

# On the Heterogeneity of Independent Learning Dynamics in Zero-sum Stochastic Games

**Muhammed O. Sayin**

SAYIN@EE.BILKENT.EDU.TR

*Bilkent University, Electrical and Electronics Engineering Department, Ankara, Turkey 06800*

**K. Alperen Cetiner**

ALPEREN.CETINER@BILKENT.EDU.TR

*Bilkent University, Electrical and Electronics Engineering Department, Ankara, Turkey 06800*

*Aselsan A.Ş., Ankara, Turkey 06800*

**Editors:** R. Firoozi, N. Mehr, E. Yel, R. Antonova, J. Bohg, M. Schwager, M. Kochenderfer

## Abstract

We analyze the convergence properties of the two-timescale fictitious play combining the classical fictitious play with the  $Q$ -learning for two-player zero-sum stochastic games with player-dependent learning rates. We show its almost sure convergence under the standard assumptions in two-timescale stochastic approximation methods when the discount factor is less than the product of the ratios of player-dependent step sizes. To this end, we formulate a novel Lyapunov function formulation and present a one-sided asynchronous convergence result.

**Keywords:**  $Q$ -learning, stochastic games, multi-agent learning, heterogenous systems

## 1. Introduction

Multi-agent reinforcement learning has become the frontier of many advancements in artificial intelligence systems, where autonomous agents make decisions in dynamic environments (e.g., see [Zhang et al. \(2021\)](#) and the references therein). Heterogeneity and independence of the learning dynamics adopted by these autonomous agents are inevitable in practical applications of multi-agent systems. However, there has been very limited progress addressing it in the context of stochastic games (introduced by [Shapley \(1953\)](#)) - a canonical model for dynamic multi-agent interactions.

Independent learning in strategic-form games played repeatedly has been studied extensively with many well-established results, e.g., see ([Fudenberg and Levine, 1998](#); [Young, 2004](#); [Fudenberg and Levine, 2009](#)). On the other hand, for stochastic games, [Arslan and Yuksel \(2017\)](#); [Wei et al. \(2017\)](#) presented learning dynamics with double-loop-like update rules necessitating coordination among players that may not be inline with their best interests. Recently, [Leslie et al. \(2020\)](#) has drawn a two-timescale learning framework in which continuous-time best-response dynamics could also convergence to an equilibrium of a two-player zero-sum stochastic game though the players' learning dynamics are not completely independent since they track a common parameter together. Within the two-timescale learning framework, [Sayin et al. \(2020\)](#) presented independent learning dynamics for stochastic games and analyzed its almost-sure convergence also in two-player zero-sum stochastic games. Note that [Ozdaglar et al. \(2022\)](#) provides an overview of studies on independent learning dynamics in stochastic games.

Heterogenous learning in games has also been studied, however, with a specific focus on multi-timescale learning, e.g., ([Leslie and Collins, 2003, 2005](#)) for strategic-form games with repeated

play and recently (Daskalakis et al., 2020) for two-player zero-sum stochastic games. Particularly, dynamics of different players evolve at different timescales, e.g., one player’s dynamics evolve slower than the others, contrary to (Leslie et al., 2020; Sayin et al., 2020) where the two-timescale framework is at player-level. On heterogenous rates that may not lead to different timescales, Zhu et al. (2011) studies heterogenous learning in a special case of zero-sum stochastic games while Chasnov et al. (2020) studies heterogenous gradient-based learning dynamics in continuous games.

In this paper, we address the heterogeneity and independence of learning in stochastic games by characterizing the convergence properties of the independent learning dynamics presented in (Sayin et al., 2020) with player-dependent learning rates. This dynamics is a new variant of fictitious play combining the classical fictitious play (Fudenberg and Levine, 1998) with the  $Q$ -learning (Watkins and Dayan, 1992) while they evolve at two different timescales. The key idea is that the underlying stochastic game can be viewed as a collection of auxiliary stage-games specific to each state whose payoff functions are the  $Q$ -functions. Though these auxiliary stage-games are not necessarily stationary, the slow evolution of  $Q$ -function estimates make them relatively stationary. The key challenge is the deviation of these auxiliary stage-games from the zero-sum structure due to the independent update of the  $Q$ -function estimates, and heterogenous learning rates boost this deviation further. We show the almost sure convergence of the dynamics under the usual two-timescale stochastic approximation assumptions when the discount factor is less than the product of the ratios of player-dependent step sizes. We elaborate on the implications and high-level interpretation of this result later in Section 3. To show this result, we formulate a novel Lyapunov function formulation which reduces to the one presented in (Sayin et al., 2020) in the homogenous case, and present a one-sided asynchronous convergence result, which has a similar flavor with (Tsitsiklis, 1994, Theorem 1).

We emphasize that the dynamics presented is different from equilibrium computation methods such as Shapley’s value iteration (Shapley, 1953) or its model-free version Minimax-Q algorithm (Littman, 1994) by not requiring players to know the opponent’s objective, i.e., the zero-sum structure of their stage-payoffs. Our dynamics also differs from multi-timescale learning schemes such as (Leslie and Collins, 2003; Daskalakis et al., 2020) by addressing possible heterogeneity in (comparable) learning rates of players rather than exploiting it to characterize its convergence behavior especially since a convergence result with common learning rates is challenging. Furthermore, Daskalakis et al. (2020) considers the repeated play of a stochastic game with indefinite termination time allowing players to revise and improve their policies across repetitions. On the other hand, here we focus on the convergence of beliefs within a single stochastic game played over infinite horizon.

