An Experiment on the Value of Structural Information in a 2x2 Repeated Game*

By Atanasios Mitropoulos°

Faculty of Economics and Management
Otto-von-Guericke-University Magdeburg

Version 1.0

November 2001

Abstract:
In experimental studies pairs that repeatedly play the simple coordination game *mutual fate control* may regularly fail to coordinate when they are given little information, i.e. when subjects are uninformed about the payoff matrix and feedback is limited to their own payoff. Our experimental study shows that the provision of a small amount of structural information prior to playing the game changes subject behaviour and significantly improves performance, even though standard adaptive learning rules do not take such information into account and optimal adaptive rules do not differ much between the two treatments. Our study calls for a more intense investigation into the cognitive processing of information.

Keywords: repeated games, experiments, information, coordination

JEL classification: C73, C92, D83

---

* We are thankful to Erwin Amann who proposed to us this experiment. Support by the Deutsche Forschungsgemeinschaft is gratefully acknowledged.
° Address: Lehrstuhl VWL III, PO Box 4120, 39016 Magdeburg; E-mail: atanasios.mitropoulos@ww.uni-magdeburg.de

1. Introduction

The impact of information on decisions has extensively been dealt with in theoretical literature (e.g. Green and Stokey 1980, Sulganik and Zilcha 1996) as well as experiments (e.g. Kagel et al. 1996). In repeated games focus has been laid on the amount of feedback (e.g. Duffy and Feltovich 1999, Feltovich 2000, Blume et al. forthcoming). However, few efforts have been made to assess the impact of information that is available *prior* to playing a repeated game. Such an investigation promises to give valuable insights into the way individuals process information.

Our investigation is based on the game of *mutual fate control*[1]. Table 1 shows the normal-form representation of this game. We chose this game for the following reason. There is only one cell that is salient to both players (*B*,*B*). This cell maximises individual as well as joint payoffs. One may expect that, if subjects know that they are playing this game, they instantly coordinate on this cell. However, preliminary studies have shown that subjects encounter considerable difficulties in coordinating when feedback is limited to their own payoff and no prior information about the strategic situation is provided[2].

		player 2: *A*	player 2: *B*
player1	*A*	0, 0	1, 0
player1	*B*	0, 1	1, 1

Table 1: Mutual Fate Control

A natural question arising from this observation is: How much information do subjects need in order to be able to coordinate quickly? Is there an intermediate level of information that constitutes a threshold after which coordination is easily established by repeatedly playing partners? We investigate this question by adding a new treatment to the data from Mitropoulos (2001a) in which we provide subjects with information prior to playing the game. The information only concerns the structure and still does not reveal any part of the payoff matrix.

In the past, adaptive learning theories have laid their focus on the usage of feedback information. As a result they do not attempt to model the impact of prior information on adaptive

[1] For a thorough introduction into this game and an overview of related literature, see Mitropoulos (2001a).
[2] Note that the Nash-equilibrium does not discriminate between any mixed-strategy profile. Furthermore, different adaptive learning rules may predict convergence to (*B*,*B*) or they may not. See Mitropoulos (2001a).

play. Therefore, they treat our two treatments as entirely different games and do not make comparative static predictions.

An investigation into optimal learning rules reveals that, while in the environment with prior structural information win-stay lose-change performs best, the variant win-stay lose-randomise is slightly superior in the little-information environment. But qualitatively, predictions do not differ much.

2. Design

The experiments were conducted in seven sessions at the MaxLab[3] (four for the little-information treatment – henceforth treatment 1, three for the treatment with prior structural information – henceforth treatment 2). In each session 8 to 10 subjects, most of them first-year business and economics students, were randomly allocated to computer terminals, which were separated by mobile cardboard devices. A custom made computer program, written in Java, helped to make decisions conveniently and to view all information gathered during play. There were 17 pairs (34 subjects) participating in treatment 1 and 14 pairs (28 subjects) in treatment 2.

