

On the Convergence of Model Free Learning in Mean Field Games

Romuald Elie,¹ Julien Pérolat,² Mathieu Laurière,³ Matthieu Geist,⁴ Olivier Pietquin⁴

¹Université Paris-Est, ²Deepmind

³ORFE, Princeton University, ⁴Google Research, Brain Team

romuald.elie@univ-mlv.fr, {perolat, mfgeist, pietquin}@google.com, lauriere@princeton.edu

Abstract

Learning by experience in Multi-Agent Systems (MAS) is a difficult and exciting task, due to the lack of stationarity of the environment, whose dynamics evolves as the population learns. In order to design scalable algorithms for systems with a large population of interacting agents (*e.g.*, swarms), this paper focuses on Mean Field MAS, where the number of agents is asymptotically infinite. Recently, a very active burgeoning field studies the effects of diverse reinforcement learning algorithms for agents with no prior information on a stationary Mean Field Game (MFG) and learn their policy through repeated experience. We adopt a high perspective on this problem and analyze in full generality the convergence of a fictitious iterative scheme using any single agent learning algorithm at each step. We quantify the quality of the computed approximate Nash equilibrium, in terms of the accumulated errors arising at each learning iteration step. Notably, we show for the first time convergence of model free learning algorithms towards non-stationary MFG equilibria, relying only on classical assumptions on the MFG dynamics. We illustrate our theoretical results with a numerical experiment in a continuous action-space environment, where the approximate best response of the iterative fictitious play scheme is computed with a deep RL algorithm.

1 Introduction

In Multi-agent systems (MAS), several autonomous robots or agents interact and cooperate, compete or coordinate in order to complete their task. The difficult nature of the task at hand combined with the large number of possible situations imply that the agents have to learn by experience. In comparison to the single-agent case, the derivation of efficient learning algorithms in this context is difficult due to the lack of stationarity of the environment, whose dynamics evolves as the population learns (Bu et al. 2008). This gives rise to research topics lying at the intersection of game theory and reinforcement learning. Nevertheless, in typical examples, the number of interacting agents can be very large (*e.g.*, swarm systems) and defies the scalability properties of most learning algorithms. For anonymous identical agents, a key simplification in game theory is the introduction of

the asymptotic limit where the number of agents is infinite, leading to the modeling intuition behind the theory of Mean Field Games (MFG). This calls for an analysis of model free learning scheme for MAS in terms of MFG.

MFG were introduced by Lasry and Lions (2006a; 2006b) and Huang, Malhamé, and Caines (2006) in order to model the dynamic equilibrium between a large number of anonymous identical agents in interactions. Such systems encompass the modeling of numerous applications such as traffic jam dynamics, swarm systems, financial market equilibrium, crowd evacuation, smart grid control, web advertising auction, vaccination dynamics, rumor spreading on social media, among others. In a sequential game theory setting, each player needs to take into account his impact on the strategies of the other players. Studying games with an infinite number of players is easier from this point of view, as the impact of one single player on the others can be neglected. Hereby, the asymptotic limit with infinite population size considered in MFG becomes highly relevant. A solution to a dynamic MFG is determined via the optimal policy of a representative agent in response to the flow of the entire population. A mean field (MF) Nash equilibrium arises when the distribution of the best response policies over the population generates the exact same population flow. In most cases, a MF Nash equilibrium provides an approximate Nash equilibrium for an analogous game with a finite number of players (Cardaliaguet 2013; Bensoussan, Frehse, and Yam 2013; Carmona and Delarue 2018).

In the abundant literature on MFG, most papers consider planning problems with fully informed agents about the game operation scheme, the reward function and the MF population dynamics. Only a few contributions focus on learning problems in MFG, see *e.g.* (Yin et al. 2010; Cardaliaguet and Lehalle 2018; Hadikhannloo 2018; Hadikhannloo and Silva 2019) for model based approaches. Very recently, a rapidly growing literature intends to approximate the solution of stationary MFG in the realistic setting where agents with no prior information on the game learn their best response policy through repeated experience. These contributions restrict to a stationary setting and focus on specific Reinforcement Learning (RL) algorithms: Q-learning (Guo et al. 2019; Yang et al. 2018b), fictitious

play (Mguni, Jennings, and Munoz de Cote 2018) or policy gradient methods (Subramanian and Mahajan 2019), and sometimes rely on hardly verifiable assumptions.

In this paper, we take a step back and adopt a general high perspective on the convergence of model free learning algorithms in possibly non-stationary MFG and emphasize their potential for MAS with a large number of agents. Our approach investigates how any single-agent learning algorithm can perform in an MFG setting, in order to learn a (possibly approximate) Nash equilibrium, via repeated experiences and without any prior knowledge. Namely, we quantify precisely how the convergence of model free iterative learning algorithms reduces to the error analysis of each learning iteration step, in analogy with how the convergence of RL algorithms reduces to the aggregation of repeated supervised learning approximation errors (Farahmand, Szepesvári, and Munos 2010; Scherrer et al. 2015). For this purpose, our approach relies on a model free Fictitious Play (FP) iterative learning scheme for repeated games, where each agent calibrates its belief to the empirical frequency of the previously observed population flows. The FP approach is very natural when agents are trying to learn how to play a game by experience, while interacting with others. Before a new round of experience, they need to anticipate the behavior of the other players, and FP ergodic averaging is nicely designed for this purpose. This algorithm is typically useful for building from experience collaboration a co-operation patterns in a MAS using a decentralized learning scheme. In our framework of interest, all agents are identical (as usual in MFG), and we consistently suppose that they use the same learning scheme. Whenever the agents can compute their exact best response to any population flow, FP is proved to reach asymptotically a Nash equilibrium in some (but not all (Shapley 1964)) classes of games, such as first order monotone MFG (Hadikhani 2018). However, in a realistic setting, the agents are not able to compute their exact best response and can only attain an approximate version of it. This induces at each iteration a learning approximation error, which propagates through the FP learning scheme.

