

A General Framework for Learning Mean-Field Games

Xin Guo*

University of California, Berkeley, IEOR, xinguo@berkeley.edu
Amazon.com, xnguo@amazon.com

Anran Hu

University of California, Berkeley, IEOR, anran.hu@berkeley.edu

Renyuan Xu

University of Southern California, Industrial Systems and Engineering, renyuanx@usc.edu
University of Oxford, Mathematical Institute, xur@maths.ox.ac.uk

Junzi Zhang*

Amazon.com, junziz@amazon.com

This paper presents a general mean-field game (GMFG) framework for simultaneous learning and decision-making in stochastic games with a large population. It first establishes the existence of a unique Nash Equilibrium to this GMFG, and demonstrates that naively combining reinforcement learning with the fixed-point approach in classical MFGs yields unstable algorithms. It then proposes value-based and policy-based reinforcement learning algorithms (GMF-V and GMF-P, respectively) with smoothed policies, with analysis of their convergence properties and computational complexities. Experiments on an equilibrium product pricing problem demonstrate that GMF-V-Q and GMF-P-TRPO, two specific instantiations of GMF-V and GMF-P, respectively, with Q-learning and TRPO, are both efficient and robust in the GMFG setting. Moreover, their performance is superior in convergence speed, accuracy, and stability when compared with existing algorithms for multi-agent reinforcement learning in the N -player setting.

1. Introduction

Motivating example. This paper is motivated by the following Ad auction problem for an advertiser. An Ad auction is a stochastic game on an Ad exchange platform among a large number of players, the advertisers. In between the time a web user requests a page and the time the page is displayed, usually within a millisecond, a Vickrey-type of second-best-price auction is run to incentivize interested advertisers to bid for an Ad slot to display advertisement. Each advertiser has limited information before each bid: first, her own valuation for a slot depends on some random conversion of clicks for the item; secondly, she, should she win the bid, only knows the reward after the user's activities on the website are finished. In addition, she has a budget constraint in this repeated auction.

The question is, how should she bid in this online sequential repeated game when there is a large population of bidders competing on the Ad platform, with random conversions of clicks and rewards?

Besides Ad auctions, there are many other real-world problems involving a large number of players and uncertain systems. Examples include massive multi-player online role-playing games Jeong et al. [48], high frequency tradings Lehalle and Mouzouni [56], and the sharing economy Hamari et al. [41].

* Work done prior to joining or outside of Amazon.

Our work. Motivated by these problems, we consider a general framework of simultaneous learning and decision-making in stochastic games with a large population. We formulate a general mean-field-game (GMFG) with incorporation of action distributions and (randomized) relaxed policies. This general framework can also be viewed as a generalized version of MFGs of extended McKean-Vlasov type Acciaio et al. [1], which is a different paradigm from the classical MFG. It is also beyond the scope of the existing reinforcement learning (RL) framework for Markov decision processes (MDP), as MDP is technically equivalent to a single player stochastic game.

On the theory front, this general framework differs from the existing MFGs. We establish under appropriate technical conditions the existence and uniqueness of the Nash equilibrium (NE) to this GMFG. On the computational front, we show that naively combining reinforcement learning with the three-step fixed-point approach in classical MFGs yields unstable algorithms. We then propose both value based and policy based reinforcement learning algorithms with smoothed policies (GMF-V and GMF-P, respectively), establish the convergence property and analyze the computational complexity (see Section 7 for all proof details). Finally, we apply GMF-V-Q and GMF-P-TRPO, which are two specific instantiations of GMF-V and GMF-P, respectively, with Q-learning and TRPO, to an equilibrium product pricing problem¹. Both algorithms have demonstrated to be efficient and robust in the GMFG setting. Their performance is superior in terms of convergence speed, accuracy and stability, when compared with existing algorithms for multi-agent reinforcement learning in the N -player setting. Note that an earlier and preliminary version Guo et al. [38] has been published in NeurIPS. Nevertheless, the conference version focuses only on GMF-V-Q, whereas this paper provides a new meta framework for learning mean-field-game which combines (1) the three-step fixed point approach, (2) the smoothing techniques, and (3) the single-agent algorithms with sample complexity guarantees in the sub-routine. This general framework incorporates both value-based algorithms and policy-based algorithms. In addition, the policy-based RL algorithm (GMF-P-TRPO) in this paper is the first globally convergent policy-based algorithm for solving mean-field-games. Numerical results show that it achieves similar performance as the Q-learning based algorithm (GMF-V-Q) in Guo et al. [38].

Related works. On learning large population games with mean-field approximations, Yang et al. [87] focuses on inverse reinforcement learning for MFGs without decision making, with its extension in Chen et al. [17] for agent-level inference; Yang et al. [88] studies an MARL problem with a first-order mean-field approximation term modeling the interaction between one player and all the other finite players, which has been generalized to the setting with partially observable states in Subramanian et al. [76]; and Kizilkale and Caines [51] and Yin et al. [89] consider model-based adaptive learning for MFGs in specific models (*e.g.*, linear-quadratic and oscillator games). More recently, Mguni et al. [60] studies the local convergence of actor-critic algorithms on finite time horizon MFGs, and Subramanian and Mahajan [74] proposes a policy-gradient based algorithm and analyzes the so-called local NE for reinforcement learning in infinite time horizon MFGs. For learning large population games without mean-field approximation, see Hernandez-Leal et al. [42], Kapoor [50] and the references therein. In the specific topic of learning auctions with a large number of advertisers, Cai et al. [13] and Jin et al. [49] explore reinforcement learning techniques to search for social optimal solutions with real-word data, and Iyer et al. [47] uses MFGs to model the auction system with unknown conversion of clicks within a Bayesian framework.

However, none of these works consider the problem of simultaneous learning and decision-making in a general MFG framework. Neither do they establish the existence and uniqueness of the (global) NE, nor do they present model-free learning algorithms with complexity analysis and convergence to the NE. Note that in principle, global results are harder to obtain compared to local results.

¹ The numerical experiments on the application of GMF-V-Q to the motivating Ad auction problem can be found in the conference version of our paper Guo et al. [38].

Following the conference version Guo et al. [38] of the current paper, various efforts have been made to extend our reinforcement learning work in Guo et al. [38] to more general MFG settings. These include linear-quadratic MFGs in both discrete-time setting Fu et al. [25], uz Zaman et al. [78, 79] and in continuous-time setting Guo et al. [39], Wang et al. [84], Delarue and Vasileiadis [21], MFGs with general continuous state and/or action spaces Anahtarci et al. [3], entropy regularized MFGs in discrete time Anahtarci et al. [4], Xie et al. [85, 86], Cui and Koepl [18] and in continuous time Guo et al. [39], and non-stationary MFGs Mishra et al. [62]. In particular, Cui and Koepl [18] interprets the softmax smoothing technique proposed in Guo et al. [38] from a smoothed equilibrium perspective. In addition, different frameworks based on monotonicity assumptions (instead of the contractivity assumption in Guo et al. [38]) have also been proposed, and fictitious play algorithms with policy and mean-field averaging Elie et al. [23], Perrin et al. [69] and online mirror descent algorithms Perolat et al. [66] have been proposed to solve MFGs under such assumptions. There are also some recent extensions to reinforcement learning of MFGs with strategic complementarity Lee et al. [55] and multiple agent types Ghosh and Aggarwal [32], Subramanian et al. [75]. These algorithms for reinforcement learning of MFGs have also been applied in economics Angiuli et al. [7], in finance de Luca et al. [20], in animal behavior simulation Perrin et al. [68], and in concave utility reinforcement learning Geist et al. [29]. In the meantime, the idea of simultaneous learning and decision making with mean-field interaction has been used for analyzing collaborative games with social optimal solution Carmona et al. [15, 16], Gu et al. [35], Luo et al. [58], Wang et al. [83], Pasztor et al. [65], Gagrani et al. [27], Cui et al. [19], Angiuli et al. [6].

Notations. Let $(\mathcal{X}, d_{\mathcal{X}})$ be a metric space and \mathcal{X} is equipped with the Borel σ -field $\mathcal{B}(\mathcal{X})$, meaning the σ -field generated by the open sets of \mathcal{X} . Denote $\mathcal{P}(\mathcal{X})$ for the set of (Borel) probability measures on \mathcal{X} . \mathcal{W}_p denotes the Wasserstein distance of order p such that

$$W_p(\mu, \mu') = \inf \left\{ \left(\int_{\mathcal{X} \times \mathcal{X}} d_{\mathcal{X}}^p(x, x') \nu(dx, dx') \right)^{1/p} : \nu \in \mathcal{P}(\mathcal{X} \times \mathcal{X}) \text{ with marginals } \mu, \mu' \in \mathcal{P}(\mathcal{X}) \right\}.$$

$\mathcal{P}(\mathcal{X})$ is always equipped with $W_1(\mu, \mu')$. The Borel σ -field of $\mathcal{P}(\mathcal{X})$ is the σ -field induced by the evaluation $\mathcal{P}(\mathcal{X}) \ni \mu \mapsto \mu(C)$ for any Borel set $C \subset \mathcal{X}$. Note that the Borel σ -field of $\mathcal{P}(\mathcal{X})$ is generated by W_1 . (See e.g. Villani [81] and Lacker [52]).

Given two measurable spaces $(\mathcal{Y}, \mathcal{B}(\mathcal{Y}))$ and $(\mathcal{X}, \mathcal{B}(\mathcal{X}))$, we say a measure-valued function $f : \mathcal{Y} \rightarrow \mathcal{P}(\mathcal{X})$ is measurable if $\Lambda_C \circ f : \mathcal{Y} \rightarrow [0, 1]$ is measurable for any $C \in \mathcal{B}(\mathcal{X})$, where $\Lambda_C : \mathcal{P}(\mathcal{X}) \ni \mu \mapsto \mu(C) \in [0, 1]$.

2. Framework of General MFG (GMFG)

2.1. Background: classical N -player Markovian game and MFG Let us first recall the classical N -player game. There are N players in a game. At each step t , the state of player i ($= 1, 2, \dots, N$) is $s_t^i \in \mathcal{S} \subseteq \mathbb{R}^d$ and she takes an action $a_t^i \in \mathcal{A} \subseteq \mathbb{R}^p$. Here d, p are positive integers. The state space $(\mathcal{S}, d_{\mathcal{S}})$ and the action space $(\mathcal{A}, d_{\mathcal{A}})$ are two compact metric spaces, including the case of \mathcal{S} and \mathcal{A} being finite. Given the current state profile of N -players $\mathbf{s}_t = (s_t^1, \dots, s_t^N) \in \mathcal{S}^N$ and the action a_t^i , player i will receive a reward $r^i(\mathbf{s}_t, a_t^i)$ sampled from a distribution $R^i(\mathbf{s}_t, a_t^i)$ and her state will change to s_{t+1}^i according to a transition probability function $P^i(\mathbf{s}_t, a_t^i)$. In particular, the probability transition $P^i : \mathcal{S}^N \times \mathcal{A} \rightarrow \mathcal{P}(\mathcal{S})$ and the distribution of the reward function $R^i : \mathcal{S}^N \times \mathcal{A} \rightarrow \mathcal{P}([0, R_{\max}])$ are both measurable functions with some constant $R_{\max} > 0$.

A Markovian game further restricts the admissible policy/control for player i to be of the form $a_t^i \sim \pi_t^i(\mathbf{s}_t)$ with π_t^i measurable. That is, $\pi_t^i : \mathcal{S}^N \rightarrow \mathcal{P}(\mathcal{A})$ maps each state profile $\mathbf{s} \in \mathcal{S}^N$ to a randomized action. The accumulated reward (a.k.a. the value function) for player i , given the initial

state profile \mathbf{s} and the policy profile sequence $\boldsymbol{\pi} := \{\boldsymbol{\pi}_t\}_{t=0}^{\infty}$ with $\boldsymbol{\pi}_t = (\pi_t^1, \dots, \pi_t^N)$, is then defined as

$$V^i(\mathbf{s}, \boldsymbol{\pi}) := \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t r^i(\mathbf{s}_t, a_t^i) \mid \mathbf{s}_0 = \mathbf{s} \right], \quad (1)$$

where $\gamma \in (0, 1)$ is the discount factor, $a_t^i \sim \pi_t^i(\mathbf{s}^t)$, and $s_{t+1}^i \sim P^i(\mathbf{s}_t, a_t^i)$. The goal of each player is to maximize her value function over all admissible policy sequences such that (1) is finite.

In general, this type of stochastic N -player game is notoriously hard to analyze, especially when N is large Papadimitriou and Roughgarden [64]. Mean field game (MFG), pioneered by Huang et al. [46] and Lasry and Lions [54] in the continuous settings and later developed in Benaim and Le Boudec [10], Gomes et al. [34], Huang and Ma [45], López [57], Saldi et al. [71] for discrete settings, provides an ingenious and tractable aggregation approach to approximate the otherwise challenging N -player stochastic games. The basic idea for an MFG goes as follows. Assume all players are identical, indistinguishable and interchangeable, when $N \rightarrow \infty$, one can view the limit of other players' states $\mathbf{s}_t^{-i} = (s_t^1, \dots, s_t^{i-1}, s_t^{i+1}, \dots, s_t^N)$ as a population state distribution μ_t with $\mu_t(s) := \lim_{N \rightarrow \infty} \frac{\sum_{j=1, j \neq i}^N \mathbf{1}_{s_t^j = s}}{N}$.² Due to the homogeneity of the players, one can then focus on a single (representative) player. At time t , after the representative player chooses her action a_t according to some policy π_t , she will receive reward $r(s_t, a_t, \mu_t)$ and her state will evolve under a *controlled stochastic dynamics* of a mean-field type $P(\cdot | s_t, a_t, \mu_t)$. Here the policy π_t depends on both the current state s_t and the current population state distribution μ_t such that $\pi_t : \mathcal{S} \times \mathcal{P}(\mathcal{S}) \rightarrow \mathcal{P}(\mathcal{A})$. Then, in mean-field limit, one may consider instead the following optimization problem,

$$\begin{aligned} & \text{maximize}_{\boldsymbol{\pi}} V(s, \boldsymbol{\pi}, \boldsymbol{\mu}) := \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t, \mu_t) \mid s_0 = s \right] \\ & \text{subject to } s_{t+1} \sim P(s_t, a_t, \mu_t), \quad a_t \sim \pi_t(s_t, \mu_t), \end{aligned}$$

where $\boldsymbol{\pi} := \{\pi_t\}_{t=0}^{\infty}$ denotes the policy sequence and $\boldsymbol{\mu} := \{\mu_t\}_{t=0}^{\infty}$ the distribution flow.

2.2. General MFG (GMFG) In the classical MFG setting, the reward and the dynamic for each player are known. They depend only on the state of the player s_t , the action of this particular player a_t , and the population state distribution μ_t . In contrast, in the motivating auction example, the reward and the dynamic are unknown; they rely on the actions of *all* players, as well as on s_t and μ_t .

We therefore define the following general MFG (GMFG) framework. At time t , after the representative player chooses her action a_t according to some measurable policy $\pi : \mathcal{S} \times \mathcal{P}(\mathcal{S}) \rightarrow \mathcal{P}(\mathcal{A})$, she will receive a (possibly random) reward $r(s_t, a_t, \mathcal{L}_t)$ sampled from distribution $R(s_t, a_t, \mathcal{L}_t)$ and her state will evolve according to $P(\cdot | s_t, a_t, \mathcal{L}_t)$, with $\mathcal{L}_t = \mathbb{P}_{s_t, a_t} m \mathcal{P}(\mathcal{S} \times \mathcal{A})$ the joint distribution of the state and the action, *i.e.*, the population state-action pair. This joint distribution \mathcal{L}_t has marginal distributions α_t for the population action and μ_t for the population state. Note the inclusion of α_t allows the reward and the dynamic to depend on all players' actions. Here $P : \mathcal{S} \times \mathcal{A} \times \mathcal{P}(\mathcal{S} \times \mathcal{A}) \rightarrow \mathcal{P}(\mathcal{S})$ and $R : \mathcal{S} \times \mathcal{A} \times \mathcal{P}(\mathcal{S} \times \mathcal{A}) \rightarrow \mathcal{P}([0, R_{\max}])$ are measurable functions with some constant $R_{\max} > 0$. The objective of the player is to solve the following control problem:

$$\begin{aligned} & \text{maximize}_{\boldsymbol{\pi}} V(s, \boldsymbol{\pi}, \boldsymbol{\mathcal{L}}) := \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t, \mathcal{L}_t) \mid s_0 = s \right] \\ & \text{subject to } s_{t+1} \sim P(s_t, a_t, \mathcal{L}_t), \quad a_t \sim \pi_t(s_t, \mu_t). \end{aligned} \quad (\text{GMFG})$$

² Here the indicator function $\mathbf{1}_{s_t^j = s} = 1$ if $s_t^j = s$ and 0 otherwise.

Here the expectation in the objective function is always taken for all randomness in the system. In addition, $\mathcal{L} := \{\mathcal{L}_t\}_{t=0}^\infty$ and may be time dependent. That is, an infinite-time horizon MFG may have time-dependent NE solutions due to the mean information process in the MFG. This is fundamentally different from the theory of MDP where the optimal control, if exists uniquely, would be time independent in an infinite time horizon setting.

In this paper, we will analyze the existence of NE to GMFG. For ease of exposition, we will first focus on stationary NEs. Accordingly, for notational brevity, we abbreviate $\boldsymbol{\pi} = \{\pi\}_{t=0}^\infty$ and $\mathcal{L} = \{\mathcal{L}\}_{t=0}^\infty$ as π and \mathcal{L} , respectively. We will show in the end how this stationary constraint can be relaxed (*cf.* Section 9).

DEFINITION 1 (STATIONARY NE FOR GMFGs). In (GMFG), a player-population profile (π^*, \mathcal{L}^*) is called a stationary NE if

1. (Single player side) For any policy π and any initial state $s \in \mathcal{S}$,

$$V(s, \pi^*, \mathcal{L}^*) \geq V(s, \pi, \mathcal{L}^*). \quad (2)$$

2. (Population side) $\mathbb{P}_{s_t, a_t} = \mathcal{L}^*$ for all $t \geq 0$, where $\{s_t, a_t\}_{t=0}^\infty$ is the dynamics under the policy π^* starting from $s_0 \sim \mu^*$, with $a_t \sim \pi^*(s_t, \mu^*)$, $s_{t+1} \sim P(\cdot | s_t, a_t, \mathcal{L}^*)$, and μ^* being the population state marginal of \mathcal{L}^* .

The single player side condition captures the optimality of π^* , when the population side is fixed. The population side condition ensures the ‘‘consistency’’ of the solution: it guarantees that the state and action distribution flow of the single player does match the population state and action sequence $\mathcal{L}^* := \{\mathcal{L}^*\}_{t=0}^\infty$.

2.3. Examples of GMFG

Here we provide three examples under the framework of GMFG.

A toy example. Take a two-state dynamic system with two choices of controls. The state space $\mathcal{S} = \{0, 1\}$, the action space $\mathcal{A} = \{L, R\}$. Here the action L means to move left and R means to move right. The dynamic of the representative agent in the mean-field system $\{s_t\}_{t \geq 1}$ goes as follows: if the agent is in state s_t and she takes action $a_t = L$ at time t , then $s_{t+1} = 0$; if she takes action $a_t = R$, then $s_{t+1} = 1$. At the end of each round, the agent will receive a reward $-W_2(\mu_t, B) - W_2(\beta_t(s_t, \cdot), B)$, which depends on all agents, where W_2 is the ℓ_2 -Wasserstein distance. Here $\mu_t(\cdot)$ denotes the state distribution of the mean-field population at time t , $\beta_t(s, \cdot) := \mathcal{L}_t(s, \cdot) / \mu_t(s)$ denotes the action distribution of the population in state s ($s = 0, 1$) at time t (set $\beta_t(s, \cdot) := (0.5, 0.5)$ when $\mu_t(s) = 0$), and B is a given Bernoulli distribution with parameter p ($0 < p < 1$).

As a demonstrating example, here we provide the calculation for one stationary NE solution. Note that $-W_2(\mu, B) \leq 0$ for any distribution μ over \mathcal{S} . Similarly, $-W_2(\alpha, B) \leq 0$ for any distribution α over \mathcal{A} . Hence for each policy $\boldsymbol{\pi}$, given population distribution flow $\mathcal{L} = \{\mathcal{L}_t\}_{t=1}^\infty$,

$$V(0, \boldsymbol{\pi}, \mathcal{L}) = - \sum_{t=1}^{\infty} \gamma^t \mathbb{E}[W_2(\mu_t, B) + W_2(\beta_t(s_t, \cdot), B) | s_0 = 0] \leq 0, \quad (3)$$

and

$$V(1, \boldsymbol{\pi}, \mathcal{L}) = - \sum_{t=1}^{\infty} \gamma^t \mathbb{E}[W_2(\mu_t, B) + W_2(\beta_t(s_t, \cdot), B) | s_0 = 1] \leq 0. \quad (4)$$

It is easy to check that $\mu^* = (p, 1-p)$ and $\pi^*(s, \mu^*) = (p, 1-p)$ ($s = 0, 1$). is a pair of stationary mean-field solution. And \mathcal{L}^* is defined with $\mathcal{L}^*(s, a) = \mu^*(s)\pi^*(a|s, \mu^*)$ for any $s \in \mathcal{S}$, $a \in \mathcal{A}$, accordingly, where $\pi(a|s, \mu)$ is defined as the probability of taking action a following the action distribution $\pi(s, \mu)$. In this case, the corresponding optimal value function is defined as

$$V(0, \pi^*, \mathcal{L}^*) = V(1, \pi^*, \mathcal{L}^*) = 0,$$

which reaches the upper bound in (3) and (4).

Repeated auction. Take a representative advertiser in the auction aforementioned in the motivating example in Section 1. Denote $s_t \in \{0, 1, 2, \dots, s_{\max}\}$ as the budget of this player at time t , where $s_{\max} \in \mathbb{N}^+$ is the maximum budget allowed on the Ad exchange with a unit bidding price. Denote $a_t \in \{0, 1, 2, \dots, a_{\max}\}$ as the bid price submitted by this player, where a_{\max} is the maximum bid set by the bidder, and α_t as the bidding/(action) distribution of the population. At time t , all advertisers are randomly divided into different groups and each group of advertisers competes for one slot to display their ads. Assuming that there are M advertisers in each group, then the representative advertiser competes with $M - 1$ other representative players whose bidding prices are independently sampled from α_t . Let w_t^M denote whether the representative player wins the bid. Then if she takes action a_t , the probability she will win the bid is $\mathbb{P}(w_t^M = 1) = F_{\alpha_t}(a_t)^{M-1}$, where F_{α_t} is the cumulative distribution function of a random variable $X \sim \alpha_t$.