In all sessions the instructions were read aloud by the same experimenter. The instruction sheet informed subjects that they were about to play 100 repetitions of a 2x2 stage game with a randomly determined anonymous partner. Even though the game matrix was not revealed to them, they were informed of some characteristics concerning the game environment. In more detail, subjects knew they were playing in fixed pairs and that the stage game would remain the same throughout all repetitions. They were also explicitly told that in each round the two actions of the players completely determined their payoffs and that no randomising device was being used. Subjects received feedback only on their own payoff from the previous round.

Subjects also knew that each received point was going to be converted to 0.30 Deutsche Mark and that they were going to be paid out in cash at the end of the experiment. Fractions of Deutsche Marks were rounded to the next higher integer. So minimum and maximum payoffs that could be earned were 0 Deutsche Mark and 30 Deutsche Mark, respectively. On average, subjects earned 18.25 Deutsche Mark in the little-information treatment and 24.76 Deutsche Mark in the treatment with structural information. Sessions lasted for 45 to 60 minutes. However, large differences between individual total payoffs indicate that considerable payoff in-

[3] The Magdeburg Experimental Laboratory, Germany.

centives were given. The maximum total payoff was 30 Deutsche Mark, while the minimum total payoff was only 2 Deutsche Mark.

In treatment 1 the information about the payoff structure of the game was minimal. Subjects were left ignorant about the entries of the payoff matrix. However, it was common knowledge that payoffs for any strategy combination of a pair could take on either of the two values, 0 or 1[4]. This was done for mainly two reasons. First, subjects would not need to implement special exploratory strategies in order to assess the range of possible payoffs. Instead, it was clear from the outset that the goal was to receive a payoff of 1 as often as possible. Second, as is explained in Mitropoulos (2001b) this restriction reduces the game environment to a situation that is very close to a *global game* (Kajii and Morris 1997) involving three coordination games and the matching pennies game. It can be shown that a simple dynamics (stay with the action after receiving a payoff of one and randomise with equal probability between the two actions otherwise) – if applied by both players – would lead subjects into quick coordination in the coordination games and to high expected payoffs overall.

Treatment 2 differed only in one respect. Subjects were now told that the payoff scheme was such that each player determined the payoff to the other player by choosing one or the other action. Subjects were only left ignorant about which action caused the other player to receive the positive payoff and which action caused the other player to receive zero payoff. In the instructions, this crucial difference between the designs reads as follows:

> "In each round two persons are playing with each other. Each player may choose among two alternative actions, A or B. The actions of the two players determine the payoffs to both players. The following payoff scheme applies: Each person has one action available which assigns one point to the other player, and one action which assigns zero points to the other player." (the author's translation from German).

Note that this information is equivalent to saying that the underlying payoff matrix may be any of the four payoff matrices that result from interchanging the labels of the actions. Compared to treatment 1, we, thus, reduce the set of possible payoff matrices to the four that correspond to the mutual fate control game. In this case, the rule win-stay lose-change – if applied by both players – generates a slightly higher expected payoff than win-stay lose-randomise. We did not provide subjects with probabilities on the occurrence of the matrices. However, the design makes fairly clear that, in order to avoid possible biases due to one label

[4] Instructions are given in the appendix.

being perceived as more salient than another[5], we are forced to implement each payoff matrix with roughly the same frequency.

At the end of the experiment, i.e. after the 100 repetitions had been played but before subjects were paid out, a questionnaire was distributed to all subjects. The experimenter made explicit that the subjects were free to answer the questions or to leave the whole sheet blank, and that, in any case, their payoff would not be affected. Nevertheless, most subjects gave enough information to allow us a rough classification of their strategies. In order to avoid artificial play by players who, in anticipation of the questionnaire, restrict behaviour to only those strategies which they can verbalise, subjects were left uninformed about the questionnaire until it was handed out.

3. Data Analysis

Our goal is to analyse the differences in the ability of subjects to coordinate. However, as yet, there is no standard of how to distinguish coordinating from non-coordinating pairs. Each game environment calls for different definitions. For our purpose we make use of two definitions[6]:

> *Definition 1:*
>
> A pair coordinated on the efficient cell, if (i) within the last 7 rounds at least 6 rounds, or (ii) within the last 10 rounds at least 8 rounds, or (iii) within the last 20 rounds at least 16 rounds resulted in the efficient cell.
>
> *Definition 2:*
>
> The starting round of coordination is determined by the first round of the first sequence of successful coordinations that's length is neither eventually exceeded by a sequence of non-coordination nor is exceeded by 3 times the length of the directly following sequence of non-coordinations[7].

