

Stability and diversity in collective adaptation

Yuzuru Sato^{a,*}, Eizo Akiyama^b, James P. Crutchfield^a

^a Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501, USA

^b Institute of Policy and Planning Sciences, University of Tsukuba, Tennodai 1-1-1, Tsukuba, Ibaraki 305-8573, Japan

Received 26 November 2004; received in revised form 6 May 2005; accepted 21 June 2005

Available online 8 August 2005

Communicated by Y. Kuramoto

Abstract

We derive a class of macroscopic differential equations that describe collective adaptation, starting from a discrete-time stochastic microscopic model. The behavior of each agent is a dynamic balance between adaptation that locally achieves the best action and memory loss that leads to randomized behavior. We show that, although individual agents interact with their environment and other agents in a purely self-interested way, macroscopic behavior can be interpreted as game dynamics. Application to several familiar, explicit game interactions shows that the adaptation dynamics exhibits a diversity of collective behaviors. The simplicity of the assumptions underlying the macroscopic equations suggests that these behaviors should be expected broadly in collective adaptation. We also analyze the adaptation dynamics from an information-theoretic viewpoint and discuss self-organization induced by the dynamics of uncertainty, giving a novel view of collective adaptation.

© 2005 Elsevier B.V. All rights reserved.

Keywords: Collective adaptation; Game theory; Information theory; Dynamical systems

1. Introduction

Collective behavior in groups of adaptive systems is an important and cross-cutting topic that appears under various guises in many fields, including biology, neuroscience, computer science, and social science. In these fields, adaptive systems consist of individual agents that interact with one another and modify their behaviors according to the information they receive through those interactions. Often, though, collective behaviors emerge that are beyond the individual agent's perceptual capabilities and that sometimes frustrate satisfying their local goals. With competitive interactions dynamic adaptation can produce rich and unexpected behaviors. This kind of mutual adaptation has been discussed, for example, in studies of biological group interaction [1–3], interactive learning [4–6], large-scale adaptive systems [7,8], and learning in games [9,10].

* Corresponding author. Tel.: +1 505 984 8800; fax: +1 505 982 0565.

E-mail address: ysato@santafe.edu (Y. Sato).

Here we develop a class of coupled differential equations for mutual adaptation in agent collectives—systems in which agents learn how to act in their environment and with other agents through reinforcement of their actions. We show that the adaptive behavior in agent collectives, in special cases, reduces to a generalized form of multipopulation replicator equations and, generally, can be viewed as an information-theoretic self-organization in a collective adaptive system.

Suppose that many agents interact with an environment and each independently attempts to adjust its behavior to the environment based on its sensory stimuli. The environment consists of other agents and other exogenous influences. The agents could be humans, animals, or machines, but we make no assumptions about their detailed internal structure. That is, the central hypothesis in the following is that collective adaptation is a dynamical behavior driven by agents' environment-mediated interactions. By separating the time scales of change in the environment, of agents' adaptation, and of agent–agent interactions, our models describe, not the deterministic decision-making itself, but the temporal change in the probability distribution of choices.

1.1. Related work

This approach should be compared and contrasted with a game-theoretic view [11]. First, classical game theory often assumes that players have knowledge of the entire environmental structure and of other players' decision-making processes. Our adaptive agents, however, have no knowledge of a game in which they might be playing. Thus, unlike classical game theory, in our setting there is no bird's eye view for the entire collective that is available to the agents. Agents have only a myopic model of the environment, since any information external to them is given implicitly via the reinforcements for their action choices. Second, although we employ game-theoretic concepts such as Nash equilibria, we focus almost exclusively on *dynamics*—transients, attractors, and so on—of collective adaptation, while, naturally, making contact with the *statics* familiar from game theory. Finally, despite the differences, game structures can be introduced as a set of parameters corresponding to approximated static environments.

While replicator dynamics were introduced originally for evolutionary game theory [12–14], the relationship between learning with reinforcement and replicator dynamics has been discussed only recently [9,10]. Briefly stated, in our model the state space represents an individual agent's probability distribution to choose actions and the adaptation equations describe the temporal evolution of choice probabilities as the agents interact. Here, we extend these considerations to collective adaptation, introducing the theory behind a previously reported model [15,16]. The overall approach, though, establishes a general framework for dynamical-systems modeling and analysis of adaptive behavior in collectives. It is important to emphasize that our framework goes beyond multipopulation replicator equations and asymmetric game dynamics since it does not require a static environment (cf. refs. [17,18] for dynamic environments) and it includes the key element of the temporal loss of memory.

We model adaptation in terms of the distribution of agents' choices, developing a set of differential equations that are a continuous-time limit of a discrete-time stochastic process; cf. ref. [19]. We spend some time discussing the origin of action probabilities, since this is necessary to understand the model variables and also to clarify the limits that we invoke to arrive at our model. One is tempted to give a game-theoretic interpretation of the model and its development. For example, mixed strategies in game play are often interpreted as weights over all (complete plans of) actions. However, the game-theoretic view is inappropriate for analyzing local, myopic adaptation and the time evolution of collective behavior.

Another interpretation of our use of action probabilities comes from regarding them as frequencies of action choices. In this view, one needs long-time trials so that the frequencies take on statistical validity for an agent. Short of this, they would be dominated by fluctuations, due to undersampling. In particular, one requires that stable limit distributions exist. Moreover, the underlying deterministic dynamics of adaptation should be ergodic and have strong mixing properties. Finally, considering agent–agent interactions, one needs to assume that their adaptation is very slow compared to interaction dynamics. For rapid, say, real-time adaptation, these assumptions would be

invalid. Nonetheless, they are appropriate for long-term reinforcement, as found in learning motion through iterated exercise and learning customs through social interaction, for example.

1.2. Synopsis

The approach we take is ultimately phenomenological. We are reminded of the reaction-diffusion models of biological morphogenesis introduced originally in ref. [20]. There, the detailed processes of biological development and pattern formation were abstracted, since their biochemical basis was (and still is) largely unknown, and a behavioral phenomenology was developed on this basis. Similarly, we abstract the detailed and unknown perceptual processes that underlie agent adaptation and construct a phenomenology that captures adaptive behavior at a larger scale in agent collectives.

The phenomenology that we develop for this is one based on communications systems. Agents in a collective are confronted with the same three problems of communication posed by Weaver in the founding work of information theory—*The Mathematical Theory of Communication* [21]: (a) “How accurately can the symbols of communication be transmitted?”, (b) “How precisely do the transmitted symbols convey the desired meaning?” and (c) “How effectively does the received meaning affect conduct in the desired way?”. Shannon solved the first problem developing his theory of error-free transmission [21]. In their vocabulary adaptive agents are *information sources*. Each (a) receives information transmitted from the external environment, which includes other agents, (b) interprets the received information and modifies its internal model accordingly, and then, (c) making decisions based on the internal model, generates future behavior.

We will show that this information-theoretic view provides useful tools for analyzing collective adaptation and also an appropriate description for our assumed frequency dynamics. Using these we derive a new state space based on the self-information of agent’s actions and this allows one to investigate the dynamics of uncertainties in collective adaptation. It will become clear, though, that the assumption of global information maximization has limited relevance here, even for simple mutual adaptation in a static environment. Instead, self-organization that derives from the dynamics of uncertainties gives us a new view of collective adaptation.

To illustrate collective adaptation, we present several simulations of example environments; in particular, those having frustrated agent–agent interactions [22]. Interestingly, for two agents with perfect memory interacting via zero-sum rock-scissors-paper interactions the dynamics exhibits Hamiltonian chaos [15]. With memory loss, though, the dynamics becomes dissipative and displays the full range of nonlinear dynamical behaviors, including limit cycles, intermittency, and deterministic chaos [16].

The examples illustrate that Nash equilibria often play little or no role in collective adaptation. They are fixed points determined by the intersections of nullclines of the adaptation dynamics and sometimes the dynamics is explicitly excluded from reaching Nash equilibria, even asymptotically. Rather, it turns out that the network describing the switching between deterministic actions is a dominant factor in structuring the state-space flows. From it, much of the dynamics, including the origins of chaos becomes intuitively clear.

In the next section (Section 2), we develop a dynamical system that models adaptive behavior in collectives. In Section 3, we introduce an information-theoretic view and coordinate-transformation for adaptation dynamics and discuss self-organization induced by the dynamics of uncertainty. To illustrate the rich range of behaviors, in Section 4, we give several examples of adaptive dynamics based on non-transitive interactions. Finally, in Section 5, we interpret our results and suggest future directions.

2. Dynamics for collective adaptation

Before developing the full equations for a collective of adaptive agents, it is helpful to first describe the dynamics of how an individual agent adapts to the constraints imposed by its environment using the memory of its past behaviors. We then build up a description of how multiple agents interact, focusing only on the additional features

that come from interaction. The result is a set of coupled differential equations that determine the behavior of adaptive agent collectives and that are amenable to various kinds of geometric, statistical, and information-theoretic analyses.

2.1. Individual agent adaptation

Here we develop a continuous-time model for adaptation in an environment with a single adaptive agent. Although the behavior in this case is relatively simple, the single-agent case allows us to explain several basic points about dynamic adaptation, without the complications of a collective and agent–agent interactions. In particular, we discuss how and why we go from a discrete-time stochastic process to a *continuous-time* limit. We also describe an agent’s effective internal model of the environment and how we model its adaptation process via the temporal evolution of a *probability distribution* of action choices.

An agent takes one of N possible *actions*: $i = 1, 2, \dots, N$ at each time step τ . Let the probability for the agent to choose action i be $x_i(\tau)$, where τ is the number of steps from the initial state $x_i(0)$. The agent’s state vector—its *choice distribution*—at time τ is $\mathbf{x}(\tau) = (x_1(\tau), x_2(\tau), \dots, x_N(\tau))$, where $\sum_{n=1}^N x_n(\tau) = 1$. In the following we call the temporal behavior of $\mathbf{x}(\tau)$ as the *dynamics of adaptation*.

Let $r_i(\tau)$ denote the reinforcement the agent receives for its taking action i at step τ . Denote the collection of these by the vector $\mathbf{r}(\tau) = (r_1(\tau), \dots, r_N(\tau))$. The agent’s memories—denoted $\mathbf{Q}(\tau) = (Q_1(\tau), \dots, Q_N(\tau))$ —of past rewards from its actions are updated according to

$$Q_i(\tau + 1) - Q_i(\tau) = \frac{1}{T}[\delta_i(\tau)r_i(\tau) - \alpha Q_i(\tau)], \quad (1)$$

where

$$\delta_i(\tau) = \begin{cases} 1, & \text{action } i \text{ chosen at step } \tau \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

with $i = 1, \dots, N$. T is a constant that sets the agent–environment interaction time scale. $\alpha \in [0, 1)$ controls the agent’s memory loss rate. For $\alpha = 0$, the agent has a perfect memory as the sum of the past reinforcements; for $\alpha > 0$ the memory is attenuated in that older reinforcements have less effect on the current Q_i s and more recent reinforcements are given larger weight. One imagines that the agent constructs a histogram of past reinforcements and this serves as a simple internal memory of its environment.

An agent chooses its next action according to its choice distribution which is updated from the reinforcement memory according to:

$$x_i(\tau) = \frac{e^{\beta Q_i(\tau)}}{\sum_{n=1}^N e^{\beta Q_n(\tau)}}, \quad (3)$$

where $i = 1, 2, \dots, N$. $\beta \in [0, \infty]$ controls the adaptation rate: how much the choice distribution is changed by the memory of past reinforcements. For example, if $\beta = 0$, the choice distribution is unaffected by past reinforcements. Specifically, it becomes independent of \mathbf{Q} and one has $x_i(\tau) = 1/N$. In this case, the agent chooses actions with uniform probability and so behaves completely randomly. In a complementary fashion, in the limit $\beta \rightarrow \infty$, an agent chooses that action i with the maximum $Q_i(\tau)$ and $x_i(\tau) \rightarrow 1$.

Given Eq. (3) the time evolution of agent’s choice distribution is:

$$x_i(\tau + 1) = \frac{x_i(\tau)e^{\beta(Q_i(\tau+1) - Q_i(\tau))}}{\sum_{n=1}^N x_n(\tau)e^{\beta(Q_n(\tau+1) - Q_n(\tau))}}, \quad (4)$$

where $i = 1, 2, \dots, N$. Eqs. (1)–(4) determine how the agent adapts its choice distribution using reinforcements it has received from the environment for its past actions.

This simple kind of adaptation was introduced as a principle of behavioral learning [23,24] and as a model of stochastic learning [25]. It is sometimes referred to as reinforcement learning [26,27]. Arguably, it is the simplest form of adaptation in which an agent develops relationships or behavior patterns through reinforcements from external stimuli.

Starting with the discrete-time model above, one can develop a continuous-time model that corresponds to the agent performing a large number of actions, iterates of Eq. (1), for each choice distribution update, iterates of Eq. (3). Thus, we recognize two different time scales: one for agent–environment interactions and one for adaptation of the agent’s internal model based on its internal memory. We assume that the adaptation dynamics is very slow compared to interactions and so \mathbf{x} is essentially constant during interactions; see Fig. 1.

Starting from Eq. (1), one can show that the continuous-time dynamics of memory updates is given by the differential equations

$$\dot{Q}_i(t) = R_i(t) - \alpha Q_i(t), \quad (5)$$

with $i = 1, 2, \dots, N$; see Appendix A. Here R_i is the reward the environment gives to the agent choosing action i : the average of $r_i(\tau)$ during the time interval between updates of \mathbf{x} at t and $t + dt$.

From Eq. (3) one sees that the map from $\mathbf{Q}(t)$ to $\mathbf{x}(t)$ at time t is given by

$$x_i(t) = \frac{e^{\beta Q_i(t)}}{\sum_{n=1}^N e^{\beta Q_n(t)}}, \quad (6)$$

where $i = 1, 2, \dots, N$. Differentiating Eq. (6) gives the continuous-time dynamics

$$\dot{x}_i(t) = \beta x_i(t) \left(\dot{Q}_i(t) - \sum_{n=1}^N \dot{Q}_n(t) x_n(t) \right), \quad (7)$$

with $i = 1, 2, \dots, N$.

Assembling Eqs. (5)–(7), one finds the basic dynamic that governs agent behavior on the adaptation time-scale:

$$\frac{\dot{x}_i}{x_i} = \beta(R_i - R) + \alpha(H_i - H) \equiv \beta \Delta R_i + \alpha \Delta H_i, \quad (8)$$

where $i = 1, 2, \dots, N$. Here

$$R = \sum_{n=1}^N x_n R_n \quad (9)$$

is the net reinforcement averaged over the agent’s possible actions. And,

$$H_i = -\log x_i \quad (10)$$




Fig. 1. The time scale (τ) of a single agent interacting with its environment and the time scale (t) of the agent’s adaptation: $\tau \ll t$.

where $i = 1, 2, \dots, N$, is the *self-information* or degree of surprise when the agent takes action i [21]. The average self-information, or *Shannon entropy* of the choice distribution, also appears as

$$H = \sum_{n=1}^N x_n H_n = - \sum_{n=1}^N x_n \log x_n. \quad (11)$$

These are the entropies of the agent's choice distribution measured, not in *bits* (binary digits), but in *nats* (natural digits), since the natural logarithm is used. The entropy measures the choice distribution's flatness, being maximized when the choices all have equal probability.

