

Replicator dynamics with turnover of players

Jeppe Juul,¹ Ardeshir Kianercy,² Sebastian Bernhardsson,³ and Simone Pigolotti⁴

¹*Niels Bohr Institute, Blegdamsvej 17, DK-2100 Copenhagen, Denmark*

²*Information Sciences Institute, University of Southern California, Marina del Rey, California 90292, USA*

³*Swedish Defence Research Agency, SE-147 25 Tumba, Sweden*

⁴*Departament de Física i Enginyeria Nuclear, Universitat Politècnica de Catalunya Edifici GAIA, Rambla Sant Nebridi 22, 08222 Terrassa, Barcelona, Spain*

(Received 22 March 2013; revised manuscript received 5 July 2013; published 8 August 2013)

We study adaptive dynamics in games where players abandon the population at a given rate and are replaced by naive players characterized by a prior distribution over the admitted strategies. We demonstrate how such a process leads macroscopically to a variant of the replicator equation, with an additional term accounting for player turnover. We study how Nash equilibria and the dynamics of the system are modified by this additional term for prototypical examples such as the rock-paper-scissors game and different classes of two-action games played between two distinct populations. We conclude by showing how player turnover can account for nontrivial departures from Nash equilibria observed in data from lowest unique bid auctions.

DOI: [10.1103/PhysRevE.88.022806](https://doi.org/10.1103/PhysRevE.88.022806)

PACS number(s): 89.65.Gh, 02.50.Le, 89.75.-k

I. INTRODUCTION

Perhaps the most important skill for competing agents is the ability to adapt to a changing environment by constantly assessing and modifying their behavior. Consequently, while game theory has been initially mostly focused on the study of equilibria [1,2], the study of adaptive dynamics has acquired more and more relevance in recent years. Various models have been put forward to capture the learning processes of competing individuals and the resulting evolution of the population [3–5]. A key result is that many agent-based algorithms of adaptive dynamics allow for a simple macroscopic description in terms of replicator equations [6–9]. In the replicator dynamics, a large population of individuals participate in a game. Let us call x_i the fraction of population playing a given strategy, where admissible strategies are labeled as $i = \{1, 2, 3, \dots\}$. The time evolution of such fractions is given by

$$\frac{d}{dt}x_i = x_i(\pi_i(\mathbf{x}) - \bar{\pi}(\mathbf{x})), \quad (1)$$

where $\pi_i(\mathbf{x})$ is the frequency-dependent payoff of strategy i , and $\bar{\pi} = \sum_i x_i \pi_i$ is the average payoff. In this setting, the initial condition \mathbf{x}_i^0 represents a prior distribution of strategy preferences before the adaptation process takes place. Approaches based on Eq. (1) have proven successful in describing systems within biology as well as economics [8,10,11]. It can be easily shown that stable equilibria of replicator dynamics correspond to Nash equilibria, where no individual can benefit from changing strategy unilaterally [7,12].

Usually, in adaptive dynamics, one has in mind a fixed population of players that acquire experience over time or biological populations where offsprings inherit strategies from their parents. However, one can think of a number of concrete examples where games are played in a more open setting, with the possibility of players to leave the game and be replaced by less experienced ones. On the market, new companies are founded while old companies collapse. Within companies, experienced employees retire so that young graduates may start their career. On general grounds, one should expect turnover to

have a profound impact on the dynamics of the game and lead to a rich phenomenology. In a standard adaptive dynamics all individuals within the population have the same degree of experience. Conversely, here each agent sees a nontrivial mixture of players with different experience levels. While the Nash equilibrium strategy will be optimal against very experienced players, it would not necessarily be the most effective to exploit naive newcomers. Therefore, one can expect adaptive dynamics in this case to converge to equilibria being different from Nash equilibria and being crucially affected by the rate of turnover, which in turn determines the steady-state structure of the population in terms of both experience and strategies. In such a framework, interesting information can be obtained by dissecting the experience composition of each strategy, for example, to assess which experience classes are performing better in the game in terms of payoff.

In this paper, we demonstrate how taking into account turnover of players leads macroscopically to a variant of the replicator dynamics. We present a derivation of the replicator equation with turnover [Eq. (6)] in Sec. II. In the remainder of the paper, we apply this equation to analyze the effect of turnover in simple evolutionary games. We begin with the simple paradigmatic cases of the rock-paper-scissors game and the set of two-action games played between two different populations. In the latter case, we show how increasing the turnover rate can lead to abrupt change in the equilibrium state caused by bifurcations in the corresponding dynamical system. We conclude the paper by showing how this approach can provide an interpretation for the observed bid distribution in online lowest unique bid auctions.

II. MODEL

We aim at generalizing the replicator equations (1) to situations in which agents are replaced with inexperienced individuals at a rate p . Let a large population of players engage in a game with strategies labeled $i = \{1, 2, 3, \dots\}$. We divide the players in experience classes: Let $n_i(\tau, t)$ be, at time t , the number of players having been in the game for a time τ

and playing strategy i . The normalization condition for the n 's reads

$$\sum_i \sum_{\tau=0}^{\infty} n_i(\tau, t) = N \quad \forall t, \quad (2)$$

where N is a fixed total population size. We also define the total fraction of players adopting strategy i at time t , $x_i(t) = \sum_{\tau=0}^{\infty} n_i(\tau, t)/N$. As in the standard replicator equation (1), we introduce the average payoff of strategy i , $\pi_i(\mathbf{x}(t))$ and the average payoff across strategies $\bar{\pi} = \sum_i x_i \pi_i$.

For the sake of simplicity, we consider a simple adaptive dynamics in which individuals learn of alternative strategies

and their average payoff at a rate equal to the fraction of the population that plays by this strategy. Further, an individual playing strategy i learning of a strategy j with a higher average payoff, changes to strategy j with a probability proportional to the payoff difference. Combining these assumptions gives an overall rate of change from strategy i to strategy j equal to $n_i x_j (\pi_j - \pi_i)$, if $\pi_j > \pi_i$.

Furthermore, agents leave the population at rate p and are replaced by inexperienced agents. The new agents play each strategy i with a probability proportional to a given distribution x_i^0 .

