# Appendix 



Figure 2: Visualization of singleton absorbing MDP technique

We provide a visualization (Figure 2) for understanding the singleton absorbing MDP technique at the beginning of Appendix. 2(a), 2(c) demonstrate the infinite horizon case and 2(b), 2(d) demonstrate the finite horizon case. In particular, it should be a $H$-dimensional hypercube $[0, H]^{H}$ (that contains $\left.\widehat{V}_{1}^{\star}, \ldots, \widehat{V}_{h}^{\star}\right)$ instead of only the square $[0, H] \times[0, H]\left(\widehat{V}_{1}^{\star}, \widehat{V}_{2}^{\star}\right)$. This is only for the ease of visualization.

The standard absorbing MDP technique Agarwal et al. [2020], Cui and Yang [2020] leverages a set of absorbing MDPs to cover the range of value functions (following the standard covering principle) to make sure $\widehat{V}^{\star}$ is close to one of the element (absorbing MDP) in the set (Figure 2(a) 2(b)). The size of the covering set (i.e. the covering number) grows exponentially in $H$ 2(b) in the finite horizon setting and this is due to the fact that there are $\widehat{V}_{1}^{\star}, \widehat{V}_{2}^{\star}, \ldots, \widehat{V}_{H}^{\star}$ quantities to cover. This results in the metric entropy (the log of the covering number) to blow up by a factor of $H$ and incurs suboptimality. On the other hand, by the nifty chosen singleton absorbing MDP $\widehat{V}_{h, u^{\star}}^{\star}$ (Figure 2(c) 2(d)), we completely get rid of the covering issue (covering the $H$-dimensional space requires exponential in $H$ size), maintain the independence and control the error propagation $\left(\left\|\widehat{V}^{\star}-\widehat{V}_{u^{\star}}^{\star}\right\|_{\infty}\right.$ is sufficiently small). See Section for all the technical details.

## A Discussion on Related works

Offline reinforcement learning. Information-theoretical considerations for offline RL are first proposed for infinite horizon discounted setting via Fitted Q-Iteration (FQI) type function approximation
algorithms [Chen and Jiang, 2019, Le et al., 2019, Xie and Jiang, 2021, 2020] which can be traced back to [Munos, 2003, Szepesvári and Munos, 2005, Antos et al.| 2008a b]. Later, Xie and Jiang [2021] considers the offline RL under only the realizability assumption and Liu et al. |2020b| considers the offline RL without good exploration. Those are all challenging problems but with they only provide suboptimal polynomial complexity in terms of $(1-\gamma)^{-1}$.

For the finite horizon case, Yin et al. [2021a] first achieves $\tilde{O}\left(H^{3} / d_{m} \epsilon^{2}\right)$ complexity under nonstationary transition but their results cannot be further improved in the stationary setting. Concurrent to our work, a recently released work Yin et al. [2021b] designs the offline variance reduction algorithm for achieving the optimal $\tilde{O}\left(H^{2} / d_{m} \epsilon^{2}\right)$ rate. Their result is for a specific algorithm that uses data splitting while our results work for any algorithms that returns a nearly empirically optimal policy via a uniform convergence guarantee. Our results on the offline task-agnostic and the rewardfree settings are entirely new. Another concurrent work Ren et al. [2021] considers the horizon-free setting but does not provide uniform convergence guarantee. Even more recently, Rashidinejad et al. [2021] considers the single concentrability coefficient $C^{\star}:=\max _{s, a} \frac{d^{\pi^{\star}}(s, a)}{d^{\mu}(s, a)}$ and obtains the sample complexity $\tilde{O}\left[(1-\gamma)^{-5} S C^{\star} / \epsilon^{2}\right]$. Concurrently, Yin and Wang [2021] first derives the instance-dependent offline RL bounds, but their result is for time-inhomogeneous MDPs.

In the linear MDP case, Jin et al. [2020c] studies the pessimism-based algorithms for offline policy optimization under the weak compliance assumption and Wang et al. [2021], Zanette [2021] provide some negative results (exponential lower bound) for offline RL with linear MDP structure.

Model-based approaches with minimaxity. It is known model-based methods are minimax-optimal for online RL with regret $\tilde{O}(\sqrt{H S A T})($ e.g. Azar et al. [2017], Efroni et al. [2019]). For linear MDP, In the generative model setting, Agarwal et al. [2020] shows model-based approach is still minimax optimal $\tilde{O}\left((1-\gamma)^{-3} S A / \epsilon^{2}\right)$ by using a $s$-absorbing MDP construction and this modelbased technique is later reused for other more general settings (e.g. Markov games [Zhang et al. 2020a] and linear MDPs [Cui and Yang, 2020|) and also for improving the sample size barrier [Li et al., 2020]. In offline RL, Yu et al. [2020], Kidambi et al. [2020] use model-based approaches for continuous policy optimization and Yin et al. [2021a] uses the model-based methods to achieve $\tilde{O}\left(H^{3} / d_{m} \epsilon^{2}\right)$ complexity.
Task-agnostic and Reward-free problems. The reward-free problem is initiated in the online RL [Jin et al., 2020a] where the agent needs to efficiently explore an MDP environment without using any reward information. It requires high probability guarantee for learning optimal policy for any reward function, which is strictly stronger than the standard learning task that one only needs to learn to optimal policy for a fixed reward. Later, Kaufmann et al. [2020], Menard et al. [2020] establish the $\tilde{O}\left(H^{3} S^{2} A / \epsilon^{2}\right)$ complexity and Zhang et al. 2020c] further tightens the dependence to $\tilde{O}\left(H^{2} S^{2} A / \epsilon^{2}\right){ }^{9}$ Recently, Zhang et al. [2020b] proposes the task-agnostic setting where one needs to use exploration data to simultaneously learn $K$ tasks and provides a upper bound with complexity $\tilde{O}\left(H^{5} S A \log (K) / \epsilon^{2}\right)$. For linear MDP setting, Wang et al. [2020] achieves the sample complexity $\tilde{O}\left(d^{3} H^{6} / \epsilon^{2}\right)$ and Liu et al. [2020a] considers such problem in the online two-player Markov game. However, although these settings remain critical in the offline regime, no statistical result has been formally derived so far.

## B Proof of optimal local uniform convergence

## B. 1 Model-based Offline Plug-in Estimator

Recall the model-based estimator uses empirical estimator $\widehat{P}$ for estimating $P$ and the estimator is calculated accordingly:

$$
\widehat{Q}_{h}^{\pi}=r+\widehat{P}^{\pi_{h+1}} Q_{h+1}^{\pi}=r+\widehat{P} V_{H+1}^{\pi}
$$

where $\widehat{P}\left(s^{\prime} \mid s, a\right)$ can be expressed as:

[^0]$$
\widehat{P}\left(s^{\prime} \mid s, a\right)=\frac{\sum_{i=1}^{n} \sum_{h=1}^{H} \mathbf{1}\left[\left(s_{h+1}^{(i)}, a_{h}^{(i)}, s_{h}^{(i)}\right)=\left(s^{\prime}, s, a\right)\right]}{n_{s, a}}, \quad n_{s, a}=\sum_{h=1}^{H} \sum_{i=1}^{n} \mathbf{1}\left[\left(s_{h}^{(i)}, a_{h}^{(i)}\right)=(s, a)\right]
$$
and $\widehat{P}\left(s^{\prime} \mid s, a\right)=\frac{1}{S}$, if $n_{s, a}=0$. The initial distribution is also constructed as $\widehat{d_{1}^{\pi}}(s)=n_{s} / n$.
First of all, we have by definition the Bellman optimality equation
\[

$$
\begin{equation*}
V_{t}^{\star}(s)=\max _{a}\left\{r(s, a)+\sum_{s^{\prime}} P\left(s^{\prime} \mid s, a\right) V_{t+1}^{\star}\left(s^{\prime}\right)\right\}, \quad \forall s \in \mathcal{S} \tag{3}
\end{equation*}
$$

\]

and similarly the empirical version

$$
\widehat{V}_{t}^{\star}(s)=\max _{a}\left\{r(s, a)+\sum_{s^{\prime}} \widehat{P}\left(s^{\prime} \mid s, a\right) \widehat{V}_{t+1}^{\star}\left(s^{\prime}\right)\right\}, \quad \forall s \in \mathcal{S}
$$

The key difficulty in obtaining the optimal dependence in stationary setting is decoupling the dependence of $P-\widehat{P}$ and $\widehat{V}^{\star}$. This issue is not encountered in the non-stationary setting due to the possibility to estimate different transition at each time [Yin et al., 2021a], but it cannot further reduce the sample complexity on $H$. Moreover, the direct use of $s$-absorbing MDP in Agarwal et al. [2020] is not sharp for finite horizon stationary setting, as it requires $s$-absorbing MDPs with $H$-dimensional cover (which has size $\approx e^{H}$ and it is not optimal). We design the singleton-absorbing MDP to get rid of the issue.

## B. 2 General absorbing MDP

The general absorbing MDP is defined as follows: for a fixed state $s$ and a sequence $\left\{u_{t}\right\}_{t=1}^{H}$, MDP $M_{s,\left\{u_{t}\right\}_{t=1}^{H}}$ is identical to $M$ for all states except $s$, and state $s$ is absorbing in the sense $P_{M_{s,\left\{u_{t}\right\}_{t=1}^{H}}}(s \mid s, a)=1$ for all $a$, and the instantaneous reward at time $t$ is $r_{t}(s, a)=u_{t}$ for all $a \in \mathcal{A}$. Also, we use the shorthand notation $V_{\left\{s, u_{t}\right\}}^{\pi}$ for $V_{s, M_{s,\left\{u_{t}\right\}_{t=1}^{H}}^{\pi}}$ and similarly for $Q_{\left\{s, u_{t}\right\}}$ and transition $P_{\left\{s, u_{t}\right\}}$. Then the following properties hold:

## Lemma B.1.

$$
V_{h,\left\{s, u_{t}\right\}}^{\star}(s)=\sum_{t=h}^{H} u_{t}
$$

Proof. We prove by backward induction. For $h=H$, under $M_{s,\left\{u_{t}\right\}_{t=1}^{H}}$ state $s$ is absorbing (and by convention $V_{H+1,\left\{s, u_{t}\right\}}^{\star}=0$ ) therefore
$V_{H,\left\{s, u_{t}\right\}}^{\star}(s)=\max _{a}\left\{r_{H,\left\{s, u_{t}\right\}}(s, a)+\sum_{s^{\prime}} P_{\left\{s, u_{t}\right\}}\left(s^{\prime} \mid s, a\right) V_{H+1,\left\{s, u_{t}\right\}}^{\star}\left(s^{\prime}\right)\right\}=\max _{a}\left\{r_{H,\left\{s, u_{t}\right\}}(s, a)\right\}=u_{H}$
for general $h$, note $\sum_{s^{\prime}} P_{\left\{s, u_{t}\right\}}\left(s^{\prime} \mid s, a\right) V_{h+1,\left\{s, u_{t}\right\}}^{\star}\left(s^{\prime}\right)=1 \cdot V_{h+1,\left\{s, u_{t}\right\}}^{\star}(s)$, therefore using induction property $V_{h+1,\left\{s, u_{t}\right\}}^{\star}(s)=\sum_{t=h+1}^{H} u_{t}$ we can similarly obtain $V_{h,\left\{s, u_{t}\right\}}^{\star}(s)=\sum_{t=h}^{H} u_{t}$.

Lemma B.2. Fix state s. For two different sequences $\left\{u_{t}\right\}_{t=1}^{H}$ and $\left\{u_{t}^{\prime}\right\}_{t=1}^{H}$, we have

$$
\max _{h}\left\|Q_{h,\left\{s, u_{t}\right\}}^{\star}-Q_{h,\left\{s, u_{t}^{\prime}\right\}}^{\star}\right\|_{\infty} \leq H \cdot \max _{t \in[H]}\left|u_{t}-u_{t}^{\prime}\right|
$$

Proof. Let $\pi_{\left\{s, u_{t}\right\}}^{\star}$ be the optimal policy in $M_{\left\{s, u_{t}\right\}}$. Then (by convention $\prod_{a=h+1}^{h} P^{\pi_{a}}=I$ )

$$
\begin{aligned}
& Q_{h,\left\{s, u_{t}\right\}}^{\star}-Q_{h,\left\{s, u_{t}^{\prime}\right\}}^{\star}=Q_{h,\left\{s, u_{t}\right\}}^{\star}-\max _{\pi} \sum_{i=h}^{H}\left(\prod_{a=h+1}^{i} P_{\left\{s, u_{t}^{\prime}\right\}}^{\pi_{a}}\right) r_{i,\left\{s, u_{t}^{\prime}\right\}} \\
& \leq Q_{h,\left\{s, u_{t}\right\}}^{\star}-\sum_{i=h}^{H}\left(\prod_{a=h+1}^{i} P_{\left\{s, u_{t}^{\prime}\right\}}^{\pi_{a,\left\{s, u_{t}\right\}}^{\star}}\right) r_{i,\left\{s, u_{t}^{\prime}\right\}}=\sum_{i=h}^{H}\left(\prod_{a=h+1}^{i} P_{\left\{s, u_{t}^{\prime}\right\}}^{\pi_{a,\left\{s, u_{t}\right\}}^{\star}}\right)\left(r_{i,\left\{s, u_{t}\right\}}-r_{i,\left\{s, u_{t}^{\prime}\right\}}\right) \\
& \quad \leq \sum_{i=h}^{H} \max _{s, a}\left\|\left(\prod_{a=h+1}^{i} P_{\left\{s, u_{t}^{\prime}\right\}}^{\pi_{a,\left\{s, u_{t}\right\}}^{\star}}\right)^{i-h}(\cdot \mid s, a)\right\|_{1} \cdot\left\|r_{i,\left\{s, u_{t}\right\}}-r_{i,\left\{s, u_{t}^{\prime}\right\}}\right\|_{\infty} \cdot \mathbf{1}=(H-h+1) \cdot \max _{t}\left|u_{t}-u_{t}^{\prime}\right| \cdot \mathbf{1}
\end{aligned}
$$

where the first equal sign uses the definition of $Q^{\star}$, the second equal sign uses $P_{\left\{s, u_{t}\right\}}$ only depends $s$ but not the specification of $u_{t}$ 's and the last equal sign comes from $r_{i,\left\{s, u_{t}\right\}}(s, a)=u_{i}$ for any $a \in \mathcal{A}$ and $r_{i,\left\{s, u_{t}\right\}}(\tilde{s}, a)=r_{i,\left\{s, u_{t}^{\prime}\right\}}(\tilde{s}, a)$ for any $\tilde{s} \neq s$. Lastly by symmetry we finish the proof.

## B. 3 Singleton-absorbing MDP

The direct of transfer of absorbing technique created in Agarwal et al. [2020] will require each $u_{t}$ to fill in the range of $[0, H]$ using evenly spaced elements. For finite horizon MDP there are $H$ layers, therefore the total number of $H$-tuples $\left(u_{1}, \ldots, u_{H}\right)$ has order $\left|U_{s}\right|=\operatorname{Poly}(H)^{H}$, therefore when apply the union bound, it will incur the additional $H$ factor. We get rid of this issue by choosing one single point in $H$-dimensional space $[0, H]^{H}$. We first give the following two lemmas.

Lemma B.3. $V_{t}^{\star}(s)-V_{t+1}^{\star}(s) \geq 0$, for all state $s \in \mathcal{S}$ and all $t \in[H]$.
Proof. Let the optimal policy for $V_{t+1}^{\star}$ be $\pi_{t+1: H}^{\star}$, i.e. $V_{t+1}^{\star}=V_{t+1}^{\pi_{t+1: H}^{\star}}$, then artificially construct a policy $\pi_{t: H}$ such that $\pi_{t: H-1}=\pi_{t+1: H}^{\star}$ and $\pi_{H}$ is arbitrary, then by the definition of optimal value

$$
\begin{aligned}
V_{t}^{\star}(s) & \geq V_{t}^{\pi_{t: H}}(s)=\mathbb{E}^{\pi_{t: H}}\left[\sum_{i=t}^{H} r\left(s_{i}, a_{i}\right) \mid s_{t}=s\right] \\
& =\mathbb{E}^{\pi_{t: H-1}}\left[\sum_{i=t}^{H-1} r\left(s_{i}, a_{i}\right) \mid s_{t}=s\right]+\mathbb{E}^{\pi_{t: H}}\left[r\left(s_{H}, a_{H}\right) \mid s_{t}=s\right] \\
& =\mathbb{E}^{\pi_{t+1: H}^{\star}}\left[\sum_{i=t+1}^{H} r\left(s_{i}, a_{i}\right) \mid s_{t+1}=s\right]+\mathbb{E}^{\pi_{t: H}}\left[r\left(s_{H}, a_{H}\right) \mid s_{t}=s\right] \\
& \geq \mathbb{E}^{\pi_{t+1: H}^{\star}}\left[\sum_{i=t+1}^{H} r\left(s_{i}, a_{i}\right) \mid s_{t+1}=s\right]+0=V_{t+1}^{\star}(s)
\end{aligned}
$$

where the third equal sign uses exactly that $P$ is a STATIONARY transition and definition $\pi_{t: H-1}=$ $\pi_{t+1: H}^{\star}$. The last inequality uses assumption that reward is always non-negative.
Remark B.4. Lemma B.3 leverages $P$ is stationary and above may not be true in the non-stationary setting. This enables us to establish the following lemma, which is the key for singleton-absorbing MDP.
Lemma B.5. Fix a state s. If we choose $u_{t}^{\star}:=V_{t}^{\star}(s)-V_{t+1}^{\star}(s) \forall t \in[H]$, then we have the following vector form equation

$$
V_{h,\left\{s, u_{t}^{\star}\right\}}^{\star}=V_{h, M}^{\star} \quad \forall h \in[H] .
$$

Similarly, if we choose $\widehat{u}_{t}^{\star}:=\widehat{V}_{t}^{\star}(s)-\widehat{V}_{t+1}^{\star}(s)$, then $\widehat{V}_{h,\left\{s, \hat{u}_{t}^{\star}\right\}}^{\star}=\widehat{V}_{h, M}^{\star}$, $\forall h \in[H]$.
Proof. We focus on the first claim. Note by Lemma B.3 the assignment of $u_{t}^{\star}\left(:=r_{t,\left\{s, u_{t}^{\star}\right\}}\right)$ is well-defined. Next recall $V_{h, M}^{\star}$ is the optimal value under true MDP $M$ and $V_{h,\left\{s, u_{t}^{\star}\right\}}^{\star}$ is the optimal value under the assimilating MDP $M_{s,\left\{u_{t}^{\star}\right\}_{t=1}^{H}}$. We prove by backward induction.