Fortunately, the basic dynamic captured by Eq. (8) is quite intuitive, being the balance of two terms on the right-hand side. The first term describes adaptation dynamics, whose time scale is controlled by β . The second describes the loss of memory with a time scale controlled by α . That is, the adaptation in choice probabilities is driven by a balance between two forces: the tendency to concentrate the choice probability based on the reinforcement $\mathbf{R} = (R_1, R_2, \dots, R_N)$ and the tendency to make choices equally likely. Finally, on the lefthand side, one has the logarithmic derivative of the choice probabilities: $\dot{x}_i/x_i = d/dt(\log x_i)$.

Note that each of the terms on the righthand side is a difference between a function of a particular choice and that function's average. Specifically, the first term $\Delta R_i \equiv R_i - R$ is the relative benefit in choosing action i compared to the mean reinforcement across all choices. Other things being held constant, if this term is positive, then action i is the better choice compared to the mean and x_i will increase. The second term $\Delta H_i \equiv H_i - H$ is the relative informativeness of taking action i compared to the average H , that is Shannon entropy. Thus, x_i decreases in proportion to the entropy at time t and so this term works to increase the uncertainty of agent's actions, flattening the choice distribution by increasing the probability of unlikely actions. When $x_i = N^{-1}$, the distribution is flat (purely random choices), $\Delta H = 0$, and memory loss effects disappear.

Mathematically, the adaptation equations have quite a bit of structure and this has important consequences, as we will see. Summarizing, the adaptation equations describe a dynamic that balances the tendency to concentrate on choices associated with the best action against the tendency to make the choices equally likely. The net result is to increase the choice uncertainty, subject to the constraints imposed by the environment via the reinforcements. Thus, the choice distribution is the least biased distribution consistent with environmental constraints and individual memory loss. We will return to discuss this mechanism in detail using information theory in Section 3 (Fig. 2).

Since the reinforcement determines the agent's interactions with the environment, there are, in fact, three different time scales operating: that for agent–environment interactions, that for each agent's adaptation, and that for changes to the environment. However, if the environment changes very slowly compared to the agent's internal adaptation, the environment $r_i(t)$ can be regarded as effectively constant, as shown in Fig. 3.

In this case $r_i(t)$ can be approximated as a static relationship between an agent's actions and the reinforcements given by the environment. Let $r_i(t) = a_i$, where $\mathbf{a} = (a_1, \dots, a_N)$ are constants that are normalized: $\sum_{n=1}^N a_n = 0$. Given this, the agent's time-averaged reinforcements are a_i ($R_i = a_i$) and the continuous-time dynamics simplifies




Fig. 2. A dynamic balance of adaptation and memory loss: adaptation concentrates the probability distribution on the best action. Memory loss of past history leads to a distribution that is flatter and has higher entropy.




Fig. 3. The time scales of dynamic adaptation: agent adaptation is slow compared to agent–environment interaction and environmental change is slower still compared to adaptation.




Fig. 4. Dynamics of single-agent adaptation: here there are three actions, labeled 1–3, and the environment gives reinforcements according to $\mathbf{a} = (\frac{2}{3}\epsilon, -1 - \frac{1}{3}\epsilon, 1 - \frac{1}{3}\epsilon)$. The figure shows two trajectories from simulations with $\epsilon = 0.5$ and $\beta = 0.1$ and with $\alpha = 0.0$ (right) and $\alpha = 0.3$ (left).

to:

$$\frac{\dot{x}_i}{x_i} = \beta \left(a_i - \sum_{n=1}^N a_n x_n \right) + \alpha \left(-\log x_i + \sum_{n=1}^N x_n \log x_n \right), \quad (12)$$

where $i = 1, 2, \dots, N$.

The behavior of single-agent adaptation given by Eq. (12) is very simple. When α is small, so that adaptation is dominant $x_i \rightarrow 1$, where i is the action with the highest reward a_i , and $x_j \rightarrow 0$ for $j \neq i$. The agent receives this information from the fixed environment and its behavior is simply to choose the action with the maximum reward and the choice distribution moves to the associated simplex vertex $\mathbf{x}^* = (0, \dots, 1, \dots, 0)$. In the special case when $\alpha = 0$, it is known that for arbitrary \mathbf{a} Eq. (12) moves \mathbf{x} to the vertex corresponding to the maximum a_i [2]. In a complementary way, when α is large enough to overcome the relative differences in reinforcements—that is, when $\beta/\alpha \rightarrow 0$ —memory loss dominates, the agent states goes to a uniform choice distribution ($x_i = N^{-1}$) and the system converges to the simplex center. Note that in machine learning this balance between local optimization and randomized behavior, which selects non-optimal actions, is referred to as the *exploitation–exploration trade-off* [27].

For instance, consider an agent that takes $N = 3$ actions $\{1, 2, 3\}$ in an environment described by $\mathbf{a} = (\frac{2}{3}\epsilon, -1 - \frac{1}{3}\epsilon, 1 - \frac{1}{3}\epsilon)$, with $\epsilon \in [-1, 1]$. In the perfect memory case ($\alpha = 0$), the choice distribution converges to a stable fixed point $(0,0,1)$. $\mathbf{x}^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ is an unstable hyperbolic fixed point. In the memory loss case ($\alpha > 0$), dynamics converges to a stable fixed point inside the simplex. (These cases are illustrated in Fig. 4.)

Even when the environment is time-dependent, the agent's behavior can track the highest-reward action as long as the time scale of environment change is slow compared to the agent's adaptation. However, the situation is more interesting when environment change occurs at a rate near the time-scale set by adaptation. Mutual adaptation in agent collectives, the subject of the following sections, corresponds to just this situation. Other agents provide, through their own adaptation, a dynamic environment to any given agent and if their time scales of adaptation are close, the dynamics can be quite rich and difficult to predict and analyze.

2.2. Two agent adaptation

To develop equations of motion for adaptation in an agent collective we initially assume, for simplicity, that there are only two agents. The agents, denoted X and Y , at each moment take one of N or M actions, respectively. The agents states at time t are $\mathbf{x} = (x_1, \dots, x_N)$ and $\mathbf{y} = (y_1, \dots, y_M)$, with $\sum_{n=1}^N x_n = \sum_{m=1}^M y_m = 1$. $\mathbf{x}(0)$ and $\mathbf{y}(0)$ are the initial conditions. We view the time evolution of each agent's state vector in the simplices $\mathbf{x} \in \Delta_X$ and $\mathbf{y} \in \Delta_Y$ and the group dynamics in the *collective* state space Δ which is the product of the agent simplices:

$$\mathbf{X} = (\mathbf{x}, \mathbf{y}) \in \Delta = \Delta_X \times \Delta_Y. \quad (13)$$

There are again three different time scales to consider: one for agent–agent interaction, one for each agent's internal adaptation, and one for the environment which now mediates agent interactions via the reinforcements given to the agents. The external environment controls, for example, the degree of coupling between the agents. In contrast with the single-agent case, in the many-agent setting each agent's behavior produces a *dynamic environment* for the other agents. This environmental dynamics is particularly important when the adaptation time scales of each agent are close.

Following the single-agent case, though, we assume that the adaptation dynamics is very slow compared to that of agent–agent interactions and that the dynamics of the external environment changes very slowly compared to that of agents' mutual adaptation. Under these assumptions the agent state vectors \mathbf{x} and \mathbf{y} are effectively constant during the agent–agent interactions that occur between adaptation updates. The immediate consequence is that one can describe the collective state space in terms of the frequencies of actions (the choice distributions). Additionally, the environment is essentially constant relative to changes in the states \mathbf{x} and \mathbf{y} .

Denote the agents' memories by $\mathbf{Q}^X = (Q_1^X, \dots, Q_N^X)$ for X and $\mathbf{Q}^Y = (Q_1^Y, \dots, Q_M^Y)$ for Y . For the dynamic governing memory updates we have

$$\begin{aligned} Q_i^X(\tau + 1) - Q_i^X(\tau) &= \frac{1}{T} [\sum_j \delta_{ij}(\tau) r_{ij}^X(\tau) - \alpha_X Q_i^X(\tau)], \\ Q_j^Y(\tau + 1) - Q_j^Y(\tau) &= \frac{1}{T} [\sum_i \delta_{ij}(\tau) r_{ji}^Y(\tau) - \alpha_Y Q_j^Y(\tau)], \end{aligned} \quad (14)$$

where

$$\delta_{ij}(\tau) = \begin{cases} 1, & \text{pair of actions } (i, j) \text{ chosen at step } \tau \\ 0, & \text{otherwise} \end{cases} \quad (15)$$

with $i = 1, \dots, N$ and $j = 1, \dots, M$. T is the time constant, as before. Then the continuous-time dynamics of memory updates for X and Y are given by the differential equations

$$\begin{aligned} \dot{Q}_i^X &= R_i^X - \alpha_X Q_i^X, \\ \dot{Q}_j^Y &= R_j^Y - \alpha_Y Q_j^Y, \end{aligned} \quad (16)$$

for $i = 1, 2, \dots, N$ and $j = 1, 2, \dots, M$. R_i^X is the reward for agent X choosing action i , averaged over agent Y 's actions between adaptive updates; and R_j^Y is Y 's. The parameters $\alpha_X, \alpha_Y \in [0, 1)$ control each agent's memory loss rate, respectively.

The map from $\mathbf{Q}^X(t)$ to $\mathbf{x}(t)$ and from $\mathbf{Q}^Y(t)$ to $\mathbf{y}(t)$ at time t is

$$\begin{aligned} x_i(t) &= \frac{e^{\beta_X Q_i^X(t)}}{\sum_{n=1}^N e^{\beta_X Q_n^X(t)}}, \\ y_j(t) &= \frac{e^{\beta_Y Q_j^Y(t)}}{\sum_{m=1}^M e^{\beta_Y Q_m^Y(t)}}, \end{aligned} \quad (17)$$

for $i = 1, \dots, N$ and $j = 1, \dots, M$. Here $\beta_X, \beta_Y \in [0, \infty]$ control the agents' adaptation rates, respectively. Differentiating Eq. (17) with respect to t , the continuous-time adaptation for two agents is governed by

$$\begin{aligned} \dot{x}_i &= \beta_X x_i \left(\dot{Q}_i^X - \sum_{n=1}^N \dot{Q}_n^X x_n \right), \\ \dot{y}_j &= \beta_Y y_j \left(\dot{Q}_j^Y - \sum_{m=1}^M \dot{Q}_m^Y y_m \right), \end{aligned} \quad (18)$$

for $i = 1, \dots, N$ and $j = 1, \dots, M$.

Putting together Eqs. (16)–(18), one finds the coupled adaptation equations for two agents:

$$\begin{aligned} \frac{\dot{x}_i}{x_i} &= \beta_X (R_i^X - R^X) + \alpha_X (H_i^X - H^X) \equiv \beta_X \Delta R_i^X + \alpha_X \Delta H_i^X, \\ \frac{\dot{y}_j}{y_j} &= \beta_Y (R_j^Y - R^Y) + \alpha_Y (H_j^Y - H^Y) \equiv \beta_Y \Delta R_j^Y + \alpha_Y \Delta H_j^Y, \end{aligned} \quad (19)$$

for $i = 1, \dots, N$ and $j = 1, \dots, M$ and where

$$\begin{aligned} R^X &= \sum_{n=1}^N x_n R_n^X, \quad R^Y = \sum_{m=1}^M y_m R_m^Y, \\ H^X &= \sum_{n=1}^N x_n H_n^X, \quad H^Y = \sum_{m=1}^M y_m H_m^Y. \end{aligned} \quad (20)$$

The interpretations of the $\Delta R = R_i - R$ and $\Delta H = H_i - H$ terms are not essentially different from those introduced to describe the single-agent case. That is, the behavior of each agent is a dynamic balance between (i) adaptation—concentrating the choice probability on the best action at t —and (ii) memory loss—increasing the choice uncertainty. What is new here is that there are two (and eventually more) agents attempting to achieve this balance together using information that comes from their interactions with the dynamic environment which includes other agents.

As given, the adaptation equations include the possibility of a time-dependent environment, which would be implemented, say, using a time-dependent reinforcement scheme. However, as with the single-agent case, it is helpful to simplify the model by assuming a static external environment and, in particular, static relationships between the agents.

Assume that the external environment changes slowly compared to the dynamics of mutual adaptation, as illustrated in Fig. 3. This implies a nearly static relationship between pairs of action choices (i, j) and reinforcements r_{ij}^X and r_{ji}^Y for both agents. Since the environmental dynamics is very slow compared to each agents' adaptation,

$r_{ij}^X(t)$ and $r_{ji}^Y(t)$ are essentially constant during adaptation. The r s can be approximated then as constant:

$$\begin{aligned} r_{ij}^X(t) &= a_{ij}, \\ r_{ji}^Y(t) &= b_{ji}, \end{aligned} \quad (21)$$

for $i = 1, \dots, N$ and $j = 1, \dots, M$. a_{ij} and b_{ji} are normalized over j and i so that when summing over all actions the reinforcements vanish:

$$\begin{aligned} \sum_{n=1}^N a_{nj} &= 0, \\ \sum_{m=1}^M b_{mi} &= 0. \end{aligned} \quad (22)$$

Given the form of ΔR in the adaptation equations, this normalization does not affect the dynamics.

