

Stochastic gain in finite populations

Torsten Röhl,¹ Arne Traulsen,^{2,3} Jens Christian Claussen,¹ and Heinz Georg Schuster¹

¹*Institute of Theoretical Physics and Astrophysics, University of Kiel, Leibnizstraße 15, D-24098 Kiel, Germany*

²*Program for Evolutionary Dynamics, Harvard University, Cambridge, Massachusetts 02138, USA*

³*Max Planck Institute for Evolutionary Biology, August-Thienemann-Strasse 2, 24306 Plön, Germany*

(Received 2 July 2007; revised manuscript received 8 April 2008; published 15 August 2008)

Flexible learning rates can lead to increased payoffs under the influence of noise. In a previous paper [Traulsen *et al.*, Phys. Rev. Lett. **93**, 028701 (2004)], we have demonstrated this effect based on a replicator dynamics model which is subject to external noise. Here, we utilize recent advances on finite population dynamics and their connection to the replicator equation to extend our findings and demonstrate the stochastic gain effect in finite population systems. Finite population dynamics is inherently stochastic, depending on the population size and the intensity of selection, which measures the balance between the deterministic and the stochastic parts of the dynamics. This internal noise can be exploited by a population using an appropriate microscopic update process, even if learning rates are constant.

DOI: [10.1103/PhysRevE.78.026108](https://doi.org/10.1103/PhysRevE.78.026108)

PACS number(s): 89.65.-s, 87.23.-n, 05.45.-a, 02.50.Ey

I. INTRODUCTION

Evolutionary biology considers the dynamics in populations in which the potential to produce offspring differs among the individuals. Individuals with higher fitness produce more offspring and spread in the population. In the simplest case, fitness is a fixed number and the fittest type spreads fastest. Evolutionary game theory describes frequency-dependent selection, i.e., the fitness depends on the types and frequencies of others in the population [1–4]. Traditionally, such systems are described by deterministic replicator equations [5–7], which corresponds to a mean field approximation in physics. Using techniques from statistical physics, a connection between replicator equations and related microscopic processes can be established [8–12]. Different microscopic dynamics lead to different forms of the replicator equation in the limit of infinite populations. However, usually this affects time scales and does not alter the stability of fixed points and the qualitative dynamics. In structured populations, the microscopic update rule that determines how individual players switch strategies can significantly change the macroscopic dynamics [13–15]. The field of evolutionary game theory has recently tackled the effects of noise in great detail; see [4,16] for recent reviews. Only if such stochastic effects are included can simple rules for the evolution of cooperation on networks [17,18] or group structured populations be derived [19]. New phenomena arise in coordination games under weak selection [20–22] or when payoffs are stochastic [23]. In systems where the replicator equation predicts cyclic behavior, extinction is certain in a stochastic setting [24,25].

Usually, evolutionary game theory addresses the dynamics within a single population and describes how different types spread in this population [13,17,20,26,27]. Bimatrix games (or asymmetric conflicts) address situations in which two different populations with different preferences interact, which is the generic case in economics. In biology, the two populations can be males and females or defenders and attackers of a territory [1]. In social and economic systems, one could think of employees and employers or shoppers and

sellors. This switching between economical and biological thinking is typical for evolutionary game theory, which originated from the application of ideas from economics in biology.

Usually, evolutionary game theory considers the competition between different individuals within a population (e.g., between timid and aggressive males). But one can also compare the average success of one population to the average success of the second population. In these systems, small changes in the details of the replicator dynamics can change the qualitative dynamics and the stability of fixed points [1]. In a previous paper, we have shown that one population that changes its adaptation or learning rate according to its current situation can in the long run outperform a population with a constant learning rate in the presence of external additive noise, even when learning faster alone is not enough [28].

Here, we investigate this “stochastic gain” effect in interacting finite populations and show that even internal noise that naturally arises from the finiteness of a population can be exploited in this way. We start by showing that the effect can be found in a finite population not subject to external noise and demonstrate that our results are qualitatively robust with respect to the microscopic update mechanism. Subsequently, we show examples that an adaptive learning rate (which forms the basis of the “stochastic gain” in Ref. [28]) is not necessary at all if both populations employ different microscopic update mechanisms. In this case, specific microscopic update rules are more successful.

II. THE BASIC MODEL

A. Payoffs

To address the stochastic gain effect in finite populations, we consider two interacting populations X and Y , both of size N . Each individual can choose between two strategies 0 and 1. The payoff in the X population depends on the number of different strategic types in the Y population and vice versa. The payoffs for a single interaction can be written as

$$\begin{array}{c} Y_0 \quad Y_1 \\ X_0 \quad \begin{pmatrix} a_{00}, b_{00} & a_{01}, b_{01} \\ a_{10}, b_{10} & a_{11}, b_{11} \end{pmatrix}. \end{array} \quad (1)$$

Thus, an X_0 individual interacting with a Y_0 individual obtains the payoff a_{00} whereas its interaction partner obtains b_{00} . The average payoff π_0^X of each of the X_0 individuals depends on the number of Y_0 individuals, j , and on the number of Y_1 individuals, $N-j$. It is given by $\pi_0^X = a_{00}j/N + a_{01}(N-j)/N$. Similarly, we have $\pi_1^X = a_{10}j/N + a_{11}(N-j)/N$. In the Y population, the payoffs depend on the number of X_0 and X_1 individuals, i and $i=N-i$, which leads to $\pi_0^Y = b_{00}i/N + b_{10}(N-i)/N$ and $\pi_1^Y = b_{01}i/N + b_{11}(N-i)/N$.

