

Coordination and Self-Organization in Minority Games: Experimental Evidence¹

Giulio Bottazzi

Sant'Anna School of Advanced Studies, Pisa, Italy

Giovanna Devetag

Department of Management and Computer Science

University of Trento, Italy

Preliminary draft prepared for the WEHIA conference, Trieste, May 30-June 1, 2002.

Please, do not cite or circulate without permission.

Abstract

This work presents experimental results on a coordination game in which agents must repeatedly choose between two sides, and a positive fixed payoff is assigned only to agents who pick the minoritarian side. The game presents a variety of asymmetric pure strategy equilibria, and a unique symmetric mixed-strategy equilibrium in which agents randomize between the two sides at every stage. The game reflects some essential features of those economic situations in which positive rewards are assigned to individuals who behave in opposition to the modal behavior in a population. We conduct laboratory experiments in which stationary groups of five players play the game for 100 periods, and manipulate two treatment variables: the amount of 'memory' M that players have regarding the game history (i.e., the length of the string of past outcomes that players can see on the screen while choosing), and the amount of information about other players' past choices: in the aggregate information treatment, players only know which side was the minority side at each period, while in the full information treatment players have information regarding the entire distribution of choices in the group at each round. We first analyze aggregate results in terms of both *allocative* and *informational* efficiency. We then analyze individual behavior in the game as compared to the theoretical benchmark provided by the mixed strategy equilibrium solution. Our results show that, first, both allocative and informational efficiency are higher on average than the benchmark value corresponding to the mixed strategy equilibrium in all treatments, suggesting that a quite remarkable degree of coordination is achieved; second, providing players with full information about other players' choice distribution does not appear to improve efficiency significantly. At the individual level, a substantial portion of subjects exhibit 'inertial' behavior, i.e., the tendency to replicate their previous round choice with a higher frequency than the one prescribed by randomizing behavior, and such inertia seems to be enhanced rather than decreased by a full information treatment.

¹We thank Marco Tecilla and the CEEL staff for support in the implementation of the experimental software and in the carrying out of the experiments. Financial support from the Ministry of Education, University and Research (MIUR) is gratefully acknowledged.

1 Introduction

Many social interaction phenomena imply some advantage that arises from acting in opposition to the modal behavior in a population. For example, if deciding when to leave for vacation with our car, we will be better off avoiding rush hours and travelling when only few other people are on the road. If we want to spend some hours surfing on the web, we will prefer to do that when few other people are surfing so that the connection is faster. If a long-awaited new movie hit is being shown at our local theater, we will prefer to go on those nights in which spending two hours in line for tickets is *least* likely, and so on.

Thomas Schelling in his book on *Micromotives and Macrobehavior* (1978) described and modelled a vast array of social situations that fall under this broad category, from micro-phenomena like the ones just mentioned to far more relevant ones such as racial segregation and house migration. Clearly, what matters in this type of interactive settings is the interplay between the individuals' expectations and iterated expectations (i.e., expectations of expectations), and aggregate behavior in terms of goodness of coordination. The key to the whole problem in fact consists not simply in trying to guess what the majority will do, and do the opposite, but rather to guess what the majority will think that the majority will do, and so on ad infinitum. Intuitively, the more heterogeneous individual expectations are, the 'smoother' - and hence efficient - the aggregate outcome of interaction is likely to be (e.g., no traffic jams on the highways, no band-width congestion, etc.), because different beliefs will induce differentiated enough behaviors so as to produce a kind of self-organization at the population level. Too similar expectations, on the contrary, may be likely to result in large overall inefficiency².

In this work we present experiments on a very simple repeated game in which the payoff to each player is based on a minority rule. In the Minority Game, first introduced by Challet and Zhang (1997, 1998) a fixed group of N players (where N is odd) must privately and independently choose each round between two actions available to them. The players choosing the action that is chosen by the minority of the players earn a fixed, positive payoff, while players who end up on the majority side earn nothing. Hence, the two alternatives are perfectly symmetric and individual payoff is solely based on how players distribute between them³. Formally, the model is an N -person non-cooperative game with multiple asymmetric Nash equilibria in pure strategies, i.e., mutual best response outcomes in which agents differentiate from one another on the basis of the side chosen, and a unique symmetric equilibrium in mixed strategies, requiring that all agents pick the two sides with equal probability. Besides the seminal contributions of Schelling (1978) and Arthur (1988), similar coordination problems have been extensively investigated within the experimental economics literature, usually going under the heading of 'market entry games' (e.g., Mayer et al. 1992; Ochs, 1990, 1995 and references therein). In these games, a group of N players must repeatedly and autonomously decide whether or not to enter one or more markets each having a fixed capacity k , with $k < N$. It is generally assumed that entrants receive a positive profit - often declining with their number - up to a capacity threshold, beyond which the incentive to enter disappears (Non entrants may or may not receive a fixed positive payoff from their choice). If too many people enter the market, they all suffer losses, while if too few people enter, opportunities for gain are left unexploited; therefore, aggregate efficiency is maximal when the average number of entrants equals market capacity, which corresponds to the outcome occurring at any of the games' Nash equilibria.

²This point had been first made by B. Arthur (1989) in his 'El Farol Bar' model of interaction.

³For a large collection of papers concerning both the analytical and numerical explorations of the original minority game and various extensions see also <http://www.unifr.ch/econophysics/>.

Several variations of the basic ‘market-entry’ setup have been studied experimentally: some involve changes in the institutional details (e.g., the payoff function, the presence or absence of irrecoverable entry fees, the payoff-asymmetry implied by the pure equilibrium solutions, etc.) which in turn imply different long-run stability properties of the resulting equilibria; other modifications involved features of the experimental design (e.g., group size, stationary vs. non-stationary sets of players, etc.; for a discussion see e.g. Zwick and Rapoport, 1999). Interestingly, the experimental evidence gathered so far appears scarcely consistent with equilibrium predictions, especially at the level of individual behavior, while revealing a high and persistent degree of heterogeneity among players’ beliefs and strategies (see also Erev and Rapoport, 1998; Rapoport, Seale, Erev and Sundali, 1998; Sundali, Rapoport and Seale, 1995). Further, players’ individual behavior, and consequently patterns of equilibrium selection, seem to largely depend on the specification of the institutional details of the market, making it hard to make direct comparisons between experimental data.