Assume further that \mathbf{x} and \mathbf{y} are independently distributed. This is equivalent to agents never having a global view of the collective or their interactions with the environment (other agents). Each agent's knowledge of the environment is uncorrelated, at each moment, with the state of the other agents. The time-average rewards for X and Y now become

$$\begin{aligned} R_i^X &= \sum_{m=1}^M a_{im} y_m = (\mathbf{A}\mathbf{y})_i, \\ R_j^Y &= \sum_{n=1}^N b_{jn} x_n = (\mathbf{B}\mathbf{x})_j, \end{aligned} \quad (23)$$

for $i = 1, \dots, N$ and $j = 1, \dots, M$. In this restricted case, the continuous-time dynamic is given by the coupled adaptation equations

$$\begin{aligned} \frac{\dot{x}_i}{x_i} &= \beta_X [(\mathbf{A}\mathbf{y})_i - \mathbf{x} \cdot \mathbf{A}\mathbf{y}] + \alpha_X \left[-\log x_i + \sum_{n=1}^N x_n \log x_n \right], \\ \frac{\dot{y}_j}{y_j} &= \beta_Y [(\mathbf{B}\mathbf{x})_j - \mathbf{y} \cdot \mathbf{B}\mathbf{x}] + \alpha_Y \left[-\log y_j + \sum_{m=1}^M y_m \log y_m \right]. \end{aligned} \quad (24)$$

for $i = 1, \dots, N$ and $j = 1, \dots, M$. A is an $N \times M$ matrix and B is a $M \times N$ matrix with $(A)_{ij} = a_{ij}$ and $(B)_{ji} = b_{ji}$, respectively. $\mathbf{x} \cdot \mathbf{A}\mathbf{y}$ is the inner product between \mathbf{x} and $\mathbf{A}\mathbf{y}$ and similarly for $\mathbf{y} \cdot \mathbf{B}\mathbf{x}$:

$$\begin{aligned} \mathbf{x} \cdot \mathbf{A}\mathbf{y} &= \sum_{n=1}^N \sum_{m=1}^M a_{nm} x_n y_m, \\ \mathbf{y} \cdot \mathbf{B}\mathbf{x} &= \sum_{m=1}^M \sum_{n=1}^N b_{mn} y_m x_n. \end{aligned} \quad (25)$$

2.3. Collective adaptation

Generalizing to an arbitrary number of agents at this point should appear straightforward. It simply requires extending Eqs. (19) to a collection of adaptive agents. Suppose there are S agents labeled $s = 1, 2, \dots, S$ and each agent can take one of N^s actions. One describes the time evolution of the agents' state vectors in the simplices $\mathbf{x}^1 \in \Delta_1$, $\mathbf{x}^2 \in \Delta_2$, \dots , and $\mathbf{x}^S \in \Delta_S$. The adaptation dynamics in the higher-dimensional *collective* state space occurs within

$$\mathbf{X} = (\mathbf{x}^1, \mathbf{x}^2, \dots, \mathbf{x}^S) \in \Delta = \Delta_1 \times \Delta_2 \times \dots \Delta_S. \quad (26)$$

Then we have the dynamics for collective adaptation as

$$\frac{\dot{x}_{is}^s}{x_{is}^s} = \beta_s(R_{is}^s - R^s) + \alpha_s(H_{is}^s - H^s) \equiv \beta_s \Delta R_{is}^s + \alpha_s \Delta H_{is}^s. \quad (27)$$

for $i^s = 1, \dots, N^s$ and $s = 1, \dots, S$. R_{is}^s and H_{is}^s are the reinforcement and the self-information for s to choose action i^s , respectively. Eqs. (27) constitute our general model for adaptation in agents collective.

With three agents X , Y , and Z , the collective state space is

$$\mathbf{X} = (\mathbf{x}, \mathbf{y}, \mathbf{z}) \in \Delta = \Delta_X \times \Delta_Y \times \Delta_Z, \quad (28)$$

one obtains:

$$\begin{aligned} \frac{\dot{x}_i}{x_i} &= \beta_X(R_i^X - R^X) + \alpha_X[H_i^X - H^X] \equiv \beta_X \Delta R_i^X + \alpha_X \Delta H_i^X, \\ \frac{\dot{y}_j}{y_j} &= \beta_Y(R_j^Y - R^Y) + \alpha_Y[H_j^Y - H^Y] \equiv \beta_Y \Delta R_j^Y + \alpha_Y \Delta H_j^Y, \\ \frac{\dot{z}_k}{z_k} &= \beta_Z(R_k^Z - R^Z) + \alpha_Z[H_k^Z - H^Z] \equiv \beta_Z \Delta R_k^Z + \alpha_Z \Delta H_k^Z, \end{aligned} \quad (29)$$

for $i = 1, \dots, N$, $j = 1, \dots, M$, and $k = 1, \dots, L$. The static environment version reduces to

$$\begin{aligned} \frac{\dot{x}_i}{x_i} &= \beta_X[(A\mathbf{y}\mathbf{z})_i - \mathbf{x} \cdot A\mathbf{y}\mathbf{z}] + \alpha_X \left[-\log x_i + \sum_{n=1}^N x_n \log x_n \right], \\ \frac{\dot{y}_j}{y_j} &= \beta_Y[(B\mathbf{z}\mathbf{x})_j - \mathbf{y} \cdot B\mathbf{z}\mathbf{x}] + \alpha_Y \left[-\log y_j + \sum_{m=1}^M y_m \log y_m \right], \\ \frac{\dot{z}_k}{z_k} &= \beta_Z[(C\mathbf{x}\mathbf{y})_k - \mathbf{z} \cdot C\mathbf{x}\mathbf{y}] + \alpha_Z \left[-\log z_k + \sum_{l=1}^L z_l \log z_l \right], \end{aligned} \quad (30)$$

for $i = 1, \dots, N$, $j = 1, \dots, M$, and $k = 1, \dots, L$, and with tensors $(A)_{ijk} = a_{ijk}$, $(B)_{jki} = b_{jki}$, and $(C)_{kij} = c_{kij}$. Here

$$(A\mathbf{y}\mathbf{z})_i = \sum_{m=1}^M \sum_{l=1}^L a_{iml} y_m z_l \quad (31)$$

and

$$\mathbf{x} \cdot A\mathbf{y}\mathbf{z} = \sum_{n=1}^N \sum_{m=1}^M \sum_{l=1}^L a_{nml} x_n y_m z_l \quad (32)$$

and similarly for Y and Z . Note that the general model includes heterogeneous network settings with local interactions rather than global interactions; see Appendix B.

2.4. Evolutionary dynamics and game theory

We now interrupt the development to discuss the connections between these models and those from population dynamics and game theory. There are interesting connections and also several important distinctions that need to be kept in mind, before we can move forward.

The special case that allows us to make contact with evolutionary dynamics and game theory is the restriction to agents with perfect memory interacting in a static environment. (For further details see Appendix C.) In the

two agent, static external environment case we set $\alpha_X = \alpha_Y = 0$ and equal adaptation rates, $\beta_X = \beta_Y$. Under these assumptions our model, Eqs. (24), reduces to what is either called *multipopulation replicator equations* [13] or *asymmetric game dynamics* [9,10,13]. The equations are:

$$\begin{aligned}\frac{\dot{x}_i}{x_i} &= (A\mathbf{y})_i - \mathbf{x} \cdot A\mathbf{y}, \\ \frac{\dot{y}_j}{y_j} &= (B\mathbf{x})_j - \mathbf{y} \cdot B\mathbf{x}.\end{aligned}\tag{33}$$

From the perspective of game theory, one regards the interactions determined by A and B as X 's and Y 's *payoff matrices*, respectively, for a linear game in which X plays action i against Y 's action j . Additionally, \mathbf{x} and \mathbf{y} , the agent state vectors, are interpreted as *mixed strategies*. In fact, $\mathbf{x} \cdot A\mathbf{y}$ and $\mathbf{y} \cdot B\mathbf{x}$ in Eqs. (33) formally satisfy von Neumann–Morgenstern utilities [11]. If they exist in the interior of the collective simplices Δ_X and Δ_Y , interior Nash equilibria of the game (A, B) are the fixed points determined by the intersections of the x - and y -nullclines of Eqs. (33).

One must be careful, though, in drawing parallels between our general dynamic setting and classical game theory. In idealized economic agents, it is often assumed that agents have knowledge of the entire game structure and of other agents' decision-making processes. Game theory's central methodology derives from how these *rational players* should act. Our adaptive agents, in contrast, have no knowledge of a game in which they might be playing, only a myopic model of the environment and, even then, this is given only implicitly via the reinforcements the agents receive from the environment. In particular, the agents do not know whether they are playing a game or not, how many agents there are beyond themselves, or even whether other agents exist or not. Our model of dynamic adaptation under such constraints is appropriate nonetheless for many real world adaptive systems, whether animal, human, or economic agent collectives [28]. The bi-matrix game (A, B) appears above as a description of the collective's global dynamic only under the assumption that the external environment changes very slowly.

The connection with evolutionary dynamics is formal and comes from the fact that Eqs. (33) are the well known replicator equations of population dynamics [2]. However, the interpretation of the variables is rather different. Population dynamics views \mathbf{x} and \mathbf{y} as two separate, but interacting (infinite population size) groups. These two populations are described as distributions of various organismal phenotypes. The equations of motion determine the evolution of these populations over generations and through interaction. In our model, in contrast, \mathbf{x} and \mathbf{y} represent the probability to choose actions for each agent. The equations of motion describe their dynamic adaptation to each other through interaction.

Despite the similarities that one can draw in this special case, it is important to emphasize that our framework goes beyond the multipopulation replicator equations and asymmetric game dynamics. First, the reinforcement scheme \mathbf{R} need not lead to linear interactions. Second, the model does not require a static environment described by a constant bi-matrix (A, B) . Finally, the occurrence of the memory loss term is entirely new and not found in game theory or evolutionary dynamics.

3. Information, uncertainty, and dynamic adaptation

We now shift away from a dynamical systems view and, as promised earlier, begin to think of the agent collective as a communication network. Although initially this will appear unrelated, we will show that there is a close connection between the dynamical and information-theoretic perspectives—connections that have both mathematical and pragmatic consequences.

We consider the adaptive agents in the collective to be *information sources*. First, each agent receives information from its environment, which includes other agents. Next, they interpret the received information and modify their

behavior (and so the structure of information source) accordingly, changing from $\mathbf{x}(t)$ to $\mathbf{x}(t + dt)$. Finally, they generate a series of messages (actions) based on their updated internal model, introducing the new behavior back into the environment. This is a different interpretation of the interaction process in the collective which we motivated up to this point only as a dynamical process. Now we discuss the adaptive dynamics from information theoretic viewpoint.

3.1. Dynamics in information space

In this section, we introduce a new state space that directly represents the uncertainties of agent actions. First, as before, for clarity we focus on the two-agent static-environment case, Eqs. (24). Since the components of the agents' states are probabilities, the quantities

$$\begin{aligned}\xi_i &= -\log x_i, \\ \eta_j &= -\log y_j,\end{aligned}\tag{34}$$

are the *self-informations* of agents X and Y choosing actions i and j , respectively. When x_i is small, for example, the self-information ξ_i is large since action i is rarely chosen by agent X . Consider the resulting change in coordinates in $\mathbf{R}_+^N \times \mathbf{R}_+^M$:

$$\mathcal{E} = (\boldsymbol{\xi}, \boldsymbol{\eta}) = (\xi_1, \dots, \xi_N) \times (\eta_1, \dots, \eta_M).\tag{35}$$

The normalization conditions— $\sum_{n=1}^N x_n = \sum_{m=1}^M y_m = 1$ —that restrict the agent states to lie in simplices become $\sum_{n=1}^N e^{-\xi_n} = \sum_{m=1}^M e^{-\eta_m} = 1$ in \mathcal{E} .

In this space the equations of motion become:

$$\begin{aligned}\dot{\xi}_i &= -\beta_X[(A \exp(-\boldsymbol{\eta}))_i - \exp(-\boldsymbol{\xi}) \cdot A \exp(-\boldsymbol{\eta})] - \alpha_X[\xi_i - \exp(-\boldsymbol{\xi}) \cdot \boldsymbol{\xi}], \\ \dot{\eta}_j &= -\beta_Y[(B \exp(-\boldsymbol{\xi}))_j - \exp(-\boldsymbol{\eta}) \cdot B \exp(-\boldsymbol{\xi})] - \alpha_Y[\eta_j - \exp(-\boldsymbol{\eta}) \cdot \boldsymbol{\eta}],\end{aligned}\tag{36}$$

for $i = 1, \dots, N$ and $j = 1, \dots, M$ and where $\exp(-\boldsymbol{\xi}) = (e^{-\xi_1}, \dots, e^{-\xi_N})$ and $\exp(-\boldsymbol{\eta}) = (e^{-\eta_1}, \dots, e^{-\eta_M})$.

Recall that both the ΔR interaction term and the ΔH memory-loss term are differences from means. This suggests yet another transformation to remove these comparisons to the mean:

$$\begin{aligned}u_i &= \xi_i - N^{-1} \sum_{n=1}^N \xi_n, \\ v_j &= \eta_j - M^{-1} \sum_{m=1}^M \eta_m,\end{aligned}\tag{37}$$

with $i = 1, \dots, N$ and $j = 1, \dots, M$. This leads to the normalized space in $\mathbf{R}^N \times \mathbf{R}^M$:

$$\mathbf{U} = (\mathbf{u}, \mathbf{v}) = (u_1, \dots, u_N) \times (v_1, \dots, v_M),\tag{38}$$

with the constraints $\sum_{n=1}^N u_n = \sum_{m=1}^M v_m = 0$. \mathbf{u} and \mathbf{v} are the *normalized self-informations* relative to their means. We refer to this state space as *information space*.

The combined coordinate transformation, Eqs. (37) composed with Eqs. (34), gives the well known *centered log-ratio* coordinates [29]. The inverse transformation is:

$$\begin{aligned} x_i &= \frac{e^{-u_i}}{\sum_{n=1}^N e^{-u_n}}, \\ y_j &= \frac{e^{-v_j}}{\sum_{m=1}^M e^{-v_m}}, \end{aligned} \quad (39)$$

or, by using the notation $\exp(-\mathbf{u}) = (e^{-u_1}, \dots, e^{-u_N})$ and $\exp(-\mathbf{v}) = (e^{-v_1}, \dots, e^{-v_M})$,

$$\begin{aligned} \mathbf{x} &= Z^{-1}(\mathbf{u}) \exp(-\mathbf{u}), \\ \mathbf{y} &= Z^{-1}(\mathbf{v}) \exp(-\mathbf{v}), \end{aligned} \quad (40)$$

where $Z(\mathbf{u}) = \sum_{n=1}^N e^{-u_n}$ and $Z(\mathbf{v}) = \sum_{m=1}^M e^{-v_m}$. The resulting transformed adaptation equations directly model the dynamics of uncertainties of agents' behavior:

$$\begin{aligned} \dot{\mathbf{u}} &= -\beta_X Z^{-1}(\mathbf{v}) \left[A \exp(-\mathbf{v}) - \sum_{n=1}^N (A \exp(-\mathbf{v}))_n \right] - \alpha_X \mathbf{u}, \\ \dot{\mathbf{v}} &= -\beta_Y Z^{-1}(\mathbf{u}) \left[B \exp(-\mathbf{u}) - \sum_{m=1}^M (B \exp(-\mathbf{u}))_m \right] - \alpha_Y \mathbf{v}. \end{aligned} \quad (41)$$

When the interaction matrices are normalized to zero mean, $\sum_{m=1}^M a_{im} = \sum_{n=1}^N b_{jn} = 0$, the equations simplify even further to

$$\begin{aligned} \dot{\mathbf{u}} &= -\beta_X A [Z^{-1}(\mathbf{v}) \exp(-\mathbf{v})] - \alpha_X \mathbf{u}, \\ \dot{\mathbf{v}} &= -\beta_Y B [Z^{-1}(\mathbf{u}) \exp(-\mathbf{u})] - \alpha_Y \mathbf{v}. \end{aligned} \quad (42)$$

The origin $\mathbf{O} = (0, 0, \dots, 0)$ of the normalized information space \mathbf{U} corresponds to random behavior: $(\mathbf{x}, \mathbf{y}) = (1/N, \dots, 1/N, 1/M, \dots, 1/M)$. The Shannon entropy of the choice distribution is maximized at this point. In contrast, when agents choose an action with probability 1, the entropy vanishes and the agent state is located in Δ at the simplex vertices and in \mathbf{U} at infinity.

Eqs. (42) directly describe the dynamics of uncertainties of individual behavior interacting with other agents' uncertainties. In Eqs. (42) the first term comes from mutual adaptation—uncertainty change by interacting with other information sources with interaction matrices (A, B) . The second term comes from internal memory loss—uncertainty gain based on past self-information of their own behavior. Eqs. (42) are useful for theory, for analysis in certain limits, as we will shortly demonstrate, and for numerical stability during simulation, which we will illustrate when considering example collectives below. Note that Eqs. (24), (36), and (41) are topologically orbit equivalent.

