

# Statistical Inference of the Value Function for Reinforcement Learning in Infinite-Horizon Settings

Chengchun Shi<sup>†</sup>, Sheng Zhang<sup>‡</sup>, Wenbin Lu<sup>‡</sup> and Rui Song<sup>‡</sup>

<sup>†</sup>*London School of Economics and Political Science, and* <sup>‡</sup>*North Carolina State University*

## Abstract

Reinforcement learning is a general technique that allows an agent to learn an optimal policy and interact with an environment in sequential decision making problems. The goodness of a policy is measured by its value function starting from some initial state. The focus of this paper is to construct confidence intervals (CIs) for a policy's value in infinite horizon settings where the number of decision points diverges to infinity. We propose to model the action-value state function (Q-function) associated with a policy based on series/sieve method to derive its confidence interval. When the target policy depends on the observed data as well, we propose a **SequentiAl Value Evaluation** (SAVE) method to recursively update the estimated policy and its value estimator. As long as either the number of trajectories or the number of decision points diverges to infinity, we show that the proposed CI achieves nominal coverage even in cases where the optimal policy is not unique. Simulation studies are conducted to back up our theoretical findings. We apply the proposed method to a dataset from mobile health studies and find that reinforcement learning algorithms could help improve patient's health status. A Python implementation of the proposed procedure is available at <https://github.com/shengzhang37/SAVE>.

**Key Words:** Confidence interval; Value function; Reinforcement learning; Infinite horizons; Bidirectional Asymptotics.

## 1 Introduction

Reinforcement learning (RL) is a general technique that allows an agent to learn and interact with an environment. A policy defines the agent's way of behaving. It maps the states of environments to a set of actions to be chosen from. RL algorithms have made tremendous achievements and found extensive applications in video games (Silver et al., 2016), robotics (Kormushev et al., 2013), bidding (Jin et al., 2018), ridesharing (Xu

et al., 2018), etc. In particular, a number of RL methods have been proposed in precision medicine, to derive an optimal policy as a set of sequential treatment decision rules that optimize patients' clinical outcomes over a fixed period of time (finite horizon). References include Murphy (2003); Zhang et al. (2013); Zhao et al. (2015); Shi et al. (2018a,b); Zhang et al. (2018), to name a few.

Mobile health (or mHealth) technology has recently emerged due to the use of mobile devices such as mobile phones, tablet computers or wearable devices in health care. It allows health-care providers to communicate with patients and manage their illness in real time. It also collects rich longitudinal data (e.g., through mobile health apps) that can be used to estimate the optimal policy. Data from mHealth applications differ from those in finite horizon settings in that the number of treatment decision points for each patient is not necessarily fixed (infinite horizon) while the total number of patients could be limited. Take the OhioT1DM dataset (Marling and Bunescu, 2018) as an example. It contains data for six patients with type 1 diabetes. For all patients, their continuous glucose monitoring (CGM) blood glucose levels, insulin doses including bolus and basal rates, self-reported times of meals and exercises are continually measured and recorded for eight weeks. Developing an optimal policy as functions of these time-varying covariates could potentially assist these patients in improving their health status.

In this paper, we focus on the infinite horizon setting where the data generating process is modeled by a Markov decision process (MDP, Puterman, 1994). Specifically, at each time point, the agent selects an action based on the observed state. The system responds by giving the decision maker a corresponding outcome and moving into a new state in the next time step. This model is generally applicable to sequential decision making, including applications from mHealth, games, robotics, ridesharing, etc. After a policy is being proposed, it is important to examine its benefit prior to recommending it for practical use. The goodness of a policy is quantified by its (state) value function, corresponding to the discounted cumulative reward that the agent receives on average, starting from some initial state. The inference of the value function helps a decision maker to evaluate the impact of implementing a policy when the environment is in a certain state. In some applications, it is also important to evaluate the integrated value of a policy aggregated over different initial states. For example, in medical studies, one might wish to know the

mean outcome of patients in the population. The integrated value could thus be used as a criterion for comparing different policies.

In statistics literature, a few methods have been proposed to estimate the optimal policy in infinite horizons. Ertefaie and Strawderman (2018) proposed a variant of gradient Q-learning method. Lockett et al. (2019) proposed a V-learning to directly search the optimal policy among a restricted class of policies. Inference of the value function under a generic (data-dependent) policy has not been studied in these papers. In the computer science literature, Thomas et al. (2015) and Jiang and Li (2016) proposed (augmented) inverse propensity-score weighted ((A)IPW) estimators for the the value function in infinite horizons and derived their associated CIs. However, these methods are not suitable for settings where only a limited number of trajectories (e.g., plays of a game or patients in medical studies) are available, since (A)IPW estimators become increasingly unstable as the number of decision points diverges to infinity. Recently, Kallus and Uehara (2019) proposed a double reinforcement learning (DRL) method that achieves consistent estimation of the value under a fixed policy even with limited number of trajectories. Their method computes a Q-function and a marginalized density ratio. Learning the density ratio is challenging in general and it remains difficult to investigate the goodness-of-fit of the estimated density ratio in practice.

The focus of this paper is to construct confidence intervals (CIs) for a (possibly data-dependent) policy's value function at a given state as well as its integrated value with respect to a given reference distribution. Our proposed CI is derived by estimating the state-action value function (Q-function) under the target policy. Similar to the value, the Q-function measures the discounted cumulative reward that the agent receives on average, starting from some initial state-action pair. We use series/sieve method to approximate the Q-function based on  $L$  basis functions, where  $L$  grows with the total number of observations. The advances of our proposed method are summarized as follows. First, the proposed inference method is generally applicable. Specifically, it can be applied to any fixed policy (either deterministic or random) and any data-dependent policy whose value converges at a certain rate. The latter includes policies estimated by general Q-learning type algorithms that learns an optimal Q-function from the observed data, such as gradient Q-learning (Maei et al., 2010; Ertefaie and Strawderman, 2018), fitted Q-iteration (see for example,

Ernst et al., 2005; Riedmiller, 2005), etc. See Section 3.2.4 for detailed illustrations.

Second, when applied to data-dependent policies, our method is valid in nonregular cases where the optimal policy is not uniquely defined. Inference without requiring the uniqueness of the optimal policy is extremely challenging even in the simpler finite-horizon settings (see the related discussions in Luedtke and van der Laan, 2016). The major challenge lies in that the estimated policy may not stabilize as sample size grows, making the variance of the value estimator difficult to estimate (see Section 3.2.1 for details). We achieve valid inference by proposing a **SequentiAl Value Evaluation** (SAVE) method that splits the data into several blocks and recursively update the estimated policy and its value estimator. It is worth mentioning that the data-splitting rule cannot be arbitrarily determined since the observations are time dependent in infinite horizon settings (see Section 3.2.2 for details).

Third, our CI is valid as long as either the number of trajectories  $n$  in the data, or the number of decision points  $T$  per trajectory diverges to infinity. It can thus be applied to a wide variety of real applications in infinite horizons ranging from the Framingham heart study (Tsao and Vasan, 2015) with over two thousand patients to the OhioT1DM dataset that contains eight weeks’ worth of data for six people. We also allow both  $n$  and  $T$  to approach infinity, which is the case in applications from video games. In contrast, CIs proposed by Thomas et al. (2015) and Jiang and Li (2016) require  $n$  to grow to infinity to achieve nominal coverage.

Lastly, we consider both off-policy and on-policy learning methods. In off-policy settings, CIs are derived based on historical data collected by a potentially different behavior policy. Off-policy evaluation is critical in situations where running the target policy could be expensive, risky or unethical. In on-policy settings, the estimated policy is recursively updated as batches of new observations arrive. To our knowledge, this is the first work on statistical inference of a data-dependent policy in on-policy settings in sequential decision making with infinite horizons.

To study the asymptotic properties of our proposed CI, we focus on tensor-product spline and wavelet series estimators. Our technical contributions are described as follows. First, we introduce a bidirectional-asymptotic framework that allows either  $n$  or  $T$  to approach infinity. Our major technical contribution is to derive a nonasymptotic error bound for the spectral norm of sums of mean zero random matrices formed by the data

transactions from MDP as a function of  $n$ ,  $T$  and  $L$  (see e.g., Lemma 3). This result is important in studying the limiting distribution of series estimators under such a theoretical framework.

Second, for policies that are estimated by Q-learning type algorithms such as the greedy gradient Q-learning, fitted Q-iteration and deep Q-network (Mnih et al., 2015), we relate the convergence rate of their values to the prediction error of the corresponding estimated Q-functions. We show in Theorems 3 and 4 that the values can converge at faster rates than the estimated Q-functions under certain margin type conditions on the optimal Q-function. To the best of our knowledge, these findings have not been discovered in the reinforcement-learning literature. Our theorems form a basis for researchers to study the value properties of Q-learning type algorithms. Moreover, our theoretical results are consistent with findings in point treatment studies where there is only one single decision point (see e.g., Qian and Murphy, 2011; Luedtke and van der Laan, 2016). However, the derivation of Theorems 3 and 4 is more involved since the value function in our settings is an infinite series involving both immediate and future rewards.

Third, when these basis functions are used, we mathematically characterize the approximation error of the Q-function as a function of  $L$ , the dimension of the state variables, and the smoothness of the Markov transition function and the conditional mean of the immediate reward as a function of the state-action pair. This offers some guidance to practitioners on the choice of the number of basis functions  $L$ , when some prior knowledge on the degree of smoothness of the aforementioned functions are available.

The rest of the paper is organized as follows. We introduce the model setup in Section 2. In Sections 3 and 4, we present the proposed off-policy and on-policy evaluation methods, respectively. Simulation studies are conducted to evaluate the empirical performance of the proposed inference methods in Section 5. We apply the proposed inference method to the OhioT1DM dataset in Section 6, Finally, we conclude our paper by a discussion section.

## 2 Optimal policy in infinite-horizon settings

We begin by introducing the notion of the optimal policy, the Q-function and the value function in infinite-horizon settings. Let  $X_{0,t} \in \mathbb{X}$  be the time-varying covariates collected

at time point  $t$ ,  $A_{0,t} \in \mathcal{A}$  denote the action taken at time  $t$ , and  $Y_{0,t}$  stand for the immediate reward observed. Here,  $\mathbb{X}$  and  $\mathcal{A}$  denote the state and action space, respectively. We assume  $\mathbb{X}$  is a subspace of  $\mathbb{R}^d$  where  $d$  is the number of state vectors and  $\mathcal{A}$  is a discrete space  $\{0, 1, \dots, m-1\}$  where  $m$  denotes the number of actions. Suppose the system satisfies the following Markov assumption (MA),

$$\Pr(X_{0,t+1} \in \mathcal{B} | X_{0,t} = x, A_{0,t} = a, \{Y_{0,j}\}_{0 \leq j < t}, \{X_{0,j}\}_{0 \leq j < t}, \{A_{0,j}\}_{0 \leq j < t}) = \mathcal{P}(\mathcal{B} | x, a),$$

for some transition function  $\mathcal{P}$ . Here,  $\mathcal{P}$  defines the next state distribution conditional on the current state-action pair. Moreover, suppose the following conditional mean independence assumption (CMIA) holds

$$\begin{aligned} \mathbb{E}(Y_{0,t} | X_{0,t} = x, A_{0,t} = a, \{Y_{0,j}\}_{0 \leq j < t}, \{X_{0,j}\}_{0 \leq j < t}, \{A_{0,j}\}_{0 \leq j < t}) \\ = \mathbb{E}(Y_{0,t} | X_{0,t} = x, A_{0,t} = a) = r(x, a), \end{aligned}$$

for some reward function  $r$ . By MA, CMIA automatically holds when  $Y_{0,t}$  is a deterministic function of  $X_{0,t}$ ,  $A_{0,t}$  and  $X_{0,t+1}$  that measures the system's status at time  $t+1$ . The latter is satisfied in our real data application (see Section 6 for details) and is commonly assumed in the reinforcement learning literature. CMIA is thus weaker than this condition. MA and CMIA are important to guarantee the existence of an optimal policy (see (2.1)) and derive the bidirectional-asymptotic theory of the proposed CI (see the discussions below Theorem 1). We assume both assumptions hold throughout this paper.

In the following, we focus on the class of stationary policies that map the covariate space  $\mathbb{X}$  to probability mass functions on  $\mathcal{A}$ . Let  $\pi(\cdot | \cdot)$  denote such a policy. It satisfies  $\pi(a|x) \geq 0$ , for any  $a \in \mathcal{A}$ ,  $x \in \mathbb{X}$  and  $\sum_{a \in \mathcal{A}} \pi(a|x) = 1$ , for any  $x \in \mathbb{X}$ . For a deterministic policy, we have  $\pi(a|x) \in \{0, 1\}$ , for any  $a \in \mathcal{A}$ ,  $x \in \mathbb{X}$ . Under  $\pi$ , a decision maker will set  $A_{0,t} = a$  with probability  $\pi(a|X_{0,t})$  at time  $t$ . For such a policy and a given discounted factor  $0 \leq \gamma < 1$ , let  $V(\pi; x)$  denote the value function

$$V(\pi; x) = \sum_{t \geq 0} \gamma^t \mathbb{E}^\pi(Y_{0,t} | X_{0,0} = x),$$

where the expectation  $\mathbb{E}^\pi$  is taken by assuming that the system follows the policy  $\pi$ . The rate  $\gamma$  reflects a trade-off between immediate and future rewards. If  $\gamma = 0$ , the agent tends

to choose actions that maximize the immediate reward. As  $\gamma$  increases, the agent will consider future rewards more seriously. Under CMIA, we have

$$V(\pi; x) = \sum_{t \geq 0} \gamma^t \mathbb{E}^\pi \{ \mathbb{E}(Y_{0,t} | X_{0,t}, A_{0,t}) | X_{0,0} = x \} = \sum_{t \geq 0} \gamma^t \sum_{a \in \mathcal{A}} \mathbb{E}^\pi \{ \pi(a | X_{0,t}) r(X_{0,t}, a) | X_{0,0} = x \}.$$

Similar to Theorem 6.2.12 of Puterman (1994), we can show under the given conditions that there exists at least one optimal policy  $\pi^{opt}$  that satisfies

$$V(\pi^{opt}; x) \geq V(\pi; x), \quad \forall \pi, x. \quad (2.1)$$

To better understand  $\pi^{opt}$ , we introduce the state-action function (Q-function) under a policy  $\pi$  as

$$Q(\pi; x, a) = \sum_{t \geq 0} \gamma^t \mathbb{E}^\pi (Y_{0,t} | X_{0,0} = x, A_{0,0} = a).$$

Let  $Q^{opt}$  denote the optimal Q-function, i.e,  $Q^{opt}(\cdot, \cdot) = \sup_\pi Q(\pi; \cdot, \cdot)$ . It can be shown that  $\pi^{opt}$  satisfies

$$\pi^{opt}(a|x) = 0 \quad \text{if } a \notin \arg \max_{a'} Q^{opt}(x, a'), \quad \forall x, a.$$

There exist infinitely many optimal policies when  $\arg \max_a Q^{opt}(x, a)$  is not unique for some  $x \in \mathbb{X}$ . Let  $\Pi^{opt}$  denote the set consisting of all these optimal policies. Define

$$\pi_0^{opt}(a|x) = \begin{cases} 1, & \text{if } a = \text{sarg} \max_{a'} Q^{opt}(x, a'), \\ 0, & \text{otherwise,} \end{cases} \quad (2.2)$$

where  $\text{sarg} \max$  denotes the smallest maximizer when the  $\arg \max$  is not unique. Such a deterministic optimal policy may be appealing in medical studies. For example, in optimal dose studies, it is preferred to assign each patient the smallest optimal dose level to avoid toxicity.

## 3 Off-policy evaluation

### 3.1 Inference of the value under a fixed policy

Let  $n$  denote the number of trajectories in the dataset. For the  $i$ -th trajectory, let  $\{A_{i,t}\}_{t \geq 0}$ ,  $\{X_{i,t}\}_{t \geq 0}$  and  $\{Y_{i,t}\}_{t \geq 0}$  denote the sequence of actions, states and rewards, respectively. It

is worth mentioning that the time points are not necessarily homogeneous across different trajectories. Suppose the data are generated according to a fixed policy  $b(\cdot|\cdot)$ , better known as the behavior policy such that

$$\{(X_{1,t}, A_{1,t}, Y_{1,t})\}_{t \geq 0}, \{(X_{2,t}, A_{2,t}, Y_{2,t})\}_{t \geq 0}, \dots, \{(X_{n,t}, A_{n,t}, Y_{n,t})\}_{t \geq 0},$$

are i.i.d copies of  $\{(X_{0,t}, A_{0,t}, Y_{0,t})\}_{t \geq 0}$ . The observed data can thus be summarized as  $\{(X_{i,t}, A_{i,t}, Y_{i,t}, X_{i,t+1})\}_{0 \leq t < T_i, 1 \leq i \leq n}$ , where  $T_i$  is the termination time of the  $i$ -th trajectory. The goal of off-policy evaluation is to learn the value under a target policy  $\pi(\cdot|\cdot)$ , possibly different from  $b(\cdot|\cdot)$ .

### 3.1.1 Modelling value or Q-function?

Luckett et al. (2019) showed that the value function satisfies

$$\mathbb{E} \left[ \frac{\pi(A_{i,t}, X_{i,t})}{b(A_{i,t}, X_{i,t})} \{Y_{i,t} + \gamma V(\pi; X_{i,t+1}) - V(\pi; X_{i,t})\} \middle| X_{i,t} \right] = 0. \quad (3.3)$$

Based on (3.3), they directly modelled the value function, constructed an estimator for the integrated value under their estimated optimal policy and proved that it is asymptotically normal (see Theorem 4.3, Luckett et al., 2019).

Following their procedure, for a fixed policy  $\pi$ , one might estimate  $V(\pi; \cdot)$  nonparametrically and construct the CI using the resulting estimate. However, such an approach might not be appropriate for policies that are discontinuous functions of the covariates. To better illustrate this, notice that  $V(\pi, \cdot)$  satisfies the following Bellman equation

$$V(\pi; x) = \sum_{a \in \mathcal{A}} \pi(a|x) \underbrace{\left\{ r(x, a) + \gamma \int_{x'} V(\pi; x') \mathcal{P}(dx'|x, a) \right\}}_{C(\pi; x, a)}. \quad (3.4)$$

When  $\mathcal{P}$  satisfies certain smoothness conditions (see Condition A1 below), we have

$$\begin{aligned} & \left| \int_{x'} V(\pi; x') \mathcal{P}(dx'|x_1, a) - \int_{x'} V(\pi; x') \mathcal{P}(dx'|x_2, a) \right| \\ & \leq \int_{x'} |V(\pi; x')| |\mathcal{P}(dx'|x_1, a) - \mathcal{P}(dx'|x_2, a)| \rightarrow 0, \quad \text{as } \|x_1 - x_2\|_2 \rightarrow 0, \end{aligned} \quad (3.5)$$

for any  $\pi$ . Suppose  $r(\cdot, a)$  is continuous for any  $a \in \mathcal{A}$ . Then  $C(\pi; x, a)$  is continuous in  $x$  for any  $\pi$  and  $a$ . When  $\pi$  is a non-continuous function of  $x$ , it follows from (3.4) that



$V(\pi; \cdot)$  is not continuous either. However, many nonparametric methods, such as kernel smoothers and series estimation, require the underlying function to possess certain degree of smoothness in order to achieve estimation consistency. Notice that any non-constant deterministic policy has jumps and is not continuous at certain points (such as the optimal policy  $\pi_0$  given in (2.2)). This poses significant challenges in performing inference to these policies.

To allow valid inference for both deterministic and random policies, we consider modelling the Q-function. Under CMIA, we have

$$Q(\pi; x, a) = \sum_{t \geq 0} \gamma^t \mathbb{E}^\pi \{r(X_{0,t}, A_{0,t}) | X_{0,0} = x, A_{0,0} = a\}.$$

This together with MA yields

$$\begin{aligned} Q(\pi; x, a) &= r(x, a) + \gamma \mathbb{E}^\pi \left[ \sum_{t \geq 0} \gamma^t \{ \mathbb{E}^\pi r(X_{0,t}, A_{0,t}) | X_{0,1}, A_{0,1} \} \middle| X_{0,0} = x, A_{0,0} = a \right] \\ &= r(x, a) + \gamma \mathbb{E}^\pi \{ Q(X_{0,1}, A_{0,1}) | X_{0,0} = x, A_{0,0} = a \}. \end{aligned}$$

As a result, the Q-function satisfies the following Bellman equation

$$Q(\pi; x, a) = r(x, a) + \gamma \sum_{a' \in \mathcal{A}} \int_{x'} Q(\pi; x', a') \pi(a' | x') \mathcal{P}(dx' | x, a). \quad (3.6)$$

Similar to (3.5), we can show the second term on the right-hand-side (RHS) of (3.5) is a smooth function of  $x$  for any  $\pi$  and  $a$ . When  $r(\cdot, a)$  is smooth, so is  $Q(\pi, \cdot, a)$ . To formally establish these results, we introduce the notion of  $p$ -smoothness (also known as Hölder smoothness with exponent  $p$ ) below.

Let  $h(\cdot)$  be an arbitrary function on  $\mathbb{X}$ . For a  $d$ -tuple  $\alpha = (\alpha_1, \dots, \alpha_d)^\top$  of nonnegative integers, let  $D^\alpha$  denote the differential operator:

$$D^\alpha h(x) = \frac{\partial^{|\alpha|} h(x)}{\partial x_1^{\alpha_1} \dots \partial x_d^{\alpha_d}}.$$

Here,  $x_j$  denotes the  $j$ -th element of  $x$ . For any  $p > 0$ , let  $\lfloor p \rfloor$  denote the largest integer that is smaller than  $p$ . Define the class of  $p$ -smooth functions as follows:

$$\Lambda(p, c) = \left\{ h : \sup_{\|\alpha\|_1 \leq \lfloor p \rfloor} \sup_{x \in \mathbb{X}} |D^\alpha h(x)| \leq c, \sup_{\|\alpha\|_1 = \lfloor p \rfloor} \sup_{\substack{x, y \in \mathbb{X} \\ x \neq y}} \frac{|D^\alpha h(x) - D^\alpha h(y)|}{\|x - y\|_2^{p - \lfloor p \rfloor}} \leq c \right\}.$$

When  $0 < p \leq 1$ , we have  $\lfloor p \rfloor = 0$ . It is equivalent to require  $h$  to satisfy  $\sup_{x,y} |h(x) - h(y)| / \|x - y\|_2^p \leq c$ . The notion of  $p$ -smoothness is thus reduced to the Hölder continuity.

For any  $x \in \mathbb{X}$ ,  $a \in \mathcal{A}$ , suppose the transition kernel  $\mathcal{P}(\cdot|x, a)$  is absolutely continuous with respect to the Lebesgue measure. Then there exists some transition density function  $q$  such that  $\mathcal{P}(dx'|x, a) = q(x'|x, a)dx'$ . We impose the following condition.

(A1.) There exist some  $p, c > 0$  such that  $r(\cdot, a), q(x'|\cdot, a) \in \Lambda(p, c)$  for any  $a \in \mathcal{A}, x' \in \mathbb{X}$ .

**Lemma 1** *Under A1, there exists some constant  $c' > 0$  such that  $Q(\pi; \cdot, a) \in \Lambda(p, c')$  for any policy  $\pi$  and  $a \in \mathcal{A}$ .*

Lemma 1 implies the Q-function has bounded derivatives up to order  $\lfloor p \rfloor$ . This motivates us to first estimate the Q-function and then derive the corresponding value estimators based on the relation  $V(\pi; x) = \sum_{a \in \mathcal{A}} \pi(a|x)Q(\pi; x, a)$ . By the Bellman equation (3.6), we can show the Q-function satisfies

$$\mathbb{E} \left[ \left\{ Y_{i,t} + \gamma \sum_{a \in \mathcal{A}} Q(\pi; X_{i,t+1}, a) \pi(a|X_{i,t+1}) - Q(\pi; X_{i,t}, A_{i,t}) \right\} \middle| X_{i,t}, A_{i,t} \right] = 0. \quad (3.7)$$

The above equation forms a basis of our methods to learn  $Q(\pi; \cdot, \cdot)$  (see details in the next section). In contrast to Equation (3.3), the sampling ratio  $\pi(a|x)/b(a|x)$  does not appear in (3.7). This is because  $A_{i,t}$  is the only sampling action and no further actions are involved in (3.7). As a result, our method does not require correct specification of the behavior policy. Nor do we need to estimate it from the observed dataset. This is another advantage of modelling the Q-function over the value.

### 3.1.2 Method

We describe our procedure in this section. We propose to approximate  $Q(\pi; \cdot, \cdot)$  based on linear sieves, which takes the form

$$Q(\pi; x, a) \approx \Phi_L^\top(x) \beta_{\pi,a}, \quad \forall x \in \mathbb{X}, a \in \mathcal{A},$$

where  $\Phi_L(\cdot) = \{\phi_{L,1}(\cdot), \dots, \phi_{L,L}(\cdot)\}^\top$  is a vector consisting of  $L$  sieve basis functions, such as splines or wavelet bases (see for example, Huang, 1998, for choices of basis functions).

We allow  $L$  to grow with the sample size to reduce the bias of the resulting estimates.

Under certain mild conditions, there exist some  $\{\beta_{\pi,a}^*\}_{a \in \mathcal{A}}$  that satisfy

$$\mathbb{E} \left\{ Y_{i,t} + \gamma \sum_{a \in \mathcal{A}} \Phi_L^\top(X_{i,t+1}) \beta_{\pi,a}^* \pi(a|X_{i,t+1}) - \Phi_L^\top(X_{i,t}) \beta_{\pi,a'}^* \right\} \Phi_L(X_{i,t}) \mathbb{I}(A_{i,t} = a') = 0,$$

for any  $a' \in \mathcal{A}$ . Recall that  $\mathcal{A} = \{0, 1, \dots, m-1\}$ . Define  $\beta_\pi^* = (\beta_{\pi,1}^{*T}, \dots, \beta_{\pi,m}^{*T})^\top$ ,

$$\boldsymbol{\xi}(x, a) = \{\Phi_L^\top(x) \mathbb{I}(a=0), \Phi_L^\top(x) \mathbb{I}(a=1), \dots, \Phi_L^\top(x) \mathbb{I}(a=m-1)\}^\top,$$

$$\mathbf{U}_\pi(x) = \{\Phi_L^\top(x) \pi(0|x), \Phi_L^\top(x) \pi(1|x), \dots, \Phi_L^\top(x) \pi(m-1|x)\}^\top,$$

$\boldsymbol{\xi}_{i,t} = \boldsymbol{\xi}(X_{i,t}, A_{i,t})$ ,  $\mathbf{U}_{\pi,i,t} = \mathbf{U}_\pi(X_{i,t})$ . The above equation can be rewritten as  $\mathbb{E} \boldsymbol{\xi}_{i,t} (\boldsymbol{\xi}_{i,t} - \gamma \mathbf{U}_{\pi,i,t+1})^\top \beta_\pi^* = \mathbb{E} \boldsymbol{\xi}_{i,t} Y_{i,t}$ . Based on the observed data, we propose to estimate  $\beta_\pi^*$  by solving

$$\hat{\beta}_\pi = \left\{ \underbrace{\frac{1}{\sum_i T_i} \sum_{i=1}^n \sum_{t=0}^{T_i-1} \boldsymbol{\xi}_{i,t} (\boldsymbol{\xi}_{i,t} - \gamma \mathbf{U}_{\pi,i,t+1})^\top}_{\hat{\boldsymbol{\Sigma}}_\pi} \right\}^{-1} \left( \frac{1}{\sum_i T_i} \sum_{i=1}^n \sum_{t=0}^{T_i-1} \boldsymbol{\xi}_{i,t} Y_{i,t} \right).$$

Let  $\hat{\boldsymbol{\beta}}_\pi = (\hat{\beta}_{\pi,1}^\top, \dots, \hat{\beta}_{\pi,m}^\top)^\top$ , we propose to estimate  $V(\pi; x)$  by

$$\hat{V}(\pi; x) = \sum_{a \in \mathcal{A}} \Phi_L^\top(x) \hat{\beta}_{\pi,a} \pi(a|x) = \mathbf{U}_\pi^\top(x) \hat{\boldsymbol{\beta}}_\pi.$$

A two-side CI is given by

$$[\hat{V}(\pi; x) - z_{\alpha/2} (\sum_i T_i)^{-1/2} \hat{\sigma}(\pi; x), \hat{V}(\pi; x) + z_{\alpha/2} (\sum_i T_i)^{-1/2} \hat{\sigma}(\pi; x)], \quad (3.8)$$

where  $z_\alpha$  denotes the upper  $\alpha$ -th quantile of a standard normal distribution, and

$$\hat{\sigma}^2(\pi; x) = \mathbf{U}_\pi^\top(x) \hat{\boldsymbol{\Sigma}}_\pi^{-1} \hat{\boldsymbol{\Omega}}_\pi (\hat{\boldsymbol{\Sigma}}_\pi^\top)^{-1} \mathbf{U}_\pi(x),$$

where

$$\hat{\boldsymbol{\Omega}}_\pi = \frac{1}{\sum_i T_i} \sum_{i=1}^n \sum_{t=0}^{T_i-1} \boldsymbol{\xi}_{i,t} \boldsymbol{\xi}_{i,t}^\top \left\{ Y_{i,t} + \gamma \sum_{a \in \mathcal{A}} \Phi_L^\top(X_{i,t+1}) \hat{\beta}_{\pi,a} \pi(a|X_{i,t+1}) - \Phi_L^\top(X_{i,t}) \hat{\beta}_{\pi,A_{i,t}} \right\}^2.$$

Let  $\mathbb{G}$  be a reference distribution on the covariate space  $\mathbb{X}$ . Define the following integrated value function

$$V(\pi; \mathbb{G}) = \int_{x \in \mathbb{X}} V(\pi; x) \mathbb{G}(dx).$$

By setting  $\mathbb{G}(\cdot)$  to be a Dirac measure  $\delta_x(\cdot)$ , i.e,  $\mathbb{G}(\mathcal{X}) = \mathbb{I}(x \in \mathcal{X}), \forall \mathcal{X} \subseteq \mathbb{X}$ ,  $V(\pi; \mathbb{G})$  is reduced to  $V(\pi; x)$ . Let  $\nu_0(\cdot)$  be the probability density function of  $X_{0,0}$ . By setting  $\mathbb{G}(dx) = \nu_0(x)dx$ , we obtain

$$V(\pi; \mathbb{G}) = \int_{x \in \mathbb{X}} V(\pi; x) \nu_0(x) dx.$$

Based on  $\widehat{\beta}_\pi$ , a two-side CI for  $V(\pi; \mathbb{G})$  is given by

$$[\widehat{V}(\pi; \mathbb{G}) - z_{\alpha/2} (\sum_i T_i)^{-1/2} \widehat{\sigma}(\pi; \mathbb{G}), \widehat{V}(\pi; \mathbb{G}) + z_{\alpha/2} (\sum_i T_i)^{-1/2} \widehat{\sigma}(\pi; \mathbb{G})], \quad (3.9)$$

where

$$\widehat{V}(\pi; \mathbb{G}) = \int_{x \in \mathbb{X}} \widehat{V}(\pi; x) \mathbb{G}(dx), \quad (3.10)$$

$$\widehat{\sigma}^2(\pi; \mathbb{G}) = \left\{ \int_{x \in \mathbb{X}} \mathbf{U}_\pi(x) \mathbb{G}(dx) \right\}^\top \widehat{\Sigma}_\pi^{-1} \widehat{\Omega}_\pi (\widehat{\Sigma}_\pi^\top)^{-1} \left\{ \int_{x \in \mathbb{X}} \mathbf{U}_\pi(x) \mathbb{G}(dx) \right\}. \quad (3.11)$$

### 3.1.3 Theory

In this section, we focus on proving the validity of the proposed CIs in (3.9). By setting  $\mathbb{G}(\cdot) = \delta_x(\cdot)$ , it implies that the CI in (3.8) achieves nominal coverage as well. To simplify the presentation, we assume  $T_1 = T_2 = \dots = T_n = T$ , all the covariates are continuous and  $\mathbb{X} = [0, 1]^d$ . Our theory is valid regardless of whether  $T$  is bounded or diverges to infinity. We remark that the boundedness of  $T$  does not mean we work on a finite-horizon setting, since  $T$  is the termination time of the study, not the final time step of each trajectory.

In addition, we restrict our attentions to two particular types of sieve basis functions, corresponding to tensor product of B-splines with degree  $r$  and dimension  $L$  or Wavelets with regularity  $r$  and dimension  $L$ . See Section 6 of Chen and Christensen (2015) for a brief review of these sieve bases. This together with A1 implies that there exists a set of vectors  $\{\beta_{\pi,a}^*\}_{a \in \mathcal{A}}$  that satisfy  $\sup_{x \in \mathbb{X}, a \in \mathcal{A}} |Q(\pi; x, a) - \Phi_L^\top(x) \beta_{\pi,a}^*| = O(L^{-p/d})$ . See Section 2.2 of Huang (1998) for detailed discussions on the approximation power of these sieve bases.

Following the behavior policy  $b(\cdot|\cdot)$ , the set of variables  $\{X_{0,t}\}_{t \geq 0}$  forms a time-homogeneous Markov chain. Its transition kernel  $\mathcal{P}_X$  is given by

$$\mathcal{P}_X(\mathcal{B}|x) = \Pr(X_{0,1} \in \mathcal{B} | X_{0,0} = x) = \sum_{a \in \mathcal{A}} \mathcal{P}(\mathcal{B}|x, a) b(a|x), \quad \forall \mathcal{B} \in \mathbb{X}.$$

We impose the following assumptions.

(A2.) The Markov chain  $\{X_{0,t}\}_{t \geq 0}$  has a unique invariant distribution with some density function  $\mu(\cdot)$ . The density functions  $\mu$  and  $\nu_0$  are uniformly bounded away from 0 and  $\infty$ .

(A3.) Suppose (i) and (ii) hold when  $T \rightarrow \infty$  and (i) holds when  $T$  is bounded.

(i)  $\lambda_{\min}[\sum_{t=0}^{T-1} \mathbb{E}\{\boldsymbol{\xi}_{0,t} \boldsymbol{\xi}_{0,t}^\top - \gamma^2 \mathbf{u}_\pi(X_{0,t}, A_{0,t}) \mathbf{u}_\pi^\top(X_{0,t}, A_{0,t})\}] \geq T\bar{c}$  for some constant  $\bar{c} > 0$ , where  $\mathbf{u}_\pi(x, a) = \mathbb{E}\{\mathbf{U}_\pi(X_{0,1}) | X_{0,0} = x, A_{0,0} = a\}$  and  $\lambda_{\min}(K)$  denotes the minimum eigenvalue of a matrix  $K$ .

(ii) The Markov chain  $\{X_{0,t}\}_{t \geq 0}$  is geometrically ergodic.

We make a few remarks. First, we do not require the limiting density function  $\mu$  to be equal to the initial state density  $\nu_0$ .

Second, Condition A3(i) guarantees the matrix  $\mathbb{E}\widehat{\boldsymbol{\Sigma}}_\pi$  is invertible. In Section C.1 of the supplementary article, we show A3(i) is automatically satisfied when  $\mu = \nu_0$ , the target policy  $\pi$  is deterministic and  $b$  is the  $\epsilon$ -greedy policy with respect to  $\pi$  that satisfies  $\epsilon \leq 1 - \gamma^2$ .

Third, we present the detailed definition of geometric ergodicity in Appendix A to save space. Suppose the Markov chain  $\{X_{0,t}\}_{t \geq 0}$  has a finite state space. Assume  $\mathcal{P}_X$  is diagonalizable. Then A3(ii) holds when the second largest eigenvalue of  $\mathbb{P}_X$  is strictly smaller than 1. When  $X_{0,t}$ 's are generated by the vector autoregressive process  $\mathbb{E}\{X_{0,t} | X_{0,t-1}\} = f(X_{0,t-1})$  for some function  $f$ , Saikkonen (2001) provided sufficient conditions that ensure the geometric ergodicity of the Markov chain.

Finally, when  $\nu_0 = \mu$ ,  $\{X_{0,t}\}_{t \geq 0}$  is stationary. Under Condition A3(ii), it follows from Theorem 3.7 of Bradley (2005) that  $\{X_{0,t}\}_{t \geq 0}$  is exponentially  $\beta$ -mixing (see the proof of Lemma 3 for details). When  $T \rightarrow \infty$ , A3(ii) enables us to derive matrix concentration inequalities for  $\widehat{\boldsymbol{\Sigma}}_\pi$ . This together with A3(i) implies that  $\widehat{\boldsymbol{\Sigma}}_\pi$  is invertible, with probability approaching 1 (wpa1). We remark that A3(ii) is not needed when  $T$  is bounded.

For any  $x \in \mathbb{X}, a \in \mathcal{A}$ , define

$$\omega_\pi(x, a) = \mathbb{E} \left[ \left\{ Y_{0,0} + \gamma \sum_{a \in \mathcal{A}} \pi(a | X_{0,1}) Q(\pi; X_{0,1}, a) - Q(\pi; X_{0,0}, A_{0,0}) \right\}^2 \mid X_{0,0} = x, A_{0,0} = a \right].$$

**Theorem 1 (bidirectional asymptotics)** *Assume A1-A3 hold. Suppose  $L$  satisfies  $L = o\{\sqrt{nT}/\log(nT)\}$ ,  $L^{2p/d} \gg nT\{1 + \|\int_x \Phi_L(x) \mathbb{G}(dx)\|_2^{-2}\}$ , and there exists some constant*

$c_0 \geq 1$  such that  $\omega_\pi(x, a) \geq c_0^{-1}$  for any  $x \in \mathbb{X}, a \in \mathcal{A}$  and  $Pr(\max_{0 \leq t \leq T-1} |Y_{0,t}| \leq c_0) = 1$ . Then as either  $n \rightarrow \infty$  or  $T \rightarrow \infty$ , we have

$$\sqrt{nT} \widehat{\sigma}^{-1}(\pi; \mathbb{G}) \{ \widehat{V}(\pi; \mathbb{G}) - V(\pi; \mathbb{G}) \} \xrightarrow{d} N(0, 1).$$

A sketch for the proof of Theorem 1 is given in Appendix E.1. Under the conditions in Theorem 1, we can show that  $\widehat{\sigma}(\pi; \mathbb{G})$  converges almost surely to some  $\sigma(\pi; \mathbb{G})$ . The form of  $\sigma(\pi; \mathbb{G})$  is given in Section E.5. In addition, we have

$$\frac{\widehat{V}(\pi; \mathbb{G}) - V(\pi; \mathbb{G})}{(nT)^{-1/2} \widehat{\sigma}(\pi; \mathbb{G})} = \frac{(nT)^{-1/2}}{\sigma(\pi; \mathbb{G})} \sum_{i=1}^n \sum_{t=1}^T \left\{ \int_{x \in \mathbb{X}} \mathbf{U}_\pi(x) \mathbb{G}(dx) \right\}^\top \boldsymbol{\Sigma}_\pi^{-1} \boldsymbol{\xi}_{i,t} \varepsilon_{i,t} + o_p(1), \quad (3.12)$$

where  $\boldsymbol{\Sigma}_\pi = \mathbb{E} \widehat{\boldsymbol{\Sigma}}_\pi$  and

$$\varepsilon_{i,t} = Y_{i,t} + \gamma \sum_{a \in \mathcal{A}} Q(\pi; X_{i,t+1}, a) \pi(a|X_{i,t+1}) - Q(\pi; X_{i,t}, A_{i,t}).$$

By MA, CMIA and (3.7), the leading term on the RHS of (3.12) forms a mean-zero martingale (details can be found in Section E.5). As either  $n$  or  $T$  grows to infinity, the asymptotic normality follows from the martingale central limit theorem.

