



Laboratory of Economics and Management

Sant'Anna School of Advanced Studies

Piazza Martiri della Libertà, 33 - 56127 PISA (Italy)

Tel. +39-050-883-343 Fax +39-050-883-344

Email: lem@sssup.it Web Page: <http://www.lem.sssup.it/>

LEM

Working Paper Series

The Minority Game Unpacked: Coordination and Competition in a Team-based Experiment

Giovanna Devetag ^{*}
Francesca Pancotto ^o
Thomas Brenner [§]

^{*} University of Perugia, Italy

^o University of Bologna, Italy

[§] Philipps University Marburg, Germany

2011/18

August 2011

ISSN (online) 2284-0400

The Minority Game Unpacked: Coordination and Competition in a Team-based Experiment*

Giovanna Devetag[†]

Francesca Pancotto[‡]

Thomas Brenner[§]

July 28, 2011

Abstract

In minority games, players in a group must decide at each round which of two available options to choose, knowing that only subjects who picked the minority option obtain a positive reward. Previous experiments on the minority and similar congestion games have shown that players interacting repeatedly are remarkably able to coordinate efficiently, despite not conforming to Nash equilibrium behavior. We conduct an experiment on a minority-of-three game in which each player is a team composed by three subjects. Each team can freely discuss its strategies in the game and decisions must be made via a majority rule. Team discussions are recorded and their content analyzed to detect evidence of strategy co-evolution among teams

*We 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 Italian Ministry of Education, University and Research (PRIN 2005) is gratefully acknowledged. The usual disclaimer applies.

[†]Department of Law and Management, University of Perugia, Italy. Corresponding author: devetag@unipg.it

[‡]Department of Economics, University of Bologna, Italy.

[§]Philipps University Marburg, Germany

playing together. Our main results of team discussion analysis show no evidence supporting the mixed strategy Nash equilibrium solution, and support a low-rationality, backward-looking approach to model behavior in the game, more consistent with reinforcement learning models than with belief-based models. Showing level-2 rationality (i.e., reasoning about others' beliefs) is positively and significantly correlated with higher than average earnings in the game, showing that a mildly sophisticated approach pays off. In addition, teams that are more successful tend to become more egocentric over time, paying more attention to their own past successes than to the behavior of other teams. Finally, we find evidence of mutual adaptation over time, as teams that are more strategic (i.e., they pay more attention to other teams' moves) induce competing teams to be more egocentric instead. Our results contribute to the understanding of coordination dynamics resting on heterogeneity and co-evolution of decision rules rather than on conformity to equilibrium behavior. In addition, they provide support at the decision process level to the validity of modeling behavior using low-rationality reinforcement learning models.

JEL codes: C72, C91, C92

Keywords: coordination, minority game, market efficiency, information, self-organization, reinforcement learning

1 Introduction

Congestion problems are at the core of many economic and social phenomena. A congestion problem arises whenever a limited resource is overexploited by a population, while at the same time other similar resources are underexploited due to a lack of coordination among users. Route choice problems belong to this category: if two alternative routes can be used to move from A to B, the average driving time will depend, among other things, on how many drivers choose each of the two options; if there is lack of coordination, one road may be congested and the other empty. Likewise, if a new market opportunity emerges, no firm will be able to profit from it if too many firms choose to enter the business. Common to these examples is the fact that the payoff accruing to each individual depends on the choices of other individuals.

Congestion problems have recently received considerable attention from economists and other social scientists, who have attempted to investigate them through the formal apparatus of game theory, treating the coordination problem as an underlying problem of selection of one out of many possible equilibria. Among the experimental studies are those dealing with market entry games (Kahneman 1988), route choice games (Selten et al., 2002), and minority games (Bottazzi and Devetag, 2007).

The minority game (Challet and Zhang, 1997, 1998) is probably the most stylized and barebone example of a congestion problem: in the game, a population of N agents (where N is an odd number) must independently choose one of two possible actions or sides, say, A or B. A positive and fixed payoff is assigned only to those agents who picked the minoritarian side, while agents on the majority side earn nothing. The game is usually repeated for several rounds, and the interesting question is whether or not a given population converges to a Pareto-optimal equilibrium without the possibility to communicate or to make binding agreements. The game is inspired by the well-known ‘El Farol’ Bar problem formulated originally by Arthur (1989), and it is a stylized representation of all those economic and social situations in which there are advantages in belonging to the minority side. In the absence of any element that may differentiate A from B, and lacking any intrinsic preference

for either option, the actors face the non trivial task of trying to imagine what the majority of other actors will choose, in order to choose the opposite. Coordination, and hence aggregate efficiency, is achieved when players distribute themselves evenly among the two available sides. The coordination problem is compounded by the fact that, unlike other coordination games such as the stag hunt, all Nash equilibria in pure strategies are payoff-equivalent, leaving no obvious ‘salient’ solution. Moreover, these equilibria are also non-strict and payoff-asymmetric, hence making them implausible stable outcomes in the long run. Therefore, coordination in the minority game must necessarily rest on some dynamic process of learning and co-evolution of strategies and behavioral rules by which players learn to differentiate.

Experiments on minority games and similar congestion games have shown that groups are generally remarkably good at achieving coordination, although individual behavior does not conform to either the pure strategy or the mixed strategy Nash equilibria. Alternative explanations of behavior that have been advanced take into account reinforcement learning models (Erev and Roth, 1998), or a quantal response equilibrium (Goeree et al., 2005). However, these explanation all imply that players randomize between the available options, which is contradicted by the data. Moreover, the role of additional information available to players seems inconsistent with all existing models of learning (Duffy and Hopkins, 2005).

Several experiments have pointed out that aggregate behavior in congestion games, rather than through standard learning models or equilibrium notions, may better be explained through an ‘ecology’ of heterogeneous rules of behavior: some of these are not contingent on the game outcomes (e. g., choosing always the same option), while others condition the choice in the current round on outcomes in previous rounds (Selten et al., 2007; Bottazzi and Devetag, 2007). These behavioral rules, moreover, seem affected to some extent by the availability of information: for example, a higher occurrence of ‘sorting’ strategies (i.e., rules prescribing to pick always the same side) and quicker convergence to pure strategy equilibria have been observed when more information is available (Bottazzi and Devetag, 2007; Duffy and Hopkins, 2005). Moreover, it has been shown that it is not only heterogeneity of rules that determines aggregate efficiency, but, rather, their

co-evolution over time, by which players modify their behavior in response to the game outcomes and to the behavior of others (Bottazzi and Devetag, 2007).

The evidence collected so far has suggested to redirect research on behavior in congestion games to a systematic study of the different response modes that players develop, ‘to identify patterns of behavior shared by subsets of players...and then try to account for aggregate behavior in terms of the behavior of the clusters of players that form these aggregates’ (Zwick and Rapoport (2002) also cited in Kets (2011)). Identifying clusters of similar behaviors and their dynamic interaction may be realized ideally in the context of a minority game, in which any response mode may be effective, provided that it recommends actions in opposition to the actions of opponents. Hence, it is only the interaction of rules that counts for aggregate efficiency, not their rationality *per se*.

This paper presents a laboratory experiment that attempts to open the ‘black box’ of the minority game and of similar congestion games, by tracking down the main motivations behind the adoption of behavioral rules and strategies, and hence help identifying the most common patterns of behavior and their interaction. In order to do so, we design and implement an experiment in which a minority game is played by teams of 3 players, who are given the opportunity to freely discuss about the most appropriate course of action throughout the game, and must decide on what to do in every round by applying a unanimity or majority rule. The team discussions are video-recorded, and the discussion content analyzed with the method of content analysis (Hennig-Schmidt, 1999).

Our interest is *not* in studying differences between individual and team behavior in the minority game, but only in trying to infer the most common behavioral rules adopted in the game and their main properties. Our use of teams instead of individual players, hence, is only instrumental in capturing the main reasons behind the observed behavior through the analysis of team discussion.

Hence, through content analysis we attempt to answer the following research questions: are choices in the game mainly based on one’s history of past payoffs, or on the behavior of others? In the former case, how long is the history that is considered when choosing? If the latter is the case, are players trying to influence the behavior of others or mainly react

to it? How often are strategies such as ‘sorting’, ‘randomization’, ‘alternation’, or others suggested? Do these tend to be used more or less frequently depending on the behavior of others during the course of the game?

In addition to our main research question, we also want to investigate the effect of information, therefore we run two different treatments: in the Partial information treatment, teams are only informed about the ‘winning’ side at the end of every round, whereas in the Full information treatment, they are informed of each team’s individual choice. We want to test whether the availability of more information radically modifies the main motivations and strategic considerations at the basis of observed behavior. Our results can be used to narrow down the set of possible response modes in congestion games, in order to better describe aggregate behavior in these games as a function of a small set of features such as, e.g., group size, availability of information about the choices of others, length of ‘history’ of past outcomes available, and so forth.

Our results show no evidence supporting the mixed strategy Nash equilibrium solution, and suggest a low-rationality, backward-looking approach to the game more consistent with reinforcement learning models than with belief-based models. Showing level-2 type of rationality (i.e., reasoning about others’ beliefs) is positively and significantly correlated with higher than average earnings in the game, suggesting that a mildly sophisticated approach pays off on average. In addition, teams that are more successful tend to become more self-centered over time, paying more attention to their own past successes than to the behavior of other teams. Finally, we find evidence of mutual adaptation over time, as teams that are more allocentric (i.e., they pay more attention to other teams’ moves) tend, on average, to induce the opposite tendency in competing teams. Our results contribute to the understanding of coordination dynamics resting on heterogeneity and co-evolution of decision rules rather than on conformity to equilibrium behavior. In addition, they confirm, at the decision process level, the validity of modeling behavior using low-rationality reinforcement learning models.

