
Finding Regularized Competitive Equilibria of Heterogeneous Agent Macroeconomic Models via Reinforcement Learning

Ruitu Xu
Yale University

Yifei Min
Yale University

Tianhao Wang
Yale University

Zhaoran Wang
Northwestern University

Michael I. Jordan
University of California, Berkeley

Zhuoran Yang
Yale University

Abstract

We study a heterogeneous agent macroeconomic model with an infinite number of households and firms competing in a labor market. Each household earns income and engages in consumption at each time step while aiming to maximize a concave utility subject to the underlying market conditions. The households aim to find the optimal saving strategy that maximizes their discounted cumulative utility given the market condition, while the firms determine the market conditions through maximizing corporate profit based on the household population behavior. The model captures a wide range of applications in macroeconomic studies, and we propose a data-driven reinforcement learning framework that finds the regularized competitive equilibrium of the model. The proposed algorithm enjoys theoretical guarantees in converging to the equilibrium of the market at a sub-linear rate.

1 INTRODUCTION

The behavior of labor markets has always been one of the key subjects of study in macroeconomics, and it is crucial to understand the underlying mechanisms that give rise to the aggregate macroeconomic indicators, such as price level and unemployment rate. In addition, detailed economic statistics such as wealth distribution have also become critical factors in analyzing the societal impact of economic policies. There is therefore a major need for a deeper understanding of the interaction between economic variables in large and complex markets, a challenging endeavor for economists that

machine learning and ubiquitous data collection promise to catalyze.

A variety of macroeconomic models have been proposed to characterize high-dimensional economic dependencies (Aiyagari, 1994; Bewley, 1986; Huggett, 1993). Under these models, the competition in the market among different types of market participants, such as households and firms, is often regarded as a sequential game, where the participants interact and eventually arrive at certain equilibrium strategies (Kuhn, 2013). Being able to find these equilibrium outcomes is critical in unraveling many important aspects of economic growth and providing insightful guidance to policymakers and corporations.

In this paper, we focus on a labor market model that is composed of two groups of participants: households and firms, each with an infinite number of heterogeneous agents. One of the classic examples of such heterogeneous agent model is Aiyagari model (Aiyagari, 1994). At every time step, each household retains certain capital holdings, *e.g.*, cash and investments, and earns a certain amount of income, *e.g.*, salary and dividends. The income of each household is subject to independent and heterogeneous exogenous shocks over time, which reflects the impact of market conditions on individual household employment. Facing the capital holdings and income at each step, the household saves part of its total wealth for the next step and spends the rest to enjoy some utility.

The competition between the households and firms lies at the essence of the model: The goal of the household is to come up with a saving strategy that maximizes its discounted cumulative utility. The utility function is concave with respect to savings, a critical assumption that reflects the economic principle of diminishing returns (Shephard and Färe, 1974). The firms make decisions that maximize corporate profit given the population behavior of the households subject to the saving strategy, and these decisions then determine the market conditions. Such competition between the two parties is characteristic of modern labor markets,

which eventually arrive at an equilibrium determined by the inherent properties of the markets. Note that the model comprises infinitely many households interacting through a population variable, and it can be cast into a mean-field game (MFG) (Lasry and Lions, 2007), such that the representative agent of the MFG stands for a typical household and the mean-field term represents the market condition that incorporates both the population behavior of the households and the optimal decisions of a typical firm (Light and Weintraub, 2022).

Existing numerical methods in economic literature suffer from a number of drawbacks in solving for such equilibria, and these drawbacks impose significant limits on the applications of the model to real markets (Kuhn, 2013; Achdou et al., 2014b, 2022, 2014a). Specifically: 1) The underlying stochastic model for income shocks is unknown in general, while existing methods assume prior knowledge of income shock transitions, which is unrealistic in practice. 2) These methods also require a discretization of the state-action space in computing the utilities despite the model featuring a continuous variable space. It results in massive computational inefficiency due to the curse of dimensionality (Bellman and Dreyfus, 2015). 3) There are few theoretical results that guarantee the convergence of these existing methods to the desired equilibrium.

In this work we present a reinforcement learning (RL) based alternative to the existing methods. Machine learning techniques have been adopted for solving economic models recently, thanks to the emergence of massive computational power and the increasingly available microeconomic data (Achdou et al., 2022; Curry et al., 2022; Min et al., 2022b). Despite some attempts to apply data-driven RL methods for macroeconomic models (Curry et al., 2022), theoretical understanding of such algorithms is still limited. In particular, the analysis of RL algorithms for heterogeneous agent models face several unique challenges: 1) The concavity assumption on the utility function of the households features the economic nature of the problem, and it induces a concave shape constraint on the value functions of the agents for any feasible policy. Such shape constraint is absent in the existing literature on MFG, and it requires special treatments to achieve better sample efficiency. 2) The model induces a continuous domain of the value functions, whereas existing literature on MFG only considers a discrete state-action space. Under such a continuous setup of state-action space, the optimal policy of an agent may not be unique, and this may compromise the uniqueness of the equilibrium as well as the stability of the learning process.

Our Contributions. We summarize the main contributions of this paper as follows:

- We propose a RL framework for a class of heterogeneous agent models in macroeconomics, by formulating it as an MFG with a shape constraint on the

value function. Our formulation generalizes the well-known Aiyagari model, extending the model into high-dimensional state-action space.

- To guarantee a provable convergence to an equilibrium, we propose a regularized policy iteration algorithm that we refer to as `ConcaveHAM`, facilitated by the combination of fitted Q -iteration (Riedmiller, 2005) and convex regression. Our algorithm provides a data-driven approach that estimates the income shock transitions, and with the introduction of convex regression, it avoids the curse of dimensionality in discretization and provides a continuous solution on value functions that better characterize the problem. It is noteworthy that our algorithm can readily incorporate other forms of shape constraints under large state spaces.
- We prove that `ConcaveHAM` converges to the desired equilibrium at a sublinear rate under standard assumptions. Especially, it learns a quantal response mean-field equilibrium (QR-MFE) of the MFG, an equivalence of the regularized recursive competitive equilibrium of the macroeconomic model (Prescott and Mehra, 2005), which incorporates the concept of bounded rationality (Selten, 1990) and enforces the uniqueness of optimal policy through entropy regularization. To the best of our knowledge, this is the first data-driven framework for heterogeneous agent models that captures economic intuitions and guarantees theoretical tractability simultaneously.
- In the MFG that we consider, the feasible action set is not independent of its current state, which is a natural consequence of the household budget constraint. This is also a unique characteristic that has not been considered in the existing literature to the best of our knowledge and might be of independent interest.

Notation. We denote the set of all positive integers by \mathbb{Z}_+ and the set of all non-negative real numbers by \mathbb{R}_+ . For functions f and g , we denote $f(n) \lesssim g(n)$ if $f(n) \leq Cg(n)$ for every $n \in \mathbb{Z}_+$ with some universal constant $C > 0$ (similarly, $f(n) \gtrsim g(n)$); further, we write $f(n) \asymp g(n)$ if $f(n) \lesssim g(n)$ and $f(n) \gtrsim g(n)$. We write $\text{diam}(\mathcal{X}) := \max_{x, y \in \mathcal{X}} \|x - y\|_\infty$ for any set \mathcal{X} . For any measurable set A , we define Δ_A as the set of all density functions supported on A . We denote the set of policies π with Π . For any measurable functions f and g supported on domain A , we denote $\langle f, g \rangle_A = \int_A f(x)g(x) dx$.

2 PRELIMINARY

In this paper, we focus on a heterogeneous agent model as a generalization of the well-known Aiyagari model of labor markets with idiosyncratic income shocks (Aiyagari, 1994). Notably, we formulate the model as an MFG, *i.e.*, an

asymptotic approximation of a multi-agent Markov game with a large number of agents. In the rest of this section, we elaborate on how to cast a generalized Aiyagari model in macroeconomics into the form of an MFG. Table 1 shows a detailed correspondence between the two interpretations.

2.1 Heterogeneous Agent Model

Consider $d_b, d_w, d_z \in \mathbb{Z}_+$. Let $\mathcal{B} \subseteq [0, 1]^{d_b}$ be the set of all capital holdings, $\mathcal{W} \subseteq [0, 1]^{d_w}$ be the set of all possible incomes, and $\mathcal{A} \subseteq [0, 1]^{d_a}$ be the set of all feasible capital savings. Further let $\mathcal{Z} \subseteq \mathbb{R}^{d_z}$ be the set of all aggregate indicators that represent possible market conditions.

Households. Within a heterogeneous agent model, each household on the market is characterized by its total assets $b \in \mathcal{B}$ and current income $w \in \mathcal{W}$ at each time step. Under any given market condition represented by some aggregate indicators $z \in \mathcal{Z}$, the income w of each household is subject to independent heterogeneous exogenous shocks, *i.e.*, the income of each household forms a Markov chain as time evolves. The household adopts a saving strategy $\pi : \mathcal{B} \times \mathcal{W} \rightarrow \mathcal{A}$, and at each time step, after making saving $a = \pi(b, w)$, the household gains utility $r_z(b, w, a)$, where $r_z : \mathcal{B} \times \mathcal{W} \times \mathcal{A} \rightarrow \mathbb{R}$ is the utility function under market condition z . The savings a made at the current time step then become the capital holdings b at the next time step.

The goal of each household is to come up with a saving strategy π_z^* that maximizes its expected cumulative utility $V_z^{\pi_z^*}$ (as a γ -discounted sum of all future rewards) subject to the underlying market condition z . The optimal saving strategy of the households, together with the aggregate indicator, gives rise to a population average of household capital retention $b_z^{\pi_z^*}$ and income level $w_z^{\pi_z^*}$.

Firms. Based on $b_z^{\pi_z^*}$ and $w_z^{\pi_z^*}$, the firms desire to make decisions that maximize their profits, which in turn gives a new market condition indicator z' following some aggregate function Ψ and production mapping Φ .

More specifically, for $d_\Xi \in \mathbb{Z}_+$ and the set Π of all feasible policies, we assume the representative firm has access to an aggregate function $\Psi : \mathcal{Z} \times \Pi \rightarrow \mathbb{R}^{d_\Xi}$ that maps any mean-field term z and policy π to a set of aggregate indicators $\Xi_z^\pi = (b_z^\pi, w_z^\pi) \in \mathbb{R}^{d_\Xi}$ (which represents the population behaviors of all households in the market). Here b_z^π denotes the aggregated indicators on capital retention and w_z^π denotes those on labor supply. The representative firm then makes corporate decisions based on the indicators Ξ_z^π . In particular, the firm takes Ξ_z^π as given and picks the corporate decisions that maximize the production function of the firm, which give rise to a new market condition represented by a new mean-field term $z' = \Phi(\Xi_z^\pi) \in \mathcal{Z}$.

Competition Between Households and Firms. Under the new market condition z' , the households then need to update their saving policy, and the competition between the

households and the firms continues iteratively.

Notice that the households and firms interact only through the aggregate terms b_z^π , w_z^π , and z ; therefore, the competition can be regarded as between one representative household and one representative firm. Such repetitive competitions reach a competitive equilibrium $(z^*, \pi_{z^*}^*)$ under mild conditions. Our goal is to learn an approximate equilibrium of $(z^*, \pi_{z^*}^*)$ from observational data, where the behavior policy may not be optimal with respect to the underlying market condition z .

2.2 Mean-Field Game on Households

We cast the above heterogeneous agent model into a framework of MFG, where infinitely many identical agents interact through a mean-field term (Lasry and Lions, 2007). In particular, we focus on a representative agent that stands for a representative household and consider its interaction with a mean-field term z , which forms a counterpart of the aggregate indicators z of the market. From now on, we describe the components of the MFG in the language of RL, as listed in Table 1.

Shape-Constrained MDP for Representative Agent. Fix any mean-field term z . The interaction between the agent and economic environment forms a discounted infinite-horizon MDP with a shape constraint, denoted by a tuple $(\mathcal{S}, \mathcal{A}, \mathcal{G}, \Gamma, \mathbb{P}_z, r_z, \gamma)$, where $\gamma \in (0, 1)$ is the discount factor. Here $\mathcal{S} := \mathcal{B} \times \mathcal{W}$ is the state space of total assets and income, and \mathcal{A} is the action space of savings. In contrast to standard definitions of MDPs, here at each state $s \in \mathcal{S}$, the agent can only take action in a feasible actions set $\Gamma(s) \subseteq \mathcal{A}$ that reflects the household budget constraint. We define the feasible state-action set as $\mathcal{G} := \{(s, a) \in \mathcal{S} \times \mathcal{A} \mid a \in \Gamma(s)\}$. In particular, we assume that $\mathcal{G} \subseteq \mathbb{R}^{d_s + d_w}$ is a convex set, corresponding to a convex household budget constraint, and that it is independent of the choice of mean-field term z . We also assume the reward function $r_z : \mathcal{G} \rightarrow [0, 1]$ is a concave function that depends on the mean-field term z .

At each time step h given state $s_h = (b_h, w_h)$, the agent takes some action $a_h \in \Gamma(s_h)$ and receives a reward $r_z(s_h, a_h)$. Then the agent transitions to the next state $s_{h+1} = (b_{h+1}, w_{h+1}) \sim \mathbb{P}_z(\cdot \mid s_h, a_h)$, where the transition probability \mathbb{P}_z represents the idiosyncratic income shock that implicitly depends on the mean-field term z . More specifically, for any $a_h \in \Gamma(b_h, w_h)$.

$$\mathbb{P}_z(b_{h+1}, w_{h+1} \mid b_h, w_h, a_h) = \begin{cases} \mathbb{P}_z[w_{h+1} \mid w_h] & b_{h+1} = a_h \\ 0 & b_{h+1} \neq a_h \end{cases} \quad (2.1)$$

where $\mathbb{P}_z[\cdot \mid \cdot]$ denotes the income shock.

Learning Goal. The representative agent aims to maximize the γ -discounted cumulative reward over an infinite time horizon. We encode the agent's strategy in a map $\pi : \mathcal{S} \rightarrow$

Table 1: A Comparison of Notation Between the Generalized Aiyagari Model and Its MFG Formulation.

Mean-field game (Representative agent)		Heterogeneous agent model (Representative household)	
Notation	Interpretation	Notation	Interpretation
$s \in \mathcal{S} \subseteq \mathbb{R}^{d_s}$	State	$(b, w) \in \mathcal{B} \times \mathcal{W} \subseteq \mathbb{R}^{d_s}$	Capital holdings and income
$a \in \mathcal{A} \subseteq \mathbb{R}^{d_w}$	Action	$a \in \mathcal{A} \subseteq \mathbb{R}^{d_w}$	Capital savings
$\pi : \mathcal{S} \rightarrow \Delta_{\mathcal{A}}$	Policy	$\pi : \mathcal{S} \rightarrow \Delta_{\mathcal{A}}$	Saving strategy
$z \in \mathcal{Z} \subseteq \mathbb{R}^{d_z}$	Mean-field term	$z \in \mathcal{Z} \subseteq \mathbb{R}^{d_z}$	Market condition (indicators)
r_z	Reward function	r_z	Utility function
V_z^π	Value function	V_z^π	Cumulative utility
Ξ_z^π	Aggregate indicator	(b_z^π, w_z^π)	Population capital and labor
(z^*, π^*)	QR-MFE	(z^*, π^*)	Market equilibrium

$\Delta_{\mathcal{A}}$, which is called a *policy*. In particular, $\pi(s)$ is only supported on $\Gamma(s)$ for any $s \in \mathcal{S}$. Then given any mean-field term z and policy π , we define the Q-function (*i.e.*, action-value function) $Q_z^\pi : \mathcal{S} \rightarrow \mathbb{R}$ as the expected discounted cumulative reward under the policy π , *i.e.*,

$$Q_z^\pi(s, a) := \mathbb{E} \left[\sum_{h=1}^{\infty} \gamma^h r_z(s_h, a_h) \mid s_0 = s, a_0 = a \right],$$

where the expectation \mathbb{E} is taken with respect to both \mathbb{P}_z and π along the trajectory. Similarly, we define the value function $V_z^\pi(s) : \mathcal{S} \rightarrow \mathbb{R}$ of π as

$$V_z^\pi(s) := \mathbb{E} \left[\sum_{h=0}^{\infty} \gamma^h r_z(s_h, a_h) \mid s_0 = s \right],$$

which we also write as $V_z^\pi(s) = \langle Q_z^\pi(s, \cdot), \pi(\cdot \mid s) \rangle_{\Gamma(s)}$. The goal of the representative agent is then to find an optimal policy π_z^* that maximizes its value functions $Q_z^*(s, a) := \max_{\pi} Q_z^\pi(s, a)$ and $V_z^*(s) := \max_{\pi} V_z^\pi(s)$. Note that such optimal policy always exists for infinite horizon Markov decision processes (MDPs) (Puterman, 2014), and it holds that $V_z^*(s) = \langle Q_z^*(s, \cdot), \pi_z^*(\cdot \mid s) \rangle_{\Gamma(s)}$.

Note that the convexity of the reward functions further induces the convexity of the value functions for any policy, and thus the policies are economically meaningful for the heterogeneous agent model.

Mean-Field Term: the Representative Firm. In our formulation of the MFG, we treat the representative firm as a function that outputs the mean-field term given any policy π from the representative agent.

Formally, we assume a mapping $\Phi \circ \Psi : \mathcal{Z} \times \Pi \rightarrow \mathcal{Z}$ that sends any policy π under a current mean-field term z to its corresponding updated mean-field term z' . We assume full knowledge of both functions, as the aggregation function Ψ can be estimated with a simulator and the production mapping Φ is generally a deterministic function given the production function.

2.3 Quantal Response Mean-Field Equilibrium

The MFG described above admits at least one equilibrium under mild conditions, and any equilibrium (z^*, π^*) of the MFG then corresponds to an equilibrium of the heterogeneous agent model, where the equilibrium reveals itself as the pair of market condition z^* and the optimal household strategy $\pi_{z^*}^*$ under such market condition.

Issues With RCE. One commonly considered equilibrium for MFG is recursive competitive equilibrium (RCE), which is represented as a tuple (z^*, Q^*, π^*, Ξ^*) where π^* is the optimal policy with respect to the MDP under transition kernel \mathbb{P}_{z^*} , the optimal Q-function $Q^* = Q_{z^*}^*$, and it holds that $\Phi(\Xi^*) = z^*$ for the aggregate indicators $\Xi^* = \Psi(z^*, \pi^*)$ (Light and Weintraub, 2022). However, such equilibrium can be non-unique and unstable with respect to Q^* and its estimation. Furthermore, policy estimation is not robust under the definition of RCE. The lack of robustness limits the capability of value-based RL algorithms from converging to the equilibria. Further discussions on the lack of uniqueness and robustness of RCE is delayed to Appendix A.2.

To overcome this challenge, we propose a regularized competitive equilibrium named Quantal Response Mean-Field Equilibrium (QR-MFE).

Definition 2.1 (Quantal Response Mean-Field Equilibrium). A representative agent policy and a mean-field term (π^*, z^*) reach quantal response mean-field equilibrium with respect to a strongly convex entropy function \mathcal{H} if $z^* = \Phi(\Psi(z^*, \pi^*))$ and for any policy π , state $s \in \mathcal{S}$, and the optimal Q-function $Q_{z^*}^*$ under mean-field term z^*

$$\begin{aligned} & \int_{\Gamma(s)} \pi^*(a \mid s) Q_{z^*}^*(s, a) da - \mathcal{H}(\pi^*) \\ & \geq \int_{\Gamma(s)} \pi(a \mid s) Q_{z^*}^*(s, a) da - \mathcal{H}(\pi). \end{aligned}$$

The quantal response equilibrium introduces the notion of bounded rationality that characterizes human decision-

making, where the agent makes stochastic ‘‘quantal’’ decisions according to a smooth probability distribution around the best response (Goeree et al., 2020).

Under this new notion of equilibrium, the equilibrium policy π^* is unique with respect to $Q_{z^*}^*$ thanks to the strongly convex entropy regularization \mathcal{H} . It is noteworthy that the pair (π^*, z^*) of policy and mean-field term alone is enough to characterize the equilibrium; the optimal Q-function $Q_{z^*}^*$ is deterministic given z^* and the aggregation indicators $\Xi^* = \Psi(z^*, \pi^*)$ are also fixed given π^* and z^* . Especially, the optimal policy can be estimated from data through RL algorithms when the mean-field term is fixed, and therefore we only consider the approximation error to the equilibrium mean-field term z^* in the rest of this paper.

3 A REINFORCEMENT LEARNING ALGORITHM

In this section, we propose a RL-based algorithm called `ConcaveHAM` for learning the QR-MFE of the heterogeneous agent model from offline data. The main algorithm is displayed in Algorithm 1, which, at a high level, consists of two major procedures: 1) a mean-field term generation from regularized optimal policy, and 2) an optimal value function estimation under concave shape constraint.

Algorithm 1 executes as follows: We initialize Algorithm 1 with an arbitrary mean-field term $z^0 \in \mathcal{Z}$, zero Q-function Q^0 , and uniform policy $\hat{\pi}^0$, which is indeed the regularized optimal policy of constant Q-function Q^0 . In each iteration $t \in [T]$, given the previous mean-field term z^{t-1} and the representative agent’s policy $\hat{\pi}^{t-1}$, the representative firm determines a new mean-field term z^t (Line 3). Then under z^t , we estimate the optimal Q-function \hat{Q}^t for the representative agent with the concave fitted Q-iteration (FQI) (Algorithm 2), using the offline dataset (Line 4). The regularized optimal policy is further solved according to (3.1) with respect to \hat{Q}^t (Line 5). Finally, Algorithm 1 returns an approximation z^T of the equilibrium mean-field term z^* after T iterations.

As we will show in Theorem 4.6, the mean-field term estimation of Algorithm 1 converges to the equilibrium quantity z^* under proper conditions. In the rest of this section, we explain the details for each component of Algorithm 1.

Regularized Optimal Policy. For any Q-function Q and any strongly convex regularizer \mathcal{H} , we define the regularized optimal policy π_Q with respect to Q as

$$\pi_Q(\cdot | s) := \arg \max_{u \in \Delta_{\Gamma(s)}} \left\{ \int_{\Gamma(s)} u(a) Q(s, a) da - \mathcal{H}(u) \right\}. \quad (3.1)$$

In particular, we assume \mathcal{H} is ζ -strongly convex in $\|\cdot\|_1$. A classic example is the negative entropy, i.e., $-\mathcal{H}(u) =$

Algorithm 1 Concave Heterogeneous Agent Model (`ConcaveHAM`)

Input: Mean-field term $z^0 \in \mathcal{Z}$, number of iterations T , number of iterations τ within each CFQI

- 1: **initialize** $Q^0(s, a) \leftarrow 0$ for all $(s, a) \in \mathcal{G}$, $\hat{\pi}^0 \leftarrow \text{Unif}(\Gamma(s))$ for all $s \in \mathcal{S}$
- 2: **for** $t = 1, \dots, T$ **do**
- 3: $z^t \leftarrow \Phi(\Psi(z^{t-1}, \hat{\pi}^{t-1}))$
- 4: $\hat{Q}^t \leftarrow \text{CFQI}(\mathcal{D}_t; \tau)$ where \mathcal{D}_t is the offline dataset under z^t
- 5: $\hat{\pi}^t \leftarrow \pi_{\hat{Q}^t}$ following Eq. (3.1)
- 6: **end for**
- 7: **Return** z^T

$\zeta \int u(a) \log u(a) da$. We further assume π_Q is known for any Q-function Q and regularizer \mathcal{H} for simplicity, and the policy optimization with entropy regularization can be found in Mei et al. (2020); Chen et al. (2021b).

Q-Function Estimation. For any fixed mean-field term z , we propose a Concave FQI algorithm in Algorithm 2, abbreviated as `CFQI`, for estimating the optimal Q-function Q_z^* of the underlying MDP from data.

Under each mean-field term z , we define the Bellman optimality operator \mathcal{T}_z as follows:

$$(\mathcal{T}_z Q)(s, a) := r_z(s, a) + \gamma \mathbb{E}_{s' \sim \mathbb{P}_z(\cdot | s, a)}[(\mathfrak{J}Q)(s')],$$

where $(\mathfrak{J}Q)(s) := \max_{a \in \Gamma(s)} Q(s, a)$. For each mean-field term $z \in \mathcal{Z}$, we assume access to an offline dataset $\mathcal{D}_z := \{(s_{z,m}, a_{z,m}, r_{z,m}, s'_{z,m})\}_{m=1}^M$ that contains M i.i.d. samples, where $(s_{z,m}, a_{z,m}) \sim \tilde{\mu} \in \Delta_{\mathcal{G}}$, the reward $r_{z,m} = r_z(s_{z,m}, a_{z,m})$, and the next state $s'_{z,m} \sim \mathbb{P}_z(\cdot | s_{z,m}, a_{z,m})$. Then given the dataset \mathcal{D}_z , FQI (Antos et al., 2007) applies an approximate value iteration on the Q-functions, by iteratively computing an estimator \tilde{Q}^ℓ of $\mathcal{T}_z \tilde{Q}^{\ell-1}$ for $\ell = 1, \dots, \tau$.

