## A Notation

Below, we introduce some notations that will be used in appendices.

| notation | definition |
| :---: | :---: |
| $\begin{gathered} \mathbf{n}^{k} \\ \mathbf{N}_{h}^{k}(a, \tilde{a}) \\ \mathbf{M}_{h}^{k}(o, a, \tilde{a}) \\ \theta \\ \theta^{\star} \\ \operatorname{POMDP}(\theta) \\ \tau_{h}{ }^{4} \\ \Gamma(\pi, h)^{5} \\ \mathbf{b}\left(\tau_{h} ; \theta\right) \\ \mathbb{P}_{\theta}^{\pi}\left(s_{h}=s\right) \end{gathered}$ | value of $\mathbf{n}$ after the update in the $k^{\text {th }}$ iteration of Algorithm 1 value of $\mathbf{N}_{h}(a, \tilde{a})$ after the update in the $k^{\text {th }}$ iteration of Algorithm 1 value of $\mathbf{M}_{h}(o, a, \tilde{a})$ after the update in the $k^{\text {th }}$ iteration of Algorithm 1 a parameter triple $\left(\mathbb{T}, \mathbb{O}, \mu_{1}\right)$ of a POMDP the groundtruth POMDP parameter triple $\operatorname{POMDP}\left(H, \mathscr{S}, \mathscr{A}, \mathscr{O}, \mathbb{T}, \mathbb{O}, r, \mu_{1}\right)$ <br> a length- $h$ trajectory: $\tau_{h}=\left[a_{h}, o_{h}, \ldots, a_{1}, o_{1}\right] \in(\mathscr{A} \times \mathscr{O})^{h}$ $\begin{gathered} \left\{\tau_{h}=\left(a_{h}, o_{h}, \ldots, a_{1}, o_{1}\right) \mid \pi\left(a_{h}, \ldots, a_{1} \mid o_{h}, \ldots, o_{1}\right)=1\right\} . \\ \mathbf{B}_{h}\left(a_{h}, o_{h} ; \theta\right) \cdots \mathbf{B}_{1}\left(a_{1}, o_{1} ; \theta\right) \cdot \mathbf{b}_{0}(\theta) \end{gathered}$ <br> probability of visiting state $s$ at $h^{\text {th }}$ step when executing policy $\pi$ on $\operatorname{POMDP}(\theta)$ |
| $\begin{gathered} \mathbf{1}(x=y) \\ \mathbf{e}_{o} \\ (\mathbf{X})_{o} \\ \mathbf{I}_{n} \\ C_{\text {poly }} \\ \iota \end{gathered}$ | $\begin{gathered} \text { equal to } 1 \text { if } x=y \text { and } 0 \text { otherwise. } \\ \text { an } O \text {-dimensional vector with }\left(\mathbf{e}_{o}\right)_{i}=\mathbf{1}(o=i) \\ \text { the } o^{\text {th }} \text { column of matrix } \mathbf{X} \\ n \times n \text { identity matrix } \\ \operatorname{poly}(S, O, A, H, 1 / \alpha, \log (1 / \delta)) \\ \log (A O H K / \delta) \end{gathered}$ |

Let $\mathbf{x} \in \mathbb{R}^{n_{x}}, \mathbf{y} \in \mathbb{R}^{n_{y}}$ and $\mathbf{z} \in \mathbb{R}^{n_{z}}$. We denote by $\mathbf{x} \otimes \mathbf{y} \otimes \mathbf{z}$ the tensor product of vectors $\mathbf{x}, \mathbf{y}$ and $\mathbf{z}$, an $n_{x} \times n_{y} \times n_{z}$ tensor with $(i, j, k)^{\text {th }}$ entry equal to $\mathbf{x}_{i} \mathbf{y}_{j} \mathbf{z}_{k}$. Let $\mathbf{X} \in \mathbb{R}^{n_{X} \times m}$, $\mathbf{Y} \in \mathbb{R}^{n_{Y} \times m}$ and $\mathbf{Z} \in \mathbb{R}^{n_{Z} \times m}$. We generalize the notation of tensor product to matrices by defining $\mathbf{X} \otimes \mathbf{Y} \otimes \mathbf{Z}=\sum_{l=1}^{m}(\mathbf{X})_{l} \otimes(\mathbf{Y})_{l} \otimes(\mathbf{Z})_{l}$, which is an $n_{X} \times n_{Y} \times n_{Z}$ tensor with $(i, j, k)^{\text {th }}$ entry equal to $\sum_{l=1}^{m} \mathbf{X}_{i l} \mathbf{Y}_{j l} \mathbf{Z}_{k l}$.
Let $X$ be a random variable taking value in $[m]$, we denote by $\mathbb{P}(X=\cdot)$ an $m$-dimensional vector whose $i^{\text {th }}$ entry is $\mathbb{P}(X=i)$.

[^0]
## B Proof of Hardness Results

The hard examples constructed below are variants of the ones used in [19].
Proposition 1. For any algorithm $\mathfrak{A}$, there exists an overcomplete POMDP $(S>O)$ with $S$ and $O$ being small constants, which satisfies $\sigma_{\min }\left(\mathbb{O}_{h}\right)=1$ for all $h \in[H]$, such that algorithm $\mathfrak{A}$ requires at least $\Omega\left(A^{H-1}\right)$ samples to ensure learning a (1/4)-optimal policy with probability at least $1 / 2$.

Proof. Consider the following $H$-step nonstationary POMDP:

1. STATE There are four states: two good states $g_{1}$ and $g_{2}$ and two bad states $b_{1}$ and $b_{2}$. The initial state is picked uniformly at random.
2. OBSERVATION There are only two different observations $u_{1}$ and $u_{2}$. At step $h \in[H-1]$, we always observe $u_{1}$ at $g_{1}$ and $b_{1}$, and observe $u_{2}$ at $g_{2}$ and $b_{2}$. At step $H$, we always observe $u_{1}$ at good states and $u_{2}$ at bad states. It's direct to verify $\sigma_{\min }\left(\mathbb{O}_{h}\right)=1$ for all $h \in[H]$.
3. REWARD There is no reward at the fist $H-1$ steps (i.e. $r_{h}=0$ for all $\left.h \in[H-1]\right)$. At step $H$, we receive reward 1 if we observe $u_{1}$ and no reward otherwise (i.e. $r_{H}(o)=\mathbf{1}\left(o=u_{1}\right)$ ).
4. TRANSITION There is one good action $a_{h}^{\star}$ and $A-1$ bad actions for each $h \in[H-1]$. At step $h \in[H-1]$, suppose we are at a good state $\left(g_{1}\right.$ or $\left.g_{2}\right)$, then we will transfer to $g_{1}$ or $g_{2}$ uniformly at random if we take $a_{h}^{\star}$ and otherwise transfer to $b_{1}$ or $b_{2}$ uniformly at random. In contrast, if we are at a bad state ( $b_{1}$ or $b_{2}$ ), we will always transfer to $b_{1}$ or $b_{2}$ uniformly at random no matter what action we take. Note that two good (bad) states are equivalent in terms of transition.

We have the following key observations:

1. Once we are at bad states, we always stay at bad states.
2. We have

$$
\begin{aligned}
& \mathbb{P}\left(o_{1: H-1}=z \mid a_{1: H-1}, o_{H}\right)=\frac{1}{2^{H-1}} \\
& \text { for any } z \in\left\{u_{1}, u_{2}\right\}^{H-1} \text { and }\left(a_{1: H-1}, o_{H}\right) \in[A]^{H-1} \times\left\{u_{1}, u_{2}\right\}
\end{aligned}
$$

Therefore, the observations at the first $H-1$ steps provide no information about the underlying transition. The only useful information is the last observation $o_{H}$ which tells us whether we end in good states or not.
3. The optimal policy is unique and is to execute the good action sequence $\left(a_{1}^{\star}, \ldots, a_{H-1}^{\star}\right)$ regardless of the obervations.

Based on the observations above, this is equivalent to a multi-arm bandits problem with $A^{H-1}$ arms. Therefore, we cannot do better than Brute-force search, which has sample complexity at least $\Omega\left(A^{H-1}\right)$.
Proposition 2. For any algorithm $\mathfrak{A}$, there exists an undercomplete POMDP $(S \leq O)$ with $S$ and $O$ being small constants, such that algorithm $\mathfrak{A}$ requires at least $\Omega\left(A^{H-1}\right)$ samples to ensure learning a (1/4)-optimal policy with probability at least $1 / 2$.

Proof. We continue to use the POMDP constructed in Proposition 1 and slightly modify it by splitting $u_{2}$ into another 4 different observations $\left\{q_{1}, q_{2}, q_{3}, q_{4}\right\}$, so in the new $\operatorname{POMDP}(O=5>S=4)$, we will observe a $q_{i}$ picked uniformly at random from $\left\{q_{1}, q_{2}, q_{3}, q_{4}\right\}$ when we are 'supposed' to observe $u_{2}$. It's easy to see the modification does not change its hardness.

## C Analysis of OMM-UCB

Throughout the proof, we use $\tau_{h}$ to denote a length- $h$ trajectory: $\left[a_{h}, o_{h}, \ldots, a_{1}, o_{1}\right] \in(\mathscr{A} \times \mathscr{O})^{h}$. Note that this definition is different from the one used in Section5, where $\tau_{h}=\left[o_{h}, \ldots, a_{1}, o_{1}\right] \in$ $\mathscr{O} \times(\mathscr{A} \times \mathscr{O})^{h-1}$ does not include the action $a_{h}$ at $h^{\text {th }}$ step. Besides, we define $\Gamma(\pi, h)=\left\{\tau_{h}=\right.$ $\left.\left(a_{h}, o_{h}, \ldots, a_{1}, o_{1}\right) \mid \pi\left(a_{h}, \ldots, a_{1} \mid o_{h}, \ldots, o_{1}\right)=1\right\}$, which is also different from the definition in Section 5 wherer $a_{h}$ is not included.
Please refer to Appendix A for definitions of frequently used notations.

## C. 1 Bounding the error in belief states

In this subsection, we will bound the error in (unnormalized) belief states, i.e., $\mathbf{b}\left(\tau_{h} ; \theta\right)-\mathbf{b}\left(\tau_{h} ; \hat{\theta}\right)$ by the error in operators reweighed by the probability distribution of visited states.
We start by proving the following lemma that helps us decompose the error in belief states inductively.
Lemma 7. Given a deterministic policy $\pi$ and two set of POMDP parameters $\hat{\theta}=\left(\hat{\mathbb{O}}, \hat{\mathbb{T}}, \hat{\mu}_{1}\right)$ and $\theta=\left(\mathbb{O}, \mathbb{T}, \mu_{1}\right)$, for all $h \geq 1$ and $\mathbf{X} \in\left\{\mathbf{I}_{O}, \hat{\mathbb{O}}_{h+1}^{\dagger}\right\}$, we have

$$
\begin{aligned}
\sum_{\tau_{h} \in \Gamma(\pi, h)}\left\|\mathbf{X}\left(\mathbf{b}\left(\tau_{h} ; \theta\right)-\mathbf{b}\left(\tau_{h} ; \hat{\theta}\right)\right)\right\|_{1} \leq & \sum_{\tau_{h-1} \in \Gamma(\pi, h-1)}\left\|\hat{\mathbb{O}}_{h}^{\dagger}\left(\mathbf{b}\left(\tau_{h-1} ; \theta\right)-\mathbf{b}\left(\tau_{h-1} ; \hat{\theta}\right)\right)\right\|_{1} \\
& +\sum_{\tau_{h} \in \Gamma(\pi, h)}\left\|\mathbf{X}\left(\mathbf{B}_{h}\left(a_{h}, o_{h} ; \hat{\theta}\right)-\mathbf{B}_{h}\left(a_{h}, o_{h} ; \theta\right)\right) \mathbf{b}\left(\tau_{h-1} ; \theta\right)\right\|_{1}
\end{aligned}
$$

