The minority game: An economics perspective

Willemien Kets*

November 18, 2008

Abstract

This paper gives a critical account of the minority game literature. The minority game is a simple congestion game: players need to choose between two options, and those who have selected the option chosen by the minority win. The learning model proposed in this literature seems to differ markedly from the learning models commonly used in economics. We relate the learning model from the minority game literature to standard game-theoretic learning models, and show that in fact it shares many features with these models. However, the predictions of the learning model differ considerably from the predictions of most other learning models. We discuss the main predictions of the learning model proposed in the minority game literature, and compare these to experimental findings on congestion games.

JEL classification: C73, C90.

Keywords: Learning, congestion games, experiments.

*Santa Fe Institute and Tilburg University. Address: 1399 Hyde Park Road, Santa Fe, NM 87501, USA. E-mail: willemien.kets@santafe.edu. Phone: +1-505-946 2782.

1 Introduction

Congestion games are ubiquitous in economics. In a congestion game (Rosenthal, 1973), players use several facilities from a common pool. The costs or benefits that a player derives from a facility depends on the number of users of that facility. A congestion game is therefore a natural game to model scarcity of common resources. Examples of such systems include vehicular traffic (Nagel et al., 1997), packet traffic in networks (Huberman and Lukose, 1997), and ecologies of foraging animals (DeAngelis and Gross, 1992). Similar coordination problems are encountered in market entry games (Selten and Güth, 1982).

Congestion games are also interesting from a theoretical point of view. In congestion games, players need to coordinate to differentiate. This seems to be more difficult than coordinating on the same action, as any commonality of expectations is broken up. For instance, when commuters have to choose between two roads A and B and all believe that the others will choose road A, nobody will choose that road, invalidating beliefs. The sorting of players predicted in the pure-strategy Nash equilibria of such games violates the common belief that in symmetric games, all rational players will evaluate the situation identically, and hence, make the same choices in similar situations (see Harsanyi and Selten, 1988, p. 73). Moreover, in congestion games, players may obtain asymmetric payoffs in equilibrium which may complicate attainment of equilibrium, as coordination cannot be achieved through tacit coordination based on historical precedent (cf. Meyer et al., 1992). Finally, congestion games often have many equilibria, so that players also face the difficulty of coordinating on the same equilibrium.

Nevertheless, the theory of learning in games provides sharp predictions on players' behavior in congestion games. As congestion games belong to the class of potential games (Monderer and Shapley, 1996), all results that have been derived for potential games apply to the class of congestion games (Hofbauer and Hopkins, 2005; Hofbauer and Sandholm, 2002; Monderer and Shapley, 1996; Sandholm, 2001, 2007). Experimental evidence, however, is not always in line with these predictions. Though several experimental studies have shown that players are remarkably successful at learning to coordinate in congestion games (e.g Kahneman, 1988; Meyer et al., 1992; Selten et al., 2007), regularities on the aggregate level generally conceal non-equilibrium behavior at the individual level. Even though aggregate play is close to the Nash equilibrium, individual players generally do not play equilibrium strategies (Meyer et al., 1992; Erev and Rapoport, 1998; Selten et al., 2007; Bottazzi and Devetag, 2004). Moreover, providing players with more information does not always lead to better outcomes. For instance, in their experiments on market entry games, Erev and Rapoport (1998) find that providing players with information on other players' actions may actually lead to *lower* average payoffs.

These experimental findings are hard to explain with standard learning models. This paper discusses the literature on the minority game, a simple congestion game based on the El Farol bar problem introduced by Arthur (1994). Players have to choose between two alternatives. Only those who have chosen the minority side get a positive payoff. The minority game literature proposes a learning model that is able to account for many of the experimental findings listed above. We relate this learning model to the standard learning models in economics, and compare its predictions to experimental results on congestion games. The contribution of the current paper is that it relates the literature on the minority game, which has been largely developed in physics, to the literature on learning in game theory and to the literature in experimental economics on congestion games.[1]

The outline of this paper is as follows. In Section 2, we introduce the minority game and discuss its equilibria. The learning model proposed in the minority game literature is discussed in Section 3. In Section 4, we discuss the main predictions from the learning model. These predictions are compared to experimental results on congestion games in Section 5. Section 6 concludes.

2 The stage game

The minority game is a game in which an odd number of players have to choose between two actions; for instance, players either go to a bar or stay home, either buy or sell an asset, etcetera. Players want to distinguish themselves from the crowd: their aim is to take a different action than the majority of players.

Following the notation of Tercieux and Voorneveld (2005), we denote the set of players by $\mathcal{N} = \{1, \ldots, 2k+1\}$, with $k \in \mathbb{N}$. Each player $i \in \mathcal{N}$ has a set of *pure strategies* $A_i = \{-1, +1\}$: agents have to choose between two options. The set of *mixed strategies* of player i is denoted by $\Delta(A_i)$. We denote a mixed strategy profile by $\alpha \in \times_{i \in \mathcal{N}} \Delta(A_i)$, and we use the standard notation $\alpha_{-i} \in \times_{j \in \mathcal{N} \setminus \{i\}} \Delta(A_j)$ to denote a strategy profile of players other than $i \in \mathcal{N}$. With each action $a \in \{-1, +1\}$, a function

$$f_a : \{1, \ldots, 2k+1\} \to \mathbb{R}$$

can be associated which indicates for each $n \in \{1, \ldots, 2k+1\}$ the payoffs to a player choosing a when the total number of players choosing a equals n. The von Neumann-Morgenstern utility function of a player is then given by

$$u_i(a) = f_{a_i}\left(|\{j \in \mathcal{N} : a_j = a_i\}|\right), \tag{2.1}$$

where $a \in \times_{j \in \mathcal{N}} A_j$. Payoffs are extended to mixed strategies in the usual way.

The function $f_a(\cdot), a \in \{-1, +1\}$ can have several forms. It is commonly assumed that congestion is costly:

[Mon] f_{-1} and f_{+1} are strictly decreasing functions,

and that the congestion effect is the same across alternatives:

[Sym] $f_{-1} = f_{+1}$.

A commonly used form is $f_{-1}(n) = f_{+1}(n) = 1$ if $n \in \{1, \ldots, k\}$ and 0 otherwise (Challet and Zhang, 1997). Alternatively, one could define payoffs in terms of the aggregate action $\sum_{i \in \mathcal{N}} a_i$ for a given action profile $a = (a_i)_{i \in \mathcal{N}}$, with $a_i \in \{-1, +1\}$ for all i. Let g be a function on $\mathbb{R}$ such that $g(-x) = -g(x)$ for all $x \in \mathbb{R}$ and $g(x) > 0$ for $x > 0$. A player $i \in \mathcal{N}$ is then assigned the payoff

$$u_i(a) = -a_i g\left(\sum_{j \in \mathcal{N}} a_j\right). \tag{2.2}$$

In our notation:

$$f_{-1}(n) = f_{+1}(n) = g\left(2(k-n)+1\right).$$

Common choices include

$$g(x) = x/(2k+1) \tag{2.3}$$

and

$$g(x) = \text{sign}(x).$$

Most of the predictions of the learning model are not affected qualitatively by the precise choice of payoff function, given that it satisfies [Mon] and [Sym] (Li et al., 2000). Notice that the minority game is a congestion game (Rosenthal, 1973) and hence a finite exact potential game (Monderer and Shapley, 1996).