This work is aimed at answering the following research questions: what is the degree of efficiency of the system when groups of players interact repeatedly compared to the benchmark value offered by the mixed strategy equilibrium solution? What is the impact - in terms of efficiency - of varying the amount of information that players have regarding the game history and the past actions of the other players? Do players randomize between rounds, as prescribed by the Nash equilibrium solution? Or are they able to detect and exploit predictable patterns in others’ behavior, when they are given more information?

In order to answer these questions, we conducted experiments in which stationary groups of five players played a minority game for 100 rounds, and we applied a 2×3 factorial design in which the two treatment variables were the amount of ‘memory’ M that players could have regarding the game history (i.e., the number of past rounds for which information was made visible on the subjects’ computer screen), and the amount of information that players had on the past choices of the other players (see section 3 for a description of the experimental design).

Our preliminary results can be roughly summarized as follows: first, efficiency is higher on average than the benchmark value corresponding to the game-theoretic solution in all treatments, suggesting that a quite remarkable degree of coordination is achieved; second, providing players with full information about other players’ choice distribution does not appear to improve efficiency significantly. At the individual level, a substantial portion of subjects exhibit ‘inertial’ behavior, i.e., the tendency to replicate their previous round choice with a higher frequency than the one prescribed by randomizing behavior, and such inertia seems to be enhanced rather than decreased by a full information treatment. Finally, analysis of individual strings of choices reveals a remarkable degree of heterogeneity in the decision rules adopted.

The paper is organized as follows: Section 2 describes the game-theoretic framework and discusses previous related work. Section 3 describes the experimental design, and section 4 analyzes the results. Finally, section 5 offers some concluding remarks and directions for future research.

2 The Minority Game

The minority game is played repeatedly by a stationary set of N players, where N is an odd number. On each period of the stage game, each player must choose privately and independently between two actions or sides which will label 0 and 1. The payoff function,

which is the same for all players is given by

$$\pi_i = \begin{cases} 1 & \text{if } k_i \leq (N-1)/2 \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

where $i \in \{0, 1\}$ and k_i is the number of players choosing i . In descriptive terms, players are rewarded with one point each whenever the side they choose is chosen by the minority, while players who end up on the majority side obtain a null payoff.

It is straightforward to see that the game has $N!/(((N-1)/2)!)^2$ asymmetric Nash equilibria in pure strategies, in which exactly $(N-1)/2$ players choose either one of the two sides. The pure strategy Nash equilibria are Pareto efficient but they are not strict, as players on the majority side are just indifferent between deviating and sticking to equilibrium play. Given the payoff asymmetry and the resulting fragility of the pure strategy equilibria, one can rule out the possibility that groups may develop a simple form of tacit coordination based on historical precedent, as it may occur in analogous games (e.g., Meyer et al., (1992)). The game also has a unique mixed-strategy Nash equilibrium, in which each player selects the two actions with equal probability. Besides the unique symmetric mixed strategy equilibrium, there are infinite asymmetric mixed strategy equilibria.

From a game-theoretic standpoint, the game is analogous to the market entry games studied experimentally by Ochs (1990), Meyer et al. (1992), Sundali et al. (1995), Rapoport et al. (1998), Seale and Rapoport (2000). In most of these games, a group of N players must decide whether or not to enter a market that has a certain fixed capacity k , where k is typically $\leq N/4$. Some of the parameters of the payoff function may vary so as to render the ‘staying out’ option more or less attractive to players (in some cases players receive a fixed payoff from not entering). In addition, the payoff to entrants may vary or not depending on their number when this is less than the capacity value. In all such games there is usually a unique symmetric mixed strategy equilibrium and several asymmetric pure strategy equilibria. The minority game, although belonging to the same class of coordination games, differs from the previously studied market entry games in some crucial aspects: its equilibrium structure is extremely unstable relative to other games, in that all pure strategy equilibria, whereby players ‘lock in’ in the choice of the same side, besides being non strict, and hence subject to being disrupted by minor deviations, imply a payoff-asymmetry between the two resulting parties which strongly increase their instability; for this reason, MG may be the ideal setup to study the emergence of collective efficiency resting on some forms of dynamic coordination between players. Moreover, while in market entry games the value of the market capacity usually changes from period to period thus changing the probability distribution implied by the mixed strategy equilibrium, in the MG this remains fixed throughout the game and it is very simple, as it assigns equal probability to the two available options; this, combined with the large numbers of rounds, the small group size and the strongly competitive nature of the game, should provide a strong incentive to random behavior and should allow for learning, thus making the unique mixed strategy equilibrium a natural ‘benchmark’ against which to confront actual behavior. Deviations from the mixed strategy hypothesis hence may be interpreted as the result of decision rules - other than randomization - which may depend on the game structure and to the history of interaction rather than as deriving from subjects’ ‘imperfect’ ability to generate random series (e.g., Budescu, 1987).

⁴In other cases, as in Meyer et al. (1992), there are actually two markets between which players must allocate their entry decisions.

3 The experiment: design and implementation

The information treatments

A total of 120 subjects participated in the experiment. Four groups of five players participated in the single treatments. A group size of five seems large enough as to prevent subjects from developing repeated game strategies aimed at achieving tacit coordination. Indeed, such a form of dynamic coordination is not trivial to achieve in the minority game even with a small group, because it would require a complex rotation in the chosen side between players⁵. The number of rounds for each group was kept constant and equal to 100.

Table 1 summarizes the experimental design.

	Memory (M)		
	$M = 1$	$M = 4$	$M = 16$
Information			
Aggregate (A)	4	4	4
Full (F)	4	4	4

Table 1: The 3X2 factorial design of the experiments. Numbers in the cells indicate the number of groups that were assigned to each treatment

As previously stated, in this work we are mainly interested in studying the impact of information on aggregate efficiency and individual behavior.

Several previous simulation studies on the MG (e.g., Challet and Zhang, 1997, 1998) have focussed on the role of the information available to players in determining aggregate efficiency. In most of these studies the amount of information given to artificial agents was varied by changing the value of the ‘memory’ parameter M , defined as the length of the string of past outcomes that agents were able to retain and process when making their choices. Indeed, these works have shown that the amount of agents’ memory (and hence the amount of information on the ‘history’ of interaction) has a significant impact on the system’s degree of self-organization when artificial agents are designed to behave according to experience-based strategies.