The main contribution of this paper is theoretical, as we provide a rigorous study of the error propagation in Approximate FP algorithms for MFGs, using an innovative line of proof in comparison to the standard two time scale approximation convergence results (Leslie and Collins 2006; Borkar 1997). Our convergence results are derived under easily verifiable assumptions on possibly non-stationary MFG dynamics and cost, which are highly classical in the MFG literature (namely 1^{st} order monotone MFGs). This allows discussing the convergence to a (possibly approximate) MF Nash equilibrium, when using any standard single-agent learning algorithm as an inner step embedded in a FP iterative scheme. Especially, our theoretical framework encompasses the convergence of RL algorithms to MFG equilibria in non stationary settings, which, as far as we know, is new in the literature. We illustrate our theoretical results on an authoritative MFG numerical experiment on crowd congestion, where the approximate best response of the iterative FP scheme is computed with a deep RL algorithm. This provides for the first time a model free learning example on

MFG in a continuous state-action environment.

2 Background

Mean Field Games. MFGs were introduced by Lasry and Lions (2006a; 2006b) and by Huang, Malhamé, and Caines (2006) and correspond to the asymptotic limit of a differential game, where the number of agents is infinite. Since all agents are assumed to be identical and indistinguishable, individual interactions are irrelevant in the limit and only the distribution of states matters (see (Carmona and Delarue 2018) for a complete overview). Most of the MFG literature is displayed in continuous time, but we choose to present our analysis in a discrete time setting in order to alleviate the presentation and emphasize the fruitful connections with the learning literature.

Finding a Mean Field Nash equilibrium boils down to identifying the equilibrium distribution dynamics of the population as well as the best response (or optimal policy) of a representative agent to this population mean field flow. Since the number of players is infinite, each agent has an infinitesimal influence on the population distribution. Yet, since all agents are rational, at equilibrium the state distribution generated by the optimal policy must coincide with the population distribution.

Notations. Let \mathcal{X} and \mathcal{A} be compact convex subsets of \mathbb{R}^ℓ and \mathbb{R}^d respectively, which represent the state and action spaces common to every agent. Let $T > 0$ be a time horizon and let \mathbb{T} denote the time sequence $\{0, 1, \dots, T\}$. We denote by $\mathcal{P}(\mathcal{X})$ the set of probability measures on \mathcal{X} and by $M_T = \mathcal{P}(\mathcal{X})^T$ the set of all possible flows of population state distributions $\mu = (\mu_0, \mu_1, \dots, \mu_T)$. The initial distribution of the population is an atomless measure on \mathcal{X} denoted by μ_0 . For $\mu \in M_T$, μ_t represents the distribution at time t of the state occupation of the entire population.

State dynamics & Mean field population flow. At any time $t \in \mathbb{T}$, each agent belongs to a state $x_t \in \mathcal{X}$ and picks an action $a_t \in \mathcal{A}$. For a sequence of actions $a := (a_t)_{t \in \mathbb{T}}$, the dynamics of x is governed by a Markov Decision Process (MDP) with (possibly non-stationary) transition density $p_t(\cdot | x_t, a_t, \mu_t)$ parameterized by the mean field flow $\mu \in M_T$ of the population. This indexation transcribes the interactions with the other agents, through their state distribution μ_t . Typically, the dynamics of x is described by an equation of the form

$$x_{t+1} = x_t + b(x_t, a_t, \mu_t) + \epsilon_{t+1}, \quad (1)$$

where $b : \mathcal{X} \times \mathcal{A} \times \mathcal{P}(\mathcal{X}) \rightarrow \mathcal{X}$ is a drift function and $(\epsilon_t)_{t \geq 1}$ is a dynamic source of noise. We stress that the mean field term μ_t represents the whole population distribution and not just the average state, as e.g. in (Yang et al. 2018b).

We denote by Π_T the set of policies (or controls) $\pi : \mathbb{T} \times \mathcal{X} \rightarrow \mathcal{A}$ which are feedback in the state: at time t , an agent using policy π while in state x_t plays the action $a_t = \pi(t, x_t)$. The process x controlled by $\pi \in \Pi_T$ is denoted x^π .

Agent’s reward scheme. An infinitesimal agent starting at time 0 in state $x_0 \in \mathcal{X}$ chooses a policy $\pi \in \Pi_T$ in order to maximize the following discounted expected sum of

rewards:

$$J(x_0, \pi, \mu) := \mathbb{E} \left[\sum_{t=0}^{T-1} \gamma^t r(x_t^\pi, \mu_t, a_t) \right], \quad (2)$$

while interacting with the population MF flow $\mu \in M_T$. At time t , the agent's rewards are impacted by μ_t , which represents the aggregate state distribution of all the other agents (*i.e.* of the whole population). Since the agents are anonymous, only the MF distribution flow μ of the states matters.

As μ_0 denotes the state distribution at time 0, the average reward for a representative agent is given by

$$\mathbb{J}(\pi, \mu) := \mathbb{E}_{x_0 \sim \mu_0} [J(x_0, \pi, \mu)], \quad (3)$$

when this agent uses policy $\pi \in \Pi_T$, while the mean field population flow is μ .

Definition 1 (Best response). *A policy π^* maximizing $\mathbb{J}(\cdot, \mu)$ is called a best response of the representative agent to the MF population dynamic flow μ .*

MF Nash Equilibrium. While interacting, the agents may or may not reach a Nash equilibrium, whose definition, based on the previous best response policy characterization, reads as follows:

Definition 2 (Mean Field Nash equilibrium). *A pair $(\pi_t^*, \mu_t^*)_{t \in \mathbb{T}}$ consisting of a policy and a MF population distribution flow is called MF Nash equilibrium if it satisfies*

- **Agent rationality:** π^* is a best response to μ^* ;
- **Population consistency:** for all $t \in \mathbb{T}$, μ_t^* is the distribution of x_t^* , starting with distribution μ_0 and controlled by policy π^* .

Namely, if the mean field population flow is μ^* , the policy π^* is optimal, and if all the agents play according to π^* , the induced mean field population flow coincides with μ^* . Hereby, (π^*, μ^*) identifies to an MF Nash equilibrium.

