

NORTHWESTERN UNIVERSITY

On the Optimality and Complexity of Reinforcement Learning

A DISSERTATION

SUBMITTED TO THE GRADUATE SCHOOL
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

for the degree

DOCTOR OF PHILOSOPHY

Field of Industrial Engineering and Management Sciences

By

Zuyue Fu

EVANSTON, ILLINOIS

August 2022

© Copyright by Zuyue Fu 2022

All Rights Reserved

ABSTRACT

On the Optimality and Complexity of Reinforcement Learning

Zuyue Fu

In this dissertation, we aim to develop algorithms that achieve optimality with provable complexity guarantees under various settings in reinforcement learning (RL). Specifically, in Markov decision processes (MDPs), we study single-agent and multi-agent online RL, respectively, and offline RL under the presence of unobserved confounders.

- Single-agent online RL. We design a single-timescale actor-critic method to solve single-agent RL, where the actor and critic are updated simultaneously. Specifically, in each iteration, the critic update is obtained by applying the Bellman evaluation operator only once while the actor is updated in the policy gradient direction computed using the critic. We prove that the actor sequence converges to a globally optimal policy at a sublinear $O(K^{-1/2})$ rate, where K is the number of iterations.
- Multi-agent online RL. We study discrete-time mean-field Markov games with infinite numbers of agents, where each agent aims to minimize its ergodic cost. Specifically, we consider a linear-quadratic case, where the agents have identical

linear state transitions and quadratic cost functions. For such a game, we provide sufficient conditions for the existence and uniqueness of its Nash equilibrium, and also propose a model-free mean-field actor-critic algorithm. In particular, we prove that our algorithm converges to the Nash equilibrium at a linear rate.

- Offline RL with unobserved confounders. We study offline RL in the face of unobserved confounders. Offline RL is typically facing the following two significant challenges: (i) the agent may be confounded by the unobserved state variables; (ii) the offline data collected a priori does not provide sufficient coverage for the environment. To tackle the above challenges, we study the policy learning in the confounded MDPs with the aid of instrumental variables (IVs). Specifically, we propose value- and ratio-based identification results for the identification of the expected total reward. Then by leveraging pessimism and our identification results, we propose various policy learning methods with the finite-sample suboptimality guarantee of finding the optimal in-class policy under minimal data coverage and modeling assumptions.

Acknowledgements

It is only with the unwavering support of many people that I can achieve this milestone in my life, and I would like to express my sincerest gratitude to them.

First and foremost, I would like to thank my advisor Professor Zhaoran Wang. It is a great honor for me to be advised by and work alongside one of the great minds in the theory of machine learning. Thank you for being a great advisor who has patiently guided me, once an outsider to the field of machine learning, into such a beautiful world. Thank you for being a great mentor and offering me so much advice, guidance, and opportunities for my academic and career development. Thank you for being a great friend, who is so caring, understanding, supporting, encouraging, and always willing to help. Your knowledge, sharpness, dedication, and passion, inspire me forward in my academic journey. You are great role models to whom I shall always look up.

I would like to thank many faculties I have had throughout my Ph.D. studies, all of whom I have learned a great deal from. Thank you, Professor Ermin Wei and Professor Zhengling Qi, for your dedicated guidance and serving on my dissertation committee. Thank you, Professor Zhuoran Yang, for teaching me so much and for your help with many challenging problems I have encountered. Thank you, Professor Yongxin Chen and Professor Yang Liu, for the valuable opportunities to collaborate with you and leading me to various fields in machine learning. I have also had the privilege of collaborating

with many other exceptional researchers: Lingxiao, Yufeng, Hongyi, Zhihan, Zhihong, and Luofeng. Thank you all for your efforts in our research projects.

I want to thank the IEMS faculty for offering me admission to the Ph.D. program several years ago, teaching me so much about industrial engineering, and being role models as exceptional researchers. I also want to thank the IEMS staff members for always being so warm and helpful.

The past five years have been incredibly happy, for which I have my dear friends to thank. Lingxiao and Boyi, when I began my Ph.D. studies five years ago, I could never have imagined that I could make such dear friends like you, with whom I would share so many memories (food and car, primarily) and from whom I would learn so much. Bing, thanks for being my roommate and helping me so much in life during my earliest year in the country. We will for sure watch more soccer games in the future. Alex, you are so interesting and helpful. I really enjoy the days that we do homework together. Shukai, thanks for tirelessly organizing various activities during the past three years. Oliver, you are always so dedicated and passionate about everything. It is always a pleasure talking with you. Michelle and Xinyi, thanks for all the help you have given me (including the fancy food). Xiaochun, thank you for being in my life. It is the greatest joy to always have you by my side. There are so many friends not only in the IEMS family but also throughout the university that I have drank and ate and played with. I could not list all the names here, but the memories with these friends will always be with me.

Last but not least, I want to thank all my family members. Thank you for bringing me to the wonderful world. Thank you for your unconditional love, trust and support.

Without your extraordinary efforts on my education since my childhood, this dissertation would be impossible to accomplish. I would like to dedicate this work to you.

Table of Contents

ABSTRACT	3
Acknowledgements	5
Table of Contents	8
List of Tables	11
List of Figures	12
Chapter 1. Introduction	13
1.1. Single-Timescale Actor-Critic Provably Finds Globally Optimal Policy	15
1.2. Actor-Critic Provably Finds Nash Equilibria of Linear-Quadratic Mean-Field Games	22
1.3. Offline Reinforcement Learning with Instrumental Variables in Confounded Markov Decision Processes	28
Chapter 2. Single-Timescale Actor-Critic Provably Finds Globally Optimal Policy	35
2.1. Background	35
2.2. Algorithms	40
2.3. Theoretical Results	47
2.4. Proof Sketch of Theorem 2.3.4	51
2.5. Conclusion	58

Chapter 3. Actor-Critic Provably Finds Nash Equilibria of Linear-Quadratic Mean-Field Games	59
3.1. Linear-Quadratic Mean-Field Game	59
3.2. Mean-Field Actor-Critic	65
3.3. Global Convergence Results	70
3.4. Conclusion	73
Chapter 4. Offline Reinforcement Learning with Instrumental Variables in Confounded Markov Decision Processes	74
4.1. Confounded Markov Decision Processes	74
4.2. Assumptions and Identification Results	77
4.3. Instrumental-Variable-Assisted RL with Pessimism	84
4.4. Theoretical Results	90
4.5. Dual Formulation	99
4.6. Identifiability	101
References	103
Appendix A. Supplemental Materials in Chapter 2	130
A.1. Details of Algorithms	130
A.2. Convergence Results of Algorithm 4	131
A.3. Proofs of Theorems	135
A.4. Supporting Results	151
A.5. Proofs of Propositions	155
A.6. Proofs of Lemmas	173

Appendix B. Supplemental Materials in Chapter 3	187
B.1. Notations in the Appendix	187
B.2. Auxiliary Algorithms and Analysis	187
B.3. General Formulation	198
B.4. Proofs of Theorems	203
B.5. Proofs of Propositions	228
B.6. Proofs of Lemmas	243
B.7. Auxiliary Results	277
Appendix C. Supplemental Materials in Chapter 4	278
C.1. Analysis of Maximum Likelihood Estimation	278
C.2. Proofs of Results in §4.2	283
C.3. Proofs of Results in §4.4.1	288
C.4. Proofs of Results in §4.4.2	298
C.5. Proof of Results in §4.4.3	311
C.6. Proof of Results in §4.5	340
C.7. Auxiliary Results	342
C.8. Proofs of Auxiliary Results	349

List of Tables

- 1.1 Assumptions required by our VF-, MIS-, and DR-based methods, where w^π is the density ratio between visitation measures induced by the policy π and the behavior policy b (see (4.2.1) for a detailed definition), and V^π is the state-value function of the policy π . Here \mathcal{V} and Π are function classes, and C_* is a positive absolute constant. 32

List of Figures

- 2.1 Illustration of the relationship among $A_{1,k}$, $A_{2,k}$, $A_{3,k}$, ϵ_{k+1}^c , e_{k+1} , and ϑ_k . Here $\{\theta_k, \omega_k\}$ and $\{\theta_{k+1}, \omega_{k+1}\}$ are two consecutive iterates of actor-critic. The red arrow from Q_{ω_k} to $Q_{\omega_{k+1}}$ represents the critic update and the red arrow from $Q^{\pi_{\theta_k}}$ to $Q^{\pi_{\theta_{k+1}}}$ represents the action-value functions associated with the two policies in any actor update. Here ϑ_k given in (2.4.7) quantifies the difference between π_{θ_k} and $\pi_{\theta_{k+1}}$ in terms of their KL distances to π^* . In addition, the cyan arrows represent quantities $A_{1,k}$, $A_{2,k}$, and $A_{3,k}$ introduced in (2.4.2)–(2.4.4), which are intermediate terms used for analyzing the error $Q^* - Q^{\pi_{k+1}}$. Finally, the blue arrows represent ϵ_{k+1}^c and e_{k+1} defined in (2.4.5) and (2.4.6), respectively. Here ϵ_{k+1}^c corresponds to the statistical error due to having finite data whereas e_{k+1} essentially quantifies the difference between π_{θ_k} and $\pi_{\theta_{k+1}}$. 54
- 4.1 A graphical illustration of the confounded MDP. 80

CHAPTER 1

Introduction

In reinforcement learning (RL, Sutton and Barto (2018)), the agent aims to make sequential decisions that maximize the expected total reward through interacting with the environment and learning from the experiences, where the environment is modeled as a Markov decision process (MDP, Puterman (2014)). RL with deep neural networks achieves tremendous successes in practice, e.g., video games (Silver et al., 2016; OpenAI, 2018), robotics (Kalashnikov et al., 2018), solving social dilemmas (de Cote et al., 2006; Leibo et al., 2017; Hughes et al., 2018), etc. However, there still lack a theoretical understanding on the optimality, i.e., how good is a learned policy compared with the optimal policy, and complexity, i.e, how large a dataset is required to learn a good policy, of various types of RL methods. In this dissertation, we aim to develop and study algorithms that provably achieve optimality and complexity guarantees under various settings in RL.

In single-agent online RL, where the agent can actively interact with the environment to collect new data, to achieve the highest possible total reward in expectation, the actor-critic method (Konda and Tsitsiklis, 2000) is the among the most commonly used algorithms. Specifically, to establish convergence guarantees for actor-critic, most existing works either focus on the bi-level setting or the two-timescale setting, which are seldom adopted in practice. In Chapter 2, we propose a single-timescale actor-critic method, and investigate its convergence and global optimality under linear and deep neural network function approximation. In particular, we focus on the family of energy-based policies

and aim to find the optimal policy within this class. In our actor-critic algorithm, the actor update follows proximal policy optimization (PPO, Schulman et al. (2017)) and the critic update is obtained by applying the Bellman evaluation operator only once to the current critic iterate, which is more closed to the practical algorithms.

Multi-agent RL extends single-agent RL to sequential decision-making problems involving multiple agents. Mean-field game (MFG, Huang et al. (2003, 2006); Lasry and Lions (2006a,b, 2007)) is a specific form of multi-agent RL, which utilizes mean-field approximation to model the strategic interactions within a large population. In Chapter 3, we develop an efficient model-free RL approach to solve MFGs, which provably attains the Nash equilibrium. In particular, we focus on discrete-time MFGs with linear state transitions and quadratic cost functions, where the aggregated effect of the population is quantified by the mean-field state. In detail, we propose a mean-field actor-critic algorithm, which alternately updates the policy and mean-field state. In theory, we prove that the sequence of policies and its corresponding sequence of mean-field states converge to the unique Nash equilibrium at a linear rate.

Since actively interacting with the environment in an MDP is usually either expansive or unethical (e.g., in healthcare (Raghu et al., 2017; Komorowski et al., 2018; Gottesman et al., 2019), autonomous driving (Shalev-Shwartz et al., 2016)), a growing body of literature focus on designing RL methods in the offline setting, where the agent aims to learn an optimal policy π^* in the infinite-horizon Markov decision process (MDP, Puterman (2014)) only through observational data. In Chapter 4, we aim to solve the following two challenges in offline RL: (i) The agent may be confounded by unobserved variables (confounders) in the observational data; (ii) Due to insufficient data coverage, typical

methods (Precup, 2000; Antos et al., 2008b; Chen and Jiang, 2019) may fail to converge without imposing strong assumptions on the data generation process. To tackle such challenges, by leveraging instrumental variables (IVs) and the principle of pessimism, we propose value-based and ratio-based estimators of the optimal policy, which are shown enjoy provable optimality and complexity guarantees.

1.1. Single-Timescale Actor-Critic Provably Finds Globally Optimal Policy

In reinforcement learning (RL, Sutton et al. (1998)), the agent aims to make sequential decisions that maximize the expected total reward through interacting with the environment and learning from the experiences, where the environment is modeled as a Markov Decision Process (MDP, (Puterman, 2014)). To learn a policy that achieves the highest possible total reward in expectation, the actor-critic method (Konda and Tsitsiklis, 2000) is among the most commonly used algorithms. In actor-critic, the actor refers to the policy and the critic corresponds to the value function that characterizes the performance of the actor. This method directly optimizes the expected total return over the policy class by iteratively improving the actor, where the update direction is determined by the critic. In particular, recently, actor-critic combined with deep neural networks (LeCun et al., 2015) achieves tremendous empirical successes in solving large-scale RL tasks, such as the game of Go (Silver et al., 2017), StarCraft (Vinyals et al., 2019), Dota (OpenAI, 2018), Rubik’s cube (Agostinelli et al., 2019; Akkaya et al., 2019), and autonomous driving (Sallab et al., 2017). See Li (2017) for a detailed survey of the recent developments of deep reinforcement learning.

Despite these great empirical successes of actor-critic, there is still an evident chasm between theory and practice. Specifically, to establish convergence guarantees for actor-critic, most existing works either focus on the bi-level setting or the two-timescale setting, which are seldom adopted in practice. In particular, under the bi-level setting (Yang et al., 2019a; Wang et al., 2019; Agarwal et al., 2019; Fu et al., 2019; Liu et al., 2019; Abbasi-Yadkori et al., 2019a,b; Cai et al., 2019; Hao et al., 2020; Mei et al., 2020; Bhandari and Russo, 2020), the actor is updated only after the critic solves the policy evaluation sub-problem completely, which is equivalent to applying the Bellman evaluation operator to the previous critic for infinite times. Consequently, actor-critic under the bi-level setting is a double-loop iterative algorithm where the inner loop is allocated for solving the policy evaluation sub-problem of the critic. In terms of theoretical analysis, such a double-loop structure decouples the analysis for the actor and critic. For the actor, the problem is essentially reduced to analyzing the convergence of a variant of the policy gradient method (Sutton et al., 2000; Kakade, 2002) where the error of the gradient estimate depends on the policy evaluation error of the critic. Besides, under the two-timescale setting (Borkar and Konda, 1997; Konda and Tsitsiklis, 2000; Xu et al., 2020; Wu et al., 2020; Hong et al., 2020), the actor and the critic are updated simultaneously, but with disparate stepsizes. More concretely, the stepsize of the actor is set to be much smaller than that of the critic, with the ratio between these stepsizes converging to zero. In an asymptotic sense, such a separation between stepsizes ensures that the critic completely solves its policy evaluation sub-problem asymptotically. In other words, such a two-timescale scheme results in a separation between actor and critic in an asymptotic sense, which leads to asymptotically unbiased policy gradient estimates. In sum, in terms of convergence analysis, the existing

theory of actor-critic hinges on decoupling the analysis for critic and actor, which is ensured via focusing on the bi-level or two-timescale settings.

However, most practical implementations of actor-critic are under the single-timescale setting (Peters and Schaal, 2008a; Schulman et al., 2015; Mnih et al., 2016; Schulman et al., 2017; Haarnoja et al., 2018), where the actor and critic are simultaneously updated, and particularly, the actor is updated without the critic reaching an approximate solution to the policy evaluation sub-problem. Meanwhile, in comparison with the two-timescale setting, the actor is equipped with a much larger stepsize in the the single-timescale setting such that the asymptotic separation between the analysis of actor and critic is no longer valid.

Furthermore, when it comes to function approximation, most existing works only analyze the convergence of actor-critic with either linear function approximation (Xu et al., 2020; Wu et al., 2020; Hong et al., 2020), or shallow-neural-network parameterization (Wang et al., 2019; Liu et al., 2019). In contrast, practically used actor-critic methods such as asynchronous advantage actor-critic (Mnih et al., 2016) and soft actor-critic (Haarnoja et al., 2018) oftentimes represent both the actor and critic using deep neural networks.

Thus, the following question is left open:

Does single-timescale actor-critic provably find a globally optimal policy under the function approximation setting, especially when deep neural networks are employed?

To answer such a question, we make the first attempt to investigate the convergence and global optimality of single-timescale actor-critic with linear and neural network function approximation. In particular, we focus on the family of energy-based policies and aim

to find the optimal policy within this class. Here we represent both the energy function and the critic as linear or deep neural network functions. In our actor-critic algorithm, the actor update follows proximal policy optimization (PPO) (Schulman et al., 2017) and the critic update is obtained by applying the Bellman evaluation operator only once to the current critic iterate. As a result, the actor is updated before the critic solves the policy evaluation sub-problem. Such a coupled updating structure persists even when the number of iterations goes to infinity, which implies that the update direction of the actor is always biased compared with the policy gradient direction. This brings an additional challenge that is absent in the bi-level and the two-timescale settings, where the actor and critic are decoupled asymptotically.

To tackle such a challenge, our analysis captures the joint effect of actor and critic updates on the objective function, dubbed as the “double contraction” phenomenon, which plays a pivotal role for the success of single-timescale actor-critic. Specifically, thanks to the discount factor of the MDP, the Bellman evaluation operator is contractive, which implies that, after each update, the critic makes noticeable progress by moving towards the value function associated with the current actor. As a result, although we use a biased estimate of the policy gradient, thanks to the contraction brought by the discount factor, the accumulative effect of the biases is controlled. Such a phenomenon enables us to characterize the progress of each iteration of joint actor and critic update, and thus yields the convergence to the globally optimal policy. In particular, for both the linear and neural settings, we prove that, single-timescale actor-critic finds a $O(K^{-1/2})$ -globally optimal policy after K iterations. To the best of our knowledge, we seem to establish the first theoretical guarantee of global convergence and global optimality for

actor-critic with function approximation in the single-timescale setting. Moreover, under the broader scope of policy optimization with nonlinear function approximation, our work seems to prove convergence and optimality guarantees for actor-critic with deep neural network for the first time.

Contribution. Our contribution is two-fold. First, in the single-timescale setting with linear function approximation, we prove that, after K iterations of actor and critic updates, actor-critic returns a policy that is at most $O(K^{-1/2})$ inferior to the globally optimal policy. Second, when both the actor and critic are represented by deep neural networks, we prove a similar $O(K^{-1/2})$ rate of convergence to the globally optimal policy when the architecture of the neural networks are properly chosen.

Related Work. Our work extends the line of works on the convergence of actor-critic under the function approximation setting. In particular, actor-critic is first introduced in Sutton et al. (2000); Konda and Tsitsiklis (2000). Later, Kakade (2002); Peters and Schaal (2008b) propose the natural actor-critic method which updates the policy via the natural gradient (Amari, 1998) direction. The convergence of (natural) actor-critic with linear function approximation are studied in Bhatnagar et al. (2008, 2009); Bhatnagar (2010); Castro and Meir (2010); Maei (2018). However, these works only characterize the asymptotic convergence of actor-critic and their proofs all resort to tools from stochastic approximation via ordinary differential equations (Borkar, 2008). As a result, these works only show that actor-critic with linear function approximation converges to the set of stable equilibria of a set of ordinary differential equations. Recently, Zhang et al. (2019a) propose a variant of actor-critic where Monte-Carlo sampling is used to ensure the critic and the policy gradient estimates are unbiased. Although they incorporate nonlinear

function approximation in the actor, they only establish finite-time convergence result to a stationary point of the expected total reward. Moreover, due to having an inner loop for solving the policy evaluation sub-problem, they focus on the bi-level setting. Moreover, under the two-timescale setting, Wu et al. (2020); Xu et al. (2020) show that actor-critic with linear function approximation finds an ε -stationary point with $\tilde{O}(\varepsilon^{-5/2})$ samples, where ε measures the squared norm of the policy gradient. All of these results establish the convergence of actor-critic, without characterizing the optimality of the policy obtained by actor-critic.

In terms of the global optimality of actor-critic, Fazel et al. (2018); Malik et al. (2018); Tu and Recht (2018); Yang et al. (2019a); Bu et al. (2019); Fu et al. (2019) show that policy gradient and bi-level actor-critic methods converge to the globally optimal policies under the linear-quadratic setting, where the state transitions follow a linear dynamical system and the reward function is quadratic. For general MDPs, Bhandari and Russo (2019) recently prove the global optimality of vanilla policy gradient under the assumption that the families of policies and value functions are both convex. In addition, our work is also related to Liu et al. (2019) and Wang et al. (2019), where they establish the global optimality of proximal policy optimization and (natural) actor-critic, respectively, where both the actor and critic are parameterized by two-layer neural networks. Our work is also related to Agarwal et al. (2019); Abbasi-Yadkori et al. (2019a,b); Cai et al. (2019); Hao et al. (2020); Mei et al. (2020); Bhandari and Russo (2020), which focus on characterizing the optimality of natural policy gradient in tabular and/or linear settings. However, these aforementioned works all focus on bi-level actor-critic, where the actor is updated only after the critic solves the policy evaluation sub-problem to an approximate optimum.

Besides, these works consider linear or two-layer neural network function approximations whereas we focus on the setting with deep neural networks. Furthermore, under the two-timescale setting, Xu et al. (2020); Hong et al. (2020) prove that linear actor-critic requires a sample complexity of $\tilde{O}(\varepsilon^{-4})$ for obtaining an ε -globally optimal policy. In comparison, our $O(K^{-1/2})$ convergence for single-timescale actor-critic can be translated into a similar $\tilde{O}(\varepsilon^{-4})$ sample complexity directly. Moreover, when reusing the data, our result leads to an improved $\tilde{O}(\varepsilon^{-2})$ sample complexity. In addition, our work is also related to Geist et al. (2019), which proposes a variant of policy iteration algorithm with Bregman divergence regularization. Without considering an explicit form of function approximation, their algorithm is shown to converge to the globally optimal policy at a similar $O(K^{-1/2})$ rate, where K is the number of policy updates. In contrast, our method is single-timescale actor-critic with linear or deep neural network function approximation, which enjoys both global convergence and global optimality. Meanwhile, our proof is based on a finite-sample analysis, which involves dealing with the algorithmic errors that track the performance of actor and critic updates as well as the statistical error due to having finite data.

Our work is also related to the literature on deep neural networks. Previous works (Daniely, 2017; Jacot et al., 2018; Wu et al., 2018; Allen-Zhu et al., 2018a,b; Du et al., 2018; Zou et al., 2018; Chizat and Bach, 2018; Jacot et al., 2018; Li and Liang, 2018; Cao and Gu, 2019a,b; Arora et al., 2019; Lee et al., 2019; Gao et al., 2019) analyze the computational and statistical rates of supervised learning methods with overparameterized neural networks. In contrast, our work employs overparameterized deep neural networks

in actor-critic for solving RL tasks, which is significantly more challenging than supervised learning due to the interplay between the actor and the critic.

1.2. Actor-Critic Provably Finds Nash Equilibria of Linear-Quadratic Mean-Field Games

In reinforcement learning (RL) (Sutton and Barto, 2018), an agent learns to make decisions that minimize its expected total cost through sequential interactions with the environment. Multi-agent reinforcement learning (MARL) (Shoham et al., 2003, 2007; Busoniu et al., 2008) aims to extend RL to sequential decision-making problems involving multiple agents. In a non-cooperative game, we are interested in the Nash equilibrium (Nash, 1951), which is a joint policy of all the agents such that each agent cannot decrease its expected total cost by unilaterally deviating from its Nash policy. The Nash equilibrium plays a critical role in understanding the social dynamics of self-interested agents (Ash, 2000; Axtell, 2002) and constructing the optimal policy of a particular agent via fictitious self-play (Bowling and Veloso, 2000; Ganzfried and Sandholm, 2009). With the recent development in deep learning (LeCun et al., 2015), MARL with function approximation achieves tremendous empirical successes in applications, including Go (Silver et al., 2016, 2017), Poker (Heinrich and Silver, 2016; Moravčík et al., 2017), Star Craft (Vinyals et al., 2019), Dota (OpenAI, 2018), autonomous driving (Shalev-Shwartz et al., 2016), multi-robotic systems (Yang and Gu, 2004), and solving social dilemmas (de Cote et al., 2006; Leibo et al., 2017; Hughes et al., 2018). However, since the capacity of the joint state and action spaces grows exponentially in the number of agents, such MARL approaches become computationally intractable when the number of agents is large, which

is common in real-world applications (Sandholm, 2010; Calderone, 2017; Wang et al., 2017a).

Mean-field game is proposed by Huang et al. (2003, 2006); Lasry and Lions (2006a,b, 2007) with the idea of utilizing mean-field approximation to model the strategic interactions within a large population. In a mean-field game, each agent has the same cost function and state transition, which depend on the other agents only through their aggregated effect. As a result, the optimal policy of each agent depends solely on its own state and the aggregated effect of the population, and such an optimal policy is symmetric across all the agents. Moreover, if the aggregated effect of the population corresponds to the Nash equilibrium, then the optimal policy of each agent jointly constitutes a Nash equilibrium. Although such a Nash equilibrium corresponds to an infinite number of agents, it well approximates the Nash equilibrium for a sufficiently large number of agents (Bensoussan et al., 2016). Also, as the aggregated effect of the population abstracts away the strategic interactions between individual agents, it circumvents the computational intractability of the MARL approaches that do not exploit symmetry.

However, most existing work on mean-field games focuses on characterizing the existence and uniqueness of the Nash equilibrium rather than designing provably efficient algorithms. In particular, most existing work considers the continuous-time setting, which requires solving a pair of Hamilton-Jacobi-Bellman (HJB) and Fokker-Planck (FP) equations, whereas the discrete-time setting is more common in practice, e.g., in the aforementioned applications. Moreover, most existing approaches, including the ones based on solving the HJB and FP equations, require knowing the model of dynamics (Bardi and Priuli, 2014), or having the access to a simulator, which generates the next state given

any state-action pair and aggregated effect of the population (Guo et al., 2019), which is often unavailable in practice.

To address these challenges, we develop an efficient model-free RL approach to mean-field game, which provably attains the Nash equilibrium. In particular, we focus on discrete-time mean-field games with linear state transitions and quadratic cost functions, where the aggregated effect of the population is quantified by the mean-field state. Such games capture the fundamental difficulties of general mean-field games and well approximates a variety of real-world systems such as power grids (Minciardi and Sacile, 2011), swarm robots (Fang, 2014; Araki et al., 2017; Doerr et al., 2018), and financial systems (Zhou and Li, 2000; Huang and Li, 2018). In detail, based on the Nash certainty equivalence (NCE) principle (Huang et al., 2006, 2007), we propose a mean-field actor-critic algorithm which, at each iteration, given the mean-field state μ , approximately attains the optimal policy π_μ^* of each agent, and then updates the mean-field state μ assuming that all the agents follow π_μ^* . We parametrize the actor and critic by linear and quadratic functions, respectively, and prove that such a parameterization encompasses the optimal policy of each agent. Specifically, we update the actor parameter using policy gradient (Sutton et al., 2000) and natural policy gradient (Kakade, 2002; Peters and Schaal, 2008a; Bhatnagar et al., 2009) and update the critic parameter using primal-dual gradient temporal difference (Sutton et al., 2009a,b). In particular, we prove that given the mean-field state μ , the sequence of policies generated by the actor converges linearly to the optimal policy π_μ^* . Moreover, when alternatingly update the policy and mean-field state, we prove that the sequence of policies and its corresponding sequence of mean-field states converge to the unique Nash equilibrium at a linear rate. Our approach can be interpreted

from both “passive” and “active” perspectives: (i) Assuming that each self-interested agent employs the single-agent actor-critic algorithm, the policy of each agent converges to the unique Nash policy, which characterizes the social dynamics of a large population of model-free RL agents. (ii) For a particular agent, our approach serves as a fictitious self-play method for it to find its Nash policy, assuming the other agents give their best responses. To the best of our knowledge, our work establishes the first efficient model-free RL approach with function approximation that provably attains the Nash equilibrium of a discrete-time mean-field game. As a byproduct, we also show that the sequence of policies generated by the single-agent actor-critic algorithm converges at a linear rate to the optimal policy of a linear-quadratic regulator (LQR) problem in the presence of drift, which may be of independent interest.

Related Work. Mean-field game is first introduced in Huang et al. (2003, 2006); Lasry and Lions (2006a,b, 2007). In the last decade, there is growing interest in understanding continuous-time mean-field games. See, e.g., Guéant et al. (2011); Bensoussan et al. (2013); Gomes et al. (2014); Carmona and Delarue (2013, 2018) and the references therein. Due to their simple structures, continuous-time linear-quadratic mean-field games are extensively studied under various model assumptions. See Li and Zhang (2008); Bardi (2011); Wang and Zhang (2012); Bardi and Priuli (2014); Huang et al. (2016a,b); Bensoussan et al. (2016, 2017); Caines and Kizilkale (2017); Huang and Huang (2017); Moon and Başar (2018); Huang and Zhou (2019) for examples of this line of work. Meanwhile, the literature on discrete-time linear-quadratic mean-field games remains relatively scarce. Most of this line of work focuses on characterizing the existence of a Nash equilibrium and the behavior of such a Nash equilibrium when the number of agents goes to infinity

(Gomes et al., 2010; Tembine and Huang, 2011; Moon and Başar, 2014; Biswas, 2015; Saldi et al., 2018a,b, 2019). See also Yang et al. (2018a), which applies maximum entropy inverse RL (Ziebart et al., 2008) to infer the cost function and social dynamics of discrete-time mean-field games with finite state and action spaces. Our work is most related to Guo et al. (2019), where they propose a mean-field Q-learning algorithm (Watkins and Dayan, 1992) for discrete-time mean-field games with finite state and action spaces. Such an algorithm requires the access to a simulator, which, given any state-action pair and mean-field state, outputs the next state. In contrast, both our state and action spaces are infinite, and we do not require such a simulator but only observations of trajectories under given mean-field state. Correspondingly, we study the mean-field actor-critic algorithm with linear function approximation, whereas their algorithm is tailored to the tabular setting. Also, our work is closely related to Mguni et al. (2018), which focuses on a more restrictive setting where the state transition does not involve the mean-field state. In such a setting, mean-field games are potential games, which is, however, not true in more general settings (Li et al., 2017; Briani and Cardaliaguet, 2018). In comparison, we allow the state transition to depend on the mean-field state. Meanwhile, they propose a fictitious self-play method based on the single-agent actor-critic algorithm and establishes its asymptotic convergence. However, their proof of convergence relies on the assumption that the single-agent actor-critic algorithm converges to the optimal policy, which is unverified therein. In addition, our work is related to Jayakumar and Aditya (2019), where the proposed algorithm is only shown to converge asymptotically to a stationary point of the mean-field game.

Our work also extends the line of work on finding the Nash equilibria of Markov games using MARL. Due to the computational intractability introduced by the large number of agents, such a line of work focuses on finite-agent Markov games (Littman, 1994, 2001; Hu and Wellman, 1998; Bowling, 2001; Lagoudakis and Parr, 2002; Hu and Wellman, 2003; Conitzer and Sandholm, 2007; Perolat et al., 2015; Pérolat et al., 2016b,a, 2018; Wei et al., 2017; Zhang et al., 2018; Zou et al., 2019; Casgrain et al., 2019). See also Shoham et al. (2003, 2007); Busoniu et al. (2008); Li (2018) for detailed surveys. Our work is related to Yang et al. (2018b), where they combine the mean-field approximation of actions (rather than states) and Nash Q-learning (Hu and Wellman, 2003) to study general-sum Markov games with a large number of agents. However, the Nash Q-learning algorithm is only applicable to finite state and action spaces, and its convergence is established under rather strong assumptions. Also, when the number of agents goes to infinity, their approach yields a variant of tabular Q-learning, which is different from our mean-field actor-critic algorithm.

For policy optimization, based on the policy gradient theorem, Sutton et al. (2000); Konda and Tsitsiklis (2000) propose the actor-critic algorithm, which is later generalized to the natural actor-critic algorithm (Peters and Schaal, 2008a; Bhatnagar et al., 2009). Most existing results on the convergence of actor-critic algorithms are based on stochastic approximation using ordinary differential equations (Bhatnagar et al., 2009; Castro and Meir, 2010; Konda and Tsitsiklis, 2000; Maei, 2018), which are asymptotic in nature. For policy evaluation, the convergence of primal-dual gradient temporal difference is studied in Liu et al. (2015); Du et al. (2017); Wang et al. (2017b); Yu (2017); Wai et al. (2018). However, this line of work assumes that the feature mapping is bounded, which is not the

case in our setting. Thus, the existing convergence results are not applicable to analyzing the critic update in our setting. To handle the unbounded feature mapping, we utilize a truncation argument, which requires more delicate analysis.

Finally, our work extends the line of work that studies model-free RL for LQR. For example, Bradtke (1993); Bradtke et al. (1994) show that policy iteration converges to the optimal policy, Tu and Recht (2017); Dean et al. (2017) study the sample complexity of least-squares temporal-difference for policy evaluation. More recently, Fazel et al. (2018); Malik et al. (2018); Tu and Recht (2018) show that the policy gradient algorithm converges at a linear rate to the optimal policy. See as also Hardt et al. (2016); Dean et al. (2018) for more in this line of work. Our work is also closely related to Yang et al. (2019b), where they show that the sequence of policies generated by the natural actor-critic algorithm enjoys a linear rate of convergence to the optimal policy. Compared with this work, when fixing the mean-field state, we use the actor-critic algorithm to study LQR in the presence of drift, which introduces significant difficulties in the analysis. As we show in §3.2, the drift causes the optimal policy to have an additional intercept, which makes the state- and action-value functions more complicated.

1.3. Offline Reinforcement Learning with Instrumental Variables in Confounded Markov Decision Processes

Reinforcement learning (RL, Sutton and Barto (2018)) with deep neural networks gains tremendous successes in practice, e.g., games (Silver et al., 2016; OpenAI, 2018), robotics (Kalashnikov et al., 2018), etc. Most RL methods heavily rely on an efficient data generator, e.g., game engines (Bellemare et al., 2013) and physics simulators (Todorov

et al., 2012), which serve as an environment to be interacted with the agent. Since interacting with such an environment is usually either expansive or unethical (e.g., in health-care (Raghu et al., 2017; Komorowski et al., 2018; Gottesman et al., 2019), autonomous driving (Shalev-Shwartz et al., 2016)), RL methods requiring actively collecting data in an online fashion is impractical under this scenario. Thus, a growing body of literature focus on designing RL methods in the offline setting, where the agent aims to learn an optimal policy π^* in the infinite-horizon Markov decision process (MDP, Puterman (2014)) only through observational data, which consists of N trajectories generated by a behavior policy b with a finite horizon T .

However, applying RL methods in the offline setting still possesses the following challenges: (i) The agent may be confounded by unobserved variables (confounders) in the observational data. We refer to the MDP with such confounders as confounded MDP. Such confounders usually comes from private data or heuristic information not recorded (Brookhart et al., 2010). In the confounded MDP, the causal effects of actions on the transitions and rewards are not identifiable from the observational data, leaving most of-line RL methods assuming unconfoundedness not applicable in our setting. (ii) To learn an optimal policy from observational data, many prior methods (Precup, 2000; Antos et al., 2008b; Chen and Jiang, 2019) require a data coverage assumption for any policy π , i.e., the density ratio between the state-action visitation measure induced by π and that induced by the behavior policy b is uniformly upper bounded for any π . However, such a data coverage assumption is hard to satisfy in practice, especially when the state or action spaces are large. Further, many existing methods developed under this assumption are

not stable or even do not converge when the assumption is violated (Wang et al., 2020, 2021c).

To tackle the above challenge (i), we study the confounded MDP via instrumental variables (IVs, Pearl (2009)). Informally, IVs are variables that affect the transitions and rewards only through actions. With the aid of IVs, we introduce two types of identification results: value function (VF)-based and marginalized importance sampling (MIS)-based. Specifically, with only finite-horizon data, VF-based identification result helps identify the state-value function in the infinite-horizon confounded MDP and establish a new Bellman equation by leveraging IVs, which relies on the memoryless assumption on the unmeasured confounders. The memoryless assumption rules out the existence of unmeasured confounders that can affect future rewards and dynamics. On the other hand, we establish the MIS-based identification result for estimating the expected total reward. Interestingly, our MIS-based identification result does not require the memoryless assumption and allow the existence of unmeasured confounders that affects future rewards and dynamics. Therefore our identification result via MIS can be applied in a more general confounded MDP.

In the meanwhile, to tackle the above challenge (ii), we employ pessimism to achieve policy learning. Specifically, when using VF-based identification, we first formulate a minimax estimator of the state-value function via the newly established estimating equation. Then, we construct a confidence set of such a minimax estimator, so that the true state-value function lies within the confidence set with a high probability. Finally, we search for the best policy that maximizes the most conservative expected total reward associated

with the estimated state-value function within the confidence set. As a theoretical contribution, under data coverage assumption only for an optimal policy π^* and realizability assumption for all policies, we show that the suboptimality of the learned best policy is upper bounded by $O(\log(NT)(NT)^{-1/2})$. It is worth noting that our theoretical analysis does not assume that the observational data are generated from stationary distribution or even independent, which is widely imposed in related literature (Farahmand et al., 2016; Nachum et al., 2019; Tang et al., 2019; Xie et al., 2021; Kallus and Uehara, 2022). Without imposing such a restrictive assumption, inspired by Wang et al. (2021a), our convergence analysis relies on novel concentration inequalities for geometrically ergodic sequences, which significantly increases the applicability of our results in practice. In the meanwhile, pessimism with MIS-based identification achieves a similar result by imposing realizability assumption only for an optimal policy π^* and data coverage assumption for all policies. Further, by combining VF- and MIS-based identification results, we propose a doubly robust (DR) estimator of the optimal policy π^* . Theoretically, for such a DR estimator, we show a similar suboptimality at a rate of $O(\log(NT)(NT)^{-1/2})$, but only requiring that either the assumptions imposed in VF-based method or those imposed in MIS-based one hold. See Table 1.1 for an overview of our theoretical results.

Contribution. As a summary, our contribution is three-fold. First, by leveraging IVs, we provide VF- and MIS-based identification results. Second, by employing pessimism, we construct estimators of the optimal policy π^* via VF- and MIS-based identification. Further, by combining VF- and MIS-based identification, we propose a DR-based algorithm for estimating π^* . Third, under mild conditions on data coverage and realizability, we show that the suboptimality of the proposed estimators are upper bounded by

Methods	Data Coverage	Realizability	Identifiability
VF-based	$\ w^{\pi^*}\ _\infty \leq C_*$	$w^{\pi^*} \in \mathcal{W},$ $V^\pi \in \mathcal{V} \forall \pi \in \Pi$	$J(\pi^*)$ is identifiable
MIS-based	$\ w^\pi\ _\infty \leq C_* \forall \pi \in \Pi$	$V^{\pi^*} \in \mathcal{V},$ $w^\pi \in \mathcal{W} \forall \pi \in \Pi$	$J(\pi)$ is identifiable $\forall \pi \in \Pi$
DR-based	$\ w^{\pi^*}\ _\infty \leq C_*, w^{\pi^*} \in \mathcal{W}, V^\pi \in \mathcal{V} \forall \pi \in \Pi,$ and $J(\pi^*)$ is identifiable; OR $\ w^\pi\ _\infty \leq C_* \forall \pi \in \Pi, V^{\pi^*} \in \mathcal{V},$ $w^\pi \in \mathcal{W} \forall \pi \in \Pi,$ and $J(\pi)$ is identifiable $\forall \pi \in \Pi$		

Table 1.1. Assumptions required by our VF-, MIS-, and DR-based methods, where w^π is the density ratio between visitation measures induced by the policy π and the behavior policy b (see (4.2.1) for a detailed definition), and V^π is the state-value function of the policy π . Here \mathcal{V} and Π are function classes, and C_* is a positive absolute constant.

$O(\log(NT)(NT)^{-1/2})$, without requiring that the observational data is generated from stationary distribution or even independent.

Related Work. Our work is related to the line of works that study RL under the presence of unobserved confounders. Zhang and Bareinboim (2019) propose an online RL method to solve dynamic treatment regimes in a finite-horizon setting with the presence of confounded observational data. Their method relies on sensitivity analysis, which constructs a set of possible models based on the confounded observational data to obtain partial identification. Also, to incorporate the observational data into the finite-horizon RL, Wang et al. (2021b) propose deconfounded optimistic value iteration, which is an online algorithm with a provable regret guarantee. To ensure the identifiability through the observational data, they impose the backdoor criterion (Pearl, 2009; Peters et al., 2017) when confounders are partially observed, and also the frontdoor criterion when confounders are unobserved. Our work is also closely related to Liao et al. (2021a), where they propose an IV-aided value iteration algorithm to study confounded MDPs in the offline setting. It is worth mentioning that they only consider the finite-horizon MDP,

where the transition dynamics is a linear function of some known feature map. In the meanwhile, to ensure identifiability, they assume that the unobserved confounders are Gaussian noise, which does not affect the immediate reward and only affect the transition dynamics in an additive manner. In contrast, we consider infinite-horizon confounded MDP without such specific assumptions on the structure of the model, which brings significant technical challenges. With unobserved confounders, Kallus and Zhou (2020) study off-policy evaluation in the infinite-horizon setting based on sensitivity analysis, which imposes additional assumptions on how strong the unobserved confounding can possibly be. Bennett et al. (2021) also study off-policy evaluation in the infinite-horizon MDP via a conditional moment restriction. In the meanwhile, to ensure identifiability, Namkoong et al. (2020) consider the case where the unmeasured confounders affect only one of the decisions made.

Our work is also related to the line of research on policy evaluation and policy learning in the offline setting assuming unconfoundedness. In terms of off-policy evaluation, most works either employs a VF-based method (Ernst et al., 2005; Shi et al., 2020; Liao et al., 2021b; Uehara et al., 2021), or an MIS-based method (Liu et al., 2018; Nachum et al., 2019; Zhang et al., 2020; Wang et al., 2021a; Uehara et al., 2021). Our work is also related to research that propose DR estimators in off-policy evaluation. See Jiang and Li (2016); Thomas and Brunskill (2016); Tang et al. (2019); Kallus and Uehara (2020); Uehara et al. (2020); Kallus and Uehara (2022) for this line of research. As for policy learning in the offline setting, Munos and Szepesvári (2008) and Antos et al. (2008b) prove that fitted value and policy iterations converge to an optimal policy under the data coverage assumption and realizability assumption for all policies, respectively. By employing

pessimism, Xie et al. (2021) guarantee a near-optimal policy under the realizability and the completeness assumptions for all policies, while Jiang and Huang (2020) provide a similar guarantee under the data coverage assumption for the optimal policy and the realizability assumption for all policies. More recently, Zhan et al. (2022) claims that they can learn a near-optimal policy under the data coverage and realizability assumptions for the optimal policy. Their method is built upon a regularized version of the LP formulation of MDPs and thus working on a class of regularized policies. However, due to regularization, the policy learned by Zhan et al. (2022) is typically suboptimal even given infinite data. Moreover, their realizability assumption is imposed on the regularized value function, making it difficult to interpret and compare with other works. In contrast, our work still focuses on a non-regularized setting under standard realizability and data coverage assumptions, even under the presence of unobserved confounders and non-stationary dependent observational data.

CHAPTER 2

Single-Timescale Actor-Critic Provably Finds Globally Optimal Policy

2.1. Background

In this section, we introduce the background on discounted Markov decision processes (MDPs) and actor-critic methods.

Notation. We denote by $[n]$ the set $\{1, 2, \dots, n\}$. For any measure ν and $1 \leq p \leq \infty$, we denote by $\|f\|_{\nu, p} = (\int_{\mathcal{X}} |f(x)|^p d\nu)^{1/p}$ and $\|f\|_p = (\int_{\mathcal{X}} |f(x)|^p d\mu)^{1/p}$, where μ is the Lebesgue measure.

2.1.1. Discounted MDP

A discounted MDP is defined by a tuple $(\mathcal{S}, \mathcal{A}, P, \zeta, \mathcal{R}, \gamma)$. Here \mathcal{S} and \mathcal{A} are the state and action spaces, respectively, $P: \mathcal{S} \times \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ is the Markov transition kernel, $\zeta: \mathcal{S} \rightarrow [0, 1]$ is the initial state distribution, $\mathcal{R}: \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is the deterministic reward function, and $\gamma \in [0, 1)$ is the discount factor. A policy $\pi(a | s)$ measures the probability of taking the action a at the state s . We focus on a family of parameterized policies defined as follows,

$$(2.1.1) \quad \Pi = \{\pi_{\theta}(\cdot | s) \in \mathcal{P}(\mathcal{A}) : s \in \mathcal{S}\},$$

where $\mathcal{P}(\mathcal{A})$ is the probability simplex on the action space \mathcal{A} and θ is the parameter of the policy π_θ . For any state-action pair $(s, a) \in \mathcal{S} \times \mathcal{A}$, we define the action-value function as follows,

$$(2.1.2) \quad Q^\pi(s, a) = (1 - \gamma) \cdot \mathbb{E}_\pi \left[\sum_{t=0}^{\infty} \gamma^t \cdot \mathcal{R}(s_t, a_t) \mid s_0 = s, a_0 = a \right],$$

where $s_{t+1} \sim P(\cdot | s_t, a_t)$ and $a_{t+1} \sim \pi(\cdot | s_{t+1})$ for any $t \geq 0$. We use $\mathbb{E}_\pi[\cdot]$ to denote that the actions follow the policy π , which further affect the transition of the states. We aim to find an optimal policy π^* such that $Q^{\pi^*}(s, a) \geq Q^\pi(s, a)$ for any policy π and state-action pair $(s, a) \in \mathcal{S} \times \mathcal{A}$. That is to say, such an optimal policy π^* attains a higher expected total reward than any other policy π , regardless of the initial state-action pair (s, a) . For notational convenience, we denote by $Q^*(s, a) = Q^{\pi^*}(s, a)$ for any $(s, a) \in \mathcal{S} \times \mathcal{A}$ hereafter.

Meanwhile, we denote by $\nu_\pi(s)$ and $\rho_\pi(s, a) = \nu_\pi(s) \cdot \pi(a | s)$ the stationary state distribution and stationary state-action distribution of the policy π , respectively, for any $(s, a) \in \mathcal{S} \times \mathcal{A}$. Correspondingly, we denote by $\nu^*(s)$ and $\rho^*(s, a)$ the stationary state distribution and stationary state-action distribution of the optimal policy π^* , respectively, for any $(s, a) \in \mathcal{S} \times \mathcal{A}$. For ease of presentation, given any functions $g_1: \mathcal{S} \rightarrow \mathbb{R}$ and $g_2: \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$, we define two operators \mathbb{P} and \mathbb{P}^π as follows,

$$(2.1.3) \quad [\mathbb{P}g_1](s, a) = \mathbb{E}[g_1(s_1) | s_0 = s, a_0 = a], \quad [\mathbb{P}^\pi g_2](s, a) = \mathbb{E}_\pi[g_2(s_1, a_1) | s_0 = s, a_0 = a],$$

where $s_1 \sim P(\cdot | s_0, a_0)$ and $a_1 \sim \pi(\cdot | s_1)$. Intuitively, given the current state-action pair (s_0, a_0) , the operator \mathbb{P} pushes the agent to its next state s_1 following the Markov

transition kernel $P(\cdot | s_0, a_0)$, while the operator \mathbb{P}^π pushes the agent to its next state-action pair (s_1, a_1) following the Markov transition kernel $P(\cdot | s_0, a_0)$ and policy $\pi(\cdot | s_1)$. These operators also relate to the Bellman evaluation operator \mathcal{T}^π , which is defined for any function $g: \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ as follows,

$$(2.1.4) \quad \mathcal{T}^\pi g = (1 - \gamma) \cdot \mathcal{R} + \gamma \cdot \mathbb{P}^\pi g.$$

The Bellman evaluation operator \mathcal{T}^π is used to characterize the actor-critic method in the following section. By the definition in (2.1.2), it is straightforward to verify that the action-value function Q^π is the fixed point of the Bellman evaluation operator \mathcal{T}^π defined in (2.1.4), that is, $Q^\pi = \mathcal{T}^\pi Q^\pi$ for any policy π . For notational convenience, we let \mathbb{P}^ℓ denote the ℓ -fold composition $\underbrace{\mathbb{P} \mathbb{P} \cdots \mathbb{P}}_\ell$. Such notation is also adopted for other linear operators such as \mathbb{P}^π and \mathcal{T}^π .

2.1.2. Actor-Critic Method

To obtain an optimal policy π^* , the actor-critic method (Konda and Tsitsiklis, 2000) aims to maximize the expected total reward as a function of the policy, which is equivalent to solving the following maximization problem,

$$(2.1.5) \quad \max_{\pi \in \Pi} J(\pi) = \mathbb{E}_{s \sim \zeta, a \sim \pi(\cdot | s)} [Q^\pi(s, a)],$$

where ζ is the initial state distribution, Q^π is the action-value function defined in (2.1.2), and the family of parameterized policies Π is defined in (2.1.1). The actor-critic method solves the maximization problem in (2.1.5) via first-order optimization using an estimator of the policy gradient $\nabla_\theta J(\pi)$. Here θ is the parameter of the policy π . In detail, by the

policy gradient theorem (Sutton et al., 2000), we have

$$(2.1.6) \quad \nabla_{\theta} J(\pi) = \mathbb{E}_{(s,a) \sim \varrho_{\pi}} [Q^{\pi}(s, a) \cdot \nabla_{\theta} \log \pi(a | s)].$$

Here ϱ_{π} is the state-action visitation measure of the policy π , which is defined as $\varrho_{\pi}(s, a) = (1 - \gamma) \cdot \sum_{t=0}^{\infty} \gamma^t \cdot \Pr[s_t = s, a_t = a]$. Based on the closed form of the policy gradient in (2.1.6), the actor-critic method consists of the following two parts: (i) the critic update, where a policy evaluation algorithm is invoked to estimate the action-value function Q^{π} , e.g., by applying the Bellman evaluation operator \mathcal{T}^{π} to the current estimator of Q^{π} , and (ii) the actor update, where a policy improvement algorithm, e.g., the policy gradient method, is invoked using the updated estimator of Q^{π} .

In this paper, we consider the following variant of the actor-critic method,

$$(2.1.7) \quad \begin{aligned} \pi_{k+1} &\leftarrow \operatorname{argmax}_{\pi \in \Pi} \mathbb{E}_{\nu_{\pi_k}} [\langle Q_k(s, \cdot), \pi(\cdot | s) \rangle - \beta \cdot \text{KL}(\pi(\cdot | s) \| \pi_k(\cdot | s))], \\ Q_{k+1}(s, a) &\leftarrow \mathbb{E}_{\pi_{k+1}} [(1 - \gamma) \cdot \mathcal{R}(s_0, a_0) + \gamma \cdot Q_k(s_1, a_1) | s_0 = s, a_0 = a], \end{aligned}$$

for any $(s, a) \in \mathcal{S} \times \mathcal{A}$, where $s_1 \sim P(\cdot | s_0, a_0)$, $a_1 \sim \pi_{k+1}(\cdot | s_1)$, and we write $\mathbb{E}_{\nu_{\pi_k}}[\cdot] = \mathbb{E}_{s \sim \nu_{\pi_k}}[\cdot]$ for notational convenience. Here Π is defined in (2.1.1) and $\text{KL}(\pi(\cdot | s) \| \pi_k(\cdot | s))$ is the Kullback-Leibler (KL) divergence between $\pi(\cdot | s)$ and $\pi_k(\cdot | s)$, which is defined for any $s \in \mathcal{S}$ as follows,

$$\text{KL}(\pi(\cdot | s) \| \pi_k(\cdot | s)) = \sum_{a \in \mathcal{A}} \log \left(\frac{\pi(a | s)}{\pi_k(a | s)} \right) \cdot \pi(a | s).$$

In (2.1.7), the actor update uses the proximal policy optimization (PPO) method (Schulman et al., 2017), while the critic update applies the Bellman evaluation operator $\mathcal{T}^{\pi_{k+1}}$

defined in (2.1.4) to Q_k only once, which is the current estimator of the action-value function. Furthermore, we remark that the updates in (2.1.7) provide a general framework in the following two aspects. First, the critic update can be extended to letting $Q_{k+1} \leftarrow (\mathcal{T}^{\pi_{k+1}})^\tau Q_k$ for any fixed $\tau \geq 1$, which corresponds to updating the value function via τ -step rollouts following π_{k+1} . Here we only focus on the case with $\tau = 1$ for simplicity. Our theory can be easily modified for any fixed τ . Moreover, the KL divergence used in the actor step can also be replaced by other Bregman divergences between probability distributions over \mathcal{A} . Second, the actor and critic updates in (2.1.7) is a general template that admits both on- and off-policy evaluation methods and various function approximators in the actor and critic. In the next section, we present an incarnation of (2.1.7) with on-policy sampling and linear and neural network function approximation.

Furthermore, for analyzing the actor-critic method, most existing works (Yang et al., 2019a; Wang et al., 2019; Agarwal et al., 2019; Fu et al., 2019; Liu et al., 2019) rely on (approximately) obtaining $Q^{\pi_{k+1}}$ at each iteration, which is equivalent to applying the Bellman evaluation operator $\mathcal{T}^{\pi_{k+1}}$ infinite times to Q_k . This is usually achieved by minimizing the mean-squared Bellman error $\|Q - \mathcal{T}^{\pi_{k+1}}Q\|_{\rho^{\pi_{k+1}}, 2}^2$ using stochastic semi-gradient descent, e.g., as in the temporal-difference method (Sutton, 1988), to update the critic for sufficiently many iterations. The unique global minimizer of the mean-squared Bellman error gives the action-value function $Q^{\pi_{k+1}}$, which is used in the actor update. Meanwhile, the two-timescale setting is also considered in existing works (Borkar and Konda, 1997; Konda and Tsitsiklis, 2000; Xu et al., 2019, 2020; Wu et al., 2020; Hong et al., 2020), which require the actor to be updated more slowly than the critic in an

asymptotic sense. Such a requirement is usually satisfied by forcing the ratio between the stepsizes of the actor and critic updates to go to zero asymptotically.

In comparison with the setting with bi-level updates, we consider the single-timescale actor and critic updates in (2.1.7), where the critic involves only one step of update, that is, applying the Bellman evaluation operator \mathcal{T}^π to Q_k only once. Meanwhile, in comparison with the two-timescale setting, where the actor and critic are updated simultaneously but with the ratio between their stepsizes asymptotically going to zero, the single-timescale setting is able to achieve a faster rate of convergence by allowing the actor to be updated with a larger stepsize, while updating the critic simultaneously. In particular, such a single-timescale setting better captures a broader range of practical algorithms (Peters and Schaal, 2008a; Schulman et al., 2015; Mnih et al., 2016; Schulman et al., 2017; Haarnoja et al., 2018), where the stepsize of the actor is not asymptotically zero. In §2.2, we discuss the implementation of the updates in (2.1.7) for different schemes of function approximation. In §2.3, we compare the rates of convergence between the two-timescale and single-timescale settings.

2.2. Algorithms

We consider two settings, where the actor and critic are parameterized using linear functions and deep neural networks, respectively. We consider the energy-based policy $\pi_\theta(a|s) \propto \exp(\tau^{-1}f_\theta(s,a))$, where the energy function $f_\theta(s,a)$ is parameterized with the parameter θ . Also, for the (estimated) action-value function, we consider the parameterization $Q_\omega(s,a)$ for any $(s,a) \in \mathcal{S} \times \mathcal{A}$, where ω is the parameter. For such parameterizations of the actor and critic, the updates in (2.1.7) have the following forms.

Actor Update. The following proposition gives the closed form of π_{k+1} in (2.1.7).

Proposition 2.2.1. Let $\pi_{\theta_k}(a | s) \propto \exp(\tau_k^{-1} f_{\theta_k}(s, a))$ be an energy-based policy and

$$\tilde{\pi}_{k+1} = \operatorname{argmax}_{\pi} \mathbb{E}_{\nu_k} [\langle Q_{\omega_k}(s, \cdot), \pi(\cdot | s) \rangle - \beta \cdot \text{KL}(\pi(\cdot | s) \| \pi_{\theta_k}(\cdot | s))].$$

Then $\tilde{\pi}_{k+1}$ has the following closed form,

$$\tilde{\pi}_{k+1}(a | s) \propto \exp(\beta^{-1} Q_{\omega_k}(s, a) + \tau_k^{-1} f_{\theta_k}(s, a)),$$

for any $(s, a) \in \mathcal{S} \times \mathcal{A}$, where $\nu_k = \nu_{\pi_{\theta_k}}$ is the stationary state distribution of π_{θ_k} .

Proof. See §A.5.1 for a detailed proof. □

Motivated by Proposition 2.2.1, to implement the actor update in (2.1.7), we update the actor parameter θ by solving the following minimization problem,

$$(2.2.1) \quad \theta_{k+1} \leftarrow \operatorname{argmin}_{\theta} \mathbb{E}_{\rho_k} [(f_{\theta}(s, a) - \tau_{k+1} \cdot (\beta^{-1} Q_{\omega_k}(s, a) + \tau_k^{-1} f_{\theta_k}(s, a)))^2],$$

where $\rho_k = \rho_{\pi_{\theta_k}}$ is the stationary state-action distribution of π_{θ_k} .

Critic Update. To implement the critic update in (2.1.7), we update the critic parameter ω by solving the following minimization problem,

$$(2.2.2) \quad \omega_{k+1} \leftarrow \operatorname{argmin}_{\omega} \mathbb{E}_{\rho_{k+1}} [(Q_{\omega} - (1 - \gamma) \cdot \mathcal{R} - \gamma \cdot \mathbb{P}^{\pi_{\theta_{k+1}}} Q_{\omega_k})(s, a))^2],$$

where $\rho_{k+1} = \rho_{\pi_{\theta_{k+1}}}$ is the stationary state-action distribution of $\pi_{\theta_{k+1}}$ and the operator \mathbb{P}^{π} is defined in (2.1.3).

2.2.1. Linear Function Approximation

In this section, we consider linear function approximation. More specifically, we parameterize the action-value function using $Q_\omega(s, a) = \omega^\top \varphi(s, a)$ and the energy function of the energy-based policy π_θ using $f_\theta(s, a) = \theta^\top \varphi(s, a)$. Here $\varphi(s, a) \in \mathbb{R}^d$ is the feature vector, where $d > 0$ is the dimension. Without loss of generality, we assume that $\|\varphi(s, a)\|_2 \leq 1$ for any $(s, a) \in \mathcal{S} \times \mathcal{A}$, which can be achieved by normalization.

Actor Update. The minimization problem in (2.2.1) admits the following closed-form solution,

$$(2.2.3) \quad \theta_{k+1} = \tau_{k+1} \cdot (\beta^{-1} \omega_k + \tau_k^{-1} \theta_k),$$

which corresponds to a step of the natural policy gradient method (Kakade, 2002).

Critic Update. The minimization problem in (2.2.2) admits the following closed-form solution,

$$(2.2.4) \quad \tilde{\omega}_{k+1} = (\mathbb{E}_{\rho_{k+1}}[\varphi(s, a)\varphi(s, a)^\top])^{-1} \cdot \mathbb{E}_{\rho_{k+1}}[(1 - \gamma) \cdot \mathcal{R} + \gamma \cdot \mathbb{P}^{\pi_{\theta_{k+1}}} Q_{\omega_k}](s, a) \cdot \varphi(s, a).$$

Since the closed-form solution $\tilde{\omega}_{k+1}$ in (2.2.4) involves the expectation over the stationary state-action distribution ρ_{k+1} of $\pi_{\theta_{k+1}}$, we use data to approximate such an expectation. More specifically, we sample $\{(s_{\ell,1}, a_{\ell,1})\}_{\ell \in [N]}$ and $\{(s_{\ell,2}, a_{\ell,2}, r_{\ell,2}, s'_{\ell,2}, a'_{\ell,2})\}_{\ell \in [N]}$ such that $(s_{\ell,1}, a_{\ell,1}) \sim \rho_{k+1}$, $(s_{\ell,2}, a_{\ell,2}) \sim \rho_{k+1}$, $r_{\ell,2} = \mathcal{R}(s_{\ell,2}, a_{\ell,2})$, $s'_{\ell,2} \sim P(\cdot | s_{\ell,2}, a_{\ell,2})$, and $a'_{\ell,2} \sim \pi_{\theta_{k+1}}(\cdot | s'_{\ell,2})$, where N is the sample size. We approximate $\tilde{\omega}_{k+1}$ using ω_{k+1} , which is

defined as follows,

$$(2.2.5) \quad \omega_{k+1} = \Gamma_R \left\{ \left(\sum_{\ell=1}^N \varphi(s_{\ell,1}, a_{\ell,1}) \varphi(s_{\ell,1}, a_{\ell,1})^\top \right)^{-1} \cdot \sum_{\ell=1}^N \left((1 - \gamma) \cdot r_{\ell,2} + \gamma \cdot Q_{\omega_k}(s'_{\ell,2}, a'_{\ell,2}) \right) \cdot \varphi(s_{\ell,2}, a_{\ell,2}) \right\}.$$

Here Γ_R is the projection operator, which projects the parameter onto the centered ball with radius R in \mathbb{R}^d . Such a projection operator stabilizes the algorithm (Konda and Tsitsiklis, 2000; Bhatnagar et al., 2009). It is worth mentioning that one may also view the update in (2.2.5) as one step of the least-squares temporal difference method (Bradtke and Barto, 1996), which can be modified for the off-policy setting (Antos et al., 2007; Yu, 2010; Liu et al., 2018; Nachum et al., 2019; Xie et al., 2019; Zhang et al., 2020; Uehara and Jiang, 2019; Nachum and Dai, 2020). Such a modification allows the data points in (2.2.5) to be reused in the subsequent iterations, which further improves the sample complexity. Specifically, let $\rho_{\text{bhv}} \in \mathcal{P}(\mathcal{S} \times \mathcal{A})$ be the stationary state-action distribution induced by a behavioral policy π_{bhv} . We replace the actor and critic updates in (2.2.1) and (2.2.2) by

$$(2.2.6) \quad \theta_{k+1} \leftarrow \underset{\theta}{\operatorname{argmin}} \mathbb{E}_{\rho_{\text{bhv}}} \left[\left(f_{\theta}(s, a) - \tau_{k+1} \cdot \left(\beta^{-1} Q_{\omega_k}(s, a) + \tau_k^{-1} f_{\theta_k}(s, a) \right) \right)^2 \right],$$

$$(2.2.7) \quad \omega_{k+1} \leftarrow \underset{\omega}{\operatorname{argmin}} \mathbb{E}_{\rho_{\text{bhv}}} \left[\left([Q_{\omega} - (1 - \gamma) \cdot \mathcal{R} - \gamma \cdot \mathbb{P}^{\pi_{\theta_{k+1}}} Q_{\omega_k}](s, a) \right)^2 \right],$$

respectively. With linear function approximation, the actor update in (2.2.6) is reduced to (2.2.3), while the critic update in (2.2.7) admits a closed form solution

$$\tilde{\omega}_{k+1} = \left(\mathbb{E}_{\rho_{\text{bhv}}} [\varphi(s, a) \varphi(s, a)^\top] \right)^{-1} \cdot \mathbb{E}_{\rho_{\text{bhv}}} \left[[(1 - \gamma) \cdot \mathcal{R} + \gamma \cdot \mathbb{P}^{\pi_{\theta_{k+1}}} Q_{\omega_k}](s, a) \cdot \varphi(s, a) \right],$$

which can be well approximated using state-action pairs drawn from ρ_{bhv} . See §2.3 for a detailed discussion.

Finally, by assembling the updates in (2.2.3) and (2.2.5), we present the linear actor-critic method in Algorithm 1 as follows.

Algorithm 1 Linear Actor-Critic Method

Input: Number of iterations K , sample size N , temperature parameter β .

Initialization: Set $\tau_0 \leftarrow \infty$, and randomly initialize the actor parameter θ_0 and the critic parameter ω_0 .

for $k = 0, 1, 2, \dots, K$ **do**

Actor Update: Update θ_{k+1} via (2.2.3) with $\tau_{k+1}^{-1} = (k + 1) \cdot \beta^{-1}$.

Critic Update: Sample $\{(s_{\ell,1}, a_{\ell,1})\}_{\ell \in [N]}$ and $\{(s_{\ell,2}, a_{\ell,2}, r_{\ell,2}, s'_{\ell,2}, a'_{\ell,2})\}_{\ell \in [N]}$ as specified in §2.2.1. Update ω_{k+1} via (2.2.5).

end for

Output: $\{\pi_{\theta_k}\}_{k \in [K+1]}$, where $\pi_{\theta_k} \propto \exp(\tau_k^{-1} f_{\theta_k})$.

2.2.2. Deep Neural Network Approximation

In this section, we consider deep neural network approximation. We first formally define deep neural networks. Then we introduce the actor-critic method under such a parameterization.

A deep neural network (DNN) $u_\theta(x)$ with the input $x \in \mathbb{R}^d$, depth H , and width m is defined as

$$(2.2.8) \quad x^{(0)} = x, \quad x^{(h)} = \frac{1}{\sqrt{m}} \cdot \sigma(W_h^\top x^{(h-1)}), \text{ for } h \in [H], \quad u_\theta(x) = b^\top x^{(H)}.$$

Here $\sigma: \mathbb{R}^m \rightarrow \mathbb{R}^m$ is the rectified linear unit (ReLU) activation function, which is defined as $\sigma(y) = (\max\{0, y_1\}, \dots, \max\{0, y_m\})^\top$ for any $y = (y_1, \dots, y_m)^\top \in \mathbb{R}^m$. Also, we have $b \in \{-1, 1\}^m$, $W_1 \in \mathbb{R}^{d \times m}$, and $W_h \in \mathbb{R}^{m \times m}$ for $2 \leq h \leq H$. Meanwhile, we denote the parameter of the DNN u_θ as $\theta = (\text{vec}(W_1)^\top, \dots, \text{vec}(W_H)^\top)^\top \in \mathbb{R}^{m_{\text{all}}}$ with $m_{\text{all}} = md + (H - 1)m^2$. We call $\{W_h\}_{h \in [H]}$ the weight matrices of θ . Without loss of generality, we normalize the input x such that $\|x\|_2 = 1$.

We initialize the DNN such that each entry of W_h follows the standard Gaussian distribution $\mathcal{N}(0, 1)$ for any $h \in [H]$, while each entry of b follows the uniform distribution $\text{Unif}(\{-1, 1\})$. Without loss of generality, we fix b during training and only optimize $\{W_h\}_{h \in [H]}$. We denote the initialization of the parameter θ as $\theta_0 = (\text{vec}(W_1^0)^\top, \dots, \text{vec}(W_H^0)^\top)^\top$. Meanwhile, we restrict θ within the ball $\mathcal{B}(\theta_0, R)$ during training, which is defined as follows,

$$(2.2.9) \quad \mathcal{B}(\theta_0, R) = \{\theta \in \mathbb{R}^{m_{\text{all}}} : \|W_h - W_h^0\|_F \leq R, \text{ for } h \in [H]\}.$$

Here $\{W_h\}_{h \in [H]}$ and $\{W_h^0\}_{h \in [H]}$ are the weight matrices of θ and θ_0 , respectively. By (2.2.9), we have $\|\theta - \theta_0\|_2 \leq R\sqrt{H}$ for any $\theta \in \mathcal{B}(\theta_0, R)$. Now, we define the family of DNNs as

$$(2.2.10) \quad \mathcal{U}(m, H, R) = \{u_\theta : \theta \in \mathcal{B}(\theta_0, R)\},$$

where u_θ is a DNN with depth H and width m .

We parameterize the action-value function using $Q_\omega(s, a) \in \mathcal{U}(m_c, H_c, R_c)$ and the energy function of the energy-based policy π_θ using $f_\theta(s, a) \in \mathcal{U}(m_a, H_a, R_a)$. Here $\mathcal{U}(m_c, H_c, R_c)$ and $\mathcal{U}(m_a, H_a, R_a)$ are the families of DNNs defined in (2.2.10). Hereafter we assume that the energy function f_θ and the action-value function Q_ω share the same architecture and initialization, i.e., $m_a = m_c$, $H_a = H_c$, $R_a = R_c$, and $\theta_0 = \omega_0$. Such shared architecture and initialization of the DNNs ensure that the parameterizations of the policy and the action-value function are approximately compatible. See Sutton et al. (2000); Konda and Tsitsiklis (2000); Kakade (2002); Peters and Schaal (2008a); Wang et al. (2019) for a detailed discussion.

Actor Update. To solve (2.2.1), we use projected stochastic gradient descent, whose n -th iteration has the following form,

$$\begin{aligned} \theta(n+1) &\leftarrow \Gamma_{\mathcal{B}(\theta_0, R_a)}(\theta(n)) \\ &\quad - \alpha \cdot (f_{\theta(n)}(s, a) - \tau_{k+1} \cdot (\beta^{-1} Q_{\omega_k}(s, a) + \tau_k^{-1} f_{\theta_k}(s, a))) \cdot \nabla_\theta f_{\theta(n)}(s, a). \end{aligned}$$

Here $\Gamma_{\mathcal{B}(\theta_0, R_a)}$ is the projection operator, which projects the parameter onto the ball $\mathcal{B}(\theta_0, R_a)$ defined in (2.2.9). The state-action pair (s, a) is sampled from the stationary state-action distribution ρ_k . We summarize the update in Algorithm 5, which is deferred to §A.1 of the appendix.

Critic Update. To solve (2.2.2), we apply projected stochastic gradient descent. More specifically, at the n -th iteration of projected stochastic gradient descent, we sample a tuple (s, a, r, s', a') , where $(s, a) \sim \rho_{k+1}$, $r = \mathcal{R}(s, a)$, $s' \sim P(\cdot | s, a)$, and $a' \sim \pi_{\theta_{k+1}}(\cdot | s')$.

We define the residual at the n -th iteration as $\delta(n) = Q_{\omega(n)}(s, a) - (1 - \gamma) \cdot r - \gamma \cdot Q_{\omega_k}(s', a')$.

Then the n -th iteration of projected stochastic gradient descent has the following form,

$$\omega(n + 1) \leftarrow \Gamma_{\mathcal{B}(\omega_0, R_c)}(\omega(n) - \eta \cdot \delta(n) \cdot \nabla_{\omega} Q_{\omega(n)}(s, a)).$$

Here $\Gamma_{\mathcal{B}(\omega_0, R_c)}$ is the projection operator, which projects the parameter onto the ball $\mathcal{B}(\omega_0, R_c)$ defined in (2.2.9). We summarize the update in Algorithm 6, which is deferred to §A.1 of the appendix.

By assembling Algorithms 5 and 6, we present the deep neural actor-critic method in Algorithm 4, which is deferred to §A.1 of the appendix.

Finally, we remark that the off-policy actor and critic updates given in (2.2.6) and (2.2.7) can also incorporate deep neural network approximation with a slight modification, which enables data reuse in the algorithm.

2.3. Theoretical Results

In this section, we upper bound the regret of the linear actor-critic method. We defer the analysis of the deep neural actor-critic method to §A.2 of the appendix. Hereafter we assume that $|\mathcal{R}(s, a)| \leq \mathcal{R}_{\max}$ for any $(s, a) \in \mathcal{S} \times \mathcal{A}$, where \mathcal{R}_{\max} is a positive absolute constant. First, we impose the following assumptions. Recall that ρ^* is the stationary state-action distribution of π^* , while ρ_k is the stationary state-action distribution of π_{θ_k} . Moreover, let $\rho \in \mathcal{P}(\mathcal{S} \times \mathcal{A})$ be a state-action distribution with respect to which we aim to characterize the performance of the actor-critic algorithm. Specifically, after $K + 1$ actor

updates, we are interest in upper bounding the following regret

$$(2.3.1) \quad \mathbb{E} \left[\sum_{k=0}^K (\|Q^* - Q^{\pi_{\theta_{k+1}}}\|_{\rho,1}) \right] = \mathbb{E} \left[\sum_{k=0}^K (Q^*(s, a) - Q^{\pi_{\theta_{k+1}}}(s, a)) \right],$$

where the expectation is taken with respect to $\{\theta_k\}_{k \in [K+1]}$ and $(s, a) \sim \rho$. Here we allow ρ to be any fixed distribution for generality, which might be different from ρ^* .

Assumption 2.3.1 (Concentrability Coefficient). The following statements hold.

- (i) There exists a positive absolute constant ϕ^* such that $\phi_k^* \leq \phi^*$ for any $k \geq 1$, where $\phi_k^* = \|\mathrm{d}\rho^*/\mathrm{d}\rho_k\|_{\rho_k,2}$.
- (ii) We assume that for any $k \geq 1$ and a sequence of policies $\{\pi_i\}_{i \geq 1}$, the k -step future-state-action distribution $\rho^{\mathbb{P}^{\pi_1} \cdots \mathbb{P}^{\pi_k}}$ is absolutely continuous with respect to ρ^* , where ρ is the same as the one in (2.3.1) Also, it holds for such ρ that

$$C_{\rho, \rho^*} = (1 - \gamma)^2 \sum_{k=1}^{\infty} k^2 \gamma^k \cdot c(k) < \infty,$$

where $c(k) = \sup_{\{\pi_i\}_{i \in [k]}} \|\mathrm{d}(\rho^{\mathbb{P}^{\pi_1} \cdots \mathbb{P}^{\pi_k}})/\mathrm{d}\rho^*\|_{\rho^*, \infty}$.

In Assumption 2.3.1, C_{ρ, ρ^*} is known as the discounted-average concentrability coefficient of the future-state-action distributions. Similar assumptions are commonly imposed in the literature (Szepesvári and Munos, 2005; Munos and Szepesvári, 2008; Antos et al., 2008a,b; Scherrer, 2013; Scherrer et al., 2015; Farahmand et al., 2016; Yang et al., 2019c; Geist et al., 2019; Chen and Jiang, 2019).

Assumption 2.3.2 (Zero Approximation Error). It holds for any $\omega, \theta \in \mathcal{B}(0, R)$ that

$$\inf_{\bar{\omega} \in \mathcal{B}(0, R)} \mathbb{E}_{\rho_{\pi_{\theta}}} \left[\left([\mathcal{T}^{\pi_{\theta}} Q_{\omega} - \bar{\omega}^{\top} \varphi](s, a) \right)^2 \right] = 0,$$

where \mathcal{T}^{π_θ} is defined in (2.1.4).

Assumption 2.3.2 states that the Bellman evaluation operator maps a linear function to a linear function. Such an assumption only aims to simplify the presentation of our results. If the approximation error is nonzero, we only need to incorporate an additional bias term into the rate of convergence.

Assumption 2.3.3 (Well-Conditioned Feature). The minimum singular value of the matrix $\mathbb{E}_{\rho_k}[\varphi(s, a)\varphi(s, a)^\top]$ is uniformly lower bounded by a positive absolute constant σ^* for any $k \geq 1$.

Assumption 2.3.3 ensures that the minimization problem in (2.2.2) admits a unique minimizer, which is used in the critic update. Similar assumptions are commonly imposed in the literature (Bhandari et al., 2018; Zou et al., 2019).

Under Assumptions 2.3.1, 2.3.2, and 2.3.3, we upper bound the regret of Algorithm 1 in the following theorem.

Theorem 2.3.4. We assume that Assumptions 2.3.1, 2.3.2, and 2.3.3 hold. Let ρ be a state-action distribution satisfying (ii) of Assumption 2.3.1. Also, for any confidence parameter $\delta \in (0, 1)$ and sufficiently large number of iterations $K > 0$, let $\beta = K^{1/2}$, $N = \Omega(KC_{\rho, \rho^*}^2 \cdot (\phi^*/\sigma^*)^2 \cdot \log^2(KN/\delta))$, and the sequence of policy parameters $\{\theta_k\}_{k \in [K+1]}$ be generated by Algorithm 1. It holds with probability at least $1 - \delta$ that

$$(2.3.2) \quad \mathbb{E}_\rho \left[\sum_{k=0}^K (Q^*(s, a) - Q^{\pi_{\theta_{k+1}}}(s, a)) \right] \leq (2(1 - \gamma)^{-3} \cdot \log |\mathcal{A}| + O(1)) \cdot K^{1/2},$$

where the expectation is taken with respect $(s, a) \sim \rho$.

Proof. We sketch the proof in §2.4. See §A.3.1 for a detailed proof. \square

Theorem 2.3.4 establishes an $O(K^{1/2})$ regret of Algorithm 1, where K is the total number of iterations. Here $O(\cdot)$ omits terms involving $(1 - \gamma)^{-1}$ and $\log |\mathcal{A}|$. To better understand Theorem 2.3.4, we consider the ideal setting, where we have access to the action-value function Q^π of any policy π . In such an ideal setting, the critic update is unnecessary. However, the natural policy gradient method, which only uses the actor update, achieves the same $O(K^{1/2})$ regret (Liu et al., 2019; Agarwal et al., 2019; Cai et al., 2019). In other words, in terms of the iteration complexity, Theorem 2.3.4 shows that in the single-timescale setting, using only one step of the critic update along with one step of the actor update is as efficient as the natural policy gradient method in the ideal setting.

Furthermore, by the regret bound in (2.3.2), to obtain an ε -globally optimal policy, it suffices to set $K \asymp (1 - \gamma)^{-6} \cdot \varepsilon^{-2} \cdot \log^2 |\mathcal{A}|$ in Algorithm 1 and output a randomized policy that is drawn from $\{\pi_{\theta_k}\}_{k=1}^{K+1}$ uniformly. Plugging such a K into $N = \Omega(KC_{\rho, \rho^*}^2 (\phi^*/\sigma^*)^2 \cdot \log^2(KN/\delta))$, we obtain that $N = \tilde{O}(\varepsilon^{-2})$, where $\tilde{O}(\cdot)$ omits the logarithmic terms. Thus, to achieve an ε -globally optimal policy, the total sample complexity of Algorithm 1 is $\tilde{O}(\varepsilon^{-4})$. This matches the sample complexity results established in Xu et al. (2020); Hong et al. (2020) for two-timescale actor-critic methods. Meanwhile, notice that here the critic updates are on-policy and we draw N new data points in each critic update. As discussed in §2.2.1, under the off-policy setting, the critic updates given in (2.2.7) can be implemented using a fixed dataset sampled from ρ_{bhv} , the stationary state-action distribution induced by the behavioral policy. Under this scenario, the total number of data points used by the algorithm is equal to N . Moreover, by imposing similar assumptions on ρ_{bhv} as in (i) of Assumption 2.3.1 and Assumption 2.3.3, we can establish

a similar $O(K^{1/2})$ regret as in (2.3.2) for the off-policy setting. As a result, with data reuse, to obtain an ε -globally optimal policy, the sample complexity of Algorithm 1 is essentially $\tilde{O}(\varepsilon^{-2})$, which demonstrates the advantage of our single-timescale actor-critic method. Besides, only focusing on the convergence to an ε -stationary point, Wu et al. (2020); Xu et al. (2020) establish the sample complexity of $\tilde{O}(\varepsilon^{-5/2})$ for two-timescale actor-critic, where ε measures the squared Euclidean norm of the policy gradient. In contrast, by adopting the natural policy gradient (Kakade, 2002) in actor updates, we achieve convergence to the globally optimal policy. To the best of our knowledge, we establish the rate of convergence and global optimality of the actor-critic method with function approximation in the single-timescale setting for the first time.

Furthermore, as we will show in Theorem A.2.5 of §A.1, when both the actor and the critic are represented using overparameterized deep neural networks, we establish a similar $O((1 - \gamma)^{-3} \cdot \log |\mathcal{A}| \cdot K^{1/2})$ regret when the architecture of the actor and critic neural networks are properly chosen. To our best knowledge, this seems the first theoretical guarantee for the actor-critic method with deep neural network function approximation in terms of the rate of convergence and global optimality.

2.4. Proof Sketch of Theorem 2.3.4

In this section, we sketch the proof of Theorem 2.3.4. Recall that ρ is a state-action distribution satisfying (ii) of Assumption 2.3.1. We first upper bound $\sum_{k=0}^K (Q^*(s, a) - Q^{\pi_{\theta_{k+1}}}(s, a))$ for any $(s, a) \in \mathcal{S} \times \mathcal{A}$ in part 1. Then by further taking the expectation over ρ in part 2, we conclude the proof of Theorem 2.3.4. See §A.3.1 for a detailed proof.

Part 1. In the sequel, we upper bound $\sum_{k=0}^K (Q^*(s, a) - Q^{\pi_{\theta_{k+1}}}(s, a))$ for any $(s, a) \in \mathcal{S} \times \mathcal{A}$.

We first decompose $Q^* - Q^{\pi_{\theta_{k+1}}}$ into the following three terms,

$$(2.4.1) \quad \sum_{k=0}^K [Q^* - Q^{\pi_{\theta_{k+1}}}] (s, a) = \sum_{k=0}^K [(I - \gamma \mathbb{P}^{\pi^*})^{-1} (A_{1,k} + A_{2,k} + A_{3,k})] (s, a),$$

the proof of which is deferred to (A.3.1) and (A.3.2) in §A.3.1 of the appendix. Here the operator \mathbb{P}^{π^*} is defined in (2.1.3), $(I - \gamma \mathbb{P}^{\pi^*})^{-1} = \sum_{i=0}^{\infty} (\gamma \mathbb{P}^{\pi^*})^i$, and $A_{1,k}$, $A_{2,k}$, and $A_{3,k}$ are defined as follows,

$$(2.4.2) \quad A_{1,k}(s, a) = [\gamma (\mathbb{P}^{\pi^*} - \mathbb{P}^{\pi_{\theta_{k+1}}}) Q_{\omega_k}] (s, a),$$

$$(2.4.3) \quad A_{2,k}(s, a) = [\gamma \mathbb{P}^{\pi^*} (Q^{\pi_{\theta_{k+1}}} - Q_{\omega_k})] (s, a),$$

$$(2.4.4) \quad A_{3,k}(s, a) = [\mathcal{T}^{\pi_{\theta_{k+1}}} Q_{\omega_k} - Q^{\pi_{\theta_{k+1}}}] (s, a).$$

To understand the intuition behind $A_{1,k}$, $A_{2,k}$, and $A_{3,k}$, we interpret them as follows.

Interpretation of $A_{1,k}$. As defined in (2.4.2), $A_{1,k}$ arises from the actor update and measures the convergence of the policy $\pi_{\theta_{k+1}}$ towards a globally optimal policy π^* , which implies the convergence of $\mathbb{P}^{\pi_{\theta_{k+1}}}$ towards \mathbb{P}^{π^*} .

Interpretation of $A_{3,k}$. Note that by (2.1.2) and (2.1.4), we have $Q^{\pi_{\theta_{k+1}}} = \mathcal{T}^{\pi_{\theta_{k+1}}} Q^{\pi_{\theta_{k+1}}}$ and $\mathcal{T}^{\pi_{\theta_{k+1}}}$ is a γ -contraction, which implies that applying the Bellman evaluation operator $\mathcal{T}^{\pi_{\theta_{k+1}}}$ to any Q , e.g., Q_{ω_k} , infinite times yields $Q^{\pi_{\theta_{k+1}}}$. As defined in (2.4.4), $A_{3,k}$ measures the error of tracking the action-value function $Q^{\pi_{\theta_{k+1}}}$ of $\pi_{\theta_{k+1}}$ by applying the Bellman evaluation operator $\mathcal{T}^{\pi_{\theta_{k+1}}}$ to Q_{ω_k} only once, which arises from the critic update. Also, as $A_{3,k} = \mathcal{T}^{\pi_{\theta_{k+1}}}(Q_{\omega_k} - Q^{\pi_{\theta_{k+1}}})$, $A_{3,k}$ measures the difference between $Q^{\pi_{\theta_k}}$, which is approximated by Q_{ω_k} as discussed subsequently, and $Q^{\pi_{\theta_{k+1}}}$. Such a difference can also

be viewed as the difference between π_{θ_k} and $\pi_{\theta_{k+1}}$, which arises from the actor update. Therefore, the convergence of $A_{3,k}$ to zero implies the contractions of not only the critic update but also the actor update, which illustrates the “double contraction” phenomenon. We establish the convergence of $A_{3,k}$ to zero in (2.4.10) subsequently.

Interpretation of $A_{2,k}$. Assuming that $A_{3,k-1}$ converges to zero, we have $\mathcal{T}^{\pi_{\theta_k}} Q_{\omega_{k-1}} \approx Q^{\pi_{\theta_k}}$. Moreover, assuming that the number of data points N is sufficiently large and ignoring the projection in (2.2.5), we have $\mathcal{T}^{\pi_{\theta_k}} Q_{\omega_{k-1}} = Q_{\tilde{\omega}_k} \approx Q_{\omega_k}$ as $\tilde{\omega}_k$ defined in (2.2.4) is an estimator of ω_k . Hence, we have $Q^{\pi_{\theta_k}} \approx Q_{\omega_k}$. Such an approximation error is characterized by ϵ_k^c defined in (2.4.5) subsequently. Hence, $A_{2,k}$ measures the difference between π_{θ_k} and $\pi_{\theta_{k+1}}$ through the difference between $Q^{\pi_{\theta_k}} \approx Q_{\omega_k}$ and $Q^{\pi_{\theta_{k+1}}}$, which relies on the convergence of $A_{3,k-1}$ to zero.

In the sequel, we upper bound $A_{1,k}$, $A_{2,k}$, and $A_{3,k}$, respectively. To establish such upper bounds, we define the following quantities,

$$(2.4.5) \quad \epsilon_{k+1}^c(s, a) = [\mathcal{T}^{\pi_{\theta_{k+1}}} Q_{\omega_k} - Q_{\omega_{k+1}}](s, a),$$

$$(2.4.6) \quad e_{k+1}(s, a) = [Q_{\omega_k} - \mathcal{T}^{\pi_{\theta_{k+1}}} Q_{\omega_k}](s, a),$$

$$(2.4.7) \quad \vartheta_k(s) = \text{KL}(\pi^*(\cdot | s) \| \pi_{\theta_k}(\cdot | s)) - \text{KL}(\pi^*(\cdot | s) \| \pi_{\theta_{k+1}}(\cdot | s)).$$

To understand the intuition behind ϵ_{k+1}^c , e_{k+1} , and ϑ_k , we interpret them as follows.

Interpretation of ϵ_{k+1}^c . Recall that $\tilde{\omega}_{k+1}$ is defined in (2.2.4), which parameterizes $\mathcal{T}^{\pi_{\theta_{k+1}}} Q_{\omega_k}$ (ignoring the projection in (2.2.5)). Here ϵ_{k+1}^c arises from approximating $\tilde{\omega}_{k+1}$ using ω_{k+1} as an estimator, which is constructed based on ω_k and the N data points. In




Figure 2.1. Illustration of the relationship among $A_{1,k}$, $A_{2,k}$, $A_{3,k}$, ϵ_{k+1}^c , e_{k+1} , and ϑ_k . Here $\{\theta_k, \omega_k\}$ and $\{\theta_{k+1}, \omega_{k+1}\}$ are two consecutive iterates of actor-critic. The red arrow from Q_{ω_k} to $Q_{\omega_{k+1}}$ represents the critic update and the red arrow from $Q^{\pi_{\theta_k}}$ to $Q^{\pi_{\theta_{k+1}}}$ represents the action-value functions associated with the two policies in any actor update. Here ϑ_k given in (2.4.7) quantifies the difference between π_{θ_k} and $\pi_{\theta_{k+1}}$ in terms of their KL distances to π^* . In addition, the cyan arrows represent quantities $A_{1,k}$, $A_{2,k}$, and $A_{3,k}$ introduced in (2.4.2)–(2.4.4), which are intermediate terms used for analyzing the error $Q^* - Q^{\pi_{k+1}}$. Finally, the blue arrows represent ϵ_{k+1}^c and e_{k+1} defined in (2.4.5) and (2.4.6), respectively. Here ϵ_{k+1}^c corresponds to the statistical error due to having finite data whereas e_{k+1} essentially quantifies the difference between π_{θ_k} and $\pi_{\theta_{k+1}}$.

particular, ϵ_{k+1}^c decreases to zero as $N \rightarrow \infty$, which is used in characterizing $A_{2,k}$ defined in (2.4.3).

Interpretation of e_{k+1} . Assuming that $A_{3,k-1}$ defined in (2.4.4) and ϵ_k^c defined in (2.4.5) converge to zero, which implies $\mathcal{T}^{\pi_{\theta_k}} Q_{\omega_{k-1}} \approx Q^{\pi_{\theta_k}}$ and $\mathcal{T}^{\pi_{\theta_k}} Q_{\omega_{k-1}} \approx Q_{\omega_k}$, respectively, we have $Q_{\omega_k} \approx Q^{\pi_{\theta_k}}$. Therefore, as defined in (2.4.6), $e_{k+1} = Q_{\omega_k} - \mathcal{T}^{\pi_{\theta_{k+1}}} Q_{\omega_k} \approx Q^{\pi_{\theta_k}} -$

$\mathcal{T}^{\pi_{\theta_{k+1}}}Q^{\pi_{\theta_k}} = (\mathcal{T}^{\pi_{\theta_k}} - \mathcal{T}^{\pi_{\theta_{k+1}}})Q^{\pi_{\theta_k}}$ measures the difference between π_{θ_k} and $\pi_{\theta_{k+1}}$, which implies the difference between $\mathcal{T}^{\pi_{\theta_k}}$ and $\mathcal{T}^{\pi_{\theta_{k+1}}}$. We remark that e_{k+1} fully characterizes $A_{3,k}$ defined in (2.4.4) as shown in (2.4.8) subsequently.

Interpretation of ϑ_k . As defined in (2.4.7), ϑ_k measures the difference between π_{θ_k} and $\pi_{\theta_{k+1}}$ in terms of their differences with π^* , which are measured by the corresponding KL-divergences. In particular, ϑ_k is used in characterizing $A_{1,k}$ and $A_{2,k}$ defined in (2.4.2) and (2.4.3), respectively.

We remark that ϵ_{k+1}^c measures the statistical error in the critic update, while ϑ_k measures the optimization error in the actor update. As discussed above, the convergence of $A_{3,k}$ to zero implies the contraction of both the actor update and the critic update, which illustrates the “double contraction” phenomenon. Meanwhile, since e_{k+1} fully characterizes $A_{3,k}$ as shown in (2.4.8) subsequently, e_{k+1} plays a key role in the “double contraction” phenomenon. In particular, the convergence of e_{k+1} to zero is established in (2.4.9) subsequently. See Figure 2.1 for an illustration of these quantities.

With the quantities defined in (2.4.5), (2.4.6), and (2.4.7), we upper bound $A_{1,k}$, $A_{2,k}$, and $A_{3,k}$ as follows,

$$\begin{aligned}
A_{1,k}(s, a) &\leq \gamma\beta \cdot [\mathbb{P}\vartheta_k](s, a), \\
A_{2,k}(s, a) &\leq [(\gamma\mathbb{P}^{\pi^*})^{k+1}(Q^* - Q_{\omega_0})](s, a) + \gamma\beta \cdot \sum_{i=0}^{k-1} [(\gamma\mathbb{P}^{\pi^*})^{k-i}\mathbb{P}\vartheta_i](s, a) \\
&\quad + \sum_{i=0}^{k-1} [(\gamma\mathbb{P}^{\pi^*})^{k-i}\epsilon_{i+1}^c](s, a), \\
(2.4.8) \quad A_{3,k}(s, a) &= [\gamma\mathbb{P}^{\pi_{\theta_{k+1}}}(I - \gamma\mathbb{P}^{\pi_{\theta_{k+1}}})^{-1}e_{k+1}](s, a),
\end{aligned}$$

the proof of which is deferred to Lemmas A.3.1, A.3.2, and A.3.3 in §A.3.1 of the appendix, respectively. Meanwhile, by recursively expanding (2.4.5) and (2.4.6), we have

$$(2.4.9) \quad e_{k+1}(s, a) \leq \left[\gamma^k \left(\prod_{s=1}^k \mathbb{P}^{\pi_{\theta_s}} \right) e_1 + \sum_{i=1}^k \gamma^{k-i} \left(\prod_{s=i+1}^k \mathbb{P}^{\pi_{\theta_s}} \right) (I - \gamma \mathbb{P}^{\pi_{\theta_i}}) \epsilon_i^c \right] (s, a),$$

the proof of which is deferred to Lemma A.3.4 in §A.3.1 of the appendix. By plugging (2.4.9) into (2.4.8), we have

$$(2.4.10) \quad A_{3,k}(s, a) \leq \left[\gamma \mathbb{P}^{\pi_{\theta_{k+1}}} (I - \gamma \mathbb{P}^{\pi_{\theta_{k+1}}})^{-1} \left(\gamma^k \left(\prod_{s=1}^k \mathbb{P}^{\pi_{\theta_s}} \right) e_1 + \sum_{i=1}^k \gamma^{k-i} \left(\prod_{s=i+1}^k \mathbb{P}^{\pi_{\theta_s}} \right) (I - \gamma \mathbb{P}^{\pi_{\theta_i}}) \epsilon_i^c \right) \right] (s, a).$$

To better understand (2.4.10) and how it relates to the convergence of $A_{3,k}$, $A_{2,k}$, and $A_{1,k}$ to zero, we discuss in the following two steps.

Step (i). We assume $\epsilon_i^c = 0$, which corresponds to the number of data points $N \rightarrow \infty$. Then (2.4.10) yields $A_{3,k} = O(\gamma^k)$, which implies that $A_{3,k}$ defined in (2.4.4) converges to zero driven by the discount factor γ . As discussed above, the convergence of $A_{3,k}$ to zero also implies the contraction between π_{θ_k} and $\pi_{\theta_{k+1}}$ of the actor update and the contraction between Q_{ω_k} and $Q^{\pi_{\theta_k}}$ of the critic update, which illustrates the “double contraction” phenomenon.

Step (ii). The convergence of $A_{3,k}$ to zero further ensures that $A_{2,k}$ converges to zero. To see this, we further assume $A_{3,k} = 0$, which together with the assumption that $\epsilon_{k+1}^c = 0$

implies $Q^{\pi_{k+1}} = \mathcal{T}^{\pi_{k+1}} Q_{\omega_k} = Q_{\omega_{k+1}}$ by their definitions in (2.4.4) and (2.4.5), respectively. Then by telescoping the sum of $A_{2,k}$ defined in (2.4.3), which cancels out $Q_{\omega_{k+1}}$ and $Q^{\pi_{k+1}}$, we obtain the convergence of $A_{2,k}$ to zero. Meanwhile, telescoping the sum of $A_{1,k}$ defined in (2.4.2) and the sum of its upper bound in (2.4.8) implies that $A_{1,k}$ converges to zero.

Now, by plugging (2.4.8) and (2.4.10) into (2.4.1), we establish an upper bound of $\sum_{k=0}^K (Q^*(s, a) - Q^{\pi_{k+1}}(s, a))$ for any $(s, a) \in \mathcal{S} \times \mathcal{A}$, which is deferred to (A.3.12) in §A.3.1 of the appendix. Hence, we conclude the proof in part 1. See part 1 of §A.3.1 for details.

Part 2. Recall that ρ is a state-action distribution satisfying (ii) of Assumption 2.3.1. In the sequel, we take the expectation over ρ in (A.3.12) and upper bound each term. We first introduce the following lemma, which upper bounds ϵ_{k+1}^c defined in (2.4.5).

Lemma 2.4.1. Under Assumptions 2.3.2 and 2.3.3, with probability at least $1 - \delta$, it holds for any $k \in \{0, 1, \dots, K\}$ that

$$\begin{aligned} \mathbb{E}_{\rho_{k+1}} [\epsilon_{k+1}^c(s, a)^2] &= \mathbb{E}[(Q_{\omega_{k+1}}(s, a) - [\mathcal{T}^{\pi_{k+1}} Q_{\omega_k}](s, a))^2] \\ &\leq \frac{32(\mathcal{R}_{\max} + R)^2}{N(\sigma^*)^4} \cdot \log^2(NK/p + dK/p), \end{aligned}$$

where the expectation is taken with respect to $(s, a) \sim \rho_{k+1}$.

Proof. See §A.6.1 for a detailed proof. □

On the right-hand side of (A.3.12) in §A.3.1 of the appendix, for the terms not involving ϵ_{k+1}^c , i.e., M_1 , M_2 , and M_3 in (A.3.13), we take the expectation over ρ and establish their upper bounds in the ℓ_∞ -norm over (s, a) in Lemma A.3.5. On the other hand, for

the terms involving ϵ_{k+1}^c , i.e., M_4 and M_5 in (A.3.14), we take the expectation over ρ and then change the measure from ρ to ρ_{k+1} . By Assumption 2.3.1 and Lemma 2.4.1, which relies on ρ_{k+1} , we establish the upper bounds in Lemma A.3.6. See part 2 of §A.3.1 for details.

Combining Lemmas A.3.5 and A.3.6 yields Theorem 2.3.4. See §A.3.1 for a detailed proof.

2.5. Conclusion

In this paper, we analyze the actor-critic method in single-timescale setting with function approximation. We theoretically show that the method achieves an $O(K^{1/2})$ -regret, which appears to be the first success to provide upper bound of regret of the actor-critic method in the single-timescale setting with function approximation. For future research, we aim to extend our analysis to other variations of actor-critic setting, including mean-field-type control tasks and games (Zhang et al., 2019b), risk-sensitive RL tasks (Prashanth and Ghavamzadeh, 2013), and partially observed MDP (Bhatnagar, 2010).

CHAPTER 3

Actor-Critic Provably Finds Nash Equilibria of Linear-Quadratic Mean-Field Games

3.1. Linear-Quadratic Mean-Field Game

A linear-quadratic mean-field N_a -player game involves $N_a \in \mathbb{N}$ agents, whose state transitions are given by

$$x_{t+1}^i = Ax_t^i + Bu_t^i + \bar{A}\bar{x}_t + d^i + \omega_t^i, \quad \forall t \geq 0, i \in [N_a].$$

Here $A \in \mathbb{R}^{m \times m}$, $B \in \mathbb{R}^{m \times k}$, and $\bar{A} \in \mathbb{R}^{m \times m}$ are matrices, $x_t^i \in \mathbb{R}^m$ and $u_t^i \in \mathbb{R}^k$ are the state and action vectors of agent i , respectively, the vector $d^i \in \mathbb{R}^m$ is a drift term, $\omega_t^i \in \mathbb{R}^m$ is an independent random noise term following the Gaussian distribution $\mathcal{N}(0, \Psi_\omega)$, and $\bar{x}_t = 1/N_a \cdot \sum_{j=1}^{N_a} x_t^j$ is the mean-field state. The agents are coupled through the mean-field state \bar{x}_t . In the linear-quadratic mean-field N_a -player game, the cost of agent $i \in [N_a]$ at time $t \geq 0$ is given by

$$c_t^i = (x_t^i)^\top Q x_t^i + (u_t^i)^\top R u_t^i + \bar{x}_t^\top \bar{Q} \bar{x}_t,$$

where $Q \in \mathbb{R}^{m \times m}$, $R \in \mathbb{R}^{k \times k}$, and $\bar{Q} \in \mathbb{R}^{m \times m}$ are matrices, and u_t^i is generated by π^i , i.e., the policy of agent i . To measure the performance of agent i following its policy π^i under

the influence of the other agents, we define the expected total cost of agent i as

$$J^i(\pi^1, \pi^2, \dots, \pi^{N_a}) = \lim_{T \rightarrow \infty} \mathbb{E} \left(\frac{1}{T} \sum_{t=0}^T c_t^i \right).$$

We are interested in finding a Nash equilibrium $(\pi^1, \pi^2, \dots, \pi^{N_a})$, which is defined by

$$J^i(\pi^1, \dots, \pi^{i-1}, \pi^i, \pi^{i+1}, \dots, \pi^{N_a}) \leq J^i(\pi^1, \dots, \pi^{i-1}, \tilde{\pi}^i, \pi^{i+1}, \dots, \pi^{N_a}), \quad \forall \tilde{\pi}^i, i \in [N_a].$$

That is, agent i cannot further decrease its expected total cost by unilaterally deviating from its Nash policy.

For the simplicity of discussion, we assume that the drift term d^i is identical for each agent. We consider taking the infinite-population limit $N_a \rightarrow \infty$, where each agent has an infinitesimal contribution to the dynamics of the system. Thus, the joint policy of all the agents except agent i can be modeled as a mean-field policy π^\dagger , and all the agents following such a mean-field policy π^\dagger generate the mean-field state $\mathbb{E}x_t^\dagger$, where $\{x_t^\dagger\}_{t \geq 0}$ is generated following the policy π^\dagger . By the symmetry of the agents in terms of their state transitions and cost functions, we focus on a fixed agent and drop the superscript i hereafter.

Before we formally present the formulation of linear-quadratic mean-field games, we first introduce the following mean-field LQR (MF-LQR) problem, which aims to find an optimal policy for the fixed agent given the mean-field policy π^\dagger .

Problem 3.1.1 (MF-LQR). Given the mean-field policy π^\dagger , we consider the following formulation,

$$\begin{aligned} x_{t+1} &= Ax_t + Bu_t + \bar{A}\mathbb{E}x_t^\dagger + d + \omega_t, \\ c(x_t, u_t) &= x_t^\top Qx_t + u_t^\top Ru_t + (\mathbb{E}x_t^\dagger)^\top \bar{Q}(\mathbb{E}x_t^\dagger), \\ J(\pi, \pi^\dagger) &= \lim_{T \rightarrow \infty} \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^T c(x_t, u_t) \right], \end{aligned}$$

where $x_t \in \mathbb{R}^m$ is the state vector, $u_t \in \mathbb{R}^k$ is the action vector generated by the policy π , $\{x_t^\dagger\}_{t \geq 0}$ is the trajectory generated by the policy π^\dagger , $\omega_t \in \mathbb{R}^m$ is an independent random noise term following the Gaussian distribution $\mathcal{N}(0, \Psi_\omega)$, and $d \in \mathbb{R}^m$ is a drift term. Here the expectation $\mathbb{E}x_t^\dagger$ is taken across all the agents. We aim to find π^* such that $J(\pi^*, \pi^\dagger) = \inf_{\pi \in \Pi} J(\pi, \pi^\dagger)$.

Note that a controllable linear system using linear quadratic optimal control is always stable. Further combining the fact that our linear closed-loop dynamics in Problem 3.1.1 is driven by the Gaussian noise term ω_t , we know that the Markov chain of states generated by the policy π^\dagger admits a stationary distribution and converges to this stationary distribution. This implies that the mean-field state $\mathbb{E}x_t^\dagger$ converges to a constant vector μ^\dagger as $t \rightarrow \infty$, which serves as a time-invariant mean-field state. As we consider the ergodic setting, it then suffices to study Problem 3.1.1 with t sufficiently large. Therefore, the influence of the mean-field policy π^\dagger is captured by the mean-field state μ^\dagger . By re-formulating Problem 3.1.1, with slight abuse of notations, we obtain the following drifted-LQR (D-LQR).

Problem 3.1.2 (D-LQR). Given a mean-field state $\mu \in \mathbb{R}^m$, we consider the following formulation,

$$\begin{aligned} x_{t+1} &= Ax_t + Bu_t + \bar{A}\mu + d + \omega_t, \\ c_\mu(x_t, u_t) &= x_t^\top Qx_t + u_t^\top Ru_t + \mu^\top \bar{Q}\mu, \\ J_\mu(\pi) &= \lim_{T \rightarrow \infty} \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^T c_\mu(x_t, u_t) \right], \end{aligned}$$

where $x_t \in \mathbb{R}^m$ is the state vector, $u_t \in \mathbb{R}^k$ is the action vector generated by the policy π , $\omega_t \in \mathbb{R}^m$ is an independent random noise term following the Gaussian distribution $\mathcal{N}(0, \Psi_\omega)$, and $d \in \mathbb{R}^m$ is a drift term. We aim to find an optimal policy π_μ^* such that $J_\mu(\pi_\mu^*) = \inf_{\pi \in \Pi} J_\mu(\pi)$.

Compared with the most studied LQR problem (Lewis et al., 2012), both the state transition and the cost function in Problem 3.1.2 have drift terms, which act as the mean-field “force” that drives the states away from zero. Such a mean-field “force” introduces additional challenges when solving Problem 3.1.2 in the model-free setting (see §3.2.3 for details). On the other hand, the unique optimal policy π_μ^* of Problem 3.1.2 admits a linear form $\pi_\mu^*(x_t) = -K_{\pi_\mu^*}x_t + b_{\pi_\mu^*}$ under certain regularity conditions (Anderson and Moore, 2007), where the matrix $K_{\pi_\mu^*} \in \mathbb{R}^{k \times m}$ and the vector $b_{\pi_\mu^*} \in \mathbb{R}^k$ are the parameters of π_μ^* . Motivated by such a linear form of the optimal policy, we define the class of linear-Gaussian policies as

$$(3.1.1) \quad \Pi = \{\pi(x) = -K_\pi x + b_\pi + \sigma \cdot \eta : K_\pi \in \mathbb{R}^{k \times m}, b_\pi \in \mathbb{R}^k\},$$

where $\sigma \in \mathbb{R}$ and the standard Gaussian noise term $\eta \in \mathbb{R}^k$ is included to encourage exploration. To solve Problem 3.1.2, it suffices to find the optimal policy π_μ^* within Π . We define $\Lambda_1(\mu) = \pi_\mu^*$ as the optimal policy under the mean-field state μ .

Assume that all the agents follow the linear policy $\pi(x) = -K_\pi x + b_\pi$ under the mean-field state μ . By plugging $u_t = \pi(x_t)$ into the state transition in Problem 3.1.2, as $t \rightarrow \infty$, we know that these agents generate a new mean-field state μ_{new} such that

$$\mu_{\text{new}} = (I - A + BK_\pi)^{-1}(Bb_\pi + \bar{A}\mu + d).$$

We define $\Lambda_2(\mu, \pi) = \mu_{\text{new}}$ as such a new mean-field state.

