

Mean Field Equilibrium in Multi-Armed Bandit Game with Continuous Reward

Xiong Wang¹, Riheng Jia²

¹The Chinese University of Hong Kong, Hong Kong SAR, China

²Zhejiang Normal University, Jinhua, China

xwang@cse.cuhk.edu.hk, rihengjia@zjnu.edu.cn

Abstract

Mean field game facilitates analyzing multi-armed bandit (MAB) for *a large number of agents* by approximating their interactions with an average effect. Existing mean field models for multi-agent MAB mostly assume a binary reward function, which leads to tractable analysis but is usually not applicable in practical scenarios. In this paper, we study the mean field bandit game with a *continuous reward function*. Specifically, we focus on deriving the existence and uniqueness of *mean field equilibrium* (MFE), thereby guaranteeing the asymptotic stability of the multi-agent system. To accommodate the continuous reward function, we encode the learned reward into an agent state, which is in turn mapped to its stochastic arm playing policy and updated using realized observations. We show that the state evolution is upper semi-continuous, based on which the existence of MFE is obtained. As the Markov analysis is mainly for the case of discrete state, we transform the stochastic continuous state evolution into a *deterministic* ordinary differential equation (ODE). On this basis, we can characterize a contraction mapping for the ODE to ensure a unique MFE for the bandit game. Extensive evaluations validate our MFE characterization, and exhibit tight empirical regret of the MAB problem.

1 Introduction

Great efforts have been devoted to multi-armed bandit (MAB) for sequential decision making, where agents can only observe limited information when pulling arms [Bubeck *et al.*, 2012]. Though principled, these techniques are mostly suitable for the single-agent scenario. As multiple agents may co-exist, multi-agent learning is proposed to investigate agent interactions [Panait and Luke, 2005], and particularly Markov game serves as the main tool to characterize the learning equilibrium in multi-agent systems [Littman, 1994]. Perfect information is usually needed in Markov game to determine the learning strategy, which makes it inefficient to analyze the equilibrium when the number of agents scales or when

dealing with MAB problems. To accommodate this bandit feedback, Hart and Mas-Colell propose a no-regret learning method [Hart and Mas-Colell, 2000], yet only the coarse correlated equilibrium can be derived for a handful of agents.

Mean field game is an effective model to approximate complex interactions among large populations [Lasry and Lions, 2007; Huang *et al.*, 2006], where the approximation error is $O(1/\sqrt{N})$ with N the number of agents [Ying, 2016]. Combined with online learning framework, mean field model is applied to multi-agent systems for deciding agents' strategic actions and characterizing system's stable state [Cardaliaguet and Hadikhannloo, 2017; Iyer *et al.*, 2014]. Existing mean field analysis often requires complete knowledge of not only the reward function, but also the historical information to obtain the mean field equilibrium (MFE), which is not applicable in scenarios with merely limited feedback available.

The problem of analyzing bandit feedback for many agents remains open until Gummadi *et al.* use the mean field model to study multi-agent MAB in a repeated game [Gummadi *et al.*, 2013], where they derive a unique MFE based on assumptions of binary reward and state regeneration. Following works have adopted their model in cellular network [Maghsudi and Hossain, 2017] and smart grid [Zhao and Liu, 2017]. However, the binary reward setting is too restricted for real-world multi-agent systems. Like in a resource competition game, agents usually share the resource with each other instead of occupying exclusively, thus the reward is a continuous value in $[0, 1]$ rather than only 0 or 1 [Yang *et al.*, 2018a; Hanif *et al.*, 2015]. Besides, agent state is assumed to regenerate with a certain probability to deduce an equilibrium in these works, whereas typical repeated games mostly entail iterative plays of each agent with no regeneration [Panait and Luke, 2005]. One critical question then is how to achieve the *equilibrium for the non-regenerated bandit game involving a large number of agents with a generalized continuous reward*.

In this paper, we propose a mean field model to tackle the bandit game in large-population multi-agent systems with a continuous reward function. We aim to resolve the *existence and uniqueness* issues of MFE, and hence are faced with the following challenges. First, characterizing the learning equilibrium is inconsistent with minimizing the regret since a non-regenerated bandit game is unstable under classical MAB algorithms. Therefore, one needs to model the agent state cautiously to strike a balance between system stability and tight-

Corresponding author: Riheng Jia

bounded regret. Second, existing Markov game is ineffective to track the state evolution, which is identified in a continuous regime instead of a discrete value due to the generalized reward function. Third, existence and uniqueness of MFE is hard to derive as we only observe a bandit feedback.

To handle these challenges, we encode and update an agent state using observed rewards, and devise a stationary policy to map the state to stochastic arm playing strategy in turn. We show that the state evolution satisfies a *fixed point* condition, thus proving the convergence of states to an existing MFE. Then, we develop a stochastic approximation to transform the stochastic state evolution into a *deterministic ordinary differential equation* (ODE), so that we can derive a contraction mapping for the ODE to obtain the unique MFE. Finally, we deduce a regret related cumulative state change, and extend the mean field model. Our contributions are summarized:

- We propose a new framework of mean field analysis to explore the multi-agent bandit game with a *continuous reward*. Our framework can *generalize* the previous binary reward function to a more universal scenario.
- We characterize the stable state for a large-population stochastic system with only limited information by deriving the existence and uniqueness of MFE. Specifically, we show there exists a MFE via *Kakutani fixed-point* theorem, and further devise both rigorous and relaxed conditions for $\|\cdot\|_\infty$ -contraction to obtain the unique MFE, thereby providing a tractable and guaranteed system performance.
- We encode the learned reward into agent state, which can both ensure the *system stability* and yield a *tight empirical regret*. Model extensions also reveal the *robustness* of our mean field analysis in different variants of MAB problems.

2 Model and Setup

We study a repeated multi-agent bandit game, where there are a large number of agents $\mathcal{N} = \{1, 2, \dots, N\}$ and a set of arms (actions) $\mathcal{M} = \{1, 2, \dots, M\}$. Time is slotted $\{0, 1, \dots, n, \dots\}$.

