Self-organization in economics: The minority game

Willemien Kets*

Abstract

This paper surveys the minority game literature. The minority game is a simple congestion game that has been analyzed using methods from statistical physics. The present paper aims to complement the existing literature, which mainly focuses on the statistical mechanics of the game, by relating the minority game literature to the economics literature, in particular the literature in economics on learning in games. In addition, the paper aims to demonstrate the novel perspectives this literature provides to understand the behavior of players in congestion games.

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 thus 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 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, as in e.g.

* Address: Department of Econometrics and Operations Research, Tilburg University, P.O. Box 90153, 5000 LE Tilburg, The Netherlands. E-mail: w.kets@uvt.nl. Tel: +31-13-466 2478. Fax: +31-13-466 3280.

Meyer et al. (1992).[1] Finally, congestion games often have many equilibria, so that players also face the difficulty of coordinating on the same equilibrium.

Surprisingly, several experimental studies have shown that players learn to coordinate in congestion games. For instance, interacting players rapidly achieve a "magical" degree of tacit coordination in market entry games, which is accounted for on the aggregate level by the Nash equilibrium solution (Kahneman, 1988; Rapoport, 1995; Sundali et al., 1995; Rapoport et al., 1998; Erev and Rapoport, 1998). Also Meyer et al. (1992) find that average behavior is consistent with equilibrium in their experiments on binary allocation games. Similarly, in experiments on route-choice behavior, Selten et al. (2003) find that the mean number of drivers on the different roads is very close to the equilibrium number.

In these studies, however, regularities on the aggregate level generally hide 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 (e.g Meyer et al., 1992; Erev and Rapoport, 1998; Selten et al., 2003; 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) and Duffy and Hopkins (2005) find that providing players with information on other players' actions may actually lead to *lower* average payoffs.

These features are hard to explain using standard learning models. This paper surveys the literature on the minority game, a simple congestion game with remarkably rich behavior. The learning model proposed in this literature may help to explain the puzzling findings mentioned above. In this literature, it has been shown that it is possible to reconcile equilibrium play at the aggregate level and non-equilibrium behavior at the individual level. Moreover, the literature proposes a variant of the basic reinforcement model (Roth and Erev, 1995; Erev and Roth, 1998) that offers a new and behaviorally plausible way to model the effect of information on other players' actions.

The minority game is based on the El Farol bar problem, first described by Brian Arthur in his seminal paper (Arthur, 1994). In the El Farol bar problem, 100 people decide independently each week whether to go the El Farol bar in Santa Fe, New Mexico. As space is limited, people only prefer going if not more than 60 people go. It is not known in advance how many will go this week. As there is no natural model to predict opponents' play, agents have to rely on inductive reasoning rather than on deductive logic to decide on their actions (cf. Rapoport et al., 1998). Arthur shows that at the aggregate level, equilibrium is attained: on average, 60 people go to the bar each week. However, individual behavior keeps fluctuating. Some weeks people go, in other weeks they stay home. This finding contradicts standard learning models that predict convergence in games similar to the minority game to an asymmetric pure strategy Nash equilibrium in which some players always go to the bar and others stay home always (e.g. Duffy and Hopkins, 2005), but is in accordance with experimental evidence on congestion games

[1] Indeed, Falk and Kosfeld (2003) find that strategic asymmetry (players have to choose different actions in equilibrium) and asymmetry with respect to payoffs inhibits play of the Nash equilibrium in experiments on network formation. Strategic asymmetry hinders attainment of equilibrium, while asymmetry in payoffs makes that players are unwilling to maintain a given equilibrium once it is reached.

(e.g. Erev and Rapoport, 1998; Selten et al., 2003).

The El Farol bar problem is simplified and further formalized in the minority game literature, initiated by Challet and Zhang (1997).[2] In this game, an odd number of agents N have to choose repeatedly between two alternatives. Those who have chosen the minority side get a positive payoff, other agents get a zero or negative payoff. After agents have chosen their actions, agents learn the winning side. On the basis of a very simple learning model, based on the model of Arthur (1994), agents learn to coordinate to differentiate. Consistent with experimental results on congestion games, the average size of the minority is equal to the equilibrium value, while individual agents do not play Nash strategies. It is shown that this behavior can be described by the formation of so-called crowds and anti-crowds (Challet and Zhang, 1998; Johnson et al., 1998; Hart et al., 2001). This is in accordance with the experimental results of Erev and Rapoport (1998), Selten et al. (2003) and Chmura and Pitz (2004), who conjecture that coordination may be achieved in their experiments by some players choosing a certain response mode, as Selten et al. (2003) term it, while others best-reply to that. However, in contrast with these models, that postulate the existence of certain response modes, the response modes in the minority game literature arise endogenously, as a result of lower-level heterogeneity.

This important result is only a first step in the minority game literature. The literature investigates many issues, such as efficiency and the role of information, which may help answering the questions posed by the experimental literature on congestion games and formulated succinctly by Brian Arthur (1994): how do agents learn to coordinate to differentiate?

This paper surveys the literature on minority games. In particular, it relates the literature on minority game to the literature in economics on learning in games. We will not cover the whole field, as an enormous amount of work on the minority game has been done (an extensive collection of papers on the minority game can be found at `http://www.unifr.ch/econophysics/minority/`). In particular, we will not go into the statistical mechanics of the game.[3] Also, we will not discuss the many extensions of the standard model.[4] In a forthcoming paper (Kets, 2005), the results from the minority game literature will be compared with predictions from standard learning models and with experimental results on congestion games.

The outline of this paper is as follows. We introduce the minority game in Section 2. We discuss the stage game in Section 2.1. We then introduce the learning model in

[2] Other formalizations of the El Farol bar problem include Johnson et al. (1998), Challet et al. (2003), Franke (2003), and Zambrano (2004).

[3] See Moro (2003) for an introduction to the field; see Challet et al. (2004) for a collection of the main papers from the minority game literature; see Coolen (2005) for a thorough mathematical treatment.

[4] Examples include models in which agents can choose to participate in the market (Challet et al., 2001; Challet and Marsili, 2003), models that allow for multiple choice instead of binary choice (see e.g. Chow and Chau (2001)), models that explicitly model supply and demand (see e.g. Savit et al. (2003)) and models in which badly performing agents are weeded out by some evolutionary mechanism (see e.g. Zhang (1998); Li et al. (2000a,b)). An important class of models extends the minority game model to include different type of traders, making it a more realistic model to describe financial markets (Zhang, 1999; Challet et al., 2000).

Section 2.2. Section 2.3 and 2.4 relate the model to the economics literature by discussing the modelling of strategies and the learning model in more detail. In Section 3, the main theoretical results from the minority game literature are discussed. Section 4 concludes.

2 The minority game model

In the minority game literature, the term 'minority game' commonly refers not solely to the game itself (players, strategies, payoffs), but to the combination of the game, the learning model and a set of rationality assumptions. Here, we do not follow that convention to avoid confusion when relating the minority game literature to the game-theoretic literature. We reserve the term 'minority game' to refer to the game itself, while we use the term 'minority game model' to denote the combination of the game, the learning model and the associated rationality assumptions, i.e. what is termed 'the minority game' in the minority game literature.

In this section, we discuss the minority game model, and relate it to the economics literature. In Section 2.1, we discuss the stage game and the game's Nash equilibria. In Section 2.2, we introduce the minority game model. In Sections 2.3 and 2.4, we discuss the modelling of strategies in the minority game model and the learning model in more detail.

