

# Episodic Logit-Q Dynamics for Efficient Learning in Stochastic Teams

Onur Unlu and Muhammed O. Sayin

**Abstract**—We present new learning dynamics combining (independent) log-linear learning and value iteration for stochastic games within the auxiliary stage game framework. The dynamics presented provably attain the efficient equilibrium (also known as optimal equilibrium) in identical-interest stochastic games, beyond the recent concentration of progress on provable convergence to some (possibly inefficient) equilibrium. The dynamics are also independent in the sense that agents take actions consistent with their local viewpoint to a reasonable extent rather than seeking equilibrium. These aspects can be of practical interest in the control applications of intelligent and autonomous systems. The key challenges are the convergence to an inefficient equilibrium and the non-stationarity of the environment from a single agent’s viewpoint due to the adaptation of others. The log-linear update plays an important role in addressing the former. We address the latter through the play-in-episodes scheme in which the agents update their Q-function estimates only at the end of the episodes.

## I. INTRODUCTION

Shapley introduced stochastic games (SGs) as a generalization of Markov decision processes (MDPs) to non-cooperative multi-agent settings [1]. Since MDPs are the basis of reinforcement learning, SGs have also attracted significant interest as an ideal model for multi-agent reinforcement learning. Examples include planning for intelligent and autonomous systems [2]. Shapley also showed that Markov stationary equilibrium exists in discounted SGs with finite state and action spaces similar to the existence of a stationary solution in MDPs. Apart from the studies focusing on the computation of the stationary equilibrium in SGs, recently, there has also been a growing interest in determining whether non-equilibrium adaptation of learning agents reach stationary equilibrium in SGs, e.g., [3], [4], [5], [6], similar to the extensive literature on learning in the repeated play of games, e.g., see [7]. However, in identical-interest SGs, also known as stochastic teams, reaching possibly inefficient equilibrium may not be desirable in control and optimization applications. For example, efficient learning in which agents reach efficient equilibrium has been studied extensively for games with repeated play, e.g., [8], [9], [10]. However, very limited results exist on efficient learning in stochastic games, as reviewed in detail later.

In the *auxiliary stage game* framework, we can view SGs as agents are playing a stage game associated with the state visited. The payoffs of these stage games, called *Q-functions*,

depend not only on the immediate reward they receive but also on the rewards they will receive in future stages, often referred to as the continuation payoff. However, the continuation payoff depends on the evolving strategies of the agents. Hence, the stage games are not necessarily stationary. The non-stationarity of stage games is acknowledged as one of the main challenges for learning in SGs [11]. There is no straightforward generalization of the existing results from repeated play setting that address efficient learning, e.g., [8], [9], [10] into SGs. Recently, the *two-timescale learning framework* addressed the non-stationarity issue for the best response and fictitious play dynamics [3], [4]. However, these learning dynamics do not necessarily reach efficient equilibrium even in identical-interest SGs.

In this paper, we present the new *logit-Q* learning dynamics, combining log-linear learning and value iteration for efficient learning in SGs. We also present the *independent logit-Q* learning dynamics to address the coordination burden in the synchronous (in turn) update of actions inherent to the log-linear learning. We show the almost sure convergence of the Q-function estimates to the *globally optimum* Q-function of the underlying identical-interest SG in both logit-Q and independent logit-Q dynamics. We also verify the convergence of the learning dynamics via numerical examples.

Independent and simultaneous adaptation of agents poses the key challenges of: *i*) achieving the *global* optimum rather than converging to an inefficient equilibrium and *ii*) the *non-stationarity* of the environment from each agent’s perspective. The former one is addressed in the log-linear learning for the repeated games yet the convergence results, such as [8], are in the sense that only the optimal solution is *stochastically stable*. On the other hand, the latter one is viewed as a core issue in multi-agent reinforcement learning [2] and in learning in SGs, as discussed above.

The key property of our dynamics is to let agents *play-in-episodes*. Within an episode, they play stage games whenever the associated state gets visited without any update on the Q-function estimates, as illustrated later in Figure 1. Therefore, their act is always consistent with the stage game they play. We consider two update schemes for the Q-function update: In the *average value update*, agents take the empirical average of the payoffs received in the repeated play of the stage games within the episode. In the *frequent value update*, agents assign the value of a state as the payoff of the corresponding stage game associated with the action profile most frequently played. Both are consistent with the (desired) uncoupled nature of the dynamics. The former ensures convergence to the optimal solution approximately with an approximation error characterized. The latter ensures convergence to the optimal

This work was supported by the TUBITAK BIDEB 2232-B International Fellowship for Early Stage Researchers under Grant Number 121C124.