2.1 State and Playing Policy

The state of each agent i encodes its *learned reward* based on the observations realized so far, which is denoted as $s_n^i = [s_n^i(1), s_n^i(2), \dots, s_n^i(M)] \in \mathbb{R}^M$ with $s_n^i(j)$ being the learned reward of arm j upon to time slot n .

Agents are assumed to follow a *stationary policy* when solving MAB problems in the repeated bandit game [Cohen *et al.*, 2017]. Denote a simplex $\Delta^{M-1} = \{z \in [0, 1]^M : \sum_{j=1}^M z(j) = 1\}$ as the probability distribution over M arms. An arm playing policy is a mapping from the state space to the simplex $\sigma : \mathbb{R}^M \rightarrow \Delta^{M-1}$, i.e., $\sigma(s_n^i, j)$ means the probability that agent i would choose arm j with $\sum_{j=1}^M \sigma(s_n^i, j) = 1$. In particular, we focus on a Hedge stationary policy below:

$$\sigma(s_n^i, j) = (1 - \eta) \frac{\text{Exp}(\beta s_n^i(j))}{\sum_{k=1}^M \text{Exp}(\beta s_n^i(k))} + \frac{\eta}{M}, \quad (1)$$

where $\text{Exp}(\cdot)$ represents the exponential function and $\beta > 0$ is the smoothing parameter. The value $\sigma(s_n^i, j)$ consists of two parts balanced by parameter $\eta \in [0, 1]$, in which the first part is a logit policy and the second is a random selection.

2.2 Mean Field Reward

Let a_n^i be the played arm of agent i following policy $\sigma(s_n^i)$. Population profile $f_n = [f_n(1), f_n(2), \dots, f_n(M)]$ indicates the proportion of agents playing the various arms in time slot n . Denote $\mathbb{1}_{\{\cdot\}}$ as the indicator function, and then $f_n(j)$ is defined:

$$f_n(j) = \frac{1}{N} \sum_{i=1}^N \mathbb{1}_{\{a_n^i=j\}}. \quad (2)$$

Each agent's reward of playing an arm is determined by the actions of all agents due to their interactions. For instance, in a resource competition game, the reward will decrease if more competitors (agents) simultaneously compete for the same resource (arm). As directly characterizing the influence of agents' actions is difficult for large populations owing to the curse of dimensionality, we employ the mean field model to approximate interactions among agents, and accordingly their rewards will depend on the population profile f_n [Yang *et al.*, 2018b; Gummadi *et al.*, 2013]. Denote $r(f_n, a_n^i)$ as the realized reward of agent i when pulling arm a_n^i in time slot n , where $r(\cdot)$ is the reward function. With the bandit feedback, any agent only observes its realized reward without knowing the reward function and the population profile. Formally, $r(\cdot)$ is a continuous function rather than a discrete binary value as in [Gummadi *et al.*, 2013; Maghsudi and Hossain, 2017]. Apparently, continuous reward is more general, also widely adopted in both the single-agent MAB [Auer *et al.*, 2002a; Auer *et al.*, 2002b] and the multi-agent learning [Yang *et al.*, 2018b; Cohen *et al.*, 2017]. By convention, we assume that the reward $r(f_n, a_n^i)$ is in the range $[0, 1]$ which can be easily extended to other arbitrary intervals.

The state is updated after an agent obtains a realized reward. If agent i observes $r(f_n, a_n^i)$, we update its state as:

$$s_{n+1}^i(j) = (1 - \gamma_n) s_n^i(j) + \gamma_n w_n^i(j), \quad (3)$$

where:

$$w_n^i(j) = \begin{cases} r(f_n, a_n^i) & \text{if } a_n^i = j, \\ s_n^i(j) & \text{otherwise.} \end{cases} \quad (4)$$

The updating rule implies that only the state to the played arm is renewed while others remain unchanged. Moreover, the stepsize γ_n satisfies the following condition:

$$\sum_n \gamma_n = \infty, \quad \sum_n \gamma_n^2 < \infty. \quad (5)$$

2.3 Objective

Define $\mathbf{s}_n = [s_n^1, s_n^2, \dots, s_n^N]$ as the state profile, i.e., the states of all agents, thus $\mathbf{s}_n \in [0, 1]^{N \times M}$ since $r(f_n, j) \in [0, 1]$. Our objective is to analyze the convergence of \mathbf{s}_n , particularly to derive the existence and uniqueness of MFE. Meanwhile, we will also deduce the cumulative state change, which entails the regret information of the MAB problem.

3 Existence and Uniqueness of MFE

Characterizing the MFE is critical to a multi-agent system, because it can provide a guaranteed and predictable system performance [Adlakha *et al.*, 2015]. To achieve this goal, we need to answer two fundamental questions: 1) Does MFE exist? 2) If so, does there exist only one MFE?

The evolution of state profile s_n can be decomposed into playing and evolving processes. The playing process is the stationary policy of Eq. (1), mapping states to population profile $\Gamma_1 : s_n \rightarrow f_n$; the evolving process amounts to the state updating of Eq. (3), mapping population profile to states in turn $\Gamma_2 : f_n \rightarrow s_{n+1}$. Let $\Gamma = \Gamma_2 \circ \Gamma_1$ be the compound mapping, so the state evolution is interpreted as $\Gamma : s_n \rightarrow s_{n+1}$. The definition of MFE under mapping Γ is now presented.

Definition 1. State profile \bar{s} is a MFE if $\bar{s} = \Gamma(\bar{s})$.

3.1 Existence of MFE

Definition 1 indicates that a MFE is indeed a *fixed point under mapping* Γ . Since agents stochastically play actions based on Eq. (1), we obtain that Γ is a set mapping. Suppose the reward $r(f_n, j)$ is continuous in the population profile f_n .

Theorem 1. There exists a MFE \bar{s} satisfying $\bar{s} \in \Gamma(\bar{s})$.

