

---

# On the Role of Discount Factor in Offline Reinforcement Learning

---

Hao Hu<sup>\*1</sup> Yiqing Yang<sup>\*2</sup> Qianchuan Zhao<sup>2</sup> Chongjie Zhang<sup>1</sup>

## Abstract

Offline reinforcement learning (RL) enables effective learning from previously collected data without exploration, which shows great promise in real-world applications when exploration is expensive or even infeasible. The discount factor,  $\gamma$ , plays a vital role in improving online RL sample efficiency and estimation accuracy, but the role of the discount factor in offline RL is not well explored. This paper examines two distinct effects of  $\gamma$  in offline RL with theoretical analysis, namely the regularization effect and the pessimism effect. On the one hand,  $\gamma$  is a regulator to trade-off optimality with sample efficiency upon existing offline techniques. On the other hand, lower guidance  $\gamma$  can also be seen as a way of pessimism where we optimize the policy’s performance in the worst possible models. We empirically verify the above theoretical observation with tabular MDPs and standard D4RL tasks. The results show that the discount factor plays an essential role in the performance of offline RL algorithms, both under small data regimes upon existing offline methods and in large data regimes without other conservative methods.

## 1. Introduction

Online reinforcement learning has achieved great success in various domains, including video games (Mnih et al., 2015), class games (Silver et al., 2016) and continuous control (Lillicrap et al., 2015). However, it requires extensive interaction with the environment to learn through trial-and-error. In many real-world problems, like personalized recommendation systems (Swaminathan & Joachims, 2015) and autonomous driving (Shalev-Shwartz et al., 2016;

---

<sup>\*</sup>Equal contribution <sup>1</sup>Institute of Interdisciplinary Information Sciences, Tsinghua University, Beijing, China <sup>2</sup>Department of Automation, Tsinghua University, Beijing, China. Correspondence to: Hao Hu <hu-h19@mails.tsinghua.edu.cn>, Yiqing Yang <yangyiqi19@mails.tsinghua.edu.cn>.

Dataset size/quality	w other pessimisms	w\o other pessimisms
Large, good coverage		pessimism effect ✓
Small or bad coverage	regularization effect ✓	

Table 1. The applicability of a lower guidance discount factor.

Singh et al., 2020), exploration can be expensive and unsafe, which requires conducting RL in a supervised manner. In contrast, offline RL (Levine et al., 2020; Fujimoto et al., 2019) enables learning from previously collected datasets, which shows great potential for real-world applications.

One of the major challenges of offline RL comes from the distributional shift between the data collection policy and the learned policy (Levine et al., 2020), and the direct application of online RL algorithms is known to lead to poor performance (Fujimoto et al., 2019). One paradigm for algorithm design in offline RL incorporates proper conservatism to the learning algorithm. There are several conservative methods in existing empirical literature, including policy regularization (Fujimoto et al., 2019; Peng et al., 2019; Kumar et al., 2020), ensemble-based uncertainty (Wu et al., 2021; An et al., 2021) and model-based penalty (Yu et al., 2021; 2020; Kidambi et al., 2020).

However, as an essential number in RL, the discount factor already provides a natural way of conservatism. The effect of  $\gamma$  is extensively discussed in online RL and specifically TD-learning. For instance, Jiang et al. (2015) shows that the discount factor can be regarded as a complexity control parameter for the class of policies. Amit et al. (2020) shows that it is equivalent to L2-regularization of Q-values in TD-learning to improve generalization. However, the analysis of the role of  $\gamma$  in the context of offline RL is missing, which naturally leads to the following question:

*What are the roles of the discount factor in the context of offline RL, and how does it contribute to proper conservative offline RL methods?*

In this paper, we examine the two roles of the discount factor in offline RL that affect the performance in two distinct ways. First, we show that  $\gamma$  plays as a regulator upon exist-

ing conservative methods and achieves a trade-off between optimality and sample efficiency. Second, we show that a lower discount factor resembles model-based pessimism. A lower discount factor is equivalent to maximizing the value function in the worst possible models among a confidence set. We summarize the applicability of the two effects in Table 1. To give a rigid characterization, we analyze the two effects above under the context of linear MDPs and derive two different performance bounds. These quantitative results also characterize how the impact of a lower guidance discount factor relies on other factors like the size of the dataset and the coverage coefficient in terms of optimal policies. We empirically verify the two effects on both tabular MDPs and the standard D4RL benchmark (Fu et al., 2020). Results indicate that the discount factor can significantly affect the performance of offline RL algorithms, both under small data regimes upon existing pessimism and in large data regimes without other pessimisms. We believe that our findings highlight the importance of setting a proper lower discount factor in offline RL practices.

### 1.1. Related Works

**Role of Discount Factor.** The discount factor is extensively analyzed in online RL. Petrik & Scherrer (2008) first shows that the approximation error bounds may be tightened with a lower discount factor in online RL. Jiang et al. (2015) gives a more rigorous analysis by analyzing the size of potentially optimal policy class; Amit et al. (2020) points out the equivalence between the L2 regularization and a smaller discount factor in TD-learning. Zhang et al. (2020) analyzes the discount factor in actor-critic algorithms from a bias-variance-representation trade-off perspective. In the practical aspect, Chen et al. (2018) observes a similar regularization effect in partial observable settings; Fedus et al. (2019) uses a geometry discount factor to learn multiple values in an ensemble way. Some works (Xu et al., 2018; Sherstan et al., 2020; Romoff et al., 2019) suggest learning a sequence of value functions with lower discount factors for online RL.

**Convervatism in Offline RL.** There are extensive works on the conservative offline RL, which can be roughly divided into policy constraint-based and uncertainty estimation-based. Policy constraint methods attempt to enforce the trained policy to be close to the behavior policy via introducing a generative model (Fujimoto et al., 2019; Fujimoto & Gu, 2021), KL-divergence (Peng et al., 2019; Nair et al., 2020; Siegel et al., 2020; Wu et al., 2019) or value function regularization (Kumar et al., 2020; Agarwal et al., 2020; Kostrikov et al., 2021). Some more recent policy constraint methods (Yang et al., 2021; Ma et al., 2021; Ghasemipour et al., 2021; Kostrikov et al., 2021; Brandfonbrener et al., 2021) also suggested that only trusting the state-action pairs

given in the dataset can solve complex tasks well. As for the uncertainty estimation of model-free methods, researchers attempt to take into account the confidence of the Q-value prediction using dropout or ensemble techniques (Wu et al., 2021; An et al., 2021). Differently, model-based offline methods incorporate the uncertainty in the model space (Yu et al., 2020; 2021; Kidambi et al., 2020).

In addition, there are also some theoretical results for pessimism in offline RL. Jin et al. (2021) proves that using a negative bonus from the online exploration is sufficient for offline RL. Rashidinejad et al. (2021) proves that a UCB-like penalty is provably sufficient in tabular settings. Uehara & Sun (2021) studies the properties of model-based pessimisms under partial coverage.

## 2. Preliminaries

We consider infinite-horizon discounted Markov Decision Processes (MDPs), defined by the tuple  $(\mathcal{S}, \mathcal{A}, \mathcal{P}, r, \gamma)$ , where  $\mathcal{S}$  is a state space,  $\mathcal{A}$  is a action space,  $\gamma \in [0, 1)$  is the discount factor and  $\mathcal{P} : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$ ,  $r : \mathcal{S} \times \mathcal{A} \rightarrow [0, r_{\max}]$  are the transition function and reward function, respectively. We also assume a fixed distribution  $\mu_0 \in \Delta(\mathcal{S})$  as the initial state distribution.

To make our analysis more general, we consider the *linear MDP* (Yang & Wang, 2019; Jin et al., 2020) as follows, where the transition kernel and expected reward function are linear with respect to a feature map. Note that tabular MDPs are linear MDPs with the canonical one-hot representation.

**Definition 2.1** (Linear MDP). We say an infinite-horizon discounted MDP  $(\mathcal{S}, \mathcal{A}, \mathcal{P}, r, \gamma)$  is a linear MDP with known feature map  $\phi : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^d$ ,  $\psi : \mathcal{S} \rightarrow \mathbb{R}^l$  if there exist an unknown matrix parameter  $M \in \mathbb{R}^{d \times l}$  and an unknown vector parameter  $\theta \in \mathbb{R}^d$  such that

$$\begin{aligned} \mathcal{P}(s' | s, a) &= \phi(s, a)^\top M \psi(s'), \\ \mathbb{E}[r(s, a)] &= \phi(s, a)^\top \theta \end{aligned} \quad (1)$$

for all  $(s, a, s') \in \mathcal{S} \times \mathcal{A} \times \mathcal{S}$ . And we assume  $\|\phi(s, a)\|_\infty \leq 1$ ,  $\|\psi(s')\|_\infty \leq 1$  for all  $(s, a, s') \in \mathcal{S} \times \mathcal{A} \times \mathcal{S}$  and  $\max\{\|M\|_2, \|\theta\|_2\} \leq \sqrt{d}$ .

A policy  $\pi : \mathcal{S} \rightarrow \Delta(\mathcal{A})$  specifies a decision-making strategy in which the agent chooses its actions based on the current state, i.e.,  $a_t \sim \pi(\cdot | s_t)$ . The value function  $V_{\gamma, M}^\pi : \mathcal{S} \rightarrow \mathbb{R}$  is defined as the  $\gamma$ -discounted sum of future rewards starting at state  $s$  for policy  $\pi$  in model  $M$ , i.e.

$$V_{\gamma, M}^\pi(s) = \mathbb{E}_\pi \left[ \sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) \mid s_0 = s \right], \quad (2)$$

where the expectation is with respect to the trajectory  $\tau$  induced by policy  $\pi$ . The action-value function  $Q^\pi : \mathcal{S} \times$

$\mathcal{A} \rightarrow \mathbb{R}$  is similarly defined as

$$Q_{\gamma, M}^{\pi}(s, a) = \mathbb{E}_{\pi} \left[ \sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) \mid s_0 = s, a_0 = a \right]. \quad (3)$$

We overload the notation and define  $V_{\gamma, M}(\pi)$  as the expected  $\gamma$ -discounted value of policy  $\pi$  under the initial distribution  $\mu_0$ , i.e.  $V_{\gamma, M}(\pi) = \mathbb{E}_{s_0 \sim \mu_0} [V_{\gamma, M}^{\pi}(s_0)]$ , and similarly we have  $Q_{\gamma, M}(\pi) = \mathbb{E}_{s_0 \sim \mu_0} [Q_{\gamma, M}^{\pi}(s_0, a_0)]$ . When it does not lead to confusion, we omit the index for  $\gamma$  and  $M$  for simplicity.

We define the Bellman operator as

$$(\mathbb{B}_{\gamma} f)(s, a) = \mathbb{E} [r(s, a) + \gamma f(s')], \quad (4)$$

for any  $f : \mathcal{S} \rightarrow \mathbb{R}$  and  $\gamma \in [0, 1)$ . The optimal Q-function  $Q^*$ , and the optimal value function  $V^*$  are related by the Bellman optimality equation

$$\begin{aligned} V_{\gamma, M}^*(s) &= \max_{a \in \mathcal{A}} Q_{\gamma, M}^*(s, a), \\ Q_{\gamma, M}^*(s, a) &= (\mathbb{B}_{\gamma} V_{\gamma, M}^*)(s, a), \end{aligned} \quad (5)$$

while the optimal policy is defined as

$$\pi_{\gamma, M}^*(\cdot | s) = \operatorname{argmax}_{\pi} \mathbb{E}_{a \sim \pi} Q_{\gamma, M}^*(s, a).$$

We define the suboptimality as the performance difference of the optimal policy  $\pi_{\gamma}^*$  and the policy  $\pi$  given the initial distribution  $\mu_0$  evaluated with discount factor  $\gamma$ . That is

$$\text{SubOpt}(\pi; \gamma) = V_{\gamma, M}(\pi_{\gamma}^*) - V_{\gamma, M}(\pi), \quad (6)$$

We also define the suboptimality for each state, that is

$$\text{SubOpt}(\pi, s; \gamma) = V_{\gamma, M}^{\pi_{\gamma}^*}(s) - V_{\gamma, M}^{\pi}(s).$$

## 2.1. Pessimistic Offline Algorithms

In this section, we sketch two offline algorithms to characterize the effect of a lower guidance discount factor. The first one is *pessimistic value iteration* (PEVI; Jin et al., 2021), as shown in Algorithm 1, which uses uncertainty as a negative bonus for value learning. The second one is *model-based pessimistic policy optimization* (MBPPO; Uehara & Sun, 2021), as shown in Algorithm 2, which optimizes the worst possible performance of a policy over a set of models.

PEVI uses negative bonus  $\Gamma(\cdot, \cdot)$  over standard Q-value estimation  $\widehat{Q}(\cdot, \cdot) = (\widehat{\mathbb{B}}_{\gamma} \widehat{V})(\cdot)$  to reduce potential bias due to finite data, where  $\widehat{\mathbb{B}}_{\gamma}$  is the empirical estimation of  $\mathbb{B}_{\gamma}$  from dataset  $\mathcal{D}$ . We use the notion of  $\xi$ -uncertainty quantifier as follows to formalize the idea of pessimism.

**Definition 2.2** ( $\xi$ -Uncertainty Quantifier). We say  $\Gamma : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$  is a  $\xi$ -uncertainty quantifier for  $\widehat{\mathbb{B}}_{\gamma}$  and  $\widehat{V}$  if with probability  $1 - \xi$ ,

$$|(\widehat{\mathbb{B}}_{\gamma} \widehat{V})(s, a) - (\mathbb{B}_{\gamma} \widehat{V})(s, a)| \leq \Gamma(s, a), \quad (7)$$

for all  $(s, a) \in \mathcal{S} \times \mathcal{A}$ .

---

### Algorithm 1 Pessimistic Value Iteration

---

- 1: **Require:** Dataset  $\mathcal{D} = \{(s_{\tau}, a_{\tau}, r_{\tau})\}_{\tau=1}^T$ .
  - 2: **Initialization:** Set  $\widehat{V}(\cdot) \leftarrow 0$  and construct  $\Gamma(\cdot, \cdot)$ .
  - 3: **while** not converged **do**
  - 4:   Construct  $(\widehat{\mathbb{B}}_{\gamma} \widehat{V})(\cdot, \cdot)$
  - 5:   Set  $\widehat{Q}(\cdot, \cdot) \leftarrow (\widehat{\mathbb{B}}_{\gamma} \widehat{V})(\cdot, \cdot) - \Gamma(\cdot, \cdot)$ .
  - 6:   Set  $\widehat{\pi}(\cdot | \cdot) \leftarrow \operatorname{argmax}_{\pi} \mathbb{E}_{\pi} [\widehat{Q}(\cdot, \cdot)]$ .
  - 7:   Set  $\widehat{V}(\cdot) \leftarrow \mathbb{E}_{\widehat{\pi}} [\widehat{Q}(\cdot, \cdot)]$ .
  - 8: **end while**
  - 9: **Return**  $\widehat{\pi}$
- 

---

### Algorithm 2 Model-Based Pessimistic Policy Optimization

---

- 1: **Require:** Dataset  $\mathcal{D}$ , discount factor  $\gamma$ , policy class  $\Pi$ , Model set  $\mathcal{M}$
- 2: Optimize policy with respect to the worst possible model:

$$\widehat{\pi} = \operatorname{argmax}_{\pi \in \Pi} \min_{M \in \mathcal{M}} V_{\gamma, M}^{\pi}. \quad (8)$$

- 3: **Return**  $\widehat{\pi}$
- 

Intuitively,  $\Gamma(s, a)$  represents the uncertainty when estimating the value function. The negative bonus ensures that we do not overestimate the value function due to finite samples with high probability, which allows us to give a performance lower bound with respect to the number of samples in the dataset.

MBPPO, on the contrary, considers the uncertainty in the model space to reduce potential sampling bias. By optimizing the performance in the worst possible model, as shown in Algorithm 2, we obtain a model-based offline algorithm, which gives better suboptimality bounds compared to the model-free counterpart when we use a proper set of models.

Both Algorithm 1 and Algorithm 2 are meta descriptions rather than detailed implementations. We provide more details of the algorithms in Appendix A.

## 3. Theoretical Analysis

This section characterizes the effects of a lower guidance discount factor in offline RL. We first show that a similar

regularization effect exists in offline algorithms as in online algorithms. We quantify this effect by providing a performance bound to analyze how other factors like the data size and the coverage coefficient affect this regularization effect. We then show an equivalence between a lower discount factor and the model-based pessimism. This equivalence leads to a performance guarantee with a lower guidance discount factor without other conservative regularizations. These two effects indicate that the discount factor plays a vital role in offline reinforcement learning.

### 3.1. Regularization Effect

As Jiang et al. (2015) suggests, the discount factor acts as a regularization coefficient to bound the complexity of the potential optimal policy class. However, it is unclear what affects the effectiveness of  $\gamma$  regularization in the offline setting, especially when the data coverage is poor and the algorithms have additional pessimism. Empirically, we found that the quality and size of the dataset are the main factors that affect the regularization from  $\gamma$ . To shed light on this observation, we first derive a performance bound in the linear MDP setting for model-free pessimistic algorithms like Algorithm 1. The analysis is analogous to (Jin et al., 2021), but we focus on the discount setting.