Observe that reaching an MF Nash equilibrium requires the computation of the exact best response policy, which can be difficult in practice. We are concerned with the design of an iterative learning scheme, where the available best response is partially accurate and typically approximated by RL through repeated experiences. For example, this realistic situation arises when agents are repeatedly optimizing their daily driving trajectories, without any prior information on the traffic jam dynamics.

3 Fictitious Play Algorithms for MFG

Fictitious play (Robinson 1951) is an iterative learning scheme for repeated games, where each agent calibrates its belief to the empirical frequency of previously observed strategies of other agents, and plays optimally according to its beliefs. This constitutes its best response. Even in simple two-player games, the convergence of FP to a Nash equilibrium is not guaranteed (Shapley 1964). However, the convergence of FP has recently been proved for some classes of MFG (Hadikhannloo 2018; Cardaliaguet and Hadikhannloo 2017).

Yet, in most cases, agents do not have access to the exact best response policy π^* but use an approximate version of it instead, in the spirit of (Leslie and Collins 2006; Pérolat, Piot, and Pietquin 2018). At iteration n , the agent has only access to an approximate version $\hat{\pi}^{n+1}$ of the best response $\pi^{*,n+1}$ to the anticipated mean field flow $\bar{\mu}^{(n)}$, defined precisely in Algorithm 1.

At iteration step n , the learning scheme induces an average additional error ℓ_n defined as

$$\ell_n := \mathbb{J}(\pi^{*,(n+1)}, \bar{\mu}^{(n)}) - \mathbb{J}(\hat{\pi}^{(n+1)}, \bar{\mu}^{(n)}) \geq 0, \quad (4)$$

for $n \in \mathbb{N}$. Observe that ℓ_n identifies to the expected loss over the entire population at step n , when replacing the exact best response $\pi^{*,(n+1)}$ by the approximate policy $\hat{\pi}^{(n+1)}$.

In Section 4 below, we quantify the propagation of approximating errors ℓ_n and clarify the convergence properties of Algorithm 1 for any type of learning procedure at each intermediate step. The specific setting where the approximate optimal policy is computed using single-agent RL algorithms is discussed in Sec 4.3.

Algorithm 1: Approximate Fictitious Play for MFG

Data: An initial distribution μ_0 ; an initial flow of distributions $\bar{\mu}^{(0)}$; a number of steps N .

Result: A flow of distribution μ and a policy π .

```

1 begin
2   for  $n = 0, 1, \dots, N - 1$  do
3     Compute  $\hat{\pi}^{(n+1)}$ , the approximate response
       policy against  $\bar{\mu}^{(n)}$ 
4     Compute  $\hat{\mu}^{(n+1)}$ , the mean field flow associated
       to  $\hat{\pi}^{(n+1)}$  starting with  $\hat{\mu}_0^{(n+1)} = \mu_0$ 
5     Update  $\bar{\pi}^{(n+1)}$ , the uniformly randomized
       policy over  $(\hat{\pi}^{(k)})_{k=1, \dots, n+1}$ 
6     Update  $\bar{\mu}^{(n+1)} = \frac{n}{n+1} \bar{\mu}^{(n)} + \frac{1}{n+1} \hat{\mu}^{(n+1)}$ 
7
8   return  $(\bar{\mu}^{(N)}, \bar{\pi}^{(N)})$ 

```

Approximate Nash equilibrium At each step n , we denote by $\bar{\pi}^{(n)}$ the representative agent belief on the aggregate population policy, defined as an equally randomized version of all previous approximate best responses $(\hat{\pi}^{(k)})_{k=1, \dots, n}$: for each $t \in \mathbb{T}$ and $x \in \mathcal{X}$, $\bar{\pi}^{(n)}(t, x) \in \mathcal{P}(\mathcal{A})$ is the probability distribution on the set of actions \mathcal{A} according to which the player picks uniformly at random an element of $\{\hat{\pi}^{(1)}(t, x), \dots, \hat{\pi}^{(n)}(t, x)\}$.

With a slight abuse of notation, we write

$$J(x_0, \bar{\pi}^{(n)}, \bar{\mu}^{(n)}) := \frac{1}{n} \sum_{k=1}^n J(x_0, \hat{\pi}^{(k)}, \bar{\mu}^{(n)}), \quad n \in \mathbb{N}.$$

and modify the definition of \mathbb{J} in (3) accordingly. Observe for later use that, by construction, $\bar{\mu}^{(n)}$ defined in Algorithm 1 coincides with the population MF flow induced by

the policy $\bar{\pi}^{(n)}$. In order to assess the quality of $(\bar{\mu}^{(n)}, \bar{\pi}^{(n)})$ as an (approximate) MF Nash equilibrium, we introduce, for $n \in \mathbb{N}$,

$$e_n := \mathbb{J}(\pi^{*,(n+1)}, \bar{\mu}^{(n)}) - \mathbb{J}(\bar{\pi}^{(n)}, \bar{\mu}^{(n)}) \geq 0.$$

The exploitability e_n quantifies at iteration n the expected gain for a typical agent, when shifting its belief $\bar{\pi}^{(n)}$ to the exact best response $\pi^{*,(n+1)}$, while interacting with the MF population flow $\bar{\mu}^{(n)}$. After n iterations in Algorithm 1, e_n is a quantitative measure of the quality of $(\bar{\mu}^{(n)}, \bar{\pi}^{(n)})$ as an MF Nash equilibrium. For the sake of clarification, let us introduce a more precise weaker notion of MF Nash equilibrium, inspired by (Carmona 2004).

Definition 3 (Approximate MF Nash equilibrium). *For $\epsilon > 0$ and $\delta \in (0, 1)$, a pair $(\pi_{\epsilon, \delta}^*, \mu_{\epsilon, \delta}^*)$ consisting of a policy and a population distribution flow is called an (ϵ, δ) -MF Nash equilibrium if*

$$\mu_0 \left(\left\{ x_0; J(x_0, \pi_{\epsilon, \delta}^*, \mu_{\epsilon, \delta}^*) \geq J(x_0, \pi', \mu_{\epsilon, \delta}^*) - \epsilon, \forall \pi' \right\} \right)$$

is at least $1 - \delta$, and $\mu_{\epsilon, \delta}^*$ coincides with the MF distribution flow starting from μ_0 , when every agent uses policy $\pi_{\epsilon, \delta}^*$.

