

This version of the article has been accepted for publication, after peer review (when applicable) and is subject to Springer Nature's AM terms of use (<https://www.springernature.com/gp/open-research/policies/accepted-manuscript-terms>), but is not the Version of Record and does not reflect post-acceptance improvements, or any corrections. The Version of Record is available online at: <http://dx.doi.org/10.1007/s00245-024-10205-7>.

Continuous time q-learning for mean-field control problems

Xiaoli Wei ^{*} Xiang Yu [†]

Abstract

This paper studies the q-learning, recently coined as the continuous time counterpart of Q-learning by [Jia and Zhou \(2023\)](#), for continuous time mean-field control problems in the setting of entropy-regularized reinforcement learning. In contrast to the single agent's control problem in [Jia and Zhou \(2023\)](#), we reveal that two different q-functions naturally arise in mean-field control problems: (i) the integrated q-function (denoted by q) as the first-order approximation of the integrated Q-function introduced in [Gu et al. \(2023\)](#), which can be learnt by a weak martingale condition using all test policies; and (ii) the essential q-function (denoted by q_e) that is employed in the policy improvement iterations. We show that two q-functions are related via an integral representation. Based on the weak martingale condition and our proposed searching method of test policies, some model-free learning algorithms are devised. In two examples, one in LQ control framework and one beyond LQ control framework, we can obtain the exact parameterization of the optimal value function and q-functions and illustrate our algorithms with simulation experiments.

Keywords: continuous time reinforcement learning, integrated q-function, mean-field control, weak martingale characterization, test policies

1 Introduction

Mean-field control (MFC) problems, also known as McKean-Vlasov control problems, concern stochastic systems of large population where the agents interact through the distribution of their states and the optimal decision is made by a social planner to attain the Pareto optimality. It has been seen rapid progress in both theories and applications in the study of MFC problems in the past two decades. The comprehensive survey of this topic and related studies can be found in [Carmona and Delarue \(2018a\)](#) and [Carmona and Delarue \(2018b\)](#).

By nature, the model parameters in the stochastic system with a large population of agents are generally difficult to observe or estimate, and it is a desirable research direction to design efficient learning algorithms for solving the MFC problems under the unknown environment. Reinforcement learning (RL) provides many well-suited learning algorithms for this purpose that enables the agent to learn the optimal control through the trial-and-error procedure. Through the exploration and exploitation in RL, the agent takes an action and observes a reward outcome that signals the influence of the agent's action such that the agent can learn to select actions based on past experiences and by making new choices. Recently, many exploration-exploitation

^{*}Email: xiaoli.wei@hit.edu.cn, Institute for Advance Study in Mathematics, Harbin Institute of Technology, China.

[†]Email: xiang.yu@polyu.edu.hk, Department of Applied Mathematics, The Hong Kong Polytechnic University, Kowloon, Hong Kong.

RL algorithms for single agent’s control problems were quickly adopted and generalized in many multi-agent and mean-field reinforcement learning applications.

Despite its substantial success in wide applications, the theoretical study of reinforcement learning has been predominantly limited to discrete time models, especially in the MFC and mean-field game problems. [Gu et al. \(2023\)](#) discussed the correct form of Q-function in an integral form in Q-learning and established the dynamic programming principle (DPP) for learning MFC problems. [Gu et al. \(2021a,b\)](#) proposed a model-free Q learning in a centralized way and an actor-critic algorithm in a decentralized manner for MFC problems with convergence analyses. [Carmona et al \(2023\)](#) studied MFC with common noise under both closed-loop and open-loop policies through the lens of mean-field Markov decision process, which also unveiled the correct form of the Q-function and its dynamic programming principle and adapted existing (deep) RL methods to the mean-field setting. [Angiuli et al. \(2022, 2023\)](#) studied unified two-timescale RL algorithms to solve infinite horizon asymptotic MFC and mean-field game problems in finite and continuous state-action space, respectively. The corresponding convergence result is studied in the subsequent work [Angigli et al. \(2023\)](#). There also exists a variety of model-based or model-free RL algorithms for MFCs that can be found, for instance in [Mondal et al. \(2022, 2023\)](#); [Pasztor et al. \(2021\)](#). On the other hand, for continuous time stochastic control problems by a single agent, [Wang et al. \(2020\)](#), [Jia and Zhou \(2022a,b, 2023\)](#) have laid the overarching theoretical foundation for reinforcement learning in continuous time with continuous state space and possibly continuous action space. In particular, [Jia and Zhou \(2023\)](#) developed a continuous time q-learning theory by considering the first order approximation of the conventional Q-function. Given a stochastic policy, the value function and the associated q-function can be characterized by martingale conditions of some stochastic processes in [Jia and Zhou \(2023\)](#) in both on-policy and off-policy settings. Several actor-critic algorithms are devised therein for solving underlying RL problem. Convergence or regret analyses of continuous-time single-agent stochastic control problems have been conducted in [Giegrich et al. \(2024\)](#); [Szpruch et al. \(2021\)](#) for linear-quadratic or linear-convex models. For continuous time mean-field LQ games, [Guo et al. \(2022\)](#) examined the theoretical justification that entropy regularization helps stabilizing and accelerating the convergence to the Nash equilibrium. [Li et al. \(2023\)](#) concerned the model-based policy iteration reinforcement learning method for continuous time infinite horizon mean-field LQ control problems. Recently, [Frikha et al. \(2023\)](#) generalized the policy gradient algorithm in [Jia and Zhou \(2022b\)](#) to continuous time MFC problems and devised actor-critic algorithms based on a gradient expectation representation of the value function, where the value function and the policy are learnt alternatively via the observed samples of the state and a model-free estimation of the population distribution.

As demonstrated in [Jia and Zhou \(2023\)](#), one important advantage of continuous time framework lies in its robustness with respect to the time discretization in the algorithm design and implementations. Inspired by [Jia and Zhou \(2023\)](#), we are particularly interested in whether, and if yes, how the continuous time q-learning can be applied in learning McKean-Vlasov control problems in the mean-field model with infinitely many interacting agents. Contrary to [Jia and Zhou \(2023\)](#), our first contribution is to reveal that two distinct q-functions are generically needed to learn continuous time McKean-Vlasov control problems. In fact, due to the feature of mean-field interactions with the population, it has been shown in [Gu et al. \(2023\)](#); [Carmona et al \(2023\)](#) in the discrete time setting that the time consistency and DPP only hold for the *so-called* integrated Q-function (or IQ) for learning MFC problems, where the IQ-function depends on

the distribution of both the state and the control, i.e., the state space of the state variable and the action variable needs to be appropriately enlarged to the Wasserstein space of probability measures so that the time-consistency can be guaranteed. Similar to the counterpart in Gu et al. (2023); Carmona et al (2023), we also show that the correct definition of the continuous time q-function, as the first order approximation of the IQ-function, needs to be defined on the distribution of the state and the control as an integral form of the Hamiltonian operator. This integral form of the q-function is called the integrated q-function in the present paper, which is crucial in establishing a weak martingale characterization of the value function and the q-function using all test policies in a neighbourhood of the target policy. However, on the other hand, from the relaxed control formulation with the entropy regularizer and its associated exploratory HJB equation, the optimal policy can be obtained as a Gibbs measure related to the Hamiltonian operator directly. Therefore, the integrated q-function actually can not be utilized directly to learn the optimal policy. Instead, a proper way is to introduce another q-function that is defined as the Hamiltonian (without the integral form) plus the temporal dispersion, which can be employed in the policy improvement iterations. To distinguish two different q-functions in our setting, we shall call the first order approximation of the IQ-function as the continuous time integrated q-function, denoted by q (see Definition 3.1); and we shall name the second q-function related to the policy improvement as the *essential* q-function, denoted by q_e (see Definition 3.3).

One key finding of the present paper is the integral representation between q and q_e involving all test policies, see the relationship in (3.13). As a result, we can employ the weak martingale characterization of the integrated q-function (see Theorems 4.3 and 4.4) to devise the learning algorithm in the following order: (i) We first parameterize the value function $J^\theta(t, \mu)$ and the essential q-function q_e^ψ , and then obtain the parameterized q-function q^ψ from the integral relationship that shares the same parameter ψ of q_e^ψ . (ii) By minimizing the weak martingale loss under all test policies, we devise updating rules for the parameters in the integrated q-function q^ψ . Here, a new challenge we encounter is to decide the searching rule of test policies in the weak martingale characterization. To address this issue, we consider the minimization of the weak martingale loss robust with respect to all test policies, i.e., the updating rule is determined by a minmax problem. We propose a method based on the average of some test policies close to the target policy, which does not create new parameters. In addition, it is worth noting that the test policy instead of the target policy will be used to generate samples and observations due to the dependence of the distribution of the action in the integrated q-function, therefore making our learning algorithms similar to the off-policy learning in the conventional Q-learning, see Remark 4.5. (iii) After the updating of the parameter ψ in q from step (ii), the parameter of q_e is also updated that gives the updating rule of the associated policy.

On the other hand, it is shown in Gu et al. (2023) that the IQ-function in the discrete time framework is generally nonlinear in the distribution of control, and hence it is difficult to characterize the distribution of the optimal policy. By contrast, it is interesting to see that our integrated q-function, as the first order derivative of the IQ-function with respect to time, admits a linear integral form of the distribution in (3.13), which is more suitable for the name of “integrated function” of the distribution of the population and allows the explicit connection between the optimal policy and the essential q-function, see Remark 3.5.

To illustrate our model-free q-learning algorithms, we study two concrete examples in financial applications, namely the optimal mean-variance portfolio problem in the LQ control

framework and the mean-field optimal R&D investment and consumption problem beyond the LQ control framework. In both examples, we can derive the explicit expressions of the optimal value function and two q-functions from the exploratory HJB equation and hence obtain their exact parameterized forms. From simulation experiments, we illustrate the satisfactory performance of our q-learning algorithm and the proposed searching method of test policies.

The remainder of this paper is organized as follows. Section 2 introduces the exploratory learning formulation for continuous time McKean-Vlasov control problems. In Section 3, the definitions of the continuous time integrated q-function and the essential q-function are given and their relationship is established. The weak martingale characterization of the value function and two q-functions using test policies are established in Section 4. Some offline and online q-learning algorithms using the average-based test policies are presented in Section 5. Section 6 examines two concrete financial applications and numerically illustrates our q-learning algorithms in some simulation experiments. Section 7 concludes the contributions of the paper and discusses some directions for future research.

2 Problem Formulation

2.1 Strong Control Formulation

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be the probability space that supports a n -dimensional Brownian motion $W = (W_s)_{s \in [0, T]}$. We denote by $\mathbb{F}^W = (\mathcal{F}_s^W)_{s \in [0, T]}$ the \mathbb{P} -completion filtration of W . We assume that there is a sub-algebra \mathcal{G} of \mathcal{F} such that \mathcal{G} is independent of \mathbb{F}^W and is “rich enough” in the sense that for any $\mu \in \mathcal{P}_2(\mathbb{R}^d)$, there exists a \mathcal{G} -measurable random variable ξ such that $\mathbb{P}_\xi = \mu$. Indeed, such \mathcal{G} exists if and only if there exists a \mathcal{G} -measurable random variable U with the uniform distribution $\mathcal{U}([0, 1])$ that is independent of $W = (W_s)_{s \geq 0}$. We denote by $\mathbb{F} = (\mathcal{F}_s)_{s \geq 0}$ the filtration defined by $\mathcal{F}_s = \mathcal{F}_s^W \vee \mathcal{G}$, see section 2 in [Cosso et al. \(2023\)](#) for more details.

The state dynamics of the controlled McKean-Vlasov SDE is given by

$$dX_s = b(s, X_s, \mathbb{P}_{X_s}, a_s)ds + \sigma(s, X_s, \mathbb{P}_{X_s}, a_s)dW_s, \quad X_t = \xi \sim \mu, \quad (2.1)$$

where \mathbb{P}_{X_s} denotes the probability distribution of X_s . In the present paper, we only consider the model without common noise, leaving the case with common noise as future study.

The goal of the McKean-Vlasov control problem is to find the optimal \mathbb{F} -progressively measurable sequence of actions $\{a_s\}_{t \leq s \leq T}$ valued in the space $\mathcal{A} \subset \mathbb{R}^m$ that maximizes the expected discounted total reward

$$V^*(t, \mu) := \sup_{\{a_s\}_{t \leq s \leq T}} \mathbb{E} \left[\int_t^T e^{-\beta(s-t)} r(s, X_s, \mathbb{P}_{X_s}, a_s) ds + e^{-\beta(T-t)} g(X_T, \mathbb{P}_{X_T}) \middle| \mathbb{P}_{X_t} = \mu \right]. \quad (2.2)$$

One can show that the optimal value function is law-invariant and $V^* : [0, T] \times \mathcal{P}_2(\mathbb{R}^d) \rightarrow \mathbb{R}$ satisfies the dynamic programming principle

$$V^*(t, \mu) = \sup_{\{a_s\}_{t \leq s \leq T}} \mathbb{E} \left[\int_t^{t+h} e^{-\beta(s-t)} r(s, X_s, \mathbb{P}_{X_s}, a_s) ds + e^{-\beta h} V^*(t+h, \mathbb{P}_{X_{t+h}}) \right].$$

Furthermore, when b , σ , r and g are known, the optimal value function satisfies the following dynamic programming equation

$$\frac{\partial V^*}{\partial t}(t, \mu) - \beta V^*(t, \mu) + \int_{\mathbb{R}^d} \sup_{a \in \mathcal{A}} H(t, x, \mu, a, \partial_\mu V^*(t, \mu)(x), \partial_x \partial_\mu V^*(t, \mu)(x)) \mu(dx) = 0,$$

with the terminal condition $V^*(T, \mu) = \int_{\mathbb{R}^d} g(x, \mu) \mu(dx) := \hat{g}(\mu)$, where $\partial_\mu V^*(t, \mu)(x)$ is defined via the lifting identification and it is called L-derivative of V^* with respect to the measure μ , see [Lions \(2006\)](#) and Definition 5.22 in [Carmona and Delarue \(2018a\)](#), and $\partial_x \partial_\mu V^*(t, \mu)(x)$ denotes the partial derivative of $\partial_\mu V^*(t, \mu)(x)$ with respect to x , and the Hamiltonian operator H is defined by

$$H(t, x, \mu, a, p, q) = b(t, x, \mu, a)^\top p + \frac{1}{2} \text{Tr}(\sigma \sigma^\top(t, x, \mu, a) q) + r(t, x, \mu, a). \quad (2.3)$$

2.2 Exploratory Formulation

We now consider the situation when the model is unknown, i.e., we do not have the exact information of the model parameters b and σ in state dynamics (2.1) and the reward function r in (2.2). To determine the optimal control in face of the unknown model, we choose to apply the reinforcement learning based on the trial-and-error procedure. That is, what the representative agent (social planner) can do is to try a sequence of actions $(a_s)_{s \in [0, T]}$ and observe the corresponding state process $(X_s^a)_{s \in [0, T]}$ and the distribution $(\mathbb{P}_{X_s^a})_{s \in [0, T]}$ along with a stream of discounted running rewards $r(s, X_s^a, \mathbb{P}_{X_s^a}, a_s)_{s \in [0, T]}$ and continuously update and improve her actions based on these observations.

To describe the exploration step in reinforcement learning, we can randomize the actions and consider its distribution as a relaxed control. To this end, we need to assume that $(\Omega, \mathcal{F}, \{\mathcal{F}_s\}_s, \mathbb{P})$ is rich enough to support a continuum of independent random variables $(Z_s)_{s \in [0, T]}$ that are uniformly distributed on $[0, 1]$ and also independent of W ; See Theorem 1 in [Sun \(2006\)](#) for the construction of $(Z_s)_{s \in [0, T]}$. Note that the additional randomness $(Z_s)_{s \in [0, T]}$ allows the agent to freely randomize his action at each time s .

Let π be a stochastic control that maps $(t, x, \mu) \in [0, T] \times \mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d) \rightarrow \mathcal{P}(\mathcal{A})$, where $\mathcal{P}(\mathcal{A})$ is the space of probability measures on \mathcal{A} . At each time $s \in [0, T]$, an action $a_s := a(s, X_s, \mathbb{P}_{X_s}, Z_s)$ is sampled from $\pi(\cdot | s, X_s, \mathbb{P}_{X_s})$. Fix a stochastic policy π and an initial pair $(t, \mu) \in [0, T] \times \mathcal{P}_2(\mathbb{R}^d)$, let us consider the controlled McKean-Vlasov SDE

$$dX_s^\pi = b(s, X_s^\pi, \mathbb{P}_{X_s^\pi}, a_s^\pi) ds + \sigma(s, X_s^\pi, \mathbb{P}_{X_s^\pi}, a_s^\pi) dW_s, \quad X_t^\pi \sim \mu, \quad a_s^\pi \sim \pi(\cdot | s, X_s^\pi, \mathbb{P}_{X_s^\pi}). \quad (2.4)$$

When coefficients b and σ are Lipschitz in $(t, x, \mu) \in [0, T] \times \mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d)$ for every action $a \in \mathcal{A}$, the existence and uniqueness of solution of (2.4) can be proved by following Theorem 5.1.1 in [Stroock and Caradhan \(1997\)](#). The solution to (2.4) is denoted as $X^{t, \mu, \pi} = \{X_s^{t, \mu, \pi}, t \leq s \leq T\}$.

To encourage the exploration in continuous time, we adopt the Shannon entropy regularizer suggested by [Wang et al. \(2020\)](#), and the value function of the exploration formulation is given by

$$J(t, \mu; \pi) = \mathbb{E} \left[\int_t^T e^{-\beta(s-t)} [r(s, X_s^\pi, \mathbb{P}_{X_s^\pi}, a_s^\pi) - \gamma \log \pi(a_s^\pi | s, X_s^\pi, \mathbb{P}_{X_s^\pi})] ds \right]$$

$$+ e^{-\beta(T-t)} g(X_T^\pi, \mathbb{P}_{X_T^\pi}) \Big| X_t^\pi \sim \mu \Big]. \quad (2.5)$$

Due to the Shannon entropy regularization in (2.5), it is clear that the value function under the policy without a density becomes $-\infty$. Therefore, we shall restrict our attention to policies that admit densities. A stochastic policy π is called admissible if π admits a density function and is jointly measurable with respect to $(t, x, \mu, a) \in [0, T] \times \mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d) \times \mathcal{A}$ and $\text{supp}(\pi) = \mathcal{A}$. We denote by Π the set of admissible (stochastic) policies π .