The payoffs now determine the probability that a new strategy is adopted. Different ways to incorporate this strategy dynamics are possible [21,29,30].

B. Population dynamics

Here, we use a microscopic update mechanism based on the pairwise comparison of individuals: One focal individual is selected at random and compares its payoff to a randomly selected role model from the same population. It adopts the strategy of the role model with a probability proportional to the payoff difference. Thus, the probabilities to change the number of X_0 individuals from i to $i \pm 1$ are given by

$$T_X^{i \rightarrow i \pm 1} = \eta_x \left(\frac{1}{2} \pm \frac{w}{2} \frac{\pi_0^X - \pi_1^X}{\Delta \pi_{\max}} \right) \frac{i}{N} \frac{N-i}{N}. \quad (2)$$

Via the payoffs, these transition probabilities depend on the number of different types of the Y population. The learning rate η_x determines how likely it is to change a strategy in the X population. For $\eta_x \ll 1$, it is very unlikely that an individual switches the strategy. In the extreme case of $\eta_x = 0$, the population does not evolve. For the maximum value $\eta_x = 1$, evolution within the X population proceeds at the fastest rate. The intensity of selection w controls how likely it is to adapt to a better (or worse) strategy [10,20]. For $w \ll 1$, selection is weak and the probability to switch to a better strategy is only slightly higher than the probability to switch to a worse strategy. For $w = 1$, the probability to adopt a better strategy reaches one if the payoff difference $\pi_0^X - \pi_1^X$ reaches the maximum possible payoff difference $\Delta \pi_{\max}$. With probability $T_X^{i \rightarrow i} = 1 - T_X^{i \rightarrow i+1} - T_X^{i \rightarrow i-1}$, the number of X_0 and X_1 individuals remains constant. In an equivalent way, we obtain

$$T_Y^{j \rightarrow j \pm 1} = \eta_y \left(\frac{1}{2} \pm \frac{w}{2} \frac{\pi_0^Y - \pi_1^Y}{\Delta \pi_{\max}} \right) \frac{j}{N} \frac{N-j}{N}. \quad (3)$$

The transition probabilities (2) and (3) define a birth-death process in each of the two populations on the state space $\{i, j\} = \{0, 1, \dots, N\}$. The two processes are coupled. For large populations, the dynamics of the system can be approximated by a Fokker-Planck equation with drift term $a(i) = T_X^{i \rightarrow i+1} - T_X^{i \rightarrow i-1}$ and diffusion term $b(i) = \sqrt{[T_X^{i \rightarrow i+1} + T_X^{i \rightarrow i-1}]/N}$ [10]. Applying the Itô calculus [31–33], this description is equivalent to a stochastic differential equation for $x \equiv i/N$ and $y \equiv j/N$,

$$\dot{x} = \frac{\eta_x w}{\Delta \pi_{\max}} (\pi_0^X - \pi_1^X) x(1-x) + \sqrt{\eta_x \frac{x(1-x)}{N}} \xi,$$

$$\dot{y} = \frac{\eta_y w}{\Delta \pi_{\max}} (\pi_0^Y - \pi_1^Y) y(1-y) + \sqrt{\eta_y \frac{y(1-y)}{N}} \xi. \quad (4)$$

Here, ξ is white Gaussian noise with variance 1. Since we describe the coupled dynamics in two populations, we need two replicator equations [34].

With this specific microscopic update mechanism, we now concentrate on a particular game for which the stochastic gain effect is most striking. We consider the payoff matrix

$$\begin{array}{c} Y_0 \quad Y_1 \\ X_0 \quad \begin{pmatrix} +1, -1 & -1, +1 \\ -1, +1 & +1, -1 \end{pmatrix}. \end{array} \quad (5)$$

This game is known as “matching pennies” (in economics) or “battle of the sexes” (in biology). A biological motivation of this game can be found in [1,10,35]. The X population is better off if it uses the same strategy as the Y population, i.e., an X_0 player obtains the highest payoff when matched with a Y_0 player and an X_1 player obtains the highest payoff when matched with a Y_1 player. Individuals from the Y population obtain the highest payoff when paired with a different strategy in the X population, i.e., Y_0 with X_1 and Y_1 with X_0 . In the limit $N \rightarrow \infty$, the multiplicative noise term vanishes and we obtain from Eq. (5) the usual replicator equation. In this case, the quantity $[x(1-x)]^{2\eta_x} [y(1-y)]^{2\eta_y}$ is a constant of motion [36]. The system cycles on closed trajectories around a neutral fixed point at $(x, y) = (1/2, 1/2)$.

C. Dynamics close to the interior fixed point

Due to the symmetry of the game, both populations obtain the same average payoff. The same holds under the influence of external additive noise [28]. In the case of a finite population, the microscopic update mechanism is a natural source of internal noise. Considering the symmetry between Eqs. (2) and (3), both populations have the same average payoff even in this case.