For $h=H$, note by convention $V_{H+1}^{\star}=0$, therefore $u_{H}^{\star}=V_{H}^{\star}(s)-V_{H+1}^{\star}(s)=V_{H}^{\star}(s)-0=V_{H}^{\star}(s)$ and Bellman optimality equation becomes

$$
V_{H}^{\star}(\tilde{s})=\max _{a}\{r(\tilde{s}, a)\}, \quad \forall \tilde{s} \in \mathcal{S}
$$

Under $M_{s,\left\{u_{t}^{\star}\right\}_{t=1}^{H}}$, for state $s$ by Lemma B.1 we have $V_{H,\left\{s, u_{t}^{\star}\right\}}^{\star}(s)=u_{H}^{\star}=V_{H}^{\star}(s)$, for other states $\tilde{s} \neq s$, reward in $M_{s,\left\{u_{t}^{\star}\right\}_{t=1}^{H}}=M$ so we also have $V_{H,\left\{s, u_{t}^{\star}\right\}}^{\star}(\tilde{s})=V_{H}^{\star}(\tilde{s})$ for all $\tilde{s} \neq s$.
Now for general $h$, for state $s$ by Lemma B. 1

$$
V_{h,\left\{s, u_{t}^{\star}\right\}}^{\star}(s)=\sum_{t=h}^{H} u_{t}^{\star}=\sum_{t=h}^{H}\left(V_{t}^{\star}(s)-V_{t+1}^{\star}(s)\right)=V_{h}^{\star}(s),
$$

for state $\tilde{s} \neq s$, by Bellman optimality equation

$$
\begin{aligned}
V_{h,\left\{s, u_{t}^{\star}\right\}}^{\star}(\tilde{s}) & =\max _{a}\left\{r_{\left\{s, u_{t}^{\star}\right\}}(\tilde{s}, a)+\sum_{s^{\prime}} P_{\left\{s, u_{t}^{\star}\right\}}\left(s^{\prime} \mid \tilde{s}, a\right) V_{h+1,\left\{s, u_{t}^{\star}\right\}}^{\star}\left(s^{\prime}\right)\right\} \\
& =\max _{a}\left\{r(\tilde{s}, a)+\sum_{s^{\prime}} P\left(s^{\prime} \mid \tilde{s}, a\right) V_{h+1,\left\{s, u_{t}^{\star}\right\}}^{\star}\left(s^{\prime}\right)\right\} \\
& =\max _{a}\left\{r(\tilde{s}, a)+\sum_{s^{\prime}} P\left(s^{\prime} \mid \tilde{s}, a\right) V_{h+1}^{\star}\left(s^{\prime}\right)\right\}=V_{h}^{\star}(\tilde{s})
\end{aligned}
$$

where the second equal sign uses when $\tilde{s} \neq s, M_{s,\left\{u_{t}^{\star}\right\}_{t=1}^{H}}$ is identical to $M$ and the third equal sign uses induction assumption that element-wisely $V_{h+1,\left\{s, u_{t}^{\star}\right\}}^{\star}=V_{h+1}^{\star}$. Similar result can be derived for $\hat{u}^{\star}$ version and this completes the proof.

The singleton MDP we used is exactly $M_{s,\left\{u_{t}^{\star}\right\}_{t=1}^{H}}\left(\right.$ or $\widehat{M}_{s,\left\{u_{t}^{\star}\right\}_{t=1}^{H}}$ ).

## B. 4 Proof for local uniform convergence

Recall the local policy class

$$
\Pi_{l}:=\left\{\pi \text { : s.t. }\left\|\widehat{V}_{h}^{\pi}-\widehat{V}_{h}^{\widehat{\pi}^{\star}}\right\|_{\infty} \leq \epsilon_{\mathrm{opt}}, \forall h \in[H]\right\} .
$$

For ease of exposition, we denote $N:=\min _{s, a} n_{s, a}$. Note $N$ itself is a random variable, therefore for the rest of proof we first conditional on $N$. Later we shall remove the conditional on $N$ (see Section B.7.

For any $\widehat{\pi} \in \Pi_{l}$, by (empirical) Bellman equation we have element-wisely:

$$
\begin{aligned}
\widehat{Q}_{h}^{\pi}-Q_{h}^{\widehat{\pi}} & =r_{h}+\widehat{P}^{\widehat{\pi}_{h+1}} \widehat{Q}_{h+1}^{\hat{\pi}}-r_{h}-P^{\widehat{\pi}_{h+1}} Q_{h+1}^{\widehat{\pi}} \\
& =\left(\widehat{P}^{\widehat{\pi}_{h+1}}-P^{\widehat{\pi}_{h+1}}\right) \widehat{Q}_{h+1}^{\widehat{\pi}}+P^{\widehat{\pi}_{h+1}}\left(\widehat{Q}_{h+1}^{\widehat{\pi}}-Q_{h+1}^{\widehat{\pi}}\right) \\
& =(\widehat{P}-P) \widehat{V}_{h+1}^{\widehat{\pi}}+P^{\widehat{\pi}_{h+1}}\left(\widehat{Q}_{h+1}^{\widehat{\pi}}-Q_{h+1}^{\widehat{\pi}}\right) \\
& =\ldots=\sum_{t=h}^{H} \Gamma_{h+1: t}^{\widehat{\pi}}(\widehat{P}-P) \widehat{V}_{t+1}^{\widehat{\pi}} \\
& \leq \underbrace{\sum_{t=h}^{H} \Gamma_{h+1: t}^{\widehat{\pi}}\left|(\widehat{P}-P) \widehat{V}_{t+1}^{\widehat{\pi}^{\star}}\right|}_{(*)}+\underbrace{\sum_{t=h}^{H} \Gamma_{h+1: t}^{\widehat{\pi}}\left|(\widehat{P}-P)\left(\widehat{V}_{t+1}^{\widehat{\pi}}-\widehat{V}_{t+1}^{\widehat{\pi}^{\star}}\right)\right|}_{(* *)}
\end{aligned}
$$

where $\Gamma_{h+1: t}^{\pi}=\prod_{i=h+1}^{t} P^{\pi_{i}}$ is multi-step state-action transition and $\Gamma_{h+1: h}:=I$.

## B. 5 Analyzing ( $\star \star$ )

Term ( $\star \star$ ) can be readily bounded using the following lemma.
Lemma B.6. Fix $N>0$, we have with probability $1-\delta$, for all $t=1, \ldots, H-1$

$$
\sum_{t=h}^{H} \Gamma_{h+1: t}^{\widehat{\pi}}\left|(\widehat{P}-P)\left(\widehat{V}_{h+1}^{\widehat{\pi}^{\star}}-\widehat{V}_{h+1}^{\widehat{\pi}}\right)\right| \leq C \epsilon_{\mathrm{opt}} \cdot \sqrt{\frac{H^{2} S \log (S A / \delta)}{N}} \cdot \mathbf{1}
$$

where $C$ absorb the higher order term and absolute constants.

Proof. First, by vector induced matrix norm ${ }^{10}$ we have

$$
\begin{aligned}
\left\|\sum_{t=h}^{H} \Gamma_{h+1: t}^{\widehat{\pi}} \cdot\left|(\widehat{P}-P)\left(\widehat{V}_{t+1}^{\widehat{\pi}^{\star}}-\widehat{V}_{t+1}^{\hat{\pi}}\right)\right|\right\|_{\infty} & \leq H \cdot \sup _{t}\left\|\Gamma_{h+1: t}^{\widehat{\pi}}\right\|_{\infty}\left\|\left|(\widehat{P}-P)\left(\widehat{V}_{t+1}^{\widehat{\pi}^{\star}}-\widehat{V}_{t+1}^{\widehat{\pi}}\right)\right|\right\|_{\infty} \\
& \leq H \cdot \sup _{t}\left\|\left|(\widehat{P}-P)\left(\widehat{V}_{t+1}^{\widehat{\pi}^{\star}}-\widehat{V}_{t+1}^{\hat{\pi}}\right)\right|\right\|_{\infty} \\
& =H \cdot \sup _{t, s, a}\left|(\widehat{P}-P)(\cdot \mid s, a)\left(\widehat{V}_{t+1}^{\widehat{\pi}^{\star}}-\widehat{V}_{t+1}^{\widehat{\pi}}\right)\right| \\
& \leq H \cdot \sup _{t, s, a}\|(\widehat{P}-P)(\cdot \mid s, a)\|_{1} \cdot\left\|\widehat{V}_{t+1}^{\widehat{\pi}^{\star}}-\widehat{V}_{t+1}^{\widehat{\pi}}\right\|_{\infty} \cdot \mathbf{1}
\end{aligned}
$$

where the second inequality uses multi-step transition $\Gamma_{t+1: h-1}^{\pi}$ is row-stochastic. Note given $N$, therefore by Lemma J. 9 and a union bound we have with probability $1-\delta$,

$$
\sup _{s, a}\|(\widehat{P}-P)(\cdot \mid s, a)\|_{1} \leq C\left(\sqrt{\frac{S \log (S A / \delta)}{N}}\right)
$$

(where $C$ absorb the higher order term and absolute constants) and using definition of $\Pi_{l}$ we have $\sup _{t}\left\|\widehat{V}_{t}^{\widehat{\pi}^{\star}}-\widehat{V}_{t}^{\hat{\pi}}\right\|_{\infty} \leq \epsilon_{\mathrm{opt}}$. This indicates

$$
\sup _{t, s, a}\|(\widehat{P}-P)(\cdot \mid s, a)\|_{1} \cdot\left\|\widehat{V}_{t+1}^{\widehat{\pi}^{\star}}-\widehat{V}_{t+1}^{\widehat{\pi}}\right\|_{\infty} \cdot \mathbf{1} \leq C\left(\epsilon_{\mathrm{opt}} \sqrt{\frac{S \log (S A / \delta)}{N}} \cdot \mathbf{1}\right)
$$

where $1 \in \mathbb{R}^{S}$ is all-one vector. Then multiple by $H$ to get the stated result.

## B. 6 Analyzing (*)

 are on longer independent, Bernstein inequality cannot be directly applied. We use the singletonabsorbing MDP $M_{s,\left\{u_{t}^{\star}\right\}_{t=1}^{H}}$ to handle the case (recall $u_{t}^{\star}:=V_{t}^{\star}(s)-V_{t+1}^{\star}(s) \forall t \in[H]$ ). Again, let us fix a state $s$ and $a \in \mathcal{A}$ be any action. Also, we use $P_{s, a}$ to denote row vector to avoid long

[^1]expression. Then we have:
\[

\left.$$
\begin{array}{rl} 
& \left(\widehat{P}_{s, a}-P_{s, a}\right) \widehat{V}_{h}^{\star}=\left(\widehat{P}_{s, a}-P_{s, a}\right)\left(\widehat{V}_{h}^{\star}-\widehat{V}_{h,\left\{s, u_{t}^{\star}\right\}}^{\star}+\widehat{V}_{h,\left\{s, u_{t}^{\star}\right\}}^{\star}\right) \\
= & \left(\widehat{P}_{s, a}-P_{s, a}\right)\left(\widehat{V}_{h}^{\star}-\widehat{V}_{h,\left\{s, u_{t}^{\star}\right\}}^{\star}\right)+\left(\widehat{P}_{s, a}-P_{s, a}\right) \widehat{V}_{h,\left\{s, u_{t}^{\star}\right\}}^{\star} \\
\leq & \left\|\widehat{P}_{s, a}-P_{s, a}\right\|_{1}\left\|\widehat{V}_{h}^{\star}-\widehat{V}_{h,\left\{s, u_{t}^{\star}\right\}}^{\star}\right\|_{\infty}+\sqrt{\frac{2 \log (4 / \delta)}{N}} \sqrt{\operatorname{Var}_{s, a}\left(\widehat{V}_{h,\left\{s, u_{t}^{\star}\right\}}^{\star}\right)}+\frac{2 H \log (1 / \delta)}{3 N} \\
\leq & \left\|\widehat{P}_{s, a}-P_{s, a}\right\|_{1}\left\|\widehat{V}_{h}^{\star}-\widehat{V}_{h,\left\{s, u_{t}^{\star}\right\}}^{\star}\right\|_{\infty}+\sqrt{\frac{2 \log (4 / \delta)}{N}}\left(\sqrt{\operatorname{Var}_{s, a}\left(\widehat{V}_{h}^{\star}\right)}+\sqrt{\operatorname{Var}_{s, a}\left(\widehat{V}_{h,\left\{s, u_{t}^{\star}\right\}}^{\star}-\widehat{V}_{h}^{\star}\right)}\right)+\frac{2 H \log (1 / \delta)}{3 N} \\
\leq & \left\|\widehat{P}_{s, a}-P_{s, a}\right\|_{1}\left\|\widehat{V}_{h}^{\star}-\widehat{V}_{h,\left\{s, u_{t}^{\star}\right\}}^{\star}\right\|_{\infty}+\sqrt{\frac{2 \log (4 / \delta)}{N}}\left(\sqrt{\operatorname{Var}_{s, a}\left(\widehat{V}_{h}^{\star}\right)}+\sqrt{\left\|\widehat{V}_{h,\left\{s, u_{t}^{\star}\right\}}-\widehat{V}_{h}^{\star}\right\|_{\infty}^{2}}\right)+\frac{2 H \log (1 / \delta)}{3 N} \\
= & \left(\left\|\widehat{P}_{s, a}-P_{s, a}\right\|_{1}+\sqrt{\frac{2 \log (4 / \delta)}{N}}\right) \| \widehat{V}_{h}^{\star}-\widehat{V}_{h,\left\{s, u_{t}^{\star}\right\}} \tag{4}
\end{array}
$$ \|_{\infty}+\sqrt{\frac{2 \log (4 / \delta)}{N}} \sqrt{\operatorname{Var}_{s, a}\left(\widehat{V}_{h}^{\star}\right)}+\frac{2 H \log (1 / \delta)}{3 N}\right)
\]

where the first inequality uses Bernstein inequality (Lemma J.3, the second inequality uses $\sqrt{\operatorname{Var}(\cdot)}$ is norm (norm triangle inequality). Now we treat $\left\|\widehat{P}_{s, a}-P_{s, a}\right\|_{1}$ and $\left\|\widehat{V}_{h}^{\star}-\widehat{V}_{h,\left\{s, u_{t}^{\star}\right\}}^{\star}\right\|_{\infty}$ separately.
For $\left\|\widehat{P}_{s, a}-P_{s, a}\right\|_{1}$. Indeed, by Lemma J.9 again $\left\|\widehat{P}_{s, a}-P_{s, a}\right\|_{1} \leq \tilde{O}\left(\sqrt{\frac{S \log (S / \delta)}{N}}\right)$ and by a union bound we obtain w.p., $1-\delta$

$$
\begin{equation*}
\sup _{s, a}\left\|\widehat{P}_{s, a}-P_{s, a}\right\|_{1} \leq C \sqrt{\frac{S \log (S A / \delta)}{N}} . \tag{5}
\end{equation*}
$$

where $C$ absorbs the higher order term and constants.
For $\left\|\widehat{V}_{h}^{\star}-\widehat{V}_{h,\left\{s, u_{t}^{\star}\right\}}^{\star}\right\|_{\infty}$. Note if we set $\widehat{u}_{t}^{\star}=\widehat{V}_{t}^{\star}(s)-\widehat{V}_{t+1}^{\star}(s)$, then by Lemma B.5

$$
\widehat{V}_{h}^{\star}=\widehat{V}_{h,\left\{s, \hat{u}_{t}^{\star}\right\}}^{\star}
$$

Next since $\widehat{V}_{h,\left\{s, \hat{u}_{t}^{\star}\right\}}^{\star}(\tilde{s})=\max _{a} \widehat{Q}_{h,\left\{s, \hat{u}_{t}^{\star}\right\}}^{\star}(\tilde{s}, a) \forall \tilde{s} \in \mathcal{S}$, by generic inequality $|\max f-\max g| \leq$ $\max |f-g|$, we have $\left|\widehat{V}_{h,\left\{s, \hat{u}_{t}^{\star}\right\}}(\tilde{s})-\widehat{V}_{h,\left\{s, u_{t}^{\star}\right\}}^{\star}(\tilde{s})\right| \leq \max _{a}\left|\widehat{Q}_{h,\left\{s, \hat{u}_{t}^{\star}\right\}}(\tilde{s}, a)-\widehat{Q}_{h,\left\{s, u_{t}^{\star}\right\}}^{\star}(\tilde{s}, a)\right|$, taking $\max _{\tilde{s}}$ on both sides, we obtain exactly

$$
\left\|\widehat{V}_{h,\left\{s, \hat{u}_{t}^{\star}\right\}}^{\star}-\widehat{V}_{h,\left\{s, u_{t}^{\star}\right\}}^{\star}\right\|_{\infty} \leq\left\|\widehat{Q}_{h,\left\{s, \hat{u}_{t}^{\star}\right\}}^{\star}-\widehat{Q}_{h,\left\{s, u_{t}^{\star}\right\}}^{\star}\right\|_{\infty}
$$

then by Lemma B.2.

$$
\begin{equation*}
\left\|\widehat{V}_{h}^{\star}-\widehat{V}_{h,\left\{s, u_{t}^{\star}\right\}}^{\star}\right\|_{\infty} \leq\left\|\widehat{Q}_{h,\left\{s, \hat{u}_{t}^{\star}\right\}}^{\star}-\widehat{Q}_{h,\left\{s, u_{t}^{\star}\right\}}^{\star}\right\|_{\infty} \leq H \max _{t}\left|\hat{u}_{t}^{\star}-u_{t}^{\star}\right|, \tag{6}
\end{equation*}
$$

Recall

$$
\hat{u}_{t}^{\star}-u_{t}^{\star}=\widehat{V}_{t}^{\star}(s)-\widehat{V}_{t+1}^{\star}(s)-\left(V_{t}^{\star}(s)-V_{t+1}^{\star}(s)\right) .
$$

Now we denote

$$
\Delta_{s}:=\max _{t}\left|\hat{u}_{t}^{\star}-u_{t}^{\star}\right|=\max _{t}\left|\widehat{V}_{t}^{\star}(s)-\widehat{V}_{t+1}^{\star}(s)-\left(V_{t}^{\star}(s)-V_{t+1}^{\star}(s)\right)\right|
$$

then $\Delta_{s}$ itself is a scalar and a random variable.
To sum up, by (4), (5) and (6) and a union bound we have
Lemma B.7. Fix $N>0$. With probability $1-\delta$, element-wisely, for all $h \in[H]$,
$\left|(\widehat{P}-P) \widehat{V}_{h}^{\star}\right| \leq C \sqrt{\frac{S \log (H S A / \delta)}{N}} \cdot H \max _{s} \Delta_{s} \cdot \mathbf{1}+\sqrt{\frac{2 \log (4 H S A / \delta)}{N}} \sqrt{\operatorname{Var}_{P}\left(\widehat{V}_{h}^{\star}\right)}+\frac{2 H \log (H S A / \delta)}{3 N} \cdot \mathbf{1}$

