Citation for published version:

Yang, Q, Rogers, T & Dawes, JHP 2017, 'Demographic noise slows down cycles of dominance', *Journal of Theoretical Biology*, vol. 432, pp. 157-168. <https://doi.org/10.1016/j.jtbi.2017.07.025>

DOI:

[10.1016/j.jtbi.2017.07.025](https://doi.org/10.1016/j.jtbi.2017.07.025)

Publication date:

2017

Document Version

Peer reviewed version

[Link to publication](#)

Publisher Rights

CC BY-NC-ND

University of Bath

Alternative formats

If you require this document in an alternative format, please contact:
openaccess@bath.ac.uk

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Demographic noise slows down cycles of dominance

Qian Yang^a, Tim Rogers^{a,*}, Jonathan H.P. Dawes^a

^a*Centre for Networks and Collective Behaviour, Department of Mathematical Sciences,
University of Bath, Bath BA2 7AY, UK*

Abstract

We study the phenomenon of cyclic dominance in the paradigmatic Rock–Paper–Scissors model, as occurring in both stochastic individual-based models of finite populations and in the deterministic replicator equations. The mean-field replicator equations are valid in the limit of large populations and, in the presence of mutation and unbalanced payoffs, they exhibit an attracting limit cycle. The period of this cycle depends on the rate of mutation; specifically, the period grows logarithmically as the mutation rate tends to zero. We find that this behaviour is not reproduced in stochastic simulations with a fixed finite population size. Instead, demographic noise present in the individual-based model dramatically slows down the progress of the limit cycle, with the typical period growing as the reciprocal of the mutation rate. Here we develop a theory that explains these scaling regimes and delineates them in terms of population size and mutation rate. We identify a further intermediate regime in which we construct a stochastic differential equation model describing the transition between stochastically-dominated and mean-field behaviour.

Keywords: cyclic dominance ecology, limit cycle, mean field model, replicator equation, stochastic differential equation, stochastic simulation

1. Introduction

Many mathematical models in ecology are well-known to be capable of generating oscillatory dynamics in time; important examples stretch right back to the initial work of Lotka and Volterra on predator-prey interactions [1, 2, 3, 4, 5, 6]. Such models, although dramatic simplifications when compared to real biological systems, have a significant impact in shaping our understanding of the modes of response of ecological systems and are helpful in understanding implications of different strategies for, for example, biodiversity management, and the structure of food webs [7].

Competition between species is a key driver of complex dynamics in ecological models. Even very simple competitive interactions can yield complex dynamical behaviour, for example the well documented example of the different strategies adopted by three distinct kinds of side-blotched lizard [8]. Similar cyclical interactions occur in bacterial colonies of competing strains of *E. coli* [9, 10, 11]. In mathematical neuroscience dynamical switches of this type have been referred to as ‘winnerless competition’ since there is no best-performing state overall [12, 13].

*Corresponding author

Email addresses: Q.Yang2@bath.ac.uk (Qian Yang), T.C.Rogers@bath.ac.uk (Tim Rogers), J.H.P.Dawes@bath.ac.uk (Jonathan H.P. Dawes)

Evolutionary Game Theory (EGT) provides a useful framework for modelling competitive interaction, in particular the replicator equations [14, 15] give a dynamical systems interpretation for models posed in game-theoretic language. Work by many authors, including in particular Hofbauer and Sigmund [16, 17] has resulted in a very good understanding of replicator equation models for competing species. Recent work has extended these deterministic approaches to consider stochastic effects that emerge from consideration of finite, rather than infinite, populations. The classic Rock–Paper–Scissors (RPS) provides an important example of stochastic phenomena in ecological dynamics. When mutation (allowing individuals to spontaneously swap strategies) is added to the replicator equations for the RPS game, the deterministic can exhibit damped oscillations that converge to a fixed point. In [18], it was shown that stochastic effects present in finite populations cause an amplification of these transient oscillations, leading to so-called quasi-cycles [19]. For smaller values of mutation rate, the deterministic system passes through a Hopf bifurcation, and a limit cycles appears. Some past studies exist on the role of noise around limit cycles, such as [20, 21], in which small-scale fluctuations around the mean-field equations are explored using Floquet theory. More recently, it has been discovered that noise can induce much stronger effects including counterrotation and bistability [22].

In this paper we combine deterministic and stochastic approaches in order to present a comprehensive description of the effect of demographic fluctuations around cycles of dominance in the RPS model. We determine three regimes, depending on the scaling of population size N and mutation rate μ . The basic link between stochastic individual-based dynamics and population-level ODEs is a theorem of Kurtz [23], allowing us to construct a consistent set of individual-level behaviours corresponding to the mean-field replicator dynamics for the RPS model that we take as our starting point. Between these two views of the dynamics lies a third: the construction of a stochastic differential equation (SDE) that captures the transition between them. Changing variables to the asymptotic phase of the ODE limit cycle reveals that the contribution of the stochasticity is to speed up some parts of the phase space dynamics and to slow down others but that the overall effect is to markedly increase the oscillation period. Our central conclusion is that as the stochastic effects become more important, the period of the oscillations increases rapidly, and this slowing down is a significant departure from the prediction of oscillation periods made on the basis of the mean-field ODE model.

The structure of the remainder of the paper is as follows. In section 2 we introduce the replicator dynamical model for the rock-paper-scissors game with mutation. The mean-field ODE version of the model is well-known and we derive a self-consistent individual-based description; this is not as straightforward as one might initially imagine. We show numerically that the two models give the same mean period for the cyclic dynamics when the mutation rate is large but disagree when μ is small.

In section 3 we summarise the computation to estimate the period of the limit cycle when μ is small. This follows the usual approach, dividing up trajectories into local behaviour near equilibrium points, and global maps valid near the unstable manifolds of these saddle points. Section 4 turns to the stochastic population model and analyses the dynamics in terms of a Markov chain. This leads to a detailed understanding of the individual-level behaviour in the limit of small mutation rate μ . Section 5 then fills the gap between the analyses of sections 3 and 4 by deriving an SDE that allows us to understand the relative contributions of the stochastic behaviour and the deterministic parts in an intermediate regime. Finally, section 6 discusses our results and concludes.

2. Models for Rock-Paper-Scissors with mutation

2.1. Rock-paper-scissors with mutation

Rock–Paper–Scissors (RPS) is a simple two-player, three-state game which illustrates the idea of cyclic dominance: a collection of strategies, or unchanging system states, in which each state in turn is unstable to the next in the cycle. In detail: playing the strategy ‘rock’ beats the strategy ‘scissors’ but loses to the strategy ‘paper’; similarly, ‘scissors’ beats ‘paper’ but loses to ‘rock’. When the two players play the same strategy the contest is a draw.

This information is summarised in the payoff matrix

$$P := \begin{pmatrix} 0 & -1 - \beta & 1 \\ 1 & 0 & -1 - \beta \\ -1 - \beta & 1 & 0 \end{pmatrix} \quad (1)$$

where $\beta \geq 0$ is a parameter that indicates that the loss incurred in losing contests is greater than the payoff gained from winning them. When $\beta = 0$, the row and column sums of P are zero: this is the simplest case. When $\beta > 0$, the game becomes more complicated, particularly when we would like to relate the behaviour at the population level to the individual level, as we discuss later in sections 2.2 and 2.3.

2.2. Deterministic rate equations

Setting the RPS game in the context of Evolutionary Game Theory (EGT), one considers a large well-mixed population of N players playing the game against opponents drawn uniformly at random from the whole population. We are then interested in the proportions of the total population who are playing different strategies at future times. The state of the system is given by the population fractions $(x_a(t), x_b(t), x_c(t)) := (N_A(t), N_B(t), N_C(t))/N$ where $N_{A,B,C}(t)$ are the numbers of players playing strategies A, B and C respectively.

The proportions of the population playing each strategy are expected to change over time according to the typical payoff received, as compared to the average over the whole population. The simplest mean field model for the resultant dynamics are the replicator equations

$$\dot{x}_i = x_i(t) \left(\sum_j P_{ij} x_j(t) - \sum_{j,k} P_{jk} x_j(t) x_k(t) \right), \quad (2)$$

where the subscripts i, j take values in $\{a, b, c\}$ and the proportions x_i sum to unity.

A common variant of the model introduces the additional mechanism of mutation between the three strategies, occurring between any pair with equal frequency. Mutation affects the rate of change of strategy i over time since the strategies other than i will contribute new players of i at rates μ while i will lose players at a rate given by $2\mu x_i$ as these players switch to a different strategy.

