

The prevalence of chaotic dynamics in games with many players

James BT Sanders,^{1,*} Tobias Galla,^{1,†} and J. Doyne Farmer^{2,3,4,‡}

¹*Theoretical Physics, School of Physics and Astronomy,
The University of Manchester, Manchester M13 9PL, UK*

²*Institute for New Economic Thinking at the Oxford Martin School, University of Oxford, Oxford OX2 6ED, UK*

³*Mathematical Institute, University of Oxford, Oxford OX1 3LP, UK*

⁴*Santa Fe Institute, Santa Fe, NM 87501, UK*

We study adaptive learning in a typical p -player game. The payoffs of the games are randomly generated and then held fixed. The strategies of the players evolve through time as the players learn. The trajectories in the strategy space display a range of qualitatively different behaviors, with attractors that include unique fixed points, multiple fixed points, limit cycles and chaos. In the limit where the game is complicated, in the sense that the players can take many possible actions, we use a generating-functional approach to establish the parameter range in which learning dynamics converge to a stable fixed point. The size of this region goes to zero as the number of players goes to infinity, suggesting that complex non-equilibrium behavior, exemplified by chaos, may be the norm for complicated games with many players.

PACS numbers: 02.50.Le, 05.45.Jn, 89.65.Gh

I. INTRODUCTION

Many branches of science are interested in systems made up of a large number of competing individuals. Examples range from financial markets and other social systems, to populations undergoing biological evolution, to networked computer systems. In many such situations individuals compete for limited resources, and the natural model is a game, which consists of a set of players who at any point in time choose from a set of possible actions in an attempt to maximize their payoff. Game theory has received a great deal of attention since the mid-20th century [1], but research has overwhelmingly focused on simple games, with only a very small number of players and pure strategies, even though real game-like systems often involve large numbers of individuals and possible strategies. While many of the observed properties of simple games carry over directly to more complicated ones, it is becoming increasingly clear that complicated games can show important types of behavior not found in simpler systems.

Early work in game theory focused on the concept of equilibria, in particular the famous *Nash equilibrium* [2], in which the players adopt strategies such that no player can improve her own payoff by unilaterally changing her own strategy. The strategies at a Nash equilibrium can be probabilistic combinations of pure strategies, called *mixed strategies*. Nash's ideas have been particularly influential in economics, where agents are often assumed to adopt Nash equilibria. It should be emphasized that this is a behavioral assumption, and that the empirical evidence for this is mixed [3].

Equilibrium models are perfectly plausible in the context of simple games when there is a unique Nash equilibrium that is easy to calculate. In complicated games, there are typically numerous distinct Nash equilibria [4, 5], and locating even one of them can be a laborious task: the computing time of the best known algorithms increases exponentially with the size of the game [6]. This seems to cast doubt on whether it is reasonable to assume that players of complicated games naturally discover equilibria. But if they don't play equilibria, what do they do instead? And is there a way to determine *a priori* whether players will converge to a unique Nash equilibrium using a reasonable learning strategy?

Opper and Diederich studied replicator dynamics, a standard model of biological evolution, in the context of complicated games. They found that the dynamics converges to a unique fixed point in some regions of parameter space, but that in other regions, the dynamics does not settle down [7, 8]. Sato and co-workers showed that adaptation learning can result in chaotic dynamics even in low-dimensional games [9–11]. Building on this earlier work, Galla and Farmer studied complicated two-player games in which the strategies are randomly generated but fixed in time, assuming that players use an experience-weighted attraction dynamics to learn their strategies [12]. They found that there are distinct regions of the parameter space with different behaviors. When the timescale for learning is short and the payoffs of the players are strongly negatively correlated, they observed convergence to unique fixed points. But when the time scale for learning is long and when the payoffs are less negatively correlated they observed limit cycles and chaos. And when the timescale for learning is long and the payoffs are positively correlated they observed a large multiplicity of stable fixed points.

In this paper we extend the work of [12] to p -player games and find that, for large numbers of players, complex dynamics is not merely frequent but ubiquitous. The region

* james.sanders-2@postgrad.manchester.ac.uk

† tobias.galla@manchester.ac.uk

‡ doyne.farmer@oxfordmartin.ox.ac.uk

of parameter space in which the players' strategies consistently converge to a unique fixed point appears to vanish as $p \rightarrow \infty$. **This suggests that complex non-equilibrium behavior may be the norm in dynamics on complicated games with many players, at least under the type of learning we study here.**

Our goal can be understood through an analogy to fluid flow. It is well known that fluid flow can be characterized *a priori* in terms of a few key parameters that can be estimated on the back of an envelope. The most famous of these is the Reynolds number, which is a non-dimensional ratio of stress and viscosity. Though the precise transition point depends on other parameters, if the Reynolds number is high then the flow is likely to be turbulent. Our goal here is similar: We seek parameters that can characterize *a priori* whether or not a game will exhibit complex dynamics in the strategy space as the players learn. Here we are particularly interested in what happens as the number of players increases. Since the presence of many players makes the game more complex, we hypothesize that it will tend to make the strategy dynamics more complex as well. (And indeed this is what we observe).

The remainder of the paper is organized as follows: In Sec. II we introduce the experience-weighted attraction learning algorithm and we define what we mean by a complex p -player game. Sec. III then contains an overview of the different types of behavior seen in the learning of such games, along with a more quantitative analysis based on numerical simulations. In Sec. IV we then turn to a semi-analytical study of p -player learning based on a generating-functional approach. This technique allows us to derive estimates for the boundaries of stable and complex behavior in parameter space for games with a large number of strategies. We then turn to a brief discussion of volatility clustering and the relevance of our result for the modeling of financial markets in Sec. VI, before we summarize our results. The Appendix contains further details of the numerical methods used to identify the different types of dynamical behavior, and of the generating functional analysis. It also contains some additional numerical results.

II. MODEL

II.1. Experience-weighted attraction

Suppose that a set of players repeatedly play a game in which they each choose from a distinct set of strategies, without conferring with each other. The players have good memories and a full understanding of the payoffs that a given combination of strategies would yield for them, and are only interested in maximizing their own payoffs. We are interested in the case where the players learn their strategies based on past experience. A common approach assumes the players adapt behavior over time based on the past success of each possible strategy

[13–16]. The basic idea is that the players calculate a numerical score, known as an ‘attraction’ or ‘propensity’, for each possible strategy, describing how successful they expect it to be. They then select a strategy based on the relative score of each possible action, play the game one or more times, then use the outcome to update their attractions for future play. This defines an adaptation process, in which agents learn from past experience and continuously try to improve their actions.

Two types of simple learning models have proved especially popular over the years. In reinforcement learning, the players calculate the attraction of a strategy based on how successful it has been when they have employed it in the past. In belief-based learning, the attractions are determined according to how successful the possible strategies would have been, had they been used in prior iterations. The experience-weighted attraction (EWA) system, introduced by Camerer and Ho [13], combines these two approaches into a single algorithm—in fact, belief-based learning can be seen as a deterministic limit of reinforcement learning in which the players sample all pure strategies at each time step. Combined with a logit model of how the players choose strategies based on their attractions, EWA is observed to be a reasonably good match for how people learn to play simple games [14, 15]. We are interested in how often EWA converges to equilibria, so we select the deterministic (belief-based) version of the model. While the noisy (reinforcement) version may well perform stochastic oscillations about equilibria in the right conditions, we would expect that in general, the introduction of noise would lead to complex dynamics being observed even more often. Both cases were studied by Galla and Farmer and the differences were not dramatic (see the Supplementary Material of [12]).

Consider a p -player game, where each player has N actions to choose from. The rewards for the players are defined by the generalised payoff matrix $\Pi_{i, i_{\mu+1}, \dots, i_{\mu-1}}^{\mu}$, which represents the payoff to player μ if they play action i , while the other players $\mu + 1, \dots, \mu - 1$ play actions $i_{\mu+1}, \dots, i_{\mu-1}$, respectively (where the subscripts labelling the players are to be interpreted modulo p).

We use updates rules similar to those of [12], but adapted to the multi-player game,

$$\begin{aligned} x_i^{\mu}(t+1) &= \frac{\exp[\beta Q_i^{\mu}(t)]}{\sum_k \exp[\beta Q_k^{\mu}(t)]}, \\ Q_i^{\mu}(t+1) &= (1 - \alpha)Q_i^{\mu}(t) + \sum_{i_{\mu+1}, \dots, i_{\mu-1}} \Pi_{i, i_{\mu+1}, \dots, i_{\mu-1}}^{\mu} \prod_{\kappa \neq \mu} x_{i_{\kappa}}^{\kappa}(t). \end{aligned} \quad (1)$$

Here \mathbf{x} represents the players' strategies, with $x_i^{\mu}(t)$ denoting the probability that player μ will choose action i at time t . The value $Q_i^{\mu}(t)$ is player μ 's *attraction* to action i at time t . The two parameters of the system are the memory loss rate α , which lies in the interval $[0, 1]$, and the *intensity of choice* β , which is non-negative. A player's attraction to an action is essentially a geometri-

cally discounted average of the payoffs that would have been achieved by playing that action in earlier time steps, with a discount factor determined by α . The intensity of choice β determines the bias with which players choose actions with higher attractions—if $\beta = 0$, then the players ignore the attractions and choose each action with equal probability, while in the limit as $\beta \rightarrow \infty$, the players each choose their most attractive action with probability 1. The intensity of choice therefore plays a similar role to the inverse temperature in thermodynamics. Our system is identical to that studied by Galla and Farmer in reference [12], except that they restricted the number of players to be $p = 2$. Our contribution is to understand how this changes as p becomes larger. The system of Camerer and Ho [14, 15], as well as allowing for noise, allows the intensity of choice β to vary over time. In the present work we assume that it has attained its long-term value. Thus the dynamics we study here are a special case of [14, 15].