2.1 The stage game and its Nash equilibria

The minority game is a simple game, in which an odd number of players have to choose between two actions, for instance, players either go to the bar or stays 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), denote the set of players by $\mathcal{N} = \{1, \ldots, N\}$, with $N = 2k + 1$ for some fixed $k \in \mathbb{N}$. Each player $i \in \mathcal{N}$ has a set of pure strategies $B_i = \{-1, +1\}$: agents have to choose between two options. The set of mixed strategies of player i is denoted by $\Delta(B_i)$, with elements $\lambda_i = (p_i, 1 - p_i)$, where $p_i \in [0, 1]$ denotes the probability with which player i plays $b = +1$. Denote a mixed strategy profile of players other than $i \in \mathcal{N}$ by $\lambda_{-i} \in \times_{j \in \mathcal{N} \setminus \{i\}} \Delta(B_j)$. With a slight abuse of notation, let $\mathbf{b} = (b_i)_{i \in \mathcal{N}}$ denote the random variable of the joint action profile in $B = \times_{i \in \mathcal{N}} B_i$, i.e. $b_i \in \{-1, +1\}$ is a realization of player i's choice when he plays according to the mixed strategy λ_i. With each action $b \in \{-1, +1\}$, a function

$$f_b : \{1, \ldots, N\} \to \mathbb{R}$$

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

$$u_i(\mathbf{b}) = f_{b_i}\left(|\{j \in \mathcal{N} : b_j = b_i\}|\right), \tag{2.1}$$

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

The characteristic feature of a minority game is that a unilateral deviation from a majority to a minority always pays off:

$$\forall a, b \in \{-1, +1\}, a \neq b, \forall n \in \{k+2, \ldots, 2k+1\}: \quad f_a(n) < f_b(2k+2-n). \tag{2.2}$$

The minority game is a congestion game (Rosenthal, 1973) and hence a finite exact potential game as defined by Monderer and Shapley (1996) (Tercieux and Voorneveld, 2005).

The function $f(\cdot)$ can have several forms. The payoff functions for both options are typically assumed to be identical:

$$f_{-1} = f_{+1}, \tag{2.3}$$

and $f(\cdot)$ often satisfies the condition

$$\forall a \in \{-1, +1\}, \forall n, n' \in \{1, \ldots, N\}: \quad n > n' \Rightarrow f_a(n) \leq f_a(n'). \tag{2.4}$$

A commonly used form for f 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 $A = \sum_{i\in\mathcal{N}} a_i$ for a given action profile $\mathbf{a} = (a_i)_{i\in\mathcal{N}}$, with $a_i \in \{-1, +1\}$. A player $i \in \mathcal{N}$ is then assigned the payoff

$$u(a_i) = -a_i g\left(\sum_{i\in\mathcal{N}} a_i\right), \tag{2.5}$$

with $g(\cdot)$ an odd function, i.e. $g(-x) = -g(x)$, with $g(x) > 0$ for $x > 0$ (Moro, 2003). In our notation:

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

Common choices include

$$g(x) = x/N \tag{2.6}$$

and

$$g(x) = \text{sign}(x). \tag{2.7}$$

In the following sections, we focus on payoff functions of the form (2.6), as the linearity of the payoff function simplifies the learning dynamics considerably (Marsili, 2004).

The stage game has a large number of Nash equilibria.[5] The set of Nash equilibria is very sensitive to the precise formulation of the payoff function. However, a few general statements can be made. As shown by Tercieux and Voorneveld (2005, Prop. 4.3(b)), a pure strategy profile $\mathbf{a} = (a_i)_{i\in\mathcal{N}}$ is a Nash equilibrium of a minority game G (in particular, under the conditions (2.2) and (2.3)) if and only if there is an alternative $a \in \{-1, +1\}$ such that exactly k players choose k. There are $2\binom{N}{k}$ of such asymmetric pure strategy equilibria. If payoffs are of the form (2.5), total payoffs are negative.

It is not possible to give a complete characterization of the set of mixed strategy Nash equilibria for general payoff functions. It can be shown (Kets, 2005), however, that a

[5] For a discussion of the uniform equilibria of the infinitely repeated game, see Renault et al. (2005).

mixed strategy profile λ such that there is a group of q players playing the pure strategy $a = -1$, a group of q players playing $a = +1$ and a group $\mathcal{C}$ of $2(k-q)+1$ players playing a strictly mixed strategy, is a Nash equilibrium if and only if each player $j \in \mathcal{C}$ plays according to the mixed strategy $\lambda_j = \left(\frac{1}{2}, \frac{1}{2}\right)$. If payoffs are of the form (2.5), expected total payoffs are zero in these equilibria. In a game with $N = 2k+1$ players, there are

$$2\sum_{q=0}^{k} \binom{2k+1}{2q}\binom{2q}{q}$$

of such equilibria. It is easy to show that there is a unique symmetric equilibrium in mixed strategies, in which each player randomizes over both options with equal probability. Moreover, it is easy to show that there is no equilibrium in which there is a group of $|\mathcal{C}| = 2$ agents playing a strictly mixed strategy.

Depending on the form of the payoff function, there may be many more equilibria. In all mixed strategy equilibria, the expected number of players choosing option $a \in \{-1, +1\}$ equals $k + \frac{1}{2}$ when the payoff of an option $a \in \{-1, +1\}$ is linear in the number of players n_a choosing that option, as in (2.6).

The large number of Nash equilibria in the minority game raises the question of equilibrium selection. Without pre-play communication, players do not have enough information to implement a pure strategy Nash equilibrium (see also Menezes and Pitchford, 2006). Players could use common knowledge of rationality and symmetry to deduce and hence select the symmetric mixed strategy equilibrium (cf. Meyer et al., 1992). However, in addition to the coordination problem, players also face an incentive problem. Players get a payoff of zero in the mixed strategy equilibrium, when payoffs are of the form (2.6), but they can earn a positive payoff if they manage to outsmart the other players. It is quite likely that players will depart from pure randomization in an attempt to outsmart the other players. As proposed by Meyer et al. (1992) and Arthur (1994), players may try to find patterns in the play of others when the game is played repeatedly.[6] The learning model proposed in the minority game literature is a way of formalizing this notion. We discuss this model in the next section.

2.2 Learning to play the minority game

In the minority game model, the stage game defined in Section 2.1 is played repeatedly. It is assumed that after each round of play of the stage game, the agents learn the winning group, $\text{sign}[A(t)]$, where $A(t) = \sum_{i=1}^{N} a_i(t)$ and $a_i(t) \in \{-1, +1\}$ is the action taken by player i in round t. The information set of each player is thus restricted to include only aggregate 'market' statistics, as in Meyer et al. (1992). Furthermore, it is assumed that agents only retain the sequence of the last m outcomes. Hence, in round t, players observe the m most recent outcomes $h_m^t = (\text{sign}[A(\tau)])_{t-m \leq \tau < t}$.

[6] Indeed, this is the strategy suggested in "The Official Rock Paper Scissors Strategy Guide" (Walker and Walker, 2004) for the Rock-Paper-Scissors game. The Rock-Paper-Scissors game is similar to the minority game in that players can obtain a higher payoff than the payoff in the symmetric mixed strategy Nash equilibrium if they manage to outsmart the other players.