If this advertiser does not win the bid, her reward $r_t = 0$. If she wins, there are several components in her reward: a_t^M , the second best bid in a Vickrey auction, paid by the winning advertiser; v_t , the conversion of clicks of the slot; and ρ , the rate of penalty for overshooting if the payment a_t^M exceeds her budget s_t . Therefore, at each time t , her reward with bid a_t and budget s_t is

$$r_t = \mathbf{I}_{\{w_t^M=1\}} \left[(v_t - a_t^M) - (1 + \rho) \mathbf{I}_{\{s_t < a_t^M\}} (a_t^M - s_t) \right], \quad (5)$$

where the first term is the profit of winning the auction and the second term is the penalty of overshooting. And the budget dynamics s_t follows,

$$s_{t+1} = \begin{cases} s_t, & w_t^M \neq 1, \\ s_t - a_t^M, & w_t^M = 1 \text{ and } a_t^M \leq s_t, \\ 0, & w_t^M = 1 \text{ and } a_t^M > s_t. \end{cases} \quad (6)$$

That is, if this player does not win the bid, the budget remains the same; if she wins and has sufficient money to pay, her budget will decrease from s_t to $s_t - a_t^M$; however, if she wins but does not have enough money to pay, her budget will be 0 after the payment and there will be a penalty in the reward function.

Notice that both distributions of w_t^M and a_t^M depend on the population distribution \mathcal{L}_t (or more specifically α_t). In fact, the reward function $r(s_t, a_t) = r_t$ and the transition probability $s_{t+1} \sim P(\cdot | s_t, a_t, \mathcal{L}_t)$ specified by (5) and (6) are fully characterized by the probabilities $\mathbb{P}(w_t^M = 1, a_t^M \leq \cdot | s_t, a_t, \mathcal{L}_t)$ and $\mathbb{P}(w_t^M = 0)$ (since $r_t = 0$ and $s_{t+1} = s_t$ whenever $w_t^M = 0$), with

$$\mathbb{P}(w_t^M = 1, a_t^M \leq x | s_t, a_t, \mathcal{L}_t) = F_{\alpha_t}(\min\{x, a_t\})^{M-1}, \quad \mathbb{P}(w_t^M = 0) = 1 - F_{\alpha_t}(a_t)^{M-1}.$$

Clearly the above model fits into the framework of (GMFG), with the following transition probability.

$$\mathbb{P}(s' | s, a, \mathcal{L}) = \begin{cases} F_{\alpha}(a)^{M-1} - F_{\alpha}(\min\{s, a\})^{M-1}, & s' = 0, \\ 1 - F_{\alpha}(a)^{M-1}, & s' = s, \\ F_{\alpha}(\min\{s - s', a\})^{M-1} - F_{\alpha}(\min\{s - s' - 1, a\})^{M-1}, & 0 < s' < s, \end{cases} \quad (7)$$

where α is the action marginal of \mathcal{L} . The reward model can be explicitly written similarly.

In practice, one may modify the dynamics of s_{t+1} with a non-negative random budget fulfillment $\Delta(s_{t+1})$ after the auction clearing such that $\hat{s}_{t+1} = s_{t+1} + \Delta(s_{t+1})$ Andelman and Mansour [5], Gummadi et al. [37].

Experiments of this repeated auction problem can be found in the conference version Guo et al. [38] of this paper, and will not be repeated here.

Equilibrium price. Another example, adapted from Guéant et al. [36, Section 3] is to consider a large number (continuum) of homogeneous firms producing the same product under perfect competition, and the price of the product is determined endogenously by the supply-demand equilibrium Bernstein and Griffin [11]. Each firm, meanwhile, maintains a certain inventory level of the raw materials for production.

Given the homogeneity of the firms, it is sufficient to focus on a representative firm paired with the population distribution. In each period t , the representative firm decides a quantity q_t to consume the raw materials for production and a quantity h_t to replenish the inventory of raw materials. For simplicity, we assume each unit of the raw material is used to produce one unit of the product. Both the new products and ordered raw materials will be available at the end of this given period t . The representative agent makes decision based on her current inventory level of the raw material, denoted as s_t , which evolves according to

$$s_{t+1} = s_t - \min\{q_t, s_t\} + h_t. \quad (8)$$

Note that if the firm overproduces and exceeds her current inventory capacity (i.e., $q_t > s_t$), then the firm will pay a cost for an emergency order of the raw material. Finally, the reward during this period t is given by

$$r_t = (p_t - c_0)q_t - c_1 q_t^2 - c_2 h_t - (c_2 + c_3) \max\{q_t - s_t, 0\} - c_4 s_t. \quad (9)$$

Here p_t is the selling price of the product of all firms; $c_0 > 0$ is the manufacturing cost and labor cost for making one unit of the product; $c_1 > 0$ is the quadratic cost which can be viewed as the transient price impact associated with the production level q_t ; $c_2 > 0$ is the cost of regular orders of the raw materials; $c_3 > 0$ is the additional cost for the emergency order of the raw materials; and finally, $c_4 > 0$ is the inventory cost.

The price p_t is determined according to the supply-demand equilibrium on the market at each moment. On one hand, the normalized demand (per producer) on the market $D(p_t)$ follows (Guéant et al. [36])

$$D(p_t) := d p_t^{-\sigma}, \quad (10)$$

where d denotes some benchmark demand level and σ is the elasticity of demand that can be interpreted as the elasticity of substitution between the given product and any other good. On the other hand, the (average) supply in this market is given by the average production of all firms which follows $\mathbb{E}_{q_t \sim \pi_t}[q_t]$ under some policy π_t . If all firms are restricted to stationary policies (denoted as π), then this leads to a stationary equilibrium price q which satisfies the supply-demand equilibrium:

$$\mathbb{E}_{q \sim \pi}[q] = d p^{-\sigma}. \quad (11)$$

To fit into the theoretical framework proposed in Section 2, we set $\mathcal{S} = \{0, 1, \dots, S\}$ and $\mathcal{A} = \{(q, h) \mid q \in \{0, 1, \dots, Q\} \text{ and } h \in \{0, 1, \dots, H\}\}$ for some positive integers S, Q and H .

3. Solution for GMFGs We now establish the existence and uniqueness of the stationary NE to (GMFG), by generalizing the classical fixed-point approach for MFGs to this GMFG setting. (See Huang et al. [46] and Lasry and Lions [54] for the classical case.) It consists of three steps.

Step A. Fix \mathcal{L} , (GMFG) becomes the classical single-player optimization problem. Indeed, with \mathcal{L} fixed, the population state distribution μ is also fixed, and hence the space of admissible policies is reduced to the single-player case. Solving (GMFG) is now reduced to finding a policy $\pi_{\mathcal{L}}^* \in \Pi := \{\pi \mid \pi : \mathcal{S} \rightarrow \mathcal{P}(\mathcal{A})\}$ to maximize

$$V(s, \pi_{\mathcal{L}}, \mathcal{L}) := \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t, \mathcal{L}) \mid s_0 = s \right],$$

subject to $s_{t+1} \sim P(s_t, a_t, \mathcal{L}), \quad a_t \sim \pi_{\mathcal{L}}(s_t).$

Notice that with \mathcal{L} fixed, one can safely suppress the dependency on μ in the admissible policies.

Now given this fixed \mathcal{L} and the solution $\pi_{\mathcal{L}}^*$ to the above optimization problem, one can define a mapping from the fixed population distribution \mathcal{L} to a chosen optimal randomized policy sequence. That is,

$$\Gamma_1 : \mathcal{P}(\mathcal{S} \times \mathcal{A}) \rightarrow \Pi,$$

such that $\pi_{\mathcal{L}}^* = \Gamma_1(\mathcal{L})$. Note that the optimal policy of an MDP in general may not be unique. To ensure that Γ_1 is a single-valued instead of set-valued mapping, here Γ_1 includes a policy selection component to select a single optimal policy from the set of optimal policies for a given \mathcal{L} , which is guaranteed to exist by Zermelo's Axiom of Choice. For example, when the action space is finite, one can utilize the **argmax-e** operator and set the "maximizing" actions with equal probabilities (see Section 4.1 for the detailed definition). In addition, for non-degenerate linear-quadratic MFGs Fu et al. [26] and general MFGs where the Bellman mappings are strongly concave in actions Anahtarci et al. [3] and the action space is convex in the Euclidean space, the optimal policy $\pi_{\mathcal{L}}^*$ for a given \mathcal{L} is unique under appropriate assumptions. Hence no policy selection is needed in such cases.

Note that this $\pi_{\mathcal{L}}^*$ satisfies the single player side condition in Definition 1 for the population state-action pair \mathcal{L} ,

$$V(s, \pi_{\mathcal{L}}^*, \mathcal{L}) \geq V(s, \pi, \mathcal{L}), \quad (12)$$

for any policy π and any initial state $s \in \mathcal{S}$.

As in the MFG literature Huang et al. [46], a feedback regularity condition is needed for analyzing Step A.

ASSUMPTION 1. *There exists a constant $d_1 \geq 0$, such that for any $\mathcal{L}, \mathcal{L}' \in \mathcal{P}(\mathcal{S} \times \mathcal{A})$,*

$$D(\Gamma_1(\mathcal{L}), \Gamma_1(\mathcal{L}')) \leq d_1 W_1(\mathcal{L}, \mathcal{L}'), \quad (13)$$

where

$$D(\pi, \pi') := \sup_{s \in \mathcal{S}} W_1(\pi(s), \pi'(s)), \quad (14)$$

and W_1 is the ℓ_1 -Wasserstein distance (a.k.a. earth mover distance) between probability measures Gibbs and Su [33], Peyré and Cuturi [70], Villani [80].

Step B. Given $\pi_{\mathcal{L}}^*$ obtained from Step A, update the initial \mathcal{L} to \mathcal{L}' following the controlled dynamics $P(\cdot \mid s_t, a_t, \mathcal{L})$.

Accordingly, for any admissible policy $\pi \in \Pi$ and a joint population state-action pair $\mathcal{L} \in \mathcal{P}(\mathcal{S} \times \mathcal{A})$, define a mapping $\Gamma_2 : \Pi \times \mathcal{P}(\mathcal{S} \times \mathcal{A}) \rightarrow \mathcal{P}(\mathcal{S} \times \mathcal{A})$ as follows:

$$\Gamma_2(\pi, \mathcal{L}) := \hat{\mathcal{L}} = \mathbb{P}_{s_1, a_1}, \quad (15)$$

where $a_1 \sim \pi(s_1)$, $s_1 \sim \mu P(\cdot \mid \cdot, a_0, \mathcal{L})$, $a_0 \sim \pi(s_0)$, $s_0 \sim \mu$, and μ is the population state marginal of \mathcal{L} .

One needs a standard assumption in this step.

ASSUMPTION 2. *There exist constants $d_2, d_3 \geq 0$, such that for any admissible policies π, π_1, π_2 and joint distributions $\mathcal{L}, \mathcal{L}_1, \mathcal{L}_2$,*

$$W_1(\Gamma_2(\pi_1, \mathcal{L}), \Gamma_2(\pi_2, \mathcal{L})) \leq d_2 D(\pi_1, \pi_2), \quad (16)$$

$$W_1(\Gamma_2(\pi, \mathcal{L}_1), \Gamma_2(\pi, \mathcal{L}_2)) \leq d_3 W_1(\mathcal{L}_1, \mathcal{L}_2). \quad (17)$$

Step C. Repeat Step A and Step B until \mathcal{L}' matches \mathcal{L} .

This step is to ensure the population side condition. To ensure the convergence of the combined step one and step two, it suffices if $\Gamma : \mathcal{P}(\mathcal{S} \times \mathcal{A}) \rightarrow \mathcal{P}(\mathcal{S} \times \mathcal{A})$ with $\Gamma(\mathcal{L}) := \Gamma_2(\Gamma_1(\mathcal{L}), \mathcal{L})$ is a contractive mapping under the W_1 distance. Then by the Banach fixed point theorem and the completeness of the related metric spaces (*cf.* Appendix A), there exists a unique stationary NE of the GMFG. That is,

THEOREM 1 (Existence and Uniqueness of stationary GMFG solution). *Given Assumptions 1 and 2, and assume $d_1 d_2 + d_3 < 1$. Then there exists a unique stationary NE to (GMFG).*

[Proof of Theorem 1] First by Definition 1 and the definitions of Γ_i ($i = 1, 2$), (π, \mathcal{L}) is a stationary NE iff $\mathcal{L} = \Gamma(\mathcal{L}) = \Gamma_2(\Gamma_1(\mathcal{L}), \mathcal{L})$ and $\pi = \Gamma_1(\mathcal{L})$, where $\Gamma(\mathcal{L}) = \Gamma_2(\Gamma_1(\mathcal{L}), \mathcal{L})$. This indicates that for any $\mathcal{L}_1, \mathcal{L}_2 \in \mathcal{P}(\mathcal{S} \times \mathcal{A})$,

$$\begin{aligned} W_1(\Gamma(\mathcal{L}_1), \Gamma(\mathcal{L}_2)) &= W_1(\Gamma_2(\Gamma_1(\mathcal{L}_1), \mathcal{L}_1), \Gamma_2(\Gamma_1(\mathcal{L}_2), \mathcal{L}_2)) \\ &\leq W_1(\Gamma_2(\Gamma_1(\mathcal{L}_1), \mathcal{L}_1), \Gamma_2(\Gamma_1(\mathcal{L}_2), \mathcal{L}_1)) + W_1(\Gamma_2(\Gamma_1(\mathcal{L}_2), \mathcal{L}_1), \Gamma_2(\Gamma_1(\mathcal{L}_2), \mathcal{L}_2)) \\ &\leq (d_1 d_2 + d_3) W_1(\mathcal{L}_1, \mathcal{L}_2). \end{aligned} \quad (18)$$

And since $d_1 d_2 + d_3 \in [0, 1)$, by the Banach fixed-point theorem, we conclude that there exists a unique fixed-point of Γ , or equivalently, a unique stationary MFG solution to (GMFG).

REMARK 1 (EXISTENCE AND UNIQUENESS OF THE GMFG SOLUTION). (1) In general, there may multiple optimal policies in Step A under a fixed mean-field information \mathcal{L} . In this case, the candidate fixed point(s) are the fixed point(s) of a set-valued map as described in Lacker [53]. To simplify the analysis, we specify a rule in Step A to select one optimal policy to ensure that Γ is an injection.

(2) In the MFG literature, the uniqueness of the MFG solution can be verified under the small parameter condition Caines et al. [14] or the monotonicity condition Lasry and Lions [54]. Our condition of $d_1 d_2 + d_3 < 1$ extends the small parameter condition in Caines et al. [14] for strict controls to relaxed controls.

(3) Finally, Theorem 1 can be extended to a non-stationary setting, as will be shown in Section 9.

REMARK 2. Assumptions 1 and 2 can be more explicit in specific problem settings.

For instance, when the action space is the Euclidean space or its convex subset, explicit conditions on P and r have been described for the linear-quadratic MFG (LQ-MFG) Fu et al. [26] and later generalized in Anahtarçı et al. [3].

When the action space is finite, the following lemma explicitly characterizes Assumption 2.

LEMMA 1. *Suppose that $\max_{s,a,\mathcal{L},s'} P(s'|s, a, \mathcal{L}) \leq c_1$, and that $P(s'|s, a, \cdot)$ is c_2 -Lipschitz in W_1 , i.e.,*

$$|P(s'|s, a, \mathcal{L}_1) - P(s'|s, a, \mathcal{L}_2)| \leq c_2 W_1(\mathcal{L}_1, \mathcal{L}_2). \quad (19)$$

Then in Assumption 2, d_2 and d_3 can be chosen as

$$d_2 = \frac{2 \text{diam}(\mathcal{S}) \text{diam}(\mathcal{A}) |\mathcal{S}| c_1}{d_{\min}(\mathcal{A})} \quad (20)$$

and $\text{diam}(\mathcal{S}) \text{diam}(\mathcal{A}) |\mathcal{S}| \left(\frac{c_2}{2} + \frac{1}{d_{\min}(\mathcal{S} \times \mathcal{A})} \right)$ respectively. Here $d_{\min}(\mathcal{A}) = \min_{a \neq a' \in \mathcal{A}} \|a - a'\|_2$, which is guaranteed to be positive when \mathcal{A} is finite.

When entropy regularization is introduced into the system (see *e.g.*, Anahtarci et al. [4], Xie et al. [85]), Assumption 1 can be reduced to boundedness and Lipschitz continuity conditions on P and r as in Lemma 1. Moreover, Theorem 1 and all subsequent theoretical results hold whenever the composed mapping Γ is contractive (in W_1), independent of Assumptions 1 or 2. In Section 8.2, we numerically verify that the Γ mapping is contractive for various choices of the model parameters in our tested problems.

4. Naive algorithm and stabilization techniques In this section, we design algorithms for the GMFG. Since the reward and transition distributions are unknown, this is simultaneously learning the system and finding the NE of the game. We will focus on the case with finite state and action spaces, *i.e.*, $|\mathcal{S}|, |\mathcal{A}| < \infty$. We will look for stationary (time independent) NEs. This stationarity property enables developing appropriate stationary reinforcement learning algorithms, suitable for an infinite time horizon game. Instead of knowing the transition probability P and the reward r explicitly, the algorithms we propose only assume access to a simulator oracle, which is described below. This is not restrictive in practice. For instance, in the ad auction example, one may adopt the bid recommendation perspective of the publisher, say Google, Facebook or Amazon, who acts as the auctioneer and owns the Ad slot inventory on its own Ad exchange platform. In this case, a high quality auction simulator is typically built and maintained by a team of the publisher. See also Subramanian and Mahajan [74] for more examples.

Simulator oracle. For any policy $\pi \in \Pi$, given the current state $s \in \mathcal{S}$, for any population distribution \mathcal{L} , one can obtain a *sample* of the next state $s' \sim P(\cdot|s, \pi(s), \mathcal{L})$, a reward $r = r(s, \pi(s), \mathcal{L})$, and the next population distribution $\mathcal{L}' = \mathbb{P}_{s', \pi(s')}$. For brevity, we denote the simulator as $(s', r, \mathcal{L}') = \mathcal{G}(s, \pi, \mathcal{L})$. This simulator oracle can be weakened to fit the N -player setting, see Section 6.

In the following, we begin with a naive algorithm that simply combines the three-step fixed point approach with general RL algorithms, and demonstrate that this algorithm can be unstable (Section 4.1). We then propose some smoothing and projection techniques to resolve the issue (Section 4.2). In Section 5.1 and Section 5.2, we design general value-based and policy-based RL algorithms, and establish the corresponding convergence and complexity results. These two algorithms include most of the RL algorithms in the literature. We then illustrate by two concrete examples based on Q-learning and trust-region policy optimization algorithms.

4.1. Naive algorithm and its issue We follow the three-step fixed-point approach described in Section 3. Notice the fact that with \mathcal{L} fixed, Step A in Section 3 becomes a standard learning problem for an infinite horizon discounted MDP. More specifically, the MDP to be solved is $\mathcal{M}_{\mathcal{L}} = (\mathcal{S}, \mathcal{A}, P_{\mathcal{L}}, r_{\mathcal{L}}, \gamma)$, where $P_{\mathcal{L}}(s'|s, a) = P(s'|s, a, \mathcal{L})$ and $r_{\mathcal{L}}(s, a) = r(s, a, \mathcal{L})$. In general, for an MDP $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, r, \gamma)$, for any policy π one can define its value functions $V_{\mathcal{M}}^{\pi}(s) = \mathbb{E}[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) | s_0 = s]$ and its Q-functions $Q_{\mathcal{M}}^{\pi}(s, a) = \mathbb{E}[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) | s_0 = s, a_0 = a]$, where s_t, a_t is the trajectory under policy π . One can also define the optimal Q-function as the unique solution of the Bellman equation:

$$Q_{\mathcal{M}}^*(s, a) = \mathbb{E}[r(s, a)] + \gamma \max_{a'} \sum_{s' \in \mathcal{S}} P(s'|s, a) Q_{\mathcal{M}}^*(s', a')$$

for all s, a and its optimal value function $V_{\mathcal{M}}^*(s) = \max_a Q_{\mathcal{M}}^*(s, a)$ for all s . We also use the shorthand $V_{\mathcal{L}}^* = V_{\mathcal{M}_{\mathcal{L}}}^*$ and $Q_{\mathcal{L}}^* = Q_{\mathcal{M}_{\mathcal{L}}}^*$ for notational brevity. Whenever the context is clear, we may omit \mathcal{M} , \mathcal{L} and $\mathcal{M}_{\mathcal{L}}$ for notational convenience.

Given the optimal Q-function $Q_{\mathcal{L}}^*$, one can obtain an optimal policy $\pi_{\mathcal{L}}^*$ with $\pi_{\mathcal{L}}^*(s) = \mathbf{argmax-e}(Q_{\mathcal{L}}^*(s, \cdot))$. Here the **argmax-e** operator is defined so that actions with equal maximum Q-values would have equal probabilities to be selected. Hereafter, we specify Γ_1 as a mapping to the

aforementioned choice of the optimal policy, *i.e.*, the s -component $\Gamma_1(\mathcal{L})_s = \mathbf{argmax-e}(Q_{\mathcal{L}}^*(s, \cdot))$ for any $s \in \mathcal{S}$.

The population update in Step B can then be directly obtained from the simulator \mathcal{G} following policy $\pi_{\mathcal{L}}^*$. Combining these two steps leads to the following naive algorithm (Algorithm 1).

Algorithm 1 Naive Reinforcement Learning for GMFGs

- 1: **Input:** Initial population state-action pair \mathcal{L}_0
 - 2: **for** $k = 0, 1, \dots$ **do**
 - 3: Obtain the optimal Q-function $Q_k(s, a) = Q_{\mathcal{L}_k}^*(s, a)$ of an MDP with dynamics $P_{\mathcal{L}_k}(s'|s, a)$ and reward distributions $R_{\mathcal{L}_k}(s, a)$.
 - 4: Compute $\pi_k \in \Pi$ with $\pi_k(s) = \mathbf{argmax-e}(Q_k(s, \cdot))$.
 - 5: Sample $s \sim \mu_k$, where μ_k is the population state marginal of \mathcal{L}_k , and obtain \mathcal{L}_{k+1} from $\mathcal{G}(s, \pi_k, \mathcal{L}_k)$.
 - 6: **end for**
-

Unfortunately, in practice, one cannot obtain the exact optimal Q-function Q_k . In fact, invoking any commonly used RL algorithm with the simulator \mathcal{G} leads to an approximation \hat{Q}_k of the actual Q_k . This approximation error is then magnified by the discontinuous and sensitive **argmax-e**, which eventually leads to an unstable algorithm (see Figure 4 for an example of divergence). To see why **argmax-e** is not continuous, consider the following simple example. Let $x = (1, 1)$, then $\mathbf{argmax-e}(x) = (1/2, 1/2)$. For any $\epsilon > 0$, let $y_\epsilon = (1, 1 - \epsilon)$, then $\mathbf{argmax-e}(y_\epsilon) = (1, 0)$. Hence $\lim_{\epsilon \rightarrow 0} y_\epsilon = x$ but

$$\lim_{\epsilon \rightarrow 0} \mathbf{argmax-e}(y_\epsilon) \neq \mathbf{argmax-e}(x).$$

This instability issue will be addressed by introducing smoothing and projection techniques.