**Lemma 3.1.** *Suppose there exists an absolute constant*

$$c^\dagger = \sup_{x \in \mathbb{R}^d} \frac{x^\top \Sigma_{\pi^*, s} x}{x^\top \Sigma_{\mathcal{D}} x} < \infty, \quad (9)$$

for all  $s \in \mathcal{S}$  with probability  $1 - \xi/2$ , where

$$\begin{aligned} \Sigma_{\mathcal{D}} &= \frac{1}{N} \sum_{\tau=1}^N [\phi(s_\tau, a_\tau) \phi(s_\tau, a_\tau)^\top], \\ \Sigma_{\pi^*, s} &= \mathbb{E}_{\pi^*} [\phi(s_t, a_t) \phi(s_t, a_t)^\top \mid s_0 = s]. \end{aligned}$$

In Algorithm 1, we follow Equation (17) and (18), and set

$$\lambda = 1, \beta = c \cdot dr_{\max} \sqrt{\zeta} / (1 - \gamma), \zeta = \log(4dN / (1 - \gamma)\xi),$$

where  $c > 0$  is an absolute constant and  $\xi \in (0, 1)$  is the confidence parameter. Then with probability  $1 - \xi$ , the policy  $\hat{\pi}$  generated by Algorithm 1 satisfies

$$\text{SubOpt}(\hat{\pi}, s; \gamma) \leq \frac{2cr_{\max}}{(1 - \gamma)^2} \sqrt{c^\dagger d^3 \zeta / N}, \forall s \in \mathcal{S}.$$

*Proof.* See Appendix B.1 for a detailed proof.  $\square$

Equation (9) defines a finite coverage coefficient, namely  $c^\dagger$ , which represents the maximum ratio between the density of empirical state-action distribution and the density induced from the optimal policy. Intuitively, it represents the quality of the dataset. For example, the `expert` dataset has a low

coverage ratio while the `random` dataset may have a high ratio. The probability  $1 - \xi/2$  for a finite coverage coefficient is measured concerning the data collection process. That is, we are only making assumptions about the data collection process rather than the specific dataset. The dependence on  $\gamma$  in Lemma 3.1 suggests that the performance bound  $\text{SubOpt}(\hat{\pi}, s; \gamma)$  decreases as the discount factor gets lower. However, a lower discount factor also biases the optimal policy, as characterized by the following lemma.

**Lemma 3.2** (Jiang et al. (2015)). *For any MDP  $M$  with rewards in  $[0, r_{\max}]$ ,  $\forall \pi : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$  and  $\gamma \leq \gamma_e$ ,*

$$V_{M, \gamma}(\pi) \leq V_{M, \gamma_e}(\pi) \leq V_{M, \gamma}(\pi) + \frac{\gamma_e - \gamma}{(1 - \gamma)(1 - \gamma_e)} r_{\max}, \quad (10)$$

where  $\gamma_e$  is the evaluation discount factor.

The bound above is tight in a trivial case, while it is usually very loose in practice. Together, we have the following bound with a guidance discount factor  $\gamma$  different from the true evaluation discount factor  $\gamma_e$ .

**Theorem 3.3.** *Based on the same assumption and definition as in Lemma 3.1, we set  $\zeta = \log(4dN / (1 - \gamma)\xi)$ ,  $\beta = c \cdot dr_{\max} \sqrt{\zeta} / (1 - \gamma)$  and  $\lambda = 1$ . Then with probability  $1 - \xi$ , the suboptimality bound of the policy  $\hat{\pi}$  generated by Algorithm 1 satisfies*

$$\begin{aligned} \text{SubOpt}(\hat{\pi}; \gamma_e) &\leq \frac{2c}{(1 - \gamma)^2} \sqrt{c^\dagger d^3 \zeta / N} \cdot r_{\max} \\ &\quad + \frac{\gamma_e - \gamma}{(1 - \gamma)(1 - \gamma_e)} \cdot r_{\max}. \end{aligned}$$

*Proof.* The result follows immediately from Lemma 3.1 and Lemma 3.2.  $\square$

Theorem 3.3 gives an upper bound for pessimistic offline algorithms with two terms. Both terms monotonically depend on the guidance discount factor  $\gamma$  but in an opposite way, which suggests *there exists an optimal trade-off  $\gamma^* \in [0, \gamma_e]$* . It also suggests that this optimal trade-off  $\gamma^*$  is dependent on other factors like the coverage ratio of the dataset and the size of the dataset. A small or poorly covered dataset (i.e., a large coverage coefficient) makes the first term's coefficient larger, requiring a lower discount factor to achieve optimal performance. We empirically observe this effect on both toy examples as well as large D4RL tasks, as shown in Section 4.

### 3.2. Pessimistic Effect

This section analyses the pessimism effect of a lower discount factor. We show that, perhaps surprisingly, learning with a lower discount factor is equivalent to one type of the model-based pessimism mechanism, as depicted in Algorithm 2. This is characterized by the following lemma.



**Lemma 3.4.** *The optimal value function with a lower discount factor is equivalent to the pessimistic value function over a set of models. Formally, let*

$$\pi_{\mathcal{M}_\varepsilon}^* \in \operatorname{argmax}_\pi \min_{M \in \mathcal{M}_\varepsilon} V_{M,\gamma}(\pi), \quad (11)$$

where

$$\mathcal{M}_\varepsilon = \{M | \mathcal{P}_M(\cdot | s, a) = (1 - \varepsilon)\mathcal{P}_{M_0}(\cdot | s, a) + \varepsilon\mathcal{P}(\cdot)\},$$

and  $\mathcal{P}(\cdot)$  is an arbitrary distribution over  $\mathcal{S}$ , then we have

$$V_{M_0, (1-\varepsilon)\gamma}^* = V_{M_0, \gamma}(\pi_{\mathcal{M}_\varepsilon}^*) + \Delta, \quad (12)$$

where  $\Delta$  is an absolute constant.

*Proof.* See Appendix B.2 for a detailed proof.  $\square$

The equality in Lemma 3.4 shows that learning with a lower discount factor itself acts as a kind of model-based pessimism, which allows us to derive a bound without any other additional regularizations. We consider the model-based pessimism, where the model parameters are learned through maximum likelihood estimation (MLE). With the techniques in (Geer et al., 2000) that allow us to estimate the concentration rate of the MLE estimator, we have the following theorem. The proof of the following theorem is analogous to the analysis in (Uehara & Sun, 2021).

**Theorem 3.5.** (informal) *Suppose there exists an absolute constant*

$$c^\dagger = \sup_{x \in \mathbb{R}^d} \frac{x^\top \Sigma_{\pi^*} x}{x^\top \Sigma_\rho x} < \infty, \quad (13)$$

$$\Sigma_\rho = \mathbb{E}_\rho[\phi(s, a)\phi(s, a)^\top], \quad \Sigma_{\pi^*} = \mathbb{E}_{d^{\pi^*}}[\phi(s, a)\phi(s, a)^\top].$$

And suppose the underlying MDP follows the regularity condition in Assumption B.4. Set

$$\gamma = (1 - \varepsilon)\gamma_e, \quad \varepsilon = c_1 \sqrt{d\zeta/N}, \quad \zeta = \log^2(c_2 N d / \xi).$$

Then, with probability  $1 - \xi$ , Learning with a guidance discount factor  $\gamma$  yields a policy  $\hat{\pi}$  such that

$$\operatorname{SubOpt}(\hat{\pi}; \gamma_e) \leq \frac{c_3}{(1 - \gamma_e)^2} \sqrt{c^\dagger d^2 \zeta / N} \cdot r_{\max}, \quad (14)$$

where  $c_1, c_2, c_3$  are constants.

*Proof.* See Appendix B.3 for a detailed description and proof.  $\square$

Similar to  $c^\dagger$  in Equation (9),  $c^\ddagger$  can be interpreted as another coverage coefficient, and the difference between  $c^\dagger$  and  $c^\ddagger$  are technical. Theorem 3.5 shows that, with a properly chosen  $\gamma$ , we have a performance guarantee without

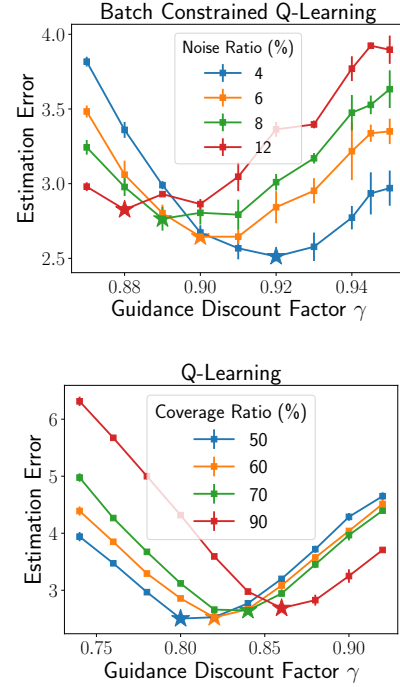


Figure 1. The estimation error of BCQ and Q-Learning in the random MDP task. The star shapes mark the minimum of the curve.

any other offline techniques. We name this effect of the discount factor the pessimistic effect. Note that to make this bound meaningful, we need the dataset size  $N$  to be large enough so that  $1 - \varepsilon > 0$ . This means that this theorem applies when the data is sufficient, contrary to the condition for the regularization effect. Compared to the bound in Theorem 3.3, we note that the bound in Theorem 3.5 is smaller with a factor of  $\sqrt{d}$  because a lower  $\gamma$  resembles the model-based pessimism rather than a model-free one.

We empirically verify this effect on the tabular MDPs and the D4RL dataset, where simple discount factor regularization is enough to derive a reasonable performance. This pessimistic effect suggests that  $\gamma$  affects the performance in offline RL differently compared to online settings.

## 4. Empirical Results

In this section we examine the role of discount factor through various experiments and aim to investigate the following questions: 1. How effective is the regularization effect of  $\gamma$  and how other factors affect its performance? 2. How effective is the pessimistic effect of  $\gamma$  and how does it contribute to the performance? 3. Is a lower guidance discount factor effective and essential in practical offline settings? We answer the questions above by experiments on both tabular MDPs and D4RL tasks. Each experiment result is averaged over three random seeds with the standard

## On the Role of Discount Factor in Offline Reinforcement Learning

Tasks	BCQ	BCQ ( $\gamma$ )	TD3+BC	TD3+BC ( $\gamma$ )	COMBO	COMBO ( $\gamma$ )
walker2d (0 noised traj)	59.6±2.7	51.5±3.6	<b>62.0±3.2</b>	52.2±1.1	26.1±3.2	<b>65.5±1.7</b>
walker2d (10 noised traj)	53.7±2.5	51.8±1.3	<b>60.9±1.2</b>	45.7±4.2	27.9±2.3	<b>63.1±1.6</b>
walker2d (50 noised traj)	20.3±3.3	<b>52.4±3.9</b>	4.3±1.2	<b>46.8±1.9</b>	27.2±1.6	<b>69.6±1.9</b>
walker2d (100 noised traj)	18.6±1.9	<b>52.1±2.2</b>	2.1±0.2	<b>46.6±1.3</b>	13.3±1.1	<b>70.7±2.3</b>
hopper (0 noised traj)	52.8±2.1	40.3±2.5	<b>52.5±1.8</b>	51.0±0.9	1.5±0.1	<b>53.5±3.2</b>
hopper (10 noised traj)	47.9±2.1	41.0±2.7	15.4±0.5	<b>47.9±0.3</b>	1.2±0.1	<b>56.5±2.5</b>
hopper (50 noised traj)	12.7±3.5	<b>44.1±1.9</b>	3.0±0.2	<b>47.0±0.5</b>	1.0±0.1	<b>48.6±4.2</b>
hopper (100 noised traj)	1.0±0.1	<b>41.6±0.6</b>	1.5±0.4	<b>46.3±0.7</b>	1.3±0.1	<b>52.3±1.7</b>
halfcheetah (0 noised traj)	40.2±1.3	<b>42.1±1.1</b>	45.3±1.5	<b>46.9±1.6</b>	32.6±1.6	27.6±1.5
halfcheetah (10 noised traj)	39.5±0.3	<b>40.2±3.3</b>	45.7±0.4	<b>47.3±1.6</b>	32.3±2.8	29.7±2.7
halfcheetah (50 noised traj)	36.5±0.9	<b>37.8±0.8</b>	45.9±0.3	<b>47.3±1.3</b>	31.1±4.7	28.0±1.6
halfcheetah (100 noised traj)	35.4±1.1	<b>36.4±1.7</b>	47.3±1.0	<b>46.1±1.8</b>	30.0±1.9	29.3±0.6

Table 2. Experimental results on noised D4RL tasks, such as walker2d ( $x$ ), hopper ( $x$ ) and halfcheetah ( $x$ ), containing 50 medium and  $x$  noised trajectories. We conduct a comparison between **original** and **lower** guidance discount factor (e.g., BCQ and BCQ ( $\gamma$ )).

deviation.

### 4.1. Regularization Effect

#### 4.1.1. TABULAR EXPERIMENTS

We first adopt the BCQ-style offline method in a tabular MDP environment to investigate the effectiveness of the lower discount factor. We consider a random MDP, where the state space consists of a  $30 \times 30$  grid, and each state has 10 actions. The reward and the transition probabilities are generated randomly. Given the tabular MDP, we compute the true optimal value  $Q_{\gamma_e}^*$ , where  $\gamma_e = 0.95$ . Then, we calculate the behavior policy according to  $\mu(\cdot | s) = \text{softmax}(Q_{\gamma_e}^*(s, \cdot)) \cdot \text{mask}(s, \cdot)$ , where  $\text{mask}(s, \cdot)$  is randomly selected in state-action space to approximate the unseen pairs in offline tasks. Further, we calculate  $\hat{\mu}(\cdot | s) = \text{softmax}(Q_{\gamma_e}^*(s, \cdot)) \cdot \overline{\text{mask}}(s, \cdot)$  to approximate the generative model in BCQ, where  $\|\overline{\text{mask}}\|_1 - \|\text{mask}\|_1 > 0$ . The Noise Ratio =  $\frac{\|\overline{\text{mask}}\|_1 - \|\text{mask}\|_1}{\|\text{mask}\|_1} * 100\%$  represents the inaccuracy of the generative model. We calculate BCQ-style value function  $\hat{Q}_\gamma$  by constraining the maximization operator into the finite state-action space:

$$\hat{Q}(s, a) = r(s, a) + \gamma \max_{a' \text{ s.t. } \hat{\mu}(a'|s') > 0} \hat{Q}(s', a') \quad (15)$$

In the experiments, the proportion of masked state-action pairs is 0.5, and the noise ratio coefficients are {4%, 6%, 8%, 12%} respectively. We compute the estimation error  $\|(\hat{Q}_\gamma - Q_{\gamma_e}^*)\|_\infty$  at seen state-action pairs for different noise ratio coefficients and different discount factors. The results are summarized in Figure 1. Each plot shows the average estimation error across 100 MDP instances. The experimental results demonstrate that the lower guidance discount factor significantly reduces the estimation error in the tabular offline task. Moreover, as the noise ratio increases, the optimal discount factor  $\gamma^*$  marked by the star shapes

becomes smaller. This indicates that discount regularization is more significant when function approximation error is large, which may result from insufficient data or poor data coverage.

#### 4.1.2. EXPERIMENTAL RESULTS ON D4RL TASKS

This section investigates the regularization effect in complex tasks. Specifically, we evaluate various offline RL algorithms (TD3+BC (Fujimoto & Gu, 2021), BCQ (Fujimoto et al., 2019) and COMBO (Yu et al., 2021)) with a lower guidance discount factor on the limited and noised D4RL benchmark of MuJoCo tasks. To test performance in low data regimes and poor coverage scenarios, the training dataset in our experiments contains 50 medium trajectories and additional noisy trajectories ranging from 0 to 100. The noised trajectories are fragments of the random datasets in D4RL. We use the author-provided implementation or the recognized code to ensure a fair and identical experimental evaluation across algorithms. The experimental results are shown from two aspects: coverage ratio and data size.

**Coverage Ratio.** We report the final performance of TD3+BC with different amounts of noisy data in Table 2 and the detailed training curves in Appendix C. The training datasets contain 50 medium and  $x$  noised trajectories, with  $x$  ranging from 0 to 100. The noised trajectories are fragments from the random dataset. Results show that the performance of current offline RL methods with original discount factor ( $\gamma = 0.99$ ) degrades with more noisy data due to the poor coverage ratio in the dataset. In contrast, offline RL methods with a lower guidance discount factor ( $\gamma = 0.95$ ) achieve stable and robust performance in most scenarios.

Further, the generated data of model-based offline RL is usually noisy since the challenges of the limited data (please see the performance gap between COMBO and COMBO ( $\gamma$ ) in walker2d and hopper task).