In particular, `CFQI` enforces a concave shape constraint on the estimated Q-function through a convex regression solver termed least square estimator (LSE). Given $\tilde{Q}^{\ell-1}$ from the previous iteration and a set of concave functions denoted by \mathcal{F} , LSE finds $\hat{f} \in \mathcal{F}$ that minimizes the empirical risk given by the Bellman error between \hat{f} and $\tilde{Q}^{\ell-1}$ evaluated on the dataset \mathcal{D}_z . Specifically, for any Q and $\hat{f} \in \mathcal{F}$, we say \hat{f} is an ϵ -approximate LSE of $\mathcal{T}_z Q$ on dataset \mathcal{D}_z if

$$\sum_{(s,a,r,s') \in \mathcal{D}_z} (\hat{f}(s, a) - r - \gamma(\mathfrak{J}Q)(s'))^2 - \min_{f \in \mathcal{F}} \sum_{(s,a,r,s') \in \mathcal{D}_z} (f(s, a) - r - \gamma(\mathfrak{J}Q)(s'))^2 \leq \epsilon.$$

We denote the regression solver used to acquire such ϵ -approximate LSE as $\text{LSE}(\mathcal{F}, \mathcal{D}_z, \mathcal{T}_z Q; \epsilon)$. Then in each iteration ℓ , we solve for $\tilde{Q}^\ell = \text{LSE}(\mathcal{F}, \mathcal{D}_z, \mathcal{T}_z \tilde{Q}^{\ell-1}; \epsilon)$,

Algorithm 2 Concave FQI (CFQI)

Input: Data $\mathcal{D}_z = \{s_{z,m}, a_{z,m}, r_{z,m}, s'_{z,m}\}_{m=1}^M$, number of iterations τ

- 1: **initialize** $\tilde{Q}^0(s, a) \leftarrow 0$ for all $(s, a) \in \mathcal{G}$
- 2: **for** $\ell = 1, \dots, \tau$ **do**
- 3: $\tilde{Q}^\ell \leftarrow \text{LSE}(\mathcal{F}, \mathcal{D}_z, \mathcal{J}_z \tilde{Q}^{\ell-1}; \epsilon)$
- 4: **end for**
- 5: **Return** \tilde{Q}^τ

and Algorithm 2 terminates after τ iterations and returns an approximate optimal Q-function \tilde{Q}^τ .

In general LSE is hard to obtain without structures on \mathcal{F} , but fortunately here we have the convex shape constraint. Recall that the reward function r_z is concave and takes values in $[0, 1]$, and it follows that Q_z^π and V_z^π are in $[0, 1/(1-\gamma)]$. Let us denote $B := 1/(1-\gamma)$ and $L := L_r/(1-\gamma)$, then the Q-function Q_z^π is B -bounded and L -Lipschitz for any $z \in \mathcal{Z}$ and policy π .

Hence, as Lemma B.5 shows, the estimation target $\mathcal{J}_z Q$ is a concave function for any concave Q , and this allows effective estimation of $\mathcal{J}_z Q$ with a simple concave function set. We also remark that if the model has finite number of discrete income levels for variable w , the regression target $\mathcal{J}_z Q(\cdot, w, \cdot)$ is concave for any concave $Q(\cdot, w, \cdot)$ and all $w \in \mathcal{W}$ without the requirement of stochastic concavity, cf. Lemma B.6.

3.1 Convex Regression

We now introduce a few regression solvers available to Algorithm 2. Without loss of generality, for any regression on a concave Q bounded from above by B , we consider an equivalent regression problem on the non-negative convex function $-Q + B$. With a slight abuse of notation, we write Q to denote this transformed convex function from now on to simplify discussions.

Bounded Lipschitz Convex Functions. Let $\mathcal{C}_{B,L}$ denote the set of all non-negative B -bounded and L -Lipschitz convex functions on \mathcal{G} , i.e.,

$$\mathcal{C}_{B,L} := \{f : \mathcal{G} \rightarrow [0, B] \mid f \text{ is convex; } \partial f(x) \neq \emptyset, \forall x \in \mathcal{G}; \|\partial f(x)\| \leq L, \forall x \in \mathcal{G}\},$$

where $\partial f(x)$ denotes the subdifferential of f at $x \in \mathcal{G}$ and $\|\partial f(x)\|$ denotes $\sup_{g \in \partial f(x)} \|g\|_\infty$. Recall that $\mathcal{J}_z Q \in \mathcal{C}_{B,L}$ for any $Q \in \mathcal{C}_{B,L}$ and mean-field term z following Lemma B.5. A convex regression solver for $\mathcal{J}_z Q$ is written as $\text{LSE}(\mathcal{C}_{B,L}, \mathcal{D}_z, \mathcal{J}_z Q; \epsilon)$.

Lipschitz Max-Affine Functions.

Further, we let $\mathcal{A}_{B,L}^{K,+}$ denote the set of all non-negative L -Lipschitz K -max-affine functions, i.e.,

$$\mathcal{A}_{B,L}^{K,+} := \{h : \mathcal{G} \rightarrow [0, B] \mid h(x) = \max_{k \in [K]} \alpha_k^\top x + c_k,$$

$$\|\alpha_k\|_\infty \leq L, \forall k \in [K]\}.$$

Note that max-affine functions are convex by definition, and therefore any estimator in $\mathcal{A}_{B,L}^{K,+}$ also lies in $\mathcal{C}_{B,L}$. We write a K -max-affine regression solver for $\mathcal{J}_z Q$ as $\text{LSE}(\mathcal{A}_{B,L}^{K,+}, \mathcal{D}_z, \mathcal{J}_z Q; \epsilon)$. It is worth noting that an approximate LSE to the convex regression on both $\mathcal{C}_{B,L}$ and $\mathcal{A}_{B,L}^{K,+}$ can be obtained by solving the corresponding convex optimization problems up to arbitrary accuracy (Mazumder et al., 2019; Balázs et al., 2015), and we therefore assume the approximate LSEs returned by the solvers admit at most ϵ excess empirical risk compared to the LSE.

Input Convex Neural Networks. Beyond the standard convex regression algorithms, the shape-constrained estimation of $\mathcal{J}_z Q$ can be achieved by minimizing the empirical risk over a set of input convex neural network (ICNN), where the network output is always convex with respect to the network input when the parameter space of the network is constrained (Amos et al., 2017).

More specifically, for any K -layer ICNN $f_K(x) : \mathbb{R}^d \rightarrow \mathbb{R}$, each of its i -th layer y_i can be expressed as

$$y_{i+1} = \sigma(W_i^{(y)} y_i + W_i^{(x)} x - \beta_i), \quad i = 0, 1, \dots, K-1,$$

where $y_i, \beta_i \in \mathbb{R}^{\ell_i}$, $W_i^{(y)} \in \mathbb{R}^{\ell_{i+1} \times \ell_i}$, $W_i^{(x)} \in \mathbb{R}^{\ell_{i+1} \times d}$, and σ denotes ReLU function. Especially, $y_0 = x \in \mathbb{R}^d$ represents the network input, $y_K = f_K(x) \in \mathbb{R}$ represents the network output, and ℓ_i denotes the number of neurons on the i -th layer (especially, $\ell_0 = d$). To guarantee the convexity of $f_K(x)$ with respect to x , we further restrict $W_i^{(y)} \in \mathbb{R}_+^{\ell_{i+1} \times \ell_i}$, so that the composition of convex and convex non-decreasing function is also convex (Amos et al., 2017).

We further show that under a convex parameter set of the network specified by \mathcal{G} , we are able to construct a set $\mathcal{N}_{B,L}^K$ of ICNNs such that it is equivalent to the set of all non-negative B -bounded and L -Lipschitz K -max-affine functions, i.e., $\mathcal{A}_{B,L}^{K,+}$. Therefore, minimizing the empirical risk on $\mathcal{N}_{B,L}^K$ also serves as an estimation of the convex function $\mathcal{J}_z Q$, assuming that the solver trains to the global minimum. This is summarized in the following lemma, and see Appendix B.7 for a proof.

Lemma 3.1. *There exists a set $\mathcal{N}_{B,L}^K$ of K -layer ICNNs, such that for any L -Lipschitz function $f_K : \mathcal{X} \rightarrow [0, B]$ defined on a convex domain $\mathcal{X} \subseteq \mathbb{R}^d$ that is the maximum of K affine functions, $\mathcal{N}_{B,L}^K$ is able to represent f_K exactly under a convex constraint set specified by $W_i^{(x)} \geq 0$, $\|\sum_{j=i-1}^{K-1} w_j\|_\infty \leq L$, and $\max_{x \in \mathcal{X}} \sum_{j=i-1}^{K-1} w_j^\top x + \sum_{j=i-1}^{K-1} b_j \leq B$ for all $i \in [K]$, where $w_i := w_{i,1} - w_{i,2}$ and $W_i^{(x)} = [w_{i,1}^\top, w_{i,2}^\top]^\top$. In addition, we set $W_0^{(y)} = 0$ and $W_i^{(y)} = 1$ fixed for all $i \in [K-1]$. Especially, $\mathcal{N}_{B,L}^K$ is equivalent to the set of all non-negative K -max-affine functions that are B -bounded and L -Lipschitz.*

4 THEORETICAL RESULTS

Next, we show the existence and uniqueness of the QR-MFE, and provide theoretical guarantees to the convergence of Algorithm 1 to the unique equilibrium. Before diving into the main results, we make a few moderate assumptions on the MFG, under which our theoretical results hold.

We first define formally the shape constrain assumption on the feasible action sets and reward functions, as well as some technical conditions required for the following theorems. It is noteworthy that we consider b and w as continuous variables.

Assumption 4.1. The set of all feasible mean-field terms $\mathcal{Z} \subseteq \mathbb{R}^{d_z}$ has bounded ℓ_1 -radius, i.e., $\max_{z, z' \in \mathcal{Z}} \|z - z'\|_1 \leq Z$ for some $Z > 0$. The feasible set $\mathcal{G} = \{(s, a) \in \mathbb{R}^{d_s + d_w} \mid a \in \Gamma(s)\}$ is a closed convex set. The reward function $r_z(s, a)$ is concave in (s, a) and L_r -Lipschitz in ℓ_∞ norm with respect to (s, a, z) , i.e., $|r_z(s, a) - r_{z'}(s', a')| \leq L_r \cdot \max\{\|s - s'\|_\infty, \|a - a'\|_\infty, \|z - z'\|_\infty\}$.

Assumption 4.2. For all $z \in \mathcal{Z}$, the transition kernel \mathbb{P}_z of the MDP is stochastically concave (Smith and McCardle, 2002), i.e., $q(b, w, a) := \mathbb{E}_{w'}[Q(b, w', a) \mid w]$ is concave on for any concave Q .

We assume that the transition kernel (2.1) is Lipschitz with respect to both the state and mean-field term. Recall that the transition is independent of the current capital holdings b and savings a (see Equation 2.1). Such assumption admits smooth transition between different mean-field environments.

Assumption 4.3. The transition kernel \mathbb{P} is Lipschitz in the sense that for any mean-field term $z, z' \in \mathcal{Z}$ and $s, s' \in \mathcal{S}$, there exists $L_{\mathbb{P}} > 0$ such that $\|\mathbb{P}_z(\cdot \mid w) - \mathbb{P}_{z'}(\cdot \mid w')\|_1 \leq L_{\mathbb{P}} \cdot (\|s - s'\| + \|z - z'\|_1)$.

To guarantee a stable iteration for Algorithm 1, we also impose a Lipschitz condition on the aggregation function Ψ and the production mapping Φ . Here for any two policies π and π' , $\|\pi - \pi'\|_{\infty, \bar{\nu}}$ is defined as $\int_{\mathcal{S}} \|\pi(\cdot \mid s) - \pi'(\cdot \mid s)\|_\infty d\nu(s)$ for any $\nu \in \Delta_{\mathcal{S}}$.

Assumption 4.4. The aggregate function Ψ is L_Ψ -Lipschitz for all policy π and mean-field term $z \in \mathcal{Z}$ and the production mapping Φ is L_Φ -Lipschitz, i.e.,

$$\begin{aligned} \|\Psi(z, \pi) - \Psi(z', \pi')\|_1 &\leq L_\Psi(\|z - z'\|_1 + \|\pi - \pi'\|_{\infty, \bar{\nu}}), \\ \|\Phi(\Xi) - \Phi(\Xi')\|_1 &\leq L_\Phi \cdot \|\Xi - \Xi'\|_1, \end{aligned}$$

for some policy-induced distribution $\bar{\nu}$. Hence, $\Phi(\Psi(\cdot, \cdot))$ is L_F -Lipschitz, where $L_F := L_\Psi L_\Phi$.

As mentioned above, we assume that for any mean-field term z , we have access to a logged dataset \mathcal{D}_z comprising

of i.i.d. samples from the trajectory of an exploratory behavior policy, where the behavior policy satisfies a coverage assumption. These are standard for FQI in the RL literature (Antos et al., 2007), and the state-action pairs needed for the training dataset can be sampled from the heterogeneous households in the market.

Assumption 4.5. For any $z \in \mathcal{Z}$, the offline data $\mathcal{D}_z = \{(s_{z,m}, a_{z,m}, r_{z,m}, s'_{z,m})\}_{m=1}^M$ follows an exploratory distribution $\tilde{\mu}$, where i.i.d. samples $(s_{z,m}, a_{z,m}) \sim \tilde{\mu}$. In particular, $\tilde{\mu} = \tilde{\nu} \times \pi_b$ follows from a sample distribution $\tilde{\nu}$ on \mathcal{S} and an exploratory behavior policy π_b , such that for any mean-field term z , non-stationary policy π^* , step h , and feasible state-action pair (s, a) , there exists a constant D such that

$$\frac{\mu_{z,h}^\pi(s, a)}{\tilde{\mu}(s, a)} \leq D,$$

where μ_z^π denotes the distribution on step h over $\mathcal{S} \times \mathcal{A}$ induced by transition kernel \mathbb{P}_z and policy π . Note that $\tilde{\nu}(s) = \int_{\Gamma(s)} \tilde{\mu}(s, a) da$, and we have $\int_{\Gamma(s)} \mu_{z,h}^\pi(s, a) da / \tilde{\nu}(s) \leq D$ for any $s \in \mathcal{S}$.

Under these assumptions, we show the existence and uniqueness of the QR-MFE for the MFG as well as the convergence guarantees of Algorithm 1 with a range of external solvers for convex regression.

To ease the notation, let us introduce some problem-dependent quantities: Let $B := 1/(1-\gamma)$, $L := L_r/(1-\gamma)$, $d := d_s + d_w$, and

$$\kappa := \frac{JL_F D^2}{\zeta} + L_F.$$

We further write

$$C_d = B^{\frac{d+8}{d+2}}(d+1)^{\frac{2}{d+2}}(dL)^{\frac{d}{d+2}}$$

and $C'_d = (d+1)B^4 + (B+dL)dL$.

First, we prove that the MFG has a unique QR-MFE. The essential idea is that if $\Phi(\Psi(\pi_{Q_z^*}, z))$ a contraction and \mathcal{Z} is complete, then the equilibrium is unique. See proof in Appendix B.1.

Theorem 4.6 (Unique QR-MFE). *Under Assumptions 4.1, 4.2, 4.4 and 4.5, further assume that $\kappa < 1$. Then, there exists a unique quantal response mean-field equilibrium to the MFG with regularized policies.*

Further, we show that Algorithm 1 converges to the QR-MFE with a range of external solvers that provides approximate LSE on the set of convex functions, max-affine functions, and ICNNs. In particular, as we will see in Theorems 4.7 and 4.8, the RL framework ConcaveHAM enjoys

*A non-stationary policy may change over time, i.e., there may not exist a policy $\pi : \mathcal{S} \rightarrow \Delta_{\mathcal{A}}$ such that $\pi_t = \pi$ for all t .

sub-linear convergence rate to the equilibrium mean-field term z^* for all three convex regression solvers if the estimation error $\epsilon = 0$. Proofs for Theorems 4.7 and 4.8 can be found in Appendices B.2 and B.3.

Theorem 4.7 (Convergence for $\mathcal{C}_{B,L}$). *Under Assumptions 4.1, 4.2, 4.3, 4.4, and 4.5, further assume that $\kappa < 1$. For $\delta \in (0, 1)$, suppose that the number of iterations within each CFQI subroutine satisfies*

$$\tau \asymp \frac{\log \left(C_d^{-1} D^{-1/2} \frac{M^{2/(d+2)}}{\log M} \right)}{\log \frac{1}{\gamma}},$$

and that the sample size M satisfies

$$\log M \gtrsim \max \left\{ \frac{B^{\frac{4d+8}{d+2}}}{C_d M^{\frac{d}{d+2}}} \log \frac{\tau T}{\delta}, \left(1 + \log \frac{R_d^* B}{C_d} \right)^{\frac{2}{d}} \right\},$$

where $R_d^* \leq \max\{8dL, 2B + 4dL\}$. Then with probability at least $1 - \delta$, Algorithm 1 with T iterations on $\mathcal{C}_{B,L}$ gives mean-field term z^T such that

$$\|z^T - z^*\|_1 \lesssim \frac{C_d L_F D^{5/2}}{(1-\gamma)(1-\kappa)\zeta} M^{-\frac{2}{d+2}} \log M + \frac{L_F D^{5/2} \epsilon}{(1-\gamma)(1-\kappa)\zeta} + \kappa^T Z + \frac{\kappa^{T-1}}{1-\gamma}.$$

Theorem 4.7 suggests that when we take convex regression on $\mathcal{F} = \mathcal{C}_{B,L}$ for the estimation of optimal value function in Algorithm 2, the estimation z^T of ConcaveHAM converges to z^* in the rate of $\tilde{O}(M^{-\frac{2}{d+2}})$ when the sample size M is large enough. Note that $\kappa^T Z + \frac{\kappa^{T-1}}{1-\gamma}$ converges to zero exponentially fast, and the number of required iterations T is in logarithmic order. It is noteworthy that the number of iterations τ in Algorithm 2 is kept in the order of $\log(M^{\frac{2}{d+2}} / \log M)$ to guarantee the convergence of ConcaveHAM.

Due to the large covering entropy of $\mathcal{C}_{B,L}$ in the order of d^d , the sample complexity with the term C_d in Theorem 4.7 suffers from an almost linear dependence on d . This issue can be avoided by taking the regression solver $\text{LSE}(\mathcal{A}_{B,L}^{K,+}, \mathcal{D}_z, \mathcal{T}_z Q; \epsilon)$ on the max-affine function set $\mathcal{A}_{B,L}^{K,+}$ with a specially designed K .

Theorem 4.8 (Convergence for $\mathcal{N}_{B,L}^K$). *Under Assumptions 4.1, 4.2, 4.3, 4.4, and 4.5, further assume that $\kappa < 1$. For $\delta \in (0, 1)$, suppose that the number of iterations within each CFQI subroutine satisfies*

$$\tau \asymp \frac{\log \frac{M^{4/(d+4)}}{d^2 L^2 \sqrt{D}}}{\log \frac{1}{\gamma}},$$

and that the sample size M satisfies

$$M \gtrsim \max \left\{ \frac{\log \frac{M^{4/(d+4)}}{d^2 L^2 \sqrt{D}}}{\delta \log \frac{1}{\gamma}}, \frac{\tau T}{\exp(2(d+1)M^{d/(d+4)})\delta} \right\},$$

$$\left(\frac{B^3(B+dL)}{d \log M} \right)^{\frac{d+4}{2d+4}}, \left(\frac{B^2}{d^2 L^2} \log \frac{\tau T}{\delta} \right)^{\frac{d+4}{d}} \}.$$

then with probability at least $1 - \delta$, Algorithm 1 with T iterations on $\mathcal{A}_{B,L}^{K,+}$ or $\mathcal{N}_{B,L}^K$ gives mean-field term z^T such that

$$\|z^T - z^*\|_1 \lesssim \frac{C_d L_F D^{5/2}}{(1-\gamma)(1-\kappa)\zeta} M^{-\frac{4}{d+4}} \log M + \frac{L_F D^{5/2} \epsilon}{(1-\gamma)(1-\kappa)\zeta} + \kappa^T Z + \frac{\kappa^{T-1}}{1-\gamma}.$$

With the regression solver on max-affine function set $\mathcal{A}_{B,L}^K$, ConcaveHAM converges in the rate of $\tilde{O}(M^{-\frac{4}{d+4}})$ when the sample size M is large enough. Compared to Theorem 4.7, it not only enjoys a better rate but also enjoys a better constant in d , compared to the one in Theorem 4.7 for $\mathcal{C}_{B,L}$. ConcaveHAM also enjoys the same rate of $\tilde{O}(M^{-\frac{4}{d+4}})$ for large M when paired with regression on the set $\mathcal{N}_{B,L}^K$ of ICNNs, and this is a natural statement following Lemma 3.1 and assuming the learning algorithms $\text{LSE}(\mathcal{N}_{B,L}^K, \mathcal{D}_z, \mathcal{T}_z Q; \epsilon)$ approach the global optimum on minimizing the empirical risk.

Therefore, the ConvHAM algorithm is provably convergent to the QR-MFE of the mean-field game at a sublinear rate, which provides a data-driven approach to a large class of heterogeneous agent models that does not require prior knowledge of the income shock transitions. By taking advantage of convex regression, it is capable of capturing the economic intuitions, and at the same time, providing a continuous solution on value functions and avoiding the curse of dimensionality in discretization.

5 NUMERICAL EXPERIMENTS

To corroborate our theoretical results, we apply the proposed ConvHAM algorithm to some of the existing heterogeneous agent models in macroeconomics and show the convergence rate on mean-field terms that matches our theory.

Experiment Setup. Our simulation features a canonical setup of the Aiyagari model where market incompleteness and idiosyncratic income shocks give rise to agent heterogeneity.

Under this setup, the household has two levels of labor opportunities $n \in \{0, 1\}$, i.e., employment if $n = 1$ and unemployment if $n = 0$. The utility function is logarithmic with respect to the amount of spending χ that is restricted by the household budget constraint $\chi \leq (1 + \varsigma - \delta)b + \omega n - a$, where the mean-field term $z = (\omega, \varsigma)$ represents the wage and the rent on the market while δ denotes a depreciation constant. We restrict feasible capital holding $0 \leq b \leq 20$ and the representative firm takes the Cobb-Douglas production function $F(b, n) = b^\alpha n^{1-\alpha}$ for any




Figure 1: Convergence of the Mean-Field Terms: Rent ζ and Wage ω . The red, blue, and purple lines show the average outputs of the mean-field terms by `ConvHAM` after 15 rounds of iterations with different number of samples M provided at each iteration. The corresponding confidence bands are plotted over 50 random trials. The black dotted line indicates the convergence rate of order $\mathcal{O}(M^{-2/5})$.

aggregated population capital b and population labor n . The training data are sampled from a uniform distribution on feasible states $s = (b, n)$ following a uniform policy π_0 .

Simulation Results. Our method is able to achieve stable convergence in a few dozens of rounds of `ConvHAM` with moderate regularization, and the convergence rate matches our theoretical guarantees as illustrated in Figure 1.

More specifically, we aggregate the numerical results over 50 random trials for three initialization of population capital, which is determined by the initial mean-field terms, with a range of sample sizes M for the concave value function estimation. We plot the mean-field terms (wage and rent) against different sample sizes for 15 rounds of `ConcaveHAM`, and compare them with the theoretically predicted convergence rate of $M^{-2/5}$ in terms of the sample size. The numerical results verify that (i) our algorithm achieves stable convergence, and (ii) the convergence rate matches our theoretical guarantees.

6 RELATED WORK

Heterogeneous agent models are a fundamental class of macroeconomic models that aim to understand the dynamics of individual participants of the market and how their population behaviors impact the overall performance of the economy (Hommes, 2006). These models often consider random economic shocks to individual participants, such as aggregate and exogenous risk (Fernández-Villaverde et al., 2019), and focus on understanding the RCE between households and firms (Kuhn, 2013; Light, 2020). Machine learning has gained momentum in analyzing a range of economic models (Lepetyuk et al., 2020; Fernandez-Villaverde et al., 2020). Especially, deep RL has claimed success in solving heterogeneous agent models with aggregate shocks and discrete-continuous choice dynamic models with heterogeneous agents (Han et al., 2021; Maliar and Maliar, 2022); it is also proven effective in finding micro-founded

general equilibria for macroeconomic models with many agents (Curry et al., 2022).