The attractions Q_i^μ can be eliminated from the update rules in Eq. (1) to yield

$$x_i^\mu(t+1) = \frac{1}{Z^\mu(t)} x_i^\mu(t)^{1-\alpha} \times \exp \left(\beta \sum_{i_{\mu+1}, \dots, i_{\mu-1}} \Pi_{i, i_{\mu+1}, \dots, i_{\mu-1}}^\mu \prod_{\kappa \neq \mu} x_{i_\kappa}^\kappa(t) \right), \quad (2)$$

where

$$Z^\mu(t) = \sum_k x_k^\mu(t)^{1-\alpha} \exp \left(\beta \sum_{i_{\mu+1}, \dots, i_{\mu-1}} \Pi_{k, i_{\mu+1}, \dots, i_{\mu-1}}^\mu \prod_{\kappa \neq \mu} x_{i_\kappa}^\kappa(t) \right) \quad (3)$$

is a normalization factor.

Following [12], we focus on the continuous-time limit of the EWA system. Letting

$$r = \beta/\alpha,$$

this limit is found by keeping r fixed while taking $\alpha \rightarrow 0$ and $\beta \rightarrow 0$ simultaneously, and rescaling time. Both the geometric discounting of the past profits of an action and the logit selection remain significant in this limit. This yields the so-called Sato-Crutchfield dynamics [9, 10]

$$\frac{\dot{x}_i^\mu(t)}{x_i^\mu(t)} = -\frac{1}{r} \ln x_i^\mu(t) + \sum_{i_{\mu+1}, \dots, i_{\mu-1}} \Pi_{i, i_{\mu+1}, \dots, i_{\mu-1}}^\mu \prod_{\kappa \neq \mu} x_{i_\kappa}^\kappa(t) - \rho^\mu(t), \quad (4)$$

where $\rho^\mu(t) = \ln Z^\mu(t)/\beta$. A detailed derivation can be found for example in the Supplementary Material of [12]. Note that, since each player satisfies the constraint that the probability of taking any given action sums to one, the resulting dynamical system for the learning dynamics is of dimension $(N-1)p$.

We assume throughout this work that α and β are small enough that the continuous limit is a good approxima-

tion. We take this limit mainly for analytical convenience; the continuous limit is easier to study, as it has only one relevant parameter, the ratio of α to β .

When considering large values of N it is convenient to rescale the probabilities $x_i^\mu(t)$ so that they sum to N for each player. This means that $x_i^\mu(t) = \mathcal{O}(1)$, and so the objects inside the exponentials in Eq. (2) remain finite as $N \rightarrow \infty$. This can be achieved by modifying the definition of the normalization factors Z^μ accordingly. The payoff matrix entries are rescaled as well as explained below.

II.2. Constructing typical complicated games with p players

We are interested in the behavior of the EWA system for generic complicated games. To that end, we draw the payoff matrix Π from a multivariate normal distribution with

$$\left\langle \Pi_{i_\mu, i_{\mu+1}, \dots, i_{\mu-1}}^\mu \Pi_{i_\nu, i_{\nu+1}, \dots, i_{\nu-1}}^\nu \right\rangle = \begin{cases} \frac{1}{N^{p-1}} & \mu = \nu \\ \frac{\Gamma}{(p-1)N^{p-1}} & \mu \neq \nu \end{cases} \quad (5)$$

and all other correlations zero. The multivariate Gaussian distribution is a natural choice because it is the maximum entropy distribution when there are constraints on the first and second moments. We have chosen the construction above because it yields the only possible multivariate Gaussian distribution that satisfies the following properties: (i) the distribution is symmetric with respect to the different players and pure strategies (that is, swapping any two players or pure strategies would leave the distribution unchanged) and (ii) the payoffs for any two distinct choices of actions are uncorrelated (by choice of actions we mean the tuple (i_1, \dots, i_p) representing the pure strategies played by the p players). As before, we wish to emphasize that once Π is chosen it remains fixed through the duration of the iterated game.

The parameter Γ can be seen as the level of ‘zero-sumness’ of the game, and must lie on the interval $[-1, p-1]$. When $\Gamma = p-1$, any outcome of the game leads to each player receiving the same payoff with probability 1. When $\Gamma = 0$, all payoffs are uncorrelated. When $\Gamma = -1$, for any given outcome, the players’ payoffs sum to zero with probability 1. Therefore Γ can be seen as a ‘competition parameter’—at large positive values, the players have common goals, while at large negative values, they are working against each other.

When $p = 2$, the possible values of Γ span the range $-1 \leq \Gamma \leq 1$. However when $p > 2$ the situation becomes more complicated. Eq. (5) indicates that the payoffs each have variance $1/N^{p-1}$, the correlations between payoffs for different outcomes are uncorrelated, and the correlation between the payoffs of two different players for any given outcome is $\Gamma/((p-1)N^{p-1})$. The scaling of the payoff matrix elements with $N^{-(p-1)}$ ensures that the objects inside the exponentials in Eq. (2) remain finite

as $N \rightarrow \infty$, so that both the memory-loss (α) and payoff (β) factors remain significant in this limit.

II.3. Strategy for exploring the parameter space

For a p -player game the payoff “matrices” each have p indices, and so are not two dimensional matrices in the normal usage of the world, but are p -dimensional. This significantly complicates the problem of exploring the parameter space numerically. For a game in which a single player can take one of N actions the payoff matrix for a single player has N^p components; for p players there are pN^p components. For $p = 10$ and $N = 10$, for example, this means that 10^{11} (a hundred billion) random numbers must be generated in order to construct the game. The sheer amount of memory needed for simulation creates a serious bottleneck.

Given the numerical constraints this forces us to rely more heavily on the analytic calculation than Galla and Farmer did when they explored two-player games. We use numerical simulations to get a feeling for the behavior of the system, with relatively small values of N and p . In parallel we perform an analytic calculation of the stability boundary between the region where there is unique convergence to a fixed point and the rest of the parameter space. We then compare the analytic and numerical simulations and demonstrate that the analytic calculation seems to be reasonably accurate, given the magnitude of the finite-size effects. Finally we use the analytic calculation to assess the behavior in the limit where $N \rightarrow \infty$ and p is large.

III. NUMERICAL EXPLORATION OF THE PARAMETER SPACE

III.1. Overview

In this section we give an overview of our numerical exploration of the parameter space. As observed by Galla and Farmer, the strategies of the players can converge to any of the possible types of attractors, including fixed points, limit cycles and chaos. In some regimes a given game may have multiple fixed points, i.e. multiple basins of attraction, but we have not observed this when the attractors are limit cycles or chaos. In Fig. 1 we show some examples. The chaotic behavior can be low dimensional, as shown in the middle row, or high dimensional, as shown in the last row. For high-dimensional chaos a given action typically has epochs in which it is almost never selected and others in which it is used frequently – the range of variation is striking.

When the system converges to a fixed point, it usually does so rather quickly, as shown in Fig. 2(a). However there are sometimes long metastable chaotic-looking transients that suddenly collapse to a fixed point, as shown in Fig. 2(b). This is particularly likely for small

values of α/β and small positive Γ (i.e., for weakly positively correlated payoffs and players with long memories). Although the boundaries can be a bit fuzzy, the parameter space divides into distinct regions. These are illustrated in the schematic diagram in Fig. 3. We briefly describe the main features:

Positively correlated payoffs: For $\Gamma > 0$ and large α/β the dynamics converge to a unique fixed point. Holding Γ constant, for small α/β the dynamics converges to one of multiple fixed points, see also Fig. A3.

Negatively correlated payoffs: For $\Gamma < 0$ and large α/β the dynamics converge to a unique fixed point (just as for $\Gamma > 0$). Holding Γ constant, for small α/β the dynamics are chaotic. This corresponds to longer memory. As Γ increases from $\Gamma = -1$ to $\Gamma = 0$ the size of the chaotic region increases. The stability boundary dividing the region with complex dynamics from the unique equilibrium shifts to the right (i.e. toward a larger value of α/β) as the number of players increases.

Uncorrelated payoffs: Limit cycles are common near the boundaries, particularly near the boundary where $\Gamma \approx 0$, see also Figs. A1 and A2 in the Appendix. In general the behaviors reported here are not strict, in the sense that generating different payoff matrices with the same values of α/β can result in different behaviors, particularly near the boundaries. We conjecture that these are finite size effects, so that the boundaries would become distinct and the behavior at a given set of parameters would become crisp in the limit as $N \rightarrow \infty$.

Prevalence of chaotic dynamics: The key result is that as the number of players increases the size of the region with complex dynamics grows. If $\Gamma > 0$, this means that the region with multiple fixed points becomes larger; if $\Gamma < 0$ this means that the region with chaotic dynamics becomes larger. More players make the system less likely to converge to a unique equilibrium. This is particularly important for zero sum games ($\Gamma = -1$); in this case for $p = 2$ chaos is only observed in the limit as $\alpha/\beta \rightarrow 0$, whereas for $p > 2$ chaos is observed over a finite interval in α/β .

As already mentioned, it is not possible to perform numerical experiments for large values of both p and N . We will present some numerical evidence for these results, but they make more sense when guided by the theory.