Onur Unlu and Muhammed O. Sayin are with the Department of Electrical & Electronics Engineering, Bilkent University, Ankara, Turkey 06800. Emails: onur.unlu@ug.bilkent.edu.tr, sayin@ee.bilkent.edu.tr

solution exactly. Furthermore, these update schemes make the stage games stationary within an episode and allows agents to converge to efficient equilibria at each stage game through the (independent) log-linear learning dynamics to a certain extent. Therefore, the update of continuation payoff approximates the value iteration in MDPs. Increasing episode lengths across episodes combined with the contraction property of the value iteration yield that agents can attain the globally optimum solution in the limit. The increasing episode lengths play a similar role with the step sizes vanishing at two different timescales in the two timescale learning framework used in [3] and [4].

There are a few studies addressing efficient learning in SGs, such as [12], [13], and [14]. In [12], the authors presented a learning dynamic for identical-interest SGs based on the adaptive play. Quite contrary to our uncoupled learning dynamics, theirs has a coupled structure due to the off-policy learning of game structure including a maximization over joint actions of all agents. Notably, in [13], the authors focused on decentralized learning of efficient equilibrium in SGs. Their key idea is to focus on stationary pure strategies so that the underlying SG can be viewed as a (huge) normal-form game in which actions correspond to stationary pure strategies (since there are only finitely many contrary to a continuum of stationary mixed strategies). This reduces the problem to the repeated play of this (huge) normal-form game. They addressed the unknown state transitions and no access to opponent actions through learning-in-phases in which there is no update of strategies. In that sense, their dynamics have a flavor more similar with the actor-critic methods where two different timescales are used in the reverse order with our approach. In other words, the approaches differ in terms of learning-in-phases vs playing-in-episodes. Lastly, in [14], we address efficient learning in stochastic teams with a vanishing step size used in the value function update different from the episodic scheme presented. However, the dynamics there provably reach a neighborhood of the efficient equilibrium with a more involved convergence analysis. Recall that logit-Q learning dynamics with the frequent value update reaches the exact efficient equilibrium.

The paper is organized as follows. We describe identical-interest SGs in Section II and the new classical and independent logit-Q dynamics in Section III. Then, we present the convergence results in Section IV, and an overview of their proofs in Section V. In Section VI, we show the performance of the algorithm with numerical examples. We conclude the paper with some remarks in Section VII. An appendix includes the proof of a technical lemma.

## II. IDENTICAL INTEREST STOCHASTIC GAMES

Formally, an  $n$ -agent SG is a dynamic game played over infinite horizon that can be characterized by the tuple  $\langle S, A, r, p \rangle$ . At each stage, the game visits a state  $s$  from a *finite* set  $S$  and each agent  $i$  simultaneously takes an action  $a^i$  from a *finite* set  $A^i$  to collect *stage-payoffs*.<sup>1</sup> Similar to

<sup>1</sup>The formulation can be extended to state-variant action sets rather straightforwardly.

MDPs, stage-payoffs and transition of the game across states depend only on the current state visited and current action profile  $a = \{a^i\}_{i=1}^n$  played. Particularly,  $p(s'|s, a)$  for each  $(s, a, s')$  denotes the probability of transition from  $s$  to  $s'$  under action profile  $a$ , and  $r^i : S \times A \rightarrow \mathbb{R}$  denotes the stage-payoff function of agent  $i$ . We specifically consider the identical-interest SGs, also known as stochastic teams, where there exists  $r : S \times A \rightarrow \mathbb{R}$  such that  $r^i(s, a) = r(s, a)$  for all  $(i, s, a)$ .

We let agents randomize their actions according to a strategy determining the probabilities of actions to be played. They choose their strategies to maximize the expected sum of discounted stage-payoffs with the discount factor  $\gamma \in [0, 1)$ . For example, the objective of player  $i$  is given by

$$\mathbb{E} \left[ \sum_{k=0}^{\infty} \gamma^k r(s_k, a_k) \right], \quad (1)$$

where  $(s_k, a_k)$  denotes the pair of state and action profile at stage  $k$  and the expectation is taken with respect to the randomness on  $(s_k, a_k)$  for each  $k \geq 0$ .

Notice that, in SGs, agents play a normal-form game associated with the state visited at each stage where the payoffs of these stage games, called *Q-functions*, are of the form

$$Q(s, a) := \mathbb{E} \left[ \sum_{k=1}^{\infty} \gamma^k r(s_k, a_k) \mid s_0 = s, a_0 = a \right], \quad (2)$$

and the corresponding value of state, called value function is defined as

$$v(s) := \mathbb{E} \left[ \sum_{k=1}^{\infty} \gamma^k r(s_k, a_k) \mid s_0 = s \right], \quad (3)$$

