

Learning to play public good games

Alex J. Bladon* and Tobias Galla†

*Theoretical Physics, School of Physics and Astronomy,
University of Manchester, Manchester M13 9PL, United Kingdom*

(Dated: July 6, 2011)

We extend recent analyses of stochastic effects in game dynamical learning to cases of multi-player games, and to games defined on networked structures. By means of an expansion in the noise strength we consider the weak-noise limit, and present an analytical computation of spectral properties of fluctuations in multi-player public good games. This extends existing work on two-player games. In particular we show that coherent cycles may emerge driven by noise in the adaptation dynamics. These phenomena are not too dissimilar from cyclic strategy switching observed in experiments of behavioural game theory.

PACS numbers: 02.50.Le, 05.10.Gg, 02.50.Ey, 87.23.Kg

I. INTRODUCTION

The theory of evolutionary dynamics is a cornerstone of modern biology. It has helped to describe a variety of systems where distinct elements compete to reproduce, for example in zoology and population dynamics [1], but also in the dynamics of cancer growth and the different possible progressions of HIV [2], or in the evolution of language [3]. Since the work of Maynard Smith [1] evolutionary ideas have also been applied in the context of game theory, and have augmented the more traditional approach to strategic decision making, based on equilibrium concepts [4–6]. There is now an established field referred to as ‘evolutionary game theory’ (see e.g. [7, 8]). This discipline is concerned with the study of populations of players, who interact in a game and who each carry a strategy and resulting fitness as a result of their success or otherwise in the game. Successful individuals then reproduce faster than those who are less successful in playing the game. Strategies are passed on from parent to offspring, with or without mutation, and the strategic content of the population of individuals hence evolves in time. Traditionally such processes have been modelled by means of replicator equations (or similar dynamical systems). These descriptions are typically deterministic in nature, and systematically neglect stochastic effects. More recently, the role of stochasticity and intrinsic noise has been considered in more detail, and the development of a mathematical theory with which to systematize these effects is very much work in progress. Phenomena brought about purely by stochasticity, and hence not captured by deterministic approaches, include for example fixation and the loss of biodiversity [9, 10], drift reversal [11, 12], or coherent stochastic oscillations [13–15].

With an increasing number of experiments involving real human players being performed in laboratories of behavioural economics, ideas from evolutionary game theory have also been applied to describe the adaptation processes of game learning, see for example [16–18]. In this context, the strategic choices themselves are ‘genetic’ elements. A ‘reproduction event’ corresponds to an instance in which a human decides to switch from one strategy to another. These processes are frequently described mathematically by models drawn from evolutionary game theory. The advantage of using evolutionary game theory rests in the fact that there is a vast body of theoretical work available, ready to be tied to experiments of behavioural game theory. This approach does have its limitations though. Models of evolutionary game theory describe *populations* of agents, subject to selection pressure and generally involving birth and death processes. For obvious reasons no such birth-death dynamics takes place in a behavioural economist’s laboratory, instead a number of fixed individuals interacts repeatedly in such experiments, without specific individuals being removed or replaced, born or killed.

It may therefore seem more appropriate to model the psychological decision making processes directly on the level of the interacting individual, rather than on a population level. Examples of such models can be found in [18–24]. As one main ingredient of these psychologically motivated models players keep so-called ‘attractions’ or ‘propensities’ for each of the possible actions. These are then converted into a mixed strategy profile, and specific moves (pure actions) are played with the corresponding frequencies. Payoffs are received, depending on one’s own move as well as on those made by the relevant opponents. In response to the outcome, players then update their propensities, increasing those of strategies that would have been successful against the observed moves, and reducing those of less successful actions. The process then iterates.

On the mathematical level, models of such pro-

*Electronic address: alex.bladon@postgrad.manchester.ac.uk

†Electronic address: tobias.galla@manchester.ac.uk

cesses define stochastic dynamical systems in discrete time. These models are not too dissimilar from learning algorithms studied in machine learning, including for example dynamical processes such as fictitious play and derivatives [24–26]. Stochasticity in learning can here have profound effects, depending on the details of the learning model adaptation can be seen to converge to Nash equilibria, or the learning process may fail to reach a stationary point due to persistent fluctuations and noise.

Existing work on the analytical characterisation of noisy trajectories in game learning [27, 28] has revealed close similarities with stochastic effects in evolutionary processes in finite populations, but differences have been found as well [29]. Analytical studies have mostly been limited to the case of two-player learning i.e., games in which two individuals interact repeatedly and adapt to each other’s moves. The correlation properties of fluctuations about limiting deterministic trajectories have been computed based on an expansion in the inverse noise strength, in good agreement with numerical simulations. The purpose of the present paper is to extend these studies and to address multi-player learning and learning of agents arranged on a (fixed) network. Stochastic effects in network learning have been considered before [29], but no systematic analytical characterisation has been attempted. As one result of our analysis we will demonstrate that amplified stochastic oscillations, commonly found in models of evolutionary game theory, also appear in multi-player and network learning. We also analyse how the networks structure and parameters of the earning rule affect the outcome of adaptation.

The remainder of the paper is organised as follows: in Sec. II we outline a simple model of reinforcement learning and define the public good game we will be studying subsequently. Sec. III contains a brief analysis of a deterministic limiting case of the learning process. We then classify the oscillations associated with stochastic multi-player learning in Sec. IV. We analyse the resulting power spectra and discuss their implications for game learning in Sec. V. Network games are then considered in Sec. VI. Sec. VII contains concluding remarks and an outlook on possible future lines of research.

II. MODEL AND DEFINITIONS

In this section we define the public goods game and the learning algorithm which we will analyze in the subsequent sections.

A. Public good game

In a typical public good game each player decides whether or not to contribute an amount c to a shared ‘pot’, for example a contribution to a common effort. The amount in the pot is then multiplied by a factor $r > 1$ and re-distributed amongst all players involved in the game, no matter whether they contributed or not. If a player does not contribute they still receive a share of the pot, however the fewer people contribute the smaller each individual’s share becomes. Thus the public goods game can be thought of as a multi-player social dilemma, similar to the celebrated two-person prisoner’s dilemma. There is a temptation to defect (i.e., not to contribute), but the reward for a group consisting only of defectors is less than that for a group of contributors. In the public goods game outlined by [17] this basic setup is extended to allow players not to participate in the game at all. Such players, called ‘loners’, instead receive a small, but guaranteed payout of σc , where σ is some non-zero constant, chose such that σc is less than $(r - 1)c$. This produces a cyclic relationship between the three actions: if everyone is contributing the best thing to do is defect; if everyone is defecting, playing ‘loner’ is the best option and if everyone is a loner then the greatest payoff comes from contributing.

The net payoff for each strategy (defined as the difference between the money received at the end and the money put in at the beginning) can be written as

$$\begin{aligned}\pi_C &= -c + rc \frac{n_C + 1}{1 + n_C + n_D}, \\ \pi_D &= rc \frac{n_C}{1 + n_C + n_D}, \\ \pi_L &= \sigma c,\end{aligned}\tag{1}$$

where n_C and n_D are numbers of contributors and defectors in the *rest* of the group interacting in a particular instance of the game. In other words, the above payoff π_C is that of a co-operator playing with a group of n_C *other* co-operators and n_D defectors, and similar for π_D . At variance with some definitions of public good games, we here assume that the above payoff relations also hold if only one single player decides to participate in the game, and if all other players decide to abstain. This choice is mainly made for analytical convenience. Unless otherwise specified we will use $r = 1.8$, $c = 1$ and $\sigma = 0.5$ (as in [17]). We will abbreviate the three pure strategies as C , D and L in the following.