In the particular case of Rock–Paper–Scissors, the combined effects of the replicator dynamics together with mutations between strategies gives rise to the ordinary differential equations

$$\begin{aligned} \dot{x}_a &= x_a[x_c - (1 + \beta)x_b + \beta(x_ax_b + x_bx_c + x_ax_c)] + \mu(x_b + x_c - 2x_a), \\ \dot{x}_b &= x_b[x_a - (1 + \beta)x_c + \beta(x_ax_b + x_bx_c + x_ax_c)] + \mu(x_a + x_c - 2x_b), \\ \dot{x}_c &= x_c[x_b - (1 + \beta)x_a + \beta(x_ax_b + x_bx_c + x_ax_c)] + \mu(x_a + x_b - 2x_c). \end{aligned} \quad (3)$$

We note that these equations are to be solved in the region of $\mathbb{R}^3$ where all coordinates are non-negative. This region is clearly invariant under the vector field (3). Moreover, the constraint $x_a + x_b + x_c = 1$ is required to hold at all times.

The system (3) possess a single interior equilibrium point $x^* = (1/3, 1/3, 1/3)$, in which the three strategies are balanced. Straightforward linear stability analysis shows that this equilibrium is stable when $\mu > m_c = \beta/18$. Previous work by Mobilia [18] has shown that the system undergoes a Hopf bifurcation as μ is lowered and that for $\mu < \mu_c$ trajectories of (3) spiral away from x^* and are attracted to a unique periodic orbit which is stable, i.e a limit cycle.

2.3. Stochastic chemical reactions

The replicator equations shown above are expected to hold in the limit of infinitely large population size. In finite populations, however, the behaviour of many competing individuals is more appropriately modelled as a Markov process describing the random timing individual events. It is common practice to specify such a stochastic model as a chemical reaction scheme, which, chosen appropriately should recover the ODEs (3) in the limit of large systems. It is interesting to note that different stochastic individual-based models may give rise to the same mean-field ODEs, so that the question of constructing a stochastic reaction scheme starting from a particular set of ODEs does not have a unique answer. Moreover, the construction of the stochastic model is subject to a number of natural constraints, for example that all reaction rates are at all times non-negative.

With the goal of studying dynamics around the fixed point x^* a reaction scheme was proposed in [18] based on consideration of the frequency-dependent Moran model with rates chosen to match the mean-field equations as required in the infinite system limit. Although that scheme is well motivated and perfectly correct for the regimes previously studied, it will not be suitable for our study of the stochastic dynamics around the limit cycle, as the reaction rates can take negative (unphysical) values near the system boundaries. Thus our first task here is to choose a reaction scheme that permits us to study the whole of state space in the case $\beta > 0$, $\mu \ll 1$.

For the “replicator” part of the dynamics we propose the chemical reactions

Mutations are included in the model through the additional six reactions

This information can be more usefully summarised in terms of an integer ‘jump matrix’ $\mathbf{S}$ describing the changes to $\mathbf{x}$ caused by the various reactions (so element S_{ij} represents the increment or decrement of species i taking place in reaction j), and vector $\mathbf{r}(\mathbf{x})$

describing the reaction rates. For the reactions in (4) we have

$$\mathbf{S}_1 = \begin{pmatrix} -1 & 0 & 1 & -1 & 1 & 0 \\ 1 & -1 & 0 & 1 & 0 & -1 \\ 0 & 1 & -1 & 0 & -1 & 1 \end{pmatrix}, \quad (9)$$

$$\mathbf{r}_1 = (x_a x_b, x_b x_c, x_c x_a, \beta x_a x_b^2, \beta x_a^2 x_c, \beta x_b x_c^2)^T. \quad (10)$$

Whilst the mutation reactions in (8) are specified by

$$\mathbf{S}_2 = \begin{pmatrix} -1 & 0 & 1 & -1 & 1 & 0 \\ 1 & -1 & 0 & 0 & -1 & 1 \\ 0 & 1 & -1 & 1 & 0 & -1 \end{pmatrix}, \quad (11)$$

$$\mathbf{r}_2 = (x_a, x_b, x_c, x_a, x_b, x_c)^T. \quad (12)$$

The jump matrix and rate vector for the full scheme are found simply by concatenation:

$$\mathbf{S} = (\mathbf{S}_1, \mathbf{S}_2), \quad \mathbf{r} = \begin{pmatrix} \mathbf{r}_1 \\ \mathbf{r}_2 \end{pmatrix}. \quad (13)$$

We choose to express our reactions in this form in order to appeal to a very useful theorem of Kurtz [23], stating that in the limit $N \rightarrow \infty$ of large populations, the stochastic process described by these reactions converges to the deterministic dynamical system

$$\dot{\mathbf{x}} = \mathbf{S}\mathbf{r}(\mathbf{x}).$$

It is easy to check that the above scheme thus reproduces the system (3). Although we will develop our analysis for this particular chemical reaction scheme, it is important to reiterate that it is not unique in reproducing the replicator equations in the large population limit. However, we argue that the important features of the stochastic slowdown we observe will be common to all reasonable models, including those with more biologically realistic properties such as involving only two-body interactions, and permitting varying population sizes.

2.4. Simulations

To gain an initial insight into the differences between the deterministic and stochastic viewpoints for the RPS game with mutation, we use the Gillespie algorithm [24], the standard and widely used stochastic simulation algorithm (SSA), to simulate the chemical reactions (4)-(8) in order to illustrate the typical dynamics in the stochastic case and to compare that with the deterministic case.

Figure 1 presents results comparing the stochastic and deterministic cases for two different values of the mutation rate μ . In each plot we show a typical realisation of the stochastic simulation algorithm, for two different finite, but large, population sizes $N = 2^8$ (green dashed line) and $N = 2^{16}$ (red dashed line), together with a trajectory of the ODEs (blue solid line).

The solution to the ODEs is shown by the blue line in 1(a). The red dashed line in 1(a) starts from a very similar initial condition and evolves similarly: spiralling out towards the boundary of phase space and at large times occupying a region of phase space near to the limit cycle but with small fluctuations around it. The green dashed

line in 1(a) shows much larger fluctuations around the limit cycle, including excursions that take the simulation onto the boundaries of the phase space, and much closer to the corners. Note that because of the mutations, the boundaries of the phase space are not absorbing states for the random process (or invariant lines for the ODE dynamics).

Figure 1(b) illustrates the behaviour for a significantly larger value of μ for which the limit cycle for the ODEs (blue curve) lies much closer to the centre of the phase space. The SSA for $N = 2^{16}$ lies close to the limit cycle but fluctuates around it; for $N = 2^8$ the fluctuations are much larger and the sample path of the stochastic process lies outside the limit cycle for a large proportion of the simulation time.

To quantify the differences between the three results shown in each part of figure 1, we focus on one specific aspect of the dynamics: the period of the oscillations around the central equilibrium point at $(1/3, 1/3, 1/3)$. For the ODEs, the period of the limit cycle can be defined to be the smallest elapsed time between successive crossings of a hyperplane in the same direction, for example the plane $x_a = 1/2$ (this is a sensible choice because, as we will later show, the effect of noise is smallest at this part of the cycle). We denote by T_{ODE} the period of the deterministic limit cycle. In stochastic simulations, a trajectory may by chance cross a hyperplane back and forth several times in quick succession, thus the time between crossings may not represent a full transit of the cycle. We avoid this complication by measuring the period as three times the transit time between a hyperplane and its $2\pi/3$ rotation. The expected value of this oscillation period in the stochastic case we denote by T_{SSA} .

Figure 1: Illustrative comparisons between trajectories of the ODEs (blue solid lines) and realisations of the stochastic simulations (red and green dashed lines). ‘Unbalanced’ means that the payoff matrix is not zero-sum, i.e. $\beta > 0$. (a) When $\mu \ll \mu_c$ the trajectories lie close to the boundary of the phase space, $\beta = 1/2$, $\mu = 1/216$. (b) When μ is only slightly smaller than μ_c , trajectories lie much closer to the central equilibrium point, $\beta = 1/2$, $\mu = 5/198$.

Figure 2 provides a quantitative comparison of the different dependencies of the average period T_{SSA} of the stochastic simulations and the period T_{ODE} of the deterministic ODEs on the mutation rate μ , for $0 < \mu < \mu_c$. The blue solid line indicates the relatively slow increase in their period T_{ODE} as μ decreases. The red error bars indicate the range of values of the oscillation period in the stochastic case, with the averages of those values shown by the red dots. The data in this figure for the stochastic simulations was obtained from simulations at a fixed value $N = 2^{17}$ with μ varying, but

for presentational reasons that will become clearer later in the paper we have chosen to plot the scaled quantity $\mu N \ln N$ on the horizontal axis. We then observe that if μ is sufficiently small then the difference between T_{ODE} and T_{SSA} is very significant: the oscillations in the stochastic simulation have a much longer period, on average, than that predicted by the ODEs, while if $\mu N \ln N$ is larger than unity, the agreement in terms of the oscillation period, between the deterministic and stochastic simulations is very good. We label these two regimes ‘Region III’ and ‘Region I’ respectively. The cross-over region where $\mu N \ln N \approx 1$ we label ‘Region II’.

