
Doubly Robust Off-Policy Value and Gradient Estimation for Deterministic Policies

Nathan Kallus, Masatoshi Uehara *
Cornell University and Cornell Tech
New York, NY
kallus@cornell.edu, mu223@cornell.edu

Abstract

Offline reinforcement learning, wherein one uses off-policy data logged by a fixed behavior policy to evaluate and learn new policies, is crucial in applications where experimentation is limited such as medicine. We study the estimation of policy value and gradient of a deterministic policy from off-policy data when actions are continuous. Targeting deterministic policies, for which action is a deterministic function of state, is crucial since optimal policies are always deterministic (up to ties). In this setting, standard importance sampling and doubly robust estimators for policy value and gradient fail because the density ratio does not exist. To circumvent this issue, we propose several new doubly robust estimators based on different kernelization approaches. We analyze the asymptotic mean-squared error of each of these under mild rate conditions for nuisance estimators. Specifically, we demonstrate how to obtain a rate that is independent of the horizon length.

1 Introduction

Offline reinforcement learning (RL), wherein one uses off-policy data logged by a fixed behavior policy to evaluate and learn new policies, is crucial in applications where experimentation is limited (Farajtabar et al., 2018; Bibaut et al., 2019; Liu et al., 2018; Kallus and Uehara, 2019b). A key application is RL for healthcare (Gottesman et al., 2019; Murphy, 2003). Since it is not possible to collect new data, it is crucial to efficiently use the available data. Recent work on off-policy evaluation (OPE; Kallus and Uehara, 2020a, 2019a) have shown how efficiently taking advantage of problem structure, such as Markovianness and ergodicity, can improve OPE and tackle the well-known issue of OPE known as the curse of horizon (Liu et al., 2018). Kallus and Uehara (2020b) applied these advances to off-policy *learning* using a policy gradient approach, i.e., proposing efficient off-policy estimators for the policy gradient and incorporating them into gradient ascent methods.

All the aforementioned methods, however, cannot be directly applied to the evaluation and learning of deterministic policies when actions are continuous since the density ratio (Radon-Nikodym derivative) does not exist: the behavior policy usually has zero mass on single actions (more generally, it can have at most countably-many atoms while the evaluation policy may take any of a continuum of actions). This question is important since the maximum-value policy is generally deterministic (up to ties between actions) so if one seeks optimal policies one should focus on deterministic ones. In the bandit setting (horizon of one action), several recent works tackle this problem (Bibaut and J. van der Laan, 2017; Kallus and Zhou, 2018; Colangelo and Lee, 2019), but applying these methods in a straightforward manner to RL may lead to a bad convergence rate that deteriorates with horizon.

In this paper, we propose several doubly robust off-policy value and gradient estimators for deterministic policies in an RL setting. We analyze the asymptotic mean-squared error (MSE) of each

*corresponding author

Table 1: Comparison of off-policy value and gradient estimators for deterministic policies. Proposed estimators are typeset in bold. “MSE” is the *convergence rate* when nuisances are estimated at rate under “Rate,” *irrespective* of choice of nuisance estimators. “–” means the convergence rate depends on the choice of estimators. “Nuis” are nuisances for OPE. “Nuis+” are *additional* nuisances for policy gradient. “D/I” is whether differentiation (D) or integration (I) of $\hat{q}_t(s_t, a_t)$ wrt a_t is required.

(a) OPE				(b) Off-policy gradient			
	MSE	Nuis	Rate		MSE	Nuis+	D/I
DM	–	q^{π^e}	–	DPG	–	$d^{q^{\pi^e}}$	D
CDRD	$n^{-\frac{4}{H+4}}$	$\lambda^{\mathcal{K}}, q^{\mathcal{K}}$	$n^{-\frac{1}{H+4}}$	CPGD	$n^{-\frac{4}{H+6}}$	$d^{\lambda^{\mathcal{K}}}, d^{q^{\mathcal{K}}}$	D
CDRK	$n^{-\frac{4}{H+4}}$	$\lambda^{\pi^e}, q^{\pi^e}$	$n^{-\frac{1}{H+4}}$	CPGK	$n^{-\frac{4}{H+6}}$	$d^{\lambda^{\pi^e}}, d^{q^{\pi^e}}$	I
MDRD	$n^{-\frac{4}{5}}$	$w^{\mathcal{K}}, q^{\mathcal{K}}$	$n^{-\frac{1}{5}}$	MPGD	$n^{-\frac{4}{7}}$	$d^{w^{\mathcal{K}}}, d^{q^{\mathcal{K}}}$	D
MDRK	$n^{-\frac{4}{5}}$	w^{π^e}, q^{π^e}	$n^{-\frac{1}{5}}$	MPGK	$n^{-\frac{4}{7}}$	$d^{w^{\pi^e}}, d^{q^{\pi^e}}$	I

estimator under extremely lax conditions that accommodate flexible learning of the nuisances that appear in the estimator (such as q -functions). Specifically, we propose estimators of policy value and gradient with MSE convergence rate that does not deteriorate with horizon and the leading term’s coefficient has only a polynomial dependence on horizon. These results are summarized in Table 1.

2 Preliminaries

Problem set up Consider an H -long time-varying Markov decision process (MDP), with states $s_t \in \mathcal{S}_t$, actions $a_t \in \mathcal{A}_t$, rewards $r_t \in [0, R_{\max}]$, initial state distribution $p_1(s_1)$, transition distributions $p_{t+1}(s_{t+1} | s_t, a_t)$, and reward distribution $p_t(r_t | s_t, a_t)$, for $t = 1, \dots, H$. A policy $\pi = (\pi_t(a_t | s_t))_{t \leq H}$ induces a distribution over trajectories $\mathcal{T} = (s_1, a_1, r_1, \dots, s_T, a_H, r_H)$:

$$p_\pi(\mathcal{T}) = p_1(s_1)\pi_1(a_1 | s_1)p_1(r_1 | s_1, a_1) \prod_{t=2}^H p_t(s_t | s_{t-1}, a_{t-1})\pi_t(a_t | s_t)p_t(r_t | s_t, a_t). \quad (1)$$

In this paper, we focus on *continuous* actions, $\mathcal{A}_t \subseteq \mathbb{R}$. For brevity we focus on the univariate case; the extension to multivariate actions is straightforward. We are interested in the value, $J = \mathbb{E}_{p_{\pi^e}}[\sum_{t=1}^H r_t]$, of a given policy, π^e , called the evaluation policy. In particular, we consider the case where π^e is *deterministic* in that it is given by maps $\tau = (\tau_t)_{t \leq H}$, $\tau_t : \mathcal{S}_t \rightarrow \mathbb{R}$ such that $\pi_t^e(a_t | s_t) = \delta(a_t - \tau_t(s_t))$ is the Dirac measure at $\tau_t(s_t)$, meaning when we follow π^e , a_t is a function τ_t of s_t . When τ is parametrized as $\tau = \tau_\theta$ by some parameter $\theta \in \Theta$, we are also often interested in the policy gradient, $Z = \nabla_\theta J$, as it can be used for policy *learning* via gradient ascent. We often drop the subscript and understand ∇ to be with respect to (wrt) θ . When studying policy gradient estimation, we will assume throughout that $\tau_{\theta,t}(s_t)$ is almost surely differentiable in θ and that $\|\nabla \tau_t(s_t)\|_{\text{op}} \leq \Upsilon$ for some $\Upsilon < \infty$, where $\|\cdot\|_{\text{op}}$ is the matrix operator norm. Additionally, in theoretical results, we will assume actions are bounded: $\text{support}(\pi_t^b(\cdot | s_t)) = [0, 1]$, $\tau_t(s_t) \in (0, 1)$.

In the *offline* setting, the data available to us for estimating J and Z consists only of trajectory observations from some *different* fixed policy, π^b , called the *behavior policy*:

$$\mathcal{T}^{(1)}, \dots, \mathcal{T}^{(n)} \sim p_{\pi^b}, \mathcal{T}^{(i)} = (S_1^{(i)}, A_1^{(i)}, R_1^{(i)}, \dots, S_H^{(i)}, A_H^{(i)}, R_H^{(i)}). \quad (\text{off-policy data})$$

Let $w_t^{\pi^e}(s_t) = p_{\pi^e}(s_t)/p_{\pi^b}(s_t)$ denote the marginal density ratio, where $p_{\pi^e}(s_t)$, $p_{\pi^b}(s_t)$ are the marginal densities of s_t under $p_{\pi^e}(\mathcal{T})$, $p_{\pi^b}(\mathcal{T})$, respectively. Let $q_t^{\pi^e}(s_t, a_t) = \mathbb{E}_{p_{\pi^e}}[\sum_{k=t}^H r_k | s_t, a_t]$, $v_t^{\pi^e}(s_t) = \mathbb{E}_{p_{\pi^e}}[\sum_{k=t}^H r_k | s_t]$ be π^e ’s q - and v -functions. We assume throughout that $1/\pi_t^b(a_t | s_t) \leq C_1$, $w_t^{\pi^e}(s_t) \leq C_2$. For clarity, we reserve capital letters for observed data, dropping superscripts $^{(i)}$ for a generic data point, and use lower case for generic MDP random variables. All expectations without subscripts are taken wrt p_π^b . The *empirical expectation* is $\mathbb{P}_n f = \frac{1}{n} \sum_{i=1}^n f(\mathcal{T}^{(i)})$. Define the L_2 norm as $\|f(\mathcal{T})\|_2 = \mathbb{E}[f(\mathcal{T}^2)]^{1/2}$. We let $\cdot^{(i)}$ denote the i -th order derivative wrt action alone and define the max Sobolev norm as $\|f\|_{j,\infty} = \max_{i=0,\dots,j} \|f^{(i)}\|_\infty$. We emphasize that, e.g., $q^{(i)}(s_t, a_t)$ refers to differentiating only wrt a_t alone.

Integral kernels In developing estimators we will use a second-order kernel $k : \mathbb{R} \rightarrow \mathbb{R}$, i.e., $\int k(u)du = 1$, $\int uk(u)du = 0$, $M_2(k) = \int u^2k(u)du < \infty$. We define $\Omega_2^{(i)}(k) = \int (k^{(i)}(u))^2 du$. Given a bandwidth h , let $K_h(u) = h^{-1}k(u/h)$. Examples of differentiable kernels include Gaussian $k(u) \propto \exp(-u^2)$, biweight $k(u) \propto \max(0, 1 - u^2)^2$, and triweight $k(u) \propto \max(0, 1 - u^2)^3$.

Background on Offline Evaluation and Policy Gradient Direct estimation of q -functions (direct method, DM; Munos and Szepesvári, 2008) and step-wise importance sampling (IS; Precup et al., 2000) are two common approaches for OPE. However, the former is known to suffer from the high variance and the latter from model misspecification. The doubly robust (DR) estimate combines the two, but its asymptotic MSE can still grow exponentially in horizon (Jiang and Li, 2016; Thomas and Brunskill, 2016). Kallus and Uehara (2020a) show that the efficient MSE in the MDP case is polynomial in $\mathcal{O}(H^2/n)$ and give an estimator achieving it by combining marginalized IS (Xie et al., 2019) and q -modeling using cross-fold estimation (Chernozhukov et al., 2018). OPE and off-policy policy gradient estimation are closely related (Huang and Jiang, 2019). Efficiency analysis and efficient estimators for off-policy policy gradients was recently given in Kallus and Uehara (2020b).

All of the aforementioned methods assume that the density ratio $\pi_t^e(a_t | s_t)/\pi_t^b(a_t | s_t)$ exists and is bounded. However, when π^e is deterministic this is generally violated as it requires that $\pi_t^b(a_t | s_t)$ has an atom at $\tau_t(s_t)$ but π^b usually has no atoms and at most can only have countably-many, while $\tau_t(s_t)$ can vary continuously, especially for policy learning (e.g., via gradient ascent). In the bandit setting ($H = 1$), Kallus and Zhou (2018) recently showed that OPE is feasible under additional smoothness assumptions on induced action densities. The core idea is to approximate the deterministic policy by stochastic policy based on kernels. Namely, under appropriate smoothness,

$$\lim_{h \rightarrow 0} \mathbb{E} \left[\frac{K_h(A_1 - \tau_1(S_1))R_1}{\pi^b(A_1 | S_1)} \right] = \mathbb{E}_{\pi^e} [r_1]. \quad (2)$$

Kallus and Zhou (2018), assuming known behavior policy, therefore propose an IS-type estimator $\mathbb{P}_n \left[\frac{K_h(A_1 - \tau_1(S_1))R_1}{\pi^b(A_1 | S_1)} \right]$ for h appropriately shrinking in n and analyze its bias and variance. This essentially amounts to approximating the deterministic policy with a stochastic one concentrated near, but not fully at, $\tau_1(S_1)$. Several works similarly deal with causal inference with continuous treatments (Imai and van Dyk, 2004; Hirano and Imbens, 2005; Galvao and Wang, 2015; Fong et al., 2018; Wu et al., 2018; Kennedy et al., 2017; Su et al., 2019) focusing on a single ($H = 1$) constant ($\tau_1(s_1) = a_1^*$) action. Although we do not explicitly use counterfactual notation, our estimand is equivalent to a counterfactual one under sequential ignorability (Hernan and Robins, 2019).

3 Bandit Setting, $H = 1$

For lucid presentation, we first develop off-policy value and gradient estimators in the bandit setting, i.e., $H = 1$. In this section, since $H = 1$, we omit the time index, e.g., letting $S = S_1$, $q = q_1$, etc.

3.1 Off-Policy Deterministic Policy Evaluation

There are actually several different ways to kernelize the deterministic policy. To motivate our estimators, note that, under appropriate smoothness (see Theorem 1 below) we have that

$$\lim_{h \rightarrow 0} \mathbb{E} \left[\frac{K_h(A - \tau(S))\{R - f_1\}}{f_2} + f_3 \right] = \mathbb{E}_{\pi^e} [r_1], \quad (3)$$

when each of f_1, f_2, f_3 takes any of the following two values, giving 2^3 possible combinations:

$$f_1 = \begin{cases} q(S, A) \\ q(S, \tau(S)) \end{cases}, \quad f_2 = \begin{cases} \pi^b(A | S) \\ \pi^b(\tau(S) | S) \end{cases}, \quad f_3 = \begin{cases} \int q(S, a)K_h(a - \tau(S))da \\ q(S, \tau(S)) \end{cases}.$$

This general form includes the form of some previously proposed estimators, including the IS-type estimator in Kallus and Zhou (2018) ($f_1 = 0, f_2 = \pi^b(A | S), f_3 = 0$), the DR-type estimator in Kallus and Zhou (2018) ($f_1 = q(S, A), f_2 = \pi^b(A | S), f_3 = q(S, \tau(S))$), and the DR-type estimator in Colangelo and Lee (2019) ($f_1 = q(S, \tau(S)), f_2 = \pi^b(\tau(S) | S), f_3 = q(S, \tau(S))$).

Based on this, we propose a general DR-type estimator with nuisance functions f_1, f_2, f_3 : first, we split the data randomly into two halves \mathcal{U}_1 and \mathcal{U}_2 ; then, we define the estimator as

$$\frac{1}{2} \mathbb{P}_{\mathcal{U}_1} \left[\frac{K_h(A - \tau(S))\{R - \hat{f}_1^{[1]}\}}{\hat{f}_2^{[1]}} + \hat{f}_3^{[1]} \right] + \frac{1}{2} \mathbb{P}_{\mathcal{U}_2} \left[\frac{K_h(A - \tau(S))\{R - \hat{f}_1^{[2]}\}}{\hat{f}_2^{[2]}} + \hat{f}_3^{[2]} \right], \quad (4)$$

where $\hat{f}_j^{[k]}$ is an estimate of f_j based on estimating π^b, q by $\hat{\pi}^{b,[k]}, \hat{q}^{[k]}$ using only the data in \mathcal{U}_{3-k} . This technique is called a cross-fitting, which is used to avoid metric entropy conditions on nuisance estimators (Chernozhukov et al., 2018). All of our nuisances can be estimated by standard nonparametric density or regression estimators (Hansen, 2009). To simplify notation, we often write this and similar estimators as $\hat{J} = \mathbb{P}_n \left[K_h(A - \tau(S)) \{R - \hat{f}_1\} / \hat{f}_2 + \hat{f}_3 \right]$, where implicitly \hat{f}_j are fit using cross-fitting on the half of the data that excludes the data point on which it is evaluated.

We next analyze two primary cases: $f_1 = q(S, A)$, $f_2 = \pi^b(A | S)$, $f_3 = \int q(S, a) K_h(a - \tau(S)) da$, which refer to as \hat{J}^K (for “kernel”), and $f_1 = q(S, \tau(S))$, $f_2 = \pi^b(\tau(S) | S)$, $f_3 = q(S, \tau(S))$, which refer to as \hat{J}^D (for “deterministic” as we plug in the deterministic policy into the nuisances).

Theorem 1. *Suppose for $i = 1, 2$, $\mathbb{E}[\|\hat{\pi}^{b,[i]} - \pi^b\|_{1,\infty}] = o(1)$, $\mathbb{E}[\|\hat{q}^{[i]} - q\|_{1,\infty}] = o(1)$, $\mathbb{E}[\|\hat{\pi}^{b,[i]} - \pi^b\|_\infty \|\hat{q}^{[i]} - q\|_\infty] = o((nh)^{-1/2})$, $nh^5 = \mathcal{O}(1)$, $nh \rightarrow \infty$, that $\pi^b(a | s)$, $q(s, a)$ are twice continuously differentiable wrt a for almost all s , and that $\hat{\pi}^{b,[i]}, \hat{q}^{[i]}$ are uniformly bounded by a constant. Then, the bias and variance of \hat{J}^D are $\mathbb{E}[\hat{J}^D] - J = 0.5M_2(k)h^2B + o((nh)^{-1/2})$, $\text{var}[\hat{J}^D] = \frac{\Omega_2(k)}{nh}(V + o(1))$, where*

$$B = \mathbb{E}[q^{(2)}(S, \tau(S)) + 2q^{(1)}(S, \tau(S))\pi^{b(1)}(\tau(S)|S)/\pi^b(\tau(S)|S)], \quad V = \mathbb{E} \left[\frac{\text{var}[R|S, \tau(S)]}{\pi^b(\tau(S)|S)} \right].$$

If additionally $\hat{\pi}^{b,[i]}(a | s)$, $\hat{q}^{[i]}(s, a)$ are twice continuously differentiable wrt a , then the same holds for \hat{J}^K with $B = \mathbb{E}[q^{(2)}(S, \tau(S))]$ and the same V as the above. In both cases, setting $h = \Theta(n^{-1/5})$ yields the minimal MSE of order $\mathcal{O}(n^{-4/5})$.

Here, we assume that nuisance estimation errors converge in expectation (i.e., in L_1). If we instead assume the weaker convergence in probability, we can obtain the same guarantee on the bias and variance, conditioned on an event that occurs with high probability. Refer to Appendix E.

Remark 1. Three things should be noted. First, the best MSE rate achievable in the result is $\mathcal{O}(n^{-4/5})$, which is slower than the usual rate $\mathcal{O}(n^{-1})$ in OPE of stochastic policies under positivity. This slow rate is expected because our estimand in the deterministic case is not regular (Kennedy et al., 2017). Since the minimax rate for density estimation in a Sobolev class of smoothness parameter 2 is $\mathcal{O}(n^{-4/5})$ (Korostelev, 2011), we expect that this rate is minimax for J among problems satisfying the conditions of Theorem 1. Establishing this formally is future work. Second, for \hat{J}^D , the only condition on nuisance estimators is a sub-parametric rate, which can, e.g., be satisfied when each nuisance converges at the rate $o(n^{-1/5})$ for $h = \Theta(n^{-1/5})$. This condition appears weaker than the $o(n^{-1/4})$ nuisance rate required in usual OPE (Kallus and Uehara, 2020a; Chernozhukov et al., 2018); however, the required norm itself is stronger since $\|\cdot\|_2 \leq \|\cdot\|_\infty$. In contrast, when just q is estimated at rate $o(n^{-1/5})$, one cannot guarantee a similar $\mathcal{O}(n^{-4/5})$ MSE rate for DM without additional assumptions; a simple triangle inequality yields only an $o(n^{-2/5})$ MSE rate. The same is true for IS when we estimate π^b only at $o(n^{-1/5})$ rate.² Third, both the variances and rates for \hat{J}^D and \hat{J}^K are the same, while the bias constant is slightly different. For brevity, in the following, we focus on deriving theoretical properties for \hat{J}^K where similar results are easily obtainable for \hat{J}^D .

Optimality of \hat{J}^K in terms of leading constant Kallus and Zhou (2018) computed the asymptotic bias and variance for the IS estimator ($f_1 = 0$, $f_2 = \pi^b(A | S)$, $f_3 = 0$) with known behavior policy (hence no nuisances). The estimator \hat{J}^K is obtained by adding the control variate $f_1 = f_3 = q(S, A)$ and while the two estimators have the *same* asymptotic bias, the leading variance term is *smaller*, having $\mathbb{E} \left[\frac{\text{var}[R|S, \tau(S)]}{\pi^b(\tau(S)|S)} \right]$ instead of the larger $\mathbb{E} \left[\frac{\mathbb{E}[R^2|S, \tau(S)]}{\pi^b(\tau(S)|S)} \right]$ in Kallus and Zhou (2018). Thus, the advantage of our DR-type estimators over the IS estimator is not only ensuring $\mathcal{O}(n^{-4/5})$ convergence with an *unknown* behavior policy but also in providing an improvement in the variance leading term. In fact, we can prove \hat{J}^K is optimal among a class of estimators.

