

Asymptotic Convergence and Performance of Multi-Agent Q-learning Dynamics

Aamal Abbas Hussain
Imperial College London
London, United Kingdom
aamal.hussain15@imperial.ac.uk

Francesco Belardinelli
Imperial College London
London, United Kingdom
francesco.belardinelli@imperial.ac.uk

Georgios Piliouras
Singapore University of Technology
and Design
Tampines, Singapore
georgios@sutd.edu.sg

ABSTRACT

Achieving convergence of multiple learning agents in general N -player games is imperative for the development of safe and reliable machine learning (ML) algorithms and their application to autonomous systems. Yet it is known that, outside the bounds of simple two-player games, convergence cannot be taken for granted.

To make progress in resolving this problem, we study the dynamics of smooth Q-Learning, a popular reinforcement learning algorithm which quantifies the tendency for learning agents to explore their state space or exploit their payoffs. We show a sufficient condition on the rate of exploration such that the Q-Learning dynamics is guaranteed to converge to a unique equilibrium in any game. We connect this result to games for which Q-Learning is known to converge with arbitrary exploration rates, including weighted Potential games and weighted zero sum polymatrix games.

Finally, we examine the performance of the Q-Learning dynamic as measured by the Time Averaged Social Welfare, and comparing this with the Social Welfare achieved by the equilibrium. We provide a sufficient condition whereby the Q-Learning dynamic will outperform the equilibrium even if the dynamics do not converge.

KEYWORDS

Multi-Agent Learning; Quantal Response Equilibrium; Learning in Games

ACM Reference Format:

Aamal Abbas Hussain, Francesco Belardinelli, and Georgios Piliouras. 2023. Asymptotic Convergence and Performance of Multi-Agent Q-learning Dynamics. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023)*, London, United Kingdom, May 29 – June 2, 2023, IFAAMAS, 9 pages.

1 INTRODUCTION

Understanding the behaviour of multi-agent learning systems has been a hallmark problem of game theory and online learning. The requirement is that agents must explore potentially sub-optimal decisions whilst interacting with other agents to ultimately maximise their long-term reward. In contrast to online learning with a single agent, this poses a fundamentally non-stationary problem, in which convergence to an equilibrium is not always guaranteed.

In fact, recent work has consistently found that, when learning on games, agents may present a wide array of behaviours. This includes cycles [7, 10, 28], and even chaos [2, 11, 36–38]. Furthermore,

the equilibria of a game need not be unique, so even if convergence is guaranteed, it may be to one of many (or even a continuum) of equilibria. Thus, predicting the behaviour of online learning in games with many players becomes a particularly challenging problem.

Yet it remains an important problem to solve. Recent advances in machine learning require training multiple neural networks for applications to generative models [5, 16]. In order for such applications to be realised, it is required that that training provably converges to an equilibrium. Of equal importance is that this equilibrium be unique, so that the outcome remains consistent regardless of initial conditions. A unique equilibrium guarantees not only the reproducibility of the system, but also ensures that desired behaviours will persist, even if the system is perturbed from its desired state. Similarly, the most complex tasks often require the interaction of multiple autonomous agents [13]. This again requires that agents are able to reliably equilibrate their behaviour.

There is a strong, and ongoing, effort in the research community to understand these learning behaviours, with positive convergence results being found in an assortment of game structures. For instance, games with two players and two actions are well understood [21, 31, 34]. Beyond this, some of the most widely studied games are potential games, in which agents collaborate to maximise a shared global function, and zero sum games (and its network variants), in which agents are in competition. Indeed, it has been found that a number of learning algorithms, including Fictitious Play [41], Q-Learning [38, 43, 46], Replicator Dynamics [17, 18, 25] all converge to equilibria (though not always unique) in potential games [14, 23]. In zero sum games, the former two converge to a unique fixed point [8, 24], whereas the latter is known to cycle, always maintaining its distance from the equilibrium [28].

Outside of this class of games, however, the story becomes much more complicated. A wide array of results show learning algorithms may be chaotic in even the simplest games [37]. Such complex behaviours become even more pronounced as the number of agents in the game increases [36]. However, they are also influenced by the structure of the game and the parameters of the learning algorithm [33]. This dichotomy between the range of possible learning behaviours and the convergence requirement of the applications motivate our central question:

Are there learning dynamics such that convergence to a unique equilibrium is guaranteed in any game?

Main Contribution. To answer this, we study the (smooth) Q-Learning dynamics, a popular online learning model which captures

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom. © 2023 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

the behaviour of agents who must balance their tendency to explore their strategy space against exploiting their payoffs.

In the first of our contributions, we answer the above question in a positive manner. Namely, we show that, through sufficient exploration, agents engaged in any game can reach a unique equilibrium. We parameterise the amount of exploration required in terms of the size of the game and the number of players. We then revisit previously established results of Q-Learning and show that convergence of the algorithm in weighted potential and weighted network zero sum games both follow as special cases of our main result.

To qualify our convergence result, we consider the *payoff performance* of learning dynamics in terms of the sum of payoffs received by all agents. We provide a condition whereby learning dynamics that do not converge can outperform the equilibrium. In such a case, it may be beneficial to assume no exploration on the parts of the agents. This result is supported by our experiments in which examples of games are considered where payoff performance degrades as agents are pushed towards an equilibrium through exploration.

Related Work. Our work applies the framework of monotone games, which encompass a number of classes of games, including potential [35] and network zero sum games [20]. Indeed recent work has considered the question of designing games which are monotone [44]. Most prominent in this is the study of network games [26]. These works relate the monotonicity of the game with properties of the network. From a design perspective, this makes the study of monotone games rather attractive, as it becomes possible to make any game, be it co-operative or competitive, monotone. This fact is also exploited in [27], in which it was shown how any game, regardless of its structure, can be made monotone by appropriate parameter tuning.

In addition, monotone games have been used to design online learning algorithms that converge to a Nash Equilibrium [9, 30]. However, many of these require the gradient (or estimates) of their cost functions at each step. Ideally, an online learning algorithm should require only obtained rewards at each time step. To resolve this, [44] derive a distributed algorithm which converges to a Nash Equilibrium in monotone games. These results are advanced in [45] in which an algorithm is developed which converges, whilst also achieving no-regret; this was the first instance of a no-regret algorithm who could provably converge in a monotone game. In a similar manner, [35] show the convergence of the Best Response Algorithm in a class of network games which satisfy the monotonicity property.

Our work departs from the above by considering Q-Learning, and by lifting strong technical assumptions on the payoff functions. Specifically, we do not assume a form of the payoffs as in [35], or the growth of the function as in [45]. In addition, we require no knowledge of the cost gradients as in [6]. Finally, our work considers the generalised class of *weighted* monotone games, rather than the unweighted case considered by the above. This class of games is also considered in [15, 48], in which variations of online gradient descent are analyzed. However, the former requires weighted strong monotonicity (which is much more restrictive even than strict monotonicity) and the latter requires strong assumptions on the parameters of the learning algorithm.

Our work also touches upon the Follow the Regularised Leader (FTRL) dynamic. The strongest convergence result regarding this dynamic in continuous time is a negative one: FTRL does not converge to a Mixed Nash Equilibrium [47], or to any equilibrium in zero sum games [28]. The strongest positive result, and the one most similar to our own is [12] in which it is shown that FTRL converges in unweighted strictly monotone games, a more restrictive class that that analysed here. Other positive results regard FTRLs local convergence to a strict Nash Equilibrium [29] and its convergence in time average [12].