To analyze the game's Nash equilibria, we introduce some more notation. A player who uses a mixed strategy that puts positive probability on both pure strategies is referred to as a *mixer*. A player that puts full probability mass on the alternative -1 is called a (-1)-*player*; similarly, a player that puts full probability mass on the alternative $+1$ is called a $(+1)$-*player*.

The stage game has a large number of Nash equilibria. Tercieux and Voorneveld (2005) show that a pure strategy profile is a Nash equilibrium if and only if one of the alternatives -1 or $+1$ is chosen by exactly k of the $2k+1$ players. Note that these Nash equilibria are *not* strict, as a player that is in the majority is indifferent between sticking to his choice or switching actions, as his deviation would shift the majority. There are $2\binom{2k+1}{k}$ of such asymmetric pure-strategy Nash equilibria.

Kets and Voorneveld (2007) characterize the game's mixed-strategy Nash equilibria. It can be shown that in any Nash equilibrium with at least one mixer, all mixers use the same mixed strategy. Moreover, player labels are irrelevant by [Sym] (if α is a Nash equilibrium, so is every permutation of α). Together, these facts imply that a Nash equilibrium with at least one mixer can be summarized by its *type* (ℓ, r, λ), where $\ell, r \in \{0, 1, \ldots, 2k+1\}$ denote the number of players choosing pure strategy -1 or $+1$, respectively, and $\lambda \in (0,1)$ the probability with which the remaining $z(\ell, r, \lambda) := (2k+1) - (\ell + r) > 0$ mixers choose -1. Let $v_{-1}(\ell, r, \lambda)$ denote the expected payoff to a player choosing -1; $v_{+1}(\ell, r, \lambda)$ is defined similarly. Then, a strategy profile of type (ℓ, r, λ) is a Nash equilibrium if and only if

$$v_{-1}(\ell+1, r, \lambda) = v_{+1}(\ell, r+1, \lambda), \tag{2.4}$$

i.e., the expected payoffs to a mixer of playing the pure strategy $a = -1$ are equal to the expected payoffs of the pure strategy $a = +1$. It can be shown that there exist Nash equilibria with exactly one mixer. These equilibria are of type (k, k, λ) with arbitrary $\lambda \in (0, 1)$, i.e., the mixer uses an arbitrary mixed strategy, whereas the remaining $2k$ players are spread evenly over the two pure strategies. In addition, there are Nash equilibria with more than one mixer. For $\ell, r \in \{0, 1, \dots, 2k+1\}$ such that $\ell + r \leq 2k - 1$, there is a Nash equilibrium of type (ℓ, r, λ) if and only if $\max\{\ell, r\} < k$. The corresponding probability $\lambda \in (0, 1)$ solves (2.4), and can be shown to be unique. It follows from these results that there is a unique symmetric mixed-strategy Nash equilibrium. In this equilibrium, each player chooses each option with probability $\frac{1}{2}$. The expected number of players choosing each option is then $k + \frac{1}{2}$.

3 Learning in the minority game

Players in the minority game face both a coordination problem and an incentive problem. The coordination problem is not easy to solve. As the equilibria in pure strategies cannot be Pareto-ranked or ordered in terms of risk-dominance, no particular pure-strategy Nash equilibrium can be singled out as being most salient (Schelling, 1960). Hence, without pre-play communication, players do not have enough information to implement a pure-strategy Nash equilibrium (cf. Menezes and Pitchford, 2006). While players could use common knowledge of rationality and symmetry to deduce and select the symmetric mixed-strategy Nash equilibrium (cf. Ochs, 1990; Meyer et al., 1992), this may raise an incentive problem, as players can earn a higher payoff than in the symmetric mixed-strategy Nash equilibrium if they manage to outsmart the other players. Hence, players may try to find patterns in the play of others when the game is played repeatedly (cf. Arthur, 1994; Meyer et al., 1992). The learning model proposed in the minority game literature provides a way of formalizing this notion. In this section, we first introduce the model, and then discuss its assumptions, relating the learning model to other learning models in the literature.

3.1 Model

Time is discrete and indexed by $t \in \{0, 1, \dots\}$. At each time t, the stage game is played. After each round of play t of the stage game, the players are informed of the aggregate action $A(t) := \sum_{i \in N} a_i(t)$, where $a_i(t) \in A_i = \{-1, +1\}$ is the action taken by player i at time t. We assume that players have a limited memory: they only retain the sequence of the minority actions

h_m	$s_{i,1}$	$s_{i,2}$	$s_{i,3}$	$s_{i,4}$
$(-1,-1,-1)$	$+1$	-1	-1	$+1$
$(-1,-1,+1)$	-1	-1	$+1$	-1
$(-1,+1,-1)$	$+1$	-1	-1	$+1$
$(-1,+1,+1)$	-1	$+1$	-1	$+1$
$(+1,-1,-1)$	$+1$	$+1$	$+1$	$+1$
$(+1,-1,+1)$	-1	-1	$+1$	$+1$
$(+1,+1,-1)$	-1	-1	-1	$+1$
$(+1,+1,+1)$	-1	$+1$	-1	$+1$

Table 1: An example of a subset of response modes with $m = 3$ and $n_S = 4$ for some player $i \in N$. For instance, if the history of outcomes is $(-1,-1,-1)$, then response mode $s_{i,1}$ prescribes action $a_i = +1$, and response mode $s_{i,2}$ prescribes action $a_i = -1$.

in the previous m rounds, where $m \in \mathbb{N}$. More specifically, in round t, players observe a *history* $h_m(t) = (-\text{sign}[A(\tau)])_{\tau\in\{t-m,t-m+1,\ldots,t-1\}}$, where we note that $-\text{sign}[A(t)]$ indicates the minority action at time t: if there are fewer players choosing -1 $(+1)$ at time t than there are players choosing $+1$ (-1), then $-\text{sign}[A(t)]$ is equal to -1 $(+1)$.

A *response mode* s assigns to each history $h_m \in \{-1,+1\}^m$ an action $a \in \{-1,+1\}$. That is, a response mode s prescribes which action $s(h_m(t)) \in \{-1,+1\}$ to take, for a given history of play $h_m(t)$ at time t. Note that a $s(h_m(t))$ does not depend on t, other than through $h_m(t)$: if $h_m(t) = h_m(t')$, then $s(h_m(t)) = s(h_m(t'))$. It is not hard to see that there are 2^{2^m} different response modes: there are 2^m possible signals h_m of length m, and for each signal, there are two possible actions. For memory length m, denote the set of all response modes by $\mathcal{S}^{(m)}$. An important assumption in the current learning model is that each player $i \in N$ is endowed with a subset S_i of the set of all possible response modes $\mathcal{S}^{(m)}$. All players are endowed with the same number of response modes: there exists $n_S \in \mathbb{N}$ such that $|S_i| = n_S$ for all $i \in N$. We assume that $n_S \geq 2$. For each player $i \in N$, the response modes in S_i are drawn uniformly at random from $\mathcal{S}^{(m)}$, independently across players. An example for $n_S = 4$ and $m = 3$ is given in Table 1.