In an SG, we say that a strategy is *stationary* if it depends only on the current state (and not on the time). Similar to the existence of a stationary optimal solution in MDPs, there also exists a *stationary equilibrium* in SGs such that agents do not have any incentive to change their stationary strategies unilaterally [15]. Furthermore, specifically in identical-interest SGs, there also exists a stationary equilibrium, called *efficient stationary equilibrium*, attaining the global maximum of the common objective (1). The (unique) value function and the (unique) Q-function associated with an efficient stationary equilibrium, respectively, denoted by  $v_* : S \rightarrow \mathbb{R}$  and  $Q_* : S \times A \rightarrow \mathbb{R}$ , satisfy the following fixed-point equations for all  $(s, a)$ , and  $s$  :

$$v_*(s) = \max_{a \in A} \left\{ r(s, a) + \gamma \sum_{s' \in S} p(s'|s, a) v_*(s') \right\}, \quad (4)$$

$$Q_*(s, a) = r(s, a) + \gamma \sum_{s' \in S} p(s'|s, a) \max_{a' \in A} \{ Q_*(s', a') \},$$

and  $v_*(s) = \max_{a \in A} \{ Q_*(s, a) \}$  for all  $s$ . The uniqueness follows from the contraction property of the right-hand side. Indeed the right-hand side in (4) is the Bellman operator we would have if it was a (single-agent) MDP.

## III. EPISODIC (INDEPENDENT) LOGIT-Q LEARNING

An important question is *how an agent would/should play*. The rich literature on learning in games has studied this question for the repeated play of the same normal-form game.




Fig. 1. A figurative explanation for the logit-Q learning dynamics in an SG with four states. The circles shaded represent the state visited at the corresponding stage and the agents play the associated stage game only at those times. The arrows represent the flow of state transitions in time. The same color is used within an episode to highlight that the same normal-form games are getting played repeatedly.

As previously discussed, agents also play a normal-form stage game associated with the state visited at each stage in SGs. However, the payoffs of these stage games,  $Q$ -functions, depend both on the stage-payoff received immediately, e.g.,  $\gamma^k r(s_k, \cdot)$ , and the expected continuation-payoff to be received in future stages, e.g.,  $\gamma^k \mathbb{E} [\sum_{l=k+1}^{\infty} \gamma^{l-k} r(s_l, a_l)]$  since  $s_{k+1} \sim p(\cdot | s_k, \cdot)$  as can be seen in (2). The latter is ambiguous and possibly non-stationary by depending on the future play. Hence, in SGs, the agents are not necessarily playing the same stage game repeatedly whenever the associated state gets visited.

To address the ambiguity of the future play, we use the standard approach to let agents estimate the value of a state, called *value function*, based on the history of their interactions as if the agents would play in the future as they played in the past. To mitigate the non-stationarity issue, we propose to let agents update their value function estimates (and therefore, the payoff of the stage games) *only* at certain stages. We call the time between two consecutive updates as an *episode*. Correspondingly, the agents play stationary stage games, i.e., follow the (independent) log-linear learning dynamics, described precisely later, within each episode, as illustrated in Figure 1. We re-emphasize that this play-in-episodes scheme has a similar flavor with the two-timescale learning framework studied in [3], [4] in addressing the non-stationarity issue. Furthermore, it is consistent with the differences in the evolution pace of choices and preferences as studied in the evolutionary game theory literature.

We denote the value function and  $Q$ -function estimates for episode  $t = 1, 2, \dots$  by  $v_{(t)} : S \rightarrow \mathbb{R}$  and  $Q_{(t)} : S \times A \rightarrow \mathbb{R}$ , respectively. Based on  $v_{(t)}(\cdot)$ , the agents can construct the  $Q$ -function estimate (and therefore, the payoff of the stage games) for the next episode ( $t + 1$ ) according to

$$Q_{(t+1)}(s, a) = r(s, a) + \gamma \sum_{s' \in S} p(s' | s, a) v_{(t)}(s'), \quad (5)$$

with  $Q_{(1)}(s, a) = r(s, a)$ , for all  $(s, a)$ . For the update of  $v_{(t)}$ , we consider two schemes. Firstly, the agents can take the empirical average of payoffs received in the associated stage games, i.e.,

$$v_{(t)}^{\text{ave}}(s) = \sum_{a \in A} \eta_{(t)}(s, a) Q_{(t)}(s, a), \quad \forall s, \quad (6)$$

where  $\eta_{(t)}(s, a) := c_{(t)}(s, a)/c_{(t)}(s)$  is the sampled frequency of  $(s, a)$ , corresponding to the number of times that  $a$  gets played in state  $s$ ,  $c_{(t)}(s, a)$ , divided by the number of times that  $s$  gets visited,  $c_{(t)}(s)$ , within episode  $t$ . Alternatively, they can estimate it as the payoff associated with the action profile played the most frequently within the latest episode, i.e.,

$$v_{(t)}^{\text{freq}}(s) = \frac{1}{|A_{(t)}^{\text{freq}}(s)|} \sum_{a \in A_{(t)}^{\text{freq}}(s)} Q_{(t)}(s, a), \quad \forall s, \quad (7)$$

where  $A_{(t)}^{\text{freq}}(s) := \text{argmax}_{a \in A} \{\eta_{(t)}(s, a)\}$ .<sup>2</sup> Both schemes (6) and (7) result in a value function estimate common among agents, and therefore, the stage games always have the identical interest structure throughout the underlying SG. Furthermore, both are also consistent with the (desired) uncoupled nature of the dynamics.