The paper is organized as follows: Section 2 reviews the related literature; Section 3 presents the game, and illustrates the experimental design and implementation; Section

4 presents experimental results on choice behavior, distinguishing between aggregate and individual (team) behavior; Section 5 presents results on content analysis; finally, Section 6 offers some conclusions and possible directions of future research.

2 Literature Review

Congestion games can be defined as participation games with negative feedback: players must decide whether or not to participate in a certain activity, and the payoff they gain decreases in the number of other participants¹. Market entry games belong to this class, and were the first to be investigated through laboratory experiments; in a market entry game, every player faces the option of not entering one or more markets and getting a ‘sure’ payoff (which can be either positive or zero), or entering a market in which payoff decreases (usually linearly) in the number of entrants, up to a certain capacity k , beyond which payoff from entry becomes negative.

The first experiment on a market entry game revealed that groups of players coordinated surprisingly well, the average number of entrants over time being very close to k (Kahneman, 1988). Meyer et al. (1992) and Ochs (1990), on the other hand, pointed out the difficulty of players to converge to pure strategy equilibria when more than one market is introduced. Several subsequent experiments (Rapoport, 1995; Sundali et al., 1995; Erev and Rapoport, 1998; Rapoport et al., 2000) have substantially confirmed that aggregate behavior in the game conforms surprisingly well to the Nash equilibrium solution, and this regularity has proven robust across a variety of different parameters, such as group size, entry costs, payoff from staying out, number of markets among which to choose, and information conditions. Other experiments (Duffy and Hopkins, 2005) show that additional information increases the speed of convergence to the game pure strategy equilibria. However, Nash equilibria often fail to account for the high between-subject variability (Rapoport et al., 1998). In addition, individual behavior analysis reveals that

¹In opposition, participation games with positive feedback are those in which individual payoffs increase in the number of participants, so that the Pareto-optimal equilibrium implies that all players participate.

players rarely conform to the predictions implied by the game Nash equilibrium, whether pure or mixed; rather, the data show a high and persistent degree of heterogeneity between subjects' decision rules combined with a large within-subject variability in the decisions taken across iterations of the same stage game (See Erev and Rapoport (1998); Rapoport et al. (1998)). Finally, while individual behavior cannot usually be accounted for by any model that assumes randomization, simple reinforcement-based decision rules that condition the choice in one round to the payoff obtained in the previous round(s) describes behavior in these games quite accurately (Erev and Rapoport, 1998; Rapoport et al., 1998, 2000). Zwick and Rapoport (2002), however, have shown that no convergence to equilibrium play is observed even at the aggregate level when the payoff function is not linear and when costs of deviations from equilibrium play are substantially increased.

Route choice games are another category of congestion games, in which, unlike market entry games, there is no 'safe' option, and all the alternatives lead to a payoff that depends on how many players have selected them. Usually, however, the pure strategy Nash equilibria of route choice games are payoff-equivalent, and hence more stable than in the minority game. Selten et al. (2002) have run a route choice game in which players had to choose between a main road (M) and a side road (S); individual travel time was inversely related to payoff and depended on how many travelers had chosen each road. In one of two information treatments, players received information only about traveling time in the road chosen, while in the other treatment players had information about traveling time in both roads. The authors find that the mean number of travelers of each road is very close to the equilibrium value, and that providing information about both traveling times significantly reduces (though not to zero) fluctuations around equilibrium. In a related experiment, Selten et al. (2007) have subjects play a route choice game for 200 rounds, again running two information treatments as in the previous experiment: they find out that, while the mean numbers of travelers on both roads tends to be always close to equilibrium values, ample fluctuations are observed that do not tend to diminish with time, despite the high number of periods. Moreover, additional information reduces these fluctuations, but not substantially. The authors also find heterogeneity in behavior, and the emergence of two

main response modes: a direct response mode, by which a road change follows a bad payoff, and a contrary response mode, by which a road change follows a good payoff. Their simulations with agents using the two response modes plus pure strategies (i.e., choosing always the same road no matter what) replicate their experimental findings quite closely.

The third class of congestion games includes minority games. The minority game differs from the other congestion games in its very unstable equilibrium structure: its Nash equilibria in pure strategies, whereby players sort themselves by side chosen, are non-strict and payoff-asymmetric. This property, combined with the absence of any symmetry-breaking feature or similar potential coordination devices, makes the unique mixed strategy equilibrium a ‘natural’ behavioral benchmark. However, randomizing in the game is rational only insofar as all players in the game randomize: otherwise, it may be more rewarding in terms of payoff to try to predict and outsmart other players’ behavior. These attempts may in turn produce efficient coordination at the aggregate level if the different rules of behavior developed somehow ‘cancel out’. Minority games were at first studied extensively within the econophysics literature, through simulations and analytical models². Many of these studies were motivated by the model of inductive reasoning developed in Arthur (1994), and were aimed at identifying the conditions that allow aggregate efficiency to emerge from the interaction of agents endowed with choice heuristics in the form of “if...then” rules (see also Kets (2007), and Kets and Voorneveld (2007), for an account of learning models from the minority game literature in econophysics, and their comparison with standard game-theoretic learning models). Experimental studies on the minority game (Bottazzi and Devetag, 2007) have shown that groups of five players are able to coordinate efficiently, although individual behavior does not conform to equilibrium; in addition, the availability of full information increases the use of pure strategies that prescribe sorting. Finally, the path of play shows evidence of mutual adaptation between players’ strategies over time.

Chmura et al. (2010) test the extent to which subjects conform to the symmetric mixed strategy equilibrium in a minority game, by implementing incentives that allow subjects to

²For 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/>.

pick strategies randomly. Their results show that only 25% of their experimental subjects decide to choose according to a coin toss, while the majority implement a decision rule that can be adequately described by a reinforcement learning model. Belief learning models on the contrary have low predictive success, which is line with the fact that participants mainly collect information about past payoffs. Chmura et al. (2010) conclude that the observed heterogeneity in players’ behavior in games where multiple equivalent equilibria are present must be ascribed to the structure of individuals’ preferences rather than to their inability to calculate mixed equilibria. Finally, Liu et al. (2010) study minority game properties by comparing behavior in the game played by humans and by fish, and discover several behavioral similarities between the two species.

3 The experiment: design and implementation

In a minority game, N agents (with N odd) have to pick one of two actions or sides, say 0 or 1, without communicating. The earnings for picking side i at time t are expressed by the following function, equal for all players:

$$\pi_{i,t} = \begin{cases} 1 & \text{if } n_{i,t} \leq (N - 1)/2 \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

where $i \in \{1,0\}$ and $n_{i,t} \in \{1, \dots, N\}$ is the number of players choosing side i . In our experiment $N = 3$, consequently participants earn 1 if they pick the choice not taken by the other two players, while in the opposite case they earn zero. In our setting, a decisional unit is composed by a three-subjects entity that we call *team*. Consequently, in our experiment, 3 teams of 3 players each played a minority game that lasted for 100 rounds. We denote the three teams expressing decisions in the same experimental session as a *group*. From now on, therefore, we will use the term ‘team’ to indicate the 3-person decisional unit or player, while we will use the term ‘group’ for the entity composed by the three teams playing together. In each round, teams had to choose between two actions that were labeled A and B, and a minority rule determined team payoff. A total of 108 experimental subjects,

divided in 12 sessions, participated in this experiment. Subjects belonging to a session showed up at the specified time and were all accommodated in a room where instructions for the experiment were distributed and read aloud (see Instructions in the Appendix). Each subject had to sign a consent form to allow video-recording of the team discussion. After this part was over, subjects were instructed to exit the room one at a time; once out of the room, each subject had to extract a card from an urn containing 9 cards reporting the names of the three different rooms that were used as locations for the different teams (for each room name, there were 3 cards reporting that name). The subject hence was accompanied to the assigned room. This procedure assured that each subject only knew the composition of his/her own team (which was determined randomly anyway) but did not know the composition of the other two teams. Once all subjects had been accompanied to the assigned rooms, the experiment could start and the video-recording device was turned on. There was only one computer and one video-recording device in each room. Choices in the game were computerized (Fig. 3 shows a sample of the graphical interface used), but the software did not allow any form of inter-team communication other than through the choices made. One laboratory staff member (not involved with the experiment) was always present in each room in case some unexpected event occurred (e.g., hardware failure).

In order to test the effect of the different amount of information available to subjects, we conducted two treatments: in the *Partial* information treatment, teams received information only about which choice had been the “winning” choice after each round; in the *Full* information treatment teams were also informed of the individual choices of all three teams after each round. This information remained visible on the computer screen for the entire game duration, so that each team had information about the entire game history up to the current round.