MFG characterizes the decision-making dynamics of infinitely many players through a mean-field term, and it provides a natural extension to multi-agent games that admits simpler asymptotic analysis (Lasry and Lions, 2007; Guéant et al., 2011). Convergence of RL algorithms to equilibria of MFGs have been analyzed in recent works (Guo et al., 2019; Subramanian and Mahajan, 2019). Particularly, Guo et al. (2022) studied MFGs with entropy regularization on discrete action space and finite time horizon.

Shape-constrained non-parametric modeling is a powerful tool that enables reliable estimation of special function classes, and it is especially useful for economic models that often demand shape constraints on utility functions (Blundell et al., 2003). In particular, theoretical guarantees have been laid out for convex and max-affine shape constraint regression in a thread of literature (Lim and Glynn, 2012; Han and Wellner, 2016; Mazumder et al., 2019; Seijo and Sen, 2011; Balázs et al., 2015). An architecture of ICNN has also been considered for the representation of convex functions with neural networks (Amos et al., 2017; Chen et al., 2018).

7 CONCLUSION

We study a generalized Aiyagari model as an MFG with concave shape constraint imposed on the agent’s value function. We propose a data-driven RL framework `ConcaveHAM` for solving the MFG that dramatically improves the computational efficiency with the help of a variety of convex regression procedures. The algorithm converges to the regularized recursive competitive equilibrium of the macroeconomic model at a sub-linear rate and can readily incorporate other forms of shape constraints. To the best of our knowledge, this is the first data-driven algorithm for heterogeneous agent models with theoretical guarantees.

Acknowledgements

Zhaoran Wang acknowledges National Science Foundation (Awards 2048075, 2008827, 2015568, 1934931), Simons Institute (Theory of Reinforcement Learning), Amazon, J.P. Morgan, and Two Sigma for their support. Zhuoran Yang acknowledges Simons Institute (Theory of Reinforcement Learning) for the support.

References

- Achdou, Y., Buera, F. J., Lasry, J.-M., Lions, P.-L., and Moll, B. (2014a). Partial differential equation models in macroeconomics. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 372(2028):20130397.
- Achdou, Y., Han, J., Lasry, J.-M., Lions, P.-L., and Moll, B. (2014b). Heterogeneous agent models in continuous time. *Preprint*, 14.
- Achdou, Y., Han, J., Lasry, J.-M., Lions, P.-L., and Moll, B. (2022). Income and wealth distribution in macroeconomics: A continuous-time approach. *The review of economic studies*, 89(1):45–86.
- Aiyagari, S. R. (1994). Uninsured idiosyncratic risk and aggregate saving. *The Quarterly Journal of Economics*, 109(3):659–684.
- Amos, B., Xu, L., and Kolter, J. Z. (2017). Input convex neural networks. In *International Conference on Machine Learning*, pages 146–155. PMLR.
- Antos, A., Munos, R., and Szepesvári, C. (2007). Fitted q-iteration in continuous action-space mdps.
- Arora, S., Du, S., Hu, W., Li, Z., and Wang, R. (2019). Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks. In *International Conference on Machine Learning*, pages 322–332. PMLR.
- Ayoub, A., Jia, Z., Szepesvári, C., Wang, M., and Yang, L. (2020). Model-based reinforcement learning with value-targeted regression. In *International Conference on Machine Learning*, pages 463–474. PMLR.
- Balázs, G., György, A., and Szepesvári, C. (2015). Near-optimal max-affine estimators for convex regression. In *Artificial Intelligence and Statistics*, pages 56–64. PMLR.
- Bellman, R. E. and Dreyfus, S. E. (2015). *Applied dynamic programming*, volume 2050. Princeton university press.
- Bertsekas, D. (2012). *Dynamic programming and optimal control: Volume I*, volume 1. Athena scientific.
- Bewley, T. (1986). Stationary monetary equilibrium with a continuum of independently fluctuating consumers. *Contributions to mathematical economics in honor of Gérard Debreu*, 79.
- Blanchet, J., Glynn, P. W., Yan, J., and Zhou, Z. (2019). Multivariate distributionally robust convex regression under absolute error loss. *Advances in Neural Information Processing Systems*, 32.
- Blundell, R. W., Browning, M., and Crawford, I. A. (2003). Nonparametric engel curves and revealed preference. *Econometrica*, 71(1):205–240.
- Cai, Q., Yang, Z., Jin, C., and Wang, Z. (2020). Provably efficient exploration in policy optimization. In *International Conference on Machine Learning*, pages 1283–1294. PMLR.
- Chen, J. and Jiang, N. (2019). Information-theoretic considerations in batch reinforcement learning. In *International Conference on Machine Learning*, pages 1042–1051. PMLR.
- Chen, L., Min, Y., Belkin, M., and Karbasi, A. (2021a). Multiple descent: Design your own generalization curve. *Advances in Neural Information Processing Systems*, 34:8898–8912.
- Chen, L., Min, Y., Zhang, M., and Karbasi, A. (2020). More data can expand the generalization gap between adversarially robust and standard models. In *International Conference on Machine Learning*, pages 1670–1680. PMLR.
- Chen, Y., Dong, J., and Wang, Z. (2021b). A primal-dual approach to constrained markov decision processes. *arXiv preprint arXiv:2101.10895*.
- Chen, Y., Shi, Y., and Zhang, B. (2018). Optimal control via neural networks: A convex approach. *arXiv preprint arXiv:1805.11835*.
- Christmann, A. and Steinwart, I. (2007). Consistency and robustness of kernel-based regression in convex risk minimization.
- Curry, M., Trott, A., Phade, S., Bai, Y., and Zheng, S. (2022). Finding general equilibria in many-agent economic simulations using deep reinforcement learning. *CoRR*, abs/2201.01163.
- Du, S., Kakade, S., Lee, J., Lovett, S., Mahajan, G., Sun, W., and Wang, R. (2021). Bilinear classes: A structural framework for provable generalization in rl. In *International Conference on Machine Learning*, pages 2826–2836. PMLR.
- Du, S. S., Hou, K., Salakhutdinov, R. R., Póczos, B., Wang, R., and Xu, K. (2019). Graph neural tangent kernel: Fusing graph neural networks with graph kernels. *Advances in neural information processing systems*, 32.
- Dubey, A. and Pentland, A. (2021). Provably efficient cooperative multi-agent reinforcement learning with function approximation. *arXiv preprint arXiv:2103.04972*.
- Fei, Y. and Xu, R. (2022a). Cascaded gaps: Towards gap-dependent regret for risk-sensitive reinforcement learning. *arXiv preprint arXiv:2203.03110*.

- Fei, Y. and Xu, R. (2022b). Cascaded gaps: Towards logarithmic regret for risk-sensitive reinforcement learning. In *International Conference on Machine Learning*, pages 6392–6417. PMLR.
- Fernández-Villaverde, J., Hurtado, S., and Nuno, G. (2019). Financial frictions and the wealth distribution. Technical report, National Bureau of Economic Research.
- Fernandez-Villaverde, J., Nuno, G., Sorg-Langhans, G., and Vogler, M. (2020). Solving high-dimensional dynamic programming problems using deep learning. *Unpublished working paper*.
- Fonteneau, R., Murphy, S. A., Wehenkel, L., and Ernst, D. (2013). Batch mode reinforcement learning based on the synthesis of artificial trajectories. *Annals of operations research*, 208:383–416.
- Goeree, J. K., Holt, C. A., and Pfaffel, T. R. (2020). Stochastic game theory for social science: A primer on quantal response equilibrium. In *Handbook of Experimental Game Theory*. Edward Elgar Publishing.
- Grunewalder, S., Lever, G., Baldassarre, L., Pontil, M., and Gretton, A. (2012). Modelling transition dynamics in mdps with rkhs embeddings. *arXiv preprint arXiv:1206.4655*.
- Guéant, O., Lasry, J.-M., and Lions, P.-L. (2011). Mean field games and applications. In *Paris-Princeton lectures on mathematical finance 2010*, pages 205–266. Springer.
- Guo, X., Hu, A., Xu, R., and Zhang, J. (2019). Learning mean-field games. *Advances in Neural Information Processing Systems*, 32.
- Guo, X., Hu, A., Xu, R., and Zhang, J. (2020). A general framework for learning mean-field games. *arXiv preprint arXiv:2003.06069*.
- Guo, X., Xu, R., and Zariphopoulou, T. (2022). Entropy regularization for mean field games with learning. *Mathematics of Operations Research*.
- Györfi, L., Kohler, M., Krzyżak, A., and Walk, H. (2002). *A distribution-free theory of nonparametric regression*, volume 1. Springer.
- Han, J., Yang, Y., et al. (2021). Deepham: A global solution method for heterogeneous agent models with aggregate shocks.
- Han, Q. and Wellner, J. A. (2016). Multivariate convex regression: global risk bounds and adaptation. *arXiv preprint arXiv:1601.06844*.
- He, J., Wang, T., Min, Y., and Gu, Q. (2022). A simple and provably efficient algorithm for asynchronous federated contextual linear bandits. *Advances in neural information processing systems*.
- Hommes, C. H. (2006). Heterogeneous agent models in economics and finance. *Handbook of computational economics*, 2:1109–1186.
- Huggett, M. (1993). The risk-free rate in heterogeneous-agent incomplete-insurance economies. *Journal of Economic Dynamics and Control*, 17(5-6):953–969.
- Jin, C., Yang, Z., Wang, Z., and Jordan, M. I. (2020). Provably efficient reinforcement learning with linear function approximation. In *Conference on Learning Theory*, pages 2137–2143. PMLR.
- Kuhn, M. (2013). Recursive equilibria in an aiyagari-style economy with permanent income shocks. *International Economic Review*, 54(3):807–835.
- Kumar, A., Fu, J., Soh, M., Tucker, G., and Levine, S. (2019). Stabilizing off-policy q-learning via bootstrapping error reduction. *Advances in Neural Information Processing Systems*, 32.
- Lasry, J.-M. and Lions, P.-L. (2007). Mean field games. *Japanese journal of mathematics*, 2(1):229–260.
- Lepetyuk, V., Maliar, L., and Maliar, S. (2020). When the us catches a cold, canada sneezes: a lower-bound tale told by deep learning. *Journal of Economic Dynamics and Control*, 117:103926.
- Light, B. (2020). Uniqueness of equilibrium in a bewley–aiyagari model. *Economic Theory*, 69(2):435–450.
- Light, B. and Weintraub, G. Y. (2022). Mean field equilibrium: uniqueness, existence, and comparative statics. *Operations Research*, 70(1):585–605.
- Lim, E. and Glynn, P. W. (2012). Consistency of multi-dimensional convex regression. *Operations Research*, 60(1):196–208.
- Ling, S., Xu, R., and Bandeira, A. S. (2019). On the landscape of synchronization networks: A perspective from nonconvex optimization. *SIAM Journal on Optimization*, 29(3):1879–1907.
- Lu, M., Min, Y., Wang, Z., and Yang, Z. (2023). Pessimism in the face of confounders: Provably efficient offline reinforcement learning in partially observable markov decision processes. In *International Conference on Learning Representation*.
- Maliar, L. and Maliar, S. (2022). Deep learning classification: Modeling discrete labor choice. *Journal of Economic Dynamics and Control*, 135:104295.
- Mazumder, R., Choudhury, A., Iyengar, G., and Sen, B. (2019). A computational framework for multivariate convex regression and its variants. *Journal of the American Statistical Association*, 114(525):318–331.
- Mei, J., Xiao, C., Szepesvari, C., and Schuurmans, D. (2020). On the global convergence rates of softmax policy gradient methods. In *International Conference on Machine Learning*, pages 6820–6829. PMLR.
- Min, Y., Chen, L., and Karbasi, A. (2021a). The curious case of adversarially robust models: More data can help, double descend, or hurt generalization. In *Uncertainty in Artificial Intelligence*, pages 129–139. PMLR.

- Min, Y., He, J., Wang, T., and Gu, Q. (2022a). Learning stochastic shortest path with linear function approximation. In *International Conference on Machine Learning*, pages 15584–15629. PMLR.
- Min, Y., Wang, T., Xu, R., Wang, Z., Jordan, M., and Yang, Z. (2022b). Learn to match with no regret: Reinforcement learning in markov matching markets. In *Advances in Neural Information Processing Systems*.
- Min, Y., Wang, T., Zhou, D., and Gu, Q. (2021b). Variance-aware off-policy evaluation with linear function approximation. *Advances in neural information processing systems*, 34:7598–7610.
- Pfrommer, S., Anderson, B. G., Piet, J., and Soudjoudi, S. (2023). Asymmetric certified robustness via feature-convex neural networks. *arXiv preprint arXiv:2302.01961*.
- Prescott, E. C. and Mehra, R. (2005). Recursive competitive equilibrium: The case of homogeneous households. In *Theory Of Valuation*, pages 357–371. World Scientific.
- Puterman, M. L. (2014). *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons.
- Riedmiller, M. (2005). Neural fitted q iteration—first experiences with a data efficient neural reinforcement learning method. In *Machine Learning: ECML 2005: 16th European Conference on Machine Learning, Porto, Portugal, October 3-7, 2005. Proceedings 16*, pages 317–328. Springer.
- Schmidt, M., Roux, N., and Bach, F. (2011). Convergence rates of inexact proximal-gradient methods for convex optimization. *Advances in neural information processing systems*, 24.
- Seijo, E. and Sen, B. (2011). Nonparametric least squares estimation of a multivariate convex regression function. *The Annals of Statistics*, 39(3):1633–1657.
- Selten, R. (1990). Bounded rationality. *Journal of Institutional and Theoretical Economics (JITE)/Zeitschrift für die gesamte Staatswissenschaft*, 146(4):649–658.
- Shalev-Shwartz, S. and Singer, Y. (2007). *Online learning: Theory, algorithms, and applications*.
- Shephard, R. W. and Färe, R. (1974). The law of diminishing returns. In *Production Theory: Proceedings of an International Seminar Held at the University at Karlsruhe May–July 1973*, pages 287–318. Springer.
- Smith, J. E. and McCardle, K. F. (2002). Structural properties of stochastic dynamic programs. *Operations Research*, 50(5):796–809.
- Song, G., Xu, R., and Lafferty, J. (2021a). Convergence and alignment of gradient descent with random backpropagation weights. *Advances in Neural Information Processing Systems*, 34:19888–19898.
- Song, Z., Mei, S., and Bai, Y. (2021b). When can we learn general-sum markov games with a large number of players sample-efficiently? *arXiv preprint arXiv:2110.04184*.
- Subramanian, J. and Mahajan, A. (2019). Reinforcement learning in stationary mean-field games. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, pages 251–259.
- Sutton, R. S. and Barto, A. G. (2018). *Reinforcement learning: An introduction*. MIT press.
- Tan, H. Y., Mukherjee, S., Tang, J., and Schönlieb, C.-B. (2022). Data-driven mirror descent with input-convex neural networks. *arXiv preprint arXiv:2206.06733*.
- Wang, Y., Liu, Q., Bai, Y., and Jin, C. (2023). Breaking the curse of multiagency: Provably efficient decentralized multi-agent rl with function approximation. *arXiv preprint arXiv:2302.06606*.
- Xu, R., Chen, L., and Karbasi, A. (2021). Meta learning in the continuous time limit. In *International Conference on Artificial Intelligence and Statistics*, pages 3052–3060. PMLR.
- Zhu, Y. (2020). A convex optimization formulation for multivariate regression. *Advances in Neural Information Processing Systems*, 33:17652–17661.

A CLARIFICATION OF NOTATION

We make a few clarifications on the problem setup of our heterogeneous agent model as well as the related notation that we intentionally omitted in the main text due to space constraints.

A.1 Bellman Operators

We define Bellman operator \mathcal{T}_z^π for any policy π and mean-field term z that maps from any Q-function Q to

$$\begin{aligned} (\mathcal{T}_z^\pi Q)(s, a) &:= r_z(s, a) + \gamma \mathbb{E}[\langle Q(s', \cdot), \pi(\cdot | s') \rangle_{\Gamma(s')} | s, a] \\ &= r_z(s, a) + \gamma (\mathcal{P}_z V_z^\pi)(s, a), \end{aligned}$$

where $(\mathcal{P}_z v)(s, a) := \mathbb{E}[v(s') | s, a]$ and $s' \sim \mathbb{P}_z(\cdot | s, a)$ be the next state for any state-value function v . Further, we also denote \mathcal{T}_z to be the Bellman optimality operator, such that for any value function Q

$$(\mathcal{T}_z Q)(s, a) := r_z(s, a) + \gamma \mathbb{E} \left[\max_{a' \in \Gamma(s')} Q(s', a') \mid s, a \right]$$

under a greedy policy with respect to Q . Note that the value function Q_z^π corresponding to any policy π is the fixed point of \mathcal{T}_z^π , i.e., $Q_z^\pi(s, a) = (\mathcal{T}_z^\pi Q_z^\pi)(s, a)$, and for any mean-field term z , the optimal Q-function $Q_z^*(s)$ is the fixed point of the Bellman optimality operator \mathcal{T}_z , i.e.,

$$Q_z^*(s, a) = (\mathcal{T}_z Q_z^*)(s, a) := r_z(s, a) + \gamma \mathbb{E} \left[\max_{a' \in \Gamma(s')} Q_z^*(s', a') \mid s, a \right]$$

and the optimal state-value function follows from $V_z^*(s) = \max_{a \in \Gamma(s)} Q_z^*(s, a)$.

A.2 Issues With RCE

For any concave value function Q , there may exist multiple actions maximizing $Q(s, \cdot)$ for any $s \in \mathcal{S}$, and this non-uniqueness in optimal policy leads to the non-uniqueness of the aggregate indicator Ξ^* and therefore the non-uniqueness of RCE.

The non-robustness of an equilibrium implies that a small estimation error on the optimal value function can lead to a large estimation error on the corresponding optimal policy.

A.3 Regularized RCE

Through adding an entropy regularization, we have a unique regularized optimal policy given any mean-field term, and for the convenience of notation, we denote the regularized optimal policy for any mean-field term z as

$$\Lambda_1(y) := \left\{ \pi^\dagger : \pi^\dagger = \arg \max_{\pi} \left\{ \int_{\Gamma(s)} \pi(a | s) Q_z^*(s, a) da - \mathcal{H}(\pi(\cdot | s)) \right\} \text{ for all } s \in \mathcal{S} \right\}.$$

For any policy π , we also denote the next-iteration mean-field term under any current mean-field term z with a shorthand

$$\Lambda_2(\pi, z) := \Phi(\Psi(z, \pi)).$$

To conclude the existence and uniqueness of the QR-MFE, we need to show the convergence of the above iteration process starting with any mean-field term $z \in \mathcal{Z}$, and it suffices if $\Lambda_2(\Lambda_1(\cdot))$ is a contractive mapping.

B PROOF OF MAIN RESULTS

B.1 Proof of Theorem 4.6

Proof. Recall that the optimal regularized policy π_z^* of the representative agent under any mean-field term z is defined as

$$\pi_z^* = \Lambda_1(y),$$

and the unique mean-field term induced by the best response of the environment to π_z^* under any mean-field term z is given by

$$\Lambda_2(\pi_z^*, z) := \Phi(\Psi(\pi_z^*, \pi)).$$

Let us further denote $\Lambda(y) := \Lambda_2(\Lambda_1(y), z)$ to be the operator that maps any mean-field term z to its corresponding next-iteration mean-field term through the regularized policy $\Lambda_1(y)$. A policy-environment pair (π^*, z^*) is a QR-MFE if and only if $z^* = \Lambda(z^*)$ and $\pi^* = \Lambda_1(z^*)$. Similar to the argument in [Guo et al. \(2020\)](#), note that for any mean-field terms z and z' , the distance between their corresponding next-iteration mean-field terms is controlled by

$$\begin{aligned} \|\Lambda(y) - \Lambda(z')\|_1 &= \|\Lambda_2(\Lambda_1(y), z) - \Lambda_2(\Lambda_1(z'), z')\|_1 \\ &\leq \|\Lambda_2(\Lambda_1(y), z) - \Lambda_2(\Lambda_1(z'), z)\|_1 + \|\Lambda_2(\Lambda_1(z'), z) - \Lambda_2(\Lambda_1(z'), z')\|_1 \\ &\leq \|\Phi(\Psi(\pi_z^*, z)) - \Phi(\Psi(\pi_{z'}^*, z))\|_1 + \|\Phi(\Psi(\pi_{z'}^*, z)) - \Phi(\Psi(\pi_{z'}^*, z'))\|_1, \end{aligned}$$

where the first equality follows from the definition of operator $\Lambda(\cdot)$, the first inequality follows from triangular inequality, and the last inequality follows from the definition of $\Lambda_1(\cdot)$ and $\Lambda_2(\cdot, \cdot)$. Following [Assumption 4.4](#), the Lipschitzness of the mappings Ψ and Φ yields

$$\|\Phi(\Psi(\pi_{z'}^*, z)) - \Phi(\Psi(\pi_{z'}^*, z'))\|_1 \leq L_F \|z - z'\|_1$$

and

$$\|\Phi(\Psi(\pi_z^*, z)) - \Phi(\Psi(\pi_{z'}^*, z))\|_1 \leq L_F \cdot \|\pi_z^* - \pi_{z'}^*\|_{\infty, \bar{\nu}}.$$

Further, it also holds for the optimal Q-functions Q_z^* and $Q_{z'}^*$, that the distance between their corresponding optimal policies under entropy regularization \mathcal{H} is bounded through the distance between the mean-field terms, *i.e.*,

$$\begin{aligned} \|\pi_z^* - \pi_{z'}^*\|_{\infty, \bar{\nu}} &\leq D \cdot \|\pi_z^* - \pi_{z'}^*\|_{\infty, \bar{\nu}} \\ &\leq \frac{D^2}{\zeta} \|Q_z^* - Q_{z'}^*\|_{1, \bar{\mu}} \\ &\leq \frac{JD^2}{\zeta} \|z - z'\|_1, \end{aligned}$$

where the first inequality is due to the concentrability of the sampling distribution $\tilde{\nu}$ in [Assumption 4.5](#), the second inequality follows by applying [Lemma B.2](#), and the last inequality is due to [Lemma B.1](#). Hence, we combine the above inequalities together to get

$$\begin{aligned} \|\Lambda(y) - \Lambda(z')\|_1 &\leq L_F \|z - z'\|_1 + L_F \cdot \|\pi_z^* - \pi_{z'}^*\|_{\infty, \bar{\nu}} \\ &\leq \left(\frac{JL_FD^2}{\zeta} + L_F \right) \|z - z'\|_1, \end{aligned}$$

and the operator $\Lambda(\cdot)$ forms a contraction if $\kappa := \frac{JL_FD^2}{\zeta} + L_F < 1$. The existence and uniqueness of the QR-MFE (π^*, z^*) follows from Banach fixed-point theorem on compact metric space \mathcal{Z} and the uniqueness of global optimizer on strongly convex optimization. \square