Our paper is structured as follows. We begin in Section 2 by outlining the setting and tools through which we analyse convergence of learning. Section 3 proceeds with our main results, namely that convergence of Q-Learning dynamics can be achieved through sufficient exploration (Theorem 1), and that non-convergent learning dynamics can outperform the equilibrium (Theorem 4). Our experiments in Section 4 validate these results by showing examples of games in which convergence occurs through increased exploration, although at the cost of decreasing payoff across all agents.

2 PRELIMINARIES

In this section, we expand on the necessary background for our main results; specifically how the game model is set up, the learning dynamics of interest, and the techniques used in our analysis.

2.1 Game Model

In our study, we consider a game $\Gamma = (\mathcal{N}, (S_k, u_k)_{k \in \mathcal{N}})$, where \mathcal{N} denotes a finite set of agents indexed by $k = 1, \dots, N$. Each agent $k \in \mathcal{N}$ is equipped with a finite set of actions denoted by S_k with the number of actions $n_k := |S_k|$ as well as a payoff function u_k . We denote a *mixed strategy* (hereafter just *strategy*) \mathbf{x}_k of an agent k as a probability vector over its actions. Then, the set of all strategies of agent k is $\Delta_k := \{\mathbf{x}_k \in \mathbb{R}^{n_k} : \sum_i x_{ki} = 1, x_{ki} \geq 0\}$ on which act the payoff functions $u_k : \Delta_k \times \Delta_{-k} \rightarrow \mathbb{R}$. We denote by $\mathbf{x} := (\mathbf{x}_k)_{k \in \mathcal{N}} \in \Delta = \times_k \Delta_k$ the *joint strategy* of all agents and, for any k , $\mathbf{x}_{-k} := (\mathbf{x}_l)_{l \in \mathcal{N} \setminus \{k\}} \in \Delta_{-k}$ the joint strategy of all agents other than k .

For any $\mathbf{x} \in \Delta$, we define the reward to agent k for playing action $i \in S_k$ as $r_{ki}(\mathbf{x}) := \frac{\partial u_{ki}(\mathbf{x})}{\partial x_{ki}}$. We write $r_k(\mathbf{x}) = (r_{ki}(\mathbf{x}))_{i \in S_k}$ as the concatenation of all rewards to agent k . Using this notation we can write $u_k(\mathbf{x}) = \langle \mathbf{x}_k, r_k(\mathbf{x}) \rangle$ in which $\langle \mathbf{x}, \mathbf{y} \rangle = \mathbf{x}^\top \mathbf{y}$ denotes the standard inner product in \mathbb{R}^n . With this in mind, we define the *Equilibrium* of a game.

Definition 1 (Equilibrium). A joint mixed strategy $\bar{\mathbf{x}} \in \Delta$ is an *Equilibrium* if, for all agents k and all $\mathbf{x}_k \in \Delta_k$

$$\langle \mathbf{x}_k, r_k(\bar{\mathbf{x}}) \rangle \leq \langle \bar{\mathbf{x}}_k, r_k(\bar{\mathbf{x}}) \rangle \quad (1)$$

$\bar{\mathbf{x}}$ is a *strict equilibrium* if the inequality (1) is strict for all $\mathbf{x}_k \neq \bar{\mathbf{x}}_k$.

The equilibrium concept which we are primarily interested in with regards to Q-Learning is the Quantal Response Equilibrium (QRE) [4], which acts as an equilibrium concept when agents have bounded rationality.

Definition 2 (Quantal Response Equilibrium (QRE)). A joint mixed strategy $\bar{\mathbf{x}} \in \Delta$ is a *Quantal Response Equilibrium* (QRE) if, for all

agents k and all actions $i \in S_k$

$$\bar{x}_{ki} = \frac{\exp(r_{ki}(\bar{x}_{-k})/T_k)}{\sum_{j \in S_k} \exp(r_{kj}(\bar{x}_{-k})/T_k)}$$

where T_k denotes the *exploration rate* of the agent: low values of T_k correspond to a higher tendency to exploit the best performing action. This, purely rational behaviour, corresponds to the *Nash Equilibrium* in the limit $T_k \rightarrow 0$. Higher values of T_k , meanwhile, implies a higher exploration rate. The two equilibria concepts are related by the following results

Proposition 1 ([27] Proposition 2). Consider a game $\Gamma = (\mathcal{N}, (S_k, u_k)_{k \in \mathcal{N}})$ and take any $T_k > 0$. Define the *perturbed game* $\Gamma^H = (\mathcal{N}, (S_k, u_k^H)_{k \in \mathcal{N}})$ with the modified payoffs

$$u_k^H(\mathbf{x}_k, \mathbf{x}_{-k}) = \langle \mathbf{x}_k, r_k(\mathbf{x}_{-k}) \rangle - T_k \langle \mathbf{x}_k, \ln \mathbf{x}_k \rangle$$

Then $\bar{\mathbf{x}} \in \Delta$ is a QRE of Γ if and only if it is an NE of Γ^H .

Finally, in the application of our results (specifically Lemma 1), we will consider the *influence bound* of a game, which gives a notion of its *size*. Formally, we apply the definition from [27]

Definition 3 (Influence Bound). Consider a finite normal form game Γ , the *influence bound* δ is given by

$$\delta = \max_{k \in \mathcal{N}, i \in S_k, s_{-k}, \tilde{s}_{-k} \in S_{-k}} \{|r_{ki}(s_{-k}) - r_{ki}(\tilde{s}_{-k})|\}$$

where the pure strategies $s_{-k}, \tilde{s}_{-k} \in S_{-k}$ differ only in the strategies of one agent $l \neq k$.

Since $|r_{ki}(s_{-k}) - r_{ki}(\tilde{s}_{-k})|$ measures the change in reward to agent k for playing action i due to a change the other players' actions, the influence bound δ can be thought of as a measure of the maximum influence (in terms of reward) that any agent could receive from their opponents. In Section 4, we consider two games and show how their influence bound can be readily determined.

2.2 Learning Model

The principal multi agent learning model that we analyse in this study is the *smooth Q-Learning* (QL) dynamic [46] which is foundational in economics [3] and artificial intelligence [43]. In [38, 46] a continuous time approximation of the Q-Learning algorithm was found which accurately captures the behaviour of learning agents through the following ODE:

$$\frac{\dot{x}_{ki}}{x_{ki}} = r_{ki}(\mathbf{x}) - \langle \mathbf{x}_k, r_k(\mathbf{x}) \rangle - T_k (\ln x_{ki} - \langle \mathbf{x}_k, \ln \mathbf{x}_k \rangle) \quad (2)$$

A full review of this dynamic is beyond the scope of this study, however a derivation of (2) can be found in [24] as well as the result that fixed points of the Q-Learning dynamic correspond to the QRE of the game Γ .

It was shown in [23] that the transformation between the game Γ and the perturbed game Γ^H as described in Sec. 2.1 also relates the Q-Learning Dynamics to the well studied *replicator dynamics* (RD) [18, 25], in the following manner.

Lemma 1 ([23]). Consider the game $\Gamma = (\mathcal{N}, (S_k, u_k)_{k \in \mathcal{N}})$ and some $T_k > 0$. Then the Q-Learning dynamics (2) can be written as

$$\frac{\dot{x}_{ki}}{x_{ki}} = r_{ki}^H(\mathbf{x}) - \langle \mathbf{x}_k, r_k^H(\mathbf{x}) \rangle \quad (3)$$

where $r_{ki}^H(\mathbf{x}) = r_{ki}(\mathbf{x}) - T_k (\ln x_{ki} + 1)$. In particular, (2) corresponds to RD in the perturbed game $\Gamma^H = (\mathcal{N}, (S_k, u_k^H)_{k \in \mathcal{N}})$.