The number of pairs that successfully coordinate on the efficient cell draws a fairly clear picture. Table 2 depicts the corresponding numbers from treatment 1 and treatment 2 on the

[5] Schelling (1960, pp. 53ff) already reports on the natural salience of certain labels in certain choice tasks.
[6] More elaborate discussions of these two definitions are provided by Mitropoulos (2001b).
[7] The two definitions can be viewed as addressing complementary aspects of coordination. While definition 1 focuses on the occurrence rate at the end of the session, definition 2 tries to capture sequences of coordination and non-coordination. Note that, whether or not the starting round of coordination is defined provides us with an alternative definition for coordination. Our results, however, are independent of which definition we use.

structure of the environment. A Fisher-exact test rejects the hypothesis of equal distributions between treatments (to the 1%-level, two-tailed)[8].

	Successes	Failures	Total
Treatment 1	7	10	17
Treatment 2	13	1	14

Table 2: Frequencies of successful and unsuccessful coordinations over treatments.

Restricting our attention to the pairs that coordinate, we find that the starting round of coordination seems to differ between the treatments. Table 3 depicts all numbers on the starting round of coordination[9]. We find the means to differ considerably between the treatment 1 and treatment 2, the mean for the pairs of the former treatment (14.0) lying lower than the mean for the pairs of the latter (26.3). However, a Mann-Whitney-U test just fails to be significant ($p = 0.104$, two-tailed).

	Treatment 1						Treatment 2												
Starting round of coordination	2	2	2	4	9	65	1	3	7	9	12	13	13	14	31	44	50	52	93
Mean	14.0						26.3												

Table 3: Starting round of convergence for the coordinating pairs over treatments

We then tried to find differences in behaviour between those subjects that coordinate in the two information treatments. We calculate the win-stay and the lose-change rates for each individual and compare the rates between the two treatments. We suspect that subjects that are given more information behave more in accordance with the win-stay lose-change dynamics, since this scheme would guide players quickly to the efficient cell. Table 4 reports the individual mean win-stay and lose-change rates of subjects until coordination sets in[10]. Each column shows rates for all subjects that coordinated in the corresponding treatment. Numbers have been ordered from low to high. By inspection, one already notices that the variances are quite high. Especially, in treatment 1, pairs coordinated rather soon (see table 3), and thus, individual averages were computed from few within-subject observations. Mann-Whitney-U tests on the equality of medians between the two treatments are significant ($p < 0.01$) for the

[8] If we use the alternative definition based on the starting round of coordination only six pairs from treatment 1 would be classified as coordinating. Rejection of the null hypothesis of equal frequencies of coordination across treatments would then become even stronger.
[9] Note that, according to definition 2, one pair of treatment 1 which we classified as coordinating, cannot be assigned a definitive starting round of coordination.
[10] Note that two subjects never lost before their coordination phase started.

win-stay rates and insignificant (p = 0.11) for the lose-change rates[11]. However, for the reasons mentioned, we deem these results not to be robust.

	Win-stay		Lose-change	
	Treatment 1	Treatment 2	Treatment 1	Treatment 2
	0.00	0.00	0.25	0.25
	0.00	0.00	0.33	0.25
	0.00	0.00	0.33	0.33
	0.00	0.00	0.50	0.33
	0.00	0.14	0.56	0.38
	0.00	0.25	0.69	0.39
	0.00	0.29	1.00	0.40
	0.00	0.33	1.00	0.43
	0.00	0.38	1.00	0.44
	0.24	0.42	1.00	0.50
	0.61	0.50	1.00	0.50
	0.86	0.50	1.00	0.50
		0.60		0.52
		0.67		0.54
		0.67		0.55
		0.71		0.56
		0.75		0.57
		0.75		0.57
		0.78		0.60
		0.80		0.66
		0.80		0.71
		0.82		0.83
		0.83		
		0.88		
Mean	0.14	0.49	0.72	0.49
St. dev.	0.29	0.31	0.31	0.14

Table 4: Individual win-stay and lose-change rates until coordination sets in

We also tried to estimate learning theories for the data on coordinating pairs. By this means, we hoped to find indicators for structural differences between the two differently informed subject groups. However, probably due to the small data sample, the estimations were uninformative.