We also consider the dynamics (2.6), viewed intuitively as the average of the sample trajectories X^π in (2.4) over randomized actions.

$$dX_s = b_\pi(s, X_s, \mathbb{P}_{X_s}, \pi(\cdot|s, X_s, \mathbb{P}_{X_s}))ds + \sigma_\pi(s, X_s, \mathbb{P}_{X_s}, \pi(\cdot|s, X_s, \mathbb{P}_{X_s}))dW_s, \quad (2.6)$$

where

$$b_\pi(s, x, \mu) := \int_{\mathcal{A}} b(s, x, \mu, a) \pi(a|s, x, \mu) da, \quad \sigma_\pi(s, x, \mu) := \sqrt{\int_{\mathcal{A}} \sigma \sigma^\top(s, x, \mu, a) \pi(a|s, x, \mu) da}.$$

The wellposedness of (2.6) is guaranteed under Assumption 2.1 (i) and its solution is denoted by $\tilde{X}^\pi = \{\tilde{X}_s^\pi, t \leq s \leq T\}$. In appendix A, we also rigorously prove the connection between (2.4) and (2.6) in the sense that X_s^π and \tilde{X}_s^π have the same law for every $s \in [t, T]$ by showing that they correspond to the same martingale problem.

As $\mathbb{P}_{X_s^\pi} = \mathbb{P}_{\tilde{X}_s^\pi}$, the objective function in (2.5) is equivalent to

$$J(t, \mu; \pi) = \mathbb{E} \left[\int_t^T e^{-\beta(s-t)} [r_\pi(s, \tilde{X}_s^\pi, \mathbb{P}_{\tilde{X}_s^\pi}) + \gamma \mathcal{E}_\pi(s, \tilde{X}_s^\pi, \mathbb{P}_{\tilde{X}_s^\pi})] ds \right. \\ \left. + e^{-\beta(T-t)} g(\tilde{X}_T^\pi, \mathbb{P}_{\tilde{X}_T^\pi}) \Big| \tilde{X}_t^\pi \sim \mu \right], \quad (2.7)$$

where we denote

$$r_\pi(t, x, \mu) := \int_{\mathcal{A}} r(t, x, \mu, a) \pi(a|t, x, \mu) da, \quad \mathcal{E}_\pi(t, x, \mu) := - \int_{\mathcal{A}} \log \pi(a|t, x, \mu) \pi(a|t, x, \mu) da.$$

Then the value function associated with the admissible policy π satisfies the dynamic programming equation:

$$\frac{\partial J}{\partial t}(t, \mu; \pi) - \beta J(t, \mu; \pi) + \mathbb{E}_{\xi \sim \mu} \left[\int_{\mathcal{A}} \left[H(t, \xi, \mu, a, \partial_\mu J(t, \mu; \pi)(\xi), \partial_x \partial_\mu J(t, \mu; \pi)(\xi)) \right. \right. \\ \left. \left. - \gamma \log \pi(a|t, \xi, \mu) \right] \pi(a|t, \xi, \mu) da \right] = 0, \quad (2.8)$$

with the terminal condition $J(T, \mu; \pi) = \hat{g}(\mu) := \int_{\mathbb{R}^d} g(x, \mu) \mu(dx)$.

In addition, let us denote

$$M_2(\mu) := \int_{\mathbb{R}^d} |x|^2 \mu(dx).$$

We shall make the following regularity assumptions on coefficients and reward functions throughout the paper (see Assumption 2.1 and Remark 2.1 in Frikha et al. (2023)).

Assumption 2.1 (i) For $f \in \{b_\pi, \sigma_\pi\}$, the derivatives $\partial_x f(t, x, \mu)$, $\partial_x^2 f(t, x, \mu)$, $\partial_\mu f(t, x, \mu)(v)$ and $\partial_v \partial_\mu f(t, x, \mu)(v)$ exist for any $(t, x, v, \mu) \in [0, T] \times \mathbb{R}^d \times \mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d)$, are bounded and locally Lipschitz continuous with respect to x, μ, v uniformly in $t \in [0, T]$. For all $(t, x, \mu) \in [0, T] \times \mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d)$, we have $|f(t, x, \mu)| \leq C(1 + |x| + M_2(\mu))$.

(ii) For any $t \in [0, T]$, $r_\pi(t, \cdot)$, $\mathcal{E}_\pi(t, \cdot)$ and $g(\cdot) \in \mathcal{C}^{2,2}(\mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d))$.

(iii) There exists some constant $C < \infty$, such that for any $(t, x, v, \mu) \in [0, T] \times \mathbb{R}^d \times \mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d)$, we have

$$\begin{aligned} |r_\pi(t, x, \mu)| + |\mathcal{E}_\pi(t, x, \mu)| + |g(x, \mu)| &\leq C(1 + |x|^2 + M_2(\mu)), \\ |\partial_x r_\pi(t, x, \mu)| + |\partial_x \mathcal{E}_\pi(t, x, \mu)| + |\partial_x g(x, \mu)| &\leq C(1 + |x| + M_2(\mu)), \\ |\partial_\mu r_\pi(t, x, \mu)(v)| + |\partial_\mu \mathcal{E}_\pi(t, x, \mu)(v)| + |\partial_\mu g(x, \mu)(v)| &\leq C(1 + |x| + |v| + M_2(\mu)), \\ |\partial_v \partial_\mu r_\pi(t, x, \mu)(v)| + |\partial_x^2 r_\pi(t, x, \mu)| + |\partial_v \partial_\mu \mathcal{E}_\pi(t, x, \mu)(v)| \\ + |\partial_x^2 \mathcal{E}_\pi(t, x, \mu)| + |\partial_v \partial_\mu g(x, \mu)(v)| + |\partial_x^2 g(x, \mu)| &\leq C(1 + M_2(\mu)). \end{aligned}$$

Under Assumption 2.1, Proposition 2.1 of Frikha et al. (2023) guarantees that the function $J(t, \mu; \pi)$ defined in (2.5) is of $C^{1,2}([0, T] \times \mathcal{P}_2(\mathbb{R}^d))$.

The objective of the social planner is to maximize $J(t, \mu; \pi)$ over all admissible policies

$$J^*(t, \mu) = \sup_{\pi \in \Pi} J(t, \mu; \pi). \quad (2.9)$$

We can derive the exploratory HJB equation for the optimal value function $J^*(t, \mu)$ by

$$\begin{aligned} \frac{\partial J^*}{\partial t}(t, \mu) + \int_{\mathbb{R}^d} \sup_{\pi \in \mathcal{P}(\mathcal{A})} \left[\int_{\mathcal{A}} \left(H(t, x, \mu, a, \partial_\mu J^*(t, \mu)(x), \partial_x \partial_\mu J^*(t, \mu)(x)) \right. \right. \\ \left. \left. - \gamma \log \pi(a) \right) \pi(a) da \right] \mu(dx) - \beta J^*(t, \mu) = 0. \end{aligned} \quad (2.10)$$

The optimal policy π satisfies a Gibbs measure or widely-used Boltzmann policy in RL after normalization that

$$\pi^*(a|t, x, \mu) = \frac{\exp \left\{ \frac{1}{\gamma} H(t, x, \mu, a, \partial_\mu J^*(t, \mu)(x), \partial_x \partial_\mu J^*(t, \mu)(x)) \right\}}{\int_{\mathcal{A}} \exp \left\{ \frac{1}{\gamma} H(t, x, \mu, a, \partial_\mu J^*(t, \mu)(x), \partial_x \partial_\mu J^*(t, \mu)(x)) \right\} da}. \quad (2.11)$$

Plugging (2.11) into (2.10), we get that

$$\begin{aligned} \frac{\partial J^*}{\partial t}(t, \mu) + \gamma \int_{\mathbb{R}^d} \log \int_{\mathcal{A}} \exp \left\{ \frac{1}{\gamma} H(t, x, \mu, a, \partial_\mu J^*(t, \mu)(x), \partial_x \partial_\mu J^*(t, \mu)(x)) \right\} da \mu(dx) \\ - \beta J^*(t, \mu) = 0, \end{aligned} \quad (2.12)$$

with the terminal condition $J^*(T, \mu) = \hat{g}(\mu)$.

Recall that the goal of MFC is to find the optimal policy that maximizes $J(t, \mu; \pi)$. Let us consider the operator $\mathcal{I} : \Pi \rightarrow \Pi$ that

$$\mathcal{I}(\pi) = \frac{\exp \left\{ \frac{1}{\gamma} H(t, x, \mu, a, \partial_\mu J(t, \mu; \pi)(x), \partial_x \partial_\mu J(t, \mu; \pi)(x)) \right\}}{\int_{\mathcal{A}} \exp \left\{ \frac{1}{\gamma} H(t, x, \mu, a, \partial_\mu J(t, \mu; \pi)(x), \partial_x \partial_\mu J(t, \mu; \pi)(x)) \right\} da}. \quad (2.13)$$

Theorem 2.2 (Policy improvement result) For any given $\pi \in \Pi$, define $\pi' = \mathcal{I}(\pi)$, with \mathcal{I} given in (2.13). Then

$$J(t, \mu; \pi') \geq J(t, \mu; \pi).$$

Moreover, if the map \mathcal{I} in (2.13) has a fixed point $\pi^* \in \Pi$, then π^* is the optimal policy of (2.9).

Proof. For two given admissible policies $\pi, \pi' \in \Pi$, and any $0 \leq t \leq T$, by applying Itô's formula to the value function $J(s, \mathbb{P}_{X_s^{\pi'}}; \pi)$ between t and T , we get that

$$\begin{aligned} & e^{-\beta T} J(T, \mathbb{P}_{X_T^{\pi'}}; \pi) - e^{-\beta t} J(t, \mathbb{P}_{X_t^{\pi'}}; \pi) \\ & + \mathbb{E} \left[\int_t^T \int_{\mathcal{A}} e^{-\beta s} \left[r(s, X_s^{\pi'}, \mathbb{P}_{X_s^{\pi'}}, a) - \gamma \log \pi'(a|s, X_s^{\pi'}, \mathbb{P}_{X_s^{\pi'}}) \right] \pi'(a|s, X_s^{\pi'}, \mathbb{P}_{X_s^{\pi'}}) dad s \right] \\ = & \int_t^T e^{-\beta s} \left[\frac{\partial J}{\partial t}(s, \mathbb{P}_{X_s^{\pi'}}; \pi) - \beta J(s, \mathbb{P}_{X_s^{\pi'}}; \pi) \right] ds \\ & + \mathbb{E} \left[\int_t^T \int_{\mathcal{A}} \left[H(s, X_s^{\pi'}, \mathbb{P}_{X_s^{\pi'}}, a, \partial_\mu J(s, \mathbb{P}_{X_s^{\pi'}}; \pi)(X_s^{\pi'}), \partial_x \partial_\mu J(s, \mathbb{P}_{X_s^{\pi'}}; \pi)(X_s^{\pi'})) \right. \right. \\ & \left. \left. - \gamma \log \pi'(a|s, \mathbb{P}_{X_s^{\pi'}}, X_s^{\pi'}) \right] \pi'(a|s, \mathbb{P}_{X_s^{\pi'}}, X_s^{\pi'}) dad s \right]. \end{aligned} \quad (2.14)$$

From Lemma 9 in Jia and Zhou (2023), it follows that for any $(s, y, \mu) \in [0, T] \times \mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d)$,

$$\begin{aligned} & \int_{\mathcal{A}} \left[H(s, y, \mu, a, \partial_\mu J(s, \mu; \pi)(y), \partial_x \partial_\mu J(s, \mu; \pi)(y)) - \gamma \log \pi'(a|s, y, \mu) \right] \pi'(a|s, y, \mu) da \\ & \geq \int_{\mathcal{A}} \left[H(s, y, \mu, a, \partial_\mu J(s, \mu; \pi)(y), \partial_x \partial_\mu J(s, \mu; \pi)(y)) - \gamma \log \pi(a|s, y, \mu) \right] \pi(a|s, y, \mu) da. \end{aligned}$$

Therefore, it holds that

$$\begin{aligned} & e^{-\beta T} J(T, \mathbb{P}_{X_T^{\pi'}}; \pi) - e^{-\beta t} J(t, \mathbb{P}_{X_t^{\pi'}}; \pi) \\ & + \mathbb{E} \left[\int_t^T \int_{\mathcal{A}} e^{-\beta s} \left[r(s, X_s^{\pi'}, \mathbb{P}_{X_s^{\pi'}}, a) - \gamma \log \pi'(a|s, X_s^{\pi'}, \mathbb{P}_{X_s^{\pi'}}) \right] \pi'(a|s, X_s^{\pi'}, \mathbb{P}_{X_s^{\pi'}}) dad s \right] \geq 0. \end{aligned} \quad (2.15)$$

Letting $\mathbb{P}_{X_t^{\pi'}} = \mu$ and rearranging terms in (2.15) yield the desired result that

$$\begin{aligned} & J(t, \mu; \pi) \leq e^{-\beta(T-t)} J(T, \mathbb{P}_{X_T^{\pi'}}; \pi) \\ & + \mathbb{E} \left[\int_t^T \int_{\mathcal{A}} e^{-\beta s} \left[r(s, X_s^{\pi'}, \mathbb{P}_{X_s^{\pi'}}, a) - \gamma \log \pi'(a|s, X_s^{\pi'}, \mathbb{P}_{X_s^{\pi'}}) \right] \pi'(a|s, X_s^{\pi'}, \mathbb{P}_{X_s^{\pi'}}) dad s \right] \\ = & e^{-\beta(T-t)} \hat{g}(\mathbb{P}_{X_T^{\pi'}}) \\ & + \mathbb{E} \left[\int_t^T \int_{\mathcal{A}} e^{-\beta s} \left[r(s, X_s^{\pi'}, \mathbb{P}_{X_s^{\pi'}}, a) - \gamma \log \pi'(a|s, X_s^{\pi'}, \mathbb{P}_{X_s^{\pi'}}) \right] \pi'(a|s, X_s^{\pi'}, \mathbb{P}_{X_s^{\pi'}}) dad s \right] \\ = & J(t, \mu; \pi'). \end{aligned}$$

The learning procedure in Theorem 2.2 starts with some policy π and produces a new policy π' by setting $\pi' = \mathcal{I}(\pi)$. We next estimate the distance between π and π' .

Lemma 2.3 Under Assumption 2.1, define $\pi' = \mathcal{I}(\pi)$ with \mathcal{I} given in (2.13) for any given $\pi \in \Pi$, then we have

$$0 \leq D_{KL}^{\text{average}}(\pi \|\pi') := \int_{\mathbb{R}^d} D_{KL}(\pi(\cdot|t, x, \mu) \|\pi'(\cdot|t, x, \mu))\mu(dx) \leq \delta(\mu) := C(1 + M_2(\mu)), \quad (2.16)$$

where C is a generic constant. Furthermore, if π is a fixed point of \mathcal{I} , then $D_{KL}^{\text{average}}(\pi \|\pi') = 0$.

Proof. By the definition of D_{KL} , we have that

$$\begin{aligned} D_{KL}(\pi(\cdot|t, x, \mu) \|\pi'(\cdot|t, x, \mu)) &= \int_{\mathcal{A}} \log \frac{\pi(a|t, x, \mu)}{\pi'(a|t, x, \mu)} \pi(a|t, x, \mu) da \\ &= -\frac{1}{\gamma} \int_{\mathcal{A}} H(t, x, \mu, \partial_{\mu} J(t, \mu; \pi)(x), \partial_x \partial_{\mu} J(t, \mu; \pi)(x)) \pi(a|t, x, \mu) da \\ &\quad + \log \int_{\mathcal{A}} \exp \left\{ \frac{1}{\gamma} H(t, x, \mu, a, \partial_{\mu} J(t, \mu; \pi)(x), \partial_x \partial_{\mu} J(t, \mu; \pi)(x)) \right\} da \\ &\quad + \int_{\mathcal{A}} \log \pi(a|t, x, \mu) \pi(a|t, x, \mu) da. \end{aligned}$$

It then follows from (2.8) that

$$\begin{aligned} D_{KL}^{\text{average}}(\pi \|\pi') &= \int_{\mathbb{R}^d} D_{KL}(\pi(\cdot|t, x, \mu) \|\pi'(\cdot|t, x, \mu))\mu(dx) \\ &= \frac{\partial J}{\partial t}(t, \mu; \pi) - \beta J(t, \mu; \pi) \\ &\quad + \int_{\mathbb{R}^d} \log \int_{\mathcal{A}} \exp \left\{ \frac{1}{\gamma} H(t, x, \mu, a, \partial_{\mu} J(t, \mu; \pi)(x), \partial_x \partial_{\mu} J(t, \mu; \pi)(x)) \right\} da \mu(dx). \end{aligned}$$

We then estimate $D_{KL}^{\text{average}}(\pi \|\pi')$. Under Assumption 2.1, the proof of Proposition 2.1 in Frikha et al. (2023) implies that

$$\sup_{t \in [0, T]} \left\{ \left| \frac{\partial J}{\partial t}(t, \mu; \pi) \right| + |J(t, \mu; \pi)| + |\partial_{\mu} J(t, \mu; \pi)(x)| + |\partial_x \partial_{\mu} J(t, \mu; \pi)(x)| \right\} \leq C(1 + M_2(\mu)).$$

Consequently, the desired result holds that $D_{KL}^{\text{average}}(\pi \|\pi') \leq C(1 + M_2(\mu))$.

3 q-Functions for Continuous Time Mean-field Control

The aim of this section is to examine the continuous time analogue of the discrete time IQ-function and Q-learning for mean-field control problems (see Gu et al. (2023); Carmona et al (2023)). In particular, in the framework of continuous time McKean-Vlasov control, it is an interesting open problem that what is the correct definition of the integrated q-function comparing with the q-function for a single agent's control problem in Jia and Zhou (2023) and the IQ-function in the discrete time framework in Gu et al. (2023); Carmona et al (2023)? In addition, it is important to explore how can one utilize the learnt integrated q-function to learn the optimal policy?

3.1 Soft Q-learning for Mean-field Control

To better elaborate our definitions of continuous time q-functions for mean-field control problems, let us first detour in this subsection to discuss the correct definition and results of Q-learning with entropy regularizer (soft Q-learning) in the discrete time framework.

We consider a mean-field Markov Decision Process (MDP) $X = \{X_t, t = 0, 1, \dots, T\}$ with a finite state space \mathcal{X} and a finite action space \mathcal{A} , and a transition probability of mean-field type $\mathbb{P}(X_{t+1} = x' | X_t = x, \mathbb{P}_{X_t} = \mu, a_t = a) =: p(x' | x, \mu, a)$. At each time step t , an action is sampled from a stochastic policy π . The social planner's expected total reward is $\mathbb{E}[\sum_{t=0}^{T-1} \beta^t r(X_t, \mathbb{P}_{X_t}, a_t) + g(X_T, \mathbb{P}_{X_T})]$, with $\beta := \exp(-\alpha) \in (0, 1)$.