Follow the Regularised Leader. It is known [29, 47] that RD is derived as a particular instance of the *Follow the Regularised Leader* (FTRL) dynamic [39]. In essence, FTRL requires that the agents maximise their cumulative payoff up to the current time t . However, it also imposes a regularisation on the agents' actions which softens the arg max function. More formally we have that for every agent k ,

$$\begin{aligned} \mathbf{y}_k(t) &= \mathbf{y}_k(0) + \int_0^t \mathbf{r}_k(\mathbf{x}(s)) ds \\ \mathbf{x}_k(t) &= Q_k(\mathbf{y}_t) := \arg \max_{\mathbf{x}_k \in \Delta_k} \{\langle \mathbf{x}_k, \mathbf{y}_k \rangle - h_k(\mathbf{x}_k)\} \end{aligned} \quad (4)$$

To make RD compatible with FTRL, we make the following assumption on the regularisers

Assumption 1. For every agent k , the regulariser h_k is:

- (1) Continuously differentiable with differential ∇h_k which itself is Lipschitz on Δ_k , with constant L_k .
- (2) *Steep* in that $\|\nabla h_k(\mathbf{x})\| \rightarrow \infty$ as $\mathbf{x} \rightarrow \partial \Delta$ (c.f. [47]).
- (3) Strongly convex on Δ_k , i.e. there exists κ such that, for any $\mathbf{x}_k, \mathbf{y}_k \in \Delta_k$

$$\langle \nabla h_k(\mathbf{y}_k), \mathbf{x}_k - \mathbf{y}_k \rangle \leq h_k(\mathbf{y}_k) - h_k(\mathbf{x}_k) - \frac{\kappa}{2} \|\mathbf{x}_k - \mathbf{y}_k\|^2$$

The choice of h_k and r_k will, of course, depend on the application scenario. Our interest in this work is to consider the case $h_k(\mathbf{x}_k) = \sum_i x_{ki} \ln x_{ki}$, which satisfies Assumption 1, and from which RD is derived. As mentioned, smooth Q-Learning describes RD in the perturbed game Γ^H . This perturbation will be instrumental in proving our results on convergence.

2.3 Variational Inequalities and Game Theory

In this work we will examine game-theoretic concepts through the lens of Variational Inequalities [9]. This branch of research modifies the problem of finding an Equilibrium of a game to that of finding a solution to a variational inequality, which is defined as follows.

Definition 4 (Variational Inequalities). Consider a set $\mathcal{X} \subset \mathbb{R}^d$ and a map $F : \mathcal{X} \rightarrow \mathbb{R}^d$. The Variational Inequality $VI(\mathcal{X}, F)$ is given as

$$\langle \mathbf{x} - \bar{\mathbf{x}}, F(\bar{\mathbf{x}}) \rangle \geq 0, \quad \text{for all } \mathbf{x} \in \mathcal{X}. \quad (5)$$

We say that $\bar{\mathbf{x}} \in \mathcal{X}$ belongs to the set of solutions to a variational inequality $VI(\mathcal{X}, F)$ if it satisfies (5).

In this work, we will be considering the state space $\mathcal{X} = \Delta = \times_k \Delta_k$ alongside the map $F : \Delta \rightarrow \mathbb{R}^{Nn}$ defined as

$$F(\mathbf{x}) = (F_k(\mathbf{x}))_{k \in \mathcal{N}} = (-r_k(\mathbf{x}))_{k \in \mathcal{N}}$$

This map is sometimes referred to as the *pseudo-gradient* of the game [44]. Its properties allow for conditions to be found under which the equilibrium $\bar{\mathbf{x}}$ is unique. To illuminate these conditions, we have the following definition.

Definition 5 (Weighted Monotone Game). A game Γ with a continuous pseudo-gradient F is *weighted monotone* if there exist positive constants w_1, \dots, w_d such that, for all $\mathbf{x}, \mathbf{y} \in \Delta$,

$$\langle \mathbf{x} - \mathbf{y}, F(\mathbf{x}; \mathbf{w}) - F(\mathbf{y}; \mathbf{w}) \rangle \geq 0 \quad (6)$$

where $F(\mathbf{x}, \mathbf{w}) = (w_1 F_1, \dots, w_d F_d)^\top$.

Naturally, the game is *weighted strictly monotone* if, for all $\mathbf{x} \neq \mathbf{y} \in \Delta$, inequality (6) is strict.

Through these elements, it is possible to explore the nature of the equilibrium of a game from a variational perspective (see for instance [9, 27, 35]). The seminal results from this analysis, outlined in depth in [9] are:

Lemma 2. Consider a game Γ with pseudo-gradient $F(\mathbf{x})$. Then, $\bar{\mathbf{x}} \in \Delta$ is an Equilibrium of Γ if and only if it satisfies

$$\langle \mathbf{x} - \bar{\mathbf{x}}, F(\mathbf{x}) \rangle \geq 0 \quad \text{for all } \mathbf{x} \in \Delta.$$

Lemma 3. If $\Gamma = (\mathcal{N}, (S_k, u_k)_{k \in \mathcal{N}})$ is a strictly monotone game, it has a unique Equilibrium $\bar{\mathbf{x}} \in \Delta$.

In our work, we will leverage the fact that both of these Lemmas extend readily to the weighted case (as all weights are assumed non-negative). By applying these properties, it is our goal to examine whether the equilibrium will be reached by learning agents.

3 LEARNING IN WEIGHTED MONOTONE GAMES

We have two main results in this section. The first is that multi-agent Q-Learning can achieve convergence to an equilibrium in *any* game through sufficient exploration, parameterised by T_k . Following this result, we consider the optimality of convergence to an equilibrium. Namely we show a sufficient condition (convergence in time-average) in which non-fixed point behaviour (e.g. cycles) outperforms the equilibrium in terms of payoff.

3.1 Convergence through Sufficient Exploration

Without further delay, we state our first main result

Theorem 1. Let δ be the influence bound of an arbitrary game Γ with N players. Then, Q-Learning converges to the unique QRE $\bar{\mathbf{x}}$ if, for all agents k ,

$$T_k > \delta(N - 1) \quad (7)$$

This result provides a general condition from which the convergence to the QRE can be guaranteed in any game. As one might expect, the amount of exploration (the size of T_k) required to achieve this will be influenced by the size of the game (parameterised by δ) and the number of players in the game (parameterised by N). An interesting point to note is that this results supports that of [36] in which it was shown that non-fixed point behaviour is more likely with low exploration rates, if the number of players is increased.

PROOF SKETCH. Theorem 1 relies on two points: first, the perturbed game Γ^H is weighted strictly monotone (which gives uniqueness of $\bar{\mathbf{x}}$) if T_k satisfies (7), and second, the QL dynamics converges to a fixed point in such games. The first point is shown through [27] Theorem 1 so we focus on the latter.

We achieve this in the following manner. First we show that, along trajectories of FTRL in weighted strictly monotone games, the distance to the equilibrium $\bar{\mathbf{x}}$ is decreasing. From this, the convergence of replicator in strictly monotone games is immediate. The reader will recall that Q-Learning describes replicator dynamics in a perturbed game. We show that, if the original game Γ is weighted monotone, then the corresponding perturbed game Γ^H is weighted strictly monotone. Putting all this together yields the convergence of QL in weighted monotone games.

Step 1: Convergence of FTRL. In order to show that the distance to the equilibrium $\bar{\mathbf{x}}$ is being decreased, we must first define what we mean by *distance*. We do this through the *Bregman Divergence* [39].