4.2. Restoring stability

Smoothing techniques. To address the instability caused, we replace **argmax-e** with a smooth function that is a good approximation to **argmax-e** while being Lipschitz continuous. One such candidate is the softmax operator $\mathbf{softmax}_c : \mathbb{R}^n \rightarrow \mathbb{R}^n$, with

$$\mathbf{softmax}_c(x)_i = \frac{\exp(cx_i)}{\sum_{j=1}^n \exp(cx_j)}, \quad i = 1, \dots, n,$$

for some positive constant c . The resulting policies are sometimes called Boltzmann policies, and are widely used in the literature of reinforcement learning Asadi and Littman [8], Haarnoja et al. [40].

The softmax operator can be generalized to a wide class of operators. In fact, for positive constants $c, c' > 0$, one can consider a parametrized family $\mathcal{F}_{c,c'} \subseteq \{f_{c,c'} : \mathbb{R}^n \rightarrow \mathbb{R}^n\}$ of all “smoothed” **argmax-e**’s, *i.e.*, all $f_{c,c'} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ that satisfies the following two conditions:

- Condition 1: $f_{c,c'}$ is c -Lipschitz, *i.e.*, $\|f_{c,c'}(x) - f_{c,c'}(y)\|_2 \leq c\|x - y\|_2$.
- Condition 2: $f_{c,c'}$ is a good approximation of **argmax-e**, *i.e.*,

$$\|f_{c,c'}(x) - \mathbf{argmax-e}(x)\|_2 \leq 2n \exp(-c'\delta),$$

where $\delta = x_{\max} - \max_{x_j < x_{\max}} x_j$, $x_{\max} = \max_{i=1, \dots, n} x_i$, and $\delta := \infty$ when all x_j are equal.

Notice that $\mathcal{F}_{c,c'}$ is closed under convex combinations, *i.e.*, if $f_{c,c'}, g_{c,c'} \in \mathcal{F}_{c,c'}$, then for any $\lambda \in [0, 1]$, $\lambda f_{c,c'} + (1 - \lambda)g_{c,c'}$ also satisfies the two conditions. Hence $\mathcal{F}_{c,c'}$ is convex.

To have a better idea of what $\mathcal{F}_{c,c'}$ looks like, we describe a subset $\mathcal{B}_{c,c'}$ of $\mathcal{F}_{c,c'}$ consisting of the generalized softmax operator $\mathbf{softmax}_h : \mathbb{R}^n \rightarrow \mathbb{R}^n$, defined as

$$\mathbf{softmax}_h(x)_i = \frac{\exp(h(x_i))}{\sum_{j=1}^n \exp(h(x_j))}, \quad i = 1, \dots, n, \quad (21)$$

where $h : \mathbb{R} \rightarrow \mathbb{R}$ satisfies $c'(x - y) \leq h(x) - h(y) \leq c(x - y)$ for any $x \geq y$. When h is continuously differentiable, a sufficient condition is that $c' \leq h'(x) \leq c$. In particular, if $h(x) \equiv cx$ for some constant $c > 0$, the operator reduces to the classical softmax operator, in which case we overload the notation to write $\mathbf{softmax}_h$ as $\mathbf{softmax}_c$.

This operator is Lipschitz continuous and close to the **argmax-e** (see Lemmas 6 and 7 in the Appendix), and in particular one can show that $\mathcal{B}_{c,c'} \subseteq \mathcal{F}_{c,c'}$. As a result, even though smoothed (*e.g.*, Boltzmann) policies are not optimal, the difference between the smoothed and the optimal one can always be controlled by choosing a function h with appropriate parameters c, c' . Note that other smoothing operators (*e.g.*, Mellowmax Asadi and Littman [8], which is a softmax operator with time-varying and problem dependent temperatures) may also be considered.

Error control in updating \mathcal{L} . Given the sub-optimality of the smoothed policy, one needs to characterize the difference between the optimal policy and the non-optimal ones. In particular, one can define the action gap between the best and the second best actions in terms of the Q-value as

$$\delta^s(\mathcal{L}) := \max_{a' \in \mathcal{A}} Q_{\mathcal{L}}^*(s, a') - \max_{a \notin \mathop{\text{argmax}}_{a \in \mathcal{A}} Q_{\mathcal{L}}^*(s, a)} Q_{\mathcal{L}}^*(s, a) > 0.$$

Action gap is important for approximation algorithms Bellemare et al. [9], and is closely related to the problem-dependent bounds for regret analysis in reinforcement learning and multi-armed bandits, and advantage learning algorithms including A3C Minh et al. [61].

The problem is: in order for the learning algorithm to converge in terms of \mathcal{L} (Theorems 2 and 3), one needs to ensure a definite differentiation between the optimal policy and the sub-optimal ones. This is problematic as the infimum of $\delta^s(\mathcal{L})$ over an infinite number of \mathcal{L} can be 0. To address this, the population distribution at step k , say \mathcal{L}_k , needs to be projected to a finite grid, called ϵ -net. The relation between the ϵ -net and action gaps is as follows:

For any $\epsilon > 0$, there exist a positive function $\phi(\epsilon)$ and an ϵ -net $S_\epsilon := \{\mathcal{L}^{(1)}, \dots, \mathcal{L}^{(N_\epsilon)}\} \subseteq \mathcal{P}(\mathcal{S} \times \mathcal{A})$, with the properties that $\min_{i=1, \dots, N_\epsilon} d_{TV}(\mathcal{L}, \mathcal{L}^{(i)}) \leq \epsilon$ for any $\mathcal{L} \in \mathcal{P}(\mathcal{S} \times \mathcal{A})$, and that $\max_{a' \in \mathcal{A}} Q_{\mathcal{L}^{(i)}}^(s, a') - Q_{\mathcal{L}^{(i)}}^*(s, a) \geq \phi(\epsilon)$ for any $i = 1, \dots, N_\epsilon$, $s \in \mathcal{S}$, and any $a \notin \mathop{\text{argmax}}_{a \in \mathcal{A}} Q_{\mathcal{L}^{(i)}}^*(s, a)$.*

Here the existence of ϵ -nets is trivial due to the compactness of the probability simplex $\mathcal{P}(\mathcal{S} \times \mathcal{A})$, and the existence of $\phi(\epsilon)$ comes from the finiteness of the action set \mathcal{A} .

In practice, $\phi(\epsilon)$ often takes the form of $D\epsilon^\alpha$ with $D > 0$ and the exponent $\alpha > 0$ characterizing the decay rate of the action gaps. In general, experiments are robust with respect to the choice of ϵ -net.

In the next section, we propose value based and policy based algorithms for learning GMFG.

5. RL Algorithms for (stationary) GMFGs

5.1. Value-based algorithms We start by introducing the following definition.

DEFINITION 2 (VALUE-BASED GUARANTEE). For an arbitrary MDP \mathcal{M} , we say that an algorithm has a *value-based guarantee* with parameters $\{C_{\mathcal{M}}^{(i)}, \alpha_1^{(i)}, \alpha_2^{(i)}, \alpha_3^{(i)}, \alpha_4^{(i)}\}_{i=1}^m$, if for any $\epsilon, \delta > 0$, after obtaining

$$T_{\mathcal{M}}(\epsilon, \delta) = \sum_{i=1}^m C_{\mathcal{M}}^{(i)} \left(\frac{1}{\epsilon}\right)^{\alpha_1^{(i)}} \left(\log \frac{1}{\epsilon}\right)^{\alpha_2^{(i)}} \left(\frac{1}{\delta}\right)^{\alpha_3^{(i)}} \left(\log \frac{1}{\delta}\right)^{\alpha_4^{(i)}} \quad (22)$$

samples from the simulator oracle \mathcal{G} , with probability at least $1 - 2\delta$, it outputs an approximate Q-function $\hat{Q}^{T_{\mathcal{M}}(\epsilon, \delta)}$ which satisfies $\|\hat{Q}^{T_{\mathcal{M}}(\epsilon, \delta)} - Q^*\|_\infty \leq \epsilon$. Here the norm $\|\cdot\|_\infty$ is understood element-wisely.

5.1.1. GMF-V We now state the first main algorithm (Algorithm 2). It applies to any algorithm Alg with a value-based guarantee.

Algorithm 2 GMF-V($Alg, f_{c,c'}$)

- 1: **Input:** Initial \mathcal{L}_0 , ϵ -net S_ϵ , temperatures $c, c' > 0$, tolerances $\epsilon_k, \delta_k > 0, k = 0, 1, \dots$
 - 2: **for** $k = 0, 1, \dots$ **do**
 - 3: Apply Alg to find the approximate Q-function $\hat{Q}_k^* = \hat{Q}^{T_k}$ of the MDP $\mathcal{M}_{\mathcal{L}_k}$, where $T_k = T_{\mathcal{M}_{\mathcal{L}_k}}(\epsilon_k, \delta_k)$.
 - 4: Compute $\pi_k(s) = f_{c,c'}(\hat{Q}_k^*(s, \cdot))$.
 - 5: Sample $s \sim \mu_k$ (μ_k is the population state marginal of \mathcal{L}_k), obtain $\tilde{\mathcal{L}}_{k+1}$ from $\mathcal{G}(s, \pi_k, \mathcal{L}_k)$.
 - 6: Find $\mathcal{L}_{k+1} = \mathbf{Proj}_{S_\epsilon}(\tilde{\mathcal{L}}_{k+1})$
 - 7: **end for**
-

Here $\mathbf{Proj}_{S_\epsilon}(\mathcal{L}) = \operatorname{argmin}_{\mathcal{L}^{(1)}, \dots, \mathcal{L}^{(N_\epsilon)}} d_{TV}(\mathcal{L}^{(i)}, \mathcal{L})$. For computational tractability, it is sufficient to choose S_ϵ as a truncation grid so that projection of $\tilde{\mathcal{L}}_k$ onto the ϵ -net reduces to truncating $\tilde{\mathcal{L}}_k$ to a certain number of digits. For instance, in experiments in Section 8, the number of digits is chosen to be 4. Appropriate choices of the hyper-parameters c, c', ϵ and tolerances ϵ_k, δ_k ($k \geq 0$) are given in Theorems 2. Our experiment shows the algorithm is robust with respect to these hyper-parameters.

We next establish the convergence of the above GMF-V algorithm to an approximate Nash equilibrium of (GMFG), with complexity analysis.

THEOREM 2 (Convergence and complexity of GMF-V). *Assume the same assumptions as Theorem 1, and $f_{c,c'} \subseteq \mathcal{F}_{c,c'}$. Suppose that Alg has a value-based guarantee with parameters*

$$\{C_{\mathcal{M}}^{(i)}, \alpha_1^{(i)}, \alpha_2^{(i)}, \alpha_3^{(i)}, \alpha_4^{(i)}\}_{i=1}^m.$$

For any $\epsilon, \delta > 0$, set $\delta_k = \delta/K_{\epsilon, \eta}$, $\epsilon_k = (k+1)^{-(1+\eta)}$ for some $\eta \in (0, 1]$ ($k = 0, \dots, K_{\epsilon, \eta} - 1$), and $c \geq c' \geq \frac{\log(1/\epsilon)}{\phi(\epsilon)}$. Then with probability at least $1 - 2\delta$,

$$W_1(\mathcal{L}_{K_{\epsilon, \eta}}, \mathcal{L}^*) \leq C_0 \epsilon.$$

Here $K_{\epsilon, \eta} := \lceil 2 \max\{(\eta\epsilon/c)^{-1/\eta}, \log_d(\epsilon/\max\{\operatorname{diam}(\mathcal{S})\operatorname{diam}(\mathcal{A}), c\}) + 1\} \rceil$ is the number of outer iterations, and the constant C_0 is independent of δ, ϵ and η .

Moreover, the total number of samples $T = \sum_{k=0}^{K_{\epsilon, \eta}-1} T_{\mathcal{M}_{\mathcal{L}_k}}(\delta_k, \epsilon_k)$ is bounded by

$$T \leq \sum_{i=1}^m \frac{2^{\alpha_2^{(i)}}}{2\alpha_1^{(i)} + 1} C_{\mathcal{M}}^{(i)} K_{\epsilon, \eta}^{2\alpha_1^{(i)} + 1} (K_{\epsilon, \eta}/\delta)^{\alpha_3^{(i)}} (\log(K_{\epsilon, \eta}/\delta))^{\alpha_2^{(i)} + \alpha_4^{(i)}}. \quad (23)$$

The proof of Theorem 2 (in Section 7.4) depends on the Lipschitz continuity of the smoothing operator $f_{c,c'}$, the closeness between $f_{c,c'}$ and the **argmax-e** (Lemma 7 in Section 7.3), and the complexity of Alg provided by the value-based guarantee.

5.1.2. GMF-V-Q: GMF-V with Q-learning As an example of the GMF-V algorithm, we describe algorithm GMF-V-Q, a Q-learning based GMF-V algorithm. For an MDP $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, r, \gamma)$, the synchronous Q-learning algorithm approximates the value iteration by stochastic approximation. At each step l , with state s and action a , the system reaches state s' according to the controlled dynamics, and the Q-function approximation Q_l is updated by

$$\hat{Q}^{l+1}(s, a) = (1 - \beta_l)\hat{Q}^l(s, a) + \beta_l \left[r(s, a) + \gamma \max_{\bar{a}} \hat{Q}^l(s', \bar{a}) \right], \quad \forall s \in \mathcal{S}, a \in \mathcal{A}, \quad (24)$$

where $\hat{Q}^0(s, a) = C$ for some constant $C \in \mathbb{R}$ for any $s \in \mathcal{S}$ and $a \in \mathcal{A}$, and the step size β_l can be chosen as (Even-Dar and Mansour [24])

$$\beta_l = |l + 1|^{-h}, \quad (25)$$

with $h \in (1/2, 1)$.

The corresponding synchronous Q-learning based algorithm with the standard **softmax** operator is GMF-V-Q (Algorithm 3), and will be used in the experiment (Section 8).

Algorithm 3 Q-learning for GMFGs (GMF-V-Q)

- 1: **Input:** Initial \mathcal{L}_0 , ϵ -net S_ϵ , tolerances $\epsilon_k, \delta_k > 0, k = 0, 1, \dots$.
 - 2: **for** $k = 0, 1, \dots$ **do**
 - 3: Perform synchronous Q-learning with stepsizes (25) for $T_k = T_{\mathcal{M}_{\mathcal{L}_k}}(\epsilon_k, \delta_k)$ iterations to find the approximate Q-function $\hat{Q}_k^* = \hat{Q}^{T_k}$ of the MDP $\mathcal{M}_{\mathcal{L}_k}$.
 - 4: Compute $\pi_k \in \Pi$ with $\pi_k(s) = \mathbf{softmax}_c(\hat{Q}_k^*(s, \cdot))$.
 - 5: Sample $s \sim \mu_k$ (μ_k is the population state marginal of \mathcal{L}_k), obtain $\tilde{\mathcal{L}}_{k+1}$ from $\mathcal{G}(s, \pi_k, \mathcal{L}_k)$.
 - 6: Find $\mathcal{L}_{k+1} = \mathbf{Proj}_{S_\epsilon}(\tilde{\mathcal{L}}_{k+1})$
 - 7: **end for**
-

Let us first recall the following sample complexity result for synchronous Q-learning method.

LEMMA 2 (Even-Dar and Mansour [24]: sample complexity of synchronous Q-learning). For an MDP, say $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, r, \gamma)$, suppose that the Q-learning algorithm takes step-sizes (25). Then $\|\hat{Q}^{T_{\mathcal{M}}(\delta, \epsilon)} - Q_{\mathcal{M}}^*\|_\infty \leq \epsilon$ with probability at least $1 - 2\delta$. Here \hat{Q}^T is the T -th update in the Q-learning updates (24), $Q_{\mathcal{M}}^*$ is the (optimal) Q-function, and

$$T_{\mathcal{M}}(\epsilon, \delta) = \Omega \left(\left(\frac{V_{\max}^2 \log \left(\frac{|\mathcal{S}||\mathcal{A}|V_{\max}}{\delta\beta\epsilon} \right)}{\beta^2\epsilon^2} \right)^{\frac{1}{h}} + \left(\frac{1}{\beta} \log \frac{V_{\max}}{\epsilon} \right)^{\frac{1}{1-h}} \right),$$

where $\beta = (1 - \gamma)/2$, $V_{\max} = R_{\max}/(1 - \gamma)$, and R_{\max} is such that a.s. $0 \leq r(s, a) \leq R_{\max}$.

This lemma implies immediately the value-based guarantee (as in Definition 2) and the convergence for GMF-V-Q. Similar results can be established for asynchronous Q-learning method, as shown in Appendix B.

COROLLARY 1. The synchronous Q-learning algorithm with appropriate choices of step-sizes (cf. (25)) satisfies the value-based guarantee with parameters $\{\tilde{C}_{\mathcal{M}}^{(i)}, \alpha_1^{(i)}, \alpha_2^{(i)}, \alpha_3^{(i)}, \alpha_4^{(i)}\}_{i=1}^3$, where $C_{\mathcal{M}}^{(i)} (i = 1, 2, 3)$ are constants depending on $|\mathcal{S}|, |\mathcal{A}|, V_{\max}, \beta$ and h , and

$$\begin{aligned} \alpha_1^{(1)} &= 2/h, \alpha_2^{(1)} = 1/h, \alpha_3^{(1)} = \alpha_4^{(1)} = 0; \\ \alpha_1^{(2)} &= 2/h, \alpha_2^{(2)} = \alpha_3^{(2)} = 0 \text{ and } \alpha_4^{(2)} = 1/h; \\ \alpha_1^{(3)} &= 0, \alpha_2^{(3)} = 1/(1-h), \alpha_3^{(3)} = 0 \text{ and } \alpha_4^{(3)} = 0. \end{aligned}$$

In addition, assume the same assumptions as Theorem 1, then for Algorithm 3 with synchronous Q-learning method, with probability at least $1 - 2\delta$, $W_1(\mathcal{L}_{K_{\epsilon,\eta}}, \mathcal{L}^*) \leq C_0\epsilon$, where $K_{\epsilon,\eta}$ is defined as in Theorem 2. And the total number of samples $T = \sum_{k=0}^{K_{\epsilon,\eta}-1} T_{\mathcal{M}_{\mathcal{L}_k}}(\epsilon_k, \delta_k)$ is bounded by

$$T \leq O \left(K_{\epsilon,\eta}^{\frac{4}{h}+1} \left(\log \frac{K_{\epsilon,\eta}}{\delta} \right)^{\frac{1}{h}} + \left(\log \frac{K_{\epsilon,\eta}}{\delta} \right)^{\frac{1}{1-h}} \right).$$

5.2. Policy-based algorithms In addition to algorithms with value-based guarantees (cf. Definition 2), there are also numerous algorithms with *policy-based guarantees*.

DEFINITION 3 (POLICY-BASED GUARANTEE). For an arbitrary MDP \mathcal{M} , we say that an algorithm has a *policy-based guarantee* with parameters $\{C_{\mathcal{M}}^{(i)}, \alpha_1^{(i)}, \alpha_2^{(i)}, \alpha_3^{(i)}, \alpha_4^{(i)}\}_{i=1}^m$, if for any $\epsilon, \delta > 0$, after obtaining

$$T_{\mathcal{M}}(\epsilon, \delta) = \sum_{i=1}^m C_{\mathcal{M}}^{(i)} \left(\frac{1}{\epsilon} \right)^{\alpha_1^{(i)}} \left(\log \frac{1}{\epsilon} \right)^{\alpha_2^{(i)}} \left(\frac{1}{\delta} \right)^{\alpha_3^{(i)}} \left(\log \frac{1}{\delta} \right)^{\alpha_4^{(i)}} \quad (26)$$

samples from the simulator oracle \mathcal{G} , with probability at least $1 - 2\delta$, it outputs an approximate policy $\pi_{T_{\mathcal{M}}(\epsilon, \delta)}$, which satisfies $V_{\mathcal{M}}^*(s) - V_{\pi_{T_{\mathcal{M}}(\epsilon, \delta)}}(s) \leq \epsilon, \forall s \in \mathcal{S}$.

5.2.1. GMF-P Before we present policy-based RL algorithms, let us first establish a connection between policy-based and value-based guarantees.

To start, take any policy $\pi \in \Pi$, consider the following synchronous temporal difference (TD) iterations:

$$\tilde{Q}_{\pi}^{l+1}(s, a) = (1 - \beta_l) \tilde{Q}_{\pi}^l(s, a) + \beta_l \left[r(s, a) + \gamma \tilde{Q}_{\pi}^l(s', a') \right], \quad \forall s \in \mathcal{S}, a \in \mathcal{A}, \quad (27)$$

where $a' \sim \pi(s')$, $\tilde{Q}_{\pi}^0(s, a) = C$ for some constant $C \in \mathbb{R}$ and any $s \in \mathcal{S}$ and $a \in \mathcal{A}$, and the step size $\beta_l = (l+1)^{-h}$ for some $h \in (1/2, 1)$.

Then we have

LEMMA 3. Suppose that the algorithm Alg satisfies a policy-based guarantee with parameters $\{C_{\mathcal{M}}^{(i)}, \alpha_1^{(i)}, \alpha_2^{(i)}, \alpha_3^{(i)}, \alpha_4^{(i)}\}_{i=1}^m$. Let \tilde{Q}_{π}^l be defined by (27). Then for any $\delta \in (0, 1)$ and $\epsilon > 0$, with probability at least $1 - 2\delta$, $\left\| \tilde{Q}_{\pi_{T_{\mathcal{M}}(\epsilon, \delta/2)}}^l - Q_{\mathcal{M}}^* \right\|_{\infty} \leq \epsilon$ if

$$l = \Omega \left(\left(\frac{V_{\max} \log \left(\frac{|\mathcal{S}||\mathcal{A}|V_{\max}}{\delta\beta^2\epsilon} \right)}{\beta^4\epsilon^2} \right)^{1/h} + \left(\frac{1}{\beta} \log \frac{V_{\max}}{\beta\epsilon} \right)^{1/(1-h)} \right), \quad (28)$$

where $V_{\max} = R_{\max}/(1 - \gamma)$ and $\beta = (1 - \gamma)/2$.