In the first and the most standard version of the log-linear learning we consider, which was originally introduced by [16], agents respond to the latest action of others in turn. At the first play of each stage game specific to a state, the agents take arbitrary actions. Then, they keep track of the latest action played by each agent  $i$  at each state  $s$  until and excluding stage  $k$ , denoted by  $\alpha_k^i(s)$ . Every agent except a randomly chosen one according to the uniform distribution over agents, say  $i$ , takes the latest action they have taken in the associated state. Then, agent  $i$  takes action  $a^i$  drawn according to the distribution  $\sigma^i(Q_{(t)}^i(s_k, \cdot), \alpha_k^{-i}(s_k))$ , where  $\sigma^i : \mathbb{R}^{|A^i|} \rightarrow \Delta(A^i)$  is the soft-max function given by

$$\sigma^i(z)[a] = \frac{e^{z[a]/\tau}}{\sum_{a'} e^{z[a']/\tau}} \quad \forall z \in \mathbb{R}^{|A^i|}, \quad (8)$$

where  $\tau > 0$  is a temperature parameter ensuring that each action gets played with some positive probability.

Though the *update-in-turn* plays an important role in the convergence of the log-linear learning dynamics, certain structured relaxations minimizing the coordination among agents are also possible. One of the most critical structural relaxations is to let agents independently decide to change their actions. Therefore, in the second version of the log-linear learning we consider, agents independently respond to

<sup>2</sup>For a finite set  $A$ , we denote its number of elements by  $|A|$ .

---

**Algorithm 1** Episodic Logit-Q Dynamics

---

**Require:** each agent keeps track of  $\{\alpha_k^j(s)\}_{(j,s) \in [n] \times S}$

- 1: **for** each episode  $t = 1, 2, \dots$  **do**
- 2:     **input:**  $Q_{(t)}(s, a)$ , described in (5) using  $v^{\text{ave}}$  or  $v^{\text{freq}}$
- 3:     **for** each stage  $l = \kappa_{(t)}, \dots, \kappa_{(t+1)} - 1$  **do**
- 4:         **input:** Current state  $s_l$
- 5:         each agent  $i$  **plays**  
$$\begin{cases} a^i(s_l) \sim \sigma^i(Q_{(t)}(s_l, \cdot, \alpha_l^{-i}(s_l))) & \text{if he/she is picked} \\ a^i(s_l) = \alpha_l^i(s_l) & \text{otherwise} \end{cases}$$
- 6:         **output:** Count of realizations:  $c_{(t)}(s), c_{(t)}(s, a)$
- 7:     **end for**
- 8:     **compute:**  $\eta_{(t)}(s, a) = c_{(t)}(s, a)/c_{(t)}(s) \quad \forall(s, a)$ ,
- 9:     **output:**  
$$\begin{cases} v_{(t)}^{\text{ave}}(s) = \sum_a \eta_{(t)}(s, a) Q_{(t)}^{\text{ave}}(s, a) \\ \text{or} \\ v_{(t)}^{\text{freq}}(s) = \frac{1}{|A_{(t)}(s)|} \sum_{a \in A_{(t)}(s)} Q_{(t)}^{\text{freq}}(s, a) \end{cases}$$
- 10: **end for**

---

the latest action of others, see [8, Section 4] for an extended introduction. As before, at the first play of each stage game specific to a state, the agents take arbitrary actions. Then starting from the first play, they again keep track of the latest action played by each agent  $i$  at each state  $s$  until and excluding stage  $k$ , denoted by  $\alpha_k^i(s)$ . Then, at each play of the stage game, in contrast to the classical log-linear learning, every agent *independently* decides on whether to change their actions with some probability  $\omega$ . An agent, say  $i$ , that decides to change their action, take action  $a^i$  drawn according to the distribution  $\sigma^i(Q_{(t)}^i(s_k, \cdot, \alpha_k^{-i}(s_k)))$ . We refer to this dynamics as *independent* logit-Q dynamics.

We emphasize that our approach differs from pure algorithmic solutions used in decentralized/distributed control or dynamic/mathematical programming applications as it is based on a classical behavioral model. Furthermore, we emphasize that both Algorithms 1 and 2 have an uncoupled nature as agents do not operate on the joint action profile contrary to the coupled learning algorithms that tackle the Markov team problems that involve maximization over joint action profiles, e.g., Team Q-learning introduced in [17].