Definition 6. Consider a set of functions $h_k : \Delta_k \rightarrow \mathbb{R}$ and a set of positive scalars $\mathbf{w} = (w_1, \dots, w_N)$ with $k \in \mathcal{N}$. The *Weighted Bregman Divergence* induced by $h = (h_k)_{k \in \mathcal{N}}$ between a set of probability vectors $\mathbf{x}, \mathbf{y} \in \Delta$ with $\mathbf{x} = (\mathbf{x}_k)_{k \in \mathcal{N}}, \mathbf{y} = (\mathbf{y}_k)_{k \in \mathcal{N}}$ is given by

$$\begin{aligned} W_B(\mathbf{x}||\mathbf{y}; h) &= \sum_k w_k D_B(\mathbf{x}_k||\mathbf{y}_k; h) \\ &= \sum_k w_k (h_k(\mathbf{y}_k) - h_k(\mathbf{x}_k) - \langle \nabla h_k(\mathbf{x}_k), \mathbf{y}_k - \mathbf{x}_k \rangle). \end{aligned}$$

Remark. In particular, with the choice $h_k(x_k) = \sum_i x_{ki} \ln x_{ki}$, the Weighted Bregman Divergence corresponds to the *Weighted Kullback-Leibler (KL) Divergence* defined by

$$W_{KL}(\mathbf{x}||\mathbf{y}) = \sum_{k \in \mathcal{N}} w_k D_{KL}(\mathbf{x}_k||\mathbf{y}_k) = \sum_{k \in \mathcal{N}} w_k \sum_{i \in S_k} x_{ki} \ln \frac{x_{ki}}{y_{ki}}. \quad (8)$$

Theorem 2. Consider a weighted strictly monotone game Γ . If each agent follows an FTRL algorithm whose regulariser satisfies Assumption 1 then, for any initial condition $\mathbf{x}(0)$, $\mathbf{x}(t)$ minimises the weighted Bregman divergence towards the unique equilibrium $\bar{\mathbf{x}}$.

Remark. The question of convergence to the equilibrium $\bar{\mathbf{x}}$ can also be established through Theorem 4.9 of [12] and showing that the assumptions required by their theorem are met by Assumption 1, which we do in the Appendix. However, in order to progress towards convergence of Q-Learning in *weighted monotone games* (i.e. to lift the strictness requirement), we take the extra step of showing that the *Bregman Divergence* is decreasing along trajectories.

Step 2: Convergence of RD. Using the specialisation $h_k(x_k) = \sum_i x_{ki} \ln x_{ki}$, and the relation between the Bregman and KL Divergences from , Theorem 2 implies that $W_{KL}(\bar{\mathbf{x}}||\mathbf{x}(t))$ is a Lyapunov function for RD in weighted strictly monotone games.

Corollary 1. Consider a weighted strictly monotone game Γ . Then trajectories $\mathbf{x}(t)$ under the replicator dynamics minimise the KL-Divergence to the unique Equilibrium.

Remark. Corollary 1 mirrors the result found by [42] (Prop. 4.5), which also showed that the KL-divergence decreases. Here, we show that this is a special case of Theorem 2.

Step 3: Convergence of Q-Learning. Finally, we recognise that the same transformation which takes Γ to Γ^H also takes Q-Learning to RD (c.f. [23] Lemma 3.1). Putting this all together, if T_k satisfies (7), the perturbed game Γ^H is strictly monotone, which yields the convergence of RD to the unique equilibrium $\bar{\mathbf{x}}$. This immediately gives the convergence of QL in the original game Γ . \square

3.2 Convergence through Arbitrary Exploration

Whilst Theorem 1 gives a condition to achieve convergence through sufficient exploration, it is known that there are games for which

convergence to a QRE can be achieved with any non-zero exploration rate T_k . These games are: weighted potential games [23] and weighted zero sum polymatrix games [24]. In this section, we point out that, under suitable game structures, the reduced requirement on exploration in these games stems from the fact that they are both examples of *monotone* games. We can, therefore, apply the following Theorem to yield convergence of QL in such games.

Theorem 3. If the game Γ is weighted monotone then, for any $T > 0$, Q-Learning converges to the unique QRE.

The proof of this statement, presented in full in the Supplementary Material, lies in recognising that the transformation which takes Γ to Γ^H is an additive term of the form

$$-T_k(\ln x_{ki} + 1)$$

which is a strictly monotone function in \mathbf{x}_k . As such, if Γ is already weighted monotone then, for any positive value of T_k , Γ^H must be weighted strictly monotone.

Weighted Potential Games. Our first application concerns the dynamics of learning in weighted potential games. These concern the cooperative setting in which the game admits a global function over all agents’ strategies. Formally, a game is called a weighted potential game if there exists a function $U : \Delta \rightarrow \mathbb{R}$ and positive weights $w_1, \dots, w_N > 0$ such that, for each player $k \in \mathcal{N}$, $U(\mathbf{x}_k, \mathbf{x}_{-k}) - U(y_k, \mathbf{x}_{-k}) = w_k (u_k(\mathbf{x}_k, \mathbf{x}_{-k}) - u_k(y_k, \mathbf{x}_{-k}))$ for all $\mathbf{x}_k, y_k \in \Delta_k$ and all $\mathbf{x}_{-k} \in \Delta_{-k}$.

Lemma 4. Consider a weighted potential game Γ with concave potential $U(\mathbf{x})$. Then, for any $T_K > 0$, the Q-Learning dynamics converges to the unique QRE $\bar{\mathbf{x}}$.

Of course, it must be noted that, in many games, the potential is not concave and so require a different approach towards showing convergence. [23] performs such an analysis and also considers the geometry of the multiple QRE which can exist due to the non-concavity of the potential.

Weighted Zero Sum Polymatrix Games. In this application, we consider the competitive setting through the *weighted zero sum network polymatrix game*. Formally, a network polymatrix game $\Gamma = ((\mathcal{N}, \mathcal{E}), (S_{kl}, A^{kl})_{(k,l) \in \mathcal{E}})$ includes a network $(\mathcal{N}, \mathcal{E})$ in which \mathcal{E} consists of pairs of agents $k, l \in \mathcal{N}$ who are connected in the network. Each edge is imbued with payoff matrices A^{kl} which denotes the reward to agent k against agent l . The total payoff received by agent k , then, is

$$u_k(\mathbf{x}_k, \mathbf{x}_{-k}) = \sum_{(k,l) \in \mathcal{E}} \langle \mathbf{x}_k, A^{kl} \mathbf{x}_l \rangle \quad (9)$$

A game is called *weighted zero sum network polymatrix game* if there exists positive constants w_1, \dots, w_N such that, for all $\mathbf{x} \in \Delta$

$$\sum_{k \in \mathcal{N}} w_k \langle \mathbf{x}_k, A^{kl} \mathbf{x}_l \rangle = 0 \quad (10)$$

Lemma 5. Consider a weighted zero sum network polymatrix game Γ . The unique QRE \mathbf{x} is globally asymptotically stable under QL for any $T_k > 0$.

The convergence of Q-Learning in this class of games is also proven through an alternate proof in [24], in which the weighted KL-divergence was also found as a Lyapunov function. Lemma 5 proves this point by showing that weighted network polymatrix zero sum games fall under the more general class of weighted monotone games. In such a manner, the results of [24], can be extended to all weighted monotone games.

3.3 Learning Outperforms the Equilibrium

In our second result, we consider the optimality of exploration as a means of reaching the equilibrium. We know, for instance, that RD, which corresponds to QL with zero exploration rates T_k , will not converge in a merely monotone game. However, in this case, we know from [12] that the *time-average* of the trajectory does reach the equilibrium. In this case we are presented with the following question

Should we apply non-zero exploration rates T_k to ensure convergence to an equilibrium, or allow the trajectory to remain non convergent?

