

Memory Asymmetry Creates Heteroclinic Orbits to Nash Equilibrium in Learning in Zero-Sum Games^{*†}

Yuma Fujimoto^{1,2,3}, Kaito Ariu^{3,4}, Kenshi Abe^{3,5}

¹SOKENDAI

²The University of Tokyo

³CyberAgent

⁴KTH

⁵The University of Electro-Communications
fujimoto_yuma_xa@cyberagent.co.jp

Abstract

Learning in games considers how multiple agents maximize their own rewards through repeated games. Memory, an ability that an agent changes his/her action depending on the history of actions in previous games, is often introduced into learning to explore more clever strategies and discuss the decision-making of real agents like humans. However, such games with memory are hard to analyze because they exhibit complex phenomena like chaotic dynamics or divergence from Nash equilibrium. In particular, how asymmetry in memory capacities between agents affects learning in games is still unclear. In response, this study formulates a gradient ascent algorithm in games with asymmetry memory capacities. To obtain theoretical insights into learning dynamics, we first consider a simple case of zero-sum games. We observe complex behavior, where learning dynamics draw a heteroclinic connection from unstable fixed points to stable ones. Despite this complexity, we analyze learning dynamics and prove local convergence to these stable fixed points, i.e., the Nash equilibria. We identify the mechanism driving this convergence: an agent with a longer memory learns to exploit the other, which in turn endows the other's utility function with strict concavity. We further numerically observe such convergence in various initial strategies, action numbers, and memory lengths. This study reveals a novel phenomenon due to memory asymmetry, providing fundamental strides in learning in games and new insights into computing equilibria.

Introduction

Learning in games discusses how players learn their strategies through repeated games (Fudenberg and Levine 1998). Here, whether or not players learn their optimal strategies, called Nash equilibrium, is a nontrivial problem because each player's optimal strategy depends on the strategy of his/her opponent. This problem becomes further exacerbated in zero-sum games, where both the players conflict in their utility functions. In order to approach this problem, various

learning algorithms have been proposed, such as replicator dynamics (Börgers and Sarin 1997; Hofbauer, Sigmund et al. 1998; Sato, Akiyama, and Farmer 2002) and gradient ascent (Singh, Kearns, and Mansour 2000; Zinkevich 2003; Bowling and Veloso 2002; Bowling 2004). The dynamics of these algorithms have undergone extensive study, with findings indicating a cycling behavior around the Nash equilibrium in zero-sum games (Mertikopoulos and Sandholm 2016; Mertikopoulos, Papadimitriou, and Piliouras 2018).

Memory, referring to an agent's ability to alter their subsequent actions based on past decisions, is often incorporated into games to investigate more sophisticated strategies and explore the decision-making processes of real agents, such as humans. Such memory introduces greater complexity and diversity into games. A celebrated study in economics has shown that players with memory can take various strategies in the Nash equilibria (Fudenberg and Maskin 2009). Indeed, in prisoner's dilemma games, agents with memory can achieve cooperation (Axelrod and Hamilton 1981) and asymmetric degree of cooperation (Fujimoto and Kaneko 2019), even though those without memories cannot cooperate in the Nash equilibrium. Memory also complicates learning dynamics such as chaotic dynamics (Barfuss, Donges, and Kurths 2019) and divergence from the Nash equilibrium (Fujimoto, Ariu, and Abe 2023) in zero-sum games. In recent years, theoretical approaches to games with memories have garnered significant interest (Barfuss 2020; Usui and Ueda 2021; Meylahn, Janssen et al. 2022; Ueda 2023).

How learning changes depending on memory length is an interesting question in games with memory. Indeed, this question has often been discussed in prisoner's dilemma games. It is reported that agents with longer memories can cooperate more cleverly (Hilbe et al. 2017; Murase and Baek 2020). Furthermore, the difference in payoffs emerges between agents with different memory lengths in various learning algorithms such as Q-learning (Sandholm and Crites 1996), coupled replicator dynamics (Fujimoto and Kaneko 2021), and evolutionary dynamics (Baek et al. 2016; Schmid et al. 2022). This suggests that an asymmetry in memory capacities could potentially lead to the exploitation of the opponent's payoff in zero-sum games. Despite this, the impact

^{*}The full version is at <https://arxiv.org/abs/2305.13619>

[†]The codes that we used are available at https://github.com/CyberAgentAILab/learning_games_with_memory_asymmetry
Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

of memory asymmetry in games has yet to be thoroughly investigated.

This study considers learning in zero-sum games between agents with memory asymmetry. First, we formulate games with asymmetry memory and a gradient ascent algorithm in the games. This algorithm corresponds to replicator dynamics, provided that the learning rate is sufficiently small, as outlined in Fujimoto, Ariu, and Abe (2023). In order to catch theoretical insights, we first focus on two-action zero-sum games between agents with memory lengths of one and zero (called one-memory and zero-memory). We analyze the Nash equilibrium in this with-memory game in comparison with the ‘‘original’’ Nash equilibrium in games without memory. Interestingly, we prove that this original Nash equilibrium can be divided into two distinct regions. The first region comprises unstable fixed points within the learning dynamics, while the second consists of stable fixed points. These stable points correspond to the Nash equilibrium in with-memory games. Thus, heteroclinic orbits (Strogatz 2018) are observed in with-memory games; Dynamics diverge from the unstable region and converge to the stable region within the original Nash equilibrium. This convergence is because the agent with a longer memory learns to concave the opponent’s utility function. Furthermore, we experimentally confirmed this convergence across various initial states, action numbers, and memory lengths.

To summarize, our study uncovers complex heteroclinic dynamics in learning in with-memory games. We also identify a novel phenomenon of convergence in learning in zero-sum games, which is achieved due to memory asymmetry. Given that replicator dynamics exhibit cyclical behavior without memory (see Panel A in Fig. 1) and diverging behavior with memory symmetry (see Panel B), our discovery is significant (see Panel C).

Preliminary

Two-Player Normal-Form Games

We consider two players of X and Y. In every round, each of them chooses its action from $\mathcal{A} = \{a_1, \dots, a_m\}$ and $\mathcal{B} = \{b_1, \dots, b_m\}$. If they choose $a \in \mathcal{A}$ and $b \in \mathcal{B}$, they immediately receive a payoff of $u_{ab} \in \mathbb{R}$ and $v_{ab} \in \mathbb{R}$, respectively. For the illustration of two-player normal-form games, see the area surrounded by the magenta dotted lines in Fig. 2-A.