Consequently, the algorithm Alg (combined with TD updates (27)) also has a value-based guarantee with parameters $\{\tilde{C}_{\mathcal{M}}^{(i)}, \alpha_1^{(i)}, \alpha_2^{(i)}, \alpha_3^{(i)}, \alpha_4^{(i)}\}_{i=1}^{m+3}$, where $\tilde{C}_{\mathcal{M}}^{(i)}$ is some constant multiple of $C_{\mathcal{M}}^{(i)}$ ($i = 1, \dots, m$), $\tilde{C}_{\mathcal{M}}^{(m+1)}$ ($i = 1, 2, 3$) are constants depending on $V_{\max}, |\mathcal{S}|, |\mathcal{A}|, \beta$ and h , and we have

$$\begin{aligned} \alpha_1^{(m+1)} &= 2/h, \alpha_2^{(m+1)} = 1/h, \alpha_3^{(m+1)} = \alpha_4^{(m+1)} = 0; \\ \alpha_1^{(m+2)} &= 2/h, \alpha_2^{(m+2)} = \alpha_3^{(m+2)} = 0 \text{ and } \alpha_4^{(m+2)} = 1/h; \\ \alpha_1^{(m+3)} &= 0, \alpha_2^{(m+3)} = 1/(1-h), \alpha_3^{(m+3)} = 0 \text{ and } \alpha_4^{(m+3)} = 0. \end{aligned} \quad (29)$$

The above lemma indicates that any algorithm with a policy-based guarantee also satisfies a value-based guarantee with similar parameters (when combined with the TD updates). The policy-based algorithm GMF-P (Algorithm 4) makes use of Lemma 3 to select the hyper-parameter l so that the resulting $\tilde{Q}_{\pi_{T_{\mathcal{M}}(\epsilon, \delta/2)}}^l$ forms a good value-based certificate.

We next present the convergence property for the GMF-P algorithm by combining the proofs of Lemma 3 and Theorem 2.

Algorithm 4 GMF-P(Alg, $f_{c,c'}$)

-
- 1: **Input:** Initial \mathcal{L}_0 , ϵ -net S_ϵ , temperatures $c, c' > 0$, tolerances $\epsilon_k, \delta_k > 0, k = 0, 1, \dots$
 - 2: **for** $k = 0, 1, \dots$ **do**
 - 3: Apply Alg to find the approximate policy $\hat{\pi}_k = \pi_{T_k}$ of the MDP $\mathcal{M}_k := \mathcal{M}_{\mathcal{L}_k}$, where $T_k = T_{\mathcal{M}_k}(\epsilon_k, \delta_k/2)$.
 - 4: Compute $\tilde{Q}_{\hat{\pi}_k}^{l_k}$ using TD updates (27) for MDP \mathcal{M}_k , with l_k satisfying (28) (with ϵ and δ replaced by ϵ_k and $\delta_k/2$, respectively).
 - 5: Compute $\pi_k(s) = f_{c,c'}(\tilde{Q}_{\hat{\pi}_k}^{l_k}(s, \cdot))$.
 - 6: Sample $s \sim \mu_k$ (μ_k is the population state marginal of \mathcal{L}_k), obtain $\tilde{\mathcal{L}}_{k+1}$ from $\mathcal{G}(s, \pi_k, \mathcal{L}_k)$.
 - 7: Find $\mathcal{L}_{k+1} = \mathbf{Proj}_{S_\epsilon}(\tilde{\mathcal{L}}_{k+1})$
 - 8: **end for**
-

THEOREM 3 (Convergence and complexity of GMF-P). *Assume the same assumptions as in Theorem 1, and in addition that $f_{c,c'} \subseteq \mathcal{F}_{c,c'}$. Suppose that Alg has a policy-based guarantee with parameters*

$$\{C_{\mathcal{M}}^{(i)}, \alpha_1^{(i)}, \alpha_2^{(i)}, \alpha_3^{(i)}, \alpha_4^{(i)}\}_{i=1}^m.$$

Then for any $\epsilon, \delta > 0$, set $\delta_k = \delta/K_{\epsilon,\eta}$, $\epsilon_k = (k+1)^{-(1+\eta)}$ for some $\eta \in (0, 1]$ ($k = 0, \dots, K_{\epsilon,\eta} - 1$), and $c \geq c' \geq \frac{\log(1/\epsilon)}{\phi(\epsilon)}$, with probability at least $1 - 2\delta$,

$$W_1(\mathcal{L}_{K_{\epsilon,\eta}}, \mathcal{L}^*) \leq C_0 \epsilon.$$

Here $K_{\epsilon,\eta} := \lceil 2 \max\{(\eta\epsilon/c)^{-1/\eta}, \log_d(\epsilon/\max\{\text{diam}(\mathcal{S})\text{diam}(\mathcal{A}), c\}) + 1\} \rceil$ is the number of outer iterations, and the constant C_0 is independent of δ, ϵ and η .

Moreover, the total number of samples $T = \sum_{k=0}^{K_{\epsilon,\eta}-1} T_{\mathcal{M}_{\mathcal{L}_k}}(\delta_k, \epsilon_k)$ is bounded by

$$T \leq \sum_{i=1}^{m+3} \frac{2^{\alpha_2^{(i)}}}{2^{\alpha_1^{(i)}} + 1} \tilde{C}_{\mathcal{M}}^{(i)} K_{\epsilon,\eta}^{2\alpha_1^{(i)}+1} (K_{\epsilon,\eta}/\delta)^{\alpha_3^{(i)}} (\log(K_{\epsilon,\eta}/\delta))^{\alpha_2^{(i)} + \alpha_4^{(i)}}, \quad (30)$$

where the parameters $\{\tilde{C}_{\mathcal{M}}^{(i)}, \alpha_1^{(i)}, \alpha_2^{(i)}, \alpha_3^{(i)}, \alpha_4^{(i)}\}_{i=1}^{m+3}$ are defined in Lemma 3.

5.2.2. GMF-P-TRPO: GMF-P with TRPO A special form of the GMF-P algorithm utilizes the trust region policy optimization (TRPO) algorithm Schulman et al. [72], Shani et al. [73]. We call it GMF-P-TRPO.

Sample-based TRPO Shani et al. [73] assumes access to a ν -restart model. That is, it can only access sampled trajectories and restarts according to the distribution ν . Here we pick ν such that $C^{\pi^*} := \left\| \frac{d_{\text{Unif}_{\mathcal{S}}}^{\pi^*}}{\nu} \right\|_{\infty} = \max_{s \in \mathcal{S}} \left| \frac{d_{\text{Unif}_{\mathcal{S}}}^{\pi^*}(s)}{\nu(s)} \right| < \infty$, where $d_{\rho}^{\pi} = (1 - \gamma)\rho(I - \gamma P^{\pi})^{-1}$ and $\text{Unif}_{\mathcal{S}}$ is the uniform distribution on set \mathcal{S} . Sample-based TRPO samples M_0 trajectories per episode. The initial state s_0 at the beginning of each episode is sampled from ν . In every trajectory m ($m = 1, 2, \dots, M_0$) of the l -th episode, it first samples $s_m \sim d_{\nu}^{\pi_l}$ and takes an action $a_m \sim \text{Unif}_{\mathcal{A}}$ where $\text{Unif}_{\mathcal{A}}$ is the uniform distribution on the set \mathcal{A} . Then, by following the current π_l , it estimates $Q^{\pi_l}(s_m, a_m)$ using a rollout. Denote this estimate as $\hat{Q}^{\pi_l}(s_m, a_m, m)$ and observe that it is (nearly) an unbiased estimator of $Q^{\pi_l}(s_m, a_m)$. We assume that each rollout runs sufficiently long so that the bias is sufficiently small. Sample-Based TRPO updates the policy at the end of the l -th episode, by the following proximal problem

$$\pi_{l+1} \in \arg \max_{\pi \in \Delta_{\mathcal{A}}^{|\mathcal{S}|}} \left\{ \frac{1}{M_0} \sum_{m=1}^{M_0} \frac{1}{t_l(1-\gamma)} B_w(s_m; \pi, \pi_l) + \langle \hat{\nabla} V^{\pi_l}[m], \pi(s_m) - \pi_l(s_m) \rangle \right\},$$

where the estimation of the gradient is

$$\hat{\nabla} V^{\pi_l}[m] := \frac{1}{1-\gamma} |\mathcal{A}| \hat{Q}^{\pi_l}(s_m, \cdot, m) \circ \mathbf{I}_{\{=a_m\}}.$$

Given two policies π_1 and π_2 , we denote their Bregman distance associated with a strongly convex function w as $B_w(s; \pi_1, \pi_2) = B_w(\pi_1(s), \pi_2(s))$, where $B_w(x, y) := w(x) - w(y) - \langle \nabla w(y), x - y \rangle$ and $\pi_i(s) \in \mathcal{P}(\mathcal{A})$ ($i = 1, 2$). Denote $B_w(\pi_1, \pi_2) \in \mathbb{R}^{|\mathcal{S}|}$ as the corresponding state-wise vector. Here we consider two common cases for w : when $w(x) = \frac{1}{2} \|x\|_2^2$ is the Euclidean distance, $B_w(x, y) = \frac{1}{2} \|x - y\|_2^2$; when $w(x) = H(x)$ is the negative entropy, $B_w(x, y) = d_{\text{KL}}(x||y)$. We refer to [73, Section 6.2] for more detailed discussion on Sample-based TRPO.

The above guarantee follows from the sample complexity result below by specifying $\mu := \text{Unif}_{\mathcal{S}}$. Notice that here for any $\mu \in \mathcal{P}(\mathcal{S})$, we define $V^*(\mu) := \sum_{s \in \mathcal{S}} \mu(s) V^*(s)$, and similarly $V^{\pi_k}(\mu) := \sum_{s \in \mathcal{S}} \mu(s) V^{\pi_k}(s)$.

The sample complexity of TRPO algorithm can be characterized as below.

LEMMA 4 (Theorem 5 in Shani et al. [73]: sample complexity of TRPO). *Let $\{\pi_l\}_{l \geq 0}$ be the sequence generated by Sample-Based TRPO, using*

$$M_0 \geq \Omega\left(\frac{|\mathcal{A}|^2 C^2 (|\mathcal{S}| \log |\mathcal{A}| + \log 1/\delta)}{(1-\gamma)^2 \epsilon^2}\right)$$

samples in each episode, with $t_l = \frac{(1-\gamma)}{C_{w,1} C \sqrt{l+1}}$. Let $\{V_{best}^N\}_{N \geq 0}$ be the sequence of best achieved values, $V_{best}^N(\mu) := \max_{l=0,1,\dots,N} V^{\pi_l}(\mu)$, where $\mu \in \mathcal{P}(\mathcal{S})$. Then with probability greater than $1 - \delta$ for every $\epsilon > 0$, the following holds for all $N \geq 1$:

$$V^*(\mu) - V_{best}^N(\mu) \leq O\left(\frac{C_{w,1} C}{(1-\gamma)^2 \sqrt{N}} + \frac{C^{\pi^*} \epsilon}{(1-\gamma)^2}\right).$$

Here $C > 0$ is the upper bound on the reward function r , $C_{w,1} = \sqrt{|\mathcal{A}|}$ in the euclidean case and $C_{w,1} = 1$ in the non-euclidean case, $C_{w,2} = 1$ for the euclidean case and $C_{w,2} = |\mathcal{A}|^2$ for the non-euclidean case. Note that unlike the case of Q-learning, here we are only guaranteed to have *some* iterate among iterations $0, \dots, N$ that satisfy the desired sub-optimality bound. Note that this is a common pattern of the theoretical results for policy optimization algorithms in the RL literature Agarwal et al. [2], Wang et al. [82], unless the (oracle) access to exact policy gradients is assumed Mei et al. [59]. For simplicity, hereafter we assume an oracle access to such an iterate after running TRPO. In practice, with additional (polynomial number of) samples, one can explicitly identify a single policy satisfying the desired bound with high probability; see *e.g.*, the two-phase technique in Ghadimi and Lan [31].

Note that [73, Theorem 5] has both regularized version and unregularized version of TRPO. Here we only adopt the unregularized version which fits the framework of Algorithm 4. For more materials on regularized MDPs and reinforcement learning, we refer the readers to Neu et al. [63], Geist et al. [30], Derman and Mannor [22].

Based on the sample complexity in Lemma 4, the following policy-based guarantee for TRPO algorithm and the convergence result for GMF-P-TRPO can be obtained.

COROLLARY 2. *Let $t_l = \frac{(1-\gamma)}{C_{w,1} C \sqrt{l+1}}$, then TRPO algorithm satisfies the policy-based guarantee with parameters $\{\tilde{C}_{\mathcal{M}}^{(i)}, \alpha_1^{(i)}, \alpha_2^{(i)}, \alpha_3^{(i)}, \alpha_4^{(i)}\}_{i=1}^2$, where $C_{\mathcal{M}}^{(i)}$ ($i = 1, 2$) are constants depending on $|\mathcal{S}|, |\mathcal{A}|, V_{\max}, \beta$ and h , and we have:*

$$\begin{aligned} \alpha_1^{(1)} &= 5/2, \quad \alpha_j^{(1)} = 0 \text{ for } j = 2, 3, 4, \\ \alpha_1^{(2)} &= 5/2, \quad \alpha_4^{(2)} = 1, \alpha_2^{(2)} = \alpha_3^{(2)} = 0. \end{aligned}$$

In addition, under same assumptions as Theorem 1, then for Algorithm 4 using TRPO method, with probability at least $1 - 2\delta$, $W_1(\mathcal{L}_{K_{\epsilon,\eta}}, \mathcal{L}^*) \leq C_0\epsilon$, where $K_{\epsilon,\eta}$ is defined as in Theorem 2. And the total number of samples $T = \sum_{k=0}^{K_{\epsilon,\eta}-1} T_{\mathcal{M}_{\mathcal{L}_k}}(\delta_k, \epsilon_k)$ is bounded by

$$T \leq O \left(K_{\epsilon,\eta}^6 \left(\log \frac{K_{\epsilon,\eta}}{\delta} \right) + K_{\epsilon,\eta}^{\frac{4}{h}+1} \left(\log \frac{K_{\epsilon,\eta}}{\delta} \right)^{\frac{1}{h}} + \left(\log \frac{K_{\epsilon,\eta}}{\delta} \right)^{\frac{1}{1-h}} \right).$$

6. Applications to N -player Games In this section, we discuss a potential application of our modeling and approach to N -player settings. To this end, we consider extensions of Algorithms 2 and 4 with weaker assumptions on the simulator access. In particular, we weaken the simulator oracle assumption in Section 4 as follows.

Weak simulator oracle. For each player i , given any policy $\pi \in \Pi$, the current state $s_i \in \mathcal{S}$, for any empirical population state-action distribution \mathcal{L}_N , one can obtain a *sample* of the next state $s'_i \sim P_{\mathcal{L}_N}(\cdot | s_i, \pi(s_i)) = P(\cdot | s_i, \pi(s_i), \mathcal{L}_N)$ and a reward $r = r_{\mathcal{L}_N}(s_i, \pi(s_i)) = r(s_i, \pi(s_i), \mathcal{L}_N)$. For brevity, we denote the simulator as $(s'_i, r) = \mathcal{G}_W(s_i, \pi, \mathcal{L}_N)$.

We say that \mathcal{L}_N is an empirical population state-action distribution of N -players if for each $s \in \mathcal{S}$, $a \in \mathcal{A}$, $\mathcal{L}_N(s, a) = \frac{1}{N} \sum_{i=1}^N \mathbf{I}_{s_i=s, a_i=a}$ for some state-action profile of $\{s_i, a_i\}_{i=1}^N$. Equivalently, this holds if $N\mathcal{L}_N(s, a)$ is a non-negative integer for each $s \in \mathcal{S}$, $a \in \mathcal{A}$, and $\sum_{s,a} \mathcal{L}_N(s, a) = 1$. We denote the set of empirical population state-action distributions as \mathbf{Emp}_N .

RL algorithms with access only to \mathcal{G}_W . Compared to the original simulator oracle \mathcal{G} , the weak simulator \mathcal{G}_W only accepts empirical population state-action distributions as inputs, and does not directly output the next (empirical) population state-action distribution.

To make use of the simulator \mathcal{G}_W , we modify Algorithm 2 and Algorithm 4 to algorithms (Algorithms 5 and 6). In particular, see Step 6 in Algorithm 5 and Step 7 in Algorithm 6 for generating empirical distributions from simulator \mathcal{G}_W .

Algorithm 5 GMF-VW(Alg, $f_{c,c'}$): weak simulator

- 1: **Input:** Initial \mathcal{L}_0 , temperatures $c, c' > 0$, tolerances $\epsilon_k, \delta_k > 0, k = 0, 1, \dots$
 - 2: **for** $k = 0, 1, \dots$ **do**
 - 3: Apply Alg to find the approximate Q-function $\hat{Q}_k^* = \hat{Q}^{T_k}$ of the MDP $\mathcal{M}_{\mathcal{L}_k}$, where $T_k = T_{\mathcal{M}_{\mathcal{L}_k}}(\epsilon_k, \delta_k)$.
 - 4: Compute $\pi_k(s) = f_{c,c'}(\hat{Q}_k^*(s, \cdot))$.
 - for** $i = 1, 2, \dots, N$ **do**
 - 6: Sample $s_i \stackrel{\text{i.i.d.}}{\sim} \mu_k$, then obtain s'_i i.i.d. from $\mathcal{G}_W(s_i, \pi_k, \mathcal{L}_k)$ and $a'_i \stackrel{\text{i.i.d.}}{\sim} \pi_k(s'_i)$.
 - end for**
 - Compute \mathcal{L}_{k+1} with $\mathcal{L}_{k+1}(s, a) = \frac{1}{N} \sum_{i=1}^N \mathbf{I}_{s'_i=s, a'_i=a}$.
 - 9: **end for**
-

One can observe that \mathbf{Emp}_N already serves as an $1/N$ -net. So one can directly use it without additional projections. The definition of \mathcal{L}_k also makes sure that $\mathcal{L}_k \in \mathbf{Emp}_N$ as required for the input of the weaker simulator.

Convergence results similar to Theorems 2 and 3 can be obtained for Algorithms 5 and 6, respectively. (See Appendix C.) Here the major difference is an additional $O(1/\sqrt{N})$ term in the finite step error bound. It is worth mentioning that $O(1/\sqrt{N})$ is consistent with the literature on MFG approximation errors of finite N -player games Huang et al. [46].

Algorithm 6 GMF-PW(*Alg*, $f_{c,c'}$): weak simulator

-
- 1: **Input:** Initial \mathcal{L}_0 , temperatures $c, c' > 0$, tolerances $\epsilon_k, \delta_k > 0, k = 0, 1, \dots$
 - 2: **for** $k = 0, 1, \dots$ **do**
 - 3: Apply *Alg* to find the approximate policy $\hat{\pi}_k = \pi_{T_k}$ of the MDP $\mathcal{M}_k := \mathcal{M}_{\mathcal{L}_k}$, where $T_k = T_{\mathcal{M}_{\mathcal{L}_k}}(\epsilon_k, \delta_k/2)$.
 - 4: Compute $\tilde{Q}_{\hat{\pi}_k}^{l_k}$ using TD updates (27) for MDP \mathcal{M}_k , with l_k satisfying (28) (with ϵ and δ replaced by ϵ_k and $\delta_k/2$, respectively).
 - 5: Compute $\pi_k(s) = f_{c,c'}(\tilde{Q}_{\hat{\pi}_k}^{l_k}(s, \cdot))$.
 - for** $i = 1, 2, \dots, N$ **do**
 - 7: Sample $s_i \stackrel{\text{i.i.d.}}{\sim} \mu_k$, then obtain s'_i i.i.d. from $\mathcal{G}_W(s_i, \pi_k, \mathcal{L}_k)$ and $a'_i \stackrel{\text{i.i.d.}}{\sim} \pi_k(s'_i)$.
 - end for**
 - Compute \mathcal{L}_{k+1} with $\mathcal{L}_{k+1}(s, a) = \frac{1}{N} \sum_{i=1}^N \mathbf{I}_{s'_i=s, a'_i=a}$.
 - 10: **end for**
-

7. Proof of the main results

7.1. Proof of Lemma 1 In this section, we provide the proof of Lemma 1.

[Proof of Lemma 1] We begin by noticing that $\mathcal{L}' = \Gamma_2(\pi, \mathcal{L})$ can be expanded and computed as follows:

$$\mu'(s') = \sum_{s \in \mathcal{S}, a \in \mathcal{A}} \mu(s) P(s'|s, a, \mathcal{L}) \pi(a|s), \quad \mathcal{L}'(s', a') = \mu'(s') \pi(a'|s'), \quad (31)$$

where μ is the state marginal distribution of \mathcal{L} .