B. Reinforcement learning

We consider the game learning model proposed by [18]. In such models, N players play a game attempting to maximise their payoffs. In each round

of the game each player uses one of the S available pure strategies, each with a probability defined by the strategy's previous performance as perceived by the player. This will be specified further below. For the case of the public goods game we have $S = 3$ pure strategies, contribute (C), defect (D) and 'loner' (L), but our theory is applicable for general S .

1. Update of propensities

The probabilities with which a given player chooses to play the three actions form a mixed strategy profile in the language of game theory. In our learning model the strategy profile will be based on so-called 'scores' (or propensities) that each player assigns to the pure strategies available to him/her. These scores measure the performance of the pure actions against the observed actions of opponents. The score given to strategy $s \in \{1, \dots, S\}$ by player i is assumed to be updated once every Ω rounds of the game, and kept constant inbetween. Real-world play corresponds to $\Omega = 1$, i.e. updating after each round. For our purposes it is however useful to introduce the more general dynamics as (see also [27, 28, 30])

$$q_{i,s}(t + \Omega) = (1 - \lambda)q_{i,s}(t) + \frac{1}{\Omega} \sum_{t'=t}^{t'+\Omega-1} u(s, \mathbf{s}_{-i}(t')), \quad (2)$$

where $q_{i,s}$ is assumed to remain unchanged between time t and time $t + \Omega - 1$. Here $u(s, \mathbf{s}_{-i}(t))$ is the payoff for player i when they play strategy s at time t given that the other players play the pure strategies denoted by $\mathbf{s}_{-i} \in \{1, \dots, S\}^{N-1}$. In our case s can take the values $s \in \{C, D, L\}$, and similarly each entry in \mathbf{s}_{-i} takes values in $\{C, D, L\}$. The payoff $u(s, \mathbf{s}_{-i}(t))$ is given by expressions (1), where n_C and n_D are the number of entries C and D respectively in the vector \mathbf{s}_{-i} . Some more explanation of the update rule of Eq. (2) is here appropriate. The above learning rule assumes that adaptation occurs only once every Ω iterations of the game. This corresponds to what is known as batch learning [31]. The last term on the RHS of Eq. (2) is indeed the *average* payoff per round obtained by player i in Ω iterations of the game, and given his or her opponents' actions $\mathbf{s}_{-i}(t')$ during those Ω rounds ($t' = t, t + 1, \dots, t + \Omega - 1$). The philosophy behind introducing batch-learning dynamics is discussed further below in the context of the deterministic limit of the stochastic dynamics (Sec. II B 4). The pre-factor $1 - \lambda$ in the first term on the RHS of the update rule describes memory loss. For $\lambda = 0$ the score $q_{i,s}(t)$ is proportional to the payoff player i would have received up to time t had he or she played action i at all times, and given his or her opponents' actions. All past play carries equal weight and is accumulated. The parameter $\lambda \in [0, 1]$ is a memory-loss parameter. For $\lambda > 0$ past play

is exponentially discounted, and observations in the distant past carry a lesser weight than more recent iterations of the game. Similar mechanisms are present in learning models of behavioural game theory, see for example [20], they have also been used in [30]. Occasionally a pre-factor λ is introduced in the second term on the RHS of Eq. (2). This is mainly a matter of notation, and amounts to a re-scaling of payoffs. We here use the notation of [18, 30].

2. Fully connected and networked setups

In writing down Eq. (2) we have implicitly assumed that the payoff of any given player i may at least potentially depend on the actions \mathbf{s}_{-i} of *all* other players. This corresponds to a fully connected (or well-mixed) population, in which each player interacts with any other player. Specifically, a well mixed model describes a group of N players, who *all* engage in the same N -player public good game by each choosing from the three actions $\{C, D, L\}$. Different setups have been considered for example in [32] where a group of N players is selected randomly from a larger populations of Z individuals.

We also consider the case of a networked arrangement of players. Here agents are placed on the nodes of a graph, and each player is then assumed to repeatedly play iterations of the public goods game with their neighbours on the graph. Evolutionary descriptions of such network public goods games have been considered for example in [33]. The analog of the adaptation rule of Eq. (2) is then given by

$$q_{i,s}(t + \Omega) = (1 - \lambda)q_{i,s}(t) + \frac{1}{\Omega} \sum_{t'=t}^{t'+\Omega-1} u(s, \mathbf{s}_{\partial i}(t')), \quad (3)$$

where ∂i denotes the set of neighbours of i in the network, and where $\mathbf{s}_{\partial i}$ accordingly is the vector of pure actions taken by those neighbours. The learning dynamics on networks will be discussed in more detail below (see Sec. VI).

3. Conversion into mixed strategies

In our model the scores (or propensities) $\{q_{i,s}\}$ determine the mixed strategy profiles of players using the so-called logit (or softmax) rule

$$x_{i,s}(t) = \frac{e^{\beta q_{i,s}(t)}}{\sum_{s'} e^{\beta q_{i,s'}(t)}}, \quad (4)$$

where $x_{i,s}(t)$ is the probability of player i using strategy s in round t . This scheme has similar features to the Fermi update rule commonly used in evolutionary dynamics [13, 34, 35]. Rules of this type have

also been used to fit data from experiments with real players, see e.g. [16, 21, 36].

The parameter β in Eq. (4) is referred to as the intensity of choice (or learning rate) and determines how much importance is given to a difference in pay-offs when calculating the mixed strategy profile. If $\beta = 0$ for example the players choose their actions completely at random and disregard their propensities entirely. If $\beta = \infty$ then they strictly play the pure action with the highest score.

Following [30] it is possible to combine Eq. (4) and Eq. (2) into map

$$x_{i,s}(t + \Omega) = \frac{x_{i,s}(t)^{1-\lambda} \exp\left(\frac{\beta}{\Omega} \sum_{t'=t}^{t'+\Omega-1} u(s, \mathbf{s}_{-i}(t'))\right)}{\sum_{s'} x_{i,s'}(t)^{1-\lambda} \exp\left(\frac{\beta}{\Omega} \sum_{t'=t}^{t'+\Omega-1} u(s', \mathbf{s}_{-i}(t'))\right)}. \quad (5)$$

This equation defines the discrete-time evolution of the mixed strategy vectors of players, no reference to the propensities is required. Eq. (5) here applies to the well-mixed case, but it is straightforward to formulate the corresponding map for a networked system.

4. Deterministic limit

Models in behavioural game theory [20, 21] are typically based on frequent update of propensities and mixed strategy profiles. Adaptation is assumed to occur after each iteration of the game, corresponding to $\Omega = 1$ in our model. Such dynamics are intrinsically stochastic, as moves of all players are drawn at random from the underlying mixed strategy profiles. The main purpose of the present work is to investigate the effects this stochasticity has on the outcome of learning. It is however generally very difficult to obtain analytical results for stochastic dynamical systems such as the one corresponding to the case of $\Omega = 1$. We therefore follow an approach similar to that used in evolutionary dynamics: (i) we first derive the *deterministic* limit dynamics. This corresponds to taking the limit $\Omega \rightarrow \infty$, and is akin to considering the limit of *infinite* populations in evolutionary game theory (resulting in deterministic differential equations of the replicator or replicator-mutator type), (ii) we then perform a systematic perturbative expansion in $\Omega^{-1/2}$ in order to capture stochastic effects to lowest non-trivial order. This is again similar to approaches taken in the context of evolutionary games, where expansions in the inverse population-size are often carried out, see for example [11–15]. While such expansions in the inverse batch size are technically only valid for large, but finite Ω we will show below that they can capture the on-line dynamics, $\Omega = 1$ in good approximation as well.

We will here briefly derive the deterministic limit of learning, and analyze the outcome in the next section.

