

---

# Future-Dependent Value-Based Off-Policy Evaluation in POMDPs

---

**Masatoshi Uehara** \*

Genentech

uehara.masatoshi@gene.com

**Haruka Kiyohara** †

Cornell University

hk844@cornell.edu

**Andrew Bennett**

Morgan Stanley ‡

Andrew.Bennett@morganstanley.com

**Victor Chernozhukov**

MIT

vchern@mit.edu

**Nan Jiang**

UIUC

nanjiang@illinois.edu

**Nathan Kallus**

Cornell University

kallus@cornell.edu

**Chengchun Shi**

LSE

c.shi7@lse.ac.uk

**Wen Sun**

Cornell University

ws455@cornell.edu

## Abstract

We study off-policy evaluation (OPE) for partially observable MDPs (POMDPs) with general function approximation. Existing methods such as sequential importance sampling estimators suffer from the curse of horizon in POMDPs. To circumvent this problem, we develop a novel model-free OPE method by introducing future-dependent value functions that take future proxies as inputs and perform a similar role to that of classical value functions in fully-observable MDPs. We derive a new off-policy Bellman equation for future-dependent value functions as conditional moment equations that use history proxies as instrumental variables. We further propose a minimax learning method to learn future-dependent value functions using the new Bellman equation. We obtain the PAC result, which implies our OPE estimator is close to the true policy value under Bellman completeness, as long as futures and histories contain sufficient information about latent states. Our code is available at <https://github.com/aiueola/neurips2023-future-dependent-ope>.

## 1 Introduction

Reinforcement learning (RL) has demonstrated success when it is possible to interact with the environment and collect data in an adaptive manner. However, for domains such as healthcare, education, robotics, and social sciences, online learning can be risky and unethical [e.g., [KHS<sup>+</sup>15](#), [KL19](#), [TDHL19](#), [SRRH21](#)]. To address this issue, a variety of offline RL methods have recently been developed for policy learning and evaluation using historical data in these domains [see [LKTF20](#), for an overview]. In this paper, we focus on the off-policy evaluation (OPE) problem, which concerns estimating the value of a new policy (called the evaluation policy) using offline log data that was

---

\*Co-first author. This work was done at Cornell University

†Co-first author. This work was done at Tokyo Institute of Technology

‡This work was done at Cornell University

generated under a different policy (called the behavior policy) [TB16, JL16, TDHL19, CQ22]. OPE is especially useful in the aforementioned high-stakes domains.

We focus on off-policy evaluation (OPE) in partially observable Markov decision processes (POMDPs). Partial observability is a common phenomenon in practical applications [Cas98, KLC98, XLSS20], and it poses a serious challenge for sample-efficient learning. Most existing OPE methods are developed for MDPs and their statistical learning guarantees rely crucially on the Markov assumption. To extend these methods to POMDPs, one may use the entire history of observations as a *state* to satisfy Markovianity. This allows us to apply sequential importance sampling [SIS, Pre00, JSLZ19, HW21] or its variant [e.g., sequential doubly robust (SDR) methods RRS98, MvdLRG01, JL16, TB16, FCG18, BMVVDL19] to the transformed data for valid policy evaluation. An alternative approach is to employ value-based methods by constructing a history-dependent Q-function that takes the entire history as input. These value functions can be estimated via fitted-Q evaluation [FQE, EGW05, MS08, LUY19] or minimax methods [ASM08, CJ19, FLL19, NCDL19, UHJ20, UIJ<sup>+</sup>21, ZW22]. However, all aforementioned estimators suffer from the curse of horizon, as their estimation errors grow exponentially with respect to the (effective) horizon and become prohibitively large in long-horizon settings [LLTZ18, KU20, KU22].

The goal of this paper is to devise practical and efficient OPE methods in large partially observable environments, while breaking the curse of horizon. As a first step, we restrict our attention to the evaluation of short memory-based policies that take several recent observations instead of the entire history as inputs. Short memory-based policies are widely used in practice, since maintaining the whole previous observation leads to a well-known computational challenge dubbed as the curse of history [PGT06, GMR22]. For example, in the deep Q-network (DQN) algorithm, four game frames are stacked together to train the optimal policy in Atari [MKS<sup>+</sup>13], while Open AI Five uses a window of length 16 [BBC<sup>+</sup>19]. Even when considering the evaluation of policies with short-memory or memory-less characteristics (i.e., policies that depend solely on current observations), naïve methods such as SIS, SDR, and FQE still suffer from the curse of horizon<sup>4</sup>.

Our proposal takes a model-free perspective and introduces a new concept called “future-dependent value functions”. The proposed value function does not involve latent states which are unobservable in POMDPs. Nor does it rely on the entire data history. It performs the role of standard value functions by incorporating future observations as proxies for latent states. We demonstrate that these future-dependent value functions satisfy a new off-policy Bellman equation for POMDPs, which involves past observations to approximate latent states. To estimate these future-dependent value functions from offline data using the new Bellman equation, we propose minimax learning methods that accommodate various function approximation, including deep neural networks, RKHS, and linear models. Notably, our method extends the classical LSTD [LP03] developed in MDPs to the POMDP setting when using linear models. To summarize, our contributions are as follows:

- We propose a novel model-free method for OPE that leverages future-dependent value functions. Our key idea is to utilize future and historical observations as proxies for latent states.
- We derive a new Bellman equation for learning future-dependent value functions. The proposed estimator can accommodate various types of function approximations.
- We provide PAC guarantees to demonstrate that our method can address the curse of horizon, and conduct numerical experiments to showcase its superiority over existing methods.

Note that similar concepts of future-dependent value functions have recently been introduced in the context of OPE in confounded POMDPs [SUHJ22] and online RL in POMDPs [USL<sup>+</sup>22]. However, we focus on OPE in *non-confounded* POMDPs. Further detailed discussions can be found in Section 1.1.

## 1.1 Related Literature

**OPE in confounded POMDPs.** OPE with unmeasured confounders has been actively studied [ZB16, NKYB20, TSM20, WYW20, KZ20, LFY<sup>+</sup>21, NJ21, GCZ<sup>+</sup>22, SZY<sup>+</sup>22, XZS<sup>+</sup>22, LMWY22, CWWY23, BSZ23]. Among them, [TSM20] adopted the POMDP model to formulate the OPE problem in the presence of unmeasured confounders. They borrowed ideas from the causal inference literature on double negative control adjustment [MGT18, MST18, CPS<sup>+</sup>20,

---

<sup>4</sup>Notice that even when restricting to short-memory policies, it is necessary to use the entire history instead of the short memory as the state to meet the Markov property.

[KMU21, SMNTT20] to derive consistent value estimators in tabular settings. Later, [BKLM21] and [SUHJ22] extend their proposal by incorporating confounding bridge functions, thereby enabling general function approximation. (see the difference between these bridge functions and the proposed future-dependent value function in Section B). However, these methods do not apply to our unconfounded POMDP setting. This is because these methods require the behavior policy to *only* depend on the latent state to ensure certain conditional independence assumptions. These assumptions are violated in our setting, where the behavior policy may depend on the observation – a common scenario in practical applications. In addition, we show that in the unconfounded setting, it is feasible to leverage multi-step future observations to relax specific rank conditions in their proposal, which is found to be difficult in the confounded setting [NJ21]. Refer to Example 1.

**Online learning in POMDPs.** In the literature, statistically efficient online learning algorithms with polynomial sample complexity have been proposed in tabular POMDPs [GDB16, ALA16, JKKL20, LCSJ22], linear quadratic Gaussian setting (LQG) [LAHA21, SSH20], latent POMDPs [KECM21] and reactive POMDPs/PSRs [KAL16, JKA<sup>+</sup>17]. All the methods above require certain model assumptions. To provide a more unified framework, researchers have actively investigated online learning in POMDPs with general function approximation, as evidenced by a vast body of work [ZUSL22, LNSJ22, CBM22, ZXZ<sup>+</sup>22]. As the most relevant work, by leveraging future-dependent value functions, [USL<sup>+</sup>22] propose an efficient PAC RL algorithm in the online setting. They require the existence of future-dependent value functions and a low-rank property of the Bellman loss. Instead, in our approach, while we similarly require the existence of future-dependent value functions, we do not require the low-rank property of Bellman loss. Instead, we use the invertibility condition in Theorem 1, i.e., we have informative history proxies as instrumental variables. As a result, the strategies to ensure efficient PAC RL in POMDPs are quite different in the offline and online settings.

**Learning dynamical systems via spectral learning.** There is a rich literature on POMDPs by representing them as predictive state representations (PSRs) [Jae00, LS01, RGT04, SJR04]. PSRs are models of dynamical systems that represent the state as a vector of predictions about future observable events. They are appealing because they can be defined directly from observable data without inferring hidden variables and are more expressive than Hidden Markov Models (HMMs) and POMDPs [SJR04]. Several spectral learning algorithms have been proposed for PSRs [BSG11, HKZ12, BGG13], utilizing conditional mean embeddings or Hilbert space embeddings. These approaches provide closed-form learning solutions, which simplifies computation compared to EM-type approaches prone to local optima and non-singularity issues. While these methods have demonstrated success in real-world applications [BSG11] and have seen subsequent improvements [KJS15, HDG15, VSH<sup>+</sup>16, DHB<sup>+</sup>17], it is still unclear how to incorporate more flexible function approximations such as neural networks in a model-free end-to-end manner with a valid PAC guarantee. Our proposed approach not only offers a new model-free method for OPE, but it also incorporates these model-based spectral learning methods when specialized to dynamical system learning with linear models (see Section F).

**Planning in POMDPs.** There is a large amount of literature on planning in POMDPs. Even if the models are known, exact (or nearly optimal) planning in POMDPs is known to be NP-hard in the sense that it requires exponential time complexity with respect to the horizon [PT87, BDRS96]. This computational challenge is often referred to as the *curse of history* [PGT06]. A natural idea to mitigate this issue is to restrict our attention to short-memory policies [AYA18, MY20, KY22, GMR22]. Practically, a short-memory policy can achieve good empirical performance, as mentioned earlier. This motivates us to focus on evaluating short-memory policies in this paper.

## 2 Preliminaries

In this section, we introduce the model setup, describe the offline data and present some notations.

**Model setup.** We consider an infinite-horizon discounted POMDP  $\mathcal{M} = \langle \mathcal{S}, \mathcal{A}, \mathcal{O}, r, \gamma, \mathbb{O}, \mathbb{T} \rangle$  where  $\mathcal{S}$  denotes the state space,  $\mathcal{A}$  denotes the action space,  $\mathcal{O}$  denotes the observation space,  $\mathbb{O} : \mathcal{S} \rightarrow \Delta(\mathcal{O})$  denotes the emission kernel (i.e., the conditional distribution of the observation given the state),  $\mathbb{T} : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$  denotes the state transition kernel (i.e., the conditional distribution of the next state given the current state-action pair),  $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$  denotes the reward function, and  $\gamma \in [0, 1)$  is the discount factor. All the three functions  $\mathbb{T}, r, \mathbb{O}$  are unknown to the learner.

For simplicity, we consider the evaluation of memory-less policies  $\pi : \mathcal{O} \rightarrow \Delta(\mathcal{A})$  that solely depend on the current observation  $O$  in the main text. The extension to policies with short-memory is discussed in Section C.

Next, given a memory-less evaluation policy  $\pi^e$ , we define the parameter of interest, i.e., the policy value. Following a policy  $\pi^e$ , the data generating process can be described as follows. First,  $S_0$  is generated according to some initial distribution  $\nu_S \in \Delta(\mathcal{S})$ . Next, the agent observes  $O_0 \sim \mathbb{O}(\cdot | S_0)$ , executes the initial action  $A_0 \sim \pi^e(\cdot | O_0)$ , receives a reward  $r(S_0, A_0)$ , the environment transits to the next state  $S_1 \sim \mathbb{T}(\cdot | S_0, A_0)$ , and this process repeats. Our objective is to estimate

$$J(\pi^e) := \mathbb{E}_{\pi^e} [\sum_{t=0}^{\infty} \gamma^t R_t],$$

where the expectation  $\mathbb{E}_{\pi^e}$  is taken by assuming the data trajectory follows the evaluation policy  $\pi^e$ .

**Offline data.** We convert the trajectory data generated by a behavior policy  $\pi^b : \mathcal{O} \rightarrow \Delta(\mathcal{A})$ , into a set of history-observation-action-reward-future transition tuples (denoted by  $\mathcal{D}_{\text{tra}}$ ) and a set of initial observations (denoted by  $\mathcal{D}_{\text{ini}}$ ). The first data subset enables us to learn the reward, emission and transition kernels whereas the second data subset allows us to learn the initial observation distribution.

Specifically, the dataset  $\mathcal{D}_{\text{tra}}$  consists of  $N$  data tuples  $\{(H^{(i)}, O^{(i)}, A^{(i)}, R^{(i)}, F'^{(i)})\}_{i=1}^N$ . We use  $(H, O, A, R, F')$  to denote a generic history-observation-action-reward-future tuple where  $H$  denotes the  $M_H$ -step historical observations obtained prior to the observation  $O$  and  $F'$  denotes the  $M_F$ -step future observations after  $(O, A)$  for some integers  $M_H, M_F \geq 1$ . Specifically, at a given time step  $t$  in the data trajectory, we use  $(O, A, R)$  to denote  $(O_t, A_t, R_t)$ , and set

$$H = (O_{t-M_H:t-1}, A_{t-M_H:t-1}) \text{ and } F' = (O_{t+1:t+M_F}, A_{t+1:t+M_F-1}).$$

We additionally set  $F = (O_{t:t+M_F-1}, A_{t:t+M_F-2})$ . Note we use the prime symbol ' to represent the next time step. See Figure 1 for details when we set  $t = 0$ .

Throughout this paper, we use uppercase letters such as  $(H, S, O, A, R, S', F')$  to denote *random variables* in the offline data, and lowercase letters such as  $(h, s, o, a, r, s', f')$  to denote their *realizations*, unless stated otherwise. To simplify the presentation, we assume the stationarity of the environment, i.e., the marginal distributions of  $(H, S, F)$  and  $(H', S', F')$  are identical.

The dataset  $\mathcal{D}_{\text{ini}}$  consists of  $N'$  data tuples  $\{O_{0:M_F-1}^{(i)}, A_{0:M_F-2}^{(i)}\}_{i=1}^{N'}$  which is generated as follows:  $S_0 \sim \nu_S$ ,  $O_0 \sim \mathbb{O}(\cdot | S_0)$ ,  $A_0 \sim \pi^b(\cdot | O_0)$ ,  $S_1 \sim \mathbb{T}(\cdot | S_0, A_0)$ ,  $\dots$ , until we observe  $O_{M_F-1}^{(i)}$  and  $A_{M_F-1}^{(i)}$ . We denote its distribution over  $\mathcal{F} = \mathcal{O}^{M_F} \times \mathcal{A}^{M_F-1}$  by  $\nu_{\mathcal{F}}(\cdot)$ .

**Remark 1** (Standard MDPs). *Consider the setting where  $S = O$  and  $M_H = 0$ ,  $M_F = 1$ . In that case, we set  $H$  to  $S$  instead of histories. Then,  $\mathcal{D}_{\text{tra}} = \{O^{(i)}, A^{(i)}, R^{(i)}, O'^{(i)}\}$ . We often assume  $\nu_{\mathcal{O}}$  ( $\nu_{\mathcal{F}}$  in our setting) is known. This yields the standard OPE setting in MDPs [CJ19, UHJ20].*

**Notation.** We streamline the notation as follows. We define a state value function under  $\pi^e$ :  $V^{\pi^e}(s) := \mathbb{E}_{\pi^e} [\sum_{k=0}^{\infty} \gamma^k R_k | S_0 = s]$  for any  $s \in \mathcal{S}$ . Let  $d_t^{\pi^e}(\cdot)$  be the marginal distribution of  $S_t$  under the policy  $\pi^e$ . Then, we define the discounted occupancy distribution  $d_{\pi^e}(\cdot) := (1 - \gamma) \sum_{t=0}^{\infty} \gamma^t d_t^{\pi^e}(\cdot)$ . We denote the domain of  $F, H$  by  $\mathcal{F} = (\mathcal{O} \times \mathcal{A})^{M_F-1} \times \mathcal{O}$ ,  $\mathcal{H} = (\mathcal{O} \times \mathcal{A})^{M_H}$ , respectively. The notations  $\mathbb{E}, \mathbb{E}_{\mathcal{D}}$  (without any subscripts) represent the population or sample average over the offline data  $\mathcal{D} = \mathcal{D}_{\text{tra}} \cup \mathcal{D}_{\text{ini}}$ , respectively. We denote the distribution of offline data by  $P_{\pi^b}(\cdot)$ . We define the marginal density ratio  $w_{\pi^e}(S) := d_{\pi^e}(S)/P_{\pi^b}(S)$ . Given a matrix  $C$ , we denote its Moore-Penrose inverse by  $C^+$  and its smallest singular value by  $\sigma_{\min}(C)$ . For a given integer  $m > 0$ , let  $I_m$  denote an  $m \times m$  identity matrix. Let  $\otimes$  denote outer product and  $[T] = \{0, \dots, T\}$  for any integer  $T > 0$ .

### 3 Identification via Future-dependent Value Functions

In this section, we present our proposal to identify policy values under partial observability by introducing future-dependent value functions. We remark that the target policy's value is identifiable from the observed data via SIS or SDR. Nonetheless, as commented earlier in Section 1, these methods suffer from the curse of horizon. Here, we propose an alternative identification approach that




Figure 1: Case with  $M_H = 1$ ,  $M_F = 2$ . An action  $A$  is generated depending on  $O$ . The extension to memory-based policy is discussed in Section C.

can possibly circumvent the curse of horizon. It serves as a building block to motivate the proposed estimation methods in Section 4.

In fully observable MDPs, estimating a policy’s value is straightforward using the value function-based method  $J(\pi^e) = \mathbb{E}_{s \sim \nu_S}[V^{\pi^e}(s)]$ . However, in partial observable environments where the latent state is inaccessible, the state value function  $V^{\pi^e}(s)$  is unidentifiable. To address this challenge, we propose the use of *future-dependent value functions* that are defined based on observed variables and serve a similar purpose to state value functions in MDPs.

**Definition 1** (Future-dependent value functions). *Future-dependent value functions  $g_V \in [\mathcal{F} \rightarrow \mathbb{R}]$  are defined such that the following holds almost surely,*

$$\mathbb{E}[g_V(F) \mid S] = V^{\pi^e}(S).$$

*Recall that the expectation is taken with respect to the offline data generated by  $\pi^b$ .*

Crucially, the future-dependent value functions mentioned above may not always exist, and they need not be unique. Existence is a vital assumption in our framework, although we don’t insist on uniqueness. We will return to the topics of existence and uniqueness after demonstrating their relevance in offline policy evaluation.

Hereafter, we explain the usefulness of future-dependent value functions. Future-dependent value functions are defined as embeddings of value functions on latent states onto multi-step futures. Notice that  $J(\pi^e) = \mathbb{E}_{s \sim \nu_S}[V^{\pi^e}(s)] = \mathbb{E}_{f \sim \nu_{\mathcal{F}}}[g_V(f)]$ . These future-dependent value functions are useful in the context of OPE as they enable us to accurately estimate the final policy value. However, the future-dependent value function itself cannot be identified since its definition relies on unobserved states. To overcome this challenge, we introduce a learnable counterpart called the *learnable future-dependent value function*. This learnable version is defined based on observed quantities and thus can be identified.

**Definition 2** (Learnable future-dependent value functions). *Define  $\mu(O, A) := \pi^e(A \mid O)/\pi^b(A \mid O)$ . Learnable future-dependent value functions  $b_V \in [\mathcal{F} \rightarrow \mathbb{R}]$  are defined such that the following holds almost surely,*

$$0 = \mathbb{E}[\mu(O, A)\{R + \gamma b_V(F')\} - b_V(F) \mid H]. \quad (1)$$

*Recall that the expectation is taken with respect to the offline data generated by  $\pi^b$ . We denote the set of solutions by  $\mathcal{B}_V$ .*

To motivate this definition, we recall that the off-policy Bellman equation in MDPs [DNP<sup>+</sup>14] can be expressed as follows:

$$V^{\pi^e}(S) = \mathbb{E}[\mu(O, A)(R + \gamma V^{\pi^e}(S')) \mid S]. \quad (2)$$

Then, by the definition of future-dependent value functions and certain conditional independence relations ( $F' \perp (O, A) \mid S'$ ), we obtain that

$$0 = \mathbb{E}[\mu(O, A)\{R + \gamma g_V(F')\} - g_V(F) \mid S]. \quad (3)$$

Since  $H \perp (O, A, F') \mid S$ , taking another conditional expectation given  $H$  on both sides yields that

$$0 = \mathbb{E}[\mu(O, A)\{R + \gamma g_V(F')\} - g_V(F) \mid H]. \quad (4)$$

Therefore, (4) can be seen as an off-policy Bellman equation for future-dependent value functions, analogous to the Bellman equation (2) in MDPs. Based on the above discussion, we present the following lemma.

**Lemma 1.** *Future-dependent value functions are learnable future-dependent value functions.*

**Remark 2** (Non-stationary case). *When the offline data is non-stationary, i.e., the pdfs of  $(H, S)$  and  $(H', S')$  are different, we need to additionally require  $\mathbb{E}[g_V(F') \mid S'] = V^{\pi^e}(S')$  in the definition.*

**Remark 3** (Comparisons with related works). *Similar concepts have been recently proposed in the context of confounded POMDPs [see e.g., SUHJ22, BK21]. However, our proposal significantly differs from theirs. First, their confounded setting does not cover our unconfounded setting because their behavior policies  $\pi^b$  is not allowed to depend on current observations as mentioned in Section 1.1. Secondly, their proposal overlooks the significant aspect of incorporating multi-step future observations, which plays a pivotal role in facilitating the existence of future-dependent value functions, as will be discussed in Example 1. For a detailed discussion, refer to Section B.*

---

**Algorithm 1** Minimax OPE on POMDPs

---

**Require:** Dataset  $\mathcal{D}$ , function classes  $\mathcal{Q} \subset [\mathcal{F} \rightarrow \mathbb{R}]$ ,  $\Xi \subset [\mathcal{H} \rightarrow \mathbb{R}]$ , hyperparameter  $\lambda \geq 0$   
1:  $\hat{b}_V = \operatorname{argmin}_{q \in \mathcal{Q}} \max_{\xi \in \Xi} \mathbb{E}_{\mathcal{D}_{\text{tra}}}[\{\mu(A, O)\{R + \gamma q(F')\} - q(F)\}\xi(H) - \lambda \xi^2(H)]$ .  
2: **return**  $\hat{J}_{\text{VM}} = \mathbb{E}_{\mathcal{D}_{\text{ini}}}[\hat{b}_V(f)]$

---

Finally, we present a theorem to identify the policy value.

**Theorem 1** (Identification). *Suppose (1a) the existence of learnable future-dependent value functions (need not be unique); (1b) the invertibility condition, i.e., any  $g : \mathcal{S} \rightarrow \mathbb{R}$  that satisfies  $\mathbb{E}[g(S) | H] = 0$  must also satisfy  $g(S) = 0$  (i.e.,  $g(s) = 0$  for almost every  $s$  that belongs to the support of  $S$ ), (1c) the overlap condition  $w_{\pi^e}(S) := d_{\pi^e}(S)/P_{\pi^b}(S) < \infty$ ,  $\mu(O, A) < \infty$ . Then, for any  $b_V \in \mathcal{B}_V$ ,*

$$J(\pi^e) = \mathbb{E}_{f \sim \nu_{\mathcal{F}}}[b_V(f)]. \quad (5)$$

We assume three key conditions: the observability condition (i.e.,  $\mathcal{B}_V \neq \emptyset$ ), the invertibility condition, and the overlap condition. We call Condition (1a) the observability condition since it is reduced to the well-known concept of observability in the LQG control theory, as we will see in Section D. While Condition (1a) itself is concerned with learnable future-dependent value functions, it is implied by the existence of *unlearnable* future-dependent value functions according to Lemma 1 which can be verified using Picard’s theorem in functional analysis [CFR07]. In general, the observability condition requires the future proxy  $F$  to contain sufficient information about  $S$  and is likely to hold when  $F$  consists of enough future observations. We will see more interpretable conditions in the tabular POMDPs (Example 1) and POMDPs with Hilbert space embeddings (HSE-POMDPs) where the underlying dynamical systems have linear conditional mean embeddings (Example 5 in Section D).

The invertibility condition is imposed to ensure that a learnable future-dependent value function  $b_V$  satisfies (3) (note that right hand side of Eq. 3 is a function of  $Z, S$  instead of  $H$ ). Again, we will present more interpretable conditions in the tabular and linear settings below and in Section D. Roughly speaking, it requires  $H$  to retain sufficient information about  $S$ . In that sense, the history proxy serves as an instrumental variable (IV), which is widely used in economics [Hor11, New13]<sup>5</sup>.

Finally, the overlap condition is a standard assumption in OPE [UIJ<sup>+</sup>21].