3.2. Self-organization induced by the dynamics of uncertainty

Eqs. (42) describe the dynamics of uncertainty between deterministic and random behavior in the information space \mathbf{U} . Uncertainties change when the agents adapt to environmental constraints and accordingly modify their choice distribution. Uncertainties increase when memory loss dominates and agents behave more randomly with less regard to the environmental constraints. The dissipation rate γ of the dynamics in \mathbf{U} is controlled entirely by




Fig. 5. Dynamics of zero-sum interaction without memory loss: constant of motion $E = \beta_X^{-1} D(\mathbf{x}^* || \mathbf{x}) + \beta_Y^{-1} D(\mathbf{y}^* || \mathbf{y})$ keeps the linear sum of the distances between the interior Nash equilibrium and each agent's state.

the memory loss rate α :

$$\gamma = \sum_{n=1}^N \frac{\partial \dot{u}_n}{\partial u_n} + \sum_{m=1}^M \frac{\partial \dot{v}_m}{\partial v_m} = -N\alpha_X - M\alpha_Y. \quad (43)$$

Therefore, Eqs. (42) are volume preserving in \mathbf{U} when $\alpha_X = \alpha_Y = 0$.

In the case that agents behave without memory loss ($\alpha_X = \alpha_Y = 0$), if the interaction specified by (A, B) is zero-sum, $B = -A^T$, and if, in addition, it determines an interior Nash equilibrium $(\mathbf{x}^*, \mathbf{y}^*)$ (see Appendix C), then the collective has a constant of motion:

$$E = \beta_X^{-1} D(\mathbf{x}^* || \mathbf{x}) + \beta_Y^{-1} D(\mathbf{y}^* || \mathbf{y}), \quad (44)$$

where $D(\mathbf{p} || \mathbf{q}) = \sum_k p_k \log(p_k/q_k)$ is the *relative entropy* or the *information gain* which measures the similarity between probability distributions \mathbf{p} and \mathbf{q} [30]. (Appendix D gives the derivation of Eq. (44).)

Since the constant of motion E is a linear sum of relative entropies, the collective maintains the information-theoretic distance between the interior Nash equilibrium and each agent's state. Thus, in the perfect memory case ($\alpha = 0$), by the inequality $D(\mathbf{p} || \mathbf{q}) \geq 0$, the interior Nash equilibrium cannot be reached unless the initial condition itself starts there (Fig. 5). This is an information-theoretic interpretation of the constant of motion noted in ref. [31]. Moreover, when $N = M$ the dynamics has a symplectic structure in \mathbf{U} with the Hamiltonian E given in Eq. (44) [31]. In this case, Eqs. (41) are described quite simply,

$$\dot{\mathbf{U}} = J \nabla_{\mathbf{U}} E, \quad (45)$$

with a Poisson structure

$$J = \begin{pmatrix} O & P \\ -P^T & O \end{pmatrix} \quad \text{with} \quad P = -\beta_X \beta_Y A. \quad (46)$$

Again, see Appendix D.

When the bi-matrix interaction (A, B) satisfies $B = A^T$, in generic cases where there's no payoff tie, E is a Lyapunov function of the dynamics and decreases to 0 over time [2]. In this case, each agent can adapt to the environment independently and the collective adaptation dynamics reaches one of the stable states. In some circumstances, though, the Nash equilibria $(\mathbf{x}^*, \mathbf{y}^*)$ may not be in the interior of the collective simplex Δ . (Note that symmetric artificial neural networks have similar properties [32].)

In some cases when neither $B = -A^T$ nor $B = A^T$, E increases non-monotonically, the dynamics in \mathbf{U} diverges, and the Shannon entropies of agents' choice distribution asymptotically decreases (see Figs. 17 and 20 below). Note that in single-agent adaptation with state \mathbf{x} and normalizing the environment's reinforcements to a probability distribution \mathbf{p}_e , $D(\mathbf{p}_e || \mathbf{x})$ is always a Lyapunov function of the dynamics and decreases monotonically. In mutual adaptation, however, agents adapt to a dynamic environment that includes the other agents. As a result, in some cases, E , a linear sum of agent relative entropies, will itself exhibit nontrivial dynamics and, in addition, the uncertainties




Fig. 6. Smale's horseshoe in mutual adaptation: mutual adaptation and memory loss produce unstable and stable directions. The non-transitive structure of interactions leads to state-space folding.

of agents' choices will asymptotically decrease. When agents adapt with memory loss ($\alpha > 0$), the dynamics is dissipative. The dynamics varies between random and deterministic behavior in the information space \mathbf{U} . Notably, when the agents attempt to achieve this balance together by interacting and, in particular, when the interaction has *non-transitive* structure, the dynamics can persistently wander in a bounded area in information space. Since, in some cases, mutual adaptation and memory loss produce successive stretching and folding, deterministic chaos can occur within a significant range of α , even with only two agents. A schematic view of Smale's horseshoe in mutual adaptation is given in Fig. 6.

In the case that the agents are completely decoupled or in the case that $B = A^T$ and $\alpha_X = \alpha_Y = 0$ for two agents, information space locally splits into two subspaces governed by the effects of mutual adaptation and memory loss. They correspond to unstable and stable flow directions as in single-agent adaptation. However, in the case that agents are coupled via non-transitive interaction, mutual adaptation and memory loss nonlinearly interact with each other and Smale's horseshoe can be produced. In a general scheme, we can say that flow of information in mutual adaptation is multidimensional since each agent obtains information from its environment, organizes its behavior based on that information, and that local adaptation is then fed back into the environment affecting other agents.

In this case, "weak" uncertainty of behavior plays an important role in the self-organization of the collective behavior. Small fluctuations in the certainty of individual decision-making are amplified through mutual adaptation via non-transitive interactions. At the same time, memory loss increases the choice uncertainty and so stabilizes the collective's behavior.

Now consider many agents interacting. In the perfect memory case, when the game is zero-sum and has an interior Nash equilibrium $(\mathbf{x}^{1*}, \mathbf{x}^{2*}, \dots, \mathbf{x}^{S*})$, following Eq. (44), the following constant of motion exists:

$$E = \sum_{s=1}^S \frac{1}{\beta_s} D(\mathbf{x}^{s*} \| \mathbf{x}^s) = \sum_{s=1}^S \frac{1}{\beta_s} \left(\sum_{n^s=1}^{N^s} x_n^{s*} \log \frac{x_n^{s*}}{x_n^s} \right). \quad (47)$$

Although, strictly speaking, Hamiltonian dynamics and the associated symplectic structure of information space occurs only for two agents, one can describe multiple agent dynamics as a generalized Hamiltonian system [33]. In the general case with $\alpha > 0$, dissipative dynamics and high-dimensional chaotic flows can give rise to several unstable directions, since information flow has a network structure relative to the other agents. At least S stable directions are expected since memory loss comes from each individual's internal dynamics.

Summarizing, in single-agent adaptation, information flows unidirectionally from the environment to the agent and the agent adapts its behavior to the environmental constraints. Adaptation leads to $D(\mathbf{p}_e \| \mathbf{x}) \rightarrow 0$. For mutual adaptation in an agent collective, however, information flow is multidimensional since each agent obtains information from an environment that includes other agents. In this situation, E need not be a Lyapunov function for the dynamics. As we will see, when the dynamics is chaotic, global *information maximization* is of doubtful utility and the dynamic view of adaptation shown in Fig. 6 is more appropriate. When dynamic memory stored in collectives emerges, collective adaptation becomes a non-trivial problem. A detailed dynamical and information theoretic analysis along these lines will be reported elsewhere.

In the next section, we will give several phenomenological examples that illustrate collective adaptation.

4. Examples

To illustrate collective adaptation, we now give several examples of the dynamics in a static environment with two and three agents interacting via versions of Matching Pennies and Rock-Scissors-Paper, games with non-transitive structures. Appendix E gives the details of the reinforcement schemes for these cases. The agents will have equal adaptation rates ($\beta_X = \beta_Y = \dots$) and the same number of actions ($N = M = L = \dots$). In these simplified cases, the equations of motion for two agents are given by

$$\begin{aligned}\frac{\dot{x}_i}{x_i} &= [(A\mathbf{y})_i - \mathbf{x} \cdot A\mathbf{y}] + \alpha_X \left[-\log x_i + \sum_{n=1}^N x_n \log x_n \right], \\ \frac{\dot{y}_j}{y_j} &= [(B\mathbf{x})_j - \mathbf{y} \cdot B\mathbf{x}] + \alpha_Y \left[-\log y_j + \sum_{m=1}^M y_m \log y_m \right],\end{aligned}\quad (48)$$

for $i, j = 1, \dots, N$. A detailed analysis of this case with zero memory loss ($\alpha = 0$) is given in ref. [2] in terms of asymmetric game dynamics. We will present results for zero and positive memory loss rates.

We then consider three agents, for which the adaptation equations are

$$\begin{aligned}\frac{\dot{x}_i}{x_i} &= [(A\mathbf{y}\mathbf{z})_i - \mathbf{x} \cdot A\mathbf{y}\mathbf{z}] + \alpha_X \left[-\log x_i + \sum_{n=1}^N x_n \log x_n \right], \\ \frac{\dot{y}_j}{y_j} &= [(B\mathbf{z}\mathbf{x})_j - \mathbf{y} \cdot B\mathbf{z}\mathbf{x}] + \alpha_Y \left[-\log y_j + \sum_{m=1}^M y_m \log y_m \right], \\ \frac{\dot{z}_k}{z_k} &= [(C\mathbf{x}\mathbf{y})_k - \mathbf{z} \cdot C\mathbf{x}\mathbf{y}] + \alpha_Z \left[-\log z_k + \sum_{l=1}^L z_l \log z_l \right],\end{aligned}\quad (49)$$

for $i, j, k = 1, \dots, N$. We also will describe cases with and without memory loss.

Computer simulations are executed in the information spaces \mathbf{U} and the results are shown in the state spaces Δ , Δ_X , Δ_Y , and Δ_Z .

4.1. Two agents adapting under Matching Pennies interaction

In the matching pennies game, agents play one of two actions: heads (H) or tails (T). Agent X wins when the plays do not agree; agent Y wins when they do. Agent X 's state space is $\Delta_X = (x_1, x_2)$ with $x_i \in (0, 1)$ and $x_1 + x_2 = 1$. That is, x_1 is the probability that agent X plays heads; x_2 , tails. Agent Y is described similarly. Thus, each agent's state space is effectively one dimensional and the collective state space $\Delta = \Delta_X \times \Delta_Y$, two dimensional.

The environment for two agents interacting via the matching pennies game leads to the following matrices for Eqs. (48):

$$A = \begin{bmatrix} -\epsilon_X & \epsilon_X \\ \epsilon_X & -\epsilon_X \end{bmatrix} \quad \text{and} \quad B = \begin{bmatrix} -\epsilon_Y & \epsilon_Y \\ \epsilon_Y & -\epsilon_Y \end{bmatrix}, \quad (50)$$

where $\epsilon_X \in (0.0, 1.0]$ and $-\epsilon_Y \in (0.0, 1.0]$.

Fig. 7 shows a heteroclinic cycle of adaptation dynamics on the boundary of Δ when the α s vanish. Flows on the border occur only when agents completely ignore an action at the initial state; that is, when $x_i(0) = 0$ or $y_j(0) = 0$ for at least one i or j . Each vertex of the simplex is a saddle since the interaction is non-transitive.




Fig. 7. Flows on the boundary in Matching Pennies interaction: actions H and T correspond to “heads” and “tails”, respectively. Arrows indicate the direction of adaptation dynamics on the boundary of the state space Δ .




Fig. 8. Adaptation dynamics in Matching Pennies interaction: here $\epsilon_X = 0.5$ and $\epsilon_Y = -0.3$ with (left) $\alpha_X = \alpha_Y = 0$ and (right) $\alpha_X = 0.02$ and $\alpha_Y = 0.01$.

The Nash equilibrium $(\mathbf{x}^*, \mathbf{y}^*)$ of the Matching Pennies game is in the center of Δ : $(\mathbf{x}^*, \mathbf{y}^*) = (1/2, 1/2, 1/2, 1/2)$ and this is also a fixed point of the adaptation dynamics. The Jacobian at $(\mathbf{x}^*, \mathbf{y}^*)$ is

$$J = \begin{pmatrix} -\frac{\alpha_X}{2}(1 + \log 2) & -\frac{\epsilon_X}{2} \\ -\frac{\epsilon_Y}{2} & -\frac{\alpha_Y}{2}(1 + \log 2) \end{pmatrix} \quad (51)$$

and its eigenvalues are

$$\frac{4\lambda_i}{1 + \log 2} = -(\alpha_X + \alpha_Y) \pm \sqrt{(\alpha_X - \alpha_Y)^2 + \frac{4\epsilon_X\epsilon_Y}{(1 + \log 2)^2}}. \quad (52)$$

In the perfect memory case ($\alpha_X = \alpha_Y = 0$), trajectories near $(\mathbf{x}^*, \mathbf{y}^*)$ are neutrally stable periodic orbits, since $\lambda_i = \pm(1/2)\sqrt{\epsilon_X\epsilon_Y}$ are pure imaginary. In the memory loss case ($\alpha_X > 0$ and $\alpha_Y > 0$), $(\mathbf{x}^*, \mathbf{y}^*)$ is globally asymptotically stable, since $\text{Re}(\lambda_1)$ and $\text{Re}(\lambda_2)$ are strictly negative. Examples of the trajectories in these two cases are given in Fig. 8.

4.2. Three agents adapting under Even–Odd interaction

Now consider extending Matching Pennies for two agents so that it determines the interactions between three. Here we introduce the *Even–Odd* interaction in which there are again two actions, H and T , but agents win according to whether or not the number of heads in the group of three plays by the agents is even or odd. The environment




Fig. 9. Flows on the state space boundary under Even–Odd interactions: H and T correspond to “heads” and “tails”, respectively. Arrows indicate the direction of adaptation dynamics on Δ ’s boundary when the α s vanish.




Fig. 10. Dynamics of adaptation in the Even–odd interaction: $\epsilon_X = 0.5$, $\epsilon_Y = 0.2$, and $\epsilon_Z = -0.3$ with $\alpha_X = \alpha_Y = \alpha_Z = 0$ in (left) and with $\alpha_X = \alpha_Y = 0$ and $\alpha_Z = 0.01$ in (right). The trajectories with several initial conditions are shown in (left). The neutral subspace is shown as the horizontal cross and the trajectory chosen illustrates the attraction to a point in this subspace in (right).

now is given by, for agent X ,

$$a_{ijk} = \begin{cases} \epsilon_X, & \text{number of } H\text{s is even} \\ -\epsilon_X, & \text{otherwise} \end{cases} \quad (53)$$

with actions for agents X , Y , and Z given by $i, j, k = \{H, T\}$ and $\epsilon_X \in (0.0, 1.0]$. The interaction matrices b_{jki} and c_{kij} for agents Y and Z , respectively, are given similarly, but with $\epsilon_Y \in (0.0, 1.0]$ and $\epsilon_Z \in [-1.0, 0.0)$. Appendix E gives the details of the reinforcement scheme.

