## A Broader Impacts

This study delves into the theoretical aspects of offline imitation learning with supplementary data, and we verify our findings through experiments on established benchmarks. While this paper does not present any immediate, direct social impacts, the potential practical applications of our research could bring about positive change. By expanding the reach of imitation learning algorithms, our work may facilitate the development of more efficient and effective solutions in fields such as robotics, autonomous vehicles, and healthcare. However, we must also acknowledge that the misuse of such technology could have negative consequences, such as the manipulation of information to influence people's behavior. Therefore, it is crucial to remain vigilant in ensuring that the benefits of imitation learning are harnessed in a responsible and ethical manner.

## B Additional Related Work

Our work builds upon previous research in IL with supplementary data, specifically the algorithms DemoDICE [22] and DWBC [47]. These studies highlight the importance of careful data selection when using a supplementary dataset. In this vein, our method ISW-BC re-weights samples based on importance sampling, which we show to be theoretically sound. Notably, a significant distinction arises between ISW-BC and these two methods in terms of the weighting rule design. While DemoDICE and DWBC employ regularized weighting rules, our method directly estimates the importance sampling ratio. This fundamental difference can be critical as regularized weighting rules may struggle to recover the expert policy exactly even with infinite samples. We provide further elaboration on this point below.

First, DemoDICE also uses the weighted BC objective in Eq. (3). But, DemoDICE uses the weighting rule of $\widetilde{w}(s, a) \propto d^{\star}(s, a) / d^{\mathrm{U}}(s, a)$ (refer to the formula between Equations (19)-(20) in [22]), where $d^{\star}(s, a)$ is computed by the expert's state-action distribution matching objective regularized by a divergence to the union data distribution (refer to [22, Equations (5)-(7)] $)^{2}$ :

$$
\begin{array}{ll} 
& d^{\star}=\underset{d}{\operatorname{argmin}} D_{\mathrm{KL}}\left(d \| d^{\mathrm{E}}\right)+\alpha D_{\mathrm{KL}}\left(d \| d^{\mathrm{U}}\right) \\
\text { s.t. } & d(s, a) \geq 0 \quad \forall s, a . \\
& \sum_{a} d(s, a)=(1-\gamma) \rho(s)+\gamma \sum_{s^{\prime}, a^{\prime}} P\left(s \mid s^{\prime}, a^{\prime}\right) d\left(s^{\prime}, a^{\prime}\right) \quad \forall s .
\end{array}
$$

where $\gamma \in[0,1)$ is the discount factor, $\alpha>0$ is a hyper-parameter. Due to the regularization term in the objective, it holds that $d^{\star}(s, a) \neq d^{\pi^{\mathrm{E}}}(s, a)$, resulting in a biased weighting rule $\widetilde{w}(s, a)$.
Second, DWBC considers a different policy learning objective (refer to [47, Equation (17)]):

$$
\begin{align*}
\min _{\pi} & \alpha \sum_{(s, a) \in \mathcal{D}^{\mathrm{E}}}[-\log \pi(a \mid s)]-\sum_{(s, a) \in \mathcal{D}^{\mathrm{E}}}\left[-\log \pi(a \mid s) \cdot \frac{\lambda}{c(1-c)}\right] \\
& +\sum_{(s, a) \in \mathcal{D}^{\mathrm{S}}}\left[-\log \pi(a \mid s) \cdot \frac{1}{1-c}\right] \tag{7}
\end{align*}
$$

where $\alpha>0, \lambda>0$ are hyper-parameters, and $c$ is the output of the discriminator that is jointly trained with $\pi$ (refer to [47, Equation (8)]):

$$
\begin{aligned}
\min _{c} & \lambda \sum_{(s, a) \in \mathcal{D}^{\mathrm{E}}}[-\log c(s, a, \log \pi(a \mid s))]+\sum_{(s, a) \in \mathcal{D}^{\mathrm{S}}}[-\log (1-c(s, a, \log \pi(a \mid s)))] \\
& -\lambda \sum_{(s, a) \in \mathcal{D}^{\mathrm{E}}}[-\log (1-c(s, a, \log \pi(a \mid s)))] .
\end{aligned}
$$

Since its input additionally incorporates $\log \pi$, the discriminator is not guaranteed to estimate the state-action distribution. Thus, the weighting in Eq. (7) loses a connection with the importance sampling ratio.

[^0]In addition to our work, [7] have also explored the use of supplementary data in the offline setting. However, their approach (called MILO) is based on adversarial imitation learning. Specifically, MILO learns a transition model from the supplementary dataset and performs adversarial imitation learning within the learned model. In contrast, our proposed method, ISW-BC, tackles the challenge of scarce expert data by identifying and utilizing expert-style samples that are hidden within the supplementary dataset. MILO has an imitation gap bound of $\mathcal{O}\left(H \sqrt{\frac{|\mathcal{S}|}{N_{\mathrm{E}}}}+H^{2}|\mathcal{S}| \sqrt{\frac{|\mathcal{A}|}{N_{\mathrm{S}} / \mu}}\right)$ in theory. However, MILO makes different assumptions about the data collection procedure compared with ISW-BC. Consequently, the imitation gap bounds of MILO and ISW-BC are incomparable.

The problem considered in this paper is related to IL with a single imperfect dataset $[44,6,41$, $42,36,25]$. In particular, the supplementary dataset in our set-up can also be viewed as imperfect demonstrations. However, our problem setting differs from IL with imperfect demonstrations in two key aspects. First, in IL with imperfect demonstrations, they either pose strong assumptions [ $41,36,25$ ] or require auxiliary information (e.g., confidence scores on imperfect trajectories) on the imperfect dataset $[44,6]$. In contrast, we assume access to a small number of expert trajectories to identify in-expert-distribution data. Second, most works [44, 6, 41, 42] in IL with imperfect demonstrations require online environment interactions while we focus on the offline setting.

## C Proof of Results in Section 4

Recall the objective of BC in Eq. (1):

$$
\pi^{\mathrm{BC}} \in \max _{\pi} \sum_{h=1}^{H} \sum_{(s, a) \in \mathcal{S} \times \mathcal{A}} \widehat{d_{h}^{\mathrm{E}}}(s, a) \log \pi_{h}(a \mid s)
$$

where $\widehat{d_{h}^{\mathrm{E}}}(s, a)=n_{h}^{\mathrm{E}}(s, a) / N_{\text {tot }}$ is the empirical state-action distribution in the expert dataset, and $n_{h}^{\mathrm{E}}(s, a)$ is the number of expert trajectories such that their state-action pairs are equal to $(s, a)$ in time step $h$. With the tabular representations, we can obtain a closed-formed solution to the above optimization problem.

$$
\pi_{h}^{\mathrm{BC}}(a \mid s)= \begin{cases}\frac{n_{h}^{\mathrm{E}}(s, a)}{n_{h}^{\mathrm{E}}(s)} & \text { if } n_{h}^{\mathrm{E}}(s)>0  \tag{8}\\ \frac{1}{|\mathcal{A}|} & \text { otherwise }\end{cases}
$$

where $n_{h}^{\mathrm{E}}(s) \triangleq \sum_{a^{\prime}} n_{h}^{\mathrm{E}}\left(s, a^{\prime}\right)$. Analogously, we also have a closed-form solution for NBCU in the tabular setting:

$$
\pi_{h}^{\mathrm{NBCU}}(a \mid s)= \begin{cases}\frac{n_{h}^{\mathrm{U}}(s, a)}{n_{h}^{\mathrm{U}}(s)} & \text { if } n_{h}^{\mathrm{U}}(s)>0  \tag{9}\\ \frac{1}{|\mathcal{A}|} & \text { otherwise }\end{cases}
$$

We will discuss the generalization performance of NBCU later.
In the proof, we frequently use the notation $\lesssim$ and $\gtrsim$. In particular, $a(n) \lesssim b(n)$ means that there exist $C, n_{0}>0$ such that $a(n) \leq C b(n)$ for all $n \geq n_{0}$. In our context, $n$ usually refers to the number of trajectories. For any two distributions $P$ and $Q$ over a finite set $\mathcal{X}$, we define the total variation distance as

$$
\mathrm{TV}(P, Q)=\frac{1}{2} \sum_{x \in \mathcal{X}}|P(x)-Q(x)|=\|P-Q\|_{1}
$$

## C. 1 Proof of Theorem 1

When $\left|\mathcal{D}^{\mathrm{E}}\right| \geq 1$, by [32, Theorem 4.2], we have the following imitation gap bound for BC:

$$
V\left(\pi^{\mathrm{E}}\right)-\mathbb{E}_{\mathcal{D}^{\mathrm{E}}}\left[V\left(\pi^{\mathrm{BC}}\right)\right] \leq \frac{4|\mathcal{S}| H^{2}}{9\left|\mathcal{D}^{\mathrm{E}}\right|}
$$

When $\left|\mathcal{D}^{\mathrm{E}}\right|=0$, we simply have that

$$
V\left(\pi^{\mathrm{E}}\right)-\mathbb{E}_{\mathcal{D}^{\mathrm{E}}}\left[V\left(\pi^{\mathrm{BC}}\right)\right] \leq H
$$

For $V\left(\pi^{\mathrm{E}}\right)-V\left(\pi^{\text {mix }}\right)$, we have that

$$
\begin{align*}
V\left(\pi^{\mathrm{E}}\right)-V\left(\pi^{\mathrm{mix}}\right) & =\sum_{h=1}^{H} \sum_{(s, a) \in \mathcal{S} \times \mathcal{A}}\left(d_{h}^{\pi^{\mathrm{E}}}(s, a)-d_{h}^{\pi^{\mathrm{mix}}}(s, a)\right) r_{h}(s, a) \\
& =\sum_{h=1}^{H} \sum_{(s, a) \in \mathcal{S} \times \mathcal{A}}\left(d_{h}^{\pi^{\mathrm{E}}}(s, a)-d_{h}^{\mathrm{mix}}(s, a)\right) r_{h}(s, a) \\
& =(1-\eta) \sum_{h=1}^{H} \sum_{(s, a) \in \mathcal{S} \times \mathcal{A}}\left(d_{h}^{\pi^{\mathrm{E}}}(s, a)-d_{h}^{\pi^{\beta}}(s, a)\right) r_{h}(s, a) \\
& =(1-\eta)\left(V\left(\pi^{\mathrm{E}}\right)-V\left(\pi^{\beta}\right)\right) . \tag{11}
\end{align*}
$$

Therefore, we have the following unified bound.

$$
V\left(\pi^{\mathrm{E}}\right)-\mathbb{E}_{\mathcal{D}^{\mathrm{E}}}\left[V\left(\pi^{\mathrm{BC}}\right)\right] \leq \frac{|\mathcal{S}| H^{2}}{\max \left\{\left|\mathcal{D}^{\mathrm{E}}\right|, 1\right\}} \leq \frac{2|\mathcal{S}| H^{2}}{\left|\mathcal{D}^{\mathrm{E}}\right|+1}
$$

The last inequality follows that $\max \{x, 1\} \geq(x+1) / 2$ for any $x \geq 0$. Finally, notice that $\left|\mathcal{D}^{\mathrm{E}}\right|$ follows a binomial distribution by Assumption 1, i.e., $\left|\mathcal{D}^{\mathrm{E}}\right| \sim \operatorname{Bin}\left(N_{\text {tot }}, \eta\right)$. By Lemma 3, we have that $\mathbb{E}\left[1 /\left(|\mathcal{D}|^{E}+1\right)\right] \leq N_{\text {tot }} \eta$, so

$$
V\left(\pi^{\mathrm{E}}\right)-\mathbb{E}\left[V\left(\pi^{\mathrm{BC}}\right)\right] \leq \mathbb{E}\left[\frac{2|\mathcal{S}| H^{2}}{\left|\mathcal{D}^{\mathrm{E}}\right|+1}\right] \leq \frac{2|\mathcal{S}| H^{2}}{N_{\mathrm{tot}} \eta}=\frac{2|\mathcal{S}| H^{2}}{N_{\mathrm{E}}}
$$

which completes the proof.

## C. 2 Proof of Theorem 2

For analysis, we first define the mixture state-action distribution as follows.

$$
\begin{aligned}
d_{h}^{\operatorname{mix}}(s, a) & \triangleq \eta d_{h}^{\pi^{\mathrm{E}}}(s, a)+(1-\eta) d_{h}^{\pi^{\beta}}(s, a), \\
d_{h}^{\mathrm{mix}}(s) & \triangleq \sum_{a \in \mathcal{A}} d_{h}^{\mathrm{mix}}(s, a), \forall(s, a) \in \mathcal{S} \times \mathcal{A}, \forall h \in[H] .
\end{aligned}
$$

By Assumption 1, in the population level, the marginal state-action distribution of union dataset $\mathcal{D}^{\mathrm{U}}$ in time step $h$ is exactly $d_{h}^{\text {mix }}$. That is, $d_{h}^{U}(s, a)=d_{h}^{\text {mix }}(s, a), \forall(s, a, h) \in \mathcal{S} \times \mathcal{A} \times[H]$. Then we define the mixture policy $\pi^{\text {mix }}$ induced by $d^{\text {mix }}$ as follows.

$$
\pi_{h}^{\operatorname{mix}}(a \mid s)=\left\{\begin{array}{ll}
\frac{d_{h}^{\operatorname{mix}}(s, a)}{d_{h}^{\mathrm{mix}}(s)} & \text { if } d_{h}^{\operatorname{mix}}(s)>0,  \tag{10}\\
\frac{1}{|\mathcal{A}|} & \text { otherwise. }
\end{array} \forall(s, a) \in \mathcal{S} \times \mathcal{A}, \forall h \in[H] .\right.
$$

From the theory of Markov Decision Processes, we know that (see, e.g., [31])

$$
\forall h \in[H], \forall(s, a) \in \mathcal{S} \times \mathcal{A}, \quad d_{h}^{\pi^{\mathrm{mix}}}(s, a)=d_{h}^{\operatorname{mix}}(s, a) .
$$

Therefore, we can obtain that the marginal state-action distribution of union dataset $\mathcal{D}^{\mathrm{U}}$ in time step $h$ is exactly $d_{h}^{\pi^{\text {mix }}}$. Then we have the following decomposition.

$$
\begin{aligned}
\mathbb{E}\left[V\left(\pi^{\mathrm{E}}\right)-V\left(\pi^{\mathrm{NBCU}}\right)\right] & =\mathbb{E}\left[V\left(\pi^{\mathrm{E}}\right)-V\left(\pi^{\text {mix }}\right)+V\left(\pi^{\text {mix }}\right)-V\left(\pi^{\mathrm{NBCU}}\right)\right] \\
& =\mathbb{E}\left[V\left(\pi^{\mathrm{E}}\right)-V\left(\pi^{\text {mix }}\right)\right]+\mathbb{E}\left[V\left(\pi^{\text {mix }}\right)-V\left(\pi^{\mathrm{NBCU}}\right)\right] \\
& =V\left(\pi^{\mathrm{E}}\right)-V\left(\pi^{\mathrm{mix}}\right)+\mathbb{E}\left[V\left(\pi^{\mathrm{mix}}\right)-V\left(\pi^{\mathrm{NBCU}}\right)\right]
\end{aligned}
$$

The last equation follows the dual formulation of policy value (see, e.g., [31]), i.e., $V(\pi)=$ $\sum_{h=1}^{H} \sum_{(s, a)} d_{h}^{\pi}(s, a) r_{h}(s, a)$ for any policy $\pi$. Besides, notice that $\mathbb{E}\left[V\left(\pi^{\text {mix }}\right)-V\left(\pi^{\mathrm{NBCU}}\right)\right]$ is exactly the imitation gap of BC when regarding $\pi^{\text {mix }}$ and $\mathcal{D}^{\mathrm{U}}$ as the expert policy and expert dataset, respectively. Note that $\pi^{\text {mix }}$ may be a stochastic policy. By [32, Theorem 4.4], we have the following imtiation gap bound

$$
\begin{equation*}
\mathbb{E}\left[V\left(\pi^{\mathrm{mix}}\right)-V\left(\pi^{\mathrm{NBCU}}\right)\right] \lesssim \frac{|\mathcal{S}| H^{2} \log \left(N_{\mathrm{tot}}\right)}{N_{\mathrm{tot}}} \tag{12}
\end{equation*}
$$

## C. 3 Proof of Proposition 1

The hard instance in Proposition 1 builds on the Standard Imitation MDP proposed in [46]; see Figure 4 for illustration. For this MDP, each state is an absorbing state, i.e., $P_{h}(s \mid s, a)=1$ for any $s$ and $a$. This property is mainly used to facilitate probability calculation and does not change the nature of our analysis. Furthermore, by only taking the action $a^{1}$ (shown in green), the agent can obtain a reward of +1 . Otherwise, the agent obtains a reward of 0 for the other action $a \neq a^{1}$. The initial state distribution is a uniform distribution, i.e., $\rho(s)=1 /|\mathcal{S}|$ for any $s \in \mathcal{S}$.