**Example 1** (Tabular Setting). *In the tabular case, abstract assumptions in Theorem 1 are reduced to certain rank assumptions. We first define  $\mathcal{S}_b = \{s \in \mathcal{S} : P_{\pi^b}(s) > 0\}$ . We define a matrix  $\Pr_{\pi^b}(\mathbf{F} | \mathbf{S}_b) \in \mathbb{R}^{|\mathcal{F}| \times |\mathcal{S}|}$  whose  $(i, j)$ -th element is  $\Pr_{\pi^b}(F = x_i | S = x'_j)$  where  $x_i$  is the  $i$ th element in  $\mathcal{F}$ , and  $x'_j$  is the  $j$ th element in  $\mathcal{S}_b$ . We similarly define another matrix  $\Pr_{\pi^b}(\mathbf{S}_b, \mathbf{H})$  whose  $(i, j)$ -th element is  $\Pr_{\pi^b}(S = x'_i, H = x''_j)$  where  $x''_j$  denotes the  $j$ th element in  $\mathbf{H}$ .*

**Lemma 2** (Sufficient conditions for observability and invertibility). *(a) When  $\operatorname{rank}(\Pr_{\pi^b}(\mathbf{F} | \mathbf{S}_b)) = |\mathcal{S}_b|$ , future-dependent value functions exist. Then, from Lemma 1, learnable future-dependent value functions exist. (b) The invertibility is satisfied when  $\operatorname{rank}(\Pr_{\pi^b}(\mathbf{S}_b, \mathbf{H})) = |\mathcal{S}_b|$ .*

*The first two conditions require that the cardinalities of  $\mathcal{F}$  and  $\mathcal{H}$  must be greater than or equal to  $\mathcal{S}_b$ , respectively. The proof of Lemma 2 is straightforward, as integral equations reduce to matrix algebra in the tabular setting. We note that similar conditions have been assumed in the literature on HMMs and POMDPs [SBS<sup>+</sup>10, BSG11, BGG13]. In particular,  $\Pr_{\pi^b}(\mathbf{O} | \mathbf{S}_b) = |\mathcal{S}_b|$  has been imposed in previous works [NJ21, SUHJ22] in the context of confounded POMDPs. Nonetheless, our condition  $\Pr_{\pi^b}(\mathbf{F} | \mathbf{S}_b) = |\mathcal{S}_b|$  is strictly weaker when  $\mathcal{F}$  includes multi-step future observations, demonstrating the advantage of incorporating multi-step future observations compared to utilizing only the current observation. A more detailed discussion can be found in Appendix D.1.*

## 4 Estimation with General Function Approximation

In this section, we demonstrate how to estimate the value of a policy based on the results presented in Section 3. We begin by outlining the proposed approach for estimating  $b_V(\cdot)$ . The key observation is

---

<sup>5</sup>The observation that history can serve as an instrumental variable in POMDPs is mentioned in [HDG15, VSH<sup>+</sup>16]. However, these works aim to learn the system dynamics instead of policy evaluation or learning. Hence, concepts like future-dependent value functions do not appear in these works.

that it satisfies  $\mathbb{E}[L(b_V, \xi)] = 0$  for any  $\xi : \mathcal{H} \rightarrow \mathbb{R}$  where  $L(q, \xi)$  is defined as

$$L(q, \xi) := \{\mu(A, O)\{R + \gamma q(F')\} - q(F)\}\xi(H)$$

for  $q : \mathcal{F} \rightarrow \mathbb{R}$  and  $\xi : \mathcal{H} \rightarrow \mathbb{R}$ . Given some constrained function classes  $\mathcal{Q} \subset [\mathcal{F} \rightarrow \mathbb{R}]$  and  $\Xi \subset [\mathcal{H} \rightarrow \mathbb{R}]$  and a hyperparameter  $\lambda \geq 0$ , the estimator is computed according to Line 1 of [Algorithm 1](#). When the realizability  $\mathcal{B}_V \cap \mathcal{Q} \neq \emptyset$  holds and  $\Xi$  is unconstrained, we can easily show that the population-level minimizers  $\operatorname{argmin}_{q \in \mathcal{Q}} \max_{\xi \in \Xi} \mathbb{E}[L(q, \xi) - 0.5\lambda\xi^2(H)]$  are all learnable future-dependent value functions. This is later formalized in [Theorem 2](#).

We can use any function classes such as neural networks, RKHS, and random forests to parameterize  $\mathcal{Q}$  and  $\Xi$ . Here, the function class  $\Xi$  plays a critical role in measuring how  $q$  deviates from the ground truth  $b_V$ . The hyperparameter  $\lambda$  is introduced to obtain a fast convergence rate. We call it a stabilizer instead of a regularizer since  $\lambda$  does not need to shrink to zero as  $n$  approaches infinity. Note that regularizers are needed when we penalize the norms of  $q$  and  $\xi$ , i.e.,

$$\hat{b}_V = \operatorname{argmin}_{q \in \mathcal{Q}} \max_{\xi \in \Xi} \mathbb{E}_{\mathcal{D}}[L(q, \xi) - 0.5\lambda\xi^2(H)] + 0.5\alpha' \|q\|_{\mathcal{Q}}^2 - 0.5\alpha \|\xi\|_{\Xi}^2,$$

for certain function norms  $\|\cdot\|_{\mathcal{Q}}$  and  $\|\cdot\|_{\Xi}$  and hyperparameters  $\alpha', \alpha > 0$ .

In the remainder of this section, we present three concrete examples using linear models, RKHSs and neural networks. Let  $\|\cdot\|_{\mathcal{Q}}$  and  $\|\cdot\|_{\Xi}$  denote  $L^2$ -norms when we use linear models and RKHS norms when using RKHSs.

**Example 2 (Linear models).** Suppose  $\mathcal{Q}, \Xi$  are linear, i.e.,  $\mathcal{G} = \{\theta^\top \phi_{\mathcal{F}}(\cdot) \mid \theta \in \mathbb{R}^{d_{\mathcal{F}}}\}, \Xi = \{\theta^\top \phi_{\mathcal{H}}(\cdot) \mid \theta \in \mathbb{R}^{d_{\mathcal{H}}}\}$  with features  $\phi_{\mathcal{F}} : \mathcal{F} \rightarrow \mathbb{R}^{d_{\mathcal{F}}}, \phi_{\mathcal{H}} : \mathcal{H} \rightarrow \mathbb{R}^{d_{\mathcal{H}}}$ . Then,

$$\hat{b}_V(\cdot) = \phi_{\mathcal{F}}^\top(\cdot) \{\mathbf{M}_2^\top \{\alpha I_{d_{\mathcal{H}}} + \lambda \mathbf{M}_3\}^{-1} \mathbf{M}_2 + \alpha' I_{d_{\mathcal{F}}}\}^{-1} \mathbf{M}_2^\top \{\alpha I_{d_{\mathcal{H}}} + \lambda \mathbf{M}_3\}^{-1} \mathbf{M}_1,$$

$$\mathbf{M}_1 = \mathbb{E}_{\mathcal{D}}[\mu(O, A)R\phi_{\mathcal{H}}(H)], \mathbf{M}_2 = \mathbb{E}_{\mathcal{D}}[\phi_{\mathcal{H}}(H)\{\phi_{\mathcal{F}}^\top(F) - \gamma\mu(O, A)\phi_{\mathcal{F}}^\top(F')\}], \mathbf{M}_3 = \mathbb{E}_{\mathcal{D}}[\phi_{\mathcal{H}}(H)\phi_{\mathcal{H}}^\top(H)].$$

When  $\alpha' = 0, \alpha = 0, \lambda = 0$ , the value estimators boil down to

$$\hat{b}_V(\cdot) = \phi_{\mathcal{F}}^\top(\cdot) \mathbf{M}_2^+ \mathbf{M}_1, \quad \hat{J}_{\text{VM}} = \mathbb{E}_{f \sim \nu_{\mathcal{F}}}[\phi_{\mathcal{F}}^\top(f)] \mathbf{M}_2^+ \mathbf{M}_1.$$

The above estimators are closely related to the LSTD estimators in MDPs [[LP03](#)]. Specifically, the off-policy LSTD estimator for state-value functions [[DNP<sup>+</sup>14](#), [UJ20](#)] is given by

$$\mathbb{E}_{s \sim \nu_S}[\phi_S^\top(s)] \mathbb{E}_{\mathcal{D}}[\phi_S^\top(S)\{\phi_S^\top(S) - \gamma\mu(S, A)\phi_S^\top(S')\}]^+ \mathbb{E}_{\mathcal{D}}[\mu(S, A)R\phi_S(S)]. \quad (6)$$

Our new proposed estimator  $\hat{J}_{\text{VM}}$  is

$$\mathbb{E}_{f \sim \nu_{\mathcal{F}}}[\phi_{\mathcal{F}}^\top(f)] \mathbb{E}_{\mathcal{D}}[\phi_{\mathcal{H}}(H)\{\phi_{\mathcal{F}}^\top(F) - \gamma\mu(O, A)\phi_{\mathcal{F}}^\top(F')\}]^+ \mathbb{E}_{\mathcal{D}}[\mu(O, A)R\phi_{\mathcal{H}}(H)],$$

which is very similar to (6). The critical difference lies in that we use futures (including current observations) and histories as proxies to infer the latent state  $S$  under partial observability.

**Example 3 (RKHSs).** Let  $\mathcal{Q}, \Xi$  be RKHSs with kernels  $k_{\mathcal{F}}(\cdot, \cdot) : \mathcal{F} \times \mathcal{F} \rightarrow \mathbb{R}, k_{\mathcal{H}}(\cdot, \cdot) : \mathcal{H} \times \mathcal{H} \rightarrow \mathbb{R}$ . Then,

$$\hat{b}_V(\cdot) = \mathbf{k}_{\mathcal{F}}(\cdot)^\top \{ \{\mathbf{K}'_{\mathcal{F}}\}^\top \{\alpha I_n + \mathbf{K}_{\mathcal{H}}\}^{-1} \mathbf{K}'_{\mathcal{F}} + \alpha' I_n \}^{-1} \{\mathbf{K}'_{\mathcal{F}}\}^\top \mathbf{K}_{\mathcal{H}}^{1/2} \{\alpha I_n + \mathbf{K}_{\mathcal{H}}\}^{-1} \mathbf{K}_{\mathcal{H}}^{1/2} \mathbf{Y}$$

where  $\mathbf{Y} \in \mathbb{R}^n, k(\cdot) \in \mathbb{R}^n, \mathbf{K}_{\mathcal{H}} \in \mathbb{R}^{n \times n}, \mathbf{K}_{\mathcal{F}} \in \mathbb{R}^{n \times n}$  such that

$$\{\mathbf{K}_{\mathcal{H}}\}_{(i,j)} = k_{\mathcal{H}}(H^{(i)}, H^{(j)}), \{\mathbf{K}_{\mathcal{F}}\}_{(i,j)} = k_{\mathcal{F}}(F^{(i)}, F^{(j)}), \{\mathbf{Y}\}_i = \mu(O^{(i)}, A^{(i)})R^{(i)}, \\ \{\mathbf{K}'_{\mathcal{F}}\}_{(i,j)} = k_{\mathcal{F}}(F^{(i)}, F^{(j)}) - \gamma k_{\mathcal{F}}(F^{(i)}, F^{(j)}), \{\mathbf{k}_{\mathcal{F}}(\cdot)\}_i = k_{\mathcal{F}}(\cdot, F^{(i)}).$$

**Example 4 (Neural Networks).** We set  $\mathcal{Q}$  to a class of neural networks and recommend to set  $\Xi$  to a linear or RKHS class so that the minimax optimization is reduced to single-stage optimization. The resulting optimization problem can be solved by off-the-shelf stochastic gradient descent (SGD) methods. Specifically, when we set  $\Xi$  to a linear model,  $\hat{b}_V$  is reduced to

$$\hat{b}_V = \operatorname{argmin}_{q \in \mathcal{Q}} Z(q)^\top \{\alpha I + \lambda \mathbb{E}_{\mathcal{D}}[\phi_{\mathcal{H}}(H)\phi_{\mathcal{H}}^\top(H)]\}^{-1} Z(q),$$

$$Z(q) := \mathbb{E}_{\mathcal{D}}[\mu(A, O)\{R + \gamma q(F')\} - q(F)]\phi_{\mathcal{H}}(H).$$

When we set  $\Xi$  to be an RKHS,  $\hat{b}_V$  is reduced to

$$\hat{b}_V = \operatorname{argmin}_{q \in \mathcal{Q}} \{Z'(q)\}^\top \mathbf{K}_{\mathcal{H}}^{1/2} \{\alpha I_n + \lambda \mathbf{K}_{\mathcal{H}}\}^{-1} \mathbf{K}_{\mathcal{H}}^{1/2} Z(q),$$

$$\{Z'(q)\}_i := \mu(A^{(i)}, O^{(i)})\{R^{(i)} + \gamma q(F'^{(i)})\} - q(F^{(i)}).$$

**Remark 4** (Comparison with minimax methods for MDPs). *Our proposed methods are closely related to the minimax learning approach (a.k.a. Bellman residual minimization) for RL [ASM08]. More specifically, in MDPs where the evaluation policy is memory-less, [UIJ<sup>+</sup>21, TFL<sup>+</sup>20] proposed minimax learning methods based on the following equations,*

$$\mathbb{E}[h(S)\{\mu(S, A)\{\gamma V^{\pi^e}(S') + R\} - V^{\pi^e}(S)\}] = 0, \quad \forall h : \mathcal{S} \rightarrow \mathbb{R}.$$

*These methods are no longer applicable in POMDPs since we are unable to directly observe the latent state. Although substituting the latent state with the entire history can still provide statistical guarantees, it is susceptible to the curse of horizon.*

**Remark 5** (Modeling of system dynamics). *Our minimax estimators can be extended to learning system dynamics. In particular, for tabular POMDPs, these results are closely related to the literature on spectral learning in HMMs and POMDPs [SBS<sup>+</sup>10, BSG11, BGG13]. Refer to Section F,G.*

**Remark 6** (Finite horizon setting). *The extension to the finite horizon setting is discussed in Section E*

## 5 Finite Sample Results

We study the finite sample convergence rate of the proposed estimators in this section. To simplify the technical analysis, we impose three assumptions. First, we assume the function classes are bounded, i.e.,  $\|\mathcal{Q}\|_\infty \leq C_{\mathcal{Q}}, \|\Xi\|_\infty \leq C_\Xi$  for some  $C_{\mathcal{Q}}, C_\Xi > 0$ . Second, we assume that the offline data are i.i.d.<sup>6</sup> Third, we assume  $\mathcal{Q}, \Xi$  are finite hypothesis classes. Meanwhile, our results can be extended to infinite hypothesis classes using the global/local Rademacher complexity theory; see e.g., [UIJ<sup>+</sup>21]. To simplify the presentation, following standard convention in OPE, we suppose the initial distribution is known, i.e.,  $|\mathcal{D}_{\text{ini}}| = \infty$ .

**Accuracy of  $\hat{b}_V$ .** We first demonstrate that  $\hat{b}_V$  consistently estimates the learnable future-dependent value bridge functions. To formalize this, we introduce the following Bellman operators.

**Definition 3** (Bellman operators). *The Bellman residual operator onto the history is defined as*

$$\mathcal{T} : [\mathcal{F} \rightarrow \mathbb{R}] \ni q(\cdot) \mapsto \mathbb{E}[\mu(O, A)\{R + \gamma q(F')\} - q(F) \mid H = \cdot],$$

*and the Bellman residual error onto the history is defined as  $\mathbb{E}[(\mathcal{T}q)^2(H)]$ . Similarly, the Bellman residual operator onto the latent state,  $\mathcal{T}^S$  is defined as*

$$\mathcal{T}^S : [\mathcal{F} \rightarrow \mathbb{R}] \ni q(\cdot) \mapsto \mathbb{E}[\mu(O, A)\{R + \gamma q(F')\} - q(F) \mid S = \cdot],$$

*and the Bellman residual error onto the latent state is defined as  $\mathbb{E}[\{\mathcal{T}^S(q)\}^2(S)]$ . The conditional expectations equal to zero when  $h$  and  $s$  lie outside the support of  $H$  and  $S$ , respectively.*

The Bellman residual error onto the history is zero for any learnable future-dependent value function, i.e.,  $\mathbb{E}[(\mathcal{T}b_V)^2(H)] = 0$ . Thus, this is a suitable measure to assess how well value function-based estimators approximate the true learnable future-dependent value functions.

**Theorem 2** (Finite sample property of  $\hat{b}_V$ ). *Set  $\lambda > 0$ . Suppose (2a)  $\mathcal{B}_V \cap \mathcal{Q} \neq \emptyset$  (realizability) and (2b)  $\mathcal{T}\mathcal{Q} \subset \Xi$  (Bellman completeness). With probability  $1 - \delta$ , we have  $\mathbb{E}[(\mathcal{T}\hat{b}_V)^2(H)]^{1/2} \leq c\{1/\lambda + \lambda\} \max(1, C_{\mathcal{Q}}, C_\Xi) \sqrt{\frac{\ln(|\mathcal{Q}||\Xi|c/\delta)}{n}}$  where  $c$  is some universal constant.*

Note (2a) and (2b) are commonly assumed in the literature on MDPs [CJ19, UIJ<sup>+</sup>21] as well. In particular, Bellman completeness means that the function class  $\Xi$  is sufficiently rich to capture the Bellman update of functions in  $\mathcal{Q}$ . For instance, these assumptions are naturally met in HSE-POMDPs. We require  $\lambda > 0$  in the statement of Theorem 2 to obtain a parametric convergence rate. When  $\lambda = 0$ , although we can obtain a rate of  $O_p(n^{-1/4})$ , it is unclear whether we can achieve  $O_p(n^{-1/2})$ .

**Accuracy of  $\hat{J}_{\text{VM}}$ .** We derive the finite sample guarantee for  $\hat{J}_{\text{VM}}$ .

<sup>6</sup>Without the independence assumption, similar results can be similarly established by imposing certain mixing conditions; see e.g., [SZLS20, LQM20, KU22].



**Theorem 3** (Finite sample property of  $\hat{J}_{\text{VM}}$ ). *Set  $\lambda > 0$ . Suppose (2a), (2b), (2c) any element in  $q \in \mathcal{Q}$  that satisfies  $\mathbb{E}[\{\mathcal{T}^S(q)\}(S) \mid H] = 0$  also satisfies  $\mathcal{T}^S(q)(S) = 0$ . (2d) the overlap  $\mu(O, A) < \infty$  and any element in  $q \in \mathcal{Q}$  that satisfies  $\mathcal{T}^S(q)(S) = 0$  also satisfies  $\mathcal{T}^S(q)(S^\circ) = 0$  where  $S^\circ \sim d_{\pi^e}(s)$ . With probability  $1 - \delta$ , we have*

$$|J(\pi^e) - \hat{J}_{\text{VM}}| \leq \frac{c(1/\lambda + \lambda)}{(1 - \gamma)^2} \max(1, C_{\mathcal{Q}}, C_{\Xi}) \text{IV}_1(\mathcal{Q}) \text{Dr}_{\mathcal{Q}}[d_{\pi^e}, P_{\pi^b}] \sqrt{\frac{\ln(|\mathcal{Q}||\Xi|c/\delta)}{n}}, \quad (7)$$

where

$$\text{IV}_1^2(\mathcal{Q}) := \sup_{\{q \in \mathcal{Q}; \mathbb{E}[\{\mathcal{T}(q)(H)\}^2] \neq 0\}} \frac{\mathbb{E}[\{\mathcal{T}^S(q)(S)\}^2]}{\mathbb{E}[\{\mathcal{T}(q)(H)\}^2]}, \quad (8)$$

$$\text{Dr}_{\mathcal{Q}}^2[d_{\pi^e}, P_{\pi^b}] := \sup_{\{q \in \mathcal{Q}; \mathbb{E}_{s \sim P_{\pi^b}}[\{\mathcal{T}^S(q)(s)\}^2] \neq 0\}} \frac{\mathbb{E}_{s \sim d_{\pi^e}}[\{\mathcal{T}^S(q)(s)\}^2]}{\mathbb{E}_{s \sim P_{\pi^b}}[\{\mathcal{T}^S(q)(s)\}^2]}. \quad (9)$$

On top of (2a) and (2b) that are assumed in Theorem 2, when  $\text{IV}_1(\mathcal{Q}) < \infty$ ,  $\text{Dr}_{\mathcal{Q}}(d_{\pi^e}, P_{\pi^b}) < \infty$ , (2c) and (2d) hold, we have the non-vacuous PAC guarantee. The condition  $\text{Dr}_{\mathcal{Q}}(d_{\pi^e}, P_{\pi^b}) < \infty$  and (2d) are the overlap conditions, which are adaptive to the function class  $\mathcal{Q}$  and are weaker than (1c). These are also used in MDPs [X CJ<sup>+</sup>21]. Here,  $\text{Dr}_{\mathcal{Q}}(d_{\pi^e}, P_{\pi^b})$  is a refined version of the density ratio and satisfies  $\text{Dr}_{\mathcal{Q}}(d_{\pi^e}, P_{\pi^b}) \leq w_{\pi^e}(S)$ . The condition  $\text{IV}_1(\mathcal{Q}) < \infty$  and (2c) are characteristic conditions in POMDPs that quantify how much errors are properly translated from on  $H$  to on  $S$ . Similar assumptions are frequently imposed in IV literature [DLMS20, CP12].

The upper error bound in (7) does not have explicit exponential dependence on the effective horizon  $1/(1 - \gamma)$ . In particular, as shown in Section D, for tabular POMDPs and HSE-POMDPs, the terms  $\text{Dr}_{\mathcal{Q}}(d_{\pi^e}, P_{\pi^b})$  and  $\kappa(\mathcal{Q})$  can be reduced to certain condition numbers associated with covariance matrices spanned by feature vectors; see (15) and (16) in Appendix D. Hence, unlike SIS-based methods, we are able to break the curse of horizon.

The numbers of future and history proxies included in  $F$  and  $H$  represents a tradeoff. Specifically, if  $F$  contains enough observations, it is likely that (2a) will hold. Meanwhile, if  $H$  contains enough observations, it is more likely that (2b) will hold. These facts demonstrate the benefits of including a sufficient number of observations in  $F$  and  $H$ . However, the statistical complexities  $\ln(|\mathcal{Q}||\Xi|)$  will increase with the number of observations in  $F$  and  $H$ .

Lastly, it's worth noting that while our methods effectively address the curse of horizon, they may incur exponential growth concerning the number of future proxies used. This also applies to history proxies, which should be longer than the length of short memory policies. Here, we focus on the explanation of future proxies. For instance, in the tabular case,  $\log |\Omega|$  might scale with  $|\mathcal{O}|^{M_F}$  when considering  $\Omega$  as the set of all functions on  $\mathcal{O}^{M_F}$ . However, this situation differs significantly from the curse of horizon, a challenge that naive methods like replacing states with the entire history encounter. These methods would necessitate the entire history to achieve Markovianity, whereas we only require a shorter length of future observations to establish the conditions outlined in Theorem 1 (existence), which can be much shorter. Specific examples are provided throughout the paper, including Example 1, which discusses the tabular setting and demonstrates that we essentially need as many future proxies as states, as long as there is sufficient statistical dependence between them.

## 6 Experiment

This section empirically evaluates the performance of the proposed method on a synthetic dataset.<sup>7</sup>

We use the CartPole environment provided by OpenAI Gym [BCP<sup>+</sup>16], which is commonly employed in other OPE studies [SUHJ22, FCG18]. By default, this non-tabular environment consists of 4-dimensional states, which are fully observable. Following [SUHJ22], we create partial observability by adding independent Gaussian noises to each dimension of the state as  $O^{(j)} = S^{(j)}(1 + \mathcal{N}(1 + 0.3^2))$ ,  $1 \leq j \leq 4$ . To define behavior and evaluation policies, we first train an expert policy using DDQN [VHGS16] on latent states  $S$ . Subsequently, we apply Behavior Cloning (BC) on two datasets, one containing pairs of latent state and action ( $S, A$ ) and the other containing pairs of observation and action ( $O, A$ ), respectively. Then, we use the base policy obtained by BC on the state-action

<sup>7</sup>Our code is available at <https://github.com/aieola/neurips2023-future-dependent-ope>




Figure 2: Logarithms of relative biases (left) and MSEs (right) of the proposed and the baseline estimators for various values of  $\epsilon$ , which specify the evaluation policy. The confidence intervals are obtained through 100 Monte Carlo simulations.

pairs  $(S, A)$  to define an  $\epsilon$ -greedy behavior policy, where we set  $\epsilon = 0.3$ .<sup>8</sup> Similarly, the evaluation policy is also an  $\epsilon$ -greedy policy, based on the base policy obtained by BC on the observation-action pairs  $(O, A)$ , with different values of  $\epsilon \in [0.1, 0.3, 0.5, 0.7]$ . We conduct the experiment with 100 random seeds, and for each simulation, we collect logged data consisting of 1000 trajectories, each containing 100 steps.