IV. GENERATING FUNCTIONAL APPROACH

We now turn to treat the learning dynamics analytically. In the language of the theory of disordered system the random payoff matrix in our problem represents quenched disorder. Techniques from spin glass physics can be applied to study the thermodynamic limit (i.e. the limit of large payoff matrices, $N \rightarrow \infty$). We use the Martin-Siggia-Rose generating functional to derive an effective dynamics. For a recent review of these methods

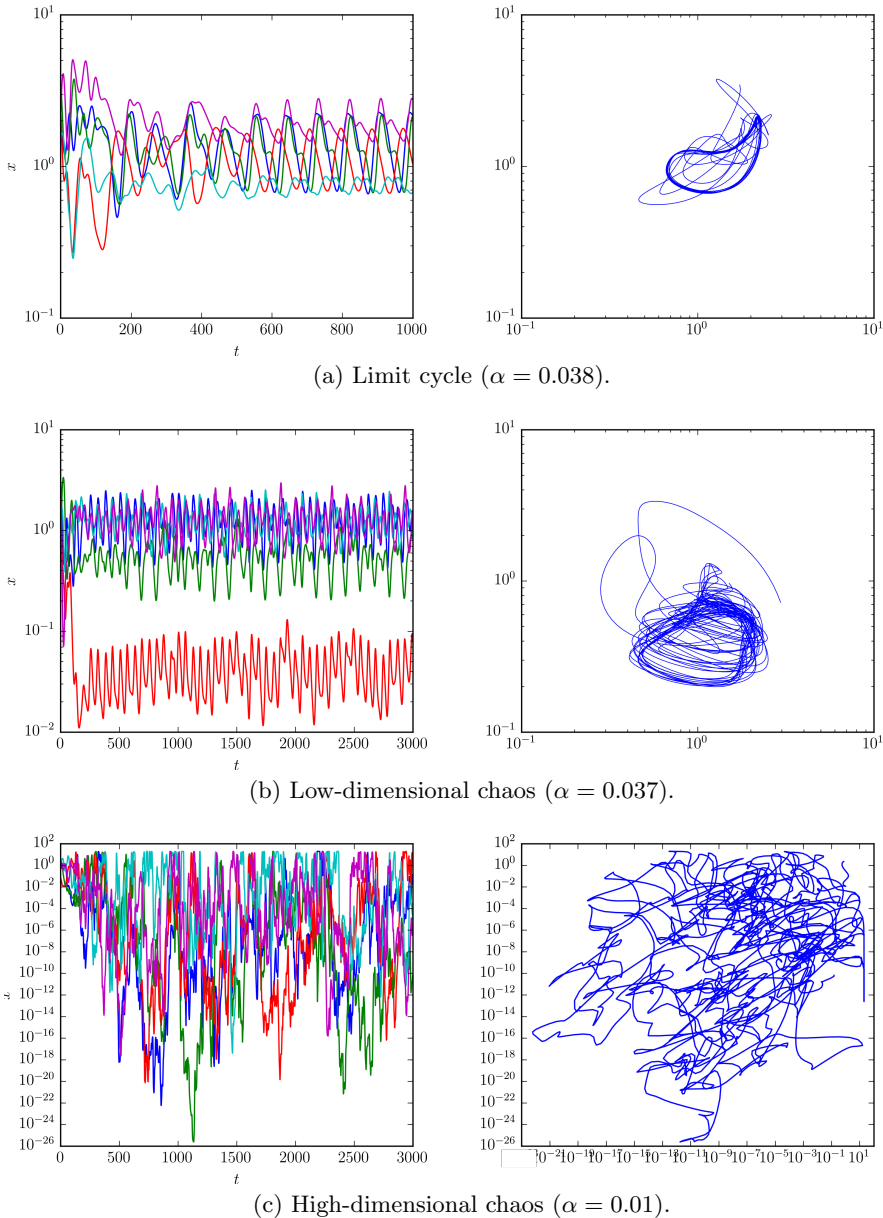


FIG. 1. Time series and phase plots showing complex dynamics under EWA learning, including (a) limit cycle, (b) low dimensional chaos, and (c) high dimensional chaos. The game has three players ($p = 3$) and $N = 20$ possible actions, with $\beta = 0.05$ and $\Gamma = -0.5$. The time series plots on the left show the probability x_i^μ for player μ to use action i as a function of time for five different actions, and the phase plots on the right shows the probability for two of the actions as a function of each other. Case (a) illustrates that limit cycles can have complicated geometric forms and long periods. For smaller values of α/β and negative Γ , chaos is very common, ranging from low-dimensional chaotic attractors as shown in (b) to high dimensional attractors as shown in (c). Note that for high dimensional chaos the probability that a given action is used at different points in time can vary by as much as a factor of 10^{20} .

see [17].

Broadly speaking the calculation proceeds as follows: in a first step the average over the disorder (the random payoff matrices) is carried out and the effective dynamics is derived. This process is subject to coloured noise, and reflects the statistics over all games within the Gaussian ensemble. In a second step, a fixed point of this dynam-

ics is assumed. We investigate the linear stability of this fixed point, and calculate the boundary to the phase in parameter space where more complex dynamics are seen. Thus we cannot calculate where the dynamics follow limit cycles or chaos, but we can calculate the boundary between the unique stable fixed point and other behaviors. As discussed later, we can only do this for $\Gamma < 0$. We note




FIG. 2. Trajectories for EWA system leading to a fixed point in a three-player game. Panel (a) shows an instance in which a fixed point is reached relatively quickly. Panel (b) illustrates a metastable chaotic transient eventually collapsing to a fixed point. In both examples each player has a choice of $N = 20$ possible actions and the intensity of choice is $\beta = 0.05$. A random sample of five of the players' strategy components x_i^μ are plotted. Remaining parameters are $\alpha = 0.1$, $\Gamma = -0.5$ in panel (a), and $\alpha = 0.01$, $\Gamma = 0.1$ in panel (b).

that similar calculations have been carried out for replicator dynamics on two-player and in multi-player games [7, 18], see also [19] for approaches to p -player random games using static replica methods.

IV.1. Effective process

The first step is to set up a generating functional to describe the probability measure of all possible paths of the dynamics. Performing the average over the assignment of payoff matrices an 'effective dynamics' can then be derived. This calculation is based on path integrals, and somewhat lengthy. We do not report it here in the main paper, instead details are relegated to Appendix A1. The outcome of the calculation is a stochastic integro-differential equation for the evolution of the distribution of components $x(t)$ of the players' strategies in

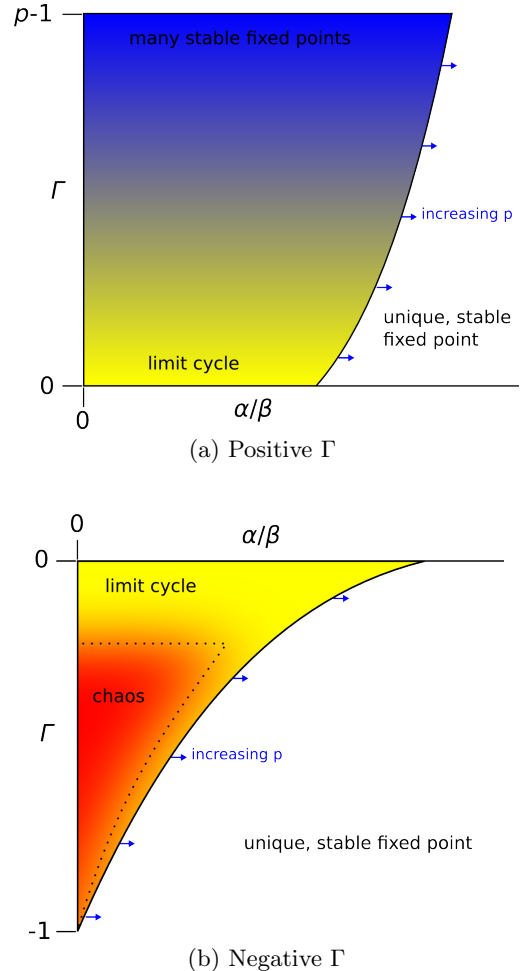


FIG. 3. Schematic phase diagrams describing the observed long-term behavior of the p -player EWA system for large but finite N . In (a) $\Gamma > 0$, meaning players' payoffs are positively correlated. Here we observe a unique stable equilibrium for large α/β and multiple stable equilibria for small α/β . In (b) $\Gamma < 0$, meaning players' payoffs are anti-correlated. Here we once again observe a unique stable equilibrium for large α/β , but we now observe chaos for small α/β . Limit cycles are common near the boundaries, particularly near $\Gamma \approx 0$. We show the case for $p = 2$; as the number of players is increased the stability boundaries move to the right, but little else changes. The heat map is subjective.

the large- N limit. This process is of the form

$$\frac{\dot{x}(t)}{x(t)} = \Gamma \int dt' G(t, t') C(t, t')^{p-2} x(t') - \frac{1}{r} \ln x(t) - \rho(t) + \eta(t), \quad (6)$$

where $\eta(t)$ is a colored Gaussian random variable satisfying $\langle \eta(t)\eta(t') \rangle_* = C(t, t')^{p-1}$ and $\langle \eta(t) \rangle_* = 1$. We use $\langle \cdot \rangle_*$ to denote an average over realizations of the effective dynamics. The dynamical order parameters $C(t, t')$ and $G(t, t')$ are correlation and response functions of the

learning dynamics. They are determined from

$$G(t, t') = \left\langle \frac{\delta x(t)}{\delta \eta(t')} \right\rangle_*, \quad C(t, t') = \langle x(t)x(t') \rangle_*. \quad (7)$$

The effective process in Eq. (6) together with Eqs. (7) define a self-consistent system for C and G . The function $\rho(t)$ in the effective process is a Lagrange multiplier ensuring normalization. It is defined such that $\langle x \rangle_* = 1$. Note that in the derivation of the effective dynamics, it is assumed that each component of each player's strategy is initially drawn from an identical distribution.