In line with previous simulation studies, we also chose to focus on the role of information; moreover, in this study we decided to jointly investigate the impact of both information about the series of past outcomes (i.e., the players’ available memory capacity) and information (or lack thereof) about the behavior of single players in the population.

As far as the first parameter is concerned, in the case of human agents memory itself cannot be directly controlled for; hence, a rough ‘proxy’ for it was devised by varying the number of past outcomes for which information was disclosed and remained visible on the subjects’ computer screens while playing. We chose three levels of memory: short ($M = 1$), medium ($M = 4$), and long ($M = 16$). In the last case, we assume that the amount of information made available is considerably more than what individuals are able to process, due to well known human working memory constraints (e.g., Miller, 1956).

The second treatment variable was the type of information provided to subjects at each round of play. In the *aggregate* information treatment, subjects were only allowed to know which side (0 or 1) was the winning side at each round (i.e., they only knew the aggregate outcome of interaction at every time step), while in the *full* information treatment, subjects could also see the entire distribution of individual choices in their group.

⁵This form of tacit cooperation was in fact never observed in our experiments.

Previous experimental studies have shown that providing information about other players' choices usually has a significant impact on individual behavior in repeated games, although such impact tends to be game-specific and not necessarily efficiency-enhancing. Duffy and Hopkins (2001) showed that when players had full information about others' choices, they converged much more quickly to equilibrium play in a repeated market entry game with strict pure strategy equilibria. On the contrary, providing subjects with full information about others' choices lowered efficiency in a large-group coordination game with multiple Pareto-ranked equilibria (Van Huyck, Battalio and Beil 1990). Hence, information about others' behavior, though used by players to modify their choices, not necessarily leads to better collective outcomes compared to the case in which only minimal or no information is available. In the case of a minority game, it would be reasonable to expect that when players have full information (and such condition is common knowledge), they might try to render their behavior less predictable, knowing that other players could exploit patterns in their choices, and might at the same time try to exploit patterns in others' behavior, provided they can detect any of them. Thus, it is plausible to expect that a full information treatment would have a beneficial effect on aggregate efficiency in a minority game.

Implementation

Subjects on each round had to choose between two actions that were labelled A and B in the experiment. The information that each player received at the end of every round was the visualization of the action chosen in that round (A or B), the payoff gained (1 or 0), the additional information prescribed by the experimental treatment, and his or her cumulative payoff, expressed as a percentage of successful rounds over the total at that point⁶

The experiment was ran at the Computable and Experimental Economics Laboratory of the University of Trento. Subjects were recruited from different departments of the University of Trento and had never participated in experiments of this type before. The experiment was conducted in six sessions of 20 players each, and it was computerized. Subjects were randomly assigned a seat at the lab computer room upon arrival, and were given written instructions⁷. They could see each other, but no form of communication was allowed during the entire session.

Subjects were told that they would be randomly divided in four groups of five players at the beginning of the experiment and that such division would be known only to the computer program. Then they would participate in a repeated decision-making experiment with the same group of people for a total of 100 periods. Each subject received a fixed show up fee plus anything he or she could earn in the experiment. Payoffs in the game were expressed in 'experimental points'. At each round a subject could earn either 1 point or zero, and the monetary exchange rate was fixed at (E.18 for each point, for a theoretical maximum of (E25.82 per subject (show up fee included). When all members of a group entered their choice, the computer calculated each player's payoff and disclosed information about the round. Sessions lasted fortyfive minutes on average, including instruction time.

⁶Subjects were explicitly told that in case all five members chose the same action, they would all earn a payoff of zero, and the computer would display the action not chosen by anyone as the winning action.

⁷A translation of the instructions is given in Appendix.

	Aggregate info			Full info		
	M=1	M=4	M=16	M=1	M=4	M=16
group 1	66	76	68	60	75	76
group 2	56	66	79	73	70	81
group 3	67	71	60	71	64	71
group 4	66	68	63	56	61	74
average	63.75	70.25	67.5	65	67.5	75.5

Table 2: Relative frequencies of equilibrium play by group and treatment

time steps	$M = 1$		$M = 4$		$M = 16$	
	A	F	A	F	A	F
25	1.31	1.05	.99	1.45	1.25	.91
50	1.25	1.29	.83	1.13	1.29	.71
75	1.19	1.31	.93	.85	.89	.97
100	.99	1.03	.91	.77	.89	.61

Table 3: Values of Σ averaged over successive 25 time steps and across groups, reported separately for values of M and for the aggregate (A) and full (F) information treatments

4 Results

4.1 Aggregate behavior

Group level test of the equilibrium solution is quite straightforward. Tab. 2 reports, for each group separately, the percentage of times over the 100 rounds in which the size of the ‘winning’ minority was exactly equal to 2 players, an allocation which would correspond to any of the game’s Nash equilibria (the pure and mixed strategy equilibrium proportions for the two parties are equal, therefore at the aggregate level we do not distinguish between the two). Table 2 reports the percentage of equilibrium play separately by group and by treatment.

Two observations are noteworthy: first, the equilibrium solution seems to describe aggregate behavior rather well, accounting for roughly 68% of the history of play. Since the percentage of equilibrium play in the case of perfectly symmetric players who follow the mixed strategy equilibrium is equal to 62.5%, in first approximation we can conclude that the level of coordination that players achieve in all treatments is at least as high as that implied by the Nash solution.

Second, there are no statistical differences between treatments (Kruskal-Wallis and Kolmogorov-Smirnov nonparametric tests). Note that the percentage of play compatible with the equilibrium solution provides likewise a measure of the system’s degree of *allocative efficiency*. In fact, the further the size of the winning minority is from the equilibrium value of $(N - 1)/2$, the larger the amount of money which is ‘left on the table’, hence the lower the resulting efficiency.

