are studying the role of *endogenously generated*

[2] While we focus on pure coordination games, the case of contract negotiations shows that coordination and shared language issues arise in a wide variety of settings. Crawford [10] mentions "negotiations between agents who share a common language but have different cultures" as being more appropriately described by a model with absence of a common language.

[3] These ideas have been pursued further by Calvert [6] and Kramarz [15] and have been critically evaluated by Goyal and Janssen [13]. Blume *et al.* [4] have studied the evolution of meaning in a setting without an *a priori* common language.

[4] The structural properties of natural language are investigated by Rubinstein [22], who argues that binary relations in language tend to have properties that are dictated by their usefulness, e.g., in facilitating the naming of nameless objects.

focal points.[5] In the environment that we consider, optimal learning implies very different behavior than adaptive learning rules, such as fictitious play in Robinson [19] or stimulus–response behavior in Roth and Erev [20].[6] Thus, experiments of this form potentially permit a sharp rejection of popular learning theories for games.

We find evidence for optimal learning when we limit both the cognitive demands on players and the information available to them. We also find that uniqueness of the optimal strategy is no guarantee for it to be used. Optimal learning can be impeded by the complexity of the coordination task. In addition, more information may obstruct optimal learning and lead to lower expected payoffs for the players.

2. THE EXPERIMENT

This section first discusses a class of simple, two-player, two-period coordination games and derives the unique optimal learning rules for games in this class. Then we describe our experimental design based on these games.

2.1. *Games and Predictions*

We consider coordination games that are played between two players for two periods. Within a period, players independently and simultaneously choose one of a finite odd number of locations that are arranged in a circular order. There is a positive payoff if both players choose the same location in a given period, otherwise their payoff is zero in that period. All locations are *a priori* identical.

Our theoretical predictions are based on work by Crawford and Haller [9] and on a recent extension by Blume [3]. A strategy is called *attainable* if it respects the symmetries in the game; i.e., every two pure strategies that differ only in terms of features of the game that are not commonly distinguished by the players enter an attainable strategy with the same probability. Thus, after every history, an attainable strategy is consistent with the remaining restrictions on a common language. For example, actions that are not commonly distinguished must be used with the same probability, and players whose positions in the game are not distinguished must use identical strategies. An attainable strategy that is ex ante efficient among attainable strategies will be called an *optimal attainable strategy* (OAS). We can think of the focus on attainable strategies as an expression of players'

[5] Focal points, whether endogenously generated or not, are important not only in common interest games, but also in divergent interest games, as was originally emphasized by Schelling [23]. Roth and Murnighan [21] find experimental evidence for the role of game descriptions in enabling social conventions that determine bargaining outcomes.

[6] For a review of this literature, see Fudenberg and Levine [12].

strategic uncertainty. OAS selects among equilibria that respect this uncertainty according to the Pareto optimality criterion.[7] OASs are particularly attractive when they are unique.

In our games, the OAS criterion makes the same prediction independent of the number of locations: The first-round choice is random by design. If players happen to coordinate in the first round, it is uniquely optimal for them to choose the same location in the second round. If they fail to coordinate, there is a location that is distinguished from all others. To see this, note that the two first-round choices-are separated from each other by an even number of locations on one side and by an odd number on the other side (the "odd side"). The midpoint of the odd side is thus uniquely distinguished from all others. This is the location that both players ought to use if they follow an optimal strategy.

2.2. *Experimental Design*

We had four experimental sessions, which differed in the number of locations on the circle and in the information feedback after the first stage. We first describe the procedure for the treatment with three sectors in which "full information" was given and then describe the differences between it and the other treatments.

Students were invited into two rooms that were separated by a curtain and given "registration numbers." There they received instructions (see Appendix A), which were read aloud to both rooms together in order to make them common knowledge. Participants were told that the experiment consisted of two parts and that the instructions for the second part would follow after the first part was over. They were told that they would interact with the same participant in the two parts of the experiment and that only the investigator would know the identity of the person with whom they were matched, as well as the order of moves (that is, who of the two would choose first).