We wish to study the time development of the $n_i(\tau, t)$. The learning dynamics encoded in the rules above reads

$$n_i(\tau + 1, t + 1) - n_i(\tau, t) = (1 - p) \left(\sum_{j: \pi_j < \pi_i} n_j(\tau, t) x_i(t) (\pi_i - \pi_j) - \sum_{j: \pi_j > \pi_i} n_i(\tau, t) x_j(t) (\pi_j - \pi_i) \right) - p n_i(\tau, t), \quad (3)$$

where the first sum represents players changing to strategy i from strategies with lower payoffs, and the second sum represents players changing from strategy i to strategies with higher payoffs. Performing a continuous time limit, the left hand side of the above equation becomes $\partial_\tau n_i(\tau, t) + \partial_t n_i(\tau, t)$. We can now integrate the continuous time version of Eq. (3) over τ and divide by the population size N to obtain a closed evolution equation for the strategies $x_i(t)$. Let us recall that $n_i(0, t) = \sigma N x_i^0$, where the proportionality constant σ can be determined by imposing that the final equation preserves the normalization condition. We also assume that, due to the effect of turnover, one has $\lim_{\tau \rightarrow \infty} n_i(\tau, t) = 0 \quad \forall i, t$. After combining the two sums, this results in

$$\frac{d}{dt} x_i(t) = (1 - p) \sum_j x_j(t) x_i(t) (\pi_i - \pi_j) + \sigma x_i^0 - p x_i. \quad (4)$$

Imposing the normalization results in $\sigma = p$. This also implies that the density of players having a given experience level at steady state is exponentially distributed, $\sum_i n_i(\tau, \infty) = N p \exp(-p\tau)$. If we furthermore use the normalization condition $\sum_j x_j = 1$ and the definition of the average payoff $\bar{\pi}$, we get

$$\frac{d}{dt} x_i = (1 - p) x_i (\pi_i - \bar{\pi}) + p (x_i^0 - x_i). \quad (5)$$

Upon rescaling time by $1 - p$ and defining a rescaled turnover rate $\chi = p/(1 - p)$, we finally obtain

$$\frac{d}{dt} x_i = x_i (\pi_i - \bar{\pi}) + \chi (x_i^0 - x_i). \quad (6)$$

Equation (6) constitutes the starting point of our analysis. It should be clear that our specific choice of adaptive dynamics is not crucial for deriving Eq. (6) and that the same macroscopic limit could be obtained for other microscopic adaptation rules leading to the replicator equation (1) in the absence of turnover. Equation (6) can be, of course, also derived (or justified) heuristically, for example by making an analogy with a (damped) driven dynamical system, where the prior

distribution x_i^0 in the right hand side acts as a forcing term. Note that Eq. (6) can be formally recast in replicator form,

$$\frac{d}{dt} x_i = x_i (\pi_i - \bar{\pi}) + \chi (x_i^0 - x_i) \equiv x_i (\tilde{\pi}_i - \bar{\pi}), \quad (7)$$

where the ‘‘effective payoff’’ of strategy i , $\tilde{\pi}_i$, is defined as

$$\tilde{\pi}_i = \pi_i + \chi \left(\frac{x_i^0}{x_i} - 1 \right). \quad (8)$$

In Eq. (7) the average payoff $\bar{\pi}$ does not contain a contribution from the second term in Eq. (8) as its average is zero, so that $\sum_i x_i \pi_i = \sum_i x_i \tilde{\pi}_i$. The mapping in Eq. (7) is valid only in the interior of the simplex $x_i > 0 \forall i$, $\sum_i x_i = 1$, while at the boundary the effective payoffs diverge. This situation has some resemblance to the case of evolutionary dynamics in the presence of mutations [7, 13]. The main difference is that in our case the ‘‘mutants’’ are not characterized by new, pure strategies or random strategies, but the mixed distribution x_i^0 of the existing strategies.

The fact that the replicator equation with turnover (6) can be rewritten in replicator form (7) implies that the two equations share several mathematical properties. Among these is that the mean effective payoff for the population cannot decrease along any trajectory. If the payoff function is continuous and bounded, it therefore serves as a Lyapunov function [14], which guarantees that all game dynamics evolve to either an equilibrium point or a closed orbit. For the replicator equation without turnover, $\chi = 0$, all Nash equilibria of a game are equilibrium points of the dynamics [15]. For positive χ , the equilibrium points are generally different from the Nash equilibria, and we call these the *turnover equilibria* of the game. In the limit $\chi \rightarrow \infty$, agents are insensitive to the rewards and the initial strategy x_i^0 is the only equilibrium point.

III. ROCK-PAPER-SCISSORS

We now apply the concept of agent turnover to the three-strategy game of rock-paper-scissors. In this well-known,

biologically relevant game [16–20], a large population of players can choose among the strategies rock, paper, and scissors. The average fraction of the population employing each strategy is denoted x_R , x_P , and x_S , respectively. The players are paired randomly in each round. Since the strategies dominate each other cyclically, their payoffs are given by

$$\pi_R = x_S - x_P, \tag{9}$$

$$\pi_P = x_R - x_S, \tag{10}$$

$$\pi_S = x_P - x_R. \tag{11}$$

Rock-paper-scissors is a zero-sum game, so the average payoff in (6) is always zero. The only Nash equilibrium of the game is when all individuals play randomly, giving them all an average payoff of zero [21].

When we introduce a player turnover, naive players that enter the game choose to play rock, paper, or scissors with probability x_R^0 , x_P^0 , and x_S^0 , respectively. Gradually, each individual changes strategy according to Eq. (3), resulting in the mean behavior given by Eq. (6).

In Figs. 1(a)–1(d) the time development of the average strategy played is shown for an initial strategy $\mathbf{x}^0 = (0.8, 0.1, 0.1)$ biased towards rock and for different values of χ . Similar results are obtained using different initial strategies. In the absence of agent turnover, the strategy of the population oscillates around the Nash equilibrium [21,22]. Upon introducing turnover, the initial overrepresentation of players playing rock makes paper a rewarding strategy. As players change strategy towards paper, scissors becomes more rewarding, and so on. In the limit of $\chi \rightarrow \infty$, all players remain naive, so the strategy of the population is equal to the initial strategy.