Now plug Lemma B. 7 back into $(\star)$ and combine Lemma B.6, we receive:

$$
\begin{aligned}
& \left|\widehat{Q}_{h}^{\hat{\pi}}-Q_{h}^{\widehat{\pi}}\right| \\
\leq & \sum_{t=h}^{H} \Gamma_{h+1: t}^{\widehat{\pi}}\left(C \sqrt{\frac{S \log (H S A / \delta)}{N}} \cdot H \max _{s} \Delta_{s} \cdot \mathbf{1}+\sqrt{\frac{2 \log (4 H S A / \delta)}{N}} \sqrt{\operatorname{Var}_{P}\left(\widehat{V}_{t+1}^{\star}\right)}+\frac{2 H \log (H S A / \delta)}{3 N} \cdot \mathbf{1}\right) \\
+ & C \epsilon_{\mathrm{opt}} \cdot \sqrt{\frac{H^{2} S \log (S A / \delta)}{N}} \cdot \mathbf{1} \\
\leq & \sum_{t=h}^{H} \Gamma_{h+1: t}^{\widehat{\pi}} \sqrt{\frac{2 \log (4 H S A / \delta)}{N}} \sqrt{\operatorname{Var}_{P}\left(\widehat{V}_{t+1}^{\star}\right)}+C H^{2} \sqrt{\frac{S \log (H S A / \delta)}{N}} \cdot \max _{s} \Delta_{s} \cdot \mathbf{1}+\frac{2 H^{2} \log (H S A / \delta)}{3 N} \cdot \mathbf{1} \\
+ & C \epsilon_{\mathrm{opt}} \cdot \sqrt{\frac{H^{2} S \log (S A / \delta)}{N}} \cdot \mathbf{1}
\end{aligned}
$$

Next note

$$
\begin{align*}
& \sqrt{\operatorname{Var}_{P}\left(\widehat{V}_{h}^{\star}\right)}:=\sqrt{\operatorname{Var}_{P}\left(\widehat{V}_{h}^{\pi^{\star}}\right)}=\sqrt{\operatorname{Var}_{P}\left(\widehat{V}_{h}^{\widehat{\pi}}-\widehat{V}_{h}^{\widehat{\pi}}+\widehat{V}_{h}^{\widehat{\pi}}\right)} \\
\leq & \sqrt{\operatorname{Var}_{P}\left(\widehat{V}_{h}^{\hat{\pi}}\right)}+\sqrt{\operatorname{Var}_{P}\left(\widehat{V}_{h}^{\widehat{\pi}^{\star}}-\widehat{V}_{h}^{\widehat{\pi}}\right)} \leq \sqrt{\operatorname{Var}_{P}\left(\widehat{V}_{h}^{\widehat{\pi}}\right)}+\left\|\widehat{V}_{h}^{\widehat{\pi}^{\star}}-\widehat{V}_{h}^{\widehat{\pi}}\right\|_{\infty}  \tag{7}\\
\leq & \sqrt{\operatorname{Var}_{P}\left(\widehat{V}_{h}^{\hat{\pi}}\right)}+\epsilon_{\mathrm{opt}} \cdot \mathbf{1} \leq \sqrt{\operatorname{Var}_{P}\left(V_{h}^{\widehat{\pi}}\right)}+\sqrt{\operatorname{Var}_{P}\left(\widehat{V}_{h}^{\widehat{\pi}}-V_{h}^{\widehat{\pi}}\right)}+\epsilon_{\mathrm{opt}} \cdot \mathbf{1} \\
\leq & \sqrt{\operatorname{Var}_{P}\left(V_{h}^{\widehat{\pi}}\right)}+\left\|\widehat{V}_{h}^{\widehat{\pi}}-V_{h}^{\widehat{\pi}}\right\|_{\infty}+\epsilon_{\mathrm{opt}} \cdot \mathbf{1} \leq \sqrt{\operatorname{Var}_{P}\left(V_{h}^{\widehat{\pi}}\right)}+\left\|\widehat{Q}_{h}^{\widehat{\pi}}-Q_{h}^{\widehat{\pi}}\right\|_{\infty}+\epsilon_{\mathrm{opt}} \cdot \mathbf{1}
\end{align*}
$$

Plug (7) back to above we obtain $\forall h \in[H]$,

$$
\begin{align*}
& \left|\widehat{Q}_{h}^{\hat{\pi}}-Q_{h}^{\widehat{\pi}}\right| \leq \sqrt{\frac{2 \log (4 H S A / \delta)}{N}} \sum_{t=h}^{H} \Gamma_{h+1: t}^{\widehat{\pi}}\left(\sqrt{\operatorname{Var}_{P}\left(V_{t+1}^{\hat{\pi}}\right)}+\left\|\widehat{Q}_{t+1}^{\hat{\pi}}-Q_{t+1}^{\widehat{\pi}}\right\|_{\infty}+\epsilon_{\mathrm{opt}} \cdot \mathbf{1}\right) \\
& +C H^{2} \sqrt{\frac{S \log (H S A / \delta)}{N}} \cdot \max _{s} \Delta_{s} \cdot \mathbf{1}+\frac{2 H^{2} \log (H S A / \delta)}{3 N} \cdot \mathbf{1}+C \epsilon_{\mathrm{opt}} \cdot \sqrt{\frac{H^{2} S \log (S A / \delta)}{N}} \cdot \mathbf{1} \\
& \leq \sqrt{\frac{2 \log (4 H S A / \delta)}{N}} \sum_{t=h}^{H} \Gamma_{h+1: t}^{\widehat{\pi}} \sqrt{\operatorname{Var}_{P}\left(V_{t+1}^{\widehat{\pi}}\right)}+\sqrt{\frac{2 \log (4 H S A / \delta)}{N}} \sum_{t=h}^{H}\left\|\widehat{Q}_{t+1}^{\widehat{\pi}}-Q_{t+1}^{\widehat{\pi}}\right\|_{\infty} \\
& +C H^{2} \sqrt{\frac{S \log (H S A / \delta)}{N}} \cdot \max _{s} \Delta_{s} \cdot \mathbf{1}+\frac{2 H^{2} \log (H S A / \delta)}{3 N} \cdot \mathbf{1}+C_{1} \epsilon_{\mathrm{opt}} \cdot \sqrt{\frac{H^{2} S \log (S A / \delta)}{N}} \cdot \mathbf{1} \tag{8}
\end{align*}
$$

Apply Lemma J.5 and the coarse uniform bound (Lemma J.10) we obtain the following lemma:
Lemma B.8. Given $N>0$ and $\epsilon_{\text {opt }} \leq \sqrt{H / S}$. With probability $1-\delta$, for all $h \in[H]$,

$$
\left\|\widehat{Q}_{h}^{\widehat{\pi}}-Q_{h}^{\widehat{\pi}}\right\|_{\infty} \leq \sqrt{\frac{C_{0} H^{3} \log (4 H S A / \delta)}{N}}+\sqrt{\frac{2 \log (4 H S A / \delta)}{N}} \sum_{t=h}^{H}\left\|\widehat{Q}_{t+1}^{\widehat{\pi}}-Q_{t+1}^{\widehat{\pi}}\right\|_{\infty}+C^{\prime} H^{4} \frac{S \log (H S A / \delta)}{N}
$$

Proof. Since

$$
\begin{align*}
\Delta_{s} & :=\max _{t}\left|\hat{u}_{t}^{\star}-u_{t}^{\star}\right|=\max _{t}\left|\widehat{V}_{t}^{\star}(s)-\widehat{V}_{t+1}^{\star}(s)-\left(V_{t}^{\star}(s)-V_{t+1}^{\star}(s)\right)\right| \\
& \leq 2 \cdot \max _{t}\left|\widehat{V}_{t}^{\star}(s)-V_{t}^{\star}(s)\right| \\
& =2 \cdot \max _{t}\left|\max _{\pi} \widehat{V}_{t}^{\pi}(s)-\max _{\pi} V_{t}^{\pi}(s)\right|  \tag{9}\\
& \leq 2 \cdot \max _{\pi \in \Pi_{g}, t \in[H]}\left\|\widehat{V}_{t}^{\pi}-V_{t}^{\pi}\right\|_{\infty} \leq C \cdot H^{2} \sqrt{\frac{S \log (H S A / \delta)}{N}}
\end{align*}
$$

where the last inequality uses Lemma J. 10 . Then apply union bound w.p. $1-\delta / 2$, we obtain $\max _{s} \Delta_{s} \leq C \cdot H^{2} \sqrt{\frac{S \log (H S A / \delta)}{N}}$. Note (8) holds with probability $1-\delta / 2$, therefore plug above into (8) we obtain w.p. $1-\delta$,

$$
\begin{aligned}
& \left|\widehat{Q}_{h}^{\widehat{\pi}}-Q_{h}^{\widehat{\pi}}\right| \leq \sqrt{\frac{2 \log (4 H S A / \delta)}{N}} \sum_{t=h}^{H} \Gamma_{h+1: t}^{\widehat{\pi}} \sqrt{\operatorname{Var}_{P}\left(V_{t+1}^{\widehat{\pi}}\right)}+\sqrt{\frac{2 \log (4 H S A / \delta)}{N}} \sum_{t=h}^{H}\left\|\widehat{Q}_{t+1}^{\widehat{\pi}}-Q_{t+1}^{\widehat{\pi}}\right\|_{\infty} \\
& +C^{\prime} H^{4} \frac{S \log (H S A / \delta)}{N} \cdot \mathbf{1}+C_{1} \epsilon_{\mathrm{opt}} \cdot \sqrt{\frac{H^{2} S \log (S A / \delta)}{N}} \cdot \mathbf{1} \\
& \leq\left[\sqrt{\frac{C_{0} H^{3} \log (4 H S A / \delta)}{N}}+\sqrt{\frac{2 \log (4 H S A / \delta)}{N}} \sum_{t=h}^{H}\left\|\widehat{Q}_{t+1}^{\hat{\pi}}-Q_{t+1}^{\widehat{\pi}}\right\|_{\infty}+C^{\prime} H^{4} \frac{S \log (H S A / \delta)}{N}\right] \cdot \mathbf{1}
\end{aligned}
$$

where the last inequality uses Lemma $J$.5 and $\epsilon_{\text {opt }} \leq \sqrt{H / S}$ and renames $C^{\prime}=C^{\prime}+C_{1}$. Take $\|\cdot\|_{\infty}$ then obtain the result.
Lemma B.9. Given $N>0$. Define $C^{\prime \prime}:=2 \cdot \max \left(\sqrt{C_{0}}, C^{\prime}\right)$ where $C^{\prime}$ is the universal constant in Lemma B. 8 When $N \geq 8 H^{2} \log (4 H S A / \delta)$, then with probability $1-\delta, \forall h \in[H]$,

$$
\begin{align*}
&\left\|\widehat{Q}_{h}^{\widehat{\pi}}-Q_{h}^{\widehat{\pi}}\right\|_{\infty} \leq C^{\prime \prime} \sqrt{\frac{H^{3} \log (4 H S A / \delta)}{N}}+C^{\prime \prime} \frac{H^{4} S \log (H S A / \delta)}{N}  \tag{10}\\
&\left\|\widehat{Q}_{h}^{\pi^{\star}}-Q_{h}^{\pi^{\star}}\right\|_{\infty} \leq C^{\prime \prime} \sqrt{\frac{H^{3} \log (4 H S A / \delta)}{N}}+C^{\prime \prime} \frac{H^{4} S \log (H S A / \delta)}{N}
\end{align*}
$$

Proof. We prove by backward induction. For $h=H$, by Lemma B. 8

$$
\begin{aligned}
\left\|\widehat{Q}_{H}^{\widehat{\pi}}-Q_{H}^{\widehat{\pi}}\right\|_{\infty} & \leq \sqrt{\frac{C_{0} H^{3} \log (4 H S A / \delta)}{N}}+\sqrt{\frac{2 \log (4 H S A / \delta)}{N}}\left\|\widehat{Q}_{H+1}^{\widehat{\pi}}-Q_{H+1}^{\widehat{\pi}}\right\|_{\infty}+C^{\prime} H^{4} \frac{S \log (H S A / \delta)}{N} \\
& =\sqrt{\frac{C_{0} H^{3} \log (4 H S A / \delta)}{N}}+0+C^{\prime} H^{4} \frac{S \log (H S A / \delta)}{N} \\
& \leq C^{\prime \prime} \sqrt{\frac{H^{3} \log (4 H S A / \delta)}{N}}+C^{\prime \prime} H^{4} \frac{S \log (H S A / \delta)}{N}
\end{aligned}
$$

for general $h$, by condition we have $H \sqrt{\frac{2 \operatorname{log(4HSA/\delta )}}{N}} \leq 1 / 2$, therefore by Lemma B. 8

$$
\begin{aligned}
& \left\|\widehat{Q}_{h}^{\widehat{\pi}}-Q_{h}^{\widehat{\pi}}\right\|_{\infty} \leq \sqrt{\frac{C_{0} H^{3} \log (4 H S A / \delta)}{N}}+\sqrt{\frac{2 \log (4 H S A / \delta)}{N}} \sum_{t=h}^{H}\left\|\widehat{Q}_{t+1}^{\widehat{\pi}}-Q_{t+1}^{\widehat{\pi}}\right\|_{\infty}+C^{\prime} H^{4} \frac{S \log (H S A / \delta)}{N} \\
& \leq \sqrt{\frac{C_{0} H^{3} \log (4 H S A / \delta)}{N}}+H \sqrt{\frac{2 \log (4 H S A / \delta)}{N}} \max _{t+1}\left\|\widehat{Q}_{t+1}^{\widehat{\pi}}-Q_{t+1}^{\widehat{\pi}}\right\|_{\infty}+C^{\prime} H^{4} \frac{S \log (H S A / \delta)}{N} \\
& \leq \sqrt{\frac{C_{0} H^{3} \log (4 H S A / \delta)}{N}}+C^{\prime} H^{4} \frac{S \log (H S A / \delta)}{N} \\
& +\frac{1}{2}\left(C^{\prime \prime} \sqrt{\frac{H^{3} \log (4 H S A / \delta)}{N}}+C^{\prime \prime} \frac{H^{4} S \log (H S A / \delta)}{N}\right. \\
& \leq C^{\prime \prime} \sqrt{\frac{H^{3} \log (4 H S A / \delta)}{N}}+C^{\prime \prime} \frac{H^{4} S \log (H S A / \delta)}{N}
\end{aligned}
$$

The proof of the second claim is even easier since $\pi^{\star}$ is no longer a random policy and it is really just a non-uniform point-wise OPE. There are multiple ways to prove it and we leave it as an exercise to avoid redundancy: 1. Follow the same proving pipeline as $\left\|\widehat{Q}_{h}^{\hat{\pi}}-Q_{h}^{\hat{\pi}}\right\|_{\infty}$ used; 2. Mimic the procedure of point-wise OPE result in Lemma 3.4. in Yin et al. [2021a].

Remark B.10. Note the higher order term has dependence $H^{4} S$, which is somewhat unsatisfactory. We use the recursion-back trick to further reduce it to $H^{3.5} S^{0.5}$.
Lemma B.11. Given $N>0$. There exists universal constants $C_{1}, C_{2}$ such that when $N \geq$ $C_{1} H^{2} \log (H S A / \delta)$, then with probability $1-\delta, \forall h \in[H]$,

$$
\begin{equation*}
\left\|\widehat{Q}_{h}^{\hat{\pi}}-Q_{h}^{\widehat{\pi}}\right\|_{\infty} \leq C_{2} \sqrt{\frac{H^{3} \log (H S A / \delta)}{N}}+C_{2} \frac{H^{3} \sqrt{H S} \log (H S A / \delta)}{N} \tag{11}
\end{equation*}
$$

and

$$
\left\|\widehat{Q}_{h}^{\pi^{\star}}-Q_{h}^{\pi^{\star}}\right\|_{\infty} \leq C_{2} \sqrt{\frac{H^{3} \log (H S A / \delta)}{N}}+C_{2} \frac{H^{3} \sqrt{H S} \log (H S A / \delta)}{N}
$$

Proof. Note

$$
\begin{align*}
\widehat{V}_{t}^{\star}(s)-V_{t}^{\star}(s) & :=\widehat{V}_{t}^{\hat{\pi}^{\star}}(s)-V_{t}^{\pi^{\star}}(s) \\
& =\widehat{V}_{t}^{\widehat{\pi}^{\star}}(s)-V_{t}^{\widehat{\pi}^{\star}}(s)+V_{t}^{\widehat{\pi}^{\star}}(s)-V_{t}^{\pi^{\star}}(s)  \tag{12}\\
& \leq \widehat{V}_{t}^{\widehat{\pi}^{\star}}(s)-V_{t}^{\widehat{\pi}^{\star}}(s) \leq\left|\widehat{V}_{t}^{\widehat{\pi}^{\star}}(s)-V_{t}^{\widehat{\pi}^{\star}}(s)\right|
\end{align*}
$$

and similarly $V_{t}^{\star}(s)-\widehat{V}_{t}^{\star}(s) \leq\left|\widehat{V}_{t}^{\pi^{\star}}(s)-V_{t}^{\pi^{\star}}(s)\right|$, therefore by Lemma B. 9 (and use $\| \widehat{V}_{t}^{\pi}-$ $\left.V_{t}^{\pi}\left\|_{\infty} \leq\right\| \widehat{Q}_{t}^{\pi}-Q_{t}^{\pi} \|_{\infty}\right)$, with probability $1-\delta$,
$\Delta_{s} \leq 2 \cdot \sup _{t}\left\|V_{t}^{\star}-\widehat{V}_{t}^{\star}\right\| \leq 2 \max _{\widehat{\pi}^{\star}, \pi^{\star}} \sup _{t}\left\|\widehat{V}_{t}^{\pi}-V_{t}^{\pi}\right\|_{\infty} \leq C_{2} \sqrt{\frac{H^{3} \log (H S A / \delta)}{N}}+C_{2} \frac{H^{4} S \log (H S A / \delta)}{N}$,
where the second inequality uses (12). This replaces the crude bound of $O\left(\sqrt{H^{4} S \log (H S A / \delta) / N}\right)$ for $\max _{s} \Delta_{s}\left(\right.$ recall 9 ) by $O\left(\sqrt{H^{3} \log (H S A / \delta) / N}\right)$.
Plug this back to (8) and repeat the similar analysis we end up with (11). The second result is similarly proved.