Previous experiments have shown that the presence of a full information treatment has an effect on individual strategies; in particular, with full information more individuals tend to follow pure strategies over time (i.e., choosing always the same action), and this behavior has been shown to have, on average, an efficiency-enhancing effect (Bottazzi and Devetag, 2007). We choose to implement the game for 100 rounds to allow teams to learn and

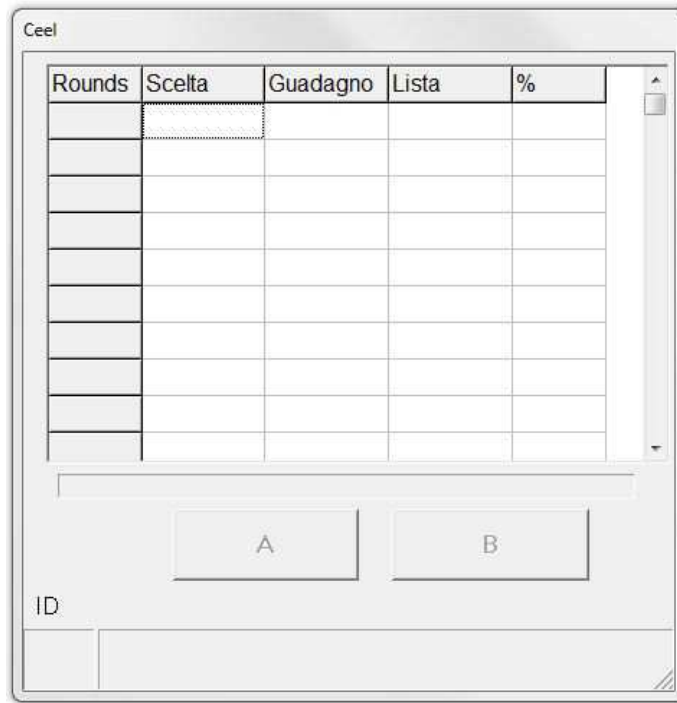


Figure 1: Experimental software screen.

change their strategies in response to the behavior of other teams. Teams were instructed to make their choices after free discussion with no time limits, and to try whenever possible to reach unanimity; in case of disagreement, the majority rule had to be applied (for this reason we fixed team size equal to 3). In each round the “minority” team was assigned one Experimental Currency Unit (ECU), which was worth 0.75 euros. Hence, each team could earn a maximum of 75 euros in the experiment, which had to be divided equally among the team members. In addition, each subject was paid a show up fee of 3 euros. All payments were made privately and in cash at the end of the experiment.

The experiment lasted two hours on average, with high variability between groups and with a longer average time observed in the full information treatment. Average earnings were equal to 18 euros.

4 Analysis of Choices

We first analyze choice data, separately for group and individual (team) behavior. We then turn our attention to the analysis of the team discussion content.

4.1 Allocative Efficiency

In this and the following sections we follow closely the analyses reported in Bottazzi and Devetag (2007). We start by analyzing allocative efficiency, which we define simply as group payoff, π_{group} .

Group payoff at time t can be defined as

$$\pi_{group} = \begin{cases} n_A(t) & \text{if } n_A(t) < N/2 \\ N - n_A(t) & \text{if } n_A(t) > N/2. \end{cases} \quad (2)$$

where $n_A(t)$ is the number of players (teams) choosing action A at time t and N is the total number of players (teams) in the game. In our case, as $N = 3$, the per-round group payoff can be either 0 or 1.

The benchmark that we use for the parameter π_{group} corresponds to the value of allocative efficiency in the case of players who behave according to the symmetric mixed strategy Nash equilibrium. In this case, the theoretical probability that exactly n players choose action A, is the expected value calculated over a binomial distribution with parameters N and n (as in Bottazzi and Devetag (2007)). In our case, $N = 3$ and $n \in \{0, \dots, N\}$, thus the benchmark value, $\bar{\pi}_{group}$ is 0.75. The value of allocative efficiency corresponding to any of the game pure strategy Nash equilibria would equal 1. We investigate both overall allocative efficiency and its variation over time.

Table 1 reports the values of allocative efficiency for each group, with the top (bottom) half reporting values for the partial (full) information treatment. First, values of allocative efficiency averaged across successive time intervals of 25 rounds are reported; the last row reports the value computed over the entire game duration.

Looking at values in the last rows only, it is noteworthy that only two groups (group

2 and 12) exhibit a level of allocative efficiency below the theoretical benchmark of 0.75. All remaining values are higher, with no significant differences between the two treatments ($p=.29$, Mann Whitney U test, two-tailed), although a higher average characterizes the Full information treatment.

Looking at remaining values, we observe an increase in allocative efficiency over time. The phenomenon is more evident in the Full Information treatment. Except for group 6, values in the last 25 rounds are always higher (and in some cases remarkably so) than values in the first 25 rounds, showing evidence of learning and - possibly - mutual adaptation between team strategies.

Table 1 about here

In order to test whether the team condition made a difference in the aggregate with respect to the individual condition, we compare our data with data from the experiment in Bottazzi and Devetag (2007), in which groups of five individuals play a minority game for 100 periods. Their experiment implements two different information treatments as well, full and partial. In order to compare two data sets obtained with different group sizes, we divide each group allocative efficiency value for the corresponding theoretical benchmark (equal to 0.75 in the group condition, and to 1.5625 in the individual condition); the normalized data are then compared by using a Mann-Whitney U test, separately for the two information treatments, so as to compare data obtained with different group size but identical information conditions: we observe no significant difference in allocative efficiency, both for the aggregate information treatment ($p\text{-value} = .206$, two-tailed, Mann Whitney U test) and for the full information treatment ($p\text{-value} = .061$, two-tailed, Mann Whitney U test). Hence, the team condition did not determine a difference in terms of aggregate efficiency with respect to the individual condition; this result is important as it legitimates extending findings from the team discussion analysis to explain behavior in minority games with individual players. Concluding this section, aggregate results show high allocative efficiency on average, and evidence of learning over time. Like in the minority game

with individual players, efficiency is higher in the full information treatment, although not significantly so.

4.2 Individual (team) Behavior

In this section we analyze team choice behavior in the game. A binomial test performed on choices pooled across treatments failed to reject the null hypothesis of no association between action label and choice³. Hence, we collapse our data across the label variable (i.e., A or B).

We then computed the number of times that teams repeated their previous round choice conditional on having been in the ‘minority’ or ‘majority’ side, i.e., on their previous round payoff having been 1 or 0. Table 2 reports the results.

The findings indicate a tendency to repeat previous round choices more often than what would be implied by the mixed strategy equilibrium solution for both treatments, in line with previous results (Bottazzi and Devetag, 2007).

Table 2 about here

Finally, Table 3 reports the relative frequency of choices of action B in successive and separate intervals of 25 rounds for each individual team. Values close to 0 or 1 signal behavior approaching a pure strategy (i.e., always choosing the same action). Previous experiments (Bottazzi and Devetag, 2007) have revealed that, generally, the tendency to follow pure strategies on the part of players increases over time, with a marked difference depending on the information available: when players have full information about individual choices, they tend to resort more often to pure strategies, and this behavior is, on average, efficiency-enhancing. Pure strategies emerge from our data as well (e.g., in group 1, 3, 4, 8, 9 10 and 11), although no significant difference in its frequency of use emerges as a function of information availability.

Table 3 about here

³Choices of A are 1811 over 3600, that is 50.3% of all choices (binomial test, $p = .5$.)

Summing up the results from our choice behavior analysis, the following facts are worth noticing: allocative efficiency is on average very high in all groups, confirming previous results and suggesting that coordination in the game is very successful with ‘small’ (i.e., $N=3,5$) groups. There is evidence of learning, since coordination improves over time, and this phenomenon appears more markedly in the full information treatment. Unlike previous experiments, there is no evidence of an increased frequency of pure strategies when full information is available. We now turn to the analysis of the team discussion content.

5 Analysis of Strategies

In this section we present a content analysis of team discussions throughout the game. We aim at understanding why certain choices are made, which strategies emerge more frequently, and if and how people react to (or anticipate) the behavior of others. To this end, we conduct a number of analyses. First, we examine what types of strategic considerations are most common, and how these differ between treatments (Section 5.1). Then, we turn to the dynamics of various arguments and examine whether any change in the team choice process over time can be identified (Section 5.2). Third, we study whether and to what extent both team payoff and group allocative efficiency depend on the strategic considerations of the teams involved and on the types of arguments most often brought up (Section 5.3). Fourth, we examine whether and how the strategic considerations and choices of one team influence the considerations and subsequent choices of the other interacting teams (Section 5.4).

We first describe the method employed in detail. All team discussions in the experiment were video-recorded and then transcribed in separate word files (there are a total of 36 video-recordings and transcript files), by reporting each complete or incomplete sentence in a separate line, numbered progressively, and by identifying the speaker identity with a numeric code (from 1 to 3). We were only interested in the spoken words; the video recording was only used to allow unambiguous identification of the speakers in case of partial voice overlapping. Then, two sets of categories were created (see below for a description of the

categories used) and listed in a coding manual that informed the subsequent classification task: each transcript was classified by two independent classifiers who were unaware of the experiment purposes, and were involved neither in the design of the experiment nor in the subsequent data analysis. Each sentence in the transcript had to be assigned to one and only one of two sets of distinct categories, to be considered in a predetermined sequence. We used two different and independent raters to control for classification quality. In the literature the Cohens Kappa coefficient is normally employed as a measure of interrater-reliability (Cohen, 1960). Usually the classifications are said to be reliable if the Cohens Kappa coefficient is greater than or equal to 0.7. The values of Cohens Kappa for our two classifications are equal to 0.699 and 0.68. Although this value is not entirely satisfactory, we believe it is sufficient for our purposes. We restrict our analysis to those classifications in which both raters agree.

Each sentence in the transcripts was classified twice: the first classification concerned the type of argument, if any, that was contained in the sentence. We use the term ‘argument’ in a very general and comprehensive sense, i.e., as a statement or a set of statements intended to provide support, justification or evidence in favor of a proposal or proposition. Arguments in our case may include statements that support a specific action in the game, statements that point at evidence that may serve as reason or justification for a specific course of action, or simply descriptive statements that may be used within the team to support or discourage a course of action (e.g., statements that describe the behavior of other teams or that try to infer the use of specific strategies on the part of other teams, or that refer to aspects of the game structure without explicit reference to actions and players).