Following Cremer *et al.* [25], we can derive a Fokker-Planck equation and approximate the dynamics of the system close to the neutral fixed. Starting from the deterministic part of the replicator equations (5), we can use the coordinate transformation $\hat{x} = x - \frac{1}{2}$ and $\hat{y} = y - \frac{1}{2}$. Further, we assume that $\eta_x = \eta_y$ and set $\hat{\omega} = w \eta_x / \Delta \pi_{\max} = w \eta_y / \Delta \pi_{\max}$. Thus, we arrive at

$$\begin{aligned} \partial_t \hat{x} &= -\hat{\omega} \hat{y} (4\hat{x}^2 - 1), \\ \partial_t \hat{y} &= +\hat{\omega} \hat{x} (4\hat{y}^2 - 1). \end{aligned} \quad (6)$$

The Fokker-Planck equation, which describes the time evolution of the probability density P (where we dropped the dependence on \hat{x} , \hat{y} , and t) is in its general form given by

$$\partial_t P = - \sum_k \partial_k a_k P + \frac{1}{2} \sum_{k,l} \partial_k \partial_l \beta_{kl} P. \quad (7)$$

In our case, the drift vector is $\mathbf{a}=(\partial_t \hat{x}, \partial_t \hat{y})$. The diffusion matrix is diagonal with $\beta_{\hat{x}\hat{x}}=(\frac{1}{4}-\hat{x}^2)/N$, $\beta_{\hat{y}\hat{y}}=(\frac{1}{4}-\hat{y}^2)/N$ and $\beta_{\hat{x}\hat{y}}=\beta_{\hat{y}\hat{x}}=0$.

Due to the symmetry of the problem we can introduce polar coordinates $\hat{x}=r \cos \varphi$ and $\hat{y}=r \sin \varphi$, so the equations read

$$\begin{aligned} \partial_t r &= -r^3 \hat{\omega} \sin(4\varphi), \\ \partial_t \varphi &= \hat{\omega} [r^2 - \cos(4\varphi)r^2 - 1]. \end{aligned} \quad (8)$$

Note that these two functions are invariant under rotation by $\pi/2$. Close to $r=0$, where we can approximate the diffusion term by $1/(4N)$, the Fokker-Planck equation becomes in polar coordinates

$$\partial_t P = \hat{\omega} r^3 \sin(4\varphi) \partial_r P + \hat{\omega} [1 - r^2 + r^2 \cos(4\varphi)] \partial_\varphi P + \frac{1}{4N} \Delta_r P. \quad (9)$$

Here $P=P(r, \varphi, t)$ and

$$\Delta_r = \frac{1}{r} \frac{\partial}{\partial r} r \frac{\partial}{\partial r} + \frac{1}{r^2} \frac{\partial^2}{\partial \varphi^2}$$

is the Laplacian in polar coordinates. We are interested in the radial dependence only close to the interior fixed point, $r \ll 1$. Therefore, we average over the angular dependence by replacing all angle-dependent terms by their mean value, for example $\langle \sin 4\varphi \rangle = \int_0^{2\pi} \sin 4\varphi d\varphi = 0$. Then the equation reduces to $\partial_t P = (1/4N) \Delta_r P$. Since the constant $P(r)=P_0$ is a solution of this differential equation, one can expect that a stationary solution will be constant close to the interior fixed point at $r=0$. Because of the absorbing boundaries of the process, we can at most expect quasistationary behavior until the boundaries are reached. This is of particular interest if the time until the boundaries are reached is large [25].

It can be shown that the quantity $H(x, y) = -x(1-x)y(1-y)$ is a constant of motion of the replicator equations of the system, Eq. (5). However, $H(x, y)$ is not constant for all update mechanisms. In general, the change in $H(x, y)$ determines whether the system spirals towards the neutral fixed point $(\frac{1}{2}, \frac{1}{2})$ or away from it. In such cases, $H(x, y)$ can often be interpreted as a Lyapunov function which determines whether the fixed point is asymptotically stable [25, 28]. In [37], it has been shown that the average drift of the local-update system away from the fixed point can be determined by averaging over $H(x, y)$. Thus, the change of $H(x, y)$ allows us to infer if the stochastic system drifts away from the fixed point or towards it. Only in the latter case can, a quasistationary solution be obtained. Averaging over the area around the fixed point in polar coordinates leads to

$$\begin{aligned} \langle \Delta H \rangle &= \frac{1}{\pi r^2} \frac{2}{N^2} \int_0^{2\pi} d\varphi \int_0^r dr' r' H \left(\frac{1}{2} - r' \cos \varphi, \frac{1}{2} - r' \sin \varphi \right) \\ &= \frac{1}{N^2} \left(\frac{1}{8} - \frac{r^2}{4} + \frac{r^4}{12} \right) > 0. \end{aligned} \quad (10)$$

Because this quantity is positive for finite N and small r , the system will leave the vicinity of the interior fixed point after a short time. This is different for the Moran process, as shown in [25]. Both results are compatible with the picture for $N \rightarrow \infty$, where the adjusted replicator dynamics associated with the Moran process predicts a stable fixed point and the usual replicator dynamics associated with the local update process discussed here predicts a neutrally stable fixed point.