Figure 4: The Standard Imitation MDP in [46] corresponding to prove Proposition 1.

We consider that the expert policy $\pi^{\mathrm{E}}$ always takes the action $a^{1}$ (shown in green) while the behavioral policy $\pi^{\beta}$ always takes another action $a^{2}$ (shown in blue). Formally, $\pi_{h}^{\mathrm{E}}\left(a^{1} \mid s\right)=1$ and $\pi_{h}^{\beta}\left(a^{2} \mid s\right)=1$ for any $s \in \mathcal{S}$ and $h \in[H]$. It is direct to calculate that $V\left(\pi^{\mathrm{E}}\right)=H$ and $V\left(\pi^{\beta}\right)=0$. The supplementary dataset $\mathcal{D}^{\mathrm{S}}$ and the expert dataset $\mathcal{D}^{\mathrm{E}}$ are collected according to Assumption 1. The mixture state-action distribution (introduced in Appendix C.2) can be calculated as for any $s \in \mathcal{S}$ and $h \in[H]:$

$$
\begin{aligned}
& d_{h}^{\operatorname{mix}}\left(s, a^{1}\right)=\eta d_{h}^{\pi^{\mathrm{E}}}\left(s, a^{1}\right)+(1-\eta) d_{h}^{\pi^{\beta}}\left(s, a^{1}\right)=\eta d_{h}^{\pi^{\mathrm{E}}}\left(s, a^{1}\right)=\eta \rho(s), \\
& d_{h}^{\operatorname{mix}}\left(s, a^{2}\right)=\eta d_{h}^{\pi^{\mathrm{E}}}\left(s, a^{2}\right)+(1-\eta) d_{h}^{\pi^{\beta}}\left(s, a^{2}\right)=(1-\eta) d_{h}^{\pi^{\beta}}\left(s, a^{2}\right)=(1-\eta) \rho(s) .
\end{aligned}
$$

Note that in the population level, the marginal distribution of the union dataset $\mathcal{D}^{\mathrm{U}}$ in time step $h$ is exactly $d_{h}^{\text {mix }}$. The mixture policy induced by $d^{\text {mix }}$ (introduced in Appendix C.2) can be formulated as

$$
\pi_{h}^{\operatorname{mix}}\left(a^{1} \mid s\right)=\eta, \pi_{h}^{\operatorname{mix}}\left(a^{2} \mid s\right)=1-\eta, \forall s \in \mathcal{S}, h \in[H] .
$$

Just like before, we have $d_{h}^{\pi^{\text {mix }}}(s, a)=d_{h}^{\text {mix }}(s, a)$. The policy value of $\pi^{\text {mix }}$ can be calculated as

$$
V\left(\pi^{\mathrm{mix}}\right)=\sum_{h=1}^{H} \sum_{(s, a) \in \mathcal{S} \times \mathcal{A}} d_{h}^{\mathrm{mix}}(s, a) r_{h}(s, a)=\sum_{h=1}^{H} \sum_{s \in \mathcal{S}} d_{h}^{\mathrm{mix}}\left(s, a^{1}\right)=\eta H .
$$

Recall from Eq. (9) that $\pi^{\mathrm{NBCU}}$ can be formulated as

$$
\forall h \in[H], \quad \pi_{h}^{\mathrm{NBCU}}(a \mid s)= \begin{cases}\frac{n_{h}^{\mathrm{U}}(s, a)}{\sum_{a^{\prime}} n_{h}^{\mathrm{U}}\left(s, a^{\prime}\right)} & \text { if } \sum_{a^{\prime}} n_{h}^{\mathrm{U}}\left(s, a^{\prime}\right)>0  \tag{13}\\ \frac{1 \mathcal{A} \mid}{} & \text { otherwise }\end{cases}
$$

We can view that the BC's policy learned on the union dataset mimics the mixture policy $\pi^{\text {mix }}$. In the following part, we analyze the lower bound on the imitation gap of $\pi^{\mathrm{NBCU}}$.

$$
\begin{aligned}
\mathbb{E}\left[V\left(\pi^{\mathrm{E}}\right)-V\left(\pi^{\mathrm{NBCU}}\right)\right] & =V\left(\pi^{\mathrm{E}}\right)-V\left(\pi^{\mathrm{mix}}\right)+\mathbb{E}\left[V\left(\pi^{\mathrm{mix}}\right)-V\left(\pi^{\mathrm{NBCU}}\right)\right] \\
& =H-\eta H+\mathbb{E}\left[V\left(\pi^{\mathrm{mix}}\right)-V\left(\pi^{\mathrm{NBCU}}\right)\right] \\
& =(1-\eta)\left(V\left(\pi^{\mathrm{E}}\right)-V\left(\pi^{\beta}\right)\right)+\mathbb{E}\left[V\left(\pi^{\mathrm{mix}}\right)-V\left(\pi^{\mathrm{NBCU}}\right)\right] .
\end{aligned}
$$

595 Then we consider the term $\mathbb{E}\left[V\left(\pi^{\text {mix }}\right)-V\left(\pi^{\mathrm{NBCU}}\right)\right]$.

$$
\begin{aligned}
& V\left(\pi^{\mathrm{mix}}\right)-V\left(\pi^{\mathrm{NBCU}}\right) \\
= & \sum_{h=1}^{H} \sum_{(s, a) \in \mathcal{S} \times \mathcal{A}}\left(d_{h}^{\pi^{\mathrm{mix}}}(s, a)-d_{h}^{\pi^{\mathrm{NBCU}}}(s, a)\right) r_{h}(s, a)
\end{aligned}
$$

$$
\begin{aligned}
= & \sum_{h=1}^{H} \sum_{(s, a) \in \mathcal{S} \times \mathcal{A}} \rho(s)\left(\pi_{h}^{\mathrm{mix}}(a \mid s)-\pi_{h}^{\mathrm{NBCU}}(a \mid s)\right) r_{h}(s, a) \\
= & \sum_{h=1}^{H} \sum_{(s, a) \in \mathcal{S} \times \mathcal{A}} \rho(s)\left(\pi_{h}^{\mathrm{mix}}(a \mid s)-\pi_{h}^{\mathrm{NBCU}}(a \mid s)\right) r_{h}(s, a) \mathbb{I}\left\{n_{h}^{\mathrm{U}}(s)>0\right\} \\
& +\sum_{h=1}^{H} \sum_{(s, a) \in \mathcal{S} \times \mathcal{A}} \rho(s)\left(\pi_{h}^{\mathrm{mix}}(a \mid s)-\pi_{h}^{\mathrm{NBCU}}(a \mid s)\right) r_{h}(s, a) \mathbb{I}\left\{n_{h}^{\mathrm{U}}(s)=0\right\} .
\end{aligned}
$$

We take expectation over the randomness in $\mathcal{D}^{\mathrm{U}}$ on both sides and obtain that

$$
\begin{align*}
& \mathbb{E}\left[V\left(\pi^{\mathrm{mix}}\right)-V\left(\pi^{\mathrm{NBCU}}\right)\right]  \tag{14}\\
= & \mathbb{E}\left[\sum_{h=1}^{H} \sum_{(s, a) \in \mathcal{S} \times \mathcal{A}} \rho(s)\left(\pi_{h}^{\mathrm{mix}}(a \mid s)-\pi_{h}^{\mathrm{NBCU}}(a \mid s)\right) r_{h}(s, a) \mathbb{I}\left\{n_{h}^{\mathrm{U}}(s)>0\right\}\right] \\
& +\mathbb{E}\left[\sum_{h=1}^{H} \sum_{(s, a) \in \mathcal{S} \times \mathcal{A}} \rho(s)\left(\pi_{h}^{\mathrm{mix}}(a \mid s)-\pi_{h}^{\mathrm{NBCU}}(a \mid s)\right) r_{h}(s, a) \mathbb{I}\left\{n_{h}^{\mathrm{U}}(s)=0\right\}\right] . \tag{15}
\end{align*}
$$

7 For the first term in RHS, we have that

$$
\begin{aligned}
& \mathbb{E}\left[\sum_{h=1}^{H} \sum_{(s, a) \in \mathcal{S} \times \mathcal{A}} \rho(s)\left(\pi_{h}^{\text {mix }}(a \mid s)-\pi_{h}^{\mathrm{NBCU}}(a \mid s)\right) r_{h}(s, a) \mathbb{I}\left\{n_{h}^{\mathrm{U}}(s)>0\right\}\right] \\
= & \sum_{h=1}^{H} \sum_{(s, a) \in \mathcal{S} \times \mathcal{A}} \rho(s) r_{h}(s, a) \mathbb{E}\left[\left(\pi_{h}^{\operatorname{mix}}(a \mid s)-\pi_{h}^{\mathrm{NBCU}}(a \mid s)\right) \mathbb{I}\left\{n_{h}^{\mathrm{U}}(s)>0\right\}\right] \\
= & \sum_{h=1}^{H} \sum_{(s, a) \in \mathcal{S} \times \mathcal{A}} \rho(s) r_{h}(s, a) \mathbb{P}\left(n_{h}^{\mathrm{U}}(s)>0\right) \mathbb{E}\left[\pi_{h}^{\operatorname{mix}}(a \mid s)-\pi_{h}^{\mathrm{NBCU}}(a \mid s) \mid n_{h}^{\mathrm{U}}(s)>0\right] \\
= & 0 .
\end{aligned}
$$

The last equation follows the fact that $\pi_{h}^{\mathrm{NBCU}}(a \mid s)$ is an unbiased estimation of $\pi_{h}^{\operatorname{mix}}(a \mid s)$, so $\mathbb{E}\left[\pi_{h}^{\text {mix }}(a \mid s)-\pi_{h}^{\mathrm{NBCU}}(a \mid s) \mid n_{h}^{\mathrm{U}}(s)>0\right]$. For the second term in Eq. (15), we have that

$$
\begin{aligned}
& \mathbb{E}\left[\sum_{h=1}^{H} \sum_{(s, a) \in \mathcal{S} \times \mathcal{A}} \rho(s)\left(\pi_{h}^{\text {mix }}(a \mid s)-\pi_{h}^{\mathrm{NBCU}}(a \mid s)\right) r_{h}(s, a) \mathbb{I}\left\{n_{h}^{\mathrm{U}}(s)=0\right\}\right] \\
&= \sum_{h=1}^{H} \sum_{(s, a) \in \mathcal{S} \times \mathcal{A}} \rho(s) r_{h}(s, a) \mathbb{E}\left[\left(\pi_{h}^{\mathrm{mix}}(a \mid s)-\pi_{h}^{\mathrm{NBCU}}(a \mid s)\right) \mathbb{I}\left\{n_{h}^{\mathrm{U}}(s)=0\right\}\right] \\
&= \sum_{h=1}^{H} \sum_{(s, a) \in \mathcal{S} \times \mathcal{A}} \rho(s) r_{h}(s, a) \mathbb{P}\left(n_{h}^{\mathrm{U}}(s)=0\right) \mathbb{E}\left[\pi_{h}^{\mathrm{mix}}(a \mid s)-\pi_{h}^{\mathrm{NBCU}}(a \mid s) \mid n_{h}^{\mathrm{U}}(s)=0\right] \\
&= \sum_{h=1}^{H} \sum_{(s, a) \in \mathcal{S} \times \mathcal{A}} \rho(s) r_{h}(s, a) \mathbb{P}\left(n_{h}^{\mathrm{U}}(s)=0\right)\left(\pi_{h}^{\mathrm{mix}}(a \mid s)-\frac{1}{|\mathcal{A}|}\right) \\
& \stackrel{(a)}{=} \sum_{h=1}^{H} \sum_{s \in \mathcal{S}} \rho(s) \mathbb{P}\left(n_{h}^{\mathrm{U}}(s)=0\right)\left(\eta-\frac{1}{|\mathcal{A}|}\right) \\
& \stackrel{(b)}{=} H\left(\eta-\frac{1}{|\mathcal{A}|}\right) \sum_{s \in \mathcal{S}} \rho(s) \mathbb{P}\left(n_{1}^{\mathrm{U}}(s)=0\right) .
\end{aligned}
$$

In the equation $(a)$, we use the fact that $r_{h}\left(s, a^{1}\right)=1$ but $r_{h}(s, a)=0$ for any $a \neq a^{1}$. In the equation (b), since each state is an absorbing state, we have that $\mathbb{P}\left(n_{h}^{U}(s)=0\right)=\mathbb{P}\left(n_{1}^{U}(s)=0\right)$ for any $h \in[H]$. We consider two cases to address RHS of equation (b). In the first case of $\eta \geq 1 /|\mathcal{A}|$,

In the inequality $(a)$, we use that

$$
\sum_{s \in \mathcal{S}} \rho(s) \mathbb{P}\left(n_{1}^{\mathrm{U}}(s)=0\right)=\sum_{s \in \mathcal{S}} \rho(s)(1-\rho(s))^{N_{\mathrm{tot}}}=\left(1-\frac{1}{|\mathcal{S}|}\right)^{N_{\mathrm{tot}}} \leq \exp \left(-\frac{N_{\mathrm{tot}}}{|\mathcal{S}|}\right)
$$

This implies that

$$
\begin{aligned}
\mathbb{E}\left[V\left(\pi^{\mathrm{E}}\right)-V\left(\pi^{\mathrm{NBCU}}\right)\right] & \geq(1-\eta)\left(V\left(\pi^{\mathrm{E}}\right)-V\left(\pi^{\beta}\right)\right)-\frac{(1-\eta) H}{2} \\
& =\frac{(1-\eta)}{2}\left(V\left(\pi^{\mathrm{E}}\right)-V\left(\pi^{\beta}\right)\right) .
\end{aligned}
$$

## D. 1 Proof of Proposition 2

In the tabular case, with the first-order optimality condition, we have $c_{h}^{\star}(s, a)=\widehat{d_{h}^{\mathrm{E}}}(s, a) / \widehat{d_{h}^{\mathrm{E}}}(s, a)+$ $\left.\widehat{d_{h}^{\mathrm{U}}}(s, a)\right)$. By Eq. (5), we have

$$
\widehat{d_{h}^{\mathrm{U}}}(s, a) w_{h}(s, a)=\widehat{d_{h}^{\mathrm{U}}}(s, a) \times \frac{\widehat{d_{h}^{\mathrm{E}}}(s, a)}{\widehat{d_{h}^{\mathrm{U}}}(s, a)}=\widehat{d_{h}^{\mathrm{E}}}(s, a) .
$$

Recall that

$$
\Delta_{h}(\theta)=\min _{(s, a) \in \mathcal{D}_{h}^{\mathrm{E}} \cup \mathcal{D}_{h}^{\mathrm{S}, 1}}\left\langle\theta, \phi_{h}(s, a)\right\rangle-\max _{\left(s^{\prime}, a^{\prime}\right) \in \mathcal{D}_{h}^{\mathrm{S}, 2}}\left\langle\theta, \phi_{h}\left(s^{\prime}, a^{\prime}\right)\right\rangle .
$$