When faced with a history, a player has to choose which of his response modes to use in the next period. Each player keeps a virtual score for each response mode in his endowment that reflects that response mode's past performance. The virtual score of each response mode is updated after each time period, regardless of whether the response mode has been used or not. When a response mode would have correctly predicted the minority action, its virtual score is increased with the payoffs it would have earned, otherwise it is decreased with the same amount. More specifically, the *virtual score* of player $i \in N$ for response mode $s_i \in S_i$ at time $t > 0$ is given by:

$$p_{i,t}(s_i) = p_{i,t-1} - \left(\frac{s_i(h_m(t))}{2k+1}\right) [2(k - \ell(s_i(h_m(t)), t)) + 1]$$

where we recall that $s_i(h)$ is the action prescribed by response mode s_i when the history is h, and where $\ell(a,t)$ is the number of players choosing action a in round t. At $t = 0$, the virtual score of player i for s_i is $p_{i,t}(s_i) = 0$.

An important thing to note is that players do *not* take the effect of their action on the aggregate outcome into account. In determining the virtual score of a response mode, players only consider whether this response mode would have predicted the actual outcome correctly, neglecting the question whether playing this response mode would have affected the outcome. To see this, suppose that at time t, player i chooses $a_i = -1$, and that the total number of players choosing this action is $k+1$, i.e., -1 is the majority action. Then, $2(k-(k+1))+1 = -1$ would be added to all response modes prescribing action $a_i = -1$ (given the current history), and $-(2(k-(k+1))+1) = +1$ would be added to all response modes prescribing $a_i = +1$. However, if player i would have chosen $a_i = +1$, the number of players choosing $a = +1$ would have been $k+1$, and $+1$ would have been the majority action. This is an important assumption of the model, and we discuss its implications in Section 3.2.

The probability that a player chooses a response mode at a given time step is determined by its virtual score at that time, with the choice probabilities following the well-known logit choice rule. For $i \in N$, denote the response mode selected by player i at time t by $s_i(t)$. Then,

$$\forall s_i \in S_i : \quad \mathbb{P}(s_i(t) = s_i) = \frac{\exp\left[\beta\, p_{i,t}(s_i)\right]}{\sum_{s_j \in S_i} \exp\left[\beta\, p_{i,t}(s_j)\right]}, \tag{3.1}$$

where $\beta > 0$ is the logit parameter. The parameter β can be interpreted as the sensitivity of choice to marginal information. In the limiting case $\beta \to \infty$, players mix uniformly among the response modes with the highest virtual score. Otherwise, players choose response modes with lower virtual scores with positive probability, with a probability increasing in the virtual scores. Perhaps surprisingly, this additional noise may actually improve collective performance, as we discuss in Section 4.1.

3.2 Discussion

In this section, we discuss two of the most important assumptions of the learning model in the minority game model: the assumption that all players are endowed with a random subset of response modes and the assumption that players update the virtual scores of response modes not used, without taking into account the effect of that response mode on the game's outcome. Although the learning model of the minority game literature seems to depart markedly from the standard evolutionary and learning models used in economics, we argue here that in fact the learning model combines different aspects of several game-theoretic models to provide a realistic model of player behavior in congestion games.

3.2.1 Response modes and heterogeneity

In the learning model proposed in the minority game literature, players base their action on the recent past, trying to discern patterns in their opponents' behavior, as in Arthur (1994). In the El Farol bar problem described by Arthur, players need to decide whether to go to a bar or not. Going to the bar is only pleasant if it is not too crowded. Arthur proposes that players condition their decision to go on attendance levels in the previous weeks: if the bar has been crowded for the last three weeks, say, they expect it to be crowded next week also. These mental models are mapped into actions: if a player expects the bar to be crowded, he will not go.

The response modes in the learning model of the minority game literature are a concise way of modeling this notion. An important question, however, is which response modes need to be included in the model. There are two possible approaches. Firstly, one could simply incorporate all possible response modes. However, if all possible response modes are included in the learning model, the strategy space becomes huge already for very simple games. Many different response modes are conceivable in a simple game such as the minority game, as illustrated by the list of examples in Arthur (1994).

A second possibility is to include only a selection of possible response modes. In that case, one could either make a selection based on behavioral assumptions, or let the subset of response modes be determined at random. In the first case, a natural choice is to include response modes that reflect beliefs about other players' actions, based on recent outcomes. The first approach is commonly taken in the economics literature (e.g. Erev and Rapoport, 1998; Selten et al., 2007), while the current learning model chooses the second approach. In the latter case, there are no restrictions on the types of response modes that players use.

This may seem to be a weak point of the model, as response modes need not have a sensible interpretation in the learning model of the minority game literature. However, in games such as the minority game, whether a response mode is reasonable *only* depends on the response modes used by others. Conversely, *any* response mode, whether it has a sensible interpretation or not, will work if opponents use response modes that recommend them to take the opposite action. For instance, in experiments on route-choice games, Selten et al. (2007) report that some subjects use a "direct" response mode, while other subjects use a "contrarian" response mode. Subject who use the former response mode will switch roads if they experienced congestion in the last period, while subjects using the contrarian response mode stick with their choice, as they expect other subjects to switch. The important point to note here is that the direct response mode is only sensible if there are players who use the contrarian response mode and vice versa. In such a case, agnosticism on the type of response modes that players use may well provide a more realistic model of players' reasoning processes than the more restrictive assumptions employed in different learning models. This offers an elegant solution to the dilemma signalled by Erev and Roth (1998, p. 873) that it is virtually impossible to include all possible behavioral rules, but that selection of specific rules bears the risk of "parameter fitting in a model with an enormous number of parameters". In the learning model proposed in the minority game literature, no response mode is ruled out on a priori grounds, while sensible behavioral rules evolve naturally, as the only criterion for a behavioral rule to be sensible in the minority game is that there are other players who follow a "contrarian" behavioral rule. Indeed, in Section 4.3, we show that under the current learning model, players will self-organize into groups that use different response modes in such a way that their actions cancel out to the extent possible.