IV.2. Fixed point solution

We now focus on the dynamics at large values of α/β . Numerical simulations of the learning dynamics suggest that one unique stable fixed point is found for any one realization of the game in this regime. We therefore make a fixed point ansatz for the effective dynamics. In such a stationary fixed point regime the response function $G(t, t')$ becomes a function of the time difference only, i.e., $G(t, t') = G(t - t')$, while the correlation function tends to a constant, $C(t, t') \equiv q$, see also [7, 12, 18] for further details. Within the fixed-point ansatz the random variable $\eta(t)$ in Eq. (6) tends to a constant value drawn from a Gaussian distribution with zero mean and variance q^{p-1} . Fixed points of the effective dynamics are then found from

$$x^* \left(\Gamma q^{p-2} \chi x^* - \frac{1}{r} \ln x^* + \eta^* - \rho^* \right) = 0, \quad (8)$$

where $q = \langle (x^*)^2 \rangle_*$ and $\chi = \int_0^\infty d\tau G(\tau)$, and x^* , η^* , and ρ^* are the fixed point values of x , η , and ρ , respectively. We can write $\eta^* = q^{(p-1)/2} z$, where z is a standard Gaussian random variable. Then, dropping the stars, we have

$$x(z) \left(\Gamma q^{p-2} \chi x(z) - \frac{1}{r} \ln x(z) + q^{\frac{p-1}{2}} z - \rho \right) = 0, \quad (9)$$

where χ , q , and ρ are to be determined from

$$\begin{aligned} \left\langle \frac{\partial x(z)}{\partial z} \right\rangle_* &= q^{\frac{p-1}{2}} \chi, \\ \langle x(z)^2 \rangle_* &= q, \\ \langle x(z) \rangle_* &= 1. \end{aligned} \quad (10)$$

These relations can be re-written as

$$\begin{aligned} \int_{-\infty}^{\infty} Dz \frac{\partial x(z)}{\partial z} &= q^{\frac{p-1}{2}} \chi, \\ \int_{-\infty}^{\infty} Dz x(z)^2 &= q, \\ \int_{-\infty}^{\infty} Dz x(z) &= 1, \end{aligned} \quad (11)$$

where $Dz = \frac{dz}{\sqrt{2\pi}} e^{-z^2/2}$.

The relation in Eq. (9) can be re-arranged to give an explicit expression for $x(z)$ in terms of the so-called Lambert W function $W(\cdot)$. The value $W(y)$ is defined as the solution of the equation $We^W = y$. Restricting W and y to the real line, the solution exists for $y \geq -1/e$. It is uniquely defined for $y \geq 0$ and double valued for $-1/e < y < 0$. We find

$$x = -\frac{1}{\Gamma r q^{p-2} \chi} W \left(-\Gamma r q^{p-2} \chi e^{r(q^{(p-1)/2} z - \rho)} \right). \quad (12)$$

We note that it is not clear that Eq. (12) has valid solutions for all choices of the model parameters. If these do not exist the fixed point ansatz is invalid, and so we do not expect the dynamics to settle down. There may also be instances in which Eq. 9 has multiple solutions for x for a given value of the standard Gaussian variable z . In principle the distribution of fixed points could be composed of any mixture of these solutions. If the argument of the Lambert function is positive however, there is a unique and well defined solution, $x(z)$, for any value of z . Throughout this discussion it is important to keep in mind that the macroscopic order parameters q , χ and ρ are to be determined self-consistently via Eqs. (11).

IV.3. Stability analysis

By numerically solving the fixed point equations, we see that for a given value of p , stable fixed points exist for large values of $1/r$ but not for small values. We now proceed to determine the boundary of stability. Suppose the effective process in Eq. (6) is perturbed from a fixed point by a small noise term $\xi(t)$. We then have

$$\begin{aligned} \dot{x}(t) = x(t) &\left[\Gamma \int dt' G(t, t') C(t, t')^{p-2} x(t') \right. \\ &\left. - \frac{1}{r} \ln x(t) - \rho(t) + \eta(t) + \xi(t) \right]. \end{aligned} \quad (13)$$

We assume that $\xi(t)$ is white Gaussian noise of unit amplitude. Writing $x(t) = x^* + \hat{x}(t)$ and $\eta(t) = \eta^* + \hat{\eta}(t)$, and keeping only linear terms in ξ , \hat{x} , and $\hat{\eta}$, we obtain

$$\begin{aligned} \dot{\hat{x}}(t) = -\frac{1}{r} \hat{x}(t) + x^* &\left[\Gamma \int dt' G(t-t') C(t-t')^{p-2} \hat{x}(t') \right. \\ &\left. + \hat{\eta}(t) + \xi(t) \right]. \end{aligned} \quad (14)$$

Defining $H(t, t') = G(t, t') C(t, t')^{p-2}$, and taking the Fourier transform of Eq. (14), yields

$$\left(\frac{i\omega + r^{-1}}{x^*} - \Gamma \tilde{H}(\omega) \right) \tilde{x}(\omega) = \tilde{\eta}(\omega) + \tilde{\xi}(\omega), \quad (15)$$

where the tildes denote Fourier transforms. This leads to the relation

$$\begin{aligned} \langle |\tilde{x}(\omega)|^2 \rangle_* &= \left((p-1) \langle (x^*)^2 \rangle_*^{p-2} \langle |\tilde{x}(\omega)|^2 \rangle_* + \langle |\tilde{\xi}(\omega)|^2 \rangle_* \right) \\ &\times \langle |A(\omega, x^*)|^{-2} \rangle_*, \end{aligned} \quad (16)$$

where

$$\mathcal{A}(\omega, x^*) = \frac{i\omega + r^{-1}}{x^*} - \Gamma \tilde{H}(\omega). \quad (17)$$

We can write this as

$$\left\langle |\tilde{x}(\omega)|^2 \right\rangle_* = \left(\left\langle |\mathcal{A}(\omega, x^*)|^{-2} \right\rangle_*^{-1} - (p-1)q^{p-2} \right)^{-1}. \quad (18)$$

The left-hand side is positive by definition so the calculation runs into a contradiction if the expression on the right-hand side turns negative. As it approaches zero (from above) the magnitude of fluctuations diverges. The fixed point can only be stable when

$$\left\langle |\mathcal{A}(\omega, x^*)|^{-2} \right\rangle_* < \frac{1}{(p-1) \langle (x^*)^2 \rangle_*^{p-2}}. \quad (19)$$

Following [7] we focus on $\omega = 0$, so stability can only be expected provided that

$$\left\langle \left| \frac{1}{rx^*} - \Gamma q^{p-2} \chi \right|^{-2} \right\rangle_* < \frac{1}{(p-1)q^{p-2}}, \quad (20)$$

i.e.,

$$\int_{-\infty}^{\infty} Dz \left(\frac{1}{rx(z)} - \Gamma q^{p-2} \chi \right)^{-2} \leq \frac{1}{(p-1)q^{p-2}}. \quad (21)$$

For negative values of Γ , the position of the stability boundary can be determined straightforwardly by numerically solving Eqs. (11,12) for a given set of model parameters, and by subsequently evaluating the stability condition Eq. (21) throughout parameter space. As already mentioned, this procedure fails for $\Gamma > 0$.

IV.4. The large- p limit

In general the location of the stability region defined by (21) cannot be determined fully analytically. However, it is possible to make progress in several limits as shown below and make a good sketch of the behavior when $\Gamma < 0$.

Taking $\Gamma \rightarrow 0$ in Eq. (9) we find $x(z) = \exp[r(q^{(p-1)/2}z - \rho)]$, and using Eq. (11) the order parameters r , q , and ρ satisfy

$$\begin{aligned} \chi &= r, \\ \exp(r^2 q^{p-1}) &= q, \\ \rho &= \frac{\ln q}{2r}, \end{aligned} \quad (22)$$

in this limit. The second expression is equivalent to

$$q = \left(-\frac{W(-(p-1)r^2)}{(p-1)r^2} \right)^{1/(p-1)}, \quad (23)$$

where $W(\cdot)$ is the Lambert W function. For $(p-1)r^2 > 1/e$ this has no solutions. For $(p-1)r^2 < 1/e$ it has two, given by the two branches of the Lambert W function. The solution corresponding to the upper branch ($W > -1$) satisfies (21) so is always stable, while the other is always unstable. Therefore, on the $\Gamma = 0$ line, there is a stable fixed point only for large values of $1/r$, with the stability boundary given by

$$\begin{aligned} r = \chi &= \frac{1}{\sqrt{(p-1)e}}, \\ q &= \exp\left(\frac{1}{p-1}\right), \\ \rho &= \frac{1}{2} \sqrt{\frac{e}{p-1}}. \end{aligned} \quad (24)$$

For $\Gamma = -1$ and $p = 2$ the system is stable as soon as there is non-zero memory loss ($\alpha > 0$), that is to say the stability line passes through the point ($1/r = 0, \Gamma = -1$) when $p = 2$, see also [12]. For larger numbers of players we find that the stability boundary never reaches the $1/r = 0$ line.

The boundary between the two regions crosses the $\Gamma = 0$ line at the location we have determined analytically, and tends to a straight line as $p \rightarrow \infty$, as shown in Fig. 4. It is in fact possible to demonstrate this analytically, see Appendix A2 for details.