Table 4 reports the main categories that were employed for Classification I. Each higher-level category was subsequently subdivided in a certain number of sub-categories. The categories and sub-categories were created with the intentions to be as comprehensive as possible and to minimize ambiguity in the classification task. The categories W and U are built for all arguments with no content, either because they are incomplete or not understandable (U) or because a statement concerns a choice suggestion without providing any

explanation (W). Categories U and W are not analyzed further as they are not informative.

Table 4 about here

Classification II concerns the type of action choice that was suggested or implied in the statement. Typical examples are the suggestion to choose action A or B. However several statements may concern more complex action choices, such as, for example, an alternating sequence of actions or the suggestion to randomize. If statements do not contain any suggestion on the action to choose, they are categorized in the second classification as belonging to class N. This class is not further investigated. A list of all categories used in the second classification is provided in Table 5. All statements unambiguously unrelated to the game were excluded from the classification.

Table 5 about here

The categories used to classify the content of the discussion reflect the main types of motives supporting the choice of a strategy in the minority game. Some of them derive from results of previous experiments (e.g., Bottazzi and Devetag (2007)), others from game-theoretic notions that refer to the pure and mixed strategy Nash equilibria of the game (for example, the strategy ‘randomize’ refers to the prescription on behavior implied by the unique symmetric mixed strategy Nash equilibrium of the game). *Irrational* motivations, such as those based on intuition, were also included.

5.1 Frequency of categories

We first present an overview of the most frequent categories from Classification I and Classification II, separately for the two treatments. Let us begin with Classification I, which concerns the motivations behind choices and strategies. 66% of all arguments contained a proposal of a specific action without providing a reason in support of that choice (Category W). Statements that could not be classified (Category U) are very rare, being only 1% of the total. Among the remaining categories, O (‘Others’), S (‘Self’), G(‘General’) and

P('Patterns') are the most frequent. Their shares in the two treatments are listed in Table 6.

Table 6 about here

A clear observation from the last column of Table 6 is that arguments of type O ('Others') and S ('Self') are by far the most frequent, accounting for 53.5% and 32.5% of all arguments, respectively. Hence, players spend most of the time discussing about others' behavior and about their own history of actions and/or payoffs, with the first class of arguments being more prominent. Roughly 10% of the arguments are about identifying patterns in the series of winning sides (P), and only 4% are devoted to a general understanding of the game being played (G).

Another observation concerns differences between treatments: the two distributions are significantly different (Chi-square test: $p = .000$) whereby in the Partial Information treatment 47% of arguments concern the team's own history of actions and payoffs, and only 40% concern the behavior of other players, while in the Full Information treatment arguments about others' behavior raise to 67% and those concerning the team's own history amount only to 18% of the total. This difference is highly significant⁴. This finding is noteworthy for two reasons: first, it shows that in the absence of detailed information about each team's individual choices, an egocentric rather than an allocentric approach prevails, providing support for all those behavior modeling exercises that neglect beliefs about others, like reinforcement learning models. Second, it shows that when more information about others is made available, teams use it extensively, changing markedly the motivations at the core of choices of actions, which, accordingly, become less egocentric and more strategic⁵.

Table 7 reports the most frequent lower-level categories, divided by treatment. The distributions in the partial and full information treatment are again significantly different (Chi-square test: $p = .000$). Players most frequently try to predict others' choices in the

⁴ $p = .004$, Mann-Whitney U test, two-tailed, 6 observations per treatment

⁵By 'strategic' we mean more focused on others.

next round (Op with 35.5%), try to infer predictable patterns in their choice sequences (Sp with 26%), and generally look at other teams' past choices (Ol, 19.5%)⁶. They also focus on their own unsuccessful choices (Sn, 14.5%).

It is noteworthy that only 4% of the statements concerns beliefs of other players, a sophisticated reasoning corresponding to level-2 types in models describing players by the number of iterated reasoning levels they can perform (Camerer et al., 2004; Costa-Gomes et al., 2001). Previous experimental evidence on iteratively dominance-solvable games have shown that individuals do not usually perform more than 2, 3 steps of iterated reasoning, with individual differences being partly explained by cognitive constraints (Devetag and Warglien, 2003). Our finding suggests that previous evidence may have overestimated the extent to which players perform reasoning steps. It is noteworthy, however, that the percentage passes from 2% in the partial information treatment to 6% in the full information treatment. Hence, full information increases the tendency to consider other players' beliefs when choosing. Finally, only 2% of the statements attempts at influencing other teams' behavior, and 3% concerns possible reactions of others to one's own behavior. Regarding differences between treatments, the table shows that all subclasses concerning others' behavior increase markedly in the presence of full information, while those related to the team own previous history decrease substantially.

Table 7 about here

Hence, we find that the vast majority of relevant arguments in the game is based either on one's own history of past successes/failures, which is consistent with reinforcement learning models and, more generally, with low-rationality, backward-looking decision rules, or on other players' behavior. In the absence of detailed feedback about others, arguments about self are the most prominent. In addition, when players do reason about others, they do not show sophisticated forms of iterated reasoning, but mostly they focus on other's observed choices.

⁶We use here the term 'player' and 'team' interchangeably, to refer to the game decisional entity.

The fact that people do rarely develop sophisticated strategies in the minority game is further supported by our findings regarding Classification II, concerning choice suggestions. Statements that do not contain a clear proposal on what to do are classified as N and not analyzed further.

Except for class N, the most frequent categories are A, B, SQ, C and F. Table 8 lists the shares of statements that fall into these categories.

Table 8 about here

Categories A and B are by far the most frequent, accounting for about 80% of choice suggestions altogether. Summing the shares of categories C and F to this value, it can be noted that nearly 90% of statements concern choice suggestions that involve *only* the successive round of play. There are significant differences between treatments here as well (Chi-square test: $p=.000$). The category SQ ('Sequence'), which contains statements that suggest more complex (i.e., multiple round) strategies, accounts only for 9% of the statements. Interestingly, the frequency of SQ statements in the partial information treatment is more than double that of SQ statements in the full information treatment, and this difference is highly significant (Chi-square test: $p=.000$). This finding suggests that when feedback about others is limited, players tend to resort more to strategies that do not condition team behavior on the behavior of other teams⁷. It is noteworthy that Category R ('Randomize'), which corresponds to the prescription implied by the unique symmetric mixed strategy Nash equilibrium contains only the 0.6% of the statements, with no significant differences between treatments ($p = .11$).

Hence, the mixed strategy equilibrium solution is unsupported by the results of our analysis: this finding confirms the results in Chmura et al. (2010) obtained by considering choice data. We can hence state that behaviors incompatible with the mixed strategy

⁷Note that with a group size equal to 3, the information about one's own payoff in a round is sufficient to determine the entire choice distribution in the group in that round. The difference between the two information treatments lies in the possibility to know which of the other two teams made a certain choice, hence to keep track of each team history throughout the game.

equilibrium do not derive from players' imperfect ability to randomize, but from players not intentionally pursuing this type of strategy.

F('Fixed') and C ('Change') prescribe 'stay with previous round choice' and 'change with respect to previous round choice', respectively. The share of F is higher in the full information treatment, and the difference is significant (Chi-square test: $p = .002$). The same holds for the share of C (Chi-square test: $p = .005$). Observing the course of action of other teams helps players in the full information treatment to position themselves, at least in some runs of the game, on a pure Nash equilibrium and to lock one of the other two players in a losing position.

Class SQ ('Sequences') is the category that contains suggestions on sequences of choices that extend for more than one round: the main subclasses of SQ and their shares are listed in Table 9.

Table 9 about here

We find that the most frequent action suggestion is that of sticking to the same side (SQf-fixed, SQa-always A and SQb-always B), which accounts for the 63% of suggestions. The proposal to alternate between A and B account for 20% (SQs), while other sequences count for the remaining 17% (SQo).

Interestingly, there are significant differences between the treatments (Chi-square test: $p = .001$). More specifically, suggestions to alternate choices (SQs) are more frequent in the partial information treatment (Chi-square test: $p = .0001$; one-tailed, Mann-Whitney U test: $p = .002$). Categories SQf, SQa and SQb (which can be summed up as they refer to the same strategy of sticking with the same choice) instead are more frequent in the full information treatment, although the difference is not statistically significant (Chi-square test: $p = .19$). The latter subcategories may be thought as supporting a Nash equilibrium in pure strategies which prescribes sorting, i.e., players differentiating on the basis of the action chosen. The fact that these suggestions appear more frequently in the presence of full information is compatible with data on higher-level classes (i.e., the difference between the shares of class F above), and with data on choice behavior from this experiment as well

as from previous experiments with individual players (Bottazzi and Devetag, 2007; Duffy and Hopkins, 2005): in both cases the emergence of sorting is more frequent when full information is available. This is confirmed by our discussion on content analysis, although it has to be noted that class SQ is less frequent in the full information treatment. This latter finding shows that when full information is available, players tend to be more ‘myopic’, not planning courses of action that extend beyond the current round.

Category D is very rare and no statement was classified in category X. Finally, almost no suggestions belonged to Category Sc, containing ‘cooperative’ strategy suggestions, such as the proposal to alternate with other teams in being on the winning side. This collusive strategy is not trivial to implement in a game like the minority game, but the full informational feedback renders it in principle more easily applicable. However, despite the presence of some such suggestions in the full information treatment, their frequency (0.1%) is negligible.