Figure 2: Comparison between the period T_{ODE} of the limit cycle in the deterministic case and the average period T_{SSA} of oscillations in the stochastic case, as a function of the parameter μ at fixed $N = 2^{17}$. The blue line shows T_{ODE} and red dots indicate T_{SSA} . As μ decreases, the two periods start to separate from each other in Region II where $\mu N \log N \approx 1$.

In the following sections of the paper we will focus on each of the three regions in turn. In section 3 we study the period of the limit cycle in the ODEs (region I). In section 4 we analyse the stochastic dynamics to determine the average period of cycles in region III. In section 5 the cross-over between the deterministic and stochastic regions is understood through the analysis of an SDE that combines both deterministic and stochastic effects, thus describing region II.

3. Analysis of the periodic orbit in region I

Region I is defined to be the right hand side of figure 2; more precisely, it is the regime in which $N \gg 1$ and $(N \log N)^{-1} \ll \mu < \mu_c$. In this region, stochastic simulation results show only very small fluctuations about a mean value, and this mean value coincides well with numerical solutions to the ODEs. This is evidence that the behaviour of the deterministic ODEs provides a very good guide to the stochastic simulation results in this region. Hence our goal in this section is to explain the asymptotic result that the period $T_{\text{ODE}} \propto -3 \log \mu$ for μ small, but for large, even infinite, N as shown in figure 3.

Figure 3: Period T_{SSA} for three different values of N , compared with the result for T_{ODE} (blue solid line) on a log-linear plot indicating the scaling law $T_{\text{ODE}} \propto -3 \log \mu$ which applies in region I where the stochastic simulations behave in a similar way to the ODEs.

The standard approach to the analysis of trajectories near the limit cycle, in the regime where it lies close to the corners of the phase space, is to construct local maps that analyse the flow near the corners, and global maps that approximate the behaviour of trajectories close to the boundaries [25]. Figure 4 shows the construction of a local map in the neighbourhood of the corner where $x_c = 1$, following by the global map from this neighbourhood to a neighbourhood of the corner where $x_a = 1$. In this section we consider local and global maps in turn, in sections 3.1 and 3.2, and then in section 3.2 we study their composition, and deduce an estimate for the period T_{ODE} of the limit cycle.

3.1. Local map

Using the notation as in figure 4, we begin by defining a neighbourhood of the corner $x_c = 1$ by setting $0 < h \ll 1$ to be a small positive constant. We assume that the trajectory for the ODEs (3) starts at the point $(x_a^{(0)}, h, x_c^{(0)})$ at time $t = 0$, and arrives at the point $(h, x_b^{(1)}, x_c^{(1)})$ at time $t = T_1 > 0$. For the whole time $0 < t < T_1$ the trajectory lies close to the corner $(0, 0, 1)$, so we suppose that $0 < x_a, x_b, u \ll 1$ where $u := 1 - x_c$. Note that we have $x_a + x_b - u = 0$ since $x_a + x_b + x_c = 1$.

Then the behaviour of the ODEs (3) in this neighbourhood is very similar to that of their linearisation obtained by dropping terms higher than linear order in the small quantities $0 < x_a, x_b, u \ll 1$. For the linearised system we obtain

$$\begin{cases} \dot{x}_a = \mu - x_a(3\mu - 1), \\ \dot{x}_b = \mu - x_b(1 + \beta + 3\mu), \\ \dot{u} = -x_b(1 + \beta) + x_a + 2\mu - 3\mu u. \end{cases} \quad (14)$$

The equations (14) are linear and constant coefficient and so can be solved analytically; note that the first and second equation are actually decoupled from each other and from

Figure 4: This picture shows the local map and the global map starting from initial position (x_a^0, h, x_c^0) .

the u equation. Integrating from $t = 0$ to the time $t = T_1$, we denote the point on the trajectory by $(x_a(T_1), x_b(T_1), x_c(T_1)) \equiv (x_a^{(1)} = h, x_b^{(1)}, x_c^{(1)})$ where

$$x_b^{(1)} = \frac{\mu}{\gamma} (1 - e^{-\gamma T_1}) + h e^{-\gamma T_1}, \quad (15)$$

$$x_a^{(1)} = h = \frac{\mu}{\hat{\gamma}} + e^{-\hat{\gamma}T_1} \left(x_a^0 - \frac{\mu}{\hat{\gamma}} \right) \quad (16)$$

where $\gamma := 1 + \beta + 3\mu$ and $\hat{\gamma} := 3\mu - 1$. Equation (16) allows us to express T_1 in terms of $x_a^{(0)}$:

$$T_1 = -\frac{1}{\hat{\gamma}} \log \left(\frac{h - \mu/\hat{\gamma}}{x_a^{(0)} - \mu/\hat{\gamma}} \right), \quad (17)$$

and we can now use (17) to eliminate T_1 from (15) and obtain a relationship between $x_a^{(0)}$ and $x_b^{(1)}$ which takes the form

$$x_b^{(1)} = \frac{\mu}{\gamma} + \left(h - \frac{\mu}{\gamma}\right) \left(\frac{h - \mu/\hat{\gamma}}{x_a^{(0)} - \mu/\hat{\gamma}}\right)^{\gamma/\hat{\gamma}}. \quad (18)$$

This relationship is the key part of the local map near the point $(x_a, x_b, x_c) = (0, 0, 1)$ that we will use in what follows.

3.2. Global map

For the global map, we observe that trajectories remain close to one of the boundaries (in this case, the boundary $x_b = 0$) and so we propose that the trajectory starting from

$(x_a^{(1)} = h, x_b^{(1)}, x_c^{(1)})$ arrives at the point $(x_a^{(2)}, x_b^{(2)}, h)$ at time $t = T_2$. Referring to (3), the ODE for x_b near the boundary can be well approximated by taking just $\dot{x}_b = \mu$ (since $x_a + x_c = 1$ when $x_b = 0$), so its solution is

$$x_b^{(2)} = x_b^{(1)} + (T_2 - T_1)\mu = x_b^{(1)} + \mathcal{C}_0\mu, \quad (19)$$

where we denote the elapsed time by $\mathcal{C}_0 := T_2 - T_1$. As is typical in these analyses, trajectories take a relatively short time to arrive at the hyperplane $x_c^{(2)} = h$ starting from $x_c^{(1)}$, compared to the time taken to move along the part of the trajectory from $x_c^{(0)}$ to $x_c^{(1)}$; this is intuitively because the absolute value of $\dot{x}_c$ on the global part of the map between T_1 and T_2 is much larger than on the local part, i.e. when $0 < t < T_1$. As a result, the time taken on the global part of the map, $\mathcal{C}_0 := T_2 - T_1$ is much less than the local travel time T_1 ; the majority of the time spent on the limit cycle is taken up with travel near the corners.

The composition of local and global maps near the corner $x_c = 1$ and boundary $x_b = 0$ is now straightforward: we combine (18) and (19) to obtain:

$$x_b^{(2)} = \mathcal{C}_0\mu + \frac{\mu}{\gamma} + \left(h - \frac{\mu}{\gamma}\right) \left(\frac{h - \mu/\hat{\gamma}}{x_a^{(0)} - \mu/\hat{\gamma}}\right)^{\gamma/\hat{\gamma}}. \quad (20)$$

We can now use the permutation symmetry inherent in the dynamics to complete the analysis. Due to the fact that the model is rotationally symmetric, the next stage of the evolution is local map again with the same parameter values. The trajectory will start from $(x_a^{(2)}, x_b^{(2)}, h)$ and stay near the corner $(1, 0, 0)$ for long time before arriving at a point, say, $(x_a^{(3)}, h, x_c^{(3)})$ where we will construct another global map, and so on:

$$x_a^{(0)} \xrightarrow{\text{local}} x_b^{(1)} \xrightarrow{\text{global}} x_b^{(2)} \xrightarrow{\text{local}} x_c^{(3)} \xrightarrow{\text{global}} x_c^{(4)} \xrightarrow{\text{local}} x_a^{(5)} \xrightarrow{\text{global}} x_a^{(6)} \rightarrow \dots \quad (21)$$

which can be summarised further as

$$x_0 \xrightarrow{\text{local \& global}} x_1 \xrightarrow{\text{local \& global}} x_2 \xrightarrow{\text{local \& global}} x_3 \rightarrow \dots \quad (22)$$

Equations (21) and (22) define a one-to-one correspondence between the points $\{(x_a^{(n)}, x_b^{(n)}, x_c^{(n)})\}$ on a trajectory and a sequence of values (selecting appropriate coordinates) $\{x_n\}$. From the previous discussion on local and global maps, the map that generates the sequence $\{x_n\}$ takes the form

$$x_{n+1} = \mathcal{C}_0\mu + \frac{\mu}{\gamma} + \left(h - \frac{\mu}{\gamma}\right) \left(\frac{h - \mu/\hat{\gamma}}{x_n - \mu/\hat{\gamma}}\right)^{\gamma/\hat{\gamma}}, \quad (23)$$

where, as before, $\gamma := 1 + \beta + 3\mu$ and $\hat{\gamma} := 3\mu - 1$.