For intermediate turnover rates, the system undergoes damped or critically damped oscillations towards the turnover equilibrium, where the gradual learning of the players staying in the game exactly balances the exchange of experienced players with naive players. Thus, agent turnover stabilizes the game dynamics in a similar way as spatial organization of agents [17,22–24]. The location of the turnover equilibrium depends nontrivially on the strategy of naive players and on the turnover rate, as shown in Fig. 1(d). When χ is increased from 0 to ∞ , the turnover equilibrium shifts from the Nash equilibrium to the strategy distribution of inexperienced players along a curved line. This means that, for a generic value of χ , the turnover equilibrium is not trivially a linear combination of the Nash equilibrium and the naive strategy.

In a traditional game of rock-paper-scissors, the time-averaged payoff of all players is zero in any steady state situation. In the turnover equilibrium, however, the average payoff of players increases monotonically with their experience, such that new, naive players tend to lose to more experienced players. Figure 1(e) shows the average strategy $n_i(\tau, \infty)/N$ as a function of the experience level τ , where the turnover is $\chi = 0.25$ and the naive strategy is the same as before. The strategy with the highest payoff is paper, which is the only strategy that very experienced players employ. However, due to the turnover of agents, the number of experienced players falls off exponentially. This causes a radical qualitative difference between the payoff of a given player having an

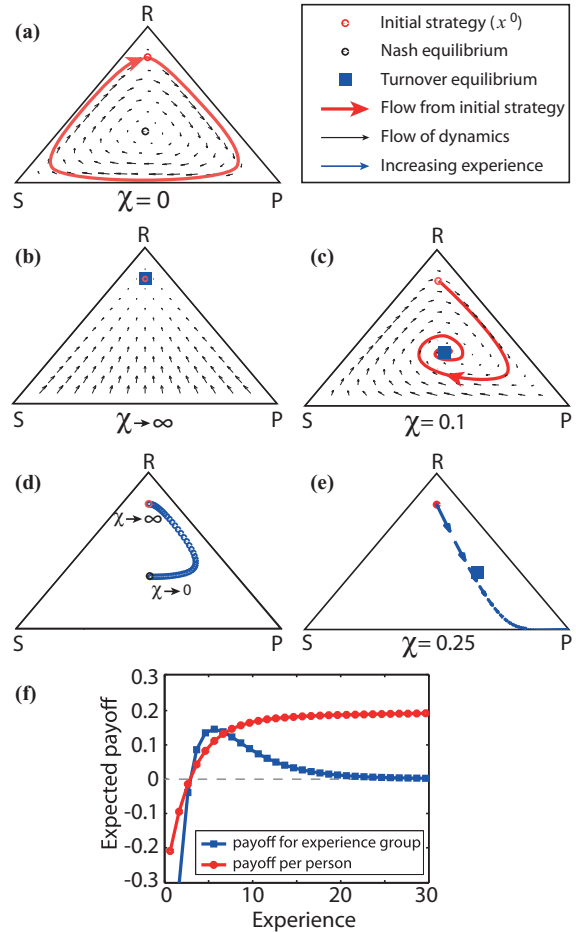


FIG. 1. (Color online) Turnover of players in the rock-paper-scissors game with an initial strategy of 80% rock and 10% paper and scissors. (a) For $\chi = 0$, the average strategy oscillates in a fixed distance from the Nash equilibrium. (b) For infinite turnover, only the initial strategy \mathbf{x}^0 is employed. (c) For intermediate turnovers, the average strategy converges to the turnover equilibrium. (d) The turnover equilibrium converges for different values of χ . The equilibrium is displaced towards paper, which dominates the initial strategy. (e) In the turnover equilibrium, inexperienced players quickly turn away from playing rock. Experienced players only play paper. (f) Inexperienced players have a large negative payoff, which monotonically increases as the players gain experience. However, since the number of experienced players falls off exponentially, most games are won by players of intermediate experience. In this example, we chose a total population size $N = 20$.

experience level τ and the total payoff of players in a given experience level. Denoting the turnover equilibrium by \mathbf{x}^* , the total payoff collected by players with experience τ is given by $\sum_i n_i(\tau, \infty) \pi_i(\mathbf{x}^*)$, while the payoff of a single player with this experience is $\sum_i n_i(\tau, \infty) \pi_i(\mathbf{x}^*) / (\sum_i n_i(\tau, \infty))$. The two different payoffs are shown in Fig. 1(f) as a function of the experience level τ for a population of 20 players. The figure demonstrates the counterintuitive fact that, while the payoff of each single player increase monotonically with the experience, the majority of the games are won by players with intermediate experience levels.

IV. TWO-AGENT TWO-ACTION GAMES

We now turn to the broad class of games, where two populations can each choose between two strategies. Let us define the mixed strategies as $(x, 1 - x)$ for the first population and $(y, 1 - y)$ for the second population. The payoffs obtained by the two populations depend on their joint actions [15] and are traditionally encoded in payoff matrices A and B ,

$$\begin{pmatrix} \pi_1^x \\ \pi_2^x \end{pmatrix} = \begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix} \begin{pmatrix} y \\ 1 - y \end{pmatrix}, \quad (12)$$

$$\begin{pmatrix} \pi_1^y \\ \pi_2^y \end{pmatrix} = \begin{pmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{pmatrix} \begin{pmatrix} x \\ 1 - x \end{pmatrix}, \quad (13)$$

where π_1^x is the payoff of the first strategy for the first population, and so forth. The two populations are learning concurrently and their turnover rates χ_x and χ_y can, in principle, be different. The learning dynamics can be obtained from Eq. (6),

$$\frac{d}{dt}x = x(1 - x)(\pi_1^x - \pi_2^x) + \chi_x(x^0 - x), \quad (14)$$

$$\frac{d}{dt}y = y(1 - y)(\pi_1^y - \pi_2^y) + \chi_y(y^0 - y), \quad (15)$$

where the equations for the second strategies can be simply obtained from the normalization conditions. In the absence of turnover, this system is known as a bimatrix replicator equation [15,25]. In the turnover equilibrium, which we denote as $(x, y) = (x^*, y^*)$, the time derivatives are equal to zero. Substituting expressions (12) and (13) for the payoffs leads to the equations

$$\alpha y^* + a_{12} - a_{22} = \chi_x \frac{x^* - x^0}{x^*(1 - x^*)}, \quad (16)$$

$$\beta x^* + b_{12} - b_{22} = \chi_y \frac{y^* - y^0}{y^*(1 - y^*)}, \quad (17)$$

where we have introduced

$$\alpha = a_{11} + a_{22} - a_{21} - a_{12}, \quad (18)$$

$$\beta = b_{11} + b_{22} - b_{21} - b_{12}. \quad (19)$$

The number of solutions to Eqs. (16) and (17) depends on the sign of the product $\alpha\beta$. When $\alpha\beta < 0$, there is always one mixed turnover equilibrium. When $\alpha\beta > 0$, the number of turnover equilibria can increase at critical values of the turnover parameter. The derivation of these results can be found in Appendix A. In the following, we provide examples for two well-known two-action games belonging to the two categories, $\alpha\beta < 0$ and $\alpha\beta > 0$: the game of matching pennies and a coordination game, respectively.