The rest of the paper is organized as follows. In Section 2, we introduce stochastic games and describe the learning dynamics. We present the assumptions and main convergence result in Section 3 and the proof of the main convergence result in Section 4. In Section 5, we provide an illustrative example. We conclude the paper with some remarks in Section 6. The extended version (Sayin and Cetiner, 2021) includes the proofs of the technical lemmas used in Section 4.

## 2. Independent Learning in (Zero-sum) Stochastic Games

Formally, a two-player stochastic game is characterized by a tuple  $\langle S, A, r^1, r^2, p, \gamma \rangle$ .<sup>1</sup> The *finite* set of states is denoted by  $S$  while  $A = A^1 \times A^2$  with  $A^i$  denoting the *finite* set of actions that player

---

1. For easy referral, we index players as player 1 and player 2. Furthermore, player  $i$  is the typical player and player  $-i$  is her opponent.

$i$  can take at any state.<sup>2</sup> The *stage payoff function* of player  $i$  is denoted by  $r^i : S \times A \rightarrow \mathbb{R}$ . In zero-sum case, we have  $r^1(s, a) + r^2(s, a) = 0$  for all  $(s, a) \in S \times A$ . At any stage  $k = 0, 1, \dots$ , if players play the action profile  $a = (a^1, a^2) \in A$ , then the state of the game,  $s$ , transits to another state,  $s'$ , according to the transition probability  $p(s'|s, a)$ . Player  $i$ 's objective is to maximize her expected sum of discounted stage-payoffs with the discount factor  $\gamma \in [0, 1)$ .<sup>3</sup>

Shapley (1953) (and Fink (1964)) showed that in two-player zero-sum (and  $n$ -player general-sum) stochastic games, there always exists a *stationary equilibrium* where players play *stationary* strategies depending only on the current state. Let  $\pi^i : S \rightarrow \Delta(A^i)$  denote the stationary strategy of player  $i$  such that  $\pi^i(s) \in \Delta(A^i)$  corresponds to her mixed strategy at state  $s$ , and  $\pi = (\pi^1, \pi^2)$  denote the (stationary) strategy profile of players.<sup>4</sup> Then, the expected discounted sum of stage payoff of player  $i$  under the strategy profile  $\pi$  is given by

$$U^i(\pi) := \mathbb{E} \left\{ \sum_{k=0}^{\infty} \gamma^k r^i(s_k, a_k) \right\}, \quad (1)$$

where  $a_k \sim \pi(s_k)$  denotes the action profile at stage  $k$  while  $\{s_k\}_{k \geq 0}$  is a stochastic process such that  $s_k$  represents the state at stage  $k$ . The expectation is taken with respect to all randomness.

**Definition 1 (Stationary Nash Equilibrium)** *We say that a stationary strategy profile  $\pi_*$  is a stationary mixed-strategy equilibrium of the two-player stochastic game provided that*

$$U^i(\pi_*) \geq U^i(\pi^i, \pi_*^{-i}), \quad \forall \pi^i \text{ and } i = 1, 2. \quad (2)$$

We consider the same (independent) learning dynamics presented in (Sayin et al., 2020) but with player-dependent learning rates to examine the robustness of the convergence result to such heterogeneity. Particularly, we can view the stage-wise interaction among players as they are playing *auxiliary stage-games* specific to current state. For example, if player  $i$  knew that the opponent will play according to the stationary strategy  $\pi^{-i}$  in future stages, then her payoff function in the auxiliary stage-game specific to state  $s$ , denoted by  $Q^i(s, \cdot) : A \rightarrow \mathbb{R}$ , would satisfy the following fixed-point condition

$$Q^i(s, a) = r^i(s, a) + \gamma \sum_{s' \in S} p(s'|s, a) \max_{\tilde{a}^i \in A^i} \mathbb{E}_{\tilde{a}^{-i} \sim \pi^{-i}(s')} \{Q^i(s', \tilde{a})\}, \quad \forall a \in A. \quad (3)$$

This follows from backward induction that player  $i$  would always look for maximizing her utility, as described in (1). Note that the dependence on  $\pi^{-i}$  is implicit for notational convenience. The function  $Q^i(\cdot)$  is known as *Q-function* in the MDP or reinforcement learning literature (Filar and Vrieze, 1997; Sutton and Barto, 2018), and it is well-defined due to the contraction property of the Bellman operator. Therefore, the auxiliary stage game is the tuple  $\langle A, Q^1(s, \cdot), Q^2(s, \cdot) \rangle$ .

Players form a belief about their  $Q$ -function and adopt fictitious play in auxiliary stage-games by forming also a belief on the opponent strategy. We denote the beliefs of player  $i$  at stage  $k$  about the opponent strategy by  $\pi_k^{-i}$  and about her  $Q$ -function by  $Q_k^i$ . At each stage  $k$ , player  $i$

2. The formulation can be extended to state-variant action sets straightforwardly.

3. Without a common discount factor, the underlying stochastic game may not be a zero-sum game even when  $r^1(s, a) + r^2(s, a) = 0$  for all  $(s, a) \in S \times A$ .