We next analyze the degree of allocative efficiency over time. In order to do this, we define a complementary measure of efficiency provided by the magnitude of fluctuations around the game’s Nash equilibrium solution. Hence, we define the *allocative efficiency* Σ as

$$\Sigma = \sum_t (N_0(t) - \frac{N}{2})^2 \quad (2)$$

time steps	A	F
25	1,18	1,14
50	1,12	1,04
75	1,00	1,04
100	0,93	0,80

Table 4: Values of Σ over time averaged across values of M for the aggregate and full information treatments

time	$M = 1$	$M = 4$	$M = 16$
25	1,18	1,22	1,08
50	1,27	0,98	1
75	1,25	0,89	0,93
100	1,01	0,84	0,75

Table 5: Values of Σ over time averaged across information treatments for different values of M

where $N_0(t)$ is the number of players choosing party 0 at time t and N is the total number of players participating in the game. The value of Σ can be comprised between 0.25 and 6.25. Clearly, the lower its value, the higher the efficiency.

Tab. 3 reports the values of Σ over time separately by the value of M and averaged across groups in both the aggregate and full information treatments.

The ‘benchmark’ value for Σ in the case of random symmetric players is equal to 1.25. Note that the degree of allocative efficiency in all treatments is significantly higher than the benchmark value. Furthermore, a general increasing trend in efficiency is clearly detectable (see also Tab. 4 and 5).

Hence, all groups in the experiment are able to achieve a quite remarkable degree of self-organization around the Nash solution, with a higher achieved efficiency than the one attainable through purely random play. Recall that the highest degree of efficiency would be obtained in correspondence of any of the pure strategy Nash equilibria, whereby players repeatedly choose the same side. However, as such equilibria are not strict they were in fact never observed.

Given the high degree of allocative efficiency obtained by all groups, and hence the high level of dynamic coordination between players, one might ask whether such coordination produces ‘predictable’ patterns in the series of aggregate outcomes. In fact, such a good level of coordination could be produced by the players’ use of repeated choice patterns that could reflect themselves into regularities at the aggregate level.

Indeed, if this was the case, the observed level of allocative efficiency would not be ‘optimal’ because it could be in principle further increased by some hypothetical arbitraging behaviors aimed at exploiting the predictability in the time series of outcomes.

In order to investigate the issue, we introduce a measure of *informational efficiency*, which describes the amount of ‘informational content’ of the history of interaction. In fact, if players behave randomly or if there is a significant degree of heterogeneity and change in their behavior, the aggregate information represented by the history of interaction - i.e., the string of successive winning sides at each round - should be perfectly described by a series of i.i.d. draws and hence render the outcome in the next period of play perfectly unpredictable. The resulting ‘market’ would be, in a sense, perfectly efficient. Otherwise, the presence of ‘structure’ in the time series would indicate the presence of some ‘slack’ in the players’ coordination patterns that could in principle be exploited to increase the system’s allocative efficiency.

Note that the degree of informational efficiency of the resulting ‘market’ too is an aggregate measure of closeness to the mixed strategy equilibrium solution. In fact, mixed strategy equilibrium behavior implies that the series of winning sides be a sequence of unrelated draws.

We introduce a measure of *informational efficiency* Λ defined as

	Aggregate Info.			Full Info.		
	Mean	St. dev.	K.S. test	Mean	St. dev.	K.S. test
All	32.6	6.67	0.039	33.15	6.41	0.008
$m = 1$	31.7	5.46	0.56	31.9	5.22	0.56
$m = 4$	33.7	9.13	0.016	32.75	4.89	0.11
$m = 16$	32.6	4.77	0.16	34.8	8.48	0.03
Theoretical	31.25	4.63				

Table 6: Payoff distribution analysis

$$\Lambda(l) = \sum_{h \in I_l} \max\{F(0|h), F(1|h)\} \quad (3)$$

where I_l is the set of all the binary strings of length l and $N(i|h)$ with $i \in \{0, 1\}$ is the number of times i has been the winning side when h was the preceding history of length l (i.e. the string of the previous l winning sides). From its definition, $\Lambda(l)$ represents the average payoff of the (ex-post defined) best strategy of a hypothetical player who has access to the last l rounds of the history of play. It follows that $\Lambda(l)$ is a non decreasing function of l .

The results for the informational efficiency are shown in Fig. 1 and Fig. 2 for the aggregate and full information treatments respectively. We calculated the values of Λ for different values of the history length l . The dotted lines represent the 2 st.dev. band relative to the usual ‘benchmark’ case. Interestingly, as the figures show, irrespectively of the information and memory amounts, the observed informational efficiency departs from the ‘benchmark’ values for less than one standard deviation, suggesting a substantial lack of structure in the outcome time series.

4.2 Analysis of payoffs distribution

The last aggregate analysis concerns the total payoff distribution over the players population in all treatments. The natural benchmark to which compare the observed distribution is again the “random players” case where all the players involved choose randomly with equal probabilities between the two sides.

If N is the number of players the winning side can be of size $0, \dots, (N-1)/2$ and the probability for any given player of belonging to a winning side of size k is k/n . Then the probability to be a winner can be computed as

$$p_w = \sum_{h=0}^{(N-1)/2} 2 \binom{N}{h} \frac{h}{N} \quad (4)$$

where the factor 2 accounts for the symmetry of the game and the probability of winning k matches out of M is given by a binomial distribution:

$$p(k; M, p_w) = \binom{M}{k} p_w^k (1 - p_w)^{M-k} \quad (5)$$

In our case $N = 5$ and since we are concerned with the total payoff in each treatment, $M = 100$. Tab. 6 shows the mean and average payoff for all treatments and for the subpopulation

obtained considering only treatments with a given m . The distributions so obtained are then compared with the theoretical distribution defined in Eq. 5 by using the Kolmogorov Smirnov statistics (von Mises, 1964). As can be seen, the null hypothesis of identity between the observed distributions and the theoretical prediction based on purely random players can never be disproved with high significance, even if the variance of the observed payoffs seems larger especially in the $m = 4$ case.

Hence, concluding the section on aggregate behavior, we can summarize the main results as follows: our groups of players were able to achieve a high degree of coordination - and hence were fully able of exploiting the game's gain opportunities - without producing structure in the times series of aggregate outcomes, i.e., producing a correspondingly high level of informational efficiency. Finally, the analysis of payoffs distribution reveals that the achieved gains were uniformly distributed among the populations, at least relative to the theoretical benchmark of perfect randomizing players. In the next section we will analyze how players implemented individually the observed collective self-organization.

4.3 Individual Behavior

In this section we analyze individual behavior in the game. First, we check whether choices are biased toward one of the two possible options A and B due to their labels. One may in fact hypothesize that players react to choice labels in various "irrational" ways or in the effort of using them as a coordination device. A binomial test performed on choices pooled across treatments failed to reject the hypothesis of no association between action label and choice⁸. Hence, the results of the subsequent analysis is collapsed across the label variable.