4. Questionnaires

Much more informative were the questionnaires that were handed out after the experiments were ended. These questionnaires asked for information on strategies and possible motives for the way the subjects made their decisions. Three questions were asked. First, subjects should briefly describe verbally according to which process (German: “Verfahren”) they had made their decisions in the first rounds. The second question asked the same for the last rounds.

[11] Corresponding Mann-Whitney-U tests using data on all subjects, i.e. including subjects in non-coordinating pairs, show no significant differences, neither in the win-stay rates, nor in the lose-change rates. This result is not surprising, since discouraging results during the first rounds may quickly have led subjects to abort playing simple round-by-round dynamics.

And the third question asked whether the subjects had changed their strategy in between, and if they had, why they had changed their strategy.

We classified responses to the questionnaires according to the seven most prominent patterns of play. Table 5 shows the aggregate numbers of players who, at some point in time, appealed to the corresponding pattern. Since some subjects changed strategies several times, the numbers within one column may sum to more than the total number of players who responded to the questionnaire in the respective treatment.

	Treatment 1	Treatment 2
Win-stay	7	16
Lose-change	3	9
Lose-change-return	0	2
Randomise	16	7
Alternate	9	2
Stay with one action	8	4
Total number of responses	32	25

Table 5: Classification of responses in the questionnaires

Table 5 shows that the subjects' intended play differs much between treatments. Even though in some cases our interpretation of the vague descriptions are surely oversimplifying, we may state that a general trend can be observed. Without information on the payoff scheme subjects make extensive use of simple exploratory multi-round patterns (randomisation, alternation between actions, and staying with one action for multiple rounds). Prior structural information drives subjects to accord much more to the win-stay lose-change dynamics. In treatment 2 two subjects even had the idea of using punishment strategies, i.e. after losing a point these subjects intended to signal discontent by changing their strategy only once and returning back to their original strategy immediately afterwards. These subjects ignored that they might start off with the action that gives zero points to the opponent, thus, initiating a cycle in punishment signals. Yet, these strategies indicate that subjects were much more determined in their seek for the goal, than are the rather "confused" subjects of treatment 1.

While, in treatment 1, none of the subjects dared to state a definitive solution concept, several did so in treatment 2. One subject who clearly stated to have used the win-stay lose-change dynamics, also wrote:

> "Whoever, after 20 rounds, still does not know which action to press, in order to be 'rewarded' by the other player, cannot think logically". (translation from German by the author)

Obviously, with inclusion of prior structural information, some subjects perceived the task as trivial. Yet, considering that a substantial number of subjects still made use of randomisation and multi-round patterns, it seems astonishing that almost all pairs eventually succeeded to coordinate.

5. Conclusion

This small experiment on the impact of structural information prior to repeatedly playing a game in which otherwise little information is provided to the players, led to one clear statement: Those subjects who are not informed about the structure of the environment encounter difficulties in coordinating on the single efficient and individually payoff maximising cell. Few pairs obtain coordination, and those who do, most often reach coordination very early in the sequence. To the contrary, subjects who are informed about the structure of the interaction most often succeed in reaching coordination, even though sometimes pairs need a considerable amount of repetitions.

We failed to attribute the increased success of the informed subjects to characteristics of individual play. In both treatments, play looks extremely variable. Also, in both treatments subjects repeatedly use multi-round patterns, which may be conceived of as investigation of the environment. This use of complicated structures of behaviour, together with the fact that data on behaviour of successful uninformed subjects is scarce, impeded a firm distinction of adaptive behaviour. Our speculation that we might find a difference between differently informed groups when looking at the data on a round-by-round basis, was partly confirmed. Informed subjects, after encountering a win, seem to stay with their action more often than the uninformed counterparts. We could not detect a similar difference between treatments in behaviour after losing. While the scarcity of data on decisions impeded firm, statistically safe statements, we found strong support for our hypothesis in the questionnaires that were handed out after the repeated game was ended.