Lemma B.1. *Let $z, z' \in \mathcal{Z}$ be two mean-field terms and Q_z^* and $Q_{z'}^*$, to be the optimal value functions under the MDPs with transition kernels \mathbb{P}_z and $\mathbb{P}_{z'}$, respectively. Define $B := 1/(1 - \gamma)$. Then under [Assumptions Assumption 4.3](#), and [4.5](#), we have*

$$\|Q_z^* - Q_{z'}^*\|_{1, \bar{\mu}} \leq J \cdot \|z - z'\|_1,$$

where

$$J := \frac{L_r + \gamma BL_{\mathbb{P}}}{1 - \gamma D}.$$

Proof. Notice that for any optimal Q-functions Q_z^* and $Q_{z'}^*$, it follows from the Bellman optimality equation that

$$\begin{aligned} \|Q_z^* - Q_{z'}^*\|_{1,\tilde{\mu}} &= \|\mathcal{T}_z Q_z^* - \mathcal{T}_{z'} Q_{z'}^*\|_{1,\tilde{\mu}} \\ &= \left\| r_z(s, a) + \gamma \int_{\mathcal{S}} \mathbb{P}_z(s' | s, a) \max_{a \in \Gamma(s')} Q_z^*(s', a) ds' \right. \\ &\quad \left. - r_{z'}(s, a) - \gamma \int_{\mathcal{S}} \mathbb{P}_{z'}(s' | s, a) \max_{a \in \Gamma(s')} Q_{z'}^*(s', a) ds' \right\|_{1,\tilde{\mu}}, \end{aligned}$$

where the first equality is due to the fact that optimal Q-functions Q_z^* and $Q_{z'}^*$ are the fixed points under the Bellman operators \mathcal{T}_z and $\mathcal{T}_{z'}$ respectively, the second equality expands the Bellman operators following the definition of \mathcal{T}_z and $\mathcal{T}_{z'}$. Further, we decompose the difference into three pieces, *i.e.*,

$$\begin{aligned} \|Q_z^* - Q_{z'}^*\|_{1,\tilde{\mu}} &\leq \gamma \left\| \int_{\mathcal{S}} \mathbb{P}_z(s' | s, a) \left[\max_{a \in \Gamma(s')} \{Q_z^*(s', a)\} - \max_{a \in \Gamma(s')} \{Q_{z'}^*(s', a)\} \right] ds' \right\|_{1,\tilde{\mu}} \\ &\quad + \gamma \left\| \int_{\mathcal{S}} [\mathbb{P}_z(s' | s, a) - \mathbb{P}_{z'}(s' | s, a)] \max_{a \in \Gamma(s')} Q_{z'}^*(s', a) ds' \right\|_{1,\tilde{\mu}} \\ &\quad + \|r_z(s, a) - r_{z'}(s, a)\|_{1,\tilde{\mu}}, \end{aligned} \tag{B.1}$$

where the inequality follows from triangular inequality and supplying an intermediate dummy term $\int_{\mathcal{S}} \mathbb{P}_z(s' | s, a) \max_{a \in \Gamma(s')} Q_{z'}^*(s', a) ds'$. Notice that the last term in (B.1) can be bounded following the Lipschitzness of the reward function, *i.e.*,

$$\|r_z(s, a) - r_{z'}(s, a)\|_{1,\tilde{\mu}} \leq L_r \cdot \|z - z'\|_1,$$

and the second term in (B.1) is controlled by the distance between transition kernels \mathbb{P}_z and $\mathbb{P}_{z'}$, *i.e.*,

$$\begin{aligned} \left\| \int_{\mathcal{S}} [\mathbb{P}_z(s' | s, a) - \mathbb{P}_{z'}(s' | s, a)] \max_{a \in \Gamma(s')} Q_{z'}^*(s', a) ds' \right\|_{1,\tilde{\mu}} &\leq L_{\mathbb{P}} \cdot \|Q^*\|_{\infty} \|z - z'\|_1 \\ &\leq BL_{\mathbb{P}} \|z - z'\|_1, \end{aligned}$$

where the first inequality is due to the Lipschitz condition on \mathbb{P}_z with respect to z in Assumption 4.3 and $\max_{a \in \Gamma(s')} Q_{z'}^*(s', a) \leq \|Q^*\|_{\infty} \leq B$ for all feasible state-action pair (s, a) . The second inequality follows from the assumption that the reward $r_h(s, a) \in [0, 1]$ for any feasible state-action pair (s, a) and step $h \in \mathbb{Z}_+$. Finally, for any optimal Q-functions Q_z^* and $Q_{z'}^*$, the first term in (B.1) is bounded by

$$\begin{aligned} &\left\| \int_{\mathcal{S}} \mathbb{P}_z(s' | s, a) \left[\max_{a' \in \Gamma(s')} \{Q_z^*(s', a')\} - \max_{a' \in \Gamma(s')} \{Q_{z'}^*(s', a')\} \right] ds' \right\|_{1,\tilde{\mu}} \\ &\leq \left\| \int_{\mathcal{S}} \mathbb{P}_z(s' | s, a) \max_{a' \in \Gamma(s')} |Q_z^*(s', a') - Q_{z'}^*(s', a')| ds' \right\|_{1,\tilde{\mu}} \\ &= \left\| \max_{a \in \Gamma(s)} |Q_z^*(s, a) - Q_{z'}^*(s, a)| \right\|_{1,\mathbb{P}_z \tilde{\mu}} \\ &= \|Q_z^*(s, a) - Q_{z'}^*(s, a)\|_{1,\mathbb{P}_z \tilde{\mu} \times \pi^\dagger(Q_z^*, Q_{z'}^*)} \end{aligned}$$

where the inequality follows from the fact that $|\max_{x \in \mathcal{X}} f(x) - \max_{y \in \mathcal{X}} g(y)| \leq \max_{x \in \mathcal{X}} |f(x) - g(x)|$ for any x, y in the shared domain \mathcal{X} , the first equality is due to the shorthand $s' \sim \mathbb{P}_z \mu$ to denote $s' \sim \mathbb{P}_z(\cdot | s, a)$ and $(s, a) \sim \mu$, and the last equality is due to defining the policy $\pi^\dagger(Q_z^*, Q_{z'}^*)$ that takes the action to maximize the difference between Q-functions Q_z^* and $Q_{z'}^*$, *i.e.*, $\pi^\dagger(Q_z^*, Q_{z'}^*)(s) := \arg \max_{a \in \mathcal{A}} |Q_z^*(s, a) - Q_{z'}^*(s, a)|$. Putting the bounds above together, we have an upper bound

$$\|Q_z^* - Q_{z'}^*\|_{1,\tilde{\mu}} \leq \gamma D \cdot \|Q_z^* - Q_{z'}^*\|_{1,\tilde{\mu}} + \gamma BL_{\mathbb{P}} \cdot \|z - z'\|_1 + L_r \cdot \|z - z'\|_1, \tag{B.2}$$

where the inequality follows from Assumption 4.5. Rearranging the term in (B.2) and following the definition of J to get

$$\|Q_z^* - Q_{z'}^*\|_{1,\tilde{\mu}} \leq \frac{L_r + \gamma BL_{\mathbb{P}}}{1 - \gamma D} \|z - z'\|_1 = J \cdot \|z - z'\|_1.$$

□

Lemma B.2. Let Q_z^* and $Q_{z'}^*$ be two optimal value functions under the MDPs with transition kernels \mathbb{P}_z and $\mathbb{P}_{z'}$, respectively. Then, under Assumption 4.5, we have

$$\|\pi_z^*(\cdot | \cdot) - \pi_{z'}^*(\cdot | \cdot)\|_{\infty, \bar{\nu}} \leq \frac{D}{\zeta} \cdot \|Q_z^* - Q_{z'}^*\|_{1, \bar{\mu}}.$$

Proof. For any mean-field term z and its corresponding optimal value function Q_z^* , the optimal regularized policy π_z^* can be written as

$$\begin{aligned} \pi_z^*(\cdot | s) &= \arg \max_{u \in \Delta_{\Gamma(s)}} \left\{ \int_{\Gamma(s)} u(a) Q_z^*(s, a) da - \mathcal{H}(u) \right\} \\ &= \arg \max_{u \in \Delta_{\Gamma(s)}} \{ \langle u, Q_z^*(s, \cdot) \rangle - \mathcal{H}(u) \} \\ &= \nabla \mathcal{H}^*(Q(s, \cdot)), \end{aligned}$$

where the first equality follows from the definition of the optimal regularized policy π_z^* , and the last equality follows from Lemma C.1 with $\nabla \mathcal{H}^*(Q(s, \cdot))$ being the functional derivative of the dual functional \mathcal{H}^* evaluated at $Q(s, \cdot)$ for $s \in \mathcal{S}$. Similarly, we write the optimal regularized policy $\pi_{z'}^*(\cdot | s) = \nabla \mathcal{H}^*(Q_{z'}^*(s, \cdot))$ for any $s \in \mathcal{S}$, and it follows that the distance between optimal regularized policies is controlled by that between optimal value functions, *i.e.*,

$$\begin{aligned} \|\pi_z^*(\cdot | \cdot) - \pi_{z'}^*(\cdot | \cdot)\|_{\infty, \bar{\nu}} &= \|\nabla \mathcal{H}^*(Q_z^*(s, \cdot)) - \nabla \mathcal{H}^*(Q_{z'}^*(s, \cdot))\|_{\infty, \bar{\nu}} \\ &\leq \frac{1}{\zeta} \|Q_z^*(s, \cdot) - Q_{z'}^*(s, \cdot)\|_{1, \bar{\nu}} \\ &\leq \frac{D}{\zeta} \|Q_z^* - Q_{z'}^*\|_{1, \bar{\mu}}, \end{aligned}$$

where the first inequality is due to the ζ -strong convexity assumption on \mathcal{H} and Lemma C.1, and the last inequality is due to Assumption 4.5 and $\|Q_z^*(s, \cdot) - Q_{z'}^*(s, \cdot)\|_{1, \bar{\nu}} = \|Q_z^* - Q_{z'}^*\|_{1, \bar{\nu} \times \pi_U}$ with π_U being a uniform policy. \square

From the proof of Lemma B.2, we bound the distance between the optimal household policies under different mean-field environments. It provides an abstraction over multi-agent games where exponentially many interactions between the agents are considered. Recall that mean-field setups, viewed as a limit of their multi-agent counterparts (Song et al., 2021b; Dubey and Pentland, 2021; He et al., 2022; Xu et al., 2021; Wang et al., 2023; Ling et al., 2019), drive the number of agents to infinity. Compared to the multi-agent setting, in our analysis we bound the distance between the optimal household policies under different mean-field environments.

B.2 Proof of Theorem 4.7

Proof. For each $t \in [T]$, the mean-field term z^t is determined by the mean-field term z^{t-1} from last iteration of Algorithm 1 and the regularized policy $\pi_{\widehat{Q}^{t-1}}$ induced by the estimator \widehat{Q}^{t-1} of the optimal value function $Q_{z^{t-1}}^*$. Hence, we follow the definition of operator Λ_1 and Λ_2 , and the distance between z^t and the regularized RCE z^* is bounded by

$$\begin{aligned} \|z^t - z^*\|_1 &= \|\Lambda_2(\pi_{\widehat{Q}^{t-1}}, z^{t-1}) - \Lambda_2(\Lambda_1(z^*), z^*)\|_1 \\ &\leq \|\Lambda_2(\Lambda_1(z^{t-1}), z^{t-1}) - \Lambda_2(\Lambda_1(z^*), z^*)\|_1 \\ &\quad + \|\Lambda_2(\Lambda_1(z^{t-1}), z^{t-1}) - \Lambda_2(\pi_{\widehat{Q}^{t-1}}, z^{t-1})\|_1 \\ &\leq \left(\frac{JL_F D^2}{\zeta} + L_F \right) \|z^{t-1} - z^*\|_1 + L_F \|\pi_{Q_{z^{t-1}}^*} - \pi_{\widehat{Q}^{t-1}}\|_{\infty, \bar{\nu}}, \end{aligned}$$

where the equality follows from the definition of Λ_1 and Λ_2 as well as Definition 2.1, the first inequality follows from triangular inequality, and the second inequality is due to Theorem 4.6 and Assumption 4.4. We denote $\kappa := JL_F D^2 / \zeta + L_F$ for the convenience of notation. Following Lemma B.2, we further have

$$\begin{aligned} \|z^t - z^*\|_1 &\leq \left(\frac{JL_F D^2}{\zeta} + L_F \right) \|z^{t-1} - z^*\|_1 + \frac{L_F D^2}{\zeta} \|\widehat{Q}^{t-1} - Q_{z^{t-1}}^*\|_{1, \bar{\mu}} \\ &\leq \kappa \cdot \|z^{t-1} - z^*\|_1 + \frac{L_F D^2}{\zeta} \|\widehat{Q}^{t-1} - Q_{z^{t-1}}^*\|_{\bar{\mu}}. \end{aligned} \tag{B.3}$$

After expanding the the first term of (B.3) recursively for $t = T$, it holds that for any $T \in \mathbb{Z}_+$, the distance between the output mean-field term z^T of Algorithm 1 and z^* is controlled by

$$\begin{aligned} \|z^T - z^*\|_1 &\leq \kappa^T \|z^0 - z^*\|_1 + \frac{L_F D^2}{\zeta} \sum_{t=1}^{T-1} \kappa^{T-1-t} \|\widehat{Q}^t - Q_{z^t}^*\|_{\bar{\mu}} + \kappa^{T-1} \|\widehat{Q}^0 - Q_{p^0}^*\|_{\bar{\mu}} \\ &\leq \kappa^T \|z^0 - z^*\|_1 + \frac{L_F D^2}{\zeta} \sup_{t \in [T]} \|\widehat{Q}^t - Q_{z^t}^*\|_{\bar{\mu}} \sum_{t=1}^{T-1} \kappa^{T-1-t} + \kappa^{T-1} \|\widehat{Q}^0 - Q_{p^0}^*\|_{\bar{\mu}} \\ &\leq \kappa^T \|z^0 - z^*\|_1 + \frac{L_F D^2}{(1-\kappa)\zeta} \sup_{t \in [T]} \|\widehat{Q}^t - Q_{z^t}^*\|_{\bar{\mu}} + \frac{\kappa^{T-1}}{1-\gamma}, \end{aligned}$$

where the first inequality follows from recursive expansion of (B.3), the second inequality is due to uniform bounding of the estimation error $\|\widehat{Q}^t - Q_{z^t}^*\|_{\bar{\mu}} \leq \sup_{t \in [T]} \|\widehat{Q}^t - Q_{z^t}^*\|_{\bar{\mu}}$ for all $t \in [T]$, and the last inequality is due to the bounded optimal value function $\|Q_z^*\|_\infty \leq \frac{1}{1-\gamma} = B$ for all $z \in \mathcal{Z}$. Therefore, for regression on the set $\mathcal{C}_{B,L}$ of concave functions, we have with probability $1 - \delta$ that

$$\|z^T - z^*\|_1 \lesssim \frac{L_F D^2 \sqrt{D}}{(1-\gamma)(1-\kappa)\zeta} \left[B^{\frac{d+8}{d+2}} (d+1)^{\frac{2}{d+2}} (dL)^{\frac{d}{d+2}} M^{-\frac{2}{d+2}} \log M + \epsilon \right] + \kappa^T Z + \frac{\kappa^{T-1}}{1-\gamma}$$

for sample size $M^{\frac{d}{d+2}} \gtrsim \log M$, $M^{\frac{d}{d+2}} \log M \gtrsim B^{\frac{3d}{d+2}} (d+1)^{\frac{-2}{d+2}} (dL)^{\frac{-d}{d+2}} \log \frac{\tau T}{\delta}$, and

$$(\log M)^{\frac{d}{2}} \gtrsim 1 + \log(R_d^* B^{\frac{6}{d+2}} (d+1)^{\frac{-2}{d+2}} (dL)^{\frac{-d}{d+2}}),$$

where $R_d^* \leq \max\{8dL, 2B + 4dL\}$ and

$$\tau \asymp \log \left(D^{-1/2} B^{-\frac{d+8}{d+2}} (d+1)^{-\frac{2}{d+2}} (dL)^{-\frac{d}{d+2}} M^{\frac{2}{d+2}} \frac{1}{\log M} \right) / \log \frac{1}{\gamma}.$$

The inequality follows from applying Lemma B.9 with a union bound. \square

B.3 Proof of Theorem 4.8

Combine the statements of Theorems B.3 and B.4 to get Theorem 4.8.

Theorem B.3. Let $B := 1/(1-\gamma)$, $L := L_r/(1-\gamma)$. Assume that $\kappa := \frac{J L_F D^2}{\zeta} + L_F < 1$, then with probability at least $1 - \delta$, Algorithm 1 with T iterations on $\mathcal{A}_{B,L}^K$ gives mean-field term z^T such that

$$\|z^T - z^*\|_1 \lesssim \frac{L_F D^2 \sqrt{D}}{(1-\gamma)(1-\kappa)\zeta} \left[((d+1)B^4 + (B+dL)dL) M^{-\frac{4}{d+4}} \log M + \epsilon \right] + \kappa^T Z + \frac{\kappa^{T-1}}{1-\gamma}$$

for sample size $M \gtrsim \frac{1}{\delta \log \frac{1}{\gamma}} \log \frac{M^{4/(d+4)}}{d^2 L^2 \sqrt{D}}$, $M^{\frac{2d+4}{d+4}} \log M \gtrsim \frac{B^3(B+4dL)}{d+1}$, $M \gtrsim \frac{\tau T}{\exp(2(d+1)M^{d/(d+4)})\delta}$, and $M^{\frac{d}{d+4}} \gtrsim \frac{B^2}{d^2 L^2} \log \frac{2\tau T}{\delta}$, where $\tau \asymp \log \frac{M^{4/(d+4)}}{d^2 L^2 \sqrt{D}} / \log \frac{1}{\gamma}$.

Proof. Following similar argument as in the proof of Theorem 4.7, the distance between the output mean-field term z^T of Algorithm 1 and z^* is controlled by

$$\begin{aligned} \|z^T - z^*\|_1 &\leq \kappa^T \|z^0 - z^*\|_1 + \frac{L_F D^2}{\zeta} \sum_{t=1}^{T-1} \kappa^{T-1-t} \|\widehat{Q}^t - Q_{z^t}^*\|_{\bar{\mu}} + \kappa^{T-1} \|\widehat{Q}^0 - Q_{p^0}^*\|_{\bar{\mu}} \\ &\leq \kappa^T \|z^0 - z^*\|_1 + \frac{L_F D^2}{\zeta} \sup_{t \in [T]} \|\widehat{Q}^t - Q_{z^t}^*\|_{\bar{\mu}} \sum_{t=1}^{T-1} \kappa^{T-1-t} + \kappa^{T-1} \|\widehat{Q}^0 - Q_{p^0}^*\|_{\bar{\mu}} \\ &\leq \kappa^T \|z^0 - z^*\|_1 + \frac{L_F D^2}{(1-\kappa)\zeta} \sup_{t \in [T]} \|\widehat{Q}^t - Q_{z^t}^*\|_{\bar{\mu}} + \frac{\kappa^{T-1}}{1-\gamma}, \end{aligned}$$

where the first inequality follows from recursive expansion of (B.3), the second inequality is due to uniform bounding of the estimation error $\|\widehat{Q}^t - Q_{z^t}^*\|_{\bar{\mu}} \leq \sup_{t \in [T]} \|\widehat{Q}^t - Q_{z^t}^*\|_{\bar{\mu}}$ for all $t \in [T]$, and the last inequality is due to the bounded optimal value function $\|Q_z^*\|_{\infty} \leq B$ for all $z \in \mathcal{Z}$. For regression on max-affine functions $\mathcal{A}_{B,L}^K$, we have with probability $1 - \delta$ that

$$\|z^T - z^*\|_1 \lesssim \frac{L_F D^2 \sqrt{D}}{(1-\gamma)(1-\kappa)\zeta} \left[((d+1)B^4 + (B+dL)dL) M^{-\frac{4}{d+4}} \log M + \epsilon \right] + \kappa^T Z + \frac{\kappa^{T-1}}{1-\gamma}$$

for

$$\begin{aligned} M &\gtrsim \frac{1}{\delta \log \frac{1}{\gamma}} \log \frac{M^{4/(d+4)}}{d^2 L^2 \sqrt{D}}, & M^{\frac{2d+4}{d+4}} \log M &\gtrsim \frac{B^3(B+4dL)}{d+1}, \\ M &\gtrsim \frac{\tau T}{\exp(2(d+1)M^{d/(d+4)})\delta}, & M^{\frac{d}{d+4}} &\gtrsim \frac{B^2}{d^2 L^2} \log \frac{2\tau T}{\delta}, \end{aligned}$$

where

$$\tau \asymp \log \frac{M^{4/(d+4)}}{d^2 L^2 \sqrt{D}} / \log \frac{1}{\gamma}.$$

The inequality follows from applying Lemma B.9 with a union bound. \square

Theorem B.4. Let $B := 1/(1-\gamma)$, $L := L_r/(1-\gamma)$. Assume that $\kappa := \frac{J L_F D^2}{\zeta} + L_F < 1$, then with probability at least $1 - \delta$, Algorithm 1 with T iterations on $\mathcal{N}_{B,L}^K$ gives mean-field term z^T such that

$$\|z^T - z^*\|_1 \lesssim \frac{L_F D^2 \sqrt{D}}{(1-\gamma)(1-\kappa)\zeta} \left[((d+1)B^4 + (B+dL)dL) M^{-\frac{4}{d+4}} \log M + \epsilon \right] + \kappa^T Z + \frac{\kappa^{T-1}}{1-\gamma}$$

for sample size $M \gtrsim \frac{1}{\delta \log \frac{1}{\gamma}} \log \frac{M^{4/(d+4)}}{d^2 L^2 \sqrt{D}}$, $M^{\frac{2d+4}{d+4}} \log M \gtrsim \frac{B^3(B+4dL)}{d+1}$, $M \gtrsim \frac{\tau T}{\exp(2(d+1)M^{d/(d+4)})\delta}$, and $M^{\frac{d}{d+4}} \gtrsim \frac{B^2}{d^2 L^2} \log \frac{2\tau T}{\delta}$, where $\tau \asymp \log \frac{M^{4/(d+4)}}{d^2 L^2 \sqrt{D}} / \log \frac{1}{\gamma}$.

Proof. Following a similar argument as in the proof of Theorem 4.7, the distance between the output mean-field term z^T of Algorithm 1 and z^* is controlled by

$$\begin{aligned} \|z^T - z^*\|_1 &\leq \kappa^T \|z^0 - z^*\|_1 + \frac{L_F D^2}{\zeta} \sum_{t=1}^{T-1} \kappa^{T-1-t} \|\widehat{Q}^t - Q_{z^t}^*\|_{\bar{\mu}} + \kappa^{T-1} \|\widehat{Q}^0 - Q_{p^0}^*\|_{\bar{\mu}} \\ &\leq \kappa^T \|z^0 - z^*\|_1 + \frac{L_F D^2}{\zeta} \sup_{t \in [T]} \|\widehat{Q}^t - Q_{z^t}^*\|_{\bar{\mu}} \sum_{t=1}^{T-1} \kappa^{T-1-t} + \kappa^{T-1} \|\widehat{Q}^0 - Q_{p^0}^*\|_{\bar{\mu}} \\ &\leq \kappa^T \|z^0 - z^*\|_1 + \frac{L_F D^2}{(1-\kappa)\zeta} \sup_{t \in [T]} \|\widehat{Q}^t - Q_{z^t}^*\|_{\bar{\mu}} + \frac{\kappa^{T-1}}{1-\gamma}, \end{aligned}$$

where the first inequality follows from recursive expansion of (B.3), the second inequality is due to uniform bounding of the estimation error $\|\widehat{Q}^t - Q_{z^t}^*\|_{\bar{\mu}} \leq \sup_{t \in [T]} \|\widehat{Q}^t - Q_{z^t}^*\|_{\bar{\mu}}$ for all $t \in [T]$, and the last inequality is due to the bounded optimal value function $\|Q_z^*\|_{\infty} \leq B$ for all $z \in \mathcal{Z}$. For regression on max-affine functions $\mathcal{N}_{B,L}^K$, we have with probability $1 - \delta$ that

$$\|z^T - z^*\|_1 \lesssim \frac{L_F D^2 \sqrt{D}}{(1-\gamma)(1-\kappa)\zeta} \left[((d+1)B^4 + (B+dL)dL) M^{-\frac{4}{d+4}} \log M + \epsilon \right] + \kappa^T Z + \frac{\kappa^{T-1}}{1-\gamma}$$

for

$$\begin{aligned} M &\gtrsim \frac{1}{\delta \log \frac{1}{\gamma}} \log \frac{M^{4/(d+4)}}{d^2 L^2 \sqrt{D}}, & M^{\frac{2d+4}{d+4}} \log M &\gtrsim \frac{B^3(B+4dL)}{d+1}, \\ M &\gtrsim \frac{\tau T}{\exp(2(d+1)M^{d/(d+4)})\delta}, & M^{\frac{d}{d+4}} &\gtrsim \frac{B^2}{d^2 L^2} \log \frac{2\tau T}{\delta}, \end{aligned}$$

where

$$\tau \asymp \log \frac{M^{4/(d+4)}}{d^2 L^2 \sqrt{D}} / \log \frac{1}{\gamma}.$$

The inequality follows from applying Lemma B.9 with a union bound. \square

In order to apply concave regression in estimating the optimal value function Q_z^* under the MDP with transition kernel \mathbb{P}_z , it is required that for any function $Q \in \mathcal{C}_{B,L}$, the Bellman optimality operator \mathcal{T}_z preserves the concavity, i.e., $\mathcal{T}_z Q \in \mathcal{C}_{B,L}$. As we show in Lemmas B.5 and B.6, this holds for the Bellman optimality operator \mathcal{T}_z and B -bounded and L -Lipschitz concave function set $\mathcal{C}_{B,L}$ for $B = 1/(1 - \gamma)$ and $L = L_r/(1 - \gamma)$.