## On the Role of Discount Factor in Offline Reinforcement Learning

Halfcheetah	random-v2	medium-v2	medium-expert-v2	medium-replay-v2	expert-v2
SAC-N ( $\gamma=0.95$ )	<b>30.0±1.6</b>	<b>65.1±0.9</b>	<b>51.4±2.2</b>	<b>28.1±1.2</b>	<b>82.7±0.8</b>
SAC-N ( $\gamma=0.99$ )	26.6±1.5	48.7±1.3	26.7±1.1	0.6±0.1	80.2±0.6
Hopper	random-v2	medium-v2	medium-expert-v2	medium-replay-v2	expert-v2
SAC-N ( $\gamma=0.95$ )	8.4±1.7	<b>22.4±2.1</b>	<b>23.1±1.9</b>	15.5±3.2	<b>14.5±2.6</b>
SAC-N ( $\gamma=0.99$ )	14.5±3.5	7.1±2.0	15.4±1.4	100.9±0.3	2.3±0.3

Table 3. Experimental results on Halfcheetah and Hopper tasks in D4RL, where the Q-ensemble size N is 2 in Halfcheetah tasks and N is 50 in Hopper tasks.

Adroit	pen-expert-v0	door-expert-v0	hammer-expert-v0	relocate-expert-v0
SAC-N (lower $\gamma$ )	<b>97.1±3.2</b>	<b>106.4±1.9</b>	<b>100.6±2.3</b>	0.5±0.1
SAC-N ( $\gamma=0.99$ )	3.6±1.1	2.2±0.2	65.5±4.2	0.4±0.1

Table 4. Experimental results on Adroit tasks in D4RL, where the Q-ensemble size N is 50. We select  $\gamma = 0.95$  in pen-expert-v0, hammer-expert-v0 and relocate-expert-v0 tasks. We select lower  $\gamma = 0.9$  in door-expert-v0 tasks.

Most scenarios in Table 2 adopt  $\gamma = 0.95$  as a lower discount factor other than BCQ ( $\gamma = 0.9$ ) and COMBO ( $\gamma = 0.9$ ) in hopper tasks.

**Data Size.** We evaluate TD3+BC on datasets containing  $x$  medium and 100 noised trajectories, with  $x$  ranging from 0 to 1000. We select the optimal discount factor  $\gamma^* \in [0.89, 0.99]$ . Experimental results in Figure 4 show that the optimal discount factor  $\gamma^*$  increases with the number of trajectories. That is, the effect of the discount factor is more significant when the size of the dataset is small, which is consistent with the analysis in Theorem 3.3. The experimental results also suggest we prefer a higher discount factor in large datasets (e.g., the standard datasets in D4RL tasks).

**Sparse Reward.** We evaluate EVL (Ma et al., 2021) with various  $\gamma$  on standard Antmaze tasks in D4RL, where the stitching (approximate dynamical programming) is necessary. Further, Antmaze is a sparse reward task that tightens the bound in Lemma 3.1 up to a  $1/(1-\gamma)$  factor, making the optimal  $\gamma$  larger. However, a lower  $\gamma$  still works better than the usual 0.99 even in large tasks (Please refer to experimental results in Table 5). These results show that the regularization effects and finetuning  $\gamma$  work generally, even in sparse reward tasks. (In the medium and large tasks, we set the terminal reward  $r_T$  as 100 and 50, and the expectile ratio is 0.96.)

play-v0	umaze	medium ( $r_T=50$ )	large ( $r_T=100$ )
$\gamma=0.98$	86.7±2.5	73.6±2.3	<b>45.7±2.2</b>
$\gamma=0.99$	87.4±2.6	71.0±1.8	37.8±3.0

Table 5. Results on Antmaze-play-v0 tasks with EVL.




Figure 2. Relationship between  $\gamma$  and  $\alpha$  of TD3+BC on halfcheetah task. The value is the normalized return metric. The optimal  $\gamma^*$  is marked with orange color.

### 4.2. Pessimistic Effect

We conduct experiments on both tabular MDPs and D4RL tasks with a lower discount factor and no other offline regularization to investigate the pessimistic effect.

**Tabular MDPs.** We adopt the same setting and evaluation metric as the toy example in Section 4.1.1. In the experiments, we define the coverage ratio as the proportion of masked state-action pairs ranging from 50% to 90%. The experimental results in Figure 1 show that a lower discount factor promotes offline RL algorithms to have a better estimation, and the effect is more significant when the data coverage is poor.

**D4RL Task.** We evaluate the standard off-policy algorithm SAC with ensemble networks on standard D4RL datasets, which is shown in Table 3 and Table 4. Adroit tasks require dynamic programming to infer the complete action

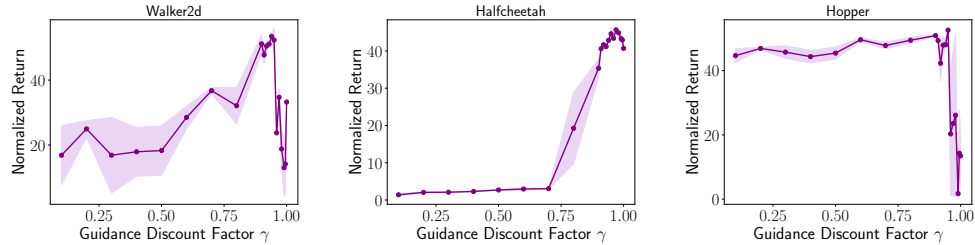


Figure 3. Performance of TD3+BC on noised D4RL tasks containing 50 medium and 25 noised random trajectories. We adopt the normalized return metric proposed by D4RL benchmark (Fu et al., 2020). Scores roughly range from 0 to 100.




Figure 4. The relationship between optimal discount factor  $\gamma^*$  and datsize on noised D4RL tasks.

sequence for finishing the task, while a lower  $\gamma$  works surprisingly well with  $Q$ -ensembles. Moreover, the experimental results on mujoco tasks show that SAC achieves a noticeable performance with a lower guidance discount factor, demonstrating the role of the discount factor as a pessimistic mechanism. While the pessimism of lower  $\gamma$  is not as good as the state-of-the-art offline algorithms, it serves as a good baseline. It significantly surpasses online algorithms with no discount factor pessimism. The pessimistic mechanism of the discount factor is coarse, in the sense that it is a single parameter that affects all state-action pairs, which explains why the pessimism of lower  $\gamma$  is not as good as the state-of-the-art offline algorithms. We leave finding methods for more fine-grained discount factor control as interesting future work.

### 4.3. Effectiveness of Discount Regularization

#### 4.3.1. SENSITIVITY OF THE DISCOUNT FACTOR

This section tests the performance sensitivity regarding  $\gamma$  and whether we need to fine-tune  $\gamma$  for each task. To this end, we evaluate TD3+BC on datasets containing 50 medium and 25 noised trajectories with various  $\gamma$ . In Figure 3, TD3+BC achieves a high performance by decreasing  $\gamma$  from 0.99 to around 0.95. Different offline RL algorithms have minor differences in selecting  $\gamma$  on various tasks (e.g., the particular case in Table 2). The suitable area of  $\gamma$  is almost identical [0.95, 0.99].

#### 4.3.2. DISCUSSION BETWEEN $\gamma$ AND $\alpha$ IN TD3+BC

Many offline RL algorithms achieve a trade-off between conservation and generalization by some other parameter (e.g., the regularized hyper-parameter  $\alpha$  in TD3+BC). This naturally leads to the following question: Is finetuning  $\gamma$  more effective than simple regularizations like behavior cloning? To answer this question, we evaluate TD3+BC on the halfcheetah task containing 100 medium and random trajectories fragments. Experimental results in Figure 2 show that by properly combining  $\alpha$  and  $\gamma$ , we can solve tasks more effectively. This indicates that  $\gamma$  offers more flexibility than the original regularization. Moreover, note that the behavior cloning performance in this task is poor (the normalized score is 0.0). Therefore, the role of the lower discount factor is not equivalent to behavior cloning.

Further, the experimental results also suggest that other conservative regularization affects the optimal guidance discount factor (shown in orange in Figure 2). In general, the stronger the other regularization, the higher the optimal discount factor, matching our intuition that there is a trade-off between two regularizations.

## 5. Conclusion

This paper examines the two distinct effects of the discount factor in offline RL, i.e., the regularization effect and the pessimistic effect. On the one hand, the discount factor acts as a regulator to trade-off optimality with sample efficiency upon existing offline techniques like a negative uncertainty-



based bonus. On the other hand, we show that a lower guidance discount factor is equivalent to the model-based pessimism, where we optimize the policy’s performance in the worst possible models. We quantify the above effects by analyzing the performance bounds with a lower guidance discount factor in linear MDPs. Moreover, we verify the above theoretical observations in tabular MDPs and D4RL tasks. Empirical results show that a lower discount factor can significantly improve performance in two scenarios. The first is when the data size is limited or there is poor data coverage, and we can apply a lower discount factor with additional pessimism. The second is when the data is sufficient, and we can directly use a lower discount factor as a proper pessimistic mechanism.

Our work suggests that the discount factor plays a vital role in offline RL, promoting current offline RL methods in complex and diverse scenarios. This leaves several interesting future works: (i) Can a lower guidance discount factor be better integrated with current offline algorithms, like fine-grained guidance  $\gamma$  at each transition? (ii) Is there a better theoretical explanation for the success of a lower discount factor in offline RL? (iii) How to develop more efficient offline algorithms in limited data size and insufficient coverage scenarios?

## 6. Acknowledgements

This work is supported in part by Science and Technology Innovation 2030 - “New Generation Artificial Intelligence” Major Project (No. 2018AAA0100904) and National Natural Science Foundation of China (62176135).

## References

- Abbasi-Yadkori, Y., Pál, D., and Szepesvári, C. Improved algorithms for linear stochastic bandits. *Advances in neural information processing systems*, 24:2312–2320, 2011.
- Agarwal, R., Schuurmans, D., and Norouzi, M. An optimistic perspective on offline reinforcement learning. In *International Conference on Machine Learning*, pp. 104–114. PMLR, 2020.
- Amit, R., Meir, R., and Ciosek, K. Discount factor as a regularizer in reinforcement learning. In *International conference on machine learning*, pp. 269–278. PMLR, 2020.
- An, G., Moon, S., Kim, J.-H., and Song, H. O. Uncertainty-based offline reinforcement learning with diversified q-ensemble. *Advances in Neural Information Processing Systems*, 34, 2021.
- Argenson, A. and Dulac-Arnold, G. Model-based offline planning. *arXiv preprint arXiv:2008.05556*, 2020.
- Brandfonbrener, D., Whitney, W. F., Ranganath, R., and Bruna, J. Offline rl without off-policy evaluation. *arXiv preprint arXiv:2106.08909*, 2021.
- Chen, Y.-C., Kochenderfer, M. J., and Spaan, M. T. Improving offline value-function approximations for pomdps by reducing discount factors. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 3531–3536. IEEE, 2018.
- Fedus, W., Gelada, C., Bengio, Y., Bellemare, M. G., and Larochelle, H. Hyperbolic discounting and learning over multiple horizons. *arXiv preprint arXiv:1902.06865*, 2019.
- Fu, J., Kumar, A., Nachum, O., Tucker, G., and Levine, S. D4rl: Datasets for deep data-driven reinforcement learning. *arXiv preprint arXiv:2004.07219*, 2020.
- Fujimoto, S. and Gu, S. S. A minimalist approach to offline reinforcement learning. *arXiv preprint arXiv:2106.06860*, 2021.
- Fujimoto, S., Meger, D., and Precup, D. Off-policy deep reinforcement learning without exploration. In *International Conference on Machine Learning*, pp. 2052–2062. PMLR, 2019.
- Geer, S. A., van de Geer, S., and Williams, D. *Empirical Processes in M-estimation*, volume 6. Cambridge university press, 2000.
- Ghasemipour, S. K. S., Schuurmans, D., and Gu, S. S. Emaq: Expected-max q-learning operator for simple yet effective offline and online rl. In *International Conference on Machine Learning*, pp. 3682–3691. PMLR, 2021.
- Hong, M., Wai, H.-T., Wang, Z., and Yang, Z. A two-timescale framework for bilevel optimization: Complexity analysis and application to actor-critic. *arXiv preprint arXiv:2007.05170*, 2020.
- Jiang, N., Kulesza, A., Singh, S., and Lewis, R. The dependence of effective planning horizon on model accuracy. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, pp. 1181–1189. Citeseer, 2015.
- Jin, C., Yang, Z., Wang, Z., and Jordan, M. I. Provably efficient reinforcement learning with linear function approximation. In *Conference on Learning Theory*, pp. 2137–2143. PMLR, 2020.
- Jin, Y., Yang, Z., and Wang, Z. Is pessimism provably efficient for offline rl? In *International Conference on Machine Learning*, pp. 5084–5096. PMLR, 2021.

- Kidambi, R., Rajeswaran, A., Netrapalli, P., and Joachims, T. Morel: Model-based offline reinforcement learning. *arXiv preprint arXiv:2005.05951*, 2020.
- Kostrikov, I., Nair, A., and Levine, S. Offline reinforcement learning with implicit q-learning. *arXiv preprint arXiv:2110.06169*, 2021.
- Kumar, A., Fu, J., Soh, M., Tucker, G., and Levine, S. Stabilizing off-policy q-learning via bootstrapping error reduction. *Advances in Neural Information Processing Systems*, 32:11784–11794, 2019.
- Kumar, A., Zhou, A., Tucker, G., and Levine, S. Conservative q-learning for offline reinforcement learning. *arXiv preprint arXiv:2006.04779*, 2020.
- Levine, S., Kumar, A., Tucker, G., and Fu, J. Offline reinforcement learning: Tutorial, review, and perspectives on open problems. *arXiv preprint arXiv:2005.01643*, 2020.
- Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., and Wierstra, D. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.
- Ma, X., Yang, Y., Hu, H., Liu, Q., Yang, J., Zhang, C., Zhao, Q., and Liang, B. Offline reinforcement learning with value-based episodic memory. *arXiv preprint arXiv:2110.09796*, 2021.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al. Human-level control through deep reinforcement learning. *nature*, 518(7540): 529–533, 2015.
- Nair, A., Dalal, M., Gupta, A., and Levine, S. Accelerating online reinforcement learning with offline datasets. *arXiv preprint arXiv:2006.09359*, 2020.
- Peng, X. B., Kumar, A., Zhang, G., and Levine, S. Advantage-weighted regression: Simple and scalable off-policy reinforcement learning. *arXiv preprint arXiv:1910.00177*, 2019.
- Petrik, M. and Scherrer, B. Biasing approximate dynamic programming with a lower discount factor. *Advances in neural information processing systems*, 21:1265–1272, 2008.
- Rashidinejad, P., Zhu, B., Ma, C., Jiao, J., and Russell, S. Bridging offline reinforcement learning and imitation learning: A tale of pessimism. *arXiv preprint arXiv:2103.12021*, 2021.
- Rathnam, S., Murphy, S. A., and Doshi-Velez, F. Comparison and unification of three regularization methods in batch reinforcement learning. *arXiv preprint arXiv:2109.08134*, 2021.
- Romoff, J., Henderson, P., Touati, A., Brunskill, E., Pineau, J., and Ollivier, Y. Separating value functions across timescales. In *International Conference on Machine Learning*, pp. 5468–5477. PMLR, 2019.
- Shalev-Shwartz, S., Shammah, S., and Shashua, A. Safe, multi-agent, reinforcement learning for autonomous driving. *arXiv preprint arXiv:1610.03295*, 2016.
- Sherstan, C., Dohare, S., MacGlashan, J., Günther, J., and Pilarski, P. M. Gamma-nets: Generalizing value estimation over timescale. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pp. 5717–5725, 2020.
- Siegel, N. Y., Springenberg, J. T., Berkenkamp, F., Abdolmaleki, A., Neunert, M., Lampe, T., Hafner, R., Heess, N., and Riedmiller, M. Keep doing what worked: Behavioral modelling priors for offline reinforcement learning. *arXiv preprint arXiv:2002.08396*, 2020.
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Van Den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., et al. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484–489, 2016.
- Singh, A., Yu, A., Yang, J., Zhang, J., Kumar, A., and Levine, S. Cog: Connecting new skills to past experience with offline reinforcement learning. *arXiv preprint arXiv:2010.14500*, 2020.
- Swaminathan, A. and Joachims, T. Batch learning from logged bandit feedback through counterfactual risk minimization. *The Journal of Machine Learning Research*, 16(1):1731–1755, 2015.
- Uehara, M. and Sun, W. Pessimistic model-based offline reinforcement learning under partial coverage. *arXiv preprint arXiv:2107.06226*, 2021.
- Wainwright, M. J. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge University Press, 2019.
- Wu, Y., Tucker, G., and Nachum, O. Behavior regularized offline reinforcement learning. *arXiv preprint arXiv:1911.11361*, 2019.
- Wu, Y., Zhai, S., Srivastava, N., Susskind, J., Zhang, J., Salakhutdinov, R., and Goh, H. Uncertainty weighted actor-critic for offline reinforcement learning. *arXiv preprint arXiv:2105.08140*, 2021.
- Xie, T., Cheng, C.-A., Jiang, N., Mineiro, P., and Agarwal, A. Bellman-consistent pessimism for offline reinforcement learning. *arXiv preprint arXiv:2106.06926*, 2021.