We compare our proposal with Sequential Importance Sampling (SIS) [Pre00] and the naive minimax OPE [UJH20], which is designed for fully-observable MDPs and does not account for partial observability. The naive minimax estimator is defined as if the environment was fully observable, replacing  $H$  and  $\bar{F}$  in Algorithm 1 with the current observation  $O$ . In contrast, our proposed method uses a 3-step history as  $H$  and a one-step future as  $F$  to address partial observability. Both our proposed method and the naive approach use two-layer neural networks for the function  $Q$  and RKHSs for  $\Xi$ , as detailed in Example 4.

We present the results in Figure 2, which demonstrate the superior accuracy of our proposed estimator compared to the baselines (SIS and naive minimax estimator) in terms of mean square errors (MSEs). Additional experimental details and ablation results, including the variations in the length of  $H$ ,  $\bar{F}$ , and the choice of RKHSs, can be found in Appendix H.3.

## 7 Conclusion

We present a novel approach for OPE of short-memory policies in POMDPs. Our method involves introducing future-dependent value functions and the associated off-policy Bellman equations, followed by proposing a minimax estimator based on these equations. This is the first model-free method that allows for general function approximation and mitigates the curse of horizon. Our proposal is grounded in three interpretable key assumptions: observability, which asserts the presence of (short) future observations retaining adequate information about latent states, invertibility, which posits the existence of (short) histories preserving ample information about latent states; and the overlap between evaluation policies and behavior policies.

We have several avenues for enhancing our proposals. Firstly, automatically determining the appropriate lengths of futures and histories holds practical significance. Additionally, exploring recent attention mechanisms that extend beyond selecting the most recent history proxies or the nearest future proxies shows promise. Secondly, while we establish Bellman equations for POMDPs and use a simple minimax approach akin to [DLMS20], we may benefit from leveraging more refined methods introduced in recent research for solving conditional moment equations [BKM<sup>+</sup>23c, BKM<sup>+</sup>23a, BKM<sup>+</sup>23b].

## Acknowledgements

This material is based upon work supported by the National Science Foundation under Grant Nos. IIS 2112471, IIS 2141781, IIS 1846210, IIS 2154711 and an Engineering and Physical Sciences Research Council grant EP/W014971/1.

<sup>8</sup>When collecting offline data from the behavior policy, the behavior policy takes the observation  $O$  as input, noting that the dimension of  $S$  and  $O$  is the same in our experimental setting.

## References

- [ALA16] Kamyar Azizzadenesheli, Alessandro Lazaric, and Animashree Anandkumar. Reinforcement learning of pomdps using spectral methods. In *Conference on Learning Theory*, pages 193–256. PMLR, 2016.
- [ASM08] András Antos, Csaba Szepesvári, and Rémi Munos. Learning near-optimal policies with bellman-residual minimization based fitted policy iteration and a single sample path. *Machine Learning*, 71(1):89–129, 2008.
- [AYA18] Kamyar Azizzadenesheli, Yisong Yue, and Animashree Anandkumar. Policy gradient in partially observable environments: Approximation and convergence. *arXiv preprint arXiv:1810.07900*, 2018.
- [BBC<sup>+</sup>19] Christopher Berner, Greg Brockman, Brooke Chan, Vicki Cheung, Przemyslaw Debiak, Christy Dennison, David Farhi, Quirin Fischer, Shariq Hashme, Chris Hesse, et al. Dota 2 with large scale deep reinforcement learning. *arXiv preprint arXiv:1912.06680*, 2019.
- [BCP<sup>+</sup>16] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.
- [BDRS96] Dima Burago, Michel De Rougemont, and Anatol Slissenko. On the complexity of partially observed markov decision processes. *Theoretical Computer Science*, 157(2):161–183, 1996.
- [Ber12] Dimitri Bertsekas. *Dynamic programming and optimal control: Volume I*, volume 1. Athena scientific, 2012.
- [BGG13] Byron Boots, Geoffrey Gordon, and Arthur Gretton. Hilbert space embeddings of predictive state representations. *arXiv preprint arXiv:1309.6819*, 2013.
- [BK21] Andrew Bennett and Nathan Kallus. Proximal reinforcement learning: Efficient off-policy evaluation in partially observed markov decision processes. 2021.
- [BKLM21] Andrew Bennett, Nathan Kallus, Lihong Li, and Ali Mousavi. Off-policy evaluation in infinite-horizon reinforcement learning with latent confounders. In *International Conference on Artificial Intelligence and Statistics*, pages 1999–2007. PMLR, 2021.
- [BKM<sup>+</sup>23a] Andrew Bennett, Nathan Kallus, Xiaojie Mao, Whitney Newey, Vasilis Syrgkanis, and Masatoshi Uehara. Inference on strongly identified functionals of weakly identified functions. In *The Thirty Sixth Annual Conference on Learning Theory*, pages 2265–2265. PMLR, 2023.
- [BKM<sup>+</sup>23b] Andrew Bennett, Nathan Kallus, Xiaojie Mao, Whitney Newey, Vasilis Syrgkanis, and Masatoshi Uehara. Minimax instrumental variable regression and  $l_2$  convergence guarantees without identification or closedness. In Gergely Neu and Lorenzo Rosasco, editors, *Proceedings of Thirty Sixth Conference on Learning Theory*, volume 195 of *Proceedings of Machine Learning Research*, pages 2291–2318, 2023.
- [BKM<sup>+</sup>23c] Andrew Bennett, Nathan Kallus, Xiaojie Mao, Whitney Newey, Vasilis Syrgkanis, and Masatoshi Uehara. Source condition double robust inference on functionals of inverse problems. *arXiv preprint arXiv:2307.13793*, 2023.
- [BKS19] Andrew Bennett, Nathan Kallus, and Tobias Schnabel. Deep generalized method of moments for instrumental variable analysis. In *Advances in Neural Information Processing Systems 32*, pages 3564–3574. 2019.
- [BMVVDL19] Aurelien Bibaut, Ivana Malenica, Nikos Vlassis, and Mark Van Der Laan. More efficient off-policy evaluation through regularized targeted learning. In *International Conference on Machine Learning*, pages 654–663. PMLR, 2019.
- [BSG11] Byron Boots, Sajid M Siddiqi, and Geoffrey J Gordon. Closing the learning-planning loop with predictive state representations. *The International Journal of Robotics Research*, 30(7):954–966, 2011.
- [BSZ23] David Bruns-Smith and Angela Zhou. Robust fitted-q-evaluation and iteration under sequentially exogenous unobserved confounders. *arXiv preprint arXiv:2302.00662*, 2023.

- [Cas98] Anthony R Cassandra. A survey of pomdp applications. In *Working notes of AAAI 1998 fall symposium on planning with partially observable Markov decision processes*, volume 1724, 1998.
- [CBM22] Fan Chen, Yu Bai, and Song Mei. Partially observable rl with b-stability: Unified structural condition and sharp sample-efficient algorithms. *arXiv preprint arXiv:2209.14990*, 2022.
- [CFR07] Marine Carrasco, Jean-Pierre Florens, and Eric Renault. Linear inverse problems in structural econometrics estimation based on spectral decomposition and regularization. *Handbook of econometrics*, 6:5633–5751, 2007.
- [CJ19] Jinglin Chen and Nan Jiang. Information-theoretic considerations in batch reinforcement learning. In *International Conference on Machine Learning*, pages 1042–1051. PMLR, 2019.
- [CKPW18] Zhiyong Cui, Ruimin Ke, Ziyuan Pu, and Yin Hai Wang. Deep bidirectional and unidirectional lstm recurrent neural network for network-wide traffic speed prediction. *arXiv preprint arXiv:1801.02143*, 2018.
- [CNS18] Victor Chernozhukov, Whitney K Newey, and Rahul Singh. Learning l2 continuous regression functionals via regularized riesz representers. *arXiv preprint arXiv:1809.05224*, 2018.
- [CNSS20] Victor Chernozhukov, Whitney Newey, Rahul Singh, and Vasilis Syrgkanis. Adversarial estimation of riesz representers. *arXiv preprint arXiv:2101.00009*, 2020.
- [CO20] Sayak Ray Chowdhury and Rafael Oliveira. No-regret reinforcement learning with value function approximation: a kernel embedding approach. *arXiv preprint arXiv:2011.07881*, 2020.
- [CP12] Xiaohong Chen and Demian Pouzo. Estimation of nonparametric conditional moment models with possibly nonsmooth generalized residuals. *Econometrica*, 80(1):277–321, 2012.
- [CPS<sup>+</sup>20] Yifan Cui, Hongming Pu, Xu Shi, Wang Miao, and Eric Tchetgen Tchetgen. Semi-parametric proximal causal inference. *arXiv preprint arXiv:2011.08411*, 2020.
- [CQ22] Xiaohong Chen and Zhengling Qi. On well-posedness and minimax optimal rates of nonparametric q-function estimation in off-policy evaluation. *arXiv preprint arXiv:2201.06169*, 2022.
- [CWYW23] Siyu Chen, Yitan Wang, Zhaoran Wang, and Zhuoran Yang. A unified framework of policy learning for contextual bandit with confounding bias and missing observations. *arXiv preprint arXiv:2303.11187*, 2023.
- [DHB<sup>+</sup>17] Carlton Downey, Ahmed Hefny, Byron Boots, Geoffrey J Gordon, and Boyue Li. Predictive state recurrent neural networks. *Advances in Neural Information Processing Systems*, 30, 2017.
- [DJW20] Yaqi Duan, Zeyu Jia, and Mengdi Wang. Minimax-optimal off-policy evaluation with linear function approximation. In *International Conference on Machine Learning*, pages 2701–2709. PMLR, 2020.
- [DLMS20] Nishanth Dikkala, Greg Lewis, Lester Mackey, and Vasilis Syrgkanis. Minimax estimation of conditional moment models. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 12248–12262, 2020.
- [DNP<sup>+</sup>14] Christoph Dann, Gerhard Neumann, Jan Peters, et al. Policy evaluation with temporal differences: A survey and comparison. *Journal of Machine Learning Research*, 15:809–883, 2014.
- [EGW05] Damien Ernst, Pierre Geurts, and Louis Wehenkel. Tree-based batch mode reinforcement learning. *Journal of Machine Learning Research*, 6:503–556, 2005.
- [FCG18] Mehrdad Farajtabar, Yinlam Chow, and Mohammad Ghavamzadeh. More robust doubly robust off-policy evaluation. In *International Conference on Machine Learning*, pages 1447–1456. PMLR, 2018.

- [FLL19] Yihao Feng, Lihong Li, and Qiang Liu. A kernel loss for solving the bellman equation. *Advances in Neural Information Processing Systems*, 32, 2019.
- [FS19] Dylan J Foster and Vasilis Syrgkanis. Orthogonal statistical learning. *arXiv preprint arXiv:1901.09036*, 2019.
- [GCZ<sup>+</sup>22] Hongyi Guo, Qi Cai, Yufeng Zhang, Zhuoran Yang, and Zhaoran Wang. Provably efficient offline reinforcement learning for partially observable markov decision processes. In *International Conference on Machine Learning*, pages 8016–8038. PMLR, 2022.
- [GDB16] Zhaohan Daniel Guo, Shayan Doroudi, and Emma Brunskill. A pac rl algorithm for episodic pomdps. In *Artificial Intelligence and Statistics*, pages 510–518. PMLR, 2016.
- [GMR22] Noah Golowich, Ankur Moitra, and Dhruv Rohatgi. Planning in observable pomdps in quasipolynomial time. *arXiv preprint arXiv:2201.04735*, 2022.
- [HDG15] Ahmed Hefny, Carlton Downey, and Geoffrey J Gordon. Supervised learning for dynamical system learning. *Advances in neural information processing systems*, 28:1963–1971, 2015.
- [HDL<sup>+</sup>21] Botao Hao, Yaqi Duan, Tor Lattimore, Csaba Szepesvári, and Mengdi Wang. Sparse feature selection makes batch reinforcement learning more sample efficient. In *International Conference on Machine Learning*, pages 4063–4073. PMLR, 2021.
- [HKZ12] Daniel Hsu, Sham M Kakade, and Tong Zhang. A spectral algorithm for learning hidden markov models. *Journal of Computer and System Sciences*, 78(5):1460–1480, 2012.
- [Hor11] Joel L Horowitz. Applied nonparametric instrumental variables estimation. *Econometrica*, 79(2):347–394, 2011.
- [HW17] David A Hirshberg and Stefan Wager. Augmented minimax linear estimation. *arXiv preprint arXiv:1712.00038*, 2017.
- [HW21] Yuchen Hu and Stefan Wager. Off-policy evaluation in partially observed markov decision processes. *arXiv preprint arXiv:2110.12343*, 2021.
- [Jae00] Herbert Jaeger. Observable operator models for discrete stochastic time series. *Neural computation*, 12(6):1371–1398, 2000.
- [JKA<sup>+</sup>17] Nan Jiang, Akshay Krishnamurthy, Alekh Agarwal, John Langford, and Robert E Schapire. Contextual decision processes with low bellman rank are pac-learnable. In *International Conference on Machine Learning*, pages 1704–1713. PMLR, 2017.
- [JKKL20] Chi Jin, Sham M Kakade, Akshay Krishnamurthy, and Qinghua Liu. Sample-efficient reinforcement learning of undercomplete pomdps. *arXiv preprint arXiv:2006.12484*, 2020.
- [JL16] Nan Jiang and Lihong Li. Doubly robust off-policy value evaluation for reinforcement learning. In *International Conference on Machine Learning*, pages 652–661. PMLR, 2016.
- [JSLZ19] Binyan Jiang, Rui Song, Jialiang Li, and Donglin Zeng. Entropy learning for dynamic treatment regimes. *Statistica Sinica*, 29(4):1633, 2019.
- [KAL16] Akshay Krishnamurthy, Alekh Agarwal, and John Langford. Pac reinforcement learning with rich observations. *Advances in Neural Information Processing Systems*, 29, 2016.
- [KB14] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [KECM21] Jeongyeol Kwon, Yonathan Efroni, Constantine Caramanis, and Shie Mannor. RL for latent mdps: Regret guarantees and a lower bound. *Advances in Neural Information Processing Systems*, 34, 2021.
- [KHS<sup>+</sup>15] Predrag Klasnja, Eric B Hekler, Saul Shiffman, Audrey Boruvka, Daniel Almirall, Ambuj Tewari, and Susan A Murphy. Microrandomized trials: An experimental design for developing just-in-time adaptive interventions. *Health Psychology*, 34(S):1220, 2015.

- [KJS15] Alex Kulesza, Nan Jiang, and Satinder Singh. Spectral learning of predictive state representations with insufficient statistics. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015.
- [KL19] Michael R Kosorok and Eric B Laber. Precision medicine. *Annual review of statistics and its application*, 6:263–286, 2019.
- [KLC98] Leslie Pack Kaelbling, Michael L Littman, and Anthony R Cassandra. Planning and acting in partially observable stochastic domains. *Artificial intelligence*, 101(1-2):99–134, 1998.
- [KMU21] Nathan Kallus, Xiaojie Mao, and Masatoshi Uehara. Causal inference under unmeasured confounding with negative controls: A minimax learning approach. *arXiv preprint arXiv:2103.14029*, 2021.
- [KU20] Nathan Kallus and Masatoshi Uehara. Double reinforcement learning for efficient off-policy evaluation in markov decision processes. *J. Mach. Learn. Res.*, 21(167):1–63, 2020.
- [KU22] Nathan Kallus and Masatoshi Uehara. Efficiently breaking the curse of horizon in off-policy evaluation with double reinforcement learning. *Operations Research*, 2022.
- [KY22] Ali Devran Kara and Serdar Yuksel. Near optimality of finite memory feedback policies in partially observed markov decision processes. *Journal of Machine Learning Research*, 23(11):1–46, 2022.
- [KZ20] Nathan Kallus and Angela Zhou. Confounding-robust policy evaluation in infinite-horizon reinforcement learning. *arXiv preprint arXiv:2002.04518*, 2020.
- [LAHA21] Sahin Lale, Kamyar Azizzadenesheli, Babak Hassibi, and Anima Anandkumar. Adaptive control and regret minimization in linear quadratic gaussian (lqg) setting. In *2021 American Control Conference (ACC)*, pages 2517–2522. IEEE, 2021.
- [LCSJ22] Qinghua Liu, Alan Chung, Csaba Szepesvári, and Chi Jin. When is partially observable reinforcement learning not scary? *arXiv preprint arXiv:2204.08967*, 2022.
- [LFY<sup>+</sup>21] Luofeng Liao, Zuyue Fu, Zhuoran Yang, Yixin Wang, Mladen Kolar, and Zhaoran Wang. Instrumental variable value iteration for causal offline reinforcement learning. *arXiv preprint arXiv:2102.09907*, 2021.
- [LKTF20] Sergey Levine, Aviral Kumar, George Tucker, and Justin Fu. Offline reinforcement learning: Tutorial, review, and perspectives on open problems. *arXiv preprint arXiv:2005.01643*, 2020.
- [LLTZ18] Qiang Liu, Lihong Li, Ziyang Tang, and Dengyong Zhou. Breaking the curse of horizon: Infinite-horizon off-policy estimation. In *Advances in Neural Information Processing Systems 31*, pages 5356–5366. 2018.
- [LMWY22] Miao Lu, Yifei Min, Zhaoran Wang, and Zhuoran Yang. Pessimism in the face of confounders: Provably efficient offline reinforcement learning in partially observable markov decision processes. *arXiv preprint arXiv:2205.13589*, 2022.
- [LNSJ22] Qinghua Liu, Praneeth Netrapalli, Csaba Szepesvari, and Chi Jin. Optimistic mle—a generic model-based algorithm for partially observable sequential decision making. *arXiv preprint arXiv:2209.14997*, 2022.
- [LP03] Michail G Lagoudakis and Ronald Parr. Least-squares policy iteration. *The Journal of Machine Learning Research*, 4:1107–1149, 2003.
- [LQM20] Peng Liao, Zhengling Qi, and Susan Murphy. Batch policy learning in average reward markov decision processes. *arXiv preprint arXiv:2007.11771*, 2020.
- [LS01] Michael Littman and Richard S Sutton. Predictive representations of state. *Advances in neural information processing systems*, 14, 2001.
- [LVY19] Hoang Le, Cameron Voloshin, and Yisong Yue. Batch policy learning under constraints. In *International Conference on Machine Learning*, pages 3703–3712. PMLR, 2019.
- [MGT18] Wang Miao, Zhi Geng, and Eric J Tchetgen Tchetgen. Identifying causal effects with proxy variables of an unmeasured confounder. *Biometrika*, 105(4):987–993, 2018.

- [MKS<sup>+</sup>13] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.
- [MS08] Rémi Munos and Csaba Szepesvári. Finite-time bounds for fitted value iteration. *Journal of Machine Learning Research*, 9(5), 2008.
- [MST18] Wang Miao, Xu Shi, and Eric Tchetgen Tchetgen. A confounding bridge approach for double negative control inference on causal effects. *arXiv preprint arXiv:1808.04945*, 2018.
- [MvdLRG01] Susan A Murphy, Mark J van der Laan, James M Robins, and Conduct Problems Prevention Research Group. Marginal mean models for dynamic regimes. *Journal of the American Statistical Association*, 96(456):1410–1423, 2001.
- [MY20] Curtis McDonald and Serdar Yüksel. Exponential filter stability via dobrushin’s coefficient. *Electronic Communications in Probability*, 25:1–13, 2020.
- [NCDL19] Ofir Nachum, Yinlam Chow, Bo Dai, and Lihong Li. Dualdice: Behavior-agnostic estimation of discounted stationary distribution corrections. *arXiv preprint arXiv:1906.04733*, 2019.
- [New13] Whitney K Newey. Nonparametric instrumental variables estimation. *American Economic Review*, 103(3):550–56, 2013.
- [NJ21] Yash Nair and Nan Jiang. A spectral approach to off-policy evaluation for pomdps. *arXiv preprint arXiv:2109.10502*, 2021.
- [NKYB20] Hongseok Namkoong, Ramtin Keramati, Steve Yadlowsky, and Emma Brunskill. Off-policy policy evaluation for sequential decisions under unobserved confounding. *arXiv preprint arXiv:2003.05623*, 2020.
- [PGT06] Joelle Pineau, Geoffrey Gordon, and Sebastian Thrun. Anytime point-based approximations for large pomdps. *Journal of Artificial Intelligence Research*, 27:335–380, 2006.
- [Pre00] Doina Precup. Eligibility traces for off-policy policy evaluation. *Computer Science Department Faculty Publication Series*, page 80, 2000.
- [PT87] Christos H Papadimitriou and John N Tsitsiklis. The complexity of markov decision processes. *Mathematics of operations research*, 12(3):441–450, 1987.
- [RGT04] Matthew Rosencrantz, Geoff Gordon, and Sebastian Thrun. Learning low dimensional predictive representations. In *Proceedings of the twenty-first international conference on Machine learning*, page 88, 2004.
- [RRS98] Andrea Rotnitzky, James M Robins, and Daniel O Scharfstein. Semiparametric regression for repeated outcomes with nonignorable nonresponse. *Journal of the american statistical association*, 93(444):1321–1339, 1998.
- [SBS<sup>+</sup>10] Le Song, Byron Boots, Sajid M Siddiqi, Geoffrey Gordon, and Alex Smola. Hilbert space embeddings of hidden markov models. In *Proceedings of the 27th International Conference on International Conference on Machine Learning*, pages 991–998, 2010.
- [SHSF09] Le Song, Jonathan Huang, Alex Smola, and Kenji Fukumizu. Hilbert space embeddings of conditional distributions with applications to dynamical systems. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pages 961–968, 2009.
- [SJR04] Satinder Singh, Michael James, and Matthew Rudary. Predictive state representations: A new theory for modeling dynamical systems. In *Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence*, 2004.
- [SMNTT20] Xu Shi, Wang Miao, Jennifer C Nelson, and Eric J Tchetgen Tchetgen. Multiply robust causal inference with double-negative control adjustment for categorical unmeasured confounding. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 82(2):521–540, 2020.
- [SRRH21] Adish Singla, Anna N Rafferty, Goran Radanovic, and Neil T Heffernan. Reinforcement learning for education: Opportunities and challenges. *arXiv preprint arXiv:2107.08828*, 2021.

- [SSH20] Max Simchowitz, Karan Singh, and Elad Hazan. Improper learning for non-stochastic control. In *Conference on Learning Theory*, pages 3320–3436. PMLR, 2020.
- [SUHJ22] Chengchun Shi, Masatoshi Uehara, Jiawei Huang, and Nan Jiang. A minimax learning approach to off-policy evaluation in confounded partially observable markov decision processes. In *International Conference on Machine Learning*, pages 20057–20094. PMLR, 2022.
- [SZLS20] Chengchun Shi, Sheng Zhang, Wenbin Lu, and Rui Song. Statistical inference of the value function for reinforcement learning in infinite horizon settings. *arXiv preprint arXiv:2001.04515*, 2020.
- [SZY<sup>+</sup>22] Chengchun Shi, Jin Zhu, Shen Ye, Shikai Luo, Hongtu Zhu, and Rui Song. Off-policy confidence interval estimation with confounded markov decision process. *Journal of the American Statistical Association*, pages 1–12, 2022.
- [TB16] Philip Thomas and Emma Brunskill. Data-efficient off-policy policy evaluation for reinforcement learning. In *International Conference on Machine Learning*, pages 2139–2148. PMLR, 2016.
- [TDHL19] Anastasios A Tsiatis, Marie Davidian, Shannon T Holloway, and Eric B Laber. *Dynamic Treatment Regimes: Statistical Methods for Precision Medicine*. CRC press, 2019.
- [TFL<sup>+</sup>20] Ziyang Tang, Yihao Feng, Lihong Li, Dengyong Zhou, and Qiang Liu. Harnessing infinitehorizon off-policy evaluation: Double robustness via duality. *ICLR 2020*, 2020.
- [TSM20] Guy Tennenholtz, Uri Shalit, and Shie Mannor. Off-policy evaluation in partially observable environments. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 10276–10283, 2020.
- [UHJ20] Masatoshi Uehara, Jiawei Huang, and Nan Jiang. Minimax weight and q-function learning for off-policy evaluation. In *International Conference on Machine Learning*, pages 9659–9668. PMLR, 2020.
- [UIJ<sup>+</sup>21] Masatoshi Uehara, Masaaki Imaizumi, Nan Jiang, Nathan Kallus, Wen Sun, and Tengyang Xie. Finite sample analysis of minimax offline reinforcement learning: Completeness, fast rates and first-order efficiency. *arXiv preprint arXiv:2102.02981*, 2021.
- [US21] Masatoshi Uehara and Wen Sun. Pessimistic model-based offline reinforcement learning under partial coverage. *arXiv preprint arXiv:2107.06226*, 2021.
- [USL<sup>+</sup>22] Masatoshi Uehara, Ayush Sekhari, Jason D Lee, Nathan Kallus, and Wen Sun. Provably efficient reinforcement learning in partially observable dynamical systems. *arXiv preprint arXiv:2206.12020*, 2022.
- [VHGS16] Hado Van Hasselt, Arthur Guez, and David Silver. Deep reinforcement learning with double q-learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 30, 2016.
- [VSH<sup>+</sup>16] Arun Venkatraman, Wen Sun, Martial Hebert, J Bagnell, and Byron Boots. Online instrumental variable regression with applications to online linear system identification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 30, 2016.
- [WYW20] Lingxiao Wang, Zhuoran Yang, and Zhaoran Wang. Provably efficient causal reinforcement learning with confounded observational data. *arXiv preprint arXiv:2006.12311*, 2020.
- [XCJ<sup>+</sup>21] Tengyang Xie, Ching-An Cheng, Nan Jiang, Paul Mineiro, and Alekh Agarwal. Bellman-consistent pessimism for offline reinforcement learning. *Advances in neural information processing systems*, 34, 2021.
- [XLSS20] Zekun Xu, Eric Laber, Ana-Maria Staicu, and Emanuel Severus. Latent-state models for precision medicine. *arXiv preprint arXiv:2005.13001*, 2020.
- [XZS<sup>+</sup>22] Yang Xu, Jin Zhu, Chengchun Shi, Shikai Luo, and Rui Song. An instrumental variable approach to confounded off-policy evaluation. *arXiv preprint arXiv:2212.14468*, 2022.