Now, we are ready to present the following linear-quadratic mean-field game (LQ-MFG).

Problem 3.1.3 (LQ-MFG). We consider the following formulation,

$$\begin{aligned} x_{t+1} &= Ax_t + Bu_t + \bar{A}\mu + d + \omega_t, \\ c(x_t, u_t) &= x_t^\top Qx_t + u_t^\top Ru_t + \mu^\top \bar{Q}\mu, \\ J(\pi, \mu) &= \lim_{T \rightarrow \infty} \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^T c(x_t, u_t) \right], \end{aligned}$$

where $x_t \in \mathbb{R}^m$ is the state vector, $u_t \in \mathbb{R}^k$ is the action vector generated by the policy π , $\mu \in \mathbb{R}^m$ is the mean-field state, $\omega_t \in \mathbb{R}^m$ is an independent random noise term following the Gaussian distribution $\mathcal{N}(0, \Psi_\omega)$, and $d \in \mathbb{R}^m$ is a drift term. We aim to find a pair (μ^*, π^*) such that (i) $J(\pi^*, \mu^*) = \inf_{\pi \in \Pi} J(\pi, \mu^*)$; (ii) $\mathbb{E}x_t^*$ converges to μ^* as $t \rightarrow \infty$, where $\{x_t^*\}_{t \geq 0}$ is the Markov chain of states generated by the policy π^* .

The formulation in Problem 3.1.3 is studied by Lasry and Lions (2007); Bensoussan et al. (2016); Saldi et al. (2018a,b). We propose a more general formulation in Problem B.3.2 (see §B.3 of the appendix for details), where an additional interaction term between the state vector x_t and the mean-field state μ is incorporated into the cost function. According to our analysis in §B.3, up to minor modification, the results in the following sections also carry over to Problem B.3.2. Therefore, for the sake of simplicity, we focus on Problem 3.1.3 in the sequel.

In Problem 3.1.3, condition (i) is equivalent to the optimality of the policy π^* under the mean-field state μ^* , namely, $\Lambda_1(\mu^*) = \pi^*$. Meanwhile, condition (ii) is equivalent to the invariance of the mean-field state μ^* given the policy π^* , namely, $\Lambda_2(\mu^*, \pi^*) = \mu^*$. Such equivalence follows from the NCE principle (Huang et al., 2006, 2007), which also motivates the following definition of the Nash equilibrium pair (Saldi et al., 2018a,b).

Definition 3.1.4 (Nash Equilibrium Pair). The pair $(\mu^*, \pi^*) \in \mathbb{R}^m \times \Pi$ constitutes a Nash equilibrium pair of Problem 3.1.3 if it satisfies $\pi^* = \Lambda_1(\mu^*)$ and $\mu^* = \Lambda_2(\mu^*, \pi^*)$. Here μ^* is called the Nash mean-field state and π^* is called the Nash policy.

By Definition 3.1.4, Problem 3.1.3 aims to find a Nash equilibrium pair (μ^*, π^*) .

Notations. We denote by $\|M\|_*$ the spectral norm, $\rho(M)$ the spectral radius, $\sigma_{\min}(M)$ the minimum singular value, and $\sigma_{\max}(M)$ the maximum singular value of a matrix M . We use $\|\alpha\|_2$ to represent the ℓ_2 -norm of a vector α , and $(\alpha)_i^j$ to denote the sub-vector $(\alpha_i, \alpha_{i+1}, \dots, \alpha_j)^\top$, where α_k is the k -th entry of the vector α . For scalars a_1, \dots, a_n , we denote by $\text{poly}(a_1, \dots, a_n)$ the polynomial of a_1, \dots, a_n , and this polynomial may vary from line to line. We use $[n]$ to denote the set $\{1, 2, \dots, n\}$ for any $n \in \mathbb{N}$.

3.2. Mean-Field Actor-Critic

We first characterize the existence and uniqueness of the Nash equilibrium pair of Problem 3.1.3 under mild regularity conditions, and then propose a mean-field actor-critic algorithm to obtain such a Nash equilibrium. As a building block of the mean-field actor-critic, we propose the natural actor-critic to solve Problem 3.1.2.

3.2.1. Existence and Uniqueness of Nash Equilibrium Pair

We now establish the existence and uniqueness of the Nash equilibrium pair defined in Definition 3.1.4. We impose the following regularity conditions.

Assumption 3.2.1. We assume that the following statements hold:

- (i) The algebraic Riccati equation $X = A^\top X A + Q - A^\top X B (B^\top X B + R)^{-1} B^\top X A$ admits a unique symmetric positive definite solution X^* ;
- (ii) It holds for $L_0 = L_1 L_3 + L_2$ that $L_0 < 1$, where

$$L_1 = \left\| \left[(I - A) Q^{-1} (I - A)^\top + B R^{-1} B^\top \right]^{-1} \bar{A} \right\|_2 \cdot \left\| [K^* Q^{-1} (I - A)^\top - R^{-1} B^\top] \right\|_2,$$

$$L_2 = [1 - \rho(A - B K^*)]^{-1} \cdot \|\bar{A}\|_2, \quad L_3 = [1 - \rho(A - B K^*)]^{-1} \cdot \|B\|_2.$$

$$\text{Here } K^* = -(B^\top X^* B + R)^{-1} B^\top X^* A.$$

The first assumption is implied by mild regularity conditions on the matrices A , B , Q , and R , which are (1) the positivity of R ; (2) the non-negativity of $Q = C^\top C$; (3) the observability of (A, C) ; (4) the stability of (A, B) . See De Souza et al. (1986); Lewis et al. (2012) for more details. The second assumption is standard in the literature (Bensoussan et al., 2016; Saldi et al., 2018b), which ensures the stability of the LQ-MFG. In the

following proposition, we show that Problem 3.1.3 admits a unique Nash equilibrium pair.

Proposition 3.2.2 (Existence and Uniqueness of Nash Equilibrium Pair). Under Assumption 3.2.1, the operator $\Lambda(\cdot) = \Lambda_2(\cdot, \Lambda_1(\cdot))$ is L_0 -Lipschitz, where L_0 is given in Assumption 3.2.1. Moreover, there exists a unique Nash equilibrium pair (μ^*, π^*) of Problem 3.1.3.

Proof. See §B.5.1 for a detailed proof. □

3.2.2. Mean-Field Actor-Critic for LQ-MFG

The NCE principle motivates a fixed-point approach to solve Problem 3.1.3, which generates a sequence of policies $\{\pi_s\}_{s \geq 0}$ and mean-field states $\{\mu_s\}_{s \geq 0}$ satisfying the following two properties: (i) Given the mean-field state μ_s , the policy π_s is optimal. (ii) The mean-field state becomes μ_{s+1} as $t \rightarrow \infty$, if all the agents follow π_s under the current mean-field state μ_s . Here (i) requires solving Problem 3.1.2 given the mean-field state μ_s , while (ii) requires simulating the agents following the policy π_s given the current mean-field μ_s . Based on such properties, we propose the mean-field actor-critic in Algorithm 2.

Algorithm 2 requires solving Problem 3.1.2 at each iteration to obtain $\pi_s = \Lambda_1(\mu_s)$ and $\mu_{s+1} = \Lambda_2(\mu_s, \pi_s)$. To this end, we introduce the natural actor-critic in §3.2.3 that solves Problem 3.1.2.

3.2.3. Natural Actor-Critic for D-LQR

Now we focus on solving Problem 3.1.2 for a fixed mean-field state μ , we thus drop the subscript μ hereafter. With slight abuse of notations, we write $\pi_{K,b}(x) = -Kx + b + \sigma \cdot \eta$

Algorithm 2 Mean-Field Actor-Critic for solving LQ-MFG.

1: **Input:**

- Initial mean-field state μ_0 and Initial policy π_0 with parameters K_0 and b_0 .
- Numbers of iterations S , $\{N_s\}_{s \in [S]}$, $\{H_s\}_{s \in [S]}$, $\{\tilde{T}_{s,n}, T_{s,n}\}_{s \in [S], n \in [N_s]}$, $\{\tilde{T}_{s,h}^b, T_{s,h}^b\}_{s \in [S], h \in [H_s]}$.
- Stepsizes $\{\gamma_s\}_{s \in [S]}$, $\{\gamma_s^b\}_{s \in [S]}$, $\{\gamma_{s,n,t}\}_{s \in [S], n \in [N_s], t \in [T_{s,n}]}$, $\{\gamma_{s,h,t}^b\}_{s \in [S], h \in [H_s], t \in [T_{s,h}^b]}$.

2: **for** $s = 0, 1, 2, \dots, S - 1$ **do**

- 3: **Policy Update:** Solve for the optimal policy π_{s+1} with parameters K_{s+1} and b_{s+1} of Problem 3.1.2 via Algorithm 3 with μ_s , π_s , N_s , H_s , $\{\tilde{T}_{s,n}, T_{s,n}\}_{n \in [N_s]}$, $\{\tilde{T}_{s,h}^b, T_{s,h}^b\}_{h \in [H_s]}$, γ_s , γ_s^b , $\{\gamma_{s,n,t}\}_{n \in [N_s], t \in [T_{s,n}]}$, and $\{\gamma_{s,h,t}^b\}_{h \in [H_s], t \in [T_{s,h}^b]}$, which gives the estimated mean-field state $\hat{\mu}_{K_{s+1}, b_{s+1}}$.

- 4: **Mean-Field State Update:** Update the mean-field state via $\mu_{s+1} \leftarrow \hat{\mu}_{K_{s+1}, b_{s+1}}$.

5: **end for**

- 6: **Output:** Pair (π_S, μ_S) .
-

to emphasize the dependence on K and b , and $J(K, b) = J(\pi_{K,b})$ consequently. Now, we propose the natural actor-critic to solve Problem 3.1.2.

For any policy $\pi_{K,b} \in \Pi$, by the state transition in Problem 3.1.2, we have

$$(3.2.1) \quad x_{t+1} = (A - BK)x_t + (Bb + \bar{A}\mu + d) + \epsilon_t, \quad \epsilon_t \sim \mathcal{N}(0, \Psi_\epsilon),$$

where $\Psi_\epsilon = \sigma BB^\top + \Psi_\omega$. It is known that if $\rho(A - BK) < 1$, then the Markov chain $\{x_t\}_{t \geq 0}$ induced by (3.2.1) has a unique stationary distribution $\mathcal{N}(\mu_{K,b}, \Phi_K)$ (Anderson and Moore, 2007), where the mean-field state $\mu_{K,b}$ and the covariance Φ_K satisfy that

$$(3.2.2) \quad \mu_{K,b} = (I - A + BK)^{-1}(Bb + \bar{A}\mu + d),$$

$$(3.2.3) \quad \Phi_K = (A - BK)\Phi_K(A - BK)^\top + \Psi_\epsilon.$$

Meanwhile, the Bellman equation for Problem 3.1.2 takes the following form

$$(3.2.4) \quad P_K = (Q + K^\top RK) + (A - BK)^\top P_K (A - BK).$$

Then by calculation (see Proposition B.2.2 in §B.2.1 of the appendix for details), it holds that the expected total cost $J(K, b)$ is decomposed as

$$(3.2.5) \quad J(K, b) = J_1(K) + J_2(K, b) + \sigma^2 \cdot \text{tr}(R) + \mu^\top \bar{Q} \mu,$$

where $J_1(K)$ and $J_2(K, b)$ are defined as

$$(3.2.6) \quad \begin{aligned} J_1(K) &= \text{tr}[(Q + K^\top RK)\Phi_K] = \text{tr}(P_K \Psi_\epsilon), \\ J_2(K, b) &= \begin{pmatrix} \mu_{K,b} \\ b \end{pmatrix}^\top \begin{pmatrix} Q + K^\top RK & -K^\top R \\ -RK & R \end{pmatrix} \begin{pmatrix} \mu_{K,b} \\ b \end{pmatrix}. \end{aligned}$$

Here $J_1(K)$ is the expected total cost in the most studied LQR problems (Yang et al., 2019b; Fazel et al., 2018), where the state transition does not have drift terms. Meanwhile, $J_2(K, b)$ corresponds to the expected cost induced by the drift terms. The following two propositions characterize the properties of $J_2(K, b)$.

First, we show that $J_2(K, b)$ is strongly convex in b .

Proposition 3.2.3. Given any K , the function $J_2(K, b)$ is ν_K -strongly convex in b . Here $\nu_K = \sigma_{\min}(Y_{1,K}^\top Y_{1,K} + Y_{2,K}^\top Y_{2,K})$, where $Y_{1,K} = R^{1/2}K(I - A + BK)^{-1}B - R^{1/2}$ and $Y_{2,K} = Q^{1/2}(I - A + BK)^{-1}B$. Also, $J_2(K, b)$ has ι_K -Lipschitz continuous gradient in b , where ι_K is upper bounded as $\iota_K \leq [1 - \rho(A - BK)]^{-2} \cdot (\|B\|_2^2 \cdot \|K\|_2^2 \cdot \|R\|_2 + \|B\|_2^2 \cdot \|Q\|_2)$.

Proof. See §B.5.4 for a detailed proof. □

Second, we show that $\min_b J_2(K, b)$ is independent of K .

Proposition 3.2.4. We define $b^K = \operatorname{argmin}_b J_2(K, b)$, where $J_2(K, b)$ is defined in (3.2.6).

It holds that

$$b^K = [KQ^{-1}(I - A)^\top - R^{-1}B^\top] \cdot [(I - A)Q^{-1}(I - A)^\top + BR^{-1}B^\top]^{-1} \cdot (\bar{A}\mu + d).$$

Moreover, $J_2(K, b^K)$ takes the form of

$$J_2(K, b^K) = (\bar{A}\mu + d)^\top [(I - A)Q^{-1}(I - A)^\top + BR^{-1}B^\top]^{-1} \cdot (\bar{A}\mu + d),$$

which is independent of K .

Proof. See §B.5.2 for a detailed proof. □

Since $\min_b J_2(K, b)$ is independent of K by Proposition 3.2.4, it holds that the optimal K^* is the same as $\operatorname{argmin}_K J_1(K)$. This motivates us to minimize $J(K, b)$ by first updating K following the gradient direction $\nabla_K J_1(K)$ to the optimal K^* , then updating b following the gradient direction $\nabla_b J_2(K^*, b)$. We now design our algorithm based on this idea.

We define Υ_K , $p_{K,b}$, and $q_{K,b}$ as

$$(3.2.7) \quad \begin{aligned} \Upsilon_K &= \begin{pmatrix} Q + A^\top P_K A & A^\top P_K B \\ B^\top P_K A & R + B^\top P_K B \end{pmatrix} = \begin{pmatrix} \Upsilon_K^{11} & \Upsilon_K^{12} \\ \Upsilon_K^{21} & \Upsilon_K^{22} \end{pmatrix}, \\ p_{K,b} &= A^\top [P_K \cdot (\bar{A}\mu + d) + f_{K,b}], \quad q_{K,b} = B^\top [P_K \cdot (\bar{A}\mu + d) + f_{K,b}], \end{aligned}$$

where $f_{K,b} = (I - A + BK)^{-\top} [(A - BK)^\top P_K (Bb + \bar{A}\mu + d) - K^\top Rb]$. By calculation (see Proposition B.2.3 in §B.2.1 of the appendix for details), the gradients of $J_1(K)$ and

$J_2(K, b)$ take the forms of

$$\nabla_K J_1(K) = 2(\Upsilon_K^{22}K - \Upsilon_K^{21}) \cdot \Phi_K, \quad \nabla_b J_2(K, b) = \Upsilon_K^{22}(-K\mu_{K,b} + b) + \Upsilon_K^{21}\mu_{K,b} + q_{K,b}.$$

Our algorithm follows the natural actor-critic method (Bhatnagar et al., 2009) and actor-critic method (Konda and Tsitsiklis, 2000). Specifically, (i) To obtain the optimal K^* , in the critic update step, we estimate the matrix Υ_K by $\widehat{\Upsilon}_K$ via a policy evaluation algorithm, e.g., Algorithm 7 or Algorithm 8 (see §B.2.2 and §B.2.3 of the appendix for details); in the actor update step, we update K via $K \leftarrow K - \gamma \cdot (\widehat{\Upsilon}_K^{22}K - \widehat{\Upsilon}_K^{21})$, where the term $\widehat{\Upsilon}_K^{22}K - \widehat{\Upsilon}_K^{21}$ is the estimated natural gradient. (ii) To obtain the optimal b^* given K^* , in the critic update step, we estimate Υ_{K^*} , $q_{K^*,b}$, and $\mu_{K^*,b}$ by $\widehat{\Upsilon}_{K^*}$, $\widehat{q}_{K^*,b}$, and $\widehat{\mu}_{K^*,b}$ via a policy evaluation algorithm; In the actor update step, we update b via $b \leftarrow b - \gamma \cdot \widehat{\nabla}_b J_2(K^*, b)$, where $\widehat{\nabla}_b J_2(K^*, b) = \widehat{\Upsilon}_{K^*}^{22}(-K^*\widehat{\mu}_{K^*,b} + b) + \widehat{\Upsilon}_{K^*}^{21}\widehat{\mu}_{K^*,b} + \widehat{q}_{K^*,b}$ is the estimated gradient. Combining the above procedure, we obtain the natural actor-critic for Problem 3.1.2, which is stated in Algorithm 3.

One may want to apply gradient method to $J(K, b)$ directly in the joint space of K and b . However, the gradient dominance property of $J_1(K)$ in the most studied LQR problem (Yang et al., 2019b) no longer holds for $J(K, b)$. Therefore, the convergence of the gradient method to $J(K, b)$ is not guaranteed in our problem.

3.3. Global Convergence Results

The following theorem establishes the rate of convergence of Algorithm 2 to the Nash equilibrium pair (μ^*, π^*) of Problem 3.1.3.

Algorithm 3 Natural Actor-Critic Algorithm for D-LQR.

1: **Input:**

- Mean-field state μ and initial policy π_{K_0, b_0} .
- Numbers of iterations $N, H, \{\tilde{T}_n, T_n\}_{n \in [N]}, \{\tilde{T}_h^b, T_h^b\}_{h \in [H]}$.
- Stepsizes $\gamma, \gamma^b, \{\gamma_{n,t}\}_{n \in [N], t \in [T_n]}, \{\gamma_{h,t}^b\}_{h \in [H], t \in [T_h^b]}$.

2: **for** $n = 0, 1, 2, \dots, N - 1$ **do**3: **Critic Update:** Compute \hat{Y}_{K_n} via Algorithm 7 with $\pi_{K_n, b_0}, \mu, \tilde{T}_n, T_n, \{\gamma_{n,t}\}_{t \in [T_n]}, K_0,$ and b_0 as inputs.4: **Actor Update:** Update the parameter via

$$K_{n+1} \leftarrow K_n - \gamma \cdot (\hat{Y}_{K_n}^{22} K_n - \hat{Y}_{K_n}^{21}).$$

5: **end for**6: **for** $h = 0, 1, 2, \dots, H - 1$ **do**7: **Critic Update:** Compute $\hat{\mu}_{K_N, b_h}, \hat{Y}_{K_N}, \hat{q}_{K_N, b_h}$ via Algorithm 7 with $\pi_{K_N, b_h}, \mu, \tilde{T}_h^b, T_h^b, \{\gamma_{h,t}^b\}_{t \in [T_h^b]}, K_0,$ and b_0 .8: **Actor Update:** Update the parameter via

$$b_{h+1} \leftarrow b_h - \gamma^b \cdot [\hat{Y}_{K_N}^{22} (-K_N \hat{\mu}_{K_N, b_h} + b_h) + \hat{Y}_{K_N}^{21} \hat{\mu}_{K_N, b_h} + \hat{q}_{K_N, b_h}].$$

9: **end for**10: **Output:** Policy $\pi_{K,b} = \pi_{K_N, b_H}$, estimated mean-field state $\hat{\mu}_{K,b} = \hat{\mu}_{K_N, b_H}$.

Theorem 3.3.1 (Convergence of Algorithm 2). For a sufficiently small tolerance $\varepsilon > 0$, we set the number of iterations S in Algorithm 2 such that

$$(3.3.1) \quad S > \frac{\log(\|\mu_0 - \mu^*\|_2 \cdot \varepsilon^{-1})}{\log(1/L_0)}.$$

For any $s \in [S]$, we define

$$\varepsilon_s = \min \left\{ [1 - \rho(A - BK^*)]^4 (\|B\|_2 + \|\bar{A}\|_2)^{-4} (\|\mu_s\|_2^{-2} + \|d\|_2^{-2}) \cdot \sigma_{\min}(\Psi_\varepsilon) \cdot \sigma_{\min}(R) \cdot \varepsilon^2, \right.$$

(3.3.2)

$$\left. \nu_{K^*} \cdot [1 - \rho(A - BK^*)]^4 \cdot \|B\|_2^{-2} \cdot M_b(\mu_s) \cdot \varepsilon^2, \varepsilon \right\} \cdot 2^{-s-10},$$

where ν_{K^*} is defined in Proposition 3.2.3 and

$$(3.3.3) \quad M_b(\mu_s) = 4 \left\| Q^{-1}(I - A)^\top \cdot [(I - A)Q^{-1}(I - A)^\top + BR^{-1}B^\top]^{-1} \cdot (\bar{A}\mu_s + d) \right\|_2 \\ \cdot [\nu_{K^*}^{-1} + \sigma_{\min}^{-1}(\Psi_\epsilon) \cdot \sigma_{\min}^{-1}(R)]^{1/2}.$$

In the s -th policy update step in Line 3 of Algorithm 2, we set the inputs via Theorem B.2.4 such that $J_{\mu_s}(\pi_{s+1}) - J_{\mu_s}(\pi_{\mu_s}^*) < \varepsilon_s$, where the expected total cost $J_{\mu_s}(\cdot)$ is defined in Problem 3.1.2, and $\pi_{\mu_s}^* = \Lambda_1(\mu_s)$ is the optimal policy under the mean-field state μ_s . Then it holds with probability at least $1 - \varepsilon^5$ that

$$\|\mu_S - \mu^*\|_2 \leq \varepsilon, \quad \|K_S - K^*\|_{\mathbb{F}} \leq \varepsilon, \quad \|b_S - b^*\|_2 \leq (1 + L_1) \cdot \varepsilon.$$

Here μ^* is the Nash mean-field state, K_S and b_S are parameters of the policy π_S , and K^* and b^* are parameters of the Nash policy π^* .

PROOF SKETCH. The proof of the theorem is based on the convergence of the natural actor-critic algorithm 3 and a contraction argument. First, we prove in Theorem B.2.4 that Algorithm 3 converges linearly to the optimal policy of Problem 3.1.2. By this, in each iteration of Algorithm 2, we control the error between μ_{s+1} and μ_{s+1}^* to be $\varepsilon_s > 0$ with high probability, where μ_{s+1}^* is the mean-field state generated by the optimal policy $\Lambda_1(\mu_s)$; in other words, $\mu_{s+1}^* = \Lambda(\mu_s)$. Combining the fact from Proposition 3.2.2 that $\Lambda(\cdot)$ is a contraction, we deduce that

$$\|\mu_{s+1} - \mu^*\|_2 \leq \|\Lambda(\mu_s) - \Lambda(\mu^*)\|_2 + \tilde{\varepsilon}_s \leq L_0 \cdot \|\mu_s - \mu^*\|_2 + \tilde{\varepsilon}_s$$

with high probability, where $\tilde{\varepsilon}_s > 0$ is some error term and is specified in the detailed proof. Moreover, by telescoping sum, and note that the sum $\sum_{s=1}^S \tilde{\varepsilon}_s$ is upper bounded by the desired error ε , we conclude the theorem. See §B.4.1 for a detailed proof. \square

We highlight that if the inputs of Algorithm 2 satisfy the conditions stated in Theorem B.2.4, it holds that $J_{\mu_s}(\pi_{s+1}) - J_{\mu_s}(\pi_{\mu_s}^*) < \varepsilon_s$ for any $s \in [S]$. See Theorem B.2.4 in §B.2.1 of the appendix for details. By Theorem 3.3.1, Algorithm 2 converges linearly to the unique Nash equilibrium pair (μ^*, π^*) of Problem 3.1.3. To the best of our knowledge, this theorem is the first successful attempt to establish that reinforcement learning with function approximation finds the Nash equilibrium pairs in mean-field games with theoretical guarantee, which lays the theoretical foundations for applying modern reinforcement learning techniques to general mean-field games.

3.4. Conclusion

For the discrete-time linear-quadratic mean-field games, we provide sufficient conditions for the existence and uniqueness of the Nash equilibrium pair. Moreover, we propose the mean-field actor-critic algorithm with linear function approximation that is shown converges to the Nash equilibrium pair with linear rate of convergence. Our algorithm can be modified to use other parametrized function classes, including deep neural networks, for solving mean-field games. For future research, we aim to extend our algorithm to other variations of mean-field games including risk-sensitive mean-field games (Saldi et al., 2018a; Tembine et al., 2014), robust mean-field games (Bauso et al., 2016), and partially observed mean-field games (Saldi et al., 2019).

CHAPTER 4

Offline Reinforcement Learning with Instrumental Variables in Confounded Markov Decision Processes

4.1. Confounded Markov Decision Processes

In this section, we introduce the framework of confounded Markov decision processes with discrete instrumental variables. We aim to leverage the batch data to find an optimal in-class policy that maximizes the expected total rewards.

Confounded MDPs In a confounded MDP, we observe $\{S_t, A_t, R_t\}_{t \geq 0}$ for each trajectory, where S_t is the observed state, A_t is the action taken after observing S_t , and R_t is the immediate reward received after making an action A_t for $t \geq 0$. We denote by \mathcal{S} and \mathcal{A} the state and action spaces, respectively. Furthermore, we assume that at each decision point $t \geq 0$, there exist some unmeasured state variables $U_t \in \mathcal{U}$, which may confound the effect of action A_t on the rewards and future transitions. Due to such unobserved confounders, the (causal) effect of the action on the immediate and future rewards may not be non-parametrically identified and directly applying standard RL algorithms for MDPs will produce sub-optimal policies.

To address this concern, we study the confounded MDP via the instrumental variable (IV) method (Angrist and Imbens, 1995), which has been widely used in the literature of causal inference (e.g., Pearl (2009); Hernán and Robins (2010)) to identify the causal effect of a treatment under unmeasured confounding. Specifically, at each decision point

t , we further assume that we also observe a time-varying IV $Z_t \in \mathcal{Z}$, which is independent of U_t and does not have a direct effect on the immediate reward R_t and all future states, actions, and rewards. With such an IV, we observe $\{S_t, Z_t, A_t, R_t\}_{t \geq 0}$ for each trajectory in the confounded MDP.

In this work, we consider finite action and IV spaces, i.e., $\mathcal{A} = \{a_j\}_{j \in [K]}$ and $\mathcal{Z} = \{z_j\}_{j \in [K]}$, where $K \geq 2$ is an integer. Furthermore, we consider a simplex encoding for both actions and IVs which enjoy a nice interpretation (Zhang and Liu, 2014). Specifically, for any $j \in [K]$, we let

$$(4.1.1) \quad a_j = z_j = \begin{cases} (K-1)^{-1/2} \mathbf{1}_{K-1} & \text{if } j = 1, \\ \frac{(1+\sqrt{K}\mathbf{1}_{K-1})}{(K-1)^{3/2}} + \sqrt{\frac{K}{K-1}} e_{j-1} & \text{if } 2 \leq j \leq K, \end{cases}$$

where $\mathbf{1}_{K-1} \in \mathbb{R}^{K-1}$ is an all-one vector and $e_j \in \mathbb{R}^{K-1}$ is a vector with all elements 0 except 1 for j -th position. By the simplex encoding in (4.1.1), one can see that $\sum_{j \in [K]} a_j = \sum_{j \in [K]} z_j = 0$ and $a_i^\top a_j = z_i^\top z_j = -\mathbb{1}\{i \neq j\}/(K-1) + \mathbb{1}\{i = j\}$ for any $i, j \in [K]$, where $\mathbb{1}\{\cdot\}$ is an indicator function. We remark that any reasonable encoding mechanisms can be incorporated here and our results can be equally applied.

Value Function and Performance Metric. In the confounded MDP, we aim to find an optimal in-class policy $\pi^* \in \Pi$ such that π^* maximizes the expected total rewards, where Π is a class of time-homogeneous policies mapping from the observed state space \mathcal{S} into the probability distribution over the action space \mathcal{A} . In particular, $\pi(a|s)$ refers to the probability of choosing action $a \in \mathcal{A}$ given the state value $s \in \mathcal{S}$. Formally, for any

$\pi \in \Pi$, we define the value function V^π and the expected total reward $J(\pi)$ as follows,

$$(4.1.2) \quad V^\pi(s) = \mathbb{E}_\pi \left[\sum_{t=0}^{\infty} \gamma^t R_t \mid S_0 = s \right], \quad J(\pi) = (1 - \gamma) \cdot \mathbb{E}_{S_0 \sim \nu} [V^\pi(S_0)],$$

where the expectation $\mathbb{E}_\pi[\cdot]$ is taken with respect to the distribution such that the action $A_t \sim \pi(\cdot \mid S_t)$ for any $t \geq 0$, and ν is a known reference distribution over \mathcal{S} . Given the definition of $J(\pi)$ in (4.1.2), our goal is to leverage the batch data to estimate π^* , where

$$\pi^* \in \operatorname{argmax}_{\pi \in \Pi} J(\pi).$$

Suppose the batch data we have collected consist of N independent and identically distributed copies of $\{S_t, Z_t, A_t, R_t\}_{t \geq 0}$ with total decision points T for each trajectory. Then we can summarize our batch data as $\mathcal{D} = \{\{S_t^i, Z_t^i, A_t^i, R_t^i, S_{t+1}^i\}_{t=0}^{T-1}\}_{i \in [N]}$. In the meanwhile, we define the performance metric as

$$\text{SubOpt}(\pi) = J(\pi^*) - J(\pi),$$

which characterizes the suboptimality of the policy π compared with the optimal in-class policy π^* .

Why is Confounded MDP Challenging? In the standard MDP (Sutton and Barto, 2018), all states are assumed fully observed and the trajectory $\{S_t, A_t, R_t\}_{t \geq 0}$ satisfies the Markovian property. By leveraging the celebrated Bellman equation, under some mild conditions, one can non-parametrically identify $J(\pi)$ for any $\pi \in \Pi$, which serves as a foundation for many existing RL algorithms. However, in the confounded MDP, the Markovian assumption on the trajectory no longer holds and Bellman equation cannot

be applied anymore due to the existence of the unmeasured confounders $\{U_t\}_{t \geq 0}$. More seriously, the effect of actions on the rewards and future states cannot be identified even we include all past history information at each decision point $t \geq 0$. Therefore, additional assumptions are needed to identify $J(\pi)$ and in this work, we rely on the IV to deal with such challenges.

Notation. Throughout the paper, we denote by c a positive absolute constant, which may vary from lines to lines. Without further explanation, we denote by $\mathbb{E}_\pi[\cdot]$ the expectation taken with respect to the trajectory generated by the policy π , $\mathbb{E}[\cdot]$ the expectation taken with respect to the trajectory generated by the behavior policy, and $\widehat{\mathbb{E}}[\cdot]$ the empirical average across all N trajectories.

4.2. Assumptions and Identification Results

In this section, we introduce several assumptions to help us identify $J(\pi)$ for any $\pi \in \Pi$ by using an IV. The first assumption is related to the trajectory $\{S_t, U_t, A_t, R_t\}_{t \geq 0}$, where we model it by a time-homogenous MDP.

Assumption 4.2.1. The following statements hold.

- (a) For any $t \geq 1$, we have $(S_{t+1}, U_{t+1}) \perp\!\!\!\perp \{S_j, U_j, A_j\}_{0 \leq j < t} \mid (S_t, U_t, A_t)$ and the transition probability is stationary;
- (b) For any $t \geq 0$, we have $R_t = R(U_t, S_t, A_t, S_{t+1}, U_{t+1})$ for some deterministic function $R: \mathcal{U} \times \mathcal{S} \times \mathcal{A} \times \mathcal{S} \times \mathcal{U} \rightarrow \mathbb{R}$. Also, we assume $|R_t| \leq 1$ almost surely for any $t \geq 0$;

- (c) The offline dataset \mathcal{D} is generated by an unknown initial distribution ζ over \mathcal{S} and a stationary policy b , which is a function mapping from $\mathcal{S} \times \mathcal{U} \times \mathcal{Z}$ into a probability distribution over \mathcal{A} .

Here b is often called behavior policy in RL literature. Assumption 4.2.1 is standard in the literature of RL, which is mild as $\{U_t\}_{t \geq 0}$ is unobserved. The uniformly bounded assumption on the reward R_t is used to simplify the technical analysis and can be relaxed by imposing some high-order moment condition on R_t instead. Due to the unobserved state variables U_t , we make the following IV assumptions.

Assumption 4.2.2. The following statements hold.

- (a) For any $t \geq 0$, we have $(S_{t+1}, U_{t+1}) \perp\!\!\!\perp Z_t \mid (S_t, U_t, A_t)$;
- (b) For any $a \in \mathcal{A}$ and $t \geq 0$, we have $\mathbb{P}(A_t = a \mid S_t, Z_t) \neq \mathbb{P}(A_t = a \mid S_t)$;
- (c) For any $t \geq 0$, we have $Z_t \perp\!\!\!\perp U_t \mid S_t$ and the probability distribution of Z_t given S_t is time-homogeneous.
- (d) For any $t \geq 0$, we have the behavior policy satisfy that

$$\begin{aligned} b(A_t = a \mid S_t, U_t, Z_t = a) &- \frac{1}{K-1} \sum_{z \in \mathcal{Z}, z \neq a} b(A_t = a \mid S_t, U_t, Z_t = z) \\ &= b(A_t = a \mid S_t, Z_t = a) - \frac{1}{K-1} \sum_{z \in \mathcal{Z}, z \neq a} b(A_t = a \mid S_t, Z_t = z) = \Delta^*(S_t, a), \end{aligned}$$

i.e., the compliance $\Delta^*(S_t, a)$ defined above is independent of the unobserved confounder U_t almost surely.

Assumption 4.2.2(a) states that there is not direct effect of the IV Z_t on the future states and rewards except through the action A_t , which is a typical assumption in the literature of causal inference with IVs (Angrist and Imbens, 1995; Angrist et al., 1996).

Note that by Assumption 4.2.1(b), we have implicitly restricted the effect of Z_t on the reward R_t only through A_t in this assumption. Assumption 4.2.2(b) requires that the IV Z_t will influence the action A_t , which is called IV relevance in the causal inference. Assumption 4.2.2(c), corresponding to IV independence, ensures that the effect of Z_t on futures states and rewards is unconfounded by adjusting the current state S_t . The homogeneous assumption on the conditional distribution of Z_t given S_t is imposed here as our target parameter $J(\pi)$ is defined over the infinite horizon. Define a function

$$\Theta^*(s, z) = \mathbb{P}(Z_t = z | S_t = s)$$

for every $(s, z) \in \mathcal{S} \times \mathcal{Z}$, which is independent of the decision point due to such time-homogeneity. In addition, Assumption 4.2.2(d) essentially indicates that there is no interaction between U_t and Z_t in affecting whether the action A_t will follow Z_t or not. This so-called independent compliance assumption has been widely adopted in identifying the average treatment effect of binary treatments in causal inference (Wang and Tchetgen Tchetgen, 2018). Here we generalize it to the setting of multiple treatments and instrumental variables, which may be of independent interest. A graphical illustration of Assumptions 4.2.1 and 4.2.2 is presented in Figure 4.1, which also illustrates how the offline data in the confounded MDP are generated. In the following section, we introduce value function (VF)-based identification and marginalized importance sampling (MIS)-based identification, respectively.




Figure 4.1. A graphical illustration of the confounded MDP.

4.2.1. Value Function-based Identification

In the unconfounded MDP, value function defined in (4.1.2) can be used to identify $J(\pi)$ and itself can be identified via the Bellman equation. However, due to the existence of unobserved confounders, the regular Bellman equation, which relies on the Markovian assumption, does not hold in general and the effect of actions on the reward cannot be identified either. Fortunately, by leveraging the IV, we are able to provide a way to identify $J(\pi)$ via the state-value function V^π , which can be identified by an IV-aided Bellman equation. Before stating our result, we make one additional assumption.

Assumption 4.2.3. We have $(Z_t, U_t) \perp\!\!\!\perp (\{S_j, U_j, A_j\}_{0 \leq j < t}) \mid S_t$ and the probability distribution of (Z_t, U_t) given S_t is time-homogeneous for any $t \geq 0$.

Assumption 4.2.3 ensures that (Z_t, U_t) is “memoryless” and only depends on the current observed state S_t . This essentially ensures that the stochastic process $\{S_t, A_t\}_{t \geq 0}$ satisfies Markov property, which rules out the scenario that U_t affects future rewards and transitions. The memoryless assumption on the unmeasured confounders has been commonly used in the confounded MDP. See Kallus and Zhou (2020); Shi et al. (2022) for more details.

Lemma 4.2.4. Under Assumptions 4.2.1 and 4.2.2, for any $s \in \mathcal{S}$ and $\pi \in \Pi$, we have

$$V^\pi(s) = \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t R_t \left(\prod_{j=0}^t \frac{Z_j^\top A_j \pi(A_j | S_j)}{\Delta^*(S_j, A_j) \Theta^*(S_j, Z_j)} \right) \middle| S_0 = s \right].$$

If additionally Assumption 4.2.3 is satisfied, it holds for any $t \geq 0$ that,

$$V^\pi(s) = \mathbb{E} \left[\frac{Z_t^\top A_t \pi(A_t | S_t)}{\Delta^*(S_t, A_t) \Theta^*(S_t, Z_t)} \cdot (R_t + \gamma V^\pi(S_{t+1})) \middle| S_t = s \right].$$

Then the policy value $J(\pi)$ for $\pi \in \Pi$ can be identified via

$$J(\pi) = (1 - \gamma) \cdot \mathbb{E}_{S_0 \sim \nu} [V^\pi(S_0)].$$

Proof. See §C.2.3 for a detailed proof. □

We remark that the Bellman equation in the unconfounded MDP takes the following form,

$$V_{\text{unconf}}^\pi(s) = \mathbb{E} \left[\frac{\pi(A_t | S_t)}{\mathbb{P}(A_t | S_t)} \cdot (R_t + \gamma V_{\text{unconf}}^\pi(S_{t+1})) \middle| S_t = s \right],$$

where V_{unconf}^π is the corresponding state-value function in the unconfounded MDP. In comparison, to deal with the unobserved confounders, our identification result in Lemma 4.2.4 incorporates the IVs into the action density ratio. It is also interesting to see that if one can observe the trajectory $\{S_t, Z_t, A_t, R_t\}_{t \geq 0}$ to the infinity, then Assumptions 4.2.1 and 4.2.2 are sufficient to identify $V^\pi(s)$ and $J(\pi)$ based on the first statement of Lemma 4.2.4. However, due to the limitation of only observing trajectories up to a finite horizon, we impose Assumption 4.2.3 so that Bellman equation is satisfied and used to break the curse of infinite-horizon. Based on Lemma 4.2.4, we introduce the following VF-based

estimating equation, which will be used later in §4.3.1 to construct an estimator of the value function V^π .

Corollary 4.2.5 (VF-based Estimating Equation). Under Assumptions 4.2.1, 4.2.2, and 4.2.3, it holds for any function $g: \mathcal{S} \rightarrow \mathbb{R}$ that

$$\mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} g(S_t) \frac{Z_t^\top A_t \pi(A_t | S_t)}{\Delta^*(S_t, A_t) \Theta^*(S_t, Z_t)} \cdot (R_t + \gamma V^\pi(S_{t+1})) \right] = \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} g(S_t) V^\pi(S_t) \right].$$

Proof. See §C.2.4 for a detailed proof. \square

4.2.2. Marginalized Importance Sampling-based Identification

In this subsection, we propose another way to identify $J(\pi)$ via the marginal importance sampling. We first introduce the following notations. For any $t \geq 0$, we denote by $p_t^\pi(\cdot)$ the marginal distribution of S_t under the known initial observed state distribution ν following the policy π . In the meanwhile, with a slight abuse of notations, we denote by $p_t^b(\cdot)$ the marginal distribution of S_t under the unknown offline data generation distribution ζ following the behavior policy b . In addition, we denote by for every $s \in \mathcal{S}$,

$$(4.2.1) \quad d^\pi(s) = (1 - \gamma) \sum_{t=0}^{\infty} \gamma^t p_t^\pi(s), \quad d^b(s) = \frac{1}{T} \sum_{t=0}^{T-1} p_t^b(s), \quad w^\pi(s) = \frac{d^\pi(s)}{d^b(s)}$$

the discounted state visitation measure under the policy π , the average state visitation measure under the behavior policy b , and their density ratio, respectively. In order to identify $J(\pi)$ via the marginal importance sampling, i.e., w^π , we make the following assumption.

Assumption 4.2.6. For any $t \geq 0$, $U_t | S_t$ is time-homogeneous.

Assumption 4.2.6 is weaker than assuming the unmeasured confounder U_t is memoryless, and *does not* imply that $\{S_t, A_t\}_{t \geq 0}$ is a Markov chain. Such an assumption will be satisfied for example when the trajectory comes from a stationary sequence or unmeasured confounders are independent of observed states and actions. Motivated by the idea of marginalized importance sampling for off-policy evaluation in the standard MDP (Liu et al., 2018), we establish the following novel identification result for the expected total reward $J(\pi)$ in the confounded MDP.

Lemma 4.2.7. Under Assumptions 4.2.1, 4.2.2, and 4.2.6, for any $\pi \in \Pi$, we have

$$J(\pi) = \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} \frac{Z_t^\top A_t \pi(A_t | S_t)}{\Delta^*(S_t, A_t) \Theta^*(S_t, Z_t)} \cdot w^\pi(S_t) R_t \right],$$

where w^π is defined in (4.2.1).

Proof. See §C.2.1 for a detailed proof. □

We remark that the expected total reward in the unconfounded MDP takes the following form,

$$J_{\text{unconf}}(\pi) = \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} \frac{\pi(A_t | S_t)}{\mathbb{P}(A_t | S_t)} \cdot w^\pi(S_t) R_t \right],$$

where $J_{\text{unconf}}(\pi)$ is the corresponding expected total reward in the unconfounded MDP, and the expectation $\mathbb{E}[\cdot]$ is taken with respect to the trajectory generated by the behavior policy. In comparison, to deal with the unobserved confounders, our identification result in Lemma 4.2.7 incorporates the IVs into the action density ratio. Based on Lemma 4.2.7, we introduce the following MIS-based estimating equation, which will be used later in §4.3.2 to construct an estimator of the density ratio w^π .

Lemma 4.2.8 (MIS-based Estimating Equation). Under Assumptions 4.2.1, 4.2.2, and 4.2.6, for any $\pi \in \Pi$, it holds for any function $f: \mathcal{S} \rightarrow \mathbb{R}$ that

$$(1 - \gamma)\mathbb{E}_{S_0 \sim \nu} [f(S_0)] = \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} \frac{Z_t^\top A_t \pi(A_t | S_t)}{\Delta^*(S_t, A_t) \Theta^*(S_t, Z_t)} \cdot w^\pi(S_t) (f(S_t) - \gamma f(S_{t+1})) \right].$$

Proof. See §C.2.2 for a detailed proof. \square

Lemma 4.2.8 is somewhat surprising in that with the help of IVs and Assumption 4.2.6, the estimating equation for the density ratio w^π holds even when the Markov condition fails on the observed trajectory, thus allowing the existence of unmeasured confounder U_t that can affect future rewards and transitions. This is different from the existing approaches such as Liu et al. (2018); Zhang et al. (2020) for estimating ratio functions in the standard MDP setting, which relies crucially on the Markovian assumption. Compared with the value function-based identification, marginalized importance sampling-based identification on the $J(\pi)$ requires fewer conditions and can be applied into more general confounded MDP problems.

4.3. Instrumental-Variable-Assisted RL with Pessimism

In this section, we introduce three pessimistic RL methods to estimate π^* in our confounded MDP. Generally, pessimistic RL first employs the offline data to construct a conservative estimate of the values for any policy, then select the policy with the highest conservative estimate of its value. Though the recently proposed pessimistic RL shows promising performance in practice (Kumar et al., 2020; Yu et al., 2020; Kidambi et al., 2020; Deng et al., 2021), its theoretical understanding is far from complete and is only

limited to fully observable MDP (Levine et al., 2020). In this section, we adapt the idea of pessimism in our case and present related theoretical results in Section 4.4.

Since both identification results in §4.2.1 and §4.2.2 require estimating the quantities $\Delta^*(s, a)$ and $\Theta^*(s, z)$, we first introduce the estimating procedure for such quantities. We assume that there exists an oracle that gives estimators of $\Delta^*(s, a)$ and $\Theta^*(s, z)$ via two loss functions $\widehat{L}_0(\Delta)$ and $\widehat{L}_1(\Theta)$ as follows,

$$\widehat{\Delta} \in \operatorname{argmin}_{\Delta \in \mathcal{F}_0} \widehat{L}_0(\Delta), \quad \widehat{\Theta} \in \operatorname{argmin}_{\Theta \in \mathcal{F}_1} \widehat{L}_1(\Theta),$$

where \mathcal{F}_0 and \mathcal{F}_1 are two function classes. We remark that we can use the negative likelihood functions for \widehat{L}_0 and \widehat{L}_1 (see §4.4 for details). In the meanwhile, we construct two confidence sets for Δ and Θ , respectively, as follows,

(4.3.1)

$$\operatorname{conf}_{\alpha_0}^0 = \left\{ \Delta \in \mathcal{F}_0 : \widehat{L}_0(\Delta) - \widehat{L}_0(\widehat{\Delta}) \leq \alpha_0 \right\}, \quad \operatorname{conf}_{\alpha_1}^1 = \left\{ \Theta \in \mathcal{F}_1 : \widehat{L}_1(\Theta) - \widehat{L}_1(\widehat{\Theta}) \leq \alpha_1 \right\},$$

where (α_0, α_1) are some constants that will be specified later. These two confidence sets are used to construct conservative estimators for $J(\pi)$ via either VF-, MIS-, or doubly robust-based estimation.

4.3.1. VF-based Pessimistic Method

We introduce VF-based pessimistic RL in this section. We first define the following quantity,

$$\widehat{\Phi}_{\text{vf}}^{\pi}(v, g; \Delta, \Theta) = \widehat{\mathbb{E}} \left[\frac{1}{T} \sum_{t=0}^{T-1} g(S_t) \left(\frac{Z_t^\top A_t \pi(A_t | S_t)}{\Delta(S_t, A_t) \Theta(S_t, Z_t)} (R_t + \gamma v(S_{t+1})) - v(S_t) \right) \right],$$

where $\widehat{\mathbb{E}}[\cdot]$ is the empirical measure defined by the offline data \mathcal{D} . In the meanwhile, we define its population counterpart as

$$\Phi_{\text{vf}}^{\pi}(v, g; \Delta, \Theta) = \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} g(S_t) \left(\frac{Z_t^{\top} A_t \pi(A_t | S_t)}{\Delta(S_t, A_t) \Theta(S_t, Z_t)} (R_t + \gamma v(S_{t+1})) - v(S_t) \right) \right]$$

for any (v, g, Δ, Θ) , where the expectation $\mathbb{E}[\cdot]$ is taken with respect to the trajectory generated by the behavior policy. Then by the VF-based estimating equation specified in Corollary 4.2.5, it is easy to see that $\Phi_{\text{vf}}^{\pi}(V^{\pi}, g; \Delta^*, \Theta^*) = 0$ for any function $g: \mathcal{S} \rightarrow \mathbb{R}$.

With the aforementioned notions, for any (Δ, Θ) , we construct an estimator of V^{π} via solving the following minimax optimization problem,

$$(4.3.2) \quad \widehat{v}_{\Delta, \Theta}^{\pi} \in \underset{v \in \mathcal{V}}{\operatorname{argmin}} \max_{g \in \mathcal{W}} \widehat{\Phi}_{\text{vf}}^{\pi}(v, g; \Delta, \Theta),$$

where \mathcal{V} and \mathcal{W} are two sets to be specified later. To obtain an estimation $\widehat{\pi}_{\text{vf}}$ for an optimal in-class policy π^* that maximizes the expected total reward $J(\pi)$ defined in (4.1.2), we formulate the following optimization problem,

$$(4.3.3)$$

$$\widehat{\pi}_{\text{vf}} = \underset{\pi \in \Pi}{\operatorname{argmax}} \min_{(\Delta, \Theta) \in \operatorname{conf}_{\alpha_0}^0 \times \operatorname{conf}_{\alpha_1}^1} \min_{v \in \operatorname{conf}_{\alpha_{\text{vf}}}^{\text{vf}}(\Delta, \Theta, \pi)} (1 - \gamma) \mathbb{E}_{S \sim \nu}[v(S)],$$

$$(4.3.4)$$

$$\text{where } \operatorname{conf}_{\alpha_{\text{vf}}}^{\text{vf}}(\Delta, \Theta, \pi) = \left\{ v \in \mathcal{V} : \max_{g \in \mathcal{W}} \widehat{\Phi}_{\text{vf}}^{\pi}(v, g; \Delta, \Theta) - \max_{g \in \mathcal{W}} \widehat{\Phi}_{\text{vf}}^{\pi}(\widehat{v}_{\Delta, \Theta}^{\pi}, g; \Delta, \Theta) \leq \alpha_{\text{vf}} \right\},$$

where $(\alpha_0, \alpha_1, \alpha_{\text{vf}})$ are constants to be specified, and $\operatorname{conf}_{\alpha_0}^0$ and $\operatorname{conf}_{\alpha_1}^1$ are confidence sets defined in (4.3.1). Intuitively, the policy $\widehat{\pi}_{\text{vf}}$ defined in (4.3.3) aims to maximize the most pessimistic estimator of the expected total reward. As we will see in Theorem 4.4.9,

such a pessimistic method provably converges to an optimal policy with data coverage assumption only for the optimal policy and some other mild conditions.

4.3.2. MIS-based Pessimistic Method

We introduce MIS-based pessimistic RL in this section. We first define the following quantity,

$$\begin{aligned} \widehat{\Phi}_{\text{mis}}^{\pi}(w, f; \Delta, \Theta) &= \mathbb{E}_{S_0 \sim \nu} [(1 - \gamma)f(S_0)] \\ &\quad - \widehat{\mathbb{E}} \left[\frac{1}{T} \sum_{t=0}^{T-1} \frac{Z_t^{\top} A_t \pi(A_t | S_t)}{\Delta(S_t, A_t) \Theta(S_t, Z_t)} w(S_t) (f(S_t) - \gamma f(S_{t+1})) \right]. \end{aligned}$$

We define its population counterpart as

$$\begin{aligned} \Phi_{\text{mis}}^{\pi}(w, f; \Delta, \Theta) &= \mathbb{E}_{S_0 \sim \nu} [(1 - \gamma)f(S_0)] \\ &\quad - \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} \frac{Z_t^{\top} A_t \pi(A_t | S_t)}{\Delta(S_t, A_t) \Theta(S_t, Z_t)} w(S_t) (f(S_t) - \gamma f(S_{t+1})) \right] \end{aligned}$$

for any (w, f, Δ, Θ) . Then by the MIS-based estimating equation specified in Lemma 4.2.8, it can be seen that $\Phi_{\text{mis}}^{\pi}(w^{\pi}, f; \Delta^*, \Theta^*) = 0$ for any function $f: \mathcal{S} \rightarrow \mathbb{R}$. With the aforementioned notions, for any (Δ, Θ) , we construct an estimator of w^{π} via solving the following minimax optimization problem,

$$(4.3.5) \quad \widehat{w}_{\Delta, \Theta}^{\pi} \in \operatorname{argmin}_{w \in \mathcal{W}} \max_{f \in \mathcal{V}} \widehat{\Phi}_{\text{mis}}^{\pi}(w, f; \Delta, \Theta).$$

With a slight abuse of notations, here \mathcal{W} and \mathcal{V} are again two sets to be specified later.

We aim to obtain an optimal policy that maximizes the expected total reward $J(\pi)$ by

utilizing the estimators constructed in (4.3.5). For this, we further define the following estimator of $J(\pi)$ via Lemma 4.2.7,

$$\widehat{L}_{\text{mis}}(w, \pi; \Delta, \Theta) = \widehat{\mathbb{E}} \left[\frac{1}{T} \sum_{t=0}^{T-1} \frac{Z_t^\top A_t \pi(A_t | S_t)}{\Delta(S_t, A_t) \Theta(S_t, Z_t)} w(S_t) R_t \right].$$

Then we aim to solve the following optimization problem,

$$(4.3.6) \quad \widehat{\pi}_{\text{mis}} \in \operatorname{argmax}_{\pi \in \Pi} \min_{(\Delta, \Theta) \in \operatorname{conf}_{\alpha_0}^0 \times \operatorname{conf}_{\alpha_1}^1} \min_{w \in \operatorname{conf}_{\alpha_{\text{mis}}}^{\text{mis}}(\Delta, \Theta, \pi)} \widehat{L}_{\text{mis}}(w, \pi; \Delta, \Theta),$$

$$(4.3.7) \quad \text{where } \operatorname{conf}_{\alpha_{\text{mis}}}^{\text{mis}}(\Delta, \Theta, \pi) = \left\{ w \in \mathcal{W} : \max_{f \in \mathcal{V}} \widehat{\Phi}_{\text{mis}}^\pi(w, f; \Delta, \Theta) \right. \\ \left. - \max_{f \in \mathcal{V}} \widehat{\Phi}_{\text{mis}}^\pi(\widehat{w}_{\Delta, \Theta}^\pi, f; \Delta, \Theta) < \alpha_{\text{mis}} \right\},$$

where $(\alpha_0, \alpha_1, \alpha_{\text{mis}})$ are constants to be specified, and $\operatorname{conf}_{\alpha_0}^0$ and $\operatorname{conf}_{\alpha_1}^1$ are confidence sets defined in (4.3.1). Similarly as in (4.3.3), the policy $\widehat{\pi}_{\text{mis}}$ defined in (4.3.6) aims to maximize the most pessimistic estimator of the expected total reward. As we will see in Theorem 4.4.13, such a pessimistic method provably converges to an optimal policy with realizability assumption only for the optimal policy.

4.3.3. Doubly Robust based Pessimistic Method

As a combination of VF-based and MIS-based policy optimization methods, we introduce a doubly robust (DR)-based pessimistic RL algorithm in this section. We define the

following DR estimator with its population counterpart,

$$\begin{aligned}\widehat{L}_{\text{dr}}(w, v, \pi; \Delta, \Theta) &= \widehat{\mathbb{E}} \left[\frac{1}{T} \sum_{t=0}^{T-1} \frac{Z_t^\top A_t \pi(A_t | S_t)}{\Delta(S_t, A_t) \Theta(S_t, Z_t)} w(S_t) (R_t + \gamma v(S_{t+1}) - v(S_t)) \right] \\ &\quad + (1 - \gamma) \mathbb{E}_{S_0 \sim \nu} [v(S_0)], \\ L_{\text{dr}}(w, v, \pi; \Delta, \Theta) &= \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} \frac{Z_t^\top A_t \pi(A_t | S_t)}{\Delta(S_t, A_t) \Theta(S_t, Z_t)} w(S_t) (R_t + \gamma v(S_{t+1}) - v(S_t)) \right] \\ &\quad + (1 - \gamma) \mathbb{E}_{S_0 \sim \nu} [v(S_0)].\end{aligned}$$

Note that $L_{\text{dr}}(w^\pi, v, \pi; \Delta^*, \Theta^*) = L_{\text{dr}}(w, V^\pi, \pi; \Delta^*, \Theta^*) = J(\pi)$ for any $(\pi, w, v) \in \Pi \times \mathcal{W} \times \mathcal{V}$. Thus, the quantity \widehat{L}_{dr} serves as a valid DR estimator of $J(\pi)$. In the follows, based on such a DR estimator of the expected total reward, we formulate the following optimization problem for estimating the optimal in-class policy,

$$\begin{aligned}\widehat{\pi}_{\text{dr}} \in \operatorname{argmax}_{\pi \in \Pi} \quad & \min_{(\Delta, \Theta) \in \operatorname{conf}_{\alpha_0}^0 \times \operatorname{conf}_{\alpha_1}^1} \min_{(w, v) \in \operatorname{conf}_{\alpha_{\text{mis}}, \alpha_{\text{vf}}}(\Delta, \Theta, \pi)} \widehat{L}_{\text{dr}}(w, v, \pi; \Delta, \Theta), \\ (4.3.8) \quad & \text{where } \operatorname{conf}_{\alpha_{\text{mis}}, \alpha_{\text{vf}}}(\Delta, \Theta, \pi) = \operatorname{conf}_{\alpha_{\text{mis}}}^{\text{mis}}(\Delta, \Theta, \pi) \times \operatorname{conf}_{\alpha_{\text{vf}}}^{\text{vf}}(\Delta, \Theta, \pi),\end{aligned}$$

where $(\alpha_0, \alpha_1, \alpha_{\text{vf}}, \alpha_{\text{mis}})$ are constants to be specified, and $\operatorname{conf}_{\alpha_{\text{vf}}}^{\text{vf}}(\Delta, \Theta, \pi)$ and $\operatorname{conf}_{\alpha_{\text{mis}}}^{\text{mis}}(\Delta, \Theta, \pi)$ are defined in (4.3.4) and (4.3.7), respectively. As we will see in Theorem 4.4.14, such a DR-based pessimistic method provably converges to an optimal policy with realizability assumption only for the optimal policy.

4.4. Theoretical Results

In this section, we investigate theoretical properties of the aforementioned three methods. We aim to derive the finite-sample upper bounds for the sub-optimality of our estimated policies, i.e., $\text{SubOpt}(\hat{\pi})$, where $\hat{\pi}$ is either $\hat{\pi}_{\text{vf}}$, $\hat{\pi}_{\text{mis}}$, or $\hat{\pi}_{\text{dr}}$. To begin with, we first introduce the following definition of covering number, and then impose some assumptions.

Definition 4.4.1 (Covering Number). Let $(\mathcal{C}, \|\cdot\|_\infty)$ be a normed space, and $\mathcal{H} \subseteq \mathcal{C}$. The set $\{x_1, x_2, \dots, x_n\}$ is a ε -covering over \mathcal{H} if $\mathcal{H} \subseteq \cup_{i=1}^n B(x_i, \varepsilon)$, where $B(x_i, \varepsilon)$ is the L_∞ -ball centered at x_i with radius ε . Then the covering number of \mathcal{H} is defined as $N(\varepsilon, \mathcal{H}, \|\cdot\|_\infty) = \min\{n: \exists \varepsilon\text{-covering over } \mathcal{H} \text{ of size } n\}$.

Assumption 4.4.2. The following statements hold.

- (a) For any set $\mathcal{H} \in \{\mathcal{F}_0, \mathcal{F}_1, \mathcal{V}, \mathcal{W}, \Pi\}$, there exists a constant $\mathfrak{C}_{\mathcal{H}}$ such that

$$N(\varepsilon, \mathcal{H}, \|\cdot\|_\infty) \leq c \cdot (1/\varepsilon)^{\mathfrak{C}_{\mathcal{H}}},$$

where $c > 0$ is a constant. Further, we denote by $\mathfrak{C}_{\mathcal{H}_1, \mathcal{H}_2, \dots, \mathcal{H}_k} = \sum_{j \in [k]} \mathfrak{C}_{\mathcal{H}_j}$ for any $\{\mathcal{H}_1, \mathcal{H}_2, \dots, \mathcal{H}_k\}$.

- (b) There exist positive constants C_{Δ^*} and C_{Θ^*} such that $|\Delta^*(s, a)| \geq C_{\Delta^*}^{-1}$ and $\Theta^*(s, z) \geq C_{\Theta^*}^{-1}$ for any $(s, z, a) \in \mathcal{S} \times \mathcal{Z} \times \mathcal{A}$, where $\Theta^*(s, z)$ and $\Delta^*(s, a)$ are defined in Assumption 4.2.2.
- (c) We have $|\Delta(s, a)| \geq C_{\Delta^*}^{-1}$ and $\Theta(s, z) \geq C_{\Theta^*}^{-1}$ for any $(\Delta, \Theta, s, a, z) \in \mathcal{F}_0 \times \mathcal{F}_1 \times \mathcal{S} \times \mathcal{A} \times \mathcal{Z}$.
- (d) We have $\sup_{s \in \mathcal{S}} |V^{\pi_1}(s) - V^{\pi_2}(s)| \leq L_\Pi \cdot \sup_{(s, a) \in \mathcal{S} \times \mathcal{A}} |\pi_1(a|s) - \pi_2(a|s)|$ for any $\pi_1, \pi_2 \in \Pi$, where L_Π is a positive constant.

- (e) We have $\|v\|_\infty \leq 1/(1-\gamma)$ and $\|w\|_\infty \leq C_*$ for any $(v, w) \in \mathcal{V} \times \mathcal{W}$, where $C_* > 0$ is a constant.

Assumption 4.4.2(a) states that the function spaces have finite log covering numbers. Assumption 4.4.2(b) states that the conditional probability Θ^* and the compliance Δ^* are uniformly lower bounded. Such an assumption ensures the identifiability of the expected total reward. With Assumption 4.4.2(b), we only need to consider a lower bounded function class to recover Θ^* and Δ^* , which is imposed in Assumption 4.4.2(c). In the meanwhile, the Lipschitz condition imposed in Assumption 4.4.2(d) aims to control the complexity of the value function class induced by Π , i.e., the class $\{V^\pi(\cdot): \pi \in \Pi\}$. Such an assumption is commonly imposed in related literature (Zhou et al., 2017; Liao et al., 2020). Finally, Assumption 4.4.2(e) states that the sets \mathcal{V} and \mathcal{W} are upper bounded.

Assumption 4.4.3. The sequence $\{S_t, Z_t, U_t, A_t\}_{t \geq 0}$ admits a unique stationary distribution G_{stat} over $\mathcal{S} \times \mathcal{Z} \times \mathcal{U} \times \mathcal{A}$ and is geometrically ergodic, i.e., there exists a function $\varphi: \mathcal{S} \times \mathcal{Z} \times \mathcal{U} \times \mathcal{A} \rightarrow \mathbb{R}^+$ and a constant $\kappa > 0$ such that

$$\|G_{\text{stat}}(\cdot) - G_t(\cdot | s_0, z_0, u_0, a_0)\|_{\text{TV}} \leq \varphi(s_0, z_0, u_0, a_0) \cdot \exp(-2\kappa t),$$

where $G_t(\cdot | s_0, z_0, u_0, a_0)$ is the marginal distribution of (S_t, Z_t, U_t, A_t) given $(S_0, Z_0, U_0, A_0) = (s_0, z_0, u_0, a_0)$ under the behavior policy b . Further, we have $\int \varphi(s, z, u, a) d\nu(s, z, u, a) \leq c$ and $\int \varphi(s, z, u, a) dG_{\text{stat}}(s, z, u, a) \leq c$ for some positive absolute constant c .

Assumption 4.4.3 states that the the Markov chain $\{S_t, Z_t, U_t, A_t\}_{t \geq 0}$ mixes geometrically. Such an assumption is widely adopted in the related literature (Van Roy, 1998; Liao et al., 2020, 2021b; Wang et al., 2021a) to deal with dependent data.

To establish the upper bounds for the sub-optimality of the resulting policies, we need to first show that in our proposed algorithms, there exists at least one feasible solution that satisfies the constraints with properly chosen constants. In the following, we focus on $\text{conf}_{\alpha_0}^0$ and $\text{conf}_{\alpha_1}^1$ for Δ^* and Θ^* respectively. Since $\text{conf}_{\alpha_0}^0$ and $\text{conf}_{\alpha_1}^1$ can be constructed by many methods, to keep our theoretical results general, we assume that there exists a proper choice of (α_0, α_1) that ensures $\Delta^* \in \text{conf}_{\alpha_0}^0$ and $\Theta^* \in \text{conf}_{\alpha_1}^1$ and then give a valid example that justifies this assumption.

Assumption 4.4.4. There exists (α_0, α_1) such that with probability at least $1 - \delta$, we have

$$\Delta^* \in \text{conf}_{\alpha_0}^0, \quad \Theta^* \in \text{conf}_{\alpha_1}^1.$$

Further, with probability at least $1 - \delta$, for any $(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1$, we have

$$\begin{aligned} \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} \|\Delta^*(S_t, \cdot) - \Delta(S_t, \cdot)\|_1^2 \right] &\leq \xi_0^2 \frac{C_{\Delta^*}}{NT\kappa} \cdot \mathfrak{C}_{\mathcal{F}_0} \log \frac{2}{\delta}, \\ \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} \|\Theta^*(S_t, \cdot) - \Theta(S_t, \cdot)\|_1^2 \right] &\leq \xi_1^2 \frac{C_{\Theta^*}}{NT\kappa} \cdot \mathfrak{C}_{\mathcal{F}_1} \log \frac{2}{\delta}. \end{aligned}$$

We now illustrate that Assumption 4.4.4 can be realized via maximum likelihood estimation (MLE) by replacing ξ_0 and ξ_1 with proper quantities. Note that the estimation of Δ^* can be decomposed into the estimation of $\mathbb{P}(A = a | S = s, Z = z)$ for all $z \in \mathcal{Z}$, which can also be obtained via MLE. This implies that estimating Δ^* is similar to estimating Θ^* . Therefore, we only show how to estimate Θ^* so that Assumption 4.4.4 holds for the simplicity of presentation. By maximum likelihood, we construct the loss

function \widehat{L}_1 and the estimator $\widehat{\Theta}$ as follows,

$$\widehat{L}_1(\Theta) = -\widehat{\mathbb{E}} \left[\frac{1}{T} \sum_{t=0}^{T-1} \log \Theta(S_t, Z_t) \right] = -\frac{1}{NT} \sum_{i \in [N]} \sum_{t=0}^{T-1} \log \Theta(S_t^i, Z_t^i), \quad \widehat{\Theta} \in \underset{\Theta \in \mathcal{F}_1}{\operatorname{argmin}} \widehat{L}_1(\Theta),$$

where for the simplicity of notations, we denote by $\widehat{\mathbb{E}}[\cdot]$ the empirical measure generated by the offline data \mathcal{D} hereafter. In addition, we assume that \mathcal{F}_1 is a parametric class such that $\mathcal{F}_1 = \{\Theta_\theta : \theta \in \mathbb{R}^d, \|\theta\|_2 \leq \theta_{\max}\}$. We introduce the following results.

Theorem 4.4.5. Suppose $\mathcal{F}_1 = \{\Theta_\theta : \theta \in \mathbb{R}^d \text{ and } \|\theta\|_2 \leq \theta_{\max}\}$, and

$$\alpha_1 = c \cdot \frac{C_{\Theta^*}}{NT\kappa} \cdot d \log \frac{\theta_{\max}}{\delta} \log(NT),$$

where $c/(N^2T^2) \cdot \log(NT) \leq \delta \leq 1$. Then under Assumptions 4.2.2, 4.4.2(b), 4.4.2(c), and 4.4.3, it holds with probability at least $1 - \delta$ that $\Theta^* \in \operatorname{conf}_{\alpha_1}^1$. Further, with probability at least $1 - \delta$, it holds for any $\Theta \in \operatorname{conf}_{\alpha_1}^1$ that

$$\sqrt{\mathbb{E} [\|\Theta(S, \cdot) - \Theta^*(S, \cdot)\|_1^2]} \leq c \sqrt{\frac{C_{\Theta^*}}{NT\kappa} \cdot d \log \frac{\theta_{\max}}{\delta}}.$$

Proof. See §C.1.1 for a detailed proof. □

Supported by Theorem 4.4.5, we assume Assumption 4.4.4 holds throughout this section.

4.4.1. Theoretical Results for VF-based Pessimistic Method

We first impose the following assumption, which assumes that V^π is realizable in \mathcal{V} for any policy π , and w^{π^*} is realizable in \mathcal{W} only for the optimal policy π^* .

Assumption 4.4.6. We have $V^\pi \in \mathcal{V}$ for any $\pi \in \Pi$ and $w^{\pi^*} \in \mathcal{W}$. Further, we have $-w \in \mathcal{W}$ for any $w \in \mathcal{W}$.

In the following lemma, we show that with a proper choice of α_{vf} , we have $V^\pi \in \text{conf}_{\alpha_{\text{vf}}}^{\text{vf}}(\Delta^*, \Theta^*, \pi)$ with a high probability.

Lemma 4.4.7. Suppose

$$\alpha_{\text{vf}} = c \cdot \frac{C_{\Delta^*} C_{\Theta^*} C_*}{1 - \gamma} \sqrt{\frac{\mathfrak{C}_{\mathcal{W}, \mathcal{V}, \Pi}}{NT\kappa} \cdot \log \frac{1}{\delta} \log(NT)}$$

and $c/(NT)^2 \leq \delta \leq 1$. Then under Assumptions 4.4.2 and 4.4.6, with probability at least $1 - \delta$, it holds for any $\pi \in \Pi$ that $V^\pi \in \text{conf}_{\alpha_{\text{vf}}}^{\text{vf}}(\Delta^*, \Theta^*, \pi)$.

Proof. See §C.3.2 for a detailed proof. \square

In the following lemma, we show that for any $v \in \cup_{(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1} \text{conf}_{\alpha_{\text{vf}}}^{\text{vf}}(\Delta, \Theta, \pi)$, we can upper bound the risk $\max_{g \in \mathcal{W}} \Phi_{\text{vf}}^\pi(v, g; \Delta^*, \Theta^*)$.

Lemma 4.4.8. Suppose that $(\alpha_0, \alpha_1, \alpha_{\text{vf}})$ is defined in Assumption 4.4.4 and Lemma 4.4.7 and $c/(NT)^2 \leq \delta \leq 1$. Then under Assumptions 4.2.2, 4.2.3, and 4.4.2–4.4.6, with probability at least $1 - \delta$, it holds for any policy $\pi \in \Pi$ and $v \in \cup_{(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1} \text{conf}_{\alpha_{\text{vf}}}^{\text{vf}}(\Delta, \Theta, \pi)$ that

$$\max_{g \in \mathcal{W}} \Phi_{\text{vf}}^\pi(v, g; \Delta^*, \Theta^*) \leq c \cdot \frac{C_{\Delta^*}^2 C_{\Theta^*}^2 C_*}{1 - \gamma} (\xi_0 + \xi_1) L_\Pi \sqrt{\frac{1}{NT\kappa} \cdot \mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}, \Pi} \cdot \log \frac{1}{\delta} \log(NT)}.$$

Proof. See §C.3.3 for a detailed proof. \square

Equipped with the above results, we introduce the following theorem, which characterizes the suboptimality of the learned policy $\widehat{\pi}_{\text{vf}}$ constructed in (4.3.3).

Theorem 4.4.9. Suppose $c/(NT)^2 \leq \delta \leq 1$. Under Assumptions 4.2.2–4.4.6, it holds with probability at least $1 - \delta$ that

$$\text{SubOpt}(\widehat{\pi}_{\text{vf}}) \leq c \cdot \frac{C_{\Delta^*}^2 C_{\Theta^*}^2 C^*}{1 - \gamma} (\xi_0 + \xi_1) L_{\Pi} \sqrt{\frac{1}{NT\kappa} \cdot \mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \nu, \Pi} \cdot \log \frac{1}{\delta} \log(NT)}.$$

PROOF SKETCH. In the proof sketch, we assume that we have full knowledge on Δ^* and Θ^* . By the definition of $J(\pi)$ in (4.1.2), we have

$$\begin{aligned} J(\pi^*) - J(\widehat{\pi}_{\text{vf}}) &= (1 - \gamma) \mathbb{E}_{S_0 \sim \nu} [V^{\pi^*}(S_0) - V^{\widehat{\pi}_{\text{vf}}}(S_0)] \\ &\leq (1 - \gamma) \mathbb{E}_{S_0 \sim \nu} [V^{\pi^*}(S_0)] - \min_{v \in \text{conf}_{\alpha_{\text{vf}}}^{\text{vf}}(\Delta^*, \Theta^*, \widehat{\pi}_{\text{vf}})} (1 - \gamma) \mathbb{E}_{S_0 \sim \nu} [v(S_0)] \\ &\leq (1 - \gamma) \mathbb{E}_{S_0 \sim \nu} [V^{\pi^*}(S_0)] - \min_{v \in \text{conf}_{\alpha_{\text{vf}}}^{\text{vf}}(\Delta^*, \Theta^*, \pi^*)} (1 - \gamma) \mathbb{E}_{S_0 \sim \nu} [v(S_0)] \\ &\leq (1 - \gamma) \cdot \max_{v \in \text{conf}_{\alpha_{\text{vf}}}^{\text{vf}}(\Delta^*, \Theta^*, \pi^*)} |\mathbb{E}_{S_0 \sim \nu} [V^{\pi^*}(S_0) - v(S_0)]|, \end{aligned}$$

where in the first inequality, we use Lemma 4.4.7 that $V^{\widehat{\pi}_{\text{vf}}} \in \text{conf}_{\alpha_{\text{vf}}}^{\text{vf}}(\Delta^*, \Theta^*, \widehat{\pi}_{\text{vf}})$ with a high probability; while in the second inequality, we use the optimality of $\widehat{\pi}_{\text{vf}}$. In the meanwhile, by Lemmas 4.2.7 and 4.2.8, we have the following decomposition,

$$(4.4.1) \quad \begin{aligned} (1 - \gamma) \mathbb{E}_{S_0 \sim \nu} [V^{\pi^*}(S_0)] &= J(\pi^*) = \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} w^{\pi^*}(S_t) \frac{Z_t^\top A_t \pi^*(A_t | S_t)}{\Delta^*(S_t, A_t) \Theta^*(S_t, Z_t)} R_t \right], \\ (1 - \gamma) \mathbb{E}_{S_0 \sim \nu} [v(S_0)] &= \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} w^{\pi^*}(S_t) \frac{Z_t^\top A_t \pi^*(A_t | S_t)}{\Delta^*(S_t, A_t) \Theta^*(S_t, Z_t)} (v(S_t) - \gamma v(S_{t+1})) \right]. \end{aligned}$$

Now, by plugging (4.4.1), we have

$$J(\pi^*) - J(\widehat{\pi}_{\text{vf}}) \leq \max_{v \in \text{conf}_{\alpha_{\text{vf}}}^{\text{vf}}(\Delta^*, \Theta^*, \pi^*)} \left| \Phi_{\text{vf}}^{\pi^*}(v, w^{\pi^*}; \Delta^*, \Theta^*) \right|.$$

We then can upper bound the above suboptimality by Lemma 4.4.8, which concludes the proof of the theorem. See §C.3.1 for a detailed proof. \square

In Theorem 4.4.9, we impose data coverage and realizability assumptions as in Assumption 4.4.6, which only requires that the offline data covers the trajectory generated by the optimal policy π^* and V^π is realizable in \mathcal{V} for any π .

4.4.2. Theoretical Results for MIS-based Pessimistic Method

We first impose the following assumption, which assumes that w^π is realizable in \mathcal{W} for any policy π , and V^{π^*} is realizable in \mathcal{V} only for the optimal policy π^* .

Assumption 4.4.10. We have $w^\pi \in \mathcal{W}$ for any $\pi \in \Pi$ and $V^{\pi^*} \in \mathcal{V}$. Further, we have $-v \in \mathcal{V}$ for any $v \in \mathcal{V}$.

In the following lemma, we show that with a proper choice of α_{mis} , we have $w^\pi \in \text{conf}_{\alpha_{\text{mis}}}^{\text{mis}}(\Delta^*, \Theta^*, \pi)$ with high probability.

Lemma 4.4.11. Suppose

$$\alpha_{\text{mis}} = c \cdot \frac{C_{\Delta^*} C_{\Theta^*} C_*}{1 - \gamma} \sqrt{\frac{1}{NT\kappa} \mathfrak{C}_{\mathcal{V}, \mathcal{W}, \Pi} \log \frac{1}{\delta} \log(NT)}$$

and $c/(NT)^2 \leq \delta \leq 1$. Then under Assumptions 4.4.2 and 4.4.10, with probability at least $1 - \delta$, it holds for any $\pi \in \Pi$ that $w^\pi \in \text{conf}_{\alpha_{\text{mis}}}^{\text{mis}}(\Delta^*, \Theta^*, \pi)$.

Proof. See §C.4.2 for a detailed proof. \square

In the following lemma, we show that for any $w \in \cup_{(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1} \text{conf}_{\alpha_{\text{mis}}}^{\text{mis}}(\Delta, \Theta, \pi)$, we can upper bound the risk $\max_{f \in \mathcal{V}} \Phi_{\text{mis}}^\pi(w, f; \Delta^*, \Theta^*)$.

Lemma 4.4.12. Suppose that $(\alpha_0, \alpha_1, \alpha_{\text{mis}})$ is defined in Assumption 4.4.4 and Lemma 4.4.11, and $c/(NT)^2 \leq \delta \leq 1$. Then under Assumptions 4.2.2, 4.2.6–4.4.4, and 4.4.10, with probability at least $1 - \delta$, it holds for any $\pi \in \Pi$ and $w \in \cup_{(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1} \text{conf}_{\alpha_{\text{mis}}}^{\text{mis}}(\Delta, \Theta, \pi)$ that

$$\max_{f \in \mathcal{V}} \Phi_{\text{mis}}^\pi(w, f; \Delta^*, \Theta^*) \leq c \cdot \frac{C_{\Delta^*}^2 C_{\Theta^*}^2 C_*}{1 - \gamma} (\xi_0 + \xi_1) \sqrt{\frac{1}{NT\kappa} \cdot \mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}, \Pi} \cdot \log \frac{1}{\delta} \log(NT)}.$$

Proof. See §C.4.3 for a detailed proof. □

Equipped with the above results, we introduce the following theorem, which characterizes the suboptimality of the learned policy $\hat{\pi}_{\text{mis}}$ constructed in (4.3.6).

Theorem 4.4.13. Suppose $c/(NT)^2 \leq \delta \leq 1$. Under Assumptions 4.2.2, 4.2.6–4.4.4, and 4.4.10, it holds with probability at least $1 - \delta$ that

$$\text{SubOpt}(\hat{\pi}_{\text{mis}}) \leq c \cdot \frac{C_{\Delta^*}^2 C_{\Theta^*}^2 C_*}{1 - \gamma} (\xi_0 + \xi_1) \sqrt{\frac{1}{NT\kappa} \cdot \mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}, \Pi} \cdot \log \frac{1}{\delta} \log(NT)}.$$

Proof. See §C.4.1 for a detailed proof. □

In Theorem 4.4.13, we impose data coverage and realizability assumptions as in Assumption 4.4.10, which only requires that V^{π^*} is realizable in \mathcal{V} and the offline data covers the trajectory generated by the policy π for any $\pi \in \Pi$.

4.4.3. Theoretical Results for DR-based Pessimistic Method

Theorem 4.4.14. Suppose that $(\alpha_0, \alpha_1, \alpha_{\text{mis}}, \alpha_{\text{vf}})$ is defined in Assumption 4.4.4, Lemmas 4.4.7, 4.4.11, and at least one of Assumptions 4.4.6 and 4.4.10 hold. Then under Assumptions 4.2.2–4.4.4, it holds with probability at least $1 - \delta$ for any $c/(NT)^2 \leq \delta \leq 1$ that

$$\text{SubOpt}(\hat{\pi}_{\text{dr}}) \leq c \cdot \frac{C_{\Delta^*}^2 C_{\Theta^*}^2 C_*}{1 - \gamma} (\xi_0 + \xi_1) \sqrt{\frac{1}{NT\kappa} \mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}, \Pi} \log \frac{1}{\delta} \log(NT)},$$

where $\hat{\pi}_{\text{dr}}$ is defined in (4.3.8).

Proof. See §C.5.1 for a detailed proof. □

Theorem 4.4.14 shows that $\hat{\pi}_{\text{dr}}$ is a doubly robust estimator of the optimal policy in the sense that either Assumption 4.4.6 or Assumption 4.4.10 ensures the convergence of $\hat{\pi}_{\text{dr}}$.

Our results in §4.3.1–§4.3.3 hinge on the data coverage and realizability assumptions. Can we obtain a similar upper bound if such assumptions are violated? In this section, we answer this question affirmatively. First, we introduce the following assumption. We denote by

$$(4.4.2) \quad \tilde{v}^\pi \in \operatorname{argmin}_{v \in \mathcal{V}} \max_{w \in \mathcal{W}} \Phi_{\text{vf}}^\pi(v, w; \Delta^*, \Theta^*), \quad \tilde{w}^\pi \in \operatorname{argmin}_{w \in \mathcal{W}} \max_{v \in \mathcal{V}} \Phi_{\text{mis}}^\pi(w, v; \Delta^*, \Theta^*).$$

Assumption 4.4.15 (Model Misspecification). The following statements hold.

- (a) We have $\|V^\pi - \tilde{v}^\pi\|_\infty \leq \varepsilon_{\text{vf}}^{\mathcal{V}}$ for any $\pi \in \Pi$ and $\|w^{\pi^*} - \tilde{w}^{\pi^*}\|_\infty \leq \varepsilon_{\text{vf}}^{\mathcal{W}}$.
- (b) We have $\|w^\pi - \tilde{w}^\pi\|_\infty \leq \varepsilon_{\text{mis}}^{\mathcal{W}}$ for any $\pi \in \Pi$ and $\|V^{\pi^*} - \tilde{v}^{\pi^*}\|_\infty \leq \varepsilon_{\text{mis}}^{\mathcal{V}}$.

Though Assumption 4.4.15 requires that (a) and (b) hold simultaneously, we remark that previous assumptions imposed in VF-, MIS-, and DR-based pessimism can be recovered by such an assumption. Specifically, Assumptions 4.4.6 and 4.4.10 can be recovered by taking $(\varepsilon_{\text{vf}}^{\mathcal{V}}, \varepsilon_{\text{vf}}^{\mathcal{W}}, \varepsilon_{\text{mis}}^{\mathcal{V}}, \varepsilon_{\text{mis}}^{\mathcal{W}}) = (0, 0, \infty, \infty)$ and $(\varepsilon_{\text{vf}}^{\mathcal{V}}, \varepsilon_{\text{vf}}^{\mathcal{W}}, \varepsilon_{\text{mis}}^{\mathcal{V}}, \varepsilon_{\text{mis}}^{\mathcal{W}}) = (\infty, \infty, 0, 0)$, respectively, in Assumption 4.4.15. Similarly, the data coverage and realizability assumptions in Theorem 4.4.14 can also be recovered by either taking $(\varepsilon_{\text{vf}}^{\mathcal{V}}, \varepsilon_{\text{vf}}^{\mathcal{W}}, \varepsilon_{\text{mis}}^{\mathcal{V}}, \varepsilon_{\text{mis}}^{\mathcal{W}}) = (0, 0, \infty, \infty)$ or taking $(\varepsilon_{\text{vf}}^{\mathcal{V}}, \varepsilon_{\text{vf}}^{\mathcal{W}}, \varepsilon_{\text{mis}}^{\mathcal{V}}, \varepsilon_{\text{mis}}^{\mathcal{W}}) = (\infty, \infty, 0, 0)$.

Theorem 4.4.16. Suppose that $(\alpha_0, \alpha_1, \alpha_{\text{mis}}, \alpha_{\text{vf}})$ is defined in Assumption 4.4.4, Lemma 4.4.7, and Lemma 4.4.11. Then under Assumptions 4.2.2–4.4.4, and 4.4.15, it holds with probability at least $1 - \delta$ for any $c/(NT)^2 \leq \delta \leq 1$ that

$$\begin{aligned} \text{SubOpt}(\widehat{\pi}_{\text{dr}}) \leq & c \cdot \frac{C_{\Delta^*}^2 C_{\Theta^*}^2 C_*}{1 - \gamma} (\xi_0 + \xi_1) \sqrt{\frac{1}{NT\kappa} \mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \nu, \Pi} \log \frac{NT}{\delta}} \\ & + 3C_{\Delta^*} C_{\Theta^*} \min \{ C_* \varepsilon_{\text{vf}}^{\mathcal{V}} + \varepsilon_{\text{vf}}^{\mathcal{W}} / (1 - \gamma), C_* \varepsilon_{\text{mis}}^{\mathcal{V}} + \varepsilon_{\text{mis}}^{\mathcal{W}} / (1 - \gamma) \}, \end{aligned}$$

where $\widehat{\pi}_{\text{dr}}$ is defined in (4.3.8).

Proof. See §C.5.2 for a detailed proof. \square

Due to the model misspecification, compared with Theorem 4.4.14, there is an additional bias term on the upper bound of the suboptimality in Theorem 4.4.16. We remark that either $(\varepsilon_{\text{vf}}^{\mathcal{V}}, \varepsilon_{\text{vf}}^{\mathcal{W}}) = (0, 0)$ or $(\varepsilon_{\text{mis}}^{\mathcal{V}}, \varepsilon_{\text{mis}}^{\mathcal{W}}) = (0, 0)$ ensures zero bias in Theorem 4.4.16.

4.5. Dual Formulation

To improve the computational efficiency of estimating the optimal in-class policy due to the confidence sets, we propose a dual formulation of the aforementioned pessimistic

methods. For illustration purpose, we only consider the dual formulation of the VF-based pessimistic method proposed in §4.3.1. Similar formulations for MIS-based and DR-based methods can also be derived accordingly.

For the ease of presentation, we assume that there exists an oracle that gives us Δ^* and Θ^* . Without the existence of such an oracle, we only need to employ two additional dual variables to consider the uncertainty induced by estimating Δ^* and Θ^* .

We consider the following dual form of (4.3.3),

$$(4.5.1) \quad \begin{aligned} \widehat{\pi}_{\text{vf}}^\dagger &= \operatorname{argmax}_{\pi \in \Pi} \max_{\lambda \geq 0} \min_{v \in \mathcal{V}} (1 - \gamma) \mathbb{E}_{S \sim \nu} [v(S)] + \lambda \cdot \left(\widehat{M}_{\text{vf}}^\pi(v) - \alpha_{\text{vf}} \right), \\ \text{s.t. } \widehat{M}_{\text{vf}}^\pi(v) &= \max_{g \in \mathcal{W}} \widehat{\Phi}_{\text{vf}}^\pi(v, g; \Delta^*, \Theta^*) - \max_{g \in \mathcal{W}} \widehat{\Phi}_{\text{vf}}^\pi(\widehat{v}_{\Delta^*, \Theta^*}^\pi, g; \Delta^*, \Theta^*), \end{aligned}$$

where $\widehat{v}_{\Delta^*, \Theta^*}^\pi = \operatorname{argmin}_{v \in \mathcal{V}} \max_{g \in \mathcal{W}} \widehat{\Phi}_{\text{vf}}^\pi(v, g; \Delta^*, \Theta^*)$ and λ is the dual variable that corresponds to the constraint $v \in \operatorname{conf}_{\alpha_{\text{vf}}}^{\text{vf}}(\Delta^*, \Theta^*, \pi)$. In comparison to the constrained optimization problem in (4.3.3), the problem in (4.5.1) can be solved efficiently using gradient methods.

In the following theorem, we characterize the suboptimality of $\widehat{\pi}_{\text{vf}}^\dagger$.

Theorem 4.5.1. Suppose that \mathcal{V} is convex, $c/(NT)^2 \leq \delta \leq 1$, and α_{vf} is defined in Lemma 4.4.8. Under Assumptions 4.2.2–4.4.6, it holds with probability at least $1 - \delta$ that

$$\operatorname{SubOpt}(\widehat{\pi}_{\text{vf}}^\dagger) \leq c \cdot \frac{C_{\Delta^*}^2 C_{\Theta^*}^2 C_*}{1 - \gamma} (\xi_0 + \xi_1) L_\Pi \sqrt{\frac{1}{NT\kappa} \cdot \mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}, \Pi} \cdot \log \frac{1}{\delta} \log(NT)}.$$

Proof. See §C.6.1 for a detailed proof. □

In Theorem 4.5.1, we show a similar suboptimality holds as in Theorem 4.4.9. Thus, to avoid computation challenges induced by the confidence sets, we only need to solve

(4.5.1) to obtain an optimal policy. We remark similar dual formulations for MIS-based and DR-based methods also hold, as well as their theoretical properties.

4.6. Identifiability

We discuss identifiability in this section. First, we use tabular MDP as an example to illustrate non-identifiability. Then we discuss the identifiability in our methods.

Non-Identifiability in Tabular MDP. We consider a tabular MDP with states $\mathcal{S} = \{s_1, s_2, \dots, s_{|\mathcal{S}|}\}$, where the behavior policy b used to generate offline data only covers states $\{s_2, \dots, s_{|\mathcal{S}|}\}$. We assume that the expected total reward under such a tabular MDP is $J(\pi)$ for any policy π . Since the offline data generated following b never covers the state s_1 , we cannot infer any information of the reward received at the state s_1 . Thus, for any policy π that arrives the state s_1 with a nonzero probability, we cannot identify the value $J(\pi)$ uniquely. In the meanwhile, the state-value function $V^\pi: \mathcal{S} \rightarrow \mathbb{R}$ is not uniquely identifiable for any policy π (even for π^*), since the value $V^\pi(s_1)$ is not identifiable.

Identifiability in Our Methods. In §4.2 and §4.3, we do not explicitly impose any identifiability assumptions. But certain identifiability assumptions are implied by the data coverage assumptions as follows.

- VF-based pessimism. As imposed in Assumption 4.4.6, we require that w^{π^*} is upper bounded. Thus, we know that the trajectory generated by the optimal policy π^* is covered by the offline data, which implies that $J(\pi^*)$ is identifiable.
- MIS-based pessimism. As imposed in Assumption 4.4.10, we require that w^π is upper bounded for any policy π . Thus, we know that the trajectory generated by

any policy π is covered by the offline data, which implies that $J(\pi)$ is identifiable for any π .

- DR-based pessimism. Since either Assumption 4.4.6 or Assumption 4.4.10 hold, we require that $J(\pi^*)$ is identifiable or $J(\pi)$ is identifiable for any π .

It is worth noting that though we impose realizability assumptions on the state-value function V^π in Assumptions 4.4.6 and 4.4.10, we do not require that V^π is identifiable for any π (even for π^*). Such non-identifiability implies that the IV-aided Bellman equation in Lemma 4.2.4 may have multiple fixed point solutions.

References

- Abbasi-Yadkori, Y., Bartlett, P., Bhatia, K., Lazic, N., Szepesvari, C. and Weisz, G. (2019a). Politex: Regret bounds for policy iteration using expert prediction. In *International Conference on Machine Learning*.
- Abbasi-Yadkori, Y., Lazic, N., Szepesvari, C. and Weisz, G. (2019b). Exploration-enhanced politex. *arXiv preprint arXiv:1908.10479*.
- Agarwal, A., Kakade, S. M., Lee, J. D. and Mahajan, G. (2019). Optimality and approximation with policy gradient methods in Markov decision processes. *arXiv preprint arXiv:1908.00261*.
- Agostinelli, F., McAleer, S., Shmakov, A. and Baldi, P. (2019). Solving the Rubik’s cube with deep reinforcement learning and search. *Nature Machine Intelligence*, **1** 356–363.
- Akkaya, I., Andrychowicz, M., Chociej, M., Litwin, M., McGrew, B., Petron, A., Paino, A., Plappert, M., Powell, G., Ribas, R. et al. (2019). Solving Rubik’s cube with a robot hand. *arXiv preprint arXiv:1910.07113*.
- Alizadeh, F., Haeberly, J.-P. A. and Overton, M. L. (1998). Primal-dual interior-point methods for semidefinite programming: convergence rates, stability and numerical results. *SIAM Journal on Optimization*, **8** 746–768.
- Allen-Zhu, Z., Li, Y. and Liang, Y. (2018a). Learning and generalization in overparameterized neural networks, going beyond two layers. *arXiv preprint arXiv:1811.04918*.

- Allen-Zhu, Z., Li, Y. and Song, Z. (2018b). A convergence theory for deep learning via over-parameterization. *arXiv preprint arXiv:1811.03962*.
- Amari, S.-I. (1998). Natural gradient works efficiently in learning. *Neural Computation*, **10** 251–276.
- Anderson, B. D. and Moore, J. B. (2007). *Optimal control: linear quadratic methods*. Courier Corporation.
- Angrist, J. and Imbens, G. (1995). Identification and estimation of local average treatment effects.
- Angrist, J. D., Imbens, G. W. and Rubin, D. B. (1996). Identification of causal effects using instrumental variables. *Journal of the American statistical Association*, **91** 444–455.
- Antos, A., Szepesvári, C. and Munos, R. (2007). Value-iteration based fitted policy iteration: learning with a single trajectory. In *IEEE International Symposium on Approximate Dynamic Programming and Reinforcement Learning*. IEEE.
- Antos, A., Szepesvári, C. and Munos, R. (2008a). Fitted Q-iteration in continuous action-space mdps. In *Advances in neural information processing systems*.
- Antos, A., Szepesvári, C. and Munos, R. (2008b). Learning near-optimal policies with Bellman-residual minimization based fitted policy iteration and a single sample path. *Machine Learning*, **71** 89–129.
- Araki, B., Strang, J., Pohorecky, S., Qiu, C., Naegeli, T. and Rus, D. (2017). Multi-robot path planning for a swarm of robots that can both fly and drive. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE.

- Arora, S., Du, S. S., Hu, W., Li, Z. and Wang, R. (2019). Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks. *arXiv preprint arXiv:1901.08584*.
- Ash, C. (2000). Social-self-interest. *Annals of public and cooperative economics*, **71** 261–284.
- Axtell, R. L. (2002). Non-cooperative dynamics of multi-agent teams. In *Autonomous Agents and Multiagent Systems*.
- Bardi, M. (2011). Explicit solutions of some linear-quadratic mean field games. *Networks and heterogeneous media*, **7** 243–261.
- Bardi, M. and Priuli, F. S. (2014). Linear-quadratic n -person and mean-field games with ergodic cost. *SIAM Journal on Control and Optimization*, **52** 3022–3052.
- Barrera, D. and Gobet, E. (2021). Generalization bounds for nonparametric regression with β -mixing samples. *arXiv preprint arXiv:2108.00997*.
- Bauso, D., Tembine, H. and Başar, T. (2016). Robust mean field games. *Dynamic games and applications*, **6** 277–303.
- Bellemare, M. G., Naddaf, Y., Veness, J. and Bowling, M. (2013). The Arcade learning environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research*, **47** 253–279.
- Bennett, A., Kallus, N., Li, L. and Mousavi, A. (2021). Off-policy evaluation in infinite-horizon reinforcement learning with latent confounders. In *International Conference on Artificial Intelligence and Statistics*. PMLR.

- Bensoussan, A., Chau, M., Lai, Y. and Yam, S. C. P. (2017). Linear-quadratic mean field Stackelberg games with state and control delays. *SIAM Journal on Control and Optimization*, **55** 2748–2781.
- Bensoussan, A., Frehse, J. and Yam, P. (2013). *Mean field games and mean field type control theory*. Springer.
- Bensoussan, A., Sung, K., Yam, S. C. P. and Yung, S.-P. (2016). Linear-quadratic mean field games. *Journal of Optimization Theory and Applications*, **169** 496–529.
- Berbee, H. C. P. (1979). Random walks with stationary increments and renewal theory.
- Bhandari, J. and Russo, D. (2019). Global optimality guarantees for policy gradient methods. *arXiv preprint arXiv:1906.01786*.
- Bhandari, J. and Russo, D. (2020). A note on the linear convergence of policy gradient methods. *arXiv preprint arXiv:2007.11120*.
- Bhandari, J., Russo, D. and Singal, R. (2018). A finite time analysis of temporal difference learning with linear function approximation. *arXiv preprint arXiv:1806.02450*.
- Bhatnagar, S. (2010). An actor-critic algorithm with function approximation for discounted cost constrained Markov Decision Processes. *Systems & Control Letters*, **59** 760–766.
- Bhatnagar, S., Ghavamzadeh, M., Lee, M. and Sutton, R. S. (2008). Incremental natural actor-critic algorithms. In *Advances in Neural Information Processing Systems*.
- Bhatnagar, S., Sutton, R. S., Ghavamzadeh, M. and Lee, M. (2009). Natural actor-critic algorithms. *Automatica*, **45** 2471–2482.
- Biswas, A. (2015). Mean field games with ergodic cost for discrete time markov processes. *arXiv preprint arXiv:1510.08968*.

- Borkar, V. S. (2008). *Stochastic Approximation: A Dynamical Systems Viewpoint*. Cambridge University Press.
- Borkar, V. S. and Konda, V. R. (1997). The actor-critic algorithm as multi-time-scale stochastic approximation. *Sadhana*, **22** 525–543.
- Borkar, V. S. and Meyn, S. P. (2000). The ODE method for convergence of stochastic approximation and reinforcement learning. *SIAM Journal on Control and Optimization*, **38** 447–469.
- Bowling, M. (2001). Rational and convergent learning in stochastic games. In *International Conference on Artificial Intelligence*.
- Bowling, M. and Veloso, M. (2000). An analysis of stochastic game theory for multiagent reinforcement learning. Tech. rep., Carnegie Mellon University.
- Bradtke, S. J. (1993). Reinforcement learning applied to linear quadratic regulation. In *Advances in Neural Information Processing Systems*.
- Bradtke, S. J. and Barto, A. G. (1996). Linear least-squares algorithms for temporal difference learning. *Machine learning*, **22** 33–57.
- Bradtke, S. J., Ydstie, B. E. and Barto, A. G. (1994). Adaptive linear quadratic control using policy iteration. In *American Control Conference*, vol. 3. IEEE.
- Briani, A. and Cardaliaguet, P. (2018). Stable solutions in potential mean field game systems. *Nonlinear Differential Equations and Applications*, **25** 1.
- Brookhart, M. A., Stürmer, T., Glynn, R. J., Rassen, J. and Schneeweiss, S. (2010). Confounding control in healthcare database research: challenges and potential approaches. *Medical care*, **48** S114.

- Bu, J., Mesbahi, A., Fazel, M. and Mesbahi, M. (2019). LQR through the lens of first order methods: Discrete-time case. *arXiv preprint arXiv:1907.08921*.
- Busoniu, L., Babuska, R. and De Schutter, B. (2008). A comprehensive survey of multi-agent reinforcement learning. *IEEE Transactions on Systems, Man, and Cybernetics, Part C: Applications and Reviews*, **38** 156–172.
- Cai, Q., Yang, Z., Jin, C. and Wang, Z. (2019). Provably efficient exploration in policy optimization. *arXiv preprint arXiv:1912.05830* 1283–1294.
- Caines, P. E. and Kizilkale, A. C. (2017). ϵ -nash equilibria for partially observed LQG mean field games with a major player. *IEEE Transactions on Automatic Control*, **62** 3225–3234.
- Calderone, D. J. (2017). *Models of Competition for Intelligent Transportation Infrastructure: Parking, Ridesharing, and External Factors in Routing Decisions*. University of California, Berkeley.
- Cao, Y. and Gu, Q. (2019a). Generalization bounds of stochastic gradient descent for wide and deep neural networks. *arXiv preprint arXiv:1905.13210*.
- Cao, Y. and Gu, Q. (2019b). A generalization theory of gradient descent for learning over-parameterized deep ReLU networks. *arXiv preprint arXiv:1902.01384*.
- Carmona, R. and Delarue, F. (2013). Probabilistic analysis of mean-field games. *SIAM Journal on Control and Optimization*, **51** 2705–2734.
- Carmona, R. and Delarue, F. (2018). *Probabilistic Theory of Mean Field Games with Applications I-II*. Springer.
- Casgrain, P., Ning, B. and Jaimungal, S. (2019). Deep Q-learning for Nash equilibria: Nash-DQN. *arXiv preprint arXiv:1904.10554*.

- Castro, D. D. and Meir, R. (2010). A convergent online single time scale actor critic algorithm. *Journal of Machine Learning Research*, **11** 367–410.
- Chen, J. and Jiang, N. (2019). Information-theoretic considerations in batch reinforcement learning. *arXiv preprint arXiv:1905.00360*.
- Chizat, L. and Bach, F. (2018). A note on lazy training in supervised differentiable programming. *arXiv preprint arXiv:1812.07956*.
- Conitzer, V. and Sandholm, T. (2007). AWESOME: A general multiagent learning algorithm that converges in self-play and learns a best response against stationary opponents. *Machine Learning*, **67** 23–43.
- Daniely, A. (2017). SGD learns the conjugate kernel class of the network. In *Advances in Neural Information Processing Systems*.
- de Cote, E. M., Lazaric, A. and Restelli, M. (2006). Learning to cooperate in multi-agent social dilemmas. In *International Conference on Autonomous Agents and Multiagent Systems*. ACM.
- De Souza, C., Gevers, M. and Goodwin, G. (1986). Riccati equations in optimal filtering of nonstabilizable systems having singular state transition matrices. *IEEE Transactions on Automatic Control*, **31** 831–838.
- Dean, S., Mania, H., Matni, N., Recht, B. and Tu, S. (2017). On the sample complexity of the linear quadratic regulator. *arXiv preprint arXiv:1710.01688*.
- Dean, S., Mania, H., Matni, N., Recht, B. and Tu, S. (2018). Regret bounds for robust adaptive control of the linear quadratic regulator. In *Advances in Neural Information Processing Systems*.

- Deng, Z., Fu, Z., Wang, L., Yang, Z., Bai, C., Wang, Z. and Jiang, J. (2021). Score: Spurious correlation reduction for offline reinforcement learning. *arXiv preprint arXiv:2110.12468*.
- Doerr, B., Linares, R., Zhu, P. and Ferrari, S. (2018). Random finite set theory and optimal control for large spacecraft swarms. *arXiv preprint arXiv:1810.00696*.
- Du, S. S., Chen, J., Li, L., Xiao, L. and Zhou, D. (2017). Stochastic variance reduction methods for policy evaluation. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org.
- Du, S. S., Lee, J. D., Li, H., Wang, L. and Zhai, X. (2018). Gradient descent finds global minima of deep neural networks. *arXiv preprint arXiv:1811.03804*.
- Ernst, D., Geurts, P. and Wehenkel, L. (2005). Tree-based batch mode reinforcement learning. *Journal of Machine Learning Research*, **6** 503–556.
- Fang, J. (2014). The LQR controller design of two-wheeled self-balancing robot based on the particle swarm optimization algorithm. *Mathematical Problems in Engineering*, **2014**.
- Farahmand, A.-m., Ghavamzadeh, M., Szepesvári, C. and Mannor, S. (2016). Regularized policy iteration with nonparametric function spaces. *The Journal of Machine Learning Research*, **17** 4809–4874.
- Farahmand, A.-m., Szepesvári, C. and Munos, R. (2010). Error propagation for approximate policy and value iteration. In *Advances in Neural Information Processing Systems*.
- Fazel, M., Ge, R., Kakade, S. M. and Mesbahi, M. (2018). Global convergence of policy gradient methods for the linear quadratic regulator. *arXiv preprint arXiv:1801.05039*.

- Fu, Z., Yang, Z., Chen, Y. and Wang, Z. (2019). Actor-critic provably finds Nash equilibria of linear-quadratic mean-field games. *arXiv preprint arXiv:1910.07498*.
- Ganzfried, S. and Sandholm, T. (2009). Computing equilibria in multiplayer stochastic games of imperfect information. In *Twenty-First International Joint Conference on Artificial Intelligence*.
- Gao, R., Cai, T., Li, H., Wang, L., Hsieh, C.-J. and Lee, J. D. (2019). Convergence of adversarial training in overparametrized networks. *arXiv preprint arXiv:1906.07916*.
- Geer, S. A., van de Geer, S. and Williams, D. (2000). *Empirical Processes in M-estimation*, vol. 6. Cambridge university press.
- Geist, M., Scherrer, B. and Pietquin, O. (2019). A theory of regularized markov decision processes. *arXiv preprint arXiv:1901.11275*.
- Gomes, D. A., Mohr, J. and Souza, R. R. (2010). Discrete time, finite state space mean field games. *Journal de mathématiques pures et appliquées*, **93** 308–328.
- Gomes, D. A. et al. (2014). Mean field games models—a brief survey. *Dynamic Games and Applications*, **4** 110–154.
- Gottesman, O., Johansson, F., Komorowski, M., Faisal, A., Sontag, D., Doshi-Velez, F. and Celi, L. A. (2019). Guidelines for reinforcement learning in healthcare. *Nature medicine*, **25** 16–18.
- Guéant, O., Lasry, J.-M. and Lions, P.-L. (2011). Mean field games and applications. In *Paris-Princeton lectures on mathematical finance 2010*. Springer, 205–266.
- Guo, X., Hu, A., Xu, R. and Zhang, J. (2019). Learning mean-field games. *arXiv preprint arXiv:1901.09585*.

- Haarnoja, T., Zhou, A., Abbeel, P. and Levine, S. (2018). Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. *arXiv preprint arXiv:1801.01290*.
- Hao, B., Lazic, N., Abbasi-Yadkori, Y., Joulani, P. and Szepesvari, C. (2020). Provably efficient adaptive approximate policy iteration. *arXiv preprint arXiv:2002.03069*.
- Hardt, M., Ma, T. and Recht, B. (2016). Gradient descent learns linear dynamical systems. *arXiv preprint arXiv:1609.05191*.
- Heinrich, J. and Silver, D. (2016). Deep reinforcement learning from self-play in imperfect-information games. *arXiv preprint arXiv:1603.01121*.
- Hernán, M. A. and Robins, J. M. (2010). Causal inference: What if.
- Hong, M., Wai, H.-T., Wang, Z. and Yang, Z. (2020). A two-timescale framework for bilevel optimization: Complexity analysis and application to actor-critic. *arXiv preprint arXiv:2007.05170*.
- Hu, J. and Wellman, M. P. (1998). Multiagent reinforcement learning: Theoretical framework and an algorithm. In *International Conference on Machine Learning*. Morgan Kaufmann Publishers Inc.
- Hu, J. and Wellman, M. P. (2003). Nash Q-learning for general-sum stochastic games. *Journal of machine learning research*, **4** 1039–1069.
- Huang, J. and Huang, M. (2017). Robust mean field linear-quadratic-gaussian games with unknown L^2 -disturbance. *SIAM Journal on Control and Optimization*, **55** 2811–2840.
- Huang, J. and Li, N. (2018). Linear-quadratic mean-field game for stochastic delayed systems. *IEEE Transactions on Automatic Control*, **63** 2722–2729.