The value function at each time t associated with a given policy π is defined as

$$J(t, \mu; \{\pi_s\}_{s \geq t}) = \mathbb{E} \left[\sum_{s=t}^{T-1} \beta^{(s-t)} [r(X_s^\pi, \mathbb{P}_{X_s^\pi}, a_s^\pi) - \gamma \log \pi(a_s^\pi | s, X_s^\pi, \mathbb{P}_{X_s^\pi})] + \beta^{(T-t)} g(X_T^\pi, \mathbb{P}_{X_T^\pi}) \right],$$

with $\mathbb{P}_{X_t^\pi} = \mu$. The Bellman equation for J is given by

$$J(t, \mu; \{\pi_s\}_{s \geq t}) = \mathbb{E} [r(X_t^\pi, \mu, a_t^\pi) - \gamma \log \pi(a_t^\pi | s, X_t^\pi, \mu)] + \beta J(t+1, \mathbb{P}_{X_{t+1}^\pi}; \{\pi_s\}_{s \geq t+1}). \quad (3.1)$$

When the social planner takes the policy \mathbf{h} at time t , and then the policy $\{\pi_s\}_{s \geq t+1}$ afterwards, the integrated Q-function (see Gu et al. (2023) without the entropy regularizer) is defined by

$$\begin{aligned} Q(t, \mu, \mathbf{h}; \{\pi_s\}_{s \geq t+1}) &= \mathbb{E} [r(X_t^{\mathbf{h}}, \mathbb{P}_{X_t^{\mathbf{h}}}, a_t^{\mathbf{h}}) - \gamma \log \mathbf{h}(a_t^{\mathbf{h}} | t, X_t^{\mathbf{h}}, \mathbb{P}_{X_t^{\mathbf{h}}})] \\ &\quad + \beta \mathbb{E} \left[\sum_{s=t+1}^{T-1} \beta^{s-(t+1)} [r(X_s^\pi, \mathbb{P}_{X_s^\pi}, a_s^\pi) - \gamma \log \pi(a_s^\pi | s, X_s^\pi, \mathbb{P}_{X_s^\pi})] \middle| X_{t+1}^\pi \sim \mathbb{P}_{X_{t+1}^{\mathbf{h}}} \right] \\ &= \mathbb{E} [r(X_t^{\mathbf{h}}, \mathbb{P}_{X_t^{\mathbf{h}}}, a_t^{\mathbf{h}}) - \gamma \log \mathbf{h}(a_t^{\mathbf{h}} | t, X_t^{\mathbf{h}}, \mathbb{P}_{X_t^{\mathbf{h}}})] + \beta J(t+1, \mathbb{P}_{X_{t+1}^{\mathbf{h}}}; \{\pi_s\}_{s \geq t+1}), \end{aligned} \quad (3.2)$$

with the terminal condition $Q(T, \mu, \mathbf{h}) = \mathbb{E}_{\xi \sim \mu} [g(\xi, \mu)]$.

Next, we consider the optimal value function $J^*(t, \mu) := \sup_{\{\pi_s\}_{s \geq t+1}} J(t, \mu; \{\pi_s\}_{s \geq t+1})$ and the optimal Q-function and $Q^*(t, \mu, \mathbf{h}) = \sup_{\{\pi_s\}_{s \geq t+1}} Q(t, \mu; \mathbf{h}; \{\pi_s\}_{s \geq t+1})$ associated with the optimal policy π^* . First, the DPP or the Bellman equation for J^* is

$$J^*(t, \mu) = \sup_{\pi} \left\{ \mathbb{E}_{\mu, \pi} [r(\xi, \mu, a_t^\pi) - \gamma \log \pi(a_t^\pi | t, \xi, \mu)] + \beta J^*(t+1, \mathbb{P}_{X_{t+1}^{t, \mu, \pi}}) \right\}, \quad (3.3)$$

where for any measurable function f , $\mu \in \mathcal{P}_2(\mathbb{R}^d)$ and $\mathbf{h} \in \Pi$, we denote

$$\mathbb{E}_{\mu, \mathbf{h}} [f(\xi, a^{\mathbf{h}})] := \mathbb{E}_{\xi \sim \mu, a^{\mathbf{h}} \sim \mathbf{h}(\cdot | t, \xi, \mu)} [f(\xi, a^{\mathbf{h}})] = \int_{\mathbb{R}^d} \int_{\mathcal{A}} f(x, a) \mathbf{h}(a | t, x, \mu) da \mu(dx). \quad (3.4)$$

From (3.2) and (3.3), we have the relation between the optimal value function and the optimal Q-function

$$J^*(t, \mu) = \sup_{\mathbf{h}} Q^*(t, \mu, \mathbf{h}). \quad (3.5)$$

Substituting (3.5) to the Bellman equation for J^* , we obtain that

$$\begin{aligned}
J^*(t, \mu) &= \sup_{\boldsymbol{\pi}} \left\{ \mathbb{E}_{\mu, \boldsymbol{\pi}} [r(\xi, \mu, a_t^{\boldsymbol{\pi}}) - \gamma \log \boldsymbol{\pi}(a_t^{\boldsymbol{\pi}} | t, \xi, \mu)] + \beta \sup_{\mathbf{h}'} Q^*(t+1, \mathbb{P}_{X_{t+1}^{t, \mu, \boldsymbol{\pi}}}; \mathbf{h}') \right\} \\
&= \sup_{\boldsymbol{\pi}} \left\{ \sum_{x \in \mathcal{X}} \sum_{a \in \mathcal{A}} (r(x, \mu, a) - \gamma \log \boldsymbol{\pi}(a | t, x, \mu)) \boldsymbol{\pi}(a | t, x, \mu) \mu(x) \right. \\
&\quad \left. + \beta \sup_{\mathbf{h}'} Q^*(t+1, \Phi(t, \mu, \boldsymbol{\pi}), \mathbf{h}') \right\}, \tag{3.6}
\end{aligned}$$

where Φ , which characterizes the evolution of $\mathbb{P}_{X_t^{\boldsymbol{\pi}}}$ over time, is defined by $\Phi(t, \mu, \mathbf{h})(x') = \sum_{x \in \mathcal{X}} \sum_{a \in \mathcal{A}} p(x' | x, \mu, a) \mathbf{h}(a | t, x, \mu) \mu(x)$.

To derive the optimal policy, we need to find the candidate policy that achieves $\sup_{\mathbf{h}} Q^*(t, \mu, \mathbf{h})$ for any $\mu \in \mathcal{P}_2(\mathcal{X})$ and $t \in 0, 1, \dots, T-1$, or equivalently we can search for the optimal policy by solving the optimization problem on the right hand side of (3.6). However, due to the non-linear dependence of $\sup_{\mathbf{h}'} Q^*(t+1, \Phi(t, \mathbb{P}_{X_t^{\boldsymbol{\pi}}}, \boldsymbol{\pi}), \mathbf{h}')$ in $\boldsymbol{\pi}$, it is impossible to obtain an explicit expression of $\boldsymbol{\pi}^*$. Therefore, for McKean-Vlasov control problems in the mean-field model, we reveal that the introduction of the entropy regularizer does not provide any help to derive the distribution of the optimal policy $\boldsymbol{\pi}^*$, which differs significantly from the soft Q-learning for the single agent's stochastic control problem, see [Jia and Zhou \(2023\)](#).

Moreover, taking the supremum over all $\{\boldsymbol{\pi}_s\}_{s \geq t+1}$ in (3.2), and substituting (3.5) to (3.2), we can heuristically obtain the DPP or the Bellman equation for the optimal IQ-function that

$$Q^*(t, \mu, \mathbf{h}) = \mathbb{E}_{\mu, \mathbf{h}} [r(\xi, \mu, a_t^{\mathbf{h}}) - \gamma \log \mathbf{h}(a_t^{\mathbf{h}} | t, \xi, \mu)] + \beta \sup_{\mathbf{h}'} Q^*(t+1, \mathbb{P}_{X_{t+1}^{t, \mu, \mathbf{h}}}, \mathbf{h}'), \tag{3.7}$$

with the terminal condition $Q^*(T, \mu, \mathbf{h}) = \mathbb{E}_{\xi \sim \mu} [g(\xi, \mu)]$.

Note that when $\gamma = 0$, (3.7) coincides with the Bellman equation for the optimal IQ-function established in [Gu et al. \(2023\)](#), in which the correct definition of the IQ-function $Q(t, \mu, \mathbf{h})$ defined on the distribution of both the population and the action is proposed for MFC such that the time consistency and DPP hold. In (3.7) for the learning MFC problem with the entropy regularizer, it is straightforward to modify arguments in [Gu et al. \(2023\)](#) and similarly prove the correct DPP for the optimal IQ-function with an additional entropy term associated with \mathbf{h} .

On the other hand, when there is no mean-field interaction, i.e., there is no population distribution in the transition probability p or in the reward r , the terminal payoff g and the stochastic policy $\boldsymbol{\pi}$, the problem degenerates to the single-agent Q learning. We can use either the single-agent Q-function denoted by Q_{single}^* or the IQ-function to learn the optimal policy. In fact, these two Q-functions are related by the equation

$$Q^*(t, \mu, \mathbf{h}) = \sum_{x \in \mathcal{X}} \sum_{a \in \mathcal{A}} (Q_{\text{single}}^*(t, x, a) - \gamma \log \mathbf{h}(a | t, x)) \mathbf{h}(a | t, x) \mu(x), \tag{3.8}$$

which implies that the optimal policy that attains $\sup_{\mathbf{h}} Q^*(t, \mu, \mathbf{h})$ can be explicitly written by

$$\boldsymbol{\pi}^*(a | t, x) = \frac{\exp\{\frac{1}{\gamma} Q_{\text{single}}^*(t, x, a)\}}{\sum_{a \in \mathcal{A}} \exp\{\frac{1}{\gamma} Q_{\text{single}}^*(t, x, a)\}}, \tag{3.9}$$

which has been shown in [Jia and Zhou \(2023\)](#).

According to (3.8), we see that (3.7) can be reformulated in terms of Q_{single}^* that

$$\begin{aligned} Q_{\text{single}}^*(t, x, a) &= r(t, x, a) + \beta \sum_{x' \in \mathcal{X}} \gamma \log \sum_{a \in \mathcal{A}} \exp \left\{ \frac{1}{\gamma} Q_{\text{single}}^*(t+1, x', a) \right\} p(x'|t, x, a) \\ &= r(t, x, a) + \beta \mathbb{E} \left[\gamma \log \sum_{a \in \mathcal{A}} \exp \left\{ \frac{1}{\gamma} Q_{\text{single}}^*(t+1, X_{t+1}^{t,x,a}, a) \right\} \middle| X_t = x, a_t = a \right], \end{aligned} \quad (3.10)$$

and that (3.5) can be expressed in terms of Q_{single}^* and J_{single}^* that

$$J_{\text{single}}^*(t, x) = \sup_{\boldsymbol{\pi}} \left\{ \sum_{a \in \mathcal{A}} (Q_{\text{single}}^*(t, x, a) - \gamma \log \boldsymbol{\pi}(a|t, x)) \boldsymbol{\pi}(a|t, x) \right\}. \quad (3.11)$$

Equations (3.9)-(3.11) are consistent with results for the single-agent soft Q-learning, see e.g. [Sutton and Barto \(2018\)](#) and [Jia and Zhou \(2023\)](#).

3.2 Two Continuous Time q-functions

Due to the large population of interacting agents as well as the technical issue discussed in Remark 4.2 below, we need to consider the randomized action $\mathbf{h} \in \Pi$ on the interval $[t, t + \Delta t)$ in the IQ-function in the McKean-Vlasov control framework (see [Gu et al. \(2023\)](#)), which differs significantly from the perturbed policy with a constant action in [Jia and Zhou \(2023\)](#) for a single agent. Therefore, let us consider a ‘‘perturbed policy’’ $\bar{\boldsymbol{\pi}} \in \Pi$, which takes $\mathbf{h} \in \Pi$ on $[t, t + \Delta t)$, and then $\boldsymbol{\pi} \in \Pi$ on $[t + \Delta t, T)$. By Lemma 2.3, it is sufficient to restrict \mathbf{h} within the ball $B_{\delta(\mu)}(\boldsymbol{\pi}) = \{\mathbf{h} \in \Pi : D_{KL}^{average}(\boldsymbol{\pi} || \mathbf{h}) \leq \delta(\mu)\}$ centered at $\boldsymbol{\pi}$ that significantly reduces the searching space of policies \mathbf{h} .

Denote $\rho_s^{t,\mu,\bar{\boldsymbol{\pi}}} := \mathbb{P}_{X_s^{t,\mu,\bar{\boldsymbol{\pi}}}}$ for any $\boldsymbol{\pi} \in \Pi$ for notational simplicity. The state process $X_s^{t,\mu,\bar{\boldsymbol{\pi}}}$ on $[t, T)$ is governed by

$$\begin{aligned} dX_s^{t,\mu,\bar{\boldsymbol{\pi}}} &= b(s, X_s^{t,\mu,\bar{\boldsymbol{\pi}}}, \rho_s^{t,\mu,\bar{\boldsymbol{\pi}}}, a_s^{\mathbf{h}}) ds + \sigma(s, X_s^{t,\mu,\bar{\boldsymbol{\pi}}}, \rho_s^{t,\mu,\bar{\boldsymbol{\pi}}}, a_s^{\mathbf{h}}) dW_s, \quad s \in [t, t + \Delta t), \quad X_t^{t,\mu,\bar{\boldsymbol{\pi}}} \sim \mu, \\ dX_s^{t,\mu,\bar{\boldsymbol{\pi}}} &= b(s, X_s^{t,\mu,\bar{\boldsymbol{\pi}}}, \rho_s^{t,\mu,\bar{\boldsymbol{\pi}}}, a_s^{\boldsymbol{\pi}}) ds + \sigma(s, X_s^{t,\mu,\bar{\boldsymbol{\pi}}}, \rho_s^{t,\mu,\bar{\boldsymbol{\pi}}}, a_s^{\boldsymbol{\pi}}) dW_s, \quad s \in [t + \Delta t, T), \quad X_{t+\Delta t}^{t,\mu,\bar{\boldsymbol{\pi}}} \sim \rho_{t+\Delta t}^{t,\mu,\mathbf{h}}. \end{aligned}$$

Based on the definition in (3.2), we can first consider the discrete time integrated Q -function defined on $[0, T] \times \mathcal{P}_2(\mathbb{R}^d) \times B_{\delta(\mu)}(\boldsymbol{\pi}) \times \Pi$ with the time interval Δt and the entropy regularizer that

$$\begin{aligned} &Q_{\Delta t}(t, \mu, \mathbf{h}; \boldsymbol{\pi}) \quad (3.12) \\ &= \mathbb{E} \left[\int_t^{t+\Delta t} \int_{\mathcal{A}} e^{-\beta(s-t)} \left[r(s, X_s^{t,\mu,\bar{\boldsymbol{\pi}}}, \rho_s^{t,\mu,\bar{\boldsymbol{\pi}}}, a) - \gamma \log \mathbf{h}(a|s, X_s^{t,\mu,\bar{\boldsymbol{\pi}}}, \rho_s^{t,\mu,\bar{\boldsymbol{\pi}}}) \right] \mathbf{h}(a|s, X_s^{t,\mu,\bar{\boldsymbol{\pi}}}, \rho_s^{t,\mu,\bar{\boldsymbol{\pi}}}) da ds \right. \\ &\quad + \int_{t+\Delta t}^T \int_{\mathcal{A}} e^{-\beta(s-t)} \left[r(s, X_s^{t,\mu,\bar{\boldsymbol{\pi}}}, \rho_s^{t,\mu,\bar{\boldsymbol{\pi}}}, a) - \gamma \log \boldsymbol{\pi}(a|s, X_s^{t,\mu,\bar{\boldsymbol{\pi}}}, \rho_s^{t,\mu,\bar{\boldsymbol{\pi}}}) \right] \boldsymbol{\pi}(a|s, X_s^{t,\mu,\bar{\boldsymbol{\pi}}}, \rho_s^{t,\mu,\bar{\boldsymbol{\pi}}}) da ds \\ &\quad \left. + e^{-\beta(T-t)} g(X_T^{t,\mu,\bar{\boldsymbol{\pi}}}, \rho_T^{t,\mu,\bar{\boldsymbol{\pi}}}) \middle| X_t^{t,\mu,\bar{\boldsymbol{\pi}}} \sim \mu \right] \\ &= \mathbb{E} \left[\int_t^{t+\Delta t} \int_{\mathcal{A}} e^{-\beta(s-t)} \left[r(s, X_s^{t,\mu,\bar{\boldsymbol{\pi}}}, \rho_s^{t,\mu,\bar{\boldsymbol{\pi}}}, a) - \gamma \log \mathbf{h}(a|s, X_s^{t,\mu,\bar{\boldsymbol{\pi}}}, \rho_s^{t,\mu,\bar{\boldsymbol{\pi}}}) \right] \mathbf{h}(a|s, X_s^{t,\mu,\bar{\boldsymbol{\pi}}}, \rho_s^{t,\mu,\bar{\boldsymbol{\pi}}}) da \right] \end{aligned}$$

$$\begin{aligned}
& + e^{-\beta\Delta t} J(t + \Delta t, \rho_{t+\Delta t}^{t, \mu, \bar{\pi}}; \boldsymbol{\pi}) - J(t, \mu; \boldsymbol{\pi}) + J(t, \mu; \boldsymbol{\pi}) \\
= & \mathbb{E} \left[\int_t^{t+\Delta t} \int_{\mathcal{A}} e^{-\beta(s-t)} \left[H(s, X_s^{t, \mu, \mathbf{h}}, \rho_s^{t, \mu, \mathbf{h}}, a, \partial_\mu J(s, \rho_s^{t, \mu, \mathbf{h}}; \boldsymbol{\pi})(X_s^{t, \mu, \mathbf{h}}), \partial_x \partial_\mu J(s, \rho_s^{t, \mu, \mathbf{h}}; \boldsymbol{\pi})(X_s^{t, \mu, \mathbf{h}})) \right. \right. \\
& \left. \left. - \gamma \log \mathbf{h}(a|s, X_s^{t, \mu, \bar{\pi}}, \rho_s^{t, \mu, \bar{\pi}}) \right] \mathbf{h}(a|s, X_s^{t, \mu, \bar{\pi}}, \rho_s^{t, \mu, \bar{\pi}}) da ds \right] \\
& + \int_t^{t+\Delta t} \left[\frac{\partial J}{\partial t}(s, \rho_s^{t, \mu, \mathbf{h}}; \boldsymbol{\pi}) - \beta J(s, \rho_s^{t, \mu, \mathbf{h}}; \boldsymbol{\pi}) \right] ds + J(t, \mu; \boldsymbol{\pi}) \\
= & J(t, \mu; \boldsymbol{\pi}) + \Delta t \left[\frac{\partial J}{\partial t}(t, \mu; \boldsymbol{\pi}) - \beta J(t, \mu; \boldsymbol{\pi}) + \mathbb{E}_{\xi \sim \mu} \left[\int_{\mathcal{A}} \left(H(t, \xi, \mu, a, \partial_\mu J(t, \mu; \boldsymbol{\pi})(\xi), \partial_x \partial_\mu J(t, \mu; \boldsymbol{\pi})(\xi)) \right. \right. \right. \\
& \left. \left. \left. - \gamma \log \mathbf{h}(a|t, \xi, \mu) \right) \mathbf{h}(a|t, \xi, \mu) da \right] \right] + o(\Delta t),
\end{aligned}$$

where we have applied Itô's formula to $e^{-\beta\Delta t} J(t + \Delta t, \rho_{t+\Delta t}^{t, \mu, \bar{\pi}}; \boldsymbol{\pi}) - J(t, \mu; \boldsymbol{\pi})$ in the third equality.