4. Given a set  $A$ , we denote the probability simplex over  $A$  by  $\Delta(A)$ .

always take the best response action in the auxiliary stage-game based on her beliefs  $\pi_k^{-i}(s_k)$  and  $Q_k^i(s_k) := Q_k^i(s_k, \cdot)$ . Therefore, her action  $a_k^i \in A^i$  always satisfies

$$a_k^i \in \operatorname{argmax}_{a^i \in A^i} \mathbb{E}_{a^{-i} \sim \pi_k^{-i}(s_k)} \{Q_k^i(s_k, a)\}. \quad (4)$$

Player  $i$  can observe the opponent's action  $a_k^{-i} \in A^{-i}$  and update her beliefs according to

$$\pi_{k+1}^{-i}(s) = \pi_k^{-i}(s) + \mathbb{I}_{\{s=s_k\}} \alpha_{c_k(s)}^i (a_k^{-i} - \pi_k^{-i}(s)), \quad (5a)$$

$$Q_{k+1}^i(s, a) = Q_k^i(s, a) + \mathbb{I}_{\{s=s_k\}} \beta_{c_k(s)}^i \left( r^i(s, a) + \gamma \sum_{s' \in S} p(s'|s, a) v_k^i(s') - Q_k^i(s, a) \right), \quad (5b)$$

for all  $(s, a)$ , where  $\mathbb{I}_{\{s=s_k\}}$  is the indicator function, we let the pure action  $a_k^{-i}$  be a deterministic strategy in the probability simplex  $\Delta(A^{-i})$ , the value function estimate  $v_k^i : S \rightarrow \mathbb{R}$  is defined by

$$v_k^i(s) = \max_{a^i \in A^i} \mathbb{E}_{a^{-i} \sim \pi_k^{-i}(s)} \{Q_k^i(s, a)\}, \quad (6)$$

and the player-dependent step sizes  $\alpha_c^i \in (0, 1)$  and  $\beta_c^i \in (0, 1)$  vanish with  $c_k(s)$ , the number of times state  $s$  is visited until stage  $k$ .

### 3. Convergence Results

The dynamics in (5) allow player-dependent and belief-dependent step sizes. In this section, we identify the conditions under which such a heterogenous two-timescale learning dynamics is guaranteed to converge to an equilibrium of the underlying two-player zero-sum stochastic game.

**Assumption 2 (Markov Chain)** *Every state is visited infinitely often with probability 1.*

Beliefs associated with a state gets updated only if that state is visited. This assumption ensures that beliefs associated with each state gets updated infinitely often. Importantly, state transitions are controlled by players while players take actions according to their beliefs. Without Assumption 2, it might be possible that they form incorrect beliefs about the value of other states such that these states would not get visited due to their greedy best responses (though they could have been visited for some non-greedy action). Then, the players are not able to revise and improve their beliefs about the values of these states. Therefore, their beliefs may converge to a *self-confirming equilibrium* (a concept introduced for learning in extensive-form games to address out-of-equilibrium actions (Fudenberg and Kreps, 1995)) rather than a Nash equilibrium. We leave it as a future research direction.

Furthermore, Assumption 2 holds if the underlying stochastic game is *irreducible*, e.g., transition probabilities between any pair of states are positive for any joint action as in Leslie et al. (2020). This can be a restrictive assumption in practical application when players take deterministic best responses. However, it can be relaxed further as discussed in Ozdaglar et al. (2022) if players choose strategies in which every action is taken with some positive probability, e.g., due to smoothed best response or exploration, especially in the model-free cases where players do not know the stage-payoff and state transition kernel. We leave them as future research directions.

**Assumption 3 (Step Sizes)** *The step sizes satisfy the following conditions:*

- (a) For each  $i = 1, 2$ , the step sizes  $\alpha_c^i \rightarrow 0$ ,  $\beta_c^i \rightarrow 0$ , and  $\beta_c^i/\alpha_c^i \rightarrow 0$  as  $c \rightarrow \infty$ , and the series  $\sum_{c=0}^{\infty} \alpha_c^i = \sum_{c=0}^{\infty} \beta_c^i = \infty$ .
- (b) Let  $i, j \in \{1, 2\}$  denote arbitrary player indices with  $i \neq -i$  and  $j \neq -j$ . Then, for some  $d_\alpha, d_\beta \in (0, 1]$ , we have  $\lim_{c \rightarrow \infty} \alpha_c^i/\alpha_c^{-i} = d_\alpha$  and  $\lim_{c \rightarrow \infty} \beta_c^j/\beta_c^{-j} = d_\beta$ .<sup>5</sup>

In Assumption 3, part-(a) is standard in multi-timescale stochastic approximation, e.g., see (Borkar, 2008, Chapter 6). Note that we do not need the assumption that the step sizes are square-summable because there is no stochastic approximation error (e.g., that can be induced from sampling noise) in the differential inclusion approximation of the discrete-time update rule. The part-(b) says that the ratios of the step sizes have a non-zero limit at each timescale, i.e., players use comparable learning rates with each other. The limit assumption can be relaxed into conditions on limit inferior and limit superior, necessitating more involved analysis. We leave it as a future research direction.

**Theorem 4 (Convergence Result)** *Given a two-player zero-sum stochastic game, suppose that players follow the heterogenous and independent learning dynamics (5). Under Assumptions 2 and 3, we have  $\pi_k^i \rightarrow \pi_*^i$  and  $Q_k^i \rightarrow Q_*^i$  for each  $i = 1, 2$  as  $k \rightarrow \infty$  with probability 1, for some stationary equilibrium  $\pi_*$  and the associated  $Q$ -functions  $(Q_*^1, Q_*^2)$  provided that  $\gamma \leq d_\alpha d_\beta$ .*

We attribute the upper bound on the discount factor as the heterogeneity increases how much the auxiliary stage-games can deviate from the zero-sum structure. However, a small discount factor will compensate this by restraining the deviation since the stage-payoffs have zero-sum structure. Furthermore, when players have (asymptotically) common step sizes, i.e.,  $d_\alpha = d_\beta = 1$ , the bound on the discount factor becomes  $\gamma < 1$ , which is inherent to the discounted stochastic games. Therefore, Theorem 4 reduces to the convergence result (Sayin et al., 2020, Theorem 4.3) for the special case of homogeneous learning rates.