Corollary 1. *Assume $\pi^b(a|s)$, $f(s, a)$, $q(s, a)$ are C^2 -functions wrt a . Then $\mathbb{P}_n \left[\frac{K_h(A - \tau(S))}{\pi^b(A|S)} \{R - f(S, A)\} + f(S, A) \right]$ has a bias independent of f and variance minimized by letting $f = q$.*

²In the special case when q (or π^b , respectively) is estimated using kernel estimators, we can obtain an MSE rate $\mathcal{O}(n^{-4/5})$ for DM (or IS, respectively) but it requires much stricter conditions on the smoothness, including smoothness in the state variable in addition to smoothness in action (Hsu et al., 2018; Lee, 2018).

This would correspond to an efficiency result in semiparametric theory, but that cannot be applied here since our estimand is not regular (Kennedy et al., 2017). It is also difficult to compare $\hat{J}^{\mathcal{K}}$ and $\hat{J}^{\mathcal{D}}$ since the bias terms are different. We leave further investigation of optimality to future work.

Remark 2 (Relation with previous literature). Bibaut and J. van der Laan (2017) proposed a general approach for the estimation of non-regular estimands using smoothing. See the estimator in Example 3; however, they assumed a behavior policy is known. Foster and Syrgkanis (2019, Section 8) also touched on the idea of case \mathcal{K} . However, they did not analyze a mathematical detail.

3.2 Off-Policy Deterministic Policy Gradient Estimation

For a deterministic policy class $\{\tau_\theta(s) : \theta \in \Theta\}$, consider estimating $Z = \nabla J$ at a given θ . Usually policy gradients involve the policy score, $\nabla \log(\pi_\theta^e(a | s))$ (Peters and Schaal, 2006; Kallus and Uehara, 2020b). However, for deterministic policies, these policy scores do not exist. However, assuming that q is differentiable in a immediately yields $Z = \mathbb{E}[q^{(1)}(S, \tau_\theta(S)) \nabla \tau_\theta(S)]$, suggesting this may still be possible under appropriate smoothness.

Deterministic Policy Gradient (DPG) and IS Policy Gradient (ISPG) By taking a derivative of the IS and DM estimators wrt θ , we obtain corresponding policy gradient estimators:

$$\hat{Z}^{\text{IS}} = \mathbb{P}_n[K_h^{(1)}(A - \tau(S))R/\hat{\pi}^b(A | S)], \quad \hat{Z}^{\text{DPG}} = \mathbb{P}_n[\hat{q}^{(1)}(S, \tau_\theta(S)) \nabla \tau_\theta(S)], \quad (5)$$

where $K_h^{(1)}(u) = -h^{-2}k^{(1)}(u/h)$. The latter estimator is a bandit version of DPG (Silver et al., 2014). Like their OPE counterparts, these estimators suffer from high dependence on the nuisance estimates and potentially slow rates.

Doubly Robust DPG By differentiating our OPE estimator $\hat{J}^{\mathcal{K}}$ wrt θ we propose a new policy gradient estimator: $\hat{Z}^{\mathcal{K}} = \mathbb{P}_n[\psi^{\mathcal{K}}(\hat{q}, \hat{\pi}^b)]$ (recall $\hat{q}, \hat{\pi}^b$ are implicitly cross-fit in this notation), where

$$\psi^{\mathcal{K}}(q, \pi^b) = \left\{ \frac{K_h^{(1)}(a - \tau_\theta(s))\{r - q(s, a)\}}{\pi^b(a|s)} + \int K_h^{(1)}(a - \tau_\theta(s))q(s, a)da \right\} \nabla_\theta \tau_\theta(s). \quad (6)$$

Notice this does not involve explicit differentiation of q . Instead the convolution with $K_h^{(1)}$ essentially acts as an estimator for the derivative and may be computationally more stable than differentiating \hat{q} .

Similarly, by differentiating our OPE estimator $\hat{J}^{\mathcal{D}}$ wrt θ , we propose $\hat{Z}^{\mathcal{D}} = \mathbb{P}_n[\psi^{\mathcal{D}}(\hat{q}, \hat{\pi}^b)]$, where

$$\psi^{\mathcal{D}}(q, \pi^b) = \left\{ \frac{K_h^{(1)}(a - \tau_\theta(s))\{r - q(s, \tau_\theta(s))\}}{\pi^b(\tau_\theta(s)|s)} + q^{(1)}(s, \tau_\theta(s)) \right\} \nabla_\theta \tau_\theta(s). \quad (7)$$

This can be understood as DPG *plus* a control variate. As in DPG, $\hat{Z}^{\mathcal{D}}$ requires differentiation of \hat{q} .

Theorem 2. Suppose for $i = 1, 2$, $\mathbb{E}[\|\hat{\pi}^{b,[i]} - \pi^b\|_{2,\infty}] = o(1)$, $\mathbb{E}[\|\hat{q}^{[i]} - q\|_{2,\infty}] = o(1)$, $\mathbb{E}[\|\hat{\pi}^{b,[i]} - \pi^b\|_{1,\infty} \|\hat{q}^{[i]} - q\|_{1,\infty}] = o(n^{-1/2}h^{-3/2})$, $nh^7 = \mathcal{O}(1)$, $nh \rightarrow \infty$, $q(s, a), \pi^b(a | s), \hat{q}^{[i]}(s, a), \hat{\pi}^{b,[i]}(a | s)$ are thrice continuously differentiable functions wrt a , and $\hat{q}^{[i]}, \hat{\pi}^{b,[i]}$ are uniformly bounded by a constant. Then, the bias and variance of $\hat{Z}^{\mathcal{K}}$ are $\mathbb{E}[\hat{Z}^{\mathcal{K}}] - Z = 0.5h^2 M_2(k)\tilde{B} + o(n^{-1/2}h^{-3/2})$, $\text{var}[\hat{Z}^{\mathcal{K}}] = \frac{\Omega_2^{(1)}(k)}{nh^3} \{\tilde{V} + o(1)\}$, where

$$\tilde{B} = \mathbb{E}[\nabla \tau_\theta(S) q^{(3)}(S, \tau_\theta(S))], \quad \tilde{V} = \mathbb{E}\left[\otimes \nabla \tau_\theta(S) \frac{\text{var}[R|S, \tau_\theta(S)]}{\pi^b(\tau_\theta(S)|S)}\right],$$

where $\otimes v = vv^\top$. Setting $h = \Theta(n^{-1/7})$ yields the minimal MSE of order $\mathcal{O}(n^{-4/7})$.

Remark 3. We can obtain a similar result for $\hat{Z}^{\mathcal{D}}$ with slightly different differentiability conditions; we omit the details. The best-achievable MSE rate in Theorem 2, $\mathcal{O}(n^{-4/7})$, matches the minimax rate for density gradient estimation in a Sobolev space of smoothness parameter 3 (Korostelev, 2011). We therefore conjecture the rate for Z is minimax optimal under the assumptions of the theorem.

Remark 4. As in the case of OPE, there are two crucial advantages of $\hat{Z}^{\mathcal{K}}, \hat{Z}^{\mathcal{D}}$ over ISPG and DPG estimators in Eq. (5). First, the required convergence rates on nuisances are weaker and we do not depend on the particular estimators. On the other hand, ISPG and DPG do not have convergence guarantees given only rate conditions on $\hat{\pi}^b$ and \hat{q} , respectively. Second, our leading constant in the variance is smaller than ISDP with an oracle behavior policy, just as in Corollary 1.

Remark 5 ($\hat{Z}^{\mathcal{K}}$ vs $\hat{Z}^{\mathcal{D}}$). Unlike $\hat{Z}^{\mathcal{K}}$, the estimator $\hat{Z}^{\mathcal{D}}$ does not involve integration, which may be computationally preferable. However, direct differentiation of q -functions can often be statistically unstable (see also Athey and Wager, 2017, Section 5.2). In our empirical results in Section 5, we indeed find $\hat{Z}^{\mathcal{K}}$ is superior. Moreover, when we use the Gaussian kernel and polynomial sieve nuisance estimators, the integration can be easily done analytically.

4 Offline RL with Deterministic Policies

We next discuss how to extend the ideas from the previous section to the RL setting where $H \geq 1$ in general. In this setting there are actually different ways to account for the IS part of the estimator, leading to different dependence on horizon. Throughout this section, we will assume the densities $p_t(r_t | s_t, a_t), p_{t+1}(s_{t+1} | s_t, a_t)$ are thrice continuously differentiable wrt action and $\|\int |p_t^{(i)}(r_t | s_t, a_t)|dr_t\|_{\infty} \leq G_1^{(i)} < \infty, \|\int |p_t^{(i)}(s_t | s_{t-1}, a_{t-1})|ds_t\|_{\infty} \leq G_2^{(i)} < \infty$ for $i = 1, 2, 3$. We also assume all of nuisance estimators introduced in this section are uniformly bounded by some constant.

4.1 Off-Policy Deterministic Policy Evaluation

Motivated by DR OPE using cumulative density ratios for the case of stochastic policies (Jiang and Li, 2016), we propose analogous extensions of $\hat{J}^{\mathcal{K}}, \hat{J}^{\mathcal{D}}$ for $H \geq 1$: Cumulative DR case \mathcal{K} (CDRK) $\hat{J}_{\text{cdr}}^{\mathcal{K}} = \mathbb{P}_n[\phi_{\text{cdr}}^{\mathcal{K}}(\hat{q}_t^{\mathcal{K}}, \hat{\pi}_t^{\mathcal{b}})]$ and Cumulative DR case \mathcal{D} (CDRD) $\hat{J}_{\text{cdr}}^{\mathcal{D}} = \mathbb{P}_n[\phi_{\text{cdr}}^{\mathcal{D}}(\hat{q}_t^{\pi^e}, \hat{\pi}_t^{\mathcal{b}})]$, where

$$\begin{aligned}\phi_{\text{cdr}}^{\mathcal{K}}(q_t^{\mathcal{K}}, \pi_t^{\mathcal{b}}) &= v_1^{\mathcal{K}} + \sum_{t=1}^H \lambda_t^{\mathcal{K}} \{r_t - q_t^{\mathcal{K}}(s_t, a_t) + v_{t+1}^{\mathcal{K}}\}, \lambda_t^{\mathcal{K}} = \prod_{k=1}^t \frac{K_h(a_k - \tau_k(s_k))}{\pi_k^{\mathcal{b}}(a_k | s_k)}, \\ \phi_{\text{cdr}}^{\mathcal{D}}(q_t^{\pi^e}, \pi_t^{\mathcal{b}}) &= v_1^{\pi^e} + \sum_{t=1}^H \lambda_t^{\pi^e} \{r_t - q_t^{\pi^e}(s_t, \tau_t(s_t)) + v_{t+1}^{\pi^e}\}, \lambda_t^{\pi^e} = \prod_{k=1}^t \frac{K_h(a_k - \tau_k(s_k))}{\pi_k^{\mathcal{b}}(\tau_k(s_k) | s_k)},\end{aligned}$$

and where $q_t^{\mathcal{K}}$ is the q -function associated with the kernelized evaluation policy, $\pi_t^{\pi^e, \mathcal{K}}(a_t | s_t) = K_h(a_t - \tau_{\theta}(s_t))$. We discuss the estimation of nuisances in Remark 7. Recall we use cross-fitting.

Theorem 3. Suppose for $j \leq H, i = 1, 2, \mathbb{E}[\|\hat{\pi}_j^{b, [i]} - \pi_j^{\mathcal{b}}\|_{1, \infty}] = o(1), \mathbb{E}[\|\hat{q}_j^{\mathcal{K}, [i]} - q_j^{\mathcal{K}}\|_{1, \infty}] = o(1), \mathbb{E}[\|\hat{\pi}_j^{b, [i]} - \pi_j^{\mathcal{b}}\|_{\infty} \|\hat{q}_j^{\mathcal{K}, [i]} - q_j^{\mathcal{K}}\|_{\infty}] = o(n^{-1/2}h^{-H/2}), nh^{H+4} = \mathcal{O}(1), nh \rightarrow \infty$. Then, we have $\mathbb{E}[\hat{J}_{\text{cdr}}^{\mathcal{K}}] - J = 0.5h^2M_2(k)B_H^{\text{cdr}} + o(n^{-1/2}h^{-H/2}), \text{var}[\hat{J}_{\text{cdr}}^{\mathcal{K}}] = \frac{\Omega_2^H(k)}{nh^H}\{V_H^{\text{cdr}} + o(1)\}$, where

$$\begin{aligned}B_H^{\text{cdr}} &= \sum_{t=1}^H \mathbb{E}_{p_{\pi}^e} \left[\frac{r_t p_t^{(2)}(r_t | s_t, \tau_t(s_t))}{p_t(r_t | s_t, \tau_t(s_t))} \right] + \sum_{j=1}^{t-1} \mathbb{E}_{p_{\pi}^e} \left[\frac{r_t p_{j+1}^{(2)}(s_{j+1} | s_j, \tau_j(s_j))}{p_{j+1}(s_{j+1} | s_j, \tau_j(s_j))} \right], \\ V_H^{\text{cdr}} &= \mathbb{E}_{p_{\pi}^e} \left[\frac{1}{\prod_{i=1}^H \pi_i^{\mathcal{b}}(\tau_i(s_i) | s_i)} \text{var}[r_H | s_H, a_H] \right].\end{aligned}$$

In the above, setting $h = \Theta(n^{-1/(H+4)})$ yields the minimal MSE of order $\mathcal{O}(n^{-4/(H+4)})$.

Remark 6 (Curse of Horizon in Rate). Notice CDRK and CDRD have convergence *rate* that deteriorates as the horizon grows. In usual OPE for stochastic policies, Liu et al. (2018); Kallus and Uehara (2019a, 2020a) show that using cumulative IS leads to MSE with leading constant that grows exponentially in horizon, but it still has rate $\mathcal{O}(1/n)$. Thus, the curse of horizon is even more detrimental for deterministic policies. However, CDRK and CDRD technically work also for *non*-Markov decision processes. Next, we will tackle the curse in rate by leveraging Markovian structure, following Kallus and Uehara (2020a).

Motivated by DR OPE using marginal density ratios for stochastic policies (Kallus and Uehara, 2020a), we propose Marginal DR case \mathcal{K} (MDRK) for deterministic OPE: $\hat{J}_{\text{mdr}}^{\mathcal{K}} = \mathbb{P}_n[\phi_{\text{mdr}}^{\mathcal{K}}(\hat{q}_t^{\mathcal{K}}, \hat{w}_t^{\mathcal{K}}, \hat{\pi}_t^{\mathcal{b}})]$, where

$$\phi_{\text{mdr}}^{\mathcal{K}}(q_t^{\mathcal{K}}, w_t^{\mathcal{K}}, \pi_t^{\mathcal{b}}) = v_1^{\mathcal{K}}(s_1) + \sum_{t=1}^H w_t^{\mathcal{K}}(s_t) \pi_t^{\mathcal{b}}(a_t | s_t) K_h(a_t - \tau_t(s_t)) (r_t - q_t^{\mathcal{K}}(s_t, a_t) + v_{t+1}^{\mathcal{K}}(s_{t+1}))$$

and $w_t^{\mathcal{K}}(s_t) = p_{\pi^e, \mathcal{K}}(s_t)/p_{\pi^{\mathcal{b}}}(s_t)$ is the marginal density ratio associated with $\pi_t^{\pi^e, \mathcal{K}}$. Again, a similar estimator, Marginal DR case \mathcal{D} (MDRD), $\hat{J}_{\text{mdr}}^{\mathcal{D}}$, is constructed by replacing $q_t^{\mathcal{K}}(s_t, a_t), w_t^{\mathcal{K}}(s_t), \pi_t^{\mathcal{b}}(a_t | s_t)$ with $q_t^{\pi^e}(s_t, \tau_t(s_t)), w_t^{\pi^e}(s_t), \pi_t^{\mathcal{b}}(\tau_t(s_t) | s_t)$.

Theorem 4. Suppose for $j \leq H, i = 1, 2, \mathbb{E}[\|\hat{\pi}_j^{b, [i]} - \pi_j^{\mathcal{b}}\|_{1, \infty}] = o(1), \mathbb{E}[\|\hat{w}_j^{\mathcal{K}, [i]} - w_j^{\mathcal{K}}\|_{\infty}] = o(1), \mathbb{E}[\|\hat{q}_j^{\mathcal{K}, [i]} - q_j^{\mathcal{K}}\|_{1, \infty}] = o(1), \mathbb{E}[\max\{\|\hat{\pi}_j^{b, [i]} - \pi_j^{\mathcal{b}}\|_{\infty}, \|\hat{w}_j^{\mathcal{K}, [i]} - w_j^{\mathcal{K}}\|_{\infty}\} \|\hat{q}_j^{\mathcal{K}, [i]} - q_j^{\mathcal{K}}\|_{\infty}] =$

$\mathcal{O}(n^{-1/2}h^{-1/2})$, $nh^5 = \mathcal{O}(1)$, $nh \rightarrow \infty$. Then, the bias of $\hat{J}_{\text{mdr}}^{\mathcal{K}}$ is the same as $\hat{J}_{\text{cdr}}^{\mathcal{K}}$ in Theorem 3 and its variance is $\text{var}[\hat{J}_{\text{mdr}}^{\mathcal{K}}] = \frac{\Omega_2(k)}{nh} V_H^{\text{mdr}} + \mathcal{O}(n^{-1}h^{-1})$, where

$$V_H^{\text{mdr}} = \sum_{t=1}^H \mathbb{E}_{p_{\pi}^e} \left[\frac{w_t^{\pi^e}(s_t)}{\pi_t^b(\tau_t(s_t)|s_t)} \text{var}[r_t + v_{t+1}^{\pi^e}(s_{t+1}) \mid s_t, \tau_t(s_t)] \right].$$

Setting $h = \Theta(n^{-1/5})$ yields the minimal MSE of order $\mathcal{O}(n^{-4/5})$. Specifically, if $h = cn^{-1/5}$, then

$$\mathbb{E}[(\hat{J}_{\text{mdr}}^{\mathcal{K}} - J)^2] \leq n^{-4/5} R_{\max}^2 H^2 \left\{ \frac{c^4 M_2^2(k)}{4} (G_1^{(2)} + \frac{(H-1)}{2} G_2^{(2)})^2 + \frac{C_1 C_2 \Omega_2(k)}{c} \right\} + \mathcal{O}(n^{-4/5}).$$

Notice the minimal MSE rate is the same as in the bandit case. We therefore conjecture the rate to be minimax optimal. More crucially, it does not suffer from the curse of horizon in rate. Moreover, the dependence of the leading constant is polynomial in horizon, $\mathcal{O}(H^4)$. The leading constant is smaller when $C_1, C_2, G_1^{(2)}, G_2^{(2)}$ are smaller, i.e., when the behavior policy is closer to the evaluation policy and the reward and transition densities are smoother.

4.2 Off-Policy Deterministic Policy Gradient Estimation

We next construct deterministic policy gradient estimators for RL. By differentiating $\phi_{\text{mdr}}^{\mathcal{K}}$, we obtain the Marginal PG case \mathcal{K} (MPGK) estimator, $\hat{Z}_{\text{mpg}}^{\mathcal{K}} = \mathbb{P}_n[\psi_{\text{mpg}}^{\mathcal{K}}(\hat{q}_t^{\mathcal{K}}, \hat{w}_t^{\mathcal{K}}, \hat{d}_t^{q^{\mathcal{K}}}, \hat{d}_t^{w^{\mathcal{K}}}, \hat{\pi}_t^b)]$, where

$$\begin{aligned} \psi_{\text{mpg}}^{\mathcal{K}}(q_t^{\mathcal{K}}, w_t^{\mathcal{K}}, d_t^{q^{\mathcal{K}}}, d_t^{w^{\mathcal{K}}}, \pi_t^b) &= d_1^{v^{\mathcal{K}}} + \sum_{t=1}^H \frac{d_t^{v^{\mathcal{K}}}}{\pi_t^b(a_t|s_t)} K_h(a_t - \tau_t(s_t)) (r_t - q_t^{\mathcal{K}} + v_{t+1}^{\mathcal{K}}) \\ &\quad + \sum_{t=1}^H \frac{w_t^{\mathcal{K}}}{\pi_t^b(a_t|s_t)} \left(K_h^{(1)}(a_t - \tau_t(s_t)) (r_t - q_t^{\mathcal{K}} + v_{t+1}^{\mathcal{K}}) \nabla \tau_t(s_t) + K_h(a_t - \tau_t(s_t)) (-d_t^{q^{\mathcal{K}}} + d_{t+1}^{v^{\mathcal{K}}}) \right), \\ q_t^{\mathcal{K}} &= q_t^{\mathcal{K}}(s_t, a_t), d_t^{q^{\mathcal{K}}} = d_t^{q^{\mathcal{K}}}(s_t, a_t) = \nabla q_t^{\mathcal{K}}(s_t, a_t), d_t^{w^{\mathcal{K}}} = \nabla w_t^{\mathcal{K}}(s_t), d_t^{v^{\mathcal{K}}} = \nabla v_t^{\mathcal{K}}(s_t), \\ v_t^{\mathcal{K}}(s_t) &= \int q_t^{\mathcal{K}}(s_t, a_t) K_h(a_t - \tau_t(s_t)) da_t, \\ d_t^{v^{\mathcal{K}}}(s_t) &= \int \left\{ d_t^{q^{\mathcal{K}}}(s_t, a_t) K_h(a_t - \tau_t(s_t)) + q_t(s_t, a_t) K_h^{(1)}(a_t - \tau_t(s_t)) \nabla \tau_t(s_t) \right\} da_t. \end{aligned}$$

Notice we only estimate the nuisances $q_t^{\mathcal{K}}, w_t^{\mathcal{K}}, d_t^{q^{\mathcal{K}}}, d_t^{w^{\mathcal{K}}}, \pi_t^b$; then estimates for $v_t^{\mathcal{K}}, d_t^{v^{\mathcal{K}}}$ are defined in terms of these. The Marginal PG case \mathcal{D} (MPGD) estimator, $\hat{Z}_{\text{mpg}}^{\mathcal{D}}$, is similarly defined by replacing the nuisances in $\hat{Z}_{\text{mpg}}^{\mathcal{K}}$ by $q_t^{\pi^e}(s_t, \tau_t(s_t)), w_t^{\pi^e}(s_t), d_t^{q^{\pi^e}}(s_t, \tau_t(s_t)), d_t^{w^{\pi^e}}(s_t), \pi_t^b(\tau_t(s_t) \mid s_t)$. Again, note $\hat{Z}_{\text{mpg}}^{\mathcal{D}}$ involves a differentiation while $\hat{Z}_{\text{mpg}}^{\mathcal{K}}$ does not but involves an integration, as in Remark 5.