Definition 3.1 Fix a stochastic policy $\boldsymbol{\pi} \in \Pi$. For any $(t, \mu, \mathbf{h}) \in [0, T] \times \mathcal{P}_2(\mathbb{R}^d) \times B_{\delta(\mu)}(\boldsymbol{\pi})$, we define the continuous time integrated q -function by

$$\begin{aligned}
q(t, \mu, \mathbf{h}; \boldsymbol{\pi}) := & \frac{\partial J}{\partial t}(t, \mu; \boldsymbol{\pi}) - \beta J(t, \mu; \boldsymbol{\pi}) + \mathbb{E}_{\mu, \mathbf{h}} \left[H(t, \xi, \mu, a^{\mathbf{h}}, \partial_\mu J(t, \mu; \boldsymbol{\pi})(\xi), \partial_x \partial_\mu J(t, \mu; \boldsymbol{\pi})(\xi)) \right] \\
& + \gamma \mathbb{E}_{\xi \sim \mu} [\mathcal{E}_{\mathbf{h}}(t, \xi, \mu)].
\end{aligned}$$

We also call the integrated q -function associated with the optimal policy $\boldsymbol{\pi}^*$ in (2.11) as the optimal integrated q -function, which is defined by

$$\begin{aligned}
q^*(t, \mu, \mathbf{h}) := & \frac{\partial J^*}{\partial t}(t, \mu) - \beta J^*(t, \mu) + \mathbb{E}_{\mu, \mathbf{h}} \left[H(t, \xi, \mu, a^{\mathbf{h}}, \partial_\mu J^*(t, \mu)(\xi), \partial_x \partial_\mu J^*(t, \mu)(\xi)) \right] \\
& + \gamma \mathbb{E}_{\xi \sim \mu} [\mathcal{E}_{\mathbf{h}}(t, \xi, \mu)].
\end{aligned}$$

Remark 3.2 In the framework of learning MFC, with a fixed small time step size Δt , the integrated Q -function $Q_{\Delta t}(t, \mu, \mathbf{h}; \boldsymbol{\pi})$ is related to our integrated q -function in the following sense:

$$Q_{\Delta t}(t, \mu, \mathbf{h}; \boldsymbol{\pi}) \approx J(t, \mu; \boldsymbol{\pi}) + q(t, \mu, \mathbf{h}; \boldsymbol{\pi}) \Delta t.$$

That is, $q(t, \mu, \mathbf{h}; \boldsymbol{\pi})$ is the first-order derivative of the integrated Q -function $Q_{\Delta t}(t, \mu, \mathbf{h}; \boldsymbol{\pi})$ with respect to Δt .

Contrary to the discrete time Q -learning algorithm that learns $Q_{\Delta t}$, the q -learning algorithm learns zeroth-order and first-order terms of $Q_{\Delta t}$ simultaneously. These two terms are independent of Δt and therefore robust to the choice of the time discretization in implementations (see the discussion and numerical comparison results in Jia and Zhou (2023) for the single agent's control problem).

Definition 3.3 If there exists a function $q_e : [0, T] \times \mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d) \times \mathcal{A} \rightarrow \mathbb{R}$ such that

$$q(t, \mu, \mathbf{h}; \boldsymbol{\pi}) - \gamma \mathbb{E}_{\xi \sim \mu} [\mathcal{E}_{\mathbf{h}}(t, \xi, \mu)] = \mathbb{E}_{\mu, \mathbf{h}} [q_e(t, \xi, \mu, a^{\mathbf{h}}; \boldsymbol{\pi})], \quad (3.13)$$

it is called the essential q -function, which plays the essential role in the policy improvement and the characterization of the optimal policy.

Contrary to the integrated q -function, the function q_e above is defined on $[0, T] \times \mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d) \times \mathcal{A}$, independent of \mathbf{h} . The following result ensures the existence of an essential q -function in our current framework.

Lemma 3.4 *There exists at least one essential q-function q_e given in Definition 3.3.*

Proof. Note from Definition 3.3 that

$$\begin{aligned} & q(t, \mu, \mathbf{h}; \boldsymbol{\pi}) - \gamma \mathbb{E}_{\xi \sim \mu} [\mathcal{E}_{\mathbf{h}}(t, \xi, \mu)] \\ &= \frac{\partial J}{\partial t}(t, \mu; \boldsymbol{\pi}) - \beta J(t, \mu; \boldsymbol{\pi}) + \mathbb{E}_{\mu, \mathbf{h}} \left[H(t, \xi, \mu, a^{\mathbf{h}}, \partial_{\mu} J(t, \mu; \boldsymbol{\pi})(\xi), \partial_x \partial_{\mu} J(t, \mu; \boldsymbol{\pi})(\xi)) \right], \end{aligned}$$

is linear in \mathbf{h} . Therefore, we can always define

$$\begin{aligned} q_e(t, x, \mu, a; \boldsymbol{\pi}) &:= \frac{\partial J}{\partial t}(t, \mu; \boldsymbol{\pi}) - \beta J(t, \mu; \boldsymbol{\pi}) \\ &\quad + H(t, x, \mu, a, \partial_{\mu} J(t, \mu; \boldsymbol{\pi})(x), \partial_x \partial_{\mu} J(t, \mu; \boldsymbol{\pi})(x)). \end{aligned} \tag{3.14}$$

Then \mathcal{I} defined in (2.13) can be written in terms of the essential q-function q_e that

$$\mathcal{I}(\boldsymbol{\pi}) = \frac{\exp(\frac{1}{\gamma} q_e(t, x, \mu, a; \boldsymbol{\pi}))}{\int_{\mathcal{A}} \exp(\frac{1}{\gamma} q_e(t, x, \mu, a; \boldsymbol{\pi})) da}.$$

Remark 3.5 *It is worth noting that the integrated Q-function $Q_{\Delta t}(t, \mu, \mathbf{h}; \boldsymbol{\pi})$ in general discrete time MFC problems has an implicit and nonlinear dependence on \mathbf{h} , which has three direct consequences: (i) in contrast to single-agent control problems where the optimal policy is in the form of Gibbs measure (3.9), the optimal policy $\boldsymbol{\pi}^*$ in the discrete time MFC Q-learning cannot be expressed in terms of Gibbs measure and does not exhibit any explicit distribution; (ii) the relationship in (3.8) does not hold in general even when there is no entropy that $\gamma = 0$; (iii) the counterpart of the essential q-function in the discrete time framework does not exist. Indeed, one cannot define a discrete time Q-function only depending on a (independent of \mathbf{h}) that satisfies the DPP, see the discussions in Gu et al. (2023).*

By contrast, as the first order derivative of $Q_{\Delta t}(t, \mu, \mathbf{h}; \boldsymbol{\pi})$ with respect to time Δt , the integrated q-function $q(t, \mu, \mathbf{h}; \boldsymbol{\pi})$ modified by the entropy term $\gamma \mathbb{E}_{\xi \sim \mu} [\mathcal{E}_{\mathbf{h}}(t, \xi, \mu)]$ has the nice linear dependence on the control policy \mathbf{h} in view of the integral relationship (3.13). As a consequence, solving the optimization problem $\sup_{\mathbf{h}} \{J(t, \mu; \boldsymbol{\pi}^*) + q(t, \mu, \mathbf{h}; \boldsymbol{\pi}^*) \Delta t\}$ leads to an explicit form of $\boldsymbol{\pi}^*$ in (2.11). This illustrates another advantage of working in continuous time framework for MFC as it is more convenient to characterize the optimal policy $\boldsymbol{\pi}^*$ and establish the relationship between the optimal policy $\boldsymbol{\pi}^*$ and the q-function to devise algorithms for policy improvement. We also note that the integrated q-function $q(t, \mu, \mathbf{h}; \boldsymbol{\pi})$ is more natural for the name “integrated function” as it (modified by the entropy) can be expressed as a linear double-integral in (3.13) that explicitly integrates the distribution of the state and the action of the population, while the discrete time integrated Q-function $Q_{\Delta t}(t, \mu, \mathbf{h}; \boldsymbol{\pi})$ aggregates the distribution in an implicit manner.

Finally, we stress that the integrated q-function is unique, but the essential q-function is not. In fact, any function $\kappa(x, \mu)$ satisfying $\int_{\mathbb{R}^d} \kappa(x, \mu) = 0$ can be added to (3.14) and still be an essential q-function.

4 Weak Martingale Characterizations

The first result below gives a characterization of the integrated q-function associated with a given policy $\boldsymbol{\pi} \in \Pi$ under the assumption that the value function J is known.

Theorem 4.1 (Characterization of the integrated q-function) *Let $\boldsymbol{\pi} \in \Pi$, its value function J and a continuous function $\hat{q} : [0, T] \times \mathcal{P}_2(\mathbb{R}^d) \times B_{\delta(\mu)}(\boldsymbol{\pi}) \rightarrow \mathbb{R}$ be given. Then \hat{q} is the integrated q-function if and only if \hat{q} satisfies*

$$\hat{q}(t, \mu, \boldsymbol{\pi}) = 0, \quad (4.1)$$

and for any $(t, \mu, \mathbf{h}) \in [0, T] \times \mathcal{P}_2(\mathbb{R}^d) \times B_{\delta(\mu)}(\boldsymbol{\pi})$, the value of

$$e^{-\beta s} J(s, \mathbb{P}_{X_s^{t, \mu, \mathbf{h}}}; \boldsymbol{\pi}) - e^{-\beta t} J(t, \mu; \boldsymbol{\pi}) + \int_t^s e^{-\beta t'} \left[\hat{r}(t', \mathbb{P}_{X_{t'}^{\mathbf{h}}}, \mathbf{h}) - \hat{q}(t', \mathbb{P}_{X_{t'}^{\mathbf{h}}}, \mathbf{h}) + \gamma \mathbb{E}[\mathcal{E}_{\mathbf{h}}(t, X_{t'}^{\mathbf{h}}, \mathbb{P}_{X_{t'}^{\mathbf{h}}})] \right] dt'$$

is 0, where $\{X_s^{\mathbf{h}}, t \leq s \leq T\}$ is the solution to (2.4) under the policy \mathbf{h} with $\mathbb{P}_{X_t^{\mathbf{h}}} = \mu$, and $\hat{r} : [0, T] \times \mathcal{P}_2(\mathbb{R}^d) \times B_{\delta(\mu)}(\boldsymbol{\pi}) \rightarrow \mathbb{R}$ is the aggregated reward defined by $\hat{r}(t, \mu, \mathbf{h}) = \mathbb{E}_{\mu, \mathbf{h}}[r(t, \xi, \mu, a^{\mathbf{h}})]$.

Proof. By (2.8) and Definition 3.1, it is a direct consequence that (4.1) holds.

Applying Itô's formula to $J(t', \mathbb{P}_{X_{t'}^{\mathbf{h}}}; \boldsymbol{\pi})$ between t and s , $0 \leq t < s \leq T$, we get that

$$\begin{aligned} & e^{-\beta s} J(s, \mathbb{P}_{X_s^{\mathbf{h}}}; \boldsymbol{\pi}) - e^{-\beta t} J(t, \mu; \boldsymbol{\pi}) + \int_t^s e^{-\beta t'} \left[\hat{r}(t', \mathbb{P}_{X_{t'}^{\mathbf{h}}}, \mathbf{h}) - \hat{q}(t', \mathbb{P}_{X_{t'}^{\mathbf{h}}}, \mathbf{h}) + \gamma \mathbb{E}[\mathcal{E}_{\mathbf{h}}(t, X_{t'}^{\mathbf{h}}, \mathbb{P}_{X_{t'}^{\mathbf{h}}})] \right] dt' \\ &= \int_t^s e^{-\beta t'} \left\{ \frac{\partial}{\partial t} J(t', \mathbb{P}_{X_{t'}^{\mathbf{h}}}; \boldsymbol{\pi}) - \beta J(t', \mathbb{P}_{X_{t'}^{\mathbf{h}}}; \boldsymbol{\pi}) \right. \\ &\quad + \int_{\mathbb{R}^d} \int_{\mathcal{A}} H(t', X_{t'}^{\mathbf{h}}, \mathbb{P}_{X_{t'}^{\mathbf{h}}}, a, \partial_{\mu} J(t', \mathbb{P}_{X_{t'}^{\mathbf{h}}}; \boldsymbol{\pi})(X_{t'}^{\mathbf{h}}), \partial_{\mathbf{x}} \partial_{\mu} J(t', \mathbb{P}_{X_{t'}^{\mathbf{h}}}; \boldsymbol{\pi})(X_{t'}^{\mathbf{h}})) \mathbf{h}(a|t', X_{t'}^{\mathbf{h}}, \mathbb{P}_{X_{t'}^{\mathbf{h}}}) da \mathbb{P}_{X_{t'}^{\mathbf{h}}}(dx) \\ &\quad \left. + \gamma \mathbb{E}[\mathcal{E}_{\mathbf{h}}(t, X_{t'}^{\mathbf{h}}, \mathbb{P}_{X_{t'}^{\mathbf{h}}})] - \hat{q}(t', \mathbb{P}_{X_{t'}^{\mathbf{h}}}, \mathbf{h}) \right\} dt' \\ &= \int_t^s e^{-\beta t'} \left[q(t', \mathbb{P}_{X_{t'}^{\mathbf{h}}}, \mathbf{h}; \boldsymbol{\pi}) - \hat{q}(t', \mathbb{P}_{X_{t'}^{\mathbf{h}}}, \mathbf{h}) \right] dt' = 0. \end{aligned}$$

Conversely, we need to show

$$\int_t^s e^{-\beta t'} \left[q(t', \mathbb{P}_{X_{t'}^{\mathbf{h}}}, \mathbf{h}; \boldsymbol{\pi}) - \hat{q}(t', \mathbb{P}_{X_{t'}^{\mathbf{h}}}, \mathbf{h}) \right] dt' = 0 \quad (4.2)$$

implies that $q(t, \mu, \mathbf{h}; \boldsymbol{\pi}) \equiv \hat{q}(t, \mu, \mathbf{h})$, which will be proved by contradiction. Denote $f(t, \mu, \mathbf{h}) = q(t, \mu, \mathbf{h}; \boldsymbol{\pi}) - \hat{q}(t, \mu, \mathbf{h})$. Then f is a continuous function that maps $[0, T] \times \mathcal{P}_2(\mathbb{R}^d) \times B_{\delta(\mu)}(\boldsymbol{\pi})$ to \mathbb{R} . Suppose that the claim does not hold. Then there exists $(t^*, \mu^*, \mathbf{h}^*) \in [0, T] \times \mathcal{P}_2(\mathbb{R}^d) \times B_{\delta(\mu)}(\boldsymbol{\pi})$ and $\epsilon > 0$ such that $f(t^*, \mu^*, \mathbf{h}^*) > \epsilon$. As f is continuous, there exists $\delta > 0$ such that when $\max\{|t - t^*|, \mathcal{W}_2(\mu, \mu^*)\} < \delta$, $f(t, \mu, \mathbf{h}^*) > \frac{\epsilon}{2}$. Now let us consider the process $X_s^{\mathbf{h}^*}$ starting from (t^*, μ^*) , i.e., $X_s^{\mathbf{h}^*}, t^* \leq s \leq T$ with $X_{t^*}^{\mathbf{h}^*} \sim \mu^*$. Define

$$\tau := \inf\{t > t^* : \mathcal{W}_2(\mathbb{P}_{X_t^{\mathbf{h}^*}}, \mu^*) > \delta\} \wedge (t^* + \delta).$$

Then it holds that $\int_{t^*}^{\tau} e^{-\beta t'} f(t', \mathbb{P}_{X_{t'}^{\mathbf{h}^*}}, \mathbf{h}^*) dt' > 0$, which contradicts with (4.2), and our conclusion holds.

Remark 4.2 *It is worth noting that if the policy \mathbf{h} in the integrated Q-function is replaced with a constant action a , the proof of Theorem 4.1 will fail. This is the technical reason for us to consider a policy \mathbf{h} in $Q_{\Delta t}$ in (3.12) instead of a constant action a .*

The following result characterizes both the integrated q-function and the value function J associated with a given policy π .

Theorem 4.3 (Characterization of the value function and the integrated q-function)

Let $\pi \in \Pi$, a continuous function $\hat{J} : [0, T] \times \mathcal{P}_2(\mathbb{R}^d) \rightarrow \mathbb{R}$ and a continuous function $\hat{q} : [0, T] \times \mathcal{P}_2(\mathbb{R}^d) \times B_{\delta(\mu)}(\pi) \rightarrow \mathbb{R}$ be given. Then \hat{J} and \hat{q} are respectively the value function and the integrated q-function associated with π if and only if \hat{J} and \hat{q} satisfy

$$\hat{J}(T, \mu) = \hat{g}(\mu), \quad \hat{q}(t, \mu, \pi) = 0, \quad (4.3)$$

and for any $(t, \mu, \mathbf{h}) \in [0, T] \times \mathcal{P}_2(\mathbb{R}^d) \times B_{\delta(\mu)}(\pi)$, the value of

$$e^{-\beta s} \hat{J}(s, \mathbb{P}_{X_s^{\mathbf{h}}}) - e^{-\beta t} \hat{J}(t, \mu) + \int_t^s e^{-\beta t'} [\hat{r}(t', \mathbb{P}_{X_{t'}^{\mathbf{h}}}, \mathbf{h}) - \hat{q}(t', \mathbb{P}_{X_{t'}^{\mathbf{h}}}, \mathbf{h}) + \gamma \mathbb{E}[\mathcal{E}_{\mathbf{h}}(t, X_{t'}^{\mathbf{h}}, \mathbb{P}_{X_{t'}^{\mathbf{h}}})]] dt' \quad (4.4)$$

is 0, where $\{X_s^{\mathbf{h}}, t \leq s \leq T\}$ is the solution to (2.4) under the stochastic policy \mathbf{h} with $\mathbb{P}_{X_t^{\mathbf{h}}} = \mu$. Furthermore, suppose that there exists a function $\hat{q}_e : [0, T] \times \mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d) \times \mathcal{A} \rightarrow \mathbb{R}$ such that $\hat{q}(t, \mu, \mathbf{h}) - \gamma \mathbb{E}_{\xi \sim \mu}[\mathcal{E}_{\mathbf{h}}(t, \xi, \mu)] = \mathbb{E}_{\mu, \mathbf{h}}[\hat{q}_e(t, \xi, \mu, a^{\mathbf{h}})]$ for any $\mathbf{h} \in B_{\delta(\mu)}(\pi)$, and it holds further that $\pi(a|t, x, \mu) = \frac{\exp\{\frac{1}{\gamma} \hat{q}_e(t, x, \mu, a)\}}{\int_{\mathcal{A}} \exp\{\frac{1}{\gamma} \hat{q}_e(t, x, \mu, a)\} da}$, then π is an optimal policy and \hat{J} is the optimal value function, i.e. $\hat{J} = J^*$.