Proof. The state profile s_n is in a nonempty, compact, and convex set $[0, 1]^{N \times M}$. The mapping Γ maps s_n to $s_{n+1} \in [0, 1]^{N \times M}$ which is a nonempty, closed, convex subset of $[0, 1]^{N \times M}$. We next prove Γ is upper semi-continuous.

Let \mathcal{P} be the set of the population profile f_n . For any state profile s_n , we have $\Gamma_1(s_n) \in \mathcal{P}$ according to the playing process. From Eq. (3), if agent i plays arm j , then its state corresponding to this arm after updating satisfies:

$$s_{n+1}^i(j) \in \bigcup_{f_n \in \mathcal{P}} \{(1 - \gamma_n)s_n^i(j) + \gamma_n r(f_n, j)\}. \quad (6)$$

Suppose there are arbitrary sequences $x_n, y_n \in [0, 1]^{N \times M}$ such that $\lim_{n \rightarrow \infty} x_n \rightarrow \bar{x}$, $\lim_{n \rightarrow \infty} y_n \rightarrow \bar{y}$, and $y_n \in \Gamma(x_n)$. Given a state profile, the set \mathcal{P} is determined. Since $r(f_n, j)$ is continuous in f_n , we can claim that $\bar{y} \in \Gamma(\bar{x})$ from Eq. (6), i.e., Γ is upper semi-continuous. By applying the Kakutani fixed-point theorem [Glicksberg, 1952], there exists a fixed point \bar{s} , or MFE, under the mapping Γ . \square

Theorem 1 ensures the existence of MFE, to which the states will ultimately converge. Since there may exist multiple MFEs, it is hard to determine which MFE the states will eventually approach. Therefore, we need to further derive a unique MFE so as to achieve a more guaranteed performance.

3.2 Uniqueness of MFE

Unique MFE implies there is only one fixed point. For mean field game, Lasry and Lions show that if the reward is a monotone function known by any agent, one will obtain a unique MFE [Lasry and Lions, 2007], which however is inconsistent with the MAB problem. To handle the bandit feedback, we use the stochastic approximation [Benaïm, 1999] to transform the discrete-time bandit game into a continuous-time ODE, and derive the unique MFE by proving that the ODE will only converge to one fixed point. Therefore, we can resort to *the deterministic ODE to figure out the stochastic bandit game*.

Stochastic Approximation

The state evolution is a stochastic process in that each agent plays arms randomly according to the stationary policy. Considering this fact, we rewrite the state updating of Eq. (3):

$$\begin{aligned} s_{n+1}^i(j) &= s_n^i(j) + \gamma_n (w_n^i(j) - s_n^i(j)) \\ &= s_n^i(j) + \gamma_n (\mathbb{E}[w_n^i(j)] - s_n^i(j) + u_n^i(j)), \end{aligned} \quad (7)$$

where $u_n^i(j) = w_n^i(j) - \mathbb{E}[w_n^i(j)]$. Denote $s_t = [s_t^1, s_t^2, \dots, s_t^N]$ as the state profile at continuous time t . In the following, we show that the state profile s_n at discrete time will *asymptotically converge* to s_t , which is characterized by a deterministic ODE. Hereinafter, we use t to index continuous time, and n to index discrete time slot.

Lemma 1. When $n \rightarrow \infty$ and $t \rightarrow \infty$, the state $s_n^i = [s_n^i(1), s_n^i(2), \dots, s_n^i(M)]$ of agent i will asymptotically converge to s_t^i specified by the following ODE:

$$\frac{ds_t^i}{dt} = \mathbb{E}[w_t^i | s_t] - s_t^i. \quad (8)$$

See Appendix A in the technical report [Wang and Jia, 2021] for the proof. With this lemma, if the ODE Eq. (8) of s_t solely converges to one fixed point, then there is a unique MFE for s_n . Note that the mapping Γ for the continuous-time state evolution now corresponds to the ODE Eq. (8). Specifically, using Eqs. (3) and (4), the ODE of $s_t^i(j)$ is obtained:

$$\frac{ds_t^i(j)}{dt} = \sigma(s_t^i, j) (r(f_t, j) - s_t^i(j)). \quad (9)$$

We emphasize that $r(f_t, j)$ in Eq. (9) is the *expected reward* over the state profile s_t as the population profile f_t is mapped from s_t by the stochastic stationary policy.

Convergence to Unique MFE

Consistent with Lemma 1, we first assume the reward function is a contraction mapping to obtain the unique fixed point, and then derive the conditions for this contraction mapping. We express the reward $r(f_t, j)$ of playing arm j as $r(f(s_t), j)$ to explicitly indicate its dependence on s_t . Also denote \bar{s} as the fixed point for ODE Eq. (9) at which the derivative is 0.

Theorem 2. Suppose that the reward function $r(f(s_t), j)$ is a $\|\cdot\|_\infty$ -contraction in the state profile s_t , then the fixed point \bar{s} is the unique MFE for the bandit game. \bar{s} is also the global attractor for the ODE Eq. (9), and s_t converges to \bar{s} with exponential rate.

Proof. Because $r(f(s_t), j)$ is a contraction mapping in s_t , $r(f(s_t), j) - s_t^i(j) = 0$ in Eq. (9) will have only one fixed point \bar{s} . From Lemma 1, we obtain that \bar{s} is also the unique MFE for the bandit game.