5.2 Evolution of arguments and actions over time

In order to analyze the evolution of arguments and strategies over time, we subsequently partitioned the whole sequence of statements within each team into five parts each containing an equal number of statements. This temporal division does not necessarily mirror an equal division of the rounds of play, since the discussion intensity throughout the game may not have been uniformly distributed. Again we only consider the most frequent categories of both classifications. We report relative frequencies of main classes and subclasses of Classification I and Classification II in the first and fifth part of the game in Table 10. We also report the results of a nonparametric Wilcoxon signed rank test that has been used to detect significant differences.

Table 10 about here

As far as the first classification is concerned, the following facts emerge: first, in both treatments a significant increase in the share of S arguments can be detected, together

with an equally significant decrease in the share of O arguments. Hence, regardless of the information available, over time teams seem to become more and more self-centered. Second, the effect of full information seems strongest at the beginning of the game, given that in the first part the share of O arguments is almost 80%, compared to 46% of the partial information treatment. The remaining two classes, G and P, do not show significant time-dependent patterns. The data on subclasses reveal that the argument mostly responsible for the decrease in the O class over time is Op (‘arguments about other teams’ behavior in the next round or other teams’ behavior or strategy without reference to a period’). However, a difference emerges between the two treatments: while in the partial information treatment the share of Op decreases significantly, in the Full Information treatment that of Op remains invariant. A significant shift happens in the S subclasses for the Full Treatment: the share of Sp almost doubles from beginning to end, whereas that of Sn decreases substantially. These results can be interpreted by taking into account the role of information: in the partial information treatment, teams’ attention toward trying to spot and predict other teams’ strategies (Op) declines because of the inherent difficulty of doing so given the absence of individual information. There is a slight increase in arguments concerning others’ behavior in the previous round, but it is not significant. In the full information treatment, on the contrary, attention to others’ behavior remains high (Op), but as the game progresses and the history of past choices and outcomes becomes longer and richer, teams start paying more and more attention to their own past behavior (Sp increases) to try to detect successful patterns. In the absence of history the attitude is less backward-looking and more strategic. One could say that the longer the past, the more it influences the future. The data from the second classification reveal no significant changes in the aggregate (last column of the table). However, in the full information treatment the share of SQ (‘Sequences that involve more than one round’) increases from 2% to 10%, indicating that the combination of full information and a longer available history of play gives rise to the elaboration of more sophisticated, multiple-round strategies, probably based on those patterns that the team identifies as having been successful in the past.

5.3 Arguments and efficiency

In this section we investigate whether the performance of teams and groups depends on the types of arguments brought up in the discussion. To this end, we use OLS regression models with team payoff and group allocative efficiency as dependent variables, and the shares of categories of arguments and actions as independent variables. We have 36 observations for team payoff and 12 observations for allocative efficiency. We apply a simple cross section regression. As done before, we only use the most frequent categories. For each dependent variable we perform three regressions. The first includes the categories of Classification I: O, S, G and P. The second includes the most frequently observed subcategories of O and S: Ol, Op, Ob, Sp and Sn. The third takes into account the categories of Classification II: A, B, SQ, C and F.

We do not find any significant dependence of team payoff on the main categories of reasoning considered. However, we observe a highly significant positive dependence of team payoff on the subcategory Ob (see Table 11) meaning that teams obtain a higher payoff if they reason more about other teams' beliefs. Hence, given the large prevalence of egocentric perspectives, the minority of players adopting a sophisticated, belief-based approach on average performs significantly better. This is by far not a trivial result for the minority game and it may not be true for all games, since in the guessing game for example (Nagel, 1995) being smarter than the average (too much) does not pay off.

Table 11 about here

Studying the categories of action, we find significant results only if we use a 90% significance level (see the first two columns of Table 12). However, these slightly significant results are very interesting: teams that frequently support the option of sticking to previous round choice (Class F) obtain higher payoffs, while teams that frequently argue to change action (Class C) obtain lower payoffs.

This result partly holds for group allocative efficiency as well. Table 12 shows a positive and weakly significant dependence of allocative efficiency on the share of statements that propose to repeat the previous period choice.

The above analysis uses statements within teams as independent variables. These statements represent intentions and beliefs but not actual choices of actions. Hence, we repeat the above regressions using the number of actual choice changes from one round to the next as an additional independent variable. The results are reported in the last two columns of Table 12.

Table 12 about here

Indeed, for team performance we find a significant negative dependence on the number of actual choice changes. However, this does only partly eliminate the dependence on team discussion. The frequency of arguments in favor of repeating previous round choice (F) is still positively and significantly correlated with team payoff. Hence, both actual choices and discussions in the team matter for the final outcome. This is confirmed by considering that the number of arguments in favor of changing choice is not significantly correlated with the number of actual choice changes (correlation coefficient: -0.161; p-value: 0.347). In the case of allocative efficiency, all significant relationships disappear if both the frequency of arguments and the frequency of choice changes are considered.

5.4 Interdependence between groups

The aim of our final analysis is to test whether arguments brought up within one team have an influence on the arguments in other teams with which the team interacts. Again we only focus on the most frequent types of arguments. Looking at categories of reasoning, O and S are the most frequent. We test whether the share of these two categories depends on the share in other teams.

Again we separate the arguments into three parts: the initial part with the first fifth of arguments, the end part with the last fifth of arguments, and the middle part with the remaining arguments. We use the share of Categories O and S in the end part as dependent variables.

To test interdependence between teams we use the share of the two categories, O and S, in the other two teams in the middle part; denoted by O_{o2} and S_{o2} , respectively. If

the reasoning in other teams has an influence we should see a reaction in the end part of the game to the reasoning of other teams in the middle part. However, teams might also have fundamental tendencies to use certain kinds of arguments more than others. In order to capture the existence of these basic tendencies, we use the shares of O and S within a team in the initial part as further independent variables, denoted by O1 and S1, respectively. Finally, we include two control variables: a dummy for the full information treatment (TD), and one for the number of wins of the considered team in the middle part (Win2). The latter control variable captures the fact that a team discussion might be influenced, among other things, by that team's previous success. We report the results of the regressions in Table 13.

Table 13 about here

In a further step, we analyze the Subclasses Ol, Op, Ob, Sp and Sn in the same way as we did for the Classes O and S. The results are presented in Table 14.

Table 14 about here

First, we find that there is, indeed, a basic tendency within teams that determines the most frequent statements. This holds especially for arguments of Class O: teams that make a lot of O statements at the beginning also make a lot of O statements at the end of the game (see Table 13). In the case of Class S we also find a high coefficient for the dependence of the share of this kind of statement at the end on the share of the same statement at the beginning. However, this coefficient is not significant. This makes it impossible to conclude whether the above finding also holds for Class S or whether Class S is different from Class O. We tend to opt for the former hypothesis, lacking any reason to consider S arguments inherently different from O arguments.

Besides this persistence effect, we find two types of dependencies on the events in the game. First, the statements at the end of the game depend on the number of wins in the middle part. Table 13 clearly shows that a higher number of wins in the middle part is correlated with more S arguments and less O arguments in the last part. There are

two alternative interpretations for this finding. On the one hand, teams that make more extensive use of S arguments in the middle part might be more successful in this part and keep a high share of S arguments in the last part. However, above we found no correlation between S arguments and team success. On the other hand, success might lead teams to focus more on their own behavior and what has made them successful, while losing teams tend to focus on other players instead. Table 14 shows that success (high winning rates) in the middle part leads to a higher number of Sp statements in the last part of the game.

Hence, the main consequence of success is that players become more egocentric, and look backward to try to identify what made them successful.

Second, we find that statements of Class O are less frequent in the end part if the other teams have frequently used the same type of statements in the middle part. This result is difficult to interpret. Again we refer to the results on the subclasses (presented in Table 14), which show that statements of Class Ob become less frequent in the end part if other teams use statements of Class O. Hence, if the other two teams are more focused on other teams, the team tends to become *less* focused on other teams.

This result is interesting for two, related reasons: first, it shows *evidence of mutual dependence between teams' arguments and motivations, showing in particular that having more strategic partners in the game (whereby strategic means more focused on others) leads to being more egocentric instead*. Second, the *mutual dependence emerges in a situation in which teams have no access to other teams' internal discussions, but only to their choices: nevertheless, teams are able to influence each other's motivations*. The complementarity relation (more O arguments leading to less O arguments on the part of competing teams) may be responsible for the mutual adaptation of choices in the minority game and, ultimately, for the high levels of efficiency observed. It is left to explain why a negative instead of a positive correlation emerges, which could be the object of future research.

6 Conclusions

The main aim of this work is to investigate motivations and intentions laying behind actions and strategies in minority games. In order to do so, we implemented a minority of three game played by teams of three players, where decisions had to be discussed in front of a camera and subsequently analyzed by means of content analysis (Hennig-Schmidt, 1999).

Several results of our analysis are noteworthy: first, a strong similarity emerges between aggregate behavior in minority games played by individual decision makers and those played by teams: as in Bottazzi and Devetag (2007), in our experiments allocative efficiency is on average remarkably high, and mostly higher than the mixed strategy equilibrium benchmark, confirming that coordination in minority games played by humans is very successful with small groups; moreover, coordination tends to improve over time, which suggests that teams learn and co-evolve as the game progresses.

The core of our experiment is the identification of the types of arguments (Classification I) and suggestions for action (Classification II), obtained through Content Analysis (Hennig-Schmidt, 1999) performed on the transcripts of the discussions that took place among team members throughout the experiment. Our results show that participants direct most of their effort and attention to themselves, unless provided with detailed information about the choices of others. Only very rarely teams direct attention at trying to identify patterns in the series of 'winning sides' or attempt a general understanding of the game that would allow them to derive general rules of behavior.