Proof. By the definition of $\mathbf{b}\left(\tau_{h} ; \theta\right)$ and $\mathbf{b}\left(\tau_{h} ; \hat{\theta}\right)$,

$$
\begin{aligned}
& \sum_{\tau_{h} \in \Gamma(\pi, h)}\left\|\mathbf{X}\left(\mathbf{b}\left(\tau_{h} ; \theta\right)-\mathbf{b}\left(\tau_{h} ; \hat{\theta}\right)\right)\right\|_{1} \\
= & \sum_{\tau_{h} \in \Gamma(\pi, h)}\left\|\mathbf{X}\left(\mathbf{B}_{h}\left(a_{h}, o_{h} ; \theta\right) \mathbf{b}\left(\tau_{h-1} ; \theta\right)-\mathbf{B}_{h}\left(a_{h}, o_{h} ; \hat{\theta}\right) \mathbf{b}\left(\tau_{h-1} ; \hat{\theta}\right)\right)\right\|_{1} \\
\leq & \sum_{\tau_{h} \in \Gamma(\pi, h)}\left\|\mathbf{X} \mathbf{B}_{h}\left(a_{h}, o_{h} ; \hat{\theta}\right)\left(\mathbf{b}\left(\tau_{h-1} ; \theta\right)-\mathbf{b}\left(\tau_{h-1} ; \hat{\theta}\right)\right)\right\|_{1} \\
+ & \sum_{\tau_{h} \in \Gamma(\pi, h)}\left\|\mathbf{X}\left(\mathbf{B}_{h}\left(a_{h}, o_{h} ; \hat{\theta}\right)-\mathbf{B}_{h}\left(a_{h}, o_{h} ; \theta\right)\right) \mathbf{b}\left(\tau_{h-1} ; \theta\right)\right\|_{1}
\end{aligned}
$$

The first term can be bounded as following,

$$
\begin{aligned}
& \sum_{\tau_{h} \in \Gamma(\pi, h)}\left\|\mathbf{X B} \mathbf{B}_{h}\left(a_{h}, o_{h} ; \hat{\theta}\right)\left(\mathbf{b}\left(\tau_{h-1} ; \theta\right)-\mathbf{b}\left(\tau_{h-1} ; \hat{\theta}\right)\right)\right\|_{1} \\
= & \sum_{\tau_{h} \in \Gamma(\pi, h)}\left\|\mathbf{X} \hat{\mathbb{O}}_{h+1} \hat{\mathbb{T}}_{h}\left(a_{h}\right) \operatorname{diag}\left(\hat{\mathbb{O}}_{h}\left(o_{h} \mid \cdot\right)\right) \hat{\mathbb{O}}_{h}^{\dagger}\left(\mathbf{b}\left(\tau_{h-1} ; \theta\right)-\mathbf{b}\left(\tau_{h-1} ; \hat{\theta}\right)\right)\right\|_{1} \\
\leq & \sum_{\tau_{h} \in \Gamma(\pi, h)} \sum_{i}\left\|\left(\mathbf{X} \hat{\mathbb{O}}_{h+1} \hat{\mathbb{T}}_{h}\left(a_{h}\right) \operatorname{diag}\left(\hat{\mathbb{O}}_{h}\left(o_{h} \mid \cdot\right)\right)\right)_{i}\right\|_{1}\left|\left(\hat{\mathbb{O}}_{h}^{\dagger}\left(\mathbf{b}\left(\tau_{h-1} ; \theta\right)-\mathbf{b}\left(\tau_{h-1} ; \hat{\theta}\right)\right)\right)_{i}\right| \\
= & \sum_{\tau_{h} \in \Gamma(\pi, h)} \sum_{i}\left\|\left(\mathbf{X}_{\widehat{\mathbb{O}}_{h+1}} \hat{\mathbb{T}}_{h}\left(a_{h}\right)\right)_{i}\right\|_{1} \hat{\mathbb{O}}_{h}\left(o_{h} \mid i\right) \mid\left(\hat{\mathbb{O}}_{h}^{\dagger}\left(\mathbf{b}\left(\tau_{h-1} ; \theta\right)-\mathbf{b}\left(\tau_{h-1} ; \hat{\theta}\right)\right)_{i} \mid\right. \\
= & \sum_{\tau_{h} \in \Gamma(\pi, h)} \sum_{i} \hat{\mathbb{O}}_{h}\left(o_{h} \mid i\right)\left|\left(\hat{\mathbb{O}}_{h}^{\dagger}\left(\mathbf{b}\left(\tau_{h-1} ; \theta\right)-\mathbf{b}\left(\tau_{h-1} ; \hat{\theta}\right)\right)\right)_{i}\right|
\end{aligned}
$$

where the inequality is by triangle inequality, and the last identity follows from $\hat{\mathbb{T}}_{h}\left(a_{h}\right)$ (when $\mathbf{X}=\hat{\mathbb{O}}_{h+1}^{\dagger}$ ) and $\hat{\mathbb{O}}_{h+1} \hat{\mathbb{T}}_{h}\left(a_{h}\right)$ (when $\mathbf{X}=\mathbf{I}_{O}$ ) having columns with $\ell_{1}$-norm equal to 1 .

As $\pi$ is deterministic, $a_{h}$ is unique given $\tau_{h-1}$ and $o_{h}$. Therefore,

$$
\begin{aligned}
& \sum_{\tau_{h} \in \Gamma(\pi, h)} \sum_{i} \hat{\mathbb{O}}_{h}\left(o_{h} \mid i\right)\left|\left(\hat{\mathbb{O}}_{h}^{\dagger}\left(\mathbf{b}\left(\tau_{h-1} ; \theta\right)-\mathbf{b}\left(\tau_{h-1} ; \hat{\theta}\right)\right)\right)_{i}\right| \\
= & \sum_{\tau_{h-1} \in \Gamma(\pi, h-1)} \sum_{o_{h}} \sum_{i} \hat{\mathbb{O}}_{h}\left(o_{h} \mid i\right)\left|\left(\hat{\mathbb{O}}_{h}^{\dagger}\left(\mathbf{b}\left(\tau_{h-1} ; \theta\right)-\mathbf{b}\left(\tau_{h-1} ; \hat{\theta}\right)\right)\right)_{i}\right| \\
= & \sum_{\tau_{h-1} \in \Gamma(\pi, h-1)} \sum_{i} \sum_{o_{h}} \hat{\mathbb{O}}_{h}\left(o_{h} \mid i\right)\left|\left(\hat{\mathbb{O}}_{h}^{\dagger}\left(\mathbf{b}\left(\tau_{h-1} ; \theta\right)-\mathbf{b}\left(\tau_{h-1} ; \hat{\theta}\right)\right)\right)_{i}\right| \\
= & \sum_{\tau_{h-1} \in \Gamma(\pi, h-1)} \sum_{i}\left|\left(\hat{\mathbb{O}}_{h}^{\dagger}\left(\mathbf{b}\left(\tau_{h-1} ; \theta\right)-\mathbf{b}\left(\tau_{h-1} ; \hat{\theta}\right)\right)\right)_{i}\right| \\
= & \sum_{\tau_{h-1} \in \Gamma(\pi, h-1)}\left\|\hat{\mathbb{O}}_{h}^{\dagger}\left(\mathbf{b}\left(\tau_{h-1} ; \theta\right)-\mathbf{b}\left(\tau_{h-1} ; \hat{\theta}\right)\right)\right\|_{1},
\end{aligned}
$$

which completes the proof.
By applying Lemma 7 inductively, we can bound the error in belief states by the projection of errors in operators on preceding belief states.
Lemma 8. Given a deterministic policy $\pi$ and two sets of undercomplete POMDP parameters $\theta=\left(\mathbb{O}, \mathbb{T}, \mu_{1}\right)$ and $\hat{\theta}=\left(\hat{\mathbb{O}}, \hat{\mathbb{T}}, \hat{\mu}_{1}\right)$ with $\sigma_{\min }(\hat{\mathbb{O}}) \geq \alpha$, for all $h \geq 1$, we have

$$
\begin{aligned}
& \sum_{\tau_{h} \in \Gamma(\pi, h)}\left\|\mathbf{b}\left(\tau_{h} ; \theta\right)-\mathbf{b}\left(\tau_{h} ; \hat{\theta}\right)\right\|_{1} \\
\leq & \frac{\sqrt{S}}{\alpha} \sum_{j=1}^{h} \sum_{\tau_{j} \in \Gamma(\pi, j)}\left\|\left(\mathbf{B}_{j}\left(a_{j}, o_{j} ; \hat{\theta}\right)-\mathbf{B}_{j}\left(a_{j}, o_{j} ; \theta\right)\right) \mathbf{b}\left(\tau_{j-1} ; \theta\right)\right\|_{1}+\frac{\sqrt{S}}{\alpha}\left\|\mathbf{b}_{0}(\theta)-\mathbf{b}_{0}(\hat{\theta})\right\|_{1} .
\end{aligned}
$$

Proof. Invoking Lemma 7 with $\mathbf{X}=\hat{\mathbb{O}}_{j+1}^{\dagger}$, we have

$$
\begin{align*}
& \sum_{\tau_{j} \in \Gamma(\pi, j)}\left\|\hat{\mathbb{O}}_{j+1}^{\dagger}\left(\mathbf{b}\left(\tau_{j} ; \theta\right)-\mathbf{b}\left(\tau_{j} ; \hat{\theta}\right)\right)\right\|_{1} \leq \sum_{\tau_{j-1} \in \Gamma(\pi, j-1)}\left\|\hat{\mathbb{O}}_{j}^{\dagger}\left(\mathbf{b}\left(\tau_{j-1} ; \theta\right)-\mathbf{b}\left(\tau_{j-1} ; \hat{\theta}\right)\right)\right\|_{1} \\
&+\sum_{\tau_{j} \in \Gamma(\pi, j)}\left\|\hat{\mathbb{O}}_{j+1}^{\dagger}\left(\mathbf{B}_{j}\left(a_{j}, o_{j} ; \hat{\theta}\right)-\mathbf{B}_{j}\left(a_{j}, o_{j} ; \theta\right)\right) \mathbf{b}\left(\tau_{j-1} ; \theta\right)\right\|_{1} \tag{8}
\end{align*}
$$

Summing (8) over $j=1, \ldots, h-1$, we obtain

$$
\begin{align*}
& \sum_{\tau_{h-1} \in \Gamma(\pi, h-1)}\left\|\hat{\mathbb{O}}_{h}^{\dagger}\left(\mathbf{b}\left(\tau_{h-1} ; \theta\right)-\mathbf{b}\left(\tau_{h-1} ; \hat{\theta}\right)\right)\right\|_{1}  \tag{9}\\
\leq & \sum_{j=1}^{h-1} \sum_{\tau_{j} \in \Gamma(\pi, j)}\left\|\hat{\mathbb{O}}_{j+1}^{\dagger}\left(\mathbf{B}_{j}\left(a_{j}, o_{j} ; \hat{\theta}\right)-\mathbf{B}_{j}\left(a_{j}, o_{j} ; \theta\right)\right) \mathbf{b}\left(\tau_{j-1} ; \theta\right)\right\|_{1}+\left\|\hat{\mathbb{O}}_{1}^{\dagger}\left(\mathbf{b}_{0}(\theta)-\mathbf{b}_{0}(\hat{\theta})\right)\right\|_{1}
\end{align*}
$$

Again, invoking Lemma 7 with $\mathbf{X}=\mathbf{I}_{O}$ gives

$$
\begin{align*}
& \sum_{\tau_{h} \in \Gamma(\pi, h)}\left\|\mathbf{b}\left(\tau_{h} ; \theta\right)-\mathbf{b}\left(\tau_{h} ; \hat{\theta}\right)\right\|_{1} \leq \sum_{\tau_{h-1} \in \Gamma(\pi, h-1)}\left\|\hat{\mathbb{O}}_{h}^{\dagger}\left(\mathbf{b}\left(\tau_{h-1} ; \theta\right)-\mathbf{b}\left(\tau_{h-1} ; \hat{\theta}\right)\right)\right\|_{1} \\
&+\sum_{\tau_{h} \in \Gamma(\pi, h)}\left\|\left(\mathbf{B}_{h}\left(a_{h}, o_{h} ; \hat{\theta}\right)-\mathbf{B}_{h}\left(a_{h}, o_{h} ; \theta\right)\right) \mathbf{b}\left(\tau_{h-1} ; \theta\right)\right\|_{1} \tag{10}
\end{align*}
$$

Plugging (9) into (10), and using the fact that $\left\|\hat{\mathbb{O}}_{h}^{\dagger}\right\|_{1 \rightarrow 1} \leq \sqrt{S}\left\|\hat{\mathbb{O}}_{h}^{\dagger}\right\|_{2} \leq \frac{\sqrt{S}}{\alpha}$ complete the proof.
The following lemma bounds the projection of any vector on belief states by its projection on the product of the observation matrix and the transition matrix, reweighed by the visitation probability of states.