Proof. The “if” part is a direct consequence by following the same argument as in the proof of Theorem 4.1. We therefore only prove the “only if” direction. Note that

$$e^{-\beta s} \hat{J}(s, \mathbb{P}_{X_s^{\mathbf{h}}}) - e^{-\beta t} \hat{J}(t, \mu) + \int_t^s [\hat{r}(t', \mathbb{P}_{X_{t'}^{\mathbf{h}}}, \mathbf{h}) - \hat{q}(t', \mathbb{P}_{X_{t'}^{\mathbf{h}}}, \mathbf{h}) + \gamma \mathbb{E}[\mathcal{E}_{\mathbf{h}}(t, X_{t'}^{\mathbf{h}}, \mathbb{P}_{X_{t'}^{\mathbf{h}}})]] dt' = 0.$$

From the proof of Theorem 4.1, it holds that

$$\begin{aligned} \frac{\partial}{\partial t} \hat{J}(t, \mu) - \beta \hat{J}(t, \mu) + \mathbb{E} \left[\int_{\mathcal{A}} H(t, \xi, \mu, a, \partial_{\mu} \hat{J}(t, \mu)(\xi), \partial_x \partial_{\mu} \hat{J}(t, \mu)(\xi)) \mathbf{h}(a|t, \xi, \mu) da \right] \\ + \gamma \mathbb{E}_{\xi \sim \mu}[\mathcal{E}_{\mathbf{h}}(t, \xi, \mu)] - \hat{q}(t, \mu, \mathbf{h}) = 0. \end{aligned}$$

Letting $\mathbf{h} = \pi$ and combining with $\hat{q}(t, \mu, \pi) = 0$ yields that

$$\begin{aligned} \frac{\partial}{\partial t} \hat{J}(t, \mu) - \beta \hat{J}(t, \mu) + \mathbb{E} \left[\int_{\mathcal{A}} \{H(t, \xi, \mu, a, \partial_{\mu} \hat{J}(t, \mu)(\xi), \partial_x \partial_{\mu} \hat{J}(t, \mu)(\xi)) \right. \\ \left. - \gamma \pi(a|t, \xi, \mu)\} \pi(a|t, \xi, \mu) da \right] = 0. \end{aligned}$$

This, together with the terminal condition $\hat{J}(T, \mu) = \hat{g}(\mu)$, yields that $\hat{J}(t, \mu) = J(t, \mu; \pi)$ by virtue of Feynman-Kac formula (2.8). Furthermore, based on Theorem 4.1, one has $\hat{q}(t, \mu, \mathbf{h}) = q(t, \mu, \mathbf{h}; \pi)$.

Finally, if $\pi(a|t, \mu, x) = \frac{\exp\{\frac{1}{\gamma} \hat{q}_e(t, x, \mu, a)\}}{\int_{\mathcal{A}} \exp\{\frac{1}{\gamma} \hat{q}_e(t, x, \mu, a)\} da}$, then $\pi = \mathcal{I}\pi$. This implies that π is an optimal policy and \hat{J} is the optimal value function J^* .

We end this subsection with the characterization of the optimal value function and optimal essential q-function q_e .

Theorem 4.4 (Characterization of the optimal value function and essential q-function)

Let $\hat{J}^* : [0, T] \times \mathcal{P}_2(\mathbb{R}^d) \rightarrow \mathbb{R}$ and $\hat{q}_e^* : [0, T] \times \mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d) \times \mathcal{A} \rightarrow \mathbb{R}$ be continuous functions.

Then \hat{J}^* , \hat{q}_e^* and $\hat{\pi}^*(a|t, x, \mu) := \frac{\exp\{\frac{1}{\gamma}\hat{q}_e^*(t, x, \mu, a)\}}{\int_{\mathcal{A}} \exp\{\frac{1}{\gamma}\hat{q}_e^*(t, x, \mu, a)\} da}$ are respectively the optimal value function J^* , the optimal essential q-function q_e^* and the associated optimal policy π^* if and only if

$$\hat{J}^*(T, \mu) = \hat{g}(\mu), \quad \int_{\mathbb{R}^d} \log \int_{\mathcal{A}} \exp\left\{\frac{1}{\gamma}\hat{q}_e^*(t, x, \mu, a)\right\} da \mu(dx) = 0, \quad (4.5)$$

and for any $(t, \mu, \mathbf{h}) \in [0, T] \times \mathcal{P}_2(\mathbb{R}^d) \times B_{\delta(\mu)}(\hat{\pi}^*)$, it holds that

$$\begin{aligned} e^{-\beta s} \hat{J}^*(s, \mathbb{P}_{X_s^{\mathbf{h}}}) - e^{-\beta t} \hat{J}^*(t, \mu) + \int_t^s e^{-\beta t'} \left\{ \hat{r}(t', \mathbb{P}_{X_{t'}^{\mathbf{h}}}, \mathbf{h}) \right. \\ \left. - \int_{\mathbb{R}^d \times \mathcal{A}} \hat{q}_e^*(t', X_{t'}^{\mathbf{h}}, \mathbb{P}_{X_{t'}^{\mathbf{h}}}, a) \mathbf{h}(a|t', X_{t'}^{\mathbf{h}}, \mathbb{P}_{X_{t'}^{\mathbf{h}}}) da \mathbb{P}_{X_{t'}^{\mathbf{h}}}(dx) \right\} dt' = 0. \end{aligned} \quad (4.6)$$

Remark 4.5 Theorem 4.4 characterizes the optimal value function and the optimal essential q-function in terms of a weak martingale condition (4.6) using all test policies and the consistency condition (4.5), which is the foundation for designing learning algorithms from the social planner’s perspective. On one hand, utilizing the martingale condition in (4.6) requires us to employ the test policy \mathbf{h} to generate samples and observations instead of the target policy π , which has a similar flavor of the off-policy q-learning for stochastic control by a single agent in [Jia and Zhou \(2023\)](#) that is based on observations of the given behavior policy. On the other hand, we highlight that the feature of mean-field interactions requires us to test all policies \mathbf{h} in the neighbourhood of the target policy π to fully characterize the optimal value function and the optimal q-function. The necessity of searching all test policies differs significantly from the standard off-policy q-learning in [Jia and Zhou \(2023\)](#) where any given behavior policy together with the consistency condition is sufficient to fully characterize the optimal value function and the optimal q-function.

Remark 4.6 We emphasize that both the integrated q-function and the essential q-function are crucial in our current setting. The integrated q-function is important from the theoretical perspective while the essential q-function is convenient for implementing the algorithms.

Firstly, recall that in single agent control problems, q-function originates from the time derivative of Q-function in the discrete time framework, which is the main reason that the proposed q-learning algorithms are regarded as the continuous-time counterpart of Q-learning algorithms. As noted in [Remark 3.5](#), our integrated q-function is also the time derivative of the discrete-time integrated Q-function in MFC problems as studied in [Gu et al. \(2023\)](#); [Carmona et al \(2023\)](#), however, the discrete-time version of the essential q-function does not exist! Therefore, without first introducing the integrated q-function, it will be difficult to introduce the definition of the essential q-function. Secondly, as the essential q-function is not unique (see [Remark 3.5](#)), we have to first show that the integrated q-function and the value function can be fully characterized by the weak martingale condition (see [Theorem 4.1](#)), which then yields the martingale condition of the essential q-function described in [Theorem 4.4](#). For the purpose of the implementation of the algorithm, we can then replace the integrated q-function by the essential q-function in [Theorem 4.4](#) using the integral representation in [\(3.13\)](#).

Remark 4.7 One relevant work is [Frikha et al. \(2023\)](#), which investigated the policy gradient algorithm for continuous time MFC problems. It is worth noting that both RL approaches have

pros and cons: In their representation of the policy gradient (see (3.3) in Theorem 3.1 in [Frikha et al. \(2023\)](#)), there is an additional term \mathcal{H} that can only be computed under the assumption when the model is separable. This term \mathcal{H} does not appear in our q-learning approach, which is a main advantage of our q-learning algorithm to cope with MFC problems. However, as a price to pay, we have to utilize infinitely many test polices \mathbf{h} to fully determine the integrated q-function, which is not needed in [Frikha et al. \(2023\)](#). When the population distribution cannot be observed, we have also established a comparison between the policy gradient algorithm in [Frikha et al. \(2023\)](#) and the q-learning algorithm from the representative agent's perspective, see proposition 5.6 in [Wei et al. \(2024\)](#).

5 Algorithms under Continuous Time q-Learning

In this subsection, we assume that the social planner has a simulator for the state distribution and we learn directly the optimal value function and the optimal q-function according to Theorem 4.4. First, the parametrized function approximators J^θ and q_e^ψ are chosen such that the consistency condition (4.5) is satisfied that

$$J^\theta(T, \mu) = \hat{g}(\mu) \text{ and } \int_{\mathbb{R}^d} \log \int_{\mathcal{A}} \exp \left\{ \frac{1}{\gamma} q_e^\psi(t, x, \mu, a) \right\} da \mu(dx) = 0. \quad (5.1)$$

and we have the parameterized policy

$$\pi^\psi(a|t, x, \mu) = \frac{\exp \left\{ \frac{1}{\gamma} q_e^\psi(t, x, \mu, a) \right\}}{\int_{\mathcal{A}} \exp \left\{ \frac{1}{\gamma} q_e^\psi(t, x, \mu, a) \right\} da}.$$

In what follows, it is assumed that the constant $\int_{\mathcal{A}} \exp \left\{ \frac{1}{\gamma} q_e^\psi(t, x, \mu, a) \right\} da$ is known. We first devise the q-learning algorithm in an offline setting. We can utilize the weak martingale condition (4.4) in Theorem 4.4 to devise the updating rules for parameters by minimizing the loss function robust with respect to all test policies in an offline setting. Denote

$$M_s^{\theta, \psi, \mathbf{h}} := e^{-\beta s} J^\theta(s, \mathbb{P}_{X_s^{\mathbf{h}}}) + \int_0^s e^{-\beta t'} \left\{ \hat{r}(t', \mathbb{P}_{X_{t'}^{\mathbf{h}}}, \mathbf{h}) - \int_{\mathbb{R}^d} \int_{\mathcal{A}} q_e^\psi(t', x, \mathbb{P}_{X_{t'}^{\mathbf{h}}}, a) \mathbf{h}(a|t', x, \mathbb{P}_{X_{t'}^{\mathbf{h}}}) da \mathbb{P}_{X_{t'}^{\mathbf{h}}} \right\} dt',$$

where $\mathbb{P}_{X_t^{\mathbf{h}}} := \mathbb{P}_{X_t^{0, \mu, \mathbf{h}}}$ is the population state distribution associated with \mathbf{h} starting from 0.

According to (4.6), the parameters θ and ψ will be updated by optimizing mean-square TD error averaged over all test policies that

$$\inf_{\theta, \psi} L(\theta, \psi) := \frac{1}{2} \inf_{\theta, \psi} \int_{\mathbf{h} \in B_{\delta(\mu)}(\pi^\psi)} |\dot{M}_s^{\theta, \psi, \mathbf{h}}|^2 ds \nu(d\mathbf{h}), \quad (5.2)$$

where $\dot{M}_s^{\theta, \psi, \mathbf{h}}$ is the time derivative of $M_s^{\theta, \psi, \mathbf{h}}$ and called TD error rate and ν is a probability measure defined on $B_{\delta(\mu)}(\pi^\psi) \subset \Pi$ with $\text{supp}(\nu) = B_{\delta(\mu)}(\pi^\psi)$. The loss function $L(\theta, \psi)$ has a similar sprit of mean-square temporal difference error in the setting of the deterministic dynamics, see [Doya \(2020\)](#) and also the discussion in section 3.1 in [Jia and Zhou \(2022b\)](#). The infimum of the loss function $\inf_{\theta, \psi} L(\theta, \psi) = 0$ implies that the inside integrand is zero for almost surely t and all \mathbf{h} , which, combined with Theorem 4.6, indicates that J^θ and q^ψ are indeed respectively the optimal value function and the optimal Q-function.

The key issue in the implementation is that we need all test policy \mathbf{h} to generate sample trajectories to achieve the purpose of learning. In practice, it is impossible to implement all test policies \mathbf{h} to ensure (5.2) to hold. In this paper, we propose a parametrization method that is based on the fact the test policy has the same parametrized form as the policy π^ψ yet with distinct parameter $\tilde{\psi} \in B_\delta(\psi)$, where $B_\delta(\psi)$ is a ball with radius δ centered at ψ . It is then sufficient to consider the average of the weak martingale loss over $\tilde{\psi} \in B_\delta(\psi)$.

$$\inf_{\theta, \psi} \int_{\tilde{\psi} \in B_\delta(\psi)} |\dot{M}_s^{\theta, \psi, \mathbf{h}^{\tilde{\psi}}}|^2 ds d\tilde{\psi}. \quad (5.3)$$

Average-based test policies: As the practical implementation is of discrete fashion, we choose a sequence of test policies $\mathbf{h}^{\tilde{\psi}^1}, \mathbf{h}^{\tilde{\psi}^2}, \dots, \mathbf{h}^{\tilde{\psi}^M}$ using the parameterized policy π^ψ . More precisely, at the beginning of each episode j , we use ψ to randomly generate M parameters $\tilde{\psi}^1, \dots, \tilde{\psi}^M$, which corresponds to M test policies $\mathbf{h}^{\tilde{\psi}^1}, \dots, \mathbf{h}^{\tilde{\psi}^M}$. For example, $\tilde{\psi}^1, \dots, \tilde{\psi}^M$ are i.i.d. drawn from $\psi \cdot \mathcal{U}([p(j), q(j)])$, where $\mathcal{U}([p(j), q(j)])$ is uniform distribution on $[p(j), q(j)]$.

In this case, we discretize $[0, T]$ on the grid $\{t_k = k\Delta t, k = 0, 1, \dots, K\}$ and hence the weak martingale loss function is averaged over all test policies

$$\frac{1}{2M} \sum_{m=1}^M \sum_{k=0}^{K-1} |\delta_k^{\theta, \psi, \mathbf{h}^{\tilde{\psi}^m}}|^2 \Delta t, \quad (5.4)$$

where $\mathbb{P}_{X_t^{\mathbf{h}^{\tilde{\psi}^m}}} := \mathbb{P}_{X_t^{0, \mu, \mathbf{h}^{\tilde{\psi}^m}}}$ is the population state distribution associated with $\mathbf{h}^{\tilde{\psi}^m}$, $1 \leq m \leq M$ starting from 0, and $\delta_k^{\theta, \psi, \mathbf{h}^{\tilde{\psi}^m}}$ is given by

$$\begin{aligned} \delta_k^{\theta, \psi, \mathbf{h}^{\tilde{\psi}^m}} &= e^{-\beta t_k} \left(\frac{J^\theta(t_{k+1}, \mu_{t_{k+1}}^m) - J^\theta(t_k, \mu_{t_k}^m)}{\Delta t} + \hat{r}_{t_k}^m - \beta J^\theta(t_k, \mu_{t_k}^m) \right. \\ &\quad \left. - \int_{\mathbb{R}^d \times \mathcal{A}} q_e^\psi(t_k, x, \mu_{t_k}^m, a) \mathbf{h}^{\tilde{\psi}^m}(a|t_k, x, \mu_{t_k}^m) da \mu_{t_k}^m(dx) \right). \end{aligned}$$

It is shown in Theorem 5.4 in [Wei et al. \(2024\)](#) that as $M \rightarrow +\infty$ and $\Delta t \rightarrow 0$, the solution of (5.4) converges to that of (5.3).

We then apply the vanilla gradient descent to update θ and ψ :

$$\theta \leftarrow \theta - \alpha_\theta \frac{1}{M} \sum_{m=1}^M \Delta^m \theta, \quad \psi \leftarrow \psi - \alpha_\psi \frac{1}{M} \sum_{m=1}^M \Delta^m \psi,$$

where $\mu_{t_k}^m$ and $\hat{r}_{t_k}^m$, $0 \leq k \leq K-1$ are observed state distribution and observed aggregated reward associated with the test policy $\mathbf{h}^{\tilde{\psi}^m}$, and α_θ and α_ψ are learning rates and

$$\begin{aligned} G_{t_k}^m &= e^{-2\beta t_k} \left(J^\theta(t_{k+1}, \mu_{t_{k+1}}^m) - J^\theta(t_k, \mu_{t_k}^m) + (\hat{r}_{t_k}^m - \beta J^\theta(t_k, \mu_{t_k}^m)) \right. \\ &\quad \left. - \int_{\mathbb{R}^d \times \mathcal{A}} q_e^\psi(t_k, x, \mu_{t_k}^m, a) \mathbf{h}^{\tilde{\psi}^m}(a|t_k, x, \mu_{t_k}^m) da \mu_{t_k}^m(dx) \right) \Delta t, \end{aligned} \quad (5.5)$$

$$\Delta^m \theta = \sum_{k=0}^{K-1} G_{t_k}^m \left(\frac{1}{\Delta t} \left(\frac{\partial J^\theta}{\partial \theta}(t_{k+1}, \mu_{t_{k+1}}^m) - \frac{\partial J^\theta}{\partial \theta}(t_k, \mu_{t_k}^m) \right) - \beta J^\theta(t_k, \mu_{t_k}^m) \right), \quad (5.6)$$

$$\Delta^m \psi = - \sum_{k=0}^{K-1} G_{t_k}^m \int_{\mathbb{R}^d \times \mathcal{A}} \frac{\partial q_e^\psi}{\partial \psi}(t_i, x, \mu_{t_i}^m, a) \mathbf{h}^{\tilde{\psi}^m}(a|t_i, x, \mu_{t_i}^m) da \mu_{t_i}^m(dx). \quad (5.7)$$

The pseudo-code is described in Algorithm 1.