The expansion in the noise strength is then carried out in Sec. IV.

In the limit $\Omega \rightarrow \infty$, Eq. (5) the mean of the pay-offs from the last Ω rounds approaches the expected payoff for player i when using a given pure strategy, where the expectation is to be taken with respect to the mixed strategies of i 's opponents. Specifically we have

$$\lim_{\Omega \rightarrow \infty} \frac{1}{\Omega} \sum_{t'=t}^{t'+\Omega-1} u(s, \mathbf{s}_{-i}(t')) = \mu_{i,s} \equiv \sum_{\mathbf{s}_{-i}} u(s, \mathbf{s}_{-i}) p_{-i, \mathbf{s}_{-i}}, \quad (6)$$

where $p_{-i, \mathbf{s}_{-i}} = \prod_{j \neq i} x_{j, s_j}$ is the probability of player i facing the pure strategies $\mathbf{s}_{-i} \in \{1, \dots, S\}^{N-1}$ being played by his or her opponents. The sum over \mathbf{s}_{-i} in the above expression extends over all S^{N-1} possible realizations of the vector \mathbf{s}_{-i} . Rescaling time by a suitable factor the update rule of Eq. (5) then becomes

$$x_{i,s}(\tau + 1) = \frac{x_{i,s}(\tau)^{1-\lambda} \exp(\beta \mu_{i,s}(\tau))}{\sum_{s'} x_{i,s'}(\tau)^{1-\lambda} \exp(\beta \mu_{i,s'}(\tau))}, \quad (7)$$

Again, this deterministic map was formulated for the well-mixed case, generalization to network learning is straightforward.

It is also possible to write Eq. (4) and Eq. (2) together in the form of an approximate continuous time dynamic if $\beta \ll 1$. This is known as the Sato-Crutchfield dynamics [30]. These dynamics are similar to the replicator-mutator equations where the learning rate is analogous to selection strength and the memory loss is analogous to the mutation rate [13, 37]. We will restrict our analysis to the map as we feel it better represents the discrete rounds used in experiments.

III. THE FULLY CONNECTED CASE: RESULTS OF THE DETERMINISTIC ANALYSIS

In this section we will first focus on the fully connected (or well-mixed) setup, networked systems are addressed below. Before discussing the outcome of deterministic learning is it useful to briefly comment on the role of initial conditions. If the mixed strategy profiles for the different players are initiated from the same point in the space of all mixed strategies i.e., $x_{i,s}(\tau = 0) = x_{j,s}(\tau = 0)$ for all players i, j and all pure actions s , then the deterministic map, Eq. (7), will preserve this symmetry, i.e., the mixed strategy profiles will remain identical across players, $x_{i,s}(\tau) = x_{j,s}(\tau)$ for later times τ as well. We will refer to such cases as a start from homogeneous initial conditions. Each player will then receive the same




FIG. 1: (Color online) Dynamics of the three-player public good game with homogeneous initial conditions (intensity of choice is $\beta = 0.1$). The black line indicates the trajectory to the fixed point at $\lambda = 10^{-3}$, the dashed red line is the limit cycle at $\lambda = 10^{-4}$ and the solid blue line is the limit cycle around the edge of the simplex at $\lambda = 10^{-7}$.

average payoffs in each round, and the strategy vectors of all players will evolve identically. Thus each player is effectively playing with $N-1$ copies of themselves. For heterogeneous initial conditions, or noisy dynamics, this is not necessarily the case.

We will begin by examining the behaviour of the deterministic system, Eq. (7), for homogeneous initial conditions. The deterministic dynamics is found to exhibit three types of behaviour: (i) for sufficiently large values of the memory-loss parameter λ trajectories approach a stable fixed point, (ii) for intermediate values of λ one finds convergence towards limit cycles in the interior of strategy space, and (iii) at small memory-loss asymptotic trajectories along the edges of the strategy simplex are seen. Examples are shown in Fig. 1.

A more systematic analysis can be found in Fig. 2 where we show phase diagrams for games of $N = 3$, 4 and 5 players. The location and existence of the fixed point is determined numerically and, as such, the boundary between stable spirals and limit cycles reported in the figure are approximate. Performing a linear stability analysis of the fixed points reveals that stable fixed points are always found to be spirals (the eigenvalues of the Jacobian form complex conjugate pairs with negative real parts). The transition to the limit cycle regime occurs through a Hopf bifurcation. The size and degree of stability of the limit cycles is determined by the distance to the Hopf bifurcation in the parameter space. As this distance increases the limit cycles move closer to the border of the simplex, eventually becoming restricted to the border itself (as seen in Fig. 1). The dynamics of the limit cycles follow a periodic pattern, with contributing, defecting and loner each becoming the most prominent strategy




FIG. 2: (Color online) Phase diagram characterizing the outcome of deterministic learning in the N -player public goods game. Lines show the location of the Hopf bifurcation in parameter space, separating stable from unstable fixed points. From bottom to top, lines corresponds to $N = 3$ (black), $N = 4$ (red) and $N = 5$ players (green).




FIG. 3: (Color online) This figure shows the components of the fixed point as a function of the memory loss parameter λ . Here, $N = 3$ and $\beta = 0.01$. The black line shows the concentration of contribution in the mixed strategy profile, red the concentration of defection and green the concentration of the loner strategy.

in turn. The effect of the memory loss parameter on the fixed point of the players is shown in Fig. 3, at low memory loss the players tend to mostly abstain (L), at large memory loss play occurs essentially at random, with all three actions being used with nearly the same frequencies.

We will mostly ignore the case of heterogeneous initial conditions in the following, in all tested cases we find numerically that the system approaches the same fixed point as for homogeneous initial conditions. Since we will be interested in features where the deterministic system is at a fixed point, restricting our investigations to the simpler case of homogeneous initial conditions will not affect our findings.

IV. BATCH-SIZE EXPANSION FOR THE WELL-MIXED CASE

Similar to what is seen in population-based models of evolutionary game theory [13–15, 38, 39], models of epidemics [40–43] and in other biological systems [44–46] the behaviour of the stochastic dynamics of learning can be quite different from that predicted by the deterministic description. In particular a mechanism of coherent amplification can lead to sustained stochastic cycles in the noisy system for values of the model parameters in which deterministic learning converges. Fig. 4 shows an example of this behaviour. In population-based models these oscillations are known as ‘quasi-cycles’, they are an effect seen in finite populations, caused by so-called ‘demographic noise’ [44], noise induced by sampling interacting agents from a finite pool of individuals. Deterministic descriptions of evolving systems are valid only in the limit of *infinite* populations. In the context of two-player learning similar cycles have been studied in [27, 28]. While the mechanism of stochastic amplification is similar to that in population dynamics, the source of noise is different. As described in Sec. IIB the players’ adaptation is intrinsically stochastic if they base their strategy updates on *finite* sets of observations of their opponents’ moves. In Eq. (2) for example we have assumed that a finite number Ω of observations is made between any two adaptation events. Learning becomes deterministic only in the limit $\Omega \rightarrow \infty$. In this sense the batch size Ω is similar to the population size in evolving systems, they both control the strength of noise in the resulting dynamics. This observation is one of the main reasons for introducing the learning algorithm at general batch size, tuning Ω allows us to interpolate between the realistic case of frequent updating ($\Omega = 1$), and the deterministic limit, $\Omega \rightarrow \infty$.