- Huang, J., Li, X. and Wang, T. (2016a). Mean-field linear-quadratic-Gaussian (LQG) games for stochastic integral systems. *IEEE Transactions on Automatic Control*, **61** 2670–2675.
- Huang, J., Wang, S. and Wu, Z. (2016b). Backward mean-field linear-quadratic-Gaussian (LQG) games: Full and partial information. *IEEE Transactions on Automatic Control*, **61** 3784–3796.
- Huang, M., Caines, P. E. and Malhamé, R. P. (2003). Individual and mass behaviour in large population stochastic wireless power control problems: centralized and nash equilibrium solutions. In *Conference on Decision and Control*. IEEE.
- Huang, M., Caines, P. E. and Malhamé, R. P. (2007). Large-population cost-coupled LQG problems with nonuniform agents: individual-mass behavior and decentralized ε -Nash equilibria. *IEEE transactions on automatic control*, **52** 1560–1571.
- Huang, M., Malhamé, R. P., Caines, P. E. et al. (2006). Large population stochastic dynamic games: closed-loop McKean-Vlasov systems and the Nash certainty equivalence principle. *Communications in Information & Systems*, **6** 221–252.
- Huang, M. and Zhou, M. (2019). Linear quadratic mean field games: Asymptotic solvability and relation to the fixed point approach. *arXiv preprint arXiv:1903.08776*.
- Hughes, E., Leibo, J. Z., Phillips, M., Tuyls, K., Dueñez-Guzman, E., Castañeda, A. G., Dunning, I., Zhu, T., McKee, K., Koster, R. et al. (2018). Inequity aversion improves cooperation in intertemporal social dilemmas. In *Advances in Neural Information Processing Systems*.

- Jacot, A., Gabriel, F. and Hongler, C. (2018). Neural tangent kernel: Convergence and generalization in neural networks. In *Advances in Neural Information Processing Systems*.
- Jayakumar, S. and Aditya, M. (2019). Reinforcement learning in stationary mean-field games. In *International Conference on Autonomous Agents and Multiagent Systems*.
- Jiang, N. and Huang, J. (2020). Minimax value interval for off-policy evaluation and policy optimization. *Advances in Neural Information Processing Systems*, **33** 2747–2758.
- Jiang, N. and Li, L. (2016). Doubly robust off-policy value evaluation for reinforcement learning. In *International Conference on Machine Learning*. PMLR.
- Kakade, S. M. (2002). A natural policy gradient. In *Advances in neural information processing systems*.
- Kalashnikov, D., Irpan, A., Pastor, P., Ibarz, J., Herzog, A., Jang, E., Quillen, D., Holly, E., Kalakrishnan, M., Vanhoucke, V. et al. (2018). Scalable deep reinforcement learning for vision-based robotic manipulation. In *Conference on Robot Learning*. PMLR.
- Kallus, N. and Uehara, M. (2020). Double reinforcement learning for efficient off-policy evaluation in markov decision processes. *J. Mach. Learn. Res.*, **21** 1–63.
- Kallus, N. and Uehara, M. (2022). Efficiently breaking the curse of horizon in off-policy evaluation with double reinforcement learning. *Operations Research*.
- Kallus, N. and Zhou, A. (2020). Confounding-robust policy evaluation in infinite-horizon reinforcement learning. *Advances in Neural Information Processing Systems*, **33** 22293–22304.

- Kidambi, R., Rajeswaran, A., Netrapalli, P. and Joachims, T. (2020). MOREL: Model-based offline reinforcement learning. *arXiv preprint arXiv:2005.05951*.
- Komorowski, M., Celi, L. A., Badawi, O., Gordon, A. C. and Faisal, A. A. (2018). The artificial intelligence clinician learns optimal treatment strategies for sepsis in intensive care. *Nature medicine*, **24** 1716–1720.
- Konda, V. R. and Tsitsiklis, J. N. (2000). Actor-critic algorithms. In *Advances in neural information processing systems*.
- Korda, N. and La, P. (2015). On TD(0) with function approximation: Concentration bounds and a centered variant with exponential convergence. In *International Conference on Machine Learning*.
- Kumar, A., Zhou, A., Tucker, G. and Levine, S. (2020). Conservative Q-learning for offline reinforcement learning. *arXiv preprint arXiv:2006.04779*.
- Kushner, H. and Yin, G. G. (2003). *Stochastic approximation and recursive algorithms and applications*, vol. 35. Springer Science & Business Media.
- Lagoudakis, M. G. and Parr, R. (2002). Value function approximation in zero-sum Markov games. In *Uncertainty in Artificial Intelligence*.
- Lasry, J.-M. and Lions, P.-L. (2006a). Jeux à champ moyen. I—le cas stationnaire. *Comptes Rendus Mathématique*, **343** 619–625.
- Lasry, J.-M. and Lions, P.-L. (2006b). Jeux à champ moyen. II—horizon fini et contrôle optimal. *Comptes Rendus Mathématique*, **343** 679–684.
- Lasry, J.-M. and Lions, P.-L. (2007). Mean field games. *Japanese journal of mathematics*, **2** 229–260.
- LeCun, Y., Bengio, Y. and Hinton, G. (2015). Deep learning. *Nature*, **521** 436–444.

- Lee, J., Xiao, L., Schoenholz, S. S., Bahri, Y., Sohl-Dickstein, J. and Pennington, J. (2019). Wide neural networks of any depth evolve as linear models under gradient descent. *arXiv preprint arXiv:1902.06720*.
- Leibo, J. Z., Zambaldi, V., Lanctot, M., Marecki, J. and Graepel, T. (2017). Multi-agent reinforcement learning in sequential social dilemmas. In *International Conference on Autonomous Agents and Multiagent Systems*. International Foundation for Autonomous Agents and Multiagent Systems.
- Levine, S., Kumar, A., Tucker, G. and Fu, J. (2020). Offline reinforcement learning: Tutorial, review, and perspectives on open problems. *arXiv preprint arXiv:2005.01643*.
- Lewis, F. L., Vrabie, D. and Syrmos, V. L. (2012). *Optimal control*. John Wiley & Sons.
- Li, S., Zhang, W. and Zhao, L. (2017). Connections between mean-field game and social welfare optimization. *arXiv preprint arXiv:1703.10211*.
- Li, T. and Zhang, J.-F. (2008). Asymptotically optimal decentralized control for large population stochastic multiagent systems. *IEEE Transactions on Automatic Control*, **53** 1643–1660.
- Li, Y. (2017). Deep reinforcement learning: An overview. *arXiv preprint arXiv:1701.07274*.
- Li, Y. (2018). Deep reinforcement learning. *arXiv preprint arXiv:1810.06339*.
- Li, Y. and Liang, Y. (2018). Learning overparameterized neural networks via stochastic gradient descent on structured data. In *Advances in Neural Information Processing Systems*.
- Liao, L., Fu, Z., Yang, Z., Wang, Y., Kolar, M. and Wang, Z. (2021a). Instrumental variable value iteration for causal offline reinforcement learning. *arXiv preprint*

arXiv:2102.09907.

- Liao, P., Klasnja, P. and Murphy, S. (2021b). Off-policy estimation of long-term average outcomes with applications to mobile health. *Journal of the American Statistical Association*, **116** 382–391.
- Liao, P., Qi, Z., Klasnja, P. and Murphy, S. (2020). Batch policy learning in average reward markov decision processes. *arXiv preprint arXiv:2007.11771*.
- Littman, M. L. (1994). Markov games as a framework for multi-agent reinforcement learning. In *Machine Learning Proceedings 1994*. Elsevier, 157–163.
- Littman, M. L. (2001). Friend-or-foe Q-learning in general-sum games. In *Proceedings of the Eighteenth International Conference on Machine Learning*.
- Liu, B., Cai, Q., Yang, Z. and Wang, Z. (2019). Neural proximal/trust region policy optimization attains globally optimal policy. *arXiv preprint arXiv:1906.10306* 10564–10575.
- Liu, B., Liu, J., Ghavamzadeh, M., Mahadevan, S. and Petrik, M. (2015). Finite-sample analysis of proximal gradient TD algorithms. In *Conference on Uncertainty in Artificial Intelligence*. AUAI Press.
- Liu, Q., Li, L., Tang, Z. and Zhou, D. (2018). Breaking the curse of horizon: Infinite-horizon off-policy estimation. In *Advances in Neural Information Processing Systems*, vol. 31.
- Luenberger, D. G. (1997). *Optimization by vector space methods*. John Wiley & Sons.
- Maei, H. R. (2018). Convergent actor-critic algorithms under off-policy training and function approximation. *arXiv preprint arXiv:1802.07842*.

- Magnus, J. R. (1979). The expectation of products of quadratic forms in normal variables: the practice. *Statistica Neerlandica*, **33** 131–136.
- Magnus, J. R. et al. (1978). *The moments of products of quadratic forms in normal variables*. Univ., Instituut voor Actuarie en Econometrie.
- Malik, D., Pananjady, A., Bhatia, K., Khamaru, K., Bartlett, P. L. and Wainwright, M. J. (2018). Derivative-free methods for policy optimization: Guarantees for linear quadratic systems. *arXiv preprint arXiv:1812.08305*.
- Mei, J., Xiao, C., Szepesvari, C. and Schuurmans, D. (2020). On the global convergence rates of softmax policy gradient methods. *arXiv preprint arXiv:2005.06392*.
- Meitz, M. and Saikkonen, P. (2021). Subgeometric ergodicity and β -mixing. *Journal of Applied Probability*, **58** 594–608.
- Mguni, D., Jennings, J. and de Cote, E. M. (2018). Decentralised learning in systems with many, many strategic agents. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Minciardi, R. and Sacile, R. (2011). Optimal control in a cooperative network of smart power grids. *IEEE Systems Journal*, **6** 126–133.
- Mnih, V., Badia, A. P., Mirza, M., Graves, A., Lillicrap, T., Harley, T., Silver, D. and Kavukcuoglu, K. (2016). Asynchronous methods for deep reinforcement learning. In *International Conference on Machine Learning*.
- Moon, J. and Başar, T. (2014). Discrete-time LQG mean field games with unreliable communication. In *Conference on Decision and Control*. IEEE.
- Moon, J. and Başar, T. (2018). Linear quadratic mean field stackelberg differential games. *Automatica*, **97** 200–213.

- Moravčík, M., Schmid, M., Burch, N., Lisý, V., Morrill, D., Bard, N., Davis, T., Waugh, K., Johanson, M. and Bowling, M. (2017). Deepstack: Expert-level artificial intelligence in heads-up no-limit poker. *Science*, **356** 508–513.
- Munos, R. and Szepesvári, C. (2008). Finite-time bounds for fitted value iteration. *Journal of Machine Learning Research*, **9** 815–857.
- Nachum, O., Chow, Y., Dai, B. and Li, L. (2019). Dualdice: Behavior-agnostic estimation of discounted stationary distribution corrections. In *Advances in Neural Information Processing Systems*, vol. 32.
- Nachum, O. and Dai, B. (2020). Reinforcement learning via Fenchel-Rockafellar duality. *arXiv preprint arXiv:2001.01866*.
- Namkoong, H., Keramati, R., Yadlowsky, S. and Brunskill, E. (2020). Off-policy policy evaluation for sequential decisions under unobserved confounding. *Advances in Neural Information Processing Systems*, **33** 18819–18831.
- Nash, J. (1951). Non-cooperative games. *Annals of mathematics* 286–295.
- OpenAI (2018). Openai five. <https://blog.openai.com/openai-five/>.
- Pearl, J. (2009). *Causality*. Cambridge university press.
- Pérolat, J., Piot, B., Geist, M., Scherrer, B. and Pietquin, O. (2016a). Softened approximate policy iteration for markov games. In *International Conference on Machine Learning*.
- Pérolat, J., Piot, B. and Pietquin, O. (2018). Actor-critic fictitious play in simultaneous move multistage games. In *International Conference on Artificial Intelligence and Statistics*.

- Pérolat, J., Piot, B., Scherrer, B. and Pietquin, O. (2016b). On the use of non-stationary strategies for solving two-player zero-sum Markov games. In *International Conference on Artificial Intelligence and Statistics*.
- Perolat, J., Scherrer, B., Piot, B. and Pietquin, O. (2015). Approximate dynamic programming for two-player zero-sum Markov games. In *International Conference on Machine Learning (ICML 2015)*.
- Peters, J., Janzing, D. and Schölkopf, B. (2017). *Elements of causal inference: foundations and learning algorithms*. The MIT Press.
- Peters, J. and Schaal, S. (2008a). Natural actor-critic. *Neurocomputing*, **71** 1180–1190.
- Peters, J. and Schaal, S. (2008b). Reinforcement learning of motor skills with policy gradients. *Neural Networks*, **21** 682–697.
- Prashanth, L. and Ghavamzadeh, M. (2013). Actor-critic algorithms for risk-sensitive mdps. In *Advances in neural information processing systems*.
- Precup, D. (2000). Eligibility traces for off-policy policy evaluation. *Computer Science Department Faculty Publication Series* 80.
- Puterman, M. L. (2014). *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons.
- Raghu, A., Komorowski, M., Celi, L. A., Szolovits, P. and Ghassemi, M. (2017). Continuous state-space models for optimal sepsis treatment: a deep reinforcement learning approach. In *Machine Learning for Healthcare Conference*. PMLR.
- Rudelson, M., Vershynin, R. et al. (2013). Hanson-wright inequality and sub-Gaussian concentration. *Electronic Communications in Probability*, **18** 82–90.

- Saldi, N., Basar, T. and Raginsky, M. (2018a). Discrete-time risk-sensitive mean-field games. *arXiv preprint arXiv:1808.03929*.
- Saldi, N., Basar, T. and Raginsky, M. (2018b). Markov–Nash equilibria in mean-field games with discounted cost. *SIAM Journal on Control and Optimization*, **56** 4256–4287.
- Saldi, N., Basar, T. and Raginsky, M. (2019). Approximate Nash equilibria in partially observed stochastic games with mean-field interactions. *Mathematics of Operations Research*.
- Sallab, A. E., Abdou, M., Perot, E. and Yogamani, S. (2017). Deep reinforcement learning framework for autonomous driving. *Electronic Imaging*, **2017** 70–76.
- Sandholm, W. H. (2010). *Population Games and Evolutionary Dynamics*. MIT Press.
- Scherrer, B. (2013). On the performance bounds of some policy search dynamic programming algorithms. *arXiv preprint arXiv:1306.0539*.
- Scherrer, B., Ghavamzadeh, M., Gabillon, V., Lesner, B. and Geist, M. (2015). Approximate modified policy iteration and its application to the game of tetris. *Journal of Machine Learning Research*, **16** 1629–1676.
- Schulman, J., Levine, S., Abbeel, P., Jordan, M. and Moritz, P. (2015). Trust region policy optimization. In *International Conference on Machine Learning*.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A. and Klimov, O. (2017). Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Shalev-Shwartz, S., Shammah, S. and Shashua, A. (2016). Safe, multi-agent, reinforcement learning for autonomous driving. *arXiv preprint arXiv:1610.03295*.

- Shi, C., Zhang, S., Lu, W. and Song, R. (2020). Statistical inference of the value function for reinforcement learning in infinite horizon settings. *arXiv preprint arXiv:2001.04515*.
- Shi, C., Zhu, J., Shen, Y., Luo, S., Zhu, H. and Song, R. (2022). Off-policy confidence interval estimation with confounded markov decision process. *arXiv preprint arXiv:2202.10589*.
- Shoham, Y., Powers, R. and Grenager, T. (2003). Multi-agent reinforcement learning: a critical survey. *Technical Report*.
- Shoham, Y., Powers, R. and Grenager, T. (2007). If multi-agent learning is the answer, what is the question? *Artificial Intelligence*, **171** 365–377.
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Van Den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M. et al. (2016). Mastering the game of Go with deep neural networks and tree search. *Nature*, **529** 484–489.
- Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A. et al. (2017). Mastering the game of Go without human knowledge. *Nature*, **550** 354.
- Sutton, R. S. (1988). Learning to predict by the methods of temporal differences. *Machine Learning*, **3** 9–44.
- Sutton, R. S. and Barto, A. G. (2018). *Reinforcement learning: An introduction*. MIT Press.
- Sutton, R. S., Barto, A. G. et al. (1998). *Introduction to Reinforcement Learning*, vol. 135. MIT press.
- Sutton, R. S., Maei, H. R., Precup, D., Bhatnagar, S., Silver, D., Szepesvári, C. and Wiewiora, E. (2009a). Fast gradient-descent methods for temporal-difference learning

- with linear function approximation. In *Proceedings of the 26th Annual International Conference on Machine Learning*. ACM.
- Sutton, R. S., Maei, H. R. and Szepesvári, C. (2009b). A convergent $o(n)$ temporal-difference algorithm for off-policy learning with linear function approximation. In *Advances in neural information processing systems*.
- Sutton, R. S., McAllester, D. A., Singh, S. P. and Mansour, Y. (2000). Policy gradient methods for reinforcement learning with function approximation. In *Advances in neural information processing systems*.
- Szepesvári, C. and Munos, R. (2005). Finite time bounds for sampling based fitted value iteration. In *International Conference on Machine Learning*. ACM.
- Sznitman, A.-S. (1991). Topics in propagation of chaos. In *Ecole d'été de probabilités de Saint-Flour XIX—1989*. Springer, 165–251.
- Tang, Z., Feng, Y., Li, L., Zhou, D. and Liu, Q. (2019). Doubly robust bias reduction in infinite horizon off-policy estimation. *arXiv preprint arXiv:1910.07186*.
- Tembine, H. and Huang, M. (2011). Mean field difference games: McKean-Vlasov dynamics. In *Conference on Decision and Control and European Control Conference*. IEEE.
- Tembine, H., Zhu, Q. and Başar, T. (2014). Risk-sensitive mean-field games. *IEEE Transactions on Automatic Control*, **59** 835–850.
- Thomas, P. and Brunskill, E. (2016). Data-efficient off-policy policy evaluation for reinforcement learning. In *International Conference on Machine Learning*. PMLR.
- Todorov, E., Erez, T. and Tassa, Y. (2012). Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ international conference on intelligent robots and systems*. IEEE.

- Tosatto, S., Pirotta, M., D'Eramo, C. and Restelli, M. (2017). Boosted fitted Q-iteration. In *International Conference on Machine Learning*. JMLR. org.
- Tropp, J. A. (2015). An introduction to matrix concentration inequalities. *arXiv preprint arXiv:1501.01571*.
- Tu, S. and Recht, B. (2017). Least-squares temporal difference learning for the linear quadratic regulator. *arXiv preprint arXiv:1712.08642*.
- Tu, S. and Recht, B. (2018). The gap between model-based and model-free methods on the linear quadratic regulator: An asymptotic viewpoint. *arXiv preprint arXiv:1812.03565* 3036–3083.
- Uehara, M., Huang, J. and Jiang, N. (2020). Minimax weight and q-function learning for off-policy evaluation. In *International Conference on Machine Learning*. PMLR.
- Uehara, M., Imaizumi, M., Jiang, N., Kallus, N., Sun, W. and Xie, T. (2021). Finite sample analysis of minimax offline reinforcement learning: Completeness, fast rates and first-order efficiency. *arXiv preprint arXiv:2102.02981*.
- Uehara, M. and Jiang, N. (2019). Minimax weight and Q-function learning for off-policy evaluation. *arXiv preprint arXiv:1910.12809* 9659–9668.
- Van Roy, B. (1998). *Learning and value function approximation in complex decision processes*. Ph.D. thesis, Massachusetts Institute of Technology.
- Vinyals, O., Babuschkin, I., Chung, J., Mathieu, M., Jaderberg, M., Czarnecki, W., Dudzik, A., Huang, A., Georgiev, P., Powell, R. et al. (2019). Alphastar: Mastering the real-time strategy game starcraft ii. <https://deepmind.com/blog/article/alphastar-mastering-real-time-strategy-game-starcraft-ii/>.

- Wai, H.-T., Yang, Z., Wang, P. Z. and Hong, M. (2018). Multi-agent reinforcement learning via double averaging primal-dual optimization. In *Advances in Neural Information Processing Systems*.
- Wang, B.-C. and Zhang, J.-F. (2012). Mean field games for large-population multiagent systems with Markov jump parameters. *SIAM Journal on Control and Optimization*, **50** 2308–2334.
- Wang, J., Qi, Z. and Wong, R. K. (2021a). Projected state-action balancing weights for offline reinforcement learning. *arXiv preprint arXiv:2109.04640*.
- Wang, J., Zhang, W., Yuan, S. et al. (2017a). Display advertising with real-time bidding (RTB) and behavioural targeting. *Foundations and Trends® in Information Retrieval*, **11** 297–435.
- Wang, L., Cai, Q., Yang, Z. and Wang, Z. (2019). Neural policy gradient methods: Global optimality and rates of convergence. *arXiv preprint arXiv:1909.01150*.
- Wang, L. and Tchetgen Tchetgen, E. (2018). Bounded, efficient and multiply robust estimation of average treatment effects using instrumental variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **80** 531–550.
- Wang, L., Yang, Z. and Wang, Z. (2021b). Provably efficient causal reinforcement learning with confounded observational data. *Advances in Neural Information Processing Systems*, **34**.
- Wang, R., Foster, D. P. and Kakade, S. M. (2020). What are the statistical limits of offline rl with linear function approximation? *arXiv preprint arXiv:2010.11895*.

- Wang, R., Wu, Y., Salakhutdinov, R. and Kakade, S. (2021c). Instabilities of offline rl with pre-trained neural representation. In *International Conference on Machine Learning*. PMLR.
- Wang, Y., Chen, W., Liu, Y., Ma, Z.-M. and Liu, T.-Y. (2017b). Finite sample analysis of the GTD policy evaluation algorithms in Markov setting. In *Advances in Neural Information Processing Systems*.
- Watkins, C. J. and Dayan, P. (1992). Q-learning. *Machine learning*, **8** 279–292.
- Wei, C.-Y., Hong, Y.-T. and Lu, C.-J. (2017). Online reinforcement learning in stochastic games. In *Advances in Neural Information Processing Systems*.
- Wu, L., Ma, C. and Weinan, E. (2018). How SGD selects the global minima in over-parameterized learning: A dynamical stability perspective. In *Advances in Neural Information Processing Systems*.
- Wu, Y., Zhang, W., Xu, P. and Gu, Q. (2020). A finite time analysis of two time-scale actor critic methods. *arXiv preprint arXiv:2005.01350*.
- Xie, T., Cheng, C.-A., Jiang, N., Mineiro, P. and Agarwal, A. (2021). Bellman-consistent pessimism for offline reinforcement learning. *Advances in neural information processing systems*, **34**.
- Xie, T., Ma, Y. and Wang, Y.-X. (2019). Towards optimal off-policy evaluation for reinforcement learning with marginalized importance sampling. In *Advances in Neural Information Processing Systems*.
- Xu, T., Wang, Z. and Liang, Y. (2020). Non-asymptotic convergence analysis of two time-scale (natural) actor-critic algorithms. *arXiv preprint arXiv:2005.03557*.

- Xu, T., Zou, S. and Liang, Y. (2019). Two time-scale off-policy TD learning: Non-asymptotic analysis over Markovian samples. In *Advances in Neural Information Processing Systems*.
- Yang, E. and Gu, D. (2004). Multiagent reinforcement learning for multi-robot systems: A survey. *Manuscript*.
- Yang, J., Ye, X., Trivedi, R., Xu, H. and Zha, H. (2018a). Deep mean field games for learning optimal behavior policy of large populations. In *International Conference on Learning Representations*.
- Yang, Y., Luo, R., Li, M., Zhou, M., Zhang, W. and Wang, J. (2018b). Mean field multi-agent reinforcement learning. *arXiv preprint arXiv:1802.05438*.
- Yang, Z., Chen, Y., Hong, M. and Wang, Z. (2019a). On the global convergence of actor-critic: A case for linear quadratic regulator with ergodic cost. *arXiv preprint arXiv:1907.06246*.
- Yang, Z., Chen, Y., Hong, M. and Wang, Z. (2019b). On the global convergence of actor-critic: A case for linear quadratic regulator with ergodic cost. *arXiv preprint arXiv:1907.06246*.
- Yang, Z., Xie, Y. and Wang, Z. (2019c). A theoretical analysis of deep Q-learning. *arXiv preprint arXiv:1901.00137*.
- Yu, H. (2010). Convergence of least squares temporal difference methods under general conditions. In *International Conference on Machine Learning*.
- Yu, H. (2017). On convergence of some gradient-based temporal-differences algorithms for off-policy learning. *arXiv preprint arXiv:1712.09652*.

- Yu, T., Thomas, G., Yu, L., Ermon, S., Zou, J., Levine, S., Finn, C. and Ma, T. (2020). MOPO: Model-based offline policy optimization. *arXiv preprint arXiv:2005.13239*.
- Zhan, W., Huang, B., Huang, A., Jiang, N. and Lee, J. D. (2022). Offline reinforcement learning with realizability and single-policy concentrability. *arXiv preprint arXiv:2202.04634*.
- Zhang, C. and Liu, Y. (2014). Multicategory angle-based large-margin classification. *Biometrika*, **101** 625–640.
- Zhang, J. and Bareinboim, E. (2019). Near-optimal reinforcement learning in dynamic treatment regimes. *Advances in Neural Information Processing Systems*, **32**.
- Zhang, K., Koppel, A., Zhu, H. and Başar, T. (2019a). Global convergence of policy gradient methods to (almost) locally optimal policies. *arXiv preprint arXiv:1906.08383*.
- Zhang, K., Yang, Z. and Başar, T. (2019b). Multi-agent reinforcement learning: A selective overview of theories and algorithms. *arXiv preprint arXiv:1911.10635*.
- Zhang, K., Yang, Z., Liu, H., Zhang, T. and Başar, T. (2018). Finite-sample analyses for fully decentralized multi-agent reinforcement learning. *arXiv preprint arXiv:1812.02783*.
- Zhang, R., Dai, B., Li, L. and Schuurmans, D. (2020). Gendice: Generalized offline estimation of stationary values. *arXiv preprint arXiv:2002.09072*.
- Zhou, X., Mayer-Hamblett, N., Khan, U. and Kosorok, M. R. (2017). Residual weighted learning for estimating individualized treatment rules. *Journal of the American Statistical Association*, **112** 169–187.
- Zhou, X. Y. and Li, D. (2000). Continuous-time mean-variance portfolio selection: A stochastic LQ framework. *Applied Mathematics and Optimization*, **42** 19–33.

- Ziebart, B. D., Maas, A. L., Bagnell, J. A. and Dey, A. K. (2008). Maximum entropy inverse reinforcement learning. In *AAAI Conference on Artificial Intelligence*, vol. 3.
- Zou, D., Cao, Y., Zhou, D. and Gu, Q. (2018). Stochastic gradient descent optimizes over-parameterized deep ReLU networks. *arXiv preprint arXiv:1811.08888*.
- Zou, S., Xu, T. and Liang, Y. (2019). Finite-sample analysis for SARSA and Q-learning with linear function approximation. *arXiv preprint arXiv:1902.02234* 8665–8675.

APPENDIX A

Supplemental Materials in Chapter 2

A.1. Details of Algorithms

We introduce the actor-critic method with DNN approximation in Algorithm 4, which relies on Algorithms 5 and 6 for the actor and critic updates.

Algorithm 4 Deep Neural Actor-Critic Method

Input: Number of iterations K, N_a, N_c , stepsizes α, η , and temperature parameter β .
Initialization: Set $\tau_0 \leftarrow \infty$ and initialize DNNs f_{θ_0} and Q_{ω_0} as specified in §2.2.2.
for $k = 0, 1, 2, \dots, K$ **do**
 Actor Update: Update θ_{k+1} via Algorithm 5 with input $\pi_{\theta_k}, \theta_0, Q_{\omega_k}, \alpha, \beta, \tau_{k+1} = (k+1)^{-1} \cdot \beta$, and N_a .
 Critic Update: Update ω_{k+1} via Algorithm 6 with input $\pi_{\theta_{k+1}}, Q_{\omega_k}, \omega_0, \eta$, and N_c .
end for
Output: $\{\pi_{\theta_k}\}_{k \in [K+1]}$, where $\pi_{\theta_k} \propto \exp(\tau_k^{-1} f_{\theta_k})$.

Algorithm 5 Actor Update for Deep Neural Actor-Critic Method

Input: Policy $\pi_{\theta} \propto \exp(\tau^{-1} f_{\theta})$, initial actor parameter θ_0 , action-value function Q_{ω} , stepsize α , temperature parameter β , temperature $\tilde{\tau}$, and number of iterations N_a .
Initialization: Set $\theta(0) \leftarrow \theta_0$.
for $n = 0, 1, 2, \dots, N_a - 1$ **do**
 Sample (s, a) as specified in §2.2.2.
 Set $\theta(n+1) \leftarrow \Gamma_{\mathcal{B}(\theta_0, R_a)}(\theta(n) - \alpha \cdot (f_{\theta(n)}(s, a) - \tilde{\tau} \cdot (\beta^{-1} Q_{\omega}(s, a) + \tau^{-1} f_{\theta}(s, a))) \cdot \nabla_{\theta} f_{\theta(n)}(s, a))$.
end for
Output: $\bar{\theta} = 1/N_a \cdot \sum_{n=1}^{N_a} \theta(n)$.

Algorithm 6 Critic Update for Deep Neural Actor-Critic Method

Input: Policy π_θ , action-value function Q_ω , initial critic parameter ω_0 , stepsize η , and number of iterations N_c .

Initialization: Set $\omega(0) \leftarrow \omega_0$.

for $n = 0, 1, 2, \dots, N_c - 1$ **do**

Sample (s, a, r, s', a') as specified in §2.2.2.

Set $\delta(n) \leftarrow Q_{\omega(n)}(s, a) - (1 - \gamma) \cdot r - \gamma \cdot Q_{\omega(n)}(s', a')$.

Set $\omega(n + 1) \leftarrow \Gamma_{\mathcal{B}(\omega_0, R_c)}(\omega(n) - \eta \cdot \delta(n) \cdot \nabla_\omega Q_{\omega(n)}(s, a))$.

end for

Output: $\bar{\omega} = 1/N_c \cdot \sum_{n=1}^{N_c} \omega(n)$.

A.2. Convergence Results of Algorithm 4

In this section, we upper bound the regret of the deep neural actor-critic method. Hereafter we assume that $|\mathcal{R}(s, a)| \leq \mathcal{R}_{\max}$ for any $(s, a) \in \mathcal{S} \times \mathcal{A}$, where \mathcal{R}_{\max} is a positive absolute constant. First, we impose the following assumptions in parallel to Assumption 2.3.1. Recall that ρ^* is the stationary state-action distribution of π^* , while ρ_k is the stationary state-action distribution of π_{θ_k} .

Assumption A.2.1 (Concentrability Coefficient). The following statements hold.

- (i) There exists a positive absolute constant ϕ^* such that $\phi_k^* \leq \phi^*$ for any $k \geq 1$, where $\phi_k^* = \|\mathrm{d}\rho^*/\mathrm{d}\rho_k\|_{\rho_k, 2}$.
- (ii) For the state-action distribution ρ used to define the regret in (2.3.1), we assume that for any $k \geq 1$ and a sequence of policies $\{\pi_i\}_{i \geq 1}$, the k -step future-state-action distribution $\rho \mathbb{P}^{\pi_1} \dots \mathbb{P}^{\pi_k}$ is absolutely continuous with respect to ρ^* . Also,

it holds that

$$C_{\rho, \rho^*} = (1 - \gamma)^2 \sum_{k=1}^{\infty} k^3 \gamma^k \cdot c(k) < \infty,$$

where $c(k) = \sup_{\{\pi_i\}_{i \in [k]}} \|\mathrm{d}(\rho \mathbb{P}^{\pi_1} \dots \mathbb{P}^{\pi_k}) / \mathrm{d}\rho^*\|_{\rho^*, \infty}$.

Meanwhile, we impose the following assumption in parallel to Assumption 2.3.2.

Assumption A.2.2 (Zero Approximation Error). For any $Q_\omega \in \mathcal{U}(m_c, H_c, R_c)$ and policy π , it holds that $\mathcal{T}^\pi Q_\omega \in \mathcal{U}(m_c, H_c, R_c)$, where \mathcal{T}^π is defined in (2.1.4).

Assumption A.2.2 states that $\mathcal{U}(m_c, H_c, R_c)$ is closed under the Bellman evaluation operator \mathcal{T}^π , which is commonly imposed in the literature (Munos and Szepesvári, 2008; Antos et al., 2008a; Farahmand et al., 2010, 2016; Tosatto et al., 2017; Yang et al., 2019c; Liu et al., 2019).

We upper bound the regret of the deep neural actor-critic method in Algorithm 4 in the sequel. To establish such an upper bound, we first establish the rates of convergence of Algorithms 5 and 6 as follows.

Proposition A.2.3. For any sufficiently large $N_a > 0$, let $m_a = \Omega(d^{3/2} R_a^{-1} H_a^{-3/2} \log(m_a^{1/2}/R_a)^{3/2})$, $H_a = O(N_a^{1/4})$, and $R_a = O(m_a^{1/2} H_a^{-6} (\log m_a)^{-3})$. We denote by $\bar{\theta}$ the output of Algorithm 5 with input $\pi_\theta \propto \exp(\tau^{-1} f_\theta)$, θ_0 , Q_ω , α , β , $\tilde{\tau} = (\tau^{-1} + \beta^{-1})^{-1}$, and N_a . Also, let $\tilde{f} = \tilde{\tau} \cdot (\beta^{-1} Q_\omega + \tau^{-1} f_\theta)$. With probability at least $1 - \exp(-\Omega(R_a^{2/3} m_a^{2/3} H_a))$ over the random initialization θ_0 , we have

$$\mathbb{E}[(f_{\bar{\theta}}(s, a) - \tilde{f}(s, a))^2] = O(R_a^2 N_a^{-1/2} + R_a^{8/3} m_a^{-1/6} H_a^7 \log m_a).$$

Here the expectation is taken over the randomness of $\bar{\theta}$ conditioning on the initialization θ_0 and $(s, a) \sim \rho_{\pi_\theta}$, where ρ_{π_θ} is the stationary state-action distribution of π_θ .

Proof. See §A.5.2 for a detailed proof. \square

Proposition A.2.4. For any sufficiently large $N_c > 0$, let $m_c = \Omega(d^{3/2}R_c^{-1}H_c^{-3/2} \log(m_c^{1/2}/R_c)^{3/2})$, $H_c = O(N_c^{1/4})$, and $R_c = O(m_c^{1/2}H_c^{-6}(\log m_c)^{-3})$. We denote by $\bar{\omega}$ the output of Algorithm 6 with input π_θ , Q_ω , ω_0 , η , and N_c . Also, let $\tilde{Q} = (1 - \gamma) \cdot \mathcal{R} + \gamma \cdot \mathbb{P}^{\pi_\theta} Q_\omega$. With probability at least $1 - \exp(-\Omega(R_c^{2/3}m_c^{2/3}H_c))$ over the random initialization ω_0 , we have

$$\mathbb{E}[(Q_{\bar{\omega}}(s, a) - \tilde{Q}(s, a))^2] = O(R_c^2 N_c^{-1/2} + R_c^{8/3} m_c^{-1/6} H_c^7 \log m_c).$$

Here the expectation is taken over the randomness of $\bar{\omega}$ conditioning on the initialization ω_0 and $(s, a) \sim \rho_{\pi_\theta}$, where ρ_{π_θ} is the stationary state-action distribution of π_θ .

Proof. See §A.5.3 for a detailed proof. \square

Propositions A.2.3 and A.2.4 characterize the errors that arise from the actor and critic updates in Algorithm 4, respectively. In particular, if the widths m_a and m_c of the DNNs f_θ and Q_ω are sufficiently large, the errors characterized in Propositions A.2.3 and A.2.4 decay to zero at the rates of $O(N_a^{-1/2})$ and $O(N_c^{-1/2})$, respectively. Propositions A.2.3 and A.2.4 act as the key ingredients to upper bounding the regret of the deep neural actor-critic method.

Based on Propositions A.2.3 and A.2.4, we upper bound the regret of Algorithm 4 in the following theorem, which is in parallel to Theorem 2.3.4.

Theorem A.2.5. We assume that Assumptions A.2.1 and A.2.2 hold. Let ρ be a state-action distribution satisfying (ii) of Assumption A.2.1. Also, for any sufficiently large

$K > 0$, let $N_a = \Omega(K^6 C_{\rho, \rho^*}^4 (\phi^* + \psi^* + 1)^4 R_a^4)$, $N_c = \Omega(K^6 C_{\rho, \rho^*}^4 \phi^{*4} R_c^4)$, $H_a = H_c = O(N_c^{1/4})$, $R_a = R_c = O(m_c^{1/2} H_c^{-6} (\log m_c)^{-3})$, $m_a = m_c = \Omega(d^{3/2} K^6 C_{\rho, \rho^*}^{12} (\phi^* + \psi^* + 1)^{12} R_c^{16} H_c^{42} \log(m_c^{1/2}/R_c)^{3/2})$, $\beta = K^{1/2}$, and the sequence $\{\theta_k\}_{k \in [K]}$ be generated by Algorithm 4. With probability at least $1 - 1/K$ over the random initialization θ_0 and ω_0 , it holds that

$$\mathbb{E} \left[\sum_{k=0}^K Q^*(s, a) - Q^{\pi_{\theta_{k+1}}}(s, a) \right] \leq (2(1 - \gamma)^{-3} \log |\mathcal{A}| + O(1)) \cdot K^{1/2},$$

where the expectation is taken over the randomness of $(s, a) \sim \rho$ and $\{\theta_{k+1}\}_{k \in [K]}$ conditioning on the initialization θ_0 and ω_0 .

Proof. See §A.3.2 for a detailed proof. □

When the architecture of the actor and critic neural networks are properly chosen, Theorem A.2.5 establishes an $O(K^{1/2})$ regret of Algorithm 4, where K is the total number of iterations. Specifically speaking, to establish such a regret upper bound, we need the widths m_a and m_c of the DNNs f_θ and Q_ω to be sufficiently large. Meanwhile, to control the errors of actor update and critic update in Algorithm 4, we also run sufficiently large numbers of iterations in Algorithms 5 and 6.

In terms of the total sample complexity, to simplify our discussion, we omit constant and logarithmic terms here. To obtain an ε -globally optimal policy, it suffices to set $K \asymp \varepsilon^{-2}$ in Algorithm 4. By plugging such a K into $N_a = \Omega(K^6 C_{\rho, \rho^*}^4 (\phi^* + \psi^* + 1)^4 R_a^4)$ and $N_c = \Omega(K^6 C_{\rho, \rho^*}^4 \phi^{*4} R_c^4)$ as required in Theorem A.2.5, we have $N_a = \tilde{O}(\varepsilon^{-12})$ and $N_c = \tilde{O}(\varepsilon^{-12})$. Thus, to achieve an ε -globally optimal policy, the total sample complexity

of Algorithm 4 is $\tilde{O}(\varepsilon^{-14})$. With the modification to off-policy setting as in §2.2.1, the total sample complexity of Algorithm 4 is $\tilde{O}(\varepsilon^{-12})$.

To the best of our knowledge, we establish the rate of convergence and global optimality of the actor-critic method under single-timescale setting with DNN approximation for the first time.

A.3. Proofs of Theorems

A.3.1. Proof of Theorem 2.3.4

Recall that ρ is a state-action distribution satisfying (ii) of Assumption 2.3.1. We first upper bound $\sum_{k=0}^K (Q^*(s, a) - Q^{\pi_{\theta_{k+1}}}(s, a))$ for any $(s, a) \in \mathcal{S} \times \mathcal{A}$ in part 1. Then by further taking the expectation over ρ and invoking Lemma 2.4.1 in part 2, we conclude the proof of Theorem 2.3.4.

Part 1. In the sequel, we upper bound $\sum_{k=0}^K (Q^*(s, a) - Q^{\pi_{\theta_{k+1}}}(s, a))$ for any $(s, a) \in \mathcal{S} \times \mathcal{A}$.

By the definition of Q^* in (2.1.2), it holds for any $(s, a) \in \mathcal{S} \times \mathcal{A}$ that

$$\begin{aligned}
& [Q^* - Q^{\pi_{\theta_{k+1}}}] (s, a) \\
&= \sum_{\ell=0}^{\infty} [(1 - \gamma) \cdot (\gamma \mathbb{P}^{\pi^*})^\ell \mathcal{R}] (s, a) - Q^{\pi_{\theta_{k+1}}}(s, a) \\
&= \sum_{\ell=0}^{\infty} [(1 - \gamma) \cdot (\gamma \mathbb{P}^{\pi^*})^\ell \mathcal{R} + (\gamma \mathbb{P}^{\pi^*})^{\ell+1} Q^{\pi_{\theta_{k+1}}} - (\gamma \mathbb{P}^{\pi^*})^{\ell+1} Q^{\pi_{\theta_{k+1}}}] (s, a) - Q^{\pi_{\theta_{k+1}}}(s, a) \\
&= \sum_{\ell=0}^{\infty} [(1 - \gamma) \cdot (\gamma \mathbb{P}^{\pi^*})^\ell \mathcal{R} + (\gamma \mathbb{P}^{\pi^*})^{\ell+1} Q^{\pi_{\theta_{k+1}}} - (\gamma \mathbb{P}^{\pi^*})^\ell Q^{\pi_{\theta_{k+1}}}] (s, a)
\end{aligned}$$

(A.3.1)

$$= \sum_{\ell=0}^{\infty} [(\gamma \mathbb{P}^{\pi^*})^\ell ((1 - \gamma) \cdot \mathcal{R} + \gamma \cdot \mathbb{P}^{\pi^*} Q^{\pi_{\theta_{k+1}}} - Q^{\pi_{\theta_{k+1}}})] (s, a),$$

where \mathbb{P}^{π^*} is defined in (2.1.3). We upper bound $[(1 - \gamma) \cdot \mathcal{R} + \gamma \cdot \mathbb{P}^{\pi^*} Q^{\pi_{\theta_{k+1}}} - Q^{\pi_{\theta_{k+1}}}] (s, a)$

on the RHS of (A.3.1) in the sequel. By calculation, we have

$$\begin{aligned}
& [(1 - \gamma) \cdot \mathcal{R} + \gamma \cdot \mathbb{P}^{\pi^*} Q^{\pi_{\theta_{k+1}}} - Q^{\pi_{\theta_{k+1}}}] (s, a) \\
&= \left[((1 - \gamma) \cdot \mathcal{R} + \gamma \cdot \mathbb{P}^{\pi^*} Q^{\pi_{\theta_{k+1}}}) - ((1 - \gamma) \cdot \mathcal{R} + \gamma \cdot \mathbb{P}^{\pi^*} Q_{\omega_k}) \right] (s, a) \\
&\quad + \left[((1 - \gamma) \cdot \mathcal{R} + \gamma \cdot \mathbb{P}^{\pi^*} Q_{\omega_k}) - ((1 - \gamma) \cdot \mathcal{R} + \gamma \cdot \mathbb{P}^{\pi_{\theta_{k+1}}} Q_{\omega_k}) \right] (s, a) \\
&\quad + \left[((1 - \gamma) \cdot \mathcal{R} + \gamma \cdot \mathbb{P}^{\pi_{\theta_{k+1}}} Q_{\omega_k}) - Q^{\pi_{\theta_{k+1}}} \right] (s, a) \\
&= A_{1,k}(s, a) + A_{2,k}(s, a) + A_{3,k}(s, a),
\end{aligned}$$

(A.3.2)

where $A_{1,k}$, $A_{2,k}$, and $A_{3,k}$ are defined as follows,

$$\begin{aligned}
A_{1,k}(s, a) &= [\gamma(\mathbb{P}^{\pi^*} - \mathbb{P}^{\pi_{\theta_{k+1}}})Q_{\omega_k}](s, a), \\
A_{2,k}(s, a) &= [\gamma\mathbb{P}^{\pi^*}(Q^{\pi_{\theta_{k+1}}} - Q_{\omega_k})](s, a), \\
(A.3.3) \quad A_{3,k}(s, a) &= [\mathcal{T}^{\pi_{\theta_{k+1}}}Q_{\omega_k} - Q^{\pi_{\theta_{k+1}}}] (s, a).
\end{aligned}$$

Here $\mathcal{T}^{\pi_{\theta_{k+1}}}$ is defined in (2.1.4). By the following three lemmas, we upper bound $A_{1,k}$, $A_{2,k}$, and $A_{3,k}$ on the RHS of (A.3.2), respectively.

Lemma A.3.1. It holds for any $(s, a) \in \mathcal{S} \times \mathcal{A}$ that

$$A_{1,k}(s, a) = [\gamma(\mathbb{P}^{\pi^*} - \mathbb{P}^{\pi_{\theta_{k+1}}})Q_{\omega_k}](s, a) \leq [\gamma\beta \cdot \mathbb{P}(\vartheta_k + \epsilon_{k+1}^a)](s, a),$$

where ϑ_k and ϵ_{k+1}^a are defined as follows,

$$(A.3.4) \quad \vartheta_k(s) = \text{KL}(\pi^*(\cdot | s) \| \pi_{\theta_k}(\cdot | s)) - \text{KL}(\pi^*(\cdot | s) \| \pi_{\theta_{k+1}}(\cdot | s)),$$

$$(A.3.5) \quad \epsilon_{k+1}^a(s) = \langle \log(\pi_{\theta_{k+1}}(\cdot | s)/\pi_{\theta_k}(\cdot | s)) - \beta^{-1} \cdot Q_{\omega_k}(s, \cdot), \pi^*(\cdot | s) - \pi_{\theta_{k+1}}(\cdot | s) \rangle.$$

Proof. See §A.6.2 for a detailed proof. □

We remark that $\epsilon_{k+1}^a = 0$ for any k in the linear actor-critic method. Meanwhile, such a term is included in Lemma A.3.1 only aiming to generalize to the deep neural actor-critic method.

Lemma A.3.2. It holds for any $(s, a) \in \mathcal{S} \times \mathcal{A}$ that

$$A_{2,k}(s, a) \leq [(\gamma\mathbb{P}^{\pi^*})^{k+1}(Q^* - Q_{\omega_0})](s, a) + \gamma\beta \cdot \sum_{i=0}^{k-1} [(\gamma\mathbb{P}^{\pi^*})^{k-i}\mathbb{P}(\vartheta_i + \epsilon_{i+1}^a)](s, a) \\ + \sum_{i=0}^{k-1} [(\gamma\mathbb{P}^{\pi^*})^{k-i}\epsilon_{i+1}^c](s, a),$$

where ϑ_i is defined in (A.3.4) of Lemma A.3.1, ϵ_{i+1}^a is defined in (A.3.5) of Lemma A.3.1, and ϵ_{i+1}^c is defined as follows,

$$(A.3.6) \quad \epsilon_{i+1}^c(s, a) = [\mathcal{T}^{\pi_{\theta_{i+1}}}Q_{\omega_i} - Q_{\omega_{i+1}}](s, a).$$

Proof. See §A.6.3 for a detailed proof. □

We remark that $\epsilon_{k+1}^a = 0$ for any k in the linear actor-critic method. Meanwhile, such a term is included in Lemma A.3.2 only aiming to generalize to the deep neural actor-critic method.

Lemma A.3.3. It holds for any $(s, a) \in \mathcal{S} \times \mathcal{A}$ that

$$A_{3,k}(s, a) = [\gamma\mathbb{P}^{\pi_{\theta_{k+1}}}(I - \gamma\mathbb{P}^{\pi_{\theta_{k+1}}})^{-1}e_{k+1}](s, a),$$

where e_{k+1} is defined as follows,

$$(A.3.7) \quad e_{k+1}(s, a) = [Q_{\omega_k} - \mathcal{T}^{\pi_{\theta_{k+1}}}Q_{\omega_k}](s, a).$$

Proof. See §A.6.4 for a detailed proof. □

We upper bound e_{k+1} in (A.3.7) of Lemma A.3.3 using Lemma A.3.4 as follows.

Lemma A.3.4. It holds for any $(s, a) \in \mathcal{S} \times \mathcal{A}$ that

$$e_{k+1}(s, a) \leq \left[\gamma^k \left(\prod_{s=1}^k \mathbb{P}^{\pi_{\theta_s}} \right) e_1 + \sum_{i=1}^k \gamma^{k-i} \left(\prod_{s=i+1}^k \mathbb{P}^{\pi_{\theta_s}} \right) (\gamma \beta \mathbb{P} \epsilon_{i+1}^b + (I - \gamma \mathbb{P}^{\pi_{\theta_i}}) \epsilon_i^c) \right] (s, a).$$

where $\epsilon_i^c(s, a)$ is defined in (A.3.6) of Lemma A.3.2 and $\epsilon_{i+1}^b(s)$ is defined as follows,

$$(A.3.8) \quad \epsilon_{i+1}^b(s) = \langle \log(\pi_{\theta_{i+1}}(\cdot | s) / \pi_{\theta_i}(\cdot | s)) - \beta^{-1} \cdot Q_{\omega_i}(s, \cdot), \pi_{\theta_i}(\cdot | s) - \pi_{\theta_{i+1}}(\cdot | s) \rangle.$$

Proof. See §A.6.5 for a detailed proof. □

We remark that $\epsilon_{i+1}^b = 0$ for any i in the linear actor-critic method. Meanwhile, such a term is included in Lemma A.3.4 only aiming to generalize to the deep neural actor-critic method.

Combining Lemmas A.3.3 and A.3.4, we obtain the following upper bound of $A_{3,k}$,

$$(A.3.9) \quad \begin{aligned} A_{3,k}(s, a) &= [\gamma \mathbb{P}^{\pi_{\theta_{k+1}}} (I - \gamma \mathbb{P}^{\pi_{\theta_{k+1}}})^{-1} e_{k+1}] (s, a) \\ &\leq \left[\gamma \mathbb{P}^{\pi_{\theta_{k+1}}} (I - \gamma \mathbb{P}^{\pi_{\theta_{k+1}}})^{-1} \left(\gamma^k \left(\prod_{s=1}^k \mathbb{P}^{\pi_{\theta_s}} \right) e_1 \right. \right. \\ &\quad \left. \left. + \sum_{i=1}^k \gamma^{k-i} \left(\prod_{s=i+1}^k \mathbb{P}^{\pi_{\theta_s}} \right) (\beta \gamma \mathbb{P} \epsilon_{i+1}^b + (I - \gamma \mathbb{P}^{\pi_{\theta_i}}) \epsilon_i^c) \right) \right] (s, a). \end{aligned}$$

Combining (A.3.1), (A.3.2), Lemma A.3.1 and Lemma A.3.2, it holds for any $(s, a) \in \mathcal{S} \times \mathcal{A}$ that

$$\begin{aligned}
& \sum_{k=0}^K [Q^* - Q^{\pi^{\theta_{k+1}}}] (s, a) \\
& \leq \sum_{k=0}^K \left[(I - \gamma \mathbb{P}^{\pi^*})^{-1} ((\gamma \mathbb{P}^{\pi^*})^{k+1} (Q^* - Q_{\omega_0}) + \sum_{i=0}^k (\gamma \mathbb{P}^{\pi^*})^{k-i} \gamma \beta \mathbb{P}(\vartheta_i + \epsilon_{i+1}^a) \right. \\
& \quad \left. + \sum_{i=0}^{k-1} (\gamma \mathbb{P}^{\pi^*})^{k-i} \epsilon_{i+1}^c + A_{3,k} \right] (s, a) \\
\text{(A.3.10)} & \\
& = \left[(I - \gamma \mathbb{P}^{\pi^*})^{-1} \left(\sum_{k=0}^K (\gamma \mathbb{P}^{\pi^*})^{k+1} (Q^* - Q_{\omega_0}) + \sum_{k=0}^K \sum_{i=0}^k (\gamma \mathbb{P}^{\pi^*})^{k-i} \gamma \beta \mathbb{P} \epsilon_{i+1}^a \right. \right. \\
& \quad \left. \left. + \sum_{k=0}^K \sum_{i=0}^{k-1} (\gamma \mathbb{P}^{\pi^*})^{k-i} \epsilon_{i+1}^c + \sum_{k=0}^K A_{3,k} + \sum_{k=0}^K \sum_{i=0}^k (\gamma \mathbb{P}^{\pi^*})^{k-i} \gamma \beta \mathbb{P} \vartheta_i \right) \right] (s, a),
\end{aligned}$$

where ϑ_i , ϵ_{i+1}^a , ϵ_{i+1}^c , and e_{k+1} are defined in (A.3.4) of Lemma A.3.1, (A.3.5) of Lemma A.3.1, (A.3.6) of Lemma A.3.2, and (A.3.7) of Lemma A.3.3, respectively. We upper

bound the last term as follows,

$$\begin{aligned}
& \left[\sum_{k=0}^K \sum_{i=0}^k (\gamma \mathbb{P}^{\pi^*})^{k-i} \gamma \beta \mathbb{P} \vartheta_i \right] (s, a) = \left[\sum_{k=0}^K \sum_{i=0}^k \gamma \beta (\gamma \mathbb{P}^{\pi^*})^i \mathbb{P} \vartheta_{k-i} \right] (s, a) \\
& = \left[\sum_{i=0}^K \gamma \beta (\gamma \mathbb{P}^{\pi^*})^i \mathbb{P} \sum_{k=i}^K \vartheta_{k-i} \right] (s, a) \\
& = \left[\sum_{i=0}^K \gamma \beta (\gamma \mathbb{P}^{\pi^*})^i \mathbb{P} \sum_{k=i}^K \left(\text{KL}(\pi^* \parallel \pi_{\theta_{k-i}}) - \text{KL}(\pi^* \parallel \pi_{\theta_{k-i+1}}) \right) \right] (s, a) \\
& = \left[\sum_{i=0}^K \gamma \beta (\gamma \mathbb{P}^{\pi^*})^i \mathbb{P} \left(\text{KL}(\pi^* \parallel \pi_{\theta_0}) - \text{KL}(\pi^* \parallel \pi_{\theta_{K-i+1}}) \right) \right] (s, a) \\
\text{(A.3.11)} \quad & \leq \left[\sum_{i=0}^K \gamma \beta (\gamma \mathbb{P}^{\pi^*})^i \mathbb{P} \text{KL}(\pi^* \parallel \pi_{\theta_0}) \right] (s, a),
\end{aligned}$$

where we use the definition of ϑ_{k-i} in (A.3.4) of Lemma A.3.1 and the non-negativity of the KL divergence in the second equality and the last inequality, respectively. By plugging

(A.3.9) and (A.3.11) into (A.3.10), we have

$$\begin{aligned}
& \sum_{k=0}^K [Q^* - Q^{\pi_{\theta_{k+1}}}] (s, a) \\
& \text{(A.3.12)} \\
& \leq \left[(I - \gamma \mathbb{P}^{\pi^*})^{-1} \left(\sum_{k=0}^K (\gamma \mathbb{P}^{\pi^*})^{k+1} (Q^* - Q_{\omega_0}) + \sum_{k=0}^K \sum_{i=0}^k (\gamma \mathbb{P}^{\pi^*})^{k-i} \gamma \beta \mathbb{P} \epsilon_{i+1}^a \right. \right. \\
& \quad + \sum_{k=0}^K \sum_{i=0}^{k-1} (\gamma \mathbb{P}^{\pi^*})^{k-i} \epsilon_{i+1}^c + \sum_{k=0}^K \gamma^{k+1} \mathbb{P}^{\pi_{\theta_{k+1}}} (I - \gamma \mathbb{P}^{\pi_{\theta_{k+1}}})^{-1} \left(\prod_{s=1}^k \mathbb{P}^{\pi_{\theta_s}} \right) e_1 \\
& \quad \left. \left. + \sum_{k=0}^K \mathbb{P}^{\pi_{\theta_{k+1}}} (I - \gamma \mathbb{P}^{\pi_{\theta_{k+1}}})^{-1} \sum_{\ell=1}^k \gamma^{k-\ell+1} \left(\prod_{s=\ell+1}^k \mathbb{P}^{\pi_{\theta_s}} \right) (\gamma \beta \mathbb{P} \epsilon_{\ell+1}^b + (I - \gamma \mathbb{P}^{\pi_{\theta_\ell}}) \epsilon_\ell^c) \right) \right] (s, a). \\
& \quad + \sum_{i=0}^K (\gamma \mathbb{P}^{\pi^*})^i \gamma \beta \mathbb{P} \text{KL}(\pi^* \parallel \pi_{\theta_0})
\end{aligned}$$

We remark that $\epsilon_{i+1}^a = \epsilon_{i+1}^b = 0$ for any i in the linear actor-critic method. Meanwhile, such terms is included in (A.3.12) only aiming to generalize to the deep neural actor-critic method. This concludes the proof in part 1.

Part 2. Recall that ρ is a state-action distribution satisfying (ii) of Assumption 2.3.1. In the sequel, we take the expectation over ρ in (A.3.12) and upper bound each term. Recall that $\epsilon_{i+1}^a = \epsilon_{i+1}^b = 0$ for any i in the linear actor-critic method. Hence, we only need to consider terms in (A.3.12) that do not involve ϵ_{i+1}^a or ϵ_{i+1}^b . We first upper bound terms on the RHS of (A.3.12) that do not involve ϵ_{i+1}^c . More specifically, for any measure ρ

satisfying (ii) of Assumption 2.3.1, we upper bound the following three terms,

$$\begin{aligned}
M_1 &= \mathbb{E}_\rho \left[(I - \gamma \mathbb{P}^{\pi^*})^{-1} \sum_{k=0}^K (\gamma \mathbb{P}^{\pi^*})^{k+1} (Q^* - Q_{\omega_0}) \right], \\
M_2 &= \mathbb{E}_\rho \left[(I - \gamma \mathbb{P}^{\pi^*})^{-1} \sum_{k=0}^K \gamma^{k+1} \mathbb{P}^{\pi_{\theta_{k+1}}} (I - \gamma \mathbb{P}^{\pi_{\theta_{k+1}}})^{-1} \left(\prod_{s=1}^k \mathbb{P}^{\pi_{\theta_s}} \right) e_1 \right], \\
(A.3.13) \quad M_3 &= \mathbb{E}_\rho \left[(I - \gamma \mathbb{P}^{\pi^*})^{-1} \sum_{i=0}^K (\gamma \mathbb{P}^{\pi^*})^i \gamma \beta \text{PKL}(\pi^* \parallel \pi_{\theta_0}) \right].
\end{aligned}$$

We upper bound M_1 , M_2 , and M_3 in the following lemma.

Lemma A.3.5. It holds that

$$\begin{aligned}
|M_1| &\leq 4(1 - \gamma)^{-2} \cdot (\mathcal{R}_{\max} + R), & |M_2| &\leq (1 - \gamma)^{-3} \cdot (2R + \mathcal{R}_{\max}), \\
|M_3| &\leq (1 - \gamma)^{-2} \cdot \log |\mathcal{A}| \cdot K^{1/2},
\end{aligned}$$

where M_1 , M_2 , and M_3 are defined in (A.3.13).

Proof. See §A.6.6 for a detailed proof. □

Now, we upper bound terms on the RHS of (A.3.12) that involve ϵ_{i+1}^c . More specifically, for any measure ρ satisfying (ii) of Assumption 2.3.1, we upper bound the following two terms,

$$\begin{aligned}
(A.3.14) \quad M_4 &= \mathbb{E}_\rho \left[(I - \gamma \mathbb{P}^{\pi^*})^{-1} \sum_{k=0}^K \sum_{i=0}^k (\gamma \mathbb{P}^{\pi^*})^{k-i} \epsilon_{i+1}^c \right], \\
M_5 &= \mathbb{E}_\rho \left[(I - \gamma \mathbb{P}^{\pi^*})^{-1} \sum_{k=0}^K \mathbb{P}^{\pi_{\theta_{k+1}}} (I - \gamma \mathbb{P}^{\pi_{\theta_{k+1}}})^{-1} \sum_{\ell=1}^k \gamma^{k-\ell+1} \left(\prod_{s=\ell+1}^k \mathbb{P}^{\pi_{\theta_s}} \right) (I - \gamma \mathbb{P}^{\pi_{\theta_\ell}}) \epsilon_\ell^c \right].
\end{aligned}$$

We upper bound M_4 and M_5 in the following lemma.

Lemma A.3.6. It holds that

$$|M_4| \leq 3KC_{\rho, \rho^*} \cdot \varepsilon_Q, \quad |M_5| \leq KC_{\rho, \rho^*} \cdot \varepsilon_Q.$$

where M_4 and M_5 are defined in (A.3.14).

Proof. See §A.6.7 for a detailed proof. □

Now, by plugging Lemmas A.3.5 and A.3.6 into (A.3.12), we have

$$(A.3.15) \quad \mathbb{E}_\rho \left[\sum_{k=0}^K Q^*(s, a) - Q^{\pi_{\theta_{k+1}}}(s, a) \right] \leq 2(1 - \gamma)^{-3} \cdot \log |\mathcal{A}| \cdot K^{1/2} + 4KC_{\rho, \rho^*} \cdot \varepsilon_Q + O(1).$$

Meanwhile, by changing measure from ρ^* to ρ_{k+1} , it holds for any k that

$$(A.3.16) \quad \mathbb{E}_{\rho^*} [|\epsilon_{k+1}^c|] \leq \sqrt{\mathbb{E}_{\rho_{k+1}} [(\epsilon_{k+1}^c(s, a))^2]} \cdot \phi_{k+1}^*,$$

where ϕ_{k+1}^* is defined in Assumption 2.3.1. Also, by Lemma 2.4.1, with probability at least $1 - \delta$, it holds for any $k \in \{0, 1, \dots, K\}$ that

$$(A.3.17) \quad \sqrt{\mathbb{E}_{\rho_{k+1}} [(\epsilon_{k+1}^c(s, a))^2]} = O(1/(\sqrt{N}\sigma^*) \cdot \log(KN/\delta)).$$

Now, by plugging (A.3.17) into (A.3.16), combining the definition of $\varepsilon_Q = \max_k \mathbb{E}_{\rho^*} [|\epsilon_{k+1}^c(s, a)|]$, it holds with probability at least $1 - \delta$ that

$$(A.3.18) \quad \varepsilon_Q = O(\phi^*/(\sqrt{N}\sigma^*) \cdot \log(KN/\delta)).$$

Combining (A.3.15), (A.3.18), and the choices of parameters stated in the theorem that

$$N = \Omega(KC_{\rho,\rho^*}^2(\phi^*/\sigma^*)^2 \cdot \log^2(KN/\delta)),$$

we have

$$\mathbb{E}_\rho \left[\sum_{k=0}^K Q^*(s, a) - Q^{\pi_{\theta_{k+1}}}(s, a) \right] \leq (2(1-\gamma)^{-3} \log |\mathcal{A}| + O(1)) \cdot K^{1/2},$$

which concludes the proof of Theorem 2.3.4.

A.3.2. Proof of Theorem A.2.5

We follow the proof of Theorem 2.3.4 in §A.3.1. Following similar arguments when deriving (A.3.12) in §A.3.1, we have

$$\begin{aligned} & \sum_{k=0}^K [Q^* - Q^{\pi_{\theta_{k+1}}}] (s, a) \\ (A.3.19) \quad & \leq \left[(I - \gamma \mathbb{P}^{\pi^*})^{-1} \cdot \left(\sum_{k=0}^K (\gamma \mathbb{P}^{\pi^*})^{k+1} (Q^* - Q_{\omega_0}) + \sum_{k=0}^K \sum_{i=0}^k (\gamma \mathbb{P}^{\pi^*})^{k-i} \cdot \gamma \beta \mathbb{P} \epsilon_{i+1}^a \right. \right. \\ & \quad \left. \left. + \sum_{k=0}^K \sum_{i=0}^{k-1} (\gamma \mathbb{P}^{\pi^*})^{k-i} \epsilon_{i+1}^c + \sum_{i=0}^K (\gamma \mathbb{P}^{\pi^*})^i \cdot \gamma \beta \mathbb{P} \cdot \text{KL}(\pi^* \parallel \pi_{\theta_0}) \right) \right. \\ & \quad \left. + \sum_{k=0}^K \gamma^{k+1} \mathbb{P}^{\pi_{\theta_{k+1}}} (I - \gamma \mathbb{P}^{\pi_{\theta_{k+1}}})^{-1} \left(\prod_{s=1}^k \mathbb{P}^{\pi_{\theta_s}} \right) e_1 \right. \\ & \quad \left. + \sum_{k=0}^K \mathbb{P}^{\pi_{\theta_{k+1}}} (I - \gamma \mathbb{P}^{\pi_{\theta_{k+1}}})^{-1} \sum_{\ell=1}^k \gamma^{k-\ell+1} \left(\prod_{s=\ell+1}^k \mathbb{P}^{\pi_{\theta_s}} \right) (\beta \gamma \mathbb{P} \epsilon_{\ell+1}^b - (I - \gamma \mathbb{P}^{\pi_{\theta_\ell}}) \epsilon_\ell^c) \right] (s, a), \end{aligned}$$

for any $(s, a) \in \mathcal{S} \times \mathcal{A}$. Here ϵ_{i+1}^a , $\epsilon_{\ell+1}^b$, ϵ_{i+1}^c , and e_1 are defined in (A.3.5), (A.3.8), (A.3.6), and (A.3.7), respectively.

Now, it remains to upper bound each term on the RHS of (A.3.19). We introduce the following error propagation lemma.

Lemma A.3.7. Suppose that

$$(A.3.20) \quad \mathbb{E}_{\rho_k} \left[\left(f_{\theta_{k+1}}(s, a) - \tau_{k+1} \cdot (\beta^{-1} Q_{\omega_k}(s, a) - \tau_k^{-1} f_{\theta_k}(s, a)) \right)^2 \right]^{1/2} \leq \varepsilon_{k+1, f}.$$

Then, we have

$$\mathbb{E}_{\nu^*} [|\epsilon_{k+1}^a(s)|] \leq \sqrt{2} \tau_{k+1}^{-1} \cdot \varepsilon_{k+1, f} \cdot (\phi_k^* + \psi_k^*), \quad \mathbb{E}_{\nu^*} [|\epsilon_{k+1}^b(s)|] \leq \sqrt{2} \tau_{k+1}^{-1} \cdot \varepsilon_{k+1, f} \cdot (1 + \psi_k^*),$$

where ϵ_{k+1}^a and ϵ_{k+1}^b are defined in (A.3.5) and (A.3.8), respectively, ϕ_k^* and ψ_k^* are defined in Assumption A.2.1.

Proof. See §A.6.8 for a detailed proof. □

Following from Lemma A.4.4, with probability at least $1 - O(H_c) \exp(-\Omega(H_c^{-1} m_c))$, we have $|Q_{\omega_0}| \leq 2$. Also, from the fact that $|\mathcal{R}(s, a)| \leq \mathcal{R}_{\max}$, we know that $|Q^*| \leq \mathcal{R}_{\max}$. Therefore, for any measure ρ , we have

$$(A.3.21) \quad \begin{aligned} & \left| \mathbb{E}_{\rho} \left[(I - \gamma \mathbb{P}^{\pi^*})^{-1} \sum_{k=0}^K (\gamma \mathbb{P}^{\pi^*})^{k+1} (Q^* - Q_{\omega_0}) \right] \right| \\ & \leq \mathbb{E}_{\rho} \left[(I - \gamma \mathbb{P}^{\pi^*})^{-1} \sum_{k=0}^K (\gamma \mathbb{P}^{\pi^*})^{k+1} |Q^* - Q_{\omega_0}| \right] \\ & \leq \mathcal{R}_{\max} (1 - \gamma)^{-1} \sum_{k=0}^K \gamma^{k+1} \leq \mathcal{R}_{\max} (1 - \gamma)^{-2}. \end{aligned}$$

Also, by changing the index of summation, we have

$$\begin{aligned}
& \left| \mathbb{E}_\rho \left[(I - \gamma \mathbb{P}^{\pi^*})^{-1} \sum_{k=0}^K \sum_{i=0}^k (\gamma \mathbb{P}^{\pi^*})^{k-i} \gamma \beta \mathbb{P} \epsilon_{i+1}^a \right] \right| \\
&= \left| \mathbb{E}_\rho \left[\sum_{k=0}^K \sum_{i=0}^k \sum_{j=0}^{\infty} (\gamma \mathbb{P}^{\pi^*})^{k-i+j} \gamma \beta \mathbb{P} \epsilon_{i+1}^a \right] \right| \\
&= \left| \mathbb{E}_\rho \left[\sum_{k=0}^K \sum_{i=0}^k \sum_{t=k-i}^{\infty} (\gamma \mathbb{P}^{\pi^*})^t \gamma \beta \mathbb{P} \epsilon_{i+1}^a \right] \right| \\
\text{(A.3.22)} \quad & \leq \sum_{k=0}^K \sum_{i=0}^k \sum_{t=k-i}^{\infty} \left| \mathbb{E}_\rho [(\gamma \mathbb{P}^{\pi^*})^t \gamma \beta \mathbb{P} \epsilon_{i+1}^a] \right|,
\end{aligned}$$

where we expand $(I - \gamma \mathbb{P}^{\pi^*})^{-1}$ into an infinite sum in the first equality. Further, by changing the measure of the expectation on the RHS of (A.3.22), we have

$$\text{(A.3.23)} \quad \sum_{k=0}^K \sum_{i=0}^k \sum_{t=k-i}^{\infty} \left| \mathbb{E}_\rho [(\gamma \mathbb{P}^{\pi^*})^t \gamma \beta \mathbb{P} \epsilon_{i+1}^a] \right| \leq \sum_{k=0}^K \sum_{i=0}^k \sum_{t=k-i}^{\infty} \beta \gamma^{t+1} c(t) \cdot \mathbb{E}_{\nu^*} [|\epsilon_{i+1}^A|],$$

where $c(t)$ is defined in Assumption A.2.1. Further, by Lemma A.3.7 and interchanging the summation on the RHS of (A.3.23), we have

$$\begin{aligned}
& \left| \mathbb{E}_\rho \left[(I - \gamma \mathbb{P}^{\pi^*})^{-1} \sum_{k=0}^K \sum_{i=0}^k (\gamma \mathbb{P}^{\pi^*})^{k-i} \gamma \beta \mathbb{P} \epsilon_{i+1}^a \right] \right| \\
& \leq 2 \sum_{k=0}^K \sum_{t=0}^{\infty} \sum_{i=\max\{0, k-t\}}^k \beta \gamma^{t+1} c(t) \cdot \tau_{i+1}^{-1} \varepsilon_f(\phi_i^* + \psi_i^*) \\
& \leq \sum_{k=0}^K \sum_{t=0}^{\infty} 4kt \gamma^{t+1} c(t) \cdot \varepsilon_f(\phi^* + \psi^*) \\
\text{(A.3.24)} \quad & \leq \gamma \sum_{k=0}^K 4C_{\rho, \rho^*} \cdot \varepsilon_f(\phi^* + \psi^*) \leq 2\gamma K^2 C_{\rho, \rho^*}(\phi^* + \psi^*) \cdot \varepsilon_f,
\end{aligned}$$

where $\varepsilon_f = \max_i \mathbb{E}_{\rho_i} [(f_{\theta_{i+1}}(s, a) - \tau_{i+1} \cdot (\beta^{-1} Q_{\omega_i}(s, a) - \tau_i^{-1} f_{\theta_i}(s, a)))^2]^{1/2}$, and C_{ρ, ρ^*} is defined in Assumption A.2.1. Here in the second inequality, we use the fact that $\tau_{i+1}^{-1} = (i+1) \cdot \beta^{-1}$, and $\phi_i^* \leq \phi^*$ and $\psi_i^* \leq \psi^*$ by Assumption A.2.1.

By similar arguments in the derivation of (A.3.24), we have

(A.3.25)

$$\begin{aligned} & \left| \mathbb{E}_{\rho} \left[(I - \gamma \mathbb{P}^{\pi^*})^{-1} \sum_{k=0}^K \sum_{i=0}^{k-1} (\gamma \mathbb{P}^{\pi^*})^{k-i} \epsilon_{i+1}^c \right] \right| \leq 2(K+1) C_{\rho, \rho^*} \phi^* \cdot \varepsilon_Q, \\ & \left| \mathbb{E}_{\rho} \left[(I - \gamma \mathbb{P}^{\pi^*})^{-1} \sum_{i=0}^K (\gamma \mathbb{P}^{\pi^*})^i \gamma \beta \mathbb{P} \text{KL}(\pi^* \parallel \pi_{\theta_0}) \right] \right| \leq \log |\mathcal{A}| \cdot K^{1/2} (1 - \gamma)^{-2}, \\ & \mathbb{E}_{\rho} \left[(I - \gamma \mathbb{P}^{\pi^*})^{-1} \sum_{k=0}^K \gamma^{k+1} \mathbb{P}^{\pi_{\theta_{k+1}}} (I - \gamma \mathbb{P}^{\pi_{\theta_{k+1}}})^{-1} \left(\prod_{s=1}^k \mathbb{P}^{\pi_{\theta_s}} \right) e_1 \right] \leq (2 + \mathcal{R}_{\max}) \cdot (1 - \gamma)^{-3}, \end{aligned}$$

where $\varepsilon_Q = \max_i \mathbb{E}_{\rho^*} [|\epsilon_{i+1}^c|]$. And we use the fact that $\beta = K^{1/2}$.

Now, it remains to upper bound the last term on the RHS of (A.3.19). We first consider the terms involving $\epsilon_{\ell+1}^b$. We have

$$\begin{aligned} & \mathbb{E}_{\rho} \left[(I - \gamma \mathbb{P}^{\pi^*})^{-1} \sum_{k=0}^K \mathbb{P}^{\pi_{\theta_{k+1}}} (I - \gamma \mathbb{P}^{\pi_{\theta_{k+1}}})^{-1} \sum_{\ell=1}^k \gamma^{k-\ell+1} \left(\prod_{s=\ell+1}^k \mathbb{P}^{\pi_{\theta_s}} \right) \beta \gamma \mathbb{P} \epsilon_{\ell+1}^b \right] \\ &= \sum_{j=0}^{\infty} \sum_{i=0}^{\infty} \sum_{k=0}^K \sum_{\ell=1}^k \mathbb{E}_{\rho} \left[(\gamma \mathbb{P}^{\pi^*})^j (\gamma \mathbb{P}^{\pi_{\theta_{k+1}}})^{i+1} \gamma^{k-\ell} \left(\prod_{s=\ell+1}^k \mathbb{P}^{\pi_{\theta_s}} \right) \beta \gamma \mathbb{P} \epsilon_{\ell+1}^b \right] \\ &\leq \beta \gamma \sum_{k=0}^K \sum_{\ell=1}^k \sum_{j=0}^{\infty} \sum_{i=0}^{\infty} \gamma^{i+j+k-\ell+1} \cdot \mathbb{E}_{\rho^*} [|\mathbb{P} \epsilon_{\ell+1}^b|] \cdot c(i+j+k-\ell+1) \\ \text{(A.3.26)} \quad &\leq 2\gamma \sum_{k=0}^K \sum_{\ell=1}^k \sum_{j=0}^{\infty} \sum_{i=0}^{\infty} \gamma^{i+j+k-\ell+1} \cdot (\ell+1) \varepsilon_f \cdot (1 + \psi_{\ell}^*) \cdot c(i+j+k-\ell+1), \end{aligned}$$

where we expand $(I - \gamma \mathbb{P}^{\pi^*})^{-1}$ and $(I - \gamma \mathbb{P}^{\pi^{\theta_{k+1}}})^{-1}$ to infinite sums in the first equality, change the measure of the expectation in the first inequality, and use Lemma A.3.7 in the last inequality. Now, by changing the index of the summation, we have

$$\begin{aligned}
& \gamma \sum_{k=0}^K \sum_{\ell=1}^k \sum_{j=0}^{\infty} \sum_{i=0}^{\infty} \gamma^{i+j+k-\ell+1} \cdot (\ell+1) \varepsilon_f \cdot (1 + \psi_{\ell}^*) \cdot c(i+j+k-\ell+1) \\
&= \gamma \sum_{k=0}^K \sum_{\ell=1}^k \sum_{j=0}^{\infty} \sum_{t=j+k-\ell+1}^{\infty} \gamma^t \cdot (\ell+1) \varepsilon_f \cdot (1 + \psi_{\ell}^*) \cdot c(t) \\
\text{(A.3.27)} \quad & \leq \gamma \sum_{k=0}^K \sum_{j=0}^{\infty} \sum_{t=j+1}^{\infty} \sum_{\ell=\max\{0, j+k-t+1\}}^k \gamma^t \cdot (\ell+1) \varepsilon_f \cdot (1 + \psi^*) \cdot c(t),
\end{aligned}$$

where we use the fact that $\psi_{\ell}^* \leq \psi^*$ from Assumption A.2.1 in the last inequality. By further manipulating the order of summations of the RHS of (A.3.27), we have

$$\begin{aligned}
& \gamma \sum_{k=0}^K \sum_{j=0}^{\infty} \sum_{t=j+1}^{\infty} \sum_{\ell=\max\{0, j+k-t+1\}}^k \gamma^t \cdot (\ell+1) \varepsilon_f (1 + \psi^*) \cdot c(t) \\
& \leq \gamma \sum_{k=0}^K \sum_{j=0}^{\infty} \left(\sum_{t=j+1}^{j+k+1} (t-j)(2k+j-k+1) \cdot \gamma^t c(t) + \sum_{t=j+k+2}^{\infty} k^2 \cdot \gamma^t c(t) \right) \cdot \varepsilon_f (1 + \psi^*) \\
& = \gamma \sum_{k=0}^K \left(\sum_{t=1}^{\infty} \sum_{j=\max\{0, t-k-1\}}^{t-1} (t-j)(2k+j-k+1) \cdot \gamma^t c(t) \right. \\
& \quad \left. + \sum_{t=k+2}^{\infty} \sum_{j=1}^{t-k-2} k^2 \cdot \gamma^t c(t) \right) \cdot \varepsilon_f (1 + \psi^*) \\
& \leq 20\gamma \sum_{k=0}^K \left(\sum_{t=1}^{\infty} k^2 \cdot t \gamma^t c(t) + \sum_{t=1}^{\infty} k^2 \cdot t \gamma^t c(t) \right) \cdot \varepsilon_f (1 + \psi^*) \\
\text{(A.3.28)} \quad & \leq 20\gamma K \cdot C_{\rho, \rho^*} \cdot \varepsilon_f (1 + \psi^*),
\end{aligned}$$

where we use the definition of C_{ρ, ρ^*} from Assumption A.2.1 in the last inequality. Now, combining (A.3.26), (A.3.27), and (A.3.28), we have

$$(A.3.29) \quad \mathbb{E}_\rho \left[(I - \gamma \mathbb{P}^{\pi^*})^{-1} \sum_{k=0}^K \mathbb{P}^{\pi_{\theta_{k+1}}} (I - \gamma \mathbb{P}^{\pi_{\theta_{k+1}}})^{-1} \sum_{\ell=1}^k \gamma^{k-\ell+1} \left(\prod_{s=\ell+1}^k \mathbb{P}^{\pi_{\theta_s}} \right) \beta \gamma \mathbb{P} \epsilon_{\ell+1}^b \right] \leq 20\gamma K \cdot C_{\rho, \rho^*} \cdot \varepsilon_f \cdot (1 + \psi^*).$$

Following from similar arguments when deriving (A.3.29), we have

$$(A.3.30) \quad \mathbb{E}_\rho \left[(I - \gamma \mathbb{P}^{\pi^*})^{-1} \sum_{k=0}^K \mathbb{P}^{\pi_{\theta_{k+1}}} (I - \gamma \mathbb{P}^{\pi_{\theta_{k+1}}})^{-1} \sum_{\ell=1}^k \gamma^{k-\ell+1} \left(\prod_{s=\ell+1}^k \mathbb{P}^{\pi_{\theta_s}} \right) (I - \gamma \mathbb{P}^{\pi_{\theta_\ell}}) \epsilon_\ell^c \right] \leq 20K \cdot C_{\rho, \rho^*} \phi^* \cdot \varepsilon_Q,$$

Now, by plugging (A.3.21), (A.3.24), (A.3.25), (A.3.29), and (A.3.30) into (A.3.19), with probability at least $1 - O(H_c) \exp(-\Omega(H_c^{-1} m_c))$, we have

$$(A.3.31) \quad \mathbb{E}_\rho \left[\sum_{k=0}^K Q^*(s, a) - Q^{\pi_{\theta_{k+1}}}(s, a) \right] \leq 2 \log |\mathcal{A}| \cdot K^{1/2} (1 - \gamma)^{-3} + 60K^2 C_{\rho, \rho^*} (\phi^* + \psi^* + 1) \cdot \varepsilon_f + 50K C_{\rho, \rho^*} \phi^* \cdot \varepsilon_Q.$$

Meanwhile, following from Propositions A.2.3 and A.2.4, it holds with probability at least $1 - 1/K$ that

$$(A.3.32) \quad \begin{aligned} \varepsilon_f &= O(R_a N_a^{-1/4} + R_a^{4/3} m_a^{-1/12} H_a^{7/2} (\log m_a)^{1/2}), \\ \varepsilon_Q &= O(R_c N_c^{-1/4} + R_c^{4/3} m_c^{-1/12} H_c^{7/2} (\log m_c)^{1/2}). \end{aligned}$$

Combining (A.3.31), (A.3.32), and the choices of parameters stated in the theorem, it holds with probability at least $1 - 1/K$ that

$$\mathbb{E}_\rho \left[\sum_{k=0}^K Q^*(s, a) - Q^{\pi_{\theta_{k+1}}}(s, a) \right] \leq (2(1 - \gamma)^{-3} \log |\mathcal{A}| + O(1)) \cdot K^{1/2},$$

which concludes the proof of Theorem A.2.5.

A.4. Supporting Results

In this section, we provide some supporting results in the proof of Theorems 2.3.4 and A.2.5. We introduce Lemma A.4.1, which applies to both Algorithms 1 and 4. To introduce Lemma A.4.1, for any policy π and action-value function Q , we define $\tilde{\pi}(a | s) \propto \exp(\beta^{-1} Q(s, a)) \cdot \pi(a | s)$.

Lemma A.4.1. For any $s \in \mathcal{S}$ and π^\dagger , we have

$$\begin{aligned} \beta^{-1} \cdot \langle Q(s, \cdot), \pi^\dagger(\cdot | s) - \tilde{\pi}(\cdot | s) \rangle &\leq \text{KL}(\pi^\dagger(\cdot | s) \| \pi(\cdot | s)) - \text{KL}(\pi^\dagger(\cdot | s) \| \tilde{\pi}(\cdot | s)) \\ &\quad + \langle \log(\tilde{\pi}(\cdot | s)/\pi(\cdot | s)) - \beta^{-1} \cdot Q(s, \cdot), \pi^\dagger(\cdot | s) - \tilde{\pi}(\cdot | s) \rangle. \end{aligned}$$

Proof. By calculation, it suffices to show that

$$\begin{aligned} & \langle \log(\tilde{\pi}(\cdot | s)/\pi(\cdot | s)), \pi^\dagger(\cdot | s) - \tilde{\pi}(\cdot | s) \rangle \\ & \leq \text{KL}(\pi^\dagger(\cdot | s) \| \pi(\cdot | s)) - \text{KL}(\pi^\dagger(\cdot | s) \| \tilde{\pi}(\cdot | s)). \end{aligned}$$

By the definition of the KL divergence, it holds for any $s \in \mathcal{S}$ that

$$\begin{aligned} & \text{KL}(\pi^\dagger(\cdot | s) \| \pi(\cdot | s)) - \text{KL}(\pi^\dagger(\cdot | s) \| \tilde{\pi}(\cdot | s)) \\ \text{(A.4.1)} \quad & = \langle \log(\tilde{\pi}(\cdot | s)/\pi(\cdot | s)), \pi^\dagger(\cdot | s) \rangle. \end{aligned}$$

Meanwhile, for the term on the RHS of (A.4.1), we have

$$\begin{aligned} & \langle \log(\tilde{\pi}(\cdot | s)/\pi_{\theta_k}(\cdot | s)), \pi^\dagger(\cdot | s) \rangle \\ & = \langle \log(\tilde{\pi}(\cdot | s)/\pi(\cdot | s)), \pi^\dagger(\cdot | s) - \tilde{\pi}(\cdot | s) \rangle \\ & \quad + \langle \log(\tilde{\pi}(\cdot | s)/\pi(\cdot | s)), \tilde{\pi}(\cdot | s) \rangle \\ & = \langle \log(\tilde{\pi}(\cdot | s)/\pi(\cdot | s)), \pi^\dagger(\cdot | s) - \tilde{\pi}(\cdot | s) \rangle + \text{KL}(\tilde{\pi}(\cdot | s) \| \pi(\cdot | s)) \\ \text{(A.4.2)} \quad & \geq \langle \log(\tilde{\pi}(\cdot | s)/\pi(\cdot | s)), \pi^\dagger(\cdot | s) - \tilde{\pi}(\cdot | s) \rangle. \end{aligned}$$

Combining (A.4.1) and (A.4.2), we obtain that

$$\begin{aligned} & \langle \log(\tilde{\pi}(\cdot | s)/\pi(\cdot | s)), \pi^\dagger(\cdot | s) - \tilde{\pi}(\cdot | s) \rangle \\ & \leq \text{KL}(\pi^\dagger(\cdot | s) \| \pi(\cdot | s)) - \text{KL}(\pi^\dagger(\cdot | s) \| \tilde{\pi}(\cdot | s)), \end{aligned}$$

which concludes the proof of Lemma A.4.1. \square

A.4.1. Local Linearization of DNNs

In the proofs of Propositions A.2.3 and A.2.4 in §A.5.2 and §A.5.3, respectively, we utilize the linearization of DNNs. We introduce some related auxiliary results here. First, we define the linearization \bar{u}_θ of the DNN $u_\theta \in \mathcal{U}(w, H, R)$ as follows,

$$\bar{u}_\theta(\cdot) = u_{\theta_0}(\cdot) + (\theta - \theta_0)^\top \nabla_{\theta_0} u_{\theta_0}(\cdot),$$

where θ_0 is the initialization of u_θ . The following lemmas characterize the linearization error.

Lemma A.4.2. Suppose that $H = O(m^{1/12} R^{-1/6} (\log m)^{-1/2})$ and $m = \Omega(d^{3/2} R^{-1} H^{-3/2} \log(m^{1/2}/R)^{3/2})$. Then with probability at least $1 - \exp(-\Omega(R^{2/3} m^{2/3} H))$ over the random initialization θ_0 , it holds for any $\theta \in \mathcal{B}(\theta_0, R)$ and any $(s, a) \in \mathcal{S} \times \mathcal{A}$ that

$$\|\nabla_\theta u_\theta(s, a) - \nabla_\theta u_{\theta_0}(s, a)\|_2 = O(R^{1/3} m^{-1/6} H^{5/2} (\log m)^{1/2})$$

and

$$\|\nabla_\theta u_\theta(s, a)\|_2 = O(H).$$

Proof. See the proof of Lemma A.5 in Gao et al. (2019) for a detailed proof. \square

Lemma A.4.3. Suppose that $H = O(m^{1/12} R^{-1/6} (\log m)^{-1/2})$ and $m = \Omega(d^{3/2} R^{-1} H^{-3/2} \log(m^{1/2}/R)^{3/2})$. Then with probability at least $1 - \exp(-\Omega(R^{2/3} m^{2/3} H))$ over the random initialization θ_0 , it holds for any $\theta \in \mathcal{B}(\theta_0, R)$ and any $(s, a) \in \mathcal{S} \times \mathcal{A}$ that

$$|u_\theta(s, a) - \bar{u}_\theta(s, a)| = O(R^{4/3} m^{-1/6} H^{5/2} (\log m)^{1/2}).$$

Proof. Recall that

$$\bar{u}_\theta(s, a) = u_{\theta_0}(s, a) + (\theta - \theta_0)^\top \nabla_\theta u_{\theta_0}(s, a).$$

By mean value theorem, there exists $t \in [0, 1]$, which depends on θ and (s, a) , such that

$$u_\theta(s, a) - \bar{u}_\theta(s, a) = (\theta - \theta_0)^\top (\nabla_\theta u_{\theta_0+t(\theta-\theta_0)}(s, a) - \nabla_\theta u_{\theta_0}(s, a)).$$

Further by Lemma A.4.2, we have

$$\begin{aligned} |u_\theta(s, a) - \bar{u}_\theta(s, a)| &\leq \|\theta - \theta_0\|_2 \cdot \left\| \nabla_\theta u_{\theta_0+t(\theta-\theta_0)}(s, a) - \nabla_\theta u_{\theta_0}(s, a) \right\|_2 \\ &= O(R^{4/3} m^{-1/6} H^{5/2} (\log m)^{1/2}), \end{aligned}$$

where we use Cauchy-Schwarz inequality in the first inequality. This concludes the proof of Lemma A.4.3. \square

We denote by $x^{(h)}$ the output of the h -th layer of the DNN $u_\theta \in \mathcal{U}(m, H, R)$, and $x^{(h),0}$ the output of the h -th layer of the DNN $u_{\theta_0} \in \mathcal{U}(m, H, R)$. The following lemma upper bounds the distance between $x^{(h)}$ and $x^{(h),0}$.

Lemma A.4.4. With probability at least $1 - \exp(-\Omega(R^{2/3} m^{2/3} H))$ over the random initialization θ_0 , for any $\theta \in \mathcal{B}(\theta_0, R)$ and any $h \in [H]$, we have

$$\|x^{(h)} - x^{(h),0}\|_2 = O(RH^{5/2} m^{-1/2} (\log m)^{1/2}).$$

Also, with probability at least $1 - O(H) \exp(-\Omega(H^{-1}m))$ over the random initialization θ_0 , for any $\theta \in \mathcal{B}(\theta_0, R)$ and any $h \in [H]$, it holds that

$$2/3 \leq \|x^{(h)}\|_2 \leq 4/3.$$

Proof. The first inequality follows from Lemma A.5 in Gao et al. (2019), and the second inequality follows from Lemma 7.1 in Allen-Zhu et al. (2018b). \square

A.5. Proofs of Propositions

A.5.1. Proof of Proposition 2.2.1

The proof follows the proof of Proposition 3.1 in Liu et al. (2019). First, we write the update $\tilde{\pi}_{k+1} \leftarrow \operatorname{argmax}_{\pi} \mathbb{E}_{\nu_k} [\langle Q_{\omega_k}(s, \cdot), \pi(\cdot | s) \rangle - \beta \cdot \text{KL}(\pi(\cdot | s) \| \pi_{\theta_k}(\cdot | s))]$ as a constrained optimization problem in the following way,

$$\begin{aligned} & \max_{\pi} \mathbb{E}_{\nu_k} [\langle \pi(\cdot | s), Q_{\omega_k}(s, \cdot) \rangle - \beta \cdot \text{KL}(\pi(\cdot | s) \| \pi_{\theta_k}(\cdot | s))] \\ & \text{s.t. } \sum_{a \in \mathcal{A}} \pi(a | s) = 1, \quad \text{for any } s \in \mathcal{S}. \end{aligned}$$

We consider the Lagrangian of the above programming problem,

$$\begin{aligned} & \int_{s \in \mathcal{S}} \left(\langle \pi(\cdot | s), Q_{\omega_k}(s, \cdot) \rangle - \beta \cdot \text{KL}(\pi(\cdot | s) \| \pi_{\theta_k}(\cdot | s)) \right) d\nu_k(s) \\ & + \int_{s \in \mathcal{S}} \left(\sum_{a \in \mathcal{A}} \pi(a | s) - 1 \right) d\lambda(s), \end{aligned}$$

where $\lambda(\cdot)$ is the dual parameter, which is a function on \mathcal{S} . Now, by plugging in

$$\pi_{\theta_k}(a | s) = \frac{\exp(\tau_k^{-1} f_{\theta_k}(s, a))}{\sum_{a' \in \mathcal{A}} \exp(\tau_k^{-1} f_{\theta_k}(s, a'))},$$

we have the following optimality condition,

$$\begin{aligned} Q_{\omega_k}(s, a) + \beta \tau_k^{-1} f_{\theta_k}(s, a) - \beta \cdot \left(\log \left(\sum_{a' \in \mathcal{A}} \exp(\tau_k^{-1} f_{\theta_k}(s, a')) \right) + \log \pi(a | s) + 1 \right) + \frac{\lambda(s)}{\nu_k(s)} \\ = 0, \end{aligned}$$

for any $(s, a) \in \mathcal{S} \times \mathcal{A}$. Note that $\log(\sum_{a' \in \mathcal{A}} \exp(\tau_k^{-1} f_{\theta_k}(s, a')))$ is only a function of s .

Thus, we have

$$\widehat{\pi}_{k+1}(a | s) \propto \exp(\beta^{-1} Q_{\omega_k}(s, a) + \tau_k^{-1} f_{\theta_k}(s, a))$$

for any $(s, a) \in \mathcal{S} \times \mathcal{A}$, which concludes the proof of Proposition 2.2.1.

A.5.2. Proof of Proposition A.2.3

We define the local linearization of f_{θ} as follows,

$$(A.5.1) \quad \bar{f}_{\theta} = f_{\theta_0} + (\theta - \theta_0)^{\top} \nabla_{\theta_0} f_{\theta}.$$

Meanwhile, we denote by

$$\begin{aligned}
g_n &= (f_{\theta(n)} - \tilde{\tau} \cdot (\beta^{-1}Q_\omega + \tau^{-1}f_\theta)) \cdot \nabla_\theta f_{\theta(n)}, & g_n^e &= \mathbb{E}_{\rho_{\pi_\theta}}[g_n], \\
\bar{g}_n &= (\bar{f}_{\theta(n)} - \tilde{\tau} \cdot (\beta^{-1}Q_\omega + \tau^{-1}f_\theta)) \cdot \nabla_\theta f_{\theta_0}, & \bar{g}_n^e &= \mathbb{E}_{\rho_{\pi_\theta}}[\bar{g}_n], \\
g_* &= (f_{\theta_*} - \tilde{\tau} \cdot (\beta^{-1}Q_\omega + \tau^{-1}f_\theta)) \cdot \nabla_\theta f_{\theta_*}, & g_*^e &= \mathbb{E}_{\rho_{\pi_\theta}}[g_*], \\
\bar{g}_* &= (\bar{f}_{\theta_*} - \tilde{\tau} \cdot (\beta^{-1}Q_\omega + \tau^{-1}f_\theta)) \cdot \nabla_\theta f_{\theta_0}, & \bar{g}_*^e &= \mathbb{E}_{\rho_{\pi_\theta}}[\bar{g}_*],
\end{aligned}
\tag{A.5.2}$$

where θ_* satisfies that

$$\theta_* = \Gamma_{\mathcal{B}(\theta_0, R_a)}(\theta_* - \alpha \cdot \bar{g}_*^e).
\tag{A.5.3}$$

By Algorithm 5, we know that

$$\theta(n+1) = \Gamma_{\mathcal{B}(\theta_0, R_a)}(\theta(n) - \alpha \cdot g_n).
\tag{A.5.4}$$

By (A.5.3) and (A.5.4), we have

$$\begin{aligned}
&\mathbb{E}_{\rho_{\pi_\theta}} [\|\theta(n+1) - \theta_*\|_2^2 \mid \theta(n)] \\
&= \mathbb{E}_{\rho_{\pi_\theta}} [\|\Gamma_{\mathcal{B}(\theta_0, R_a)}(\theta(n) - \alpha \cdot g_n) - \Gamma_{\mathcal{B}(\theta_0, R_a)}(\theta_* - \alpha \cdot \bar{g}_*^e)\|_2^2 \mid \theta(n)] \\
&\leq \mathbb{E}_{\rho_{\pi_\theta}} [\|(\theta(n) - \alpha \cdot g_n) - (\theta_* - \alpha \cdot \bar{g}_*^e)\|_2^2 \mid \theta(n)] \\
&= \|\theta(n) - \theta_*\|_2^2 + 2\alpha \cdot \underbrace{\langle \theta_* - \theta(n), g_n^e - \bar{g}_*^e \rangle}_{(i)} + \alpha^2 \cdot \underbrace{\mathbb{E}_{\rho_{\pi_\theta}} [\|g_n - \bar{g}_*^e\|_2^2 \mid \theta(n)]}_{(ii)},
\end{aligned}
\tag{A.5.5}$$

where we use the fact that $\Gamma_{\mathcal{B}(\theta_0, R_a)}$ is a contraction mapping in the first inequality. We upper bound term (i) and term (ii) on the RHS of (A.5.5) in the sequel.

Upper Bound of Term (i). By Cauchy–Schwarz inequality, it holds that

$$\begin{aligned}
\langle \theta_* - \theta(n), g_n^e - \bar{g}_*^e \rangle &= \langle \theta_* - \theta(n), g_n^e - \bar{g}_n^e \rangle + \langle \theta_* - \theta(n), \bar{g}_n^e - \bar{g}_*^e \rangle \\
&\leq \|\theta_* - \theta(n)\|_2 \cdot \|g_n^e - \bar{g}_n^e\|_2 + \langle \theta_* - \theta(n), \bar{g}_n^e - \bar{g}_*^e \rangle \\
\text{(A.5.6)} \quad &\leq 2R_a \cdot \|g_n^e - \bar{g}_n^e\|_2 + \langle \theta_* - \theta(n), \bar{g}_n^e - \bar{g}_*^e \rangle,
\end{aligned}$$

where we use the fact that $\theta(n), \theta_* \in \mathcal{B}(\theta_0, R_a)$ in the last inequality. Further, by the definitions in (A.5.2), it holds that

$$\begin{aligned}
\langle \theta_* - \theta(n), \bar{g}_n^e - \bar{g}_*^e \rangle &= \mathbb{E}_{\rho_{\pi_\theta}} [(\bar{f}_{\theta(n)} - \bar{f}_{\theta_*}) \cdot \langle \theta_* - \theta(n), \nabla_{\theta} f_{\theta_0} \rangle] \\
&= \mathbb{E}_{\rho_{\pi_\theta}} [(\bar{f}_{\theta(n)} - \bar{f}_{\theta_*}) \cdot (\bar{f}_{\theta_*} - \bar{f}_{\theta(n)})] \\
\text{(A.5.7)} \quad &= -\mathbb{E}_{\rho_{\pi_\theta}} [(\bar{f}_{\theta(n)} - \bar{f}_{\theta_*})^2],
\end{aligned}$$

where we use (A.5.1) in the second equality. Combining (A.5.6) and (A.5.7), we obtain the following upper bound of term (i),

$$\text{(A.5.8)} \quad \langle \theta_* - \theta(n), g_n^e - \bar{g}_*^e \rangle \leq 2R_a \cdot \|g_n^e - \bar{g}_n^e\|_2 - \mathbb{E}_{\rho_{\pi_\theta}} [(\bar{f}_{\theta(n)} - \bar{f}_{\theta_*})^2].$$

Upper Bound of Term (ii). We now upper bound term (ii) on the RHS of (A.5.5). It holds by Cauchy-Schwarz inequality that

$$\begin{aligned}
\mathbb{E}_{\rho_{\pi_\theta}} [\|g_n - \bar{g}_*^e\|_2^2 | \theta(n)] &\leq 2\mathbb{E}_{\rho_{\pi_\theta}} [\|g_n - g_n^e\|_2^2 | \theta(n)] + 2\|g_n^e - \bar{g}_*^e\|_2^2 \\
\text{(A.5.9)} \quad &\leq 2 \underbrace{\mathbb{E}_{\rho_{\pi_\theta}} [\|g_n - g_n^e\|_2^2 | \theta(n)]}_{\text{(ii).a}} + 4 \underbrace{\|g_n^e - \bar{g}_n^e\|_2^2}_{\text{(ii).b}} + 4 \underbrace{\|\bar{g}_n^e - \bar{g}_*^e\|_2^2}_{\text{(ii).c}}.
\end{aligned}$$

We upper bound term (ii).a, term (ii).b, and term (ii).c in the sequel.

Upper Bound of Term (ii).a. Note that

$$(A.5.10) \quad \mathbb{E}_{\rho_{\pi_\theta}} [\|g_n - g_n^e\|_2^2 | \theta(n)] = \mathbb{E}_{\rho_{\pi_\theta}} [\|g_n\|_2^2 - \|g_n^e\|_2^2 | \theta(n)] \leq \mathbb{E}_{\rho_{\pi_\theta}} [\|g_n\|_2^2 | \theta(n)].$$

Meanwhile, by the definition of g_n in (A.5.2), it holds that

$$(A.5.11) \quad \|g_n\|_2^2 = (f_{\theta(n)} - \tilde{\tau} \cdot (\beta^{-1}Q_\omega + \tau^{-1}f_\theta))^2 \cdot \|\nabla_\theta f_{\theta(n)}\|_2^2.$$

We first upper bound f_θ as follows,

$$f_\theta^2 = x^{(H_a)\top} b b^\top x^{(H_a)} = x^{(H_a)\top} x^{(H_a)} = \|x^{(H_a)}\|_2^2,$$

where $x^{(H_a)}$ is the output of the H_a -th layer of the DNN f_θ . Further combining Lemma A.4.4, it holds with probability at least $1 - O(H_a) \exp(-\Omega(H_a^{-1}m_a))$ that

$$(A.5.12) \quad |f_\theta| \leq 2.$$

Following from similar arguments, with probability at least $1 - O(H_a) \exp(-\Omega(H_a^{-1}m_a))$, we have

$$(A.5.13) \quad |Q_\omega| \leq 2, \quad |f_{\theta(n)}| \leq 2.$$

Combining Lemma A.4.2, (A.5.10), (A.5.11), (A.5.12), and (A.5.13), it holds with probability at least $1 - \exp(-\Omega(R_a^{2/3}m_a^{2/3}H_a))$ that

$$(A.5.14) \quad \mathbb{E}_{\rho_{\pi_\theta}} [\|g_n - g_n^e\|_2^2 | \theta(n)] = O(H_a^2),$$

which establishes an upper bound of term (ii).a.

Upper Bound of Term (ii).b. It holds that

$$\begin{aligned}
\|g_n^e - \bar{g}_n^e\|_2 &= \left\| \mathbb{E}_{\rho_{\pi_\theta}} \left[(f_{\theta(n)} - \tilde{\tau} \cdot (\beta^{-1}Q_\omega + \tau^{-1}f_\theta)) \cdot \nabla_\theta f_{\theta(n)} \right. \right. \\
&\quad \left. \left. - (\bar{f}_{\theta(n)} - \tilde{\tau} \cdot (\beta^{-1}Q_\omega + \tau^{-1}f_\theta)) \cdot \nabla_\theta f_{\theta_0} \right] \right\|_2 \\
&\leq \mathbb{E}_{\rho_{\pi_\theta}} \left[\|f_{\theta(n)} \nabla_\theta f_{\theta(n)} - \bar{f}_{\theta(n)} \nabla_\theta f_{\theta_0}\|_2 \right] \\
&\quad + \tilde{\tau} \cdot \mathbb{E}_{\rho_{\pi_\theta}} \left[\|(\beta^{-1}Q_\omega + \tau^{-1}f_\theta) \cdot (\nabla_\theta f_{\theta_0} - \nabla_\theta f_{\theta(n)})\|_2 \right] \\
\text{(A.5.15)} \quad &\leq \mathbb{E}_{\rho_{\pi_\theta}} \left[\|f_{\theta(n)} \nabla_\theta f_{\theta_0} - \bar{f}_{\theta(n)} \nabla_\theta f_{\theta_0}\|_2 \right] + \mathbb{E}_{\rho_{\pi_\theta}} \left[\|f_{\theta(n)} \nabla_\theta f_{\theta(n)} - f_{\theta(n)} \nabla_\theta f_{\theta_0}\|_2 \right] \\
&\quad + \mathbb{E}_{\rho_{\pi_\theta}} \left[\|\tilde{\tau} \cdot (\beta^{-1}Q_\omega + \tau^{-1}f_\theta) \cdot (\nabla_\theta f_{\theta_0} - \nabla_\theta f_{\theta(n)})\|_2 \right].
\end{aligned}$$

We upper bound the three terms on the RHS of (A.5.15) in the sequel, respectively.

For the term $\|f_{\theta(n)} \nabla_\theta f_{\theta_0} - \bar{f}_{\theta(n)} \nabla_\theta f_{\theta_0}\|_2$ on the RHS of (A.5.15), following from Lemmas A.4.2 and A.4.3, it holds with probability at least $1 - \exp(-\Omega(R_a^{2/3} m_a^{2/3} H_a))$ that

$$\text{(A.5.16)} \quad \|f_{\theta(n)} \nabla_\theta f_{\theta_0} - \bar{f}_{\theta(n)} \nabla_\theta f_{\theta_0}\|_2 = O(R_a^{4/3} m_a^{-1/6} H_a^{7/2} (\log m_a)^{1/2}).$$

For the term $\|f_{\theta(n)} \nabla_\theta f_{\theta(n)} - f_{\theta(n)} \nabla_\theta f_{\theta_0}\|_2$ on the RHS of (A.5.15), following from (A.5.13) and Lemma A.4.2, with probability at least $1 - \exp(-\Omega(R_a^{2/3} m_a^{2/3} H_a))$, we have

$$\text{(A.5.17)} \quad \|f_{\theta(n)} \nabla_\theta f_{\theta(n)} - f_{\theta(n)} \nabla_\theta f_{\theta_0}\|_2 = O(R_a^{1/3} m_a^{-1/6} H_a^{5/2} (\log m_a)^{1/2}).$$

For the term $\|\tilde{\tau} \cdot (\beta^{-1}Q_\omega + \tau^{-1}f_\theta) \cdot (\nabla_\theta f_{\theta_0} - \nabla_\theta f_{\theta(n)})\|_2$ on the RHS of (A.5.15), we first upper bound $\tilde{\tau} \cdot (\beta^{-1}Q_\omega + \tau^{-1}f_\theta)$ as follows,

$$|\tilde{\tau} \cdot (\beta^{-1}Q_\omega + \tau^{-1}f_\theta)| \leq 2,$$

where we use (A.5.12), (A.5.13), and the fact that $\tilde{\tau}^{-1} = \beta^{-1} + \tau^{-1}$. Further combining Lemma A.4.2, it holds with probability at least $1 - \exp(-\Omega(R_a^{2/3}m_a^{2/3}H_a))$ that

$$(A.5.18) \quad \|\tilde{\tau} \cdot (\beta^{-1}Q_\omega + \tau^{-1}f_\theta) \cdot (\nabla_\theta f_{\theta_0} - \nabla_\theta f_{\theta(n)})\|_2 = O(R_a^{1/3}m_a^{-1/6}H_a^{5/2}(\log m_a)^{1/2}).$$

Now, combining (A.5.15), (A.5.16), (A.5.17), and (A.5.18), it holds with probability at least $1 - \exp(-\Omega(R_a^{2/3}m_a^{2/3}H_a))$ that

$$(A.5.19) \quad \|g_n^e - \bar{g}_n^e\|_2^2 = O(R_a^{8/3}m_a^{-1/3}H_a^7 \log m_a),$$

which establishes an upper bound of term (ii).b.