Construct Lyapunov function $V(s_t) = \|s_t - \bar{s}\|_\infty$, and assume $V(s_t)$ gets its maxima at $s_t^i(j)$, that is $V(s_t) = |s_t^i(j) - \bar{s}^i(j)|$. If $s_t^i(j) > \bar{s}^i(j)$, we have $V(s_t) = s_t^i(j) - \bar{s}^i(j)$. Take derivative over time t and use $r(f(\bar{s}), j) = \bar{s}^i(j)$:

$$\begin{aligned} \frac{dV(s_t)}{dt} &= \sigma(s_t^i, j) [r(f(s_t), j) - s_t^i(j) - r(f(\bar{s}), j) + \bar{s}^i(j)] \\ &\leq \sigma(s_t^i, j) [C_1 \|s_t - \bar{s}\|_\infty - |s_t^i(j) - \bar{s}^i(j)|] \\ &= \sigma(s_t^i, j) (C_1 - 1) |s_t^i(j) - \bar{s}^i(j)| \\ &= \sigma(s_t^i, j) (C_1 - 1) V(s_t), \end{aligned} \quad (10)$$

where $C_1 < 1$ because of the $\|\cdot\|_\infty$ -contraction mapping. As a result, $\frac{dV(s_t)}{dt} \leq 0$ and $\frac{dV(s_t)}{dt} = 0$ only at \bar{s} . If $s_t^i(j) \leq \bar{s}^i(j)$, following the same approach, we can also obtain $\frac{dV(s_t)}{dt} \leq \sigma(s_t^i, j) (C_1 - 1) V(s_t)$. Therefore, the fixed point \bar{s} is a global attractor for the ODE.

As Lyapunov function satisfies $\frac{dV(s_t)}{dt} \leq \sigma(s_t^i, j)(C_1 - 1)V(s_t)$, which implies $\frac{dV(s_t)}{V(s_t)} \leq \sigma(s_t^i, j)(C_1 - 1)dt$, we have $V(s_t) \leq C_2 \text{Exp}[\sigma(s_t^i, j)(C_1 - 1)t]$ where C_2 is a constant. Therefore, s_t converges to \bar{s} exponentially fast. \square

According to Lemma 1 and Theorem 2, we know that $s_n \rightarrow \bar{s}$ when $n \rightarrow \infty$. Next, we discuss the convergence rate of the discrete-time s_n . Recall from the stepsize γ_n in Eq. (5), we can set $\gamma_n = 1/(n+1)^\alpha$, $\alpha \in (1/2, 1]$.

Theorem 3. Suppose that the reward function $r(f(s_n), j)$ is a $\|\cdot\|_\infty$ -contraction in the state profile s_n . Denote the distance $e_n = \|s_n - \bar{s}\|_\infty$:

- 1) if $\alpha \in (\frac{1}{2}, 1)$, given $n = \Omega\left(\left(\frac{\ln \frac{1}{\epsilon^2}}{\epsilon^2}\right)^{\frac{1}{\alpha}} + \left(\ln \frac{1}{\epsilon}\right)^{\frac{1}{1-\alpha}}\right)$, then $e_n \leq \epsilon$ with probability at least $1 - \delta$;
- 2) if $\alpha = 1$, given $n = \Omega\left((2 + \Psi)^{\ln \frac{1}{\epsilon}} \frac{\ln \frac{1}{\delta\Psi\epsilon}}{\Psi^2\epsilon^2}\right)$, then $e_n \leq \epsilon$ with probability at least $1 - \delta$ for any positive constant Ψ .

Proof. Based on Theorems 2 and 3 in [Even-Dar and Mansour, 2003], we only need to prove $\|\mathbb{E}[w_n(s_n)] - \bar{s}\|_\infty \leq C\|s_n - \bar{s}\|_\infty$ with $C \in [0, 1)$, to draw the two conclusions. Assume $\|\mathbb{E}[w_n] - \bar{s}\|_\infty = \|\mathbb{E}[w_n^i(j)] - \bar{s}^i(j)\|$. According to Eq. (4) and the fact \bar{s} satisfies $r(f(\bar{s}), j) = \bar{s}^i(j)$, we have:

$$\begin{aligned} \|\mathbb{E}[w_n] - \bar{s}\|_\infty &= \|\mathbb{E}[w_n^i(j)] - \bar{s}^i(j)\| \\ &= |\sigma(s_n^i, j)r(f(s_n), j) + (1 - \sigma(s_n^i, j))s_n^i(j) - \bar{s}^i(j)| \\ &\leq |\sigma(s_n^i, j)(r(f(s_n), j) - r(f(\bar{s}), j))| \\ &\quad + |(1 - \sigma(s_n^i, j))(s_n^i(j) - \bar{s}^i(j))| \\ &\leq \sigma(s_n^i, j)C_1\|s_n - \bar{s}\|_\infty + (1 - \sigma(s_n^i, j))|s_n^i(j) - \bar{s}^i(j)| \\ &\leq \sigma(s_n^i, j)C_1\|s_n - \bar{s}\|_\infty + (1 - \sigma(s_n^i, j))\|s_n - \bar{s}\|_\infty \\ &\leq C\|s_n - \bar{s}\|_\infty. \end{aligned} \tag{11}$$

The second inequality holds as $r(f(s_n), j)$ is a contraction, and the third inequality is from $\|\cdot\|_\infty$ -definition. \square

Contraction Mapping Condition

The foundation of the unique MFE lies in that the reward function $r(f(s_t), j)$ is a contraction mapping in the state profile s_t . Suppose the reward function $r(f_n, j)$ is θ -Lipschitz continuous in the population profile f_n with regard to $\|\cdot\|_1$ -norm: $|r(f_n, j) - r(f'_n, j)| \leq \theta\|f_n - f'_n\|_1$. Combining with the idea in [Cominetti et al., 2010] and the mean field model, we now characterize the contraction mapping condition.

Theorem 4. If parameters β, η in Eq. (1) and θ satisfy the condition $4\theta(1 - \eta)\beta < 1$, then the reward function $r(f(s_t), j)$ is a $\|\cdot\|_\infty$ -contraction in the state profile s_t .