- [ZB16] Junzhe Zhang and Elias Bareinboim. Markov decision processes with unobserved confounders: A causal approach. Technical report, Technical Report R-23, Purdue AI Lab, 2016.
- [ZLKB20] Andrea Zanette, Alessandro Lazaric, Mykel Kochenderfer, and Emma Brunskill. Learning near optimal policies with low inherent bellman error. In *International Conference on Machine Learning*, pages 10978–10989. PMLR, 2020.
- [ZUSL22] Wenhao Zhan, Masatoshi Uehara, Wen Sun, and Jason D Lee. Pac reinforcement learning for predictive state representations. *arXiv preprint arXiv:2207.05738*, 2022.
- [ZW22] Andrea Zanette and Martin J Wainwright. Bellman residual orthogonalization for offline reinforcement learning. *arXiv preprint arXiv:2203.12786*, 2022.
- [ZWB21] Andrea Zanette, Martin J Wainwright, and Emma Brunskill. Provable benefits of actor-critic methods for offline reinforcement learning. *Advances in neural information processing systems*, 34, 2021.
- [ZXZ<sup>+</sup>22] Han Zhong, Wei Xiong, Sirui Zheng, Liwei Wang, Zhaoran Wang, Zhuoran Yang, and Tong Zhang. A posterior sampling framework for interactive decision making. *arXiv preprint arXiv:2211.01962*, 2022.

## A Additional Literature review

**Minimax learning.** Minimax learning (also known as adversarial learning) has been widely applied to a large variety of problems ranging from instrumental variable estimation [BKS19, DLMS20] to policy learning/evaluation in contextual bandits/MDPs [HW17, CNS18, FS19, FLL19, LLTZ18, CNSS20, UIJ+21]. For example, in OPE problems with fully-observable environments, minimax learning methods have been developed to learn q-functions and marginal ratios that are characterized as solutions to certain conditional moment equations [UJ20]. The solutions to these conditional moment equations are uniquely defined. On the other hand, our case is more challenging since the solutions to conditional moment equations are not uniquely defined. Although the solutions are *not* uniquely defined and hence cannot be identified, our estimands, i.e., policy values, *can* be still identified. This requires significantly distinctive analysis, which is not seen in standard IV settings or MDP settings.

## B Comparison to Analogue of Future-dependent Value Functions

Analogs of our future-dependent value functions have been introduced in confounded contextual bandits and confounded POMDPs as bridge functions [MGT18, CPS+20, TSM20, SUHJ22, BK21]. These works consider confounded settings where actions depend on unobserved states and introduce bridge functions to deal with confounding. Instead, we introduce future-dependent value functions to deal with the curse of the horizon while there is no confounding issue. Existing definitions of bridge functions in confounded POMDPs do not work in standard POMDPs. In the definition of existing bridge functions, behavior policies cannot depend on observations  $O$  since observations  $O$  are used as so-called negative controls, which do not affect action  $A$  and are not affected by action  $A$ . In our setting,  $O$  does *not* serve as negative controls unlike their works since  $A$  clearly depends on  $O$ . Instead,  $O$  just play a role in covariates. See Figure 3. Due to this fact, we can further add  $F$  as input domains of future-dependent value functions, unlike bridge functions by regarding  $F$  as just covariates. This is impossible in the definition of existing bridge functions without further assumptions as mentioned in [NJ21]. In this sense, our setting does *not* fall into the category of the so-called proximal causal inference framework. At the same time, our definition does not work in these confounded settings since Definition 2 explicitly includes behavior policies.

We finally remark the observation that history can serve as an instrumental variable in POMDPs is mentioned in [HDG15, VSH+16]. However, they did not propose future-dependent value functions; their goal is to learn system dynamics.

## C Off-Policy Evaluation for Memory-Based Policies

So far, we have discussed how to evaluate memoryless policies to simplify the notation. In this section, we will now turn our attention to the evaluation of memory-based policies.

### C.1 Settings

We consider  $M$ -memory policies  $\pi : \mathcal{Z} \times \mathcal{O} \rightarrow \Delta(\mathcal{A})$  that are functions of the current observation  $O_t$  and past observation-action pairs at time point  $t, t-1, \dots, t_M$  denoted by  $Z_t = (O_{t-M:t-1}, A_{t-M:t-1}) \in \mathcal{Z} = \mathcal{O}^M \times \mathcal{A}^M$ , for some integer  $M > 1$ . We assume the existence of  $M$  observation pairs obtained prior to the initial time point (denoted by  $Z_0$ ). Following an  $M$ -memory policy  $\pi$ , the data generating process can be described as follows. First,  $Z_0$  and  $S_0$  are generated according to some initial distribution  $\nu_{\bar{\mathcal{S}}} \in \Delta(\bar{\mathcal{S}})$  where  $\bar{\mathcal{S}} = \mathcal{Z} \times \mathcal{S}$ . Next, the agent observes  $O_0 \sim \mathbb{O}(\cdot | S_0)$ , executes the initial action  $A_0 \sim \pi(\cdot | Z_0, O_0)$ , receives a reward  $r(S_0, A_0)$ , the environment transits to the next state  $S_1 \sim \mathbb{T}(\cdot | S_0, A_0)$ , and this process repeats. See Figure 4 for a graphical illustration of the data-generating process. We assume that both the behavior and evaluation policies are  $M$ -memory.

Our goal is to estimate a policy value  $J(\pi^e)$  for an  $M$ -memory evaluation policy  $\pi^e$ . Toward this end, we define a state-value function under  $\pi^e$ :

$$V^{\pi^e}(z, s) := \mathbb{E}_{\pi^e} \left[ \sum_{k=0}^{\infty} \gamma^k R_k \mid Z_0 = z, S_0 = s \right]$$

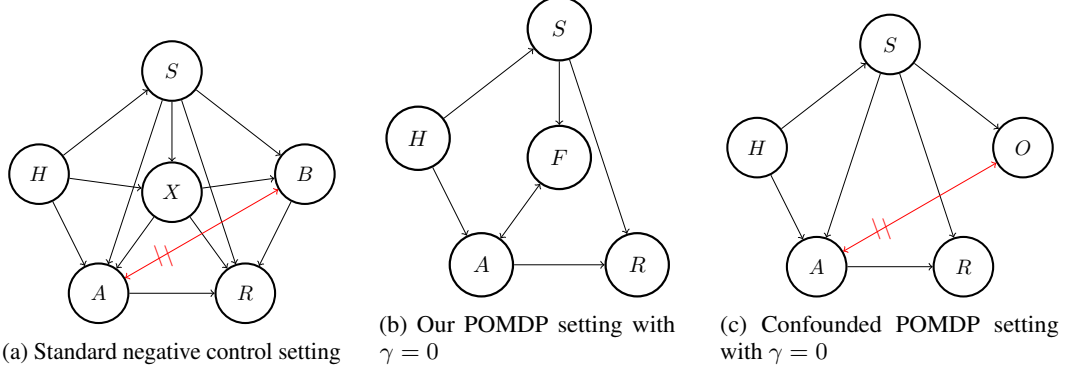


Figure 3: Comparisons of three DAGs.  $\parallel$  means no arrows. For simplicity, we consider the case with  $\gamma = 0$ , i.e., we do not need the next transitions. The first graph is a graph used in [CPS+20]. Note  $H, B, X$  are an action negative control, a reward negative control, and a covariate, respectively. We need  $(H, A) \perp B \mid S, X$  and  $H \perp Y \mid S, X, A$ . The graph is one instance satisfying this condition. The second graph corresponds to the contextual bandit version ( $\gamma = 0$ ) of our setting. Future proxies  $F$  just serve as a covariate and  $H$  serves as an action negative control. There are no nodes that correspond to reward negative controls. The third graph corresponds to the contextual bandit version ( $\gamma = 0$ ) of confounded POMDPs [TSM20, SUHJ22, BK21]. A node  $O$  corresponds to a reward negative control, and  $H$  corresponds to an action negative control that satisfies  $(H, A) \perp O \mid S$ . Thus,  $O$  cannot include futures proxies ( $F$ ) since then we cannot ensure  $(H, A) \perp F \mid S$  since there is an arrow from  $A$  to  $F$ .




Figure 4: POMDPs when  $M = 2$ . Note  $S_0, O_0, A_0, R_0$  correspond to  $(S, O, A, R)$ , respectively, and  $(O_{-2}, A_{-2}, O_{-1}, A_{-1}) = Z_0$ . We cannot observe  $S$  in the offline data.

for any  $z \in \mathcal{Z}, s \in \mathcal{S}$ . Compared to the memory-less case, the input additionally includes  $z$ . Let  $\bar{S}_t = (Z_t, S_t)$ , and  $d_t^{\pi^e}(\cdot)$  be the marginal distribution of  $\bar{S}_t$  under the policy  $\pi^e$ .

Next, we explain how the offline data is collected when behavior policies are  $M$ -memory. Specifically, the dataset  $\mathcal{D}_{\text{tra}}$  consists of  $n$  data tuples  $\{(H^{(i)}, O^{(i)}, A^{(i)}, R^{(i)}, F'^{(i)})\}_{i=1}^n$ . We use  $(H, O, A, R, F')$  to denote a generic history-observation-action-reward-future tuple where  $H$  denotes the  $M_H$ -step historical observations obtained prior to the observation  $O$  and  $F'$  denotes the  $M_F$ -step future observations after  $(O, A)$  for some integers  $M_H > M$  and  $M_F \geq 1$ . Hence, given some time step  $t$  in the trajectory data, we set  $(O, A, R) = (O_t, A_t, R_t)$ ,

$$H = (O_{-M_H:t-1}, A_{-M_H:t-1}) \text{ and } F' = (O_{t+1:t+M_F}, A_{t+1:t+M_F-1}).$$

We additionally set  $F = (O_{t:t+M_F-1}, A_{t:t+M_F-2})$ . We use the prime symbol  $'$  to represent the next time step. Then,  $Z' = (O_{t-M+1:t}, A_{t-M+1:t})$ . See Figure 5 for details when we set  $t = 0$ .

Throughout this paper, uppercase letters such as  $(H, S, O, A, R, S', F')$  are reserved for *random variables* and lower case letters such as  $(h, s, o, a, r, s', f')$  are reserved for their *realizations*. For




Figure 5: Case with  $M_H = 3, M = 2, M_F = 2$ . A 2-memory policy determines action  $A$  based on  $(Z, O)$ .

---

### Algorithm 2 Minimax OPE on POMDPs

---

**Require:** Dataset  $\mathcal{D}$ , function classes  $\mathcal{Q} \subset [\bar{\mathcal{F}} \rightarrow \mathbb{R}], \Xi \subset [\mathcal{H} \rightarrow \mathbb{R}]$ , hyperparameter  $\lambda \geq 0$

1:  $\hat{b}_V = \operatorname{argmin}_{q \in \mathcal{Q}} \max_{\xi \in \Xi} \mathbb{E}_{\mathcal{D}}[\{\mu(Z, A, O)\{R + \gamma q(\bar{F}')\} - q(\bar{F}')\}\xi(H) - \lambda \xi^2(H)]$ .

2: **return**  $\hat{J}_{\text{VM}} = \mathbb{E}_{\mathcal{D}_{\text{ini}}}[\hat{b}_V(\bar{f})]$

---

simplicity, we impose the stationarity assumption, i.e., the marginal distributions of  $(H, S, F)$  and  $(H', S', F')$  are the same.

The dataset  $\mathcal{D}_{\text{ini}}$  consists of  $N'$  data tuples  $\{Z_0^{(i)}, O_{0:M_F-1}^{(i)}, A_{0:M_F-1}^{(i)}\}_{i=1}^{N'}$  which is generated as follows:  $\bar{S}_0 \sim \nu_S, O_0 \sim \mathbb{O}(\cdot | S_0), A_0 \sim \pi^b(\cdot | Z_0, O_0), S_1 \sim \mathbb{T}(\cdot | S_0, A_0), \dots$ , until we observe  $O_{M_F-1}^{(i)}$  and  $A_{M_F-1}^{(i)}$ . We denote its distribution over  $\bar{\mathcal{F}} = \mathcal{Z} \times \mathcal{F}$  by  $\nu_{\bar{\mathcal{F}}}(\cdot)$ .

**Notation.** We denote the domain of  $Z$  by  $\mathcal{Z} = \mathcal{O}^M \times \mathcal{A}^M$ . We define  $\bar{S} = (Z, S), \bar{F} = (Z, F)$ .

### C.2 Required changes in Section 3

Every definition and statement holds by replacing  $F, S, F', \mathcal{F}, \mathcal{S}$  with  $\bar{F}, \bar{S}, \bar{F}', \bar{\mathcal{F}}, \bar{\mathcal{S}}$ , respectively. For completeness, we show these definitions and theorems tailored to  $M$ -memory policies.

**Definition 4** (Future-dependent value functions). *Future-dependent value functions*  $g_V \in [\bar{\mathcal{F}} \rightarrow \mathbb{R}]$  are defined such that the following holds almost surely,

$$\mathbb{E}[g_V(\bar{F}) | \bar{S}] = V^{\pi^e}(\bar{S}).$$

**Definition 5** (Learnable future-dependent value functions). Define  $\mu(Z, O, A) := \pi^e(A | Z, O) / \pi^b(A | Z, O)$ . *Learnable future-dependent value functions*  $b_V \in [\bar{\mathcal{F}} \rightarrow \mathbb{R}]$  are defined such that the following holds almost surely,

$$0 = \mathbb{E}[\mu(Z, O, A)\{R + \gamma b_V(\bar{F}')\} - b_V(\bar{F}) | H]. \quad (10)$$

We denote the set of solutions by  $\mathcal{B}_V$ .

**Theorem 4** (Identification Theorem). *Suppose (a) the existence of learnable future-dependent value functions (need not be unique); (b) the invertibility condition, i.e., any  $g : \bar{\mathcal{S}} \rightarrow \mathbb{R}$  that satisfies  $\mathbb{E}[g(\bar{S}) | H] = 0$  must also satisfy  $g(\bar{S}) = 0$  (i.e.,  $g(\bar{s}) = 0$  for almost every  $\bar{s}$  that belongs to the support of  $\bar{S}$ ), (c) the overlap condition  $w_{\pi^e}(\bar{S}) < \infty, \mu(Z, O, A) < \infty$ . Then, for any  $b_V \in \mathcal{B}_V$ , we have*

$$J(\pi^e) = \mathbb{E}_{\bar{f} \sim \nu_{\bar{\mathcal{F}}}}[b_V(\bar{f})]. \quad (11)$$

### C.3 Required changes in Section 4

Every algorithm holds by replacing  $F, S, F', \mathcal{F}, \mathcal{S}$  with  $\bar{F}, \bar{S}, \bar{F}', \bar{\mathcal{F}}, \bar{\mathcal{S}}$ , respectively. For completeness, we show the modified version of [Algorithm 1](#) in [Algorithm 2](#).

## C.4 Required change in Section 5

We present the modified version of [Theorem 3](#) tailored to  $M$ -memory policies.

**Definition 6** (Bellman operators). *The Bellman residual operator onto the history is defined as*

$$\mathcal{T} : [\bar{\mathcal{F}} \rightarrow \mathbb{R}] \ni q(\cdot) \mapsto \mathbb{E}[\mu(Z, O, A)\{R + \gamma q(\bar{F}')\} - q(\bar{F}) \mid H = \cdot],$$

and the Bellman residual error onto the history is defined as  $\mathbb{E}[(\mathcal{T}q)^2(H)]$ . Similarly, the Bellman residual operator onto the latent state,  $\mathcal{T}^S$  is defined as

$$\mathcal{T}^S : [\bar{\mathcal{F}} \rightarrow \mathbb{R}] \ni q(\cdot) \mapsto \mathbb{E}[\mu(Z, O, A)\{R + \gamma q(\bar{F}')\} - q(\bar{F}) \mid \bar{S} = \cdot],$$

and the Bellman residual error onto the latent state is defined as  $\mathbb{E}[\{\mathcal{T}^S(q)\}^2(\bar{S})]$ .

**Theorem 5** (Finite sample property of  $\hat{b}_V$ ). *Set  $\lambda > 0$ . Suppose (5a)  $\mathcal{B}_V \cap \mathcal{Q} \neq \emptyset$  (realizability) and (5b)  $\mathcal{T}\mathcal{Q} \subset \Xi$  (Bellman completeness). With probability  $1 - \delta$ ,*

$$\mathbb{E}[\{\mathcal{T}\hat{b}_V\}^2(H)] \leq c\{1/\lambda + \lambda\} \max(1, C_{\mathcal{Q}}, C_{\Xi}) \sqrt{\frac{\ln(|\mathcal{Q}||\Xi|c/\delta)}{n}},$$

where  $c$  is some universal constant.

**Theorem 6** (Finite sample property of  $\hat{J}_{VM}$ ). *Set  $\lambda > 0$ . Suppose (5a), (5b), (5c) any element in  $q \in \mathcal{Q}$  that satisfies  $\mathbb{E}[\{\mathcal{T}^S(q)\}(\bar{S}) \mid H] = 0$  also satisfies  $\mathcal{T}^S(q)(\bar{S}) = 0$ . (5d) the overlap  $\mu(Z, O, A) < \infty$  and any element in  $q \in \mathcal{Q}$  that satisfies  $\mathcal{T}^S(q)(\bar{S}) = 0$  also satisfies  $\mathcal{T}^S(q)(\bar{S}^\circ) = 0$  where  $\bar{S}^\circ \sim d_{\pi^e}(\bar{s})$ . With probability  $1 - \delta$ , we have*

$$|J(\pi^e) - \hat{J}_{VM}| \leq c(1 - \gamma)^{-2}(1/\lambda + \lambda) \max(1, C_{\mathcal{Q}}, C_{\Xi}) \text{IV}_1(\mathcal{Q}) \text{Dr}_{\mathcal{Q}}[d_{\pi^e}, P_{\pi^b}] \sqrt{\frac{\ln(|\mathcal{Q}||\Xi|c/\delta)}{n}}, \quad (12)$$

where

$$\text{IV}_1^2(\mathcal{Q}) := \sup_{\{q \in \mathcal{Q}; \mathbb{E}[\{\mathcal{T}(q)(H)\}^2] \neq 0\}} \frac{\mathbb{E}[\{\mathcal{T}^S(q)(\bar{S})\}^2]}{\mathbb{E}[\{\mathcal{T}(q)(H)\}^2]}, \quad (13)$$

$$\text{Dr}_{\mathcal{Q}}[d_{\pi^e}, P_{\pi^b}] := \sup_{\{q \in \mathcal{Q}; \mathbb{E}_{\bar{s} \sim P_{\pi^b}}[\{\mathcal{T}^S(q)(\bar{s})\}^2] \neq 0\}} \frac{\mathbb{E}_{\bar{s} \sim d_{\pi^e}}[\{\mathcal{T}^S(q)(\bar{s})\}^2]}{\mathbb{E}_{\bar{s} \sim P_{\pi^b}}[\{\mathcal{T}^S(q)(\bar{s})\}^2]}. \quad (14)$$

## D Examples

### D.1 Tabular POMDPs

We have seen that in [Lemma 2](#),  $\text{rank}(\text{Pr}_{\pi^b}(\mathbf{S}_b, \mathbf{H})) = |\mathcal{S}_b|$  and  $\text{rank}(\text{Pr}_{\pi^b}(\mathbf{F} \mid \mathbf{S}_b)) = |\mathcal{S}_b|$  are sufficient conditions for the identification in the tabular setting. The following theorem show that the abovementioned two conditions are equivalent to  $\text{rank}(\text{Pr}_{\pi^b}(\mathbf{F}, \mathbf{H})) = |\mathcal{S}_b|$ .

**Lemma 3.**  *$\text{rank}(\text{Pr}_{\pi^b}(\mathbf{S}_b, \mathbf{H})) = |\mathcal{S}_b|$  and  $\text{rank}(\text{Pr}_{\pi^b}(\mathbf{F} \mid \mathbf{S}_b)) = |\mathcal{S}_b|$  holds if and only if  $\text{rank}(\text{Pr}_{\pi^b}(\mathbf{F}, \mathbf{H})) = |\mathcal{S}_b|$ .*

We again make a few remarks. First,  $\text{rank}(\text{Pr}_{\pi^b}(\mathbf{F}, \mathbf{H})) = |\mathcal{S}|$  is often imposed to model HMMs and POMDPs [[HKZ12](#), [BSG11](#)]. Here, our condition  $\text{rank}(\text{Pr}_{\pi^b}(\mathbf{F}, \mathbf{H})) = |\mathcal{S}_b|$  is weaker than this assumption. We discuss the connection to the aforementioned works in [Section F](#). Second, in the literature of online RL,  $\text{rank}(\text{Pr}_{\pi^b}(\mathbf{F}, \mathbf{S})) = |\mathcal{S}|$  is frequently imposed as well [[JKKL20](#), [LCSJ22](#)] although they don't impose assumptions associated with the history proxy  $H$ . In confounded POMDPs, [[NJ21](#), [SUHJ22](#)] imposed a closely-related assumption, namely,  $\text{rank}(\text{Pr}_{\pi^b}(\mathbf{F}, \mathbf{H}, a)) = |\mathcal{S}|$  for any  $a \in \mathcal{A}$  where  $\text{Pr}_{\pi^b}(\mathbf{F}, \mathbf{H}, a)$  is a matrix whose  $(i, j)$ -th element is  $\text{Pr}_{\pi^b}(F = x_i, H = x'_j, A = a)$  ( $\mathcal{F} = \{x_i\}, \mathcal{H} = \{x'_j\}$ ).

### D.2 HSE-POMDPs and LQGs

In this section, we primarily emphasize the identification results in HSE-POMDPs. Extending these results to the final sample result is straightforward.

**Refined identification theorem.** First, we describe the population version of [Theorem 3](#) as follows.

**Theorem 7** (Refined identification theorem). *Suppose (7a)  $\mathcal{B}_V \cap \mathcal{Q} \neq \emptyset$ , (7b) any element in  $q \in \mathcal{Q}$  that satisfies  $\mathbb{E}[\{\mathcal{T}^S(q)\}(S) \mid H] = 0$  also satisfies  $\mathcal{T}^S(q)(S) = 0$ . (7c) the overlap  $\mu(O, A) < \infty$  and any element in  $q \in \mathcal{Q}$  that satisfies  $\mathcal{T}^S(q)(S) = 0$  also satisfies  $\mathcal{T}^S(q)(S^\circ) = 0$  where  $S^\circ \sim d_{\pi^e}(s)$ . Under the above three conditions, for any  $b_V \in \mathcal{B}_V \cap \mathcal{Q}$ , we have*

$$J(\pi^e) = \mathbb{E}_{f \sim \nu_{\mathcal{F}}}[b_V(f)].$$

The proof is deferred to [Section L](#)

**HSE-POMDPS.** By using [Theorem 7](#), we can obtain a useful identification formula when we set  $\mathcal{Q}, \Xi$  to be linear models in HSE-POMDPS. We start with the definition.

**Example 5** (HSE-POMDPS with linear models). *Introduce features  $\phi_{\mathcal{F}} : \mathcal{F} \rightarrow \mathbb{R}^{d_{\mathcal{F}}}, \phi_{\mathcal{S}} : \mathcal{S}_b \rightarrow \mathbb{R}^{d_{\mathcal{S}}}, \phi_{\mathcal{H}} : \mathcal{H} \rightarrow \mathbb{R}^{d_{\mathcal{H}}}$  such that  $\|\phi_{\mathcal{F}}(\cdot)\| \leq 1, \|\phi_{\mathcal{S}}(\cdot)\| \leq 1, \|\phi_{\mathcal{H}}(\cdot)\| \leq 1$ . Letting  $\mathcal{Q}$  and  $\Xi$  be linear models, the existence of future-dependent value functions in  $\mathcal{Q}$  is ensured as follows under certain conditions. Then, the existence of learnable future-dependent value functions in  $\mathcal{Q}$ , (7a) is automatically satisfied.*

Next, we provide sufficient conditions for the realizability (7a).