Algorithm 1 and Algorithm 2 provide a description of the classical and independent logit-Q dynamics respectively. There, the agents know the stage-payoff function and transition probabilities but the extension into the case where they do not know but estimate the stage-payoffs and transition probabilities can be found in the preliminary version of this paper [18].

#### IV. MAIN RESULTS

In this section, we characterize the convergence properties of Algorithms 1 and 2. The following theorem characterizes the convergence of Q-function estimates in irreducible identical-interest SGs. We say that an SG is irreducible if

---

**Algorithm 2** Episodic Independent Logit-Q Dynamics

---

**Require:** each agent keeps track of  $\{\alpha_k^j(s)\}_{(j,s) \in [n] \times S}$

- 1: **for** each episode  $t = 1, 2, \dots$  **do**
- 2:     **input:**  $Q_{(t)}(s, a)$ , described in (5) using  $v^{\text{ave}}$  or  $v^{\text{freq}}$
- 3:     **for** each stage  $l = \kappa_{(t)}, \dots, \kappa_{(t+1)} - 1$  **do**
- 4:         **input:** Current state  $s_l$
- 5:         each agent  $i$  **independently plays**  
$$\begin{cases} a^i(s_l) \sim \sigma^i(Q_{(t)}(s_l, \cdot, \alpha_l^{-i}(s_l))) & \text{w. p. } \omega \\ a^i(s_l) = \alpha_l^i(s_l) & \text{w. p. } (1 - \omega) \end{cases}$$
- 6:         **output:** Count of realizations:  $c_{(t)}(s), c_{(t)}(s, a)$
- 7:     **end for**
- 8:     **compute:**  $\eta_{(t)}(s, a) = c_{(t)}(s, a)/c_{(t)}(s) \quad \forall(s, a)$ ,
- 9:     **output:**  
$$\begin{cases} v_{(t)}^{\text{ave}}(s) = \sum_a \eta_{(t)}(s, a) Q_{(t)}^{\text{ave}}(s, a) \\ \text{or} \\ v_{(t)}^{\text{freq}}(s) = \frac{1}{|A_{(t)}(s)|} \sum_{a \in A_{(t)}(s)} Q_{(t)}^{\text{freq}}(s, a) \end{cases}$$
- 10: **end for**

---

$p(s' | s, a) > 0$  for all  $(s, a, s')$  as in [3].<sup>3</sup>

**Theorem 1.** *Given an irreducible identical-interest SG, consider that agents follow Algorithm 1. Suppose that the episode lengths go to infinity, i.e.,  $L_{(t)} := \kappa_{(t+1)} - \kappa_{(t)} \rightarrow \infty$  as  $t \rightarrow \infty$ , e.g.,  $L_{(t)} \propto t^2$ .*

(i) *If the agents follow the average value update scheme, then we have*

$$\limsup_{t \rightarrow \infty} |Q_{(t)}(s, a) - Q_*(s, a)| \leq \tau \frac{\log |A|}{1 - \gamma} \quad \forall(s, a)$$

*with probability 1.*

(ii) *If the agents follow the frequent value update scheme, then we have*

$$\limsup_{t \rightarrow \infty} |Q_{(t)}(s, a) - Q_*(s, a)| = 0, \quad \forall(s, a)$$

*with probability 1.*

Correspondingly, the following theorem is the counterpart of Theorem 1 for Algorithm 2.

**Theorem 2.** *Given an irreducible identical-interest SG, consider that agents follow Algorithm 2. Suppose that the episode lengths go to infinity, i.e.,  $L_{(t)} := \kappa_{(t+1)} - \kappa_{(t)} \rightarrow \infty$  as  $t \rightarrow \infty$ , e.g.,  $L_{(t)} \propto t^2$ .*

(i) *If the agents follow the average value update scheme, then we have*

$$\limsup_{t \rightarrow \infty} |Q_{(t)}(s, a) - Q_*(s, a)| \leq \frac{\tau \log |A| + H(\omega)}{1 - \gamma}$$

*$\forall(s, a)$ , with probability 1, where  $H(\omega) \in \mathcal{O}(\omega^n)$  is induced by the differences between the stationary distributions of classical and independent log-linear updates. Note that  $H(\omega)$  vanishes as  $\omega \rightarrow 0$ .*

<sup>3</sup>Irreducibility assumption can be relaxed by focusing on the recurrent class that the MG reaches eventually, see the preliminary version of this paper [18] for detailed discussion.

(ii) If the agents follow the frequent value update scheme, then we have

$$\limsup_{t \rightarrow \infty} |Q_{(t)}(s, a) - Q_*(s, a)| \leq \frac{\tau M(\omega)}{1 - \gamma},$$

$\forall (s, a)$ , with probability 1, where  $M(\omega) \in \mathcal{O}(n \log(\omega))$  is induced by the differences in the maximizers of the stationary distributions of classical and independent log-linear updates. Note that  $M(\omega)$  vanishes as  $\omega \rightarrow 0$ .