Information			Action			
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 behavioral strategies with $m = 3$ and $n_S = 4$ for some agent $i \in \mathcal{N}$.*

A *behavioral strategy* $s_i \in S_i$ for player $i \in \mathcal{N}$ assigns to each information set $h_m^t \in \mathcal{H}_m = \{\{x_k\}_{k=1}^m | x_k \in \{-1, +1\}\}$ an action $a \in \{-1, +1\}$.[7] There are thus 2^{2^m} possible behavioral strategies: there are 2^m possible signals h_m of length m, and for each signal, there are two possible actions. For each agent i, a behavioral strategy $s \in S_i$ thus prescribes which action $a_i(t) = s(h_m^t) \in \{-1, +1\}$ to take, for a given history of play h_m^t at time t. Note that for $m = 0$, we have the standard repeated stage game, and the set of behavioral strategies S_i coincides with the set of strategies B_i of the stage game.

An important assumption in the minority game model is that agents are endowed only with a random subset of $n_S \geq 2$ behavioral strategies of all possible behavioral strategies. In the standard model, this subset is fixed for each agent, and all agents are endowed with the same number of behavioral strategies. An example of such a subset of behavioral strategies for $n_S = 4$ and $m = 3$ is given in Table 1.

When faced with a history h_m^t, an agent has to choose which of his n_S behavioral strategies to use in the next round. It is assumed that each agent i keeps a virtual score $p_i^\ell(t)$ for each behavioral strategy $s_i^\ell \in S_i$ that reflects that behavioral strategy's past performance. The virtual score of each behavioral strategy is updated after each round, *regardless of whether the behavioral strategy has been used or not.* When the behavioral strategy would have correctly predicted the winning side, the virtual score is increased, otherwise it is decreased:

$$p_i^\ell(t+1) = p_i^\ell(t) - s_i^\ell(h_m^t) \cdot \frac{A(t)}{N} \tag{2.8}$$

with $\ell \in \{1, \ldots, n_S\}$ and $A(t) = \sum_{i \in \mathcal{N}} a_i(t)$, and where we have used Equation (2.6) for the payoffs.

[7] Note that our definition of a behavioral strategy differs somewhat from the common definition in game theory. In the standard definition, a behavioral strategy for player $i \in \mathcal{N}$ in a given extensive form game specifies, for every information set $h \in \mathcal{H}_i$ and action $a \in A_h$ a probability $\gamma(a, h) \geq 0$, with $\sum_{a \in A_h} \gamma(a, h) = 1$ for all $h \in \mathcal{H}_i$, with $\mathcal{H}_i$ the collection of player i's information sets and A_h the set of possible actions at information set h (e.g Mas-Colell et al., 1995).

The probability that agent $i \in \mathcal{N}$ chooses the behavioral strategy $s_i^\ell \in S_i$ in the next round is given by the logit choice rule (Cavagna et al., 1999):

$$\text{Prob}\{s_i(t) = s_i^\ell\} = \frac{e^{\beta \cdot p_i^\ell(t)}}{\sum_j e^{\beta \cdot p_i^j(t)}}. \tag{2.9}$$

Actions, outcomes and performance are thus linked by a complex feedback system, as illustrated in Figure 1.[8]

Probabilistic rather than deterministic choice rules are well supported by experimental evidence on learning in games (see e.g. Mookherjee and Sopher (1994, 1997), Ho and Weigelt (1996) and Fudenberg and Levine (1998)). The parameter β can be interpreted as the sensitivity of choice to marginal information. The early literature focuses on the case $\beta \to \infty$, i.e. play is fully deterministic in the sense that players choose the behavioral strategy with the highest virtual score. Allowing for $\beta < \infty$ adds noise at the individual level as well as it introduces additional heterogeneity. When $\beta \to \infty$, all agents endowed with a certain behavioral strategy keep the same virtual score for that behavioral strategy. By contrast, for a finite β, agents differ in their ranking of behavioral strategies, as their endowment of behavioral strategies determines the denominator of Equation (2.9). Perhaps surprisingly, this added heterogeneity and noise actually improves collective performance, as discussed in Section 3.1.

An important point to note is that for all behavioral strategies s_i^ℓ, whether they are actually played or not, players use the same updating rule (2.8). In particular, $A(t)$ does not depend on the behavioral strategy updated. In other words, players do not take the effect of their action on the aggregate outcome $A(t)$ into account. In determining the virtual score of behavioral strategy $s_i^\ell \in S_i$, agents only consider whether this behavioral strategy would have predicted the actual outcome correctly, neglecting the question whether playing s_i^ℓ would have affected the outcome. This assumption has far-reaching implications. We will come back to this issue in Section 2.4.

Agents are thus assumed to be boundedly rational in the minority game model. They have a limited memory m, and they only use a subset of all possible behavioral strategies ($n_S \leq 2^{2^m}$). In addition, agents behave as if they were playing against some exogenous signal $A(t)$, rather than against $N-1$ opponents. Agents thus do not take the effect of their own action into account. The minority game model has some features that may seem awkward at first sight. In the following two sections, we relate the minority game literature to the economics literature.

[8] The dependence of actions on history h_m introduces a feedback loop. As h_m is not Markovian, this leads to complicated dynamics. Cavagna (1999) shows, however, that for many practical purposes the endogenous information h_m may well be replaced by some exogenous signal drawn independently for each time step from some distribution. This exogenous signal can be interpreted as common information (market shocks) to which agents react.

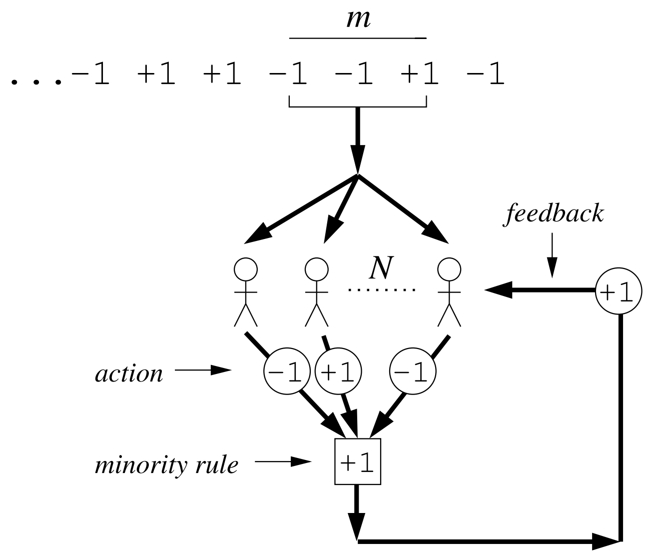

Figure 1: *A schematic overview of the minority game model. Agents observe the recent outcomes, and choose a behavioral strategy with probability specified in Equation* (2.9)*, resulting in an action* $a_i(t) \in \{-1, +1\}$*. The collective action determines the winning side through the minority rule; this information is then fed back to the agents and adds to the sequence of outcomes. Figure taken from Moro (2003).*

2.3 Strategies and heterogeneity

In the minority game model, players base their action on the recent past, trying to discern patterns in their opponents' behavior. Arthur (1994) proposes that agents condition their decision to go to a bar on attendance levels in the previous weeks. Arthur employs the terms 'predictor' or 'hypothesis' rather than strategy: if the bar has been crowded for the last three weeks, I expect it to be crowded next week also. These mental models are mapped into actions: if I expect the bar to be crowded, I will not go. The behavioral strategies in the minority game model are a concise way of modelling this notion.

The assumption that players base their strategy on the history of some aggregate statistic, rather than on the history of their opponents' actions has been discussed by Meyer et al. (1992) in the context of market allocations. In an anonymous allocation game with a large number of players, this is a natural assumption. The history of public signals may then help players to anticipate their opponents' play in the next round, as suggested by Arthur (1994). Indeed, Selten et al. (2003) find in their experiments on route choice behavior that some subjects employ a direct response mode, while others choose the contrarian response mode when choosing between two roads. Players that play according to the direct response mode revise their choice if the road of their choice turned out to be congested, while players that employ the latter mode stick to their choice in that case, as they expect others to switch. Interestingly, providing players with information on their opponents' actions does not seem to change this considerably, suggesting that agents react merely to the public signal. Hence, the public signal h_m^t is translated into beliefs on opponents' play in the next round.