**Lemma 4.** *Suppose (LM1):  $\mathbb{E}[\phi_{\mathcal{F}}(F) \mid S] = K_1 \phi_{\mathcal{S}}(S)$  for some  $K_1 \in \mathbb{R}^{d_{\mathcal{F}} \times d_{\mathcal{S}}}$ , (LM2):  $V^{\pi^e}(S)$  is linear in  $\phi_{\mathcal{S}}(S)$ , i.e.,  $V^{\pi^e}(S) \in \{w^\top \phi_{\mathcal{S}}(S) : w \in \mathbb{R}^{d_{\mathcal{S}}}, \|w\| \leq C_{\text{LM}}\}$  for some  $C_{\text{LM}} \in \mathbb{R}$ , (LM3) for any  $b \in \mathbb{R}^S$  such that  $\|b\| \leq C_{\text{LM}}$ , there exists  $a \in \mathbb{R}^{d_{\mathcal{F}}}, \|a\| \leq C_{\mathcal{Q}}$  such that  $a^\top K_1 \phi_{\mathcal{S}}(S) = b^\top \phi_{\mathcal{S}}(S)$ . Then, future-dependent value functions exist and belong to  $\mathcal{Q} = \{w^\top \phi_{\mathcal{F}}(\cdot) : w \in \mathbb{R}^{d_{\mathcal{F}}}, \|w\| \leq C_{\mathcal{Q}}\}$  for some  $C_{\mathcal{Q}} \in \mathbb{R}$ .*

The condition (LM1) requires the existence of a conditional mean embedding operator between  $\mathcal{F}$  and  $\mathcal{S}$ . This assumption is widely used to model PSRs, which include POMDPs and HMMs [[SHSF09](#), [BGG13](#)]. In addition, assumptions of this type are frequently imposed to model MDPs as well [[ZLKB20](#), [DJW20](#), [CO20](#), [HDL<sup>+</sup>21](#)]. (LM2) is realizability on the latent state space. (LM3) says the information of the latent space is not lost on the observation space. The condition  $\text{rank}(K_1) = d_{\mathcal{S}}$  is a sufficient condition; then, we can take  $C_{\mathcal{Q}} = C_{\text{LM}}/\sigma_{\min}(K_1)$ . In the tabular case, we set  $\phi_{\mathcal{F}}, \phi_{\mathcal{S}}, \phi_{\mathcal{H}}$  be one-hot encoding vectors over  $\mathcal{F}, \mathcal{S}_b, \mathcal{H}$ , respectively. Here, we remark that  $S$  in  $\phi_{\mathcal{S}}(S)$  is a random variable in the offline data; thus,  $\phi_{\mathcal{S}}(\cdot)$  needs to be just defined on the support of the offline data. Hence, (LM3) is satisfied when  $\text{rank}(\text{Pr}_{\pi^b}(\mathbf{F} \mid \mathbf{S}_b)) = |\mathcal{S}_b|$ .

Next, we see (7b) is satisfied as follows under certain conditions.

**Lemma 5.** *Suppose (LM1), (LM2), (LM4):  $\mathbb{E}[\mu(O, A)\phi_{\mathcal{S}}(S') \mid S]$  is linear in  $\phi_{\mathcal{S}}(S)$  and (LM5):*

$$\sup_{x \in \mathbb{R}^d, x^\top \mathbb{E}[\mathbb{E}[\phi_{\mathcal{S}}(S) \mid H] \mathbb{E}[\phi_{\mathcal{S}}(S) \mid H]^\top] x \neq 0} \frac{x^\top \mathbb{E}[\phi_{\mathcal{S}}(S)\phi_{\mathcal{S}}(S)^\top] x}{x^\top \mathbb{E}[\mathbb{E}[\phi_{\mathcal{S}}(S) \mid H] \mathbb{E}[\phi_{\mathcal{S}}(S) \mid H]^\top] x} < \infty, \quad (15)$$

hold. Then, (7b) is satisfied.

Condition (LM4) requires the existence of conditional mean embedding between  $S'$  and  $S$  under the distribution induced by a policy  $\pi^e$ . The condition (LM5) is satisfied when  $\text{rank}(\text{Pr}_{\pi^b}(\mathbf{S}_b, \mathbf{H})) = |\mathcal{S}_b|$  in the tabular setting.

Combining [Lemma 4](#) and [Lemma 5](#) with [Theorem 7](#), we obtain the following handy formula.

**Lemma 6** (Formula with linear models in HSE-POMDPS). *Suppose (LM1)-(LM5), (LM6): there exists a matrix  $K_2 \in \mathbb{R}^{d_{\mathcal{S}} \times d_{\mathcal{H}}}$  such that  $\mathbb{E}[\phi_{\mathcal{S}}(S) \mid H] = K_2 \phi_{\mathcal{H}}(H)$ , (LM7):  $\mu(Z, O, A) < \infty$  and*

$$\sup_{x \in \mathbb{R}^d, 0 \neq x^\top \mathbb{E}_{s \sim P_{\pi^b}}[\phi_{\mathcal{S}}(s)\phi_{\mathcal{S}}(s)^\top] x} \frac{x^\top \mathbb{E}_{s \sim d_{\pi^e}}[\phi_{\mathcal{S}}(s)\phi_{\mathcal{S}}(s)^\top] x}{x^\top \mathbb{E}_{s \sim P_{\pi^b}}[\phi_{\mathcal{S}}(s)\phi_{\mathcal{S}}(s)^\top] x} < \infty, \quad (16)$$

hold. Then, we have

$$J(\pi^e) = \mathbb{E}_{f \sim \nu_{\mathcal{F}}}[\phi_{\mathcal{F}}(f)]^\top \mathbb{E}[\phi_{\mathcal{H}}(H) \{ \phi_{\mathcal{F}}(F) - \gamma \mu(O, A) \phi_{\mathcal{F}}(F') \}^\top] + \mathbb{E}[\mu(O, A) R \phi_{\mathcal{H}}(H)]. \quad (17)$$

We imposed two additional assumptions in [Lemma 6](#). (LM6) is used to ensure the Bellman completeness assumption  $\mathcal{T}\mathcal{Q} \subset \Xi$ . (LM7) is similar to the overlap condition (7c) in linear models. It is characterized based on a relative condition number whose value is smaller than the density ratio. Similar assumptions are imposed in offline RL with fully observable MDPs as well [[XCJ<sup>+</sup>21](#), [ZWB21](#), [US21](#)].

**LQG.** Finally, we extend our result to the case of LQG.

**Example 1 (LQG).** *Linear Quadratic Gaussian (LQG) falls in the category of HSE-POMDPs. Suppose*

$$s_{t+1} = As_t + Ba_t + \epsilon_{1t}, \quad r_t = -s_t^\top Qs_t - a_t^\top Ra_t, \quad o_t = Cs_t + \epsilon_{2t}$$

where  $A, B, C, Q, R$  are matrices that parametrize models and  $\epsilon_{1t}$  and  $\epsilon_{2t}$  are Gaussian noises. Consider a linear evaluation policy  $\pi^e(a_t | o_t, z_t) = \mathbb{I}(a_t = F[o_t, z_t]^\top)$  for certain matrix  $F$ . Notice that linear policies are commonly used in LQG since the globally optimal policy is linear [Ber12]. Then, defining  $\phi_S(s) = (1, \{s \otimes s\}^\top)^\top$ ,  $\phi_{\mathcal{F}}(f) = (1, \{f \otimes f\}^\top)^\top$  and  $\phi_{\mathcal{H}}(h) = (1, \{h \otimes h\}^\top)^\top$ , the following holds.

**Lemma 7.** *In LQG, (LM1),(LM2), (LM4) are satisfied. When  $C$  is left-invertible, (LM3) holds.*

Thus, what we additionally need to assume is only (LM5), (LM6) and (LM7) in LQG.

## E Finite Horizon Off-Policy Evaluation

For completeness, we also consider estimation of  $J_T(\pi^e) = \mathbb{E}_{\pi^e}[\sum_{k=0}^{T-1} \gamma^k R_k]$  when the horizon is finite and the system dynamics are nonstationary. We first define value and learnable future-dependent value functions following Section 3. Let  $V_t^{\pi^e}(s) = \mathbb{E}_{\pi^e}[\sum_{k=t}^{\infty} \gamma^{k-t} R_k | S_t = s]$  denote the state value function.

**Definition 7** (Future-dependent value functions). *For  $t \in [T - 1]$ , future-dependent value functions  $\{g_V^{[t]}\}_{t=0}^{T-1}$  are defined as solutions to*

$$0 = \mathbb{E}[g_V^{[t]}(F) | S] - V_t^{\pi^e}(S)$$

and  $g_V^{[T]} = 0$ . We denote the set of  $g_V^{[t]}$  by  $\mathcal{G}_V^{[t]}$ .

**Definition 8** (Learnable future-dependent value functions). *For  $t \in [T - 1]$ , learnable future-dependent value functions  $\{b_V^{[t]}\}_{t=0}^{T-1}$  are defined as solutions to*

$$\mathbb{E}[\mu(O, A)\{R + \gamma b_V^{[t+1]}(F')\} - b_V^{[t]}(F) | H]$$

where  $b_V^{[T]} = 0$ . We denote the set of  $b_V^{[t]}$  by  $\mathcal{B}_V^{[t]}$ .

We define the following Bellman operator:

$$\mathcal{T}^{\mathcal{S}, t} : \prod_{t=0}^{T-1} [\mathcal{F} \rightarrow \mathbb{R}] \ni \{q_t(\cdot)\} \mapsto \mathbb{E}[\mu(O, A)\{R + \gamma q_{t+1}(F')\} - q_t(F) | S = \cdot] \in [\mathcal{S} \rightarrow \mathbb{R}].$$

We again remark that while the conditional expectation of the offline data is not defined on the outside of  $\mathcal{S}_b$  (the support of  $\mathcal{S}$ ) above, we just set 0 on the outside of  $\mathcal{S}_b$ .

Here are the analogs statements of Theorem 7 in the finite horizon setting.

**Lemma 8.** *Future-dependent value functions are learnable future-dependent value functions.*

**Theorem 8** (Identification for finite horizon OPE). *Suppose for any  $t \in [T - 1]$ , (8a)  $\mathcal{B}_V^{[t]} \cap \mathcal{Q}_t \neq \emptyset$ , (8b) for any  $q \in \mathcal{Q}_t$  that satisfies  $\mathbb{E}[\{\mathcal{T}^{\mathcal{S}, t}(q)\}(S) | H] = 0$  also satisfies  $\{\mathcal{T}^{\mathcal{S}, t}(q)\}(S) = 0$ , (8c) overlap  $\mu(O, A) < \infty$  and for any  $q \in \mathcal{Q}_t$  that satisfies  $\{\mathcal{T}^{\mathcal{S}, t}(q)\}(S) = 0$  also satisfies  $\{\mathcal{T}^{\mathcal{S}, t}(q)\}(S_t^\circ) = 0$  where  $S_t^\circ \sim d_t^{\pi^e}(\cdot)$ . Then, for any  $b_V^{[0]} \in \mathcal{B}_V^{[0]} \cap \mathcal{Q}_0$ , we have*

$$J_T(\pi^e) = \mathbb{E}_{f \sim \nu_{\mathcal{F}}}[b_V^{[0]}(f)].$$

Here, (8a), (8b), (8c) correspond to (7a), (7b), (7c), respectively. Lemma 8 is often useful to ensure (8a).

In the tabular setting, when we have  $\text{rank}(\text{Pr}_{\pi_b}(\mathbf{F}, \mathbf{H})) = |\mathcal{S}_b|$ , conditions (a) and (b) are satisfied. This is the same condition imposed in Example 1. When we use linear models for  $\mathcal{Q}_t$  and  $\Xi_t$ , we have the following corollary. This is the finite-horizon version of Lemma 6.

**Corollary 1** (Formula with linear models in HSE-POMDPs ). *Suppose (LM1), (LM2f)  $V_t^{\pi^e}(S)$  is linear in  $\phi_S(S)$  for any  $t \in [T-1]$ , (LM3), (LM4), (LM5), (LM6). Then under the overlap  $\text{Dr}_{\mathcal{Q}}(d_t^{\pi^e}, P_{\pi^b}) < \infty$  and  $\mu(O, A) < \infty$  for any  $t \in [T-1]$ . Starting from  $\theta_T = 0$ , we recursively define*

$$\theta_t = \mathbb{E}[\phi_{\mathcal{H}}^{\top}(H)\phi_{\mathcal{F}}(F)]^+ \mathbb{E}[\mu(O, A)\phi_{\mathcal{H}}(H)\{R + \gamma\theta_{t+1}^{\top}\phi_{\mathcal{F}}(F')\}].$$

Then,  $J_T(\pi^e) = \mathbb{E}_{f \sim \nu_{\mathcal{F}}}[\theta_0^{\top}\phi_{\mathcal{F}}(f)]$ .

**Remark 7** (Comparison to SIS). *When we have finite samples, the estimator is defined as in Section 4. Then, we can obtain the finite sample result as in Section 5. In this case, we can again possibly circumvent the curse of horizon. The error scales with the marginal ratio  $\max_{t \in [T-1]} \max_{s \in S}(d_t^{\pi^e}(s)/P_{\pi^b}(s))^{1/2}$ . Compared to SIS, the finite sample error does not directly incur the exponential dependence on the horizon.*

## F Modeling of System Dynamics

We have so far discussed how to estimate cumulative rewards under evaluation policies. In the literature on POMDPs [SBS<sup>+</sup>10, BSG11, BGG13, KJS15], we are often interested in learning of system dynamics. In this section, we discuss how our methods are extended to achieve this goal. We ignore rewards in this section. We assume policies are memory-less.

### F.1 Tabular Setting

Here, let  $S_0^{\diamond}, O_0^{\diamond}, A_0^{\diamond}, \dots$  be random variables under a memory-less evaluation policy  $\pi^e : \mathcal{O} \rightarrow \Delta(\mathcal{A})$ . Following [SBS<sup>+</sup>10, BSG11, HKZ12], we consider two common estimands:

$$\Pr_{\pi^e}(o_0, a_0, \dots, o_{T-1}, a_{T-1}) := \Pr_{\pi^e}(O_0^{\diamond} = o_0, A_0^{\diamond} = a_0, O_1^{\diamond} = o_1, \dots), \quad (18)$$

$$\Pr_{\pi^e}(\mathbf{O}_T \mid o_0, a_0, \dots, a_{T-1}) := \{\Pr_{\pi^e}(O_T^{\diamond} = x_i \mid O_0^{\diamond} = o_0, \dots, A_{T-1}^{\diamond} = a_{T-1})\}_{i=1}^{|\mathcal{O}|}, \quad (19)$$

given a sequence  $o_0 \in \mathcal{O}, \dots, a_{T-1} \in \mathcal{A}$ . Our goal is to estimate (18) and (19) from the offline data. To simplify the discussion, we first consider the tabular case. If we can model a  $|\mathcal{O}|$ -dimensional vector  $\Pr_{\pi^e}(o_0, \dots, a_{T-1}, \mathbf{O}_T) \in \mathbb{R}^{|\mathcal{O}|}$  where the entry indexed by  $x_i \in \mathbf{O}$  is  $\Pr_{\pi^e}(o_0, \dots, a_{T-1}, x_i)$ , the latter estimand is computed by normalizing a vector, i.e., dividing it over the sum of all elements. Therefore, we have  $\Pr_{\pi^e}(\mathbf{O}_T \mid o_0, a_0, \dots, a_{T-1}) \propto \Pr_{\pi^e}(o_0, \dots, a_{T-1}, \mathbf{O}_T)$ . Hereafter, we consider modeling  $\Pr_{\pi^e}(o_0, \dots, a_{T-1}, \mathbf{O}_T)$  instead of  $\Pr_{\pi^e}(\mathbf{O}_T \mid o_0, \dots, a_{T-1})$ .

To identify estimands without suffering from the curse of horizon, we would like to apply our proposal in the previous sections. Toward that end, we set rewards as the product of indicator functions

$$\Pr_{\pi^e}(o_0, a_0, \dots, o_{T-1}, a_{T-1}) = \mathbb{E} \left( \prod_{k=0}^{T-1} \mathbf{I}(O_k^{\diamond} = o_k, A_k^{\diamond} = a_k) \right).$$

Since this is a product but not a summation, we cannot directly use our existing results. Nevertheless, the identification strategy is similar.

We first introduce learnable future-dependent value functions. These are analogs of Definition 2 tailored to the modeling of system dynamics. Let  $\Xi \subset [\mathcal{H} \rightarrow \mathbb{R}]$  and  $\mathcal{Q} \subset [\mathcal{F} \rightarrow \mathbb{R}]$ . Below, we fix  $o_0 \in \mathcal{O}, \dots, a_{T-1} \in \mathcal{A}$ .

**Definition 9** (Learnable future-dependent value functions for modeling dynamics). *Learnable future-dependent value functions  $\{b_D^{[t]}\}_{t=0}^{T-1}$  where  $b_D^{[t]} : \mathcal{F} \rightarrow \mathbb{R}$ , for joint observational probabilities are defined as solutions to*

$$\begin{aligned} 0 \leq t \leq T-1; \mathbb{E}[b_D^{[t]}(F) - \mathbf{I}(O = o_t, A = a_t)\mu(O, A)b_D^{[t+1]}(F') \mid H] &= 0, \\ \mathbb{E}[b_D^{[T]}(F) - 1 \mid H] &= 0. \end{aligned}$$

We denote the set of solutions  $b_D^{[t]}$  by  $\mathcal{B}_D^{[t]}$ . *Learnable future-dependent value functions  $\{b_P^{[t]}\}_{t=0}^{T-1}$  where  $b_P^{[t]} : \mathcal{F} \rightarrow \mathbb{R}^{|\mathcal{O}|}$  for conditional observational probabilities are defined as solutions to*

$$\begin{aligned} 0 \leq t \leq T-1; \mathbb{E}[b_P^{[t]}(F) - \mathbf{I}(O = o_t, A = a_t)\mu(O, A)b_P^{[t+1]}(F') \mid H] &= \mathbf{0}_{|\mathcal{O}|}, \\ \mathbb{E}[b_P^{[T]}(F) - \phi_{\mathcal{O}}(O) \mid H] &= \mathbf{0}_{|\mathcal{O}|}, \end{aligned}$$



where  $\phi_{\mathcal{O}}(\cdot)$  is a  $|\mathcal{O}|$ -dimensional one-hot encoding vector over  $\mathcal{O}$  and  $\mathbf{0}_{|\mathcal{O}|}$  is a  $|\mathcal{O}|$ -dimensional vector consisting of 0. We denote the set of solutions  $b_P^{[t]}$  by  $\mathcal{B}_P^{[t]}$ .

Next, we define the Bellman operator.

**Definition 10** (Bellman operators for modeling systems).

$$\mathcal{T}_t^{\mathcal{S}} : \prod_{t=0}^{T-1} [\mathcal{F} \rightarrow \mathbb{R}] \ni \{q_t(\cdot)\} \mapsto \mathbb{E}[q_t(F) - \mathbb{I}(O = o_t, A = a_t)\mu(O, A)q_{t+1}(F') \mid S = \cdot] \in [\mathcal{S} \rightarrow \mathbb{R}].$$

Following [Theorem 7](#), we can identify estimands. Here, let  $\dot{d}_t^{\pi}(\cdot) \in \Delta(\mathcal{S})$  be a probability density function of  $S_t^{\circ}$  conditional on  $O_0^{\circ} = o_0, \dots, A_T^{\circ} = a_{T-1}$ .

**Theorem 9** (Identification of joint probabilities). *Suppose (9a)  $\mathcal{B}_D^{[t]} \cap \mathcal{Q}_t \neq \emptyset$ , (9b) for any  $q \in \mathcal{Q}_t$  that satisfies  $\mathbb{E}[(\mathcal{T}_t^{\mathcal{S}} q)(S) \mid H] = 0$  also satisfies  $(\mathcal{T}_t^{\mathcal{S}} q)(S) = 0$ , (9c) for any  $q \in \mathcal{Q}_t$  that satisfies  $(\mathcal{T}_t^{\mathcal{S}} q)(S) = 0$  also satisfies  $(\mathcal{T}_t^{\mathcal{S}} q)(\dot{S}_t) = 0$  where  $\dot{S}_t \sim \dot{d}_t^{\pi}(\cdot)$  and  $\mu(O, A) < \infty$ . Then, for any  $b_D^{[0]} \cap \mathcal{Q}_0 \in \mathcal{B}_D^{[0]}$ , we have*

$$\Pr_{\pi^e}(o_0, a_0, \dots, o_{T-1}, a_{T-1}) = \mathbb{E}_{f \sim \nu_{\mathcal{F}}}[b_D^{[0]}(f)].$$

**Theorem 10** (Identification of conditional probabilities). *Suppose (10a)  $\mathcal{B}_P^{[t]} \cap \mathcal{Q}_t \neq \emptyset$ , (10b) for any  $q \in \mathcal{Q}_t$  that satisfies  $\mathbb{E}[\{\mathcal{T}_t^{\mathcal{S}}(q)\}(S) \mid H] = 0$  also satisfies  $(\mathcal{T}_t^{\mathcal{S}} q)(S) = 0$ , (10c) for any  $q \in \mathcal{Q}_t$  that satisfies  $(\mathcal{T}_t^{\mathcal{S}} q)(S) = 0$  also satisfies  $(\mathcal{T}_t^{\mathcal{S}} q)(\dot{S}_t)$  where  $\dot{S}_t \sim \dot{d}_t^{\pi}(\cdot)$  and  $\mu(O, A) < \infty$ . Then, for any  $b_P^{[0]} \cap \mathcal{Q}_0 \in \mathcal{B}_P^{[0]}$ , we have*

$$\Pr_{\pi^e}(o_0, a_0, \dots, o_{T-1}, a_{T-1}, \mathbf{O}_T) = \mathbb{E}_{f \sim \nu_{\mathcal{F}}}[b_P^{[0]}(f)].$$

The following corollary is an immediate application of [Theorem 9](#) and [Theorem 10](#).

**Corollary 2** (Tabular Models). *Let  $\phi_{\mathcal{F}}(\cdot), \phi_{\mathcal{H}}(\cdot)$  be one-hot encoding vectors over  $\mathcal{F}$  and  $\mathcal{H}$ , respectively. Suppose  $\text{rank}(\Pr(\mathbf{F}, \mathbf{H})) = |\mathcal{S}_b|$  and  $\dot{d}_t^{\pi^e}(\cdot)/P_{\pi^b}(\cdot) < \infty$  for any  $t \in [T-1]$  where  $P_{\pi^b}(\cdot)$  is a pdf of  $\mathcal{S}$  in the offline data. Then, we have*

$$\Pr_{\pi^e}(o_0, a_0, \dots, o_{T-1}, a_{T-1}) = \Pr_{\pi^b}(\mathbf{H})^{\top} B^+ \left\{ \prod_{t=T-1}^0 D_t B^+ \right\} C, \quad (20)$$

$$\Pr_{\pi^e}(\mathbf{O}_T \mid o_0, a_0, \dots, o_{T-1}, a_{T-1}) \propto \Pr_{\pi^b}(\mathbf{O}, \mathbf{H}) B^+ \left\{ \prod_{t=T-1}^0 D_t B^+ \right\} C, \quad (21)$$

where  $B = \Pr_{\pi^b}(\mathbf{F}, \mathbf{H})$ ,  $D_t = \mathbb{E}[\mathbb{I}(O = o_t, A = a_t)\mu(O, A)\phi_{\mathcal{F}}(F')\phi_{\mathcal{H}}^{\top}(H)]$ ,  $C = \Pr_{\nu_{\mathcal{F}}}$ .

In particular, when behavior policies are uniform policies and evaluation policies are atomic i.e.,  $\pi_t^e(a) = \mathbb{I}(a = a_t)$  for some  $a_t$  and any  $t$ , we have  $D_t = \Pr_{\pi^b}(O = o_t, \mathbf{F}', \mathbf{H} \mid A = a_t)$ . In addition, the rank assumption is reduced to  $\text{rank}(\Pr_{\pi^b}(\mathbf{F}, \mathbf{H})) = |\mathcal{S}_b|$ .

**Linear models.** Next, we consider cases where  $\mathcal{Q}_t, \Xi_t$  are linear models as in [Example 2](#). We first define value functions:

$$V_{D, [t]}^{\pi^e}(\cdot) = \Pr_{\pi^e}(O_t^{\circ} = o_t, A_t^{\circ} = a_t, \dots, O_T^{\circ} = O_T, A_T^{\circ} = a_T \mid S_t^{\circ} = \cdot),$$

$$V_{P, [t]}^{\pi^e}(\cdot) = \{\Pr_{\pi^e}(O_t^{\circ} = o_t, A_t^{\circ} = a_t, \dots, O_T^{\circ} = x_t \mid S_t^{\circ} = \cdot)\}_{i=1}^{|\mathcal{O}|}.$$

Then, we can obtain the following formula as in [Lemma 6](#). Here,  $\text{Dr}_{\mathcal{Q}}(\dot{d}_t^{\pi^e}, P_{\pi^b})$  is the condition number in [Lemma 6](#).