Games with Memory Asymmetry

We assume that X and Y can memorize their actions in the newest $n_X \in \mathbb{N}$ and $n_Y \in \mathbb{N}$ rounds, respectively. We set that X has a longer memory ($n_X > n_Y$). Let $\mathcal{S} = \prod_{k=1}^{n_X} (\mathcal{A} \times \mathcal{B})$ be the set of all states that X can memorize. Notably, each state $s \in \mathcal{S}$ is given by a string of actions of length $2n_X$. Under state s , X can choose its action a with the probability of $x_{a|s} \in (0, 1)$. X’s strategy is the set of $x_{a|s}$ for all $a \in \mathcal{A}$ and $s \in \mathcal{S}$ and is denoted as $|\mathcal{S}| (= m^{2n_X})$ -numbers of $(m - 1)$ -dimension simplexes, $\mathbf{x} \in \prod_{s \in \mathcal{S}} \Delta^{m-1}$. Fig. 2-A shows the introduction of memories into normal-form games.

For each $n \in \{0, \dots, n_X - 1\}$, let s_n denote the newest n substring of s . We also define the set \mathcal{S}_n as $\mathcal{S}_n = \prod_{k=1}^n (\mathcal{A} \times$

$\mathcal{B})$, which represents the set of all possible s_n . Since the length of Y’s memory is $n_Y (< n_X)$, the set of states memorized by Y can be represented by \mathcal{S}_{n_Y} . We assume that under each state s_{n_Y} , Y chooses its action b with probability $y_{b|s_{n_Y}} \in (0, 1)$. If $n_Y = 0$, $\mathcal{S}_{n_Y} = \emptyset$ and we simply denote $y_{b|s_{n_Y}}$ as y_b . Y’s strategy is given by $\mathbf{y} \in \prod_{s_{n_Y} \in \mathcal{S}_{n_Y}} \Delta^{m-1}$. Because the above immediate payoff is determined by both the players’ actions in the last round ($s_1 \in \mathcal{A} \times \mathcal{B}$), we can rewrite the payoff as u_{s_1} and v_{s_1} .

Formulation as Markov Transition Processes

The above setting is described by a discrete-time Markov transition process (see the illustration of Fig. 2-B), where the transition rate from state s to s' is

$$M_{s's}(\mathbf{x}, \mathbf{y}) := \begin{cases} x_{a|s} y_{b|s_{n_Y}} & (s' = abs_{n_X-1}) \\ 0 & (\text{otherwise}) \end{cases}. \quad (1)$$

We formulate the Nash equilibrium based on the Markov transition process. Then, strategies \mathbf{x} and \mathbf{y} may be regarded as fixed in the above Markov transition process. For fixed \mathbf{x} and \mathbf{y} within the interior of the simplexes, the stationary state for the above Markov transition processes is uniquely denoted as $p_s^{\text{st}}(\mathbf{x}, \mathbf{y})$. The stationary state satisfies $p_s^{\text{st}}(\mathbf{x}, \mathbf{y}) = \sum_s M_{s's}(\mathbf{x}, \mathbf{y}) p_s^{\text{st}}(\mathbf{x}, \mathbf{y})$. If the game continues a sufficiently long time, each player’s payoff is given under this stationary state; $u^{\text{st}}(\mathbf{x}, \mathbf{y}) = \sum_s p_s^{\text{st}}(\mathbf{x}, \mathbf{y}) u_{s_1}$ and $v^{\text{st}}(\mathbf{x}, \mathbf{y}) = \sum_s p_s^{\text{st}}(\mathbf{x}, \mathbf{y}) v_{s_1}$. Therefore, the above payoffs reflect that every player learns slowly enough, i.e., their learning rates are sufficiently small. The Nash equilibrium is defined as the strategies that maximize their own payoff functions in the stationary state;

$$\begin{cases} \mathbf{x}^* \in \operatorname{argmax}_{\mathbf{x}} u^{\text{st}}(\mathbf{x}, \mathbf{y}^*) \\ \mathbf{y}^* \in \operatorname{argmax}_{\mathbf{y}} v^{\text{st}}(\mathbf{x}^*, \mathbf{y}) \end{cases}. \quad (2)$$

The objective of this study is to learn this Nash equilibrium.

Algorithm

Referring to Fujimoto, Ariu, and Abe (2023), we now formulate the algorithm of multi-memory gradient ascent (MMGA) for asymmetric memory lengths between the players. This discretized algorithm can be implemented even when we do not know the analytical expression of $u^{\text{st}}(\operatorname{Norm}(\mathbf{x}), \mathbf{y})$.

Algorithm 1: Discretized MMGA

Input: η, γ

```

1: for  $t = 0, 1, 2, \dots$  do
2:   for  $a \in \mathcal{A}, s \in \mathcal{S}$  do
3:      $\mathbf{x}' \leftarrow \mathbf{x} + \gamma \mathbf{e}_{a|s}$ 
4:      $\Delta_{a|s} \leftarrow \frac{u^{\text{st}}(\operatorname{Norm}(\mathbf{x}'), \mathbf{y}) - u^{\text{st}}(\mathbf{x}, \mathbf{y})}{\gamma}$ 
5:   end for
6:   for  $a \in \mathcal{A}, s \in \mathcal{S}$  do
7:      $x_{a|s} \leftarrow x_{a|s} (1 + \eta \Delta_{a|s})$ 
8:   end for
9:    $\mathbf{x} \leftarrow \operatorname{Norm}(\mathbf{x})$ 
10: end for
```




Figure 1: Learning dynamics illustrated for three memory-configuration scenarios involving two agents. Learning dynamics show a cycling behavior around the Nash equilibrium when the agents have no memory (left panel). Learning dynamics diverge from the Nash equilibrium when the agents have the same memory capacity (center). Learning dynamics draw heteroclinic orbits and eventually converge to the Nash equilibrium when the agents have different memory lengths (right). In all the panels, the horizontal and vertical axes indicate the probabilities of the agents choosing “head” in matching-pennies games (see Fig. 2). In the center and right panels, the color gradient indicates the passage of time (blue represents older data, and red represents newer data).