620 Then we have that

$$
\Delta_{h}\left(\bar{\theta}_{h}\right)-\Delta_{h}(\theta)=\min _{(s, a) \in \mathcal{D}_{h}^{\mathrm{E}} \cup \mathcal{D}_{h}^{\mathrm{S}, 1}}\left\langle\bar{\theta}_{h}, \phi_{h}(s, a)\right\rangle-\max _{\left(s^{\prime}, a^{\prime}\right) \in \mathcal{D}_{h}^{\mathrm{S}, 2}}\left\langle\bar{\theta}_{h}, \phi_{h}\left(s^{\prime}, a^{\prime}\right)\right\rangle
$$

$$
\begin{aligned}
& -\min _{(s, a) \in \mathcal{D}_{h}^{\mathrm{E}} \cup \mathcal{D}_{h}^{\mathrm{S}, 1}}
\end{aligned}\left\langle\theta, \phi_{h}(s, a)\right\rangle+\max _{\left(s^{\prime}, a^{\prime}\right) \in \mathcal{D}_{h}^{\mathrm{S}, 2}}\left\langle\theta, \phi_{h}\left(s^{\prime}, a^{\prime}\right)\right\rangle, \begin{aligned}
& (a) \\
& \leq
\end{aligned}\left\langle\bar{\theta}_{h}, \phi_{h}\left(s^{1}, a^{1}\right)\right\rangle-\left\langle\bar{\theta}_{h}, \phi_{h}\left(s^{2}, a^{2}\right)\right\rangle-\left\langle\theta, \phi_{h}\left(s^{1}, a^{1}\right)\right\rangle+\left\langle\theta, \phi_{h}\left(s^{2}, a^{2}\right)\right\rangle, \begin{aligned}
& = \\
& = \\
& \left.\stackrel{\left(\bar{\theta}_{h}\right.}{ }-\theta, \phi_{h}\left(s^{1}, a^{1}\right)-\phi_{h}\left(s^{2}, a^{2}\right)\right\rangle \\
& \leq
\end{aligned}\left\|\bar{\theta}_{h}-\theta\right\|\left\|\phi_{h}\left(s^{1}, a^{1}\right)-\phi_{h}\left(s^{2}, a^{2}\right)\right\| . \quad .
$$

In inequality $(a)$, we utilize the facts that $\left(s^{1}, a^{1}\right) \in \operatorname{argmin}_{(s, a) \in \mathcal{D}_{h}^{\mathrm{E}} \cup \mathcal{D}_{h}^{\mathrm{S}, 1}}\left\langle\theta_{h}, \phi_{h}(s, a)\right\rangle$ and $\left(s^{2}, a^{2}\right) \in \operatorname{argmax}_{(s, a) \in \mathcal{D}_{h}^{\mathrm{s}, 2}}\left\langle\theta_{h}, \phi_{h}(s, a)\right\rangle$. Inequality $(b)$ follows the Cauchy-Schwarz inequality. Let $L_{h}=\left\|\phi_{h}\left(s^{1}, a^{1}\right)-\stackrel{\phi}{\phi}_{h}\left(s^{2}, a^{2}\right)\right\|$ and we finish the proof.

## D. 3 Proof of Lemma 2

First, by Taylor's Theorem, there exists $\theta_{h}^{\prime} \in\left\{\theta \in \mathbb{R}^{d}: \theta^{t}=\theta_{h}^{\star}+t\left(\bar{\theta}_{h}-\theta_{h}^{\star}\right), \forall t \in[0,1]\right\}$ such that

$$
\begin{align*}
\mathcal{L}_{h}\left(\bar{\theta}_{h}\right) & =\mathcal{L}_{h}\left(\theta_{h}^{\star}\right)+\left\langle\nabla \mathcal{L}_{h}\left(\theta_{h}^{\star}\right), \bar{\theta}_{h}-\theta_{h}^{\star}\right\rangle+\frac{1}{2}\left(\bar{\theta}_{h}-\theta_{h}^{\star}\right)^{\top} \nabla^{2} \mathcal{L}_{h}\left(\theta_{h}^{\prime}\right)\left(\bar{\theta}_{h}-\theta_{h}^{\star}\right) \\
& =\mathcal{L}_{h}\left(\theta_{h}^{\star}\right)+\frac{1}{2}\left(\bar{\theta}_{h}-\theta_{h}^{\star}\right)^{\top} \nabla^{2} \mathcal{L}_{h}\left(\theta_{h}^{\prime}\right)\left(\bar{\theta}_{h}-\theta_{h}^{\star}\right) . \tag{16}
\end{align*}
$$

The last equality follows the optimality condition that $\nabla \mathcal{L}_{h}\left(\theta_{h}^{\star}\right)=0$. Then, our strategy is to prove that the smallest eigenvalue of the Hessian matrix $\nabla^{2} \mathcal{L}_{h}\left(\theta_{h}^{\prime}\right)$ is positive, i.e., $\lambda_{\min }\left(\nabla^{2} \mathcal{L}_{h}\left(\theta_{h}^{\prime}\right)\right)>0$. We first calculate the Hessian matrix $\nabla^{2} \mathcal{L}_{h}\left(\theta_{h}^{\prime}\right)$. Given $\mathcal{D}^{\mathrm{E}}$ and $\mathcal{D}^{\mathrm{U}}$, we define the function $G$ : $\mathbb{R}^{\left(\left|\mathcal{D}^{\mathrm{E}}\right|+\left|\mathcal{D}^{\mathrm{U}}\right|\right)} \rightarrow \mathbb{R}$ as

$$
G(v) \triangleq \frac{1}{\left|\mathcal{D}^{\mathrm{E}}\right|} \sum_{i=1}^{\left|\mathcal{D}^{\mathrm{E}}\right|} g\left(v_{i}\right)+\frac{1}{\left|\mathcal{D}^{\mathrm{U}}\right|} \sum_{j=1}^{\left|\mathcal{D}^{\mathrm{U}}\right|} g\left(v_{j}\right)
$$

where $v_{i}$ is the $i$-th element in the vector $v \in \mathbb{R}^{\left(\left|\mathcal{D}^{\mathrm{E}}\right|+\left|\mathcal{D}^{\mathrm{U}}\right|\right)}$ and $g(x)=\log (1+\exp (x))$ is a real-valued function. Besides, we use $B_{h} \in \mathbb{R}^{\left(\left|\mathcal{D}^{\mathrm{E}}\right|+\left|\mathcal{D}^{\mathrm{U}}\right|\right) \times d}$ to denote the matrix whose $i$-th row $B_{h, i}=-y_{i} \phi_{h}\left(s^{i}, a^{i}\right)^{\top}$, and $y_{i}=1$ if $\left(s^{i}, a^{i}\right) \in \mathcal{D}_{h}^{\mathrm{E}}, y_{i}=-1$ if $\left(s^{i}, a^{i}\right) \notin \mathcal{D}_{h}^{\mathrm{E}}$. Then the objective function can be reformulated as

$$
\begin{aligned}
& \mathcal{L}_{h}\left(\theta_{h}\right) \\
= & \sum_{(s, a)} \widehat{d_{h}^{\mathrm{E}}}(s, a)\left[\log \left(1+\exp \left(-\left\langle\phi_{h}(s, a), \theta_{h}\right\rangle\right)\right)\right]+\sum_{(s, a)} \widehat{d_{h}^{\mathrm{U}}}(s, a)\left[\log \left(1+\exp \left(\left\langle\phi_{h}(s, a), \theta_{h}\right\rangle\right)\right)\right] \\
= & \frac{1}{\left|\mathcal{D}^{\mathrm{E}}\right|} \sum_{(s, a) \in \mathcal{D}^{\mathrm{E}}} \log \left(1+\exp \left(-\left\langle\phi_{h}(s, a), \theta_{h}\right\rangle\right)\right)+\frac{1}{\left|\mathcal{D}^{\mathrm{U}}\right|} \sum_{(s, a) \in \mathcal{D}^{\mathrm{U}}} \log \left(1+\exp \left(\left\langle\phi_{h}(s, a), \theta_{h}\right\rangle\right)\right) \\
= & G\left(B_{h} \theta_{h}\right)
\end{aligned}
$$

Then we have that $\nabla^{2} \mathcal{L}_{h}\left(\theta_{h}\right)=B_{h}^{\top} \nabla^{2} G\left(B_{h} \theta_{h}\right) B_{h}$, where

$$
\begin{aligned}
& \nabla^{2} G\left(B_{h} \theta_{h}\right) \\
= & \operatorname{diag}\left(\frac{g^{\prime \prime}\left(\left(B_{h} \theta_{h}\right)_{1}\right)}{\left|\mathcal{D}^{\mathrm{E}}\right|}, \ldots, \frac{g^{\prime \prime}\left(\left(B_{h} \theta_{h}\right)_{\left|\mathcal{D}^{\mathrm{E}}\right|}\right)}{\left|\mathcal{D}^{\mathrm{E}}\right|}, \frac{g^{\prime \prime}\left(\left(B_{h} \theta_{h}\right)_{\left|\mathcal{D}^{\mathrm{E}}\right|+1}\right)}{\left|\mathcal{D}^{\mathrm{E}}\right|+\left|\mathcal{D}^{\mathrm{U}}\right|}, \ldots, \frac{g^{\prime \prime}\left(\left(B_{h} \theta_{h}\right)_{\left|\mathcal{D}^{\mathrm{E}}\right|+\left|\mathcal{D}^{\mathrm{U}}\right|}\right)}{\left|\mathcal{D}^{\mathrm{E}}\right|+\left|\mathcal{D}^{\mathrm{U}}\right|}\right) .
\end{aligned}
$$

Here $g^{\prime \prime}(x)=\sigma(x)(1-\sigma(x))$, where $\sigma(x)=1 /(1+\exp (-x))$ is the sigmoid function. The eigenvalues of $\nabla^{2} G\left(B_{h} \theta_{h}\right)$ are

$$
\left\{\frac{g^{\prime \prime}\left(\left(B_{h} \theta_{h}\right)_{1}\right)}{\left|\mathcal{D}^{\mathrm{E}}\right|}, \ldots, \frac{g^{\prime \prime}\left(\left(B_{h} \theta_{h}\right)_{\left|\mathcal{D}^{\mathrm{E}}\right|}\right)}{\left|\mathcal{D}^{\mathrm{E}}\right|}, \frac{g^{\prime \prime}\left(\left(B_{h} \theta_{h}\right)_{\left|\mathcal{D}^{\mathrm{E}}\right|+1}\right)}{\left|\mathcal{D}^{\mathrm{E}}\right|+\left|\mathcal{D}^{\mathrm{U}}\right|}, \ldots, \frac{g^{\prime \prime}\left(\left(B_{h} \theta_{h}\right)_{\left|\mathcal{D}^{\mathrm{E}}\right|+\left|\mathcal{D}^{\mathrm{U}}\right|}\right)}{\left|\mathcal{D}^{\mathrm{E}}\right|+\left|\mathcal{D}^{\mathrm{U}}\right|}\right\}
$$

Notice that $\theta_{h}^{\prime} \in\left\{\theta \in \mathbb{R}^{d}: \theta^{t}=\theta_{h}^{\star}+t\left(\bar{\theta}_{h}-\theta_{h}^{\star}\right), \forall t \in[0,1]\right\}$. For a matrix $A$, we use $\lambda_{\min }(A)$ to denote the minimal eigenvalue of $A$. Here we claim that the minimum of the minimal eigenvalues of $\nabla^{2} G\left(B_{h} \theta^{t}\right)$ over $t \in[0,1]$ is achieved at $t=0$ or $t=1$. That is,

$$
\min \left\{\lambda_{\min }\left(\nabla^{2} G\left(B_{h} \theta^{t}\right)\right): \forall t \in[0,1]\right\}=\min \left\{\lambda_{\min }\left(\nabla^{2} G\left(B_{h} \theta^{0}\right)\right), \lambda_{\min }\left(\nabla^{2} G\left(B_{h} \theta^{1}\right)\right)\right\}
$$

We prove this claim as follows. For any $t \in[0,1]$, we use $\left\{\lambda_{1}(t), \ldots, \lambda_{\left|\mathcal{D}^{\mathrm{E}}\right|+\left|\mathcal{D}^{\mathrm{U}}\right|}(t)\right\}$ to denote the eigenvalues of $\nabla^{2} G\left(B_{h} \theta^{t}\right)$. For each $i \in\left[\left|\mathcal{D}^{\mathrm{E}}\right|+\left|\mathcal{D}^{\mathrm{U}}\right|\right]$, we consider $\lambda_{i}(t):[0,1] \rightarrow \mathbb{R}$ as a function of $t$. Specifically,

$$
\lambda_{i}(t)= \begin{cases}\frac{g^{\prime \prime}\left(\left(B_{h} \theta_{h}^{\star}\right)_{i}+t\left(B_{h}\left(\bar{\theta}_{h}-\theta_{h}^{\star}\right)\right)_{i}\right)}{\left|\mathcal{D}^{\mathrm{E}}\right|}, & \text { if } i \in\left[\left|\mathcal{D}^{\mathrm{E}}\right|\right] \\ \frac{g^{\prime \prime}\left(\left(B_{h} \theta_{h}^{\star}\right)_{i}+t\left(B_{h}\left(\bar{\theta}_{h}-\theta_{h}^{\star}\right)\right)_{i}\right)}{\left|\mathcal{D}^{\mathrm{E}}+\left|\mathcal{D}^{\mathrm{U}}\right|\right.}, & \text { otherwise. }\end{cases}
$$

We observe that $g^{\prime \prime \prime}(x)=\sigma(x)(1-\sigma(x))(1-2 \sigma(x))$ which satisfies that $\forall x \leq 0, g^{\prime \prime \prime}(x) \geq 0$, and $\forall x \geq 0, g^{\prime \prime \prime}(x) \leq 0$. Therefore, we have that the minimum of $\lambda_{i}(t)$ over $t \in[0,1]$ must be achieved at $t=0$ or $t=1$. That is,

$$
\begin{equation*}
\min _{t \in[0,1]} \lambda_{i}(t)=\min \left\{\lambda_{i}(0), \lambda_{i}(1)\right\} . \tag{17}
\end{equation*}
$$

For any $t \in[0,1]$, we define $i^{t} \in\left[\left|\mathcal{D}^{\mathrm{E}}\right|+\left|\mathcal{D}^{\mathrm{U}}\right|\right]$ as the index of the minimal eigenvalue of $\nabla^{2} G\left(B_{h} \theta^{t}\right)$, i.e., $\lambda_{i^{t}}(t)=\lambda_{\text {min }}\left(\nabla^{2} G\left(B_{h} \theta^{t}\right)\right)$. Then we have that

$$
\begin{aligned}
\min \left\{\lambda_{\min }\left(\nabla^{2} G\left(B_{h} \theta^{t}\right)\right): \forall t \in[0,1]\right\} & =\min \left\{\lambda_{i^{t}}(t): \forall t \in[0,1]\right\} \\
& \stackrel{(a)}{=} \min \left\{\min \left\{\lambda_{i^{t}}(0), \lambda_{i^{t}}(1)\right\}: \forall t \in[0,1]\right\} \\
& =\min \left\{\lambda_{i^{0}}(0), \lambda_{i^{1}}(1)\right\} \\
& \stackrel{(b)}{=} \min \left\{\lambda_{\min }\left(\nabla^{2} G\left(B_{h} \theta^{0}\right)\right), \lambda_{\min }\left(\nabla^{2} G\left(B_{h} \theta^{1}\right)\right)\right\}
\end{aligned}
$$

Equality (a) follows (17) and equality (b) follows that $\lambda_{i^{0}}(0)$ and $\lambda_{i^{1}}(1)$ are the minimal eigenvalues of $\nabla^{2} G\left(B_{h} \theta^{0}\right)$ and $\nabla^{2} G\left(B_{h} \theta^{1}\right)$, respectively.
In summary, we derive that

$$
\begin{equation*}
\min \left\{\lambda_{\min }\left(\nabla^{2} G\left(B_{h} \theta^{t}\right)\right): \forall t \in[0,1]\right\}=\min \left\{\lambda_{\min }\left(\nabla^{2} G\left(B_{h} \theta^{0}\right)\right), \lambda_{\min }\left(\nabla^{2} G\left(B_{h} \theta^{1}\right)\right)\right\} \tag{18}
\end{equation*}
$$

which proves the previous claim.
Further, we consider $\lambda_{\text {min }}\left(\nabla^{2} \mathcal{L}_{h}\left(\theta_{h}\right)\right)$.