Lemma B.5. *Let $B := 1/(1 - \gamma)$ and $L := L_r/(1 - \gamma)$. For any mean-field term $z \in \mathcal{Z}$ and the MDP with stochastic concave transition kernel \mathbb{P}_z , if any value function $Q \in \mathcal{C}_{B,L}$, then $\mathcal{T}_z Q \in \mathcal{C}_{B,L}$. Consequently, the optimal value function $Q_z^* \in \mathcal{C}_{B,L}$.*

Proof. For any bounded Lipschitz value function $Q \in \mathcal{C}_{B,L}$, we fix $b, b' \in \mathcal{B}$, $w, w' \in \mathcal{W}$, $a \in \Gamma_z(b, w)$, and $a' \in \Gamma_z(b', w')$. For any $\lambda \in [0, 1]$, we define $\bar{b} := (1 - \lambda)b + \lambda b'$, $\bar{w} := (1 - \lambda)w + \lambda w'$, and $\bar{a} := (1 - \lambda)a + \lambda a'$. Note that the state-action pair $(\bar{b}, \bar{w}, \bar{a})$ lies in the feasible set \mathcal{G} due to the concavity assumption. It then follows that under Bellman optimality operator \mathcal{T}_z that

$$\begin{aligned} (1 - \lambda)(\mathcal{T}_z Q)(b, w, a) + \lambda(\mathcal{T}_z Q)(b', w', a') &= (1 - \lambda)r_z(b, w, a) + \lambda r_z(b', w', a') \\ &\quad + \gamma(1 - \lambda) \int_{\mathcal{W}} \mathbb{P}_z(\tilde{w} | w) Q(a, \tilde{w}, a_{1,\tilde{w}}) d\tilde{w} \\ &\quad + \gamma\lambda \int_{\mathcal{W}} \mathbb{P}_z(\tilde{w} | w') Q(a', \tilde{w}, a_{2,\tilde{w}}) d\tilde{w}, \end{aligned}$$

where we define $a_{1,\tilde{w}} := \arg \max_{a^\dagger \in \Gamma(a, \tilde{w})} Q(a, \tilde{w}, a^\dagger)$ to be the optimal action under state (a, \tilde{w}) and $a_{2,\tilde{w}} := \arg \max_{a^\dagger \in \Gamma(a', \tilde{w})} Q(a', \tilde{w}, a^\dagger)$ to be the optimal action under state (a', \tilde{w}) for any $\tilde{w} \in \mathcal{W}$. Given the optimal actions $a_{1,\tilde{w}}$ and $a_{2,\tilde{w}}$, we further denote $\hat{a} := (1 - \lambda)a_{1,\tilde{w}} + \lambda a_{2,\tilde{w}}$. Notice that

$$(1 - \lambda)r_z(b, w, a) + \lambda r_z(b', w', a') \leq r_z(\bar{b}, \bar{w}, \bar{a})$$

due to the concavity of the reward function r_z . It also follows from the stochastic concavity of the transition kernel \mathbb{P}_z that

$$\begin{aligned} (1 - \lambda) \int_{\mathcal{W}} \mathbb{P}_z(\tilde{w} | w) Q(a, \tilde{w}, a_{1,\tilde{w}}) d\tilde{w} + \lambda \int_{\mathcal{W}} \mathbb{P}_z(\tilde{w} | w') Q(a', \tilde{w}, a_{2,\tilde{w}}) d\tilde{w} \\ \leq \int_{\mathcal{W}} \mathbb{P}_z(\tilde{w} | \bar{w}) Q(\bar{a}, \tilde{w}, \hat{a}) d\tilde{w}. \end{aligned}$$

Hence, we have the concavity of $\mathcal{T}_z Q$ following

$$\begin{aligned} (1 - \lambda)(\mathcal{T}_z Q)(b, w, a) + \lambda(\mathcal{T}_z Q)(b', w', a') &\leq r_z(\bar{b}, \bar{w}, \bar{a}) + \gamma \int_{\mathcal{W}} \mathbb{P}_z(\tilde{w} | w) Q(\bar{a}, \tilde{w}, \hat{a}) d\tilde{w} \\ &\leq r_z(\bar{b}, \bar{w}, \bar{a}) + \gamma \int_{\mathcal{W}} \mathbb{P}_z(\tilde{w} | w) Q(\bar{a}, \tilde{w}, a^\dagger) d\tilde{w} \\ &= (\mathcal{T}_z Q)(\bar{b}, \bar{w}, \bar{a}), \end{aligned}$$

where the second inequality is due to the greedy action $Q(\bar{a}, \tilde{w}, \hat{a}) \leq \max_{a^\dagger} Q(\bar{a}, \tilde{w}, a^\dagger)$, and the last equality is due to $a^\dagger := \max_{a^\dagger \in \Gamma(\bar{a}, \tilde{w})} Q(\bar{a}, \tilde{w}, a^\dagger)$. The Lipschitzness of $\mathcal{T}_z Q$ follows from $L = L_r/(1 - \gamma)$.

Recall that Q_z^* is the fixed point of \mathcal{T}_z , and our desired conclusion follows from the closeness of $\mathcal{C}_{B,L}$ and the operation \mathcal{T}_z that preserves concavity and Lipschitzness of the value function. \square

Lemma B.6. *Let $B := 1/(1 - \gamma)$ and $L := L_r/(1 - \gamma)$. For any MDP with transition kernel \mathbb{P}_z and any fixed income variable $w \in \mathcal{W}$, if any value function $Q(\cdot, w, \cdot) \in \mathcal{C}_{B,L}$, then $\mathcal{T}_z Q(\cdot, w, \cdot) \in \mathcal{C}_{B,L}$. Consequently, $Q_z^*(\cdot, w, \cdot) \in \mathcal{C}_{B,L}$ for any $w \in \mathcal{W}$.*

Proof. For any bounded Lipschitz value function $Q(\cdot, w, \cdot) \in \mathcal{C}_{B,L}$ with fixed $w \in \mathcal{W}$, we pick any $b, b' \in \mathcal{B}$, $a \in \Gamma_z(b, w)$, and $a' \in \Gamma_z(b', w)$. For any $\lambda \in [0, 1]$, we define $\bar{b} := (1 - \lambda)b + \lambda b'$, and $\bar{a} := (1 - \lambda)a + \lambda a'$. Note that the state-action pair (\bar{b}, w, \bar{a}) lies in the feasible set \mathcal{G} due to the concavity assumption. It then follows that under Bellman optimality operator \mathcal{T}_z

$$\begin{aligned} (1 - \lambda)(\mathcal{T}_z Q)(b, w, a) + \lambda(\mathcal{T}_z Q)(b', w, a') &= (1 - \lambda)r_z(b, w, a) + \lambda r_z(b', w, a') \\ &\quad + \gamma(1 - \lambda) \int_{\mathcal{W}} \mathbb{P}_z(\tilde{w} | w) Q(a, \tilde{w}, a_{1,\tilde{w}}) d\tilde{w} \\ &\quad + \gamma\lambda \int_{\mathcal{W}} \mathbb{P}_z(\tilde{w} | w) Q(a', \tilde{w}, a_{2,\tilde{w}}) d\tilde{w}, \end{aligned}$$

where we define $a_{1,\tilde{w}} := \arg \max_{a^* \in \Gamma(a, \tilde{w})} Q(a, \tilde{w}, a^*)$ to be the optimal action under state (a, \tilde{w}) and $a_{2,\tilde{w}} := \arg \max_{a^* \in \Gamma(a', \tilde{w})} Q(a', \tilde{w}, a^*)$ to be the optimal action under state (a', \tilde{w}) for any $\tilde{w} \in \mathcal{W}$. Given the optimal actions $a_{1,\tilde{w}}$ and $a_{2,\tilde{w}}$, we further denote $\hat{a}_{\tilde{w}} := (1 - \lambda)a_{1,\tilde{w}} + \lambda a_{2,\tilde{w}}$. Notice that

$$(1 - \lambda)r_z(b, w, a) + \lambda r_z(b', w, a') \leq r_z(\bar{b}, w, \bar{a})$$

due to the concavity of the reward function r_z . It also holds that

$$\begin{aligned} (1 - \lambda) \int_{\mathcal{W}} \mathbb{P}_z(\tilde{w} | w) Q(a, \tilde{w}, a_{1,\tilde{w}}) d\tilde{w} + \lambda \int_{\mathcal{W}} \mathbb{P}_z(\tilde{w} | w) Q(a', \tilde{w}, a_{2,\tilde{w}}) d\tilde{w} \\ = \int_{\mathcal{W}} \mathbb{P}_z(\tilde{w} | w) [(1 - \lambda)Q(a, \tilde{w}, a_{1,\tilde{w}}) + \lambda Q(a', \tilde{w}, a_{2,\tilde{w}})] d\tilde{w} \\ \leq \int_{\mathcal{W}} \mathbb{P}_z(\tilde{w} | w) Q(\bar{a}, \tilde{w}, \hat{a}_{\tilde{w}}) d\tilde{w}, \end{aligned}$$

where the inequality is due to the concavity of the value function Q . Hence, we have the concavity of $\mathcal{T}_z Q$ following

$$\begin{aligned} (\mathcal{T}_z Q)(\bar{b}, w, \bar{a}) &= r_z(\bar{b}, w, \bar{a}) + \gamma \int_{\mathcal{W}} \mathbb{P}_z(\tilde{w} | w) Q(\bar{a}, \tilde{w}, a^\dagger) d\tilde{w} \\ &\geq r_z(\bar{b}, w, \bar{a}) + \gamma \int_{\mathcal{W}} \mathbb{P}_z(\tilde{w} | w) Q(\bar{a}, \tilde{w}, \hat{a}_{\tilde{w}}) d\tilde{w}, \end{aligned}$$

where the equality is due to the greedy action $a^\dagger := \max_{a \in \Gamma(\bar{a}, \tilde{w})} Q(\bar{a}, \tilde{w}, a)$, and the inequality is due to $Q(\bar{a}, \tilde{w}, \hat{a}) \leq Q(\bar{a}, \tilde{w}, a^\dagger)$.

Recall that $r_z(b, w, a) \in [0, 1]$ for all $(b, w, a) \in \mathcal{G}$. For any value function Q such that $\|Q\|_\infty \leq \frac{1}{1-\gamma} = B$, we have $\|\mathcal{T}_z Q\|_\infty \leq 1 + \frac{\gamma}{1-\gamma} = \frac{1}{1-\gamma} = B$ following the Bellman equation. The Lipschitzness of $\mathcal{T}_z Q$ similarly follows from $L = L_r/(1 - \gamma)$. Hence, we have $(\mathcal{T}_z Q)(\cdot, w, \cdot) \in \mathcal{C}_{B,L}$ if any value function $Q(\cdot, w, \cdot) \in \mathcal{C}_{B,L}$.

Notice that Q_z^* is the fixed point of \mathcal{T}_z , and the operation \mathcal{T}_z preserving concavity and Lipschitzness of any value function Q in $\mathcal{C}_{B,L}$ and the closeness of $\mathcal{C}_{B,L}$ lead to $Q_z^* \in \mathcal{C}_{B,L}$. \square

B.4 Proof of FQI with Concave Regression

FQI is a popular offline RL algorithm, together with its online value iteration counterpart, has been analyzed by a large body of literature under different settings (Bertsekas, 2012; Grunewalder et al., 2012; Fonteneau et al., 2013; Sutton and Barto, 2018; Kumar et al., 2019; Zhu, 2020; Min et al., 2021b; Fei and Xu, 2022b; Min et al., 2022a; Lu et al., 2023). In particular, our method features FQI with concave regression that integrates the economic insights to the algorithm.

We would like to introduce some notation for the proofs of Algorithm 2 that follows. For any function f and dataset $\mathcal{D}_f = \{(X_m, Y_m)\}_{m=1}^M$ where data pairs (X, Y) are i.i.d. sampled from a distribution \mathcal{D} such that $\mathbb{E}_{\mathcal{D}}[Y | X = x] = f(x)$, we define the empirical risk of any estimator f' of the ground truth f as

$$\mathcal{L}_{\mathcal{D}}(f', f) := \frac{1}{M} \sum_{m=1}^M (f'(X_m) - Y_m)^2,$$

where the subscript f on \mathcal{D} is omitted for simplicity. Further, we define the true risk of the estimator f' as

$$\mathcal{L}(f', f) := \mathbb{E}[(f'(X) - Y)^2] = \mathbb{E}_{\mathcal{D}}[\mathcal{L}_{\mathcal{D}}(f', f)],$$

where the expectation is taken with respect to the underlying distribution \mathcal{D} . For any function class \mathcal{F} and estimator $f' \in \mathcal{F}$, we denote

$$\mathcal{E}_{\mathcal{D}}(f', f) := \mathcal{L}_{\mathcal{D}}(f', f) - \inf_{g \in \mathcal{F}} \mathcal{L}_{\mathcal{D}}(g, f)$$

as the sub-optimality of f' compared to the best empirical risk estimator of f in \mathcal{F} . We say f' is an ϵ -approximate LSE if $\mathcal{E}_{\mathcal{D}}(f', f) \leq \epsilon$, and we denote the set of all ϵ -approximate LSEs $f' \in \mathcal{F}$ on data \mathcal{D}_f as $\text{LSE}(\mathcal{F}, \mathcal{D}_f; \epsilon)$. Similarly, we also define

$$\mathcal{E}(f', f) := \mathcal{L}(f', f) - \inf_{g \in \mathcal{F}} \mathcal{L}(g, f)$$

to be the sub-optimality of f' in terms of true risk over distribution \mathcal{D} . In particular, for any offline dataset $\mathcal{D}_{\mathcal{T}Q} := \{(s_m, a_m, r_m, s'_m)\}_{m=1}^M$ i.i.d. sampled from $(s, a) \sim \tilde{\mu}$ under the transition kernel $s' \sim \mathbb{P}(\cdot | s, a)$, we aim to estimate $\mathcal{T}Q$ with concave regression given any value function $Q \in \mathcal{C}_{B,L}$. It follows that

$$\mathcal{L}_{\mathcal{D}}(f', \mathcal{T}Q) = \frac{1}{M} \sum_{m=1}^M (f'(s_m, a_m) - r_m - \gamma V(s'_m))^2$$

and for any function set \mathcal{F} ,

$$\mathcal{E}_{\mathcal{D}}(f', \mathcal{T}Q) := \mathcal{L}_{\mathcal{D}}(f', \mathcal{T}Q) - \inf_{g \in \mathcal{F}} \mathcal{L}_{\mathcal{D}}(g, \mathcal{T}Q),$$

where $V(s'_m) := \max_{a \in \Gamma(s'_m)} Q(s'_m, a)$ is the the action-value function under the greedy policy with respect to Q evaluated at s'_m . Further, we have $\mathcal{L}(f', \mathcal{T}Q) := \mathbb{E}_{\mathcal{D}}[\mathcal{L}_{\mathcal{D}}(f', \mathcal{T}Q)]$. We write $\mathcal{D}_{\mathcal{T}Q}$ as \mathcal{D} for simplicity when the context is clear. Before getting into the proof for regression, we first bound the estimation error of FQI using a similar argument as in [Chen and Jiang \(2019\)](#).

Lemma B.7. *Let Q_z^* be the optimal value function under the MDP with transition kernel \mathbb{P}_z and \tilde{Q}_z^ℓ be the estimation of Q_z^* at the ℓ -th iteration of Algorithm 2. Then for any policy-induced measure μ , it holds that the distance between Q_z^* and the estimator \tilde{Q}_z^τ of Algorithm 2 is bounded by*

$$\|\tilde{Q}_z^\tau - Q_z^*\|_{\mu} \leq \sum_{\ell=1}^{\tau} \gamma^{\tau-\ell} \sqrt{D} \cdot \|\tilde{Q}_z^\ell - \mathcal{T}_z \tilde{Q}_z^{\ell-1}\|_{\mu} + \frac{\gamma^{\tau}}{1-\gamma}.$$

Proof. Under the MDP induced by mean-field term z and policy-induced measure $\mu \in \Delta_{\mathcal{S}}$, it holds for any $\ell \in [\tau]$ that

$$\|\tilde{Q}_z^\ell - Q_z^*\|_{\mu} \leq \|\tilde{Q}_z^\ell - \mathcal{T}_z \tilde{Q}_z^{\ell-1}\|_{\mu} + \|\mathcal{T}_z \tilde{Q}_z^{\ell-1} - Q_z^*\|_{\mu}, \quad (\text{B.4})$$

where $\mathcal{T}_z Q$ for any value function Q denotes the Bellman operator induced by the mean-field term z and the greedy policy with respect to Q . Notice that the first term of (B.4) represents the estimation error of \tilde{Q}_z^ℓ to the target function $\mathcal{T}_z \tilde{Q}_z^{\ell-1}$. For any measure μ , the second term of (B.4) can be further bounded through

$$\begin{aligned} \|\mathcal{T}_z \tilde{Q}_z^{\ell-1} - Q_z^*\|_{\mu}^2 &= \|\mathcal{T}_z \tilde{Q}_z^{\ell-1} - \mathcal{T}_z Q_z^*\|_{\mu}^2 \\ &= \mathbb{E}_{\mu}[(\mathcal{T}_z \tilde{Q}_z^{\ell-1}(s, a) - \mathcal{T}_z Q_z^*(s, a))^2] \\ &= \mathbb{E}_{\mu}[(\gamma \mathbb{P}[\tilde{V}_z^{\ell-1}(s') - V_z^*(s')])^2], \end{aligned}$$

where the first equality is due to the fact that the optimal value function Q_z^* is the fixed point of \mathcal{T}_z , and the second equality follows from expanding the terms $\mathcal{T}_z \tilde{Q}_z^{\ell-1}(s, a) = r_z(s, a) + \gamma \int_{\mathcal{S}} \mathbb{P}_z(s' | s, a) \tilde{V}_z^{\ell-1}(s') ds$ and $\mathcal{T}_z Q_z^*(s, a) = r_z(s, a) + \gamma \int_{\mathcal{S}} \mathbb{P}_z(s' | s, a) V_z^*(s') ds$. Apply Jensen's inequality to the right-hand side of the second inequality, we have

$$\|\mathcal{T}_z \tilde{Q}_z^{\ell-1} - Q_z^*\|_{\mu}^2 \leq \gamma^2 \mathbb{E}_{\mu} \mathbb{P}_z[(\tilde{V}_z^{\ell-1}(s') - V_z^*(s'))^2]$$

$$= \gamma^2 \mathbb{E}_{\mathbb{P}_z \mu} [(\tilde{V}_z^{\ell-1}(s') - V_z^*(s'))^2],$$

where we define a shorthand $s' \sim \mathbb{P}_z \mu$ that implies $s' \sim \mathbb{P}_z(\cdot | s, a)$ and $(s, a) \sim \mu$. In particular, we further expand according to the definitions of $\tilde{V}_z^{\ell-1}(s')$ and $V_z^*(s')$ to get

$$\begin{aligned} \mathbb{E}_{\mathbb{P}_z \mu} [(\tilde{V}_z^{\ell-1}(s') - V_z^*(s'))^2] &= \mathbb{E}_{\mathbb{P}_z \mu} \left[\left(\max_{a \in \mathcal{A}} \tilde{Q}_z^{\ell-1}(s', a) - \max_{a \in \mathcal{A}} Q_z^*(s', a) \right)^2 \right] \\ &\leq \mathbb{E}_{\mathbb{P}_z \mu} \left[\left(\tilde{Q}_z^{\ell-1}(s', \pi^\dagger(\tilde{Q}_z^{\ell-1}, Q_z^*)) - Q_z^*(s', \pi^\dagger(\tilde{Q}_z^{\ell-1}, Q_z^*)) \right)^2 \right] \\ &= \|\tilde{Q}_z^{\ell-1} - Q_z^*\|_{\mathbb{P}_z \mu \times \pi^\dagger(\tilde{Q}_z^{\ell-1}, Q_z^*)}^2, \end{aligned}$$

where the inequality follows from the inequality follows from the fact that $|\max_{x \in \mathcal{X}} f(x) - \max_{y \in \mathcal{X}} g(y)| \leq \max_{x \in \mathcal{X}} |f(x) - g(x)|$ for any x, y in the shared domain \mathcal{X} . Especially, the policy $\pi^\dagger(\tilde{Q}_z^{\ell-1}, Q_z^*)$ is defined to take the action to maximize the difference between value functions $\tilde{Q}_z^{\ell-1}$ and Q_z^* , i.e., $\pi^\dagger(\tilde{Q}_z^{\ell-1}, Q_z^*)(s) := \arg \max_{a \in \mathcal{A}} |\tilde{Q}_z^{\ell-1}(s, a) - Q_z^*(s, a)|$. Combine the upper bounds above together, for any $\ell \in [\tau]$, we have

$$\begin{aligned} \|\tilde{Q}_z^\ell - Q_z^*\|_\mu &\leq \|\tilde{Q}_z^\ell - \mathcal{T}_z \tilde{Q}_z^{\ell-1}\|_\mu + \gamma \|\tilde{Q}_z^{\ell-1} - Q_z^*\|_{\mathbb{P}_z \mu \times \pi^\dagger(\tilde{Q}_z^{\ell-1}, Q_z^*)} \\ &\leq \sqrt{D} \|\tilde{Q}_z^\ell - \mathcal{T}_z \tilde{Q}_z^{\ell-1}\|_{\tilde{\mu}} + \gamma \|\tilde{Q}_z^{\ell-1} - Q_z^*\|_{\mathbb{P}_z \mu \times \pi^\dagger(\tilde{Q}_z^{\ell-1}, Q_z^*)}, \end{aligned} \quad (\text{B.5})$$

where the inequality follows from Assumption 4.5. Notice that $\mathbb{P}_z \mu \times \pi^\dagger(\tilde{Q}_z^{\ell-1}, Q_z^*)$ is also a policy-induced measure, and we apply the above argument recursively on $\|\tilde{Q}_z^{i-1} - Q_z^*\|_{\mathbb{P}_z \mu \times \pi_{\tilde{Q}_z^{i-1}, Q_z^*}}$ for all $i \in [\ell]$. Take $\ell = \tau$ and expand the inequality (B.5) for τ times, we obtain

$$\|\tilde{Q}_z^\tau - Q_z^*\|_\mu \leq \sum_{\ell=1}^{\tau} \gamma^{\tau-\ell} \sqrt{D} \cdot \|\tilde{Q}_z^\ell - \mathcal{T}_z \tilde{Q}_z^{\ell-1}\|_{\tilde{\mu}} + \frac{\gamma^\tau}{1-\gamma},$$

where $\|\tilde{Q}_z^\ell - \mathcal{T}_z \tilde{Q}_z^{\ell-1}\|_{\tilde{\mu}}$ is the estimation error for target function $\mathcal{T}_z \tilde{Q}_z^{\ell-1}$. \square

Lemma B.8. Let $B := 1/(1-\gamma)$ and $L := L_r/(1-\gamma)$, and for any mean-field term z , $Q \in \mathcal{C}_{B,L}$, and data $\mathcal{D}_{\mathcal{T}_z Q} := \{(s_m, a_m, r_m, s'_m)\}_{m=1}^M$, we define $\hat{f} \in \text{LSE}(\mathcal{C}_{B,L}, \mathcal{D}_{\mathcal{T}_z Q}; \epsilon)$. Then for any measure $\tilde{\mu} \in \Delta_{\mathcal{G}}$ and any training set size $M := |\mathcal{D}_{\mathcal{T}_z Q}|$ such that $M^{\frac{d}{d+2}} \gtrsim \log M$, $M^{\frac{d+8}{d+2}} \log M \gtrsim B^{\frac{3d}{d+2}} (d+1)^{\frac{-2}{d+2}} (dL)^{\frac{-d}{d+2}} \log \frac{1}{\delta}$, and

$$(\log M)^{\frac{d}{2}} \gtrsim 1 + \log(R_d^* B^{\frac{-6}{d+2}} (d+1)^{\frac{-2}{d+2}} (dL)^{\frac{-d}{d+2}}),$$

with probability at least $1 - \delta$

$$\|\hat{f} - \mathcal{T}_z Q\|_{\tilde{\mu}}^2 \lesssim B^{\frac{d+8}{d+2}} (d+1)^{\frac{2}{d+2}} (dL)^{\frac{d}{d+2}} M^{-\frac{2}{d+2}} \log M + \epsilon,$$

where $R_d^* \leq \max\{8dL, 2B + 4dL\}$.