In Alg. 1, the initial states of \mathbf{x} and \mathbf{y} are set. In each time step of t , \mathbf{X} shifts each component of \mathbf{x} by γ from the original strategy \mathbf{x} and define \mathbf{x}' (line 3). Here, $\mathbf{e}_{a|s}$ is defined as the unit vector for the direction of $x_{a|s}$. Furthermore, we normalize \mathbf{x}' for each state as $\text{Norm}(\mathbf{x}') := \{x'_{a|s} / (\sum_a x'_{a|s})\}_{a,s}$. By comparing $\text{Norm}(\mathbf{x}')$ with \mathbf{x} , \mathbf{X} gets the gradient $\Delta_{a|s}$ (line 4). Using all of $\Delta_{a|s}$, \mathbf{X} updates its strategy $x_{a|s}$ with learning rate η (line 7) and normalizes it (line 8). Here, note that the update is weighted by $x_{a|s}$ itself. We also implement Alg. 1 for agent \mathbf{Y} similarly. However, because \mathbf{Y} has a shorter memory than \mathbf{X} , its strategy is updated for $s_{n_Y} \in \mathcal{S}_{n_Y}$ instead of $s \in \mathcal{S}$. \mathbf{Y} 's action $b \in \mathcal{B}$ and payoff v^{st} can be different from $a \in \mathcal{A}$ and u^{st} .

If the analytical expression of $u^{\text{st}}(\mathbf{x}, \mathbf{y})$ and $v^{\text{st}}(\mathbf{x}, \mathbf{y})$ are known, we can use the following dynamics (continualized MMGA (Fujimoto, Ariu, and Abe 2023)) instead of Alg. 1 by taking the limit of $\gamma \rightarrow 0$ and $\eta \rightarrow 0$;

$$\dot{x}_{a|s}(\mathbf{x}, \mathbf{y}) = x_{a|s} \frac{\partial}{\partial x_{a|s}} u^{\text{st}}(\text{Norm}(\mathbf{x}), \mathbf{y}), \quad (3)$$

$$\dot{y}_{b|s_{n_Y}}(\mathbf{x}, \mathbf{y}) = y_{b|s_{n_Y}} \frac{\partial}{\partial y_{b|s_{n_Y}}} v^{\text{st}}(\mathbf{x}, \text{Norm}(\mathbf{y})). \quad (4)$$

These continuous-time dynamics of the MMGA algorithm are useful for the theoretical treatment as they allow for classical stability analysis and numerical integration.

Theoretical Results

In this section, we analyze learning dynamics under asymmetric memory lengths between the players. To simplify theoretical treatments, we first establish several assumptions for the games, such as payoff matrix, action number, and memory length. Next, we derive the Nash equilibrium in these games and compare this equilibrium with the original Nash equilibrium in games without memory. Finally, we analyze the dynamics of learning. We prove that the fixed points

of learning dynamics are divided into stable and unstable ones. Thus, heteroclinic dynamics diverging from unstable fixed points and converging to stable ones are drawn. In this section, we present only the proof sketch for each theorem, while detailed proofs can be found in Technical Appendix A.


Assumptions

First, we consider two-action zero-sum games as follows.

Assumption 1 (Two-action zero-sum game). *We denote the two actions of two agents as $\mathcal{A} = \{a_1, a_2\}$ and $\mathcal{B} = \{b_1, b_2\}$. There are four states $\mathcal{S}_1 = \mathcal{A} \times \mathcal{B} = \{a_1b_1, a_1b_2, a_2b_1, a_2b_2\}$. If each state occurs, \mathbf{X} receives the payoff corresponding to $\mathbf{u} = (u_1, \dots, u_4) := (u_{a_1b_1}, u_{a_1b_2}, u_{a_2b_1}, u_{a_2b_2})$. Zero-sum games assume that \mathbf{Y} 's payoff is given by $v_{ab} = -u_{ab}$ for all $a \in \mathcal{A}$ and $b \in \mathcal{B}$. To ensure a nontrivial game, we assume that u_1 and u_4 are both larger than u_2 and u_3 .*

In Asm. 1, we consider a specific class of zero-sum games where the Nash equilibrium exists in the interior of the strategy spaces of two players. This is because if the Nash equilibrium is on the boundary of the strategy spaces, a dominant strategy exists and learning dynamics are trivial. An example that satisfies Asm. 1 is matching-pennies games ($u_1 = u_4 = +1$ and $u_2 = u_3 = -1$). Indeed, the Nash equilibrium is defined as follows.

Definition 1 (Original Nash equilibrium in two-action zero-sum normal-form game). *Under Asm. 1 (two-action zero-sum game), we consider $(n_X, n_Y) = (0, 0)$ (no memory = normal-form game). Then, the mixed strategy of \mathbf{X} is given by a single variable $x := x_{a_1} \in (0, 1)$. Similarly, \mathbf{Y} 's strategy is given by $y := y_{b_1} \in (0, 1)$. Then, the Nash equilibrium (x°, y°) and the payoffs in the equilibrium (u°, v°) are*



given by

$$x^o = \frac{-u_3 + u_4}{u_1 - u_2 - u_3 + u_4}, \quad y^o = \frac{-u_2 + u_4}{u_1 - u_2 - u_3 + u_4}, \quad (5)$$

$$u^o = \frac{u_1 u_4 - u_2 u_3}{u_1 - u_2 - u_3 + u_4}, \quad v^o = -u^o. \quad (6)$$

The Nash equilibrium in no-memory games has been known for a long time (Nash Jr 1950). On the other hand, when agents with memories are considered, the region of the Nash equilibrium is generally extended (Fudenberg and Maskin 2009). Therefore, we call such a no-memory Nash equilibrium (x^o, y^o) as the “original” Nash equilibrium for distinction.

We now introduce a specific class of with-memory games under Asm. 1 and define notations.