We next compare individual data with the theoretical benchmark represented by the symmetric mixed strategy equilibrium solution. There is considerable evidence that individuals have difficulties in generating random sequences, even when explicitly asked to do so. Many of them, in fact, exhibit a negative recency bias, generating series with too many alternations compared to those of a purely random series (e.g., Budescu, 1987). This phenomenon seems related to the representative heuristic (Kahneman, Slovic and Tversky, 1982) according to which a binary series with many alternations is generally judged as very representative of a random series, although it is not random. However, in competitive game-theoretic settings in which subjects repeatedly play games with a unique, mixed strategy equilibrium, the performance is generally better, as subjects have an incentive to render their future actions unpredictable (Rapoport and Budescu, 1992), although the negative recency bias does not disappear completely (Rapoport and Boebel, 1992).

In the minority game, the amount of information available to players might influence the extent to which they behave randomly; in particular, in the full information treatment, in which everybody's choice is disclosed to all group members, players may be more induced to render their choice unpredictable and may try to exploit predictable patterns in others' choices, and this is more likely to occur the higher the value of M (the longer the string of the game history). However, a longer 'memory' of one's own past performance associated to each choice may induce players to take more into account their personal history of gains and losses when choosing, which goes against randomizing behavior (which would require a 'memoryless' selection procedure).

In the aggregate information treatment, a binomial test performed on each individual

⁸Choices of A are 50.2% of all choices (binomial test, $p = .05$). From now on, all statements on statistical significance based on nonparametric tests will always be based on the conventional significance level of .05.

sequence of choices allows to reject the null hypothesis of a binomial distribution in the 35% of cases (5 of 20 players in both the $m = 1$ and $m = 16$ treatments, and 6 of the 20 players in the $m = 4$ treatment). Few of these players followed almost pure strategies, i.e., they almost always chose the same side. The vast majority of the players did not follow pure strategies. However, only few of the mixed strategies can be accounted for by the mixed-strategy equilibrium.

In the full information treatment, the percentage of players' strings for which the binomial distribution can be rejected rise up to 52% (9 of 20 players in the $M=1$ treatment, 5 of 20 players in the $M=4$ treatment, and 17 of 20 players in the $M=16$ treatment), a datum that seem to go against the intuition that increased information would increase randomizing behavior.

Let us then introduce a finer-grained investigation of individual choice patterns. We will next detect the presence of 'structure' in the time series of players' choices, where by 'structure' we mean any deviation from the null hypothesis of mixed strategy equilibrium.

In order to provide a synthetical description of possible general relations between a player's choice and the information he possesses about the game (i.e., his previous choices and the information about past aggregate outcomes and, in the full information treatment, the individual choices of others) we use the measure of *entropy*, which defines the informational content of a probability distribution.

In order to clarify the analysis that follows, let us first review a few formal definitions:

Given a random variable X taking discrete values with probabilities $\{p_1, \dots, p_N\}$ its entropy is defined as

$$H(X) = - \sum_{i=1}^N p_i \log(p_i) \quad (6)$$

and characterizes its information content (see [17]). In the same way, considering two random variables X and Y , one can compute the entropy of the the joint distribution $H(X, Y)$ and the entropy of one variable with respect to the other, for instance of the variable Y with respect to the variable X which is defined as $H(Y|X) = H(X, Y) - H(X)$. This quantity can be thought to represent the increase in information obtained observing the realizations of the variable Y when the variable X is already known. If $H(Y|X) = H(Y)$ the increment in information is exactly the information content of the variable itself, in other words, the previous knowledge of X doesn't reduce the uncertainty on Y and hence the two variables can be considered to be unrelated. Indeed, it is easy to show that the relation $H(X, Y) = H(X) + H(Y)$ holds if and only if X and Y are independently distributed random variables.

Following from the previous properties the "uncertainty coefficient" of a variable Y with respect to a variable X can be defined as:

$$U(Y|X) = \frac{H(Y) - H(Y|X)}{H(Y)} = \frac{H(X) + H(Y) - H(X, Y)}{H(Y)} \quad (7)$$

This "coefficient" represents the fractional contribution to the knowledge of Y given by the previous knowledge of X , i.e. which fraction of the possible realizations of the variable Y can be predicted knowing the actual realizations of X . If $U(Y|X) = 0$ the variables are independent, while if $U(Y|X) = 1$ they are identically distributed; therefore U can be considered a measure of probabilistic dependence between X and Y . Notice that in this method no ad-hoc assumption is made on the nature of the probability distributions and on the form of the relation between them.

l	m=1	m=4	m=16	bench.
1	-4.39 (-6.44,-2.95)	-6.41 (-7.39,-4.53)	-4.74 (-5.24,-3.23)	-6.06 (-7.22,-4.7)
2	-3.22 (-4.19,-2.53)	-4.00 (-4.99,-3.18)	-3.42 (-4.02, -2.38)	-4.14 (-4.68,-3.44)
3	-2.49 (-3.04,-1.89)	-2.84 (-3.33,-2.41)	-2.52 (-3.16,-1.95)	-3.05 (-3.4,-2.64)
4	-1.77 (-2.14,-1.50)	-1.97 (-2.34,-1.69)	-1.74 (-2.08,-1.52)	-2.09 (-2.33,-1.83)
1	-4.94 (-5.45,-3.20)	-4.61 (-5.53,-2.65)	-4.84 (-6.01,-3.60)	-6.06 (-7.22,-4.7)
2	-3.56 (-4.40,-2.99)	-3.41 (-4.11,-2.44)	-3.51 (-4.42,-2.83)	-4.14 (-4.68,-3.44)
3	-2.71 (-3.16,-1.97)	-2.43 (-3.02,-1.94)	-2.34 (-2.78,-1.88)	-3.05 (-3.4,-2.64)
4	-1.94 (-2.31,-1.56)	-1.78 (-2.19,-1.48)	-1.65 (-1.92,-1.50)	-2.09 (-2.33,-1.83)

Table 7: Mean and st. dev. for the informational content U of the string of choices with respect to the string of past outcomes for the aggregate info. (top) and the full info. (bottom) cases. The values for four different depths $l = 1, 2, 3, 4$ are reported.