Proof. Before getting into the proof of the lemma, we first provide a few useful definitions. We denote $g_z^\dagger(Q) := \arg \min_{f \in \mathcal{C}_{B,L}} \mathcal{L}(f, \mathcal{T}_z Q)$ to be the LSE of $\mathcal{T}_z Q$ that minimizes the true risk under $\tilde{\mu}$ for any mean-field term z and value function $Q \in \mathcal{C}_{B,L}$. We also define $g_z(Q; \mathcal{D}) := \arg \min_{f \in \mathcal{C}_{B,L}} \mathcal{L}_{\mathcal{D}}(f, \mathcal{T}_z Q)$ to be the LSE of $\mathcal{T}_z Q$ that minimizes the empirical risk on \mathcal{D} . Note that $\mathcal{T}_z Q \in \mathcal{C}_{B,L}$ following Lemma B.5, and we have $g_z^\dagger(Q) = \mathcal{T}_z Q$. Further recall that

$$\mathcal{L}_{\mathcal{D}}(\hat{f}, \mathcal{T}_z Q) = \frac{1}{M} \sum_{m=1}^M (f'(s_m, a_m) - r_m - \gamma \max_{a \in \Gamma(s'_m)} Q(s'_m, a))^2$$

and $\mathcal{L}(\hat{f}, \mathcal{T}_z Q) = \mathbb{E}_{\mathcal{D}}[\mathcal{L}_{\mathcal{D}}(\hat{f}, \mathcal{T}_z Q)]$. For any $\hat{f} \in \text{LSE}(\mathcal{C}_{B,L}, \mathcal{D}; \epsilon)$, it follows by definition that $\mathcal{E}_{\mathcal{D}}(\hat{f}, \mathcal{T}_z Q) = \mathcal{L}_{\mathcal{D}}(\hat{f}, \mathcal{T}_z Q) - \mathcal{L}_{\mathcal{D}}(g_z(Q; \mathcal{D}), \mathcal{T}_z Q) \leq \epsilon$.

It is noteworthy that $\|\hat{f} - \mathcal{T}_z Q\|_{\tilde{\mu}}^2 - \|g_z^\dagger(Q) - \mathcal{T}_z Q\|_{\tilde{\mu}}^2 = \mathcal{L}(\hat{f}, \mathcal{T}_z Q) - \mathcal{L}(g_z^\dagger(Q), \mathcal{T}_z Q)$, which is followed from the fact that for any $f \in \mathcal{C}_{B,L}$,

$$\mathcal{L}(f, \mathcal{T}_z Q) = \mathbb{E}_{\tilde{\mu}}[(f(s, a) - r_z(s, a) - \gamma(\mathcal{T}_z Q)(s'))^2]$$

$$= \mathbb{E}_{\tilde{\mu}}[(f(s, a) - (\mathcal{T}_z Q)(s, a))^2] + \text{Var}_{\tilde{\mu}}[\mathbb{E}[r(s, a) + \gamma(\mathfrak{Y}Q)(s') \mid s, a]],$$

where $(s, a) \sim \tilde{\mu}$ and $s' \sim \mathbb{P}_z(\cdot \mid s, a)$; the variance terms cancel with each other and yield the equality. It follows that

$$\|\hat{f} - \mathcal{T}_z Q\|_{\tilde{\mu}}^2 = \mathcal{L}(\hat{f}, \mathcal{T}_z Q) - \mathcal{L}(\mathcal{T}_z Q, \mathcal{T}_z Q)$$

where the equality is due to $\|g_z^\dagger(Q) - \mathcal{T}_z Q\|_{\tilde{\mu}}^2 = 0$ and $g_z^\dagger(Q) = \mathcal{T}_z Q$. Following Theorem C.2, it holds that for any $\alpha, \beta > 0$ and $0 < \varrho \leq \frac{1}{2}$,

$$\begin{aligned} & \mathbb{P} \left\{ \exists f \in \mathcal{C}_{B,L} : \mathcal{L}(f, \mathcal{T}_z Q) - \mathcal{L}(\mathcal{T}_z Q, \mathcal{T}_z Q) - (\mathcal{L}_{\mathcal{D}}(f, \mathcal{T}_z Q) - \mathcal{L}_{\mathcal{D}}(\mathcal{T}_z Q, \mathcal{T}_z Q)) \right. \\ & \quad \left. \geq \varrho \cdot (\alpha + \beta + \mathcal{L}(f, \mathcal{T}_z Q) - \mathcal{L}(\mathcal{T}_z Q, \mathcal{T}_z Q)) \right\}, \\ & \leq 14 \sup_{\mathcal{D}} \mathcal{N}_1 \left(\frac{\beta \varrho}{20B}, \mathcal{C}_{B,L}, \mathcal{D} \right) \exp \left(-\frac{\varrho^2(1-\varrho)\alpha M}{214(1+\varrho)B^4} \right), \end{aligned}$$

where the inequality follows from $|f| \leq \frac{1}{1-\gamma} = B$ for all $f \in \mathcal{C}_{B,L}$, and $\mathcal{N}_1(\varepsilon, \mathcal{F}, \mathcal{D})$ denotes the ε -covering number of \mathcal{F} on \mathcal{D} with respect to ℓ_1 metric. The supremum is taken with respect to all random dataset \mathcal{D} . This provides a high probability upper bound on $\|\hat{f} - \mathcal{T}_z Q\|_{\tilde{\mu}}^2$ as

$$\begin{aligned} & \mathbb{P} \left\{ (1-\varrho)[\mathcal{L}(\hat{f}, \mathcal{T}_z Q) - \mathcal{L}(\mathcal{T}_z Q, \mathcal{T}_z Q)] \geq \varrho \cdot (\alpha + \beta) + [\mathcal{L}_{\mathcal{D}}(\hat{f}, \mathcal{T}_z Q) - \mathcal{L}_{\mathcal{D}}(\mathcal{T}_z Q, \mathcal{T}_z Q)] \right\} \\ & \leq \mathbb{P} \left\{ \exists f \in \mathcal{C}_{B,L} : \mathcal{L}(f, \mathcal{T}_z Q) - \mathcal{L}(g^\dagger, \mathcal{T}_z Q) - (\mathcal{L}_{\mathcal{D}}(f, \mathcal{T}_z Q) - \mathcal{L}_{\mathcal{D}}(g^\dagger, \mathcal{T}_z Q)) \right. \\ & \quad \left. \geq \varrho \cdot (\alpha + \beta + \mathcal{L}(f, \mathcal{T}_z Q) - \mathcal{L}(g^\dagger, \mathcal{T}_z Q)) \right\}. \end{aligned}$$

Further, for the set of all B -bounded, L -Lipschitz concave functions, its covering entropy is bounded by

$$\begin{aligned} \sup_{\mathcal{D}} \log \mathcal{N}_1 \left(\frac{\beta \varrho}{20B}, \mathcal{C}_{B,L}, \mathcal{D} \right) & \leq \log \mathcal{N}_{\infty} \left(\frac{\beta \varrho}{20B}, \mathcal{C}_{B,L} \right) \\ & \leq 2(d+1) \left(\frac{1600BdL}{\beta \varrho} \right)^{d/2} \ln \left(\frac{200BR_d^*}{\beta \varrho} \right), \end{aligned}$$

where the first inequality follows from upper bounding ℓ_1 covering number with ℓ_{∞} covering number, and the second inequality follows from Lemma C.3. More specifically, for Lipschitz constant L and uniform ℓ_{∞} bound B , the covering entropy of $\mathcal{C}_{B,L}$ is bounded through

$$\log \mathcal{N}_{\infty}(\varepsilon, \mathcal{C}_{B,L}) \leq 2(d+1) \left(\frac{80Ld}{\varepsilon} \right)^{d/2} \ln \left(\frac{10R_d^*}{\varepsilon} \right)$$

for any $\varepsilon \in (0, 80Ld]$ where $L_d := dL \text{diam}(\mathcal{G}) \leq dL$ and $R_d^* \leq \max\{8dL, 2B + 4dL\}$. Thus, for any $\beta \leq \frac{1600}{\varrho} BdL$, we have

$$\begin{aligned} & \mathbb{P} \left\{ (1-\varrho)[\mathcal{L}(\hat{f}, \mathcal{T}_z Q) - \mathcal{L}(\mathcal{T}_z Q, \mathcal{T}_z Q)] \geq \varrho \cdot (\alpha + \beta) + [\mathcal{L}_{\mathcal{D}}(\hat{f}, \mathcal{T}_z Q) - \mathcal{L}_{\mathcal{D}}(\mathcal{T}_z Q, \mathcal{T}_z Q)] \right\} \\ & \leq 14 \exp \left(2(d+1) \left(\frac{1600BdL}{\beta \varrho} \right)^{d/2} \ln \left(\frac{200BR_d^*}{\beta \varrho} \right) - \frac{\varrho^2(1-\varrho)\alpha M}{214(1+\varrho)B^4} \right). \end{aligned}$$

Take $\varrho = \frac{1}{2}$ and for the estimator $\hat{f} \in \mathcal{C}_{B,L}$ that with probability at least

$$1 - 14 \exp \left(2(d+1) \left(\frac{3200BdL}{\beta} \right)^{d/2} \ln \left(\frac{400BR_d^*}{\beta} \right) - \frac{\alpha M}{2568B^4} \right),$$

it holds that the estimation error is bounded by

$$\|\hat{f} - \mathcal{T}_z Q\|_{\tilde{\mu}}^2 = \mathcal{L}(\hat{f}, \mathcal{T}_z Q) - \mathcal{L}(\mathcal{T}_z Q, \mathcal{T}_z Q)$$

$$\begin{aligned}
 &\leq \alpha + \beta + 2(\mathcal{L}_{\mathcal{D}}(\widehat{f}, \mathcal{T}_z Q) - \mathcal{L}_{\mathcal{D}}(\mathcal{T}_z Q, \mathcal{T}_z Q)) \\
 &= \alpha + \beta + 2(\mathcal{E}_{\mathcal{D}}(\widehat{f}, \mathcal{T}_z Q) + \mathcal{L}_{\mathcal{D}}(g_z(Q; \mathcal{D}), \mathcal{T}_z Q) - \mathcal{L}_{\mathcal{D}}(\mathcal{T}_z Q, \mathcal{T}_z Q)) \\
 &\leq \alpha + \beta + 2\epsilon
 \end{aligned}$$

where the second equality is due to $\mathcal{E}_{\mathcal{D}}(\widehat{f}, \mathcal{T}_z Q) = \mathcal{L}_{\mathcal{D}}(\widehat{f}, \mathcal{T}_z Q) - \mathcal{L}_{\mathcal{D}}(g_z(Q; \mathcal{D}), \mathcal{T}_z Q)$, and the last inequality is due to the fact that $g_z(Q; \mathcal{D})$ is defined to be the LSE that minimizes the empirical risk, i.e., $\mathcal{L}_{\mathcal{D}}(g_z(Q; \mathcal{D}), \mathcal{T}_z Q) \leq \mathcal{L}_{\mathcal{D}}(\mathcal{T}_z Q, \mathcal{T}_z Q)$. Consequently, solve for proper α and β to get our conclusion: if $M^{\frac{d}{d+2}} \gtrsim \log M$, $M^{\frac{d}{d+2}} \log M \gtrsim B^{\frac{3d}{d+2}} (d+1)^{\frac{-2}{d+2}} (dL)^{\frac{-d}{d+2}} \log \frac{1}{\delta}$, and

$$(\log M)^{\frac{d}{2}} \gtrsim 1 + \log(R_d^* B^{\frac{-6}{d+2}} (d+1)^{\frac{-2}{d+2}} (dL)^{\frac{-d}{d+2}}),$$

then for any $d \geq 2$, with probability at least $1 - \delta$

$$\|\widehat{f} - \mathcal{T}_z Q\|_{\tilde{\mu}}^2 \lesssim B^{\frac{d+8}{d+2}} (d+1)^{\frac{2}{d+2}} (dL)^{\frac{d}{d+2}} M^{-\frac{2}{d+2}} \log M + \epsilon.$$

□

Lemma B.9. For the MDP with transition kernel \mathbb{P}_z , if the regression of Algorithm 2 finds ϵ -approximate LSE in $\mathcal{C}_{B,L}$ at all τ iterations, then the total approximation error of \widetilde{Q}_z^τ with respect to the optimal value function Q_z^* under any measure $\tilde{\mu}$ is upper bounded by

$$\|\widetilde{Q}_z^\tau - Q_z^*\|_{\tilde{\mu}} \lesssim \frac{\sqrt{D}}{1-\gamma} B^{\frac{d+8}{d+2}} (d+1)^{\frac{2}{d+2}} (dL)^{\frac{d}{d+2}} M^{-\frac{2}{d+2}} \log M + \frac{\sqrt{D}\epsilon}{1-\gamma}$$

with probability at least $1 - \delta$ for the number of iterations

$$\tau \asymp \log \left(D^{-1/2} B^{-\frac{d+8}{d+2}} (d+1)^{-\frac{2}{d+2}} (dL)^{-\frac{d}{d+2}} M^{\frac{2}{d+2}} \frac{1}{\log M} \right) / \log \frac{1}{\gamma}$$

for any sample size $M^{\frac{d}{d+2}} \gtrsim \log M$, $M^{\frac{d}{d+2}} \log M \gtrsim B^{\frac{3d}{d+2}} (d+1)^{\frac{-2}{d+2}} (dL)^{\frac{-d}{d+2}} \log \frac{\tau}{\delta}$, and

$$(\log M)^{\frac{d}{2}} \gtrsim 1 + \log(R_d^* B^{\frac{-6}{d+2}} (d+1)^{\frac{-2}{d+2}} (dL)^{\frac{-d}{d+2}}),$$

where $R_d^* \leq \max\{8dL, 2B + 4dL\}$.

Proof. Combining Lemma B.7 and Lemma B.8, we have with probability at least $1 - \delta$ that the total approximation error of the output \widetilde{Q}_z^τ to the optimal value function Q_z^* under any measure $\tilde{\mu}$ is upper bounded by

$$\begin{aligned}
 \|\widetilde{Q}_z^\tau - Q_z^*\|_{\tilde{\mu}} &\leq \sum_{\ell=1}^{\tau} \gamma^{\tau-\ell} \sqrt{D} \cdot \|\widetilde{Q}_z^\ell - \mathcal{T} \widetilde{Q}_z^{\ell-1}\|_{\tilde{\mu}} + \frac{\gamma^\tau}{1-\gamma} \\
 &\lesssim \sum_{\ell=1}^{\tau} \gamma^{\tau-\ell} \sqrt{D} \cdot \left(B^{\frac{d+8}{d+2}} (d+1)^{\frac{2}{d+2}} (dL)^{\frac{d}{d+2}} M^{-\frac{2}{d+2}} \log M + \epsilon \right) + \frac{\gamma^\tau}{1-\gamma} \\
 &\lesssim \frac{1}{1-\gamma} \sqrt{D} \cdot \left(B^{\frac{d+8}{d+2}} (d+1)^{\frac{2}{d+2}} (dL)^{\frac{d}{d+2}} M^{-\frac{2}{d+2}} \log M + \epsilon \right) + \frac{\gamma^\tau}{1-\gamma},
 \end{aligned}$$

for $M^{\frac{d}{d+2}} \gtrsim \log M$, $M^{\frac{d}{d+2}} \log M \gtrsim B^{\frac{3d}{d+2}} (d+1)^{\frac{-2}{d+2}} (dL)^{\frac{-d}{d+2}} \log \frac{\tau}{\delta}$, and

$$(\log M)^{\frac{d}{2}} \gtrsim 1 + \log(R_d^* B^{\frac{-6}{d+2}} (d+1)^{\frac{-2}{d+2}} (dL)^{\frac{-d}{d+2}}).$$

The first inequality follows from Lemma B.7, the second inequality is due to applying Lemma B.8 such that the upper bound on $\|\widetilde{Q}_z^\ell - \mathcal{T} \widetilde{Q}_z^{\ell-1}\|_{\tilde{\mu}}$ holds with probability at least $1 - \delta/\tau$ for all $\ell \in [\tau]$, and the last inequality follows from a sum over geometric series. If the number of iteration τ is large enough, such that

$$\tau \gtrsim \log \left(D^{-1/2} B^{-\frac{d+8}{d+2}} (d+1)^{-\frac{2}{d+2}} (dL)^{-\frac{d}{d+2}} M^{\frac{2}{d+2}} \frac{1}{\log M} \right) / \log \frac{1}{\gamma},$$

it follows that

$$\|\widetilde{Q}_z^\tau - Q_z^*\|_{\tilde{\mu}} \lesssim \frac{\sqrt{D}}{1-\gamma} B^{\frac{d+8}{d+2}} (d+1)^{\frac{2}{d+2}} (dL)^{\frac{d}{d+2}} M^{-\frac{2}{d+2}} \log M + \frac{\sqrt{D}\epsilon}{1-\gamma}.$$

□

B.5 Proof of FQI With Max-Affine Functions

In this subsection, we present the analysis of our FQI algorithm when the underlying feasible functions are max-affine. Max-affine functions generalize the commonly studied linear functions in RL (Jin et al., 2020; Ayoub et al., 2020; Cai et al., 2020; Du et al., 2021; Min et al., 2022b; Fei and Xu, 2022a), and more importantly, all max-affine functions constitute a subset of all convex functions. Let $\mathcal{A}_{B,L}^K$ denote the set of all bounded L -Lipschitz K -max-affine functions defined as follows

$$\mathcal{A}_{B,L}^K := \{h : \mathcal{G} \rightarrow \mathbb{R} \mid h(x) = \max_{k \in [K]} \alpha_k^\top x + c_k, \|\alpha_k\|_\infty \leq L, h(x) \in [-dL \text{diam}(\mathcal{G}), B]\}.$$

Lemma B.10. *Let $B := 1/(1 - \gamma)$ and $L := L_r/(1 - \gamma)$, and for any mean-field term z , $Q \in \mathcal{C}_{B,L}$, and dataset $\mathcal{D}_{\mathcal{T}_z Q} := \{(s_m, a_m, r_m, s'_m)\}_{m=1}^M$ with $(s_m, a_m) \in \mathbb{R}^d$ i.i.d. sampled from any measure $\tilde{\mu} \in \Delta_{\mathcal{G}}$ for all $m \in [M]$, we define the ϵ -approximate estimator $\hat{f} \in \text{LSE}(\mathcal{A}_{B,L}^{K,+}, \mathcal{D}_{\mathcal{T}_z Q}; \epsilon)$ of $\mathcal{T}_z Q$ for $K \in \mathbb{Z}_+$. Then for any training set size $M := |\mathcal{D}_{\mathcal{T}_z Q}|$ such that*

$$M \gtrsim \frac{1}{e^{2(d+1)M^{d/(d+4)}} \delta}, \quad M^{\frac{2d+4}{d+4}} \log M \gtrsim \frac{B^3(B + 4dL)}{d + 1},$$

then with probability at least $1 - \delta$

$$\begin{aligned} \|\hat{f} - \mathcal{T}_z Q\|_{\tilde{\mu}}^2 &\lesssim (d + 1)B^4 M^{-4/(d+4)} \log M + d^2 L^2 M^{-4/(d+4)} \\ &\quad + BdLM^{-4/(d+4)} \sqrt{\log \frac{2}{\delta} + \frac{B^2 \log \frac{2}{\delta}}{M}} + \epsilon. \end{aligned}$$

Proof. Let us denote $g_z^\dagger(Q) := \arg \min_{f \in \mathcal{A}_{B,L}^{K,+}} \mathcal{L}(f, \mathcal{T}_z Q)$ to be the estimator of $\mathcal{T}_z Q$ in $\mathcal{A}_{B,L}^{K,+}$ that minimizes the true risk under $\tilde{\mu}$ for any mean-field term z . We also define $g_z(Q; \mathcal{D}) := \arg \min_{f \in \mathcal{A}_{B,L}^{K,+}} \mathcal{L}_{\mathcal{D}}(f, \mathcal{T}_z Q)$ to be the LSE of $\mathcal{T}_z Q$ in $\mathcal{A}_{B,L}^{K,+}$ that minimizes the empirical risk on $\mathcal{D}_{\mathcal{T}_z Q}$. Note that for any $Q \in \mathcal{C}_{B,L}$, $\mathcal{T}_z Q \in \mathcal{C}_{B,L}$ following Lemma B.5, but it does not have to be in $\mathcal{A}_{B,L}^{K,+}$. Hence, the estimator $g_z^\dagger(Q) = \mathcal{T}_z Q$ may not hold. For any $\hat{f} \in \text{LSE}(\mathcal{A}_{B,L}^{K,+}, \mathcal{D}_{\mathcal{T}_z Q}; \epsilon)$, it follows by definition that $\mathcal{E}_{\mathcal{D}}(\hat{f}, \mathcal{T}_z Q) = \mathcal{L}_{\mathcal{D}}(\hat{f}, \mathcal{T}_z Q) - \mathcal{L}_{\mathcal{D}}(g_z(Q; \mathcal{D}), \mathcal{T}_z Q) \leq \epsilon$.