Remark 5.1 *Generally speaking, it is computationally costly to evaluate the integrals in (5.5)-(5.7). However, in numerical examples in Section 6, we can leverage the model structure to explicitly compute these integrals. In particular, in the first example in subsection 6.1, q_e^ψ is quadratic in x and a , and the parameterized test policy $\mathbf{h}^{\tilde{\psi}}$ is taken as the normal distribution so that the integrals become functions of expectation and variance of the population distribution μ . In general cases, we may have to compute these integrals by a Monte Carlo approach or the approximation of q_e^ψ using simple functions, such as polynomials of x and a .*

As the social planner has access to the full information of the population distribution, the population distribution is generated by a simulator $(\mu', \hat{r}) = \text{Environment}_{\Delta t}(t, \mu, \mathbf{h})$ in Algorithms 1-2 that is based on the time discretized version of Fokker-Planck equation. More precisely, consider the time discretized SDE

$$X_{t_{k+1}} \approx X_{t_k} + b(t_k, X_{t_k}, \mu_{t_k}, a^{\mathbf{h}}) \Delta t + \sigma(t_k, X_{t_k}, \mu_{t_k}, a^{\mathbf{h}}) (W_{t_{k+1}} - W_{t_k}), \quad a^{\mathbf{h}} \sim \mathbf{h}(\cdot|t_k, X_{t_k}, \mu_{t_k}).$$

Denote the probability transition function of by $p(dx'|t, X_{t_k}, \mu_{t_k}, a^{\mathbf{h}}) := \mathbb{P}(X_{t_{k+1}} \in dx'|X_{t_k})$, then p is a normal distribution $\mathcal{N}(X_{t_k} + b(t_k, X_{t_k}, \mu_{t_k}, a^{\mathbf{h}}) \Delta t, \sigma \sigma^\top(t_k, X_{t_k}, \mu_{t_k}, a^{\mathbf{h}}) \Delta t)$. By the law of total probability

$$\mu_{t_{k+1}}(dx') = \int_{\mathbb{R}^d} \int_{\mathcal{A}} p(dx'|t, x, \mu_{t_k}, a) \mathbf{h}(a|t, x, \mu_{t_k}) da \mu_{t_k}(dx). \quad (5.8)$$

(5.8) implies the evolution of population state distribution μ_{t_k} over the time.

Remark 5.2 *In practice, instead of computing (5.8), we assume that the social planner has access to an environment simulator that updates moments of the population distribution. Precisely, this is represented by*

$$(\bar{\mu}_{t_{k+1}}, \bar{\mu}_{2,t_{k+1}}, \dots, \bar{\mu}_{m,t_{k+1}}, \hat{r}) = \text{Environment}_{\Delta t}(t_k, \bar{\mu}_{t_k}, \bar{\mu}_{2,t_k}, \dots, \bar{\mu}_{m,t_k}, \mathbf{h}), \quad m \in \mathbb{N}^*,$$

where $\bar{\mu}_{t,m} = \int_{\mathbb{R}^d} x^m \mu_t(dx)$ denotes the m -th moment of μ . For example, moments up to second order are sufficient for the linear quadratic framework. See (6.7)-(6.8) for the simulator of the mean-variance example and (6.12) for mean-field optimal consumption problem.

Remark 5.3 *The information structure (i.e. what an agent can observe) plays a significant role in the reinforcement learning of MFC problems. In this study, we consider MFC from the perspective of the social planner, who can observe macroscopic quantities, including population distributions μ_t and aggregated rewards \hat{r}_t . In contrast, Wei et al. (2024) studies MFC from the viewpoint of the representative agent, who lacks the access to population distributions μ_t and can only observe his own states x_t and rewards r_t . It is assumed therein that the representative agent updates the population distributions based on his observed states. Consequently, different information structures will lead to distinct q -functions with different martingale characterizations and environment simulators.*

We can also similarly devise the q-learning algorithm in an online-setting where the parameters are updated in real-time. Similar to the previous algorithm, we can minimize the loss function over all average-based test policies $\{\mathbf{h}^{\psi^1}, \dots, \mathbf{h}^{\psi^M}\}$ between t_k and t_{k+1} that

$$\frac{1}{2M} \sum_{m=1}^M \left| J^\theta(t_{k+1}, \mathbb{P}_{X_{t_{k+1}}^{\mathbf{h}^{\psi^m}}}) - J^\theta(t_k, \mathbb{P}_{X_{t_k}^{\mathbf{h}^{\psi^m}}}) + \left(\hat{r}(t_k, \mathbb{P}_{X_{t_k}^{\mathbf{h}^{\psi^m}}}, \mathbf{h}^{\psi^m}) - \int_{\mathbb{R}^d} \int_{\mathcal{A}} q_e^\psi(t_k, x, \mathbb{P}_{X_{t_k}^{\mathbf{h}^{\psi^m}}}, a) \mathbf{h}^{\psi^m}(a|t_k, x, \mathbb{P}_{X_{t_k}^{\mathbf{h}^{\psi^m}}}) \mathbb{P}_{X_{t_k}^{\mathbf{h}^{\psi^m}}}(dx) - \beta J^\theta(t_k, \mathbb{P}_{X_{t_k}^{\mathbf{h}^{\psi^m}}}) \right) \Delta t \right|^2.$$

We can then apply gradient descent (GD) to the above loss function after the time discretization of $[0, T]$ on the grids $\{t_k = k\Delta t, k = 0, \dots, K-1\}$ and obtain the updating rules for θ and ψ that

$$\theta \leftarrow \theta - \alpha_\theta \frac{1}{M} \sum_{m=1}^M \Delta^m \theta, \quad \psi \leftarrow \psi - \alpha_\psi \frac{1}{M} \sum_{m=1}^M \Delta^m \psi,$$

where $\mu_{t_k}^m, 0 \leq k \leq K-1$ and $\hat{r}_{t_k}^m, 0 \leq k \leq K-1$ are observed state distribution and reward associated to the test policy \mathbf{h}^{ψ^m} , and α_θ and α_ψ are learning rates and

$$\begin{aligned} \delta_k^m &= J^\theta(t_{k+1}, \mu_{t_{k+1}}^m) - J^\theta(t_k, \mu_{t_k}^m) + \left(\hat{r}_{t_k} - \beta J^\theta(t_k, \mu_{t_k}^m) \right. \\ &\quad \left. - \int_{\mathbb{R}^d} \int_{\mathcal{A}} q_e^\psi(t_k, x, \mu_{t_k}^m, a) \mathbf{h}^{\psi^m}(a|t_k, x, \mu_{t_k}^m) \mu_{t_k}^m(dx) \right) \Delta t, \end{aligned} \quad (5.9)$$

$$\Delta^m \theta = \delta_k^m \left(\frac{\partial J^\theta}{\partial \theta}(t_{k+1}, \mu_{t_{k+1}}^m) - \frac{\partial J^\theta}{\partial \theta}(t_k, \mu_{t_k}^m) - \beta J^\theta(t_k, \mu_{t_k}^m) \Delta t \right), \quad (5.10)$$

$$\Delta^m \psi = -\delta_k^m \int_{\mathbb{R}^d} \int_{\mathcal{A}} \frac{\partial q_e^\psi}{\partial \psi}(t_k, x, \mu_{t_k}^m, a) \mathbf{h}^{\psi^m}(a|t, x, \mu_{t_k}^m) \mu_{t_k}^m(dx) \Delta t. \quad (5.11)$$

The pseudo-code is described in Algorithm 2. We remark that the dynamic is deterministic from the perspective of the social planner and thus the updating rule has a similar spirit as Doya's Temporal Difference (TD) algorithm for deterministic dynamics [Doya \(2020\)](#).

Remark 5.4 *For Algorithms 1 and 2, we again emphasize that the test policy \mathbf{h} is used to the environment simulator instead of the target policy π , which makes our algorithms similar to the so-called off-policy in conventional Q-learning. However, the feature of mean-field interactions requires us to explore all test policies to meet the weak martingale condition, and hence makes our algorithms different from the standard off-policy learning. The parameters in the policy (or q-function) are used and updated in the learning procedure but the resulting policy π does not participate in the procedure directly. All samples and observations are based on our chosen test policy \mathbf{h} , which can be adaptively updated by our proposed two different methods.*

Algorithm 1 Offline q-Learning Algorithm

Inputs: initial state distribution μ_0 , horizon T , time step Δt , number of mesh grids K , number of test policies M , initial learning rates α_θ and α_ψ , functional forms of parameterized value function $J^\theta(\cdot, \cdot)$ and $q_e^\psi(\cdot, \cdot, \cdot, \cdot)$ satisfying (4.5) and temperature parameter γ .

Required program: Moments simulator or environment simulator $(\mu', \hat{r}) = \text{Environment}_{\Delta t}(t, \mu, \mathbf{h}^{\tilde{\psi}})$ that takes current time–state distribution pair (t, μ) and the test policy $\mathbf{h}^{\tilde{\psi}}$ as inputs and generates state distribution μ' at time $t + \Delta t$ and the aggregated reward \hat{r} at time t as outputs.

Learning procedure:

Initialize θ and ψ .

for episode $j = 1$ **to** N **do**

Observe the initial state distribution μ_0 and store $\mu_{t_k}^m \leftarrow \mu_0$.

for $m = 1$ **to** M **do**

Draw $\tilde{\psi}^m$ from $\psi \cdot \mathcal{U}([p(j), q(j)])$ and set the test policy $\mathbf{h}^{\tilde{\psi}^m}$.

Initialize $k = 0$.

while $k < K$ **do**

Apply the test policy $\mathbf{h}^{\tilde{\psi}^m}$ to the environment simulator $(\mu, \hat{r}) = \text{Environment}_{\Delta t}(t_k, \mu_{t_k}^m, \mathbf{h}^{\tilde{\psi}^m})$, and observe the new state distribution μ and the aggregated reward \hat{r} as output. Store $\mu_{t_{k+1}}^m \leftarrow \mu$ and $\hat{r}_{t_k}^m \leftarrow \hat{r}$.

Update $k \leftarrow k + 1$.

end while

For every $k = 0, 1, \dots, K - 1$, compute $G_{t_k:T}^m$, $\Delta^m \theta$ and $\Delta^m \psi$ according to (5.5)-(5.7).

end for

Update θ and ψ by $\theta \leftarrow \theta - \alpha_\theta \frac{1}{M} \sum_{m=1}^M \Delta^m \theta$, $\psi \leftarrow \psi - \alpha_\psi \frac{1}{M} \sum_{m=1}^M \Delta^m \psi$.

end for

Algorithm 2 Online q-Learning Algorithm

Inputs: initial state distribution μ_0 , horizon T , time step Δt , number of mesh grids K , number of behavior policies M , initial learning rates α_θ and α_ψ , functional forms of parameterized value function $J^\theta(\cdot, \cdot)$ and $q_e^\psi(\cdot, \cdot, \cdot, \cdot)$ satisfying (4.5) and temperature parameter γ .

Required program: Moments simulator or environment simulator $(\mu', \hat{r}) = Environment_{\Delta t}(t, \mu, \mathbf{h}^{\tilde{\psi}})$ that takes current time–state distribution pair (t, μ) and the test policy $\mathbf{h}^{\tilde{\psi}}$ as inputs and generates the state distribution μ' at time $t + \Delta t$ and the aggregated reward \hat{r} at time t as outputs.

Learning procedure:

Initialize θ and ψ .

for episode $j = 1$ **to** N **do**

Initialize $k = 0$. Observe the initial state distribution μ_0 and store $\mu_{t_k}^m \leftarrow \mu_0$, $m = 1, \dots, M$.

while $k < K$ **do**

for $m = 1$ **to** M **do**

Draw $\tilde{\psi}^m$ from $\psi \cdot \mathcal{U}([p(j), q(j)])$ and set the test policy $\mathbf{h}^{\tilde{\psi}^m}$. Apply the test policy $\mathbf{h}^{\tilde{\psi}^m}$ to environment simulator $(\mu, \hat{r}) = Environment_{\Delta t}(t_k, \mu_{t_k}^m, \mathbf{h}^{\tilde{\psi}^m})$, and observe new state distribution μ and the aggregated reward \hat{r} as output. Store $\mu_{t_{k+1}}^m \leftarrow \mu$ and $\hat{r}_{t_k}^m \leftarrow \hat{r}$.

Compute δ_k^m , $\Delta^m \theta$ and $\Delta^m \psi$ according to (5.9)-(5.11).

end for

Update θ and ψ by $\theta \leftarrow \theta - \alpha_\theta \frac{1}{M} \sum_{m=1}^M \Delta^m \theta$, $\psi \leftarrow \psi - \alpha_\psi \frac{1}{M} \sum_{m=1}^M \Delta^m \psi$.

Update $k \leftarrow k + 1$.

end while

end for

6 Financial Applications

6.1 Mean-Variance Portfolio Optimization

Let us consider the wealth process that satisfies the SDE that

$$dX_s^\pi = a_s^\pi (bds + \sigma dW_s), a_s^\pi \sim \pi(\cdot | s, X_s^\pi, \mathbb{P}_{X_s^\pi}), s \geq t, X_0 = x, \quad (6.1)$$

where a_s^π is the wealth amount invested in the risky asset at time s , b is the excess return and σ is the volatility. The learning mean-variance portfolio optimization problem with entropy regularizer is defined by

$$J(t, \mu; \pi) = \mathbb{E}[X_T^\pi] - \lambda \text{Var}(X_T^\pi) - \gamma \mathbb{E} \left[\int_t^T \int_{\mathbb{R}} \log \pi(a | s, X_s^\pi, \mathbb{P}_{X_s^\pi}) \pi(a | s, X_s^\pi, \mathbb{P}_{X_s^\pi}) da \right].$$

The Hamiltonian operator is given by $H(t, x, \mu, a, p, q) = bap + \frac{1}{2} \sigma^2 a^2 q$. It then follows that

$$\begin{aligned} \int_{\mathbb{R}} \exp \left\{ \frac{1}{\gamma} H(t, x, \mu, a, p, q) \right\} da &= \exp \left(-\frac{b^2 p^2}{2\sigma^2 q \gamma} \right) \int_{\mathbb{R}} \exp \left(\frac{1}{2\gamma} q \sigma^2 \left(a + \frac{bp}{q\sigma^2} \right)^2 \right) da \\ &= \exp \left(-\frac{b^2 p^2}{2\gamma \sigma^2 q} \right) \sqrt{-\frac{2\pi\gamma}{q\sigma^2}}. \end{aligned}$$

As $J^*(t, \mu)$ satisfies (2.10), we have $J^*(T, \mu) = -\lambda \text{Var}(\mu) + \bar{\mu}$ and the exploratory HJB equation is given by

$$\begin{aligned} \frac{\partial J^*}{\partial t}(t, \mu) + \gamma \int_{\mathbb{R}} \log \int_{\mathbb{R}} \exp \left\{ \frac{ba}{\gamma} \partial_\mu J^*(t, \mu)(x) + \frac{\sigma^2 a^2}{2\gamma} \partial_x \partial_\mu J^*(t, \mu)(x) \right\} da \mu(dx) & \quad (6.2) \\ = \frac{\partial J^*}{\partial t}(t, \mu) - \frac{b^2}{\sigma^2} \int_{\mathbb{R}} \frac{|\partial_\mu J^*(t, \mu)(x)|^2}{2\partial_x \partial_\mu J^*(t, \mu)(x)} \mu(dx) + \frac{\gamma}{2} \int_{\mathbb{R}} \log \frac{2\pi\gamma}{-\sigma^2 \partial_x \partial_\mu J^*(t, \mu)(x)} \mu(dx) & = 0. \end{aligned}$$

Denote $\text{Var}(\mu) := \int_{\mathbb{R}} (x - \bar{\mu})^2 \mu(dx)$, $\bar{\mu} := \int_{\mathbb{R}} x \mu(dx)$. We conjecture that $J^*(t, \mu)$ satisfies the quadratic form

$$J^*(t, \mu) = A(t) \text{Var}(\mu) + C(t) \bar{\mu} + D(t). \quad (6.3)$$

It then holds that

$$\begin{aligned} \frac{\partial J^*}{\partial t}(t, \mu) &= \dot{A}(t) \text{Var}(\mu) + \dot{C}(t) \bar{\mu} + \dot{D}(t), \\ \partial_\mu J^*(t, \mu)(x) &= 2A(t)(x - \bar{\mu}) + C(t), \quad \partial_x \partial_\mu J^*(t, \mu)(x) = 2A(t). \end{aligned}$$

Plugging these into the exploratory HJB equation (6.2), we get that

$$\left[\dot{A}(t) - \frac{b^2}{\sigma^2} A(t) \right] \text{Var}(\mu) + \dot{C}(t) \bar{\mu} + \left[\dot{D}(t) - \frac{b^2}{\sigma^2} \frac{C(t)^2}{4A(t)} + \frac{\gamma}{2} \log \frac{\pi\gamma}{-\sigma^2 A(t)} \right] = 0.$$

By the terminal conditions $A(T) = -\lambda$, $C(T) = 1$ and $D(T) = 0$, we can obtain the explicit solution of the ODEs that

$$A(t) = -\lambda \exp \left(\frac{b^2}{\sigma^2} (t - T) \right), \quad C(t) = 1,$$

$$D(t) = \frac{\gamma b^2}{4\sigma^2}(t-T)^2 - (t-T)\frac{\gamma}{2}\log\frac{\pi\gamma}{\sigma^2\lambda} + \frac{1}{4\lambda}\exp\left(-\frac{b^2}{\sigma^2}(t-T)\right) - \frac{1}{4\lambda}.$$

Then, we can take the essential q-function q_e^* as

$$q_e^*(t, x, \mu, a) = -\lambda\sigma^2 \exp\left(\frac{b^2}{\sigma^2}(t-T)\right) \left(a + \frac{b}{\sigma^2}(x - \bar{\mu}) - \frac{b}{2\lambda\sigma^2}\exp\left(-\frac{b^2}{\sigma^2}(t-T)\right)\right)^2 \\ - \frac{b^2\lambda}{\sigma^2} \exp\left(\frac{b^2}{\sigma^2}(t-T)\right) [(x - \bar{\mu})^2 - \text{Var}(\mu)] - \frac{b^2}{\sigma^2}(x - \bar{\mu}) - \frac{\gamma}{2}\log\frac{\pi\gamma}{\sigma^2\lambda} + \frac{\gamma b^2}{2\sigma^2}(t-T).$$

It then follows that the optimal policy π^* is

$$\pi^*(\cdot|t, x, \mu) = \mathcal{N}\left(-\frac{b}{\sigma^2}(x - \bar{\mu}) - \frac{1}{2\lambda}\exp\left(-\frac{b^2}{\sigma^2}(t-T)\right), \frac{\gamma}{2\lambda\sigma^2}\exp\left(-\frac{b^2}{\sigma^2}(t-T)\right)\right). \quad (6.4)$$

When model parameters b , σ and λ are unknown, we can derive the parameterized functions J^θ and q_e^ψ by

$$J^\theta(t, \mu) = -\frac{1}{4\theta_3}\exp(\theta_1(t-T))\text{Var}(\mu) + \bar{\mu} + \frac{\gamma\theta_1}{4}(t-T)^2 + \theta_2(t-T) \\ + \theta_3\exp(-\theta_1(t-T)) - \theta_3, \quad (6.5)$$

$$q_e^\psi(t, x, \mu, a) = -\frac{\exp(\psi_1 + \psi_2(t-T))}{2}(a + \psi_3(x - \bar{\mu}) + \psi_4\exp(-\psi_2(t-T)))^2 - \frac{\gamma}{2}\log(2\pi\gamma) + \frac{\gamma}{2}\psi_1 \\ + \frac{\psi_2\gamma}{2}(t-T) - \psi_2(x - \bar{\mu}) + \psi_5\exp(\psi_2(t-T))((x - \bar{\mu})^2 - \text{Var}(\mu)), \quad (6.6)$$

where $\theta = (\theta_1, \theta_2, \theta_3)^\top \in \mathbb{R}^3$, $\psi = (\psi_1, \dots, \psi_5) \in \mathbb{R}^5$. Note that the above parameterizations satisfy the constraints in (5.1). It then follows that $\pi^\psi(\cdot|t, x, \mu) = \mathcal{N}(-\psi_3(x - \bar{\mu}) - \psi_4\exp(-\psi_2(t-T)), \gamma\exp(-\psi_1 - \psi_2(t-T)))$. Note that ψ_5 does not make any contribution to learn π^ψ and is redundant to be learnt by the q-learning algorithm.