- Xu, Z., van Hasselt, H. P., and Silver, D. Meta-gradient reinforcement learning. *Advances in Neural Information Processing Systems*, 31:2396–2407, 2018.
- Yang, L. and Wang, M. Sample-optimal parametric q-learning using linearly additive features. In *International Conference on Machine Learning*, pp. 6995–7004. PMLR, 2019.
- Yang, Y., Ma, X., Li, C., Zheng, Z., Zhang, Q., Huang, G., Yang, J., and Zhao, Q. Believe what you see: Implicit constraint approach for offline multi-agent reinforcement learning. *arXiv preprint arXiv:2106.03400*, 2021.
- Yu, T., Thomas, G., Yu, L., Ermon, S., Zou, J., Levine, S., Finn, C., and Ma, T. Mopo: Model-based offline policy optimization. *arXiv preprint arXiv:2005.13239*, 2020.
- Yu, T., Kumar, A., Rafailov, R., Rajeswaran, A., Levine, S., and Finn, C. Combo: Conservative offline model-based policy optimization. *arXiv preprint arXiv:2102.08363*, 2021.
- Zhang, S., Laroche, R., van Seijen, H., Whiteson, S., and Combes, R. T. d. A deeper look at discounting mismatch in actor-critic algorithms. *arXiv preprint arXiv:2010.01069*, 2020.

## A. Algorithm Details

In this section, we describe some details of the algorithms mentioned in Section 2.1.

### A.1. Pessimistic Value Iteration

In linear MDPs, we can construct  $\widehat{\mathbb{B}}_\gamma \widehat{V}$  and  $\Gamma$  based on  $\mathcal{D}$  as follows, where  $\widehat{\mathbb{B}}_\gamma \widehat{V}$  is the empirical estimation for  $\mathbb{B}_\gamma \widehat{V}$ . For a given dataset  $\mathcal{D} = \{(s_\tau, a_\tau, r_\tau)\}_{\tau=1}^N$ , we define the empirical mean squared Bellman error (MSBE) as

$$M(w) = \sum_{\tau=1}^N (r_\tau + \gamma \widehat{V}(s_{\tau+1}) - \phi(s_\tau, a_\tau)^\top w)^2 + \lambda \|w\|_2^2$$

Here  $\lambda > 0$  is the regularization parameter. Note that  $\widehat{w}$  has the closed form

$$\begin{aligned} \widehat{w} &= \Lambda^{-1} \left( \sum_{\tau=1}^N \phi(s_\tau, a_\tau) \cdot (r_\tau + \gamma \widehat{V}(s_{\tau+1})) \right), \\ \text{where } \Lambda &= \lambda I + \sum_{\tau=1}^N \phi(s_\tau, a_\tau) \phi(s_\tau, a_\tau)^\top. \end{aligned} \quad (16)$$

Then we simply let

$$\widehat{\mathbb{B}}_\gamma \widehat{V} = \langle \phi, \widehat{w} \rangle. \quad (17)$$

Meanwhile, we construct  $\Gamma$  based on  $\mathcal{D}$  as

$$\Gamma(s, a) = \beta \cdot (\phi(s, a)^\top \Lambda^{-1} \phi(s, a))^{1/2}. \quad (18)$$

Here  $\beta > 0$  is the scaling parameter.

### A.2. Model-based Pessimistic Policy Optimization

To give a proper performance bound, we consider the following model set

$$\mathcal{M}_{\mathcal{D}} = \left\{ P(\cdot | s, a) \in \mathcal{M} \mid \mathbb{E}_{\mathcal{D}} \left[ D_{\text{TV}}(\widehat{P}(\cdot | s, a), P(\cdot | s, a))^2 \right] \leq \xi \right\}, \quad (19)$$

where  $\widehat{P} = \operatorname{argmax}_P \mathbb{E}_{\mathcal{D}} [\ln P(s' | s, a)]$  and  $\mathcal{M}$  is the set of linear models. In practice, we can parameterize the model  $P_\theta(\cdot | s, a)$  and train the model via the following loss

$$\mathcal{L}(\theta, \mathcal{D}) = \frac{1}{N} \sum_{\tau=1}^N \ln P_\theta(s_{\tau+1} | s_\tau, a_\tau). \quad (20)$$

When assuming the transitions are Gaussian, the MLE objective can be reduced to prediction loss as follows

$$\mathcal{L}(\theta, \mathcal{D}) = \frac{1}{N} \sum_{\tau=1}^N \|f_\theta(s_\tau, a_\tau) - s_{\tau+1}\|_2. \quad (21)$$

As to the minimax optimization in (8), we can use techniques like bi-level optimization (Hong et al., 2020) to get the approximate solution.

## B. Additional Lemmas and Missing Proofs

### B.1. Proof of Lemma 3.1

*Proof.* For a sufficiently large  $\lambda$ , it is easy to see that  $\mathcal{T}\widehat{V} := \max_a (\widehat{\mathbb{B}}_\gamma \widehat{V} - \Gamma)$  is a contraction. Without loss of generality, we assume  $\lambda = 1$ . Then Algorithm 1 converges and we have

$$\begin{aligned} \widehat{V}(\cdot) &= \max_a \widehat{Q}(\cdot, a), \\ \widehat{Q}(\cdot, \cdot) &= \widehat{\mathbb{B}}_\gamma \widehat{V} - \Gamma(\cdot, \cdot). \end{aligned}$$

From the definition of  $\delta(\cdot, \cdot)$ , we have

$$\delta(s, a) = \mathbb{B}_\gamma \widehat{V}(s) - \widehat{Q}(s, a) = \mathbb{B}_\gamma \widehat{V}(s) - \widehat{\mathbb{B}}_\gamma \widehat{V} + \Gamma(s, a). \quad (22)$$

Under the condition of Lemma B.2, it holds that

$$0 \leq \delta(s, a) \leq 2\Gamma(s, a), \text{ for all } s, a. \quad (23)$$

From Lemma B.1, we have

$$\begin{aligned} \text{SubOpt}(\widehat{\pi}, s; \gamma) &= -\mathbb{E}_{\widehat{\pi}} \left[ \sum_{t=0}^{\infty} \gamma^t \delta(s_t, a_t) \mid s_0 = s \right] + \mathbb{E}_{\pi^*} \left[ \sum_{t=0}^{\infty} \gamma^t \delta(s_t, a_t) \mid s_0 = s \right] \\ &\quad + \mathbb{E}_{\pi^*} \left[ \sum_{t=0}^{\infty} \gamma^t \left\langle \widehat{Q}(s_t, \cdot), \pi^*(\cdot | s_t) - \widehat{\pi}(\cdot | s_t) \right\rangle \mid s_0 = s \right] \\ &\leq -\mathbb{E}_{\widehat{\pi}} \left[ \sum_{t=0}^{\infty} \gamma^t \delta(s_t, a_t) \mid s_0 = s \right] + \mathbb{E}_{\pi^*} \left[ \sum_{t=0}^{\infty} \gamma^t \delta(s_t, a_t) \mid s_0 = s \right] \\ &\leq 2\mathbb{E}_{\pi^*} \left[ \sum_{t=0}^{\infty} \gamma^t \Gamma(s_t, a_t) \mid s_0 = s \right] \\ &= 2\beta \mathbb{E}_{\pi^*} \left[ \sum_{t=0}^{\infty} \gamma^t (\phi(s_t, a_t)^\top \Lambda^{-1} \phi(s_t, a_t))^{1/2} \mid s_0 = s \right]. \end{aligned} \quad (24)$$

Here the first inequality follows from the fact that  $\widehat{\pi}(\cdot | s) = \operatorname{argmax}_\pi \langle \widehat{Q}(\cdot, \cdot), \pi(\cdot | s) \rangle$  and the second inequality follows from Equation (23). Then the following event

$$\mathcal{E} = \left\{ \text{SubOpt}(\widehat{\pi}, s; \gamma) \leq 2\beta \mathbb{E}_{\pi^*} \left[ \sum_{t=0}^{\infty} \gamma^t (\phi(s_t, a_t)^\top \Lambda^{-1} \phi(s_t, a_t))^{1/2} \mid s_0 = s \right] \text{ for all } s \in \mathcal{S} \right\} \quad (25)$$

holds with probability  $1 - \xi/2$ . From the assumption in Equation (9), the following event

$$\mathcal{E}^\dagger = \left\{ c^\dagger \cdot \frac{1}{N} \sum_{\tau=1}^N \phi(s_\tau, a_\tau) \phi(s_\tau, a_\tau)^\top \succeq \mathbb{E}_{\pi^*} [\phi(s_t, a_t) \phi(s_t, a_t)^\top \mid s_0 = s] \text{ for all } s \in \mathcal{S} \right\}$$

also holds with probability  $1 - \xi/2$ . Then from the union bound, the event  $\mathcal{E} \cap \mathcal{E}^\dagger$  holds with probability  $1 - \xi$ . We condition on this event here after. By the Cauchy-Schwarz inequality, we have

$$\begin{aligned} &\mathbb{E}_{\pi^*} \left[ \sum_{t=0}^{\infty} \gamma^t (\phi(s_t, a_t)^\top \Lambda^{-1} \phi(s_t, a_t))^{1/2} \mid s_0 = s \right] \\ &= \frac{1}{1-\gamma} \mathbb{E}_{d^{\pi^*}} \left[ \sqrt{\operatorname{Tr} (\phi(s, a)^\top \Lambda^{-1} \phi(s, a))} \mid s_0 = s \right] \\ &= \frac{1}{1-\gamma} \mathbb{E}_{d^{\pi^*}} \left[ \sqrt{\operatorname{Tr} (\phi(s, a) \phi(s, a)^\top \Lambda^{-1})} \mid s_0 = s \right] \\ &\leq \frac{1}{1-\gamma} \sqrt{\operatorname{Tr} \left( \mathbb{E}_{d^{\pi^*}} [\phi(s, a) \phi(s, a)^\top \mid s_0 = s] \Lambda^{-1} \right)} \\ &= \frac{1}{1-\gamma} \sqrt{\operatorname{Tr} \left( \Sigma_{\pi^*, s}^\top \Lambda^{-1} \right)}, \end{aligned} \quad (26)$$



for all  $s \in \mathcal{S}$ . On the event  $\mathcal{E} \cap \mathcal{E}^\dagger$ , we have

$$\begin{aligned} \text{SubOpt}(\hat{\pi}, s; \gamma) &\leq 2\beta \mathbb{E}_{\pi^*} \left[ \sum_{t=0}^{\infty} \gamma^t (\phi(s_t, a_t)^\top \Lambda^{-1} \phi(s_t, a_t))^{1/2} \mid s_0 = s \right] \\ &\leq \frac{2\beta}{1-\gamma} \sqrt{\text{Tr} \left( \Sigma_{\pi^*, s} \cdot \left( I + \frac{1}{c^\dagger} \cdot N \cdot \Sigma_{\pi^*, s} \right)^{-1} \right)} \\ &= \frac{2\beta}{1-\gamma} \sqrt{\sum_{j=1}^d \frac{\lambda_j(s)}{1 + \frac{1}{c^\dagger} \cdot N \cdot \lambda_j(s)}}. \end{aligned}$$

Here  $\{\lambda_j(s)\}_{j=1}^d$  are the eigenvalues of  $\Sigma_{\pi^*, s}$  for all  $s \in \mathcal{S}$ , the first inequality follows from the definition of  $\mathcal{E}$  in Equation (25), and the second inequality follows from Equation (26) and the definition of  $\mathcal{E}^\dagger$  in Equation (9). Meanwhile, by Definition 2.1, we have  $\|\phi(s, a)\| \leq 1$  for all  $(s, a) \in \mathcal{S} \times \mathcal{A}$ . By Jensen's inequality, we have

$$\|\Sigma_{\pi^*, s}\|_{\text{op}} \leq \mathbb{E}_{\pi^*} [\|\phi(s, a)\phi(s, a)^\top\|_{\text{op}} \mid s_0 = s] \leq 1 \quad (27)$$

for all  $s \in \mathcal{S}$ . As  $\Sigma_{\pi^*, s}$  is positive semidefinite, we have  $\lambda_j(s) \in [0, 1]$  for all  $s \in \mathcal{S}$  and all  $j \in [d]$ . Hence, on  $\mathcal{E} \cap \mathcal{E}^\dagger$ , we have

$$\begin{aligned} \text{SubOpt}(\hat{\pi}, s; \gamma) &\leq \frac{2\beta}{1-\gamma} \sqrt{\sum_{j=1}^d \frac{\lambda_j(s)}{1 + \frac{1}{c^\dagger} \cdot N \cdot \lambda_j(s)}} \\ &\leq \frac{2\beta}{1-\gamma} \sqrt{\sum_{j=1}^d \frac{1}{1 + \frac{1}{c^\dagger} \cdot N}} \leq \frac{2c r_{\max}}{(1-\gamma)^2} \sqrt{c^\dagger d^3 \zeta / N} \end{aligned}$$

for all  $x \in \mathcal{S}$ , where the second inequality follows from the fact that  $\lambda_j(s) \in [0, 1]$  for all  $s \in \mathcal{S}$  and all  $j \in [d]$ , while the third inequality follows from the choice of the scaling parameter  $\beta > 0$ . Then we have the conclusion in Lemma 3.1.  $\square$

**Lemma B.1** (Suboptimality Decomposition). *We have*

$$\begin{aligned} \text{SubOpt}(\hat{\pi}, s; \gamma) &= -\mathbb{E}_{\hat{\pi}} \left[ \sum_{t=0}^{\infty} \gamma^t \delta(s_t, a_t) \mid s_0 = s \right] + \mathbb{E}_{\pi^*} \left[ \sum_{t=0}^{\infty} \gamma^t \delta(s_t, a_t) \mid s_0 = s \right] \\ &\quad + \mathbb{E}_{\pi^*} \left[ \sum_{t=0}^{\infty} \gamma^t \langle \hat{Q}(s_t, \cdot), \pi^*(\cdot | s_t) - \hat{\pi}(\cdot | s_t) \rangle \mid s_0 = s \right], \end{aligned} \quad (28)$$

where  $\langle f, g \rangle = \int_{a \in \mathcal{A}} f(a)g(a)da$ .

*Proof.* We have

$$\text{SubOpt}(\hat{\pi}, s; \gamma) = V^{\pi^*}(s) - V^{\hat{\pi}}(s) = V^{\pi^*}(s) - \hat{V}(s) + \hat{V}(s) - V^{\hat{\pi}}(s).$$

The first term satisfies

$$\begin{aligned} \hat{V}(s) - V^{\hat{\pi}}(s) &= \mathbb{E}_{a \sim \hat{\pi}} [\hat{Q}(s, a)] - \mathbb{E}_{a \sim \hat{\pi}, s' \sim \mathcal{P}(\cdot | s, a)} [r(s, a) + \gamma V^{\hat{\pi}}(s')] \\ &= \mathbb{E}_{a \sim \hat{\pi}, s' \sim \mathcal{P}(\cdot | s, a)} [\hat{Q}(s, a) - r(s, a) - \gamma \hat{V}(s')] + \gamma \mathbb{E}_{a \sim \hat{\pi}, s' \sim \mathcal{P}(\cdot | s, a)} [\hat{V}(s') - V^{\hat{\pi}}(s')] \\ &= \mathbb{E}_{\hat{\pi}} [\delta(s, a)] + \gamma \mathbb{E}_{a \sim \hat{\pi}, s' \sim \mathcal{P}(\cdot | s, a)} [\hat{V}(s') - V^{\hat{\pi}}(s')] \\ &= \mathbb{E}_{\hat{\pi}} [\delta(s, a)] + \dots \\ &= \mathbb{E}_{\hat{\pi}} \left[ \sum_{t=0}^{\infty} \gamma^t \delta(s_t, a_t) \mid s_0 = s \right], \end{aligned}$$

while the second term satisfies

$$\begin{aligned}
 V^{\pi^*}(s) - \widehat{V}(s) &= \mathbb{E}_{a \sim \pi^*, s' \sim \mathcal{P}(\cdot | s, a)} \left[ r(s, a) + \gamma V^{\widehat{\pi}}(s') \right] - \mathbb{E}_{a \sim \widehat{\pi}} \left[ \widehat{Q}(s, a) \right] \\
 &= \mathbb{E}_{a \sim \pi^*, s' \sim \mathcal{P}(\cdot | s, a)} \left[ r(s, a) + \gamma V^{\widehat{\pi}}(s') - \widehat{Q}(s, a) \right] + \mathbb{E}_{a \sim \pi^*} \left[ \widehat{Q}(s, a) \right] - \mathbb{E}_{a \sim \widehat{\pi}} \left[ \widehat{Q}(s, a) \right] \\
 &= \mathbb{E}_{a \sim \pi^*, s' \sim \mathcal{P}(\cdot | s, a)} \left[ r(s, a) + \gamma \widehat{V}(s') - \widehat{Q}(s, a) \right] + \gamma \mathbb{E}_{a \sim \pi^*, s' \sim \mathcal{P}(\cdot | s, a)} \left[ V^{\widehat{\pi}}(s') - \widehat{V}(s') \right] \\
 &\quad + \left\langle \widehat{Q}(s, \cdot), \pi^*(\cdot | s) - \widehat{\pi}(\cdot | s) \right\rangle_{\mathcal{A}} \\
 &= -\mathbb{E}_{a \sim \pi^*, s' \sim \mathcal{P}(\cdot | s, a)} [\delta(s, a)] + \left\langle \widehat{Q}(s, \cdot), \pi^*(\cdot | s) - \widehat{\pi}(\cdot | s) \right\rangle_{\mathcal{A}} + \dots \\
 &= -\mathbb{E}_{\pi^*} \left[ \sum_{t=0}^{\infty} \gamma^t \delta(s_t, a_t) \mid s_0 = s \right] + \mathbb{E}_{\pi^*} \left[ \sum_{t=0}^{\infty} \gamma^t \left\langle \widehat{Q}(s_t, \cdot), \pi^*(\cdot | s_t) - \widehat{\pi}(\cdot | s_t) \right\rangle_{\mathcal{A}} \mid s_0 = s \right].
 \end{aligned}$$