III. ADAPTIVE LEARNING RATE

Usually, the learning rates η_x and η_y have the same fixed value. Szolnoki and Szabó have analyzed a system in which different players have different, but fixed values of η [38]. Here, we assume that the rate with which a population adapts to a new situation is different depending on the average success of a population. If the overall success is unsatisfactory, one is more likely to try something new.

We follow our previous work based on the replicator equation [28] and introduce an adaptive learning rate η_x for the X population. This rate η_x that increases above a base value η_0 when the average payoff difference to the Y population is negative. It decreases below η_0 when the payoff difference is positive. In our case, the adaptive learning rate is given by

$$\eta_x = \eta_0 [1 - \tanh(\alpha_x \Delta \pi)]. \quad (11)$$

Here, $\Delta \pi = \langle \pi^X \rangle - \langle \pi^Y \rangle$ is the difference between the average payoffs, which are defined by $\langle \pi^X \rangle = [\pi_0^X i + \pi_1^X (N-i)]/N$ and $\langle \pi^Y \rangle = [\pi_0^Y j + \pi_1^Y (N-j)]/N$. The parameter η_0 is set to $\frac{1}{2}$, and α_x determines how sensitive this change in the learning rate is. If the X population has a larger average payoff than the Y population, adaptation of strategies becomes slower. If the X population has a smaller average payoff than the Y population, adaptation of strategies becomes faster. Thus, we introduce an adaptive learning rate that follows the “win stay–lose shift” paradigm in the comparison between the two populations.

A. Local update mechanism

Our system reduces to the one analyzed in [28] in the limit of $N \rightarrow \infty$, but then includes no noise term. Moreover, the nature of the noise is very different. In [28], the noise is additive and arises from an external source. Here, the noise is multiplicative and arises internally from the system. In Fig. 1, we demonstrate that for finite populations the stochastic gain effect can still be found with the additional twist that now internal noise is exploited. The noise intensity is controlled by w : Small w implies that the system is close to neutral selection where all individuals have the same fitness. In this case, it is difficult to increase payoffs solely by changing the rate of adaption. Large w means that the system is

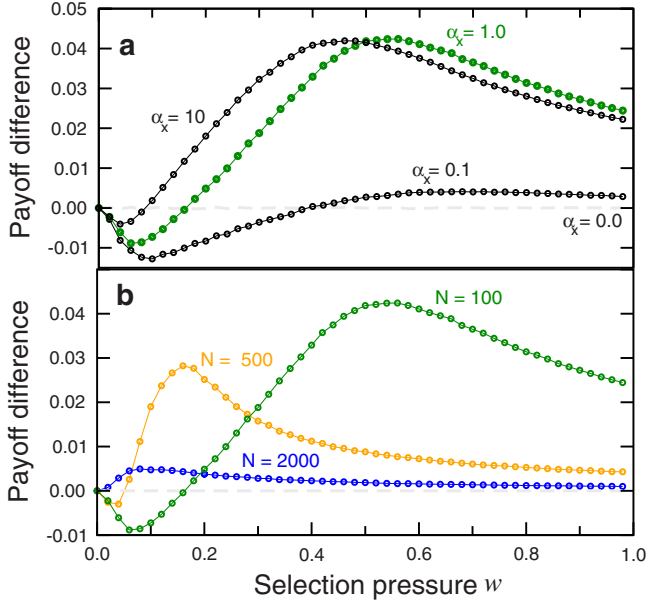


FIG. 1. (Color online) Average payoff difference of a population with an adaptive learning rate $\eta_x = 1/2 - (1/2)\tanh(\alpha_x \Delta \pi)$ compared to a population with a constant learning rate of $\eta_y = 0.5$ in the matching pennies game using the local update rule (see text). (a) Payoff differences for a constant population size of $N = 100$ and for different α . The average payoff difference is small if the intensity of selection w is small otherwise the population with an adaptive learning rate has a higher average payoff. The stochastic gain effect becomes more pronounced with increasing α (averages over 2×10^4 random initial conditions in the interior until the maximum time $T = 10^5$ or until the absorbing boundaries of the system are reached). (b) Payoff differences for three different population sizes for fixed $\alpha = 1.0$. The total payoff decreases with higher population size N because the noise intensity decreases. In the limit of $N \rightarrow \infty$ we obtain the equation for the stochastic replicator dynamic without external noise. Thus the payoff difference converges to zero. Averages over 7000 random initial conditions until the absorbing boundaries are reached or until the maximum time $T = N \times 10^3$.

very unlikely to leave the area close to the point $(x, y) = (0.5, 0.5)$ where both populations have the same payoff. Thus, there exists an optimal w for which the payoff difference becomes highest. For $N \rightarrow \infty$, the internal noise vanishes and the effect is no longer present. We have started with a random initial condition and simulated the system until one strategy in one population has reached extinction.

Figure 1 shows the mean payoff difference averaged over a large number of such runs. Figure 1(a) shows that the effect becomes more pronounced with increasing α (for very large α , the effect becomes weaker again, see below for an explanation). Figure 1(b) shows that, with increasing N , the position of the maximum changes to smaller intensity of selection. This is consistent with the fact that often the effective selection pressure is given by Nw [20]. In addition, the size of the maximum decreases.