Lemma 9. For any deterministic policy $\pi$, fixed $a_{h+1} \in \mathscr{A}, \mathbf{u} \in \mathbb{R}^{O}$, and $h \geq 0$, we have

$$
\sum_{o_{h+1} \in \mathscr{O}} \sum_{\tau_{h} \in \Gamma(\pi, h)}\left|\mathbf{u}^{\top} \mathbf{b}\left(\left[a_{h+1}, o_{h+1}, \tau_{h}\right] ; \theta\right)\right| \leq \sum_{s=1}^{S}\left|\mathbf{u}^{\top}\left(\mathbb{O}_{h+2} \mathbb{T}_{h+1}\left(a_{h+1}\right)\right)_{s}\right| \mathbb{P}_{\theta}^{\pi}\left(s_{h+1}=s\right)
$$

Proof. By definition, for any $\left[a_{h+1}, o_{h+1}, \tau_{h}\right] \in \mathscr{A} \times \mathscr{O} \times \Gamma(\pi, h)$, we have

$$
\mathbf{b}\left(\left[a_{h+1}, o_{h+1}, \tau_{h}\right] ; \theta\right)=\mathbb{O}_{h+2} \mathbb{T}_{h+1}\left(a_{h+1}\right) \mathbb{P}_{\theta}^{\pi}\left(s_{h+1}=\cdot,\left[o_{h+1}, \tau_{h}\right]\right)
$$

where $\mathbb{P}_{\theta}^{\pi}\left(s_{h+1}=\cdot,\left[o_{h+1}, \tau_{h}\right]\right)$ is an $s$-dimensional vector, whose $i^{\text {th }}$ entry is equal to the probability of observing $\left[o_{h+1}, \tau_{h}\right]$ and reaching state $i$ at step $h+1$ when executing policy $\pi$ in $\operatorname{POMDP}(\theta)$.
Therefore,

$$
\begin{aligned}
& \sum_{\tau_{h} \in \Gamma(\pi, h)} \sum_{o_{h+1} \in \mathscr{O}}\left|\mathbf{u}^{\top} \mathbf{b}\left(\left[a_{h+1}, o_{h+1}, \tau_{h}\right] ; \theta\right)\right| \\
= & \sum_{\tau_{h} \in \Gamma(\pi, h)} \sum_{o_{h+1} \in \mathscr{O}}\left|\mathbf{u}^{\top} \mathbb{O}_{h+2} \mathbb{T}_{h+1}\left(a_{h+1}\right) \mathbb{P}_{\theta}^{\pi}\left(s_{h+1}=\cdot,\left[o_{h+1}, \tau_{h}\right]\right)\right| \\
\leq & \sum_{\tau_{h} \in \Gamma(\pi, h)} \sum_{o_{h+1} \in \mathscr{O}} \sum_{s=1}^{S}\left|\mathbf{u}^{\top}\left(\mathbb{O}_{h+2} \mathbb{T}_{h+1}\left(a_{h+1}\right)\right)_{s}\right| \mathbb{P}_{\theta}^{\pi}\left(s_{h+1}=s,\left[o_{h+1}, \tau_{h}\right]\right) \\
= & \sum_{s=1}^{S}\left|\mathbf{u}^{\top}\left(\mathbb{O}_{h+2} \mathbb{T}_{h+1}\left(a_{h+1}\right)\right)_{s}\right|\left(\sum_{\tau_{h} \in \Gamma(\pi, h)} \sum_{o_{h+1} \in \mathscr{O}} \mathbb{P}_{\theta}^{\pi}\left(s_{h+1}=s,\left[o_{h+1}, \tau_{h}\right]\right)\right) \\
= & \sum_{s=1}^{S}\left|\mathbf{u}^{\top}\left(\mathbb{O}_{h+2} \mathbb{T}_{h+1}\left(a_{h+1}\right)\right)_{s}\right| \mathbb{P}_{\theta}^{\pi}\left(s_{h+1}=s\right) .
\end{aligned}
$$

Combining Lemma 8 and Lemma9, we obtain the target bound.
Lemma 10. Given a deterministic policy $\pi$ and two sets of undercomplete POMDP parameters $\theta=\left(\mathbb{O}, \mathbb{T}, \mu_{1}\right)$ and $\hat{\theta}=\left(\hat{\mathbb{O}}, \hat{\mathbb{T}}, \hat{\mu}_{1}\right)$ with $\sigma_{\min }(\hat{\mathbb{O}}) \geq \alpha$, for all $h \geq 1$, we have

$$
\begin{aligned}
& \sum_{\tau_{h} \in \Gamma(\pi, h)}\left\|\mathbf{b}\left(\tau_{h} ; \theta\right)-\mathbf{b}\left(\tau_{h} ; \hat{\theta}\right)\right\|_{1} \\
\leq & \frac{\sqrt{S}}{\alpha}\left\|\mathbf{b}_{0}(\theta)-\mathbf{b}_{0}(\hat{\theta})\right\|_{1}+\frac{\sqrt{S}}{\alpha} \sum_{(a, o) \in \mathscr{A} \times \mathscr{O}}\left\|\left(\mathbf{B}_{1}(a, o ; \hat{\theta})-\mathbf{B}_{1}(a, o ; \theta)\right) \mathbf{b}_{0}(\theta)\right\|_{1} \\
& +\frac{\sqrt{S}}{\alpha} \sum_{j=2}^{h} \sum_{(a, \tilde{a}, o) \in \mathscr{A}^{2} \times \mathscr{O}} \sum_{s=1}^{S}\left\|\left(\mathbf{B}_{j}(a, o ; \hat{\theta})-\mathbf{B}_{j}(a, o ; \theta)\right)\left(\mathbb{O}_{j} \mathbb{T}_{j-1}(\tilde{a})\right)_{s}\right\|_{1} \mathbb{P}_{\theta}^{\pi}\left(s_{j-1}=s\right) .
\end{aligned}
$$

Proof. By Lemma 8 ,

$$
\begin{align*}
& \sum_{\tau_{h} \in \Gamma(\pi, h)}\left\|\mathbf{b}\left(\tau_{h} ; \theta\right)-\mathbf{b}\left(\tau_{h} ; \hat{\theta}\right)\right\|_{1} \\
\leq & \frac{\sqrt{S}}{\alpha} \sum_{j=2}^{h} \sum_{\tau_{j} \in \Gamma(\pi, j)}\left\|\left(\mathbf{B}_{j}\left(a_{j}, o_{j} ; \hat{\theta}\right)-\mathbf{B}_{j}\left(a_{j}, o_{j} ; \theta\right)\right) \mathbf{b}\left(\tau_{j-1} ; \theta\right)\right\|_{1} \\
& +\frac{\sqrt{S}}{\alpha} \sum_{\tau_{1} \in \Gamma(\pi, 1)}\left\|\left(\mathbf{B}_{1}\left(a_{1}, o_{1} ; \hat{\theta}\right)-\mathbf{B}_{1}\left(a_{1}, o_{1} ; \hat{\theta}\right)\right) \mathbf{b}_{0}(\theta)\right\|_{1}+\frac{\sqrt{S}}{\alpha}\left\|\mathbf{b}_{0}(\theta)-\mathbf{b}_{0}(\hat{\theta})\right\|_{1} . \tag{11}
\end{align*}
$$

Bounding the first term: note that $\Gamma(\pi, j) \subseteq \Gamma(\pi, j-2) \times(\mathscr{O} \times \mathscr{A})^{2}$, so we have

$$
\begin{align*}
& \sum_{\tau_{j} \in \Gamma(\pi, j)}\left\|\left(\mathbf{B}_{j}\left(a_{j}, o_{j} ; \hat{\theta}\right)-\mathbf{B}_{j}\left(a_{j}, o_{j} ; \theta\right)\right) \mathbf{b}\left(\tau_{j-1} ; \theta\right)\right\|_{1} \\
& \leq \sum_{\tau_{j-2} \in \Gamma(\pi, j-2)} \sum_{o_{j-1} \in \mathscr{O}} \sum_{a_{j-1} \in \mathscr{A}} \sum_{o_{j} \in \mathscr{O}} \sum_{a_{j} \in \mathscr{A}} \\
&\left\|\left(\mathbf{B}_{j}\left(a_{j}, o_{j} ; \hat{\theta}\right)-\mathbf{B}_{j}\left(a_{j}, o_{j} ; \theta\right)\right) \mathbf{b}\left(\left[a_{j-1}, o_{j-1}, \tau_{j-2}\right] ; \theta\right)\right\|_{1} \\
&= \sum_{\left(a_{j}, a_{j-1}, o_{j}\right) \in \mathscr{A}^{2} \times \mathscr{O}} \\
& \underbrace{}_{\underbrace{\tau_{j-2} \in \Gamma(\pi, j-2)}} \sum_{o_{j-1} \in \mathscr{O}}\left\|\left(\mathbf{B}_{j}\left(a_{j}, o_{j} ; \hat{\theta}\right)-\mathbf{B}_{j}\left(a_{j}, o_{j} ; \theta\right)\right) \mathbf{b}\left(\left[a_{j-1}, o_{j-1}, \tau_{j-2}\right] ; \theta\right)\right\|_{1} \tag{12}
\end{align*} .
$$

$(\diamond)$
We can bound $(\diamond)$ by Lemma 9 and obtain,

$$
\begin{align*}
& \sum_{\tau_{j} \in \Gamma(\pi, j)}\left\|\left(\mathbf{B}_{j}\left(a_{j}, o_{j} ; \hat{\theta}\right)-\mathbf{B}_{j}\left(a_{j}, o_{j} ; \theta\right)\right) \mathbf{b}\left(\tau_{j-1} ; \theta\right)\right\|_{1} \\
\leq & \sum_{\left(a_{j}, a_{j-1}, o_{j}\right) \in \mathscr{A}^{2} \times \mathscr{O}} \sum_{s=1}^{S}\left\|\left(\mathbf{B}_{j}\left(a_{j}, o_{j} ; \hat{\theta}\right)-\mathbf{B}_{j}\left(a_{j}, o_{j} ; \theta\right)\right)\left(\mathbb{O}_{j} \mathbb{T}_{j-1}\left(a_{j-1}\right)\right)_{s}\right\|_{1} \mathbb{P}_{\theta}^{\pi}\left(s_{j-1}=s\right) \\
= & \sum_{(a, \tilde{a}, o) \in \mathscr{A}^{2} \times \mathscr{O}} \sum_{s=1}^{S}\left\|\left(\mathbf{B}_{j}(a, o ; \hat{\theta})-\mathbf{B}_{j}(a, o ; \theta)\right)\left(\mathbb{O}_{j} \mathbb{T}_{j-1}(\tilde{a})\right)_{s}\right\|_{1} \mathbb{P}_{\theta}^{\pi}\left(s_{j-1}=s\right), \tag{13}
\end{align*}
$$

where the identity only changes the notations $\left(a_{j}, a_{j-1}, o_{j}\right) \rightarrow(a, \tilde{a}, o)$ to make the expression cleaner.