A plastic plate that had been divided by lines on both sides into three equal sectors was then shown to all participants. In the first part of the experiment each participant in one room was asked to choose one of the sectors on one side of the plate. Using a sticker, we marked this sector with the letter "A." The matched participant in the other room was then asked (without knowing the choice of his counterpart) to choose one sector, but on the other side of the plate. Participants were told that if the sectors chosen matched, then each would receive f 10, otherwise f 0.[8]

After participants in both rooms had made their choices, we marked each side of the plate using a sticker with the letter "B" to indicate the

[7] Harsanyi and Selten [14] use Pareto optimality as a selection criterion.
[8] At the time of the experiment $1 = f 1.9.

TABLE I

Two-by-Two Experimental Design

	Full information	Partial information
Three sectors	3FI	3PI
Nine sectors	9FI	9PI

choice of the respective counterpart. So now each side of the plate indicated for the participant what his or her choice in round one was as well as what the choice of his or her counterpart was.

Participants then received the instructions for part 2 (see Appendix B). The rules in the second part were the same as in the first, apart from the information about the choices made in stage 1: participants were asked to choose one sector, and this sector was marked by a sticker with the letter "C." At the end of the second stage, all payoffs were calculated and participants were paid.

The two-by-two experimental design is described in Table I.

The difference in procedures for the 3FI treatment, which was described above, and the 3PI treatment concerned the stickers: Those that marked the choices in the first stage of the PI treatment were blank. As a result, participants in the second stage knew what choices had been made in the first stage, but not who had made each choice. The 9FI treatment and the 9PI treatment were similar to the 3FI and 3PI treatment, respectively, apart from the number of sectors on the plate, which was nine instead of three.

The experiment was conducted at Tilburg University and Amsterdam University (CREED); in all, 86 undergraduate students from all fields of study (24, 30, 16, and 16 students in treatments 3FI, 3PI, 9FI, and 9PI, respectively) participated.

3. RESULTS

The results of our experiment are reported in Tables II and III. The discussion will focus on second-round play and distinguish between continuation play conditional on whether there was a first-round match or not.

Table II reports on second-round play following a first-round match. There appears to be a tendency toward OAS play conditional on a first-round match. Pooling the data for first-round matches, we find that 13 of 18 participants (72%) chose to repeat their first-round choice. Using the binomial test, we can reject the null hypothesis that the probability of a repetition of the first-round choice is one-third ($p = 0.00765$). On the other

TABLE II

Second Stage Play If Matched in First Stage

	Fraction choosing the same sector	Fraction not choosing the same sector
3FI (8 Observations)	0.63	0.37
3PI (8 Observations)	0.75	0.25
3FI (2 Observations)	1	0
3PI (0 Observations)	—	—

hand, one notes that play far from matches the OAS Prediction perfectly. Moreover, the observations could be rationalized as being the result of stimulus–response learning (which here would amount to increasing the probability of the successful first-period action in the second period) or of stochastic fictitious play (which here would amount to increasing the probability of the best reply to the counterpart's first-period action).

More interesting than the case of first-round matches is that in which initial choices happen not to be coordinated. In both the three- and the nine-sector full-information treatments, we can clearly reject the OAS hypothesis. Rather than choosing the unique distinct sector, participants appear to prefer their own first-round choice. Combining the data from 3FI and 9FI when players were not matched in the first round, we see that 27 of the 30 participants chose one of the two first-round choices. We can reject the null hypothesis that both first-round choices were equally likely ($p = 0.009579$). Note that this behavior is also inconsistent with stochastic fictitious play, which would tend to favor the counterpart's first-round choice, and not easily explained as stimulus–response learning, as with zero first-round payoffs there is no reinforcement of the first-round choice.

For the nine-sector partial information treatment we can also reject OAS, as the unique distinct sector is never chosen. The overwhelming

TABLE III

Second Stage Play If Not Matched in First Stage

	Fraction choosing one of the 1st-stage choices	Fraction choosing own 1st-stage choice	Fraction choosing other's 1st-stage choice	Fraction choosing the "optimal sector"	Other sectors
3FI (16 Obs.)	0.875	0.625	0.25	0.125	—
3PI (22 Obs.)	0.363	—	—	0.637	—
9FI (14 Obs.)	0.928	0.714	0.214	0	0.072
9PI (16 Obs.)	0.875	—	—	0	0.125

number of participants chose one of the first-round choices. This behavior is consistent with a stochastic version of fictitious play, but is not easily explained by stimulus–response learning. Note that this behavior is close to the unique OAS if we ignore the circular structure on locations. The participants would then (nearly) be achieving a "second best" optimality subject to a cognitive constraint.