Next, we show that for $\alpha = \alpha_x > 0$ and $\alpha_y = 0$, the rotational symmetry of the system is changed. By considering $\alpha \ll 1$, we can replace the hyperbolic tangent with its argument ($\tanh \alpha \approx \alpha$). We also set $\omega = w \eta_0 / \Delta \pi_{\max}$. In this case, Eq. (8) becomes

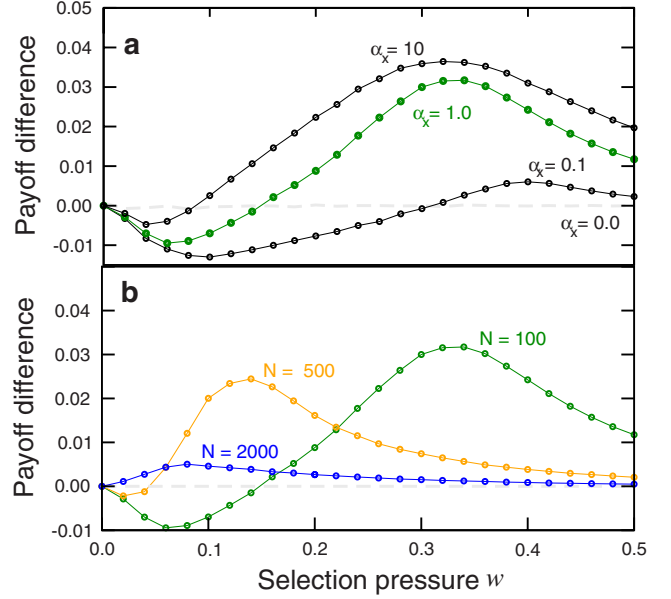


FIG. 2. (Color online) Average payoff difference of a population with an adaptive learning rate $\eta_x = 1/2 - (1/2)\tanh(\alpha_x \Delta \pi)$ against a population with a constant learning rate of $\eta = 0.5$ in the matching pennies game using the frequency-dependent Moran process in both populations. The parameter values in (a) and (b) are identical to those of Fig. 1, but in the Moran process shown here the maximum intensity of selection is given by $w = 0.5$. Qualitatively, the stochastic gain effect does not depend on the details of the update mechanism in finite populations: With α increasing from $\alpha = 0.0$, the payoff advantage of the adaptive population increases. However, there is an optimal α for which the stochastic gain effect is most pronounced; see text. With increasing N , the system approaches a deterministic replicator system and the intrinsic noise vanishes. Thus, increasing N leads to smaller payoff differences. Moreover, the finite-size effect of a negative payoff difference for low intensity of selection vanishes.

$$\begin{aligned} \partial_t r &= -r^3 \omega \sin(4\varphi) + 2r^3 \alpha \omega [2r^2 + 2r^2 \cos(2\varphi) - 1] \sin(2\varphi)^2, \\ \partial_t \varphi &= \omega [-\cos(4\varphi)r^2 + r^2 - 1] - 8r^2 \alpha \omega \cos(\varphi) \\ &\quad \times [4r^2 \cos(\varphi)^2 - 1] \sin(\varphi)^3. \end{aligned} \quad (12)$$

The α -independent terms are invariant under rotation by $\pi/2$, whereas the α -dependent terms are only invariant under rotation by π . This results from the asymmetry of the system, as the population with the adaptive learning rate leads to a higher stationary density in the areas where it has a higher payoff. In other words, $\alpha > 0$ breaks the rotational symmetry of the system. In principle, a Fokker-Planck equation for $\alpha > 0$ can now be derived. However, neglecting the angular dependence recovers the result for $\alpha = 0$. The full solution of the Fokker-Planck equation showing the effect of $\alpha > 0$ can only be done numerically.

B. Moran process

Next, we show that the stochastic gain effect in finite populations does not depend on the details of the microscopic update mechanism. A standard approach for game dy-




FIG. 3. (Color online) Payoff distribution and stationary distribution for the Moran process in the strategy space spanned by the state space for a constant population size of $N=100$, encoded by a color scale where bright colors indicate high values. (a) Average payoffs $\langle \pi^X \rangle = [\pi_0^X i + \pi_1^X (N-i)]/N$ of the adaptive populations (X) are shown. In the bottom left and top right areas (green, +) the payoffs of the adaptive population (X) are positive, whereas in the bottom right and top left areas (red, -) the payoffs are negative. The adaptive population can obtain a higher stationary probability density in the bottom left and top right areas, leading to the stochastic gain effect. (b) Stationary distribution for $\alpha=0.0$, invariant under rotation by $\pi/2$ and approximately rotationally invariant close to the fixed point of the replicator dynamics at $(x,y)=(0.5,0.5)$. The population dynamics drives the system around this point. (c) With increasing $\alpha=1.0$ the system is driven to the interior. Now the areas where the average payoff for the adaptive population is higher shows a larger stationary probability density. (d) Same as in (c) but the stationary distribution for $\alpha=10.0$ is shown. Increasing α further leads to a smaller probability density in the areas where the payoff difference is high and thus, the payoff difference decreases again for large α (for all panels, parameter values are $N=100$, $R=5 \times 10^5$ independent realizations, $w=0.35$, maximum number of time steps $T=N \times 10^3$).

namics in finite populations is the frequency-dependent Moran process [19,21,39,40]. In each of the two populations, the following update process takes place.