We will now proceed to obtain an analytical characterization of the stochastic quasi-cycles shown in Fig. 4. To this end we carry out a systematic expansion in the noise strength, more precisely in powers $\Omega^{-1/2}$ of the inverse batch size. This is conceptually similar to the van Kampen expansion [47] in powers of the inverse system size of population-based models. We here discuss the key steps of the calculation for the general N -player game, choosing a suitably compact notation. In order to make the mathematical details




FIG. 4: An example of the stochastic oscillations in the mixed strategy profiles of a fixed player for $\beta = 0.1$ and $\lambda = 0.005$. The three oscillating curves show the mixed strategy obtained from one simulation run of the stochastic learning dynamics (L, C and D from top to bottom). Solid lines indicate the deterministic fixed point.

more transparent we also detail the resulting expressions for the explicit case $N = 3$ in the Appendix. The first step of the expansion is to rewrite the last term on the RHS of Eq. (2). This term represents the average payoff per round to player i in Ω iterates of the game. Separating stochastic fluctuations from the expected *mean* payoff this term can be written as follows

$$\frac{1}{\Omega} \sum_{t'=t}^{t'+\Omega-1} u(s, \mathbf{s}_{-i}(t')) = \mu_{i,s} + \tilde{\xi}_{i,s} \quad (8)$$

where $\mu_{i,s} = \sum_{\mathbf{s}_{-i}} u(s, \mathbf{s}_{-i}) p_{-i, \mathbf{s}_{-i}}$ is the expected payoff (per round) for player i given his opponents’ mixed strategy profiles, and provided player i plays pure strategy s . The second contribution, $\tilde{\xi}_{i,s}$, represents the fluctuations about the deterministic trajectory. The correlations between these noise terms are found by rearranging Eq. (8), isolating $\tilde{\xi}_{i,s}$ and averaging over all possible actions of all players. Correlations between noise variables associated with strategies belonging to the same player are given by

$$\langle \tilde{\xi}_{i,s}(\tau) \tilde{\xi}_{i,s'}(\tau') \rangle = \frac{\delta_{\tau\tau'}}{\Omega} \sum_{\mathbf{s}_{-i}} (u(s, \mathbf{s}_{-i}) - \mu_{i,s}) \times (u(s', \mathbf{s}_{-i}) - \mu_{i,s'}) p_{-i, \mathbf{s}_{-i}}. \quad (9)$$

Unlike in the two-player case [27, 28], there are now correlations also between the noise variables associated with different players. This is because the payoffs of any two players both depend on the actions of the remaining players. These correlations take the

form

$$\begin{aligned} \langle \tilde{\xi}_{i,s}(\tau) \tilde{\xi}_{j,s'}(\tau') \rangle &= \quad (10) \\ &= \frac{\delta_{\tau\tau'}}{\Omega} \sum_{\mathbf{s}} (u(\mathbf{s}, \mathbf{s}_{-i}) - \mu_{i,s}) \\ &\quad \times (u(\mathbf{s}', \mathbf{s}_{-j}) - \mu_{j,s'}) p_{\mathbf{s}}, \end{aligned}$$

where $\mathbf{s} = (s_1, \dots, s_N) \in \{1, \dots, S\}^N$ describes the vector of actions taken by the whole group of players, and where $p_{\mathbf{s}} = \prod_k x_{k,s_k}$. It is worth pointing out that Eq. (9) can be seen as a special case of Eq. (10), if $i = j$ in the latter relation then the sum over actions of player i can be factored out, and one obtains Eq. (9).

Given that the payoffs, Eq. (8), are now random variables, the same will be true for the components $x_{i,s}$ of the players' mixed strategy profiles. In order to separate fluctuations from the deterministic dynamics we use the ansatz

$$x_{i,s}(\tau) = \bar{x}_{i,s}(\tau) + \tilde{x}_{i,s}(\tau), \quad (11)$$

where the first term on the RHS is the deterministic trajectory, and where the second term is of order $\Omega^{-1/2}$ and describes fluctuations about the deterministic solution. Using this ansatz, together with Eq. (8), and after the re-scaling of time discussed above, we can rewrite Eq. (5) as

$$\begin{aligned} x_{i,s}(\tau + 1) &= \\ &= \frac{(\bar{x}_{i,s} + \tilde{x}_{i,s})^{1-\lambda} e^{\beta(\bar{\mu}_{i,s} + \tilde{\mu}_{i,s} + \tilde{\xi}_{i,s})}}{\sum_{s'} (\bar{x}_{i,s'} + \tilde{x}_{i,s'})^{1-\lambda} e^{\beta(\bar{\mu}_{i,s'} + \tilde{\mu}_{i,s'} + \tilde{\xi}_{i,s'})}}, \quad (12) \end{aligned}$$

where

$$\bar{\mu}_{i,s} = \sum_{\mathbf{s}_{-i}} u(\mathbf{s}, \mathbf{s}_{-i}) \bar{p}_{-i, \mathbf{s}_{-i}}, \quad (13)$$

$$\tilde{\mu}_{i,s} = \sum_{\mathbf{s}_{-i}} u(\mathbf{s}, \mathbf{s}_{-i}) \tilde{p}_{-i, \mathbf{s}_{-i}}, \quad (14)$$

and where all quantities on the RHS of Eq. (12) are evaluated at time τ . We have here introduced $\bar{p}_{-i, \mathbf{s}_{-i}} = \prod_{k \neq i} \bar{x}_{k, s_k}$ and

$$\begin{aligned} \tilde{p}_{-i, \mathbf{s}_{-i}} &= \sum_{k \neq i} \frac{\partial \bar{p}_{-i, \mathbf{s}_{-i}}}{\partial \bar{x}_{k, s_k}} \tilde{x}_{k, s_k} \\ &= \sum_{k \neq i} \left(\prod_{j \notin \{i, k\}} \bar{x}_{j, s_j} \right) \tilde{x}_{k, s_k}. \quad (15) \end{aligned}$$

We now expand Eq. (12) in powers of $\Omega^{-1/2}$, taking into account that $\tilde{x}_{i,s}$ and $\tilde{\xi}_{i,s}$ scale as $\Omega^{-1/2} \forall i, s$. Naturally, from the leading-order terms in this one expansion recovers the deterministic dynamics

$$\bar{x}_{i,s}(\tau + 1) = \frac{\bar{x}_{i,s}^{1-\lambda} e^{\beta \bar{\mu}_{i,s}}}{\sum_{s'} \bar{x}_{i,s'}^{1-\lambda} e^{\beta \bar{\mu}_{i,s'}}}. \quad (16)$$

At next-leading order in the expansion one finds

$$\begin{aligned} \tilde{x}_{i,s}(\tau + 1) &= \sum_k \sum_{s'} \left(\frac{\partial g_{i,s}}{\partial \bar{x}_{k, s'}} \right) \Big|_{\tilde{\mathbf{x}} = \tilde{\boldsymbol{\xi}} = 0} \tilde{x}_{k, s'} \\ &\quad + \sum_{s'} \left(\frac{\partial g_{i,s}}{\partial \tilde{\xi}_{i, s'}} \right) \Big|_{\tilde{\mathbf{x}} = \tilde{\boldsymbol{\xi}} = 0} \tilde{\xi}_{i, s'}, \quad (17) \end{aligned}$$

where $g_{i,s}$ is a short-hand for the RHS of Eq. (12). The notation $\tilde{\mathbf{x}} = \tilde{\boldsymbol{\xi}} = 0$ indicates that all variables $\tilde{x}_{i,s}$ and $\tilde{\xi}_{i,s}$ ($i = 1, \dots, N, s = 1, \dots, S$) are to be taken to zero when evaluating the derivatives. The derivative of $g_{i,s}$ with respect to $\tilde{\xi}_{k, s'}$ vanishes for $k \neq i$. The resulting coefficients in Eq. (17) multiplying the variables $\{\tilde{x}_{k, s'}\}$ are the entries of the Jacobian of Eq. (12), to be evaluated at $\tilde{x}_{i,s} = \tilde{\xi}_{i,s} = 0$. We now further simplify these terms and find