Upper Bound of Term (ii).c. It holds that

$$\|\bar{g}_n^e - \bar{g}_*^e\|_2^2 = \|\mathbb{E}_{\rho_{\pi_\theta}}[(\bar{f}_{\theta(n)} - \bar{f}_{\theta_*})\nabla_\theta f_{\theta_0}]\|_2^2 \leq \mathbb{E}_{\rho_{\pi_\theta}}[(\bar{f}_{\theta(n)} - \bar{f}_{\theta_*})^2 \cdot \|\nabla_\theta f_{\theta_0}\|_2^2].$$

Further combining Lemma A.4.2, it holds with probability at least $1 - \exp(-\Omega(R_a^{2/3}m_a^{2/3}H_a))$ that

$$(A.5.20) \quad \|\bar{g}_n^e - \bar{g}_*^e\|_2^2 \leq O(H_a^2) \cdot \mathbb{E}_{\rho_{\pi_\theta}}[(\bar{f}_{\theta(n)} - \bar{f}_{\theta_*})^2],$$

which establishes an upper bound of term (ii).c.

Now, combining (A.5.9), (A.5.14), (A.5.19), and (A.5.20), we have

$$(A.5.21) \quad \mathbb{E}_{\rho_{\pi_\theta}} [\|g_n - \bar{g}_*\|_2^2 | \theta(n)] \leq O(R_a^{8/3} m_a^{-1/3} H_a^7 \log m_a) + O(H_a^2) \cdot \mathbb{E}_{\rho_{\pi_\theta}} [(\bar{f}_{\theta(n)} - \bar{f}_{\theta_*})^2],$$

which is an upper bound of term (ii) on the RHS of (A.5.5).

By plugging the upper bound of term (i) in (A.5.8) and the upper bound of term (ii) in (A.5.21) into (A.5.5), combining (A.5.19), with probability at least $1 - \exp(-\Omega(R_a^{2/3} m_a^{2/3} H_a))$, we have

$$(A.5.22) \quad \begin{aligned} & \mathbb{E}_{\rho_{\pi_\theta}} [\|\theta(n+1) - \theta_*\|_2^2 | \theta(n)] \\ & \leq \|\theta(n) - \theta_*\|_2^2 + 2\alpha \cdot \left(O(R_a^{7/3} m_a^{-1/6} H_a^{7/2} (\log m_a)^{1/2}) - \mathbb{E}_{\rho_{\pi_\theta}} [(\bar{f}_{\theta(n)} - \bar{f}_{\theta_*})^2] \right) \\ & \quad + \alpha^2 \cdot \left(O(R_a^{8/3} m_a^{-1/3} H_a^7 \log m_a) + O(H_a^2) \cdot \mathbb{E}_{\rho_{\pi_\theta}} [(\bar{f}_{\theta(n)} - \bar{f}_{\theta_*})^2] \right). \end{aligned}$$

Rearranging terms in (A.5.22), it holds with probability at least $1 - \exp(-\Omega(R_a^{2/3} m_a^{2/3} H_a))$ that

$$(A.5.23) \quad \begin{aligned} & (2\alpha - \alpha^2 \cdot O(H_a^2)) \cdot \mathbb{E}_{\rho_{\pi_\theta}} [(\bar{f}_{\theta(n)} - \bar{f}_{\theta_*})^2] \\ & \leq \|\theta(n) - \theta_*\|_2^2 - \mathbb{E}_{\rho_{\pi_\theta}} [\|\theta(n+1) - \theta_*\|_2^2 | \theta(n)] + \alpha \cdot O(R_a^{8/3} m_a^{-1/6} H_a^7 \log m_a). \end{aligned}$$

By telescoping the sum and using Jensen's inequality in (A.5.23), we have

$$\begin{aligned} \mathbb{E}_{\rho_{\pi_\theta}} [(\bar{f}_{\bar{\theta}} - \bar{f}_{\theta_*})^2] &\leq \frac{1}{N_a} \cdot \sum_{n=0}^{N_a-1} \mathbb{E}_{\rho_{\pi_\theta}} [(\bar{f}_{\theta^{(n)}} - \bar{f}_{\theta_*})^2] \\ &\leq 1/N_a \cdot (2\alpha - \alpha^2 \cdot O(H_a^2))^{-1} \cdot (\|\theta_0 - \theta_*\|_2^2 + \alpha N_a \cdot O(R_a^{8/3} m_a^{-1/6} H_a^7 \log m_a)) \\ &\leq N_a^{-1/2} \cdot \|\theta_0 - \theta_*\|_2^2 + O(R_a^{8/3} m_a^{-1/6} H_a^7 \log m_a), \end{aligned}$$

where the last line comes from the choices that $\alpha = N_a^{-1/2}$ and $H_a = O(N_a^{1/4})$. Further combining Lemma A.4.3 and using triangle inequality, we have

$$(A.5.24) \quad \mathbb{E}_{\rho_{\pi_\theta}} [(f_{\bar{\theta}} - \bar{f}_{\theta_*})^2] = O(R_a^2 N_a^{-1/2} + R_a^{8/3} m_a^{-1/6} H_a^7 \log m_a).$$

By the definition of θ_* in (A.5.3), we know that

$$(A.5.25) \quad \langle \bar{g}_*^e, \theta - \theta_* \rangle \geq 0, \quad \text{for any } \theta \in \mathcal{B}(\theta_0, R_a).$$

By plugging the definition of \bar{g}_*^e into (A.5.25), we have

$$\mathbb{E}_{\rho_{\pi_\theta}} [\langle \bar{f}_{\theta_*} - \tilde{\tau} \cdot (\beta^{-1} Q_\omega + \tau^{-1} f_\theta), \bar{f}_{\theta^\dagger} - \bar{f}_{\theta_*} \rangle] \geq 0, \quad \text{for any } \theta^\dagger \in \mathcal{B}(\theta_0, R_a),$$

which is equivalent to

$$(A.5.26) \quad \theta_* = \operatorname{argmin}_{\theta^\dagger \in \mathcal{B}(\theta_0, R_a)} \mathbb{E}_{\rho_{\pi_\theta}} [(\bar{f}_{\theta^\dagger} - \tilde{\tau} \cdot (\beta^{-1} Q_\omega + \tau^{-1} f_\theta))^2].$$

Meanwhile, by the fact that $\theta_0 = \omega_0$, we have

$$\begin{aligned}\tilde{\tau} \cdot (\beta^{-1}\bar{Q}_\omega + \tau^{-1}\bar{f}_\theta) &= \tilde{\tau} \cdot (\beta^{-1} \cdot (Q_{\omega_0} + (\omega - \omega_0)^\top \nabla_\omega Q_{\omega_0}) + \tau^{-1} \cdot (f_{\theta_0} + (\theta - \theta_0)^\top \nabla_\theta f_{\theta_0})) \\ &= f_{\theta_0} + (\tilde{\tau} \cdot (\beta^{-1}\omega + \tau^{-1}\theta) - \theta_0)^\top \nabla_\theta f_{\theta_0},\end{aligned}$$

where the second line comes from $\tilde{\tau}^{-1} = \beta^{-1} + \tau^{-1}$. Note that $\theta \in \mathcal{B}(\theta_0, R_a)$, $\omega \in \mathcal{B}(\omega_0, R_c)$, $\theta_0 = \omega_0$, and $R_a = R_c$, we know that $\tilde{\tau} \cdot (\beta^{-1}\omega + \tau^{-1}\theta) \in \mathcal{B}(\theta_0, R_a)$. Therefore, with probability at least $1 - \exp(-\Omega(R_a^{2/3} m_a^{2/3} H_a))$ we have

$$\begin{aligned}&\mathbb{E}_{\rho_{\pi_\theta}} [(\bar{f}_{\theta_*} - \tilde{\tau} \cdot (\beta^{-1}Q_\omega + \tau^{-1}f_\theta))^2] \\ &\leq \mathbb{E}_{\rho_{\pi_\theta}} [(\tilde{\tau} \cdot (\beta^{-1}\bar{Q}_\omega + \tau^{-1}\bar{f}_\theta) - \tilde{\tau} \cdot (\beta^{-1}Q_\omega + \tau^{-1}f_\theta))^2] \\ &\leq \tilde{\tau}^2 \cdot \beta^{-2} \cdot \mathbb{E}_{\rho_{\pi_\theta}} [(\bar{Q}_\omega - Q_\omega)^2] + \tilde{\tau}^2 \cdot \tau^{-2} \cdot \mathbb{E}_{\rho_{\pi_\theta}} [(\bar{f}_\theta - f_\theta)^2] \\ (A.5.27) \quad &= O(R_a^{8/3} m_a^{-1/3} H_a^5 \log m_a),\end{aligned}$$

where the first inequality comes from (A.5.26), and the last inequality comes from Lemma A.4.3 and the fact that $R_c = R_a$, $m_c = m_a$, and $H_c = H_a$. Combining (A.5.24) and (A.5.27), by triangle inequality, we have

$$\mathbb{E}_{\rho_{\pi_\theta}} [(f_{\bar{\theta}}(s, a) - \tilde{\tau} \cdot (\beta^{-1}Q_\omega(s, a) + \tau^{-1}f_\theta(s, a)))^2] = O(R_a^2 N_a^{-1/2} + R_a^{8/3} m_a^{-1/6} H_a^7 \log m_a),$$

which finishes the proof of Proposition A.2.3.

A.5.3. Proof of Proposition A.2.4

The proof is similar to that of Proposition A.2.3 in §A.5.2. For the completeness of the paper, we present it here. We define the local linearization of Q_ω as follows,

$$(A.5.28) \quad \bar{Q}_\omega = Q_{\omega_0} + (\omega - \omega_0)^\top \nabla_{\omega_0} Q_\omega.$$

We denote by

$$(A.5.29) \quad \begin{aligned} g_n &= (Q_{\omega(n)}(s_0, a_0) - \gamma \cdot Q_\omega(s_1, a_1) - (1 - \gamma) \cdot r_0) \cdot \nabla_\omega Q_{\omega(n)}(s_0, a_0), & g_n^e &= \mathbb{E}_{\pi_\theta}[g_n], \\ \bar{g}_n &= (\bar{Q}_{\omega(n)}(s_0, a_0) - \gamma \cdot Q_\omega(s_1, a_1) - (1 - \gamma) \cdot r_0) \cdot \nabla_\omega Q_{\omega_0}(s_0, a_0), & \bar{g}_n^e &= \mathbb{E}_{\pi_\theta}[\bar{g}_n], \\ g_* &= (Q_{\omega_*}(s_0, a_0) - \gamma \cdot Q_\omega(s_1, a_1) - (1 - \gamma) \cdot r_0) \cdot \nabla_\omega Q_{\omega_*}(s_0, a_0), & g_*^e &= \mathbb{E}_{\pi_\theta}[g_*], \\ \bar{g}_* &= (\bar{Q}_{\omega_*}(s_0, a_0) - \gamma \cdot Q_\omega(s_1, a_1) - (1 - \gamma) \cdot r_0) \cdot \nabla_\omega Q_{\omega_0}(s_0, a_0), & \bar{g}_*^e &= \mathbb{E}_{\pi_\theta}[\bar{g}_*], \end{aligned}$$

where ω_* satisfies that

$$(A.5.30) \quad \omega_* = \Gamma_{\mathcal{B}(\omega_0, R_c)}(\omega_* - \alpha \cdot \bar{g}_*^e).$$

Here the expectation $\mathbb{E}_{\pi_\theta}[\cdot]$ is taken following $(s_0, a_0) \sim \rho_{\pi_\theta}(\cdot)$, $s_1 \sim P(\cdot | s_0, a_0)$, $a_1 \sim \pi_\theta(\cdot | s_1)$, and $r_0 = \mathcal{R}(s_0, a_0)$. By Algorithm 6, we know that

$$\omega(n+1) = \Gamma_{\mathcal{B}(\omega_0, R_c)}(\omega(n) - \eta \cdot g_n).$$

Note that

$$\begin{aligned}
& \mathbb{E}_{\pi_\theta} [\|\omega(n+1) - \omega_*\|_2^2 \mid \omega(n)] \\
&= \mathbb{E}_{\pi_\theta} [\|\Gamma_{\mathcal{B}(\omega_0, R_c)}(\omega(n) - \eta \cdot g_n) - \Gamma_{\mathcal{B}(\omega_0, R_c)}(\omega_* - \eta \cdot \bar{g}_*)\|_2^2 \mid \omega(n)] \\
&\leq \mathbb{E}_{\pi_\theta} [\|(\omega(n) - \eta \cdot g_n) - (\omega_* - \eta \cdot \bar{g}_*)\|_2^2 \mid \omega(n)] \\
\text{(A.5.31)} \quad &= \|\omega(n) - \omega_*\|_2^2 + 2\eta \cdot \underbrace{\langle \omega_* - \omega(n), g_n^e - \bar{g}_*^e \rangle}_{\text{(iii)}} + \eta^2 \cdot \underbrace{\mathbb{E}_{\pi_\theta} [\|g_n - \bar{g}_*^e\|_2^2 \mid \omega(n)]}_{\text{(iv)}}.
\end{aligned}$$

We upper bound term (iii) and term (iv) on the RHS of (A.5.31) in the sequel.

Upper Bound of Term (iii). By Hölder's inequality, it holds that

$$\begin{aligned}
& \langle \omega_* - \omega(n), g_n^e - \bar{g}_*^e \rangle \\
&= \langle \omega_* - \omega(n), g_n^e - \bar{g}_n^e \rangle + \langle \omega_* - \omega(n), \bar{g}_n^e - \bar{g}_*^e \rangle \\
&\leq \|\omega_* - \omega(n)\|_2 \cdot \|g_n^e - \bar{g}_n^e\|_2 + \langle \omega_* - \omega(n), \bar{g}_n^e - \bar{g}_*^e \rangle \\
\text{(A.5.32)} \quad &\leq 2R_c \cdot \|g_n^e - \bar{g}_n^e\|_2 + \langle \omega_* - \omega(n), \bar{g}_n^e - \bar{g}_*^e \rangle,
\end{aligned}$$

where we use the fact that $\omega(n), \omega_* \in \mathcal{B}(\omega_0, R_c)$ in the last line. Further, by the definitions in (A.5.29), it holds that

$$\begin{aligned}
& \langle \omega_* - \omega(n), \bar{g}_n^e - \bar{g}_*^e \rangle \\
&= \mathbb{E}_{\pi_\theta} [(\bar{Q}_{\omega(n)}(s_0, a_0) - \bar{Q}_{\omega_*}(s_0, a_0)) \cdot \langle \omega_* - \omega(n), \nabla_\omega Q_{\omega_0}(s_0, a_0) \rangle] \\
&= \mathbb{E}_{\pi_\theta} [(\bar{Q}_{\omega(n)}(s_0, a_0) - \bar{Q}_{\omega_*}(s_0, a_0)) \cdot (\bar{Q}_{\omega_*}(s_0, a_0) - \bar{Q}_{\omega(n)}(s_0, a_0))] \\
\text{(A.5.33)} \quad &= -\mathbb{E}_{\pi_\theta} [(\bar{Q}_{\omega(n)}(s_0, a_0) - \bar{Q}_{\omega_*}(s_0, a_0))^2] = -\mathbb{E}_{\rho_{\pi_\theta}} [(\bar{Q}_{\omega(n)} - \bar{Q}_{\omega_*})^2],
\end{aligned}$$

where the second equality comes from (A.5.28), and the last equality comes from the fact that the expectation is only taken to the state-action pair (s_0, a_0) . Combining (A.5.32) and (A.5.33), we obtain the following upper bound of term (i),

$$(A.5.34) \quad \langle \omega_* - \omega(n), g_n^e - \bar{g}_*^e \rangle \leq 2R_c \cdot \|g_n^e - \bar{g}_n^e\|_2 - \mathbb{E}_{\rho_{\pi_\theta}} [(\bar{Q}_{\omega(n)} - \bar{Q}_{\omega_*})^2].$$

Upper Bound of Term (iv). We now upper bound term (iv) on the RHS of (A.5.31).

It holds by Cauchy-Schwarz inequality that

$$(A.5.35) \quad \begin{aligned} & \mathbb{E}_{\pi_\theta} [\|g_n - \bar{g}_*^e\|_2^2 | \omega(n)] \\ & \leq 2\mathbb{E}_{\pi_\theta} [\|g_n - g_n^e\|_2^2 | \omega(n)] + 2\|g_n^e - \bar{g}_*^e\|_2^2 \\ & \leq 2\underbrace{\mathbb{E}_{\pi_\theta} [\|g_n - g_n^e\|_2^2 | \omega(n)]}_{(iv).a} + 4\underbrace{\|g_n^e - \bar{g}_n^e\|_2^2}_{(iv).b} + 4\underbrace{\|\bar{g}_n^e - \bar{g}_*^e\|_2^2}_{(iv).c}. \end{aligned}$$

We upper bound term (iv).a, term (iv).b, and term (iv).c in the sequel.

Upper Bound of Term (iv).a. We now upper bound term (iv).a on the RHS of (A.5.35). By expanding the square, we have

$$(A.5.36) \quad \mathbb{E}_{\pi_\theta} [\|g_n - g_n^e\|_2^2 | \omega(n)] = \mathbb{E}_{\pi_\theta} [\|g_n\|_2^2 - \|g_n^e\|_2^2 | \omega(n)] \leq \mathbb{E}_{\pi_\theta} [\|g_n\|_2^2 | \omega(n)].$$

Meanwhile, by the definition of g_n in (A.5.29), it holds that

$$(A.5.37) \quad \|g_n\|_2^2 = (Q_{\omega(n)}(s_0, a_0) - \gamma \cdot Q_{\omega}(s_1, a_1) - (1 - \gamma) \cdot r_0)^2 \cdot \|\nabla_{\omega} Q_{\omega(n)}(s_0, a_0)\|_2^2.$$

We first upper bound Q_ω as follows,

$$Q_\omega^2 = x^{(H_c)\top} b b^\top x^{(H_c)} = x^{(H_c)\top} x^{(H_c)} = \|x^{(H_c)}\|_2^2,$$

where $x^{(H_c)}$ is the output of the H_c -th layer of the DNN Q_ω . Further combining Lemma A.4.4, it holds that

$$(A.5.38) \quad |Q_\omega| \leq 2.$$

Similarly, we have

$$(A.5.39) \quad |Q_{\omega(n)}| \leq 2.$$

Combining Lemma A.4.2, (A.5.36), (A.5.37), (A.5.38), and (A.5.39), we have

$$(A.5.40) \quad \mathbb{E}_{\pi_\theta} [\|g_n - g_n^e\|_2^2 | \omega(n)] = O(H_c^2).$$

Upper Bound of Term (iv).b. We now upper bound term (iv).b on the RHS of (A.5.35). It holds that

$$\begin{aligned}
& \|g_n^e - \bar{g}_n^e\|_2 \\
&= \left\| \mathbb{E}_{\pi_\theta} \left[\left(Q_{\omega(n)}(s_0, a_0) - \gamma \cdot Q_{\omega}(s_1, a_1) - (1 - \gamma) \cdot r_0 \right) \cdot \nabla_{\omega} Q_{\omega(n)}(s_0, a_0) \right. \right. \\
&\quad \left. \left. - \left(\bar{Q}_{\omega(n)}(s_0, a_0) - \gamma \cdot Q_{\omega}(s_1, a_1) - (1 - \gamma) \cdot r_0 \right) \cdot \nabla_{\omega} Q_{\omega_0}(s_0, a_0) \right] \right\|_2 \\
&\leq \mathbb{E}_{\pi_\theta} \left[\left\| \left(\gamma \cdot Q_{\omega}(s_1, a_1) + (1 - \gamma) \cdot r_t \right) \cdot \left(\nabla_{\omega} Q_{\omega_0}(s_0, a_0) - \nabla_{\omega} Q_{\omega(n)}(s_0, a_0) \right) \right\|_2 \right] \\
&\quad + \mathbb{E}_{\rho_{\pi_\theta}} \left[\left\| Q_{\omega(n)} \nabla_{\omega} Q_{\omega(n)} - \bar{Q}_{\omega(n)} \nabla_{\omega} Q_{\omega_0} \right\|_2 \right] \\
\text{(A.5.41)} \quad & \leq \mathbb{E}_{\pi_\theta} \left[\left\| \left(\gamma \cdot Q_{\omega}(s_1, a_1) + (1 - \gamma) \cdot r_0 \right) \cdot \left(\nabla_{\omega} Q_{\omega_0}(s_0, a_0) - \nabla_{\omega} Q_{\omega(n)}(s_0, a_0) \right) \right\|_2 \right] \\
&\quad + \mathbb{E}_{\rho_{\pi_\theta}} \left[\left\| \left(Q_{\omega(n)} - \bar{Q}_{\omega(n)} \right) \cdot \nabla_{\omega} Q_{\omega_0} \right\|_2 \right] + \mathbb{E}_{\rho_{\pi_\theta}} \left[\left\| Q_{\omega(n)} \cdot \left(\nabla_{\omega} Q_{\omega(n)} - \nabla_{\omega} Q_{\omega_0} \right) \right\|_2 \right].
\end{aligned}$$

We now upper bound the three terms on the RHS of (A.5.41) in the sequel, respectively.

For the term $\mathbb{E}_{\rho_{\pi_\theta}} \left[\left\| \left(Q_{\omega(n)} - \bar{Q}_{\omega(n)} \right) \cdot \nabla_{\omega} Q_{\omega_0} \right\|_2 \right]$ on the RHS of (A.5.41), following from Lemmas A.4.2 and A.4.3, it holds with probability at least $1 - \exp(-\Omega(R_c^{2/3} m_c^{2/3} H_c))$ that

$$\text{(A.5.42)} \quad \mathbb{E}_{\rho_{\pi_\theta}} \left[\left\| \left(Q_{\omega(n)} - \bar{Q}_{\omega(n)} \right) \cdot \nabla_{\omega} Q_{\omega_0} \right\|_2 \right] = O(R_c^{4/3} m_c^{-1/6} H_c^{7/2} (\log m_c)^{1/2}).$$

For the term $\mathbb{E}_{\rho_{\pi_\theta}} \left[\left\| Q_{\omega(n)} \cdot \left(\nabla_{\omega} Q_{\omega(n)} - \nabla_{\omega} Q_{\omega_0} \right) \right\|_2 \right]$ on the RHS of (A.5.41), following from (A.5.39) and Lemma A.4.2, with probability at least $1 - \exp(-\Omega(R_c^{2/3} m_c^{2/3} H_c))$, we have

$$\text{(A.5.43)} \quad \mathbb{E}_{\rho_{\pi_\theta}} \left[\left\| Q_{\omega(n)} \cdot \left(\nabla_{\omega} Q_{\omega(n)} - \nabla_{\omega} Q_{\omega_0} \right) \right\|_2 \right] = O(R_c^{1/3} m_c^{-1/6} H_c^{5/2} (\log m_c)^{1/2}).$$

For the term $\mathbb{E}_{\pi_\theta} [\|(\gamma \cdot Q_\omega(s_1, a_1) + (1 - \gamma) \cdot r_0) \cdot (\nabla_\omega Q_{\omega_0}(s_0, a_0) - \nabla_\omega Q_{\omega(n)}(s_0, a_0))\|_2]$ on the RHS of (A.5.41), we first upper bound $|\gamma \cdot Q_\omega(s_1, a_1) + (1 - \gamma) \cdot r_0|$ as follows,

$$|\gamma \cdot Q_\omega(s_1, a_1) + (1 - \gamma) \cdot r_0| \leq 2 + \mathcal{R}_{\max},$$

where we use (A.5.38) and the fact that $|\mathcal{R}(s, a)| \leq \mathcal{R}_{\max}$ for any $(s, a) \in \mathcal{S} \times \mathcal{A}$. Further combining Lemma A.4.2, with probability at least $1 - \exp(-\Omega(R_c^{2/3} m_c^{2/3} H_c))$, we have

$$\begin{aligned} & \mathbb{E}_{\pi_\theta} [\|(\gamma \cdot Q_\omega(s_1, a_1) + (1 - \gamma) \cdot r_0) \cdot (\nabla_\omega Q_{\omega_0}(s_0, a_0) - \nabla_\omega Q_{\omega(n)}(s_0, a_0))\|_2] \\ (A.5.44) \quad & = O(R_c^{1/3} m_c^{-1/6} H_c^{5/2} (\log m_c)^{1/2}). \end{aligned}$$

Now, combining (A.5.41), (A.5.42), (A.5.43), and (A.5.44), it holds with probability at least $1 - \exp(-\Omega(R_c^{2/3} m_c^{2/3} H_c))$ that

$$(A.5.45) \quad \|\bar{g}_n^e - \bar{g}_n^e\|_2^2 = O(R_c^{8/3} m_c^{-1/3} H_c^7 \log m_c).$$

Upper Bound of Term (iv).c. We now upper bound term (iv).c on the RHS of (A.5.35).

It holds that

$$\|\bar{g}_n^e - \bar{g}_*^e\|_2^2 = \|\mathbb{E}_{\rho_{\pi_\theta}} [(\bar{Q}_{\omega(n)} - \bar{Q}_{\omega_*}) \nabla_\omega Q_{\omega_0}]\|_2^2 \leq \mathbb{E}_{\rho_{\pi_\theta}} [(\bar{Q}_{\omega(n)} - \bar{Q}_{\omega_*})^2 \cdot \|\nabla_\omega Q_{\omega_0}\|_2^2].$$

Further combining Lemma A.4.2, it holds that

$$(A.5.46) \quad \mathbb{E}_{\pi_\theta} [\|\bar{g}_n^e - \bar{g}_*^e\|_2^2 \mid \omega(n)] \leq O(H_c^2) \cdot \mathbb{E}_{\rho_{\pi_\theta}} [(\bar{Q}_{\omega(n)} - \bar{Q}_{\omega_*})^2].$$

Combining (A.5.35), (A.5.40), (A.5.45), and (A.5.46), we obtain the following upper bound for term (iv) on the RHS of (A.5.31),

(A.5.47)

$$\mathbb{E}_{\pi_\theta} [\|g_n - \bar{g}_*^e\|_2^2 | \omega(n)] \leq O(R_c^{8/3} m_c^{-1/3} H_c^7 \log m_c) + O(H_c^2) \cdot \mathbb{E}_{\rho_{\pi_\theta}} [(\bar{Q}_{\omega(n)} - \bar{Q}_{\omega_*})^2].$$

We continue upper bounding (A.5.31). By plugging (A.5.34) and (A.5.47) into (A.5.31), it holds with probability at least $1 - \exp(-\Omega(R_c^{2/3} m_c^{2/3} H_c))$ that

$$\begin{aligned} & \mathbb{E}_{\pi_\theta} [\|\omega(n+1) - \omega_*\|_2^2 | \omega(n)] \\ & \leq \|\omega(n) - \omega_*\|_2^2 + 2\eta \cdot \left(O(R_c^{7/3} m_c^{-1/6} H_c^{7/2} (\log m_c)^{1/2}) - \mathbb{E}_{\rho_{\pi_\theta}} [(\bar{Q}_{\omega(n)} - \bar{Q}_{\omega_*})^2] \right) \\ & \quad + \eta^2 \cdot \left(O(R_c^{8/3} m_c^{-1/3} H_c^7 \log m_c) + O(H_c^2) \cdot \mathbb{E}_{\rho_{\pi_\theta}} [(\bar{Q}_{\omega(n)} - \bar{Q}_{\omega_*})^2] \right). \end{aligned}$$

(A.5.48)

Rearranging terms in (A.5.48), it holds with probability at least $1 - \exp(-\Omega(R_c^{2/3} m_c^{2/3} H_c))$ that

$$\begin{aligned} & (2\eta - \eta^2 \cdot O(H_c^2)) \cdot \mathbb{E}_{\rho_{\pi_\theta}} [(\bar{Q}_{\omega(n)} - \bar{Q}_{\omega_*})^2] \\ & \leq \|\omega(n) - \omega_*\|_2^2 - \mathbb{E}_{\rho_{\pi_\theta}} [\|\omega(n+1) - \omega_*\|_2^2 | \omega(n)] + \eta \cdot O(R_c^{8/3} m_c^{-1/3} H_c^7 \log m_c). \end{aligned}$$

(A.5.49)

By telescoping the sum and using Jensen's inequality in (A.5.49), we have

$$\begin{aligned}
\mathbb{E}_{\rho_{\pi_\theta}} [(\bar{Q}_{\bar{\omega}} - \bar{Q}_{\omega_*})^2] &\leq \frac{1}{N_c} \cdot \sum_{n=0}^{N_c-1} \mathbb{E}_{\rho_{\pi_\theta}} [(\bar{Q}_{\omega^{(n)}} - \bar{Q}_{\omega_*})^2] \\
&\leq 1/N_c \cdot (2\eta - \eta^2 \cdot O(H_c^2))^{-1} \cdot (\|\omega_0 - \omega_*\|_2^2 + \eta N_c \cdot O(R_c^{8/3} m_c^{-1/6} H_c^7 \log m_c)) \\
&\leq N_c^{-1/2} \cdot \|\theta_0 - \theta_*\|_2^2 + O(R_c^{8/3} m_c^{-1/6} H_c^7 \log m_c),
\end{aligned}$$

where the last line comes from the choices that $\eta = N_c^{-1/2}$ and $H_c = O(N_c^{1/4})$. Further combining Lemma A.4.3 and using triangle inequality, we have

$$(A.5.50) \quad \mathbb{E}_{\rho_{\pi_\theta}} [(Q_{\bar{\omega}} - \bar{Q}_{\omega_*})^2] = O(R_c^2 N_c^{-1/2} + R_c^{8/3} m_c^{-1/6} H_c^7 \log m_c).$$

To establish the upper bound of $\mathbb{E}_{\rho_{\pi_\theta}} [(\bar{Q}_{\omega_*} - \tilde{Q})^2]$, we upper bound $\mathbb{E}_{\rho_{\pi_\theta}} [(\bar{Q}_{\omega_*} - \tilde{Q})^2]$ in the sequel. By the definition of ω_* in (A.5.30), following a similar argument to derive (A.5.26), we have

$$(A.5.51) \quad \omega_* = \operatorname{argmin}_{\omega^\dagger \in \mathcal{B}(\omega_0, R_c)} \mathbb{E}_{\rho_{\pi_\theta}} [(\bar{Q}_{\omega^\dagger}(s_0, a_0) - \tilde{Q}(s_0, a_0))^2].$$

From the fact that $\tilde{Q} \in \mathcal{U}(m_c, H_c, R_c)$ by Assumption A.2.2, we know that $\tilde{Q} = Q_{\tilde{\omega}}$ for some $\tilde{\omega} \in \mathcal{B}(\omega_0, R_c)$. Therefore, by (A.5.51), with probability at least $1 - \exp(-\Omega(R_c^{2/3} m_c^{2/3} H_c))$, we have

$$(A.5.52) \quad \mathbb{E}_{\rho_{\pi_\theta}} [(\bar{Q}_{\omega_*} - \tilde{Q})^2] \leq \mathbb{E}_{\rho_{\pi_\theta}} [(\bar{Q}_{\tilde{\omega}} - \tilde{Q})^2] = O(R_c^{8/3} m_c^{-1/3} H_c^5 \log m_c),$$

where we use Lemma A.4.3 in the last inequality. Now, combining (A.5.50) and (A.5.52), by triangle inequality, with probability at least $1 - \exp(-\Omega(R_c^{2/3} m_c^{2/3} H_c))$, we have

$$\begin{aligned} \mathbb{E}_{\rho_{\pi_\theta}} [(Q_{\bar{\omega}} - \tilde{Q})^2] &\leq 2\mathbb{E}_{\rho_{\pi_\theta}} [(Q_{\bar{\omega}} - \bar{Q}_{\omega_*})^2] + 2\mathbb{E}_{\rho_{\pi_\theta}} [(\bar{Q}_{\omega_*} - \tilde{Q})^2] \\ &= O(R_c^2 N_c^{-1/2} + R_c^{8/3} m_c^{-1/6} H_c^7 \log m_c), \end{aligned}$$

which concludes the proof of Proposition A.2.4.

A.6. Proofs of Lemmas

A.6.1. Proof of Lemma 2.4.1

We denote by $\tilde{Q} = \mathcal{T}^{\pi_{\theta_k}} Q_{\omega_k}$. In the sequel, we upper bound $\mathbb{E}_{\rho_{k+1}} [(Q_{\omega_{k+1}} - Q_{\bar{\omega}_{k+1}})^2]$, where $\bar{\omega}_{k+1} = \Gamma_R(\tilde{\omega}_{k+1})$ and $\tilde{\omega}_{k+1}$ is defined in (2.2.4). Note that by the fact that $\|\varphi(s, a)\|_2 \leq 1$ uniformly, it suffices to upper bound $\|\omega_{k+1} - \tilde{\omega}_{k+1}\|_2$. By the definitions of ω_{k+1} and $\tilde{\omega}_{k+1}$ in (2.2.5) and (2.2.4), respectively, we have

$$(A.6.1) \quad \|\omega_{k+1} - \bar{\omega}_{k+1}\|_2 \leq \|\hat{\Phi}\hat{v} - \Phi v\|_2 \leq \|\Phi\|_2 \cdot \|\hat{v} - v\|_2 + \|\hat{\Phi} - \Phi\|_2 \cdot \|\hat{v}\|_2.$$

Here, we use the fact that the projection $\Gamma_R(\cdot)$ is a contraction in the first inequality, and triangle inequality in the second inequality. Also, for notational convenience, we denote

by $\widehat{\Phi}$, Φ , \widehat{v} , and v in (A.6.1) as follows,

$$\begin{aligned}\widehat{\Phi} &= \left(\frac{1}{N} \sum_{\ell=1}^N \varphi(s_{\ell,1}, a_{\ell,1}) \varphi(s_{\ell,1}, a_{\ell,1})^\top \right)^{-1}, \quad \Phi = \left(\mathbb{E}_{\rho_{k+1}} [\varphi(s, a) \varphi(s, a)^\top] \right)^{-1}, \\ \widehat{v} &= \frac{1}{N} \sum_{\ell=1}^N \left((1 - \gamma) r_{\ell,2} + \gamma Q_{\omega_k}(s'_{\ell,2}, a'_{\ell,2}) \right) \cdot \varphi(s_{\ell,2}, a_{\ell,2}), \\ v &= \mathbb{E}_{\rho_{k+1}} \left[\left((1 - \gamma) \mathcal{R} + \gamma \mathbb{P}^{\pi_{\theta_{k+1}}} Q_{\omega_k} \right) (s, a) \cdot \varphi(s, a) \right].\end{aligned}$$

By the fact that $\|\varphi(s, a)\|_2 \leq 1$, $|\mathcal{R}(s, a)| \leq \mathcal{R}_{\max}$, and $\|\omega_k\|_2 \leq R$ we have

$$(A.6.2) \quad \|\Phi\|_2 \leq 1/\sigma^*, \quad \|\widehat{v}\|_2 \leq \mathcal{R}_{\max} + R.$$

Now, following from matrix Bernstein inequality (Tropp, 2015) and Assumption 2.3.3, with probability at least $1 - p/2$, we have

$$(A.6.3) \quad \|\widehat{\Phi} - \Phi\|_2 \leq \frac{4}{\sqrt{N}(\sigma^*)^2} \cdot \log(N/p + d/p),$$

where σ^* is defined in Assumption 2.3.3. Similarly, with probability at least $1 - p/2$, we have

$$(A.6.4) \quad \|\widehat{v} - v\|_2 \leq 4(\mathcal{R}_{\max} + R)/\sqrt{N} \cdot \log(N/p + d/p).$$

Now, combining (A.6.1), (A.6.2), (A.6.3), and (A.6.4), we have

$$\|\omega_{k+1} - \bar{\omega}_{k+1}\|_2 \leq \frac{16(\mathcal{R}_{\max} + R)}{\sqrt{N}(\sigma^*)^2} \cdot \log(N/p + d/p).$$

Therefore, it holds with probability at least $1 - p$ that

$$(A.6.5) \quad (Q_{\omega_{k+1}} - Q_{\bar{\omega}_{k+1}})^2 \leq \frac{32(\mathcal{R}_{\max} + R)^2}{N(\sigma^*)^2} \cdot \log^2(N/p + d/p).$$

Meanwhile, by Assumption 2.3.2 and the definition of $\bar{\omega}_{k+1}$, we have

$$(A.6.6) \quad \tilde{Q}(s, a) = Q_{\bar{\omega}_{k+1}}(s, a)$$

for any $(s, a) \in \mathcal{S} \times \mathcal{A}$. Combining (A.6.5) and (A.6.6) and a union bound argument, with probability at least $1 - \delta$, it holds for any $k \in \{0, 1, \dots, K\}$ that

$$\mathbb{E}_{\rho_{k+1}} [(Q_{\omega_{k+1}}(s, a) - \tilde{Q}(s, a))^2] \leq \frac{32(\mathcal{R}_{\max} + R)^2}{N(\sigma^*)^4} \cdot \log^2(NK/p + dK/p),$$

which concludes the proof of Lemma 2.4.1.

A.6.2. Proof of Lemma A.3.1

Following from the definitions of \mathbb{P}^π and \mathbb{P} in (2.1.3), we have

$$(A.6.7) \quad A_{1,k}(s, a) = [\gamma(\mathbb{P}^{\pi^*} - \mathbb{P}^{\pi_{\theta_{k+1}}})Q_{\omega_k}](s, a) = [\gamma\mathbb{P}\langle Q_{\omega_k}, \pi^* - \pi_{\theta_{k+1}} \rangle](s, a).$$

By invoking Lemma A.4.1 and combining (A.6.7), it holds for any $(s, a) \in \mathcal{S} \times \mathcal{A}$ that

$$A_{1,k}(s, a) = [\gamma(\mathbb{P}^{\pi^*} - \mathbb{P}^{\pi_{\theta_{k+1}}})Q_{\omega_k}](s, a) \leq [\gamma\beta \cdot \mathbb{P}(\vartheta_k + \epsilon_{k+1}^a)](s, a),$$

where ϑ_k and ϵ_{k+1}^a are defined in (A.3.4) and (A.3.5) of Lemma A.3.1, respectively. We conclude the proof of Lemma A.3.1.

A.6.3. Proof of Lemma A.3.2

By the definition that Q^* is the action-value function of an optimal policy π^* , we know that $Q^*(s, a) \geq Q^\pi(s, a)$ for any policy π and state-action pair $(s, a) \in \mathcal{S} \times \mathcal{A}$. Therefore, for any $(s, a) \in \mathcal{S} \times \mathcal{A}$, we have

$$(A.6.8) \quad A_{2,k}(s, a) = [\gamma \mathbb{P}^{\pi^*}(Q^{\pi_{k+1}} - Q_{\omega_k})](s, a) \leq [\gamma \mathbb{P}^{\pi^*}(Q^* - Q_{\omega_k})](s, a).$$

In the sequel, we upper bound $Q^*(s, a) - Q_{\omega_k}(s, a)$ for any $(s, a) \in \mathcal{S} \times \mathcal{A}$. We define

$$\tilde{Q}_{k+1} = (1 - \gamma) \cdot \mathcal{R} + \gamma \cdot \mathbb{P}^{\pi_{k+1}} Q_{\omega_k}.$$

By its definition, we know that $\tilde{Q}_{k+1} = \mathcal{T}^{\pi_{k+1}} Q_{\omega_k}$. It holds for any $(s, a) \in \mathcal{S} \times \mathcal{A}$ that

$$\begin{aligned} & Q^*(s, a) - Q_{\omega_{k+1}}(s, a) \\ &= Q^*(s, a) - \tilde{Q}_{k+1}(s, a) + \tilde{Q}_{k+1}(s, a) - Q_{\omega_{k+1}}(s, a) \\ &= [((1 - \gamma) \cdot \mathcal{R} + \gamma \cdot \mathbb{P}^{\pi^*} Q^*) - ((1 - \gamma) \cdot \mathcal{R} + \gamma \cdot \mathbb{P}^{\pi_{k+1}} Q_{\omega_k})](s, a) + \epsilon_{k+1}^c(s, a) \\ &= \gamma \cdot [\mathbb{P}^{\pi^*} Q^* - \mathbb{P}^{\pi_{k+1}} Q_{\omega_k}](s, a) + \epsilon_{k+1}^c(s, a) \\ &= \gamma \cdot [\mathbb{P}^{\pi^*} Q^* - \mathbb{P}^{\pi^*} Q_{\omega_k}](s, a) + \gamma \cdot [\mathbb{P}^{\pi^*} Q_{\omega_k} - \mathbb{P}^{\pi_{k+1}} Q_{\omega_k}](s, a) + \epsilon_{k+1}^c(s, a) \\ &= \gamma \cdot [\mathbb{P}^{\pi^*}(Q^* - Q_{\omega_k})](s, a) + A_{1,k}(s, a) + \epsilon_{k+1}^c(s, a) \\ (A.6.9) \quad & \leq \gamma \cdot [\mathbb{P}^{\pi^*}(Q^* - Q_{\omega_k})](s, a) + \gamma\beta \cdot [\mathbb{P}(\vartheta_k + \epsilon_{k+1}^a)](s, a) + \epsilon_{k+1}^c(s, a), \end{aligned}$$

where ϵ_{k+1}^c and $A_{1,k}$ are defined in (A.3.6) and (A.3.3), respectively. Here, we use Lemma A.3.1 to upper bound $A_{1,k}$ in the last line. We remark that (A.6.9) upper bounds $Q^* - Q_{\omega_{k+1}}$ using $Q^* - Q_{\omega_k}$. By recursively applying a similar argument as in (A.6.9), we have

$$\begin{aligned}
& Q^*(s, a) - Q_{\omega_k}(s, a) \\
\text{(A.6.10)} \quad & \leq [(\gamma\mathbb{P}^{\pi^*})^k(Q^* - Q_{\omega_0})](s, a) + \gamma\beta \cdot \sum_{i=0}^{k-1} [(\gamma\mathbb{P}^{\pi^*})^{k-i-1}\mathbb{P}(\vartheta_i + \epsilon_{i+1}^a)](s, a) \\
& \quad + \sum_{i=0}^{k-1} [(\gamma\mathbb{P}^{\pi^*})^{k-i-1}\epsilon_{i+1}^c](s, a).
\end{aligned}$$

Combining (A.6.8) and (A.6.10), it holds for any $(s, a) \in \mathcal{S} \times \mathcal{A}$ that

$$\begin{aligned}
A_{2,k}(s, a) & \leq [\gamma\mathbb{P}^{\pi^*}(Q^* - Q_{\omega_k})](s, a) \\
& \leq [(\gamma\mathbb{P}^{\pi^*})^{k+1}(Q^* - Q_{\omega_0})](s, a) + \gamma\beta \cdot \sum_{i=0}^{k-1} [(\gamma\mathbb{P}^{\pi^*})^{k-i}\mathbb{P}(\vartheta_i + \epsilon_{i+1}^a)](s, a) \\
& \quad + \sum_{i=0}^{k-1} [(\gamma\mathbb{P}^{\pi^*})^{k-i}\epsilon_{i+1}^c](s, a),
\end{aligned}$$

where ϑ_i , ϵ_{i+1}^a , and ϵ_{i+1}^c are defined in (A.3.4) of Lemma A.3.1, (A.3.5) of Lemma A.3.1, and (A.3.6) of Lemma A.3.2, respectively. We conclude the proof of Lemma A.3.2.

A.6.4. Proof of Lemma A.3.3

Note that for any $(s, a) \in \mathcal{S} \times \mathcal{A}$, we have

$$\begin{aligned}
A_{3,k}(s, a) &= [\mathcal{T}^{\pi_{\theta_{k+1}}} Q_{\omega_k} - Q^{\pi_{\theta_{k+1}}}] (s, a) \\
&= \left[((1 - \gamma) \cdot \mathcal{R} + \gamma \mathbb{P}^{\pi_{\theta_{k+1}}} Q_{\omega_k}) - Q^{\pi_{\theta_{k+1}}} \right] (s, a) \\
&= \left[((1 - \gamma) \cdot \mathcal{R} + \gamma \mathbb{P}^{\pi_{\theta_{k+1}}} Q_{\omega_k}) - \sum_{t=0}^{\infty} (1 - \gamma) (\gamma \mathbb{P}^{\pi_{\theta_{k+1}}})^t \mathcal{R} \right] (s, a) \\
&= \left[\sum_{t=1}^{\infty} ((\gamma \mathbb{P}^{\pi_{\theta_{k+1}}})^t Q_{\omega_k} - (\gamma \mathbb{P}^{\pi_{\theta_{k+1}}})^{t+1} Q_{\omega_k}) - \sum_{t=1}^{\infty} (1 - \gamma) (\gamma \mathbb{P}^{\pi_{\theta_{k+1}}})^t \mathcal{R} \right] (s, a) \\
&= \sum_{t=1}^{\infty} \left[(\gamma \mathbb{P}^{\pi_{\theta_{k+1}}})^t (Q_{\omega_k} - \gamma \mathbb{P}^{\pi_{\theta_{k+1}}} Q_{\omega_k} - (1 - \gamma) \cdot \mathcal{R}) \right] (s, a) \\
&= \sum_{t=1}^{\infty} \left[(\gamma \mathbb{P}^{\pi_{\theta_{k+1}}})^t (Q_{\omega_k} - \mathcal{T}^{\pi_{\theta_{k+1}}} Q_{\omega_k}) \right] (s, a) \\
&= \sum_{t=1}^{\infty} \left[(\gamma \mathbb{P}^{\pi_{\theta_{k+1}}})^t e_{k+1} \right] (s, a) = [\gamma \mathbb{P}^{\pi_{\theta_{k+1}}} (I - \gamma \mathbb{P}^{\pi_{\theta_{k+1}}})^{-1} e_{k+1}] (s, a),
\end{aligned}$$

where the term e_{k+1} in the last line is defined in (A.3.7). We conclude the proof of Lemma A.3.3.

A.6.5. Proof of Lemma A.3.4

We invoke Lemma A.4.1 in §A.4, which gives

$$\begin{aligned}
& \beta^{-1} \cdot \langle Q_{\omega_k}(s, \cdot), \pi_{\theta_k}(\cdot | s) - \pi_{\theta_{k+1}}(\cdot | s) \rangle \\
& \leq \langle \log(\pi_{\theta_{k+1}}(\cdot | s) / \pi_{\theta_k}(\cdot | s)) - \beta^{-1} \cdot Q_{\omega_k}(s, \cdot), \pi_{\theta_k}(\cdot | s) - \pi_{\theta_{k+1}}(\cdot | s) \rangle \\
& \quad - \text{KL}(\pi_{\theta_k}(\cdot | s) \| \pi_{\theta_{k+1}}(\cdot | s)) \\
\text{(A.6.11)} \quad & \leq \langle \log(\pi_{\theta_{k+1}}(\cdot | s) / \pi_{\theta_k}(\cdot | s)) - \beta^{-1} \cdot Q_{\omega_k}(s, \cdot), \pi_{\theta_k}(\cdot | s) - \pi_{\theta_{k+1}}(\cdot | s) \rangle = \epsilon_{k+1}^b(s).
\end{aligned}$$

Combining (A.6.11) and the definition of \mathbb{P}^π in (2.1.3), we have

$$\text{(A.6.12)} \quad [\mathbb{P}^{\pi_{\theta_k}} Q_{\omega_k} - \mathbb{P}^{\pi_{\theta_{k+1}}} Q_{\omega_k}](s, a) \leq \beta [\mathbb{P} \epsilon_{k+1}^b](s).$$

By the definition of e_{k+1} in (A.3.7), we have

$$\begin{aligned}
e_{k+1}(s, a) &= [Q_{\omega_k} - \gamma \cdot \mathbb{P}^{\pi_{\theta_{k+1}}} Q_{\omega_k} - (1 - \gamma) \cdot \mathcal{R}](s, a) \\
\text{(A.6.13)} \quad &\leq [Q_{\omega_k} - \gamma \cdot \mathbb{P}^{\pi_{\theta_k}} Q_{\omega_k} - (1 - \gamma) \cdot \mathcal{R}](s, a) + \beta \gamma \cdot [\mathbb{P} \epsilon_{k+1}^b](s, a) \\
&= [\tilde{Q}_k - \gamma \cdot \mathbb{P}^{\pi_{\theta_k}} \tilde{Q}_k - (1 - \gamma) \cdot \mathcal{R}](s, a) + [\beta \gamma \mathbb{P} \epsilon_{k+1}^b - (I - \gamma \mathbb{P}^{\pi_{\theta_k}}) \epsilon_k^c](s, a),
\end{aligned}$$

where we use (A.6.12) in the first inequality, and

$$\text{(A.6.14)} \quad \tilde{Q}_k = (1 - \gamma) \cdot \mathcal{R} + \gamma \cdot \mathbb{P}^{\pi_{\theta_k}} Q_{\omega_{k-1}}.$$

For the first term on the RHS of (A.6.13), by (A.6.14), it holds that

$$\begin{aligned}
& \tilde{Q}_k - \gamma \cdot \mathbb{P}^{\pi_{\theta_k}} \tilde{Q}_k - (1 - \gamma) \cdot \mathcal{R} \\
&= (1 - \gamma) \cdot \mathcal{R} + \gamma \cdot \mathbb{P}^{\pi_{\theta_k}} Q_{\omega_{k-1}} - \gamma(1 - \gamma) \cdot \mathbb{P}^{\pi_{\theta_k}} \mathcal{R} - (\gamma \mathbb{P}^{\pi_{\theta_k}})^2 Q_{\omega_{k-1}} - (1 - \gamma) \cdot \mathcal{R} \\
\text{(A.6.15)} \quad &= \gamma \cdot \mathbb{P}^{\pi_{\theta_k}} (Q_{\omega_{k-1}} - \gamma \mathbb{P}^{\pi_{\theta_k}} Q_{\omega_{k-1}} - (1 - \gamma) \mathcal{R}) = \gamma \cdot \mathbb{P}^{\pi_{\theta_k}} e_k.
\end{aligned}$$

Combining (A.6.13) and (A.6.15), we have for any $(s, a) \in \mathcal{S} \times \mathcal{A}$ that

$$\text{(A.6.16)} \quad e_{k+1}(s, a) \leq [\gamma \mathbb{P}^{\pi_{\theta_k}} e_k](s, a) + [\beta \gamma \mathbb{P} \epsilon_{k+1}^b - (I - \gamma \mathbb{P}^{\pi_{\theta_k}}) \epsilon_k^c](s, a).$$

By telescoping (A.6.16), it holds that

$$e_{k+1}(s, a) \leq \left[\left(\prod_{s=1}^k \gamma \mathbb{P}^{\pi_{\theta_s}} \right) e_1 + \sum_{i=1}^k \gamma^{k-i} \left(\prod_{s=i+1}^k \mathbb{P}^{\pi_{\theta_s}} \right) (\beta \gamma \mathbb{P} \epsilon_{i+1}^b - (I - \gamma \mathbb{P}^{\pi_{\theta_i}}) \epsilon_i^c) \right](s, a).$$

This finishes the proof of the lemma.

A.6.6. Proof of Lemma A.3.5

Note that $\|\omega_0\|_2 \leq R$ and $|r(s, a)| \leq r_{\max}$ for any $(s, a) \in \mathcal{S} \times \mathcal{A}$, which implies that

$|Q_{\omega_0}(s, a)| \leq R$ and $|Q^*(s, a)| \leq r_{\max}$ by their definitions. Thus, for M_1 , we have

$$\begin{aligned}
\text{(A.6.17)} \quad |M_1| &\leq \mathbb{E}_\rho \left[(I - \gamma \mathbb{P}^{\pi^*})^{-1} \sum_{k=0}^K (\gamma \mathbb{P}^{\pi^*})^{k+1} |Q^* - Q_{\omega_0}| \right] \\
&\leq 4(1 - \gamma)^{-1} \sum_{k=0}^K \gamma^{k+1} \cdot (\mathcal{R}_{\max} + R) \leq 4(1 - \gamma)^{-2} \cdot (\mathcal{R}_{\max} + R).
\end{aligned}$$

For M_2 , by the definition of e_1 in (A.3.7), $|\omega_k| \leq R$, $|\phi(s, a)| \leq 1$, and $|r(s, a)| \leq r_{\max}$, we have

$$\begin{aligned}
|e_1(s, a)| &= |[Q_{\omega_k} - \mathcal{T}^{\pi_{\theta_{k+1}}} Q_{\omega_k}](s, a)| \\
&= |\omega_k^\top \phi(s, a) - \gamma \cdot \omega_k^\top [\mathbb{P}^{\pi_{\theta_{k+1}}} \phi](s, a) - (1 - \gamma) \cdot r(s, a)| \\
\text{(A.6.18)} \quad &\leq 2R + r_{\max}
\end{aligned}$$

for any $(s, a) \in \mathcal{S} \times \mathcal{A}$. Therefore, we have

$$\text{(A.6.19)} \quad |M_2| \leq (1 - \gamma)^{-3} \cdot (2R + \mathcal{R}_{\max}).$$

Meanwhile, by the initialization $\tau_0 = \infty$ in Algorithm 1, the initial policy $\pi_{\theta_0}(\cdot | s)$ is a uniform distribution over \mathcal{A} . Therefore, it holds for any $s \in \mathcal{S}$ that

$$\begin{aligned}
\text{KL}(\pi^*(\cdot | s) \parallel \pi_{\theta_0}(\cdot | s)) &= \int_{\mathcal{A}} \pi^*(a | s) \log \frac{\pi^*(a | s)}{\pi_{\theta_0}(a | s)} da \\
&= \int_{\mathcal{A}} \pi^*(a | s) \log \pi^*(a | s) da - \int_{\mathcal{A}} \pi^*(a | s) \log \pi_{\theta_0}(a | s) da \\
&\leq - \int_{\mathcal{A}} \pi^*(a | s) \log \pi_{\theta_0}(a | s) da \\
\text{(A.6.20)} \quad &= \int_{\mathcal{A}} \pi^*(a | s) \log |\mathcal{A}| da = \log |\mathcal{A}|.
\end{aligned}$$

Therefore, by (A.6.20), we have

$$\text{(A.6.21)} \quad M_3 \leq (1 - \gamma)^{-2} \cdot \log |\mathcal{A}| \cdot K^{1/2},$$

where we use $\beta = K^{1/2}$. We see that (A.6.17), (A.6.19), and (A.6.21) upper bound M_1 , M_2 , and M_3 , respectively. We conclude the proof of Lemma A.3.5.

A.6.7. Proof of Lemma A.3.6

For M_4 , by changing the index of summation, we have

$$\begin{aligned}
|M_4| &= \left| \mathbb{E}_\rho \left[\sum_{k=0}^K \sum_{i=0}^k \sum_{j=0}^{\infty} (\gamma \mathbb{P}^{\pi^*})^{k-i+j} \epsilon_{i+1}^c \right] \right| \\
&= \left| \mathbb{E}_\rho \left[\sum_{k=0}^K \sum_{i=0}^k \sum_{t=k-i}^{\infty} (\gamma \mathbb{P}^{\pi^*})^t \epsilon_{i+1}^c \right] \right| \\
(A.6.22) \quad &\leq \sum_{k=0}^K \sum_{i=0}^k \sum_{t=k-i}^{\infty} \left| \mathbb{E}_\rho [(\gamma \mathbb{P}^{\pi^*})^t \epsilon_{i+1}^c] \right|,
\end{aligned}$$

where we expand $(I - \gamma \mathbb{P}^{\pi^*})^{-1}$ into an infinite sum in the first equality. Further, by changing the measure of the expectation from ρ to ρ^* on the RHS of (A.6.22), we have

$$(A.6.23) \quad \sum_{k=0}^K \sum_{i=0}^k \sum_{t=k-i}^{\infty} \left| \mathbb{E}_\rho [(\gamma \mathbb{P}^{\pi^*})^t \epsilon_{i+1}^c] \right| \leq \sum_{k=0}^K \sum_{i=0}^k \sum_{t=k-i}^{\infty} \gamma^t c(t) \cdot \mathbb{E}_{\rho^*} [|\epsilon_{i+1}^c|],$$

where $c(t)$ is defined in Assumption 2.3.1. Further, by changing the index of summation on the RHS of (A.6.23), combining (A.6.22), we have

$$\begin{aligned}
|M_4| &\leq \sum_{k=0}^K \sum_{t=0}^{\infty} \sum_{i=\max\{0, k-t\}}^k \gamma^t c(t) \cdot \varepsilon_Q \\
&\leq \sum_{k=0}^K \sum_{t=0}^{\infty} 2t \gamma^t c(t) \cdot \varepsilon_Q \\
(A.6.24) \quad &\leq \gamma \sum_{k=0}^K 2C_{\rho, \rho^*} \cdot \varepsilon_Q \leq 3KC_{\rho, \rho^*} \cdot \varepsilon_Q,
\end{aligned}$$

where $\varepsilon_Q = \max_i \mathbb{E}_{\rho^*} [|\epsilon_{i+1}^c|]$, and C_{ρ, ρ^*} is defined in Assumption 2.3.1.

Now, for M_5 , by a similar argument as in the derivation of (A.6.24), we have

$$\begin{aligned}
 M_5 &\leq \sum_{i=0}^{\infty} \sum_{k=0}^K \sum_{j=0}^{\infty} \sum_{\ell=1}^k \gamma^{i+j+k-\ell+1} c(i+j+k-\ell+1) \cdot \varepsilon_Q \\
 (A.6.25) \quad &= \sum_{i=0}^{\infty} \sum_{k=0}^K \sum_{j=0}^{\infty} \sum_{t=i+j+1}^{i+j+k} \gamma^t c(t) \cdot \varepsilon_Q \leq \sum_{k=0}^K \sum_{t=1}^{\infty} t^2 \gamma^t c(t) \cdot \varepsilon_Q \leq K C_{\rho, \rho^*} \cdot \varepsilon_Q.
 \end{aligned}$$

We see that (A.6.24) and (A.6.25) upper bound M_4 and M_5 , respectively. We conclude the proof of Lemma A.3.6.

A.6.8. Proof of Lemma A.3.7

Part 1. We first show that the first inequality holds. Note that

$$\pi_{\theta_k}(a | s) = \exp(\tau_k^{-1} f_{\theta_k}(s, a)) / Z_{\theta_k}(s), \quad \pi_{\theta_{k+1}}(a | s) = \exp(\tau_{k+1}^{-1} f_{\theta_{k+1}}(s, a)) / Z_{\theta_{k+1}}(s),$$

Here $Z_{\theta_k}(s), Z_{\theta_{k+1}}(s) \in \mathbb{R}$ are normalization factors, which are defined as

$$Z_{\theta_k}(s) = \sum_{a' \in \mathcal{A}} \exp(\tau_k^{-1} f_{\theta_k}(s, a')), \quad Z_{\theta_{k+1}}(s) = \sum_{a' \in \mathcal{A}} \exp(\tau_{k+1}^{-1} f_{\theta_{k+1}}(s, a')).$$

Thus, we have

$$\begin{aligned}
 &\langle \log(\pi_{\theta_{k+1}}(\cdot | s) / \pi_{\theta_k}(\cdot | s)) - \beta^{-1} Q_{\omega_k}(s, \cdot), \pi^*(\cdot | s) - \pi_{\theta_{k+1}}(\cdot | s) \rangle \\
 (A.6.26) \quad &= \langle \tau_{k+1}^{-1} f_{\theta_{k+1}}(s, \cdot) - (\beta^{-1} Q_{\omega_k}(s, \cdot) + \tau_k^{-1} f_{\theta_k}(s, \cdot)), \pi^*(\cdot | s) - \pi_{\theta_k}(\cdot | s) \rangle,
 \end{aligned}$$

where we use the fact that

$$\begin{aligned} & \langle \log Z_{\theta_{k+1}}(s) - \log Z_{\theta_k}(s), \pi^*(\cdot | s) - \pi_{\theta_{k+1}}(\cdot | s) \rangle \\ &= (\log Z_{\theta_{k+1}}(s) - \log Z_{\theta_k}(s)) \cdot \sum_{a' \in \mathcal{A}} (\pi^*(a' | s) - \pi_{\theta_{k+1}}(a' | s)) = 0. \end{aligned}$$

Thus, it remains to upper bound the right-hand side of (A.6.26). We have

(A.6.27)

$$\begin{aligned} & \langle \tau_{k+1}^{-1} f_{\theta_{k+1}}(s, \cdot) - (\beta_k^{-1} Q_{\omega_k}(s, \cdot) + \tau_k^{-1} f_{\theta_k}(s, \cdot)), \pi^*(\cdot | s) - \pi_{\theta_{k+1}}(\cdot | s) \rangle \\ &= \left\langle \tau_{k+1}^{-1} f_{\theta_{k+1}}(s, \cdot) - (\beta_k^{-1} Q_{\omega_k}(s, \cdot) + \tau_k^{-1} f_{\theta_k}(s, \cdot)), \pi_{\theta_k}(\cdot | s) \cdot \left(\frac{\pi^*(\cdot | s)}{\pi_{\theta_k}(\cdot | s)} - \frac{\pi_{\theta_{k+1}}(\cdot | s)}{\pi_{\theta_k}(\cdot | s)} \right) \right\rangle. \end{aligned}$$

Taking expectation with respect to $s \sim \nu^*$ on the both sides of (A.6.27) and using the Cauchy-Schwarz inequality, we obtain

$$\begin{aligned} & \mathbb{E}_{\nu^*} \left[\left| \langle \tau_{k+1}^{-1} f_{\theta_{k+1}}(s, \cdot) - (\beta_k^{-1} Q_{\omega_k}(s, \cdot) + \tau_k^{-1} f_{\theta_k}(s, \cdot)), \pi^*(\cdot | s) - \pi_{\theta_{k+1}}(\cdot | s) \rangle \right| \right] \\ &= \int_{\mathcal{S}} \left| \left\langle \tau_{k+1}^{-1} f_{\theta_{k+1}}(s, \cdot) - (\beta_k^{-1} Q_{\omega_k}(s, \cdot) + \tau_k^{-1} f_{\theta_k}(s, \cdot)), \right. \right. \\ & \quad \left. \left. \pi_{\theta_k}(\cdot | s) \cdot \nu_k(s) \cdot \left(\frac{\pi^*(\cdot | s)}{\pi_{\theta_k}(\cdot | s)} - \frac{\pi_{\theta_{k+1}}(\cdot | s)}{\pi_{\theta_k}(\cdot | s)} \right) \right\rangle \right| \cdot \left| \frac{\nu^*(s)}{\nu_k(s)} \right| ds \\ &= \int_{\mathcal{S} \times \mathcal{A}} \left| \tau_{k+1}^{-1} f_{\theta_{k+1}}(s, a) - (\beta_k^{-1} Q_{\omega_k}(s, a) + \tau_k^{-1} f_{\theta_k}(s, a)) \right| \\ & \quad \cdot \left| \frac{\rho^*(a | s)}{\rho_k(a | s)} - \frac{\pi_{\theta_{k+1}}(a | s) \cdot \nu^*(s)}{\rho_k(a | s)} \right| d\rho_k(s, a) \\ &\leq \sqrt{\mathbb{E}_{\rho_k} \left[(\tau_{k+1}^{-1} f_{\theta_{k+1}}(s, a) - (\beta_k^{-1} Q_{\omega_k}(s, a) + \tau_k^{-1} f_{\theta_k}(s, a)))^2 \right] \mathbb{E}_{\rho_k} \left[\left| \frac{d\rho^*}{d\rho_k} - \frac{d(\pi_{\theta_{k+1}} \nu^*)}{d\rho_k} \right|^2 \right]} \\ &\leq \sqrt{2} \tau_{k+1}^{-1} \cdot \varepsilon_{k+1, f} \cdot (\phi_k^* + \psi_k^*), \end{aligned}$$

where in the last inequality we use the error bound in (A.3.20) and the definition of ϕ_k^* and ψ_k^* in Assumption A.2.1. This finishes the proof of the first inequality.

Part 2. The proof of the second inequality follows from a similar argument as above. We have

$$\begin{aligned}
 & \langle \log(\pi_{\theta_{k+1}}(\cdot | s) / \pi_{\theta_k}(\cdot | s)) - \beta^{-1} Q_{\omega_k}(s, \cdot), \pi_{\theta_k}(\cdot | s) - \pi_{\theta_{k+1}}(\cdot | s) \rangle \\
 \text{(A.6.28)} \quad & = \langle \tau_{k+1}^{-1} f_{\theta_{k+1}}(s, \cdot) - (\beta^{-1} Q_{\omega_k}(s, \cdot) + \tau_k^{-1} f_{\theta_k}(s, \cdot)), \pi_{\theta_k}(\cdot | s) - \pi_{\theta_{k+1}}(\cdot | s) \rangle,
 \end{aligned}$$

where we use the fact that

$$\begin{aligned}
 & \langle \log Z_{\theta_{k+1}}(s) - \log Z_{\theta_k}(s), \pi_{\theta_k}(\cdot | s) - \pi_{\theta_{k+1}}(\cdot | s) \rangle \\
 & = (\log Z_{\theta_{k+1}}(s) - \log Z_{\theta_k}(s)) \cdot \sum_{a' \in \mathcal{A}} (\pi_{\theta_k}(a' | s) - \pi_{\theta_{k+1}}(a' | s)) = 0.
 \end{aligned}$$

Thus, it remains to upper bound the right-hand side of (A.6.28). We have

$$\begin{aligned}
 \text{(A.6.29)} \quad & \langle \tau_{k+1}^{-1} f_{\theta_{k+1}}(s, \cdot) - (\beta_k^{-1} Q_{\omega_k}(s, \cdot) + \tau_k^{-1} f_{\theta_k}(s, \cdot)), \pi_{\theta_k}(\cdot | s) - \pi_{\theta_{k+1}}(\cdot | s) \rangle \\
 & = \left\langle \tau_{k+1}^{-1} f_{\theta_{k+1}}(s, \cdot) - (\beta_k^{-1} Q_{\omega_k}(s, \cdot) + \tau_k^{-1} f_{\theta_k}(s, \cdot)), \pi_{\theta_k}(\cdot | s) \cdot \left(1 - \frac{\pi_{\theta_{k+1}}(\cdot | s)}{\pi_{\theta_k}(\cdot | s)}\right) \right\rangle.
 \end{aligned}$$

Taking expectation with respect to $s \sim \nu^*$ on the both sides of (A.6.29) and using the Cauchy-Schwarz inequality, we obtain

$$\begin{aligned}
& \mathbb{E}_{\nu^*} \left[\left| \left\langle \tau_{k+1}^{-1} f_{\theta_{k+1}}(s, \cdot) - (\beta_k^{-1} Q_{\omega_k}(s, \cdot) + \tau_k^{-1} f_{\theta_k}(s, \cdot)), \pi_{\theta_k}(\cdot | s) - \pi_{\theta_{k+1}}(\cdot | s) \right\rangle \right| \right] \\
&= \int_{\mathcal{S}} \left| \left\langle \tau_{k+1}^{-1} f_{\theta_{k+1}}(s, \cdot) - (\beta_k^{-1} Q_{\omega_k}(s, \cdot) + \tau_k^{-1} f_{\theta_k}(s, \cdot)), \pi_{\theta_k}(\cdot | s) \nu_k(s) \left(1 - \frac{\pi_{\theta_{k+1}}(\cdot | s)}{\pi_{\theta_k}(\cdot | s)} \right) \right\rangle \right| \\
&\quad \cdot \left| \frac{\nu^*(s)}{\nu_k(s)} \right| ds \\
&= \int_{\mathcal{S} \times \mathcal{A}} \left| \tau_{k+1}^{-1} f_{\theta_{k+1}}(s, a) - (\beta_k^{-1} Q_{\omega_k}(s, a) + \tau_k^{-1} f_{\theta_k}(s, a)) \right| \left| 1 - \frac{\pi_{\theta_{k+1}}(a | s) \cdot \nu^*(s)}{\rho_k(s, a)} \right| d\rho_k \\
&\leq \mathbb{E}_{\rho_k} \left[(\tau_{k+1}^{-1} f_{\theta_{k+1}}(s, a) - (\beta_k^{-1} Q_{\omega_k}(s, a) + \tau_k^{-1} f_{\theta_k}(s, a)))^2 \right]^{1/2} \mathbb{E}_{\rho_k} \left[\left| 1 - \frac{d(\pi_{\theta_{k+1}} \nu^*)}{d\rho_k} \right|^2 \right]^{1/2} \\
&\leq \sqrt{2} \tau_{k+1}^{-1} \cdot \varepsilon_{k+1, f} \cdot (1 + \psi_k^*),
\end{aligned}$$

where in the last inequality we use the error bound in (A.3.20) and the definition of ψ_k^* in Assumption A.2.1. This finishes the proof of the second inequality.

APPENDIX B

Supplemental Materials in Chapter 3**B.1. Notations in the Appendix**

In the proof, for convenience, for any invertible matrix M , we denote by $M^{-\top} = (M^{-1})^\top = (M^\top)^{-1}$ and $\|M\|_F$ the Frobenius norm. We also denote by $\text{svec}(M)$ the symmetric vectorization of the symmetric matrix M , which is the vectorization of the upper triangular matrix of the symmetric matrix M , with off-diagonal entries scaled by $\sqrt{2}$. We denote by $\text{smat}(\cdot)$ the inverse operation. For any matrices G and H , we denote by $G \otimes H$ the Kronecker product, and $G \otimes_s H$ the symmetric Kronecker product, which is defined as a mapping on a vector $\text{svec}(M)$ such that $(G \otimes_s H)\text{svec}(M) = 1/2 \cdot \text{svec}(HMG^\top + GMH^\top)$.

For notational simplicity, we write $\mathbb{E}_\pi(\cdot)$ to emphasize that the expectation is taken following the policy π .

B.2. Auxiliary Algorithms and Analysis**B.2.1. Results in D-LQR**

In this section, we provide auxiliary results in analyzing Problem 3.1.2. First, we introduce the value functions of the Markov decision process (MDP) induced by Problem 3.1.2. We

define the state- and action-value functions $V_{K,b}(x)$ and $Q_{K,b}(x, u)$ as follows

$$(B.2.1) \quad V_{K,b}(x) = \sum_{t=0}^{\infty} \left\{ \mathbb{E}[c(x_t, u_t) \mid x_0 = x] - J(K, b) \right\},$$

$$(B.2.2) \quad Q_{K,b}(x, u) = c(x, u) - J(K, b) + \mathbb{E}[V_{K,b}(x_1) \mid x_0 = x, u_0 = u],$$

where x_t follows the state transition, and u_t follows the policy $\pi_{K,b}$ given x_t . In other words, we have $u_t = -Kx_t + b + \sigma\eta_t$, where $\eta_t \sim \mathcal{N}(0, I)$. The following proposition establishes the close forms of these value functions.

Proposition B.2.1. The state-value function $V_{K,b}(x)$ takes the form of

$$(B.2.3) \quad V_{K,b}(x) = x^\top P_K x - \text{tr}(P_K \Phi_K) + 2f_{K,b}^\top (x - \mu_{K,b}) - \mu_{K,b}^\top P_K \mu_{K,b},$$

and the action-value function $Q_{K,b}(x, u)$ takes the form of

$$(B.2.4) \quad \begin{aligned} Q_{K,b}(x, u) = & \begin{pmatrix} x \\ u \end{pmatrix}^\top \Upsilon_K \begin{pmatrix} x \\ u \end{pmatrix} + 2 \begin{pmatrix} p_{K,b} \\ q_{K,b} \end{pmatrix}^\top \begin{pmatrix} x \\ u \end{pmatrix} - \text{tr}(P_K \Phi_K) - \sigma^2 \cdot \text{tr}(R + P_K B B^\top) \\ & - b^\top R b + 2b^\top R K \mu_{K,b} - \mu_{K,b}^\top (Q + K^\top R K + P_K) \mu_{K,b} \\ & + 2f_{K,b}^\top [(\bar{A}\mu + d) - \mu_{K,b}] + (\bar{A}\mu + d)^\top P_K (\bar{A}\mu + d), \end{aligned}$$

where $f_{K,b} = (I - A + BK)^{-\top} [(A - BK)^\top P_K (Bb + \bar{A}\mu + d) - K^\top R b]$, and Υ_K , $p_{K,b}$, and $q_{K,b}$ are defined in (3.2.7).

Proof. See §B.5.6 for a detailed proof. \square

By Proposition B.2.1, we know that $V_{K,b}(x)$ is quadratic in x , while $Q_{K,b}(x, u)$ is quadratic in $(x^\top, u^\top)^\top$. Now, we show that (3.2.5) holds.

Proposition B.2.2. The expected total cost $J(K, b)$ defined in Problem 3.1.2 takes the form of

$$J(K, b) = J_1(K) + J_2(K, b) + \sigma^2 \cdot \text{tr}(R) + \mu^\top \bar{Q} \mu,$$

where

$$J_1(K) = \text{tr}[(Q + K^\top RK)\Phi_K] = \text{tr}(P_K \Psi_\epsilon),$$

$$J_2(K, b) = \begin{pmatrix} \mu_{K,b} \\ b \end{pmatrix}^\top \begin{pmatrix} Q + K^\top RK & -K^\top R \\ -RK & R \end{pmatrix} \begin{pmatrix} \mu_{K,b} \\ b \end{pmatrix}.$$

Here $\mu_{K,b}$ is defined in (3.2.2), Φ_K is defined in (3.2.3), and P_K is defined in (3.2.4).

Proof. See §B.5.3 for a detailed proof. □

The following proposition establishes the gradients of $J_1(K)$ and $J_2(K, b)$, respectively.

Proposition B.2.3. The gradient of $J_1(K)$ and the gradient of $J_2(K, b)$ with respect to b take the forms of

$$\nabla_K J_1(K) = 2(\Upsilon_K^{22} K - \Upsilon_K^{21}) \cdot \Phi_K, \quad \nabla_b J_2(K, b) = 2[\Upsilon_K^{22}(-K\mu_{K,b} + b) + \Upsilon_K^{21}\mu_{K,b} + q_{K,b}],$$

where Υ_K and $q_{K,b}$ are defined in (3.2.7).

Proof. See §B.5.5 for a detailed proof. □

The following theorem establishes the convergence of Algorithm 3.

Theorem B.2.4 (Convergence of Algorithm 3). Assume that $\rho(A - BK_0) < 1$. Let $\varepsilon > 0$ be a sufficiently small tolerance. We set

$$\begin{aligned} \gamma &\leq [\|R\|_2 + \|B\|_2^2 \cdot J(K_0, b_0) \cdot \sigma_{\min}^{-1}(\Psi_\varepsilon)]^{-1}, \\ N &\geq C \cdot \|\Phi_{K^*}\|_2 \cdot \gamma^{-1} \cdot \log\left\{4[J(K_0, b_0) - J(K^*, b^*)] \cdot \varepsilon^{-1}\right\}, \\ T_n &\geq \text{poly}(\|K_n\|_F, \|b_0\|_2, \|\mu\|_2, J(K_0, b_0)) \cdot \lambda_{K_n}^{-4} \cdot [1 - \rho(A - BK_n)]^{-9} \cdot \varepsilon^{-5}, \\ \tilde{T}_n &\geq \text{poly}(\|K_n\|_F, \|b_0\|_2, \|\mu\|_2, J(K_0, b_0)) \cdot \lambda_{K_n}^{-2} \cdot [1 - \rho(A - BK_n)]^{-12} \cdot \varepsilon^{-12}, \\ \gamma_{n,t} &= \gamma_0 \cdot t^{-1/2}, \\ \gamma^b &\leq \min\left\{1 - \rho(A - BK_N), \right. \\ &\quad \left.[1 - \rho(A - BK_N)]^{-2} \cdot (\|B\|_2^2 \cdot \|K_N\|_2^2 \cdot \|R\|_2 + \|B\|_2^2 \cdot \|Q\|_2)\right\}, \\ H &\geq C_0 \cdot \nu_{K_N}^{-1} \cdot (\gamma^b)^{-1} \cdot \log\left\{4[J(K_N, b_0) - J(K_N, b^{K_N})] \cdot \varepsilon^{-1}\right\}, \\ T_h^b &\geq \text{poly}(\|K_N\|_F, \|b_h\|_2, \|\mu\|_2, J(K_N, b_0)) \cdot \lambda_{K_N}^{-4} \cdot \nu_{K_N}^{-4} \cdot [1 - \rho(A - BK_N)]^{-11} \cdot \varepsilon^{-5}, \\ \tilde{T}_h^b &\geq \text{poly}(\|K_N\|_F, \|b_h\|_2, \|\mu\|_2, J(K_N, b_0)) \cdot \lambda_{K_N}^{-4} \cdot \nu_{K_N}^{-2} \cdot [1 - \rho(A - BK_N)]^{-17} \cdot \varepsilon^{-8}, \\ \gamma_{h,t}^b &= \gamma_0 \cdot t^{-1/2}, \end{aligned}$$

where C , C_0 , and γ_0 are positive absolute constants, $\{K_n\}_{n \in [N]}$ and $\{b_h\}_{h \in [H]}$ are the sequences generated by Algorithm 3, λ_{K_n} is specified in Proposition B.2.6, and ν_{K_N} is specified in Proposition 3.2.3. Then it holds with probability at least $1 - \varepsilon^{10}$ that

$$\begin{aligned} J(K_N, b_H) - J(K^*, b^*) &< \varepsilon, \quad \|b_H - b^*\|_2 \leq M_b(\mu) \cdot \varepsilon^{1/2}, \\ \|K_N - K^*\|_F &\leq [\sigma_{\min}^{-1}(\Psi_\varepsilon) \cdot \sigma_{\min}^{-1}(R) \cdot \varepsilon]^{1/2}, \quad \|\hat{\mu}_{K_N, b_H} - \mu_{K^*, b^*}\|_2 \leq \varepsilon, \end{aligned}$$

where $M_b(\mu)$ is defined in (3.3.3).

Proof. See §B.4.2 for a detailed proof. □

By Theorem B.2.4, given any mean-field state μ , Algorithm 3 converges linearly to the optimal policy π_μ^* of Problem 3.1.2.

B.2.2. Primal-Dual Policy Evaluation Algorithm

Note that the critic update steps in Algorithm 3 are built upon the estimators of the matrix Υ_K and the vector $q_{K,b}$. We now derive a policy evaluation algorithm to establish the estimators of Υ_K and $q_{K,b}$, which is based on gradient temporal difference algorithm (Sutton et al., 2009a).

We define the feature vector as

$$(B.2.5) \quad \psi(x, u) = \begin{pmatrix} \varphi(x, u) \\ x - \mu_{K,b} \\ u - (-K\mu_{K,b} + b) \end{pmatrix},$$

where

$$\varphi(x, u) = \text{svec} \left[\begin{pmatrix} x - \mu_{K,b} \\ u - (-K\mu_{K,b} + b) \end{pmatrix} \begin{pmatrix} x - \mu_{K,b} \\ u - (-K\mu_{K,b} + b) \end{pmatrix}^\top \right].$$

Recall $\text{svec}(M)$ gives the symmetric vectorization of the symmetric matrix M . We also define

$$(B.2.6) \quad \alpha_{K,b} = \begin{pmatrix} \text{svec}(\Upsilon_K) \\ \Upsilon_K \begin{pmatrix} \mu_{K,b} \\ -K\mu_{K,b} + b \end{pmatrix} + \begin{pmatrix} p_{K,b} \\ q_{K,b} \end{pmatrix} \end{pmatrix},$$

where Υ_K , $p_{K,b}$, and $q_{K,b}$ are defined in (3.2.7). To estimate Υ_K and $q_{K,b}$, it suffices to estimate $\alpha_{K,b}$. Meanwhile, we define

$$(B.2.7) \quad \Theta_{K,b} = \mathbb{E}_{\pi_{K,b}} \left\{ \psi(x, u) [\psi(x, u) - \psi(x', u')]^\top \right\},$$

where (x', u') is the state-action pair after (x, u) following the policy $\pi_{K,b}$ and the state transition. The following proposition characterizes the connection between $\Theta_{K,b}$ and $\alpha_{K,b}$.

Proposition B.2.5. It holds that

$$\begin{pmatrix} 1 & 0 \\ \mathbb{E}_{\pi_{K,b}}[\psi(x, u)] & \Theta_{K,b} \end{pmatrix} \begin{pmatrix} J(K, b) \\ \alpha_{K,b} \end{pmatrix} = \begin{pmatrix} J(K, b) \\ \mathbb{E}_{\pi_{K,b}}[c(x, u)\psi(x, u)] \end{pmatrix},$$

where $\psi(x, u)$ is defined in (B.2.5), $\alpha_{K,b}$ is defined in (B.2.6), and $\Theta_{K,b}$ is defined in (B.2.7).

Proof. See §B.5.7 for a detailed proof. \square

By Proposition B.2.5, to obtain $\alpha_{K,b}$, it suffices to solve the following linear system in $\zeta = (\zeta_1, \zeta_2^\top)^\top$,

$$(B.2.8) \quad \tilde{\Theta}_{K,b} \cdot \zeta = \begin{pmatrix} J(K, b) \\ \mathbb{E}_{\pi_{K,b}}[c(x, u)\psi(x, u)] \end{pmatrix},$$

where for notational convenience, we define

$$(B.2.9) \quad \tilde{\Theta}_{K,b} = \begin{pmatrix} 1 & 0 \\ \mathbb{E}_{\pi_{K,b}}[\psi(x, u)] & \Theta_{K,b} \end{pmatrix}.$$

The following proposition shows that $\Theta_{K,b}$ is invertible.

Proposition B.2.6. If $\rho(A - BK) < 1$, then the matrix $\Theta_{K,b}$ is invertible, and $\|\Theta_{K,b}\|_2 \leq 4(1 + \|K\|_{\mathbb{F}}^2) \cdot \|\Phi_K\|_2^2$. Also, $\sigma_{\min}(\tilde{\Theta}_{K,b}) \geq \lambda_K$, where λ_K only depends on $\|K\|_2$ and $\rho(A - BK)$.

Proof. See §B.5.8 for a detailed proof. □

By Proposition B.2.6, $\Theta_{K,b}$ is invertible. Therefore, (B.2.8) admits the unique solution $\zeta_{K,b} = (J(K, b), \alpha_{K,b}^\top)^\top$.

Now, we present the primal-dual gradient temporal difference algorithm.

Primal-Dual Gradient Method. Instead of solving (B.2.8) directly, we minimize the following loss function with respect to $\zeta = ((\zeta^1)^\top, (\zeta^2)^\top)$,

$$(B.2.10) \quad [\zeta^1 - J(K, b)]^2 + \left\| \mathbb{E}_{\pi_{K,b}}[\psi(x, u)]\zeta^1 + \Theta_{K,b}\zeta^2 - \mathbb{E}_{\pi_{K,b}}[c(x, u)\psi(x, u)] \right\|_2^2.$$

By Fenchel's duality, the minimization of (B.2.10) is equivalent to the following primal-dual min-max problem,

$$(B.2.11) \quad \min_{\zeta \in \mathcal{V}_\zeta} \max_{\xi \in \mathcal{V}_\xi} F(\zeta, \xi) = \left\{ \mathbb{E}_{\pi_{K,b}}[\psi(x, u)]\zeta^1 + \Theta_{K,b}\zeta^2 - \mathbb{E}_{\pi_{K,b}}[c(x, u)\psi(x, u)] \right\}^\top \xi^2 \\ + [\zeta^1 - J(K, b)] \cdot \xi^1 - \|\xi\|_2^2/2,$$

where we restrict the primal variable ζ in a compact set \mathcal{V}_ζ and the dual variable ξ in a compact set \mathcal{V}_ξ , which are specified in Definition B.2.7. It holds that

$$\nabla_{\zeta^1} F = \xi^1 + \mathbb{E}_{\pi_{K,b}}[\psi(x, u)]^\top \xi^2, \quad \nabla_{\zeta^2} F = \Theta_{K,b}^\top \xi^2, \quad \nabla_{\xi^1} F = \zeta^1 - J(K, b) - \xi^1, \quad (\text{B.2.12})$$

$$\nabla_{\xi^2} F = \mathbb{E}_{\pi_{K,b}}[\psi(x, u)] \zeta^1 + \Theta_{K,b} \zeta^2 - \mathbb{E}_{\pi_{K,b}}[c(x, u)\psi(x, u)] - \xi^2.$$

The primal-dual gradient method updates ζ and ξ via

$$\begin{aligned} \zeta^1 &\leftarrow \zeta^1 - \gamma \cdot \nabla_{\zeta^1} F(\zeta, \xi), & \zeta^2 &\leftarrow \zeta^2 - \gamma \cdot \nabla_{\zeta^2} F(\zeta, \xi) \\ \xi^1 &\leftarrow \xi^1 - \gamma \cdot \nabla_{\xi^1} F(\zeta, \xi), & \xi^2 &\leftarrow \xi^2 - \gamma \cdot \nabla_{\xi^2} F(\zeta, \xi). \end{aligned} \quad (\text{B.2.13})$$

Estimation of Mean-Field State $\mu_{K,b}$. To utilize the primal-dual gradient method in (B.2.13), it remains to evaluate the feature vector $\psi(x, u)$. Note that by (B.2.5), the evaluation of the feature vector $\psi(x, u)$ requires the mean-field state $\mu_{K,b}$. In what follows, we establish the estimator $\hat{\mu}_{K,b}$ of the mean-field state $\mu_{K,b}$ by simulating the MDP following the policy $\pi_{K,b}$ for \tilde{T} steps, and calculate the estimated feature vector $\hat{\psi}(x, u)$ by

$$\hat{\psi}(x, u) = \begin{pmatrix} \hat{\varphi}(x, u) \\ x - \hat{\mu}_{K,b} \\ u - (-K\hat{\mu}_{K,b} + b) \end{pmatrix}, \quad (\text{B.2.14})$$

where $\widehat{\varphi}(x, u)$ takes the form of

$$\widehat{\varphi}(x, u) = \text{svec} \left[\begin{pmatrix} x - \widehat{\mu}_{K,b} \\ u - (-K\widehat{\mu}_{K,b} + b) \end{pmatrix} \begin{pmatrix} x - \widehat{\mu}_{K,b} \\ u - (-K\widehat{\mu}_{K,b} + b) \end{pmatrix}^\top \right].$$

We now define the sets \mathcal{V}_ζ and \mathcal{V}_ξ in (B.2.11).

Definition B.2.7. Given K_0 and b_0 such that $\rho(A - BK_0) < 1$ and $J(K_0, b_0) < \infty$, we define the sets \mathcal{V}_ζ and \mathcal{V}_ξ as

$$\begin{aligned} \mathcal{V}_\zeta &= \left\{ \zeta: 0 \leq \zeta^1 \leq J(K_0, b_0), \|\zeta^2\|_2 \leq M_{\zeta,1} + M_{\zeta,2} \cdot (1 + \|K\|_F) \cdot [1 - \rho(A - BK)]^{-1} \right\}, \\ \mathcal{V}_\xi &= \left\{ \xi: |\xi^1| \leq J(K_0, b_0), \|\xi^2\|_2 \leq M_\xi \cdot (1 + \|K\|_F^2)^3 \cdot [1 - \rho(A - BK)]^{-1} \right\}. \end{aligned}$$

Here $M_{\zeta,1}$, $M_{\zeta,2}$, and M_ξ are constants independent of K and b , which take the forms of

$$\begin{aligned} M_{\zeta,1} &= \left[(\|Q\|_F + \|R\|_F) + (\|A\|_F^2 + \|B\|_F^2) \cdot \sqrt{d} \cdot J(K_0, b_0) \cdot \sigma_{\min}^{-1}(\Psi_\omega) \right] \\ &\quad + (\|A\|_2 + \|B\|_2) \cdot J(K_0, b_0)^2 \cdot \sigma_{\min}^{-1}(\Psi_\omega) \cdot \sigma_{\min}^{-1}(Q), \\ &\quad + \left[(\|Q\|_2 + \|R\|_2) + (\|A\|_2 + \|B\|_2)^2 \cdot J(K_0, b_0) \cdot \sigma_{\min}^{-1}(\Psi_\omega) \right] \\ &\quad \cdot J(K_0, b_0) \cdot [\sigma_{\min}^{-1}(Q) + \sigma_{\min}^{-1}(R)] \end{aligned}$$

$$M_{\zeta,2} = (\|A\|_2 + \|B\|_2) \cdot (\kappa_Q + \kappa_R), \quad M_\xi = C \cdot (M_{\zeta,1} + M_{\zeta,2}) \cdot J(K_0, b_0)^2 \cdot \sigma_{\min}^{-2}(Q),$$

where C is a positive absolute constant, and κ_Q and κ_R are condition numbers of Q and R , respectively.

We summarize the primal-dual gradient temporal difference algorithm in Algorithm 7. Hereafter, for notational convenience, we denote by $\widehat{\psi}_t$ the estimated feature vector $\widehat{\psi}(x_t, u_t)$.

Algorithm 7 Primal-Dual Gradient Temporal Difference Algorithm.

- 1: **Input:** Policy $\pi_{K,b}$, mean-field state μ , numbers of iteration \widetilde{T} and T , stepsizes $\{\gamma_t\}_{t \in [T]}$, parameters K_0 and b_0 .
 - 2: Define the sets \mathcal{V}_ζ and \mathcal{V}_ξ via Definition B.2.7 with K_0 and b_0 .
 - 3: Initialize the parameters by $\zeta_0 \in \mathcal{V}_\zeta$ and $\xi_0 \in \mathcal{V}_\xi$.
 - 4: Sample \widetilde{x}_0 from the stationary distribution $\mathcal{N}(\mu_{K,b}, \Phi_K)$.
 - 5: **for** $t = 0, \dots, \widetilde{T} - 1$ **do**
 - 6: Given the mean-field state μ , take action \widetilde{u}_t following $\pi_{K,b}$ and generate the next state \widetilde{x}_{t+1} .
 - 7: **end for**
 - 8: Set $\widehat{\mu}_{K,b} \leftarrow 1/\widetilde{T} \cdot \sum_{t=1}^{\widetilde{T}} \widetilde{x}_t$ and compute the estimated feature vector $\widehat{\psi}$ via (B.2.14).
 - 9: Sample x_0 from the stationary distribution $\mathcal{N}(\mu_{K,b}, \Phi_K)$.
 - 10: **for** $t = 0, \dots, T - 1$ **do**
 - 11: Given the mean-field state μ , take action u_t following $\pi_{K,b}$, observe the cost c_t , and generate the next state x_{t+1} .
 - 12: Set $\delta_{t+1} \leftarrow \zeta_t^1 + (\widehat{\psi}_t - \widehat{\psi}_{t+1})^\top \zeta_t^2 - c_t$.
 - 13: Update parameters via

$$\begin{aligned} \zeta_{t+1}^1 &\leftarrow \zeta_t^1 - \gamma_{t+1} \cdot (\zeta_t^1 + \widehat{\psi}_t^\top \zeta_t^2), & \zeta_{t+1}^2 &\leftarrow \zeta_t^2 - \gamma_{t+1} \cdot \widehat{\psi}_t (\widehat{\psi}_t - \widehat{\psi}_{t+1})^\top \zeta_t^2, \\ \xi_{t+1}^1 &\leftarrow (1 - \gamma_{t+1}) \cdot \xi_t^1 + \gamma_{t+1} \cdot (\zeta_t^1 - c_t), & \xi_{t+1}^2 &\leftarrow (1 - \gamma_{t+1}) \cdot \xi_t^2 + \gamma_{t+1} \cdot \delta_{t+1} \cdot \widehat{\psi}_t. \end{aligned}$$
 - 14: Project ζ_{t+1} and ξ_{t+1} to \mathcal{V}_ζ and \mathcal{V}_ξ , respectively.
 - 15: **end for**
 - 16: Set $\widehat{\alpha}_{K,b} \leftarrow (\sum_{t=1}^T \gamma_t)^{-1} \cdot (\sum_{t=1}^T \gamma_t \cdot \zeta_t^2)$, and

$$\widehat{\Upsilon}_K \leftarrow \text{smat}(\widehat{\alpha}_{K,b,1}), \quad \begin{pmatrix} \widehat{p}_{K,b} \\ \widehat{q}_{K,b} \end{pmatrix} \leftarrow \widehat{\alpha}_{K,b,2} - \widehat{\Upsilon}_K \begin{pmatrix} \widehat{\mu}_{K,b} \\ -K\widehat{\mu}_{K,b} + b \end{pmatrix},$$
 where $\widehat{\alpha}_{K,b,1} = (\widehat{\alpha}_{K,b})_1^{(k+d+1)(k+d)/2}$ and $\widehat{\alpha}_{K,b,2} = (\widehat{\alpha}_{K,b})_{(k+d+1)(k+d)/2+1}^{(k+d+3)(k+d)/2}$.
 - 17: **Output:** Estimators $\widehat{\mu}_{K,b}$, $\widehat{\Upsilon}_K$, and $\widehat{q}_{K,b}$.
-

We now characterize the rate of convergence of Algorithm 7.

Theorem B.2.8 (Convergence of Algorithm 7). Given K_0 , b_0 , K , and b such that $\rho(A - BK_0) < 1$ and $J(K, b) \leq J(K_0, b_0)$, we define the sets \mathcal{V}_ζ and \mathcal{V}_ξ through Definition B.2.7.

Let $\gamma_t = \gamma_0 t^{-1/2}$, where γ_0 is a positive absolute constant. Let $\rho \in (\rho(A - BK), 1)$. For $\tilde{T} \geq \text{poly}_0(\|K\|_{\mathbb{F}}, \|b\|_2, \|\mu\|_2, J(K_0, b_0)) \cdot (1 - \rho)^{-6}$ and a sufficiently large T , it holds with probability at least $1 - T^{-4} - \tilde{T}^{-6}$ that

$$\begin{aligned} & \|\hat{\alpha}_{K,b} - \alpha_{K,b}\|_2^2 \\ & \leq \lambda_K^{-2} \cdot \text{poly}_1(\|K\|_{\mathbb{F}}, \|b\|_2, \|\mu\|_2, J(K_0, b_0)) \cdot \left[\frac{\log^6 T}{T^{1/2} \cdot (1 - \rho)^4} + \frac{\log \tilde{T}}{\tilde{T}^{1/4} \cdot (1 - \rho)^2} \right], \end{aligned}$$

where λ_K is defined in Proposition B.2.6. Same bounds for $\|\hat{\Upsilon}_K - \Upsilon_K\|_{\mathbb{F}}^2$, $\|\hat{p}_{K,b} - p_{K,b}\|_2^2$, and $\|\hat{q}_{K,b} - q_{K,b}\|_2^2$ hold. Meanwhile, it holds with probability at least $1 - \tilde{T}^{-6}$ that

$$\|\hat{\mu}_{K,b} - \mu_{K,b}\|_2 \leq \frac{\log \tilde{T}}{\tilde{T}^{1/4}} \cdot (1 - \rho)^{-2} \cdot \text{poly}_2(\|\Phi_K\|_2, \|K\|_{\mathbb{F}}, \|b\|_2, \|\mu\|_2, J(K_0, b_0)).$$

Proof. See §B.4.3 for a detailed proof. □

B.2.3. Temporal Difference Policy Evaluation Algorithm

Besides the primal-dual gradient temporal difference algorithm, we can also evaluate $\alpha_{K,b}$ by TD(0) method (Sutton and Barto, 2018) in practice, which is presented in Algorithm 8.

Note that in related literature (Bhandari et al., 2018; Korda and La, 2015), non-asymptotic convergence analysis of TD(0) method with linear function approximation is only applied to discounted MDP. As for our ergodic setting, the convergence of TD(0) method is only shown asymptotically (Borkar and Meyn, 2000; Kushner and Yin, 2003)

Algorithm 8 Temporal Difference Policy Evaluation Algorithm.

- 1: **Input:** Policy $\pi_{K,b}$, number of iteration \tilde{T} and T , stepsizes $\{\gamma_t\}_{t \in [T]}$.
 - 2: Sample \tilde{x}_0 from the stationary distribution $\mathcal{N}(\mu_{K,b}, \Phi_K)$.
 - 3: **for** $t = 0, \dots, \tilde{T} - 1$ **do**
 - 4: Take action \tilde{u}_t under the policy $\pi_{K,b}$ and generate the next state \tilde{x}_{t+1} .
 - 5: **end for**
 - 6: Set $\hat{\mu}_{K,b} \leftarrow 1/\tilde{T} \cdot \sum_{t=1}^{\tilde{T}} \tilde{x}_t$.
 - 7: Sample x_0 from the stationary distribution $\mathcal{N}(\mu_{K,b}, \Phi_K)$.
 - 8: **for** $t = 0, \dots, T$ **do**
 - 9: Given the mean-field state μ , take action u_t following $\pi_{K,b}$, observe the cost c_t , and generate the next state x_{t+1} .
 - 10: Set $\delta_{t+1} \leftarrow \zeta_t^1 + (\hat{\psi}_t - \hat{\psi}_{t+1})^\top \zeta_t^2 - c_t$.
 - 11: Update parameters via $\zeta_{t+1}^1 \leftarrow (1 - \gamma_{t+1}) \cdot \zeta_t^1 + \gamma_{t+1} \cdot c_t$ and $\zeta_{t+1}^2 \leftarrow \zeta_t^2 - \gamma_{t+1} \cdot \delta_{t+1} \cdot \hat{\psi}_t$.
 - 12: Project ζ_t to \mathcal{V}'_ζ , where \mathcal{V}'_ζ is a compact set.
 - 13: **end for**
 - 14: Set $\hat{\alpha}_{K,b} \leftarrow (\sum_{t=1}^T \gamma_t)^{-1} \cdot (\sum_{t=1}^T \gamma_t \cdot \zeta_t^2)$, and
$$\hat{\Upsilon}_K \leftarrow \text{smat}(\hat{\alpha}_{K,b,1}), \quad \begin{pmatrix} \hat{p}_{K,b} \\ \hat{q}_{K,b} \end{pmatrix} \leftarrow \hat{\alpha}_{K,b,2} - \hat{\Upsilon}_K \begin{pmatrix} \hat{\mu}_{K,b} \\ -K\hat{\mu}_{K,b} + b \end{pmatrix},$$

where $\hat{\alpha}_{K,b,1} = (\hat{\alpha}_{K,b})_1^{(k+d+1)(k+d)/2}$ and $\hat{\alpha}_{K,b,2} = (\hat{\alpha}_{K,b})_{(k+d+1)(k+d)/2+1}^{(k+d+3)(k+d)/2}$.
 - 15: **Output:** Estimators $\hat{\mu}_{K,b}$, $\hat{\Upsilon}_K$, and $\hat{q}_{K,b}$.
-

using ordinary differential equation method. Therefore, in the convergence theorem proposed in §3.2, we only focus on the primal-dual gradient temporal difference method (Algorithm 7) to establish non-asymptotic convergence result.

B.3. General Formulation

Compared with Problem 3.1.3, a more general formulation includes an additional term $x_t^\top P\mu$ in the cost function. For the completeness of this paper, we define this general formulation here. Following from a same argument as in §3.1, it suffices to study the setting where t is sufficiently large. First, we propose the following general drifted LQR (general D-LQR) problem, which is parallel to Problem 3.1.2.

Problem B.3.1 (General D-LQR). For any given mean-field state $\mu \in \mathbb{R}^m$, consider the following formulation

$$\begin{aligned} x_{t+1} &= Ax_t + Bu_t + \bar{A}\mu + d + \omega_t, \\ \tilde{c}_\mu(x_t, u_t) &= x_t^\top Qx_t + u_t^\top Ru_t + \mu^\top \bar{Q}\mu + 2x_t^\top P\mu, \\ \tilde{J}_\mu(\pi) &= \lim_{T \rightarrow \infty} \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^T \tilde{c}_\mu(x_t, u_t) \right], \end{aligned}$$

where $x_t \in \mathbb{R}^m$ is the state vector, $u_t \in \mathbb{R}^k$ is the action vector generated by the policy π , $\omega_t \in \mathbb{R}^m$ is an independent random noise term following the Gaussian distribution $\mathcal{N}(0, \Psi_\omega)$, and $d \in \mathbb{R}^m$ is a drift term. We aim to find an optimal policy π_μ^* such that $\tilde{J}_\mu(\pi_\mu^*) = \inf_{\pi \in \Pi} \tilde{J}_\mu(\pi)$.

In Problem B.3.1, the unique optimal policy $\pi_\mu^*(\cdot)$ still admits a linear form $\pi_\mu^*(x_t) = -K_{\pi_\mu^*}x_t + b_{\pi_\mu^*}$ (Anderson and Moore, 2007), where the matrix $K_{\pi_\mu^*} \in \mathbb{R}^{k \times m}$ and the vector $b_{\pi_\mu^*} \in \mathbb{R}^k$ are the parameters of the policy π . It then suffices to find the optimal policy in the class Π defined in (3.1.1).

Parallel to Problem 3.1.3, we define the general LQ-MFG problem as follows.

Problem B.3.2 (General LQ-MFG). We consider the following formulation

$$\begin{aligned} x_{t+1} &= Ax_t + Bu_t + \bar{A}\mu + d + \omega_t, \\ \tilde{c}(x_t, u_t) &= x_t^\top Qx_t + u_t^\top Ru_t + \mu^\top \bar{Q}\mu + 2x_t^\top P\mu, \\ \tilde{J}(\pi, \mu) &= \lim_{T \rightarrow \infty} \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^T \tilde{c}(x_t, u_t) \right], \end{aligned}$$

where $x_t \in \mathbb{R}^m$ is the state vector, $u_t \in \mathbb{R}^k$ is the action vector generated by the policy π , $\mu \in \mathbb{R}^m$ is the mean-field state, $\omega_t \in \mathbb{R}^m$ is an independent random noise term following the Gaussian distribution $\mathcal{N}(0, \Psi_\omega)$, and $d \in \mathbb{R}^m$ is a drift term. We aim to find a pair (μ^*, π^*) such that (i) $\tilde{J}(\pi^*, \mu^*) = \inf_{\pi \in \Pi} \tilde{J}(\pi, \mu^*)$; (ii) $\mathbb{E}x_t^*$ converges to μ^* as $t \rightarrow \infty$, where $\{x_t^*\}_{t \geq 0}$ is the Markov chain of states generated by the policy π^* .

One can see that Problem B.3.2 aims to find a Nash equilibrium pair (μ^*, π^*) .

Similar to the discussions in §3.2.2, to solve Problem B.3.2, one can design an algorithm similar to Algorithm 2, which solves Problem B.3.1 and obtain the new mean-field state at each iteration. We omit the detailed algorithm here. Now, we focus on solving Problem B.3.1 in the sequel.

Similar to §3.2.3, we drop the subscript μ when we focus on Problem B.3.1 for a fixed μ . We write $\pi_{K,b}(x) = -Kx + b + \sigma \cdot \eta$ to emphasize the dependence on K and b , and $\tilde{J}(K, b) = \tilde{J}(\pi_{K,b})$ consequently. We derive an explicit form of the expected total cost $\tilde{J}(K, b)$ in the following proposition.

Proposition B.3.3. The expected total cost $\tilde{J}(K, b)$ in Problem B.3.1 is decomposed as

$$\tilde{J}(K, b) = \tilde{J}_1(K) + \tilde{J}_2(K, b) + \sigma^2 \cdot \text{tr}(R) + \mu^\top \bar{Q} \mu,$$

where $\tilde{J}_1(K)$ and $\tilde{J}_2(K, b)$ take the forms of

$$\begin{aligned} \tilde{J}_1(K) &= \text{tr}[(Q + K^\top RK)\Phi_K] = \text{tr}(P_K \Psi_\epsilon), \\ \tilde{J}_2(K, b) &= \begin{pmatrix} \mu_{K,b} \\ b \end{pmatrix}^\top \begin{pmatrix} Q + K^\top RK & -K^\top R \\ -RK & R \end{pmatrix} \begin{pmatrix} \mu_{K,b} \\ b \end{pmatrix} + 2\mu^\top P \mu_{K,b}. \end{aligned}$$

Here $\mu_{K,b}$ is given in (3.2.2), Φ_K is given in (3.2.3), and P_K is given in (3.2.4).

Proof. The proof is similar to the one of Proposition B.2.2. Thus we omit it here. \square

Compared with the form of $J(K, b)$ given in (3.2.5), we see that the only difference is that $\tilde{J}(K, b)$ contains an extra term $2\mu^\top P\mu_{K,b}$ in $\tilde{J}_2(K, b)$, which is only a linear term in b (recall that $\mu_{K,b}$ is linear in b by (3.2.2)). Thus, $\tilde{J}_2(K, b)$ is still strongly convex in b , as shown in the proposition below.