When  $\sigma(\pi; \mathbb{G})$  is bounded away from zero, it can be seen from (3.12) that  $\widehat{V}(\pi; \mathbb{G}) - V(\pi; \mathbb{G}) = O_p(n^{-1/2} T^{-1/2})$ . That is, the proposed value estimator converges at a rate of  $(nT)^{-1/2}$ . In contrast, AIPW-type estimators typically converge at a rate of  $n^{-1/2}$  and are thus not suitable for settings with only a few trajectories.

### 3.2 Inference of the value under an (estimated) optimal policy

For simplicity, we assume  $T_1 = T_2 = \dots = T_n = T$  throughout this section. Consider an estimated policy  $\widehat{\pi}$ , computed based on the data  $\{(X_{i,t}, A_{i,t}, Y_{i,t}, X_{i,t+1})\}_{0 \leq t < T, 1 \leq i \leq n}$ . The integrated value under  $\widehat{\pi}$  is given by

$$V(\widehat{\pi}; \mathbb{G}) = \int_{x \in \mathbb{X}} V(\widehat{\pi}; x) \mathbb{G}(dx).$$

We will require the value of  $\widehat{\pi}$  to converge to some fixed policy  $\pi^*$  (possibly different from  $\pi^{opt}$ ), i.e.,

$$V(\widehat{\pi}; \mathbb{G}) - V(\pi^*; \mathbb{G}) \xrightarrow{P} 0. \quad (3.13)$$

In this section, we focus on constructing CIs for  $V(\widehat{\pi}; \mathbb{G})$  and  $V(\pi^*; \mathbb{G})$ .

### 3.2.1 The challenge

We begin by outlining the challenge of obtaining inference in the nonregular cases. Suppose  $\pi^* \in \Pi^{opt}$ . When the optimal policy is not unique,  $\hat{\pi}$  might not converge to a fixed policy, despite that its value converges (see (3.13)). To better illustrate this, suppose  $\hat{\pi}$  is computed by some Q-learning type algorithms, i.e.,

$$\hat{\pi}(a|x) = \begin{cases} 1, & \text{if } a = \text{sarg} \max_{a' \in \mathcal{A}} \hat{Q}(x, a'), \\ 0, & \text{otherwise,} \end{cases} \quad (3.14)$$

where  $\hat{Q}(\cdot, \cdot)$  denotes some consistent estimator for  $Q^{opt}(\cdot, \cdot)$ . Assume there exists a subset  $\mathbb{X}_0$  of  $\mathbb{X}$  with positive Lebesgue measure such that the argmax of  $Q^{opt}(x, \cdot)$  is not uniquely defined for any  $x \in \mathbb{X}_0$ . Then  $\hat{\pi}(\cdot|x)$  might not converge to a fixed quantity for any  $x \in \mathbb{X}_0$ .

Consider the plug-in estimator  $\hat{V}(\hat{\pi}; \mathbb{G})$  for  $V(\hat{\pi}; \mathbb{G})$ . Similar to (3.12), we can show

$$\frac{\hat{V}(\hat{\pi}; \mathbb{G}) - V(\hat{\pi}; \mathbb{G})}{(nT)^{-1/2} \hat{\sigma}(\hat{\pi}; \mathbb{G})} = \frac{(nT)^{-\frac{1}{2}}}{\sigma(\hat{\pi}; \mathbb{G})} \sum_{i,t} \left\{ \int_{x \in \mathbb{X}} \mathbf{U}_{\hat{\pi}}(x) \mathbb{G}(dx) \right\}^\top \Sigma_{\hat{\pi}}^{-1} \boldsymbol{\xi}_{i,t} \hat{\varepsilon}_{i,t} + o_p(1), \quad (3.15)$$

where

$$\hat{\varepsilon}_{i,t} = Y_{i,t} + \gamma \sum_{a \in \mathcal{A}} Q(\hat{\pi}; X_{i,t+1}, a) \hat{\pi}(a|X_{i,t+1}) - Q(\hat{\pi}; X_{i,t}, A_{i,t}).$$

When  $\hat{\pi}$  does not converge,  $\sigma(\hat{\pi}; \mathbb{G})$ ,  $\mathbf{U}_{\hat{\pi}}(x)$ ,  $\Sigma_{\hat{\pi}}$  and  $\hat{\varepsilon}_{i,t}$  will fluctuate randomly and might not stabilize. Since these quantities depend on the data as well, the martingale structure is violated. As a result, the leading term on the RHS of (3.15) does not have a well tabulated limiting distribution. Thus, CIs based on  $\hat{V}(\hat{\pi}; \mathbb{G})$  will fail to maintain the nominal coverage probability.

To allow for valid inference, we use a sequential value evaluation procedure to construct the CI. That is, we propose sequentially estimating the optimal policy and evaluating its value using different data subsets. This allows us to treat the estimated optimal policy as known conditional on past observations (see Equation (E.30) in Appendix E.2). The martingale CLT can thus be applied to obtain the limiting distribution for our estimator (see (E.32) and the related discussions). We detail our procedure in the next section.

### 3.2.2 SAVE for the value under an (estimated) optimal policy

We begin by dividing  $\mathcal{I}_0 = \{(i, t) : 1 \leq i \leq n, 0 \leq t < T\}$  into  $K$  non-overlapping subsets, denoted by  $\mathcal{I}_1, \mathcal{I}_2, \dots, \mathcal{I}_K$ . At the  $k$ -th step, we use the sub-dataset

$$O_k = \{(X_{i,t}, A_{i,t}, Y_{i,t}, X_{i,t+1}) : (i, t) \in \bar{\mathcal{I}}_k = \mathcal{I}_1 \cup \mathcal{I}_2 \cup \dots \cup \mathcal{I}_k\},$$

to compute an estimated optimal policy (denoted by  $\hat{\pi}_{\bar{\mathcal{I}}_k}$ ). Then we apply the proposed procedure in Section 3.1 to dataset in the  $(k+1)$ -th block  $O_{k+1} = \{(X_{i,t}, A_{i,t}, Y_{i,t}, X_{i,t+1}) : (i, t) \in \mathcal{I}_{k+1}\}$  to compute its value estimator  $\hat{V}_{\mathcal{I}_{k+1}}(\hat{\pi}_{\bar{\mathcal{I}}_k}; \mathbb{G})$  and the associated standard error  $\hat{\sigma}_{\mathcal{I}_{k+1}}(\hat{\pi}_{\bar{\mathcal{I}}_k}; \mathbb{G})$  (Details are given in Appendix B.1).

As commented in the introduction, the data-splitting rule cannot be arbitrary. For any of the two tuples  $(i_1, t_1)$  and  $(i_2, t_2)$ , define an order  $(i_2, t_2) \succ (i_1, t_1)$  if either  $t_2 > t_1$  or  $i_2 > i_1$ . For any  $(i_2, t_2) \in \mathcal{I}_{k+1}$ , we require the following:

$$(i_2, t_2) \succ (i_1, t_1), \quad \forall (i_1, t_1) \in \bar{\mathcal{I}}_k. \quad (3.16)$$

Then  $\hat{\pi}_{\bar{\mathcal{I}}_k}$  depends on the  $i_2$ -th patent's trajectory only through  $\{(X_{i_2,j}, A_{i_2,j}, Y_{i_2,j})\}_{j < t_2}$  and  $X_{i_1,t_2}, A_{i_1,t_2}$ . Under MA and CMIA, (3.7) still holds with  $\pi = \hat{\pi}_{\bar{\mathcal{I}}_k}$  any  $(i, t) = (i_2, t_2)$ . Similar to (3.12), we can show conditional on the observations in  $\bar{\mathcal{I}}_k$ ,  $\sqrt{nT/K} \hat{V}_{\mathcal{I}_{k+1}}(\hat{\pi}_{\bar{\mathcal{I}}_k}; \mathbb{G})$  is asymptotically normal with variance consistently estimated by  $\hat{\sigma}_{\mathcal{I}_{k+1}}^2(\hat{\pi}_{\bar{\mathcal{I}}_k}; \mathbb{G})$ .

Our final estimator  $\tilde{V}(\mathbb{G})$  is defined as a weighted average of these  $\hat{V}_{\mathcal{I}_{k+1}}(\hat{\pi}_{\bar{\mathcal{I}}_k}; \mathbb{G})$ 's. Specifically, we set

$$\tilde{V}(\mathbb{G}) = \left\{ \sum_{k=1}^{K-1} \frac{1}{\hat{\sigma}_{\mathcal{I}_{k+1}}(\hat{\pi}_{\bar{\mathcal{I}}_k}; \mathbb{G})} \right\}^{-1} \left\{ \sum_{k=1}^{K-1} \frac{\hat{V}_{\mathcal{I}_{k+1}}(\hat{\pi}_{\bar{\mathcal{I}}_k}; \mathbb{G})}{\hat{\sigma}_{\mathcal{I}_{k+1}}(\hat{\pi}_{\bar{\mathcal{I}}_k}; \mathbb{G})} \right\}.$$

The inclusion of the inverse weight  $1/\hat{\sigma}_{\mathcal{I}_{k+1}}(\hat{\pi}_{\bar{\mathcal{I}}_k}; \mathbb{G})$  is necessary for the theoretical development of asymptotic normality of  $\tilde{V}(\mathbb{G})$  (see (E.32)). Our CI is given by

$$[\tilde{V}(\mathbb{G}) - z_{\alpha/2} \{nT(K-1)/K\}^{-1/2} \tilde{\sigma}(\mathbb{G}), \tilde{V}(\mathbb{G}) + z_{\alpha/2} \{nT(K-1)/K\}^{-1/2} \tilde{\sigma}(\mathbb{G})], \quad (3.17)$$

where  $\tilde{\sigma}(\mathbb{G}) = (K-1) \{ \sum_{k=1}^{K-1} \hat{\sigma}_k^{-1}(\hat{\pi}_{\bar{\mathcal{I}}_k}; \mathbb{G}) \}^{-1}$ .

It remains to specify  $\mathcal{I}_1, \mathcal{I}_2, \dots, \mathcal{I}_K$  that satisfy (3.16). Consider some positive integers  $n_{\min} \leq n$ ,  $T_{\min} \leq T$ . Assume  $n$  and  $T$  are divisible by  $n_{\min}$  and  $T_{\min}$ , respectively. Let  $K_n = n/n_{\min}$  and  $K_T = T/T_{\min}$ . We set  $K = K_n K_T$ . For any  $1 \leq k_n \leq K_n$ ,  $1 \leq k_T \leq K_T$ , define a set  $\mathcal{I}(k_n, k_T)$  by

$$\{(i, t) : (k_n - 1)n_{\min} < i \leq k_n n_{\min}, (k_T - 1)T_{\min} \leq t < k_T T_{\min}\}.$$



Thus, each block  $\mathcal{I}(k_n, k_T)$  contains data from  $n_{\min}$  trajectories with  $T_{\min}$  decision time points. Below, we introduce two special examples.

1. When only a few trajectories are available, we may set  $n_{\min} = n$ . Then, the blocks are constructed according to the times that decisions were being made.
2. When each trajectory contains a very short time period, we may set  $T_{\min} = T$ . Then, the observations are divided according to the trajectories they belong to.

We order these blocks by

$$\mathcal{I}(1, 1), \mathcal{I}(2, 1), \dots, \mathcal{I}(K_n, 1), \mathcal{I}(1, 2), \mathcal{I}(2, 2), \dots, \mathcal{I}(K_n, 2), \dots, \mathcal{I}(1, k_T), \dots, \mathcal{I}(K_n, K_T).$$

Based on this order, we set  $\mathcal{I}_k = \mathcal{I}(n(k), T(k))$  where  $n(k)$  and  $T(k)$  are the unique positive integers that satisfy  $k = n(k) + (T(k) - 1)K_n$ . For any  $k_2 > k_1$ , we have either  $n(k_2) > n(k_1)$  or  $T(k_2) > T(k_1)$ . Thus, the proposed data-splitting rule guarantees (3.16) holds for any  $k$ .

In Theorem 2 below, we establish the validity of our CI in (3.17). It relies on Condition A3\* and A4. A3\* is very similar to A3 and we present the detailed definition in Appendix A to save space.

$$(A4) \quad \mathbb{E}|V(\hat{\pi}_{\bar{\mathcal{I}}_k}; \mathbb{G}) - V(\pi^*; \mathbb{G})| = O(|\bar{\mathcal{I}}_k|^{-b_0}) \text{ for some } b_0 > 1/2 \text{ such that } (nT)^{b_0-1/2} \gg \|\int_x \Phi_L(x) \mathbb{G}(dx)\|_2^{-1}, \text{ where the big-}O \text{ term is uniform in } k.$$

Set  $\mathcal{I} = \mathcal{I}_0$ . By Markov's inequality, it is immediate to see A4 implies that Condition (3.13) holds. When the tensor-product B-splines are used, we have  $\liminf_L \|\int_x \Phi_L(x) \mathbb{G}(dx)\|_2 > 0$ . Thus, it is equivalent to require  $\mathbb{E}|V(\hat{\pi}_{\bar{\mathcal{I}}_k}; \mathbb{G}) - V(\pi^*; \mathbb{G})| = O(|\bar{\mathcal{I}}_k|^{-b_0})$  for some  $b_0 > 1/2$ . In Section 3.2.3, we discuss the rate  $b_0$  in detail when  $\pi^* = \pi^{opt}$  and  $\hat{\pi}$  is a greedy policy derived based on some Q-learning algorithms.

**Theorem 2 (bidirectional asymptotics)** *Assume A1-A2, A3\* and A4 hold. Suppose  $K = O(1)$  and  $L$  satisfies  $L = o\{\sqrt{nT}/\log(nT)\}$ ,  $L^{2p/d} \gg nT\{1 + \|\int_x \Phi_L(x) \mathbb{G}(dx)\|_2^{-2}\}$ . Suppose  $T_{\min} = T$  if  $T$  is bounded. Assume there exists some constant  $c_0 \geq 1$  such that  $\omega_\pi(x, a) \geq c_0^{-1}$  for any  $x, a, \pi$  and  $\Pr(\max_{0 \leq t \leq T-1} |Y_{0,t}| \leq c_0) = 1$ . Then as either  $n \rightarrow \infty$  or  $T \rightarrow \infty$ , we have*

$$\begin{aligned} \sqrt{nT(K-1)/K\tilde{\sigma}^{-1}(\mathbb{G})}\{\tilde{V}(\mathbb{G}) - V(\hat{\pi}; \mathbb{G})\} &\xrightarrow{d} N(0, 1), \\ \sqrt{nT(K-1)/K\tilde{\sigma}^{-1}(\mathbb{G})}\{\tilde{V}(\mathbb{G}) - V(\pi^*; \mathbb{G})\} &\xrightarrow{d} N(0, 1). \end{aligned}$$

We provide a sketch for the proof of Theorem 2 in Appendix E.2.

### 3.2.3 Convergence of the value under an estimated optimal policy

For any  $\mathcal{I} \subseteq \mathcal{I}_0$ , we use  $\hat{\pi}_{\mathcal{I}}$  to denote an estimated optimal policy based on observations in  $\mathcal{I}$ . Let  $\hat{Q}_{\mathcal{I}}(\cdot, \cdot)$  denote some consistent estimator for  $Q^{opt}(\cdot, \cdot)$  and  $\hat{\pi}_{\mathcal{I}}$  denote the greedy policy with respect to  $\hat{Q}_{\mathcal{I}}(\cdot, \cdot)$  (see Equation (3.14)).

In the following, we focus on relating  $|V(\pi^{opt}; \mathbb{G}) - V(\hat{\pi}_{\mathcal{I}}; \mathbb{G})|$  to the prediction loss  $\hat{Q}_{\mathcal{I}} - Q^{opt}$ . By definition,  $V(\pi^{opt}; x) \geq V(\hat{\pi}_{\mathcal{I}}; x), \forall x \in \mathbb{X}$ . Hence,  $V(\pi^{opt}; \mathbb{G}) \geq V(\hat{\pi}_{\mathcal{I}}; \mathbb{G})$ . It suffices to provide an upper bound for  $V(\pi^{opt}; \mathbb{G}) - V(\hat{\pi}_{\mathcal{I}}; \mathbb{G})$ . We introduce a margin-type condition A5 below.

(A5) Assume there exist some constants  $\alpha, \delta_0 > 0$  such that

$$\lambda \left\{ x \in \mathbb{X} : \max_a Q^{opt}(x, a) - \max_{a' \in \mathcal{A} - \arg \max_a Q^{opt}(x, a)} Q^{opt}(x, a') \leq \varepsilon \right\} = O(\varepsilon^\alpha), \quad (3.18)$$

$$\mathbb{G} \left\{ x \in \mathbb{X} : \max_a Q^{opt}(x, a) - \max_{a' \in \mathcal{A} - \arg \max_a Q^{opt}(x, a)} Q^{opt}(x, a') \leq \varepsilon \right\} = O(\varepsilon^\alpha), \quad (3.19)$$

where  $\lambda$  denotes the Lebesgue measure, the big- $O$  terms are uniform in  $0 < \varepsilon \leq \delta_0$ , and  $\max_{a' \in \mathcal{A} - \arg \max_a Q^{opt}(x, a)} Q^{opt}(x, a') = -\infty$  if the set  $\mathcal{A} - \arg \max_a Q^{opt}(x, a) = \emptyset$ .

For each  $x$ , the quantity  $\max_a Q^{opt}(x, a) - \max_{a' \in \mathcal{A} - \arg \max_a Q^{opt}(x, a)} Q^{opt}(x, a')$  measures the difference in value between  $\pi^{opt}$  and the policy that assigns the best suboptimal treatment(s) at the first decision point and follows  $\pi^{opt}$  subsequently. In point treatment studies, Qian and Murphy (2011) imposed a similar condition (see Equation (3.3), Qian and Murphy, 2011) to derive sharp convergence rate for the value under an estimated optimal individualized treatment regime. Here, we generalize their condition in infinite-horizon settings. A5 is also closely related to the margin condition commonly used to bound the excess misclassification error (Tsybakov, 2004; Audibert and Tsybakov, 2007).

The margin-type condition is mild. In Appendix A.3, we present detailed examples and show the condition holds under these examples. The following theorems summarize our results.

**Theorem 3** *Assume A1, (3.18) and (3.19) hold. Suppose the following event occurs with probability at least  $1 - O(|\mathcal{I}|^{-\kappa})$  for any finite  $\kappa > 0$ ,*

$$\sup_{x \in \mathbb{X}, a \in \mathcal{A}} |\hat{Q}_{\mathcal{I}}(x, a) - Q^{opt}(x, a)| = O(|\mathcal{I}|^{-b_*}),$$

for some  $b_* > 0$ . Then  $\mathbb{E}|V(\pi^{opt}; \mathbb{G}) - V(\widehat{\pi}_{\mathcal{I}}; \mathbb{G})| = O(|\mathcal{I}|^{-b_*(1+\alpha)})$ .

In Theorem 3, we require the estimated Q-function to satisfy certain uniform convergence rate. In Theorem 4 below, we relax this condition by assuming that the integrated loss converges to zero at certain rate.

**Theorem 4** *Assume A1 and A5 hold. Suppose*

$$\left( \mathbb{E} \int_{x \in \mathbb{X}} \sum_{a \in \mathcal{A}} |\widehat{Q}_{\mathcal{I}}(x, a) - Q^{opt}(x, a)|^2 dx \right)^{1/2} = O(|\mathcal{I}|^{-b_*}),$$

for some  $b_* > 0$ . Then  $\mathbb{E}|V(\pi^{opt}; \mathbb{G}) - V(\widehat{\pi}_{\mathcal{I}}; \mathbb{G})| = O(|\mathcal{I}|^{-b_*(2+2\alpha)/(2+\alpha)})$ .

It can be seen from Theorems 3 and 4 that the integrated value converges faster compared to the Q-function. We provide a sketch for the proofs of both theorems in Appendix E.3.

### 3.2.4 Applications

In this section, we provide several examples to illustrate the convergence rate of  $\widehat{Q}_{\mathcal{I}}$ . The proposed methods can be applied to evaluating the values under these estimated policies. The algorithm in Example 1 requires to impose a linear model assumption for the optimal Q-function. The algorithm in Example 2 allows more general nonlinear and nonparametric models for the optimal Q-function.

**Example 1 (Greedy gradient Q-learning)** *The optimal Q-function satisfies*

$$Q^{opt}(x, a) = r(x, a) + \gamma \int_{x'} \max_{a' \in \mathcal{A}} Q^{opt}(x', a') \mathcal{P}(dx' | x, a),$$

for any  $a$  and  $x$ , and hence

$$\mathbb{E} \left[ \left\{ Y_{i,t} + \gamma \max_{a' \in \mathcal{A}} Q^{opt}(X_{i,t+1}, a') - Q^{opt}(X_{i,t}, A_{i,t}) \right\} \middle| X_{i,t}, A_{i,t} \right] = 0.$$

Suppose we model  $Q^{opt}(x, a)$  by linear sieves  $\Phi_L^\top(x)\theta_a$ . Then we can compute  $\{\widehat{\theta}_{a,\mathcal{I}}\}_{a \in \mathcal{A}}$  by minimizing the following projected Bellman error:

$$\arg \min_{\{\theta_a\}_{a \in \mathcal{A}}} \left( \sum_{(i,t) \in \mathcal{I}} \delta_{i,t}(\{\theta_a\}_{a \in \mathcal{A}}) \boldsymbol{\xi}_{i,t} \right)^\top \left( \sum_{(i,t) \in \mathcal{I}} \boldsymbol{\xi}_{i,t} \boldsymbol{\xi}_{i,t}^\top \right)^{-1} \left( \sum_{(i,t) \in \mathcal{I}} \delta_{i,t}(\{\theta_a\}_{a \in \mathcal{A}}) \boldsymbol{\xi}_{i,t} \right),$$

where  $\delta_{i,t}(\{\theta_a\}_{a \in \mathcal{A}}) = Y_{i,t} + \gamma \max_{a' \in \mathcal{A}} \Phi_L^\top(X_{i,t+1})\theta_{a'} - \Phi_L^\top(X_{i,t})\theta_{A_{i,t}}$ . The above loss is non-smooth and non-convex as a function of  $\{\theta_a\}_{a \in \mathcal{A}}$ . The estimator  $\{\hat{\theta}_{a,\mathcal{I}}\}_{a \in \mathcal{A}}$  can be computed based on the greedy gradient  $Q$ -learning algorithm.

Assuming the optimal  $Q$ -function is correctly specified, Ertefaie and Strawderman (2018) established the consistency and asymptotic normality of the parameter estimates under the scenario where both  $L$  and  $T$  are fixed. Set  $\hat{Q}_{\mathcal{I}}(x, a) = \Phi^\top(x)\hat{\theta}_{a,\mathcal{I}}$ . Using similar arguments in proving Theorem 1, we can show that with proper choice of  $L$ ,  $\sup_{x \in \mathbb{X}, a \in \mathcal{A}} |\hat{Q}_{\mathcal{I}}(x, a) - Q^{\text{opt}}(x, a)|$  coverages at a rate of  $O(|\mathcal{I}|^{-p/(2p+d)})$  up to some logarithmic factors, with probability at least  $1 - O(n^{-2}T^{-2})$ . The condition in Theorem 3 thus holds for any  $b_* < p/(2p+d)$ .

**Example 2 (Fitted  $Q$ -iteration)** In fitted  $Q$ -iteration (FQI), the optimal  $Q$ -function is approximated by some nonparametric models  $Q(\cdot, \cdot, \theta)$  indexed by  $\theta$ . The parameter  $\theta$  is iteratively updated by

$$\hat{\theta}_{k+1} = \arg \min_{\theta} \sum_{(i,t) \in \mathcal{I}_k} \{Y_{i,t} + \gamma \max_a Q(X_{i,t+1}, a; \hat{\theta}_k) - Q(X_{i,t}, A_{i,t}; \theta)\}^2,$$

for  $k = 0, 1, 2, \dots, K-1$ , where  $\mathcal{I}_k$ 's are some subsets of  $\mathcal{I}$ . When  $\mathcal{I}_1 = \dots = \mathcal{I}_K = \mathcal{I}$  and  $Q(\cdot, \cdot, \theta)$  is the family of neural networks, this algorithm is the neural FQI proposed by Riedmiller (2005). Fan et al. (2020) studied a variant of neural FQI by assuming  $\mathcal{I}_k$ 's are disjoint and the training samples in  $\cup_{k=1}^K \mathcal{I}_k$  are independent. Using similar arguments in the proof of Theorem 4.4 in Fan et al. (2020), we can show  $\mathbb{E} \int_{x \in \mathbb{X}} \sum_{a \in \mathcal{A}} |\hat{Q}_{\mathcal{I}}(x, a) - Q^{\text{opt}}(x, a)|^2 dx$  coverages at a rate of  $O(|\mathcal{I}|^{-(2p)/(2p+d)})$  up to some logarithmic factors. The conditions in Theorem 4 thus hold for any  $b_* < p/(2p+d)$ .

## 4 Extensions to on-policy evaluation

We now extend our methodology in Section 3 to on-policy settings. The proposed CI is similar to that presented in Section 3.2.2 and applies to any reinforcement learning algorithms that iteratively update the estimated policy based on batches of observations. Let  $\{T(k)\}_{k \geq 1}$  be a monotonically increasing sequence that diverges to infinity. At the  $k$ -th iteration, define  $\bar{\mathcal{I}}_k = \{(i, t) : 1 \leq i \leq n, 0 \leq t < \sum_{j=1}^k T(j)\}$ . The data observed so far can be summarized as  $\{(X_{i,t}, A_{i,t}, Y_{i,t}, X_{i,t+1})\}_{1 \leq i \leq n, 0 \leq t < \sum_{j=1}^k T(j)}$ . We compute the

estimated policy  $\widehat{\pi}_{\bar{\mathcal{I}}_k}$  based on these data. Then we determine the behavior policy  $\widehat{b}_{\bar{\mathcal{I}}_k}$  as a function of  $\widehat{\pi}_{\bar{\mathcal{I}}_k}$  and generate new observations

$$\{(A_{i,t}, Y_{i,t}, X_{i,t+1})\}_{1 \leq i \leq n, \sum_{j=1}^k T(j) \leq t < \sum_{j=1}^{k+1} T(j)}, \quad (4.20)$$

according to  $\widehat{b}_{\bar{\mathcal{I}}_k}$ . To balance the exploration-exploitation trade-off, a common choice of  $\widehat{b}_{\bar{\mathcal{I}}_k}$  is the  $\epsilon$ -greedy policy with respect to  $\widehat{\pi}_{\bar{\mathcal{I}}_k}$ .

Let  $\mathcal{I}_{k+1} = \{1 \leq i \leq n, \sum_{j=1}^k T(j) \leq t < \sum_{j=1}^{k+1} T(j)\}$ . The new observations in (4.20) are conditionally independent of  $\widehat{\pi}_{\bar{\mathcal{I}}_k}$  given those in  $\bar{\mathcal{I}}_k$ . So the Bellman equation in (3.7) is valid with  $\pi = \widehat{\pi}_{\bar{\mathcal{I}}_k}$  for any  $(i, t) \in \mathcal{I}_{k+1}$ . We compute  $\widehat{V}_{\mathcal{I}_{k+1}}(\widehat{\pi}_{\bar{\mathcal{I}}_k}; \mathbb{G})$  and  $\widehat{\sigma}_{\mathcal{I}_{k+1}}(\widehat{\pi}_{\bar{\mathcal{I}}_k}; \mathbb{G})$  as in Appendix B.1 of the supplementary article, where the number of basis  $L(k+1)$  depends on both  $n$  and  $T(k+1)$ . We iterate this procedure for  $k = 1, 2, \dots, K-1$ . The estimated value and CI for  $V(\widehat{\pi}_{\bar{\mathcal{I}}_k}; \mathbb{G})$  are given by

$$\widetilde{V}(\mathbb{G}) = \left\{ \sum_{k=1}^{K-1} \frac{\sqrt{T(k)}}{\widehat{\sigma}_{\mathcal{I}_{k+1}}(\widehat{\pi}_{\bar{\mathcal{I}}_k}; \mathbb{G})} \right\}^{-1} \left\{ \sum_{k=1}^{K-1} \frac{\sqrt{T(k)} \widehat{V}_{\mathcal{I}_{k+1}}(\widehat{\pi}_{\bar{\mathcal{I}}_k}; \mathbb{G})}{\widehat{\sigma}_{\mathcal{I}_{k+1}}(\widehat{\pi}_{\bar{\mathcal{I}}_k}; \mathbb{G})} \right\},$$

and

$$\left[ \widetilde{V}(\mathbb{G}) - z_{\alpha/2} \left\{ \sum_{k=2}^K \sqrt{\frac{nT(k)}{K-1}} \right\}^{-1/2} \widetilde{\sigma}(\mathbb{G}), \widetilde{V}(\mathbb{G}) + z_{\alpha/2} \left\{ \sum_{k=2}^K \sqrt{\frac{nT(k)}{K-1}} \right\}^{-1/2} \widetilde{\sigma}(\mathbb{G}) \right],$$

where  $\widetilde{\sigma}(\mathbb{G}) = \{\sum_{k=2}^K \sqrt{T(k)}\} \{\sum_{k=2}^K \sqrt{T(k)} \widehat{\sigma}_k^{-1}(\widehat{\pi}_{\bar{\mathcal{I}}_{k-1}}; \mathbb{G})\}^{-1}$ . Similar to Theorem 2, we can show such a CI achieves nominal coverage under certain conditions. To save space, we provide our technical results in Section C.3 of the supplementary article.

## 5 Simulations

In this section, we conduct Monte Carlo simulations to examine the finite sample performance of the proposed CI. We consider off-policy settings in Sections 5.1 and 5.2, where CIs for values under both fixed and optimal policies are reported. In Section 5.3, we report CIs computed in on-policy settings. The state vector  $X_{0,t}$  in our settings might not have bounded supports. For  $j = 1, \dots, d$ , we define  $X_{0,t}^{(j)*} = \Phi(X_{0,t}^{(j)})$  where  $X_{0,t}^{(j)}$  stands for the  $j$ -th element of  $X_{0,t}$  and  $\Phi(\cdot)$  is the cumulative distribution function of a standard normal random variable. This gives us a transformed state vector with bounded




Figure 1: Illustration of Cliff Walking

support. The basis functions are constructed from the tensor product of  $d$  one-dimensional cubic B-spline sets where knots are placed at equally spaced sample quantiles of the transformed state variables. For discrete state space  $\mathbb{X} = \{x_1, \dots, x_M\}$ , we set  $L = M$  and  $\Phi_M(\cdot) = \{\mathbb{I}(\cdot = x_1), \dots, \mathbb{I}(\cdot = x_M)\}^\top$ . We set the discount factor  $\gamma = 0.5$  in all settings, and set  $L = \lfloor (nT)^\eta \rfloor$  with  $\eta = 3/7$ . Here, for any  $z \in \mathbb{R}$ ,  $\lfloor z \rfloor$  denotes the largest integer that is smaller or equal to  $z$ . We tried several other values of the parameter  $\eta$ , and the resulting CIs are very similar and not sensitive to the choice of  $\eta$ . We also tried several other values of  $\gamma$ . Overall, the proposed CI achieves nominal coverage and performs better than other baseline methods. More details can be found in Appendix D.2 of the supplementary article.

## 5.1 Off-policy evaluation with a fixed target policy

We consider three scenarios. In Scenarios (A) and (B), the system dynamics are given by

$$\begin{aligned} X_{0,t+1} &= \begin{bmatrix} \frac{3}{4}(2A_{0,t} - 1) & 0 \\ 0 & \frac{3}{4}(1 - 2A_{0,t}) \end{bmatrix} X_{0,t} + z_t, \\ Y_{0,t} &= X_{0,t+1}^\top \begin{bmatrix} 2 \\ 1 \end{bmatrix} - \frac{1}{4}(2A_{0,t} - 1), \end{aligned} \quad (5.21)$$

for  $t \geq 0$ , where  $\{z_t\}_{t \geq 0} \stackrel{iid}{\sim} N(0_2, I_2/4)$  and  $X_{0,0} \sim N(0_2, I_2)$ . In Scenario (A), we consider a completely randomized study and set  $\{A_{0,t}\}_{t \geq 0}$  to i.i.d Bernoulli random variables with expectation 0.5. In Scenario (B), we allow the treatment assignment mechanism to depend on the observed state. Specifically, we set  $m = 2$  and  $\Pr(A_{0,t} = 1 | X_{0,t}) = 0.5 \text{sigmoid}(X_{0,t,1}) + 0.5 \text{sigmoid}(X_{0,t,2})$  where  $X_{0,t,i}$  denotes the  $i$ th element in  $X_{0,t}$ . The target policy we consider is designed as follows,

$$\pi(a|x) = \begin{cases} 0, & x_1 > 0 \text{ and } x_2 > 0; \\ 1, & \text{otherwise,} \end{cases}$$

where  $x_i$  denotes the  $i$ th element of  $x$ . The reference distribution  $\mathbb{G}$  is set to  $N(0_2, I_2)$ .

In Scenario (C), we consider a standard RL setting included in OpenAI Gym (Brockman et al., 2016): Cliff Walking. This RL example is detailed in Example 6.6 in Sutton and Barto (2018). The objective is to identify the optimal path from the starting point S to the destination point G without falling off the cliff (see Figure 1). This scenario corresponds to an episodic task where the agent will be sent instantly to the starting point wherever it steps into the cliff or arrives at the destination. We manually add some noises to the immediate rewards simulated by the OpenAI Gym to ensure that the system dynamics are not deterministic. We remark that this task is considered in Kallus and Uehara (2020) as well. The target policy is the optimal policy and the behavior policy is a 50-50 mixture of the optimal and uniform random policies.


The true value function  $V(\pi, \mathbb{G})$  is computed by Monte Carlo approximations. Specifically, we simulate  $N = 10^5$  independent trajectories with initial state variable distributed according to  $\mathbb{G}$ . The action at each decision point is chosen according to  $\pi$ . Then we approximate  $V(\pi, \mathbb{G})$  by  $\sum_{j=1}^N \sum_{t=0}^{T_i-1} \gamma^t Y_{j,t} / N$  where  $T_i$  is set to 500 in Scenarios (A), (B) and the termination time of each episode in Scenario (C). The integrals in (3.10) and (3.11) are computed via Monte Carlo methods. For Scenarios (A) and (B), we further consider 9 cases by setting  $n = 25, 50, 100$  and  $T = 30, 50, 70$ . For Scenario (C), we consider 3 cases by setting  $n = 500, 1000, 1500$ . Each trajectory have 13 time points on average, under the behavior policy.

The DRL estimator has been shown to be much more efficient than AIPW or IPW estimators (Thomas et al., 2015; Jiang and Li, 2016). So we focus on comparing our approach with DRL. DRL requires the calculation of the Q-function, the marginalized density ratio and the behavior policy. Here, we treat the behavior policy as known and estimate the Q-function and the density ratio based on nonparametric sieve regression.

Table 1: Empirical coverage probabilities (ECP) and average lengths (AL) of CIs constructed by the proposed method and DRL as well as the mean-squared errors (MSE) of the corresponding value estimators under Scenario (C).

n	SAVE			DRL		
	ECP	AL	log(MSE)	ECP	AL	log(MSE)
500	0.96	0.11	-8.11	0.89	0.13	-6.50
1000	0.93	0.08	-7.60	0.85	0.09	-7.26
1500	0.96	0.06	-8.11	0.87	0.07	-7.82

Figure 2: Empirical coverage probabilities (ECP) and average lengths (AL) of CIs constructed by the proposed method (colored in blue) and the DRL method (colored in red) as well as the mean-squared errors (MSE) of the corresponding value estimators, with different choices of  $n$  and  $T$ . Settings correspond to Scenario (A) and Scenario (B), from top plots to bottom plots.



In Figure 2 and Table 1, we report the empirical coverage probabilities (ECPs) and average lengths (ALs) of CIs constructed by the proposed method, with different choices of  $n$  and  $T$ . It can be seen that our CI achieves nominal coverage in all cases. Its length decreases as  $nT$  increases. This is consistent with our theoretical findings where we show the proposed value estimator converges at a rate of  $n^{-1/2}T^{-1/2}$  under certain conditions (see the discussions below Theorem 1).

Comparing our method with DRL, it is clear that our CIs are in general narrower than those constructed by DRL. In addition, MSEs of the proposed value estimates are smaller than those based on DRL. This is consistent with our theoretical analysis in Appendix C.2 where we show the variance of our value estimator is strictly smaller than that based on DRL under certain conditions. In addition, it can be seen from Table 1 that ECPs of DRL are below 90% in Scenario (C).

In Appendix D.3, we conduct some additional simulation studies under Scenario (A) by setting the reference distribution  $\mathbb{G}$  to a Dirac measure. The proposed CI achieves nominal



coverage under these settings as well.

## 5.2 Off-policy evaluation with an (estimated) optimal policy

In this section, we focus on constructing the CI for value under an optimal policy. Specifically, we use a version of fitted Q-iteration (double FQI) to compute the estimated optimal policy. Detailed algorithm can be found in Section B.3 of the supplementary article. To implement the proposed CI in Section 3.2, we set  $K_n = 2$ ,  $K_T = 2$  in Scenarios (A), (B) and  $K_n = 3$ ,  $K_T = 1$  in Scenario (C). To evaluate our CI, we generate a very large dataset to compute an estimated optimal policy  $\hat{\pi}^*$  based on double FQI and use the Monte Carlo methods described in Section 5.1 to evaluate its value  $V(\hat{\pi}^*; \mathbb{G})$ . Then we treat  $V(\hat{\pi}^*; \mathbb{G})$  as the true optimal value  $V(\pi^*; \mathbb{G})$ . We consider the same three scenarios detailed in Section 5.1. For Scenarios (A) and (B), we fix  $\gamma = 0.5$  and consider 6 cases by setting  $n = 100, 200$  and  $T = 60, 100, 140$ . For Scenario (C), we consider 4 cases by setting  $n = 3000, 4500$  and  $\gamma = 0.5, 0.7$ . The ECP, AL and MSE of our CI are reported in the top two panels of Figure 3 in Table 2. It can be seen that these ECPs are close to the nominal level in most cases. ALs and MSEs decay as either  $n$  or  $T$  increases.

In addition, we design a non-regular setting Scenario (D) where the actions do not have effects on the transition dynamics or the immediate rewards. Specifically, for any  $t \geq 0$ , we set

$$X_{0,t+1} = \begin{bmatrix} -\frac{3}{4} & 0 \\ 0 & \frac{3}{4} \end{bmatrix} X_{0,t} + z_t,$$

and

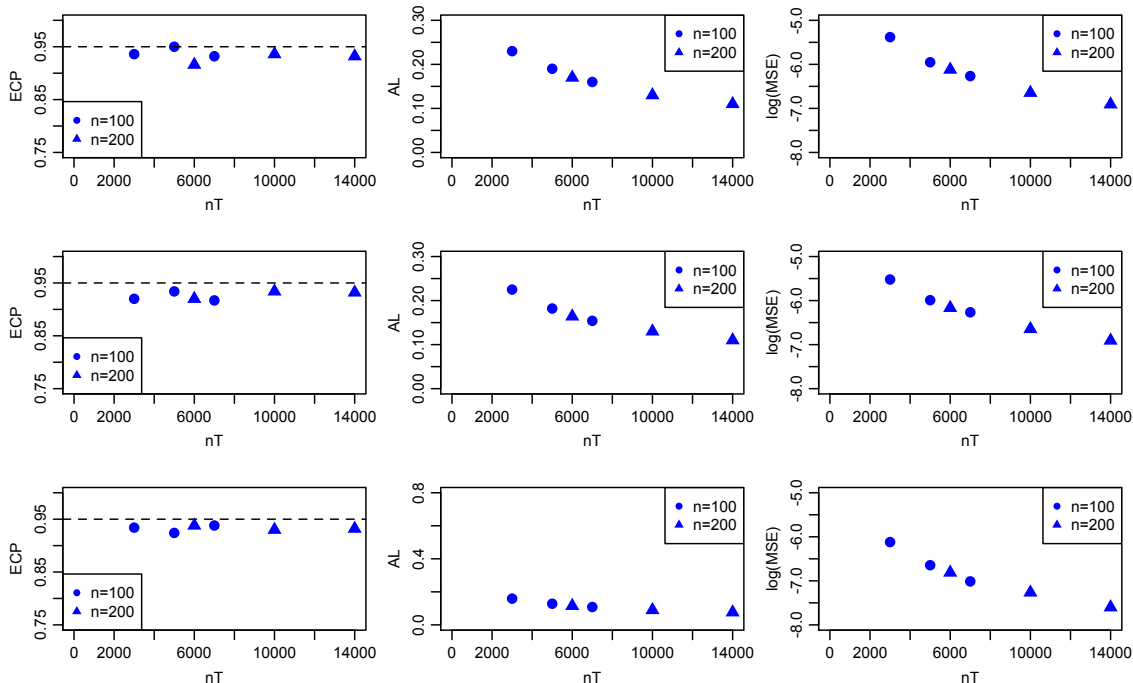
$$Y_{0,t} = X_{0,t+1}^\top \begin{bmatrix} 2 \\ 1 \end{bmatrix},$$

where  $\{z_t\}_{t \geq 0} \stackrel{iid}{\sim} N(0_2, I_2/4)$  and  $\{A_{0,t}\}_{t \geq 0}$  are i.i.d Bernoulli random variables with expectation 0.5. Under this setup, any policy will achieve the same value function. As a result,

Table 2: Empirical coverage probabilities (ECP) and average lengths (AL) of CIs constructed by the proposed method as well as the mean-squared errors (MSE) of the corresponding value estimators under Scenario (C), with different combinations of  $n$  and  $\gamma$ .