This way of modelling raises the question which behavioral strategies to include in the model. If all possible behavioral strategies are included in the learning model, i.e. if strategies are allowed to depend on the past T rounds of play, where T can be varied from zero up to the full length of history of play, the strategy space becomes huge already for very simple games. Many different behavioral strategies are conceivable, as illustrated by the list of examples in Arthur (1994). On the other hand, in the words of Erev and Roth (1998, p. 873), simply selecting a subset of all possible behavioral strategies is "like parameter fitting in a model with an enormous number of parameters".

If only a subset of all possible behavioral strategies is to be included in the model, a natural choice would be to include behavioral strategies that reflect beliefs about other players' actions, based on recent outcomes, as is done in for instance Erev and Rapoport (1998) and Selten et al. (2003). Similarly, in models of cognitive hierarchy (see e.g. Stahl and Wilson (1995) and Camerer et al. (2004)),[9] players form beliefs about the cognitive abilities of their opponents, and best-rely to that. Indeed, cognitive hierarchy models have been quite successful in explaining the 'magical' coordination observed in experiments on the market entry game (Camerer et al., 2004).[10]

Whereas Erev and Rapoport (1998) and Selten et al. (2003) and models of cognitive hierarchy make explicit assumptions on the types of strategies players use, the minority game model places no restriction on the types of behavioral strategy used. At first sight, this may seem to be a weak point of the model, as behavioral strategies do not need to have a sensible interpretation in the minority game model. However, it can be shown that no matter what behavioral strategies players are endowed with, players will self-organize into several crowds and anti-crowds, groups of players whose actions cancel out (see Challet and Zhang, 1998; Hart et al., 2001, see also Section 3.3).

Hence, the strong point of the minority game model is exactly that no assumptions regarding the form of heterogeneity are needed. In games such as the minority game, whether a behavioral strategy is reasonable *only* depends on the behavioral strategies employed by others. The contrarian response mode of Selten et al. (2003) is only sensible if there are players that employ the direct response mode and vice versa. Conversely, *any* behavioral strategy, whether it has a sensible interpretation or not, will work if opponents use behavioral strategies that recommend them to take the opposite action.

The minority game differs in this respect from games such as the p-beauty contest[11]

---

[9] In models of cognitive hierarchy, it is assumed that there are players that differ in their depth of reasoning. So-called level-0 players simply randomize over strategies, while level-1 players best respond to this, and so on.

[10] An interesting real-world illustration of the success of cognitive hierarchy players in the Rock-Paper-Scissors game can be found in The New York Times (2005).

[11] In the p-beauty contest, players have to choose a number in a certain interval. Players have to guess what the average choice will be; the player that picks the number that is closest to some fraction $\varphi < 1$ of the average choice will win. Suppose players have to choose a number between 0 and 100, and will win with their choice is closest to $\varphi = 2/3$ of the average choice. Then nobody will choose a number higher

(Keynes, 1936, p. 156). Playing the p-beauty contest with game theorists is not much fun, as all will choose the equilibrium action. By contrast, in the minority game, players base their actions on their beliefs about other players' actions, who in turn base their actions on ..., etcetera, giving rise to a never-ending recursion of actions and beliefs. In the p-beauty contents, on the other hand, the recursion ends at a well-defined limit point, the equilibrium. In both games the payoff to an action depends solely on the action of opponents. However, whereas the assumptions of perfect rationality, common knowledge of the rationality of the players, and utility maximization makes the opponents' actions predictable in the p-beauty contest, these assumptions get you nowhere in the minority game. If one could make such a prediction in the latter game, it would immediately be invalidated: if all think that $a = -1$ will be the minority choice, then all will choose that action. Therefore, the minority game model's agnosticism on the type of behavioral strategies that players use may well provide a more realistic model of players' reasoning processes than the more restrictive assumptions employed in different learning models.

Which behavioral strategy agents choose from the set of behavioral strategies they are endowed with, is determined by the virtual score of each behavioral strategy, as modelled in the learning model. In the next section, we discuss the learning model in the minority game literature in relation with other learning models in the literature.

2.4 Bounded rationality and the law of simulated effect

In the minority game model, agents learn by keeping a virtual score for each behavioral strategy in their subset. The learning process proposed in the minority game literature (Equation (2.8)) 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 not played. In the basic reinforcement learning model, the scores of these strategies are not updated. By contrast, in the minority game model, the scores of all behavioral strategies are updated in every period, as in hypothetical reinforcement learning and stochastic fictitious play (Fudenberg and Levine, 1998).

The assumption that agents also consider the payoffs to strategies 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 results that agents 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 agents to play according to the act of simulated effect, they need more information than for standard reinforcement learning. For instance, in market entry games, players need to know the payoff rule as well as the number of entrants in order to play according to fictitious play, in addition to their own payoff. By contrast, they only

than 66.6667, so nobody should pick a number higher than $2/3 \cdot 66.6667$, and so on. The equilibrium choice is thus 0 in this case.

need to know their own payoff to play according to the standard reinforcement learning model.

In many cases, this information may not be available. Alternatively, calculating foregone payoffs of strategies not used may be too hard in complex situations for boundedly rational agents. In the minority game model, 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, agents can account for foregone payoffs of strategies not used, while they do not need detailed information on the number of players choosing each side.

At first sight, one may think that for a large number of players, it will not matter whether agents account for their own impact. However, due to the minority rule, there remains a systematic bias in the rewarding of behavioral strategies, even if the number of players goes to infinity. The reason is that the virtual score of a behavioral strategy that is currently played is systematically lower than that of the behavioral strategies that are not used. These latter behavioral strategies 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 actually played, so that they would have guessed wrong in reality. As the behavioral strategy that is actually played does not have this advantage, the behavioral strategies that are not played are systematically favored and hence results depend on whether agents take their market impact into account (Marsili et al., 2000; Marsili and Challet, 2001).

As discussed in Section 3.1, play does *not* converge to a Nash equilibrium of the game[12] if players do not take the effect of their own action on the aggregate outcome into account. Hence, updating the scores of strategies not used by players without information on opponents' actions only seems to work counterproductively, as it is known that play converges to a Nash equilibrium in the basic reinforcement learning model of Roth and Erev (1995) in games similar to the minority game (Duffy and Hopkins, 2005). Indeed, when the updating rule (2.8) is modified so that the effect of a player's own action is taken into account, i.e. if players learn according to the hypothetical reinforcement learning rule, the dynamics converge to one of the pure Nash equilibria of the game (Challet et al., 2000b; Marsili et al., 2000; De Martino and Marsili, 2001).