Based on this result we can estimate that the size, A , of the unstable region in the $\alpha/\beta - \Gamma$ plane as shown in Fig. 5 (restricted to $-1 \leq \Gamma \leq 0$). As the number of players p increases this area grows as

$$A \approx \sqrt{e(p-1)} \quad (25)$$

as shown in Fig. 5.

Thus, in the case where $\Gamma < 0$ and $p \rightarrow \infty$, the unstable region takes over the entire parameter space. That is, in situations where the players' payoffs are anticorrelated, their behavior will always be complex and will never settle down on an equilibrium.

V. COMPARISON BETWEEN THEORY AND NUMERICAL EXPERIMENTS

We now compare the numerical results to the theoretical predictions. We measure the stability boundary in the numerical experiments by determining whether or not the system converges to a unique fixed point, independent of initial conditions. To do this we choose a set of parameters and initial conditions, iterate the dynamics for a large number of time steps, and apply heuristics to check whether the players' strategies have converged to a fixed point. If we find a fixed point, we repeat this for many different initial conditions and check to see whether we always find the same fixed point. We also perform a similar procedure to check for limit cycles. The precise methods we use are discussed in detail in appendix A3.




FIG. 4. Stability boundaries of the effective dynamics for several values of p as a function of α/β and Γ , for the case where $\Gamma < 0$. Each curve is the stability boundary for the stated value of p . To the left of any curve the fixed point of the effective dynamics is unstable, to the right it is stable. The key result is that the stability boundary moves to the left as p increases, so the size of the regime with complex dynamics grows.

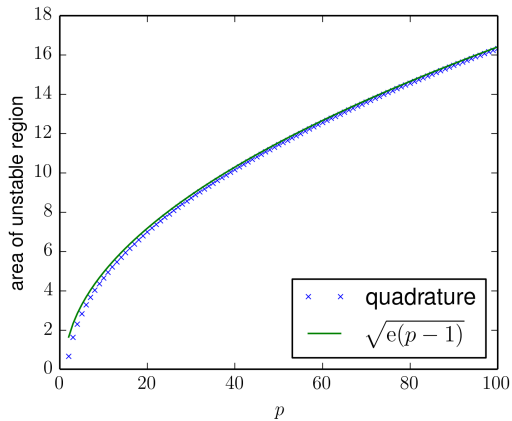




FIG. 5. Plot showing the area of the unstable region for negative Γ as a function of the number of players, p . This area is estimated numerically using Gaussian quadrature on results obtained for $\beta = 0.01$; this is valid in the limit as $N \rightarrow \infty$. The analytic estimate of the area is $\sqrt{e(p-1)}$, see Eq. (25). This indicates that the area of the parameter space with complex dynamics goes to infinity proportional to \sqrt{p} as $p \rightarrow \infty$.


Fig. 6 shows how the likelihood of converging to a unique stable fixed point varies throughout the negative- Γ region of parameter space for different values of N and p . The values we choose are roughly at the limit of what was computationally feasible. We investigate $p = 2, 3, 4, 5$, which constrains the corresponding values of N to be $N = 50, 12, 6, 4$. We then sweep Γ and α with $\beta = 0.05$ and construct a heat map showing the likelihood of convergence to a unique stable fixed point. We then compare




(a) $p = 2, N = 50$.



(b) $p = 3, N = 12$.



(c) $p = 4, N = 6$.



(d) $p = 5, N = 4$.

FIG. 6. Probability of convergence to a fixed point as a function of the memory parameter α and the competition parameter Γ . For each set of parameters we iterate the system from 500 random initial conditions. The heat maps show the fraction that converged to a fixed point (convergence criteria discussed in Appendix A3). Black means 100% convergence, red (grey) indicates the majority converge, yellow a minority, and white no convergence. The unstable region extends to larger values of α as the number of players is increased. The solid curves are derived from a generating functional analysis ($N \rightarrow \infty$, see Sec. IV.3), and separate the region in which a unique stable fixed point is to be expected in this limit (to the right of the green curves) from regions in parameters space where the behavior is more complex.

this to the stability line predicted by the generating functional approach described in the previous section.

The heat map of Fig. 6 is constructed so that black corresponds to convergence to a stable fixed point 100% of the time, red (grey) to convergence roughly 50% – 70%, yellow (light grey) 10% – 35%, and white to the case in which unique stable fixed points are never found. The behavior is consistent with what we described schematically in Fig. 3(b): Unique stable fixed points are more likely for higher α (i.e. short memory) and the size of the stable region grows with increasing Γ . The region in which complex dynamics are observed grows as p increases; in particular for the zero-sum case where $\Gamma = -1$ the size of the interval corresponding to complex dynamics is finite and growing with p .

The correspondence between the predicted vs. the observed stability boundary gets better as N increases. For $p = 2$, where we can make $N = 50$, the correspondence is quite good (Fig. 6(a)); for $p = 5$, where we are only able to make $N = 4$, the correspondence is not as good (Fig. 6(d)); the stability line scales more or less tracks the numerically-observed boundary, but is consistently to the right of it. Given that $N = 4 \ll N = \infty$, it is not surprising that the approximation is not perfectly accurate. We hypothesize that this is due to finite size effects. To test this, in Fig. A4 in the Appendix we hold the number of players constant at $p = 2$ and systematically vary N . We find the correspondence between theory and experiment improving with increasing N . In addition the behavior becomes crisper in the sense that the transition from certain convergence to a unique fixed point to never converging to a unique fixed point happens more suddenly when N becomes large. This indicates that the generating function methods gives reasonably good predictions for large N , lending confidence to its reliability in the limit as $N \rightarrow \infty$.

VI. CHAOS AND VOLATILITY CLUSTERING

Before concluding we would like to make a few notes about chaos and volatility clustering. We have so far asserted that much of the behavior in the competitive region where $\Gamma < 0$ is chaotic, without presenting any evidence. In fact we have done extensive computation of Lyapunov exponents using the procedures described in Galla and Farmer [12]. While we experience some numerical problems we can nonetheless state with confidence that the preponderance of the complex dynamics to the right of the stability line corresponds to chaos. Problems arise because it can sometimes be difficult to numerically distinguish chaos and limit cycles without making very long simulations, and because of the metastable chaos observed in Fig. 2, which means that in any given simulation there is a small but nonzero probability that the simulation will eventually collapse to a fixed point. Nonetheless, most of the time we observe chaos, and as p increases it tends to be of higher dimension. To prove our main point




FIG. 7. Time series of the changes in the sum of the players’ payoffs for a game with three-players. This corresponds to high dimensional chaos and clustered volatility. By this we mean the tendency for time variability to be positively autocorrelated, with periods of relative calm and periods of relative variability.

here and compare to the theory we only needed to determine whether or not we observe convergence to a unique fixed point, which is much easier computationally, so we have chosen not to present evidence based on Lyapunov exponents.

In the chaotic regime we consistently observe clustered volatility, similar to that reported for $p = 2$ by Galla and Farmer. By this we mean that the fluctuation in payoffs to the players fluctuates in time in a way that is “clustered”, i.e. positively autocorrelated. There are epochs in which the payoffs are relatively steady and other epochs in which they are highly variable, as shown in Fig. 7. For $p > 2$ the chaos tends to be higher dimensional and the clustered volatility stronger. Clustered volatility is common in many real-world situations, including financial time series; our work here suggests that this may be a generic result for games in which players learn their strategies using procedures similar to EWA. We conjecture that this is connected to the tendency for a given action to vary from being used frequently for long periods of time to being almost never used, as observed in Fig. 1(c), but this remains to be investigated.

VII. CONCLUSIONS

In summary we have characterized the outcome of adaptive learning in complex multi-player games with Gaussian random payoff matrices. The learning dynamics we have simulated is a special case of experience-weighted attraction learning used in behavioral economics. We see different types of dynamical behavior: convergence to fixed points, limit cycles and chaotic trajectories. Broadly speaking, convergence to a unique stable equilibrium is observed when players’ learning neglects out-

comes from the distant past, i.e. when they forget quickly, corresponding to large values of α/β . In contrast, when they have a long memory, i.e. small α/β , we observed more complex dynamics. The nature of this dynamics depends on how competitive the game is. For competitive games (i.e. with $\Gamma < 0$) the complex dynamics exhibits itself as a limit cycle or chaotic attractor. For cooperative games (i.e. with $\Gamma > 0$) the complex dynamics exhibits itself as a multiplicity of fixed points. The boundaries between these behaviors become sharper as N increases.

The main focus of this paper was to study the properties of games with more than two players. For two players we replicate the results of Galla and Farmer [12]. For more than two players we are not able to simulate situations with a very large number of actions due to computational constraints, e.g. for $p = 3$ players the largest number of actions we simulated was $N = 12$. To clearly understand the behavior for large p and large N we rely on an analytic treatment, which allows us to estimate the stability boundary in the limit as $N \rightarrow \infty$. Based on tools from the theory of disordered systems, we have carried out a generating-functional analysis of the continuous-time limit of EWA learning, and we have derived approximate semi-analytical results for the onset of stability for games with an infinite number of strategies and with an arbitrary finite number of players, p . These results reveal that the parameter range in which learning cannot be expected to settle down to fixed points increases as the number of players in the game grows. This is summarized in Figs. 4 and 5. In contrast to the Galla and Farmer paper, where the analytic results were just making the numerical results more rigorous, here the analytic methods were essential to understand the behavior for large p .