Proposition B.3.4. Given any K , the function $\tilde{J}_2(K, b)$ is ν_K -strongly convex in b , here $\nu_K = \sigma_{\min}(Y_{1,K}^\top Y_{1,K} + Y_{2,K}^\top Y_{2,K})$, where $Y_{1,K} = R^{1/2}K(I - A + BK)^{-1}B - R^{1/2}$ and $Y_{2,K} = Q^{1/2}(I - A + BK)^{-1}B$. Also, $\tilde{J}_2(K, b)$ has ι_K -Lipschitz continuous gradient in b , where ι_K is upper bounded such that $\iota_K \leq [1 - \rho(A - BK)]^{-2} \cdot (\|B\|_*^2 \cdot \|K\|_*^2 \cdot \|R\|_* + \|B\|_*^2 \cdot \|Q\|_*)$.

Proof. The proof is similar to the one of Proposition 3.2.3. Thus we omit it here. \square

Parallel to Proposition 3.2.4, we derive a similar proposition in the sequel.

Proposition B.3.5. Denote by $\tilde{b}^K = \operatorname{argmin}_b \tilde{J}_2(K, b)$, then $\tilde{J}_2(K, \tilde{b}^K)$ takes the form

$$\begin{aligned} & \tilde{J}_2(K, \tilde{b}^K) \\ &= \begin{pmatrix} \bar{A}\mu + d \\ P^\top \mu \end{pmatrix}^\top \begin{pmatrix} S & S(I - A)Q^{-1} \\ Q^{-1}(I - A)^\top S & 3Q^{-1}(I - A)^\top S(I - A)Q^{-1} - Q^{-1} \end{pmatrix} \begin{pmatrix} \bar{A}\mu + d \\ P^\top \mu \end{pmatrix}, \end{aligned}$$

which is independent of K . Here $S = [(I - A)Q^{-1}(I - A)^\top + BR^{-1}B^\top]^{-1}$. And \tilde{b}^K takes the form

$$\tilde{b}^K = [KQ^{-1}(I - A)^\top - R^{-1}B^\top] \cdot S \cdot [(\bar{A}\mu + d) + (I - A)Q^{-1}P^\top \mu] - KQ^{-1}P^\top \mu.$$

Proof. The proof is similar to the one of Proposition 3.2.4. Thus we omit it here. \square

Similar to Problem 3.1.2, we define the state- and action-value functions as

$$\begin{aligned}\tilde{V}_{K,b}(x) &= \sum_{t=0}^{\infty} \left\{ \mathbb{E}[\tilde{c}(x_t, u_t) \mid x_0 = x, u_t = -Kx_t + b + \sigma\eta_t] - \tilde{J}(K, b) \right\}, \\ \tilde{Q}_{K,b}(x, u) &= \tilde{c}(x, u) - \tilde{J}(K, b) + \mathbb{E}[\tilde{V}_{K,b}(x') \mid x, u],\end{aligned}$$

where the x' is the state generated by the state transition after the state-action pair (x, u) .

A slight modification of Proposition B.2.1 gives the proposition below.

Proposition B.3.6. For Problem B.3.1, the state-value function $\tilde{V}_{K,b}(x)$ takes the form

$$\tilde{V}_{K,b}(x) = x^\top P_K x - \text{tr}(P_K \Phi_K) + 2\tilde{f}_{K,b}^\top (x - \mu_{K,b}) - (\mu_{K,b})^\top P_K \mu_{K,b},$$

and the action-value function $\tilde{Q}_{K,b}(x, u)$ takes the form

$$\begin{aligned}\tilde{Q}_{K,b}(x, u) &= \begin{pmatrix} x \\ u \end{pmatrix}^\top \Upsilon_K \begin{pmatrix} x \\ u \end{pmatrix} + 2 \begin{pmatrix} \tilde{p}_{K,b} \\ \tilde{q}_{K,b} \end{pmatrix}^\top \begin{pmatrix} x \\ u \end{pmatrix} \\ &\quad - \text{tr}(P_K \Phi_K) - \sigma^2 \cdot \text{tr}(R + P_K B B^\top) - b^\top R b \\ &\quad + 2b^\top R K \mu_{K,b} - (\mu_{K,b})^\top (Q + K^\top R K + P_K) \mu_{K,b} + 2\tilde{f}_{K,b}^\top [(\bar{A}\mu + d) - \mu_{K,b}] \\ &\quad + (\bar{A}\mu + d)^\top P_K (\bar{A}\mu + d) - 2\mu^\top P \mu_{K,b}.\end{aligned}$$

Here the matrix Υ_K is given in (3.2.7), and the vectors $\tilde{p}_{K,b}, \tilde{q}_{K,b}$ are given as

$$(B.3.1) \quad \begin{pmatrix} \tilde{p}_{K,b} \\ \tilde{q}_{K,b} \end{pmatrix} = \begin{pmatrix} A^\top [P_K \cdot (\bar{A}\mu + d) + \tilde{f}_{K,b}] + P\mu \\ B^\top [P_K \cdot (\bar{A}\mu + d) + \tilde{f}_{K,b}] \end{pmatrix},$$

where the vector $\tilde{f}_{K,b} = (I - A + BK)^{-\top}[(A - BK)^\top P_K(Bb + \bar{A}\mu + d) - K^\top Rb + P\mu]$.

Proof. The proof is similar to the one of Proposition B.2.1. Thus we omit it here. \square

Now we establish the gradients of $\tilde{J}(K, b)$ for Problem B.3.1.

Proposition B.3.7. The gradient of $\tilde{J}_1(K)$ and the gradient of $\tilde{J}_2(K, b)$ w.r.t. b takes the form

$$\nabla_K \tilde{J}_1(K) = 2(\Upsilon_K^{22} K - \Upsilon_K^{21}) \cdot \Phi_K,$$

$$\nabla_b \tilde{J}_2(K, b) = 2[\Upsilon_K^{22}(-K\mu_{K,b} + b) + \Upsilon_K^{21}\mu_{K,b} + \tilde{q}_{K,b}],$$

where the matrix Υ_K is given in (3.2.7), and the vector $\tilde{q}_{K,b}$ is given in (B.3.1).

Proof. The proof is similar to the one of Proposition B.2.3. Thus we omit it here. \square

Equipped with above results, parallel to the analysis in §3.2, it is clear that by slight modification of Algorithms 2, 3, and 7, we derive similar actor-critic algorithms to solve both Problem B.3.2 and Problem B.3.1, where all the non-asymptotic convergence results hold. We omit the algorithms and the convergence results here.

B.4. Proofs of Theorems

B.4.1. Proof of Theorem 3.3.1

We define $\mu_{s+1}^* = \Lambda(\mu_s)$, which is the mean-field state generated by the optimal policy $\pi_{K^*(\mu_s), b^*(\mu_s)} = \Lambda_1(\mu_s)$ under the current mean-field state μ_s . By Proposition 3.2.4, the optimal $K^*(\mu)$ is independent of the mean-field state μ . Therefore, we write $K^* = K^*(\mu)$

hereafter for notational convenience. By (3.2.2), we know that

$$\mu_{s+1}^* = (I - A + BK^*)^{-1} \cdot [Bb^*(\mu_s) + \bar{A}\mu_s + d].$$

We define

$$\tilde{\mu}_{s+1} = (I - A + BK_s)^{-1}(Bb_s + \bar{A}\mu_s + d),$$

which is the mean-field state generated by the policy π_s under the current mean-field state μ_s , where K_s and b_s are the parameters of the policy π_s . By triangle inequality, we have

$$(B.4.1) \quad \|\mu_{s+1} - \mu^*\|_2 \leq \underbrace{\|\mu_{s+1} - \tilde{\mu}_{s+1}\|_2}_{E_1} + \underbrace{\|\tilde{\mu}_{s+1} - \mu_{s+1}^*\|_2}_{E_2} + \underbrace{\|\mu_{s+1}^* - \mu^*\|_2}_{E_3},$$

where μ_{s+1} is generated by Algorithm 2. We upper bound E_1 , E_2 , and E_3 in the sequel.

Upper Bound of E_1 . By Theorem B.2.4, it holds with probability at least $1 - \varepsilon^{10}$ that

$$(B.4.2) \quad E_1 = \|\mu_{s+1} - \tilde{\mu}_{s+1}\|_2 < \varepsilon_s \leq \varepsilon/8 \cdot 2^{-s},$$

where ε_s is given in (3.3.2).

Upper Bound of E_2 . By the triangle inequality, we have

$$(B.4.3) \quad \begin{aligned} E_2 &= \left\| (I - A + BK_s)^{-1}(Bb_s + \bar{A}\mu_s + d) - (I - A + BK^*)^{-1} \cdot [Bb^*(\mu_s) + \bar{A}\mu_s + d] \right\|_2 \\ &\leq \|Bb^*(\mu_s) + \bar{A}\mu_s + d\|_2 \cdot \left\| [I - A + BK^* + B(K_s - K^*)]^{-1} - (I - A + BK^*)^{-1} \right\|_2 \\ &\quad + \|(I - A + BK_s)^{-1}\|_2 \cdot \|B\|_2 \cdot \|b_s - b^*(\mu_s)\|_2. \end{aligned}$$

By Taylor's expansion, we have

$$\begin{aligned}
& \left\| [I - A + BK^* + B(K_s - K^*)]^{-1} - (I - A + BK^*)^{-1} \right\|_2 \\
&= \left\| (I - A + BK^*)^{-1} [I + (I - A + BK^*)^{-1} B(K_s - K^*)]^{-1} - (I - A + BK^*)^{-1} \right\|_2 \\
\text{(B.4.4)} \quad & \leq 2 \left\| (I - A + BK^*)^{-1} B(K_s - K^*) (I - A + BK^*)^{-1} \right\|_2.
\end{aligned}$$

Meanwhile, by Taylor's expansion, it holds with probability at least $1 - \varepsilon^{10}$ that

$$\begin{aligned}
& \left\| (I - A + BK_s)^{-1} \right\|_2 \\
&= \left\| (I - A + BK^* + B(K_s - K^*))^{-1} \right\|_2 \\
&= \left\| (I - A + BK^*)^{-1} (I + (I - A + BK^*)^{-1} B(K_s - K^*))^{-1} \right\|_2 \\
&\leq [1 - \rho(A - BK^*)]^{-1} \cdot \left(1 + \left\| (I - A + BK^*)^{-1} B \right\|_2 \cdot \|K^* - K_s\|_2 \right) \\
\text{(B.4.5)} \quad & \leq 2 [1 - \rho(A - BK^*)]^{-2},
\end{aligned}$$

where the last inequality comes from Theorem B.2.4. By plugging (B.4.4) and (B.4.5) in (B.4.3), it holds with probability at least $1 - \varepsilon^{10}$ that

$$\begin{aligned}
E_2 &\leq 2 \left\| Bb^*(\mu_s) + \bar{A}\mu_s + d \right\|_2 \cdot \left\| (I - A + BK^*)^{-1} B(K_s - K^*) (I - A + BK^*)^{-1} \right\|_2 \\
&\quad + \left\| (I - A + BK_s)^{-1} \right\|_2 \cdot \|B\|_2 \cdot \|b_s - b^*(\mu_s)\|_2 \\
\text{(B.4.6)} \quad &\leq 2 \left\| Bb^*(\mu_s) + \bar{A}\mu_s + d \right\|_2 \cdot [1 - \rho(A - BK^*)]^{-2} \cdot \|B\|_2 \cdot \|K_s - K^*\|_2 \\
&\quad + 2 [1 - \rho(A - BK^*)]^{-2} \cdot \|B\|_2 \cdot \|b_s - b^*(\mu_s)\|_2.
\end{aligned}$$

By Proposition 3.2.4, it holds that

$$(B.4.7) \quad \begin{aligned} \|Bb^*(\mu_s) + \bar{A}\mu_s + d\|_2 &\leq L_1 \cdot \|B\|_2 \cdot \|\mu_s\|_2 + \|\bar{A}\|_2 \cdot \|\mu_s\|_2 + \|d\|_2 \\ &\leq (L_1 \cdot \|B\|_2 + \|\bar{A}\|_2) \cdot \|\mu_s\|_2 + \|d\|_2, \end{aligned}$$

where the scalar L_1 is defined in Assumption 3.2.1. Meanwhile, by Theorem B.2.4, it holds with probability at least $1 - \varepsilon^{10}$ that

$$(B.4.8) \quad \|K_s - K^*\|_F \leq [\sigma_{\min}^{-1}(\Psi_\varepsilon) \cdot \sigma_{\min}^{-1}(R) \cdot \varepsilon_s]^{1/2}, \quad \|b_s - b^*(\mu_s)\|_2 \leq M_b(\mu_s) \cdot \varepsilon_s^{1/2},$$

where $M_b(\mu_s)$ is defined in (3.3.3). Combining (B.4.6), (B.4.7), (B.4.8), and the choice of ε_s in (3.3.2), it holds with probability at least $1 - \varepsilon^{10}$ that

$$(B.4.9) \quad E_2 \leq \varepsilon/8 \cdot 2^{-s}.$$

Upper Bound of E_3 . By Proposition 3.2.2, we have

$$(B.4.10) \quad E_3 = \|\mu_{s+1}^* - \mu^*\|_2 = \|\Lambda(\mu_s) - \Lambda(\mu^*)\|_2 \leq L_0 \cdot \|\mu_s - \mu^*\|_2,$$

where $L_0 = L_1L_3 + L_2$ by Assumption 3.2.1.

By plugging (B.4.2), (B.4.9), and (B.4.10) in (B.4.1), we know that

$$(B.4.11) \quad \|\mu_{s+1} - \mu^*\|_2 \leq L_0 \cdot \|\mu_s - \mu^*\|_2 + \varepsilon \cdot 2^{-s-2},$$

which holds with probability at least $1 - \varepsilon^{10}$. Following from (B.4.11) and a union bound argument with $S = \mathcal{O}(\log(1/\varepsilon))$, it holds with probability at least $1 - \varepsilon^5$ that

$$\|\mu_S - \mu^*\|_2 \leq L_0^S \cdot \|\mu_0 - \mu^*\|_2 + \varepsilon/2,$$

where we use the fact that $L_0 < 1$ by Assumption 3.2.1. By the choice of S in (3.3.1), it further holds with probability at least $1 - \varepsilon^6$ that

$$(B.4.12) \quad \|\mu_S - \mu^*\| \leq \varepsilon.$$

By Theorem B.2.4 and the choice of ε_s in (3.3.2), it holds with probability at least $1 - \varepsilon^5$ that

$$(B.4.13) \quad \|K_S - K^*\|_F = \|K_S - K^*(\mu_S)\|_F \leq [\sigma_{\min}^{-1}(\Psi_\varepsilon) \cdot \sigma_{\min}^{-1}(R) \cdot \varepsilon_S]^{1/2} \leq \varepsilon.$$

Meanwhile, by the triangle inequality and the choice of ε_s in (3.3.2), it holds with probability at least $1 - \varepsilon^5$ that

$$(B.4.14) \quad \begin{aligned} \|b_S - b^*\|_2 &\leq \|b_S - b^*(\mu_S)\|_2 + \|b^*(\mu_S) - b^*\|_2 \\ &\leq M_b(\mu_S) \cdot \varepsilon_S^{1/2} + L_1 \cdot \|\mu_S - \mu^*\|_2 \\ &\leq (1 + L_1) \cdot \varepsilon, \end{aligned}$$

where the second inequality comes from Theorem B.2.4 and Proposition 3.2.4, and the last inequality comes from (B.4.12). By (B.4.12), (B.4.13), and (B.4.14), we conclude the proof of the theorem.

B.4.2. Proof of Theorem B.2.4

Proof. We first show that $J_1(K_N) - J_1(K^*) < \varepsilon/2$ with a high probability, then show that $J_2(K_N, b_H) - J_2(K^*, b^*) < \varepsilon/2$ with a high probability. Then we have

$$J(K_N, b_N) - J(K^*, b^*) = J_1(K_N) + J_2(K_N, b_H) - J_1(K^*) - J_2(K^*, b^*) < \varepsilon$$

with a high probability, which proves Theorem B.2.4.

Part 1. We show that $J_1(K_N) - J_1(K^*) < \varepsilon/2$ with a high probability.

We first bound $J_1(K_1) - J_1(K_2)$ for any K_1 and K_2 . By Proposition B.2.2, $J_1(K)$ takes the form of

$$(B.4.15) \quad J_1(K) = \text{tr}(P_K \Psi_\epsilon) = \mathbb{E}_{y \sim \mathcal{N}(0, \Psi_\epsilon)}(y^\top P_K y).$$

The following lemma calculates $y^\top P_{K_1} y - y^\top P_{K_2} y$ for any K_1 and K_2 .

Lemma B.4.1. Assume that $\rho(A - BK_1) < 1$ and $\rho(A - BK_2) < 1$. For any state vector y , we denote by $\{y_t\}_{t \geq 0}$ the sequence generated by the state transition $y_{t+1} = (A - BK_2)y_t$ with initial state $y_0 = y$. It holds that

$$y^\top P_{K_2} y - y^\top P_{K_1} y = \sum_{t \geq 0} D_{K_1, K_2}(y_t),$$

where

$$D_{K_1, K_2}(y) = 2y^\top (K_2 - K_1)(\Upsilon_{K_1}^{22} K_1 - \Upsilon_{K_1}^{21})y + y^\top (K_2 - K_1)^\top \Upsilon_{K_1}^{22} (K_2 - K_1)y.$$

Here Υ_K is defined in (3.2.7).

Proof. See §B.6.1 for a detailed proof. \square

The following lemma shows that $J_1(K)$ is gradient dominant.

Lemma B.4.2. Let K^* be the optimal parameter and K be a parameter such that $J_1(K) < \infty$, then it holds that

$$(B.4.16) \quad J_1(K) - J_1(K^*) \leq \sigma_{\min}^{-1}(R) \cdot \|\Phi_{K^*}\|_2 \cdot \text{tr}[(\Upsilon_K^{22}K - \Upsilon_K^{21})^\top (\Upsilon_K^{22}K - \Upsilon_K^{21})],$$

$$(B.4.17) \quad J_1(K) - J_1(K^*) \geq \sigma_{\min}(\Psi_\omega) \cdot \|\Upsilon_K^{22}\|_2^{-1} \cdot \text{tr}[(\Upsilon_K^{22}K - \Upsilon_K^{21})^\top (\Upsilon_K^{22}K - \Upsilon_K^{21})].$$

Proof. See §B.6.2 for a detailed proof. \square

Recall that from Algorithm 3, the parameter K is updated via

$$(B.4.18) \quad K_{n+1} = K_n - \gamma \cdot (\hat{\Upsilon}_{K_n}^{22} K_n - \hat{\Upsilon}_{K_n}^{21}),$$

where $\hat{\Upsilon}_{K_n}$ is the output of Algorithm 7. We upper bound $|J_1(K_{n+1}) - J_1(K^*)|$ in the sequel. First, we show that if $J_1(K_n) - J_1(K^*) \geq \varepsilon/2$ holds for any $n \leq N$, we obtain that

$$(B.4.19) \quad J_1(K_N) \leq J_1(K_{N-1}) \leq \cdots \leq J_1(K_0),$$

which holds with probability at least $1 - \varepsilon^{13}$. We prove (B.4.19) by mathematical induction. Suppose that

$$(B.4.20) \quad J_1(K_n) \leq J_1(K_{n-1}) \leq \cdots \leq J_1(K_0),$$

which holds for $n = 0$. In what follows, we define \tilde{K}_{n+1} as

$$(B.4.21) \quad \tilde{K}_{n+1} = K_n - \gamma \cdot (\Upsilon_{K_n}^{22} K_n - \Upsilon_{K_n}^{21}),$$

where Υ_{K_n} is given in (3.2.7). By (B.4.21), we have

$$(B.4.22) \quad \begin{aligned} J_1(\tilde{K}_{n+1}) - J_1(K_n) &= \mathbb{E}_{y \sim \mathcal{N}(0, \Psi_\epsilon)} [y^\top (P_{\tilde{K}_{n+1}} - P_{K_n}) y] \\ &= -2\gamma \cdot \text{tr} [\Phi_{\tilde{K}_{n+1}} \cdot (\Upsilon_{K_n}^{22} K_n - \Upsilon_{K_n}^{21})^\top (\Upsilon_{K_n}^{22} K_n - \Upsilon_{K_n}^{21})] \\ &\quad + \gamma^2 \cdot \text{tr} [\Phi_{\tilde{K}_{n+1}} \cdot (\Upsilon_{K_n}^{22} K_n - \Upsilon_{K_n}^{21})^\top \Upsilon_{K_n}^{22} (\Upsilon_{K_n}^{22} K_n - \Upsilon_{K_n}^{21})] \\ &\leq -2\gamma \cdot \text{tr} [\Phi_{\tilde{K}_{n+1}} \cdot (\Upsilon_{K_n}^{22} K_n - \Upsilon_{K_n}^{21})^\top (\Upsilon_{K_n}^{22} K_n - \Upsilon_{K_n}^{21})] \\ &\quad + \gamma^2 \cdot \|\Upsilon_{K_n}^{22}\|_2 \cdot \text{tr} [\Phi_{\tilde{K}_{n+1}} \cdot (\Upsilon_{K_n}^{22} K_n - \Upsilon_{K_n}^{21})^\top (\Upsilon_{K_n}^{22} K_n - \Upsilon_{K_n}^{21})], \end{aligned}$$

where the first equality comes from (B.4.15), the second equality comes from Lemma B.4.1, and the last inequality comes from the trace inequality. By the definition of Υ_K in (3.2.7), we obtain that

$$(B.4.23) \quad \begin{aligned} \|\Upsilon_{K_n}^{22}\|_2 &\leq \|R\|_2 + \|B\|_2^2 \cdot \|P_{K_n}\|_2 \leq \|R\|_2 + \|B\|_2^2 \cdot J_1(K_n) \cdot \sigma_{\min}^{-1}(\Psi_\epsilon) \\ &\leq \|R\|_2 + \|B\|_2^2 \cdot J_1(K_0) \cdot \sigma_{\min}^{-1}(\Psi_\epsilon), \end{aligned}$$

where the second inequality comes from Proposition B.2.2. By plugging (B.4.23) and the choice of stepsize $\gamma \leq [\|R\|_2 + \|B\|_2^2 \cdot J_1(K_0) \cdot \sigma_{\min}^{-1}(\Psi_\epsilon)]^{-1}$ into (B.4.22), we obtain that

$$\begin{aligned}
J_1(\tilde{K}_{n+1}) - J_1(K_n) &\leq -\gamma \cdot \text{tr}[\Phi_{\tilde{K}_{n+1}} \cdot (\Upsilon_{K_n}^{22} K_n - \Upsilon_{K_n}^{21})^\top (\Upsilon_{K_n}^{22} K_n - \Upsilon_{K_n}^{21})] \\
&\leq -\gamma \cdot \sigma_{\min}(\Psi_\epsilon) \cdot \text{tr}[(\Upsilon_{K_n}^{22} K_n - \Upsilon_{K_n}^{21})^\top (\Upsilon_{K_n}^{22} K_n - \Upsilon_{K_n}^{21})] \\
\text{(B.4.24)} \quad &\leq -\gamma \cdot \sigma_{\min}(\Psi_\epsilon) \cdot \sigma_{\min}(R) \cdot \|\Phi_{K^*}\|_2^{-1} \cdot [J_1(K_n) - J_1(K^*)] < 0,
\end{aligned}$$

where the last inequality comes from Lemma B.4.2.

The following lemma upper bounds $|J_1(\tilde{K}_{n+1}) - J_1(K_{n+1})|$.

Lemma B.4.3. Assume that $J_1(K_n) \leq J_1(K_0)$. It holds with probability at least $1 - \varepsilon^{15}$ that

$$|J_1(\tilde{K}_{n+1}) - J_1(K_{n+1})| \leq \gamma \cdot \sigma_{\min}(\Psi_\epsilon) \cdot \sigma_{\min}(R) \cdot \|\Phi_{K^*}\|_2^{-1} \cdot \varepsilon/4,$$

where K_{n+1} and \tilde{K}_{n+1} are defined in (B.4.18) and (B.4.21), respectively.

Proof. See §B.6.3 for a detailed proof. □

Combining (B.4.24) and Lemma B.4.3, if $J_1(K_n) - J_1(K^*) \geq \varepsilon/2$, it holds with probability at least $1 - \varepsilon^{15}$ that

$$\begin{aligned}
J_1(K_{n+1}) - J_1(K_n) &\leq J_1(\tilde{K}_{n+1}) - J_1(K_n) + |J_1(\tilde{K}_{n+1}) - J_1(K_{n+1})| \\
\text{(B.4.25)} \quad &\leq -\gamma \cdot \sigma_{\min}(\Psi_\epsilon) \cdot \sigma_{\min}(R) \cdot \|\Phi_{K^*}\|_2^{-1} \cdot \varepsilon/4 < 0.
\end{aligned}$$

Combining (B.4.20) and (B.4.25), it holds with probability at least $1 - \varepsilon^{15}$ that

$$J_1(K_{n+1}) \leq J_1(K_n) \leq \cdots \leq J_1(K_0).$$

Finally, following from a union bound argument and the choice of N in Theorem B.2.4, if $J_1(K_n) - J_1(K^*) \geq \varepsilon/2$ holds for any $n \leq N$, we have

$$J_1(K_N) \leq J_1(K_{N-1}) \leq \cdots \leq J_1(K_0),$$

which holds with probability at least $1 - \varepsilon^{13}$. Thus, we complete the proof of (B.4.19).

Combining (B.4.24) and (B.4.25), for $J_1(K_n) - J_1(K^*) \geq \varepsilon/2$, we have

$$J_1(K_{n+1}) - J_1(K^*) \leq [1 - \gamma \cdot \sigma_{\min}(\Psi_\varepsilon) \cdot \sigma_{\min}(R) \cdot \|\Phi_{K^*}\|_2^{-1}] \cdot [J_1(K_n) - J_1(K^*)],$$

which holds with probability at least $1 - \varepsilon^{13}$. Meanwhile, following from a union bound argument and the choice of N in Theorem B.2.4, it holds with probability at least $1 - \varepsilon^{11}$ that

$$(B.4.26) \quad J_1(K_N) - J_1(K^*) \leq \varepsilon/2.$$

The following lemma upper bounds $\|K_N - K^*\|_F$.

Lemma B.4.4. For any K , we have

$$\|K - K^*\|_F^2 \leq \sigma_{\min}^{-1}(\Psi_\varepsilon) \cdot \sigma_{\min}^{-1}(R) \cdot [J_1(K) - J_1(K^*)].$$

Proof. See §B.6.4 for a detailed proof. □

Combining (B.4.26) and Lemma B.4.4, we have

$$(B.4.27) \quad \|K_N - K^*\|_F \leq [\sigma_{\min}^{-1}(\Psi_\epsilon) \cdot \sigma_{\min}^{-1}(R) \cdot \epsilon/2]^{1/2},$$

which holds with probability $1 - \epsilon^{11}$.

Part 2. We show that $J_2(K_N, b_H) - J_2(K^*, b^*) < \epsilon/2$ with high probability. Following from Proposition 3.2.4, it holds that $J_2(K^*, b^*) = J_2(K_N, b^{K_N})$. Therefore, it suffices to show that $J_2(K_N, b_H) - J_2(K_N, b^{K_N}) < \epsilon/2$.

First, we show that if $J_2(K_N, b_h) - J_2(K_N, b^{K_N}) \geq \epsilon/2$ for any $h \leq H$, we obtain that

$$(B.4.28) \quad J_2(K_N, b_H) \leq J_2(K_N, b_{H-1}) \leq \cdots \leq J_2(K_N, b_1) \leq J_2(K_N, b_0),$$

which holds with probability at least $1 - \epsilon^{13}$. We prove (B.4.28) by mathematical induction. Suppose that

$$(B.4.29) \quad J_2(K_N, b_h) \leq J_2(K_N, b_{h-1}) \leq \cdots \leq J_2(K_N, b_0),$$

Recall that by Algorithm 3, the parameter b is updated via

$$(B.4.30) \quad b_{h+1} = b_h - \gamma^b \cdot \widehat{\nabla}_b J_2(K_N, b_h).$$

Here

$$(B.4.31) \quad \widehat{\nabla}_b J_2(K_N, b_h) = \widehat{\Upsilon}_{K_N}^{22}(-K_N \widehat{\mu}_{K_N, b_h} + b_h) + \widehat{\Upsilon}_{K_N}^{21} \widehat{\mu}_{K_N, b_h} + \widehat{q}_{K_N, b_h},$$

where $\widehat{\Upsilon}_{K_N}$ and \widehat{q}_{K_N, b_h} are the outputs of Algorithm 7. We define \widetilde{b}_{h+1} as

$$(B.4.32) \quad \widetilde{b}_{h+1} = b_h - \gamma^b \cdot \nabla_b J_2(K_N, b_h).$$

Here

$$(B.4.33) \quad \nabla_b J_2(K_N, b_h) = \Upsilon_{K_N}^{22} (-K_N \mu_{K_N, b_h} + b_h) + \Upsilon_{K_N}^{21} \mu_{K_N, b_h} + q_{K_N, b_h},$$

where Υ_{K_N} and q_{K_N, b_h} are defined in (3.2.7). We upper bound $J_2(K_N, b_{h+1}) - J_2(K_N, b^{K_N})$ in the sequel. Following from (B.4.32) and Proposition 3.2.3, we have

$$(B.4.34) \quad \begin{aligned} J_2(K_N, \widetilde{b}_{h+1}) - J_2(K_N, b_h) &\leq -\gamma^b/2 \cdot \|\nabla_b J_2(K_N, b_h)\|_2^2 \\ &\leq -\nu_{K_N} \cdot \gamma^b \cdot [J_2(K_N, b_h) - J_2(K_N, b^{K_N})] \\ &\leq -\nu_{K_N} \cdot \gamma^b \cdot \varepsilon < 0, \end{aligned}$$

where ν_{K_N} is specified in Proposition 3.2.3. The following lemma upper bounds $|J_2(K_N, b_{h+1}) - J_2(K_N, \widetilde{b}_{h+1})|$.

Lemma B.4.5. Assume that $J_2(K_N, b_h) \leq J_2(K_N, b_0)$. It holds with probability at least $1 - \varepsilon^{15}$ that

$$|J_2(K_N, b_{h+1}) - J_2(K_N, \widetilde{b}_{h+1})| \leq \nu_{K_N} \cdot \gamma^b \cdot \varepsilon/2,$$

where b_{h+1} and \widetilde{b}_{h+1} are defined in (B.4.30) and (B.4.32), respectively.

Proof. See §B.6.5 for a detailed proof. □

Combining (B.4.34) and Lemma B.4.5, we know that if $J_2(K_N, b_h) - J_2(K_N, b^{K_N}) \geq \varepsilon$, it holds with probability at least $1 - \varepsilon^{15}$ that

$$\begin{aligned} J_2(K_N, b_{h+1}) - J_2(K_N, b_h) &\leq J_2(K_N, \tilde{b}_{h+1}) - J_2(K_N, b_h) + |J_2(K_N, b_{h+1}) - J_2(K_N, \tilde{b}_{h+1})| \\ (B.4.35) \qquad \qquad \qquad &\leq -\nu_{K_N} \cdot \gamma^b \cdot \varepsilon/2 < 0. \end{aligned}$$

Combining (B.4.29) and (B.4.35), it holds with probability at least $1 - \varepsilon^{15}$ that

$$J_2(K_N, b_{h+1}) \leq J_2(K_N, b_h) \leq \cdots \leq J_2(K_N, b_0).$$

Following from a union bound argument and the choice of H in Theorem B.2.4, if $J_2(K_N, b_h) - J_2(K_N, b^{K_N}) \geq \varepsilon$ holds for any $h \leq H$, we have

$$J_2(K_N, b_H) \leq J_2(K_N, b_{H-1}) \leq \cdots \leq J_2(K_N, b_0),$$

which holds with probability at least $1 - \varepsilon^{13}$. Thus, we finish the proof of (B.4.28).

Combining (B.4.34) and Lemma B.4.5, for $J_2(K_N, b_h) - J_2(K_N, b^{K_N}) \geq \varepsilon/2$, we have

$$J_2(K_N, b_{h+1}) - J_2(K_N, b^{K_N}) \leq (1 - \nu_{K_N} \cdot \gamma^b) \cdot [J_2(K_N, b_h) - J_2(K_N, b^{K_N})],$$

which holds with probability at least $1 - \varepsilon^{13}$. Meanwhile, following from a union bound argument and the choice of H in Theorem B.2.4, it holds with probability at least $1 - \varepsilon^{11}$ that

$$(B.4.36) \qquad \qquad \qquad J_2(K_N, b_H) - J_2(K_N, b^{K_N}) \leq \varepsilon/2.$$

By Proposition 3.2.3 and (B.4.36), it holds with probability at least $1 - \varepsilon^{11}$ that

$$(B.4.37) \quad \|b_H - b^{K_N}\|_2 \leq (2\varepsilon/\nu_{K^*})^{1/2}.$$

Following from Proposition 3.2.4, we know that

$$(B.4.38) \quad \begin{aligned} b^{K_N} - b^* &= (K_N - K^*)Q^{-1}(I - A)^\top \\ &\cdot [(I - A)Q^{-1}(I - A)^\top + BR^{-1}B^\top]^{-1} \cdot (\bar{A}\mu + d). \end{aligned}$$

Combining (B.4.27), (B.4.37), and (B.4.38), it holds with probability $1 - \varepsilon^{10}$ that

$$\|b_H - b^{K_N}\|_2 \leq M_b \cdot \varepsilon^{1/2},$$

where

$$\begin{aligned} M_b(\mu) &= 4 \left\| Q^{-1}(I - A)^\top \cdot [(I - A)Q^{-1}(I - A)^\top + BR^{-1}B^\top]^{-1} \cdot (\bar{A}\mu + d) \right\|_2 \\ &\cdot [\nu_{K^*}^{-1} + \sigma_{\min}^{-1}(\Psi_\varepsilon) \cdot \sigma_{\min}^{-1}(R)]^{1/2}. \end{aligned}$$

We finish the proof of the theorem. □

B.4.3. Proof of Theorem B.2.8

Proof. We follow the proof of Theorem 4.2 in Yang et al. (2019b), where they only consider LQR without drift terms. Since our proof requires much more delicate analysis, we present it here.

Part 1. We denote by $\widehat{\zeta}$ and $\widehat{\xi}$ the primal and dual variables generated by Algorithm 7.

We define the primal-dual gap of (B.2.11) as

$$(B.4.39) \quad \text{gap}(\widehat{\zeta}, \widehat{\xi}) = \max_{\xi \in \mathcal{V}_\xi} F(\widehat{\zeta}, \xi) - \min_{\zeta \in \mathcal{V}_\zeta} F(\zeta, \widehat{\xi}).$$

In the sequel, we upper bound $\|\widehat{\alpha}_{K,b} - \alpha_{K,b}\|_2$ using (B.4.39).

We define $\zeta_{K,b}$ and $\xi(\zeta)$ as

$$(B.4.40) \quad \zeta_{K,b} = (J(K, b), \alpha_{K,b}^\top)^\top, \quad \xi(\zeta) = \underset{\xi}{\operatorname{argmax}} F(\zeta, \xi).$$

Following from (B.2.12), we know that

$$(B.4.41) \quad \xi^1(\zeta) = \zeta^1 - J(K, b), \quad \xi^2(\zeta) = \mathbb{E}_{\pi_{K,b}}[\psi(x, u)]\zeta^1 + \Theta_{K,b}\zeta^2 - \mathbb{E}_{\pi_{K,b}}[c(x, u)\psi(x, u)].$$

The following lemma shows that $\zeta_{K,b} \in \mathcal{V}_\zeta$ and $\xi(\zeta) \in \mathcal{V}_\xi$ for any $\zeta \in \mathcal{V}_\zeta$.

Lemma B.4.6. Under the assumptions in the statement of Theorem B.2.8, we have

$\zeta_{K,b} = (J(K, b), \alpha_{K,b}^\top)^\top \in \mathcal{V}_\zeta$. Also, for any $\zeta \in \mathcal{V}_\zeta$, the vector $\xi(\zeta)$ defined in (B.4.40) satisfies that $\xi(\zeta) \in \mathcal{V}_\xi$.

Proof. See §B.6.6 for a detailed proof. □

By (B.2.12), we know that $\nabla_\zeta F(\zeta_{K,b}, 0) = 0$ and $\nabla_\xi F(\zeta_{K,b}, 0) = 0$. Combining Lemma B.4.6, it holds that $(\zeta_{K,b}, 0)$ is a saddle point of the function $F(\zeta, \xi)$ defined in (B.2.11).

Following from (B.4.39), it holds that

$$\begin{aligned}
& \left\| \mathbb{E}_{\pi_{K,b}} [\psi(x, u)] \widehat{\zeta}^1 + \Theta_{K,b} \widehat{\zeta}^2 - \mathbb{E}_{\pi_{K,b}} [c(x, u) \psi(x, u)] \right\|_2^2 + |\widehat{\zeta}^1 - J(K, b)|^2 \\
\text{(B.4.42)} \quad & = F(\widehat{\zeta}, \xi(\widehat{\zeta})) = \max_{\xi \in \mathcal{V}_\xi} F(\widehat{\zeta}, \xi) = \text{gap}(\widehat{\zeta}, \widehat{\xi}) + \min_{\zeta \in \mathcal{V}_\zeta} F(\zeta, \widehat{\xi}),
\end{aligned}$$

where the first equality comes from (B.4.41), and the second equality comes from the fact that $\xi(\widehat{\zeta}) = \text{argmax}_{\xi \in \mathcal{V}_\xi} F(\widehat{\zeta}, \xi)$ by (B.4.40) and Lemma B.4.6. We upper bound the RHS of (B.4.42) and lower bound the LHS of (B.4.42) in the sequel.

As for the RHS of (B.4.42), it holds for any $\xi \in \mathcal{V}_\xi$ that

$$\begin{aligned}
\min_{\zeta \in \mathcal{V}_\zeta} F(\zeta, \xi) & \leq \min_{\zeta \in \mathcal{V}_\zeta} \max_{\xi \in \mathcal{V}_\xi} F(\zeta, \xi) = \min_{\zeta \in \mathcal{V}_\zeta} F(\zeta, \xi(\zeta)) \\
& = \frac{1}{2} \min_{\zeta \in \mathcal{V}_\zeta} \left\{ \left\| \mathbb{E}_{\pi_{K,b}} [\psi(x, u)] \zeta^1 + \Theta_{K,b} \zeta^2 - \mathbb{E}_{\pi_{K,b}} [c(x, u) \psi(x, u)] \right\|_2^2 + |\zeta^1 - J(K, b)|^2 \right\} \\
\text{(B.4.43)} \quad & = 0,
\end{aligned}$$

where the first equality comes from the fact that $\xi(\zeta) = \text{argmax}_{\xi \in \mathcal{V}_\xi} F(\zeta, \xi)$ by (B.4.40) and Lemma B.4.6, the second equality comes from (B.4.41), and the last equality holds by taking $\zeta = \zeta_{K,b} \in \mathcal{V}_\zeta$. Meanwhile, we lower bound the LHS of (B.4.42) as

$$\begin{aligned}
& \left\| \mathbb{E}_{\pi_{K,b}} [\psi(x, u)] \widehat{\zeta}^1 + \Theta_{K,b} \widehat{\zeta}^2 - \mathbb{E}_{\pi_{K,b}} [c(x, u) \psi(x, u)] \right\|_2^2 + |\widehat{\zeta}^1 - J(K, b)|^2 \\
\text{(B.4.44)} \quad & = \|\widetilde{\Theta}_{K,b}(\widehat{\zeta} - \zeta_{K,b})\|_2^2 \geq \lambda_K^2 \cdot \|\widehat{\zeta} - \zeta_{K,b}\|_2^2 \geq \lambda_K^2 \cdot \|\widehat{\alpha}_{K,b} - \alpha_{K,b}\|_2^2,
\end{aligned}$$

where the first equality comes from the definition of $\widetilde{\Theta}_{K,b}$ in (B.2.9), and the first inequality comes from Proposition B.2.6. Here λ_K is defined in Proposition B.2.6. Combining

(B.4.42), (B.4.43), and (B.4.44), it holds that

$$(B.4.45) \quad \|\widehat{\alpha}_{K,b} - \alpha_{K,b}\|_2^2 \leq \lambda_K^{-2} \cdot \text{gap}(\widehat{\zeta}, \widehat{\xi}),$$

which finishes the proof of this part.

Part 2. We now upper bound $\text{gap}(\widehat{\zeta}, \widehat{\xi})$. We denote by $\widetilde{z}_t = (\widetilde{x}_t^\top, \widetilde{u}_t^\top)^\top$ for $t \in [\widetilde{T}]$, where \widetilde{x}_t and \widetilde{u}_t are generated in Line 6 of Algorithm 7. Following from the state transition in Problem 3.1.3 and the form of the linear policy, $\{\widetilde{z}_t\}_{t \in [\widetilde{T}]}$ follows the following transition,

$$(B.4.46) \quad \widetilde{z}_{t+1} = L\widetilde{z}_t + \nu + \delta_t,$$

where

$$\nu = \begin{pmatrix} \bar{A}\mu + d \\ -K(\bar{A}\mu + d) + b \end{pmatrix}, \quad \delta_t = \begin{pmatrix} \omega_t \\ -K\omega_t + \sigma\eta \end{pmatrix}, \quad L = \begin{pmatrix} A & B \\ -KA & -KB \end{pmatrix}.$$

Note that we have

$$L = \begin{pmatrix} A & B \\ -KA & -KB \end{pmatrix} = \begin{pmatrix} I \\ -K \end{pmatrix} \begin{pmatrix} A & B \end{pmatrix}.$$

Then by the property of spectral radius, it holds that

$$\rho(L) = \rho\left(\begin{pmatrix} A & B \\ -K \end{pmatrix} \begin{pmatrix} I \\ -K \end{pmatrix}\right) = \rho(A - BK) < 1.$$

Thus, the Markov chain generated by (B.4.46) admits a unique stationary distribution $\mathcal{N}(\mu_z, \Sigma_z)$, where

$$(B.4.47) \quad \mu_z = (I - L)^{-1}\nu, \quad \Sigma_z = L\Sigma_z L^\top + \begin{pmatrix} \Psi_\omega & -\Psi_\omega K^\top \\ -K\Psi_\omega & K\Psi_\omega K^\top + \sigma^2 I \end{pmatrix}.$$

The following lemma characterizes the average

$$(B.4.48) \quad \hat{\mu}_z = 1/\tilde{T} \cdot \sum_{t=1}^{\tilde{T}} \tilde{z}_t.$$

Lemma B.4.7. It holds that

$$\hat{\mu}_z \sim \mathcal{N}\left(\mu_z + \frac{1}{\tilde{T}}\mu_{\tilde{T}}, \frac{1}{\tilde{T}}\tilde{\Sigma}_{\tilde{T}}\right),$$

where $\|\mu_{\tilde{T}}\|_2 \leq M_\mu \cdot (1 - \rho)^{-2} \cdot \|\mu_z\|_2$ and $\|\tilde{\Sigma}_{\tilde{T}}\|_F \leq M_\Sigma \cdot (1 - \rho)^{-1} \cdot \|\Sigma_z\|_F$. Here M_μ and M_Σ are positive absolute constants. Moreover, it holds with probability at least $1 - \tilde{T}^{-6}$ that

$$\|\hat{\mu}_z - \mu_z\|_2 \leq \frac{\log \tilde{T}}{\tilde{T}^{1/4}} \cdot (1 - \rho)^{-2} \cdot \text{poly}(\|\Phi_K\|_2, \|K\|_F, \|b\|_2, \|\mu\|_2).$$

Proof. See §B.6.7 for a detailed proof. □

Lemma B.4.7 gives that

$$\|\hat{\mu}_{K,b} - \mu_{K,b}\|_2 \leq \frac{\log \tilde{T}}{\tilde{T}^{1/4}} \cdot (1 - \rho)^{-2} \cdot \text{poly}(\|\Phi_K\|_2, \|K\|_F, \|b\|_2, \|\mu\|_2),$$

which holds with probability at least $1 - \tilde{T}^{-6}$.

We now apply a truncation argument to show that $\text{gap}(\widehat{\zeta}, \widehat{\xi})$ is upper bounded. We define the event \mathcal{E} in the sequel. Following from Lemma B.4.7, it holds for any $z \sim \mathcal{N}(\mu_z, \Sigma_z)$ that

$$z - \widehat{\mu}_z + 1/\widetilde{T} \cdot \mu_{\widetilde{T}} \sim \mathcal{N}(0, \Sigma_z + 1/\widetilde{T} \cdot \widetilde{\Sigma}_{\widetilde{T}}).$$

By Lemma B.7.3, there exists a positive absolute constant C_0 such that

(B.4.49)

$$\mathbb{P}\left[\left|\|z - \widehat{\mu}_z + 1/\widetilde{T} \cdot \mu_{\widetilde{T}}\|_2^2 - \text{tr}(\widetilde{\Sigma}_z)\right| > \tau\right] \leq 2 \exp\left[-C_0 \cdot \min(\tau^2 \|\widetilde{\Sigma}_z\|_{\text{F}}^{-2}, \tau \|\widetilde{\Sigma}_z\|_2^{-1})\right],$$

where we write $\widetilde{\Sigma}_z = \Sigma_z + 1/\widetilde{T} \cdot \widetilde{\Sigma}_{\widetilde{T}}$ for notational convenience. By taking $\tau = C_1 \cdot \log T \cdot \|\widetilde{\Sigma}_z\|_{\text{F}}$ in (B.4.49) for a sufficiently large positive absolute constant C_1 , it holds that

$$(B.4.50) \quad \mathbb{P}\left[\left|\|z - \widehat{\mu}_z + 1/\widetilde{T} \cdot \mu_{\widetilde{T}}\|_2^2 - \text{tr}(\widetilde{\Sigma}_z)\right| > C_1 \cdot \log T \cdot \|\widetilde{\Sigma}_z\|_{\text{F}}\right] \leq T^{-6}.$$

We define the event $\mathcal{E}_{t,1}$ for any $t \in [T]$ as

$$\mathcal{E}_{t,1} = \left\{ \left| \|z_t - \widehat{\mu}_z + 1/\widetilde{T} \cdot \mu_{\widetilde{T}}\|_2^2 - \text{tr}(\widetilde{\Sigma}_z) \right| \leq C_1 \cdot \log T \cdot \|\widetilde{\Sigma}_z\|_{\text{F}} \right\}.$$

Then by (B.4.50), it holds for any $t \in [T]$ that

$$(B.4.51) \quad \mathbb{P}(\mathcal{E}_{t,1}) \geq 1 - T^{-6}.$$

Also, we define

$$(B.4.52) \quad \mathcal{E}_1 = \bigcap_{t \in [T]} \mathcal{E}_{t,1}.$$

Following from a union bound argument and (B.4.51), it holds that

$$(B.4.53) \quad \mathbb{P}(\mathcal{E}_1) \geq 1 - T^{-5}.$$

Also, conditioning on \mathcal{E}_1 , it holds for sufficiently large \tilde{T} that

$$\begin{aligned} & \max_{t \in [T]} \|z_t - \hat{\mu}_z\|_2^2 \\ & \leq C_1 \cdot \log T \cdot \|\tilde{\Sigma}_z\|_F + \text{tr}(\tilde{\Sigma}_z) + \|1/\tilde{T} \cdot \mu_{\tilde{T}}\|_2^2 \\ & \leq 2\tilde{C}_1 \cdot [1 + M_\Sigma(1 - \rho)^{-1}/\tilde{T}^2] \cdot \log T \cdot \|\Sigma_z\|_2 + M_\mu(1 - \rho)^{-2}/\tilde{T}^2 \cdot \|\mu_z\|_2^2 \\ & \leq C_2 \cdot \log T \cdot (1 + \|K\|_F^2) \cdot \|\Phi_K\|_2 \cdot (1 - \rho)^{-1} + C_3 \cdot (\|b\|_2^2 + \|\mu\|_2^2) \cdot (1 - \rho)^{-4} \cdot \tilde{T}^{-2} \end{aligned}$$

(B.4.54)

$$\leq 2C_2 \cdot \log T \cdot (1 + \|K\|_F^2) \cdot \|\Phi_K\|_2 \cdot (1 - \rho)^{-1},$$

where \tilde{C}_1 , C_2 , and C_3 are positive absolute constants. Here, the first inequality comes from the definition of \mathcal{E}_1 in (B.4.52), the second inequality comes from Lemma B.4.7, and the third inequality comes from (B.4.47). Also, we define the following event

$$(B.4.55) \quad \mathcal{E}_2 = \{\|\hat{\mu}_z - \mu_z + 1/\tilde{T} \cdot \mu_{\tilde{T}}\|_2 \leq C_1\}.$$

Then by Lemma B.4.7, we know that

$$(B.4.56) \quad \mathbb{P}(\mathcal{E}_2) \geq 1 - \tilde{T}^{-6}$$

for \tilde{T} sufficiently large. We define the event \mathcal{E} as

$$\mathcal{E} = \mathcal{E}_1 \cap \mathcal{E}_2.$$

Then following from (B.4.53), (B.4.56), and a union bound argument, we know that

$$\mathbb{P}(\mathcal{E}) \geq 1 - T^{-5} - \tilde{T}^{-6}.$$

Now, we define the truncated feature vector $\tilde{\psi}(x, u)$ as $\tilde{\psi}(x, u) = \hat{\psi}(x, u) \mathbf{1}_{\mathcal{E}}$, the truncated cost function $\tilde{c}(x, u)$ as $\tilde{c}(x, u) = c(x, u) \mathbf{1}_{\mathcal{E}}$, and also the truncated objective function $\tilde{F}(\zeta, \xi)$ as

(B.4.57)

$$\tilde{F}(\zeta, \xi) = \left\{ \mathbb{E}(\tilde{\psi})\zeta^1 + \mathbb{E}[(\tilde{\psi} - \tilde{\psi}')\tilde{\psi}^\top]\zeta^2 - \mathbb{E}(\tilde{c}\tilde{\psi}) \right\}^\top \xi^2 + [\zeta^1 - \mathbb{E}(\tilde{c})] \cdot \xi^1 - \|\xi\|_2^2/2,$$

where we write $\tilde{\psi} = \tilde{\psi}(x, u)$ and $\tilde{c} = \tilde{c}(x, u)$ for notational convenience. Here the expectation is taken following the policy $\pi_{K,b}$ and the state transition. The following lemma establishes the upper bound of $|F(\zeta, \xi) - \tilde{F}(\zeta, \xi)|$, where $F(\zeta, \xi)$ and $\tilde{F}(\zeta, \xi)$ are defined in (B.2.11) and (B.4.57), respectively.

Lemma B.4.8. It holds with probability at least $1 - \tilde{T}^{-6}$ that

$$|F(\zeta, \xi) - \tilde{F}(\zeta, \xi)| \leq \left(\frac{1}{2T} + \frac{\log \tilde{T}}{\tilde{T}^{1/4}} \right) \cdot (1 - \rho)^{-2} \cdot \text{poly}(\|K\|_F, \|b\|_2, \|\mu\|_2, J(K_0, b_0)).$$

Proof. See §B.6.8 for a detailed proof. □

Following from (B.4.39) and Lemma B.4.8, it holds with probability at least $1 - \tilde{T}^{-6}$ that

$$(B.4.58) \quad \begin{aligned} & \left| \text{gap}(\widehat{\zeta}, \widehat{\xi}) - \widetilde{\text{gap}}(\widehat{\zeta}, \widehat{\xi}) \right| \\ & \leq \left(\frac{1}{2T} + \frac{\log \tilde{T}}{\tilde{T}^{1/4}} \right) \cdot (1 - \rho)^{-2} \cdot \text{poly}(\|K\|_{\mathbb{F}}, \|b\|_2, \|\mu\|_2, J(K_0, b_0)). \end{aligned}$$

where we define $\widetilde{\text{gap}}(\widehat{\zeta}, \widehat{\xi})$ as

$$\widetilde{\text{gap}}(\widehat{\zeta}, \widehat{\xi}) = \max_{\xi \in \mathcal{V}_{\widehat{\xi}}} \widetilde{F}(\widehat{\zeta}, \xi) - \min_{\zeta \in \mathcal{V}_{\widehat{\zeta}}} \widetilde{F}(\zeta, \widehat{\xi}).$$

Therefore, to upper bound of $\text{gap}(\zeta, \xi)$, we only need to upper bound $\widetilde{\text{gap}}(\zeta, \xi)$.

Part 3. We upper bound $\widetilde{\text{gap}}(\zeta, \xi)$ in the sequel. We first show that the trajectory generated by the policy $\pi_{K,b}$ and the state transition in Problem 3.1.2 is β -mixing.

Lemma B.4.9. Consider a linear system $y_{t+1} = Dy_t + \vartheta + v_t$, where $\{y_t\}_{t \geq 0} \subset \mathbb{R}^m$, the matrix $D \in \mathbb{R}^{m \times m}$ satisfying $\rho(D) < 1$, the vector $\vartheta \in \mathbb{R}^m$, and $v_t \sim \mathcal{N}(0, \Sigma)$ is the Gaussians. We denote by ϖ_t the marginal distribution of y_t for any $t \geq 0$. Meanwhile, assume that the stationary distribution of $\{y_t\}_{t \geq 0}$ is a Gaussian distribution $\mathcal{N}((I - D)^{-1}\vartheta, \Sigma_{\infty})$, where Σ_{∞} is the covariance matrix. We define the β -mixing coefficients for any $n \geq 1$ as follows

$$\beta(n) = \sup_{t \geq 0} \mathbb{E}_{y \sim \varpi_t} \left[\left\| \mathbb{P}_{y_n}(\cdot | y_0 = y) - \mathbb{P}_{\mathcal{N}((I-D)^{-1}\vartheta, \Sigma_{\infty})}(\cdot) \right\|_{\text{TV}} \right].$$

Then, for any $\rho \in (\rho(D), 1)$, the β -mixing coefficients satisfy that

$$\beta(n) \leq C_{\rho, D, \vartheta} \cdot [\text{tr}(\Sigma_{\infty}) + m \cdot (1 - \rho)^{-2}]^{1/2} \cdot \rho^n,$$

where $C_{\rho, D, \vartheta}$ is a constant, which only depends on ρ , D , and ϑ . We say that the sequence $\{y_t\}_{t \geq 0}$ is β -mixing with parameter ρ .

Proof. See Proposition 3.1 in Tu and Recht (2017) for details. \square

Recall that by (3.2.1), the sequence $\{x_t\}_{t \geq 0}$ follows

$$x_{t+1} = (A - BK)x_t + (Bb + \bar{A}\mu + d) + \epsilon_t, \quad \epsilon_t \sim \mathcal{N}(0, \Psi_\epsilon),$$

where the matrix $A - BK$ satisfies that $\rho(A - BK) < 1$. Therefore, by Lemma B.4.9, the sequence $\{z_t\}_{t \geq 0}$ is β -mixing with parameter $\rho \in (\rho(A - BK), 1)$, where $z_t = (x_t^\top, u_t^\top)^\top$. The following lemma upper bounds the primal-dual gap for a convex-concave problem.

Lemma B.4.10. Let \mathcal{X} and \mathcal{Y} be two compact and convex sets such that $\|x - x'\|_2 \leq M$ and $\|y - y'\|_2 \leq M$ for any $x, x' \in \mathcal{X}$ and $y, y' \in \mathcal{Y}$. We consider solving the following minimax problem

$$\min_{x \in \mathcal{X}} \max_{y \in \mathcal{Y}} H(x, y) = \mathbb{E}_{\epsilon \sim \varpi_\epsilon} [G(x, y; \epsilon)],$$

where the objective function $H(x, y)$ is convex in x and concave in y . In addition, we assume that the distribution ϖ_ϵ is β -mixing with $\beta(n) \leq C_\epsilon \cdot \rho^n$, where C_ϵ is a constant. Meanwhile, we assume that it holds almost surely that $G(x, y; \epsilon)$ is \tilde{L}_0 -Lipschitz in both x and y , the gradient $\nabla_x G(x, y; \epsilon)$ is \tilde{L}_1 -Lipschitz in x for any $y \in \mathcal{Y}$, the gradient $\nabla_y G(x, y; \epsilon)$ is \tilde{L}_1 -Lipschitz in y for any $x \in \mathcal{X}$, where $C_\epsilon, \tilde{L}_0, \tilde{L}_1 > 1$. Each step of our gradient-based method takes the following forms,

$$x_{t+1} = \Gamma_{\mathcal{X}} [x_t - \gamma_{t+1} \cdot \nabla_x G(x_t, y_t; \epsilon_t)], \quad y_{t+1} = \Gamma_{\mathcal{Y}} [y_t - \gamma_{t+1} \cdot \nabla_y G(x_t, y_t; \epsilon_t)],$$

where the operators $\Gamma_{\mathcal{X}}$ and $\Gamma_{\mathcal{Y}}$ projects the variables back to \mathcal{X} and \mathcal{Y} , respectively, and the stepsizes take the form $\gamma_t = \gamma_0 \cdot t^{-1/2}$ for a constant $\gamma_0 > 0$. Moreover, let $\hat{x} = (\sum_{t=1}^T \gamma_t)^{-1}(\sum_{t=1}^T \gamma_t x_t)$ and $\hat{y} = (\sum_{t=1}^T \gamma_t)^{-1}(\sum_{t=1}^T \gamma_t y_t)$ be the final output of the gradient method after T iterations, then there exists a positive absolute constant C , such that for any $\delta \in (0, 1)$, the primal-dual gap to the minimax problem is upper bounded as

$$\max_{x \in \mathcal{X}} H(\hat{x}, y) - \min_{y \in \mathcal{Y}} H(x, \hat{y}) \leq \frac{C \cdot (M^2 + \tilde{L}_0^2 + \tilde{L}_0 \tilde{L}_1 M)}{\log(1/\rho)} \cdot \frac{\log^2 T + \log(1/\delta)}{\sqrt{T}} + \frac{C \cdot C_\epsilon \tilde{L}_0 M}{T},$$

which holds with probability at least $1 - \delta$.

Proof. See Theorem 5.4 in Yang et al. (2019b) for details. \square

To use Lemma B.4.10, we define the function $G(\zeta, \xi; \tilde{\psi}, \tilde{\psi}')$ as

$$G(\zeta, \xi; \tilde{\psi}, \tilde{\psi}') = [\tilde{\psi} \zeta^1 + (\tilde{\psi} - \tilde{\psi}') \tilde{\psi}^\top \zeta^2 - \tilde{c} \tilde{\psi}]^\top \xi^2 + (\zeta^1 - \tilde{c}) \cdot \xi^1 - 1/2 \cdot \|\xi\|_2^2,$$

where $\tilde{\psi} = \tilde{\psi}(x, u)$ and $\tilde{\psi}' = \tilde{\psi}(x', u')$. Note that the gradients of $G(\zeta, \xi; \tilde{\psi}, \tilde{\psi}')$ take the form

$$\begin{aligned} \nabla_\zeta G(\zeta, \xi; \tilde{\psi}, \tilde{\psi}') &= \begin{pmatrix} \tilde{\psi}^\top \xi^2 + \xi^1 \\ \tilde{\psi}(\tilde{\psi} - \tilde{\psi}')^\top \xi^2 \end{pmatrix}, \\ \nabla_\xi G(\zeta, \xi; \tilde{\psi}, \tilde{\psi}') &= \begin{pmatrix} \zeta^1 - \tilde{c} - \xi^1 \\ \tilde{\psi} \zeta^1 + (\tilde{\psi} - \tilde{\psi}') \tilde{\psi}^\top \zeta^2 - \tilde{c} \tilde{\psi} - \xi^2 \end{pmatrix}. \end{aligned}$$

By Definition B.2.7 and Lemma B.4.6, we know that

$$(B.4.59) \quad \begin{aligned} \|\nabla_{\zeta} G(\zeta, \xi; \tilde{\psi}, \tilde{\psi}')\|_2 &\leq \text{poly}(\|K\|_{\mathbb{F}}, J(K_0, b_0)) \cdot \log^2 T \cdot (1 - \rho)^{-2}, \\ \|\nabla_{\xi} G(\zeta, \xi; \tilde{\psi}, \tilde{\psi}')\|_2 &\leq \text{poly}(\|K\|_{\mathbb{F}}, \|\mu\|_2, J(K_0, b_0)) \cdot \log^2 T \cdot (1 - \rho)^{-2}. \end{aligned}$$

This gives the Lipschitz constant \tilde{L}_0 in Lemma B.4.10 for $G(\zeta, \xi; \tilde{\psi}, \tilde{\psi}')$. Also, the Hessians of $G(\zeta, \xi; \tilde{\psi}, \tilde{\psi}')$ take the forms of

$$\nabla_{\zeta\zeta}^2 G(\zeta, \xi; \tilde{\psi}, \tilde{\psi}') = 0, \quad \nabla_{\xi\xi}^2 G(\zeta, \xi; \tilde{\psi}, \tilde{\psi}') = -I,$$

which follows that

$$(B.4.60) \quad \|\nabla_{\zeta\zeta}^2 G(\zeta, \xi; \tilde{\psi}, \tilde{\psi}')\|_2 = 0, \quad \|\nabla_{\xi\xi}^2 G(\zeta, \xi; \tilde{\psi}, \tilde{\psi}')\|_2 = 1.$$

This gives the Lipschitz constant \tilde{L}_1 in Lemma B.4.10 for $\nabla_{\zeta} G(\zeta, \xi; \tilde{\psi}, \tilde{\psi}')$ and $\nabla_{\xi} G(\zeta, \xi; \tilde{\psi}, \tilde{\psi}')$. Moreover, note that (B.4.54) provides an upper bound of M , combining (B.4.59), (B.4.60) and Lemma B.4.10, it holds with probability at least $1 - T^{-5}$ that

$$(B.4.61) \quad \widetilde{\text{gap}}(\hat{\zeta}, \hat{\xi}) \leq \frac{\text{poly}(\|K\|_{\mathbb{F}}, \|\mu\|_2, J(K_0, b_0)) \cdot \log^6 T}{(1 - \rho)^4 \cdot \sqrt{T}}.$$

Combining (B.4.45), (B.4.58), and (B.4.61), we know that

$$\begin{aligned} &\|\hat{\alpha}_{K,b} - \alpha_{K,b}\|_2^2 \\ &\leq \lambda_K^{-2} \cdot \text{poly}_1(\|K\|_{\mathbb{F}}, \|b\|_2, \|\mu\|_2, J(K_0, b_0)) \cdot \left[\frac{\log^6 T}{T^{1/2} \cdot (1 - \rho)^4} + \frac{\log \tilde{T}}{\tilde{T}^{1/4} \cdot (1 - \rho)^2} \right]. \end{aligned}$$

Same bounds for $\|\widehat{\Upsilon}_K - \Upsilon_K\|_{\mathbb{F}}^2$, $\|\widehat{p}_{K,b} - p_{K,b}\|_2^2$, and $\|\widehat{q}_{K,b} - q_{K,b}\|_2^2$ hold. We finish the proof of the theorem. \square

B.5. Proofs of Propositions

B.5.1. Proof of Proposition 3.2.2

Proof. We follow a similar proof as in the one of Theorem 1.1 in Sznitman (1991) and Theorem 3.2 in Bensoussan et al. (2016). Note that for any policy $\pi_{K,b} \in \Pi$, the parameters K and b uniquely determine the policy. We define the following metric on Π .

Definition B.5.1. For any $\pi_{K_1,b_1}, \pi_{K_2,b_2} \in \Pi$, we define the following metric,

$$\|\pi_{K_1,b_1} - \pi_{K_2,b_2}\|_2 = c_1 \cdot \|K_1 - K_2\|_2 + c_2 \cdot \|b_1 - b_2\|_2,$$

where c_1 and c_2 are positive constants.

One can verify that Definition B.5.1 satisfies the requirement of being a metric. We first evaluate the forms of the operators $\Lambda_1(\cdot)$ and $\Lambda_2(\cdot, \cdot)$.

Forms of the operators $\Lambda_1(\cdot)$ and $\Lambda_2(\cdot, \cdot)$. By the definition of $\Lambda_1(\mu)$, which gives the optimal policy under the mean-field state μ , it holds that

$$\Lambda_1(\mu) = \pi_{\mu}^*,$$

where π_{μ}^* solves Problem 3.1.2. This gives the form of $\Lambda_1(\cdot)$. We now turn to $\Lambda_2(\mu, \pi)$, which gives the mean-field state μ_{new} generated by the policy π under the current mean-field state μ . In Problem 3.1.2, the sequence of states $\{x_t\}_{t \geq 0}$ constitutes a Markov chain, which admits a unique stationary distribution. Thus, by the state transition in Problem

3.1.2 and the form of the linear-Gaussian policy, we have

$$(B.5.1) \quad \mu_{\text{new}} = (A - BK_{\pi})\mu_{\text{new}} + (Bb_{\pi} + \bar{A}\mu + d),$$

where K_{π} and b_{π} are parameters of the policy π . By solving (B.5.1) for μ_{new} , it holds that

$$\Lambda_2(\mu, \pi) = \mu_{\text{new}} = (I - A + BK_{\pi})^{-1}(Bb_{\pi} + \bar{A}\mu + d).$$

This gives the form of $\Lambda_2(\cdot, \cdot)$.

Next, we compute the Lipschitz constants for $\Lambda_1(\cdot)$ and $\Lambda_2(\cdot, \cdot)$.

Lipschitz constant for $\Lambda_1(\cdot)$. By Proposition 3.2.4, for any $\mu_1, \mu_2 \in \mathbb{R}^m$, the optimal K^* is fixed for Problem 3.1.2. Therefore, by the form of the optimal b^K given in Proposition 3.2.4, it holds that

$$(B.5.2) \quad \begin{aligned} \|\Lambda_1(\mu_1) - \Lambda_1(\mu_2)\|_2 &\leq c_2 \cdot \left\| \left[(I - A)Q^{-1}(I - A)^{\top} + BR^{-1}B^{\top} \right]^{-1} \bar{A} \right\|_2 \\ &\quad \cdot \left\| \left[K^*Q^{-1}(I - A)^{\top} - R^{-1}B^{\top} \right] \right\|_2 \cdot \|\mu_1 - \mu_2\|_2 \\ &= c_2 L_1 \cdot \|\mu_1 - \mu_2\|_2, \end{aligned}$$

where L_1 is defined in Assumption 3.2.1.

Lipschitz constants for $\Lambda_2(\cdot, \cdot)$. By Proposition 3.2.4, for any $\mu_1, \mu_2 \in \mathbb{R}^m$, the optimal K^* is fixed for Problem 3.1.2. Thus, for any $\pi \in \Pi$ such that π is an optimal policy under

some $\mu \in \mathbb{R}^m$, it holds that

$$\begin{aligned}
\|\Lambda_2(\mu_1, \pi) - \Lambda_2(\mu_2, \pi)\|_2 &= \|(I - A + BK_\pi)^{-1} \cdot \bar{A} \cdot (\mu_1 - \mu_2)\|_2 \\
&\leq [1 - \rho(A - BK^*)]^{-1} \cdot \|\bar{A}\|_2 \cdot \|\mu_1 - \mu_2\|_2 \\
\text{(B.5.3)} \qquad \qquad \qquad &= L_2 \cdot \|\mu_1 - \mu_2\|_2,
\end{aligned}$$

where L_2 is defined in Assumption 3.2.1, and $K_\pi = K^*$ is the parameter of the policy π . Meanwhile, for any mean-field state $\mu \in \mathbb{R}^m$, and any policies $\pi_1, \pi_2 \in \Pi$ that are optimal under some mean-field states μ_1, μ_2 , respectively, we have

$$\begin{aligned}
\|\Lambda_2(\mu, \pi_1) - \Lambda_2(\mu, \pi_2)\|_2 &= \|(I - A + BK^*)^{-1} B \cdot (b_{\pi_1} - b_{\pi_2})\|_2 \\
&\leq [1 - \rho(A - BK^*)]^{-1} \cdot \|B\|_2 \cdot \|b_{\pi_1} - b_{\pi_2}\|_2 \\
\text{(B.5.4)} \qquad \qquad \qquad &= c_2^{-1} L_3 \cdot \|\pi_1 - \pi_2\|_2,
\end{aligned}$$

where in the last equality, we use the fact that $K_{\pi_1} = K_{\pi_2} = K^*$ by Proposition 3.2.4. Here L_3 is defined in Assumption 3.2.1, and b_{π_1} and b_{π_2} are the parameters of the policies π_1 and π_2 .

Now we show that the operator $\Lambda(\cdot)$ is a contraction. For any $\mu_1, \mu_2 \in \mathbb{R}^m$, it holds that

$$\begin{aligned}
\|\Lambda(\mu_1) - \Lambda(\mu_2)\|_2 &= \left\| \Lambda_2(\mu_1, \Lambda_1(\mu_1)) - \Lambda_2(\mu_2, \Lambda_1(\mu_2)) \right\|_2 \\
&\leq \left\| \Lambda_2(\mu_1, \Lambda_1(\mu_1)) - \Lambda_2(\mu_1, \Lambda_1(\mu_2)) \right\|_2 + \left\| \Lambda_2(\mu_1, \Lambda_1(\mu_2)) - \Lambda_2(\mu_2, \Lambda_1(\mu_2)) \right\|_2 \\
&\leq c_2^{-1} L_3 \cdot \|\Lambda_1(\mu_1) - \Lambda_1(\mu_2)\|_2 + L_2 \cdot \|\mu_1 - \mu_2\|_2 \\
&\leq c_2^{-1} L_3 \cdot c_2 L_1 \cdot \|\mu_1 - \mu_2\|_2 + L_2 \cdot \|\mu_1 - \mu_2\|_2 = (L_1 L_3 + L_2) \cdot \|\mu_1 - \mu_2\|_2,
\end{aligned}$$

where the first inequality comes from triangle inequality, the second inequality comes from (B.5.3) and (B.5.4), and the last inequality comes from (B.5.2). By Assumption 3.2.1, we know that $L_0 = L_1 L_3 + L_2 < 1$, which shows that the operator $\Lambda(\cdot)$ is a contraction. Moreover, by Banach fixed-point theorem, we obtain that $\Lambda(\cdot)$ has a unique fixed point, which gives the unique equilibrium pair of Problem 3.1.3. We finish the proof of the proposition. \square

B.5.2. Proof of Proposition 3.2.4

Proof. By the definition of $J_2(K, b)$ in (3.2.6) and the definition of $\mu_{K,b}$ in (3.2.2), the problem

$$\min_b J_2(K, b)$$

is equivalent to the following constrained problem,

$$(B.5.5) \quad \min_{\mu, b} \begin{pmatrix} \mu \\ b \end{pmatrix}^\top \begin{pmatrix} Q + K^\top RK & -K^\top R \\ -RK & R \end{pmatrix} \begin{pmatrix} \mu \\ b \end{pmatrix}$$

s.t. $(I - A + BK)\mu - (Bb + \bar{A}\mu + d) = 0.$

Following from the KKT conditions of (B.5.5), it holds that

$$(B.5.6) \quad 2M_K \begin{pmatrix} \mu \\ b \end{pmatrix} + N_K \lambda = 0, \quad N_K^\top \begin{pmatrix} \mu \\ b \end{pmatrix} + \bar{A}\mu + d = 0,$$

where

$$M_K = \begin{pmatrix} Q + K^\top RK & -K^\top R \\ -RK & R \end{pmatrix}, \quad N_K = \begin{pmatrix} -(I - A + BK)^\top \\ B^\top \end{pmatrix}.$$

By solving (B.5.6), the minimizer of (B.5.5) takes the form of

$$(B.5.7) \quad \begin{pmatrix} \mu_{K, b^K} \\ b^K \end{pmatrix} = -M_K^{-1} N_K (N_K^\top M_K^{-1} N_K)^{-1} (\bar{A}\mu + d).$$

By substituting (B.5.7) into the definition of $J_2(K, b)$ in (3.2.6), we have

$$(B.5.8) \quad J_2(K, b^K) = (\bar{A}\mu + d)^\top (N_K^\top M_K^{-1} N_K)^{-1} (\bar{A}\mu + d).$$

Meanwhile, by calculation, we have

$$M_K^{-1} = \begin{pmatrix} Q^{-1} & Q^{-1}K^\top \\ KQ^{-1} & KQ^{-1}K^\top + R^{-1} \end{pmatrix}.$$

Therefore, the term $N_K^\top M_K^{-1} N_K$ in (B.5.8) takes the form of

$$(B.5.9) \quad N_K^\top M_K^{-1} N_K = (I - A)Q^{-1}(I - A^\top) + BR^{-1}B^\top.$$

By plugging (B.5.9) into (B.5.8), we have

$$J_2(K, b^K) = (\bar{A}\mu + d)^\top [(I - A)Q^{-1}(I - A^\top) + BR^{-1}B^\top]^{-1} (\bar{A}\mu + d).$$

Also, by plugging (B.5.9) into (B.5.7), we have

$$\begin{pmatrix} \mu_{K, b^K} \\ b^K \end{pmatrix} = \begin{pmatrix} Q^{-1}(I - A)^\top \\ KQ^{-1}(I - A)^\top - R^{-1}B^\top \end{pmatrix} [(I - A)Q^{-1}(I - A)^\top + BR^{-1}B^\top]^{-1} (\bar{A}\mu + d).$$

We finish the proof of the proposition. \square

B.5.3. Proof of Proposition B.2.2

Proof. By the definition of the cost function $c(x, u)$ in Problem 3.1.2 (recall that we drop the subscript μ when we focus on Problem 3.1.2), we have

$$\begin{aligned} \mathbb{E}c_t &= \mathbb{E}(x_t^\top Qx_t + u_t^\top Ru_t + \mu^\top \bar{Q}\mu) \\ &= \mathbb{E}(x_t^\top Qx_t + x_t^\top K^\top RKx_t - 2b^\top RKx_t + b^\top Rb + \sigma^2 \eta_t^\top R\eta_t + \mu^\top \bar{Q}\mu) \\ (B.5.10) \quad &= \mathbb{E}[x_t^\top (Q + K^\top RK)x_t - 2b^\top RKx_t] + b^\top Rb + \sigma^2 \cdot \text{tr}(R) + \mu^\top \bar{Q}\mu, \end{aligned}$$

where we write $c_t = c(x_t, u_t)$ for notational convenience. Here in the second line we use $u_t = \pi_{K,b}(x_t) = -Kx_t + b + \sigma\eta_t$. Therefore, combining (B.5.10) and the definition of $J(K, b)$ in Problem 3.1.2, we have

$$\begin{aligned}
J(K, b) &= \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^T \left\{ \mathbb{E} [x_t^\top (Q + K^\top RK)x_t - 2b^\top RKx_t] + b^\top Rb + \sigma^2 \cdot \text{tr}(R) + \mu^\top \bar{Q}\mu \right\} \\
&= \mathbb{E}_{x \sim \mathcal{N}(\mu_{K,b}, \Phi_K)} [x^\top (Q + K^\top RK)x - 2b^\top RKx] + b^\top Rb + \sigma^2 \cdot \text{tr}(R) + \mu^\top \bar{Q}\mu \\
\text{(B.5.11)} \quad &= \text{tr}[(Q + K^\top RK)\Phi_K] + \mu_{K,b}^\top (Q + K^\top RK)\mu_{K,b} - 2b^\top RK\mu_{K,b} \\
&\quad + b^\top Rb + \sigma^2 \cdot \text{tr}(R) + \mu^\top \bar{Q}\mu.
\end{aligned}$$

Now, by iteratively applying (3.2.3) and (3.2.4), we have

$$\text{(B.5.12)} \quad \text{tr}[(Q + K^\top RK)\Phi_K] = \text{tr}(P_K \Psi_\epsilon),$$

where P_K is given in (3.2.4). Combining (B.5.11) and (B.5.12), we know that

$$J(K, b) = J_1(K) + J_2(K, b) + \sigma^2 \cdot \text{tr}(R) + \mu^\top \bar{Q}\mu,$$

where

$$\begin{aligned}
J_1(K) &= \text{tr}[(Q + K^\top RK)\Phi_K] = \text{tr}(P_K \Psi_\epsilon), \\
J_2(K, b) &= \begin{pmatrix} \mu_{K,b} \\ b \end{pmatrix}^\top \begin{pmatrix} Q + K^\top RK & -K^\top R \\ -RK & R \end{pmatrix} \begin{pmatrix} \mu_{K,b} \\ b \end{pmatrix}.
\end{aligned}$$

We finish the proof of the proposition. \square

B.5.4. Proof of Proposition 3.2.3

Proof. By calculating the Hessian matrix of $J_2(K, b)$, we have

$$\begin{aligned}\nabla_{bb}^2 J_2(K, b) &= B^\top (I - A + BK)^{-\top} (Q + K^\top RK) (I - A + BK)^{-1} B \\ &\quad - [RK(I - A + BK)^{-1} B + B^\top (I - A + BK)^{-\top} K^\top R] + R \\ &= [R^{1/2} K (I - A + BK)^{-1} B - R^{1/2}]^\top [R^{1/2} K (I - A + BK)^{-1} B - R^{1/2}] \\ &\quad + B^\top (I - A + BK)^{-\top} Q (I - A + BK)^{-1} B,\end{aligned}$$

which is a positive definite matrix independent of b . We denote by its minimum singular value as ν_K . Also, note that $\|\nabla_{bb}^2 J_2(K, b)\|_2$ is upper bounded as

$$\|\nabla_{bb}^2 J_2(K, b)\|_2 \leq [1 - \rho(A - BK)]^{-2} \cdot (\|B\|_2^2 \cdot \|K\|_2^2 \cdot \|R\|_2 + \|B\|_2^2 \cdot \|Q\|_2).$$

Therefore, it holds that

$$\iota_K \leq [1 - \rho(A - BK)]^{-2} \cdot (\|B\|_2^2 \cdot \|K\|_2^2 \cdot \|R\|_2 + \|B\|_2^2 \cdot \|Q\|_2),$$

where ι_K is the maximum singular value of $\nabla_{bb}^2 J_2(K, b)$. We finish the proof of the proposition. \square

B.5.5. Proof of Proposition B.2.3

Proof. Following from Proposition B.2.2, it holds that

$$(B.5.13) \quad J_1(K) = \text{tr}(P_K \Psi_\epsilon) = \mathbb{E}_{y \sim \mathcal{N}(0, \Psi_\epsilon)}(y^\top P_K y) = \mathbb{E}_{y \sim \mathcal{N}(0, \Psi_\epsilon)}[f_K(y)],$$

where $f_K(y) = y^\top P_K y$. By the definition of P_K in (3.2.4), we obtain that

$$\begin{aligned} \nabla_K f_K(y) &= \nabla_K \left\{ y^\top (Q + K^\top R K) y + [(A - BK)y]^\top P_K [(A - BK)y]^\top \right\} \\ (B.5.14) \quad &= 2RKyy^\top + \nabla_K \left[f_K((A - BK)y) \right]. \end{aligned}$$

Also, we have

$$(B.5.15) \quad \nabla_K \left[f_K((A - BK)y) \right] = \nabla_K f_K((A - BK)y) - 2B^\top P_K (A - BK)yy^\top.$$

By plugging (B.5.15) into (B.5.14), we have

$$(B.5.16) \quad \nabla_K f_K(y) = 2[(R + B^\top P_K B)K - B^\top P_K A]yy^\top + \nabla_K f_K((A - BK)y).$$

By iteratively applying (B.5.16), it holds that

$$(B.5.17) \quad \nabla_K f_K(y) = 2[(R + B^\top P_K B)K - B^\top P_K A] \cdot \sum_{t=0}^{\infty} y_t y_t^\top,$$

where $y_{t+1} = (A - BK)y_t$ with $y_0 = y$. Now, combining (B.5.13) and (B.5.17), it holds that

$$\nabla_K J_1(K) = 2[(R + B^\top P_K B)K - B^\top P_K A] \Phi_K = 2(\Upsilon_K^{22} K - \Upsilon_K^{21}) \cdot \Phi_K,$$

where Υ_K is defined in (3.2.7). Meanwhile, combining the form of $\mu_{K,b}$ in (3.2.2), it holds by calculation that

$$\nabla_b J_2(K, b) = 2[\Upsilon_K^{22}(-K\mu_{K,b} + b) + \Upsilon_K^{21}\mu_{K,b} + q_{K,b}],$$

where $q_{K,b}$ is defined in (3.2.7). We finish the proof of the proposition. \square

B.5.6. Proof of Proposition B.2.1

Proof. From the definition of $V_{K,b}(x)$ in (B.2.1) and the definition of the cost function $c(x, u)$ in Problem 3.1.2, it holds that

$$V_{K,b}(x) = \sum_{t=0}^{\infty} \left\{ \mathbb{E} \left[x_t^\top (Q + K^\top RK) x_t - 2b^\top RK x_t + b^\top Rb + \sigma^2 \eta_t^\top R \eta_t + \mu^\top \bar{Q} \mu \mid x_0 = x \right] - J(K, b) \right\}.$$

Combining (3.2.1), we know that $V_{K,b}(x)$ is a quadratic function taking the form of $V_{K,b}(x) = x^\top Gx + r^\top x + h$, where G , r , and h are functions of K and b . Note that $V_{K,b}(x)$ satisfies that

$$(B.5.18) \quad V_{K,b}(x) = c(x, -Kx + b) - J(K, b) + \mathbb{E}[V_{K,b}(x') \mid x],$$

by substituting the form of $c(x, -Kx + b)$ in Problem 3.1.2 and $J(K, b)$ in (3.2.5) into (B.5.18), we obtain that

$$\begin{aligned}
& x^\top Gx + r^\top x + h \\
\text{(B.5.19)} \quad & = x^\top (Q + K^\top RK)x - 2b^\top RKx + b^\top Rb + \mu^\top \bar{Q}\mu \\
& \quad - [\text{tr}(P_K \Psi_\epsilon) + \mu_{K,b}^\top (Q + K^\top RK)\mu_{K,b} - 2b^\top RK\mu_{K,b} + \mu^\top \bar{Q}\mu + b^\top Rb] \\
& \quad + [(A - BK)x + (Bb + \bar{A}\mu + d)]^\top G[(A - BK)x + (Bb + \bar{A}\mu + d)] \\
& \quad + \text{tr}(G\Psi_\epsilon) + r^\top [(A - BK)x + (Bb + \bar{A}\mu + d)] + h - \sigma^2 \cdot \text{tr}(R).
\end{aligned}$$

By comparing the quadratic terms and linear terms on both the LHS and RHS in (B.5.19), we obtain that

$$G = P_K, \quad r = 2f_{K,b},$$

where $f_{K,b} = (I - A + BK)^{-\top} [(A - BK)^\top P_K (Bb + \bar{A}\mu + d) - K^\top Rb]$. Also, by the definition of $V_{K,b}(x)$ in (B.2.1), we know that $\mathbb{E}[V_{K,b}(x)] = 0$, where the expectation is taken following the stationary distribution generated by the policy $\pi_{K,b}$ and the state transition. Therefore, we have

$$h = -2f_{K,b}\mu_{K,b} - \mu_{K,b}^\top P_K \mu_{K,b} - \text{tr}(P_K \Phi_K),$$

which shows that

$$(B.5.20) \quad V_{K,b}(x) = x^\top P_K x - \text{tr}(P_K \Phi_K) + 2f_{K,b}^\top(x - \mu_{K,b}) - \mu_{K,b}^\top P_K \mu_{K,b}.$$

For the action-value function $Q_{K,b}(x, u)$, by plugging (B.5.20) into (B.2.2), we obtain that

$$\begin{aligned} Q_{K,b}(x, u) &= \begin{pmatrix} x \\ u \end{pmatrix}^\top \Upsilon_K \begin{pmatrix} x \\ u \end{pmatrix} + 2 \begin{pmatrix} p_{K,b} \\ q_{K,b} \end{pmatrix}^\top \begin{pmatrix} x \\ u \end{pmatrix} - \text{tr}(P_K \Phi_K) - \sigma^2 \cdot \text{tr}(R + P_K B B^\top) \\ &\quad - b^\top R b + 2b^\top R K \mu_{K,b} - \mu_{K,b}^\top (Q + K^\top R K + P_K) \mu_{K,b} \\ &\quad + 2f_{K,b}^\top [(\bar{A}\mu + d) - \mu_{K,b}] + (\bar{A}\mu + d)^\top P_K (\bar{A}\mu + d). \end{aligned}$$

We finish the proof of the proposition. □

B.5.7. Proof of Proposition B.2.5

Proof. By Proposition B.2.1, it holds that $Q_{K,b}$ takes the following linear form

$$(B.5.21) \quad Q_{K,b}(x, u) = \psi(x, u)^\top \alpha_{K,b} + \beta_{K,b},$$

where $\beta_{K,b}$ is a scalar independent of x and u . Note that $Q_{K,b}(x, u)$ satisfies that

$$(B.5.22) \quad Q_{K,b}(x, u) = c(x, u) - J(K, b) + \mathbb{E}_{\pi_{K,b}}[Q_{K,b}(x', u') \mid x, u],$$

where (x', u') is the state-action pair after (x, u) following the policy $\pi_{K,b}$ and the state transition. Combining (B.5.21) and (B.5.22), we obtain that

$$(B.5.23) \quad \psi(x, u)^\top \alpha_{K,b} = c(x, u) - J(K, b) + \mathbb{E}_{\pi_{K,b}}[\psi(x', u') | x, u]^\top \alpha_{K,b}.$$

By left multiplying $\psi(x, u)$ to both sides of (B.5.23), and taking the expectation, we have

$$\mathbb{E}_{\pi_{K,b}} \left\{ \psi(x, u) [\psi(x, u) - \psi(x', u')]^\top \right\} \cdot \alpha_{K,b} + \mathbb{E}_{\pi_{K,b}} [\psi(x, u)] \cdot J(K, b) = \mathbb{E}_{\pi_{K,b}} [c(x, u) \psi(x, u)].$$

Combining the definition of the matrix $\Theta_{K,b}$ in (B.2.7), we have

$$\begin{pmatrix} 1 & 0 \\ \mathbb{E}_{\pi_{K,b}} [\psi(x, u)] & \Theta_{K,b} \end{pmatrix} \begin{pmatrix} J(K, b) \\ \alpha_{K,b} \end{pmatrix} = \begin{pmatrix} J(K, b) \\ \mathbb{E}_{\pi_{K,b}} [c(x, u) \psi(x, u)] \end{pmatrix},$$

which concludes the proof of the proposition. \square

B.5.8. Proof of Proposition B.2.6

Proof. Invertibility and Upper Bound. We denote by $z_t = (x_t^\top, u_t^\top)^\top$ for any $t \geq 0$.

Then following from the state transition and the policy $\pi_{K,b}$, the transition of $\{z_t\}_{t \geq 0}$ takes the form of

$$(B.5.24) \quad z_{t+1} = Lz_t + \nu + \delta_t,$$

where L , ν and δ are defined as

$$L = \begin{pmatrix} A & B \\ -KA & -KB \end{pmatrix}, \quad \nu = \begin{pmatrix} \bar{A}\mu + d \\ -K(\bar{A}\mu + d) + b \end{pmatrix}, \quad \delta_t = \begin{pmatrix} \omega_t \\ -K\omega_t + \sigma\eta_t \end{pmatrix}.$$

Note that L also takes the form of

$$L = \begin{pmatrix} I \\ -K \end{pmatrix} \begin{pmatrix} A & B \end{pmatrix}.$$

Combining the fact that $\rho(UV) = \rho(VU)$ for any matrices U and V , we know that $\rho(L) = \rho(A - BK) < 1$, which verifies the stability of (B.5.24). Following from the stability of (B.5.24), we know that the Markov chain generated by (B.5.24) admits a unique stationary distribution $\mathcal{N}(\mu_z, \Sigma_z)$, where μ_z and Σ_z satisfy that

$$\mu_z = L\mu_z + \nu, \quad \Sigma_z = L\Sigma_z L^\top + \Psi_\delta.$$

where

$$\Psi_\delta = \begin{pmatrix} \Psi_\omega & -\Psi_\omega K^\top \\ -K\Psi_\omega & K\Psi_\omega K^\top + \sigma^2 I \end{pmatrix}.$$

Also, we know that Σ_z takes the form of

(B.5.25)

$$\Sigma_z = \text{Cov} \left[\begin{pmatrix} x \\ u \end{pmatrix} \right] = \begin{pmatrix} \Phi_K & -\Phi_K K^\top \\ -K\Phi_K & K\Phi_K K^\top + \sigma^2 I \end{pmatrix} = \begin{pmatrix} 0 & 0 \\ 0 & \sigma^2 I \end{pmatrix} + \begin{pmatrix} I \\ -K \end{pmatrix} \Phi_K \begin{pmatrix} I \\ -K \end{pmatrix}^\top,$$

where Φ_K is defined in (3.2.3).

The following lemma establishes the form of $\Theta_{K,b}$.

Lemma B.5.2. The matrix $\Theta_{K,b}$ in (B.2.7) takes the form of

$$\Theta_{K,b} = \begin{pmatrix} 2(\Sigma_z \otimes_s \Sigma_z)(I - L \otimes_s L)^\top & 0 \\ 0 & \Sigma_z(I - L)^\top \end{pmatrix}.$$

Proof. See §B.6.9 for a detailed proof. \square

Note that since $\rho(L) < 1$, both $I - L \otimes_s L$ and $I - L$ are positive definite. Therefore, by Lemma B.5.2, the matrix $\Theta_{K,b}$ is invertible. This finishes the proof of the invertibility of $\Theta_{K,b}$. Moreover, by (B.5.25) and Lemma B.5.2, we upper bound the spectral norm of $\Theta_{K,b}$ as

$$\|\Theta_{K,b}\|_2 \leq 2 \max\left\{ \|\Sigma_z\|_2^2 \cdot (1 + \|L\|_2^2), \|\Sigma_z\|_2 \cdot (1 + \|L\|_2) \right\} \leq 4(1 + \|K\|_F^2) \cdot \|\Phi_K\|_2^2,$$

which proves the upper bound of $\|\Theta_{K,b}\|_2$.

Minimum singular value. To lower bound $\sigma_{\min}(\tilde{\Theta}_{K,b})$, we only need to upper bound $\sigma_{\max}(\tilde{\Theta}_{K,b}^{-1}) = \|\tilde{\Theta}_{K,b}^{-1}\|_2$. We first calculate $\tilde{\Theta}_{K,b}^{-1}$. Recall that the matrix $\tilde{\Theta}_{K,b}$ in (B.2.8) takes the form of

$$\tilde{\Theta}_{K,b} = \begin{pmatrix} 1 & 0 \\ \mathbb{E}_{\pi_{K,b}}[\psi(x, u)] & \Theta_{K,b} \end{pmatrix}.$$

By the definition of the feature vector $\psi(x, u)$ in (B.2.5), the vector $\tilde{\sigma}_z = \mathbb{E}_{\pi_{K,b}}[\psi(x, u)]$ takes the form of

$$\tilde{\sigma}_z = \mathbb{E}_{\pi_{K,b}}[\psi(x, u)] = \begin{pmatrix} \text{svec}(\Sigma_z) \\ \mathbf{0}_{k+m} \end{pmatrix},$$

where $\mathbf{0}_{k+m}$ denotes the all-zero column vector with dimension $k + m$. Also, since $\Theta_{K,b}$ is invertible, the matrix $\tilde{\Theta}_{K,b}$ is also invertible, whose inverse takes the form of

$$\tilde{\Theta}_{K,b}^{-1} = \begin{pmatrix} 1 & 0 \\ -\Theta_{K,b}^{-1} \cdot \tilde{\sigma}_z & \Theta_{K,b}^{-1} \end{pmatrix}.$$

The following lemma upper bounds the spectral norm of $\tilde{\Theta}_{K,b}^{-1}$.

Lemma B.5.3. The spectral norm of the matrix $\tilde{\Theta}_{K,b}^{-1}$ is upper bounded by a positive constant $\tilde{\lambda}_K$, where $\tilde{\lambda}_K$ only depends on $\|K\|_2$ and $\rho(A - BK)$.

Proof. See §B.6.10 for a detailed proof. \square

By Lemma B.5.3, we know that $\sigma_{\min}(\tilde{\Theta}_{K,b})$ is lower bounded by a positive constant $\lambda_K = 1/\tilde{\lambda}_K$, which only depends on $\|K\|_2$ and $\rho(A - BK)$. This concludes the proof of the proposition. \square

B.6. Proofs of Lemmas

B.6.1. Proof of Lemma B.4.1

Proof. Following from (3.2.4), it holds that

$$(B.6.1) \quad y^\top P_{K_2} y = \sum_{t \geq 0} y^\top [(A - BK_2)^t]^\top (Q + K_2^\top R K_2) (A - BK_2)^t y.$$

Meanwhile, by the state transition $y_{t+1} = (A - BK_2)y_t$, we know that

$$(B.6.2) \quad y_t = (A - BK_2)^t y_0 = (A - BK_2)^t y.$$

By plugging (B.6.2) into (B.6.1), it holds that

$$(B.6.3) \quad y^\top P_{K_2} y = \sum_{t \geq 0} y_t^\top (Q + K_2^\top R K_2) y_t = \sum_{t \geq 0} (y_t^\top Q y_t + y_t^\top K_2^\top R K_2 y_t).$$

Also, it holds that

$$(B.6.4) \quad y^\top P_{K_1} y = \sum_{t \geq 0} (y_{t+1}^\top P_{K_1} y_{t+1} - y_t^\top P_{K_1} y_t)$$

Combining (B.6.3) and (B.6.4), we have

$$(B.6.5) \quad y^\top P_{K_2} y - y^\top P_{K_1} y = \sum_{t \geq 0} (y_t^\top Q y_t + y_t^\top K_2^\top R K_2 y_t + y_{t+1}^\top P_{K_1} y_{t+1} - y_t^\top P_{K_1} y_t).$$

Also, by the state transition $y_{t+1} = (A - BK_2)y_t$, it holds for any $t \geq 0$ that

$$\begin{aligned} & y_t^\top Q y_t + y_t^\top K_2^\top R K_2 y_t + y_{t+1}^\top P_{K_1} y_{t+1} - y_t^\top P_{K_1} y_t \\ &= y_t^\top [Q + (K_2 - K_1 + K_1)^\top R (K_2 - K_1 + K_1)] y_t \\ & \quad + y_t^\top [A - BK_1 - B(K_2 - K_1)]^\top P_{K_1} [A - BK_1 - B(K_2 - K_1)] y_t - y_t^\top P_{K_1} y_t \\ &= 2y_t^\top (K_2 - K_1)^\top [(R + B^\top P_{K_1} B)K_1 - B^\top P_{K_1} A] y_t \\ & \quad + y_t^\top (K_2 - K_1)^\top (R + B^\top P_{K_1} B)(K_2 - K_1) y_t \\ (B.6.6) \quad &= 2y_t^\top (K_2 - K_1)^\top (\Upsilon_{K_1}^{22} K_1 - \Upsilon_{K_1}^{21}) y_t + y_t^\top (K_2 - K_1)^\top \Upsilon_{K_1}^{22} (K_2 - K_1) y_t, \end{aligned}$$

where the matrix Υ_{K_1} is defined in (3.2.7). By plugging (B.6.6) into (B.6.5), we have

$$\begin{aligned} & y^\top P_{K_2} y - y^\top P_{K_1} y \\ &= \sum_{t \geq 0} 2y_t^\top (K_2 - K_1)^\top (\Upsilon_{K_1}^{22} K_1 - \Upsilon_{K_1}^{21}) y_t + y_t^\top (K_2 - K_1)^\top \Upsilon_{K_1}^{22} (K_2 - K_1) y_t \\ &= \sum_{t \geq 0} D_{K_1, K_2}(y_t), \end{aligned}$$

where $D_{K_1, K_2}(y) = 2y^\top (K_2 - K_1) (\Upsilon_{K_1}^{22} K_1 - \Upsilon_{K_1}^{21}) y + y^\top (K_2 - K_1)^\top \Upsilon_{K_1}^{22} (K_2 - K_1) y$. We finish the proof of the lemma. \square

B.6.2. Proof of Lemma B.4.2

Proof. We prove (B.4.16) and (B.4.17) separately in the sequel.

Proof of (B.4.16). From the definition of $J_1(K)$ in (3.2.6), we have

$$\begin{aligned} (B.6.7) \quad J_1(K) - J_1(K^*) &= \text{tr}(P_K \Psi_\epsilon - P_{K^*} \Psi_\epsilon) = \mathbb{E}_{y \sim \mathcal{N}(0, \Psi_\epsilon)} (y^\top P_K y - y^\top P_{K^*} y) \\ &= -\mathbb{E} \left[\sum_{t \geq 0} D_{K, K^*}(y_t) \right], \end{aligned}$$

where in the last equality, we apply Lemma B.4.1 and the expectation is taken following the transition $y_{t+1} = (A - BK^*)y_t$ with initial state $y_0 \sim \mathcal{N}(0, \Psi_\epsilon)$. Here we denote by $D_{K, K^*}(y)$ as

$$D_{K, K^*}(y) = 2y^\top (K^* - K) (\Upsilon_K^{22} K - \Upsilon_K^{21}) y + y^\top (K^* - K)^\top \Upsilon_K^{22} (K^* - K) y.$$

Also, we write $D_{K,K^*}(y)$ as

(B.6.8)

$$\begin{aligned} D_{K,K^*}(y) &= 2y^\top (K^* - K)(\Upsilon_K^{22}K - \Upsilon_K^{21})y + y^\top (K^* - K)^\top \Upsilon_K^{22}(K^* - K)y \\ &= y^\top [K^* - K + (\Upsilon_K^{22})^{-1}(\Upsilon_K^{22}K - \Upsilon_K^{21})]^\top \Upsilon_K^{22}[K^* - K + (\Upsilon_K^{22})^{-1}(\Upsilon_K^{22}K - \Upsilon_K^{21})]y \\ &\quad - y^\top (\Upsilon_K^{22}K - \Upsilon_K^{21})^\top (\Upsilon_K^{22})^{-1}(\Upsilon_K^{22}K - \Upsilon_K^{21})y. \end{aligned}$$

Note that the first term on the RHS of (B.6.8) is positive, due to the fact that it is a quadratic form of a positive definite matrix, we lower bound $D_{K,K^*}(y)$ as

$$(B.6.9) \quad D_{K,K^*}(y) \geq -y^\top (\Upsilon_K^{22}K - \Upsilon_K^{21})^\top (\Upsilon_K^{22})^{-1}(\Upsilon_K^{22}K - \Upsilon_K^{21})y.$$

Combining (B.6.7) and (B.6.9), it holds that

$$\begin{aligned} J_1(K) - J_1(K^*) &\leq \left\| \mathbb{E} \left(\sum_{t \geq 0} y_t y_t^\top \right) \right\|_2 \cdot \text{tr}[(\Upsilon_K^{22}K - \Upsilon_K^{21})^\top (\Upsilon_K^{22})^{-1}(\Upsilon_K^{22}K - \Upsilon_K^{21})] \\ &= \|\Phi_{K^*}\|_2 \cdot \text{tr}[(\Upsilon_K^{22}K - \Upsilon_K^{21})^\top (\Upsilon_K^{22})^{-1}(\Upsilon_K^{22}K - \Upsilon_K^{21})] \\ &\leq \|(\Upsilon_K^{22})^{-1}\|_2 \cdot \|\Phi_{K^*}\|_2 \cdot \text{tr}[(\Upsilon_K^{22}K - \Upsilon_K^{21})^\top (\Upsilon_K^{22}K - \Upsilon_K^{21})] \\ &\leq \sigma_{\min}^{-1}(R) \cdot \|\Phi_{K^*}\|_2 \cdot \text{tr}[(\Upsilon_K^{22}K - \Upsilon_K^{21})^\top (\Upsilon_K^{22}K - \Upsilon_K^{21})], \end{aligned}$$

where the last line comes from the fact that $\Upsilon_K^{22} = R + B^\top P_K B \succeq R$. This complete the proof of (B.4.16).

Proof of (B.4.17). Note that for any \tilde{K} , it holds by the optimality of K^* that

$$(B.6.10) \quad J_1(K) - J_1(K^*) \geq J_1(K) - J_1(\tilde{K}) = -\mathbb{E}\left[\sum_{t \geq 0} D_{K, \tilde{K}}(y_t)\right],$$

where the expectation is taken following the transition $y_{t+1} = (A - B\tilde{K})y_t$ with initial state $y_0 \sim \mathcal{N}(0, \Psi_\epsilon)$. By taking $\tilde{K} = K - (\Upsilon_K^{22})^{-1}(\Upsilon_K^{22}K - \Upsilon_K^{21})$ and following from a similar calculation as in (B.6.8), the function $D_{K, \tilde{K}}(y)$ takes the form of

$$(B.6.11) \quad D_{K, \tilde{K}}(y) = -y^\top (\Upsilon_K^{22}K - \Upsilon_K^{21})^\top (\Upsilon_K^{22})^{-1} (\Upsilon_K^{22}K - \Upsilon_K^{21}) y.$$

Combining (B.6.10) and (B.6.11), it holds that

$$\begin{aligned} J(K) - J(K^*) &\geq \text{tr}\left[\Phi_{\tilde{K}}(\Upsilon_K^{22}K - \Upsilon_K^{21})^\top (\Upsilon_K^{22})^{-1} (\Upsilon_K^{22}K - \Upsilon_K^{21})\right] \\ &\geq \sigma_{\min}(\Psi_\epsilon) \cdot \|\Upsilon_K^{22}\|_2^{-1} \cdot \text{tr}\left[(\Upsilon_K^{22}K - \Upsilon_K^{21})^\top (\Upsilon_K^{22}K - \Upsilon_K^{21})\right], \end{aligned}$$

where we use the fact that $\Phi_{\tilde{K}} = (A - B\tilde{K})\Phi_{\tilde{K}}(A - B\tilde{K})^\top + \Psi_\epsilon \succeq \Psi_\epsilon$ in the last line.

This finishes the proof of (B.4.17). \square

B.6.3. Proof of Lemma B.4.3

Proof. By Proposition B.2.2, we have

$$(B.6.12) \quad |J_1(\tilde{K}_{n+1}) - J_1(K_{n+1})| = \text{tr}\left[(P_{\tilde{K}_{n+1}} - P_{K_{n+1}})\Psi_\epsilon\right] \leq \|P_{\tilde{K}_{n+1}} - P_{K_{n+1}}\|_2 \cdot \|\Psi_\epsilon\|_F.$$

The following lemma upper bounds the term $\|P_{\tilde{K}_{n+1}} - P_{K_{n+1}}\|_2$.