Closest to the OAS prediction are the observations for the three-sector, partial information treatment. In particular, 14 of the 22 participants (64%) chose the "optimal sector:" we can reject the null hypothesis that all sectors were equally likely ($p = 0.003138$) in favor of the alternative that the unique distinct sector was more likely than the others. Perhaps, since participants are unable to locate their own first-round choices (as in 3PI), they spend more effort in analyzing the game. It is also possible that the symmetry of the two first-round choices is more apparent, thereby making the unique distinct action more conspicuous.[9] Note also that with full information, the set of "cognitive strategies" is larger than that with only partial information; e.g., the rule to adopt the choice of the other player is simply not feasible with only partial information. This sharp reduction in cognitive strategies in the 3PI treatment may provide one explanation for why we find more coordination. In addition, the unique-distinct sector may be easier to identify in 3PI than in 9PI. With nine sectors, there is a large set of cognitive strategies even if there is only partial information. This might account for the difference in the results between these two treatments. Thus, at least in a cognitively simple environment, OAS appears to help in explaining the data. Note that this is so despite the fact that stimulus–response learning provides no convincing explanation and stochastic fictitious play favors play of one of the first-round choices.

Finally, there is the paradox that participants make better choices with less information. The strategy of repeating one's first-round choice in the absence of a first-round match leads to the worst possible outcome if

[9] Following Bacharach's [1] study of *variable universe* games, we can think of different conceptualizations of a given game. Here such conceptualizations may depend on play in the first round and on whether players can recall their own first-round choices. Suppose there was no coordination in the first round. If a player can recall his or her own first-round choice, one possible conceptualization of the strategy space is that of *own choice*, *other's choice*, or *unchosen action*. Unless players realize that there is a symmetry between *own choice* and *other's choice* once the role symmetry of the players is taken into account, there are three equally prominent actions. Without recollection of own choice, on the other hand, the only possible conceptualization is that of *chosen action* or *unchosen action*. Since there are two *chosen actions*, the coordination probability in the second round is higher (in fact, it equals one) if both players pick the unchosen action. Using Pareto optimality as a selection criterion, players with the latter conceptualization would choose the "unchosen action" in the second round, provided they were sufficiently confident that their counterparts think about the game in the same way.

everyone uses it, worse even than randomizing uniformly over all possible choices in the second round.

4. RELATED LITERATURE

Our paper is concerned with learning and the optimal use of endogenously generated focal points. There is an extensive literature on both focal points and learning. This section briefly examines this literature in connection with the present paper.

Focal points were first discussed by Schelling [23], who gives an intuitive discussion and also reports the results of some informal experiments. According to Schelling the two prime characteristics of focal points are *conspicuousness* and *uniqueness*. He suggests that finding them may depend more on imagination than on logic. This suggests that Schelling is skeptical of a formal game theoretical investigation of focal points. In particular, he expresses his reservations about the "empirical relevance of mathematical foci." One should not ascribe to the players in a game the mathematical sophistication of the analyst. For a sophisticated mathematical solution to be focal for a player, that player not only needs to be a mathematician, but also must view his on her playing partners as such.