If iterates of the map (23) converge to a fixed point then this corresponds to a stable limit cycle for the ODE dynamics. We now estimate the location of this fixed point and deduce an estimate for the period of the resulting limit cycle.

Let $y_n := x_n/\mu$ be a scaled version of x_n , then (23) can be written as:

$$\begin{aligned} y_{n+1} &= \mathcal{C}_0 + \frac{1}{\gamma} + \left(h - \frac{\mu}{\gamma}\right) \frac{1}{\mu} \left(\frac{\mu y_n + \frac{\mu}{1-3\mu}}{h + \frac{\mu}{1-3\mu}}\right)^{\frac{1+3\mu+\beta}{1-3\mu}} \\ &= \mathcal{C}_0 + \frac{1}{\gamma} + \left(h - \frac{\mu}{\gamma}\right) \mu^{\frac{1+3\mu+\beta}{1-3\mu}-1} \left(\frac{y_n(1-3\mu) + 1}{h(1-3\mu) + \mu}\right)^{\frac{1+3\mu+\beta}{1-3\mu}}, \end{aligned} \quad (24)$$

where $\gamma := 1 + 3\mu + \beta$. Denoting the fixed point of the map by y^* , from (24) we see that $y^* = \mathcal{C}_0 + 1/\gamma + o(\mu)$ in the limit $\mu \rightarrow 0$. This observation is crucial in order to ensure that we obtain the correct leading-order behaviour and distinguish carefully between the various small quantities in the problem. Then it follows that $x^* = \mu y^* \approx \mathcal{C}_1 \mu$, where $\mathcal{C}_1 = \mathcal{C}_0 + 1/(1 + \beta)$, in the limit $\mu \rightarrow 0$.

Introducing this leading-order approximation for x^* into (17), we obtain an estimate of the time spent in a neighbourhood of the corner T_1 of the stable limit cycle:

$$T_1 = \frac{1}{1 - 3\mu} \log \left(\frac{h + \mu(1 - 3\mu)}{\mathcal{C}_1 \mu + \mu(1 - 3\mu)} \right). \quad (25)$$

Since in this case μ is assumed to be very small, (25) takes the form, at leading order,

$$T_1 = \frac{1}{1 - 3\mu} \log \left(\frac{1}{\mu} \right) + \mathcal{B}_0 \text{ as } \mu \rightarrow 0. \quad (26)$$

where $\mathcal{B}_0$ is a constant.

Finally, as remarked on above, because trajectories remain near each corner for large parts of the period of the orbit, the local map travel time T_1 is the dominant contribution, compared to the time spent on the global map. Hence the period of the limit cycle is given at leading order by considering only the contribution from the three local maps required in one full period of the limit cycle. Hence our estimate for T_{ODE} becomes

$$T_{\text{ODE}} \approx 3T_1 = -3 \log \mu + \mathcal{B}_1, \quad (27)$$

as $\mu \rightarrow 0$.

4. Analysis of the periodic orbit in region III

We now turn our attention to Region III, where the period of the orbit in the stochastic simulations increases much more rapidly as μ decreases, at fixed finite N , than the prediction from the analysis of the ODEs in section 3 above suggests. Figure 5 illustrates this by plotting the period T_{SSA} as a function of $\mu N \ln N$ for three different values of N . By plotting, on a log-log scale, the mean values of the periods and the error bars from an ensemble of stochastic simulations we observe that the period $T_{\text{SSA}} \approx (\mu N \ln N)^{-1}$ for small μ , with a constant that does not demonstrate any systematic dependence on N . In fact, the range of values of N presented here is small: we cannot distinguish from these numerical results the precise form of the dependence on N .

As well as the dependence of the period T_{SSA} on μ it is also of interest to determine the extent of Region III in which this scaling behaviour applies; in other words how small is μ required to be in order to move into this regime? The numerical data in figure 5 indicate that the cross-over from the ODE result to this new stochastic scaling arises when $\mu N \log N = 1$.

In this section, then, our aim is to explain firstly why this new scaling regime in Region III exists, and secondly, why it extends as far as $\mu N \ln N = 1$. We will find that in fact the period should scale as $T_{\text{SSA}} \approx (\mu N)^{-1}$ for small μ , but that the observation that the cross-over occurs when $\mu N \ln N = 1$ is correct; this explains why we have chosen to plot figure 5 in the form that it appears.

Careful examination of the numerical simulations in this regime show that their behaviour is qualitatively different from that in Region I discussed above. In stochastic

Figure 5: Log-log plot of the period T_{SSA} as a function of the mutation rate μ for three different values of N , showing that in region III the mean period scales roughly as $T_{\text{SSA}} \approx (\mu N \ln N)^{-1}$.

simulations the system becomes strongly attracted to the corner states (which would be absorbing states in the absence of mutation) and then can only escape from a corner when a mutation occurs; mutations are rare when μ is small. Our analysis later in this section shows that the **mean period of oscillation** is dominated by the contribution from the time needed to escape one step from a corner. Trajectories then typically move along a boundary towards the next corner. We show below that, although this latter part requires at least $N - 1$ steps, the expected total time required is less than the waiting time to escape from the corner. Our schematic approach is illustrated in figure 6 while sketches this separation into a first step away from a corner, following by movement along the adjoining boundary.

In the following subsections we consider in detail three issues: firstly, in section 4.1 we show that the probability that the motion along a boundary is towards the next corner in the sequence, and that the systems state hits this next corner with a probability that tends to 1 as $N \rightarrow \infty$. Since we are interested in the regime where the mutation rate is very small, we carry out this calculation in the limit $\mu = 0$. In section 4.2 we compute the expected time until the system hits this next corner, again setting $\mu = 0$. Finally, in section 4.3 we compare this expected time for the system state to evolve along the boundary with the expected time until a mutation occurs. Together, this analysis confirms the intuitive picture outlined in the previous paragraph.

4.1. Probability of hitting the absorbing status

In this section we consider the state-discrete and time-continuous Markov Chain dynamics of the system evolving along a boundary, ignoring the effect of mutation. In this case the system is a one dimensional chain of states labelled 0 to N , corresponding to, for example, the number of strategy B individuals invading a population of size N initially composed of all strategy A individuals. Let $h_i = \mathbb{P}_i(\text{hit } 0)$ be the probability that the system hits (and is therefore absorbed by) state 0 having started at node i , and similarly let $\bar{h}_i = \mathbb{P}_i(\text{hit } N)$ be the probability of hitting state N having started at node i .

Figure 6: Sketch of the dynamics in the small- μ stochastic limit. The system is strongly attracted to corner states where the population is all in one state. Mutation is then the only mechanism for escape, and the expected time for the first mutation (curved arrows) is longer than the expected time required for subsequent steps along a boundary towards the next corner (straight lines with arrows).