Lemma B.6.1. Suppose that the parameters K and \tilde{K} satisfy that

$$(B.6.13) \quad \|\tilde{K} - K\|_2 \cdot (\|A - BK\|_2 + 1) \cdot \|\Phi_K\|_2 \leq \sigma_{\min}(\Psi_\omega)/4 \cdot \|B\|_2^{-1},$$

then it holds that

$$(B.6.14) \quad \|P_{\tilde{K}} - P_K\|_2 \leq 6 \cdot \sigma_{\min}^{-1}(\Psi_\omega) \cdot \|\Phi_K\|_2 \cdot \|K\|_2 \cdot \|R\|_2 \cdot \|\tilde{K} - K\|_2 \\ \cdot (\|B\|_2 \cdot \|K\|_2) \cdot \|A - BK\|_2 + \|B\|_2 \cdot \|K\|_2 + 1).$$

Proof. See Lemma 5.7 in Yang et al. (2019b) for a detailed proof. \square

To use Lemma B.6.1, it suffices to verify that \tilde{K}_{n+1} and K_{n+1} satisfy (B.6.13). Note that from the definitions of K_{n+1} and \tilde{K}_{n+1} in (B.4.18) and (B.4.21), respectively, we have

$$(B.6.15) \quad \|\tilde{K}_{n+1} - K_{n+1}\|_2 \cdot (\|A - B\tilde{K}_{n+1}\|_2 + 1) \cdot \|\Phi_{\tilde{K}_{n+1}}\|_2 \\ \leq \gamma \cdot \|\hat{\Upsilon}_{K_n} - \Upsilon_{K_n}\|_F \cdot (1 + \|K_n\|_2) \cdot (\|A - B\tilde{K}_{n+1}\|_2 + 1) \cdot \|\Phi_{\tilde{K}_{n+1}}\|_2.$$

Now, we upper bound the RHS of (B.6.15). For the term $\|A - B\tilde{K}_{n+1}\|_2$, it holds by the definition of \tilde{K}_{n+1} in (B.4.21) that

$$(B.6.16) \quad \|A - B\tilde{K}_{n+1}\|_2 \leq \|A - BK_n\|_2 + \gamma \cdot \|B\|_2 \cdot \|\Upsilon_{K_n}^{22} K_n - \Upsilon_{K_n}^{21}\|_2 \\ \leq \|A - BK_n\|_2 + \gamma \cdot \|B\|_2 \cdot \|\Upsilon_{K_n}\|_2 \cdot (1 + \|K_n\|_2).$$

By the definition of Υ_{K_n} in (3.2.7), we upper bound $\|\Upsilon_{K_n}\|_2$ as

$$\begin{aligned} \|\Upsilon_{K_n}\|_2 &\leq \|Q\|_2 + \|R\|_2 + (\|A\|_F + \|B\|_F)^2 \cdot \|P_{K_n}\|_2 \\ (B.6.17) \quad &\leq \|Q\|_2 + \|R\|_2 + (\|A\|_F + \|B\|_F)^2 \cdot J_1(K_0) \cdot \sigma_{\min}^{-1}(\Psi_\epsilon), \end{aligned}$$

where the last line comes from the fact that

$$J_1(K_0) \geq J_1(K_n) = \text{tr}[(Q + K_n^\top R K_n) \Phi_{K_n}] = \text{tr}(P_{K_n} \Psi_\epsilon) \geq \|P_{K_n}\|_2 \cdot \sigma_{\min}(\Psi_\epsilon).$$

As for the term $\|\Phi_{\tilde{K}_{n+1}}\|_2$ in (B.6.15), from the fact that

$$J_1(K_0) \geq J_1(\tilde{K}_{n+1}) = \text{tr}[(Q + \tilde{K}_{n+1}^\top R \tilde{K}_{n+1}) \Phi_{\tilde{K}_{n+1}}] \geq \|\Phi_{\tilde{K}_{n+1}}\|_2 \cdot \sigma_{\min}(Q),$$

it holds that

$$(B.6.18) \quad \|\Phi_{\tilde{K}_{n+1}}\|_2 \leq J_1(K_0) \cdot \sigma_{\min}^{-1}(Q).$$

Therefore, combining (B.6.15), (B.6.16), (B.6.17), and (B.6.18), we know that

$$\begin{aligned} &\|\tilde{K}_{n+1} - K_{n+1}\|_2 \cdot (\|A - B\tilde{K}_{n+1}\|_2 + 1) \cdot \|\Phi_{\tilde{K}_{n+1}}\|_2 \\ (B.6.19) \quad &\leq \text{poly}_1(\|K_n\|_2) \cdot \|\hat{\Upsilon}_{K_n} - \Upsilon_{K_n}\|_F. \end{aligned}$$

From Theorem B.2.8, it holds with probability at least $1 - T_n^{-4} - \tilde{T}_n^{-6}$ that

$$\begin{aligned} (B.6.20) \quad \|\hat{\Upsilon}_{K_n} - \Upsilon_{K_n}\|_F &\leq \frac{\text{poly}_3(\|K_n\|_F, \|\mu\|_2)}{\lambda_{K_n} \cdot (1 - \rho)^2} \cdot \frac{\log^3 T_n}{T_n^{1/4}} \\ &\quad + \frac{\text{poly}_4(\|K_n\|_F, \|b_0\|_2, \|\mu\|_2)}{\lambda_{K_n}} \cdot \frac{\log^{1/2} \tilde{T}_n}{\tilde{T}_n^{1/8} \cdot (1 - \rho)}, \end{aligned}$$

which holds for any $\rho \in (\rho(A - BK_n), 1)$. Note that from the choice of T_n and \tilde{T}_n in the statement of Theorem B.2.4 that

$$\begin{aligned} T_n &\geq \text{poly}_5(\|K_n\|_{\mathbb{F}}, \|b_0\|_2, \|\mu\|_2) \cdot \lambda_{K_n}^{-4} \cdot [1 - \rho(A - BK_n)]^{-9} \cdot \varepsilon^{-5}, \\ \tilde{T}_n &\geq \text{poly}_6(\|K_n\|_{\mathbb{F}}, \|b_0\|_2, \|\mu\|_2) \cdot \lambda_{K_n}^{-2} \cdot [1 - \rho(A - BK_n)]^{-12} \cdot \varepsilon^{-12}, \end{aligned}$$

it holds that

$$\begin{aligned} &\frac{\text{poly}_3(\|K_n\|_{\mathbb{F}}, \|\mu\|_2)}{\lambda_{K_n} \cdot (1 - \rho)^2} \cdot \frac{\log^3 T_n}{T_n^{1/4}} + \frac{\text{poly}_4(\|K_n\|_{\mathbb{F}}, \|b_0\|_2, \|\mu\|_2)}{\lambda_{K_n}} \cdot \frac{\log^{1/2} \tilde{T}_n}{\tilde{T}_n^{1/8} \cdot (1 - \rho)} \\ \text{(B.6.21)} \quad &\leq \min \left\{ \left[\text{poly}_1(\|K_n\|_2) \right]^{-1} \cdot \sigma_{\min}(\Psi_\omega)/4 \cdot \|B\|_2^{-1}, \right. \\ &\quad \left. \left[\text{poly}_2(\|K_n\|_2) \right]^{-1} \cdot \varepsilon/8 \cdot \gamma \cdot \sigma_{\min}(\Psi_\varepsilon) \cdot \sigma_{\min}(R) \cdot \|\Phi_{K^*}\|_2^{-1} \cdot \|\Psi_\varepsilon\|_{\mathbb{F}}^{-1} \right\}. \end{aligned}$$

Combining (B.6.19), (B.6.20), and (B.6.21), we know that (B.6.13) holds with probability at least $1 - \varepsilon^{15}$ for sufficiently small $\varepsilon > 0$. Meanwhile, by (B.6.16), (B.6.17), and (B.6.18), the RHS of (B.6.14) is upper bounded as

$$\begin{aligned} &6 \cdot \sigma_{\min}^{-1}(\Psi_\omega) \cdot \|\Phi_{\tilde{K}_{n+1}}\|_2 \cdot \|\tilde{K}_{n+1}\|_2 \cdot \|R\|_2 \cdot \|\tilde{K}_{n+1} - K_{n+1}\|_2 \\ &\quad \cdot (\|B\|_2 \cdot \|\tilde{K}_{n+1}\|_2) \cdot \|A - B\tilde{K}_{n+1}\|_2 + \|B\|_2 \cdot \|\tilde{K}_{n+1}\|_2 + 1) \\ \text{(B.6.22)} \quad &\leq \text{poly}_2(\|K_n\|_2) \cdot \|\hat{\Upsilon}_{K_n} - \Upsilon_{K_n}\|_{\mathbb{F}}. \end{aligned}$$

Now, by Lemma B.6.1, it holds with probability at least $1 - \varepsilon^{15}$ that

$$\begin{aligned}
\|P_{\tilde{K}_{n+1}} - P_{K_{n+1}}\|_2 &\leq 6 \cdot \sigma_{\min}^{-1}(\Psi_\omega) \cdot \|\Phi_{\tilde{K}_{n+1}}\|_2 \cdot \|\tilde{K}_{n+1}\|_2 \cdot \|R\|_2 \cdot \|\tilde{K}_{n+1} - K_{n+1}\|_2 \\
&\quad \cdot (\|B\|_2 \cdot \|\tilde{K}_{n+1}\|_2) \cdot \|A - B\tilde{K}_{n+1}\|_2 + \|B\|_2 \cdot \|\tilde{K}_{n+1}\|_2 + 1) \\
&\leq \text{poly}_2(\|K_n\|_2) \cdot \|\hat{\Upsilon}_{K_n} - \Upsilon_{K_n}\|_F \\
\text{(B.6.23)} \quad &\leq \varepsilon/8 \cdot \gamma \cdot \sigma_{\min}(\Psi_\epsilon) \cdot \sigma_{\min}(R) \cdot \|\Phi_{K^*}\|_2^{-1} \cdot \|\Psi_\epsilon\|_F^{-1},
\end{aligned}$$

where the second inequality comes from (B.6.22), and the last inequality comes from (B.6.20) and (B.6.21). Combining (B.6.12) and (B.6.23), it holds with probability at least $1 - \varepsilon^{15}$ that

$$|J_1(\tilde{K}_{n+1}) - J_1(K_{n+1})| \leq \gamma \cdot \sigma_{\min}(\Psi_\epsilon) \cdot \sigma_{\min}(R) \cdot \|\Phi_{K^*}\|_2^{-1} \cdot \varepsilon/4,$$

which concludes the proof of the lemma. \square

B.6.4. Proof of Lemma B.4.4

Proof. Note that $\Upsilon_{K^*}^{22} K^* - \Upsilon_{K^*}^{21}$ is the natural gradient of J_1 at the minimizer K^* , which implies that

$$\text{(B.6.24)} \quad \Upsilon_{K^*}^{22} K^* - \Upsilon_{K^*}^{21} = 0.$$

By Lemma B.4.1, it holds that

$$\begin{aligned}
J_1(K) - J_1(K^*) &= \text{tr}(P_K \Psi_\epsilon - P_{K^*} \Psi_\epsilon) = \mathbb{E}_{y \sim \mathcal{N}(0, \Psi_\epsilon)}(y^\top P_K y - y^\top P_{K^*} y) \\
&= \mathbb{E} \left\{ \sum_{t \geq 0} \left[2y_t^\top (K - K^*) (\Upsilon_{K^*}^{22} K^* - \Upsilon_{K^*}^{21}) y_t + y_t^\top (K - K^*)^\top \Upsilon_{K^*}^{22} (K - K^*) y_t \right] \right\} \\
\text{(B.6.25)} \quad &= \mathbb{E} \left\{ \sum_{t \geq 0} y_t^\top (K - K^*)^\top \Upsilon_{K^*}^{21} (K - K^*) y_t \right\},
\end{aligned}$$

where we use (B.6.24) in the last line. Here the expectations are taken following the transition $y_{t+1} = (A - BK)y_t$ with initial state $y_0 \sim \mathcal{N}(0, \Psi_\epsilon)$. Also, we have

$$\begin{aligned}
&\mathbb{E} \left\{ \sum_{t \geq 0} y_t^\top (K - K^*)^\top \Upsilon_{K^*}^{22} (K - K^*) y_t \right\} \\
&= \text{tr}[\Phi_K (K - K^*)^\top \Upsilon_{K^*}^{22} (K - K^*)] \\
&\geq \|\Phi_K\|_2 \cdot \|\Upsilon_{K^*}^{22}\|_2 \cdot \text{tr}[(K - K^*)^\top (K - K^*)] \\
\text{(B.6.26)} \quad &\geq \sigma_{\min}(\Psi_\epsilon) \cdot \sigma_{\min}(R) \cdot \|K - K^*\|_{\mathbb{F}}^2,
\end{aligned}$$

where we use the fact that $\Phi_K = (A - BK)\Phi_K(A - BK) + \Psi_\epsilon \succeq \Psi_\epsilon$ and $\Upsilon_{K^*}^{22} = R + B^\top P_{K^*} B \succeq R$ in the last line. Combining (B.6.25) and (B.6.26), we have

$$J_1(K) - J_1(K^*) \geq \sigma_{\min}(\Psi_\epsilon) \cdot \sigma_{\min}(R) \cdot \|K - K^*\|_{\mathbb{F}}^2.$$

We conclude the proof of the lemma. □

B.6.5. Proof of Lemma B.4.5

Proof. Following from Proposition 3.2.3, we have

$$\begin{aligned}
& J_2(K_N, b_{h+1}) - J_2(K_N, \tilde{b}_{h+1}) \\
& \leq \gamma^b \cdot \nabla_b J_2(K_N, \tilde{b}_{h+1})^\top [\nabla_b J_2(K_N, b_h) - \widehat{\nabla}_b J_2(K_N, b_h)] \\
& \quad + (\gamma^b)^2 \cdot \nu_{K_N}/2 \cdot \|\nabla_b J_2(K_N, b_h) - \widehat{\nabla}_b J_2(K_N, b_h)\|_2^2, \\
& J_2(K_N, \tilde{b}_{h+1}) - J_2(K_N, b_{h+1}) \\
\text{(B.6.27)} \quad & \leq -\gamma^b \cdot \nabla_b J_2(K_N, \tilde{b}_{h+1})^\top [\nabla_b J_2(K_N, b_h) - \widehat{\nabla}_b J_2(K_N, b_h)] \\
& \quad - (\gamma^b)^2 \cdot \iota_{K_N}/2 \cdot \|\nabla_b J_2(K_N, b_h) - \widehat{\nabla}_b J_2(K_N, b_h)\|_2^2,
\end{aligned}$$

where ν_{K_N} and ι_{K_N} are defined in Proposition 3.2.3. Also, following from Proposition B.2.3, it holds that

$$\text{(B.6.28)} \quad \|\nabla_b J_2(K_N, \tilde{b}_{h+1})\|_2 \leq \text{poly}_1(\|K_N\|_F, \|b_h\|_2, \|\mu\|_2, J(K_N, b_0)) \cdot [1 - \rho(A - BK_N)]^{-1}.$$

Combining (B.6.27), (B.6.28), and the fact that $\nu_{K_N} \leq \iota_{K_N} \leq [1 - \rho(A - BK_N)]^{-2} \cdot \text{poly}_2(\|K_N\|_2)$, we know that

(B.6.29)

$$\begin{aligned}
& |J_2(K_N, b_{h+1}) - J_2(K_N, \tilde{b}_{h+1})| \\
& \leq (\gamma^b)^2 \cdot \text{poly}_2(\|K_N\|_2) \cdot \|\nabla_b J_2(K_N, b_h) - \widehat{\nabla}_b J_2(K_N, b_h)\|_2^2 \cdot [1 - \rho(A - BK_N)]^{-2} \\
& \quad + \gamma^b \cdot \text{poly}_1(\|K_N\|_F, \|b_h\|_2, \|\mu\|_2, J(K_N, b_0)) \cdot \|\nabla_b J_2(K_N, b_h) - \widehat{\nabla}_b J_2(K_N, b_h)\|_2 \\
& \quad \cdot [1 - \rho(A - BK_N)]^{-1}.
\end{aligned}$$

Note that from the definition of $\widehat{\nabla}_b J_2(K_N, b_h)$ and $\nabla_b J_2(K_N, b_h)$ in (B.4.31) and (B.4.33), respectively, it holds by triangle inequality that

$$\begin{aligned}
& \|\nabla_b J_2(K_N, b_h) - \widehat{\nabla}_b J_2(K_N, b_h)\|_2 \\
& \leq \|\widehat{\Upsilon}_{K_N}^{22} - \Upsilon_{K_N}^{22}\|_2 \cdot \|K_N\|_2 \cdot \|\widehat{\mu}_{K_N, b_h}\|_2 + \|\Upsilon_{K_N}^{22}\|_2 \cdot \|K_N\|_2 \cdot \|\widehat{\mu}_{K_N, b_h} - \mu_{K_N, b_h}\|_2 \\
& \quad + \|\widehat{\Upsilon}_{K_N}^{21} - \Upsilon_{K_N}^{21}\|_2 \cdot \|b_h\|_2 + \|\widehat{\Upsilon}_{K_N}^{21} - \Upsilon_{K_N}^{21}\|_2 \cdot \|\widehat{\mu}_{K_N, b_h}\|_2 + \|\widehat{q}_{K_N, b_h} - q_{K_N, b_h}\|_2 \\
& \quad + \|\Upsilon_{K_N}^{21}\|_2 \cdot \|\widehat{\mu}_{K_N, b_h} - \mu_{K_N, b_h}\|_2.
\end{aligned}$$

By Theorem B.2.8, combining the fact that $J_2(K_N, b_h) \leq J_2(K_N, b_0)$ and the fact that $\|\mu_{K_N, b}\|_2 \leq J(K_N, b_0)/\sigma_{\min}(Q)$, we know that with probability at least $1 - (T_n^b)^{-4} - (\widetilde{T}_n^b)^{-6}$,

it holds for any $\rho \in (\rho(A - BK_N), 1)$ that

(B.6.30)

$$\begin{aligned} & \|\nabla_b J_2(K_N, b_h) - \widehat{\nabla}_b J_2(K_N, b_h)\|_2 \\ & \leq \lambda_{K_N}^{-1} \cdot \text{poly}_3(\|K_N\|_F, \|b_h\|_2, \|\mu\|_2, J_2(K_N, b_0)) \cdot \left[\frac{\log^3 T_n^b}{(T_n^b)^{1/4}(1-\rho)^2} + \frac{\log^{1/2} \widetilde{T}_n^b}{(\widetilde{T}_n^b)^{1/8} \cdot (1-\rho)} \right]. \end{aligned}$$

Following from the choices of γ^b , T_n^b , and \widetilde{T}_n^b in the statement of Theorem B.2.4, it holds that

$$\begin{aligned} & \gamma^b \cdot \text{poly}_1(\|K_N\|_F, \|b_h\|_2, \|\mu\|_2, J(K_N, b_0)) \cdot \lambda_{K_N}^{-1} \cdot \text{poly}_3(\|K_N\|_F, \|b_h\|_2, \|\mu\|_2, J_2(K_N, b_0)) \\ & \cdot \left[\frac{\log^3 T_n^b}{(T_n^b)^{1/4}(1-\rho)^2} + \frac{\log^{1/2} \widetilde{T}_n^b}{(\widetilde{T}_n^b)^{1/8} \cdot (1-\rho)} \right] \cdot [1 - \rho(A - BK_N)]^{-1} + [1 - \rho(A - BK_N)]^{-2} \\ & \cdot \text{poly}_3(\|K_N\|_F, \|b_h\|_2, \|\mu\|_2, J_2(K_N, b_0)) \cdot \left[\frac{\log^6 T_n^b}{(T_n^b)^{1/2}(1-\rho)^4} + \frac{\log \widetilde{T}_n^b}{(\widetilde{T}_n^b)^{1/4} \cdot (1-\rho)^2} \right] \\ & \cdot (\gamma^b)^2 \cdot \text{poly}_2(\|K_N\|_2) \cdot \lambda_{K_N}^{-1} \\ & \leq \nu_{K_N} \cdot \gamma^b \cdot \varepsilon/2. \end{aligned}$$

Further combining (B.6.29) and (B.6.30), it holds with probability at least $1 - \varepsilon^{15}$ that

$$|J_2(K_N, b_{h+1}) - J_2(K_N, \widetilde{b}_{h+1})| \leq \nu_{K_N} \cdot \gamma^b \cdot \varepsilon/2.$$

We then finish the proof of the lemma. \square

B.6.6. Proof of Lemma B.4.6

Proof. We show that $\zeta_{K,b} \in \mathcal{V}_\zeta$ and $\xi(\zeta) \in \mathcal{V}_\xi$ for any $\zeta \in \mathcal{V}_\zeta$ separately.

Part 1. First we show that $\zeta_{K,b} \in \mathcal{V}_\zeta$. Note that from Definition B.2.7, we know that $\zeta_{K,b}^1 = J(K, b)$ satisfies that $0 \leq \zeta_{K,b}^1 \leq J(K_0, b_0)$. It remains to show that $\zeta_{K,b}^2 = \alpha_{K,b}$ satisfies that $\|\zeta_{K,b}^2\|_2 \leq M_\zeta$. By the definition of $\alpha_{K,b}$ in (B.2.6), we know that

$$(B.6.31) \quad \begin{aligned} \|\alpha_{K,b}\|_2^2 &\leq \|\Upsilon_K\|_F^2 + \|\Upsilon_K\|_2^2 \cdot (\|\mu_{K,b}\|_2^2 + \|\mu_{K,b}^u\|_2^2) \\ &\quad + (\|A\|_2 + \|B\|_2)^2 \cdot (\|P_K\|_2 \cdot \|\bar{A}\mu + d\|_2 + \|f_{K,b}\|_2)^2 \end{aligned}$$

where $f_{K,b} = (I - A + BK)^{-\top}[(A - BK)^\top P_K(Bb + \bar{A}\mu + d) - K^\top Rb]$ and for notational simplicity, we denote by $\mu_{K,b}^u = -K\mu_{K,b} + b$. We only need to bound Υ_K , $\mu_{K,b}$, $\mu_{K,b}^u$, P_K , and $f_{K,b}$. Note that by Proposition B.2.2, the expected total cost $J(K, b)$ takes the form of

$$J(K, b) = \text{tr}(P_K \Psi_\epsilon) + \mu_{K,b}^\top Q \mu_{K,b} + (\mu_{K,b}^u)^\top R \mu_{K,b}^u + \sigma^2 \cdot \text{tr}(R) + \mu^\top \bar{Q} \mu.$$

Thus, we have

$$J(K_0, b_0) \geq J(K, b) \geq \sigma_{\min}(\Psi_\omega) \cdot \text{tr}(P_K) \geq \sigma_{\min}(\Psi_\omega) \cdot \|P_K\|_2,$$

$$J(K_0, b_0) \geq J(K, b) \geq \mu_{K,b}^\top Q \mu_{K,b} \geq \sigma_{\min}(Q) \cdot \|\mu_{K,b}\|_2,$$

$$J(K_0, b_0) \geq J(K, b) \geq (\mu_{K,b}^u)^\top R \mu_{K,b}^u \geq \sigma_{\min}(R) \cdot \|\mu_{K,b}^u\|_2,$$

which imply that

$$\begin{aligned}
\|P_K\|_2 &\leq J(K_0, b_0)/\sigma_{\min}(\Psi_\omega), \\
\|\mu_{K,b}\|_2 &\leq J(K_0, b_0)/\sigma_{\min}(Q), \\
\|\mu_{K,b}^u\|_2 &\leq J(K_0, b_0)/\sigma_{\min}(R).
\end{aligned}
\tag{B.6.32}$$

For Υ_K , it holds that

$$\Upsilon_K = \begin{pmatrix} Q & \mathbf{0} \\ \mathbf{0} & R \end{pmatrix} + \begin{pmatrix} A^\top \\ B^\top \end{pmatrix} P_K \begin{pmatrix} A & B \end{pmatrix},$$

which gives

$$\begin{aligned}
\|\Upsilon_K\|_F &\leq (\|Q\|_F + \|R\|_F) + (\|A\|_F^2 + \|B\|_F^2) \cdot \|P_K\|_F, \\
\|\Upsilon_K\|_2 &\leq (\|Q\|_2 + \|R\|_2) + (\|A\|_2 + \|B\|_2)^2 \cdot \|P_K\|_2.
\end{aligned}$$

Combining (B.6.32) and the fact that $\|P_K\|_F \leq \sqrt{m} \cdot \|P_K\|_2$, we know that

$$\begin{aligned}
\|\Upsilon_K\|_F &\leq (\|Q\|_F + \|R\|_F) + (\|A\|_F^2 + \|B\|_F^2) \cdot \sqrt{m} \cdot J(K_0, b_0)/\sigma_{\min}(\Psi_\omega), \\
\|\Upsilon_K\|_2 &\leq (\|Q\|_2 + \|R\|_2) + (\|A\|_2 + \|B\|_2)^2 \cdot J(K_0, b_0)/\sigma_{\min}(\Psi_\omega).
\end{aligned}
\tag{B.6.33}$$

Now, we upper bound the vector $f_{K,b}$. Note that by algebra, the vector $f_{K,b}$ takes the form of

$$f_{K,b} = -P_K \mu_{K,b} + (I - A + BK)^{-T} (Q \mu_{K,b} - K^\top R \mu_{K,b}^u).$$

Therefore, we upper bound $f_{K,b}$ as

(B.6.34)

$$\|f_{K,b}\|_2 \leq J(K_0, b_0)^2 \cdot \sigma_{\min}^{-1}(\Psi_\omega) \cdot \sigma_{\min}^{-1}(Q) + [1 - \rho(A - BK)]^{-1} \cdot (\kappa_Q + \kappa_R \cdot \|K\|_F)$$

Combining (B.6.31), (B.6.32), (B.6.33), and (B.6.34), it holds that

$$\|\zeta_{K,b}^2\|_2 = \|\alpha_{K,b}\|_2 \leq M_{\zeta,1} + M_{\zeta,2} \cdot (1 + \|K\|_F) \cdot [1 - \rho(A - BK)]^{-1}.$$

Therefore, it holds that $\zeta_{K,b} \in \mathcal{V}_\zeta$.

Part 2. Now we show that for any $\zeta \in \mathcal{V}_\zeta$, we have $\xi(\zeta) \in \mathcal{V}_\xi$. Recall that from (B.4.41), it holds that

(B.6.35)

$$\xi^1(\zeta) = \zeta^1 - J(K, b), \quad \xi^2(\zeta) = \mathbb{E}_{\pi_{K,b}}[\psi(x, u)]\zeta^1 + \Theta_{K,b}\zeta^2 - \mathbb{E}_{\pi_{K,b}}[c(x, u)\psi(x, u)].$$

Then we have

$$(B.6.36) \quad |\xi^1(\zeta)| = |\zeta^1 - J(K, b)| \leq J(K_0, b_0),$$

where we use the fact that since $\zeta \in \mathcal{V}_\zeta$, we have $0 \leq \zeta^1 \leq J(K_0, b_0)$ by Definition B.2.7.

Also, by (B.6.35), we have

$$(B.6.37) \quad \|\xi^2(\zeta)\|_2 \leq \underbrace{\left\| \mathbb{E}_{\pi_{K,b}}[\psi(x, u)]\zeta^1 \right\|_2}_{B_1} + \underbrace{\|\Theta_{K,b}\|_2 \cdot \|\zeta^2\|_2}_{B_2} + \underbrace{\left\| \mathbb{E}_{\pi_{K,b}}[c(x, u)\psi(x, u)] \right\|_2}_{B_3}.$$

Note that we upper bound B_1 as

$$(B.6.38) \quad B_1 \leq J(K_0, b_0) \cdot \left\| \mathbb{E}_{\pi_{K,b}} [\psi(x, u)] \right\|_2.$$

Following from the definition of $\psi(x, u)$ in (B.2.5), we know that

$$(B.6.39) \quad \left\| \mathbb{E}_{\pi_{K,b}} [\psi(x, u)] \right\|_2 \leq \|\Sigma_z\|_F,$$

where Σ_z is defined as

$$\Sigma_z = \text{Cov} \left[\begin{pmatrix} x \\ u \end{pmatrix} \right] = \begin{pmatrix} \Phi_K & -\Phi_K K^\top \\ -K \Phi_K & K \Phi_K K^\top + \sigma^2 I \end{pmatrix} = \begin{pmatrix} 0 & 0 \\ 0 & \sigma^2 I \end{pmatrix} + \begin{pmatrix} I \\ -K \end{pmatrix} \Phi_K \begin{pmatrix} I \\ -K \end{pmatrix}^\top.$$

Combining (B.6.38) and (B.6.39), we have

$$(B.6.40) \quad B_1 \leq J(K_0, b_0) \cdot \|\Sigma_z\|_F.$$

By Proposition B.2.6, we upper bound B_2 as

$$(B.6.41) \quad B_2 \leq 4(1 + \|K\|_F^2)^3 \cdot \|\Phi_K\|_2^2 \cdot (M_{\zeta,1} + M_{\zeta,2}) \cdot [1 - \rho(A - BK)]^{-1},$$

where we use the fact that $\zeta \in \mathcal{V}_\zeta$ and Definition B.2.7. As for the term B_3 in (B.6.37), we utilize the following lemma to provide an upper bound.

Lemma B.6.2. The vector $\mathbb{E}_{\pi_{K,b}}[c(x, u)\psi(x, u)]$ takes the following form,

$$\mathbb{E}_{\pi_{K,b}}[c(x, u)\psi(x, u)] = \begin{pmatrix} 2\text{svec}[\Sigma_z \text{diag}(Q, R)\Sigma_z + \langle \Sigma_z, \text{diag}(Q, R) \rangle \Sigma_z] \\ \Sigma_z \begin{pmatrix} 2Q\mu_{K,b} \\ 2R\mu_{K,b}^u \end{pmatrix} \end{pmatrix} + [\mu_{K,b}^\top Q \mu_{K,b} + (\mu_{K,b}^u)^\top R \mu_{K,b}^u + \mu^\top \bar{Q} \mu] \begin{pmatrix} \text{svec}(\Sigma_z) \\ \mathbf{0}_m \\ \mathbf{0}_k \end{pmatrix}.$$

Here the matrix Σ_z takes the form of

$$\Sigma_z = \begin{pmatrix} \Phi_K & -\Phi_K K^\top \\ -K \Phi_K & K \Phi_K K^\top + \sigma^2 \cdot I \end{pmatrix}.$$

Proof. See §B.6.11 for a detailed proof. □

From Lemma B.6.2 and (B.6.32), it holds that

$$(B.6.42) \quad B_3 \leq 3[\|Q\|_F + \|R\|_F + J(K_0, b_0) \cdot \|Q\|_2 / \sigma_{\min}(Q) + J(K_0, b_0) \cdot \|R\|_2 / \sigma_{\min}(R)] \cdot \|\Sigma_z\|_2^2.$$

Moreover, by the definition of Σ_z in (B.5.25), combining the triangle inequality, we have the following bounds for $\|\Sigma_z\|_F$ and $\|\Sigma_z\|_2$,

$$(B.6.43) \quad \|\Sigma_z\|_F \leq 2(d + \|K\|_F^2) \cdot \|\Phi_K\|_2, \quad \|\Sigma_z\|_2 \leq 2(1 + \|K\|_F^2) \cdot \|\Phi_K\|_2.$$

Also, we have

$$J(K_0, b_0) \geq J(K, b) \geq \text{tr}[(Q + K^\top RK)\Phi_K] \geq \|\Phi_K\|_2 \cdot \sigma_{\min}(Q),$$

which gives the upper bound for Φ_K as follows,

$$(B.6.44) \quad \|\Phi_K\|_2 \leq J(K_0, b_0)/\sigma_{\min}(Q).$$

Therefore, combining (B.6.37), (B.6.40), (B.6.41), (B.6.42), (B.6.43), and (B.6.44), we know that

$$(B.6.45) \quad \begin{aligned} \|\xi^2(\zeta)\|_2 &\leq C \cdot (M_{\zeta,1} + M_{\zeta,2}) \cdot J(K_0, b_0)^2 / \sigma_{\min}^2(Q) \\ &\quad \cdot (1 + \|K\|_{\mathbb{F}}^2)^3 \cdot [1 - \rho(A - BK)]^{-1}. \end{aligned}$$

By (B.6.36) and (B.6.45), we know that $\xi(\zeta) \in \mathcal{V}_\xi$ for any $\zeta \in \mathcal{V}_\zeta$. We conclude the proof of the lemma. \square

B.6.7. Proof of Lemma B.4.7

Proof. Assume that $\tilde{z}_0 \sim \mathcal{N}(\mu_\dagger, \Sigma_\dagger)$. Following from the fact that

$$\tilde{z}_{t+1} = L\tilde{z}_t + \nu + \delta_t,$$

it holds that

$$(B.6.46) \quad \tilde{z}_t \sim \mathcal{N}\left(L^t \mu_\dagger + \sum_{i=0}^{t-1} L^i \cdot \nu, (L^\top)^t \Sigma_\dagger L^t + \sum_{i=0}^{t-1} (L^\top)^i \Psi_\delta L^i\right),$$

where

$$\Psi_\delta = \begin{pmatrix} \Psi_\omega & K\Psi_\omega \\ K\Psi_\omega & K\Psi_\omega K^\top + \sigma^2 I \end{pmatrix}.$$

From (B.4.47), we know that μ_z takes the form of

$$(B.6.47) \quad \mu_z = (I - L)^{-1}\nu = \sum_{j=0}^{\infty} L^j \nu.$$

Therefore, combining (B.6.46) and (B.6.47), we have

$$(B.6.48) \quad \mathbb{E}(\widehat{\mu}_z) = \mu_z + \frac{1}{\widetilde{T}} \sum_{t=1}^{\widetilde{T}} L^t \mu_\dagger - \frac{1}{\widetilde{T}} \sum_{t=1}^{\widetilde{T}} \sum_{i=t}^{\infty} L^i \nu.$$

We denote by

$$\mu_{\widetilde{T}} = \sum_{t=1}^{\widetilde{T}} L^t \mu_\dagger - \sum_{t=1}^{\widetilde{T}} \sum_{i=t}^{\infty} L^i \nu.$$

Meanwhile, it holds that

$$\begin{aligned} & \left\| \sum_{t=1}^{\widetilde{T}} L^t \mu_\dagger - \sum_{t=1}^{\widetilde{T}} \sum_{i=t}^{\infty} L^i \nu \right\|_2 \\ & \leq \sum_{t=1}^{\widetilde{T}} \rho(L)^t \cdot \|\mu_\dagger\|_2 + \sum_{t=1}^{\widetilde{T}} \sum_{i=t}^{\infty} \rho(L)^i \cdot \|\nu\|_2 \\ & \leq [1 - \rho(L)]^{-1} \cdot \|\mu_\dagger\|_2 + [1 - \rho(L)]^{-2} \cdot \|\nu\|_2 \\ (B.6.49) \quad & \leq M_\mu \cdot (1 - \rho)^{-2} \cdot \|\mu_z\|_2, \end{aligned}$$

where M_μ is a positive absolute constant.

For the covariance, note that for any random variables $X \sim \mathcal{N}(\mu_1, \Sigma_1)$ and $Y \sim \mathcal{N}(\mu_2, \Sigma_2)$, we know that $Z = X + Y \sim \mathcal{N}(\mu_1 + \mu_2, \Sigma)$, where $\|\Sigma\|_{\mathbb{F}} \leq 2\|\Sigma_1\|_{\mathbb{F}} + 2\|\Sigma_2\|_{\mathbb{F}}$. Combining (B.6.46), we know that $\hat{\mu}_z \sim \mathcal{N}(\mathbb{E}\hat{\mu}_z, \tilde{\Sigma}_{\tilde{T}}/\tilde{T})$, where $\tilde{\Sigma}_{\tilde{T}}$ satisfies that

$$\begin{aligned} \tilde{T}/2 \cdot \|\tilde{\Sigma}_{\tilde{T}}\|_{\mathbb{F}} &\leq \sum_{t=1}^{\tilde{T}} \rho(L)^{2t} \cdot \|\Sigma_{\dagger}\|_{\mathbb{F}} + \sum_{t=1}^{\tilde{T}} \sum_{i=0}^{t-1} \rho(L)^{2i} \cdot \|\Psi_{\delta}\|_{\mathbb{F}} \\ &\leq [1 - \rho(L)^2]^{-1} \cdot \|\Sigma_{\dagger}\|_{\mathbb{F}} + \tilde{T} \cdot [1 - \rho(L)^2]^{-1} \cdot \|\Psi_{\delta}\|_{\mathbb{F}}, \end{aligned}$$

which implies that

$$(B.6.50) \quad \|\tilde{\Sigma}_{\tilde{T}}\|_{\mathbb{F}} \leq M_{\Sigma} \cdot (1 - \rho)^{-1} \cdot \|\Sigma_z\|_{\mathbb{F}},$$

where M_{Σ} is a positive absolute constant. Combining (B.6.48), (B.6.49), and (B.6.50), we obtain that

$$\hat{\mu}_z \sim \mathcal{N}\left(\mu_z + \frac{1}{\tilde{T}}\mu_{\tilde{T}}, \frac{1}{\tilde{T}}\tilde{\Sigma}_{\tilde{T}}\right),$$

where $\|\mu_{\tilde{T}}\|_2 \leq M_{\mu} \cdot (1 - \rho)^{-2} \cdot \|\mu_z\|_2$ and $\|\tilde{\Sigma}_{\tilde{T}}\|_{\mathbb{F}} \leq M_{\Sigma} \cdot (1 - \rho)^{-1} \cdot \|\Sigma_z\|_{\mathbb{F}}$. Moreover, by the Gaussian tail inequality, it holds with probability at least $1 - \tilde{T}^{-6}$ that

$$\|\hat{\mu}_z - \mu_z\|_2 \leq \frac{\log \tilde{T}}{\tilde{T}^{1/4}} \cdot (1 - \rho)^{-2} \cdot \text{poly}(\|\Phi_K\|_2, \|K\|_{\mathbb{F}}, \|b\|_2, \|\mu\|_2).$$

Then we finish the proof of the lemma. \square

B.6.8. Proof of Lemma B.4.8

Proof. We continue using the notations given in §B.4.3. We define

$$\widehat{F}(\zeta, \xi) = \left\{ \mathbb{E}(\widehat{\psi})\zeta^1 + \mathbb{E}[(\widehat{\psi} - \widehat{\psi}')\widehat{\psi}^\top]\zeta^2 - \mathbb{E}(c\widehat{\psi}) \right\}^\top \xi^2 + [\zeta^1 - \mathbb{E}(c)] \cdot \xi^1 - 1/2 \cdot \|\xi\|_2^2,$$

where $\widehat{\psi} = \widehat{\psi}(x, u)$ is the estimated feature vector. Here the expectation is only taken over the trajectory generated by the state transition and the policy $\pi_{K,b}$, conditioning on the randomness induced when calculating the estimated feature vectors. Thus, the function $\widehat{F}(\zeta, \xi)$ is still random, where the randomness comes from the estimated feature vectors. Note that $|F(\zeta, \xi) - \widetilde{F}(\zeta, \xi)| \leq |F(\zeta, \xi) - \widehat{F}(\zeta, \xi)| + |\widehat{F}(\zeta, \xi) - \widetilde{F}(\zeta, \xi)|$. Thus, we only need to upper bound $|F(\zeta, \xi) - \widehat{F}(\zeta, \xi)|$ and $|\widehat{F}(\zeta, \xi) - \widetilde{F}(\zeta, \xi)|$.

Part 1. First we upper bound $|F(\zeta, \xi) - \widehat{F}(\zeta, \xi)|$. Note that by algebra, we have

$$\begin{aligned} & |F(\zeta, \xi) - \widehat{F}(\zeta, \xi)| \\ &= \left| \left\{ \mathbb{E}(\psi - \widehat{\psi})\zeta^1 + \mathbb{E}[(\psi - \psi')\psi^\top - (\widehat{\psi} - \widehat{\psi}')\widehat{\psi}^\top]\zeta^2 - \mathbb{E}[c(\psi - \widehat{\psi})] \right\}^\top \xi^2 \right| \\ (B.6.51) \quad & \leq \mathbb{E}(\|\psi - \widehat{\psi}\|_2) \cdot \left[|\zeta^1| + \mathbb{E}(\|\psi - \psi'\|_2 + 2\|\widehat{\psi}\|_2) \cdot \|\zeta^2\|_2 + \mathbb{E}(c) \right] \cdot \|\xi^2\|_2, \end{aligned}$$

where the expectation is only taken over the trajectory generated by the state transition and the policy $\pi_{K,b}$. From Lemma B.4.7, it holds that

$$(B.6.52) \quad \mathbb{P}(\|\widehat{\mu}_z - \mu_z + 1/\widetilde{T} \cdot \mu_{\widetilde{T}}\|_2 \leq C_1) \geq 1 - \widetilde{T}^{-6}.$$

Therefore, combining (B.6.52), it holds with probability at least $1 - \tilde{T}^{-6}$ that

$$(B.6.53) \quad \mathbb{E}(\|\psi - \psi'\|_2 + 2\|\widehat{\psi}\|_2) \leq \text{poly}(\|\Phi_K\|_2, \|K\|_{\text{F}}, \|b\|_2, \|\mu\|_2, J(K_0, b_0)),$$

where the expectation is conditioned on the randomness induced when calculating the estimated feature vectors. Also, we know that

$$(B.6.54) \quad \mathbb{E}(c) \leq \text{poly}(\|\Phi_K\|_2, \|K\|_{\text{F}}, \|b\|_2, \|\mu\|_2, J(K_0, b_0)).$$

Therefore, combining (B.6.51), (B.6.53), (B.6.54), and Definition B.2.7, it holds with probability at least $1 - \tilde{T}^{-6}$ that

$$(B.6.55) \quad |F(\zeta, \xi) - \widehat{F}(\zeta, \xi)| \leq \mathbb{E}(\|\psi - \widehat{\psi}\|_2) \cdot \text{poly}(\|\Phi_K\|_2, \|K\|_{\text{F}}, \|b\|_2, \|\mu\|_2, J(K_0, b_0)).$$

Following from the definitions of $\psi(x, u)$ in (B.2.5) and $\widehat{\psi}(x, u)$ in (B.2.14), we upper bound $\|\psi(x, u) - \widehat{\psi}(x, u)\|_2$ for any x and u as

$$(B.6.56) \quad \begin{aligned} \|\psi(x, u) - \widehat{\psi}(x, u)\|_2^2 &= \|\widehat{\mu}_z - \mu_z\|_2^2 + \|z(\widehat{\mu}_z - \mu_z)^\top + (\widehat{\mu}_z - \mu_z)z^\top\|_{\text{F}}^2 + \|\mu_z\mu_z^\top - \widehat{\mu}_z\widehat{\mu}_z^\top\|_{\text{F}}^2 \\ &\leq \text{poly}(\|\Phi_K\|_2, \|K\|_{\text{F}}, \|b\|_2, \|\mu\|_2, J(K_0, b_0)) \cdot \|\widehat{\mu}_z - \mu_z\|_2^2, \end{aligned}$$

where μ_z is defined in (B.4.47), $\widehat{\mu}_z$ is defined in (B.4.48), and $z = (x^\top, u^\top)^\top$. Also, by Lemma B.4.7, we know that

$$(B.6.57) \quad \|\widehat{\mu}_z - \mu_z\|_2 \leq \frac{\log \tilde{T}}{\tilde{T}^{1/4}} \cdot (1 - \rho)^{-2} \cdot \text{poly}(\|\Phi_K\|_2, \|K\|_{\text{F}}, \|b\|_2, \|\mu\|_2, J(K_0, b_0)),$$

which holds with probability at least $1 - \tilde{T}^{-6}$. Combining (B.6.55), (B.6.56), and (B.6.57), it holds with probability at least $1 - \tilde{T}^{-6}$ that

$$(B.6.58) \quad |F(\zeta, \xi) - \widehat{F}(\zeta, \xi)| \leq \frac{\log \tilde{T}}{\tilde{T}^{1/4}} \cdot (1 - \rho)^{-2} \cdot \text{poly}(\|K\|_F, \|b\|_2, \|\mu\|_2, J(K_0, b_0)).$$

Part 2. We now upper bound $|\widehat{F}(\zeta, \xi) - \widetilde{F}(\zeta, \xi)|$ in the sequel. By definitions, we have

$$(B.6.59) \quad \begin{aligned} & |\widetilde{F}(\zeta, \xi) - \widehat{F}(\zeta, \xi)| \\ &= \left| \left\{ \mathbb{E}(\widetilde{\psi} - \widehat{\psi})\zeta^1 + \mathbb{E}[(\widetilde{\psi} - \widetilde{\psi}')\widetilde{\psi}^\top - (\widehat{\psi} - \widehat{\psi}')\widehat{\psi}^\top]\zeta^2 - \mathbb{E}(\widetilde{c}\widetilde{\psi} - \widehat{c}\widehat{\psi}) \right\}^\top \xi^2 + \mathbb{E}(\widehat{c} - \widetilde{c})\xi^1 \right| \\ &\leq \left| \left\{ \mathbb{E}(\widehat{\psi})\zeta^1 + \mathbb{E}(\widehat{\psi}\widehat{\psi}^\top)\zeta^2 - \mathbb{E}(\widehat{c}\widehat{\psi}) \right\}^\top \xi^2 + \mathbb{E}(\widehat{c})\xi^1 \right| \cdot \mathbf{1}_{\mathcal{E}^c} \\ &\quad + \left| \left[\mathbb{E}(\widehat{\psi}'\widehat{\psi}^\top)\zeta^2 \right]^\top \xi^2 \right| \cdot \mathbf{1}_{(\mathcal{E}' \cap \mathcal{E})^c}, \end{aligned}$$

where we define the event \mathcal{E}' as

$$\mathcal{E}' = \left(\bigcap_{t \in [T]} \left\{ \left| \|z'_t - \mu_z + 1/\tilde{T} \cdot \mu_{\tilde{T}}\|_2^2 - \text{tr}(\widetilde{\Sigma}_z) \right| \leq C_1 \cdot \log T \cdot \|\widetilde{\Sigma}_z\|_2 \right\} \right) \cap \mathcal{E}_2,$$

where \mathcal{E}_2 is defined in (B.4.55). Combining the fact that $\mathbb{P}(\mathcal{E}_2) \geq 1 - \tilde{T}^{-6}$ and Lemma B.7.3, it holds that $\mathbb{P}(\mathcal{E}') \geq 1 - T^{-5} - \tilde{T}^{-6}$. Following a similar argument as in **Part 1**, it holds from (B.6.59) that

$$(B.6.60) \quad |\widetilde{F}(\zeta, \xi) - \widehat{F}(\zeta, \xi)| \leq \left(\frac{1}{T} + \frac{1}{\tilde{T}^{1/4}} \right) \cdot \text{poly}(\|K\|_F, \|b\|_2, \|\mu\|_2, J(K_0, b_0))$$

for sufficiently large T and \tilde{T} .

Now, combining (B.6.58) and (B.6.60), by triangle inequality, it holds with probability at least $1 - \tilde{T}^{-6}$ that

$$|F(\zeta, \xi) - \tilde{F}(\zeta, \xi)| \leq \left(\frac{1}{2\tilde{T}} + \frac{\log \tilde{T}}{\tilde{T}^{1/4}} \right) \cdot (1 - \rho)^{-2} \cdot \text{poly}(\|K\|_F, \|b\|_2, \|\mu\|_2, J(K_0, b_0)).$$

We finish the proof of the lemma. □

B.6.9. Proof of Lemma B.5.2

Proof. Recall that the feature vector $\psi(x, u)$ takes the following form

$$\psi(x, u) = \begin{pmatrix} \text{svec}[(z - \mu_z)(z - \mu_z)^\top] \\ z - \mu_z \end{pmatrix}.$$

We then have

$$(B.6.61) \quad \psi(x, u) - \psi(x', u') = \begin{pmatrix} \text{svec}[yy^\top - (Ly + \delta)(Ly + \delta)^\top] \\ y - (Ly + \delta) \end{pmatrix},$$

where we denote by $y = z - \mu_z$, and (x', u') is the state-action pair after (x, u) following the state transition and the policy $\pi_{K,b}$. Therefore, for any symmetric matrices M, N and

any vectors m, n , it holds from (B.2.7) and (B.6.61) that

$$\begin{aligned}
& \begin{pmatrix} \text{svec}(M) \\ m \end{pmatrix}^\top \Theta_{K,b} \begin{pmatrix} \text{svec}(N) \\ n \end{pmatrix} \\
&= \mathbb{E}_{y,\delta} \left\{ \begin{pmatrix} \text{svec}(M) \\ m \end{pmatrix}^\top \begin{pmatrix} \text{svec}(yy^\top) \\ y \end{pmatrix} \begin{pmatrix} \text{svec}[yy^\top - (Ly + \delta)(Ly + \delta)^\top] \\ y - (Ly + \delta) \end{pmatrix}^\top \begin{pmatrix} \text{svec}(N) \\ n \end{pmatrix} \right\} \\
&= \mathbb{E}_{y,\delta} \left\{ (\langle M, yy^\top \rangle + m^\top y) \cdot [\langle N, yy^\top - (Ly + \delta)(Ly + \delta)^\top \rangle + n^\top (y - Ly - \delta)] \right\} \\
\text{(B.6.62)} & \\
&= \underbrace{\mathbb{E}_y(\langle yy^\top, M \rangle \cdot \langle yy^\top - Lyy^\top L^\top - \Psi_\delta, N \rangle)}_{A_1} + \underbrace{\mathbb{E}_y(\langle yy^\top, M \rangle \cdot n^\top (y - Ly))}_{A_2} \\
&\quad + \underbrace{\mathbb{E}_y(m^\top y \cdot \langle yy^\top - Lyy^\top L^\top - \Psi_\delta, N \rangle)}_{A_3} + \underbrace{\mathbb{E}_y[m^\top y \cdot n^\top (y - Ly)]}_{A_4},
\end{aligned}$$

where the expectations are taken over $y \sim \mathcal{N}(0, \Sigma_z)$ and $\delta \sim \mathcal{N}(0, \Psi_\delta)$. We evaluate the terms A_1, A_2, A_3 , and A_4 in the sequel.

For the terms A_2 and A_3 in (B.6.62), by the fact that $y = z - \mu_z \sim \mathcal{N}(0, \Sigma_z)$, we know that these two terms vanish. For A_4 , it holds that

$$\text{(B.6.63)} \quad A_4 = \mathbb{E}_y[m^\top y \cdot (y - Ly)^\top n] = \mathbb{E}_y[m^\top yy^\top (I - L)^\top n] = m^\top \Sigma_z (I - L)^\top n.$$

For A_1 , by algebra, we have

$$\begin{aligned}
A_1 &= \mathbb{E}_y(\langle yy^\top, M \rangle \cdot \langle yy^\top - Lyy^\top L^\top - \Psi_\delta, N \rangle) \\
&= \mathbb{E}_y(\langle yy^\top, M \rangle \cdot \langle yy^\top - Lyy^\top L^\top, N \rangle) - \mathbb{E}_y(\langle yy^\top, M \rangle \cdot \langle \Psi_\delta, N \rangle) \\
&= \mathbb{E}_y[y^\top M y \cdot y^\top (N - L^\top N L) y] - \langle \Sigma_z, M \rangle \cdot \langle \Psi_\delta, N \rangle \\
\text{(B.6.64)} \quad &= \mathbb{E}_{u \sim \mathcal{N}(0, I)}[u^\top \Sigma_z^{1/2} M \Sigma_z^{1/2} u \cdot u^\top \Sigma_z^{1/2} (N - L^\top N L) \Sigma_z^{1/2} u] - \langle \Sigma_z, M \rangle \cdot \langle \Psi_\delta, N \rangle.
\end{aligned}$$

Now, by applying Lemma B.7.1 to the first term on the RHS of (B.6.64), we know that

$$\begin{aligned}
A_1 &= 2 \operatorname{tr}[\Sigma_z^{1/2} M \Sigma_z^{1/2} \cdot \Sigma_z^{1/2} (N - L^\top N L) \Sigma_z^{1/2}] \\
&\quad + \operatorname{tr}(\Sigma_z^{1/2} M \Sigma_z^{1/2}) \cdot \operatorname{tr}[\Sigma_z^{1/2} (N - L^\top N L) \Sigma_z^{1/2}] - \langle \Sigma_z, M \rangle \cdot \langle \Psi_\delta, N \rangle \\
&= 2 \langle M, \Sigma_z (N - L^\top N L) \Sigma_z \rangle + \langle \Sigma_z, M \rangle \cdot \langle \Sigma_z - L \Sigma_z L^\top - \Psi_\delta, N \rangle \\
&= 2 \langle M, \Sigma_z (N - L^\top N L) \Sigma_z \rangle,
\end{aligned}$$

where we use the fact that $\Sigma_z = L \Sigma_z L^\top + \Psi_\delta$ in the last equality. By using the property of the operator $\operatorname{svec}(\cdot)$ and the definition of the symmetric Kronecker product, we obtain that

$$\begin{aligned}
A_1 &= 2 \operatorname{svec}(M)^\top \operatorname{svec}[\Sigma_z (N - L^\top N L) \Sigma_z] \\
&= 2 \operatorname{svec}(M)^\top [\Sigma_z \otimes_s \Sigma_z - (\Sigma_z L^\top) \otimes_s (\Sigma_z L^\top)] \operatorname{svec}(N) \\
\text{(B.6.65)} \quad &= 2 \operatorname{svec}(M)^\top [(\Sigma_z \otimes_s \Sigma_z)(I - L \otimes_s L)^\top] \operatorname{svec}(N).
\end{aligned}$$

Combining (B.6.62), (B.6.63), and (B.6.65), we obtain that

$$\begin{aligned}
& \begin{pmatrix} \text{svec}(M) \\ m \end{pmatrix}^\top \Theta_{K,b} \begin{pmatrix} \text{svec}(N) \\ n \end{pmatrix} \\
&= \text{svec}(M)^\top [2(\Sigma_z \otimes_s \Sigma_z)(I - L \otimes_s L)^\top] \text{svec}(N) + m^\top \Sigma_z (I - L)^\top n \\
&= \begin{pmatrix} \text{svec}(M) \\ m \end{pmatrix}^\top \begin{pmatrix} 2(\Sigma_z \otimes_s \Sigma_z)(I - L \otimes_s L)^\top & 0 \\ 0 & \Sigma_z (I - L)^\top \end{pmatrix} \begin{pmatrix} \text{svec}(N) \\ n \end{pmatrix}.
\end{aligned}$$

Thus, the matrix $\Theta_{K,b}$ takes the following form,

$$\Theta_{K,b} = \begin{pmatrix} 2(\Sigma_z \otimes_s \Sigma_z)(I - L \otimes_s L)^\top & 0 \\ 0 & \Sigma_z (I - L)^\top \end{pmatrix},$$

which concludes the proof of the lemma. \square

B.6.10. Proof of Lemma B.5.3

Proof. From the definition of $\tilde{\Theta}_{K,b}$ in (B.2.9), it holds that

$$(B.6.66) \quad \|\tilde{\Theta}_{K,b}^{-1}\|_2^2 \leq 1 + \|\Theta_{K,b}^{-1}\|_2^2 + \|\Theta_{K,b}^{-1} \tilde{\sigma}_z\|_2^2,$$

where $\tilde{\sigma}_z$ is defined as

$$\tilde{\sigma}_z = \mathbb{E}_{\pi_{K,b}} [\psi(x, u)] = \begin{pmatrix} \text{svec}(\Sigma_z) \\ \mathbf{0}_{k+m} \end{pmatrix}.$$

We bound the RHS of (B.6.66) in the sequel. For the term $\Theta_{K,b}^{-1}\tilde{\sigma}_z$, combining Lemma B.5.2, we have

$$\begin{aligned}
\Theta_{K,b}^{-1}\tilde{\sigma}_z &= \begin{pmatrix} 1/2 \cdot (I - L \otimes_s L)^{-\top} (\Sigma_z \otimes_s \Sigma_z)^{-1} \cdot \text{svec}(\Sigma_z) \\ \mathbf{0}_{k+m} \end{pmatrix} \\
&= \begin{pmatrix} 1/2 \cdot (I - L \otimes_s L)^{-\top} (\Sigma_z^{-1} \otimes_s \Sigma_z^{-1}) \cdot \text{svec}(\Sigma_z) \\ \mathbf{0}_{k+m} \end{pmatrix} \\
\text{(B.6.67)} \quad &= \begin{pmatrix} 1/2 \cdot (I - L \otimes_s L)^{-\top} \cdot \text{svec}(\Sigma_z^{-1}) \\ \mathbf{0}_{k+m} \end{pmatrix},
\end{aligned}$$

where we use the property of the symmetric Kronecker product in the second and last line. By taking the spectral norm on both sides of (B.6.67), it holds that

$$\begin{aligned}
\|\Theta_{K,b}^{-1}\tilde{\sigma}_z\|_2 &= 1/2 \cdot \|(I - L \otimes_s L)^{-\top} \cdot \text{svec}(\Sigma_z^{-1})\|_2 \\
&\leq 1/2 \cdot \|(I - L \otimes_s L)^{-\top}\|_2 \cdot \|\text{svec}(\Sigma_z^{-1})\|_2 \\
&\leq 1/2 \cdot [1 - \rho^2(L)]^{-1} \cdot \|\Sigma_z^{-1}\|_{\text{F}} \\
&\leq 1/2 \cdot \sqrt{k+m} \cdot [1 - \rho^2(L)]^{-1} \cdot \|\Sigma_z^{-1}\|_2 \\
\text{(B.6.68)} \quad &= 1/2 \cdot \sqrt{k+m} \cdot [1 - \rho^2(L)]^{-1} \cdot \sigma_{\min}^{-1}(\Sigma_z),
\end{aligned}$$

where in the third line we use Lemma B.7.2 to the matrix $L \otimes_s L$. Similarly, we upper bound $\|\Theta_{K,b}^{-1}\|_2$ in the sequel

$$\text{(B.6.69)} \quad \|\Theta_{K,b}^{-1}\|_2 \leq \min \left\{ 1/2 \cdot [1 - \rho^2(L)]^{-1} \sigma_{\min}^{-2}(\Sigma_z), [1 - \rho(L)]^{-1} \sigma_{\min}^{-1}(\Sigma_z) \right\}.$$

Thus, combining (B.6.66), (B.6.68), and (B.6.69), we obtain that

$$(B.6.70) \quad \begin{aligned} \|\tilde{\Theta}_{K,b}^{-1}\|_2^2 &\leq 1 + 1/2 \cdot \sqrt{k+m} \cdot [1 - \rho^2(L)]^{-1} \cdot \sigma_{\min}^{-1}(\Sigma_z) \\ &+ \min\left\{1/2 \cdot [1 - \rho^2(L)]^{-1} \sigma_{\min}^{-2}(\Sigma_z), [1 - \rho(L)]^{-1} \sigma_{\min}^{-1}(\Sigma_z)\right\}. \end{aligned}$$

Now it remains to characterize $\sigma_{\min}(\Sigma_z)$. For any vectors $s \in \mathbb{R}^m$ and $r \in \mathbb{R}^k$, we have

$$(B.6.71) \quad \begin{aligned} \begin{pmatrix} s \\ r \end{pmatrix}^\top \Sigma_z \begin{pmatrix} s \\ r \end{pmatrix} &= \mathbb{E}_{x \sim \mathcal{N}(\mu_{K,b}, \Phi_K), u \sim \pi_{K,b}(\cdot | x)} \left\{ [s^\top (x - \mu_{K,b}) + r^\top (u + K\mu_{K,b} - b)]^2 \right\} \\ &= \mathbb{E}_{x \sim \mathcal{N}(\mu_{K,b}, \Phi_K), \eta \sim \mathcal{N}(0, I)} \left\{ [(s - K^\top r)^\top (x - \mu_{K,b}) + \sigma r^\top \eta]^2 \right\} \\ &= \mathbb{E}_{x \sim \mathcal{N}(\mu_{K,b}, \Phi_K)} \left\{ [(s - K^\top r)^\top (x - \mu_{K,b})]^2 \right\} + \mathbb{E}_{\eta \sim \mathcal{N}(0, I)} [(\sigma r^\top \eta)^2]. \end{aligned}$$

The first term on the RHS of (B.6.71) is lower bounded as

$$(B.6.72) \quad \begin{aligned} \mathbb{E}_{x \sim \mathcal{N}(\mu_{K,b}, \Phi_K)} \left\{ [(s - K^\top r)^\top (x - \mu_{K,b})]^2 \right\} &= (s - K^\top r)^\top \Phi_K (s - K^\top r) \\ &\geq \|s - K^\top r\|_2^2 \cdot \sigma_{\min}(\Phi_K) \geq \|s - K^\top r\|_2^2 \cdot \sigma_{\min}(\Psi_\omega), \end{aligned}$$

where the last inequality comes from the fact that $\sigma_{\min}(\Phi_K) \geq \sigma_{\min}(\Psi_\omega)$ by (3.2.3). The second term on the RHS of (B.6.71) takes the form of

$$(B.6.73) \quad \mathbb{E}_{\eta \sim \mathcal{N}(0, I)} [(\sigma r^\top \eta)^2] = \sigma^2 \|r\|_2^2.$$

Therefore, combining (B.6.71), (B.6.72), and (B.6.73), we have

$$\begin{aligned} \begin{pmatrix} s \\ r \end{pmatrix}^\top \Sigma_z \begin{pmatrix} s \\ r \end{pmatrix} &\geq \|s - K^\top r\|_2^2 \cdot \sigma_{\min}(\Psi_\omega) + \sigma^2 \|r\|_2^2 \\ &\geq \sigma_{\min}(\Psi_\omega) \cdot \|s\|_2^2 + [\sigma^2 - \|K\|_2^2 \cdot \sigma_{\min}(\Psi_\omega)] \cdot \|r\|_2^2. \end{aligned}$$

From this, we know that

$$(B.6.74) \quad \sigma_{\min}(\Sigma_z) \geq \min\{\sigma_{\min}(\Psi_\omega), \sigma^2 - \|K\|_2^2 \cdot \sigma_{\min}(\Psi_\omega)\}.$$

Thus, combining (B.6.70) and (B.6.74), we know that $\|\tilde{\Theta}_{K,b}^{-1}\|_2$ is upper bounded by a constant $\tilde{\lambda}_K$, where $\tilde{\lambda}_K$ only depends on $\|K\|_2$ and $\rho(L) = \rho(A - BK)$. This finishes the proof of the lemma. \square

B.6.11. Proof of Lemma B.6.2

Proof. First, note that the cost function $c(x, u)$ takes the following form,

$$c(x, u) = \psi(x, u)^\top \begin{pmatrix} \text{svec}[\text{diag}(Q, R)] \\ 2Q\mu_{K,b} \\ 2R\mu_{K,b}^u \end{pmatrix} + [\mu_{K,b}^\top Q \mu_{K,b} + (\mu_{K,b}^u)^\top R \mu_{K,b}^u + \mu^\top \bar{Q} \mu].$$

For any matrix V and vectors v_x, v_u , it holds that

$$\begin{aligned}
 & \mathbb{E}_{\pi_{K,b}} [c(x, u) \psi(x, u)]^\top \begin{pmatrix} \text{svec}(V) \\ v_x \\ v_u \end{pmatrix} \\
 \text{(B.6.75)} \quad & = \underbrace{\mathbb{E}_{\pi_{K,b}} \left\{ \psi(x, u)^\top \begin{pmatrix} \text{svec}[\text{diag}(Q, R)] \\ 2Q\mu_{K,b} \\ 2R\mu_{K,b}^u \end{pmatrix} \psi(x, u)^\top \begin{pmatrix} \text{svec}(V) \\ v_x \\ v_u \end{pmatrix} \right\}}_{D_1} \\
 & \quad + \underbrace{\mathbb{E}_{\pi_{K,b}} \left\{ \psi(x, u)^\top (\mu_{K,b}^\top Q \mu_{K,b} + (\mu_{K,b}^u)^\top R \mu_{K,b}^u + \mu^\top \bar{Q} \mu) \begin{pmatrix} \text{svec}(V) \\ v_x \\ v_u \end{pmatrix} \right\}}_{D_2}.
 \end{aligned}$$

In the sequel, we calculate D_1 and D_2 respectively.

Calculation of D_1 . Note that by the definition of $\psi(x, u)$ in (B.2.5), it holds that

$$\begin{aligned}
D_1 &= \mathbb{E}_{\pi_{K,b}} \left\{ \left[(z - \mu_z)^\top \text{diag}(Q, R)(z - \mu_z) + (z - \mu_z)^\top \begin{pmatrix} 2Q\mu_{K,b} \\ 2R\mu_{K,b}^u \end{pmatrix} \right] \right. \\
&\quad \left. \cdot \left[(z - \mu_z)^\top V(z - \mu_z) + (z - \mu_z)^\top \begin{pmatrix} v_x \\ v_u \end{pmatrix} \right] \right\} \\
\text{(B.6.76)} \quad &= \mathbb{E}_{\pi_{K,b}} \left[(z - \mu_z)^\top \text{diag}(Q, R)(z - \mu_z) \cdot (z - \mu_z)^\top V(z - \mu_z) \right] \\
&\quad + \mathbb{E}_{\pi_{K,b}} \left[\begin{pmatrix} 2Q\mu_{K,b} \\ 2R\mu_{K,b}^u \end{pmatrix}^\top (z - \mu_z)(z - \mu_z)^\top \begin{pmatrix} v_x \\ v_u \end{pmatrix} \right].
\end{aligned}$$

Here $z = (x^\top, u^\top)^\top$ and $\mu_z = \mathbb{E}_{\pi_{K,b}}(z)$. For the first term on the RHS of (B.6.76), note that $z - \mu_z \sim \mathcal{N}(0, \Sigma_z)$. Therefore, by Lemma B.7.1, we obtain that

$$\begin{aligned}
&\mathbb{E}_{\pi_{K,b}} \left[(z - \mu_z)^\top \text{diag}(Q, R)(z - \mu_z) \cdot (z - \mu_z)^\top V(z - \mu_z) \right] \\
&= 2\langle \Sigma_z \text{diag}(Q, R) \Sigma_z, V \rangle + \langle \Sigma_z, \text{diag}(Q, R) \rangle \cdot \langle \Sigma_z, V \rangle \\
\text{(B.6.77)} \quad &= \text{svec} \left[2\Sigma_z \text{diag}(Q, R) \Sigma_z + \langle \Sigma_z, \text{diag}(Q, R) \rangle \cdot \Sigma_z \right]^\top \text{svec}(V).
\end{aligned}$$

Meanwhile, the second term on the RHS of (B.6.76) takes the form of

$$\text{(B.6.78)} \quad \mathbb{E}_{\pi_{K,b}} \left[\begin{pmatrix} 2Q\mu_{K,b} \\ 2R\mu_{K,b}^u \end{pmatrix}^\top (z - \mu_z)(z - \mu_z)^\top \begin{pmatrix} v_x \\ v_u \end{pmatrix} \right] = \left[\Sigma_z \begin{pmatrix} 2Q\mu_{K,b} \\ 2R\mu_{K,b}^u \end{pmatrix} \right]^\top \begin{pmatrix} v_x \\ v_u \end{pmatrix}.$$

Combining (B.6.76), (B.6.77), and (B.6.78), we obtain that

$$(B.6.79) \quad D_1 = \begin{pmatrix} 2\text{svec}[\Sigma_z \text{diag}(Q, R)\Sigma_z + \langle \Sigma_z, \text{diag}(Q, R) \rangle \Sigma_z] \\ \Sigma_z \begin{pmatrix} 2Q\mu_{K,b} \\ 2R\mu_{K,b}^u \end{pmatrix} \end{pmatrix}^\top \begin{pmatrix} \text{svec}(V) \\ v_x \\ v_u \end{pmatrix}.$$

Calculation of D_2 . By the definition of the feature vector $\psi(x, u)$ in (B.2.5), we know that

$$(B.6.80) \quad D_2 = (\mu_{K,b}^\top Q \mu_{K,b} + (\mu_{K,b}^u)^\top R \mu_{K,b}^u + \mu^\top \bar{Q} \mu) \begin{pmatrix} \text{svec}(\Sigma_z) \\ \mathbf{0}_m \\ \mathbf{0}_k \end{pmatrix}^\top \begin{pmatrix} \text{svec}(V) \\ v_x \\ v_u \end{pmatrix}.$$

Now, combining (B.6.75), (B.6.79), and (B.6.80), it holds that

$$\begin{aligned} \mathbb{E}_{\pi_{K,b}}[c(x, u)\psi(x, u)] &= \begin{pmatrix} 2\text{svec}[\Sigma_z \text{diag}(Q, R)\Sigma_z + \langle \Sigma_z, \text{diag}(Q, R) \rangle \Sigma_z] \\ \Sigma_z \begin{pmatrix} 2Q\mu_{K,b} \\ 2R\mu_{K,b}^u \end{pmatrix} \end{pmatrix} \\ &\quad + [\mu_{K,b}^\top Q \mu_{K,b} + (\mu_{K,b}^u)^\top R \mu_{K,b}^u + \mu^\top \bar{Q} \mu] \begin{pmatrix} \text{svec}(\Sigma_z) \\ \mathbf{0}_m \\ \mathbf{0}_k \end{pmatrix}, \end{aligned}$$

which concludes the proof of the lemma. \square

B.7. Auxiliary Results

Lemma B.7.1. Assume that the random variable $w \sim \mathcal{N}(0, I)$, and let U and V be two symmetric matrices, then it holds that

$$\mathbb{E}(w^\top U w \cdot w^\top V w) = 2 \operatorname{tr}(UV) + \operatorname{tr}(U) \cdot \operatorname{tr}(V).$$

Proof. See Magnus et al. (1978) and Magnus (1979) for a detailed proof. \square

Lemma B.7.2. Let M, N be commuting symmetric matrices, and let $\alpha_1, \dots, \alpha_n, \beta_1, \dots, \beta_n$ denote their eigenvalues with v_1, \dots, v_n a common basis of orthogonal eigenvectors. Then the $n(n+1)/2$ eigenvalues of $M \otimes_s N$ are given by $(\alpha_i \beta_j + \alpha_j \beta_i)/2$, where $1 \leq i \leq j \leq n$.

Proof. See Lemma 2 in Alizadeh et al. (1998) for a detailed proof. \square

Lemma B.7.3. For any integer $m > 0$, let $A \in \mathbb{R}^{m \times m}$ and $\eta \sim \mathcal{N}(0, I_m)$. Then, there exists some absolute constant $C > 0$ such that for any $t \geq 0$, we have

$$\mathbb{P}\left[|\eta^\top A \eta - \mathbb{E}(\eta^\top A \eta)| > t\right] \leq 2 \cdot \exp\left[-C \cdot \min(t^2 \|A\|_{\text{F}}^{-2}, t \|A\|_2^{-1})\right].$$

Proof. See Rudelson et al. (2013) for a detailed proof. \square

APPENDIX C

Supplemental Materials in Chapter 4

C.1. Analysis of Maximum Likelihood Estimation

We denote by

$$L_1(\Theta) = -\mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} \log \Theta(S_t, Z_t) \right]$$

the population counterpart of \widehat{L}_1 . We define

$$(C.1.1) \quad H^2(\Theta_1, \Theta_2) = \frac{1}{2} \cdot \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} \int \left(\sqrt{\Theta_1(S_t, z)} - \sqrt{\Theta_2(S_t, z)} \right)^2 dz \right].$$

We have the following supporting result.

Lemma C.1.1. Under Assumption (b), for any $\Theta_1, \Theta_2 \in \mathcal{F}_1$, it holds with probability at least $1 - \delta$ for any $c/(NT)^2 \leq \delta \leq 1$ that

$$\begin{aligned} & \left| (L_1(\Theta_1) - L_1(\Theta_2)) - (\widehat{L}_1(\Theta_1) - \widehat{L}_1(\Theta_2)) \right| \\ & \leq c \cdot \frac{\log C_{\Theta^*}}{NT\kappa} \log \frac{1}{\delta} \log(NT) + c \cdot \sqrt{\frac{C_{\Theta^*}}{NT\kappa} H^2(\Theta_1, \Theta_2) \log \frac{1}{\delta} \log(NT)}, \end{aligned}$$

where $H^2(\Theta_1, \Theta_2)$ is defined in (C.1.1).

Proof. By Theorem C.7.7, it holds with probability at least $1 - \delta$ that

$$\begin{aligned}
& \left| (L_1(\Theta_1) - L_1(\Theta_2)) - (\widehat{L}_1(\Theta_1) - \widehat{L}_1(\Theta_2)) \right| \\
\text{(C.1.2)} \quad & \leq c \cdot \frac{\log C_{\Theta^*}}{NT\kappa} \log \frac{1}{\delta} \log(NT) + c \cdot \sqrt{\frac{1}{NT\kappa} \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} \left(\log \frac{\Theta_1(S_t, Z_t)}{\Theta_2(S_t, Z_t)} \right)^2 \right]} \log \frac{1}{\delta} \log(NT)}.
\end{aligned}$$

Now, it suffices to upper bound the variance term on the RHS of the above inequality.

Note that $\log x \leq 2(\sqrt{x} - 1)$ for any $x > 0$. Thus, for any $s \in \mathcal{S}$, we have

$$\begin{aligned}
& \int \Theta^*(s, z) \left(\log \frac{\Theta_1(s, z)}{\Theta_2(s, z)} \right)^2 dz \\
& \leq 4 \int \Theta^* \max \left\{ \left(\sqrt{\frac{\Theta_2}{\Theta_1}} - 1 \right)^2, \left(\sqrt{\frac{\Theta_1}{\Theta_2}} - 1 \right)^2 \right\} dz \\
& = 4 \int \max \left\{ \frac{\Theta^*}{\Theta_1} \left(\sqrt{\Theta_2} - \sqrt{\Theta_1} \right)^2, \frac{\Theta^*}{\Theta_2} \left(\sqrt{\Theta_1} - \sqrt{\Theta_2} \right)^2 \right\} dz \\
\text{(C.1.3)} \quad & \leq 4C_{\Theta^*} \int \left(\sqrt{\Theta_1(s, z)} - \sqrt{\Theta_2(s, z)} \right)^2 dz,
\end{aligned}$$

which implies that

$$\text{(C.1.4)} \quad \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} \left(\log \frac{\Theta_1(S_t, Z_t)}{\Theta_2(S_t, Z_t)} \right)^2 \right] \leq 8C_{\Theta^*} H^2(\Theta_1, \Theta_2).$$

By plugging (C.1.4) into (C.1.2), we conclude the proof of the lemma. \square

C.1.1. Proof of Theorem 4.4.5

Proof. Proof of the first statement. It suffices to show that with probability at least $1 - \delta$, we have

$$\widehat{L}_1(\Theta^*) - \widehat{L}_1(\widehat{\Theta}) \leq \alpha_1.$$

By Corollary C.7.10, it holds with probability at least $1 - \delta$ that

$$(C.1.5) \quad H^2(\Theta^*, \widehat{\Theta}) \leq c \cdot \frac{d}{NT\kappa} \log \frac{\theta_{\max}}{\delta},$$

where $c > 0$ is an absolute constant, which may vary from lines to lines. Thus, by Lemma C.1.1, it holds with probability at least $1 - \delta$ that

$$(C.1.6) \quad \begin{aligned} & \left| \left(L_1(\Theta^*) - L_1(\widehat{\Theta}) \right) - \left(\widehat{L}_1(\Theta^*) - \widehat{L}_1(\widehat{\Theta}) \right) \right| \\ & \leq c \cdot \frac{\log C_{\Theta^*}}{NT\kappa} d \log \frac{\theta_{\max}}{\delta} \log(NT) + c \cdot \sqrt{\frac{C_{\Theta^*}}{NT\kappa} H^2(\Theta^*, \widehat{\Theta}) d \log \frac{\theta_{\max}}{\delta} \log(NT)} \\ & \leq c \cdot \frac{C_{\Theta^*} d}{NT\kappa} \log \frac{\theta_{\max}}{\delta} \log(NT), \end{aligned}$$

where we use a covering argument and (C.1.5) in the first and last inequalities, respectively.

Further, by a similar idea as in (C.1.3), we upper bound $|L_1(\Theta^*) - L_1(\widehat{\Theta})|$ as follows,

$$(C.1.7) \quad \begin{aligned} |L_1(\Theta^*) - L_1(\widehat{\Theta})| &= \left| \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} \int \Theta^*(S_t, z) \log \frac{\Theta^*(S_t, z)}{\widehat{\Theta}(S_t, z)} dz \right] \right| \\ &\leq 2C_{\Theta^*} H^2(\Theta^*, \widehat{\Theta}) \leq c \cdot \frac{C_{\Theta^*} d}{NT\kappa} \log \frac{\theta_{\max}}{\delta}, \end{aligned}$$

where we use (C.1.5) and Corollary C.7.10 in the first and last inequalities, respectively.

Now, by combining (C.1.6) and (C.1.7), it holds with probability at least $1 - \delta$ that

$$\widehat{L}_1(\Theta^*) - \widehat{L}_1(\widehat{\Theta}) \leq c \cdot \frac{C_{\Theta^*} d}{NT\kappa} \log \frac{\theta_{\max}}{\delta} \log(NT) = \alpha_1,$$

which concludes the proof of the first statement.

Proof of the second statement. By Lemma C.1.1, with probability at least $1 - \delta$, for any $\Theta \in \text{conf}_{\alpha_1}^1$, we have

$$(C.1.8) \quad \begin{aligned} & \left| (L_1(\Theta^*) - L_1(\Theta)) - (\widehat{L}_1(\Theta^*) - \widehat{L}_1(\Theta)) \right| \\ & \leq c \cdot \frac{\log C_{\Theta^*}}{NT\kappa} d \log \frac{\theta_{\max}}{\delta} \log(NT) + c \cdot \sqrt{\frac{C_{\Theta^*}}{NT\kappa} H^2(\Theta^*, \Theta) d \log \frac{\theta_{\max}}{\delta} \log(NT)}, \end{aligned}$$

where we use a covering argument. In the meanwhile, by the first statement, we have

$\Theta^* \in \text{conf}_{\alpha_1}^1$ with probability at least $1 - \delta$. Thus, we have

$$(C.1.9) \quad \left| \widehat{L}_1(\Theta^*) - \widehat{L}_1(\Theta) \right| \leq \left| \widehat{L}_1(\Theta^*) - \widehat{L}_1(\widehat{\Theta}) \right| + \left| \widehat{L}_1(\widehat{\Theta}) - \widehat{L}_1(\Theta) \right| \leq 2\alpha_1,$$

where we use the fact that $\Theta \in \text{conf}_{\alpha_1}^1$. By combining (C.1.8) and (C.1.9), with probability at least $1 - \delta$, it holds for any $\Theta \in \text{conf}_{\alpha_1}^1$ that

$$(C.1.10) \quad \begin{aligned} & L_1(\Theta) - L_1(\Theta^*) \\ & \leq c \cdot \frac{\log C_{\Theta^*}}{NT\kappa} d \log \frac{\theta_{\max}}{\delta} \log(NT) + c \cdot \sqrt{\frac{C_{\Theta^*}}{NT\kappa} H^2(\Theta^*, \Theta) d \log \frac{\theta_{\max}}{\delta} \log(NT)}. \end{aligned}$$

On the other hand, it holds for any $s \in \mathcal{S}$ that

$$\begin{aligned} - \int \Theta^*(s, z) \log \frac{\Theta(s, z)}{\Theta^*(s, z)} dz &\geq -2 \int \Theta^*(s, z) \left(\sqrt{\frac{\Theta(s, z)}{\Theta^*(s, z)}} - 1 \right) dz \\ &= \int \left(\Theta^*(s, z) + \Theta(s, z) - 2\sqrt{\Theta(s, z)\Theta^*(s, z)} \right) dz \\ &= \int \left(\sqrt{\Theta^*(s, z)} - \sqrt{\Theta(s, z)} \right)^2 dz, \end{aligned}$$

which implies that

$$(C.1.11) \quad L_1(\Theta) - L_1(\Theta^*) \geq 2H^2(\Theta^*, \Theta).$$

By combining (C.1.10) and (C.1.11), we have

$$H^2(\Theta^*, \Theta) \leq c \cdot \frac{\log C_{\Theta^*}}{NT\kappa} d \log \frac{\theta_{\max}}{\delta} \log(NT) + c \cdot \sqrt{\frac{C_{\Theta^*}}{NT\kappa} H^2(\Theta^*, \Theta) d \log \frac{\theta_{\max}}{\delta} \log(NT)},$$

which implies that

$$H^2(\Theta^*, \Theta) \leq c \cdot \frac{C_{\Theta^*} d}{NT\kappa} \log \frac{\theta_{\max}}{\delta} \log(NT).$$

Now, by Lemma C.7.8, with probability at least $1 - \delta$, it holds for any $\Theta \in \text{conf}_{\alpha_1}^1$ that

$$\sqrt{\mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} \|\Theta(S_t, \cdot) - \Theta^*(S_t, \cdot)\|_1^2 \right]} \leq c \cdot \sqrt{\frac{C_{\Theta^*} d}{NT\kappa} \log \frac{\theta_{\max}}{\delta} \log(NT)},$$

which concludes the proof of the second statement. \square

C.2. Proofs of Results in §4.2

We provide proofs of results in §4.2. We first present proofs for §4.2.2, then we present proofs for §4.2.1.

C.2.1. Proof of Lemma 4.2.7

Proof. By Assumption (b), we know that

$$\mathbb{E} \left[\frac{Z_t^\top A_t \pi(A_t | S_t)}{\Delta^*(S_t, A_t) \Theta^*(S_t, Z_t)} w^\pi(S_t) R_t \mid S_t \right]$$

is well-defined. Further, we observe for any $t \in \{0, 1, \dots, T-1\}$ that

$$\begin{aligned} & \mathbb{E} \left[\frac{Z_t^\top A_t \pi(A_t | S_t)}{\Delta^*(S_t, A_t) \Theta^*(S_t, Z_t)} w^\pi(S_t) R_t \mid S_t \right] \\ &= \sum_{a \in \mathcal{A}} \mathbb{E} \left[\frac{Z_t^\top a \pi(a | S_t) d^\pi(S_t) \mathbf{1}\{A_t = a\}}{\Delta^*(S_t, a) \Theta^*(S_t, Z_t) d^b(S_t)} R(S_t, U_t, a, S_{t+1}, U_{t+1}) \mid S_t \right] \\ &= \sum_{a \in \mathcal{A}} \mathbb{E} \left[\frac{Z_t^\top a \pi(a | S_t) d^\pi(S_t) \mathbf{1}\{A_t = a\}}{\Delta^*(S_t, a) \Theta^*(S_t, Z_t) d^b(S_t)} r(S_t, U_t, a) \mid S_t \right] \\ &= \sum_{a \in \mathcal{A}} \mathbb{E} \left[\frac{Z_t^\top a \pi(a | S_t) d^\pi(S_t) \mathbb{P}(A_t = a | S_t, U_t, Z_t)}{\Delta^*(S_t, a) d^b(S_t) \Theta^*(S_t, Z_t)} r(S_t, U_t, a) \mid S_t \right] \\ &= \sum_{a \in \mathcal{A}} \mathbb{E} \left[\frac{\pi(a | S_t) d^\pi(S_t) \mathbb{P}(A_t = a | S_t, U_t, Z_t = a)}{\Delta^*(S_t, a) d^b(S_t)} r(S_t, U_t, a) \mid S_t \right] \\ &\quad - \sum_{a \in \mathcal{A}} \sum_{z \in \mathcal{Z}, z \neq a} \frac{1}{K-1} \mathbb{E} \left[\frac{\pi(a | S_t) d^\pi(S_t) \mathbb{P}(A_t = a | S_t, U_t, Z_t = z)}{\Delta^*(S_t, a) d^b(S_t)} r(S_t, U_t, a) \mid S_t \right] \\ &= \sum_{a \in \mathcal{A}} \frac{\pi(a | S_t) d^\pi(S_t)}{d^b(S_t)} \mathbb{E}_{U_t} [r(S_t, U_t, a) \mid S_t], \end{aligned}$$

where in the first equality, we use the definition of $w^\pi(s)$; in the second equality, we use Assumption (a); in the third equality, we use Assumption (b); in the fourth equality, we use

Assumption (c); while in the fifth equality, we use Assumption (d). Now, by Assumption 4.2.6, we know that $\mathbb{E}_{U_t}[r(S_t, U_t, a) | S_t] = \tilde{r}(S_t, a)$, where the function \tilde{r} is independent of t . Therefore, we have

$$\begin{aligned}
& \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} \frac{Z_t^\top A_t \pi(A_t | S_t)}{\Delta^*(S_t, A_t) \Theta^*(S_t, Z_t)} w^\pi(S_t) R_t \right] \\
&= \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \left[\sum_{a \in \mathcal{A}} \frac{\pi(a | S_t) d^\pi(S_t)}{d^b(S_t)} \tilde{r}(S_t, a) \right] \\
&= \frac{1}{T} \sum_{t=0}^{T-1} \sum_{a \in \mathcal{A}} \int \frac{\pi(a | s) d^\pi(s)}{d^b(s)} \tilde{r}(s, a) p_t^b(s) ds \\
&= \sum_{a \in \mathcal{A}} \int \pi(a | s) d^\pi(s) \tilde{r}(s, a) ds = J(\pi),
\end{aligned}$$

which concludes the proof of the lemma. \square

C.2.2. Proof of Lemma 4.2.8

Proof. Similar to the proof of Lemma 4.2.7 in §C.2.1, we observe that

$$(C.2.1) \quad \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} \frac{Z_t^\top A_t d^\pi(S_t) \pi(A_t | S_t)}{\Delta(S_t, A_t) d^b(S_t) P(Z_t | S_t)} f(S_t) \right] = \int f(s') d^\pi(s') ds'.$$

Similarly, we have

$$\begin{aligned}
(C.2.2) \quad & \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} \frac{Z_t^\top A_t d^\pi(S_t) \pi(A_t | S_t)}{\Delta(S_t, A_t) d^b(S_t) P(Z_t | S_t)} f(S_{t+1}) \right] \\
&= \int f(s') d^\pi(s) \pi(a | s) \mathbb{P}(S' = s' | S = s, A = a) ds' da ds.
\end{aligned}$$

Meanwhile, by the definition of $d^\pi(s, a)$, we have

$$\begin{aligned}
d^\pi(s') &= (1 - \gamma) \sum_{t=0}^{\infty} \gamma^t p_t^\pi(s') \\
&= (1 - \gamma) \nu(s') + (1 - \gamma) \sum_{t=0}^{\infty} \gamma^{t+1} p_{t+1}^\pi(s') \\
&= (1 - \gamma) \left(\nu(s') + \gamma \sum_{t=0}^{\infty} \gamma^t \int \mathbb{P}(S_{t+1} = s' \mid S_t = s, A_t = a) \pi(a \mid s) p_t^\pi(s) ds da \right) \\
\text{(C.2.3)} \quad &= (1 - \gamma) \nu(s') + \gamma \int \mathbb{P}(S' = s' \mid S = s, A = a) \pi(a \mid s) d^\pi(s) ds da,
\end{aligned}$$

where we use the assumption that $S_{t+1} \mid (S_t, A_t)$ is time-homogeneous. Combining (C.2.1) and (C.2.2), we have

$$\begin{aligned}
&\mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} \frac{Z_t^\top A_t d^\pi(S_t) \pi(A_t \mid S_t)}{\Delta(S_t, A_t) d^b(S_t) P(Z_t \mid S_t)} (f(S_t) - \gamma f(S_{t+1})) \right] \\
&= \int f(s') \left(d^\pi(s') - \gamma \int d^\pi(s) \pi(a \mid s) \mathbb{P}(S' = s' \mid S = s, A = a) ds da \right) ds' \\
&= (1 - \gamma) \int f(s') \nu(s') ds' = (1 - \gamma) \mathbb{E}_{S \sim \nu} [f(S)],
\end{aligned}$$

where we use (C.2.3) in the fourth equality. This concludes the proof of the lemma. \square

C.2.3. Proof of Lemma 4.2.4

Proof. Similar to the proof of Lemma 4.2.7 in §C.2.1, we observe that

$$\begin{aligned}
&\mathbb{E} \left[\frac{Z_0^\top A_0 \pi(A_0 \mid S_0)}{\Delta^*(S_0, A_0) P(Z_0 \mid S_0)} R_0 \mid S_0 = s \right] \\
\text{(C.2.4)} \quad &= \mathbb{E}_{U_0} \left[\sum_{a \in \mathcal{A}} \pi(a \mid S_0) r(S_0, U_0, a) \mid S_0 = s \right] = \mathbb{E}_\pi [R_0 \mid S_0 = s],
\end{aligned}$$

where $\mathbb{E}_\pi[\cdot]$ denotes that the expectation is taken w.r.t. $A_0 \sim \pi(\cdot | S_0)$. Similarly, by Assumption 4.2.3, we observe that

$$\begin{aligned}
& \mathbb{E} \left[\frac{Z_0^\top A_0 \pi(A_0 | S_0)}{\Delta^*(S_0, A_0) P(Z_0 | S_0)} \frac{Z_1^\top A_1 \pi(A_1 | S_1)}{\Delta^*(S_1, A_1) P(Z_1 | S_1)} R_1 \mid S_0 = s \right] \\
&= \mathbb{E} \left[\sum_{a \in \mathcal{A}} \pi(a | S_0) \mathbb{E} \left[\frac{Z_1^\top A_1 \pi(A_1 | S_1)}{\Delta^*(S_1, A_1) P(Z_1 | S_1)} R_1 \mid S_0, U_0, A_0 = a \right] \mid S_0 = s \right] \\
&= \mathbb{E} \left[\sum_{a \in \mathcal{A}} \pi(a | S_0) \mathbb{E} \left[\mathbb{E} \left[\frac{Z_1^\top A_1 \pi(A_1 | S_1)}{\Delta^*(S_1, A_1) P(Z_1 | S_1)} R_1 \mid S_1, U_1 \right] \mid S_0, U_0, A_0 = a \right] \mid S_0 = s \right] \\
&= \mathbb{E} \left[\sum_{a \in \mathcal{A}} \pi(a | S_0) \mathbb{E} [\mathbb{E}_\pi[R_1 | S_1, U_1] | S_0, U_0, A_0 = a] \mid S_0 = s \right] = \mathbb{E}_\pi[R_1 | S_0 = s].
\end{aligned}$$

Now, by induction, it holds for any $t \geq 0$ that

$$\mathbb{E} \left[R_t \left(\prod_{j=0}^t \frac{Z_j^\top A_j \pi(A_j | S_j)}{\Delta^*(S_j, A_j) P(Z_j | S_j)} \right) \mid S_0 = s \right] = \mathbb{E}_\pi[R_t | S_0 = s],$$

which implies that

$$V^\pi(s) = \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t R_t \left(\prod_{j=0}^t \frac{Z_j^\top A_j \pi(A_j | S_j)}{\Delta^*(S_j, A_j) P(Z_j | S_j)} \right) \mid S_0 = s \right].$$

To show that the IV-aided Bellman equation holds, by a similar argument as in (C.2.4), we observe that

$$\begin{aligned}
& \mathbb{E} \left[\frac{Z_0^\top A_0 \pi(A_0 | S_0)}{\Delta^*(S_0, A_0) P(Z_0 | S_0)} \gamma V^\pi(S_1) \mid S_0 = s \right] \\
&= \mathbb{E} \left[\frac{Z_0^\top A_0 \pi(A_0 | S_0)}{\Delta^*(S_0, A_0) P(Z_0 | S_0)} \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^{t+1} R_{t+1} \left(\prod_{j=1}^{t+1} \frac{Z_j^\top A_j \pi(A_j | S_j)}{\Delta^*(S_j, A_j) P(Z_j | S_j)} \right) \mid S_1 \right] \mid S_0 = s \right] \\
&= \mathbb{E} \left[\sum_{t=1}^{\infty} \gamma^t R_t \left(\prod_{j=0}^t \frac{Z_j^\top A_j \pi(A_j | S_j)}{\Delta^*(S_j, A_j) P(Z_j | S_j)} \right) \mid S_0 = s \right].
\end{aligned}$$

Thus, we have

$$\begin{aligned}
& \mathbb{E} \left[\frac{Z_0^\top A_0 \pi(A_0 | S_0)}{\Delta^*(S_0, A_0) P(Z_0 | S_0)} (R_0 + \gamma V^\pi(S_1)) \mid S_0 = s \right] \\
&= \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^{t+1} R_{t+1} \left(\prod_{j=0}^{t+1} \frac{Z_j^\top A_j \pi(A_j | S_j)}{\Delta^*(S_j, A_j) P(Z_j | S_j)} \right) \mid S_0 = s \right] = V^\pi(s).
\end{aligned}$$

Similar results also hold for any $t \geq 0$. This concludes the proof of the lemma. \square

C.2.4. Proof of Corollary 4.2.5

Proof. We have

$$\begin{aligned}
& \mathbb{E} \left[f(S_t) \frac{Z_t^\top A_t \pi(A_t | S_t)}{\Delta^*(S_t, A_t) \Theta^*(S_t, Z_t)} (R_t + \gamma V^\pi(S_{t+1})) \right] \\
&= \mathbb{E} \left[f(S_t) \mathbb{E} \left[\frac{Z_t^\top A_t \pi(A_t | S_t)}{\Delta^*(S_t, A_t) \Theta^*(S_t, Z_t)} (R_t + \gamma V^\pi(S_{t+1})) \mid S_t \right] \right] \\
&= \mathbb{E} [f(S_t) V^\pi(S_t)],
\end{aligned}$$

where the last equality comes from Lemma 4.2.4. By summing all $t \in \{0, 1, \dots, T-1\}$, we conclude the proof of the corollary. \square

C.3. Proofs of Results in §4.4.1

C.3.1. Proof of Theorem 4.4.9

Proof. By the definition of $J(\pi)$ in (4.1.2), we proceed as follows,

$$\begin{aligned}
& J(\pi^*) - J(\widehat{\pi}_{\text{vf}}) \\
&= (1 - \gamma) \mathbb{E}_{S_0 \sim \nu} [V^{\pi^*}(S_0) - V^{\widehat{\pi}_{\text{vf}}}(S_0)] \\
&\leq (1 - \gamma) \mathbb{E}_{S_0 \sim \nu} [V^{\pi^*}(S_0)] - \min_{(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1} \min_{v \in \text{conf}_{\alpha_{\text{vf}}}^{\text{vf}}(\Delta, \Theta, \widehat{\pi})} (1 - \gamma) \mathbb{E}_{S_0 \sim \nu} [v(S_0)] \\
\text{(C.3.1)} \quad &\leq (1 - \gamma) \mathbb{E}_{S_0 \sim \nu} [V^{\pi^*}(S_0)] - \min_{(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1} \min_{v \in \text{conf}_{\alpha_{\text{vf}}}^{\text{vf}}(\Delta, \Theta, \pi^*)} (1 - \gamma) \mathbb{E}_{S_0 \sim \nu} [v(S_0)],
\end{aligned}$$

where in the first inequality, we use Lemma 4.4.7; while in the last inequality, we use the optimality of $\widehat{\pi}_{\text{vf}}$. It suffices to characterize the RHS of the above. We proceed (C.3.1) as follows,

$$\begin{aligned}
& J(\pi^*) - J(\widehat{\pi}_{\text{vf}}) \\
\text{(C.3.2)} \quad &\leq \max_{(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1} \max_{v \in \text{conf}_{\alpha_{\text{vf}}}^{\text{vf}}(\Delta, \Theta, \pi^*)} \left| (1 - \gamma) \mathbb{E}_{S_0 \sim \nu} [v(S_0)] - (1 - \gamma) \mathbb{E}_{S_0 \sim \nu} [V^{\pi^*}(S_0)] \right|.
\end{aligned}$$

Meanwhile, by Lemmas 4.2.7 and 4.2.8, we have

$$(1 - \gamma)\mathbb{E}_{S_0 \sim \nu} [V^{\pi^*}(S_0)] = J(\pi^*) = \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} w^{\pi^*}(S_t) \frac{Z_t^\top A_t \pi^*(A_t | S_t)}{\Delta^*(S_t, A_t) \Theta^*(S_t, Z_t)} R_t \right],$$

(C.3.3)

$$(1 - \gamma)\mathbb{E}_{S_0 \sim \nu} [v(S_0)] = \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} w^{\pi^*}(S_t) \frac{Z_t^\top A_t \pi^*(A_t | S_t)}{\Delta^*(S_t, A_t) \Theta^*(S_t, Z_t)} (v(S_t) - \gamma v(S_{t+1})) \right].$$

Now, by plugging (C.3.3) into the RHS of (C.3.2), we obtain

$$J(\pi^*) - J(\widehat{\pi}_{\text{vf}}) \leq \max_{(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1} \max_{v \in \text{conf}_{\alpha_{\text{vf}}}^{\text{vf}}(\Delta, \Theta, \pi^*)} |\Phi_{\text{vf}}^{\pi^*}(v, w^{\pi^*}; \Delta^*, \Theta^*)|.$$

By continuing the above inequality, we have

$$\begin{aligned} & J(\pi^*) - J(\widehat{\pi}_{\text{vf}}) \\ & \leq \max_{(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1} \max_{v \in \text{conf}_{\alpha_{\text{vf}}}^{\text{vf}}(\Delta, \Theta, \pi^*)} \max_{g \in \mathcal{W}} |\Phi_{\text{vf}}^{\pi^*}(v, g; \Delta^*, \Theta^*)| \\ & = \max_{(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1} \max_{v \in \text{conf}_{\alpha_{\text{vf}}}^{\text{vf}}(\Delta, \Theta, \pi^*)} \max_{g \in \mathcal{W}} \max \{ \Phi_{\text{vf}}^{\pi^*}(v, g; \Delta^*, \Theta^*), -\Phi_{\text{vf}}^{\pi^*}(v, g; \Delta^*, \Theta^*) \} \\ & = \max_{(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1} \max_{v \in \text{conf}_{\alpha_{\text{vf}}}^{\text{vf}}(\Delta, \Theta, \pi^*)} \max_{g \in \mathcal{W}} \max \{ \Phi_{\text{vf}}^{\pi^*}(v, g; \Delta^*, \Theta^*), \Phi_{\text{vf}}^{\pi^*}(v, -g; \Delta^*, \Theta^*) \} \\ & = \max_{(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1} \max_{v \in \text{conf}_{\alpha_{\text{vf}}}^{\text{vf}}(\Delta, \Theta, \pi^*)} \max_{g \in \mathcal{W}} \Phi_{\text{vf}}^{\pi^*}(v, g; \Delta^*, \Theta^*) \\ & \leq c \cdot \frac{C_{\Delta^*}^2 C_{\Theta^*}^2 C_*}{1 - \gamma} (\xi_0 + \xi_1) L_{\Pi} \sqrt{\frac{1}{NT\kappa} \cdot \mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \nu, \Pi} \cdot \log \frac{1}{\delta} \log(NT)}, \end{aligned}$$

where in the first inequality, we use Assumption 4.4.6; in the third equality, we use the fact that \mathcal{W} is symmetric; while in the last inequality, we use Lemma 4.4.8. This concludes the proof of the theorem. \square

C.3.2. Proof of Lemma 4.4.7

Proof. By Assumption 4.4.6, we know that $V^\pi \in \mathcal{V}$. Thus, to show that $V^\pi \in \text{conf}_{\alpha_{\text{vf}}}^{\text{vf}}(\Delta^*, \Theta^*, \pi)$ with a high probability, it suffices to show that

$$(C.3.4) \quad \max_{g \in \mathcal{W}} \widehat{\Phi}_{\text{vf}}^\pi(V^\pi, g; \Delta^*, \Theta^*) - \max_{g \in \mathcal{W}} \widehat{\Phi}_{\text{vf}}^\pi(\widehat{v}_{\Delta^*, \Theta^*}^\pi, g; \Delta^*, \Theta^*) \leq \alpha_{\text{vf}}.$$

In the follows, we show that (C.3.4) holds with a high probability. For the simplicity of notations, we denote by $\Phi_{\text{vf}}^\pi(v, g; *) = \Phi_{\text{vf}}^\pi(v, g; \Delta^*, \Theta^*)$ and $\widehat{v}_*^\pi = \widehat{v}_{\Delta^*, \Theta^*}^\pi$ for any (π, v, g) .

Note that

$$\begin{aligned} & \max_{g \in \mathcal{W}} \widehat{\Phi}_{\text{vf}}^\pi(V^\pi, g; *) - \max_{g \in \mathcal{W}} \widehat{\Phi}_{\text{vf}}^\pi(\widehat{v}_*^\pi, g; *) \\ &= \max_{g \in \mathcal{W}} \widehat{\Phi}_{\text{vf}}^\pi(V^\pi, g; *) - \max_{g \in \mathcal{W}} \Phi_{\text{vf}}^\pi(V^\pi, g; *) + \max_{g \in \mathcal{W}} \Phi_{\text{vf}}^\pi(V^\pi, g; *) - \max_{g \in \mathcal{W}} \Phi_{\text{vf}}^\pi(\widehat{v}_*^\pi, g; *) \\ & \quad + \max_{g \in \mathcal{W}} \Phi_{\text{vf}}^\pi(\widehat{v}_*^\pi, g; *) - \max_{g \in \mathcal{W}} \widehat{\Phi}_{\text{vf}}^\pi(\widehat{v}_*^\pi, g; *) \\ &\leq \max_{g \in \mathcal{W}} \widehat{\Phi}_{\text{vf}}^\pi(V^\pi, g; *) - \max_{g \in \mathcal{W}} \Phi_{\text{vf}}^\pi(V^\pi, g; *) + \max_{g \in \mathcal{W}} \Phi_{\text{vf}}^\pi(\widehat{v}_*^\pi, g; *) - \max_{g \in \mathcal{W}} \widehat{\Phi}_{\text{vf}}^\pi(\widehat{v}_*^\pi, g; *) \\ &\leq 2 \max_{v \in \mathcal{V}} \left| \max_{g \in \mathcal{W}} \widehat{\Phi}_{\text{vf}}^\pi(v, g; *) - \max_{g \in \mathcal{W}} \Phi_{\text{vf}}^\pi(v, g; *) \right| \\ (C.3.5) \quad &\leq 2 \max_{v \in \mathcal{V}} \max_{g \in \mathcal{W}} \left| \widehat{\Phi}_{\text{vf}}^\pi(v, g; *) - \Phi_{\text{vf}}^\pi(v, g; *) \right|, \end{aligned}$$

where in the first inequality, we use $\max_{g \in \mathcal{W}} \Phi_{\text{vf}}^\pi(V^\pi, g; *) = 0$ while $\max_{g \in \mathcal{W}} \Phi_{\text{vf}}^\pi(v, g; *) \geq 0$ for any v . In the meanwhile, by Theorem C.7.6, with probability at least $1 - \delta$, it holds

for any $(\pi, v, g) \in \Pi \times \mathcal{V} \times \mathcal{W}$ that

$$(C.3.6) \quad \left| \widehat{\Phi}_{\text{vf}}^\pi(v, g; *) - \Phi_{\text{vf}}^\pi(v, g; *) \right| \leq c \cdot \frac{C_{\Delta^*} C_{\Theta^*} C_*}{1 - \gamma} \sqrt{\frac{\mathfrak{C}_{\mathcal{W}, \mathcal{V}, \Pi}}{NT\kappa} \cdot \log \frac{1}{\delta} \log(NT)},$$

where we use Assumption (b) and $\|g\|_\infty \leq C_*$ for any $g \in \mathcal{W}$. Now, combining (C.3.5) and (C.3.6), with probability at least $1 - \delta$, we have

$$\max_{g \in \mathcal{W}} \widehat{\Phi}_{\text{vf}}^\pi(V^\pi, g; *) - \max_{g \in \mathcal{W}} \widehat{\Phi}_{\text{vf}}^\pi(\widehat{v}_*^\pi, g; *) \leq c \cdot \frac{C_{\Delta^*} C_{\Theta^*} C_*}{1 - \gamma} \sqrt{\frac{\mathfrak{C}_{\mathcal{W}, \mathcal{V}, \Pi}}{NT\kappa} \cdot \log \frac{1}{\delta} \log(NT)} = \alpha_{\text{vf}},$$

which implies that $V^\pi \in \text{conf}_{\alpha_{\text{vf}}}^{\text{vf}}(\Delta^*, \Theta^*, \pi)$ for any $\pi \in \Pi$. This concludes the proof of the lemma. \square

C.3.3. Proof of Lemma 4.4.8

Proof. Since $v \in \cup_{(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1} \text{conf}_{\alpha_{\text{vf}}}^{\text{vf}}(\Delta, \Theta, \pi)$, there exists a pair $(\widetilde{\Delta}, \widetilde{\Theta}) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1$ such that $v \in \text{conf}_{\alpha_{\text{vf}}}^{\text{vf}}(\widetilde{\Delta}, \widetilde{\Theta}, \pi)$. For the simplicity of notations, we denote by

$$(C.3.7) \quad \widetilde{v} \in \underset{v \in \mathcal{V}}{\text{argmin}} \max_{g \in \mathcal{W}} \widehat{\Phi}_{\text{vf}}^\pi(v, g; \widetilde{\Delta}, \widetilde{\Theta}),$$

i.e., $\widetilde{v} = \widehat{v}_{\widetilde{\Delta}, \widetilde{\Theta}}^\pi$, which is defined in (4.3.2). By the definition of \widetilde{v} and $v \in \text{conf}_{\alpha_{\text{vf}}}^{\text{vf}}(\widetilde{\Delta}, \widetilde{\Theta}, \pi)$, we know that

$$(C.3.8) \quad \max_{g \in \mathcal{W}} \widehat{\Phi}_{\text{vf}}^\pi(v, g; \widetilde{\Delta}, \widetilde{\Theta}) - \max_{g \in \mathcal{W}} \widehat{\Phi}_{\text{vf}}^\pi(\widetilde{v}, g; \widetilde{\Delta}, \widetilde{\Theta}) \leq \alpha_{\text{vf}}.$$

Note that

$$\begin{aligned}
& \max_{g \in \mathcal{W}} \Phi_{\text{vf}}^\pi(v, g; \Delta^*, \Theta^*) \\
&= \max_{g \in \mathcal{W}} \Phi_{\text{vf}}^\pi(v, g; \Delta^*, \Theta^*) - \max_{g \in \mathcal{W}} \widehat{\Phi}_{\text{vf}}^\pi(v, g; \Delta^*, \Theta^*) + \max_{g \in \mathcal{W}} \widehat{\Phi}_{\text{vf}}^\pi(v, g; \Delta^*, \Theta^*) \\
&\quad - \max_{g \in \mathcal{W}} \widehat{\Phi}_{\text{vf}}^\pi(v, g; \widetilde{\Delta}, \widetilde{\Theta}) + \max_{g \in \mathcal{W}} \widehat{\Phi}_{\text{vf}}^\pi(v, g; \widetilde{\Delta}, \widetilde{\Theta}) - \max_{g \in \mathcal{W}} \widehat{\Phi}_{\text{vf}}^\pi(\widetilde{v}, g; \widetilde{\Delta}, \widetilde{\Theta}) \\
&\quad + \max_{g \in \mathcal{W}} \widehat{\Phi}_{\text{vf}}^\pi(\widetilde{v}, g; \widetilde{\Delta}, \widetilde{\Theta}) - \max_{g \in \mathcal{W}} \Phi_{\text{vf}}^\pi(\widetilde{v}, g; \widetilde{\Delta}, \widetilde{\Theta}) + \max_{g \in \mathcal{W}} \Phi_{\text{vf}}^\pi(\widetilde{v}, g; \widetilde{\Delta}, \widetilde{\Theta}) \\
&\leq 2 \underbrace{\max_{(v, g, \Delta, \Theta) \in (\mathcal{V}, \mathcal{W}, \mathcal{F}_0, \mathcal{F}_1)} \left| \Phi_{\text{vf}}^\pi(v, g; \Delta, \Theta) - \widehat{\Phi}_{\text{vf}}^\pi(v, g; \Delta, \Theta) \right|}_{\text{Term (I)}} + \underbrace{\max_{g \in \mathcal{W}} \Phi_{\text{vf}}^\pi(\widetilde{v}, g; \widetilde{\Delta}, \widetilde{\Theta})}_{\text{Term (II)}} \\
&\quad + \underbrace{\max_{g \in \mathcal{W}} \left| \widehat{\Phi}_{\text{vf}}^\pi(v, g; \Delta^*, \Theta^*) - \widehat{\Phi}_{\text{vf}}^\pi(v, g; \widetilde{\Delta}, \widetilde{\Theta}) \right|}_{\text{Term (III)}} + \alpha_{\text{vf}},
\end{aligned}
\tag{C.3.9}$$

where we use (C.3.8) in the last inequality. Now we upper bound terms (I), (II), and (III) on the RHS of (C.3.9).

Upper Bounding Term (I). By Theorem C.7.6, with probability at least $1 - \delta$, it holds for any $(v, g, \Delta, \Theta, \pi) \in (\mathcal{V}, \mathcal{W}, \mathcal{F}_0, \mathcal{F}_1, \Pi)$ that

$$\left| \widehat{\Phi}_{\text{vf}}^\pi(v, g; \Delta, \Theta) - \Phi_{\text{vf}}^\pi(v, g; \Delta, \Theta) \right| \leq c \cdot \frac{C_{\Delta^*} C_{\Theta^*} C_*}{1 - \gamma} \sqrt{\frac{\mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}, \Pi}}{NT\kappa} \log \frac{1}{\delta} \log(NT)},$$

which implies that with probability at least $1 - \delta$, we have

$$\text{Term (I)} \leq c \cdot \frac{C_{\Delta^*} C_{\Theta^*} C_*}{1 - \gamma} \sqrt{\frac{\mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}, \Pi}}{NT\kappa} \log \frac{1}{\delta} \log(NT)}.
\tag{C.3.10}$$

Upper Bounding Term (II). We introduce the following lemma to help upper bound term (II).