Schelling's distinction between mathematical and psychological foci is potentially relevant for interpreting the results of our experiment. If players do not coordinate in the first round, their two first-round choices become conspicuously distinct from all locations not chosen by either player. On the other hand, the fact that these locations can be used to uniquely define a further location on which to coordinate requires moderately sophisticated inference. That location may thus be somewhat less conspicuous than the two others. Even if it is just as conspicuous to a given player, that player may not be confident that the location is equally conspicuous to the playing partner.[10] Sugden [24] develops a formal theory of focal points by explicitly introducing the labeling of strategies into the analysis. He aims at a "general theory of how labels can influence decisions in games" (Sugden [24, p. 534]). In a pure coordination game, his theory prescribes that players use decision rules, maps from their private descriptions to a labeled choice, that induce a distribution over choices that maximizes the coordination probability. He argues that in environments with a common culture this prescription often leads to a unique optimal decision rule because of the skewed distribution of the different items mentioned. In our

[10] Bacharach [1] constructs a formal model along these lines. Bacharach and Bernasconi [2] test the predictions from this model experimentally.

setting, this rule defines a unique optimal choice in the second round *if* the unique distinct location of the second round is recognized by the participants as such.

Mehta *et al.* [16] examine the concept of a focal point experimentally in pure coordination games. Their objectives are to replicate Schelling's informal experiments and to discriminate among alternative explanations for coordination success being more frequent than accounted for by pure chance. They distinguish among primary, secondary, and Schelling salience. Primary salience of an action means that (for whatever reason) it is likely to come to mind. An action has secondary salience if it is the optimal reply to one deemed to have primary salience for the playing partner(s). An action has Schelling salience if there is a selection rule that, if used by both players, unambiguously singles out that action as guaranteeing coordination success. They confirm the observation that coordination success is often more frequent than would be suggested by pure chance, and they reject the explanation that this is due to a combination of primary salience and shared cultural experience. They suggest that both secondary and Schelling salience play a role.

There has been a recent resurgence of interest in adaptive learning theories, both theoretically (see Fudenberg and Levine [12] for a review of the literature) and empirically (e.g., Mookherjee and Sopher [17], Roth and Erev [20], Camerer and Ho [7], Cheung and Friedman [8], and Blume *et al.* [5]). These theories seem to be quite successful in explaining certain salient characteristics of experimental data in some domains. For example, Roth and Erev find simulated stimulus–response learning to be similar to observed behavior in games like the ultimatum game and find a parallelism between observed and simulated behavior across different games. Blume *et al.* [5] find that both stimulus–response and belief-based learning models fit their learning data for sender–receiver games well.

Erev and Roth [11] also suggest that behaviors that are not explainable in terms of learning stage-game strategies (like alternation in repeated play of "Chicken" (see Rapoport *et al.* [18]) could be accounted for by adaptive learning that updates repeated game strategies instead. This avenue is not available for reconciling our results with adaptive learning theories. In our setting, the only learning occurs between the first and second periods and thus must be in terms of stage-game strategies.[11] If agents are not coordinated in the first round, then stimulus–response learning does not favor any of the actions in the second round, since there is no positive payoff

[11] Note that in our setting, what is commonly known about players' descriptions of the game changes between the first and the second round. If, in contrast, players do not have access to a common history, as in a turnpike design in which each player meets a different player in every round, then it is likely that stimulus–response learning will be more successful at explaining the data.

reinforcement.[12] Belief-based learning would tend to favor the action taken by the other player. Neither behavior is consistent with our data for the 3PI treatment, in which a majority chose the "unchosen action" of the first round.

APPENDIX A

Instructions for Part 1

Welcome to this experiment in decision making. In the experiment, we will ask you to make some very simple decisions and you may earn some money that will be paid to you, in cash, at the end.

The experiment consists of two parts. At the beginning of the first part we will randomly match you with another participant from a group of students that stays in another room. You will interact with the same participant in the two parts of the experiment. Only we will know the identity of the person with whom you are matched.

Procedure

A plate, like the one shown to you now, is also shown to the participants in the other room. As you can see, the plate is divided into three equal sectors on both sides. In the first part of the experiment we ask you to choose one of the sectors on one side of the plate. We will mark this sector using a sticker with the letter "A." The participant you are matched with is also asked to choose one sector, but on the other side of the plate (he can either do this before or after you did). None of you will see the choice of the other person before choosing.

Payment: If the sectors chosen by both of you match, then each of you will receive f 10. Otherwise each of you will receive f 0. Instructions for part two of the experiment will follow.

Do you have any questions?

APPENDIX B

Instructions for Part 2

This is the last part of the experiment. We now show you your plate again. A sticker with the letter "B" marks the choice of the other person.