However, this approach raises some questions. Firstly, one may ask why players are heterogeneous in their endowment of response modes. Perhaps more importantly, one could ask why players only consider a fixed number of response modes. Indeed, individual players have an incentive to increase the number of response modes they use, as that gives them an advantage over other players (Marsili et al., 2000). However, these assumptions are not uncommon in game-theoretic models of learning and bounded rationality. Possible justifications for such assumptions include that each player has different experiences prior to playing the minority game and therefore deems different response modes more reasonable than others (cf. Aumann, 1997; Fudenberg and Levine, 1998), and that boundedly rational players may prefer to just consider a subset of response modes that have worked well in the past, rather than considering all 2^{2^m} response modes (cf. Ellison and Fudenberg, 1993).

3.2.2 The law of simulated effect and boundedly rational players

Which response mode players choose from the set of response modes they are endowed with, is determined by the virtual score of each response mode. The learning process proposed in the minority game literature is closely related to the reinforcement learning model of Roth and Erev (1995) and Erev and Roth (1998). The main difference between the basic reinforcement learning model of Roth and Erev and the learning model of the minority game literature lies in the updating of the score of strategies or response modes not played. In the basic reinforcement learning model, the scores of these strategies are not updated, while in the learning model proposed in the minority game literature, the scores of all response modes are updated every period, as in hypothetical reinforcement learning or stochastic fictitious play (Fudenberg and Levine, 1998). The assumption that players also consider the payoffs to strategies or response modes not played seems to be reasonable. Camerer and Ho (1999) argue on the basis of theoretical arguments as well as on the basis of experimental findings that players obey not only the "law of *actual* effect", but also the "law of *simulated* effect", meaning that in reinforcement, not only payoffs from strategies that are actually used count, but also foregone payoffs from strategies not played.

However, for players to play according to the law of simulated effect, they need more information than for standard reinforcement learning. Under standard reinforcement learning, players only need to know the payoff to the action they choose. By contrast, to play according to stochastic fictitious play, players additionally need to know the payoff rule as well as the actions chosen by their opponents. Even in a game such as the minority game, where the players only need to know the aggregate choice of other players (and not their individual choices), calculating foregone payoffs of strategies not used may be too hard for players that are boundedly rational. In the learning model proposed in the minority game literature, players' bounded rationality is reconciled with the law of simulated effect by assuming that players do not take the effect of their own action on the global outcome into account. In that way, players can account for foregone payoffs of response modes not used, without having to do complicated calculations.

One may think that for a large number of players, it will not matter much whether players account for their own impact. However, due to the minority rule, there remains a systematic bias in the rewarding of response modes, even if the number of players is arbitrarily large. The reason is that the reward for a response mode that is currently played is systematically lower than that for the response modes that are not used. These latter response modes get a point if they prescribe the current minority side, even if they would have tipped the minority to the other side if they would have been played, so that they would have guessed wrong in reality. As the response mode that is actually played does not have this advantage, the response modes that are not played are systematically favored and hence results depend on whether players take the effect of their action on the aggregate outcome into account. This makes that players keep switching between response modes: over time, a response mode that is not played for some time will gather sufficiently many virtual points so as to be selected to be played, thus losing its advantage, until another response mode takes over again.

Interestingly, if players correct for this bias by allocating a small additional reward to response modes currently played, this does not happen, and in the long run, players use the same response mode in every period (in the limit $\beta \to \infty$). When m is sufficiently small, in each round, k players will choose one action, and $k+1$ the other, as in the pure Nash equilibria of the game (Marsili et al., 2000).

The learning model proposed in the minority game literature thus combines features from several learning models in the literature on learning in games. However, it makes distinctly different predictions than game-theoretical learning models. We discuss these predictions in the next section.

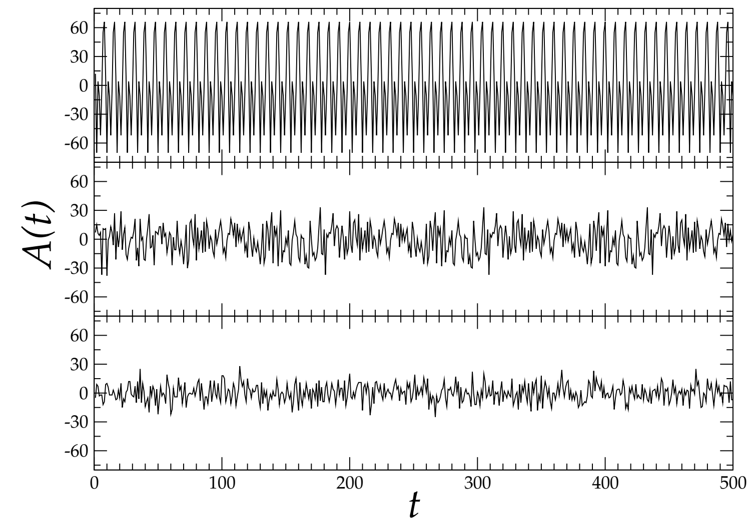

Figure 4.1: Time evolution of the aggregate action $A(t)$, with $2k + 1 = 301$ and $n_S = 2$. Panels correspond to $m = 2, 7, 15$ from top to bottom. Figure taken from Moro (2003).

4 Predictions of the learning model

In this section, we discuss the main predictions on the learning model proposed in the minority game literature. Some results are obtained analytically, others by simulations. In the simulations, a given number of agents is endowed with a random subset of response modes, and results are obtained by averaging over different assignments. In the first two sections, we characterize the behavior of the model in terms of social efficiency and informational efficiency, and show that the two are intimately linked in this learning model. In Section 4.3, we discuss how the predictions of the model can be understood in terms of the formation of groups who use counteracting response modes.

4.1 Volatility

Simulations show that the aggregate action $A(t) := \sum_{i \in N} a_i(t)$ keeps fluctuating around 0, as can be seen in Figure 4.1. As the game is symmetric, the time average of $A(t)$ will be 0 in the long run (e.g. Challet and Zhang, 1997). More interesting is the behavior of the *volatility* $\sigma^2 := \langle A^2 \rangle$, where $\langle \cdot \rangle$ denotes the (time) average of a quantity. The volatility is a measure of the degree of efficiency (measured in terms of aggregate payoffs) achieved in a population. The higher the volatility, the larger the loss in aggregate payoffs: large fluctuations around 0 imply that the size of the minority is only small, as aggregate payoffs are proportional to $-\sum_i a_i(t)A(t) = -(A(t))^2$.