## B. 7 Proof of Theorem 4.1

Proof of Theorem 4.1. Note $n_{s, a}=\sum_{i=1}^{n} \sum_{t=1}^{H} \mathbf{1}\left[s_{t}^{(i)}=s, a_{t}^{(i)}=a\right]$, which implies

$$
\mathbb{E}\left[n_{s, a}\right]=\mathbb{E}\left[\sum_{i=1}^{n} \sum_{t=1}^{H} \mathbf{1}\left[s_{t}^{(i)}=s, a_{t}^{(i)}=a\right]\right]=n \cdot \sum_{t=1}^{H} d_{t}^{\mu}(s, a)
$$

Or equivalently, $n_{s, a}$ follows $\operatorname{Binomial}\left(n, \sum_{t=1}^{H} d_{t}^{\mu}(s, a)\right)$. Then apply the first result of Lemma J. 1 by taking $\theta=1 / 2$, we have when $n>1 / d_{m} \cdot \log (H S A / \delta)^{12}$, then with probability $1-\delta$,

$$
n_{s, a} \geq \frac{1}{2} n \cdot \sum_{t=1}^{H} d_{t}^{\mu}(s, a), \quad \forall s \in \mathcal{S}, a \in \mathcal{A} .
$$

This further implies w.p. $1-\delta, n_{s, a} \geq \frac{1}{2} n \cdot \sum_{t=1}^{H} d_{t}^{\mu}(s, a)=\frac{1}{2} n \cdot H \cdot d^{\mu}(s, a) \geq \frac{1}{2} n H \cdot d_{m}$ and further ensures

$$
N:=\min _{s, a} n_{s, a} \geq \frac{1}{2} n H \cdot d_{m}
$$

Finally, apply above to Lemma B.11, we can get over with the condition on $N$ and obtain the stated result.

[^2]
## C Proof of minimax lower bound for model-based global uniform OPE

Proof of Theorem 3.1. In particular, we first focus on the case where $H=2$ and extend the result of $H=2$ to the general $H \geq 3$ at the end.

First of all, by Definition 2.1 let $\widehat{P}$ be the learned transition by certain model-based method. Since we assume $r_{h}$ is known and by convention $Q_{H+1}^{\pi}=0$ for any $\pi$, then by Bellman equation

$$
\widehat{Q}_{h}^{\pi}=r_{h}+\widehat{P}^{\pi_{h+1}} \widehat{Q}_{h+1}^{\pi}, \forall h \in[H] .
$$

In particular, $\widehat{Q}_{H+1}^{\pi}=Q_{H+1}^{\pi}=0$, and this implies

$$
\widehat{Q}_{H}^{\pi}=r_{H}+\widehat{P}^{\pi_{H+1}} \widehat{Q}_{H+1}^{\pi}=r_{H} ; \quad Q_{H}^{\pi}=r_{H}+P^{\pi_{H+1}} Q_{H+1}^{\pi}=r_{H}+0=r_{H}
$$

Now, again by definition of Bellman equation

$$
\begin{aligned}
\widehat{Q}_{H-1}^{\pi} & =r_{H-1}+\widehat{P}^{\pi_{H}} \widehat{Q}_{H}^{\pi}=r_{H-1}+\widehat{P}^{\pi_{H}} r_{H} \\
Q_{H-1}^{\pi} & =r_{H-1}+P^{\pi_{H}} Q_{H}^{\pi}=r_{H-1}+P^{\pi_{H}} r_{H}
\end{aligned}
$$

Therefore

$$
\begin{aligned}
& \sup _{\pi \in \Pi_{g}}\left\|\widehat{Q}_{H-1}^{\pi}-Q_{H-1}^{\pi}\right\|_{\infty}=\sup _{\pi \in \Pi_{g}}\left\|\left(\widehat{P}^{\pi_{H}}-P^{\pi_{H}}\right) r_{H}\right\|_{\infty} \\
= & \sup _{\pi \in \Pi_{g}}\left\|(\widehat{P}-P) r_{H}^{\pi_{H}}\right\|_{\infty}=\sup _{\pi \in \Pi_{g}} \sup _{s, a}\left|(\widehat{P}(\cdot \mid s, a)-P(\cdot \mid s, a)) r_{H}^{\pi_{H}}\right| \\
= & \sup _{s, a} \sup _{\pi \in \Pi_{g}}\left|(\widehat{P}(\cdot \mid s, a)-P(\cdot \mid s, a)) r_{H}^{\pi_{H}}\right|
\end{aligned}
$$

where $P^{\pi_{H}} \in \mathbb{R}^{S \cdot A \times S \cdot A}, r_{H} \in \mathbb{R}^{S \cdot A}, P \in \mathbb{R}^{S \cdot A \times S}$ and $r_{H}^{\pi_{H}} \in \mathbb{R}^{S}$. Note $A \geq 2$, so we can choose an instance of $r_{H}$ as (there are at least two actions since $A \geq 2$ )

$$
\left(r_{H}\left(s, a_{1}\right), r_{H}\left(s, a_{2}\right), \ldots\right):=(1,0, \ldots) \quad \forall s \in \mathcal{S}
$$

Above implies: if $\pi_{H}(s)=a_{1}$, then $r_{H}^{\pi_{H}}(s)=1$; if $\pi_{H}(s)=a_{2}$, then $r_{H}^{\pi_{H}}(s)=0 ; \ldots$
Hence, if $\Pi_{g}$ is the global deterministic policy class, then $r_{H}^{\pi_{H}}$ can traverse all the $S$-dimensional vectors with either 0 or 1 in each coordinate, which is exactly

$$
\left\{r_{H}^{\pi_{H}} \in \mathbb{R}^{S}: \pi_{H} \in \Pi_{g}\right\} \supset\{0,1\}^{S}
$$

Now let us first consider fixed $s, a$. Then with this choice of $r$, above implies

$$
\begin{aligned}
& \sup _{\pi \in \Pi_{g}}\left|(\widehat{P}(\cdot \mid s, a)-P(\cdot \mid s, a)) r_{H}^{\pi_{H}}\right| \geq \sup _{r \in\{0,1\}^{S}}|(\widehat{P}(\cdot \mid s, a)-P(\cdot \mid s, a)) \cdot r| \\
= & \sup _{r \in\{0,1\}^{S}}\left|\sum_{i: r_{i}=1}\left(\widehat{P}\left(s_{i} \mid s, a\right)-P\left(s_{i} \mid s, a\right)\right)\right|
\end{aligned}
$$

Let $I_{+}:=\left\{i \in[S]\right.$ : s.t. $\left.\widehat{P}\left(s_{i} \mid s, a\right)-P\left(s_{i} \mid s, a\right)>0\right\}$ be the set of indices where $\widehat{P}\left(s_{i} \mid s, a\right)-$ $P\left(s_{i} \mid s, a\right)$ are positive and $I_{-}:=\left\{i \in[S]: s . t . \widehat{P}\left(s_{i} \mid s, a\right)-P\left(s_{i} \mid s, a\right)<0\right\}$ be the set of indices where $\widehat{P}\left(s_{i} \mid s, a\right)-P\left(s_{i} \mid s, a\right)$ are negative, then we further have

$$
\begin{aligned}
& \sup _{r \in\{0,1\}^{S}}\left|\sum_{i: r_{i}=1}\left(\widehat{P}\left(s_{i} \mid s, a\right)-P\left(s_{i} \mid s, a\right)\right)\right| \\
\geq & \max \left\{\left|\sum_{i \in I+}\left[\widehat{P}\left(s_{i} \mid s, a\right)-P\left(s_{i} \mid s, a\right)\right]\right|,\left|\sum_{i \in I-}\left[\widehat{P}\left(s_{i} \mid s, a\right)-P\left(s_{i} \mid s, a\right)\right]\right|\right\} \\
= & \max \left\{\sum_{i \in I+}\left|\widehat{P}\left(s_{i} \mid s, a\right)-P\left(s_{i} \mid s, a\right)\right|, \sum_{i \in I-}\left|\widehat{P}\left(s_{i} \mid s, a\right)-P\left(s_{i} \mid s, a\right)\right|\right\}
\end{aligned}
$$

On the other hand, we have

$$
\sum_{i \in I+}\left|\widehat{P}\left(s_{i} \mid s, a\right)-P\left(s_{i} \mid s, a\right)\right|+\sum_{i \in I-}\left|\widehat{P}\left(s_{i} \mid s, a\right)-P\left(s_{i} \mid s, a\right)\right|=\|\widehat{P}(\cdot \mid s, a)-P(\cdot \mid s, a)\|_{1}
$$

since $\widehat{P}\left(s_{i} \mid s, a\right)-P\left(s_{i} \mid s, a\right)=0$ contributes nothing to the $l_{1}$ norm. Combine all the steps together, we obtain

$$
\begin{equation*}
\sup _{\pi \in \Pi_{g}}\left\|\widehat{Q}_{H-1}^{\pi}-Q_{H-1}^{\pi}\right\|_{\infty} \geq \sup _{s, a} \frac{1}{2}\|\widehat{P}(\cdot \mid s, a)-P(\cdot \mid s, a)\|_{1} \underset{\substack{1}}{\geq} c \cdot \sup _{s, a} \sqrt{\frac{S}{n_{s, a}}} \geq c^{\prime} \sqrt{\frac{S}{n d_{m}}} \tag{13}
\end{equation*}
$$

holds with constant probability $p$. Here $n_{s, a}=\sum_{h=1}^{H} \sum_{i=1}^{n} \mathbf{1}\left[s_{h}^{(i)}=s, a_{h}^{(i)}=a\right]$ is the number of data pieces visited $(s, a)$ in $n$ episodes. Now we explain how to obtain (1) and (2). In particular, we first explain (2).

Explain (2. Recall we consider the case $H=2$. Then

$$
\mathbb{E}\left[n_{s, a}\right]=\mathbb{E}\left[\sum_{h=1}^{H} \sum_{i=1}^{n} \mathbf{1}\left[s_{h}^{(i)}=s, a_{h}^{(i)}=a\right]\right]=n \sum_{i=1}^{2} \mathbb{E}\left[\mathbf{1}\left[s_{h}^{(1)}=s, a_{h}^{(1)}=a\right]\right]=n \sum_{h=1}^{2} d_{h}^{\mu}(s, a)
$$

i.e. $n_{s, a}$ is a Binomial random variable with parameter $n$ and $\sum_{h=1}^{2} d_{h}^{\mu}(s, a)$. Then by Lemma J. 1 . choose $\theta=\frac{1}{2}$, apply the second result, we obtain when $n>\left(1 / 2 d_{m}\right) \cdot \log (S A / \delta)^{13}$ with probability $1-\delta$

$$
n_{s, a} \leq \frac{3}{2} n \cdot \sum_{h=1}^{2} d_{h}^{\mu}(s, a), \quad \forall s, a
$$

Next, similar to the lower bound proof (Theorem G.2.) of Yin et al. [2021a], we can choose $\mu$ and $M$ (near uniform but not exact uniform) such that $d_{h}^{\mu}(s, a) \leq C \cdot d_{m}$, which further implies $n_{s, a} \leq C \cdot n \cdot d_{m}, \forall s, a$. Summarize above we end up with the following Lemma:

Lemma C.1. Suppose $n \geq\left(1 / 2 d_{m}\right) \cdot \log (S A / \delta)$, then

$$
\sup _{\mu, M} \mathbb{P}\left[\sqrt{\frac{1}{n_{s, a}}} \geq C \cdot \sqrt{\frac{1}{n \cdot d_{m}}}, \forall s, a\right] \geq 1-\delta
$$

Explain (1). To make the explanation rigorous, we first fix a pair $(s, a)$ and conditional on $n_{s, a}$. Then by a direct translation of Lemma J.7, we have

$$
\inf _{\widehat{P}} \sup _{P(\cdot \mid s, a) \in \mathcal{M}_{S}} \mathbb{P}\left[\left.\|\widehat{P}(\cdot \mid s, a)-P(\cdot \mid s, a)\|_{1} \geq \frac{1}{8} \sqrt{\frac{e S}{2 n_{s, a}}}-o(\cdot) \right\rvert\, n_{s, a} \geq \frac{e}{32} S\right] \geq p
$$

where $o(\cdot)$ is some exponentially small term in $S, n$. Now we consider everything under the condition $n \geq \frac{e}{32} \cdot S / d_{m} \log (S A / \delta)$. Next again take $\theta=1 / 2$, then by the first result of Lemma J.1, with probability $1-\delta$,

$$
n_{s, a} \geq \frac{1}{2} n \cdot \sum_{h=1}^{2} d_{h}^{\mu}(s, a) \geq n \cdot d_{m} \geq \frac{e}{32} S \log (S A / \delta)
$$

[^3]where the last inequality uses the condition $n \geq \frac{e}{32} \cdot S / d_{m} \log (S A / \delta)$. Therefore this implies
\[

$$
\begin{aligned}
& \inf _{\widehat{P}} \sup _{P(\cdot \mid s, a) \in \mathcal{M}_{S}} \mathbb{P}\left[\|\widehat{P}(\cdot \mid s, a)-P(\cdot \mid s, a)\|_{1} \geq \frac{1}{8} \sqrt{\frac{e S}{2 n_{s, a}}}-o(\cdot)\right] \\
= & \inf _{\widehat{P}} \sup _{P(\cdot \mid s, a) \in \mathcal{M}_{S}}\left(\mathbb{P}\left[\left.\|\widehat{P}(\cdot \mid s, a)-P(\cdot \mid s, a)\|_{1} \geq \frac{1}{8} \sqrt{\frac{e S}{2 n_{s, a}}}-o(\cdot) \right\rvert\, n_{s, a} \geq \frac{e}{32} S\right] \cdot \mathbb{P}\left[n_{s, a} \geq \frac{e}{32} S\right]\right. \\
+ & \left.\mathbb{P}\left[\left.\|\widehat{P}(\cdot \mid s, a)-P(\cdot \mid s, a)\|_{1} \geq \frac{1}{8} \sqrt{\frac{e S}{2 n_{s, a}}}-o(\cdot) \right\rvert\, n_{s, a} \leq \frac{e}{32} S\right] \cdot \mathbb{P}\left[n_{s, a} \leq \frac{e}{32} S\right]\right) \\
\geq & \inf _{\widehat{P}} \sup _{P(\cdot \mid s, a) \in \mathcal{M}_{S}} \mathbb{P}\left[\left.\|\widehat{P}(\cdot \mid s, a)-P(\cdot \mid s, a)\|_{1} \geq \frac{1}{8} \sqrt{\frac{e S}{2 n_{s, a}}}-o(\cdot) \right\rvert\, n_{s, a} \geq \frac{e}{32} S\right] \cdot \mathbb{P}\left[n_{s, a} \geq \frac{e}{32} S\right] \\
\geq & p \cdot(1-\delta),
\end{aligned}
$$
\]

To sum up, we have the following lemma:
Lemma C.2. Let $n \geq \frac{e}{32} S / d_{m} \cdot \log (S A / \delta)$, then there exists a $0<p<1$,

$$
\inf _{\widehat{P}} \sup _{P(\cdot \mid s, a) \in \mathcal{M}_{S}} \mathbb{P}\left[\|\widehat{P}(\cdot \mid s, a)-P(\cdot \mid s, a)\|_{1} \geq \frac{1}{8} \sqrt{\frac{e S}{2 n_{s, a}}}-o(\cdot)\right] \geq p \cdot(1-\delta) .
$$

Now we finish the proof for the case where $H=2$. First note by (13),

$$
\sup _{\pi \in \Pi_{g}}\left\|\widehat{Q}_{H-1}^{\pi}-Q_{H-1}^{\pi}\right\|_{\infty} \geq \sup _{s, a} \frac{1}{2}\|\widehat{P}(\cdot \mid s, a)-P(\cdot \mid s, a)\|_{1}
$$

with probability 1 , therefore by (13), Lemma C.1, Lemma C. 2 we have

$$
\inf _{\widehat{P}} \sup _{P \in \mathcal{M}_{S}} \mathbb{P}\left[\sup _{\pi \in \Pi_{g}}\left\|\widehat{Q}_{H-1}^{\pi}-Q_{H-1}^{\pi}\right\|_{\infty} \geq C \cdot \sqrt{\frac{S}{n d_{m}}}\right] \geq p(1-\delta)-\delta
$$

when $n \geq c \cdot S / d_{m} \log (S A / \delta)$ for some $c \geq \frac{e}{32}$. Above holds for any $\delta$.
It is easy to check $\frac{3}{2} \frac{p}{1+p} \leq 1$, therefore, in particular we set $\delta=\frac{3}{2} \frac{p}{1+p}$, direct calculation shows

$$
p(1-\delta)-\delta=\frac{p}{2}
$$

which completes the proof for $H=2$.
Extend to the general $H \geq 3$.
Step 1. Similar to the decomposition in section B.4, we also have:

$$
\hat{Q}_{t}^{\pi}-Q_{t}^{\pi}=\sum_{h=t}^{H} \hat{\Gamma}_{t+1: h}^{\pi}(\widehat{P}-P) V_{h+1}^{\pi}
$$

Step 2. Now choosing rewards recursively from back (with $\left\|r_{H}\right\|_{\infty}=c$ sufficiently small) such that $1 \geq r_{h} \geq\left(\left\|r_{h+1}\right\|_{\infty}+\ldots+\left\|r_{H}\right\|_{\infty}\right)$ element-wisely $\forall h$, and $\max _{s, a} r_{h}(s, a)=3 \min _{s, a} r_{h}(s, a)$. We denote $r_{h, \max }:=\max _{s, a} r_{h}(s, a)$ and $r_{h, \min }:=\min _{s, a} r_{h}(s, a)$. This choice guarantees:

$$
r_{h, \min }:=\min _{s, a} r_{h}(s, a)>\left\|P^{\pi_{h+1}} r_{h+1}+. .+P^{\pi_{h+1: H}} r_{H}\right\|_{\infty}
$$

since $P^{\pi_{h}}$ is row-stochastic.
Step 3. Next note $V_{h}^{\pi}=r_{h}+P^{\pi_{h+1}} r_{h+1}+. .+P^{\pi_{h+1: H}} r_{H}$, so set $\left(r_{h}\left(s, a_{1}\right), r_{h}\left(s, a_{2}\right), \ldots\right):=$ $\left(\max _{s, a} r_{h}(s, a), \min _{s, a} r_{h}(s, a), \ldots\right)$, then choose $\pi_{h}$ similar to the $H=2$ case and use Step 1 and
Step 2 we have

$$
\begin{aligned}
\left|\left(\hat{P}_{s, a}-P_{s, a}\right) V_{h}^{\pi}\right| & \geq \frac{1}{2}\left\|\hat{P}_{s, a}-P_{s, a}\right\|_{1} \cdot\left(r_{h, \text { max }}-r_{h, \text { min }}-\left(P^{\pi_{h+1}} r_{h+1}+. .+P^{\pi_{h+1: H}} r_{H}\right)\right) \\
& \geq \frac{1}{2}\left\|\hat{P}_{s, a}-P_{s, a}\right\|_{1} \cdot r_{h, \min } \geq \frac{1}{2}\left\|\hat{P}_{s, a}-P_{s, a}\right\|_{1} \cdot c
\end{aligned}
$$

where the reasoning of the first inequality is similar to the case of $H=2$. Next use $\hat{\Gamma}_{t+1: h}^{\pi}$ is row-stochastic then from Step 1 and take the sum we have

$$
\left\|\hat{Q}_{1}^{\pi}-Q_{1}^{\pi}\right\|_{\infty} \geq \frac{1}{2} c \cdot H \min _{s, a}\left\|\hat{P}_{s, a}-P_{s, a}\right\|_{1}
$$

for such choice of rewards and $\pi$.
Step 4. However, in the above construction $c$ actually depends on $H$ due to the design $1 \geq r_{h} \geq$ $\left(\left\|r_{h+1}\right\|_{\infty}+\ldots+\left\|r_{H}\right\|_{\infty}\right)$. To get a universal constant $c$ we could use the bound $\left\|\hat{Q}_{1}^{\pi}-Q_{1}^{\pi}\right\|_{\infty} \gtrsim$ $r_{\frac{H}{2}, \text { min }} \cdot \frac{H}{2} \min _{s, a}\left\|\hat{P}_{s, a}-P_{s, a}\right\|_{1}$ instead, where $r_{\frac{H}{2}, \text { min }}$ in Step 2 is universally lower bounded. Then we apply $\left\|\hat{P}_{s, a}-P_{s, a}\right\|_{1} \gtrsim \Omega\left(\sqrt{S / n d_{m}}\right)$ to obtain the lower bound $\Omega\left(\sqrt{H^{2} S / n d_{m}}\right)$.