	$n = 3000$			$n = 4500$		
$\gamma$	ECP	AL	$\log(\text{MSE})$	ECP	AL	$\log(\text{MSE})$
0.5	0.94	0.12	-6.91	0.94	0.10	-7.51
0.7	0.95	0.23	-5.81	0.96	0.19	-6.35

Figure 3: Empirical coverage probabilities and average lengths of CIs constructed by the proposed method as well as the mean squared errors of the value estimates, with different choices of  $n$  and  $T$ . Settings correspond to Scenarios (A), (B) and (D), from top panels to bottom panels.



the optimal policy is not unique. We consider the same reference distribution  $\mathbb{G}$ , and the same combinations of  $n$  and  $T$  as in the regular setting. ECPs and ALs of the proposed CIs are plotted in the bottom panels of Figure 3. It can be seen that our CIs achieve nominal coverage in the non-regular setting as well.

In Appendix D.3, we conduct some additional simulation studies under Scenario (A) by setting the reference distribution  $\mathbb{G}$  to a Dirac measure. Findings are very similar to those in cases where  $\mathbb{G} = N(0_2, I_2)$ .

### 5.3 On-policy evaluation with an (estimated) optimal policy

We consider a setting where the transition dynamics and immediate rewards are defined by (5.21). In the first block of data  $\{(X_{i,t}, A_{i,t}, Y_{i,t}, X_{i,t+1})\}_{0 \leq t < T(1), 1 \leq i \leq n}$ , the actions are generated according to i.i.d Bernoulli random variables with expectation 0.5. For  $k = 2, \dots, K$ , we use double FQI to estimate the optimal policy based on the data observed so

Table 3: Empirical coverage probabilities and average lengths of CIs constructed in on-policy settings with different choices of  $T$  and  $\mathbb{G}$

	ECPs			ALs		
	T = 120	200	280	T = 120	200	280
$\mathbb{G} = \delta_{(0.5,0.5)}$	0.947	0.914	0.931	0.51	0.46	0.38
$\mathbb{G} = \delta_{-(0.5,0.5)}$	0.925	0.942	0.940	0.72	0.59	0.48
$\mathbb{G} = N(0_2, I_2)$	0.914	0.926	0.948	0.29	0.21	0.17

far  $\{(X_{i,t}, A_{i,t}, Y_{i,t}, X_{i,t+1})\}_{0 \leq t < T(k-1), 1 \leq i \leq n}$  and use an  $\epsilon$ -greedy method to generate actions in the next block of data. In our experiments, we set  $\epsilon = 0.2$ ,  $K = 4$  and  $T(1) = T(2) = \dots = T(K) = T$ . We fix  $n = 25$  and consider three choices of  $T$ , corresponding to  $T = 120$ , 200 and 280. We consider three choices of  $\mathbb{G}$ , corresponding to  $\mathbb{G} = \delta_{(0.5,0.5)}$ ,  $\delta_{(-0.5,-0.5)}$  and  $N(0_2, I_2)$ . The true optimal value function is approximated by Monte Carlo methods, as in off-policy settings. ECPs and ALs of the proposed CIs are reported in Table 3. It can be seen that ECPs are close to the nominal level in almost all cases and ALs decrease as  $T$  increases.

## 6 Application to the OhioT1DM dataset

As commented in the introduction, this dataset contains eight weeks' records of CGM blood glucose levels, insulin doses and self-reported life-event data for each of six patients with type 1 diabetes. To analyze this data, we divide these eight weeks into three hour intervals. The state variable  $X_{i,t}$  is set to be a three-dimensional vector. Specifically, its first element  $X_{i,t}^{(1)}$  is the average CGM blood glucose levels during the three hour interval  $[t-1, t)$ . The second covariate  $X_{i,t}^{(2)}$  is constructed based on the  $i$ -patient's self-reported time and the carbohydrate estimate for the meal. Suppose the patient has meals at time  $t_1, t_2, \dots, t_N \in [t-1, t)$  with the carbohydrate estimates  $CE_1, CE_2, \dots, CE_N$ . Define

$$X_{i,t}^{(2)} = \sum_{j=1}^N CE_j \gamma_c^{36(t_j - t + 1)},$$

where  $\gamma_c$  corresponds to the decay rate every five minutes. Here, we set  $\gamma_c = 0.5$ . The third covariate  $X_{i,t}^{(3)}$  is defined as an average of the basal rate during the three hour interval.

We discretize the action according to the amount of insulin injected in the three hour interval. Specifically,  $A_{i,t} = 1$  when the total amount of insulin delivered to the  $i$ -th patient

Figure 4: Confidence interval of the value difference between an estimated optimal policy and the behavior policy for each of the six patients, with  $\gamma = 0.5$ .





Figure 5: Confidence intervals of the value difference between an estimated optimal policy and the behavior policy for each of the six patients, with different  $\gamma$ .



is greater than one unit. Otherwise, we set  $A_{i,t} = 0$ . The immediate reward  $Y_{i,t}$  is defined according to the Index of Glycemic Control (IGC, Rodbard, 2009), which is a non-linear function of the blood glucose levels. Specifically, we set

$$Y_{i,t} = \begin{cases} -\frac{1}{30}(80 - X_{i,t+1}^{(1)})^2, & X_{i,t+1}^{(1)} < 80; \\ 0, & 80 \leq X_{i,t+1}^{(1)} < 140; \\ -\frac{1}{30}(X_{i,t+1}^{(1)} - 140)^{1.35}, & 140 \leq X_{i,t+1}^{(1)}. \end{cases}$$

A large IGC indicates the patient is in good health status. We set the discount factor  $\gamma = 0.5$ , as in simulations.

For the  $i$ -th patient, we apply the double FQI algorithm to the data

$$\{(X_{i,t}, A_{i,t}, Y_{i,t}, X_{i,t+1})\}_{0 \leq t < T},$$

to estimate a patient-specific optimal policy. Then we compute the estimator for the value function  $V(\pi^{opt}; X_{i,0})$  starting from the initial state variable  $X_{i,0}$ . In addition, we extend our methodology in Section 5.2 to construct the confidence interval for the value difference  $V(\pi^{opt}; X_{i,0}) - V(b; X_{i,0})$  where  $V(b; X_{i,0})$  corresponds to the value under the behavior policy. See Appendix B.2 for details. In Figure 4, we plot our proposed CI for the value difference, for each of the six patients, when the initial starting time is either 8:00 am or 2:00 pm in Day 1. It can be seen that the estimated value differences are strictly positive in all cases. This implies that the optimal value is strictly larger than the observed discounted cumulative reward. In some cases, the lower bound of our CI is larger than zero. The difference is thus significant. In Figure 5, we fix the starting time to 8:00 am and plot the CI of the value difference with different  $\gamma$ . Results show a similar qualitative pattern. This suggests applying reinforcement learning algorithms could potentially improve some patients' health status.

## 7 Discussion

### 7.1 Comparison between DRL

We discuss the advantages and limitations among the proposed method and DRL when inferring the value under a fixed decision rule. Generally speaking, the proposed method results in narrower CIs and would be preferred in cases where  $(nT)^{-1/2}$ -consistent estimation of the Q-function is feasible. This includes settings where the dimension of the state-vector is not large, as in our real data applications. In contrast, DRL would be preferred in ergodic environments with high-dimensional covariates where  $(nT)^{-1/2}$ -consistent estimation of the Q-function is infeasible.

Specifically, in Appendix C.2, we consider settings where both the behavior policy and the target policy are nondynamic. Under certain conditions, we prove that the variance of the DRL estimator is strictly larger than that of the proposed estimator. This in turn implies that our method yields a narrower CI in general.

In addition, we remark that the CI constructed by DRL requires the data to be generated from an ergodic environment. In the Cliff Walking example, the data are generated by a mixture of the optimal and random policy. Since the agent will be sent instantly to the starting point wherever it steps into the cliff or arrives at the destination, the Markov chain formed by the state-action pair is no longer ergodic. It can be seen from Table 1 where ECP of the CI is well below the nominal level in the Cliff Walking example. Although our procedure also requires the ergodicity assumption (see Condition (A3)(ii)), this assumption is not necessary. It can be seen from the proof of Theorem 1 that our CI is valid as long as the random matrix  $\widehat{\Sigma}_\pi$  stabilizes. This is consistent with the findings in Table 1 where our CI achieves nominal coverage in all cases.

However, to ensure the proposed CI is valid, we require the bias of our Q-estimator to decay at a rate of  $o((nT)^{-1/2})$ . Consequently, our estimator converges at a rate of  $(nT)^{-1/2}$ . This rate might not be achievable in high-dimensions. In contrast, DRL requires a weaker condition. The CI based on DRL is valid when both the Q-estimator and the estimated marginalized density ratio converge at a rate of  $o((nT)^{-1/4})$ .

Another potential limitation of our method is that in cases where  $\widehat{\Sigma}_\pi$  is close to a singular matrix, the resulting Q-estimator might suffer from over-fitting, leading to an unbounded outcome. In practice, we could add a ridge penalty to reduce over-fitting. We discuss in detail in Appendix D.4.

## 7.2 Number of basis functions

We outline a procedure to choose the number of basis function  $L$  in this section. The idea is to simulate the model dynamics and select  $L$  such that the resulting confidence interval achieves nominal coverage under the simulated model. Specifically, given the observed data  $\{(X_{i,t}, A_{i,t}, Y_{i,t}, X_{i,t+1})\}_{0 \leq t < T_i, 1 \leq i \leq n}$ , we propose to learn the conditional density function of  $(Y_{i,t}, X_{i,t+1})$  given  $(A_{i,t}, X_{i,t})$ . Following Janner et al. (2019), we recommend to use a Gaussian distribution to model the conditional density function in practice,

$$(Y_{i,t}, X_{i,t+1}) | (A_{i,t}, X_{i,t}) = (a, x) \sim N(\mu(a, x), \Sigma(a, x)). \quad (7.22)$$

The conditional mean  $\mu$  can be estimated via nonparametric regression (e.g., random forest). Let  $\widehat{\mu}$  denote the corresponding estimator. Given the set of estimated residuals

$\{(Y_{i,t}, X_{i,t+1})^\top - \widehat{\mu}(A_{i,t}, X_{i,t})\}_{i,t}$  the conditional covariance function  $\Sigma$  can be estimated via regression as well. We remark that in addition to the Gaussian function, other density functions could be used to model the system dynamics as well.

The behavior policy  $b$  can be similarly estimated via regression. Given an estimated behavior policy  $\widehat{b}$  and  $\widehat{\mu}, \widehat{\Sigma}$ , we generate simulated trajectories to investigate the performance of the proposed CI with different choices of  $L$ .

Finally, we choose  $L$  such that the resulting CI is the shortest among all CIs whose coverage probabilities are above certain level (e.g., 93%) under the simulated environment.

In Appendix D.1, we investigate the finite sample performance of such a method and find that it performs reasonably well. We remark that alternative to the aforementioned method, cross-validation could be applied to select  $L$ .

### 7.3 Sensitivity to the ordering of trajectories

The proposed sequential value evaluation procedure in Section 3.2 divides the data into blocks defined both by trajectories and by time. While there is a natural order in time, there does not appear to be a natural order in the trajectories. In Appendix D.2.3, we conduct additional simulation studies to investigate the sensitivity of our CI to the ordering of trajectories under Scenario (B). Results suggest that our CI is not overly sensitive under our simulation setting.

As suggested by one of the referees, we may aggregate CIs over multiple orderings in cases where the results depend strongly on the ordering of the trajectories. Dezeure et al. (2015) derived a CI for the regression coefficients in high-dimensional models by aggregating results over multiple sample splits using a quantile function. We can adopt their method to aggregate our CIs over multiple orderings. Alternatively, one may average the value estimates over sufficiently many orderings and apply similar methods developed in Wang et al. (2020); Shi et al. (2020a) to derive the CI. However, these algorithms are much more time-consuming.

### 7.4 More on value-based method

In Section 3.1.1, we discuss a potential drawback of using nonparametric methods to directly model the value function. We remark that a kernel-type importance sampling estimator

for the  $V(\pi; x)$  will not suffer from this issue, since it does not directly model the value function, but uses an inverse propensity-score weighted method instead. Both IPW and regression type estimators have their own merits. In general, IPWEs might suffer from a large variance whereas regression-based estimators might suffer from a large bias. There exist methods that combine both for more robust off-policy evaluation (see e.g., Kallus and Uehara, 2019; Uehara et al., 2020; Tang et al., 2020; Shi et al., 2021). However, as commented in Section 7.1, they might yield larger CIs compared to our method.

## 7.5 Rate of convergence of Q-learning type algorithms

Through authors' communication, we found a recent independent work by Hu et al. (2021) that derived a nonasymptotic error bound on the value of the estimated optimal policy computed by Q-learning type algorithms under the margin condition. Their results are consistent with our theoretical findings in Theorems 3 and 4 that show the value of the estimated optimal policy converges to the optimal value at a faster rate than the estimated Q-function.

## References

- Audibert, J.-Y. and Tsybakov, A. B. (2007) Fast learning rates for plug-in classifiers. *Ann. Statist.*, **35**, 608–633.
- Bradley, R. C. (2005) Basic properties of strong mixing conditions. A survey and some open questions. *Probab. Surv.*, **2**, 107–144. Update of, and a supplement to, the 1986 original.
- Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J. and Zaremba, W. (2016) Openai gym. *arXiv preprint arXiv:1606.01540*.
- Burman, P. and Chen, K.-W. (1989) Nonparametric estimation of a regression function. *Ann. Statist.*, **17**, 1567–1596.
- Chen, X. and Christensen, T. M. (2015) Optimal uniform convergence rates and asymptotic normality for series estimators under weak dependence and weak conditions. *J. Econometrics*, **188**, 447–465.



- Davydov, Y. A. (1973) Mixing conditions for markov chains. *Teoriya Veroyatnostei i ee Primeneniya*, **18**, 321–338.
- Dezeure, R., Bühlmann, P., Meier, L. and Meinshausen, N. (2015) High-dimensional inference: Confidence intervals, p-values and r-software hdi. *Statistical science*, 533–558.
- Ernst, D., Geurts, P. and Wehenkel, L. (2005) Tree-based batch mode reinforcement learning. *J. Mach. Learn. Res.*, **6**, 503–556.
- Ertefaie, A. and Strawderman, R. L. (2018) Constructing dynamic treatment regimes over indefinite time horizons. *Biometrika*, **105**, 963–977.
- Fan, J., Wang, Z., Xie, Y. and Yang, Z. (2020) A theoretical analysis of deep q-learning. In *Learning for Dynamics and Control*, 486–489. PMLR.
- Hasselt, H. V. (2010) Double q-learning. In *Advances in Neural Information Processing Systems*, 2613–2621.
- Hu, Y., Kallus, N. and Uehara, M. (2021) Fast rates for the regret of offline reinforcement learning. *arXiv preprint arXiv:2102.00479*.
- Huang, J. Z. (1998) Projection estimation in multiple regression with application to functional ANOVA models. *Ann. Statist.*, **26**, 242–272.
- Janner, M., Fu, J., Zhang, M. and Levine, S. (2019) When to trust your model: Model-based policy optimization. In *Advances in Neural Information Processing Systems*, 12519–12530.
- Jiang, N. and Li, L. (2016) Doubly robust off-policy value evaluation for reinforcement learning. In *International Conference on Machine Learning*, 652–661.
- Jin, J., Song, C., Li, H., Gai, K., Wang, J. and Zhang, W. (2018) Real-time bidding with multi-agent reinforcement learning in display advertising. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, 2193–2201. ACM.
- Kallus, N. and Uehara, M. (2019) Efficiently breaking the curse of horizon in off-policy evaluation with double reinforcement learning. *arXiv preprint arXiv:1909.05850*.

- (2020) Double reinforcement learning for efficient off-policy evaluation in markov decision processes. *Journal of Machine Learning Research*, **21**, 1–63.
- Kormushev, P., Calinon, S. and Caldwell, D. (2013) Reinforcement learning in robotics: Applications and real-world challenges. *Robotics*, **2**, 122–148.
- Luckett, D. J., Laber, E. B., Kahkoska, A. R., Maahs, D. M., Mayer-Davis, E. and Kosorok, M. R. (2019) Estimating dynamic treatment regimes in mobile health using v-learning. *J. Amer. Statist. Assoc.*
- Luedtke, A. R. and van der Laan, M. J. (2016) Statistical inference for the mean outcome under a possibly non-unique optimal treatment strategy. *Ann. Statist.*, **44**, 713–742.
- (2017) Evaluating the impact of treating the optimal subgroup. *Statistical methods in medical research*, **26**, 1630–1640.
- Maei, H. R., Szepesvári, C., Bhatnagar, S. and Sutton, R. S. (2010) Toward off-policy learning control with function approximation. In *ICML*, 719–726.
- Marling, C. and Bunescu, R. C. (2018) The ohiot1dm dataset for blood glucose level prediction. In *KHD@ IJCAI*, 60–63.
- McLeish, D. L. (1974) Dependent central limit theorems and invariance principles. *Ann. Probability*, **2**, 620–628.
- Meitz, M. and Saikkonen, P. (2019) Subgeometric ergodicity and  $\beta$ -mixing. *arXiv preprint arXiv:1904.07103*.
- Meyer, Y. (1992) *Wavelets and operators*, vol. 37 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge. Translated from the 1990 French original by D. H. Salinger.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G. et al. (2015) Human-level control through deep reinforcement learning. *Nature*, **518**, 529.

- Murphy, S. A. (2003) Optimal dynamic treatment regimes. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, **65**, 331–366.
- Puterman, M. L. (1994) *Markov decision processes: discrete stochastic dynamic programming*. Wiley Series in Probability and Mathematical Statistics: Applied Probability and Statistics. John Wiley & Sons, Inc., New York. A Wiley-Interscience Publication.
- Qian, M. and Murphy, S. A. (2011) Performance guarantees for individualized treatment rules. *Ann. Statist.*, **39**, 1180–1210.
- Riedmiller, M. (2005) Neural fitted q iteration—first experiences with a data efficient neural reinforcement learning method. In *European Conference on Machine Learning*, 317–328. Springer.
- Rodbard, D. (2009) Interpretation of continuous glucose monitoring data: glycemic variability and quality of glycemic control. *Diabetes technology & therapeutics*, **11**, S–55.
- Saikkonen, P. (2001) Stability results for nonlinear vector autoregressions with an application to a nonlinear error correction model. *Tech. rep.*, Discussion Papers, Interdisciplinary Research Project 373: Quantification and Simulation of Economic Processes.
- Schumaker, L. L. (1981) *Spline functions: basic theory*. John Wiley & Sons, Inc., New York. Pure and Applied Mathematics, A Wiley-Interscience Publication.
- Shi, C., Fan, A., Song, R. and Lu, W. (2018a) High-dimensional  $A$ -learning for optimal dynamic treatment regimes. *Ann. Statist.*, **46**, 925–957.
- Shi, C., Lu, W. and Song, R. (2020a) Breaking the curse of nonregularity with subagging— inference of the mean outcome under optimal treatment regimes. *Journal of Machine Learning Research*, **21**, 1–67.
- (2020b) A sparse random projection-based test for overall qualitative treatment effects. *Journal of the American Statistical Association*, **115**, 1201–1213.
- Shi, C., Song, R., Lu, W. and Fu, B. (2018b) Maximin projection learning for optimal treatment decision with heterogeneous individualized treatment effects. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **80**, 681–702.

- Shi, C., Wan, R., Chernozhukov, V. and Song, R. (2021) Deeply-debiased off-policy interval estimation. In *International Conference on Machine Learning*, accepted.
- 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.
- Sutton, R. S. and Barto, A. G. (2018) *Reinforcement learning: an introduction*. Adaptive Computation and Machine Learning. MIT Press, Cambridge, MA, second edn.
- Tang, Z., Feng, Y., Li, L., Zhou, D. and Liu, Q. (2020) Doubly robust bias reduction in infinite horizon off-policy estimation. In *International Conference on Learning Representations*.
- Thomas, P. S., Theodorou, G. and Ghavamzadeh, M. (2015) High-confidence off-policy evaluation. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*.
- Tropp, J. A. (2011) Freedman’s inequality for matrix martingales. *Electron. Commun. Probab.*, **16**, 262–270.
- (2012) User-friendly tail bounds for sums of random matrices. *Found. Comput. Math.*, **12**, 389–434.
- Tsao, C. W. and Vasani, R. S. (2015) Cohort profile: The framingham heart study (fhs): overview of milestones in cardiovascular epidemiology. *International journal of epidemiology*, **44**, 1800–1813.
- Tsybakov, A. B. (2004) Optimal aggregation of classifiers in statistical learning. *Ann. Statist.*, **32**, 135–166.
- Uehara, M., Huang, J. and Jiang, N. (2020) Minimax weight and q-function learning for off-policy evaluation. In *International Conference on Machine Learning*.
- Wang, J., He, X. and Xu, G. (2020) Debiased inference on treatment effect in a high-dimensional model. *Journal of the American Statistical Association*, **115**, 442–454.

- Xu, Z., Li, Z., Guan, Q., Zhang, D., Li, Q., Nan, J., Liu, C., Bian, W. and Ye, J. (2018) Large-scale order dispatch in on-demand ride-hailing platforms: A learning and planning approach. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 905–913. ACM.
- Zhang, B., Tsiatis, A. A., Laber, E. B. and Davidian, M. (2013) Robust estimation of optimal dynamic treatment regimes for sequential treatment decisions. *Biometrika*, **100**, 681–694.
- Zhang, Y., Laber, E. B., Davidian, M. and Tsiatis, A. A. (2018) Estimation of optimal treatment regimes using lists. *J. Amer. Statist. Assoc.*, **113**, 1541–1549.
- Zhao, Y.-Q., Zeng, D., Laber, E. B. and Kosorok, M. R. (2015) New statistical learning methods for estimating optimal dynamic treatment regimes. *J. Amer. Statist. Assoc.*, **110**, 583–598.

## A Some technical conditions

### A.1 More on Conditions A3

Define  $\mathcal{P}_X^t$  as the  $t$ -step transition kernel, i.e.  $\mathcal{P}_X^t(\mathcal{B}|x) = \Pr(X_{0,t} \in \mathcal{B}|X_{0,0} = x)$ . Geometric ergodicity implies that there exists some function  $M(\cdot)$  on  $\mathbb{X}$  and some constant  $\rho < 1$  such that  $\int_{x \in \mathbb{X}} M(x)\mu(x)dx < +\infty$  and

$$\|\mathcal{P}_X^t(\cdot|x) - \mu(\cdot)\|_{TV} \leq M(x)\rho^t, \quad \forall t \geq 0,$$

where  $\|\cdot\|_{TV}$  denotes the total variation norm.

### A.2 Conditions A3\*

We present the technical condition (A3\*) below. We assume the estimated policy satisfies  $\hat{\pi}_{\mathcal{I}} \in \Pi$  with probability 1, for any  $\mathcal{I}$ . For example, if Q-learning type algorithms are used and we approximate the optimal Q-function based on a linear model  $\Phi^\top(x, a)\beta$  with some basis function  $\Phi(x, a) \in \mathbb{R}^M$ . Then for any  $\beta$ , we can define a policy  $\pi(\beta)$  as follows:

$$\pi_\beta(a|x) = \begin{cases} 1, & \text{if } a = \text{sarg} \max_{a' \in \mathcal{A}} \Phi^\top(x, a')\beta, \\ 0, & \text{otherwise.} \end{cases}$$

Then we have  $\Pi = \{\pi_\beta : \beta \in \mathbb{R}^M\}$ .

(A3\*.) Assume (i) and (ii) hold if  $T \rightarrow \infty$  and (i) holds if  $T$  is bounded.

(i)  $\inf_{\pi \in \Pi} \lambda_{\min}[\sum_{t=0}^{T-1} \mathbb{E}\{\boldsymbol{\xi}_{0,t}\boldsymbol{\xi}_{0,t}^\top - \gamma^2 \mathbf{u}_\pi(X_{0,t}, A_{0,t})\mathbf{u}_\pi^\top(X_{0,t}, A_{0,t})\}] \geq \bar{c}T$  for some constant  $\bar{c} > 0$ .

(ii) The Markov chain  $\{X_{0,t}\}_{t \geq 0}$  is geometrically ergodic.

We remark that Condition A3\*(ii) is the same as A3(ii).

### A.3 More on the margin condition

To better understand Condition A5, we consider a simple scenario where  $\mathcal{A} = \{0, 1\}$ . Define

$\tau(x) = Q^{opt}(x, 1) - Q^{opt}(x, 0)$ . It follows that

$$\max_a Q^{opt}(x, a) - \max_{a' \in \mathcal{A} - \arg \max_a Q^{opt}(x, a)} Q^{opt}(x, a') = \begin{cases} |\tau(x)|, & \text{if } \tau(x) \neq 0, \\ +\infty, & \text{otherwise.} \end{cases}$$

As a result, (3.18) and (3.19) are equivalent to the followings:

$$\lambda \{x \in \mathbb{X} : 0 < |\tau(x)| \leq \varepsilon\} = O(\varepsilon^\alpha), \quad (\text{A.23})$$

$$\mathbb{G} \{x \in \mathbb{X} : 0 < |\tau(x)| \leq \varepsilon\} = O(\varepsilon^\alpha). \quad (\text{A.24})$$

Apparently, these two conditions hold when  $\inf_{x \in \mathbb{X}} |\tau(x)| > 0$ . They are satisfied in many other cases. For example, let  $d = 1$ . Consider

$$\tau(x) = \begin{cases} x^{1/\alpha}, & \text{if } x > 0, \\ 0, & \text{otherwise,} \end{cases}$$

for some  $\alpha > 0$ . Then, with some calculations, we can show

$$\lambda \{x \in \mathbb{X} : 0 < |\tau(x)| \leq \varepsilon\} \leq \lambda\{x : 0 < x < \varepsilon^\alpha\} = \varepsilon^\alpha.$$

This verifies (A.23). When  $\mathbb{G}$  has a bounded density function on  $\mathbb{X}$ , (A.24) is reduced to (A.23). If  $\mathbb{G}(\cdot)$  equals the Dirac measure  $\delta_x(\cdot)$ , then (A.24) automatically holds for any  $\alpha > 0$ .

## B Additional details regarding the method

### B.1 More on the CI in (3.17)

We begin by providing more details on the estimators  $\widehat{V}_{\mathcal{I}_{k+1}}(\widehat{\pi}_{\mathcal{I}_k}; \mathbb{G})$  and its standard error  $\widehat{\sigma}_{\mathcal{I}_{k+1}}(\widehat{\pi}_{\mathcal{I}_k}; \mathbb{G})$ . In general, for a given  $\mathcal{I} \subseteq \mathcal{I}_0$  and any policy  $\pi$ , we define  $\widehat{V}_{\mathcal{I}}(\pi; \mathbb{G})$  and

$\widehat{\sigma}_{\mathcal{I}}(\pi; \mathbb{G})$  as

$$\begin{aligned}\widehat{V}_{\mathcal{I}}(\pi; \mathbb{G}) &= \left\{ \int_{x \in \mathbb{X}} \mathbf{U}_{\pi}(x) \mathbb{G}(dx) \right\}^{\top} \widehat{\boldsymbol{\beta}}_{\mathcal{I}, \pi}, \\ \widehat{\sigma}_{\mathcal{I}}^2(\pi; \mathbb{G}) &= \left\{ \int_{x \in \mathbb{X}} \mathbf{U}_{\pi}(x) \mathbb{G}(dx) \right\}^{\top} \widehat{\boldsymbol{\Sigma}}_{\mathcal{I}, \pi}^{-1} \widehat{\boldsymbol{\Omega}}_{\mathcal{I}, \pi} (\widehat{\boldsymbol{\Sigma}}_{\mathcal{I}, \pi}^{\top})^{-1} \left\{ \int_{x \in \mathbb{X}} \mathbf{U}_{\pi}(x) \mathbb{G}(dx) \right\},\end{aligned}$$

where

$$\begin{aligned}\widehat{\boldsymbol{\beta}}_{\mathcal{I}, \pi} &= (\widehat{\beta}_{\mathcal{I}, \pi, 1}^{\top}, \dots, \widehat{\beta}_{\mathcal{I}, \pi, m}^{\top})^{\top} = \frac{1}{|\mathcal{I}|} \sum_{(i,t) \in \mathcal{I}} \widehat{\boldsymbol{\Sigma}}_{\mathcal{I}, \pi}^{-1} \boldsymbol{\xi}_{i,t} Y_{i,t}, \quad \widehat{\boldsymbol{\Sigma}}_{\mathcal{I}, \pi} = \frac{1}{|\mathcal{I}|} \sum_{(i,t) \in \mathcal{I}} \boldsymbol{\xi}_{i,t} (\boldsymbol{\xi}_{i,t} - \gamma \mathbf{U}_{\pi, i, t+1})^{\top}, \\ \widehat{\boldsymbol{\Omega}}_{\mathcal{I}, \pi} &= \frac{1}{|\mathcal{I}|} \sum_{(i,t) \in \mathcal{I}} \boldsymbol{\xi}_{i,t} \boldsymbol{\xi}_{i,t}^{\top} \left\{ Y_{i,t} + \gamma \sum_{a \in \mathcal{A}} \Phi_L^{\top}(X_{i, t+1}) \widehat{\beta}_{\mathcal{I}, \pi, a} \pi(a | X_{i, t+1}) - \Phi_L^{\top}(X_{i,t}) \widehat{\beta}_{\pi, A_{i,t}} \right\}^2,\end{aligned}$$

and  $|\mathcal{I}|$  stands for the number of elements in  $\mathcal{I}$ .

## B.2 Value difference between the target and behavior policy

In this section, we outline a method to evaluate the value difference function between the target and behavior policy. We first consider the scenario where the target policy is a fixed policy. We next consider the scenario where the target policy is an estimated optimal policy. To simplify the presentation, we assume  $T_1 = \dots = T_n = T$ . The proposed method can be similarly extended to on-policy settings.

### B.2.1 Inference of the value difference under a fixed policy

Consider a data-independent policy  $\pi$ . We aim to evaluate the value difference function  $\text{VD}(\pi; x) = V(\pi; x) - V(b; x)$  where  $b$  is the unknown behavior policy. We first apply our method in Section 3.1.2 to compute an estimator value function  $\widehat{V}(\pi; x)$  for  $V(\pi; x)$ .

To estimate  $V(b; x)$ , we observe that the Q-function  $Q(b; x, a)$  satisfies the Bellman equation,  $\mathbb{E}[\{Y_{i,t} + \gamma Q(b; X_{i, t+1}, A_{i, t+1}) - Q(b; X_{i,t}, A_{i,t})\} | X_{i,t}, A_{i,t}] = 0$ . We approximate  $Q(b; x, a)$  based on linear sieves  $\Phi_L^{\top}(x) \boldsymbol{\beta}_{b,a}^*$ . Similar to Section 3.1.2,  $\boldsymbol{\beta}_b = (\beta_{b,1}^*, \dots, \beta_{b,m}^*)^{\top}$  can be estimated by

$$\widehat{\boldsymbol{\beta}}_b = (\widehat{\beta}_{b,1}^{\top}, \dots, \widehat{\beta}_{b,m}^{\top})^{\top} = \underbrace{\left\{ \frac{1}{nT} \sum_{i=1}^n \sum_{t=0}^{T-1} \boldsymbol{\xi}_{i,t} (\boldsymbol{\xi}_{i,t} - \gamma \boldsymbol{\xi}_{i, t+1})^{\top} \right\}^{-1}}_{\widehat{\boldsymbol{\Sigma}}_b} \left( \frac{1}{nT} \sum_{i=1}^n \sum_{t=0}^{T-1} \boldsymbol{\xi}_{i,t} Y_{i,t} \right).$$

The resulting estimates for  $Q(b; x, a)$  can be derived as  $\Phi_L^\top(x)\widehat{\beta}_{b,a}$ . The corresponding estimator for  $V(b, x)$  is given by  $\widehat{V}(b, x) = \sum_a \widehat{b}(a|x)\Phi_L^\top(x)\widehat{\beta}_{b,a}$  where  $\widehat{b}(a|x)$  denotes the sieve estimator  $\Phi_L^\top(x)\widehat{\alpha}_a$  for  $b(a|x)$  where

$$\widehat{\alpha}_a = \underbrace{\left\{ \frac{1}{nT} \sum_{i=1}^n \sum_{t=0}^{T-1} \Psi_L(X_{i,t}) \Psi_L^\top(X_{i,t}) \right\}^{-1}}_{\widehat{\Psi}} \left\{ \frac{1}{nT} \sum_{i=1}^n \sum_{t=0}^{T-1} \Psi_L(X_{i,t}) \mathbb{I}(A_{i,t} = a) \right\}.$$

This yields the estimator for the value difference  $\widehat{\text{VD}}(\pi; x) = \widehat{V}(\pi; x) - \widehat{V}(b; x)$ .

We next derive a confidence interval for  $\text{VD}(\pi; x)$  based on  $\widehat{\text{VD}}(\pi; x)$ . Similar to the proof of Theorem 1, we can show  $\sqrt{nT}\{\widehat{\text{VD}}(\pi; x) - \text{VD}(\pi; x)\}$  is equivalent to

$$\begin{aligned} & \mathbf{U}_\pi^\top(x) \Sigma_\pi^{-1} \left( \frac{1}{\sqrt{nT}} \sum_{i=1}^n \sum_{t=0}^{T-1} \boldsymbol{\xi}_{i,t} \varepsilon_{i,t} \right) - \mathbf{U}_b^\top(x) \Sigma_b^{-1} \left( \frac{1}{\sqrt{nT}} \sum_{i=1}^n \sum_{t=0}^{T-1} \boldsymbol{\xi}_{i,t} \varepsilon_{i,t}^* \right) \\ & - \sum_a Q(b; a, x) \Psi_L^\top(x) \Psi^{-1} \left[ \frac{1}{\sqrt{nT}} \sum_{i=1}^n \sum_{t=0}^{T-1} \Psi_L(X_{i,t}) \{ \mathbb{I}(A_{i,t} = a) - b(a|X_{i,t}) \} \right], \end{aligned} \quad (\text{B.25})$$

where  $\varepsilon_{i,t}^*$  denotes the temporal difference error  $Y_{i,t} + \gamma Q(b; X_{i,t+1}, A_{i,t+1}) - Q(b; X_{i,t}, A_{i,t})$  and  $\Psi$  denotes the population limit of  $\widehat{\Psi}$ . Note that the RHS can be rewritten as  $(nT)^{-1/2} \sum_{i=1}^n \sum_{t=0}^{T-1} \psi_{i,t}$  that corresponds to a sum of martingale difference. Its variance can be consistently estimated by  $\widehat{\sigma}^{*2}(\pi; x) = (nT)^{-1} \sum_{i=1}^n \sum_{t=0}^{T-1} \widehat{\psi}_{i,t}^2$  where  $\widehat{\psi}_{i,t}$  denotes some consistent estimator for  $\psi_{i,t}$  based on  $\widehat{Q}(\pi; \cdot, \cdot)$ ,  $\widehat{Q}(b; \cdot, \cdot)$  and  $\widehat{b}$ . The confidence interval for  $\text{VD}(\pi; x)$  is given by

$$[\widehat{\text{VD}}(\pi; x) - z_{\alpha/2}(nT)^{-1/2}\widehat{\sigma}^*(\pi; x), \widehat{\text{VD}}(\pi; x) + z_{\alpha/2}(nT)^{-1/2}\widehat{\sigma}^*(\pi; x)].$$

## B.2.2 Inference of the value difference under an estimated optimal policy

We begin by dividing the data into  $K$  non-overlapping subsets  $\cup_{k=1}^K \mathcal{I}_k$ . Similar to Section 3.2.2, we construct the value difference estimator by

$$\widetilde{\text{VD}}(x) = \left\{ \sum_{k=1}^{K-1} \frac{1}{\widehat{\sigma}_{\mathcal{I}_{k+1}}^*(\widehat{\pi}_{\mathcal{I}_k}; x)} \right\}^{-1} \left\{ \sum_{k=1}^{K-1} \frac{\widehat{\text{VD}}_{\mathcal{I}_{k+1}}(\widehat{\pi}_{\mathcal{I}_k}; x)}{\widehat{\sigma}_{\mathcal{I}_{k+1}}^*(\widehat{\pi}_{\mathcal{I}_k}; x)} \right\},$$

where  $\widehat{\text{VD}}_{\mathcal{I}_k}$  and  $\widehat{\sigma}_{\mathcal{I}_k}^*$  denote the versions of  $\text{VD}$  and  $\widehat{\sigma}^*$  based on samples in  $\mathcal{I}_k$  only. The corresponding confidence interval is given by

$$[\widetilde{\text{VD}}(x) - z_{\alpha/2}\{nT(K-1)/K\}^{-1/2}\widetilde{\sigma}^*(x), \widetilde{\text{VD}}(x) + z_{\alpha/2}\{nT(K-1)/K\}^{-1/2}\widetilde{\sigma}^*(x)],$$



where  $\tilde{\sigma}^*(x) = (K - 1)\{\sum_{k=1}^{K-1} \hat{\sigma}_k^{*-1}(\hat{\pi}_{\mathcal{I}_k}; x)\}^{-1}$ .

Finally, we remark that such a confidence interval might not be valid in the extreme case where the behavior policy is equal to a deterministic optimal policy. To elaborate, notice that when the behavior policy is deterministic, the second line of (B.25) equal zero. In addition, when  $\pi = b = \pi^{opt}$  for some optimal policy  $\pi^{opt}$ , the first line equals zero as well. In that case,  $\widehat{\text{VD}}(\pi; x)$  would have a degenerate distribution. Suppose the estimated optimal policy is consistent for  $\pi^{opt}$ . Then  $\widetilde{\text{VD}}(x)$  might not have a tractable limiting distribution, leading to an invalid confidence interval.