When the system is in state i there are two possible moves: jumps to the left or to the right. The rate at which jumps to the left occur is $p_i = (N - i)i^2/N^2$; the rate at which jumps to the right occur is $q_i = (N - i)i^2(1 + \beta)/N^2 + (N - i)i/N$. The transition probabilities of moving to the left and to the right are then $\ell_i = p_i/(p_i + q_i)$ and $r_i = q_i/(p_i + q_i)$; clearly $\ell_i + r_i = 1$. We remark that the ratio of transition probabilities can be simplified to be

$$\frac{\ell_i}{r_i} = \frac{p_i}{q_i} = \frac{1}{(1 + \beta) + \frac{N}{i}}. \quad (28)$$

This allows us to derive a recurrence relation for the probability h_i that the system hits the state 0 starting from state i :

$$\begin{cases} h_0 = 1, \\ h_i = \ell_i h_{i-1} + r_i h_{i+1}, & \text{for } i = 1 \dots N - 1, \\ h_N = 0. \end{cases} \quad (29)$$

Recurrence relations of this type are straightforward to solve by standard methods [26]. In our case we find

$$h_1 = \frac{\sum_{i=1}^{N-1} \gamma_i}{1 + \sum_{i=1}^{N-1} \gamma_i}, \quad (30)$$

and of course $\bar{h}_1 = 1 - h_1$. Here we define γ_i as

$$\gamma_i := \prod_{j=1}^i \frac{\ell_j}{r_j} = \prod_{j=1}^i \frac{j}{j(1 + \beta) + N} < \frac{i!}{N^i} \quad (31)$$

We now examine the limiting behaviour of this result and prove that $h_1 \rightarrow 1$ as $N \rightarrow \infty$. Since

$$0 < \gamma_{i+1} = \frac{p_{i+1}}{q_{i+1}} \gamma_i,$$

and $p_{i+1}/q_{i+1} < 1$, it follows that $\gamma_{i+1} < \gamma_i$. Also, since $\gamma_i < i!/N^i$, we conclude that

$$\begin{aligned} 0 < \sum_{i=1}^{N-1} \gamma_i &= \gamma_1 + \sum_{i=2}^{N-1} \gamma_i \\ &\leq \frac{1}{N} + (N-2) \frac{2}{N^2}, \end{aligned}$$

which clearly tends to zero as $N \rightarrow \infty$. Hence

$$\lim_{N \rightarrow \infty} \sum_{i=1}^{N-1} \gamma_i = 0. \quad (32)$$

In conclusion, we have shown that, given that it starts at state 1, the probability of the system hitting the absorbing state at N tends to 1 as $N \rightarrow \infty$.

4.2. Average hitting time

Having shown that the system reaches state N with high probability, we now consider the expected time until this happens. Let T_i be the first time at which the system hits an absorbing state starting from state i at time 0, i.e. the time at which the system hits either 0 or N . Define $\tau_i = \mathbb{E}(T_i)$ for $0 \leq i \leq N$, the expected hitting time starting from state i . Clearly we have that $\tau_0 = \mathbb{E}(T_0) = 0$ and $\tau_N = \mathbb{E}(T_N) = 0$. Through a similar calculation to section 4.1, we now compute τ_1 , the mean time taken to reach either 0 or N , starting from 1.

The recurrence formula for expected hitting times τ_i can be computed as the sum of the expected time spent in state i before jumping either left or right, plus the expected future time required when in that new state:

$$\tau_i = \frac{1}{p_i + q_i} + \ell_i \tau_{i-1} + r_i \tau_{i+1}, \quad \text{for } i = 1 \dots N-1. \quad (33)$$

This three-term recurrence can be rearranged to give a two term recurrence on the variable $\tau_{i-1} - \tau_i$, whose solution yields

$$\tau_1 = \left(1 + \sum_{i=1}^{N-1} \gamma_i \right)^{-1} \sum_{i=1}^{N-1} \left(\frac{1}{q_i} \sum_{j=i}^{N-1} \frac{\gamma_j}{\gamma_i} \right). \quad (34)$$

We now wish to examine the asymptotic behaviour of this expression for the expected hitting time, in the limit when $N \gg 1$. Initial numerical explorations lead us to propose that $\tau_1 \propto \ln N$ when N is sufficiently large. In the remainder of this section we will deduce this estimate systematically.

First, note that from (32) we know that

$$\lim_{N \rightarrow \infty} \left(1 + \sum_{i=1}^{N-1} \gamma_i \right) = 1. \quad (35)$$

From the discussion before (28) we have that

$$\begin{aligned} q_i &= i(N-i) \frac{1}{N} \left[\frac{i}{N} (1 + \beta) + 1 \right] \\ \Rightarrow \frac{1}{q_i} &= \frac{N}{i(N-i)} \frac{1}{1 + (1 + \beta) \frac{i}{N}}. \end{aligned} \quad (36)$$

Since

$$\frac{1}{2+\beta} < \frac{1}{1+(1+\beta)\frac{i}{N}} < 1,$$

(as $1 \leq i \leq N-1$), we can bound $\sum_{i=1}^{N-1} 1/q_i$ as follows

$$\begin{aligned} \frac{1}{2+\beta} \sum_{i=1}^{N-1} \frac{N-i+i}{(N-i)i} &\leq \sum_{i=1}^{N-1} \frac{1}{q_i} \leq \sum_{i=1}^{N-1} \frac{N-i+i}{(N-i)i} \\ \implies \frac{2}{2+\beta} \ln(N-1) &\leq \sum_{i=1}^{N-1} \frac{1}{q_i} \leq 2(1 + \ln(N-1)), \end{aligned} \quad (37)$$

Next we prove that, for any $1 \leq i \leq N-1$, $\sum_{j=i}^{N-1} \gamma_j/\gamma_i$ is bounded as $N \rightarrow \infty$. From the definition of γ_j we see that

$$\frac{\gamma_j}{\gamma_i} = \prod_{k=i+1}^j \frac{1}{(1+\beta) + \frac{N}{k}} < \prod_{k=i+1}^j \frac{k}{N+k} = \frac{j!}{i!} \frac{(N+i)!}{(N+j)!}. \quad (38)$$

Therefore,

$$\sum_{j=i}^{N-1} \frac{\gamma_j}{\gamma_i} < \sum_{j=i}^{N-1} \frac{j!}{i!} \frac{(N+i)!}{(N+j)!} = \frac{1}{N-1} \left[i + N - 2N \frac{N!(N+i)!}{i!(2N)!} \right] \leq 2.$$

We note that if we set $i = N-1$ in the above, then the limiting value is small, due to the influence of the negative term that is a ratio of factorials, but the limit must still be positive; in the case $i = 1$ we have a limit that is closer to 1. In all cases, since $1 \leq i \leq N-1$ the limiting value must be non-negative, and at most 2. Hence we see that

$$\lim_{N \rightarrow \infty} \sum_{j=i}^{N-1} \frac{\gamma_j}{\gamma_i} \leq 2. \quad (39)$$

We now apply the results (35), (37) and (39) to (34) in order to deduce the result

$$\tau_1 \leq 4(1 + \ln(N-1)) \quad (40)$$

which is an upper bound on the expected time, starting in state 1, until the system hits one of the two absorbing states 0 or N . Together with the conclusion of the previous section, in which we computed that the probability that the system arrives in state N tends to 1 as $N \rightarrow \infty$, we can conclude that the expected time required for the system to hit state N is no greater than $4(1 + \ln N)$. In using this result later, we will omit the subdominant constant term 1 since we are concerned primarily with values of N for which $\ln N \gg 1$, in the limit $N \rightarrow \infty$.

4.3. The effect of mutation on the stochastic dynamics

The analysis in the previous two subsections ignored the role of mutation in order to understand the dynamics on the boundary of phase space and, in particular, to estimate the time required to travel along the boundary to a corner.

Since in the absence of mutation the corners are absorbing states, mutation plays an important role in moving the system from a corner (state 0) to state 1 on the boundary, allowing it then to travel further towards the next corner. The mutation rate μ , together with our assumptions on the stochastic dynamics, imply that, for any system state, the time until the next mutation event M occurs is exponentially distributed:

$$\mathbb{P}(M > t) = e^{-\mu N t}$$

For system states on the boundary of phase space, we would expect that mutations would move them away, into the interior, where the analysis in the previous sections might become less useful. This is unlikely to happen if mutations are not expected during the time taken for the system to move along the whole boundary, i.e. if

$$\mathbb{P}(\text{hit state } N \text{ before mutation}) = \mathbb{P}(M > 4 \ln N) = e^{-4\mu N \ln N} \quad (41)$$

is close to 1. From the form of (41), if N is fixed and $\mu \rightarrow 0$, then this probability of hitting state N before any mutation occurs tends to 1, which implies that the system remains in boundary states and the analysis of sections 4.1 and 4.2 applies. On the other hand, if $N \rightarrow \infty$, and μ remains fixed, then mutation is expected to occur before the system reaches the next corner, and so the system state tends to leave the boundary. The intermediate balance between these two regimes occurs when $\mu \sim (N \ln N)^{-1}$.

An equivalent discussion can be framed in terms of the sketch of the dynamics indicated in figure 6. The time required for a full period of the oscillation is composed of two contributions on each boundary piece: the first contribution τ_m is the time required to jump, via mutation, from a corner to a state with one new individual of the appropriate kind. This mutation takes an expected time $\tau_m \sim 1/(N\mu)$ since there are N individuals and each mutates independently at a rate μ . The second contribution τ_1 is the time required to traverse the boundary starting from state 1. This is approximately $\tau_1 = 4 \ln N$. If $\tau_m \gg \tau_1$, i.e. $\mu \ll (4N \ln N)^{-1}$, then the largest contribution to the **mean period of oscillation** is from the mutation events, and so we expect in this regime to have the period of the orbit being dominated by the time required for three independent mutations to occur, i.e. $T_{SSA} \sim 3/(N\mu)$.