**Corollary 3** (Formula with linear models in HSE-POMDPs). *Suppose (LM1), (LM2D)  $V_{D, [t]}^{\pi^e}(S)$  is linear in  $\phi_{\mathcal{S}}(S)$ , (LM3), (LM4D)  $\mathbb{E}[\mu(O, A)\mathbb{I}(O = o_t, A = a_t)\phi_{\mathcal{F}}(F') \mid S]$ ,  $\mathbb{E}[1 \mid S]$  is linear in  $\phi_{\mathcal{S}}(S)$ , (LM5) and (LM6D)  $\text{Dr}_{\mathcal{Q}}(\dot{d}_t^{\pi^e}, P_{\pi^b}) < \infty$ . Then, we have*

$$\Pr_{\pi^e}(o_0, a_0, \dots, o_{T-1}, a_{T-1}) = \mathbb{E}[\phi_{\mathcal{H}}(H)]^{\top} B^+ \left\{ \prod_{t=T-1}^0 D_t B^+ \right\} C \quad (22)$$

---

**Algorithm 3** Minimax Modeling of Dynamics on POMDPs
 

---

**Require:** Dataset  $\mathcal{D}$ , function classes  $\mathcal{Q} \subset [\mathcal{F} \rightarrow \mathbb{R}]$ ,  $\Xi \subset [\mathcal{H} \rightarrow \mathbb{R}]$ , hyperparameter  $\lambda \geq 0$ , Horizon  $T$

1: Set

$$\hat{b}_D^{[T]} = \operatorname{argmin}_{q \in \mathcal{Q}} \max_{\xi \in \Xi} \mathbb{E}_{\mathcal{D}} [\{1 - q(F)\}\xi(H) - 0.5\lambda\xi^2(H)].$$

2: **for**  $t = T - 1$  **do**

3:

$$\hat{b}_D^{[t]} = \operatorname{argmin}_{q \in \mathcal{Q}} \max_{\xi \in \Xi} \mathbb{E}_{\mathcal{D}} [\{\mathbb{I}(O = o_t, A = a_t)\mu(O, A)\hat{b}_D^{[t+1]}(F') - q(F)\}\xi(H) - 0.5\lambda\xi^2(H)]. \quad (24)$$

4:  $t \leftarrow t - 1$

5: **end for**

6: **return**  $\hat{J}_{\text{VM}} = \mathbb{E}_{f \sim \nu_{\mathcal{F}}} [\hat{b}_D^{[0]}(f)]$

---

where

$$B = \mathbb{E}[\phi_{\mathcal{F}}(F)\phi_{\mathcal{H}}(H)^{\top}], D_t = \mathbb{E}[\mathbb{I}(O = o_t, A = a_t)\mu(O, A)\phi_{\mathcal{F}}(F)\phi_{\mathcal{H}}(H)^{\top}], C = \mathbb{E}_{f \sim \nu_{\mathcal{F}}}[\phi_{\mathcal{F}}(f)].$$

Suppose (LM1), (LM2P)  $V_{P, [t]}^{\pi^e}(\cdot)$  is linear in  $\phi_{\mathcal{S}}(\cdot)$ , (LM3), (LM4P)  $\mathbb{E}[\mu(O, A)\mathbb{I}(O_t = o_t, A_t = a_t)\phi_{\mathcal{F}}(F') \mid S]$ ,  $\mathbb{E}[\phi_{\mathcal{O}}(O) \mid S]$  is linear in  $\phi_{\mathcal{S}}(S)$ , (LM5) and (LM6P)  $\text{Dr}_{\mathcal{Q}}(d_t^{\pi^e}, P_{\pi^b}) < \infty$ . Then,

$$\Pr_{\pi^e}(\mathbf{O}_T \mid o_0, a_0, \dots, o_{T-1}, a_{T-1}) \propto \mathbb{E}[\phi_{\mathcal{O}}(O)\phi_{\mathcal{H}}^{\top}(H)]B^+ \left\{ \prod_{t=T-1}^0 D_t B^+ \right\} C. \quad (23)$$

When behavior policies are uniform, the formulas in (20) and (22) are essentially the same to those obtained via spectral learning [HKZ12, BSG11]. We emphasize that (20)–(23) appear to be novel to the literature since we consider the offline setting.

**General function approximation.** Finally, we introduce an algorithm to estimate joint probabilities in the tabular case with general function approximation, summarized in Algorithm 3. The conditional probabilities can be similarly estimated. We remark that function approximation is extremely useful in large-scale RL problems.

## E.2 Non-Tabular Setting

We have so far focused on the tabular case. In this section, we consider the non-tabular case. Our goal is to estimate joint probabilities  $\Pr_{\pi^e}(o_0, a_0, \dots, a_{T-1})$  in (18) and

$$\mathbb{E}_{\pi^e}[\phi_{\mathcal{O}}(O_T) \mid o_0, a_0, \dots, a_{T-1}]$$

where  $\phi_{\mathcal{O}} : \mathcal{O} \rightarrow \mathbb{R}$ . When  $\phi_{\mathcal{O}}(\cdot)$  is a one-hot encoding vector, this is equivalent to estimating  $\Pr_{\pi^e}(\mathbf{O}_T \mid o_0, a_0, \dots, a_{T-1})$  in (19).

In the non-tabular case, Theorem 9 and Theorem 10 still hold by defining learnable future-dependent value functions  $b^{[t]} : \mathcal{F} \rightarrow \mathbb{R}$  as solutions to

$$\mathbb{E}[b^{[t]}(F) \mid H] = \mathbb{E}[\mu(O, A)b^{[t]}(F') \mid H, O = o_t, A = a_t]\Pr_{\pi^b}(O = o_t, A = a_t).$$

where  $b^{[t]}$  is either  $b_D^{[t]}$  or  $b_P^{[t]}$ . Then, Corollary 3 holds by just replacing  $D_t$  with

$$\mathbb{E}[\mu(O, A)\phi_{\mathcal{F}}(F)\phi_{\mathcal{H}}^{\top}(H) \mid O = o_t, A = a_t]\Pr_{\pi^b}(O = o_t, A = a_t).$$

When we have a finite sample of data, we need to perform density estimation for  $\Pr_{\pi^b}(O = o_t, A = a_t)$ . This practically leads to instability of estimators. However, when our goal is just to

estimate  $\mathbb{E}[\phi_{\mathcal{O}}(O) \mid o_0, a_0, \dots, a_{T-1}]$  up to some scaling constant as in [SBS<sup>+</sup>10], we can ignore  $\Pr_{\pi^b}(O = o_t, A = a_t)$ . Then, we obtain the following formula:

$$\mathbb{E}_{\pi^e}[\phi_{\mathcal{O}}(O_T) \mid o_0, a_0, \dots, a_{T-1}] \propto \mathbb{E}[\phi_{\mathcal{O}}(O)\phi_{\mathcal{H}}^\top(H)]B^+ \left\{ \prod_{t=T-1}^0 \mathbb{E}[\mu(O, A)\phi_{\mathcal{F}}(F)\phi_{\mathcal{H}}^\top(H) \mid O = o_t, A = a_t]B^+ \right\} C. \quad (25)$$

This formula is known in HMMs [SBS<sup>+</sup>10] and in POMDPs [BGG13] where behavior policies are open-loop.

## G Connection with Literature on Spectral Learning

We discuss the connection with previous literature in detail. [SBS<sup>+</sup>10, HKZ12] consider the modeling of HMMs with no action. In this case, our task is to model

$$\Pr(o_0, \dots, o_{T-1}) := \Pr(O_0 = o_0, \dots, O_{T-1} = o_{T-1})$$

and the predictive distribution:

$$\Pr(\mathbf{O}_T \mid o_0, \dots, o_{T-1}) := \{\Pr(O_T = x_i \mid O_0 = o_0, \dots, O_{T-1} = o_{T-1})\}_{i=1}^{|\mathcal{O}|}.$$

In the tabular case, the Corollary 2 is reduced to

$$\Pr(o_0, \dots, o_{T-1}) = \Pr(\mathbf{H})^\top \Pr(\mathbf{F}, \mathbf{H})^+ \left\{ \prod_{t=T-1}^0 \Pr(O = o_t, \mathbf{F}', \mathbf{H})\Pr(\mathbf{F}, \mathbf{H})^+ \right\} \Pr_{\nu_{\mathcal{F}}}(\mathbf{F}_0),$$

$$\Pr(\mathbf{O}_T \mid o_0, \dots, o_{T-1}) \propto \Pr(\mathbf{O}, \mathbf{H})\Pr(\mathbf{F}, \mathbf{H})^+ \left\{ \prod_{t=T-1}^0 \Pr(O = o_t, \mathbf{F}', \mathbf{H})\Pr(\mathbf{F}, \mathbf{H})^+ \right\} \Pr_{\nu_{\mathcal{F}}}(\mathbf{F}_0).$$

This formula is reduced to the one in [HKZ12] when  $\mathcal{F} = \mathcal{H}$  and  $\Pr(\mathbf{H}) = \Pr_{\nu_{\mathcal{F}}}(\mathbf{F}_0)$  (stationarity). Here, the offline data consists of three random variables  $\{O_{-1}, O_0, O_1\}$ . In this case, the above formulae are

$$\Pr(o_0, \dots, o_{T-1}) = \Pr(\mathbf{O}_{-1})^\top \Pr(\mathbf{O}_0, \mathbf{O}_{-1})^+ \left\{ \prod_{t=T-1}^0 \Pr(O_0 = o_t, \mathbf{O}_1, \mathbf{O}_{-1})\Pr(\mathbf{O}_0, \mathbf{O}_{-1})^+ \right\} \Pr(\mathbf{O}_{-1}),$$

$$\Pr(\mathbf{O}_T \mid o_0, \dots, o_{T-1}) \propto \Pr(\mathbf{O}_0, \mathbf{O}_{-1})\Pr(\mathbf{O}_0, \mathbf{O}_{-1})^+ \left\{ \prod_{t=T-1}^0 \Pr(O_0 = o_t, \mathbf{O}_1, \mathbf{O}_{-1})\Pr(\mathbf{O}_0, \mathbf{O}_{-1})^+ \right\} \Pr(\mathbf{O}_{-1}).$$

Next, we consider the case when we use linear models to estimate

$$\mathbb{E}[\phi_{\mathcal{O}}(O_T) \mid o_0, \dots, o_{T-1}]$$

up to some constant scaling. When there are no actions, the formula (25) reduces to

$$\mathbb{E}[\phi_{\mathcal{O}}(O)\phi_{\mathcal{H}}^\top(H)]\mathbb{E}[\phi_{\mathcal{F}}(F)\phi_{\mathcal{H}}^\top(H)]^+ \left\{ \prod_{t=T-1}^0 \mathbb{E}[\phi_{\mathcal{F}}(F')\phi_{\mathcal{H}}^\top(H) \mid O = o_t]\mathbb{E}[\phi_{\mathcal{F}}(F)\phi_{\mathcal{H}}^\top(H)]^+ \right\} \mathbb{E}_{f \sim \nu_{\mathcal{F}}}[\phi_{\mathcal{F}}(f)].$$

When  $\mathcal{F} = \mathcal{H}$ , the pdf of  $O_{-1}$  is the same as  $\nu_{\mathcal{F}}(\cdot)$  and the offline data consists of three random variables  $\{O_{-1}, O_0, O_1\}$ , the above is reduced to the one in [SBS<sup>+</sup>10] as follows:

$$\mathbb{E}[\phi_{\mathcal{O}}(O_0)\phi_{\mathcal{O}}^\top(O_{-1})]\mathbb{E}[\phi_{\mathcal{O}}(O_0)\phi_{\mathcal{O}}^\top(O_{-1})]^+ \left\{ \prod_{t=T-1}^0 \mathbb{E}[\phi_{\mathcal{O}}(O_1)\phi_{\mathcal{O}}^\top(O_{-1}) \mid O_0 = o_t]\mathbb{E}[\phi_{\mathcal{O}}(O_0)\phi_{\mathcal{O}}^\top(O_{-1})]^+ \right\} \mathbb{E}[\phi_{\mathcal{O}}(O_{-1})].$$

## H Omitted Experiment Details and Additional Results

This section provides additional implementation details and ablation results of the synthetic experiment.

### H.1 Sequential Importance Sampling (SIS)

We compare SIS [Pre00] as a baseline estimator. SIS is a non-parametric approach, which corrects the distribution shift between the behavior and evaluation policies by applying importance sampling as follows.

$$\hat{J}_{\text{SIS}}(\pi^e; \mathcal{D}) := \mathbb{E}_{\mathcal{D}} \left[ \sum_{t=1}^{\infty} \gamma^{t-1} \left( \prod_{t'=1}^t \frac{\pi^e(O_{t'} \mid A_{t'})}{\pi^b(O_{t'} \mid A_{t'})} \right) R_t \right]$$

SIS enables unbiased estimation when absolute continuity (i.e.,  $\forall(O, A) \in \mathcal{O} \times \mathcal{A}, \pi^e(O|A) > 0 \rightarrow \pi^b(O|A) > 0$ ) holds. However, as the importance weight grows exponentially as  $t$  becomes large, SIS suffers from high variance [LLTZ18, UHJ20, SUHJ22]. We also empirically verify that SIS incurs high estimation error due to variance in the experimental results.

## H.2 Evaluation Metrics

We use the same evaluation metrics with [SUHJ22]. Given the i.i.d. dataset  $\mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_m$ , the values estimated on them  $\hat{J}_1, \hat{J}_2, \dots, \hat{J}_m$ , and the true value  $J^{\pi^e}$ , we define the relative bias and relative MSE as follows:

$$\text{Bias}(\hat{J}; \mathcal{D}, \pi) := \left| \frac{1}{m} \sum_{j=1}^m \left( \frac{\hat{J}_j - J^{\pi^e}}{J^{\pi^e}} \right) \right|, \quad \text{MSE}(\hat{J}; \mathcal{D}, \pi) := \frac{1}{m} \sum_{j=1}^m \left( \frac{\hat{J}_j - J^{\pi^e}}{J^{\pi^e}} \right)^2$$

We use the above metrics to compare the performance of SIS, the naive baseline, and our proposal.

## H.3 CartPole Setting

**Environment.** Here, we describe the state, action, and reward settings of the CartPole environment. First, the 4-dimensional states of CartPole represent the position and velocity of the cart and the angle and angle velocity of the pole. The action space is  $\{0, 1\}$ , either pushing the cart to the left or right. To better distinguish the values among different policies, we used modified reward following [SUHJ22]. Specifically, we define the reward as

$$R = \frac{1}{2} \left( \left| 2.0 - \frac{x}{x_{\text{clip}}} \right| \cdot \left| 2.0 - \frac{\theta}{\theta_{\text{clip}}} \right| - 1.0 \right),$$

where  $x$  and  $\theta$  are the positions of Cart and angle of Pole.  $x_{\text{clip}}$  and  $\theta_{\text{clip}}$  are the thresholds such that the episode will terminate when either  $|x| \geq x_{\text{clip}}$  or  $|\theta| \geq \theta_{\text{clip}}$ . Under this definition, we observe a larger reward when the cart is closer to the center, and the pole stands straight. We also set the discount factor  $\gamma = 0.95$ , and the values of the policies used in our experiment are somewhere between 20 and 40.

**Estimator Implementation.** We parametrize the value function  $b_v(\cdot)$  of our proposal and the naive estimator with a two-layer neural network. The network uses a 100-dimensional hidden layer with ReLU as its activation function, and Adam [KB14] is its optimizer. Both the naive estimator and our proposal optimize  $b_v(\cdot)$  with the loss function defined in Example 5, but the naive one replaces  $\bar{F}$  and  $H$  with  $O$ . Specifically, the naive estimator takes 4-dimensional observation  $O$  as input. On the other side, our proposed estimator additionally inputs the concatenated vector of observation  $O$  and one-hot representation of  $A$  for several future steps (i.e.,  $M_F$ ) to consider. The convergence is based on the test loss evaluated on the test dataset, which is independent of the datasets to train value functions and estimate the policy value. Specifically, to find a global convergence point, we first run 20000 epochs with 10 gradient steps per each. Then, we report the results on the epoch which achieves the minimum loss. The convergence point is usually less than 10000 epochs.

For the adversarial function space  $\Xi$  of both methods, we use the following RBF kernel  $K(x_i; x_2)$ :

$$K(x_1; x_2) := \exp \left( -\frac{\|x_1 - x_2\|_2}{2\sigma^2} \right),$$

where  $\|x_1 - x_2\|_2$  is the l2-distance between  $x_1$  and  $x_2$ , and  $\sigma$  is a bandwidth hyperparameter. We use  $\sigma = 1.0$  in the main text and provide ablation results with varying values of  $\sigma$  in the following section. Finally, the naive estimator uses  $O$  as the input of the kernel. The proposed method first predicts the latent spaces from the historical observations as  $\hat{S} := f_{\text{LSTM}}(H)$  and then use  $\hat{S}$  as the input of the kernel.  $f_{\text{LSTM}}(H)$  is a bi-directional LSTM [CKPW18] with 10-dimensional hidden dimension. We train the LSTM with MSE loss in predicting the noisy state (i.e.,  $O$ ) with Adam [KB14] as its optimizer.

## H.4 Additional Results

Figure 6 shows the experimental results in the case of using the behavior policy with  $\epsilon = 0.1$ . The result suggests that the proposed estimator reduces MSEs of the baseline estimators as observed in




Figure 6: Logarithms of relative biases (left) and MSEs (right) of the proposed and the baseline estimators for various values of  $\epsilon$ , which specify the evaluation policy. The confidence intervals are obtained through 100 Monte Carlo simulations.

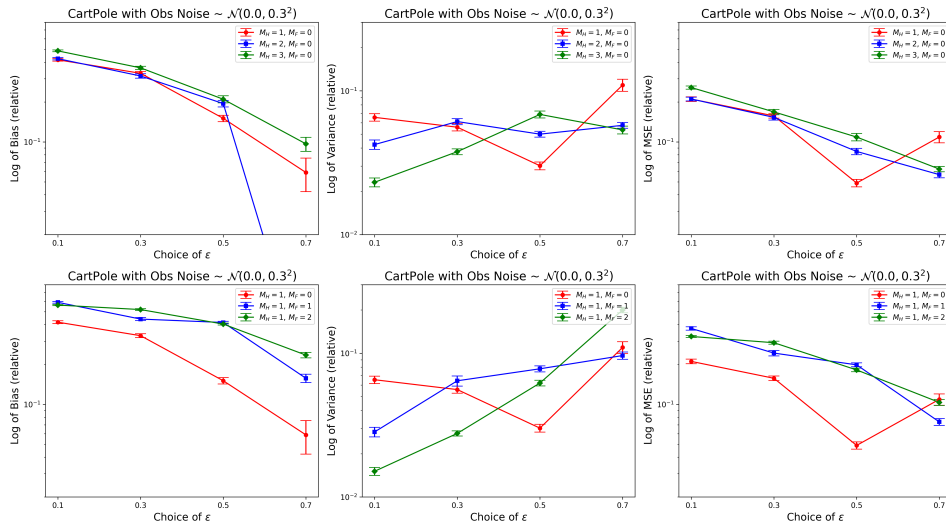


Figure 7: Logarithms of relative biases (left), variances (center), and MSEs (right) of the proposed estimator with varying lengths of history  $M_H$  (top) and varying lengths of future steps  $M_F$  (bottom). The x-axis corresponds to the varying values of  $\epsilon$  of the evaluation policy, and the associated confidence interval is based on 20 simulations.

the  $\epsilon = 0.3$  case in the main text. Note that experimental settings other than  $\epsilon$  of the behavior policy are the same as those used in Section 6.

## H.5 Ablation Results

Here, we provide ablation results with (1) varying choices of  $\bar{F}$  and  $H$ , and (2) varying values of bandwidth hyperparameter  $\sigma$  of RKHSs. We provide the results with 20 random seeds in the following to conduct comprehensive ablations.

The first set of experiments aims to study how the use of history and future observations help improve the accuracy of value estimation. For this, we compare our proposed method with varying values of history length  $M_H$  and future length  $M_F$ . Figure 7 (top) shows the result of varying history lengths  $M_H \in \{1, 2, 3\}$  with a fixed future length ( $M_F = 0$ ). We observe that the performance of the proposed method does not change greatly with the choice of history lengths. This result suggests that 1-step history is almost sufficient to identify the latent state in our experimental settings. Next, we report the results with varying future lengths  $M_F \in \{0, 1, 2\}$  with a fixed history length ( $M_H = 1$ ) in Figure 7 (bottom). The result suggests that the increased number of future steps can slightly increase bias. We attribute this to the increased estimation difficulty of the value function due to the increase in the dimensionality of inputs. However, we should also note that future observations may help




Figure 8: Logarithms of relative biases (left), variances (center), and MSEs (right) of the proposed estimator with varying bandwidth hyperparameter  $\sigma$  of RKHSs. The x-axis corresponds to the varying values of  $\epsilon$  of the evaluation policy, and the associated confidence interval is based on 20 simulations.

improve the performance of the proposed method when the history itself is insufficient to identify the latent state.

The second set of experiments is to see how robust the proposed method is to the choice of hyperparameter  $\sigma$ . We thus vary the values of  $\sigma \in [0.1, 0.5, 1.0]$ , and report the result in Figure 8. The result shows that the estimation accuracy is almost the same between  $\sigma = 1.0$  and  $\sigma = 0.5$ , suggesting that the proposed value learning method is robust to the change of bandwidth hyperparameter of RKHSs to some extent. On the other hand, we observe that a very small value (i.e.,  $\sigma = 0.1$ ) can increase the variance of estimation. Therefore, our recommendation is to avoid a (too) small value for the bandwidth hyperparameter  $\sigma$ .

## I Proof of Section 3

We often use

$$F' \perp S, O, A \mid S' \quad \text{and} \quad (H) \perp (A, O, R, F') \mid S.$$

This is easily checked by graphical models Figure 4.

### I.1 Proof of Lemma 1

From the Bellman equation, we have

$$\mathbb{E}[\mu(O, A)\{R + \gamma V^{\pi^e}(S')\} - V^{\pi^e}(S) \mid S] = 0.$$

Then, from the definition of future-dependent value functions,

$$\mathbb{E}[\mu(O, A)\{R + \gamma \mathbb{E}[g_V(F') \mid S']\} - \mathbb{E}[g_V(F) \mid S] \mid S] = 0.$$

Here, we use the stationarity assumption to ensure  $\mathbb{E}[g_V(F') \mid S'] = V^{\pi^e}(S')$ .

Next, by using  $F' \perp S, O, A \mid S'$ , we have

$$\mathbb{E}[\mu(O, A)\{R + \gamma g_V(F')\} - g_V(F) \mid S] = 0. \quad (26)$$

More specifically,

$$\begin{aligned} & \mathbb{E}[\mu(O, A)\{R + \gamma g_V(F')\} - g_V(F) \mid S] \\ &= \mathbb{E}[\mu(O, A)\{R + \gamma \mathbb{E}[g_V(F') \mid S', S, O, A, Z]\} - \mathbb{E}[g_V(F) \mid S] \mid S] \quad (\text{Law of total expectation}) \\ &= \mathbb{E}[\mu(O, A)\{R + \gamma \mathbb{E}[g_V(F') \mid S']\} - \mathbb{E}[g_V(F) \mid S] \mid S] \quad (\text{Use } F' \perp S, O, A \mid S') \\ &= 0. \end{aligned}$$

Hence,

$$\begin{aligned} & \mathbb{E}[\mu(O, A)\{R + \gamma g_V(F')\} - g_V(F) \mid H] \\ &= \mathbb{E}[\mathbb{E}[\mu(O, A)\{R + \gamma g_V(F')\} - g_V(F) \mid S, H] \mid H] \quad (\text{Law of total expectation}) \\ &= \mathbb{E}[\mathbb{E}[\mu(O, A)\{R + \gamma g_V(F')\} - g_V(F) \mid S] \mid H] \quad (\text{Use } R, O, A, F' \perp (H) \mid S) \\ &= 0. \quad (\text{From (26)}) \end{aligned}$$

Thus,  $g_V$  is a learnable future-dependent value function.

## I.2 Proof of Theorem 1

It follows from Theorem 7, which is an improved version of Theorem 1.

## I.3 Proof of Lemma 2

The first statement is straightforward noting the equation is equal to solving

$$x^\top \Pr_{\pi^b}(\mathbf{F} \mid \mathbf{S}_b) = y$$

for  $x$  given  $y$ .

The second statement is straightforward noting it is equivalent to

$$x^\top \Pr_{\pi^b}(\mathbf{S}_b \mid \mathbf{H}) = 0$$

implies  $x = 0$ . This is satisfied when  $\text{rank}(\Pr_{\pi^b}(\mathbf{S}_b, \mathbf{H})) = |\mathcal{S}_b|$ .