Theorem 5. Suppose for $j \leq H, i = 1, 2$, $\mathbb{E}[\|\hat{w}_j^{\mathcal{K}, [i]} - w_j^{\mathcal{K}}\|_{\infty}] = \mathcal{O}(1)$, $\mathbb{E}[\|\hat{\pi}_j^{b, [i]} - \pi_j^b\|_{2, \infty}] = \mathcal{O}(1)$, $\mathbb{E}[\|\hat{d}_j^{w^{\mathcal{K}, [i]}} - d_j^{w^{\mathcal{K}}}\|_{\infty}] = \mathcal{O}(1)$, $\mathbb{E}[\|\hat{q}_j^{q^{\mathcal{K}, [i]}} - q_j^{q^{\mathcal{K}}}\|_{1, \infty}] = \mathcal{O}(1)$, $\mathbb{E}[\max\{\|\hat{w}_j^{\mathcal{K}, [i]} - w_j^{\mathcal{K}}\|_{\infty}, \|\hat{\pi}_j^{b, [i]} - \pi_j^b\|_{1, \infty}, \|\hat{d}_j^{w^{\mathcal{K}, [i]}} - d_j^{w^{\mathcal{K}}}\|_{\infty}\} \max\{\|\hat{q}_j^{q^{\mathcal{K}, [i]}} - q_j^{q^{\mathcal{K}}}\|_{1, \infty}, \|\hat{d}_j^{q^{\mathcal{K}, [i]}} - d_j^{q^{\mathcal{K}}}\|_{\infty}\}] = \mathcal{O}(n^{-1/2}h^{-3/2})$, $nh^7 = \mathcal{O}(1)$, $nh \rightarrow \infty$. Then, we have $\mathbb{E}[\hat{Z}_{\text{mpg}}^{\mathcal{K}}] - Z = 0.5h^2 M_2(k) \tilde{B}_H^{\text{mpg}} + \mathcal{O}(n^{-1/2}h^{-3/2})$, $\text{var}[\hat{Z}_{\text{mpg}}^{\mathcal{K}}] = \frac{\Omega_2^{(1)}(k)}{nh^3} \tilde{V}_H^{\text{mpg}} + \mathcal{O}(n^{-1}h^{-3})$, where

$$\begin{aligned} \tilde{B}_H^{\text{mpg}} &= \sum_{t=1}^H \left(\nabla \mathbb{E}_{p_{\pi}^e} \left[\frac{r_t p_t^{(2)}(r_t | s_t, \tau_t(s_t))}{p_t(r_t | s_t, \tau_t(s_t))} \right] + \sum_{l=1}^{t-1} \nabla \mathbb{E}_{p_{\pi}^e} \left[\frac{r_t p_{l+1}^{(2)}(s_{l+1} | s_l, \tau_l(s_l))}{p_{l+1}(s_{l+1} | s_l, \tau_l(s_l))} \right] \right), \\ \tilde{V}_H^{\text{mpg}} &= \sum_{t=1}^H \mathbb{E}_{p_{\pi}^e} \left[\frac{w_t^{\pi^e}(s_t)}{\pi_t^b(\tau_t(s_t)|s_t)} \text{var}_{p_{\pi}^e} [r_t + q_{t+1}^{\pi^e}(s_{t+1}, \tau_{t+1}(s_{t+1})) \mid s_t, \tau_t(s_t)] \otimes \nabla \tau_t(s_t) \right]. \end{aligned}$$

Setting $h = \Theta(n^{-1/7})$ yields the minimal MSE of order $\mathcal{O}(n^{-4/7})$. Specifically, if $h = cn^{-1/7}$, the operator norm of the MSE is bounded by

$$\frac{R_{\max}^2 H^2 \Upsilon^2}{n^{4/7}} \left(\frac{c^4 M_2^2(k)}{4} \left\{ G_1^{(3)} + \frac{(H-1)}{2} \{G_1^{(2)} G_2^{(1)} + G_1^{(1)} G_2^{(2)} + G_2^{(3)}\} + \frac{(H-1)(H-2)}{3} G_2^{(1)} G_2^{(2)} \right\}^2 + \frac{C_1 C_2 \Omega_2^{(1)}(k)}{c^3} \right).$$

Note again that the MSE rate $\mathcal{O}(n^{-4/7})$ is slower than the usual $\mathcal{O}(1/n)$ efficient rate for off-policy gradient estimation with stochastic policies (Kallus and Uehara, 2020b). Nonetheless, it matches the bandit case and we therefore conjecture it is minimax optimal. More importantly, we note that it alleviates the curse of horizon, both in rate and in leading constant. We can also derive corresponding CPGK and CPGD estimators by differentiating the CDRK and CDRD estimating functions, but these will suffer from the curse of horizon in rate, as in Theorem 3.

Remark 7 (Estimation of nuisance functions). Our OPE and off-policy policy gradient estimators depend on estimating some nuisances. A unique and new feature of our estimators and analysis compared to previous deterministic off-policy estimators is that the MSE guarantees do not depend on the particular nuisance estimator used and we make no assumptions except for their (slow) convergence rate. The estimation of q_t, w_t for stochastic policies is discussed in Kallus and Uehara (2020a) and of d_t^q, d_t^w in Kallus and Uehara (2020b). These can be applied directly to estimate $q_t^K, w_t^K, d_t^{q^K}, d_t^{w^K}$ since the kernelized evaluation policy, $\pi^{e,K}$, is stochastic. The estimation of $q_t^{\pi^e}$ is the same for deterministic policies and a small adjustment can also be made for $d_t^{q^{\pi^e}}$ as we explain in Appendix B. The estimation $w_t^{\pi^e}, d_t^{w^{\pi^e}}$ for deterministic policies is difficult, but we can simply use estimates of $w_t^K, d_t^{w^K}$ as estimates for $w_t^{\pi^e}, d_t^{w^{\pi^e}}$, which is essentially a kernel density estimation approach for the densities in the latter. For additional detail, refer to Appendix B.

Remark 8 (Policy learning algorithms). To do offline RL to learn a deterministic policy, we can combine any type of gradient-based optimization algorithm with our estimated gradients. A simple gradient ascent is given as an example in Appendix C and used in the experiments in the next section. Following Kallus and Uehara (2020b) we can also combine standard results for gradient ascent with our error bounds to obtain a regret guarantee. Since the proof is exactly the same, simply plugging in our error bounds instead, we omit the details and refer the reader to Kallus and Uehara (2020b).

5 Experiments

We next conduct an experiment in a very simple environment to confirm the theoretical guarantees of the proposed estimators. More extensive experimentation remains future work. The setting is as follows. Set $\mathcal{S}_t = \mathbb{R}, \mathcal{A}_t = \mathbb{R}, s_0 = 0$. Then, set the transition dynamics as $s_t = a_{t-1} - s_{t-1} + \mathcal{N}(0, 0.3^2)$, the reward as $r_t = -s_t^2$, the behavior policy as $\pi^b(a | s) = \mathcal{N}(0.8s, 1.0^2)$, the deterministic evaluation policy as $\tau_t(s_t) = \theta s_t$, and the horizon as $H = 20$. Note that in this setting, the optimal policy is given by $\theta^* = 1$. We compare CPGK, CPGD, MP GK, MP GD using the Gaussian kernel with PG. The nuisance functions q, w, d^q, d^w (and their case \mathcal{K} equivalents) are estimated using polynomial sieve regressions (Chen, 2007). We assume the behavior policy is known. Since q is estimated by polynomials and k is Gaussian, we can compute the integrals in MP GK and CPGK analytically. We use the same estimated q in PG. We choose h by bootstrapping the estimator for each of $h \in \{0.05, 0.1, 0.25, 0.5\}$ and choosing that with smallest bootstrap variance.

First, in Fig. 1, we compare the MSE of gradient estimators at $\theta = 1.0$ over 100 replications for each of $n = 200, 400, 600, 800$. We find that the performance of MP GK is far superior to all other estimators in terms of MSE, which confirms our theoretical results. Interestingly, the performance of MP GD is slightly worse than CPGD. The possible reason is it is more difficult to estimate w than w^K . The reasonably good performance of CDGD and CDGK can be attributed to the known λ_t^D, λ_t^K , which ensures less sensitivity to the q -estimation due to the doubly robust error structure.

Second, in Fig. 2, we apply gradient ascent (see Appendix C) with $\alpha_t = 0.05, T = 50$, and $\hat{\theta}_1$ randomly chosen from $[0.8, 1.2]$. We only run the bootstrap for $\hat{\theta}_1$ and then keep the same h for the next iterations. We compare the regret of the final policy for the different policy gradient estimators, i.e., $J(\theta^*) - J(\hat{\theta}_{50})$, averaging over 100 replications of the experiment for each of $n = 200, 400, 600, 800$. Again, the performance of MP GK is superior to other estimators also in terms of regret.

6 Conclusion and Future work

We developed doubly robust versions of DPG and showed that they can circumvent issues of curse of horizon and of dependence on nuisances such as q -estimates. Theoretically, a next question may be showing the rates we obtain are minimax optimal by appealing to minimax theory for nonparametric density estimation (Korostelev, 2011). A more practical next step may be to apply this in larger-scale RL environments. Offline RL in large-scale environments is notoriously difficult (Fujimoto et al., 2019). We therefore expect it necessary to combine several heuristics, such as gradient updates to nuisance estimates and adaptive step sizes, to make the algorithm work well in practice.





Fig. 1: MSE of gradient estimation with 95% CF Fig. 2: Regret after gradient ascent with 95% CF

Acknowledgement

This material is based upon work supported by the National Science Foundation under Grant No. 1846210

Masatoshi Uehara was supported in part by MASASON Foundation.

Broader Impact

As a contribution to offline RL, our work is of particular importance for RL in the context of social and medical sciences, where experimentation is limited and observational data must be used. Deterministic policies are particularly important for real applications of offline RL because optimal policies are deterministic (up to ties). There exist various examples of the use of deterministic policies in epidemiology (Wu et al., 2018), political science (Imai and van Dyk, 2004), and economics (Colangelo and Lee, 2019). We provide some of the first theoretical results on convergence rates in sequential RL settings. Our work contributes both to the theoretical understanding of deterministic policies in RL and to practical methods for evaluating and learning these.

That said, it is well understood that there is generally a gap between theory and practice in RL as it is applied to very complex and large-scale systems, and valid convergence rate guarantees need not directly translate by themselves into practical success in real, large-scale settings. It is therefore important to keep in mind practical heuristics, stopgaps, and approximations from applied RL when translating this work into practice. The systematic investigation of the use of these in the context of deterministic-policy offline RL may require additional future work.

There are also several potential dangers to be cognizant of when applying offline RL tools in general. First, the presence of large unobserved confounding in the data (in our setting this would manifest as large violations of the MDP model) can bias evaluation and, unchecked, may potentially hide harms done by the policy being evaluated or the policy learned by off-policy optimization. Second, if the observational data is not representative of the population, that is, there is a covariate shift in the initial state distribution, then the evaluation both of value and of gradient will reflect these biases and correspondingly be unrepresentative, under-emphasizing value to some parts of the population and over-emphasizing value to others, leading to learned policies that possibly benefit over-represented subpopulations more than others. More generally, even without covariate shift, here we focused on evaluation and optimization of *average* welfare, which may average the harms to some and the benefits to others; it may therefore be important in some applications to also conduct auxiliary evaluations on certain protected subgroups to ensure equal impact, which can be done by segmenting the data.

References

- Antos, A., C. Szepesvári, and R. Munos (2008). Learning near-optimal policies with bellman-residual minimization based fitted policy iteration and a single sample path. *Machine Learning* 71, 89–129.
- Athey, S. and S. Wager (2017). Efficient policy learning. *arXiv preprint arXiv:1702.02896*.
- Bibaut, A. and M. J. van der Laan (2017). Data-adaptive smoothing for optimal-rate estimation of possibly non-regular parameters. *arXiv preprint arxiv:1706.07408*.

- Bibaut, A., I. Malenica, N. Vlassis, and M. Van Der Laan (2019). More efficient off-policy evaluation through regularized targeted learning. In *Proceedings of the 36th International Conference on Machine Learning*, Volume 97, pp. 654–663.
- Chen, X. (2007). Chapter 76 large sample sieve estimation of semi-nonparametric models. *Handbook of Econometrics* 6, 5549–5632.
- Chernozhukov, V., D. Chetverikov, M. Demirer, E. Duflo, C. Hansen, W. Newey, and J. Robins (2018). Double/debiased machine learning for treatment and structural parameters. *Econometrics Journal* 21, C1–C68.
- Colangelo, K. and Y.-Y. Lee (2019). Double debiased machine learning nonparametric inference with continuous treatments. *cemmap*.
- Farajtabar, M., Y. Chow, and M. Ghavamzadeh (2018). More robust doubly robust off-policy evaluation. In *Proceedings of the 35th International Conference on Machine Learning*, 1447–1456.
- Fong, C., C. Hazlett, and K. Imai (2018). Covariate balancing propensity score for a continuous treatment: Application to the efficacy of political advertisements. *Ann. Appl. Stat.* 12, 156–177.
- Foster, D. J. and V. Syrgkanis (2019). Orthogonal statistical learning. *arXiv preprint arXiv:1901.09036*.
- Fujimoto, S., D. Meger, and D. Precup (2019). Off-policy deep reinforcement learning without exploration. In *Proceedings of the 36th International Conference on Machine Learning*, Volume 97, pp. 2052–2062.
- Galvao, A. F. and L. Wang (2015). Uniformly semiparametric efficient estimation of treatment effects with a continuous treatment. 110, 1528–1542.
- Gottesman, O., F. Johansson, M. Komorowski, A. Faisal, D. Sontag, F. Doshi-Velez, and L. A. Celi (2019). Guidelines for reinforcement learning in healthcare. *Nat Med* 25, 16–18.
- Hansen, B. (2009). Lecture notes on nonparametrics. Technical report, University of Wisconsin.
- Hernan, M. and J. Robins (2019). *Causal Inference*. Boca Raton: Chapman & Hall/CRC.
- Hirano, K. and G. Imbens (2005). The propensity score with continuous treatments. In *Wiley Series in Probability and Statistics*, pp. 73–84. Chichester, UK: John Wiley & Sons, Ltd.
- Hsu, Y.-C., M. Huber, Y.-Y. Lee, and L. Pipoz (2018). Direct and indirect effects of continuous treatments based on generalized propensity score weighting. *IDEAS Working Paper Series from RePEc*.
- Huang, J. and N. Jiang (2019). From importance sampling to doubly robust policy gradient. *ICML 2020 (To appear)*.
- Imai, K. and D. A. van Dyk (2004). Causal inference with general treatment regimes: Generalizing the propensity score. *Journal of the American Statistical Association* 99, 854–866.
- Jiang, N. and L. Li (2016). Doubly robust off-policy value evaluation for reinforcement learning. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning-Volume*, 652–661.
- Kallus, N. and M. Uehara (2019a). Efficiently breaking the curse of horizon: Double reinforcement learning in infinite-horizon processes. *arXiv preprint arXiv:1909.05850*.
- Kallus, N. and M. Uehara (2019b). Intrinsically efficient, stable, and bounded off-policy evaluation for reinforcement learning. In *Advances in Neural Information Processing Systems* 32, pp. 3320–3329.
- Kallus, N. and M. Uehara (2020a). Double reinforcement learning for efficient off-policy evaluation in markov decision processes. *Journal of Machine Learning Research* 21, 1–63.
- Kallus, N. and M. Uehara (2020b). Statistically efficient off-policy policy gradients. *ICML 2020 (To appear)*.

- Kallus, N. and A. Zhou (2018). Policy evaluation and optimization with continuous treatments. In *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*, Volume 84, pp. 1243–1251.
- Kennedy, E. H., Z. Ma, M. D. McHugh, and D. S. Small (2017). Non-parametric methods for doubly robust estimation of continuous treatment effects. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 79, 1229–1245.
- Korostelev, Alexander; Korosteleva, O. (2011). *Mathematical Statistics : Asymptotic Minimax Theory*. Providence: American Mathematical Society.
- Le, H., C. Voloshin, and Y. Yue (2019). Batch policy learning under constraints. In *Proceedings of the 36th International Conference on Machine Learning*, pp. 3703–3712.
- Lee, Y.-Y. (2018). Partial mean processes with generated regressors: Continuous treatment effects and nonseparable models. *arXiv preprint arXiv: 1811.00157*.
- Li, Q. and J. S. Racine (2007). *Nonparametric econometrics : theory and practice*. Princeton, N.J.: Princeton University Press.
- Liu, Q., L. Li, Z. Tang, and D. Zhou (2018). Breaking the curse of horizon: Infinite-horizon off-policy estimation. In *Advances in Neural Information Processing Systems 31*, pp. 5356–5366.
- Liu, Y., O. Gottesman, A. Raghu, M. Komorowski, A. A. Faisal, F. Doshi-Velez, and E. Brunskill (2018). Representation balancing mdps for off-policy policy evaluation. In *Advances in Neural Information Processing Systems 31*, pp. 2644–2653.
- Munos, R. and C. Szepesvári (2008). Finite-time bounds for fitted value iteration. *Journal of Machine Learning Research* 9, 815–857.
- Murphy, S. A. (2003). Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 65, 331–355.
- Peters, J. and S. Schaal (2006). Policy gradient methods for robotics. In *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 2219–2225. IEEE.
- Precup, D., R. S. Sutton, and S. P. Singh (2000). Eligibility Traces for Off-Policy Policy Evaluation. In *Proceedings of the 17th International Conference on Machine Learning*, pp. 759–766.
- Silver, D., G. Lever, N. Heess, T. Degris, D. Wierstra, and M. Riedmiller (2014). Deterministic policy gradient algorithms. In *Proceedings of the 31st International Conference on Machine Learning*, pp. 387–395.
- Su, L., T. Ura, and Y. Zhang (2019). Non-separable models with high-dimensional data. *Journal of Econometrics* 212, 646–677.
- Thomas, P. and E. Brunskill (2016). Data-efficient off-policy policy evaluation for reinforcement learning. In *Proceedings of the 33rd International Conference on Machine Learning*, 2139–2148.
- Wu, X., F. Mealli, M.-A. Kioumourtoglou, F. Dominici, and D. Braun (2018). Matching on generalized propensity scores with continuous exposures. *arXiv preprint arXiv:1812.06575*.
- Xie, T., Y. Ma, and Y.-X. Wang (2019). Towards optimal off-policy evaluation for reinforcement learning with marginalized importance sampling. In *Advances in Neural Information Processing Systems 32*, pp. 9665–9675.
- Yin, M. and Y.-X. Wang (2020). Asymptotically efficient off-policy evaluation for tabular reinforcement learning. In *Proceedings of the 23rd International Workshop on Artificial Intelligence and Statistics (To appear)*.