Proof. Let s_a and s_b be two state profiles, and define a sequence $\mathbf{A}_k = [s_a^1, s_a^2, \dots, s_a^k, s_b^{k+1}, \dots, s_b^N]$, $k = 1, 2, \dots, N$ with the first k elements from s_a and the rest from s_b . Denote A_k^i as the i -th state in \mathbf{A}_k , and \mathbf{a}_k as the played arms following the stationary policy $\sigma(A_k^i)$. Let \mathbf{a}_k^{-i} be the arm set except agent i , and $f(\mathbf{a}_k)$ or $f(\mathbf{A}_k)$ be the population profile. When $\mathbf{a}_k^i = j$, $|r(f(s_a), j) - r(f(s_b), j)| = |\sum_{k=1}^N [r(f(\mathbf{A}_k), j) -$

$$r(f(\mathbf{A}_{k-1}), j)]| \leq \sum_{k=1}^N |r(f(\mathbf{A}_k), j) - r(f(\mathbf{A}_{k-1}), j)| = \sum_{k=1}^N |r(f(\mathbf{a}_k), j) - r(f(\mathbf{a}_{k-1}), j)|.$$

Let $\Delta_k = r(f(\mathbf{a}_k), j) - r(f(\mathbf{a}_{k-1}), j)$, so we can express $\Delta_k = \Delta r_k \cdot (\sigma(s_a^k) - \sigma(s_b^k))$ where Δr_k is a M -length vector with the l -th element $\Delta r_k(l)$ as $\Delta r_k(l) = \sum_{z \in \mathbf{a}_k^{-i}, z^k=l} r(f(\{j, z\}), j) \prod_{m \neq k, i} \sigma(A_k^m, z^m)$. Note that $\Delta r_k(h)\mathbf{1} \cdot (\sigma(s_a^k) - \sigma(s_b^k)) = 0, \forall h \in \mathcal{M}$, where $\mathbf{1}$ is a M -length vector with each element equal to 1. Therefore, we can rewrite $\Delta_k = (\Delta r_k - \Delta r_k(h)\mathbf{1}) \cdot (\sigma(s_a^k) - \sigma(s_b^k))$. If $z_a, z_b \in \mathbf{a}_k^{-i}$ with only the k -th arm being different: $z_a^k = l, z_b^k = h$, we obtain $\|f(\{j, z_a\}) - f(\{j, z_b\})\|_1 \leq \frac{2}{N}$ from Eq. (2). As $r(f_n, j)$ is θ -Lipschitz continuous, then:

$$|\Delta r_k(l) - \Delta r_k(h)| \leq \left| \sum_{z \in \mathbf{a}^{-i}} \prod_{m \neq k, i} \sigma(A_k^m, z^m) \right| \frac{2\theta}{N} \leq \frac{2\theta}{N}. \tag{12}$$

As a result, we obtain:

$$|r(f(s_a), j) - r(f(s_b), j)| \leq \frac{2\theta}{N} \sum_{k=1}^N \|\sigma(s_a^k) - \sigma(s_b^k)\|_1. \tag{13}$$

From mean value theorem, there is a $x \in [s_a^k, s_b^k]$ such that:

$$\sigma(s_a^k, j) - \sigma(s_b^k, j) = \nabla \sigma(x, j) \cdot (s_a^k - s_b^k). \tag{14}$$

Hence, $|\sigma(s_a^k, j) - \sigma(s_b^k, j)| \leq \|\nabla \sigma(x, j)\|_1 \|s_a^k - s_b^k\|_\infty$ and $\|\sigma(s_a^k) - \sigma(s_b^k)\|_1 \leq \sum_{j=1}^M \|\nabla \sigma(x, j)\|_1 \|s_a^k - s_b^k\|_\infty$. Considering the stationary policy of Eq. (1), we have $\frac{d\sigma(x, j)}{dx(l)} = (1 - \eta)\beta\sigma(x, j)(\mathbf{1}_{\{j=l\}} - \sigma(x, l))$. Therefore:

$$\begin{aligned} \|\nabla \sigma(x, j)\|_1 &= (1 - \eta)\beta\sigma(x, j) \sum_{l=1}^M |(\mathbf{1}_{\{j=l\}} - \sigma(x, l))| \\ &= 2(1 - \eta)\beta\sigma(x, j)(1 - \sigma(x, j)) \leq 2(1 - \eta)\beta\sigma(x, j). \end{aligned} \tag{15}$$

Combining with Eqs. (14) and (15), we attain:

$$\begin{aligned} \|\sigma(s_a^k) - \sigma(s_b^k)\|_1 &\leq \sum_{j=1}^M 2(1 - \eta)\beta\sigma(x, j)\|s_a^k - s_b^k\|_\infty \\ &\leq 2(1 - \eta)\beta\|s_a - s_b\|_\infty. \end{aligned} \tag{16}$$

In line with Eq. (13), we present the final result:

$$|r(f(s_a), j) - r(f(s_b), j)| \leq 4\theta(1 - \eta)\beta\|s_a - s_b\|_\infty. \tag{17}$$

\square

The condition $4\theta(1 - \eta)\beta < 1$ in Theorem 4 is a little stringent for two reasons. First, the reward function $r(f(s_t), j)$ depends on *all* M elements of $f(s_t)$. Second, $r(f(s_t), j)$ is *non-linear* in $f(s_t)$, so calculating the expected reward needs multiple scaling operations. In fact, the reward $r(f(s_t), j)$ of playing arm j is often only impacted by the number of agents who select arm j [Gummadi et al., 2013], such as an agent playing arm j only competes with those making the same choice in a resource competition game. Hence, we assume $r(f(s_t), j)$ merely depends on the j -th element, denoted as $f(s_t, j)$. Besides, we further presume the reward is a linear function in the population profile. With these two assumptions, we recharacterize a relaxed contraction condition.