l	m=1	m=4	m=16	bench.
1	-3.14 (-3.49,-2.27)	-3.96 (-4.80,-3.21)	-4.36 (-5.12,-2.73)	-6.06 (-7.22,-4.7)
2	-2.44 (-2.99,-2.06)	-2.81 (-3.55,-2.18)	-3.00 (-3.84,-2.40)	-4.14 (-4.68,-3.44)
3	-2.11 (-2.57,-1.80)	-2.13 (-2.76,-1.70)	-2.36 (-2.69,-1.93)	-3.05 (-3.4,-2.64)
4	-1.54 (-1.71,-1.48)	-1.55 (-1.81,-1.42)	-1.75 (-2.02,-1.40)	-2.09 (-2.33,-1.83)
1	-4.94 (-5.45,-3.20)	-4.61 (-5.53,-2.65)	-4.84 (-6.01,-3.60)	-6.06 (-7.22,-4.7)
2	-3.56 (-4.40,-2.99)	-3.41 (-4.11,-2.44)	-3.51 (-4.42,-2.83)	-4.14 (-4.68,-3.44)
3	-2.71 (-3.16,-1.97)	-2.43 (-3.02,-1.94)	-2.34 (-2.78,-1.88)	-3.05 (-3.4,-2.64)
4	-1.94 (-2.31, -1.56)	-1.78 (-2.19,-1.48)	-1.65 (-1.92,-1.50)	-2.09 (-2.33,-1.83)

Table 8: Mean and st. dev. for the informational content U of the string of choices with respect to the string of past choices for the aggregate info. (top) and the full info. (bottom) cases. The values for four different depths $l = 1, 2, 3, 4$ are reported.

For what concerns our analysis, we fix a “memory depth” l and consider the joint probability distribution $p(c, h)$ where $c \in \{0, 1\}$ is a player’s choice at a given round and $h \in \times^l \{0, 1\}$ is the past history of winning sides of length l at that time. This distribution can be computed for each player in each treatment from the observed frequencies and analogously can be computed the probability distribution of choices $p(c)$ and of past histories $p(h)$.

Using Eq. (7) we can define an “uncertainty coefficient” $U_{\text{hist}}(l) = U(c|h)$ of choices with respect to previous outcomes for each player in each treatment and obtain, for a given memory m , a distribution of $5 \times 4 = 20$ points. In Tab.7 we compare their means with a “benchmark” distribution obtained by a Monte Carlo simulation using group of 5 perfectly random player. The same procedure is repeated using, instead of the history of past aggregate outcomes, the string of the l past choices of the player himself, obtaining a distribution for $U_{\text{self}}(l)$ compared with the Monte Carlo distribution in Tab.8. Notice that, due to the relatively short length of the time series, the length parameter l cannot take too large values. Therefore we limit our analysis to $l \leq 4$.

The first thing to notice in Tab.7 and Tab. 8 is that, irrespectively of l and m , the relationships between a player’s present choice and his own past choices results stronger that

	M=1		M=4		M=16	
	A	F	A	F	A	F
$c(t+1)=c(t)$.63 (.11)	.68 (.12)	.62 (.14)	.65 (.14)	.61 (.12)	.74 (.13)
$c(t+1)=c(t)$ if $\pi(t) = 1$.70 (.18)	.77 (.19)	.70 (.18)	.70 (.20)	.70 (.21)	.83 (.14)
$c(t+1)=c(t)$ if $\pi(t) = 0$.59 (.14)	.65 (.13)	.58 (.13)	.61 (.15)	.55 (.11)	.70 (.15)
$c(t+1)=\text{winner}(t)$.49 (.12)	.49 (.08)	.51 (.06)	.49 (.12)	.52 (.09)	.51 (.07)

Table 9: Frequencies of choices of actions that were equal to the previous period action and to the previous period winning action separately for each treatment (standard deviations are reported in parenthesis)

$P\{c(t+1) = c(t)\}$	if $\pi(t) = 1$		if $\pi(t) = 0$	
	A	F	A	F
mean	.70	.76	.57	.65
st.dev.	.13	.18	.07	.15

Table 10: Mean and standard deviation for the probability of choosing an action equal to the previous period action on the pooled population in the aggregate and full information treatments

the one between present choices and previous outcomes. It might appear that players tend to repeat, quite irrationally, their previous choices rather than trying to spot and exploit possible patterns in the series of group outcomes, a behavior that would be suggested by any “sophisticated” approach to the game. Moreover, no substantial difference can be noticed between treatments, with one exception: the $m = 4$ aggregate information treatment, which is the one characterized by the highest allocative efficiency, is characterized by a substantial lack of “short range” ($l = 1, 2$) relationship between players’ present and past outcomes (see Tab.7).

A qualitative analysis of choices over time confirms these observations. No statistical differences were found at this level of analysis between treatments⁹. Players in the game choose an action equal to their previous period action, on average, 65% of times, and such percentage is significantly different from the mixed strategy equilibrium hypothesis (binomial test). Hence, players exhibit a certain degree of inertia in their behavior. Such persistence is reinforced when the previous period action has been successful (i.e., when players chose the “winning side”). In fact, in this case players repeat their last period action 73% of the times. The percentage after a “losing” round decreases to 61%, but it is still significantly different from the mixed strategy equilibrium hypothesis. Moreover, the ‘inertia’ and ‘reinforcement’ effects appear even strongly in the full information treatments, as Tab. 10 shows, and they reach their maximal values in correspondence of the F($M = 16$) treatment, i.e., in the treatment in which the highest amount of information was given to subjects, as shown in Tab.9. Hence, increasing information seem to increase the tendency toward ‘inertial’ behavior.

On the other hand, players do not seem to react to the previous period “aggregate information” , i.e., to what action was actually the winning one. In fact, players choosing an action equal to the last round winning action is overall equal to 50.3%, which is perfectly consistent

⁹Kruskal-Wallis and Kolmogorov-Smirnov tests. Each group value was treated as a single independent observation.

with the hypothesis of randomizing behavior.