To address this concern, we could redefine the inverse weights  $\hat{\sigma}_{\mathcal{I}_{k+1}}^*(\hat{\pi}_{\mathcal{I}_k}; x)$  by  $\hat{\sigma}_{\mathcal{I}_{k+1}}^*(\hat{\pi}_{\mathcal{I}_k}; x, \delta) = \max\{\hat{\sigma}_{\mathcal{I}_{k+1}}^*(\hat{\pi}_{\mathcal{I}_k}; x), \delta\}$  for some  $\delta > 0$ , as in Luedtke and van der Laan (2017). This guarantees that these inverse weights are strictly greater than zero. A similar approach is employed by Shi et al. (2020b) for testing the overall qualitative treatment effects in single-stage decision making. In addition, one could allow  $\delta$  to depend on  $N$  and  $T$ . The resulting confidence interval would be valid as long as  $\delta \gg (NT)^{-1/6}$  (see e.g., Theorem 3.1 of Shi et al., 2020b). However, a potential limitation is that it would yield a conservative confidence interval when the truncation is active, as discussed in Luedtke and van der Laan (2017).

### B.3 Double fitted Q-iteration

In this section, we introduce our algorithm for computing the estimated optimal policy in our numerical studies. The proposed algorithm is based on FQI that recursively updates the estimated optimal Q-function by some supervised learning method (see Example 2 in Section 3.2.3). In FQI, at each iteration, a maximization over estimated Q-function is used as an estimate of the maximum of the true Q-function. This can lead to a significant positive bias (Sutton and Barto, 2018). Hasselt (2010) proposed a double Q-learning method to reduce the maximization bias. Here, we apply similar ideas to FQI to compute the estimated optimal policy. We use a pseudocode to summarize our algorithm below.

In Algorithm 1, we can apply any non-parametric models  $Q(\cdot, \cdot; \theta)$  indexed by  $\theta$  to model the optimal Q-function. In our implementation, we set  $Q(\cdot, \cdot; \cdot)$  to be a linear combination of tensor product B-spline basis functions.

---

**Algorithm 1:** Double Fitted Q-iteration Algorithm
 

---

- Input:**  $\{(X_{i,t}, A_{i,t}, Y_{i,t}, X_{i,t+1})\}_{0 \leq t < T_i, 1 \leq i \leq n}$ , initialize parameters  $\widehat{\theta}_A, \widehat{\theta}_B$ .
- 1 **while** *not convergence* **do**
  - 2     For  $0 \leq t < T_i, 1 \leq i \leq n$ , construct the target value:
 
$$\widehat{Q}_{i,t}^A = Y_{i,t} + \gamma Q(X_{i,t+1}, \arg \max_{a'} Q(X_{i,t+1}, a'; \widehat{\theta}_A), \widehat{\theta}_B)$$

$$\widehat{Q}_{i,t}^B = Y_{i,t} + \gamma Q(X_{i,t+1}, \arg \max_{a'} Q(X_{i,t+1}, a'; \widehat{\theta}_B), \widehat{\theta}_A)$$
  - 3     Update  $\widehat{\theta}_A$  :  $\widehat{\theta}_A \leftarrow \arg \min_{\theta_A} \sum_i \sum_t \|Q(X_{i,t}, A_{i,t}; \theta_A) - \widehat{Q}_{i,t}^A\|^2$
  - 4     Update  $\widehat{\theta}_B$  :  $\widehat{\theta}_B \leftarrow \arg \min_{\theta_B} \sum_i \sum_t \|Q(X_{i,t}, A_{i,t}; \theta_B) - \widehat{Q}_{i,t}^B\|^2$
- 

## C Additional technical details

### C.1 Additional details regarding Condition A4

When  $\nu_0 = \mu$ , the density function of  $X_{0,1}$  equals  $\mu$  as well. By Jensen's inequality, we have for any  $\mathbf{v} \in \mathbb{R}^{mL}$  that

$$\{\mathbf{v}^\top \mathbf{u}_\pi(x, a)\}^2 \leq \mathbb{E}[\{\mathbf{v}^\top \mathbf{U}_\pi(X_{0,1})\}^2 | X_{0,0} = x, A_{0,0} = a],$$

and hence

$$\int_{x \in \mathbb{X}} \{\mathbf{v}^\top \mathbf{u}_\pi(x, a)\}^2 b(a|x) \mu(x) dx \leq \int_{x \in \mathbb{X}} \{\mathbf{v}^\top \mathbf{U}_\pi(x)\}^2 \mu(x) dx.$$

The matrix

$$\int_{x \in \mathbb{X}} \sum_{a \in \mathcal{A}} \{\mathbf{U}_\pi(x) \mathbf{U}_\pi^\top(x) - \mathbf{u}_\pi(x, a) \mathbf{u}_\pi^\top(x, a)\} b(a|x) \mu(x) dx$$

is positive semidefinite. It follows that

$$\begin{aligned} & \lambda_{\min} \left[ \int_{x \in \mathbb{X}} \sum_{a \in \mathcal{A}} \{\boldsymbol{\xi}(x, a) \boldsymbol{\xi}^\top(x, a) - \gamma^2 \mathbf{u}_\pi(x, a) \mathbf{u}_\pi^\top(x, a)\} b(a|x) \mu(x) dx \right] \\ & \geq \lambda_{\min} \left[ \int_{x \in \mathbb{X}} \sum_{a \in \mathcal{A}} \{\boldsymbol{\xi}(x, a) \boldsymbol{\xi}^\top(x, a) - \gamma^2 \mathbf{U}_\pi(x) \mathbf{U}_\pi^\top(x)\} b(a|x) \mu(x) dx \right]. \end{aligned}$$

When  $\pi$  is a deterministic policy,  $\sum_{a \in \mathcal{A}} \boldsymbol{\xi}(x, a) \boldsymbol{\xi}^\top(x, a) b(a|x) - \gamma^2 \mathbf{U}_\pi(x) \mathbf{U}_\pi^\top(x)$  is a block diagonal matrix. To show A4(i) holds, it suffices to show

$$\lambda_{\min}[\Phi_L(x) \Phi_L^\top(x) \{b(a|x) - \gamma^2 \pi(a|x)\} \mu(x) dx] > 0, \quad \forall a \in \mathcal{A}.$$

Suppose  $b$  is the  $\epsilon$ -greedy policy with respect to  $\pi$ , i.e,  $b(a|x) = \epsilon m^{-1} + (1 - \epsilon)\pi(a|x)$ , for any  $a \in \{1, \dots, m\}$  and  $\epsilon$  satisfies  $\epsilon \leq 1 - \gamma^2$ , we have

$$\lambda_{\min} \left[ \int_{x \in \mathbb{X}} \Phi_L(x) \Phi_L^\top(x) \{b(a|x) - \gamma^2 \pi(a|x)\} \mu(x) dx \right] \geq \frac{1}{m} \lambda_{\min} \left\{ \int_{x \in \mathbb{X}} \Phi_L(x) \Phi_L^\top(x) \mu(x) dx \right\}.$$

Suppose A3 holds. It suffices to require

$$\lambda_{\min} \left\{ \int_{x \in \mathbb{X}} \Phi_L(x) \Phi_L^\top(x) dx \right\} > 0. \quad (\text{C.26})$$

The condition in (C.26) is automatically satisfied when A2 holds (see, e.g., Burman and Chen, 1989; Chen and Christensen, 2015).

## C.2 Additional details on the variance comparison

We consider a randomized study where  $b(a|x)$  is a constant function of  $x$ . In addition, we assume the target policy is nondynamic, i.e.,  $\pi(a^*|x) = 1$  for some  $1 \leq a^* \leq m$  and any  $x$ .

We impose the following conditions.

(C1) The process  $\{X_{0,t}\}_{t \geq 0}$  is stationary.

(C2) The temporal difference error  $\varepsilon_{0,t}$  is independent of  $(X_{0,t}, A_{0,t})$ .

(C3)  $\mathbb{E}\{\Phi_L(X_{0,t+1})\Phi_L^\top(X_{0,t+1})|A_{0,t} = a^*\} = \mathbb{E}\Phi_L(X_{0,t+1})\Phi_L^\top(X_{0,t+1})$ .

We make some remarks. First, Condition (C1) is imposed to simplify the presentation. The same results hold as long as  $\{X_{0,t}\}_t$  will converge to its stationary distribution. Second, the variances of our estimator and DRL are very difficult to analyse in general. Conditions (C2)-(C3) are imposed to simplify the calculation. Even when these conditions are violated, we expect the variance of the proposed estimator will be smaller in general, as reflected in our numerical study.

**Theorem 5** *Assume (C1)-(C3) hold. Then the asymptotic variance of the DRL estimator is at least  $h^{-1}(\Phi_L)\{1 + \text{Var}(w(\pi, X_{0,t}))\}$  times larger than the proposed estimator where  $h(\Phi_L) = \int_x \Phi_L(x) \mathbb{G}(dx) \{\mathbb{E}\Phi_L(X_{0,t})\Phi_L^\top(X_{0,t})\}^{-1} \int_x \Phi_L(x) \mathbb{G}(dx)$  and  $w$  is the marginalized density ratio (Kallus and Uehara, 2019).*

We next investigate the value of the factor  $h(\Phi_L)$ . We set  $\mathbb{G}$  and the stationary distribution of  $X_{0,t}$ , i.e.,  $\mu$  to a uniform distribution on  $[0, 1]$ . We consider a polynomial basis




Figure 6:  $h(\Phi_{0,L})$  with different choices of  $\Phi_{0,L}$ .

function and a B-spline basis function. Figure 6 depicts the value of this factor with different choices of  $L$ . We also tried several other combinations of  $\mathbb{G}$  and  $\mu$ , and find this factor is in general smaller than or very close to 1. Since  $\text{Var}(w(\pi, X_{0,t})) > 0$ , Theorem 5 implies that the proposed estimator achieves smaller variance.

We next sketch a few lines to prove Theorem 5. Based on Theorem 1, the asymptotic variance of our estimator is given by

$$(nT)^{-1} \left\{ \int_{x \in \mathbb{X}} \mathbf{U}_\pi(x) \mathbb{G}(dx) \right\}^\top \boldsymbol{\Sigma}_\pi^{-1} \boldsymbol{\Omega} (\boldsymbol{\Sigma}_\pi^{-1})^\top \left\{ \int_{x \in \mathbb{X}} \mathbf{U}_\pi(x) \mathbb{G}(dx) \right\}, \quad (\text{C.27})$$

where  $\boldsymbol{\Omega}$  is defined in Step 2 of the proof of Theorem 1. Under (C1) and (C2), we have  $\boldsymbol{\Omega} = \mathbb{E} \boldsymbol{\xi}_{0,t} \boldsymbol{\xi}_{0,t}^\top \varepsilon_{0,t}^2 = \sigma_*^2 \mathbb{E} \boldsymbol{\xi}_{0,t} \boldsymbol{\xi}_{0,t}^\top$  where  $\sigma_*^2$  is the variance of  $\varepsilon_{0,t}$ .

Using similar arguments in the proof of Theorem 16 of Kallus and Uehara (2019), we can show that the asymptotic variance of the DRL estimator equals

$$\frac{1}{nT(1-\gamma)^2} \mathbb{E} w^2(X_{0,0}) \frac{\mathbb{I}(A_{0,0} = a^*)}{b^2(A_{0,0}|X_{0,0})} \varepsilon_{0,0}^2.$$

Under the given conditions, the above variance is equal to  $\{nT(1-\gamma)^2 p_{a^*}\}^{-1} \sigma_*^2 \{1 + \text{Var}(\omega(X_{0,0}))\}$  where  $p_{a^*} = \Pr(A_{0,t} = a^*)$ . Consequently, it suffices to show

$$\left\{ \int_{x \in \mathbb{X}} \mathbf{U}_\pi(x) \mathbb{G}(dx) \right\}^\top \boldsymbol{\Sigma}_\pi^{-1} \mathbb{E} \boldsymbol{\xi}_{0,t} \boldsymbol{\xi}_{0,t}^\top (\boldsymbol{\Sigma}_\pi^{-1})^\top \left\{ \int_{x \in \mathbb{X}} \mathbf{U}_\pi(x) \mathbb{G}(dx) \right\} \leq \frac{h(\Phi_{0,L})}{(1-\gamma)^2 p_{a^*}}. \quad (\text{C.28})$$

Since  $\pi$  is a nondynamic policy,  $\int_{x \in \mathbb{X}} \mathbf{U}_\pi(x) \mathbb{G}(dx)$  is a sparse vector that takes the following form:

$$\left\{ \underbrace{0, \dots, 0}_{(a^*-1)L \text{ many } 0\text{s}}, \int_{\mathbb{X}} \Phi_L(x) \mathbb{G}(dx), \underbrace{0, \dots, 0}_{(m-a^*)L \text{ many } 0\text{s}} \right\}.$$

By the definition of  $\Sigma_\pi$  and  $\xi_{0,t}$ , the left-hand-side of (C.28) is equal to

$$\begin{aligned} & \left\{ \int_{\mathbb{X}} \Phi_L(x) \mathbb{G}(dx) \right\}^\top [\mathbb{E} \Phi_L(X_{0,t}) \mathbb{I}(A_{0,t} = a^*) \{\Phi_L(X_{0,t}) - \gamma \Phi_L(X_{0,t+1})\}^\top]^{-1\top} \Omega_{a^*} \\ & \times [\mathbb{E} \Phi_L(X_{0,t}) \mathbb{I}(A_{0,t} = a^*) \{\Phi_L(X_{0,t}) - \gamma \Phi_L(X_{0,t+1})\}^\top]^{-1} \left\{ \int_{\mathbb{X}} \Phi_L(x) \mathbb{G}(dx) \right\}, \end{aligned}$$

where  $\Omega_{a^*} = \mathbb{E} \Phi_L(X_{0,t}) \mathbb{I}(A_{0,t} = a^*) \Phi_L^\top(X_{0,t})$ . Using similar arguments in C.1, we can show

$$\begin{aligned} & \mathbf{a}^\top \Omega_{a^*}^{-1/2} [\mathbb{E} \Phi_L(X_{0,t}) \mathbb{I}(A_{0,t} = a^*) \{\Phi_L(X_{0,t}) - \gamma \Phi_L(X_{0,t+1})\}] \Omega_{a^*}^{-1/2} \mathbf{a} \\ & = \mathbf{a}^\top \Omega_{a^*}^{-1/2} \mathbb{E} \Phi_L(X_{0,t}) \mathbb{I}(A_{0,t} = a^*) \Phi_L^\top(X_{0,t}) \Omega_{a^*}^{-1/2} \mathbf{a} \\ & - \gamma \mathbf{a}^\top \Omega_{a^*}^{-1/2} \mathbb{E} \Phi_L(X_{0,t}) \mathbb{I}(A_{0,t} = a^*) \Phi_L^\top(X_{0,t+1}) \Omega_{a^*}^{-1/2} \mathbf{a} \\ & \geq \|\mathbf{a}\|_2^2 - \gamma \|\mathbf{a}\|_2 \sqrt{\mathbf{a}^\top \Omega_{a^*}^{-1/2} \mathbb{E} \mathbb{I}(A_{0,t} = a^*) \Phi_L(X_{0,t+1}) \Phi_L^\top(X_{0,t+1}) \Omega_{a^*}^{-1/2} \mathbf{a}} = (1 - \gamma) \|\mathbf{a}\|_2^2, \end{aligned}$$

under (C3). Consequently, we have

$$\|\Omega_{a^*}^{1/2} [\mathbb{E} \Phi_L(X_{0,t}) \mathbb{I}(A_{0,t} = a^*) \{\Phi_L(X_{0,t}) - \gamma \Phi_L(X_{0,t+1})\}]^{-1} \Omega_{a^*}^{1/2}\|_2 \leq (1 - \gamma).$$

As such, the left-hand-side of (C.28) is upper bounded by

$$\begin{aligned} (1 - \gamma)^{-2} \left\{ \mathbb{E} \int_{\mathbb{X}} \Phi_L(x) \mathbb{G}(dx) \right\}^\top \Omega_{a^*}^{-1} \left\{ \int_{\mathbb{X}} \Phi_L(x) \mathbb{G}(dx) \right\} & = (1 - \gamma)^{-2} p_{a^*}^{-1} \left\{ \int_{\mathbb{X}} \Phi_L(x) \mathbb{G}(dx) \right\}^\top \\ & \times \{\mathbb{E} \Phi_{0,L}(X_{0,t}) \Phi_{0,L}^\top(X_{0,t})\}^{-1} \left\{ \int_{\mathbb{X}} \Phi_L(x) \mathbb{G}(dx) \right\} = (1 - \gamma)^{-2} p_{a^*}^{-1} h(\Phi_{0,L}). \end{aligned}$$

This completes the proof.

### C.3 Additional details on on-policy evaluation

In this section, we show our proposed CI in Section 4 achieves nominal converge. To simplify the analysis, we focus on the setting where  $K$  is finite,  $T(1) = \dots = T(K) = T$  and  $L(1) = \dots = L(K) = L$ . When  $K$  diverges, the sequences  $\{T(k)\}_{k \geq 1}$  and  $\{L(k)\}_{k \geq 1}$

shall be properly chosen to reduce the bias of the value estimates. We leave this for future research.

Similar to Appendix A.2, we assume the estimated policy  $\widehat{\pi}_{\mathcal{I}} \in \Pi$  with probability 1, for any  $\mathcal{I}$ . In on-policy settings, the behavior policy  $b_{\pi}$  is a function of the estimated policy  $\pi \in \Pi$ . For instance, when an  $\epsilon$ -greedy policy is used to determine the behavior policy, then we have  $b_{\pi} = (1 - \epsilon)\pi + \epsilon\pi^*$  where  $\pi^*$  denotes a uniform random policy. Let  $\mathcal{B} = \{b_{\pi} : \pi \in \Pi\}$ .

For any behavior policy  $b \in \mathcal{B}$ , consider a Markov chain  $\{(X_{0,t,b}, A_{0,t,b})\}_{t \geq 0}$  generated by this behavior policy. Let  $Y_{0,t,b}$  be the realization of the immediate reward at time  $t$ . Let  $\mu_b$  denote the limiting distribution of the Markov chain  $\{X_{0,t,b}\}_{t \geq 0}$ , and  $\mathcal{P}_X^{t,b}(\cdot|x)$  be its  $t$ -step transition kernel. For any  $x \in \mathbb{X}, a \in \mathcal{A}$ , define  $\omega_{\pi,b}(x, a)$  as

$$\mathbb{E} \left[ \left\{ Y_{0,0,b} + \gamma \sum_{a \in \mathcal{A}} \pi(a|X_{0,1,b}) Q(\pi; X_{0,1,b}, a) - Q(\pi; X_{0,0,b}, A_{0,0,b}) \right\}^2 \mid X_{0,0,b} = x, A_{0,0,b} = a \right].$$

We introduce the following conditions.

(A2'.) Assume  $\nu_0$  and  $q$  are uniformly bounded away from 0 and  $\infty$  on their supports.

(A3'.) Assume (i) and (ii) hold if  $T \rightarrow \infty$  and (iii) holds if  $T$  is bounded.

(i)  $\inf_{\pi \in \Pi, b \in \mathcal{B}} \lambda_{\min}[\int_{x \in \mathbb{X}} \sum_{a \in \mathcal{A}} \{\boldsymbol{\xi}(x, a)\boldsymbol{\xi}^{\top}(x, a) - \gamma^2 \mathbf{u}_{\pi}(x, a)\mathbf{u}_{\pi}^{\top}(x, a)\} b(a|x)\mu_b(x)dx] \geq \bar{c}$  for some constant  $\bar{c} > 0$ .

(ii) There exists some function  $M(\cdot)$  on  $\mathbb{X}$  and some constant  $\rho < 1$  such that

$$\sup_{\pi \in \Pi} \int_{x \in \mathbb{X}} M(x)\mu_{b_{\pi}}(x)dx < +\infty,$$

and

$$\sup_{\pi \in \Pi} \|\mathcal{P}_X^{t,b_{\pi}}(\cdot|x) - \mu_{b_{\pi}}(\cdot)\|_{TV} \leq M(x)\rho^t, \quad \forall t \geq 0.$$

(iii) There exists some constant  $\bar{c} > 0$  such that

$$\inf_{\pi \in \Pi, b \in \mathcal{B}} \lambda_{\min} \left[ \sum_{t=0}^{T-1} \mathbb{E} \{ \boldsymbol{\xi}(X_{0,t,b}, A_{0,t,b})\boldsymbol{\xi}(X_{0,t,b}, A_{0,t,b})^{\top} - \gamma^2 \mathbf{u}_{\pi}(X_{0,t,b}, A_{0,t,b})\mathbf{u}_{\pi}^{\top}(X_{0,t,b}, A_{0,t,b}) \} \right] \geq \bar{c}T.$$

(A4') For any  $1 \leq k \leq K$ , we have  $\mathbb{E}|V(\widehat{\pi}_{\bar{\mathcal{I}}_k}; \mathbb{G}) - V(\pi^*; \mathbb{G})| = O(|\bar{\mathcal{I}}_k|^{-b_0})$ , for some  $b_0 > 1/2$  such that  $(nT)^{b_0-1/2} \gg \|\int_x \Phi_L(x)\mathbb{G}(dx)\|_2^{-1}$ , where the big- $O$  term is uniform in  $\mathcal{I}$ .

**Theorem 6** *Assume A1, A2'-A4' hold. Suppose  $L = o\{\sqrt{nT}/\log(nT)\}$  and  $L^{2p/d} \gg nT\{1 + \|\int_x \Phi_L(x)\mathbb{G}(dx)\|_2^{-2}\}$ . Assume there exists some constant  $c_0 \geq 1$  such that  $\omega_{\pi,b}(x, a) \geq c_0^{-1}$  for any  $x, a, \pi, b$  and  $Pr(\max_{0 \leq t \leq T-1} |Y_{0,t}| \leq c_0) = 1$ . Then as either  $n \rightarrow \infty$  or  $T \rightarrow \infty$ ,*

$$\begin{aligned} \sqrt{nT(K-1)}\tilde{\sigma}^{-1}(\mathbb{G})\{\tilde{V}(\mathbb{G}) - V(\hat{\pi}; \mathbb{G})\} &\xrightarrow{d} N(0, 1), \\ \sqrt{nT(K-1)}\tilde{\sigma}^{-1}(\mathbb{G})\{\tilde{V}(\mathbb{G}) - V(\pi^*; \mathbb{G})\} &\xrightarrow{d} N(0, 1). \end{aligned}$$

Proof of Theorem 6 is omitted for brevity.

## D Additional numerical results

### D.1 Data-adaptive selection of $L$

We apply the proposed method detailed in Section 7.2 to Scenario (B) where the treatment assignment mechanism depends on the observed state, to investigate the finite sample performance of the resulting CI. Specifically, we apply the random forest algorithm to learn the conditional mean function  $\mu$  and the behavior policy  $b$ . We assume  $\Sigma(a, x)$  is a constant function of  $(a, x)$  and estimated it by

$$\hat{\Sigma} = \frac{1}{nT} \sum_{i,t} \{X_{i,t+1} - \hat{\mu}(A_{i,t}, X_{i,t})\} \{X_{i,t+1} - \hat{\mu}(A_{i,t}, X_{i,t})\}^\top.$$

We use the tensor product B-spline basis for  $\Phi_L$ , as in Section 5. Note that the state is a two-dimensional vector,  $L$  is selected among the set  $\{4^2, 5^2, 6^2, 7^2, 8^2, 9^2\}$ . Specifically, we choose  $L$  such that the resulting CI is the shortest among all CIs whose coverage probabilities are above 93%. If no such CI exists, we select the CI with the highest coverage probability.

We report the ECP and AL of the resulting CI in the left and middle panels of Figure 7. It can be seen that ECP is close to the nominal level in all cases and AL decays as either  $n$  or  $T$  increases. In the right panel of Figure 7, we report the number of basis functions that is being selected most by the proposed method as a function of  $n$  and  $T$  (denote by  $L(n, T)$ ). It is clear from Figure 7 that  $L(n, T)$  increases with the total number of observations  $nT$ . This is consistent with the following intuition: as  $nT$  increases, more basis functions are needed to reduce the approximation error and guarantee the nominal coverage of the resulting CI.




Figure 7: Empirical coverage probabilities (ECP) and average lengths (AL) of CIs constructed by the proposed method detailed in Section 7.2 as well as the number of basis functions being selected most, with different combinations of  $n$  and  $T$ .

## D.2 Sensitivity analysis

### D.2.1 Sensitivity test for $\eta$

In this section, we conduct the sensitivity test for the parameter  $\eta$  in the number of basis  $L = \lfloor (nT)^\eta \rfloor$ . We consider the simulation of the off-policy evaluation with a fixed target policy in Section 5.1. For scenario (A), (B) and (C), we set  $n = 100$ ,  $T = 100$  and the different  $\eta$ 's are chosen from  $(0.25, 0.30, 0.35, 0.40, 3/7, 0.45)$ . The result of the ECPs are plotted in Figure 8 where all the ECPs are close to the nominal coverage rate 0.95. It shows that the results of the coverage are not sensitive to the different choices of  $\eta$ .

### D.2.2 Sensitivity test for $\gamma$

In Figure 9, we report the ECP and AL of the proposed CI and the MSE of our value estimate under Scenario B where the target policy is fixed, with  $\gamma = 0.3$  and  $0.7$ . In Figure 10, we report the ECP and AL of the proposed CI and the MSE of our value estimate under Scenario B where the target policy is an estimated optimal policy, with  $\gamma = 0.3$  and  $0.7$ . It can be seen that findings are very similar to those with  $\gamma = 0.5$ .

In Tables 4 and 5, we report the ECP, AL and MSE of the proposed method and DRL under Scenario (C) where the target policy is fixed, with  $\gamma = 0.3$  and  $0.7$ . It can be seen that the proposed CI achieves nominal coverage in all cases. When  $\gamma = 0.3$ , ECP of the DRL method is well below the nominal level in all cases. When  $\gamma = 0.7$ , the AL and MSE of the proposed CI are much smaller than those based on DRL.



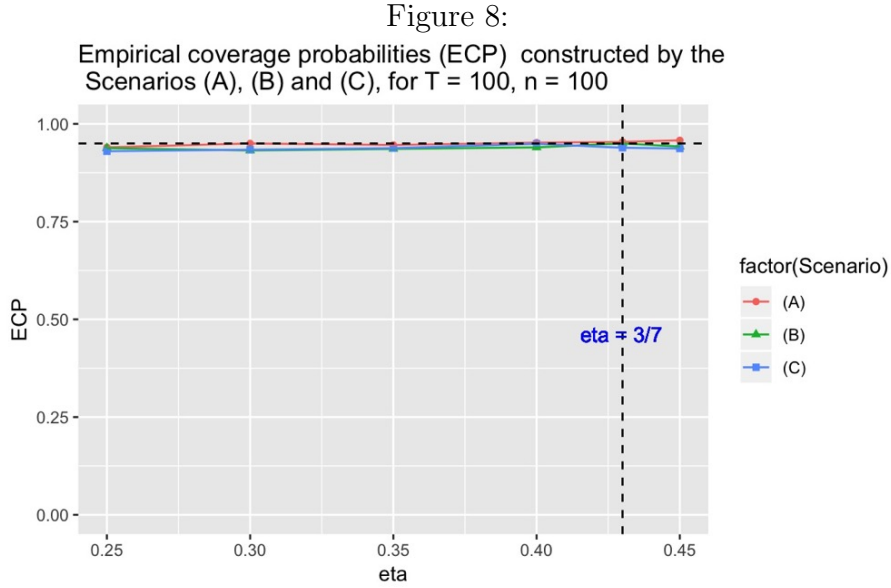



Figure 9: Empirical coverage probabilities and average lengths of CIs constructed by the proposed method as well as the mean squared errors of the value estimates, under Scenario (B), with different choices of  $n$  and  $T$ . The target policy is fixed.  $\gamma = 0.3$  and  $0.7$ , from top to bottom panels.



### D.2.3 Sensitivity test for the ordering of trajectories

We focus on Scenario (B), detailed in Section 5.1, to examine the sensitivity of the proposed CI to the ordering of trajectories. Specifically, we first randomly permute all trajectories

Figure 10: Empirical coverage probabilities and average lengths of CIs constructed by the proposed method as well as the mean squared errors of the value estimates, under Scenario (B), with different choices of  $n$  and  $T$ . The target policy is an estimated optimal policy.  $\gamma = 0.3$  and  $0.7$ , from top to bottom panels.

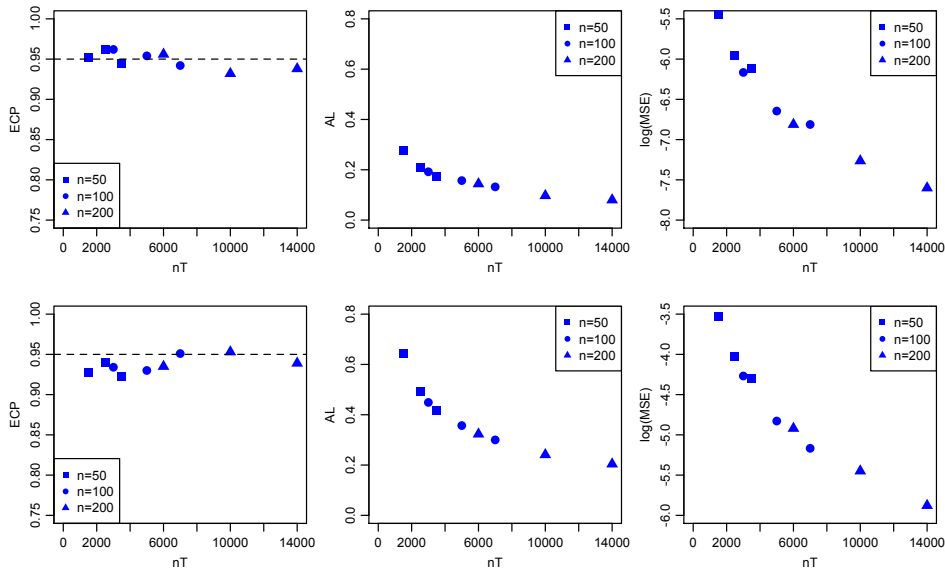


Table 4: Empirical coverage probabilities (ECP) and average lengths (AL) of CIs constructed by the proposed method and DRL as well as the mean-squared errors (MSE) of the corresponding value estimators under Scenario (C), with  $\gamma = 0.3$ .

n	SAVE			DRL		
	ECP	AL	log(MSE)	ECP	AL	log(MSE)
500	0.93	0.09	-7.13	0.76	0.08	-7.01
1000	0.95	0.07	-8.22	0.76	0.06	-7.60
1500	0.94	0.05	-8.52	0.79	0.05	-7.82

with some fixed random seed. We next apply our SAVE procedure to construct the CI. We use three random seeds to generate different random permutations and depict the corresponding results in Figure 11. It can be seen that our method is not overly sensitive to the ordering of trajectories.

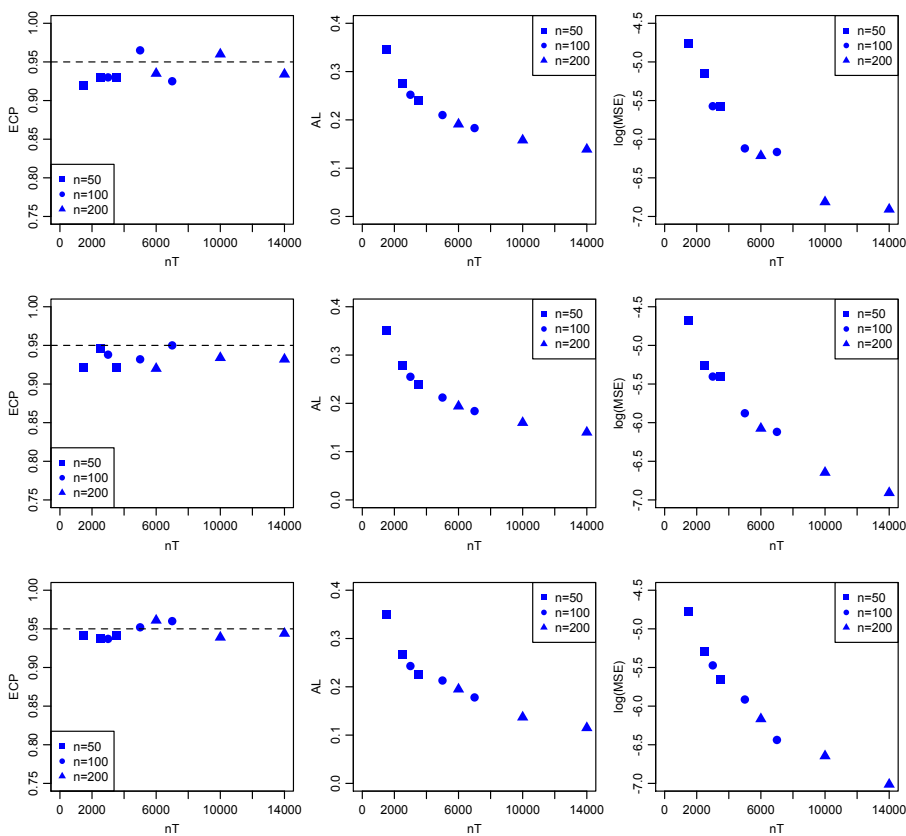
### D.3 Additional settings

In this section, we conduct additional simulation studies to investigate the finite sample performance of the proposed method under settings where the reference distribution  $\mathbb{G}$  is a Dirac delta function. Specifically, we consider the settings in Scenario (A) and set  $\mathbb{G}$  to

Table 5: Empirical coverage probabilities (ECP) and average lengths (AL) of CIs constructed by the proposed method and DRL as well as the mean-squared errors (MSE) of the corresponding value estimators under Scenario (C), with  $\gamma = 0.7$ .

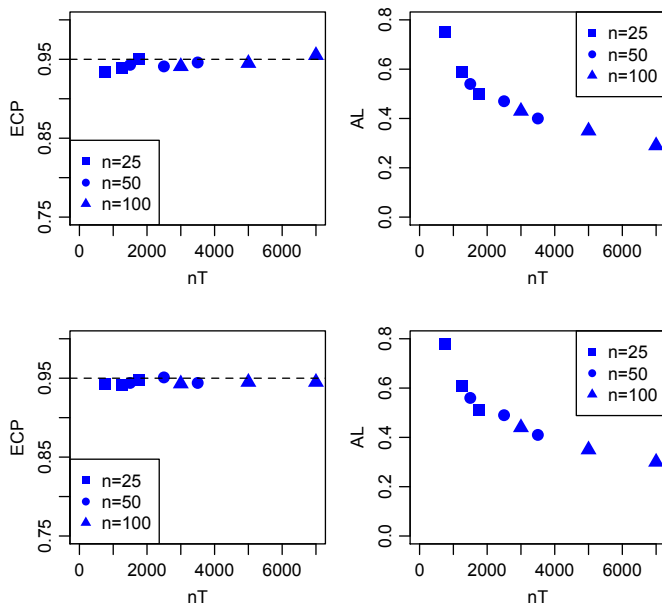
n	SAVE			DRL		
	ECP	AL	log(MSE)	ECP	AL	log(MSE)
500	0.94	0.11	-8.11	0.99	0.34	-5.36
1000	0.96	0.11	-7.60	0.99	0.24	-5.99
1500	0.94	0.08	-8.11	0.99	0.19	-6.81

Figure 11: Empirical coverage probabilities and average lengths of CIs constructed by the proposed method under Scenarios (B), with different choices of  $n$  and  $T$ . From top plots to bottom plots, we use three different random seeds to generate the random permutation applied to all trajectories.



$(0.5, 0.5)^\top$  and  $(-0.5, -0.5)^\top$ . It can be seen from Figures 12 and 13 that our CIs achieve nominal coverage and their lengths decrease as  $nT$  increases.

Figure 12: Empirical coverage probabilities and average lengths of CIs constructed by the proposed method under Scenarios (A), with different choices of  $n$  and  $T$ . The target policy is fixed. The reference distribution  $\mathbb{G}$  corresponds to  $(0.5, 0.5)^\top$  and  $(-0.5, -0.5)^\top$ , from top plots to bottom plots.



## D.4 Additional real data results

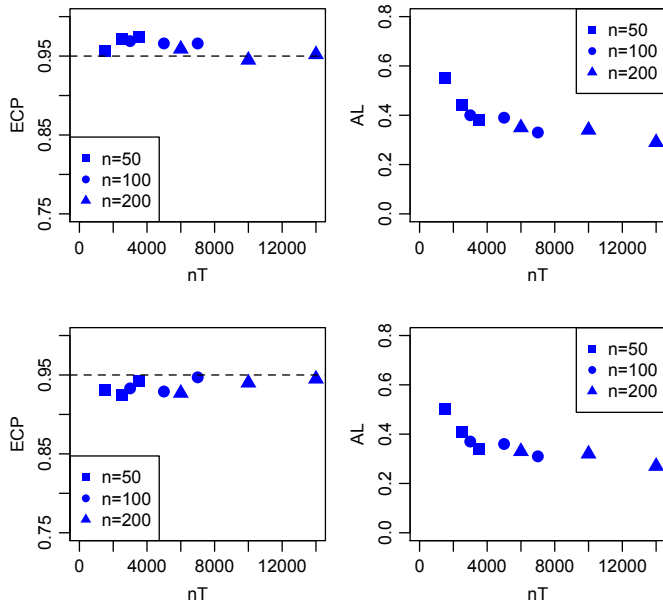
We use our real data example to discuss the issue of over-fitting in this section. Specifically, we apply the proposed method in Section 3.2.2 to evaluate the optimal value  $V(\pi^{opt}; X_{i,0})$  starting from the initial state variable  $X_{i,0}$ , for  $i = 1, 2, \dots, 6$ . When the initial starting time is 8:00 am in Day 1, CIs for Patient 5 and Patient 6 are  $[-6.288, 3.287]$  and  $[-6.313, 10.644]$ , respectively. Both upper bounds are positive. However, according to our definition, the immediate reward is nonpositive. As such, the value and Q-function shall be nonpositive as well. This reflects one of the drawback of the proposed method. The resulting Q-estimator might suffer from over-fitting, leading to an unbounded outcome.

Specifically, it is due to that the matrix  $\widehat{\Sigma}_\pi$  is close to singular. Note that the regression coefficients  $\widehat{\beta}_\pi$  are computed by solving the linear equation

$$\widehat{\Sigma}_\pi \widehat{\beta}_\pi = \frac{1}{\sum_i T_i} \sum_{i=1}^n \sum_{t=0}^{T_i-1} \boldsymbol{\xi}_{i,t} Y_{i,t}.$$

In our data example, the number of basis function equals 12. As such,  $\widehat{\Sigma}_\pi$  is a 12 by 12

Figure 13: Empirical coverage probabilities and average lengths of CIs constructed by the proposed method under Scenarios (A), with different choices of  $n$  and  $T$ . The target policy is an estimated optimal policy. The reference distribution  $\mathbb{G}$  corresponds to  $(0.5, 0.5)^\top$  and  $(-0.5, -0.5)^\top$ , from top plots to bottom plots.



matrix. When it is close to singular, the resulting Q-estimator might be unbounded.

To avoid over-fitting, we note that in theory,  $\widehat{\Sigma}_\pi$  is a positive definite matrix under Condition (A3)(i). This motivates us to compute  $\widehat{\beta}_\pi$  by solving

$$(\widehat{\Sigma}_\pi + \lambda I)\widehat{\beta}_\pi = \frac{1}{\sum_i T_i} \sum_{i=1}^n \sum_{t=0}^{T_i-1} \xi_{i,t} Y_{i,t},$$

where  $I$  denotes the identity matrix. As long as  $\lambda$  satisfies  $\lambda = O(N^{-1}T^{-1})$ , the proposed CI remains valid. In our real data example, we set  $\lambda = 1 \times 10^{-9}$ . The resulting CIs for Patient 5 and Patient 6 are  $[-8.034, -5.667]$  and  $[-9.866, -7.104]$ . Both upper bounds are strictly negative.