Definition 2 (One-memory and zero-memory strategies and vector notation). *Under Asm. 1, consider $(n_X, n_Y) = (1, 0)$. Because a constraint $x_{a_1|s} + x_{a_2|s} = 1$ holds for all $s \in \mathcal{S} = \{a_1 b_1, a_1 b_2, a_2 b_1, a_2 b_2\}$, X’s strategy is determined by a four-variable vector $\mathbf{x} = (x_1, \dots, x_4) := (x_{a_1|a_1 b_1}, x_{a_1|a_1 b_2}, x_{a_1|a_2 b_1}, x_{a_1|a_2 b_2}) \in (0, 1)^4$. On the other hand, Y’s strategy is determined by $y := y_{b_1} \in (0, 1)$. We also use the vector notation of the stationary state $\mathbf{p}^{\text{st}} = (p_1^{\text{st}}, \dots, p_4^{\text{st}}) := (p_{a_1 b_1}^{\text{st}}, p_{a_1 b_2}^{\text{st}}, p_{a_2 b_1}^{\text{st}}, p_{a_2 b_2}^{\text{st}})$. Expected payoffs in the stationary state are also described as $u^{\text{st}}(\mathbf{x}, y) := \mathbf{p}^{\text{st}}(\mathbf{x}, y) \cdot \mathbf{u}$ and $v^{\text{st}}(\mathbf{x}, y) = -u^{\text{st}}(\mathbf{x}, y)$.*

Analysis of Nash Equilibrium

First, we provide an important theorem characterizing games between one-memory and zero-memory agents. In the following, we define $\tilde{\mathcal{X}} := 1 - \mathcal{X}$ for any function or variable \mathcal{X} .

Theorem 1 (Stationary state). *Under Def. 2, the stationary state can be described as $\mathbf{p}^{\text{st}}(\mathbf{x}, y) = (x^{\text{st}}, \tilde{x}^{\text{st}}) \otimes (y, \tilde{y}) := (x^{\text{st}} y, x^{\text{st}} \tilde{y}, \tilde{x}^{\text{st}} y, \tilde{x}^{\text{st}} \tilde{y})$. Here, x^{st} is called X’s “marginalized” strategy, a function of (\mathbf{x}, y) ;*

$$x^{\text{st}}(\mathbf{x}, y) = \frac{x_3 y + x_4 \tilde{y}}{\tilde{x}_1 y + \tilde{x}_2 \tilde{y} + x_3 y + x_4 \tilde{y}}. \quad (7)$$

Proof Sketch. We consider the stationary state condition $p_{s'}^{\text{st}} = \sum_s M_{s's} p_s^{\text{st}}$. Because Y uses zero-memory strategies, we derive $p_1^{\text{st}} + p_3^{\text{st}} = y$, $p_2^{\text{st}} + p_4^{\text{st}} = \tilde{y}$, and $p_1^{\text{st}}/p_2^{\text{st}} = p_3^{\text{st}}/p_4^{\text{st}}$. These three equations show that the stationary state is described as $\mathbf{p}^{\text{st}} = (x^{\text{st}}, \tilde{x}^{\text{st}}) \otimes (y, \tilde{y})$ with a function $x^{\text{st}}(\mathbf{x}, y)$. By substituting this equation for the stationary state condition, we obtain the mathematical expression of $x^{\text{st}}(\mathbf{x}, y)$. \square

Thm. 1 shows that in the stationary state, how each action is chosen by any one-memory strategy $\mathbf{x} \in (0, 1)^4$ can be given by a zero-memory strategy $x = x^{\text{st}}(\mathbf{x}, y)$. Here, because $x^{\text{st}}(\mathbf{x}, y)$ is obtained by compressing \mathbf{x} and y , we call it a “marginalized” strategy (following the terminology used in Press and Dyson (2012)). The representation of the marginalized strategy is as if X is using the no-memory strategy in the stationary state, but this is because opponent Y has no memory. However, note that this marginalized strategy $x^{\text{st}}(\mathbf{x}, y)$ is not a variable but a function that changes depending on the other’s strategy y . Based on Thm. 1, we obtain the Nash equilibrium as follows.

Theorem 2 (With-memory Nash equilibrium). *In with-memory games under Def. 2, $(\mathbf{x}, y) = (\mathbf{x}^*, y^*)$ is the with-memory Nash equilibrium if and only if $x^{\text{st}}(\mathbf{x}^*, y^*) = x^o$, $y^* = y^o$, and $-\tilde{x}_1^* x_4^* + \tilde{x}_2^* x_3^* \geq 0$ are satisfied.*

Proof Sketch. We consider the extreme value conditions of X and Y. Using the notation of $\mathcal{F}|_* := \mathcal{F}|_{(x,y)=(x^*,y^*)}$ for

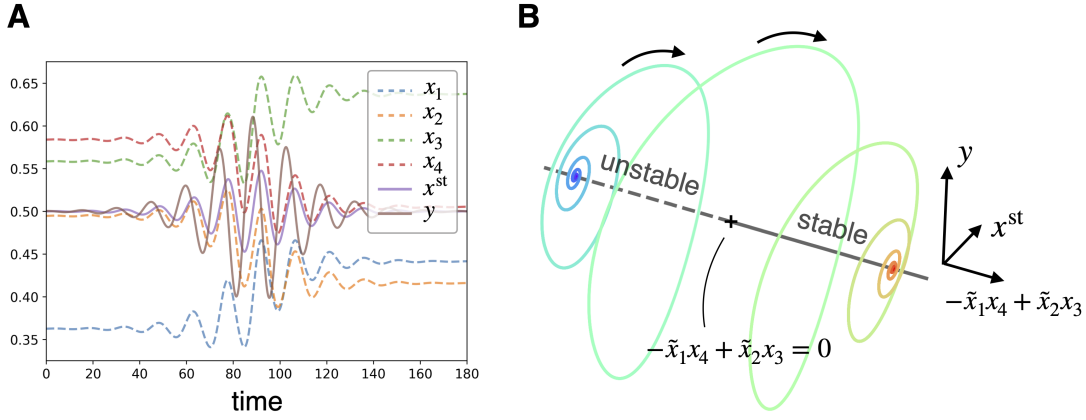


Figure 3: An example of heteroclinic orbit. A: Time series of $\mathbf{x} = (x_1, x_2, x_3, x_4)$, $x^{\text{st}}(\mathbf{x}, y)$, and y . B: Illustration of trajectories of learning dynamics. The trajectory is plotted in a three-dimensional space consisting of $x^{\text{st}}(\mathbf{x}, y)$, y , and $-\tilde{x}_1x_4 + \tilde{x}_2x_3$. Only the solid line and the cross mark (i.e., $-\tilde{x}_1x_4 + \tilde{x}_2x_3 \geq 0$) is the Nash equilibrium, but the solid and dashed gray lines indicate the states that correspond to the original Nash equilibrium, i.e., $(x^{\text{st}}(\mathbf{x}, y), y) = (x^\circ, y^\circ)$. Note that this line should be a three-dimensional manifold in the five-dimensional space of \mathbf{x} and y in practice. The solid (resp. dashed) line satisfies the stable condition $-\tilde{x}_1x_4 + \tilde{x}_2x_3 > 0$ (resp. < 0). The trajectory is plotted from blue (time is 0) to red (time is 180) by the time series in panel A.