A. Matching pennies

A simple example of $\alpha\beta < 0$ is the zero sum game of matching pennies. Two players secretly turn a penny each to heads or tails and reveal the coins simultaneously. If the pennies match, player one receives a reward r from player two. If they do not, player one pays the reward to player two.

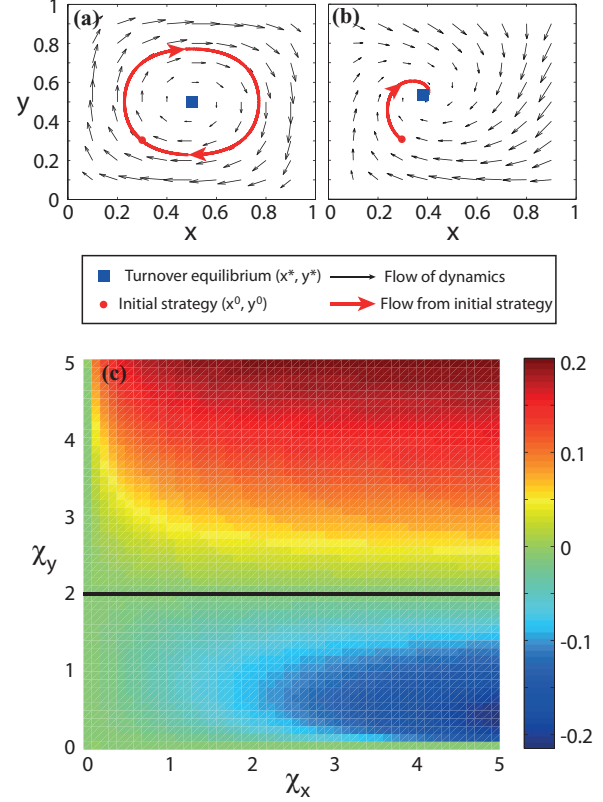


FIG. 2. (Color online) Dynamics of the matching pennies game with $r = 2$ and initial strategy $x^0 = y^0 = 0.3$. Other choices of parameter values yield similar results. (a) For $\chi = 0$, the average strategy oscillates at a fixed distance from the Nash equilibrium. (b) When a turnover is introduced, the average strategy converges to the turnover equilibrium. Here $\chi_x = \chi_y = 1$. (c) Expected payoff for population one in the turnover equilibrium for varying turnovers. Above a critical value of $\chi_y = 2$, the first population will win most games.

This game is characterized by the payoff matrices

$$A = \begin{pmatrix} r & -r \\ -r & r \end{pmatrix}, \quad B = \begin{pmatrix} -r & r \\ r & -r \end{pmatrix}. \quad (20)$$

The learning dynamics (14) and (15) with no turnover leads to a marginally stable Nash equilibrium surrounded by concentric closed orbits [see Fig. 2(a)]. When a positive turnover of players is introduced, the Nash equilibrium is perturbed to a stable turnover equilibrium (x^*, y^*) , as displayed in Fig. 2(b). Inserting the payoff matrices into the conditions for turnover equilibrium (16) and (17) gives

$$y^* = \frac{1}{2} + \frac{\chi_x}{4r} \frac{x^* - x^0}{x^*(1 - x^*)}, \quad (21)$$

$$x^* = \frac{1}{2} + \frac{\chi_y}{4r} \frac{y^* - y^0}{y^*(1 - y^*)}. \quad (22)$$

Matching pennies is a zero sum game. The expected payoff of population one is positive if the two populations have a tendency of both playing heads or both playing tails, while it is zero if and only if one of the populations chooses their strategy randomly. From Eqs. (21) and (22) we see that this happens in the turnover equilibrium if $(x^*, y^*) = (1/2, y^0)$

or $(x^*, y^*) = (x^0, 1/2)$. For a given set of initial strategies of the populations, this corresponds to the critical turnover rates

$$\chi_{x,c} = -r \frac{1 - 2y^0}{1 - 2x^0}, \tag{23}$$

$$\chi_{y,c} = r \frac{1 - 2x^0}{1 - 2y^0}. \tag{24}$$

Equation (23) can only be satisfied if population two wins most games when the initial strategies are employed. In this case, the expected payoff of population one is negative if it has a high turnover, since its equilibrium strategy is close to its initial strategy. If the turnover is decreased below the critical value (23), the first population starts winning more games than it loses regardless of the turnover of population two. Even if population two has an even lower turnover, and thus has an initially superior strategy and more time to gain experience, the equilibrium state remains advantageous for population one.

Likewise, Eq. (24) gives a positive critical turnover only if population one wins most games when the initial strategies are employed. Here, the payoff for population two goes from being positive to negative when its turnover is increased past the critical value given by (24), regardless of the turnover of the first population. In Fig. 2(c), the average payoff for population one is shown for varying turnovers of both populations and the initial strategies $x^0 = y^0 = 0.3$. In this case, if the turnover of population two is larger than $\chi_{y,c} = 2$, population one wins most games.

B. Coordination game

When the payoff matrices of the two populations equal each other, $A = B$, one must necessarily have $\alpha\beta \geq 0$. One example of this is a coordination game, where the two players strive to play the same strategy. We consider the payoff matrices

$$A = B = \begin{pmatrix} 6 & 0 \\ 3 & 2 \end{pmatrix}, \tag{25}$$

where $\alpha = \beta = 5$. Coordination games with different payoff matrices will have similar dynamical properties. Without player turnover, this game has two pure Nash equilibria where the populations employ the same strategy. In addition, there is a mixed Nash equilibrium at $(x^*, y^*) = (2/5, 2/5)$ with one stable and one unstable manifold. The stable manifold constitutes the boundary between the basins of attraction of the two pure Nash equilibria [see Fig. 3(a)]. For sufficiently small turnovers, there are three turnover equilibria close to these three Nash equilibria. The initial strategies (x^0, y^0) determine which equilibrium state the system goes to.