By simulations, it has been found that σ^2 is only a function of $\alpha := 2^m/(2k+1)$ and n_S, where we recall that n_S is the number of response modes of each player (Savit et al., 1999). Figure 4.2

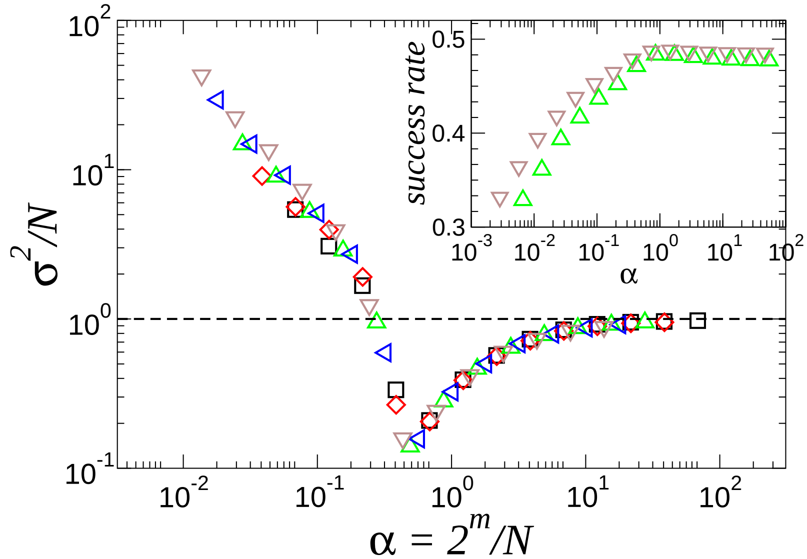

Figure 4.2: Volatility as a function of the order parameter α for $n_S = 2$ and different number of players $N := 2k+1 = 101, 201, 301, 501, 701$ (□, ◊, △, ◁, ▽, respectively). The critical value α_c is the value of α for which the volatility is at a minimum. Inset: players' average success rate as function of α (not discussed here, see Moro (2003)). Figure taken from Moro (2003).

shows the volatility as a function of α. As can be seen in the figure, the volatility converges to the volatility exhibited in the symmetric mixed Nash equilibrium for $\alpha \to \infty$. With a large number of players (α small), overall performance is much worse. In fact, the volatility is of order $(2k+1)^2$, so that the size of the group of players choosing the minority action is much smaller than k. At intermediate values of α, volatility is low, and it attains a minimum at $\alpha_c(n_S) \cong n_S/2 - 0.66$ (Marsili et al., 2000). Hence, at intermediate values of α, players are able to coordinate their actions and perform better collectively than under the symmetric mixed Nash equilibrium. This means that players can somehow exploit the available information to reduce σ^2 relative to the symmetric mixed Nash equilibrium. Note that this is not the result of some form of cooperative behavior of the players: agents are selfishly maximizing their own payoffs, and for intermediate values of α, this leads to higher aggregate payoffs. However, coordination is not complete under the current learning model. The aggregate payoff is maximized if players play according to one of the pure Nash equilibria of the game, with k players choosing the minority action. In that case, almost half of the players are in the minority, and $\sigma^2/(2k+1) = 1/(2k+1)$. Players come close to this optimum at $\alpha = \alpha_c$, but they never reach it.

Strikingly, global efficiency is enhanced for certain values of α when players do not always choose the response mode with the highest number of virtual points, i.e., when $\beta < \infty$ in Equation (3.1). It can be shown that for $\alpha < \alpha_c$, when the volatility is much higher than under the benchmark of the symmetric Nash equilibrium, volatility *decreases* when the noise level $1/\beta$ *increases*. For $\alpha > \alpha_c$, the value of β does not affect the level of volatility (Cavagna et al., 1999). The explanation is that under the current learning model, players form herds when $\alpha < \alpha_c$. The parameter α is a measure of the total number of response modes relative to the number of players. When $\alpha < \alpha_c$, there are few response modes relative to the number of players. In that case, players have to herd at a limited number of response modes, leading to a large number of players choosing the same alternative (see Section 4.3). Decreasing β is then equivalent to slowing down the updating of virtual scores for response modes. A finite β therefore acts as a brake against overreaction.

4.2 Information and efficiency

As discussed in the previous section, players seem to be able to coordinate reasonably well for some parameter configurations. The only way players can interact is through the history of play. This observation led some authors to study the information contained in the history of play. The information content of the history of play, or the degree of predictability can be measured by (Challet and Marsili, 1999):

$$H := \frac{1}{2^m} \sum_{\nu=1}^{2^m} \langle A(t+1) \mid h_m(t) = \nu \rangle^2,$$

where $\langle A(t+1) | h_m(t) = \nu \rangle$ is the time average of the aggregate action conditional on a given history of play. Loosely speaking, H measures the information in the time series of $A(t)$. If $A(t+1)$ and $h_m(t)$ are independent, then $H = 0$. If $H > 0$, then the signal $A(t)$ contains information. It can be shown that players under the current learning model minimize the degree of predictability (Marsili et al., 2000). Depending on the value of α, players are more or less successful in doing that. At α_c, the system changes from an informationally efficient phase with low aggregate payoffs ($H = 0$, σ^2 large) to an information-rich phase with high aggregate payoffs ($H > 0$, σ^2 small). In the informationally efficient phase, aggregate payoffs are lower than under the symmetric mixed Nash-equilibrium. By contrast, in the information rich phase, players manage to coordinate and aggregate payoffs are higher than under the symmetric mixed Nash equilibrium.

This transition between the informationally efficient and the information rich phase, or equivalently between the phase with low aggregate payoffs and the phase with high aggregate payoffs, is central to the current learning model. At this transition, there is a qualitative change in collective behavior, while the principles behind the behavior of individual players remain unchanged. For all values of α, players try to outsmart each other, but for low values of α, they are on average less successful. In the next section, we discuss the interpretation of α.

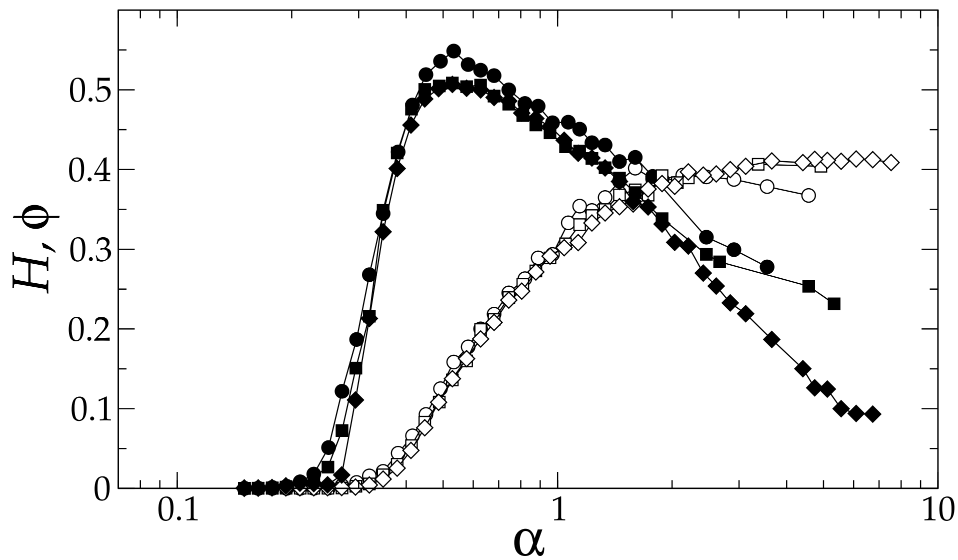

Figure 4.3: Information H (open symbols) and fraction of frozen players ϕ (full symbols; not discussed here, see Moro, 2003) as a function of the control parameter $\alpha = 2^m/(2k+1)$ for $n_S = 2$ and $m = 5, 6, 7$ (circles, squares and diamonds, respectively). Figure taken from Moro (2003).