Since the Wasserstein distance W_1 can be related to the total variation distance via the following inequalities Gibbs and Su [33]:

$$d_{\min}(\mathcal{X}) d_{TV}(\nu, \nu') \leq W_1(\nu, \nu') \leq \text{diam}(\mathcal{X}) d_{TV}(\nu, \nu'), \quad (32)$$

where $d_{\min}(\mathcal{X}) = \min_{x \neq y \in \mathcal{X}} \|x - y\|_2$, we have

$$\begin{aligned}
W_1(\Gamma_2(\pi_1, \mathcal{L}), \Gamma_2(\pi_2, \mathcal{L})) &\leq \text{diam}(\mathcal{S} \times \mathcal{A}) d_{TV}(\Gamma_2(\pi_1, \mathcal{L}), \Gamma_2(\pi_2, \mathcal{L})) \\
&= \frac{\text{diam}(\mathcal{S} \times \mathcal{A})}{2} \sum_{s' \in \mathcal{S}, a' \in \mathcal{A}} \left| \sum_{s \in \mathcal{S}, a \in \mathcal{A}} \mu(s) P(s'|s, a, \mathcal{L}) (\pi_1(a|s) \pi_1(a'|s') - \pi_2(a|s) \pi_2(a'|s')) \right| \\
&\leq \frac{\text{diam}(\mathcal{S} \times \mathcal{A})}{2} \max_{s, a, \mathcal{L}, s'} P(s'|s, a, \mathcal{L}) \sum_{s, a, s', a'} \mu(s) (\pi_1(a|s) + \pi_2(a|s)) |\pi_1(a'|s') - \pi_2(a'|s')| \\
&\leq \frac{\text{diam}(\mathcal{S} \times \mathcal{A})}{2} \max_{s, a, \mathcal{L}, s'} P(s'|s, a, \mathcal{L}) \sum_{s', a'} |\pi_1(a'|s') - \pi_2(a'|s')| \cdot (1 + 1) \\
&= 2 \text{diam}(\mathcal{S} \times \mathcal{A}) \max_{s, a, \mathcal{L}, s'} P(s'|s, a, \mathcal{L}) \sum_{s'} d_{TV}(\pi_1(s'), \pi_2(s')) \\
&\leq \frac{2 \text{diam}(\mathcal{S} \times \mathcal{A}) \max_{s, a, \mathcal{L}, s'} P(s'|s, a, \mathcal{L}) |\mathcal{S}|}{d_{\min}(\mathcal{A})} D(\pi_1, \pi_2) = \frac{2 \text{diam}(\mathcal{S}) \text{diam}(\mathcal{A}) |\mathcal{S}| c_1}{d_{\min}(\mathcal{A})} D(\pi_1, \pi_2).
\end{aligned} \quad (33)$$

Similarly, we have

$$\begin{aligned}
W_1(\Gamma_2(\pi, \mathcal{L}_1), \Gamma_2(\pi, \mathcal{L}_2)) &\leq \text{diam}(\mathcal{S} \times \mathcal{A}) d_{TV}(\Gamma_2(\pi, \mathcal{L}_1), \Gamma_2(\pi, \mathcal{L}_2)) \\
&= \frac{\text{diam}(\mathcal{S} \times \mathcal{A})}{2} \sum_{s' \in \mathcal{S}, a' \in \mathcal{A}} \left| \sum_{s \in \mathcal{S}, a \in \mathcal{A}} \pi(a|s) \pi(a'|s') (\mu_1(s) P(s'|s, a, \mathcal{L}_1) - \mu_2(s) P(s'|s, a, \mathcal{L}_2)) \right| \\
&\leq \frac{\text{diam}(\mathcal{S} \times \mathcal{A})}{2} \sum_{s' \in \mathcal{S}, a' \in \mathcal{A}} \left| \sum_{s \in \mathcal{S}, a \in \mathcal{A}} \pi(a|s) \pi(a'|s') \mu_1(s) (P(s'|s, a, \mathcal{L}_1) - P(s'|s, a, \mathcal{L}_2)) \right| \\
&\quad + \frac{\text{diam}(\mathcal{S} \times \mathcal{A})}{2} \sum_{s' \in \mathcal{S}, a' \in \mathcal{A}} \left| \sum_{s \in \mathcal{S}, a \in \mathcal{A}} \pi(a|s) \pi(a'|s') (\mu_1(s) - \mu_2(s)) P(s'|s, a, \mathcal{L}_2) \right| \\
&\leq \frac{\text{diam}(\mathcal{S} \times \mathcal{A})}{2} \sum_{s, a, s', a'} \mu_1(s) \pi(a|s) \pi(a'|s') |P(s'|s, a, \mathcal{L}_1) - P(s'|s, a, \mathcal{L}_2)| \\
&\quad + \frac{\text{diam}(\mathcal{S} \times \mathcal{A})}{2} \sum_{s, a, s', a'} |\mu_1(s) - \mu_2(s)| \pi(a|s) \pi(a'|s') P(s'|s, a, \mathcal{L}_2) \\
&\leq \frac{\text{diam}(\mathcal{S}) \text{diam}(\mathcal{A}) |\mathcal{S}| c_2}{2} W_1(\mathcal{L}_1, \mathcal{L}_2) + \frac{\text{diam}(\mathcal{S}) \text{diam}(\mathcal{A}) |\mathcal{S}|}{2} 2d_{TV}(\mathcal{L}_1, \mathcal{L}_2) \\
&\leq \text{diam}(\mathcal{S}) \text{diam}(\mathcal{A}) |\mathcal{S}| \left(\frac{c_2}{2} + \frac{1}{d_{\min}(\mathcal{S} \times \mathcal{A})} \right) W_1(\mathcal{L}_1, \mathcal{L}_2).
\end{aligned} \tag{34}$$

Here μ_1 and μ_2 are the state marginals of \mathcal{L}_1 and \mathcal{L}_2 , respectively.

This completes the proof. Q.E.D.

7.2. Proof of Lemma 3 For notation simplicity, in the following analysis we fix the MDP and omit the notation \mathcal{M} .

We begin by establishing the convergence rate of the synchronous TD updates (27).

LEMMA 5. *Take \tilde{Q}_π^l from (27). Then for any $\delta \in (0, 1)$ and $\epsilon > 0$, with probability at least $1 - \delta$, $\|\tilde{Q}_{\pi_{T(\epsilon, \delta)}}^l - Q^{\pi_{T(\epsilon, \delta)}}\|_\infty \leq \epsilon$ if*

$$l = \Omega \left(\left(\frac{V_{\max} \log \left(\frac{|\mathcal{S}| |\mathcal{A}| V_{\max}}{\delta \beta \epsilon} \right)}{\beta^2 \epsilon^2} \right)^{1/h} + \left(\frac{1}{\beta} \log \frac{V_{\max}}{\epsilon} \right)^{1/(1-h)} \right), \tag{35}$$

where $V_{\max} = R_{\max}/(1 - \gamma)$ and $\beta = (1 - \gamma)/2$.

The proof is adapted from that of [24, Theorem 2], with the max term in the Bellman operator modified to actions sampled from the current policy π . The details are omitted.

[Proof of Lemma 3] First, if $V^*(s') - V^{\pi_{T(\epsilon, \delta/2)}}(s') \leq \epsilon$, then

$$\begin{aligned}
|Q^{\pi_{T(\epsilon, \delta/2)}}(s, a) - Q^*(s, a)| &= \gamma \left| \sum_{s' \in \mathcal{S}} P(s'|s, a) V^{\pi_{T(\epsilon, \delta/2)}}(s') - \sum_{s' \in \mathcal{S}} P(s'|s, a) V^*(s') \right| \\
&\leq \gamma \sum_{s' \in \mathcal{S}} P(s'|s, a) |V^{\pi_{T(\epsilon, \delta/2)}}(s') - V^*(s')| \leq \gamma \epsilon < \epsilon.
\end{aligned} \tag{36}$$

for any $s \in \mathcal{S}$, $a \in \mathcal{A}$. Since Alg is assumed to satisfying the policy-based guarantee, (36) holds with probability at least $1 - \delta$.

In addition, by Lemma 5, whenever l satisfies (28), with probability at least $1 - \delta/2 \geq 1 - \delta$,

$$\|\tilde{Q}_{\pi_{T(\epsilon, \delta/2)}}^l - Q^{\pi_{T(\epsilon, \delta/2)}}\|_\infty \leq (1 - \gamma)\epsilon. \tag{37}$$

Combining (36) and (37), then for any l satisfying (28), with probability at least $1 - 2\delta$, we have

$$\|\tilde{Q}_{\pi_{T(\epsilon, \delta/2)}}^l - Q^*\|_\infty \leq \gamma\epsilon + (1 - \gamma)\epsilon = \epsilon.$$

The above result shows that for any $\delta \in (0, 1)$ and $\epsilon > 0$, after obtaining $T(\epsilon, \delta/2) + |\mathcal{S}||\mathcal{A}|l$ samples (with l satisfying the lower bound (28)) from the simulator, with probability at least $1 - 2\delta$, it outputs an approximate Q -function $\tilde{Q}_{\pi_{T(\epsilon, \delta/2)}}^l$ which satisfies $\|\tilde{Q}_{\pi_{T(\epsilon, \delta/2)}}^l - Q^*\|_\infty \leq \epsilon$. Thus Alg also has a value-based guarantee with parameters

$$\{\tilde{C}_{\mathcal{M}}^{(i)}, \alpha_1^{(i)}, \alpha_2^{(i)}, \alpha_3^{(i)}, \alpha_4^{(i)}\}_{i=1}^{m+3}, \quad (38)$$

specified in (29). Here the first m groups of parameters come from $T(\epsilon, \delta/2)$ while the last three groups of parameters come from $|\mathcal{S}||\mathcal{A}|l$ (with the lower bound (28) of l plugged in here). Q.E.D.

7.3. Proof of $\mathcal{B}_{c,c'} \subseteq \mathcal{F}_{c,c'}$

LEMMA 6. *Suppose that $h: \mathbb{R} \rightarrow \mathbb{R}$ satisfies $h(a) - h(b) \leq c(a - b)$ for any $a \geq b \in \mathbb{R}$. Then the softmax function $\mathbf{softmax}_h$ is c -Lipschitz, i.e., $\|\mathbf{softmax}_h(x) - \mathbf{softmax}_h(y)\|_2 \leq c\|x - y\|_2$ for any $x, y \in \mathbb{R}^n$.*

[Proof of Lemma 6] Notice that $\mathbf{softmax}_h(x) = \mathbf{softmax}(\tilde{h}(x))$, where

$$\mathbf{softmax}(x)_i = \frac{\exp(x_i)}{\sum_{j=1}^n \exp(x_j)} \quad (i = 1, \dots, n)$$

is the standard softmax function and $\tilde{h}(x)_i = h(x_i)$ for $i = 1, \dots, n$. Now since $\mathbf{softmax}$ is 1-Lipschitz continuous (cf. [28, Proposition 4]), and \tilde{h} is c -Lipschitz continuous, we conclude that the composition $\mathbf{softmax} \circ \tilde{h}$ is c -Lipschitz continuous. Q.E.D.

Notice that for a finite set $\mathcal{X} \subseteq \mathbb{R}^k$ and any two (discrete) distributions ν, ν' over \mathcal{X} , we have

$$W_1(\nu, \nu') \leq \text{diam}(\mathcal{X})d_{TV}(\nu, \nu') = \frac{\text{diam}(\mathcal{X})}{2}\|\nu - \nu'\|_1 \leq \frac{\text{diam}(\mathcal{X})\sqrt{|\mathcal{X}|}}{2}\|\nu - \nu'\|_2, \quad (39)$$

where in computing the ℓ_1 -norm, ν, ν' are viewed as vectors of length $|\mathcal{X}|$.

Lemma 6 implies that for any $x, y \in \mathbb{R}^{|\mathcal{X}|}$, when $\mathbf{softmax}_c(x)$ and $\mathbf{softmax}_c(y)$ are viewed as probability distributions over \mathcal{X} , we have

$$W_1(\mathbf{softmax}_c(x), \mathbf{softmax}_c(y)) \leq \frac{\text{diam}(\mathcal{X})\sqrt{|\mathcal{X}|}c}{2}\|x - y\|_2 \leq \frac{\text{diam}(\mathcal{X})|\mathcal{X}|c}{2}\|x - y\|_\infty.$$

LEMMA 7. *Suppose that $h: \mathbb{R} \rightarrow \mathbb{R}$ satisfies $c'(a - b) \leq h(a) - h(b)$ for any $a \leq b \in \mathbb{R}$. Then for any $x \in \mathbb{R}^n$, the distance between the $\mathbf{softmax}_h$ and the $\mathbf{argmax-e}$ mapping is bounded by*

$$\|\mathbf{softmax}_h(x) - \mathbf{argmax-e}(x)\|_2 \leq 2n \exp(-c'\delta),$$

where $\delta = x_{\max} - \max_{x_j < x_{\max}} x_j$, $x_{\max} = \max_{i=1, \dots, n} x_i$, and $\delta := \infty$ when all x_j are equal.

Similar to Lemma 6, Lemma 7 implies that for any $x \in \mathbb{R}^{|\mathcal{X}|}$, viewing $\mathbf{softmax}_h(x)$ as probability distributions over \mathcal{X} leads to

$$W_1(\mathbf{softmax}_h(x), \mathbf{argmax-e}(x)) \leq \text{diam}(\mathcal{X})|\mathcal{X}| \exp(-c\delta).$$

[Proof of Lemma 7] Without loss of generality, assume that $x_1 = x_2 = \dots = x_m = \max_{i=1, \dots, n} x_i = x^* > x_j$ for all $m < j \leq n$. Then

$$\begin{aligned} \mathbf{argmax}\text{-e}(x)_i &= \begin{cases} \frac{1}{m}, & i \leq m, \\ 0, & \text{otherwise.} \end{cases} \\ \mathbf{softmax}_h(x)_i &= \begin{cases} \frac{e^{h(x^*)}}{me^{h(x^*)} + \sum_{j=m+1}^n e^{h(x_j)}}, & i \leq m, \\ \frac{e^{h(x_i)}}{me^{h(x^*)} + \sum_{j=m+1}^n e^{h(x_j)}}, & \text{otherwise.} \end{cases} \end{aligned}$$

Therefore

$$\begin{aligned} \|\mathbf{softmax}_h(x) - \mathbf{argmax}\text{-e}(x)\|_2 &\leq \|\mathbf{softmax}_h(x) - \mathbf{argmax}\text{-e}(x)\|_1 \\ &= m \left(\frac{1}{m} - \frac{e^{h(x^*)}}{me^{h(x^*)} + \sum_{j=m+1}^n e^{h(x_j)}} \right) + \frac{\sum_{i=m+1}^n e^{h(x_i)}}{me^{h(x^*)} + \sum_{j=m+1}^n e^{h(x_j)}} \\ &= \frac{2 \sum_{i=m+1}^n e^{h(x_i)}}{me^{h(x^*)} + \sum_{i=m+1}^n e^{h(x_i)}} = \frac{2 \sum_{i=m+1}^n e^{-c'\delta_i}}{m + \sum_{i=m+1}^n e^{-c\delta_i}} \\ &\leq \frac{2}{m} \sum_{i=m+1}^n e^{-c'\delta_i} \leq \frac{2(n-m)}{m} e^{-c'\delta} \leq 2ne^{-c'\delta}, \end{aligned}$$

with $\delta_i = x^* - x^i$. Q.E.D.

We are now ready to present the proofs of Theorems 2 and 3.

7.4. Proof of Theorems 2 and 3

[Proof of Theorem 2] Here we prove the case when we are using GMF-V and Alg has a value-based guarantee. Define $\hat{\Gamma}_1^k(\mathcal{L}_k) := f_{c,c'}(\hat{Q}_k^*)$. In the following, $\pi = f_{c,c'}(Q_{\mathcal{L}})$ is understood as the policy π with $\pi(s) = f_{c,c'}(Q_{\mathcal{L}}(s, \cdot))$. Let \mathcal{L}^* be the population state-action pair in a stationary NE of (GMFG). Then $\pi_k = \hat{\Gamma}_1^k(\mathcal{L}_k)$. Denoting $d := d_1 d_2 + d_3$, we see

$$\begin{aligned} W_1(\tilde{\mathcal{L}}_{k+1}, \mathcal{L}^*) &= W_1(\Gamma_2(\pi_k, \mathcal{L}_k), \Gamma_2(\Gamma_1(\mathcal{L}^*), \mathcal{L}^*)) \\ &\leq W_1(\Gamma_2(\Gamma_1(\mathcal{L}_k), \mathcal{L}_k), \Gamma_2(\Gamma_1(\mathcal{L}^*), \mathcal{L}^*)) + W_1(\Gamma_2(\Gamma_1(\mathcal{L}_k), \mathcal{L}_k), \Gamma_2(\hat{\Gamma}_1^k(\mathcal{L}_k), \mathcal{L}_k)) \\ &\leq W_1(\Gamma(\mathcal{L}_k), \Gamma(\mathcal{L}^*)) + d_2 D(\Gamma_1(\mathcal{L}_k), \hat{\Gamma}_1^k(\mathcal{L}_k)) \\ &\leq (d_1 d_2 + d_3) W_1(\mathcal{L}_k, \mathcal{L}^*) + d_2 D(\mathbf{argmax}\text{-e}(Q_{\mathcal{L}_k}^*), f_{c,c'}(\hat{Q}_k^*)) \\ &\leq d W_1(\mathcal{L}_k, \mathcal{L}^*) + d_2 D(f_{c,c'}(\hat{Q}_k^*), f_{c,c'}(Q_{\mathcal{L}_k}^*)) \\ &\quad + d_2 D(\mathbf{argmax}\text{-e}(Q_{\mathcal{L}_k}^*), f_{c,c'}(Q_{\mathcal{L}_k}^*)) \\ &\leq d W_1(\mathcal{L}_k, \mathcal{L}^*) + \frac{cd_2 \text{diam}(\mathcal{A})|\mathcal{A}|}{2} \|\hat{Q}_k^* - Q_{\mathcal{L}_k}^*\|_{\infty} + d_2 D(\mathbf{argmax}\text{-e}(Q_{\mathcal{L}_k}^*), f_{c,c'}(Q_{\mathcal{L}_k}^*)). \end{aligned}$$

Since $\mathcal{L}_k \in S_{\epsilon}$ by the projection step, by Lemma 7 and the algorithm Alg has a policy-based guarantee, with the choice of $T_k = T_{\mathcal{M}_{\mathcal{L}_k}}(\delta_k, \epsilon_k)$, we have, with probability at least $1 - 2\delta_k$,

$$W_1(\tilde{\mathcal{L}}_{k+1}, \mathcal{L}^*) \leq d W_1(\mathcal{L}_k, \mathcal{L}^*) + \frac{cd_2 \text{diam}(\mathcal{A})|\mathcal{A}|}{2} \epsilon_k + d_2 \text{diam}(\mathcal{A})|\mathcal{A}| e^{-c'\phi(\epsilon)}. \quad (40)$$

Finally, with probability at least $1 - 2\delta_k$,

$$\begin{aligned} W_1(\mathcal{L}_{k+1}, \mathcal{L}^*) &\leq W_1(\tilde{\mathcal{L}}_{k+1}, \mathcal{L}^*) + W_1(\tilde{\mathcal{L}}_{k+1}, \mathbf{Proj}_{S_{\epsilon}}(\tilde{\mathcal{L}}_{k+1})) \\ &\leq d W_1(\mathcal{L}_k, \mathcal{L}^*) + \frac{cd_2 \text{diam}(\mathcal{A})|\mathcal{A}|}{2} \epsilon_k + d_2 \text{diam}(\mathcal{A})|\mathcal{A}| e^{-c'\phi(\epsilon)} + \epsilon. \end{aligned}$$

This implies that with probability at least $1 - 2 \sum_{k=0}^{K-1} \delta_k$,

$$\begin{aligned} W_1(\mathcal{L}_K, \mathcal{L}^*) &\leq d^K W_1(\mathcal{L}_0, \mathcal{L}^*) + \frac{cd_2 \text{diam}(\mathcal{A}) |\mathcal{A}|}{2} \sum_{k=0}^{K-1} d^{K-k} \epsilon_k \\ &\quad + \frac{(d_2 \text{diam}(\mathcal{A}) |\mathcal{A}| e^{-c' \phi(\epsilon)} + \epsilon)(1 - d^K)}{1 - d}. \end{aligned} \quad (41)$$

Since ϵ_k is summable, we have $\sup_{k \geq 0} \epsilon_k < \infty$,

$$\sum_{k=0}^{K-1} d^{K-k} \epsilon_k \leq \frac{\sup_{k \geq 0} \epsilon_k}{1 - d} d^{\lfloor (K-1)/2 \rfloor} + \sum_{k=\lceil (K-1)/2 \rceil}^{\infty} \epsilon_k.$$

Now plugging in $K = K_{\epsilon, \eta}$, with the choice of δ_k and $c = \frac{\log(1/\epsilon)}{\phi(\epsilon)}$, and noticing that $d \in [0, 1)$, we have with probability at least $1 - 2\delta$,

$$\begin{aligned} W_1(\mathcal{L}_{K_{\epsilon, \eta}}, \mathcal{L}^*) &\leq d^{K_{\epsilon, \eta}} W_1(\mathcal{L}_0, \mathcal{L}^*) \\ &\quad + \frac{cd_2 \text{diam}(\mathcal{A}) |\mathcal{A}|}{2} \left(\frac{\sup_{k \geq 0} \epsilon_k}{1 - d} d^{\lfloor (K_{\epsilon, \eta}-1)/2 \rfloor} + \sum_{k=\lceil (K_{\epsilon, \eta}-1)/2 \rceil}^{\infty} \epsilon_k \right) \\ &\quad + \frac{(d_2 \text{diam}(\mathcal{A}) |\mathcal{A}| + 1)\epsilon}{1 - d}. \end{aligned} \quad (42)$$

Setting $\epsilon_k = (k+1)^{-(1+\eta)}$, then when $K_{\epsilon, \eta} \geq 2(\log_d(\epsilon/c) + 1)$,

$$\frac{\sup_{k \geq 0} \epsilon_k}{1 - d} d^{\lfloor (K_{\epsilon, \eta}-1)/2 \rfloor} \leq \frac{\epsilon/c}{1 - d}.$$

Similarly, when $K_{\epsilon, \eta} \geq 2(\eta\epsilon/c)^{-1/\eta}$,

$$\sum_{k=\lceil \frac{K_{\epsilon, \eta}-1}{2} \rceil}^{\infty} \epsilon_k \leq \epsilon/c.$$

Finally, when $K_{\epsilon, \eta} \geq \log_d(\epsilon/(\text{diam}(\mathcal{S})\text{diam}(\mathcal{A})))$, $d^{K_{\epsilon, \eta}} W_1(\mathcal{L}_0, \mathcal{L}^*) \leq \epsilon$, since $W_1(\mathcal{L}_0, \mathcal{L}^*) \leq \text{diam}(\mathcal{S} \times \mathcal{A}) = \text{diam}(\mathcal{S})\text{diam}(\mathcal{A})$.