Let us consider a situation of strong initial disagreement $(x^0, y^0) = (0.9, 0.1)$, where inexperienced players of populations one and two tend to employ the first and second strategies, respectively. We fix $\chi_y = 2$, while varying χ_x as a control parameter. In this case, the point $(x, y) = (0.4, 0.9)$ is always an equilibrium of the dynamics. At a critical value, $\chi_x \approx 0.71$, the boundary between the basins of attractions passes the initial strategy (x^0, y^0) [see Figs. 3(b) and 3(c)], and the steady state changes discontinuously from being dominated by strategy two to strategy one. At the same point, the payoff of both populations increases drastically.

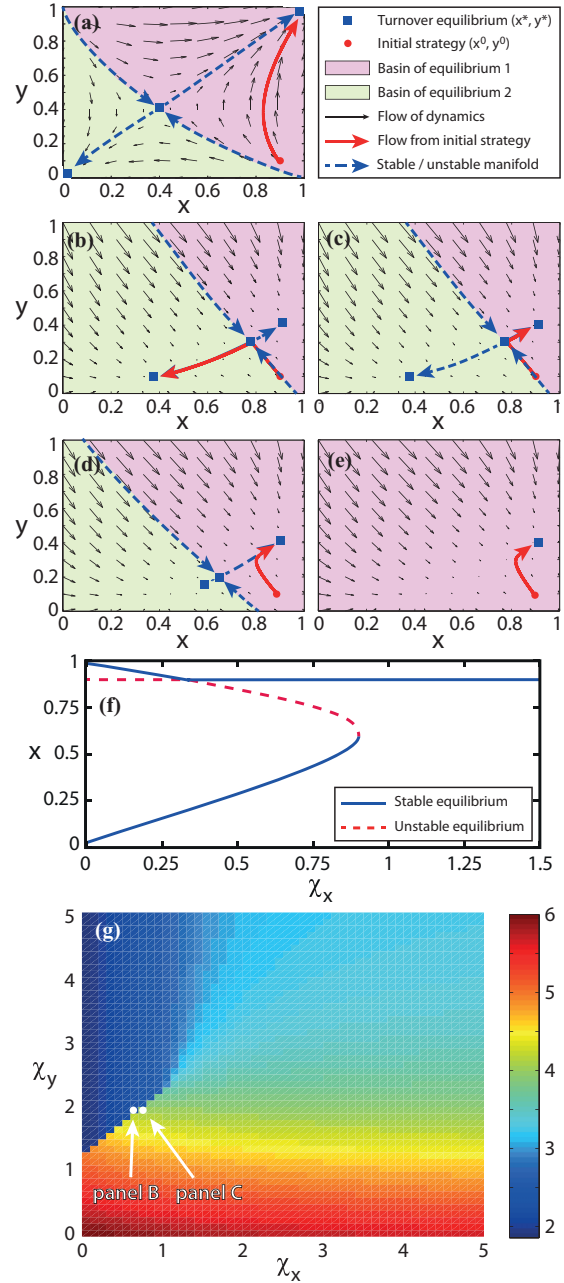


FIG. 3. (Color online) Dynamics of a coordination game with initial strategy $(x^0, y^0) = (0.9, 0.1)$. (a) Without turnover, the game has one saddle point Nash equilibrium with stable manifolds that separate the basins of attraction of two pure Nash equilibria. (b), (c) When the initial strategy goes from one basin of attraction to another, the resulting equilibrium state changes discontinuously. Here $\chi_y = 2$ and $\chi_x = 0.7$ and 0.72 , respectively. (d), (e) At a critical set of turnover rates two turnover equilibria annihilate in a saddle node bifurcation. Here $\chi_y = 2$ and $\chi_x = 0.91$ and 0.92 , respectively. (f) Bifurcation diagram for turnover equilibria. The point $(x, y) = (0.4, 0.9)$ is always an equilibrium. We observe a transcritical bifurcation at $\chi_x \approx 0.4$ and the saddle node bifurcation from panels (d) and (e) at $\chi_x \approx 0.9$. (G) Expected payoff of population one in the turnover equilibrium for varying turnovers. The dramatic change in payoffs between panels B and C can clearly be seen.

Upon increasing the turnover of the first population even further, the basin of attraction of strategy two suddenly disappears at $\chi_x \approx 0.91$. For this value, the saddle node equilibrium annihilates with the turnover equilibrium of strategy two [see Figs. 3(d) and 3(e)]. For larger values of χ_x , there is only one turnover equilibrium. In this case the number of turnover equilibria changes through a saddle node bifurcation, but for other payoff matrices [(A),(B)], a pitchfork bifurcation could occur (see Appendix A).

A full bifurcation diagram is shown in Fig. 3(f), and the expected payoff of population one is shown in Fig. 3(g) as a function of the turnover rates. The abrupt change in payoff at a critical set of turnover rates can clearly be seen. The expected payoff of population two follows a similar pattern. In general, a low turnover of either population results in higher payoffs of both populations, since a larger proportion of the players has enough experience to mainly bid on the dominating strategy. However, increasing the turnover may change which equilibrium the system goes to, leading to an increase in the payoff of both players.

V. LOWEST-UNIQUE-BID AUCTIONS

Lowest unique bid auctions have become popular online games and have recently attracted attention from the scientific community [26–28]. In this game, N players pay an entrance fee to independently bid on an item, e.g., a car. The winner of the auction is the player with the lowest bid not also played by another player. The strategies available to each player are the possible integer bids $i = \{1, 2, 3, \dots\}$.

If a population of players place their bids stochastically according to the distribution x_i , the average payoff of each strategy is, in the large N limit, proportional to [28]

$$\pi_i = e^{-Nx_i} \prod_{j=1}^{i-1} (1 - Nx_j e^{-Nx_j}). \quad (26)$$

The Nash equilibrium corresponding to the payoffs of Eq. (26) was compared with data from the auctions web site *auctionair.co.uk* in [28]. An excellent agreement between the Nash equilibrium and the data was found for low to intermediate values of N . In larger auctions, the bidding distribution departed from the Nash equilibrium and was more resemblant of an exponential distribution,

$$x_i^0 = \frac{e^{-\beta i}}{\sum_i e^{-\beta i}}, \quad (27)$$

where the fitted value of the exponential constant was $\beta \approx 0.02$. On general theoretical grounds, one can argue that the exponential bidding strategy is the expected prior distribution for inexperienced players, with limited information about the game, that want to avoid the cost of a large bid [28,29].