$$\begin{aligned} &\sum_{s'} \left(\frac{\partial g_{i,s}}{\partial \tilde{\xi}_{i, s'}} \right) \Big|_{\tilde{\mathbf{x}} = \tilde{\boldsymbol{\xi}} = 0} \tilde{\xi}_{i, s'} \\ &= \beta \bar{x}_{i,s} \left(\tilde{\xi}_{i,s} - \sum_{s'} \bar{x}_{i, s'} \tilde{\xi}_{i, s'} \right), \quad (18) \end{aligned}$$

similar to what has been reported in [28]. The resulting dynamics for the fluctuations $\tilde{x}_{i,s}$ about the deterministic trajectory can be compactly written as a set of Langevin equations for the $N \times S$ variables, $\tilde{\mathbf{x}}_{i,s}$,

$$\boldsymbol{\delta}(\tau + 1) = \boldsymbol{\delta}(\tau) + \mathcal{J} \boldsymbol{\delta}(\tau) + \boldsymbol{\zeta}. \quad (19)$$

We have here introduced $\boldsymbol{\delta} = (\tilde{\mathbf{x}}_1, \dots, \tilde{\mathbf{x}}_N)$, where $\tilde{\mathbf{x}}_i = (\tilde{x}_{i,1}, \dots, \tilde{x}_{i,S})$. We have also written $\boldsymbol{\zeta} = (\tilde{\boldsymbol{\gamma}}_1, \dots, \tilde{\boldsymbol{\gamma}}_N)$ with $\tilde{\boldsymbol{\gamma}}_i = (\tilde{\gamma}_{i,1}, \dots, \tilde{\gamma}_{i,S})$, and where

$$\tilde{\gamma}_{i,s} = \beta \bar{x}_{i,s} \left(\tilde{\xi}_{i,s} - \sum_{s'} \bar{x}_{i, s'} \tilde{\xi}_{i, s'} \right). \quad (20)$$

The $(NS) \times (NS)$ -matrix \mathcal{J} is the Jacobian of Eq. (12), evaluated at the fixed point of the deterministic dynamics. The correlations between the components of $\boldsymbol{\gamma}$, can be expressed in terms of those of the $\{\tilde{\xi}_{i,s}\}$, specifically we have

$$\begin{aligned} &\langle \tilde{\gamma}_{i,s} \tilde{\gamma}_{j, s'} \rangle \\ &= \beta^2 \bar{x}_{i,s} \bar{x}_{j, s'} \left(\langle \tilde{\xi}_{i,s} \tilde{\xi}_{j, s'} \rangle - \bar{x}_{i,s} \sum_{s''} \bar{x}_{j, s''} \langle \tilde{\xi}_{i,s} \tilde{\xi}_{j, s''} \rangle \right. \\ &\quad \left. - \bar{x}_{j, s'} \sum_{s''} \bar{x}_{i, s''} \langle \tilde{\xi}_{i, s''} \tilde{\xi}_{j, s'} \rangle \right. \\ &\quad \left. + \sum_{s'' s'''} \bar{x}_{i, s''} \bar{x}_{j, s'''} \langle \tilde{\xi}_{i, s''} \tilde{\xi}_{j, s'''} \rangle \right). \quad (21) \end{aligned}$$

Fourier transforming Eq. (19) gives

$$((e^{i\omega} - 1)\mathcal{I} - \mathcal{J}) \hat{\boldsymbol{\delta}} = \hat{\boldsymbol{\zeta}}, \quad (22)$$

where $\widehat{\delta}$ is the Fourier transform of δ and $\widehat{\zeta}$ that of ζ . Defining $\mathcal{M}(\omega) = ((e^{i\omega} - 1)\mathcal{I} - \mathcal{J})$, with \mathcal{I} the identity matrix, and using the notation \mathcal{B} for the matrix of correlations of γ we find

$$\langle \widehat{\delta}(\omega)\widehat{\delta}(\omega')^T \rangle = \mathcal{M}(\omega)^{-1}\mathcal{B}(\mathcal{M}^\dagger(\omega))^{-1}\delta(\omega + \omega'). \quad (23)$$

We will refer to the matrix $P(\omega) = \mathcal{M}(\omega)^{-1}\mathcal{B}(\mathcal{M}^\dagger(\omega))^{-1}$ as the set of power spectra of fluctuations in the following. This is the final result of the batch-size expansion. Although formulated in Fourier space the expressions in Eq. (23) contain full information about the autocorrelation functions and crosscorrelation functions of the variables $\tilde{x}_{i,s}$. They hence characterize the fluctuations about the deterministic fixed point (during the course of the calculation we have assumed that the deterministic trajectory approaches a stable fixed point, generalization to limit cycles is possible, but tedious, see [48, 49]). It is straightforward to carry out the required matrix inversions, and to evaluate the RHS of Eq. (23) numerically. We will here mostly focus on quasi-cycles, it is therefore convenient to operate in Fourier space. In particular we will compare the power spectra of quasi-cycles, as obtained from Eq. (23) against numerical simulations in the next section. Given that we have carried out an expansion in powers of $\Omega^{-1/2}$ and that we have only retained the next-to-leading order we expect our results to be valid for large, but finite values of the batch size Ω .

V. CHARACTERIZATION OF QUASI-CYCLES IN THE FULLY CONNECTED CASE

We now proceed to test the theoretical results obtained in the previous section against numerical simulations of the multi-player learning process. Taking a three-player game as an example, we find good agreement between simulation and theory, see Fig. 5. This figure also demonstrates the finite-size effects at different batch sizes. The case $\Omega = 1$ is the so-called ‘on-line’ learning limit where players update their scores after every round [31]. Here the agreement with the theory is only approximate, which is no surprise given that higher-order terms in the expansion have been discarded. Nevertheless the analytical theory is able to predict the dominating frequency of the quasi-cycles to a good approximation. At lower noise strengths, the theory becomes increasingly more accurate. The batch size required to obtain a pre-set degree of agreement between simulations theory will generally increase the closer a given choice of parameters is to the Hopf bifurcation, similar to what has been seen in [48]. Although the mixed strategy profiles of individual players will evolve differently in any




FIG. 5: This figure shows the power spectrum of the contributory strategy in a three player game at $\lambda = 0.005$, $\beta = 0.1$. The different symbols shows simulations at $\Omega = 1, 10, 100$, averaged over in excess of 100 independent runs. The solid line is the theoretical spectrum from Eq. (23).

one run of the learning dynamics, there are statistically equivalent when an average over sufficiently many runs is taken. Power spectra such as those shown in Fig. 5 in particular are identical for different players, due to the symmetry of the well-mixed setup with respect to permutations of players. As we will see below, this is no longer the case in networked systems when different players have different connectivities.

Fig. 6 shows a comparison between the spectrum for a three, four and five player game. Higher numbers of players lead to smaller amplitudes of oscillations, greater distance between the two peaks and an increase in the ratio of the heights of the peaks. The difference in amplitudes of fluctuations may here be due to the fact that there are more possible payoffs in between the maximum and the minimum payoff as the number of players is increased. Smaller jumps between payoffs may therefore reduce the effects of stochasticity, keeping all other model parameters fixed.