In summary, if $K_{\epsilon, \eta} = \lceil 2 \max\{(\eta\epsilon/c)^{-1/\eta}, \log_d(\epsilon/\max\{\text{diam}(\mathcal{S})\text{diam}(\mathcal{A}), c\}) + 1\} \rceil$, then with probability at least $1 - 2\delta$,

$$W_1(\mathcal{L}_{K_{\epsilon, \eta}}, \mathcal{L}^*) \leq \left(1 + \frac{d_2 \text{diam}(\mathcal{A}) |\mathcal{A}| (2-d)}{2(1-d)} + \frac{(d_2 \text{diam}(\mathcal{A}) |\mathcal{A}| + 1)}{1-d} \right) \epsilon = O(\epsilon). \quad (43)$$

Finally, if we are using GMF-V and have assumed that Alg satisfies a value-based guarantee with parameters $\{C_{\mathcal{M}}^{(i)}, \alpha_1^{(i)}, \alpha_2^{(i)}, \alpha_3^{(i)}, \alpha_4^{(i)}\}_{i=1}^m$, plugging in ϵ_k and δ_k into $T_{\mathcal{M}_{\mathcal{L}}}(\delta_k, \epsilon_k)$, and noticing that $k \leq K_{\epsilon, \eta}$ and $\sum_{k=0}^{K_{\epsilon, \eta}-1} (k+1)^\alpha \leq \frac{K_{\epsilon, \eta}^{\alpha+1}}{\alpha+1}$, we have

$$\begin{aligned} T &= \sum_{k=0}^{K_{\epsilon, \eta}} \sum_{i=1}^m C_{\mathcal{M}}^{(i)} \left(\frac{1}{\epsilon_k} \right)^{\alpha_1^{(i)}} \left(\log \frac{1}{\epsilon_k} \right)^{\alpha_2^{(i)}} \left(\frac{1}{\delta_k} \right)^{\alpha_3^{(i)}} \left(\log \frac{1}{\delta_k} \right)^{\alpha_4^{(i)}} \\ &= \sum_{k=0}^{K_{\epsilon, \eta}} \sum_{i=1}^m (1+\eta)^{\alpha_2^{(i)}} C_{\mathcal{M}}^{(i)} (k+1)^{\alpha_1^{(i)}(1+\eta)} (\log(k+1))^{\alpha_2^{(i)}} (K_{\epsilon, \eta}/\delta)^{\alpha_3^{(i)}} (\log(K_{\epsilon, \eta}/\delta))^{\alpha_4^{(i)}} \\ &\leq \sum_{i=1}^m \frac{(1+\eta)^{\alpha_2^{(i)}}}{\alpha_1^{(i)}(1+\eta) + 1} C_{\mathcal{M}}^{(i)} K_{\epsilon, \eta}^{\alpha_1^{(i)}(1+\eta)+1} (\log(K_{\epsilon, \eta} + 1))^{\alpha_2^{(i)}} (K_{\epsilon, \eta}/\delta)^{\alpha_3^{(i)}} (\log(K_{\epsilon, \eta}/\delta))^{\alpha_4^{(i)}} \\ &\leq \sum_{i=1}^m \frac{2^{\alpha_2^{(i)}}}{2\alpha_1^{(i)} + 1} C_{\mathcal{M}}^{(i)} K_{\epsilon, \eta}^{2\alpha_1^{(i)}+1} (K_{\epsilon, \eta}/\delta)^{\alpha_3^{(i)}} (\log(K_{\epsilon, \eta}/\delta))^{\alpha_2^{(i)} + \alpha_4^{(i)}}, \end{aligned} \quad (44)$$

which completes the proof of the value-based case.

[Proof of Theorem 3] If we use GMF-P and assume that Alg has the policy-based guarantee, then by Lemma 3,

$$\mathbb{P}\left(\left\|\tilde{Q}_{\hat{\pi}_k}^{l_k} - Q_{\mathcal{L}_k}^*\right\|_{\infty} > \epsilon\right) \leq 2\delta. \quad (45)$$

Hence one can simply replace \hat{Q}_k^* by $\tilde{Q}_{\hat{\pi}_k}^{l_k}$ in the proof of Theorem 2, and obtain the same bound on $W_1(\mathcal{L}_{K_{\epsilon,\eta}}, \mathcal{L}^*)$ (cf. (43)). The only difference is that in each iteration, the required number of samples $T_{\mathcal{M}_{\mathcal{L}}}$ now has parameters $\{\tilde{C}_{\mathcal{M}}^{(i)}, \alpha_1^{(i)}, \alpha_2^{(i)}, \alpha_3^{(i)}, \alpha_4^{(i)}\}_{i=1}^{m+3}$ as defined in Lemma 3. Hence repeating the proof of (44) leads to (30). Q.E.D.

8. Experiments In this section, we report the performance of the proposed GMF-V-Q Algorithm and GMF-P-TRPO Algorithm with an equilibrium pricing model (see Section 2.3). The objectives of the experiments include 1) testing the convergence and stability of both GMF-V-Q and GMF-P-TRPO in the GMFG setting, 2) empirically verifying the contractive property of mapping Γ , and 3) comparing GMF-V-Q and GMF-P-TRPO with existing multi-agent reinforcement learning algorithms, including the Independent Learner (IL) algorithm Tan [77], Hu et al. [44] and the MF-Q³ algorithm Yang et al. [88]. Another set of experiments for the repeated auction model (see Section 2.3) is demonstrated in the short version Guo et al. [38].

8.1. Set-up and parameter configuration We introduce two testing environments in our numerical experiments, one is the GMFG environment with a continuum of agents (i.e., infinite number of agents) described in Section 2.3 and the other one is an N -player environment with a weak simulator.

Equilibrium price as an N -player game. We also consider an N -player game version of the equilibrium price model, which is the GMFG version described above with an N -player weak simulator oracle as described in Section 6. In particular, Take N companies. At each time t , company i decides a quantity q_t^i for production and a quantity h_t^i to replenish the inventory. Let s_t^i denote the current inventory level of company i at time t . Then similar to Section 2.3, the inventory level evolves according to

$$s_{t+1}^i = s_t^i - \min\{q_t^i, s_t^i\} + h_t^i$$

and the reward of company i at time t is given by

$$r_t^i = (p_t - c_0)q_t^i - c_1(q_t^i)^2 - c_2h_t^i - (c_2 + c_3)\max\{q_t^i - s_t^i, 0\} - c_4s_t^i.$$

Here p_t , the price of the product at time t , is determined according to the supply-demand equilibrium on the market. The total supply is $\sum_{i=1}^N q_t^i$, while the total demand is assumed to be $d_N p_t^{-\sigma}$, where $d_N = dN$ is supposed to be linearly growing as N grows, i.e., the number of customers grows proportionally to the number of producers in the market. Then by equating supply and demand, we obtain that

$$\frac{1}{N} \sum_{i=1}^N q_t^i = d p_t^{-\sigma},$$

and by taking the limit $N \rightarrow \infty$, we obtain the mean-field counterpart (11).

In this setting, accordingly, we test the performance of GMF-VW-Q, which is GMF-VQ (Algorithm 5) with synchronous Q-learning and the standard **softmax** operator (cf. Algorithm 7) and

³ Note that MF-Q is designed for global states and coupled local actions, while in our equilibrium price example we have coupled local (private) states and decoupled local actions. To suit this setting, we adapt MF-Q by replacing the mean-field action term with the mean-field state term.

GMF-PW-TRPO, which is GMF-PW (Algorithm 6) with TRPO and the standard **softmax** operator.⁴

Algorithm 7 Q-learning for GMFGs (GMF-VW-Q): weak simulator

- 1: **Input:** Initial \mathcal{L}_0 , ϵ -net S_ϵ , tolerances $\epsilon_k, \delta_k > 0, k = 0, 1, \dots$
 - 2: **for** $k = 0, 1, \dots$ **do**
 - 3: Perform Q-learning with hyper-parameters in Lemma 2 for $T_k = T_{\mathcal{M}_{\mathcal{L}_k}}(\epsilon_k, \delta_k)$ iterations to find the approximate Q-function $\hat{Q}_k^* = \hat{Q}^{T_k}$ of the MDP $\mathcal{M}_{\mathcal{L}_k}$.
 - 4: Compute $\pi_k \in \Pi$ with $\pi_k(s) = \mathbf{softmax}_c(\hat{Q}_k^*(s, \cdot))$.
 - 5: **for** $i = 1, 2, \dots, N$ **do**
 - 6: Sample $s_i \stackrel{\text{i.i.d.}}{\sim} \mu_k$, then obtain s'_i i.i.d. from $\mathcal{G}_W(s_i, \pi_k, \mathcal{L}_k)$ and $a'_i \stackrel{\text{i.i.d.}}{\sim} \pi_k(s'_i)$.
 - 7: **end for**
 - 8: Compute \mathcal{L}_{k+1} with $\mathcal{L}_{k+1}(s, a) = \frac{1}{N} \sum_{i=1}^N \mathbf{I}_{s'_i=s, a'_i=a}$.
 - 9: **end for**
-

Parameters. The model parameters are (unless otherwise specified): $\gamma = 0.2$, $d = 50$ and $\sigma = 2$. $S = Q = H = 10$ and hence $|\mathcal{S}| = 10$ and $|\mathcal{A}| = 100$. $c_0 = 0.5$, $c_1 = 0.1$, $c_2 = 0.5$, $c_3 = 0.2$ and $c_4 = 0.2$.

The algorithm parameters are (unless otherwise specified): the temperature parameter is set as $c = 4.0$ and the learning rate is set as $\eta = 0.01$ ⁵. For simplicity, we set the inner iteration T_k to be $100 \times |\mathcal{S}| \times |\mathcal{A}|$. The 90%-confidence intervals are calculated with 20 sample paths.

8.2. Performance evaluation in the GMFG setting. Our experiments show that GMF-V-Q and GMF-P-TRPO Algorithms are efficient and robust.

Performance metric. We adopt the following metric to measure the difference between a given policy π and an NE (here $\epsilon_0 > 0$ is a safeguard, and is taken as 0.1 in the experiments):

$$C_{MF}(\pi) = \frac{\max_{\pi'} \mathbb{E}_{s \sim \mu} [V(s, \pi', \mathcal{L})] - V(s, \pi, \mathcal{L})}{|\max_{\pi'} \mathbb{E}_{s \sim \mu} [V(s, \pi', \mathcal{L})]| + \epsilon_0}.$$

Here μ is the invariant distribution of the transition matrix P^π , where $P^\pi(s, s') = \sum_{a \in \mathcal{A}} P(s'|s, a)\pi(a|s)$ for $s, s' \in \mathcal{S}$, and $\mathcal{L}(s, a) = \mu(s)\pi(a|s)$ for $s, a \in \mathcal{S} \times \mathcal{A}$. Note that in the equilibrium product pricing model we are considering here, the transition model P is independent of the mean-field term \mathcal{L} , and hence we write $P(s'|s, a) = P(s'|s, a, \mathcal{L})$. In general, an additional mean-field matching error term needs to be added into the definition of $C_{MF}(\pi)$. Clearly $C_{MF}(\pi) \geq 0$, and $C_{MF}(\pi^*) = 0$ if and only if (π^*, \mathcal{L}^*) is an NE where \mathcal{L}^* is the invariant distribution of P^{π^*} . A similar metric without normalization has been adopted in Cui and Koepl [18].

Contractiveness of mapping Γ . As explained in Remark 1 from Section 3, the contractiveness property of Γ is the key for establishing the uniqueness of MFG solution and hence the convergence of the GMFG algorithm. To empirically verify whether this property holds for the equilibrium price example, we plot the value of $\frac{\|\Gamma(\mathcal{L}_1) - \Gamma(\mathcal{L}_2)\|_1}{\|\mathcal{L}_1 - \mathcal{L}_2\|_1}$ for randomly generated state-action distributions \mathcal{L}_1 and \mathcal{L}_2 . Technically speaking, Γ is contractive and there exists a unique MFG solution if the value of $\frac{\|\Gamma(\mathcal{L}_1) - \Gamma(\mathcal{L}_2)\|_1}{\|\mathcal{L}_1 - \mathcal{L}_2\|_1}$ is smaller than one for all choices of \mathcal{L}_1 and \mathcal{L}_2 .

⁴ For the sake of brevity, we omit the algorithm frame for GMF-PW-TRPO.

⁵ Lemma 2 indicates that the learning rate should be inverse proportional to the current visitation number of a given state-action pair, we observe that constant learning rate works well in practice which is easier to implement.

We observe from Figure 1 that, with various choices of different model parameters, the quantity $\frac{\|\Gamma(\mathcal{L}_1) - \Gamma(\mathcal{L}_2)\|_1}{\|\mathcal{L}_1 - \mathcal{L}_2\|_1}$ is always smaller than 0.3 indicating that Γ is contractive.

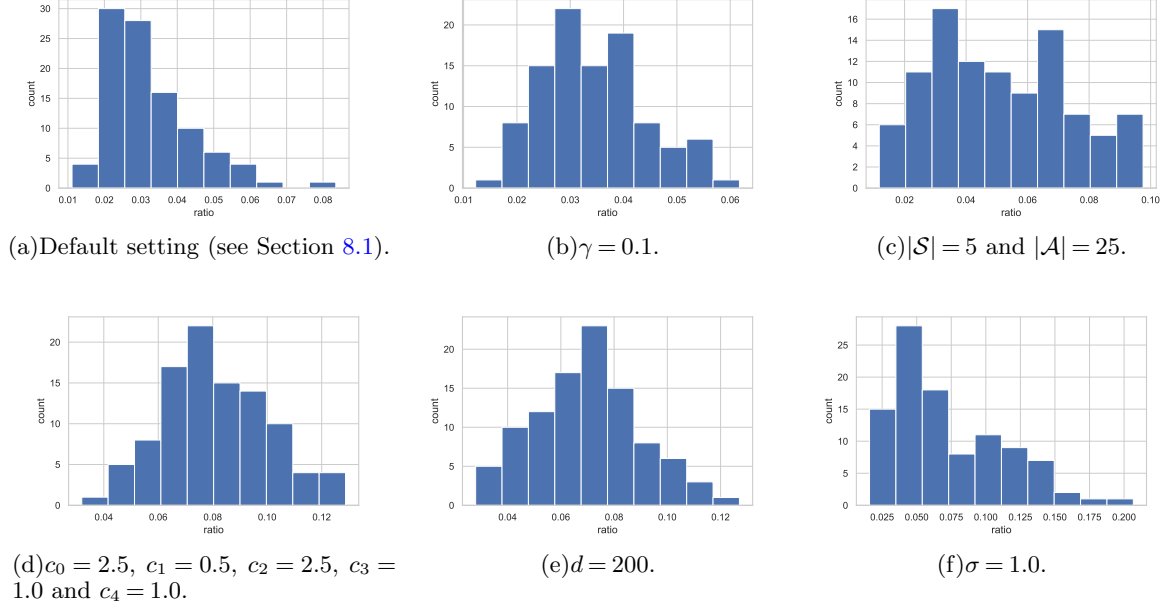


FIGURE 1. Histogram of $\frac{\|\Gamma(\mathcal{L}_1) - \Gamma(\mathcal{L}_2)\|_1}{\|\mathcal{L}_1 - \mathcal{L}_2\|_1}$ under various settings (\mathcal{L}_1 and \mathcal{L}_2 are randomly sampled according to the uniform distribution).

Convergence and stability. Both GMF-V-Q and GMF-P-TRPO are efficient and robust. First, both GMF-V-Q and GMF-P-TRPO converge within about 5 outer iterations; secondly, as the number of inner iterations increases, the error decreases (Figure 2); and finally, the convergence is robust with respect to both the change of number of states and actions (Figure 3). The performance of GMF-V-Q is (slightly) more stable than GMF-P-TRPO with a smaller variance across 20 repeated experiments (see Figure 2(a) versus Figure 2(b) or Figure 3(a) versus Figure 3(b)). This is due to the fact that GMF-P-TRPO uses asynchronous updates, which leads to slightly less stable performance compared to GMF-V-Q, which uses synchronous updates.

In contrast, the Naive algorithms, i.e., GMF-V-Q without smoothing (denoted as GMF-V-Q-nonsmoothing) and GMF-P-TRPO without smoothing (denoted as GMF-P-TRPO-nonsmoothing), do not converge even with 50 outer iterations and $200 \times |\mathcal{S}| \times |\mathcal{A}|$ inner iterations within each outer iteration. In particular, GMF-V-Q-nonsmoothing and GMF-P-TRPO-nonsmoothing present different unstable behaviors (see Figure 4). The joint distribution \mathcal{L}_t from GMF-V-Q-nonsmoothing keeps fluctuating (Figure 4(a)) whereas the joint distribution \mathcal{L}_t from GMF-P-TRPO (without smoothing) is trapped around the initialization which is far away from the true equilibrium distribution (Figure 4(b)).

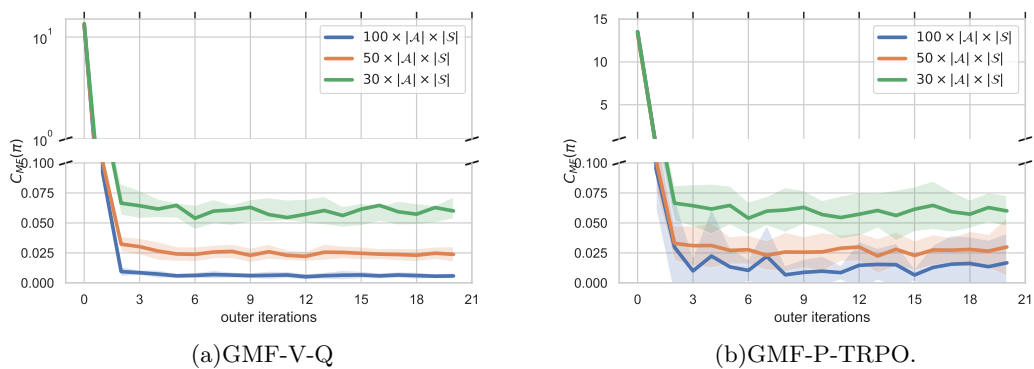


FIGURE 2. Convergence with different number of inner iterations ($|\mathcal{A}| = 100$ and $|\mathcal{S}| = 10$).

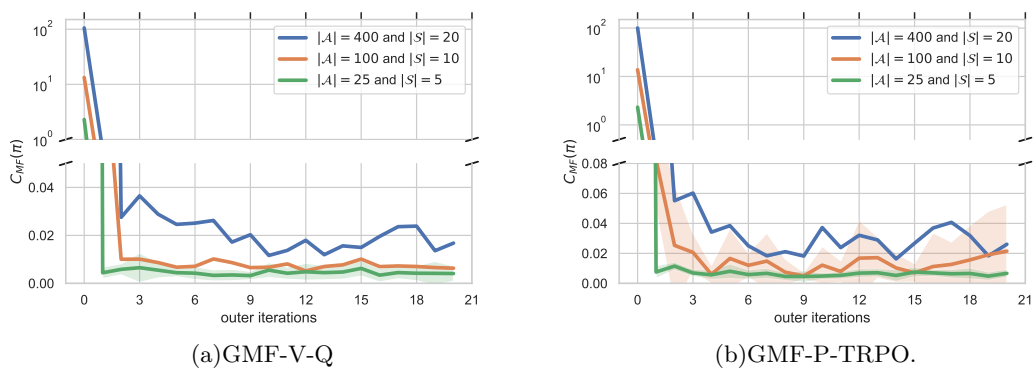


FIGURE 3. Convergence with different size of state space and action space .

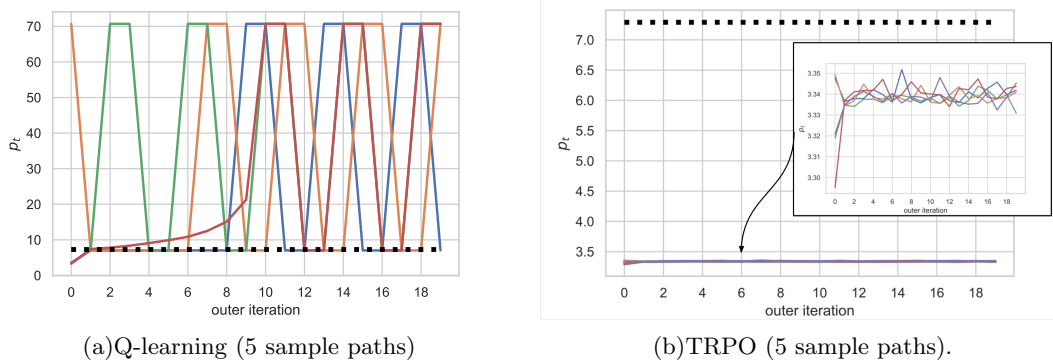
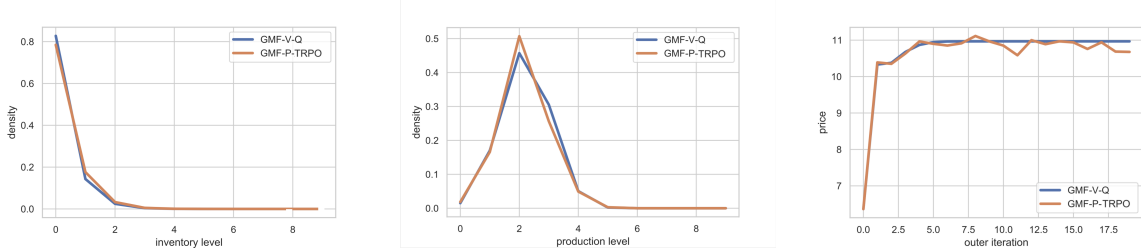



FIGURE 4. Fluctuations of algorithms without smoothing (Dotted black line: theoretical value of the equilibrium price).

Model verification and interpretation of equilibrium scenario. In Figures 5 and 6, we run both algorithms for 20 outer iterations with the same number of inner iterations ($100,000 = 100 \times |\mathcal{A}| \times |\mathcal{S}|$) within each outer iteration. The final equilibrium inventory distribution and production distribution from both algorithms are close to each other.



(a)Equilibrium inventory distribu- (b)Equilibrium production distribu- (c)Equilibrium price at the end of
tion. tion. each outer iteration.

FIGURE 5. GMF-V-Q versus GMF-P-TRPO ($\sigma = 1.3$ and one trajectory).



(a)Equilibrium inventory distribu- (b)Equilibrium production distribu- (c)Equilibrium price at the end of
tion. tion. each outer iteration.

FIGURE 6. GMF-V-Q versus GMF-P-TRPO ($\sigma = 2.0$ and one trajectory).

Note that the demand elasticity σ captures how sensitive demand for a product is compared to the changes in other economic factors, such as price or income. When σ is increased from 1.3 to 2.0 indicating that the demand is more sensitive to price rise, the equilibrium price decreases from 10.9 to 7.3 (see Figures 5(c) and 6(c)) and the distribution of the equilibrium production level is centered towards smaller values (see Figures 5(b) and 6(b)). The equilibrium inventory level has a huge mass at 0 when $\sigma = 1.3$. This implies that producers do not keep large inventories and pay the inventory cost in the equilibrium. On the other hand, the equilibrium inventory is more uniformly distributed when $\sigma = 2$.

8.3. Performance evaluation in the N-player setting

Performance metric. Similar to the performance metric introduced in Section 8.2 for the GMFG setting, we adopt the following metric to measure the difference between a given policy π and an NE under THE N-player setting (here $\epsilon_0 > 0$ is a safeguard, and is taken as 0.1 in the experiments):

$$C(\boldsymbol{\pi}) = \frac{1}{N|\mathcal{S}|^N} \sum_{i=1}^N \sum_{\mathbf{s} \in \mathcal{S}^N} \frac{\max_{\pi^i} V^i(\mathbf{s}, (\boldsymbol{\pi}^{-i}, \pi^i)) - V^i(\mathbf{s}, \boldsymbol{\pi})}{|\max_{\pi^i} V^i(\mathbf{s}, (\boldsymbol{\pi}^{-i}, \pi^i))| + \epsilon_0}.$$

Clearly $C(\boldsymbol{\pi}) \geq 0$, and $C(\boldsymbol{\pi}^*) = 0$ if and only if $\boldsymbol{\pi}^*$ is an NE. Policy $\arg \max_{\pi_i} V_i(\mathbf{s}, (\boldsymbol{\pi}^{-i}, \pi_i))$ is called the best response to $\boldsymbol{\pi}^{-i}$. A similar metric without normalization has been adopted in P erolat et al. [67].