## E Technical proofs

For any two positive sequences  $\{a_t\}_{t \geq 1}$  and  $\{b_t\}_{t \geq 1}$ , we write  $a_t \preceq b_t$  if there exists some constant  $C > 0$  such that  $a_t \leq Cb_t$  for any  $t$ . The notation  $a_t \preceq 1$  means  $a_t = O(1)$ . We will use  $C, \bar{C} > 0$  to denote some universal constants whose values are allowed to change

from place to place. Let  $q_X(\cdot|x)$  denote the density function of  $\mathcal{P}_X(\cdot|x)$ . Define  $\mathbb{S}^{mL-1}$  as the unit sphere  $\{\mathbf{v} \in \mathbb{R}^{mL} : \|\mathbf{v}\|_2 = 1\}$ . When splines are used to estimate the Q-function, we assume the internal knots are equally spaced.

The rest of the section is organized as follows. We first present the proof sketches for Theorems 1-4. We next present the detailed technical proofs.

## E.1 A sketch for the proof of Theorem 1

We provide an outline for the proof in this section. The detailed proof can be found in Section E.5 of the supplementary article. We break the proof into three steps. In the first step, we show the estimator  $\widehat{\beta}_\pi$  satisfies

$$\widehat{\beta} - \beta^* = \Sigma_\pi^{-1} \left( \frac{1}{nT} \sum_{i=1}^n \sum_{t=0}^{T-1} \boldsymbol{\xi}_{i,t} \varepsilon_{i,t} \right) + O_p(L^{-p/d}) + O_p\{L(nT)^{-1} \log(nT)\}, \quad (\text{E.29})$$

where  $\Sigma_\pi = \mathbb{E}\widehat{\Sigma}_\pi$ . The proof of (E.29) relies on some random matrix inequalities established in Lemma 3 of the supplementary article.

In the second step, we show the linear representation in (3.12) holds. The proof of (3.12) relies on the convergence rate of  $\widehat{\beta}$  established in the first step and some additional random matrix nequalities in Lemma 4.

In the last step, we show the leading term on the RHS of (3.12) is asymptotically normal, based on the martingale central limit theorem. The completes the proof of Theorem 1.

## E.2 A sketch for the proof of Theorem 2

Similar to (3.15), for each  $1 \leq k \leq K$ , we have

$$\begin{aligned} & \frac{\widehat{V}_{\mathcal{I}_{k+1}}(\widehat{\pi}_{\mathcal{I}_k}; \mathbb{G}) - V(\widehat{\pi}_{\mathcal{I}_k}; \mathbb{G})}{(nT/K)^{-1/2} \widehat{\sigma}_{\mathcal{I}_{k+1}}(\widehat{\pi}_{\mathcal{I}_k}; \mathbb{G})} \\ &= \frac{(nT/K)^{-\frac{1}{2}}}{\sigma(\widehat{\pi}_{\mathcal{I}_k}; \mathbb{G})} \sum_{(i,t) \in \mathcal{I}_{k+1}} \left\{ \int_{x \in \mathbb{X}} \mathbf{U}_{\widehat{\pi}_{\mathcal{I}_k}}(x) \mathbb{G}(dx) \right\}^\top \Sigma_{\widehat{\pi}_{\mathcal{I}_k}}^{-1} \boldsymbol{\xi}_{i,t} \varepsilon_{i,t}(\widehat{\pi}_{\mathcal{I}_k}) + R_k^{(1)}, \end{aligned} \quad (\text{E.30})$$

where  $R_k^{(1)}$  denotes the remainder term and

$$\varepsilon_{i,t}(\widehat{\pi}_{\mathcal{I}_k}) = Y_{i,t} + \gamma \sum_{a \in \mathcal{A}} Q(\widehat{\pi}_{\mathcal{I}_k}; X_{i,t+1}, a) \widehat{\pi}_{\mathcal{I}_k}(a|X_{i,t+1}) - Q(\widehat{\pi}_{\mathcal{I}_k}; X_{i,t}, A_{i,t}).$$

Since (3.16) is satisfied, we have  $\mathbb{E}\{\varepsilon_{i,t}(\widehat{\pi}_{\mathcal{I}_k})|O_k\} = 0$ . Conditional on the data in  $O_k$ ,  $\widehat{\pi}_{\mathcal{I}_k}$  is a deterministic rule. The RHS of (E.30) is thus equivalent to a mean-zero martingale.

When (A5) is satisfied, we have

$$\frac{\widehat{V}_{\mathcal{I}_{k+1}}(\widehat{\pi}_{\mathcal{I}_k}; \mathbb{G}) - V(\widehat{\pi}_{\mathcal{I}_k}; \mathbb{G})}{(nT/K)^{-1/2}\widehat{\sigma}_{\mathcal{I}_{k+1}}(\widehat{\pi}_{\mathcal{I}_k}; \mathbb{G})} = \frac{\widehat{V}_{\mathcal{I}_{k+1}}(\widehat{\pi}_{\mathcal{I}_k}; \mathbb{G}) - V(\pi^*; \mathbb{G})}{(nT/K)^{-1/2}\widehat{\sigma}_{\mathcal{I}_{k+1}}(\widehat{\pi}_{\mathcal{I}_k}; \mathbb{G})} + R_k^{(2)}, \quad (\text{E.31})$$

for some remainder term  $R_k^{(2)}$ . Suppose  $R_k^{(1)}$  and  $R_k^{(2)}$  satisfy certain convergence rates. Combining (E.31) together with (E.30) yields,

$$\begin{aligned} & \frac{\widetilde{V}(\mathbb{G}) - V(\pi^*; \mathbb{G})}{\{nT(K-1)/K\}^{-1/2}\widetilde{\sigma}(\mathbb{G})} \\ &= \sqrt{\frac{nT(K-1)}{K}} \sum_{k=1}^{K-1} \sum_{(i,t) \in \mathcal{I}_{k+1}} \frac{1}{\sigma(\widehat{\pi}_{\mathcal{I}_k}; \mathbb{G})} \left\{ \int_{x \in \mathbb{X}} \mathbf{U}_{\widehat{\pi}_{\mathcal{I}_k}}(x) \mathbb{G}(dx) \right\}^\top \boldsymbol{\Sigma}_{\widehat{\pi}_{\mathcal{I}_k}}^{-1} \boldsymbol{\xi}_{i,t \varepsilon_{i,t}}(\widehat{\pi}_{\mathcal{I}_k}) + o_p(1), \end{aligned} \quad (\text{E.32})$$

due to our use of the inverse weighting trick. Theorem 2 thus follows from the martingale central limit theorem.

### E.3 A sketch for the proofs of Theorems 3 and 4

Proofs of Theorems 3 and 4 are divided into two steps. In the first step, we decompose the value difference  $V(\pi^{opt}; \mathbb{G}) - V(\widehat{\pi}_{\mathcal{I}}; \mathbb{G})$  into the sum of an infinite series and provide upper bounds for all the terms in the series. In the second step, we use the margin-type condition A5 to further characterize these upper bounds. We only present the first step in this section.

For  $j = 0, 1, 2, \dots$ , define a time-dependent policy  $\widehat{\pi}_{\mathcal{I}}^{(j)}$  that executes  $\widehat{\pi}_{\mathcal{I}}$  at the first  $j$  time points and then follows  $\pi^{opt}$ . By definition, we have  $\pi^{opt} = \widehat{\pi}_{\mathcal{I}}^{(0)}$  and  $\widehat{\pi}_{\mathcal{I}} = \widehat{\pi}_{\mathcal{I}}^{(\infty)}$ . Notice that

$$V(\pi^{opt}; \mathbb{G}) - V(\widehat{\pi}_{\mathcal{I}}^{(1)}; \mathbb{G}) = \int_x \sum_{a \in \mathcal{A}} Q^{opt}(x, a) \{ \pi^{opt}(a|x) - \widehat{\pi}_{\mathcal{I}}(a|x) \} \mathbb{G}(dx). \quad (\text{E.33})$$

Moreover, for any  $j \geq 1$ ,

$$\begin{aligned} V(\widehat{\pi}_{\mathcal{I}}^{(j)}; \mathbb{G}) - V(\widehat{\pi}_{\mathcal{I}}^{(j+1)}; \mathbb{G}) &= \int_x \sum_{t \geq 0} \gamma^t \{ \mathbb{E}^{\widehat{\pi}_{\mathcal{I}}^{(j)}}(Y_{0,t}|X_{0,0} = x) - \mathbb{E}^{\widehat{\pi}_{\mathcal{I}}^{(j+1)}}(Y_{0,t}|X_{0,0} = x) \} \mathbb{G}(dx) \\ &= \int_x \sum_{t \geq j} \gamma^t \{ \mathbb{E}^{\widehat{\pi}_{\mathcal{I}}^{(j)}}(Y_{0,t}|X_{0,0} = x) - \mathbb{E}^{\widehat{\pi}_{\mathcal{I}}^{(j+1)}}(Y_{0,t}|X_{0,0} = x) \} \mathbb{G}(dx). \end{aligned}$$

Let  $q_X^{(j)}(\cdot|x)$  be the density function of  $X_{0,j}$  conditional on  $X_{0,0} = x$ , following the estimated policy  $\widehat{\pi}_{\mathcal{I}}$  at the first  $j$  time points, we have

$$\begin{aligned}\sum_{t \geq j} \gamma^t \mathbb{E}^{\widehat{\pi}_{\mathcal{I}}^{(j)}}(Y_{0,t}|X_{0,j}) &= \gamma^j \sum_{a \in \mathcal{A}} Q^{opt}(X_{0,j}, a) \pi^{opt}(a|X_{0,j}), \\ \sum_{t \geq j} \gamma^t \mathbb{E}^{\widehat{\pi}_{\mathcal{I}}^{(j+1)}}(Y_{0,t}|X_{0,j}) &= \gamma^j \sum_{a \in \mathcal{A}} Q^{opt}(X_{0,j}, a) \widehat{\pi}_{\mathcal{I}}(a|X_{0,j}).\end{aligned}$$

It follows that

$$V(\widehat{\pi}_{\mathcal{I}}^{(j)}; \mathbb{G}) - V(\widehat{\pi}_{\mathcal{I}}^{(j+1)}; \mathbb{G}) = \gamma^j \int_{x, x' \in \mathbb{X}} \sum_{a \in \mathcal{A}} Q^{opt}(x', a) \{ \pi^{opt}(a|x') - \widehat{\pi}_{\mathcal{I}}(a|x') \} q_X^{(j)}(x'|x) dx' \mathbb{G}(dx).$$

By A1, we have  $\sup_{x, x', a} q(x'|x, a) \leq c$ . Under the Markov assumption,

$$q_X^{(j)}(x'|x) = \int_y \sum_{a \in \mathcal{A}} q(x'|y, a) \widehat{\pi}_{\mathcal{I}}(a|y) q_X^{(j-1)}(y|x) dy \leq c \int_y q_X^{(j-1)}(y|x) dy = c.$$

In addition,  $\sum_{a \in \mathcal{A}} Q^{opt}(x', a) \{ \pi^{opt}(a|x') - \widehat{\pi}_{\mathcal{I}}(a|x') \} \geq 0$  for any  $x'$ , by the definition of  $\pi^{opt}$ . Therefore, we obtain

$$V(\widehat{\pi}_{\mathcal{I}}^{(j)}; \mathbb{G}) - V(\widehat{\pi}_{\mathcal{I}}^{(j+1)}; \mathbb{G}) \leq c \gamma^j \int_x \sum_{a \in \mathcal{A}} Q^{opt}(x, a) \{ \pi^{opt}(a|x) - \widehat{\pi}_{\mathcal{I}}(a|x) \} dx. \quad (\text{E.34})$$

By A1, the reward function  $r(\cdot, \cdot)$  is uniformly bounded. This further implies that  $Q^{opt}(\cdot, \cdot)$  is uniformly bounded. Therefore,  $\sum_{j \geq t} V(\widehat{\pi}_{\mathcal{I}}^{(j)}; \mathbb{G}) - V(\widehat{\pi}_{\mathcal{I}}^{(j+1)}; \mathbb{G}) \rightarrow 0$  as  $t \rightarrow \infty$ . It follows from (E.33) and (E.34) that

$$\begin{aligned}V(\pi^{opt}; \mathbb{G}) - V(\widehat{\pi}_{\mathcal{I}}; \mathbb{G}) &= \sum_{j=0}^{+\infty} \{ V(\widehat{\pi}_{\mathcal{I}}^{(j)}; \mathbb{G}) - V(\widehat{\pi}_{\mathcal{I}}^{(j+1)}; \mathbb{G}) \} \\ &\leq \int_x \sum_{a \in \mathcal{A}} Q^{opt}(x, a) \{ \pi^{opt}(a|x) - \widehat{\pi}_{\mathcal{I}}(a|x) \} \mathbb{G}(dx)\end{aligned} \quad (\text{E.35})$$

$$+ \frac{c\gamma}{1-\gamma} \int_x \sum_{a \in \mathcal{A}} Q^{opt}(x, a) \{ \pi^{opt}(a|x) - \widehat{\pi}_{\mathcal{I}}(a|x) \} dx. \quad (\text{E.36})$$

Let  $\lambda$  denote the Lebesgue measure on  $\mathbb{R}^d$ . In Sections E.11 and E.12, we use A5 to further bound (E.35) and (E.36).



## E.4 Proof of Lemma 1

Since  $\mathbb{X}$  is compact, Condition (A1) implies that  $\sup_{x \in \mathbb{X}, a \in \mathcal{A}} |r(x, a)| \leq R$  for some  $0 < R < +\infty$ . Under CMIA, we have

$$Q(\pi; x, a) = \sum_{t \geq 0} \gamma^t \mathbb{E}^\pi(Y_{0,t}, X_{0,0} = x, A_{0,0} = a) = \sum_{t \geq 0} \gamma^t \mathbb{E}^\pi\{r(X_{0,t}, A_{0,t}) | X_{0,0} = x, A_{0,0} = a\}.$$

As a result, we obtain

$$\sup_{\pi, x, a} |Q(\pi; x, a)| \leq \frac{R}{1 - \gamma}. \quad (\text{E.37})$$

By the Bellman equation, we obtain

$$Q(\pi; x, a) = r(x, a) + \gamma \int_{x'} \sum_{a' \in \mathcal{A}} Q(\pi; x', a') \pi(a' | x') q(x' | x, a) dx'.$$

Since  $r(\cdot, a)$  is  $p$ -smooth for any  $a \in \mathcal{A}$ , it suffices to show

$$T(\pi; x, a) = \int_{x'} \sum_{a' \in \mathcal{A}} Q(\pi; x', a') \pi(a' | x') q(x' | x, a) dx',$$

is  $p$ -smooth for any  $a \in \mathcal{A}$  and any policy  $\pi$ .

For any function  $h(\cdot)$  defined on  $\mathbb{X}$ , let  $\partial_j h(x)$  denote the partial derivative  $\partial h(x) / \partial x_j$ . Without loss of generality, suppose  $p > 1$  such that  $\partial_j p(x' | x, a)$  exists for any  $j$ . In the following, we show  $\partial_j T(\pi; x, a)$  exists for any  $j$ . Let

$$e_j = (\underbrace{0, \dots, 0}_{j-1}, 1, \underbrace{0, \dots, 0}_{d-j})^\top, \quad \forall j = 1, \dots, d.$$

For any  $\delta \in \mathbb{R}$ , consider the limit

$$\text{Re}(j, \delta) = \frac{T(\pi; x + e_j \delta, a) - T(\pi; x, a)}{\delta} - \int_{x'} \sum_{a' \in \mathcal{A}} Q(\pi; x', a') \pi(a' | x') \partial_j q(x' | x, a) dx',$$

as  $j \rightarrow \infty$ . By the mean value theorem, we have

$$\begin{aligned} |\text{Re}(j, \delta)| &\leq \int_{x'} \sum_{a' \in \mathcal{A}} |Q(\pi; x', a')| \pi(a' | x') \left| \frac{q(x' | x + e_j \delta, a) - q(x' | x, a)}{\delta} - \partial_j q(x' | x, a) \right| dx' \\ &\leq \int_{x'} \sum_{a' \in \mathcal{A}} |Q(\pi; x', a')| \pi(a' | x') |\partial_j q(x' | x + e_j \theta_x \delta, a) - \partial_j q(x' | x, a)| dx', \end{aligned}$$

where  $0 \leq \theta_x \leq 1$  for all  $x$ . When  $1 < p \leq 2$ , we have  $\lfloor p \rfloor = 1$ . It follows from Condition A1 that

$$|\partial_j q(x'|x + e_j \theta_x \delta, a) - \partial_j q(x'|x, a)| \leq c \delta^{p-1}.$$

When  $p > 2$ ,  $|\partial_j \partial_j q(x'|x, a)|$  exists and is bounded by  $c$  for any  $x'$ ,  $x$  and  $a$ . It follows from the mean value theorem that

$$|\partial_j q(x'|x + e_j \theta_x \delta, a) - \partial_j q(x'|x, a)| \leq c \delta.$$

In either case, we have that

$$|\operatorname{Re}(j, \delta)| \leq c \max(\delta, \delta^{p-1}) \int_{x'} \sum_{a' \in \mathcal{A}} |Q(\pi; x', a')| \pi(a'|x') dx'.$$

By (E.37) and that  $\mathbb{X}$  is compact, we obtain  $\operatorname{Re}(j, \delta) \rightarrow 0$  as  $\delta \rightarrow 0$ . This implies that  $\partial_j T(\pi; x, a)$  exists for any  $x \in \mathbb{X}$ ,  $a \in \mathcal{A}$  and equals

$$\int_{x'} \sum_{a' \in \mathcal{A}} Q(\pi; x', a') \pi(a'|x') \partial_j q(x'|x, a) dx'.$$

In addition, it follows from A1 and (E.37) that

$$|\partial_j T(\pi; x, a)| \leq \frac{Rc}{1-\gamma} \lambda(\mathbb{X}), \quad \forall j, x, a, \pi,$$

where  $\lambda(\cdot)$  denotes the Lebesgue measure. Using the same arguments, we can show for any  $d$ -tuple  $\alpha = (\alpha_1, \dots, \alpha_d)^\top$  of nonnegative integers that satisfies  $\|\alpha\|_1 \leq \lfloor p \rfloor$ ,

$$D^\alpha T(\pi; \cdot, a) = \int_{x'} \sum_{a' \in \mathcal{A}} Q(\pi; x', a') \pi(a'|x') D^\alpha q(x'|\cdot, a) dx', \quad (\text{E.38})$$

and

$$\sup_{\pi} \sup_{x \in \mathbb{X}, a \in \mathcal{A}} |D^\alpha T(\pi; x, a)| \leq \frac{Rc}{1-\gamma} \lambda(\mathbb{X}). \quad (\text{E.39})$$

Moreover, by A1, (E.37) and (E.38), we have for any  $d$ -tuple  $\alpha$  with  $\|\alpha\|_1 = \lfloor p \rfloor$  that

$$\begin{aligned} & \|D^\alpha T(\pi; x, a) - D^\alpha T(\pi; y, a)\|_2 \\ & \leq \int_{x'} \sum_{a' \in \mathcal{A}} |Q(\pi; x', a')| \pi(a'|x') |D^\alpha q(x'|x, a) - D^\alpha q(x'|y, a)| dx' \leq \frac{Rc}{1-\gamma} \lambda(\mathbb{X}) \|x - y\|_2^{p-\lfloor p \rfloor}. \end{aligned}$$

This together with (E.38) implies that  $T(\pi; \cdot, a) \in \Lambda(p, Rc\lambda(\mathbb{X})(1-\gamma)^{-1})$  for any  $\pi$  and  $a$ .

The proof is thus completed.

## E.5 Proof of Theorem 1

We introduce the following lemmas before proving Theorem 1. In the proof of Theorem 1 and Lemma 2-5, we will omit the subscript  $\pi$  in  $\mathbf{U}_\pi(\cdot)$ ,  $\mathbf{u}_\pi$ ,  $\boldsymbol{\Sigma}_\pi$ ,  $\widehat{\boldsymbol{\Sigma}}_\pi$ ,  $\widehat{\boldsymbol{\beta}}_\pi, \boldsymbol{\beta}_\pi^*$ ,  $\omega_\pi$ , etc, for brevity.

**Lemma 2** *There exists some constant  $c^* \geq 1$  such that*

$$(c^*)^{-1} \leq \lambda_{\min} \left\{ \int_{x \in \mathbb{X}} \Phi_L(x) \Phi_L^\top(x) dx \right\} \leq \lambda_{\max} \left\{ \int_{x \in \mathbb{X}} \Phi_L(x) \Phi_L^\top(x) dx \right\} \leq c^*, \quad (\text{E.40})$$

and  $\sup_{x \in \mathbb{X}} \|\Phi_L(x)\|_2 \leq c^* \sqrt{L}$ .

**Lemma 3** *Suppose the conditions in Theorem 1 hold. We have as either  $n \rightarrow \infty$  or  $T \rightarrow \infty$  that  $\|\boldsymbol{\Sigma}^{-1}\|_2 \leq 3\bar{c}^{-1}$ ,  $\|\boldsymbol{\Sigma}\|_2 = O(1)$ ,  $\|\widehat{\boldsymbol{\Sigma}} - \boldsymbol{\Sigma}\|_2 = O_p\{L^{1/2}(nT)^{-1/2} \log(nT)\}$ ,  $\|\widehat{\boldsymbol{\Sigma}}^{-1} - \boldsymbol{\Sigma}^{-1}\|_2 = O_p\{L^{1/2}(nT)^{-1/2} \log(nT)\}$  and  $\|\widehat{\boldsymbol{\Sigma}}^{-1}\| \leq 6\bar{c}^{-1}$  wpa1.*

**Lemma 4** *Suppose the conditions in Theorem 1 hold. We have as either  $n \rightarrow \infty$  or  $T \rightarrow \infty$  that  $\lambda_{\max}(T^{-1} \sum_{t=0}^{T-1} \mathbb{E} \boldsymbol{\xi}_{0,t} \boldsymbol{\xi}_{0,t}^\top) = O(1)$ ,  $\lambda_{\max}\{(nT)^{-1} \sum_{i=1}^n \sum_{t=0}^{T-1} \boldsymbol{\xi}_{i,t} \boldsymbol{\xi}_{i,t}^\top\} = O_p(1)$ ,  $\lambda_{\min}(T^{-1} \sum_{t=0}^{T-1} \mathbb{E} \boldsymbol{\xi}_{0,t} \boldsymbol{\xi}_{0,t}^\top) \geq \bar{c}/2$  and  $\lambda_{\min}\{(nT)^{-1} \sum_{i=1}^n \sum_{t=0}^{T-1} \boldsymbol{\xi}_{i,t} \boldsymbol{\xi}_{i,t}^\top\} \geq \bar{c}/3$  wpa1.*

**Lemma 5**  $\|\int_x \mathbf{U}(x) \mathbb{G}(dx)\|_2 \geq m^{-1/2} \|\int_x \Phi_L(x) \mathbb{G}(dx)\|_2$ .

*Step 1.* Since  $L^{2p/d} \gg nT\{1 + \|\int_x \Phi_L(x) \mathbb{G}(dx)\|_2^{-2}\}$ , it follows from Lemma 5 that  $L^{2p/d} \gg nT\{1 + \|\int_x \mathbf{U}(x) \mathbb{G}(dx)\|_2^{-2}\}$ . By Lemma 1, there exist a set of vector  $\{\beta_a^*\}$  that satisfy (see Section 2.2 of Huang, 1998, for details)

$$\sup_{x \in \mathbb{X}, a \in \mathcal{A}} |Q(\pi; x, a) - \Phi_L^\top(x) \beta_a^*| \leq CL^{-p/d}, \quad (\text{E.41})$$

for some constant  $C > 0$ . Let  $\boldsymbol{\beta}^* = (\beta_1^{*T}, \dots, \beta_m^{*T})^\top$ , and

$$r_{i,t} = \gamma \sum_{a \in \mathcal{A}} \{\Phi_L^\top(X_{i,t+1}) \beta_a^* - Q(\pi; X_{i,t+1}, a)\} \pi(a|X_{i,t+1}) - \{\Phi_L^\top(X_{i,t}) \beta_{A_{i,t}}^* - Q(\pi; X_{i,t}, A_{i,t})\}.$$

The condition  $\Pr(\max_{0 \leq t \leq T-1} |Y_{i,t}| \leq c_0) = 1$  implies that  $|Y_{i,t}| \leq c_0, \forall i, t$ , almost surely. By Lemma 1 and the definition of  $p$ -smooth functions, we obtain that  $|Q(\pi; x, a)| \leq c'$  for any  $\pi, x, a$ . It follows that

$$\max_{0 \leq t \leq T-1, 0 \leq i \leq n} |\varepsilon_{i,t}| \leq c_0 + (\gamma + 1)c' \leq c_0 + 2c', \quad (\text{E.42})$$

almost surely. In addition, it follows from (E.41) that

$$\max_{0 \leq t \leq T-1, 1 \leq i \leq n} |r_{i,t}| \leq 2 \sup_{x \in \mathbb{X}, a \in \mathcal{A}} |Q(\pi; x, a) - \Phi_L^\top(x) \beta_a^*| \leq 2CL^{-p/d}. \quad (\text{E.43})$$

By definition, we have

$$\begin{aligned} \widehat{\beta} - \beta^* &= \widehat{\Sigma}^{-1} \left[ \frac{1}{nT} \sum_{i=1}^n \sum_{t=0}^{T-1} \boldsymbol{\xi}_{i,t} \{Y_{i,t} - (\boldsymbol{\xi}_{i,t} - \gamma \mathbf{U}_{i,t+1})^\top \beta^*\} \right] \\ &= \widehat{\Sigma}^{-1} \left[ \frac{1}{nT} \sum_{i=1}^n \sum_{t=0}^{T-1} \boldsymbol{\xi}_{i,t} \left\{ Y_{i,t} - \Phi_L^\top(X_{i,t}) \beta_{A_{i,t}}^* + \gamma \sum_{a \in \mathcal{A}} \Phi_L^\top(X_{i,t+1}) \beta_a^* \pi(a|X_{i,t+1}) \right\} \right] \\ &= \widehat{\Sigma}^{-1} \left\{ \frac{1}{nT} \sum_{i=1}^n \sum_{t=0}^{T-1} \boldsymbol{\xi}_{i,t} (\varepsilon_{i,t} - r_{i,t}) \right\} = \underbrace{\Sigma^{-1} \left( \frac{1}{nT} \sum_{i=1}^n \sum_{t=0}^{T-1} \boldsymbol{\xi}_{i,t} \varepsilon_{i,t} \right)}_{\zeta_1} \\ &+ \underbrace{(\widehat{\Sigma}^{-1} - \Sigma^{-1}) \left( \frac{1}{nT} \sum_{i=1}^n \sum_{t=0}^{T-1} \boldsymbol{\xi}_{i,t} \varepsilon_{i,t} \right)}_{\zeta_2} - \underbrace{\widehat{\Sigma}^{-1} \left( \frac{1}{nT} \sum_{i=1}^n \sum_{t=0}^{T-1} \boldsymbol{\xi}_{i,t} r_{i,t} \right)}_{\zeta_3}. \end{aligned}$$

In the following, we show  $\zeta_2 = O_p\{L(nT)^{-1} \log(nT)\}$  and  $\zeta_3 = O_p(L^{-p/d})$  as either  $n \rightarrow \infty$ , or  $T \rightarrow \infty$ .

*Error bound for  $\|\zeta_2\|_2$ :* Let  $\mathcal{F}_{i,t}$  denote the sub-dataset  $\{X_{i,t}, A_{i,t}\} \cup \{(Y_{i,j}, A_{i,j}, X_{i,j})\}_{1 \leq j < t}$ . By the Bellman equation in (3.7), MA and CMIA, we have

$$\mathbb{E}(\varepsilon_{i,t} | \mathcal{F}_{i,t}) = \mathbb{E}(\varepsilon_{i,t} | X_{i,t}, A_{i,t}) = 0.$$

Notice that  $\boldsymbol{\xi}_{i,t}$  is a function of  $X_{i,t}$  and  $A_{i,t}$  only, we have for any  $0 \leq t_1 < t_2 \leq T-1$  that

$$\mathbb{E} \varepsilon_{i,t_1} \varepsilon_{i,t_2} \boldsymbol{\xi}_{i,t_1}^\top \boldsymbol{\xi}_{i,t_2} = \mathbb{E} \{ \varepsilon_{i,t_1} \boldsymbol{\xi}_{i,t_1}^\top \boldsymbol{\xi}_{i,t_2} \mathbb{E}(\varepsilon_{i,t_2} | \mathcal{F}_{i,t_2}) \} = 0.$$

By the independence assumption, we have for any  $0 \leq t_1 < t_2 \leq T-1$  and  $1 \leq i_1 < i_2 \leq n$  that  $\mathbb{E} \varepsilon_{i_1,t_1} \varepsilon_{i_2,t_2} \boldsymbol{\xi}_{i_1,t_1}^\top \boldsymbol{\xi}_{i_2,t_2} = 0$ . It follows that

$$\mathbb{E} \left\| \sum_{i=1}^n \sum_{t=0}^{T-1} \boldsymbol{\xi}_{i,t} \varepsilon_{i,t} \right\|_2^2 = \sum_{i=1}^n \sum_{t=0}^{T-1} \mathbb{E} \varepsilon_{i,t}^2 \boldsymbol{\xi}_{i,t}^\top \boldsymbol{\xi}_{i,t} = n \sum_{t=0}^{T-1} \mathbb{E} \varepsilon_{0,t}^2 \boldsymbol{\xi}_{0,t}^\top \boldsymbol{\xi}_{0,t}.$$

By (E.42) and Lemma 2, we obtain

$$\mathbb{E} \left\| \sum_{i=1}^n \sum_{t=0}^{T-1} \boldsymbol{\xi}_{i,t} \varepsilon_{i,t} \right\|_2^2 \leq (c_0 + 2c')^2 n \sum_{t=0}^{T-1} \mathbb{E} \boldsymbol{\xi}_{0,t}^\top \boldsymbol{\xi}_{0,t} \leq (c_0 + 2c')^2 nT \sup_{x \in \mathbb{X}} \|\Phi_L(x)\|_2^2 \leq nTL.$$

By Markov's inequality, we obtain  $(nT)^{-1} \sum_{i=1}^n \sum_{t=0}^{T-1} \boldsymbol{\xi}_{i,t} \varepsilon_{i,t} = O_p\{\sqrt{L/(nT)}\}$ . Combining this together with Lemma 3 yields that  $\zeta_2 = O_p\{\sqrt{L/nT} \log(nT)\} O_p\{\sqrt{L/(nT)}\} = O_p\{L(nT)^{-1} \log(nT)\}$ .

*Error bound for  $\|\zeta_3\|_2$ :* For any  $\mathbf{a} \in \mathbb{R}^{mL}$ , it follows from (E.43) and Cauchy-Schwarz inequality that

$$\begin{aligned} \left| \mathbf{a}^\top \left( \frac{1}{nT} \sum_{i=1}^n \sum_{t=0}^{T-1} \boldsymbol{\xi}_{i,t} r_{i,t} \right) \right| &\leq \frac{1}{nT} \sum_{i=1}^n \sum_{t=0}^{T-1} |\mathbf{a}^\top \boldsymbol{\xi}_{i,t}| |r_{i,t}| \leq \max_{i,t} |r_{i,t}| \left( \frac{1}{nT} \sum_{i=1}^n \sum_{t=0}^{T-1} |\mathbf{a}^\top \boldsymbol{\xi}_{i,t}| \right) \\ &\leq 2CL^{-p/d} \left( \frac{1}{nT} \sum_{i=1}^n \sum_{t=0}^{T-1} |\mathbf{a}^\top \boldsymbol{\xi}_{i,t}| \right) \leq 2CL^{-p/d} \left( \frac{1}{nT} \sum_{i=1}^n \sum_{t=0}^{T-1} \mathbf{a}^\top \boldsymbol{\xi}_{i,t} \boldsymbol{\xi}_{i,t}^\top \mathbf{a} \right)^{1/2}. \end{aligned}$$

Therefore, we obtain

$$\left\| \frac{1}{nT} \sum_{i=1}^n \sum_{t=0}^{T-1} \boldsymbol{\xi}_{i,t} r_{i,t} \right\|_2 \leq 2CL^{-p/d} \lambda_{\max}^{1/2} \left( \frac{1}{nT} \sum_{i=1}^n \sum_{t=0}^{T-1} \boldsymbol{\xi}_{i,t} \boldsymbol{\xi}_{i,t}^\top \right).$$

By Lemma 4, we obtain

$$\left\| \frac{1}{nT} \sum_{i=1}^n \sum_{t=0}^{T-1} \boldsymbol{\xi}_{i,t} r_{i,t} \right\|_2 = O_p(L^{-p/d}). \quad (\text{E.44})$$

By Lemma 3, we have  $\|\widehat{\boldsymbol{\Sigma}}^{-1}\|_2 = O_p(1)$ . Combining this together with (E.44) yields that  $\zeta_3 = O_p(L^{-p/d})$ .

To summarize, we have shown

$$\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}^* = \zeta_1 + O_p(L^{-p/d}) + O_p\{L(nT)^{-1} \log(nT)\}. \quad (\text{E.45})$$

This completes the first step of the proof.

*Step 2:* Using similar arguments in bounding  $\|\zeta_2\|_2$  in Step 1, we can show that  $\|\zeta_1\|_2 = O_p\{L^{1/2}(nT)^{-1/2}\}$  and thus

$$\|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*\|_2 = O_p(L^{-p/d}) + O_p\{L^{1/2}(nT)^{-1/2}\}, \quad (\text{E.46})$$

under the condition that  $L \ll \sqrt{nT}/\log(nT)$ . Notice that  $\widehat{V}(\pi; \mathbb{G})$  can be presented as  $\{\int_x \mathbf{U}(x) \mathbb{G}(dx)\}^\top \widehat{\boldsymbol{\beta}}$ . As a result,

$$\left| \widehat{V}(\pi; \mathbb{G}) - \left\{ \int_x \mathbf{U}(x) \mathbb{G}(dx) \right\}^\top (\boldsymbol{\beta}^* + \zeta_1) \right| \leq \left\| \int_x \mathbf{U}(x) \mathbb{G}(dx) \right\|_2 \|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}^* - \zeta_1\|_2. \quad (\text{E.47})$$

By (E.41), we have

$$\begin{aligned} |\mathbf{U}^\top(x)\boldsymbol{\beta}^* - V(\pi; x)| &= \left| \mathbf{U}^\top(x)\boldsymbol{\beta}^* - \sum_{a \in \mathcal{A}} Q(\pi; x, a)\pi(a|x) \right| \\ &\leq \sum_{a \in \mathcal{A}} |Q(\pi; x, a) - \Phi_L^\top(x)\boldsymbol{\beta}_a^*|\pi(a|x) \leq CL^{-p/d}, \end{aligned}$$

and hence

$$\left| \left\{ \int_x \mathbf{U}(x)\mathbb{G}(dx) \right\}^\top \boldsymbol{\beta}^* - \int_x V(\pi; x)\mathbb{G}(dx) \right| \leq CL^{-p/d}.$$

This together with (E.47) yields that

$$|\widehat{V}(\pi; \mathbb{G}) - V(\pi; \mathbb{G}) - \mathbf{U}^\top(x)\zeta_1| \leq \left\| \int_x \mathbf{U}(x)\mathbb{G}(dx) \right\|_2 \|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}^* - \zeta_1\|_2 + CL^{-p/d}. \quad (\text{E.48})$$

Let  $\boldsymbol{\Omega} = T^{-1} \sum_{t=0}^{T-1} \mathbb{E}\omega(X_{0,t}, A_{0,t})\boldsymbol{\xi}_{0,t}\boldsymbol{\xi}_{0,t}^\top$  and

$$\sigma^2(\pi; \mathbb{G}) = \left\{ \int_x \mathbf{U}(x)\mathbb{G}(dx) \right\}^\top \boldsymbol{\Sigma}^{-1}\boldsymbol{\Omega}(\boldsymbol{\Sigma}^\top)^{-1} \left\{ \int_x \mathbf{U}(x)\mathbb{G}(dx) \right\}.$$

Since  $\inf_{x,a} \omega(x, a) \geq c_0^{-1}$ , it follows from Lemma 4 that

$$\lambda_{\min}(\boldsymbol{\Omega}) \geq c_0^{-1} \lambda_{\min} \left( \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}\boldsymbol{\xi}_{0,t}\boldsymbol{\xi}_{0,t}^\top \right) \geq 3^{-1}c_0^{-1}\bar{c}.$$

Hence, we obtain

$$\sigma^2(\pi; \mathbb{G}) \geq \frac{\bar{c}}{3c_0} \left\{ \int_x \mathbf{U}(x)\mathbb{G}(dx) \right\}^\top \boldsymbol{\Sigma}^{-1}(\boldsymbol{\Sigma}^\top)^{-1} \left\{ \int_x \mathbf{U}(x)\mathbb{G}(dx) \right\}. \quad (\text{E.49})$$

By Lemma 3, we have  $\|\boldsymbol{\Sigma}\|_2 = O(1)$ , or equivalently,  $\lambda_{\max}(\boldsymbol{\Sigma}^\top\boldsymbol{\Sigma}) = O(1)$ . This implies that  $\lambda_{\min}\{\boldsymbol{\Sigma}^{-1}(\boldsymbol{\Sigma}^\top)^{-1}\} \geq \bar{C}$  for some constant  $\bar{C} > 0$  and hence

$$\sigma^2(\pi; \mathbb{G}) \geq (3c_0)^{-1}\bar{c}\bar{C} \left\| \int_x \mathbf{U}(x)\mathbb{G}(dx) \right\|_2^2, \quad (\text{E.50})$$

by (E.49). Combining (E.50) together with (E.48) yields that

$$\begin{aligned} \frac{1}{\sigma(\pi; \mathbb{G})} \left| \widehat{V}(\pi; \mathbb{G}) - V(\pi; \mathbb{G}) - \left\{ \int_x \mathbf{U}(x)\mathbb{G}(dx) \right\}^\top \zeta_1 \right| &\leq \frac{\sqrt{3c_0}}{\sqrt{\bar{c}\bar{C}}} \|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}^* - \zeta_1\|_2 \\ &\quad + \frac{\sqrt{3c_0}CL^{-p/d}}{\sqrt{\bar{c}\bar{C}} \left\| \int_x \mathbf{U}(x)\mathbb{G}(dx) \right\|_2}. \end{aligned}$$

By (E.45) and that  $L \ll \sqrt{nT}/\log(nT)$ ,  $L^{2p/d} \gg nT\{1 + \|\int_x \mathbf{U}(x)\mathbb{G}(dx)\|_2^{-2}\}$ , we obtain

$$\frac{\sqrt{nT}\{\widehat{V}(\pi; \mathbb{G}) - V(\pi; \mathbb{G})\}}{\sigma(\pi; \mathbb{G})} = \frac{\sqrt{nT}\{\int_x \mathbf{U}(x)\mathbb{G}(dx)\}^\top \zeta_1}{\sigma(\pi; x)} + o_p(1). \quad (\text{E.51})$$

This completes the second step of the proof.