$$
\begin{aligned}
\lambda_{\min }\left(\nabla^{2} \mathcal{L}_{h}\left(\theta_{h}\right)\right) & =\inf _{x \in \mathbb{R}^{d}:\|x\|=1} x^{\top} \nabla^{2} \mathcal{L}_{h}\left(\theta_{h}\right) x \\
& =\inf _{x \in \mathbb{R}^{d}:\|x\|=1}\left(B_{h} x\right)^{\top} \nabla^{2} G\left(B_{h} \theta_{h}\right)\left(B_{h} x\right) \\
& =\inf _{z \in \operatorname{Im}\left(B_{h}\right)} z^{\top} \nabla^{2} G\left(B_{h} \theta_{h}\right) z \\
& =\left(\inf _{z \in \operatorname{Im}\left(B_{h}\right)}\|z\|\right)^{2} \lambda_{\min }\left(\nabla^{2} G\left(B_{h} \theta_{h}\right)\right) \\
& \geq\left(\inf _{z \in \operatorname{Im}\left(B_{h}\right)}\|z\|\right)^{2} \min \left\{\lambda_{\min }\left(\nabla^{2} G\left(B_{h} \theta^{0}\right)\right), \lambda_{\min }\left(\nabla^{2} G\left(B_{h} \theta^{1}\right)\right)\right\}
\end{aligned}
$$

Here $\operatorname{Im}\left(B_{h}\right)=\left\{z \in \mathbb{R}^{d}: z=B_{h} x,\|x\|=1\right\}$. The last inequality follows Eq. (18).
Recall we assume that $\operatorname{rank}\left(A_{h}\right)=d$, so we have that $\operatorname{rank}\left(B_{h}\right)=d$. Thus, $\operatorname{Im}\left(B_{h}\right)$ is a set of vectors with positive norms, i.e., $\inf _{z \in \operatorname{Im}\left(B_{h}\right)}\|z\|>0$. Besides, since $g^{\prime \prime}(x)=\sigma(x)(1-\sigma(x))>0$, we also have that

$$
\min \left\{\lambda_{\min }\left(\nabla^{2} G\left(B_{h} \theta^{0}\right)\right), \lambda_{\min }\left(\nabla^{2} G\left(B_{h} \theta^{1}\right)\right)\right\}>0
$$

In summary, we obtain that

$$
\lambda_{\min }\left(\nabla^{2} \mathcal{L}_{h}\left(\theta_{h}\right)\right) \geq\left(\inf _{z \in \operatorname{Im}\left(B_{h}\right)}\|z\|\right)^{2} \min \left\{\lambda_{\min }\left(\nabla^{2} G\left(B_{h} \theta^{0}\right)\right), \lambda_{\min }\left(\nabla^{2} G\left(B_{h} \theta^{1}\right)\right)\right\}>0
$$

Then, with Eq. (16), there exists

$$
\tau_{h}=\left(\inf _{z \in \operatorname{Im}\left(B_{h}\right)}\|z\|\right)^{2} \min \left\{\lambda_{\min }\left(\nabla^{2} G\left(B_{h} \theta^{0}\right)\right), \lambda_{\min }\left(\nabla^{2} G\left(B_{h} \theta^{1}\right)\right)\right\}>0
$$

such that

$$
\mathcal{L}_{h}\left(\bar{\theta}_{h}\right) \geq \mathcal{L}_{h}\left(\theta_{h}^{\star}\right)+\frac{\tau_{h}}{2}\left\|\bar{\theta}_{h}-\theta_{h}^{\star}\right\|^{2}
$$

## D. 4 Proof of Theorem 3

First, invoking Lemma 1 with $\theta=\theta_{h}^{\star}$ yields that

$$
\Delta_{h}\left(\theta_{h}^{\star}\right) \geq \Delta_{h}\left(\bar{\theta}_{h}\right)-L_{h}\left\|\bar{\theta}_{h}-\theta_{h}^{\star}\right\| .
$$

Here $L_{h}=\left\|\phi_{h}(s, a)-\phi_{h}\left(s^{\prime}, a^{\prime}\right)\right\|$ with $(s, a) \in \operatorname{argmin}_{(s, a) \in \mathcal{D}_{h}^{\mathrm{E}} \cup \mathcal{D}_{h}^{\mathrm{S}, 1}}\left\langle\theta_{h}^{\star}, \phi_{h}(s, a)\right\rangle$ and $\left(s^{\prime}, a^{\prime}\right) \in \operatorname{argmax}_{(s, a) \in \mathcal{D}_{h}^{\mathrm{s}, 2}}\left\langle\theta_{h}^{\star}, \phi_{h}(s, a)\right\rangle$. Then, by Lemma 2, there exists $\tau_{h}>0$ such that

$$
\mathcal{L}_{h}\left(\bar{\theta}_{h}\right) \geq \mathcal{L}_{h}\left(\theta_{h}^{\star}\right)+\frac{\tau_{h}}{2}\left\|\bar{\theta}_{h}-\theta_{h}^{\star}\right\|^{2} .
$$

This directly implies an upper bound of the distance between $\bar{\theta}_{h}$ and $\theta_{h}^{\star}$.

$$
\left\|\bar{\theta}_{h}-\theta_{h}^{\star}\right\| \leq \sqrt{\frac{2\left(\mathcal{L}_{h}\left(\bar{\theta}_{h}\right)-\mathcal{L}_{h}\left(\theta_{h}^{\star}\right)\right)}{\tau_{h}}}
$$

If the feature is designed such that $\sqrt{\frac{2\left(\mathcal{L}_{h}\left(\bar{\theta}_{h}\right)-\mathcal{L}_{h}\left(\theta_{h}^{\star}\right)\right)}{\tau_{h}}}<\frac{\Delta_{h}\left(\bar{\theta}_{h}\right)}{L_{h}}$ holds, we further have that $\left\|\bar{\theta}_{h}-\theta_{h}^{\star}\right\|<\Delta_{h}\left(\bar{\theta}_{h}\right) / L_{h}$. Then we get that

$$
\Delta_{h}\left(\theta_{h}^{\star}\right) \geq \Delta_{h}\left(\bar{\theta}_{h}\right)-L_{h}\left\|\bar{\theta}_{h}-\theta_{h}^{\star}\right\|>0
$$

which completes the proof of the first statement.
Then we proceed to prove the imitation gap bound. We first identify the property of $\pi^{\mathrm{ISW}-\mathrm{BC}}$. Recall the objective of WBCU.

$$
\pi^{\mathrm{ISW}-\mathrm{BC}} \in \underset{\pi}{\operatorname{argmax}} \sum_{h=1}^{H} \sum_{(s, a) \in \mathcal{S} \times \mathcal{A}}\left\{\widehat{d_{h}^{\mathrm{U}}}(s, a) \times\left[w_{h}(s, a) \log \pi_{h}(a \mid s)\right] \times \mathbb{I}\left[w_{h}(s, a) \geq \delta\right]\right\}
$$

For any state $s$ with $\sum_{a \in \mathcal{A}} \widehat{d_{h}^{\mathrm{U}}}(s, a) w_{h}(s, a) \mathbb{I}\left[w_{h}(s, a) \geq \delta\right]>0$, with the first-order optimality condition, we have

$$
\pi_{h}^{\mathrm{ISW}-\mathrm{BC}}(a \mid s)=\frac{\widehat{d_{h}^{\mathrm{U}}}(s, a) w_{h}(s, a) \mathbb{I}\left[w_{h}(s, a) \geq \delta\right]}{\sum_{a \in \mathcal{A}} \widehat{d_{h}^{\mathrm{U}}}(s, a) w_{h}(s, a) \mathbb{I}\left[w_{h}(s, a) \geq \delta\right]} .
$$

For an expert state $s$ with $d_{h}^{\pi^{\mathrm{E}}}(s)>0$, if $\left(s, \pi_{h}^{\mathrm{E}}(s)\right) \in \mathcal{D}_{h}^{\mathrm{E}} \cup \mathcal{D}_{h}^{\mathrm{S}, 1}$, we have that

$$
\left\langle\theta_{h}^{\star}, \phi_{h}\left(s, \pi_{h}^{\mathrm{E}}(s)\right)\right\rangle>\left\langle\theta_{h}^{\star}, \phi_{h}(s, a)\right\rangle, \quad \forall(s, a) \in \mathcal{D}_{h}^{\mathrm{S}, 2} .
$$

This is due to the first statement that $\Delta_{h}\left(\theta_{h}^{\star}\right)>0$ in this theorem. Recall that

$$
c_{h}\left(s, a ; \theta_{h}^{\star}\right)=\frac{1}{1+\exp \left(-\left\langle\phi_{h}(s, a), \theta_{h}^{\star}\right\rangle\right)} \quad \text { and } \quad w_{h}(s, a)=\frac{c_{h}\left(s, a ; \theta_{h}^{\star}\right)}{1-c_{h}\left(s, a ; \theta_{h}^{\star}\right)} .
$$

We can further obtain that $w_{h}\left(s, \pi_{h}^{\mathrm{E}}(s)\right)>w_{h}(s, a)$ for any $(s, a) \in \mathcal{D}_{h}^{\mathrm{S}, 2}$. This implies that we can find a $\delta$ such that $\mathbb{I}\left[w_{h}\left(s, \pi_{h}^{\mathrm{E}}(s)\right) \geq \delta\right]=1$ for any $\left(s, \pi_{h}^{\mathrm{E}}(s)\right) \in \mathcal{D}_{h}^{\mathrm{E}} \cup \mathcal{D}_{h}^{\mathrm{S}, 1}$ and $\mathbb{I}\left[w_{h}(s, a) \geq \delta\right]=0$ for any $(s, a) \in \mathcal{D}_{h}^{\mathrm{S}, 2}$. Based on the above analytical form of $\pi^{\text {ISW-BC }}$, we have that $\pi^{\mathrm{ISW}-\mathrm{BC}}\left(\pi_{h}^{\mathrm{E}}(s) \mid s\right)=1$ for any $\left(s, \pi_{h}^{\mathrm{E}}(s)\right) \in \mathcal{D}_{h}^{\mathrm{E}} \cup \mathcal{D}_{h}^{\mathrm{S}, 1}$. In summary, for any state $s$ with $\left(s, \pi_{h}^{\mathrm{E}}(s)\right) \in \mathcal{D}_{h}^{\mathrm{E}} \cup \mathcal{D}_{h}^{\mathrm{S}, 1}$, we have that $\pi_{h}^{\mathrm{ISW}-\mathrm{BC}}\left(\pi_{h}^{\mathrm{E}}(s) \mid s\right)=1$.
With the above property of $\pi^{\text {ISW-BC }}$, we proceed to analyze the policy value gap. According to [32, Lemma 4.3], we have

$$
V\left(\pi^{\mathrm{E}}\right)-V\left(\pi^{\mathrm{ISW}-\mathrm{BC}}\right) \leq H \sum_{h=1}^{H} \mathbb{E}_{s \sim d_{h}^{\pi^{\mathrm{E}}}(\cdot)}\left[\mathrm{TV}\left(\pi_{h}^{\mathrm{E}}(\cdot \mid s), \pi_{h}^{\mathrm{ISW}-\mathrm{BC}}(\cdot \mid s)\right)\right]
$$

Since $\pi^{\mathrm{E}}$ is assumed to be deterministic, we have

$$
V\left(\pi^{\mathrm{E}}\right)-V\left(\pi^{\mathrm{ISW}-\mathrm{BC}}\right) \leq H \sum_{h=1}^{H} \mathbb{E}_{s \sim d_{h}^{\mathrm{E}}(\cdot)}\left[\mathbb{E}_{a \sim \pi_{h}^{\mathrm{ISW}-\mathrm{BC}}(\cdot \mid s)}\left[\mathbb{I}\left\{a \neq \pi_{h}^{\mathrm{E}}(s)\right\}\right]\right]
$$

$$
\begin{aligned}
& \stackrel{(a)}{\leq} H \sum_{h=1}^{H} \mathbb{E}_{s \sim d_{h}^{\pi^{\mathrm{E}}}(\cdot)}\left[\mathbb{I}\left\{\left(s, \pi_{h}^{\mathrm{E}}(s)\right) \notin \mathcal{D}_{h}^{\mathrm{E}} \cup \mathcal{D}_{h}^{\mathrm{S}, 1}\right\}\right] \\
& \stackrel{(b)}{=} H \sum_{h=1}^{H} \mathbb{E}_{s \sim d_{h}^{\pi^{\mathrm{E}}}(\cdot)}\left[\mathbb{I}\left\{\left(s, \pi_{h}^{\mathrm{E}}(s)\right) \notin \mathcal{D}_{h}^{\mathrm{U}}\right\}\right]
\end{aligned}
$$

Inequality $(a)$ follows the property of $\pi^{\mathrm{ISW}-\mathrm{BC}}$ derived above. In particular, for any state $s$ with $\left(s, \pi_{h}^{\mathrm{E}}(s)\right) \in \mathcal{D}_{h}^{\mathrm{E}} \cup \mathcal{D}_{h}^{\mathrm{S}, 1}$, we have that $\pi_{h}^{\mathrm{ISW}-\mathrm{BC}}\left(\pi_{h}^{\mathrm{E}}(s) \mid s\right)=1$. Equation $(b)$ holds due to the Assumption 2. In particular, for an expert state $s$ that $d_{h}^{\pi^{\mathrm{E}}}(s)>0$, the events of $\left(s, \pi_{h}^{\mathrm{E}}(s)\right) \notin$ $\mathcal{D}_{h}^{\mathrm{E}} \cup \mathcal{D}_{h}^{\mathrm{S}, 1}$ and $\left(s, \pi_{h}^{\mathrm{E}}(s)\right) \notin \mathcal{D}_{h}^{\mathrm{U}}$ are equivalent.