We answer this question by considering the *payoff performance* of the dynamic, where performance is measured by the Time Averaged Social Welfare (TSW)

$$TSW = \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t SW(\mathbf{x}(s)) ds. \quad (11)$$

where $SW(\mathbf{x}) = \sum_k u_k(\mathbf{x}_k, \mathbf{x}_{-k})$. Intuitively, the Social Welfare $SW(\mathbf{x})$ measures the sum of payoffs received by all agents at some state $\mathbf{x} \in \Delta$. In (11), this quantity is averaged along trajectories of the learning dynamic. From the following theorem we see that, in monotone games, the social welfare is higher if agents play according to FTRL (and therefore do not converge), than if they were to play according to the equilibrium.

Theorem 4. Consider a polymatrix game $\Gamma = ((\mathcal{N}, \mathcal{E}), (S_{kl}, A^{kl})_{(k,l) \in \mathcal{E}})$. It is the case that

$$u_k(\mathbf{x}_k, \mathbf{x}_{-k}) = \sum_{(k,l) \in \mathcal{E}} \langle \mathbf{x}_k, A^{kl} \mathbf{x}_l \rangle.$$

Supposing that the learning dynamic is such that $\lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t R(s) ds = 0$, where $R(t) = \sum_k R_k(t)$ and

$$R_k(t) = \max_{\mathbf{x}'_k \in \Delta_k} \int_0^t [u_k(\mathbf{x}'_k, \mathbf{x}_{-k}(s)) - u_k(\mathbf{x}_k(s), \mathbf{x}_{-k}(s))] ds,$$

is the *regret* of agent k [39] and supposing also that the time averaged strategy $\mu(t) := \frac{1}{t} \int_0^t \mathbf{x}(s) ds$ along the learning dynamics converges to the equilibrium $\bar{\mathbf{x}}$, then TSW is asymptotically greater than or equal to the Social Welfare of the equilibrium $\bar{\mathbf{x}}$, i.e. $TSW \geq SW(\bar{\mathbf{x}})$

Remark. In particular this result holds for FTRL dynamics in monotone games, for which the time average approaches the equilibrium $\bar{\mathbf{x}}$ asymptotically [12]. In addition, FTRL is a *no-regret* dynamic, in that $\lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t R_k(s) ds = 0$ for all agents k .

PROOF. Define $\mu_k(t) \in \Delta_k$ as the time averaged strategy of agent k , i.e.

$$\mu_k(t) = \frac{1}{t} \int_0^t \mathbf{x}_k(s) ds. \quad (12)$$

Then,

$$\begin{aligned} R_k(t) &= \max_{x'_k \in \Delta_k} \int_0^t u_k(x'_k, x_{-k}(s)) ds - \int_0^t u_k(x_k(s), x_{-k}(s)) ds \\ &\geq \int_0^t u_k(\mu_k(t), x_{-k}(s)) ds - \int_0^t u_k(x_k(s), x_{-k}(s)) ds \\ &= \int_0^t \mu_k(t) \cdot \sum_{(k,l) \in \mathcal{E}} A^{kl} x_l(s) ds - \int_0^t x_k(s) \cdot \sum_{(k,l) \in \mathcal{E}} A^{kl} x_l(s) ds \end{aligned}$$

Then, dividing by t and applying (12) to x_l , the inequality reads

$$\frac{1}{t} R_k(t) \geq \mu_k(t) \cdot \sum_{(k,l) \in \mathcal{E}} A^{kl} \mu_l(t) - \frac{1}{t} \int_0^t x_k(s) \cdot \sum_{(k,l) \in \mathcal{E}} A^{kl} x_l(s) ds$$

Taking the sum over all agents k and defining $R(t) = \sum_k R_k(t)$ we have

$$\begin{aligned} \frac{1}{t} R(t) &\geq \sum_k \left(\mu_k(t) \cdot \sum_{(k,l) \in \mathcal{E}} A^{kl} \mu_l(t) - \frac{1}{t} \int_0^t x_k(s) \cdot \sum_{(k,l) \in \mathcal{E}} A^{kl} x_l(s) ds \right) \\ &= SW(\mu(t)) - \frac{1}{t} \int_0^t SW(x(s)) ds \\ &\implies \frac{1}{t} \int_0^t SW(x(s)) ds \geq SW(\mu(t)) - \frac{1}{t} R(t) \\ &\implies \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t SW(x(s)) ds \geq \lim_{t \rightarrow \infty} SW(\mu(t)) \end{aligned}$$

where, in the final equality, we use the assumption that $\lim_{t \rightarrow \infty} \frac{1}{t} R(t) = 0$. Applying now the assumption that $\mu(t) \rightarrow \bar{x}$, we have the final result that

$$\lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t SW(x(s)) ds \geq SW(\bar{x}).$$

□

Remark. We point out here that Theorem 4 defines payoff performance as a *time-average* and as a sum over all agents. As such, it is possible that a particular agent is losing out on payoff consistently throughout learning. Similarly, it is possible that at some time t , all agents are performing worse than the equilibrium. Finally, the result concerns asymptotic behaviour, so it could be the case that the equilibrium initially outperforms the learning dynamic. However, Theorem 4 shows that, eventually, the cumulative, time-averaged performance (11) will be better than the equilibrium.

Similar results on the performance of non-fixed point learning dynamics have been found for Fictitious Play [32], Replicator Dynamics [22] and no-regret learning [1]. In the latter most of these, TSW was also applied as a measure of payoff performance and it was found that the payoff in discrete time *online mirror descent* outperforms the Nash Equilibrium. Theorem 4 adds to the body of literature studying payoff performance of learning by showing a sufficient condition (convergence in time-average) under which any no-regret learning algorithm (including Fictitious Play and Replicator) will outperform the equilibrium.

3.4 Discussion of Results

Theorem 1 presents a concrete approach for achieving convergence to a unique equilibrium in any game. This allows us to abstract beyond the typical arena of study: namely weighted potential or zero sum network games. In fact, we showed convergence in

both of these cases due to the fact that they satisfy the monotonicity assumption and, therefore, satisfy the assumptions of Theorem 3.

Though general in nature, these results are not without their limitations. They rely on the assumption of a discrete action set, so that agent strategies all evolve on Δ . This allows us to assume the existence of an Equilibrium, through the compactness of Δ . However, generalising to arbitrary continuous action sets would widen the range of applications which our work encompasses. In addition, Theorem 3 is derived for continuous time QL. This is a reasonable stance to take as it has been shown repeatedly that continuous time approximations of algorithms provide a strong basis for analysing the algorithms themselves [19, 46]. However, the accuracy of discrete time algorithms is always dependent on parameters, most notably step sizes. Such an analysis of the discrete variants presents a fruitful avenue for further research.

To qualify Theorem 1, we point out that it does not give any indication regarding the performance of the system, merely its behaviour. Furthermore, its success relies on the increasing exploration rates of the agents and, therefore, at the cost of their exploitation. We showed in Theorem 4 that, under certain conditions, agents following a non-convergent dynamic may actually outperform the equilibrium in terms of payoff. Our experiments monitor this effect further by showing that, in certain games, as T_k is increased the performance of the system decreases.