One individual is selected for reproduction with a probability proportional to the fitness of the individual. We define fitness f as a convex combination of a background fitness set to 1 and the payoff, e.g., $f_0^X = 1 - w + w\pi_0^X$. Since our payoffs vary between -1 and 1 , the intensity of selection w has an upper limit, $w < 0.5$ to ensure that fitness is positive. The selected individual produces identical offspring, which replaces a randomly chosen individual. The transition probabilities in the X population are thus given by

$$T_X^{i \rightarrow i+1} = \frac{\eta_x}{2} \frac{f_0^X i}{f_0^X i + f_1^X (N-i)} \frac{N-i}{N}, \quad (13)$$

$$T_X^{i \rightarrow i-1} = \frac{\eta_x}{2} \frac{f_1^X (N-i)}{f_0^X i + f_1^X (N-i)} \frac{i}{N}, \quad (14)$$

where we have introduced a factor $1/2$ to make both considered processes identical for $w \rightarrow 0$. Similar equations hold for the Y populations.

In the Moran process, strategies with higher fitness are more likely to be selected for reproduction, whereas selection at death is the same for all strategies. Thus, the average abundance of fitter strategies increases over time. The frequency-dependent Moran process reduces to the adjusted replicator dynamics in the limit $N \rightarrow \infty$ [10]. For the adjusted replicator dynamics it is known that the game defined by the payoff matrix Eq. (5) has an asymptotically stable fixed point at $(x,y)=(0.5,0.5)$ [1]. Moreover, in finite populations the probability density is centered around this stable fixed point

for the Moran process, in contrast to the local update process discussed above. Thus, it is not *a priori* clear that the stochastic gain effect can be observed even in the frequency-dependent Moran process. However, Fig. 2 shows that a variable rate of adaptation also leads to an increased payoff in this process.

To illustrate the effect further, we consider the probability density for the different states of the system numerically. Figure 3(b) shows that for $\alpha_x=0$ this probability density is symmetric in the state space $(i, j) \in (\{0, \dots, N\}, \{0, \dots, N\})$ leading to a vanishing average payoff difference. With increasing α_x , the time the system spends in states in which the X population has the higher average payoff increases. However, if α_x becomes too large, then adaptation of the X population is so fast that random fluctuations no longer take the system to states in which the payoff difference is large, see Fig. 3(a). Thus, there exists an optimal α_x for which the payoff difference becomes highest for any given w .

IV. COMPETITION OF UPDATE MECHANISMS

The choice of the adaptive learning rate in Eq. (11) not only incorporates the state of the population, but requires also information on the competing population. In reality, such information might not be available and information access is restricted to the own population. However, the exploitation of internal (or external) noise is still possible if both populations update their strategies by different learning schemes.

So far, we have assumed that both populations use the same update mechanism and that only the adaptive learning in one population is different. However, one can also consider two different update mechanisms in the two populations. This is motivated by the following observation [10]. The local update mechanism leads to the standard replicator dynamics $\dot{x} = x(\pi^x - \langle \pi \rangle)$ in the limit $N \rightarrow \infty$. The frequency-dependent Moran process leads to the adjusted replicator dynamics $\dot{x} = x(\pi^x - \langle \pi \rangle) / \langle \pi \rangle$ in this limit. The only difference is that the right-hand side is divided by the average payoff. However, this can be interpreted as a change in the learning rate, as a large average payoff leads to slow dynamics and a small average payoff to fast dynamics. Hence, a population using the frequency dependent Moran process outperforms a population using the local update mechanisms, without requiring any knowledge about the difference between the two populations; see Fig. 4. This result becomes more pronounced for higher intensities of selection. However, in the processes discussed here there is an upper limit in the intensity of selection. To address this issue, one has to resort to different microscopic update mechanisms [41].

V. DISCUSSION

Here, we have demonstrated that the stochastic gain effect discussed in [28] can also be found in finite populations. In this case, the internal noise arising from the finiteness of the population can be exploited. The intensity of this noise can be controlled by the intensity of selection, which determines

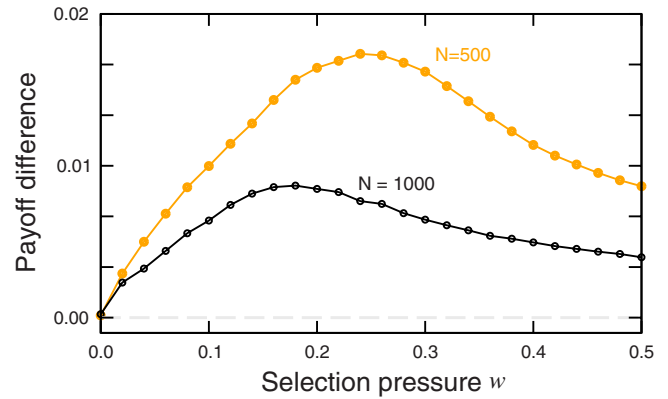


FIG. 4. (Color online) Average payoff difference of a population using the Moran process against a population using the local update rule for two different sizes of population N . Both populations do not change their learning rate ($\eta_0=0.5$). Independent of the intensity of selection w the population using the Moran process obtains a higher payoff. Thus, the stochastic gain effect can also be observed in the absence of variable learning rates. The total payoff decreases with higher population size N . In the limit of $N \rightarrow \infty$ both dynamics result in the equation for the deterministic replicator dynamic, without external noise thus the payoff difference tends to zero. Parameter values are $N=500, 1000$, $T=N \times 10^3$ maximum number of time steps, and $R=15\,000$ realizations.

how likely it is that individuals adopt better (or worse) strategies.