We can interpret Theorem 4 at a high level as players could reach to an equilibrium through heterogenous and independent learning dynamics if they are sufficiently myopic so that they discount the impact of future stages more in their utilities. We can also interpret the discount factor as the continuation probability of the stochastic game with indefinite horizon length (Shapley, 1953). Therefore, the heterogenous learning dynamics is guaranteed to converge to an equilibrium if the game has sufficiently short expected termination time.

Note also that stochastic games turn into strategic-form games with repeated play if there is only one state and  $\gamma = 0$ . In that case, the assumption on  $\gamma$  is always satisfied and the auxiliary stage-games are always zero-sum. Therefore, we have the following corollary to Theorem 4.

**Corollary 5 (Heterogenous Fictitious Play)** *Consider a two-player zero-sum strategic-form game played repeatedly. Suppose that players follow the fictitious play dynamics with player-dependent learning rates  $\alpha_k^1, \alpha_k^2$  such that their ratio has non-zero limit. Then, the beliefs formed about the opponent strategy converge to an equilibrium of the game.*

5. This is without loss of generality because if  $\lim_{c \rightarrow \infty} \alpha_c^1/\alpha_c^2 \geq 1$  then we have  $\lim_{c \rightarrow \infty} \alpha_c^2/\alpha_c^1 \leq 1$ .

#### 4. Proof of the Convergence Result - Theorem 4

In two-player zero-sum stochastic games, [Shapley \(1953\)](#) provided a (minimax) value iteration to compute equilibrium values associated with a stationary equilibrium. The operator used in Shapley's value iteration can be transformed into

$$(\mathcal{F}^i Q^i)(s, a) = r^i(s, a) + \gamma \sum_{s' \in S} p(s'|s, a) \text{val}^i(Q^i(s', \cdot)), \quad \forall (s, a) \in S \times A, \quad (7)$$

as in ([Szepesvari and Littman, 1999](#)), where the minimax value function  $\text{val}^i(\cdot)$  is defined by

$$\text{val}^i(Q^i(s, \cdot)) := \max_{\mu^i \in \Delta(A^i)} \min_{\mu^{-i} \in \Delta(A^{-i})} \mathbb{E}_{(a^i, a^{-i}) \sim (\mu^i, \mu^{-i})} \{Q^i(s, a)\}. \quad (8)$$

The operator  $\mathcal{F}^i$  is a contraction and its unique fixed point  $Q_*^i = \mathcal{F}^i Q_*^i$  is the equilibrium  $Q$ -function of the underlying stochastic game.

In this proof, we look for identifying the conditions under which

$$e_k^i(s) := v_k^i(s) - \text{val}^i(Q_k^i(s)) \rightarrow 0 \quad \text{and} \quad \tilde{Q}_k^i(s, a) := Q_k^i(s, a) - Q_*^i(s, a) \rightarrow 0, \quad (9)$$

for each  $(s, a)$  and  $i = 1, 2$ , by using stochastic differential inclusion theory while formulating a novel Lyapunov function and one-sided convergence result to address the heterogeneity. To this end, firstly, the following lemma establishes the connection between (5) and its (continuous-time) limiting differential inclusion based on ([Benaim et al., 2005](#)).<sup>6</sup> The proof is deferred to Appendix A in the extended version ([Sayin and Cetiner, 2021](#)).

**Lemma 6** *For each state  $s$ , the limiting differential inclusion of (5) is given by*

$$\dot{\pi}^1 + d_\alpha \pi^1 \in d_\alpha \operatorname{argmax}_{a^1 \in A^1} \mathbb{E}_{a^2 \sim \pi^2} \{Q^1(a)\} \quad \text{and} \quad \dot{\pi}^2 + \pi^2 \in \operatorname{argmax}_{a^2 \in A^2} \mathbb{E}_{a^1 \sim \pi^1} \{Q^2(a)\} \quad (10)$$

and  $\dot{Q}^i(a) = 0$  for all  $(i, a)$ , where  $\pi^i : [0, \infty) \rightarrow \Delta(A^i)$  and  $Q^i(a) : [0, \infty) \rightarrow \mathbb{R}$ , for each  $i = 1, 2$ , are continuous-time functions where we drop the dependence on  $s$  for notational convenience.

We note that there always exists an absolutely continuous solution to (10) since the best response satisfies the conditions listed in ([Benaim et al., 2005](#), Hypothesis 1.1). Then, ([Benaim et al., 2005](#), Theorem 3.6 and Proposition 3.27) yield that we can characterize the convergence properties of the discrete-time dynamics (5) in terms of the zero-set of a Lyapunov function to the differential inclusion (10).