In the introduction we posed our objective as seeking a ‘‘Reynolds number’’ for estimating the *a priori* likeli-

hood of complex dynamics, in much the same way that the Reynolds number characterizes turbulence in fluid flow. Indeed the parameter $r = \beta/\alpha$ characterizing the timescale of the learning process does a reasonably good job in this context: As r gets bigger, complex dynamics becomes more likely. The transition also depends on the competitiveness of the game, characterized by Γ , as well as the number of players. Our key result here is that complex dynamics become more likely as the number of players increases. This is not surprising given that games with more players are more complicated, and perhaps also more complex, in the sense that there are more factors to take into account and more inherent degrees of freedom.

The standard theoretical approach in economics assumes convergence to an equilibrium from the outset. Our results here suggest that under circumstances where the players have a long memory of the past, this approach may be inherently flawed. This is particularly true when there are many agents. Our results suggest that there may be large regimes in which the assumption of a unique equilibrium is completely invalid, and where approaches that can accommodate chaotic dynamics, such as agent-based modeling, are needed. Of course in this paper we have only studied one family of learning algorithms, and we have focused on games with many actions. More work is needed to give a definitive answer to the question above.

ACKNOWLEDGEMENTS

JBTS thanks the Engineering and Physical Sciences Research Council (EPSRC) for support and JDF thanks the Institute for New Economic Thinking.

-
- [1] J. von Neumann, O. Morgenstern, *Theory of Games and Economic behavior*, Princeton University Press, Princeton NJ (2007)
 - [2] J. Nash, Equilibrium points in n-person games, *Proceedings of the National Academy of Sciences* **36** (1) 48-49 (1950)
 - [3] J. D. Farmer, J. Geneakoplos, ‘‘The Virtues and Vices of Equilibrium and the Future of Financial Economics.’’ *Complexity* (2008)
 - [4] A. McLennan, J. Berg, *Games and Economic Behavior* **51**(2), 264-295 (2005)
 - [5] J. Berg, M. Weigt, *Europhys. Lett.* **48**(2), 129-135 (1999).
 - [6] C. Daskalakis, P. W. Goldberg, C. H. Papadimitriou, *Commun. ACM* **52** 89 (2009).
 - [7] M. Opper, S. Diederich, *Phys. Rev. Lett.* **69** 1616-1619 (1992)
 - [8] S. Diederich, M. Opper, *Phys. Rev. A* **39** 4333-4336 (1989).
 - [9] Y. Sato, E. Akiyama, J. D. Farmer, *Proc. Nat. Acad. Sci. USA* **99** 4748-4751 (2002)
 - [10] Y. Sato, J.-P. Crutchfield, *Phys. Rev. E* **67** 015206(R) (2003)
 - [11] Y. Sato, E. Akiyama, J.-P. Crutchfield, *Physica D* **210**, 21 (2005)
 - [12] T. Galla, J. D. Farmer, *Proc. Nat. Acad. Sci (USA)* **110** 1232 (2013)
 - [13] T. H. Ho, C. F. Camerer, J.-K. Chong, *J. Econ. Theor.* **133** 177-198 (2007)
 - [14] C. Camerer, T.H. Ho, *Econometrica* **67** (1999) 827
 - [15] C. Camerer, *Behavioral Game Theory: Experiments in Strategic Interaction (The Roundtable Series in Behavioral Economics)*, Princeton University Press, Princeton NJ, 2003
 - [16] D. Fudenberg, D.K. Levine, *Theory of Learning in Games*, MIT Press, Cambridge MA (1998)
 - [17] J. A. Hertz, Y. Roudi, P. Sollich, preprint arXiv:1604.05775

- [18] T. Galla, J. Phys. A: Math. Gen. **39** 3853 (2006)
- [19] V. M. de Oliveira, J. F. Fontanari, Phys. Rev. Lett. **85** 4984 (2000)
- [20] H. P. Young, *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions*, Princeton University Press, Princeton NJ (1998)
- [21] W. A. Brock, C. H. Hommes, J. Econ. Dyn. and Contr. **22** 1235-1274 (1998)
- [22] B. Skyrms, J. of Logic, Language and Information **1** 111-130 (1992)
- [23] M. A. Nowak, *Evolutionary dynamics*, Harvard University Press, Cambridge MA (2006)
- [24] J. Hofbauer, K. Sigmund, *Evolutionary games and population dynamics*, Cambridge University Press, Cambridge, 1998
- [25] R. M. May, Will a Large Complex System be Stable? Nature **238**, 413 - 414 (1972);
- [26] E. T. Jaynes, Information Theory and Statistical Mechanics, Phys. Rev. **106** 620-630 (1957)
- [27] E. T. Jaynes, Information Theory and Statistical Mechanics II, Phys. Rev. **108** 171-190 (1957)
- [28] T. M. Cover, J. A. Thomas, *Elements of information theory* (2nd edition), Wiley-Interscience, Hoboken, NJ (2006)
- [29] Note that the fixed point reached in the stable regime is only a Nash equilibrium in the limit $\alpha \rightarrow 0$. When $\alpha > 0$ the players are effectively assuming their opponent's behavior is non-stationary, and that more recent moves are more useful than moves in the distant past.
- [30] De Dominicis, C., Phys. Rev. B **18** 4913-4919 (1978)
- [31] Ghashghaie, S., Breymann, W., Peinke, J., Talkner, P., Dodge, Y., Nature **381** 767-770 (1996)
- [32] In contrast to financial markets, for the behavior we observe here the distribution of heavy tails decay exponentially (as opposed to following a power law). We hypothesize that this is because the players in financial markets use a variety of different timescales α .
- [33] J. B. T. Sanders, T. Galla and J. D. Farmer, unpublished research.
- [34] Lorenz, E. N., J. Atmos. Sci. **26**, 636-646 (1969)
- [35] J. D. Farmer, J. J. Sidorowich, Phys. Rev. Lett. **59** 845-848 (1987)
- [36] D. P. Forster, H. P. Young, Proc. Nat. Acad. Sci. USA **98** 12848-12853 (2001)
- [37] S. Hart, A. Mas-Colell, American Economic Review **93** (5) 1830-1836 (2003)

APPENDIX

Appendix A1: Generating functional analysis

Following [7, 12], we perform a generating functional analysis of the Sato-Crutchfield dynamics (i.e., the continuous limit of EWA). This will lead to an effective dynamics that is representative of the continuous limit of the EWA system, for large values of N , for typical realizations of the payoffs, and after averaging over the ensemble of random games. The fixed points of the effective dynamics are far easier to study analytically than those of the the Sato-Crutchfield equations for any particular random game.

Consider the dynamics

$$\frac{\dot{x}_i^\mu(t)}{x_i^\mu(t)} = -r^{-1} \ln x_i^\mu(t) + \sum_{i_{\mu+1}, \dots, i_{\mu-1}} \Pi_{i, i_{\mu+1}, \dots, i_{\mu-1}}^\mu \prod_{\kappa \neq \mu} x_{i_\kappa}^\kappa(t) - \rho^\mu(t) + h_i^\mu(t). \quad (\text{A1})$$

This is identical to the Sato-Crutchfield dynamics (4), except that we have added arbitrary functions $h_i^\mu(t)$ to generate response functions—these will later be set to zero. Recall that the normalization term $\rho^\mu(t)$ is defined such that the $x_i^\mu(t)$ have mean 1.

We define a generating functional

$$Z[\psi] = \int \mathcal{D}[x] \delta(\text{equations of motion}) \exp \left(i \sum_{\mu, i} \int dt x_i^\mu(t) \psi_i^\mu(t) \right), \quad (\text{A2})$$

where $\delta(\text{equations of motion})$ is used to mean that the integral is performed over realizations of (A1). Writing these delta functions in Fourier form yields

$$Z[\psi] = \int \mathcal{D}[x, \hat{x}] \exp \left(i \sum_{\mu, i} \int dt \left\{ \hat{x}_i^\mu(t) \left(\frac{\dot{x}_i^\mu(t)}{x_i^\mu(t)} + r^{-1} \ln x_i^\mu(t) - \sum_{i_{\mu+1}, \dots, i_{\mu-1}} \Pi_{i, i_{\mu+1}, \dots, i_{\mu-1}}^\mu \prod_{\kappa \neq \mu} x_{i_\kappa}^\kappa(t) + \rho^\mu(t) - h_i^\mu(t) \right) + x_i^\mu(t) \psi_i^\mu(t) \right\} \right). \quad (\text{A3})$$

The factor in this expression depending on the payoff elements is

$$Z_{\Pi} = \exp \left(-i \sum_{\mu, i_1, \dots, i_p} \int dt \Pi_{i_\mu, i_{\mu+1}, \dots, i_{\mu-1}}^\mu \hat{x}_{i_\mu}^\mu(t) \prod_{\kappa \neq \mu} x_{i_\kappa}^\kappa(t) \right). \quad (\text{A4})$$