Bounding the second term: note that $\Gamma(\pi, 1) \subseteq \mathscr{O} \times \mathscr{A}$, we have

$$
\begin{align*}
& \sum_{\tau_{1} \in \Gamma(\pi, 1)}\left\|\left(\mathbf{B}_{1}\left(a_{1}, o_{1} ; \theta\right)-\mathbf{B}_{1}\left(a_{1}, o_{1} ; \hat{\theta}\right)\right) \mathbf{b}_{0}(\theta)\right\|_{1} \\
\leq & \sum_{(a, o) \in \mathscr{A} \times \mathscr{O}}\left\|\left(\mathbf{B}_{1}(a, o ; \theta)-\mathbf{B}_{1}(a, o ; \hat{\theta})\right) \mathbf{b}_{0}(\theta)\right\|_{1} \tag{14}
\end{align*}
$$

Plugging (13) and 14 into 11 completes the proof.

## C. 2 A hammer for studying confidence sets

In this subsection, we develop a martingale concentration result, which forms the basis of analyzing confidence sets.
We start by giving the following basic fact about POMDP. The proof is just some basic algebraic calculation so we omit it here.

Fact 11. In $\operatorname{POMDP}(\theta)$, suppose $s_{h-1}$ is sampled from $\mu_{h-1}$, fix $a_{h-1} \equiv \tilde{a}$, and $a_{h} \equiv a$. Then the joint distribution of $\left(o_{h+1}, o_{h}, o_{h-1}\right)$ is

$$
\mathbb{P}\left(o_{h+1}=\cdot, o_{h}=\cdot, o_{h-1}=\cdot\right)=\left(\mathbb{O}_{h+1} \mathbb{T}_{h}(a)\right) \otimes \mathbb{O}_{h} \otimes\left(\mathbb{O}_{h-1} \operatorname{diag}\left(\mu_{h-1}\right) \mathbb{T}_{h-1}(\tilde{a})^{\top}\right)
$$

By slicing the tensor, we can further obtain

$$
\left\{\begin{aligned}
\mathbb{P}\left(o_{h-1}=\cdot\right) & =\mathbb{O}_{h-1} \mu_{h-1} \\
\mathbb{P}\left(o_{h}=\cdot, o_{h-1}=\cdot\right) & =\mathbb{O}_{h} \mathbb{T}_{h-1}(\tilde{a}) \operatorname{diag}\left(\mu_{h-1}\right) \mathbb{O}_{h-1}^{\top} \\
\mathbb{P}\left(o_{h+1}=\cdot, o_{h}=o, o_{h-1}=\cdot\right) & =\mathbb{O}_{h+1} \mathbb{T}_{h}(a) \operatorname{diag}\left(\mathbb{O}_{h}(o \mid \cdot)\right) \mathbb{T}_{h-1}(\tilde{a}) \operatorname{diag}\left(\mu_{h-1}\right) \mathbb{O}_{h-1}^{\top}
\end{aligned}\right.
$$

A simple implication of Fact 11 is that if we execute policy $\pi$ from step 1 to step $h-2$, take $\tilde{a}$ and $a$ at step $h-1$ and $h$ respectively, then the joint distribution of $\left(o_{h+1}, o_{h}, o_{h-1}\right)$ is the same as above except for replacing $\mu_{h-1}$ with $\mathbb{P}_{\theta}^{\pi}\left(s_{h-1}=\cdot\right)$.

Suppose we are given a set of sequential data $\left\{\left(o_{h+1}^{(t)}, o_{h}^{(t)}, o_{h-1}^{(t)}\right)\right\}_{t=1}^{N}$ generated from $\operatorname{POMDP}(\theta)$ in the following way: at time $t$, execute policy $\pi_{t}$ from step 1 to step $h-2$, take action $\tilde{a}$ at step $h-1$, and action $a$ at step $h$ respectively, and observe $\left(o_{h+1}^{(t)}, o_{h}^{(t)}, o_{h-1}^{(t)}\right)$. Here, we allow the policy $\pi_{t}$ to be adversarial, in the sense that $\pi_{t}$ can be chosen based on $\left\{\left(\pi_{i}, o_{h+1}^{(i)}, o_{h}^{(i)}, o_{h-1}^{(i)}\right)\right\}_{i=1}^{t-1}$. Define $\mu_{h-1}^{a d v}=\frac{1}{N} \sum_{t=1}^{N} \mathbb{P}_{\theta}^{\pi_{t}}\left(s_{h-1}=\cdot\right)$. Based on Fact 11 we define the following probability vector, matrices and tensor,

$$
\left\{\begin{aligned}
P_{h-1} & =\mathbb{O}_{h-1} \mu_{h-1}^{a d v} \\
P_{h, h-1} & =\mathbb{O}_{h} \mathbb{T}_{h-1}(\tilde{a}) \operatorname{diag}\left(\mu_{h-1}^{a d v}\right) \mathbb{O}_{h-1}^{\top}, \\
P_{h+1, h, h-1} & =\left(\mathbb{O}_{h+1} \mathbb{T}_{h}(a)\right) \otimes \mathbb{O}_{h} \otimes\left(\mathbb{O}_{h-1} \operatorname{diag}\left(\mu_{h-1}^{a d v}\right) \mathbb{T}_{h-1}(\tilde{a})^{\top}\right) \\
P_{h+1, o, h-1} & =\mathbb{O}_{h+1} \mathbb{T}_{h}(a) \operatorname{diag}\left(\mathbb{O}_{h}(o \mid \cdot)\right) \mathbb{T}_{h-1}(\tilde{a}) \operatorname{diag}\left(\mu_{h-1}^{a d v}\right) \mathbb{O}_{h-1}^{\top}, \quad o \in \mathscr{O} .
\end{aligned}\right.
$$

Accordingly, we define their empirical estimates as below

$$
\left\{\begin{aligned}
\hat{P}_{h-1} & =\frac{1}{N} \sum_{t=1}^{N} \mathbf{e}_{o_{h-1}^{(t)}}, \\
\hat{P}_{h, h-1} & =\frac{1}{N} \sum_{t=1}^{N} \mathbf{e}_{o_{h}^{(t)}} \otimes \mathbf{e}_{o_{h-1}^{(t)}}, \\
\hat{P}_{h+1, h, h-1} & =\frac{1}{N} \sum_{t=1}^{N} \mathbf{e}_{o_{h+1}^{(t)}} \otimes \mathbf{e}_{o_{h}^{(t)}} \otimes \mathbf{e}_{o_{h-1}^{(t)}}, \\
\hat{P}_{h+1, o, h-1} & =\frac{1}{N} \sum_{t=1}^{N} \mathbf{e}_{o_{h+1}^{(t)}} \otimes \mathbf{e}_{o_{h-1}^{(t)}} \mathbf{1}\left(o_{h}^{(t)}=o\right), \quad o \in \mathscr{O}
\end{aligned}\right.
$$

Lemma 12. There exists an absolute constant $c_{1}$, s.t. the following concentration bound holds with probability at least $1-\delta$

$$
\begin{aligned}
\max \{ & \left\|\hat{P}_{h+1, h, h-1}-P_{h+1, h, h-1}\right\|_{F},\left\|\hat{P}_{h, h-1}-P_{h, h-1}\right\|_{F}, \\
& \left.\max _{o \in \mathscr{O}}\left\|\hat{P}_{h+1, o, h-1}-P_{h+1, o, h-1}\right\|_{F},\left\|\hat{P}_{h-1}-P_{h-1}\right\|_{2}\right\} \leq c_{1} \sqrt{\frac{\log (O N / \delta)}{N}} .
\end{aligned}
$$

Proof. We start with proving that with probability at least $1-\delta / 2$,

$$
\left\|\hat{P}_{h+1, h, h-1}-P_{h+1, h, h-1}\right\|_{F} \leq c_{1} \sqrt{\frac{\log (O N / \delta)}{N}}
$$

Let $\mathcal{F}_{t}$ be the $\sigma$-algebra generated by $\left\{\left\{\pi_{i}\right\}_{i=1}^{t+1},\left\{\left(o_{h+1}^{(i)}, o_{h}^{(i)}, o_{h-1}^{(i)}\right)\right\}_{i=1}^{t}\right\} \cdot\left(\mathcal{F}_{t}\right)$ is a filtration. Define

$$
X_{t}=e_{o_{h+1}^{(t)}} \otimes e_{o_{h}^{(t)}} \otimes e_{o_{h-1}^{(t)}}-\left(\mathbb{O}_{h+1} \mathbb{T}_{h}(a)\right) \otimes \mathbb{O}_{h} \otimes\left(\mathbb{O}_{h-1} \operatorname{diag}\left(\mathbb{P}_{\theta}^{\pi_{t}}\left(s_{h-1}=\cdot\right)\right) \mathbb{T}_{h-1}(\tilde{a})^{\top}\right)
$$

We have $X_{t} \in \mathcal{F}_{t}$ and $\mathbb{E}\left[X_{t} \mid \mathcal{F}_{t-1}\right]=\mathbb{E}\left[X_{t} \mid \pi_{t}\right]=0$, where the second identity follows from Fact 11 Moreover,

$$
\begin{align*}
& \left\|X_{t}\right\|_{F} \leq\left\|X_{t}\right\|_{1} \leq\left\|e_{o_{h+1}^{(t)}} \otimes e_{o_{h}^{(t)}} \otimes e_{o_{h-1}^{(t)}}\right\|_{1}+ \\
& \left\|\left(\mathbb{O}_{h+1} \mathbb{T}_{h}(a)\right) \otimes \mathbb{O}_{h} \otimes\left(\mathbb{O}_{h-1} \operatorname{diag}\left(\mathbb{P}_{\theta}^{\pi_{t}}\left(s_{h-1}=\cdot\right)\right) \mathbb{T}_{h-1}(\tilde{a})^{\top}\right)\right\|_{1}=2, \tag{15}
\end{align*}
$$

where $\|\cdot\|_{1}$ denotes the entry-wise $\ell_{1}$-norm of the tensor.
Now, we can bound $\left\|\hat{P}_{h+1, h, h-1}-P_{h+1, h, h-1}\right\|_{F}$ by writing $\hat{P}_{h+1, h, h-1}-P_{h+1, h, h-1}$ as the sum of a sequence of tensor-valued martingale difference, vectorizing the tensors, and applying the standard
vector-valued martingale concentration inequality (e.g. see Corollary 7 in [17]):

$$
\begin{aligned}
& \left\|\hat{P}_{h+1, h, h-1}-P_{h+1, h, h-1}\right\|_{F} \\
= & \| \frac{1}{N} \sum_{t=1}^{N}\left(e_{o_{h+1}^{(t)}} \otimes e_{o_{h}^{(t)}} \otimes e_{o_{h-1}^{(t)}}-\right. \\
& \left.\left(\mathbb{O}_{h+1} \mathbb{T}_{h}(a)\right) \otimes \mathbb{O}_{h} \otimes\left(\mathbb{O}_{h-1} \operatorname{diag}\left(\mathbb{P}_{\theta}^{\pi_{t}}\left(s_{h-1}=\cdot\right)\right) \mathbb{T}_{h-1}(\tilde{a})^{\top}\right)\right) \|_{F} \\
= & \left\|\frac{1}{N} \sum_{t=1}^{N} X_{t}\right\|_{F} \leq \mathcal{O}\left(\sqrt{\frac{\log (O N / \delta)}{N}}\right),
\end{aligned}
$$

with probability at least $1-\delta / 2$. We remark that when vectoring a tensor, its Frobenius norm will become the $\ell_{2}$-norm the vector. So the upper bound of the norm of the vectorized martingales directly follows from (15).

Similarly, we can show that with probability at least $1-\delta / 2$,
$\left\|\hat{P}_{h, h-1}-P_{h, h-1}\right\|_{F} \leq \mathcal{O}\left(\sqrt{\frac{\log (O N / \delta)}{N}}\right) \quad$ and $\quad\left\|\hat{P}_{h-1}-P_{h-1}\right\|_{F} \leq \mathcal{O}\left(\sqrt{\frac{\log (O N / \delta)}{N}}\right)$.
Using the fact $\left\|\hat{P}_{h+1, o, h-1}-P_{h+1, o, h-1}\right\|_{F} \leq\left\|\hat{P}_{h+1, h, h-1}-P_{h+1, h, h-1}\right\|_{F}$ completes the whole proof.