Though (10) resembles to continuous-time best response dynamics in the auxiliary stage-game with time-invariant payoff functions  $Q^1$  and  $Q^2$ , there are two challenges: (i) the heterogeneity when  $d_\alpha \in (0, 1)$  and (ii) the deviation from the zero-sum structure since  $Q_k^1(s, a) + Q_k^2(s, a)$  is not necessarily zero for all  $(s, a)$  and  $k$  when they are updated independently according to (5b). Our candidate Lyapunov function is defined by

$$V(\pi, Q) := (d_\alpha \Delta^1(\pi^2, Q^1) + \Delta^2(\pi^1, Q^2) - \Xi(Q^1, Q^2))_+, \quad (11)$$

6. Without loss of generality, suppose that  $\lim_{c \rightarrow \infty} \alpha_c^1 / \alpha_c^2 = d_\alpha \in (0, 1]$ .

where we define

$$\Delta^i(\pi^{-i}, Q^i) := \max_{a^i \in A^i} \mathbb{E}_{a^{-i} \sim \pi^{-i}} \{Q^i(a)\} - \text{val}^i(Q^i) \geq 0 \quad (12)$$

$$\Xi(Q^1, Q^2) := \lambda \|Q^1 + Q^2\| - (\text{val}^1(Q^1) + \text{val}^2(Q^2)), \quad (13)$$

where  $\lambda \in (1, d_\alpha d_\beta / \gamma)$ ,  $\|\cdot\|$  is the maximum norm, i.e.,  $\|Q\| = \max_a |Q(a)|$  and the positive function  $(x)_+ = \max\{0, x\}$ . The candidate function is non-negative by its definition. Its zero-set  $\{(\pi, Q) : V(\pi, Q) = 0\}$  is given by

$$\{(\pi, Q) : \Delta^1(\pi^2, Q^1) + \Delta^2(\pi^1, Q^2) \leq \Xi(Q^1, Q^2) + (1 - d_\alpha) \Delta^1(\pi^2, Q^1)\} \quad (14)$$

and  $\Delta^i$  is the continuous-time counterpart of the tracking error (9). Therefore, the convergence to the zero-set (14) would provide an (asymptotic) upper bound on the sum of tracking errors. Note also that when  $d_\alpha = 1$ , minimax values disappear and  $V(\cdot)$  reduces to the one presented in (Sayin et al., 2020).

Since  $Q^i$ 's are time-invariant,  $\text{val}^i(Q^i)$ 's and  $\Xi$  are also time-invariant. As shown in the following lemma, these time-invariant terms together with the positive function play an important role for the validity of the candidate  $V(\cdot)$  as a Lyapunov function to (10) for the zero-set (14) and later in characterizing the convergence properties of the tracking error. The proof is deferred to Appendix B in the extended version (Sayin and Cetiner, 2021).<sup>7</sup>

**Lemma 7 (Lyapunov Function)** *The candidate function  $V(\cdot)$  is a Lyapunov function of the differential inclusion (10) for the zero-set  $\{(\pi, Q) : V(\pi, Q) = 0\}$ . In other words, for any absolutely continuous solution to (10), we have  $V(\pi(t'), Q(t')) < V(\pi(t), Q(t))$  for all  $t' > t$  if  $V(\pi(t), Q(t)) > 0$ , and  $V(\pi(t'), Q(t')) = 0$  for all  $t' > t$  if  $V(\pi(t), Q(t)) = 0$ .*

Based on the stochastic differential inclusion theory Benaim et al. (2005), Lemma 7, the definition of  $v_k^i$ , as described in (6), and the zero-set (14), we can conclude that

$$(\bar{v}_k(s) - \lambda \|\bar{Q}_k(s)\| - (1 - d_\alpha) e_k^1(s))_+ \rightarrow 0, \quad (15)$$

where we define  $\bar{v}_k(s) := v_k^1(s) + v_k^2(s)$  and  $\bar{Q}_k(s) := Q_k^1(s) + Q_k^2(s)$  with the tracking error  $e_k^1$ , as described in (9). Since we let  $i = 1$  arbitrarily for notational convenience, (15) can be written as

$$\bar{v}_k(s) \leq \lambda \|\bar{Q}_k(s)\| + (1 - d_\alpha) e_k^i(s) + \epsilon_k(s), \quad (16)$$

where  $\epsilon_k(s) \rightarrow 0$  as  $k \rightarrow \infty$  almost surely. The upper bound on  $\bar{v}_k(s)$  is in terms of  $\|\bar{Q}_k(s)\|$  and the tracking error  $e_k^i$ . The following lemma provides an upper bound on the tracking error so that we can obtain an upper bound in terms of  $\bar{Q}_k(s)$  only. The proof is deferred to Appendix C in the extended version (Sayin and Cetiner, 2021).

**Lemma 8** *We have  $0 \leq e_k^i(s) = v_k^i(s) - \text{val}^i(Q_k^i(s)) \leq \bar{v}_k(s) - \min_{a \in A} \bar{Q}_k(s, a)$  for every  $s$ .*

Based on (16) and Lemma 8, we obtain

$$\min_{a \in A} \bar{Q}_k(s, a) \leq \bar{v}_k(s) \leq \frac{\lambda}{d_\alpha} \|\bar{Q}_k(s)\| - \frac{1 - d_\alpha}{d_\alpha} \min_{a \in A} \bar{Q}_k(s, a) + \frac{1}{d_\alpha} \epsilon_k(s), \quad (17)$$

7. The condition that  $r^1(s, a) + r^2(s, a) = 0$  for all  $(s, a)$  plays an important role in the proof of Lemma 7.

where the lower bound follows since  $v_k^i(s) \geq \mathbb{E}_{a \sim \pi_k(s)} \{Q_k^i(s, a)\}$  for  $i = 1, 2$ . We emphasize that in the homogenous case, the evolution of  $\bar{Q}_k(s, a)$  is given by

$$\bar{Q}_{k+1}(s, a) = \bar{Q}_k(s, a) + \mathbb{I}_{\{s=s_k\}} \beta_{c_k(s)} \left( \gamma \sum_{s' \in S} p(s'|s, a) \bar{v}_k(s') - \bar{Q}_k(s, a) \right) \quad (18)$$

due to the symmetry that  $\beta_c^i = \beta_c$ . However, the characterization of the convergence properties of  $\bar{Q}_k(s, a)$  in the heterogenous case requires more involved analysis where we will address the limit inferior and limit superior of  $\bar{Q}_k(s, a)$  separately.