Using (26) and (27), we test here the hypothesis that the apparent change of behavior from small to large auction sizes is caused by agent turnover. Note that changing the turnover rate χ in (6) is equivalent to multiplying the payoff function (26) with a constant and rescaling time. As it is also problematic to infer a natural time scale of adaptation from the data, we leave the value of χ as a fitting parameter.

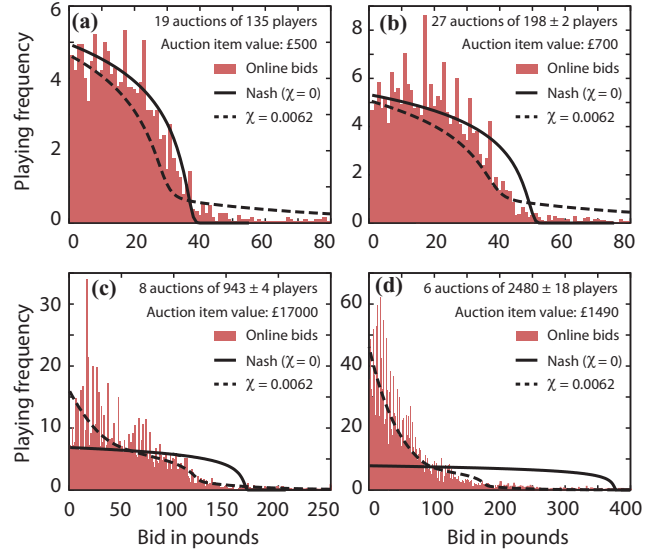


FIG. 4. (Color online) Bidding frequencies in online lowest unique bid auctions compared to the theoretical Nash equilibrium and the turnover equilibrium for a constant $\chi = 0.0062$. Different panels show auctions with items of different value and different numbers of players. For small auction sizes both the Nash and the turnover equilibria fit the empirical bidding distributions. In larger auctions the turnover equilibrium fits the data much better than the Nash equilibrium. This is due to a long adaptation time for players in large auctions.

Figure 4 compares empirical data from 60 online auctions with varying number of players to the respective Nash equilibria and turnover equilibria for a least squares fitted value of $\chi = 0.0062$. We define the squared distance d between the empirical data and the turnover and Nash equilibria as

$$d_{\text{turn}} = N^2 \sum_i (x_i^{\text{emp}} - x_i^*)^2, \quad (28)$$

$$d_{\text{Nash}} = N^2 \sum_i (x_i^{\text{emp}} - x_i^{\text{Nash}})^2, \quad (29)$$

where \mathbf{x}^{emp} is obtained from the frequencies, \mathbf{x}^{Nash} is the theoretical Nash equilibrium, and the sum runs over all bids and all auctions. Using these definitions we get $d_{\text{turn}} = 89\,084$ and $d_{\text{Nash}} = 253\,526$. In [28] it is shown that, if the bids were randomly drawn from the theoretical distribution, the squared distance is expected to be equal to the combined number of bidders $d_{\text{exp}} = \sum N = 30\,335$. This shows that the turnover equilibria describe the bidding distributions much better than the Nash equilibria, introducing just a single free parameter.

In small auctions, both the Nash equilibrium and the turnover equilibrium fit the data well, but the latter also captures the fat tail of the empirical distribution. In larger auctions, the Nash equilibrium fails to predict the observed distribution of bids, while the turnover equilibrium fits the data remarkably well.

It should be recalled that, while the exponential strategy distribution for large N is plugged into the model empirically as a prior, it is still remarkable that, for the same value of χ , the resulting equilibrium is relatively close to the prior for large N and closer to the Nash equilibrium for smaller values of N . The reason stems from the growth of the adaptation time

of the replicator dynamics with the auction size N (see [28]). The consequence is, even at equal turnover rate χ , players have sufficient time to adapt for small auction size but not for larger auctions. Clearly, one could get an even better fit by letting χ vary as a function of N , as one could have different kinds of players and therefore different turnovers for different auction types, as exemplified in Appendix B. We do not present a systematic study of the dependence of χ on N , as it is beyond the scope of this paper. Our main point here is to show that, even in the simple case of a fixed χ , the replicator equation with turnover is useful to model a real-life game and can help in conceptualizing nontrivial discrepancies between observed strategy distributions and Nash equilibria. Finally, in Appendix B we show that, in the case of larger auctions, the distance from the Nash equilibrium does not seem to depend on time. This fact further supports the presence of a steady effect, like the turnover mechanism proposed here, preventing the population to reach the Nash equilibrium.

VI. DISCUSSION AND CONCLUSION

In this paper we have extended the theoretical framework of the replicator equation by considering the effect of agent turnover, a phenomenon that is, in principle, present in many game theoretical contexts. Our model provides a simple and compact description of the dynamics of the average strategies in the population and at the same time allows for dissecting the behavior of players having different experience levels. The study of experience-dependent payoffs can bring to counterintuitive results, as we demonstrated in the simple case of a rock-paper-scissors game. When turnover is introduced, the game dynamics become richer. We have shown that it is possible to encounter bifurcations of turnover equilibria and that the equilibrium strategy can change discontinuously with the turnover rate.

The simplicity and flexibility of the approach proposed in this paper can be of use to analyze several games of practical importance. For example, we have shown here how turnover of agents can help in explaining adaptation in online games, where experienced players may exploit the naive strategies of new players. Our comparison with a dataset of bidding distributions in online lowest unique bid auctions shows that our model is able to describe the empirical steady state distributions far better than the Nash equilibrium. We expect that similar results can be obtained by studying data from other games.