Fortunately, agents do not need to have full information to correct for the bias in their score keeping. As shown by Marsili et al. (2000) and De Martino and Marsili (2001), it suffices if agents correct their virtual scores by over-rewarding the behavioral strategy they are currently using. It is easy to see that it is possible to 'tune' this correction factor in such a way that the modified version of (2.8) amounts to the hypothetical reinforcement learning rule for payoff functions of the form (2.6). However, over-rewarding the behavioral

[12] In the following, if we speak of Nash equilibria, we refer to the Nash equilibria of the repeated minority game for which $m = 0$, i.e. $S_i = \{-1, +1\}$ for all players $i \in \mathcal{N}$, to remain close to the game-theoretic literature on learning in games. When $m = 0$, each player is endowed with all possible behavioral strategies: $n_S = 2^{2^0} = 2$. It is also possible to consider Nash equilibria for the case $m > 0$, both for the case where the endowment set of players is a proper subset of the set of all possible behavioral strategies as for the case $n_S = 2^{2^m}$ (Marsili et al., 2000; De Martino and Marsili, 2001).

strategy currently played by an amount $\eta \leq \eta^*$, where η^* is the correction factor that makes the modified version of (2.8) equal to the hypothetical reinforcement learning rule, is always advantageous, both in terms of individual payoffs as in terms of total payoffs (Marsili et al., 2000). Hence, also agents with limited information can learn to play according to a Nash equilibrium in the minority game model, under the behaviorally plausible assumption that they also consider foregone payoffs of strategies not used.

The minority game model thus combines features from several learning models in the literature on learning in games. This makes that the minority game model makes distinctly different predictions than other learning models in the literature. To these prediction we now turn.

3 Results

In this section, we discuss the main theoretical results. 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 the minority game model. In Section 3.3, we show that the behavior of agents can be interpreted in terms of crowds and anti-crowds, and discuss the implications for efficiency.

3.1 Volatility and attendance

Typically, dust never settles down in the minority game model: the aggregate 'attendance' $A(t) = \sum_{i \in \mathcal{N}} a_i(t)$ keeps fluctuating, as can be seen in Figure 2. As the game is symmetric, the average of $A(t)$ will be 0 in the steady state, as borne out by simulations (see e.g Challet and Zhang, 1997, 1998; Johnson et al., 1998; Manuca et al., 2000). More interesting is the behavior of the variance $\sigma^2 \equiv \langle A^2 \rangle$, where $< \cdot >$ denotes the (time) average of a quantity in the stationary state. The variance, or volatility is a measure of the degree of efficiency achieved in a population. The higher the variance, the larger the aggregate welfare loss: large fluctuations imply that the size of the minority is only small. When payoffs are linear in $A(t)$, as in Equation (2.6), this is easy to see: in that case, total payoffs are proportional to $-\sum_{i \in \mathcal{N}} a_i A(t) = -A^2(t)$.

It has been found that σ^2 is only a function of $\alpha = 2^m/N$ for a given value of n_S, the number of behavioral strategies of each agent (Savit et al., 1999). As can be seen in Figure 3, the volatility converges to the volatility exhibited in random play for $\alpha \to \infty$, i.e. the game in which each agents independently chooses one of the alternatives with equal probability. With a large number of players (α small), overall performance is much worse; in fact, the volatility is of the order of N^2, so that the size of the winning group is much smaller than $N/2$. 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, agents are able to coordinate their actions and perform better collectively than under random play. In the market formalism, this means that agents can exploit the available information to predict future market movements so

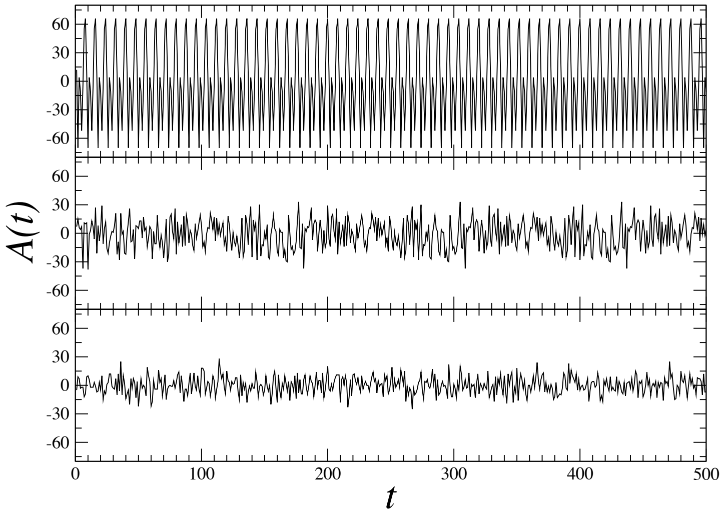

Figure 2: *Time evolution of the attendance* $A(t)$ *with* $g[A(t)] = A(t)$ *and* $N = 301$ *and* $n_S = 2$. *Panels correspond to* $m = 2, 7, 15$ *from top to bottom. Figure taken from Moro (2003).*

that the aggregate welfare loss σ^2 is reduced. Note that this is not the result of some form of cooperative behavior of the agents.[13] Agents are selfishly maximizing their own return, and in doing that, they come closer to global efficiency. However, coordination is not complete. In the socially efficient outcome, the minority would consist of k agents, and agents would have a 'success rate' of 50%, so that $\sigma^2/N = 1/N$. Players achieve a success rate close to 50% at $\alpha = \alpha_c$, although they never reach the social optimum, while their success rate is substantially lower for smaller values of α, as can be seen in the inset to Figure 3. By contrast, if players take the effect of their own action on the aggregate outcome, play converges to one of the pure strategy Nash equilibria of the game (Challet et al., 2000b; Marsili et al., 2000; Marsili and Challet, 2001; De Martino and Marsili, 2001).

Interestingly, global efficiency is enhanced for certain values of α when agents do not always choose the behavioral strategy s with the highest number of virtual points, i.e. when $\beta < \infty$ in Equation (2.9). It can be shown that for $\alpha < \alpha_c$ (the socially inefficient

[13] Agents could engage in a 'collusive' dynamic strategy aimed at maximizing joint payoffs. In the minority game, such a collusive strategy amounts to 'taking turns': players alternate in choosing sides, so that each player is in the minority k periods out of $2k + 1$ periods and there are always k players in the minority. However, this requires a complicated coordination mechanism, especially if the number of players is large. In addition, it asks for a sophisticated punishment strategy as players can only observe the aggregate action and not the individual actions (Renault et al., 2005). Whether agents engage in such collusive strategies is an open question. Duffy and Hopkins (2005) find no support for play of such a dynamic strategy in their experiments on market entry games, but Helbing et al. (2005) find evidence of coherent oscillatory behavior in experiments on a congestion game, consistent with collusive dynamic strategies. However, the number of subjects in their experiments is only small, while the minority game literature focuses on the case of a large number of players.

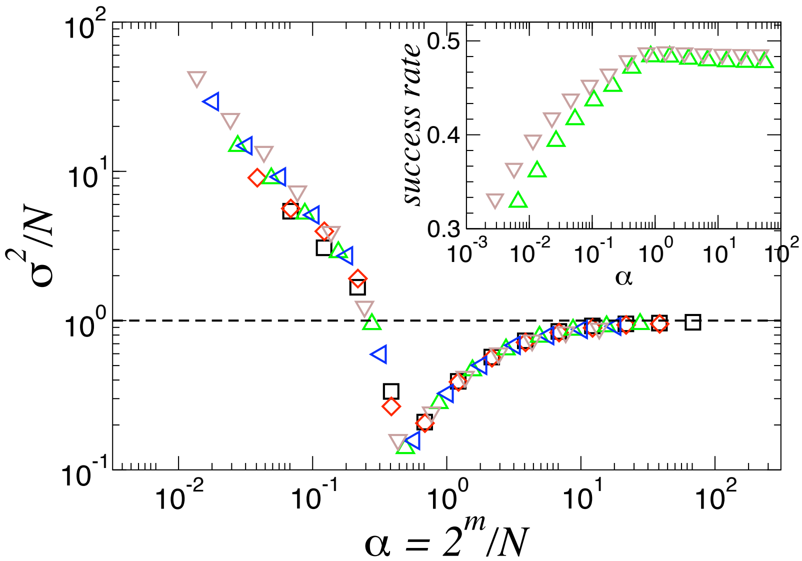

Figure 3: *Volatility as a function of the order parameter* $\alpha = 2^m/N$ *for* $n_S = 2$ *and different number of agents* $N = 101, 201, 301, 501, 701$ *(□, ◇, △, ◁, ▽, respectively). The critical value* α_c *is the value of* α *for which the volatility is at a minimum. Inset: Agent's mean success rate as function of* α*. Figure taken from Moro (2003).*