Recall the definition of  $\tilde{Q}_k^i$  in (9) and note that the fixed points  $Q_*^1$  and  $Q_*^2$  satisfy  $Q_*^1(s, a) + Q_*^2(s, a) = 0$  for every  $(s, a)$ . Therefore, we also have  $\bar{Q}_k = \tilde{Q}_k^1 + \tilde{Q}_k^2$ . Based on (5b) and the definition of the fixed point  $Q_*^i$ , the evolution of  $\tilde{Q}_k^i$  can be written as

$$\tilde{Q}_{k+1}^i(s, a) = \tilde{Q}_k^i(s, a) + \bar{\beta}_k^i(s) \left( \gamma \sum_{s' \in S} p(s'|s, a) (v_k^i(s') - \text{val}^i(Q_*^i(s'))) - \tilde{Q}_k^i(s, a) \right), \quad (19)$$

where  $\bar{\beta}_k^i(s) := \mathbb{I}_{\{s=s_k\}} \beta_{c_k(s)}^i$  and  $r^i(s, a)$  disappears since  $Q_*^i(s, a) = (\mathcal{F}^i Q_*^i)(s, a)$ . We are interested in the limit inferior of  $\tilde{Q}_k^i$ , for each  $i$ , so that we can formulate the limit inferior of  $\bar{Q}_k$ , which will play an important role in (17).

The following lemma characterizes the limit inferior of an iterate whose evolution satisfies one-sided contraction-like condition. The proof is deferred to Appendix D of the extended version (Sayin and Cetiner, 2021).

**Lemma 9 (One-sided Asynchronous Discrete-time Convergence)** *Consider a sequence of vectors  $\{y_k\}_{k=0}^\infty$  such that the  $n$ th entry, denoted by  $y_k(n)$ , satisfies the following lower bound:*

$$y_{k+1}(n) \geq y_k(n) + \beta_k(n) \left( \gamma \min_m y_k(m) - y_k(n) + \epsilon_k(n) \right), \quad (20)$$

where  $\gamma \in (0, 1)$ , the vanishing (possibly random) step size  $\beta_k(n) \in [0, 1]$  satisfies  $\beta_k(n) \rightarrow 0$  and  $\sum_{k=0}^\infty \beta_k(n) = \infty$ , and  $\liminf_k \epsilon_k(n) \geq 0$  for each  $n$ , with probability 1. Suppose that  $\min_n y_k(n) \geq M$  for all  $k$ . Then, we have  $\liminf_k y_k(n) \geq 0$  for all  $n$ , with probability 1.

By the definition of the tracking error (9), we have

$$\sum_{s' \in S} p(s'|s, a) (v_k^i(s') - \text{val}^i(Q_*^i(s'))) = \sum_{s' \in S} p(s'|s, a) (e_k^i(s') + \text{val}^i(Q_k^i(s')) - \text{val}^i(Q_*^i(s')))$$

which is greater than  $\min_{(s,a)} \tilde{Q}_k^i(s, a)$  since  $e_k^i(s') \geq 0$  and  $\text{val}^i(Q_k^i(s')) - \text{val}^i(Q_*^i(s')) \geq \min_{a \in A} \tilde{Q}_k^i(s', a)$  for all  $s'$ . Under Assumptions 2 and 3, the step size  $\bar{\beta}_k(s) \in [0, 1]$  vanishes and  $\sum_{k=0}^\infty \beta_k(n) = \infty$  with probability 1. Therefore, we can invoke Lemma 9, for every  $i = 1, 2$ , and obtain

$$\liminf_{k \rightarrow \infty} \tilde{Q}_k^i(s, a) \geq 0 \quad \Rightarrow \quad \liminf_{k \rightarrow \infty} \bar{Q}_k(s, a) \geq 0, \quad (21)$$

almost surely for every  $(s, a)$ . Combined with (17), the bound (21) yields that

$$\underline{\epsilon}_k(s) \leq \bar{v}_k(s) \leq \frac{\lambda}{d_\alpha} \|\bar{Q}_k(s)\| + \bar{\epsilon}_k(s), \quad \forall s, \quad (22)$$



for some error terms  $\underline{\epsilon}_k(s) \rightarrow 0$  and  $\bar{\epsilon}_k(s) \rightarrow 0$  as  $k \rightarrow \infty$  for each  $s$  almost surely.