Averaging this over the payoff elements gives

$$\begin{aligned} \overline{Z_{\Pi}} = & \prod_{i_1, \dots, i_p} \exp \left\{ -\frac{1}{2N^{p-1}} \sum_{\mu} \int dt \int dt' \left[\hat{x}_{i_\mu}^\mu(t) \hat{x}_{i_\mu}^\mu(t') \left(\prod_{\kappa \neq \mu} x_{i_\kappa}^\kappa(t) \right) \left(\prod_{\lambda \neq \mu} x_{i_\lambda}^\lambda(t') \right) \right. \right. \\ & \left. \left. + \Gamma \sum_{\nu \neq \mu} \hat{x}_{i_\mu}^\mu(t) \hat{x}_{i_\nu}^\nu(t') \left(\prod_{\kappa \neq \mu} x_{i_\kappa}^\kappa(t) \right) \left(\prod_{\lambda \neq \nu} x_{i_\lambda}^\lambda(t') \right) \right] \right\}, \end{aligned} \quad (\text{A5})$$

which can be written as

$$\overline{Z_{\Pi}} = \exp \left\{ -\frac{N}{2} \int dt \int dt' \sum_{\mu} \left(L^\mu(t, t') \prod_{\kappa \neq \mu} C^\kappa(t, t') + \Gamma \sum_{\nu \neq \mu} K^\mu(t, t') K^\nu(t', t) \prod_{\kappa \notin \{\mu, \nu\}} C^\kappa(t, t') \right) \right\}, \quad (\text{A6})$$

where we have introduced the functions

$$\begin{aligned} C^\mu(t, t') &= \frac{1}{N} \sum_i x_i^\mu(t) x_i^\mu(t'), \\ K^\mu(t, t') &= \frac{1}{N} \sum_i x_i^\mu(t) \hat{x}_i^\mu(t'), \\ L^\mu(t, t') &= \frac{1}{N} \sum_i \hat{x}_i^\mu(t) \hat{x}_i^\mu(t'). \end{aligned} \quad (\text{A7})$$

We can use the expression (A6) in (A3), introducing the functions C^μ , K^μ , and L^μ into the integral using delta functions, for example

$$\begin{aligned} 1 &= \int \mathcal{D}[C^\mu] \prod_{t,t'} \delta \left(C^\mu(t,t') - \frac{1}{N} \sum_i x_i^\mu(t) x_i^\mu(t') \right) \\ &= \int \mathcal{D}[C^\mu, \widehat{C}^\mu] \exp \left(iN \int dt \int dt' \widehat{C}^\mu(t,t') \left(C^\mu(t,t') - \frac{1}{N} \sum_i x_i^\mu(t) x_i^\mu(t') \right) \right). \end{aligned} \quad (\text{A8})$$

The generating functional becomes

$$\overline{Z[\psi]} = \int \mathcal{D}[C, \widehat{C}, K, \widehat{K}, L, \widehat{L}] \exp(N(\psi + \Phi + \Omega)), \quad (\text{A9})$$

where

$$\Psi = i \sum_\mu \int dt \int dt' \left(\widehat{C}^\mu(t,t') C^\mu(t,t') + \widehat{K}^\mu(t,t') K^\mu(t,t') + \widehat{L}^\mu(t,t') L^\mu(t,t') \right) \quad (\text{A10})$$

results from the introduction of C , K , and L into the integral,

$$\Phi = -\frac{1}{2} \sum_\mu \int dt \int dt' \left(L^\mu(t,t') \prod_{\kappa \neq \mu} C^\kappa(t,t') + \Gamma \sum_{\nu \neq \mu} K^\mu(t,t') K^\nu(t',t) \prod_{\kappa \notin \{\mu,\nu\}} C^\kappa(t,t') \right) \quad (\text{A11})$$

results from the average over the payoff elements, and

$$\begin{aligned} \Omega &= \frac{1}{N} \sum_{\mu,i} \ln \left\{ \int \mathcal{D}[x_i^\mu, \widehat{x}_i^\mu] p_{i,0}^\mu(x_i^\mu(0)) \exp \left(i \int dt x_i^\mu(t) \psi_i^\mu(t) \right) \exp \left(i \int dt \widehat{x}_i^\mu(t) \left(\frac{\dot{x}_i^\mu(t)}{x_i^\mu(t)} + \frac{1}{r} \ln x_i^\mu(t) + \rho^\mu(t) - h_i^\mu(t) \right) \right) \right. \\ &\quad \left. \times \exp \left[-i \int dt \int dt' \left(\widehat{C}^\mu(t,t') x_i^\mu(t) x_i^\mu(t') + \widehat{K}^\mu(t,t') x_i^\mu(t) \widehat{x}_i^\mu(t') + \widehat{L}^\mu(t,t') \widehat{x}_i^\mu(t) \widehat{x}_i^\mu(t') \right) \right] \right\} \end{aligned} \quad (\text{A12})$$

contains the integral over x and \widehat{x} . Here, $p_{i,0}^\mu(\cdot)$ represents the initial distribution of x_i^μ .

In the limit as $N \rightarrow \infty$, the integrals in (A9) can be performed using the saddle-point method. Extremising the exponent with respect to C^μ , K^μ , and L^μ gives the relations

$$\begin{aligned} i\widehat{C}^\mu(t,t') &= \frac{1}{2} \sum_{\nu \neq \mu} \left(L^\nu(t,t') \prod_{\kappa \notin \{\mu,\nu\}} C^\kappa(t,t') + \Gamma \sum_{\kappa \notin \{\mu,\nu\}} K^\nu(t,t') K^\kappa(t',t) \prod_{\lambda \notin \{\mu,\nu,\kappa\}} C^\lambda(t,t') \right), \\ i\widehat{K}^\mu(t,t') &= \Gamma \sum_{\nu \neq \mu} K^\nu(t,t') \prod_{\kappa \notin \{\mu,\nu\}} C^\kappa(t,t'), \\ i\widehat{L}^\mu(t,t') &= \frac{1}{2} \prod_{\kappa \neq \mu} C^\kappa(t,t'), \end{aligned} \quad (\text{A13})$$

while extremisation with respect to \widehat{C}^μ , \widehat{K}^μ , and \widehat{L}^μ leads to

$$\begin{aligned} C^\mu(t,t') &= \lim_{N \rightarrow \infty} \frac{1}{N} \sum_i \langle x_i^\mu(t) x_i^\mu(t') \rangle_\Omega, \\ K^\mu(t,t') &= \lim_{N \rightarrow \infty} \frac{1}{N} \sum_i \langle x_i^\mu(t) \widehat{x}_i^\mu(t') \rangle_\Omega, \\ L^\mu(t,t') &= \lim_{N \rightarrow \infty} \frac{1}{N} \sum_i \langle \widehat{x}_i^\mu(t) \widehat{x}_i^\mu(t') \rangle_\Omega, \end{aligned} \quad (\text{A14})$$

where $\langle \cdot \rangle_\Omega$ represents a mean taken against a measure defined by Ω , see for example the Supplemental Material of [12] for details in a similar calculation for $p = 2$.

It can also be seen, from the definition of the generating functional, that we have

$$\begin{aligned}
C^\mu(t, t') &= - \lim_{N \rightarrow \infty} \frac{1}{N} \sum_i \frac{\delta^2 \overline{Z[\psi]}}{\delta \psi_i^\mu(t) \delta \psi_i^\mu(t')} \Big|_{\psi=h=0}, \\
K^\mu(t, t') &= - \lim_{N \rightarrow \infty} \frac{1}{N} \sum_i \frac{\delta^2 \overline{Z[\psi]}}{\delta \psi_i^\mu(t) \delta h_i^\mu(t')} \Big|_{\psi=h=0}, \\
L^\mu(t, t') &= - \lim_{N \rightarrow \infty} \frac{1}{N} \sum_i \frac{\delta^2 \overline{Z[\psi]}}{\delta h_i^\mu(t) \delta h_i^\mu(t')} \Big|_{\psi=h=0}.
\end{aligned} \tag{A15}$$

Because of normalization, $Z[\psi = 0] = 1$ for all h , so $L^\mu(t, t') = 0 \forall t, t'$. Due to causality, we have $K^\mu(t, t') = 0$ for $t' > t$, so that $K^\mu(t, t')K^\nu(t', t) = 0$.

This leaves $\Psi + \Phi = 0$, and if we set $\psi = 0$, and assume identical perturbations $h_i^\mu(t) = h(t)$ and initial distributions $p_{i,0}^\mu(x) = p_0(x)$ for all players and strategy components, then we have

$$\begin{aligned}
\Omega &= p \ln \left\{ \int \mathcal{D}[x, \hat{x}] p_0(x(0)) \exp \left(i \int dt \hat{x}(t) \left(\frac{\dot{x}(t)}{x(t)} + \frac{1}{r} \ln x(t) + \rho(t) - h(t) \right) \right) \right. \\
&\quad \left. \times \exp \left[- \int dt \int dt' \left(\Gamma(p-1) K(t, t') C(t, t')^{p-2} x(t) \hat{x}(t') + \frac{1}{2} C(t, t')^{p-1} \hat{x}(t) \hat{x}(t') \right) \right] \right\}
\end{aligned} \tag{A16}$$

where we have dropped the distinction between different players and strategy components. Each degree of freedom then has an effective generating functional

$$\begin{aligned}
Z_{\text{eff}} &= \int \mathcal{D}[x, \hat{x}] p_0(x(0)) \exp \left(i \int dt \hat{x}(t) \left(\frac{\dot{x}(t)}{x(t)} + \frac{1}{r} \ln x(t) + \rho(t) - h(t) \right) \right) \\
&\quad \times \exp \left[- \int dt \int dt' \left(\Gamma K(t, t') C(t, t')^{p-2} x(t) \hat{x}(t') + \frac{1}{2} C(t, t')^{p-1} \hat{x}(t) \hat{x}(t') \right) \right].
\end{aligned} \tag{A17}$$