The concepts presented here suggest an interesting analogy with the idea of *cognitive hierarchy* [30–32] in human learning, where players are classified according to the number of reasoning steps they are able to make. In this scheme, nonreasoning players are reminiscent of inexperienced players in the present model, and players with more reasoning steps can be compared more to experienced players. In this perspective, the model proposed here provides a framework in which a cognitive hierarchy emerges naturally from the dynamical equilibrium between adaptation and turnover. Indeed, a cognitive hierarchy model has previously been applied to describe the bidding distributions in lowest unique bid auctions [27]. In the authors of Ref. [27]’s own words, “the cognitive hierarchy model ... should be viewed as a potential stepping stone to an

investigation using a formal learning model.” Our work is one proposal of such a learning model.

It is known that a fraction of nonrational, influenceable players may change the game dynamics considerably [33–35]. In a recent study [36], a population of fish were made up of three subgroups: a large group of fish with a small preference for going to one place in the aquarium, a smaller group with a strong preference for another place, and a group of untrained fish with no prior preference. The study showed, both through simulations and experiment, that there exists a critical size of the untrained group above which the entire population goes to the place preferred by group one, and below which the minority of group two dictates where the population goes. This resembles our results for two-agent two-action coordination games, where a critical turnover of agents changes the turnover equilibrium from being dominated by one strategy to being dominated by the other.

In our model we assume a constant turnover of agents in time, which leaves one free fitting parameter. However, changing the turnover is mathematically equivalent to scaling the expected payoff of all strategies together with a time scaling. Since the absolute value of the payoff is often treated as a free parameter when analyzing experiments of learning dynamics, this does not increase, in general, the number of fitting parameters in such an analysis. Furthermore, if the absolute values of the payoffs are known, one can, in principle, measure the turnover rate of agents and avoid having a free fitting parameter.

In conclusion, we have shown how a steady turnover of agents, which is present in most real-life games, can change the qualitative dynamics of a game. The simplicity and generality of the framework presented here makes it a natural candidate to describe adaptation in a population participating in a game and subject to turnover.

ACKNOWLEDGMENTS

We thank A. Galstyan and C. Strandkvist for comments on a preliminary version of the manuscript. A.K.’s research was supported in part by the National Science Foundation under Grant No. 0916534 and US AFOSR MURI Grant No. FA9550-10-1-0569. S.P. acknowledges partial support from the Spanish Research Ministry through Grant No. FIS2012-37655-C02-01.

APPENDIX A: NUMBER OF TURNOVER EQUILIBRIA IN TWO-AGENT TWO-ACTION GAMES

In this appendix, we show that the number of possible turnover equilibria in two-agent two-action games is dependent on the sign of the product $\alpha\beta$, where α and β are given by (18) and (19).

The conditions for turnover equilibrium are given by (16) and (17). Expressing y^* as a function of x^* , this can be written as

$$y^*(x^*) = \frac{\chi_x}{\alpha} \frac{x^* - x^0}{x^*(1 - x^*)} - \frac{a_{12} - a_{22}}{\alpha}, \quad (A1)$$

$$\beta x^* + b_{12} - b_{22} = \chi_y \frac{y^*(x^*) - y^0}{y^*(x^*)(1 - y^*(x^*))} \equiv g(x^*), \quad (A2)$$

where we defined the composite function $g(x^*)$. Notice that $\lim_{x^* \rightarrow 0} g(x^*) = \lim_{x^* \rightarrow 1} g(x^*) = 0$ and that $g(x^*)$ has two singularities for values of x^* such that $y^* = 0$ and $y^* = 1$. Since $0 \leq y^* \leq 1$, we focus on the interval of $g(x^*)$ between the two singularities.

Let us investigate the derivative $\partial g / \partial x^*$. From Eqs. (A1) and (A2) we get

$$\frac{\partial y^*}{\partial x^*} = \frac{\chi_x (x^* - x^0)^2 + x^0 - x^{02}}{\alpha (1 - x^*)^2 x^{*2}}, \quad (\text{A3})$$

$$\frac{\partial g}{\partial x^*} = \chi_y \frac{(y^* - y^0)^2 + y^0 - y^{02}}{(1 - y^*)^2 y^{*2}} \frac{\partial y^*}{\partial x^*}. \quad (\text{A4})$$

Since x^0 and y^0 are probabilities, we always have $x^0 > x^{02}$ and $y^0 > y^{02}$. The fractions in Eqs. (A3) and (A4) are therefore




FIG. 5. (Color online) The number of turnover equilibria in a two-action two-agent game as determined by a graphical solution of Eq. (A2). (a) When $\alpha\beta < 0$ the two functions have opposite slopes and hence intersect exactly once in the relevant interval $0 < y^* < 1$. Here the matching pennies game is illustrated with the parameters of Fig. 2(b). (b), (c) When $\alpha\beta > 0$ the slope of the functions have the same sign, and new turnover equilibria can appear through bifurcations. Here the bifurcation in the coordination game of Figs. 5(b) and 5(c) is shown, except that χ_x is changed from to 0.8 to 1 between panels (b) and (c).

positive, so the derivative $\partial g / \partial x^*$ must have the same sign as α . This means that $g(x^*)$ is either increasing or decreasing monotonically, depending on the sign of α .

If $\alpha\beta < 0$, the line $\beta x^* + b_{12} - b_{22}$ and the function $g(x^*)$ have opposite slopes and therefore intersect exactly once. It follows from (A2) that there is always exactly one turnover equilibrium. In Fig. 5(a), this is illustrated for the matching pennies game with the same parameters as Fig. 2(b).

If $\alpha\beta > 0$, the line $\beta x^* + b_{12} - b_{22}$ and the function $g(x^*)$ are either both increasing or both decreasing. Therefore, they can, in principle, intersect any odd number of times. Hence, it is possible to have multiple turnover equilibria, and these can appear or annihilate in pairs through either saddle point bifurcations or pitchfork bifurcations. Figures 5(b) and 5(c) shows a saddle node bifurcation in the coordination game with the same parameters as Figs. 3(d) and 3(e), except we have set χ_x equal to 0.8 and 1, respectively, to show a clearer bifurcation.