Combining the two equations above, we have the desired result.  $\square$

**Lemma B.2** ( $\xi$ -Quantifiers). *Let*

$$\lambda = 1, \quad \beta = c \cdot dV_{\max} \sqrt{\zeta}, \quad \zeta = \log(2dN/(1-\gamma)\xi). \quad (29)$$

Then  $\Gamma = \beta \cdot (\phi(s, a)^\top \Lambda^{-1} \phi(s, a))^{1/2}$  specified in Equation (18) are  $\xi$ -quantifiers. That is, with probability at least  $1 - \xi$ ,

$$|(\mathbb{B}\widehat{V})(s, a) - (\widehat{\mathbb{B}}\widehat{V})(s, a)| \leq \Gamma = \beta \sqrt{\phi(s, a)^\top \Lambda^{-1} \phi(s, a)}, \forall (s, a) \in \mathcal{S} \times \mathcal{A}. \quad (30)$$

*Proof.* we have

$$\begin{aligned}
 \mathbb{B}\widehat{V} - \widehat{\mathbb{B}}\widehat{V} &= \phi(s, a)^\top (w - \widehat{w}) \\
 &= \phi(s, a)^\top w - \phi(s, a)^\top \Lambda^{-1} \left( \sum_{\tau=1}^N \phi_\tau (r_\tau + \gamma \widehat{V}(s_{\tau+1})) \right) \\
 &= \underbrace{\phi(s, a)^\top w - \phi(s, a)^\top \Lambda^{-1} \left( \sum_{\tau=1}^N \phi_\tau \phi_\tau^\top w \right)}_{(i)} + \underbrace{\phi(s, a)^\top \Lambda^{-1} \left( \sum_{\tau=1}^N \phi_\tau \phi_\tau^\top w - \sum_{\tau=1}^N \phi_\tau (r_\tau + \gamma \widehat{V}(s_{\tau+1})) \right)}_{(ii)}, \quad (31)
 \end{aligned}$$

Then we bound (i) and (ii), respectively.

For (i), we have

$$\begin{aligned}
 (i) &= \phi(s, a)^\top w - \phi(s, a)^\top \Lambda^{-1} (\Lambda - \lambda I) w \\
 &= \lambda \phi(s, a)^\top \Lambda^{-1} w \\
 &\leq \lambda \|\phi(s, a)\|_{\lambda^{-1}} \|w\|_{\lambda^{-1}} \\
 &\leq V_{\max} \sqrt{d\lambda} \sqrt{\phi(s, a)^\top \Lambda^{-1} \phi(s, a)}, \quad (32)
 \end{aligned}$$

where the first inequality follows from Cauchy-Schwartz inequality. The second inequality follows from the fact that  $\|\Lambda^{-1}\|_{\text{op}} \leq \lambda^{-1}$  and Lemma B.3.

For notation simplicity, let  $\epsilon_\tau = r_\tau + \gamma \widehat{V}(s_{\tau+1}) - \phi_\tau^\top w$ , then we have

$$\begin{aligned}
 \text{(ii)} &= \phi(s, a) \Lambda^{-1} \sum_{\tau=1}^N \phi_\tau \epsilon_\tau \\
 &\leq \left\| \sum_{\tau=1}^N \phi_\tau \epsilon_\tau \right\|_{\Lambda^{-1}} \cdot \|\phi(s, a)\|_{\Lambda^{-1}} \\
 &= \underbrace{\left\| \sum_{\tau=1}^N \phi_\tau \epsilon_\tau \right\|_{\Lambda^{-1}}}_{\text{(iii)}} \cdot \sqrt{\phi(s, a)^\top \Lambda^{-1} \phi(s, a)}. \tag{33}
 \end{aligned}$$

The term (iii) is depend on the randomness of the data collection process of  $\mathcal{D}$ . To bound this term, we resort to uniform concentration inequalities to upper bound

$$\sup_{V \in \mathcal{V}(R, B, \lambda)} \left\| \sum_{\tau=1}^N \phi(x_\tau, a_\tau) \cdot \epsilon_\tau(V) \right\|,$$

where

$$\mathcal{V}(R, B, \lambda) = \{V(s; w, \beta, \Sigma) : \mathcal{S} \rightarrow [0, V_{\max}] \text{ with } \|w\| \leq R, \beta \in [0, B], \Sigma \succeq \lambda \cdot I\}, \tag{34}$$

where  $V(s; w, \beta, \Sigma) = \max_a \{\phi(s, a)^\top w - \beta \cdot \sqrt{\phi(s, a)^\top \Sigma^{-1} \phi(s, a)}\}$ . For all  $\epsilon > 0$ , let  $\mathcal{N}(\epsilon; R, B, \lambda)$  be the minimal cover if  $\mathcal{V}(R, B, \lambda)$ . That is, for any function  $V \in \mathcal{V}(R, B, \lambda)$ , there exists a function  $V^\dagger \in \mathcal{N}(\epsilon; R, B, \lambda)$ , such that

$$\sup_{s \in \mathcal{S}} |V(s) - V^\dagger(s)| \leq \epsilon. \tag{35}$$

Let  $R_0 = V_{\max} \sqrt{Nd/\lambda}$ ,  $B_0 = 2\beta$ , it is easy to show that at each iteration,  $\widehat{V}^u \in \mathcal{V}(R_0, B_0, \lambda)$ . From the definition of  $\mathbb{B}$ , we have

$$\|\mathbb{B}\widehat{V} - \mathbb{B}V^\dagger\| = \gamma \left| \int (\widehat{V}(s') - V^\dagger(s')) \langle \phi(s, a), \mu(s') \rangle ds' \right| \leq \gamma \epsilon. \tag{36}$$

Then we have

$$|(r + \gamma V - \mathbb{B}V) - (r + \gamma V^\dagger - \mathbb{B}V^\dagger)| \leq 2\gamma \epsilon. \tag{37}$$

Let  $\epsilon_\tau^\dagger = r(s_\tau, a_\tau) + \gamma V^\dagger(s_{\tau+1}) - \mathbb{B}V^\dagger(s, a)$ , we have

$$\begin{aligned}
 \text{(iii)}^2 &= \left\| \sum_{\tau=1}^N \phi_\tau \epsilon_\tau \right\|_{\Lambda^{-1}}^2 \leq 2 \left\| \sum_{\tau=1}^N \phi_\tau \epsilon_\tau^\dagger \right\|_{\Lambda^{-1}}^2 + 2 \left\| \sum_{\tau=1}^N \phi_\tau (\epsilon_\tau^\dagger - \epsilon_\tau) \right\|_{\Lambda^{-1}}^2 \\
 &\leq 2 \left\| \sum_{\tau=1}^N \phi_\tau \epsilon_\tau^\dagger \right\|_{\Lambda^{-1}}^2 + 8\gamma^2 \epsilon^2 \sum_{\tau=1}^N |\phi_\tau \Lambda^{-1} \phi_\tau| \\
 &\leq 2 \left\| \sum_{\tau=1}^N \phi_\tau \epsilon_\tau^\dagger \right\|_{\Lambda^{-1}}^2 + 8\gamma^2 \epsilon^2 N^2 / \lambda
 \end{aligned}$$

It remains to bound  $\left\| \sum_{\tau=1}^N \phi_\tau \epsilon_\tau^\dagger \right\|_{\Lambda^{-1}}^2$ . From the assumption for data collection process, it is easy to show that  $\mathbb{E}_{\mathcal{D}}[\epsilon_\tau | \mathcal{F}_{\tau-1}] = 0$ , where  $F_{\tau-1} = \sigma(\{(s_i, a_i)_{i=1}^\tau \cup (r_i, s_{i+1})_{i=1}^\tau\})$  is the  $\sigma$ -algebra generated by the variables from the first  $\tau$  step. Moreover, since  $\epsilon_\tau \leq 2V_{\max}$ , we have  $\epsilon_\tau$  are  $2V_{\max}$ -sub-Gaussian conditioning on  $F_{\tau-1}$ . Then we invoke Lemma B.6 with  $M_0 = \lambda \cdot I$  and  $M_k = \lambda \cdot I + \sum_{\tau=1}^k \phi(x_\tau, a_\tau) \phi(x_\tau, a_\tau)^\top$ . For the fixed function  $V : \mathcal{S} \rightarrow [0, V_{\max}]$ , we have

$$\mathbb{P}_{\mathcal{D}} \left( \left\| \sum_{\tau=1}^N \phi(x_\tau, a_\tau) \cdot \epsilon_\tau(V) \right\|_{\Lambda^{-1}}^2 > 8V_{\max}^2 \cdot \log \left( \frac{\det(\Lambda)^{1/2}}{\delta \cdot \det(\lambda \cdot I)^{1/2}} \right) \right) \leq \delta \tag{38}$$

for all  $\delta \in (0, 1)$ . Note that  $\|\phi(s, a)\| \leq 1$  for all  $(s, a) \in \mathcal{S} \times \mathcal{A}$  by Definition 2.1. We have

$$\|\Lambda\|_{\text{op}} = \left\| \lambda \cdot I + \sum_{\tau=1}^N \phi(s_\tau, a_\tau) \phi(s_\tau, a_\tau)^\top \right\|_{\text{op}} \leq \lambda + \sum_{\tau=1}^N \|\phi(s_\tau, a_\tau) \phi(s_\tau, a_\tau)^\top\|_{\text{op}} \leq \lambda + N,$$

where  $\|\cdot\|_{\text{op}}$  denotes the matrix operator norm. Hence, it holds that  $\det(\Lambda) \leq (\lambda + N)^d$  and  $\det(\lambda \cdot I) = \lambda^d$ , which implies

$$\begin{aligned} \mathbb{P}_{\mathcal{D}} \left( \left\| \sum_{\tau=1}^N \phi(s_{\tau}, a_{\tau}) \cdot \epsilon_{\tau}(V) \right\|_{\Lambda^{-1}}^2 > 4V_{\max}^2 \cdot (2 \cdot \log(1/\delta) + d \cdot \log(1 + N/\lambda)) \right) \\ \leq \mathbb{P}_{\mathcal{D}} \left( \left\| \sum_{\tau=1}^N \phi(s_{\tau}, a_{\tau}) \cdot \epsilon_{\tau}(V) \right\|_{\Lambda^{-1}}^2 > 8V_{\max}^2 \cdot \log \left( \frac{\det(\Lambda)^{1/2}}{\delta \cdot \det(\lambda \cdot I)^{1/2}} \right) \right) \leq \delta. \end{aligned}$$

Therefore, we conclude the proof of Lemma B.2.

Applying Lemma B.2 and the union bound, we have

$$\mathbb{P}_{\mathcal{D}} \left( \sup_{V \in \mathcal{N}(\varepsilon)} \left\| \sum_{\tau=1}^N \phi(x^{\tau}, a^{\tau}) \cdot \epsilon_{\tau}(V) \right\|_{\Lambda^{-1}}^2 > 4V_{\max}^2 \cdot (2 \cdot \log(1/\delta) + d \cdot \log(1 + N/\lambda)) \right) \leq \delta \cdot |\mathcal{N}(\varepsilon)|. \quad (39)$$

Recall that

$$\widehat{V} \in \mathcal{V}(R_0, B_0, \lambda), \quad \text{where } R_0 = V_{\max} \sqrt{Nd/\lambda}, \quad B_0 = 2\beta, \quad \lambda = 1, \quad \beta = c \cdot dV_{\max} \sqrt{\zeta}. \quad (40)$$

Here  $c > 0$  is an absolute constant,  $\xi \in (0, 1)$  is the confidence parameter, and  $\zeta = \log(2dV_{\max}/\xi)$  is specified in Algorithm 1. Applying Lemma B.5 with  $\varepsilon = dV_{\max}/N$ , we have

$$\begin{aligned} \log |\mathcal{N}(\varepsilon)| &\leq d \cdot \log(1 + 4d^{-1/2}N^{3/2}) + d^2 \cdot \log(1 + 32c^2 \cdot d^{1/2}N^2\zeta) \\ &\leq d \cdot \log(1 + 4d^{1/2}N^2) + d^2 \cdot \log(1 + 32c^2 \cdot d^{1/2}N^2\zeta). \end{aligned} \quad (41)$$

By setting  $\delta = \xi/|\mathcal{N}(\varepsilon)|$ , we have that with probability at least  $1 - \xi$ ,

$$\begin{aligned} \left\| \sum_{\tau=1}^N \phi(s_{\tau}, a_{\tau}) \cdot \epsilon_{\tau}(\widehat{V}) \right\|_{\Lambda^{-1}}^2 \\ \leq 8V_{\max}^2 \cdot (2 \cdot \log(V_{\max}/\xi) + 4d^2 \cdot \log(64c^2 \cdot d^{1/2}N^2\zeta) + d \cdot \log(1 + N) + 4d^2) \\ \leq 8V_{\max}^2 d^2 \zeta (4 + \log(64c^2)). \end{aligned} \quad (42)$$

Here the last inequality follows from simple algebraic inequalities. We set  $c \geq 1$  to be sufficiently large, which ensures that  $36 + 8 \cdot \log(64c^2) \leq c^2/4$  on the right-hand side of Equation (42). By Equations (33) and (42), it holds that

$$|\text{(ii)}| \leq c/2 \cdot dV_{\max} \sqrt{\zeta} \cdot \sqrt{\phi(x, a)^{\top} \Lambda^{-1} \phi(s, a)} = \beta/2 \cdot \sqrt{\phi(x, a)^{\top} \Lambda^{-1} \phi(s, a)} \quad (43)$$

By Equations (18), (31), (32), and (43), for all  $(s, a) \in \mathcal{S} \times \mathcal{A}$ , it holds that

$$|(\mathbb{B}\widehat{V})(x, a) - (\widehat{\mathbb{B}}\widehat{V})(s, a)| \leq (V_{\max} \sqrt{d} + \beta/2) \cdot \sqrt{\phi(s, a)^{\top} \Lambda^{-1} \phi(s, a)} \leq \Gamma(s, a) \quad (44)$$

with probability at least  $1 - \xi$ . Therefore, we conclude the proof of Lemma B.2.  $\square$

**Lemma B.3** (Bounded weight of value function). *Let  $V_{\max} = r_{\max}/(1 - \gamma)$ . For any function  $V : \mathcal{S} \rightarrow [0, V_{\max}]$ , we have*

$$\|w\| \leq V_{\max} \sqrt{d}, \quad \|\widehat{w}\| \leq V_{\max} \sqrt{\frac{Nd}{\lambda}}.$$

*Proof.* since

$$w^{\top} \phi(s, a) = \langle M, \phi(s, a) \rangle + \gamma \int V(s') \psi(s')^{\top} M \phi(s, a) ds',$$

We have

$$\begin{aligned}
 w &= M + \gamma \int V(s') \psi(s')^\top M ds' \\
 &= r_{\max} \sqrt{d} + \gamma V_{\max} \sqrt{d} \\
 &= V_{\max} \sqrt{d}.
 \end{aligned}$$

For  $\hat{w}$ , we have

$$\begin{aligned}
 \|\hat{w}\| &= \left\| \Lambda^{-1} \sum_{\tau=1}^N \phi_\tau (r_\tau + \gamma V(s_{\tau+1})) \right\| \\
 &\leq \sum_{\tau=1}^N \left\| \Lambda^{-1} \phi_\tau (r_\tau + \gamma V(s_{\tau+1})) \right\| \\
 &\leq V_{\max} \sum_{\tau=1}^N \left\| \Lambda^{-1} \phi_\tau \right\| \\
 &\leq V_{\max} \sum_{\tau=1}^N \sqrt{\phi_\tau^\top \Lambda^{-1/2} \Lambda^{-1} \Lambda^{-1/2} \phi_\tau} \\
 &\leq \frac{V_{\max}}{\sqrt{\lambda}} \sum_{\tau=1}^N \sqrt{\phi_\tau^\top \Lambda^{-1} \phi_\tau} \\
 &\leq V_{\max} \sqrt{\frac{N}{\lambda}} \sqrt{\text{Tr}(\Lambda^{-1} \sum_{\tau=1}^T \phi_\tau \phi_\tau^\top)} \\
 &\leq V_{\max} \sqrt{\frac{Nd}{\lambda}}.
 \end{aligned}$$