Corollary 1. *If the reward $r(f(s_t, j), j)$ is a θ -Lipschitz continuous linear function in the j -th element $f(s_t, j)$, then $r(f(s_t, j), j)$ is a $\|\cdot\|_\infty$ -contraction in the state profile s_t under the condition $\frac{\theta(1-\eta)\beta}{2} < 1$, where β, η are from Eq. (1).*

Proof. From Eqs. (1)-(2), the expected population profile, with a slight abuse of notations, is $f(s_t, j) = \mathbb{E}[\sum_{i=1}^N \frac{\mathbb{1}_{\{a_i^t=j\}}}{N}] = \frac{\sum_{i=1}^N \sigma(s_t^i, j)}{N}$. Since $r(f(s_t, j), j)$ is θ -Lipschitz continuous in $f(s_t, j)$, we obtain $|\frac{dr(f(s_t, j), j)}{df(s_t, j)}| \leq \theta$. For two state profiles s_a, s_b , there is a $x \in [s_a, s_b]$ satisfying $r(f(s_a, j), j) - r(f(s_b, j), j) = \nabla r(f(x, j), j) \cdot (s_a - s_b)$ based on mean value theorem. Hence, $|r(f(s_a, j), j) - r(f(s_b, j), j)| \leq \|\nabla r(f(x, j), j)\|_1 \|s_a - s_b\|_\infty$.

Considering $r(f(x, j), j)$ is linear in $f(x, j)$, we have $\frac{dr(f(x, j), j)}{dx^i(l)} = \frac{dr(f(x, j), j)}{df(x, j)} \frac{df(x, j)}{dx^i(l)}, \forall i \in \mathcal{N}, l \in \mathcal{M}$. As $|\frac{dr(f(x, j), j)}{df(x, j)}| \leq \theta$, we only need to handle $\frac{df(x, j)}{dx^i(l)}$. Moreover, $\frac{df(x, j)}{dx^i(l)} = \frac{1}{N} \frac{d\sigma(x^i, j)}{dx^i(l)} = \frac{(1-\eta)\beta\sigma(x^i, j)(\mathbb{1}_{\{j=l\}} - \sigma(x^i, l))}{N}$.

Consequently, we acquire the following result:

$$\begin{aligned} \|\nabla r(f(x, j), j)\|_1 &= \sum_{i=1}^N \sum_{l=1}^M \left| \frac{dr(f(x, j), j)}{dx^i(l)} \right| \\ &\leq \frac{\theta(1-\eta)\beta}{N} \sum_{i=1}^N \sum_{l=1}^M |\sigma(x^i, j)(\mathbb{1}_{\{j=l\}} - \sigma(x^i, l))| \\ &= \frac{\theta(1-\eta)\beta}{N} \sum_{i=1}^N 2\sigma(x^i, j)(1 - \sigma(x^i, j)) \\ &\leq \frac{\theta(1-\eta)\beta}{N} \sum_{i=1}^N 2 \times \frac{1}{4} \\ &= \frac{\theta(1-\eta)\beta}{2}, \end{aligned} \quad (18)$$

where the second inequality is because $y(1-y) \leq \frac{1}{4}, \forall y \in [0, 1]$. Finally, it results in:

$$|r(f(s_a, j), j) - r(f(s_b, j), j)| \leq \frac{\theta(1-\eta)\beta}{2} \|s_a - s_b\|_\infty. \quad (19)$$

□

4 State Change and Model Extension

In this section, we derive the cumulative state change to infer the regret information and extend the mean field model to demonstrate its effectiveness in various scenarios.

4.1 Cumulative State Change

After state updating according to Eq. (3), the state change is $\Delta s_n^i = \gamma_n(w_n^i - s_n^i)$. Let $\sigma(s_n^i) \cdot \Delta s_n^i$ be the inner product between the simplex $\sigma(s_n^i)$ in Eq. (1) and Δs_n^i . Also denote $(\Delta s_n^i)^2$ as the vector with square on each element of Δs_n^i . The following theorem provides the cumulative state change.

Theorem 5. *Denote $\mathbf{1}$ as a M -length vector with each element equal to 1. For agent i and an arbitrary arm j , we have:*

$$\begin{aligned} \beta s_0^i(j) + \sum_{n=0}^K \beta \Delta s_n^i(j) - \ln \left(\sum_{j=1}^M \text{Exp}(\beta s_0^i(j)) \right) &\leq \\ \sum_{n=0}^K \left[\frac{\beta(\sigma(s_n^i) - \frac{\eta}{M} \mathbf{1}) \cdot \Delta s_n^i}{1-\eta} + \frac{(e-2)\beta^2 \sigma(s_n^i) \cdot (\Delta s_n^i)^2}{1-\eta} \right]. \end{aligned} \quad (20)$$

See Appendix B in the technical report [Wang and Jia, 2021] for the proof. Cumulative state change entails the regret of arm playing, i.e., $\max_j \sum_{n=0}^{T-1} (r(f_n, j) - \mathbb{E}[r(f_n, a_n^i)])$. Specifically, if regarding the left-hand side of Eq. (20) as $\max_j r(f_n, j)$, and expanding the right-hand side via $r(f_n, a_n^i)$, one can obtain a theoretically loose bound of a scaled regret, which serves as a *tradeoff between regret minimization and system stability*. We will show in the evaluations that the stationary policy actually has a small empirical regret.

Remark: The bandit game will not converge to an equilibrium when applying the traditional MAB algorithms which may have tight-bounded regret, like UCB [Auer *et al.*, 2002a] and EXP3 [Auer *et al.*, 2002b]. This is because they need to model an agent state as the cumulative reward. Hence, the state will consistently “increase” since realized rewards are positive, so that the system is unstable. This is also the underlying reason why works [Gummadi *et al.*, 2013; Maghsudi and Hossain, 2017; Zhao and Liu, 2017] have to assume a state regeneration to achieve system stability.

4.2 Heterogeneous Learning Parameter

So far, we have analyzed the bandit game with homogeneous learning parameters β, η in stationary policy Eq. (1). Next, we explore the heterogeneous β, η for different agents.