Following the reasoning used in Matching Pennies, the collective state space $\Delta = \Delta_X \times \Delta_Y \times \Delta_Z$ is now a solid three-dimensional cube. Fig. 9 shows a heteroclinic network of adaptation dynamics on the boundary of Δ when α s vanish. Flow on Δ ’s boundary is also indicated there.

Δ is partitioned into four prism-shaped subspaces. Each prism subspace has a heteroclinic cycle on the face that is also a face of Δ .

The Nash equilibrium of the Even–Odd interaction is $(\mathbf{x}^*, \mathbf{y}^*, \mathbf{z}^*) = (1/2, 1/2, 1/2, 1/2, 1/2, 1/2)$ at the center of Δ and this is also a fixed point of the adaptation dynamics. The Jacobian there is simply

$$J = \begin{pmatrix} -\alpha_X & 0 & 0 \\ 0 & -\alpha_Y & 0 \\ 0 & 0 & -\alpha_Z \end{pmatrix}. \quad (54)$$

Its eigenvalues are $\lambda = -\alpha_X, -\alpha_Y, -\alpha_Z$. Thus, in the complete memory case ($\alpha_X = \alpha_Y = \alpha_Z = 0$), trajectories near $(\mathbf{x}^*, \mathbf{y}^*, \mathbf{z}^*)$ are neutrally stable periodic orbits. With memory loss ($\alpha_X, \alpha_Y, \alpha_Z > 0$), the $(\mathbf{x}^*, \mathbf{y}^*, \mathbf{z}^*)$ is globally asymptotically stable. The hyperbolic fixed points in the top and bottom faces are unstable in all cases. Examples of the trajectories are given in Fig. 10, where the prism partitioning is clear.

Notably, when a single agent (say, Z) has memory loss and others have perfect memory, the crossed lines given by $\{z = x = 0.5, z = y = 0.5\}$ become an invariant subspace and trajectories are attracted to points in this subspace. Thus, there are infinitely many neutrally stable points. This is shown in the right diagram of Fig. 10.

With $\alpha_X = \alpha_Y = 0$ and $\alpha_Z = 0.01$, as in Fig. 10 for example, the adaptive dynamics alternates between a Matching Pennies interaction between agents X and Z by one between agents Y and Z during the transient relaxation to a point on the invariant subspace.

4.3. Two agents adapting under Rock–Scissors–Paper interaction

In this subsection we give an example of an environment in which agents have three actions. One of the most commonly studied games with three actions is the Rock–Scissors–Paper (RSP) game, in which an agent playing Rock beats one playing Scissors, which in turn beats an agent playing Paper, which finally beats Rock.

First we examine two agents, which is a straightforward implementation of the RSP game and then extend the RSP interaction to three agents and analyze the higher-dimensional behavior. The interaction matrices for these cases are given in Appendix E.

Under the RSP interaction each agent has the option of playing one of three actions: “rock” (R), “scissors” (S), and “paper” (P). Agent X ’s probability of playing these are denoted x_1, x_2 , and x_3 and $x_1 + x_2 + x_3 = 1$. Agent Y probabilities are given similarly. Thus, the agent state spaces, Δ_X and Δ_Y , are each two dimensional simplices, and the collective state space $\Delta = \Delta_X \times \Delta_Y$ is four dimensional.

For two agents the environment is given by the interaction matrices

$$A = \begin{bmatrix} \epsilon_X & 1 & -1 \\ -1 & \epsilon_X & 1 \\ 1 & -1 & \epsilon_X \end{bmatrix} \quad \text{and} \quad B = \begin{bmatrix} \epsilon_Y & 1 & -1 \\ -1 & \epsilon_Y & 1 \\ 1 & -1 & \epsilon_Y \end{bmatrix}, \quad (55)$$

where $\epsilon_X, \epsilon_Y \in [-1.0, 1.0]$ are the rewards for ties and normalized to

$$A' = \begin{bmatrix} \frac{2}{3}\epsilon_X & 1 - \frac{1}{3}\epsilon_X & -1 - \frac{1}{3}\epsilon_X \\ -1 - \frac{1}{3}\epsilon_X & \frac{2}{3}\epsilon_X & 1 - \frac{1}{3}\epsilon_X \\ 1 - \frac{1}{3}\epsilon_X & -1 - \frac{1}{3}\epsilon_X & \frac{2}{3}\epsilon_X \end{bmatrix} \quad (56)$$

and

$$B' = \begin{bmatrix} \frac{2}{3}\epsilon_Y & 1 - \frac{1}{3}\epsilon_Y & -1 - \frac{1}{3}\epsilon_Y \\ -1 - \frac{1}{3}\epsilon_Y & \frac{2}{3}\epsilon_Y & 1 - \frac{1}{3}\epsilon_Y \\ 1 - \frac{1}{3}\epsilon_Y & -1 - \frac{1}{3}\epsilon_Y & \frac{2}{3}\epsilon_Y \end{bmatrix}. \quad (57)$$

Note that the reinforcements are normalized to zero mean and that this does not affect the dynamics.




Fig. 11. Flows on the boundary of the simplex in the Rock–Scissors–Paper interaction for two agents: R , S , and P denote “rock”, “scissors”, and “paper”, respectively. The arrows indicate the direction of the adaptation dynamics on the boundary of the collective state space Δ when the α s vanish.

The flow on Δ 's boundary is shown in Fig. 11. This represents the heteroclinic network of adaptation dynamics on Δ 's edges when the α s vanish. Each vertex is a saddle since the interaction has non-transitive structure.

The Nash equilibrium $(\mathbf{x}^*, \mathbf{y}^*)$ is given by the centers of the simplex:

$$(\mathbf{x}^*, \mathbf{y}^*) = \left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}, \frac{1}{3}, \frac{1}{3}, \frac{1}{3} \right). \quad (58)$$

This is also a fixed point of the adaptation dynamics. The Jacobian there is

$$J = \begin{pmatrix} -\alpha_X & 0 & \frac{1+\epsilon_X}{3} & \frac{2}{3} \\ 0 & -\alpha_X & -\frac{2}{3} & \frac{-1+\epsilon_X}{3} \\ \frac{1+\epsilon_Y}{3} & \frac{2}{3} & -\alpha_Y & 0 \\ -\frac{2}{3} & \frac{-1+\epsilon_Y}{3} & 0 & -\alpha_Y \end{pmatrix}. \quad (59)$$

Its eigenvalues are

$$2\lambda_i = -(\alpha_X + \alpha_Y) \pm \sqrt{(\alpha_X - \alpha_Y)^2 + \frac{4(\epsilon_X \epsilon_Y - 3 \pm \sqrt{-3(\epsilon_X + \epsilon_Y)^2})}{9}}. \quad (60)$$

Thus, when (A, B) is zero-sum ($\epsilon_X + \epsilon_Y = 0$) and agents have complete memory ($\alpha_X = \alpha_Y = 0$), trajectories near $(\mathbf{x}^*, \mathbf{y}^*)$ are neutrally stable periodic orbits since all λ 's are pure imaginary. The dynamics is Hamiltonian in this case. With memory loss ($\alpha_X, \alpha_Y > 0$) and $|\alpha_X - \alpha_Y| < \frac{2}{3}(\epsilon_X^2 + 3)$, $(\mathbf{x}^*, \mathbf{y}^*)$ is globally asymptotically stable.

For the nonzero-sum case, we will give examples of dynamics with $\epsilon_X = 0.5$, $\epsilon_Y = -0.3$, $\alpha_Y = 0.01$. In this case, when $\alpha_X > \alpha_c$, $(\mathbf{x}^*, \mathbf{y}^*)$ is globally asymptotically stable. At the point $\alpha_c \sim 0.055008938$, a period-doubling bifurcation occurs. The example of two agents adapting in the Rock–Scissors–Paper interaction adaptation dynamics illustrates various types of low-dimensional chaos. We now explore several cases.

4.3.1. Hamiltonian limit

When the agent memories are perfect ($\alpha_X = \alpha_Y = 0$) and the game is zero-sum ($\epsilon_X = -\epsilon_Y$), the dynamics in the information space \mathbf{U} is Hamiltonian with an energy $E = D(\mathbf{x}^* \|\mathbf{x}) + D(\mathbf{y}^* \|\mathbf{y})$. The left columns of Figs. 12 and 13 give trajectories in the collective state space Δ , while the plots given in the middle and right columns are these trajectories projected onto the individual agent simplices, Δ_X and Δ_Y . The trajectories were generated using a 4th-order symplectic integrator [34] in \mathbf{U} .




Fig. 12. Quasiperiodic tori: collective dynamics in Δ (left column) and individual dynamics projected onto Δ_X and Δ_Y , respectively (right two columns). Here $\epsilon_X = -\epsilon_Y = 0.0$ and $\alpha_X = \alpha_Y = 0$. The initial condition is (A): $(\mathbf{x}, \mathbf{y}) = (0.26, 0.113333, 0.626667, 0.165, 0.772549, 0.062451)$ for the top and (B): $(\mathbf{x}, \mathbf{y}) = (0.05, 0.35, 0.6, 0.1, 0.2, 0.7)$ for the bottom. The constant of motion (Hamiltonian) is $E = 0.74446808 \equiv E_0$. The Poincaré section used for Fig. 14 is given by $x_1 = x_2$ and $y_1 < y_2$ and is indicated here as the straight diagonal line in agent X 's simplex Δ_X .

When $\epsilon_X = -\epsilon_Y = 0.0$ it appears that the dynamics is integrable since only quasiperiodic tori exist for almost all initial conditions in our simulation. For some initial conditions, tori are knotted and form a trefoil. Otherwise, when $\epsilon_X = -\epsilon_Y > 0.0$, Hamiltonian chaos occurs with positive–negative pairs of Lyapunov exponents (see Table 1). The game-theoretic behavior of this example was investigated briefly in ref. [15]. The dynamics is very rich. For example, there are infinitely many distinct behaviors near the fixed point at the center—the interior Nash equilibrium—and a periodic orbit arbitrarily close to any chaotic one.

A more detailed view of the complex dynamics is given in Fig. 14 which shows Poincaré sections of Eq. (48)'s trajectories. The Poincaré section is given by $\dot{u}_3 > 0$ and $\dot{v}_3 = 0$. In (\mathbf{x}, \mathbf{y}) space the section is determined by the constraints:

$$\begin{aligned} (1 - \epsilon_X)y_1 - (1 + \epsilon_X)y_2 + \frac{2}{3}\epsilon_X &< 0, \\ (1 - \epsilon_Y)x_1 - (1 + \epsilon_Y)x_2 + \frac{2}{3}\epsilon_Y &= 0. \end{aligned} \quad (61)$$

These sections are indicated as the straight lines drawn in the Δ_X simplices of Figs. 12 and 13. In Fig. 14, when $\epsilon_X = -\epsilon_Y = 0.0$, closed loops depending on the initial conditions exhibit tori in the Poincaré section. When $\epsilon_X = -\epsilon_Y = 0.5$, some tori collapse and become chaotic. The scatter of dots among the remaining tori show characteristic Hamiltonian chaos.

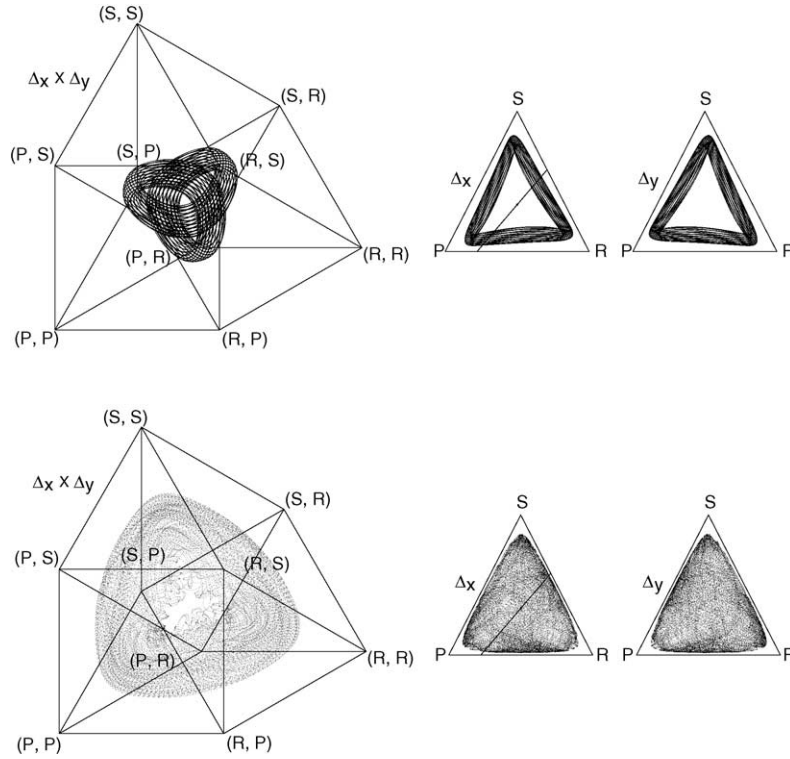


Fig. 13. Quasiperiodic tori and chaos: collective dynamics in Δ (left column) and individual dynamics projected onto Δ_X and Δ_Y , respectively (right two columns). Here $\epsilon_X = -\epsilon_Y = 0.5$ and $\alpha_X = \alpha_Y = 0$. The initial conditions are the same as in Fig. 12, (A) for top row and (B) for bottom rows, respectively. Also, the constant of motion is the same: $E = E_0$. The Poincaré section is given by $3x_1 - x_2 - 2/3 = 0$ and $y_1 - 3y_2 + 2/3 < 0$ and this is indicated as a straight line in Δ_X .

Table 1

Lyapunov spectra for different initial conditions (columns) and different values of the tie-breaking parameter ϵ_X

| ϵ_X | λ | $x_1(0) = 0.05$ | 0.06 | 0.07 | 0.08 | 0.09 | 0.10 |
|--------------|-------------|-----------------|--------------|--------|--------------|--------|--------|
| 0.0 | λ_1 | +0.881 | +0.551 | +0.563 | +0.573 | +0.575 | +0.589 |
| | λ_2 | 0.436 | +0.447 | +0.464 | +0.467 | +0.460 | +0.461 |
| | λ_3 | -0.436 | -0.447 | -0.464 | -0.467 | -0.460 | -0.461 |
| | λ_4 | -0.881 | -0.551 | -0.563 | -0.573 | -0.575 | -0.589 |
| 0.5 | λ_1 | +36.4 | +41.5 | +0.487 | +26.3 | +0.575 | +0.487 |
| | λ_2 | +0.543 | +0.666 | +0.204 | +0.350 | +0.460 | +0.460 |
| | λ_3 | -0.637 | -0.666 | -0.197 | -0.338 | -0.460 | -0.467 |
| | λ_4 | -36.3 | -41.5 | -0.494 | -26.3 | -0.575 | -0.480 |

The initial conditions are $(x_1, x_2, x_3, y_1, y_2, y_3) = (x_1, 0.35, 0.65 - x_1, 0.1, y_2, 0.9 - y_2)$ with $E = E_0 = 0.74446808$ fixed. We choose the initial conditions $(x_1, y_2) = (0.05, 0.2), (0.06, 0.160421), (0.07, 0.135275), (0.08, 0.117743), (0.09, 0.104795), (0.10, 0.0948432)$. The Lyapunov exponents are multiplied by 10^3 . Note that $\lambda_2 \simeq 0.0$, $\lambda_3 \simeq 0.0$ and $\lambda_4 \simeq -\lambda_1$ as expected. The Lyapunov exponents indicating chaos are shown in boldface.