A Notation

Table 2: Notation

$p_1(s)$	Initial distributions
$p_\pi(\cdot)$	Induced distribution by a MDP and a policy π
	$p_\pi(j_{r_{H+1}}) = p_1(s_1) \prod_{t=1}^H \pi(a_t s_t) p(r_t s_t, a_t) p(s_{t+1} s_t, a_t).$
$p_\tau(\cdot)$	$p_\pi(j_{r_{H+1}}) = p_1(s_1) \prod_{t=1}^H p(r_t s_t, \tau_t(s_t)) p(s_{t+1} s_t, \tau_t(s_t))$
$p_\kappa(\cdot)$	$p_\pi(j_{r_{H+1}}) = p_1(s_1) \prod_{t=1}^H K_h(a_t - \tau_t(s_t)) p(r_t s_t, a_t) p(s_{t+1} s_t, a_t)$
H	Horizon
h	Bandwidth
$\mathcal{T}^{(i)}$	i -th data
$\pi^b, \pi^e, \pi^{e, \kappa}$	Behavior policy, Evaluation policy, Kernelized policy
$J(\theta), Z(\theta)$	Value, Gradient
$k(\cdot), K_h(x)$	Kernel, Normalized kernel $h^{-1}k(h^{-1}x)$
τ, τ_θ	Deterministic policy with a parameter $\theta \in \Theta$
$\mathbb{E}[\cdot]$	Expectation wrt random variable generated by MDP and a behavior policy
$\mathbb{E}_{\pi^e}[\cdot]$	Expectation wrt random variable generated by a MDP and a policy π^e
$\mathbb{E}_{\pi^e}[\tau]$	Expectation wrt random variable generated by a MDP and a policy $\delta(a = \tau_t(s))$
$\mathcal{H}_{S_t}, \mathcal{H}_{A_t}$	History $(S_0, A_0, S_1, A_1, \dots, S_t), (S_0, A_0, S_1, A_1, \dots, A_t)$
$h_{s_t}^u$	History $(s_0, u_0, s_1, u_1, \dots, s_t)$
$M_2(k)$	Second moment of kernel, $\int u^2 k(u) du$
$\Omega^{(i)}(k)$	Roughness of kernel, $\int k^{2(i)}(u) du$
v^κ, q^κ	Value, Q-function wrt a kernelized policy and a MDP
v^{π^e}, q^{π^e}	Value, Q-function wrt a deterministic policy and a MDP
$q^{(i)}(s, a)$	i -th Differentiation of $q(s, a)$ wrt actions
$w_t^\kappa(s_t), w_t^{\pi^e}(s_t)$	$p_\kappa(s_t)/p_{\pi^b}(s_t), p_\tau(s_t)/p_{\pi^b}(s_t)$
Υ	$\ \otimes \tau_t(S)\ _\infty \ _{\text{op}} \leq \Upsilon$
R_{\max}	$\ R_t\ _\infty \leq R_{\max}$
C_1	$\ 1/\pi^b(a s)\ \leq C_1$
C_2	$\ w_t^{\pi^e}(s)\ \leq C_2$
∇	Differentiation wrt θ
$\ \cdot\ _2$	L^2 -norm $\{\mathbb{E}[f^2]\}^{1/2}$
$\ \cdot\ _{\text{op}}$	Operator norm
$\ \cdot\ _{i, \infty}$	$\max_{j=0, \dots, i} \ \cdot^{(j)}\ _\infty$
$d_t^{w^\kappa}(s_t), d_t^{w^{\pi^e}}(s_t)$	$\nabla w_t^\kappa(s_t), \nabla w_t^{\pi^e}(s_t)$
$d_t^{q^\kappa}(s_t, a_t), d_t^{q^{\pi^e}}(s_t, a_t)$	$\nabla q_t^{\pi^e}(s_t, a_t)$
$G_1^{(i)}$	$\ \int p^{(i)}(r s, a) dr\ _\infty \leq G_1^{(i)}$
$G_2^{(i)}$	$\ \int p^{(i)}(s' s, a) ds'\ _\infty \leq G_2^{(i)}$
$\otimes a$	aa^\top
$f(n) = \mathcal{O}(n^a)$	$f(n)$ is bounded above by n^a asymptotically
$f(n) = \Theta(n^a)$	$f(n)$ is bounded both above and below by n^a asymptotically
$f(n) = \mathcal{O}_p(n^a)$	$f(n)/n^a$ is bounded in probability
$f(n) = \mathcal{o}_p(n^a)$	$f(n)/n^a$ converges to 0 in probability
$A_n \lesssim B_n$	$\exists C, A_n < CB_n, C$ is a universal problem independent constant

B Nuisance Estimations

Our algorithm allows any estimators for q -functions and marginal ratios to be used. In this section we discuss some standard ways to estimate these nuisance functions.

B.1 Estimation of q -functions and their θ -gradients

In the tabular case, a model-based approach is the most common way to estimate q -functions of π^e from off-policy data. In the non-tabular case, we have to rely on some function approximation. The key equation to derive these methods is the Bellman equation:

$$q_t(s_t, a_t) = \mathbb{E}[r_t + q_{t+1}(s_{t+1}, \pi^e) \mid s_t, a_t].$$

There is also an equivalent equation for d^q . If π^e is stochastic we may use

$$d_j^q(s_j, a_j) = \mathbb{E}[d_{j+1}^v(s_{j+1}) \mid s_j, a_j], \quad d_j^v(s_j) = \mathbb{E}_{\pi^e}[d_j^q + g_j q_j \mid s_j],$$

where $q_t(s_t, \pi) = \int q_t(s_t, a_t) \pi(a_t \mid s_t) da_t$ and $g_t = \log \pi_t^e(a_t \mid s_t)$. When π^e is deterministic, we instead have

$$d_j^q(s_j, a_j) = \mathbb{E}[d_{j+1}^v(s_{j+1}) \mid s_j, a_j], \quad d_j^v(s_j) = d_j^q(s_j, \tau_j(s_j)) + q_j^{(1)}(s_j, \tau_j(s_j)) \nabla \tau_j(s_j).$$

One of the most common ways to operationalize this is using fitted q -iteration (Antos et al., 2008; Le et al., 2019):

- Set $\hat{q}_{H+1} \equiv 0$.
- For $t = H, \dots, 1$:
 - Estimate \hat{q}_t by regressing $r_t + \hat{q}_{t+1}(s_{t+1}, \pi^e)$ onto s_t, a_t .

Similarly, Kallus and Uehara (2020b) proposed an analogous estimation method for d_j^q :

- Set $\hat{d}_H^q = 0$.
- For $t = H, \dots, 1$:
 - Estimate \hat{d}_j^q by regressing $\hat{d}_{j+1}^v(s_{j+1})$ onto s_j, a_j .

The above approaches can be regarded as a dynamic programming approach. When π^e is stochastic, another approach is a Monte Carlo approach based on the equations:

$$q_j(s_j, a_j) = \mathbb{E} \left[\sum_{t=j}^H r_t \mid s_j, a_j \right], \quad d_j^q(s_j, a_j) = \mathbb{E} \left[\sum_{t=j+1}^H r_t \lambda_{j+1:t} \sum_{\ell=j+1}^t g_\ell \mid s_j, a_j \right].$$

Based on this, we can simply estimate q by regressing $\sum_{t=j}^H r_t$ on s_j, a_j and d^q by regressing $\sum_{t=j+1}^H r_t \lambda_{j+1:t} \sum_{\ell=j+1}^t g_\ell$ on a_j, s_j .

B.2 Estimation of Marginal Density Ratios

When \mathcal{S} is finite, a model based approach (Yin and Wang, 2020) would be a competitive way to estimate marginal density ratios:

$$\hat{w}_t^{\pi^e}(s_t) = \frac{1}{\hat{p}_{\pi_t^b}(s_t)} \int \hat{p}_t(s_t \mid s_{t-1}, \tau(s_{t-1})) \prod_{k=0}^{t-1} (\hat{p}_k(s_k \mid s_{k-1}, \tau(s_{k-1}))) d(\mathcal{H}_{a_{t-1}}),$$

where $\hat{p}_t, \hat{p}_{\pi_t^b}$ is each an empirical frequency (histogram) estimator. For general state space, we have to rely on some function approximation methods. When the target policy is stochastic, we have the following equations: $j \leq H$,

$$w_j(s_j) = \mathbb{E}[\lambda_{j-1} \mid s_j], \quad d_j^w = \mathbb{E} \left[\lambda_{0:j-1} \sum_{\ell=0}^{j-1} g_\ell \mid s_j \right].$$

Thus, for example, w_j is estimated by regressing λ_{j-1} onto s_j , and d_j^w is estimated by regressing $\lambda_{0:j-1} \sum_{\ell=0}^{j-1} g_\ell$ onto s_j (Kallus and Uehara, 2020a,b). When the evaluation policy is deterministic, it is difficult to estimate directly.

C Off-policy Optimization

Our estimated policy gradients can be used in any gradient-based optimization algorithm in order to do off-policy optimization to learn a policy. One example, which we also use in our experiment, is the simple gradient ascent algorithm, given below in Algorithm 1. Here, Proj_Θ is the projection onto Θ .

Algorithm 1 Off-policy projected gradient ascent

Input: An initial point $\theta_1 \in \Theta$ and step size schedule α_t
for $t = 1, 2, \dots$ **do**
 $\tilde{\theta}_{t+1} = \theta_t + \alpha_t \hat{Z}(\theta_t)$
 $\theta_{t+1} = \text{Proj}_\Theta(\tilde{\theta}_{t+1})$
end for

D Proof

We often drop a time index j for simplicity, like $p(s_j|s_{j-1}, a_{j-1}), p(r_j|s_j, a_j)$ instead of $p_j(s_j|s_{j-1}, a_{j-1}), p_j(r_j|s_j, a_j)$ for the ease of the notation.

D.1 Proof of Theorem 1

We prove Theorem 1. Before that, we prove Theorem 6.

Theorem 6. *Let us define*

$$\phi_1(s, a, r; q, \pi^b) = \frac{K_h(a - \tau(s))\{r - q(s, a)\}}{\pi^b(a|s)} + v(s), v(s) = \int K_h(a - \tau(s))q(s, a)d(a).$$

In the case \mathcal{K} , assuming (b), we have

$$\mathbb{E}[A_n] = o(1/\sqrt{nh}), \text{var}[A_n] = o(1/(nh)) \quad (8)$$

$$A_n = \mathbb{P}_{\mathcal{U}_1}[\phi_1(S, A, R; \hat{q}^{[1]}, \hat{\pi}^{b[1]})] + \mathbb{P}_{\mathcal{U}_2}[\phi_1(S, A, R; \hat{q}^{[2]}, \hat{\pi}^{b[2]})] - \mathbb{P}_n[\phi_1(S, A, R; q, \pi^b)].$$

In the case \mathcal{D} , the above (8) also holds assuming (a).

D.1.1 Proof of Theorem 6, Case \mathcal{K}

What we will prove is

$$\mathbb{E}[\phi_1(S, A, R; \hat{q}^{[1]}, \hat{\pi}^{b[1]}) - \phi_1(S, A, R; q, \pi^b) | \mathcal{U}_2] = o((nh)^{-1/2}), \quad (9)$$

$$\mathbb{E}[\{\phi_1(S, A, R; \hat{q}^{[1]}, \hat{\pi}^{b[1]}) - \phi_1(S, A, R; q, \pi^b)\}^2 | \mathcal{U}_2] = o(h^{-1}). \quad (10)$$

Then, the same argument holds for $\phi_1(S, A, R; \hat{q}^{[2]}, \hat{\pi}^{b[2]}) - \phi_1(S, A, R; q, \pi^b)$. The desired statement is concluded since

$$\begin{aligned} & \mathbb{E}[\mathbb{P}_{\mathcal{U}_1}[\phi_1(S, A, R; \hat{q}^{[1]}, \hat{\pi}^{b[1]})] - \mathbb{P}_{\mathcal{U}_1}[\phi_1(S, A, R; q, \pi^b)]] \\ &= \mathbb{E}[\mathbb{E}[\mathbb{P}_{\mathcal{U}_1}[\phi_1(S, A, R; \hat{q}^{[1]}, \hat{\pi}^{b[1]})] - \mathbb{P}_{\mathcal{U}_1}[\phi_1(S, A, R; q, \pi^b)] \mid \mathcal{U}_2]], \\ &= \mathbb{E}[\mathbb{E}[\phi_1(S, A, R; \hat{q}, \hat{\pi}^b) - \phi_1(S, A, R; q, \pi^b) \mid \mathcal{U}_2]]. \end{aligned}$$

In addition,

$$\begin{aligned} & \mathbb{E}[\{\mathbb{P}_{\mathcal{U}_1}[\phi_1(S, A, R; \hat{q}^{[1]}, \hat{\pi}^{b[1]})] - \mathbb{P}_{\mathcal{U}_1}[\phi_1(S, A, R; q, \pi^b)]\}^2] \\ &= \mathbb{E}[\mathbb{E}[\{\mathbb{P}_{\mathcal{U}_1}[\phi_1(S, A, R; \hat{q}^{[1]}, \hat{\pi}^{b[1]})] - \mathbb{P}_{\mathcal{U}_1}[\phi_1(S, A, R; q, \pi^b)]\}^2 \mid \mathcal{U}_2]], \\ &= n^{-1} K \mathbb{E}[\mathbb{E}[\{\phi_1(S, A, R; \hat{q}^{[1]}, \hat{\pi}^{b[1]}) - \phi_1(S, A, R; q, \pi^b)\}^2 \mid \mathcal{U}_2]]. \end{aligned}$$

Proof of Eq. (9) and Eq. (10) In this subsection, we remove $\{[1]\}$ for the ease of the notation. To prove (9), we show

$$\mathbb{E} \left[K_h(A - \tau(S)) \left\{ \frac{1}{\hat{\pi}^b(A|S)} - \frac{1}{\pi^b(A|S)} \right\} \{q(S, A) - \hat{q}(S, A)\} \right] = o(nh)^{-1/2}.$$

This is proved by

$$\begin{aligned} & \mathbb{E} \left[\mathbb{E} \left[K_h(A - \tau(S)) \left\{ \frac{1}{\hat{\pi}^b(A|S)} - \frac{1}{\pi^b(A|S)} \right\} \{q(S, A) - \hat{q}(S, A)\} \mid \mathcal{U}_2 \right] \right] \\ &= \mathbb{E} \left[\int \frac{1}{h} k((a - \tau(s))h^{-1}) \left\{ \frac{1}{\hat{\pi}^b(a|s)} - \frac{1}{\pi^b(a|s)} \right\} \{q(s, a) - \hat{q}(s, a)\} \pi^b(a | s) p(s) d(s, a) \right] \\ &= \mathbb{E} \left[\int k(u) \left\{ \frac{1}{\hat{\pi}^b(\tau(s) + uh|s)} - \frac{1}{\pi^b(\tau(s) + uh|s)} \right\} \{q(s, \tau(s) + uh) - \hat{q}(s, \tau(s) + uh)\} \pi^b(\tau(s) + uh | s) p(s) d(u, s) \right] \\ &= \mathbb{E} \left[\int k(u) \left\{ \frac{1}{\hat{\pi}^b(\tau(s)|s)} - \frac{1}{\pi^b(\tau(s)|s)} + uh \left\{ -\frac{\hat{\pi}^{b(1)}(\tau(s)|s)}{\hat{\pi}^{2b}(\tau(s)|s)} + \frac{\pi^{b(1)}(\tau(s)|s)}{\pi^{2b}(\tau(s)|s)} \right\} + \mathcal{O}((uh^2)) \right\} \times \right. \\ & \quad \{q(s, \tau(s)) - \hat{q}(s, \tau(s)) + uh\{q^{(1)}(s, \tau(s)) - \hat{q}^{(1)}(s, \tau(s))\} + \mathcal{O}((uh^2))\} \times \\ & \quad \left. \{\pi^b(\tau(s) | s) + uh\pi^{b(1)}(\tau(s) | s) + \mathcal{O}((uh^2))\} p(s) d(u, s) \right] \\ &= M_2(k) \mathbb{E} \left[\int \left\{ \frac{1}{\hat{\pi}^b(\tau(s) | s)} - \frac{1}{\pi^b(\tau(s) | s)} \right\} \{q(s, \tau(s)) - \hat{q}(s, \tau(s))\} \pi^b(\tau(s) | s) p(s) d(s) \right] + \mathcal{O}(h^2) \times o(1) \\ &= o((nh)^{-1/2}). \end{aligned}$$

In the last line, we use the assumptions that $q(a, x)$, $\pi^b(a|x)$, $\hat{q}(a, x)$, $\hat{\pi}^b(a|x)$ are C^2 -functions wrt actions, and

$$\begin{aligned} & \mathbb{E} \left[\left\| \frac{1}{\hat{\pi}^b(A | S)} - \frac{1}{\pi^b(A | S)} \right\|_\infty \|q(S, A) - \hat{q}(S, A)\|_\infty \right] = o(n^{-1/2}h^{-1/2}), \\ & \mathbb{E} \left[\left\| \frac{1}{\hat{\pi}^b(A | S)} - \frac{1}{\pi^b(A | S)} \right\|_{1,\infty} \right] = o(1), \mathbb{E} [\|\hat{q}(S, A) - \hat{q}(S, A)\|_{1,\infty}] = o(1), \\ & \mathcal{O}(h^2) \times o(1) = o((nh)^{-1/2}), nh^5 = \mathcal{O}(1). \end{aligned}$$

In addition, Eq. (10) is proved since

$$\begin{aligned} & \mathbb{E}[\mathbb{E}[\{\phi_1(S, A, R; \hat{q}^{[1]}, \hat{\pi}^{b[1]}) - \phi_1(S, A, R; q, \pi^b)\}^2 \mid \mathcal{U}_2]] \\ & \lesssim \mathbb{E} \left[\mathbb{E} \left[K_h(A - \tau(S))^2 \left\{ \frac{1}{\hat{\pi}^b(A | S)} - \frac{1}{\pi^b(A | S)} \right\}^2 \{q(S, A) - \hat{q}(S, A)\}^2 \mid \mathcal{U}_2 \right] \right] \\ & + \mathbb{E} \left[\mathbb{E} \left[K_h(A - \tau(S))^2 \left\{ \frac{1}{\hat{\pi}^b(A | S)} - \frac{1}{\pi^b(A | S)} \right\}^2 \{R - q(S, A)\}^2 \mid \mathcal{U}_2 \right] \right] \\ & + \mathbb{E} \left[\mathbb{E} \left[\frac{K_h(A - \tau(S))^2}{\pi^b(A | S)^2} \{q(S, A) - \hat{q}(S, A)\}^2 \mid \mathcal{U}_2 \right] + \mathbb{E} [\{\hat{v}(S) - v(S)\}^2 \mid \mathcal{U}_2] \right] \\ & \lesssim h^{-1} \max \left\{ \mathbb{E}[\|\hat{q}(S, \tau(S)) - q(S, \tau(S))\|_2^2], \mathbb{E} \left[\left\| \frac{1}{\hat{\pi}^b(\tau(S)|S)} - \frac{1}{\pi^b(\tau(S)|S)} \right\|_2^2 \right] \right\} + \mathcal{O}(1) = o(h^{-1}). \end{aligned}$$

D.1.2 Proof of Theorem 6, Case \mathcal{D}

Essentially, the same proof is seen in Colangelo and Lee (2019). For completeness, we also write the proof here with our notation. Let us define

$$\phi_2(s, a, r; q, \pi^b) = \frac{K_h(a - \tau(s))\{r - q(s, \tau(s))\}}{\pi^b(a|s)} + q(S, \tau(S)).$$

As in the argument similar to the case \mathcal{K} , what we have to prove is

$$\mathbb{E}[\phi_2(S, A, R; \hat{q}^{[1]}, \hat{\pi}^{b[1]}) - \phi_2(S, A, R; q, \pi^b) | \mathcal{U}_2] = o((nh)^{-1/2}), \quad (11)$$

$$\mathbb{E}[\{\phi_2(S, A, R; \hat{q}^{[1]}, \hat{\pi}^{b[1]}) - \phi_2(S, A, R; q, \pi^b)\}^2 | \mathcal{U}_2] = o(h^{-1}). \quad (12)$$

In this subsection, we remove $\{[1]\}$ for the ease of the notation.

Eq. (11) is proved since

$$\begin{aligned} & \mathbb{E} [\mathbb{E}[\phi_2(S, A, R; \hat{q}, \hat{\pi}^b) - \phi_2(S, A, R; q, \pi^b) \mid \mathcal{U}_2]] \\ &= \mathbb{E} \left[\mathbb{E} \left[K_h(A - \tau(S)) \left\{ \frac{1}{\hat{\pi}^b(\tau(S) \mid S)} - \frac{1}{\pi^b(\tau(S) \mid S)} \right\} \{q(S, \tau(S)) - \hat{q}(S, \tau(S))\} \mid \mathcal{U}_2 \right] \right] + \end{aligned} \quad (13)$$

$$+ \mathbb{E} \left[\mathbb{E} \left[K_h(A - \tau(S)) \left\{ \frac{1}{\hat{\pi}^b(\tau(S) \mid S)} - \frac{1}{\pi^b(\tau(S) \mid S)} \right\} \{R - q(S, \tau(S))\} \mid \mathcal{U}_2 \right] \right] \quad (14)$$

$$+ \mathbb{E} \left[\mathbb{E} \left[\frac{K_h(A - \tau(S))}{\pi^b(\tau(S) \mid S)} \{q(S, \tau(S)) - \hat{q}(S, \tau(S))\} + \hat{q}(S, \tau(S)) - q(S, \tau(S)) \mid \mathcal{U}_2 \right] \right] \quad (15)$$

$$= \mathcal{O}((nh)^{-1/2}) + \mathcal{O}(1) \times \mathcal{O}(h^2) + \mathcal{O}(1) \times \mathcal{O}(h^2) = \mathcal{O}(nh)^{-1/2}.$$

Here, we use the facts that (13) is $\mathcal{O}((nh)^{-1/2})$, (14) is $\mathcal{O}(1) \times \mathcal{O}(h^2)$, (15) is $\mathcal{O}(1) \times \mathcal{O}(h^2)$, which we will prove soon. In the last line, we use $nn^5 = \mathcal{O}(1)$. From now on, we prove (14) is $\mathcal{O}(1) \times \mathcal{O}(h^2)$:

$$\begin{aligned} & \mathbb{E} \left[\mathbb{E} \left[K_h(A - \tau(S)) \left\{ \frac{1}{\hat{\pi}^b(\tau(S) \mid S)} - \frac{1}{\pi^b(\tau(S) \mid S)} \right\} \{R - q(S, \tau(S))\} \mid \mathcal{U}_2 \right] \right] \\ &= \mathbb{E} \left[\mathbb{E} \left[\left\{ \frac{1}{\hat{\pi}^b(\tau(S) \mid S)} - \frac{1}{\pi^b(\tau(S) \mid S)} \right\} \{\mathbb{E}[K_h(A - \tau(S))q(S, A) \mid S] - K_h(A - \tau(S))q(S, \tau(S))\} \mid \mathcal{U}_2 \right] \right] \\ &= \mathbb{E} \left[\mathbb{E} \left[\left\{ \frac{1}{\hat{\pi}^b(\tau(S) \mid S)} - \frac{1}{\pi^b(\tau(S) \mid S)} \right\} \{\mathcal{O}(h^2)\} \mid \mathcal{U}_2 \right] \right] \\ &= \mathcal{O}(1) \times \mathcal{O}(h^2). \end{aligned}$$

More specifically,

$$\begin{aligned} & \mathbb{E}[K_h(A - \tau(S))\{q(S, A) - q(S, \tau(S))\} \mid S] = \int \frac{1}{h} k \left\{ \frac{a - \tau(s)}{h} \right\} \pi^b(a \mid s) \{q(s, a) - q(s, \tau(s))\} da \\ &= \int k(u) \pi^b(\tau(s) + uh \mid s) \{q(s, \tau(s) + uh) - q(s, \tau(s))\} du \\ &= \int k(u) \{\pi^b(\tau(s) \mid s) + \mathcal{O}(uh)\} \{uhq^{(1)}(s, \tau(s)) + \mathcal{O}(h^2)\} du = \mathcal{O}(h^2). \end{aligned}$$

noting $\int uk(u)du = 0$. Next, we prove (15) is $\mathcal{O}(1) \times \mathcal{O}(h^2)$:

$$\begin{aligned} & \mathbb{E} \left[\mathbb{E} \left[\frac{K_h(A - \tau(S))}{\pi^b(\tau(S) \mid S)} \{q(S, \tau(S)) - \hat{q}(S, \tau(S))\} + \hat{q}(S, \tau(S)) - q(S, \tau(S)) \mid \mathcal{U}_2 \right] \right] \\ &= \mathbb{E} \left[\mathbb{E} \left[\left\{ \frac{K_h(A - \tau(S))}{\pi^b(\tau(S) \mid S)} - 1 \right\} \{q(S, \tau(S)) - \hat{q}(S, \tau(S))\} \mid \mathcal{U}_2 \right] \right] \\ &= \mathbb{E} [\mathbb{E} [\{\mathcal{O}(h^2)\} \{q(S, \tau(S)) - \hat{q}(S, \tau(S))\} \mid \mathcal{U}_2]] = \mathcal{O}(1) \times \mathcal{O}(h^2). \end{aligned}$$

Eq. (12) is similarly proved as in the case \mathcal{K} .