[12] If we also had positive payoffs in the event of a coordination failure, stimulus–response learning would favor the action taken in the first round, which is consistent with some of our data.

The rules in this part are the same as in the previous one: we ask you to choose one sector. We will mark this sector using a sticker with the letter "C." The person you are matched with (the same person from part 1) receives the same information and task as you do.

The same rules for payment apply as in part 1. Do you have any questions? Thank you for your cooperation!

REFERENCES

1. M. Bacharach, Variable universe games, *in* "Frontiers of Game Theory" (K. Binmore, A. Kirman, and P. Tani, Eds.), pp. 255–275, MIT Press, Cambridge, MA, 1993.
2. M. Bacharach and M. Bernasconi, The variable frame theory of focal points: An experimental study, *Games Econ. Behav.* **19** (1997), 1–45.
3. A. Blume, "Coordination and Learning with a Partial Language," working paper, University of Iowa, Iowa City, 1999.
4. A. Blume, D. DeJong, Y.-G. Kim, and G. Sprinkle, Experimental evidence on the evolution of meaning of messages in sender–receiver games, *Amer. Econ. Rev.* **88** (1998), 1323–1340.
5. A. Blume, D. V. DeJong, G. R. Neumann, and N. E. Savin, "Learning in Sender–Receiver Games," working paper, University of Iowa, Iowa City, 1999.
6. R. L. Calvert, "Elements of a Theory of Society among Rational Actors," Department of Political Science Working Paper, University of Rochester, Rochester, 1991.
7. C. Camerer and T.-H. Ho, Experience-weighted attraction learning in games: A unifying approach, *Econometrica*, in press (1999).
8. Y. W. Cheung and D. Friedman, Individual learning in normal form games, *Games Econ. Behav.* **19** (1997), 46–79.
9. V. Crawford and H. Haller, Learning how to cooperate: Optimal play in repeated coordination games, *Econometrica* **58** (1990), 581–596.
10. V. Crawford, A survey of experiments of communication via cheap talk, *J. Econ. Theory* **78** (1998), 286–298.
11. I. Erev and A. Roth, Predicting how people play games: Reinforcement learning in experimental games with unique mixed strategy equilibria, *Amer. Econ. Rev.* **88** (1998), 848–881.
12. D. Fudenberg and D. K. Levine, "The Theory of Learning in Games," MIT Press, Cambridge, MA, 1998.
13. S. Goyal and M. Janssen, Can we rationally learn to coordinate?, *Theory and Decision* **40** (1996), 29–49.
14. J. C. Harsanyi and R. Selten, "A General Theory of Equilibrium Selection in Games," MIT Press, Cambridge, MA, 1988.
15. F. Kramarz, Dynamical focal points in N-person coordination games, *Theory and Decision* **40** (1996), 277–313.
16. J. Mehta, C. Starmer, and R. Sugden, The nature of salience: An experimental investigation of pure coordination games, *Amer. Econ. Rev.* **84** (1994), 658–673.
17. D. Mookherjee and B. Sopher, Learning behavior in an experimental matching pennies game, *Games Econ. Behav.* **7** (1994), 62–91.
18. A. Rapoport, M. J. Guyer, and D. G. Gordon, "The 2×2 Game," University of Michigan Press, Ann Arbor, 1976.

19. J. Robinson, An iterative method of solving a game, *Annals of Mathematics* **54** (1951), 296–301.
20. A. E. Roth and I. Erev, Learning in extensive form games: Experimental data and simple dynamic models in the intermediate term, *Games Econ. Behav.* **8** (1995), 164–212.
21. A. E. Roth and K. Murnighan, The role of information in bargaining: An experimental study, *Econometrica* **50** (1982), 1123–1142.
22. A. Rubinstein, Why are certain properties of binary relations relatively more common in natural language, *Econometrica* **64** (1996), 343–355.
23. T. Schelling, "The Strategy of Conflict," Harvard Univ. Press, Cambridge, MA, 1960.
24. R. Sugden, A theory of focal points, *Econ. J.* **105** (1995), 533–550.