4.3 Response modes and their antagonists

We have seen that the qualitative behavior of the system depends mainly on $\alpha = 2^m/(2k+1)$. Moreover, for some values of α, players are much more successful in coordinating behavior than for other values. What is the feature of the model underlying this behavior? We address this question in the current section. The answer to this question points to an intuitive interpretation of the model's results in terms of groups using counteracting response modes.

The minority rule forces players to differentiate: if all players choose the same response mode, all players obtain negative payoffs. As there are 2^{2^m} possible response modes for $2k+1$ players, one would expect that players succeed in differentiating if $2k+1$ is much smaller than 2^{2^m}, and be unsuccessful when $2k+1$ exceeds 2^{2^m}. Hence, one would expect a qualitative change when $2k+1$ is of order 2^{2^m}, rather than of order 2^m, as observed. To explain why the transition occurs when $2k+1$ is of order 2^m, we need some more definitions. For each $s, s' \in \mathcal{S}^{(m)}$, define

$$H^{(m)}(s, s') := \sum_{h \in \{-1,+1\}^m} |s(h) - s'(h)|$$

to be the Hamming distance between s and s'. Then, let

$$D^{(m)}(s, s') := \frac{1}{2^m} H^{(m)}(s, s')$$

be the *normalized distance* between s and s'. If $D^{(m)}(s, s') = 1$, i.e., s and s' prescribe different actions for each possible history of play, we say that s and s' are *anti-correlated*. When $D^{(m)}(s, s') = 1/2$, we say that they are *uncorrelated*. The reason that the transition occurs when $2k + 1$ is of order 2^m is that two response modes only give rise to distinctively different behavior if they are anti-correlated or uncorrelated. It can be shown analytically that for every response mode, the number of response modes that are anti-correlated or uncorrelated with that response mode is $2 \cdot 2^m / n_S$ (Challet and Zhang, 1998). Hence, the qualitative behavior depends on 2^m, not 2^{2^m}.

This leads us to an intuitive interpretation of the model's results in terms of the interplay between groups using different response modes. Let s be a response mode, and let $\bar{s}$ be the response mode that is anti-correlated with s. Suppose N_s players use the response mode s in a given time period, while $N_{\bar{s}}$ players use the anti-correlated response mode $\bar{s}$ in that period. If N_s is approximately equal to $N_{\bar{s}}$ for all anti-correlated pairs $(s, \bar{s})$ of response modes, then the actions of players using these response modes effectively cancel and the volatility will be small.

Hence, it would be optimal if the group of players that use a certain response mode is of about the same size as the group that uses the "antagonistic" response mode. However, this is not always possible, as the dimension of the space of response modes is fixed by the parameter m. Hence, when the number of players is large, the number of response modes they use will be larger than 2^m, so that players are forced to use response modes that are positively correlated. This gives rise to herding effects and large volatility when α is small. For somewhat larger values of m (for a fixed number of players), players use response modes that are either uncorrelated or mutually anti-correlated. In that case, players spread more or less evenly over both actions in each period. Finally, when m is very large relative to the number of players, the number of players using a given response mode will only be small, so that players act more or less independently (Moro, 2003). However, aggregate payoffs are still higher than under the benchmark of the symmetric mixed Nash equilibrium, as there always exist pairs of players that follow anti-correlated response modes (Challet and Zhang, 1998).

5 Comparison to experimental results

In this section, we discuss some experiments on the minority game and related congestion games. In addition to the minority game, we focus on market entry games and route-choice games. First, we briefly introduce the latter two classes of games. We then present some experimental results, and discuss whether the learning model proposed in the minority game literature could explain these results.

The market entry game (Selten and Güth, 1982) has been studied extensively in economics (see Ochs (1999) and references therein; see Duffy and Hopkins (2005) for a recent contribution).

In a market entry game, N players must decide independently and simultaneously to enter a market with a fixed capacity $c < N$ or to stay out. Players who enter the market receive a payoff that decreases in the number of entrants. The payoff of players who stay out of the market is commonly taken to be constant. The game generally has a large number of Nash equilibria, both in pure and in mixed strategies. Depending on the exact form of the payoff function, there may even be a continuum of equilibria. Pure Nash equilibria may be payoff-symmetric or payoff-asymmetric, and strict or non-strict, depending on the choice of parameters. For the payoff functions commonly studied, the (expected) number of entrants is between $c - 1$ and c in equilibrium. An important difference between the market entry game and the minority game is that in the latter game, the payoffs of both actions are subject to congestion, while in the market entry game, players can choose between an action with constant payoffs (staying out) and an action whose payoffs are subject to congestion.

Route-choice games are closer to the minority game in that the payoffs of all actions are subject to congestion. In a route-choice game, players choose between two or more roads. The payoffs of choosing one of these roads decrease in the number of other players who choose that road. In equilibrium, players divide themselves over the roads in such a way that travelling times and hence payoffs are equalized. These games have been studied experimentally by a number of authors (see Selten et al., 2007, and references therein). An important difference with the minority game is that the pure Nash equilibria of the route-choice game are payoff-symmetric, and that they are strict.

We now turn to some experimental work on market entry games, route-choice games and the minority game, and discuss whether experimental findings can be explained by the learning model proposed in the minority game literature. We focus on two issues: aggregate behavior versus individual play and the effect of information on players' behavior. We discuss these issues in turn.

5.0.1 Aggregate behavior versus individual play

A robust finding in experiments on games in these classes is that subjects quickly achieve a "magical" degree of coordination (in the sense that aggregate payoffs are high). However, individual players generally do not play equilibrium strategies. For instance, while Erev and Rapoport (1998) find that the number of entrants in a market entry game rapidly converges to the equilibrium value, they also observe large between- and within-subject variability, which does not diminish with experience. This is a common finding in experiments on market entry games (Ochs, 1999, p. 169).[2] Similarly, in their experiments on route-choice games, Selten et al. (2007) observe that the mean number of drivers on the different roads is very close to the equilibrium number, while large fluctuations in individual behavior persist until the end of the experiment. Similar experimental results have been reported for the minority game (Bottazzi and Devetag, 2007; Chmura and Pitz, 2006). In all cases, the hypothesis that fluctuations can be explained by a symmetric mixed Nash strategy equilibrium of the game can be rejected. These results cannot easily be explained with standard learning models, as these models typically predict convergence to the pure Nash equilibria of such games or to Nash equilibria with at most one mixer (see Duffy and Hopkins (2005) and the previous chapter for a discussion of the predictions of different learning models for the market entry game and the minority game, respectively).