An increased average payoff can be obtained from a different microscopic update mechanism in such situations. In principle, this could lead to higher-level selection: If competition does not only occur within populations, but also between groups of individuals [19], then groups using a more successful update mechanism will perform better. In this way, a more successful update mechanism might be selected in the long run. But this result can also be interpreted as a warning: In multiagent simulations, one should choose an update mechanism carefully and be aware of the consequences of this choice.

Modeling such systems with deterministic techniques such as the replicator equations becomes relevant only when noise has a small influence on the system, as, e.g., in large populations or under strong selection. Recently, there has been an increased interest in stochastic effects in evolutionary games. Perc and Marhl have shown that in spatial games coherence resonance can be observed [42,43]. It has been shown that noise can even enhance cooperation in many circumstances [44–46]. Internal as well as external noise has been considered in these systems and it has been shown that the nature of these disturbances can be very different [47,48].

The stochastic gain effect in finite populations becomes weaker with increasing population size, as the noise intensity of the internal noise decreases. Moreover, there is an optimal speed of adaption. Too slow adaptation cannot lead to an exploitation of internal noise. For too strong adaptation, the player with adaptive learning rate prevents the system from reaching regions where his payoff is highest by trying to exploit even small differences. The existence of the stochastic gain effect without any adaptive learning rates ($\eta_x = \eta_y$

=0.5) is based on the application of different microscopic update rules, which seems to be natural when two different populations or players interact.

One way to model such interactions is the minority game, which became a paradigm as a simplified market model [49–52]. In this game, there is an odd number of players, who choose between buying and selling. The group that is in the minority wins, since they have a better position negotiating the price. Players are usually modeled as heterogeneous agents who have different ways to incorporate available information about the game and convert it into optimal actions for the future. Thus, memory and the differences between agents form the basis for the complex dynamics. Common practice in evolutionary game theory takes very different assumptions: The population consists of identical players (since everyone can buy or sell) who do not possess any memory. Only the last interaction determines the future. From this perspective, the minority games would reduce to a special case of a 2×2 game with a mixed evolutionary stable strategy. However, the key features of the minority

game are lost in this way. Thus, there is no simple mapping of the minority game to standard evolutionary game dynamics.

A different way to analyze learning rules or update mechanisms has been discussed in [53]. There, a population of buyers and sellers interact and the performance of the different learning mechanisms are analyzed. Our results highlight the importance of the detailed consideration of such microscopic update mechanisms or learning schemes, as they can significantly alter the macroscopic dynamics of a system.

ACKNOWLEDGMENTS

We thank Jorge M. Pacheco for helpful comments on an earlier version of this manuscript. A.T. acknowledges support by the Deutsche Akademie der Naturforscher Leopoldina (Grant No. BMBF-LPD 9901/8-134) and the Emmy-Noether program of the German Research Foundation.