Existing algorithms for N -player games. To test the effectiveness of GMF-VW-Q for approximating N -player games, we next compare GMF-VW-Q with the IL algorithm and the MF-Q algorithm. The IL algorithm Tan [77] considers N independent players and each player solves a decentralized reinforcement learning problem ignoring other players in the system. The MF-Q algorithm Yang et al. [88] extends the NASH-Q Learning algorithm for the N -player game introduced in Hu and Wellman [43], adds the aggregate actions ($\bar{\mathbf{a}}_{-i} = \frac{\sum_{j \neq i} a_j}{N-1}$) from the opponents, and works for the class of games where the interactions are only through the average actions of N players.

Results and analysis. Our experiment (Figure 7) shows that GMF-VW-Q and GMF-PW-TRPO achieve similar performance, and both of them are superior in terms of convergence rate, accuracy, and stability for approximating an N -player game. In general, both algorithms converge faster than IL and MF-Q and achieve the smallest errors.

For instance, when $N = 20$, IL Algorithm converges with the largest error 0.220. The error from MF-Q is 0.101, smaller than IL but still bigger than the error from GMF-VW-Q. The GMF-VW-Q and GMF-PW-TRPO converge with the lowest error 0.065. Moreover, as N increases, the error of GMF-VW-Q and GMF-PW-TRPO decrease while the errors of both MF-Q and IL increase significantly. As $|\mathcal{S}|$ and $|\mathcal{A}|$ increase, GMF-VW-Q and GMF-PW-TRPO are robust with respect to this increase of dimensionality, while both MF-Q and IL clearly suffer from the increase of the dimensionality with decreased convergence rate and accuracy. Therefore, GMF-VW-Q and GMF-PW-TRPO are more scalable than IL and MF-Q, when the system is complex and the number of players N is large.




FIGURE 7. Learning accuracy based on $C(\boldsymbol{\pi})$.

9. Extension: Existence and uniqueness for non-stationary NE of GMFGs In this section, we describe the setting of non-stationary NE for GMFGs and establish the corresponding results of existence and uniqueness.

DEFINITION 4 (NE FOR GMFGs). In (GMFG), a player-population profile $(\boldsymbol{\pi}^*, \mathcal{L}^*) := (\{\pi_t^*\}_{t=0}^\infty, \{\mathcal{L}_t^*\}_{t=0}^\infty)$ is called an NE if

- (Single player side) Fix \mathcal{L}^* , for any policy sequence $\boldsymbol{\pi} := \{\pi_t\}_{t=0}^\infty$ and initial state $s \in \mathcal{S}$,

$$V(s, \boldsymbol{\pi}^*, \mathcal{L}^*) \geq V(s, \boldsymbol{\pi}, \mathcal{L}^*). \quad (46)$$

- (Population side) $\mathbb{P}_{s_t, a_t} = \mathcal{L}_t^*$ for all $t \geq 0$, where $\{s_t, a_t\}_{t=0}^\infty$ is the dynamics under the policy sequence $\boldsymbol{\pi}^*$ starting from $s_0 \sim \mu_0^*$, with $a_t \sim \pi_t^*(s_t, \mu_t^*)$, $s_{t+1} \sim P(\cdot | s_t, a_t, \mathcal{L}_t^*)$, and μ_t^* being the population state marginal of \mathcal{L}_t^* .

Step A. Fix $\mathcal{L} := \{\mathcal{L}_t\}_{t=0}^\infty$, (GMFG) becomes the classical optimization problem. Indeed, with \mathcal{L} fixed, the population state distribution sequence $\boldsymbol{\mu} := \{\mu_t\}_{t=0}^\infty$ is also fixed, hence the space of admissible policies is reduced to the single-player case. Solving (GMFG) is now reduced to finding a policy sequence $\boldsymbol{\pi}_{t,\mathcal{L}} \in \Pi := \{\pi \mid \pi : \mathcal{S} \rightarrow \mathcal{P}(\mathcal{A})\}$ over all admissible $\boldsymbol{\pi}_{\mathcal{L}} = \{\pi_{t,\mathcal{L}}\}_{t=0}^\infty$, to maximize

$$V(s, \boldsymbol{\pi}_{\mathcal{L}}, \mathcal{L}) := \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t, \mathcal{L}_t) \mid s_0 = s \right],$$

subject to $s_{t+1} \sim P(s_t, a_t, \mathcal{L}_t), \quad a_t \sim \pi_{t,\mathcal{L}}(s_t).$

Notice that with \mathcal{L} fixed, one can safely suppress the dependency on μ_t in the admissible policies. Moreover, given this fixed \mathcal{L} sequence and the solution $\boldsymbol{\pi}_{\mathcal{L}}^* := \{\pi_{t,\mathcal{L}}^*\}_{t=0}^\infty$, one can define a mapping from the fixed population distribution sequence \mathcal{L} to an optimal randomized policy sequence. That is,

$$\Gamma_1 : \{\mathcal{P}(\mathcal{S} \times \mathcal{A})\}_{t=0}^\infty \rightarrow \{\Pi\}_{t=0}^\infty,$$

such that $\boldsymbol{\pi}_{\mathcal{L}}^* = \Gamma_1(\mathcal{L})$. Note that this $\boldsymbol{\pi}_{\mathcal{L}}^*$ sequence satisfies the single player side condition in Definition 4 for the population state-action pair sequence \mathcal{L} . That is, $V(s, \boldsymbol{\pi}_{\mathcal{L}}^*, \mathcal{L}) \geq V(s, \boldsymbol{\pi}, \mathcal{L})$, for any policy sequence $\boldsymbol{\pi} = \{\pi_t\}_{t=0}^\infty$ and any initial state $s \in \mathcal{S}$.

Accordingly, a similar feedback regularity condition is needed in this step.

ASSUMPTION 3. *There exists a constant $d_1 \geq 0$, such that for any $\mathcal{L}, \mathcal{L}' \in \{\mathcal{P}(\mathcal{S} \times \mathcal{A})\}_{t=0}^\infty$,*

$$D(\Gamma_1(\mathcal{L}), \Gamma_1(\mathcal{L}')) \leq d_1 \mathcal{W}_1(\mathcal{L}, \mathcal{L}'), \quad (47)$$

where

$$\begin{aligned} D(\boldsymbol{\pi}, \boldsymbol{\pi}') &:= \sup_{s \in \mathcal{S}} \mathcal{W}_1(\boldsymbol{\pi}(s), \boldsymbol{\pi}'(s)) = \sup_{s \in \mathcal{S}} \sup_{t \in \mathbb{N}} \mathcal{W}_1(\pi_t(s), \pi'_t(s)), \\ \mathcal{W}_1(\mathcal{L}, \mathcal{L}') &:= \sup_{t \in \mathbb{N}} \mathcal{W}_1(\mathcal{L}_t, \mathcal{L}'_t), \end{aligned} \quad (48)$$

and \mathcal{W}_1 is the ℓ_1 -Wasserstein distance between probability measures.

Step B. Based on the analysis in Step A and $\boldsymbol{\pi}_{\mathcal{L}}^* = \{\pi_{t,\mathcal{L}}^*\}_{t=0}^\infty$, update the initial sequence \mathcal{L} to \mathcal{L}' following the controlled dynamics $P(\cdot \mid s_t, a_t, \mathcal{L}_t)$.

Accordingly, for any admissible policy sequence $\boldsymbol{\pi} \in \{\Pi\}_{t=0}^\infty$ and a joint population state-action pair sequence $\mathcal{L} \in \{\mathcal{P}(\mathcal{S} \times \mathcal{A})\}_{t=0}^\infty$, define a mapping $\Gamma_2 : \{\Pi\}_{t=0}^\infty \times \{\mathcal{P}(\mathcal{S} \times \mathcal{A})\}_{t=0}^\infty \rightarrow \{\mathcal{P}(\mathcal{S} \times \mathcal{A})\}_{t=0}^\infty$ as follows:

$$\Gamma_2(\boldsymbol{\pi}, \mathcal{L}) := \hat{\mathcal{L}} = \{\mathbb{P}_{s_t, a_t}\}_{t=0}^\infty, \quad (49)$$

where $s_{t+1} \sim \mu_t P(\cdot \mid \cdot, a_t, \mathcal{L}_t)$, $a_t \sim \pi_t(s_t)$, $s_0 \sim \mu_0$, and μ_t is the population state marginal of \mathcal{L}_t .

One also needs a similar assumption in this step.

ASSUMPTION 4. *There exist constants $d_2, d_3 \geq 0$, such that for any admissible policy sequences $\boldsymbol{\pi}, \boldsymbol{\pi}^1, \boldsymbol{\pi}^2$ and joint distribution sequences $\mathcal{L}, \mathcal{L}^1, \mathcal{L}^2$,*

$$\mathcal{W}_1(\Gamma_2(\boldsymbol{\pi}^1, \mathcal{L}), \Gamma_2(\boldsymbol{\pi}^2, \mathcal{L})) \leq d_2 D(\boldsymbol{\pi}^1, \boldsymbol{\pi}^2), \quad (50)$$

$$\mathcal{W}_1(\Gamma_2(\boldsymbol{\pi}, \mathcal{L}^1), \Gamma_2(\boldsymbol{\pi}, \mathcal{L}^2)) \leq d_3 \mathcal{W}_1(\mathcal{L}^1, \mathcal{L}^2). \quad (51)$$

Similarly, Assumption 4 can be reduced to Lipschitz continuity and boundedness of the transition dynamics P under certain conditions.

Step C. Repeat Step A and Step B until \mathcal{L}' matches \mathcal{L} .

This step is to take care of the population side condition. To ensure the convergence of the combined step A and step B, it suffices if $\Gamma : \{\mathcal{P}(\mathcal{S} \times \mathcal{A})\}_{t=0}^{\infty} \rightarrow \{\mathcal{P}(\mathcal{S} \times \mathcal{A})\}_{t=0}^{\infty}$ is a contractive mapping under the \mathcal{W}_1 distance, with $\Gamma(\mathcal{L}) := \Gamma_2(\Gamma_1(\mathcal{L}), \mathcal{L})$. Then by the Banach fixed point theorem and the completeness of the related metric spaces, there exists a unique NE to the GMFG.

In summary, we have

THEOREM 4 (Existence and Uniqueness of GMFG solution). *Given Assumptions 3 and 4, and assuming that $d_1 d_2 + d_3 < 1$, there exists a unique NE to (GMFG).*

The proof of Theorem 4 can be established by modifying appropriately the fixed-point approach for the stationary GMFG in Theorem 1.

References

- [1] Acciaio B, Backhoff J, Carmona R (2018) Extended mean field control problems: stochastic maximum principle and transport perspective. *Arxiv Preprint:1802.05754* .
- [2] Agarwal A, Kakade SM, Lee JD, Mahajan G (2021) On the theory of policy gradient methods: Optimality, approximation, and distribution shift. *Journal of Machine Learning Research* 22(98):1–76.
- [3] Anahtarçı B, Karıksız CD, Saldi N (2019) Fitted Q-learning in mean-field games. *arXiv preprint arXiv:1912.13309* .
- [4] Anahtarçı B, Karıksız CD, Saldi N (2020) Q-learning in regularized mean-field games. *arXiv preprint arXiv:2003.12151* .
- [5] Andelman N, Mansour Y (2004) Auctions with budget constraints. *Scandinavian Workshop on Algorithm Theory*, 26–38 (Springer).
- [6] Angiuli A, Fouque JP, Laurière M (2020) Unified reinforcement q-learning for mean field game and control problems. *arXiv preprint arXiv:2006.13912* .
- [7] Angiuli A, Fouque JP, Lauriere M (2021) Reinforcement learning for mean field games, with applications to economics. *arXiv preprint arXiv:2106.13755* .
- [8] Asadi K, Littman ML (2017) An alternative softmax operator for reinforcement learning. *Proceedings of the 34th International Conference on Machine Learning*, volume 70, 243–252.
- [9] Bellemare MG, Ostrovski G, Guez A, Thomas PS, Munos R (2016) Increasing the action gap: new operators for reinforcement learning. *AAAI Conference on Artificial Intelligence*, 1476–1483.
- [10] Benaim M, Le Boudec JY (2008) A class of mean field interaction models for computer and communication systems. *Performance evaluation* 65(11-12):823–838.
- [11] Bernstein MA, Griffin J (2006) Regional differences in the price-elasticity of demand for energy. Technical report, National Renewable Energy Lab.(NREL), Golden, CO (United States).
- [12] Bolley F (2008) Separability and completeness for the Wasserstein distance. *Séminaire de Probabilités XLI* 371–377.
- [13] Cai H, Ren K, Zhang W, Malialis K, Wang J, Yu Y, Guo D (2017) Real-time bidding by reinforcement learning in display advertising. *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*, 661–670 (ACM).
- [14] Caines PE, Huang M, Malhamé RP (2017) Mean field games. Basar T, Zaccour G, eds., *Handbook of Dynamic Game Theory* (Springer, Berlin).
- [15] Carmona R, Laurière M, Tan Z (2019) Linear-quadratic mean-field reinforcement learning: convergence of policy gradient methods. *arXiv preprint arXiv:1910.04295* .
- [16] Carmona R, Laurière M, Tan Z (2019) Model-free mean-field reinforcement learning: mean-field MDP and mean-field Q-learning. *arXiv preprint arXiv:1910.12802* .
- [17] Chen Y, Liu J, Khousainov B (2021) Maximum entropy inverse reinforcement learning for mean field games. *arXiv preprint arXiv:2104.14654* .
- [18] Cui K, Koepl H (2021) Approximately solving mean field games via entropy-regularized deep reinforcement learning. *International Conference on Artificial Intelligence and Statistics*, 1909–1917 (PMLR).
- [19] Cui K, Tahir A, Sinzger M, Koepl H (2021) Discrete-time mean field control with environment states. *arXiv preprint arXiv:2104.14900* .
- [20] de Luca MB, Vittori E, Trovò F, Restelli M (2021) Dealer markets: a reinforcement learning mean field approach .
- [21] Delarue F, Vasileiadis A (2021) Exploration noise for learning linear-quadratic mean field games. *arXiv preprint arXiv:2107.00839* .
- [22] Derman E, Mannor S (2020) Distributional robustness and regularization in reinforcement learning. *arXiv preprint arXiv:2003.02894* .
- [23] Elie R, Perolat J, Laurière M, Geist M, Pietquin O (2020) On the convergence of model free learning in mean field games. *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 7143–7150.

- [24] Even-Dar E, Mansour Y (2003) Learning rates for Q-learning. *Journal of Machine Learning Research* 5(Dec):1–25.
- [25] Fu Z, Yang Z, Chen Y, Wang Z (2019) Actor-critic provably finds nash equilibria of linear-quadratic mean-field games. *arXiv preprint arXiv:1910.07498* .
- [26] Fu Z, Yang Z, Chen Y, Wang Z (2019) Actor-critic provably finds nash equilibria of linear-quadratic mean-field games. *arXiv preprint arXiv:1910.07498* .
- [27] Gagrani M, Sudhakara S, Mahajan A, Nayyar A, Ouyang Y (2020) Thompson sampling for linear quadratic mean-field teams. *arXiv preprint arXiv:2011.04686* .
- [28] Gao B, Pavel L (2017) On the properties of the softmax function with application in game theory and reinforcement learning. *Arxiv Preprint:1704.00805* .
- [29] Geist M, Pérolat J, Laurière M, Elie R, Perrin S, Bachem O, Munos R, Pietquin O (2021) Concave utility reinforcement learning: the mean-field game viewpoint. *arXiv preprint arXiv:2106.03787* .
- [30] Geist M, Scherrer B, Pietquin O (2019) A theory of regularized Markov decision processes. *arXiv preprint arXiv:1901.11275* .
- [31] Ghadimi S, Lan G (2013) Stochastic first-and zeroth-order methods for nonconvex stochastic programming. *SIAM Journal on Optimization* 23(4):2341–2368.
- [32] Ghosh A, Aggarwal V (2020) Model free reinforcement learning algorithm for stationary mean field equilibrium for multiple types of agents. *arXiv preprint arXiv:2012.15377* .
- [33] Gibbs AL, Su FE (2002) On choosing and bounding probability metrics. *International Statistical Review* 70(3):419–435.
- [34] Gomes DA, Mohr J, Souza RR (2010) Discrete time, finite state space mean field games. *Journal de mathématiques pures et appliquées* 93(3):308–328.
- [35] Gu H, Guo X, Wei X, Xu R (2020) Mean-field controls with q-learning for cooperative marl: Convergence and complexity analysis. *arXiv preprint arXiv:2002.04131* .
- [36] Guéant O, Lasry JM, Lions PL (2011) Mean field games and applications. *Paris-Princeton lectures on mathematical finance 2010*, 205–266 (Springer).
- [37] Gummadi R, Key P, Proutiere A (2012) Repeated auctions under budget constraints: Optimal bidding strategies and equilibria. *the Eighth Ad Auction Workshop*.
- [38] Guo X, Hu A, Xu R, Zhang J (2019) Learning mean-field games. *Advances in Neural Information Processing Systems*, 4967–4977.
- [39] Guo X, Xu R, Zariphopoulou T (2020) Entropy regularization for mean field games with learning. *arXiv preprint arXiv:2010.00145* .
- [40] Haarnoja T, Tang H, Abbeel P, Levine S (2017) Reinforcement learning with deep energy-based policies. *Arxiv Preprint:1702.08165* .
- [41] Hamari J, Sjöklint M, Ukkonen A (2016) The sharing economy: Why people participate in collaborative consumption. *Journal of the Association for Information Science and Technology* 67(9):2047–2059.
- [42] Hernandez-Leal P, Kartal B, Taylor ME (2018) Is multiagent deep reinforcement learning the answer or the question? A brief survey. *Arxiv Preprint:1810.05587* .
- [43] Hu J, Wellman MP (2003) Nash Q-learning for general-sum stochastic games. *Journal of Machine Learning Research* 4(Nov):1039–1069.
- [44] Hu S, Leung CW, Leung Hf, Soh H (2020) The evolutionary dynamics of independent learning agents in population games. *arXiv preprint arXiv:2006.16068* .
- [45] Huang M, Ma Y (2017) Mean field stochastic games with binary action spaces and monotone costs. *ArXiv Preprint:1701.06661* .
- [46] Huang M, Malhamé RP, Caines PE (2006) Large population stochastic dynamic games: closed-loop McKean-Vlasov systems and the Nash certainty equivalence principle. *Communications in Information & Systems* 6(3):221–252.

- [47] Iyer K, Johari R, Sundararajan M (2011) Mean field equilibria of dynamic auctions with learning. *ACM SIGecom Exchanges* 10(3):10–14.
- [48] Jeong SH, Kang AR, Kim HK (2015) Analysis of game bot’s behavioral characteristics in social interaction networks of MMORPG. *ACM SIGCOMM Computer Communication Review* 45(4):99–100.
- [49] Jin J, Song C, Li H, Gai K, Wang J, Zhang W (2018) Real-time bidding with multi-agent reinforcement learning in display advertising. *Arxiv Preprint:1802.09756* .
- [50] Kapoor S (2018) Multi-agent reinforcement learning: A report on challenges and approaches. *Arxiv Preprint:1807.09427* .
- [51] Kizilkale AC, Caines PE (2013) Mean field stochastic adaptive control. *IEEE Transactions on Automatic Control* 58(4):905–920.
- [52] Lacker D (2015) Mean field games via controlled martingale problems: existence of Markovian equilibria. *Stochastic Processes and their Applications* 125(7):2856–2894.
- [53] Lacker D (2015) Mean field games via controlled martingale problems: existence of markovian equilibria. *Stochastic Processes and their Applications* 125(7):2856–2894.
- [54] Lasry JM, Lions PL (2007) Mean field games. *Japanese Journal of Mathematics* 2(1):229–260.
- [55] Lee K, Rengarajan D, Kalathil D, Shakkottai S (2021) Reinforcement learning for mean field games with strategic complementarities. *International Conference on Artificial Intelligence and Statistics*, 2458–2466 (PMLR).
- [56] Lehalle CA, Mouzouni C (2019) A mean field game of portfolio trading and its consequences on perceived correlations. *ArXiv Preprint:1902.09606* .
- [57] López JPM (2015) Discrete time mean field games: The short-stage limit. *Journal of Dynamics & Games* 2(1):89–101.
- [58] Luo Y, Yang Z, Wang Z, Kolar M (2019) Natural actor-critic converges globally for hierarchical linear quadratic regulator. *arXiv preprint arXiv:1912.06875* .
- [59] Mei J, Xiao C, Szepesvari C, Schuurmans D (2020) On the global convergence rates of softmax policy gradient methods. *International Conference on Machine Learning*, 6820–6829 (PMLR).
- [60] Mguni D, Jennings J, de Cote EM (2018) Decentralised learning in systems with many, many strategic agents. *Thirty-Second AAAI Conference on Artificial Intelligence*.
- [61] Minh VM, Badia AP, Mirza M, Graves A, Lillicrap TP, Harley T, Silver D, Kavukcuoglu K (2016) Asynchronous methods for deep reinforcement learning. *International Conference on Machine Learning*.
- [62] Mishra RK, Vasal D, Vishwanath S (2020) Model-free reinforcement learning for non-stationary mean field games. *2020 59th IEEE Conference on Decision and Control (CDC)*, 1032–1037 (IEEE).
- [63] Neu G, Jonsson A, Gómez V (2017) A unified view of entropy-regularized Markov decision processes. *arXiv preprint arXiv:1705.07798* .
- [64] Papadimitriou CH, Roughgarden T (2005) Computing equilibria in multi-player games. *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, 82–91.
- [65] Pasztor B, Bogunovic I, Krause A (2021) Efficient model-based multi-agent mean-field reinforcement learning. *arXiv preprint arXiv:2107.04050* .
- [66] Perolat J, Perrin S, Elie R, Laurière M, Piliouras G, Geist M, Tuyls K, Pietquin O (2021) Scaling up mean field games with online mirror descent. *arXiv preprint arXiv:2103.00623* .
- [67] Pérolat J, Piot B, Pietquin O (2018) Actor-critic fictitious play in simultaneous move multistage games. *International Conference on Artificial Intelligence and Statistics*.
- [68] Perrin S, Laurière M, Pérolat J, Geist M, Élie R, Pietquin O (2021) Mean field games flock! the reinforcement learning way. *arXiv preprint arXiv:2105.07933* .
- [69] Perrin S, Pérolat J, Laurière M, Geist M, Elie R, Pietquin O (2020) Fictitious play for mean field games: Continuous time analysis and applications. *arXiv preprint arXiv:2007.03458* .
- [70] Peyré G, Cuturi M (2019) Computational optimal transport. *Foundations and Trends in Machine Learning* 11(5-6):355–607.