Some authors attempt to reconcile aggregate "equilibrium" behavior in experiments with individual non-equilibrium play by conjecturing that subjects may use counteracting behavioral rules (Bottazzi and Devetag, 2007; Chmura and Pitz, 2006; Erev and Rapoport, 1998; Rapoport et al., 2000; Selten et al., 2007; Zwick and Rapoport, 2002). For instance, Bottazzi and Devetag (2007) find that there is considerable heterogeneity in players' behavior in their experiments on the minority game. They show that it is not the heterogeneity per se which determines the players' success in coordinating, rather, it is the interaction between these different behavioral rules that players can successfully coordinate on choosing different actions. These findings are in line with the predictions of the learning model of the minority game literature that players self-organize in groups that use counteracting response modes, thus reconciling aggregate equilibrium behavior and individual non-equilibrium play.

In most experiments, it is not fully clear which behavioral rules subjects employ. For example, Selten et al. (2007) are unable to classify 42% of the subjects in terms of the behavioral rules they use in their route-choice experiments, while Zwick and Rapoport (2002) cannot classify the behavior of some 60% of their subjects in their experiments on the market entry game. This leaves open the possibility that subjects use some response modes that may not have an intuitive interpretation (and are thus not recognized by the experimenters) but that nevertheless perform well, since there are players using counteracting response modes, as predicted by the current learning model (see Section 3.2 and 4.3). A systematic study of the different response modes used by experimental subjects seems needed. Indeed, Zwick and Rapoport (2002) conclude that there is a need "to re-orient research on interactive decision making to individual differences, identify patterns of behavior shared by subsets of players . . . , and then attempt to account for aggregate behavior in terms of the behavior of the clusters of players that form these aggregates".

5.0.2 Effect of information

The effect of information on players' behavior in such games remains a puzzle. Two dimensions of information have been investigated in the experimental literature.

A first dimension that has been studied is how behavior depends on the information subjects have on others' choices. Players can be provided with information only on the payoff rule and aggregate behavior in the past rounds or may be informed additionally of the individual choices of all other players. Although for a wide range of learning models including the reinforcement learning model and the learning model studied in the minority game literature, this should not affect results, in many experimental studies, behavior differs qualitatively depending on the information players have. For instance, in experiments on market entry games, Duffy and Hopkins (2005) find that behavior becomes less random when players are informed of other players' choices. In market entry games, this could be explained by the fact that providing players with more information allows them to use repeated game strategies, as the additional information allows players to signal their commitment to entering the market. While for the market entry game, such a signalling strategy pays off, this is not the case for the minority game and route-choice games. For instance, suppose that k players in the minority game commit to action $a = -1$, and k players commit to action $a = +1$. The remaining player will not be deterred from choosing either of those actions by the commitment of other players, nor does the commitment of these players guarantee them a positive payoff.[1] Nevertheless, also in the minority game and route-choice games, players switch less often between different actions when they are provided with information on the choices of others (Bottazzi and Devetag, 2007; Selten et al., 2007). This could be explained under the learning model of the minority game literature if the additional information induces players to account somehow for their impact on the aggregate action (see Section 3.2), but it is not clear why this would be the case.

A second dimension of information that has been studied in the literature refers to the salience of information on the recent history of play. Bottazzi and Devetag (2007) provide players with a string of past outcomes of varying length. When players are provided with information on play in more rounds than just the previous one, aggregate payoffs are significantly higher. Bottazzi and Devetag find that a longer history allows players to correlate their behavior over a longer time period. Notably, aggregate payoffs are highest in a treatment where players are provided with information on several rounds, and play is characterized by a substantial lack of short-range correlations between own current and past actions. Hence, players seem to exploit the additional information to improve their payoffs.

All together, these experimental studies lend some support to the learning model proposed in the minority game literature. However, the question how information influences play in congestion games has still not been satisfactorily answered. It would be interesting to compare players' behavior under different informational treatments in different congestion games. While most learning models make similar predictions for the different congestion games discussed here, intuitively, one would expect that information will play a different role in these games, as emotions like envy and regret will be more important in some games than in others, and also the scope for

[1] A repeated-game strategy that is effective in the minority game is one in which players "take turns": players alternately choose each of the two actions in such a way that each player is in the minority roughly half of the time. Indeed, Helbing et al. (2005) find some evidence of such behavior in their experiments on route-choice games with small groups, but it is unlikely that players will be able to successfully play according to such a repeated-game equilibrium when the number of players is large.

repeated-game strategies differs across games. Such a systematic comparison would allow one to better separate the learning effects from possible repeated-game and behavioral effects.

6 Conclusions

In this chapter, we have given a critical account of the learning model proposed in the minority game literature, and related it to standard learning and evolutionary models in economics, showing that it shares quite a few features with these models. Still, the predictions of this learning model are markedly different from the predictions from other models. We have argued that these predictions are in line with a number of experimental results on the minority game and related games which cannot be explained by other learning models.

However, our understanding of learning in such games is still incomplete. For instance, the effect of information on play is unclear. An interesting direction for further research would be to systematically vary players' information in experiments on different congestion games such as the minority game and the market entry game, and to compare play under the different information treatments and across games. While most learning models provide similar predictions for these games, intuitively, one would expect that information may have different effects in these games, as in some games, repeated-game strategies or emotions may play a larger role than in others. Such an experiment may help shed light on the question which learning model is appropriate in such games.

Acknowledgements

I am indebted to Ginestra Bianconi, George Ehrhardt, Doyne Farmer, Matteo Marsili, Esteban Moro, Jan Potters, Dolf Talman, and Mark Voorneveld for inspiring discussions and helpful comments and suggestions. In addition, I would like to thank Esteban Moro for his kind permission for reproducing some of the figures from Moro (2003). All remaining errors are of course my own.

Notes

[1] We have no intention of giving a comprehensive survey of the minority game literature, as an enormous amount of work on the minority game has been done. For an extensive collection of papers on the minority game, see `http://www.unifr.ch/econophysics/minority/`. See Moro (2003); Challet et al. (2004) or Coolen (2005) for an introduction to the field. Papers in economics on the minority game include Bottazzi and Devetag (2007), Chmura and Pitz (2004), and Renault et al. (2005). Blonski (1999) and Kojima and Takahashi (2004) study learning in games very similar to the minority game.

[2] An exception is Duffy and Hopkins (2005) who find that subjects coordinate on one of the pure Nash equilibria of the market entry game after a large number of rounds when they are given feedback on others' choices.

References

Arthur, W. B. (1994). Inductive reasoning and bounded rationality. *American Economic Review 84*, 406–411.

Aumann, R. (1997). Rationality and bounded rationality. In S. Hart and A. Mas-Colell (Eds.), *Cooperation: Game-Theoretic Approaches*, Volume 155 of *NATO ASI Serie F*. Berlin: Springer.

Blonski, M. (1999). Anonymous games with binary actions. *Games and Economic Behavior 28*, 171–180.