Remark C.3. We point out while our lower bound of $\Omega\left(H^{2} S / d_{m} \epsilon^{2}\right)$ for uniform OPE appears to be qualitatively similar to the lower bound of $\Omega\left(H^{2} S^{2} A / \epsilon^{2}\right)$ derived for the online reward-free RL setting [Jin et al. 2020a], our result is not implied by theirs and cannot be proven by directly adapting their construction. Those two results are in principle different since: the result in [Jin et al., 2020a] is learning-oriented where they define the problem class on $O(S)$ states and forcing $\Omega\left(S A / \epsilon^{2}\right)$ episodes in each state and end up with $O\left(S^{2} A / \epsilon^{2}\right)$ complexity; our result is evaluation-oriented where we need reduce the uniform evaluation problem to estimating probability distribution in $\ell_{1}$-error. The global uniform OPE and the reward-free setting are also different tasks (one cannot imply the other): the former deals with uniform convergence over all policies but with a fixed reward while the latter aims at learning simultaneously over all rewards.

## D Proof for optimal offline learning (Corollary 4.2)

Proof. This is a corollary of Theorem4.1. Indeed, by taking $\widehat{\pi}=\widehat{\pi}^{\star}$, we first have

$$
\left\|\widehat{V}_{1}^{\widehat{\pi}^{\star}}-V_{1}^{\widehat{\pi}^{\star}}\right\|_{\infty} \leq\left\|\widehat{Q}_{1}^{\widehat{\pi}^{\star}}-Q_{1}^{\widehat{\pi}^{\star}}\right\|_{\infty} \leq C\left[\sqrt{\frac{H^{2} \iota}{n d_{m}}}+\frac{H^{2.5} S^{0.5} \iota}{n d_{m}}\right]
$$

Similar to the second result in Lemma B.11 we also have

$$
\left\|\widehat{V}_{1}^{\pi^{\star}}-V_{1}^{\pi^{\star}}\right\|_{\infty} \leq\left\|\widehat{Q}_{1}^{\pi^{\star}}-Q_{1}^{\pi^{\star}}\right\|_{\infty} \leq C\left[\sqrt{\frac{H^{2} \iota}{n d_{m}}}+\frac{H^{2.5} S^{0.5} \iota}{n d_{m}}\right]
$$

Next, recall the definition of $\widehat{\pi} \in \Pi_{l}$ that

$$
\left\|\widehat{V}_{1}^{\widehat{\pi}^{\star}}-\widehat{V}_{1}^{\hat{\pi}}\right\|_{\infty} \leq \epsilon_{\mathrm{opt}}
$$

and Theorem 4.1 again that

$$
\left\|\widehat{V}_{1}^{\widehat{\pi}}-V_{1}^{\widehat{\pi}}\right\|_{\infty} \leq\left\|\widehat{Q}_{1}^{\widehat{\pi}}-Q_{1}^{\widehat{\pi}}\right\|_{\infty} \leq C\left[\sqrt{\frac{H^{2} \iota}{n d_{m}}}+\frac{H^{2.5} S^{0.5} \iota}{n d_{m}}\right]
$$

Therefore

$$
\begin{aligned}
V_{1}^{\pi^{\star}}-V_{1}^{\widehat{\pi}} & =V_{1}^{\pi^{\star}}-V_{1}^{\widehat{\pi}^{\star}}+V_{1}^{\widehat{\pi}^{\star}}-V_{1}^{\widehat{\pi}} \\
& \leq \max _{\widehat{\pi}^{\star}, \pi^{\star}}\left\|\widehat{V}_{1}^{\pi}-V_{1}^{\pi}\right\|_{\infty}+V_{1}^{\widehat{\pi}^{\star}}-V_{1}^{\widehat{\pi}} \\
& =\max _{\widehat{\pi}^{\star}, \pi^{\star}}\left\|\widehat{V}_{1}^{\pi}-V_{1}^{\pi}\right\|_{\infty}+\left(V_{1}^{\widehat{\pi}^{\star}}-\widehat{V}_{1}^{\widehat{\pi}^{\star}}\right)+\left(\widehat{V}_{1}^{\widehat{\pi}^{\star}}-\widehat{V}_{1}^{\widehat{\pi}}\right)+\left(\widehat{V}_{1}^{\widehat{\pi}}-V_{1}^{\widehat{\pi}}\right) \\
& \leq 3 C\left[\sqrt{\frac{H^{2} \iota}{n d_{m}}}+\frac{H^{2.5} S^{0.5} \iota}{n d_{m}}\right]+\left\|\widehat{V}_{1}^{\widehat{\pi}^{\star}}-\widehat{V}_{1}^{\widehat{\pi}}\right\|_{\infty} \cdot \mathbf{1} \\
& \leq 3 C\left[\sqrt{\frac{H^{2} \iota}{n d_{m}}}+\frac{H^{2.5} S^{0.5} \iota}{n d_{m}}\right]+\epsilon_{\mathrm{opt}} \cdot \mathbf{1} .
\end{aligned}
$$

This completes the proof.

## E Proof for optimal offline Task-agnostic learning (Theorem 5.3)

Proof. Recall the definition of offline task-agnostic setting, where $K$ tasks corresponds to $K$ MDPs $M_{k}=\left(\mathcal{S}, \mathcal{A}, P, r_{k}, H, d_{1}\right)$ with different mean reward functions $r_{k}$ 's. Since the incremental number of rewards do not incur randomness, therefore by Corollary 4.2 , choose $\widehat{\pi}_{k}=\widehat{\pi}_{k}^{\star}$ and apply a union bound we obtain with probability $1-\delta$,

$$
\begin{aligned}
\sup _{k \in[K]}\left\|V_{1, M_{k}}^{\star}-V_{1, M_{k}}^{\widehat{\pi}_{k}^{\star}}\right\|_{\infty} & \leq O\left[\sqrt{\frac{H^{2} \log (H S A K / \delta)}{n d_{m}}}+\frac{H^{2.5} S^{0.5} \log (H S A K / \delta)}{n d_{m}}\right] \\
& =O\left[\sqrt{\frac{H^{2}(\iota+\log (K))}{n d_{m}}}+\frac{H^{2.5} S^{0.5}(\iota+\log (K))}{n d_{m}}\right]
\end{aligned}
$$

which completes the proof.

Remark E.1. We stress that Section 3 of Zhang et al. [2020b] claims the definition of task-agnostic $R L$ setting embraces one challenge that $r_{k}^{(i)}$ 's are the observed random realizations and the need to accurately estimate mean rewards $r_{k}$ 's causes the additional $\log (K)$ dependence. However, for offline case, this is not essential since, by straightforward calculation, estimating $r_{k}^{(i)}$ 's accurately only requires $\tilde{O}\left(\log (K) / d_{m} \epsilon^{2}\right)$ samples, which is of lower order comparing to $\tilde{O}\left(H^{2} \log (K) / d_{m} \epsilon^{2}\right)$ learning bound. Therefore, in Definition 5.1] we do not incorporate the random version statement for reward $r_{k}$.

## E. 1 Offline Learning in the Constrained MDPs (CMDP)

Recently, there is a line of studies in the Constrained Markov Decision Processes (CMDP) (e.g. ?), where the $\operatorname{MDP} M=\left(\mathcal{S}, \mathcal{A}, P, H, d_{1}\right)$. When the reward is set to be $r$, it defines the objective function $V_{r}^{\pi}$ and there is another utility function $g$ that defines the constraint. To be concrete, the objective formualted as:

$$
\begin{equation*}
\underset{\pi \in \Delta(\mathcal{A} \mid \mathcal{S}, H)}{\operatorname{maximize}} V_{r, 1}^{\pi}\left(x_{1}\right) \text { subject to } V_{g, 1}^{\pi}\left(x_{1}\right) \geq b \tag{14}
\end{equation*}
$$

where $b \in(0, H]$ is some constraint threshold. In addition, the formulation needs a Slater condition that: there exists $\gamma>0$ and $\bar{\pi} \in \Delta(\mathcal{A} \mid \mathcal{S}, H)$ such that $V_{g, 1}^{\bar{\pi}}\left(x_{1}\right) \geq b+\gamma$.

Let $\pi^{\star}$ be the optimal solution that is compatible with the programming (14) (note this is different from the optimal policy that maximizes $V_{r, 1}^{\pi}$ only), then by feasibility it satisfies $V_{g, 1}^{\pi^{\star}} \geq b$.
Now let $\hat{\pi}^{\star}$ be the solution of the empirical program:

$$
\begin{equation*}
\underset{\pi \in \Delta(\mathcal{A} \mid \mathcal{S}, H)}{\operatorname{maximize}} \widehat{V}_{r, 1}^{\pi}\left(x_{1}\right) \text { subject to } \widehat{V}_{g, 1}^{\pi}\left(x_{1}\right) \geq b \tag{15}
\end{equation*}
$$

then we can show $\hat{\pi}^{\star}$ is a near-optimal solution for via the local uniform convergence guarantee (Theorem4.1).
Indeed, define a surrogate program:

$$
\begin{equation*}
\underset{\pi \in \Delta(\mathcal{A} \mid \mathcal{S}, H)}{\operatorname{maximize}} \widehat{V}_{r, 1}^{\pi}\left(x_{1}\right) \text { subject to } V_{g, 1}^{\pi}\left(x_{1}\right) \geq b \tag{16}
\end{equation*}
$$

and let $\bar{\pi}^{\star}$ be the solution for (16). Then apparently $\bar{\pi}^{\star}$ satisfies $V_{g, 1}^{\bar{\pi}^{\star}}\left(x_{1}\right) \geq b$. Moreover, we have

$$
\begin{aligned}
V_{r, 1}^{\pi^{\star}}-V_{r, 1}^{\bar{\pi}^{\star}} & =V_{r, 1}^{\pi^{\star}}-\widehat{V}_{r, 1}^{\pi^{\star}}+\widehat{V}_{r, 1}^{\pi^{\star}}-\widehat{V}_{r, 1}^{\bar{\pi}^{\star}}+\widehat{V}_{r, 1}^{\bar{\pi}^{\star}}-V_{r, 1}^{\bar{\pi}^{\star}} \\
& \leq V_{r, 1}^{\pi^{\star}}-\widehat{V}_{r, 1}^{\pi^{\star}}+0+\widehat{V}_{r, 1}^{\bar{\pi}^{\star}}-V_{r, 1}^{\bar{\pi}^{\star}} \\
& \leq 2 \sup _{\pi}\left|V_{r, 1}^{\pi}-\widehat{V}_{r, 1}^{\pi}\right|
\end{aligned}
$$

On the other hand, by local uniform convergence guarantee, $\left|V_{g, 1}^{\pi}-\widehat{V}_{g, 1}^{\pi}\right| \leq \tilde{O}\left(\sqrt{H^{2} / n d_{m}}\right)$ for all $\pi$ in the $\sqrt{H / S}$-neighborhood of $\hat{\pi}^{\star}$ (w.r.t $g$ ). This implies

$$
V_{r, 1}^{\pi^{\star}}-V_{r, 1}^{\hat{\pi}^{\star}} \leq 2 \sup _{\pi}\left|V_{r, 1}^{\pi}-\widehat{V}_{r, 1}^{\pi}\right|+\tilde{O}\left(\sqrt{H^{2} / n d_{m}}\right)
$$

and the violation of the constraint is bounded by $\tilde{O}\left(\sqrt{H^{2} / n d_{m}}\right)$. This means any approach that solves (15) is near-optimal for the original constrained MDP task given the uniform convergence guarantee.

## F Proof for optimal offline Reward-free learning (Theorem 5.4)

Similar to before, recall $n_{s, a}=\sum_{h=1}^{H} \sum_{i=1}^{n} \mathbf{1}\left[s_{h}^{(i)}=s, a_{h}^{(i)}=a\right]$. We first prove two lemmas which essentially provide a version of "Maximal Bernstein inequality". We first fix a pair $(s, a)$ and then conditional on $n_{s, a}$.
Lemma F.1. We define $\epsilon_{1}=\sqrt{\frac{1}{H S^{2}}}$. Let $\mathcal{G}=\left\{\left[i_{1} \epsilon_{1}, i_{2} \epsilon_{1}, \ldots, i_{S} \epsilon_{1}\right]^{\top} \mid i_{1}, i_{2}, \ldots, i_{S} \in \mathbb{Z}\right\} \cap[0, H]^{S}$ be the $S$-dimensional grid. Next define $\iota_{1}=\log \left[\left(\sqrt{H^{3} S^{2}}\right)^{S} / \delta\right]$. Then with probability $1-\delta$,

$$
\left|\left(P_{s, a}-\widehat{P}_{s, a}\right) w\right| \leq \sqrt{\frac{2 \operatorname{Var}_{s, a}(w) \iota_{1}}{n_{s, a}}}+\frac{2 H \iota_{1}}{3 n_{s, a}}, \quad \forall w \in \mathcal{G} .
$$

This is by the direct application of Bernstein inequality with a union bound, where the cardinality of $\mathcal{G}$ is

$$
\left(\frac{H}{\epsilon_{1}}\right)^{S}=\left(\sqrt{H^{3} S^{2}}\right)^{S}
$$

Lemma F.2. Let the $S$-dimensional grid be $\mathcal{G}=\left\{\left[i_{1} \epsilon_{1}, i_{2} \epsilon_{1}, \ldots, i_{S} \epsilon_{1}\right]^{\top} \mid i_{1}, i_{2}, \ldots, i_{S} \in \mathbb{Z}\right\} \cap$ $[0, H]^{S}$ and define $\iota_{1}=\log \left[\left(\sqrt{H^{3} S^{2}}\right)^{S} / \delta\right]$. It holds with probability $1-\delta$,

$$
\left|\left(P_{s, a}-\widehat{P}_{s, a}\right) v\right| \leq \sqrt{\frac{2 \operatorname{Var}_{s, a}(v) \iota_{1}}{n_{s, a}}}+C \sqrt{\frac{\iota_{1}}{n_{s, a} H S}}+\frac{2 H \iota_{1}}{3 n_{s, a}}, \quad \forall v \in[0, H]^{S}
$$

Proof. Let $z:=\operatorname{Proj}_{\mathcal{G}}(v)$. Then by design of $\mathcal{G}$ we have

$$
\|z-v\|_{\infty} \leq \epsilon_{1}=\sqrt{\frac{1}{H S^{2}}}
$$

Therefore we obtain $\forall v \in[0, H]^{S}$,

$$
\begin{aligned}
\left|\left(P_{s, a}-\widehat{P}_{s, a}\right) v\right| & \leq\left|\left(P_{s, a}-\widehat{P}_{s, a}\right)(v-z)\right|+\left|\left(P_{s, a}-\widehat{P}_{s, a}\right) z\right| \\
& \leq\left\|P_{s, a}-\widehat{P}_{s, a}\right\|_{1}\|z-v\|_{\infty}+\left|\left(P_{s, a}-\widehat{P}_{s, a}\right) z\right| \\
& \leq c \sqrt{\frac{S}{n_{s, a}}}\|z-v\|_{\infty}+\sqrt{\frac{2 \operatorname{Var}_{s, a}(z) \iota_{1}}{n_{s, a}}}+\frac{2 H \iota_{1}}{3 n_{s, a}} \\
& \leq c \sqrt{\frac{S}{n_{s, a}}}\|z-v\|_{\infty}+\sqrt{\frac{2\|z-v\|_{\infty}^{2} \iota_{1}}{n_{s, a}}}+\sqrt{\frac{2 \operatorname{Var}_{s, a}(v) \iota_{1}}{n_{s, a}}}+\frac{2 H \iota_{1}}{3 n_{s, a}} \\
& \leq C \sqrt{\frac{S \iota_{1}}{n_{s, a}}}\|z-v\|_{\infty}+\sqrt{\frac{2 \operatorname{Var}_{s, a}(v) \iota_{1}}{n_{s, a}}}+\frac{2 H \iota_{1}}{3 n_{s, a}} \\
& \leq C \sqrt{\frac{\iota_{1}}{n_{s, a} H S}}+\sqrt{\frac{2 \operatorname{Var}_{s, a}(v) \iota_{1}}{n_{s, a}}}+\frac{2 H \iota_{1}}{3 n_{s, a}}
\end{aligned}
$$

where the third inequality uses Lemma F. 1 and Lemma J. 9.

Then recall $N:=\min _{s, a} n_{s, a}$, by Lemma F. 2 and a union bound we obtain with probability $1-\delta$, element-wisely,

$$
\begin{equation*}
|(P-\widehat{P}) v| \leq C \cdot\left(\sqrt{\frac{2 \operatorname{Var}_{s, a}(v) \iota_{2}}{N}}+2 \sqrt{\frac{\iota_{2}}{N \cdot H S}}+\frac{2 H \iota_{2}}{3 N}\right) \cdot \mathbf{1}, \quad \forall v \in[0, H]^{S}, \tag{17}
\end{equation*}
$$

where $\iota_{2}=S \log (H S A / \delta)$.
Remark F.3. Equation 17 is a form of maximal Bernstein inequality as it keeps validity for all $v \in[0, H]^{S}$. The price for this stronger result is the extra $S$ factor (coming from $\iota_{2}$ ) in the dominate term.