However, the presence of heterogeneity in players' behavior is likewise revealed, both by the variances reported in Tab. 9 and 10 and, finally, by the analysis reported in Fig. 3, where $p_i(a|b, c)$, the probability that player i chooses $a \in \{0, 1\}$ given his choice was $b \in \{0, 1\}$ in the previous round and the winning side was $c \in \{0, 1\}$, computed starting from the choices string of each player, is plotted. The two tendencies of 'inertia' and 'reinforcement' clearly appear from the plots; however, despite these central tendencies, heterogeneity is relatively high, especially in the (b) and (c) cases.

4.4 Monte Carlo simulations of agents' behavior

In the previous analysis we stressed the presence of structure in the agents' individual paths of play, generated both by their "inertia" in changing the chosen side and by their "reinforcing" attitude toward preferring (on average) the last winning side.

A natural question to ask at this point would be whether these features of players' decision behaviors are enough to account for the observed values of allocative and informational efficiency. In other words, are "inertial" and "reinforcing" behaviors the sole ingredients required in order to obtain the aggregate properties we actually observe? If this was the case, we then could conclude that the "self-organizing" properties revealed by the aggregate dynamics and discussed in section 4.1 can be simply explained by the appearance of heterogeneity in the population resulting from differing degrees of "inertia" and "reinforcement" in the agents' choice paths.

On the other hand, one could alternatively imagine that this simple 'parametric' heterogeneity in behaviors is not enough to explain the observed dynamics, and that heterogeneity does in fact possess a stronger structure in terms of emerging relationships among the players' strategies.

In order to provide an answer, we performed extensive numerical simulations whereby we tried to reproduce the observed agents' behaviors. We considered artificial agents whose strategies of choice between the two actions were designed as a mixture of reinforcing and inertial effects following the "behavioral" distribution observed among real agents. In particular, we have endowed each agent from the beginning with a probability p_w of repeating his previous round choice in case it was successful, and a different probability p_l of repeating the same choice if this was the losing choice. These probabilities were randomly extracted from distributions $N(\mu_w, \sigma_w)$ and $N(\mu_l, \sigma_l)$ shaped so as to mimic the distribution shown in Fig. 3.

We repeatedly simulated groups of 5 artificial agents and recorded both their choices and the resulting aggregate outcomes. We explored the parameters space performing a Monte Carlo simulation of 50 independent games of 1000 time steps for values in the range $[\cdot 5, \cdot 8]$ for the means and $[0.0, \cdot 2]$ for the standard deviations.

We found out that, on average, both the allocative and informational efficiency of these non symmetric games (i.e., with "inertial" and "reinforcing" effects) are lower than the symmetric random games characterized by an allocative efficiency equal to $\Sigma = 1.25$ and an informational efficiency equal to $I = \cdot 5$. Fig. 5 and Fig. 4 show examples of typical results.

We can then conclude that characterizing agents' behaviors in terms of average parameters distributions does not suffice in accounting for the degree of efficiency observed in the laboratory, in that the performance obtained in the simulations is even lower than the 'benchmark' case of random, symmetric players. As a consequence, not only players' heterogeneity in beliefs and behaviors seems to constitute an important component of the groups' ability to reach high levels of efficiency but it also seems to be generated by the players' capability

of adapting their own strategies to the ones played by others, i.e. by their “self-organizing” ability.

5 Conclusion and outlooks

The main research question of this work was to investigate, first, the extent to which groups of human players are able to coordinate efficiently in a minority game, and, second, if and to what extent the amount of information available to them makes a difference in terms of both their individual choice patterns and the resulting aggregate efficiency. Our main results can be summarized as follows: at the aggregate level, the degree of both allocative and informational efficiency achieved by our groups of subjects are remarkably high in all treatments when compared to the theoretical benchmark represented by the case of hyper-rational players who follow the unique mixed strategy equilibrium strategy. The lack of a significant difference between treatments suggests that players only need minimal information to coordinate efficiently, and disclosing more information does not necessarily lead to improved aggregate performance.

Second, analysis of individual behavior in the game reveals that such efficiency is not achieved by means of ‘rational’ (i.e., randomizing) behaviors, but rather through a co-evolution in players’ - ‘suboptimal’ - choice rules. In fact, analysis of individual strings of choices reveals the presence of structure; in particular, the majority of players in all groups exhibit ‘inertial’ behavior, i.e., the tendency to replicate their previous round choice more often than what would be implied by pure randomization, and this tendency increases when the previous round choice was successful, suggesting the presence of a ‘reinforcement’ effect (e.g., Roth and Erev, 1998). More over, both inertia and reinforcement behaviors get strengthened in the full information treatment, while there is no significant relation between players’ choices and the previous rounds’ ‘winning’ sides, suggesting that players focus relatively more on their personal series of choices than on the behavior of others. Evidence on reinforcement effects confirms previous experimental studies on market entry games showing that subjects seemed to be affected in their current choices by their personal histories of gains and losses (Zwick and Rapoport, 1999).

Finally, simulations performed by endowing artificial agents with behavioral parameters drawn from the experimental data reveal that the two tendencies of ‘inertia’ and ‘reinforcement’ are not sufficient to account for the observed efficiency. Therefore, the good level of coordination observed in the lab presumably rests upon forms of dynamic mutual adaptation among players’ strategies. An additional treatment whereby groups’ compositions vary from round to round (random assignment) would allow to measure the magnitude of such ‘self-organizing’ property.

In addition, further research might involve a sound confrontation with experimental data from similar coordination games (like the aforementioned market entry games), tests of learning models and further simulations studies in order to pin down the behavioral and institutional conditions that favor the emergence of spontaneous, decentralized coordination in this large class of interactive settings.

Appendix: Instructions

This is a translation of instructions given to subjects in one of the treatments. The complete set of instructions is available from the authors upon request.

Welcome!

You are about to participate in an experiment on decision in contexts of interaction. Your earnings in the experiment will depend on your own decisions and on other participants' decisions. If you make appropriate decisions, you will be able to earn an appreciable amount of money which will be paid to you privately and in cash at the end of the experiment.

At the beginning of the experiment you will be randomly divided in four groups of five participants. The composition of the groups will remain fixed throughout the experiment; therefore, each of you will interact with the same four people for the whole duration of the experiment and will never interact with people from other groups. In this experiment you will participate in a 'market'.

The rules of the market are the following:

the market will last for a total of 100 periods. In each period of the market, you will have to choose, independently from the others and without possibility to communicate, only and only one among two actions named A and B. After all people in your group have made their choice, the computer will calculate your payoffs as follows:

- those of you who chose the action that has been chosen by the minority of persons in his/her group will earn one 'experimental' point.
- those of you who chose the action chosen by the majority will earn nothing.