any functions \mathcal{F} , we derive $\partial u^{\text{st}}/\partial x_i|_* = 0 \Leftrightarrow y^* = y^\circ$ for all i . Here, when $y = y^\circ$, u^{st} is always constant independent of \mathbf{x} . We also derive $\partial v^{\text{st}}/\partial y|_* = 0 \Leftrightarrow x^{\text{st}}(\mathbf{x}^*, y^*) = x^\circ$. Furthermore, the concavity condition of v^{st} for y is given by

$$\begin{aligned} \frac{\partial^2 v^{\text{st}}(\mathbf{x}^*, y)}{\partial y^2} \leq 0 &\Leftrightarrow \frac{\partial x^{\text{st}}(\mathbf{x}^*, y)}{\partial y} \geq 0 \\ &\Leftrightarrow -\tilde{x}_1^*x_4^* + \tilde{x}_2^*x_3^* \geq 0, \end{aligned} \quad (8)$$

completing the proof. \square

Let us discuss this with-memory Nash equilibrium, which is simply called Nash equilibrium below. First, (\mathbf{x}^*, y^*) is a three-dimensional manifold in the space of (\mathbf{x}, y) . In the space of $(x^{\text{st}}(\mathbf{x}, y), y)$, however, $(x^{\text{st}}(\mathbf{x}^*, y^*), y^*)$ is unique and corresponds to the original Nash equilibrium (x°, y°) . Here, note that not all of (\mathbf{x}, y) such that $(x^{\text{st}}(\mathbf{x}, y), y) = (x^\circ, y^\circ)$ are the Nash equilibrium (\mathbf{x}^*, y^*) . This is because one more condition $-\tilde{x}_1^*x_4^* + \tilde{x}_2^*x_3^* \geq 0$ is necessary for the Nash equilibrium.

We also discuss this condition $-\tilde{x}_1^*x_4^* + \tilde{x}_2^*x_3^* \geq 0$. First, recall that X uses four variables \mathbf{x} to construct a function $x^{\text{st}}(\mathbf{x}, y)$. As shown in Thm. 1, this function means the probability that X chooses action a_1 in the stationary state. Practically, however, X can change this probability depending on the other's strategy y because it has memory. $-\tilde{x}_1^*x_4^* + \tilde{x}_2^*x_3^* \geq 0$ is satisfied when x_1^* and x_3^* are large while x_2^* and x_4^* are small. In other words, X tends to use a_1 (resp. a_2) in response to b_1 (resp. b_2), meaning that X, in the next round, tries to use the advantageous action in response to the opponent's action. Briefly said, X exploits Y's payoff if Y biasedly chooses an action. Thus, it is best for Y to use its minimax strategy $y = y^*$. Indeed, $-\tilde{x}_1^*x_4^* + \tilde{x}_2^*x_3^* \geq 0$ is equivalent to $\partial^2 v^{\text{st}}(\mathbf{x}^*, y)/\partial y^2 \leq 0$, meaning that Y's function is concave. Y cannot increase its own payoff even if it uses other strategies $y \neq y^*$. Thus, this condition is necessary for the Nash equilibrium.

Analysis of Learning Dynamics

Having discussed the appearance of a Nash equilibrium in a game with asymmetric memory, our attention will now be directed towards the dynamics of the game. Specifically, We now analyze the dynamics of Eqs. (3) and (4) around the equilibrium. Our first finding is that the fixed points of the learning dynamics correspond to the original Nash equilibrium.

Theorem 3 (Fixed points of learning dynamics). *Under Def. 2, all the fixed points of learning dynamics are given by $(x^{\text{st}}(\mathbf{x}, y), y) = (x^\circ, y^\circ)$.*

Proof Sketch. Under Def. 2, Eqs. (3) and (4) are calculated as

$$\begin{aligned} \dot{x}_i &= x_i \tilde{x}_i (u_1 - u_2 - u_3 + u_4) (y - y^\circ) \frac{\partial x^{\text{st}}(\mathbf{x}, y)}{\partial x_i}, \quad (9) \\ \dot{y} &= -y \tilde{y} (u_1 - u_2 - u_3 + u_4) \\ &\quad \times \left\{ (y - y^\circ) \frac{\partial x^{\text{st}}(\mathbf{x}, y)}{\partial y} + (x^{\text{st}}(\mathbf{x}, y) - x^\circ) \right\}. \end{aligned} \quad (10)$$

We can calculate $(\dot{x}, \dot{y}) = (0, 0) \Leftrightarrow (x^{\text{st}}(\mathbf{x}, y), y) = (x^\circ, y^\circ)$. \square

Next, the following theorem and corollary specify whether dynamics converge to or diverge from these fixed points.

Theorem 4 (Local convergence to Nash equilibrium). *Under Def. 2, the learning dynamics are locally asymptotically stable for each Nash equilibrium (\mathbf{x}^*, y^*) such that $-\tilde{x}_1^*x_4^* + \tilde{x}_2^*x_3^* > 0$.*

Proof Sketch. We consider a linear stability analysis in the neighborhoods of the obtained fixed points. Let \mathbf{J} denote the Jacobian matrix for learning dynamics with $\mathbf{z} = (\mathbf{x}, y)$ and $J_{ij} := \partial \dot{z}_i / \partial z_j|_*$. Each fixed point is locally stable when the maximum eigenvalue of the Jacobian except for 0 (denoted