Recall that $\|\hat{f} - \mathcal{T}_z Q\|_{\tilde{\mu}}^2 = \mathcal{L}(\hat{f}, \mathcal{T}_z Q) - \mathcal{L}(\mathcal{T}_z Q, \mathcal{T}_z Q)$. Following Theorem C.2, it holds that for any $\alpha, \beta > 0$ and $0 < \varrho \leq \frac{1}{2}$,

$$\begin{aligned} &\mathbb{P} \left\{ \exists f \in \mathcal{A}_{B,L}^{K,+} : \mathcal{L}(f, \mathcal{T}_z Q) - \mathcal{L}(\mathcal{T}_z Q, \mathcal{T}_z Q) - (\mathcal{L}_{\mathcal{D}}(f, \mathcal{T}_z Q) - \mathcal{L}_{\mathcal{D}}(\mathcal{T}_z Q, \mathcal{T}_z Q)) \right. \\ &\quad \left. \geq \varrho \cdot (\alpha + \beta + \mathcal{L}(f, \mathcal{T}_z Q) - \mathcal{L}(\mathcal{T}_z Q, \mathcal{T}_z Q)) \right\}, \\ &\leq 14 \sup_{\mathcal{D}} \mathcal{N}_1 \left(\frac{\beta \varrho}{20B}, \mathcal{A}_{B,L}^{K,+}, \mathcal{D} \right) \exp \left(-\frac{\varrho^2(1 - \varrho)\alpha M}{214(1 + \varrho)B^4} \right), \end{aligned}$$

where the inequality follows from $|f| \leq B$ for all $f \in \mathcal{A}_{B,L}^{K,+}$, and $\mathcal{N}_1(\epsilon, \mathcal{F}, \mathcal{D})$ denotes the ϵ -covering number of \mathcal{F} on \mathcal{D} with respect to ℓ_1 metric. The supremum is taken with respect to all possible data \mathcal{D} . This provides a high probability upper bound on $\|\hat{f} - \mathcal{T}_z Q\|_{\tilde{\mu}}^2$ as

$$\begin{aligned} &\mathbb{P} \left\{ (1 - \varrho)[\mathcal{L}(\hat{f}, \mathcal{T}_z Q) - \mathcal{L}(\mathcal{T}_z Q, \mathcal{T}_z Q)] \geq \varrho \cdot (\alpha + \beta) + [\mathcal{L}_{\mathcal{D}}(\hat{f}, \mathcal{T}_z Q) - \mathcal{L}_{\mathcal{D}}(\mathcal{T}_z Q, \mathcal{T}_z Q)] \right\} \\ &\leq \mathbb{P} \left\{ \exists f \in \mathcal{A}_{B,L}^{K,+} : \mathcal{L}(f, \mathcal{T}_z Q) - \mathcal{L}(\mathcal{T}_z Q, \mathcal{T}_z Q) - (\mathcal{L}_{\mathcal{D}}(f, \mathcal{T}_z Q) - \mathcal{L}_{\mathcal{D}}(\mathcal{T}_z Q, \mathcal{T}_z Q)) \right. \\ &\quad \left. \geq \varrho \cdot (\alpha + \beta + \mathcal{L}(f, \mathcal{T}_z Q) - \mathcal{L}(\mathcal{T}_z Q, \mathcal{T}_z Q)) \right\}. \end{aligned}$$

Further, for the set of all B -bounded, L -Lipschitz K -piece max-affine functions, its covering entropy is bounded by

$$\begin{aligned} \sup_{\mathcal{D}} \log \mathcal{N}_1 \left(\frac{\beta \varrho}{20B}, \mathcal{A}_{B,L}^{K,+}, \mathcal{D} \right) &\leq \log \mathcal{N}_\infty \left(\frac{\beta \varrho}{20B}, \mathcal{A}_{B,L}^K \right) \\ &\leq (d + 1)K \log \frac{(20B + 80dL)B}{\beta \varrho}, \end{aligned}$$

where the first inequality follows from upper bounding ℓ_1 covering number with ℓ_∞ covering number, and the second inequality follows from Lemma C.4 with $L_d \leq dL$. More specifically, for Lipschitz constant L and uniform ℓ_∞ bound B , the covering entropy of $\mathcal{A}_{B,L}^K$ is bounded through

$$\log \mathcal{N}_\infty(\varepsilon, \mathcal{A}_{B,L}^K) \leq (d+1)K \log \frac{B+4Ld}{\varepsilon}$$

for any $\varepsilon \leq B+4Ld$. For our regression problem, d is the dimension of \mathcal{G} . Thus, for any $\beta \leq \frac{20BB}{\varrho}(4dL+B)$, we have

$$\begin{aligned} & \mathbb{P} \left\{ (1-\varrho)[\mathcal{L}(\hat{f}, \mathcal{T}_z Q) - \mathcal{L}(\mathcal{T}_z Q, \mathcal{T}_z Q)] \geq \varrho \cdot (\alpha + \beta) + [\mathcal{L}_{\mathcal{D}}(\hat{f}, \mathcal{T}_z Q) - \mathcal{L}_{\mathcal{D}}(\mathcal{T}_z Q, \mathcal{T}_z Q)] \right\} \\ & \leq 14 \exp \left((d+1)K \log \frac{(20B+80dL)B}{\beta\varrho} - \frac{\varrho^2(1-\varrho)\alpha M}{214(1+\varrho)B^4} \right). \end{aligned}$$

Take $\varrho = \frac{1}{2}$ and for the estimator $\hat{f} \in \mathcal{A}_{B,L}^K$ that with probability at least

$$1 - 14 \exp \left((d+1)K \log \frac{(40B+160dL)B}{\beta} - \frac{\alpha M}{2568B^4} \right),$$

it holds that the estimation error is bounded by

$$\begin{aligned} \|\hat{f} - \mathcal{T}_z Q\|_{\tilde{\mu}}^2 &= \mathcal{L}(\hat{f}, \mathcal{T}_z Q) - \mathcal{L}(\mathcal{T}_z Q, \mathcal{T}_z Q) \\ &\leq \alpha + \beta + 2(\mathcal{L}_{\mathcal{D}}(\hat{f}, \mathcal{T}_z Q) - \mathcal{L}_{\mathcal{D}}(\mathcal{T}_z Q, \mathcal{T}_z Q)) \\ &= \alpha + \beta + 2(\mathcal{E}_{\mathcal{D}}(\hat{f}, \mathcal{T}_z Q) + \mathcal{L}_{\mathcal{D}}(g_z(Q; \mathcal{D}), \mathcal{T}_z Q) - \mathcal{L}_{\mathcal{D}}(\mathcal{T}_z Q, \mathcal{T}_z Q)) \\ &\leq \alpha + \beta + 2\epsilon + 2(\mathcal{L}_{\mathcal{D}}(g_z(Q; \mathcal{D}), \mathcal{T}_z Q) - \mathcal{L}_{\mathcal{D}}(\mathcal{T}_z Q, \mathcal{T}_z Q)) \\ &\leq \alpha + \beta + 2\epsilon + 2(\mathcal{L}_{\mathcal{D}}(g_z^\dagger(Q), \mathcal{T}_z Q) - \mathcal{L}_{\mathcal{D}}(\mathcal{T}_z Q, \mathcal{T}_z Q)), \end{aligned}$$

where the second equality is due to $\mathcal{E}_{\mathcal{D}}(\hat{f}, \mathcal{T}_z Q) = \mathcal{L}_{\mathcal{D}}(\hat{f}, \mathcal{T}_z Q) - \mathcal{L}_{\mathcal{D}}(g_z(Q; \mathcal{D}), \mathcal{T}_z Q)$ and the last inequality is due to the fact that $g_z(Q; \mathcal{D})$ is defined to be the LSE that minimizes the empirical risk, *i.e.*, $\mathcal{L}_{\mathcal{D}}(g_z(Q; \mathcal{D}), \mathcal{T}_z Q) \leq \mathcal{L}_{\mathcal{D}}(g_z^\dagger(Q), \mathcal{T}_z Q)$. Consequently, solve for proper α and β to get that if

$$M \gtrsim \frac{1}{e^{2(d+1)K\delta}}, \quad M \log M \gtrsim \frac{(B+4dL)B^3}{(d+1)K},$$

then for any $d \geq 2$, with probability at least $1 - \delta/2$

$$\|\hat{f} - \mathcal{T}_z Q\|_{\tilde{\mu}}^2 \lesssim \frac{(d+1)K \log M}{B^4} \frac{1}{M} + 2\epsilon + 2(\mathcal{L}_{\mathcal{D}}(g_z^\dagger(Q), \mathcal{T}_z Q) - \mathcal{L}_{\mathcal{D}}(\mathcal{T}_z Q, \mathcal{T}_z Q)).$$

Following Lemma C.5, under any measure $\tilde{\mu}$ it holds that

$$\|g_z^\dagger(Q) - \mathcal{T}_z Q\|_{\tilde{\mu}}^2 \leq \|\bar{g}_z^\dagger(Q) - \mathcal{T}_z Q\|_{\tilde{\mu}}^2 \leq 72^2 L_d^2 (K-1)^{-4/d},$$

where $\bar{g}_z^\dagger(Q) := \sigma(g_z^\dagger(Q))$ is the truncated estimator, where σ denotes ReLU function and $g_z^\dagger(Q) := \arg \min_{f \in \mathcal{A}_{B,L}^{K-1}} \mathcal{L}(f, \mathcal{T}_z Q)$. Note that $\|g_z^\dagger(Q) - \mathcal{T}_z Q\|_{\tilde{\mu}}^2 = \mathcal{L}(g_z^\dagger(Q), \mathcal{T}_z Q) - \mathcal{L}(\mathcal{T}_z Q, \mathcal{T}_z Q)$, and following Bernstein's inequality, we have

$$\begin{aligned} & |\mathcal{L}(g_z^\dagger(Q), \mathcal{T}_z Q) - \mathcal{L}(\mathcal{T}_z Q, \mathcal{T}_z Q) - \mathcal{L}_{\mathcal{D}}(g_z^\dagger(Q), \mathcal{T}_z Q) + \mathcal{L}_{\mathcal{D}}(\mathcal{T}_z Q, \mathcal{T}_z Q)| \\ &= \left| \mathcal{L}(g_z^\dagger(Q), \mathcal{T}_z Q) - \mathcal{L}(\mathcal{T}_z Q, \mathcal{T}_z Q) - \frac{1}{n} \sum_{(s,a,r,s') \in \mathcal{D}} [(g_z^\dagger(Q)(s,a) - (\mathcal{T}_z Q)(s,a)) \right. \\ &\quad \left. \cdot (g_z^\dagger(Q)(s,a) + (\mathcal{T}_z Q)(s,a) - 2r - 2\gamma(\mathfrak{I}Q)(s'))] \right| \\ &\leq \frac{2B^2 \log \frac{1}{\delta}}{3M} + 2\sqrt{\frac{\log \frac{1}{\delta}}{M}} \cdot \text{Var}([(g_z^\dagger(Q)(s,a) - (\mathcal{T}_z Q)(s,a)) \end{aligned}$$

$$\cdot (g_z^\dagger(Q)(s, a) + (\mathcal{T}_z Q)(s, a) - 2r - 2\gamma(\mathfrak{S}Q)(s'))^{1/2}.$$

Since we have $|g_z^\dagger(Q)(s, a) + (\mathcal{T}_z Q)(s, a) - 2r - 2\gamma(\mathfrak{S}Q)(s')| \leq 2B$ and $\text{Var}(X) = \mathbb{E}X^2 - (\mathbb{E}X)^2$ for any random variable X , we further simplify the upper bound into

$$\begin{aligned} & |\mathcal{L}(g_z^\dagger(Q), \mathcal{T}_z Q) - \mathcal{L}(\mathcal{T}_z Q, \mathcal{T}_z Q) - \mathcal{L}_{\mathcal{D}}(g_z^\dagger(Q), \mathcal{T}_z Q) + \mathcal{L}_{\mathcal{D}}(\mathcal{T}_z Q, \mathcal{T}_z Q)| \\ & \leq \frac{2B^2 \log \frac{1}{\delta}}{3M} + 4B \|g_z^\dagger(Q) - \mathcal{T}_z Q\|_{\tilde{\mu}} \sqrt{\frac{\log \frac{1}{\delta}}{M}}. \end{aligned}$$

It follows that with probability at least $1 - \delta/2$

$$\begin{aligned} \mathcal{L}_{\mathcal{D}}(g_z^\dagger(Q), \mathcal{T}_z Q) - \mathcal{L}_{\mathcal{D}}(\mathcal{T}_z Q, \mathcal{T}_z Q) & \leq \|g_z^\dagger(Q) - \mathcal{T}_z Q\|_{\tilde{\mu}}^2 + 4B \|g_z^\dagger(Q) - \mathcal{T}_z Q\|_{\tilde{\mu}} \sqrt{\frac{\log \frac{2}{\delta}}{M}} \\ & \quad + \frac{2B^2 \log \frac{2}{\delta}}{3M} \\ & \leq 72^2 L_d^2 (K-1)^{-4/d} + 288BL_d K^{-2/d} \sqrt{\frac{\log \frac{2}{\delta}}{M}} + \frac{2B^2 \log \frac{2}{\delta}}{3M}. \end{aligned}$$

Combine the upper bounds above together, we have with probability at least $1 - \delta$

$$\|\hat{f} - \mathcal{T}_z Q\|_{\tilde{\mu}}^2 \lesssim (d+1)KB^4 \frac{\log M}{M} + d^2 L^2 (K-1)^{-4/d} + BdLK^{-2/d} \sqrt{\frac{\log \frac{2}{\delta}}{M}} + \frac{B^2 \log \frac{2}{\delta}}{M} + \epsilon.$$

Take $K = \lceil M^{d/(d+4)} \rceil$, then it holds that

$$\begin{aligned} \|\hat{f} - \mathcal{T}_z Q\|_{\tilde{\mu}}^2 & \lesssim (d+1)B^4 M^{-4/(d+4)} \log M + d^2 L^2 M^{-4/(d+4)} \\ & \quad + BdLM^{-4/(d+4)} \sqrt{\log \frac{2}{\delta}} + \frac{B^2 \log \frac{2}{\delta}}{M} + \epsilon. \end{aligned}$$

□

Lemma B.11. Let $B := 1/(1 - \gamma)$ and $L := L_r/(1 - \gamma)$. For the MDP with transition kernel \mathbb{P}_z , if the regression of Algorithm 2 finds ϵ -approximate LSE in $\mathcal{A}_{B,L}^K$ with $K = \lceil M^{d/(d+4)} \rceil$ at all τ iterations, then the approximation error of \tilde{Q}_z^τ with respect to the optimal value function Q_z^* under any measure $\tilde{\mu}$ is upper bounded by

$$\|\tilde{Q}_z^\tau - Q_z^*\|_{\tilde{\mu}} \lesssim \frac{\sqrt{D}}{1 - \gamma} ((d+1)B^4 + (B + dL)dL) M^{-\frac{4}{d+4}} \log M + \frac{\sqrt{D}\epsilon}{1 - \gamma}$$

with probability at least $1 - \delta$ for the number of iterations

$$\tau \asymp \log \frac{M^{4/(d+4)}}{d^2 L^2 \sqrt{D}} / \log \frac{1}{\gamma}$$

for any sample size $M = |\mathcal{D}|$ such that

$$\begin{aligned} M & \gtrsim \frac{1}{\delta \log \frac{1}{\gamma}} \log \frac{M^{4/(d+4)}}{d^2 L^2 \sqrt{D}}, & M^{\frac{2d+4}{d+4}} \log M & \gtrsim \frac{B^3(B + 4dL)}{d+1}, \\ M & \gtrsim \frac{\tau}{\exp(2(d+1)M^{d/(d+4)})\delta}, & M^{\frac{d}{d+4}} & \gtrsim \frac{B^2}{d^2 L^2} \log \frac{2\tau}{\delta}. \end{aligned}$$

Proof. Combining Lemma B.7 and Lemma B.10, we have with probability at least $1 - \delta$ that the total approximation error of the output \tilde{Q}_z^τ to the optimal value function Q_z^* under any measure $\tilde{\mu}$ is upper bounded by

$$\|\tilde{Q}_z^\tau - Q_z^*\|_{\tilde{\mu}} \leq \sum_{\ell=1}^{\tau} \gamma^{\tau-\ell} \sqrt{D} \cdot \|\tilde{Q}_z^\ell - \mathcal{T}\tilde{Q}_z^{\ell-1}\|_{\tilde{\mu}} + \frac{\gamma^\tau}{1 - \gamma}$$

$$\begin{aligned}
 &\lesssim \frac{1}{1-\gamma} \sqrt{D} \cdot \left((d+1)B^4 M^{-4/(d+4)} \log M + d^2 L^2 M^{-4/(d+4)} \right. \\
 &\quad \left. + BdLM^{-4/(d+4)} \sqrt{\log \frac{2\tau}{\delta}} + \frac{B^2 \log \frac{2\tau}{\delta}}{M} + \epsilon \right) + \frac{\gamma^\tau}{1-\gamma} \\
 &\lesssim \frac{1}{1-\gamma} \sqrt{D} \cdot \left((d+1)B^4 M^{-4/(d+4)} \log M + 2d^2 L^2 M^{-4/(d+4)} \right. \\
 &\quad \left. + BdLM^{-4/(d+4)} \sqrt{\log \frac{2\tau}{\delta}} + \epsilon \right) + \frac{\gamma^\tau}{1-\gamma},
 \end{aligned}$$

for $M \gtrsim \frac{\tau}{\exp(2(d+1)M^{d/(d+4)})\delta}$ and $M^{\frac{2d+4}{d+4}} \log M \gtrsim \frac{B^3(B+4dL)}{d+1}$. The first inequality follows from Lemma B.7, the second inequality is due to applying Lemma B.8 such that the upper bound on $\|\tilde{Q}_z^\ell - \mathcal{T}\tilde{Q}_z^{\ell-1}\|_{\tilde{\mu}}$ holds with probability at least $1 - \delta/\tau$ for all $\ell \in [\tau]$, and the last inequality follows from $M^{d/(d+4)} \geq \frac{B^2}{d^2 L^2} \log \frac{2\tau}{\delta}$. Take the number of iteration

$$\tau = \left\lceil \log \frac{M^{4/(d+4)}}{d^2 L^2 \sqrt{D}} / \log \frac{1}{\gamma} \right\rceil,$$

it follows that for $M \geq \frac{1}{\delta \log \frac{1}{\gamma}} \log \frac{M^{4/(d+4)}}{d^2 L^2 \sqrt{D}}$, the approximation error

$$\|\tilde{Q}_z^\tau - Q_z^*\|_{\tilde{\mu}} \lesssim \frac{\sqrt{D}}{1-\gamma} \left((d+1)B^4 + (B+dL)dL \right) M^{-\frac{4}{d+4}} \log M + \frac{\sqrt{D}\epsilon}{1-\gamma}.$$

□

B.6 Proof of FQI With ICNN

In this subsection, we present our theoretical work for the ICNN function family. As a convex object, ICNN function is robust against outliers and input perturbations (Min et al., 2021a), such merits have been discussed in related works (Christmann and Steinwart, 2007; Blanchet et al., 2019; Chen et al., 2020; Pfrommer et al., 2023). We combine Lemma 3.1 and Lemma C.4 to prove Lemma B.12 and then Lemma B.13. More specifically, Lemma C.4 provides an upper bound on the covering entropy for L -Lipschitz K -piece max-affine functions, and the covering entropy of such function set can provide an upper bound on the covering entropy of $\mathcal{N}_{B,L}^K$ as $\mathcal{N}_{B,L}^K \subseteq \mathcal{A}_{B,L}^{K+1}$. We note that our argument assumes finding the global minimum of the functions represented by the ICNN. We do not elaborate on the convergence properties of such neural network functions under different backpropagation schemes, and they have been discussed in a vast body of literature (Chen et al., 2018; Schmidt et al., 2011; Arora et al., 2019; Du et al., 2019; Chen et al., 2021a; Song et al., 2021a; Tan et al., 2022).

Lemma B.12. *Let $B := 1/(1-\gamma)$ and $L := L_r/(1-\gamma)$, and for any mean-field term z , $Q \in \mathcal{C}_{B,L}$, and data $\mathcal{D}_{\mathcal{T}_z Q} := \{(s_m, a_m, r_m, s'_m)\}_{m=1}^M$ with $(s_m, a_m) \in \mathbb{R}^d$ i.i.d. sampled from any measure $\tilde{\mu} \in \Delta_{\mathcal{G}}$ for all $m \in [M]$, we define the ϵ -approximate estimator $\hat{f} \in \text{LSE}(\mathcal{N}_{B,L}^K, \mathcal{D}_{\mathcal{T}_z Q}; \epsilon)$ of $\mathcal{T}_z Q$ for $K \in \mathbb{Z}_+$. Then for any training set size $M := |\mathcal{D}_{\mathcal{T}_z Q}|$ such that*

$$M \gtrsim \frac{1}{e^{2(d+1)K}\delta}, \quad M^{\frac{2d+4}{d+4}} \log M \gtrsim \frac{B^3(B+4dL)}{d+1},$$

it holds with probability at least $1 - \delta$

$$\begin{aligned}
 \|\hat{f} - \mathcal{T}_z Q\|_{\tilde{\mu}}^2 &\lesssim (d+1)B^4 M^{-4/(d+4)} \log M + d^2 L^2 M^{-4/(d+4)} \\
 &\quad + BdLM^{-4/(d+4)} \sqrt{\log \frac{2}{\delta}} + \frac{B^2 \log \frac{2}{\delta}}{M} + \epsilon.
 \end{aligned}$$

Proof. Let us denote $g_z^\dagger(Q) := \arg \min_{f \in \mathcal{N}_{B,L}^K} \mathcal{L}(f, \mathcal{T}_z Q)$ to be the estimator of $\mathcal{T}_z Q$ in $\mathcal{N}_{B,L}^K$ that minimizes the true risk under $\tilde{\mu}$ for any mean-field term z and value function $Q \in \mathcal{C}_{B,L}$. We also define $g_z(Q; \mathcal{D}) := \arg \min_{f \in \mathcal{N}_{B,L}^K} \mathcal{L}_{\mathcal{D}}(f, \mathcal{T}_z Q)$ to be the LSE of $\mathcal{T}_z Q$ in $\mathcal{N}_{B,L}^K$ that minimizes the empirical risk on \mathcal{D} . For any $\hat{f} \in \text{LSE}(\mathcal{N}_{B,L}^K, \mathcal{D}; \epsilon)$, it follows by definition that $\mathcal{E}_{\mathcal{D}}(\hat{f}, \mathcal{T}_z Q) = \mathcal{L}_{\mathcal{D}}(\hat{f}, \mathcal{T}_z Q) - \mathcal{L}_{\mathcal{D}}(g_z(Q; \mathcal{D}), \mathcal{T}_z Q) \leq \epsilon$.

Recall that $\|\widehat{f} - \mathcal{T}_z Q\|_{\mu}^2 = \mathcal{L}(\widehat{f}, \mathcal{T}_z Q) - \mathcal{L}(\mathcal{T}_z Q, \mathcal{T}_z Q)$. Following Theorem C.2, it holds that for any $\alpha, \beta > 0$ and $0 < \varrho \leq \frac{1}{2}$,

$$\begin{aligned} & \mathbb{P} \left\{ \exists f \in \mathcal{N}_{B,L}^K : \mathcal{L}(f, \mathcal{T}_z Q) - \mathcal{L}(\mathcal{T}_z Q, \mathcal{T}_z Q) - (\mathcal{L}_{\mathcal{D}}(f, \mathcal{T}_z Q) - \mathcal{L}_{\mathcal{D}}(\mathcal{T}_z Q, \mathcal{T}_z Q)) \right. \\ & \quad \left. \geq \varrho \cdot (\alpha + \beta + \mathcal{L}(f, \mathcal{T}_z Q) - \mathcal{L}(\mathcal{T}_z Q, \mathcal{T}_z Q)) \right\}, \\ & \leq 14 \sup_{\mathcal{D}} \mathcal{N}_1 \left(\frac{\beta \varrho}{20B}, \mathcal{N}_{B,L}^K, \mathcal{D} \right) \exp \left(-\frac{\varrho^2(1-\varrho)\alpha B^4 M}{214(1+\varrho)} \right), \end{aligned}$$

where the inequality follows from $|f| \leq B + dL = B$ for all $f \in \mathcal{N}_{B,L}^K$, and $\mathcal{N}_1(\varepsilon, \mathcal{F}, \mathcal{D})$ denotes the ε -covering number of \mathcal{F} on \mathcal{D} with respect to ℓ_1 metric. The supremum is taken with respect to all possible data \mathcal{D} . This provides a high probability upper bound on $\|\widehat{f} - \mathcal{T}_z Q\|_{\mu}^2$ as

$$\begin{aligned} & \mathbb{P} \left\{ (1-\varrho)[\mathcal{L}(\widehat{f}, \mathcal{T}_z Q) - \mathcal{L}(\mathcal{T}_z Q, \mathcal{T}_z Q)] \geq \varrho \cdot (\alpha + \beta) + [\mathcal{L}_{\mathcal{D}}(\widehat{f}, \mathcal{T}_z Q) - \mathcal{L}_{\mathcal{D}}(\mathcal{T}_z Q, \mathcal{T}_z Q)] \right\} \\ & \leq \mathbb{P} \left\{ \exists f \in \mathcal{N}_{B,L}^K : \mathcal{L}(f, \mathcal{T}_z Q) - \mathcal{L}(g^\dagger, \mathcal{T}_z Q) - (\mathcal{L}_{\mathcal{D}}(f, \mathcal{T}_z Q) - \mathcal{L}_{\mathcal{D}}(g^\dagger, \mathcal{T}_z Q)) \right. \\ & \quad \left. \geq \varrho \cdot (\alpha + \beta + \mathcal{L}(f, \mathcal{T}_z Q) - \mathcal{L}(g^\dagger, \mathcal{T}_z Q)) \right\}. \end{aligned}$$

Further, for the set of all functions in $\mathcal{N}_{B,L}^K$, its covering entropy is bounded by

$$\begin{aligned} \sup_{\mathcal{D}} \log \mathcal{N}_1 \left(\frac{\beta \varrho}{20B}, \mathcal{N}_{B,L}^K, \mathcal{D} \right) & \leq \log \mathcal{N}_{\infty} \left(\frac{\beta \varrho}{20B}, \mathcal{N}_{B,L}^K \right) \\ & \leq (d+1)(K+1) \log \frac{(40B + 80dL)B}{\beta \varrho}, \end{aligned}$$

where the first inequality follows from upper bounding ℓ_1 covering number with ℓ_{∞} covering number and Lemma 3.1, and the second inequality follows from Lemma C.4 and $L_d \leq dL$. More specifically, for Lipschitz constant L uniform ℓ_{∞} bound B , the covering entropy of $\mathcal{N}_{B,L}^K$ is bounded through

$$\log \mathcal{N}_{\infty}(\varepsilon, \mathcal{N}_{B,L}^K) \leq (d+1)(K+1) \log \frac{B + 4L_d}{\varepsilon}$$

for any $\varepsilon \leq B + 4L_d$. Thus, for any $\beta \leq \frac{20}{\varrho}(4dL_r + B)B$, we have

$$\begin{aligned} & \mathbb{P} \left\{ (1-\varrho)[\mathcal{L}(\widehat{f}, \mathcal{T}_z Q) - \mathcal{L}(\mathcal{T}_z Q, \mathcal{T}_z Q)] \geq \varrho \cdot (\alpha + \beta) + [\mathcal{L}_{\mathcal{D}}(\widehat{f}, \mathcal{T}_z Q) - \mathcal{L}_{\mathcal{D}}(\mathcal{T}_z Q, \mathcal{T}_z Q)] \right\} \\ & \leq 14 \exp \left((d+1)(K+1) \log \frac{(20B + 80dL)B}{\beta \varrho} - \frac{\varrho^2(1-\varrho)\alpha B^4 M}{214(1+\varrho)} \right). \end{aligned}$$