-
- [1] J. Maynard Smith, *Evolution and the Theory of Games* (Cambridge University Press, Cambridge, U.K., 1982).
- [2] M. A. Nowak and K. Sigmund, *Science* **303**, 793 (2004).
- [3] M. A. Nowak, *Evolutionary Dynamics* (Harvard University Press, Cambridge, MA, 2006).
- [4] G. Szabó and G. Fáth, *Phys. Rep.* **446**, 97 (2007).
- [5] P. D. Taylor and L. Jonker, *Math. Biosci.* **40**, 145 (1978).
- [6] E. C. Zeeman, in *Population dynamics from game theory*, edited by A. Nitecki and C. Robinson, *Lecture Notes in Mathematics* (Springer, Berlin, 1980), Vol. 819.
- [7] J. Hofbauer and K. Sigmund, *Evolutionary Games and Population Dynamics* (Cambridge University Press, Cambridge, U.K., 1998).
- [8] D. Helbing, *Physica A* **193**, 241 (1993).
- [9] D. Helbing, *Theory Decis.* **40**, 149 (1996).
- [10] A. Traulsen, J. C. Claussen, and C. Hauert, *Phys. Rev. Lett.* **95**, 238701 (2005).
- [11] A. Traulsen, J. C. Claussen, and C. Hauert, *Phys. Rev. E* **74**, 011901 (2006).
- [12] F. A. C. C. Chalub and M. O. Souza, *Math. Comput. Modell.* **47**, 743 (2008).
- [13] H. Ohtsuki, M. A. Nowak, and J. M. Pacheco, *Phys. Rev. Lett.* **98**, 108106 (2007).
- [14] F. C. Santos, J. F. Rodrigues, and J. M. Pacheco, *Proc. R. Soc. London, Ser. B* **273**, 51 (2006).
- [15] F. C. Santos, J. M. Pacheco, and T. Lenaerts, *Proc. Natl. Acad. Sci. U.S.A.* **103**, 3490 (2006).
- [16] J. Miekisz, e-print arXiv:q-bio.PE/0703062v1.
- [17] H. Ohtsuki, C. Hauert, E. Lieberman, and M. A. Nowak, *Nature (London)* **441**, 502 (2006).
- [18] H. Ohtsuki, J. Pacheco, and M. A. Nowak, *J. Theor. Biol.* **246**, 681 (2007).
- [19] A. Traulsen and M. A. Nowak, *Proc. Natl. Acad. Sci. U.S.A.* **103**, 10952 (2006).
- [20] A. Traulsen, J. M. Pacheco, and L. A. Imhof, *Phys. Rev. E* **74**, 021905 (2006).
- [21] M. A. Nowak, A. Sasaki, C. Taylor, and D. Fudenberg, *Nature (London)* **428**, 646 (2004).
- [22] H. Ohtsuki, P. Bordalo, and M. A. Nowak, *J. Theor. Biol.* **249**, 289 (2007).
- [23] C. P. Roca, J. A. Cuesta, and A. Sánchez, *Phys. Rev. Lett.* **97**, 158701 (2006).
- [24] T. Reichenbach, M. Mobilia, and E. Frey, *Phys. Rev. E* **74**, 051907 (2006).
- [25] J. Cremer, T. Reichenbach, and E. Frey, *Eur. Phys. J. B* **63**, 373 (2008).
- [26] M. A. Nowak and K. Sigmund, *Nature (London)* **355**, 250 (1992).
- [27] F. C. Santos and J. M. Pacheco, *Phys. Rev. Lett.* **95**, 098104 (2005).
- [28] A. Traulsen, T. Röhl, and H. G. Schuster, *Phys. Rev. Lett.* **93**, 028701 (2004).
- [29] L. A. Imhof and M. A. Nowak, *J. Math. Biol.* **52**, 667 (2006).
- [30] A. Traulsen, M. A. Nowak, and J. M. Pacheco, *Phys. Rev. E* **74**, 011909 (2006).
- [31] C. W. Gardiner, *Handbook of Stochastic Methods for Physics, Chemistry and the Natural Sciences*, *Springer Series in Synergetics* Vol. 13 (Springer, Berlin, 1983).
- [32] N. G. v. Kampen, *Stochastic Processes in Physics and Chemistry*, 2nd ed. (Elsevier, Amsterdam, 1997).
- [33] H. G. Schuster, *Complex Adaptive Systems* (Scator, Saarbrücken, 2002).
- [34] T. Galla, *Europhys. Lett.* **78**, 20005 (2007).
- [35] R. Dawkins, *The Selfish Gene* (Oxford University Press, New York, 1976).
- [36] Note that $\Delta\pi_{\max}=2$.
- [37] J. C. Claussen, *Eur. Phys. J. B* **60**, 391 (2007).
- [38] A. Szolnoki and G. Szabó, *Europhys. Lett.* **77**, 30004 (2007).
- [39] C. Taylor, D. Fudenberg, A. Sasaki, and M. A. Nowak, *Bull. Math. Biol.* **66**, 1621 (2004).

- [40] T. Antal and I. Scheuring, *Bull. Math. Biol.* **68**, 1923 (2006).
- [41] This can be done by choosing the pairwise comparison process [30,54] instead of the local update process and a Moran process with exponential rather than linear payoff to fitness mapping [55].
- [42] M. Perc, *New J. Phys.* **8**, 22 (2006).
- [43] M. Perc and M. Marhl, *New J. Phys.* **8**, 142 (2006).
- [44] G. Szabó, J. Vukov, and A. Szolnoki, *Phys. Rev. E* **72**, 047107 (2005).
- [45] A. Sánchez and J. A. Cuesta, *J. Theor. Biol.* **235**, 233 (2005).
- [46] M. Perc, *New J. Phys.* **8**, 183 (2006).
- [47] J. C. Claussen and A. Traulsen, *Phys. Rev. E* **71**, 025101(R) (2005).
- [48] M. Perc, *Phys. Rev. E* **75**, 022101 (2007).
- [49] D. Challet and Y.-C. Zhang, *Physica A* **246**, 407 (1997).
- [50] D. Challet, M. Marsili, and Y.-C. Zhang, *Minority Games: Interacting Agents in Financial Markets* (Oxford University Press, Oxford, 2004).
- [51] A. C. C. Coolen, *The Mathematical Theory of Minority Games: Statistical Mechanics of Interacting Agents* (Oxford University Press, Oxford, 2005).
- [52] T. Galla and D. Sherrington, *Eur. Phys. J. B* **46**, 153 (2005).
- [53] J. Jost and W. Li, *Physica A* **345**, 245 (2005).
- [54] G. Szabó and C. Tóke, *Phys. Rev. E* **58**, 69 (1998).
- [55] A. Traulsen, N. Shresh, and M. A. Nowak, *Bull. Math. Biol.* **63**, 363 (2008).