4 EXPERIMENTS

In the following experiments, we consider the limitation discussed in the previous section. Namely, we test to see whether the optimality of the learning algorithm may decrease as T_k is increased. To do this we consider two examples: the Mismatching Pennies and the Shapley games. The former, proposed in [22], is composed of a network of three agents, each equipped with two actions, Heads and Tails. The payoff given to each agent k is given by

$$\begin{aligned} u_k &= u_k(x_k, x_j) = x_k \mathbf{A} x_{k-1} \\ \mathbf{A} &= \begin{pmatrix} 0 & 1 \\ M & 0 \end{pmatrix}, \quad M \geq 1. \end{aligned}$$

For the latter [40], we examine a network variation on the original two player game towards a three player network variant in which payoffs, parameterised by $\beta \in (0, 1)$ are defined as

$$\begin{aligned} u_k &= u_k(x_k, x_{-k}) = x_k \mathbf{A} x_{k-1} + x_k \mathbf{B}^\top x_{k+1} \\ \mathbf{A} &= \begin{pmatrix} 1 & 0 & \beta \\ \beta & 1 & 0 \\ 0 & \beta & 1 \end{pmatrix}, \quad \mathbf{B} = \begin{pmatrix} -\beta & 1 & 0 \\ 0 & -\beta & 1 \\ 1 & 0 & -\beta \end{pmatrix}, \end{aligned}$$

Remark. In the case of a network polymatrix game, the influence bound of the game is

$$\max_{k \in \mathcal{N}, i \in \mathcal{S}_k, s_{-k}, \tilde{s}_{-k} \in \mathcal{S}_{-k}} | (A^k)_{i,s_{-k}} - (A^k)_{i,\tilde{s}_{-k}} |.$$

In other words, it is the maximum difference between any row elements across the payoff matrices for all agents. In the Mismatching game, the influence bound is M whilst the Shapley game has influence bound $1 + \beta$.

These games were analysed in [22] and [32] respectively, and it was shown that while the given learning algorithm (RD in the

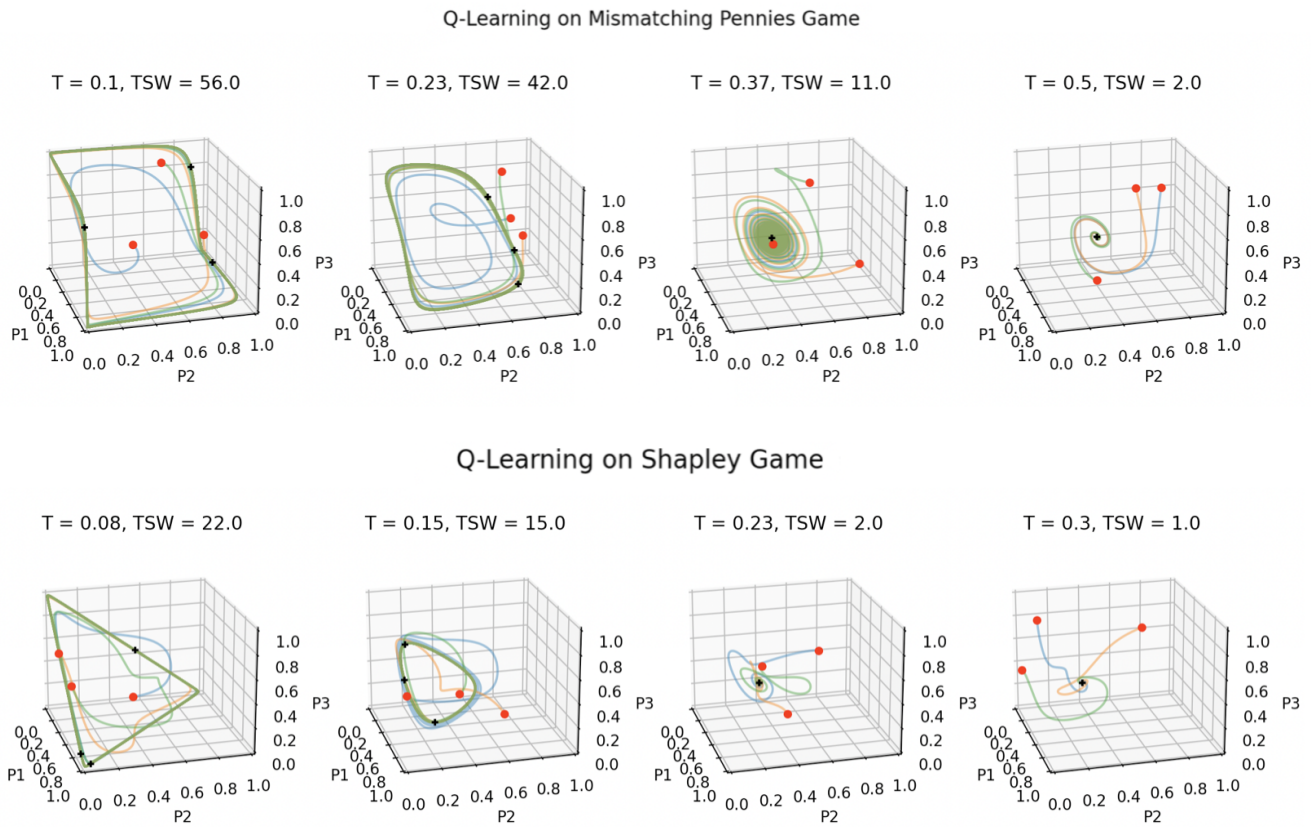


Figure 1: Trajectories of Q-Learning, red dots indicate initial conditions whilst black markers indicate final positions. TSW is averaged over all three initial conditions.

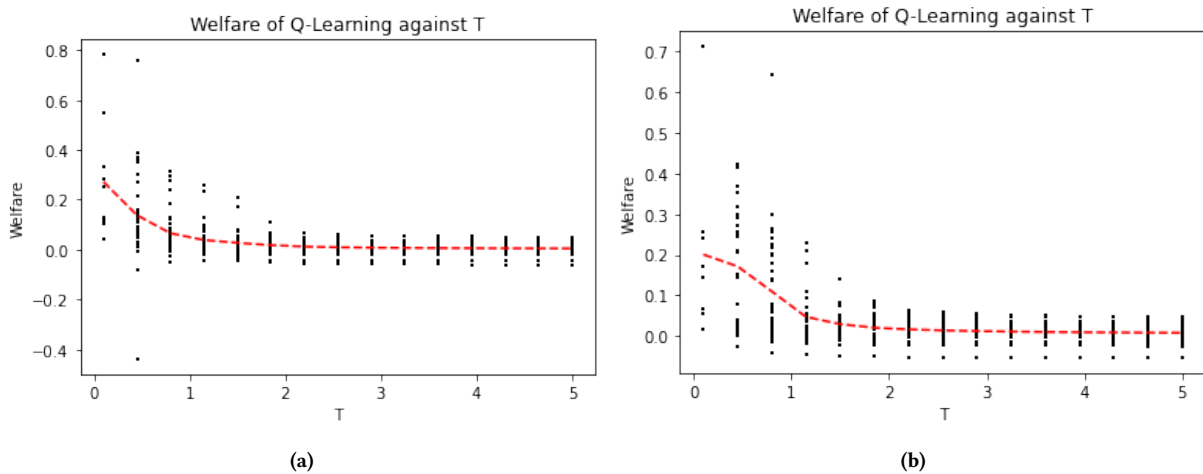


Figure 2: Normalised TSW against T for 35 randomly generated games. TSW is normalised to lie between $[-1, 1]$ in each game. Results are averaged over 10 initial conditions. The red line denotes the mean TSW across all 35 games. a. Five Players, Five Actions. b. Seven Players, Five Actions

former, Fictitious Play in the latter) did not converge to the NE, the agents actually received greater payoff through learning than they would have if they had only played the equilibrium strategy. The implication is that the agents were better off (in terms of payoff) by not converging to an equilibrium.

In Fig. 1 we plot the trajectories of QL for varying choices of T_k in each of the games. For the sake of simplicity, we enforce that all agents have the same T_k so we drop the k notation. The trajectories are displayed on the space $(0, 1)^3$ with each axis corresponding to the probability with which each agent plays their first action. Above each figure is displayed the choice of T for which Q-Learning is run, as well as the Time-averaged Social Welfare (TSW) along the trajectory, given by (11).