Lemma C.3.1. Suppose α_0 and α_1 are defined in Assumption 4.4.4. With probability at least $1 - \delta$, for any $(\Delta, \Theta, \pi) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1 \times \Pi$, we have

$$\max_{g \in \mathcal{W}} \Phi_{\text{vf}}^\pi(\widehat{v}_{\Delta, \Theta}^\pi, g; \Delta, \Theta) \leq c \cdot \frac{C_{\Delta^*}^2 C_{\Theta^*}^2 C_*}{1 - \gamma} (\xi_0 + \xi_1) L_\Pi \sqrt{\frac{\mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}, \Pi}}{NT\kappa} \log \frac{1}{\delta} \log(NT)},$$

where $\widehat{v}_{\Delta, \Theta}^\pi$ is defined in (4.3.2), ξ_0 and ξ_1 are constants defined in Assumption 4.4.4.

Proof. See §C.3.4 for a detailed proof. □

By the definition of \tilde{v} in (C.3.7) and Lemma C.3.1, with probability at least $1 - \delta$, we have

$$(C.3.11) \quad \text{Term (II)} \leq c \cdot \frac{C_{\Delta^*}^2 C_{\Theta^*}^2 C_*}{1 - \gamma} (\xi_0 + \xi_1) L_\Pi \sqrt{\frac{\mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}, \Pi}}{NT\kappa} \log \frac{1}{\delta} \log(NT)}.$$

Upper Bounding Term (III). Note that

$$(C.3.12) \quad \begin{aligned} & \left| \widehat{\Phi}_{\text{vf}}^\pi(v, g; \Delta^*, \Theta^*) - \widehat{\Phi}_{\text{vf}}^\pi(v, g; \tilde{\Delta}, \tilde{\Theta}) \right| \\ & \leq \left| \left(\widehat{\mathbb{E}} - \mathbb{E} \right) \left[\frac{1}{T} \sum_{t=0}^{T-1} g(S_t) \left(\frac{Z_t^\top A_t \pi(A_t | S_t)}{\Delta^*(S_t, A_t) \Theta^*(S_t, Z_t)} - \frac{Z_t^\top A_t \pi(A_t | S_t)}{\tilde{\Delta}(S_t, A_t) \tilde{\Theta}(S_t, Z_t)} \right) (R_t + \gamma v(S_{t+1})) \right] \right| \\ & \quad + \left| \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} g(S_t) \left(\frac{Z_t^\top A_t \pi(A_t | S_t)}{\Delta^*(S_t, A_t) \Theta^*(S_t, Z_t)} - \frac{Z_t^\top A_t \pi(A_t | S_t)}{\tilde{\Delta}(S_t, A_t) \tilde{\Theta}(S_t, Z_t)} \right) (R_t + \gamma v(S_{t+1})) \right] \right|. \end{aligned}$$

For the first term on the RHS of (C.3.12), by Theorem C.7.6, with probability at least $1 - \delta$, it holds for any $(v, g, \pi) \in \mathcal{V} \times \mathcal{W} \times \Pi$ that

$$\begin{aligned} & \left| \left(\widehat{\mathbb{E}} - \mathbb{E} \right) \left[\frac{1}{T} \sum_{t=0}^{T-1} g(S_t) \left(\frac{Z_t^\top A_t \pi(A_t | S_t)}{\Delta^*(S_t, A_t) \Theta^*(S_t, Z_t)} - \frac{Z_t^\top A_t \pi(A_t | S_t)}{\widetilde{\Delta}(S_t, A_t) \widetilde{\Theta}(S_t, Z_t)} \right) (R_t + \gamma v(S_{t+1})) \right] \right| \\ (C.3.13) \quad & \leq c \cdot \frac{C_{\Delta^*} C_{\Theta^*} C_*}{1 - \gamma} \sqrt{\frac{\mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}, \Pi}}{NT\kappa} \log \frac{1}{\delta} \log(NT)}. \end{aligned}$$

For the second term on the RHS of (C.3.12), with probability at least $1 - \delta$, it holds that

$$\begin{aligned}
& \left| \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} g(S_t) \left(\frac{Z_t^\top A_t \pi(A_t | S_t)}{\Delta^*(S_t, A_t) \Theta^*(S_t, Z_t)} - \frac{Z_t^\top A_t \pi(A_t | S_t)}{\tilde{\Delta}(S_t, A_t) \tilde{\Theta}(S_t, Z_t)} \right) (R_t + \gamma v(S_{t+1})) \right] \right| \\
& \leq \frac{C_*}{1 - \gamma} \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} \left| \frac{1}{\tilde{\Delta}(S_t, A_t) \tilde{\Theta}(S_t, Z_t)} - \frac{1}{\Delta^*(S_t, A_t) \Theta^*(S_t, Z_t)} \right| \right] \\
& = \frac{C_*}{1 - \gamma} \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} \left| \frac{\Theta^*(S_t, Z_t) - \tilde{\Theta}(S_t, Z_t)}{\tilde{\Delta}(S_t, A_t) \tilde{\Theta}(S_t, Z_t) \Theta^*(S_t, Z_t)} - \frac{\Delta^*(S_t, A_t) - \tilde{\Delta}(S_t, A_t)}{\Delta^*(S_t, A_t) \tilde{\Delta}(S_t, A_t) \Theta^*(S_t, Z_t)} \right| \right] \\
& \leq \frac{C_{\Delta^*} C_{\Theta^*} C_*}{1 - \gamma} \left(C_{\Theta^*} \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} \left\| \Delta^*(S_t, \cdot) - \tilde{\Delta}(S_t, \cdot) \right\|_1 \right] \right. \\
& \quad \left. + C_{\Delta^*} \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} \left\| \Theta^*(S_t, \cdot) - \tilde{\Theta}(S_t, \cdot) \right\|_1 \right] \right) \\
& \leq \frac{C_{\Delta^*} C_{\Theta^*} C_*}{1 - \gamma} \left(C_{\Theta^*} \sqrt{\mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} \left\| \Delta^*(S_t, \cdot) - \tilde{\Delta}(S_t, \cdot) \right\|_1^2 \right]} \right. \\
& \quad \left. + C_{\Delta^*} \sqrt{\mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} \left\| \Theta^*(S_t, \cdot) - \tilde{\Theta}(S_t, \cdot) \right\|_1^2 \right]} \right)
\end{aligned}$$

(C.3.14)

$$\leq \frac{C_{\Delta^*} C_{\Theta^*} C_*}{1 - \gamma} \left(\xi_0 C_{\Theta^*} \sqrt{\frac{C_{\Delta^*}}{NT\kappa} \mathfrak{C}_{\mathcal{F}_0} \log \frac{1}{\delta} \log(NT)} + \xi_1 C_{\Delta^*} \sqrt{\frac{C_{\Theta^*}}{NT\kappa} \mathfrak{C}_{\mathcal{F}_1} \log \frac{1}{\delta} \log(NT)} \right),$$

where in the first inequality, we use the fact that $\|v\|_\infty \leq 1/(1 - \gamma)$ and $\|g\|_\infty \leq C_*$; in the third inequality, we use Cauchy-Schwarz inequality; while in the last inequality, we use Assumption 4.4.4 with the fact that $(\tilde{\Delta}, \tilde{\Theta}) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1$. Now, by plugging (C.3.13) and (C.3.14) into (C.3.12), with probability at least $1 - \delta$, it holds for any

$v \in \cup_{(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1} \text{conf}_{\alpha_{\text{vf}}}^{\text{vf}}(\Delta, \Theta, \pi)$, $g \in \mathcal{W}$, and $\pi \in \Pi$ that

$$(C.3.15) \quad \begin{aligned} & \left| \widehat{\Phi}_{\text{vf}}^{\pi}(v, g; \Delta^*, \Theta^*) - \widehat{\Phi}_{\text{vf}}^{\pi}(v, g; \widetilde{\Delta}, \widetilde{\Theta}) \right| \\ & \leq c \cdot \frac{C_{\Delta^*}^2 C_{\Theta^*}^2 C_*}{1 - \gamma} (\xi_0 + \xi_1) \sqrt{\frac{1}{NT\kappa} \cdot \mathfrak{e}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}, \Pi} \cdot \log \frac{1}{\delta} \log(NT)}. \end{aligned}$$

Now, by plugging (C.3.10), (C.3.11), and (C.3.15) into (C.3.9), with probability at least $1 - \delta$, it holds for any $v \in \cup_{(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1} \text{conf}_{\alpha_{\text{vf}}}^{\text{vf}}(\Delta, \Theta, \pi)$ and $\pi \in \Pi$ that

$$\max_{g \in \mathcal{W}} \Phi_{\text{vf}}^{\pi}(v, g; \Delta^*, \Theta^*) \leq c \cdot \frac{C_{\Delta^*}^2 C_{\Theta^*}^2 C_*}{1 - \gamma} (\xi_0 + \xi_1) L_{\Pi} \sqrt{\frac{1}{NT\kappa} \cdot \mathfrak{e}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}, \Pi} \cdot \log \frac{1}{\delta} \log(NT)},$$

which concludes the proof of the lemma. \square

C.3.4. Proof of Lemma C.3.1

Proof. Note that

$$(C.3.16) \quad \begin{aligned} & \max_{g \in \mathcal{W}} \Phi_{\text{vf}}^{\pi}(\widehat{v}_{\Delta, \Theta}^{\pi}, g; \Delta, \Theta) \\ & = \max_{g \in \mathcal{W}} \Phi_{\text{vf}}^{\pi}(\widehat{v}_{\Delta, \Theta}^{\pi}, g; \Delta, \Theta) - \max_{g \in \mathcal{W}} \widehat{\Phi}_{\text{vf}}^{\pi}(\widehat{v}_{\Delta, \Theta}^{\pi}, g; \Delta, \Theta) + \max_{g \in \mathcal{W}} \widehat{\Phi}_{\text{vf}}^{\pi}(\widehat{v}_{\Delta, \Theta}^{\pi}, g; \Delta, \Theta) \\ & \quad - \max_{g \in \mathcal{W}} \widehat{\Phi}_{\text{vf}}^{\pi}(V^{\pi}, g; \Delta, \Theta) + \max_{g \in \mathcal{W}} \widehat{\Phi}_{\text{vf}}^{\pi}(V^{\pi}, g; \Delta, \Theta) - \max_{g \in \mathcal{W}} \Phi_{\text{vf}}^{\pi}(V^{\pi}, g; \Delta, \Theta) \\ & \quad + \max_{g \in \mathcal{W}} \Phi_{\text{vf}}^{\pi}(V^{\pi}, g; \Delta, \Theta) - \max_{g \in \mathcal{W}} \Phi_{\text{vf}}^{\pi}(V^{\pi}, g; \Delta^*, \Theta^*) \\ & \leq 2 \max_{v \in \mathcal{V}} \max_{g \in \mathcal{W}} \left| \Phi_{\text{vf}}^{\pi}(v, g; \Delta, \Theta) - \widehat{\Phi}_{\text{vf}}^{\pi}(v, g; \Delta, \Theta) \right| \\ & \quad + \max_{g \in \mathcal{W}} \left| \Phi_{\text{vf}}^{\pi}(V^{\pi}, g; \Delta, \Theta) - \Phi_{\text{vf}}^{\pi}(V^{\pi}, g; \Delta^*, \Theta^*) \right|, \end{aligned}$$

where we use the fact that $\widehat{v}_{\Delta, \Theta}^\pi \in \operatorname{argmin}_{v \in \mathcal{V}} \max_{g \in \mathcal{W}} \widehat{\Phi}_{\text{vf}}^\pi(v, g; \Delta, \Theta)$ in the last inequality. In the meanwhile, by Theorem C.7.6, with probability at least $1 - \delta$, it holds for any $(v, g, \pi) \in \mathcal{V} \times \mathcal{W} \times \Pi$ that

$$(C.3.17) \quad \left| \widehat{\Phi}_{\text{vf}}^\pi(v, g; \Delta, \Theta) - \Phi_{\text{vf}}^\pi(v, g; \Delta, \Theta) \right| \leq c \cdot \frac{C_{\Delta^*} C_{\Theta^*} C_*}{1 - \gamma} \sqrt{\frac{\mathfrak{C}_{\mathcal{V}, \mathcal{W}, \Pi}}{NT\kappa} \log \frac{1}{\delta} \log(NT)}.$$

Also, we upper bound the second term on the RHS of (C.3.16) with probability at least $1 - \delta$ by a similar argument as in (C.3.14),

$$(C.3.18) \quad \begin{aligned} & \left| \Phi_{\text{vf}}^\pi(V^\pi, g; \Delta, \Theta) - \Phi_{\text{vf}}^\pi(V^\pi, g; \Delta^*, \Theta^*) \right| \\ & \leq \frac{C_{\Delta^*} C_{\Theta^*} C_*}{1 - \gamma} \left(\xi_0 C_{\Theta^*} \sqrt{\frac{C_{\Delta^*}}{NT\kappa} \mathfrak{C}_{\mathcal{F}_0} \log \frac{1}{\delta} \log(NT)} + \xi_1 C_{\Delta^*} \sqrt{\frac{C_{\Theta^*}}{NT\kappa} \mathfrak{C}_{\mathcal{F}_1} \log \frac{1}{\delta} \log(NT)} \right), \end{aligned}$$

where in the first inequality, we use the fact that $\|V^\pi\|_\infty \leq 1/(1 - \gamma)$ and $\|g\|_\infty \leq C_*$; in the third inequality, we use Cauchy Schwarz inequality; while in the last inequality, we use Assumption 4.4.4 with $(\Delta, \Theta) \in \operatorname{conf}_{\alpha_0}^0 \times \operatorname{conf}_{\alpha_1}^1$. Now, by plugging (C.3.17) and (C.3.18) into (C.3.16), with probability at least $1 - \delta$, it holds for any $(\Delta, \Theta, \pi) \in \operatorname{conf}_{\alpha_0}^0 \times \operatorname{conf}_{\alpha_1}^1 \times \Pi$ that

$$\max_{g \in \mathcal{W}} \Phi_{\text{vf}}^\pi(\widehat{v}_{\Delta, \Theta}^\pi, g; \Delta, \Theta) \leq c \cdot \frac{C_{\Delta^*}^2 C_{\Theta^*}^2 C_*}{1 - \gamma} (\xi_0 + \xi_1) L_\Pi \sqrt{\frac{\mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}, \Pi}}{NT\kappa} \log \frac{1}{\delta} \log(NT)},$$

which concludes the proof of the lemma. \square

C.4. Proofs of Results in §4.4.2

C.4.1. Proof of Theorem 4.4.13

Proof. Before the proof of the theorem, we first introduce some supporting results as follows. We define the population counterpart of $\widehat{L}_{\text{mis}}(w, \pi; \Delta, \Theta)$ as

$$L_{\text{mis}}(w, \pi; \Delta, \Theta) = \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} \frac{Z_t^\top A_t \pi(A_t | S_t)}{\Delta(S_t, A_t) \Theta(S_t, Z_t)} w(S_t) R_t \right]$$

for any (w, π, Δ, Θ) .

Lemma C.4.1. It holds for any $(\pi, w) \in \Pi \times \mathcal{W}$ that

$$L_{\text{mis}}(w^\pi, \pi; \Delta^*, \Theta^*) - L_{\text{mis}}(w, \pi; \Delta^*, \Theta^*) = \Phi_{\text{mis}}^\pi(w, V^\pi; \Delta^*, \Theta^*),$$

where V^π is the state-value function defined in (4.1.2).

Proof. See §C.4.4 for a detailed proof. □

Lemma C.4.2. Suppose that $(\alpha_0, \alpha_1, \alpha_{\text{mis}})$ is defined in Lemmas 4.4.11. With probability at least $1 - \delta$, it holds for any $(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1$ that

$$\begin{aligned} & \left| \min_{w \in \text{conf}_{\alpha_{\text{mis}}}^{\text{mis}}(\Delta^*, \Theta^*, \pi^*)} L_{\text{mis}}(w, \pi^*; \Delta^*, \Theta^*) - \min_{w \in \text{conf}_{\alpha_{\text{mis}}}^{\text{mis}}(\Delta, \Theta, \pi^*)} L_{\text{mis}}(w, \pi^*; \Delta, \Theta) \right| \\ & \leq c \cdot \frac{C_{\Delta^*}^2 C_{\Theta^*}^2 C_*}{1 - \gamma} (\xi_0 + \xi_1) \sqrt{\frac{1}{NT\kappa} \mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}} \log \frac{1}{\delta} \log(NT)} = \varepsilon_L^*. \end{aligned}$$

Proof. See §C.4.5 for a detailed proof. □

Lemma C.4.3. With probability at least $1 - \delta$, it holds for any $(w, \Delta, \Theta, \pi) \in \mathcal{W} \times \mathcal{F}_0 \times \mathcal{F}_1 \times \Pi$ that

$$\begin{aligned} & \left| L_{\text{mis}}(w, \pi; \Delta, \Theta) - \widehat{L}_{\text{mis}}(w, \pi; \Delta, \Theta) \right| \\ & \leq c \cdot C_{\Delta^*} C_{\Theta^*} C_* \sqrt{\frac{1}{NT\kappa} \mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \Pi} \log \frac{1}{\delta} \log(NT)} = \widehat{\varepsilon}_L. \end{aligned}$$

Proof. See §C.4.6 for a detailed proof. \square

Now we start the proof of the theorem. By the definition of $J(\pi)$, it holds with probability at least $1 - \delta$ that

$$\begin{aligned} J(\pi^*) - J(\widehat{\pi}_{\text{mis}}) &= L_{\text{mis}}(w^{\pi^*}, \pi^*; \Delta^*, \Theta^*) - L_{\text{mis}}(w^{\widehat{\pi}_{\text{mis}}}, \widehat{\pi}_{\text{mis}}; \Delta^*, \Theta^*) \\ &\leq L_{\text{mis}}(w^{\pi^*}, \pi^*; \Delta^*, \Theta^*) - \min_{w \in \text{conf}_{\alpha_{\text{mis}}}^{\text{mis}}(\Delta^*, \Theta^*, \widehat{\pi}_{\text{mis}})} L_{\text{mis}}(w, \widehat{\pi}_{\text{mis}}; \Delta^*, \Theta^*) \\ &\leq L_{\text{mis}}(w^{\pi^*}, \pi^*; \Delta^*, \Theta^*) \\ &\quad - \min_{(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1} \min_{w \in \text{conf}_{\alpha_{\text{mis}}}^{\text{mis}}(\Delta, \Theta, \widehat{\pi}_{\text{mis}})} L_{\text{mis}}(w, \widehat{\pi}_{\text{mis}}; \Delta, \Theta) \\ &\leq L_{\text{mis}}(w^{\pi^*}, \pi^*; \Delta^*, \Theta^*) \\ &\quad - \min_{(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1} \min_{w \in \text{conf}_{\alpha_{\text{mis}}}^{\text{mis}}(\Delta, \Theta, \widehat{\pi}_{\text{mis}})} \widehat{L}_{\text{mis}}(w, \widehat{\pi}_{\text{mis}}; \Delta, \Theta) + \widehat{\varepsilon}_L \\ \text{(C.4.1)} \quad &\leq L_{\text{mis}}(w^{\pi^*}, \pi^*; \Delta^*, \Theta^*) \\ &\quad - \min_{(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1} \min_{w \in \text{conf}_{\alpha_{\text{mis}}}^{\text{mis}}(\Delta, \Theta, \pi^*)} \widehat{L}_{\text{mis}}(w, \pi^*; \Delta, \Theta) + \widehat{\varepsilon}_L, \end{aligned}$$

where we use Lemma 4.4.11 in the first inequality; we use Assumption 4.4.4 that $(\Delta^*, \Theta^*) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1$ with probability at least $1 - \delta$ in the second inequality; we use Lemma

C.4.3 in the third inequality; while we use the optimality of $\widehat{\pi}_{\text{mis}}$ in the last inequality. Now, by applying Lemmas C.4.2 and C.4.3, we obtain from (C.4.1) that

$$\begin{aligned}
J(\pi^*) - J(\widehat{\pi}_{\text{mis}}) &\leq L_{\text{mis}}(w^{\pi^*}, \pi^*; \Delta^*, \Theta^*) - \min_{w \in \text{conf}_{\alpha_{\text{mis}}}^{\text{mis}}(\Delta^*, \Theta^*, \pi^*)} L_{\text{mis}}(w, \pi^*; \Delta^*, \Theta^*) + \varepsilon_L^* + 2\widehat{\varepsilon}_L \\
&\leq \max_{w \in \text{conf}_{\alpha_{\text{mis}}}^{\text{mis}}(\Delta^*, \Theta^*, \pi^*)} \left| \Phi_{\text{mis}}^{\pi^*}(w, V^{\pi^*}; \Delta^*, \Theta^*) \right| + \varepsilon_L^* + 2\widehat{\varepsilon}_L \\
&\leq \max_{w \in \text{conf}_{\alpha_{\text{mis}}}^{\text{mis}}(\Delta^*, \Theta^*, \pi^*)} \max_{f \in \mathcal{V}} \left| \Phi_{\text{mis}}^{\pi^*}(w, f; \Delta^*, \Theta^*) \right| + \varepsilon_L^* + 2\widehat{\varepsilon}_L \\
&\leq \max_{w \in \text{conf}_{\alpha_{\text{mis}}}^{\text{mis}}(\Delta^*, \Theta^*, \pi^*)} \max_{f \in \mathcal{V}} \max \left\{ \Phi_{\text{mis}}^{\pi^*}(w, f; \Delta^*, \Theta^*), -\Phi_{\text{mis}}^{\pi^*}(w, f; \Delta^*, \Theta^*) \right\} \\
&\quad + \varepsilon_L^* + 2\widehat{\varepsilon}_L \\
&\leq \max_{w \in \text{conf}_{\alpha_{\text{mis}}}^{\text{mis}}(\Delta^*, \Theta^*, \pi^*)} \max_{f \in \mathcal{V}} \max \left\{ \Phi_{\text{mis}}^{\pi^*}(w, f; \Delta^*, \Theta^*), \Phi_{\text{mis}}^{\pi^*}(w, -f; \Delta^*, \Theta^*) \right\} \\
&\quad + \varepsilon_L^* + 2\widehat{\varepsilon}_L \\
\text{(C.4.2)} \quad &\leq \max_{w \in \text{conf}_{\alpha_{\text{mis}}}^{\text{mis}}(\Delta^*, \Theta^*, \pi^*)} \max_{f \in \mathcal{V}} \Phi_{\text{mis}}^{\pi^*}(w, f; \Delta^*, \Theta^*) + \varepsilon_L^* + 2\widehat{\varepsilon}_L,
\end{aligned}$$

where in the second inequality, we use Lemma C.4.1; in the third inequality, we use Assumption 4.4.10; while in the last inequality, we use the fact that \mathcal{V} is symmetric. Now, by Lemma 4.4.12 and plugging the definition of $\widehat{\varepsilon}_L$ and ε_L^* into (C.4.2), it holds with probability at least $1 - \delta$ that

$$J(\pi^*) - J(\widehat{\pi}_{\text{mis}}) \leq c \cdot \frac{C_{\Delta^*}^2 C_{\Theta^*}^2 C_*}{1 - \gamma} (\xi_0 + \xi_1) \sqrt{\frac{1}{NT\kappa}} \cdot \mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}, \Pi} \cdot \log \frac{NT}{\delta},$$

which concludes the proof of the theorem. \square

C.4.2. Proof of Lemma 4.4.11

Proof. First, by Assumption 4.4.10, we know that $w^\pi \in \mathcal{W}$. For notation simplicity, we denote by $\Phi_{\text{mis}}^\pi(w, f; *) = \Phi_{\text{mis}}^\pi(w, f; \Delta^*, \Theta^*)$ and $\hat{w}_*^\pi = \hat{w}_{\Delta^*, \Theta^*}^\pi$ for any (π, w, f) . Note that

$$\begin{aligned}
& \max_{f \in \mathcal{V}} \hat{\Phi}_{\text{mis}}^\pi(w^\pi, f; *) - \max_{f \in \mathcal{V}} \hat{\Phi}_{\text{mis}}^\pi(\hat{w}_*^\pi, f; *) \\
&= \max_{f \in \mathcal{V}} \hat{\Phi}_{\text{mis}}^\pi(w^\pi, f; *) - \max_{f \in \mathcal{V}} \Phi_{\text{mis}}^\pi(w^\pi, f; *) + \max_{f \in \mathcal{V}} \Phi_{\text{mis}}^\pi(w^\pi, f; *) - \max_{f \in \mathcal{V}} \Phi_{\text{mis}}^\pi(\hat{w}_*^\pi, f; *) \\
&\quad + \max_{f \in \mathcal{V}} \Phi_{\text{mis}}^\pi(\hat{w}_*^\pi, f; *) - \max_{f \in \mathcal{V}} \hat{\Phi}_{\text{mis}}^\pi(\hat{w}_*^\pi, f; *) \\
&\leq \max_{f \in \mathcal{V}} \hat{\Phi}_{\text{mis}}^\pi(w^\pi, f; *) - \max_{f \in \mathcal{V}} \Phi_{\text{mis}}^\pi(w^\pi, f; *) + \max_{f \in \mathcal{V}} \Phi_{\text{mis}}^\pi(\hat{w}_*^\pi, f; *) - \max_{f \in \mathcal{V}} \hat{\Phi}_{\text{mis}}^\pi(\hat{w}_*^\pi, f; *) \\
&\leq 2 \max_{w \in \mathcal{W}} \left| \max_{f \in \mathcal{V}} \hat{\Phi}_{\text{mis}}^\pi(w, f; *) - \max_{f \in \mathcal{V}} \Phi_{\text{mis}}^\pi(w, f; *) \right|
\end{aligned}$$

(C.4.3)

$$\leq 2 \max_{w \in \mathcal{W}} \max_{f \in \mathcal{V}} \left| \hat{\Phi}_{\text{mis}}^\pi(w, f; *) - \Phi_{\text{mis}}^\pi(w, f; *) \right|,$$

where in the first inequality, we use the fact that $w^\pi = \operatorname{argmin}_{w \in \mathcal{W}} \max_{f \in \mathcal{V}} \Phi_{\text{mis}}^\pi(w, f; *)$; while in the second inequality, we use $w^\pi \in \mathcal{W}$ by Assumption 4.4.10. In the meanwhile, by Theorem C.7.6, with probability at least $1 - \delta$, it holds for any $(w, f, \pi) \in \mathcal{W} \times \mathcal{V} \times \Pi$ that

$$\text{(C.4.4)} \quad \left| \hat{\Phi}_{\text{mis}}^\pi(w, f; *) - \Phi_{\text{mis}}^\pi(w, f; *) \right| \leq c \cdot \frac{C_{\Delta^*} C_{\Theta^*} C_*}{1 - \gamma} \sqrt{\frac{1}{NT\kappa} \mathfrak{C}_{\mathcal{V}, \mathcal{W}, \Pi} \log \frac{1}{\delta} \log(NT)},$$

where we use Assumption (b). Now, combining (C.4.3) and (C.4.4), with probability at least $1 - \delta$, we have

$$\begin{aligned} & \max_{f \in \mathcal{V}} \widehat{\Phi}_{\text{mis}}^\pi(w^\pi, f; *) - \max_{f \in \mathcal{V}} \widehat{\Phi}_{\text{mis}}^\pi(\widehat{w}_*^\pi, f; *) \\ & \leq c \cdot \frac{C_{\Delta^*} C_{\Theta^*} C_*}{1 - \gamma} \sqrt{\frac{1}{NT\kappa} \mathfrak{C}_{\mathcal{V}, \mathcal{W}, \Pi} \log \frac{1}{\delta} \log(NT)} = \alpha_{\text{mis}}, \end{aligned}$$

which implies that $w^\pi \in \text{conf}_{\alpha_{\text{mis}}}^{\text{mis}}(\Delta^*, \Theta^*, \pi)$. This concludes the proof of the lemma. \square

C.4.3. Proof of Lemma 4.4.12

Proof. Since $w \in \cup_{(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1} \text{conf}_{\alpha_{\text{mis}}}^{\text{mis}}(\Delta, \Theta, \pi)$, there exists a pair $(\widetilde{\Delta}, \widetilde{\Theta}) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1$ such that $w \in \text{conf}_{\alpha_{\text{mis}}}^{\text{mis}}(\widetilde{\Delta}, \widetilde{\Theta}, \pi)$. For the simplicity of notations, we denote by

$$(C.4.5) \quad \widetilde{w} \in \underset{w \in \mathcal{W}}{\text{argmin}} \max_{f \in \mathcal{V}} \widehat{\Phi}_{\text{mis}}^\pi(w, f; \widetilde{\Delta}, \widetilde{\Theta}),$$

i.e., $\widetilde{w} = \widehat{w}_{\widetilde{\Delta}, \widetilde{\Theta}}^\pi$, which is defined in (4.3.5). By the definition of \widetilde{w} and $w \in \text{conf}_{\alpha_{\text{mis}}}^{\text{mis}}(\widetilde{\Delta}, \widetilde{\Theta}, \pi)$, with probability at least $1 - \delta$, it holds for any $\pi \in \Pi$ and $w \in \text{conf}_{\alpha_{\text{mis}}}^{\text{mis}}(\widetilde{\Delta}, \widetilde{\Theta}, \pi)$ that

$$(C.4.6) \quad \max_{f \in \mathcal{V}} \widehat{\Phi}_{\text{mis}}^\pi(w, f; \widetilde{\Delta}, \widetilde{\Theta}) - \max_{f \in \mathcal{V}} \widehat{\Phi}_{\text{mis}}^\pi(\widetilde{w}, f; \widetilde{\Delta}, \widetilde{\Theta}) \leq \alpha_{\text{mis}}.$$

Further, we observe that

$$\begin{aligned}
& \max_{f \in \mathcal{V}} \Phi_{\text{mis}}^\pi(w, f; \Delta^*, \Theta^*) \\
\text{(C.4.7)} \quad & \leq \underbrace{\max_{(w, f, \Delta, \Theta) \in (\mathcal{W}, \mathcal{V}, \mathcal{F}_0, \mathcal{F}_1)} \left| \Phi_{\text{mis}}^\pi(w, f; \Delta, \Theta) - \widehat{\Phi}_{\text{mis}}^\pi(w, f; \Delta, \Theta) \right|}_{\text{Term (I)}} + \underbrace{\max_{f \in \mathcal{V}} \Phi_{\text{mis}}^\pi(\tilde{w}, f; \tilde{\Delta}, \tilde{\Theta})}_{\text{Term (II)}} \\
& \quad + \underbrace{\max_{f \in \mathcal{V}} \left| \widehat{\Phi}_{\text{mis}}^\pi(w, f; \Delta^*, \Theta^*) - \widehat{\Phi}_{\text{mis}}^\pi(w, f; \tilde{\Delta}, \tilde{\Theta}) \right|}_{\text{Term (III)}} + \alpha_{\text{mis}},
\end{aligned}$$

where we use (C.4.6) in the last inequality. Now we upper bound terms (I), (II), and (III) on the RHS of (C.4.7).

Upper Bounding Term (I). By Theorem C.7.6, with probability at least $1 - \delta$, it holds for any $(w, f, \Delta, \Theta, \pi) \in (\mathcal{W}, \mathcal{V}, \mathcal{F}_0, \mathcal{F}_1, \Pi)$ that

$$\left| \widehat{\Phi}_{\text{mis}}^\pi(w, f; \Delta, \Theta) - \Phi_{\text{mis}}^\pi(w, f; \Delta, \Theta) \right| \leq c \cdot \frac{C_{\Delta^*} C_{\Theta^*} C_*}{1 - \gamma} \sqrt{\frac{1}{NT\kappa} \mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}, \Pi} \log \frac{1}{\delta} \log(NT)},$$

which implies that with probability at least $1 - \delta$, we have

$$\text{(C.4.8)} \quad \text{Term (I)} \leq c \cdot \frac{C_{\Delta^*} C_{\Theta^*} C_*}{1 - \gamma} \sqrt{\frac{1}{NT\kappa} \mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}, \Pi} \log \frac{1}{\delta} \log(NT)}.$$

Upper Bounding Term (II). We introduce the following lemma to help upper bound term (II).

Lemma C.4.4. Suppose (α_0, α_1) is defined in Assumption 4.4.4. With probability at least $1 - \delta$, for any $(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1$ and $\pi \in \Pi$, we have

$$\max_{f \in \mathcal{V}} \Phi_{\text{mis}}^\pi(\widehat{w}_{\Delta, \Theta}^\pi, f; \Delta, \Theta) \leq c \cdot \frac{C_{\Delta^*}^2 C_{\Theta^*}^2 C_*}{1 - \gamma} (\xi_0 + \xi_1) \sqrt{\frac{1}{NT\kappa} \mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \nu, \Pi} \log \frac{1}{\delta} \log(NT)},$$

where $\widehat{w}_{\Delta, \Theta}^\pi$ is defined in (4.3.5), ξ_0 and ξ_1 are constants defined in Assumption 4.4.4.

Proof. See §C.4.7 for a detailed proof. □

By the definition of \widetilde{w} in (C.4.5) and Lemma C.4.4, with probability at least $1 - \delta$, we have

$$(C.4.9) \quad \text{Term (II)} \leq c \cdot \frac{C_{\Delta^*}^2 C_{\Theta^*}^2 C_*}{1 - \gamma} (\xi_0 + \xi_1) \sqrt{\frac{1}{NT\kappa} \mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \nu, \Pi} \log \frac{1}{\delta} \log(NT)}.$$

Upper Bounding Term (III). Note that

$$(C.4.10) \quad \begin{aligned} & \left| \widehat{\Phi}_{\text{mis}}^\pi(w, f; \Delta^*, \Theta^*) - \widehat{\Phi}_{\text{mis}}^\pi(w, f; \widetilde{\Delta}, \widetilde{\Theta}) \right| \\ & \leq \left| \left(\widehat{\mathbb{E}} - \mathbb{E} \right) \left[\frac{1}{T} \sum_{t=0}^{T-1} \left(\frac{Z_t^\top A_t \pi(A_t | S_t) w(S_t)}{\Delta^*(S_t, A_t) \Theta^*(S_t, Z_t)} - \frac{Z_t^\top A_t \pi(A_t | S_t) w(S_t)}{\widetilde{\Delta}(S_t, A_t) \widetilde{\Theta}(S_t, Z_t)} \right) (f(S_t) - \gamma f(S_{t+1})) \right] \right| \\ & \quad + \left| \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} \left(\frac{Z_t^\top A_t \pi(A_t | S_t) w(S_t)}{\Delta^*(S_t, A_t) \Theta^*(S_t, Z_t)} - \frac{Z_t^\top A_t \pi(A_t | S_t) w(S_t)}{\widetilde{\Delta}(S_t, A_t) \widetilde{\Theta}(S_t, Z_t)} \right) (f(S_t) - \gamma f(S_{t+1})) \right] \right|. \end{aligned}$$

For the first term on the RHS of (C.4.10), by Theorem C.7.6, with probability at least $1 - \delta$, it holds for any $(w, f, \pi) \in \mathcal{W} \times \mathcal{V} \times \Pi$ that

$$(C.4.11) \quad \left| \left(\widehat{\mathbb{E}} - \mathbb{E} \right) \left[\frac{1}{T} \sum_{t=0}^{T-1} \left(\frac{Z_t^\top A_t \pi(A_t | S_t) w(S_t)}{\Delta^*(S_t, A_t) \Theta^*(S_t, Z_t)} - \frac{Z_t^\top A_t \pi(A_t | S_t) w(S_t)}{\widetilde{\Delta}(S_t, A_t) \widetilde{\Theta}(S_t, Z_t)} \right) (f(S_t) - \gamma f(S_{t+1})) \right] \right| \\ \leq c \cdot \frac{C_{\Delta^*} C_{\Theta^*} C_*}{1 - \gamma} \sqrt{\frac{\mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}, \Pi}}{NT\kappa} \log \frac{1}{\delta} \log(NT)}.$$

For the second term on the RHS of (C.4.10), by a similar argument as in (C.3.14), it holds with probability at least $1 - \delta$ that

$$(C.4.12) \quad \left| \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} \left(\frac{Z_t^\top A_t \pi(A_t | S_t) w(S_t)}{\Delta^*(S_t, A_t) \Theta^*(S_t, Z_t)} - \frac{Z_t^\top A_t \pi(A_t | S_t) w(S_t)}{\widetilde{\Delta}(S_t, A_t) \widetilde{\Theta}(S_t, Z_t)} \right) (f(S_t) - \gamma f(S_{t+1})) \right] \right| \\ \leq \frac{2C_{\Delta^*} C_{\Theta^*} C_*}{1 - \gamma} \left(\xi_0 C_{\Theta^*} \sqrt{\frac{C_{\Delta^*}}{NT\kappa} \mathfrak{C}_{\mathcal{F}_0} \log \frac{1}{\delta} \log(NT)} + \xi_1 C_{\Delta^*} \sqrt{\frac{C_{\Theta^*}}{NT\kappa} \mathfrak{C}_{\mathcal{F}_1} \log \frac{1}{\delta} \log(NT)} \right),$$

where in the first inequality, we use the fact that $\|f\|_\infty \leq 1/(1 - \gamma)$ and $\|w\|_\infty \leq C_*$; in the third inequality, we use Cauchy Schwarz inequality; while in the last inequality, we use Assumption 4.4.4 with the fact that $(\widetilde{\Delta}, \widetilde{\Theta}) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1$. Now, by plugging (C.4.11) and (C.4.12) into (C.4.10), with probability at least $1 - \delta$, it holds for any $w \in \cup_{(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1} \text{conf}_{\alpha_{\text{mis}}}^{\text{mis}}(\Delta, \Theta, \pi)$ and $(f, \pi) \in \mathcal{V} \times \Pi$ that

$$(C.4.13) \quad \left| \widehat{\Phi}_{\text{mis}}^\pi(w, f; \Delta^*, \Theta^*) - \widehat{\Phi}_{\text{mis}}^\pi(w, f; \widetilde{\Delta}, \widetilde{\Theta}) \right| \\ \leq c \cdot \frac{C_{\Delta^*}^2 C_{\Theta^*}^2 C_*}{1 - \gamma} (\xi_0 + \xi_1) \sqrt{\frac{1}{NT\kappa} \cdot \mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}, \Pi} \cdot \log \frac{1}{\delta} \log(NT)}.$$

Now, by plugging (C.4.8), (C.4.9), and (C.4.13) into (C.4.7), with probability at least $1 - \delta$, it holds for any $\pi \in \Pi$ and $w \in \cup_{(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1} \text{conf}_{\alpha_{\text{mis}}}^{\text{mis}}(\Delta, \Theta, \pi)$ that

$$\max_{f \in \mathcal{V}} \Phi_{\text{mis}}^\pi(w, f; \Delta^*, \Theta^*) \leq c \cdot \frac{C_{\Delta^*}^2 C_{\Theta^*}^2 C^*}{1 - \gamma} (\xi_0 + \xi_1) \sqrt{\frac{1}{NT\kappa} \cdot \mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}, \Pi} \cdot \log \frac{1}{\delta} \log(NT)},$$

which concludes the proof of the lemma. \square

C.4.4. Proof of Lemma C.4.1

Proof. Since $\Phi_{\text{mis}}^\pi(w^\pi, V^\pi; \Delta^*, \Theta^*) = 0$, we have

$$\begin{aligned} & \Phi_{\text{mis}}^\pi(w, V^\pi; \Delta^*, \Theta^*) \\ &= \Phi_{\text{mis}}^\pi(w, V^\pi; \Delta^*, \Theta^*) - \Phi_{\text{mis}}^\pi(w^\pi, V^\pi; \Delta^*, \Theta^*) \\ &= \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} \frac{Z_t^\top A_t \pi(A_t | S_t)}{\Delta^*(S_t, A_t) \Theta^*(S_t, Z_t)} (w^\pi(S_t) - w(S_t)) (V^\pi(S_t) - \gamma V^\pi(S_{t+1})) \right] \\ &= \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} (w^\pi(S_t) - w(S_t)) \mathbb{E}_\pi [V^\pi(S_t) - \gamma V^\pi(S_{t+1}) | S_t] \right] \\ &= \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} (w^\pi(S_t) - w(S_t)) \mathbb{E}_\pi [R_t | S_t] \right] \\ &= \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} \frac{Z_t^\top A_t \pi(A_t | S_t)}{\Delta^*(S_t, A_t) \Theta^*(S_t, Z_t)} (w^\pi(S_t) - w(S_t)) R_t \right] \\ &= L_{\text{mis}}(w^\pi, \pi; \Delta^*, \Theta^*) - L_{\text{mis}}(w, \pi; \Delta^*, \Theta^*), \end{aligned}$$

which concludes the proof of the lemma. \square

C.4.5. Proof of Lemma C.4.2

Proof. With a slight abuse of notations, we denote by

$$w_0 \in \underset{w \in \text{conf}_{\alpha_{\text{mis}}}^{\text{mis}}(\Delta^*, \Theta^*, \pi^*)}{\text{argmin}} L_{\text{mis}}(w, \pi^*; \Delta^*, \Theta^*), \quad w_1 \in \underset{w \in \text{conf}_{\alpha_{\text{mis}}}^{\text{mis}}(\Delta, \Theta, \pi^*)}{\text{argmin}} L_{\text{mis}}(w, \pi^*; \Delta, \Theta).$$

Then we have

$$\begin{aligned}
& \left| \min_{w \in \text{conf}_{\alpha_{\text{mis}}}^{\text{mis}}(\Delta^*, \Theta^*, \pi^*)} L_{\text{mis}}(w, \pi^*; \Delta^*, \Theta^*) - \min_{w \in \text{conf}_{\alpha_{\text{mis}}}^{\text{mis}}(\Delta, \Theta, \pi^*)} L_{\text{mis}}(w, \pi^*; \Delta, \Theta) \right| \\
&= |L_{\text{mis}}(w_0, \pi^*; \Delta^*, \Theta^*) - L_{\text{mis}}(w_1, \pi^*; \Delta, \Theta)| \\
&\leq |L_{\text{mis}}(w_0, \pi^*; \Delta^*, \Theta^*) - L_{\text{mis}}(w^{\pi^*}, \pi^*; \Delta^*, \Theta^*)| \\
&\quad + |L_{\text{mis}}(w^{\pi^*}, \pi^*; \Delta^*, \Theta^*) - L_{\text{mis}}(w_1, \pi^*; \Delta^*, \Theta^*)| \\
&\quad + |L_{\text{mis}}(w_1, \pi^*; \Delta^*, \Theta^*) - L_{\text{mis}}(w_1, \pi^*; \Delta, \Theta)| \\
\text{(C.4.14)} \quad &= \underbrace{|\Phi_{\text{mis}}^{\pi^*}(w_0, V^{\pi^*}; \Delta^*, \Theta^*)|}_{\text{Term (I)}} + \underbrace{|\Phi_{\text{mis}}^{\pi^*}(w_1, V^{\pi^*}; \Delta^*, \Theta^*)|}_{\text{Term (II)}} \\
&\quad + \underbrace{|L_{\text{mis}}(w_1, \pi^*; \Delta^*, \Theta^*) - L_{\text{mis}}(w_1, \pi^*; \Delta, \Theta)|}_{\text{Term (III)}}.
\end{aligned}$$

We upper bound terms (I), (II), and (III) on the RHS of (C.4.14), respectively.

Upper Bounding Term (I). Note that with probability at least $1 - \delta$, we have

$$\begin{aligned}
|\Phi_{\text{mis}}^{\pi^*}(w_0, V^{\pi^*}; \Delta^*, \Theta^*)| &\leq \max_{f \in \mathcal{V}} |\Phi_{\text{mis}}^{\pi^*}(w_0, f; \Delta^*, \Theta^*)| \\
&= \max_{f \in \mathcal{V}} \max \{ \Phi_{\text{mis}}^{\pi^*}(w_0, f; \Delta^*, \Theta^*), -\Phi_{\text{mis}}^{\pi^*}(w_0, f; \Delta^*, \Theta^*) \} \\
&= \max_{f \in \mathcal{V}} \max \{ \Phi_{\text{mis}}^{\pi^*}(w_0, f; \Delta^*, \Theta^*), \Phi_{\text{mis}}^{\pi^*}(w_0, -f; \Delta^*, \Theta^*) \} \\
&= \max_{f \in \mathcal{V}} \Phi_{\text{mis}}^{\pi^*}(w_0, f; \Delta^*, \Theta^*) \\
\text{(C.4.15)} \quad &\leq c \cdot \frac{C_{\Delta^*}^2 C_{\Theta^*}^2 C_*}{1 - \gamma} (\xi_0 + \xi_1) \sqrt{\frac{1}{NT\kappa} \mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}} \log \frac{1}{\delta} \log(NT)},
\end{aligned}$$

where in the first inequality, we use the fact that $V^{\pi^*} \in \mathcal{V}$; in the third equality, we use the fact that \mathcal{V} is symmetric; in the last inequality, by noting that $w_0 \in \cup_{(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1} \text{conf}_{\alpha_{\text{mis}}}^{\text{mis}}(\Delta, \Theta, \pi)$, we use Lemma 4.4.12. This upper bounds term (I) on the RHS of (C.4.14).

Upper Bounding Term (II). Similar to (C.4.15), note that $w_1 \in \cup_{(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1} \text{conf}_{\alpha_{\text{mis}}}^{\text{mis}}(\Delta, \Theta, \pi)$, it holds with probability at least $1 - \delta$ that

$$\text{(C.4.16)} \quad |\Phi_{\text{mis}}^{\pi^*}(w_1, V^{\pi^*}; \Delta^*, \Theta^*)| \leq c \cdot \frac{C_{\Delta^*}^2 C_{\Theta^*}^2 C_*}{1 - \gamma} (\xi_0 + \xi_1) \sqrt{\frac{1}{NT\kappa} \mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}} \log \frac{1}{\delta} \log(NT)},$$

which upper bounds term (II) on the RHS of (C.4.14).

Upper Bounding Term (III). Note that with probability at least $1 - \delta$, it holds for any $(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1$ that

$$\begin{aligned}
& |L_{\text{mis}}(w_1, \pi^*; \Delta^*, \Theta^*) - L_{\text{mis}}(w_1, \pi^*; \Delta, \Theta)| \\
&= \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} \left(\frac{Z_t^\top A_t \pi^*(A_t | S_t)}{\Delta^*(S_t, A_t) \Theta^*(S_t, Z_t)} - \frac{Z_t^\top A_t \pi^*(A_t | S_t)}{\Delta(S_t, A_t) \Theta(S_t, Z_t)} \right) w_1(S_t) R_t \right] \\
&\text{(C.4.17)} \\
&\leq C_{\Delta^*} C_{\Theta^*} C_* \left(\xi_0 C_{\Theta^*} \sqrt{\frac{C_{\Delta^*}}{NT\kappa} \mathfrak{C}_{\mathcal{F}_0} \log \frac{1}{\delta} \log(NT)} + \xi_1 C_{\Delta^*} \sqrt{\frac{C_{\Theta^*}}{NT\kappa} \mathfrak{C}_{\mathcal{F}_1} \log \frac{1}{\delta} \log(NT)} \right),
\end{aligned}$$

where we use Cauchy-Schwarz inequality and Assumption 4.4.4 in the last inequality.

Now, by plugging (C.4.15), (C.4.16), and (C.4.17) into (C.4.14), with probability at least $1 - \delta$, it holds for any $(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1$ that

$$\begin{aligned}
& \left| \min_{w \in \text{conf}_{\alpha_{\text{mis}}}^{\text{mis}}(\Delta^*, \Theta^*, \pi^*)} L_{\text{mis}}(w, \pi^*; \Delta^*, \Theta^*) - \min_{w \in \text{conf}_{\alpha_{\text{mis}}}^{\text{mis}}(\Delta, \Theta, \pi^*)} L_{\text{mis}}(w, \pi^*; \Delta, \Theta) \right| \\
&\leq c \cdot \frac{C_{\Delta^*}^2 C_{\Theta^*}^2 C_*}{1 - \gamma} (\xi_0 + \xi_1) \sqrt{\frac{1}{NT\kappa} \mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}} \log \frac{1}{\delta} \log(NT)},
\end{aligned}$$

which concludes the proof of the lemma. \square

C.4.6. Proof of Lemma C.4.3

Proof. By Theorem C.7.6, with probability at least $1 - \delta$, it holds for any $(w, \Delta, \Theta, \pi) \in \mathcal{W} \times \mathcal{F}_0 \times \mathcal{F}_1 \times \Pi$ that

$$\left| L_{\text{mis}}(w, \pi; \Delta, \Theta) - \widehat{L}_{\text{mis}}(w, \pi; \Delta, \Theta) \right| \leq c \cdot C_{\Delta^*} C_{\Theta^*} C_* \sqrt{\frac{1}{NT\kappa} \mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \Pi} \log \frac{1}{\delta} \log(NT)},$$

which concludes the proof of the lemma. \square

C.4.7. Proof of Lemma C.4.4

Proof. Note that

$$\begin{aligned}
& \max_{f \in \mathcal{V}} \Phi_{\text{mis}}^\pi(\widehat{w}_{\Delta, \Theta}^\pi, f; \Delta, \Theta) \\
&= \max_{f \in \mathcal{V}} \Phi_{\text{mis}}^\pi(\widehat{w}_{\Delta, \Theta}^\pi, f; \Delta, \Theta) - \max_{f \in \mathcal{V}} \widehat{\Phi}_{\text{mis}}^\pi(\widehat{w}_{\Delta, \Theta}^\pi, f; \Delta, \Theta) + \max_{f \in \mathcal{V}} \widehat{\Phi}_{\text{mis}}^\pi(\widehat{w}_{\Delta, \Theta}^\pi, f; \Delta, \Theta) \\
&\quad - \max_{f \in \mathcal{V}} \widehat{\Phi}_{\text{mis}}^\pi(w^\pi, f; \Delta, \Theta) + \max_{f \in \mathcal{V}} \widehat{\Phi}_{\text{mis}}^\pi(w^\pi, f; \Delta, \Theta) - \max_{f \in \mathcal{V}} \Phi_{\text{mis}}^\pi(w^\pi, f; \Delta, \Theta) \\
&\quad + \max_{f \in \mathcal{V}} \Phi_{\text{mis}}^\pi(w^\pi, f; \Delta, \Theta) - \max_{f \in \mathcal{V}} \Phi_{\text{mis}}^\pi(w^\pi, f; \Delta^*, \Theta^*) \\
&\leq 2 \max_{w \in \mathcal{W}} \max_{f \in \mathcal{V}} \left| \Phi_{\text{mis}}^\pi(w, f; \Delta, \Theta) - \widehat{\Phi}_{\text{mis}}^\pi(w, f; \Delta, \Theta) \right|
\end{aligned}$$

(C.4.18)

$$+ \max_{f \in \mathcal{V}} |\Phi_{\text{mis}}^\pi(w^\pi, f; \Delta, \Theta) - \Phi_{\text{mis}}^\pi(w^\pi, f; \Delta^*, \Theta^*)|,$$

where we use the fact that $\widehat{w}_{\Delta, \Theta}^\pi \in \arg\min_{w \in \mathcal{W}} \max_{f \in \mathcal{V}} \widehat{\Phi}_{\text{mis}}^\pi(w, f; \Delta, \Theta)$ in the last inequality. In the meanwhile, by Theorem C.7.6, with probability at least $1 - \delta$, it holds for any $(w, f, \pi) \in \mathcal{W} \times \mathcal{V} \times \Pi$ that

(C.4.19)

$$\left| \widehat{\Phi}_{\text{mis}}^\pi(w, f; \Delta, \Theta) - \Phi_{\text{mis}}^\pi(w, f; \Delta, \Theta) \right| \leq c \cdot \frac{C_{\Delta^*} C_{\Theta^*} C_*}{1 - \gamma} \sqrt{\frac{1}{NT\kappa} \mathfrak{C}_{\mathcal{V}, \mathcal{W}, \Pi} \log \frac{1}{\delta} \log(NT)}.$$

Also, we upper bound the second term on the RHS of (C.4.18) with probability at least $1 - \delta$ as follows,

$$\begin{aligned}
& |\Phi_{\text{mis}}^\pi(w^\pi, f; \Delta, \Theta) - \Phi_{\text{mis}}^\pi(w^\pi, f; \Delta^*, \Theta^*)| \\
&= \left| \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} \left(\frac{Z_t^\top A_t \pi(A_t | S_t) w(S_t)}{\Delta^*(S_t, A_t) \Theta^*(S_t, Z_t)} - \frac{Z_t^\top A_t \pi(A_t | S_t) w(S_t)}{\Delta(S_t, A_t) \Theta(S_t, Z_t)} \right) (f(S_t) - \gamma f(S_{t+1})) \right] \right| \\
\text{(C.4.20)} \quad &\leq \frac{C_{\Delta^*} C_{\Theta^*} C_*}{1 - \gamma} \left(\xi_0 C_{\Theta^*} \sqrt{\frac{C_{\Delta^*}}{NT\kappa} \mathfrak{C}_{\mathcal{F}_0} \log \frac{1}{\delta} \log(NT)} + \xi_1 C_{\Delta^*} \sqrt{\frac{C_{\Theta^*}}{NT\kappa} \mathfrak{C}_{\mathcal{F}_1} \log \frac{1}{\delta} \log(NT)} \right),
\end{aligned}$$

where we use Cauchy-Schwarz inequality and Assumption 4.4.4 in the last inequality.

Now, by plugging (C.4.19) and (C.4.20) into (C.4.18), with probability at least $1 - \delta$, it holds for any $(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1$ and $\pi \in \Pi$ that

$$\max_{f \in \mathcal{V}} \Phi_{\text{mis}}^\pi(\widehat{w}_{\Delta, \Theta}^\pi, f; \Delta, \Theta) \leq c \cdot \frac{C_{\Delta^*}^2 C_{\Theta^*}^2 C_*}{1 - \gamma} (\xi_0 + \xi_1) \sqrt{\frac{1}{NT\kappa} \mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}, \Pi} \log \frac{1}{\delta} \log(NT)},$$

which concludes the proof of the lemma. \square

C.5. Proof of Results in §4.4.3

C.5.1. Proof of Theorem 4.4.14

Proof. We split the proof into two case: (i) Assumption 4.4.6 holds; (ii) Assumption 4.4.10 holds.

Case (i): Assumption 4.4.6 holds. We introduce the following supporting lemmas.

Lemma C.5.1. For any policy π , with probability at least $1 - \delta$ with $c/(NT)^2 \leq \delta \leq 1$, it holds for any $(w, v, \Delta, \Theta, \pi) \in \mathcal{W} \times \mathcal{V} \times \mathcal{F}_0 \times \mathcal{F}_1 \times \Pi$ that

$$\begin{aligned} & \left| L_{\text{dr}}(w, v, \pi; \Delta, \Theta) - \widehat{L}_{\text{dr}}(w, v, \pi; \Delta, \Theta) \right| \\ & \leq c \cdot \frac{C_{\Delta^*} C_{\Theta^*} C_*}{1 - \gamma} \sqrt{\frac{1}{NT\kappa} \mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}, \Pi} \log \frac{1}{\delta} \log(NT)} = \widehat{\epsilon}_L. \end{aligned}$$

Proof. See §C.5.3 for a detailed proof. □

Lemma C.5.2. Suppose that $(\alpha_0, \alpha_1, \alpha_{\text{mis}}, \alpha_{\text{vf}})$ is defined in Assumption 4.4.4, Lemmas 4.4.11, and 4.4.7. With probability at least $1 - \delta$, it holds for any $(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1$ that

$$\begin{aligned} & \left| \min_{(w, v) \in \text{conf}_{\alpha_{\text{mis}}, \alpha_{\text{vf}}}(\Delta^*, \Theta^*, \pi^*)} L_{\text{dr}}(w, v, \pi^*; \Delta^*, \Theta^*) - \min_{(w, v) \in \text{conf}_{\alpha_{\text{mis}}, \alpha_{\text{vf}}}(\Delta, \Theta, \pi^*)} L_{\text{dr}}(w, v, \pi^*; \Delta, \Theta) \right| \\ & \leq c \cdot \frac{C_{\Delta^*}^2 C_{\Theta^*}^2 C_*}{1 - \gamma} (\xi_0 + \xi_1) \sqrt{\frac{1}{NT\kappa} \mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}} \log \frac{1}{\delta} \log(NT)} = \epsilon_L^*. \end{aligned}$$

Proof. See §C.5.4 for a detailed proof. □

By the definition of L_{dr} , it holds with probability at least $1 - \delta$ that

$$\begin{aligned}
J(\pi^*) - J(\widehat{\pi}_{\text{dr}}) &= J(\pi^*) - L_{\text{dr}}(w, V^{\widehat{\pi}_{\text{dr}}}, \widehat{\pi}_{\text{dr}}; \Delta^*, \Theta^*) \\
&\leq J(\pi^*) - \min_{(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1} \min_{(w, v) \in \text{conf}_{\alpha_{\text{mis}}, \alpha_{\text{vf}}}(\Delta, \Theta, \widehat{\pi}_{\text{dr}})} L_{\text{dr}}(w, v, \widehat{\pi}_{\text{dr}}; \Delta, \Theta) \\
&\leq J(\pi^*) - \min_{(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1} \min_{(w, v) \in \text{conf}_{\alpha_{\text{mis}}, \alpha_{\text{vf}}}(\Delta, \Theta, \widehat{\pi}_{\text{dr}})} \widehat{L}_{\text{dr}}(w, v, \widehat{\pi}_{\text{dr}}; \Delta, \Theta) + \widehat{\epsilon}_L \\
&\leq J(\pi^*) - \min_{(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1} \min_{(w, v) \in \text{conf}_{\alpha_{\text{mis}}, \alpha_{\text{vf}}}(\Delta, \Theta, \pi^*)} \widehat{L}_{\text{dr}}(w, v, \pi^*; \Delta, \Theta) + \widehat{\epsilon}_L \\
\text{(C.5.1)} \quad &\leq J(\pi^*) - \min_{(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1} \min_{(w, v) \in \text{conf}_{\alpha_{\text{mis}}, \alpha_{\text{vf}}}(\Delta, \Theta, \pi^*)} L_{\text{dr}}(w, v, \pi^*; \Delta, \Theta) + 2\widehat{\epsilon}_L,
\end{aligned}$$

where in the first inequality, we use Assumption 4.4.4 that $(\Delta^*, \Theta^*) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1$ with probability at least $1 - \delta$, and Lemma 4.4.7 with Assumption 4.4.6 that $V^{\widehat{\pi}_{\text{dr}}} \in \text{conf}_{\alpha_{\text{vf}}}^{\text{vf}}(\Delta^*, \Theta^*, \widehat{\pi}_{\text{dr}})$ with probability at least $1 - \delta$; in the second inequality, we use Lemma C.5.1; in the third inequality, we use the optimality of $\widehat{\pi}_{\text{dr}}$; while in the last inequality, we use Lemma C.5.1 again. By combining Lemma C.5.2 and (C.5.1), we have

$$\begin{aligned}
&J(\pi^*) - J(\widehat{\pi}_{\text{dr}}) \\
&\leq J(\pi^*) - \min_{(w, v) \in \text{conf}_{\alpha_{\text{mis}}, \alpha_{\text{vf}}}(\Delta^*, \Theta^*, \pi^*)} L_{\text{dr}}(w, v, \pi^*; \Delta^*, \Theta^*) + 2\widehat{\epsilon}_L + \epsilon_L^* \\
&= L_{\text{dr}}(w, V^{\pi^*}, \pi^*; \Delta^*, \Theta^*) - \min_{(w, v) \in \text{conf}_{\alpha_{\text{mis}}, \alpha_{\text{vf}}}(\Delta^*, \Theta^*, \pi^*)} L_{\text{dr}}(w, v, \pi^*; \Delta^*, \Theta^*) + 2\widehat{\epsilon}_L + \epsilon_L^* \\
&= \max_{(w, v) \in \text{conf}_{\alpha_{\text{mis}}, \alpha_{\text{vf}}}(\Delta^*, \Theta^*, \pi^*)} \left| L_{\text{dr}}(w, V^{\pi^*}, \pi^*; \Delta^*, \Theta^*) - L_{\text{dr}}(w, v, \pi^*; \Delta^*, \Theta^*) \right| + 2\widehat{\epsilon}_L + \epsilon_L^* \\
\text{(C.5.2)} \quad &= \max_{(w, v) \in \text{conf}_{\alpha_{\text{mis}}, \alpha_{\text{vf}}}(\Delta^*, \Theta^*, \pi^*)} \left| \Phi_{\text{vf}}^{\pi^*}(v, w^{\pi^*}; \Delta^*, \Theta^*) - \Phi_{\text{vf}}^{\pi^*}(v, w; \Delta^*, \Theta^*) \right| + 2\widehat{\epsilon}_L + \epsilon_L^*,
\end{aligned}$$

where we use the following fact in the last equality,

$$L_{\text{dr}}(w, V^{\pi^*}, \pi^*; \Delta^*, \Theta^*) - L_{\text{dr}}(w, v, \pi^*; \Delta^*, \Theta^*) = \Phi_{\text{vf}}^{\pi^*}(v, w^{\pi^*}; \Delta^*, \Theta^*) - \Phi_{\text{vf}}^{\pi^*}(v, w; \Delta^*, \Theta^*).$$

In the meanwhile, note that by Assumption 4.4.6 that $w^{\pi^*} \in \mathcal{W}$, we obtain from (C.5.2) that

$$\begin{aligned} J(\pi^*) - J(\widehat{\pi}_{\text{dr}}) &\leq 2 \max_{(w,v) \in \text{conf}_{\alpha_{\text{mis}}, \alpha_{\text{vf}}}(\Delta^*, \Theta^*, \pi^*)} |\Phi_{\text{vf}}^{\pi^*}(v, w; \Delta^*, \Theta^*)| + 2\widehat{\epsilon}_L + \epsilon_L^* \\ \text{(C.5.3)} \quad &\leq c \cdot \frac{C_{\Delta^*}^2 C_{\Theta^*}^2 C^*}{1 - \gamma} (\xi_0 + \xi_1) \sqrt{\frac{1}{NT\kappa} \mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \nu, \Pi} \log \frac{NT}{\delta}}, \end{aligned}$$

where we use Lemma 4.4.8 and plug in the definition of ϵ_L and ϵ_L^* in the last inequality.

This concludes the proof of case (i).

Case (ii): Assumption 4.4.10 holds. It holds with probability at least $1 - \delta$ that

$$\begin{aligned} J(\pi^*) - J(\widehat{\pi}_{\text{dr}}) &= J(\pi^*) - L_{\text{dr}}(w^{\widehat{\pi}_{\text{dr}}}, v, \widehat{\pi}_{\text{dr}}; \Delta^*, \Theta^*) \\ &\leq J(\pi^*) - \min_{(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1} \min_{(w,v) \in \text{conf}_{\alpha_{\text{mis}}, \alpha_{\text{vf}}}(\Delta, \Theta, \widehat{\pi}_{\text{dr}})} L_{\text{dr}}(w, v, \widehat{\pi}_{\text{dr}}; \Delta, \Theta) \\ &\leq J(\pi^*) - \min_{(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1} \min_{(w,v) \in \text{conf}_{\alpha_{\text{mis}}, \alpha_{\text{vf}}}(\Delta, \Theta, \widehat{\pi}_{\text{dr}})} \widehat{L}_{\text{dr}}(w, v, \widehat{\pi}_{\text{dr}}; \Delta, \Theta) + \widehat{\epsilon}_L \\ &\leq J(\pi^*) - \min_{(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1} \min_{(w,v) \in \text{conf}_{\alpha_{\text{mis}}, \alpha_{\text{vf}}}(\Delta, \Theta, \pi^*)} \widehat{L}_{\text{dr}}(w, v, \pi^*; \Delta, \Theta) + \widehat{\epsilon}_L, \end{aligned}$$

where we use Assumption 4.4.4 that $(\Delta^*, \Theta^*) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1$ with probability at least $1 - \delta$ and Assumption 4.4.10 that $w^\pi \in \mathcal{W}$ for any $\pi \in \Pi$ in the first inequality, we use Lemma C.5.1 in the second inequality, and we use the optimality of $\widehat{\pi}_{\text{dr}}$ in the last

inequality. Further, by Lemmas C.5.1 and C.5.2, we have

$$\begin{aligned}
& J(\pi^*) - J(\widehat{\pi}_{\text{dr}}) \\
& \leq J(\pi^*) - \min_{(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1} \min_{(w, v) \in \text{conf}_{\alpha_{\text{mis}}, \alpha_{\text{vf}}}(\Delta, \Theta, \pi^*)} L_{\text{dr}}(w, v, \pi^*; \Delta, \Theta) + 2\widehat{\epsilon}_L \\
& \leq J(\pi^*) - \min_{(w, v) \in \text{conf}_{\alpha_{\text{mis}}, \alpha_{\text{vf}}}(\Delta^*, \Theta^*, \pi^*)} L_{\text{dr}}(w, v, \pi^*; \Delta^*, \Theta^*) + 2\widehat{\epsilon}_L + \epsilon_L^* \\
& \leq \max_{(w, v) \in \text{conf}_{\alpha_{\text{mis}}, \alpha_{\text{vf}}}(\Delta^*, \Theta^*, \pi^*)} \left| L_{\text{dr}}(w^{\pi^*}, v, \pi^*; \Delta^*, \Theta^*) - L_{\text{dr}}(w, v, \pi^*; \Delta^*, \Theta^*) \right| + 2\widehat{\epsilon}_L + \epsilon_L^* \\
& \leq \max_{(w, v) \in \text{conf}_{\alpha_{\text{mis}}, \alpha_{\text{vf}}}(\Delta^*, \Theta^*, \pi^*)} \left| \Phi_{\text{mis}}^{\pi^*}(w, V^{\pi^*}; \Delta^*, \Theta^*) - \Phi_{\text{mis}}^{\pi^*}(w, v; \Delta^*, \Theta^*) \right| + 2\widehat{\epsilon}_L + \epsilon_L^* \\
& \leq 2 \max_{(w, v) \in \text{conf}_{\alpha_{\text{mis}}, \alpha_{\text{vf}}}(\Delta^*, \Theta^*, \pi^*)} \left| \Phi_{\text{mis}}^{\pi^*}(w, v; \Delta^*, \Theta^*) \right| + 2\widehat{\epsilon}_L + \epsilon_L^* \\
& \leq 2 \max_{w \in \text{conf}_{\alpha_{\text{mis}}}^{\text{mis}}(\Delta^*, \Theta^*, \pi^*)} \max_{v \in \mathcal{V}} \left| \Phi_{\text{mis}}^{\pi^*}(w, v; \Delta^*, \Theta^*) \right| + 2\widehat{\epsilon}_L + \epsilon_L^*
\end{aligned}$$

where in the third inequality, we use the fact that $J(\pi^*) = L_{\text{dr}}(w^{\pi^*}, v, \pi^*; \Delta^*, \Theta^*)$; in the fourth inequality, we use the following fact

$$L_{\text{dr}}(w, v, \pi^*; \Delta^*, \Theta^*) - L_{\text{dr}}(w^{\pi^*}, v, \pi^*; \Delta^*, \Theta^*) = -\Phi_{\text{mis}}^{\pi^*}(w, V^{\pi^*}; \Delta^*, \Theta^*) + \Phi_{\text{mis}}^{\pi^*}(w, v; \Delta^*, \Theta^*)$$

for any $(w, v) \in \mathcal{W} \times \mathcal{V}$; in the fifth inequality, we use the fact that $V^{\pi^*} \in \mathcal{V}$ by Assumption 4.4.10 and $V^{\pi^*} \in \text{conf}_{\alpha_{\text{vf}}}^{\text{vf}}(\Delta^*, \Theta^*, \pi^*)$ with probability at least $1 - \delta$ by Lemma 4.4.7. Now, by Lemma 4.4.12 and the fact that \mathcal{V} is symmetric, we obtain that

$$(C.5.4) \quad J(\pi^*) - J(\widehat{\pi}_{\text{dr}}) \leq c \cdot \frac{C_{\Delta^*}^2 C_{\Theta^*}^2 C_{\pi^*}}{1 - \gamma} (\xi_0 + \xi_1) \sqrt{\frac{1}{NT\kappa} \mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}, \Pi} \log \frac{NT}{\delta}},$$

which concludes the proof of case (ii).

By combining (C.5.3) and (C.5.4), we conclude the proof of the theorem. \square

C.5.2. Proof of Theorem 4.4.16

Proof. Recall that

$$\tilde{v}^\pi \in \operatorname{argmin}_{v \in \mathcal{V}} \max_{w \in \mathcal{W}} \Phi_{\text{vf}}^\pi(v, w; \Delta^*, \Theta^*), \quad \tilde{w}^\pi \in \operatorname{argmin}_{w \in \mathcal{W}} \max_{v \in \mathcal{V}} \Phi_{\text{mis}}^\pi(w, v; \Delta^*, \Theta^*).$$

We split the proof into the following two parts.

Part (i). We first introduce the following lemmas.

Lemma C.5.3. Suppose α_{vf} is defined in Lemma 4.4.7 and $c/(NT)^2 \leq \delta \leq 1$. Then under Assumptions (b) and 4.4.3, with probability at least $1 - \delta$, it holds for any $\pi \in \Pi$ that $\tilde{v}^\pi \in \operatorname{conf}_{\alpha_{\text{vf}}}^{\text{vf}}(\Delta^*, \Theta^*, \pi)$.

Proof. See §C.5.5 for a detailed proof. □

Lemma C.5.4. Suppose that $(\alpha_0, \alpha_1, \alpha_{\text{vf}})$ is defined in Assumption 4.4.4 and Lemma 4.4.7 and $c/(NT)^2 \leq \delta \leq 1$. Then under Assumptions 4.2.2, 4.2.3, 4.4.3 and 4.4.4, with probability at least $1 - \delta$, it holds for any policy $\pi \in \Pi$ and $v \in \cup_{(\Delta, \Theta) \in \operatorname{conf}_{\alpha_0}^0 \times \operatorname{conf}_{\alpha_1}^1} \operatorname{conf}_{\alpha_{\text{vf}}}^{\text{vf}}(\Delta, \Theta, \pi)$ that

$$\begin{aligned} \max_{g \in \mathcal{W}} \Phi_{\text{vf}}^\pi(v, g; \Delta^*, \Theta^*) &\leq c \cdot \frac{C_{\Delta^*}^2 C_{\Theta^*}^2 C_*}{1 - \gamma} (\xi_0 + \xi_1) \sqrt{\frac{1}{NT\kappa} \cdot \mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}, \Pi} \cdot \log \frac{1}{\delta} \log(NT)} \\ &\quad + \max_{g \in \mathcal{W}} \Phi_{\text{vf}}^\pi(\tilde{v}^\pi, g; \Delta^*, \Theta^*). \end{aligned}$$

Proof. See §C.5.6 for a detailed proof. □

By the definition of L_{dr} , it holds that

$$\begin{aligned}
& J(\pi^*) - J(\widehat{\pi}_{\text{dr}}) \\
&= J(\pi^*) - L_{\text{dr}}(w, V^{\widehat{\pi}_{\text{dr}}}, \widehat{\pi}_{\text{dr}}; \Delta^*, \Theta^*) \\
\text{(C.5.5)} \quad &= J(\pi^*) - L_{\text{dr}}(w, \widetilde{v}^{\widehat{\pi}_{\text{dr}}}, \widehat{\pi}_{\text{dr}}; \Delta^*, \Theta^*) + L_{\text{dr}}(w, \widetilde{v}^{\widehat{\pi}_{\text{dr}}}, \widehat{\pi}_{\text{dr}}; \Delta^*, \Theta^*) - L_{\text{dr}}(w, V^{\widehat{\pi}_{\text{dr}}}, \widehat{\pi}_{\text{dr}}; \Delta^*, \Theta^*).
\end{aligned}$$

Note that

$$\text{(C.5.6)} \quad \left| L_{\text{dr}}(w, \widetilde{v}^{\widehat{\pi}_{\text{dr}}}, \widehat{\pi}_{\text{dr}}; \Delta^*, \Theta^*) - L_{\text{dr}}(w, V^{\widehat{\pi}_{\text{dr}}}, \widehat{\pi}_{\text{dr}}; \Delta^*, \Theta^*) \right| \leq C_* C_{\Delta^*} C_{\Theta^*} \varepsilon_{\text{vf}}^{\mathcal{V}}.$$

In the meanwhile, by Assumption 4.4.4 and Lemma C.5.3, it holds with probability at least $1 - \delta$ that

$$\begin{aligned}
L_{\text{dr}}(w, \widetilde{v}^{\widehat{\pi}_{\text{dr}}}, \widehat{\pi}_{\text{dr}}; \Delta^*, \Theta^*) &\geq \min_{(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1} \min_{(w, v) \in \text{conf}_{\alpha_{\text{mis}}, \alpha_{\text{vf}}}(\Delta, \Theta, \widehat{\pi}_{\text{dr}})} L_{\text{dr}}(w, v, \widehat{\pi}_{\text{dr}}; \Delta, \Theta) \\
&\geq \min_{(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1} \min_{(w, v) \in \text{conf}_{\alpha_{\text{mis}}, \alpha_{\text{vf}}}(\Delta, \Theta, \widehat{\pi}_{\text{dr}})} \widehat{L}_{\text{dr}}(w, v, \widehat{\pi}_{\text{dr}}; \Delta, \Theta) - \widehat{\epsilon}_L \\
&\geq \min_{(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1} \min_{(w, v) \in \text{conf}_{\alpha_{\text{mis}}, \alpha_{\text{vf}}}(\Delta, \Theta, \pi^*)} \widehat{L}_{\text{dr}}(w, v, \pi^*; \Delta, \Theta) - \widehat{\epsilon}_L \\
\text{(C.5.7)} \quad &\geq \min_{(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1} \min_{(w, v) \in \text{conf}_{\alpha_{\text{mis}}, \alpha_{\text{vf}}}(\Delta, \Theta, \pi^*)} L_{\text{dr}}(w, v, \pi^*; \Delta, \Theta) - 2\widehat{\epsilon}_L,
\end{aligned}$$

where in the second inequality, we use Lemma C.5.1; in the third inequality, we use the optimality of $\widehat{\pi}_{\text{dr}}$; in the fourth inequality, we again use Lemma C.5.1. Now, by plugging

(C.5.6) and (C.5.7) into (C.5.5), we have

$$(C.5.8) \quad J(\pi^*) - J(\widehat{\pi}_{\text{dr}}) \leq J(\pi^*) - \min_{(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1} \min_{(w, v) \in \text{conf}_{\alpha_{\text{mis}}, \alpha_{\text{vf}}}(\Delta, \Theta, \pi^*)} L_{\text{dr}}(w, v, \pi^*; \Delta, \Theta)$$

$$(C.5.9) \quad + 2\widehat{\epsilon}_L + C_* C_{\Delta^*} C_{\Theta^*} \epsilon_{\text{vf}}^{\mathcal{V}}.$$

By combining Lemma C.5.2 and (C.5.8), we have

$$\begin{aligned} & J(\pi^*) - J(\widehat{\pi}_{\text{dr}}) \\ & \leq J(\pi^*) - \min_{(w, v) \in \text{conf}_{\alpha_{\text{mis}}, \alpha_{\text{vf}}}(\Delta^*, \Theta^*, \pi^*)} L_{\text{dr}}(w, v, \pi^*; \Delta^*, \Theta^*) + 2\widehat{\epsilon}_L + \epsilon_L^* + C_* C_{\Delta^*} C_{\Theta^*} \epsilon_{\text{vf}}^{\mathcal{V}} \\ & = L_{\text{dr}}(w, V^{\pi^*}, \pi^*; \Delta^*, \Theta^*) - \min_{(w, v) \in \text{conf}_{\alpha_{\text{mis}}, \alpha_{\text{vf}}}(\Delta^*, \Theta^*, \pi^*)} L_{\text{dr}}(w, v, \pi^*; \Delta^*, \Theta^*) \\ & \quad + 2\widehat{\epsilon}_L + \epsilon_L^* + C_* C_{\Delta^*} C_{\Theta^*} \epsilon_{\text{vf}}^{\mathcal{V}} \\ & = \max_{(w, v) \in \text{conf}_{\alpha_{\text{mis}}, \alpha_{\text{vf}}}(\Delta^*, \Theta^*, \pi^*)} |L_{\text{dr}}(w, V^{\pi^*}, \pi^*; \Delta^*, \Theta^*) - L_{\text{dr}}(w, v, \pi^*; \Delta^*, \Theta^*)| \\ & \quad + 2\widehat{\epsilon}_L + \epsilon_L^* + C_* C_{\Delta^*} C_{\Theta^*} \epsilon_{\text{vf}}^{\mathcal{V}} \\ (C.5.10) \quad & = \max_{(w, v) \in \text{conf}_{\alpha_{\text{mis}}, \alpha_{\text{vf}}}(\Delta^*, \Theta^*, \pi^*)} |\Phi_{\text{vf}}^{\pi^*}(v, w^{\pi^*}; \Delta^*, \Theta^*) - \Phi_{\text{vf}}^{\pi^*}(v, w; \Delta^*, \Theta^*)| \\ & \quad + 2\widehat{\epsilon}_L + \epsilon_L^* + C_* C_{\Delta^*} C_{\Theta^*} \epsilon_{\text{vf}}^{\mathcal{V}}, \end{aligned}$$

where we use the following fact in the last equality,

$$L_{\text{dr}}(w, V^{\pi^*}, \pi^*; \Delta^*, \Theta^*) - L_{\text{dr}}(w, v, \pi^*; \Delta^*, \Theta^*) = \Phi_{\text{vf}}^{\pi^*}(v, w^{\pi^*}; \Delta^*, \Theta^*) - \Phi_{\text{vf}}^{\pi^*}(v, w; \Delta^*, \Theta^*).$$

We upper bound the first term on the RHS of (C.5.10) as follows,

$$\begin{aligned}
& \max_{(w,v) \in \text{conf}_{\alpha_{\text{mis}}, \alpha_{\text{vf}}}(\Delta^*, \Theta^*, \pi^*)} \left| \Phi_{\text{vf}}^{\pi^*}(v, w^{\pi^*}; \Delta^*, \Theta^*) - \Phi_{\text{vf}}^{\pi^*}(v, w; \Delta^*, \Theta^*) \right| \\
& \leq \max_{v \in \text{conf}_{\alpha_{\text{vf}}}^{\text{f}}(\Delta^*, \Theta^*, \pi^*)} \max_{w \in \mathcal{W}} \left| \Phi_{\text{vf}}^{\pi^*}(v, \tilde{w}^{\pi^*}; \Delta^*, \Theta^*) - \Phi_{\text{vf}}^{\pi^*}(v, w; \Delta^*, \Theta^*) \right| \\
& \quad + \max_{v \in \mathcal{V}} \left| \Phi_{\text{vf}}^{\pi^*}(v, w^{\pi^*}; \Delta^*, \Theta^*) - \Phi_{\text{vf}}^{\pi^*}(v, \tilde{w}^{\pi^*}; \Delta^*, \Theta^*) \right| \\
\text{(C.5.11)} \quad & \leq 2 \max_{v \in \text{conf}_{\alpha_{\text{vf}}}^{\text{f}}(\Delta^*, \Theta^*, \pi^*)} \max_{w \in \mathcal{W}} \left| \Phi_{\text{vf}}^{\pi^*}(v, w; \Delta^*, \Theta^*) \right| \\
& \quad + \max_{v \in \mathcal{V}} \left| \Phi_{\text{vf}}^{\pi^*}(v, w^{\pi^*}; \Delta^*, \Theta^*) - \Phi_{\text{vf}}^{\pi^*}(v, \tilde{w}^{\pi^*}; \Delta^*, \Theta^*) \right|,
\end{aligned}$$

where in the first inequality, we use triangle inequality; in the second inequality, we use the definition of \tilde{w}^{π^*} that $\tilde{w}^{\pi^*} \in \mathcal{W}$. By Lemma C.5.4 and Assumption 4.4.15, we obtain from (C.5.11) that

$$\begin{aligned}
& \max_{(w,v) \in \text{conf}_{\alpha_{\text{mis}}, \alpha_{\text{vf}}}(\Delta^*, \Theta^*, \pi^*)} \left| \Phi_{\text{vf}}^{\pi^*}(v, w^{\pi^*}; \Delta^*, \Theta^*) - \Phi_{\text{vf}}^{\pi^*}(v, w; \Delta^*, \Theta^*) \right| \\
\text{(C.5.12)} \quad & \leq c \cdot \frac{C_{\Delta^*}^2 C_{\Theta^*}^2 C_*}{1 - \gamma} (\xi_0 + \xi_1) \sqrt{\frac{1}{NT\kappa} \cdot \mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \nu, \Pi} \cdot \log \frac{1}{\delta} \log(NT)} \\
& \quad + 2 \max_{g \in \mathcal{W}} \Phi_{\text{vf}}^{\pi^*}(\tilde{v}^{\pi^*}, g; \Delta^*, \Theta^*) + C_{\Delta^*} C_{\Theta^*} \varepsilon_{\text{vf}}^{\mathcal{W}} / (1 - \gamma).
\end{aligned}$$

Also, we have

$$\begin{aligned}
& \max_{g \in \mathcal{W}} \Phi_{\text{vf}}^{\pi^*}(\tilde{v}^{\pi^*}, g; \Delta^*, \Theta^*) \\
&= \max_{g \in \mathcal{W}} \Phi_{\text{vf}}^{\pi^*}(\tilde{v}^{\pi^*}, g; \Delta^*, \Theta^*) - \max_{g \in \mathcal{W}} \Phi_{\text{vf}}^{\pi^*}(V^{\pi^*}, g; \Delta^*, \Theta^*) + \max_{g \in \mathcal{W}} \Phi_{\text{vf}}^{\pi^*}(V^{\pi^*}, g; \Delta^*, \Theta^*) \\
&= \max_{g \in \mathcal{W}} \Phi_{\text{vf}}^{\pi^*}(\tilde{v}^{\pi^*}, g; \Delta^*, \Theta^*) - \max_{g \in \mathcal{W}} \Phi_{\text{vf}}^{\pi^*}(V^{\pi^*}, g; \Delta^*, \Theta^*) \\
&= \max_{g \in \mathcal{W}} \left| \Phi_{\text{vf}}^{\pi^*}(\tilde{v}^{\pi^*}, g; \Delta^*, \Theta^*) - \Phi_{\text{vf}}^{\pi^*}(V^{\pi^*}, g; \Delta^*, \Theta^*) \right| \\
\text{(C.5.13)} \quad & \leq C_* C_{\Delta^*} C_{\Theta^*} \varepsilon_{\text{vf}}^{\mathcal{V}},
\end{aligned}$$

where we use Assumption 4.4.15 in the last inequality. By plugging (C.5.12) and (C.5.13) into (C.5.10), we have

$$\begin{aligned}
\text{(C.5.14)} \quad J(\pi^*) - J(\hat{\pi}_{\text{dr}}) &\leq c \cdot \frac{C_{\Delta^*}^2 C_{\Theta^*}^2 C_*}{1 - \gamma} (\xi_0 + \xi_1) \sqrt{\frac{1}{NT\kappa} \mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}, \Pi} \log \frac{NT}{\delta}} \\
&\quad + 3C_{\Delta^*} C_{\Theta^*} (C_* \varepsilon_{\text{vf}}^{\mathcal{V}} + \varepsilon_{\text{vf}}^{\mathcal{W}} / (1 - \gamma)),
\end{aligned}$$

where we plug in the definition of ε_L and ε_L^* in the last inequality.

Part (ii). We first introduce the following lemmas.

Lemma C.5.5. Suppose α_{mis} is defined in Lemma 4.4.11 and $c/(NT)^2 \leq \delta \leq 1$.

Then under Assumptions (b) and 4.4.3, with probability at least $1 - \delta$, it holds for any

$\pi \in \Pi$ that $\tilde{w}^{\pi} \in \text{conf}_{\alpha_{\text{mis}}}^{\text{mis}}(\Delta^*, \Theta^*, \pi)$.

Proof. See §C.5.7 for a detailed proof. □

Lemma C.5.6. Suppose that $(\alpha_0, \alpha_1, \alpha_{\text{mis}})$ is defined in Assumption 4.4.4 and Lemma 4.4.11, and $c/(NT)^2 \leq \delta \leq 1$. Then under Assumptions 4.2.2, 4.2.6, 4.4.3, and 4.4.4, with probability at least $1-\delta$, it holds for any $\pi \in \Pi$ and $w \in \cup_{(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1} \text{conf}_{\alpha_{\text{mis}}}^{\text{mis}}(\Delta, \Theta, \pi)$ that

$$\begin{aligned} \max_{f \in \mathcal{V}} \Phi_{\text{mis}}^\pi(w, f; \Delta^*, \Theta^*) &\leq c \cdot \frac{C_{\Delta^*}^2 C_{\Theta^*}^2 C_*}{1-\gamma} (\xi_0 + \xi_1) \sqrt{\frac{1}{NT\kappa} \cdot \mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}, \Pi} \cdot \log \frac{1}{\delta} \log(NT)} \\ &\quad + \max_{f \in \mathcal{V}} \Phi_{\text{mis}}^\pi(\tilde{w}^\pi, f; \Delta^*, \Theta^*). \end{aligned}$$

Proof. See §C.5.8 for a detailed proof. \square

By the definition of L_{dr} , we have

$$\begin{aligned} J(\pi^*) - J(\hat{\pi}_{\text{dr}}) &= J(\pi^*) - L_{\text{dr}}(w^{\hat{\pi}_{\text{dr}}}, v, \hat{\pi}_{\text{dr}}; \Delta^*, \Theta^*) \\ \text{(C.5.15)} \quad &= J(\pi^*) - L_{\text{dr}}(\tilde{w}^{\hat{\pi}_{\text{dr}}}, v, \hat{\pi}_{\text{dr}}; \Delta^*, \Theta^*) + L_{\text{dr}}(\tilde{w}^{\hat{\pi}_{\text{dr}}}, v, \hat{\pi}_{\text{dr}}; \Delta^*, \Theta^*) - L_{\text{dr}}(w^{\hat{\pi}_{\text{dr}}}, v, \hat{\pi}_{\text{dr}}; \Delta^*, \Theta^*). \end{aligned}$$

By Assumption 4.4.15, we have

$$\text{(C.5.16)} \quad \left| L_{\text{dr}}(\tilde{w}^{\hat{\pi}_{\text{dr}}}, v, \hat{\pi}_{\text{dr}}; \Delta^*, \Theta^*) - L_{\text{dr}}(w^{\hat{\pi}_{\text{dr}}}, v, \hat{\pi}_{\text{dr}}; \Delta^*, \Theta^*) \right| \leq C_{\Delta^*} C_{\Theta^*} \varepsilon_{\text{mis}}^{\mathcal{W}} / (1-\gamma).$$

In the meanwhile, by Assumption 4.4.4 and Lemma C.5.5, it holds with probability at least $1 - \delta$ that

$$\begin{aligned}
L_{\text{dr}}(\tilde{w}^{\hat{\pi}_{\text{dr}}}, v, \hat{\pi}_{\text{dr}}; \Delta^*, \Theta^*) &\geq \min_{(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1} \min_{(w, v) \in \text{conf}_{\alpha_{\text{mis}}, \alpha_{\text{vf}}}(\Delta, \Theta, \hat{\pi}_{\text{dr}})} L_{\text{dr}}(w, v, \hat{\pi}_{\text{dr}}; \Delta, \Theta) \\
&\geq \min_{(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1} \min_{(w, v) \in \text{conf}_{\alpha_{\text{mis}}, \alpha_{\text{vf}}}(\Delta, \Theta, \hat{\pi}_{\text{dr}})} \hat{L}_{\text{dr}}(w, v, \hat{\pi}_{\text{dr}}; \Delta, \Theta) - \hat{\epsilon}_L \\
&\geq \min_{(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1} \min_{(w, v) \in \text{conf}_{\alpha_{\text{mis}}, \alpha_{\text{vf}}}(\Delta, \Theta, \pi^*)} \hat{L}_{\text{dr}}(w, v, \pi^*; \Delta, \Theta) - \hat{\epsilon}_L \\
\text{(C.5.17)} \quad &\geq \min_{(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1} \min_{(w, v) \in \text{conf}_{\alpha_{\text{mis}}, \alpha_{\text{vf}}}(\Delta, \Theta, \pi^*)} L_{\text{dr}}(w, v, \pi^*; \Delta, \Theta) - 2\hat{\epsilon}_L,
\end{aligned}$$

where in the second inequality, we use Lemma C.5.1; in the third inequality, we use the optimality of $\hat{\pi}_{\text{dr}}$; in the fourth inequality, we again use Lemma C.5.1. Further, combining Lemma C.5.2 and (C.5.17), we have

$$\text{(C.5.18)} \quad L_{\text{dr}}(\tilde{w}^{\hat{\pi}_{\text{dr}}}, v, \hat{\pi}_{\text{dr}}; \Delta^*, \Theta^*) \geq \min_{(w, v) \in \text{conf}_{\alpha_{\text{mis}}, \alpha_{\text{vf}}}(\Delta^*, \Theta^*, \pi^*)} L_{\text{dr}}(w, v, \pi^*; \Delta^*, \Theta^*) - 2\hat{\epsilon}_L - \epsilon_L^*.$$

Now, by plugging (C.5.16) and (C.5.18) into (C.5.15), we have

$$\begin{aligned}
& J(\pi^*) - J(\widehat{\pi}_{\text{dr}}) \\
& \leq J(\pi^*) - \min_{(w,v) \in \text{conf}_{\alpha_{\text{mis}}, \alpha_{\text{vf}}}(\Delta^*, \Theta^*, \pi^*)} L_{\text{dr}}(w, v, \pi^*; \Delta^*, \Theta^*) \\
& \quad + 2\widehat{\epsilon}_L + \epsilon_L^* + C_{\Delta^*} C_{\Theta^*} \varepsilon_{\text{mis}}^{\mathcal{W}} / (1 - \gamma) \\
& \leq \max_{(w,v) \in \text{conf}_{\alpha_{\text{mis}}, \alpha_{\text{vf}}}(\Delta^*, \Theta^*, \pi^*)} \left| L_{\text{dr}}(w^{\pi^*}, v, \pi^*; \Delta^*, \Theta^*) - L_{\text{dr}}(w, v, \pi^*; \Delta^*, \Theta^*) \right| \\
& \quad + 2\widehat{\epsilon}_L + \epsilon_L^* + C_{\Delta^*} C_{\Theta^*} \varepsilon_{\text{mis}}^{\mathcal{W}} / (1 - \gamma) \\
\text{(C.5.19)} \quad & = \max_{(w,v) \in \text{conf}_{\alpha_{\text{mis}}, \alpha_{\text{vf}}}(\Delta^*, \Theta^*, \pi^*)} \left| \Phi_{\text{mis}}^{\pi^*}(w, V^{\pi^*}; \Delta^*, \Theta^*) - \Phi_{\text{mis}}^{\pi^*}(w, v; \Delta^*, \Theta^*) \right| \\
& \quad + 2\widehat{\epsilon}_L + \epsilon_L^* + C_{\Delta^*} C_{\Theta^*} \varepsilon_{\text{mis}}^{\mathcal{W}} / (1 - \gamma),
\end{aligned}$$

where in the second inequality, we use the fact that $J(\pi^*) = L_{\text{dr}}(w^{\pi^*}, v, \pi^*; \Delta^*, \Theta^*)$; in the last equality, we use the following fact

$$L_{\text{dr}}(w, v, \pi^*; \Delta^*, \Theta^*) - L_{\text{dr}}(w^{\pi^*}, v, \pi^*; \Delta^*, \Theta^*) = -\Phi_{\text{mis}}^{\pi^*}(w, V^{\pi^*}; \Delta^*, \Theta^*) + \Phi_{\text{mis}}^{\pi^*}(w, v; \Delta^*, \Theta^*)$$

for any $(w, v) \in \mathcal{W} \times \mathcal{V}$. We upper bound the first term on the RHS of (C.5.19) as follows,

$$\begin{aligned}
& \max_{(w,v) \in \text{conf}_{\alpha_{\text{mis}}, \alpha_{\text{vf}}}(\Delta^*, \Theta^*, \pi^*)} \left| \Phi_{\text{mis}}^{\pi^*}(w, V^{\pi^*}; \Delta^*, \Theta^*) - \Phi_{\text{mis}}^{\pi^*}(w, v; \Delta^*, \Theta^*) \right| \\
& \leq \max_{w \in \text{conf}_{\alpha_{\text{mis}}}^{\text{mis}}(\Delta^*, \Theta^*, \pi^*)} \max_{v \in \mathcal{V}} \left| \Phi_{\text{mis}}^{\pi^*}(w, v; \Delta^*, \Theta^*) - \Phi_{\text{mis}}^{\pi^*}(w, \tilde{v}^{\pi^*}; \Delta^*, \Theta^*) \right| \\
& \quad + \max_{w \in \mathcal{W}} \left| \Phi_{\text{mis}}^{\pi^*}(w, V^{\pi^*}; \Delta^*, \Theta^*) - \Phi_{\text{mis}}^{\pi^*}(w, \tilde{v}^{\pi^*}; \Delta^*, \Theta^*) \right| \\
& \leq 2 \max_{w \in \text{conf}_{\alpha_{\text{mis}}}^{\text{mis}}(\Delta^*, \Theta^*, \pi^*)} \max_{v \in \mathcal{V}} \left| \Phi_{\text{mis}}^{\pi^*}(w, v; \Delta^*, \Theta^*) \right| \\
& \quad + \max_{w \in \mathcal{W}} \left| \Phi_{\text{mis}}^{\pi^*}(w, V^{\pi^*}; \Delta^*, \Theta^*) - \Phi_{\text{mis}}^{\pi^*}(w, \tilde{v}^{\pi^*}; \Delta^*, \Theta^*) \right| \\
& \leq c \cdot \frac{C_{\Delta^*}^2 C_{\Theta^*}^2 C_*}{1 - \gamma} (\xi_0 + \xi_1) \sqrt{\frac{1}{NT\kappa} \cdot \mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}, \Pi} \cdot \log \frac{1}{\delta} \log(NT)} \\
& \quad + 2 \max_{f \in \mathcal{V}} \Phi_{\text{mis}}^{\pi^*}(\tilde{w}^{\pi^*}, f; \Delta^*, \Theta^*) + C_* C_{\Delta^*} C_{\Theta^*} \varepsilon_{\text{mis}}^{\mathcal{V}},
\end{aligned} \tag{C.5.20}$$

where in the first inequality, we use triangle inequality; in the second inequality, we use the definition of \tilde{v}^{π^*} that $\tilde{v}^{\pi^*} \in \mathcal{V}$; in the last inequality, we use Lemma C.5.6 and Assumption 4.4.15. In the meanwhile, by Assumption 4.4.15, we have

$$\begin{aligned}
& \max_{f \in \mathcal{V}} \Phi_{\text{mis}}^{\pi^*}(\tilde{w}^{\pi^*}, f; \Delta^*, \Theta^*) \\
& \leq \max_{f \in \mathcal{V}} \Phi_{\text{mis}}^{\pi^*}(\tilde{w}^{\pi^*}, f; \Delta^*, \Theta^*) - \max_{f \in \mathcal{V}} \Phi_{\text{mis}}^{\pi^*}(w^{\pi^*}, f; \Delta^*, \Theta^*) + \max_{f \in \mathcal{V}} \Phi_{\text{mis}}^{\pi^*}(w^{\pi^*}, f; \Delta^*, \Theta^*) \\
& = \max_{f \in \mathcal{V}} \Phi_{\text{mis}}^{\pi^*}(\tilde{w}^{\pi^*}, f; \Delta^*, \Theta^*) - \max_{f \in \mathcal{V}} \Phi_{\text{mis}}^{\pi^*}(w^{\pi^*}, f; \Delta^*, \Theta^*) \\
& \leq \max_{f \in \mathcal{V}} \left| \Phi_{\text{mis}}^{\pi^*}(\tilde{w}^{\pi^*}, f; \Delta^*, \Theta^*) - \Phi_{\text{mis}}^{\pi^*}(w^{\pi^*}, f; \Delta^*, \Theta^*) \right|
\end{aligned} \tag{C.5.21}$$

$$\leq C_{\Delta^*} C_{\Theta^*} \varepsilon_{\text{mis}}^{\mathcal{W}} / (1 - \gamma).$$

Now, by plugging (C.5.20) and (C.5.21) into (C.5.19), we have

$$(C.5.22) \quad J(\pi^*) - J(\widehat{\pi}_{\text{dr}}) \leq c \cdot \frac{C_{\Delta^*}^2 C_{\Theta^*}^2 C_*}{1 - \gamma} (\xi_0 + \xi_1) \sqrt{\frac{1}{NT\kappa} \mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}, \Pi} \log \frac{NT}{\delta}} \\ + 3C_{\Delta^*} C_{\Theta^*} (C_* \varepsilon_{\text{mis}}^{\mathcal{V}} + \varepsilon_{\text{mis}}^{\mathcal{W}} / (1 - \gamma)).$$

By combining (C.5.14) and (C.5.22), we have

$$J(\pi^*) - J(\widehat{\pi}_{\text{dr}}) \leq c \cdot \frac{C_{\Delta^*}^2 C_{\Theta^*}^2 C_*}{1 - \gamma} (\xi_0 + \xi_1) \sqrt{\frac{1}{NT\kappa} \mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}, \Pi} \log \frac{NT}{\delta}} \\ + 3C_{\Delta^*} C_{\Theta^*} \min \{ C_* \varepsilon_{\text{vf}}^{\mathcal{V}} + \varepsilon_{\text{vf}}^{\mathcal{W}} / (1 - \gamma), C_* \varepsilon_{\text{mis}}^{\mathcal{V}} + \varepsilon_{\text{mis}}^{\mathcal{W}} / (1 - \gamma) \}$$

which conclude the proof. \square

C.5.3. Proof of Lemma C.5.1

Proof. By Theorem C.7.6, with probability at least $1 - \delta$, it holds for any $(w, v, \Delta, \Theta, \pi) \in \mathcal{W} \times \mathcal{V} \times \mathcal{F}_0 \times \mathcal{F}_1 \times \Pi$ that

$$\left| L_{\text{dr}}(w, v, \pi; \Delta, \Theta) - \widehat{L}_{\text{dr}}(w, v, \pi; \Delta, \Theta) \right| \leq c \cdot \frac{C_{\Delta^*} C_{\Theta^*} C_*}{1 - \gamma} \sqrt{\frac{1}{NT\kappa} \mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}, \Pi} \log \frac{1}{\delta} \log(NT)},$$

which concludes the proof of the lemma. \square

C.5.4. Proof of Lemma C.5.2

Proof. With a slight abuse of notations, we denote by

$$(w_0, v_0) \in \underset{(w,v) \in \text{conf}_{\alpha_{\text{mis}}, \alpha_{\text{vf}}}(\Delta^*, \Theta^*, \pi^*)}{\text{argmin}} L_{\text{dr}}(w, v, \pi^*; \Delta^*, \Theta^*),$$

$$(w_1, v_1) \in \underset{(w,v) \in \text{conf}_{\alpha_{\text{mis}}, \alpha_{\text{vf}}}(\Delta, \Theta, \pi^*)}{\text{argmin}} L_{\text{dr}}(w, v, \pi^*; \Delta, \Theta).$$

Then we have

$$\begin{aligned} & \left| \min_{(w,v) \in \text{conf}_{\alpha_{\text{mis}}, \alpha_{\text{vf}}}(\Delta^*, \Theta^*, \pi^*)} L_{\text{dr}}(w, v, \pi^*; \Delta^*, \Theta^*) - \min_{(w,v) \in \text{conf}_{\alpha_{\text{mis}}, \alpha_{\text{vf}}}(\Delta, \Theta, \pi^*)} L_{\text{dr}}(w, v, \pi^*; \Delta, \Theta) \right| \\ &= |L_{\text{dr}}(w_0, v_0, \pi^*; \Delta^*, \Theta^*) - L_{\text{dr}}(w_1, v_1, \pi^*; \Delta, \Theta)| \\ &\leq |L_{\text{dr}}(w_0, v_0, \pi^*; \Delta^*, \Theta^*) - L_{\text{dr}}(w^\pi, v_0, \pi^*; \Delta^*, \Theta^*)| \\ &\quad + |L_{\text{dr}}(w^\pi, v_1, \pi^*; \Delta^*, \Theta^*) - L_{\text{dr}}(w_1, v_1, \pi^*; \Delta^*, \Theta^*)| \\ &\quad + |L_{\text{dr}}(w_1, v_1, \pi; \Delta^*, \Theta^*) - L_{\text{dr}}(w_1, v_1, \pi; \Delta, \Theta)| \end{aligned}$$

(C.5.23)

$$\begin{aligned} &= \underbrace{\left| \Phi_{\text{mis}}^{\pi^*}(w_0, V^{\pi^*}; \Delta^*, \Theta^*) - \Phi_{\text{mis}}^{\pi^*}(w_0, v_0; \Delta^*, \Theta^*) \right|}_{\text{Term (I)}} \\ &\quad + \underbrace{\left| \Phi_{\text{mis}}^{\pi^*}(w_1, V^{\pi^*}; \Delta^*, \Theta^*) - \Phi_{\text{mis}}^{\pi^*}(w_1, v_1; \Delta^*, \Theta^*) \right|}_{\text{Term (II)}} \\ &\quad + \underbrace{\left| L_{\text{dr}}(w_1, v_1, \pi^*; \Delta^*, \Theta^*) - L_{\text{dr}}(w_1, v_1, \pi^*; \Delta, \Theta) \right|}_{\text{Term (III)}}, \end{aligned}$$

where in the first inequality, we use triangle inequality and the fact that $L_{\text{dr}}(w^{\pi^*}, v, \pi^*; \Delta^*, \Theta^*) = J(\pi^*)$ for any function v ; while in the last equality, we use the following equality for any

(w, v) ,

$$L_{\text{dr}}(w, v, \pi^*; \Delta^*, \Theta^*) - L_{\text{dr}}(w^{\pi^*}, v, \pi^*; \Delta^*, \Theta^*) = -\Phi_{\text{mis}}^{\pi^*}(w, V^{\pi^*}; \Delta^*, \Theta^*) + \Phi_{\text{mis}}^{\pi^*}(w, v; \Delta^*, \Theta^*).$$

We upper bound terms (I), (II), and (III) on the RHS of (C.5.23), respectively.

Upper Bounding Term (I). Note that with probability at least $1 - \delta$, we have

$$\begin{aligned} |\Phi_{\text{mis}}^{\pi^*}(w_0, V^{\pi^*}; \Delta^*, \Theta^*)| &\leq \max_{f \in \mathcal{V}} |\Phi_{\text{mis}}^{\pi^*}(w_0, f; \Delta^*, \Theta^*)| \\ &= \max_{f \in \mathcal{V}} \max \{ \Phi_{\text{mis}}^{\pi^*}(w_0, f; \Delta^*, \Theta^*), -\Phi_{\text{mis}}^{\pi^*}(w_0, f; \Delta^*, \Theta^*) \} \\ &= \max_{f \in \mathcal{V}} \max \{ \Phi_{\text{mis}}^{\pi^*}(w_0, f; \Delta^*, \Theta^*), \Phi_{\text{mis}}^{\pi^*}(w_0, -f; \Delta^*, \Theta^*) \} \\ &= \max_{f \in \mathcal{V}} \Phi_{\text{mis}}^{\pi^*}(w_0, f; \Delta^*, \Theta^*) \\ &\leq c \cdot \frac{C_{\Delta^*}^2 C_{\Theta^*}^2 C_*}{1 - \gamma} (\xi_0 + \xi_1) \sqrt{\frac{1}{NT\kappa} \mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}} \log \frac{1}{\delta} \log(NT)}, \end{aligned}$$

where in the first inequality, we use the fact that $V^{\pi^*} \in \mathcal{V}$; in the third equality, we use the fact that \mathcal{V} is symmetric; in the last inequality, by noting that $w_0 \in \cup_{(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1} \text{conf}_{\alpha_{\text{mis}}}^{\pi}(\Delta, \Theta)$, we use Lemma 4.4.12. Similarly, we have

$$|\Phi_{\text{mis}}^{\pi^*}(w_0, v_0; \Delta^*, \Theta^*)| \leq c \cdot \frac{C_{\Delta^*}^2 C_{\Theta^*}^2 C_*}{1 - \gamma} (\xi_0 + \xi_1) \sqrt{\frac{1}{NT\kappa} \mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}} \log \frac{1}{\delta} \log(NT)},$$

which implies that

$$(C.5.24) \quad \text{Term (I)} \leq c \cdot \frac{C_{\Delta^*}^2 C_{\Theta^*}^2 C_*}{1 - \gamma} (\xi_0 + \xi_1) \sqrt{\frac{1}{NT\kappa} \mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}} \log \frac{1}{\delta} \log(NT)}.$$

Upper Bounding Term (II). Similar to (C.5.24), with probability at least $1 - \delta$, we have

$$(C.5.25) \quad \text{Term (II)} \leq c \cdot \frac{C_{\Delta^*}^2 C_{\Theta^*}^2 C_*}{1 - \gamma} (\xi_0 + \xi_1) \sqrt{\frac{1}{NT\kappa} \mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}} \log \frac{1}{\delta} \log(NT)}.$$

Upper Bounding Term (III). Note that with probability at least $1 - \delta$, it holds for any $(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1$ that

$$(C.5.26) \quad \begin{aligned} & \text{Term (III)} \\ &= \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} \left(\frac{Z_t^\top A_t \pi^*(A_t | S_t)}{\Delta^*(S_t, A_t) \Theta^*(S_t, Z_t)} - \frac{Z_t^\top A_t \pi^*(A_t | S_t)}{\Delta(S_t, A_t) \Theta(S_t, Z_t)} \right) w_1(S_t) (R_t + \gamma v_1(S_{t+1}) - v_1(S_t)) \right] \\ &\leq \frac{C_{\Delta^*} C_{\Theta^*} C_*}{1 - \gamma} \left(\xi_0 C_{\Theta^*} \sqrt{\frac{C_{\Delta^*}}{NT\kappa} \mathfrak{C}_{\mathcal{F}_0} \log \frac{1}{\delta} \log(NT)} + \xi_1 C_{\Delta^*} \sqrt{\frac{C_{\Theta^*}}{NT\kappa} \mathfrak{C}_{\mathcal{F}_1} \log \frac{1}{\delta} \log(NT)} \right), \end{aligned}$$

where we use triangle inequality and Assumption 4.4.4 in the last inequality.