Table 1 shows Lyapunov spectra in U for dynamics with $\epsilon_X = -\epsilon_Y = 0.0$ and $\epsilon_X = -\epsilon_Y = 0.5$, initial condition $(\mathbf{x}(0), \mathbf{y}(0)) = (x_1, 0.35, 0.65 - x_1, 0.1, y_2, 0.9 - y_2)$ and energy $E = E_0 = 0.74446808$ fixed. (x_1, y_2) satisfies

$$\frac{e^{-3(E_0 + 2 \log 3)}}{0.035} = x_1(0.65 - x_1)y_2(0.9 - y_2). \quad (62)$$




Fig. 14. Poincaré sections of the behavior in the preceding two figures. That is, $\epsilon_X = -\epsilon_Y = 0.0$ (left) and $\epsilon_X = -\epsilon_Y = 0.5$ (right). The Poincaré section is given by $x_1 = x_2$ and $y_1 < y_2$ (left) and $3x_1 - x_2 - 2/3 = 0$ and $y_1 - 3y_2 + 2/3 < 0$ (right). There are 25 randomly selected initial conditions, including the two (A) and (B), used in Figs. 12 and 13. The constant of motion ($E = E_0$) forms the outer border of the Poincaré sections.

When $x_1(0) = 0.05$, the initial condition is (B): $(\mathbf{x}, \mathbf{y}) = (0.05, 0.35, 0.6, 0.1, 0.2, 0.7)$, which we gave in the preceding examples. When $\epsilon_X = 0.5$, the Lyapunov exponents indicate positive-negative pairs for $x_1(0) = 0.05, 0.06$ and 0.08 , which clearly show Hamiltonian chaos. Note that $\lambda_2 \simeq 0.0$, $\lambda_3 \simeq 0.0$, and $\lambda_4 \simeq -\lambda_1$, as expected.

4.3.2. Conservative dynamics

With perfect memory ($\alpha_X = \alpha_Y = 0$) and a game that is not zero-sum ($\epsilon_X \neq -\epsilon_Y$) the dynamics is conservative in \mathbf{U} and one observes transients that are attracted to heteroclinic networks in the state space \mathbf{X} (see Fig. 15).

When $\epsilon_X + \epsilon_Y < 0$, the behavior is intermittent and orbits are guided by the flow on Δ 's edges, which describes a network of possible heteroclinic cycles. Since action ties are not rewarded there is only one such cycle. It is shown in the top row of Fig. (15): $(R, P) \rightarrow (S, P) \rightarrow (S, R) \rightarrow (P, R) \rightarrow (P, S) \rightarrow (R, S) \rightarrow (R, P)$. Note that during the cycle each agent switches between almost deterministic actions in the order $R \rightarrow S \rightarrow P$. The agents are out of phase with respect to each other and they alternate winning each turn.

With $\epsilon_X + \epsilon_Y > 0$, however, the orbit is an infinitely persistent chaotic transient [35]. Since, in this case, agent X can choose a tie, the cycles are not closed. For example, with $\epsilon_X > 0$, at (R, P) , X has the option of moving to (P, P) instead of (S, P) with a positive probability. This embeds an instability along the heteroclinic cycle and so orbits are chaotic (see Fig. 15, bottom row).

Fig. 16 shows the time series for these behaviors. Usually, in transient relaxation to heteroclinic cycle, the duration over which orbits stay near saddle vertices increases exponentially. However, for our case, it appears to increase subexponentially. This is because of the very small exponent; $(1 + \delta)^n \sim 1 + n\delta + \dots$ ($\delta \ll 1$). In the second chaotic transient case, it still increases subexponentially, but the visited vertices change irregularly.

Fig. 17 shows the behavior of the entropies H^X and H^Y and the energy E . For both cases E eventually increases monotonically and H^X and H^Y asymptotically decrease. The agents show a tendency to decrease choice uncertainty and to switch between almost deterministic actions. H^X and H^Y oscillate over the range $[0, \log 2]$ for $\epsilon_X = -0.1$ and $\epsilon_Y = 0.05$ and over $[0, \log 3]$ for $\epsilon_X = 0.1$ and $\epsilon_Y = -0.05$.

4.3.3. Dissipative dynamics

If the memory loss rates (α_X and α_Y) are positive, the dynamics becomes dissipative in information space \mathbf{U} and exhibits limit cycles and chaotic attractors (see Fig. 18). Fig. 19 (top) shows a diverse range of bifurcations as a function of α_X . It shows the dynamics on the surface specified by $\dot{u}_3 < 0$ and $\dot{v}_3 = 0$ projected onto v_3 . The fixed point $(\mathbf{x}^*, \mathbf{y}^*)$ becomes unstable when α_X is larger than $\alpha_c \approx 0.055008938$. Typically, period-doubling bifurcation to chaos occurs with decreasing α_X . Chaos can occur only when $\epsilon_X + \epsilon_Y > 0$ [16].

Fig. 19 (bottom) shows that the largest Lyapunov exponent in \mathbf{U} is positive across a significant fraction of the parameter space; indicating that chaos is common. The dual aspects of chaos, coherence and irregularity, imply that agents may behave cooperatively or competitively (or switch between both). This ultimately derives from the agents' successive mutual adaptation and memory loss in non-transitive interactions, such as in the RSP game; as was explained in Section 3. Note that such global behavior organization is induced by each agents' self-interested and myopic adaptation and “weak” uncertainty of their environment.

Fig. 20 shows dynamics of H^X , H^Y , and E in dissipative adaptive dynamics. For both cases shown E does not diverge due to memory loss. When $\alpha_X = 0.025$, H^X and H^Y converge to oscillations over the range $[\log 2, \log 3]$. When $\alpha_X = 0.01$, H^X and H^Y exhibit chaotic behavior over the range $[0, \log 3]$.




Fig. 15. Heteroclinic cycle with $\epsilon_X = -0.1$ and $\epsilon_Y = 0.05$ (top row). Chaotic transient to a heteroclinic network (bottom row) with $\epsilon_X = 0.1$ and $\epsilon_Y = -0.05$. For both $\alpha_X = \alpha_Y = 0$.




Fig. 16. Time series of action probabilities during the heteroclinic cycles of Fig. 15. $\epsilon_X = -0.1$ and $\epsilon_Y = 0.05$ for the left column. The right column shows the chaotic transient to a possible heteroclinic cycles when $\epsilon_X = 0.1$ and $\epsilon_Y = -0.05$. For both $\alpha_X = \alpha_Y = 0$.

4.4. Three agents adapting under Rock–Scissors–Paper interaction

Consider three agents adapting via (an extension of) the RSP interaction. Here the environment is given by the following interaction

$$a_{ijk} = \begin{cases} 2 & \text{Win over the others.} \\ -2 & \text{Lose to the other two.} \\ 1 & \text{Win over one other.} \\ -1 & \text{Lose to one other.} \\ \epsilon_X & \text{Tie.} \end{cases} \quad (63)$$

and similarly for b_{jki} and c_{kij} , with $i, j, k = \{R, S, P\}$. Here $\epsilon_X, \epsilon_Y, \epsilon_Z \in (-1.0, 1.0)$ (see Appendix E for the detailed listing of the reinforcement scheme).




Fig. 17. Dynamics of H^X, H^Y and E in conservative adaptive dynamics: $\epsilon_X = -0.1$ and $\epsilon_Y = 0.05$ for the left plot and $\epsilon_X = 0.1$ and $\epsilon_Y = -0.05$ for the right. For both $\alpha_X = \alpha_Y = 0$. Note that E increases asymptotically and H^X and H^Y tend to decrease.

As before we use normalized a'_{ijk} , b'_{jki} , and c'_{kij} :

$$a'_{ijk} = \begin{cases} 2 - \frac{\epsilon_X}{5} & \text{Win over the others.} \\ -2 - \frac{\epsilon_X}{5} & \text{Lose to the other two.} \\ 1 - \frac{\epsilon_X}{5} & \text{Win over one other.} \\ -1 - \frac{\epsilon_X}{5} & \text{Lose to one other.} \\ \frac{4}{5}\epsilon_X & \text{Tie.} \end{cases} \quad (64)$$

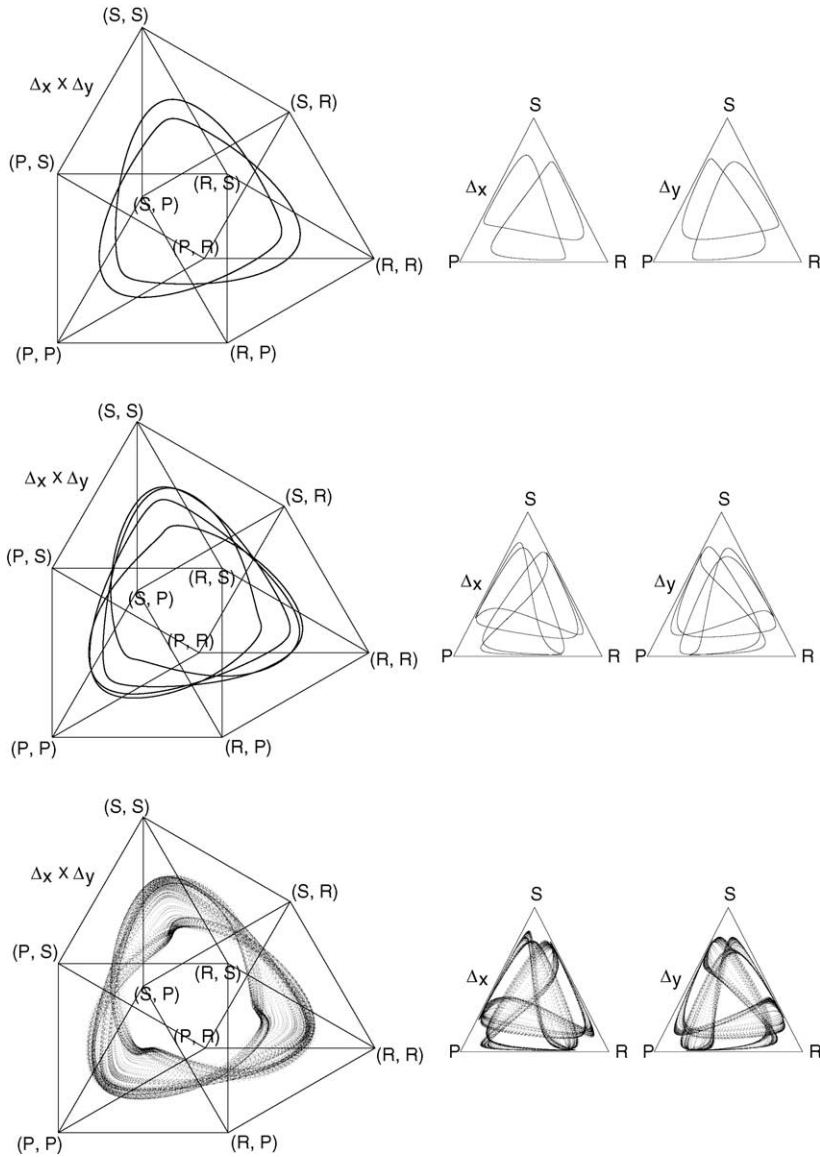


Fig. 18. Dissipative adaptive dynamics: stable limit cycle for $\alpha_X = 0.025$ (top), $\alpha_X = 0.021$ (middle) and chaotic attractors with $\alpha_X = 0.0198$ (bottom). All cases have $\epsilon_X = 0.5$, $\epsilon_Y = -0.3$ and $\alpha_Y = 0.01$. Period-doubling bifurcation to chaos occurs with decreasing α_X .




Fig. 19. Bifurcation diagram (top) of dissipative dynamics (adapting with memory loss) projected onto coordinate v_3 from the Poincaré section ($\dot{u}_3 > 0, v_3 = 0$) and the largest two Lyapunov exponents λ_1 and λ_2 (bottom) as a function of $\alpha_Y \in [0.01, 0.03]$. Here with $\epsilon_X = 0.5, \epsilon_Y = -0.3$ and $\alpha_Y = 0.01$. Simulations show that λ_3 and λ_4 are always negative.

The normalization does not affect the dynamics.

The Nash equilibrium $(\mathbf{x}^*, \mathbf{y}^*, \mathbf{z}^*)$ is at the simplex center:

$$(\mathbf{x}^*, \mathbf{y}^*, \mathbf{z}^*) = \left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}, \frac{1}{3}, \frac{1}{3}, \frac{1}{3}, \frac{1}{3}, \frac{1}{3}, \frac{1}{3} \right). \quad (65)$$




Fig. 20. Dynamics of H^X, H^Y , and E in dissipative adaptive dynamics: $\epsilon_X = 0.5, \epsilon_Y = -0.3$, and $\alpha_Y = 0.01$ for both. $\alpha_X = 0.025$ for the left plot and $\alpha_X = 0.01$ for the right. $t^* \approx 10^8$ in the right figure is the (rather long) transient time. In both cases E does not diverge due to memory loss.

It is also a fixed point of the adaptation dynamics. The Jacobian there is

$$J = \begin{pmatrix} -\alpha_X & 0 & \frac{1}{3} & \frac{2}{3} & \frac{1}{3} & \frac{2}{3} \\ 0 & -\alpha_X & -\frac{2}{3} & -\frac{1}{3} & -\frac{2}{3} & -\frac{1}{3} \\ \frac{1}{3} & \frac{2}{3} & -\alpha_Y & 0 & \frac{1}{3} & \frac{2}{3} \\ -\frac{2}{3} & -\frac{1}{3} & 0 & -\alpha_Y & -\frac{2}{3} & -\frac{1}{3} \\ \frac{1}{3} & \frac{2}{3} & \frac{1}{3} & \frac{2}{3} & -\alpha_Z & 0 \\ -\frac{2}{3} & -\frac{1}{3} & -\frac{2}{3} & -\frac{1}{3} & 0 & -\alpha_Z \end{pmatrix}. \quad (66)$$

When $\alpha_X = \alpha_Y = \alpha_Z = \alpha$, its eigenvalues are

$$\lambda_i + \alpha = \frac{i}{\sqrt{3}}(-1, -1, -2, 1, 1, 2). \quad (67)$$

In the perfect memory case ($\alpha_X = \alpha_Y = \alpha_Z = 0$), trajectories near $(\mathbf{x}^*, \mathbf{y}^*, \mathbf{z}^*)$ are neutrally stable periodic orbits, since the λ s are pure imaginary. In the memory loss case ($\alpha_X, \alpha_Y, \alpha_Z > 0$), $(\mathbf{x}^*, \mathbf{y}^*, \mathbf{z}^*)$ is asymptotically stable, since all $\text{Re}(\lambda_i)$ are strictly negative. One expects multiple attractors in this case.