Now, for any reward $r$, by (empirical) Bellman equation we have element-wisely:

$$
\begin{aligned}
\widehat{Q}_{h}^{\widehat{\pi}^{\star}}-Q_{h}^{\widehat{\pi}^{\star}} & =r_{h}+\widehat{P}^{\widehat{\pi}_{h+1}^{\star}} \widehat{Q}_{h+1}^{\pi^{\star}}-r_{h}-P^{\widehat{\pi}_{h+1}^{\star}} Q_{h+1}^{\widehat{\pi}^{\star}} \\
& =\left(\widehat{P}^{\pi_{n+1}^{\star}}-P^{\widehat{\pi}_{h+1}^{\star}}\right) \widehat{Q}_{h+1}^{\widehat{\pi}^{\star}}+P^{\widehat{\pi}_{h+1}^{\star}}\left(\widehat{Q}_{h+1}^{\pi^{\star}}-Q_{h+1}^{\widehat{\pi}^{\star}}\right) \\
& =(\widehat{P}-P) \widehat{V}_{h+1}^{\widehat{\pi}^{\star}}+P^{\widehat{\pi}_{h+1}^{\star}}\left(\widehat{Q}_{h+1}^{\widehat{\pi}^{\star}}-Q_{h+1}^{\widehat{\pi}^{\star}}\right) \\
& =\ldots=\sum_{t=h}^{H} \Gamma_{h+1: t}^{\widehat{\pi}^{\star}}(\widehat{P}-P) \widehat{V}_{t+1}^{\widehat{\pi}^{\star}}
\end{aligned}
$$

where $\Gamma_{h+1: t}^{\pi}=\prod_{i=h+1}^{t} P^{\pi_{i}}$ is multi-step state-action transition and $\Gamma_{h+1: h}:=I$.

Concentration on $(\widehat{P}-P) \widehat{V}_{h}^{\star}$. Now by (17), we have the following:

$$
\left.\begin{array}{l}
\left(\widehat{P}_{s, a}-P_{s, a}\right) \widehat{V}_{h}^{\star} \\
\leq C \cdot\left(\sqrt{\frac{2 \operatorname{Var}_{s, a}\left(\widehat{V}_{h}^{\star}\right) \iota_{2}}{N}}+2 \sqrt{\frac{\iota_{2}}{N \cdot H S}}+\frac{2 H \iota_{2}}{3 N}\right) \\
\leq C \cdot\left(\sqrt{\frac{2 \operatorname{Var}_{s, a}\left(V_{h}^{\widehat{\pi}^{\star}}\right) \iota_{2}}{N}}+2 \sqrt{\frac{\iota_{2}}{N \cdot H S}}+\sqrt{\frac{2 \iota_{2}}{N}} \cdot\left\|\widehat{V}_{h}^{\widehat{\pi}^{\star}}-V_{h}^{\widehat{\pi}^{\star}}\right\|_{\infty}+\frac{2 H \iota_{2}}{3 N}\right. \tag{18}
\end{array}\right)
$$

where the third inequality uses Lemma $J$ J.10 $\left.\right|^{44}$. Then above implies

$$
\begin{aligned}
& \widehat{Q}_{h}^{\hat{\pi}^{\star}}-Q_{h}^{\widehat{\pi}^{\star}} \\
& \leq C^{\prime} \sum_{t=h}^{H} \Gamma_{h+1: t}^{\widehat{\pi}^{\star}} \cdot\left(\sqrt{\frac{2 \operatorname{Var}_{s, a}\left(V_{h}^{\widehat{\pi}^{\star}}\right) \iota_{2}}{N}}+2 \sqrt{\frac{\iota_{2}}{N \cdot H S}}+\frac{2 H^{2} S \log (H S A / \delta)}{N}\right) \\
& \leq C^{\prime}\left[\sum_{t=h}^{H} \Gamma_{h+1: t}^{\widehat{\pi}^{\star}} \cdot \sqrt{\frac{2 \operatorname{Var}_{s, a}\left(V_{h}^{\widehat{\pi}^{\star}}\right) \iota_{2}}{N}}+2 \sqrt{\frac{H \log (H S A / \delta)}{N}}+\frac{2 H^{3} S \log (H S A / \delta)}{N}\right. \\
& \leq C^{\prime}\left[\sqrt{\frac{2 H^{3} S \log (H S A / \delta)}{N}}+2 \sqrt{\frac{H \log (H S A / \delta)}{N}}+\frac{2 H^{3} S \log (H S A / \delta)}{N}\right] \\
& \leq C^{\prime \prime}\left[\sqrt{\frac{H^{3} S \log (H S A / \delta)}{N}}+\frac{H^{3} S \log (H S A / \delta)}{N}\right] \\
& \leq O\left[\sqrt{\frac{H^{2} S \log (H S A / \delta)}{n d_{m}}}+\frac{H^{2} S \log (H S A / \delta)}{n d_{m}}\right]
\end{aligned}
$$

where the third inequality uses Lemma J.5 and the last one uses $N \geq \frac{1}{2} n d_{m}$ with high probability. Similar result holds for $\widehat{Q}_{h}^{\pi^{\star}}-Q_{h}^{\pi^{\star}}$. Combing those results we have reward-free bound (for any reward simultaneously)

$$
O\left[\sqrt{\frac{H^{2} S \log (H S A / \delta)}{n d_{m}}}+\frac{H^{2} S \log (H S A / \delta)}{n d_{m}}\right]
$$

which finishes the proof of Theorem 5.4
Remark F.4. Note above result is tight in both the dominate term AND the higher order term. Therefore this result cannot be further improved even in the higher order term.

[^4]
## G Discussion of Section 5

In this section we explain why Theorem 5.3 and Theorem 5.4 are optimal in the offline RL.
We begin with the offline task-agnostic setting. For the exquisite readers who check the proof of Theorem 5 of Zhang et al. [2020b], the proving procedure of their lower bound follows the standard reduction to best-arm identification in multi-armed bandit problems. More specifically, to incorporate the dependence of $\log (K)$, they rely on the Theorem 10 of Zhang et al. [2020b] (which is originated from Mannor and Tsitsiklis [2004]) to show in order to be ( $\epsilon, \delta$ )-correct for a problem with $A$ arms and with $K$ tasks, it need at least $\Omega\left(\frac{A}{\epsilon^{2}} \log \left(\frac{K}{\delta}\right)\right)$ samples. Such a result updates the Lemma G.1. in Yin et al. [2021b] by the extra factor $\log (K)$ for the bandit problem with $K$ tasks. With no modification, the rest of the proof in Section $E$ of Yin et al. [2021b] follows though and one can end up with the lower bound $\Omega\left(H^{2} \log (K) / d_{m} \epsilon^{2}\right)$ over the problem class $\mathcal{M}_{d_{m}}:=\left\{(\mu, M) \mid \min _{t, s_{t}, a_{t}} d_{t}^{\mu}\left(s_{t}, a_{t}\right) \geq d_{m}\right\}$. The case for the offline reward-free setting is also similar. Indeed, the $\Omega\left(S A / \epsilon^{2}\right)$ trajectories in Lemma 4.2 in Jin et al. [2020a] could be replaced by $\Omega\left(1 / d_{m} \epsilon^{2}\right)$ by choosing some hard near-uniform behavior policy instances (see Section E. 2 in Yin et al. [2021b]) and the rest follows since by forcing $S$ such instances (Section 4.2 of Jin et al. [2020a]) to obtain $\Omega\left(S / d_{m} \epsilon^{2}\right)$ and create a chain of $\Omega(H)$ rewards for $\Omega\left(H^{2} S / d_{m} \epsilon^{2}\right)$.

## H Proof of the linear MDP with anchor representations (Section 6)

Recall that we assume a generative oracle here. Sometimes we abuse the notation $\mathcal{K}$ for either anchor point set or the anchor point indices set. The meaning should be clear in each context.

## H. 1 Model-based Plug-in Estimator for Anchor Representations

Step 1: For each $\left(s_{k}, a_{k}\right)$ where index $k \in \mathcal{K}$, collect $N$ samples from $P\left(\cdot \mid s_{k}, a_{k}\right)$; compute

$$
\widehat{P}_{\mathcal{K}}\left(s^{\prime} \mid s_{k}, a_{k}\right)=\frac{\operatorname{count}\left(s, a, s^{\prime}\right)}{N}
$$

Step 2: Compute the linear combination coefficients $\lambda_{k}^{s, a}$ satisfies $\phi(s, a)=$ $\sum_{k \in \mathcal{K}} \lambda_{k}^{s, a} \phi\left(s_{k}, a_{k}\right)$;
Step 3: Estimate transition distribution

$$
\widehat{P}\left(s^{\prime} \mid s, a\right)=\sum_{k \in \mathcal{K}} \lambda_{k}^{s, a} \cdot \widehat{P}_{\mathcal{K}}\left(s^{\prime} \mid s_{k}, a_{k}\right)
$$

We need to check such $\widehat{P}\left(s^{\prime} \mid s, a\right)$ is a valid distribution. This is due to:

$$
\begin{aligned}
\sum_{k \in \mathcal{K}} \lambda_{k}^{s, a} & =\sum_{k \in \mathcal{K}} \sum_{s^{\prime}} \lambda_{k}^{s, a} P\left(s^{\prime} \mid s_{k}, a_{k}\right)=\sum_{s^{\prime}} \sum_{k \in \mathcal{K}} \lambda_{k}^{s, a} P\left(s^{\prime} \mid s_{k}, a_{k}\right) \\
& =\sum_{s^{\prime}} \sum_{k \in \mathcal{K}} \lambda_{k}^{s, a}\left\langle\phi\left(s_{k}, a_{k}\right), \psi\left(s^{\prime}\right)\right\rangle=\sum_{s^{\prime}}\left\langle\phi(s, a), \psi\left(s^{\prime}\right)\right\rangle=\sum_{s^{\prime}} P\left(s^{\prime} \mid s, a\right)=1
\end{aligned}
$$

and

$$
\begin{aligned}
\sum_{s^{\prime}} \widehat{P}\left(s^{\prime} \mid s, a\right) & =\sum_{s^{\prime}} \sum_{k \in \mathcal{K}} \lambda_{k}^{s, a} \widehat{P}_{\mathcal{K}}\left(s^{\prime} \mid s_{k}, a_{k}\right)=\sum_{k \in \mathcal{K}} \sum_{s^{\prime}} \lambda_{k}^{s, a} \widehat{P}_{\mathcal{K}}\left(s^{\prime} \mid s_{k}, a_{k}\right) \\
& =\sum_{k \in \mathcal{K}} \lambda_{k}^{s, a} \frac{N}{N}=1
\end{aligned}
$$

Step 4: construct empirical model $\widehat{M}=(\mathcal{S}, \mathcal{A}, \widehat{P}, r, H)$ and output $\widehat{\pi}^{\star}=\operatorname{argmax}_{\pi} \widehat{V}_{1}^{\pi}$.

Similarly, Bellman (optimality) equations hold ${ }^{15}$

$$
\begin{array}{ll}
V_{t}^{\star}(s)=\max _{a}\left\{r(s, a)+\int_{s^{\prime}} V_{t+1}^{\star}\left(s^{\prime}\right) d P\left(s^{\prime} \mid s, a\right)\right\}, & \forall s \in \mathcal{S} \\
\widehat{V}_{t}^{\star}(s)=\max _{a}\left\{r(s, a)+\int_{s^{\prime}} \widehat{V}_{t+1}^{\star}\left(s^{\prime}\right) d \widehat{P}\left(s^{\prime} \mid s, a\right)\right\}, & \forall s \in \mathcal{S}
\end{array}
$$

## H. 2 General absorbing MDP

The definition of the general absorbing MDP remains the same: i.e. for a fixed state $s$ and a sequence $\left\{u_{t}\right\}_{t=1}^{H}$, MDP $M_{s,\left\{u_{t}\right\}_{t=1}^{H}}$ is identical to $M$ for all states except $s$, and state $s$ is absorbing in the sense $P_{M_{s,\left\{u_{t}\right\}}^{t=1}}(s \mid s, a)=1$ for all $a$, and the instantaneous reward at time $t$ is $r_{t}(s, a)=u_{t}$ for all $a \in \mathcal{A}$. Also, we use the shorthand notation $V_{\left\{s, u_{t}\right\}}^{\pi}$ for $V_{s, M_{s,\left\{u_{t}\right\}_{t=1}^{H}}^{\pi}}$ and similarly for $Q_{\left\{s, u_{t}\right\}}$ and transition $P_{\left\{s, u_{t}\right\}}$. Then the following properties mirroring the Lemma B. 1 and Lemma B. 2 with nearly identical proof but for the integral version (which we skip):

## Lemma H.1.

$$
V_{h,\left\{s, u_{t}\right\}}^{\star}(s)=\sum_{t=h}^{H} u_{t} .
$$

Lemma H.2. Fix state s. For two different sequences $\left\{u_{t}\right\}_{t=1}^{H}$ and $\left\{u_{t}^{\prime}\right\}_{t=1}^{H}$, we have

$$
\max _{h}\left\|Q_{h,\left\{s, u_{t}\right\}}^{\star}-Q_{h,\left\{s, u_{t}^{\prime}\right\}}^{\star}\right\|_{\infty} \leq H \cdot \max _{t \in[H]}\left|u_{t}-u_{t}^{\prime}\right|
$$

## H. 3 Singleton-absorbing MDP

The well-definedness of singleton-absorbing MDP for linear MDP with anchor points depends on the following two lemmas whose proofs are still nearly identical to Lemma B. 3 and Lemma B. 5 which we skip.
Lemma H.3. $V_{t}^{\star}(s)-V_{t+1}^{\star}(s) \geq 0$, for all state $s \in \mathcal{S}$ and all $t \in[H]$.
Lemma H.4. Fix a state s. If we choose $u_{t}^{\star}:=V_{t}^{\star}(s)-V_{t+1}^{\star}(s) \forall t \in[H]$, then we have the following vector form equation

$$
V_{h,\left\{s, u_{t}^{\star}\right\}}^{\star}=V_{h, M}^{\star} \quad \forall h \in[H] .
$$

Similarly, if we choose $\widehat{u}_{t}^{\star}:=\widehat{V}_{t}^{\star}(s)-\widehat{V}_{t+1}^{\star}(s)$, then $\widehat{V}_{h,\left\{s, \hat{u}_{t}^{\star}\right\}}^{\star}=\widehat{V}_{h, M}^{\star}, \forall h \in[H]$.
The singleton MDP we used is exactly $M_{s,\left\{u_{t}^{\star}\right\}_{t=1}^{H}}$ (or $\widehat{M}_{s,\left\{u_{t}^{\star}\right\}_{t=1}^{H}}$ ).

## H. 4 Proof for the optimal sample complexity

For $\widehat{\pi}^{\star}$, by (empirical) Bellman equation we have element-wisely:

$$
\begin{aligned}
\widehat{Q}_{h}^{\widehat{\pi}^{\star}}-Q_{h}^{\widehat{\pi}^{\star}} & =r_{h}+\widehat{P}^{\widehat{\pi}_{h+1}^{\star}} \widehat{Q}_{h+1}^{\widehat{\pi}^{\star}}-r_{h}-P^{\widehat{\pi}_{h+1}^{\star}} Q_{h+1}^{\widehat{\pi}^{\star}} \\
& =\left(\widehat{P}^{\widehat{\pi}_{h+1}^{\star}}-P^{\widehat{\pi}_{h+1}^{\star}}\right) \widehat{Q}_{h+1}^{\pi^{\star}}+P^{\widehat{\pi}_{h+1}^{\star}}\left(\widehat{Q}_{h+1}^{\widehat{\pi}^{\star}}-Q_{h+1}^{\widehat{\pi}^{\star}}\right) \\
& =(\widehat{P}-P) \widehat{V}_{h+1}^{\widehat{\pi}^{\star}}+P^{\widehat{\pi}_{h+1}^{\star}}\left(\widehat{Q}_{h+1}^{\widehat{\pi}^{\star}}-Q_{h+1}^{\widehat{\pi}^{\star}}\right) \\
& =\ldots=\sum_{t=h}^{H} \Gamma_{h+1: t}^{\widehat{\pi}^{\star}}(\widehat{P}-P) \widehat{V}_{t+1}^{\widehat{\pi}^{\star}} \leq \underbrace{\sum_{t=h}^{H} \Gamma_{h+1: t}^{\widehat{\pi}^{\star}}\left|(\widehat{P}-P) \widehat{V}_{t+1}^{\widehat{\pi}^{\star}}\right|}_{(\star)}
\end{aligned}
$$

where $\Gamma_{h+1: t}^{\pi^{\star}}=\prod_{i=h+1}^{t} P^{\pi_{i}^{\star}}$ is multi-step state-action transition and $\Gamma_{h+1: h}:=I$.