An (ϵ, δ) -MF Nash equilibrium identifies to a weak equilibrium which reveals ϵ -optimal for at least a fraction $(1 - \delta)$ of the population. We are now in position to clarify how the exploitability e_n quantifies the quality of $(\bar{\mu}^{(n)}, \bar{\pi}^{(n)})$ as an MF Nash equilibrium.

Theorem 4. *If $0 \leq e_n \leq \epsilon^2$ for some $n \in \mathbb{N}$, then $(\bar{\mu}^{(n)}, \bar{\pi}^{(n)})$ is an (ϵ, ϵ) -MF Nash equilibrium in the sense of Definition 3. If e_n goes to 0 as $n \rightarrow +\infty$, any accumulation point of $\bar{\mu}^{(n)}$ is a MF Nash equilibrium*

Proof. Fix $n \in \mathbb{N}$ and assume $0 \leq e_n \leq \epsilon^2$. Let us introduce

$$\varphi(x_0) := J(x_0, \pi^{*,(n+1)}, \bar{\mu}^{(n)}) - J(x_0, \bar{\pi}^{(n)}, \bar{\mu}^{(n)}) \geq 0,$$

as $\pi^{*,(n+1)}$ is the best response to the MF flow $\bar{\mu}^{(n)}$.

Using Markov's inequality and the bound on e_n we obtain

$$\begin{aligned} \mu_0 \left(\{x_0 \in \mathcal{X} : \varphi(x_0) \geq \epsilon\} \right) &= \mathbb{P}_{x_0 \sim \mu_0} [\varphi(x_0) \geq \epsilon] \\ &\leq \frac{\mathbb{E}_{x_0 \sim \mu_0} [\varphi(x_0)]}{\epsilon} = \frac{e_n}{\epsilon}, \end{aligned}$$

which is smaller than ϵ . Collecting the terms and using the definition of φ , we deduce that

$$\mu_0 \left(\left\{ x_0; J(x_0, \bar{\pi}^{(n)}, \bar{\mu}^{(n)}) \geq J(x_0, \pi^{*,(n+1)}, \bar{\mu}^{(n)}) - \epsilon \right\} \right)$$

is at least $1 - \epsilon$, so that $(\bar{\pi}^{(n)}, \bar{\mu}^{(n)})$ is an (ϵ, ϵ) -MF Nash equilibrium.

The second part of the theorem follows directly. \square

4 Error propagation & Nash equilibrium approximation for first order MFG

Since the exploitability e_n identifies to a relevant quality measure of Algorithm 1 after n iterations, we now evaluate how the individual learning errors $(\ell_k)_{0 \leq k \leq n}$ aggregate over e_n . For the sake of simplicity, we focus our discussion

on 1st order MFG, *i.e.* without source of noise in the dynamics. This allows us to build our reasoning on the analysis of (Hadikhanloo 2018, Chapter 3) and to avoid a limitative restriction to second order games with a potential structure, for which similar results should hold in that setting as well, see Cardaliaguet and Hadikhanloo (2017).

4.1 First order mean field game

The state $(x_t)_t$ evolves in \mathbb{R}^d with dynamics (1), where we take $b(x, a, \mu) = a$, and $(\epsilon_t)_t = 0$. In other words, each agent controls exactly its state variation between two time steps and does not endure any noise. While interacting with a MF flow $\mu \in M_T$, each agent intends to maximize the classical reward scheme given by (2) with a running reward at time t of the form:

$$r(x_t^\pi, \mu_t, a_t) \mapsto \tilde{r}(x_t^\pi, a_t) + \bar{r}(x_t^\pi, \mu_t), \quad (5)$$

where the extra \bar{r} captures the impact of the other agents' positions. In Sec. 5, we provide in particular a congestion example where \bar{r} models an appeal for non-crowded regions. This type of conditions translates into the so-called Lasry-Lions monotonicity condition Lasry and Lions (2006a; 2006b) which ensures uniqueness of MF Nash equilibrium. More precisely, existence and uniqueness of solution to the 1st order MFG of interest hold under the following classical set of assumptions.

Assumption 1. *For some constant C , the reward functions \tilde{r} and \bar{r} satisfy:*

- For any $x \in \mathcal{X}$, the map $\tilde{r}(x, \cdot)$ is twice differentiable and

$$\frac{1}{C} I_d \leq D_{aa} \tilde{r}(x, \cdot) \leq C I_d,$$

- The function \bar{r} is continuous on $\mathcal{X} \times \mathcal{P}(\mathcal{X})$ and $\bar{r}(\cdot, m)$ is C^1 on \mathcal{X} ,

- We have

$$\|\tilde{r}(\cdot, \cdot)\|_\infty + \|\bar{r}(\cdot, \cdot)\|_\infty \leq C,$$

- The Lasry-Lions monotonicity condition holds: for all $m_1, m_2 \in \mathcal{P}(\mathcal{X})$,

$$\int_{\mathcal{X}} [\bar{r}(\cdot, m_1) - \bar{r}(\cdot, m_2)] d[m_1 - m_2] < 0. \quad (6)$$

4.2 Error propagation in the Fictitious Play algorithm

We now investigate how the learning error $(\ell_n)_n$ propagates through FP for any learning algorithm, while Sec 4.3 focuses on the specific case where the best response is approximated via RL.

The key ingredient of FP iterative learning schemes is the quick stabilization of the sequence of beliefs $(\bar{\mu}^{(n)})_n$.

Lemma 5. *Under Assumption 1, the FP MF flow $\bar{\mu}^{(n)}$ satisfies:*

$$d_1(\bar{\mu}^{(n)}, \bar{\mu}^{(n+1)}) \leq \frac{C}{n}, \quad n \in \mathbb{N}, \quad \text{for some } C > 0,$$

where d_1 is the Wasserstein distance.

The proof follows from a straightforward adaptation of (Hadikhanloo 2018, Lemma 3.3.2) to our setting.