*Step 3:* In the following, we show  $\sqrt{nT}\sigma^{-1}(\pi; \mathbb{G})\{\int_x \mathbf{U}(x)\mathbb{G}(dx)\}^\top \zeta_1 \xrightarrow{d} N(0, 1)$ . For any integer  $1 \leq g \leq nT$ , let  $i(g)$  and  $t(g)$  be the quotient and the remainder of  $g + T - 1$  divided by  $T$  that satisfy

$$g = \{i(g) - 1\}T + t(g) + 1 \quad \text{and} \quad 0 \leq t(g) < T.$$

Let  $\mathcal{F}^{(0)} = \{X_{1,0}, A_{1,0}\}$ . Then we iteratively define  $\{\mathcal{F}^{(g)}\}_{1 \leq g \leq nT}$  as follows:

$$\begin{aligned} \mathcal{F}^{(g)} &= \mathcal{F}^{(g-1)} \cup \{Y_{i(g),t(g)}, X_{i(g),t(g)+1}, A_{i(g),t(g)+1}\}, \quad \text{if } t(g) < T - 1, \\ \mathcal{F}^{(g)} &= \mathcal{F}^{(g-1)} \cup \{Y_{i(g),T-1}, X_{i(g),T}, X_{i(g)+1,0}, A_{i(g)+1,0}\}, \quad \text{otherwise.} \end{aligned}$$

Let  $\boldsymbol{\xi}^{(g)} = \boldsymbol{\xi}_{i(g),t(g)}$  and  $\varepsilon^{(g)} = \varepsilon_{i(g),t(g)}$ . It follows that

$$\sqrt{nT} \frac{\{\int_x \mathbf{U}(x)\mathbb{G}(dx)\}^\top \zeta_1}{\sigma(\pi; x)} = \sum_{g=1}^{nT} \frac{\{\int_x \mathbf{U}(x)\mathbb{G}(dx)\}^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{\xi}^{(g)} \varepsilon^{(g)}}{\sqrt{nT}\sigma(\pi; x)}. \quad (\text{E.52})$$

By MA, CMIA and the Bellman equation in (3.7), we obtain that

$$\mathbb{E}\{\varepsilon^{(g)} | \mathcal{F}^{(g-1)}\} = \mathbb{E}\{\varepsilon^{(g)} | X_{i(g),t(g)}, A_{i(g),t(g)}\} = 0.$$

Hence, the RHS of (E.52) forms a martingale with respect to the filtration  $\{\sigma(\mathcal{F}^{(g)})\}_{g \geq 0}$ , where  $\sigma(\mathcal{F}^{(g)})$  stands for the  $\sigma$ -algebra generated by  $\mathcal{F}^{(g)}$ . To show the asymptotic normality, we use a martingale central limit theorem for triangular arrays (Corollary 2.8 of McLeish, 1974). This requires to verify the following two conditions:

- (a)  $\max_{1 \leq g \leq nT} |\{\int_x \mathbf{U}(x)\mathbb{G}(dx)\}^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{\xi}^{(g)} \varepsilon^{(g)}| / \{\sqrt{nT}\sigma(\pi; x)\} \xrightarrow{p} 0$ ;
- (b)  $(nT)^{-1} \sum_{g=1}^{nT} |\{\int_x \mathbf{U}(x)\mathbb{G}(dx)\}^\top \boldsymbol{\Sigma} \boldsymbol{\xi}^{(g)} \varepsilon^{(g)}|^2 / \{\sigma^2(\pi; x)\} \xrightarrow{p} 1$ .

Notice that

$$\begin{aligned} & \left| \frac{\{\int_x \mathbf{U}(x)\mathbb{G}(dx)\}^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{\xi}^{(g)} \varepsilon^{(g)}}{\sqrt{nT}\sigma(\pi; x)} \right| \leq \frac{\|\{\int_x \mathbf{U}(x)\mathbb{G}(dx)\}^\top \boldsymbol{\Sigma}^{-1}\|_2 \|\boldsymbol{\xi}^{(g)}\|_2 |\varepsilon^{(g)}|}{\sqrt{nT}\sigma(\pi; x)} \\ & \leq (c_0 + 2c') \frac{\|\{\int_x \mathbf{U}(x)\mathbb{G}(dx)\}^\top \boldsymbol{\Sigma}^{-1}\|_2 \|\boldsymbol{\xi}^{(g)}\|_2}{\sqrt{nT}\sigma(\pi; x)} \leq (c_0 + 2c')c^* \sqrt{L} \frac{\|\{\int_x \mathbf{U}(x)\mathbb{G}(dx)\}^\top \boldsymbol{\Sigma}^{-1}\|_2}{\sqrt{nT}\sigma(\pi; x)} \\ & \leq \frac{\sqrt{3c_0}(c_0 + 2c')c^*}{\sqrt{\bar{c}}} \frac{\sqrt{L}}{\sqrt{nT}}, \end{aligned}$$

where the first inequality follows from Cauchy-Schwarz inequality, the second inequality is due to (E.42), the third inequality is due to Lemma 2 and the fact that  $\|\boldsymbol{\xi}^{(g)}\|_2 \leq \sup_x \|\Phi_L(x)\|_2$ , and the last inequality follows from (E.49). Since  $L \ll \sqrt{nT}/\log(nT)$ , (a) is proven. To verify (b), notice that

$$\begin{aligned} & \left| \frac{1}{nT} \sum_{g=1}^{nT} \frac{|\{\int_x \mathbf{U}(x) \mathbb{G}(dx)\}^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{\xi}^{(g)} \varepsilon^{(g)}|^2}{\sigma^2(\pi; x)} - 1 \right| = \frac{1}{\sigma^2(\pi; x)} \\ & \times \left| \left\{ \int_x \mathbf{U}(x) \mathbb{G}(dx) \right\}^\top \boldsymbol{\Sigma}^{-1} \left\{ \frac{1}{nT} \sum_{g=1}^{nT} (\varepsilon^{(g)})^2 \boldsymbol{\xi}^{(g)} (\boldsymbol{\xi}^{(g)})^\top - \boldsymbol{\Omega} \right\} (\boldsymbol{\Sigma}^\top)^{-1} \left\{ \int_x \mathbf{U}(x) \mathbb{G}(dx) \right\} \right| \\ & \leq \frac{\|\{\int_x \mathbf{U}(x) \mathbb{G}(dx)\}^\top \boldsymbol{\Sigma}^{-1}\|_2^2}{\sigma^2(\pi; x)} \left\| \frac{1}{nT} \sum_{g=1}^{nT} (\varepsilon^{(g)})^2 \boldsymbol{\xi}^{(g)} (\boldsymbol{\xi}^{(g)})^\top - \boldsymbol{\Omega} \right\|_2. \end{aligned}$$

In view of (E.49), it suffices to show

$$\left\| \frac{1}{nT} \sum_{g=1}^{nT} (\varepsilon^{(g)})^2 \boldsymbol{\xi}^{(g)} (\boldsymbol{\xi}^{(g)})^\top - \boldsymbol{\Omega} \right\|_2 = o_p(1). \quad (\text{E.53})$$

This can be proven using similar arguments in bounding  $\|\widehat{\boldsymbol{\Sigma}} - \boldsymbol{\Sigma}\|_2$  in the proof of Lemma 3. In view of (E.51) and (E.52), we have by Slutsky's theorem that

$$\frac{\sqrt{nT} \{\widehat{V}(\pi; \mathbb{G}) - V(\pi; \mathbb{G})\}}{\sigma(\pi; \mathbb{G})} \xrightarrow{d} N(0, 1).$$

To complete the proof, it remains to show  $\widehat{\sigma}(\pi; \mathbb{G})/\sigma(\pi; \mathbb{G}) \xrightarrow{p} 1$ . Using similar arguments in verifying (b), it suffices to show  $\|\widehat{\boldsymbol{\Sigma}}^{-1} \widehat{\boldsymbol{\Omega}} (\widehat{\boldsymbol{\Sigma}}^\top)^{-1} - \boldsymbol{\Sigma}^{-1} \boldsymbol{\Omega} (\boldsymbol{\Sigma}^\top)^{-1}\|_2 = o_p(1)$ . By (E.42) and Lemma 4, we have

$$\lambda_{\max}(\boldsymbol{\Omega}) \leq (c_0 + 2c')^2 \lambda_{\max} \left( \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \boldsymbol{\xi}_{0,t} \boldsymbol{\xi}_{0,t}^\top \right) = O(1),$$

and hence  $\|\boldsymbol{\Omega}\|_2 = O(1)$ . This together with Lemma 3 and the condition  $L \ll \sqrt{nT}/\log(nT)$  yields that

$$\begin{aligned} \|\widehat{\boldsymbol{\Sigma}}^{-1} \widehat{\boldsymbol{\Omega}} (\widehat{\boldsymbol{\Sigma}}^\top)^{-1} - \boldsymbol{\Sigma}^{-1} \boldsymbol{\Omega} (\boldsymbol{\Sigma}^\top)^{-1}\|_2 & \leq \|\widehat{\boldsymbol{\Sigma}}^{-1} - \boldsymbol{\Sigma}^{-1}\|_2 \|\boldsymbol{\Omega}\|_2 \|(\widehat{\boldsymbol{\Sigma}}^\top)^{-1}\|_2 + \|\boldsymbol{\Sigma}^{-1}\|_2 \|\boldsymbol{\Omega}\|_2 \|\widehat{\boldsymbol{\Sigma}}^{-1} - \boldsymbol{\Sigma}^{-1}\|_2 \\ & = O_p\{L^{1/2}(nT)^{-1/2} \log(nT)\} = o_p(1). \end{aligned}$$

Thus, it remains to show  $\|\widehat{\boldsymbol{\Sigma}}^{-1} \widehat{\boldsymbol{\Omega}} (\widehat{\boldsymbol{\Sigma}}^\top)^{-1} - \widehat{\boldsymbol{\Sigma}}^{-1} \boldsymbol{\Omega} (\widehat{\boldsymbol{\Sigma}}^\top)^{-1}\|_2 = o_p(1)$ , or  $\|\widehat{\boldsymbol{\Omega}} - \boldsymbol{\Omega}\|_2 = o_p(1)$ , by Lemma 3. In view of (E.53), it suffices to show  $\|(nT)^{-1} \sum_{g=1}^{nT} (\varepsilon^{(g)})^2 \boldsymbol{\xi}^{(g)} (\boldsymbol{\xi}^{(g)})^\top - \widehat{\boldsymbol{\Omega}}\|_2 =$



$o_p(1)$ , or equivalently,

$$\sup_{a \in \mathbb{S}^{mL-1}} \left| \frac{1}{nT} \sum_{g=1}^{nT} \mathbf{a}^\top \boldsymbol{\xi}^{(g)} (\boldsymbol{\xi}^{(g)})^\top \mathbf{a} \{(\varepsilon^{(g)})^2 - (\widehat{\varepsilon}^{(g)})^2\} \right| = o_p(1),$$

where

$$\widehat{\varepsilon}^{(g)} = Y_{i(g),t(g)} + \gamma \sum_{a \in \mathcal{A}} \Phi_L^\top(X_{i(g),t(g)+1}) \widehat{\beta}_a \pi(a|X_{i(g),t(g)+1}) - \Phi_L^\top(X_{i(g),t(g)}) \widehat{\beta}_{A_{i(g),t(g)}}.$$

By Lemma 4, we have  $\sup_{a \in \mathbb{S}^{mL-1}} (nT)^{-1} \sum_{g=1}^{nT} \mathbf{a}^\top \boldsymbol{\xi}^{(g)} (\boldsymbol{\xi}^{(g)})^\top \mathbf{a} = O_p(1)$ . Hence, it suffices to show  $\max_{1 \leq g \leq nT} |(\varepsilon^{(g)})^2 - (\widehat{\varepsilon}^{(g)})^2| = o_p(1)$ . Suppose we have shown that  $\max_{1 \leq g \leq nT} |\varepsilon^{(g)} - \widehat{\varepsilon}^{(g)}| = o_p(1)$ . By (E.42),  $\varepsilon^{(g)}$ s are uniformly bounded with probability 1 and thus we have  $\max_{1 \leq g \leq nT} |\varepsilon^{(g)} + \widehat{\varepsilon}^{(g)}| = O_p(1)$ . It follows that

$$\max_{1 \leq g \leq nT} |(\varepsilon^{(g)})^2 - (\widehat{\varepsilon}^{(g)})^2| \leq \max_{1 \leq g \leq nT} |\varepsilon^{(g)} - \widehat{\varepsilon}^{(g)}| \max_{1 \leq g \leq nT} |\varepsilon^{(g)} + \widehat{\varepsilon}^{(g)}| = o_p(1).$$

Therefore, it remains to show  $\max_{1 \leq g \leq nT} |\varepsilon^{(g)} - \widehat{\varepsilon}^{(g)}| = o_p(1)$ , or equivalently,

$$\begin{aligned} & \max_{1 \leq i \leq n, 0 \leq t \leq T-1} \left| \gamma \sum_{a \in \mathcal{A}} Q(\pi; X_{i,t+1}, a) \pi(a|X_{i,t+1}) - Q(\pi; X_{i,t}, A_{i,t}) \right. \\ & \quad \left. - \gamma \sum_{a \in \mathcal{A}} \Phi_L^\top(X_{i,t+1}) \widehat{\beta}_a \pi(a|X_{i,t+1}) - \Phi_L^\top(X_{i,t}) \widehat{\beta}_{A_{i,t}} \right| = o_p(1). \end{aligned} \quad (\text{E.54})$$

The LHS of (E.54) is upper bound by

$$(1 + \gamma) \sup_{x \in \mathbb{X}, a \in \mathcal{A}} |Q(\pi; x, a) - \Phi_L^\top(x) \widehat{\beta}_a|.$$

By (E.41), (E.46) and Lemma 2, we have

$$\begin{aligned} \sup_{x \in \mathbb{X}, a \in \mathcal{A}} |Q(\pi; x, a) - \Phi_L^\top(x) \widehat{\beta}_a| & \leq \sup_{x \in \mathbb{X}, a \in \mathcal{A}} |Q(\pi; x, a) - \Phi_L^\top(x) \beta_a^*| + \sup_{x \in \mathbb{X}} |\Phi_L^\top(x) \widehat{\beta}_a - \Phi_L^\top(x) \beta_a^*| \\ & \leq CL^{-p/d} + \sup_{x \in \mathbb{X}} \|\Phi_L(x)\|_2 \sup_{a \in \mathcal{A}} \|\widehat{\beta}_a - \beta_a^*\|_2 = O(L^{-p/d}) + O(L^{1/2}) \sup_{a \in \mathcal{A}} \|\widehat{\beta}_a - \beta_a^*\|_2 \\ & = O_p(L^{1/2-p/d}) + O_p(Ln^{-1/2}T^{-1/2}). \end{aligned}$$

Under the given conditions, we have  $L^{p/d} \gg \sqrt{nT}$  and  $L \ll \sqrt{nT}/\log(nT)$ . This implies  $O_p(L^{1/2-p/d}) = o_p(1)$ , and  $O_p(Ln^{-1/2}T^{-1/2}) = o_p(1)$ . Therefore, we have

$$\sup_{x \in \mathbb{X}, a \in \mathcal{A}} |Q(\pi; x, a) - \Phi_L^\top(x) \widehat{\beta}_a| = o_p(1).$$

The proof is hence completed.

## E.6 Proof of Lemma 2

For B-spline basis, the assertion in (E.40) follows from the arguments used in the proof of Theorem 3.3 of Burman and Chen (1989). For wavelet basis, the assertion in (E.40) follows from the arguments used in the proof of Theorem 5.1 of Chen and Christensen (2015).

For either B-spline or wavelet sieve and any  $L \geq 1$ ,  $x \in \mathbb{X}$ , the number of nonzero elements in the vector  $\Phi_L(x)$  is bounded by some constant. Moreover, each of the basis function is uniformly bounded by  $O(\sqrt{L})$ . This proves that the second assertion.

## E.7 Proof of Lemma 3

We consider two scenarios: (i)  $T$  grows to infinity; (ii)  $T$  is bounded. The proof is divided into four parts. In the first part, we show  $\mathbf{a}^\top \boldsymbol{\Sigma} \mathbf{a} \geq \bar{c} \|\mathbf{a}\|_2^2 / 2$  for any  $\mathbf{a} \in \mathbb{R}^{mL}$  and  $\|\boldsymbol{\Sigma}^{-1}\|_2 \leq 2/\bar{c}$ , as either  $n \rightarrow \infty$ , or  $T \rightarrow \infty$ . In the second part, we bound  $\|\hat{\boldsymbol{\Sigma}} - \boldsymbol{\Sigma}\|_2$ . In the third part, we bound  $\|\hat{\boldsymbol{\Sigma}}^{-1} - \boldsymbol{\Sigma}^{-1}\|_2$ . Finally, we show  $\|\boldsymbol{\Sigma}\|_2 = O(1)$ .

*Part 1:* It follows from Cauchy-Schwarz inequality that

$$\int_{x \in \mathbb{X}} \sum_{a \in \mathcal{A}} \{\mathbf{a}^\top \boldsymbol{\xi}(x, a)\} [\mathbb{E}\{\mathbf{U}(X_{0,1}) | X_{0,0} = x, A_{0,0} = a\}^\top \mathbf{a}] b(a|x) \bar{\mu}(x) dx \leq \eta_1^{1/2}(\mathbf{a}) \eta_2^{1/2}(\mathbf{a}),$$

where  $\bar{\mu} = T^{-1} \sum_{t=0}^{T-1} \mu_t$ ,  $\mu_t$  is the marginal density of  $X_{0,t}$  and

$$\eta_1(\mathbf{a}) = \int_{x \in \mathbb{X}} \sum_{a \in \mathcal{A}} \{\mathbf{a}^\top \mathbf{u}(x, a)\}^2 b(a|x) \bar{\mu}(x) dx, \quad \eta_2(\mathbf{a}) = \int_{x \in \mathbb{X}} \sum_{a \in \mathcal{A}} \{\mathbf{a}^\top \boldsymbol{\xi}(x, a)\}^2 b(a|x) \bar{\mu}(x) dx.$$

Therefore,

$$\mathbf{a}^\top \boldsymbol{\Sigma} \mathbf{a} \geq \eta_2(\mathbf{a}) - \sqrt{\eta_2(\mathbf{a})} \sqrt{\gamma^2 \eta_1(\mathbf{a})} = \sqrt{\eta_2(\mathbf{a})} \frac{\eta_2(\mathbf{a}) - \gamma^2 \eta_1(\mathbf{a})}{\sqrt{\eta_2(\mathbf{a})} + \gamma \sqrt{\eta_1(\mathbf{a})}}.$$

Under A3(i), we obtain  $\eta_2(\mathbf{a}) - \gamma^2 \eta_1(\mathbf{a}) \geq \bar{c} \|\mathbf{a}\|_2^2$  for  $\mathbf{a} \in \mathbb{R}^{mL}$ . It follows that  $\gamma \sqrt{\eta_1(\mathbf{a})} \leq \sqrt{\eta_2(\mathbf{a})}$  and

$$\mathbf{a}^\top \boldsymbol{\Sigma} \mathbf{a} \geq \frac{\bar{c} \|\mathbf{a}\|_2^2 \sqrt{\eta_2(\mathbf{a})}}{\sqrt{\eta_2(\mathbf{a})} + \gamma \sqrt{\eta_1(\mathbf{a})}} \geq \frac{\bar{c}}{2} \|\mathbf{a}\|_2^2, \quad \forall \mathbf{a} \in \mathbb{R}^{mL}. \quad (\text{E.55})$$

We now show

$$\|\boldsymbol{\Sigma} \mathbf{a}\|_2 \geq \frac{\bar{c}}{2} \|\mathbf{a}\|_2, \quad \forall \mathbf{a} \in \mathbb{R}^{mL}. \quad (\text{E.56})$$

Otherwise, there exists some  $\mathbf{a}_0 \in \mathbb{R}^{mL}$  such that  $\|\boldsymbol{\Sigma}\mathbf{a}_0\|_2 < 2^{-1}\bar{c}\|\mathbf{a}_0\|_2$ . By Cauchy-Schwarz inequality, we obtain  $\mathbf{a}_0^\top \boldsymbol{\Sigma}\mathbf{a}_0 \leq \|\mathbf{a}_0\|_2 \|\boldsymbol{\Sigma}\mathbf{a}_0\|_2 < 2^{-1}\bar{c}\|\mathbf{a}_0\|_2^2$ . However, this violates the assertion in (E.55). (E.56) is thus proven.

According to the singular value decomposition, we have  $\boldsymbol{\Sigma} = \mathbf{V}_1^\top \boldsymbol{\Lambda} \mathbf{V}_2$  for some orthogonal matrices  $\mathbf{V}_1, \mathbf{V}_2$  and some diagonal matrix  $\boldsymbol{\Lambda}$ . By orthogonality, we obtain  $\|\boldsymbol{\Sigma}\mathbf{a}\|_2 = \|\boldsymbol{\Lambda}\mathbf{V}_2\mathbf{a}\|_2$  and  $\|\mathbf{a}\|_2 = \|\mathbf{V}_2\mathbf{a}\|_2$ . In view of (E.56), we have  $\|\boldsymbol{\Lambda}\mathbf{a}\|_2 \geq 2^{-1}\bar{c}\|\mathbf{a}\|_2, \forall \mathbf{a} \in \mathbb{R}^{mL}$ . This implies that the absolute value of each diagonal element in  $\boldsymbol{\Lambda}$  is at least  $\bar{c}$ . Thus, we obtain  $\|\boldsymbol{\Lambda}^{-1}\|_2 \leq 2\bar{c}^{-1}$  and hence  $\|\boldsymbol{\Sigma}^{-1}\|_2 \leq 2\bar{c}^{-1}$ .

*Part 2:* We first consider Scenario (ii). Define the random matrix

$$\mathbf{R}_i = \frac{1}{T} \sum_{t=0}^{T-1} \boldsymbol{\xi}(X_{i,t}, A_{i,t}) \{\boldsymbol{\xi}(X_{i,t}, A_{i,t}) - \gamma \mathbf{U}(X_{i,t+1})\}^\top.$$

By Lemma 2, we have  $\max_{1 \leq i \leq n, 0 \leq t \leq T-1} \|\boldsymbol{\xi}(X_{i,t}, A_{i,t})\|_2 \leq \sup_x \|\Phi_L(x)\|_2 \leq c^* \sqrt{L}$  and  $\max_{1 \leq i \leq n, 1 \leq t \leq T} \|\mathbf{U}(X_{i,t+1})\|_2 \leq \sup_x \|\Phi_L(x)\|_2 \leq c^* \sqrt{L}$ . It follows that

$$\max_{1 \leq i \leq n} \|\mathbf{R}_i - \mathbb{E}\mathbf{R}_i\|_2 \leq \frac{2}{T} \sum_{t=0}^{T-1} c^* \sqrt{L} (c^* \sqrt{L} + \gamma c^* \sqrt{L}) \leq 4L(c^*)^2. \quad (\text{E.57})$$

Let

$$\begin{aligned} \sigma_n^2 &= \max \left\{ \left\| \sum_{i=1}^n \mathbb{E}(\mathbf{R}_i - \mathbb{E}\mathbf{R}_i)(\mathbf{R}_i - \mathbb{E}\mathbf{R}_i)^\top \right\|_2, \left\| \sum_{i=1}^n \mathbb{E}(\mathbf{R}_i - \mathbb{E}\mathbf{R}_i)^\top (\mathbf{R}_i - \mathbb{E}\mathbf{R}_i) \right\|_2 \right\} \\ &= n \max \left\{ \left\| \mathbb{E}(\mathbf{R}_0 - \mathbb{E}\mathbf{R}_0)(\mathbf{R}_0 - \mathbb{E}\mathbf{R}_0)^\top \right\|_2, \left\| \mathbb{E}(\mathbf{R}_0 - \mathbb{E}\mathbf{R}_0)^\top (\mathbf{R}_0 - \mathbb{E}\mathbf{R}_0) \right\|_2 \right\}. \end{aligned}$$

For any  $\mathbf{v} \in \mathbb{R}^{mL}$ , we have

$$\mathbf{v}^\top \mathbb{E}(\mathbf{R}_0 - \mathbb{E}\mathbf{R}_0)(\mathbf{R}_0 - \mathbb{E}\mathbf{R}_0)^\top \mathbf{v} = \mathbf{v}^\top \mathbb{E}\mathbf{R}_0 \mathbf{R}_0^\top \mathbf{v} - \mathbf{v}^\top (\mathbb{E}\mathbf{R}_0)(\mathbb{E}\mathbf{R}_0)^\top \mathbf{v} \leq \mathbf{v}^\top \mathbb{E}\mathbf{R}_0 \mathbf{R}_0^\top \mathbf{v}.$$

Moreover, using similar arguments in proving (E.57), we can show

$$\begin{aligned} \mathbf{v}^\top \mathbb{E}\mathbf{R}_0 \mathbf{R}_0^\top \mathbf{v} &\leq \mathbb{E} \left\| \frac{1}{T} \sum_{t=0}^{T-1} \mathbf{v}^\top \boldsymbol{\xi}(X_{0,t}, A_{0,t}) \{\boldsymbol{\xi}(X_{0,t}, A_{0,t}) - \gamma \mathbf{U}(X_{0,t+1})\}^\top \right\|_2^2 \\ &\leq 4L(c^*)^2 \mathbb{E} \left\{ \frac{1}{T} \sum_{t=0}^{T-1} \mathbf{v}^\top \boldsymbol{\xi}(X_{0,t}, A_{0,t}) \right\}^2. \end{aligned}$$

By Cauchy-Schwarz inequality, we obtain

$$\begin{aligned} \mathbf{v}^\top \mathbb{E} \mathbf{R}_0 \mathbf{R}_0^\top \mathbf{v} &\leq \frac{4L(c^*)^2}{T} \sum_{t=0}^{T-1} \mathbb{E} \{ \mathbf{v}^\top \boldsymbol{\xi}(X_{0,t}, A_{0,t}) \}^2 \\ &\leq \frac{4L(c^*)^2 \|\mathbf{v}\|_2^2}{T} \sum_{t=0}^{T-1} \lambda_{\max} \{ \mathbb{E} \boldsymbol{\xi}(X_{0,t}, A_{0,t}) \boldsymbol{\xi}(X_{0,t}, A_{0,t})^\top \}. \end{aligned}$$

Similarly, we can show

$$\mathbf{v}^\top \mathbb{E} \mathbf{R}_0^\top \mathbf{R}_0 \mathbf{v} \leq \frac{4L(c^*)^2 \|\mathbf{v}\|_2^2}{T} \sum_{t=0}^{T-1} \lambda_{\max} \{ \mathbb{E} \boldsymbol{\xi}(X_{0,t}, A_{0,t}) \boldsymbol{\xi}(X_{0,t}, A_{0,t})^\top \},$$

and hence

$$\sigma_n^2 \leq \frac{4Ln(c^*)^2}{T} \sum_{t=0}^{T-1} \lambda_{\max} \{ \mathbb{E} \boldsymbol{\xi}(X_{0,t}, A_{0,t}) \boldsymbol{\xi}(X_{0,t}, A_{0,t})^\top \}. \quad (\text{E.58})$$

Consider  $\lambda_{\max} \{ \mathbb{E} \boldsymbol{\xi}(X_{0,0}, A_{0,0}) \boldsymbol{\xi}(X_{0,0}, A_{0,0})^\top \}$  first. Notice that  $\mathbb{E} \boldsymbol{\xi}(X_{0,0}, A_{0,0}) \boldsymbol{\xi}(X_{0,0}, A_{0,0})^\top$  is a block diagonal matrix. For any  $\mathbf{v} \in \mathbb{R}^{mL}$ , let  $\mathbf{v} = (a_1^\top, \dots, a_m^\top)^\top$  where all the subvectors  $a_j$ s have the same length. With some calculations, we have

$$\begin{aligned} \mathbf{v}^\top \mathbb{E} \boldsymbol{\xi}(X_{0,0}, A_{0,0}) \boldsymbol{\xi}(X_{0,0}, A_{0,0})^\top \mathbf{v} &= \sum_{j=1}^m \mathbb{E} \{ a_j^\top \Phi_L(X_{0,0}) \}^2 b(j|X_{0,0}) \\ &\leq \lambda_{\max} \{ \mathbb{E} \Phi_L(X_{0,0}) \Phi_L^\top(X_{0,0}) \} \|\mathbf{v}\|_2^2 \leq \lambda_{\max} \left\{ \int_{x \in \mathbb{X}} \Phi_L(x) \Phi_L^\top(x) \nu_0(x) dx \right\} \|\mathbf{v}\|_2^2. \end{aligned}$$

By Condition A2 and Lemma 2, we obtain

$$\lambda_{\max} \left\{ \int_{x \in \mathbb{X}} \Phi_L(x) \Phi_L^\top(x) \nu_0(x) dx \right\} \leq \sup_{x \in \mathbb{X}} \nu_0(x) \lambda_{\max} \left\{ \int_{x \in \mathbb{X}} \Phi_L(x) \Phi_L^\top(x) dx \right\} \leq 1.$$

This yields

$$\lambda_{\max} \{ \mathbb{E} \boldsymbol{\xi}(X_{0,0}, A_{0,0}) \boldsymbol{\xi}(X_{0,0}, A_{0,0})^\top \} \leq 1. \quad (\text{E.59})$$

For any  $t > 0$ , the marginal density function of  $X_{0,t}$  is given by

$$\mu_t(x) = \int_{x_0, \dots, x_{t-1} \in \mathbb{X}} \nu_0(x_0) q_X(x_1|x_0) \cdots q_X(x_{t-1}|x_{t-2}) q_X(x|x_{t-1}) dx_0 dx_1 \cdots dx_{t-1}. \quad (\text{E.60})$$

Thus, we have  $\mu_t(x) \leq \sup_{x', x''} q_X(x'|x'')$  for any  $t \geq 1$  and  $x \in \mathbb{X}$ . Under Condition A1, we can show the density function  $q_X(x|x')$  is uniformly bounded for any  $x$  and  $x'$ . It follows

that  $\mu_t(x)$  is uniformly bounded for any  $t \geq 1$  and  $x$ . Using similar arguments in proving (E.59), we can show

$$\lambda_{\max}\{\mathbb{E}\boldsymbol{\xi}(X_{0,t}, A_{0,t})\boldsymbol{\xi}(X_{0,t}, A_{0,t})^\top\} \preceq 1, \quad \forall 1 \leq t \leq T. \quad (\text{E.61})$$

This together with (E.58) and (E.59) yields

$$\sigma_n^2 \leq CnL, \quad (\text{E.62})$$

for some constant  $C > 0$ . Combining this together with (E.57), an application of the matrix concentration inequality (see Theorem 1.6 in Tropp, 2012) yields that

$$\Pr\left(\left\|\sum_{i=1}^n \mathbf{R}_i - n\mathbb{E}\mathbf{R}_0\right\|_2 \geq \tau\right) \leq 2mL \exp\left(-\frac{\tau^2}{CnL + 8L(c^*)^2\tau/3}\right), \quad \forall \tau > 0.$$

Set  $\tau = 3\sqrt{CnL \log n}$ . Since  $T$  is bounded, under the given conditions,  $n$  will grow to infinity. For sufficiently large  $n$ , we have  $8L(c^*)^2\tau/3 \ll \tau^2$  and hence

$$\Pr\left(\left\|\sum_{i=1}^n \mathbf{R}_i - n\mathbb{E}\mathbf{R}_0\right\|_2 \geq 3\sqrt{CnL \log n}\right) \leq \frac{2mL}{n^4}.$$

Since  $L \ll n$  and  $T$  is bounded, we obtain  $2mL/n^4 \ll 1/(n^2T^2)$ . Thus, we can show that the following event occurs with probability at least  $1 - O(n^{-2}T^{-2})$ ,

$$\left\|\widehat{\boldsymbol{\Sigma}} - \boldsymbol{\Sigma}\right\|_2 = \frac{1}{n} \left\|\sum_{i=1}^n \mathbf{R}_i - n\mathbb{E}\mathbf{R}_0\right\|_2 \preceq \sqrt{(nT)^{-1}L \log(nT)}, \quad (\text{E.63})$$

since  $T$  is bounded.

Now let's consider Scenario (i). Let

$$\mathbf{R}_{i,t} = \boldsymbol{\xi}(X_{i,t}, A_{i,t})\{\boldsymbol{\xi}(X_{i,t}, A_{i,t}) - \gamma\mathbf{U}(X_{i,t+1})\}^\top, \quad \forall i, t.$$

We aim to apply the matrix concentration inequality to the sum of independent random matrix (regardless of whether  $n$  is bounded or not),

$$\frac{1}{n} \sum_{i=1}^n \left\{ \frac{1}{T} \sum_{t=0}^{T-1} (\mathbf{R}_{i,t} - \boldsymbol{\Sigma}) \right\}.$$

We begin by providing an upper error bound for  $\max_{1 \leq i \leq n} \|T^{-1} \sum_{t=0}^{T-1} (\mathbf{R}_{i,t} - \boldsymbol{\Sigma})\|_2$ . Let  $\mathcal{F}_{t-1} = \{(X_{0,j}, A_{0,j})\}_{0 \leq j \leq t}$ , for all  $t \geq 0$ , and  $\sigma(\mathcal{F}_t)$  be the  $\sigma$ -algebra generated by  $\mathcal{F}_t$ . Define

$$\mathbf{R}_{0,t}^* = \boldsymbol{\xi}(X_{0,t}, A_{0,t})\{\boldsymbol{\xi}(X_{0,t}, A_{0,t}) - \gamma\mathbf{u}(X_{0,t}, A_{0,t})\}^\top.$$

The sum  $\sum_{t=0}^{T-1}(\mathbf{R}_{0,t}^* - \mathbf{R}_{0,t})$  forms a mean zero matrix martingale with respect to the filtration  $\{\sigma(\mathcal{F}_t) : t \geq -1\}$ . Similar to (E.57) and (E.62), we can show

$$\begin{aligned} & \max_{0 \leq t \leq T-1} \|\mathbf{R}_{0,t} - \mathbf{R}_{0,t}^*\|_2 \leq 4L(c^*)^2, \\ & \max \left\{ \left\| \sum_{t=0}^{T-1} \mathbb{E}\{(\mathbf{R}_{0,t} - \mathbf{R}_{0,t}^*)^\top (\mathbf{R}_{0,t} - \mathbf{R}_{0,t}^*) | \mathcal{F}_{t-1}\} \right\|_2, \right. \\ & \left. \left\| \sum_{t=0}^{T-1} \mathbb{E}\{(\mathbf{R}_{0,t} - \mathbf{R}_{0,t}^*)(\mathbf{R}_{0,t} - \mathbf{R}_{0,t}^*)^\top | \mathcal{F}_{t-1}\} \right\|_2 \right\} \preceq LT. \end{aligned}$$

By the matrix martingale concentration inequality (Corollary 1.3, Tropp, 2011), we obtain the following occurs with probability at least  $1 - O(n^{-3}T^{-2})$ ,

$$\left\| \sum_{t=0}^{T-1} (\mathbf{R}_{0,t} - \mathbf{R}_{0,t}^*) \right\|_2 \preceq \sqrt{LT \log(nT)}. \quad (\text{E.64})$$

Define

$$\mathbf{R}_{0,t}^{**} = \sum_{a \in \mathcal{A}} \boldsymbol{\xi}(X_{0,t}^*, a) \{ \boldsymbol{\xi}(X_{0,t}^*, a) - \gamma \mathbf{u}(X_{0,t}^*, a) \}^\top b(a | X_{0,t}^*). \quad (\text{E.65})$$

Conditional on  $\{X_{0,t}^*\}_{t \geq 0}$ ,  $\{\mathbf{R}_{0,t}^* - \mathbf{R}_{0,t}^{**}\}_{t \geq 0}$  are independent mean zero random variables. Using similar arguments in proving (E.63), we can show that

$$\Pr \left( \left\| \sum_{t=0}^{T-1} (\mathbf{R}_{0,t}^* - \mathbf{R}_{0,t}^{**}) \right\|_2 \geq C \sqrt{LT \log(nT)} \mid \{X_{0,t}^*\}_{t \geq 0} \right) = O(n^{-3}T^{-2}),$$

for some constant  $C > 0$ , where the big- $O$  term is independent of  $\{X_{0,t}^*\}_{t \geq 0}$ . Thus, we obtain

$$\Pr \left( \left\| \sum_{t=0}^{T-1} (\mathbf{R}_{0,t}^* - \mathbf{R}_{0,t}^{**}) \right\|_2 \geq C \sqrt{LT \log(nT)} \right) = O(n^{-3}T^{-2}),$$

This together with (E.64) implies that the following event occurs with probability at least  $1 - O(n^{-3}T^{-2})$ ,

$$\left\| \sum_{t=0}^{T-1} (\mathbf{R}_{0,t} - \mathbf{R}_{0,t}^{**}) \right\|_2 \preceq \sqrt{LT \log(nT)}. \quad (\text{E.66})$$

Notice that each  $\mathbf{R}_{0,t}^{**}$  is a function of  $X_{0,t}$  only, with mean  $\boldsymbol{\Sigma}$ . Following Davydov (1973), define the  $\beta$ -mixing coefficient of the stationary Markov chain  $\{X_{0,t}\}_{t \geq 0}$  as

$$\beta(q) = \int_{x \in \mathbb{X}} \sup_{0 \leq \varphi \leq 1} |\mathbb{E}\{\varphi(X_{0,q}) | X_{0,0} = x\} - \mathbb{E}\varphi(X_{0,0})| \mu(x) dx.$$

Under the geometric ergodicity assumption in A3(ii) and , it follows from Lemma 1 of Meitz and Saikkonen (2019) that  $\{X_{0,t}\}_{t \geq 0}$  is exponentially  $\beta$ -mixing. That is,  $\beta(t) = O(\rho^t)$  for some  $\rho < 1$  and any  $t \geq 0$ . Using similar arguments in proving (E.57), we can show

$$\max_{0 \leq t \leq T-1} \|\mathbf{R}_{0,t}^{**} - \Sigma\|_2 \leq 4L(c^*)^2. \quad (\text{E.67})$$

Moreover, for any  $0 \leq t_1 \leq t_2 \leq T-1$  and any  $\mathbf{v}_1, \mathbf{v}_2 \in \mathbb{R}^{mL}$ , we have by Cauchy-Schwarz inequality that

$$\begin{aligned} & |\mathbf{v}_1^\top \mathbb{E}(\mathbf{R}_{0,t_1}^{**} - \Sigma)(\mathbf{R}_{0,t_2}^{**} - \Sigma)^\top \mathbf{v}_2| \leq \sqrt{\mathbb{E}\|\mathbf{v}_1^\top (\mathbf{R}_{0,t_1}^{**} - \Sigma)\|_2^2} \sqrt{\mathbb{E}\|\mathbf{v}_2^\top (\mathbf{R}_{0,t_2}^{**} - \Sigma)\|_2^2} \\ & \leq \sqrt{\lambda_{\max}\{\mathbb{E}(\mathbf{R}_{0,t_1}^{**} - \Sigma)(\mathbf{R}_{0,t_1}^{**} - \Sigma)^\top\}} \sqrt{\lambda_{\max}\{\mathbb{E}(\mathbf{R}_{0,t_2}^{**} - \Sigma)(\mathbf{R}_{0,t_2}^{**} - \Sigma)^\top\}} \|\mathbf{v}_1\|_2 \|\mathbf{v}_2\|_2. \end{aligned}$$

Using similar arguments in proving (E.62), we can show

$$\max_{0 \leq t \leq T-1} \lambda_{\max}\{\mathbb{E}(\mathbf{R}_{0,t}^{**} - \Sigma)(\mathbf{R}_{0,t}^{**} - \Sigma)^\top\} \preceq L.$$

This implies

$$\max_{0 \leq t_1, t_2 \leq T-1} \sup_{\mathbf{v}_1 \neq 0, \mathbf{v}_2 \neq 0} \frac{|\mathbf{v}_1^\top \mathbb{E}(\mathbf{R}_{0,t_1}^{**} - \Sigma)(\mathbf{R}_{0,t_2}^{**} - \Sigma)^\top \mathbf{v}_2|}{\|\mathbf{v}_1\|_2 \|\mathbf{v}_2\|_2} \preceq L,$$

and hence,

$$\max_{0 \leq t_1, t_2 \leq T-1} \sup_{\mathbf{v} \neq 0} \frac{\|\mathbb{E}(\mathbf{R}_{0,t_1}^{**} - \Sigma)(\mathbf{R}_{0,t_2}^{**} - \Sigma)^\top \mathbf{v}\|_2}{\|\mathbf{v}\|_2} \preceq L,$$

or equivalently,

$$\max_{0 \leq t_1, t_2 \leq T-1} \|\mathbb{E}(\mathbf{R}_{0,t_1}^{**} - \Sigma)(\mathbf{R}_{0,t_2}^{**} - \Sigma)^\top\|_2 \preceq L.$$