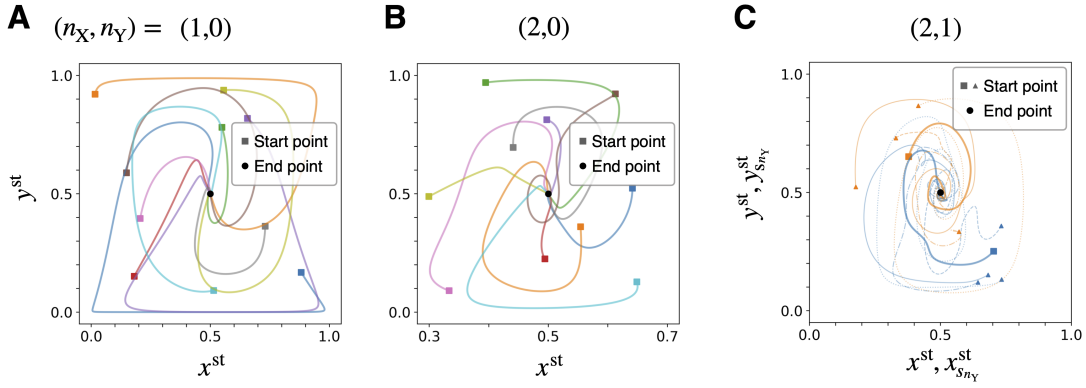


Figure 4: Learning dynamics for various asymmetric numbers of memories by removing Def. 2. Panels A and B show the cases of $(n_X, n_Y) = (1, 0)$ and $(2, 0)$, respectively. All the trajectories from various 10 random initial states (represented by different colors) converge to the same Nash equilibrium. Panel C illustrates the case of $(n_X, n_Y) = (2, 1)$. The horizontal and vertical axes indicate x^{st} and y^{st} for the wide thick lines, and X’s marginalized strategy $x_{s_{n_Y}}^{st}$ and Y’s strategy $y_{s_{n_Y}}^{st}$ for the other thin lines. The solid, dashed, dot, and dash-dot lines of the thin ones represent the cases of $s_{n_Y} (= s_1) = a_1b_1, a_1b_2, a_2b_1,$ and a_2b_2 , respectively. Both samples (represented by blue and orange) are observed to converge to the Nash equilibrium.

as λ_1) is negative. We derive $\lambda_1 < 0 \Leftrightarrow -\tilde{x}_1^*x_4^* + \tilde{x}_2^*x_3^* > 0$, completing the proof. \square

Corollary 1 (Divergence from fixed points). *Under Def. 2, if strategies \mathbf{x} and \mathbf{y} satisfy $(x^{st}(\mathbf{x}, \mathbf{y}), \mathbf{y}) = (x^o, y^o)$, the strategies are fixed points of learning dynamics even in the region of $-\tilde{x}_1x_4 + \tilde{x}_2x_3 < 0$. However, the strategies are unstable, and learning dynamics diverge from there.*

Let us explain the property of learning dynamics from Thms. 3 and 4. First, Thm. 3 means that the fixed points of learning dynamics correspond to the original Nash equilibrium. However, learning dynamics do not converge to all of these fixed points. Thm. 4 shows that learning can converge to only roughly half of the fixed points where the condition $-\tilde{x}_1x_4 + \tilde{x}_2x_3 > 0$ should be satisfied. Because the equivalence holds when all the inequalities become equalities in Eq. (8), this condition is equivalent to the strict concavity of Y’s utility function $v^{st}(\mathbf{x}, \mathbf{y})$.

Visualization of Heteroclinic Dynamics

Fig. 3 visualizes a heteroclinic orbit of learning dynamics. Panel A shows the time series of learning dynamics which is based on the continuous-time algorithm, i.e., the Runge-Kutta fourth-order method of Eqs. (3) and (4) with the step-size of 2×10^{-2} . To facilitate interpretation, Panel B shows the trajectory of the time series of five-dimensional learning dynamics of (\mathbf{x}, \mathbf{y}) in an appropriate three-dimensional space $(x^{st}, y, -\tilde{x}_1x_4 + \tilde{x}_2x_3)$. In panel B, all the gray solid and dashed lines correspond to the original Nash equilibrium, i.e., $(x^{st}, \mathbf{y}) = (x^o, y^o)$ and to the fixed points of learning dynamics. In particular, the solid (resp. dashed) line shows stable (resp. unstable) fixed points of learning dynamics. Note that these one-dimensional lines represent a three-dimensional manifold in the five-dimensional space consisting of (\mathbf{x}, \mathbf{y}) . Now, we can see a heteroclinic orbit diverging from an unstable fixed point and converging to a stable one. Panel A illustrates an increase in x_1 and x_3 , while x_2 and x_4

decrease through learning, indicating that X is learning to exploit Y. Consequently, $-\tilde{x}_1x_4 + \tilde{x}_2x_3$ increases, leading to a concave utility function for Y with respect to y .

Experimental Results

So far, our theoretical analyses of learning dynamics have revealed the existence of heteroclinic orbits that converge to stable fixed points. In addition, these fixed points are included in the original Nash equilibrium of games without memory. In the following, we numerically confirm that learning dynamics converge to the original Nash equilibrium in various initial states, action numbers, and memory lengths. All the following experimental results are based on the discrete-time algorithm, i.e., Alg. 1. In the following, we set the inputs $\eta = 10^{-3}$ and $\gamma = 10^{-6}$ in Alg. 1. In computing gradients of payoff, i.e., $\Delta_{a|s}$, we calculate the equilibrium payoff accurately enough; The analytical solution u^{st} and computational solution \hat{u}^{st} satisfy $|u^{st} - \hat{u}^{st}| \leq 10^{-9}$.

Dynamics for Various Memory Lengths. Fig. 4 shows that the learning dynamics converge to the Nash equilibrium in various numbers of memories. Panel A is for $(n_X, n_Y) = (1, 0)$, B is for $(2, 0)$, and C is for $(2, 1)$. In Panel B, because $n_X = 2$, the number of variables that are necessary to describe X’s strategy \mathbf{x} is larger than those in Def. 2. However, just like Thm. 1 shows, \mathbf{x} construct the marginalized strategy of one function $x^{st}(\mathbf{x}, \mathbf{y})$ in the stationary state. Furthermore, the equilibrium corresponds to the original Nash equilibrium, i.e., $(x^{st}(\mathbf{x}, \mathbf{y}), \mathbf{y}) = (x^o, y^o)$. Notably, Panel C of $(n_X, n_Y) = (2, 1)$ is interesting, where both the players memorize previous states as different from both A and B. X uses two-memory strategies given by sixteen variables $\mathbf{x} \in (0, 1)^{16}$, while Y uses one-memory strategies given by four variables $\mathbf{y} \in (0, 1)^4$. Then, \mathbf{x} construct the marginalized strategy of four functions and X behaves as if X employs one-memory strategies. In conclusion, learning dynamics converge to the Nash equilibrium in games between