Moreover, we take the expectation over $\mathcal{D}^{\mathrm{U}}$ on both sides and obtain that

$$
\begin{aligned}
\mathbb{E}\left[V\left(\pi^{\mathrm{E}}\right)-V\left(\pi^{\mathrm{ISW}-\mathrm{BC}}\right)\right] & \leq H \sum_{h=1}^{H} \mathbb{E}_{s \sim d_{h}^{\pi^{\mathrm{E}}}(\cdot)}\left[\mathbb{P}\left(\left(s, \pi_{h}^{\mathrm{E}}(s)\right) \notin \mathcal{D}_{h}^{\mathrm{U}}\right)\right] \\
& =H \sum_{h=1}^{H} \sum_{s \in \mathcal{S}} d_{h}^{\pi^{\mathrm{E}}}(s) \mathbb{P}\left(\left(s, \pi_{h}^{\mathrm{E}}(s)\right) \notin \mathcal{D}_{h}^{\mathrm{U}}\right)
\end{aligned}
$$

According to Assumption 1, we have that

$$
\begin{aligned}
d_{h}^{\mathrm{U}}\left(s, \pi_{h}^{\mathrm{E}}(s)\right) & =\eta d_{h}^{\pi^{\mathrm{E}}}\left(s, \pi_{h}^{\mathrm{E}}(s)\right)+(1-\eta) d_{h}^{\pi^{\beta}}\left(s, \pi_{h}^{\mathrm{E}}(s)\right) \\
& \stackrel{(a)}{\geq} \eta d_{h}^{\pi^{\mathrm{E}}}\left(s, \pi_{h}^{\mathrm{E}}(s)\right)+\frac{(1-\eta)}{\mu} d_{h}^{\pi^{\mathrm{E}}}\left(s, \pi_{h}^{\mathrm{E}}(s)\right) \\
& =\left(\eta+\frac{(1-\eta)}{\mu}\right) d_{h}^{\pi^{\mathrm{E}}}\left(s, \pi_{h}^{\mathrm{E}}(s)\right)
\end{aligned}
$$

Inequality ( $a$ ) follows the definition of $\mu$ in Theorem 3: for any $(s, h) \in \mathcal{S} \times[H]$, we have $d_{h}^{\pi^{\mathrm{E}}}\left(s, \pi_{h}^{\mathrm{E}}(s)\right) / d_{h}^{\pi^{\beta}}\left(s, \pi_{h}^{\mathrm{E}}(s)\right) \leq \mu$. Then we obtain that

$$
\begin{aligned}
\mathbb{E}\left[V\left(\pi^{\mathrm{E}}\right)-V\left(\pi^{\mathrm{ISW}-\mathrm{BC}}\right)\right] & \leq H \sum_{h=1}^{H} \sum_{s \in \mathcal{S}} d_{h}^{\pi^{\mathrm{E}}}(s)\left(1-d_{h}^{\mathrm{U}}\left(s, \pi_{h}^{\mathrm{E}}(s)\right)\right)^{N_{\mathrm{tot}}} \\
& \leq\left(\frac{1}{\eta+(1-\eta) / \mu}\right) H \sum_{h=1}^{H} \sum_{s \in \mathcal{S}} d_{h}^{\mathrm{U}}\left(s, \pi_{h}^{\mathrm{E}}(s)\right) \mathbb{P}\left(\left(s, \pi_{h}^{\mathrm{E}}(s)\right) \notin \mathcal{D}_{h}^{\mathrm{U}}\right)
\end{aligned}
$$

For each $(s, h) \in \mathcal{S} \times[H]$, we observe that

$$
d_{h}^{\mathrm{U}}\left(s, \pi_{h}^{\mathrm{E}}(s)\right) \mathbb{P}\left(\left(s, \pi_{h}^{\mathrm{E}}(s)\right) \notin \mathcal{D}_{h}^{\mathrm{U}}\right)=d_{h}^{\mathrm{U}}\left(s, \pi_{h}^{\mathrm{E}}(s)\right)\left(1-d_{h}^{\mathrm{U}}\left(s, \pi_{h}^{\mathrm{E}}(s)\right)\right)^{N_{\mathrm{tot}}} \leq \frac{4}{9 N_{\mathrm{tot}}}
$$

Here the last inequality follows Lemma 5. Consequently, we can derive that

$$
\sum_{h=1}^{H} \sum_{s \in \mathcal{S}} d_{h}^{\mathrm{U}}\left(s, \pi_{h}^{\mathrm{E}}(s)\right) \mathbb{P}\left(\left(s, \pi_{h}^{\mathrm{E}}(s)\right) \notin \mathcal{D}_{h}^{\mathrm{U}}\right) \leq \frac{4 H|\mathcal{S}|}{9 N_{\mathrm{tot}}}
$$

which further implies that

$$
\mathbb{E}\left[V\left(\pi^{\mathrm{E}}\right)-V\left(\pi^{\mathrm{ISW}-\mathrm{BC}}\right)\right] \leq\left(\frac{1}{\eta+(1-\eta) / \mu}\right) \frac{4 H^{2}|\mathcal{S}|}{9 N_{\mathrm{tot}}}=\frac{4 H^{2}|\mathcal{S}|}{9\left(N_{\mathrm{E}}+N_{\mathrm{S}} / \mu\right)}
$$

We complete the proof.

## D. 5 An Example Corresponding to Theorem 3

In this section, we provide an example that illustrates the required feature design in Theorem 3 can hold.
Example 1. To illustrate Theorem 3, we consider an example in the feature space $\mathbb{R}^{2}$. In particular, for time step $h \in[H]$, we have the expert dataset and supplementary dataset as follows.

$$
\mathcal{D}_{h}^{\mathrm{E}}=\left\{\left(s^{(1)}, a^{(1)}\right),\left(s^{(4)}, a^{(4)}\right)\right\}, \mathcal{D}_{h}^{\mathrm{S}}=\left\{\left(s^{(2)}, a^{(2)}\right),\left(s^{(3)}, a^{(3)}\right)\right\}
$$

$$
\mathcal{D}_{h}^{\mathrm{S}, 1}=\left\{\left(s^{(2)}, a^{(2)}\right)\right\}, \mathcal{D}_{h}^{\mathrm{S}, 2}=\left\{\left(s^{(3)}, a^{(3)}\right)\right\}
$$

$$
\nabla^{2} \mathcal{L}_{h}\left(\theta_{h}\right)=\left(\begin{array}{cc}
\frac{3}{4} f\left(\theta_{h, 1}\right)+\frac{1}{16} f\left(\frac{1}{2} \theta_{h, 1}\right) & 0 \\
0 & \frac{3}{4} f\left(\theta_{h, 2}\right)+\frac{1}{16} f\left(\frac{1}{2} \theta_{h, 2}\right)
\end{array}\right)
$$

Then we calculate the parameter of strong convexity $\tau_{h}$ appears in Lemma 2. Based on the proof of Lemma 2, our strategy is to calculate the minimal eigenvalue of the Hessian matrix.
First, for $\theta_{h}=\left(\theta_{h, 1}, \theta_{h, 2}\right)^{\top}$, the gradient of $\mathcal{L}_{h}\left(\theta_{h}\right)$ is

$$
\begin{aligned}
& \nabla \mathcal{L}_{h}\left(\theta_{h}\right) \\
= & -\sum_{(s, a) \in \mathcal{S} \times \mathcal{A}} \widehat{d_{h}^{\mathrm{E}}}(s, a) \sigma\left(-\left\langle\phi_{h}(s, a), \theta_{h}\right\rangle\right)+\sum_{(s, a) \in \mathcal{S} \times \mathcal{A}} \widehat{d_{h}^{\mathrm{U}}}(s, a) \sigma\left(\left\langle\phi_{h}(s, a), \theta_{h}\right\rangle\right) \\
= & \left(\frac{1}{2} \sigma\left(\theta_{h, 1}\right)-\frac{1}{4} \sigma\left(-\theta_{h, 1}\right)-\frac{1}{8} \sigma\left(-\frac{1}{2} \theta_{h, 1}\right), \frac{1}{4} \sigma\left(\theta_{h, 2}\right)-\frac{1}{2} \sigma\left(-\theta_{h, 2}\right)-\frac{1}{8} \sigma\left(-\frac{1}{2} \theta_{h, 2}\right)\right)^{\top} .
\end{aligned}
$$

Here $\sigma(x)=1 /(1+\exp (-x))$ for $x \in \mathbb{R}$ is the sigmoid function. Then the Hessian matrix at $\theta_{h}$ is
where $f(x)=\sigma(x)(1-\sigma(x))$ and $f(x)=f(-x)$. For any $t \in[0,1]$, the eigenvalues of the Hessian matrix at $\theta_{h}^{t}=\bar{\theta}_{h}+t\left(\theta_{h}^{\star}-\bar{\theta}_{h}\right)$ are

$$
\frac{3}{4} f\left(\theta_{h, 1}^{t}\right)+\frac{1}{16} f\left(\frac{1}{2} \theta_{h, 1}^{t}\right), \frac{3}{4} f\left(\theta_{h, 2}^{t}\right)+\frac{1}{16} f\left(\frac{1}{2} \theta_{h, 2}^{t}\right) .
$$

Now, we calculate the minimal eigenvalues of $\nabla^{2} \mathcal{L}_{h}\left(\theta_{h}^{t}\right)$. We consider the function

$$
g(x)=\frac{3}{4} f(x)+\frac{1}{16} f\left(\frac{1}{2} x\right), \forall x \in[a, b] .
$$

The gradient is

$$
g^{\prime}(x)=\frac{3}{4} \sigma(x)(1-\sigma(x))(1-2 \sigma(x))+\frac{1}{32} \sigma\left(\frac{1}{2} x\right)\left(1-\sigma\left(\frac{1}{2} x\right)\right)\left(1-2 \sigma\left(\frac{1}{2} x\right)\right) .
$$

We observe that $\forall x \leq 0, g^{\prime}(x) \geq 0$, and $\forall x \geq 0, g^{\prime}(x) \leq 0$. Thus, we have that the minimum of $g(x)$ must be achieved at $x=a$ or $x=b$. Besides, we have that $g(x)=g(-x)$. With the above arguments, we know that the minimal eigenvalue is $g(0.993) \approx 0.163$ and $\tau_{h} \approx 0.163$. Then we can calculate that

$$
\sqrt{\frac{2\left(\mathcal{L}_{h}\left(\bar{\theta}_{h}\right)-\mathcal{L}_{h}\left(\theta_{h}^{\star}\right)\right)}{\tau_{h}}} \approx 0.520, \frac{\Delta_{h}\left(\bar{\theta}_{h}\right)}{L_{h}}=1
$$

The inequality in Theorem 3 holds.

## E Discussion

In the main text, we focus on the tabular representations for policies. Furthermore, we consider a trajectory sampling procedure for behavior policy in collecting the supplementary dataset. We present two possible extensions in this section.

## E. 1 Function Approximation of Policies

Assume that the learner is access to a finite function class $\Pi=\left\{\pi=\left(\pi_{1}, \pi_{2}, \ldots, \pi_{h}\right)\right\}$, where $\pi_{h}: \mathcal{S} \rightarrow \Delta(\mathcal{A})$ could be any function (e.g., neural networks). For simplicity of analysis, we assume that $\Pi$ is a finite class. Notice that the algorithms considered in this paper are BC and its variants, which all take the principle of maximum likelihood estimation (MLE). The theoretical analysis of these algorithms is based on the following inequality:

$$
V\left(\pi^{\mathrm{E}}\right)-V(\pi) \leq H \sum_{h=1}^{H} \mathbb{E}_{s \sim d_{h}^{\pi}}[\cdot)\left[\mathrm{TV}\left(\pi_{h}^{\mathrm{E}}(\cdot \mid s), \pi_{h}(\cdot \mid s)\right)\right]
$$

Therefore, the key is to upper bound the TV distance. Take BC as an example (i.e., $\pi=\pi^{\mathrm{BC}}$ ). By using the concentration inequality in [1, Theorem 21], we obtain that for any $\delta \in(0,1)$, when $\left|\mathcal{D}^{\mathrm{E}}\right| \geq 1$, with probability at least $1-\delta$ over the randomness within $\mathcal{D}^{\mathrm{E}}$,

$$
\begin{equation*}
\mathbb{E}_{s \sim d_{h}^{\pi^{\mathrm{E}}}(\cdot)}\left[\mathrm{TV}^{2}\left(\pi_{h}^{\mathrm{E}}(\cdot \mid s), \pi_{h}^{\mathrm{BC}}(\cdot \mid s)\right)\right] \leq 2 \frac{\log (|\Pi| / \delta)}{\left|\mathcal{D}^{\mathrm{E}}\right|} \tag{19}
\end{equation*}
$$

With additional efforts (by using union bound and Jensen's inequality), we have the following result.
Theorem 4 (BC with Function Approximation). Under Assumption 1. In the general function approximation setting, additionally assume that $\pi^{\mathrm{E}} \in \Pi$. If we apply $B C$ on the expert data, we have

$$
\mathbb{E}\left[V\left(\pi^{\mathrm{E}}\right)-V\left(\pi^{\mathrm{BC}}\right)\right]=\mathcal{O}\left(H^{2} \sqrt{\frac{\log \left(|\Pi| H N_{\mathrm{E}}\right)}{N_{\mathrm{E}}}}\right)
$$

where the expectation is taken over the randomness in the dataset collection.
The detailed proof is deferred to Appendix F. Compared with Theorem 1, we notice that the change in theoretical bound is that $\mathcal{O}\left(|\mathcal{S}| / N_{\mathrm{E}}\right)$ is replaced by $\mathcal{O}\left(\sqrt{\log \left(|\Pi| H N_{\mathrm{E}}\right) / N_{\mathrm{E}}}\right)$.
NBCU can be analyzed in a similar way in the function approximation setting.
Theorem 5 (NBCU with Function Approximation). Under Assumption 1. In the general function approximation setting, additionally assume that the realizable policy class $\Pi$ is realizable, i.e., $\pi^{\text {mix }} \in \Pi$, where $\pi^{\text {mix }}$ is defined in Eq. (10). If we apply $B C$ on the union dataset, we have

$$
\mathbb{E}\left[V\left(\pi^{\mathrm{E}}\right)-V\left(\pi^{\mathrm{NBCU}}\right)\right]=\mathcal{O}\left((1-\eta)\left(V\left(\pi^{\mathrm{E}}\right)-V\left(\pi^{\beta}\right)\right)+H^{2} \sqrt{\frac{\log \left(|\Pi| H N_{\mathrm{tot}}\right)}{N_{\mathrm{tot}}}}\right)
$$

Unfornatunetly, the analysis of ISW-BC with function approximation is much more complicated since the maximum likelihood estimation is performed in a weighted manner.In the following part, we make a conjecture on the theoretical guarantee of the weighted maximum likelihood estimation. With such a conjecture, we can derive the imitation gap of ISW-BC with general function approximation. We leave the proof of the conjecture and other proof possibilities for future works.
Recall the objective of ISW-BC.

$$
\pi^{\mathrm{ISW}-\mathrm{BC}} \in \underset{\pi \in \Pi}{\operatorname{argmax}} \sum_{h=1}^{H} \sum_{(s, a) \in \mathcal{S} \times \mathcal{A}}\left\{\widehat{d_{h}^{\mathrm{U}}}(s, a) \times\left[w_{h}(s, a) \log \pi_{h}(a \mid s)\right] \times \mathbb{I}\left[w_{h}(s, a) \geq \delta\right]\right\}
$$

Notice that the analysis of the discriminators is independent of the function approximation of policies. Therefore, we can follow the analysis of the discriminators in the proof of Theorem 3. Importantly, we can derive that there exists $\delta$ such that $\mathbb{I}\left[w_{h}\left(s, \pi_{h}^{\mathrm{E}}(s)\right) \geq \delta\right]=1$ for any $\left(s, \pi_{h}^{\mathrm{E}}(s)\right) \in \mathcal{D}_{h}^{\mathrm{E}} \cup \mathcal{D}_{h}^{\mathrm{S}, 1}$ and $\mathbb{I}\left[w_{h}(s, a) \geq \delta\right]=0$ for any $\forall(s, a) \in \mathcal{D}_{h}^{\mathrm{S}, 2}$. Then we can obtain that

$$
\begin{aligned}
& \sum_{h=1}^{H} \sum_{(s, a) \in \mathcal{S} \times \mathcal{A}}\left\{\widehat{d_{h}^{\mathrm{U}}}(s, a) \times\left[w_{h}(s, a) \log \pi_{h}(a \mid s)\right] \times \mathbb{I}\left[w_{h}(s, a) \geq \delta\right]\right\} \\
= & \sum_{h=1}^{H} \sum_{s \in \mathcal{S}_{h}^{\mathrm{E}}, a=\pi_{h}^{\mathrm{E}}(s)} \widehat{d_{h}^{\mathrm{U}}}(s, a) \times\left[w_{h}(s, a) \log \pi_{h}(a \mid s)\right] .
\end{aligned}
$$

Here $\mathcal{S}_{h}^{\mathrm{E}}=\left\{s \in \mathcal{S}: d_{h}^{\pi^{\mathrm{E}}}(s)>0\right\}$. Then we have that

$$
\pi^{\mathrm{ISW}-\mathrm{BC}} \in \underset{\pi \in \Pi}{\operatorname{argmax}} \sum_{h=1}^{H} \sum_{s \in \mathcal{S}_{h}^{\mathrm{E}}, a=\pi_{h}^{\mathrm{E}}(s)} \widehat{d_{h}^{\mathrm{U}}}(s, a) \times\left[w_{h}(s, a) \log \pi_{h}(a \mid s)\right]
$$

We conjecture that $\pi^{\mathrm{ISW}-\mathrm{BC}}$ learned by the above weighted maximum likelihood holds the following theoretical guarantee. For any $\delta \in(0,1)$, with probability at least $1-\delta$, we have that

$$
\begin{equation*}
\sum_{s \in \mathcal{S}_{h}^{\mathrm{E}}} d_{h}^{\mathrm{U}}\left(s, \pi_{h}^{\mathrm{E}}(s)\right) \mathrm{TV}^{2}\left(\pi_{h}^{\mathrm{E}}(\cdot \mid s), \pi_{h}^{\mathrm{ISW}-\mathrm{BC}}(\cdot \mid s)\right)=\mathcal{O}\left(\frac{\log \left(|\Pi| N_{\mathrm{tot}}\right)}{N_{\mathrm{tot}}}\right) \tag{20}
\end{equation*}
$$

This conjecture corresponds to (19) in the unweighted maximum likelihood estimation. With this conjecture, we can derive the imitation gap of ISW-BC with function approximation.