The heterogeneity of  $\beta_c^i$ 's has not played a role up to this point. Suppose that  $\lim_{c \rightarrow \infty} \beta_c^1 / \beta_c^2 = d_\beta \in (0, 1]$  without loss of generality. Then, the iteration (19) can be written as

$$\begin{bmatrix} \tilde{Q}_{k+1}^1 \\ \tilde{Q}_{k+1}^2 \end{bmatrix} = \begin{bmatrix} \tilde{Q}_k^1 \\ \tilde{Q}_k^2 \end{bmatrix} + \bar{\beta}_k(s) \begin{bmatrix} d_\beta \left( \gamma \sum_{s'} p(s'|\cdot) (v_k^1(s') - \text{val}^1(Q_*^1(s'))) - \tilde{Q}_k^1 + \zeta_k \right) \\ \gamma \sum_{s'} p(s'|\cdot) (v_k^2(s') - \text{val}^2(Q_*^2(s'))) - \tilde{Q}_k^2 \end{bmatrix}, \quad (23)$$

where we dropped the argument  $(s, a)$  for notational convenience and the error term is given by

$$\zeta_k(s, a) := \frac{1}{d_\beta} \left( \frac{\beta_{c_k(s)}^1}{\beta_{c_k(s)}^2} - d_\beta \right) \left( \gamma \sum_{s'} p(s'|s, a) (v_k^1(s') - \text{val}^1(Q_*^1(s'))) - \tilde{Q}_k^1(s, a) \right), \quad (24)$$

which is asymptotically negligible by Assumptions 2 and 3, and by the boundedness of the iterates. For notational convenience, we also define  $\Gamma_k(s, a) := \tilde{Q}_k^1(s, a) + d_\beta \tilde{Q}_k^2(s, a)$ . Then, the weighted combination of the entries in (23) yield that

$$\Gamma_{k+1}(s, a) = \Gamma_k(s, a) + \bar{\beta}_k(s) d_\beta \left( \gamma \sum_{s' \in S} p(s'|s, a) \bar{v}_k(s') - \bar{Q}_k(s, a) \right) \quad (25)$$

since  $\text{val}^1(Q_*^1(s)) + \text{val}^2(Q_*^2(s)) = 0$  for each  $s$ . By (21) and (25), we also have

$$\liminf_{k \rightarrow \infty} \Gamma_k(s, a) \geq 0, \quad \forall (s, a) \quad (26)$$

with probability 1. In order to characterize the convergence properties of  $\Gamma_k$  within the framework of Lemma 9, we introduce the following lemma formulating a bound on  $\bar{Q}_k$  in terms of  $\Gamma_k$ . The proof is deferred to Appendix E in the extended version (Sayin and Cetiner, 2021).

**Lemma 10** *For some  $\underline{\eta}_k(s, a) \rightarrow 0$  and  $\bar{\eta}_k(s, a) \rightarrow 0$  with probability 1, we have*

$$\bar{\eta}_k(s, a) + \frac{1}{d_\beta} \Gamma_k(s, a) \geq |\bar{Q}_k(s, a)| \geq \bar{Q}_k(s, a) \geq \Gamma_k(s, a) + \underline{\eta}_k(s, a), \quad \forall (s, a). \quad (27)$$

Based on (22), (25), and Lemma 10, we obtain

$$\Gamma_{k+1}(s, a) \leq \Gamma_k(s, a) + \bar{\beta}_k(s) d_\beta \left( \frac{\gamma \lambda}{d_\alpha d_\beta} \max_{(s', a')} \Gamma_k(s', a') - \Gamma_k(s, a) + \varepsilon_k(s, a) \right), \quad (28)$$

where  $\varepsilon_k(s, a)$  is an asymptotically negligible error almost surely. We can invoke Lemma 9 for  $\{-\Gamma_k\}_{k \geq 0}$  and obtain that

$$\liminf_{k \rightarrow \infty} -\Gamma_k(s, a) = -\limsup_{k \rightarrow \infty} \Gamma_k(s, a) \geq 0 \quad \Rightarrow \quad \limsup_{k \rightarrow \infty} \Gamma_k(s, a) \leq 0. \quad (29)$$

Combined with (26), the bound (29) yields that  $\lim_{k \rightarrow \infty} \Gamma_k(s, a) = 0$ . Hence,  $\lim_{k \rightarrow \infty} \tilde{Q}_k^i(s, a) = 0$ , for all  $(s, a)$  and  $i = 1, 2$  with probability 1. Since  $\bar{Q}_k = \tilde{Q}_k^1 + \tilde{Q}_k^2$ , we also have  $\bar{Q}_k(s, a) \rightarrow 0$  for each  $(s, a)$ . Then, we obtain  $\bar{v}_k(s) \rightarrow 0$  for each  $s$  by (22). Therefore, Lemma 8 yields that the tracking error  $e_k^i(s) \rightarrow 0$  for each  $s$ , with probability 1.




Figure 1: Evolution of  $v_k^1$ ,  $v_k^2$ , and  $\bar{v}_k = v_k^1 + v_k^2$  for each  $s$ , respectively, in **red**, **blue**, and **green**. The dotted lines  $\dots/\dots$  are the actual equilibrium values associated with each state.

## 5. Numerical Example

In this section, we examine the convergence properties of the heterogenous learning dynamics numerically in an (irreducible) zero-sum stochastic game whose configuration is selected arbitrarily as in (Sayin et al., 2020). For example, there are three states, four actions per state, and the discount factor is 0.8. The player-dependent step sizes are set as  $\alpha_c^1 = (1 + c)^{0.5}$  and  $\alpha_c^2 = (1 + 0.81 \cdot c)^{0.5}$  while  $\beta_c^1 = (1 + c)^{-1}$  and  $\beta_c^2 = (1 + 0.95 \cdot c)^{-1}$ . In Figure 1, we plot the evolution of the value function estimates of both player in addition to  $\bar{v}_k$  to illustrate that the auxiliary stage-games become zero-sum also in the heterogenous case. Since the underlying game is irreducible, Assumption 2 holds. For these step sizes, we have  $d_\alpha = 0.9$  and  $d_\beta = 0.95$ . Assumption 3 also holds. Therefore, Theorem 4 says that the dynamics should converge to an equilibrium of the game and we have observed the convergence of the value function estimates to the equilibrium values of the game, as expected from Theorem 4.