D.1.3 Proof of Theorem 1

We prove the statement for the case \mathcal{K} . The statement for the case \mathcal{D} is similarly proved as in (Colangelo and Lee, 2019).

Bias term The bias term is calculated as follows:

$$\begin{aligned} & \mathbb{E}[\mathbb{P}_n[\phi_1(S, A, R; q, \pi^b)]] - J \\ &= \mathbb{E} \left[\int K_h(a - \tau(S)) q(S, a) da \right] - J \\ &= \mathbb{E} \left[\int k(u) \{q(S, \tau(S) + uh) - q(S, \tau(S))\} du \right] \\ &= 0.5h^2 \mathbb{E} \left[\int k(u) u^2 q^{(2)}(S, \tau(S)) du \right] + \mathcal{O}(h^2) = 0.5h^2 M_2(k) \mathbb{E} [q^{(2)}(S, \tau(S))] + \mathcal{O}(h^2). \end{aligned}$$

Here, we use a smoothness assumption. More formally, from the third line to the fourth line, based on the function $a \rightarrow q(s, a)$ is a C^2 -function on the compact space, we use

$$q(s, \tau(s) + uh) - q(s, \tau(s)) = uhq^{(1)}(s, \tau(s)) + 0.5(uh)^2q^{(2)}(s, \tau(s)) + o((uh)^2).$$

Refer to Li and Racine (2007, Exercise 1.5). Then, the all of the bias is

$$\begin{aligned}\mathbb{E}[\mathbb{P}_n[\phi_1(S, A, R; q, \pi^b)]] - J &= \mathbb{E}[\mathbb{P}_n[\phi_1(S, A, R; q, \pi^b)]] - J + o(n^{-1/2}h^{-1/2}) \\ &= 0.5h^2M_2(k)\mathbb{E}\left[q^{(2)}(S, \tau(S))\right] + o(n^{-1/2}h^{-1/2}).\end{aligned}$$

Finally, the statement is concluded since

$$\begin{aligned}\mathbb{E}[\hat{J}^K] - J &= \mathbb{E}[\mathbb{P}_n[\phi_1(S, A, R; q, \pi^b)]] - J + o((nh)^{-1/2}) \\ &= 0.5h^2M_2(k)\mathbb{E}\left[q^{(2)}(S, \tau(S))\right] + o(n^{-1/2}h^{-1/2}).\end{aligned}$$

Variance term The variance term is calculated as follows. First, we have

$$\begin{aligned}\text{var}[\mathbb{P}_n[\phi_1(S, A, R; q, \pi^b)]] &= \frac{1}{n} (\mathbb{E}[\phi_1(S, A, R; q, \pi^b)^2] - \{\mathbb{E}[\phi_1(S, A, R; q, \pi^b)]\}^2) \\ &= \frac{1}{nh^2} \left[\int \left\{ \frac{k((a - \tau(s))h^{-1})}{\pi^b(a | s)} \right\}^2 \{r - q(s, a)\}^2 p(r | a, s) \pi^b(a | s) p(s) da, s, r \right. \\ &\quad \left. + \int \left\{ \int k((a - \tau(s))h^{-1}) q(s, a) da \right\}^2 p(s) ds + O(h^2) \right] \\ &= \frac{1}{nh^2} \left[\int h \left\{ \frac{k^2(u)}{\pi^b(\tau(s) + uh | s)} \right\} \{r - q(s, \tau(s) + uh)\}^2 p(r | s, \tau(s) + uh) p(s) du, s, r + O(h^2) \right] \\ &= \frac{1}{nh^2} \left[\int h \left\{ \frac{k^2(u)}{\pi^b(\tau(s) | s)} \right\} \{r - q(s, \tau(s))\}^2 p(r | s, \tau(s)) p(s) du, s, r + O(h^2) \right] \\ &= \frac{1}{nh} \{\Omega_2(k)V + o(h)\}, \quad V = \int \left\{ \frac{\{r - q(s, \tau(s))\}^2}{\pi^b(\tau(s) | s)} \right\} p(r | s, \tau(s)) p(s) ds, r.\end{aligned}$$

Here, we use smoothness assumptions, and $\left\{ \int k((a - \tau(s))h^{-1}) q(s, a) da \right\}^2 = O(h^2)$, which is proved by a standard algebra. Then,

$$\begin{aligned}\text{var}[\mathbb{P}_n[\phi_1(S, A, R; \hat{q}, \hat{\pi}^b)]] &= \text{var}[\mathbb{P}_n[\phi_1(S, A, R; q, \pi^b)]] + \text{var}[\mathbb{P}_n[\phi_1(S, A, R; q, \pi^b) - \phi_1(S, A, R; \hat{q}, \hat{\pi}^b)]] \\ &\quad + 2\{\text{var}[\mathbb{P}_n[\phi_1(S, A, R; q, \pi^b)]] \text{var}[\mathbb{P}_n[\phi_1(S, A, R; q, \pi^b) - \phi_1(S, A, R; \hat{q}, \hat{\pi}^b)]]\}^{1/2} \\ &= \frac{\Omega_2(k)}{nh} \{V + o(h)\} + 2\left\{ \frac{1}{nh} \{V + o(h)\} \right\}^{1/2} o(n^{-1/2}h^{-1/2}) + o(n^{-1}h^{-1}) \\ &= \frac{\Omega_2(k)}{nh} \{V + o(1)\}.\end{aligned}$$

Remark 9. Colangelo and Lee (2019, Theorem 1) showed that the constant in the bias term of case \mathcal{D} is

$$\mathbb{E}[0.5q^{(2)}(s, \tau(s)) + q^{(1)}(s, \tau(s))\pi^{b(1)}(a|s)/\pi^b(a|s)].$$

D.2 Proof of Corollary 1

Obvious from Theorem 1.

D.3 Proof of Theorem 2

Replacing parts Let us define

$$\psi_1^K(s, a, r; q, \pi^b) = \left\{ \frac{\nabla K_h(a - \tau(s)) \{r - q(s, a)\}}{\pi^b(a|s)} + \int \nabla K_h(a - \tau(s)) q(s, a) d(a) \right\}.$$

Here, we prove that nuisance estimators can be replaced with true functions in the sense that

$$\mathbb{E}[A_n] = o(n^{-1/2}h^{-3/2}), \quad \mathbb{E}[A_n^2] = o(n^{-1}h^{-3}),$$

where

$$A_n = \mathbb{P}_{\mathcal{U}_1}[\psi_1^K(S, A, R; \hat{q}^{[1]}, \hat{\pi}^{b[1]}) | \mathcal{U}_2] + \mathbb{P}_{\mathcal{U}_2}[\psi_1^K(S, A, R; \hat{q}^{[2]}, \hat{\pi}^{b[2]}) | \mathcal{U}_1] - \mathbb{P}_n[\psi_1^K(S, A, R; q, \pi^b)].$$

Then, what we have to prove is

$$\mathbb{E}[\mathbb{E}[\psi_1^K(S, A, R; \hat{q}^{[1]}, \hat{\pi}^{b[1]}) - \psi_1^K(S, A, R; q, \pi^b) | \mathcal{U}_2]] = o(n^{-1/2}h^{-3/2}), \quad (16)$$

$$\mathbb{E}[\mathbb{E}[\{\psi_1^K(S, A, R; \hat{q}^{[1]}, \hat{\pi}^{b[1]}) - \psi_1^K(S, A, R; q, \pi^b)\}^2 | \mathcal{U}_2]] = o(h^{-3}). \quad (17)$$

In this subsection, we remove $\{[1]\}$ for the ease of the notation. We write $1/\pi^b(A|S)$ as $\eta(A|S)$. Eq. (16) is proved as follows:

$$\mathbb{E}[\mathbb{E}[\psi_1^K(S, A, R; \hat{q}, \hat{\pi}^b) - \psi_1^K(S, A, R; q, \pi^b) | \mathcal{U}_2]] \quad (18)$$

$$\begin{aligned} &= \mathbb{E}[\mathbb{E}[\nabla K_h(A - \tau(S)) \{\hat{\eta}(A|S) - \eta(A|S)\} \{q(S, A) - \hat{q}(S, A)\} | \mathcal{U}_2]] \\ &= \mathbb{E} \left[\mathbb{E} \left[-h^{-2}k^{(1)} \left(\frac{A - \tau(S)}{h} \right) \nabla \tau(S) \{\hat{\eta}(A|S) - \eta(A|S)\} \{q(S, A) - \hat{q}(S, A)\} | \mathcal{U}_2 \right] \right] \\ &= \mathbb{E} \left[\mathbb{E} \left[-h^{-2}k^{(1)} \left(\frac{A - \tau(S)}{h} \right) \nabla \tau(S) \{\hat{\eta}(A|S) - \eta(A|S)\} \{q(S, A) - \hat{q}(S, A)\} | \mathcal{U}_2 \right] \right] \\ &= \mathbb{E}[\mathbb{E}[\int -h^{-1}k^{(1)}(u) \nabla \tau(S) \{\hat{\eta}(\tau(S) + uh|S) - \eta(\tau(S) + uh|S)\} \times \\ &\quad \{q(S, \tau(S) + uh) - \hat{q}(S, \tau(S) + uh)\} \pi^b(\tau(S) + uh|S) du | \mathcal{U}_2]] \quad (19) \end{aligned}$$

$$\begin{aligned} &= \mathbb{E}[\mathbb{E}[\int k(u) \nabla \tau(S) \{\hat{\eta}^{(1)}(\tau(S) + uh|S) - \eta^{(1)}(\tau(S) + uh|S)\} \times \\ &\quad \{q(S, \tau(S) + uh) - \hat{q}(S, \tau(S) + uh)\} \pi^b(\tau(S) + uh|S) du | \mathcal{U}_2]] \quad (20) \\ &+ \mathbb{E}[\mathbb{E}[\int k(u) \nabla \tau(S) \{\hat{\eta}(\tau(S) + uh|S) - \eta(\tau(S) + uh|S)\} \times \\ &\quad \{q^{(1)}(S, \tau(S) + uh) - \hat{q}^{(1)}(S, \tau(S) + uh)\} \pi^b(\tau(S) + uh|S) du | \mathcal{U}_2]] \\ &+ \mathbb{E}[\mathbb{E}[\int k(u) \nabla \tau(S) \{\hat{\eta}(\tau(S) + uh|S) - \eta(\tau(S) + uh|S)\} \times \\ &\quad \{q(S, \tau(S) + uh) - \hat{q}(S, \tau(S) + uh)\} \pi^{b(1)}(\tau(S) + uh|S) du | \mathcal{U}_2]] \end{aligned}$$

Then, this is equal to

$$M_2(k) \mathbb{E}[\mathbb{E}[\nabla \tau(S) \{\hat{\eta}^{(1)}(\tau(S)|S) - \eta^{(1)}(\tau(S)|S)\} \{q(S, \tau(S)) - \hat{q}(S, \tau(S))\} \pi^b(\tau(S)|S) | \mathcal{U}_2]] \quad (21)$$

$$\begin{aligned} &+ M_2(k) \mathbb{E}[\mathbb{E}[\nabla \tau(S) \{\hat{\eta}(\tau(S)|S) - \eta(\tau(S)|S)\} \{q^{(1)}(S, \tau(S)) - \hat{q}^{(1)}(S, \tau(S))\} \pi^b(\tau(S)|S) | \mathcal{U}_2]] \\ &+ M_2(k) \mathbb{E}[\mathbb{E}[\nabla \tau(S) \{\hat{\eta}(\tau(S)|S) - \eta(\tau(S)|S)\} \{q(S, \tau(S)) - \hat{q}(S, \tau(S))\} \pi^b(\tau(S)|S) | \mathcal{U}_2]] \\ &+ o(h^2) \\ &\lesssim \mathbb{E}[\|\nabla \tau(S)\|_{\text{op}} \|\hat{q}(S, A) - q(S, A)\|_{1,\infty} \|\hat{\pi}^b(S, A) - \pi^b(S, A)\|_{1,\infty}] + o(n^{-1/2}h^{-3/2}) \quad (22) \\ &= o(n^{-1/2}h^{-3/2}). \end{aligned}$$

Here, from (19) to (20), we have used a partial integration. From (20) to (21), we have used $a \rightarrow \eta(s, a)$ and $a \rightarrow q(s, a)$ are C^3 -functions, and

$$\begin{aligned} &\mathbb{E} \left[\left\{ \hat{\eta}^{(1)}(\tau(S) + uh|S) - \eta^{(1)}(\tau(S) + uh|S) \right\} \{q(S, \tau(S) + uh) - \hat{q}(S, \tau(S) + uh)\} | \mathcal{U}_2 \right] \\ &= \mathbb{E} \left[\left\{ \hat{\eta}^{(1)}(\tau(S)|S) - \eta^{(1)}(\tau(S)|S) + uh \{ \hat{\eta}^{(2)}(\tau(S)|S) - \eta^{(2)}(\tau(S)|S) \} + \mathcal{O}((uh)^2) \right\} \times \right. \\ &\quad \left. \{q(S, \tau(S)) - \hat{q}(S, \tau(S)) + uh \{ q^{(1)}(S, \tau(S)) - \hat{q}^{(1)}(S, \tau(S)) \} + \mathcal{O}((uh)^2) \} | \mathcal{U}_2 \right] \\ &\lesssim \|\hat{\eta}(A|S) - \eta(A|S)\|_{1,\infty} \times \|\hat{q}(A|S) - q(A|S)\|_{1,\infty} \\ &+ \{\|\hat{\eta}(A|S) - \eta(A|S)\|_{2,\infty} \|\hat{q}(A|S) - q(A|S)\|_{2,\infty} + \|\hat{\eta}(A|S) - \eta(A|S)\|_{1,\infty} + \|\hat{q}(A|S) - q(A|S)\|_{1,\infty}\} \times \mathcal{O}(h^2). \end{aligned}$$

From (21) to (22), we use an assumption $nh^7 = \mathcal{O}(1)$.

Eq. (17) is proved as follows:

$$\begin{aligned}
& \mathbb{E} \left[\mathbb{E}[\{\psi_1^{\mathcal{K}}(S, A, R; \hat{q}, \hat{\eta}) - \psi_1^{\mathcal{K}}(S, A, R; q, \eta)\}^2 \mid \mathcal{U}_2] \right] \\
& \lesssim \mathbb{E} \left[\mathbb{E} \left[h^{-2} K_h^{(1)}(A - \tau(S))^2 \left\{ \frac{1}{\hat{\pi}^b(A \mid S)} - \frac{1}{\pi^b(A \mid S)} \right\}^2 \{q(S, A) - \hat{q}(S, A)\}^2 \{\nabla \tau(S)\}^2 \mid \mathcal{U}_2 \right] \right] \\
& + \mathbb{E} \left[\mathbb{E} \left[h^{-2} K_h^{(1)}(A - \tau(S))^2 \left\{ \frac{1}{\hat{\pi}^b(A \mid S)} - \frac{1}{\pi^b(A \mid S)} \right\}^2 \{R - q(S, A)\}^2 \{\nabla \tau(S)\}^2 \mid \mathcal{U}_2 \right] \right] \\
& + \mathbb{E} \left[\mathbb{E} \left[h^{-2} \frac{K_h^{(1)}(A - \tau(S))^2}{\pi^b(A \mid S)^2} \{q(S, A) - \hat{q}(S, A)\}^2 \{\nabla \tau(S)\}^2 \mid \mathcal{U}_2 \right] \right] \\
& + \mathbb{E} \left[\mathbb{E} \left[h^{-2} \left\{ \int K_h^{(1)}(a - \tau(S)) \hat{q}(S, a) da - \int K_h^{(1)}(a - \tau(S)) q(S, a) da \right\}^2 \{\nabla \tau(S)\}^2 \mid \mathcal{U}_2 \right] \right] \\
& \lesssim h^{-3} \times \mathbb{E}[\max\{\|\hat{\eta}(S, \tau(S)) - \eta(S, \tau(S))\|_2^2, \|\hat{\eta}^{(1)}(S, \tau(S)) - \eta^{(1)}(S, \tau(S))\|_2^2, \\
& \|\hat{q}(S, \tau(S)) - q(S, \tau(S))\|_2^2, \|\hat{q}^{(1)}(S, \tau(S)) - q^{(1)}(S, \tau(S))\|_2^2\}] + \mathcal{O}(h^{-2}) \\
& = \mathcal{O}(h^{-3}).
\end{aligned}$$

Calculation of the bias and variance term The bias term is calculated as

$$\begin{aligned}
& \mathbb{E}[\mathbb{P}_n[\psi_1^{\mathcal{K}}(S, A, R; q, \pi^b)] - Z(\theta)] \\
& = \mathbb{E} \left[-\frac{\nabla \tau_\theta(S)}{h^2} \int k^{(1)} \left(\frac{a - \tau_\theta(S)}{h} \right) q(S, a) da \right] - Z(\theta) \\
& = \mathbb{E} \left[-\frac{\nabla \tau_\theta(S)}{h} \int k^{(1)}(u) q(S, \tau_\theta(S) + uh) du \right] - Z(\theta) \\
& = \mathbb{E} \left[\nabla \tau_\theta(S) \int k(u) q^{(1)}(S, \tau_\theta(S) + uh) du \right] - Z(\theta) \\
& = \mathbb{E} \left[\nabla \tau_\theta(S) \int k(u) \{q^{(1)}(S, \tau_\theta(S) + uh) - q^{(1)}(S, \tau_\theta(S))\} du \right] \\
& = 0.5h^2 \mathbb{E} \left[\nabla \tau_\theta(S) q^{(3)}(S, \tau_\theta(S)) \right] \int u^2 k(u) du + \mathcal{O}(h^2).
\end{aligned}$$

In the last line, we have used that the function $a \rightarrow q(s, a)$ is a C^3 -function. In addition, we also use a fact $nh^7 = \mathcal{O}(1)$ to say

$$\begin{aligned}
\hat{Z} - Z(\theta) & = \mathbb{E}[\mathbb{P}_n[\psi_1^{\mathcal{K}}(S, A, R; q, \pi^b)] - Z(\theta)] + \mathcal{O}(h^2) \\
& = 0.5h^2 \mathbb{E} \left[\nabla \tau_\theta(S) q^{(3)}(S, \tau_\theta(S)) \right] \int u^2 k(u) du + \mathcal{O}(n^{-1/2} h^{-3/2}).
\end{aligned}$$

Remark 10. Heuristically, this $0.5 \mathbb{E} [\nabla \tau_\theta(S) q^{(3)}(S, \tau_\theta(S))]$ is calculated by differentiating the bias term of the OPE estimator: $0.5 \mathbb{E} [q^{(2)}(S, \tau_\theta(S))]$.