□

## B.2. Proof of Lemma 3.4

*Proof.* We consider the following iteration:

$$\begin{aligned}
 V_{\min} &\leftarrow \min_{s'} V(s'), \\
 Q(s, a) &\leftarrow r(s, a) + \gamma(1 - \varepsilon) \mathbb{E}_{s' \sim P_0} V(s') + \gamma \varepsilon V_{\min}, \\
 V(s) &\leftarrow \max_a Q(s, a).
 \end{aligned} \tag{45}$$

It is easy to see that if the iteration in (45) converges, it is the value function for the policies specified in Equation (11). In fact, the iteration above has a unique stationary solution follows from the fact that it is a  $\gamma$ -contraction for  $V(s)$ . Then it suffices to show that the solution to the value iteration with discount factor  $(1 - \varepsilon)\gamma$  is the same as the above stationary solution up to a constant. Let  $Q(s, a)$  and  $V(s)$  be the solution to the value iteration with discount factor  $(1 - \varepsilon)\gamma$ . Then we have

$$Q(s, a) = r(s, a) + (1 - \varepsilon)\gamma \mathbb{E}_{s'} V(s'),$$

Let  $\Delta = \gamma \varepsilon \min_s [\max_a Q(s, a)] / (1 - \gamma)$  and  $\tilde{Q}(\cdot, \cdot) = Q(\cdot, \cdot) + \Delta$ ,  $\tilde{V}(\cdot) = V(\cdot) + \Delta$ , then we have

$$\min_s [\max_a \tilde{Q}(s, a)] = \frac{(1 - \gamma + \gamma \varepsilon) \Delta}{\gamma \varepsilon}.$$



This leads to

$$\begin{aligned}
 & \tilde{Q}(s, a) \\
 &= r(s, a) + \gamma(1 - \varepsilon)\mathbb{E}_{s'} V(s') + \Delta \\
 &= r(s, a) + \gamma(1 - \varepsilon)\mathbb{E}_{s'} \tilde{V}(s') + (1 - \gamma + \gamma\varepsilon)\Delta \\
 &= r(s, a) + \gamma(1 - \varepsilon)\mathbb{E}_{s'} \tilde{V}(s') + \gamma\varepsilon \min_s [\max_a \tilde{Q}(s, a)].
 \end{aligned}$$

This means that  $\tilde{Q}$  is the unique stationary solution to the iteration in (45). Then we have the value function for policies in Equation (11) has the same value function as the policies in Equation (12) up to a constant. Then we finish the proof of Lemma 3.4.  $\square$

### B.3. Proof of Theorem 3.5

*Proof.* We first specialize the algorithm for offline value iteration with a lower discount factor and an estimated model, as depicted in Algorithm 3.

---

#### Algorithm 3 Generalized Value Iteration

---

- 1: **Require:** Dataset  $\mathcal{D}$ , discount factor  $\gamma$ , discount factor coefficient  $\varepsilon$
  - 2: Estimated the model by MLE:  $\hat{P} = \operatorname{argmax}_{P \in \mathcal{M}} \mathbb{E}_{\mathcal{D}} [\ln P(s' | s, a)]$ .
  - 3: Obtain the estimated Bellman operator  $\hat{\mathbb{B}}_{\gamma}$  from learned model  $\hat{P}$ .
  - 4: **while** not converge **do**
  - 5:   Set  $\hat{Q}(\cdot, \cdot) \leftarrow \min_{e \in \mathcal{E}(\varepsilon)} [(\hat{\mathbb{B}}_{(1-e)\gamma} \hat{V})(\cdot, \cdot)]$ , where  $\mathcal{E}(\varepsilon) = \{e(s, a) | \mathbb{E}_{\mathcal{D}}[e(s, a)] \leq \varepsilon\}$ .
  - 6:   Set  $\hat{\pi}(\cdot | \cdot) \leftarrow \operatorname{argmax}_{\pi} \mathbb{E}_{\pi} [\hat{Q}(\cdot, \cdot)]$ .
  - 7:   Set  $\hat{V}(\cdot) \leftarrow \mathbb{E}_{\hat{\pi}} [\hat{Q}(\cdot, \cdot)]$ .
  - 8: **end while**
  - 9: **Return**  $\hat{\pi}$
- 

Algorithm 3 is similar to the idea of a lower discount factor  $(1 - \varepsilon)\gamma$ , with only technical differences. It is easy to show that Algorithm 3 with discount factor  $\gamma$  and a lower yields the same policy as the following optimization problem

$$\operatorname{argmax}_{\pi \in \Pi} \operatorname{argmin}_{M \in \mathcal{M}_{\varepsilon}} V_{M, \gamma}(\pi),$$

where

$$\mathcal{M}_{\varepsilon} = \left\{ M \in \mathcal{M} \mid \exists \mathcal{P}(\cdot | s, a), e \in \mathcal{E}(\varepsilon), \mathcal{P}_M(\cdot | s, a) = (1 - e)\mathcal{P}(\cdot | s, a; \hat{M}) + e\mathcal{P}(\cdot | s, a), \forall (s, a) \in \mathcal{D} \right\}.$$

Here  $\hat{M}$  is the model obtained from MLE estimator and  $\mathcal{M}$  is the set of all linear models. The proof is similar to Lemma 3.4 and we omit it for simplicity.

Then we prove the theorem with the following steps.

#### 1. Bounding $\mathcal{M}_{\varepsilon}$ .

Let  $\mathcal{M}_{\text{TV}}(M_0, \varepsilon) = \left\{ M \mid \mathbb{E}_{\mathcal{D}} \left[ \text{D}_{\text{TV}}(\mathcal{P}(\cdot | s, a; M_0), \mathcal{P}(\cdot | s, a; M))^2 \right] \leq \varepsilon \right\}$ . To establish the equivalence between  $\mathcal{M}_{\varepsilon}$  and  $\mathcal{M}_{\text{TV}}(M_0, \varepsilon)$ , we need the follow assumption.

**Assumption B.4** (Regularity). We assume that the underlying linear MDP satisfies

$$\tilde{p} = \min\{p_{\min}, 1 - p_{\max}\} > 0,$$

where  $p_{\min} = \inf_{\mathcal{P}(s' | s, a) > 0} \mathcal{P}(s' | s, a)$ ,  $p_{\max} = \sup_{\mathcal{P}(s' | s, a) < 1} \mathcal{P}(s' | s, a)$ .

Note that this assumption is always true for tabular MDPs. It only rules out the case when there exists a sequence  $\{s'_n\}_{n=1}^\infty$  such that  $\lim_{n \rightarrow \infty} \mathcal{P}(s'_n | s, a) = 0$  or  $\lim_{n \rightarrow \infty} \mathcal{P}(s'_n | s, a) = 1$ . Then we have

$$\mathcal{M}_{\text{TV}}(\widehat{M}, \tilde{p}^2 \varepsilon^2 / 4) \subseteq \mathcal{M}_\varepsilon \subseteq \mathcal{M}_{\text{TV}}(\widehat{M}, \varepsilon^2), \quad (46)$$

Recall that  $\mathcal{M}_\varepsilon = \left\{ M | \mathcal{P}(\cdot | s, a; M) = (1 - \varepsilon)\mathcal{P}(\cdot | s, a; \widehat{M}) + \varepsilon\mathcal{P}(\cdot) \right\}$ . On the one hand, it is easy to see that the largest deviation in  $\mathcal{M}_\varepsilon$  from  $\widehat{M}$  happens when  $\mathcal{P}(\cdot | s, a; \widehat{M})$  is close to 0 or close to 1. On the other hand, we have  $D_{\text{TV}}((1 - \varepsilon)\mathcal{P}(\cdot | s, a; M) + \varepsilon\mathcal{P}(\cdot), \mathcal{P}(\cdot | s, a; M)) \leq \varepsilon$ . Then we have the result in Equation (46). The following steps largely follows (Uehara & Sun, 2021).

**2. Upper-bounding**  $\mathbb{E}_{(s,a) \sim \rho} [D_{\text{TV}}(\mathcal{P}(\cdot | s, a; M^*), \mathcal{P}(\cdot | s, a; \widehat{M}))^2]$ .

Let

$$\mathcal{M} = \left\{ \mathcal{P}(\cdot | s, a; M) \mid M \in \mathbb{R}^{d \times r}, \|M\|_2 \leq \sqrt{d}, \int \phi(s, a)^\top M \psi(s, a) d(s') = 1, \forall (s, a) \right\},$$

$$\text{and } \mathcal{H} = \left\{ \sqrt{\frac{\mathcal{P} + \mathcal{P}^*}{2}} \mid \mathcal{P} \in \mathcal{M} \right\}.$$

By invoking Theorem B.7, we first show

$$\mathbb{E}_{(s,a) \sim \rho} [D_{\text{TV}}(\mathcal{P}(\cdot | s, a; M^*), \mathcal{P}(\cdot | s, a; \widehat{M}))^2] \leq c \{ (d/N) \ln^2(Nd) + \ln(c/\delta)/N \}.$$

To do that, we calculate the entropy integral with bracketing. First, we have

$$\mathcal{N}_{[]}(\varepsilon, \mathcal{H}, d) \leq \mathcal{N}_{[]}(\varepsilon, \mathcal{M}, d'). \quad (47)$$

where

$$d'(a, b) = \mathbb{E}_{(s,a) \sim \rho} \left[ \int (a(s, a, s') - b(s, a, s'))^2 d(s') \right]^{1/2}, \quad (48)$$

$$d(a, b) = \mathbb{E}_{(s,a) \sim \rho} \left[ \int (\sqrt{a(s, a, s')} - \sqrt{b(s, a, s')})^2 d(s') \right]^{1/2}. \quad (49)$$

Here, we use two observations. The first observation is

$$d^2 \left( \sqrt{\frac{\mathcal{P}(M') + \mathcal{P}^*}{2}}, \sqrt{\frac{\mathcal{P}(M'') + \mathcal{P}^*}{2}} \right) \leq c_1 d'^2(\mathcal{P}(M'), \mathcal{P}(M''))$$

due to the mean-value theorem

$$\sqrt{a} - \sqrt{b} \leq \max(1/\sqrt{a}, 1/\sqrt{b})(a - b)$$

and Assumption B.4 that  $\mathcal{P}^*(s' | s, a) \geq c_0 > 0$ . The second observation is when we have  $P' < g < P''$ , we also have  $\sqrt{(P' + P^*)/2} < \sqrt{(g + P^*)/2} < \sqrt{(P'' + P^*)/2}$ . Then, 47 is concluded.

Next, by letting  $M^{(1)}, \dots, M^{(K)}$  be an  $\varepsilon$ -cover of the  $d$ -dimensional ball with a radius  $\sqrt{d}$ , i.e.,  $B_d(\sqrt{d})$ , we have the brackets  $\{\mathcal{P}(M^{(i)}) - \varepsilon, \mathcal{P}(M^{(i)}) + \varepsilon\}_{i=1}^K$  which cover  $\mathcal{M}$ . This is because for any  $\mathcal{P}(M) \in \mathcal{M}$ , we can take  $M^{(i)}$  s.t.  $\|M - M^{(i)}\|_2 \leq \varepsilon/\sqrt{d}$ , then,

$$\mathcal{P}(M^{(i)}) - \varepsilon < \mathcal{P}(M) < \mathcal{P}(M^{(i)}) + \varepsilon, \quad \forall (s, a, s')$$

noting

$$|\mathcal{P}(M)(s, a, s') - \mathcal{P}(M^{(i)})(s, a, s')| \leq \sqrt{d} \|M - M^{(i)}\|_2 \leq \varepsilon, \quad \forall (s, a, s') \quad (50)$$

The last equality follows from the fact that  $\|M\|_2 \leq \sqrt{d}$  and  $\|\psi\|_2 \leq 1$ . Therefore, we have

$$\mathcal{N}_{[]}(\epsilon, \mathcal{M}, \|\cdot\|_2) \leq \mathcal{N}(\epsilon/\sqrt{d}, B_d(c\sqrt{d}), \|\cdot\|_2),$$

where  $\mathcal{N}(\epsilon/\sqrt{d}, B_d(c\sqrt{d}), \|\cdot\|_2)$  is a covering number of  $B_d(c\sqrt{d})$  w.r.t  $\|\cdot\|_2$ . This is upper-bounded by  $(c\sqrt{d}/\epsilon)^d$  (Wainwright, 2019, Lemma 5.7). Thus, we can calculate the upper bound of the entropy integral  $J_B(\delta, \mathcal{M}, \|\cdot\|_2)$ :

$$\begin{aligned} \int_0^\delta d^{1/2} \ln^{1/2}(cd/u) du &\leq \int_0^\delta d^{1/2} \ln(1/u) du + \delta d^{1/2} \ln(c_1\sqrt{d}) \\ &= cd^{1/2}(\delta + \delta \ln(1/\delta)) + \delta d^{1/2} \ln(cd) \\ &\leq cd^{1/2} \delta \ln(cd/\delta). \end{aligned}$$

By taking  $G(x) = d^{1/2}x \ln(c\sqrt{d}/x)$  in B.7,  $\delta_n = (d/n)^{1/2} \ln(nd)$  satisfies the critical inequality

$$\sqrt{n}\delta_n^2 \geq d^{1/2}\delta_n \ln(cd/\delta_n).$$

Finally, with probability  $1 - \delta$

$$\mathbb{E}_{(s,a)\sim\rho}[\text{D}_{\text{TV}}(P(\cdot | s, a; M^*), P(\cdot | s, a; \widehat{M}))^2] \leq \xi', \quad \xi' := \{(d/N) \ln^2(Nd) + \ln(c/\delta)/N\}. \quad (51)$$

Hereafter, we condition on this event.

**3. Upper bounding**  $\mathbb{E}_{\mathcal{D}}[\|P(\cdot | s, a; M^*) - P(\cdot | s, a; \widehat{M})\|_1^2]$ .

We take an  $\epsilon$ -cover of the ball  $B_d(R)$  in terms of  $\|\cdot\|_2$ , i.e.,  $\widehat{\mathcal{M}} = \{M^{(1)}, \dots, M^{(K)}\}$ , where  $K = (cR/\epsilon)^d$ . Then, for any  $M \in B_d(R)$ , there exists  $M^{(i)}$  s.t.  $\forall (s, a) \in \mathcal{S} \times \mathcal{A}$ ,

$$\begin{aligned} &\left| \|P(\cdot | s, a; M) - P(\cdot | s, a; M^*)\|_1^2 - \|P(\cdot | s, a; M^{(i)}) - P(\theta^*)(\cdot | s, a)\|_1^2 \right| \\ &\leq 4\|P(\cdot | s, a; M) - P(\cdot | s, a; M^*)\|_1 - \|P(\cdot | s, a; M^{(i)}) - P(\cdot | s, a; M^*)\|_1 \quad (a^2 - b^2 = (a-b)(a+b)) \\ &\leq 4\|P(\cdot | s, a; M) - P(\cdot | s, a; M^{(i)})\|_1 \quad (\|a\| - \|b\| \leq \|a-b\|) \\ &\leq 4\|M - M^{(i)}\|_2 \quad (\text{From (50)}) \\ &\leq 4\epsilon. \quad (52) \end{aligned}$$

Here, note for  $\widehat{M}$ , we have  $M^{(i)}$  s.t.  $\|\widehat{M} - M^{(i)}\|_2 \leq \epsilon$ . Then, we have

$$\begin{aligned} &\mathbb{E}_{\mathcal{D}}\|P(\cdot | s, a; M^*) - P(\cdot | s, a; \widehat{M})\|_1^2 \\ &\lesssim \mathbb{E}_{\mathcal{D}}\|P(\cdot | s, a; M^*) - P(\cdot | s, a; M^{(i)})\|_1^2 + \epsilon \quad (\text{From 52}) \\ &\lesssim (\mathbb{E}_{\mathcal{D}} - \mathbb{E}_{(s,a)\sim\rho})\|P(\cdot | s, a; M^*) - P(\cdot | s, a; M^{(i)})\|_1^2 + \epsilon + \mathbb{E}_{(s,a)\sim\rho}\|P(\cdot | s, a; M^*) - P(\cdot | s, a; M^{(i)})\|_1^2 \\ &\lesssim \sqrt{\frac{\text{var}_{(s,a)\sim\rho}[\|P(\cdot | s, a; M^*) - P(\cdot | s, a; M^{(i)})\|_1^2] \ln(K/\delta)}{n}} + \frac{\ln(K/\delta)}{n} \\ &+ \epsilon + \mathbb{E}_{(s,a)\sim\rho}\|P(\cdot | s, a; M^*) - P(\cdot | s, a; M^{(i)})\|_1^2 \quad (\text{Bernstein inequality}) \\ &\lesssim \sqrt{\frac{\mathbb{E}_{(s,a)\sim\rho}[\|P(\cdot | s, a; M^*) - P(\cdot | s, a; M^{(i)})\|_1^2] \ln(K/\delta)}{n}} + \frac{\ln(K/\delta)}{n} \\ &+ \epsilon + \mathbb{E}_{(s,a)\sim\rho}\|P(\cdot | s, a; M^*) - P(\cdot | s, a; M^{(i)})\|_1^2. \quad (\|P(\cdot | s, a; M^*) - P(\cdot | s, a; M^{(i)})\|_1^2 \leq 4) \end{aligned}$$

Then,

$$\begin{aligned} &\mathbb{E}_{\mathcal{D}}\|P(\cdot | s, a; M^*) - P(\cdot | s, a; \widehat{M})\|_1^2 \\ &\lesssim \sqrt{\frac{\{\mathbb{E}_{(s,a)\sim\rho}[\|P(\cdot | s, a; M^*) - P(\cdot | s, a; \widehat{M})\|_1^2] + \epsilon\} \ln(K/\delta)}{n}} + \frac{\ln(K/\delta)}{n} \\ &+ \epsilon + \mathbb{E}_{(s,a)\sim\rho}\|P(\cdot | s, a; M^*) - P(\cdot | s, a; \widehat{M})\|_1^2 \\ &\lesssim \sqrt{\frac{\{\xi' + \epsilon\} \ln(K/\delta)}{n}} + \frac{\ln(K/\delta)}{n} + \epsilon + \xi'. \quad (\text{From (51)}) \end{aligned}$$

In the end, by taking  $\epsilon = 1/n$ , we have with probability  $1 - \delta$ ,

$$\mathbb{E}_{\mathcal{D}} \|P(\cdot | s, a; M^*) - P(\cdot | s, a; \widehat{M})\|_1^2 \leq \xi, \quad \xi = c\{(d/n) \ln^2(nR) + \ln(c/\delta)/n\}.$$

This implies with probability  $1 - \delta$ ,  $P^* \in \mathcal{M}_{\text{TV}}(\widehat{M}, \xi)$ . Note that  $\mathcal{M}_{\text{TV}}(\widehat{M}, \tilde{p}^2 \epsilon^2/4) \subseteq \mathcal{M}_\epsilon$ , this also implies that with probability  $1 - \delta$ ,  $P(\cdot | s, a; M^*) \in \mathcal{M}_\epsilon$ , where  $\epsilon = 2\sqrt{\xi/\tilde{p}}$ .