As the sequence of beliefs $\bar{\mu}_n$ stabilizes, the impact of recent learning errors (ℓ_n) reduces and we are in position to quantify the global error e_n of the algorithm after n iteration steps. This is the main result of the paper, whose proof interestingly differs from the more classical two-time scale approximation argument (Borkar 1997).

Theorem 6. *Under Assumption 1, the Nash equilibrium quality $(e_n)_n$ satisfies both estimates: for all $n \in \mathbb{N}$*

$$e_n \leq \frac{C_1}{n} + \frac{C_1}{n} \sum_{i=1}^n d_1(\mu^{*,(i+1)}, \hat{\mu}^{(i+1)}) + \frac{1}{n} \sum_{i=1}^n \ell_i, \quad (7)$$

$$e_n \leq \ell_n + \frac{C_2}{n} + \frac{C_2}{n} \sum_{i=1}^n d_1(\hat{\mu}^{(i+1)}, \hat{\mu}^{(i+2)}) + \sum_{i=1}^n \frac{i+1}{n} \ell_i, \quad (8)$$

for some constants C_1 and C_2 .

Sketch of Proof. Our argumentation builds up on the exact FP analysis of (Hadikhanloo 2018, Theorem 3.3.1), which hereby extends to the approximate best response setting.

Let us introduce the approximate exploitability

$$\hat{e}_n := \mathbb{J}(\hat{\pi}^{(n+1)}, \bar{\mu}^{(n)}) - \mathbb{J}(\bar{\pi}^{(n)}, \bar{\mu}^{(n)}) \geq 0,$$

so that $e_n = \ell_n + \hat{e}_n$, for $n \in \mathbb{N}$. In order to control the exploitability e_n , we focus our analysis on \hat{e}_n . Denoting $J_n : x \in \mathcal{X}^T \mapsto J(x_0, (x_{t+1} - x_t)_{t=0, \dots, T-1}, \bar{\mu}^{(n)})$, we get:

$$\begin{aligned} \hat{e}_{n+1} - \frac{n}{n+1} \hat{e}_n &= \int_{\mathcal{X}^T} J_{n+1} d(\hat{\mu}^{(n+2)} - \bar{\mu}^{(n+1)}) \\ &\quad - \frac{n}{n+1} \int_{\mathcal{X}^T} J_n d(\hat{\mu}^{(n+1)} - \bar{\mu}^{(n)}) \\ &= n \int_{\mathcal{X}^T} (J_{n+1} - J_n) d(\bar{\mu}^{(n+1)} - \bar{\mu}^{(n)}) \\ &\quad + \int_{\mathcal{X}^T} J_{n+1} d(\hat{\mu}^{(n+2)} - \hat{\mu}^{(n+1)}), \end{aligned}$$

where the last equality follows from the definition of $\bar{\mu}^{(n)}$. The monotonicity of the reward in Assumption 1 implies

$$\hat{e}_{n+1} - \frac{n}{n+1} \hat{e}_n \leq \int_{\mathcal{X}^T} J_{n+1} d(\hat{\mu}^{(n+2)} - \hat{\mu}^{(n+1)}).$$

Besides, Assumption 1 together with the compactness of \mathcal{X} and (Hadikhanloo 2018, Lemma 3.5.2) and Lemma 5 imply that $J_{n+1} - J_n$ is C/n -Lipschitz, leading to

$$\begin{aligned} \hat{e}_{n+1} - \frac{n}{n+1} \hat{e}_n &\leq \int_{\mathcal{X}^T} J_n d(\hat{\mu}^{(n+2)} - \hat{\mu}^{(n+1)}) \\ &\quad + \frac{C}{n} d_1(\hat{\mu}^{(n+2)}, \hat{\mu}^{(n+1)}). \end{aligned}$$

As $\pi^{*,(n+1)}$ is the best response to the mean field flow $\bar{\mu}^{(n)}$, recalling the definition of ℓ_n in (4), we deduce

$$\hat{e}_{n+1} - \frac{n}{n+1} \hat{e}_n \leq \ell_n + \frac{C}{n} d_1(\hat{\mu}^{(n+2)}, \hat{\mu}^{(n+1)}).$$

Together with estimate $e_n = \ell_n + \hat{e}_n$ and (Hadikhanloo 2018, Lemma 3.3.1), we derive (8) and conclude the proof. \square

Bound (7) indicates a nice averaging aggregation of the learning errors $(\ell_n)_n$, but requires a strong additional control on the Wasserstein distance between the MF flows generated by both approximate and exact best responses. Such estimate is readily available for the numerical approximation of convex stochastic control problems (Kushner and Dupuis 2013) but less classical in the RL literature, as discussed in Sec 4.3. When such an estimate is not available, Bound (8) provides a slower $n\ell_n$ convergence rate, up to a weak d_1 -regularity of the approximate best response in terms of the mean field flow $\bar{\mu}^{(n)}$, recall Lemma 5. Such estimate is highly classical in the setting of convex stochastic control problems with Lipschitz rewards (Fleming and Rishel 2012; Kushner and Dupuis 2013).

At finite distance, the following corollary sums up these properties in terms of MF Nash equilibrium.

Corollary 7. *Under Assumption 1, if ever $\frac{1}{n} \sum_{i=1}^n (\ell_i + C_1 d_1(\mu^{*,(i)}, \hat{\mu}^{(i)}))$ or $\frac{1}{n} \sum_{i=1}^n ((i+1)\ell_i + C_2 d_1(\hat{\mu}^{(i+1)}, \hat{\mu}^{(i)}))$ is bounded by $\epsilon^2/2$, $(\bar{\mu}^{(n)}, \bar{\pi}^{(n)})$ is an (ϵ, ϵ) -MF Nash equilibrium, for n large enough.*

In a similar fashion, we can conclude on the general asymptotic convergence of Algorithm 1 to the unique MF-Nash equilibrium, before discussing the specific implications for RL best response approximation schemes.