The variance term is calculated as

$$\begin{aligned}
\text{var}[\mathbb{P}_n[\psi_1^{\mathcal{K}}(S, A, R; q, \pi^b)]] &= \frac{1}{n} \mathbb{E} [\psi_1^{\mathcal{K}}(S, A, R; q, \pi^b)^2] + o(\frac{1}{nh^3}) \\
&= \frac{1}{n} \mathbb{E} \left[\otimes \nabla \tau_\theta(S) \frac{\{R - q(S, A)\}^2}{\{\pi^b(A | S)\}^2} K_h^{(1)}(A - \tau_\theta(S))^2 \right] + o(\frac{1}{nh^3}) \\
&= \frac{1}{n} \mathbb{E} [\otimes \nabla \tau_\theta(S) \int \frac{\{r - q(S, a)\}^2}{\pi^b(a | S)} K_h^{(1)}(a - \tau_\theta(S))^2 p(r | S, a) d(a, r)] \\
&= \frac{1}{nh^4} \mathbb{E} [\otimes \nabla \tau_\theta(S) \int \frac{\{r - q(S, a)\}^2}{\pi^b(a | S)} k^{(1)} \left(\frac{a - \tau_\theta(S)}{h} \right)^2 p(r | S, a) d(a, r)] + o(\frac{1}{nh^3}) \\
&= \frac{1}{nh^3} \mathbb{E} [\otimes \nabla \tau_\theta(S) \int \frac{\{r - q(S, \tau_\theta(S) + hu)\}^2}{\pi^b(\tau_\theta(S) + hu | S)} k^{(1)}(u)^2 p(r | S, \tau_\theta(S) + hu) d(u, r)] + o(\frac{1}{nh^3}) \\
&= \frac{1}{nh^3} \left\{ \mathbb{E} [\otimes \nabla \tau_\theta(S) \int \frac{\{r - q(S, \tau_\theta(S))\}^2}{\pi^b(\tau_\theta(S) | S)} k^{(1)}(u)^2 p(r | S, \tau_\theta(S)) d(u, r)] + o(1) \right\} \\
&= \frac{1}{nh^3} \int k^{(1)}(u)^2 d(u) \left\{ \mathbb{E} \left[\otimes \nabla \tau_\theta(S) \frac{\text{var}[R | S, \tau_\theta(S)]}{\pi^b(\tau_\theta(S) | S)} \right] + o(1) \right\}.
\end{aligned}$$

Thus,

$$\begin{aligned}
\text{var}[\mathbb{P}_n[\psi_1^{\mathcal{K}}(S, A, R; \hat{q}, \hat{\pi}^b)]] &= \text{var}[\mathbb{P}_n[\psi_1^{\mathcal{K}}(S, A, R; q, \pi^b)]] + o(n^{-1}h^{-3}) \\
&= \frac{1}{nh^3} \Omega_2^{(1)}(k) \left\{ \mathbb{E} \left[\otimes \nabla \tau_\theta(S) \frac{\text{var}[R | S, \tau_\theta(S)]}{\pi^b(\tau_\theta(S) | S)} \right] + o(1) \right\}.
\end{aligned}$$

D.4 Proof of Theorem 3

Replacing estimators with true functions We define

$$\phi^{\mathcal{K}}(\pi^b, q^{\mathcal{K}}) = v_0^{\mathcal{K}} + \sum_{t=1}^H \lambda_t^{\mathcal{K}} \{r_t - q_t^{\mathcal{K}} + v_{t+1}^{\mathcal{K}}\}.$$

Here, we prove that nuisance estimators can be replace with true functions:

$$\begin{aligned}
\mathbb{E}[A_n] &= o((nh^H)^{-1/2}), \quad \text{var}[A_n] = o((nh^H)^{-1}), \\
A_n &= \mathbb{P}_{\mathcal{U}_1}[\phi^{\mathcal{K}}(\hat{\pi}^{b[1]}, \hat{q}^{\mathcal{K}[1]}) | \mathcal{U}_2] + \mathbb{P}_{\mathcal{U}_2}[\phi^{\mathcal{K}}(\hat{\pi}^{b[2]}, \hat{q}^{\mathcal{K}[2]}) | \mathcal{U}_1] - \mathbb{E}[\phi^{\mathcal{K}}(\pi^b, q^{\mathcal{K}})].
\end{aligned}$$

Then, what we have to prove is

$$\begin{aligned}
\mathbb{E}[\mathbb{E}[\phi^{\mathcal{K}}(\hat{\pi}^b, \hat{q}^{\mathcal{K}}) - \phi^{\mathcal{K}}(\pi^b, q^{\mathcal{K}}) | \mathcal{U}_2]] &= o((nh^H)^{-1/2}), \\
\mathbb{E}[\mathbb{E}[\{\phi^{\mathcal{K}}(\hat{\pi}^b, \hat{q}^{\mathcal{K}}) - \phi^{\mathcal{K}}(\pi^b, q^{\mathcal{K}})\}^2 | \mathcal{U}_2]] &= o(h^{-H}).
\end{aligned}$$

The rest of the part is proved as Theorem 4. Therefore, we omit the proof here.

Next, we analyze the bias and variance.

Bias part First, we have

$$\mathbb{E}[\mathbb{P}_n[\phi^{\mathcal{K}}(\pi^b, q^{\mathcal{K}})]] = \mathbb{E}[\sum_t \lambda_t^{\mathcal{K}} r_t] - J(\theta) = \sum_t \mathbb{E}[\{\lambda_t^{\mathcal{K}} - \lambda_t\} r_t].$$

Here, we use a doubly robust property of $\phi^{\mathcal{K}}$. Then, by defining $c = q_t(s_t, a_t)$, the above is equal to

$$\begin{aligned}
&\sum_t \left\{ \int \mathbb{E}[R_t | S_t = s_t, A_t = a_t] \prod_{i=1}^t \frac{K_h(a_i - \tau(s_i))}{\pi^b(a_i | s_i)} \{ \prod_{i=1}^t \pi^b(a_i | s_i) p(s_i | s_{i-1}, a_{i-1}) \} d(h_{a_t}) - \mathbb{E}[\lambda_t r_t] \right\} \\
&= \sum_t \left\{ \int c(s_t, \tau_t(s_t) + hu_t) \prod_{i=1}^t k(u_i) \prod_{i=1}^t p(s_i | s_{i-1}, \tau(s_{i-1}) + hu_i) d(h_{a_t}^u) - \mathbb{E}[\lambda_t r_t] \right\} \\
&= \sum_t \left\{ \int \{c(s_t, \tau_t(s_t) + hu_t) - c(s_t, \tau_t(s_t))\} \prod_{i=1}^t k(u_i) \prod_{i=1}^t p(s_i | s_{i-1}, \tau(s_{i-1}) + hu_i) d(h_{a_t}^u) - \right. \\
&\quad \left. + \int c(s_t, \tau_t(s_t)) \prod_{i=1}^t k(u_i) \prod_{i=1}^t \{p(s_i | s_{i-1}, \tau(s_{i-1}) + hu_i) - p(s_i | s_{i-1}, \tau(s_{i-1}))\} d(h_{a_t}^u) \right\}
\end{aligned}$$

where $c(s_t, a_t) = \mathbb{E}[Y_t | s_t, a_t]$, $h_{a_t}^u = \{s_1, u_1, s_2, \dots\}$, $h_{a_t}^{u\tau} = \{s_1, \tau(s_1) + hu_1, s_2, \dots\}$, $h_{a_t}^\tau = \{s_1, \tau(s_1), s_2, \dots\}$. Then, we have

$$\sum_t \left\{ \int \{h^2 u_t^2 c_t^{(2)}(s_t, \tau(s_t))\} \left\{ \prod_{i=1}^t k(u_i) p(s_i | s_{i-1}, \tau(s_{i-1})) \right\} d(h_{a_t}^u) - \right. \\ \left. \int c(s_t, a_t) \left\{ \prod_{i=1}^t k(u_i) \right\} \sum_{l=1}^{t-1} \{h^2 u_l^2 p^{(2)}(s_{l+1} | s_l, \tau(s_l)) \prod_{j \neq l} p(s_j | s_{j-1}, \tau(s_{j-1}))\} d(h_{a_t}^u) \right\} + o(h^2).$$

Here, $h_{a_t}^{-u} = \{s_1, s_2, \dots\}$. Finally, it is equal to

$$0.5h^2 M_2(k) \sum_{t=1}^H \left\{ \mathbb{E}_\tau \left[\frac{r_t p^{(2)}(r_t | s_t, \tau_t(s_t))}{p(r_t | s_t, \tau_t(s_t))} \right] + \sum_{l=1}^{t-1} \mathbb{E}_\tau \left[\frac{r_t p^{(2)}(s_{l+1} | s_l, \tau(s_l))}{p(s_{l+1} | s_l, \tau(s_l))} \right] \right\} + o(h^2).$$

Then, we have

$$\mathbb{E}[\phi^K(\hat{\pi}^b, \hat{q}^K)] = \mathbb{E}[\phi^K(\pi^b, q^K)] + o((nh^H)^{-1/2}) \\ = 0.5h^2 M_2(k) \sum_{t=1}^H \left\{ \mathbb{E}_\tau \left[\frac{r_t p^{(2)}(r_t | s_t, \tau_t(s_t))}{p(r_t | s_t, \tau_t(s_t))} \right] + \sum_{l=1}^{t-1} \mathbb{E}_\tau \left[\frac{r_t p^{(2)}(s_{l+1} | s_l, \tau(s_l))}{p(s_{l+1} | s_l, \tau(s_l))} \right] \right\} + o(h^2) + o((nh^H)^{-1/2}) \\ = 0.5h^2 M_2(k) \sum_{t=1}^H \left\{ \mathbb{E}_\tau \left[\frac{r_t p^{(2)}(r_t | s_t, \tau_t(s_t))}{p(r_t | s_t, \tau_t(s_t))} \right] + \sum_{l=1}^{t-1} \mathbb{E}_\tau \left[\frac{r_t p^{(2)}(s_{l+1} | s_l, \tau(s_l))}{p(s_{l+1} | s_l, \tau(s_l))} \right] \right\} + o((nh^H)^{-1/2}).$$

Variance part The variance is

$$\text{var}[\mathbb{P}_n[\phi^K(\pi^b, q^K)]] = \frac{1}{n} \text{var}[\phi^K(\pi^b, q^K)] \\ = \frac{1}{n} \sum_{t=1}^H \mathbb{E} \left[\left\{ \prod_{i=1}^t \frac{h^{-1} k(h^{-1}(A_i - S_i))}{\pi^b(A_i | S_i)} \right\}^2 \text{var}[R_t + v^K(S_{t+1}) | A_t, S_t] \right]. \quad (23)$$

Here, $\forall f$, we have

$$\mathbb{E} \left[\left\{ \prod_{i=1}^t \frac{h^{-1} k(h^{-1}(A_i - \tau(S_i)))}{\pi^b(A_i | S_i)} \right\}^2 f(\mathcal{H}_{A_t}) \right] \\ = \frac{1}{h^{2t}} \int \prod_{i=1}^t \frac{k^2(h^{-1}(a_i - \tau(s_i)))}{\pi^b(a_i | s_i)} f(h_{a_t}) \{ \prod_{i=1}^t p(s_i | s_{i-1}, a_{i-1}) \} d(h_{a_t}) \\ = \frac{1}{h^t} \int \prod_{i=1}^t \frac{k^2(u_i)}{\pi^b(hu_i + \tau_i(s_i) | s_i)} f(h_{a_t}^{u\tau}) \{ \prod_{i=1}^t p(s_i | s_{i-1}, hu_{i-1} + \tau(s_{i-1})) \} d(h_{a_t}^u) \\ = \frac{1}{h^t} \left\{ \int \prod_{i=1}^t \frac{k^2(u_i)}{\pi^b(\tau_i(s_i) | s_i)} f(h_{a_t}^{u\tau}) \{ \prod_{i=1}^t p(s_i | s_{i-1}, \tau(s_{i-1})) \} d(h_{a_t}^u) + o(1) \right\} \\ = \frac{1}{h^t} \left\{ \prod_{i=1}^t \int k^2(u_i) d(u_i) \right\} \left\{ \int \prod_{i=1}^t \left\{ \frac{1}{\pi^b(\tau(s_i) | s_i)} \right\} f(h_{a_t}^\tau) \{ \prod_{i=1}^t p(s_i | s_{i-1}, \tau(s_{i-1})) \} d(h_{a_t}^\tau) + o(1) \right\}.$$

Therefore, (23) is

$$\frac{1}{nh^H} \left\{ \left\{ \prod_{i=1}^H \int k^2(u_i) d(u_i) \right\} \mathbb{E}_\tau \left[\prod_{i=1}^H \left\{ \frac{1}{\pi^b(\tau_i(s_i) | s_i)} \right\} \text{var}[r_H | s_H, \tau(s_H)] \right] + o(1) \right\},$$

Finally,

$$\text{var}[\mathbb{P}_n[\phi^K(\hat{\pi}^b, \hat{q}^K)]] = \text{var}[\mathbb{P}_n[\phi^K(\pi^b, q^K)]] + o(n^{-1}h^{-H}) \\ = \frac{1}{nh^H} \left\{ \Omega_2^H(k) \mathbb{E}_\tau \left[\prod_{i=1}^H \left\{ \frac{1}{\pi^b(\tau_i(s_i) | s_i)} \right\} \text{var}[r_H | s_H, \tau(s_H)] \right] + o(1) \right\}.$$

D.5 Proof of Theorem 4

Replacing estimators with true functions We define

$$\phi^K(w^K, \pi^b, q^K) = v_0^K + \sum_{t=1}^H \frac{w_t^K K_h(a_t - \tau_t(s_t))}{\pi^b(a_t|s_t)} \{r_t - q_t^K + v_{t+1}^K\}.$$

Here, we prove that nuisance estimators can be replace with true functions:

$$\begin{aligned} & \mathbb{E}[\mathbb{P}_{\mathcal{U}_1}[\phi^K(\hat{w}^{K[1]}, \hat{\pi}^{b[1]}, \hat{q}^{K[1]}|\mathcal{U}_2) + \mathbb{P}_{\mathcal{U}_2}[\phi^K(\hat{w}^{K[2]}, \hat{\pi}^{b[2]}, \hat{q}^{K[2]}|\mathcal{U}_1) - \mathbb{E}[\phi^K(w^K, \pi^b, q^K)]] = o((nh)^{-1/2}) \\ & \mathbb{E}[\{\mathbb{P}_{\mathcal{U}_1}[\phi^K(\hat{w}^{K[1]}, \hat{\pi}^{b[1]}, \hat{q}^{K[1]}|\mathcal{U}_2) + \mathbb{P}_{\mathcal{U}_2}[\phi^K(\hat{w}^{K[2]}, \hat{\pi}^{b[2]}, \hat{q}^{K[2]}|\mathcal{U}_1) - \mathbb{E}[\phi^K(w^K, \pi^b, q^K)]]^2\} = o((nh)^{-1/2}). \end{aligned}$$

Then, what we have to prove is

$$\mathbb{E}[\mathbb{E}[\phi^K(\hat{w}^{K[1]}, \hat{\pi}^{b[1]}, \hat{q}^{K[1]}|\mathcal{U}_2) - \phi^K(w^K, \pi^b, q^K)|\mathcal{U}_2]] = o((nh)^{-1/2}), \quad (24)$$

$$\mathbb{E}[\mathbb{E}[\{\phi^K(\hat{w}^{K[1]}, \hat{\pi}^{b[1]}, \hat{q}^{K[1]}|\mathcal{U}_2) - \phi^K(w^K, \pi^b, q^K)\}^2|\mathcal{U}_2]] = o((h)^{-1}). \quad (25)$$

From now on, we omit $[1], \mathcal{U}_2$ for simplicity. By defining $\tilde{w}_t^K(s_t, a_t) = w_t^K(s_t)/\pi^b(a_t|s_t)$, Eq. (24) is proved by

$$\begin{aligned} & \mathbb{E}[\phi^K(\hat{w}^K, \hat{\pi}^b, \hat{q}^K) - \phi^K(w^K, \pi^b, q^K)|\mathcal{U}_2] \\ &= \mathbb{E}[\sum_{t=1}^H (\hat{w}_t^K(S_t, A_t) - \tilde{w}_t^K(S_t, A_t)) K_h((A_t - \tau(S_t))h^{-1}) (-\hat{q}_t^K(S_t, A_t) + q_t^K(S_t, A_t) + \hat{v}_{t+1}^K(S_{t+1}) - v_{t+1}^K(S_{t+1}))] \end{aligned} \quad (26)$$

$$= M_2(k) \mathbb{E}[\sum_{t=1}^H \left\{ \hat{w}_t^K(S_t, \tau(S_t)) - \tilde{w}_t^K(S_t, \tau(S_t)) \right\} \left\{ -\hat{q}_t^K(S_t, \tau(S_t)) + q_t^K(S_t, \tau(S_t)) \right\} \pi^b(\tau(S_t)|S_t)] \quad (27)$$

$$\begin{aligned} & + M_2(k) \mathbb{E}[\sum_{t=1}^H \left\{ \hat{w}_t^K(S_t, \tau(S_t)) - \tilde{w}_t^K(S_t, \tau(S_t)) \right\} (\hat{v}_{t+1}^K(S_{t+1}) - v_{t+1}^K(S_{t+1})) \pi^b(\tau(S_t)|S_t)] + o(h^2) \\ &= M_2(k) \mathbb{E}[\sum_{t=1}^H \left\{ \hat{w}_t^K(S_t, \tau(S_t)) - \tilde{w}_t^K(S_t, \tau(S_t)) \right\} \left\{ -\hat{q}_t^K(S_t, \tau(S_t)) + q_t^K(S_t, \tau(S_t)) \right\} \pi^b(\tau(S_t)|S_t)] \end{aligned} \quad (28)$$

$$\begin{aligned} & + M_2(k) \mathbb{E}[\sum_{t=1}^H \left\{ \hat{w}_t^K(S_t, \tau(S_t)) - \tilde{w}_t^K(S_t, \tau(S_t)) \right\} \left\{ \hat{q}_{t+1}^K(S_{t+1}, \tau(S_{t+1})) - q_{t+1}^K(S_{t+1}, \tau(S_{t+1})) \right\} \pi^b(\tau(S_t)|S_t)] + o(h^2) \\ &= M_2(k) \sum_{t=1}^H \mathbb{E}[\|\hat{w}_t^K - \tilde{w}_t^K\|_\infty \left\{ \|\hat{q}_t^K - q_t^K\|_\infty + \|\hat{q}_{t+1}^K - q_{t+1}^K\|_\infty \right\}] + o(h^2) \end{aligned} \quad (29)$$

$$= o(n^{-1/2}h^{-1/2}). \quad (30)$$

Here, from (26) to (27),

$$\mathbb{E}[\max\{\|\hat{\pi}^b - \pi^b\|_{1,\infty}, \|\hat{w}_j^K - w_j^K\|_\infty, \|\hat{q}_j^K - q_j^K\|_{1,\infty}\}] = o(1).$$

From (27) to (28), $v_t^K(s) = q_t^K(s, \tau(s)) + o(1)$. From (28) to (29), we use

$$\mathbb{E}[\max\{\|\hat{\pi}^b - \pi^b\|_\infty, \|\hat{w}_j^K - w_j^K\|_\infty\} \|\hat{q}_j^K - q_j^K\|_\infty] = o(n^{-1/2}h^{-1/2}).$$

From (29) to (30), we use $nh^5 = \mathcal{O}(1)$. Eq. (25) is proved by

$$\begin{aligned} & \mathbb{E}[\mathbb{E}[\{\phi^K(\hat{w}^K, \hat{\pi}^b, \hat{q}^K) - \phi^K(w^K, \pi^b, q^K)\}^2|\mathcal{U}_2]] \\ &= \mathbb{E}[h^{-1} \max\{\|\hat{q}_t^K(S_t, \tau(S_t)) - q_t^K(S_t, \tau(S_t))\|_2^2, \|\hat{w}^K(S_t, \tau(S_t)) - \tilde{w}^K(S_t, \tau(S_t))\|_2^2\}] = o(h^{-1}). \end{aligned}$$

Calculation of bias and variance term Bias and variance terms are bounded as follows.

Bias part

Bias is the same as the proof of Theorem 3. It is reduced to

$$0.5h^2 M_2(k) \sum_{t=1}^H \left\{ \mathbb{E}_\tau [c^{(2)}(s_t, \tau_t(s_t))] + \sum_{j=1}^{t-1} \mathbb{E}_\tau \left[\frac{r_j p^{(2)}(s_{j+1} | s_j, \tau_j(s_j))}{p(s_j | s_j, \tau_j(s_j))} \right] \right\} + o(n^{-1/2} h^{-1/2}).$$

where $c(s, a) = \mathbb{E}[r | s, a]$.