Conjecture 1 (Imitation Gap of ISW-BC with Function Approximation). Under Assumptions 1 and 2 , let $\mu=\max _{(s, h) \in \mathcal{S} \times[H]} d_{h}^{\pi^{\mathrm{E}}}\left(s, \pi_{h}^{\mathrm{E}}(s)\right) / d_{h}^{\pi^{\beta}}\left(s, \pi_{h}^{\mathrm{E}}(s)\right)$. In the general function approximation setting with the realizable policy class $\Pi$, i.e., $\pi^{\mathrm{E}} \in \Pi$. Furthermore, assume that the feature is designed such that $\sqrt{\frac{2\left(\mathcal{L}_{h}\left(\bar{\theta}_{h}\right)-\mathcal{L}_{h}\left(\theta_{h}^{\star}\right)\right)}{\tau_{h}}}<\frac{\Delta_{h}\left(\bar{\theta}_{h}\right)}{L_{h}}$ holds and the conjecture in (20) holds. Then, we have the imitation gap bound

$$
\mathbb{E}\left[V\left(\pi^{\mathrm{E}}\right)-V\left(\pi^{\mathrm{ISW}-\mathrm{BC}}\right)\right]=\mathcal{O}\left(H^{2} \sqrt{\frac{\log \left(|\Pi| H N_{\mathrm{tot}}\right)}{N_{\mathrm{E}}+N_{\mathrm{S}} / \mu}}\right) .
$$

## E. 2 Supplementary Data with Corruption

In the main text, we consider the trajectory sampling procedure in Assumption 1. However, in some cases, the supplementary data can be poisoned and corrupted by an adversary. For example, although the human expert demonstrates an optimal trajectory, the recorder or the recording system possibly corrupts the data by accident or on purpose. Data corruption is one of the main security threats to imitation learning methods [25]. Therefore, it is valuable to investigate the robustness of the presented algorithms in this poison setting. Supplementary data with corruption is partially investigated in our experiments under the noisy expert setting, which we argue have a large state-action distribution shift.
Assumption 3 (Poison Setting). The supplementary dataset $\mathcal{D}^{\mathrm{S}}$ and expert dataset $\mathcal{D}^{\mathrm{E}}$ are collected in the following way: each time, with probability $\eta$, we rollout the expert policy to collect a trajectory. With probability $1-\eta$, we still rollout the expert policy to collect a trajectory but with probability
$1-\eta^{\prime}$, the actions along the sampled trajectory are replaced with actions uniformly sampled from the action space. Such an experiment is independent and identically conducted by $N_{\text {tot }}$ times.

Theorem 6 (NBCU in the Poison Setting). Under Assumption 3. In the tabular case, for any $\eta \in(0,1]$, we have

$$
\mathbb{E}\left[V\left(\pi^{\mathrm{E}}\right)-V\left(\pi^{\mathrm{NBCU}}\right)\right]=\mathcal{O}\left((1-\eta)\left(1-\eta^{\prime}\right) H^{2}\left(1-\frac{1}{|\mathcal{A}|}\right)+H^{2} \sqrt{\frac{|\mathcal{S}||\mathcal{A}|}{N_{\mathrm{tot}}}}\right)
$$

where the expectation is taken over the randomness in the dataset collection.
Theorem 7 (ISW-BC in the Poison Setting). Under Assumptions 2 and 3, if the feature is designed such that $\sqrt{\frac{2\left(\mathcal{L}_{h}\left(\bar{\theta}_{h}\right)-\mathcal{L}_{h}\left(\theta_{h}^{\star}\right)\right)}{\tau_{h}}}<\frac{\Delta_{h}\left(\bar{\theta}_{h}\right)}{L_{h}}$ holds, we have the imitation gap bound

$$
\mathbb{E}\left[V\left(\pi^{\mathrm{E}}\right)-V\left(\pi^{\mathrm{ISW}-\mathrm{BC}}\right)\right]=\mathcal{O}\left(\frac{H^{2}|\mathcal{S}|}{N_{\mathrm{E}}+N_{\mathrm{S}} \eta^{\prime}}\right)
$$

Proofs of Theorem 6 and Theorem 7 can be found in Appendix F. Compared with the imitation gap of NBCU, there is no non-vanishing gap due to the corrupted actions in the imitation gap of ISW-BC. This means that ISW-BC is still robust in this setting.

## F Proof of Results in Section E

## F. 1 Proof of Theorem 4

According to [32, Lemma 4.3], we have

$$
V\left(\pi^{\mathrm{E}}\right)-V\left(\pi^{\mathrm{BC}}\right) \leq H \sum_{h=1}^{H} \mathbb{E}_{s \sim d_{h}^{\pi \mathrm{E}}(\cdot)}\left[\mathrm{TV}\left(\pi_{h}^{\mathrm{E}}(\cdot \mid s), \pi_{h}^{\mathrm{BC}}(\cdot \mid s)\right)\right]
$$

With [1, Theorem 21], when $\left|\mathcal{D}^{\mathrm{E}}\right| \geq 1$, for any $\delta \in(0,1)$, with probability at least $1-\delta$ over the randomness within $\mathcal{D}^{\mathrm{E}}$, we have that

$$
\mathbb{E}_{s \sim d_{h}^{\pi}(\cdot)}\left[\operatorname{TV}^{2}\left(\pi_{h}^{\mathrm{E}}(\cdot \mid s), \pi_{h}^{\mathrm{BC}}(\cdot \mid s)\right)\right] \leq 2 \frac{\log (|\Pi| / \delta)}{\left|\mathcal{D}^{\mathrm{E}}\right|}
$$

With union bound, with probability at least $1-\delta$, for all $h \in[H]$, it holds that

$$
\mathbb{E}_{s \sim d_{h}^{\pi \mathrm{E}}(\cdot)}\left[\mathrm{TV}^{2}\left(\pi_{h}^{\mathrm{E}}(\cdot \mid s), \pi_{h}^{\mathrm{BC}}(\cdot \mid s)\right)\right] \leq 2 \frac{\log (|\Pi| H / \delta)}{\left|\mathcal{D}^{\mathrm{E}}\right|}
$$

which implies that

$$
\begin{aligned}
V\left(\pi^{\mathrm{E}}\right)-V\left(\pi^{\mathrm{BC}}\right) & \leq H \sum_{h=1}^{H} \mathbb{E}_{s \sim d_{h}^{\pi}}(\cdot) \\
& \stackrel{\left.\mathrm{TV}\left(\pi_{h}^{\mathrm{E}}(\cdot \mid s), \pi_{h}^{\mathrm{BC}}(\cdot \mid s)\right)\right]}{\leq H \sum_{h=1}^{H} \sqrt{\mathbb{E}_{s \sim d_{h}^{\pi^{\mathrm{E}}}(\cdot)}\left[\mathrm{TV}^{2}\left(\pi_{h}^{\mathrm{E}}(\cdot \mid s), \pi_{h}^{\mathrm{BC}}(\cdot \mid s)\right)\right]}} \\
& \leq \sqrt{2} H^{2} \sqrt{\frac{\log (|\Pi| H / \delta)}{\left|\mathcal{D}^{\mathrm{E}}\right|}}
\end{aligned}
$$

Inequality $(a)$ follows Jensen's inequality. Taking expectation over the randomness within $\mathcal{D}^{\mathrm{E}}$ yields that

$$
\begin{aligned}
\mathbb{E}_{\mathcal{D}^{\mathrm{E}}}\left[V\left(\pi^{\mathrm{E}}\right)-V\left(\pi^{\mathrm{BC}}\right)\right] & \leq \delta H+(1-\delta) \sqrt{2} H^{2} \sqrt{\frac{\log (|\Pi| H / \delta)}{\left|\mathcal{D}^{\mathrm{E}}\right|}} \\
& \stackrel{(a)}{=} \frac{H}{2\left|\mathcal{D}^{\mathrm{E}}\right|}+\left(1-\frac{1}{2\left|\mathcal{D}^{\mathrm{E}}\right|}\right) \sqrt{2} H^{2} \sqrt{\frac{\log \left(2|\Pi| H\left|\mathcal{D}^{\mathrm{E}}\right|\right)}{\left|\mathcal{D}^{\mathrm{E}}\right|}}
\end{aligned}
$$

$$
\begin{aligned}
& \leq(\sqrt{2}+1) H^{2} \sqrt{\frac{\log \left(2|\Pi| H\left|\mathcal{D}^{\mathrm{E}}\right|\right)}{\left|\mathcal{D}^{\mathrm{E}}\right|}} \\
& \leq 4 H^{2} \sqrt{\frac{\log \left(4|\Pi| H\left|\mathcal{D}^{\mathrm{E}}\right|\right)}{\left|\mathcal{D}^{\mathrm{E}}\right|}}
\end{aligned}
$$ Equation (a) holds due to the choice that $\delta=1 /\left(2\left|\mathcal{D}^{\mathrm{E}}\right|\right)$. For $\left|\mathcal{D}^{\mathrm{E}}\right|=0$, we directly have that

$$
\mathbb{E}_{\mathcal{D}^{\mathrm{E}}}\left[V\left(\pi^{\mathrm{E}}\right)-V\left(\pi^{\mathrm{BC}}\right)\right] \leq H
$$

Therefore, for any $\left|\mathcal{D}^{\mathrm{E}}\right| \geq 0$, we have that

$$
\mathbb{E}_{\mathcal{D}^{\mathrm{E}}}\left[V\left(\pi^{\mathrm{E}}\right)-V\left(\pi^{\mathrm{BC}}\right)\right] \leq 4 H^{2} \sqrt{\frac{\log \left(4|\Pi| H \max \left\{\left|\mathcal{D}^{\mathrm{E}}\right|, 1\right\}\right)}{\max \left\{\left|\mathcal{D}^{\mathrm{E}}\right|, 1\right\}}}
$$

We consider a real-valued function $f(x)=\log (c x) / x$ for $x \geq 1$, where $c=4|\Pi| H>4$. Its gradient function is $f^{\prime}(x)=(1-\log (c x)) / x^{2} \leq 0$ when $x \geq 1$. Then we know that $f(x)$ is decreasing as $x$ increases. Furthermore, we have that $\max \left\{\left|\mathcal{D}^{\mathrm{E}}\right|, 1\right\} \geq\left(\left|\mathcal{D}^{\mathrm{E}}\right|+1\right) / 2$ when $\left|\mathcal{D}^{\mathrm{E}}\right| \geq 0$. Then we obtain

$$
\begin{aligned}
\mathbb{E}_{\mathcal{D}^{\mathrm{E}}}\left[V\left(\pi^{\mathrm{E}}\right)-V\left(\pi^{\mathrm{BC}}\right)\right] & \leq 4 H^{2} \sqrt{\frac{\log \left(4|\Pi| H \max \left\{\left|\mathcal{D}^{\mathrm{E}}\right|, 1\right\}\right)}{\max \left\{\left|\mathcal{D}^{\mathrm{E}}\right|, 1\right\}}} \\
& \leq 4 H^{2} \sqrt{\frac{2 \log \left(4|\Pi| H\left(\left|\mathcal{D}^{\mathrm{E}}\right|+1\right)\right)}{\left|\mathcal{D}^{\mathrm{E}}\right|+1}}
\end{aligned}
$$

Taking expectation over the random variable $\left|\mathcal{D}^{\mathrm{E}}\right| \sim \operatorname{Bin}\left(N_{\text {tot }}, \eta\right)$ yields that

$$
\begin{aligned}
\mathbb{E}\left[V\left(\pi^{\mathrm{E}}\right)-V\left(\pi^{\mathrm{BC}}\right)\right] & \leq 4 H^{2} \mathbb{E}\left[\sqrt{\frac{2 \log \left(4|\Pi| H\left(\left|\mathcal{D}^{\mathrm{E}}\right|+1\right)\right)}{\left|\mathcal{D}^{\mathrm{E}}\right|+1}}\right] \\
& \leq 4 H^{2} \sqrt{\mathbb{E}\left[\frac{2 \log \left(4|\Pi| H\left(\left|\mathcal{D}^{\mathrm{E}}\right|+1\right)\right)}{\left|\mathcal{D}^{\mathrm{E}}\right|+1}\right]}
\end{aligned}
$$

Inequality ( $a$ ) follows Jensen's inequality. We consider the function $g(x)=-x \log (x / c)$ for $x \in(0,1]$, where $c=4|\Pi| H$.

$$
g^{\prime}(x)=-(\log (x / c)+1) \geq 0, g^{\prime \prime}(x)=-\frac{1}{x} \leq 0, \quad \forall x \in(0,1]
$$

Thus, $g(x)$ is a concave function. By Jensen's inequality, we have that $\mathbb{E}[g(x)] \leq g(\mathbb{E}[x])$. Then we can derive that

$$
\begin{aligned}
\mathbb{E}\left[V\left(\pi^{\mathrm{E}}\right)-V\left(\pi^{\mathrm{BC}}\right)\right] & \leq 4 H^{2} \sqrt{\mathbb{E}\left[\frac{2 \log \left(4|\Pi| H\left(\left|\mathcal{D}^{\mathrm{E}}\right|+1\right)\right)}{\left|\mathcal{D}^{\mathrm{E}}\right|+1}\right]} \\
& =4 \sqrt{2} H^{2} \sqrt{\mathbb{E}\left[g\left(\frac{1}{\left|\mathcal{D}^{\mathrm{E}}\right|+1}\right)\right]} \\
& \leq 4 \sqrt{2} H^{2} \sqrt{g\left(\mathbb{E}\left[\frac{1}{\left|\mathcal{D}^{\mathrm{E}}\right|+1}\right]\right)} \\
& \stackrel{(a)}{\leq} 4 \sqrt{2} H^{2} \sqrt{g\left(\frac{1}{N_{\mathrm{E}}}\right)} \\
& \leq 4 \sqrt{2} H^{2} \sqrt{\frac{\log \left(4|\Pi| H N_{\mathrm{E}}\right)}{N_{\mathrm{E}}}} .
\end{aligned}
$$

In inequality $(a)$, we use the facts that $g^{\prime}(x) \geq 0$ and $\mathbb{E}\left[1 /\left(\left|\mathcal{D}^{\mathrm{E}}\right|+1\right)\right] \leq 1 / N_{\mathrm{E}}$ from Lemma 3 . We complete the proof.

## F. 2 Proof of Theorem 5

Despite the function approximation scheme, we can perform the same decomposition analysis as in the proof of Theorem 2. Therefore, we can obtain that

$$
\mathbb{E}\left[V\left(\pi^{\mathrm{E}}\right)-V\left(\pi^{\mathrm{NBCU}}\right)\right]=(1-\eta)\left(V\left(\pi^{\mathrm{E}}\right)-V\left(\pi^{\beta}\right)\right)+\mathbb{E}\left[V\left(\pi^{\mathrm{mix}}\right)-V\left(\pi^{\mathrm{NBCU}}\right)\right] .
$$

Recall that

$$
\pi^{\mathrm{NBCU}} \in \max _{\pi \in \Pi} \sum_{h=1}^{H} \sum_{(s, a) \in \mathcal{S} \times \mathcal{A}} \widehat{d_{h}^{\mathrm{U}}}(s, a) \log \pi_{h}(a \mid s)
$$

In the proof of Theorem 2, we have shown that $d_{h}^{\mathrm{U}}(s, a)=d_{h}^{\pi^{\text {mix }}}(s, a)$, meaning that the state-action distribution of the union dataset equals the state-action distribution of the policy $\pi^{\text {mix }}$. Therefore, we can regard that $\pi^{\mathrm{NBCU}}$ is obtained by performing BC on the dataset generated by $\pi^{\mathrm{mix}}$. Consequently, we can apply Theorem 4 to obtain that ${ }^{3}$

$$
\mathbb{E}\left[V\left(\pi^{\mathrm{mix}}\right)-V\left(\pi^{\mathrm{NBCU}}\right)\right] \leq 4 \sqrt{2} H^{2} \sqrt{\frac{\log \left(4|\Pi| H N_{\mathrm{tot}}\right)}{N_{\mathrm{tot}}}}
$$

Finally, we arrive at

$$
\mathbb{E}\left[V\left(\pi^{\mathrm{E}}\right)-V\left(\pi^{\mathrm{NBCU}}\right)\right]=(1-\eta)\left(V\left(\pi^{\mathrm{E}}\right)-V\left(\pi^{\beta}\right)\right)+4 \sqrt{2} H^{2} \sqrt{\frac{\log \left(4|\Pi| H N_{\mathrm{tot}}\right)}{N_{\mathrm{tot}}}},
$$

which completes the proof.