Corollary 8. *Under Assumption 1, the approximate FP algorithm converges to the unique MF Nash equilibrium whenever one of the following two conditions holds:*

1. *The approximate best response update procedure $\bar{\mu}^{(n)} \mapsto \hat{\mu}^{(n+1)}$ is continuous in d_1 , and $n\ell_n \rightarrow 0$, as $n \rightarrow \infty$;*
2. *The learning and policy approximation errors ℓ_n and $(d_1(\mu^{*,(n)}, \hat{\mu}^{(n)}))_n$ converge to 0.*

The convergence of the sequence $(\bar{\mu}^{(n)})_n$ follows from the tightness and pre-compactness property of this collection of measures with respect to the Wasserstein distance, see e.g. Remark 3.5.3 in (Hadikhanloo 2018).

4.3 Discussion on the convergence for Best Response RL approximation

The result in Theorem 6 is general and relies on standard assumptions of MFGs. It also relies on a good enough control of the approximation error on the best response at each iteration. Here, we discuss to what extent existing theoretical results for RL algorithms allow satisfying this assumption.

As stated in Corollary 8, in order for the approximate FP to converge to the exact MF Nash equilibrium, the approximate best response should converge quickly enough to the best one, depending on the number of iterations. From an RL perspective, this would require being able to compute the approximate optimal policy to an arbitrary precision, with high probability. As far as we know, such a result is possible only when an exact representation of any value function is possible, that is, in the tabular setting which imposes finite state and action spaces. Notably, convergence and rate of convergence of Q-learning-like algorithms have been studied in the literature, see e.g. (Szepesvári 1998; Kearns and Singh 1999; Even-Dar and Mansour 2003; Azar

et al. 2011). For example, the speedy Q-learning algorithm requires $\mathcal{O}(\ln(K)/(\epsilon^2(1-\gamma)^4))$ steps to learn an ϵ -optimal state-action value function with high probability, with K the number of state-action couples. According to Corollary 8, if the error is in $\mathcal{O}(n^{-\alpha})$ with $\alpha > 1$ (and if we have continuity in d_1), then the scheme converges to the Nash equilibrium. This suggests using $\mathcal{O}(\ln(K)n^{2\alpha}/(1-\gamma)^4)$ steps for the RL agent at iteration n . Yet, this kind of results does not provide guarantees on the continuity in d_1 .

According to Corollary 7, bounding the learning errors and the distance between two iterates of the distribution is sufficient to reach an approximate Nash equilibrium. As approximate FP can be seen as repeated RL problems, RL (or approximate dynamic programming) can be seen as repeated supervised learning problems, and the propagation of errors from supervised to RL is a well studied field, see e.g. (Farahmand, Szepesvári, and Munos 2010; Scherrer et al. 2015). Basically, if the supervised learning steps are bounded by some $\mathcal{O}(\epsilon)$, then the learning error of the RL algorithm is bounded by $\mathcal{O}(C\epsilon/(1-\gamma)^2)$, where C is a so-called concentrability coefficient, measuring the mismatch between some measures. In principle, we could then propagate the learning error of the supervised learning part up to the FP error, through the RL error. However, these results do not provide any guarantees on the proximity between the estimated optimal policy and the actual one (which would be a sufficient condition for the proximity between population distributions); it only provides a guarantee on the distance between their respective returns. This is due to the fact that in RL, the optimal value function is unique, but not the optimal policy. A perspective would be to consider regularized MDPs (Geist, Scherrer, and Pietquin 2019), where the optimal policy is unique (and greediness is Lipschitz). Yet, this would come at the cost of a bias in the Nash equilibria. The approach in (Guo et al. 2019) somehow builds partially on this idea in their specific learning scheme.

5 Numerical illustration

As an illustration, we consider a stylized authoritative MFG model with congestion in the spirit of Alnulla, Ferreira, and Gomes (2017). This application should be seen as a proof of concept showing that the method described above can be applied beyond the framework used for our theoretical results. We compute a model free approximation of the MFG solution combining Algorithm 1 with Deep Deterministic Policy Gradient (DDPG) (Lillicrap et al. 2016). As far as we know, this is the first numerical illustration of model free deep RL Algorithm for MFG with continuous states and actions. Our numerical results also demonstrate the empirical convergence of the Fictitious RL scheme in a larger setting, even when the MFG is of not first order type.

As usual in RL, instead of (2) we consider the problem in infinite horizon with the following discounted reward:

$$J(x_0, \pi, \mu) := \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t r(x_t^\pi, \mu_t, a_t) \right], \quad (9)$$

when an infinitesimal player interacts with the population MF flow $\mu = (\mu_0, \mu_1, \mu_2, \dots)$. The goal is to learn the pol-

icy which is optimal in the long run, i.e., when the behavior of the population becomes stationary.

Environment Each agent has a position x located on the torus $\mathbf{T} = [0, 1]$ with periodic boundary conditions (for simplicity of explicit solution), whose dynamics is governed by $x_{t+1} = x_t + a_t \Delta_t + \sqrt{\Delta_t} \epsilon_t$, $t = 0, 1, \dots$, where Δ_t is the time step of the continuous time process. It receives the per-step reward

$$r(x_t, \mu_t, a_t) = \tilde{r}(x_t) - \frac{1}{2}|a_t|^2 - \log(\mu_t),$$

where the last term motivates agents to avoid congestion, i.e. the proximity to a region with a large population density. In the continuous time setting with no discounting, a direct PDE argument provides the ergodic solution (a^*, μ^*) in closed form (Alnulla, Ferreira, and Gomes 2017)

$$a^* : x \mapsto \pi \cos(2\pi x) \quad \text{and} \quad \mu^* : x \mapsto \frac{e^{2\sin(2\pi x)}}{\int_{\mathbf{T}} e^{2\sin(2\pi y)} dy}, \quad (10)$$

when the geographic reward is of the form $\tilde{r}(x_t) = 2\pi^2 \sin(2\pi x_t) - 2\pi^2 \cos(2\pi x_t)^2 + 2\sin(2\pi x_t)$. This closed form solution offers a nice benchmark for our experiments and allows to measure the errors made by our algorithm.