The collective state space Δ is now six-dimensional, being the product of three two-dimensional agent simplices $\Delta = \Delta_X \times \Delta_Y \times \Delta_Z$. The flow on Δ 's boundary is shown in Fig. 21, giving the adaptation dynamics on the edges of Δ when the α s vanish.

We give two examples with $\alpha_X = \alpha_Y = \alpha_Z = 0.01$, $\epsilon_X = 0.5$, $\epsilon_Y = -0.365$, $\epsilon_Z = 0.8$ (top: limit cycle) and $\epsilon_X = 0.5$, $\epsilon_Y = -0.3$, $\epsilon_Z = 0.6$ (bottom: chaos) in Fig. 22. Chaos is typically observed when $\epsilon_X + \epsilon_Y + \epsilon_Z > 0$. Limit cycles are highly complex manifolds depending on the six-dimensional heteroclinic network on the simplex boundary. The Lyapunov spectrum for the chaotic dynamics is $(\lambda_1, \dots, \lambda_6) =$




Fig. 21. Flows on the simplex edges in three-agent RSP: arrows indicate the direction of adaptation dynamics on Δ 's boundary when the α s vanish.




Fig. 22. Periodic orbit (top: $\epsilon_X = 0.5, \epsilon_Y = -0.365, \epsilon_Z = 0.8$) and chaotic orbit (bottom: $\epsilon_X = 0.5, \epsilon_Y = -0.3, \epsilon_Z = 0.6$); the other parameters are $\alpha_X = \alpha_Y = \alpha_Z = 0.01$. The Lyapunov spectrum for chaotic dynamics is $(\lambda_1, \dots, \lambda_6) = (+45.2, +6.48, -0.336, -19.2, -38.5, -53.6) \times 10^{-3}$.

$(+45.2, +6.48, -0.336, -19.2, -38.5, -53.6) \times 10^{-3}$. The dynamics has two positive Lyapunov exponents. Note that this dynamics could have many neutrally stable subspaces in three or more dimensions. These subspaces act as quasistable attractors and may even have symplectic structure. These properties of high-dimensional dynamics will be reported elsewhere.

5. Concluding remarks

We developed a class of dynamical systems to model collective adaptation. We started with very simple agents, whose adaptation was a dynamic balance between adaptation to environmental constraints and memory loss. A macroscopic description of a network of adaptive agents was produced. In one special

case we showed that the dynamical system reduces to replicator equations, familiar in evolutionary game theory and population biology. In a more general setting, we investigated several of the resulting periodic, intermittent, and chaotic behaviors in which agent–agent interactions were explicitly given as game interactions.

Self-organization induced by the dynamics of uncertainty has been discussed from an information-theoretic viewpoint. We pointed out that unlike single-agent adaptation, the flow of information is multidimensional in collective adaptation. In this case global information maximization is of doubtful utility and a dynamic view of adaptation is more appropriate. We also noted that with only two agents interacting via a non-transitive game, a Smale horseshoe can be produced in information space. This occurs due to the agents' mutual adaptation, which amplifies the fluctuations in certainty of an individual's decision-making, and to memory loss which increases the choice uncertainties and so stabilizes the collective behavior.

Since deterministic chaos occurs even in this simple setting, one expects that in higher-dimensional and heterogeneous adaptive systems intrinsic unpredictability is a dominant collective behavior. When dynamic memory stored in a collective emerges, collective adaptation becomes a non-trivial problem. A detailed information theoretic and dynamical systems theoretic analysis will be reported elsewhere.

We close by indicating some future directions in which to extend the model.

First, as we alluded to during the development, there are difficulties of scaling the model to large numbers of agents. We focused on collectives with global coupling between all agents. However, in this case, the complexity of interaction terms grows exponentially with the number of agents, which is both impractical from the viewpoints of analysis and simulation, and unrealistic for natural systems that are large collectives. The solution to this, given in Appendix B, is to develop either spatially distributed agent collective or to extend the equations to include explicit communication networks between agents. Both of these extensions will be helpful in modeling the many adaptive collectives noted in the introduction.

Second, important for applications, is to develop the stochastic generalization of the deterministic equations of motion which accounts for the effects of finite and fluctuating numbers of agents and also finite histories for adaptation. Each of these introduces its own kind of sampling stochasticity and will require a statistical dynamics analysis reminiscent of that found in population genetics [36]. It is also important to consider the effects of asynchrony of adaptive behavior in this case.

Third, one necessary and possibly difficult extension will be to agents that adapt continuous-valued actions—say, learning the spatial location of objects—to their environments. Mathematically, this requires a continuous-space extension of the adaptation equations (Eq. (19)) and this results in models that are described by PDEs [37].

Finally, another direction, especially useful if one attempts to quantify global function in large collectives, will be structural and information-theoretic analyses of local and global adaptive behaviors [38,39]. Analyzing the stored information and the causal architecture [40,41], and quantifying “information influx” and “information dissipation” in individual adaptive behavior versus that in the collective are projects now made possible using this framework.

Acknowledgments

The authors thank D. Albers, N. Ay, C.R. Shalizi, D.C. Krakauer, J.D. Farmer and an anonymous referee for helpful discussions. YS was supported by the Special Postdoctoral Researchers Program at RIKEN and the Postdoctoral Fellowship at Santa Fe Institute. EA was supported by the Grant-in-Aid for Young Scientists (B) No. 14780342, The Ministry of Education, Culture, Sports, Science and Technology, Japan. JPC was supported at the Santa Fe Institute under the Network Dynamics Program by Intel Corporation and under DARPA contract F30602-00-2-0583.

Appendix A. Continuous time

Here we give the derivation of the continuous-time limits that lead to the deterministic differential equations from the original stochastic discrete-time adaptation model.

Denote the agent–agent interaction time scale, number of interactions per adaptation interval, and adaptation time scale as $d\tau$, T , and t , respectively. We assume that adaptation is very slow compared to agent–agent interactions and take the limits $d\tau \rightarrow 0$ and $T \rightarrow \infty$, keeping $dt = Td\tau$ finite. Then we take the limit $dt \rightarrow 0$ to get the derivative of the vector $\mathbf{Q}^X(t)$.

With Eqs. (14) and (15), we have

$$Q_i^X(T) = \frac{1}{T} \sum_{k=1}^T \left[\sum_{m=1}^M \delta_{im}(k) r_{im}^X(k) - \alpha_X Q_i^X(k) \right]. \quad (\text{A.1})$$

Thus, for continuous-time,

$$\frac{Q_i^X(t+dt) - Q_i^X(t)}{dt} = \frac{1}{T dt} \sum_{k=Tt}^{T(t+dt)} \left[\sum_{m=1}^M \delta_{im} \left(\frac{k}{T} \right) r_{im}^X \left(\frac{k}{T} \right) - \alpha_X Q_i^X \left(\frac{k}{T} \right) \right]. \quad (\text{A.2})$$

Taking $T \rightarrow \infty$ and $d\tau \rightarrow 0$, we have

$$\frac{Q_i^X(t+dt) - Q_i^X(t)}{dt} = \frac{1}{dt} \int_t^{t+dt} \left[\sum_{m=1}^M \delta_{im}(s) r_{im}^X(s) \right] ds - \alpha_X \frac{1}{dt} \int_t^{t+dt} Q_i^X(s) ds. \quad (\text{A.3})$$

Assuming $r_{ij}^X(t)$ changes as slowly as the adaptive dynamics, $r_{ij}^X(t)$ is constant during the adaptation interval $t \sim t + dt$. If we assume, in addition, that the behaviors of two agents X and Y are statistically independent at time t , then the law of the large numbers gives

$$\frac{1}{dt} \int_t^{t+dt} \left[\sum_{m=1}^M \delta_{im}(s) r_{im}^X(s) \right] ds \rightarrow \sum_{m=1}^M r_{im}(t) y_m(t) \equiv R_i^X(t). \quad (\text{A.4})$$

Now take $dt \rightarrow 0$. Eqs. (A.3) and (A.4) together give

$$\dot{Q}_i^X(t) = R_i^X(t) - \alpha_X Q_i^X(t), \quad (\text{A.5})$$

for the continuous-time updating of the reinforcement memory. If the environment is static, then $r_{ij}^X(t) = a_{ij}$ and

$$R_i^X(t) = \sum_{n=1}^N a_{in} y_i(t). \quad (\text{A.6})$$

The single-agent case is given by setting $\mathbf{y} = (1, 0, 0, \dots, 0)$ and $a_{i1} = a_i$, $i = 1, \dots, N$.

Appendix B. Network interactions

We can describe heterogeneous network interactions within our model. We give an example of a model for lattice interactions here. Agents $s = 1, 2, \dots, S$ are on a spatial lattice: agent s interacts with agent $s-1$ through bi-matrices (A^s, B^{s-1}) and agent $s+1$ through (B^s, A^{s+1}) . Each bi-matrix is 2×2 (see Fig. 23).




Fig. 23. Agent s interacts with agent $s - 1$ through bi-matrices (A^s, B^{s-1}) and agent $s + 1$ through (B^s, A^{s+1}) .

Agents choose actions among the 2×2 action pairs for both the right and left neighboring agents. The action pairs are (1, 1), (1, 2), (2, 1), (2, 2) and are weighted with probabilities x_1, \dots, x_4 . Inserting the interaction bi-matrices into the S-agent adaptive dynamics of Eq. (27) gives

$$\frac{\dot{x}_i^s}{x_i^s} = \beta_s [(A^s \mathbf{x}^{s-1})_i - \mathbf{p}^s \cdot A^s \mathbf{x}^{s-1} + (B^s \mathbf{x}^{s+1})_i - \mathbf{q}^s \cdot B^s \mathbf{x}^{s+1}] + \alpha_s \left(-\log x_i^s - \sum_{n=1}^4 x_n^s \log x_n^s \right), \quad (\text{B.1})$$

where $\sum x_i^s = 1$ and $\mathbf{p}^s = (x_1^s + x_2^s, x_3^s + x_4^s)$, $\mathbf{q}^s = (x_1^s + x_3^s, x_2^s + x_4^s)$. In a similar way, arbitrary network interactions can be described by our adaptive dynamics given in Eq. (27).

Appendix C. Nash equilibria

The *Nash equilibria* $(\mathbf{x}^*, \mathbf{y}^*)$ of the bi-matrix game (A, B) are those states in which all players can do no better by changing state; that is,

$$\mathbf{x}^* A \mathbf{y}^* \geq \mathbf{x} A \mathbf{y}^* \quad \text{and} \quad \mathbf{y}^* B \mathbf{x}^* \geq \mathbf{y} B \mathbf{x}^*, \quad (\text{C.1})$$

for all $(\mathbf{x}, \mathbf{y}) \in \Delta_X \times \Delta_Y$. If they exist in the interior, the solutions of the following simultaneous equations are Nash equilibria:

$$(\mathbf{A}\mathbf{y})_i = (\mathbf{A}\mathbf{y})_1 \quad \text{and} \quad (\mathbf{B}\mathbf{x})_j = (\mathbf{B}\mathbf{x})_1 \iff (\mathbf{A}\mathbf{y})_i - \mathbf{x} A \mathbf{y} = 0 \quad \text{and} \quad (\mathbf{B}\mathbf{x})_j - \mathbf{y} B \mathbf{x} = 0, \quad (\text{C.2})$$

where $\sum_{n=1}^N x_n = \sum_{m=1}^M y_m = 1$.

It is known that $N = M$ is a necessary condition for the existence of a unique Nash equilibrium in the interior of Δ . With $N = M$ in the perfect memory case ($\alpha_X = \alpha_Y = 0$), the unique Nash equilibrium, if it exists, is the fixed point given by the intersection of the x - and y -nullclines of Eqs. (24).

This Nash equilibrium is not asymptotically stable, but the time average of trajectories converges to it. To see this, suppose that $x_i(t) > \delta$ for all t sufficiently large, we have

$$\begin{aligned} \frac{d}{dt}(\log x_i) &= \frac{\dot{x}_i}{x_i} = (\mathbf{A}\mathbf{y})_i - \mathbf{x} A \mathbf{y}, \\ \frac{d}{dt}(\log y_j) &= \frac{\dot{y}_j}{y_j} = (\mathbf{B}\mathbf{x})_j - \mathbf{y} B \mathbf{x}. \end{aligned} \quad (\text{C.3})$$

Integrating the both sides from 0 to T and dividing by T , we get

$$\begin{aligned} \frac{\log x_i(T) - \log x_i(0)}{T} &= \sum_{m=1}^M a_{im} \bar{y}_m - S_A, \\ \frac{\log y_j(T) - \log y_j(0)}{T} &= \sum_{n=1}^N b_{jn} \bar{x}_n - S_B, \end{aligned} \quad (\text{C.4})$$

where

$$\bar{x}_i = T^{-1} \int_0^T x_i dt \quad \text{and} \quad \bar{y}_j = T^{-1} \int_0^T y_j dt, \quad (\text{C.5})$$

and

$$S_A = T^{-1} \int_0^T \mathbf{x} A \mathbf{y} dt \quad \text{and} \quad S_B = T^{-1} \int_0^T \mathbf{y} B \mathbf{x} dt. \quad (\text{C.6})$$

Letting $T \rightarrow \infty$, the left-hand sides converge to 0. Thus, $\bar{\mathbf{x}}$ and $\bar{\mathbf{y}}$ are a solution of Eq. (C.2). (This proof follows Ref. [42].)