For example, let's assume that in a given round two people have chosen A and three people have chosen B. The two people choosing A will earn one point each, while the three B choosers will earn nothing. Or, assume that one person in the group chooses A and the other four choose B. The A chooser earns one point, while all the others earn nothing.

Since the groups are formed by an odd number of people, ties are excluded. In the case in which all members of the group make the same choice, no one will earn anything.

The experiment will be carried out as follows:

Please, take a look at the screen in front of you. At the beginning of each period, each one of you will make his/her choice via computer by simply clicking with the mouse on his/her preferred choice (A or B). After you have chosen, please wait in silence for other people in your group to make their choices. After all five members in the group have made their choices, the computer will count the number of A's and B's and will automatically calculate everybody's earnings.

Information:

At the end of each round, the screen will show some information relative to that round. Each of you can see on his/her screen a table with five columns. The first column on the left (round)

will report the round number. The second column (choice) will report the choice that you made in that round (A or B). The third column (earnings) will report your earnings in that round. The fourth column reports the ‘winning’ choice for that round, i.e., the action that was chosen by the group minority. Finally, the fifth column will report the percentage of your ‘profit-making’ rounds up to that point, i.e., the percentage of the rounds in which you have earned one point. This information will be visible on your screen only for the last four rounds of the ‘market’.

Warnings:

You will not be allowed to make your choice for the next round before the current round has been completed. When information will appear on your computer, you will be allowed to make your choice for the next round.

During the experiment you will not be allowed to use paper and pencil.

It is very important that you do the experiment in silence and that you don’t try to look at other participants’ choices. If the experimenter should notice any form of improper behavior, he will interrupt the session and no one will be paid. If during the experiment you have questions, please raise your hand.

References

- [1] Arthur, W.B., (1994), Inductive reasoning and bounded rationality, *Am. Econ. Review Papers and Proc.*, 84, 406-412.
- [2] Budescu, D.V., A Markov model for generation of binary random series, *Journal of Experimental Psychology: Human Perception and Performance*, 1987, 13, 25-39.
- [3] Challet, D. and Y.-C. Zhang, (1997), Emergence of cooperation and organization in an evolutionary game, *Physica A*, 226.
- [4] Challet, D. and Y.-C. Zhang, (1998), On the minority Game: Analytical and Numerical Studies, *Physica A*, 256.
- [5] Duffy, J. and Hopkins, E. (2001), *Learning, Information and Sorting in Market Entry Games: Theory and Evidence*, 01/01 University of Edinburgh discussion paper.
- [6] Erev, I., Rapoport, A., (1998), Coordination, “magic”, and reinforcement learning in a market entry game, *Games and Economic Behavior*, 23, 146-175.
- [7] Erev, I. and A. Roth, (1998), Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria, *Am. Econ. Review*, 88, 848-881.
- [8] Kahneman, D., Slovic, P., and Tversky, A. (Eds.) (1982), *Judgment under uncertainty: Heuristics and biases*, Cambridge: Cambridge University Press.
- [9] Mayer, D.J., Huyck, J., Battalio, R. and Soring T. (1992), Histories role in coordinating decentralized allocation decision: laboratory evidence on repeated binary allocation games, *Journal of Political Economy*, 100, 292-316.

- [10] Miller, G. (1956). The magical number seven, plus or minus two. Some limits of our capacity for processing information, *Psychological Review* **63**, 81-97
- [11] Ochs, J. (1990) The coordination problem in decentralized markets: An experiment, *Quarterly Journal of Economics*, 105, 545-559.
- [12] Ochs, J. (1995) Coordination problems, in J.H.Kagel and A.E. Roth (Eds.) *Handbook of experimental economics*, Princeton: princeton University Press.
- [13] Rapoport, A., and Boebel, R. B. (1992), Mixed strategies in strictly competitive games: A further test of the Minimax hypothesis, *Games and Economic Behavior*, 4, 261-283.
- [14] Rapoport, A. and Budescu, D. V. (1992), Generation of random series by humans: A two-person zerosum game paradigm, *Journal of Experimental Psychology: General*, 121, 352-363.
- [15] Rapoport, A., Seale, D.A., Erev, I., Sundali, J.A., (1998), Equilibrium play in large group market entry games, *Management Science*, 44, 119-141.
- [16] Schelling, T.C., (1978) , *Micromotives and Macrobehavior*, New York: W.W. Norton and Company.
- [17] Shannon, C.E. 1948, "A Mathematical Theory of Comunication", *The Bell System Technical Journal*, Vol.27 pp.329-423
- [18] Sundali, J.A., Rapoport, A., Seale, D.A., (1995), Coordination in market entry games with symmetric players, *Organizational Behavior and Human Decision Processes*, 64, 203-218.
- [19] Van Huyck, J., Battalio, R., and Beil, R. (1990). "Tacit Cooperation Games, Strategic Uncertainty, and Coordination Failure". *American Economic Review*, 80, 234 – 248.
- [20] von Mises, R. 1964, "Mathematical Theory of Probability and Statistics", New York, Academic Press
- [21] Zwick, R., and Rapoport, A. (1999), Tacit Coordination in a Decentralized Market Entry Game with Fixed Capacity, Hong Kong University of Science and Technology, Working Paper.




Figure 1: Average informational efficiency for $m = 1, 4, 16$ computed with different values of l in the aggregate information case.




Figure 2: Average informational efficiency for $m = 1, 4, 16$ computed with different values of l in the full information case.




Figure 3: Binned frequency for $p(0|a; b)$. (a) When the last choice was 0 and the winning side 0. (b) When the last choice was 1 and the winning side 0. (c) When the last choice was 0 and the winning side 1. (d) When the last choice was 1 and the winning side 1. The persistence is made stronger by a “reinforcement” effect




Figure 4: Average allocative efficiency Σ (as defined in Eq. 2) computed on Monte Carlo with $\sigma_w = .13$ and $\sigma_l = .07$ for various values of μ_w and μ_l




Figure 5: Average informative efficiency I (as defined in Eq. 3) computed on Monte Carlo with $\sigma_w = .13$ and $\sigma_l = .07$ for various values of μ_w and μ_l