The following simulator will be used in Algorithms 1-2 to generate sample trajectories.

Simulator Under the policy $\mathbf{h}^{\tilde{\psi}}$, the equation (6.1) becomes

$$d\tilde{X}_t = b \int_{\mathbb{R}} a \mathbf{h}^{\tilde{\psi}}(a|t, \tilde{X}_t, \mathbb{P}_{\tilde{X}_t}) da dt + \sigma \sqrt{a^2 \mathbf{h}^{\tilde{\psi}}(a|t, \tilde{X}_t, \mathbb{P}_{\tilde{X}_t})} dW_t \\ = -b\left(\tilde{\psi}_3(\tilde{X}_t - \mathbb{E}[\tilde{X}_t]) + \tilde{\psi}_4\exp(-\tilde{\psi}_2(t-T))\right) dt \\ + \sigma \sqrt{\left(\tilde{\psi}_3(\tilde{X}_t - \mathbb{E}[\tilde{X}_t]) + \tilde{\psi}_4\exp(-\tilde{\psi}_2(t-T))\right)^2 + \gamma\exp(-\tilde{\psi}_1 - \tilde{\psi}_2(t-T))} dW_t.$$

First we calculate the mean of \tilde{X}_t that $\bar{\mu}_t = \bar{\mathbb{P}}_{\tilde{X}_t} \mathbf{h}^{\tilde{\psi}}$ by taking the expectation on both sides of the above SDE, which yields that $d\mathbb{E}[\tilde{X}_t] = -b\tilde{\psi}_4\exp(-\tilde{\psi}_2(t-T))dt$. We thus deduce that

$$\bar{\mu}_{t_{k+1}} \simeq \bar{\mu}_{t_k} - b\tilde{\psi}_4\exp(-\tilde{\psi}_2(t_k - T))\Delta t. \quad (6.7)$$

We next compute the variance of \tilde{X}_t : $\text{Var}(\mathbb{P}_{\tilde{X}_t}) = \text{Var}(\tilde{X}_t)$ by applying Itô's formula to $(\tilde{X}_t - \mathbb{E}[\tilde{X}_t])^2$ and then taking expectation that

$$d\text{Var}(\tilde{X}_t) = \left((\sigma^2\tilde{\psi}_3^2 - 2b\tilde{\psi}_3)\text{Var}(\tilde{X}_t) + \sigma^2\tilde{\psi}_4^2\exp(-2\tilde{\psi}_2(t-T)) + \gamma\exp(-\tilde{\psi}_1 - \tilde{\psi}_2(t-T))\right) dt.$$

	θ_1	θ_2	θ_3	ψ_1	ψ_2	ψ_3	ψ_4
true value	0.25	-0.358	0.1667	-0.288	0.25	1	-0.333
learnt by Algorithm 1	0.249	-0.356	0.162	-0.288	0.234	1.012	-0.334

The Euler approximation of $\text{Var}(\tilde{X}_t)$ is given by

$$\begin{aligned} \text{Var}(\tilde{X}_{t_{k+1}}) = & \text{Var}(\tilde{X}_{t_k}) + \left((\sigma^2 \tilde{\psi}_3^2 - 2b\tilde{\psi}_3) \text{Var}(\tilde{X}_{t_k}) + \sigma^2 \tilde{\psi}_4^2 \exp(-2\tilde{\psi}_2(t_k - T)) \right. \\ & \left. + \gamma \sigma^2 \exp(-\tilde{\psi}_1 - \tilde{\psi}_2(t_k - T)) \right) \Delta t. \end{aligned} \quad (6.8)$$

Note that the simulator in terms of $(\mathbb{E}[\tilde{X}_{t_k}^{\tilde{\psi}}], \text{Var}(\tilde{X}_{t_k}^{\tilde{\psi}}))$ is deterministic. The aggregate reward is simulated according to $\hat{r}_{t_k} = 0$ and $\hat{r}_{t_K} = \mathbb{E}[\tilde{X}_T] - \lambda \text{Var}(\tilde{X}_T)$.

We first set the coefficients of the simulator to $T = 1, b = 0.25, \sigma = 0.5, \lambda = 1.5, \beta = 0$. Next, we set the known model parameters as: $\gamma = 0.5, \beta = 0$, the time step $\Delta t = 0.05$, the number of episodes $N = 2500$, the number of test policies $M = 10$, the lower bound of uniform distribution $p(j) = 0$, the upper bound of uniform distribution $q(j) = \frac{2}{j^{0.25}}$. We set the initialization of $\bar{\mu}_0^m \sim \mathcal{N}(0, 1)$ and $\text{Var}^m(\mu_0) \sim \mathcal{U}([0, 1])$ respectively, and choose the initialization of $\theta = (-0.5, 0.5, 0.5)^\top, \psi = (0.5, -0.5, 1.5, -0.5)^\top$. The learning rates $(\alpha_\theta, \alpha_\psi)$ are chosen by

$$\begin{aligned} \alpha_\theta(j) &= \left(\frac{0.015}{j^{0.22}}, \frac{0.01}{j^{0.1}}, \frac{0.025}{j^{0.11}} \right), \\ \alpha_\psi(j) &= \left(\frac{0.035}{j^{0.09}}, \frac{0.11}{j^{0.1}}, \frac{0.02}{j^{0.2}}, \frac{0.01}{j^{0.15}} \right). \end{aligned}$$

Based on the offline q-learning Algorithm 1, we plot in Figure 1 the numerical results on the convergence of parameters θ and ψ for the optimal value function and the optimal essential q-function q_e , and summarize the learnt parameters for the optimal value function and the optimal essential q-function q_e in Table 1.

6.2 Mean-Field Optimal Consumption Problem

We consider the mean-field R&D project process that is governed by the McKean-Vlasov SDE

$$dX_s = a_s b \mathbb{E}[X_s] ds + \sigma \mathbb{E}[X_s] dW_s - c_s ds, (a_s, c_s) \sim \pi(\cdot | s, X_s, \mathbb{P}_{X_s}), X_t = x > 0, \quad (6.9)$$

with $b > 0$ and $\sigma > 0$. Here, a_s is the drift control of the R&D investment and $c_s > 0$ is the consumption rate. Let us consider the problem of expected utility on consumption subjecting to the quadratic cost of control in the following form that

$$J(t, \mu; \{a_s\}_{s \geq t}, \{c_s\}_{s \geq t}) = \mathbb{E} \left[\int_t^T e^{-\beta(s-t)} U(c_s) ds - \int_t^T e^{-\beta(s-t)} a_s^2 ds \right]$$

with the logarithmic utility function $U(x) = \log x$. The exploratory control-consumption performance after the entropy regularizer is defined by

$$J(t, \mu; \pi) = \mathbb{E} \left[\int_t^T \int_{\mathbb{R} \times \mathbb{R}_+} \left\{ e^{-\beta(s-t)} \left(U(c) - e^{-\beta(s-t)} a^2 - \gamma \log \pi(a, c | s, X_s^\pi, \mathbb{P}_{X_s^\pi}) \right) \right\} \pi(a, c | s, X_s^\pi, \mathbb{P}_{X_s^\pi}) da dc ds \right].$$

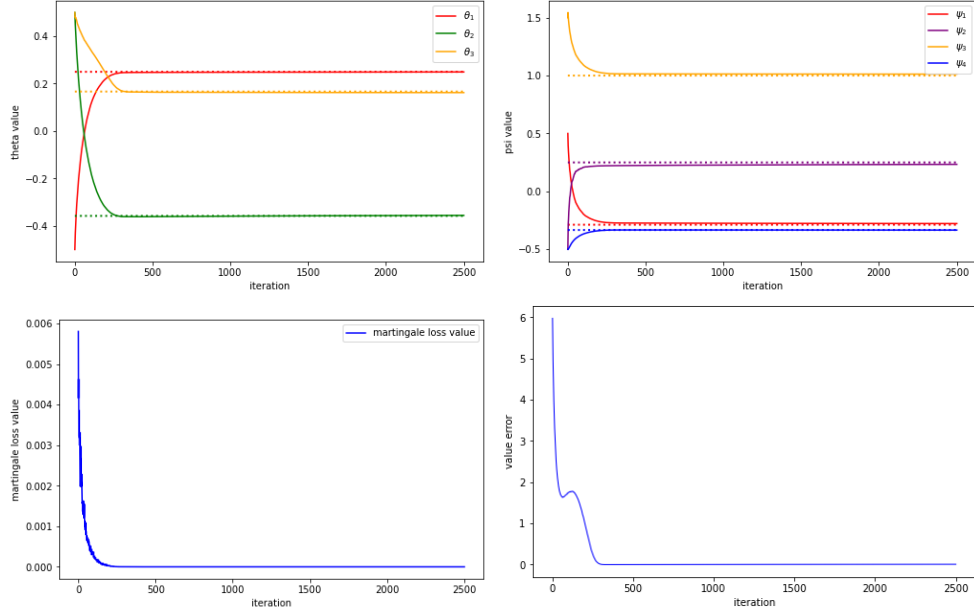


Figure 1: **Convergence of value function and q-function for Algorithm 1.** Paths of learnt parameters for value function (top left) and paths of learnt parameters for q-function (top right) vs optimal parameters shown in the dashed line. The change of the weak martingale loss value over iterations (bottom left) and the change of L^2 - error over iterations (bottom right) along a trajectory $(\bar{\mu}_{t_k}, \text{Var}(\mu)_{t_k})_k$ with $\bar{\mu}_0 = 0$ and $\text{Var}(\mu)_0 = 0.5$ controlled by the learnt policy and by the optimal policy.

The Hamiltonian operator is given by

$$H(t, x, \mu, a, c, p, q) = (b\bar{\mu}a - c)p + \frac{1}{2}\bar{\mu}^2\sigma^2q + \log(c) - a^2.$$

It then follows that

$$\begin{aligned} & \int_{\mathbb{R} \times \mathbb{R}_+} \exp\left(\frac{1}{\gamma}H(t, x, \mu, a, c, p, q)\right) dadc \\ &= \exp\left(\frac{1}{2\gamma}\bar{\mu}^2\sigma^2q\right) \int_{\mathbb{R}} \exp\left(\frac{1}{\gamma}abp\bar{\mu} - \frac{1}{\gamma}a^2\right) da \int_{\mathbb{R}_+} \exp\left(\frac{1}{\gamma}\log(c) - \frac{p}{\gamma}c\right) dc \\ &= \sqrt{\gamma\pi} \exp\left(\frac{1}{2\gamma}\bar{\mu}^2\sigma^2q\right) \exp\left(\frac{b^2p^2\bar{\mu}^2}{4\gamma}\right) \frac{\Gamma(1 + 1/\gamma)}{\left(\frac{p}{\gamma}\right)^{1+1/\gamma}} \end{aligned}$$

With the terminal condition $J^*(T, \mu) = 0$, we can write the exploratory HJB equation by

$$\begin{aligned} & \frac{\partial J^*}{\partial t}(t, \mu) - \beta J^*(t, \mu) + \gamma \int_{\mathbb{R}_+} \log \int_{\mathbb{R} \times \mathbb{R}_+} \exp\left(\frac{1}{\gamma}H(t, x, \mu, a, c, \partial_\mu J^*(t, \mu)(x), \partial_x \partial_\mu J^*(t, \mu)(x))\right) dadc \mu(dx) \\ &= \frac{\partial J^*}{\partial t}(t, \mu) - \beta J^*(t, \mu) + \frac{1}{2}\bar{\mu}^2\sigma^2 \int_{\mathbb{R}_+} \partial_x \partial_\mu J^*(t, \mu)(x) \mu(dx) + \frac{1}{4}b^2\bar{\mu}^2 \int_{\mathbb{R}_+} (\partial_\mu J^*(t, \mu)(x))^2 \mu(dx) \\ & \quad + \frac{\gamma}{2} \log(\gamma\pi) + \gamma \log \Gamma(1 + \frac{1}{\gamma}) - (1 + \gamma) \int_{\mathbb{R}_+} \log\left(\frac{\partial_\mu J^*(t, \mu)(x)}{\gamma}\right) \mu(dx). \end{aligned} \quad (6.10)$$

Suppose that $J^*(t, \mu)$ satisfies the following form

$$J^*(t, \mu) = B(t) \log \bar{\mu} + D(t),$$

with the terminal conditions $B(T) = D(T) = 0$. It holds that

$$\frac{\partial J^*(t, \mu)}{\partial t} = \dot{B}(t) \log \bar{\mu} + \dot{D}(t), \quad \partial_\mu J^*(t, \mu)(v) = B(t) \frac{1}{\bar{\mu}}, \quad \partial_v \partial_\mu J^*(t, \mu)(v) = 0.$$

Plugging these derivatives back into the exploratory HJB equation, we get that

$$\begin{aligned} & \left(\dot{B}(t) - \beta B(t) + 1 + \gamma \right) \log \bar{\mu} + \dot{D}(t) - \beta D(t) + \frac{1}{4} b^2 B^2(t) + \frac{\gamma}{2} \log(\gamma \pi) + \gamma \log \Gamma(1 + 1/\gamma) \\ & - (1 + \gamma) \log \left(\frac{B(t)}{\gamma} \right) = 0. \end{aligned}$$

Together with $B(T) = 0$, we first get that $B(t) = \frac{1+\gamma}{\beta} (1 - e^{-\beta(T-t)})$ and

$$\dot{D}(t) - \beta D(t) + \frac{1}{4} b^2 \frac{(1 + \gamma)^2}{\beta^2} (e^{-2\beta(T-t)} - 2e^{-\beta(T-t)} + 1) + K - (1 + \gamma) \log(1 - e^{-\beta(T-t)}) = 0, \quad (6.11)$$

where the constant $K := \frac{\gamma}{2} \log(\gamma \pi) + \gamma \log \Gamma(1 + 1/\gamma) - (1 + \gamma) \log \left(\frac{1+\gamma}{\beta \gamma} \right)$.

We can then obtain the explicit solution $D(t)$ of the ODE (6.11) with $D(T) = 0$ as

$$D(t) = A_1 e^{-2\beta(T-t)} + A_2 e^{-\beta(T-t)} + A_3 (1 - e^{-\beta(T-t)}) \log(1 - e^{-\beta(T-t)}) + A_4 t e^{-\beta(T-t)} + A_5,$$

where

$$\begin{aligned} A_1 &:= -\frac{b^2(1 + \gamma)^2}{4\beta^3}, \quad A_2 := (1 + \gamma)T - \frac{b^2(1 + \gamma)^2}{2\beta^2}T - \frac{K}{\beta}, \\ A_3 &:= -\frac{1 + \gamma}{\beta}, \quad A_4 := -(1 + \gamma) + \frac{b^2(1 + \gamma)^2}{2\beta^2}, \quad A_5 := \frac{b^2(1 + \gamma)^2}{4\beta^3} + \frac{K}{\beta}. \end{aligned}$$

As a result, we have the explicit essential q-function q_e^* as

$$\begin{aligned} q_e^*(t, x, \mu, a, c) &= \frac{\partial J^*(t, \mu)}{\partial t} - \beta J^*(t, \mu) + H(t, x, \mu, a, c, \partial_\mu J^*(t, \mu)(x), \partial_v \partial_\mu J^*(t, \mu)(x)) \\ &= -a^2 + \frac{1 + \gamma}{\beta} b(1 - e^{-\beta(T-t)})a - \frac{1 + \gamma}{\beta \bar{\mu}} (1 - e^{-\beta(T-t)})c + \log c - (1 + \gamma) \log \bar{\mu} \\ &\quad - \frac{b^2(1 + \gamma)^2}{4\beta^2} (e^{-\beta(T-t)} - 1)^2 - K + (1 + \gamma) \log(1 - e^{-\beta(T-t)}). \end{aligned}$$

Moreover, the separation form holds that $\pi^*(a, c|t, x, \mu) = \pi_1^*(a|t, x, \mu) \pi_2^*(c|t, x, \mu)$, where $\pi_1^*(a|t, x, \mu) := \mathcal{N}\left(\frac{b(1+\gamma)}{2\beta}(1 - e^{-\beta(T-t)}), \frac{\gamma}{2}\right)$ and $\pi_2^*(c|t, x, \mu) := \text{Gamma}(1 + \frac{1}{\gamma}, \frac{1+\gamma}{\gamma\beta\bar{\mu}}(1 - e^{-\beta(T-t)}))$.

Therefore, when the model parameters are unknown, we can parameterize the optimal value function J , the optimal essential q-function q_e and the optimal policy π^* respectively that

$$J^\theta(t, \mu) = \frac{1 + \gamma}{\beta} (1 - e^{-\beta(T-t)}) \log \bar{\mu} + \theta_1 e^{-2\beta(T-t)} + \theta_2 e^{-\beta(T-t)}$$

$$\begin{aligned}
& -\frac{1+\gamma}{\beta}(1-e^{-\beta(T-t)})\log(1-e^{-\beta(T-t)})+\theta_3te^{-\beta(T-t)}+\theta_4, \\
q_e^\psi(t,x,\mu,a,c) &= -(1+\gamma)\log\bar{\mu}+\psi_1(1-e^{-\beta(T-t)})a-a^2+\psi_2(e^{-\beta(T-t)}-1)^2 \\
& -\frac{1+\gamma}{\beta\bar{\mu}}(1-e^{-\beta(T-t)})c+\log c-K+(1+\gamma)\log(1-e^{-\beta(T-t)}) \\
\pi_1^\psi(a|t,x,\mu) &= \mathcal{N}\left(\frac{\psi_1}{2}(1-e^{-\beta(T-t)}),\frac{\gamma}{2}\right), \quad \pi_2^\psi(c|t,x,\mu) = \text{Gamma}\left(1+\frac{1}{\gamma},\frac{1+\gamma}{\gamma\beta\bar{\mu}}(1-e^{-\beta(T-t)})\right).
\end{aligned}$$

The parameters are θ_i , $i = 1, \dots, 4$ and ψ_j , $j = 1, 2$. We also require q_2^ψ to satisfy (5.1). In this case, $\psi_2 = -\frac{\psi_1^2}{4}$.