## C. 3 Properties of confidence sets

For convenience of discussion, we divide the constraints in $\Theta_{k}$ into three categories as following

## Type-0 constraint:

$$
\left.\left\|k \cdot \mathbf{b}_{0}(\hat{\theta})-\mathbf{n}^{k}\right\|_{2} \leq \beta_{k}\right\}
$$

Type-I constraint:

$$
\left\|\mathbf{B}_{1}(a, o ; \hat{\theta}) \mathbf{N}_{1}^{k}(a, \tilde{a})-\mathbf{M}_{1}^{k}(o, a, \tilde{a})\right\|_{F} \leq \gamma_{k}
$$

where $\mathbf{M}_{1}^{k}$ and $\mathbf{N}_{1}^{k}$ are actually equivalent to $O$-dimensional counting vectors because there is no observation (or only a dummy observation) at step 0 , which implies each of them has only one non-zero column. With slight abuse of notation, we use $\mathbf{M}_{1}^{k}$ and $\mathbf{N}_{1}^{k}$ to denote their non-zero columns in the following proof.
Type-II constraint: for $2 \leq h \leq H-1$,

$$
\left\|\mathbf{B}_{h}(a, o ; \hat{\theta}) \mathbf{N}_{h}^{k}(a, \tilde{a})-\mathbf{M}_{h}^{k}(o, a, \tilde{a})\right\|_{F} \leq \gamma_{k}
$$

Recalling the definition of $\mathbf{n}^{k}(\theta), \mathbf{N}_{h}^{k}(a, \tilde{a})$ and $\mathbf{M}_{h}^{k}(o, a, \tilde{a})$ and applying Lemma 12 , we get the following concentration results.
Corollary 13. Let $\theta^{\star}=\left(\mathbb{T}, \mathbb{O}, \mu_{1}\right)$. By applying Lemma 12 directly, with probability at least $1-\delta$, for all $k \in[K]$ and $(o, a, \tilde{a}) \in \mathscr{O} \times \mathscr{A}^{2}$, we have

$$
\left\{\begin{array}{l}
\left\|\frac{1}{k} \mathbf{n}^{k}-\mathbb{O}_{1} \mu_{1}\right\|_{2} \leq \mathcal{O}\left(\sqrt{\frac{\iota}{k}}\right), \\
\left\|\frac{1}{k} \mathbf{N}_{1}^{k}(a, \tilde{a})-\mathbb{O}_{1} \mu_{1}\right\|_{2} \leq \mathcal{O}\left(\sqrt{\frac{\iota}{k}}\right), \\
\left\|\frac{1}{k} \mathbf{M}_{1}^{k}(o, a, \tilde{a})-\left(\mathbb{O}_{2} \mathbb{T}_{1}(\tilde{a}) \operatorname{diag}\left(\mu_{1}\right) \mathbb{O}_{1}^{\top}\right)_{o}\right\|_{2} \leq \mathcal{O}\left(\sqrt{\frac{\iota}{k}}\right), \\
\|\frac{1}{k} \mathbf{N}_{h}^{k}(a, \tilde{a})-\underbrace{\mathbb{O}_{h} \mathbb{T}_{h-1}(\tilde{a}) \operatorname{diag}\left(\mu_{h-1}^{k}\right) \mathbb{O}_{h-1}^{\top}}_{\mathbf{V}}\|_{F} \leq \mathcal{O}\left(\sqrt{\frac{\iota}{k}}\right), \\
\|\frac{1}{k} \mathbf{M}_{h}^{k}(o, a, \tilde{a})-\underbrace{\mathbb{O}_{h+1} \mathbb{T}_{h}(a) \operatorname{diag}\left(\mathbb{O}_{h}(o \mid \cdot)\right) \mathbb{T}_{h-1}(\tilde{a}) \operatorname{diag}\left(\mu_{h-1}^{k}\right) \mathbb{O}_{h-1}^{\top}}_{\mathbf{W}}\|_{F} \leq \mathcal{O}\left(\sqrt{\frac{\iota}{k}}\right),
\end{array}\right.
$$

where

$$
\iota=\log (K A O H / \delta) \quad \text { and } \quad \mu_{h-1}^{k}=\frac{1}{k} \sum_{t=1}^{k} \mathbb{P}_{\theta^{\star}}^{\pi_{t}}\left(s_{h-1}=\cdot\right) \quad 2 \leq h \leq H-1
$$

Note that for all $k \in[K], \mu_{1}^{k}=\mu_{1}$ independent of $\pi_{1}, \ldots, \pi_{k}$.
Now, with Corollary 13 , we can prove the true parameter $\theta^{\star}$ always lies in the confidence sets for $k \in[K]$ with high probability.
Lemma 14. Denote by $\theta^{\star}=\left(\mathbb{T}, \mathbb{O}, \mu_{1}\right)$ the the ground truth parameters of the POMDP. With probability at least $1-\delta$, we have $\theta^{\star} \in \Theta_{k}$ for all $k \in[K]$.

Proof. By the definition of $\mathbf{b}_{0}\left(\theta^{\star}\right)$ and $\mathbf{B}_{h}\left(a, o ; \theta^{\star}\right)$, we have

$$
(*)\left\{\begin{array}{l}
\mathbf{b}_{0}\left(\theta^{\star}\right)=\mathbb{O}_{1} \mu_{1}, \\
\left(\mathbb{O}_{2} \mathbb{T}_{1}(\tilde{a}) \operatorname{diag}\left(\mu_{1}\right) \mathbb{O}_{1}^{\top}\right)_{o}=\mathbf{B}_{1}\left(\tilde{a}, o ; \theta^{\star}\right) \mathbb{O}_{1} \mu_{1}, \\
\mathbf{W}=\mathbf{B}_{h}\left(a, o ; \theta^{\star}\right) \cdot \mathbf{V}, \quad h \geq 2,
\end{array}\right.
$$

where $\mathbf{W}$ and $\mathbf{V}$ are shorthands defined in Corollary 13
It's easy to see $(*)$ and Corollary 13 directly imply $\left\|\mathbf{n}^{k}-\mathbf{b}_{0}\left(\theta^{\star}\right)\right\|_{2} \leq \mathcal{O}(\sqrt{k \iota})$ and thus $\theta^{\star}$ satisfies Type-0 constraint. For other constraints, we have

## Type-I constraint:

$$
\begin{aligned}
& \left\|\mathbf{M}_{1}^{k}(o, a, \tilde{a})-\mathbf{B}_{1}\left(\tilde{a}, o ; \theta^{\star}\right) \mathbf{N}_{1}^{k}(a, \tilde{a})\right\|_{2} \\
\leq & \left\|\mathbf{M}_{1}^{k}(o, a, \tilde{a})-k\left(\mathbb{O}_{2} \mathbb{T}_{1}(\tilde{a}) \operatorname{diag}\left(\mu_{1}\right) \mathbb{O}_{1}^{\top}\right)_{o}\right\|_{2}+\left\|\mathbf{B}_{1}\left(\tilde{a}, o ; \theta^{\star}\right)\left(k \mathbb{O}_{1} \mu_{1}-\mathbf{N}_{1}^{k}(a, \tilde{a})\right)\right\|_{2} \\
& +k\left\|\left(\mathbb{O}_{2} \mathbb{T}_{1}(\tilde{a}) \operatorname{diag}\left(\mu_{1}\right) \mathbb{O}_{1}^{\top}\right)_{o}-\mathbf{B}_{1}\left(\tilde{a}, o ; \theta^{\star}\right) \mathbb{O}_{1} \mu_{1}\right\|_{2} \\
= & \left\|\mathbf{M}_{1}^{k}(o, a, \tilde{a})-k\left(\mathbb{O}_{2} \mathbb{T}_{1}(\tilde{a}) \operatorname{diag}\left(\mu_{1}\right) \mathbb{O}_{1}^{\top}\right)_{o}\right\|_{2}+\left\|\mathbf{B}_{1}\left(\tilde{a}, o ; \theta^{\star}\right)\left(k \mathbb{O}_{1} \mu_{1}-\mathbf{N}_{1}^{k}(a, \tilde{a})\right)\right\|_{2} \\
\leq & \left\|\mathbf{M}_{1}^{k}(o, a, \tilde{a})-k\left(\mathbb{O}_{2} \mathbb{T}_{1}(\tilde{a}) \operatorname{diag}\left(\mu_{1}\right) \mathbb{O}_{1}^{\top}\right)_{o}\right\|_{2}+\left\|\mathbf{B}_{1}\left(\tilde{a}, o ; \theta^{\star}\right)\right\|_{2}\left\|k \mathbb{O}_{1} \mu_{1}-\mathbf{N}_{1}^{k}(a, \tilde{a})\right\|_{2} \\
\leq & \mathcal{O}\left(\frac{\sqrt{k S \iota}}{\alpha}\right)
\end{aligned}
$$

where the identity follows from $(*)$, and the last inequality follows from Corollary 13 and

$$
\begin{aligned}
\left\|\mathbf{B}_{h}\left(a, o ; \theta^{\star}\right)\right\|_{2} & =\left\|\mathbb{O}_{h+1} \mathbb{T}_{h}(a) \operatorname{diag}\left(\mathbb{O}_{h}(o \mid \cdot)\right) \mathbb{O}_{h}^{\dagger}\right\|_{2} \\
& \leq \frac{1}{\alpha}\left\|\mathbb{O}_{h+1} \mathbb{T}_{h}(a) \operatorname{diag}\left(\mathbb{O}_{h}(o \mid \cdot)\right)\right\|_{2} \\
& \leq \frac{\sqrt{S}}{\alpha}\left\|\mathbb{O}_{h+1} \mathbb{T}_{h}(a) \operatorname{diag}\left(\mathbb{O}_{h}(o \mid \cdot)\right)\right\|_{1 \rightarrow 1} \leq \frac{\sqrt{S}}{\alpha} .
\end{aligned}
$$

Type-II constraint: similarly, for $h \geq 2$, we have

$$
\begin{aligned}
& \left\|\mathbf{B}_{h}\left(a, o ; \theta^{\star}\right) \mathbf{N}_{h}^{k}(a, \tilde{a})-\mathbf{M}_{h}^{k}(o, a, \tilde{a})\right\|_{F} \\
\leq & k\left\|\mathbf{B}_{h}\left(a, o ; \theta^{\star}\right) \cdot \mathbf{V}-\mathbf{W}\right\|_{F}+\left\|\mathbf{B}_{h}\left(a, o ; \theta^{\star}\right)\left(\mathbf{N}_{h}^{k}(a, \tilde{a})-k \mathbf{V}\right)\right\|_{F}+\left\|k \mathbf{W}-\mathbf{M}_{h}^{k}(o, a, \tilde{a})\right\|_{F} \\
= & \left\|\mathbf{B}_{h}\left(a, o ; \theta^{\star}\right)\left(\mathbf{N}_{h}^{k}(a, \tilde{a})-k \mathbf{V}\right)\right\|_{F}+\left\|k \mathbf{W}-\mathbf{M}_{h}^{k}(o, a, \tilde{a})\right\|_{F} \\
\leq & \left\|\mathbf{B}_{h}\left(a, o ; \theta^{\star}\right)\right\|_{2}\left\|\mathbf{N}_{h}^{k}(a, \tilde{a})-k \mathbf{V}\right\|_{F}+\left\|k \mathbf{W}-\mathbf{M}_{h}^{k}(o, a, \tilde{a})\right\|_{F} \\
\leq & \mathcal{O}\left(\frac{\sqrt{k S \iota}}{\alpha}\right)
\end{aligned}
$$

Therefore, we conclude that $\theta^{\star} \in \Theta_{k}$ for all $k \in[K]$ with probability at least $1-\delta$.
Furthermore, with Corollary 13, we can prove the following bound for operator error.