## 6. Conclusion

We showed the almost sure convergence of two-timescale fictitious play with heterogeneous learning rates in two-player zero-sum stochastic games under the standard assumptions in two-timescale stochastic approximation methods when the discount factor is less than the product of the ratios of the player-dependent step-sizes. Since strategic-form games played repeatedly is a special case of stochastic games, this result also implied the almost sure convergence of fictitious play with heterogeneous rates in zero-sum strategic-form games. We attributed the bound on the discount factor to the deviation of the auxiliary stage-games from the zero-sum structure, which becomes multifold with the player-dependent rates while sufficiently small discount rates can compensate it. Future research directions include characterizing the convergence properties of heterogenous and independent learning dynamics in stochastic games other than zero-sum, and in model-free and minimal information cases, e.g., as in (Sayin et al., 2021).

## References

- G. Arslan and S. Yüksel. Decentralized Q-learning for stochastic teams and games. *IEEE Transactions on Automatic Control*, 62(4):1545–1558, 2017.
- M. Benaïm, J. Hofbauer, and S. Sorin. Stochastic approximations and differential inclusions. *SIAM J. Control Optim.*, 44(1):328–348, 2005.
- V. S. Borkar. *Stochastic Approximation: A Dynamical Systems Viewpoint*. Hindustan Book Agency, 2008.
- B. Chasnov, L. Ratliff, E. Mazumdar, and S. Burden. Convergence analysis of gradient-based learning in continuous games. In *Proceedings of The 35th Uncertainty in Artificial Intelligence Conference*, volume 115, pages 935–944, 2020.
- C. Daskalakis, D. J. Foster, and N. Golowich. Independent policy gradient methods for competitive reinforcement learning. In *Advances in Neural Information Processing Systems*, 2020.
- J. Filar and K. Vrieze. *Competitive Markov Decision Processes*. Springer Verlag, 1997.
- A. M. Fink. Equilibrium in stochastic n-person game. *Journal of Science Hiroshima University Series A-I*, 28:89–93, 1964.
- D. Fudenberg and D. Kreps. Learning in extensive-form games I. Self-confirming equilibria. *Games and Economic Behavior*, 8:20–55, 1995.
- D. Fudenberg and D. K. Levine. *The Theory of Learning in Games*. MIT Press, Cambridge, MA, 1998.
- D. Fudenberg and D. K. Levine. Learning and equilibrium. *The Annual Review of Economics*, 1: 385–419, 2009.
- D. S. Leslie and E. J. Collins. Convergent multi-timescales reinforcement learning algorithms in normal form games. *The Annals of Applied Probability*, 13(4):1231–1251, 2003.
- D. S. Leslie and E. J. Collins. Individual Q-learning in normal form games. *SIAM J. Control Optim.*, 44(2):495–514, 2005.
- D. S. Leslie, S. Perkins, and Z. Xu. Best-response dynamics in zero-sum stochastic games. *Journal of Economic Theory*, 189, 2020.
- M. L. Littman. Markov games as a framework for multi-agent reinforcement learning. In *Proceedings of the 11th International Conference on Machine Learning (ICML)*, 1994.
- A. Ozdaglar, M. O. Sayin, and K. Zhang. Independent learning in stochastic games. In *International Congress of Mathematicians*, 2022.
- M. O. Sayin and K. A. Cetiner. On the heterogeneity of independent learning dynamics in zero-sum stochastic games. *Extended version available at <https://arxiv.org/abs/2112.06181>*, 2021.
- M. O. Sayin, F. Parise, and A. Ozdaglar. Fictitious play in zero-sum stochastic games. *ArXiv:2010.04223*, 2020.

- M. O. Sayin, K. Zhang, D. Leslie, A. Ozdaglar, and T. Başar. Decentralized Q-learning in zero-sum markov games. In *Advances in Neural Information Processing Systems*, 2021.
- L. S. Shapley. Stochastic games. *Proceedings of National Academy of Science USA*, 39(10):1095–1100, 1953.
- R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, MA, 2018.
- C. Szepesvari and M. Littman. A unified analysis of value-function-based reinforcement-learning algorithms. *Neural Computing*, 11:2017–2060, 1999.
- J. N. Tsitsiklis. Asynchronous stochastic approximation and Q-learning. *Machine Learning*, 16:185–202, 1994.
- C. J. C. H. Watkins and P. Dayan. Q-learning. *Machine Learning*, 8(3):279–292, 1992.
- C.-Y. Wei, Y.-T. Hong, and C.-J. Lu. Online reinforcement learning in stochastic games. In *Proceedings of the 30th Conference on Neural Information Processing Systems (NIPS)*, 2017.
- H. P. Young. *Strategic Learning and Its Limits*. Oxford University Press, 2004.
- K. Zhang, Z. Yang, and T. Başar. Multi-agent reinforcement learning: A selective overview of theories and algorithms. In *Handb. Rein. Learn. Cont.*, volume 325. Springer, 2021.
- Q. Zhu, H. Tembine, and T. Başar. Heterogeneous learning in zero-sum stochastic games with incomplete information. In *IEEE Conf. Decision and Control*, pages 219–224, 2011.