Two points become immediately clear from Fig. 1. The first is that as T is increased, the dynamics break no longer cycle around the equilibrium but rather converge to a unique equilibrium. While this occurs, however, TSW is decreasing. In fact, even in the case of equilibration, trajectories which take longer to reach the QRE gain a larger TSW. It is clear then, that it is in the agents' benefit if the dynamics remain unstable, at least as far as payoff is concerned.

In Fig. 2, we move beyond these indicative examples by evaluating TSW on 35 randomly generated games as T is increased. In order to accurately compare games with different payoff functions, we divide $\sum_k u_k(\mathbf{x}_k(t), \mathbf{x}_k(t))$ by the maximum possible cumulative payoff that an agent could receive in the game. This ensures that TSW remains within $[-1, 1]$ in all games. It is clear once again that, in general, TSW decreases as T increases, i.e. as the game moves towards more convergent behaviour. Of course, this does not hold in every game. However, the red line, which denotes the mean TSW across all games, suggests that this trend is the expected behaviour for a randomly selected game.

5 CONCLUSION

Our community has made strong strides in showing that online learning in games does not always reach an equilibrium. At the same time, the rising use of multiple interacting agents in machine learning applications necessitates placing guarantees on learning. In this paper, we make a step towards resolving this dichotomy by considering how the structure of a game, beyond the correlation between agent payoffs, affects online learning.

Specifically, we considered the asymptotic convergence to unique fixed points through Q-Learning (QL). Our analysis shows that the convergence in this popular learning dynamic can be guaranteed through sufficient exploration on the part of all agents. We also subsume convergence results in co-ordination (potential) games and competitive (network zero sum) games for which any positive rate of exploration is required. We then consider the impact of convergence through the lens of payoff performance and show that no-regret algorithms will outperform the equilibrium in terms of payoff, so long as the time-average trajectory reaches an equilibrium. In our experiments we show that this behaviour holds for a large number of games. An interesting point for future work would be to develop an analytical understanding for how often non-convergent learning dynamics outperform the equilibria of the game. As our study has shown, convergence of dynamics is inextricably linked to exploration. As such, by studying the optimality

of non-convergent dynamics, one may assess quantitatively the trade-off between exploration and exploitation.

ACKNOWLEDGMENTS

Aamal Hussain and Francesco Belardinelli are partly funded by the UKRI Centre for Doctoral Training in Safe and Trusted Artificial Intelligence (grant number EP/S023356/1). This research/project is supported in part by the National Research Foundation, Singapore and DSO National Laboratories under its AI Singapore Program (AISG Award No: AISG2-RP-2020-016), NRF 2018 Fellowship NRF-NRFF2018-07, NRF2019-NRF-ANR095 ALIAS grant, grant PIESGP-AI-2020-01, AME Programmatic Fund (Grant No.A20H6b0151) from the Agency for Science, Technology and Research (A*STAR) and Provost's Chair Professorship grant RGEPPV2101.

REFERENCES

- [1] Ioannis Anagnostides, Ioannis Panageas, Gabriele Farina, and Tuomas Sandholm. 2022. On Last-Iterate Convergence Beyond Zero-Sum Games. In *Proceedings of the 39th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 162)*, Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato (Eds.). PMLR, 536–581. <https://proceedings.mlr.press/v162/anagnostides22a.html>
- [2] Gabriel P Andrade, Rafael Frongillo, Mikhail Belkin, and Samory Kpotufe. 2021. Learning in Matrix Games can be Arbitrarily Complex. , 159–185 pages. <https://proceedings.mlr.press/v134/andrade21a.html>
- [3] Colin Camerer and Teck Hua Ho. 1999. Experience-weighted attraction learning in normal form games. *Econometrica* 67, 4 (7 1999), 827–874. <https://doi.org/10.1111/1468-0262.00054>
- [4] Colin F. Camerer, Teck Hua Ho, and Juin Kuan Chong. 2004. Behavioural game theory: Thinking, learning and teaching. *Advances in Understanding Strategic Behaviour: Game Theory, Experiments and Bounded Rationality* (1 2004), 120–180. https://doi.org/10.1057/9780230523371_8/COVER
- [5] Tong Che, Yanran Li, Athul Paul Jacob, Yoshua Bengio, and Wenjie Li. 2017. Mode Regularized Generative Adversarial Networks. In *International Conference on Learning Representations*.
- [6] Pierre Coucheny, Bruno Gaujal, and Panayotis Mertikopoulos. 2014. Penalty-Regulated Dynamics and Robust Learning Procedures in Games. <https://doi.org/10.1287/moor.2014.0687> 40, 3 (11 2014), 611–633. <https://doi.org/10.1287/MOOR.2014.0687>
- [7] Aleksander Czechowski and Georgios Piliouras. 2022. Poincaré-Bendixson Limit Sets in Multi-Agent Learning; Poincaré-Bendixson Limit Sets in Multi-Agent Learning. In *International Conference on Autonomous Agents and Multiagent Systems*. www.ifaamas.org
- [8] Christian Ewerhart and Kremena Valkanova. 2020. Fictitious play in networks. *Games and Economic Behavior* 123 (9 2020), 182–206. <https://doi.org/10.1016/j.geb.2020.06.006>
- [9] Francisco Facchinei and Jong Shi Pang. 2004. Finite-Dimensional Variational Inequalities and Complementarity Problems. *Finite-Dimensional Variational Inequalities and Complementarity Problems* (2004). <https://doi.org/10.1007/B97543>
- [10] Tobias Galla. 2011. Cycles of cooperation and defection in imperfect learning. *Journal of Statistical Mechanics: Theory and Experiment* 2011, 8 (8 2011). <https://doi.org/10.1088/1742-5468/2011/08/P08007>
- [11] Tobias Galla and J. Dooyne Farmer. 2013. Complex dynamics in learning complicated games. *Proceedings of the National Academy of Sciences of the United States of America* 110, 4 (2013), 1232–1236. <https://doi.org/10.1073/pnas.1109672110>
- [12] Saeed Hadikhanloo, Rida Laraki, Panayotis Mertikopoulos, and Sylvain Sorin. 2022. Learning in nonatomic games part I Finite action spaces and population games. *Journal of Dynamics and Games*. 2022 0, 0 (2022), 0. <https://doi.org/10.3934/JDG.2022018>
- [13] Heiko Hamann. 2018. *Swarm Robotics: A Formal Approach*. Springer International Publishing. <https://doi.org/10.1007/978-3-319-74528-2>
- [14] Christopher Harris. 1998. On the Rate of Convergence of Continuous-Time Fictitious Play. *Games and Economic Behavior* 22, 2 (2 1998), 238–259. <https://doi.org/10.1006/game.1997.0582>
- [15] Amélie Hélieu, Panayotis Mertikopoulos, and Zhengyuan Zhou. 2020. Gradient-Free Online Learning in Games with Delayed Rewards. In *Proceedings of the 37th International Conference on Machine Learning (ICML '20)*. JMLR.org.
- [16] Quan Hoang, Tu Dinh Nguyen, Trung Le, and Dinh Phung. 2018. MGAN: Training Generative Adversarial Nets with Multiple Generators. In *International Conference on Learning Representations*.