The second result worth noticing, is that the introduction of a full information treatment produces a significant shift to a moderately allocentric and strategic approach to the game, in line with previous results (Bottazzi and Devetag, 2007; Duffy and Hopkins, 2005): In the full treatment in fact participants tend to use more arguments related to others' behavior and less arguments related with their own, showing as expected, that more information about others is actively used in that it reflects in a larger consideration of their moves. Nonetheless, this difference does not imply a higher level of rationality on the part of subjects: as a matter of fact, looking at the lower level classification (Classification I

sub-classes), subjects limit themselves to searching predictable patterns in others' past choices (Op) or in their own past successful choices (Sp), according to a backward-looking, low-rationality approach.

The latter statement is confirmed also by the second classification, where 90 per cent of statements relate to choices that involve only the next round of play, while multiple-round strategies (Sequences, 'SQ') account for at most 9 per cent. It is worth noticing how multiple-round strategies tend to decrease from partial to full treatment. When full information is available, players tend to be more 'myopic', not planning courses of action that extend beyond the current round; alternatively, one could argue that with full information players tend to condition choices in each round to the observation of others' choices in the previous round, whereas the partial information treatment promotes the use of unconditional rules (i.e., rules that do not condition one's behavior on the behavior of others).

The third result refers to the 'Randomization' category, the strategy implied by the unique symmetric mixed strategy Nash equilibrium: such a randomization is only proposed in a negligible percentage of statements which decreases further when full information is introduced, as in Bottazzi and Devetag (2007). This result confirms previous findings - obtained from analyzing choice data in previous minority games experiments (Chmura et al., 2010). Hence, we can conclude that the absence of evidence compatible with the mixed strategy equilibrium does not derive from players' imperfect ability to randomize but from players intentionally not pursuing this type of strategy.

Another important results concerns players' level of sophistication. Only 4% of the statements concerns other teams' beliefs, i.e., is in accordance with level-2 type of rationality as in k -level models. This very low percentage is not compatible with previous evidence in other game types; moreover, regression results show that being more sophisticated than the average pays off in terms of earnings, probably because sophistication is pursued by a small minority of players only.

Other important findings concern the time dimension: first, regardless of the information treatment, teams tend to become more egocentric over time, and more so the more

successful a team has been. The length of one’s history (i.e., the number of past rounds of play) determines the extent to which a team will be focused on its past behavior as opposed to being focused on the behavior of others. This tendency is further reinforced by success; the more successful a team has been, the more it will resort to analysis of its past behavior to find ways to be successful in the future.

The last finding points at the evidence of a temporal co-evolution between arguments developed by teams playing together: more specifically, teams that tend to be more strategic (i.e., more allocentric) tend to produce the opposite attitude in other teams over time. Future research will have to disentangle the motivations behind this finding.

Overall, our results are the first attempt at ‘unpacking’ the ecology of rules of behavior likely to emerge in a dynamic coordination game, by pinning down the principal motivations of behavior at the decision process level. Our approach, in other words, tries to go beyond the implicit ‘as if’ assumption on which most models of learning are forced to rely. We show that, despite the extremely competitive nature of the payoff function and the strong incentives to predict competitors’ behavior, there is very little evidence of behavior compatible with belief-based models and with k learning models assuming sophisticated iterated reasoning capabilities. Rather, all the evidence points at the existence of an ego-centric, low-rationality approach to the game consistent with that implied by reinforcement learning models. The interesting finding is that, despite such a low level of individual rationality, aggregate efficiency levels are remarkably high. Hence, our evidence shows with data from choice processes that the right direction to model behavior in congestion games starts from the interaction of a small set of simple behavioral rules like the ones identified here.

References

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

- Arthur, B. (1994, 2). Inductive reasoning and bounded rationality. *The American Economic Review* 84, 406 – 411.
- Bottazzi, G. and G. Devetag (2007, June). Competition and coordination in experimental minority games. *Journal of Evolutionary Economics* 17(3), 241–275.
- Camerer, C., T. H. Ho, and J. K. Chong (2004, 3). A cognitive hierarchy model of games. *The Quarterly Journal of Economics* 119, 861 – 898.
- Challet, D. and Y. C. Zhang (1997). Emergence of cooperation and organization in an evolutionary game. *Physica A: Statistical and Theoretical Physics* 246(3-4), 407 – 418.
- Challet, D. and Y.-C. Zhang (1998). On the minority game: Analytical and numerical studies. *Physica A: Statistical and Theoretical Physics* 256(3-4), 514 – 532.
- Chmura, T., W. Guth, T. Pitz, and A. Ziegelmeyer (2010). The minority of three-game: An experimental and theoretical analysis. *Jena Economic Research Paper*.
- Cohen, J. (1960, 20). A coefficient of agreement for nominal scales. in: educational and psychological measurement. *Psychological Measurement*, 37 – 46.
- Costa-Gomes, M., V. P. Crawford, and B. Broseta (2001). Cognition and behavior in normal-form games: An experimental study. *Econometrica* 69(5), 1193–1235.
- Devetag, G. and M. Warglien (2003). Games and phone numbers: Do short-term memory bounds affect strategic behavior? *Journal of Economic Psychology* 24(2), 189 – 202.
The Economic Psychology of Herbert A. Simon.
- Duffy, J. and E. Hopkins (2005). Learning, information, and sorting in market entry games: theory and evidence. *Games and Economic Behavior* 51(1), 31 – 62.
- Erev, I. and A. Rapoport (1998). Coordination, "magic," and reinforcement learning in a market entry game. *Games and Economic Behavior* 23, 146– 175.

- Erev, I. and A. E. Roth (1998, September). Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *American Economic Review* 88(4), 848–81.
- Goeree, J., C. Holt, and T. Palfrey (2005). Regular quantal response equilibrium. *Experimental Economics* 8, 347–367. 10.1007/s10683-005-5374-7.
- Hennig-Schmidt, H. (1999). *Bargaining in a Video Experiment - Determinants of Boundedly Rational Behavior*, Volume 467. Lecture Notes in Economics and Mathematical Systems: Springer-Verlag Berlin - Heidelberg - New York.
- Kahneman, D. (1988). *Experimental economics: A psychological perspective*, pp. 11– 18. Bounded Rational Behavior in Experimental Games and Markets. New York: Springer-Verlag.
- Kets, W. (2007). The minority game: An economics perspective. *CentER Discussion Paper, Tilburg University* 53.
- Kets, W. (2011). Learning with fixed rule: The minority game. *Journal of Economic Surveys*, no–no.
- Kets, W. and M. Voorneveld (2007). Congestion, equilibrium and learning: The minority game. *CentER Discussion Paper, Tilburg University* 40.
- Liu, R. T., F. Chung, and S. S. Liaw (2010). Fish play minority game as humans do. *Nature Proceedings*.
- Meyer, D. J., J. B. V. Huyck, R. C. Battalio, and T. R. Saving (1992). History’s role in coordinating decentralized allocation decisions. *Journal of Political Economy* 100, 292 – 316.
- Nagel, R. (1995, December). Unraveling in guessing games: An experimental study. *American Economic Review* 85(5), 1313–26.

- Ochs, J. (1990). The coordination problem in decentralized markets: An experiment. *Quarterly Journal of Economics* 105, 545 – 559.
- Rapoport, A. (1995). Individual strategies in a market-entry game. group decision and negotiation. *Group Decision and Negotiation* 4(1), 117 – 133.
- Rapoport, A., D. A. Seale, I. Erev, and J. A. Sundali (1998). Equilibrium play in large groups market entry games. *Management Science* 44(1), 119 – 141.
- Rapoport, A., D. A. Seale, and E. Winter (2000). An experimental study of coordination and learning in iterated two-market entry games. *Economic Theory* 16(1), 661 – 687.
- Selten, R., T. Chmura, T. Pitz, S. Kube, and M. Schreckenberg (2007). Commuters route choice behaviour. *Games and Economic Behavior* 58(2), 394 – 406.
- Selten, R., M. Schreckenberg, T. Pitz, T. Chmura, and S. Kube (2002, Nov). Experiments and simulations on day-to-day route choice-behaviour. (35).
- Sundali, J. A., A. Rapoport, and D. A. Seale (1995). Coordination in market entry games with symmetric players. *Organizational Behavior and Human Decision Processes* 64(2), 203 – 218.
- Zwick, R. and A. Rapoport (2002). Tacit coordination in a decentralized market entry game with fixed capacity. *Experimental Economics* 5, 253–272. 10.1023/A:1020892405622.

Appendix 1: Instructions

The following are an English translations of the instructions given to experimental subjects in the Full information treatment.

You are about to participate in an experiment on interactive decision making. If you follow the rules and make appropriate decisions, you will be able to earn a considerable sum of money that will be paid to you privately and in cash.

At the beginning of the experiment you will be randomly divided into three groups of three participants. You will only know the composition of your group whereas will not know the composition of the other two groups. The three groups, including yours, will be placed in three different rooms and will communicate with each other only through their computer-based choices. Each group discussion and decision process will be recorded by video cameras installed on each computer screen. By signing the consent form that we handed in you accept to be videotaped.

Each group will be able to earn a maximum of 75 euros during the experiment; the group earnings will be divided equally among participants, so each of you can earn a maximum of 25 euros. Earnings in the experiment will be expressed in experimental points. Each experimental point is worth 0.75 euros.

The three groups will participate in an interactive 'market' lasting 100 periods and with the following rules:

in every period each group will have to choose, in isolation from other groups, one and only one out of two possible actions, named A and B.

As soon as all groups have chosen for that period, earnings will be calculated as follows:

- the group that has chosen the MINORITY action, earns 1 experimental point
- the group that has chosen the MAJORITY action, earns 0 experimental points.