Variance part

The variance is

$$\begin{aligned} \text{var}[\mathbb{P}_n[\phi^\mathcal{K}(w^\mathcal{K}, \pi^\mathbf{b}, q^\mathcal{K})]] &= \frac{1}{n} \text{var}[\phi^\mathcal{K}(w^\mathcal{K}, \pi^\mathbf{b}, q^\mathcal{K})] \\ &= \frac{1}{n} \sum_{t=0}^H \mathbb{E} \left[\left\{ \mathbb{E} \left[\prod_{i=1}^t \frac{h^{-1}k(h^{-1}(A_i - S_i))}{\pi^\mathbf{b}(A_i | S_i)} \mid S_t \right] \right\}^2 \left\{ \frac{h^{-1}k(h^{-1}(A_t - \tau(S_t)))}{\pi^\mathbf{b}(A_t | S_t)} \right\}^2 \text{var}[R_t + v^\mathcal{K}(S_{t+1}) \mid A_t, S_t] \right]. \end{aligned} \quad (31)$$

First, we have

$$\begin{aligned} &\mathbb{E} \left[\prod_{i=1}^{t-1} \frac{h^{-1}k(h^{-1}(A_i - \tau_i(S_i)))}{\pi^\mathbf{b}(A_i | S_i)} \mid S_t \right] \\ &= \int \prod_{i=1}^{t-1} \frac{h^{-1}k(h^{-1}(a_i - \tau_i(s_i)))}{\pi^\mathbf{b}(a_i | s_i)} \left\{ \prod_{i=1}^{t-1} p(s_{i+1} | s_i, a_i) \pi^\mathbf{b}(a_i | s_i) \right\} \frac{p_1(s_1)}{p_{\pi^\mathbf{b}}(s_t)} d(h_{a_{t-1}}) \\ &= \int \prod_{i=1}^{t-1} k(u_i) \prod_{i=1}^{t-1} p(s_{i+1} | s_i, \tau_i(s_i) + hu_i) \frac{p_1(s_1)}{p_{\pi^\mathbf{b}}(s_t)} d(h_{a_{t-1}}^u) \\ &= \left\{ \int \prod_{i=1}^{t-1} k(u_i) \prod_{i=1}^{t-1} p(s_{i+1} | s_i, \tau_i(s_i)) \frac{p_1(s_1)}{p_{\pi^\mathbf{b}}(s_t)} d(h_{a_{t-1}}^u) \right\} + o(1) = \frac{p_{\tau_\theta}(s_t)}{p_{\pi^\mathbf{b}}(s_t)} + o(1). \end{aligned}$$

Therefore, $\forall f$ we have

$$\begin{aligned} &\mathbb{E} \left[\left\{ \mathbb{E} \left[\prod_{i=1}^t \frac{h^{-1}k(h^{-1}(A_i - \tau_i(S_i)))}{\pi^\mathbf{b}(A_i | S_i)} \mid S_t \right] \right\}^2 \left\{ \frac{h^{-1}k(h^{-1}(A_t - \tau(S_t)))}{\pi^\mathbf{b}(A_t | S_t)} \right\}^2 f(S_t, A_t) \right] \\ &= \mathbb{E} \left[\left\{ \frac{p_{\tau_\theta}(S_t)}{p_{\pi^\mathbf{b}}(S_t)} + o(1) \right\}^2 \left\{ \frac{h^{-1}k(h^{-1}(A_t - \tau(S_t)))}{\pi^\mathbf{b}(A_t | S_t)} \right\}^2 f(S_t, A_t) \right] + o(1) \\ &= \frac{1}{h} \left\{ \int \frac{p_{\tau_\theta}^2(s_t)}{p_{\pi^\mathbf{b}}(s_t)} \left\{ \frac{k(u_t)}{\pi^\mathbf{b}(\tau_t(s_t) + u_t | s_t)} \right\}^2 f(s_t, a_t) \pi^\mathbf{b}(\tau_t(s_t) + u_t | s_t) d(s_t, u_t) + o(1) \right\} \\ &= \frac{1}{h} \left\{ \Omega_2(k) \int \frac{p_{\tau_\theta}^2(s_t)}{p_{\pi^\mathbf{b}}(s_t) \pi^\mathbf{b}(\tau_t(s_t) | s_t)} f(s_t, \tau_t(s_t)) d(s_t) + o(1) \right\}. \end{aligned}$$

In addition, noting

$$\begin{aligned} v_t^\mathcal{K}(s_t) &= \int h^{-1}k(h^{-1}(a_t - \tau_t(s_t))) q_t^\mathcal{K}(s_t, a_t) d(a_t) \\ &= \int k(u_t) q_t^\mathcal{K}(s_t, \tau_t(s_t)) d(u_t) + o(1) = q_t^\mathcal{K}(s_t, \tau_t(s_t)) + o(1), \end{aligned}$$

by induction, we have $v_t^\mathcal{K}(s_t) = v_t^{\pi^\mathbf{e}}(s_t) + o(1) = q_t^{\pi^\mathbf{e}}(s_t, \tau_t(s_t)) + o(1)$. Then, noting

$$\text{var}[R_t + v_{t+1}^\mathcal{K}(S_{t+1}) \mid S_t = s_t, A_t = a_t] = \text{var}[R_t + v_{t+1}^{\pi^\mathbf{e}}(S_{t+1}) \mid S_t = s_t, A_t = a_t] + o(1).$$

Therefore, the variance term is

$$\begin{aligned} \text{var}[\mathbb{P}_n[\phi^\mathcal{K}(\hat{w}^\mathcal{K}, \hat{\pi}^\mathbf{b}, \hat{q}^\mathcal{K})]] &= \text{var}[\mathbb{P}_n[\phi^\mathcal{K}(w^\mathcal{K}, \pi^\mathbf{b}, q^\mathcal{K})]] + o(n^{-1} h^{-1}) \\ &= \frac{1}{nh} \sum_{t=1}^H \mathbb{E}_{\tau_\theta} \left[\frac{p_{\tau_\theta}(s_t)}{p_{\pi^\mathbf{b}}(s_t) \pi^\mathbf{b}(\tau_t(s_t) | s_t)} \text{var}[r_t + v_{t+1}^{\pi^\mathbf{e}}(s_{t+1}) \mid s_t, \tau_\theta(s_t)] \right] \Omega_2(k) + o(1/nh). \end{aligned}$$

Order of main constants in the bias and variance terms Bias and variance terms are upper-bounded as follows.

Bound of B_H

$$\begin{aligned} & \sum_{t=1}^H \left\{ \mathbb{E}_\tau \left[\frac{r_t p^{(2)}(r_t | s_t, \tau_t(s_t))}{p(r_t | s_t, \tau_t(s_t))} \right] + \sum_{j=1}^{t-1} \mathbb{E}_\tau \left[\frac{r_t p^{(2)}(s_{j+1} | s_j, \tau_j(s_j))}{p(s_{j+1} | s_j, \tau_j(s_j))} \right] \right\} \\ & \leq R_{\max} \left\{ HG_1 + \frac{H(H-1)}{2} G_1 \right\}. \end{aligned}$$

we used an argument:

$$\begin{aligned} |\mathbb{E}_\tau \left[\frac{r_t p^{(2)}(r_t | s_t, \tau_t(s_t))}{p(r_t | s_t, \tau_t(s_t))} \right]| & \leq \mathbb{E}_\tau \left[\frac{|r_t p^{(2)}(r_t | s_t, \tau_t(s_t))|}{p(r_t | s_t, \tau_t(s_t))} \right] \leq R_{\max} \mathbb{E}_\tau \left[\frac{|p^{(2)}(r_t | s_t, \tau_t(s_t))|}{p(r_t | s_t, \tau_t(s_t))} \right] \\ & = R_{\max} \int |p^{(2)}(r_t | s_t, \tau_t(s_t))| p_\tau(s_t) d(r_t, s_t) \\ & \leq R_{\max} \int |p^{(2)}(r_t | s_t, \tau_t(s_t))| dr_t d(s_t) \leq R_{\max} G_1^{(2)}, \\ |\mathbb{E}_\tau \left[\frac{r_t p^{(2)}(s_{j+1} | s_j, \tau_j(s_j))}{p(s_{j+1} | s_j, \tau_j(s_j))} \right]| & \leq \mathbb{E}_\tau \left[\frac{|r_t p^{(2)}(s_{j+1} | s_j, \tau_j(s_j))|}{p(s_{j+1} | s_j, \tau_j(s_j))} \right] \leq R_{\max} \mathbb{E}_\tau \left[\frac{|p^{(2)}(s_{j+1} | s_j, \tau_j(s_j))|}{p(s_{j+1} | s_j, \tau_j(s_j))} \right] \\ & = R_{\max} \int |p^{(2)}(s_{j+1} | s_j, \tau_j(s_j))| p_\tau(s_j) d(s_{j+1}, s_j) \\ & \leq R_{\max} \int |p^{(2)}(s_{j+1} | s_j, \tau_j(s_j))| p_\tau(s_j) ds_{j+1} d(s_j) \leq R_{\max} G_2^{(2)}. \end{aligned}$$

Bound of V_H

$$\begin{aligned} & \sum_{t=1}^H \mathbb{E}_{\tau_\theta} \left[\frac{p_{\tau_\theta}(s_t)}{p_{\pi^b}(s_t) \pi^b(\tau_t(s_t) | s_t)} \text{var}[r_t + v_{t+1}^{\pi^e}(s_{t+1}) | s_t, \tau_\theta(s_t)] \right] \\ & \leq \sum_{t=1}^H \mathbb{E}_{\tau_\theta} \left[C_1 C_2 \text{var}[r_t + v_{t+1}^{\pi^e}(s_{t+1}) | s_t, \tau_\theta(s_t)] \right] \\ & = C_1 C_2 \text{var}_{\tau_\theta} \left[\sum_{t=1}^H r_t \right] \leq C_1 C_2 R_{\max}^2 H^2. \end{aligned}$$

D.6 Proof of Theorem 5

Replacing estimators with true functions Here, we prove that nuisance estimators can be replaced with true functions:

$$\begin{aligned} \mathbb{E}[A_n] & = o((nh^3)^{-1/2}), \quad \text{var}[A_n] = o(n^{-1}h^{-1}), \\ A_n & = \mathbb{P}_{\mathcal{U}_1}[\psi^{\mathcal{K}}(\hat{w}^{\mathcal{K}[1]}, \hat{q}^{\mathcal{K}[1]}, \hat{d}^{w^{\mathcal{K}[1]}}, \hat{d}^{q^{\mathcal{K}[1]}}, \hat{\pi}^b) | \mathcal{U}_2] + \mathbb{P}_{\mathcal{U}_2}[\psi^{\mathcal{K}}(\hat{w}^{\mathcal{K}[2]}, \hat{q}^{\mathcal{K}[2]}, \hat{d}^{w^{\mathcal{K}[2]}}, \hat{d}^{q^{\mathcal{K}[2]}}, \hat{\pi}^b) | \mathcal{U}_1] \\ & \quad - \mathbb{E}[\psi^{\mathcal{K}}(w^{\mathcal{K}}, q^{\mathcal{K}}, d^{w^{\mathcal{K}}}, d^{q^{\mathcal{K}}}, \pi^b)]. \end{aligned}$$

Then, what we have to prove is

$$\mathbb{E}[\mathbb{E}[\psi^{\mathcal{K}}(\hat{w}^{\mathcal{K}[1]}, \hat{q}^{\mathcal{K}[1]}, \hat{d}^{w^{\mathcal{K}[1]}}, \hat{d}^{q^{\mathcal{K}[1]}}, \hat{\pi}^{b[1]}) - \psi^{\mathcal{K}}(w^{\mathcal{K}}, q^{\mathcal{K}}, d^{w^{\mathcal{K}}}, d^{q^{\mathcal{K}}}, \pi^b) | \mathcal{U}_2]] = o((nh^3)^{-1/2}), \quad (32)$$

$$\mathbb{E}[\mathbb{E}[\{\psi^{\mathcal{K}}(\hat{w}^{\mathcal{K}[1]}, \hat{q}^{\mathcal{K}[1]}, \hat{d}^{w^{\mathcal{K}[1]}}, \hat{d}^{q^{\mathcal{K}[1]}}, \hat{\pi}^{b[1]}) - \psi^{\mathcal{K}}(w^{\mathcal{K}}, q^{\mathcal{K}}, d^{w^{\mathcal{K}}}, d^{q^{\mathcal{K}}}, \pi^b)\}^2 | \mathcal{U}_2]] = o((h)^{-1}). \quad (33)$$

From now on, we omit $[1], \mathcal{U}_2$.

Eq. (32) is proved by

$$\begin{aligned}
& \mathbb{E}[\mathbb{E}[\psi^\mathcal{K}(\hat{w}^\mathcal{K}, \hat{q}^\mathcal{K}, \hat{d}^{w^\mathcal{K}}, \hat{d}^{q^\mathcal{K}}, \hat{\pi}^b) - \psi^\mathcal{K}(w^\mathcal{K}, q^\mathcal{K}, d^{w^\mathcal{K}}, d^{q^\mathcal{K}}, \pi^b) | \mathcal{U}_1]] \\
&= \mathbb{E}[\sum_{t=1}^H \{ \hat{d}_t^{w^\mathcal{K}}(S_t) \hat{\eta}(S_t, A_t) - d_t^{w^\mathcal{K}}(S_t) \eta(S_t, A_t) \} K_h((A_t - \tau(S_t))) \\
&\quad \times (-\hat{q}_t^\mathcal{K}(S_t, A_t) + q_t^\mathcal{K}(S_t, A_t) + \hat{v}_{t+1}^\mathcal{K}(S_{t+1}) - v_{t+1}^\mathcal{K}(S_{t+1}))] \\
&+ \mathbb{E}[\sum_{t=1}^H \{ \hat{w}_t^\mathcal{K}(S_t) \hat{\eta}(S_t, A_t) - w_t^\mathcal{K}(S_t) \eta(S_t, A_t) \} K_h((A_t - \tau(S_t))) \\
&\quad \times (-\hat{d}_t^{q^\mathcal{K}}(S_t, A_t) + d_t^{q^\mathcal{K}}(S_t, A_t) + \hat{d}_{t+1}^{v^\mathcal{K}}(S_{t+1}) - d_{t+1}^{v^\mathcal{K}}(S_{t+1}))] \\
&- \mathbb{E}[\sum_{t=1}^H \{ \hat{w}_t^\mathcal{K}(S_t) \hat{\eta}(S_t, A_t) - w_t^\mathcal{K}(S_t) \eta(S_t, A_t) \} \times \\
&\quad h^{-2} k^{(1)}((A_t - \tau(S_t)) h^{-1}) \nabla \tau(S_t) (-\hat{q}_t^\mathcal{K}(S_t, A_t) + q_t^\mathcal{K}(S_t, A_t) + \hat{v}_{t+1}^\mathcal{K}(S_{t+1}) - v_{t+1}^\mathcal{K}(S_{t+1}))] \\
&= o(h^{-2}) + \mathbb{E}[\sum_{t=1}^H \max\{\|\hat{w}_t^{\pi^e} - w_t^{\pi^e}\|_\infty, \|\hat{\eta}_t - \eta_t\|_{1,\infty}, \|\hat{d}_t^{w^\mathcal{K}} - d_t^{w^\mathcal{K}}\|_\infty\} \\
&\quad \times \max\{\|\hat{q}_t - q_t\|_{1,\infty}, \|\hat{q}_{t+1} - q_{t+1}\|_{1,\infty}, \|\hat{d}_t^{q^\mathcal{K}} - d_t^{q^\mathcal{K}}\|_\infty, \|\hat{d}_{t+1}^{q^\mathcal{K}} - d_{t+1}^{q^\mathcal{K}}\|_\infty\}] \\
&= o(n^{-1/2} h^{-3/2}).
\end{aligned}$$

Eq. (33) is similarly proved.

Bias part First, we have

$$\mathbb{E}[\mathbb{P}_n[\psi(w^\mathcal{K}, q^\mathcal{K}, d^{w^\mathcal{K}}, d^{q^\mathcal{K}}, \pi^b)]] = \nabla \mathbb{E}[\sum_t \lambda_t^\mathcal{K} r_t] - \nabla J(\theta) = \nabla \{\sum_t \mathbb{E}[\{\lambda_t^\mathcal{K} - \lambda_t\} r_t]\}.$$

Here, we use a doubly robust property of ϕ . Then, the above is equal to

$$\begin{aligned}
& - \sum_t \left\{ \int c(s_t, a_t) \sum_{i=1}^t \frac{k^{(1)}((a_i - \tau(s_i)) h^{-1})}{h^2} \nabla \tau_i(s_i) \prod_{j \neq i}^t \frac{k((a_j - \tau(s_j)) h^{-1})}{h} \prod_{l=1}^t p(s_l | s_{l-1}, a_{l-1}) \right\} d(h_{a_t}) - \mathbb{E}[\lambda_t r_t] \Big\} \\
&= - \sum_{t=1}^H \left\{ \int c(s_t, \tau_t(s_t) + hu_t) \sum_{i=1}^t h^{-1} k^{(1)}(u_i) \nabla \tau_i(s_i) \prod_{j \neq i}^t k(u_j) \prod_{l=1}^t p(s_l | s_{l-1}, \tau(s_{l-1}) + hu_l) d(h_{a_t}^u) - \mathbb{E}[\lambda_t r_t] \right\} \\
&= \sum_{t=1}^H \left\{ \int c_t^{(1)}(s_t, \tau_t(s_t) + hu_t) \nabla \tau_t(s_t) \prod_{j=1}^t k(u_j) \prod_{i=1}^t p(s_l | s_{i-1}, \tau(s_{i-1}) + hu_i) d(h_{a_t}^u) + \right. \\
&\quad \left. + \int c(s_t, \tau_t(s_t) + hu_t) \sum_{i=1}^t \nabla \tau_i(s_i) \frac{p^{(1)}(s_i | s_{i-1}, \tau(s_{i-1}) + hu_i)}{p(s_i | s_{i-1}, \tau(s_{i-1}) + hu_i)} \prod_{j=1}^t k(u_j) \prod_{l=1}^t p(s_l | s_{l-1}, \tau(s_{l-1}) + hu_l) d(h_{a_t}^u) \right. \\
&\quad \left. - \mathbb{E}[\lambda_t r_t] \right\}.
\end{aligned}$$

This is equal to

$$\begin{aligned}
& 0.5h^2 M_2(\mathcal{K}) \sum_t \left\{ \int c_t^{(3)}(s_t, \tau_t(s_t)) \prod_{i=1}^t p(s_i | s_{i-1}, \tau(s_{i-1})) d(h_{s_t}) + \right. \\
& \int c_t^{(1)}(s_t, \tau_t(s_t)) \nabla \tau_t(s_t) \sum_{j=2}^t \frac{p^{(2)}(s_j | s_{j-1}, \tau(s_{j-1}))}{p(s_l | s_{i-1}, \tau(s_{i-1}))} \prod_{l=1}^t p(s_l | s_{i-1}, \tau(s_{i-1})) d(h_{s_t}^{-u}) + \\
& + \int c_t^{(2)}(s_t, \tau_t(s_t)) \sum_{i=2}^t \nabla \tau_i(s_i) \frac{p^{(1)}(s_i | s_{i-1}, \tau(s_{i-1}))}{p(s_i | s_{i-1}, \tau(s_{i-1}))} \prod_{l=1}^t p(s_l | s_{i-1}, \tau(s_{i-1})) d(h_{s_t}^{-u}) \\
& + \left. \int c(s_t, \tau_t(s_t)) \sum_{i=2}^t \nabla \tau_i(s_i) \left\{ \frac{p^{(3)}(s_i | s_{i-1}, \tau(s_{i-1}))}{p(s_i | s_{i-1}, \tau(s_{i-1}))} + \frac{p^{(1)}(s_i | s_{i-1}, \tau(s_{i-1}))}{p(s_i | s_{i-1}, \tau(s_{i-1}))} \sum_{j \neq i} \frac{p^{(2)}(s_j | s_{j-1}, \tau(s_{j-1}))}{p(s_j | s_{j-1}, \tau(s_{j-1}))} \right\} \right. \\
& \times \prod_{l=1}^t p(s_l | s_{l-1}, \tau(s_{l-1})) d(h_{s_t}^{-u}) \\
& = 0.5h^2 M_2(k) \tilde{B}_H.
\end{aligned}$$

Here, $h_{a_t}^{-u} = \{s_1, s_2, \dots\}$. In the end, \tilde{B}_H is equal to

$$\begin{aligned}
& \sum_{t=1}^H \left\{ \mathbb{E}_\tau \left[r_t \frac{p^{(3)}(r_t | s_t, \tau_t(s_t))}{p(r_t | s_t, \tau_t(s_t))} \nabla \tau_t(s_t) \right] + \sum_{j=1}^{t-1} \mathbb{E}_\tau \left[r_t \frac{p^{(2)}(r_t | s_t, \tau_t(s_t)) p^{(1)}(s_{j+1} | s_j, \tau_j(s_j))}{p(r_t | s_t, \tau_t(s_t)) p(s_{j+1} | s_j, \tau_j(s_j))} \nabla \tau_j(s_j) \right] \right\} + \\
& + \sum_{t=1}^H \sum_{j=1}^{t-1} \mathbb{E}_\tau \left[r_t \frac{p^{(1)}(r_t | s_t, \tau_t(s_t)) p^{(2)}(s_{j+1} | s_j, \tau_j(s_j))}{p(r_t | s_t, \tau_t(s_t)) p(s_{j+1} | s_j, \tau_j(s_j))} \nabla \tau_t(s_t) \right] + \\
& + \sum_{t=1}^H \left\{ \sum_{j=1}^{t-1} \mathbb{E}_\tau \left[r_t \frac{p^{(3)}(s_{j+1} | s_j, \tau_j(s_j))}{p(s_{j+1} | s_j, \tau_j(s_j))} \nabla \tau_j(s_j) \right] + \sum_{j \neq i}^{t-1} \mathbb{E}_\tau \left[r_t \frac{p^{(1)}(s_{i+1} | s_i, \tau_i(s_i)) p^{(2)}(s_{j+1} | s_j, \tau_j(s_j))}{p(s_{i+1} | s_i, \tau_i(s_i)) p(s_{j+1} | s_j, \tau_j(s_j))} \nabla \tau_i(s_i) \right] \right\}.
\end{aligned}$$

The operator norm of $\otimes \tilde{B}_H$ is upper bounded by

$$R_{\max}^2 \Upsilon^2 \left\{ H G_1^{(3)} + \frac{H(H-1)}{2} \{ G_1^{(2)} G_2^{(1)} + G_1^{(1)} G_2^{(2)} + G_2^{(3)} \} + \frac{H(H-1)(H-2)}{3} G_2^{(1)} G_2^{(2)} \right\}^2.$$