## J Proof of Section 5

### J.1 Proof of Theorem 2 and Theorem 3

By simple algebra, the estimator is written as

$$\hat{b}_V = \inf_{q \in \mathcal{Q}} \sup_{\xi \in \Xi} \mathbb{E}_{\mathcal{D}}[(\mathcal{Z}f)^2 - \{\mathcal{Z}f - \lambda\xi(H)\}^2]$$

where

$$\mathcal{Z}f = \mu(A, O)\{R + \gamma q(F')\} - q(F).$$

Noting this form similarly appears in the proof of [SUHJ22, Theorem 3], the following is similarly completed by [SUHJ22, Theorem 3]:

$$\mathbb{E}[\{\mathcal{T}(\hat{b}_V)\}^2(H)]^{1/2} = \tilde{O} \left( \max(1, C_{\mathcal{Q}}, C_{\Xi}) (1/\lambda + \lambda) \sqrt{\frac{\ln(|\mathcal{G}||\Xi|/\delta)}{n}} \right).$$

using realizability and the Bellman completeness. Then, we have

$$\begin{aligned} & (J(\pi^e) - \mathbb{E}_{f \sim \nu_{\mathcal{F}}}[\hat{b}_V(f)])^2 \\ &= \left\{ (1 - \gamma)^{-1} \mathbb{E}_{(\tilde{s}) \sim d_{\pi^e}} [\mathbb{E}[\mu(A, O)R \mid S = \tilde{s}]] - \mathbb{E}_{f \sim \nu_{\mathcal{F}}}[\hat{b}_V(f)] \right\}^2 \\ &= \left\{ (1 - \gamma)^{-1} \mathbb{E}_{(\tilde{s}) \sim d_{\pi^e}} \left[ \mathbb{E} \left[ \mu(A, O)\{R + \gamma \hat{b}_V(F')\} - \hat{b}_V(F) \mid S = \tilde{s} \right] \right] \right\}^2 \quad (\text{Use Lemma 9}) \\ &\leq (1 - \gamma)^{-2} \mathbb{E}[\{\mathcal{T}^S(\hat{b}_V)\}^2 s] \quad (\text{Jensen's inequality}) \\ &\leq (1 - \gamma)^{-2} \mathbb{E}[\{\mathcal{T}(\hat{b}_V)\}^2(H)] \times \sup_{q \in \mathcal{Q}} \frac{\mathbb{E}_{s \sim d_{\pi^e}}[(\mathcal{T}^S q)^2 s]}{\mathbb{E}[(\mathcal{T}q)^2(H)]}. \end{aligned}$$

## K Proof of Section C

In this section, most of the proof follows by slightly modifying the proof for memoryless policies. For completeness, we provide the proof of Theorem 4. As we did in the proof of Theorem 1, we prove the following stronger statement.

**Theorem 11** (Refined identification theorem). *Suppose (7a)  $\mathcal{B}_V \cap \mathcal{Q} \neq \emptyset$ , (7b) any element in  $q \in \mathcal{Q}$  that satisfies  $\mathbb{E}[\{\mathcal{T}^S(q)\}(\bar{S}) \mid H] = 0$  also satisfies  $\mathcal{T}^S(q)(\bar{S}) = 0$ . (7c) the overlap  $\mu(Z, O, A) < \infty$  and any element in  $q \in \mathcal{Q}$  that satisfies  $\mathcal{T}^S(q)(\bar{S}) = 0$  also satisfies  $\mathcal{T}^S(q)(\bar{S}^\circ) = 0$  where  $S^\circ \sim d_{\pi^e}(s)$ . Under the above three conditions, for any  $b_V \in \mathcal{B}_V \cap \mathcal{Q}$ , we have*

$$J(\pi^e) = \mathbb{E}_{\bar{f} \sim \nu_{\bar{\mathcal{F}}}}[b_V(\bar{f})].$$

### K.1 Proof of Theorem 11

Note for any  $b_V \in \mathcal{B}_V \cap \mathcal{Q}$ , we have

$$\begin{aligned} 0 &= \mathbb{E} [\mu(Z, A, O) (R + \gamma b_V(Z', F')) - b_V(Z, F) \mid H] \\ &= \mathbb{E} [\mathbb{E} [\mu(Z, A, O) (R + \gamma b_V(Z', F')) - b_V(Z, F) \mid Z, S, H] \mid H] \quad (\text{Law of total expectation}) \\ &= \mathbb{E} [\mathbb{E} [\mu(Z, A, O) (R + \gamma b_V(Z', F')) - b_V(Z, F) \mid Z, S] \mid H]. \end{aligned}$$

In the last line, we use  $(H \setminus Z) \perp (A, O, F') \mid Z, S$ . From (7b), we have

$$\mathbb{E}[b_V(Z, F) \mid Z, S] = \mathbb{E} [\mu(Z, A, O) (R + \gamma b_V(Z', F')) \mid Z, S]. \quad (27)$$

Hence,  $\mathcal{T}^S(b_V)(S) = 0$ . Then, from the overlap condition (7c),

$$\mathcal{T}^S(b_V)(S^\circ) = 0. \quad (28)$$

Finally, for any  $b_V \in \mathcal{B}_V \cap \mathcal{Q}$ , we have

$$\begin{aligned} &(J(\pi^e) - \mathbb{E}_{\bar{f} \sim \nu_{\bar{F}}} [b_V(\bar{f})]) \\ &= (1 - \gamma)^{-1} \mathbb{E}_{(\bar{s}) \sim d_{\pi^e}} [\mathbb{E} [\mu(Z, A, O) R \mid \bar{S} = \bar{s}]] - \mathbb{E}_{\bar{f} \sim \nu_{\bar{F}}} [b_V(\bar{f})] \\ &= (1 - \gamma)^{-1} \mathbb{E}_{(\bar{s}) \sim d_{\pi^e}} [\mathbb{E} [\mu(Z, A, O) \{R + \gamma b_V(Z', F')\} - b_V(Z, F) \mid \bar{S} = \bar{s}]] \\ &\hspace{15em} (\text{Use Lemma 9}) \\ &= 0. \hspace{15em} (\text{Use } \mathcal{T}^S(S^\circ) = 0.) \end{aligned}$$

From the first line to the second line, we use

$$\begin{aligned} J(\pi^e) &= (1 - \gamma)^{-1} \int d_{\pi^e}(z, s) r(s, a) \pi^e(a \mid z, o) d(z, s) \\ &= (1 - \gamma)^{-1} \mathbb{E}_{(\bar{s}) \sim d_{\pi^e}} [\mathbb{E} [\mu(Z, A, O) R \mid \bar{S} = \bar{s}]]. \end{aligned}$$

## L Proof of Section D

### L.1 Proof of Theorem 7

Note for any  $b_V \in \mathcal{B}_V \cap \mathcal{Q}$ , we have

$$\begin{aligned} 0 &= \mathbb{E} [\mu(A, O) (R + \gamma b_V(F')) - b_V(F) \mid H] \\ &= \mathbb{E} [\mathbb{E} [\mu(A, O) (R + \gamma b_V(F')) - b_V(F) \mid S, H] \mid H] \quad (\text{Law of total expectation}) \\ &= \mathbb{E} [\mathbb{E} [\mu(A, O) (R + \gamma b_V(F')) - b_V(F) \mid S] \mid H]. \end{aligned}$$

In the last line, we use  $(H) \perp (A, O, F') \mid S$ . From (7b), we have

$$\mathbb{E}[b_V(F) \mid S] = \mathbb{E} [\mu(A, O) (R + \gamma b_V(F')) \mid S]. \quad (29)$$

Hence,  $\mathcal{T}^S(b_V)(S) = 0$ . Then, from the overlap condition (7c),

$$\mathcal{T}^S(b_V)(S^\circ) = 0. \quad (30)$$

Finally, for any  $b_V \in \mathcal{B}_V \cap \mathcal{Q}$ , we have

$$\begin{aligned} &(J(\pi^e) - \mathbb{E}_{f \sim \nu_{\mathcal{F}}} [b_V(f)]) \\ &= (1 - \gamma)^{-1} \mathbb{E}_{(\bar{s}) \sim d_{\pi^e}} [\mathbb{E} [\mu(A, O) R \mid S = \bar{s}]] - \mathbb{E}_{f \sim \nu_{\mathcal{F}}} [b_V(f)] \\ &= (1 - \gamma)^{-1} \mathbb{E}_{(\bar{s}) \sim d_{\pi^e}} [\mathbb{E} [\mu(A, O) \{R + \gamma b_V(F')\} - b_V(F) \mid S = \bar{s}]] \quad (\text{Use Lemma 9}) \\ &= 0. \hspace{15em} (\text{Use } \mathcal{T}^S(S^\circ) = 0.) \end{aligned}$$

From the first line to the second line, we use

$$\begin{aligned} J(\pi^e) &= (1 - \gamma)^{-1} \int d_{\pi^e}(z, s) r(s, a) \pi^e(a \mid z, o) d(z, s) \\ &= (1 - \gamma)^{-1} \mathbb{E}_{(\bar{s}) \sim d_{\pi^e}} [\mathbb{E} [\mu(A, O) R \mid S = \bar{s}]]. \end{aligned}$$



## L.2 Proof of Lemma 4

From (LM2), there exists  $w_1$  such that  $V^{\pi^e}(S) = w_1^\top \phi_S(S)$ . Then, from (LM1),

$$\mathbb{E}[w_2^\top \phi_{\mathcal{F}}(F) \mid S] = w_2^\top K_1 \phi_S(S) = w_1^\top \phi_S(S).$$

From (LM3), the above equation has a solution with respect to  $w_2$ . This concludes the proof.

## L.3 Proof of Lemma 5

From (LM1), (LM4) and the statement of Lemma 4, there exists  $w_4 \in \mathbb{R}^{d_S}$  such that

$$\mathbb{E}[\mu(O, A)\{R + \gamma q(F')\} - q(F) \mid S] = w_4^\top \phi_S(S).$$

for any  $q(\cdot) = w^\top \phi_{\mathcal{F}}(\cdot) \in \mathcal{Q}$ . Letting  $b_F(\cdot) = \{w^*\}^\top \phi_{\mathcal{F}}(\cdot)$ , this is because

$$\begin{aligned} & \mathbb{E}[\mu(O, A)\{R + \gamma q(F')\} - q(F) \mid S] & (31) \\ &= \mathbb{E}[\mu(O, A)\{\gamma q(F') - \gamma b_V(F')\} - q(F) + b_V(F) \mid S] & \text{(Statement of Lemma 4)} \\ &= \mathbb{E}[\mu(O, A)\{\gamma(w^\top - \{w^*\}^\top)\mathbb{E}[\phi_{\mathcal{F}}(F') \mid S']\} - (w^\top - \{w^*\}^\top)\mathbb{E}[\phi_{\mathcal{F}}(F) \mid S] \mid S] \\ &= \mathbb{E}[\mu(O, A)\{\gamma(w^\top - \{w^*\}^\top)K_1 \phi_S(S')\} - (w^\top - \{w^*\}^\top)K_1 \phi_S(S) \mid S] & \text{(LM1)} \\ &= w_4^\top \phi_S(S). & \text{(LM4)} \end{aligned}$$

for some  $w_4$ .

Then, from (LM5), when  $w_4^\top \mathbb{E}[\phi_S(S) \mid H] = 0$ , we have  $w_4^\top \phi_S(S) = 0$ . This is because first  $\mathbb{E}[w_4^\top \phi_S(S) \mid H] = 0$  implies  $\mathbb{E}[\{w_4^\top \phi_S(S) \mid H\}^2] = 0$ . Then, to make the ratio

$$\sup_{w \in \mathbb{R}^d} \frac{\mathbb{E}[\{w^\top \phi_S(S)\}^2]}{\mathbb{E}[\{w^\top \mathbb{E}[\phi_S(S) \mid H]\}^2]}$$

finite, we need  $\mathbb{E}[\{w_4^\top \phi_S(S)\}^2] = 0$ . This implies  $w_4^\top \phi_S(S) = 0$ .

Here, when we have

$$\begin{aligned} 0 &= \mathbb{E}[\mathbb{E}[\mu(O, A)\{R + \gamma q(F')\} - q(F) \mid S] \mid H] \\ &= \mathbb{E}[w_4^\top \phi_S(S) \mid H], \end{aligned} \quad \text{(Just plug-in)}$$

we get  $w_4^\top \phi_S(S) = 0$ , i.e.,

$$\mathbb{E}[\mu(O, A)\{R + \gamma q(F')\} - q(F) \mid S] = 0.$$

## L.4 Proof of Lemma 6

Letting

$$\tilde{w} = \mathbb{E}[\phi_{\mathcal{H}}(H)\{\gamma \phi_{\mathcal{F}}(F') - \phi_{\mathcal{F}}(F)\}]^+ \mathbb{E}[\mu(O, A)R \phi_{\mathcal{H}}(H)],$$

we want to prove  $\tilde{w}^\top \phi_{\mathcal{F}}(\cdot)$  is a learnable future-dependent value function. Then, by invoking Theorem 7, the statement is proved.

**First step.** Here, for  $q \in \mathcal{Q}$ , we have  $(\mathcal{T}q)(H) = a^\top \phi_{\mathcal{H}}(H)$  for some vector  $a \in \mathbb{R}^{d_{\mathcal{H}}}$ . Here,  $\mathbb{E}[(\mathcal{T}q)^2(H)] = 0$  is equivalent to  $(\mathcal{T}q)(H) = 0$ . Besides, the condition  $\mathbb{E}[(\mathcal{T}q)^2(H)] = 0$  is equivalent to

$$a^\top \mathbb{E}[\phi_{\mathcal{H}}(H)\phi_{\mathcal{H}}^\top(H)]a = 0.$$

Thus,  $\mathbb{E}[(\mathcal{T}q)(H)\phi_{\mathcal{H}}^\top(H)] = a^\top \mathbb{E}[\phi_{\mathcal{H}}(H)\phi_{\mathcal{H}}^\top(H)] = \mathbf{0}$ , where  $\mathbf{0} \in \mathbb{R}^{d_{\mathcal{H}}}$  is a vector consisting of 0, is a sufficient condition to satisfy  $(\mathcal{T}q)(H) = \mathbf{0}$ . Hence, if  $q(\cdot) = w^\top \phi_{\mathcal{F}}(\cdot) \in \mathcal{Q}$  satisfies

$$\mathbb{E}[\phi_{\mathcal{H}}(H)\{\mu(O, A)\{R + \gamma q(F')\} - q(F)\}] = \mathbf{0}, \quad (32)$$

$q(\cdot)$  is a learnable bridge function. Note the above equation is equal to

$$\mathbb{E}[\phi_{\mathcal{H}}(H)\{\gamma \mu(O, A)\phi_{\mathcal{F}}(F') - \phi_{\mathcal{F}}(F)\}^\top]w = \mathbb{E}[\mu(O, A)R \phi_{\mathcal{H}}(H)].$$

Vice versa, i.e., any learnable future-dependent value function satisfies the above (32) is similarly proved.

**Second step.** Since there exists a linear learnable future-dependent value function in  $\mathcal{Q}$  from [Lemma 4](#), we have a solution to

$$\mathbb{E}[\phi_{\mathcal{H}}(H)\{\gamma\mu(O, A)\phi_{\mathcal{F}}(F') - \phi_{\mathcal{F}}(F)\}^{\top}]w = \mathbb{E}[\mu(O, A)R\phi_{\mathcal{H}}(H)]$$

with respect to  $w$ . We denote it by  $w$ . Thus,  $\tilde{w}$  is also a solution since

$$B\tilde{w} = BB^+Bw = Bw = \mathbb{E}[\mu(O, A)R\phi_{\mathcal{H}}(H)]$$

where  $B = \mathbb{E}[\phi_{\mathcal{H}}(H)\{\gamma\phi_{\mathcal{F}}(F') - \phi_{\mathcal{F}}(F)\}^{\top}]$ . We use  $B = BB^+B$ . Note  $\tilde{w}$  and  $w$  can be different.

## L.5 Proof of [Lemma 7](#)

Refer to [[USL<sup>+</sup>22](#), Chapter J].

## M Proof of [Section E](#)

### M.1 Proof of [Theorem 8](#)

Take  $b_V^{[t]} \in \mathcal{B}_V^{[t]} \cap \mathcal{Q}_t$ . Then, we have

$$\mathbb{E}[\mu(O, A)\{R + \gamma b_V^{[t+1]}(F')\} - b_V^{[t]}(F) \mid H] = 0.$$

Here, this implies

$$\mathbb{E}[\mathbb{E}[\mu(O, A)\{R + \gamma b_V^{[t+1]}(F')\} - b_V^{[t]}(F) \mid S] \mid H] = 0$$

using  $H \perp (A, O, R, F') \mid S$ .

Then, from the we have  $\mathcal{T}_t^{\mathcal{S}}(b_V)(S) = 0$ , i.e.,

$$\mathbb{E}[\mu(O, A)\{R + \gamma b_V^{[t+1]}(F')\} - b_V^{[t]}(F) \mid S] = 0.$$

using the assumption (b). Next, from the overlap assumption (c), we have  $\{\mathcal{T}^{\mathcal{S}, t}(b_V)\}(S_t^{\circ}) = 0$  where  $S_t^{\circ} \sim d_t^{\pi^e}(\cdot)$ .

Therefore, we have

$$\begin{aligned} & J(\pi^e) - \mathbb{E}_{f \sim \nu_{\mathcal{F}}}[b_V^{[0]}(f)] \\ &= \left( \sum_{t=0}^{T-1} \mathbb{E}_{s \sim d_t^{\pi^e}}[\gamma^t \mathbb{E}[\mu(O, A)R \mid S = s]] \right) + \left( \sum_{t=0}^{T-1} \gamma^t \mathbb{E}_{s \sim d_{t+1}^{\pi^e}}[\mathbb{E}[b_V^{[t+1]}(F) \mid S = s]] - \mathbb{E}_{s \sim d_t^{\pi^e}}[\mathbb{E}[b_V^{[t]}(F) \mid S = s]] \right) \end{aligned}$$

(Telescoping sum)

from telescoping sum. Besides, we have

$$\begin{aligned} & \mathbb{E}_{s \sim d_{t+1}^{\pi^e}}[\mathbb{E}[b_V^{[t+1]}(F) \mid S = s]] \\ &= \mathbb{E}_{s \sim d_t^{\pi^e}}[\mathbb{E}[\mu(O, A)\mathbb{E}[b_V^{[t+1]}(F) \mid S'] \mid S = s]] \\ &= \mathbb{E}_{s \sim d_t^{\pi^e}}[\mathbb{E}[\mu(O, A)\mathbb{E}[b_V^{[t+1]}(F) \mid S', O, A] \mid S = s]] \quad (F \perp (O, A) \setminus S' \mid S') \\ &= \mathbb{E}_{s \sim d_t^{\pi^e}}[\mathbb{E}[\mu(O, A)b_V^{[t+1]}(F) \mid S = s]]. \quad (\text{Total law of expectation}) \end{aligned}$$

Therefore,

$$\begin{aligned} & J(\pi^e) - \mathbb{E}_{f \sim \nu_{\mathcal{F}}}[b_V^{[0]}(f)] \\ &= \left( \sum_{t=0}^{T-1} \mathbb{E}_{s \sim d_t^{\pi^e}}[\gamma^t \mathbb{E}[\mu(O, A)R \mid S = s]] \right) + \left( \sum_{t=0}^{T-1} \gamma^t \mathbb{E}_{s \sim d_t^{\pi^e}}[\mathbb{E}[\gamma\mu(O, A)b_V^{[t+1]}(F') - b_V^{[t]}(F) \mid S = s]] \right) \\ &= \sum_{t=0}^{T-1} \gamma^t \mathbb{E}_{s \sim d_t^{\pi^e}}[\mathbb{E}[\mu(O, A)\{R + \gamma b_V^{[t+1]}(F')\} - b_V^{[t]}(F) \mid S = s]] \\ &= \sum_{t=0}^{T-1} \mathbb{E}_{s \sim d_t^{\pi^e}}[\mathcal{T}_t^{\mathcal{S}} s] \\ &= 0. \end{aligned}$$

(Recall we derive  $\{\mathcal{T}^{\mathcal{S}, t}(b_V)\}(S_t^{\circ}) = 0$ .)

## M.2 Proof of Corollary 1

The proof consists of three steps. We use [Theorem 8](#).

**First step: verify existence of learnable future-dependent value functions (Show (8a)).** From (LM1), we need to find a solution to

$$w_1^\top \mathbb{E}[\phi_{\mathcal{F}}(F) \mid S] = V_t^{\pi^e}(S)$$

with respect to a value  $w_1$ . Then, from (LM2f), there exists  $w_2$  and  $K_1$  such that

$$w_1^\top K_1 \phi_{\mathcal{S}}(S) = w_2^\top \phi_{\mathcal{S}}(S).$$

Thus, from (LM3), the above has a solution with respect to  $w_1$ .

**Second step: verify invertibility condition (Show (8b)).** Take a function  $q^{[t]}(F)$  linear in  $\phi_{\mathcal{F}}(F)$ . From (LM1), (LM2) and (LM4), there exists  $w_3 \in \mathbb{R}^{d_S}$  such that

$$\mathbb{E}[\mu(O, A)\{R + \gamma q^{[t+1]}(F') - q^{[t]}(F)\} \mid S] = w_3^\top \phi_{\mathcal{S}}(S).$$

This is proved as in (31). Then,  $w_3^\top \mathbb{E}[\phi_{\mathcal{S}}(S) \mid H] = 0$  implies  $w_3^\top \phi_{\mathcal{S}}(S) = 0$  from (LM5). Hence, when we have

$$\begin{aligned} 0 &= \mathbb{E}[\mathbb{E}[\mu(O, A)\{R + \gamma q^{[t+1]}(F') - q^{[t]}(F)\} \mid S] \mid H] \\ &= \mathbb{E}[w_3^\top \phi_{\mathcal{S}}(S) \mid H], \end{aligned}$$

this implies  $w_3^\top \phi_{\mathcal{S}}(S) = 0$ , i.e.,

$$\mathbb{E}[\mu(O, A)\{R + \gamma q^{[t+1]}(F') - q^{[t]}(F)\} \mid S] = 0.$$

**Third step: show learnable future-dependent value functions are future-dependent value functions.** Take a learnable future-dependent value function  $b_V$ . Then, from the condition, we have

$$\mathbb{E}[\mu(O, A)\{R + \gamma b_V^{[t+1]}(F')\} - b_V^{[t]}(F) \mid S] = 0.$$

We want to prove

$$\mathbb{E}[b_V^{[t]}(F) \mid S] = V_t^{\pi^e}(S).$$

We use induction. When  $t = T - 1$ , this is clear. Next, supposing the statement is true at  $t + 1$ , we prove it at a horizon  $t$ . Here, we have

$$\begin{aligned} \mathbb{E}[b_V^{[t]}(F) \mid S] &= \mathbb{E}[\mu(O, A)\{R + \gamma b_V^{[t+1]}(F')\} \mid S] \\ &= \mathbb{E}[\mu(O, A)\{R + \gamma \mathbb{E}[b_V^{[t+1]}(F') \mid S']\} \mid S] \\ &= \mathbb{E}[\mu(O, A)\{R + \gamma V_{t+1}^{\pi^e}(S')\} \mid S] = V_t^{\pi^e}(S). \end{aligned}$$

Thus, from induction, we have  $\mathbb{E}[b_V^{[t]}(F) \mid S] = V_t^{\pi^e}(S)$  for any  $t \in [T - 1]$  for any learnable future-dependent value function  $b_V(\cdot)$ .

**Fourth step: show the final formula.** Recall we define

$$\theta_t = \mathbb{E}[\phi_{\mathcal{H}}(H)\phi_{\mathcal{F}}(F)^\top] + \mathbb{E}[\mu(O, A)\phi_{\mathcal{H}}(H)\{R + \gamma \phi_{\mathcal{F}}^\top(F')\theta_{t+1}}]$$

We want to show  $\theta_t^\top \phi_{\mathcal{F}}(\cdot)$  is a learnable future-dependent value function. Here, we need to say

$$\mathbb{E}[\mu(O, A)\{R + \gamma \theta_{t+1}^\top \phi_{\mathcal{F}}(F')\} - \theta_t^\top \phi_{\mathcal{F}}(F) \mid H] = 0.$$

This is satisfied if we have

$$\mathbb{E}(\phi_{\mathcal{H}}(H) (\mu(O, A)\{R + \gamma \phi_{\mathcal{F}}^\top(F')\theta_{t+1}\} - \phi_{\mathcal{F}}^\top(F)\theta_t)) = 0. \quad (33)$$

This is because  $\mathbb{E}[\mu(O, A)\{R + \gamma \theta_{t+1}^\top \phi_{\mathcal{F}}(F')\} - \theta_t^\top \phi_{\mathcal{F}}(F) \mid H] = w_5^\top \phi_{\mathcal{H}}(H)$  for some vector  $w_5$ . Besides,  $\mathbb{E}[w_5^\top \phi_{\mathcal{H}}(H)\phi_{\mathcal{H}}(H)] = 0$  implies  $\mathbb{E}[w_5^\top \phi_{\mathcal{H}}(H)\phi_{\mathcal{H}}(H)w_5] = 0$ , which results in

$$w_5^\top \phi_{\mathcal{H}}(H) = 0.$$

On top of that, since future-dependent value functions  $(\tilde{\theta}_t, \phi_{\mathcal{F}}(F))$  exist from the first statement, we have a solution:

$$\mathbb{E} \left( \phi_{\mathcal{H}}(H) \left( \mu(O, A) \{R + \gamma \phi_{\mathcal{F}}^{\top}(F') \tilde{\theta}_{t+1}\} - \phi_{\mathcal{F}}^{\top}(F) \tilde{\theta}_t \right) \right) = 0.$$

with respect to  $\tilde{\theta}_t$ . In the following, we use this fact.