Algorithm 2: Fictitious Play for continuous state and action Mean Field Games

Data: Number of FP iterations N_{FP} , an empty buffer of trajectories R_{FP} , the number of trajectories $N_{\text{trajectories FP}}$ to add to the replay buffer R_{FP} per steps

Result: a density $\bar{\mu}^{(N_{FP})}$ and a strategy $\hat{\pi}^{(N_{FP})}$.

```

1 begin
2   for  $n = 1, \dots, N_{FP}$  do
3     Compute a best response  $\hat{\pi}^{(n)}$  using DDPG
4     Collect trajectories  $N_{\text{trajectories FP}}$  of the strategy
        $\hat{\pi}^{(n)}$  and add them to the buffer  $R_{FP}$ 
5     Learn the population mean field distribution
        $\bar{\mu}^{(n)}$  by Monte Carlo samples on  $R_{FP}$ 
6     Compute the uniformly randomized policy  $\bar{\pi}^{(n)}$ 
       on  $(\hat{\pi}^{(1)}, \dots, \hat{\pi}^{(n)})$ 
7   return  $\bar{\mu}^{(N_{FP})}$  and  $\bar{\pi}^{(N_{FP})}$ 

```

Implemented Algorithm Model free FP for MFGs takes a somehow similar approach as Lanctot et al. (2017) in the sense that we estimate the best response using a model free RL algorithm (namely DDPG). However we do not maintain those best responses as in (Lanctot et al. 2017) but rather learn the population MF flow of the distribution of the representative agents. The best response approximation through DDPG and the estimation of the population MF are left in the Algorithm. We ran 30000 trajectories of DDPG with a trajectory length of 300. The noise used for exploration is a




Figure 1: From left to right: 1) the continuous time explicit solution for $\gamma = 1$, 2) the discrete time learned distribution of the policy with $\gamma = 0.95$ and $\Delta_t = 0.05$. 3) the L_2 error between the explicit solution and the density for each FP iterations, 4) the L_2 error between the control and the optimal control and for each FP iterations on a trajectory.

centered normal noise with variance 0.02 and we used Adam optimizers with 0.001 starting learning rate and $\tau = 0.01$. At each iteration of FP, we added $N_{\text{trajectories FP}} = 3000$ trajectories of length 1000 to the replay buffer. Finally, we estimated the density using 100 classes and doing 30000 steps of Adam (with 0.001 initial learning rate).

Results. Figure 1 presents the learned equilibrium computed for $\gamma = 0.95$, $N = 90$ and uniform initial distribution, as well as the continuous time closed form ergodic solution for $\gamma = 1$, see (10). We emphasize that the variation in γ together with the discrete/continuous time difference setting implies that the theoretical solutions to both problems are close but do not exactly coincide. We keep this benchmark since no ergodic closed form solution is available for $\gamma \neq 1$. As observed on Figure 1, both ergodic explicit and learned distributions and controls are close. As expected, the density of players is larger around the point of maximum of the reward but the distribution is not highly concentrated due to the logarithmic penalty encoding aversion for congested regions. More precisely, Figure 1 indicates that the L_2 errors between the distributions and the controls decrease with the number n of iterations. The convergence of control distributions echoes to the discussion on error propagation in Section 4.2. This clearly illustrates the numerical convergence of the Deep RL FP mean field algorithm.

6 Related work

The related literature is as follows. Recently, model free RL algorithms for solving MFGs were analyzed in the following papers: Guo et al. (2019) and Tiwari, Ghosh, and Aggarwal (2019) study Q -learning, Mguni, Jennings, and Munoz de Cote (2018) consider FP but contains several inaccuracies, as already pointed out in Subramanian and Mahajan (2019), which focuses on policy gradient methods. However, their studies are restricted to a stationary setting and focus on particular RL algorithms. Their convergence results hold under assumptions that are often hard to verify in practice. Although not focusing on an MFG, Yang et al. (2018b) uses the idea of MF approximation by considering interactions through the empirical mean action. Numerical illustrations provided in all these papers are in a finite state-action set-

ting, while we present a numerical example in a continuous state-action setting. On a different note, Yang et al. (2018a) studies the link between MFG and *inverse* RL. Some authors also study “learning” algorithms which use the full knowledge of the model (and hence are not model-free): Yin et al. (2010) studied a MF oscillator game while Hu (2019) proposed a decentralized deep FP learning architecture for large MARL, whose convergence holds on linear quadratic MFG examples with explicit solution and small maturity.

7 Conclusion and future research

In comparison to the existing literature focusing on specific RL algorithms for MFGs, we took a step back and offer a general perspective on the error propagation in iterative scheme for MFG, using any learning algorithm. We presented a rigorous convergence analysis of model free FP learning algorithm for MF Agent systems, encompassing cases where the best response is approximated using any single agent learning algorithm as well as non-stationary settings. We showed how the convergence of model free iterative FP algorithm reduces to the error analysis of each learning iteration step, as the convergence of RL algorithm reduces to the aggregation of repeated supervised learning approximation errors (Farahmand, Szepesvári, and Munos 2010; Scherrer et al. 2015). Our theoretical setting covers for the first time the consideration of non-stationary MFG and relies on reasonable and verifiable assumptions on the MFG of interest. The convergence is illustrated for the first time by numerical experiments in a continuous state-action setting, based on deep RL algorithm. Our analysis motivates and properly justifies the use of asymptotic Mean Field approximation for the study of learning by experience schemes in Multi-Agent systems, with a large number of agents. For RL approximation schemes, our analysis suggests a much faster convergence rate, whenever the best response approximation quality can be controlled in Wasserstein distance. This kind of estimate is classical in the numerical approximation of stochastic control literature but currently not available in the RL literature. The derivation of such estimate deserves to be addressed in future research papers. Finally, we focused on convergence properties for a centralized Multi-Agent learning algorithm, paving the way for addressing such property for a more relevant decentralized one.