The following simulator will be used in Algorithms 1-2 to generate sample trajectories.

Simulator Under the stochastic policy $\mathbf{h}^{\tilde{\psi}} = \mathbf{h}_1^{\tilde{\psi}}\mathbf{h}_2^{\tilde{\psi}}$, the dynamics (6.9) becomes

$$\begin{aligned}
d\tilde{X}_s &= b\mathbb{E}[\tilde{X}_s] \int_{\mathbb{R}} a\mathbf{h}_1^{\tilde{\psi}}(a|s,\tilde{X}_s,\mathbb{P}_{\tilde{X}_s})dad s - \int_{\mathbb{R}_+} c\mathbf{h}_2^{\tilde{\psi}}(c|s,\tilde{X}_s,\mathbb{P}_{\tilde{X}_s})dc ds + \sigma\mathbb{E}[\tilde{X}_s]dW_s \\
&= \frac{\tilde{\psi}_1 b}{2}(1-e^{-\beta(T-s)})\mathbb{E}[\tilde{X}_s]ds - \frac{\beta}{1-e^{-\beta(T-s)}}\mathbb{E}[\tilde{X}_s]ds + \sigma\mathbb{E}[\tilde{X}_s]dW_s.
\end{aligned}$$

By taking the expectation on both sides of the above SDE, we obtain that

$$d\mathbb{E}[\tilde{X}_s] = \frac{\tilde{\psi}_1 b}{2}(1-e^{-\beta(T-s)})\mathbb{E}[\tilde{X}_s]ds - \frac{\beta}{1-e^{-\beta(T-s)}}\mathbb{E}[\tilde{X}_s]ds.$$

The Euler approximation of $\bar{\mu}_t = d\mathbb{E}[\tilde{X}_s]$ is

$$\log \bar{\mu}_{t_{k+1}} = \log \bar{\mu}_{t_k} + \left(\frac{\tilde{\psi}_1 b}{2}(1-e^{-\beta(T-s)}) - \frac{\beta}{1-e^{-\beta(T-s)}} \right) \Delta t. \quad (6.12)$$

The simulated reward is then given by

$$\begin{aligned}
\hat{r}_{t_k} &= \int_{\mathbb{R}} \int_{\mathbb{R}_+} \log c\mathbf{h}_2^{\tilde{\psi}}(c|t_k,x,\mathbb{P}_{\tilde{X}_{t_k}})dc\mathbb{P}_{\tilde{X}_{t_k}}(dx) - \int_{\mathbb{R}} \int_{\mathbb{R}} a^2\mathbf{h}_1^{\tilde{\psi}}(a|t_k,x,\mathbb{P}_{\tilde{X}_{t_k}})da\mathbb{P}_{\tilde{X}_{t_k}}(dx) \\
&= \frac{\Gamma'(1+\frac{1}{\gamma})}{\Gamma(1+\frac{1}{\gamma})} - \log\left(\frac{1+\gamma}{\gamma\beta}(1-e^{-\beta(T-t)})\right) + \log\bar{\mu} + \frac{\tilde{\psi}_1^2}{4}(1-e^{-\beta(T-t)})^2 + \frac{\gamma^2}{4}.
\end{aligned}$$

We first set coefficients of the simulator to $T = 1, b = 0.5, \sigma = 0.5$. We next set the known parameters as: $\gamma = 0.25, \beta = 2$, the time step $\Delta t = 0.1$, the number of episodes $N = 8000$, the number of test policies $M = 10$, the lower bound of uniform distribution $p(j) = 0$, the upper bound of uniform distribution $q(j) = \frac{2}{j^{0.6}}$, the log mean $\log \mu_0^m, 1 \leq m \leq M$ that is initialed to be 0, and the learning rates $(\alpha_\theta, \alpha_\psi)$ are chosen by

$$\alpha_\theta(j) = \begin{cases} (\frac{0.05}{j^{0.22}}, \frac{0.2}{j^{0.15}}, \frac{0.2}{j^{0.25}}, 0.5), & \text{if } j \leq 4500, \\ (\frac{0.05}{j^{0.4}}, \frac{0.2}{j^{0.15}}, \frac{0.2}{j^{0.55}}, \frac{0.5}{j^{0.01}}), & \text{if } 4500 < j \leq 8000, \end{cases}$$

and

$$\alpha_\psi(j) = \begin{cases} \frac{0.15}{j^{0.31}}, & \text{if } j \leq 4500, \\ \frac{0.15}{j^{0.61}}, & \text{if } 4500 < j \leq 8000. \end{cases}$$

Table 2:

	θ_1	θ_2	θ_3	θ_4	ψ_1
true value	-9.765×10^{-5}	1.085	-1.248	0.163	0.0625
learnt by Algorithm 1	0.0796	1.107	-1.281	0.119	0.0667

Based on the offline q-learning Algorithm 1, we plot in Figure 2 the numerical results on the convergence of parameters θ and ψ for the value function and q-function, and also summarize the learnt parameters for the optimal value function and the optimal essential q-function q_e in Table 2.

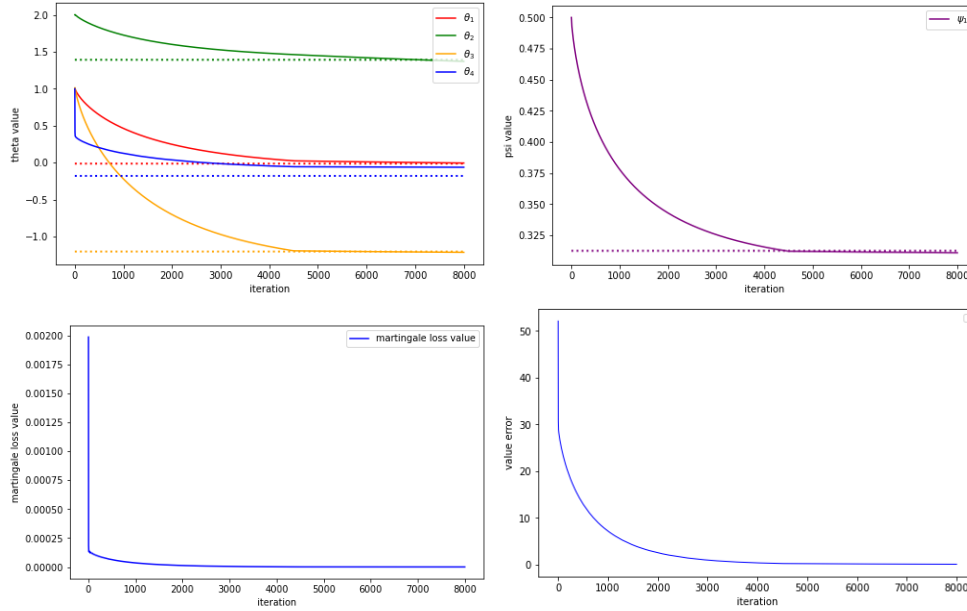


Figure 2: **Convergence of value function and q-function for Algorithm 1.** Paths of learnt parameters for value function (top left) and path of learnt parameters for q-function (top right) vs optimal parameters shown in the dashed line. The change of the weak martingale loss value over iterations (bottom left) and the change of L^2 -error over iterations (bottom right) along a trajectory $(\log(\bar{\mu}_{t_k}))_k$ with $\log(\bar{\mu}_0) = 0$ controlled by the learnt policy and by the optimal policy.

7 Conclusion

This paper aims to lay the theoretical foundation of the continuous time q-learning for mean-field control problems, which can be viewed as the bridge between the discrete time Q-learning with the integrated Q-function for MFC problems studied in Gu et al. (2023) and the continuous time q-learning with the q-function for single agent’s control problems studied in Jia and Zhou (2023). In our framework, two different q-functions are introduced for the purpose of learning, namely the integrated q-function q that is defined as the first order derivative of the integrated Q-function with respect to time and the essential q-function q_e that is used for policy improvement iterations. Comparing with two counterparts in Gu et al. (2023) and Jia and Zhou (2023), we establish the weak martingale characterization of the value function and the integrated q-function through test policies in the neighbourhood of the target policy in a similar flavor of

off-policy learning. To learn the essential q-function from the double integral representation, we propose the average-based loss function by searching test policies in the same form of the parameterized target policy but with randomized parameters. In two examples, one in the LQ control framework and one beyond the LQ control framework, we can illustrate the effectiveness of our q-learning algorithm.

Several interesting future extensions can be considered. Firstly, we may consider the controlled common noise in the mean-field control problem, where the optimal policy no longer admits the explicit form as a Gibbs measure. Instead, we can provide the first order condition of the optimal policy using the linear functional derivative with respect to probability measures. It is an interesting open problem to investigate the correct form of the continuous time integrated q-function and the associated q-learning. Secondly, it will also be appealing to investigate the decentralized continuous time q-learning from the representative agent's perspective and provide the correct q-function and policy iteration rules based on observations of the representative agent's individual state dynamics. At last, we are also interested in establishing some theoretical convergence results on the policy improvement iterations and continuous time q-learning in the mean field model.

Acknowledgement: We sincerely thank two referees for their constructive comments and suggestions. X. Wei is supported by National Natural Science Foundation of China grant under no.12201343. X. Yu is supported by the Hong Kong Polytechnic University research grant under no. P0045654.

Declarations

Competing interests The authors have no competing interests to declare that are relevant to the content of this article.

A Connection between formulations (2.4) and (2.6)

Lemma A.1 X_s^π and \tilde{X}_s^π , as solutions to (2.4) and (2.6) respectively, have the same distribution for each $s \in [t, T]$ as they correspond to the same martingale problem.

Proof. Let $\mathcal{C}_c^\infty(\mathbb{R}^d)$ denote the set of infinitely differentiable function $\varphi : \mathbb{R}^d \rightarrow \mathbb{R}$ with compact set, and let $\nabla\phi$ and $\nabla^2\phi$ denote the gradient and Hessian of φ , respectively. We define the infinitesimal generator

$$L_s^a\varphi := b(s, x, \mu, a)^\top \nabla\varphi(x) + \frac{1}{2} \text{Tr}(\sigma\sigma^\top(s, x, \mu, a)\nabla^2\varphi). \quad (\text{A.1})$$

Given $(t, \mu) \in [0, T] \times \mathcal{P}_2(\mathbb{R}^d)$ and fixed $\pi \in \Pi$, a solution to the martingale problem for the generator L^a in (A.1) is a probability measure m such that

$$\varphi(X_s) - \int_t^s \int_{\mathcal{A}} L_u^a\varphi(u, X_u, m_{X_u}, a)\pi(a|u, X_u, m_{X_u})dadu, \quad s \geq t \quad (\text{A.2})$$

is a m -martingale for all $\varphi \in \mathcal{C}_c^\infty(\mathbb{R}^d)$. The martingale formulation (A.2) was studied in Karoui and Méléard (1990) and then generalized to MFC theory in Lacker (2017) for the limit theory of N -player dynamics to the controlled McKean-Vlasov dynamics.

Note that when \tilde{X}_s^π satisfies (2.6), Itô's formula implies that \tilde{X}^π satisfies (A.2). Furthermore, we can also apply Itô's formula to X_s^π satisfying (2.4) and take the conditional expectation given \mathcal{F}_r , $t \leq r \leq s$,

$$\begin{aligned} \mathbb{E}[\varphi(X_s^\pi)|\mathcal{F}_r] - \varphi(\xi) &= \mathbb{E}\left[\int_t^s L_u^{a^\pi} \varphi(u, X_u^\pi, \mathbb{P}_{X_u}^\pi) du | \mathcal{F}_r\right] \\ &= \mathbb{E}\left[\int_t^s L_u^a \varphi(u, X_u^\pi, \mathbb{P}_{X_u}^\pi) \pi(a|u, X_u, \mathbb{P}_{X_u}) du | \mathcal{F}_r\right], \end{aligned}$$

where in the last equality we have used the fact the action a^π is sampled from π independent of W , and hence $\mathbb{E}[\varphi(X_u, \mathbb{P}_{X_u}, a_u) | \mathcal{F}_r] = \mathbb{E}[\int_{\mathcal{A}} \varphi(X_u, \mathbb{P}_{X_u}, a) \pi(a|u, X_u, \mathbb{P}_{X_u}) da | \mathcal{F}_r]$ for all φ . Therefore, X_s^π also satisfies (A.2). It follows from the uniqueness of the martingale problem that \tilde{X}_s^π and X_s^π are the same in law.

References

- A. Angiuli, J. P. Fouque and M. Laurière. (2022). Unified reinforcement Q-learning for mean field game and control problems. *Mathematics of Control, Signals, and Systems*, 34(2), 217-271.
- A. Angiuli, J.P. Fouque, R. Hu and A.Raydan (2023). Deep reinforcement learning for infinite horizon mean field problems in continuous spaces. Preprint, available at arXiv:2309.10953.
- A. Angiuli, J.P. Fouque, M. Laurière and M. Zhang (2023). Convergence of multi-scale reinforcement Q-learning algorithms for mean field game and control problems. Preprint, available at arXiv:2312.06659.
- R. Carmona and F. Delarue (2018a): Probabilistic Theory of Mean Field Games with Applications, Vol I. Springer.
- R. Carmona and F. Delarue (2018b): Probabilistic Theory of Mean Field Games with Applications, Vol II. Springer.
- R. Carmona, J. P. Fouque and L. H. Sun (2015): Mean field games and systemic risk. *Communications in Mathematical Sciences*, 13(4):911-933.
- R. Carmona, M. Laurière. and Z. Tan. (2023): Model-free mean-field reinforcement learning: mean-field MDP and mean-field Q-learning. *Annals of Applied Probability*, 33(6B), 5334-5381.
- A. Cosso, F Gozzi, I. Kharroubi, H. Pham and M. Rosestolato (2023): Optimal control of path-dependent McKean-Vlasov SDEs in infinite dimension. *The Annals of Applied Probability*, 33(4):2863-918.
- K. Doya (2020). Reinforcement learning in continuous time and space. *Neural Computation*, 12(1):219–245.
- M. Giegrich, C. Reisinger and Y. Zhang (2024): Convergence of policy gradient methods for finite-horizon exploratory linear-quadratic control problems. *SIAM Journal on Control and Optimization*, 62(2):1060-92.
- H. Gu, X. Guo, X. Wei and R. Xu (2021): Mean-field controls with Q-learning for cooperative MARL: convergence and complexity analysis. *SIAM Journal on Mathematics of Data Science*, 3(4):1168-96.
- H. Gu, X. Guo, X. Wei, R. Xu (2021): Mean-field multi-agent reinforcement learning: A decentralized network approach. Forthcoming in *Mathematics of Operations Research*.
- H. Gu, X. Guo, X. Wei and R. Xu (2023): Dynamic programming principles for mean-field controls with learning. *Operations Research*, 71(4), 1040-1054.

- X. Guo, R. Xu and T. Zariphopoulou (2022): Entropy regularization for mean field games with learning. *Mathematics of Operations Research*, 47(4), 3239-3260.
- X. Guo, A. Hu and Y. Zhang (2023): Reinforcement learning for linear-convex models with jumps via stability analysis of feedback controls *SIAM Journal on Control and Optimization*, 61(2), 755-787.
- N. Frikha, M. Germain, M. Laurière, H. Pham. and X. Song (2023). Actor-Critic learning for mean-field control in continuous time. Preprint, available at arXiv:2303.06993.
- Y. Jia and X. Y. Zhou (2022a): Policy gradient and actor-critic learning in continuous time and space: Theory and algorithms. *Journal of Machine Learning Research*. 23, 1-50.
- Y. Jia and X. Y. Zhou (2022b): Policy evaluation and temporal-difference learning in continuous time and space: A martingale approach. *Journal of Machine Learning Research*. 23, 1-55.
- Y. Jia and X. Y. Zhou (2023): q-learning in continuous time. *Journal of Machine Learning Research*. 24, 1-61.
- N. E. Karoui and S. Méléard (1990): Martingale measures and stochastic calculus. *Probability theory and related fields*. 84(1), 83-101.
- D. Lacker (2017): Limit theory for controlled McKean-Vlasov dynamics. *SIAM Journal on Control and Optimization*, 55(3):1641-1672.
- N. Li, X. Li and Z. Xu (2023): Policy iteration reinforcement learning method for continuous time mean field linear-quadratic optimal problem. Preprint, available at arXiv:2305.00424.
- P.-L. Lions (2006): Cours au collège de france: Théorie des jeux à champ moyens. Audio Conference.
- W.U. Mondal, M. Agarwal, V. Aggarwal and S.V. Ukkusuri (2022): On the approximation of cooperative heterogeneous multi-agent reinforcement learning (MARL) using mean field control (MFC). *Journal of Machine Learning Research*, 23(129), 1-46.
- W.U. Mondal, V. Aggarwal and S. V. Ukkusuri (2023): Mean-field control based approximation of multi-agent reinforcement learning in presence of a non-decomposable shared global state. Preprint, available at arXiv:2301.06889.
- B. Pasztor, I. Bogunovic and A. Krause(2021): Efficient model-based multi-agent mean-field reinforcement learning. Preprint, available at arXiv:2107.04050.
- L. Szpruch, T. Treetanthiploet and Y. Zhang. (2021): Exploration-exploitation trade-off for continuous-time episodic reinforcement learning with linear-convex models. *To appear in Annals of Applied Probability*.
- J. Schulman, X. Chen and P. Abbeel (2017): Equivalence between policy gradients and soft Q-learning. Preprint arXiv:1704.06440.
- Y. Sun (2006): The exact law of large numbers via Fubini extension and characterization of insurable risks. *Journal of Economic Theory*, 126(1), pp.31-69.
- R. S. Sutton and A. G. Barto (2018): Reinforcement learning: An introduction. MIT press.
- D. Stroock and S. Varadhan (1997): Multidimensional diffusion processes. volume 233 of *Grundlehren der mathematischen Wissenschaften*. Springer-Verlag Berlin Heidelberg.
- H. Wang, T. Zariphopoulou and X. Y. Zhou (2020): Reinforcement learning in continuous time and space: A stochastic control approach. *Journal of Machine Learning Research*, 21(1):8145-8178.
- X. Wei, X. Yu and F. Yuan (2024): Unified continuous-time q-learning for mean-field game and mean-field control problems. Preprint, available at arXiv:2407.04521.