- [71] Saldi N, Basar T, Raginsky M (2018) Markov–Nash equilibria in mean-field games with discounted cost. *SIAM Journal on Control and Optimization* 56(6):4256–4287.
- [72] Schulman J, Levine S, Abbeel P, Jordan M, Moritz P (2015) Trust region policy optimization. *International conference on machine learning*, 1889–1897.
- [73] Shani L, Efroni Y, Mannor S (2019) Adaptive trust region policy optimization: Global convergence and faster rates for regularized mdps. *arXiv preprint arXiv:1909.02769* .
- [74] Subramanian J, Mahajan A (2019) Reinforcement learning in stationary mean-field games. *18th International Conference on Autonomous Agents and Multiagent Systems*, 251–259.
- [75] Subramanian SG, Poupart P, Taylor ME, Hegde N (2020) Multi type mean field reinforcement learning. *arXiv preprint arXiv:2002.02513* .
- [76] Subramanian SG, Taylor ME, Crowley M, Poupart P (2020) Partially observable mean field reinforcement learning. *arXiv preprint arXiv:2012.15791* .
- [77] Tan M (1993) Multi-agent reinforcement learning: independent vs. cooperative agents. *International Conference on Machine Learning*, 330–337.
- [78] uz Zaman MA, Zhang K, Miehling E, Başar T (2020) Approximate equilibrium computation for discrete-time linear-quadratic mean-field games. *2020 American Control Conference (ACC)*, 333–339 (IEEE).
- [79] uz Zaman MA, Zhang K, Miehling E, Başar T (2020) Reinforcement learning in non-stationary discrete-time linear-quadratic mean-field games. *2020 59th IEEE Conference on Decision and Control (CDC)*, 2278–2284 (IEEE).
- [80] Villani C (2008) *Optimal transport: old and new*, volume 338 (Springer Science & Business Media).
- [81] Villani C (2009) *Optimal transport: old and new*, volume 338 (Springer).
- [82] Wang L, Cai Q, Yang Z, Wang Z (2019) Neural policy gradient methods: Global optimality and rates of convergence. *arXiv preprint arXiv:1909.01150* .
- [83] Wang L, Yang Z, Wang Z (2020) Breaking the curse of many agents: Provable mean embedding q-iteration for mean-field reinforcement learning. *International Conference on Machine Learning*, 10092–10103 (PMLR).
- [84] Wang W, Han J, Yang Z, Wang Z (2021) Global convergence of policy gradient for linear-quadratic mean-field control/game in continuous time. *International Conference on Machine Learning*, 10772–10782 (PMLR).
- [85] Xie Q, Yang Z, Wang Z, Minca A (2020) Provable fictitious play for general mean-field games. *arXiv preprint arXiv:2010.04211* .
- [86] Xie Q, Yang Z, Wang Z, Minca A (2021) Learning while playing in mean-field games: Convergence and optimality. *International Conference on Machine Learning*, 11436–11447 (PMLR).
- [87] Yang J, Ye X, Trivedi R, Xu H, Zha H (2017) Deep mean field games for learning optimal behavior policy of large populations. *Arxiv Preprint:1711.03156* .
- [88] Yang Y, Luo R, Li M, Zhou M, Zhang W, Wang J (2018) Mean field multi-agent reinforcement learning. *Arxiv Preprint:1802.05438* .
- [89] Yin H, Mehta PG, Meyn SP, Shanbhag UV (2013) Learning in mean-field games. *IEEE Transactions on Automatic Control* 59(3):629–644.

Appendix A: Distance metrics and completeness This section reviews some basic properties of the Wasserstein distance. It then proves that the metrics defined in the main text are indeed distance functions and define complete metric spaces.

ℓ_1 -Wasserstein distance and dual representation. The ℓ_1 Wasserstein distance over $\mathcal{P}(\mathcal{X})$ for $\mathcal{X} \subseteq \mathbb{R}^k$ is defined as

$$W_1(\nu, \nu') := \inf_{M \in \mathcal{M}(\nu, \nu')} \int_{\mathcal{X} \times \mathcal{X}} \|x - y\|_2 dM(x, y). \quad (52)$$

where $\mathcal{M}(\nu, \nu')$ is the set of all measures (couplings) on $\mathcal{X} \times \mathcal{X}$, with marginals ν and ν' on the two components, respectively.

The Kantorovich duality theorem enables the following equivalent dual representation of W_1 :

$$W_1(\nu, \nu') = \sup_{\|f\|_L \leq 1} \left| \int_{\mathcal{X}} f d\nu - \int_{\mathcal{X}} f d\nu' \right|, \quad (53)$$

where the supremum is taken over all 1-Lipschitz functions f , i.e., f satisfying $|f(x) - f(y)| \leq \|x - y\|_2$ for all $x, y \in \mathcal{X}$.

The Wasserstein distance W_1 can also be related to the total variation distance via the following inequalities Gibbs and Su [33]:

$$d_{\min}(\mathcal{X})d_{TV}(\nu, \nu') \leq W_1(\nu, \nu') \leq \text{diam}(\mathcal{X})d_{TV}(\nu, \nu'), \quad (54)$$

where $d_{\min}(\mathcal{X}) = \min_{x \neq y \in \mathcal{X}} \|x - y\|_2$, which is guaranteed to be positive when \mathcal{X} is finite.

When \mathcal{S} and \mathcal{A} are compact, for any compact subset $\mathcal{X} \subseteq \mathbb{R}^k$, and for any $\nu, \nu' \in \mathcal{P}(\mathcal{X})$, $W_1(\nu, \nu') \leq \text{diam}(\mathcal{X})d_{TV}(\nu, \nu') \leq \text{diam}(\mathcal{X}) < \infty$, where $\text{diam}(\mathcal{X}) = \sup_{x, y \in \mathcal{X}} \|x - y\|_2$ and d_{TV} is the total variation distance. Moreover, one can verify

LEMMA 8. *Both D and W_1 are distance functions, and they are finite for any input distribution pairs. In addition, both $(\{\Pi\}_{t=0}^{\infty}, D)$ and $(\{\mathcal{P}(\mathcal{S} \times \mathcal{A})\}_{t=0}^{\infty}, W_1)$ are complete metric spaces.*

These facts enable the usage of Banach fixed-point mapping theorem for the proof of existence and uniqueness (Theorems 4 and 1).

[Proof of Lemma 8] It is known that for any compact set $\mathcal{X} \subseteq \mathbb{R}^k$, $(\mathcal{P}(\mathcal{X}), W_1)$ defines a complete metric space Bolley [12]. Since $W_1(\nu, \nu') \leq \text{diam}(\mathcal{X})$ is uniformly bounded for any $\nu, \nu' \in \mathcal{P}(\mathcal{X})$, we know that $W_1(\mathcal{L}, \mathcal{L}') \leq \text{diam}(\mathcal{X})$ and $D(\boldsymbol{\pi}, \boldsymbol{\pi}') \leq \text{diam}(\mathcal{X})$ as well, so they are both finite for any input distribution pairs. It is clear that they are distance functions based on the fact that W_1 is a distance function.

Finally, we show the completeness of the two metric spaces $(\{\Pi\}_{t=0}^{\infty}, D)$ and $(\{\mathcal{P}(\mathcal{S} \times \mathcal{A})\}_{t=0}^{\infty}, W_1)$. Take $(\{\Pi\}_{t=0}^{\infty}, D)$ for example. Suppose that $\boldsymbol{\pi}^k$ is a Cauchy sequence in $(\{\Pi\}_{t=0}^{\infty}, D)$. Then for any $\epsilon > 0$, there exists a positive integer N , such that for any $m, n \geq N$,

$$D(\boldsymbol{\pi}^n, \boldsymbol{\pi}^m) \leq \epsilon \implies W_1(\pi_t^n(s), \pi_t^m(s)) \leq \epsilon \text{ for any } s \in \mathcal{S}, t \in \mathbb{N}, \quad (55)$$

which implies that $\pi_t^k(s)$ forms a Cauchy sequence in $(\mathcal{P}(\mathcal{A}), W_1)$, and hence by the completeness of $(\mathcal{P}(\mathcal{A}), W_1)$, $\pi_t^k(s)$ converges to some $\pi_t(s) \in \mathcal{P}(\mathcal{A})$. As a result, $\boldsymbol{\pi}^n \rightarrow \boldsymbol{\pi} \in \{\Pi\}_{t=0}^{\infty}$ under metric D , which shows that $(\{\Pi\}_{t=0}^{\infty}, D)$ is complete.

The completeness of $(\{\mathcal{P}(\mathcal{S} \times \mathcal{A})\}_{t=0}^{\infty}, W_1)$ can be proved similarly. Q.E.D.

The same argument for Lemma 8 shows that both D and W_1 are distance functions and are finite for any input distribution pairs, with both (Π, D) and $(\mathcal{P}(\mathcal{S} \times \mathcal{A}), W_1)$ again complete metric spaces.

Appendix B: Bounds for GMF-V-Q using asynchronous Q-learning In the main text, we have shown the results by using synchronous Q-learning algorithm. Here for the completeness, we also show the corresponding results for asynchronous Q-learning algorithm.

For asynchronous Q-learning algorithm, at each step l with the state s and an action a , the system reaches state s' according to the controlled dynamics and the Q-function approximation Q_l is updated according to

$$\hat{Q}^{l+1}(s, a) = (1 - \beta_l(s, a))\hat{Q}^l(s, a) + \beta_l(s, a) \left[r(s, a) + \gamma \max_{\bar{a}} \hat{Q}^l(s', \bar{a}) \right], \quad (56)$$

where $\hat{Q}^0(s, a) = C$ for some constant $C \in \mathbb{R}$ for any $s \in \mathcal{S}$ and $a \in \mathcal{A}$, and the step size $\beta_l(s, a)$ can be chosen as (Even-Dar and Mansour [24])

$$\beta_l(s, a) = \begin{cases} |\#(s, a, l) + 1|^{-h}, & (s, a) = (s_l, a_l), \\ 0, & \text{otherwise.} \end{cases} \quad (57)$$

with $h \in (1/2, 1)$. Here $\#(s, a, l)$ is the number of times up to time l that one visits the state-action pair (s, a) . The algorithm then proceeds to choose action a' based on \hat{Q}^{l+1} with appropriate exploration strategies, including the ϵ -greedy strategy.

LEMMA 9 (Even-Dar and Mansour [24]: sample complexity of asynchronous Q-learning).

For an MDP, say $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, r, \gamma)$, suppose that the Q-learning algorithm takes step-sizes (25). Also suppose that the covering time of the state-action pairs is bounded by L with probability at least $1 - p$ for some $p \in (0, 1)$. Then $\|\hat{Q}^{T_{\mathcal{M}}(\delta, \epsilon)} - Q_{\mathcal{M}}^*\|_{\infty} \leq \epsilon$ with probability at least $1 - 2\delta$. Here \hat{Q}^T is the T -th update in the Q-learning updates (24), $Q_{\mathcal{M}}^*$ is the (optimal) Q-function, and

$$T_{\mathcal{M}}(\delta, \epsilon) = \Omega \left(\left(\frac{L \log_p(\delta)}{\beta} \log \frac{V_{\max}}{\epsilon} \right)^{\frac{1}{1-h}} + \left(\frac{(L \log_p(\delta))^{1+3h} V_{\max}^2 \log \left(\frac{|\mathcal{S}||\mathcal{A}| V_{\max}}{\delta \beta \epsilon} \right)}{\beta^2 \epsilon^2} \right)^{\frac{1}{h}} \right),$$

where $\beta = (1 - \gamma)/2$, $V_{\max} = R_{\max}/(1 - \gamma)$, and R_{\max} is such that a.s. $0 \leq r(s, a) \leq R_{\max}$.

Here the covering time L of a state-action pair sequence is defined to be the number of steps needed to visit all state-action pairs starting from any arbitrary state-action pair. Also notice that the l_{∞} norm above is defined in an element-wise sense, i.e., for $M \in \mathbb{R}^{|\mathcal{S}| \times |\mathcal{A}|}$, we have $\|M\|_{\infty} = \max_{s \in \mathcal{S}, a \in \mathcal{A}} |M(s, a)|$.

COROLLARY 3 (Value-based guarantee of asynchronous Q-learning algorithm). The asynchronous Q-learning algorithm with appropriate choices of step-sizes (cf. (25)) satisfies the following value-based guarantee, where $C_{\mathcal{M}}^{(i)}$ ($i = 1, 2, 3$) are constants depending on $|\mathcal{S}|, |\mathcal{A}|, V_{\max}, \beta$ and h , and we have:

$$\begin{aligned} \alpha_2^{(1)} &= \alpha_4^{(1)} = \frac{1}{1-h}, \quad \alpha_1^{(1)} = \alpha_3^{(2)} = 0, \\ \alpha_1^{(2)} &= \frac{2}{h}, \quad \alpha_4^{(2)} = \frac{2+3h}{h}, \quad \alpha_j^{(2)} = 0 \text{ for } j = 2, 3, \\ \alpha_1^{(3)} &= \frac{2}{h}, \quad \alpha_2^{(3)} = \frac{1}{h}, \quad \alpha_4^{(3)} = \frac{1+3h}{h}, \quad \alpha_3^{(3)} = 0. \end{aligned}$$

In addition, assume the same assumptions as Theorem 1, then for Algorithm 3 with asynchronous Q-learning method, with probability at least $1 - 2\delta$, $W_1(\mathcal{L}_{K_{\epsilon, \eta}}, \mathcal{L}^*) \leq C_0 \epsilon$, where $K_{\epsilon, \eta}$ is defined as in Theorem 2. And the total number of samples $T = \sum_{k=0}^{K_{\epsilon, \eta} - 1} T_{\mathcal{M}_{\mathcal{L}_k}}(\delta_k, \epsilon_k)$ is bounded by

$$T \leq O \left(K_{\epsilon, \eta}^{\frac{4}{h}+1} \left(\log \frac{K_{\epsilon, \eta}}{\delta} \right)^{3+\frac{2}{h}} + \left(\log \frac{K_{\epsilon, \eta}}{\delta} \right)^{\frac{2}{1-h}} \right).$$

Appendix C: Weak simulator In this section, we state the counterpart of Theorems 2 and 3 for Algorithms 5 and 6, respectively. Notice that here the major difference is the additional $O(1/\sqrt{N})$ term.

We first (re)state the relation between \mathbf{Emp}_N (which serves as a $1/N$ -net) and action gaps:

For any positive integer N , there exist a positive constant $\phi_N > 0$, with the property that $\max_{a' \in \mathcal{A}} Q_{\mathcal{L}}^*(s, a') - Q_{\mathcal{L}}^*(s, a) \geq \phi_N$ for any $\mathcal{L} \in \mathbf{Emp}_N$, $s \in \mathcal{S}$, and any $a \notin \arg\max_{a \in \mathcal{A}} Q_{\mathcal{L}}^*(s, a)$.

Now we are ready to state the convergence results.

THEOREM 5 (Convergence and complexity of GMF-VW). Assume the same assumptions as Theorem 1. Suppose that Alg has a value-based guarantee with parameters

$$\{C_{\mathcal{M}}^{(i)}, \alpha_1^{(i)}, \alpha_2^{(i)}, \alpha_3^{(i)}, \alpha_4^{(i)}\}_{i=1}^m.$$

For any $\epsilon, \delta > 0$, set $\delta_k = \delta/K_{\epsilon, \eta}$, $\epsilon_k = (k+1)^{-(1+\eta)}$ for some $\eta \in (0, 1]$ ($k = 0, \dots, K_{\epsilon, \eta} - 1$), and $c' = c = \frac{\log(1/\epsilon)}{\phi_N}$.⁶ Then with probability at least $1 - 4\delta$,

$$W_1(\mathcal{L}_{K_{\epsilon, \eta}}, \mathcal{L}^*) \leq C\epsilon + \frac{\text{diam}(\mathcal{S})\text{diam}(\mathcal{A})|\mathcal{S}||\mathcal{A}|}{2(1-d)} \sqrt{\frac{1}{2N} \log(|\mathcal{S}||\mathcal{A}|K_{\epsilon, \eta}/\delta)}.$$

Here $K_{\epsilon, \eta} := \lceil 2 \max\{(\eta\epsilon/c)^{-1/\eta}, \log_d(\epsilon/\max\{\text{diam}(\mathcal{S})\text{diam}(\mathcal{A}), c\}) + 1\} \rceil$ is the number of outer iterations, and the constant C is independent of δ, ϵ and η .

Moreover, the total number of samples $T = \sum_{k=0}^{K_{\epsilon, \eta}-1} T_{\mathcal{M}_{\mathcal{L}_k}}(\delta_k, \epsilon_k)$ is bounded by

$$T \leq \sum_{i=1}^m \frac{2^{\alpha_2^{(i)}}}{2\alpha_1^{(i)} + 1} C_{\mathcal{M}}^{(i)} K_{\epsilon, \eta}^{2\alpha_1^{(i)} + 1} (K_{\epsilon, \eta}/\delta)^{\alpha_3^{(i)}} (\log(K_{\epsilon, \eta}/\delta))^{\alpha_2^{(i)} + \alpha_4^{(i)}}. \quad (58)$$

THEOREM 6 (Convergence and complexity of GMF-PW). Assume the same assumptions as in Theorem 1. Suppose that Alg has a policy-based guarantee with parameters

$$\{C_{\mathcal{M}}^{(i)}, \alpha_1^{(i)}, \alpha_2^{(i)}, \alpha_3^{(i)}, \alpha_4^{(i)}\}_{i=1}^m.$$

Then for any $\epsilon, \delta > 0$, set $\delta_k = \delta/K_{\epsilon, \eta}$, $\epsilon_k = (k+1)^{-(1+\eta)}$ for some $\eta \in (0, 1]$ ($k = 0, \dots, K_{\epsilon, \eta} - 1$), and $c' = c = \frac{\log(1/\epsilon)}{\phi_N}$,⁷ with probability at least $1 - 4\delta$,

$$W_1(\mathcal{L}_{K_{\epsilon, \eta}}, \mathcal{L}^*) \leq C\epsilon + \frac{\text{diam}(\mathcal{S})\text{diam}(\mathcal{A})|\mathcal{S}||\mathcal{A}|}{2(1-d)} \sqrt{\frac{1}{2N} \log(|\mathcal{S}||\mathcal{A}|K_{\epsilon, \eta}/\delta)}.$$

Here $K_{\epsilon, \eta} := \lceil 2 \max\{(\eta\epsilon/c)^{-1/\eta}, \log_d(\epsilon/\max\{\text{diam}(\mathcal{S})\text{diam}(\mathcal{A}), c\}) + 1\} \rceil$ is the number of outer iterations, and the constant C is independent of δ, ϵ and η .

Moreover, the total number of samples $T = \sum_{k=0}^{K_{\epsilon, \eta}-1} T_{\mathcal{M}_{\mathcal{L}_k}}(\delta_k, \epsilon_k)$ is bounded by

$$T \leq \sum_{i=1}^{m+1} \frac{2^{\alpha_2^{(i)}}}{2\alpha_1^{(i)} + 1} \tilde{C}_{\mathcal{M}}^{(i)} K_{\epsilon, \eta}^{2\alpha_1^{(i)} + 1} (K_{\epsilon, \eta}/\delta)^{\alpha_3^{(i)}} (\log(K_{\epsilon, \eta}/\delta))^{\alpha_2^{(i)} + \alpha_4^{(i)}}, \quad (59)$$

where the parameters $\{\tilde{C}_{\mathcal{M}}^{(i)}, \alpha_1^{(i)}, \alpha_2^{(i)}, \alpha_3^{(i)}, \alpha_4^{(i)}\}_{i=1}^{m+1}$ are defined in Lemma 3.

⁶ Here we actually only need $c' = \Omega(\frac{\log(1/\epsilon)}{\phi_N})$ and $c = O(\frac{\log(1/\epsilon)}{\phi_N})$, and the corresponding result will differ only in some absolute constants.

⁷ Here again we actually only need $c' = \Omega(\frac{\log(1/\epsilon)}{\phi_N})$ and $c = O(\frac{\log(1/\epsilon)}{\phi_N})$, and the corresponding result will differ only in some absolute constants.

The key to the proof of Theorems 5 and 6 is the following lemma, which follows from the Hoeffding inequality.

LEMMA 10. *The expectation $\mathbb{E}[\mathcal{L}_{k+1}(s', a')] = \pi_k(s', a') \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \mu_k(s) P(s' | s, a, \mathcal{L}_k) \pi_k(s, a) = \Gamma_2(\pi_k, \mathcal{L}_k)$. In addition, we have that for any $t > 0$, $s' \in \mathcal{S}$ and $a' \in \mathcal{A}$,*

$$\mathbb{P}(|\mathcal{L}_{k+1}(s', a') - \mathbb{E}[\mathcal{L}_{k+1}(s', a')]| \geq t) \leq 2 \exp(-2Nt^2). \quad (60)$$

The above lemma essentially states that the iterates \mathcal{L}_{k+1} of Algorithms 5 and 6 are very close to the “ $\tilde{\mathcal{L}}_{k+1}$ ” obtained from the (strong) simulator $\mathcal{G}(s, \pi_k, \mathcal{L}_k)$ with $s \sim \mu_k$ following line 5 in Algorithm 2 and line 6 in Algorithm 4. This bridges the gap between the weak and the strong simulators. In particular, by noticing that

$$\begin{aligned} W_1(\mathcal{L}_{k+1}, \mathcal{L}^*) &\leq W_1(\mathcal{L}_{k+1}, \Gamma_2(\pi_k, \mathcal{L}_k)) + W_1(\Gamma_2(\pi_k, \mathcal{L}_k), \Gamma_2(\Gamma_1(\mathcal{L}^*), \mathcal{L}^*)) \\ &= W_1(\mathcal{L}_{k+1}, \mathbb{E}[\mathcal{L}_{k+1}]) + W_1(\Gamma_2(\pi_k, \mathcal{L}_k), \Gamma_2(\Gamma_1(\mathcal{L}^*), \mathcal{L}^*)) \\ &\leq \frac{\text{diam}(\mathcal{S})\text{diam}(\mathcal{A})|\mathcal{S}||\mathcal{A}|}{2} \|\mathcal{L}_{k+1} - \mathbb{E}[\mathcal{L}_{k+1}]\|_\infty + W_1(\Gamma_2(\pi_k, \mathcal{L}_k), \Gamma_2(\Gamma_1(\mathcal{L}^*), \mathcal{L}^*)), \end{aligned}$$

one can bound the first term with high probability via (60). The second term is then bounded in exactly the same way as the proofs for Theorems 2 and 3, and hence we omit the details.