Our results call for an increased consideration of subjects' ability to process structural information. Past experimental as well as theoretical models of adaptive play, by focusing on the use of feedback, have neglected this point. The question of which pieces of information are processed has to be raised again. The tentative early answer by Smith (1990) that the most important piece of information is the own payoff does not satisfy any longer.

Appendix: Instructions for the little-information treatment

Preliminaries:
You are participating in a study on decision making within the framework of experimental economics. After you have read these instructions we will come to you in order to clarify open questions. If you encounter any further questions during play please raise your hand, so a member of the staff can come to you.
Please do not touch the computer until you are asked to start the game. Tips on the usage of the computer are given on the back page.
During the session you will face a sequence of decision tasks. At each decision you may earn money. How much money you will get depends on your decisions. The aggregated payoff will be paid to you in cash at the end of the experiment. Your decisions as well as your payoff is known only to you, that is we take care that no other participant will get any information on your decisions or payoffs.

Decision task:
You are one of 8 (10) persons. All persons are faced with the same task, that is, all persons have got identical instruction sheets and sit, separated from each other, in front of a computer terminal.
Before the start of the first round, all 8 (10) persons are matched randomly into four (five) pairs. The pairs remain fixed throughout the whole session, that is, the person you are playing with will be the same throughout the whole session. You will never be told whom you are paired with.
Soon, you will be required to make the following identical decision in 100 successive rounds:
In each round you have got the choice between two alternatives: action A and action B. After you and the person you are paired with have entered the decisions upon your choice the computer will calculate your payoff according to a predetermined scheme and subsequently informs you about your payoff.

Payoff scheme:
The payoff scheme will not be made public. What is known is that the payoffs within your pair solely depend on your decision and the decision of your opponent of that round. Particularly, no random process will be used to calculate payoffs. Furthermore, it is known that payoffs can take either of the values, 0 or 1, where 0 is equivalent to 0 DM and 1 is equivalent to 0.30 DM.

Your MAXLAB team

References

Blume, Andreas, Douglas V. DeJong, Gearge R. Neumann, N.E. Savin (fortchcoming) "Learning and Communication in Sender-Receiver Games: An Econometric Investigation" forthcoming in the *Journal of Applied Econometrics*.

Duffy, John and Nick Feltovich (1999) "Does observation of others affect learning in strategic environments? An experimental study" *International Journal of Game Theory* 28, 131 - 152.

Feltovich, Nick (2000) "Reinforcement-Based vs. Belief-Based Learning Models in Experimental Asymmetric-Information Games" *Econometrica* 68, 605 – 642.

Green, Jerry R., Stokey, Nancy L. (1980) "A Two-Person Game of Information Transmission" H.I.E.R. Discussion Paper No.751.

Kagel, John H., Chung Kim, Donald Moser (1996) "Fairness in Ultimatum Games with Asymmetric Information and Asymmetric Payoffs" *Games and Economic Behavior* 13, 100 – 110.

Kajii, Atsushi and Stephen Morris (1997) "The Robustness of Equilibria to Incomplete Information" *Econometrica* 65, 1283 – 1309.

Mitropoulos, Atanasios (2001a) "Learning Under Minimal Information: An Experiment on Mutual Fate Control" *Journal of Economic Psychology* 22, 523 – 557.

Mitropoulos, Atanasios (2001b) "Little Information, Efficiency, and Learning – An Experimental Study" Working Paper No. 8/2001, Faculty of Economics and Management, Magdeburg.

Schelling, Thomas (1960) *The Strategy of Conflict*, Harvard University Press.

Smith, Vernon L. (1990) "Experimental Economics: Behavioral Lessons for Microeconomic Theory and Policy" Nancy L. Schwartz Memorial Lecture, Northwestern University.

Sulganik, Eyal and Itzhak Zilcha (1996) "The Value of Information in the Presence of Futures Markets" *Journal of Futures Markets* 16, 227 – 240.