**4. Show**  $\mathbb{E}_{(s,a)\sim\rho} [\text{D}_{\text{TV}}(\mathcal{P}(\cdot | s, a; M^*), \mathcal{P}(\cdot | s, a; M))^2] \lesssim \xi, \forall \mathcal{P}(M) \in \mathcal{M}_\epsilon$ .

We show for any  $\mathcal{P} \in \mathcal{M}_\epsilon$ , the distance between  $P^*$  is controlled in terms of TV distance. Formally, we need

$$\mathbb{E}_{(s,a)\sim\rho} [\text{D}_{\text{TV}}(\mathcal{P}(\cdot | s, a; M^*), \mathcal{P}(\cdot | s, a; M))^2] \lesssim \xi, \quad \forall \mathcal{P}(M) \in \mathcal{M}_\epsilon.$$

Since  $\mathcal{M}_\epsilon \subseteq \mathcal{M}_{\text{TV}}(\widehat{M}, \epsilon^2)$ , it suffices to show that

$$\mathbb{E}_{(s,a)\sim\rho} [\text{D}_{\text{TV}}(\mathcal{P}(\cdot | s, a; M^*), \mathcal{P}(\cdot | s, a; M))^2] \lesssim \xi, \quad \forall \mathcal{P}(M) \in \mathcal{M}_{\text{TV}}(\widehat{M}, 2\epsilon).$$

For any  $P \in \mathcal{M}_{\text{TV}}(\widehat{M}, \epsilon^2)$ , we have

$$\begin{aligned} & \mathbb{E}_{\mathcal{D}} [\text{D}_{\text{TV}}(P(\cdot | s, a), P^*(\cdot | s, a))^2] \\ & \leq 2\mathbb{E}_{\mathcal{D}} [\text{D}_{\text{TV}}(\widehat{P}(\cdot | s, a), P(\cdot | s, a))^2] + 2\mathbb{E}_{\mathcal{D}} [\text{D}_{\text{TV}}(\widehat{P}(\cdot | s, a), P^*(\cdot | s, a))^2] \leq 16\xi/\tilde{p}^2. \end{aligned}$$

Thus, we have

$$\begin{aligned} & \mathbb{E}_{s,a\sim\rho} [\text{D}_{\text{TV}}(P(\cdot | s, a), P^*(\cdot | s, a))^2] \\ & = \mathbb{E}_{s,a\sim\rho} [\text{D}_{\text{TV}}(P(\cdot | s, a), P^*(\cdot | s, a))^2] - \mathbb{E}_{\mathcal{D}} [\text{D}_{\text{TV}}(P(\cdot | s, a), P^*(\cdot | s, a))^2] + \mathbb{E}_{\mathcal{D}} [\text{D}_{\text{TV}}(P(\cdot | s, a), P^*(\cdot | s, a))^2] \\ & \leq A(M) + c\xi, \end{aligned} \tag{53}$$

where  $A(M) := |(\mathbb{E}_{\mathcal{D}} - \mathbb{E}_{(s,a)\sim\rho}) [\text{D}_{\text{TV}}(\mathcal{P}(\cdot | s, a; M^*), \mathcal{P}(\cdot | s, a; M))^2]|$ .

We again consider an  $\epsilon/\sqrt{d}$ -cover of the ball  $B_d(\sqrt{d})$  in terms of  $\|\cdot\|_2$ , i.e.,  $M' = \{M^{(1)}, \dots, M^{(K)}\}$ , where  $K = (c_1 d/\epsilon)^d$  ( $\epsilon = 1/N$ ). Then  $M'$  is also an  $\epsilon/\sqrt{d}$ -cover for  $\mathcal{M}_{\text{TV}}(\widehat{M}, \epsilon^2)$ . That is for any  $M$  s.t.  $\forall \mathcal{P}(M) \in \mathcal{M}_{\text{TV}}(\widehat{M}, \epsilon^2)$ , we can take  $M' \in \mathcal{M}'$  s.t.  $\|M - M'\|_2 \leq \epsilon/\sqrt{d}$ .

Then, we have

$$\mathbb{E}_{(s,a)\sim\rho} [\text{D}_{\text{TV}}(\mathcal{P}(\cdot | s, a; M^*), \mathcal{P}(\cdot | s, a; M))^2] \leq A(M) + c\xi, \quad \forall M \in \mathcal{M}'. \tag{54}$$

This is because for any  $M^{(i)} \in \mathcal{M}'$ , we can take  $P(M) \in \mathcal{M}_{\text{TV}}(\widehat{M}, \epsilon^2)$  such that

$$\begin{aligned} & \mathbb{E}_{(s,a)\sim\rho} [\text{D}_{\text{TV}}(\mathcal{P}(\cdot | s, a; M^*), \mathcal{P}(M^{(i)})(\cdot | s, a))^2] \\ & \leq \mathbb{E}_{(s,a)\sim\rho} [\text{D}_{\text{TV}}(\mathcal{P}(\cdot | s, a; M^*), \mathcal{P}(M^{(i)})(\cdot | s, a))^2] - \text{D}_{\text{TV}}(\mathcal{P}(\cdot | s, a; M^*), P(\cdot | s, a; M))^2 \\ & \quad + \mathbb{E}_{(s,a)\sim\rho} [\text{D}_{\text{TV}}(\mathcal{P}(\cdot | s, a; M^*), \mathcal{P}(\cdot | s, a; M))^2] \\ & \leq 4\epsilon + \mathbb{E}_{(s,a)\sim\rho} [\text{D}_{\text{TV}}(\mathcal{P}(\cdot | s, a; M^*), \mathcal{P}(\cdot | s, a; M))^2] \\ & \lesssim A(M) + \xi. \end{aligned}$$

From Bernstein's inequality, we have that with probability  $1 - \delta$ ,

$$\begin{aligned} A(M) & = |(\mathbb{E}_{\mathcal{D}} - \mathbb{E}_{(s,a)\sim\rho}) [\text{D}_{\text{TV}}(\mathcal{P}(\cdot | s, a; M), \mathcal{P}(\cdot | s, a; M^*))^2]| \\ & \lesssim \sqrt{\frac{\text{var}_{(s,a)\sim\rho} [\text{D}_{\text{TV}}(\mathcal{P}(\cdot | s, a; M), \mathcal{P}(\cdot | s, a; M^*))^2] \ln(K/\delta)}{N}} + \frac{\ln(K/\delta)}{N}, \quad \forall M \in \mathcal{M}'. \end{aligned} \tag{55}$$

Based on the construction of  $\mathcal{M}'$  and Equation (54), we have

$$\text{var}_{(s,a)\sim\rho} [\text{D}_{\text{TV}}(\mathcal{P}(\cdot | s, a; M^*), \mathcal{P}(\cdot | s, a; M))^2] \lesssim A(M) + \xi, \quad \forall M \in \mathcal{M}'. \tag{56}$$

Taking Equation (56) into Equation (55), we have  $A(M)$  satisfying

$$A^2(M) - A(M)B_1 - B_2 \leq 0, \quad B_1 = \frac{\ln(K/\delta)}{N}, B_2 = \xi \frac{\ln(K/\delta)}{N} + \left( \frac{\ln(K/\delta)}{N} \right)^2.$$

Then, we have

$$A(M) \leq \frac{\ln(K/\delta)}{N} + \xi^{1/2} \sqrt{\frac{\ln(K/\delta)}{N}} \lesssim \xi, \quad \forall M \in \mathcal{M}'. \quad (57)$$

We combine all steps. Recall for any  $\forall P(M) \in \mathcal{M}_\varepsilon$ , we can take  $M' \in \mathcal{M}'$  s.t.  $\|M - M'\|_2 \leq 1/n$ . Then, for any  $P(M) \in \mathcal{M}_\varepsilon$ , we have

$$\begin{aligned} A(M) &= |(\mathbb{E}_{\mathcal{D}} - \mathbb{E}_{(s,a) \sim \rho}) [\text{D}_{\text{TV}}(\mathcal{P}(\cdot | s, a; M), \mathcal{P}(\cdot | s, a; M^*))^2] \\ &\leq |(\mathbb{E}_{\mathcal{D}} - \mathbb{E}_{(s,a) \sim \rho}) [\text{D}_{\text{TV}}(\mathcal{P}(\cdot | s, a; M), \mathcal{P}(\cdot | s, a; M^*))^2 - \text{D}_{\text{TV}}(\mathcal{P}(\cdot | s, a; M'), \mathcal{P}(\cdot | s, a; M^*))^2] \\ &\quad + |(\mathbb{E}_{\mathcal{D}} - \mathbb{E}_{(s,a) \sim \rho}) [\text{D}_{\text{TV}}(\mathcal{P}(\cdot | s, a; M'), \mathcal{P}(\cdot | s, a; M^*))^2] \\ &\lesssim 8\varepsilon + |(\mathbb{E}_{\mathcal{D}} - \mathbb{E}_{(s,a) \sim \rho}) [\text{D}_{\text{TV}}(\mathcal{P}(\cdot | s, a; M'), \mathcal{P}(\cdot | s, a; M^*))^2] \\ &\lesssim \xi. \end{aligned} \quad (\text{From (57) and } M' \in \mathcal{M}')$$

Then, we have with probability  $1 - \delta$ ,

$$A(M) \lesssim \xi, \quad \forall P(M) \in \mathcal{M}_\varepsilon. \quad (58)$$

Finally, for any  $P(M) \in \mathcal{M}_\varepsilon$ , with probability  $1 - \delta$ , we have

$$\begin{aligned} \mathbb{E}_{(s,a) \sim \rho} [\text{D}_{\text{TV}}(P(\cdot | s, a; M^*), P(\cdot | s, a; M))^2] &\leq A(M) + c\xi && (\text{From (54)}) \\ &\lesssim \xi. && (\text{From (58)}) \end{aligned}$$

## 5. Bounding the performance of $\pi^*$ .

We first prove

$$V_{P^*}^{\pi^*} - V_P^{\pi^*} \lesssim (1 - \gamma)^{-2} \sqrt{c^\ddagger d \xi} \cdot r_{\max}, \quad (59)$$

for all  $P \in \mathcal{M}_\varepsilon$ . Recall from the third step, for  $P(M) \in \mathcal{M}_\varepsilon$ , we have

$$\mathbb{E}_{(s,a) \sim \rho} [\text{D}_{\text{TV}}(P(\cdot | s, a; M^*), P(\cdot | s, a; M))^2] \lesssim \xi.$$

From the second statement of Lemma B.8,

$$\forall V : \mathcal{S} \rightarrow [0, V_{\max}], \quad (M - M^*)^\top \Sigma_{\rho, V} (M - M^*) \lesssim V_{\max}^2 \xi, \quad \Sigma_{\rho, V} = \mathbb{E}_{(s,a) \sim \rho} [\psi_V(s, a) \psi_V^\top(s, a)].$$

Here, we have

$$\begin{aligned} V_{P^*}^{\pi^*} - V_P^{\pi^*} &\leq (1 - \gamma)^{-1} \left| \mathbb{E}_{(s,a) \sim d^{\pi^*}} \left[ \int \{P(s' | s, a) - P^*(s' | s, a)\} V_P^{\pi^*}(s') d(s') \right] \right| && (\text{Simulation lemma}) \\ &\leq (1 - \gamma)^{-1} \left| \mathbb{E}_{(s,a) \sim d^{\pi^*}} \left[ (M - M^*) \psi_{V_P^{\pi^*}}(s, a) \right] \right| \\ &\leq (1 - \gamma)^{-1} \underbrace{\|M - M^*\|_{\lambda I + \Sigma_{\rho, V_P^{\pi^*}}}}_{(a)} \underbrace{\mathbb{E}_{(s,a) \sim d^{\pi^*}} \left[ \|\psi_{V_P^{\pi^*}}(s, a)\|_{(\Sigma_{\rho, V_P^{\pi^*}} + \lambda I)^{-1}} \right]}_{(b)}. && (\text{C-S inequality}) \end{aligned}$$



The first term (a) is upper-bounded by  $\sqrt{V_{\max}^2 \xi + \lambda d}$  noting  $0 \leq V_P^{\pi^*} \leq V_{\max}$ . The term (b) is upper-bounded by

$$\begin{aligned}
 \mathbb{E}_{(s,a) \sim d^{\pi^*}} \left[ \|\psi_{V_P^{\pi^*}}(s, a)\|_{(\Sigma_{\rho, V_P^{\pi^*}} + \lambda I)^{-1}} \right] &\leq \mathbb{E}_{(s,a) \sim d^{\pi^*}} \left[ \|\psi_{V_P^{\pi^*}}(s, a)\|_{(\Sigma_{\rho, V_P^{\pi^*}} + \lambda I)^{-1}}^2 \right]^{1/2} && \text{(Jensen's inequality)} \\
 &= \sqrt{\text{Tr}(\Sigma_{d^{\pi^*}, V_P^{\pi^*}} (\lambda I + \Sigma_{\rho, V_P^{\pi^*}})^{-1})} \\
 &\leq \sqrt{c_V \text{Tr}(\Sigma_{\rho, V_P^{\pi^*}} (\lambda I + \Sigma_{\rho, V_P^{\pi^*}})^{-1})} \\
 &\leq \sqrt{c_V \text{rank}(\Sigma_{\rho, V_P^{\pi^*}})} \\
 &\leq \sqrt{c_V d} = \sqrt{c^\dagger d},
 \end{aligned}$$

where  $c_V = \sup_{V \in \{S \rightarrow [0, V_{\max}]\}} \sup_x \frac{x^\top \Sigma_V(d^{\pi^*})x}{x^\top \Sigma_V(\rho)x}$ ,  $c^\dagger = \sup_x \frac{x^\top \Sigma(d^{\pi^*})x}{x^\top \Sigma(\rho)x}$ , and

$$\Sigma_V(\mu) = \mathbb{E}_{(s,a) \sim \mu} [\psi_V(s, a) \psi_V^\top(s, a)], \quad \Sigma(\mu) = \mathbb{E}_{(s,a) \sim \mu} [\phi(s, a) \phi^\top(s, a)].$$

The second inequality follows from the definition of  $c_V$ , and the last inequality follows from the third statement in Lemma B.8. By taking  $\lambda$  s.t.  $\lambda d \lesssim V_{\max} \xi$ , We have the desired Equation (59).