regime), volatility *decreases* if the 'noise level' *increases*. For $\alpha > \alpha_c$, the value of β does not affect the level of volatility (Cavagna et al., 1999; Challet et al., 2000a; Bottazzi et al., 2001; Marsili, 2004). This result is not so surprising, however, if one recalls that in the minority game model, rational agents herd in the socially inefficient regime. For $\alpha < \alpha_c$, the dimension of the space of behavioral strategies is small relative to the number of agents: agents have to 'crowd' at a limited number of behavioral strategies, leading to a large number of agents choosing the same alternative (see also Section 3.3). Setting $\beta < \infty$ is equivalent by making that agents update their scores of behavioral strategies more slowly. Hence, a finite β thus acts as a brake against overreaction (Bottazzi et al., 2001). This result is reminiscent of the findings of Goeree et al. (2004) who show that payoff-dependent noise in the decision process is able to break the cascades that would result otherwise in a social learning model.

The minority game model is thus characterized by competition and coordination: agents coordinate to differentiate. Agents compete in trying to exploit asymmetries in the games outcome, but at the same time, they try to reduce volatility, as volatility harms all agents. In trying to exploit the asymmetries in the outcome, agents in fact minimize the information present in the game's history, rather than the volatility (Marsili et al., 2000; Marsili and Challet, 2001). Hence, there is a tension between competition and coordination. These two are intimately linked in the minority game model, as are information and efficiency. We will discuss this in more detail in the next section.

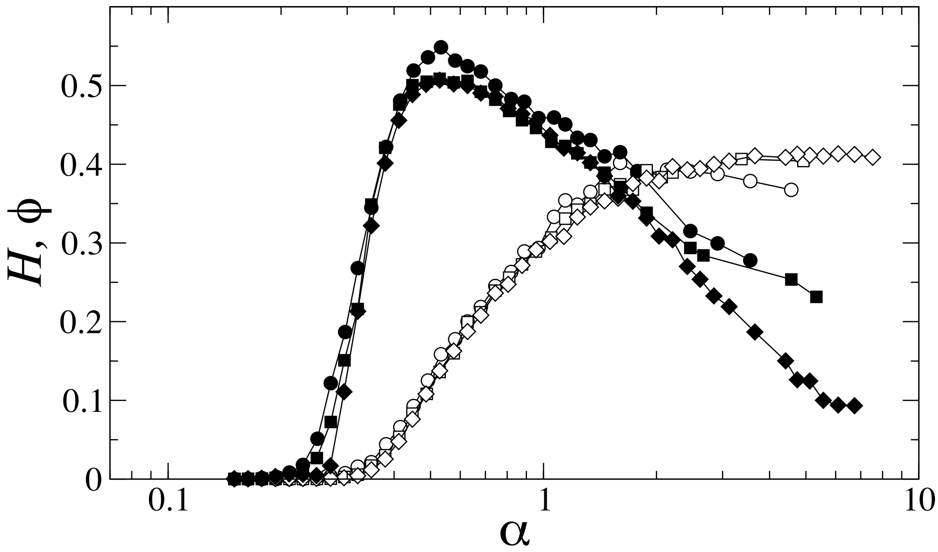

Figure 4: *Information H (open symbols) and fraction of frozen agents ϕ (full symbols) as a function of the control parameter $\alpha = 2^m/N$ for $n_S = 2$ and $m = 5, 6, 7$ (circles, squares and diamonds, respectively). Figure taken from Moro (2003).*

3.2 Information and efficiency

As shown 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 virtual scores of their behavioral strategies, implying that there is some information in these values (Challet and Zhang, 1998). This observation has 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)

$$H = \frac{1}{2^m} \sum_{\nu=1}^{2^m} < A | h_m = \nu >^2, \tag{3.1}$$

where the average is conditioned on play in the last m periods. It can be shown that agents in the minority game model minimize the degree of predictability. Depending on the value of α, they are more or less successful in doing that. At α_c, the system undergoes a phase transition from an informationally efficient and socially inefficient phase ($H = 0$, σ^2 large) to an information-rich and socially efficient phase ($H > 0$, σ^2 small). In the informationally efficient phase, players do worse than random players. By contrast, in the information rich phase, agents manage to coordinate and do better than random players.

At the phase transition, the symmetry between the two actions is broken. In the so-called symmetric phase ($\alpha < \alpha_c$), both actions are equivalent. Both actions are taken by the players with the same frequency. For $\alpha > \alpha_c$, one of the actions is preferred, i.e. the outcome is asymmetric. An asymmetry in the game's outcome represents an opportunity that could in principle be exploited. Hence, this is just a concomitant feature of the presence or absence of information in the history of play.

As an alternative to H, one could also consider the fraction of *frozen agents* as a measure of the order in the system (Challet and Marsili, 1999). Frozen agents are agents who never change their behavioral strategy in the stationary state (in the limit of $\beta \to \infty$ in Equation (2.9)). These agents thus have one behavioral strategy that outperforms all others. Note that this does not imply that these agents take the same *action* always: a behavioral strategy is a function of past play, hence the actions vary with h_m. As can be seen from Figure 4, this fraction ϕ is zero in the informationally efficient phase, while it first rises for intermediate values of α and then declines again when α goes to infinity. The intuition is that, in the informationally efficient phase, both actions are equivalent, so that there is little variation in the virtual scores of the different behavioral strategies. Agents thus switch behavioral strategies easily. For very large values of α, agents behave more or less randomly, so they switch behavioral strategies frequently. Only at intermediate values of α the fraction of frozen agents is large, as a considerable number of agents have a behavioral strategy that is superior to other behavioral strategies. Note that the success of a behavioral strategy depends on the behavioral strategies used by opponents: a behavioral strategy *per se* is not superior, it is the collective of behavioral strategies that is successful (Savit et al., 1999). In particular, it only pays to be predictable if others are predictable as well.

This phase transition between the informationally efficient and the information rich phase, or equivalently between the socially inefficient and the socially efficient phase, is central to the minority game model. At the phase transition, there is a qualitative change in collective behavior, while the principles behind the behavior of individuals remain unchanged. For all values of α, agents in the minority game model try to outsmart each other, but for low values of α, they are on average less successful.

Interestingly, this phase diagram may also serve to illustrate the mechanism behind market arbitrage. Suppose there is room for arbitrage. The 'minority market' performs better than it would with truly random players, and there is information in the history of play that can be exploited by market participants ($\alpha > \alpha_c$). Economic theory then predicts that agents will enter the market. This drives down $\alpha \propto 1/N$, until there is no information left in the market, i.e. to α_c, so that the market is hovering at the border of efficiency (Marsili, 2004).

3.3 Crowds and anti-crowds

The former sections have shown that the qualitative behavior of the system depends only on $\alpha = 2^m/N$, not on other variables such as n_S. Moreover, for some values of this parameter, agents 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 crowds and anti-crowds.