5. Analysis of the periodic orbit in region II

Region II is the cross-over region between ‘large μ ’ where the ODE approximation is valid, and ‘small μ ’ where the stochastic approach based on a Markov Chain, is appropriate. Using the analysis of the previous section we see mathematically speaking that region II arises where $\mu \sim (N \ln N)^{-1}$. In region II the stochastic simulations show large fluctuations around the ODE predictions, but the system does not spend time always near the boundaries of phase space, so it is not clear that the analysis of region III should apply directly.

In this section we examine region II, and explain why the fluctuations in the stochastic system act to increase the period of the oscillations rather than to decrease it. This involves a third approach to the dynamics, using a stochastic differential equation derived from the chemical reaction model and which is valid for large, but not infinite, system sizes. We show that the SDE approach captures, in some detail, the transition between the deterministic and the fully stochastic regimes described in previous sections of the paper.

5.1. Stochastic differential equation

It was proved by Kurtz [23] that trajectories of Markov jump processes specified in terms of a stoichiometric matrix $\mathbf{S}$ and rate vector $\mathbf{r}(\mathbf{x})$, as discussed in section 2.3, are well approximated for large N by the trajectories of the SDE

$$d\mathbf{x} = \mathbf{A}(\mathbf{x})dt + \frac{1}{\sqrt{N}}\mathbf{G}(\mathbf{x})d\mathbf{W}(t). \quad (42)$$

Here $d\mathbf{W}_t$ is a vector, each element of which is an independent Wiener process [27], and the vector $\mathbf{A}$ and matrix $\mathbf{G}$ are given by

$$\mathbf{A}(\mathbf{x}) = \mathbf{S}\mathbf{r}(\mathbf{x}), \quad \mathbf{G}_{ij}(\mathbf{x}) = \mathbf{S}_{ij}\sqrt{\mathbf{r}_j(\mathbf{x})}. \quad (43)$$

In the calculation of the statistical properties of this equation, it is more useful to consider the matrix $\mathbf{B} = \mathbf{G}\mathbf{G}^T$. For our system we obtain the explicit formulas

$$\begin{aligned} \mathbf{A}_1 &= x_a[x_c - (1 + \beta)x_b + \beta(x_ax_b + x_bx_c + x_ax_c)] + \mu(x_b + x_c - 2x_a), \\ \mathbf{A}_2 &= x_b[x_a - (1 + \beta)x_c + \beta(x_ax_b + x_bx_c + x_ax_c)] + \mu(x_a + x_c - 2x_b), \end{aligned} \quad (44)$$

$$\begin{aligned} \mathbf{B}_{11} &= x_ax_b + x_cx_a + (x_ax_b^2 + x_cx_a^2)(2 + \beta) + \mu(2x_a + x_b + x_c), \\ \mathbf{B}_{12} &= \mathbf{B}_{21} = -(x_ax_b + x_ax_b^2(2 + \beta) + \mu(x_a + x_b)), \\ \mathbf{B}_{22} &= x_ax_b + x_bx_c + (x_ax_b^2 + x_bx_c^2)(2 + \beta) + \mu(x_a + 2x_b + x_c). \end{aligned} \quad (45)$$

5.2. Asymptotic phase of points near the periodic orbit

In order to explain the increase in the period of the orbit as the fluctuations grow, we use the SDE (42) to derive an equation for the angular velocity around the limit cycle, and then compute the period of the limit cycle by integrating the angular velocity. In this section we extend this idea by deriving an SDE for the angular velocity. This allows us to investigate the effects of noise on the period of the orbit. A fully analytic approach is unfortunately not possible, so our approach is a combination of numerical and analytic methods. Identifying the correct scalings for features of the limit cycle however enables us to confirm the various asymptotic scalings found in regions I and III and to see how they both contribute in this region, region II.

Since the phase space is two dimensional, we follow the presentation started in section 2.4, showing the periodic orbit in the plane $\mathbb{R}^2$ using coordinates

$$\mathbf{y} = (y_1, y_2) = \left(x_a + \frac{1}{2}x_b, \frac{\sqrt{3}}{2}x_b \right). \quad (46)$$

These, and other definitions for our coordinate systems, are illustrated in figure 7. Let $\Gamma \subset \mathbb{R}^2$ be the set of points on the limit cycle, i.e. the red circle in figure 7. We further define the polar coordinates (ρ, θ) based on the centre of the triangle, in the $\mathbf{y}$ coordinates, i.e. let

$$\rho(\mathbf{y}) = \sqrt{\left(y_1 - \frac{1}{2}\right)^2 + \left(y_2 - \frac{\sqrt{3}}{6}\right)^2},$$

Figure 7: The red circle is the limit cycle when $\beta = 1/2, \mu = 1/216$, the blue triangle is the limit cycle when $\beta = 1/2, \mu = 10^{-6}$. When μ is very small, it is hard to tell the difference between triangle boundary and limit cycle without zooming-in.

and

$$\theta(\mathbf{y}) = \arctan \left(\frac{y_1 - 1/2}{\sqrt{3}/6 - y_2} \right).$$

Let T be the period of limit cycle, i.e. T is the smallest positive value such that $\forall \mathbf{y} \in \Gamma$, $\mathbf{y}(t) = \mathbf{y}(t + T)$ but for any $0 < T' < T$, $\mathbf{y}(t) \neq \mathbf{y}(t + T')$. All points in the interior of the limit cycle (except the equilibrium point at $\mathbf{y} = (1/2, \sqrt{3}/6)$) are attracted to the limit cycle, enabling us to define a ‘landing point’ on Γ to which they are asymptotically attracted. Although the ω -limit set of a point $\mathbf{y}_0 \in \mathbb{R}^2$ would clearly be the entire orbit Γ , by looking at the sequence defined by advancing for multiples of the period T we can identify a single limit point $p_\infty(\mathbf{y}_0)$. Specifically we write the time-evolution map for the ODEs as $\phi_t(\mathbf{y}_0) := \mathbf{y}(t)$ where $\mathbf{y}(t)$ solves the ODEs (3) subject to the initial condition $\mathbf{y}(0) = \mathbf{y}_0$. Then we define the sequence $\{p_n\}_{n \geq 0}$ by

$$p_n = \phi_{nT}(\mathbf{y}_0), \quad \text{and} \quad p_0 = \mathbf{y}_0.$$

and the limit point

$$p_\infty(\mathbf{y}_0) := \lim_{n \rightarrow \infty} \phi_{nT}(\mathbf{y}_0).$$

Note that $p_\infty(\mathbf{y}_0) \in \Gamma$ always, and that if $\mathbf{y}_0 \in \Gamma$ then $p_\infty(\mathbf{y}_0) = \mathbf{y}_0$.

We can now define the asymptotic phase $\varphi(\mathbf{y}_0)$ of a point near, but not necessarily on, the limit cycle by setting

$$\varphi(\mathbf{y}_0) := \theta(p_\infty(\mathbf{y}_0)). \quad (47)$$

So that if $\mathbf{y}_0 \in \Gamma$, then $\varphi(\mathbf{y}_0) = \theta(\mathbf{y}_0)$, and curves on which $\varphi(\mathbf{y}_0)$ is constant cross through Γ at these points. We can use $\varphi(\mathbf{y}_0)$ to consider the influence of noise, which pushes trajectories off the limit cycle, causing time advances or delays, as illustrated in figure 8. This concept has been explored recently for a more general class of limit cycles in [22], whose approach we follow here. It is analogous to the noise-induced drift observed in various ecological models including invasions [28] and the evolution of altruism [29]; see [30] for an introduction.

The three-fold rotation symmetry of the problem suggests that it is enough to focus on the interval $\varphi \in [-\pi/3, \pi/3]$.

Figure 8: Illustration of the idea of the asymptotic phase of small perturbations, using the lower part of Γ as shown in the (y_1, y_2) plane. Each of the three black points lies on the limit cycle Γ : trajectories evolve along Γ from left to right in the figure. Perturbations of these points (in green) towards the boundary lead to trajectories that have longer periods than Γ has, hence the green arrows indicate the asymptotic convergence of trajectories back to the limit cycle to points that lie to the left of the black dots. In contrast, perturbations (red dots) towards the interior of Γ lead to states that are accelerated by trajectories and converge asymptotically to points on Γ that are ahead of the black dots.