Lemma 15. With probability at least $1-\delta$, for all $k \in[K], \hat{\theta}=\left(\hat{\mathbb{O}}, \hat{\mathbb{T}}, \hat{\mu}_{1}\right) \in \Theta_{k+1}$ and $(o, a, \tilde{a}, h) \in$ $\mathscr{O} \times \mathscr{A}^{2} \times\{2, \ldots, H-1\}$, we have

$$
\left\{\begin{array}{l}
\left\|\mathbf{b}_{0}\left(\theta^{\star}\right)-\mathbf{b}_{0}(\hat{\theta})\right\|_{2} \leq \mathcal{O}\left(\sqrt{\frac{\iota}{k}}\right) \\
\left\|\left(\mathbf{B}_{1}(\tilde{a}, o ; \hat{\theta})-\mathbf{B}_{1}\left(\tilde{a}, o ; \theta^{\star}\right)\right) \mathbf{b}_{0}\left(\theta^{\star}\right)\right\|_{2} \leq \mathcal{O}\left(\sqrt{\frac{S \iota}{k \alpha^{2}}}\right) \\
\sum_{s=1}^{S}\left\|\left(\mathbf{B}_{h}(a, o ; \hat{\theta})-\mathbf{B}_{h}\left(a, o ; \theta^{\star}\right)\right)\left(\mathbb{O}_{h} \mathbb{T}_{h-1}(\tilde{a})\right)_{s}\right\|_{1} \sum_{t=1}^{k} \mathbb{P}_{\theta^{\star}}^{\pi_{t}}\left(s_{h-1}=s\right) \leq \mathcal{O}\left(\sqrt{\frac{k S^{2} O \iota}{\alpha^{4}}}\right)
\end{array}\right.
$$

where $\iota=\log (K A O H / \delta)$.

Proof. For readability, we copy the following set of identities from Lemma 14 here,

$$
(*)\left\{\begin{array}{l}
\mathbf{b}_{0}\left(\theta^{\star}\right)=\mathbb{O}_{1} \mu_{1}, \\
\left(\mathbb{O}_{2} \mathbb{T}_{1}(\tilde{a}) \operatorname{diag}\left(\mu_{1}\right) \mathbb{O}_{1}^{\top}\right)_{o}=\mathbf{B}_{1}\left(\tilde{a}, o ; \theta^{\star}\right) \mathbb{O}_{1} \mu_{1}, \\
\mathbf{W}=\mathbf{B}_{h}\left(a, o ; \theta^{\star}\right) \cdot \mathbf{V}, \quad h \geq 2 .
\end{array}\right.
$$

## Type-0 closeness:

$$
\left\|\mathbf{b}_{0}\left(\theta^{\star}\right)-\mathbf{b}_{0}(\hat{\theta})\right\|_{2} \leq\left\|\frac{1}{k} \mathbf{n}^{k}-\mathbf{b}_{0}\left(\theta^{\star}\right)\right\|_{2}+\left\|\frac{1}{k} \mathbf{n}^{k}-\mathbf{b}_{0}(\hat{\theta})\right\|_{2} \leq \mathcal{O}\left(\sqrt{\frac{\iota}{k}}\right)
$$

where the last inequality follows from (*), Corollary 13 and $\hat{\theta} \in \Theta_{k+1}$.
Type-I closeness: similarly, we have

$$
\begin{aligned}
& \left\|\left(\mathbf{B}_{1}(\tilde{a}, o ; \hat{\theta})-\mathbf{B}_{1}\left(\tilde{a}, o ; \theta^{\star}\right)\right) \mathbf{b}_{0}\left(\theta^{\star}\right)\right\|_{2} \\
\leq & \left\|\left(\mathbb{O}_{2} \mathbb{T}_{1}(\tilde{a}) \operatorname{diag}\left(\mu_{1}\right) \mathbb{O}_{1}^{\top}\right)_{o}-\mathbf{B}_{1}\left(\tilde{a}, o ; \theta^{\star}\right) \mathbf{b}_{0}\left(\theta^{\star}\right)\right\|_{2} \\
& +\left\|\left(\mathbb{O}_{2} \mathbb{T}_{1}(\tilde{a}) \operatorname{diag}\left(\mu_{1}\right) \mathbb{O}_{1}^{\top}\right)_{o}-\mathbf{B}_{1}(\tilde{a}, o ; \hat{\theta}) \mathbf{b}_{0}\left(\theta^{\star}\right)\right\|_{2} \\
= & \left\|\left(\mathbb{O}_{2} \mathbb{T}_{1}(\tilde{a}) \operatorname{diag}\left(\mu_{1}\right) \mathbb{O}_{1}^{\top}\right)_{o}-\mathbf{B}_{1}(\tilde{a}, o ; \hat{\theta}) \mathbf{b}_{0}\left(\theta^{\star}\right)\right\|_{2} \\
\leq & \left\|\left(\mathbb{O}_{2} \mathbb{T}_{1}(\tilde{a}) \operatorname{diag}\left(\mu_{1}\right) \mathbb{O}_{1}^{\top}\right)_{o}-\frac{1}{k} \mathbf{M}_{1}^{k}(o, a, \tilde{a})\right\|_{2}+\frac{1}{k}\left\|\mathbf{M}_{1}^{k}(o, a, \tilde{a})-\mathbf{B}_{1}(\tilde{a}, o ; \hat{\theta}) \mathbf{N}_{1}^{k}(a, \tilde{a})\right\|_{2} \\
& +\left\|\mathbf{B}_{1}(\tilde{a}, o ; \hat{\theta})\left(\frac{1}{k} \mathbf{N}_{1}^{k}(a, \tilde{a})-\mathbf{b}_{0}\left(\theta^{\star}\right)\right)\right\|_{2} \\
\leq & \left\|\left(\mathbb{O}_{2} \mathbb{T}_{1}(\tilde{a}) \operatorname{diag}\left(\mu_{1}\right) \mathbb{O}_{1}^{\top}\right)_{o}-\frac{1}{k} \mathbf{M}_{1}^{k}(o, a, \tilde{a})\right\|_{2}+\frac{1}{k}\left\|\mathbf{M}_{1}^{k}(o, a, \tilde{a})-\mathbf{B}_{1}(\tilde{a}, o ; \hat{\theta}) \mathbf{N}_{1}^{k}(a, \tilde{a})\right\|_{2} \\
& +\left\|\mathbf{B}_{1}(\tilde{a}, o ; \hat{\theta})\right\|_{2}\left\|\frac{1}{k} \mathbf{N}_{1}^{k}(a, \tilde{a})-\mathbb{O}_{1} \mu_{1}\right\|_{2} \\
\leq & \mathcal{O}\left(\sqrt{\frac{S \iota}{k \alpha^{2}}}\right)
\end{aligned}
$$

where the identity follows from $(*)$ and the last inequality follows from Corollary 13 and $\hat{\theta} \in \Theta_{k+1}$.

Type-II closeness: we continue to use the same proof strategy, for $h \geq 2$

$$
\begin{align*}
& \left\|\left(\mathbf{B}_{h}(a, o ; \hat{\theta})-\mathbf{B}_{h}\left(a, o ; \theta^{\star}\right)\right) \mathbf{V}\right\|_{F} \\
\leq & \left\|\mathbf{W}-\mathbf{B}_{h}\left(a, o ; \theta^{\star}\right) \mathbf{V}\right\|_{F}+\left\|\frac{1}{k} \mathbf{M}_{h}^{k}(o, a, \tilde{a})-\mathbf{W}\right\|_{F} \\
& +\frac{1}{k}\left\|\mathbf{B}_{h}(a, o ; \hat{\theta}) \mathbf{N}_{h}^{k}(a, \tilde{a})-\mathbf{M}_{h}^{k}(o, a, \tilde{a})\right\|_{F}+\left\|\mathbf{B}_{h}(a, o ; \hat{\theta})\left(\mathbf{V}-\frac{1}{k} \mathbf{N}_{h}^{k}(a, \tilde{a})\right)\right\|_{F} \\
= & \left\|\frac{1}{k} \mathbf{M}_{h}^{k}(o, a, \tilde{a})-\mathbf{W}\right\|_{F}+\frac{1}{k}\left\|\mathbf{B}_{h}(a, o ; \hat{\theta}) \mathbf{N}_{h}^{k}(a, \tilde{a})-\mathbf{M}_{h}^{k}(o, a, \tilde{a})\right\|_{F} \\
& +\left\|\mathbf{B}_{h}(a, o ; \hat{\theta})\left(\mathbf{V}-\frac{1}{k} \mathbf{N}_{h}^{k}(a, \tilde{a})\right)\right\|_{F} \\
\leq & \mathcal{O}\left(\sqrt{\frac{S \iota}{k \alpha^{2}}}\right) \tag{16}
\end{align*}
$$

where the identity follows from $(*)$ and the last inequality follows from Corollary 13 and $\hat{\theta} \in \Theta_{k+1}$. Recall $\mathbf{V}=\mathbb{O}_{h} \mathbb{T}_{h-1}(\tilde{a}) \operatorname{diag}\left(\mu_{h-1}^{k}\right) \mathbb{O}_{h-1}^{\top}$ and utilize Assumption 1 .

$$
\begin{aligned}
& \left\|\left(\mathbf{B}_{h}(a, o ; \hat{\theta})-\mathbf{B}_{h}\left(a, o ; \theta^{\star}\right)\right) \mathbf{V}\right\|_{F} \\
\geq & \alpha\left\|\left(\mathbf{B}_{h}(a, o ; \hat{\theta})-\mathbf{B}_{h}\left(a, o ; \theta^{\star}\right)\right) \mathbb{O}_{h} \mathbb{T}_{h-1}(\tilde{a}) \operatorname{diag}\left(\mu_{h-1}^{k}\right)\right\|_{F} \\
\geq & \frac{\alpha}{\sqrt{S O}}\left\|\left(\mathbf{B}_{h}(a, o ; \hat{\theta})-\mathbf{B}_{h}\left(a, o ; \theta^{\star}\right)\right) \mathbb{O}_{h} \mathbb{T}_{h-1}(\tilde{a}) \operatorname{diag}\left(\mu_{h-1}^{k}\right)\right\|_{1} \\
= & \frac{\alpha}{k \sqrt{S O}} \sum_{s=1}^{S}\left\|\left(\mathbf{B}_{h}(a, o ; \hat{\theta})-\mathbf{B}_{h}\left(a, o ; \theta^{\star}\right)\right)\left(\mathbb{O}_{h} \mathbb{T}_{h-1}(\tilde{a})\right)_{s}\right\|_{1} \sum_{t=1}^{k} \mathbb{P}_{\theta^{\star}}^{\pi_{t}}\left(s_{h-1}=s\right) .
\end{aligned}
$$

Plugging it back into (16) completes the whole proof.

## C. 4 Proof of Theorem 3

In order to utilize Lemma 15 to bound the operator error in Lemma 10, we need the following algebraic transformation. Its proof is postponed to Appendix E
Lemma 16. Let $z_{k} \in\left[0, C_{z}\right]$ and $w_{k} \in\left[0, C_{w}\right]$ for $k \in \mathbb{N}$. Define $S_{k}=\sum_{j=1}^{k} w_{j}$ and $S_{0}=0$. If $z_{k} S_{k-1} \leq C_{z} C_{w} C_{0} \sqrt{k}$ for any $k \in[K]$, we have

$$
\sum_{k=1}^{K} z_{k} w_{k} \leq 2 C_{z} C_{w}\left(C_{0}+1\right) \sqrt{K} \log (K)
$$

Moreover, there exists some hard case where we have a almost matching lower bound $O\left(C_{z} C_{w} C_{0} \sqrt{K}\right)$.

Now, we are ready to prove the main theorem based on Lemma 10. Lemma 15 and Lemma 16
Theorem 3. For any $\varepsilon \in(0, H]$, there exists $K_{\max }=\operatorname{poly}\left(H, S, A, O, \alpha^{-1}\right) / \varepsilon^{2}$ and an absolute constant $c_{1}$, such that for any POMDP that satisfies Assumption 17 if we set hyperparameters $\beta_{k}=c_{1} \sqrt{k \log (K A O H)}, \gamma_{k}=\sqrt{S} \beta_{k} / \alpha$, and $K \geq K_{\max }$, then the output policy $\hat{\pi}$ of Algorithm 1 will be $\varepsilon$-optimal with probability at least $2 / 3$.