[^5]
## H. 5 Analyzing ( $\star$ )

Concentration on $(\widehat{P}-P) \widehat{V}_{h}^{\star}$. Since $\widehat{P}$ aggregates all data from different step so that $\widehat{P}$ and $\widehat{V}_{h}^{\star}$ are on longer independent. We use the singleton-absorbing MDP $M_{s,\left\{u_{t}^{\star}\right\}_{t=1}^{H}}$ to handle the case (recall $u_{t}^{\star}:=V_{t}^{\star}(s)-V_{t+1}^{\star}(s) \forall t \in[H]$ ). Here, we fix the state action $(s, a) \in \mathcal{K}$. Then we have:

$$
\begin{align*}
& \left(\widehat{P}_{s, a}-P_{s, a}\right) \widehat{V}_{h}^{\star}=\left(\widehat{P}_{s, a}-P_{s, a}\right)\left(\widehat{V}_{h}^{\star}-\widehat{V}_{h,\left\{s, u_{t}^{\star}\right\}}^{\star}+\widehat{V}_{h,\left\{s, u_{t}^{\star}\right\}}^{\star}\right) \\
& =\left(\widehat{P}_{s, a}-P_{s, a}\right)\left(\widehat{V}_{h}^{\star}-\widehat{V}_{h,\left\{s, u_{t}^{\star}\right\}}^{\star}\right)+\left(\widehat{P}_{s, a}-P_{s, a}\right) \widehat{V}_{h,\left\{s, u_{t}^{\star}\right\}} \\
& \leq\left\|\widehat{P}_{s, a}-P_{s, a}\right\|_{1}\left\|\widehat{V}_{h}^{\star}-\widehat{V}_{h,\left\{s, u_{t}^{\star}\right\}}^{\star}\right\|_{\infty}+\sqrt{\frac{2 \log (4 / \delta)}{N}} \sqrt{\operatorname{Var}_{s, a}\left(\widehat{V}_{h,\left\{s, u_{t}^{\star}\right\}}^{\star}\right)}+\frac{2 H \log (1 / \delta)}{3 N} \\
& \leq\left\|\widehat{P}_{s, a}-P_{s, a}\right\|_{1}\left\|\widehat{V}_{h}^{\star}-\widehat{V}_{h,\left\{s, u_{t}^{\star}\right\}}^{\star}\right\|_{\infty}+\sqrt{\frac{2 \log (4 / \delta)}{N}}\left(\sqrt{\operatorname{Var}_{s, a}\left(\widehat{V}_{h}^{\star}\right)}+\sqrt{\operatorname{Var}_{s, a}\left(\widehat{V}_{h,\left\{s, u_{t}^{\star}\right\}}-\widehat{V}_{h}^{\star}\right)}\right)+\frac{2 H \log (1 / \delta)}{3 N} \\
& \leq\left\|\widehat{P}_{s, a}-P_{s, a}\right\|_{1}\left\|\widehat{V}_{h}^{\star}-\widehat{V}_{h,\left\{s, u_{t}^{\star}\right\}}^{\star}\right\|_{\infty}+\sqrt{\frac{2 \log (4 / \delta)}{N}}\left(\sqrt{\operatorname{Var}_{s, a}\left(\widehat{V}_{h}^{\star}\right)}+\sqrt{\left\|\widehat{V}_{h,\left\{s, u_{t}^{\star}\right\}}-\widehat{V}_{h}^{\star}\right\|_{\infty}^{2}}\right)+\frac{2 H \log (1 / \delta)}{3 N} \\
& =\left(\left\|\widehat{P}_{s, a}-P_{s, a}\right\|_{1}+\sqrt{\frac{2 \log (4 / \delta)}{N}}\right)\left\|\widehat{V}_{h}^{\star}-\widehat{V}_{h,\left\{s, u_{t}^{\star}\right\}}\right\|_{\infty}+\sqrt{\frac{2 \log (4 / \delta)}{N}} \sqrt{\operatorname{Var}_{s, a}\left(\widehat{V}_{h}^{\star}\right)}+\frac{2 H \log (1 / \delta)}{3 N} \tag{19}
\end{align*}
$$

where the first inequality uses Bernstein inequality (Lemma J.3) (note here $P_{s, a} V=$ $\int_{s^{\prime}} V\left(s^{\prime}\right) d P\left(s^{\prime} \mid s, a\right)$ since $\mathcal{S}$ could be continuous space, but this does not affect the availability of Bernstein inequality!), the second inequality uses $\sqrt{\operatorname{Var}(\cdot)}$ is norm (norm triangle inequality). Now we treat $\left\|\widehat{P}_{s, a}-P_{s, a}\right\|_{1}$ and $\left\|\widehat{V}_{h}^{\star}-\widehat{V}_{h,\left\{s, u_{t}^{\star}\right\}}^{\star}\right\|_{\infty}$ separately.
For $\left\|\widehat{P}_{s, a}-P_{s, a}\right\|_{1}$. Recall here $(s, a) \in \mathcal{K}$. By Lemma J.9 we obtain w.p. $1-\delta$

$$
\begin{equation*}
\left\|\widehat{P}_{s, a}-P_{s, a}\right\|_{1} \leq C \sqrt{\frac{|\mathcal{S}| \log (1 / \delta)}{N}} . \tag{20}
\end{equation*}
$$

where $C$ absorbs the higher order term and constants.
For $\left\|\widehat{V}_{h}^{\star}-\widehat{V}_{h,\left\{s, u_{t}^{\star}\right\}}^{\star}\right\|_{\infty}$. Note if we set $\widehat{u}_{t}^{\star}=\widehat{V}_{t}^{\star}(s)-\widehat{V}_{t+1}^{\star}(s)$, then by Lemma H. 4

$$
\widehat{V}_{h}^{\star}=\widehat{V}_{h,\left\{s, \hat{u}_{t}^{\star}\right\}}^{\star}
$$

Next since $\widehat{V}_{h,\left\{s, \hat{u}_{t}^{\star}\right\}}(\tilde{s})=\max _{a} \widehat{Q}_{h,\left\{s, \hat{u}_{t}^{\star}\right\}}(\tilde{s}, a) \forall \tilde{s} \in \mathcal{S}$, by generic inequality $|\max f-\max g| \leq$ $\max |f-g|$, we have $\left|\widehat{V}_{h,\left\{s, \hat{u}_{t}^{\star}\right\}}^{\star}(\tilde{s})-\widehat{V}_{h,\left\{s, u_{t}^{\star}\right\}}^{\star}(\tilde{s})\right| \leq \max _{a}\left|\widehat{Q}_{h,\left\{s, \hat{u}_{t}^{\star}\right\}}^{\star}(\tilde{s}, a)-\widehat{Q}_{h,\left\{s, u_{t}^{\star}\right\}}^{\star}(\tilde{s}, a)\right|$, taking $\max _{\tilde{s}}$ on both sides, we obtain exactly

$$
\left\|\widehat{V}_{h,\left\{s, \hat{u}_{t}^{\star}\right\}}^{\star}-\widehat{V}_{h,\left\{s, u_{t}^{\star}\right\}}^{\star}\right\|_{\infty} \leq\left\|\widehat{Q}_{h,\left\{s, \hat{u}_{t}^{\star}\right\}}^{\star}-\widehat{Q}_{h,\left\{s, u_{t}^{\star}\right\}}^{\star}\right\|_{\infty}
$$

then by Lemma H. 2 .

$$
\begin{equation*}
\left\|\widehat{V}_{h}^{\star}-\widehat{V}_{h,\left\{s, u_{t}^{\star}\right\}}^{\star}\right\|_{\infty} \leq\left\|\widehat{Q}_{h,\left\{s, \hat{u}_{t}^{\star}\right\}}^{\star}-\widehat{Q}_{h,\left\{s, u_{t}^{\star}\right\}}^{\star}\right\|_{\infty} \leq H \max _{t}\left|\hat{u}_{t}^{\star}-u_{t}^{\star}\right| \tag{21}
\end{equation*}
$$

Recall

$$
\hat{u}_{t}^{\star}-u_{t}^{\star}=\widehat{V}_{t}^{\star}(s)-\widehat{V}_{t+1}^{\star}(s)-\left(V_{t}^{\star}(s)-V_{t+1}^{\star}(s)\right) .
$$

Now we denote

$$
\Delta_{s}:=\max _{t}\left|\hat{u}_{t}^{\star}-u_{t}^{\star}\right|=\max _{t}\left|\widehat{V}_{t}^{\star}(s)-\widehat{V}_{t+1}^{\star}(s)-\left(V_{t}^{\star}(s)-V_{t+1}^{\star}(s)\right)\right|
$$

then $\Delta_{s}$ itself is a scalar and a random variable.
To sum up, by (19), (5) and (21) and a union bound over all $(s, a) \in \mathcal{K}$ we have

Lemma H.5. Fix $N>0$. With probability $1-\delta$, element-wisely, for all $h \in[H]$ and all $\left(s_{k}, a_{k}\right) \in \mathcal{K}$,

$$
\begin{aligned}
\left|\left(\widehat{P}_{s_{k}, a_{k}}-P_{s_{k}, a_{k}}\right) \widehat{V}_{h}^{\star}\right| & \leq C \sqrt{\frac{|\mathcal{S}| \log (H K / \delta)}{N}} \cdot H \max _{s_{k}} \Delta_{s_{k}} \\
& +\sqrt{\frac{2 \log (4 H K / \delta)}{N}} \sqrt{\operatorname{Var}_{P_{s_{k}, a_{k}}}\left(\widehat{V}_{h}^{\star}\right)}+\frac{2 H \log (H K / \delta)}{3 N}
\end{aligned}
$$

Now we extend Lemma H. 5 to any arbitrary $(s, a)$ by proving the following lemma:
Lemma H. 6 (recover lemma). For any function $V$ and any state action $(s, a)$, we have

$$
\sum_{k \in \mathcal{K}} \lambda_{k}^{s, a} \sqrt{\operatorname{Var}_{P_{s_{k}, a_{k}}}(V)} \leq \sqrt{\operatorname{Var}_{P_{s, a}}(V)}
$$

Proof of Lemma H.6. Since $\lambda_{k}^{s, a}$ are probability distributions, by Jensen's inequality twice

$$
\begin{aligned}
& \sum_{k \in \mathcal{K}} \lambda_{k}^{s, a} \sqrt{\operatorname{Var}_{P_{s_{k}, a_{k}}}(V)} \leq \sqrt{\sum_{k \in \mathcal{K}} \lambda_{k}^{s, a} \operatorname{Var}_{P_{s_{k}, a_{k}}}(V)} \\
= & \sqrt{\sum_{k \in \mathcal{K}} \lambda_{k}^{s, a} \operatorname{Var}_{P_{s_{k}, a_{k}}}(V)}=\sqrt{\sum_{k \in \mathcal{K}} \lambda_{k}^{s, a}\left(P_{s_{k}, a_{k}} V^{2}-\left(P_{s_{k}, a_{k}} V\right)^{2}\right)} \\
\leq & \sqrt{\sum_{k \in \mathcal{K}} \lambda_{k}^{s, a} \cdot P_{s_{k}, a_{k}} V^{2}-\left(\sum_{k \in \mathcal{K}} \lambda_{k}^{s, a} P_{s_{k}, a_{k}} V\right)^{2}} \\
= & \sqrt{P_{s, a} V^{2}-\left(P_{s, a} V\right)^{2}}=\sqrt{\operatorname{Var}_{P_{s, a}}(V)},
\end{aligned}
$$

where we use $P_{s, a}=\sum_{k \in \mathcal{K}} \lambda_{k}^{s, a} P_{s_{k}, a_{k}}$.
Therefore for all $(s, a)$, using Lemma H.5 and Lemma H.6 we obtain w.p. $1-\delta$,

$$
\begin{aligned}
& \left|\left(\widehat{P}_{s, a}-P_{s, a}\right) \widehat{V}_{h}^{\star}\right| \leq \sum_{k \in \mathcal{K}} \lambda_{k}^{s, a}\left|\left(\widehat{P}_{s_{k}, a_{k}}-P_{s_{k}, a_{k}}\right) \widehat{V}_{h}^{\star}\right| \\
& \leq C \sum_{k \in \mathcal{K}} \lambda_{k}^{s, a} \sqrt{\frac{S \log (H K / \delta)}{N}} \cdot H \max _{s_{k}} \Delta_{s_{k}}+\sum_{k \in \mathcal{K}} \lambda_{k}^{s, a} \sqrt{\frac{2 \log (4 H K / \delta)}{N}} \sqrt{\operatorname{Var}_{P_{s_{k}, a_{k}}\left(\widehat{V}_{h}^{\star}\right)}^{N}} \\
& +\sum_{k \in \mathcal{K}} \lambda_{k}^{s, a} \frac{2 H \log (H K / \delta)}{3 N} \\
& =C \sqrt{\frac{S \log (H K / \delta)}{N}} \cdot H \max _{s_{k}} \Delta_{s_{k}}+\sum_{k \in \mathcal{K}} \lambda_{k}^{s, a} \sqrt{\frac{2 \log (4 H K / \delta)}{N}} \sqrt{\operatorname{Var}_{P_{s_{k}, a_{k}}\left(\widehat{V}_{h}^{\star}\right)}^{N}} \\
& +\frac{2 H \log (H K / \delta)}{3 N} \\
& \leq C \sqrt{\frac{S \log (H K / \delta)}{N}} \cdot H \max _{s_{k}} \Delta_{s_{k}}+\sqrt{\frac{2 \log (4 H K / \delta)}{N}} \sqrt{\operatorname{Var}_{P_{s, a}}\left(\widehat{V}_{h}^{\star}\right)}+\frac{2 H \log (H K / \delta)}{3 N}
\end{aligned}
$$

Now plug above back into $(\star)$, we receive:

$$
\begin{aligned}
& \left|\widehat{Q}_{h}^{\widehat{\pi}^{\star}}-Q_{h}^{\widehat{\pi}^{\star}}\right| \\
\leq & \sum_{t=h}^{H} \Gamma_{h+1: t}^{\widehat{\pi}^{\star}}\left(C \sqrt{\frac{S \log (H K / \delta)}{N}} \cdot H \max _{s_{k}} \Delta_{s_{k}} \cdot \mathbf{1}+\sqrt{\frac{2 \log (4 H K / \delta)}{N}} \sqrt{\operatorname{Var}_{P}\left(\widehat{V}_{t+1}^{\star}\right)}+\frac{2 H \log (H K / \delta)}{3 N} \cdot \mathbf{1}\right) \\
\leq & \sum_{t=h}^{H} \Gamma_{h+1: t}^{\widehat{\pi}} \sqrt{\frac{2 \log (4 H K / \delta)}{N}} \sqrt{\operatorname{Var}_{P}\left(\widehat{V}_{t+1}^{\star}\right)}+C H^{2} \sqrt{\frac{S \log (H K / \delta)}{N}} \cdot \max _{s} \Delta_{s} \cdot \mathbf{1}+\frac{2 H^{2} \log (H K / \delta)}{3 N} \cdot \mathbf{1}
\end{aligned}
$$

Similar to before, we get

$$
\begin{equation*}
\sqrt{\operatorname{Var}_{P}\left(\widehat{V}_{h}^{\star}\right)}:=\sqrt{\operatorname{Var}_{P}\left(\widehat{V}_{h}^{\hat{\pi}^{\star}}\right)} \leq \sqrt{\operatorname{Var}_{P}\left(V_{h}^{\widehat{\pi}^{\star}}\right)}+\left\|\widehat{Q}_{h}^{\hat{\pi}^{\star}}-Q_{h}^{\widehat{\pi}^{\star}}\right\|_{\infty} \tag{22}
\end{equation*}
$$

Plug (22) back to above we obtain $\forall h \in[H]$,

$$
\begin{align*}
& \left|\widehat{Q}_{h}^{\widehat{\pi}^{\star}}-Q_{h}^{\widehat{\pi}^{\star}}\right| \leq \sqrt{\frac{2 \log (4 H K / \delta)}{N}} \sum_{t=h}^{H} \Gamma_{h+1: t}^{\widehat{\pi}^{\star}}\left(\sqrt{\operatorname{Var}_{P}\left(V_{t+1}^{\widehat{\pi}^{\star}}\right)}+\left\|\widehat{Q}_{t+1}^{\widehat{\pi}^{\star}}-Q_{t+1}^{\widehat{\pi}^{\star}}\right\|_{\infty}\right) \\
& +C H^{2} \sqrt{\frac{S \log (H K / \delta)}{N}} \cdot \max _{s_{k}} \Delta_{s_{k}} \cdot \mathbf{1}+\frac{2 H^{2} \log (H K / \delta)}{3 N} \cdot \mathbf{1} \\
& \leq \sqrt{\frac{2 \log (4 H K / \delta)}{N} \sum_{t=h}^{H} \Gamma_{h+1: t}^{\widehat{\pi}^{\star}} \sqrt{\operatorname{Var}_{P}\left(V_{t+1}^{\hat{\pi}^{\star}}\right)}+\sqrt{\frac{2 \log (4 H K / \delta)}{N}} \sum_{t=h}^{H}\left\|\widehat{Q}_{t+1}^{\widehat{\pi}^{\star}}-Q_{t+1}^{\widehat{\pi}^{\star}}\right\|_{\infty}} \\
& +C H^{2} \sqrt{\frac{S \log (H K / \delta)}{N}} \cdot \max _{s_{k}} \Delta_{s_{k}} \cdot \mathbf{1}+\frac{2 H^{2} \log (H K / \delta)}{3 N} \cdot \mathbf{1} \tag{23}
\end{align*}
$$

Apply Lemma J.5 and the (anchor version using recover lemma H.6 coarse uniform bound (Lemma $\sqrt{J .10}$ ) we obtain the following lemma:
Lemma H.7. With probability $1-\delta$, for all $h \in[H]$,

$$
\left\|\widehat{Q}_{h}^{\widehat{\pi}^{\star}}-Q_{h}^{\widehat{\pi}^{\star}}\right\|_{\infty} \leq \sqrt{\frac{C_{0} H^{3} \log (4 H K / \delta)}{N}}+\sqrt{\frac{2 \log (4 H K / \delta)}{N}} \sum_{t=h}^{H}\left\|\widehat{Q}_{t+1}^{\widehat{\pi}^{\star}}-Q_{t+1}^{\widehat{\pi}^{\star}}\right\|_{\infty}+C^{\prime} H^{4} \frac{S \log (H K / \delta)}{N}
$$

Proof. Since

$$
\begin{align*}
\Delta_{s_{k}} & :=\max _{t}\left|\hat{u}_{t}^{\star}-u_{t}^{\star}\right|=\max _{t}\left|\widehat{V}_{t}^{\star}\left(s_{k}\right)-\widehat{V}_{t+1}^{\star}\left(s_{k}\right)-\left(V_{t}^{\star}\left(s_{k}\right)-V_{t+1}^{\star}\left(s_{k}\right)\right)\right| \\
& \leq 2 \cdot \max _{t}\left|\widehat{V}_{t}^{\star}\left(s_{k}\right)-V_{t}^{\star}\left(s_{k}\right)\right| \\
& =2 \cdot \max _{t}\left|\max _{\pi} \widehat{V}_{t}^{\pi}\left(s_{k}\right)-\max _{\pi} V_{t}^{\pi}\left(s_{k}\right)\right|  \tag{24}\\
& \leq 2 \cdot \max _{\pi \in \Pi_{g}, t \in[H]}| | \widehat{V}_{t}^{\pi}-\left.V_{t}^{\pi}\right|_{\infty} \leq C \cdot H^{2} \sqrt{\frac{|\mathcal{S}| \log (H K / \delta)}{N}}
\end{align*}
$$

where the last inequality uses (the anchor version) of Lemma J.10 ${ }^{16}$ Then apply union bound w.p. $1-\delta / 2$, we obtain $\max _{s_{k}} \Delta_{s_{k}} \leq C \cdot H^{2} \sqrt{\frac{|\mathcal{S}| \log \left(H K^{2} / \delta\right)}{N}}$. Note 23) holds with probability $1-\delta / 2$, therefore plug above into 23 and uses Lemma J.5 and take $\|\cdot\|_{\infty}$ we obtain w.p. $1-\delta$, the result holds.

Lemma H.8. Given $N>0$. Define $C^{\prime \prime}:=2 \cdot \max \left(\sqrt{C_{0}}, C^{\prime}\right)$ where $C^{\prime}$ is the universal constant in Lemma H.7. When $N \geq 8 H^{2}|\mathcal{S}| \log (4 H K / \delta)$, then with probability $1-\delta, \forall h \in[H]$,

$$
\begin{align*}
& \left\|\widehat{Q}_{h}^{\hat{\pi}^{\star}}-Q_{h}^{\widehat{\pi}^{\star}}\right\|_{\infty} \leq C^{\prime \prime} \sqrt{\frac{H^{3} \log (4 H K / \delta)}{N}}+C^{\prime \prime} \frac{H^{4} S \log (H K / \delta)}{N}  \tag{25}\\
& \left\|\widehat{Q}_{h}^{\pi^{\star}}-Q_{h}^{\pi^{\star}}\right\|_{\infty} \leq C^{\prime \prime} \sqrt{\frac{H^{3} \log (4 H K / \delta)}{N}}+C^{\prime \prime} \frac{H^{4} S \log (H K / \delta)}{N}
\end{align*}
$$

Proof. The proof is the same as Lemma B. 9
Remark H.9. Note the higher order term has dependence $H^{4} S$. Use the same self-bounding trick, we can reduce it to $H^{3.5} S^{0.5}$.