5.3. A stochastic differential equation for φ

Since φ is a function on the phase space, we can derive an SDE for the evolution of φ using (42) and Ito's formula [27]. We obtain

$$\begin{aligned} d\varphi(\mathbf{x}) = & \left\{ \sum_i \mathbf{A}_i(\mathbf{x}, t) \partial_i \varphi(\mathbf{x}) + \frac{1}{2N} \sum_{i,j} [\mathbf{B}(\mathbf{x}, t)]_{ij} \partial_i \partial_j \varphi(\mathbf{x}) \right\} dt \\ & + \frac{1}{\sqrt{N}} \sum_{i,j} \mathbf{G}_{ij}(\mathbf{x}, t) \partial_i \varphi(\mathbf{x}) d\mathbf{W}_j(t). \end{aligned} \quad (48)$$

Note that our notation $\varphi(\mathbf{x})$ really means $\varphi(\mathbf{y}(\mathbf{x}))$ since φ is defined by (47) which uses the coordinates $\mathbf{y}$ defined in (46). The advection vector $\mathbf{A}$ and the matrix $\mathbf{B} \equiv \mathbf{G}\mathbf{G}^T$ are those given previously in (44) and (45).

Although analytic expressions for φ and its first and second derivatives in phase space, i.e. $\partial_a \varphi$, $\partial_b \varphi$, $\partial_a^2 \varphi$, $\partial_a \partial_b \varphi$, and $\partial_b^2 \varphi$ are unknown, they can be estimated numerically. Since we are interested in the behaviour of perturbations near Γ , we have

estimated these derivatives at points on Γ and then used them to define two functions of φ :

$$\begin{aligned}\omega_\infty(\varphi(\mathbf{x})) &:= \sum_i \mathbf{A}_i(\mathbf{x}, t) \partial_i \varphi(\mathbf{x}), \quad \text{and} \\ \omega_1(\varphi(\mathbf{x})) &:= -\frac{\mu}{2} \sum_{i,j} [\mathbf{B}(\mathbf{x}, t)]_{ij} \partial_i \partial_j \varphi(\mathbf{x}).\end{aligned}\tag{49}$$

Note that we have included in the definition of ω_1 a prefactor of μ . This is necessary since the stochastic slowdown effect is stronger for smaller μ , as indicated by the previous stochastic analysis, and shown numerically in the left panel of figure 10. We can conclude that, curiously, whilst the noise itself has negligible strength (order $1/N$) in region II, the noise-induced slowdown has an order one impact on the dynamics. Specifically, the leading order form of equation (48) in large N and small μ is simply

$$\frac{d}{dt}\varphi = \omega_\infty(\varphi) - \frac{1}{N\mu}\omega_1(\varphi),\tag{50}$$

which enables us investigate the contributions to the period of the limit cycle of the advection term $\mathbf{A}$ and the fluctuation-related contribution $\mathbf{B}$, separately.

Figures 9 and 10 show numerically-computed approximations to the two functions $\omega_\infty(\varphi)$ and $\omega_1(\varphi)$, respectively. We observe that ω_∞ remains positive as φ increases,

Figure 9: Numerically-computed function $\omega_\infty(\varphi)$ for five different values of μ from 5×10^{-5} up to 8×10^{-4} .

with a maximum near, but not exactly at, $\varphi = 0$. This shows that points on Γ move in the direction of increasing φ ; there is no reason for any symmetry about $\varphi = 0$ due to the cyclic nature of the dynamics. The shape of the curve $\omega_\infty(\varphi)$ varies little with μ except near the equilibrium points in the corners, near $\varphi = \pm\pi/3$. Figure 10 shows the variation in $\omega_1(\varphi)$ as μ increases. Note that the definition of ω_1 in (49) contains a factor of μ : including this factor of μ yields curves for $\omega_1(\varphi)$ that are extremely close to each other even as μ increases by a factor of 16. Note also that the curves tend to zero at both ends, indicating that the changes in μ do not change the contribution from $\omega_1(\varphi)$ to the time spent near each equilibrium point. In the first part of the plot, $\omega_1(\varphi)$ is negative,

Figure 10: Left: Numerically-computed functions $\sum_{i,j} [\mathbf{B}(\mathbf{x}, t)]_{ij} \partial_i \partial_j \varphi$ for five different values of μ . Right: Numerically-computed function $\omega_1(\varphi)$, the data collapse visible here justifies the inclusion of the explicit factor of μ in the definition (49). Note that ω_1 appears to cross through zero close to $\varphi = \pm\pi/3$.

indicating that the angular velocity here overall, from (50), is increased, so that the period of the orbit would be decreased. But this effect is more than compensated for by the behaviour in the second part of the plot where $\omega_1(\varphi)$ becomes more strongly positive, leading to a decrease in the overall angular velocity here, and hence a larger increase in the period of the orbit. The combined effect is therefore to increase the period of the orbit, and this increase is driven directly by the fluctuations described by the matrix $\mathbf{B}$ in the original SDE.

Equation (48) can be used to quantify these contributions to the period of the orbit since the expected period of the limit cycle can be computed from the SDE (48) by writing

$$T_{\text{SDE}} = 3 \int_{-\pi/3}^{\pi/3} \frac{d\varphi}{\dot{\varphi}} = 3 \int_{-\pi/3}^{\pi/3} \frac{d\varphi}{\omega_{\infty}(\varphi) - \frac{1}{N\mu} \omega_1(\varphi)}, \quad (51)$$

which, considering the case $N \gg 1$, can be approximated by

$$T_{\text{SDE}} = 3 \int_{-\pi/3}^{\pi/3} \frac{d\varphi}{\omega_{\infty}(\varphi)} + \frac{3}{N\mu} \int_{-\pi/3}^{\pi/3} \frac{\omega_1(\varphi)}{\omega_{\infty}^2(\varphi)} d\varphi + O\left(\frac{1}{N^2}\right).$$

The first integral on the right hand side corresponds to the ODE approximation $T_{\text{ODE}} \approx 3 \ln(1/\mu)$, valid for larger μ ; we infer that $\omega_{\infty}(\varphi)$ depends (relatively weakly) on μ . The second integral contains the leading-order N -dependent behaviour and shows that this is a contribution to the period that scales as $T \propto 1/(N\mu)$. Although ω_{∞} depends on μ , as observed above, this dependence is visible most obviously near to the equilibrium points at which points $\omega_1(\varphi)$ is close to zero. Overall we might therefore expect the μ -dependence of the second integral to be even weaker than the $\ln \mu$ dependence of the $\omega_{\infty}(\varphi)$ function itself. If this were the case, we would be left with just the $T \propto 1/(N\mu)$ dependence that would come to dominate the expression for T_{SDE} for small μ .

In this way we observe that the SDE is able to capture the effects of fluctuations near the limit cycle and yields an expression for the period of oscillations that indicates that the period increases as μ decreases, leading into region III-type scaling at small μ . Moreover, at larger μ (and larger N at fixed μ) we observe that the period of oscillations is given by the deterministic expression computed for region I. Hence this analysis of region II is able to capture the cross-over in scalings for the periodic orbit as we move between regions I and III.

6. Discussion

We have presented a detailed analysis of the simple Rock–Paper–Scissors game, played in a well-mixed population, where the effects of cyclic competition and mutation lead to oscillatory dynamics. In the mean-field ODE model where the strategy mix in the population is described by replicator equations, there is a stable limit cycle, produced in a Hopf bifurcation, when the mutation rate μ is small enough compared to the asymmetry β in the pay-off matrix. For large finite population sizes N , the dynamics of the finite population, as given by stochastic ‘chemical reactions’ between the species, closely follow the ODE dynamics. Here we investigated a specific chemical reaction scheme that was chosen to be straightforward to implement in the parameter regime we are interested in. It should be noted that this choice is not unique and may well not be optimal; for example, it does not allow for the possibility of $\beta < 0$.

Stochastic simulations in the regime of very small mutation rates shows qualitatively, and quantitatively, different dynamics for the oscillations, with the system remaining at corners for most of the time. Mutations are rare, but an essential part of the dynamics. We presented a detailed description of the dynamics in this regime using a Markov chain model for the dynamics on the boundary of phase space, and then arguing that in this regime the assumption that the dynamics took place on the boundary was valid, and led to a self-consistent picture.

Our third approach to the problem was to construct a stochastic differential equation, which, according to a theorem of Kurtz [23] approximates the individual-based dynamics for large but finite populations. This was successful in capturing the cross-over between the mean-field and fully stochastic regimes as the mutation rate μ decreased for a fixed population size. In general the effects of noise on nonlinear oscillators are complex and topic of significant current research interest, see [22]. Our work here has concentrated on understanding the cross-over between regions I and III. The cross-over occurs when $\mu N \ln N \approx 1$, and the period of the oscillations changes from the mean-field approximation $T_{\text{ODE}} \sim 3 \ln(1/\mu)$ to the stochastic approximation $T_{\text{SSA}} \sim 3/(N\mu)$. Figure 11 summarises these three regions of the behaviour, plotting the period of the oscillations as a function of $\mu N \ln N$, comparing the expected period of stochastic simulations with $N = 2^{17}$ with the theoretical results presented above.