## F. 3 Proof of Theorem 6

We first analyze the data distribution in $\mathcal{D}^{\mathrm{U}}$. According to Assumption 3, we summarize the sampling procedure of trajectories in $\mathcal{D}^{\mathrm{U}}$ as follows. Each time, we rollout the expert policy to collect a trajectory. Furthermore, with the probability of $(1-\eta)\left(1-\eta^{\prime}\right)$, the actions along the sampled expert trajectory are replaced with actions uniformly sampled from the action space. Then we put this poisoned expert trajectory into $\mathcal{D}^{\mathrm{U}}$. Otherwise, with the probability of $1-(1-\eta)\left(1-\eta^{\prime}\right)$, we directly put the original expert trajectory into $\mathcal{D}^{\mathrm{U}}$. Therefore, we can formulate the marginal distribution of the state-action pairs in time step $h$ in $\mathcal{D}^{\mathrm{U}}$. For each $(s, a, h) \in \mathcal{S} \times \mathcal{A} \times[H]$,

$$
\begin{aligned}
& d_{h}^{\mathrm{U}}(s, a)=\left(1-(1-\eta)\left(1-\eta^{\prime}\right)\right) d_{h}^{\pi^{\mathrm{E}}}(s, a)+(1-\eta)\left(1-\eta^{\prime}\right) d_{h}^{\pi^{\mathrm{E}}}(s) \frac{1}{|\mathcal{A}|} \\
& d_{h}^{\mathrm{U}}(s)=\sum_{a \in \mathcal{A}} d_{h}^{\mathrm{U}}(s, a)=d_{h}^{\pi^{\mathrm{E}}}(s)
\end{aligned}
$$

Then we proceed to analyze the imitation gap. Similar to the proof of Theorem 2, according to [32, Lemma 4.3], we have

$$
V\left(\pi^{\mathrm{E}}\right)-V\left(\pi^{\mathrm{NBCU}}\right) \leq H \sum_{h=1}^{H} \mathbb{E}_{s \sim d_{h}^{\pi \mathrm{E}}(\cdot)}\left[\mathrm{TV}\left(\pi_{h}^{\mathrm{E}}(\cdot \mid s), \pi_{h}^{\mathrm{NBCU}}(\cdot \mid s)\right)\right]
$$

Again, we introduce the definition of the policy $\pi^{\text {mix }}$.

$$
\forall(s, a) \in \mathcal{S} \times \mathcal{A}, \forall h \in[H], \pi_{h}^{\text {mix }}(a \mid s)= \begin{cases}\frac{d_{h}^{\mathrm{U}}(s, a)}{d_{h}^{U}(s)} & \text { if } d_{h}^{\mathrm{U}}(s)=d_{h}^{\pi^{\mathrm{E}}}(s)>0 \\ \frac{1}{|\mathcal{A}|} & \text { otherwise }\end{cases}
$$

In particular, if $d_{h}^{U}(s)>0$, we have that

$$
\pi_{h}^{\operatorname{mix}}(a \mid s)=\frac{d_{h}^{\mathrm{U}}(s, a)}{d_{h}^{\mathrm{U}}(s)}=\left(1-(1-\eta)\left(1-\eta^{\prime}\right)\right) \pi_{h}^{\mathrm{E}}(a \mid s)+(1-\eta)\left(1-\eta^{\prime}\right) \frac{1}{|\mathcal{A}|}
$$

831 Then we decompose the imitation gap into two parts.

$$
V\left(\pi^{\mathrm{E}}\right)-V\left(\pi^{\mathrm{NBCU}}\right)
$$

[^1]\[

$$
\begin{aligned}
& \leq H \sum_{h=1}^{H} \mathbb{E}_{s \sim d_{h}^{\pi^{\mathrm{E}}}(\cdot)}\left[\operatorname{TV}\left(\pi_{h}^{\mathrm{E}}(\cdot \mid s), \pi_{h}^{\mathrm{NBCU}}(\cdot \mid s)\right)\right] \\
& \leq H \sum_{h=1}^{H} \mathbb{E}_{s \sim d_{h}^{\pi^{\mathrm{E}}}(\cdot)}\left[\operatorname{TV}\left(\pi_{h}^{\mathrm{E}}(\cdot \mid s), \pi_{h}^{\operatorname{mix}}(\cdot \mid s)\right)\right]+H \sum_{h=1}^{H} \mathbb{E}_{s \sim d_{h}^{\pi^{\mathrm{E}}}(\cdot)}\left[\operatorname{TV}\left(\pi_{h}^{\operatorname{mix}}(\cdot \mid s), \pi_{h}^{\mathrm{NBCU}}(\cdot \mid s)\right)\right] .
\end{aligned}
$$
\]

Therefore, we can derive that

$$
H \sum_{h=1}^{H} \mathbb{E}_{s \sim d_{h}^{\pi}}[\cdot)\left[\operatorname{TV}\left(\pi_{h}^{\mathrm{E}}(\cdot \mid s), \pi_{h}^{\mathrm{mix}}(\cdot \mid s)\right)\right] \leq(1-\eta)\left(1-\eta^{\prime}\right) H^{2}\left(1-\frac{1}{|\mathcal{A}|}\right)
$$

Now we analyze the second term of

$$
H \sum_{h=1}^{H} \mathbb{E}_{s \sim d_{h}^{\pi \mathrm{E}}(\cdot)}\left[\operatorname{TV}\left(\pi_{h}^{\operatorname{mix}}(\cdot \mid s), \pi_{h}^{\mathrm{NBCU}}(\cdot \mid s)\right)\right]
$$

Recall the formula of $\pi^{\mathrm{NBCU}}$.

$$
\pi_{h}^{\mathrm{NBCU}}(a \mid s)= \begin{cases}\frac{n_{h}^{\mathrm{U}}(s, a)}{n_{h}^{U}(s)} & \text { if } n_{h}^{\mathrm{U}}(s)>0 \\ \frac{1}{|\mathcal{A}|} & \text { otherwise }\end{cases}
$$

We first analyze the first term in RHS. For certain $(s, h)$ such $d_{h}^{\mathrm{U}}(s)=d_{h}^{\pi^{\mathrm{E}}}(s)>0$, we have that

$$
\begin{aligned}
\operatorname{TV}\left(\pi_{h}^{\mathrm{E}}(\cdot \mid s), \pi_{h}^{\mathrm{mix}}(\cdot \mid s)\right) & =\sum_{a \neq \pi_{h}^{\mathrm{E}}(s)} \pi_{h}^{\mathrm{mix}}(a \mid s) \\
& =\sum_{a \neq \pi_{h}^{\mathrm{E}}(s)}\left(1-(1-\eta)\left(1-\eta^{\prime}\right)\right) \pi_{h}^{\mathrm{E}}(a \mid s)+(1-\eta)\left(1-\eta^{\prime}\right) \frac{1}{|\mathcal{A}|} \\
& =(1-\eta)\left(1-\eta^{\prime}\right)\left(1-\frac{1}{|\mathcal{A}|}\right)
\end{aligned}
$$

Recall the formula of $\pi$,

Notice that $\pi^{\mathrm{NBCU}}$ is the maximum likelihood estimation of $\pi^{\text {mix }}$. According to the concentration inequality of total variation [43], for each $(s, h) \in \mathcal{S} \times[H]$, for any fixed $\delta \in(0,1)$, when $n_{h}^{U}(s)>0$, with probability at least $1-\delta$, we have

$$
\mathrm{TV}\left(\pi_{h}^{\operatorname{mix}}(\cdot \mid s), \pi_{h}^{\mathrm{NBCU}}(\cdot \mid s)\right) \leq \sqrt{\frac{|\mathcal{A}| \log (3 / \delta)}{n_{h}^{\mathrm{U}}(s)}}
$$

When $n_{h}^{\mathrm{U}}(s)=0$, we have that

$$
\operatorname{TV}\left(\pi_{h}^{\operatorname{mix}}(\cdot \mid s), \pi_{h}^{\mathrm{NBCU}}(\cdot \mid s)\right) \leq 1 \leq \sqrt{|\mathcal{A}| \log (3 / \delta)}
$$

By combining the above two inequalities, for each $(s, h) \in \mathcal{S} \times[H]$, with probability at least $1-\delta$, we have

$$
\operatorname{TV}\left(\pi_{h}^{\operatorname{mix}}(\cdot \mid s), \pi_{h}^{\mathrm{NBCU}}(\cdot \mid s)\right) \leq \sqrt{\frac{|\mathcal{A}| \log (3 / \delta)}{\max \left\{n_{h}^{\mathrm{U}}(s), 1\right\}}}
$$

842
Applying union bound yields that with probability at least $1-\delta / 2$, for all $(s, h) \in \mathcal{S} \times[H]$,

$$
\operatorname{TV}\left(\pi_{h}^{\text {mix }}(\cdot \mid s), \pi_{h}^{\mathrm{NBCU}}(\cdot \mid s)\right) \leq \sqrt{\frac{|\mathcal{A}| \log (6|\mathcal{S}| H / \delta)}{\max \left\{n_{h}^{\mathrm{U}}(s), 1\right\}}}
$$

843 Then we have that

$$
\left.\left.\begin{array}{rl} 
& H \sum_{h=1}^{H} \mathbb{E}_{s \sim d_{h}^{\pi^{\mathrm{E}}}(\cdot)}\left[\mathrm{TV}\left(\pi_{h}^{\operatorname{mix}}(\cdot \mid s), \pi_{h}^{\mathrm{NBCU}}(\cdot \mid s)\right)\right] \\
\leq & H \sum_{h=1}^{H} \mathbb{E}_{s \sim d_{h}^{\pi^{\mathrm{E}}}(\cdot)}\left[\sqrt{\frac{|\mathcal{A}| \log (6|\mathcal{S}| H / \delta)}{\max \left\{n_{h}^{\mathrm{U}}(s), 1\right\}}}\right] \\
= & H \sqrt{|\mathcal{A}| \log (6|\mathcal{S}| H / \delta)} \sum_{h=1}^{H} \mathbb{E}_{s \sim d_{h}^{\pi \mathrm{E}}}[\cdot)
\end{array}\right] \sqrt{\frac{1}{\max \left\{n_{h}^{\mathrm{U}}(s), 1\right\}}}\right]
$$

$$
\begin{aligned}
& =H \sqrt{|\mathcal{A}| \log (6|\mathcal{S}| H / \delta)} \sum_{h=1}^{H} \sum_{s \in \mathcal{S}} \sqrt{d_{h}^{\pi^{\mathrm{E}}}(s)} \sqrt{\frac{d_{h}^{\pi^{\mathrm{E}}}(s)}{\max \left\{n_{h}^{\mathrm{U}}(s), 1\right\}}} \\
& \leq H \sqrt{|\mathcal{A}| \log (6|\mathcal{S}| H / \delta)} \sum_{h=1}^{H} \sqrt{\sum_{s \in \mathcal{S}} \frac{d_{h}^{\pi^{\mathrm{E}}}(s)}{\max \left\{n_{h}^{\mathrm{U}}(s), 1\right\}}}
\end{aligned}
$$

Here the last inequality follows Cauchy-Swartz inequality. Notice that $n_{h}^{U}(s)$ is the number of times that the state $s$ appears in $\mathcal{D}^{\mathrm{U}}$ in time step $h$ and thus follows the Binomial distribution of $\operatorname{Bin}\left(N_{\text {tot }}, d_{h}^{\pi^{\mathrm{E}}}(s)\right)$. By applying Lemma 4, for each $(s, h)$, with probability at least $1-\delta$, we have

$$
\frac{d_{h}^{\pi^{\mathrm{E}}}(s)}{\max \left\{n_{h}^{\mathrm{U}}(s), 1\right\}} \leq \frac{8 \log (1 / \delta)}{N_{\mathrm{tot}}}
$$

By union bound, with probability at least $1-\delta / 2$, for all $(s, h) \in \mathcal{S} \times[H]$,

$$
\frac{d_{h}^{\pi^{\mathrm{E}}}(s)}{\max \left\{n_{h}^{\mathrm{U}}(s), 1\right\}} \leq \frac{8 \log (2|\mathcal{S}| H / \delta)}{N_{\text {tot }}}
$$

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

$$
H \sum_{h=1}^{H} \mathbb{E}_{s \sim d_{h}^{\pi \mathrm{E}}}(\cdot)\left[\mathrm{TV}\left(\pi_{h}^{\operatorname{mix}}(\cdot \mid s), \pi_{h}^{\mathrm{NBCU}}(\cdot \mid s)\right)\right] \leq H^{2} \sqrt{\frac{8|\mathcal{S} \| \mathcal{A}| \log ^{2}(6|\mathcal{S}| H / \delta)}{N_{\text {tot }}}}
$$

Finally, we upper bound the imitation gap. With probability at least $1-\delta$, we have

$$
V\left(\pi^{\mathrm{E}}\right)-V\left(\pi^{\mathrm{NBCU}}\right) \leq(1-\eta)\left(1-\eta^{\prime}\right)\left(1-\frac{1}{|\mathcal{A}|}\right)+H^{2} \sqrt{\frac{8|\mathcal{S}||\mathcal{A}| \log ^{2}(6|\mathcal{S}| H / \delta)}{N_{\mathrm{tot}}}}
$$

We set $\delta=H / N_{\text {tot }}$ and obtain that

$$
\begin{aligned}
& \mathbb{E}\left[V\left(\pi^{\mathrm{E}}\right)-V\left(\pi^{\mathrm{NBCU}}\right)\right] \\
\leq & \delta H+(1-\delta)\left((1-\eta)\left(1-\eta^{\prime}\right)\left(1-\frac{1}{|\mathcal{A}|}\right)+H^{2} \sqrt{\frac{8|\mathcal{S}||\mathcal{A}| \log ^{2}(6|\mathcal{S}| H / \delta)}{N_{\mathrm{tot}}}}\right) \\
\leq & \frac{H^{2}}{N_{\text {tot }}}+(1-\eta)\left(1-\eta^{\prime}\right)\left(1-\frac{1}{|\mathcal{A}|}\right)+H^{2} \sqrt{\frac{8|\mathcal{S}||\mathcal{A}| \log ^{2}\left(6|\mathcal{A}| N_{\mathrm{tot}}\right)}{N_{\mathrm{tot}}}} \\
\leq & (1-\eta)\left(1-\eta^{\prime}\right)\left(1-\frac{1}{|\mathcal{A}|}\right)+4 H^{2} \sqrt{\frac{2|\mathcal{S}||\mathcal{A}| \log ^{2}\left(6|\mathcal{A}| N_{\mathrm{tot}}\right)}{N_{\mathrm{tot}}}}
\end{aligned}
$$

On the other hand, we directly have $\mathbb{E}\left[V\left(\pi^{\mathrm{E}}\right)-V\left(\pi^{\mathrm{NBCU}}\right)\right] \leq H$. We complete the proof.