For agent i , parameter β^i will keep unchanged, while η_n^i is diminishing and satisfies $\lim_{n \rightarrow \infty} \eta_n^i = 0$ to give less weight to the random choice when the reward information is accurately learned. Therefore, the stationary policy changes to $\sigma(s_n^i, j) = (1 - \eta_n^i) \frac{\text{Exp}(\beta^i s_n^i(j))}{\sum_{k=1}^M \text{Exp}(\beta^i s_n^i(k))} + \frac{\eta_n^i}{M}$. Except the playing policy, the updating rule is the same, as described in Eq. (3). Following a similar analysis in the homogeneous situation, it can be asserted that the results in Theorems 1-3 and Lemma 1 still hold. As for the contraction mapping conditions in Theorem 4 and Corollary 1, they turn out to be $4\theta\beta_{\max} < 1$ and $\frac{\theta\beta_{\max}}{2} < 1$, respectively, where $\beta_{\max} = \max\{\beta^i | i \in \mathcal{N}\}$. The cumulative state change in Theorem 5 is specified by replacing β, η with β^i, η_n^i .

4.3 Overlapping Arms

Another extension is about overlapping arms. Mathematically, there are M arms, and each agent i pulls an arm from a subset $\mathcal{M}^i, \mathcal{M}^i \subseteq \mathcal{M}$ with $M^i = |\mathcal{M}^i|$. Overlapping means there exist i, k such that $\mathcal{M}^i \cap \mathcal{M}^k \neq \emptyset$. The conclusion is: previous results still hold after we make several adjustments. The proofs are the same, thus we skip them to save space.

One adjustment is to choose arms from \mathcal{M}^i for agent i , i.e., $\sigma(s_n^i, j) = (1 - \eta) \frac{\text{Exp}(\beta s_n^i(j))}{\sum_{k \in \mathcal{M}^i} \text{Exp}(\beta s_n^i(k))} + \frac{\eta}{M^i}$. The state s_n^i is now a M^i -length vector, and the state profile s_n is in $[0, 1]^{\sum_{i \in \mathcal{N}} M^i}$. Other parameters, especially the cardinality of sets, are adapted accordingly. Similar to Section 3, we could obtain the existence and uniqueness of MFE.

5 Performance Evaluation

In this section, we carry out the evaluations where results are smoothed via LOWESS in Python for better exhibition.




Figure 1: Contraction mapping: state evolution.




Figure 2: Contraction mapping: population profile evolution.

5.1 Existence and Uniqueness of MFE

Reward function. We consider that agents will compete for different resources (arms), so the reward $r(f_n, j)$ is a non-linear decreasing function in $f_n(j)$ [Gummadi *et al.*, 2013]:

$$r(f_n, j) = \frac{1}{1 + \theta(j)f_n(j)}, \quad (21)$$

where $\theta(j) \in [0.8\theta, \theta]$, $\forall j \in \mathcal{M}$. Hence, $r(f_n, j) \in [0, 1]$ is θ -Lipschitz continuous. Set $\gamma_n = 1/(n+1)$ in Eq. (5).

Contraction mapping. Let (θ, β, η) be $(0.5, 0.5, 0.2)$, respectively, that is the contraction condition $4\theta(1-\eta)\beta < 1$ holds. Given the number of arms $M = 4$, we run the bandit game for four times and display the state evolution of arm 2 in Figure 1. We can see that the state will converge to a unique MFE for different number of agents N .

Figure 2 shows the population profile of arm 2, which is unique and tends to be stable as N increases. Because $\mathbb{E}[\mathbb{1}_{\{\bar{a}^i=j\}}] = \sigma(\bar{s}^i, j)$ and $\text{var}[\mathbb{1}_{\{\bar{a}^i=j\}}] = \sigma(\bar{s}^i, j)(1 - \sigma(\bar{s}^i, j)) \leq \frac{1}{4}$, we have $\text{var}[f(\bar{s}, j)] = \frac{1}{N^2} \sum_{i=1}^N \text{var}[\mathbb{1}_{\{\bar{a}^i=j\}}] \leq \frac{1}{4N}$. The Chebyshev's inequality states that $\Pr(|f(\bar{s}, j) - \mathbb{E}[f(\bar{s}, j)]| \geq \epsilon) \leq \frac{1}{4N\epsilon^2}$, i.e., if N increases, $f(\bar{s}, j)$ will be more stable around $\mathbb{E}[f(\bar{s}, j)]$.

Non-contraction mapping. Let θ, η, M stay the same, while β changes to 30 so the contraction mapping $4\theta(1-\eta)\beta < 1$ is violated. We run the bandit game for four times and depict the state evolution in Figure 3, which shows that the state converges to multiple distinct MFEs. The population profile is displayed in Figure 4, where the fluctuation around $\mathbb{E}[f(\bar{s}, j)]$ is also impaired for larger N . Due to multiple MFEs, the population profile has various steady values.

5.2 Empirical Regret

For the general reward, we compute the regret when contraction mapping holds, i.e., $(\theta, \beta, \eta) = (0.5, 0.5, 0.2)$. Furthermore, we implement a linear reward: $r(f_n(j), j) = 1 -$




Figure 3: Non-contraction mapping: state evolution.




Figure 4: Non-contraction mapping: population profile evolution.

$\theta(j)f_n(j)$ where $\theta(j) \in [0.8\theta, \theta]$, $\forall j \in \mathcal{M}$ and $(\theta, \beta, \eta) = (1, 2, 0.2)$, so the reward is a contraction from Corollary 1. We run the evaluation for six times with each operating for $T = 2000$ time slots, and show the average regret and cumulative rewards in Table 1. The regret bound for EXP3 is $O(\sqrt{T})$, and here $\sqrt{T} = 44.721$. We can see that both regrets are much smaller than \sqrt{T} for each N . Besides, the regret of general reward is less than that of linear reward, as changes in $f_n(j)$ have smaller impact when appearing in the denominator of Eq. (21). Moreover, the regrets decrease as N grows large for both cases due to a more stable population profile. In summary, the stationary policy has a tight empirical regret.

6 Conclusion

We propose a mean field model to study a large-population bandit game with a continuous reward. Concretely, we characterize the existence and uniqueness of MFE by showing the state evolution is upper semi-continuous, and deriving contraction mapping conditions based on the stochastic approximation, respectively. Extensive evaluations are performed to validate our mean field analysis and tight empirical regret.