Experimental situations often observe oscillations with much higher frequencies than the examples given here so far. An experiment by Milinski et al [50], although constructed differently to the situation this model mimics, shows oscillations of frequency $\omega \approx 2$. Fig. 7 shows an example with a high learning rate and rapid memory loss that reproduces such frequencies. We stress though that no claim is made that such memory-loss rates or intensity of choice are necessarily realistic. Experimental data for other games is available in [21]. Instead the purpose of Fig. 7 is mainly to show that the characteristic frequency of




FIG. 6: This figure shows the power spectrum for a player’s contributory strategy at $\lambda = 0.005$, $\beta = 0.1$ for different numbers of players. The different colours indicate different numbers of players as described in the legend. Symbols are from simulations run at $\Omega = 100$, averaged over in excess of 100 independent runs. Solid lines are from the theory.




FIG. 7: This figure shows a power spectrum for the contributing strategy which exhibits frequencies similar to those seen in experiments. The parameters for this five player game are $\Omega = 100$, $\beta = 2$, $\lambda = 1$.

the observed quasi-cycles can vary over a wide range.

VI. PUBLIC GOOD GAMES ON NETWORKS

In real social systems populations are not well-mixed, instead people will only interact with certain others. It is no surprise that the study of games

on networks, and the properties of networks in general, has become so popular, see e.g. [51–54]. Analytical progress in studying games on networks has been made [55–57], these methods typically assume meta-population models, with a large group of player placed at each node of the underlying network. We here take a different approach, and consider single individuals, connected via a static network. Within the game learning approach we can then apply the expansion outlined in Sec. IV, treating the strategies on each node as separate variables. The network structure is reflected in the calculation of the average payoffs and in the correlations of the noise. This approach produces a dynamical system whose dimensionality increases linearly with the number of nodes, making the study of large networks difficult although technically possible. However, up to recently, experimental situations are usually also restricted to small numbers of players [21, 36, 50], with some notable exceptions, e.g. [58].

A. Batch-size expansion on a network

In the networked model we assume that each player at each iteration of the game chooses a move (C , D or L), according to his or her mixed strategy profile. Payoffs are then determined by Eq. (1), where the group of opponents faced by player i is the set of his neighbors ∂i . The calculation then proceeds along similar steps to that of the well-mixed case. If there are N nodes (players) in the network, and if each one of them chooses between S pure strategies, then the resulting dynamics has $N(S - 1)$ degrees of freedom, as before. In the limit of infinite batch size this leads to an $N(S - 1)$ dimensional deterministic dynamical system in discrete time. As one crucial difference to the well mixed case the permutation symmetry between players is no longer present, even for homogeneous initial conditions. For general networks, different players will typically have different degrees, i.e., they face different numbers of opponents. Hence the mixed strategies played at deterministic fixed points will generally vary across the set of players (as before we here focus on the case of parameters in which deterministic learning converges). Carrying out the expansion in the inverse batch size one also finds that the network structure affects the noise correlators and the structure of the Jacobian of the deterministic dynamics. It is worth noting that we do not necessarily expect to see a breaking of the permutation symmetry for regular networks, i.e., when all players have the same degree k . For homogeneous initial conditions the mixed strategies of all players will then evolve identically under the deterministic dynamics. For heterogenous initial conditions and/or noisy dynamics this may not necessarily be the case though, as shown for the two-player case in [29].




FIG. 8: The network chosen for analysis. We will look at two players: the ‘hub’ player (left) and the ‘ring’ player (right).

In the networked setting the expected payoff for player i is given by

$$\mu_{i,s} = \sum_{\mathbf{s}_{\partial_i}} p_{i,\mathbf{s}_{\partial_i}} u(s, \mathbf{s}_{\partial_i}) \quad (24)$$

where $\mathbf{s}_{\partial_i} \in \{1, \dots, S\}^{|\partial_i|}$ are the actions of those players connected to player i (in the fully connected case this is simply the set of all other players), and where $p_{i,\mathbf{s}_{\partial_i}}$ is the probability for their joint action \mathbf{s}_{∂_i} to occur. The analog of Eqs. (9) and (10) can be written as

$$\begin{aligned} & \langle \tilde{\xi}_{i,s} \tilde{\xi}_{j,s'} \rangle \\ &= \sum_{\mathbf{s}_{\partial_i \cup \partial_j}} p_{\partial_i \cup \partial_j, \mathbf{s}_{\partial_i \cup \partial_j}} (u(s, \mathbf{s}_{\partial_i}) - \mu_{i,s}) \\ & \quad \times (u(s', \mathbf{s}_{\partial_j}) - \mu_{j,s'}), \end{aligned} \quad (25)$$

where $\partial_i \cup \partial_j$ is the union of the sets of neighbours of i and of j . The remaining steps are as in Sec. IV, with the appropriate modifications to take account of the changed payoff structure in the networked arrangement, Eq. (24).

B. Effect of degree

We will restrict our discussion to the network shown in Fig. 8, and we investigate the dynamics both for the central player, and for those on the outer spokes. The choice of this particular shape of network is to a certain extent arbitrary, it is important to keep in mind though that our theoretical approach applies to general graphs. The choice made here is hence mainly for purposes of illustration.

1. Deterministic features

Let us first examine the effect of degree on the deterministic behaviour of the system, we focus on




FIG. 9: This figure shows the concentrations of the three strategies for the central player in networks of the type shown in Fig. 8, as a function of the player’s degree k . The parameter values are $\beta = 0.1$, $\lambda = 0.005$. Results are from the deterministic map, lines connecting the symbols are guides to the eye.

the fixed point regime, i.e., sufficiently quick memory loss. As discussed above the mixed strategies played by different players at deterministic fixed points will generally depend on their degree and the degree of their neighbours. To illustrate this we have considered networks of the type shown in Fig. 8, varying the degree k of the central player. The resulting mixed strategy profile at convergence of the deterministic dynamics is shown in Fig. 9 for the central player, and in Fig. 10 for a player on the outer ring of the network. As the degree of the central players increases, so does the probability that he or she defects. For the outer player it is mostly the probability of abstaining (i.e., to play ‘loner’) which increases, as the connectivity of the central player is increased. This is what one would expect given that their central neighbour is becoming increasingly likely to defect.

2. Stochastic features

As a final part of the analysis we study the effect of connectivity on oscillations induced by intrinsic noise in the learning process. In Fig. 11 we depict the power spectra of fluctuations of the probability with which the central player in our sample network plays strategy C . An analogous plot for a player on the outer ring of the network is shown in Fig. 12. In both cases a good match between theory and simulations is found, even at relatively moderate batch sizes of $\Omega = 10$. The theoretical approach based on a batch-size expansion is therefore successful in predicting the characteristic frequency of the observed quasi-cycles.

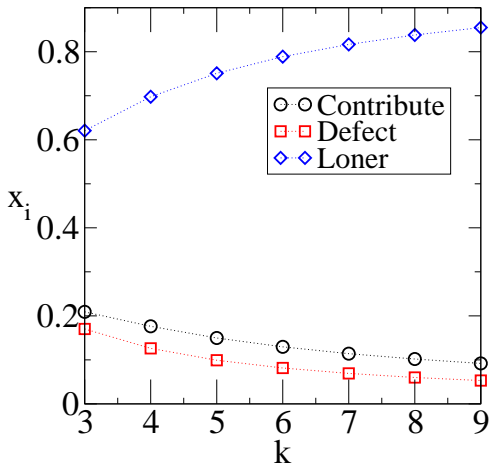


FIG. 10: This figure illustrates how the mixed strategy profile of a player on the outer ring of networks of the type shown in Fig. 8 changes, as the degree k of the central player is varied. The parameter values are $\beta = 0.1$, $\lambda = 0.005$. Results are from the deterministic map, lines connecting the symbols are guides to the eye.