Now, we go back to the proof of the main statement, i.e., we prove (33). We use the induction. This is immediately proved when  $t = T - 1$ . Here, suppose  $\phi_{\mathcal{F}}^{\top}(\cdot) \theta_{t+1}$  is a learnable future-dependent value function at  $t + 1$ . Then, we have

$$\begin{aligned} & \mathbb{E}[\phi_{\mathcal{H}}(H) \phi_{\mathcal{F}}(F)^{\top}] \theta_t \\ &= \mathbb{E}[\phi_{\mathcal{H}}(H) \phi_{\mathcal{F}}(F)^{\top}] \mathbb{E}[\phi_{\mathcal{H}}(H) \phi_{\mathcal{F}}(F)^{\top}]^+ \mathbb{E}[\mu(O, A) \phi_{\mathcal{H}}(H) \{R + \gamma \phi_{\mathcal{F}}^{\top}(F') \theta_{t+1}\}] \quad (\text{Definition}) \\ &= \mathbb{E}[\phi_{\mathcal{H}}(H) \phi_{\mathcal{F}}(F)^{\top}] \mathbb{E}[\phi_{\mathcal{H}}(H) \phi_{\mathcal{F}}(F)^{\top}]^+ \mathbb{E}[\mu(O, A) \phi_{\mathcal{H}}(H) \{R + \gamma V_{t+1}^{\pi^e}(S')\}]. \end{aligned}$$

In the last line, we use the inductive hypothesis and the third step. Recall we showed in the previous step  $\mathbb{E}[b_V^{[t]}(F) \mid S] = V_t^{\pi^e}(S)$  for any learnable future-dependent value functions  $b_V^{[t]}(\cdot)$ . Then, we have

$$\begin{aligned} & \mathbb{E}[\phi_{\mathcal{H}}(H) \phi_{\mathcal{F}}(F)^{\top}] \theta_t \\ &= \mathbb{E}[\phi_{\mathcal{H}}(H) \phi_{\mathcal{F}}(F)^{\top}] \mathbb{E}[\phi_{\mathcal{H}}(H) \phi_{\mathcal{F}}(F)^{\top}]^+ \mathbb{E}[\mu(O, A) \phi_{\mathcal{H}}(H) \{R + \gamma V_{t+1}^{\pi^e}(S')\}] \\ &= \mathbb{E}[\phi_{\mathcal{H}}(H) \phi_{\mathcal{F}}(F)^{\top}] \mathbb{E}[\phi_{\mathcal{H}}(H) \phi_{\mathcal{F}}(F)^{\top}]^+ \mathbb{E}[\mu(O, A) \phi_{\mathcal{H}}(H) \{R + \gamma \phi_{\mathcal{F}}^{\top}(F') \tilde{\theta}_{t+1}\}] \\ &= \mathbb{E}[\phi_{\mathcal{H}}(H) \phi_{\mathcal{F}}(F)^{\top}] \mathbb{E}[\phi_{\mathcal{H}}(H) \phi_{\mathcal{F}}(F)^{\top}]^+ \mathbb{E}[\phi_{\mathcal{H}}(H) \phi_{\mathcal{F}}(F)^{\top}] \tilde{\theta}_t \\ &= \mathbb{E}[\phi_{\mathcal{H}}(H) \phi_{\mathcal{F}}(F)^{\top}] \tilde{\theta}_t \quad (\text{Property of Moore-Penrose Inverse}) \\ &= \mathbb{E}[\mu(O, A) \phi_{\mathcal{H}}(H) \{R + \gamma \phi_{\mathcal{F}}^{\top}(F') \tilde{\theta}_{t+1}\}] \quad (\text{Definition of } \tilde{\theta}_t) \\ &= \mathbb{E}[\mu(O, A) \phi_{\mathcal{H}}(H) \{R + \gamma V_{t+1}^{\pi^e}(S')\}] \quad (\text{We showed in the previous step}) \\ &= \mathbb{E}[\mu(O, A) \phi_{\mathcal{H}}(H) \{R + \gamma \phi_{\mathcal{F}}^{\top}(F') \theta_{t+1}\}]. \quad (\text{From the induction}) \end{aligned}$$

Hence,  $\phi_{\mathcal{F}}^{\top}(\cdot) \theta_t$  is a learnable future-dependent value function at  $t$ .

## N Proof of Section F

### N.1 Proof of Theorem 9

We take a value bridge function  $b_D^{[t]} \in \mathcal{Q}_t \cap \mathcal{B}_D^{[t]}$  for any  $[T]$ . Then, we define

$$l_D^{[t]}(\cdot) = \mathbb{E}[b_D^{[t]}(F) \mid S = \cdot].$$

In this section,  $\mathbb{E}[\cdot; A_{0:T-1} \sim \pi^e]$  means taking expectation when we execute a policy  $\pi_e$  from  $t = 0$  to  $T - 1$ . Note  $\mathbb{E}[\cdot; A_{0:T-1} \sim \pi^b]$  is just  $\mathbb{E}[\cdot]$ . Here, we have

$$\begin{aligned} & \Pr_{\pi^e}(o_0, a_0, \dots, a_{T-1}) - \mathbb{E}_{f \sim \nu_{\mathcal{F}}}[b_D^{[0]}(f)] \\ &= \sum_{t=0}^{T-1} \mathbb{E} \left[ \left\{ \prod_{k=0}^{t-1} \mathbb{I}(O_k = o_k, A_k = a_k) \right\} \left\{ \mathbb{I}(O_t = o_t, A_t = a_t) l_D^{[t+1]}(S_{t+1}) - l_D^{[t]}(S_t) \right\}; A_{0:T-1} \sim \pi^e \right] \\ &= \sum_{t=0}^{T-1} \mathbb{E} \left[ \left\{ \prod_{k=0}^{t-1} \mathbb{I}(O_k = o_k, A_k = a_k) \right\} \left\{ \mathbb{E}[\mathbb{I}(O_t = o_t, A_t = a_t) l_D^{[t+1]}(S_{t+1}) \mid S_t, (O_0, A_0, \dots, A_{t-1}) = (o_0, a_0, \dots, a_{t-1})] \right. \right. \\ & \quad \left. \left. - l_D^{[t]}(S_t) \right\}; A_{0:T-1} \sim \pi^e \right] \\ &= \sum_{t=0}^{T-1} \mathbb{E} \left[ \left\{ \prod_{k=0}^{t-1} \mathbb{I}(O_k = o_k, A_k = a_k) \right\} \left\{ \mathbb{E}[\mathbb{I}(O_t = o_t, A_t = a_t) l_D^{[t+1]}(S_{t+1}) \mid S_t] - l_D^{[t]}(S_t) \right\}; A_{0:T-1} \sim \pi^e \right]. \end{aligned}$$

In the first line, we use a telescoping sum trick noting

$$\Pr_{\pi^e}(o_0, a_0, \dots, a_{T-1}) = \mathbb{E} \left[ \prod_{k=0}^{T-1} \mathbb{I}(O_k = o_k, A_k = a_k); A_{0:T-1} \sim \pi^e \right].$$

From the second line to the third line, we use  $S_{t+1}, O_t, A_t \perp O_1, \dots, A_t \mid S_t$ .

Here, we have

$$\mathbb{E}[\mathbb{I}(O_t = o_t, A_t = a_t) \tilde{l}_D^{[t+1]}(S_{t+1}) \mid S_t; A_t \sim \pi^e] = \tilde{l}_D^{[t]}(S_t)$$

where  $\tilde{l}_D^{[t+1]}(\cdot) = \mathbb{E}[\mathbb{I}(O = o_t, A = a_t) \mu(O, A) \tilde{l}_D^{[t+1]}(S') \mid S = \cdot; A_t \sim \pi^b]$  using importance sampling. Hence, the following holds:

$$\begin{aligned} & \Pr_{\pi^e}(o_0, a_0, \dots, a_{T-1}) - \mathbb{E}_{f \sim \nu_{\mathcal{F}}}[b_D^{[0]}(f)] \\ &= \sum_{t=0}^{T-1} \mathbb{E} \left[ \left\{ \prod_{k=0}^{t-1} \mathbb{I}(O_k = o_k, A_k = a_k) \right\} \left\{ \tilde{l}_D^{[t+1]}(S_t) - l_D^{[t]}(S_t) \right\}; A_{0:T-1} \sim \pi^e \right] \\ &= \sum_{t=0}^{T-1} \mathbb{E} \left[ \left\{ \prod_{k=0}^{t-1} \mathbb{I}(O_k = o_k, A_k = a_k) (\mathcal{T}_t^S b_D)(S_t) \right\}; A_{0:T-1} \sim \pi^e \right] \tag{34} \\ &= \sum_{t=0}^{T-1} \mathbb{E} \left[ \left\{ \prod_{k=0}^{t-1} \mathbb{I}(O_k = o_k, A_k = a_k) \mathbb{E}[(\mathcal{T}_t^S b_D)(S_t) \mid (O_1, A_1, \dots, A_K) = (o_1, a_1, \dots, o_K)] \right\}; A_{0:T-1} \sim \pi^e \right] \tag{35} \end{aligned}$$

From the second line to the third line, we use

$$\begin{aligned} & \mathbb{E}[\mathbb{I}(O = o_t, A = a_t) \mu(O, A) \tilde{l}_D^{[t+1]}(S') \mid S] \\ &= \mathbb{E}[\mathbb{I}(O = o_t, A = a_t) \mu(O, A) \mathbb{E}[b_D^{[t+1]}(F') \mid S'] \mid S] \tag{Definition} \\ &= \mathbb{E}[\mathbb{I}(O = o_t, A = a_t) \mu(O, A) \mathbb{E}[b_D^{[t+1]}(F') \mid S', O = o_t, A = a_t] \mid S] \\ & \tag{Law of total expectation} \\ &= \mathbb{E}[\mathbb{I}(O = o_t, A = a_t) \mu(O, A) b_D^{[t+1]}(F') \mid S]. \tag{F' \perp O, A \mid S'} \end{aligned}$$

Besides, we know for any learnable bridge function  $b_D$ , we have

$$\mathbb{E}[\mathbb{I}(O = o_t, A = a_t) \mu(O, A) b_D^{[t+1]}(F') - b_D^{[t]}(F) \mid H] = 0.$$

Then, since

$$\mathbb{E}[\mathbb{E}[\mathbb{I}(O = o_t, A = a_t) \mu(O, A) b_D^{[t+1]}(F') - b_D^{[t]}(F) \mid S] \mid H] = 0$$

from the invertibility condition (b), we have

$$\mathbb{E}[\mathbb{I}(O = o_t, A = a_t) \mu(O, A) b_D^{[t+1]}(F') - b_D^{[t]}(F) \mid S] = 0.$$

From the overlap (c), this implies

$$(\mathcal{T}_t^S b_D)(\dot{S}_t) = 0.$$

Finally, from (35), we can conclude  $\Pr_{\pi^e}(o_0, a_0, \dots, a_{T-1}) - \mathbb{E}_{f \sim \nu_{\mathcal{F}}}[b_D^{[0]}(f)] = 0$ .

## N.2 Proof of Theorem 10

This is proved as in Theorem 9 noting

$$\Pr_{\pi^e}(o_0, a_0, \dots, a_{T-1}, \mathbf{O}_T) = \mathbb{E} \left[ \phi_{\mathcal{O}}(\mathbf{O}_T) \prod_{k=0}^{T-1} \mathbb{I}(O_k^\diamond = o_k, A_k^\diamond = a_k) \right].$$

## N.3 Proof of Corollary 2

This is proved following Corollary 3.

## N.4 Proof of Corollary 3

The proof consists of four steps. The proof largely follows the proof of Corollary 1.

**First step: verify the existence of bridge functions (9a).** We show there exists a bridge function linear in  $\phi_{\mathcal{F}}(F)$ . We need to find a solution to  $w_1^\top \mathbb{E}[\phi_{\mathcal{F}}(F) \mid S] = V_{D,[t]}^{\pi^e}(S)$  with respect to  $w_1$ . Then, from (LM2D), there exists  $w_2$  and  $K_1$  such that

$$w_1^\top K_1 \phi_S(S) = w_2^\top \phi_S(S).$$

Thus, from (LM3), the above has a solution with respect to  $w_1$ .

**Second step: verify invertibility conditions (9b).** We take  $\mathcal{Q}_t$  to be a linear model in  $\phi_{\mathcal{F}}(F)$ . For  $q_D^{[t]} \in \mathcal{Q}_t$ , from (LM1), (LM2), (LM4), there exists  $w_3 \in \mathbb{R}^{d_S}$  such that

$$\mathbb{E}[q_D^{[t]}(F) - \mathbb{I}(O = o_t, A = a_t)\mu(O, A)q_D^{[t+1]}(F') \mid S] = w_3^\top \phi_S(S).$$

Then,  $w_3^\top \mathbb{E}[\phi_S(S) \mid H] = 0$  implies  $w_3^\top \phi_S(S) = 0$  from (LM5). Hence,

$$\begin{aligned} 0 &= \mathbb{E}[q_D^{[t]}(F) - \mathbb{I}(O = o_t, A = a_t)\mu(O, A)q_D^{[t+1]}(F') \mid H] \\ &= \mathbb{E}[\mathbb{E}[q_D^{[t]}(F) - \mathbb{I}(O = o_t, A = a_t)\mu(O, A)q_D^{[t+1]}(F') \mid S] \mid H] \end{aligned}$$

implies  $\mathbb{E}[q_D^{[t]}(F) - \mathbb{I}(O = o_t, A = a_t)\mu(O, A)q_D^{[t+1]}(F') \mid S] = 0$ . Thus, the invertibility is concluded.

**Third step: show learnable value bridge functions are value bridge functions.** From the previous discussion, learnable value bridge functions in  $\mathcal{Q}$  need to satisfy

$$0 = \mathbb{E}[b_D^{[t]}(F) - \mathbb{I}(O = o_t, A = a_t)\mu(O, A)b_D^{[t+1]}(F') \mid S].$$

Here, we want to prove

$$\mathbb{E}[b_D^{[t]}(F) \mid S] = V_{D,[t]}^{\pi^e}(S).$$

We use induction. When  $t = T - 1$ , this is clear. Next, supposing the statement is true at  $t + 1$ , we prove it at horizon  $t$ . Here, we have

$$\begin{aligned} \mathbb{E}[b_D^{[t]}(F) \mid S] &= \mathbb{E}[\mathbb{I}(O = o_t, A = a_t)\mu(O, A)b_D^{[t+1]}(F') \mid S] \\ &= \mathbb{E}[\mathbb{I}(O = o_t, A = a_t)\mu(O, A)\mathbb{E}[b_D^{[t+1]}(F') \mid S'] \mid S] \\ &= \mathbb{E}[\mathbb{I}(O = o_t, A = a_t)\mu(O, A)V_{D,[t+1]}^{\pi^e}(S') \mid S] \\ &= V_{D,[t]}^{\pi^e}(S). \end{aligned}$$

**Fourth step: show the final formula.** We recursively define

$$\theta_t^\top = \theta_{t+1}^\top D_t B^+.$$

starting from  $\tilde{w}_T^\top B = \mathbb{E}[\phi_{\mathcal{H}}(H)]$ . We want to show  $\theta_t^\top \phi_{\mathcal{F}}(\cdot)$  is a learnable value bridge function. Here, we want to say

$$\mathbb{E}[\theta_t^\top \phi_{\mathcal{F}}(F) - \mathbb{I}(O = o_t, A = a_t)\mu(O, A)\theta_{t+1}^\top \phi_{\mathcal{F}}(F') \mid H] = 0.$$

This is satisfied if we have

$$\mathbb{E}[(\theta_t^\top \phi_{\mathcal{F}}(F) - \mathbb{I}(O = o_t, A = a_t)\mu(O, A)\theta_{t+1}^\top \phi_{\mathcal{F}}(F')) \phi_{\mathcal{H}}(H)] = 0$$

Here, refer to the fourth step in the proof of Corollary 1. Besides, as we already show linear value bridge functions exist  $(\tilde{\theta}_t, \phi_{\mathcal{F}}(\cdot))$ , we have a solution to:

$$\mathbb{E}\left[(\tilde{\theta}_t^\top \phi_{\mathcal{F}}(F) - \mathbb{I}(O = o_t, A = a_t)\mu(O, A)\tilde{\theta}_{t+1}^\top \phi_{\mathcal{F}}(F')) \phi_{\mathcal{H}}(H)\right] = 0$$

with respect to  $\tilde{\theta}_t$ .

Hereafter, we use the induction. This is immediately proved when  $t = T - 1$ . Here, suppose  $\phi_{\mathcal{F}}^\top(F)\theta_{t+1}$  is a learnable value bridge function at  $t + 1$ . Then, we have

$$\begin{aligned} &\mathbb{E}[\phi_{\mathcal{H}}(H)\phi_{\mathcal{F}}(F)^\top]\theta_t \\ &= \mathbb{E}[\phi_{\mathcal{H}}(H)\phi_{\mathcal{F}}(F)^\top]\mathbb{E}[\phi_{\mathcal{H}}(H)\phi_{\mathcal{F}}(F)^\top]^+ D_t^\top \theta_{t+1} \\ &= \mathbb{E}[\phi_{\mathcal{H}}(H)\phi_{\mathcal{F}}(F)^\top]\mathbb{E}[\phi_{\mathcal{H}}(H)\phi_{\mathcal{F}}(F)^\top]^+ \mathbb{E}[\mathbb{I}(O = O_t, A = a_t)\mu(O, A)\phi_{\mathcal{H}}(H)\phi_{\mathcal{F}}^\top(F')\theta_{t+1}] \\ &= \mathbb{E}[\phi_{\mathcal{H}}(H)\phi_{\mathcal{F}}(F)^\top]\mathbb{E}[\phi_{\mathcal{H}}(H)\phi_{\mathcal{F}}(F)^\top]^+ \mathbb{E}[\mathbb{I}(O = O_t, A = a_t)\mu(O, A)\phi_{\mathcal{H}}(H)V_{D,[t+1]}^{\pi^e}(S')]. \end{aligned}$$

In the last line, we use inductive hypothesis. We showed in the previous step  $\mathbb{E}[b_D^{[t]}(F) \mid S] = V_{D, [t]}^{\pi^e}(S)$  for any learnable value bridge function  $b_D^{[t]}(F)$ . Then, we have

$$\begin{aligned}
& \mathbb{E}[\phi_{\mathcal{H}}(H)\phi_{\mathcal{F}}(F)^\top] \mathbb{E}[\phi_{\mathcal{H}}(H)\phi_{\mathcal{F}}(F)^\top] + \mathbb{E}[\mathbb{I}(O = O_t, A = a_t)\mu(O, A)\phi_{\mathcal{H}}(H)V_{D, [t+1]}^{\pi^e}(S')] \\
&= \mathbb{E}[\phi_{\mathcal{H}}(H)\phi_{\mathcal{F}}(F)^\top] \mathbb{E}[\phi_{\mathcal{H}}(H)\phi_{\mathcal{F}}(F)^\top] + \mathbb{E}[\mathbb{I}(O = O_t, A = a_t)\mu(O, A)\phi_{\mathcal{H}}(H)\{\phi_{\mathcal{F}}^\top(F')\tilde{\theta}_{t+1}\}] \\
&= \mathbb{E}[\phi_{\mathcal{H}}(H)\phi_{\mathcal{F}}(F)^\top] \mathbb{E}[\phi_{\mathcal{H}}(H)\phi_{\mathcal{F}}(F)^\top] + \mathbb{E}[\phi_{\mathcal{H}}(H)\phi_{\mathcal{F}}(F)^\top]\tilde{\theta}_t \\
&= \mathbb{E}[\phi_{\mathcal{H}}(H)\phi_{\mathcal{F}}(F)^\top]\tilde{\theta}_t && \text{(Property of Moore-Penrose Inverse)} \\
&= \mathbb{E}[\mathbb{I}(O = O_t, A = a_t)\mu(O, A)\phi_{\mathcal{H}}(H)V_{D, [t+1]}^{\pi^e}(S')] && \text{(We show in the previous step)} \\
&= \mathbb{E}[\mathbb{I}(O = O_t, A = a_t)\mu(O, A)\phi_{\mathcal{H}}(H)\{\phi_{\mathcal{F}}^\top(F')\theta_{t+1}\}]. && \text{(From the induction)}
\end{aligned}$$

So far, we prove  $\theta_t^\top \phi_{\mathcal{F}}(\cdot)$  is a learnable value bridge function for  $t$ . Finally, by considering a time-step  $t = 0$ , we can prove the target estimand is

$$\mathbb{E}[\tilde{w}_0^\top \phi_{\mathcal{F}}(F)] = \mathbb{E}[\phi_{\mathcal{H}}(H)]^\top B^+ \left\{ \prod_{t=T-1}^0 D_t B^+ \right\} C.$$

## O Auxiliary Lemmas

**Lemma 9.** Take  $g \in [S \rightarrow \mathbb{R}]$ . Then, we have

$$0 = (1 - \gamma)^{-1} \mathbb{E}_{s \sim d_{\pi^e}} [\mathbb{E}[\mu(O, A)\gamma g(S') - g(S) \mid S = s]] + \mathbb{E}_{s \sim \nu_S} [gs].$$

*Proof.* Let  $d_{\pi^e}(z, s) = d_0^{\pi^e}(\cdot)$ . Then, we have

$$\begin{aligned}
& \int g(z, s) d_{\pi^e}(z, s) d(z, s) \\
&= (1 - \gamma) \int g(z, s) \sum_{t=0}^{\infty} \gamma^t d_t^{\pi^e}(z, s) d(z, s) \\
&= \underbrace{(1 - \gamma) \int g(z, s) d_0^{\pi^e}(z, s) d(z, s)}_{(a)} + \underbrace{\gamma(1 - \gamma) \int g(z', s') \sum_{t=1}^{\infty} \gamma^{t-1} d_t^{\pi^e}(z', s') d(z', s')}_{(b)}.
\end{aligned}$$

We analyze the first term (a) and the second term (b). The first term (a) is  $\mathbb{E}_{s \sim \nu_S} [gs]$ . In the following, we analyze the second term.

Here, note

$$d_t^{\pi^e}(z', s') = \int \mathbb{T}(s' \mid s, a) \mathbb{O}(o \mid s) \pi^e(a \mid z, o) \delta^\dagger(z' = z) d_{t-1}^{\pi^e}(z, s) d(z, s).$$

where  $\delta^\dagger(z' = z) = \delta(f_{-1}(z') = f_{+1}(z))$ . Here,  $f_{-1}$  is a transformation removing the most recent tuple  $(o, a, r)$  and  $f_{+1}$  is a transformation removing the oldest tuple  $(o, a, r)$ . Hence,

$$\begin{aligned}
& (1 - \gamma) \sum_{t=1}^{\infty} \gamma^{t-1} d_t^{\pi^e}(z', s') \\
&= (1 - \gamma) \int \sum_{t=1}^{\infty} \gamma^{t-1} \mathbb{T}(s' \mid s, a) \mathbb{O}(o \mid s) \pi^e(a \mid z, o) \delta^\dagger(z' = z) d_{t-1}^{\pi^e}(z, s) d(z, s) \\
&= (1 - \gamma) \int \sum_{k=0}^{\infty} \gamma^k \mathbb{T}(s' \mid s, a) \mathbb{O}(o \mid s) \mu(z, o, a) \pi^b(a \mid z, o) \delta^\dagger(z' = z) d_k^{\pi^e}(z, s) d(z, s) \\
&= \int \mathbb{T}(s' \mid s, a) \mathbb{O}(o \mid s) \mu(z, o, a) \pi^b(a \mid z, o) \delta^\dagger(z' = z) d^{\pi^e}(z, s) d(z, s).
\end{aligned}$$

Therefore, the term (b) is

$$\begin{aligned}
& \gamma(1 - \gamma) \int g(z', s') \sum_{t=1}^{\infty} \gamma^{t-1} d_t^{\pi^e}(z', s') d(z', s') \\
&= \gamma \int g(z', s') \mathbb{T}(s' | s, a) \mathbb{O}(o | s) \mu(z, o, a) \pi^b(a | z, o) \delta^\dagger(z' = z) d^{\pi^e}(z, s) d(z', s, s') \\
&= \gamma \mathbb{E}_{s \sim d_{\pi^e}} [\mathbb{E}[\mu(O, A) \gamma g(S') | S = s]].
\end{aligned}$$

□