References

- Almulla, N.; Ferreira, R.; and Gomes, D. 2017. Two numerical approaches to stationary mean-field games. *Dyn. Games Appl.* 7(4):657–682.
- Azar, M. G.; Ghavamzadeh, M.; Kappen, H. J.; and Munos, R. 2011. Speedy q-learning. In *Advances in neural information processing systems*, 2411–2419.
- Bensoussan, A.; Frehse, J.; and Yam, S. C. P. 2013. *Mean field games and mean field type control theory*. Springer Briefs in Mathematics. Springer, New York.
- Borkar, V. S. 1997. Stochastic approximation with two time scales. *Systems & Control Letters* 29(5):291–294.
- Bu, L.; Babu, R.; De Schutter, B.; et al. 2008. A comprehensive survey of multiagent reinforcement learning. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)* 38(2):156–172.
- Cardaliaguet, P., and Hadikhanloo, S. 2017. Learning in mean field games: the fictitious play. *ESAIM: Control, Optimisation and Calculus of Variations* 23(2):569–591.
- Cardaliaguet, P., and Lehalle, C.-A. 2018. Mean field game of controls and an application to trade crowding. *Mathematics and Financial Economics* 12(3):335–363.
- Cardaliaguet, P. 2013. Notes on mean field games.
- Carmona, R., and Delarue, F. 2018. *Probabilistic Theory of Mean Field Games with Applications I-II*. Springer.
- Carmona, G. 2004. Nash equilibria of games with a continuum of players.
- Even-Dar, E., and Mansour, Y. 2003. Learning rates for q-learning. *Journal of Machine Learning Research* 5(Dec):1–25.
- Farahmand, A.-M.; Szepesvári, C.; and Munos, R. 2010. Error propagation for approximate policy and value iteration. In *Advances in Neural Information Processing Systems*, 568–576.
- Fleming, W. H., and Rishel, R. W. 2012. *Deterministic and stochastic optimal control*, volume 1. Springer Science & Business Media.
- Geist, M.; Scherrer, B.; and Pietquin, O. 2019. A theory of regularized markov decision processes. In *International Conference on Machine Learning*.
- Guo, X.; Hu, A.; Xu, R.; and Zhang, J. 2019. Learning mean-field games. *arXiv preprint arXiv:1901.09585*.
- Hadikhanloo, S., and Silva, F. J. 2019. Finite mean field games: fictitious play and convergence to a first order continuous mean field game. *J. de Math. Pures et Appliquées*.
- Hadikhanloo, S. 2018. *Learning in Mean Field Games*. Ph.D. Dissertation, University Paris-Dauphine.
- Hu, R. 2019. Deep fictitious play for stochastic differential games. *arXiv preprint arXiv:1903.09376*.
- Huang, M.; Malhamé, R. P.; and Caines, P. E. 2006. Large population stochastic dynamic games: closed-loop McKean-Vlasov systems and the Nash certainty equivalence principle. *Commun. Inf. Syst.* 6(3):221–251.
- Kearns, M. J., and Singh, S. P. 1999. Finite-sample convergence rates for q-learning and indirect algorithms. In *Advances in neural information processing systems*, 996–1002.
- Kushner, H., and Dupuis, P. G. 2013. *Numerical methods for stochastic control problems in continuous time*, volume 24. Springer Science & Business Media.
- Lanctot, M.; Zambaldi, V.; Gruslys, A.; Lazaridou, A.; Tuyls, K.; Perolat, J.; Silver, D.; and Graepel, T. 2017. A unified game-theoretic approach to multiagent reinforcement learning. In *Advances in Neural Information Processing Systems*.
- Lasry, J.-M., and Lions, P.-L. 2006a. Jeux à champ moyen. I. Le cas stationnaire. *C. R. Math. Acad. Sci. Paris* 343(9):619–625.
- Lasry, J.-M., and Lions, P.-L. 2006b. Jeux à champ moyen. II. Horizon fini et contrôle optimal. *C. R. Math. Acad. Sci. Paris* 343(10):679–684.
- Leslie, D. S., and Collins, E. J. 2006. Generalised weakened fictitious play. *Games and Economic Behavior* 56(2):285–298.
- Lillicrap, T. P.; Hunt, J. J.; Pritzel, A.; Heess, N.; Erez, T.; Tassa, Y.; Silver, D.; and Wierstra, D. 2016. Continuous control with deep reinforcement learning. In *Proceedings of the International Conference on Learning Representations (ICLR 2016)*.
- Mguni, D.; Jennings, J.; and Munoz de Cote, E. 2018. Decentralised learning in systems with many, many strategic agents. In *AAAI*.
- Pérolat, J.; Piot, B.; and Pietquin, O. 2018. Actor-critic fictitious play in simultaneous move multistage games. In *AISTATS 2018-21st International Conference on Artificial Intelligence and Statistics*.
- Robinson, J. 1951. An iterative method of solving a game. *Annals of mathematics* 296–301.
- Scherrer, B.; Ghavamzadeh, M.; Gabillon, V.; Lesner, B.; and Geist, M. 2015. Approximate modified policy iteration and its application to the game of tetris. *Journal of Machine Learning Research* 16:1629–1676.
- Shapley, L. 1964. Some topics in two-person games. *Advances in game theory* 52:1–29.
- Subramanian, J., and Mahajan, A. 2019. Reinforcement learning in stationary mean-field games. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.
- Szepesvári, C. 1998. The asymptotic convergence-rate of q-learning. In *Advances in Neural Information Processing Systems*, 1064–1070.
- Tiwari, N.; Ghosh, A.; and Aggarwal, V. 2019. Reinforcement learning for mean field game. *arXiv preprint arXiv:1905.13357*.
- Yang, J.; Ye, X.; Trivedi, R.; Xu, H.; and Zha, H. 2018a. Deep mean field games for learning optimal behavior policy of large populations. In *International Conference on Learning Representations*.
- Yang, Y.; Luo, R.; Li, M.; Zhou, M.; Zhang, W.; and Wang, J. 2018b. Mean field multi-agent reinforcement learning. In *International Conference on Machine Learning*, 5567–5576.
- Yin, H.; Mehta, P. G.; Meyn, S. P.; and Shanbhag, U. V. 2010. Learning in mean-field oscillator games. In *49th IEEE Conference on Decision and Control (CDC)*, 3125–3132. IEEE.