Acknowledgments

This work was supported by NSF China under Grant 61902358, Zhejiang Provincial Natural Science Foundation of China under Grant LQ19F020007.

| Reward | Term | $N = 100$ | $N = 200$ |
|---------|---------|-----------|-----------|
| General | Regret | 13.758 | 7.787 |
| | Rewards | 1796.961 | 1791.061 |
| Linear | Regret | 19.023 | 17.932 |
| | Rewards | 1554.948 | 1560.613 |

Table 1: Empirical regret

References

- [Adlakha *et al.*, 2015] Sachin Adlakha, Ramesh Johari, and Gabriel Y Weintraub. Equilibria of dynamic games with many players: Existence, approximation, and market structure. *Journal of Economic Theory*, 156:269–316, 2015.
- [Auer *et al.*, 2002a] Peter Auer, Nicolo Cesa Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2-3):235–256, 2002.
- [Auer *et al.*, 2002b] Peter Auer, Nicolo Cesa-Bianchi, Yoav Freund, and Robert E Schapire. The nonstochastic multiarmed bandit problem. *SIAM journal on computing*, 32(1):48–77, 2002.
- [Benaïm, 1999] Michel Benaïm. Dynamics of stochastic approximation algorithms. In *Seminaire de probabilités XXXIII*, pages 1–68. Springer, 1999.
- [Bubeck *et al.*, 2012] Sébastien Bubeck, Nicolò Cesa-Bianchi, et al. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends in Machine Learning*, 5(1):1–122, 2012.
- [Cardaliaguet and Hadikhanloo, 2017] Pierre Cardaliaguet and Saeed Hadikhanloo. Learning in mean field games: the fictitious play. *ESAIM: Control, Optimisation and Calculus of Variations*, 23(2):569–591, 2017.
- [Cohen *et al.*, 2017] Johanne Cohen, Amélie Hélieu, and Panayotis Mertikopoulos. Learning with bandit feedback in potential games. In *Proceedings of the 31th International Conference on Neural Information Processing Systems*, 2017.
- [Cominetti *et al.*, 2010] Roberto Cominetti, Emerson Melo, and Sylvain Sorin. A payoff-based learning procedure and its application to traffic games. *Games and Economic Behavior*, 70(1):71–83, 2010.
- [Even-Dar and Mansour, 2003] Eyal Even-Dar and Yishay Mansour. Learning rates for q-learning. *Journal of machine learning Research*, 5(Dec):1–25, 2003.
- [Glicksberg, 1952] Irving L Glicksberg. A further generalization of the kakutani fixed point theorem, with application to nash equilibrium points. *Proceedings of the American Mathematical Society*, 3(1):170–174, 1952.
- [Gummadi *et al.*, 2013] Ramki Gummadi, Ramesh Johari, Sven Schmit, and Jia Yuan Yu. Mean field analysis of multi-armed bandit games. *Available at SSRN 2045842*, 2013.
- [Hanif *et al.*, 2015] Ahmed Farhan Hanif, Hamidou Tembine, Mohamad Assaad, and Djamel Zeghlache. Mean-field games for resource sharing in cloud-based networks. *IEEE/ACM Transactions on Networking*, 24(1):624–637, 2015.
- [Hart and Mas-Colell, 2000] Sergiu Hart and Andreu Mas-Colell. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68(5):1127–1150, 2000.
- [Huang *et al.*, 2006] Minyi Huang, Roland P Malhamé, Peter E Caines, et al. Large population stochastic dynamic games: closed-loop mckean-vlasov systems and the nash certainty equivalence principle. *Communications in Information & Systems*, 6(3):221–252, 2006.
- [Iyer *et al.*, 2014] Krishnamurthy Iyer, Ramesh Johari, and Mukund Sundararajan. Mean field equilibria of dynamic auctions with learning. *Management Science*, 60(12):2949–2970, 2014.
- [Lasry and Lions, 2007] Jean-Michel Lasry and Pierre-Louis Lions. Mean field games. *Japanese journal of mathematics*, 2(1):229–260, 2007.
- [Littman, 1994] Michael L Littman. Markov games as a framework for multi-agent reinforcement learning. In *Machine learning proceedings 1994*, pages 157–163. Elsevier, 1994.
- [Maghsudi and Hossain, 2017] Setareh Maghsudi and Ekram Hossain. Distributed user association in energy harvesting dense small cell networks: A mean-field multi-armed bandit approach. *IEEE Access*, 5:3513–3523, 2017.
- [Panait and Luke, 2005] Liviu Panait and Sean Luke. Co-operative multi-agent learning: The state of the art. *Autonomous agents and multi-agent systems*, 11(3):387–434, 2005.
- [Wang and Jia, 2021] Xiong Wang and Riheng Jia. Mean field equilibrium in multi-armed bandit game with continuous reward. *arXiv preprint arXiv:2105.00767*, 2021.
- [Yang *et al.*, 2018a] Pu Yang, Krishnamurthy Iyer, and Peter Frazier. Mean field equilibria for resource competition in spatial settings. *Stochastic Systems*, 8(4):307–334, 2018.
- [Yang *et al.*, 2018b] Yaodong Yang, Rui Luo, Minne Li, Ming Zhou, Weinan Zhang, and Jun Wang. Mean field multi-agent reinforcement learning. In *International Conference on Machine Learning*, pages 5571–5580, 2018.
- [Ying, 2016] Lei Ying. On the approximation error of mean-field models. *ACM SIGMETRICS Performance Evaluation Review*, 44(1):285–297, 2016.
- [Zhao and Liu, 2017] Zibo Zhao and Andrew L Liu. Intelligent demand response for electricity consumers: A multi-armed bandit game approach. In *2017 19th International Conference on Intelligent System Application to Power Systems (ISAP)*, pages 1–6. IEEE, 2017.