- [17] Josef Hofbauer and Karl Sigmund. 1998. *Evolutionary Games and Population Dynamics*. Cambridge University Press. <https://doi.org/10.1017/CBO9781139173179>
- [18] Josef Hofbauer and Karl Sigmund. 2003. Evolutionary Game Dynamics. *BULLETIN (New Series) OF THE AMERICAN MATHEMATICAL SOCIETY* 40, 4 (2003), 479–519.
- [19] Josef Hofbauer and Sylvain Sorin. 2005. Best response dynamics for continuous zero-sum games. *Discrete and Continuous Dynamical Systems - B*, 2006, Volume 6, Pages 215–224 6, 1 (10 2005), 215. <https://doi.org/10.3934/DCDSB.2006.6.215>
- [20] Amit Kadan and Hu Fu. 2021. Exponential Convergence of Gradient Methods in Concave Network Zero-Sum Games. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* 12458 LNAI (2021), 19–34. https://doi.org/10.1007/978-3-030-67661-2_2/FIGURES/3
- [21] Ardeshir Kianercy and Aram Galstyan. 2012. Dynamics of Boltzmann Q learning in two-player two-action games. *Physical Review E - Statistical, Nonlinear, and Soft Matter Physics* 85, 4 (4 2012), 041145. <https://doi.org/10.1103/PhysRevE.85.041145>
- [22] Robert Kleinberg, Katrina Ligett, Georgios Piliouras, and Eva Tardos. 2011. Beyond the Nash Equilibrium Barrier. *Innovations in Computer Science* (2011).
- [23] Stefanos Leonardos and Georgios Piliouras. 2022. Exploration-exploitation in multi-agent learning: Catastrophe theory meets game theory. *Artificial Intelligence* 304 (2022), 103653. <https://doi.org/10.1016/j.artint.2021.103653>
- [24] Stefanos Leonardos, Georgios Piliouras, and Kelly Spendlove. 2021. Exploration-Exploitation in Multi-Agent Competition: Convergence with Bounded Rationality. *Advances in Neural Information Processing Systems* 34 (12 2021), 26318–26331.
- [25] J. Maynard Smith. 1974. The theory of games and the evolution of animal conflicts. *Journal of Theoretical Biology* 47, 1 (9 1974), 209–221. [https://doi.org/10.1016/0022-5193\(74\)90110-6](https://doi.org/10.1016/0022-5193(74)90110-6)
- [26] Emerson Melo. 2018. A Variational Approach to Network Games. *SSRN Electronic Journal* (11 2018). <https://doi.org/10.2139/SSRN.3143468>
- [27] Emerson Melo. 2021. On the Uniqueness of Quantal Response Equilibria and Its Application to Network Games. *SSRN Electronic Journal* (6 2021). <https://doi.org/10.2139/SSRN.3631575>
- [28] Panayotis Mertikopoulos, Christos Papadimitriou, and Georgios Piliouras. 2018. Cycles in adversarial regularized learning. *Proceedings* (2018), 2703–2717. <https://doi.org/10.1137/1.9781611975031.172>
- [29] Panayotis Mertikopoulos and William H. Sandholm. 2016. Learning in Games via Reinforcement and Regularization. <https://doi.org/10.1287/moor.2016.0778>, 4 (8 2016), 1297–1324. <https://doi.org/10.1287/MOOR.2016.0778>
- [30] Panayotis Mertikopoulos and Zhengyuan Zhou. 2019. Learning in games with continuous action sets and unknown payoff functions. *Mathematical Programming* 173 (2019), 465–507. <https://doi.org/10.1007/s10107-018-1254-8>
- [31] Andrew I Metrick and Ben Polak. 1994. Fictitious play in 2×2 games: a geometric proof of convergence*. *Econ. Theory* 4 (1994), 923–933.
- [32] Georg Ostrovski and Sebastian van Strien. 2014. Payoff performance of fictitious play. *Journal of Dynamics and Games* 1, 4 (8 2014), 621–638. <https://doi.org/10.3934/jdg.2014.1.621>
- [33] Marco Pangallo, Torsten Heinrich, and J. Doyne Farmer. 2019. Best reply structure and equilibrium convergence in generic games. *Science Advances* 5, 2 (2 2019). https://doi.org/10.1126/SCIADV.AAT1328/SUPPL_FILE/AAT1328_SJM.PDF
- [34] Marco Pangallo, James B.T. Sanders, Tobias Galla, and J. Doyne Farmer. 2022. Towards a taxonomy of learning dynamics in 2×2 games. *Games and Economic Behavior* 132 (3 2022), 1–21. <https://doi.org/10.1016/J.GEB.2021.11.015>
- [35] Francesca Parise and Asuman Ozdaglar. 2019. A variational inequality framework for network games: Existence, uniqueness, convergence and sensitivity analysis. *Games and Economic Behavior* 114 (3 2019), 47–82. <https://doi.org/10.1016/j.geb.2018.11.012>
- [36] James B T Sanders, J Doyne Farmer, and Tobias Galla. 2018. The prevalence of chaotic dynamics in games with many players. *Scientific Reports* 8, 1 (2018), 4902. <https://doi.org/10.1038/s41598-018-22013-5>
- [37] Yuzuru Sato, Eizo Akiyama, and J. Doyne Farmer. 2002. Chaos in learning a simple two-person game. *Proceedings of the National Academy of Sciences of the United States of America* 99, 7 (4 2002), 4748–4751. <https://doi.org/10.1073/pnas.032086299>
- [38] Yuzuru Sato and James P. Crutchfield. 2003. Coupled replicator equations for the dynamics of learning in multiagent systems. *Physical Review E* 67, 1 (1 2003), 015206. <https://doi.org/10.1103/PhysRevE.67.015206>
- [39] Shai Shalev-Shwartz. 2011. Online Learning and Online Convex Optimization. *Foundations and Trends in Machine Learning* 4, 2 (2011). <https://doi.org/10.1561/22000000018>
- [40] L. S. Shapley. 2016. Some Topics in Two-Person Games. In *Advances in Game Theory. (AM-52)*. Princeton University Press, 1–28. <https://doi.org/10.1515/9781400882014-002>
- [41] Yoav Shoham and Kevin Leyton-Brown. 2008. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511811654>
- [42] Sylvain Sorin and Cheng Wan. 2016. Finite composite games: Equilibria and dynamics. *Journal of Dynamics and Games* 3, 1 (2016), 101–120.
- [43] R Sutton and A Barto. 2018. *Reinforcement Learning: An Introduction*. MIT Press. <http://incompleteideas.net/book/the-book-2nd.html>
- [44] Tatiana Tatarenko and Maryam Kamgarpour. 2019. Learning Nash Equilibria in Monotone Games. *Proceedings of the IEEE Conference on Decision and Control* 2019-December (12 2019), 3104–3109. <https://doi.org/10.1109/CDC40024.2019.9029659>
- [45] Tatiana Tatarenko and Maryam Kamgarpour. 2020. Bandit Learning in Convex Non-Strictly Monotone Games. *arXiv e-prints* (9 2020), arXiv:2009.04258.
- [46] Karl Tuyls, Pieter Jan T Hoen, and Bram Vanschoenwinkel. 2006. An evolutionary dynamical analysis of multi-agent learning in iterated games. , 115–153 pages. <https://doi.org/10.1007/s10458-005-3783-9>
- [47] Emmanouil-Vasileios Vlatakis-Gkaragkounis, Lampros Flokas, Thanasis Lianeas, Panayotis Mertikopoulos, and Georgios Piliouras. 2020. No-Regret Learning and Mixed Nash Equilibria: They Do Not Mix. *Advances in Neural Information Processing Systems* 33 (2020), 1380–1391.
- [48] Zhengyuan Zhou, Panayotis Mertikopoulos, Aris I Moustakas, Nicholas Bambos, and Peter Glynn. 2021. Robust Power Management via Learning and Game Design. *Operations Research* 69, 1 (2021), 331–345. <https://doi.org/10.1287/opre.2020.1996>