Proof. There always exist an optimal deterministic policy $\pi^{\star}$ for the ground truth $\operatorname{POMDP}\left(\theta^{\star}\right)$, i.e., $V^{\star}\left(\theta^{\star}\right)=V^{\pi^{\star}}\left(\theta^{\star}\right)$. WLOG, we can always choose the greedy policy $\pi_{k}$ to be deterministic, i.e., the probability to take any action given a history is either 0 or 1 .

By Lemma 14, we have $\theta^{\star} \in \Theta_{k}$ for all $k \in[K]$ with probability at least $1-\delta$. Recall that $\left(\pi_{k}, \theta_{k}\right)=\arg \max _{\pi, \theta \in \Theta_{k}} V^{\pi}(\theta)$, so with probability at least $1-\delta$, we have

$$
\begin{align*}
& \sum_{k=1}^{K}\left(V^{\pi^{\star}}\left(\theta^{\star}\right)-V^{\pi_{k}}\left(\theta^{\star}\right)\right) \\
\leq & \sum_{k=1}^{K}\left(V^{\pi_{k}}\left(\theta_{k}\right)-V^{\pi_{k}}\left(\theta^{\star}\right)\right) \\
\leq & H \sum_{k=1}^{K} \sum_{\left[o_{H}, \tau_{H-1}\right] \in \mathscr{O} \times \Gamma\left(\pi_{k}, H-1\right)}\left\|\mathbb{P}_{\theta^{\star}}^{\pi_{k}}\left(\left[o_{H}, \tau_{H-1}\right]\right)-\mathbb{P}_{\theta_{k}}^{\pi_{k}}\left(\left[o_{H}, \tau_{H-1}\right]\right)\right\|_{1} \\
= & H \sum_{k=1}^{K} \sum_{\tau_{H-1} \in \Gamma\left(\pi_{k}, H-1\right)}\left\|\mathbf{b}\left(\tau_{H-1} ; \theta^{\star}\right)-\mathbf{b}\left(\tau_{H-1} ; \theta_{k}\right)\right\|_{1}, \tag{17}
\end{align*}
$$

where the identity follows from Fact 18 .
Applying Lemma 10, we have

$$
\begin{align*}
& \sum_{\tau_{H-1} \in \Gamma\left(\pi_{k}, H-1\right)}\left\|\mathbf{b}\left(\tau_{H-1} ; \theta^{\star}\right)-\mathbf{b}\left(\tau_{H-1} ; \theta_{k}\right)\right\|_{1} \\
& \leq \underbrace{\frac{\sqrt{S}}{\alpha}\left\|\mathbf{b}_{0}\left(\theta^{\star}\right)-\mathbf{b}_{0}\left(\theta_{k}\right)\right\|_{1}+\frac{\sqrt{S}}{\alpha} \sum_{(a, o) \in \mathscr{A} \times \mathscr{O}}\left\|\left(\mathbf{B}_{1}\left(a, o ; \theta_{k}\right)-\mathbf{B}_{1}\left(a, o ; \theta^{\star}\right)\right) \mathbf{b}_{0}\left(\theta^{\star}\right)\right\|_{1}}_{J_{k}} \\
&+\frac{\sqrt{S}}{\alpha} \sum_{h=2}^{H-1} \sum_{(a, \tilde{a}, o) \in \mathscr{A}^{2} \times \mathscr{O}} \sum_{s=1}^{S}\left\|\left(\mathbf{B}_{h}\left(a, o ; \theta_{k}\right)-\mathbf{B}_{h}\left(a, o ; \theta^{\star}\right)\right)\left(\mathbb{O}_{h} \mathbb{T}_{h-1}(\tilde{a})\right)_{s}\right\|_{1} \mathbb{P}_{\theta^{\star}}^{\pi_{k}}\left(s_{h-1}=s\right) \tag{18}
\end{align*}
$$

We can bound the first two terms by Lemma 15 , and obtain that with probability at least $1-\delta$,

$$
\begin{equation*}
H \sum_{k=1}^{K} J_{k} \leq \mathcal{O}\left(\frac{H S A O}{\alpha^{2}} \sqrt{K \iota}\right) \tag{19}
\end{equation*}
$$

Plugging 18 and 19 into (17), we obtain

$$
\begin{align*}
& \sum_{k=1}^{K}\left(V^{\pi^{\star}}\left(\theta^{\star}\right)-V^{\pi_{k}}\left(\theta^{\star}\right)\right) \leq \mathcal{O}\left(\frac{H S A O}{\alpha^{2}} \sqrt{K \iota}\right)+ \\
& \frac{H^{2} S^{1.5} A^{2} O}{\alpha} \max _{s, o, a, \tilde{a}, h} \sum_{k=1}^{K}\left\|\left(\mathbf{B}_{h}\left(a, o ; \theta_{k}\right)-\mathbf{B}_{h}\left(a, o ; \theta^{\star}\right)\right)\left(\mathbb{O}_{h} \mathbb{T}_{h-1}(\tilde{a})\right)_{s}\right\|_{1} \mathbb{P}_{\theta^{\star}}^{\pi_{k}}\left(s_{h-1}=s\right) \tag{20}
\end{align*}
$$

It remains to bound the second term.
By Lemma 15, with probability at least $1-\delta$, for all $k \in[K], \theta_{k} \in \Theta_{k}$ and $(s, o, a, \tilde{a}, h) \in$ $\mathscr{S} \times \mathscr{O} \times \mathscr{A}^{2} \times\{2, \ldots, H-1\}$, we have

$$
\begin{equation*}
\underbrace{\left\|\left(\mathbf{B}_{h}\left(a, o ; \theta_{k}\right)-\mathbf{B}_{h}\left(a, o ; \theta^{\star}\right)\right)\left(\mathbb{O}_{h} \mathbb{T}_{h-1}(\tilde{a})\right)_{s}\right\|_{1}}_{z_{k}} \sum_{t=1}^{k-1} \underbrace{\mathbb{P}_{\theta^{\star}}^{\pi_{t}}\left(s_{h-1}=s\right)}_{w_{t}} \leq \mathcal{O}\left(\sqrt{\frac{k S^{2} O \iota}{\alpha^{4}}}\right) \tag{21}
\end{equation*}
$$

By simple calculation, we have $z_{k} \leq \sqrt{S} / \alpha$. Invoking Lemma 16 with 21, we obtain

$$
\begin{equation*}
\sum_{k=1}^{K} w_{k} z_{k} \leq \mathcal{O}\left(\frac{\sqrt{S^{3} O \iota}}{\alpha^{3}} \sqrt{K} \log (K)\right) \tag{22}
\end{equation*}
$$

Plugging 22, back into 20) gives

$$
\begin{equation*}
\sum_{k=1}^{K}\left(V^{\pi^{\star}}\left(\theta^{\star}\right)-V^{\pi_{k}}\left(\theta^{\star}\right)\right) \leq \mathcal{O}\left(\frac{H^{2} S^{3} A^{2} O^{1.5} \sqrt{\iota}}{\alpha^{4}} \sqrt{K} \log (K)\right) \tag{23}
\end{equation*}
$$

Finally, choosing

$$
K_{\max }=\mathcal{O}\left(\frac{H^{4} S^{6} A^{4} O^{3} \log (H S A O / \varepsilon)}{\alpha^{8} \varepsilon^{2}}\right)
$$

and outputting a policy from $\left\{\pi_{1}, \ldots, \pi_{K}\right\}$ uniformly at random complete the proof.

## D Learning POMDPs with Deterministic Transition

In this section, we introduce a computationally and statistically efficient algorithm for POMDPs with deterministic transition. A sketched proof is provided.

We comment that some previous works have studied POMDPs with deterministic transitions or deterministic emission process assuming the model is known (e.g. [4, 5, 6]); their results mainly focus on the planning aspect. In contrast, we assume unknown models which requires to learn the transition and emission process first. It is also worth mentioning that the (quasi)-deterministic POMDPs defined in these works are not exactly the same as ours. For example, the deterministic POMDPs in [6] refer to those with stochastic initial state but deterministic emission process, while we assume deterministic initial state but stochastic emission process. Therefore, their computational hardness results do not conflict with the efficient algorithm in this section.

```
Algorithm 2 Learning POMDPs with Deterministic Transition
    initialize \(N=C \log (H S A / p) /(\min \{\epsilon /(\sqrt{O} H), \xi\})^{2}, n_{h}=\mathbf{1}(h=1)\) for all \(h \in[H]\).
    for \(h=1, \ldots, H-1\) do
        for \((s, a) \in\left[n_{h}\right] \times \mathscr{A}\) do
            Reset \(z \leftarrow \mathbf{0}_{O \times 1}\) and \(t \leftarrow n_{h+1}+1\)
            for \(i \in[N]\) do
                execute policy \(\pi_{h}(s)\) from step 1 to step \(h-1\), take action \(a\) at \(h^{\text {th }}\) step and observe \(o_{h+1}\)
                \(z \leftarrow z+\frac{1}{N} e_{o_{h+1}}\)
            for \(s^{\prime} \in\left[n_{h+1}\right]\) do
                if \(\left\|\phi_{h+1, s^{\prime}}-z\right\|_{2} \leq 0.5 \xi\) then
                \(t \leftarrow s^{\prime}\)
            if \(t=n_{h+1}+1\) then
                \(n_{h+1} \leftarrow n_{h+1}+1\)
                \(\phi_{h+1, n_{h+1}} \leftarrow z\) and \(\pi_{h+1}\left(n_{h+1}\right) \leftarrow a \circ \pi_{h}(s)\)
            Set the \(s^{\text {th }}\) column of \(\hat{\mathbb{T}}_{h, a}\) to be \(e_{t}\)
    output \(\hat{\mu}_{0}=e_{1}\) and \(\left\{n_{h},\left\{\hat{\mathbb{T}}_{h, a}\right\}_{a \in \mathscr{A}}\right.\) and \(\left.\left\{\phi_{h, i}\right\}_{i \in\left[n_{h}\right]}: h \in[H]\right\}\)
```

Theorem 4. For any $p \in(0,1]$, there exists an algorithm such that for any deterministic transition POMDP satisfying Assumption 2, within $\mathcal{O}\left(H^{2} S A \log (H S A / p) /(\min \{\varepsilon /(\sqrt{O} H), \xi\})^{2}\right)$ samples and computations, the output policy of the algorithm is $\varepsilon$-optimal with probability at least $1-p$.

Proof. The algorithm works by inductively finding all the states we can reach at each step, utilizing the property of deterministic transition and good separation between different observation vectors. We sketch a proof based on induction below.
We say a state $s$ is $h$-step reachable if there exists a policy $\pi$ s.t. $\mathbb{P}^{\pi}\left(s_{h}=s\right)=1$. In our algorithm, we use $n_{h}$ to denote the number of $h$-step reachable states. All the policies mentioned below is a sequence of fixed actions (independent of observations).
Suppose at step $h$, there are $n_{h} h$-step reachable states and we can reach the $s^{\text {th }}$ one of them at the $h^{\text {th }}$ step by executing a known policy $\pi_{h}(s)$. Note that for every state $s^{\prime}$ that is $(h+1)$-step reachable, there must exist some state $s$ and action $a$ s.t. $s$ is $h$-step reachable and $\mathbb{T}_{h}\left(s^{\prime} \mid s, a\right)=1$. Therefore, based on our induction assumption, we can reach all the $(h+1)$-step reachable states by executing all $a \circ \pi_{h}(s)$ for $(a, s) \in \mathscr{A} \times\left[n_{h}\right]$.
Now the problem is how to tell if we reach the same state by executing two different $a \circ \pi_{h}(s)$ 's. The solution is to look at the distribution of $o_{h+1}$. Because the POMDP has deterministic transition, we always reach the same state when executing the same $a \circ \pi_{h}(s)$ and hence the distribution of $o_{h+1}$ is exactly the distribution of observation corresponding to that state. By Hoeffding's inequality, for each fixed $a \circ \pi_{h}(s)$, we can estimate the distribution of $o_{h+1}$ with $\ell_{2}$-error smaller than $\xi / 8$ with high probability using $N \geq \tilde{\Omega}\left(1 / \xi^{2}\right)$ samples. Since the observation distributions of two different states have $\ell_{2}$-separation no smaller than $\xi$, we can tell if two different $a \circ \pi_{h}(s)$ 's reach the same
state by looking at the distance between their distributions of $o_{h+1}$. For those policies reaching the same state, we only need to keep one of them, so there are at most $S$ policies kept ( $n_{h+1} \leq S$ ).
By repeating the argument inductively from $h=1$ to $h=H$, we can recover the exact transition dynamics $\mathbb{T}_{h}(\cdot \mid s, a)$ and get an high-accuary estimate of $\mathbb{O}_{h}(\cdot \mid s)$ for every $h$-step reachable state $s$ and all $(h, a) \in[H] \times \mathscr{A}$ up to permutation of states. Since the POMDP has deterministic transition, we can easily find the optimal policy of the estimated model by dynamic programming.