Both Theorems 1 and 2 show the almost sure convergence of the Q-function estimates to a small neighborhood around  $Q_*$  in both average and frequent value update schemes.

## V. PROOFS OF THEOREMS 1 & 2

We defer the reader to the extended version of this paper [19] for the detailed proof due to the space constraints and provide an overview of the proof in this section. The proofs are built upon the observation that if value function estimates  $v_{(t)}(s)$  for each  $s$  track the maximum of the Q-function estimates over action profiles, i.e.,  $\max_{a \in A} Q_{(t)}(s, a)$ , then the evolution of  $v_{(t)}$  across episodes  $t = 1, 2, \dots$  approximates the value iteration for MDPs (with single agent). Hence, the convergence to the global optimum follows. Therefore, we define a tracking error term,  $e_{(t)}(s) := v_{(t)}(s) - \max_{a \in A} \{Q_{(t)}(s, a)\}$ , on the deviation between the value function estimates and the maximum of the Q-function estimates. Thus, our goal is to characterize  $c \geq \limsup_{t \rightarrow \infty} \max_s |e_{(t)}(s)|$  for Algorithm 1 and Algorithm 2 with both average and frequent value update schemes. To do so, we use the fact that the action profiles played at every visit to a specific state forms an irreducible and aperiodic Markov chain. Particularly, the ergodic theorem for Markov chains, e.g., see [20, Theorem C.1], yields that sampled frequency of action profiles played at the specific state  $s$  converges to the stationary distribution of the Markov chain formed by them. This allows us to characterize the error term for Algorithm 1 since the stationary distribution of the Markov chain induced by the classical log-linear learning is well-known, e.g., see [8]. However, stationary distribution of the Markov chain induced by the independent log-linear learning is unknown and very challenging to characterize, see [8] for a detailed discussion. Therefore, we rather focus on bounding the distance between the stationary distributions of the Markov chains induced by classical and independent log-linear learning. We characterize a bound by coupling the chains formed by classical and independent log-linear learning using our novel result on the analysis of perturbed Markov chains with perturbations. Since this result can be of independent interest we provide it in the following lemma. Its proof follows from the proof of Lemma 2 in [14] with certain modifications. We still provide the entire proof in [19].

**Lemma 3.** Consider two chains  $\{a_k\}$  and  $\{\hat{a}_k\}$  over the same state space  $A$ . Both form homogeneous, irreducible, and aperiodic Markov chains with transition kernels  $P$  and  $\hat{P}$ , respectively. Suppose that:

(i) *Small Perturbation:* There exists some  $\lambda \in (0, 1]$  such that

$$\left\| \hat{P}(\cdot | \hat{a}) - P(\cdot | \hat{a}) \right\|_{\text{TV}} \leq 1 - \lambda \quad \forall \hat{a} \in A.$$

(ii) *Matching Probability:* There exists some  $\varepsilon > 0$  such that

$$\Pr(a_k = \hat{a}_k \mid a_{k-1} \neq \hat{a}_{k-1}) \geq \varepsilon,$$

Then, there exists a coupling over the sequences  $\{\hat{a}_k\}$  and  $\{a_k\}$  such that we have

$$\Pr(\hat{a}_k \neq a_k) \leq (1 - \varepsilon)^k + \frac{1 - \lambda}{\varepsilon} \quad \forall k.$$

## VI. NUMERICAL EXAMPLES

Consider an identical-interest SG with five states and five agents, and each agent has five actions at each state. We choose the stage-payoff function and transition probabilities randomly. To this end, we first draw  $r(s, a) \propto \bar{r} \cdot \mathcal{O}(s^2)$  for each  $s$ , where  $\bar{r}$  is chosen uniformly from  $[0, 1]$ , and then normalize them by  $\max_{(s, a)} \{r(s, a)\}$  to obtain rewards  $|r(s, a)| \leq 1$  for all  $(s, a)$ . On the other hand, we first sample  $p(s'|s, a)$  uniformly from  $[0.2, 1]$  and then normalize them to have  $\sum_{s'} p(s'|s, a) = 1$ . We set the discount factor  $\gamma$ ,  $\tau$  and  $\omega$  as 0.6,  $10^{-3}$  and  $10^{-2}$ , respectively.