For example, suppose that in a round two groups choose A and one group chooses B. The two groups that have chosen A earn nothing, while the group that has chosen B earns 1 point. In other words, the earning rule awards the group who takes the minority side.

Since there is an odd number of groups, ties are automatically avoided. Whenever all three groups choose the same action, no experimental point is awarded.

At the end of the experiment, the total number of experimental points earned by each group during the experiment will be converted in euros, divided by three and the corresponding amount paid to each group component individually, privately and in cash. Moreover, each of you will receive a show-up fee of 3 euros for the participation.

Each decision needs to be discussed within the group. The decision upon the action to select in a round (A or B), must be taken according to either a unanimity rule or an alternative rule that has been unanimously approved. Each member of the group needs to put his/her signature next to each group decision on the operations register that will be provided by the experimenters. Groups will communicate their choices through an interactive software as follows (a screen shot of the software screen is shown below):



In each period, after the group decides which action to choose for that period, a member of the group will transmit the choice to the server by clicking with the mouse on the selected option (A or B). When one group has made its decision, it needs to wait for other groups to decide. After all groups have transmitted their decisions to the server, the server will automatically calculate the points earned by each group and will send the information to your monitors.

Information

At the end of each period and before the next period begins, you will receive information that will remain on your computer screen. The computer screen will show you a table with five columns.

ROUND	CHOICE	EARNING	LIST	%

The first column on the left indicates the round number. The second column reports the choice of your group in that round (A or B). The third column shows your earnings in that round (1 or 0). The fourth column, named LIST, reports the groups' choices (yours included) in that round, in the order corresponding to the group number, which appears as your ID at the bottom left corner of the screen. In this way you can keep track of every group choice in every round. The fifth column, finally, reports the percentage of rounds (up to that point) in which your choice has been a minority choice, thus the percentage of points earned on the total up to that point. When the information for a round will appear, you will be able to make your choices for the following round. During the experiment, you need to report the relevant information in the operations register, which is an Excel spreadsheet with a column in which all group members have to sign every group choice.

- It is strictly forbidden to launch or open any software other than the experimental software.
- It is strictly forbidden to communicate in any way with other groups' members, via cell phones or internet.

When you finish reading these instructions, you will be asked to pick a numbered card out of a cards deck. Participants who pick the same number will be assigned to the same group. Groups will be assigned to their rooms separately.

Are there any questions?

Appendix 2: Tables

Table 1: Allocative Efficiency calculated over successive time intervals of 25 rounds and reported separately per each group and treatment. The top (bottom) panel reports data for the Partial (Full) information treatment. The last row of each panel reports the value of allocative efficiency calculated over the entire 100 rounds of play.

Allocative Efficiency in Time and per group

Partial Info	Group 1	Group 2	Group 3	Group 4	Group 5	Group 6
1-25	0.68	0.68	0.8	0.72	0.64	0.84
26-50	0.76	0.64	0.68	1	0.88	0.8
51-75	0.92	0.8	0.88	1	0.76	0.76
76-100	0.88	0.76	1	0.96	0.88	0.88
1-100	0.81	0.72	0.84	0.92	0.79	0.82

Full Info	Group 7	Group 8	Group 9	Group 10	Group 11	Group 12
1-25	0.68	0.68	0.88	0.68	0.84	0.64
26-50	0.88	0.96	0.96	0.96	0.88	0.56
51-75	0.84	1	0.96	0.96	0.76	0.76
76-100	0.88	1	1	1	0.84	0.92
1-100	0.82	0.91	0.96	0.9	0.88	0.72

Table 2: In this table we report the number of times that teams repeated their previous round choice conditional on having been in the ‘minority’ or ‘majority’ side, i.e. the relative frequency of choices at time t , c_t , equal to the previous period, i.e. $c_t=c_{t-1}$ conditional on the payoff in the previous period (π_{t-1}) being equal to 1, $\pi_{t-1} = 1$ or equal to zero, $\pi_{t-1}=0$.

Partial Info				
Teams	$c_t = c_{t-1}$ <i>if</i> $\pi_{t-1} = 1$	Stand. Dev .	$c_t = c_{t-1}$ <i>if</i> $\pi_{t-1} = 0$	Stand. Dev.
1	0.79	0.09	0.63	0.23
2	0.74	0.14	0.67	0.17
3	0.72	0.21	0.73	0.17
4	0.79	0.14	0.81	0.11
5	0.71	0.18	0.71	0.03
6	0.58	0.18	0.61	0.13
Avg.	0.72	0.15	0.69	0.14
Full Info				
Teams	$c_t = c_{t-1}$ <i>if</i> $\pi_{t-1} = 1$	Stand. Dev .	$c_t = c_{t-1}$ <i>if</i> $\pi_{t-1} = 0$	Stand. Dev.
7	0.69	0.12	0.70	0.12
8	0.78	0.06	0.87	0.03
9	0.83	0.07	0.85	0.04
10	0.88	0.09	0.81	0.04
11	0.80	0.13	0.77	0.10
12	0.56	0.08	0.55	0.07
Avg.	0.76	0.09	0.76	0.07

Table 3: In this table we report the relative frequency of choices of action B in successive and separate intervals of 25 rounds for each individual team. The left (right) panel reports values for the Partial (Full) information treatment.

	Partial Information			Full Information		
	Group 1			Group 7		
1-25	0.56	0.40	0.44	0.36	0.32	0.72
26-50	0.56	0.08	0.48	0.24	0.52	0.68
51-75	0.72	0	0.64	0.36	0.72	0.68
76-100	0.76	0	0.68	0.48	0.56	0.56
	Group 2			Group 8		
1-25	0.44	0.36	0.56	0.44	0.4	0.44
26-50	0.4	0.88	0.76	0.72	0.48	0.16
51-75	0.56	0.68	0.24	0.36	1	0
76-100	0.44	0.44	0.72	0	0.96	0.6
	Group 3			Group 9		
1-25	0.44	0.36	0.48	0.28	0.64	0.44
26-50	0.28	0.64	0.6	0	0.52	0.88
51-75	0.28	0.52	0.88	0.8	0.52	0.24
76-100	0	0.4	1	0.96	0.2	0.2
	Group 4			Group 10		
1-25	0.56	0.40	0.48	0.48	0.72	0.48
26-50	0.48	0.24	0.64	0.2	1	0.6
51-75	0	0.6	0.92	0.04	0.96	0.52
76-100	0.04	0.48	0.92	0	0.96	0.4
	Group 5			Group 11		
1-25	0.16	0.6	0.44	0.64	0.44	0.6
26-50	0.52	0.28	0.64	0.28	0.48	0.52
51-75	0.48	0.4	0.88	0.32	0.8	0.44
76-100	0.36	0.64	0.52	0.48	0.96	0.2
	Group 6			Group 12		
1-25	0.4	0.56	0.56	0.6	0.56	0.56
26-50	0.24	0.76	0.52	0.48	0.6	0.48
51-75	0.52	0.4	0.52	0.44	0.52	0.52
76-100	0.16	0.6	0.6	0.6	0.52	0.52

Table 4: Classification I: types of arguments and relative subcategories with relative codification and a brief description of the content.

Cat	Argument
W-W/out	<p>Suggestions on the action to choose without providing any reason</p> <p>Ww: Without a clear relation to former reasoning</p> <p>Wr: With a clear relation to former reasoning</p> <p>Wc: Statements that show agreement with a previously stated argument</p> <p>Wd: Statements that show disagreement with a previously stated argument w/o providing a reason for disagreement</p>
U-n/und.	Arguments in which the reasoning is incomplete or not understandable
S-Self	<p>Argument concerning a pay-off or success that the team has achieved in the past</p> <p>Sp-self past success: Arguments in which a strategy applied in the past is observed to be successful</p> <p>Sn-self past not succ.: Arguments in which a strategy applied in the past is observed to be unsuccessful</p> <p>So-more(Sm): Other arguments about past successes or failures</p>
O-Others	<p>Arguments that concern other teams</p> <p>Oi-influence: Argument that suggest to try to influence the behavior of others by making a choice</p> <p>Or-react: Statements about how other groups react on one's own team behavior</p> <p>Ol-last round: Statements about the choice of other groups in the last round not used to make a prediction for future behavior</p> <p>Op-next period: Statements on the behavior of other teams in the next period, or on other teams' strategy or behavior without reference to a specific period</p> <p>Ob-beliefs: Statements about the beliefs of other teams about one's own team</p> <p>Ot-theorizing: Theorizing on the reasons why other teams behave in a certain way.</p>
P-Patterns	<p>Argument that concerns the past sequence of 'winning' choices</p> <p>Pp: w/out providing an explanation</p> <p>Pt-theorizing : theorizing about the structure of the sequence of the winning choice</p> <p>Pl-likelihood: Statements about the likelihood of each option to be winning</p> <p>Po: Statements about which team always wins or loses</p>
G-General	<p>Argument in which a general understanding of the game is developed</p> <p>Gs-self: consideration that have the aim to obtain the best results for themselves</p> <p>Gc-coop.: Considerations with the aim to make all the groups similarly well-off</p> <p>Go-more(Gm): Other considerations that reveal a general understanding or interpretation of the game</p>
I-Intuition	<p>Argument based on Intuition, gut feeling, or just guessing</p> <p>Is: Arguments about guesses on what would be a good choice or the winning choice</p>
M-More	<p>Further motives</p> <p>Mb-boredom: Arguments that propose a change of strategy or a choice just because of boredom</p> <p>Mt: Statements that propose a decision to just speed up the game</p> <p>Mi: arguments that propose a choice since the person does not care about the game.</p>

Table 5: Classification II: Categories of the classification concerning the type of action choice that was suggested or implied in the sentence with a brief description.