Finally, combining all things together, with probability  $1 - 2\delta$ , for any  $\pi^* \in \Pi$ , we have

$$\begin{aligned}
 V_{P^*}^{\pi^*} - V_{P^*}^{\widehat{\pi}} &\leq V_{P^*}^{\pi^*} - \min_{P \in \mathcal{M}_\varepsilon} V_P^{\pi^*} + \min_{P \in \mathcal{M}_\varepsilon} V_P^{\pi^*} - V_{P^*}^{\widehat{\pi}} \\
 &\leq V_{P^*}^{\pi^*} - \min_{P \in \mathcal{M}_\varepsilon} V_P^{\pi^*} + \min_{P \in \mathcal{M}_\varepsilon} V_P^{\widehat{\pi}} - V_{P^*}^{\widehat{\pi}} && \text{(definition of } \widehat{\pi}) \\
 &\leq V_{P^*}^{\pi^*} - \min_{P \in \mathcal{M}_\varepsilon} V_P^{\pi^*} && \text{(Second step, } P^* \in \mathcal{M}_\varepsilon) \\
 &\lesssim (1 - \gamma)^{-2} \sqrt{c^\dagger d \xi} \cdot r_{\max}. && \text{(From (59))}
 \end{aligned}$$

Finally, recall from the relationship in the second and third step, we have  $\varepsilon = \sqrt{\xi}/2\tilde{p}$ ,  $\xi = c\{(d/N) \ln^2(Nd) + \ln(c/\delta)/N\}$ , which leads to

$$V_{P^*}^{\pi^*} - V_{P^*}^{\widehat{\pi}} \lesssim \frac{c_3}{(1 - \gamma)^{-2}} \sqrt{c^\dagger d^2 \zeta / N} \cdot r_{\max}, \quad \zeta = \log^2(c_2 N d / \delta).$$

□

## B.4. Technical Lemmas

**Lemma B.5** ( $\varepsilon$ -Covering Number (Jin et al., 2020)). *For all  $h \in [H]$  and all  $\varepsilon > 0$ , we have*

$$\log |\mathcal{N}(\varepsilon; R, B, \lambda)| \leq d \cdot \log(1 + 4R/\varepsilon) + d^2 \cdot \log(1 + 8d^{1/2} B^2 / (\varepsilon^2 \lambda)).$$

*Proof of Lemma B.5.* See Lemma D.6 in (Jin et al., 2020) for a detailed proof. □

**Lemma B.6** (Concentration of Self-Normalized Processes (Abbasi-Yadkori et al., 2011)). *Let  $\{\mathcal{F}_t\}_{t=0}^\infty$  be a filtration and  $\{\epsilon_t\}_{t=1}^\infty$  be an  $\mathbb{R}$ -valued stochastic process such that  $\epsilon_t$  is  $\mathcal{F}_t$ -measurable for all  $t \geq 1$ . Moreover, suppose that conditioning on  $\mathcal{F}_{t-1}$ ,  $\epsilon_t$  is a zero-mean and  $\sigma$ -sub-Gaussian random variable for all  $t \geq 1$ , that is,*

$$\mathbb{E}[\epsilon_t | \mathcal{F}_{t-1}] = 0, \quad \mathbb{E}[\exp(\lambda \epsilon_t) | \mathcal{F}_{t-1}] \leq \exp(\lambda^2 \sigma^2 / 2), \quad \forall \lambda \in \mathbb{R}.$$

*Meanwhile, let  $\{\phi_t\}_{t=1}^\infty$  be an  $\mathbb{R}^d$ -valued stochastic process such that  $\phi_t$  is  $\mathcal{F}_{t-1}$ -measurable for all  $t \geq 1$ . Also, let  $M_0 \in \mathbb{R}^{d \times d}$  be a deterministic positive-definite matrix and*

$$M_t = M_0 + \sum_{s=1}^t \phi_s \phi_s^\top$$

for all  $t \geq 1$ . For all  $\delta > 0$ , it holds that

$$\left\| \sum_{s=1}^t \phi_s \epsilon_s \right\|_{M_t^{-1}}^2 \leq 2\sigma^2 \cdot \log \left( \frac{\det(M_t)^{1/2} \cdot \det(M_0)^{-1/2}}{\delta} \right)$$

for all  $t \geq 1$  with probability at least  $1 - \delta$ .

*Proof.* See Theorem 1 of (Abbasi-Yadkori et al., 2011) for a detailed proof.  $\square$

Given a function class  $\mathcal{F}$ , let  $\mathcal{N}_{[]}(\delta, \mathcal{F}, d)$  be the bracketing number of  $\mathcal{F}$  w.r.t the metric  $d(a, b)$  given by

$$d(a, b) = \mathbb{E}_{(s,a) \sim \rho} \left[ \int (a(s' | s, a) - b(s' | s, a))^2 d(s') \right]^{1/2}.$$

Then, the entropy integral of  $\mathcal{F}$  is given by

$$J_B(\delta, \mathcal{F}, d) = \max \left( \int_{\delta^2/2}^{\delta} (\log \mathcal{N}_{[]}(\delta, \mathcal{F}, d))^{1/2} du, \delta \right). \quad (60)$$

We also define the localized class of  $\mathcal{H}$ :

$$\mathcal{H}(\delta) = \{h \in \mathcal{H} : \mathbb{E}_{(s,a) \sim \rho} [h^2(P(\cdot | s, a) \| P^*(\cdot | s, a))] \leq \delta^2\},$$

where  $h(P(\cdot | s, a) \| P^*(\cdot | s, a))$  denotes Hellinger distance defined by

$$\left( 0.5 \int \{ \sqrt{P(s' | s, a)} - \sqrt{P^*(s' | s, a)} \}^2 d(s') \right)^{1/2}.$$

**Theorem B.7** (MLE guarantee with general function approximation, Uehara & Sun (2021)). *We take a function  $G(\epsilon) : [0, 1] \rightarrow \mathbb{R}$  s.t.  $G(\epsilon) \geq J_B[\epsilon, \mathcal{H}(\epsilon), d]$  and  $G(\epsilon)/\epsilon^2$  is a non-increasing function w.r.t  $\epsilon$ . Then, letting  $\xi_n$  be a solution to  $\sqrt{n}\epsilon^2 \geq cG(\epsilon)$  w.r.t  $\epsilon$ . With probability  $1 - \delta$ , we have*

$$\mathbb{E}_{(s,a) \sim \rho} [\| \hat{P}_{\text{MLE}}(\cdot | s, a) - P(\cdot | s, a) \|_1^2] \leq c_1 \left\{ \xi_n + \sqrt{\log(c_2/\delta)/n} \right\}^2.$$

*Proof.* The proof follows directly by adapting to conditional distribution from Theorem 7.4 in (Geer et al., 2000). Please refer to (Geer et al., 2000) for more details.  $\square$

**Lemma B.8** (Property of linear MDPs). *Let  $P(M) = P(s' | s, a; M) = \phi(s, a)^\top M \psi(s')$ . Suppose  $P(M) \in \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$ . For any function  $V \in \mathcal{S} \rightarrow [0, V_{\max}]$ , letting  $\psi_V(s, a) = \int \text{vec}(\phi(s, a) \psi(s')^\top) V(s') d(s')$ , we suppose  $\|\phi(s, a)\|_2 \leq 1$  and  $\|\psi(s')\|_2 \leq 1$ . The following theorems hold:*

1. For any  $(s, a, s')$ , we have  $|P(M)(s, a, s') - P(M')(s, a, s')| \leq \|M - M'\|_2$ .
2. For any  $(s, a)$ , we have  $\text{TV}(P(M)(s, a, \cdot), P(M')(s, a, \cdot)) \leq \|M - M'\|_2$ . Besides, for any  $V : \mathcal{S} \rightarrow [0, 1]$ , we have

$$|(M - M')\psi_V(s, a)| \leq V_{\max} \text{TV}(P(M)(s, a, \cdot), P(M')(s, a, \cdot)).$$

3.

$$\sup_{V \in \{\mathcal{S} \rightarrow [0, V_{\max}]\}} \sup_x \frac{x^\top \mathbb{E}_{(s,a) \sim d^{\pi^*}} [\psi_V(s, a) \psi_V^\top(s, a)] x}{x^\top \mathbb{E}_{(s,a) \sim \rho} [\psi_V(s, a) \psi_V^\top(s, a)] x} = \sup_x \frac{x^\top \mathbb{E}_{d^{\pi^*}} [\phi(s, a) \phi(s, a)^\top] x}{x^\top \mathbb{E}_\rho [\phi(s, a) \phi(s, a)^\top] x}.$$

*Proof.* See Lemma 12 in (Uehara & Sun, 2021) for a detailed proof.  $\square$

C. Complete experimental results on noised D4RL tasks

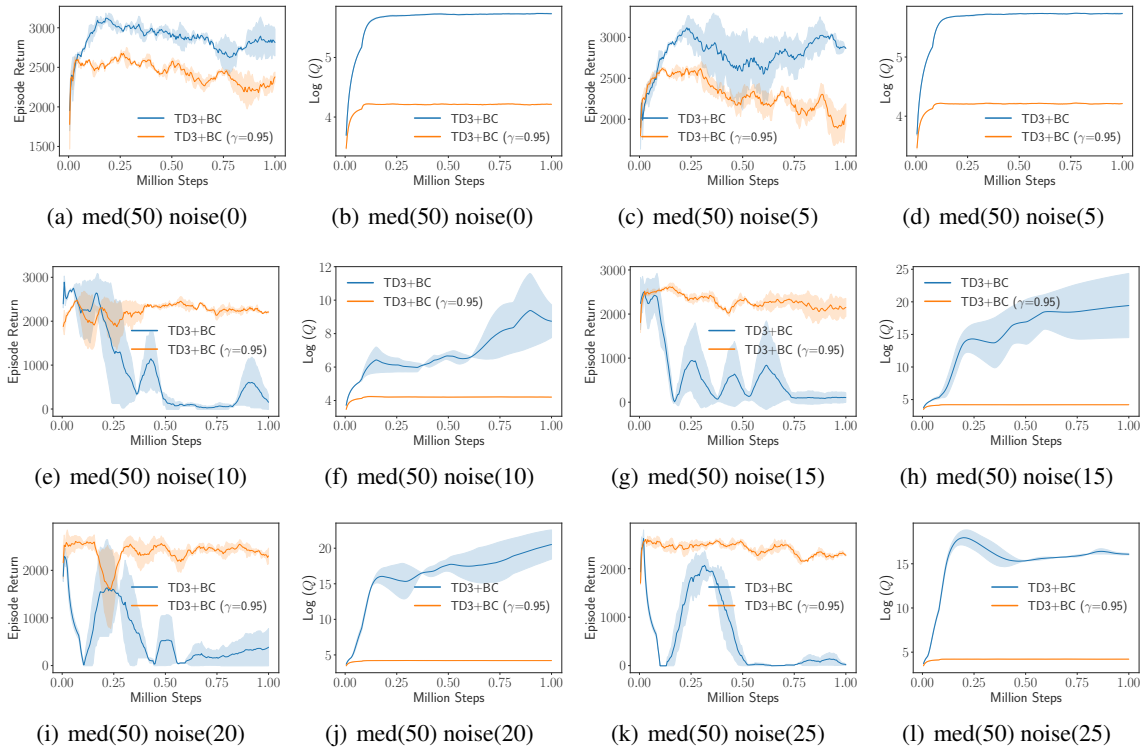


Figure 5. Experimental results on walker2d task consisting of 50 medium trajectories and  $x$  noised trajectories. The evaluation metric is the episode return and log  $Q$ -value.

## On the Role of Discount Factor in Offline Reinforcement Learning

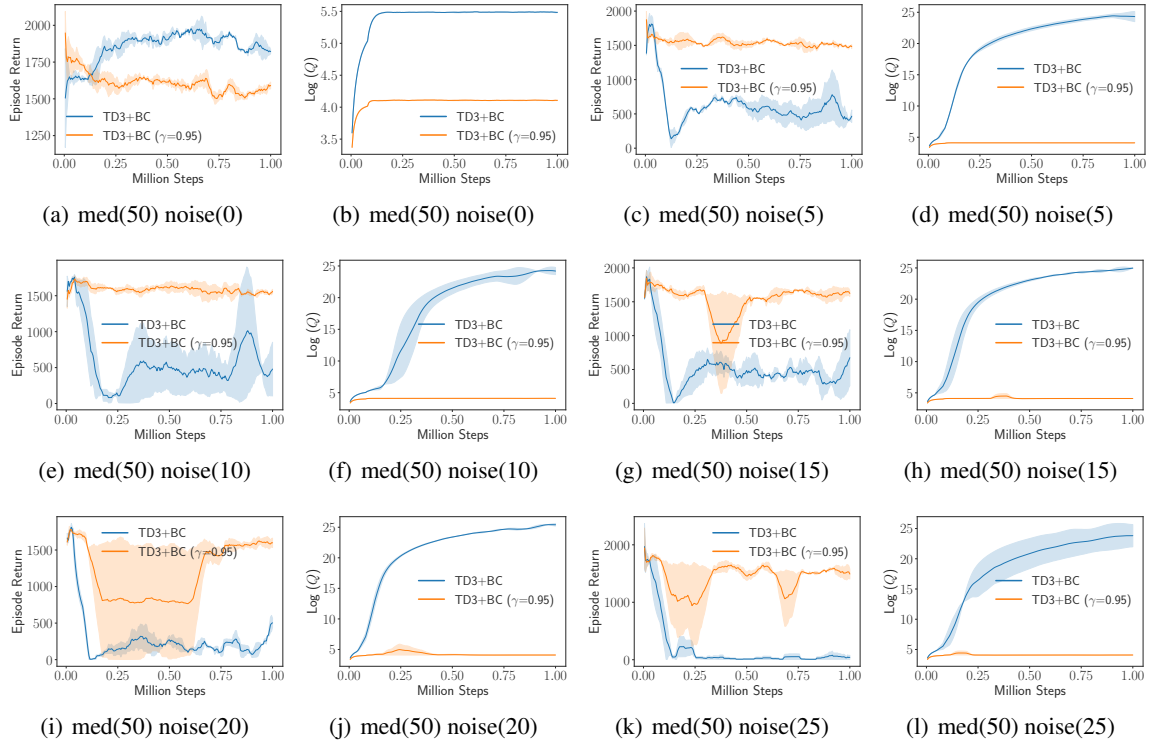


Figure 6. Experimental results on hopper task consisting of 50 medium trajectories and  $x$  noised trajectories. The evaluation metric is the episode return and log  $Q$ -value.

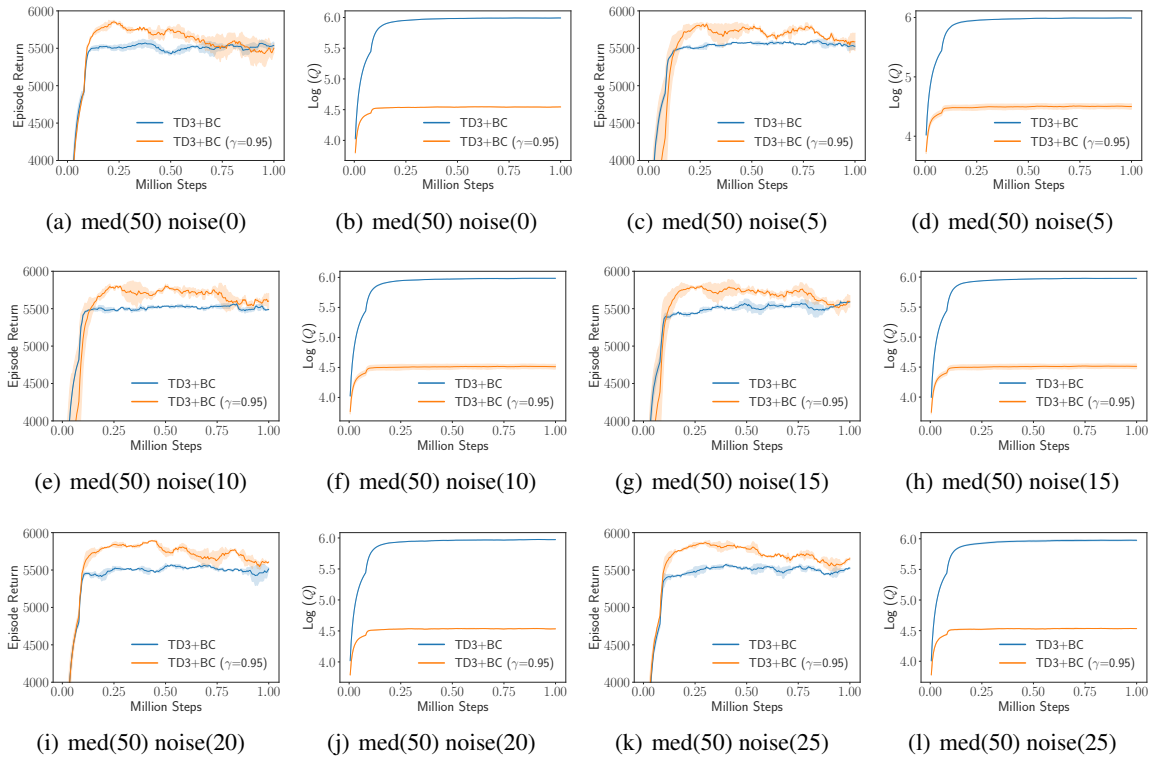


Figure 7. Experimental results on halfcheetah task consisting of 50 medium trajectories and  $x$  noised trajectories. The evaluation metric is the episode return and log  $Q$ -value.