The $\epsilon$-optimality simply follows from the fact that when $N \geq \tilde{\Omega}\left(H^{2} O / \epsilon^{2}\right)$, we have the estimated distribution of observation for each state being $\mathcal{O}(\epsilon / H)$ accurate in $\ell_{1}$-distance for all reachable states. This implies that the optimal policy of the estimated model is at most $\mathcal{O}(\epsilon / H) \times H=\mathcal{O}(\epsilon)$ suboptimal. The overall sample complexity follows from our requirement $N \geq \max \left\{\tilde{\Omega}\left(H^{2} O / \epsilon^{2}\right), \tilde{\Omega}\left(1 / \xi^{2}\right)\right\}$, and the fact we need to run $N$ episodes for each $h \in[H], s \in \mathscr{S}, a \in \mathscr{A}$.

## E Auxiliary Results

## E. 1 Derivation of equation (2)

When conditioning on a fixed action sequence $\left\{a_{H-1}, \ldots, a_{1}\right\}$, a POMDP will reduce to a non-stationary HMM, whose transition matrix and observation matrix at $h^{\text {th }}$ step are $\mathbb{T}_{h}\left(a_{h}\right)$ and $\mathbb{O}_{h}$, respectively. So $\mathbb{P}\left(o_{H}, \ldots, o_{1} \mid a_{H-1}, \ldots, a_{1}\right)$ is equal to the probability of observing $\left\{o_{H}, \ldots, o_{1}\right\}$ in this particular HMM. Using the basic properties of HMMs, we can easily express $\mathbb{P}\left(o_{H}, \ldots, o_{1} \mid a_{H-1}, \ldots, a_{1}\right)$ in terms of the transition and observation matrices

$$
\mathbb{O}_{H}\left(o_{H} \mid \cdot\right) \cdot\left[\mathbb{T}_{H-1}\left(a_{H-1}\right) \operatorname{diag}\left(\mathbb{O}_{H-1}\left(o_{H-1} \mid \cdot\right)\right)\right] \cdots\left[\mathbb{T}_{1}\left(a_{1}\right) \operatorname{diag}\left(\mathbb{O}_{1}\left(o_{1} \mid \cdot\right)\right)\right] \cdot \mu_{1}
$$

Recall the definition of operators

$$
\mathbf{B}_{h}(a, o)=\mathbb{O}_{h+1} \mathbb{T}_{h}(a) \operatorname{diag}\left(\mathbb{O}_{h}(o \mid \cdot)\right) \mathbb{O}_{h}^{\dagger}, \quad \mathbf{b}_{0}=\mathbb{O}_{1} \mu_{1}
$$

and $\mathbb{O}_{h}^{\dagger} \mathbb{O}_{h}=\mathbf{I}_{S}$, we conclude that

$$
\mathbb{P}\left(o_{H}, \ldots, o_{1} \mid a_{H-1}, \ldots, a_{1}\right)=\mathbf{e}_{o_{H}}^{\top} \cdot \mathbf{B}_{H-1}\left(a_{H-1}, o_{H-1}\right) \cdots \mathbf{B}_{1}\left(a_{1}, o_{1}\right) \cdot \mathbf{b}_{0}
$$

## E. 2 Derivation of equation (6)

Note that $\pi$ is a deterministic policy and $\Gamma(\pi, H)$ is a set of all the observation and action sequences of length $H$ that could occur under policy $\pi$, i.e., for any $\tau_{H}=\left(o_{H}, \ldots, a_{1}, o_{1}\right) \in \Gamma(\pi, H)$, we have $\pi\left(a_{H-1} \ldots, a_{1} \mid o_{H}, \ldots, o_{1}\right)=1$, and $\pi\left(a_{H-1}^{\prime} \ldots, a_{1}^{\prime} \mid o_{H}, \ldots, o_{1}\right)=0$ for any action sequence $\left(a_{H-1}^{\prime} \ldots, a_{1}^{\prime}\right) \neq\left(a_{H-1} \ldots, a_{1}\right)$. Therefore, for $\tau_{H} \in \Gamma(\pi, H)$, we have:

$$
\begin{aligned}
\mathbb{P}_{\theta}^{\pi}\left(o_{H}, \ldots, o_{1}\right) & =\sum_{a_{H-1}^{\prime} \in \mathscr{A}} \ldots \sum_{a_{1}^{\prime} \in \mathscr{A}} \mathbb{P}_{\theta}^{\pi}\left(o_{H}, a_{H-1}^{\prime}, \ldots, a_{1}^{\prime}, o_{1}\right) \\
& =\mathbb{P}_{\theta}^{\pi}\left(o_{H}, a_{H-1}, \ldots, a_{1}, o_{1}\right) \\
& =\left[\prod_{h=1}^{H-1} \pi\left(a_{h} \mid o_{h}, a_{h-1}, \ldots, a_{1}, o_{1}\right)\right] \cdot\left[\prod_{h=1}^{H} \mathbb{P}_{\theta}\left(o_{h} \mid a_{h-1}, o_{h-1} \ldots, a_{1}, o_{1}\right)\right] \\
& =\prod_{h=1}^{H} \mathbb{P}_{\theta}\left(o_{h} \mid a_{h-1}, o_{h-1}, \ldots, a_{1}, o_{1}\right) \\
& =\mathbb{P}_{\theta}\left(o_{H}, \ldots, o_{1} \mid a_{H-1}, \ldots, a_{1}\right) .
\end{aligned}
$$

## E. 3 Boosting the success probability

We can run Algorithm 1 independently for $n=\mathcal{O}(\log (1 / \delta))$ times and obtain $n$ policies. Each policy is independent of others and is $\varepsilon$-optimal with probability at least $2 / 3$. So with probability at least $1-\delta / 2$, at least one of them will be $\varepsilon$-optimal. In order to evaluate their performance, it suffices to run each policy for $\mathcal{O}\left(\log (n / \delta) H^{2} / \varepsilon^{2}\right)$ episodes and use the empirical average of the cumulative reward as an estimate. By standard concentration argument, with probability at least $1-\delta / 2$, the estimation error for each policy is smaller than $\varepsilon$. Therefore, if we pick the policy with the best empirical performance, then with probability at least $1-\delta$, it is $3 \varepsilon$-optimal. Rescaling $\varepsilon$ gives the desired accuracy. It is direct to see that the boosting procedure will only incur an additional polylog $(1 / \delta)$ factor in the final sample complexity, and thus will not hurt the optimal dependence on $\varepsilon$.

## E. 4 Basic facts about POMDPs and the operators

In this section, we provide some useful facts about POMDPs. Since their proofs are quite straightforward, we omit them here.
The following fact gives two linear equations the operators always satisfy. Its proof simply follows from the definition of the operators and Fact 11

Fact 17. In the same setting as Fact 11 suppose Assumption 1 holds, then we have

$$
\left\{\begin{aligned}
\mathbb{P}\left(o_{h}=\cdot, o_{h-1}=\cdot\right) \mathbf{e}_{o} & =\mathbf{B}_{h}(\tilde{a}, o ; \theta) \mathbb{P}\left(o_{h-1}=\cdot\right), \\
\mathbb{P}\left(o_{h+1}=\cdot, o_{h}=o, o_{h-1}=\cdot\right) & =\mathbf{B}_{h}(a, o ; \theta) \mathbb{P}\left(o_{h}=\cdot, o_{h-1}=\cdot\right)
\end{aligned}\right.
$$

The following fact relates (unnormalized) belief states to distributions of observable sequences. Its proof follows from simple computation using conditional probability formula and $\mathbb{O}_{h}^{\dagger} \mathbb{O}_{h}=\mathbf{I}_{S}$.
Fact 18. For any $\operatorname{POMDP}(\theta)$ satisfying Assumption 1 deterministic policy $\pi$ and $\left[o_{h}, \tau_{h-1}\right] \in$ $\mathscr{O} \times \Gamma(\pi, h-1)$, we have

$$
\mathbf{e}_{o_{h}}^{\top} \mathbf{b}\left(\tau_{h-1} ; \theta\right)=\mathbb{P}_{\theta}^{\pi}\left(\left[o_{h}, \tau_{h-1}\right]\right),
$$

where $\mathbb{P}_{\theta}^{\pi}\left(\left[o_{h}, \tau_{h-1}\right]\right)$ is the probability of observing $\left[o_{h}, \tau_{h-1}\right]$ when executing policy $\pi$ in $\operatorname{POMDP}(\theta)$.

## E. 5 Proof of Lemma 16

Proof. WLOG, assume $C_{z}=C_{w}=1$. Let $n=\min \left\{k \in[K]: S_{k} \geq 1\right\}$. We have

$$
\begin{aligned}
\sum_{k=1}^{K} z_{k} w_{k}=\sum_{k=1}^{n} z_{k} w_{k}+\sum_{k=n+1}^{K} z_{k} w_{k} & \leq \sum_{k=1}^{n} w_{k}+\sum_{k=n+1}^{K} z_{k} w_{k} \\
& =S_{n}+\sum_{k=n+1}^{K} z_{k} w_{k} \\
& \leq 2+\sum_{k=n+1}^{K} z_{k} w_{k}
\end{aligned}
$$

It remains to bound the second term. Using the condition that $z_{k} S_{k-1} \leq C_{0} \sqrt{k}$ for all $k \in[K]$, we have $z_{k} \leq \frac{C_{0} \sqrt{k}}{S_{k-1}}$ for all $k \in[K]$ and $i \in[m]$. Therefore,

$$
\begin{aligned}
\sum_{k=n+1}^{K} z_{k} w_{k} & \leq \sum_{k=n+1}^{K} C_{0} \sqrt{k} \frac{w_{k}}{S_{k-1}} \\
& \leq C_{0} \sqrt{K} \sum_{k=n+1}^{K} \frac{w_{k}}{S_{k-1}} \\
& \stackrel{(a)}{\leq} 2 C_{0} \sqrt{K} \sum_{k=n+1}^{K} \log \left(\frac{S_{k}}{S_{k-1}}\right) \\
& =2 C_{0} \sqrt{K} \log \left(\frac{S_{K}}{S_{n}}\right) \leq 2 C_{0} \sqrt{K} \log (K)
\end{aligned}
$$

where $(a)$ follows from $x \leq 2 \log (x+1)$ for $x \in[0,1]$.


[^0]:    ${ }^{3}$ Note that this definition is different from the one used in Section 5 , where $\tau_{h}=\left[o_{h}, \ldots, a_{1}, o_{1}\right] \in$ $\mathscr{O} \times(\mathscr{A} \times \mathscr{O})^{h-1}$ does not include the action $a_{h}$ at $h^{\text {th }}$ step.
    ${ }^{4}$ WLOG, all the polices considered in this paper are deterministic. Also note that the trajectory in $\Gamma(\pi, h)$ contains $a_{h}$, which is different from the definition in Section 5