Our main conclusion is that as demographic fluctuations become more important (i.e. at small mutation rates and for small population sizes) the period of the cyclic oscillations increases significantly above that which would be predicted on the basis of the mean-field ODE model.

To fully assess the importance of this effect in real populations will require further work. The presence of three-body interactions in our chemical reaction scheme could be viewed as less biologically realistic, and some theorists would prefer a stochastic description in terms of a birth-death process with fitness and weak selection. Similarly, we considered here a model with fixed population size, yet, in the wild, population size is of course variable. One promising route to addressing these modelling concerns has recently been presented in [31], providing a robust mapping from Lotka-Volterra population models with two-body interactions and varying populations into their stochastic dynamics in frequency space. We would argue, however, that the salient features of the stochastic slowdown we observed should be common to all reasonable model specifications: in region I the dynamics are dominated by waiting for rare mutation events, and in regions II and III the format of the noise is subdominant to the deterministic features

Figure 11: A summary of the three regions for the dynamics, comparing the theoretical predictions for the mean period T of oscillations each region with the results of stochastic simulations for $N = 2^{17}$.

of the limit cycle. A more exciting direction for future work is to explore the role of the spatial distribution of populations, for example, one might ask if stochastic slowdown also affects spiral waves (see e.g. [32, 33]).

Our very detailed work here should also provide a basis for investigation of many similar, more complex, kinds of cyclic interaction. We also intend to examine situations where the mean-field behaviours are known to produce more complicated dynamics, for example through the occurrence of heteroclinic networks, or other bifurcations from the heteroclinic cycles involved. This work then will provide a starting point to guide the investigation of these situations when stochastic effects due to large but finite populations of individuals become important.

Acknowledgement

QY is supported by the University of Bath through a 50th Anniversary Excellence Award. TR and JHPD acknowledge support from the Royal Society.

References

- [1] R. May, W. Leonard, Nonlinear Aspects of Competition Between Three Species, SIAM Journal on Applied Mathematics 29 (2) (1975) 243–253. doi:10.1137/0129022. URL <http://epubs.siam.org/doi/abs/10.1137/0129022>
- [2] J. Guckenheimer, P. Holmes, Structurally stable heteroclinic cycles, Mathematical Proceedings of the Cambridge Philosophical Society 103 (1988) 189–192.
- [3] M. Krupa, Robust heteroclinic cycles, Journal of Nonlinear Science 7 (1997) 129–176.

- [4] N. F. Britton, *Essential Mathematical Biology*, Springer, 2005.
- [5] M. A. Nowak, *Evolutionary Dynamics*, Harvard University Press, 2006.
- [6] D. F. P. Toupo, S. H. Strogatz, Nonlinear dynamics of the rock-paper-scissors game with mutations, *Physical Review E* 91 (5) (2015) 052907. doi:10.1103/PhysRevE.91.052907.
URL <http://link.aps.org/doi/10.1103/PhysRevE.91.052907>
- [7] T. Reichenbach, M. Mobilia, E. Frey, Coexistence versus extinction in the stochastic cyclic Lotka-Volterra, *Physical Review E* 74 (5). doi:10.1103/PhysRevE.74.051907.
URL <http://arxiv.org/abs/q-bio/0605042>
- [8] B. Sinervo, C. M. Lively, The rock-paper-scissors game and the evolution of alternative male strategies *Nature* 380 (6571) (1996) 240–243. doi:10.1038/380240a0.
URL <http://www.nature.com/nature/journal/v380/n6571/abs/380240a0.html>
- [9] B. Kerr, M. A. Riley, M. W. Feldman, B. J. M. Bohannan, Local dispersal promotes biodiversity in a real-life game of rock-paper-scissors, *Nature* 418 (6894) (2002) 171–174. doi:10.1038/nature00823.
URL <http://www.nature.com/nature/journal/v418/n6894/full/nature00823.html>
- [10] B. C. Kirkup, M. A. Riley, Antibiotic-mediated antagonism leads to a bacterial game of rock-paper-scissors in vivo, *Nature* 428 (6981) (2004) 412–414. doi:10.1038/nature02429.
- [11] M. F. Weber, G. Poxleitner, E. Hebisch, E. Frey, M. Opitz, Chemical warfare and survival strategies in bacterial range expansions, *Journal of The Royal Society Interface* 11 (96) (2014) 20140172. doi:10.1098/rsif.2014.0172.
URL <http://rsif.royalsocietypublishing.org/content/11/96/20140172>
- [12] M. Rabinovich, A. Volkvskii, P. Lecanda, R. Huerta, H. Abarbanel, G. Laurent, Dynamical encoding by networks of competing neuron groups: winnerless competition, *Physical Review Letters* 87 (2001) 068102.
- [13] T. Tsai, J. Dawes, Dynamics near a periodically-perturbed robust heteroclinic cycle, *Physica D* 262 (2013) 14–34.
- [14] P. D. Taylor, L. B. Jonker, Evolutionary stable strategies and game dynamics, *Mathematical biosciences* 40 (1-2) (1978) 145–156.
- [15] P. Schuster, K. Sigmund, Replicator dynamics, *Journal of theoretical biology* 100 (3) (1983) 533–538.
- [16] J. Hofbauer, K. Sigmund, *Evolutionary Games and Population Dynamics*, Cambridge University Press, 1988.
- [17] J. Hofbauer, *Heteroclinic cycles in ecological differential equations*, Tatra Mountains Mathematical Publications 4.
- [18] M. Mobilia, Oscillatory dynamics in rock-paper-scissors games with mutations, *Journal of Theoretical Biology* 264 (1) (2010) 1–10. doi:10.1016/j.jtbi.2010.01.008.
URL <http://www.sciencedirect.com/science/article/pii/S0022519310000123>

- [19] A. J. McKane, T. J. Newman, Predator-prey cycles from resonant amplification of demographic stochasticity, *Physical review letters* 94 (21) (2005) 218102.
- [20] R. P. Boland, T. Galla, A. J. McKane, How limit cycles and quasi-cycles are related in systems with intrinsic noise, *Journal of Statistical Mechanics: Theory and Experiment* 2008 (09) (2008) P09001.
- [21] R. P. Boland, T. Galla, A. J. McKane, Limit cycles, complex floquet multipliers, and intrinsic noise, *Physical Review E* 79 (5) (2009) 051131.
- [22] J. M. Newby, M. A. Schwemmer, Effects of moderate noise on a limit cycle oscillator: counter rotation and
Physical Review Letters 112 (11), arXiv: 1310.2227.
doi:10.1103/PhysRevLett.112.114101.
URL <http://arxiv.org/abs/1310.2227>
- [23] T. G. Kurtz, Strong approximation theorems for density dependent Markov chains, *Stochastic Processes and their Applications* 6 (3) (1978) 223–240.
doi:10.1016/0304-4149(78)90020-0.
URL <http://www.sciencedirect.com/science/article/pii/0304414978900200>
- [24] D. T. Gillespie, Exact stochastic simulation of coupled chemical reactions, *The journal of physical chemistry* 81 (25) (1977) 2340–2361.
- [25] J. Guckenheimer, P. Holmes, *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields*, Second edition. Springer, New York, 1986.
- [26] J. R. Norris, *Markov Chains*, Cambridge University Press, 1998.
- [27] C. W. Gardiner, *Stochastic methods: a handbook for the natural and social sciences*, Springer, Berlin, 2009.
- [28] T. L. Parsons, C. Quince, Fixation in haploid populations exhibiting density dependence II: The quasi-neutral case, *Theor. Popul. Biol.* 72 (4) (2007) 468–479.
- [29] G. W. Constable, T. Rogers, A. J. McKane, C. E. Tarnita, Demographic noise can reverse the direction of deterministic selection, *Proceedings of the National Academy of Sciences* (2016) 201603693.
- [30] T. L. Parsons, T. Rogers, Dimension reduction via timescale separation in stochastic dynamical systems, arXiv preprint arXiv:1510.07031.
- [31] G. W. Constable, A. J. McKane, A mapping of the stochastic Lotka-Volterra model to models of population genetics and game theory, arXiv preprint arXiv:1704.07565.
- [32] B. Szczesny, M. Mobilia, A. M. Rucklidge, When does cyclic dominance lead to stable spiral waves?, *EPL (Europhysics Letters)* 102 (2) (2013) 28012.
- [33] B. Szczesny, M. Mobilia, A. M. Rucklidge, Characterization of spiraling patterns in spatial rock-paper-scissors games, *Physical Review E* 90 (3) (2014) 032704.