[^6]Lemma H.10. Given $N>0$. There exists universal constants $C_{1}, C_{2}$ such that when $N \geq$ $C_{1} H^{2}|\mathcal{S}| \log (H K / \delta)$, then with probability $1-\delta, \forall h \in[H]$,

$$
\begin{equation*}
\left\|\widehat{Q}_{h}^{\pi^{\star}}-Q_{h}^{\widehat{\pi}^{\star}}\right\|_{\infty} \leq C_{2} \sqrt{\frac{H^{3} \log (H K / \delta)}{N}}+C_{2} \frac{H^{3} \sqrt{H S} \log (H K / \delta)}{N} \tag{26}
\end{equation*}
$$

and

$$
\left\|\widehat{Q}_{h}^{\pi^{\star}}-Q_{h}^{\pi^{\star}}\right\|_{\infty} \leq C_{2} \sqrt{\frac{H^{3} \log (H K / \delta)}{N}}+C_{2} \frac{H^{3} \sqrt{H S} \log (H K / \delta)}{N}
$$

Proof. The proof is similar to Lemma B. 11 .

## H. 6 Proof of Theorem 6.2

Proof. By the direct computing of the suboptimality,

$$
Q_{1}^{\star}-Q_{1}^{\widehat{\pi}^{\star}}=Q_{1}^{\star}-\widehat{Q}_{1}^{\pi^{\star}}+\widehat{Q}_{1}^{\pi^{\star}}-\widehat{Q}_{1}^{\widehat{\pi}^{\star}}+\widehat{Q}_{1}^{\widehat{\pi}^{\star}}-Q_{1}^{\widehat{\pi}^{\star}} \leq\left|Q_{1}^{\star}-\widehat{Q}_{1}^{\pi^{\star}}\right|+\left|\widehat{Q}_{1}^{\widehat{\pi}^{\star}}-Q_{1}^{\widehat{\pi}^{\star}}\right|
$$

then by LemmaH. 10 we can finish the proof.

## H. 7 Take-away in the linear MDP with anchor setting.

Under the setting $S$ could be exponential large, $\mathcal{A}$ could be infinite (or even continuous space), with anchor representations $(K \ll|\mathcal{S}|)$, our Theorem 6.2 has order $\widetilde{O}\left(\sqrt{H^{3} / N}\right)$ when $N$ is sufficiently large. This translate to $N=\widetilde{O}\left(H^{3} / \epsilon^{2}\right)$ and the total sample used is $K N=\widetilde{O}\left(K H^{3} / \epsilon^{2}\right)$. This improves the total complexity $\widetilde{O}\left(K H^{4} / \epsilon^{2}\right)$ in Cui and Yang 2020 and is optimal.

## I The computational efficiency for the model-based offline plug-in estimators

For completeness, we discuss the computational and storage aspect of our model-based method. Its computational cost is $\widetilde{O}\left(H^{4} / d_{m} \epsilon^{2}\right)$ for computing $\widehat{P}$, the same as its sample complexity in steps ( $H$ steps is an episode), and running value iteration causes $O\left(H S^{2} A\right.$ ) time (here we assume the bit complexity $L(P, r, H)=1$, see Agarwal et al. 2019] Section 1.3). The total computational complexity is $\widetilde{O}\left(H^{4} / d_{m} \epsilon^{2}\right)+O\left(H S^{2} A\right)$. The memory cost is $O\left(H S^{2} A\right)$.

## J Assisting lemmas

Lemma J. 1 (Multiplicative Chernoff bound Chernoff et al.[1952]). Let X be a Binomial random variable with parameter $p, n$. For any $1 \geq \theta>0$, we have that

$$
\mathbb{P}[X<(1-\theta) p n]<e^{-\frac{\theta^{2} p n}{2}} . \quad \text { and } \quad \mathbb{P}[X \geq(1+\theta) p n]<e^{-\frac{\theta^{2} p n}{3}}
$$

Lemma J. 2 (Hoeffding's Inequality Sridharan [2002]). Let $x_{1}, \ldots, x_{n}$ be independent bounded random variables such that $\mathbb{E}\left[x_{i}\right]=0$ and $\left|x_{i}\right| \leq \xi_{i}$ with probability 1 . Then for any $\epsilon>0$ we have

$$
\mathbb{P}\left(\frac{1}{n} \sum_{i=1}^{n} x_{i} \geq \epsilon\right) \leq e^{-\frac{2 n^{2} \epsilon^{2}}{\sum_{i=1}^{n} \xi_{i}^{2}}}
$$

Lemma J. 3 (Bernstein's Inequality). Let $x_{1}, \ldots, x_{n}$ be independent bounded random variables such that $\mathbb{E}\left[x_{i}\right]=0$ and $\left|x_{i}\right| \leq \xi$ with probability 1 . Let $\sigma^{2}=\frac{1}{n} \sum_{i=1}^{n} \operatorname{Var}\left[x_{i}\right]$, then with probability $1-\delta$ we have

$$
\frac{1}{n} \sum_{i=1}^{n} x_{i} \leq \sqrt{\frac{2 \sigma^{2} \cdot \log (1 / \delta)}{n}}+\frac{2 \xi}{3 n} \log (1 / \delta)
$$

Lemma J. 4 (Freedman's inequality Tropp et al. [2011]). Let $X$ be the martingale associated with a filter $\mathcal{F}$ (i.e. $\left.\quad X_{i}=\mathbb{E}\left[X \mid \mathcal{F}_{i}\right]\right)$ satisfying $\left|X_{i}-X_{i-1}\right| \leq M$ for $i=1, \ldots, n$. Denote $W:=\sum_{i=1}^{n} \operatorname{Var}\left(X_{i} \mid \mathcal{F}_{i-1}\right)$ then we have

$$
\mathbb{P}\left(|X-\mathbb{E}[X]| \geq \epsilon, W \leq \sigma^{2}\right) \leq 2 e^{-\frac{\epsilon^{2}}{2\left(\sigma^{2}+M \epsilon / 3\right)}}
$$

Or in other words, with probability $1-\delta$,

$$
|X-\mathbb{E}[X]| \leq \sqrt{8 \sigma^{2} \cdot \log (1 / \delta)}+\frac{2 M}{3} \cdot \log (1 / \delta), \quad \text { Or } \quad W \geq \sigma^{2}
$$

Lemma J. 5 (Sum of expectation of conditional variance of value; Lemma F. 3 of Yin et al. [2021a]).

$$
\begin{aligned}
& \operatorname{Var}_{\pi}\left[\sum_{t=h}^{H} r_{t}^{(1)} \mid s_{h}^{(1)}=s_{h}, a_{h}^{(1)}=a_{h}\right] \\
= & \sum_{t=h}^{H}\left(\mathbb{E}_{\pi}\left[\operatorname{Var}\left[r_{t}^{(1)}+V_{t+1}^{\pi}\left(s_{t+1}^{(1)}\right) \mid s_{t}^{(1)}, a_{t}^{(1)}\right] \mid s_{h}^{(1)}=s_{h}, a_{h}^{(1)}=a_{h}\right]\right. \\
+ & \left.\mathbb{E}_{\pi}\left[\operatorname{Var}\left[\mathbb{E}\left[r_{t}^{(1)}+V_{t+1}^{\pi}\left(s_{t+1}^{(1)}\right) \mid s_{t}^{(1)}, a_{t}^{(1)}\right] \mid s_{t}^{(1)}\right] \mid s_{h}^{(1)}=s_{h}, a_{h}^{(1)}=a_{h}\right]\right)
\end{aligned}
$$

By apply above, one can show

$$
\sum_{t=h}^{H} \Gamma_{h+1: t}^{\pi} \sqrt{\operatorname{Var}_{P}\left(V_{t+1}^{\pi}\right)} \leq \sqrt{(H-h)^{3}} \cdot \mathbf{1}
$$

Remark J.6. The infinite horizon discounted setting counterpart result is $\left(I-\gamma P^{\pi}\right)^{-1} \sigma_{V^{\pi}} \leq$ $(1-\gamma)^{-3 / 2}$.

## J. 1 Minimax rate of discrete distributions under $l_{1}$ loss.

This Section provides the minimax rate for $\|\widehat{P}-P\|_{1}$ for any model-based algorithms and is based on Han et al. [2015]. Let $P$ be $S$ dimensional distribution.
Lemma J. 7 (Minimax lower bound for $\|\widehat{P}-P\|_{1}$ ). Let $n$ be the number of data-points sampled from $P$. If $n>\frac{e}{32} S$, then there exists a constant $p>0$, such that

$$
\inf _{\widehat{P}} \sup _{P \in \mathcal{M}_{S}} \mathbb{P}\left[\|\widehat{P}-P\|_{1} \geq \frac{1}{8} \sqrt{\frac{e S}{2 n}}-o\left(e^{-n}\right)-o\left(e^{-S}\right)\right] \geq p
$$

where $\mathcal{M}_{S}$ denotes the set of distributions with support size $S$ and the infimum is taken over $\boldsymbol{A L L}$ estimators.
Remark J.8. Note the $\widehat{P}$ in above carries over all estimators but not just empirical estimator. This provides the minimax result.

Proof. The proof comes from Theorem 2 of Han et al. [2015], where we pick $\zeta=1$. Note they establish the minimax result for $\mathbb{E}_{P}\|\hat{P}-P\|_{1}$. However, by a simple contradiction we can get the above. Indeed, suppose

$$
\inf _{\widehat{P}} \sup _{P \in \mathcal{M}_{S}} \mathbb{P}\left[\|\widehat{P}-P\|_{1}<\frac{1}{8} \sqrt{\frac{e S}{2 n}}-o\left(e^{-n}\right)-o\left(e^{-S}\right)\right]=1
$$

then this implies $\inf _{\widehat{P}} \sup _{P \in \mathcal{M}_{S}} \mathbb{E}_{P}\|\hat{P}-P\|_{1}<\frac{1}{8} \sqrt{\frac{e S}{2 n}}-o\left(e^{-n}\right)-o\left(e^{-S}\right)$ which contradicts Theorem 2 of Han et al. [2015].

Lemma J. 9 (Upper bound for $\|\widehat{P}-P\|_{1}$ ). Let $n$ be the number of data-points sampled from $P$. Then with probability $1-\delta$

$$
\|\widehat{P}-P\|_{1} \leq C\left(\sqrt{\frac{S \log (S / \delta)}{n}}+\frac{S \log (S / \delta)}{n}\right)
$$

for any $P \in \mathcal{M}_{S}$. Here $\widehat{P}$ is the empirical (MLE) estimator.

Proof. First fix a state $s$. Let $X_{i}=\mathbf{1}\left[s_{i}=s\right]$, then $X_{i} \sim \operatorname{Bern}\left(p_{s}\left(1-p_{s}\right)\right)$ and $X_{s}=\sum_{i=1}^{n} X_{i} \sim$ $\operatorname{Binomial}\left(n, p_{i}\right)$. By Bernstein inequality,

$$
\left|\frac{X_{s}}{n}-P_{s}\right| \leq \sqrt{\frac{2 p_{s}\left(1-p_{s}\right) \log (1 / \delta)}{n}}+\frac{3}{n} \log (1 / \delta)
$$

Apply a union bound we obtain w.p. $1-\delta$

$$
\left|\frac{X_{s}}{n}-P_{s}\right| \leq \sqrt{\frac{2 p_{s}\left(1-p_{s}\right) \log (S / \delta)}{n}}+\frac{3}{n} \log (S / \delta) \quad \forall s \in \mathcal{S}
$$

which implies

$$
\begin{aligned}
\|\widehat{P}-P\|_{1} & =\sum_{s \in \mathcal{S}}\left|\frac{X_{s}}{n}-P_{s}\right| \\
& \leq \sum_{s \in \mathcal{S}} \sqrt{\frac{2 p_{s}\left(1-p_{s}\right) \log (S / \delta)}{n}}+\frac{3 S}{n} \log (S / \delta) \\
& =\sqrt{\frac{1}{n}} \sum_{s \in \mathcal{S}} \frac{1}{S} \cdot \sqrt{2 S^{2} p_{s}\left(1-p_{s}\right) \log (S / \delta)}+\frac{3 S}{n} \log (S / \delta) \\
& \leq \sqrt{\frac{1}{n}} \sqrt{2 S^{2} \cdot \frac{\sum_{s \in \mathcal{S}} p_{s}}{S}}\left(1-\frac{\sum_{s \in \mathcal{S}} p_{s}}{S}\right) \log (S / \delta)
\end{aligned}+\frac{3 S}{n} \log (S / \delta) .
$$

where the last inequality uses the concavity of $\sqrt{x(1-x)}$.
Finally, we can absorb the higher order term using the mild condition $n>c \cdot S \log (S / \delta)$.

## J. 2 A crude uniform convergence bound

Here we provide a crude bound for $\sup _{\pi \in \Pi_{g}}\left\|\widehat{V}_{1}^{\pi}-V_{1}^{\pi}\right\|_{\infty}$, which is the finite horizon counterpart of Section 2.2 of Jiang [2018] and is a form of simulation lemma.
Lemma J. 10 (Crude bound by Simulation Lemma). Fix $N>0$ to be number of samples for each coordinates. Recall $\Pi_{g}$ is the global policy class. Then w.p. $1-\delta$,

$$
\sup _{\pi \in \Pi_{g}, h \in[H]}\left\|\widehat{Q}_{h}^{\pi}-Q_{h}^{\pi}\right\|_{\infty} \leq C \cdot H^{2} \sqrt{\frac{S \log (S A / \delta)}{N}}
$$

which further implies

$$
\sup _{\pi \in \Pi_{g}, h \in[H]}\left\|\widehat{V}_{h}^{\pi}-V_{h}^{\pi}\right\|_{\infty} \leq C \cdot H^{2} \sqrt{\frac{S \log (S A / \delta)}{N}},
$$

Proof.

$$
\begin{aligned}
\widehat{Q}_{h}^{\pi}-Q_{h}^{\pi} & =r_{h}+\widehat{P}^{\pi_{h+1}} \widehat{Q}_{h+1}^{\pi}-r_{h}-P^{\pi_{h+1}} Q_{h+1}^{\pi} \\
& =\left(\widehat{P}^{\pi_{h+1}}-P^{\pi_{h+1}}\right) \widehat{Q}_{h+1}^{\pi}+P^{\pi_{h+1}}\left(Q_{h+1}^{\pi}-Q_{h+1}^{\pi}\right) \\
& =(\widehat{P}-P) \widehat{V}_{h+1}^{\pi}+P^{\pi_{h+1}}\left(\widehat{Q}_{h+1}^{\pi}-Q_{h+1}^{\pi}\right) \\
& =\ldots=\sum_{t=h}^{H} \Gamma_{h+1: t}^{\pi}(\widehat{P}-P) \widehat{V}_{t+1}^{\pi} \\
& \leq \sum_{t=h}^{H} \Gamma_{h+1: t}^{\pi}\left|(\widehat{P}-P) \widehat{V}_{t+1}^{\pi}\right| \\
& \leq \sum_{t=h}^{H} 1 \cdot \max _{s, a}\|(\widehat{P}-P)(\cdot \mid s, a)\|_{1} \cdot\left\|\widehat{V}_{t+1}^{\pi}\right\|_{\infty} \cdot \mathbf{1} \\
& \leq H^{2} \cdot \max _{s, a}\|(\widehat{P}-P)(\cdot \mid s, a)\|_{1} \cdot \mathbf{1} \leq C \cdot H^{2} \sqrt{\frac{S \log (S A / \delta)}{N}} \mathbf{1}
\end{aligned}
$$

with probability $1-\delta$, where the last inequality is by Lemma J.9. By symmetry and taking the $\|\cdot\|_{\infty}$, we obtain w.p. $1-\delta$

$$
\sup _{\pi \in \Pi_{g}, h \in[H]}\left\|\widehat{Q}_{h}^{\pi}-Q_{h}^{\pi}\right\|_{\infty} \leq C \cdot H^{2} \sqrt{\frac{S \log (S A / \delta)}{N}}
$$

The above holds for $\forall \pi \in \Pi_{g}$ since Lemma J.9 acts on $\|\widehat{P}-P\|_{1}$ and is irrelevant to $\pi$.


[^0]:    ${ }^{9}$ We translate [Zhang et al. 2020c] their dimension-free result to $\tilde{O}\left(H^{2} S^{2} A / \epsilon^{2}\right)$ under the standard assumption $r \in[0,1]$.

[^1]:    ${ }^{10}$ For $A$ a matrix and $x$ a vector we have $\|A x\|_{\infty} \leq\|A\|_{\infty}\|x\|_{\infty}$.
    ${ }^{11}$ Here we use $\widehat{V}_{h}^{\star}$ instead of $\widehat{V}_{t}^{\star}$ since we later have $\widehat{V}_{h,\left\{s, u_{t}^{\star}\right\}}^{\star}$. We avoid the same $t$ twice in the expression to prevent confusion.

[^2]:    ${ }^{12}$ The exact sufficient condition for applying Lemma J. 1 is $n>1 / \sum_{t=1}^{H} d_{t}(s, a) \cdot \log (H S A / \delta)$ for all $s, a$. However, since $\sum_{t=1}^{H} d_{t}(s, a) \geq H d_{m} \geq d_{m}$, our condition $n>1 / d_{m} \cdot \log (H S A / \delta)$ used here is a much stronger version thus Lemma J. 1 apply.

[^3]:    ${ }^{13}$ By Lemma J.1 the inequality holds as long as $n \geq 1 / \sum_{h=1}^{2} d_{h}^{\mu}(s, a) \log (S A / \delta)$, here $n>\left(1 / 2 d_{m}\right)$. $\log (S A / \delta)$ is a stronger sufficient condition.

[^4]:    ${ }^{14}$ Note the use of Lemma J.10 also works for any rewards since the only high probability result they used is for $\|P-\hat{P}\|_{1}$. Therefore conditional on the concentration for $\|P-\hat{P}\|_{1}$, the argument follows for any arbitrary reward as well.

[^5]:    ${ }^{15}$ We use the integral only to denote $\mathcal{S}$ could be exponentially large.

[^6]:    ${ }^{16}$ Here the anchor version means for any $(s, a)$ we can apply $\left\|\widehat{P}_{s, a}-P_{s, a}\right\|_{1}=\left\|\sum_{k} \lambda_{k}^{s, a}\left(\widehat{P}_{s, a}-P_{s, a}\right)\right\|_{1} \leq$ $\sum_{k} \lambda_{k}^{s, a}\left\|\widehat{P}_{s, a}-P_{s, a}\right\|_{1}$.