Comparison of Figs. 11 and 12 with Fig. 6 reveals an intriguing behaviour: (i) increasing the overall connectivity in the ‘hub-and-ring’ network seems to reduce fluctuations of mixed strategies of players on the outer ring of the network. This is similar to the effect we have seen in the well-mixed case (see Fig. 6), where fluctuations are suppressed as the number of players in the group (or equivalently the total number of links in the graph) increases. In-line with this behaviour we find that the player on the outer ring of the network is more likely to abstain as the degree of the central node is increased, similar behaviour is found in the well-mixed case as the overall number of players is increased (not shown); (ii) The central player in the network however shows a different type of behaviour as their degree is increased. He or she is less likely to abstain (see Fig. 9), and fluctuations of the central player’s mixed strategy increase with increasing degree (Fig. 11).

Care needs to be taken though in making direct comparisons between the network shown in Fig. 8 and the well-mixed case. Increasing the degree k of the central player in the network does not imply an increased connectivity of the players on the outer ring, and the situation is therefore different from that of a regular random network with uniform degree across all nodes. The juxtaposition between the behaviour of the central agent in the network and that of a player in the fully connected case indeed suggests that it is not only the degree of a player itself that determines his or her mixed strategy and the magnitude of fluctuations about it, but that on the contrary the degrees of the players he or she is connected to play




FIG. 11: Shown are the power spectra for the contributing strategy of a player in the centre of a ring-and-hub network. Different colours represent different degrees as shown in the legend. The parameter values are $\beta = 0.1$, $\lambda = 0.005$. Simulations are from 500 runs at $\Omega = 10$.




FIG. 12: The power spectra for the contributing strategy on the edge of a ring-and-hub network. Different colours represent different degrees of the ‘hub’ neighbour. The parameter values are $\beta = 0.1$, $\lambda = 0.005$. Simulations are from 500 runs at $\Omega = 10$.

an important role as well.

VII. CONCLUSION AND OUTLOOK

In summary we have investigated the deterministic and stochastic learning dynamics in multi-player games. Our approach builds on models from behavioural game theory [20, 21], and is similar to the Sato-Crutchfield formulation of game learning

[18, 19, 30]. The approach by Sato et al is however based on deterministic ordinary differential equations, and ignores effects of noise on the outcome of learning. Such stochastic effects have recently been studied for two-player games in [27, 28], the main contribution of the current work is to extend these approaches to the case of multi-player games, and to games on networks. We have here successfully carried out expansions in the inverse noise strength. On the deterministic level Sato-Crutchfield learning, and its extensions to discrete time, exhibit features similar to those of the replicator, or replicator-mutator equations. On a stochastic level we confirm the presence of quasi-cycles in a wide range of model parameters, in particular stochastic learning can exhibit persistent oscillations in parameter regimes where deterministic learning converges. Based on the expansion in the inverse noise strength we are able to make analytical predictions on the existence or otherwise of quasi-cycles as a function of parameters of the learning dynamics and of the underlying game or network structure.

We believe that taking a modelling approach based on adaptation dynamics intrinsic to individual players, inspired by models of behavioural game theory, has a number of advantages over population-based approaches repying on birth-death processes. Learning models allow the incorporation of memory loss which has been shown experimentally to be an important factor. By contrast, since the replicator approach models the evolution of strategies rather than the behaviour of players it is often not straightforward to include such psychological or cognitive effects directly. On the other hand, both approaches are not mutually exclusive, and there are indeed a number of similarities between them. Memory loss and mutation for example play similar roles in the respective modelling frameworks. Future research may therefore focus on drawing further analogies, and to use the existing knowledge on evolutionary models in finite populations to elucidate the mathematical structures and outcomes of stochastic learning models in further detail.

Acknowledgments

AJB acknowledges an EPSRC studentship. TG is supported by an RCUK Fellowship (RCUK reference EP/E500048/1), and by the Engineering and Physical Sciences Research Council (EPSRC, references EP/I005765/1 and EP/I019200/1). The authors would like to thank Yuzuru Sato for useful discussions.

Appendix A: Batch-size expansion for a well-mixed, three-player game

As an example of the application of the equations derived in Sec. IV, we detail the steps involved in the calculation for the well-mixed case with $N = 3$ and $S = 3$. The average payoff for strategy s used by player i then becomes

$$\mu_{i,s} = \sum_{s' s''} a_{s s' s''} x_{j,s'} x_{k,s''}, \quad (\text{A1})$$

where j and k indicate the other two players and $a_{s s' s''}$ is an 3-dimensional tensor containing the payoff for strategy s when the other players use strategies s' and s'' . Average payoffs for the other players are found through cyclic permutation of i, j and k . Evaluating Eq. (9) we find

$$\begin{aligned} & \langle \tilde{\xi}_{i,s}(\tau) \tilde{\xi}_{i,s'}(\tau') \rangle \\ &= \frac{\delta_{\tau\tau'}}{\Omega} \sum_{s'', s'''} x_{i,s''} x_{i,s'''} (a_{s s'' s'''} - \mu_{i,s}) \\ & \quad \times (a_{s' s'' s'''} - \mu_{j,s'}). \end{aligned} \quad (\text{A2})$$

From Eq. (10) we find

$$\begin{aligned} \langle \tilde{\xi}_{i,s}(\tau) \tilde{\xi}_{j,s'}(\tau') \rangle &= \frac{\delta_{\tau\tau'}}{\Omega} \sum_{s'', s''', s''''} x_{i,s''} x_{j,s'''} x_{k,s''''} \\ & \quad \times (a_{s s'' s'''} - \mu_{i,s}) (a_{s' s'' s''''} - \mu_{j,s'}) \end{aligned} \quad (\text{A3})$$

Eq. (10) differs from the two player-case studied in [27, 28] where the correlations of the noise associated with strategies belonging to different players is zero. In the three-player case, this noise is correlated through the actions of the third player. If we introduce the ansatz $x_{i,s} = \bar{x}_{i,s} + \tilde{x}_{i,s}$ as in Sec. IV, the average payoffs given by Eq. (A1) can be separated into a term of order Ω^0 and a term of order $\Omega^{-1/2}$

$$\bar{\mu}_{i,s} = \sum_{s' s''} a_{s s' s''} \bar{x}_{j,s'} \bar{x}_{k,s''} \quad (\text{A4})$$

$$\tilde{\mu}_{i,s} = \sum_{s' s''} a_{s s' s''} (\bar{x}_{j,s'} \tilde{x}_{k,s''} + \tilde{x}_{j,s'} \bar{x}_{k,s''}). \quad (\text{A5})$$

The expansion of Eq. (12) for the three player case then gives the following for terms of order $\Omega^{-1/2}$.

$$\begin{aligned} \tilde{x}_{i,s}(\tau+1) - \tilde{x}_{i,s}(\tau) &= \sum_{s'} \left(\frac{\partial g_{i,s}}{\partial \tilde{x}_{i,s'}} \right) \Big|_{\tilde{x}=\tilde{\xi}=0} \tilde{x}_{i,s'} \quad (\text{A6}) \\ &+ \sum_{s'} \left(\frac{\partial g_{i,s}}{\partial \tilde{x}_{j,s'}} \right) \Big|_{\tilde{x}=\tilde{\xi}=0} \tilde{x}_{j,s'} + \sum_{s'} \left(\frac{\partial g_{i,s}}{\partial \tilde{x}_{k,s'}} \right) \Big|_{\tilde{x}=\tilde{\xi}=0} \tilde{x}_{k,s'} \\ &+ \sum_{s'} \left(\frac{\partial g_{i,s}}{\partial \tilde{\xi}_{i,s'}} \right) \Big|_{\tilde{x}=\tilde{\xi}=0} \tilde{\xi}_{i,s'}. \end{aligned}$$

For player i , the coefficients multiplying $\tilde{\xi}_{j,s'}$ and $\tilde{\xi}_{k,s'}$ vanish. The remainder of the calculation then pro-

ceeds exactly as in Sec. IV.