Bottazzi, G. and G. Devetag (2004). Coordination and self-organization in minority games: Experimental evidence. In M. Gallegati, A. Kirman, and M. Marsili (Eds.), *The Complex Dynamics of Economic Interaction: Essays in Economics and Econophysics*, Volume 521 of *Lecture Notes in Economics and Mathematical Systems*. Berlin: Springer-Verlag.

Bottazzi, G. and G. Devetag (2007). Competition and coordination in experimental minority games. *Journal of Evolutionary Economics 17*, 241 – 275.

Camerer, C. and T.-H. Ho (1999). Experience-weighted attraction learning in normal form games. *Econometrica 67*, 827–874.

Cavagna, A., J. P. Garrahan, I. Giardina, and D. Sherrington (1999). A thermal model for adaptive competition in a market. *Physical Review Letters 83*, 4429–4432.

Challet, D. and M. Marsili (1999). Phase transition and symmetry breaking in the minority game. *Physical Review E 60*, R6271–R6274.

Challet, D., M. Marsili, and Y.-C. Zhang (2004). *Minority Games: Interacting Agents in Financial Markets*. Oxford: Oxford University Press.

Challet, D. and Y. C. Zhang (1997). Emergence of cooperation and organization in an evolutionary game. *Physica A 246*, 407–418.

Challet, D. and Y. C. Zhang (1998). On the minority game: Analytical and numerical studies. *Physica A 256*, 514–532.

Chmura, T. and T. Pitz (2004). Minority game - Experiments and simulations of traffic scenarios. *Bonn Econ Discussion Papers 23/2004*.

Chmura, T. and T. Pitz (2006). Successful strategies in repeated minority games. *Physica A 363*, 477–480.

Coolen, A. (2005). *The Mathematical Theory of Minority Games: Statistical Mechanics of Interacting Agents*. Oxford: Oxford University Press.

DeAngelis, D. and L. Gross (1992). *Individual-Based Models and Approaches in Ecology: Populations, Communities, and Ecosystems*. New York: Chapman and Hall.

Duffy, J. and E. Hopkins (2005). Learning, information, and sorting in market-entry games: theory and evidence. *Games and Economic Behavior 51*, 31–62.

Ellison, G. and D. Fudenberg (1993). Rules of thumb for social learning. *Journal of Political Economy 101*, 612 – 643.

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. Roth (1998). Predicting how people play games: Reinforcement learning in experimental games with unique mixed strategy equilibria. *American Economic Review 88*, 848–881.

Fudenberg, D. and D. Levine (1998). *The Theory of Learning in Games*. Cambridge, MA: MIT Press.

Harsanyi, J. C. and R. Selten (1988). *A General Theory of Equilibrium Selection in Games*. Cambridge, MA: MIT Press.

Helbing, D., M. Schönhof, H.-U. Stark, and J. A. Holust (2005). How individuals learn to take turns: Emergence of alternating cooperation in a congestion game and the Prisoners Dilemma. *Advances in Complex Systems 8*, 87 – 116.

Hofbauer, J. and E. Hopkins (2005). Learning in perturbed asymmetric games. *Games and Economic Behavior 52*, 133–152.

Hofbauer, J. and W. H. Sandholm (2002). On the global convergence of stochastic fictitious play. *Econometrica 70*, 2265–2294.

Huberman, B. and R. Lukose (1997). Social dilemmas and internet congestion. *Science 277*, 535–537.

Kahneman, D. (1988). Experimental economics: A psychological perspective. In R. Tietz, W. Albers, and R. Selten (Eds.), *Bounded Rational Behavior in Experimental Games and Markets*, pp. 11–18. New York: Springer-Verlag.

Kets, W. and M. Voorneveld (2007). Learning in a simple congestion game. *Discussion Paper, Tilburg University*.

Kojima, F. and S. Takahashi (2004). Anti-coordination games and dynamic stability. *Working Paper, Harvard University*. Forthcoming in International Game Theory Review.

Li, Y., A. VanDeemen, and R. Savit (2000). The minority game with variable payoffs. *Physica A 284*, 461 – 477.

Marsili, M., D. Challet, and R. Zecchina (2000). Exact solution of a modified El Farol's bar problem: Efficiency and the role of market impact. *Physica A 280*, 522–553.

Menezes, F. M. and R. Pitchford (2006). Binary games with many players. *Economic Theory 28*, 125–143.

Meyer, D. J., J. B. Van Huyck, R. C. Battalio, and T. R. Saving (1992). History's role in coordinating decentralized allocation decisions. *Journal of Political Economy 100*, 292–316.

Monderer, D. and L. Shapley (1996). Potential games. *Games and Economic Behavior 14*, 124–143.

Moro, E. (2003). The minority game: An introductory guide. In E. Korutcheva and R. Cuerno (Eds.), *Advances in Condensed Matter and Statistical Physics*. New York: Nova Science Publishers.

Nagel, K., S. Rasmussen, and C. Barrett (1997). Network traffic as a self-organized critical phenomenon. In F. Schweitzer (Ed.), *Self-organization of Complex Structures: From Individual to Collective Dynamics*, pp. 579. London: Gordon and Breach.

Ochs, J. (1990). The coordination problem in decentralized markets: An experiment. *Quarterly Journal of Economics 105*, 545 – 559.

Ochs, J. (1999). Coordination in market entry games. In D. Budescu, I. Erev, and R. Zwick (Eds.), *Games and Human Behavior: Essays in Honor of Amnon Rapoport*, pp. 143–172. Mahwah, NJ: Erlbaum.

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*, 661 – 687.

Renault, J., S. Scarlatti, and M. Scarsini (2005). A folk theorem for minority games. *Games and Economic Behavior 53*, 208 – 230.

Rosenthal, R. (1973). A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory 2*, 65–67.

Roth, A. E. and I. Erev (1995). Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term. *Games and Economic Behavior 8*, 164–212.

Sandholm, W. H. (2001). Potential games with continuous player sets. *Journal of Economic Theory 97*, 81–108.

Sandholm, W. H. (2007). *Population Games and Evolutionary Dynamics*. Cambridge, MA: MIT Press. Forthcoming.

Savit, R., R. Manuca, and R. Riolo (1999). Adaptive competition, market efficiency and phase transitions. *Physical Review Letters 82*, 22032206.

Schelling, T. (1960). *The Strategy of Conflict*. Cambridge, MA: Harvard University Press.

Selten, R. and W. Güth (1982). Equilibrium point selection in a class of market entry games. In M. Diestler, E. Fürst, and G. Schwödiauer (Eds.), *Games, Economic Dynamics, and Time Series Analysis*. Wien-Würzburg: Physica-Verlag.

Selten, R., M. Schreckenberg, T. Pitz, T. Chmura, and S. Kube (2007). Commuters route choice behavior. *Games and Economic Behavior 58*, 394 – 406.

Tercieux, O. and M. Voorneveld (2005). The cutting power of preparation. *SSE/EFI Working Paper 583*.

Zwick, R. and A. Rapoport (2002). Tacit coordination in a decentralized market entry game with fixed capacity. *Experimental Economics 5*, 253 – 272.