Similarly, we can show

$$\max_{0 \leq t_1, t_2 \leq T-1} \|\mathbb{E}(\mathbf{R}_{0,t_1}^{**} - \Sigma)^\top (\mathbf{R}_{0,t_2}^{**} - \Sigma)\|_2 \preceq L. \quad (\text{E.68})$$

Let  $\Sigma_t = \mathbb{E}\mathbf{R}_{0,t}$ . Notice that  $\Sigma = T^{-1} \sum_{t=0}^{T-1} \Sigma_t$  and we have  $\sum_{t=0}^{T-1} (\mathbf{R}_{0,t} - \Sigma) = \sum_{t=0}^{T-1} (\mathbf{R}_{0,t} - \Sigma_t)$ . Similar to Theorem 4.2 of Chen and Christensen (2015), we can show there exist some constant  $C > 0$  such that for any  $\tau \geq 0$  and integer  $1 < q < T$ ,

$$\begin{aligned} \Pr \left( \left\| \sum_{t=0}^{T-1} (\mathbf{R}_{0,t}^{**} - \Sigma_t) \right\|_2 \geq 6\tau \right) & \leq \frac{T}{q} \beta(q) + \Pr \left( \left\| \sum_{t \in \mathcal{I}_r} (\mathbf{R}_{0,t}^{**} - \Sigma_t) \right\|_2 \geq \tau \right) \\ & \quad + 4mL \exp \left( -\frac{\tau^2/2}{CTqL + 4qL\tau(c^*)^2/3} \right), \quad (\text{E.69}) \end{aligned}$$

where  $\mathcal{I}_r = \{q\lfloor(T+1)/q\rfloor, q\lfloor(T+1)/q\rfloor + 1, \dots, T-1\}$ . Suppose  $\tau \geq 5qL(c^*)^2$ . Notice that  $|\mathcal{I}_r| \leq q$ . It follows from (E.67) that

$$\Pr\left(\left\|\sum_{t \in \mathcal{I}_r} (\mathbf{R}_{0,t}^{**} - \Sigma_t)\right\|_2 \geq \tau\right) = 0. \quad (\text{E.70})$$

Since  $\beta(q) = O(\rho^q)$ , set  $q = -3 \log(nT)/\log \rho$ , we obtain  $T\beta(q)/q = O(n^{-3}T^{-2})$ . Set  $\tau = \max\{4\sqrt{CTqL \log(Tn)}, 11qL(c^*)^2 \log(nT)\}$ , we obtain that

$$\frac{\tau^2}{4} \geq 4CTqL \log(Tn) \quad \text{and} \quad \frac{\tau^2}{4} \geq 16qL\tau(c^*)^2/3 \log(Tn) \quad \text{and} \quad \tau \geq 5qL(c^*)^2,$$

as either  $n \rightarrow \infty$  or  $T \rightarrow \infty$ . It follows from (E.69), (E.70) and the condition  $L \ll nT$  that the following event occurs with probability at least  $1 - O(n^{-3}T^{-2})$ ,

$$\left\|\sum_{t=0}^{T-1} (\mathbf{R}_{0,t}^{**} - \Sigma)\right\|_2 \leq \max\{\sqrt{TL} \log(Tn), L \log^2(Tn)\}.$$

Combining this together with (E.66) yields that the following event occurs with probability at least  $1 - O(n^{-3}T^{-2})$ ,

$$\left\|\sum_{t=0}^{T-1} (\mathbf{R}_{0,t} - \Sigma)\right\|_2 \leq \max\{\sqrt{TL} \log(Tn), L \log^2(Tn)\}. \quad (\text{E.71})$$

By Bonferroni's inequality, we obtain with probability at least  $1 - O(n^{-2}T^{-2})$  that

$$\max_{i \in \{1, \dots, n\}} \left\|\sum_{t=0}^{T-1} (\mathbf{R}_{i,t} - \Sigma)\right\|_2 \leq \bar{C} \max\{\sqrt{TL} \log(Tn), L \log^2(Tn)\}, \quad (\text{E.72})$$

for some constant  $\bar{C} > 0$ . For  $i = 0, 1, \dots, n$ , let  $\mathcal{A}_i$  denote the event

$$\mathcal{A}_i = \left\{\left\|\sum_{t=0}^{T-1} (\mathbf{R}_{i,t} - \Sigma)\right\|_2 \leq \bar{C} \max\{\sqrt{TL} \log(Tn), L \log^2(Tn)\}\right\}.$$

It follows from (E.72) that the following event occurs with probability at least  $1 - O(n^{-2}T^{-2})$ ,

$$\sum_{i=1}^n \sum_{t=0}^{T-1} (\mathbf{R}_{i,t} - \Sigma) = \sum_{i=1}^n \left\{\sum_{t=0}^{T-1} (\mathbf{R}_{i,t} - \Sigma)\right\} \mathbb{I}(\mathcal{A}_i). \quad (\text{E.73})$$

Now we provide an upper error bound for

$$\left\|\mathbb{E}\left\{\sum_{t=0}^{T-1} (\mathbf{R}_{0,t} - \Sigma)\right\}^\top \left\{\sum_{t=0}^{T-1} (\mathbf{R}_{0,t} - \Sigma)\right\}\right\|_2 = \sup_{\mathbf{v} \in \mathbb{S}^{mL-1}} \mathbb{E}\left\|\left\{\sum_{t=0}^{T-1} (\mathbf{R}_{0,t} - \Sigma)\right\} \mathbf{v}\right\|_2^2.$$



Note that

$$\begin{aligned}
& \sup_{\mathbf{v} \in \mathbb{S}^{mL-1}} \mathbb{E} \mathbf{v}^\top \left\{ \sum_{t=0}^{T-1} (\mathbf{R}_{0,t} - \boldsymbol{\Sigma}) \right\}^\top \left\{ \sum_{t=0}^{T-1} (\mathbf{R}_{0,t} - \boldsymbol{\Sigma}) \right\} \mathbf{v} \\
& \leq \underbrace{\sum_{t=0}^{T-1} \sup_{\mathbf{v} \in \mathbb{S}^{mL-1}} \mathbb{E} \mathbf{v}^\top (\mathbf{R}_{0,t} - \boldsymbol{\Sigma}_t)^\top (\mathbf{R}_{0,t} - \boldsymbol{\Sigma}_t) \mathbf{v}}_{\eta_1} \\
& \quad + \underbrace{\sum_{\substack{0 \leq t_1, t_2 \leq T-1 \\ 1 \leq |t_1 - t_2| \leq 3}} \sup_{\mathbf{v} \in \mathbb{S}^{mL-1}} \mathbb{E} \mathbf{v}^\top (\mathbf{R}_{0,t_1} - \boldsymbol{\Sigma}_{t_1})^\top (\mathbf{R}_{0,t_2} - \boldsymbol{\Sigma}_{t_2}) \mathbf{v}}_{\eta_2} \\
& \quad + \underbrace{\sum_{\substack{0 \leq t_1, t_2 \leq T-1 \\ |t_1 - t_2| \geq 4}} \sup_{\mathbf{v} \in \mathbb{S}^{mL-1}} \mathbb{E} \mathbf{v}^\top (\mathbf{R}_{0,t_1} - \boldsymbol{\Sigma}_{t_1})^\top (\mathbf{R}_{0,t_2} - \boldsymbol{\Sigma}_{t_2}) \mathbf{v}}_{\eta_3}.
\end{aligned} \tag{E.74}$$

By (E.68), we obtain that

$$\eta_1 \leq LT \quad \text{and} \quad \eta_2 \leq LT. \tag{E.75}$$

For any  $0 \leq t_1 < t_2 \leq T-1$  with  $t_2 - t_1 \geq 4$ , it follows from MA that  $(X_{0,t_2}, A_{0,t_2}, X_{0,t_2+1})$  is independent of  $(X_{0,t_1}, A_{0,t_1}, X_{0,t_1+1})$  given  $X_{0,t_2-1}$ . Thus, we have

$$\mathbb{E}(\mathbf{R}_{0,t_1} - \boldsymbol{\Sigma}_{t_1})^\top (\mathbf{R}_{0,t_2} - \boldsymbol{\Sigma}_{t_2}) = \mathbb{E}(\mathbf{R}_{0,t_1} - \boldsymbol{\Sigma}_{t_1})^\top \mathbb{E}\{(\mathbf{R}_{0,t_2} - \boldsymbol{\Sigma}_{t_2}) | X_{0,t_2-1}\}.$$

Similarly, conditional on  $X_{0,t_1+2}$ ,  $(X_{0,t_1}, A_{0,t_1}, X_{0,t_1+1})$  and  $X_{0,t_2-1}$  are independent. It follows that

$$\begin{aligned}
& \mathbb{E}(\mathbf{R}_{0,t_1} - \boldsymbol{\Sigma}_{t_1})^\top \mathbb{E}\{(\mathbf{R}_{0,t_2} - \boldsymbol{\Sigma}_{t_2}) | X_{0,t_2-1}\} \\
& = \mathbb{E}[\mathbb{E}\{(\mathbf{R}_{0,t_1} - \boldsymbol{\Sigma}_{t_1}) | X_{0,t_1+2}\}]^\top [\mathbb{E}\{(\mathbf{R}_{0,t_2} - \boldsymbol{\Sigma}_{t_2}) | X_{0,t_2-1}\}],
\end{aligned}$$

and hence

$$\mathbb{E}(\mathbf{R}_{0,t_1} - \boldsymbol{\Sigma}_{t_1})^\top (\mathbf{R}_{0,t_2} - \boldsymbol{\Sigma}_{t_2}) = \mathbb{E}[\mathbb{E}\{(\mathbf{R}_{0,t_1} - \boldsymbol{\Sigma}_{t_1}) | X_{0,t_1+2}\}]^\top [\mathbb{E}\{(\mathbf{R}_{0,t_2} - \boldsymbol{\Sigma}_{t_2}) | X_{0,t_2-1}\}]. \tag{E.76}$$

Define  $\boldsymbol{\Theta}_{1,t_1}(x) = \mathbb{E}(\mathbf{R}_{0,t_1} | X_{0,t_1+2} = x)$  and  $\boldsymbol{\Theta}_{2,t_2}(x) = \mathbb{E}(\mathbf{R}_{0,t_2} | X_{0,t_2-1} = x)$ . Let  $\mathbb{E}^{X_{0,t}}$

denote the conditional expectation given  $X_{0,t}$ . With some calculations, we can show

$$\begin{aligned}
\Theta_{1,t_1}(x) &= \mathbb{E} \left[ \mathbb{E}^{X_{0,t_1+1}} \boldsymbol{\xi}(X_{0,t_1}, A_{0,t_1}) \{ \boldsymbol{\xi}(X_{0,t_1}, A_{0,t_1}) - \gamma \mathbf{U}(X_{0,t_1+1}) \}^\top \middle| X_{0,t_1+2} = x \right] \\
&= \mathbb{E} \left[ \int_{y \in \mathbb{X}} \sum_{a \in \mathcal{A}} \boldsymbol{\xi}(y, a) \{ \boldsymbol{\xi}(y, a) - \gamma \mathbf{U}(X_{0,t_1+1}) \}^\top \frac{b(a|y)q(X_{0,t_1+1}|y, a)\mu_{t_1}(y)}{\mu_{t_1+1}(X_{0,t_1+1})} dy \middle| X_{0,t_1+2} = x \right] \\
&= \int_{y_1, y_2 \in \mathbb{X}} \sum_{a \in \mathcal{A}} \boldsymbol{\xi}(y_1, a) \{ \boldsymbol{\xi}(y_1, a) - \gamma \mathbf{U}(y_2) \}^\top b(a|y_1)q(y_2|y_1, a) \frac{\mu_{t_1}(y_1)q_X(x|y_2)}{\mu_{t_1+2}(x)} dy_1 dy_2,
\end{aligned}$$

and

$$\begin{aligned}
\Theta_{2,t_2}(x) &= \mathbb{E} \left[ \sum_{a \in \mathcal{A}} \boldsymbol{\xi}(X_{0,t_2}, a) \{ \boldsymbol{\xi}(X_{0,t_2}, a) - \gamma \mathbf{u}(X_{0,t_2}, a) \}^\top b(a|X_{0,t_2}) \middle| X_{0,t_2-1} = x \right] \\
&= \int_{y \in \mathbb{X}} \sum_{a \in \mathcal{A}} \boldsymbol{\xi}(y, a) \{ \boldsymbol{\xi}(y, a) - \gamma \mathbf{u}(y, a) \}^\top b(a|y)q_X(y|x) dy.
\end{aligned}$$

It follows from (E.76) that

$$\begin{aligned}
\mathbf{v}^\top \mathbb{E}(\mathbf{R}_{0,t_1} - \boldsymbol{\Sigma})^\top (\mathbf{R}_{0,t_2} - \boldsymbol{\Sigma}) \mathbf{v} &= \mathbf{v}^\top \mathbb{E} \boldsymbol{\Theta}_1^\top(X_{0,t_1+2}) \boldsymbol{\Theta}_2(X_{0,t_2-1}) \mathbf{v} - \mathbf{v}^\top (\boldsymbol{\Sigma})^\top \boldsymbol{\Sigma} \mathbf{v} \\
&= \mathbf{v}^\top \mathbb{E} \boldsymbol{\Theta}_{1,t_1}^\top(X_{0,t_1+2}) \boldsymbol{\Theta}_{2,t_2}(X_{0,t_2-1}) \mathbf{v} - \mathbf{v}^\top \{ \mathbb{E} \boldsymbol{\Theta}_{1,t_1}^\top(X_{0,t_1+2}) \}^\top \mathbb{E} \boldsymbol{\Theta}_{2,t_2}(X_{0,t_2-1}) \mathbf{v} \\
&= \sum_{j=1}^{mL} \text{cov} \{ \mathbf{v}^\top \boldsymbol{\Theta}_{1,t_1, \cdot, j}(X_{0,t_1+2}), \boldsymbol{\Theta}_{2,t_2, j, \cdot}(X_{0,t_2-1}) \mathbf{v} \}
\end{aligned}$$

where  $\boldsymbol{\Theta}_{l,t,j,\cdot}(x)$  and  $\boldsymbol{\Theta}_{l,t,\cdot,j}(x)$  denote the  $j$ -th row and  $j$ -column of  $\boldsymbol{\Theta}_{l,t}(x)$ , respectively.

Let  $\xi_j(\cdot, \cdot)$  be the  $j$ -th element of  $\boldsymbol{\xi}(\cdot, \cdot)$ . By Lemma 2 and the definitions of  $\boldsymbol{\xi}$  and  $\mathbf{U}$ ,

$$\sup_{\substack{\mathbf{v} \in \mathbb{S}^{mL-1}, x \in \mathbb{X} \\ a \in \mathcal{A}}} |\boldsymbol{\xi}^\top(x, a) \mathbf{v}| \preceq L^{1/2} \quad \text{and} \quad \sup_{\mathbf{v} \in \mathbb{S}^{mL-1}, x \in \mathbb{X}} |\mathbf{U}^\top(x) \mathbf{v}| \preceq L^{1/2}.$$

It follows that

$$\begin{aligned}
&\max_t \max_{j \in \{1, \dots, mL\}} \sup_{\mathbf{v} \in \mathbb{S}^{mL-1}, x \in \mathbb{X}} |\mathbf{v}^\top \boldsymbol{\Theta}_{1,t, \cdot, j}(x)| \\
&\preceq \int_{y_1, y_2 \in \mathbb{X}} \sum_{a \in \mathcal{A}} |\xi_j(y_1, a)| |\mathbf{v}^\top \{ \boldsymbol{\xi}(y_1, a) - \gamma \mathbf{U}(y_2) \} b(a|y_1)| dy_1 dy_2 \\
&\preceq L^{1/2} \int_{y_1 \in \mathbb{X}} \sum_{a \in \mathcal{A}} |\xi_j(y_1, a)| b(a|y_1) dy_1 \leq L^{1/2} \max_{j \in \{1, \dots, L\}} \int_{y \in \mathbb{X}} |\phi_{L,j}(y)| dy.
\end{aligned}$$

Similarly to Lemma 2, we can show  $\max_{j \in \{1, \dots, L\}} \int_{y \in \mathbb{X}} |\phi_{L,j}(y)| dy \preceq L^{-1/2}$ , and hence

$$\max_{j \in \{1, \dots, mL\}} \sup_{\mathbf{v} \in \mathbb{S}^{mL-1}, x \in \mathbb{X}} |\mathbf{v}^\top \boldsymbol{\Theta}_{1,t, \cdot, j}(x)| \preceq 1. \tag{E.77}$$

It follows that

$$\begin{aligned}
& \sup_{\mathbf{v} \in \mathbb{S}^{mL-1}} |\mathbf{v}^\top \mathbb{E}(\mathbf{R}_{0,t_1} - \boldsymbol{\Sigma}_{t_1})^\top (\mathbf{R}_{0,t_2} - \boldsymbol{\Sigma}_{t_2}) \mathbf{v}| \\
& \leq \sup_{\mathbf{v} \in \mathbb{S}^{mL-1}} \sum_{j=1}^{mL} \int_{x \in \mathbb{X}} |\mathbb{E}\{\boldsymbol{\Theta}_{2,t,j,\cdot}(X_{0,t_2-1}) \mathbf{v} | X_{0,t_1+2}\} - \mathbb{E}\{\boldsymbol{\Theta}_{2,t,j,\cdot}(X_{0,t_2-1}) \mathbf{v} | \mathbf{v}^\top \boldsymbol{\Theta}_{1,\cdot,j}(x)\} \mu(x) dx \\
& \preceq \sup_{\mathbf{v} \in \mathbb{S}^{mL-1}} \sum_{j=1}^{mL} \int_{x \in \mathbb{X}} |\mathbb{E}\{\boldsymbol{\Theta}_{2,t,j,\cdot}(X_{0,t_2-1}) \mathbf{v} | X_{0,t_1+2}\} - \mathbb{E}\{\boldsymbol{\Theta}_{2,t,j,\cdot}(X_{0,t_2-1}) \mathbf{v} | \mu(x)\} dx.
\end{aligned}$$

Similar to (E.77), we can show  $\max_t \max_{j \in \{1, \dots, mL\}} \sup_{\mathbf{v} \in \mathbb{S}^{mL-1}, x \in \mathbb{X}} |\boldsymbol{\Theta}_{2,j,\cdot}(x) \mathbf{v}| \preceq 1$ . It follows from the definition of  $\beta$ -mixing coefficients and the geometric ergodicity that

$$\sup_{\mathbf{v} \in \mathbb{S}^{mL-1}} |\mathbf{v}^\top \mathbb{E}(\mathbf{R}_{0,t_1} - \boldsymbol{\Sigma}_{t_1})^\top (\mathbf{R}_{0,t_2} - \boldsymbol{\Sigma}_{t_2}) \mathbf{v}| \preceq \sum_{j=1}^{mL} \beta(t_2 - t_1 - 3) \preceq L \rho^{t_2 - t_1 - 3},$$

where the above bound is uniform for any pair  $(t_1, t_2)$  that satisfies  $0 \leq t_1 \leq t_2 - 4$ . Therefore, we have

$$\sum_{\substack{0 \leq t_1 < t_2 \leq T-1 \\ |t_2 - t_1| \geq 4}} \sup_{\mathbf{v} \in \mathbb{S}^{mL-1}} |\mathbf{v}^\top \mathbb{E}(\mathbf{R}_{0,t_1} - \boldsymbol{\Sigma}_{t_1})^\top (\mathbf{R}_{0,t_2} - \boldsymbol{\Sigma}_{t_2}) \mathbf{v}| \preceq L \sum_{\substack{0 \leq t_1 < t_2 \leq T-1 \\ |t_2 - t_1| \geq 4}} \rho^{t_2 - t_1 - 3} \preceq LT.$$

Similarly, we can show

$$\sum_{\substack{0 \leq t_1 < t_2 \leq T-1 \\ |t_2 - t_1| \geq 4}} \sup_{\mathbf{v} \in \mathbb{S}^{mL-1}} |\mathbf{v}^\top \mathbb{E}(\mathbf{R}_{0,t_1} - \boldsymbol{\Sigma}_{t_1})^\top (\mathbf{R}_{0,t_2} - \boldsymbol{\Sigma}_{t_2}) \mathbf{v}| \preceq LT,$$

and hence  $\eta_3 \preceq LT$ . This together with (E.74) and (E.75) yields that

$$\sup_{\mathbf{v} \in \mathbb{S}^{mL-1}} \mathbb{E} \mathbf{v}^\top \left\{ \sum_{t=0}^{T-1} (\mathbf{R}_{0,t} - \boldsymbol{\Sigma}) \right\}^\top \left\{ \sum_{t=0}^{T-1} (\mathbf{R}_{0,t} - \boldsymbol{\Sigma}) \right\} \mathbf{v} \preceq LT. \quad (\text{E.78})$$

By Cauchy-Schwarz inequality, we obtain

$$\sup_{\mathbf{v} \in \mathbb{S}^{mL-1}} \mathbb{E} \left\| \left\{ \sum_{t=0}^{T-1} (\mathbf{R}_{0,t} - \boldsymbol{\Sigma}) \right\} \mathbb{I}(\mathcal{A}_0) \mathbf{v} - \mathbb{E} \left\{ \sum_{t=0}^{T-1} (\mathbf{R}_{0,t} - \boldsymbol{\Sigma}) \mathbb{I}(\mathcal{A}_0) \right\} \mathbf{v} \right\|_2^2 \preceq LT.$$

Combining this with (E.72), an application of the matrix Bernstein inequality (Theorem 1.6 in Tropp, 2012) yields that

$$\left\| \sum_{i=1}^n \left\{ \sum_{t=0}^{T-1} (\mathbf{R}_{i,t} - \boldsymbol{\Sigma}) \right\} \mathbb{I}(\mathcal{A}_i) - n \mathbb{E} \left\{ \sum_{t=0}^{T-1} (\mathbf{R}_{0,t} - \boldsymbol{\Sigma}) \right\} \mathbb{I}(\mathcal{A}_0) \right\|_2 = O_p\{\sqrt{nTL} \log(nT)\},$$

under the assumption that  $L = o\{nT/\log^2(nT)\}$ . This together with (E.72) yields that

$$\left\| \sum_{i=1}^n \left\{ \sum_{t=0}^{T-1} (\mathbf{R}_{i,t} - \Sigma) \right\} - n\mathbb{E} \left\{ \sum_{t=0}^{T-1} (\mathbf{R}_{0,t} - \Sigma) \right\} \mathbb{I}(\mathcal{A}_0) \right\|_2 = O_p\{\sqrt{nTL} \log(nT)\}. \quad (\text{E.79})$$

By Cauchy-Schwarz inequality, we have for any  $\mathbf{v}_1, \mathbf{v}_2 \in \mathbb{R}^{mL}$  with  $\|\mathbf{v}_1\|_2 = \|\mathbf{v}_2\|_2 = 1$  that

$$\begin{aligned} \left| \mathbf{v}_1^\top \mathbb{E} \left\{ \sum_{t=0}^{T-1} (\mathbf{R}_{0,t} - \Sigma) \right\} \mathbb{I}(\mathcal{A}_0^c) \mathbf{v}_2 \right| &\leq \sqrt{\mathbb{E} \left[ \mathbf{v}_1^\top \left\{ \sum_{t=0}^{T-1} (\mathbf{R}_{0,t} - \Sigma) \right\} \mathbf{v}_2 \right]^2} \sqrt{\Pr(\mathcal{A}_0^c)} \\ &\leq \sqrt{\mathbb{E} \left\| \left\{ \sum_{t=0}^{T-1} (\mathbf{R}_{0,t} - \Sigma) \right\} \mathbf{v}_2 \right\|_2^2} \sqrt{\Pr(\mathcal{A}_0^c)} \leq \sqrt{LT} n^{-3/2} T^{-1} = O(n^{-1}), \end{aligned}$$

by (E.71), (E.78) and the condition that  $L \ll Tn/\log^2(Tn)$ . This together with (E.79) yields that

$$\left\| \sum_{i=1}^n \left\{ \sum_{t=0}^{T-1} (\mathbf{R}_{i,t} - \mathbb{E}\mathbf{R}_{0,t}) \right\} \right\|_2 = O_p\{\sqrt{nTL} \log(nT)\},$$

and hence  $\|\widehat{\Sigma} - \Sigma\|_2 = O_p\{L^{1/2}(nT)^{-1/2} \log(nT)\}$ .

*Part 3:* In either scenario, we have shown  $\|\widehat{\Sigma} - \Sigma\|_2 = O_p\{L^{1/2}(nT)^{-1/2} \log(nT)\}$ . Under the condition that  $L = o\{nT/\log(nT)\}$ , we have  $\|\widehat{\Sigma} - \Sigma\|_2 = o_p(1)$ . By definition, this implies that  $\|\widehat{\Sigma} - \Sigma\|_2 \leq \bar{c}/6$ , with probability approaching 1. In *Part 1*, we have shown that  $\mathbf{v}^\top \Sigma \mathbf{v} \geq \bar{c}\|\mathbf{v}\|_2/2$ , for any  $\mathbf{v} \in \mathbb{R}^{mL}$ . By Cauchy-Schwarz inequality, the following event occurs with probability approaching 1,

$$\mathbf{v}^\top \widehat{\Sigma} \mathbf{v} \geq \bar{c}\|\mathbf{v}\|_2^2/3, \quad \forall \mathbf{v} \in \mathbb{R}^{mL}.$$

Using similar arguments in *Part 1*, this implies  $\widehat{\Sigma}$  is invertible and satisfies  $\|\widehat{\Sigma}^{-1}\|_2 \leq 3\bar{c}^{-1}$ , with probability tending to 1. Therefore

$$\|\widehat{\Sigma}^{-1} - \Sigma^{-1}\|_2 = \|\widehat{\Sigma}^{-1}(\widehat{\Sigma} - \Sigma)\Sigma^{-1}\|_2 \leq \|\widehat{\Sigma}^{-1}\|_2 \|\widehat{\Sigma} - \Sigma\|_2 \|\Sigma^{-1}\|_2 \leq 6\bar{c}^{-2} \|\widehat{\Sigma} - \Sigma\|_2,$$

with probability tending to 1. Since  $\|\widehat{\Sigma} - \Sigma\|_2 = O_p\{L^{1/2}(nT)^{-1/2} \log(nT)\}$ , we obtain  $\|\widehat{\Sigma}^{-1} - \Sigma^{-1}\|_2 = O_p\{L^{1/2}(nT)^{-1/2} \log(nT)\}$ . The proof is hence completed.

*Part 4:* It suffices to show

$$\eta_4 \equiv \max_t \sup_{\mathbf{v}_1, \mathbf{v}_2 \in \mathbb{S}^{mL-1}} |\mathbf{v}_1^\top (\mathbb{E}\mathbf{R}_{0,t}^{**}) \mathbf{v}_2| \leq 1.$$

With some calculations, we have

$$\begin{aligned} \eta_4 &\leq \max_t \sup_{\mathbf{v}_1, \mathbf{v}_2 \in \mathbb{S}^{mL-1}} \sum_{a \in \mathcal{A}} \mathbb{E} |\mathbf{v}_1^\top \boldsymbol{\xi}(X_{0,t}, a)| |\mathbf{v}_2^\top \boldsymbol{\xi}(X_{0,t}, a)| b(a|X_{0,t}) \\ &+ \max_t \sup_{\mathbf{v}_1, \mathbf{v}_2 \in \mathbb{S}^{mL-1}} \sum_{a \in \mathcal{A}} \mathbb{E} |\mathbf{v}_1^\top \boldsymbol{\xi}(X_{0,t}, a)| |\mathbf{v}_2^\top \mathbf{u}(X_{0,t}, a)| b(a|X_{0,t}) \leq \frac{3}{2} \eta_4^{(1)} + \frac{1}{2} \eta_4^{(2)}, \end{aligned}$$

by Cauchy-Schwarz inequality, where

$$\begin{aligned} \eta_4^{(1)} &= \max_t \sup_{\mathbf{v} \in \mathbb{S}^{mL-1}} \sum_{a \in \mathcal{A}} \mathbb{E} |\mathbf{v}^\top \boldsymbol{\xi}(X_{0,t}, a)|^2 b(a|X_{0,t}), \\ \eta_4^{(2)} &= \max_t \sup_{\mathbf{v} \in \mathbb{S}^{mL-1}} \sum_{a \in \mathcal{A}} \mathbb{E} |\mathbf{v}^\top \mathbf{u}(X_{0,t}, a)|^2 b(a|X_{0,t}). \end{aligned}$$

By definition, we have

$$\eta_4^{(1)} = \max_t \lambda_{\max} \left\{ \sum_{a \in \mathcal{A}} \mathbb{E} \boldsymbol{\xi}(X_{0,t}, a) \boldsymbol{\xi}^\top(X_{0,t}, a) b(a|X_{0,t}) \right\}.$$

Since the matrix  $\sum_{a \in \mathcal{A}} \boldsymbol{\xi}(X_{0,t}, a) \boldsymbol{\xi}^\top(X_{0,t}, a) b(a|X_{0,t})$  is block diagonal with the main-diagonal blocks  $\{\Phi_L(X_{0,t}) \Phi_L(X_{0,t})^\top b(j|X_{0,t})\}_{j=1, \dots, m}$ . By Lemma 2 and Condition A2, we can show  $\eta_4^{(1)} \preceq 1$ . As for  $\eta_4^{(2)}$ , we have

$$\begin{aligned} \eta_4^{(2)} &\leq \max_t \sup_{\mathbf{v} \in \mathbb{S}^{mL-1}} \mathbb{E} |\mathbf{v}^\top \mathbf{U}(X_{0,t+1})|^2 = \lambda_{\max} \{ \mathbb{E} \mathbf{U}(X_{0,t+1}) \mathbf{U}^\top(X_{0,t+1}) \} \\ &\leq \sup_{\substack{v_1, \dots, v_m \in \mathbb{R}^L \\ \|v_1\|_2^2 + \dots + \|v_m\|_2^2 = 1}} \sum_{1 \leq j_1, j_2 \leq m} v_{j_1}^\top \mathbb{E} \Phi_L(X_{0,t+1}) \Phi_L(X_{0,t+1})^\top v_{j_2} \pi(j_1|X_{0,t+1}) \pi(j_2|X_{0,t+1}) \\ &\leq \lambda_{\max} \{ \mathbb{E} \Phi_L(X_{0,t+1}) \Phi_L(X_{0,t+1}^*)^\top \}, \end{aligned}$$

where the first inequality follows from Jensen's inequality. By Lemma 2, we can similarly show that  $\eta_4^{(2)} \preceq 1$ . Thus, we obtain  $\|\boldsymbol{\Sigma}\|_2 = \eta_4 \preceq 1$ . The proof is hence completed.

## E.8 Proof of Lemma 4

The proof of Lemma 4 is very similar to that of Lemma 3. By Condition A3(i), we obtain  $\lambda_{\min}(T^{-1} \sum_{t=0}^{T-1} \mathbb{E} \boldsymbol{\xi}_{0,t} \boldsymbol{\xi}_{0,t}^\top) \geq \bar{c}$ , when  $T$  is bounded. A4(iii) further implies that

$$\lambda_{\min} \left[ \int_{x \in \mathbb{X}} \left\{ \sum_{a \in \mathcal{A}} \boldsymbol{\xi}(x, a) \boldsymbol{\xi}^\top(x, a) b(a|x) \right\} \right] \geq \bar{c},$$

as  $T \rightarrow \infty$ . Using similar arguments as in the first part of the proof of Lemma 3, we can show that

$$\lambda_{\min} \left( \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \boldsymbol{\xi}_{0,t} \boldsymbol{\xi}_{0,t}^\top \right) \geq \frac{\bar{c}}{2}, \quad (\text{E.80})$$

as either  $n \rightarrow \infty$ , or  $T \rightarrow \infty$ . Using similar arguments in the third part of the proof of Lemma 3, we can show that

$$\left\| \frac{1}{nT} \sum_{i=1}^n \sum_{t=0}^{T-1} \boldsymbol{\xi}_{i,t} \boldsymbol{\xi}_{i,t}^\top - \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \boldsymbol{\xi}_{0,t} \boldsymbol{\xi}_{0,t}^\top \right\|_2 = O_p \{ L^{1/2} (nT)^{-1/2} \log(nT) \}. \quad (\text{E.81})$$

Since  $L = o\{\sqrt{nT}/\log(nT)\}$ , it follows from (E.80) and (E.81) that the following event occurs with probability tending to 1,

$$\begin{aligned} \lambda_{\min} \left( \frac{1}{nT} \sum_{i=1}^n \sum_{t=0}^{T-1} \boldsymbol{\xi}_{i,t} \boldsymbol{\xi}_{i,t}^\top \right) &\geq \lambda_{\min} \left( \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \boldsymbol{\xi}_{0,t} \boldsymbol{\xi}_{0,t}^\top \right) \\ &- \left\| \frac{1}{nT} \sum_{i=1}^n \sum_{t=0}^{T-1} \boldsymbol{\xi}_{i,t} \boldsymbol{\xi}_{i,t}^\top - \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \boldsymbol{\xi}_{0,t} \boldsymbol{\xi}_{0,t}^\top \right\|_2 \geq \frac{\bar{c}}{2} - \frac{\bar{c}}{6} = \frac{\bar{c}}{3}. \end{aligned}$$

It remains to show  $\lambda_{\max}\{(nT)^{-1} \sum_{i=1}^n \sum_{t=0}^{T-1} \boldsymbol{\xi}_{i,t} \boldsymbol{\xi}_{i,t}^\top\} = O_p(1)$  and

$$\lambda_{\max} \left( \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \boldsymbol{\xi}_{0,t} \boldsymbol{\xi}_{0,t}^\top \right) = O(1), \quad (\text{E.82})$$

Suppose (E.82) holds. By (E.81) and the condition that  $L = o\{\sqrt{nT}/\log(nT)\}$ , we have  $\lambda_{\max}\{(nT)^{-1} \sum_{i=1}^n \sum_{t=0}^{T-1} \boldsymbol{\xi}_{i,t} \boldsymbol{\xi}_{i,t}^\top\} = O_p(1)$ . Thus, it suffices to show (E.82). This can be proven using similar arguments in Part 2 of the proof of Lemma 3. We omit the details for brevity. The proof is hence completed.

## E.9 Proof of Lemma 5

It follows from Cauchy-Schwarz inequality and triangle inequality that

$$\begin{aligned} &\left\| \int_{x \in \mathbb{X}} \mathbf{U}(x) \mathbb{G}(dx) \right\|_2^2 = \sum_{a=1}^m \left\| \int_{x \in \mathbb{X}} \Phi_L(x) \pi(a|x) \mathbb{G}(dx) \right\|_2^2 \\ &\geq m^{-1} \left\{ \sum_{a=1}^m \left\| \int_{x \in \mathbb{X}} \Phi_L(x) \pi(a|x) \mathbb{G}(dx) \right\|_2 \right\}^2 \geq m^{-1} \left\| \int_{x \in \mathbb{X}} \sum_{a=1}^m \Phi_L(x) \pi(a|x) \mathbb{G}(dx) \right\|_2^2 \\ &= m^{-1} \left\| \int_{x \in \mathbb{X}} \Phi_L(x) \mathbb{G}(dx) \right\|_2^2. \end{aligned}$$

The proof is hence completed.

## E.10 Proof of Theorem 2

Without loss of generality, we assume  $n = K_n n_{\min}$  and  $T = K_T T_{\min}$  such that  $|\mathcal{I}_k| = n_{\min} T_{\min}$  for any  $k$ . Under the given conditions,  $K$  is bounded. Similar to Lemma 3, we can show under A4\* that

$$\max_{2 \leq k \leq K} \|\Sigma_{\hat{\pi}_{\bar{\mathcal{I}}_{k-1}}}^{-1}\|_2 \leq 2\bar{c}^{-1}. \quad (\text{E.83})$$

We next bound the difference between  $\Sigma_{\hat{\pi}_{\bar{\mathcal{I}}_{k-1}}}$  and  $\hat{\Sigma}_{\mathcal{I}_k, \hat{\pi}_{\bar{\mathcal{I}}_{k-1}}}$ . Consider the scenario where  $T$  is bounded first. Since  $T_{\min} = T$ , the data are divided according to the trajectories they belong to. Thus, for any  $k = 2, \dots, K$ , variables  $\{(X_{i,t}, A_{i,t}, Y_{i,t}, X_{i,t+1})\}_{(i,t) \in \mathcal{I}_k}$  are independent of  $\{(X_{i,t}, A_{i,t}, Y_{i,t}, X_{i,t+1})\}_{(i,t) \in \bar{\mathcal{I}}_{k-1}}$ . Let

$$\Sigma_{\hat{\pi}_{\bar{\mathcal{I}}_{k-1}}} = \mathbb{E}[\hat{\Sigma}_{\mathcal{I}_k, \hat{\pi}_{\bar{\mathcal{I}}_{k-1}}} | \{(X_{i,t}, A_{i,t}, Y_{i,t}, X_{i,t+1})\}_{(i,t) \in \bar{\mathcal{I}}_{k-1}}].$$

Using similar arguments in Part 2 of the proof of Lemma 3, we can show

$$\max_{2 \leq k \leq K} \|\Sigma_{\hat{\pi}_{\bar{\mathcal{I}}_{k-1}}} - \hat{\Sigma}_{\mathcal{I}_k, \hat{\pi}_{\bar{\mathcal{I}}_{k-1}}}\|_2 \preceq \sqrt{(nT)^{-1} L \log(nT)}, \quad (\text{E.84})$$

with probability at least  $1 - O(n^{-2}T^{-2}) = 1 - o(1)$ .