Defining $G(t, t') = -iK(t, t')$, we have

$$\begin{aligned}
Z_{\text{eff}} &= \int \mathcal{D}[x, \hat{x}] p_0(x(0)) \exp \left(i \int dt \hat{x}(t) \left(\frac{\dot{x}(t)}{x(t)} + \frac{1}{r} \ln x(t) + \rho(t) - h(t) \right) \right) \\
&\quad \times \exp \left[- \int dt \int dt' \left(i \Gamma G(t, t') C(t, t')^{p-2} x(t) \hat{x}(t') + \frac{1}{2} C(t, t')^{p-1} \hat{x}(t) \hat{x}(t') \right) \right],
\end{aligned} \tag{A18}$$

which is identical to the generating functional of the effective dynamics

$$\frac{\dot{x}(t)}{x(t)} = \Gamma \int dt' G(t, t') C(t, t')^{p-2} x(t') - \frac{1}{r} \ln x(t) - \rho(t) + \eta(t) + h(t), \tag{A19}$$

where $\eta(t)$ is a Gaussian random variable satisfying $\langle \eta(t) \eta(t') \rangle_* = C(t, t')^{p-1}$ and $\langle \eta(t) \rangle_* = 1$, and the functions G and C are determined by

$$\begin{aligned}
G(t, t') &= \left\langle \frac{\delta x(t)}{\delta h(t')} \right\rangle_*, \\
C(t, t') &= \langle x(t) x(t') \rangle_*,
\end{aligned} \tag{A20}$$

with $\langle \cdot \rangle_*$ used to denote an average over the effective dynamics (A19). Finally setting h to zero, the effective system is

$$\frac{\dot{x}(t)}{x(t)} = \Gamma \int dt' G(t, t') C(t, t')^{p-2} x(t') - \frac{1}{r} \ln x(t) - \rho(t) + \eta(t). \tag{A21}$$

with G , C , and η defined as above.

Appendix A2: Onset of instability in the large- p limit

Writing $n = p - 1$ for convenience, the boundary of the stable region is given by the solution of the following equations:

$$\begin{aligned}
\frac{1}{r} \ln x - \Gamma q^{n-1} \chi x - q^{n/2} z + \rho &= 0, \\
\int_{-\infty}^{\infty} Dz \frac{\partial x}{\partial z} &= q^{n/2} \chi, \\
\int_{-\infty}^{\infty} Dz x^2 &= q \\
\int_{-\infty}^{\infty} Dz x &= 1 \\
\int_{-\infty}^{\infty} Dz \left(\frac{\partial x}{\partial z} \right)^2 &= \frac{q}{n},
\end{aligned} \tag{A1}$$

where Dz is a shorthand for the standard Gaussian measure $Dz = \frac{dz}{\sqrt{2\pi}} e^{-z^2/2}$.

For $\Gamma = 0$ the order parameters at the boundary of the stable region are given by Eq. (24). As an ansatz for the region with $\Gamma < 0$ we assume that the order parameters and the value of r on the phase boundary scale with n in the same way as they do for $\Gamma = 0$. We can write

$$\begin{aligned}
q &= 1 + n^{-1} q', \\
\chi &= n^{-\frac{1}{2}} \chi', \\
r &= n^{-\frac{1}{2}} r', \\
\rho &= n^{-\frac{1}{2}} \rho',
\end{aligned} \tag{A2}$$

where all primed variables are of order $\mathcal{O}(n^0)$.

If we also write $x = 1 + n^{-\frac{1}{2}} x'$, and retain only leading-order terms in q' , χ' , r' , and ρ' we obtain from Eq. (A1):

$$\begin{aligned}
\frac{1}{r'} \ln \left(1 + n^{-\frac{1}{2}} x' \right) - n^{-\frac{1}{2}} (1 + q'/2) z + n^{-1} \rho' \\
-n^{-1} \Gamma (1 + q') \chi' - n^{-\frac{3}{2}} \Gamma (1 + q') \chi' x' &= 0, \\
\int_{-\infty}^{\infty} Dz \frac{\partial x'}{\partial z} &= \left(1 + \frac{q'}{2} \right) \chi', \\
\int_{-\infty}^{\infty} Dz x'^2 &= q', \\
\int_{-\infty}^{\infty} Dz x' &= 0, \\
\int_{-\infty}^{\infty} Dz \left(\frac{\partial x'}{\partial z} \right)^2 &= 1 + \frac{q'}{n}.
\end{aligned} \tag{A3}$$

The linear term in x' in the first of these equations is dominated by the log term except at large x . Specifically, by using the approximation $W_{-1}(y) \approx \ln(-y)$ as $y \rightarrow \infty$, it can be seen that the linear term reaches the size of the log term when the value of x' , to leading order, is

$$x' \approx x_l = \frac{n^{\frac{3}{2}} \ln n}{-\Gamma r' (1 + q') \chi'}, \tag{A4}$$

while to leading order z is

$$z \approx z_l = \frac{n^{\frac{1}{2}} \ln n}{r' \left(1 + \frac{q'}{2} \right)} \tag{A5}$$

It remains only to show that the region of the real line beyond z_l makes a vanishing contribution to the integrals in Eqs. (A3). By ignoring the linear term for $z < z_l$, and the log term for $z > z_l$, the integrals over these two regions can be approximated analytically. In each case, the $z > z_l$ contribution shrinks more quickly as n grows.

Neglecting the linear term in the first relation in Eq. (A3) altogether is equivalent to making the approximation $x = 1$ in the linear term in the first equation of (A1). This yields a system of the same form as the $\Gamma = 0$ special case, except for an additional constant term, which can be solved exactly in the same manner. So, to leading order, the parameters r , χ , and q take constant values along the stability curve for large p , while ρ takes the value

$$\rho = \left(\frac{1}{2} + \Gamma\right) \sqrt{\frac{e}{n}}. \quad (\text{A6})$$

This value for ρ scales with n in the same way as it does in our ansatz, so the ansatz is indeed valid for all negative values of Γ . This demonstrates that $r = \sqrt{e(p-1)}$ is a solution of the equations for the onset of instability in the limit of large p .

Appendix A3: Heuristic classification of the dynamic behavior

It is not necessarily straightforward to classify the long-term behavior of even low-dimensional dynamical systems using empirical data. For high-dimensional systems such as the EWA dynamics for games with large number of players and/or strategies this task can be extremely challenging.

We experience two major difficulties. Firstly, in some regions of parameter space, transient behavior can last for a very long time, and can appear chaotic for all intents and purposes even though the system eventually reaches a stable fixed point. Secondly, characterising chaos using Lyapunov exponents or measures of dimension can be problematic for such large systems. the Jacobians of the system for example can be badly conditioned.

However, these difficult cases are not the norm, and we can use heuristics to classify behavior as convergence to a stable fixed point, convergence to a limit cycle, or chaos, with a high degree of accuracy.

For a given set of parameters and initial conditions, we iterate the EWA system for a maximum of 500000 time steps, split into batches of 10000 steps. After each batch, we explicitly check for the appearance of fixed points or limit cycles. If the relative difference between the maximum and minimum values of each strategy component was less than 1%, we assume a stable fixed point has been found. If there is a τ such that all of $x_i^\mu(t + \tau)$, $x_i^\mu(t + 2\tau)$, etc. (where t was the time at the start of the batch) have components within 0.1% of the components of $x_i^\mu(t)$, then we assume a stable limit cycle has been found. Otherwise, we continue to the next batch. If convergence has not been detected after 500000 time steps, we assume the system is chaotic.

This heuristic was used to produce the plots in Figs. 6 and in Appendix A4.

Appendix A4: Further numerical results: limit cycles and multiplicity of fixed points

1. Competitive games ($\Gamma < 0$)

In Fig. A1 we show the likelihood of converging to a limit cycle for games with negatively correlated payoff matrix elements, i.e. games in which players compete against each other ($\Gamma < 0$). For intermediate values of α , just smaller than those for which stable fixed points are ubiquitous, limit cycles are seen very commonly. However, at small values of α , fixed points or limit cycles are achieved only rarely—chaos is the norm.

2. Positively correlated payoffs ($\Gamma > 0$)

For positive values of the competition parameter, chaotic dynamics is rarely observed (though chaotic-appearing transients are frequently seen). Instead, for smaller values of α and Γ , limit cycles are very common, as shown in Fig. A2. In the rest of this region, EWA consistently converges to a fixed point. However, for small values of α and large values of Γ , there are many distinct fixed points that the dynamics can converge to for a given payoff matrix. This is shown in Fig. A3.




FIG. A1. Heat maps showing the fraction of 500 random initial conditions for which the EWA system converged to a limit cycle according to the heuristic in Appendix A3, for negative Γ . Limit cycles appear in a narrow band at intermediate values of α . The green curves show the boundaries of the stable region as derived in section IV.3.




FIG. A2. Heat maps showing the fraction of 500 random initial conditions for which the EWA system converged to a limit cycle according to the heuristic in Appendix A3, for positive Γ . Limit cycles appear most commonly when the payoffs are weakly correlated.




FIG. A3. Heat maps showing the fraction of 20 independent payoff matrices for which the EWA dynamics converged to multiple distinct fixed points for different initial conditions. For each payoff matrix, the EWA system was iterated for 100 different initial conditions, with fixed points being detected as explained in Appendix A3. Fixed points were considered to be identical if the relative distance between each component was less than 0.1.





(a) $p = 2, N = 10$.(b) $p = 2, N = 20$. $p = 2, N = 50$  $p = 2, N = 100$

FIG. A4. Heat maps showing the dependence on N of the stable region for $p = 2$ and $\Gamma < 0$. For each set of parameters the system was iterated for 500 random initial conditions. The heat maps show the fraction that converged to a fixed point (numerical convergence criteria are described in Appendix A3). The size of the unstable region grows with the size N of the payoff matrix, but begins to converge around $N = 50$. The green curves are the stability curves as derived in section IV.3.