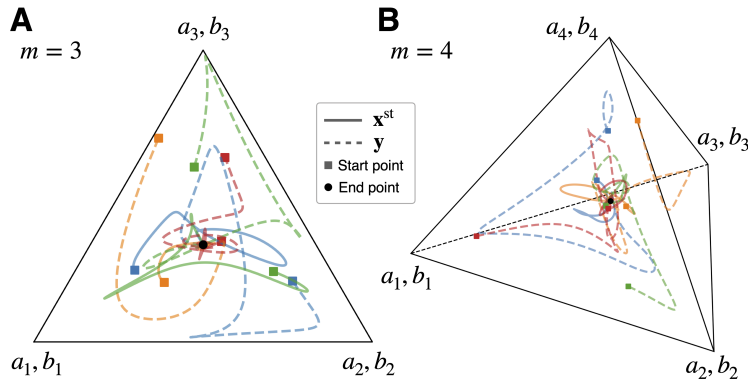


Figure 5: Learning dynamics for various numbers of actions are illustrated in two cases: (A) a rock-paper-scissors game ($m = 3$), and (B) an extended rock-paper-scissors game with four actions ($m = 4$). The dynamics of $\mathbf{x}^{\text{st}} \in \Delta^{m-1}$ and $\mathbf{y} \in \Delta^{m-1}$ are plotted in each panel. The vertices of simplexes correspond to pure strategies for choosing actions a_1, \dots, a_m in \mathbf{x}^{st} or b_1, \dots, b_m in \mathbf{y} . In both cases, the samples of 4 random initial conditions (indicated by different colors) converge to the Nash equilibrium.

one-memory strategies (Fujimoto, Ariu, and Abe 2023).

Dynamics for Various Action Numbers. Next, we remove the assumption of two-action games (Asm. 1). Fig. 5 demonstrates that learning dynamics converge in various numbers of actions, where Panel A considers a Rock-Paper-Scissors game ($m = 3$), while Panel B considers an Extended Rock-Paper-Scissors game ($m = 4$) with $(n_X, n_Y) = (1, 0)$. From these panels, we can observe that learning dynamics converge in these games as in the two-action games. We now discuss where learning dynamics converge. Just like the proof in Thm. 1, we define the probability that X chooses action a in the stationary state under m -action games as $x_a^{\text{st}}(\mathbf{x}, \mathbf{y}) := \sum_{s \in \mathcal{S}} x_{a|s} p_s^{\text{st}}(\mathbf{x}, \mathbf{y})$. We also define the set of $x_a^{\text{st}}(\mathbf{x}, \mathbf{y})$ as $\mathbf{x}^{\text{st}}(\mathbf{x}, \mathbf{y}) := \{x_a^{\text{st}}(\mathbf{x}, \mathbf{y}) \mid a \in \mathcal{A}\}$. Fig. 5 shows the last-iterate convergence to the Nash equilibrium, i.e., $(\mathbf{x}^{\text{st}}(\mathbf{x}, \mathbf{y}), \mathbf{y}) = (\mathbf{x}^{\circ}, \mathbf{y}^{\circ})$. The mechanism for this convergence is similar to that in two-action games. Indeed, even in m -action zero-sum games, the player with 1-memory endows a strict concavity to the utility function of the 0-memory player (see the Technical Appendix C for a detailed explanation).

Experimental Results with Many Samples. In order to ensure the reliability of our results, we have simulated the learning dynamics with a large number of samples with random initial strategies. In all the samples, we observed convergence to the Nash equilibrium. Refer to the Technical Appendix B for detailed data.

Discussion

This study explores the dynamics of learning in games where agents possess different memory capacities, both theoretically and experimentally. In our theoretical contributions, we assumed that two agents with memory lengths of one and zero play two-action zero-sum games. Under this assumption, we proved that the original Nash equilibrium in games without memory is divided into stable (Thm. 4) and unstable fixed points (Cor. 1) for learning. Here, hetero-

clinic orbits are discovered in learning in games with memory asymmetry; their strategies diverge from these unstable fixed points and converge to stable ones. In our experimental contributions, we relaxed the assumption on memory lengths and the number of actions, and then we demonstrated that memory asymmetry triggers this convergence to the original Nash equilibrium across a wider range of action numbers and memory lengths. Considering that dynamics show cycling behavior in games without memory and divergence in games with memory symmetry, this convergence is a non-trivial phenomenon.

Note that we have discovered new insights into computing equilibria, with memory asymmetry serving as one example. Convergent algorithms and dynamics, often referred to as last-iterate convergence, are extensively studied within the community of learning in games. The focus is mainly on computationally efficient algorithms with faster rates of convergence or global convergence guarantees (Daskalakis and Panageas 2019; Golowich, Pattathil, and Daskalakis 2020; Mertikopoulos et al. 2019; Wei et al. 2021; Lei et al. 2021; Nguyen et al. 2021; Anagnostides et al. 2022; Abe, Sakamoto, and Iwasaki 2022). In the majority of these studies, convergence is ensured through modifications to the learning algorithm. A typical example is to incorporate “optimism” symmetrically into discretized gradient ascent. Each agent then refines its strategy by naively predicting the opponent’s strategy update from past data (Daskalakis and Panageas 2019). Our method to achieve convergence differs significantly. We introduced memory asymmetry to alter the strategy spaces of one or both agents. By leveraging memory asymmetry, an agent with a longer memory learns to exploit the opponent’s payoff. Consequently, the opponent agent with a shorter memory achieves a strictly concave utility function, leading to convergence. To the best of our knowledge, this is the first study to induce convergence through such asymmetry of the strategy spaces. Our findings may open up new possibilities for achieving convergence in learning games and inspire further research in this field.

Acknowledgements

Y.F. is supported by JSPS KAKENHI Grant No. 21J01393. K.Ariu is supported by JSPS KAKENHI Grant No. 23K19986.