APPENDIX B: ADDITIONAL ANALYSIS OF THE LOWEST UNIQUE BID AUCTION DATA

In this appendix, we present additional results on the lowest unique bid auction dataset used in Sec. V. Figure 6 shows how letting χ as an independent fitting parameter for each auction size leads to very good fits of the bidding distributions in all cases. The fitted values suggest a higher turnover rate for larger auction sizes, which is reasonable considering, for example, that larger auctions typically involve larger bidding fees. Furthermore, we see from the combined distances (28) and (29) presented in Fig. 6 that we always have $d_{\text{turn}} < d_{\text{Nash}}$, so the turnover equilibria always fit the empirical data better than the Nash equilibria, but that this is most significant




FIG. 6. (Color online) Same data as Fig. 4, but this time the turnover parameter χ is fitted independently in each plot. The fitted value of χ and the total distances between the empirical bidding distributions and the turnover and Nash equilibria, respectively, are indicated in each legend.




FIG. 7. (Color online) Distance from the Nash equilibrium for 14 auctions having $N = 1398 \pm 74$, as a function of the time order of the auctions. The straight line is a fitted linear regression, with slope $m = 1 \pm 140$.

for large auction sizes. Both equilibria describe the bidding distributions well in small auctions. Notice, however, that the prior distribution x_0 has been obtained from the larger auction sizes, so that estimates of χ in the large auction cases can be affected by overfitting.

Finally, in Fig. 7 we show the evolution of the distance (29) to the Nash equilibrium in series of auctions having similar number of players $N = 1398 \pm 74$. The exact date of each auction was not available, so data are plotted simply as a function of the time ordering of the auctions.

It is easy to demonstrate that, if players bid according to the Nash equilibrium, the expected value of d is N (see [28]). A value of d significantly larger than N signals a departure from the Nash equilibrium. Fitting the data with a linear regression yields a slope $m = 1 \pm 140$ that clearly does not allow for refuting the null hypothesis of no temporal trend. This suggests the presence of a mechanism keeping the population at a finite distance from the Nash equilibrium, as the turnover mechanism proposed in this paper.

- [1] J. Von Neumann and O. Morgenstern, *Theory of Games and Economic Behavior* (Princeton University Press, New Jersey, 1944).
- [2] J. Nash, *Proc. Natl. Acad. Sci. USA* **36**, 48 (1950).
- [3] R. Bush and F. Mosteller, *Psychol. Rev.* **58**, 313 (1951).
- [4] C. Watkins and P. Dayan, *Mach. Learn.* **8**, 279 (1992).
- [5] A. Kianercy and A. Galstyan, *Phys. Rev. E* **85**, 041145 (2012).
- [6] R. Fisher, *The Theory of Natural Selection* (Oxford University Press, London, 1930).
- [7] J. Hofbauer, *J. Math. Biol.* **23**, 41 (1985).
- [8] M. Nowak, *Evolutionary Dynamics: Exploring the Equations of Life* (Belknap Press, Cambridge, MA, 2006).
- [9] A. Traulsen, J. C. Claussen, and C. Hauert, *Phys. Rev. Lett.* **95**, 238701 (2005).
- [10] J. Maynard Smith, *Evolution and the Theory of Games* (Cambridge University Press, Cambridge, UK, 1982).
- [11] G. Silverberg, *Evolutionary Modeling in Economics: Recent History and Immediate Prospects*, Research Memorandum (MERIT, Maastricht, 1997).
- [12] J. Nash, *Ann. Math.* **54**, 286 (1951).
- [13] C. E. Tarnita, T. Antal, and M. A. Nowak, *J. Theor. Biol.* **261**, 50 (2009).
- [14] S. Strogatz, *Nonlinear Dynamics and Chaos: With Applications to Physics, Biology, Chemistry and Engineering* (Perseus Books Group, New York, 2001).
- [15] J. Hofbauer and K. Sigmund, *Bull. Am. Math. Soc.* **40**, 479 (2003).
- [16] M. Frean and E. R. Abraham, *Proc. R. Soc. London B* **268**, 1323 (2001).
- [17] B. Kerr, M. A. Riley, M. W. Feldman, and B. J. M. Bohannan, *Nature (London)* **418**, 171 (2002).
- [18] B. Kirkup and M. Riley, *Nature (London)* **428**, 412 (2004).
- [19] J. Nahum, B. Harding, and B. Kerr, *Proc. Natl. Acad. Sci. USA* **108**, 10831 (2011).
- [20] B. Sinervo and C. Lively, *Nature (London)* **380**, 240 (1996).
- [21] G. Szabó and G. Fáth, *Phys. Rep.* **446**, 97 (2007).
- [22] T. Reichenbach, M. Mobilia, and E. Frey, *Phys. Rev. E* **74**, 051907 (2006).
- [23] J. Juul, K. Sneppen, and J. Mathiesen, *Phys. Rev. E* **85**, 061924 (2012).
- [24] J. Juul, K. Sneppen, and J. Mathiesen, *Phys. Rev. E* **87**, 042702 (2013).
- [25] J. Hofbauer, *J. Math. Biol.* **34**, 675 (1996).
- [26] S. K. Baek and S. Bernhardtsson, *Fluctuation Noise Lett.* **1**, 61 (2010).
- [27] R. Ostling, J. Wang, E. Chou, and C. Camerer, *Am. Econ. J.: Microeconomics* **3**, 1 (2011).
- [28] S. Pigolotti, S. Bernhardtsson, J. Juul, G. Galster, and P. Vivo, *Phys. Rev. Lett.* **108**, 088701 (2012).
- [29] J. K. Goeree and C. A. Holt, *Proc. Natl. Acad. Sci. USA* **96**, 10564 (1999).
- [30] M. Costa-Gomes, V. Crawford, and B. Broseta, *Econometrica* **68**, 1193 (2001).
- [31] C. F. Camerer, T.-H. Ho, and J.-K. Chong, *Q. J. Econ.* **119**, 861 (2004).
- [32] J. R. Wright and K. Leyton-Brown, *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence* (AAI Press, Menlo Park, 2010), pp. 901–907.
- [33] J. Halpern and R. Stern, *Debating Rationality: Non Rational Aspects of Organizational Decision Making*, Frank W. Pierce Memorial Lectureship and Conference Series (Cornell University Press, Ithaca, New York, 1998).
- [34] M. Nowak, K. Page, and K. Sigmund, *Science* **289**, 1773 (2000).
- [35] P. L. Lorentziadis, *Eur. J. Oper. Res.* **217**, 653 (2012).
- [36] I. Couzin, C. Ioannou, G. Demirel, T. Gross, C. Torney, A. Hartnett, L. Conradt, S. Levin, and N. Leonard, *Science* **334**, 1578 (2011).