Take $\varrho = \frac{1}{2}$ and for the estimator $\widehat{f} \in \mathcal{N}_{B,L}^K$ that with probability at least

$$1 - 14 \exp \left((d+1)(K+1) \log \frac{(40B + 160dL)B}{\beta} - \frac{\alpha B^4 M}{2568} \right),$$

it holds that the estimation error is bounded by

$$\begin{aligned} \|\widehat{f} - \mathcal{T}_z Q\|_{\mu}^2 & = \mathcal{L}(\widehat{f}, \mathcal{T}_z Q) - \mathcal{L}(\mathcal{T}_z Q, \mathcal{T}_z Q) \\ & \leq \alpha + \beta + 2(\mathcal{L}_{\mathcal{D}}(\widehat{f}, \mathcal{T}_z Q) - \mathcal{L}_{\mathcal{D}}(\mathcal{T}_z Q, \mathcal{T}_z Q)) \\ & = \alpha + \beta + 2(\mathcal{E}_{\mathcal{D}}(\widehat{f}, \mathcal{T}_z Q) + \mathcal{L}_{\mathcal{D}}(g_z(Q; \mathcal{D}), \mathcal{T}_z Q) - \mathcal{L}_{\mathcal{D}}(\mathcal{T}_z Q, \mathcal{T}_z Q)) \\ & \leq \alpha + \beta + 2\epsilon + 2(\mathcal{L}_{\mathcal{D}}(g_z(Q; \mathcal{D}), \mathcal{T}_z Q) - \mathcal{L}_{\mathcal{D}}(\mathcal{T}_z Q, \mathcal{T}_z Q)) \\ & \leq \alpha + \beta + 2\epsilon + 2(\mathcal{L}_{\mathcal{D}}(g_z^\dagger(Q), \mathcal{T}_z Q) - \mathcal{L}_{\mathcal{D}}(\mathcal{T}_z Q, \mathcal{T}_z Q)), \end{aligned}$$

where the second equality is due to $\mathcal{E}_{\mathcal{D}}(\widehat{f}, \mathcal{T}_z Q) = \mathcal{L}_{\mathcal{D}}(\widehat{f}, \mathcal{T}_z Q) - \mathcal{L}_{\mathcal{D}}(g_z(Q; \mathcal{D}), \mathcal{T}_z Q)$ and the last inequality is due to the fact that $g_z(Q; \mathcal{D})$ is defined to be the LSE that minimizes the empirical risk, i.e., $\mathcal{L}_{\mathcal{D}}(g_z(Q; \mathcal{D}), \mathcal{T}_z Q) \leq \mathcal{L}_{\mathcal{D}}(g_z^\dagger(Q), \mathcal{T}_z Q)$. Consequently, solve for proper α and β to get that if

$$M \gtrsim \frac{1}{e^{2(d+1)K} \delta}, \quad M \log M \gtrsim \frac{B + 4dL}{B^3(d+1)K},$$

then for any $d \geq 2$, with probability at least $1 - \delta/2$

$$\|\widehat{f} - \mathcal{T}_z Q\|_{\tilde{\mu}}^2 \lesssim (d+1)(K+1)B^4 \frac{\log M}{M} + 2\epsilon + 2(\mathcal{L}_{\mathcal{D}}(g_z^\dagger(Q), \mathcal{T}_z Q) - \mathcal{L}_{\mathcal{D}}(\mathcal{T}_z Q, \mathcal{T}_z Q)).$$

Following Lemma C.5 and the conclusion that $\mathcal{A}_{B,L}^{K,+} \subseteq \mathcal{N}_{B,L}^K$ from Lemma 3.1, under any measure $\tilde{\mu}$ it holds that

$$\|g_z^\dagger(Q) - \mathcal{T}_z Q\|_{\tilde{\mu}}^2 \leq \|\bar{g}_z^\dagger(Q) - \mathcal{T}_z Q\|_{\tilde{\mu}}^2 \leq 72^2 L_d^2 (K-1)^{-4/d},$$

where $\bar{g}_z^\dagger(Q) := \sigma(g_z^\dagger(Q))$ is the truncated estimator, where σ denotes ReLU function and $g_z^\dagger(Q) := \arg \min_{f \in \mathcal{A}_{B,L}^{K-1}} \mathcal{L}(f, \mathcal{T}_z Q)$. Note that $\|g_z^\dagger(Q) - \mathcal{T}_z Q\|_{\tilde{\mu}}^2 = \mathcal{L}(g_z^\dagger(Q), \mathcal{T}_z Q) - \mathcal{L}(\mathcal{T}_z Q, \mathcal{T}_z Q)$, and following Bernstein's inequality, we have

$$\begin{aligned} & |\mathcal{L}(g_z^\dagger(Q), \mathcal{T}_z Q) - \mathcal{L}(\mathcal{T}_z Q, \mathcal{T}_z Q) - \mathcal{L}_{\mathcal{D}}(g_z^\dagger(Q), \mathcal{T}_z Q) + \mathcal{L}_{\mathcal{D}}(\mathcal{T}_z Q, \mathcal{T}_z Q)| \\ & \leq \frac{2B^2 \log \frac{1}{\delta}}{3M} + 4B \|g_z^\dagger(Q) - \mathcal{T}_z Q\|_{\tilde{\mu}} \sqrt{\frac{\log \frac{1}{\delta}}{M}}. \end{aligned}$$

It follows that with probability at least $1 - \delta/2$

$$\mathcal{L}_{\mathcal{D}}(g_z^\dagger(Q), \mathcal{T}_z Q) - \mathcal{L}_{\mathcal{D}}(\mathcal{T}_z Q, \mathcal{T}_z Q) \leq 72^2 d^2 L^2 (K-1)^{-4/d} + 288BdLK^{-2/d} \sqrt{\frac{\log \frac{2}{\delta}}{M}} + \frac{2B^2 \log \frac{2}{\delta}}{3M}.$$

Combine the upper bounds above together, we have with probability at least $1 - \delta$

$$\|\widehat{f} - \mathcal{T}_z Q\|_{\tilde{\mu}}^2 \lesssim \frac{(d+1)(K+1) \log M}{(1-\gamma)^4} \frac{1}{M} + d^2 L^2 (K-1)^{-4/d} + BdLK^{-2/d} \sqrt{\frac{\log \frac{2}{\delta}}{M}} + \frac{B^2 \log \frac{2}{\delta}}{M} + \epsilon.$$

Take $K = \lceil M^{d/(d+4)} \rceil$, then it holds that

$$\begin{aligned} \|\widehat{f} - \mathcal{T}_z Q\|_{\tilde{\mu}}^2 & \lesssim (d+1)B^4 M^{-4/(d+4)} \log M + d^2 L^2 M^{-4/(d+4)} \\ & \quad + BdLM^{-4/(d+4)} \sqrt{\log \frac{2}{\delta}} + \frac{B^2 \log \frac{2}{\delta}}{M} + \epsilon. \end{aligned}$$

□

Lemma B.13. Let $B := 1/(1-\gamma)$ and $L := L_r/(1-\gamma)$. For the MDP induced by any mean-field term z , if Algorithm 2 achieves ϵ -approximate LSE in $\mathcal{N}_{B,L}^K$ with $K = \lceil M^{d/(d+4)} \rceil$ at all τ iterations, then the approximation error of \tilde{Q}_z^τ with respect to the optimal value function Q_z^* under any measure $\tilde{\mu}$ is upper bounded by

$$\|\tilde{Q}_z^\tau - Q_z^*\|_{\tilde{\mu}} \lesssim \frac{\sqrt{D}}{1-\gamma} ((d+1)B^4 + (B+dL)dL) M^{-\frac{4}{d+4}} \log M + \frac{\sqrt{D}\epsilon}{1-\gamma}$$

with probability at least $1 - \delta$ for the number of iterations

$$\tau \asymp \log \frac{M^{4/(d+4)}}{d^2 L^2 \sqrt{D}} / \log \frac{1}{\gamma}$$

for any sample size $M = |\mathcal{D}|$ such that

$$\begin{aligned} M & \gtrsim \frac{1}{\delta \log \frac{1}{\gamma}} \log \frac{M^{4/(d+4)}}{d^2 L^2 \sqrt{D}}, & M^{\frac{2d+4}{d+4}} \log M & \gtrsim \frac{B^3(B+4dL)}{d+1}, \\ M & \gtrsim \frac{\tau}{\exp(2(d+1)M^{d/(d+4)})\delta}, & M^{\frac{d}{d+4}} & \gtrsim \frac{B^2}{d^2 L^2} \log \frac{2\tau}{\delta}. \end{aligned}$$

Proof. Combining Lemma B.7 and Lemma B.10, we have with probability at least $1 - \delta$ that the total approximation error of the output \tilde{Q}_z^τ to the optimal value function Q_z^* under any measure $\tilde{\mu}$ is upper bounded by

$$\begin{aligned} \|\tilde{Q}_z^\tau - Q_z^*\|_{\tilde{\mu}} &\leq \sum_{\ell=1}^{\tau} \gamma^{\tau-\ell} \sqrt{D} \cdot \|\tilde{Q}_z^\ell - \mathcal{T}\tilde{Q}_z^{\ell-1}\|_{\tilde{\mu}} + \frac{\gamma^\tau}{1-\gamma} \\ &\lesssim \frac{1}{1-\gamma} \sqrt{D} \cdot \left((d+1)B^4 M^{-4/(d+4)} \log M + d^2 L^2 M^{-4/(d+4)} \right. \\ &\quad \left. + BdLM^{-4/(d+4)} \sqrt{\log \frac{2\tau}{\delta}} + \frac{B^2 \log \frac{2\tau}{\delta}}{M} + \epsilon \right) + \frac{\gamma^\tau}{1-\gamma} \\ &\lesssim \frac{1}{1-\gamma} \sqrt{D} \cdot \left((d+1)B^4 M^{-4/(d+4)} \log M + 2d^2 L^2 M^{-4/(d+4)} \right. \\ &\quad \left. + BdLM^{-4/(d+4)} \sqrt{\log \frac{2\tau}{\delta}} + \epsilon \right) + \frac{\gamma^\tau}{1-\gamma}, \end{aligned}$$

for $M \gtrsim \frac{\tau}{\exp(2(d+1)M^{d/(d+4)})\delta}$ and $M^{\frac{2d+4}{d+4}} \log M \gtrsim \frac{B^3(B+4dL)}{d+1}$. The first inequality follows from Lemma B.7, the second inequality is due to applying Lemma B.8 such that the upper bound on $\|\tilde{Q}_z^\ell - \mathcal{T}\tilde{Q}_z^{\ell-1}\|_{\tilde{\mu}}$ holds with probability at least $1 - \delta/\tau$ for all $\ell \in [\tau]$, and the last inequality follows from $M^{d/(d+4)} \geq \frac{B^2}{d^2 L^2} \log \frac{2\tau}{\delta}$. Take the number of iteration

$$\tau = \left\lceil \log \frac{M^{4/(d+4)}}{d^2 L^2 \sqrt{D}} / \log \frac{1}{\gamma} \right\rceil,$$

it follows that for $M \geq \frac{1}{\delta \log \frac{1}{\gamma}} \log \frac{M^{4/(d+4)}}{d^2 L^2 \sqrt{D}}$, the approximation error

$$\|\tilde{Q}_z^\tau - Q_z^*\|_{\tilde{\mu}} \lesssim \frac{\sqrt{D}}{1-\gamma} \left((d+1)B^4 + (B+dL)dL \right) M^{-\frac{4}{d+4}} \log M + \frac{\sqrt{D}\epsilon}{1-\gamma}.$$

□

B.7 Proof of Lemma 3.1

Lemma 3.1 shows the existence of a set $\mathcal{N}_{B,L}^K$ of ICNNs that covers the set of all non-negative L -Lipschitz K -piece max-affine functions bounded by B with a convex parameter set. Under such parameter set, we show that $\mathcal{N}_{B,L}^K$ is equivalent to the set of all truncated L -Lipschitz K -piece max-affine functions bounded by B , which is a subset of all positive L -Lipschitz $(K+1)$ -piece max-affine functions bounded by B .

Proof. The proof is based on the construction of a K -layer ICNN similar to that in Chen et al. (2018, Theorem 1), where they showed that any K -piece max-affine function can be represented exactly by a K -layer ICNN. However, Chen et al. (2018) did not provide a parameter set for the neural network such that the K -layer ICNN is able to represent non-negative L -Lipschitz functions bounded by B exclusively, which is what we are going to show in this proof.

Especially, $\mathcal{N}_{B,L}^K$ is a set of K -layer ICNNs with input dimension $2d$ and all hidden layer dimension 1. For any input $x \in \mathcal{X}$ and L -Lipschitz K -max-affine function $f_K(x) = \max\{\alpha_1^\top x + c_1, \dots, \alpha_K^\top x + c_K\}$ that maps from \mathcal{X} to $[0, B]$, the corresponding input vector for the network representation of f_K in $\mathcal{N}_{B,L}^K$ is given by $x^\ddagger := [x^\top, -x^\top]^\top$. More specifically, we have all layers y_i of the network representation of f_K defined as

$$y_i = \sigma(y_{i-1} + (\alpha_i - \alpha_{i+1})_+^\top x - (\alpha_{i+1} - \alpha_i)_+^\top x + (c_i - c_{i+1}))$$

for all $i = 2, \dots, K-1$ and

$$\begin{aligned} y_1 &= \sigma((\alpha_1 - \alpha_2)_+^\top x - (\alpha_2 - \alpha_1)_+^\top x + (c_1 - c_2)), \\ y_K &= \sigma(y_{K-1} + (\alpha_K)_+^\top x - (-\alpha_K)_+^\top x + c_K). \end{aligned}$$

Notice that we have $W_0^{(y)} = 0$ and $W_i^{(y)} = 1$ fixed for all $i \in [K-1]$; the shifting terms $\beta_{K-1} = c_K$ and $\beta_i = c_{i+1} - c_{i+2}$ for all $i = 0, \dots, K-2$. Further, we have for all $i = 0, \dots, K-2$

$$W_i^{(x)} = [(\alpha_{i+1} - \alpha_{i+2})_+^\top, (\alpha_{i+2} - \alpha_{i+1})_+^\top]^\top \in \mathbb{R}_+^{2d}$$

and $W_{K-1}^{(x)} = [(\alpha_K)_+^\top, (-\alpha_K)_+^\top]^\top \in \mathbb{R}_+^{2d}$, such that $W_i^{(x)\top} x^\ddagger = (\alpha_{i+1} - \alpha_{i+2})_+^\top x - (\alpha_{i+2} - \alpha_{i+1})_+^\top x$ and $\|W_i^{(x)}\|_\infty = \|\alpha_{i+1} - \alpha_{i+2}\|_\infty$ for all $i = 0, \dots, K-2$ as well as $W_{K-1}^{(x)\top} x^\ddagger = (\alpha_K)_+^\top x - (-\alpha_K)_+^\top x$ and $\|W_{K-1}^{(x)}\|_\infty = \|\alpha_K\|_\infty$.

On the other hand, for any set of parameters $\{(W_i^{(x)}, \beta_i)\}_{i=0}^{K-1}$ with $W_i^{(x)} \in \mathbb{R}_+^{2d}$ and $\beta_i \in \mathbb{R}$ for all $i = 0, \dots, K-1$, we deem $W_i^{(x)} = [w_{i,1}^\top, w_{i,2}^\top]^\top$ as concatenations of vectors $w_{i,1}, w_{i,2} \in \mathbb{R}^d$, then for any input x^\ddagger , we have

$$\langle W_i^{(x)}, x^\ddagger \rangle = w_{i,1}^\top x - w_{i,2}^\top x.$$

Further notice that for a concise representation $w_i := w_{i,1} - w_{i,2}$, we have $w_i^\top x = \langle W_i^{(x)}, x^\ddagger \rangle$. There exists a set of vectors $\{\alpha_i\}_{i=1}^K$ such that we have $\alpha_K = w_{K-1}$ and $\alpha_i = w_{i-1} + \alpha_{i+1}$ for all $i \in [K-1]$; similarly, there exists a set of scalars $\{c_i\}_{i=1}^K$ such that $c_K = \beta_{K-1}$ and $c_i = \beta_{i-1} + c_{i+1}$ for all $i \in [K-1]$. Solve the system of equations to get $\alpha_i = \sum_{j=i-1}^{K-1} w_j$ and $c_i = \sum_{j=i-1}^{K-1} \beta_j$ for all $i \in [K]$, and the ICNN with parameterization $\{(W_i^{(x)}, \beta_i)\}_{i=0}^{K-1}$ represents function $f_K := \sigma(\max\{\alpha_1^\top x + c_1, \dots, \alpha_K^\top x + c_K\})$. Under the convex constraint set $\|\sum_{j=i-1}^{K-1} w_j\|_\infty \leq L$, it holds that the represented max-affine function f_K is L -Lipschitz. Moreover, the represented max-affine function f_K is upper bounded by B on \mathcal{X} if and only if the composing affine functions $\alpha_i^\top x + c_i$ for all $i \in [K]$ are upper bounded; this is due to $f_K \geq \alpha_i^\top x + c_i$ for all $i \in [K]$ by definition. It requires that $\alpha_i^\top x + c_i \leq B$ for all $i \in [K]$ and $x \in \mathcal{X}$, which can be equivalently written as

$$\max_{x \in \mathcal{X}} \sum_{j=i-1}^{K-1} w_j^\top x + \sum_{j=i-1}^{K-1} \beta_j \leq B.$$

Note that for any $\alpha_1, \alpha_2 \in \mathbb{R}^d$, it holds that $\max_{x \in \mathcal{X}} (\alpha_1^\top x + \alpha_2^\top x) \leq \max_{x \in \mathcal{X}} \alpha_1^\top x + \max_{x \in \mathcal{X}} \alpha_2^\top x$, and the boundedness requirement forms a convex constraint set following the convexity of \mathcal{X} . It is not hard to see that under the convex constraint set described by all the aforementioned conditions, the set $\mathcal{N}_{B,L}^K$ ICNN represents only the truncated K -max-affine functions that are L -Lipschitz and upper bounded by B . \square

C SUPPORTING LEMMAS

Lemma C.1. *Let $f : \mathfrak{F} \rightarrow \mathbb{R} \cap \{\infty\}$ be a differentiable ζ -strongly convex function with respect to a norm $\|\cdot\|$, where \mathfrak{F} is the set of all measurable functions on $\mathcal{A} \subseteq \mathbb{R}^{d_w}$. Let the effective domain of f be $S = \{x \in \mathfrak{F} : f(x) \in \mathbb{R}\}$ and f^* be the Fenchel conjugate of f . Then, we have*

1. f^* is differentiable on \mathfrak{F} ;
2. $\nabla f^*(y) = \arg \max_{x \in S} \{\langle x, y \rangle - f(x)\}$;
3. f^* is $\frac{1}{\zeta}$ -smooth with respect to the dual norm $\|\cdot\|_*$. That is, for all $y_1, y_2 \in \mathfrak{F}$, we have

$$\|\nabla f^*(y) - \nabla f^*(y')\| \leq \frac{1}{\zeta} \|y - y'\|_*.$$

Proof. This lemma is adopted from [Shalev-Shwartz and Singer \(2007\)](#), and we provide a brief proof here for completeness. Suppose we assume there exist $x, x' \in S$ such that $x \in \partial f^*(y)$ and $x' \in \partial f^*(y')$ for any $y \in S$. For any $x \in \partial f^*(y)$, it holds

$$f^*(g) \geq f^*(y) + \langle y, g - x \rangle$$

for any $g \in S$. Rearrange the terms to get Since $(f^*)^* = f$ and the fact that for any y'

$$\begin{aligned} f^*(y') &= \sup_{g \in S} \{\langle g, y' \rangle - f(g)\} \\ &\geq \langle x, y' \rangle - f(x) \\ &= f^*(y) + \langle x, y' - y \rangle, \end{aligned}$$

we conclude that $y \in \partial f(x)$ and $y \in \partial f(x')$. Since f is ζ -strongly convex function, then for any $x, x' \in \mathfrak{X}$ and $y \in \partial f(x)$, $y' \in \partial f(x')$, we have $\|x - x'\|^2 \leq \frac{1}{\zeta} \langle x - x', y - y' \rangle$. Following Holder's inequality, we conclude that

$$\|\nabla f^*(y) - \nabla f^*(y')\| \leq \frac{1}{\zeta} \|y - y'\|_*.$$

□

Theorem C.2. We assume a data set $\{(X_i, Y_i)\}_{i=1}^n$ generated from distribution $(X, Y) \sim \mathcal{D}$ such that there exists $B > 1$ and $|Y| \leq B$ almost surely. We further define $m(x) := \mathbb{E}[Y \mid X = x]$. Let \mathcal{F} be a set of functions $f : \mathbb{R}^d \rightarrow \mathbb{R}$ and let $|f(x)| \leq B$. Then, it holds for each $n \geq 1$ that

$$\begin{aligned} & \mathbb{P} \left\{ \exists f \in \mathcal{F} : \mathbb{E} |f(X) - Y|^2 - \mathbb{E} |m(X) - Y|^2 - \frac{1}{n} \sum_{i=1}^n \{|f(X_i) - Y_i|^2 - |m(X_i) - Y_i|^2\} \right. \\ & \quad \left. \geq \varrho \cdot (\alpha + \beta + \mathbb{E} |f(X) - Y|^2 - \mathbb{E} |m(X) - Y|^2) \right\} \\ & \leq 14 \sup_{x_1^n} \mathcal{N}_1 \left(\frac{\beta \varrho}{20B}, \mathcal{F}, x_1^n \right) \exp \left(-\frac{\varrho^2 (1 - \varrho) \alpha n}{214(1 + \varrho) B^4} \right), \end{aligned}$$

where $x_1^n := (x_1, \dots, x_n)$ denotes n fixed points in \mathbb{R}^d , $\alpha, \beta > 0$, and $0 < \varrho \leq 1/2$.

Proof. This is Györfi et al. (2002, Theorem 11.4). □

Lemma C.3. Define the class of uniformly B -bounded, subdifferentiable, and uniformly L -Lipschitz functions on $\mathcal{X} \subseteq \mathbb{R}^d$ as

$$\mathcal{C}_{\mathcal{X}, B, L} := \{f : \mathcal{X} \rightarrow \mathbb{R} \mid f \text{ is convex, } \|f\|_\infty \leq B, \forall x \in \mathcal{X} : \partial f(x) \neq \emptyset, \forall g \in \partial f(x) : \|g\|_\infty \leq L\}.$$

Let $R_d^* := \max\{8L_d, 2B + 4L_d\}$ where $L_d := dL \text{diam}(\mathcal{X})$, then for all $\varepsilon \in (0, 80L_d]$, the covering entropy of $\mathcal{C}_{\mathcal{X}, B, L}$ is bounded by

$$\log \mathcal{N}_\infty(\varepsilon, \mathcal{C}_{\mathcal{X}, B, L}) \leq 2(d+1) \left(\frac{80L_d}{\varepsilon} \right)^{d/2} \ln \left(\frac{10R_d^*}{\varepsilon} \right).$$

Proof. This is Balázs et al. (2015, Lemma 4.3). □

Lemma C.4. For any $K \in \mathbb{Z}_+$, define the class of uniformly B -bounded and L -Lipschitz K -piece max-affine functions on $\mathcal{X} \subseteq \mathbb{R}^d$ as

$$\mathcal{A}_{\mathcal{X}, B, L}^K := \{h : \mathcal{X} \rightarrow \mathbb{R} \mid h(x) = \max_{k \in [K]} \alpha_k^\top x + c_k, \|\alpha_k\|_\infty \leq L, h(x) \in [-B_d, B]\}$$

where $B_d := B + dL \text{diam}(\mathcal{X})$. Then for any $R_d := 2B + 4L_d$ with $L_d := dL \text{diam}(\mathcal{X})$ and any $\varepsilon \in (0, R_d]$,

$$\log \mathcal{N}_\infty(\varepsilon, \mathcal{A}_{\mathcal{X}, B, L}^K) \leq (d+1)K \log \frac{R_d}{\varepsilon}.$$

Proof. This is Balázs et al. (2015, Lemma 4.2). □

Lemma C.5. For all B -bounded and L -Lipschitz convex function $f \in \mathcal{C}_{\mathcal{X}, B, L}$ on $\mathcal{X} \subseteq \mathbb{R}^d$ and any $K \in \mathbb{Z}_+$, it holds that the best approximation h of f in the set of $\mathcal{A}_{\mathcal{X}, B, L}^K$ satisfies

$$\inf_{h \in \mathcal{A}_{\mathcal{X}, B, L}^K} \|f - h\|_\infty \leq 72L_d K^{-2/d},$$

where $L_d := dL \text{diam}(\mathcal{X})$.

Proof. This is Balázs et al. (2015, Lemma 4.1). □