For example,

$$\begin{aligned}
\left\| \otimes \mathbb{E}_\tau \left[r_t \frac{p^{(3)}(s_{j+1} | s_j, \tau_j(s_j))}{p(s_{j+1} | s_j, \tau_j(s_j))} \nabla \tau_j(s_j) \right] \right\|_{\text{op}} & \leq \left\| \otimes \mathbb{E}_\tau \left[\frac{|r_t p^{(3)}(s_{j+1} | s_j, \tau_j(s_j))|}{p(s_{j+1} | s_j, \tau_j(s_j))} \nabla \tau_j(s_j) \right] \right\|_{\text{op}} \\
& \leq R_{\max}^2 G_2^{2(3)} \|\otimes \mathbb{E}_\tau [\nabla \tau_j(s_j)]\|_{\text{op}} \\
& \leq R_{\max}^2 G_2^{2(3)} \|\mathbb{E}_\tau [\otimes \nabla \tau_j(s_j)]\|_{\text{op}} \leq R_{\max}^2 G_2^{2(3)} \Upsilon^2.
\end{aligned}$$

Variance part The variance part is calculated as

$$\begin{aligned}
& \frac{1}{n} \sum_{h=1}^H \mathbb{E} \left[\frac{w_t^{2\mathcal{K}}(S_t) k^{2(1)}(\{A_t - \tau(S_t)\} h^{-1})}{h^4 \pi_t^{2b}(A_t | S_t)} \text{var}[R_t + v_{t+1}^{\mathcal{K}}(S_{t+1}) | S_t, A_t] \otimes \nabla \tau(S_t) \right] + o(h^3/n) \\
& = \frac{1}{nh^3} \sum_{h=1}^H \int k^{2(1)}(u_h) d(u_h) \mathbb{E} \left[\frac{w_t^{2\mathcal{K}}(S_t)}{\pi_t^{2b}(\tau(S_t) | S_t)} \text{var}[R_t + v_{t+1}^{\pi^e}(S_{t+1}) | S_t, \tau(S_t)] \otimes \nabla \tau(S_t) \right] + o(h^3/n) \\
& = \frac{\Omega_2^{(1)}(k)}{nh^3} \sum_{h=1}^H \mathbb{E}_\tau \left[\frac{w^{\pi^e}(s_t)}{\pi_t^b(\tau_t(s_t) | s_t)} \text{var}[r_t + v_{t+1}^{\pi^e}(s_{t+1}) | s_t, \tau_t(s_t)] \otimes \nabla \tau_t(s_t) \right] + o(h^3/n) \\
& = \frac{\Omega_2^{(1)}(k)}{nh^3} \tilde{V}_H + o(h^3/n).
\end{aligned}$$

Then, the operator norm of \tilde{V}_H is upper bounded as

$$\begin{aligned}
& \left\| \sum_{t=1}^H \mathbb{E}_\tau \left[\frac{w^{\pi^e}(s_t)}{\pi_t^b(\tau_t(s_t)|s_t)} \text{var}[r_t + v_{t+1}^{\pi^e}(s_{t+1})|s_t, \tau_t(s_t)] \otimes \nabla \tau_t(s_t) \right] \right\|_{\text{op}} \\
& \leq \|C_1 C_2 \sum_{t=1}^H \mathbb{E}_\tau [\text{var}[r_t + v_{t+1}^{\pi^e}(s_{t+1})|s_t, \tau_t(s_t)] \otimes \nabla \tau_t(s_t)]\|_{\text{op}} \\
& \leq \|C_1 C_2\| \otimes \nabla \tau(s) \|\infty \sum_{t=1}^H \mathbb{E}_\tau [\text{var}[r_t + v_{t+1}^{\pi^e}(s_{t+1})|s_t, \tau_t(s_t)]\|_{\text{op}} \\
& \leq \|C_1 C_2\| \otimes \nabla \tau(s) \|\infty R_{\max}^2 H^2\|_{\text{op}} \leq C_1 C_2 R_{\max} H^2 \Upsilon^2.
\end{aligned}$$

E Different Representation of Theorem 1

Theorem 7. Suppose for $i = 1, 2$, $\|\hat{\pi}^{b,[i]}(A | S) - \pi^b(A | S)\|_{1,\infty} = o_p(1)$, $\|\hat{q}^{[i]}(S, A) - q(S, A)\|_{1,\infty} = o_p(1)$, $\|\hat{\pi}^{b,[i]}(A | S) - \pi^b(A | S)\|_\infty \|\hat{q}^{[i]}(S, A) - q(S, A)\|_\infty = o_p((nh)^{-1/2})$, $nh^5 = O(1)$, $nh \rightarrow \infty$, that $\pi^b(a | s)$, $q(s, a)$ are twice continuously differentiable wrt a for almost all s , and that $\hat{\pi}^{b,[i]}$, $\hat{q}^{[i]}$ are uniformly bounded by a constant. Then, there exists F_n s.t. $\hat{J}^D - J = \mathbb{E}_n \left[\frac{K_h(A - \tau(S))\{R - f_1\}}{f_2} + f_3 \right] - J + o_p((nh)^{-1/2})$, where

$$\mathbb{E}[\mathbb{E}_n \left[\frac{K_h(A - \tau(S))\{R - f_1\}}{f_2} + f_3 \right] - J] = B + o((nh)^{-1/2}),$$

$$\text{var}[\mathbb{E}_n \left[\frac{K_h(A - \tau(S))\{R - f_1\}}{f_2} + f_3 \right]] = V + o((nh)^{-1}),$$

$$B = \mathbb{E}[q^{(2)}(S, \tau(S)) + 2q^{(1)}(S, \tau(S))\pi^{b(1)}(\tau(S)|S)/\pi^b(\tau(S)|S)], \quad V = \frac{1}{nh} \mathbb{E} \left[\frac{\text{var}[R | S, \tau(S)]}{\pi^b(\tau(S) | S)} \right].$$

If additionally $\hat{\pi}^{b,[i]}(a | s)$, $\hat{q}^{[i]}(s, a)$ are twice continuously differentiable wrt a , then the same holds for \hat{J}^K with $B = \mathbb{E}[q^{(2)}(S, \tau(S))]$ and the same V as the above.

We prove Theorem 7. To do that, we prove Theorem 8. The rest of the proof is the same the that of Theorem 1.

Theorem 8. In the case \mathcal{K} , assuming (b), we have

$$\hat{J} = \mathbb{E}_n \left[\frac{K_h(A - \tau(S))\{R - f_1\}}{f_2} + f_3 \right] + o_p((nh)^{-1/2}). \quad (34)$$

In the case \mathcal{D} , the above (34) also holds assuming (a).

E.0.1 Proof of Theorem 8, Case \mathcal{K}

Let us define

$$\phi_1(s, a, r; q, \pi^b) = \frac{K_h(a - \tau(s))\{r - q(s, a)\}}{\pi^b(a|s)} + v(s), \quad v(s) = \int K_h(a - \tau(s))q(s, a)d(a).$$

As in a standard argument (Chernozhukov et al., 2018; Kallus and Uehara, 2020a), what we have to prove is

$$\mathbb{E}[\phi_1(S, A, R; \hat{q}^{[1]}, \hat{\pi}^{b[1]}) - \phi_1(S, A, R; q, \pi^b) | \mathcal{U}_2] = o_p((nh)^{-1/2}), \quad (35)$$

$$\mathbb{E}[\{\phi_1(S, A, R; \hat{q}^{[1]}, \hat{\pi}^{b[1]}) - \phi_1(S, A, R; q, \pi^b)\}^2 | \mathcal{U}_2] = o_p(h^{-1}). \quad (36)$$

Then, the same argument holds for $\phi_1(S, A, R; \hat{q}^{[2]}, \hat{\pi}^{b[2]}) - \phi_1(S, A, R; q, \pi^b)$. The desired statement is concluded since

$$\begin{aligned}
& \mathbb{P}_n[\phi_1(S, A, R; \hat{q}, \hat{\pi}^b)] - \mathbb{P}_n[\phi_1(S, A, R; q, \pi^b)] \\
& = \mathbb{G}_n[\phi_1(S, A, R; \hat{q}, \hat{\pi}^b) - \phi_1(S, A, R; q, \pi^b) | \mathcal{U}_2] \\
& + \mathbb{E}[\phi_1(S, A, R; \hat{q}, \hat{\pi}^b) - \phi_1(S, A, R; q, \pi^b) | \mathcal{U}_2] \\
& = o_p((nh)^{-1/2}) + o_p((nh)^{-1/2}).
\end{aligned} \quad (37)$$

Especially, (37) is $\mathcal{O}_p((nh)^{-1/2})$ since for $\epsilon > 0$,

$$\begin{aligned} & \mathbb{P}(n^{1/2}h^{1/2} \times \mathbb{G}_{\mathcal{U}_2}[\phi_1(S, A, R; \hat{q}^{[1]}, \hat{\pi}^{b[1]}) - \phi_1(S, A, R; q, \pi^b)] > \epsilon | \mathcal{U}_2) \\ & \leq \mathbb{E}[h\{\phi_1(S, A, R; \hat{q}^{[1]}, \hat{\pi}^{b[1]}) - \phi_1(S, A, R; q, \pi^b)\}^2 / \epsilon^2 | \mathcal{U}_2] = \mathcal{O}_p(1). \end{aligned}$$

Therefore, noting $\mathbb{P}(n^{1/2}h^{1/2}\mathbb{G}_{\mathcal{U}_2}[\phi_1(S, A, R; \hat{q}^{[1]}, \hat{\pi}^{b[1]}) - \phi_1(S, A, R; q, \pi^b)] > \epsilon | \mathcal{U}_2)$ is uniformly integrable,

$$\mathbb{P}(n^{1/2}h^{1/2}\mathbb{G}_{\mathcal{U}_2}[\phi_1(S, A, R; \hat{q}^{[1]}, \hat{\pi}^{b[1]}) - \phi_1(S, A, R; q, \pi^b)] > \epsilon | \mathcal{U}_2) = \mathcal{O}_p(1)$$

implies

$$\mathbb{G}_{\mathcal{U}_2}[\phi_1(S, A, R; \hat{q}^{[1]}, \hat{\pi}^{b[1]}) - \phi_1(S, A, R; q, \pi^b)] = \mathcal{O}_p(n^{-1/2}h^{-1/2}).$$

Proof of Eq. (35) and Eq. (36) In this subsection, we remove $\{[1]\}$ for the ease of the notation. To prove (9), we show

$$\mathbb{E} \left[K_h(A - \tau(S)) \left\{ \frac{1}{\hat{\pi}^b(A|S)} - \frac{1}{\pi^b(A|S)} \right\} \{q(S, A) - \hat{q}(S, A)\} \mid \mathcal{U}_2 \right] = \mathcal{O}_p((nh)^{-1/2}).$$

This is proved by

$$\begin{aligned} & \mathbb{E} \left[K_h(A - \tau(S)) \left\{ \frac{1}{\hat{\pi}^b(A|S)} - \frac{1}{\pi^b(A|S)} \right\} \{q(S, A) - \hat{q}(S, A)\} \mid \mathcal{U}_2 \right] \\ &= \int \frac{1}{h} k((a - \tau(s))h^{-1}) \left\{ \frac{1}{\hat{\pi}^b(a|s)} - \frac{1}{\pi^b(a|s)} \right\} \{q(s, a) - \hat{q}(s, a)\} \pi^b(a | s) p(s) ds \\ &= \int k(u) \left\{ \frac{1}{\hat{\pi}^b(\tau(s) + uh | s)} - \frac{1}{\pi^b(\tau(s) + uh | s)} \right\} \{q(s, \tau(s) + uh) - \hat{q}(s, \tau(s) + uh)\} \pi^b(\tau(s) + uh | s) p(s) du \\ &= \int k(u) \left\{ \frac{1}{\hat{\pi}^b(\tau(s)|s)} - \frac{1}{\pi^b(\tau(s)|s)} + uh \left\{ -\frac{\hat{\pi}^{b(1)}(\tau(s)|s)}{\hat{\pi}^{2b}(\tau(s)|s)} + \frac{\pi^{b(1)}(\tau(s)|s)}{\pi^{2b}(\tau(s)|s)} \right\} + \mathcal{O}((uh^2)) \right\} \times \\ & \quad \{q(s, \tau(s)) - \hat{q}(s, \tau(s)) + uh\{q^{(1)}(s, \tau(s)) - \hat{q}^{(1)}(s, \tau(s))\} + \mathcal{O}((uh^2))\} \times \\ & \quad \{\pi^b(\tau(s) | s) + uh\pi^{b(1)}(\tau(s) | s) + \mathcal{O}((uh^2))\} p(s) du \\ &= M_2(k) \int \left\{ \frac{1}{\hat{\pi}^b(\tau(s) | s)} - \frac{1}{\pi^b(\tau(s) | s)} \right\} \{q(s, \tau(s)) - \hat{q}(s, \tau(s))\} \pi^b(\tau(s) | s) p(s) ds + \mathcal{O}(h^2) \times \mathcal{O}_p(1) \\ &= \mathcal{O}_p((nh)^{-1/2}). \end{aligned}$$

In the last line, we use the assumptions that $q(a, x), \pi^b(a|x), \hat{q}(a, x), \hat{\pi}^b(a|x)$ are C^2 -functions wrt actions, and

$$\begin{aligned} & \left\| \frac{1}{\hat{\pi}^b(A | S)} - \frac{1}{\pi^b(A | S)} \right\|_{\infty} \|q(S, A) - \hat{q}(S, A)\|_{\infty} = \mathcal{O}_p(n^{-1/2}h^{-1/2}), \\ & \left\| \frac{1}{\hat{\pi}^b(A | S)} - \frac{1}{\pi^b(A | S)} \right\|_{1, \infty} = \mathcal{O}_p(1), \|\hat{q}(S, A) - \hat{q}(S, A)\|_{1, \infty} = \mathcal{O}_p(1), \\ & \mathcal{O}(h^2) \times \mathcal{O}_p(1) = \mathcal{O}_p((nh)^{-1/2}), nh^5 = \mathcal{O}(1). \end{aligned}$$

In addition, Eq. (36) is proved since

$$\begin{aligned} & \mathbb{E}[\{\phi_1(S, A, R; \hat{q}^{[1]}, \hat{\pi}^{b[1]}) - \phi_1(S, A, R; q, \pi^b)\}^2 | \mathcal{U}_2] \\ & \lesssim \mathbb{E} \left[K_h(A - \tau(S))^2 \left\{ \frac{1}{\hat{\pi}^b(A | S)} - \frac{1}{\pi^b(A | S)} \right\}^2 \{q(S, A) - \hat{q}(S, A)\}^2 \mid \mathcal{U}_2 \right] \\ & + \mathbb{E} \left[K_h(A - \tau(S))^2 \left\{ \frac{1}{\hat{\pi}^b(A | S)} - \frac{1}{\pi^b(A | S)} \right\}^2 \{R - q(S, A)\}^2 \mid \mathcal{U}_2 \right] \\ & + \mathbb{E} \left[\frac{K_h(A - \tau(S))^2}{\pi^b(A | S)^2} \{q(S, A) - \hat{q}(S, A)\}^2 \mid \mathcal{U}_2 \right] + \mathbb{E} [\{\hat{v}(S) - v(S)\}^2 | \mathcal{U}_2] \\ & \lesssim h^{-1} \max \left\{ \|\hat{q}(S, \tau(S)) - q(S, \tau(S))\|_2^2, \left\| \frac{1}{\hat{\pi}^b(\tau(S)|S)} - \frac{1}{\pi^b(\tau(S)|S)} \right\|_2^2 \right\} + \mathcal{O}(1) = \mathcal{O}_p(h^{-1}). \end{aligned}$$

E.0.2 Proof of Theorem 8, Case \mathcal{D}

Essentially, the same proof is seen in Colangelo and Lee (2019). For completeness, we also write the proof here with our notation. Let us define

$$\phi_2(s, a, r; q, \pi^b) = \frac{K_h(a - \tau(s))\{r - q(s, \tau(s))\}}{\pi^b(a|s)} + q(S, \tau(S)).$$

As in a standard argument similar to the case \mathcal{K} , what we have to prove is

$$\mathbb{E}[\phi_2(S, A, R; \hat{q}^{[1]}, \hat{\pi}^{b[1]}) - \phi_2(S, A, R; q, \pi^b) | \mathcal{U}_2] = o_p((nh)^{-1/2}), \quad (38)$$

$$\mathbb{E}[\{\phi_2(S, A, R; \hat{q}^{[1]}, \hat{\pi}^{b[1]}) - \phi_2(S, A, R; q, \pi^b)\}^2 | \mathcal{U}_2] = o_p(h^{-1}). \quad (39)$$

In this subsection, we remove $\{[1]\}$ for the ease of the notation.

Eq. (35) is proved since

$$\begin{aligned} & \mathbb{E}[\phi_2(S, A, R; \hat{q}, \hat{\pi}^b) - \phi_2(S, A, R; q, \pi^b) | \mathcal{U}_2] \\ &= \mathbb{E} \left[K_h(A - \tau(S)) \left\{ \frac{1}{\hat{\pi}^b(\tau(S) | S)} - \frac{1}{\pi^b(\tau(S) | S)} \right\} \{q(S, \tau(S)) - \hat{q}(S, \tau(S))\} | \mathcal{U}_2 \right] + \end{aligned} \quad (40)$$

$$+ \mathbb{E} \left[K_h(A - \tau(S)) \left\{ \frac{1}{\hat{\pi}^b(\tau(S) | S)} - \frac{1}{\pi^b(\tau(S) | S)} \right\} \{R - q(S, \tau(S))\} | \mathcal{U}_2 \right] \quad (41)$$

$$\begin{aligned} &+ \mathbb{E} \left[\frac{K_h(A - \tau(S))}{\pi^b(\tau(S) | S)} \{q(S, \tau(S)) - \hat{q}(S, \tau(S))\} + \hat{q}(S, \tau(S)) - q(S, \tau(S)) | \mathcal{U}_2 \right] \quad (42) \\ &= o_p((nh)^{-1/2}) + o_p(1) \times \mathcal{O}(h^2) + o_p(1) \times \mathcal{O}(h^2) = o_p(nh)^{-1/2}. \end{aligned}$$

Here, we use the facts that (40) is $o_p((nh)^{-1/2})$, (41) is $o_p(1) \times \mathcal{O}(h^2)$, (42) is $o_p(1) \times \mathcal{O}(h^2)$, which we will prove soon. In the last line, we use $nm^5 = \mathcal{O}(1)$. From now on, we prove (41) is $o_p(1) \times \mathcal{O}(h^2)$:

$$\begin{aligned} & \mathbb{E} \left[K_h(A - \tau(S)) \left\{ \frac{1}{\hat{\pi}^b(\tau(S) | S)} - \frac{1}{\pi^b(\tau(S) | S)} \right\} \{R - q(S, \tau(S))\} | \mathcal{U}_2 \right] \\ &= \mathbb{E} \left[\left\{ \frac{1}{\hat{\pi}^b(\tau(S) | S)} - \frac{1}{\pi^b(\tau(S) | S)} \right\} \{ \mathbb{E}[K_h(A - \tau(S))q(S, A) | S] - K_h(A - \tau(S))q(S, \tau(S)) \} | \mathcal{U}_2 \right] \\ &= \mathbb{E} \left[\left\{ \frac{1}{\hat{\pi}^b(\tau(S) | S)} - \frac{1}{\pi^b(\tau(S) | S)} \right\} \{ \mathcal{O}(h^2) \} | \mathcal{U}_2 \right] \\ &= o_p(1) \times \mathcal{O}(h^2). \end{aligned}$$

More specifically,

$$\begin{aligned} & \mathbb{E}[K_h(A - \tau(S))\{q(S, A) - q(S, \tau(S))\} | S] = \int \frac{1}{h} k \left\{ \frac{a - \tau(s)}{h} \right\} \pi^b(a|s) \{q(s, a) - q(s, \tau(s))\} da \\ &= \int k(u) \pi^b(\tau(s) + uh | s) \{q(s, \tau(s) + uh) - q(s, \tau(s))\} du \\ &= \int k(u) \{\pi^b(\tau(s) | s) + \mathcal{O}(uh)\} \{uhq^{(1)}(s, \tau(s)) + \mathcal{O}(h^2)\} du = \mathcal{O}(h^2). \end{aligned}$$

noting $\int uk(u)du = 0$. Next, we prove (15) is $o_p(1) \times \mathcal{O}(h^2)$:

$$\begin{aligned} & \mathbb{E} \left[\frac{K_h(A - \tau(S))}{\pi^b(\tau(S) | S)} \{q(S, \tau(S)) - \hat{q}(S, \tau(S))\} + \hat{q}(S, \tau(S)) - q(S, \tau(S)) | \mathcal{U}_2 \right] \\ &= \mathbb{E} \left[\left\{ \frac{K_h(A - \tau(S))}{\pi^b(\tau(S) | S)} - 1 \right\} \{q(S, \tau(S)) - \hat{q}(S, \tau(S))\} | \mathcal{U}_2 \right] \\ &= \mathbb{E} [\{ \mathcal{O}(h^2) \} \{q(S, \tau(S)) - \hat{q}(S, \tau(S))\} | \mathcal{U}_2] = o_p(1) \times \mathcal{O}(h^2). \end{aligned}$$

Eq. (36) is similarly proved as in the case \mathcal{K} .