The minority rule forces agents to differentiate: if all agents choose the same behavioral strategy, all will loose. Agents want to be as far apart in the space of behavioral strategies as possible. However, there are only 2^{2^m} possible behavioral strategies for N agents. Hence, one would expect that agents succeed in differentiating if $N \ll 2^{2^m}$, while they behave more like a crowd when $N > 2^{2^m}$. So, one would expect a phase transition at $N \sim 2^{2^m}$, rather than at $N \sim 2^m$, as observed.

How can this be? Whereas there are 2^{2^m} different behavioral strategies, there are only 2^m *really* different behavioral strategies (Challet and Zhang, 1998; Johnson et al., 1998; Hart et al., 2001). The dimension of the space of behavioral strategies is 2^m, as there are 2^m different input signals. One can define a distance $D_m(s^1, s^2)$ on this hypercube, the so-called Hamming distance, which is a measure of the number of different components of a pair s^1, s^2 of behavioral strategies for a memory length m. The probability that two behavioral strategies do *not* prescribe the same action for a given history of play h_m is then proportional to their Hamming distance. Hence, the behavioral strategy with the highest virtual value is the one with the lowest probability of prescribing the action taken by other players.

Agents thus want to be different from the others. When is a behavioral strategy different from other strategies? Of course, for every behavioral strategy s there is only one behavioral strategy $\overline{s}$ which prescribes a different action in all possible circumstances, i.e. s and $\overline{s}$ are anti-correlated. However, also behavioral strategies that are uncorrelated with s are significantly different from s. It can be shown that the number of behavioral strategies that are really different is given by $2 \cdot 2^m / n_S$ (Challet and Zhang, 1998; Hart et al., 2001). Hence, α is proportional to the inverse of the density of the system in the reduced space of behavioral strategies, i.e. it determines how 'far' agents can be apart in the reduced behavioral strategy space.

This might seem to be a rather esoterical exercise without much practical use. However, this approach allows for an intuitive interpretation in terms of 'crowds' and 'anti-crowds'. If N_s agents use the same the same behavioral strategy s at time t, they act like a crowd. The $N_{\overline{s}}$ agents that use the anti-correlated behavioral strategy $\overline{s}$ at time t then act like an anti-crowd. If $N_s \approx N_{\overline{s}}$ for all anti-correlated pairs $(s, \overline{s})$ of behavioral strategies, then the actions of the crowds and the anti-crowds effectively cancel and the volatility will be small.

Hence, it would be optimal to have crowds and anti-crowds of similar sizes. However, this is not always possible, as the dimension of the space of behavioral strategies is fixed by the parameter m. Hence, agents can only be 'far apart' in terms of behavioral strategies if the number of agents is not too large relative to the dimension of the behavioral strategy space. For a given number of agents N, agents cannot differentiate if m is small, as the space of behavioral strategies is too crowded. The agents display herding behavior: for a pair of anti-correlated behavioral strategies $(s, \overline{s})$, almost all agents herd at one of them, with very few agents choosing the other. Hence, the actions of the anti-crowd do not cancel those of the crowd, so that σ^2 will be large.

For somewhat larger m (for N fixed), agents can differentiate, and the actions of the crowd and the anti-crowd effectively cancel. Hence, the system is quite successful collectively at intermediate values of α, although the minority rule prevents the system from attaining full efficiency (not all agents can be on the minority side). For a given h_m, the behavioral strategies of most players are uncorrelated, but a small share of agents uses behavioral strategies that are mutually anti-correlated. This 'coordinated avoidance' is beneficial for everybody, as it helps to get a more even division of players over both alternatives (Zhang, 1998).

Now, for very large m at fixed N, crowds will only be small, so that agents act more or less independently (Moro, 2003). However, the system still performs better than random players would. If the system would consists of N truly uncorrelated players, σ^2/N would equal one. However, in the minority game model, there is always a positive number of pairs of anti-correlated players, so that the agents' actions are never truly independent and σ^2/N is smaller than 1 (Challet and Zhang, 1998).

4 From economics to physics and back

Economists wonder at times whether all this work in physics is not just a lengthy exercise in physicists learning what economists know. It is certainly not [...]. Modern physics can offer much to economics. Not just different tools and different methods of analysis, but different concepts such as phase transitions, critical values, and power laws. And not just the analysis of pattern at stasis, but the analysis of patterns in formation. Economics needs this.
Brian W. Arthur (2004)

As observed by Arthur (2004), the reformulation of the El Farol problem in the minority game literature is not innocuous in the sense that it shifted the focus from the limitations of deduction to concepts like phase transitions, pattern formation and heterogeneity. As such, the research on the minority game literature might have more to offer than 'just' an analysis of a stylized congestion game. The minority game model, i.e. the game combined with the specific bounded rationality assumptions, is therefore of interest to economists, not because it is a toy model for the financial market, but in its own right. The core elements, frustration, individual heterogeneity and collective coordination, all important in economic problems, give rise to a number of fascinating properties.

Loosely speaking, the term frustration refers in this context to the fact that agents cannot be all on the minority side. In combination with the random assignment of strategy subsets to different agents, this makes that dust never settles:[14] agents keep switching between actions, in accordance with experimental evidence on congestion games (Erev and Rapoport, 1998; Selten et al., 2003; Bottazzi and Devetag, 2003, 2004; Chmura and Pitz, 2004).

[14] In physics, a disordered system is said to be frustrated if it is subject to numerous internal, conflicting constraints, so that a substantial fraction of the constraint cannot be satisfied by any arrangement of the constituting agents of the system (see e.g. Stein, 2003). The peculiar property of a system under frustration is that the system never gets stable in just one configuration since it costs no energy to switch from one onto another. It means in particular the existence of several ground states with exactly the same energy (Galam, 1996; Young, 1998).

The fact that the system exhibits frustration and that agents are heterogeneous in their endowment of strategies, has implications for the social and informational efficiency of the system. Agents' aim is to choose the strategy that on average minimizes their chance of being in the majority. As a result, strategies self-organize themselves so as to reduce collective dissatisfaction. Whether players are successful in self-organizing, depends on dimension of the strategy space and the number of players. With few players (relative to the dimension of the strategy space which is a function of memory length m), the collective of agents performs better than a collective of random players who randomly choose between alternatives. Moreover, the recent history h_m contains information that players can exploit. Hence, while the market is informationally inefficient, social inefficiencies are reduced. By contrast, when the number of players N increases (for constant memory length m), agents display herding behavior and overall performance is poor. In this case, informational efficiency is combined with social inefficiencies. Hence, information is intrinsically linked with efficiency in the minority game model: players act as if to maximize efficiency which is itself a measure of available information. Agents exploit all available information to their gain, until no information is left in the market.

All this illustrates the key point of the seminal work of Schelling (1978). Collective behavior need not be a simple summation of individual actions: it is the interaction between agents that gives rise to qualitatively different properties. In the minority game model, individual agents compete, while the collective displays a form of coordination. It is this combination of competition and coordination which gives rise to interesting behavior. Moreover, because of this nontrivial link between individual and collective behavior, it makes sense to consider collective properties like informational efficiency and volatility. In the words of Philip Anderson (1972) *"A gold atom is not yellow, malleable and shiny in any sense, metallicity only makes sense in terms of 10 to the 20th atoms; equally an amino acid is not alive"*.

The minority game model has thus many interesting new perspectives to offer for economics. At the same time, the minority game model is firmly rooted in the game-theoretical literature and the literature on learning, as shown in the previous sections. The dual value of the minority game model is perhaps best illustrated by the experiments on congestion games described throughout the paper: agents learn, act strategically and self-organize. It is in this synthesis of perspectives in which the strength of the minority game literature lies.