Now let us consider the scenario where  $T \rightarrow \infty$ . For  $k = 1, \dots, K$ , define  $(i_0(k), t_0(k))$  to be the tuple in  $\mathcal{I}_k$  such that  $i \geq i_0(k), t \geq t_0(k)$  for any  $(i, t) \in \mathcal{I}_k$ . Then, we have

$$\mathcal{I}_k = \{(i, t) : i_0(k) \leq i < i_0(k) + n_{\min}, t_0(k) \leq t < t_0(k) + T_{\min}\}.$$

Consider any  $k \in \{2, \dots, K\}$  with  $t_0(k) = 0$ ,  $\{(X_{i,t}, A_{i,t}, Y_{i,t}, X_{i,t+1})\}_{(i,t) \in \mathcal{I}_k}$  are independent of  $\{(X_{i,t}, A_{i,t}, Y_{i,t}, X_{i,t+1})\}_{(i,t) \in \bar{\mathcal{I}}_{k-1}}$ . Using similar arguments in Part 2 of the proof of Lemma 3, we can show wpa1 that,

$$\max_{\substack{k \in \{2, \dots, K\} \\ t_0(k) = 0}} \|\hat{\Sigma}_{\mathcal{I}_k, \hat{\pi}_{\bar{\mathcal{I}}_{k-1}}} - \Sigma_{\hat{\pi}_{\bar{\mathcal{I}}_{k-1}}}\|_2 \preceq \sqrt{\frac{L}{nT}} \log(nT). \quad (\text{E.85})$$

Consider  $k \in \{2, \dots, K\}$  with  $t_0(k) > 0$ . We decompose  $\widehat{\Sigma}_{\mathcal{I}_k, \widehat{\pi}_{\bar{\mathcal{I}}_{k-1}}}$  as

$$\underbrace{\frac{1}{n_{\min} T_{\min}} \sum_{i=i_0(k)}^{i_0(k)+n_{\min}-1} \xi_{i,t_0(k)} \{ \xi_{i,t_0(k)} - \gamma \mathbf{U}_{\widehat{\pi}_{\bar{\mathcal{I}}_{k-1}}, i, t_0(k)+1} \}^\top}_{\widehat{\Sigma}_{\mathcal{I}_k, \widehat{\pi}_{\bar{\mathcal{I}}_{k-1}}}^{(1)}} + \underbrace{\frac{1}{n_{\min} T_{\min}} \sum_{\substack{(i,t) \in \mathcal{I}_k \\ t > t_0(k)}} [\xi_{i,t} \{ \xi_{i,t} - \gamma \mathbf{U}_{\widehat{\pi}_{\bar{\mathcal{I}}_{k-1}}, i, t} \}^\top]}_{\widehat{\Sigma}_{\mathcal{I}_k, \widehat{\pi}_{\bar{\mathcal{I}}_{k-1}}}^{(2)}}.$$

It follows that

$$\begin{aligned} \max_{\substack{k \in \{2, \dots, K\} \\ t_0(k) > 0}} \|\widehat{\Sigma}_{\mathcal{I}_k, \widehat{\pi}_{\bar{\mathcal{I}}_{k-1}}} - \Sigma_{\widehat{\pi}_{\bar{\mathcal{I}}_{k-1}}}\|_2 &\leq \max_{\substack{k \in \{2, \dots, K\} \\ t_0(k) > 0}} \|\widehat{\Sigma}_{\mathcal{I}_k, \widehat{\pi}_{\bar{\mathcal{I}}_{k-1}}}^{(1)} - \Sigma_{\widehat{\pi}_{\bar{\mathcal{I}}_{k-1}}}^{(1)}\|_2 \\ &\quad + \max_{\substack{k \in \{2, \dots, K\} \\ t_0(k) > 0}} \|\widehat{\Sigma}_{\mathcal{I}_k, \widehat{\pi}_{\bar{\mathcal{I}}_{k-1}}}^{(2)} - \Sigma_{\widehat{\pi}_{\bar{\mathcal{I}}_{k-1}}}^{(2)}\|_2. \end{aligned}$$

*Error bound for  $\max_k \|\widehat{\Sigma}_{\mathcal{I}_k, \widehat{\pi}_{\bar{\mathcal{I}}_{k-1}}}^{(1)} - \Sigma_{\widehat{\pi}_{\bar{\mathcal{I}}_{k-1}}}^{(1)}\|_2$ :* Given  $\{(X_{j,t}, A_{j,t}, Y_{j,t}, X_{j,t+1})\}_{(j,t) \in \bar{\mathcal{I}}_{k-1}}$ ,  $(A_{i_0(k), t_0(k)}, X_{i_0(k), t_0(k)+1}), \dots, (A_{i_0(k)+n_{\min}-1, t_0(k)}, X_{i_0(k)+n_{\min}-1, t_0(k)+1})$  conditionally independent. Using the matrix concentration inequality, we can show wpa1 that

$$\max_{\substack{k \in \{2, \dots, K\} \\ t_0(k) > 0}} \|\widehat{\Sigma}_{\mathcal{I}_k, \widehat{\pi}_{\bar{\mathcal{I}}_{k-1}}}^{(1)} - \Sigma_{\widehat{\pi}_{\bar{\mathcal{I}}_{k-1}}}^{(1)}\|_2 \preceq \sqrt{\frac{L}{nT}} \log(nT). \quad (\text{E.86})$$

*Error bound for  $\max_k \|\widehat{\Sigma}_{\mathcal{I}_k, \widehat{\pi}_{\bar{\mathcal{I}}_{k-1}}}^{(2)} - \Sigma_{\widehat{\pi}_{\bar{\mathcal{I}}_{k-1}}}^{(2)}\|_2$ :* Given  $\{(X_{j,t}, A_{j,t}, Y_{j,t}, X_{j,t+1})\}_{(j,t) \in \bar{\mathcal{I}}_{k-1}}$ , the sets of variables  $\{(X_{i_0(k), t}, A_{i_0(k), t}, Y_{i_0(k), t}, X_{i_0(k), t+1}) : t_0(k) + 1 \leq t < t_0(k) + T_{\min}\}, \dots, \{(X_{i_0(k)+n_{\min}-1, t}, A_{i_0(k)+n_{\min}-1, t}, Y_{i_0(k)+n_{\min}-1, t}, X_{i_0(k)+n_{\min}-1, t+1}) : t_0(k) + 1 \leq t < t_0(k) + T_{\min}\}$  are conditionally independent. Moreover, for any  $i$  such that  $(i, t_0(k)) \in \mathcal{I}_k$ , the density function of  $X_{i, t_0(k)+1}$  conditional on  $\{(X_{j,t}, A_{j,t}, Y_{j,t}, X_{j,t+1})\}_{(j,t) \in \bar{\mathcal{I}}_{k-1}}$  is uniformly bounded under A3. Using similar arguments in bounding  $\|\widehat{\Sigma} - \Sigma\|_2$  in Part 2 of the proof of Lemma 3, we can show wpa1 that

$$\max_{\substack{k \in \{2, \dots, K\} \\ t_0(k) > 0}} \|\widehat{\Sigma}_{\mathcal{I}_k, \widehat{\pi}_{\bar{\mathcal{I}}_{k-1}}}^{(2)} - \Sigma_{\widehat{\pi}_{\bar{\mathcal{I}}_{k-1}}}^{(2)}\|_2 \preceq \sqrt{\frac{L}{nT}} \log(nT). \quad (\text{E.87})$$

Combining (E.86) with (E.87) and (E.85), we obtain wpa1 that

$$\max_{k \in \{2, \dots, K\}} \|\widehat{\Sigma}_{\mathcal{I}_k, \widehat{\pi}_{\bar{\mathcal{I}}_{k-1}}} - \Sigma_{\widehat{\pi}_{\bar{\mathcal{I}}_{k-1}}}\|_2 \preceq \sqrt{\frac{L}{nT}} \log(nT). \quad (\text{E.88})$$



To summarize, we have shown that

$$\max_{k \in \{2, \dots, K\}} \|\Sigma_{\widehat{\pi}_{\mathcal{I}_{k-1}}}^{-1}\|_2 \leq 2\bar{c}^{-1} \quad \text{and} \quad \max_{k \in \{2, \dots, K\}} \|\widehat{\Sigma}_{\mathcal{I}_k, \widehat{\pi}_{\mathcal{I}_{k-1}}} - \Sigma_{\widehat{\pi}_{\mathcal{I}_{k-1}}}\|_2 \preceq \sqrt{\frac{L}{nT}} \log(nT),$$

wp1, regardless of whether  $T$  is bounded or not. Under the given conditions, we have  $\sqrt{L/(nT)} \log(nT) = o(1)$ . Using similar arguments in the proof of Lemma 3, we can show wp1 that

$$\max_{k \in \{2, \dots, K\}} \|\widehat{\Sigma}_{\mathcal{I}_k, \widehat{\pi}_{\mathcal{I}_{k-1}}}^{-1}\|_2 \leq \frac{3}{\bar{c}}, \quad \max_{k \in \{2, \dots, K\}} \|\widehat{\Sigma}_{\mathcal{I}_k, \widehat{\pi}_{\mathcal{I}_{k-1}}}^{-1} - \Sigma_{\widehat{\pi}_{\mathcal{I}_{k-1}}}^{-1}\|_2 \preceq \sqrt{\frac{L}{nT}} \log(nT). \quad (\text{E.89})$$

Notice that  $\widehat{\pi} = \widehat{\pi}_{\mathcal{I}(K)}$ . By Lemma 1, we have  $Q(\widehat{\pi}_{\mathcal{I}(k)}; \cdot, a) \in \Lambda(p, c')$  for any  $k \in \{1, \dots, K\}$ . Using similar arguments in the proof of Theorem 12.8 of Schumaker (1981) and proof of Proposition 5 of Meyer (1992), there exist some vectors  $\{\beta_{\widehat{\pi}_{\mathcal{I}(k)}, a}^*\}_{a \in \mathcal{A}, 1 \leq k \leq K}$  that satisfy

$$\max_{1 \leq k \leq K} \sup_{x \in \mathbb{X}, a \in \mathcal{A}} |Q(\widehat{\pi}_{\mathcal{I}(k)}; x, a) - \Phi_L^\top(x) \beta_{\widehat{\pi}_{\mathcal{I}(k)}, a}^*| \leq CL^{-p/d}, \quad (\text{E.90})$$

for some constant  $C > 0$ . Similar to (E.43), we have by (E.90) that

$$\max_{(i,t) \in \mathcal{I}_k} |r_{i,t}| \leq 2CL^{-p/d}, \quad \forall k = 2, \dots, K, \quad (\text{E.91})$$

where

$$\begin{aligned} r_{i,t} &= \gamma \sum_{a \in \mathcal{A}} \{\Phi_L^\top(X_{i,t+1}) \beta_{\widehat{\pi}_{\mathcal{I}_{k-1}}, a}^* - Q(\widehat{\pi}_{\mathcal{I}_{k-1}}; X_{i,t+1}, a)\} \widehat{\pi}_{\mathcal{I}_{k-1}}(a | X_{i,t+1}) \\ &\quad - \{\Phi_L^\top(X_{i,t}) \beta_{\widehat{\pi}_{\mathcal{I}_{k-1}}, A_{i,t}}^* - Q(\widehat{\pi}_{\mathcal{I}_{k-1}}; X_{i,t}, A_{i,t})\}. \end{aligned}$$

Similar to the proof of Theorem 1, we have by (E.89) and (E.91) that

$$\widehat{\beta}_{\mathcal{I}_k, \widehat{\pi}_{\mathcal{I}_{k-1}}} - \beta_{\widehat{\pi}_{\mathcal{I}_{k-1}}}^* = \underbrace{\Sigma_{\widehat{\pi}_{\mathcal{I}_{k-1}}}^{-1} \left( \frac{K}{nT} \sum_{(i,t) \in \mathcal{I}_k} \xi_{i,t} \varepsilon_{i,t} \right)}_{\zeta_{1,k}} + O(L^{-p/d}) + O_p\{L(nT)^{-1} \log(nT)\},$$

where

$$\varepsilon_{i,t} = Y_{i,t} + \gamma \sum_{a \in \mathcal{A}} Q(\widehat{\pi}_{\mathcal{I}_{k-1}}; X_{i,t+1}, a) \widehat{\pi}_{\mathcal{I}_{k-1}}(a | X_{i,t+1}) - Q(\widehat{\pi}_{\mathcal{I}_{k-1}}; X_{i,t}, A_{i,t}),$$

for any  $(i, t) \in \mathcal{I}_k$ .

To prove the asymptotic normality of  $\sqrt{nT(K-1)/K}\tilde{\sigma}^{-1}(\mathbb{G})\{\tilde{V}(\mathbb{G}) - V(\hat{\pi}; \mathbb{G})\}$ , it suffices to show

$$\sqrt{\frac{nTK}{K-1}} \sum_{k=2}^K \frac{\widehat{V}_{\mathcal{I}_k}(\hat{\pi}_{\bar{\mathcal{I}}_{k-1}}; \mathbb{G}) - V(\hat{\pi}; \mathbb{G})}{\widehat{\sigma}_{\mathcal{I}_k}(\hat{\pi}_{\bar{\mathcal{I}}_{k-1}}; \mathbb{G})} \xrightarrow{d} N(0, 1). \quad (\text{E.92})$$

Using similar arguments in the proof of Theorem 1, we can show

$$\begin{aligned} & \left| \sqrt{\frac{nTK}{K-1}} \sum_{k=2}^K \frac{\widehat{V}_{\mathcal{I}_k}(\hat{\pi}_{\bar{\mathcal{I}}_{k-1}}; \mathbb{G}) - V(\hat{\pi}; \mathbb{G})}{\widehat{\sigma}_{\mathcal{I}_k}(\hat{\pi}_{\bar{\mathcal{I}}_{k-1}}; \mathbb{G})} \right. \\ & \left. - \sqrt{\frac{nTK}{K-1}} \sum_{k=2}^K \frac{V(\hat{\pi}_{\bar{\mathcal{I}}_{k-1}}; \mathbb{G}) + \{\int_x \mathbf{U}_{\hat{\pi}_{\bar{\mathcal{I}}_{k-1}}}(x) \mathbb{G}(dx)\}^\top \zeta_{1,k} - V(\hat{\pi}; \mathbb{G})}{\widehat{\sigma}_{\mathcal{I}_k}(\hat{\pi}_{\bar{\mathcal{I}}_{k-1}}; \mathbb{G})} \right| = o_p(1). \end{aligned} \quad (\text{E.93})$$

Notice that  $\hat{\pi} = \hat{\pi}_{\bar{\mathcal{I}}_K}$ . Under A5, we have

$$\begin{aligned} & \frac{1}{K-1} \sum_{k=2}^K \mathbb{E} |V(\hat{\pi}_{\bar{\mathcal{I}}_{k-1}}; \mathbb{G}) - V(\hat{\pi}; \mathbb{G})| \\ & \leq \frac{1}{K-1} \sum_{k=2}^K \mathbb{E} |V(\hat{\pi}_{\bar{\mathcal{I}}_{k-1}}; \mathbb{G}) - V(\pi^*; \mathbb{G})| + \mathbb{E} |V(\hat{\pi}; \mathbb{G}) - V(\pi^*; \mathbb{G})| \\ & \leq O(1) \left\{ \frac{1}{K-1} \sum_{k=1}^{K-1} (nT)^{-b_0} k^{-b_0} K^{b_0} + (nT)^{-b_0} \right\}, \end{aligned}$$

where  $O(1)$  denotes some positive constant. Since  $\sum_{k=1}^{K-1} k^{-b_0} \leq 1 + \int_1^K x^{-b_0} dx \leq K^{1-b_0}$ , we obtain that

$$\frac{1}{K-1} \sum_{k=2}^K \mathbb{E} |V(\hat{\pi}_{\bar{\mathcal{I}}_{k-1}}; \mathbb{G}) - V(\hat{\pi}; \mathbb{G})| = O\{(nT)^{-b_0}\},$$

and hence

$$\sqrt{\frac{nTK}{K-1}} \sum_{k=2}^K \mathbb{E} |V(\hat{\pi}_{\bar{\mathcal{I}}_{k-1}}; \mathbb{G}) - V(\hat{\pi}; \mathbb{G})| = O\{(nT)^{1/2-b_0}\} = o\left\{ \left\| \int_x \Phi_L(x) dx \right\|_2 \right\},$$

By A6. By Markov's inequality, we obtain that

$$\left\| \int_x \Phi_L(x) dx \right\|_2^{-1} \frac{\sqrt{nTK}}{K-1} \sum_{k=2}^K |V(\hat{\pi}_{\bar{\mathcal{I}}_{k-1}}; \mathbb{G}) - V(\hat{\pi}; \mathbb{G})| = o_p(1).$$

Similar to Lemma 5, we can show that for any  $k \in \{2, \dots, K\}$ ,

$$\left\| \int_x \Phi_L(x) dx \right\|_2^{-1} \geq m^{-1/2} \left\| \int_x \mathbf{U}_{\hat{\pi}_{\bar{\mathcal{I}}_{k-1}}}(x) dx \right\|_2^{-1}.$$

It follows that

$$\sqrt{\frac{nTK}{K-1}} \sum_{k=2}^K \left\| \int_x \mathbf{U}_{\widehat{\pi}_{\mathcal{I}_{k-1}}}(x) dx \right\|_2^{-1} |V(\widehat{\pi}_{\mathcal{I}_{k-1}}; \mathbb{G}) - V(\widehat{\pi}; \mathbb{G})| = o_p(1). \quad (\text{E.94})$$

Using similar arguments in the proof of Theorem 1, we can show

$$\max_{k \in \{2, \dots, K\}} \left| \frac{\widehat{\sigma}_{\mathcal{I}_k}^2(\widehat{\pi}_{\mathcal{I}_{k-1}}; \mathbb{G})}{\sigma^2(\widehat{\pi}_{\mathcal{I}_{k-1}}; \mathbb{G})} - 1 \right| = o_p(1).$$

Since

$$\left| \frac{\widehat{\sigma}_{\mathcal{I}_k}(\widehat{\pi}_{\mathcal{I}_{k-1}}; \mathbb{G})}{\sigma(\widehat{\pi}_{\mathcal{I}_{k-1}}; \mathbb{G})} - 1 \right| = \left| \frac{\widehat{\sigma}_{\mathcal{I}_k}^2(\widehat{\pi}_{\mathcal{I}_{k-1}}; \mathbb{G})}{\sigma^2(\widehat{\pi}_{\mathcal{I}_{k-1}}; \mathbb{G})} - 1 \right| \left| \frac{\widehat{\sigma}_{\mathcal{I}_k}(\widehat{\pi}_{\mathcal{I}_{k-1}}; \mathbb{G})}{\sigma(\widehat{\pi}_{\mathcal{I}_{k-1}}; \mathbb{G})} + 1 \right|^{-1},$$

we obtain

$$\max_{k \in \{2, \dots, K\}} \left| \frac{\widehat{\sigma}_{\mathcal{I}_k}(\widehat{\pi}_{\mathcal{I}_{k-1}}; \mathbb{G})}{\sigma(\widehat{\pi}_{\mathcal{I}_{k-1}}; \mathbb{G})} - 1 \right| = o_p(1). \quad (\text{E.95})$$

Similar to (E.50), we can show there exists some constant  $c_6 > 0$ , such that the following occurs wpa1,

$$\sigma(\widehat{\pi}_{\mathcal{I}_{k-1}}; \mathbb{G}) \geq c_6 \left\| \int_x \mathbf{U}_{\widehat{\pi}_{\mathcal{I}_{k-1}}}(x) \mathbb{G}(dx) \right\|_2, \quad \forall k \in \{2, \dots, K\}.$$

This together with (E.95) yields that

$$\widehat{\sigma}_{\mathcal{I}_k}(\widehat{\pi}_{\mathcal{I}_{k-1}}; \mathbb{G}) \geq c_6 \left\| \int_x \mathbf{U}_{\widehat{\pi}_{\mathcal{I}_{k-1}}}(x) \mathbb{G}(dx) \right\|_2, \quad \forall k \in \{2, \dots, K\},$$

wpa1. In view of (E.94), we obtain

$$\left| \sqrt{\frac{nTK}{K-1}} \sum_{k=2}^K \frac{V(\widehat{\pi}_{\mathcal{I}_{k-1}}; \mathbb{G}) - V(\widehat{\pi}; \mathbb{G})}{\widehat{\sigma}_{\mathcal{I}_k}(\widehat{\pi}_{\mathcal{I}_{k-1}}; \mathbb{G})} \right| = o_p(1).$$

Combining this together with (E.93), we obtain that

$$\sqrt{\frac{nTK}{K-1}} \sum_{k=2}^K \frac{\widehat{V}_{\mathcal{I}_k}(\widehat{\pi}_{\mathcal{I}_{k-1}}; \mathbb{G}) - V(\widehat{\pi}; \mathbb{G})}{\widehat{\sigma}_{\mathcal{I}_k}(\widehat{\pi}_{\mathcal{I}_{k-1}}; \mathbb{G})} - \sqrt{\frac{nTK}{K-1}} \sum_{k=2}^K \frac{\{\int_x \mathbf{U}_{\widehat{\pi}_{\mathcal{I}_{k-1}}}(x) \mathbb{G}(dx)\}^\top \zeta_{1,k}}{\widehat{\sigma}_{\mathcal{I}_k}(\widehat{\pi}_{\mathcal{I}_{k-1}}; \mathbb{G})} = o_p(1).$$

To prove (E.92), it suffices to show

$$\sqrt{\frac{nTK}{K-1}} \sum_{k=2}^K \frac{\{\int_x \mathbf{U}_{\widehat{\pi}_{\mathcal{I}_{k-1}}}(x) \mathbb{G}(dx)\}^\top \zeta_{1,k}}{\widehat{\sigma}_{\mathcal{I}_k}(\widehat{\pi}_{\mathcal{I}_{k-1}}; \mathbb{G})} \xrightarrow{d} N(0, 1). \quad (\text{E.96})$$

The LHS of (E.96) can be further decomposed as

$$\begin{aligned} & \underbrace{\sqrt{\frac{nTK}{K-1}} \sum_{k=2}^K \frac{\{\int_x \mathbf{U}_{\widehat{\pi}_{\mathcal{I}_{k-1}}}(x) \mathbb{G}(dx)\}^\top \zeta_{1,k}}{\sigma(\widehat{\pi}_{\mathcal{I}_{k-1}}; \mathbb{G})}}_{\eta_7} \\ & + \underbrace{\sqrt{\frac{nTK}{K-1}} \sum_{k=2}^K \frac{\{\int_x \mathbf{U}_{\widehat{\pi}_{\mathcal{I}_{k-1}}}(x) \mathbb{G}(dx)\}^\top \zeta_{1,k}}{\sigma(\widehat{\pi}_{\mathcal{I}_{k-1}}; \mathbb{G})}}_{\eta_8} \left\{ \frac{\sigma(\widehat{\pi}_{\mathcal{I}_{k-1}}; \mathbb{G})}{\widehat{\sigma}_{\mathcal{I}_k}(\widehat{\pi}_{\mathcal{I}_{k-1}}; \mathbb{G})} - 1 \right\}. \end{aligned}$$

In the following, we show  $\eta_7 \xrightarrow{d} N(0, 1)$ . Based on (E.94), one can show  $\eta_8 \xrightarrow{P} 0$ . Assertion (E.96) thus follows from Slutsky's theorem.

Notice that  $\eta_7$  equals

$$\eta_7 = \sqrt{\frac{K}{nT(K-1)}} \sum_{k=2}^K \sum_{(i,t) \in \mathcal{I}_k} \frac{\{\int_x \mathbf{U}_{\widehat{\pi}_{\mathcal{I}_{k-1}}}(x) \mathbb{G}(dx)\}^\top \Sigma_{\widehat{\pi}_{\mathcal{I}_{k-1}}}^{-1} \boldsymbol{\xi}_{i,t} \varepsilon_{i,t}}{\sigma(\widehat{\pi}_{\mathcal{I}_{k-1}}; \mathbb{G})}.$$

For any  $1 \leq g \leq nT$ , there exists some integer  $k(g)$  that satisfies  $\{k(g) - 1\}n_{\min}T_{\min} + 1 \leq g \leq k(g)n_{\min}T_{\min}$ . Let  $t(g)$  and  $i(g)$  be the integers that satisfy

$$g - \{k(g) - 1\}n_{\min}T_{\min} = \{i(g) - i_0(k(g))\}T_{\min} + t(g) - t_0(k(g)) + 1.$$

Let  $\mathcal{F}^{(0)} = \{X_{1,0}, A_{1,0}\}$ . Then we iteratively define  $\{\mathcal{F}^{(g)}\}_{1 \leq g \leq nT}$  as follows:

$$\begin{aligned} \mathcal{F}^{(g)} &= \mathcal{F}^{(g-1)} \cup \{Y_{i(g), T-1}, X_{i(g), T}\}, \quad \text{if } t(g+1) > 0 \text{ \& } t(g) = T-1, \\ \mathcal{F}^{(g)} &= \mathcal{F}^{(g-1)} \cup \{Y_{i(g), t(g)}, X_{i(g), t(g)+1}, A_{i(g), t(g)+1}\}, \quad \text{if } t(g+1) > 0 \text{ \& } t(g) < T-1, \\ \mathcal{F}^{(g)} &= \mathcal{F}^{(g-1)} \cup \{Y_{i(g), T-1}, X_{i(g), T}, X_{i(g+1), 0}, A_{i(g+1), 0}\}, \quad \text{if } t(g+1) = 0 \text{ \& } t(g) = T-1, \\ \mathcal{F}^{(g)} &= \mathcal{F}^{(g-1)} \cup \{Y_{i(g), t(g)}, X_{i(g), t(g)+1}, A_{i(g), t(g)+1}, X_{i(g+1), 0}, A_{i(g+1), 0}\}, \quad \text{otherwise.} \end{aligned}$$

Let  $\boldsymbol{\xi}^{(g)} = \boldsymbol{\xi}_{i(g), t(g)}$  and  $\varepsilon^{(g)} = \varepsilon_{i(g), t(g)}$ . We rewrite  $\eta_7$  as

$$\eta_7 = \sqrt{\frac{K}{nT(K-1)}} \sum_{g=nT/K+1}^{nT} \frac{\{\int_x \mathbf{U}_{\widehat{\pi}_{\mathcal{I}(k(g)-1)}}(x) \mathbb{G}(dx)\}^\top \Sigma_{\widehat{\pi}_{\mathcal{I}(k(g)-1)}}^{-1} \boldsymbol{\xi}^{(g)} \varepsilon^{(g)}}{\sigma(\widehat{\pi}_{\mathcal{I}(k(g)-1)}; \mathbb{G})}.$$

One can show that  $\eta_7$  forms a mean-zero martingale with respect to the filtration  $\{\sigma(\mathcal{F}^{(g)})\}_{g \geq nT/K}$ .

Using similar arguments in the proof of Theorem 1, we can show that

$$\max_{g \in \{nT/K+1, \dots, nT\}} \left| \frac{\{\int_x \mathbf{U}_{\widehat{\pi}_{\mathcal{I}(k(g)-1)}}(x) \mathbb{G}(dx)\}^\top \Sigma_{\widehat{\pi}_{\mathcal{I}(k(g)-1)}}^{-1} \boldsymbol{\xi}^{(g)} \varepsilon^{(g)}}{\sigma(\widehat{\pi}_{\mathcal{I}(k(g)-1)}; \mathbb{G})} \right| \xrightarrow{P} 0 \quad (\text{E.97})$$

Moreover, we have

$$\begin{aligned}
& \left| \frac{K}{nT(K-1)} \sum_{g=nT/K+1}^{nT} \frac{[\{\int_x \mathbf{U}_{\widehat{\pi}_{\mathcal{I}(k(g)-1)}}(x) \mathbb{G}(dx)\}^\top \Sigma_{\widehat{\pi}_{\mathcal{I}(k(g)-1)}}^{-1} \boldsymbol{\xi}^{(g)} \varepsilon^{(g)}\}^2]}{\sigma^2(\widehat{\pi}_{\mathcal{I}(k(g)-1)}; \mathbb{G})} - 1 \right| \\
& \leq \max_{k \in \{2, \dots, K\}} \left| \frac{K}{nT} \sum_{(i,t) \in \mathcal{I}_k} \frac{[\{\int_x \mathbf{U}_{\widehat{\pi}_{\mathcal{I}_{k-1}}}(x) \mathbb{G}(dx)\}^\top \Sigma_{\widehat{\pi}_{\mathcal{I}_{k-1}}}^{-1} \boldsymbol{\xi}_{i,t} \varepsilon_{i,t}\}^2]}{\sigma^2(\widehat{\pi}_{\mathcal{I}_{k-1}}; \mathbb{G})} - 1 \right| \\
& \leq \max_{k \in \{2, \dots, K\}} \frac{\|\{\int_x \mathbf{U}_{\widehat{\pi}_{\mathcal{I}_{k-1}}}(x) \mathbb{G}(dx)\}^\top \Sigma_{\widehat{\pi}_{\mathcal{I}_{k-1}}}^{-1}\|_2^2 \|\widehat{\boldsymbol{\Omega}}_{\mathcal{I}_k, \widehat{\pi}_{\mathcal{I}_{k-1}}}^* - \boldsymbol{\Omega}_{\widehat{\pi}_{\mathcal{I}_{k-1}}}\|_2}{\sigma^2(\widehat{\pi}_{\mathcal{I}_{k-1}}; \mathbb{G})},
\end{aligned}$$

where  $\widehat{\boldsymbol{\Omega}}_{\mathcal{I}_k, \widehat{\pi}_{\mathcal{I}_{k-1}}}^* = |\mathcal{I}_k|^{-1} \sum_{(i,t) \in \mathcal{I}_k} \boldsymbol{\xi}_{i,t} \boldsymbol{\xi}_{i,t}^\top \varepsilon_{i,t}^2$ .

Similar to the proof of Theorem 1, we can show  $\max_{k \in \{2, \dots, K\}} \|\widehat{\boldsymbol{\Omega}}_{\mathcal{I}_k, \widehat{\pi}_{\mathcal{I}_{k-1}}}^* - \boldsymbol{\Omega}_{\widehat{\pi}_{\mathcal{I}_{k-1}}}\|_2 = o_p(1)$ . Similar to (E.49), we can show there exists some constant  $c_7 > 0$  such that

$$\max_{k \in \{2, \dots, K\}} \frac{\|\{\int_x \mathbf{U}_{\widehat{\pi}_{\mathcal{I}_{k-1}}}(x) \mathbb{G}(dx)\}^\top \Sigma_{\widehat{\pi}_{\mathcal{I}_{k-1}}}^{-1}\|_2^2}{\sigma^2(\widehat{\pi}_{\mathcal{I}_{k-1}}; \mathbb{G})} \leq c_7. \quad (\text{E.98})$$

It follows that

$$\left| \frac{K}{nT(K-1)} \sum_{g=nT/K+1}^{nT} \frac{[\{\int_x \mathbf{U}_{\widehat{\pi}_{\mathcal{I}(k(g)-1)}}(x) \mathbb{G}(dx)\}^\top \Sigma_{\widehat{\pi}_{\mathcal{I}(k(g)-1)}}^{-1} \boldsymbol{\xi}^{(g)} \varepsilon^{(g)}\}^2]}{\sigma^2(\widehat{\pi}_{\mathcal{I}(k(g)-1)}; \mathbb{G})} - 1 \right| \xrightarrow{P} 0. \quad (\text{E.99})$$

Using a martingale central limit theorem for triangular arrays (Corollary 2.8 of McLeish, 1974), we have by (E.97) and (E.99) that  $\eta_7 \xrightarrow{d} N(0, 1)$ . The proof is hence completed.

## E.11 Proof of Theorem 3

Based on the discussions in Section 3.2.3, it suffices to show

$$\mathbb{E} \int_x \sum_{a \in \mathcal{A}} Q^{opt}(x, a) \{\pi^{opt}(a|x) - \widehat{\pi}_{\mathcal{I}}(a|x)\} \mathbb{G}(dx) = O(|\mathcal{I}|^{-b_*(1+\alpha)}),$$

and

$$\mathbb{E} \int_x \sum_{a \in \mathcal{A}} Q^{opt}(x, a) \{\pi^{opt}(a|x) - \widehat{\pi}_{\mathcal{I}}(a|x)\} dx = O(|\mathcal{I}|^{-b_*(1+\alpha)}). \quad (\text{E.100})$$

We only prove (E.100) for brevity. Under the given conditions, we have  $\Pr(\mathcal{A}_0) \geq 1 - O(|\mathcal{I}|^{-\kappa})$ , where

$$\mathcal{A}_0 = \left\{ \sup_{x \in \mathbb{X}, a \in \mathcal{A}} |\widehat{Q}_{\mathcal{I}}(x, a) - Q^{opt}(x, a)| \leq C |\mathcal{I}|^{-b_*} \right\}.$$

Notice that

$$\begin{aligned}
& \mathbb{E} \int_x \sum_{a \in \mathcal{A}} Q^{opt}(x, a) \{ \pi^{opt}(a|x) - \widehat{\pi}_{\mathcal{I}}(a|x) \} dx \\
&= \mathbb{E} \int_x \sum_{a \in \mathcal{A}} Q^{opt}(x, a) \{ \pi^{opt}(a|x) - \widehat{\pi}_{\mathcal{I}}(a|x) \} dx \mathbb{I}(\mathcal{A}_0^c) \\
&+ \mathbb{E} \int_x \sum_{a \in \mathcal{A}} Q^{opt}(x, a) \{ \pi^{opt}(a|x) - \widehat{\pi}_{\mathcal{I}}(a|x) \} dx \mathbb{I}(\mathcal{A}_0).
\end{aligned} \tag{E.101}$$

Under A1,  $Q^{opt}$  is uniformly bounded. Therefore, the first term on the RHS of (E.101) is upper bounded by  $O\{\Pr(\mathcal{A}_0^c)\} = O(|\mathcal{I}|^{-\kappa})$ . Since  $\kappa$  can be chosen arbitrarily large, it suffices to show

$$\mathbb{E} \int_x \sum_{a \in \mathcal{A}} Q^{opt}(x, a) \{ \pi^{opt}(a|x) - \widehat{\pi}_{\mathcal{I}}(a|x) \} dx \mathbb{I}(\mathcal{A}_0) = O\{|\mathcal{I}|^{-b_*(1+\alpha)}\}. \tag{E.102}$$

For any  $x \in \mathbb{X}$ , suppose

$$\max_a Q^{opt}(x, a) - \max_{a \in \mathcal{A} - \arg \max_{a'} Q^{opt}(x, a')} Q^{opt}(x, a) > 2C|\mathcal{I}|^{-b_*}. \tag{E.103}$$

Under the event defined in  $\mathcal{A}_0$ , we have

$$\max_a \widehat{Q}_{\mathcal{I}}(x, a) - \max_{a \in \mathcal{A} - \arg \max_{a'} Q^{opt}(x, a')} \widehat{Q}_{\mathcal{I}}(x, a) > 0,$$

and hence

$$\{a \in \mathcal{A} : \widehat{\pi}(a|x) = 1\} \subseteq \arg \max_{a \in \mathcal{A}} Q^{opt}(x, a).$$

Thus, we have

$$\sum_{a \in \mathcal{A}} Q^{opt}(x, a) \{ \pi^{opt}(a|x) - \widehat{\pi}_{\mathcal{I}}(a|x) \} = 0,$$

when (E.103) holds. Let  $\mathbb{X}_*$  denote the set of  $x$  that satisfies (E.103). It follows that

$$\begin{aligned}
& \mathbb{E} \int_x \sum_{a \in \mathcal{A}} Q^{opt}(x, a) \{ \pi^{opt}(a|x) - \widehat{\pi}_{\mathcal{I}}(a|x) \} dx \mathbb{I}(\mathcal{A}_0) \\
&= \mathbb{E} \int_x \sum_{a \in \mathcal{A}} Q^{opt}(x, a) \{ \pi^{opt}(a|x) - \widehat{\pi}_{\mathcal{I}}(a|x) \} \mathbb{I}(x \in \mathbb{X}_*^c) dx \mathbb{I}(\mathcal{A}_0).
\end{aligned} \tag{E.104}$$

Let  $\widehat{a}_{\mathcal{I}}(x) = \text{sarg} \max_a \widehat{Q}_{\mathcal{I}}(a, x)$ . Similarly, we can show the event  $\widehat{a}_{\mathcal{I}}(x) \notin \arg \max_{a \in \mathcal{A}} Q^{opt}(x, a)$  occurs only when

$$\max_a Q^{opt}(x, a) - Q^{opt}(x, \widehat{a}_{\mathcal{I}}(x)) \leq 2C|\mathcal{I}|^{-b_*}. \quad (\text{E.105})$$

Since  $\max_a Q^{opt}(x, a) - Q^{opt}(x, \widehat{a}_{\mathcal{I}}(x)) = \sum_a Q^{opt}(x, a) \{\pi^{opt}(a|x) - \widehat{\pi}_{\mathcal{I}}(a|x)\}$ , we obtain

$$\begin{aligned} & \mathbb{E} \int_x \sum_{a \in \mathcal{A}} Q^{opt}(x, a) \{\pi^{opt}(a|x) - \widehat{\pi}_{\mathcal{I}}(a|x)\} \mathbb{I}(x \in \mathbb{X}_*^c) dx \mathbb{I}(\mathcal{A}_0) \\ & \leq 2C|\mathcal{I}|^{-b_*} \int_x \mathbb{I}(x \in \mathbb{X}_*^c) dx = 2C|\mathcal{I}|^{-b_*} O(|\mathcal{I}|^{-b_*\alpha}) = O(|\mathcal{I}|^{-b_*(1+\alpha)}), \end{aligned}$$

where the first equality follows from A5. Combining this together with (E.104) yields (E.102). The proof is hence completed.

## E.12 Proof of Theorem 4

For a given  $\varepsilon > 0$ , let  $\mathcal{A}_* = \{\max_a Q^{opt}(x, a) - Q^{opt}(x, \widehat{a}_{\mathcal{I}}(x)) \leq \varepsilon\}$ . Notice that

$$\begin{aligned} & \mathbb{E} \int_x \sum_{a \in \mathcal{A}} Q^{opt}(x, a) \{\pi^{opt}(a|x) - \widehat{\pi}_{\mathcal{I}}(a|x)\} dx \quad (\text{E.106}) \\ & = \mathbb{E} \int_x \sum_{a \in \mathcal{A}} Q^{opt}(x, a) \{\pi^{opt}(a|x) - \widehat{\pi}_{\mathcal{I}}(a|x)\} \mathbb{I}(x \in \mathcal{A}_*) dx \\ & + \mathbb{E} \int_x \sum_{a \in \mathcal{A}} Q^{opt}(x, a) \{\pi^{opt}(a|x) - \widehat{\pi}_{\mathcal{I}}(a|x)\} \mathbb{I}(x \in \mathcal{A}_*^c) dx. \end{aligned}$$

Using similar arguments in the proof of Theorem 3, we can show

$$\begin{aligned} & \mathbb{E} \int_x \sum_{a \in \mathcal{A}} Q^{opt}(x, a) \{\pi^{opt}(a|x) - \widehat{\pi}_{\mathcal{I}}(a|x)\} \mathbb{I}(x \in \mathcal{A}_*) dx \leq \varepsilon \mathbb{E} \int_x \mathbb{I}(x \in \mathcal{A}_*) dx \\ & \leq \varepsilon \lambda \left\{ x : \max_a Q^{opt}(x, a) - \max_{a \in \mathcal{A} - \arg \max_{a'} Q^{opt}(x, a')} Q^{opt}(x, a) \leq \varepsilon \right\} = O(\varepsilon^{1+\alpha}). \quad (\text{E.107}) \end{aligned}$$

Moreover, similar to (E.105), we can show the event  $\widehat{a}_{\mathcal{I}}(x) \notin \arg \max_{a \in \mathcal{A}} Q^{opt}(x, a)$  occurs only when

$$\max_a Q^{opt}(x, a) - Q^{opt}(x, \widehat{a}_{\mathcal{I}}(x)) \leq 2 \max_{a \in \mathcal{A}} |\widehat{Q}_{\mathcal{I}}(x, a) - Q^{opt}(x, a)|.$$

It follows that

$$\begin{aligned}
& \mathbb{E} \int_x \sum_{a \in \mathcal{A}} Q^{opt}(x, a) \{\pi^{opt}(a|x) - \widehat{\pi}_{\mathcal{I}}(a|x)\} \mathbb{I}(x \in \mathcal{A}_*^c) dx \\
& \leq \mathbb{E} \int_x \frac{4 \max_{a \in \mathcal{A}} |\widehat{Q}_{\mathcal{I}}(x, a) - Q^{opt}(x, a)|^2}{\sum_{a \in \mathcal{A}} Q^{opt}(x, a) \{\pi^{opt}(a|x) - \widehat{\pi}_{\mathcal{I}}(a|x)\}} \mathbb{I}(x \in \mathcal{A}_*^c) dx \\
& \leq \frac{4}{\varepsilon} \mathbb{E} \int_x \max_{a \in \mathcal{A}} |\widehat{Q}_{\mathcal{I}}(x, a) - Q^{opt}(x, a)|^2 dx \leq \frac{4}{\varepsilon} \mathbb{E} \int_x \sum_{a \in \mathcal{A}} |\widehat{Q}_{\mathcal{I}}(x, a) - Q^{opt}(x, a)|^2 dx \\
& = O(\varepsilon^{-1} |\mathcal{I}|^{-2b_*}).
\end{aligned}$$

Combining this together with (E.106) and (E.107) yields that

$$\mathbb{E} \int_x \sum_{a \in \mathcal{A}} Q^{opt}(x, a) \{\pi^{opt}(a|x) - \widehat{\pi}_{\mathcal{I}}(a|x)\} dx = O(\varepsilon^{1+\alpha}) + O(\varepsilon^{-1} |\mathcal{I}|^{-2b_*}).$$

The proof is hence completed by setting  $\varepsilon = |\mathcal{I}|^{-2b_*/(2+\alpha)}$ .