References

- Abe, K.; Sakamoto, M.; and Iwasaki, A. 2022. Mutation-Driven Follow the Regularized Leader for Last-Iterate Convergence in Zero-Sum Games. In *UAI*, 1–10.
- Anagnostides, I.; Panageas, I.; Farina, G.; and Sandholm, T. 2022. On last-iterate convergence beyond zero-sum games. In *ICML*, 536–581.
- Axelrod, R.; and Hamilton, W. D. 1981. The evolution of cooperation. *Science*, 211(4489): 1390–1396.
- Baek, S. K.; Jeong, H.-C.; Hilbe, C.; and Nowak, M. A. 2016. Comparing reactive and memory-one strategies of direct reciprocity. *Scientific reports*, 6(1): 25676.
- Barfuss, W. 2020. Reinforcement learning dynamics in the infinite memory limit. In *AAMAS*, 1768–1770.
- Barfuss, W.; Donges, J. F.; and Kurths, J. 2019. Deterministic limit of temporal difference reinforcement learning for stochastic games. *Physical Review E*, 99(4): 043305.
- Börger, T.; and Sarin, R. 1997. Learning through reinforcement and replicator dynamics. *Journal of Economic Theory*, 77(1): 1–14.
- Bowling, M. 2004. Convergence and no-regret in multiagent learning. In *NeurIPS*, 209–216.
- Bowling, M.; and Veloso, M. 2002. Multiagent learning using a variable learning rate. *Artificial Intelligence*, 136(2): 215–250.
- Daskalakis, C.; and Panageas, I. 2019. Last-iterate convergence: Zero-sum games and constrained min-max optimization. In *ITCS*, 27:1–27:18.
- Fudenberg, D.; and Levine, D. K. 1998. *The theory of learning in games*, volume 2. MIT press.
- Fudenberg, D.; and Maskin, E. 2009. The folk theorem in repeated games with discounting or with incomplete information. In *A long-run collaboration on long-run games*, 209–230. World Scientific.
- Fujimoto, Y.; Ariu, K.; and Abe, K. 2023. Learning in Multi-Memory Games Triggers Complex Dynamics Diverging from Nash Equilibrium. In *IJCAI*.
- Fujimoto, Y.; and Kaneko, K. 2019. Emergence of exploitation as symmetry breaking in iterated prisoner’s dilemma. *Physical Review Research*, 1(3): 033077.
- Fujimoto, Y.; and Kaneko, K. 2021. Exploitation by asymmetry of information reference in coevolutionary learning in prisoner’s dilemma game. *Journal of Physics: Complexity*, 2(4): 045007.
- Golowich, N.; Pattathil, S.; and Daskalakis, C. 2020. Tight last-iterate convergence rates for no-regret learning in multi-player games. In *NeurIPS*, 20766–20778.
- Hilbe, C.; Martinez-Vaquero, L. A.; Chatterjee, K.; and Nowak, M. A. 2017. Memory-n strategies of direct reciprocity. *Proceedings of the National Academy of Sciences*, 114(18): 4715–4720.
- Hofbauer, J.; Sigmund, K.; et al. 1998. *Evolutionary games and population dynamics*. Cambridge university press.
- Lei, Q.; Nagarajan, S. G.; Panageas, I.; et al. 2021. Last iterate convergence in no-regret learning: constrained min-max optimization for convex-concave landscapes. In *AISTATS*, 1441–1449.
- Mertikopoulos, P.; Lecouat, B.; Zenati, H.; Foo, C.-S.; Chandrasekhar, V.; and Piliouras, G. 2019. Optimistic mirror descent in saddle-point problems: Going the extra(-gradient) mile. In *ICLR*.
- Mertikopoulos, P.; Papadimitriou, C.; and Piliouras, G. 2018. Cycles in adversarial regularized learning. In *SODA*, 2703–2717.
- Mertikopoulos, P.; and Sandholm, W. H. 2016. Learning in games via reinforcement and regularization. *Mathematics of Operations Research*, 41(4): 1297–1324.
- Meylahn, J. M.; Janssen, L.; et al. 2022. Limiting dynamics for Q-learning with memory one in symmetric two-player, two-action games. *Complexity*, 2022.
- Murase, Y.; and Baek, S. K. 2020. Five rules for friendly rivalry in direct reciprocity. *Scientific reports*, 10(1): 16904.
- Nash Jr, J. F. 1950. Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences*, 36(1): 48–49.
- Nguyen, T.-D.; Zemhoho, A. B.; Tran-Thanh, L.; et al. 2021. Last round convergence and no-dynamic regret in asymmetric repeated games. In *ALT*, 553–577.
- Press, W. H.; and Dyson, F. J. 2012. Iterated Prisoner’s Dilemma contains strategies that dominate any evolutionary opponent. *Proceedings of the National Academy of Sciences*, 109(26): 10409–10413.
- Sandholm, T. W.; and Crites, R. H. 1996. Multiagent reinforcement learning in the iterated prisoner’s dilemma. *Biosystems*, 37(1-2): 147–166.
- Sato, Y.; Akiyama, E.; and Farmer, J. D. 2002. Chaos in learning a simple two-person game. *Proceedings of the National Academy of Sciences*, 99(7): 4748–4751.
- Schmid, L.; Hilbe, C.; Chatterjee, K.; and Nowak, M. A. 2022. Direct reciprocity between individuals that use different strategy spaces. *PLoS Computational Biology*, 18(6): e1010149.
- Singh, S.; Kearns, M. J.; and Mansour, Y. 2000. Nash Convergence of Gradient Dynamics in General-Sum Games. In *UAI*, 541–548.
- Strogatz, S. H. 2018. *Nonlinear dynamics and chaos with student solutions manual: With applications to physics, biology, chemistry, and engineering*. CRC press.
- Ueda, M. 2023. Memory-two strategies forming symmetric mutual reinforcement learning equilibrium in repeated prisoners’ dilemma game. *Applied Mathematics and Computation*, 444: 127819.
- Usui, Y.; and Ueda, M. 2021. Symmetric equilibrium of multi-agent reinforcement learning in repeated prisoner’s dilemma. *Applied Mathematics and Computation*, 409: 126370.

- Wei, C.-Y.; Lee, C.-W.; Zhang, M.; and Luo, H. 2021. Linear Last-iterate Convergence in Constrained Saddle-point Optimization. In *ICLR*.
- Zinkevich, M. 2003. Online convex programming and generalized infinitesimal gradient ascent. In *ICML*, 928–936.