Acknowledgments

I would like to thank Ginestra Bianconi, George Ehrhardt, Doyne Farmer, Ben Ho, Matteo Marsili, Esteban Moro and Fernando Vega-Redondo for inspiring discussions. In addition, I would like to thank Esteban Moro for his kind permission for reproducing some of the figures from Moro (2003). Finally, I would like to thank Jan Potters, Dolf Talman and Mark Voorneveld for helpful comments and suggestions.

References

Anderson, P. W. (1972). More is different; broken symmetry and the nature of the hierarchichal structure of science. *Science 177*, 393–396.

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

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

Bottazzi, G. and G. Devetag (2003). A laboratory experiment on the Minority Game. *Physica A 324*, 124–132.

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., G. Devetag, and G. Dosi (2001). Adaptive learning and emergent coordination in Minority Games. *LEM Working Paper Series, Sant'Anna School of Advanced Studies 1999/24*.

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

Camerer, C., T.-H. Ho, and J.-K. Chong (2004). A cognitive hierarchy model of one-shot games. *Quarterly Journal of Economics 119*, 861–898.

Cavagna, A. (1999). Irrelevance of memory in the Minority Game. *Physical Review E 59*, R3783–R3786.

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. and M. Marsili (2003). Criticality and market efficiency in a simple realistic model of the stock market. *Physical Review E 68*, 036132.

Challet, D., M. Marsili, and G. Ottino (2003). Shedding light on El Farol. *Working paper (arXiv:cond-mat/0306445)*.

Challet, D., M. Marsili, and R. Zecchina (2000a). Comment on: Thermal model for adaptive competition in a market. *Physical Review Letters 85*, 5008.

Challet, D., M. Marsili, and R. Zecchina (2000b). Statistical mechanics of systems with heterogeneous agents: Minority games. *Physical Review Letters 84*, 1824–1827.

Challet, D., M. Marsili, and Y. C. Zhang (2000). Modeling market mechanism with Minority Game. *Physica A 276*, 284–315.

Challet, D., M. Marsili, and Y.-C. Zhang (2001). Stylized facts of financial markets and market crashes in Minority Games. *Physica A 294*, 514–524.

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*.

Chow, F. K. and H. F. Chau (2001). Multiple choice Minority Game. *Working paper (arXiv:cond-mat/0109166)*.

Coolen, A. (2005). *The mathematical theory of Minority Games: statistical mechanics of interacting agents*. Oxford: Oxford University Press.

De Martino, A. and M. Marsili (2001). Replica symmetry breaking in the Minority Game. *Journal of Physics A 34*, 2525–2537.

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.

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.

Falk, A. and M. Kosfeld (2003). It's all about connections: Evidence on network formation. *IZA Discussion Paper 777*.

Franke, R. (2003). Reinforcement learning in the El Farol model. *Journal of Economic Behavior and Organization 51*, 367–388.

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

Galam, S. (1996). Fragmentation versus stability in bimodal coalitions. *Physica A 230*, 174–188.

Goeree, J. K., T. R. Palfrey, B. W. Rogers, and R. D. McKelvey (2004). Self-correcting information cascades. *California Institute of Technology Social Science Working Paper 1197.*

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

Hart, M., P. Jefferies, N. F. Johnson, and P. M. Hui (2001). Crowd-anticrowd theory of the Minority Game. *Physica A 298*, 537–544.

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*. Forthcoming.

Ho, T.-H. and K. Weigelt (1996). Task complexity, equilibrium selection, and learning: An experimental study. *Management Science 42*, 659–679.

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

Johnson, N. F., S. Jarvis, R. Jonson, P. Cheung, Y. R. Kwong, and P. M. Hui (1998). Volatility and agent adaptability in a self-organizing market. *Physica A 258*, 230–236.

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. (2005). Coordinating to differentiate: The minority game. *CentER Discussion paper*. Forthcoming.

Keynes, J. M. (1936). *The General Theory of Employment, Interest and Money.* Macmillan Cambridge University Press.

Li, Y., R. Riolo, and R. Savit (2000a). Evolution in Minority Games. (I). Games with a fixed strategy space. *Physica A 276*, 234–264.

Li, Y., R. Riolo, and R. Savit (2000b). Evolution in Minority Games. (II). Games with variable strategy spaces. *Physica A 276*, 265–283.

Manuca, R., Y. Li, R. Riolo, and R. Savit (2000). The structure of adaptive competition in Minority Games. *Physica A 282*, 559–608.

Marsili, M. (2004). Statistical physics of interacting models: Minority Games. *mimeo*.

Marsili, M. and D. Challet (2001). Trading behavior and excess volatility in toy markets. *Advances in Complex Systems 4*, 3–17.

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.

Mas-Colell, A., M. D. Whinston, and J. R. Green (1995). *Microeconomic Theory*. Oxford: Oxford University Press.

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.

Mookherjee, D. and B. Sopher (1994). Learning behavior in an experimental matching pennies game. *Games and Economic Behavior 7*, 62–91.

Mookherjee, D. and B. Sopher (1997). Learning and decision costs in experimental constant sum games. *Games and Economic Behavior 19*, 97–132.

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.

Rapoport, A. (1995). Individual strategies in a market-entry game. *Group Decision and Negotiation 4*, 117–133.

Rapoport, A., D. Seale, I. Erev, and J. Sundali (1998). Equilibrium play in large group market entry games. *Management Science 44*, 119–141.

Renault, J., S. Scarlatti, and M. Scarsini (2005). A folk theorem for minority games. *Games and Economic Behavior*. In press.

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.

Savit, R., S. Brueckner, H. Van Dyke Parunak, and J. Sauter (2003). Phase structure of resource allocation games. *Physics Letters A 311*, 359–364.

Savit, R., R. Manuca, and R. Riolo (1999). Adaptive competition, market efficiency and phase transitions. *Phys. Rev. Lett. 82*, 22032206.

Schelling, T. C. (1978). *Micromotives and Macrobehavior*. W.W. Norton.

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 Seriess Analysis*. Wien-Würzburg: Physica-Verlag.

Selten, R., M. Schreckenberg, T. Pitz, T. Chmura, and S. Kube (2003). Experiments and simulations on day-to-day route choice-behaviour. *CESifo Working Paper 900*.

Stahl, D. and P. Wilson (1995). On players' models of other players: Theory and experimental evidence. *Games and Economic Behavior 10*, 218–254.

Stein, D. (2003). Spin glasses: Still complex after all these years? *Santa Fe Institute Working Paper 03-01-001*.

Sundali, J., A. Rapoport, and D. A. Seale (1995). Coordination in market entry games with symmetric players. *Organizational Behavior and Human Decision Processes 64*, 203–218.

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

The New York Times (2005). Rock, paper, payoff: Child's play wins auction house an art sale. (April 29).

Walker, G. and D. Walker (2004). *The Official Rock Paper Scissors Strategy Guide*. Simon and Schuster.

Young, A. (Ed.) (1998). *Spin Glasses and Random Fields*. Singapore: World Scientific.

Zambrano, E. (2004). The interplay between analytics and computation in the study of congestion externalities: The case of the El Farol problem. *Journal of Public Economic Theory 6*, 375–395.

Zhang, Y.-C. (1998). Modeling market mechanism with evolutionary games. *Europhys. News 29*, 51–54.

Zhang, Y.-C. (1999). Towards a theory of marginally efficient markets. *Physica A 269*, 30–44.