We examine the evolution of value function estimates for both update scheme  $v^{\text{ave}}$ , as described in (6), and  $v^{\text{freq}}$ , as described in (7), under Algorithm 1 and Algorithm 2. We use the suffix **-ave** if agents use the average value update scheme, and **-freq** if agents use the frequent value update scheme. We set the episode lengths proportional to episode index squared, i.e.,  $L_{(t)} \propto t^2$  and note that a run for both **Algorithm 1** and **Algorithm 2** approximately takes a minute on a laptop computer equipped with a 2.6 GHz Intel i7-9750H processor, 16 GB RAM and NVIDIA GeForce RTX 2070 graphic card using Python 3.7.10 in Jupyter Notebook.

In Figure 2, we plot the evolution of the value function estimates as a result of 20 independent runs. As we have shown in Theorem 1 and Theorem 2 the value function estimates converge to the optimal ones in all cases in Figure 2. Observe that  $v^{\text{ave}}$  converges to a small neighborhood around the solution, as characterized in Theorem 1, while  $v^{\text{freq}}$  converges to the exact solution. We deliberately plot the small neighborhood from Theorem 1 also for Algorithm 2 to show that for small values of  $\omega$  Algorithm 2 has almost identical performance with Algorithm 1.

## VII. CONCLUSIONS

We have presented the new logit-Q and independent logit-Q learning dynamics provably attaining the social optimum in identical-interest SGs. Unlike many recent results on learning in SGs attaining possibly inefficient equilibrium, the dynamics presented converge to an efficient equilibrium almost surely, by addressing the challenge of equilibrium selection. Unlike many results on efficient learning in repeated games, our learning dynamics have convergence guarantees for SGs, by addressing the non-stationarity of the stage games induced