-
- [1] J. M. Smith, *Evolution and the Theory of Games* (Cambridge University Press, Cambridge, UK, 1982).
- [2] M. A. Nowak, *Evolutionary Dynamics: Exploring the Equations of Life* (Harvard University Press, Cambridge, MA, 2006).
- [3] M. A. Nowak, N. L. Komarova, and P. Niyogi, *Science* **291**, 114 (2001).
- [4] J. von Neumann and O. Morgenstern, *Theory of games and economic behaviour* (Princeton University Press, Princeton, NJ, 1944).
- [5] J. F. Nash, *Proc. Natl. Acad. Sci. U.S.A.* **36**, 48 (1950).
- [6] J. Nash, *Annals of Mathematics* **54**, 286 (1951).
- [7] H. Gintis, *Game Theory Evolving (first edition)* (Princeton University Press, Princeton NJ, 2000).
- [8] J. Hofbauer and K. Sigmund, *Evolutionary Games and Population Dynamics* (Cambridge University Press, Cambridge, UK, 1998).
- [9] T. Antal and I. Scheuring, *B. Math. Biol.* **68**, 1923 (2006).
- [10] P. M. Altrock and A. Traulsen, *New J. Phys.* **11**, 013012 (2009).
- [11] A. Traulsen, J. C. Claussen, and C. Hauert, *Phys. Rev. Lett.* **95**, 238701 (2005).
- [12] J. C. Claussen, *Eur. Phys. J. B* **60**, 391 (2007).
- [13] A. J. Bladon, T. Galla, and A. J. McKane, *Phys. Rev. E* **81**, 066122 (2010).
- [14] M. Mobilia, *J. Theor. Biol.* **264**, 1 (2010).
- [15] T. Reichenbach, M. Mobilia, and E. Frey, *Phys. Rev. E* **74**, 051907 (2006).
- [16] A. Traulsen, D. Semmann, R. D. Sommerfeld, H.-J. Krambeck, and M. Milinski, *Proc. Natl. Acad. Sci. U. S. A.* **107**, 2962 (2010).
- [17] C. Hauert, S. D. Monte, J. Hofbauer, and K. Sigmund, *J. Theor. Biol.* **218**, 187 (2002).
- [18] Y. Sato and J. P. Crutchfield, *Phys. Rev. E* **67**, 015206 (2003).
- [19] Y. Sato, E. Akiyama, and J. D. Farmer, *Proc. Natl. Acad. Sci. U.S.A.* **99**, 4748 (2002).
- [20] C. Camerer and T. H. Ho, *Econometrica* **64**, 827 (1999).
- [21] T. H. Ho, C. F. Camerer, and J.-K. Chong, *J. Econ. Theory* **133**, 177 (2007).
- [22] M. W. Macy and A. Flache, *Proc. Natl. Acad. Sci. U.S.A.* **99**, 7229 (2002).
- [23] P. H. Young, *Strategic Learning and Its Limits* (Oxford University Press, Oxford, UK, 2004).
- [24] D. Fudenberg and D. K. Levine, *The Theory of Learning Games* (MIT Press, Cambridge, MA, 1998).
- [25] G. W. Brown, *Activity analysis of production and allocation* **13**, 374 (1951).
- [26] D. Fudenberg and D. M. Kreps, *Games and Economic Behavior* **5**, 320 (1993).
- [27] T. Galla, *Phys. Rev. Lett.* **103**, 198702 (2009).
- [28] T. Galla, pre-print (2011), URL <http://arxiv.org/abs/1101.4378>.
- [29] J. Realpe-Gomez, B. Szczesny, L. Dall'Asta, and T. Galla, pre-print (2011), URL <http://arxiv.org/abs/1102.0876>.
- [30] Y. Sato, E. Akiyama, and J. P. Crutchfield, *Physica D* **210**, 21 (2005).
- [31] D. Saad, *On-Line Learning in Neural Networks* (Cambridge University Press, Cambridge, UK, 1999).
- [32] J. M. Pacheco, F. C. Santos, M. O. Souza, and B. Skyrms, *Proc. R. Soc. Lond. B* **276**, 315 (2009).
- [33] F. C. Santos, M. D. Santos, and J. M. Pacheco, *Nature* **454**, 213 (2008).
- [34] A. Traulsen, M. A. Nowak, and J. M. Pacheco, *J. Theor. Biol.* **244**, 349 (2007).
- [35] J. M. Pacheco, A. Traulsen, and M. A. Nowak, *Phys. Rev. Lett.* **97**, 258103 (2006).
- [36] C. Camerer, *Behavioral game theory: experiments in strategic interaction* (Princeton University Press, Princeton, NJ, 2003).
- [37] M. A. Nowak, N. L. Komarova, and P. Niyogi, *Science* **291**, 114 (2001).
- [38] A. Traulsen and C. Hauert, *Stochastic evolutionary game dynamics* (Wiley-VCH Verlag GmbH & Co. KGaA, Weinheim, Germany, 2010), chap. 2.
- [39] A. Traulsen, C. Hauert, H. De Silva, M. A. Nowak, and K. Sigmund, *Proc. Natl. Acad. Sci. U.S.A.* **106**, 709 (2009).
- [40] A. Black, A. McKane, A. Nunes, and A. Parisi, *Phys. Rev. E* **80**, 21922 (2009), ISSN 1539-3755.
- [41] D. Alonso, A. J. McKane, and M. Pascual, *J. R. Soc. Interface* **4**, 575 (2007).
- [42] T. Coulson, P. Rohani, and M. Pascual, *Trends Ecol. Evol.* **19**, 359 (2004).
- [43] M. Simões, M. M. T. Da Gama, and A. Nunes, *J. R. Soc. Interface* **5**, 29 (2007).
- [44] A. J. McKane and T. J. Newman, *Phys. Rev. Lett.* **94**, 218102 (2005).
- [45] M. Scott, B. Ingalls, and M. Kaern, *Chaos* **16**, 026107 (2006).
- [46] D. Gonze, J. Halloy, and P. Gaspard, *J. Chem. Phys* **116**, 10997 (2002).
- [47] N. G. van Kampen, *Stochastic Processes in Physics and Chemistry* (Elsevier Science, North Holland, 1997).
- [48] R. P. Boland, T. Galla, and A. J. McKane, *J. Stat. Mech.* P09001 (2008).
- [49] R. P. Boland, T. Galla, and A. J. McKane, *Phys. Rev. E* **79**, 051131 (2009).
- [50] D. Semmann, H. Krambeck, and M. Milinski, *Nature* **425**, 390 (2003).
- [51] M. Newman, *Networks: an introduction* (Oxford University Press, 2010).
- [52] M. Newman, A.-L. Barabási, and D. J. Watts, *The structure and dynamics of networks* (Princeton University Press, 2006).
- [53] E. Lieberman, C. Hauert, and M. A. Nowak, *Nature* **433**, 312 (2005).
- [54] G. Szabó and G. Fáth, *Physics Reports* **446**, 97 (2007).
- [55] C. E. Tarnita, H. Ohtsuki, T. Antal, F. Fu, and M. A. Nowak, *J. Theor. Biol.* **259**, 570 (2009).

- [56] F. Fu, L. Wang, M. Nowak, and C. Hauert, Phys. Rev. E **79**, 046707 (2009).
- [57] T. Gross, C. D. D'Lima, and B. Blasius, Phys. Rev. Lett. **96**, 208701 (2005).
- [58] J. Grujic, C. Fosco, L. Araujo, J. A. Cuesta, and A. Sánchez, PLoS ONE **5**, e13749 (2010).