Now, by plugging (C.5.24), (C.5.25), and (C.5.26) into (C.5.23), we conclude the proof of the lemma. \square

C.5.5. Proof of Lemma C.5.3

Proof. By the definition of \tilde{v}^π in (4.4.2), we know that $\tilde{v}^\pi \in \mathcal{V}$. Thus, to show that $\tilde{v}^\pi \in \text{conf}_{\alpha_{\text{vf}}}^{\text{vf}}(\Delta^*, \Theta^*, \pi)$ with a high probability, it suffices to show that

$$(C.5.27) \quad \max_{g \in \mathcal{W}} \widehat{\Phi}_{\text{vf}}^\pi(\tilde{v}^\pi, g; \Delta^*, \Theta^*) - \max_{g \in \mathcal{W}} \widehat{\Phi}_{\text{vf}}^\pi(\tilde{v}_{\Delta^*, \Theta^*}^\pi, g; \Delta^*, \Theta^*) \leq \alpha_{\text{vf}}.$$

In the follows, we show that (C.5.27) holds with a high probability. For the simplicity of notations, we denote by $\Phi_{\text{vf}}^\pi(v, g; *) = \Phi_{\text{vf}}^\pi(v, g; \Delta^*, \Theta^*)$ and $\widehat{v}_*^\pi = \widehat{v}_{\Delta^*, \Theta^*}^\pi$ for any (π, v, g) . Note that

$$\begin{aligned}
& \max_{g \in \mathcal{W}} \widehat{\Phi}_{\text{vf}}^\pi(\widetilde{v}^\pi, g; *) - \max_{g \in \mathcal{W}} \widehat{\Phi}_{\text{vf}}^\pi(\widehat{v}_*^\pi, g; *) \\
&= \max_{g \in \mathcal{W}} \widehat{\Phi}_{\text{vf}}^\pi(\widetilde{v}^\pi, g; *) - \max_{g \in \mathcal{W}} \Phi_{\text{vf}}^\pi(\widetilde{v}^\pi, g; *) + \max_{g \in \mathcal{W}} \Phi_{\text{vf}}^\pi(\widetilde{v}^\pi, g; *) - \max_{g \in \mathcal{W}} \Phi_{\text{vf}}^\pi(\widehat{v}_*^\pi, g; *) \\
&\quad + \max_{g \in \mathcal{W}} \Phi_{\text{vf}}^\pi(\widehat{v}_*^\pi, g; *) - \max_{g \in \mathcal{W}} \widehat{\Phi}_{\text{vf}}^\pi(\widehat{v}_*^\pi, g; *) \\
&\leq \max_{g \in \mathcal{W}} \widehat{\Phi}_{\text{vf}}^\pi(\widetilde{v}^\pi, g; *) - \max_{g \in \mathcal{W}} \Phi_{\text{vf}}^\pi(\widetilde{v}^\pi, g; *) + \max_{g \in \mathcal{W}} \Phi_{\text{vf}}^\pi(\widehat{v}_*^\pi, g; *) - \max_{g \in \mathcal{W}} \widehat{\Phi}_{\text{vf}}^\pi(\widehat{v}_*^\pi, g; *) \\
&\leq 2 \max_{v \in \mathcal{V}} \left| \max_{g \in \mathcal{W}} \widehat{\Phi}_{\text{vf}}^\pi(v, g; *) - \max_{g \in \mathcal{W}} \Phi_{\text{vf}}^\pi(v, g; *) \right| \\
\text{(C.5.28)} \quad &\leq 2 \max_{v \in \mathcal{V}} \max_{g \in \mathcal{W}} \left| \widehat{\Phi}_{\text{vf}}^\pi(v, g; *) - \Phi_{\text{vf}}^\pi(v, g; *) \right|,
\end{aligned}$$

where in the first inequality, we use the fact that $\max_{g \in \mathcal{W}} \Phi_{\text{vf}}^\pi(\widetilde{v}^\pi, g; *) \leq \max_{g \in \mathcal{W}} \Phi_{\text{vf}}^\pi(v, g; *)$ for any v due to the definition of \widetilde{v}^π in (4.4.2); while in the second inequality, we use the fact that $\widetilde{v}^\pi, \widehat{v}_*^\pi \in \mathcal{V}$. In the meanwhile, by Theorem C.7.6, with probability at least $1 - \delta$, it holds for any $(\pi, v, g) \in \Pi \times \mathcal{V} \times \mathcal{W}$ that

$$\text{(C.5.29)} \quad \left| \widehat{\Phi}_{\text{vf}}^\pi(v, g; *) - \Phi_{\text{vf}}^\pi(v, g; *) \right| \leq c \cdot \frac{C_{\Delta^*} C_{\Theta^*} C_*}{1 - \gamma} \sqrt{\frac{\mathfrak{C}_{\mathcal{W}, \mathcal{V}, \Pi}}{NT\kappa} \cdot \log \frac{1}{\delta} \log(NT)},$$

where we use Assumption (b) and $\|g\|_\infty \leq C_*$ for any $g \in \mathcal{W}$. Now, combining (C.5.28) and (C.5.29), with probability at least $1 - \delta$, we have

$$\max_{g \in \mathcal{W}} \widehat{\Phi}_{\text{vf}}^\pi(\widetilde{v}^\pi, g; *) - \max_{g \in \mathcal{W}} \widehat{\Phi}_{\text{vf}}^\pi(\widehat{v}_*^\pi, g; *) \leq c \cdot \frac{C_{\Delta^*} C_{\Theta^*} C_*}{1 - \gamma} \sqrt{\frac{\mathfrak{C}_{\mathcal{W}, \mathcal{V}, \Pi}}{NT\kappa} \cdot \log \frac{1}{\delta} \log(NT)} = \alpha_{\text{vf}},$$

which implies that $\tilde{v}^\pi \in \text{conf}_{\alpha_{\text{vf}}}^{\text{vf}}(\Delta^*, \Theta^*, \pi)$ for any $\pi \in \Pi$. This concludes the proof of the lemma. \square

C.5.6. Proof of Lemma C.5.4

Proof. Since $v \in \cup_{(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1} \text{conf}_{\alpha_{\text{vf}}}^{\text{vf}}(\Delta, \Theta, \pi)$, there exists a pair $(\tilde{\Delta}, \tilde{\Theta}) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1$ such that $v \in \text{conf}_{\alpha_{\text{vf}}}^{\text{vf}}(\tilde{\Delta}, \tilde{\Theta}, \pi)$. For the simplicity of notations, we denote by

$$v^\dagger \in \underset{v \in \mathcal{V}}{\text{argmin}} \max_{g \in \mathcal{W}} \hat{\Phi}_{\text{vf}}^\pi(v, g; \tilde{\Delta}, \tilde{\Theta}),$$

i.e., $v^\dagger = \hat{v}_{\tilde{\Delta}, \tilde{\Theta}}^\pi$, which is defined in (4.3.2). By the definition of v^\dagger and $v \in \text{conf}_{\alpha_{\text{vf}}}^{\text{vf}}(\tilde{\Delta}, \tilde{\Theta}, \pi)$, we know that

$$(C.5.30) \quad \max_{g \in \mathcal{W}} \hat{\Phi}_{\text{vf}}^\pi(v, g; \tilde{\Delta}, \tilde{\Theta}) - \max_{g \in \mathcal{W}} \hat{\Phi}_{\text{vf}}^\pi(v^\dagger, g; \tilde{\Delta}, \tilde{\Theta}) \leq \alpha_{\text{vf}}.$$

Note that

$$(C.5.31) \quad \begin{aligned} & \max_{g \in \mathcal{W}} \Phi_{\text{vf}}^\pi(v, g; \Delta^*, \Theta^*) \\ & \leq 2 \underbrace{\max_{(v, g, \Delta, \Theta) \in (\mathcal{V}, \mathcal{W}, \mathcal{F}_0, \mathcal{F}_1)} \left| \Phi_{\text{vf}}^\pi(v, g; \Delta, \Theta) - \hat{\Phi}_{\text{vf}}^\pi(v, g; \Delta, \Theta) \right|}_{\text{Term (I)}} + \underbrace{\max_{g \in \mathcal{W}} \Phi_{\text{vf}}^\pi(v^\dagger, g; \tilde{\Delta}, \tilde{\Theta})}_{\text{Term (II)}} \\ & \quad + \underbrace{\max_{g \in \mathcal{W}} \left| \hat{\Phi}_{\text{vf}}^\pi(v, g; \Delta^*, \Theta^*) - \hat{\Phi}_{\text{vf}}^\pi(v, g; \tilde{\Delta}, \tilde{\Theta}) \right|}_{\text{Term (III)}} + \alpha_{\text{vf}}, \end{aligned}$$

where we use (C.5.30) in the last inequality. Now we upper bound terms (I), (II), and (III) on the RHS of (C.5.31).

Upper Bounding Term (I). By Theorem C.7.6, with probability at least $1 - \delta$, it holds for any $(v, g, \Delta, \Theta, \pi) \in (\mathcal{V}, \mathcal{W}, \mathcal{F}_0, \mathcal{F}_1, \Pi)$ that

$$\left| \widehat{\Phi}_{\text{vf}}^\pi(v, g; \Delta, \Theta) - \Phi_{\text{vf}}^\pi(v, g; \Delta, \Theta) \right| \leq c \cdot \frac{C_{\Delta^*} C_{\Theta^*} C_*}{1 - \gamma} \sqrt{\frac{\mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}, \Pi}}{NT\kappa} \log \frac{1}{\delta} \log(NT)},$$

which implies that with probability at least $1 - \delta$, we have

$$(C.5.32) \quad \text{Term (I)} \leq c \cdot \frac{C_{\Delta^*} C_{\Theta^*} C_*}{1 - \gamma} \sqrt{\frac{\mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}, \Pi}}{NT\kappa} \log \frac{1}{\delta} \log(NT)}.$$

Upper Bounding Term (II). Recall that $\tilde{v}^\pi \in \operatorname{argmin}_{v \in \mathcal{V}} \max_{w \in \mathcal{W}} \Phi_{\text{vf}}^\pi(v, w; \Delta^*, \Theta^*)$.

Note that

$$\begin{aligned} & \max_{g \in \mathcal{W}} \Phi_{\text{vf}}^\pi(v^\dagger, g; \tilde{\Delta}, \tilde{\Theta}) \\ & \leq 2 \max_{v \in \mathcal{V}} \max_{g \in \mathcal{W}} \left| \Phi_{\text{vf}}^\pi(v, g; \tilde{\Delta}, \tilde{\Theta}) - \widehat{\Phi}_{\text{vf}}^\pi(v, g; \tilde{\Delta}, \tilde{\Theta}) \right| \\ & \quad + \max_{g \in \mathcal{W}} \widehat{\Phi}_{\text{vf}}^\pi(v^\dagger, g; \tilde{\Delta}, \tilde{\Theta}) - \max_{g \in \mathcal{W}} \widehat{\Phi}_{\text{vf}}^\pi(\tilde{v}^\pi, g; \tilde{\Delta}, \tilde{\Theta}) + \max_{g \in \mathcal{W}} \Phi_{\text{vf}}^\pi(\tilde{v}^\pi, g; \tilde{\Delta}, \tilde{\Theta}) \\ (C.5.33) \quad & \leq 2 \max_{v \in \mathcal{V}} \max_{g \in \mathcal{W}} \left| \Phi_{\text{vf}}^\pi(v, g; \tilde{\Delta}, \tilde{\Theta}) - \widehat{\Phi}_{\text{vf}}^\pi(v, g; \tilde{\Delta}, \tilde{\Theta}) \right| \\ & \quad + \max_{g \in \mathcal{W}} \left| \Phi_{\text{vf}}^\pi(\tilde{v}^\pi, g; \tilde{\Delta}, \tilde{\Theta}) - \Phi_{\text{vf}}^\pi(\tilde{v}^\pi, g; \Delta^*, \Theta^*) \right| + \max_{g \in \mathcal{W}} \Phi_{\text{vf}}^\pi(\tilde{v}^\pi, g; \Delta^*, \Theta^*), \end{aligned}$$

where we use triangle inequality and the fact that $v^\dagger \in \operatorname{argmin}_{v \in \mathcal{V}} \max_{g \in \mathcal{W}} \widehat{\Phi}_{\text{vf}}^\pi(v, g; \Delta, \Theta)$ in the last inequality. In the meanwhile, by Theorem C.7.6, with probability at least $1 - \delta$,

it holds for any $(v, g, \pi) \in \mathcal{V} \times \mathcal{W} \times \Pi$ that

$$(C.5.34) \quad \left| \Phi_{\text{vf}}^\pi(v, g; \tilde{\Delta}, \tilde{\Theta}) - \hat{\Phi}_{\text{vf}}^\pi(v, g; \tilde{\Delta}, \tilde{\Theta}) \right| \leq c \cdot \frac{C_{\Delta^*} C_{\Theta^*} C_*}{1 - \gamma} \sqrt{\frac{\mathfrak{C}_{\mathcal{V}, \mathcal{W}, \Pi}}{NT\kappa} \log \frac{1}{\delta} \log(NT)}.$$

Also, we upper bound the second term on the RHS of (C.5.33) with probability at least $1 - \delta$ by a similar argument as in (C.3.14),

$$(C.5.35) \quad \left| \Phi_{\text{vf}}^\pi(\tilde{v}^\pi, g; \tilde{\Delta}, \tilde{\Theta}) - \Phi_{\text{vf}}^\pi(\tilde{v}^\pi, g; \Delta^*, \Theta^*) \right| \\ \leq \frac{C_{\Delta^*} C_{\Theta^*} C_*}{1 - \gamma} \left(\xi_0 C_{\Theta^*} \sqrt{\frac{C_{\Delta^*}}{NT\kappa} \mathfrak{C}_{\mathcal{F}_0} \log \frac{1}{\delta} \log(NT)} + \xi_1 C_{\Delta^*} \sqrt{\frac{C_{\Theta^*}}{NT\kappa} \mathfrak{C}_{\mathcal{F}_1} \log \frac{1}{\delta} \log(NT)} \right),$$

where in the first inequality, we use the fact that $\|v\|_\infty \leq 1/(1 - \gamma)$ and $\|g\|_\infty \leq C_*$; in the third inequality, we use Cauchy Schwarz inequality; while in the last inequality, we use Assumption 4.4.4 with $(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1$. Now, by plugging (C.5.34) and (C.5.35) into (C.5.33), with probability at least $1 - \delta$, it holds for any $(\Delta, \Theta, \pi) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1 \times \Pi$ that

$$(C.5.36) \quad \text{Term (II)} \leq c \cdot \frac{C_{\Delta^*}^2 C_{\Theta^*}^2 C_*}{1 - \gamma} (\xi_0 + \xi_1) \sqrt{\frac{\mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}, \Pi}}{NT\kappa} \log \frac{1}{\delta} \log(NT)} + \max_{g \in \mathcal{W}} \Phi_{\text{vf}}^\pi(\tilde{v}^\pi, g; \Delta^*, \Theta^*).$$

Upper Bounding Term (III). Note that

$$\begin{aligned}
& \left| \widehat{\Phi}_{\text{vf}}^\pi(v, g; \Delta^*, \Theta^*) - \widehat{\Phi}_{\text{vf}}^\pi(v, g; \widetilde{\Delta}, \widetilde{\Theta}) \right| \\
& \leq \left| \left(\widehat{\mathbb{E}} - \mathbb{E} \right) \left[\frac{1}{T} \sum_{t=0}^{T-1} g(S_t) \left(\frac{Z_t^\top A_t \pi(A_t | S_t)}{\Delta^*(S_t, A_t) \Theta^*(S_t, Z_t)} - \frac{Z_t^\top A_t \pi(A_t | S_t)}{\widetilde{\Delta}(S_t, A_t) \widetilde{\Theta}(S_t, Z_t)} \right) (R_t + \gamma v(S_{t+1})) \right] \right| \\
& \quad \text{(C.5.37)} \\
& \quad + \left| \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} g(S_t) \left(\frac{Z_t^\top A_t \pi(A_t | S_t)}{\Delta^*(S_t, A_t) \Theta^*(S_t, Z_t)} - \frac{Z_t^\top A_t \pi(A_t | S_t)}{\widetilde{\Delta}(S_t, A_t) \widetilde{\Theta}(S_t, Z_t)} \right) (R_t + \gamma v(S_{t+1})) \right] \right|.
\end{aligned}$$

For the first term on the RHS of (C.3.12), by Theorem C.7.6, with probability at least $1 - \delta$, it holds for any $(v, g, \pi) \in \mathcal{V} \times \mathcal{W} \times \Pi$ that

$$\begin{aligned}
& \left| \left(\widehat{\mathbb{E}} - \mathbb{E} \right) \left[\frac{1}{T} \sum_{t=0}^{T-1} g(S_t) \left(\frac{Z_t^\top A_t \pi(A_t | S_t)}{\Delta^*(S_t, A_t) \Theta^*(S_t, Z_t)} - \frac{Z_t^\top A_t \pi(A_t | S_t)}{\widetilde{\Delta}(S_t, A_t) \widetilde{\Theta}(S_t, Z_t)} \right) (R_t + \gamma v(S_{t+1})) \right] \right| \\
& \quad \text{(C.5.38)} \\
& \leq c \cdot \frac{C_{\Delta^*} C_{\Theta^*} C_*}{1 - \gamma} \sqrt{\frac{\mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}, \Pi}}{NT\kappa} \log \frac{1}{\delta} \log(NT)}.
\end{aligned}$$

For the second term on the RHS of (C.5.37), by a similar argument as in (C.3.14), with probability at least $1 - \delta$, it holds that

$$\begin{aligned}
& \left| \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} g(S_t) \left(\frac{Z_t^\top A_t \pi(A_t | S_t)}{\Delta^*(S_t, A_t) \Theta^*(S_t, Z_t)} - \frac{Z_t^\top A_t \pi(A_t | S_t)}{\widetilde{\Delta}(S_t, A_t) \widetilde{\Theta}(S_t, Z_t)} \right) (R_t + \gamma v(S_{t+1})) \right] \right| \\
& \quad \text{(C.5.39)} \\
& \leq \frac{C_{\Delta^*} C_{\Theta^*} C_*}{1 - \gamma} \left(\xi_0 C_{\Theta^*} \sqrt{\frac{C_{\Delta^*}}{NT\kappa} \mathfrak{C}_{\mathcal{F}_0} \log \frac{1}{\delta} \log(NT)} + \xi_1 C_{\Delta^*} \sqrt{\frac{C_{\Theta^*}}{NT\kappa} \mathfrak{C}_{\mathcal{F}_1} \log \frac{1}{\delta} \log(NT)} \right),
\end{aligned}$$

where in the first inequality, we use the fact that $\|v\|_\infty \leq 1/(1 - \gamma)$ and $\|g\|_\infty \leq C_*$; in the third inequality, we use Cauchy Schwarz inequality; while in the last inequality, we use

Assumption 4.4.4 with the fact that $(\tilde{\Delta}, \tilde{\Theta}) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1$. Now, by plugging (C.5.38) and (C.5.39) into (C.5.37), it holds with probability at least $1 - \delta$ that

$$(C.5.40) \quad \begin{aligned} & \left| \widehat{\Phi}_{\text{vf}}^\pi(v, g; \Delta^*, \Theta^*) - \widehat{\Phi}_{\text{vf}}^\pi(v, g; \tilde{\Delta}, \tilde{\Theta}) \right| \\ & \leq c \cdot \frac{C_{\Delta^*}^2 C_{\Theta^*}^2 C_*}{1 - \gamma} (\xi_0 + \xi_1) \sqrt{\frac{1}{NT\kappa} \cdot \mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}, \Pi} \cdot \log \frac{1}{\delta} \log(NT)}. \end{aligned}$$

Now, by plugging (C.5.32), (C.5.36), and (C.5.40) into (C.5.31), with probability at least $1 - \delta$, it holds for any $v \in \cup_{(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1} \text{conf}_{\alpha_{\text{vf}}}^{\text{vf}}(\Delta, \Theta, \pi)$ and $\pi \in \Pi$ that

$$\begin{aligned} \max_{g \in \mathcal{W}} \Phi_{\text{vf}}^\pi(v, g; \Delta^*, \Theta^*) & \leq c \cdot \frac{C_{\Delta^*}^2 C_{\Theta^*}^2 C_*}{1 - \gamma} (\xi_0 + \xi_1) \sqrt{\frac{1}{NT\kappa} \cdot \mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}, \Pi} \cdot \log \frac{1}{\delta} \log(NT)} \\ & \quad + \max_{g \in \mathcal{W}} \Phi_{\text{vf}}^\pi(\tilde{v}^\pi, g; \Delta^*, \Theta^*), \end{aligned}$$

which concludes the proof of the lemma. \square

C.5.7. Proof of Lemma C.5.5

Proof. First, by the definition of \tilde{w}^π in (4.4.2), we know that $\tilde{w}^\pi \in \mathcal{W}$. For notation simplicity, we denote by $\Phi_{\text{mis}}^\pi(w, f; *) = \Phi_{\text{mis}}^\pi(w, f; \Delta^*, \Theta^*)$ and $\hat{w}_*^\pi = \hat{w}_{\Delta^*, \Theta^*}^\pi$ for any

(π, w, f) . Note that

$$\begin{aligned}
& \max_{f \in \mathcal{V}} \widehat{\Phi}_{\text{mis}}^\pi(\tilde{w}^\pi, f; *) - \max_{f \in \mathcal{V}} \widehat{\Phi}_{\text{mis}}^\pi(\widehat{w}_*^\pi, f; *) \\
&= \max_{f \in \mathcal{V}} \widehat{\Phi}_{\text{mis}}^\pi(\tilde{w}^\pi, f; *) - \max_{f \in \mathcal{V}} \Phi_{\text{mis}}^\pi(\tilde{w}^\pi, f; *) + \max_{f \in \mathcal{V}} \Phi_{\text{mis}}^\pi(\tilde{w}^\pi, f; *) - \max_{f \in \mathcal{V}} \Phi_{\text{mis}}^\pi(\widehat{w}_*^\pi, f; *) \\
&\quad + \max_{f \in \mathcal{V}} \Phi_{\text{mis}}^\pi(\widehat{w}_*^\pi, f; *) - \max_{f \in \mathcal{V}} \widehat{\Phi}_{\text{mis}}^\pi(\widehat{w}_*^\pi, f; *) \\
&\leq \max_{f \in \mathcal{V}} \widehat{\Phi}_{\text{mis}}^\pi(\tilde{w}^\pi, f; *) - \max_{f \in \mathcal{V}} \Phi_{\text{mis}}^\pi(\tilde{w}^\pi, f; *) + \max_{f \in \mathcal{V}} \Phi_{\text{mis}}^\pi(\widehat{w}_*^\pi, f; *) - \max_{f \in \mathcal{V}} \widehat{\Phi}_{\text{mis}}^\pi(\widehat{w}_*^\pi, f; *) \\
&\leq 2 \max_{w \in \mathcal{W}} \left| \max_{f \in \mathcal{V}} \widehat{\Phi}_{\text{mis}}^\pi(w, f; *) - \max_{f \in \mathcal{V}} \Phi_{\text{mis}}^\pi(w, f; *) \right|
\end{aligned}$$

(C.5.41)

$$\leq 2 \max_{w \in \mathcal{W}} \max_{f \in \mathcal{V}} \left| \widehat{\Phi}_{\text{mis}}^\pi(w, f; *) - \Phi_{\text{mis}}^\pi(w, f; *) \right|,$$

where in the first inequality, we use the fact that $\max_{f \in \mathcal{V}} \Phi_{\text{mis}}^\pi(\tilde{w}^\pi, f; *) \leq \max_{f \in \mathcal{V}} \Phi_{\text{mis}}^\pi(\widehat{w}_*^\pi, f; *)$ by the definition of \tilde{w}^π in (4.4.2); while in the second inequality, we use the fact that $\tilde{w}^\pi, \widehat{w}_*^\pi \in \mathcal{W}$. In the meanwhile, by Theorem C.7.6, with probability at least $1 - \delta$, it holds for any $(w, f, \pi) \in \mathcal{W} \times \mathcal{V} \times \Pi$ that

$$\text{(C.5.42)} \quad \left| \widehat{\Phi}_{\text{mis}}^\pi(w, f; *) - \Phi_{\text{mis}}^\pi(w, f; *) \right| \leq c \cdot \frac{C_{\Delta^*} C_{\Theta^*} C_*}{1 - \gamma} \sqrt{\frac{1}{NT\kappa} \mathfrak{C}_{\mathcal{V}, \mathcal{W}, \Pi} \log \frac{1}{\delta} \log(NT)},$$

where we use Assumption (b). Now, combining (C.5.41) and (C.5.42), with probability at least $1 - \delta$, we have

$$\begin{aligned}
& \max_{f \in \mathcal{V}} \widehat{\Phi}_{\text{mis}}^\pi(\tilde{w}^\pi, f; *) - \max_{f \in \mathcal{V}} \widehat{\Phi}_{\text{mis}}^\pi(\widehat{w}_*^\pi, f; *) \\
&\leq c \cdot \frac{C_{\Delta^*} C_{\Theta^*} C_*}{1 - \gamma} \sqrt{\frac{1}{NT\kappa} \mathfrak{C}_{\mathcal{V}, \mathcal{W}, \Pi} \log \frac{1}{\delta} \log(NT)} = \alpha_{\text{mis}},
\end{aligned}$$

which implies that $\tilde{w}^\pi \in \text{conf}_{\alpha_{\text{mis}}}^{\text{mis}}(\Delta^*, \Theta^*, \pi)$. This concludes the proof of the lemma. \square

C.5.8. Proof of Lemma C.5.6

Proof. Since $w \in \cup_{(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1} \text{conf}_{\alpha_{\text{mis}}}^{\text{mis}}(\Delta, \Theta, \pi)$, there exists a pair $(\tilde{\Delta}, \tilde{\Theta}) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1$ such that $w \in \text{conf}_{\alpha_{\text{mis}}}^{\text{mis}}(\tilde{\Delta}, \tilde{\Theta}, \pi)$. For the simplicity of notations, we denote by

$$w^\dagger \in \underset{w \in \mathcal{W}}{\text{argmin}} \max_{f \in \mathcal{V}} \widehat{\Phi}_{\text{mis}}^\pi(w, f; \tilde{\Delta}, \tilde{\Theta}),$$

i.e., $w^\dagger = \widehat{w}_{\tilde{\Delta}, \tilde{\Theta}}^\pi$, which is defined in (4.3.5). By the definition of w^\dagger and $w \in \text{conf}_{\alpha_{\text{mis}}}^{\text{mis}}(\tilde{\Delta}, \tilde{\Theta}, \pi)$, with probability at least $1 - \delta$, it holds for any $\pi \in \Pi$ and $w \in \text{conf}_{\alpha_{\text{mis}}}^{\text{mis}}(\tilde{\Delta}, \tilde{\Theta}, \pi)$ that

$$(C.5.43) \quad \max_{f \in \mathcal{V}} \widehat{\Phi}_{\text{mis}}^\pi(w, f; \tilde{\Delta}, \tilde{\Theta}) - \max_{f \in \mathcal{V}} \widehat{\Phi}_{\text{mis}}^\pi(w^\dagger, f; \tilde{\Delta}, \tilde{\Theta}) \leq \alpha_{\text{mis}}.$$

Further, we observe that

$$(C.5.44) \quad \begin{aligned} & \max_{f \in \mathcal{V}} \Phi_{\text{mis}}^\pi(w, f; \Delta^*, \Theta^*) \\ & \leq \underbrace{\max_{(w, f, \Delta, \Theta) \in (\mathcal{W}, \mathcal{V}, \mathcal{F}_0, \mathcal{F}_1)} \left| \Phi_{\text{mis}}^\pi(w, f; \Delta, \Theta) - \widehat{\Phi}_{\text{mis}}^\pi(w, f; \Delta, \Theta) \right|}_{\text{Term (I)}} + \underbrace{\max_{f \in \mathcal{V}} \Phi_{\text{mis}}^\pi(w^\dagger, f; \tilde{\Delta}, \tilde{\Theta})}_{\text{Term (II)}} \\ & \quad + \underbrace{\max_{f \in \mathcal{V}} \left| \widehat{\Phi}_{\text{mis}}^\pi(w, f; \Delta^*, \Theta^*) - \widehat{\Phi}_{\text{mis}}^\pi(w, f; \tilde{\Delta}, \tilde{\Theta}) \right|}_{\text{Term (III)}} + \alpha_{\text{mis}}, \end{aligned}$$

where we use (C.5.43) in the last inequality. Now we upper bound terms (I), (II), and (III) on the RHS of (C.5.44).

Upper Bounding Term (I). By Theorem C.7.6, with probability at least $1 - \delta$, it holds for any $(w, f, \Delta, \Theta, \pi) \in (\mathcal{W}, \mathcal{V}, \mathcal{F}_0, \mathcal{F}_1, \Pi)$ that

$$\left| \widehat{\Phi}_{\text{mis}}^\pi(w, f; \Delta, \Theta) - \Phi_{\text{mis}}^\pi(w, f; \Delta, \Theta) \right| \leq c \cdot \frac{C_{\Delta^*} C_{\Theta^*} C_*}{1 - \gamma} \sqrt{\frac{1}{NT\kappa} \mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}, \Pi} \log \frac{1}{\delta} \log(NT)},$$

which implies that with probability at least $1 - \delta$, we have

$$(C.5.45) \quad \text{Term (I)} \leq c \cdot \frac{C_{\Delta^*} C_{\Theta^*} C_*}{1 - \gamma} \sqrt{\frac{1}{NT\kappa} \mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}, \Pi} \log \frac{1}{\delta} \log(NT)}.$$

Upper Bounding Term (II). Recall that $\tilde{w}^\pi \in \operatorname{argmin}_{w \in \mathcal{W}} \max_{v \in \mathcal{V}} \Phi_{\text{mis}}^\pi(w, v; \Delta^*, \Theta^*)$.

Note that

$$\begin{aligned} & \max_{f \in \mathcal{V}} \Phi_{\text{mis}}^\pi(w^\dagger, f; \tilde{\Delta}, \tilde{\Theta}) \\ &= \max_{f \in \mathcal{V}} \Phi_{\text{mis}}^\pi(w^\dagger, f; \tilde{\Delta}, \tilde{\Theta}) - \max_{f \in \mathcal{V}} \widehat{\Phi}_{\text{mis}}^\pi(w^\dagger, f; \tilde{\Delta}, \tilde{\Theta}) + \max_{f \in \mathcal{V}} \widehat{\Phi}_{\text{mis}}^\pi(w^\dagger, f; \tilde{\Delta}, \tilde{\Theta}) \\ & \quad - \max_{f \in \mathcal{V}} \widehat{\Phi}_{\text{mis}}^\pi(\tilde{w}^\pi, f; \tilde{\Delta}, \tilde{\Theta}) + \max_{f \in \mathcal{V}} \widehat{\Phi}_{\text{mis}}^\pi(\tilde{w}^\pi, f; \tilde{\Delta}, \tilde{\Theta}) - \max_{f \in \mathcal{V}} \Phi_{\text{mis}}^\pi(\tilde{w}^\pi, f; \tilde{\Delta}, \tilde{\Theta}) \\ & \quad + \max_{f \in \mathcal{V}} \Phi_{\text{mis}}^\pi(\tilde{w}^\pi, f; \tilde{\Delta}, \tilde{\Theta}) \\ & \leq 2 \max_{w \in \mathcal{W}} \max_{f \in \mathcal{V}} \left| \Phi_{\text{mis}}^\pi(w, f; \tilde{\Delta}, \tilde{\Theta}) - \widehat{\Phi}_{\text{mis}}^\pi(w, f; \tilde{\Delta}, \tilde{\Theta}) \right| \\ (C.5.46) \quad & + \max_{f \in \mathcal{V}} \left| \Phi_{\text{mis}}^\pi(\tilde{w}^\pi, f; \tilde{\Delta}, \tilde{\Theta}) - \Phi_{\text{mis}}^\pi(\tilde{w}^\pi, f; \Delta^*, \Theta^*) \right| + \max_{f \in \mathcal{V}} \Phi_{\text{mis}}^\pi(\tilde{w}^\pi, f; \Delta^*, \Theta^*), \end{aligned}$$

where we use triangle inequality and the fact that $w^\dagger \in \operatorname{argmin}_{w \in \mathcal{W}} \max_{f \in \mathcal{V}} \widehat{\Phi}_{\text{mis}}^\pi(w, f; \Delta, \Theta)$ in the last inequality. In the meanwhile, by Theorem C.7.6, with probability at least $1 - \delta$,

it holds for any $(w, f, \pi) \in \mathcal{W} \times \mathcal{V} \times \Pi$ that

(C.5.47)

$$\left| \Phi_{\text{mis}}^{\pi}(w, f; \tilde{\Delta}, \tilde{\Theta}) - \hat{\Phi}_{\text{mis}}^{\pi}(w, f; \tilde{\Delta}, \tilde{\Theta}) \right| \leq c \cdot \frac{C_{\Delta^*} C_{\Theta^*} C_*}{1 - \gamma} \sqrt{\frac{1}{NT\kappa} \mathfrak{C}_{\mathcal{V}, \mathcal{W}, \Pi} \log \frac{1}{\delta} \log(NT)}.$$

Also, we upper bound the second term on the RHS of (C.5.46) with probability at least $1 - \delta$ as follows,

$$\begin{aligned} & \left| \Phi_{\text{mis}}^{\pi}(\tilde{w}^{\pi}, f; \tilde{\Delta}, \tilde{\Theta}) - \Phi_{\text{mis}}^{\pi}(\tilde{w}^{\pi}, f; \Delta^*, \Theta^*) \right| \\ &= \left| \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} \left(\frac{Z_t^{\top} A_t \pi(A_t | S_t) w(S_t)}{\Delta^*(S_t, A_t) \Theta^*(S_t, Z_t)} - \frac{Z_t^{\top} A_t \pi(A_t | S_t) w(S_t)}{\tilde{\Delta}(S_t, A_t) \tilde{\Theta}(S_t, Z_t)} \right) (f(S_t) - \gamma f(S_{t+1})) \right] \right| \\ & \text{(C.5.48)} \end{aligned}$$

$$\leq \frac{C_{\Delta^*} C_{\Theta^*} C_*}{1 - \gamma} \left(\xi_0 C_{\Theta^*} \sqrt{\frac{C_{\Delta^*}}{NT\kappa} \mathfrak{C}_{\mathcal{F}_0} \log \frac{1}{\delta} \log(NT)} + \xi_1 C_{\Delta^*} \sqrt{\frac{C_{\Theta^*}}{NT\kappa} \mathfrak{C}_{\mathcal{F}_1} \log \frac{1}{\delta} \log(NT)} \right),$$

where we use Cauchy-Schwarz inequality and Assumption 4.4.4 in the last inequality.

Now, by plugging (C.5.47) and (C.5.48) into (C.5.46), it holds with probability at least $1 - \delta$ that

(C.5.49)

$$\text{Term (II)} \leq c \cdot \frac{C_{\Delta^*}^2 C_{\Theta^*}^2 C_*}{1 - \gamma} (\xi_0 + \xi_1) \sqrt{\frac{\mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}, \Pi}}{NT\kappa} \log \frac{1}{\delta} \log(NT)} + \max_{f \in \mathcal{V}} \Phi_{\text{mis}}^{\pi}(\tilde{w}^{\pi}, f; \Delta^*, \Theta^*).$$

Upper Bounding Term (III). Note that

$$\begin{aligned}
& \left| \widehat{\Phi}_{\text{mis}}^{\pi}(w, f; \Delta^*, \Theta^*) - \widehat{\Phi}_{\text{mis}}^{\pi}(w, f; \widetilde{\Delta}, \widetilde{\Theta}) \right| \\
& \leq \left| \left(\widehat{\mathbb{E}} - \mathbb{E} \right) \left[\frac{1}{T} \sum_{t=0}^{T-1} \left(\frac{Z_t^{\top} A_t \pi(A_t | S_t) w(S_t)}{\Delta^*(S_t, A_t) \Theta^*(S_t, Z_t)} - \frac{Z_t^{\top} A_t \pi(A_t | S_t) w(S_t)}{\widetilde{\Delta}(S_t, A_t) \widetilde{\Theta}(S_t, Z_t)} \right) (f(S_t) - \gamma f(S_{t+1})) \right] \right| \\
\text{(C.5.50)} \quad & + \left| \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} \left(\frac{Z_t^{\top} A_t \pi(A_t | S_t) w(S_t)}{\Delta^*(S_t, A_t) \Theta^*(S_t, Z_t)} - \frac{Z_t^{\top} A_t \pi(A_t | S_t) w(S_t)}{\widetilde{\Delta}(S_t, A_t) \widetilde{\Theta}(S_t, Z_t)} \right) (f(S_t) - \gamma f(S_{t+1})) \right] \right|.
\end{aligned}$$

For the first term on the RHS of (C.5.50), by Theorem C.7.6, with probability at least $1 - \delta$, it holds for any $(w, f, \pi) \in \mathcal{W} \times \mathcal{V} \times \Pi$ that

$$\begin{aligned}
& \left| \left(\widehat{\mathbb{E}} - \mathbb{E} \right) \left[\frac{1}{T} \sum_{t=0}^{T-1} \left(\frac{Z_t^{\top} A_t \pi(A_t | S_t) w(S_t)}{\Delta^*(S_t, A_t) \Theta^*(S_t, Z_t)} - \frac{Z_t^{\top} A_t \pi(A_t | S_t) w(S_t)}{\widetilde{\Delta}(S_t, A_t) \widetilde{\Theta}(S_t, Z_t)} \right) (f(S_t) - \gamma f(S_{t+1})) \right] \right| \\
\text{(C.5.51)} \quad & \leq c \cdot \frac{C_{\Delta^*} C_{\Theta^*} C_*}{1 - \gamma} \sqrt{\frac{\mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}, \Pi}}{NT\kappa} \log \frac{1}{\delta} \log(NT)}.
\end{aligned}$$

For the second term on the RHS of (C.5.50), by a similar argument as in (C.3.14), it holds with probability at least $1 - \delta$ that

$$\begin{aligned}
& \left| \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} \left(\frac{Z_t^{\top} A_t \pi(A_t | S_t) w(S_t)}{\Delta^*(S_t, A_t) \Theta^*(S_t, Z_t)} - \frac{Z_t^{\top} A_t \pi(A_t | S_t) w(S_t)}{\widetilde{\Delta}(S_t, A_t) \widetilde{\Theta}(S_t, Z_t)} \right) (f(S_t) - \gamma f(S_{t+1})) \right] \right| \\
\text{(C.5.52)} \quad & \leq \frac{2C_{\Delta^*} C_{\Theta^*} C_*}{1 - \gamma} \left(\xi_0 C_{\Theta^*} \sqrt{\frac{C_{\Delta^*}}{NT\kappa} \mathfrak{C}_{\mathcal{F}_0} \log \frac{1}{\delta} \log(NT)} + \xi_1 C_{\Delta^*} \sqrt{\frac{C_{\Theta^*}}{NT\kappa} \mathfrak{C}_{\mathcal{F}_1} \log \frac{1}{\delta} \log(NT)} \right),
\end{aligned}$$

where in the first inequality, we use the fact that $\|f\|_{\infty} \leq 1/(1 - \gamma)$ and $\|w\|_{\infty} \leq C_*$; in the third inequality, we use Cauchy Schwarz inequality; while in the last inequality,

we use Assumption 4.4.4 with the fact that $(\tilde{\Delta}, \tilde{\Theta}) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1$. Now, by plugging (C.5.51) and (C.5.52) into (C.5.50), with probability at least $1 - \delta$, it holds for any $w \in \cup_{(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1} \text{conf}_{\alpha_{\text{mis}}}^{\text{mis}}(\Delta, \Theta, \pi)$ and $(f, \pi) \in \mathcal{V} \times \Pi$ that

$$(C.5.53) \quad \begin{aligned} & \left| \widehat{\Phi}_{\text{mis}}^\pi(w, f; \Delta^*, \Theta^*) - \widehat{\Phi}_{\text{mis}}^\pi(w, f; \tilde{\Delta}, \tilde{\Theta}) \right| \\ & \leq c \cdot \frac{C_{\Delta^*}^2 C_{\Theta^*}^2 C_*}{1 - \gamma} (\xi_0 + \xi_1) \sqrt{\frac{1}{NT\kappa} \cdot \mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}, \Pi} \cdot \log \frac{1}{\delta} \log(NT)}. \end{aligned}$$

Now, by plugging (C.5.45), (C.5.49), and (C.5.53) into (C.5.44), with probability at least $1 - \delta$, it holds for any $\pi \in \Pi$ and $w \in \cup_{(\Delta, \Theta) \in \text{conf}_{\alpha_0}^0 \times \text{conf}_{\alpha_1}^1} \text{conf}_{\alpha_{\text{mis}}}^{\text{mis}}(\Delta, \Theta, \pi)$ that

$$\begin{aligned} \max_{f \in \mathcal{V}} \Phi_{\text{mis}}^\pi(w, f; \Delta^*, \Theta^*) & \leq c \cdot \frac{C_{\Delta^*}^2 C_{\Theta^*}^2 C_*}{1 - \gamma} (\xi_0 + \xi_1) \sqrt{\frac{1}{NT\kappa} \cdot \mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \mathcal{V}, \Pi} \cdot \log \frac{1}{\delta} \log(NT)} \\ & \quad + \max_{f \in \mathcal{V}} \Phi_{\text{mis}}^\pi(\tilde{w}^\pi, f; \Delta^*, \Theta^*), \end{aligned}$$

which concludes the proof of the lemma. \square

C.6. Proof of Results in §4.5

C.6.1. Proof of Theorem 4.5.1

Proof. We denote by $(\widehat{\pi}, \widehat{\lambda}, \widehat{v})$ the solution of (4.5.1). Note that $(1 - \gamma)\mathbb{E}_{S_0 \sim \nu}[v(S_0)]$ is a lower bounded real-valued convex functional w.r.t. v , and $\widehat{M}_{\mathcal{V}f}^{\pi^*}$ is also a convex functional. In the meanwhile, we have $\widehat{M}_{\mathcal{V}f}^{\pi^*}(\widehat{v}_{\Delta^*, \Theta^*}^{\pi^*}) = 0$. Thus, by Theorem 1 of §8.6 in

Luenberger (1997), strong duality holds, i.e.,

$$\begin{aligned}
& \max_{\lambda \geq 0} \min_{v \in \mathcal{V}} \left\{ (1 - \gamma) \mathbb{E}_{S_0 \sim \nu} [v(S_0)] + \lambda \cdot \left(\widehat{M}_{\text{vf}}^{\pi^*}(v) - \alpha_{\text{vf}} \right) \right\} \\
\text{(C.6.1)} \quad & = \min_{v \in \mathcal{V}} \max_{\lambda \geq 0} \left\{ (1 - \gamma) \mathbb{E}_{S_0 \sim \nu} [v(S_0)] + \lambda \cdot \left(\widehat{M}_{\text{vf}}^{\pi^*}(v) - \alpha_{\text{vf}} \right) \right\}.
\end{aligned}$$

By Lemma 4.4.7, it holds with probability at least $1 - \delta$ that $\widehat{M}_{\text{vf}}^{\widehat{\pi}}(V^{\widehat{\pi}}) \leq \alpha_{\text{vf}}$. Thus, with probability at least $1 - \delta$, we have

$$\begin{aligned}
& J(\pi^*) - J(\widehat{\pi}) \\
& = (1 - \gamma) \mathbb{E}_{S_0 \sim \nu} [V^{\pi^*}(S_0) - V^{\widehat{\pi}}(S_0)] \\
& \leq (1 - \gamma) \mathbb{E}_{S_0 \sim \nu} [V^{\pi^*}(S_0)] - \left((1 - \gamma) \mathbb{E}_{S_0 \sim \nu} [V^{\widehat{\pi}}(S_0)] + \widehat{\lambda} \cdot \left(\widehat{M}_{\text{vf}}^{\widehat{\pi}}(V^{\widehat{\pi}}) - \alpha_{\text{vf}} \right) \right) \\
& \leq (1 - \gamma) \mathbb{E}_{S_0 \sim \nu} [V^{\pi^*}(S_0)] - \min_{v \in \mathcal{V}} \left\{ (1 - \gamma) \mathbb{E}_{S_0 \sim \nu} [v(S_0)] + \widehat{\lambda} \cdot \left(\widehat{M}_{\text{vf}}^{\widehat{\pi}}(v) - \alpha_{\text{vf}} \right) \right\} \\
& = (1 - \gamma) \mathbb{E}_{S_0 \sim \nu} [V^{\pi^*}(S_0)] - \max_{\pi \in \Pi} \max_{\lambda \geq 0} \min_{v \in \mathcal{V}} \left\{ (1 - \gamma) \mathbb{E}_{S_0 \sim \nu} [v(S_0)] + \lambda \cdot \left(\widehat{M}_{\text{vf}}^{\pi}(v) - \alpha_{\text{vf}} \right) \right\} \\
\text{(C.6.2)} \quad & = (1 - \gamma) \mathbb{E}_{S_0 \sim \nu} [V^{\pi^*}(S_0)] - \max_{\lambda \geq 0} \min_{v \in \mathcal{V}} \left\{ (1 - \gamma) \mathbb{E}_{S_0 \sim \nu} [v(S_0)] + \lambda \cdot \left(\widehat{M}_{\text{vf}}^{\pi^*}(v) - \alpha_{\text{vf}} \right) \right\},
\end{aligned}$$

where in the second inequality, we use the fact that $V^{\widehat{\pi}} \in \mathcal{V}$; in the third inequality, we use the definition of $\widehat{\pi}$ and $\widehat{\lambda}$; in the last inequality, we use the fact that $\pi^* \in \Pi$. By

combining (C.6.1) and (C.6.2), we have

$$\begin{aligned}
& J(\pi^*) - J(\widehat{\pi}) \\
& \leq (1 - \gamma)\mathbb{E}_{S_0 \sim \nu} [V^{\pi^*}(S_0)] - \min_{v \in \mathcal{V}} \max_{\lambda \geq 0} \left\{ (1 - \gamma)\mathbb{E}_{S_0 \sim \nu} [v(S_0)] + \lambda \cdot \left(\widehat{M}_{\text{vf}}^{\pi^*}(v) - \alpha_{\text{vf}} \right) \right\} \\
& \leq (1 - \gamma)\mathbb{E}_{S_0 \sim \nu} [V^{\pi^*}(S_0)] - \min_{v \in \mathcal{V}: \widehat{M}_{\text{vf}}^{\pi^*}(v) \leq \alpha_{\text{vf}}} (1 - \gamma)\mathbb{E}_{S_0 \sim \nu} [v(S_0)] \\
\text{(C.6.3)} \quad & \leq \max_{v \in \mathcal{V}: \widehat{M}_{\text{vf}}^{\pi^*}(v) \leq \alpha_{\text{vf}}} \left| (1 - \gamma)\mathbb{E}_{S_0 \sim \nu} [V^{\pi^*}(S_0) - v(S_0)] \right|.
\end{aligned}$$

Note that by Lemmas 4.2.7 and 4.2.8, we have

$$\text{(C.6.4)} \quad (1 - \gamma)\mathbb{E}_{S_0 \sim \nu} [V^{\pi^*}(S_0) - v(S_0)] = \Phi_{\text{vf}}^{\pi^*}(v, w^{\pi^*}; \Delta^*, \Theta^*).$$

By plugging (C.6.4) into (C.6.3), we have

$$\begin{aligned}
J(\pi^*) - J(\widehat{\pi}) & \leq \max_{v \in \mathcal{V}: \widehat{M}_{\text{vf}}^{\pi^*}(v) \leq \alpha_{\text{vf}}} \left| \Phi_{\text{vf}}^{\pi^*}(v, w^{\pi^*}; \Delta^*, \Theta^*) \right| \\
& \leq c \cdot \frac{C_{\Delta^*}^2 C_{\Theta^*}^2 C_*}{1 - \gamma} (\xi_0 + \xi_1) L_{\Pi} \sqrt{\frac{1}{NT\kappa} \cdot \mathfrak{C}_{\mathcal{F}_0, \mathcal{F}_1, \mathcal{W}, \nu, \Pi} \cdot \log \frac{1}{\delta} \log(NT)},
\end{aligned}$$

where in the last inequality, we use Lemma 4.4.8 with the fact that $w^{\pi^*} \in \mathcal{W}$. This concludes the proof. \square

C.7. Auxiliary Results

We introduce auxiliary results used in the paper. We provide the proofs of these results in §C.8. We first introduce the following definition of β -mixing coefficient.

Definition C.7.1. Let $\{Z_t\}_{t \geq 0}$ be a sequence of random variables. For any $i, j \in \mathbb{N} \cup \{\infty\}$, we denote by σ_i^j the sigma algebra generated by $\{Z_k\}_{i \leq k \leq j}$. The β -mixing coefficient of $\{Z_t\}_{t \geq 0}$ is defined as $\beta(t) = \sup_n \mathbb{E}_{B \in \sigma_0^n} [\sup_{A \in \sigma_{n+t}^\infty} |\mathbb{P}(A | B) - \mathbb{P}(A)|]$.

We introduce the following form of β -mixing coefficient for Markov chains.

Lemma C.7.2. Suppose $\{Z_t\}_{t \geq 0}$ is a Markov chain with initial distribution ζ . It holds that

$$\beta(t) \leq \frac{1}{2} \int \|p_{t'}(\cdot | z) - p_{\text{stat}}(\cdot)\|_{\text{TV}} dp_{\text{stat}}(z) + \frac{3}{2} \int \|p_{t'}(\cdot | z) - p_{\text{stat}}(\cdot)\|_{\text{TV}} d\zeta(z),$$

where $t' = \lfloor t/2 \rfloor$ and $p_n(\cdot | z)$ is the marginal the distribution of Z_n given $Z_0 = z$ for any $n \in [N]$.

Proof. See the proof of Lemma 1 in Meitz and Saikkonen (2021) for a detailed proof. □

Following from Lemma C.7.2, we can upper bound the β -mixing coefficient for a Markov chain $\{Z_t\}_{t \geq 0}$. Before that, we impose the following assumption on $\{Z_t\}_{t \geq 0}$.

Assumption C.7.3. The Markov chain $\{Z_t\}_{t \geq 0}$ with initial distribution ζ admits a unique stationary distribution p_{stat} over \mathcal{Z} and is geometrically ergodic, i.e., there exists a function $\varphi: \mathcal{Z} \rightarrow [0, \infty)$ and a constant $\kappa > 0$ such that

$$\|p_{\text{stat}}(\cdot) - p_t(\cdot | z_0)\|_{\text{TV}} \leq \varphi(z_0) \cdot \exp(-2\kappa t),$$

where $p_t(\cdot | z_0)$ is the marginal distribution of Z_t given $Z_0 = z_0$ and there exists a positive absolute constant c such that $\int \varphi(z) d\zeta(z) \leq c$ and $\int \varphi(z) dp_{\text{stat}}(z) \leq c$.

Lemma C.7.4. Suppose $\{Z_t\}_{t \geq 0}$ is a Markov chain satisfying Assumption C.7.3. Then we have $\beta(t) \leq c \cdot \exp(-\kappa t)$ for any $t \geq 0$.

Proof. For any $t \geq 0$, by Lemma C.7.2, we have

$$\begin{aligned} \beta(t) &\leq \frac{1}{2} \int \|p_{t'}(\cdot | z) - p_{\text{stat}}(\cdot)\|_{\text{TV}} d p_{\text{stat}}(z) + \frac{3}{2} \int \|p_{t'}(\cdot | z) - p_{\text{stat}}(\cdot)\|_{\text{TV}} d \zeta(z) \\ &\leq \frac{1}{2} \int \varphi(z) \cdot \exp(-\kappa t) d p_{\text{stat}}(z) + \frac{3}{2} \int \varphi(z) \cdot \exp(-\kappa t) d \zeta(z) \\ &\leq c \cdot \exp(-\kappa t), \end{aligned}$$

where in the second and last inequalities, we use Assumption C.7.3. This concludes the proof of the lemma. \square

C.7.1. Concentration Inequality for Geometrically Ergodic Non-Stationary Sequence

We first introduce the following lemma, which is a straight-forward generalization of Berbee's lemma (Berbee, 1979).

Lemma C.7.5. For any $k > 0$ and a random sequence $\{Y_\ell\}_{\ell=1}^k$, there exists a random sequence $\{\tilde{Y}_\ell\}_{\ell=1}^k$ such that

- (1) $\{\tilde{Y}_\ell\}_{\ell=1}^k$ are independent;
- (2) for any $1 \leq \ell \leq k$, \tilde{Y}_ℓ and Y_ℓ have the same distribution;
- (3) for any $1 \leq \ell \leq k$, $\mathbb{P}(\tilde{Y}_\ell \neq Y_\ell) = \beta(\sigma(\{Y_{\ell'}\}_{\ell'=1}^{\ell-1}), \sigma(\{Y_\ell\}))$.

Proof. See Lemma 2.10 in Barrera and Gobet (2021) for a detailed proof. \square

We introduce the following Hoeffding's Inequality and Bernstein's Inequality for geometrically ergodic non-stationary sequences.

Theorem C.7.6 (Hoeffding's Inequality for geometrically ergodic non-stationary sequences). We denote by $\{X_t\}_{t \geq 0} \subseteq \mathcal{X}$ a Markov chain satisfying Assumption C.7.3. Then for any function $f: \mathcal{X} \rightarrow [-f_{\max}, f_{\max}]$, it holds with probability at least $1 - \delta$ with $c/(NT)^2 \leq \delta \leq 1$ that

$$\left| \frac{1}{NT} \sum_{i \in [N]} \sum_{t=0}^{T-1} f(X_t^i) - \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} f(X_t) \right] \right| \leq c \cdot f_{\max} \sqrt{\frac{1}{NT\kappa} \log \frac{2}{\delta} \log(NT)},$$

where $\{\{X_t^i\}_{t=0}^{T-1}\}_{i \in [N]}$ consists of N i.i.d. trajectories with length $T > 0$ generated from the same distribution as $\{X_t\}_{t \geq 0}$.

Proof. See §C.8.1 for a detailed proof. □

Theorem C.7.7 (Bernstein's Inequality for geometrically ergodic non-stationary sequences). We denote by $\{X_t\}_{t \geq 0} \subseteq \mathcal{X}$ a Markov chain satisfying Assumption C.7.3. Then for any function $f: \mathcal{X} \rightarrow [-f_{\max}, f_{\max}]$, it holds with probability at least $1 - \delta$ with $c/(NT)^2 \leq \delta \leq 1$ that

$$\begin{aligned} & \frac{1}{NT} \sum_{i \in [N]} \sum_{t=0}^{T-1} f(X_t^i) - \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} f(X_t) \right] \\ & \leq c_1 \cdot \frac{f_{\max}}{NT\kappa} \log \frac{2}{\delta} \log(NT) + c_2 \cdot \sqrt{\frac{1}{NT\kappa} \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} f(X_t)^2 \right] \log \frac{2}{\delta} \log(NT)}, \end{aligned}$$

where c_1 and c_2 are positive absolute constants, and $\{\{X_t^i\}_{t=0}^{T-1}\}_{i \in [N]}$ consists of N i.i.d. trajectories with length $T > 0$ generated from the same distribution as $\{X_t\}_{t \geq 0}$.

Proof. See §C.8.2 for a detailed proof. \square

C.7.2. Empirical Processes for Geometrically Ergodic Non-Stationary Sequence

For any conditional probabilities $p_1(y|x)$ and $p_2(y|x)$ such that $(x, y) \in \mathcal{X} \times \mathcal{Y}$, we define the squared Hellinger distance as follows,

$$h^2(p_1(\cdot|x), p_2(\cdot|x)) = \frac{1}{2} \int \left(\sqrt{p_1(y|x)} - \sqrt{p_2(y|x)} \right)^2 dy.$$

We further assume that \mathcal{Y} is a discrete space. We denote by $p^*(y|x)$ the true conditional probability of $y \in \mathcal{Y}$ given $x \in \mathcal{X}$. Also, let $\{(X_t, Y_t)\}_{t \geq 0} \subset \mathcal{X} \times \mathcal{Y}$ be a Markov chain such that $Y_t \sim p^*(\cdot|X_t)$ and satisfies Assumption C.7.3. Further, we denote by μ_t the marginal distribution of X_t for any $t \geq 0$. In the meanwhile, with $\mu = 1/T \cdot \sum_{t=0}^{T-1} \mu_t$, we define the generalized squared Hellinger distance over μ as follows,

$$H^2(p_1, p_2) = \mathbb{E}_{X \sim \mu} \left[h^2(p_1(\cdot|X), p_2(\cdot|X)) \right].$$

In the meanwhile, we are given a data set $\{ \{(X_t^i, Y_t^i)\}_{t=0}^{T-1} \}_{i \in [N]}$ consisting of N independent trajectories of length T , where $\{(X_t^i, Y_t^i)\}_{t=0}^{T-1}$ is generated from the same distribution as $\{(X_t, Y_t)\}_{t \geq 0}$. We construct the following maximum likelihood estimator for p^* ,

$$(C.7.1) \quad \hat{p} \in \operatorname{argmax}_{p \in \mathcal{P}} \widehat{\mathbb{E}} [\log p(Y|X)] = \frac{1}{NT} \sum_{i \in [N]} \sum_{t=0}^{T-1} \log p(Y_t^i | X_t^i).$$

We also define

$$g_p(x, y) = \frac{1}{2} \log \frac{p(y|x) + p^*(y|x)}{2p^*(y|x)}$$

for any $(x, y) \in \mathcal{X} \times \mathcal{Y}$.

Now, we are ready to introduce the following lemma.

Lemma C.7.8. We have

$$\begin{aligned} H^2\left(\frac{\widehat{p} + p^*}{2}, p^*\right) &\leq \left(\widehat{\mathbb{E}} - \mathbb{E}\right) [g_{\widehat{p}}(X, Y)], \\ H^2\left(\frac{p_1 + p^*}{2}, \frac{p_2 + p^*}{2}\right) &\leq \frac{1}{2}H^2(p_1, p_2), \\ H^2(p, p^*) &\leq 16H^2\left(\frac{p + p^*}{2}, p^*\right), \\ \|p_1(\cdot | x) - p_2(\cdot | x)\|_1 &\leq 2\sqrt{2}h(p_1(\cdot | x), p_2(\cdot | x)). \end{aligned}$$

Proof. See §C.8.3 for a detailed proof. □

We define the entropy integral as follows,

$$J_B(\delta, \overline{\mathcal{P}}^{1/2}(\delta)) = \max \left\{ \int_{\delta^2/2^{10}}^{\delta} \left(H_B(u, \overline{\mathcal{P}}^{1/2}(\delta)) \right)^{1/2} du, \delta \right\},$$

where $H_B(u, \overline{\mathcal{P}}^{1/2}(\delta))$ is the entropy of the space $\overline{\mathcal{P}}^{1/2}(\delta)$ with bracketing, and $\overline{\mathcal{P}}^{1/2}(\delta)$ is defined as follows,

$$\overline{\mathcal{P}}^{1/2}(\delta) = \{\overline{p}^{1/2} : p \in \mathcal{P} \text{ and } H^2(\overline{p}, p^*) \leq \delta^2\}.$$

Now, we introduce the following theorem, which upper bounds the distance between \widehat{p} and p^* .

Theorem C.7.9. We take $\Psi(\delta) \geq J_B(\delta, \overline{\mathcal{P}}^{1/2}(\delta))$ in such a way that $\Psi(\delta)/\delta^2$ is a non-increasing function of δ . Then for a universal constant c and any $\delta \geq \delta_{NT}$, where δ_{NT} satisfies that $\sqrt{NT}\delta_{NT}^2 \geq c\Psi(\delta_{NT})$, it holds with probability at least $1 - c/\kappa$.

$\exp(-NT\kappa\delta^2/(c^2\log(NT))) - c/(N^2T^2) \cdot \log(4/\delta)$ that

$$H^2(\widehat{p}, p^*) \leq \delta^2,$$

where \widehat{p} is defined in (C.7.1).

Proof. See §C.8.4 for a detailed proof. □

We study the following case, where \mathcal{P} is a parametric class.

Corollary C.7.10 (Parametric Class). Suppose $\mathcal{P} = \{p_\theta: \theta \in \mathbb{R}^d \text{ and } \|\theta\|_2 \leq \theta_{\max}\}$.

Then with probability at least $1 - \delta$ with $c/(N^2T^2) \cdot \log(NT) \leq \delta \leq 1$, we have

$$H^2(\widehat{p}, p^*) \leq c \cdot \frac{d}{NT\kappa} \log \frac{\theta_{\max}}{\delta} \log(NT),$$

where $c > 0$ is an absolute constant, which may vary from lines to lines.

Proof. Note that

$$J_B(\delta, \overline{\mathcal{P}}^{1/2}(\delta), d) \leq \delta \sqrt{d \log \frac{\theta_{\max}}{\delta}}.$$

By taking $\Psi(\delta) = \delta \sqrt{d \log(\theta_{\max}/\delta)}$, we have

$$\mathbb{P} \left(H^2(\widehat{p}, p^*) \leq c \cdot \frac{d}{NT\kappa} \log \frac{\theta_{\max}}{\delta} \log(NT) \right) \geq 1 - \delta$$

with $c/(N^2T^2) \cdot \log(NT) \leq \delta \leq 1$, which concludes the proof of the corollary. □

C.8. Proofs of Auxiliary Results

C.8.1. Proof of Theorem C.7.6

Proof. We take $\tau = \min\{T, 3/\kappa \cdot \log(NT)\}$, and denote by $\mathbb{E}_{\text{stat}}[\cdot]$ the expectation taken with respect to the stationary distribution of $\{X_t\}_{t \geq 0}$. We observe the following decomposition,

$$\begin{aligned}
& \frac{1}{NT} \sum_{i \in [N]} \sum_{t=0}^{T-1} f(X_t^i) - \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} f(X_t) \right] \\
&= \frac{\tau}{T} \underbrace{\left(\frac{1}{N\tau} \sum_{i \in [N]} \sum_{t=0}^{\tau-1} f(X_t^i) - \mathbb{E} \left[\frac{1}{\tau} \sum_{t=0}^{\tau-1} f(X_t) \right] \right)}_{\text{(I)}} \\
&+ \frac{T-\tau}{T} \underbrace{\left(\frac{1}{N(T-\tau)} \sum_{i \in [N]} \sum_{t=\tau}^{T-1} f(X_t^i) - \mathbb{E}_{\text{stat}} [f(X)] \right)}_{\text{(II)}} \\
&+ \frac{T-\tau}{T} \underbrace{\left(\mathbb{E}_{\text{stat}} [f(X)] - \mathbb{E} \left[\frac{1}{T-\tau} \sum_{t=\tau}^{T-1} f(X_t) \right] \right)}_{\text{(III)}}.
\end{aligned}$$

In the follows, we upper bound terms (I), (II), and (III), respectively.

Upper Bounding Term (I). Since $\{\{X_t^i\}_{t=0}^{T-1}\}_{i \in [N]}$ are i.i.d. accross each trajectory, by standard Hoeffding's inequality, with probability at least $1 - \delta$, we have

$$(C.8.1) \quad |(\text{I})| \leq f_{\max} \sqrt{\frac{2}{N} \log \frac{2}{\delta}}.$$

Upper Bounding Term (II). We consider an auxiliary Markov chain $\{\{\tilde{X}_t^i\}_{t=0}^{T-1}\}_{i \in [N]}$, where the i -th trajectory $\{\tilde{X}_t^i\}_{t=0}^{T-1}$ is sampled such that $\tilde{X}_0^i \sim p_{\text{stat}}$. Here p_{stat} is the stationary distribution of $\{X_t\}_{t \geq 0}$. Similarly, we define the following quantity,

$$(\tilde{\text{II}}) = \frac{1}{N(T-\tau)} \sum_{i \in [N]} \sum_{t=\tau}^{T-1} f(\tilde{X}_t^i) - \mathbb{E}_{\text{stat}} [f(X)].$$

Now, for any $x \geq 0$, we upper bound the difference $\mathbb{P}(\text{II}) \geq x) - \mathbb{P}(\tilde{\text{II}}) \geq x)$ as follows,

(C.8.2)

$$\mathbb{P}(\text{II}) \geq x) - \mathbb{P}(\tilde{\text{II}}) \geq x) \leq N \sum_{t=\tau}^{T-1} \mathbb{E} [\|p_t(\cdot | X_0) - p_{\text{stat}}(\cdot)\|_{\text{TV}}] \leq c \cdot NT \exp(-\kappa\tau).$$

Thus, by (C.8.2), to upper bound $\mathbb{P}(\text{II}) \geq x)$, it suffices to upper bound $\mathbb{P}(\tilde{\text{II}}) \geq x)$.

To upper bound $\mathbb{P}(\tilde{\text{II}}) \geq x)$, we take $T - \tau = 2ks$, where k and s are two positive integers for the simplicity of presentation. We partition the set $\{\tau, \tau + 1, \dots, T - 1\}$ as follows,

$$J_1 = \{\tau, \tau + 1, \dots, \tau + s - 1\}, \quad J_2 = \{\tau + s, \tau + s + 1, \dots, \tau + 2s - 1\}, \quad \dots, \\ J_{2k-1} = \{T - 2s, T - 2s + 1, \dots, T - s - 1\}, \quad J_{2k} = \{T - s, T - s + 1, \dots, T - 1\}.$$

Under such a partition, we see that $\cup_{\ell \in [2k]} J_\ell = \{\tau, \tau + 1, \dots, T - 1\}$ and $J_\ell \cap J_{\ell'} = \emptyset$ for any $\ell \neq \ell'$. Also, for any $i \in [N]$, we define

$$Z_\ell^i = (\tilde{X}_t^i)_{t \in J_\ell}$$

for any $\ell \in [2k]$. Now, for any $i \in [N]$, by Lemma C.7.5, there exists a sequence $\{W_\ell^i\}_{\ell \in [2k]}$, where $W_\ell^i = (\tilde{Y}_t^i)_{t \in J_\ell}$ such that

- (1) $\{W_\ell^i\}_{\ell \in [2k]}$ are independent;
- (2) for any $\ell \in [2k]$, W_ℓ^i and Z_ℓ^i have the same distribution;
- (3) for any $\ell \in [2k]$, $\mathbb{P}(W_\ell^i \neq Z_\ell^i) = \beta(\sigma(\{Z_{\ell'}^i\}_{\ell' \in [\ell-1]}), \sigma(\{Z_\ell^i\}))$.

Note that the following inclusion relation holds,

$$\left\{ \frac{1}{s} \sum_{t \in J_\ell} f(\tilde{X}_t^i) - \mathbb{E}_{\text{stat}}[f(X)] \geq x_\ell \right\} \subseteq \left\{ \frac{1}{s} \sum_{t \in J_\ell} f(\tilde{Y}_t^i) - \mathbb{E}_{\text{stat}}[f(X)] \geq x_\ell \right\} \cup \{W_\ell^i \neq Z_\ell^i\}$$

for any $x_\ell \in \mathbb{R}$. Thus, we have

$$\begin{aligned} \mathbb{P}(\tilde{\text{II}} \geq x) &\leq \mathbb{P}\left(\frac{1}{kN} \sum_{i \in [N], \ell \text{ is odd}} \frac{1}{s} \sum_{t \in J_\ell} f(\tilde{X}_t^i) - \mathbb{E}_{\text{stat}}[f(X)] \geq x\right) \\ &\quad + \mathbb{P}\left(\frac{1}{kN} \sum_{i \in [N], \ell \text{ is even}} \frac{1}{s} \sum_{t \in J_\ell} f(\tilde{X}_t^i) - \mathbb{E}_{\text{stat}}[f(X)] \geq x\right) \\ &\leq \mathbb{P}\left(\frac{1}{kN} \sum_{i \in [N], \ell \text{ is odd}} \frac{1}{s} \sum_{t \in J_\ell} f(\tilde{Y}_t^i) - \mathbb{E}_{\text{stat}}[f(X)] \geq x\right) \\ &\quad + \mathbb{P}\left(\frac{1}{kN} \sum_{i \in [N], \ell \text{ is even}} \frac{1}{s} \sum_{t \in J_\ell} f(\tilde{Y}_t^i) - \mathbb{E}_{\text{stat}}[f(X)] \geq x\right) + \sum_{\ell \in [2k]} \mathbb{P}(W_\ell^i \neq Z_\ell^i) \\ &\leq 2 \exp\left(-\frac{kNx^2}{2f_{\max}^2}\right) + 2k\beta(s), \end{aligned}$$

where we use Hoeffding's inequality in the last line. Similarly, we have

$$\mathbb{P}(\tilde{\text{II}} \leq -x) \leq 2 \exp\left(-\frac{kNx^2}{2f_{\max}^2}\right) + 2k\beta(s).$$

Thus, we have

$$\mathbb{P}(|(\text{II})| \geq x) \leq 4 \exp\left(-\frac{kNx^2}{2f_{\max}^2}\right) + 4k\beta(s) + c \cdot NT \exp(-\kappa\tau).$$

Now, by taking $s = 3 \log(NT)/\kappa$, it holds with probability at least $1 - \delta$ with $c/(NT)^2 \leq \delta \leq 1$ that

$$(C.8.3) \quad |(\text{II})| \leq f_{\max} \sqrt{\frac{24}{NT\kappa} \log \frac{4}{\delta} \log(NT)}.$$

Upper Bounding Term (III). We observe that

$$(C.8.4) \quad |(\text{III})| \leq f_{\max} \cdot \sum_{t=\tau}^{T-1} c \cdot \exp(-\kappa t) \leq c \cdot \frac{f_{\max}}{N^2 T^2}.$$

Combining Everything. Now, by combining (C.8.1), (C.8.3), and (C.8.4), it holds with probability at least $1 - \delta$ with $c/(NT)^2 \leq \delta \leq 1$ that

$$\frac{1}{NT} \sum_{i \in [N]} \sum_{t=0}^{T-1} f(X_t^i) - \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} f(X_t) \right] \leq c \cdot f_{\max} \sqrt{\frac{48}{NT\kappa} \log \frac{4}{\delta} \log(NT)},$$

which concludes the proof of the theorem. \square

C.8.2. Proof of Theorem C.7.7

Proof. The proof follows from the proof of Theorem C.7.6 in §C.8.1. For the completeness of the paper, we present it here. We take $\tau = \min\{T, 3/\kappa \cdot \log(NT)\}$, and denote by $\mathbb{E}_{\text{stat}}[\cdot]$ the expectation taken with respect to the stationary distribution of $\{X_t\}_{t \geq 0}$.

We observe the following decomposition,

$$\begin{aligned}
& \frac{1}{NT} \sum_{i \in [N]} \sum_{t=0}^{T-1} f(X_t^i) - \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} f(X_t) \right] \\
&= \underbrace{\frac{\tau}{T} \left(\frac{1}{N\tau} \sum_{i \in [N]} \sum_{t=0}^{\tau-1} f(X_t^i) - \mathbb{E} \left[\frac{1}{\tau} \sum_{t=0}^{\tau-1} f(X_t) \right] \right)}_{\text{(I)}} \\
&\quad + \underbrace{\frac{T-\tau}{T} \left(\frac{1}{N(T-\tau)} \sum_{i \in [N]} \sum_{t=\tau}^{T-1} f(X_t^i) - \mathbb{E}_{\text{stat}} [f(X)] \right)}_{\text{(II)}} \\
&\quad + \underbrace{\frac{T-\tau}{T} \left(\mathbb{E}_{\text{stat}} [f(X)] - \mathbb{E} \left[\frac{1}{T-\tau} \sum_{t=\tau}^{T-1} f(X_t) \right] \right)}_{\text{(III)}}.
\end{aligned}$$

In the follows, we upper bound terms (I), (II), and (III), respectively.

Upper Bounding Term (I). Since $\{\{X_t^i\}_{t=0}^{T-1}\}_{i \in [N]}$ are i.i.d. across each trajectory, by standard Bernstein's inequality, with probability at least $1 - \delta$, we have

$$\begin{aligned}
\text{(I)} &\leq \frac{2f_{\max}}{3N} \log \frac{2}{\delta} + 4 \sqrt{\frac{4}{N} \mathbb{E} \left[\left(\frac{1}{\tau} \sum_{t=0}^{\tau-1} f(X_t) \right)^2 \right] \log \frac{2}{\delta}} \\
&= \frac{2f_{\max}}{3N} \log \frac{2}{\delta} + 4 \sqrt{\frac{4}{N} \mathbb{E} \left[\frac{1}{\tau} \sum_{t=0}^{\tau-1} f(X_t)^2 \right] \log \frac{2}{\delta}}.
\end{aligned} \tag{C.8.5}$$

Upper Bounding Term (II). We consider an auxiliary dataset $\{\{\tilde{X}_t^i\}_{t=0}^{T-1}\}_{i \in [N]}$, where the i -th trajectory $\{\tilde{X}_t^i\}_{t=0}^{T-1}$ is sampled such that $\tilde{X}_0^i \sim p_{\text{stat}}$. Here p_{stat} is the stationary

distribution of $\{X_t\}_{t \geq 0}$. Similarly, we define the following quantity,

$$(\tilde{\text{II}}) = \frac{1}{N(T - \tau)} \sum_{i \in [N]} \sum_{t=\tau}^{T-1} f(\tilde{X}_t^i) - \mathbb{E}_{\text{stat}} [f(X)].$$

Now, for any $x \geq 0$, we upper bound the difference $\mathbb{P}(\text{II}) \geq x) - \mathbb{P}(\tilde{\text{II}}) \geq x)$ as follows,

(C.8.6)

$$\mathbb{P}(\text{II}) \geq x) - \mathbb{P}(\tilde{\text{II}}) \geq x) \leq N \sum_{t=\tau}^{T-1} \mathbb{E} [\|p_t(\cdot | X_0) - p_{\text{stat}}(\cdot)\|_{\text{TV}}] \leq c \cdot NT \exp(-\kappa\tau).$$

Thus, by (C.8.6), to upper bound $\mathbb{P}(\text{II}) \geq x)$, it suffices to upper bound $\mathbb{P}(\tilde{\text{II}}) \geq x)$.

To upper bound $\mathbb{P}(\tilde{\text{II}}) \geq x)$, we take $T - \tau = 2ks$, where k and s are two positive integers for the simplicity of presentation. We partition the set $\{\tau, \tau + 1, \dots, T - 1\}$ as follows,

$$J_1 = \{\tau, \tau + 1, \dots, \tau + s - 1\}, \quad J_2 = \{\tau + s, \tau + s + 1, \dots, \tau + 2s - 1\}, \quad \dots,$$

$$J_{2k-1} = \{T - 2s, T - 2s + 1, \dots, T - s - 1\}, \quad J_{2k} = \{T - s, T - s + 1, \dots, T - 1\}.$$

Under such a partition, we see that $\cup_{\ell \in [2k]} J_\ell = \{\tau, \tau + 1, \dots, T - 1\}$ and $J_\ell \cap J_{\ell'} = \emptyset$ for any $\ell \neq \ell'$. Also, for any $i \in [N]$, we define

$$Z_\ell^i = (\tilde{X}_t^i)_{t \in J_\ell}$$

for any $\ell \in [2k]$. Now, for any $i \in [N]$, by Lemma C.7.5, there exists a sequence $\{W_\ell^i\}_{\ell \in [2k]}$, where $W_\ell^i = (\tilde{Y}_t^i)_{t \in J_\ell}$ such that

- (1) $\{W_\ell^i\}_{\ell \in [2k]}$ are independent;

- (2) for any $\ell \in [2k]$, W_ℓ^i and Z_ℓ^i have the same distribution;
- (3) for any $\ell \in [2k]$, $\mathbb{P}(W_\ell^i \neq Z_\ell^i) = \beta(\sigma(\{Z_{\ell'}^i\}_{\ell' \in [\ell-1]}), \sigma(\{Z_\ell^i\}))$.

Note that the following inclusion relation holds,

$$\left\{ \frac{1}{s} \sum_{t \in J_\ell} f(\tilde{X}_t^i) - \mathbb{E}_{\text{stat}}[f(X)] \geq x_\ell \right\} \subseteq \left\{ \frac{1}{s} \sum_{t \in J_\ell} f(\tilde{Y}_t^i) - \mathbb{E}_{\text{stat}}[f(X)] \geq x_\ell \right\} \cup \{W_\ell^i \neq Z_\ell^i\}$$

for any $x_\ell \in \mathbb{R}$. Thus, we have

$$\begin{aligned} \mathbb{P}(\tilde{\text{II}} \geq x) &\leq \mathbb{P}\left(\frac{1}{kN} \sum_{i \in [N], \ell \text{ is odd}} \frac{1}{s} \sum_{t \in J_\ell} f(\tilde{X}_t^i) - \mathbb{E}_{\text{stat}}[f(X)] \geq x\right) \\ &\quad + \mathbb{P}\left(\frac{1}{kN} \sum_{i \in [N], \ell \text{ is even}} \frac{1}{s} \sum_{t \in J_\ell} f(\tilde{X}_t^i) - \mathbb{E}_{\text{stat}}[f(X)] \geq x\right) \\ &\leq \mathbb{P}\left(\frac{1}{kN} \sum_{i \in [N], \ell \text{ is odd}} \frac{1}{s} \sum_{t \in J_\ell} f(\tilde{Y}_t^i) - \mathbb{E}_{\text{stat}}[f(X)] \geq x\right) \\ &\quad + \mathbb{P}\left(\frac{1}{kN} \sum_{i \in [N], \ell \text{ is even}} \frac{1}{s} \sum_{t \in J_\ell} f(\tilde{Y}_t^i) - \mathbb{E}_{\text{stat}}[f(X)] \geq x\right) + \sum_{\ell \in [2k]} \mathbb{P}(W_\ell^i \neq Z_\ell^i) \\ &\leq \exp\left(-\frac{3k^2 N^2 x^2}{6 \sum_{i \in [N], \ell \text{ is odd}} \mathbb{E}\left[\left(\frac{1}{s} \sum_{t \in J_\ell} f(\tilde{Y}_t^i)\right)^2\right] + 2f_{\max} N k x}\right) \\ &\quad + \exp\left(-\frac{3k^2 N^2 x^2}{6 \sum_{i \in [N], \ell \text{ is even}} \mathbb{E}\left[\left(\frac{1}{s} \sum_{t \in J_\ell} f(\tilde{Y}_t^i)\right)^2\right] + 2f_{\max} N k x}\right) + 2k\beta(s), \end{aligned} \tag{C.8.7}$$

where we use Bernstein's inequality in the last line. Note that for any $(i, \ell) \in [N] \times [2k]$, we have

$$(C.8.8) \quad \mathbb{E} \left[\left(\frac{1}{s} \sum_{t \in J_\ell} f(\tilde{Y}_t^i) \right)^2 \right] = \mathbb{E}_{\text{stat}} [f(X)^2] \leq \mathbb{E} \left[\frac{1}{T-\tau} \sum_{t=\tau}^{T-1} f(X_t)^2 \right] + f_{\max}^2 T c \cdot \exp(-\kappa\tau).$$

Combining (C.8.7) and (C.8.8), we have

$$\mathbb{P} \left((\tilde{\text{II}}) \geq x \right) \leq 2 \exp \left(- \frac{3k^2 N^2 x^2}{6kN \left(\mathbb{E} \left[\frac{1}{T-\tau} \sum_{t=\tau}^{T-1} f(X_t)^2 \right] + f_{\max}^2 \beta(\tau) \right) + 2f_{\max} N k x} \right) + 2k\beta(s).$$

Similarly, we have

$$\mathbb{P} \left((\tilde{\text{II}}) \leq -x \right) \leq 2 \exp \left(- \frac{3k^2 N^2 x^2}{6kN \left(\mathbb{E} \left[\frac{1}{T-\tau} \sum_{t=\tau}^{T-1} f(X_t)^2 \right] + f_{\max}^2 \beta(\tau) \right) + 2f_{\max} N k x} \right) + 2k\beta(s).$$

Thus, we have

$$\begin{aligned} \mathbb{P} (|\text{II}| \leq x) &\leq 4 \exp \left(- \frac{3k^2 N^2 x^2}{6kN \left(\mathbb{E} \left[\frac{1}{T-\tau} \sum_{t=\tau}^{T-1} f(X_t)^2 \right] + f_{\max}^2 \beta(\tau) \right) + 2f_{\max} N k x} \right) \\ &\quad + 4k\beta(s) + c \cdot NT \exp(-\kappa\tau). \end{aligned}$$

Now, by taking $s = 3 \log(NT)/\kappa$, it holds with probability at least $1 - \delta$ with $c/(NT)^2 \leq \delta \leq 1$ that

$$(C.8.9) \quad |\text{II}| \leq \frac{2f_{\max}}{3NT\kappa} \log \frac{2}{\delta} + 4 \sqrt{\frac{4}{NT\kappa} \mathbb{E} \left[\frac{1}{T-\tau} \sum_{t=\tau}^{T-1} f(X_t)^2 \right] \log \frac{2}{\delta}}.$$

Upper Bounding Term (III). We observe that

$$(C.8.10) \quad |(\text{III})| \leq f_{\max} \cdot \sum_{t=\tau}^{T-1} c \cdot \exp(-\kappa t) \leq c \cdot \frac{f_{\max}}{N^2 T^2}.$$

Combining Everything. Now, by combining (C.8.5), (C.8.9), and (C.8.10), it holds with probability at least $1 - \delta$ with $c/(NT)^2 \leq \delta \leq 1$ that

$$\begin{aligned} & \frac{1}{NT} \sum_{i \in [N]} \sum_{t=0}^{T-1} f(X_t^i) - \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} f(X_t) \right] \\ & \leq \frac{24f_{\max}}{NT\kappa} \log \frac{2}{\delta} \log(NT) + 48 \sqrt{\frac{1}{NT\kappa} \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} f(X_t)^2 \right] \log \frac{2}{\delta} \log(NT)}, \end{aligned}$$

which concludes the proof of the theorem. \square

C.8.3. Proof of Lemma C.7.8

Proof. The proof follows from the proofs of Lemmas 4.1 and 4.2 in Geer et al. (2000).

First Inequality. By the optimality of \hat{p} , we have

$$\widehat{\mathbb{E}} [\log \hat{p}(Y | X)] \geq \widehat{\mathbb{E}} [\log p^*(Y | X)],$$

which implies that

$$\int \log \frac{\hat{p}}{p^*} d\tilde{p}^* d\tilde{\mu} \geq 0,$$

where \widehat{p}^* and $\widetilde{\mu}$ are empirical counterparts of p^* and μ . Now, by the concavity of $\log(\cdot)$, we have

$$\log \frac{\widehat{p} + p^*}{2p^*} \geq \frac{1}{2} \log \frac{\widehat{p}}{p^*} + \frac{1}{2} \log \frac{p^*}{p^*} = \frac{1}{2} \log \frac{\widehat{p}}{p^*}.$$

By the above two inequalities, we have

$$\begin{aligned} 0 &\leq \frac{1}{4} \int \log \frac{\widehat{p}}{p^*} d\widehat{p}^* d\widetilde{\mu} \leq \frac{1}{2} \int \log \frac{\widehat{p} + p^*}{2p^*} d\widehat{p}^* d\widetilde{\mu} \\ &= \frac{1}{2} \int \log \frac{\widehat{p} + p^*}{2p^*} (d\widehat{p}^* d\widetilde{\mu} - dp^* d\mu) + \frac{1}{2} \int \log \frac{\widehat{p} + p^*}{2p^*} dp^* d\mu \\ (C.8.11) \quad &= \left(\widehat{\mathbb{E}} - \mathbb{E} \right) [g_{\widehat{p}}] + \frac{1}{2} \int \log \frac{\widehat{p} + p^*}{2p^*} dp^* d\mu. \end{aligned}$$

In the meanwhile, by the fact that $\log z \leq 2(\sqrt{z} - 1)$ for any $z > 0$, we have

$$\begin{aligned} \frac{1}{2} \int \log \frac{\widehat{p} + p^*}{2p^*} dp^* d\mu &\leq \int \left(\sqrt{\frac{\widehat{p} + p^*}{2p^*}} - 1 \right) dp^* d\mu \\ &= \int \left(\sqrt{\frac{\widehat{p} + p^*}{2}} \cdot p^* - \frac{1}{2} p^* - \frac{1}{2} \cdot \frac{\widehat{p} + p^*}{2} \right) dy d\mu \\ &= - \int \frac{1}{2} \left(\sqrt{\frac{\widehat{p} + p^*}{2}} - \sqrt{p^*} \right)^2 dy d\mu \\ (C.8.12) \quad &= -H^2 \left(\frac{\widehat{p} + p^*}{2}, p^* \right). \end{aligned}$$

By combining (C.8.11) and (C.8.12), we conclude the proof of the first inequality.

Second&Third Inequality. We denote by $\bar{p} = (p + p^*)/2$ for any p . We note the following two facts,

$$\frac{p_1^{1/2} + p_2^{1/2}}{\bar{p}_1^{1/2} + \bar{p}_2^{1/2}} \leq \sqrt{2},$$

$$\left| \frac{p_1^{1/2}}{\bar{p}_1^{1/2}} - \frac{p_2^{1/2}}{\bar{p}_2^{1/2}} \right| \left(\bar{p}_1^{1/2} + \bar{p}_2^{1/2} \right) = |\bar{p}_1 - \bar{p}_2| = \left| \frac{p_1 - p_2}{2} \right| = \frac{1}{2} \left| p_1^{1/2} - p_2^{1/2} \right| \left(p_1^{1/2} + p_2^{1/2} \right).$$

Thus, we have

$$\left| \frac{p_1^{1/2}}{\bar{p}_1^{1/2}} - \frac{p_2^{1/2}}{\bar{p}_2^{1/2}} \right| = \frac{1}{2} \cdot \frac{p_1^{1/2} + p_2^{1/2}}{\bar{p}_1^{1/2} + \bar{p}_2^{1/2}} \cdot \left| p_1^{1/2} - p_2^{1/2} \right| \leq \frac{\sqrt{2}}{2} \cdot \left| p_1^{1/2} - p_2^{1/2} \right|,$$

which implies the second inequality. The third inequality can be proved in a similar way.

Forth Inequality. We note that

$$\begin{aligned} \|p_1(\cdot | x) - p_2(\cdot | x)\|_1 &= \int |p_1(y | x) - p_2(y | x)| \, dy \\ &= \int (p_1(y | x)^{1/2} - p_2(y | x)^{1/2}) (p_1(y | x)^{1/2} + p_2(y | x)^{1/2}) \, dy \\ &= \sqrt{\int (p_1(y | x)^{1/2} - p_2(y | x)^{1/2})^2 \, dy} \sqrt{\int (p_1(y | x)^{1/2} + p_2(y | x)^{1/2})^2 \, dy} \\ &\leq 2 \sqrt{\int (p_1(y | x)^{1/2} - p_2(y | x)^{1/2})^2 \, dy} = 2\sqrt{2}h(p_1(\cdot | x), p_2(\cdot | x)), \end{aligned}$$

which concludes the proof. \square

C.8.4. Proof of Theorem C.7.9

Proof. The proof follows from the proof of Theorem 7.4 in Geer et al. (2000). We define the events

$$\mathcal{E} = \{\omega \in \Omega: H^2(\widehat{p}, p^*) > \delta^2\}.$$

Conditioning on \mathcal{E} , we have

$$(C.8.13) \quad \left(\widehat{\mathbb{E}} - \mathbb{E}\right)[g_{\widehat{p}}] \geq H^2(\widehat{p}, p^*) \geq \frac{1}{16}H^2(\widehat{p}, p^*) > \frac{\delta^2}{16},$$

where the first two inequalities come from Lemma C.7.8. We further define

$$\mathcal{E}^\dagger = \left\{ \omega \in \Omega: \sup_{p \in \mathcal{P}: H^2(\bar{p}, p^*) > \delta^2/16} \left(\widehat{\mathbb{E}} - \mathbb{E}\right)[g_p] - H^2(\bar{p}, p^*) \geq 0 \right\}.$$

By (C.8.13) and the definitions of \mathcal{E} and \mathcal{E}^\dagger , we observe that $\mathcal{E} \subseteq \mathcal{E}^\dagger$. Thus, we only need to upper bound $\mathbb{P}(\mathcal{E}^\dagger)$. To do so, we use a peeling argument as follows. Let $L = \min\{\ell: 2^{\ell+1}\delta^2/16 > 1\}$. We observe that

$$(C.8.14) \quad \mathbb{P}(\mathcal{E}^\dagger) \leq \sum_{\ell=0}^L \mathbb{P}(\mathcal{E}_\ell^\dagger),$$

where

$$\mathcal{E}_\ell^\dagger = \left\{ \omega \in \Omega: \sup_{p \in \mathcal{P}_\ell} \left(\widehat{\mathbb{E}}[g_p] - \mathbb{E}[g_p]\right) \geq 2^\ell \delta^2/16 \right\}.$$

Here $\mathcal{P}_\ell = \{p \in \mathcal{P}: H^2(\bar{p}, p^*) \leq 2^{\ell+1}\delta^2/16\}$. To upper bound $\mathbb{P}(\mathcal{E}_\ell^\dagger)$, we introduce the following result.

Theorem C.8.1. Given a Markov chain $\{Z_t\}_{t \geq 0} \subset \mathcal{Z}$ satisfying Assumption C.7.3, and take

$$(C.8.15) \quad v \leq C_1 \sqrt{NTR^2}/K,$$

$$(C.8.16) \quad v \leq 8\sqrt{NTR},$$

$$(C.8.17) \quad v \geq C_0 \cdot \max \left\{ \int_{v/(2^6\sqrt{NT})}^R (\mathcal{H}_{B,K}(u, \mathcal{G}, P))^{1/2} du, R \right\},$$

$$(C.8.18) \quad v \geq C_2/(NT)^2,$$

$$(C.8.19) \quad C_0^2 \geq C^2(C_1 + 1),$$

where $\mathcal{H}_{B,K}(u, \mathcal{G}, P)$ is the generalized entropy with bracketing. Then we have

$$\begin{aligned} & \mathbb{P} \left(\sup_{g \in \mathcal{G}} \sqrt{NT} \left(\frac{1}{NT} \sum_{i \in [N]} \sum_{t=0}^{T-1} g(Z_t^i) - \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} g(Z_t) \right] \right) \geq v \right) \\ & \leq \frac{4C}{\kappa} \exp \left(-\frac{v^2 \kappa}{18C^2(C_1 + 1)R^2 \log(NT)} \right) + \frac{2}{N^2 T^2}, \end{aligned}$$

where $\{Z_t^i\}_{t=0}^{T-1}$ is generated from the same distribution as $\{Z_t\}_{t \geq 0}$ for any $i \in [N]$.

Proof. See §C.8.5 for a detailed proof. □

To invoke Theorem C.8.1, we take

$$v = \sqrt{NT} \cdot 2^\ell \delta^2 / 16, \quad K = 1, \quad R = 2^{\ell/2} \delta, \quad C_1 = 15, \quad C = c/64, \quad C_0 = c/16, \quad C_2 = c.$$

It is easy to verify that (C.8.15), (C.8.16), (C.8.18), and (C.8.19) hold. For (C.8.17), since $\sqrt{NT}\delta_{NT}^2 \geq c\Psi(\delta_{NT})$, which implies that

$$\sqrt{NT} \geq c \cdot \frac{\Psi(\delta_{NT})}{\delta_{NT}^2} \geq c \cdot \frac{\Psi(2^{\ell/2}\delta)}{2^\ell \delta^2},$$

where we use the fact that $\Psi(\delta)/\delta^2$ is a non-increasing function of δ . Thus, we have

$$16a \geq c \cdot \max \left\{ \int_{v/(2^6\sqrt{NT})}^R (\mathcal{H}_{B,1}(u, \{g_p: p \in \mathcal{P}_\ell\}, \mu_0))^{1/2} du, R \right\},$$

which justifies (C.8.17) by noting that $K = 1$. Here, we use the fact that

$$\mathcal{H}_{B,1}(u, \{g_p: p \in \mathcal{P}_\ell\}, P) \leq H_B \left(\frac{u}{\sqrt{2}}, \{\bar{p}^{1/2}: p \in \mathcal{P}_\ell\} \right).$$

Thus, by using Theorem C.8.1, we have

$$\mathbb{P}(\mathcal{E}_\ell^\dagger) \leq \frac{c}{\kappa} \exp \left(-\frac{NT\kappa 2^\ell \delta^2}{c^2 \log(NT)} \right) + \frac{2}{N^2 T^2}.$$

Further, by combining (C.8.14), we have

$$\mathbb{P}(\mathcal{E}^\dagger) \leq \frac{c}{\kappa} \exp \left(-\frac{NT\kappa \delta^2}{c^2} \right) + \frac{c}{N^2 T^2} \log \frac{4}{\delta},$$

which concludes the proof of the theorem. \square

C.8.5. Proof of Theorem C.8.1

Proof. We take $\tau = \min\{T, 3/\kappa \cdot \log(\mathcal{G}_{\max} NT)\}$, where $\mathcal{G}_{\max} = \max\{\max_{g \in \mathcal{G}} \max_{z \in \mathcal{Z}} g(z), 1\}$, and denote by $\mathbb{E}_{\text{stat}}[\cdot]$ the expectation taken with respect to the stationary distribution of

$\{Z_t\}_{t \geq 0}$. We have the following decomposition,

$$\begin{aligned}
& \mathbb{P} \left(\sup_{g \in \mathcal{G}} \sqrt{NT} \left(\frac{1}{NT} \sum_{i \in [N]} \sum_{t=0}^{T-1} g(Z_t^i) - \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} g(Z_t) \right] \right) \geq v \right) \\
&= \mathbb{P} \left(\sup_{g \in \mathcal{G}} \frac{\tau}{T} \left(\frac{1}{N\tau} \sum_{i \in [N]} \sum_{t=0}^{\tau-1} g(Z_t^i) - \mathbb{E} \left[\frac{1}{\tau} \sum_{t=0}^{\tau-1} g(Z_t) \right] \right) \right. \\
&\quad \left. + \frac{T-\tau}{T} \left(\frac{1}{N(T-\tau)} \sum_{i \in [N]} \sum_{t=\tau}^{T-1} g(Z_t^i) - \mathbb{E}_{\text{stat}}[g(Z)] \right) \right. \\
&\quad \left. + \frac{T-\tau}{T} \left(\mathbb{E}_{\text{stat}}[g(Z)] - \mathbb{E} \left[\frac{1}{T-\tau} \sum_{t=\tau}^{T-1} g(Z_t) \right] \right) \geq \frac{v}{\sqrt{NT}} \right) \\
&\leq \mathbb{P} \left(\sup_{g \in \mathcal{G}} \frac{\tau}{T} \left(\frac{1}{N\tau} \sum_{i \in [N]} \sum_{t=0}^{\tau-1} g(Z_t^i) - \mathbb{E} \left[\frac{1}{\tau} \sum_{t=0}^{\tau-1} g(Z_t) \right] \right) \geq \frac{v}{3\sqrt{NT}} \right) \\
&\tag{C.8.20} \\
&\quad + \mathbb{P} \left(\sup_{g \in \mathcal{G}} \frac{T-\tau}{T} \left(\frac{1}{N(T-\tau)} \sum_{i \in [N]} \sum_{t=\tau}^{T-1} g(Z_t^i) - \mathbb{E}_{\text{stat}}[g(Z)] \right) \geq \frac{v}{3\sqrt{NT}} \right),
\end{aligned}$$

where the last inequality comes from the fact that

$$\begin{aligned}
& \sup_{g \in \mathcal{G}} \frac{T-\tau}{T} \left(\mathbb{E}_{\text{stat}}[g(Z)] - \mathbb{E} \left[\frac{1}{T-\tau} \sum_{t=\tau}^{T-1} g(Z_t) \right] \right) \\
&\leq \sup_{g \in \mathcal{G}} \frac{T-\tau}{T} \int g(z) \frac{1}{T-\tau} \sum_{t=\tau}^{T-1} \left(p_{\text{stat}}(z) - \int p_t(z | z_0) d\zeta(z_0) \right) dz \\
&\leq \frac{1}{T} \mathcal{G}_{\max} \sum_{t=\tau}^{T-1} c \cdot \exp(-\kappa t) \leq \mathcal{G}_{\max} c \exp(-\kappa\tau) \leq \frac{v}{3\sqrt{NT}},
\end{aligned}$$

where we use the fact that $v \geq C_2/(NT)^2$ for some constant C_2 . Thus, we only need to upper bound the two terms on the RHS of (C.8.20). We first introduce the following supporting results.

Lemma C.8.2. Take

$$\begin{aligned} v &\leq C_1\sqrt{n}R^2/K, \\ v &\leq 8\sqrt{n}R, \\ v &\geq C_0 \cdot \max \left\{ \int_{v/(2^6\sqrt{n})}^R (\log \mathcal{N}_{B,K}(u, \mathcal{G}, P))^{1/2} du, R \right\}, \\ C_0^2 &\geq C^2(C_1 + 1). \end{aligned}$$

Then we have

$$\mathbb{P} \left(\sup_{g \in \mathcal{G}} \left| \sqrt{n} \left(\frac{1}{n} \sum_{i \in [n]} g(Z_i) - \mathbb{E}[g(Z)] \right) \right| \geq v \right) \leq C \exp \left(-\frac{v^2}{C^2(C_1 + 1)R^2} \right),$$

where $\{Z_i\}_{i \in [n]}$ are i.i.d. samples drawn the same distribution as Z .

Proof. See Theorem 5.11 in Geer et al. (2000) for a detailed proof. \square

Lemma C.8.3. Given a β -mixing sequence $\{Z_t\}_{t \geq 0} \subset \mathcal{Z}$ with coefficient $\beta(t)$ for any $t \geq 0$. There exists a sequence $\{Z_t^*\}_{t=0}^{T-1} \subset \mathcal{Z}$ and a set \mathcal{J} such that

- (1) \mathcal{J} is a partition of $\{0, 1, \dots, T-1\}$, i.e., $\cup_{J \in \mathcal{J}} J = \{0, 1, \dots, T-1\}$ and $J_1 \cap J_2 = \emptyset$ for any $J_1, J_2 \in \mathcal{J}$;
- (2) for any $0 \leq t \leq T-1$, Z_t^* and Z_t have the same distribution;
- (3) for any $J \in \mathcal{J}$, $\{Z_t^*\}_{t \in J}$ is an independent sequence;

(4) it holds for any $u \in \mathbb{R}$ that

$$\begin{aligned} & \mathbb{P} \left(\sup_{g \in \mathcal{G}} \frac{1}{T} \sum_{t=0}^{T-1} g(Z_t) - \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} g(Z_t) \right] \geq u \right) \\ & \leq \sum_{J \in \mathcal{J}} \mathbb{P} \left(\sup_{g \in \mathcal{G}} \frac{1}{|J|} \sum_{t \in J} g(Z_t^*) - \mathbb{E} \left[\frac{1}{|J|} \sum_{t \in J} g(Z_t) \right] \geq u \right) \\ & \quad + \sum_{J \in \mathcal{J}} |J| \cdot \beta(\min\{|t_1 - t_2| : t_1 \neq t_2 \in J\}). \end{aligned}$$

Proof. See Theorem 2.11 in Barrera and Gobet (2021) for a detailed proof. \square

We upper bound two terms on the RHS of (C.8.20) as follows.

Upper Bounding the First Term on the RHS of (C.8.20). To upper bound the first term, we invoke Lemma C.8.2. Since the sequence

$$\left\{ \frac{1}{\tau} \sum_{t=0}^{\tau-1} g(Z_t^i) \right\}_{i \in [n]}$$

is i.i.d., we have

$$\begin{aligned} & \mathbb{P} \left(\sup_{g \in \mathcal{G}} \frac{\tau}{T} \left(\frac{1}{N\tau} \sum_{i \in [N]} \sum_{t=0}^{\tau-1} g(Z_t^i) - \mathbb{E} \left[\frac{1}{\tau} \sum_{t=0}^{\tau-1} g(Z_t) \right] \right) \geq \frac{v}{3\sqrt{NT}} \right) \\ \text{(C.8.21)} \quad & \leq C \exp \left(-\frac{v^2 T}{9\tau^2 C^2 (C_1 + 1) R^2} \right) \leq C \exp \left(-\frac{v^2}{C^2 (C_1 + 1) R^2} \right). \end{aligned}$$

Upper Bounding the Second Term on the RHS of (C.8.20). To upper bound the second term, we note that

$$\mathbb{P} \left(\sup_{g \in \mathcal{G}} \frac{T - \tau}{T} \left(\frac{1}{N(T - \tau)} \sum_{i \in [N]} \sum_{t=\tau}^{T-1} g(Z_t^i) - \mathbb{E}_{\text{stat}} [g(Z)] \right) \geq \frac{v}{3\sqrt{NT}} \right) = \text{(i)} + \text{(ii)},$$

where

$$\begin{aligned}
\text{(i)} &= \mathbb{P} \left(\sup_{g \in \mathcal{G}} \frac{T - \tau}{T} \left(\frac{1}{N(T - \tau)} \sum_{i \in [N]} \sum_{t=\tau}^{T-1} g(Z_t^i) - \mathbb{E}_{\text{stat}} [g(Z)] \right) \geq \frac{v}{3\sqrt{NT}} \right) \\
&\quad - \mathbb{P} \left(\sup_{g \in \mathcal{G}} \frac{T - \tau}{T} \left(\frac{1}{N(T - \tau)} \sum_{i \in [N]} \sum_{t=\tau}^{T-1} g(\tilde{Z}_t^i) - \mathbb{E}_{\text{stat}} [g(Z)] \right) \geq \frac{v}{3\sqrt{NT}} \right), \\
\text{(ii)} &= \mathbb{P} \left(\sup_{g \in \mathcal{G}} \frac{T - \tau}{T} \left(\frac{1}{N(T - \tau)} \sum_{i \in [N]} \sum_{t=\tau}^{T-1} g(\tilde{Z}_t^i) - \mathbb{E}_{\text{stat}} [g(Z)] \right) \geq \frac{v}{3\sqrt{NT}} \right).
\end{aligned}$$

Here $\{\tilde{Z}_t^i\}_{t=0}^{T-1}$ are an auxiliary sequence for any $i \in [N]$, where \tilde{Z}_0^i is sampled from the stationary distribution of the sequence $\{Z_t\}_{t \geq 0}$. To upper bound (i), we note that

$$\text{(C.8.22)} \quad \text{(i)} \leq \sum_{i \in [N]} \sum_{t=\tau}^{T-1} c \cdot \exp(-\kappa t) \leq NTc \exp(-\kappa\tau) \leq \frac{1}{N^2 T^2}.$$

To upper bound (ii), we invoke Lemma C.8.3 by taking

$$\mathcal{J} = \{J_1, J_2, \dots, J_s\}, \quad J_j = \{\tau + j - 1, \tau + j + s - 1, \dots, T - s + j\} \text{ for any } j \in [s].$$

Then there exists a sequence $\{\{\tilde{Z}_t^{i*}\}_{t=\tau}^{T-1}\}_{i \in [N]}$ such that \tilde{Z}_t^{i*} and \tilde{Z}_t^i have the same distribution for any (i, t) ; $\{\{\tilde{Z}_t^*\}_{t \in J}\}_{i \in [N]}$ are independent; and it holds for any $u \in \mathbb{R}$ that

$$\begin{aligned}
\text{(ii)} &\leq \sum_{j=1}^s \mathbb{P} \left(\sup_{g \in \mathcal{G}} \frac{s}{N(T - \tau)} \sum_{i \in [N]} \sum_{t \in J_j} g(\tilde{Z}_t^{i*}) - \mathbb{E}_{\text{stat}} [g(Z)] \geq \frac{v}{3\sqrt{NT}} \frac{T}{T - \tau} \right) + (T - \tau) \cdot \beta(s) \\
&\leq s \cdot C \exp \left(-\frac{v^2}{9sC^2(C_1 + 1)R^2} \right) + (T - \tau)\beta(s),
\end{aligned}$$

where we use Lemma C.8.2 in the last inequality. Now, by taking $s = \min\{T, 3/\kappa \cdot \log(NT)\}$, we have

$$(C.8.23) \quad (ii) \leq \frac{3C}{\kappa} \exp\left(-\frac{v^2\kappa}{18C^2(C_1+1)R^2 \log(NT)}\right) + \frac{1}{N^2T^2},$$

where we use Lemma C.7.4 to upper bound $\beta(s)$. Now, by combining (C.8.22) and (C.8.23), we have

$$(C.8.24) \quad \begin{aligned} & \mathbb{P}\left(\sup_{g \in \mathcal{G}} \frac{T-\tau}{T} \left(\frac{1}{N(T-\tau)} \sum_{i \in [N]} \sum_{t=\tau}^{T-1} g(Z_t^i) - \mathbb{E}_{\text{stat}}[g(Z)]\right) \geq \frac{v}{3\sqrt{NT}}\right) \\ & \leq \frac{3C}{\kappa} \exp\left(-\frac{v^2\kappa}{18C^2(C_1+1)R^2 \log(NT)}\right) + \frac{2}{N^2T^2}. \end{aligned}$$

Combining Everything. By plugging (C.8.21) and (C.8.24) into (C.8.20), we have

$$\begin{aligned} & \mathbb{P}\left(\sup_{g \in \mathcal{G}} \sqrt{NT} \left(\frac{1}{NT} \sum_{i \in [N]} \sum_{t=0}^{T-1} g(Z_t^i) - \mathbb{E}\left[\frac{1}{T} \sum_{t=0}^{T-1} g(Z_t)\right]\right) \geq v\right) \\ & \leq \frac{4C}{\kappa} \exp\left(-\frac{v^2\kappa}{18C^2(C_1+1)R^2 \log(NT)}\right) + \frac{2}{N^2T^2}, \end{aligned}$$

which concludes the proof of the theorem. \square