Appendix D. Hamiltonian dynamics

Consider a game (A, B) that admits an interior Nash equilibrium $(\mathbf{x}^*, \mathbf{y}^*) \in \Delta_X \times \Delta_Y$ and is zero-sum ($B = -A^T$), then

$$E = \beta_X^{-1} D(\mathbf{x}^* \| \mathbf{x}) + \beta_Y^{-1} D(\mathbf{y}^* \| \mathbf{y}) \quad (\text{D.1})$$

is a constant of the motion. This follows by direct calculation:

$$\begin{aligned} \frac{dE}{dt} &= -\frac{1}{\beta_X} \sum_{n=1}^N x_n^* \frac{\dot{x}_n}{x_n} - \frac{1}{\beta_Y} \sum_{m=1}^M y_m^* \frac{\dot{y}_m}{y_m} \\ &= -(\mathbf{x}^* A \mathbf{y} - \mathbf{x} A \mathbf{y}) - (\mathbf{y}^* B \mathbf{x} - \mathbf{y} B \mathbf{x}) = (\mathbf{x}^* - \mathbf{x}) A (\mathbf{y}^* - \mathbf{y}) + (\mathbf{y}^* - \mathbf{y}) B (\mathbf{x}^* - \mathbf{x}) = 0. \end{aligned} \quad (\text{D.2})$$

This holds for any number of agents. Give the agents equal numbers of actions ($N = M$) and set α to zero (perfect memory) and make all β s finite and positive. Then the adaptive dynamics is Hamiltonian in the information space $\mathbf{U} = (\mathbf{u}, \mathbf{v})$ with the above constant of motion E ,

$$\dot{\mathbf{U}} = J \nabla_{\mathbf{U}} E, \quad (\text{D.3})$$

with Poisson structure J ,

$$J = \begin{pmatrix} O & P \\ -P^T & O \end{pmatrix} \quad \text{with} \quad P = -\beta_X \beta_Y A. \quad (\text{D.4})$$

Proof.

$$\begin{aligned} \frac{\partial E}{\partial u_i} &= \frac{\partial}{\partial u_i} \left[\beta_X^{-1} \sum_{n=1}^N x_n^* \log x_n^* + \beta_Y^{-1} \sum_{n=1}^N y_n^* \log y_n^* - \beta_X^{-1} \right. \\ &\quad \times \left(\sum_{n=1}^N x_n^* u_n - \log \left(\sum_{n=1}^N e^{-u_n} \right) \right) - \beta_Y^{-1} \left(\sum_{n=1}^N y_n^* v_n - \log \left(\sum_{n=1}^N e^{-v_n} \right) \right) \left. \right] \\ &= \beta_X^{-1} \left(x_i^* - \frac{e^{-u_i}}{\sum_{n=1}^N e^{-u_n}} \right) = \beta_X^{-1} (x_i^* - x_i), \end{aligned} \quad (\text{D.5})$$

$$\frac{\partial E}{\partial v_j} = \beta_Y^{-1} (y_j^* - y_j). \quad (\text{D.6})$$

Since $(\mathbf{x}^*, \mathbf{y}^*)$ is an interior Nash equilibrium, with Eq. (21), $(A\mathbf{y}^*)_i = (B\mathbf{x}^*)_j = 0$. Thus,

$$\begin{aligned} A \frac{\partial E}{\partial \mathbf{v}} &= -\frac{1}{\beta_Y} A\mathbf{y}, \\ B \frac{\partial E}{\partial \mathbf{u}} &= -\frac{1}{\beta_X} B\mathbf{x}, \end{aligned} \quad (\text{D.7})$$

and

$$J\nabla_{\mathbf{U}} E = \begin{bmatrix} -\beta_X \beta_Y A \frac{\partial E}{\partial \mathbf{v}} \\ -(-\beta_X \beta_Y A)^T \frac{\partial E}{\partial \mathbf{u}} \end{bmatrix} = \begin{bmatrix} -\beta_X A\mathbf{y} \\ -\beta_Y B\mathbf{x} \end{bmatrix} = \begin{bmatrix} \dot{\mathbf{u}} \\ \dot{\mathbf{v}} \end{bmatrix} = \dot{\mathbf{U}}. \quad (\text{D.8})$$

We can transform $\mathbf{U} = (\mathbf{u}, \mathbf{v})$ to canonical coordinates $\mathbf{U}' = (\mathbf{p}, \mathbf{q})$:

$$\dot{\mathbf{U}}' = S \nabla_{\mathbf{U}'} E, \quad (\text{D.9})$$

with

$$S = \begin{pmatrix} O & -I \\ I & O \end{pmatrix} \quad (\text{D.10})$$

where I is an $N \times N$ identity matrix and with a linear transformation $\mathbf{U}' = M\mathbf{U}$ to the Hamiltonian form. \square

Appendix E. Reinforcement schemes and interaction matrices

Here we give the reinforcement scheme interaction matrices for the constant-environment collectives investigated in Section 4.

E.1. Matching Pennies

This game describes a non-transitive competition. Each agent chooses a coin, which turns up either heads (H) or tails (T). Agent X wins when the coins differ, otherwise agent Y wins. Table E.1 gives the reinforcement scheme for the various possible plays. Note that the ϵ s determine the size of the winner's rewards. When $\epsilon_X + \epsilon_Y = 0$, the game is zero-sum. The Nash equilibrium is $\mathbf{x}^* = \mathbf{y}^* = (1/2, 1/2)$.

Various extensions of Matching Pennies to more than two players are known. We give the *Even-Odd* game as an example for three agents X , Y , and Z in a collective. All flip a coin. Agents X and Y win when the number of heads is even, otherwise Z wins. Table E.2 gives the reinforcement scheme. When the ϵ s add to zero, the game is zero-sum. The unique mixed Nash equilibrium is $\mathbf{x}^* = \mathbf{y}^* = \mathbf{z}^* = (1/2, 1/2, 1/2)$ —the simplex center.

Table E.1

The two-person Matching Pennies game: $\epsilon_X \in (0.0, 1.0]$ and $\epsilon_Y \in [-1.0, 0.0)$

| X | Y | r^X | r^Y |
|---|---|---------------|---------------|
| H | H | $-\epsilon_X$ | $-\epsilon_Y$ |
| H | T | ϵ_X | ϵ_Y |
| T | H | ϵ_X | ϵ_Y |
| T | T | $-\epsilon_X$ | $-\epsilon_Y$ |

Table E.2

The three-player Even–Odd game: $\epsilon_X \in (0.0, 1.0]$ and $\epsilon_Y, \epsilon_Z \in [-1.0, 0.0)$

| X | Y | Z | r^X | r^Y | r^Z |
|---|---|---|---------------|---------------|---------------|
| H | H | H | $-\epsilon_X$ | $-\epsilon_Y$ | $-\epsilon_Z$ |
| H | H | T | ϵ_X | ϵ_Y | ϵ_Z |
| H | T | H | ϵ_X | ϵ_Y | ϵ_Z |
| H | T | T | $-\epsilon_X$ | $-\epsilon_Y$ | $-\epsilon_Z$ |
| T | H | H | ϵ_X | ϵ_Y | ϵ_Z |
| T | H | T | $-\epsilon_X$ | $-\epsilon_Y$ | $-\epsilon_Z$ |
| T | T | H | $-\epsilon_X$ | $-\epsilon_Y$ | $-\epsilon_Z$ |
| T | T | T | ϵ_X | ϵ_Y | ϵ_Z |

E.2. Rock–Scissors–Paper

This game describes a non-transitive three-sided competition between two agents: rock (R) beats scissors (S), scissors beats paper (P), but paper beats rock. Table E.3 gives the reinforcement scheme. The ϵ s here control the rewards for ties. When they add to zero, the game is zero-sum. The unique mixed Nash equilibrium is $\mathbf{x}^* = \mathbf{y}^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ —again, the center of the simplex.

The extension of RSP interaction to three agents is straightforward. The reinforcement scheme is given in Table E.4. When $\epsilon_X + \epsilon_Y + \epsilon_Z = 0$, the game is zero-sum. The Nash equilibrium is $\mathbf{x}^* = \mathbf{y}^* = \mathbf{z}^* = (1/3, 1/3, 1/3)$.

Table E.3

The two-person Rock–Scissors–Paper game: $\epsilon_X, \epsilon_Y \in (-1.0, 1.0)$

| X | Y | r^X | r^Y |
|---|---|--------------|--------------|
| R | R | ϵ_X | ϵ_Y |
| R | S | 1 | -1 |
| R | P | -1 | 1 |
| S | R | -1 | 1 |
| S | S | ϵ_X | ϵ_Y |
| S | P | 1 | -1 |
| P | R | 1 | -1 |
| P | S | -1 | 1 |
| P | P | ϵ_X | ϵ_Y |

Table E.4

The three-person Rock–Scissors–Paper game: $\epsilon_X, \epsilon_Y, \epsilon_Z \in (-1.0, 1.0)$

| X | Y | Z | r^X | r^Y | r^Z | X | Y | Z | r^X | r^Y | r^Z | X | Y | Z | r^X | r^Y | r^Z |
|---|---|---|--------------|--------------|--------------|---|---|---|--------------|--------------|--------------|---|---|---|--------------|--------------|--------------|
| R | R | R | ϵ_X | ϵ_Y | ϵ_Z | S | R | R | -2 | 1 | 1 | P | R | R | 2 | -1 | -1 |
| R | R | S | 1 | 1 | -2 | S | R | S | -1 | 2 | -1 | P | R | S | ϵ_X | ϵ_Y | ϵ_Z |
| R | R | P | -1 | -1 | 2 | S | R | P | ϵ_X | ϵ_Y | ϵ_Z | P | R | P | 1 | -2 | 1 |
| R | S | R | 1 | -2 | 1 | S | S | R | -1 | -1 | 2 | P | S | R | ϵ_X | ϵ_Y | ϵ_Z |
| R | S | S | 2 | -1 | -1 | S | S | S | ϵ_X | ϵ_Y | ϵ_Z | P | S | S | -2 | 1 | 1 |
| R | S | P | ϵ_X | ϵ_Y | ϵ_Z | S | S | P | 1 | 1 | -2 | P | S | P | -1 | 2 | -1 |
| R | P | R | -1 | 2 | -1 | S | P | R | ϵ_X | ϵ_Y | ϵ_Z | P | P | R | 1 | 1 | -2 |
| R | P | S | ϵ_X | ϵ_Y | ϵ_Z | S | P | S | 1 | -2 | 1 | P | P | S | -1 | -1 | 2 |
| R | P | P | -2 | 1 | 1 | S | P | P | 2 | -1 | -1 | P | P | P | ϵ_X | ϵ_Y | ϵ_Z |

References

- [1] A.T. Winfree, *The Geometry of Biological Time*, Springer Verlag, 1980.
- [2] J. Hofbauer, K. Sigmund, *Evolutionary Games and Population Dynamics*, Cambridge University Press, Cambridge, 1988.
- [3] S. Camazine, J.-L. Deneubourg, N.R. Franks, J. Sneyd, G. Theraulaz, E. Bonabeau (Eds.), *Self-Organization in Biological Systems*, Princeton University Press, Princeton, 2001.
- [4] G. Bateson, *Steps to an Ecology of Mind: Collected Essays in Anthropology, Psychiatry, Evolution, and Epistemology*, Aronson, Northvale, N.J., 1987.
- [5] O.E. Rossler, Chaos in coupled optimizers, *Ann. NY Acad. Sci.* 504 (1987) 229–240.
- [6] M. Tajji, T. Ikegami, Dynamics of internal models in game players, *Physica D* 134 (1999) 253–266.
- [7] H.A. Simon, *The Sciences of the Artificial*, Karl Taylor Compton Lectures, first ed. 1969, MIT Press, Cambridge, 1996.
- [8] J.H. Holland, *Adaptation in Natural and Artificial Systems*, Second edition (First edition, 1975), MIT Press, Cambridge, MA, 1992.
- [9] T. Borgers, R. Sarin, Learning through reinforcement and replicator dynamics, *J. Econ. Th.* 77 (1997) 1–14.
- [10] D. Fudenberg, D.K. Levine, *Theory of Learning in Games*, MIT Press, 1998.
- [11] J. von Neumann, O. Morgenstern, *Theory of Games and Economic Behavior*, Princeton University Press, 1944.
- [12] P.D. Taylor, L.B. Jonker, Evolutionary stable strategies and game dynamics, *Math. Biosci.* 40 (1978) 145–156.
- [13] P.D. Taylor, Evolutionary stable strategies with two types of players, *J. Appl. Probab.* 16 (1979) 76–83.
- [14] J.W. Weibull, *Evolutionary Game Theory*, MIT Press, 1995.
- [15] Y. Sato, E. Akiyama, J.D. Farmer, Chaos in learning a simple two-person game, *Proc. Natl. Acad. Sci. U.S.A.* 99 (2002) 4748–4751.
- [16] Y. Sato, J.P. Crutchfield, Coupled replicator equations for the dynamics of learning in multiagent systems, *Phys. Rev. E* 67 (2003) 015206(R).
- [17] E. Akiyama, K. Kaneko, Dynamical systems game theory and dynamics of games, *Physica D* 147 (3–4) (2000) 221–258.
- [18] E. Akiyama, K. Kaneko, Dynamical systems game theory II—a new approach to the problem of the social dilemma, *Physica D* 167 (1–2) (2002) 36–71.
- [19] Y. Itoh, Random collision models in oriented graphs, *J. Appl. Prob.* 16 (1979) 36.
- [20] A.M. Turing, The chemical basis of morphogenesis, *Phil. Trans. R. Soc. Lond. B* 237 (1952) 37–72.
- [21] C.E. Shannon, *The Mathematical Theory of Communication*, The University of Illinois Press, 1949.
- [22] W.S. McCulloch, A heterarchy of values determined by the topology of nervous nets, *Bull. Mathe. Biophys.* 5 (1945) 115–133.
- [23] B.F. Skinner, *The Behavior of Organisms: An Experimental Analysis*, Appleton-Century, New York, 1938.
- [24] D.O. Hebb, *The Organization of Behaviour*, John Wiley and Sons, New York, 1949.
- [25] M.F. Norman, *Markov Processes and Learning Models*, Academic Press, New York, 1972.
- [26] A.L. Samuel, Some studies in machine learning using the game of checkers, *IBM J. Res. Dev* 11 (1967) 601.
- [27] R.S. Sutton, A.G. Barto, *Reinforcement Learning: An Introduction*, MIT Press, 1998.
- [28] D. Kahneman, A. Tversky (Eds.), *Choices, Values, and Frames*, Cambridge University Press, 2000.
- [29] J. Aitchison, *The Statistical Analysis of Compositional Data*, Blackburn Press, Caldwell, New Jersey, 1986.
- [30] T.M. Cover, J.A. Thomas, *Elements of Information Theory*, John Wiley and Sons Inc., 1991.
- [31] J. Hofbauer, Evolutionary dynamics for bimatrix games: a hamiltonian system?, *J. Math. Biol.* 34 (1996) 675–688.
- [32] J.J. Hopfield, Neural networks and physical systems with emergent collective computational abilities, *Proc. Natl. Acad. Sci. U.S.A.* 79 (1982) 2554–2558.
- [33] A.M. Perelomov, *Integrable Systems of Classical Mechanics and Lie Algebras*, Birkhauser, 1990.
- [34] H. Yoshida, Construction of higher order symplectic integrators, *Phys. Lett. A* 150 (1990) 262–268.
- [35] T. Chawanya, A new type of irregular motion in a class of game dynamics systems, *Prog. Theo. Phys.* 94 (1995) 163–179.
- [36] E. van Nimwegen, J.P. Crutchfield, M. Mitchell, Statistical dynamics of the royal road genetic algorithm, *Theor. Comp. Sci.* 229 (1999) 41–102.
- [37] J. Hofbauer, V. Hutson, G.T. Vickers, Travelling waves for games in economics and biology, *Nonlinear Anal. Theory, Methods, Appl.* 30 (2) (1997) 1235–1244.
- [38] R. Shaw, *The Dripping Faucet as a Model Chaotic System*, Aerial Press, Santa Cruz, California, 1984.
- [39] J.P. Crutchfield, K. Young, Inferring statistical complexity, *Phys. Rev. Lett.* 63 (1989) 105–108.
- [40] J.P. Crutchfield, C.R. Shalizi, Thermodynamic depth of causal states: objective complexity via minimal representations, *Phys. Rev. E* 59 (1) (1999) 275–283.
- [41] J.P. Crutchfield, D.P. Feldman, Regularities unseen, randomness observed: levels of entropy convergence, *Chaos* 13 (1) (2003) 25–54.
- [42] P. Schuster, K. Sigmund, J. Hofbauer, R. Wolff, Selfregulation of behavior in animal societies. Part I. Symmetric contests, *Biol. Cybern.* 40 (1981) 1–8.