Cat.	Argument
A-Choose A	Choosing A in this round without any statement about what to do in the next rounds
B-Choose B	Choosing B in this round without any statement about what to do in the next rounds
C-Change	Change choice next round without a statement about the further strategy or behavior
	Ca-Change to A: Change of choice to A. In contrast to A and C, in this case both ‘A’ and ‘changing behavior’ are mentioned in the statement
	Cb-Change to B: Change of choice to B. ‘B’ and ‘changing behavior’ are mentioned
F-Fixed	Stick with previous choice this time without a statement about the further strategy or behavior
SQ-Sequences	Sequences (proposals for more than one round)
	SQs-alternating: Alternating choices
	SQf-fixed: Fixed choices without a statement about which choice should be repeated
	SQa-always A: Fixed choice with always choosing A
	SQb-always B: Fixed choice with always choosing B
	SQc-cooperate: Cooperative strategy that causes each group to win alternatively
	SQo-other sequences(m more?): Other sequences, such as A-B-B-A-B-B-...
R-Randomize	Intentionally randomise
CC-Copycat	Do the same as the other team(s)
D-Dependence	Dependence on history
	Dw-past win: dependence on past winning choice
	Do-past others: dependence on the past decisions of others
	Ds-past self: dependence on the decisions made by themselves in the past
	Dp-past payoffs: dependence on past payoffs, for example ‘we change if we lost’
	Dn-past wins: dependence on the number of times in which one’s own team has won
	Dm-past wins of others: dependence on the number of times in which other teams have won
	Dc-past comparison: dependence on the comparison between number of times in which one’s own team won and the other groups won
X-Null suggestion	Statement of somebody who does not know what and suggests to just make a random decision
N-Not clear	No clear proposal for what ⁴² do either now or in the future

Table 6: Share of statements of Classification I that are categorized into the classes O, S, G and P for the two treatments. Values are expressed in percentages.

Classification I			
Category	Partial Info	Full Info	Total
O - Others	40	67	53.5
S - Self	47	18	32.5
G - General	4.5	4.2	4.35
P - Patterns	8.6	10	9.3

Table 7: Share of statements of Classification I that are categorized into the subclasses Ol, Op, Ob, Sp and Sn for the two treatments. Values are expressed in percentages.

Classification I - Sub-categories			
Category	Partial Info	Full Info	Total
Ol	13	26	19.5
Op	27	44	35.5
Ob	2	6	4
Sp	37	15	26
Sn	21	8	14.5

Table 8: Share of statements of Classification II that are categorized into the classes A, B, SQ, C and F for the two treatments. Values are expressed in percentages.

Classification II			
Category	Partial Info	Full Info	Total
A - choose A	38	39	38.5
B - choose B	40	42	41
SQ - sequences	13	6	9.5
C - change	3.2	5.2	4.2
F - fixed	4.9	7.5	5.25
R - randomize	0.8	0.4	0.6

Table 9: Share of statements of Classification II belonging to category SQ ('Sequences'), detailed for its subclasses Ss, Sf, Sa, Sb and So.

Classification II - sub-categories			
Category	Partial Info	Full Info	Total
Ss	29	9.7	19.35
Sf	8.9	10.7	9.8
Sa	27	30.1	28.5
Sb	23	26	24.5
So	11.6	23.3	17.45

Table 10: Relative frequencies of main classes and subclasses of Classification I and Classification II within the first and the last fifth part of the game. P-values show the significance of changes across time, calculated with the Chi-square test.

Classification I									
	Partial info			Full info			Both		
Cat.	1st	5th	p-val	1st	5th	p-val	1st	5th	p-val
O	0.46	0.29	.04	0.78	0.53	.02	0.62	0.41	.027
S	0.33	0.55	.04	0.13	0.29	.02	0.23	0.42	.028
G	0.09	0.03	.14	0.07	0.09	.6	0.08	0.06	.46
P	0.04	0.07	.89	0.02	0.02	.7	0.03	0.05	.89
Sub. Cat.	1st	5th	p-val	1st	5th	p-val	1st	5th	p-val
Ol	0.18	0.26	.3	0.34	0.32	.2	0.26	0.29	.75
Op	0.56	0.24	.02	0.49	0.46	.7	0.52	0.35	.028
Ob	0.03	0	.6	0.09	0.02	.08	0.06	0.01	.080
Sp	0.48	0.4	.7	0.33	0.64	.04	0.41	0.52	.17
Sn	0.4	0.28	.4	0.31	0.08	.02	0.35	0.18	.075
Classification II									
	Partial info			Full info			Both		
Cat.	1st	5th	p-val	1st	5th	p-val	1st	5th	p-val
A	0.41	0.47	.6	0.4	0.31	.1	0.40	0.39	.75
B	0.35	0.39	.2	0.48	0.38	.4	0.42	0.38	.34
SQ	0.12	0.08	.1	0.02	0.1	.04	0.07	0.09	.17
C	0.03	0.02	.06	0.06	0.06	.8	0.04	0.04	.75
F	0.04	0.04	.9	0.03	0.09	.06	0.03	0.06	.17

Table 11: Results of the OLS regressions with team payoffs and group allocative efficiency as dependent variables and the subcategories of reasoning as independent variables. Coefficients values are reported with standard errors in parenthesis. Standard significance values for asterisks: *** is 99%, ** is 95%, * is 90%.

	Dependent Variable	
	Team payoff	All. eff.
Const.	23.45*** (4.99)	65.05 (39.89)
Ol	17.98 (13.42)	-37.47 (36.85)
Op	-11.90 (12.43)	60.66 (73.02)
Ob	124.30*** (44.86)	157.25 (129.97)
Sp	5.51 (12.94)	62.34 77.15
Sn	12.55 (16.44)	-58.11 (39.58)

Table 12: Results of the OLS regressions with team payoffs and group allocative efficiency as dependent variables and the categories of action as independent variables. Coefficients values are reported with standard errors in parenthesis. Standard significance values for asterisks: *** is 99%, ** is 95%, * is 90%.

	Dependent Variable			
	Team payoff	All. eff.	Team payoff	All. eff.
Const.	38.59*** (12.13)	87.36* (36.26)	33.64*** (11.13)	100.42** (34.66)
A	-9.84 (13.78)	10.74 (53.35)	3.03 (13.33)	-0.64 (49.84)
B	-17.50 (14.33)	-45.78 (39.32)	-1.79 (14.18)	-36.11 (36.89)
C	-60.70* (33.61)	5.20 (58.09)	-45.86 (30.91)	-5.13 (54.06)
F	57.76* (29.13)	161.86* (69.24)	54.72** (26.41)	98.74 (77.58)
R	-64.46 (133.20)	342.62 (341.17)	-31.48 (121.25)	237.30 (323.14)
Choice change			-0.28** (0.102)	-0.10 (0.071)

Table 13: Results of the OLS regressions with the shares of O and S arguments in the last part of the game as dependent variables. Independent variables are the shares of O and S arguments in the first part (O1 and S1), the shares of O and S arguments in the other teams in the middle part (Oo2 and So2), a treatment dummy (TD) and the number of rounds in which the teams has been on the minority side in the middle part (Win2). Coefficient values are reported with standard errors in parenthesis. Standard significance values for asterisks: *** is 99%, ** is 95%, * is 90%.

	Dependent Variable	
	O arguments (last part)	S arguments (last part)
Const.	1.29 (0.79)	-0.96 (0.95)
O1	0.93*** (0.33)	-0.10 (0.40)
S1	0.02 (0.46)	0.86 (0.56)
Oo2	-1.56* (0.87)	1.26 (1.05)
So2	-1.10 (0.85)	0.96 (1.03)
TD	0.15 (0.15)	-0.16 (0.17)
Win2	-0.02* (0.009)	0.02* (0.011)

Table 14: Results of the OLS regressions with the shares of Ol, Op, Ob, Sp and Sn arguments in the last part of the game as dependent variables. Independent variables are the shares of arguments of the same type in the first part (O11, Op1, Ob1, Sp1 and Sn1), the shares of O and S arguments in the other teams in the middle part (Oo2 and So2), a treatment dummy (TD) and the number of rounds in which the team has been n the minority side in the middle part (Win2). Coefficient values are reported with standard errors in parenthesis. Standard significance values for asterisks: *** is 99%, ** is 95%, * is 90%.

	Dependent Variable				
	Ol arg. (last p.)	Op arg. (last p.)	Ob arg. (last p.)	Sp arg. (last p.)	Sn arg. (last p.)
Const.	0.011 (0.620)	0.833 (0.491)	0.072 (0.051)	0.002 (0.656)	0.503 (0.583)
O11	0.102 (0.129)	-	-	-	-
Op1	-	-0.030 (0.061)	-	-	-
Ob1	-	-	0.000 (0.020)	-	-
Sp1	-	-	-	-0.072 (0.168)	-
Sn1	-	-	-	-	0.018 (0.088)
Oo2	0.045 (0.699)	-0.793 (0.553)	-0.097* (0.057)	-0.036 (0.736)	0.020 (0.658)
So2	0.333 (0.699)	-0.784 (0.553)	-0.079 (0.057)	-0.032 (0.734)	-0.393 (0.660)
TD	0.087 (0.115)	0.134 (0.088)	0.009 (0.009)	-0.049 (0.118)	-0.302*** (0.105)
Win2	-0.002 (0.007)	-0.001 (0.006)	0.000 (0.001)	0.019** (0.008)	-0.006 (0.007)