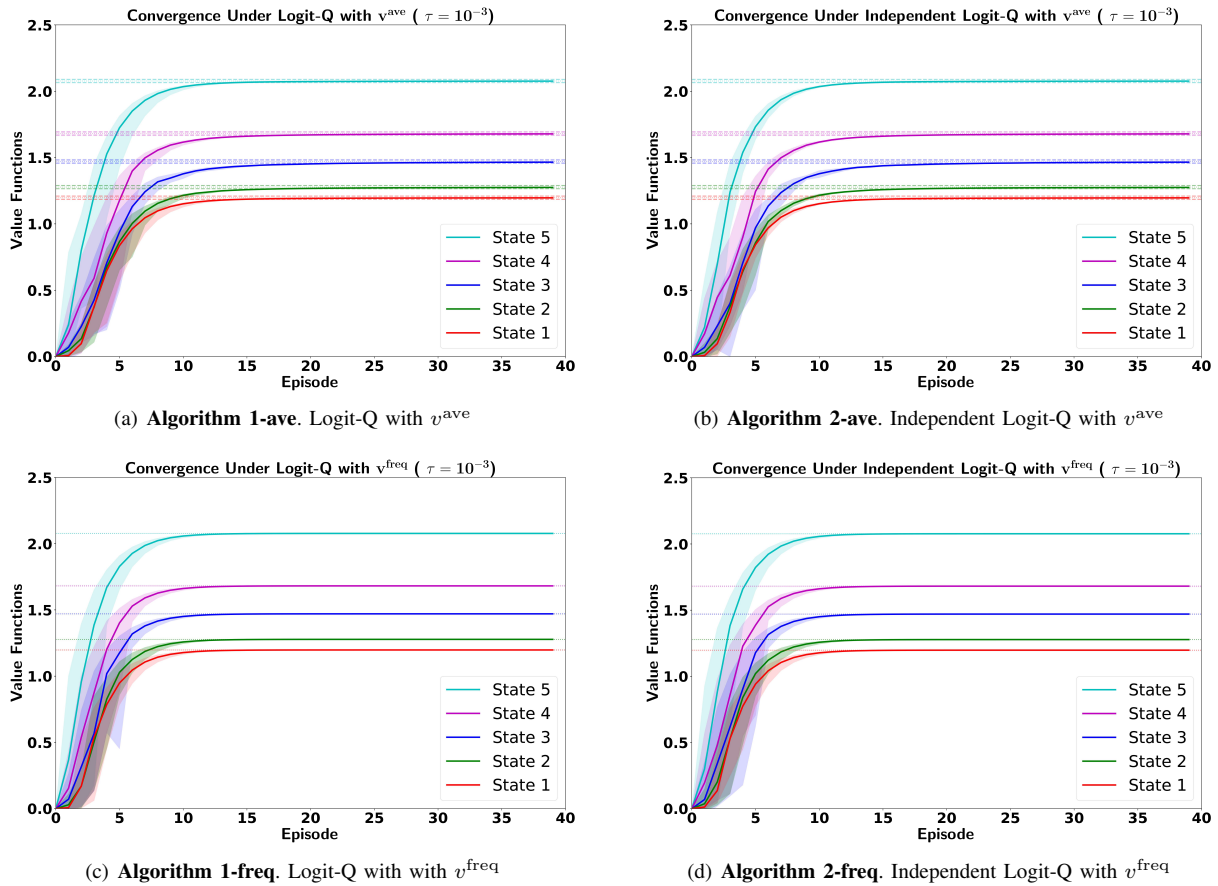


Fig. 2. Convergence of the value function estimates after 20 runs of **Algorithm 1-ave**, **Algorithm 2-ave**, **Algorithm 1-freq**, and **Algorithm 2-freq**, respectively. From bottom to top, the solid red, green, blue, magenta, and cyan curves represent the evolution of the value function estimates  $v_{(t)}(s)$  for states indexed from 1 to 5, respectively. The error bar around the solid curves represent the maximum and the minimum values of value function estimates as a result of 20 independent runs in **Algorithm 1**'s and **Algorithm 2**'s. The dotted lines ..... represent the actual optimal state values whereas the dashed lines - - - - represent the bound given by the Theorem 1 for the update scheme with  $v^{\text{ave}}$ . For all of the cases, value function estimates are observed to converge to the optimal values of states.

from agents' evolving strategies. Furthermore, the dynamics presented are online in the sense that agents play and learn simultaneously in a single SG over infinite horizon without any repeated play.

#### REFERENCES

- [1] L. S. Shapley, "Stochastic games," *Proceedings of National Academy of Science USA*, vol. 39, no. 10, pp. 1095–1100, 1953.
- [2] K. Zhang, Z. Yang, and T. Başar, "Multi-agent reinforcement learning: A selective overview of theories and algorithms," in *Handbook on RL and Control*, ser. Studies in Systems, Decision and Control, K. G. Vamvoudakis, Y. Wan, F. L. Lewis, and D. Cansever, Eds. Springer, Cham, 2021, vol. 325.
- [3] D. S. Leslie, S. Perkins, and Z. Xu, "Best-response dynamics in zero-sum stochastic games," *Journal of Economic Theory*, vol. 189, 2020.
- [4] M. O. Sayin, F. Parise, and A. Ozdaglar, "Fictitious play in zero-sum stochastic games," *SIAM Journal on Control and Optimization*, vol. 60, no. 4, pp. 2095–2114, 2022.
- [5] L. Baudin, "Best-response dynamics and fictitious play in identical interest stochastic games," in *Proc. Internat. Conf. Machine Learn. (ICML)*, 2022.
- [6] M. O. Sayin, K. Zhang, and A. Ozdaglar, "Fictitious play in Markov games with single controller," in *ACM Conference on Economics and Computation (EC)*, 2022, pp. 919–936.
- [7] D. Fudenberg and D. K. Levine, *The Theory of Learning in Games*. MIT press, 1998, vol. 2.
- [8] J. Marden and J. S. Shamma, "Revisiting log-linear learning: Asynchrony, completeness and payoff-based implementation," *Games and Economic Behavior*, vol. 75, pp. 788–808, 2012.
- [9] B. S. Pradelski and H. P. Young, "Learning efficient Nash equilibria in distributed systems," *Games and Economic Behavior*, vol. 75, no. 2, pp. 882–897, 2012.
- [10] T. Tatarenko, "Proving convergence of log-linear learning in potential games," in *2014 American Control Conference*, June 2014, pp. 972–977.
- [11] A. Ozdaglar, M. O. Sayin, and K. Zhang, "Independent learning in stochastic games," in *International Congress of Mathematicians*, 2021.
- [12] X. Wang and T. Sandholm, "Reinforcement learning to play an optimal Nash equilibrium in team Markov games," in *Advances in Neural Information Processing Systems*, 2002.
- [13] B. Yongacoglu, G. Arslan, and S. Yüksel, "Decentralized learning for optimality in stochastic dynamic teams and games with local control and global state information," *IEEE Transactions on Automatic Control*, vol. 67, no. 10, pp. 5230–5245, 2022.
- [14] M. O. Sayin and O. Unlu, "Efficient-Q learning for stochastic games," *arXiv:2302.09806*, 2023.
- [15] A. M. Fink, "Equilibrium in stochastic n-person game," *Journal of Science Hiroshima University Series A-I*, vol. 28, pp. 89–93, 1964.
- [16] L. Blume, "The statistical mechanics of strategic interaction," *Games and Economic Behavior*, vol. 5, pp. 387–424, 1993.
- [17] M. L. Littman, "Value-function reinforcement learning in Markov games," *Cognitive Systems Research*, vol. 2, no. 1, pp. 55–66, 2001.
- [18] M. O. Sayin and O. Unlu, "Logit-Q dynamics in Markov games," *arXiv:2205.13266v1*, 2022.
- [19] O. Unlu and M. O. Sayin, "Episodic Logit-Q dynamics for efficient learning in stochastic teams," *arXiv:2205.13266v2*, 2023.
- [20] D. A. Levin and Y. Peres, *Markov Chains and Mixing Times*, 2nd ed. American Mathematical Society, 2017.