## F. 4 Proof of Theorem 7

In the poison setting, we can conduct the same analysis as in the proof of Theorem 3 and demonstrate that $\pi^{\mathrm{ISW}-\mathrm{BC}}\left(\pi_{h}^{\mathrm{E}}(s) \mid s\right)=1, \forall\left(s, \pi_{h}^{\mathrm{E}}(s)\right) \in \mathcal{D}_{h}^{\mathrm{E}} \cup \mathcal{D}_{h}^{\mathrm{S}, 1}$, where $\mathcal{D}_{h}^{\mathrm{E}}$ is the set of state-action pairs in $\mathcal{D}^{\mathrm{E}}$ in time step $h$ and $\mathcal{D}_{h}^{\mathrm{S}, 1}=\left\{(s, a) \in \mathcal{D}_{h}^{\mathrm{S}}: d_{h}^{\pi^{\mathrm{E}}}(s)>0, a=\pi_{h}^{\mathrm{E}}(s)\right\}$. According to [32, Lemma 4.3], we have

$$
V\left(\pi^{\mathrm{E}}\right)-V\left(\pi^{\mathrm{ISW}-\mathrm{BC}}\right) \leq H \sum_{h=1}^{H} \mathbb{E}_{s \sim d_{h}^{\mathrm{E}^{\mathrm{E}}}(\cdot)}\left[\mathrm{TV}\left(\pi_{h}^{\mathrm{E}}(\cdot \mid s), \pi_{h}^{\mathrm{ISW}-\mathrm{BC}}(\cdot \mid s)\right)\right]
$$

857
Since the expert policy is assumed to be deterministic, we can obtain

$$
\begin{aligned}
& V\left(\pi^{\mathrm{E}}\right)-V\left(\pi^{\mathrm{ISW}-\mathrm{BC}}\right) \leq H \sum_{h=1}^{H} \mathbb{E}_{s \sim d_{h}^{\pi}}(\cdot) \\
&\left.\leq H \mathbb{E}_{a \sim \pi_{h}^{\mathrm{ISW}-\mathrm{BC}}(\cdot \mid s)}\left[\mathbb{I}\left\{a \neq \pi_{h=1}^{\mathrm{E}}(s)\right\}\right]\right] \\
& \mathbb{E}_{s \sim d_{h}^{\pi^{\mathrm{E}}}(\cdot)}\left[\mathbb{I}\left\{\left(s, \pi_{h}^{\mathrm{E}}(s)\right) \notin \mathcal{D}_{h}^{\mathrm{E}} \cup \mathcal{D}_{h}^{\mathrm{S}, 1}\right\}\right]
\end{aligned}
$$



Let $\mathcal{D}^{\text {S,clean }}$ denote the non-corrupted dataset in $\mathcal{D}^{\text {S }}$. Then we can obtain that

$$
\begin{aligned}
V\left(\pi^{\mathrm{E}}\right)-V\left(\pi^{\mathrm{ISW}-\mathrm{BC}}\right) & \stackrel{(a)}{\leq} H \sum_{h=1}^{H} \mathbb{E}_{s \sim d_{h}^{\pi^{\mathrm{E}}}(\cdot)}\left[\mathbb{I}\left\{\left(s, \pi_{h}^{\mathrm{E}}(s)\right) \notin \mathcal{D}_{h}^{\mathrm{E}} \cup \mathcal{D}_{h}^{\mathrm{S}, \text { clean }}\right\}\right] \\
& =H \sum_{h=1}^{H} \sum_{s \in \mathcal{S}} d_{h}^{\pi^{\mathrm{E}}}(s) \mathbb{I}\left\{\left(s, \pi_{h}^{\mathrm{E}}(s)\right) \notin \mathcal{D}_{h}^{\mathrm{E}} \cup \mathcal{D}_{h}^{\mathrm{S}, \text { clean }}\right\}
\end{aligned}
$$ randomness in $\mathcal{D}^{\mathrm{E}}$ and $\mathcal{D}^{\mathrm{S}, \text { clean }}$ on both sides yields that

$$
\mathbb{E}_{\mathcal{D}^{\mathrm{E}}, \mathcal{D}^{\mathrm{S}, \text { clean }}}\left[V\left(\pi^{\mathrm{E}}\right)-V\left(\pi^{\mathrm{ISW}-\mathrm{BC}}\right)\right] \leq H \sum_{h=1}^{H} \sum_{s \in \mathcal{S}} d_{h}^{\pi^{\mathrm{E}}}(s) \mathbb{P}\left(\left(s, \pi_{h}^{\mathrm{E}}(s)\right) \notin \mathcal{D}_{h}^{\mathrm{E}} \cup \mathcal{D}_{h}^{\mathrm{S}, \text { clean }}\right)
$$

Notice that both $\mathcal{D}^{\mathrm{E}}$ and $\mathcal{D}^{\mathrm{S}, \text { clean }}$ are collected by the expert policy. Then if $\left|\mathcal{D}^{\mathrm{E}}\right|+\left|\mathcal{D}^{\mathrm{S}, \text { clean }}\right| \geq 1$, we can calculate that for each $(s, h) \in \mathcal{S} \times[H]$,

$$
\begin{aligned}
d_{h}^{\pi^{\mathrm{E}}}(s) \mathbb{P}\left(\left(s, \pi_{h}^{\mathrm{E}}(s)\right) \notin \mathcal{D}_{h}^{\mathrm{E}} \cup \mathcal{D}_{h}^{\mathrm{S}, \text { clean }}\right) & =d_{h}^{\pi^{\mathrm{E}}}(s)\left(1-d_{h}^{\pi^{\mathrm{E}}}(s)\right)^{\left|\mathcal{D}^{\mathrm{E}}\right|+\left|\mathcal{D}^{\mathrm{S}, \text { clean }}\right|} \\
& \leq \frac{4}{9\left(\left|\mathcal{D}^{\mathrm{E}}\right|+\mid \mathcal{D}^{\mathrm{S}, \text { clean } \mid)}\right.}
\end{aligned}
$$

where the last inequality follows Lemma 5. If $\left|\mathcal{D}^{\mathrm{E}}\right|+\left|\mathcal{D}^{\text {S,clean }}\right|=0$, we directly have that

$$
d_{h}^{\pi^{\mathrm{E}}}(s) \mathbb{P}\left(\left(s, \pi_{h}^{\mathrm{E}}(s)\right) \notin \mathcal{D}_{h}^{\mathrm{E}} \cup \mathcal{D}_{h}^{\mathrm{S}, \text { clean }}\right) \leq 1=\frac{1}{\max \left\{\left|\mathcal{D}^{\mathrm{E}}\right|+\mid \mathcal{D}^{\mathrm{S}, \text { clean } \mid, 1\}}\right.}
$$

We unify the above two inequalities and get that

$$
d_{h}^{\pi^{\mathrm{E}}}(s) \mathbb{P}\left(\left(s, \pi_{h}^{\mathrm{E}}(s)\right) \notin \mathcal{D}_{h}^{\mathrm{E}} \cup \mathcal{D}_{h}^{\mathrm{S}, \text { clean }}\right) \leq \frac{1}{\max \left\{\left|\mathcal{D}^{\mathrm{E}}\right|+\mid \mathcal{D}^{\mathrm{S}, \text { clean } \mid, 1\}}\right.}
$$

Now we proceed to upper bound the imitation gap.

$$
\begin{aligned}
\mathbb{E}_{\mathcal{D}^{\mathrm{E}}, \mathcal{D}^{\mathrm{S}, \text { clean }}}\left[V\left(\pi^{\mathrm{E}}\right)-V\left(\pi^{\mathrm{ISW}-\mathrm{BC}}\right)\right] & \leq H \sum_{h=1}^{H} \sum_{s \in \mathcal{S}} d_{h}^{\pi^{\mathrm{E}}}(s) \mathbb{P}\left(\left(s, \pi_{h}^{\mathrm{E}}(s)\right) \notin \mathcal{D}_{h}^{\mathrm{E}} \cup \mathcal{D}_{h}^{\mathrm{S}, \text { clean }}\right) \\
& \leq \frac{|\mathcal{S}| H^{2}}{\max \left\{\left|\mathcal{D}^{\mathrm{E}}\right|+\left|\mathcal{D}^{\mathrm{S}, \text { clean }}\right|, 1\right\}}
\end{aligned}
$$

Note that $\left|\mathcal{D}^{\mathrm{E}}\right|+\left|\mathcal{D}^{\text {S,clean }}\right| \sim \operatorname{Bin}\left(N_{\text {tot }}, \eta+(1-\eta) \eta^{\prime}\right)$. Taking expectation with respect to $\left|\mathcal{D}^{\mathrm{E}}\right|+\left|\mathcal{D}^{\mathrm{S}, \text { clean }}\right|$ yields that

$$
\begin{aligned}
\mathbb{E}\left[V\left(\pi^{\mathrm{E}}\right)-V\left(\pi^{\mathrm{ISW}-\mathrm{BC}}\right)\right] & \leq \mathbb{E}\left[\frac{|\mathcal{S}| H^{2}}{\max \left\{\left|\mathcal{D}^{\mathrm{E}}\right|+\left|\mathcal{D}^{\mathrm{S}, \text { clean }}\right|, 1\right\}}\right] \\
& \leq \mathbb{E}\left[\frac{2|\mathcal{S}| H^{2}}{\left|\mathcal{D}^{\mathrm{E}}\right|+\mid \mathcal{D}^{\mathrm{S}, \text { clean } \mid+1}}\right] \\
& \stackrel{(a)}{\leq} \frac{2|\mathcal{S}| H^{2}}{N_{\text {tot }}\left(\eta+(1-\eta) \eta^{\prime}\right)} \\
& =\frac{2|\mathcal{S}| H^{2}}{N_{\mathrm{E}}+\eta^{\prime} N_{\mathrm{S}}}
\end{aligned}
$$

Inequality $(a)$ follows Lemma 3. We finish the proof.

## G Technical Lemmas

Lemma 3. For any $N \in \mathbb{N}_{+}$and $p \in(0,1)$, if the random variable $X$ follows the binomial distribution, i.e., $X \sim \operatorname{Bin}(N, p)$, then we have that

$$
\mathbb{E}\left[\frac{1}{X+1}\right] \leq \frac{1}{N p}
$$

Proof.

$$
\begin{aligned}
\mathbb{E}\left[\frac{1}{X+1}\right] & =\sum_{x=0}^{N}\left(\frac{1}{x+1}\right) \frac{N!}{x!(N-x)!} p^{x}(1-p)^{N-x} \\
& =\frac{1}{(N+1) p} \sum_{x=1}^{N+1}\left(\frac{(N+1)!}{x!(N+1-x)!}\right) p^{x}(1-p)^{N+1-x} \\
& =\frac{1}{(N+1) p}\left(1-(1-p)^{N+1}\right) \leq \frac{1}{N p}
\end{aligned}
$$




Figure 5: Training curves of online SAC on 4 locomotion control environments.
In our experimental setup, we utilize an expert dataset comprising of 1 expert trajectory collected by the trained SAC agent. Additionally, all algorithms are provided with a supplementary dataset. There are two setting of the supplementary data.

- Full Replay. The supplementary dataset is obtained from the replay buffer of the online SAC agent, which has over one million samples, equivalent to $1000+$ trajectories. The rapid convergence of online SAC, as illustrated in Figure 5, implies that the replay buffer is enriched with a substantial number of expert-level trajectories. As a result, we expect that utilizing the supplementary data without any modification may lead to desirable results.
- Noisy Expert. The supplementary dataset comprises of 10 clean expert trajectories and 5 noisy expert trajectories. In this case, we replace the action labels in the noisy trajectories with random actions drawn from $[-1,1]$. This replacement creates noisy action labels for the expert states, leading to a significant distribution shift at the state-action level, as noted in Remark 1. The high degree of distribution shift makes it challenging for using the supplementary data.

We use a 2-hidden-layer multi-layer perceptron (MLP) with hidden size 256 and ReLU activation for all algorithms, as the state information in locomotion control tasks is informative by design. The codebase of DemoDICE is based on the original authors' work, which can be accessed at https://github.com/KAIST-AILab/imitation-dice. For DWBC, we also use the authors' codebase, which is available at https://github.com/ryanxhr/DWBC. We experimented with different hyper-parameters for both algorithms but found that the default parameters provided by the authors work well. We normalize state observations in the dataset before training all algorithms, following [22]. This is crucial for achieving satisfactory performance.

In training the discriminator of ISW-BC, we use the gradient penalty (GP) regularization, as recommended by [22]. We add the following loss to the original loss (4) to enforce 1-Lipschitz continuity:

$$
\min _{\theta} \sum_{(s, a) \in \mathcal{B}}(\|g(s, a ; \theta)\|-1)^{2},
$$

where $g$ is the gradient of the discriminator $c(s, a ; \theta)$, and $\mathcal{B}$ is a mini-batch. This promotes the learning of smooth features and can improve generalization performance.

In our implementation of ISW-BC, we employ 2-hidden-layer MLPs with 256 hidden units and ReLU activation for both the discriminator and policy networks. We use a batch size of 256 and Adam optimizer with a learning rate of 0.0003 for training both networks. The training objective is to maximize the log-likelihood. We set $\delta$ to 0 and use a gradient penalty coefficient of 8 by default, unless otherwise stated. The training process is carried out for 1 million iterations. We evaluate the performance every 10k iterations with 10 episodes. The normalized score in the last column of Table 2 is computed in the following way:

$$
\begin{equation*}
\text { Normalized score }=\frac{\text { Expert performance }- \text { Agent performance }}{\text { Expert performance }- \text { Random policy performance }} . \tag{21}
\end{equation*}
$$

## H.1.2 Atari Games

We evaluate algorithms on a set of 5 Atari games from the standard benchmark: Alien, MsPacman, Phoenix, Qbert, and SpaceInvaders. We preprocess the game environments using a standard set of procedures, including sticky actions with a probability of 0.25 , grayscaling, downsampling to an image size of [84, 84], and stacking frames of 4 . These procedures follow the instructions provided by the dopamine codebase, which is available at https://github.com/google/dopamine/blob/ master/dopamine/discrete_domains/atari_lib.py. The final image inputs are of shape (84, 84, 4).

We use the replay buffer data from an online DQN agent, which is publicly available at https: //console.cloud.google.com/storage/browser/atari-replay-datasets, thanks to the work of [2]. The dataset consists of 200 million frames, divided into 50 indexed buckets (ranging from 0 to 49). However, using the entire dataset is computationally infeasible ${ }^{4}$ and unnecessary for our task. Therefore, we select specific buffer buckets for imitation learning.

We choose the expert data from bucket index 49, using only the first 400 K frames for training. This makes the task challenging (we find that BC performs well with 1 M frames of expert data). For the full replay setting, we select supplementary data from buffer indices 45 to 48 , using the first 400 K frames from each bucket. This yields a supplementary dataset that is 4 times larger than the expert

[^2]data. In the noisy task setting, we follow the same procedure for selecting supplementary data, but replace the action labels with random labels on buffer index 45 .

All agents employ the same convolutional neural network (CNN) architecture as the DQN agent, consisting of three convolutional blocks. The first block applies a filter size of 8, a stride of 4, and has a channel size of 32 . The second block uses a filter size of 4 , a stride of 4 , and a channel size of 64 , while the third block applies a filter size of 3, a stride of 4 , and has a channel size of 64 . All blocks use the ReLU activation function. The feature representations are flattened to a vector, on which a 1-hidden-layer MLP with a hidden size of 512 and ReLU activation function is applied. Finally, the outputs are passed through a softmax function to obtain a probability distribution.

Atari games are not considered in [22,47] and public implementations of DemoDICE and DWBC for Atari games are not available. To use these methods in the Atari environment, we extend their original implementation by replacing the MLP used in locomotion control with the CNN described earlier. Implementing ISW-BC is a little more complicated. We use the same CNN policy network as in the other methods, but find that directly training the discriminator from scratch is less effective. This is because the discriminator tends to focus on irrelevant background information instead of the decision-centric part. To overcome this issue, we build the discriminator upon the feature extractor of the policy network, leveraging its ability to extract useful information. The discriminator is an MLP with ReLU activation and a hidden size of 1024: the image feature representation has a dimension 512 and the action feature representation also has a dimension 512 (we randomly project one-hot discrete actions to a 512-dimension space). We find that the depth of the MLP is crucial for performance, using a depth of 1 for the full replay setting and 3 for the noisy expert setting. We clip the importance sampling ratio for numerical stability, using a minimum value of 0 and a maximum value of 5 for the full replay setting, and a minimum value of 0.2 and a maximum value of 5 for the noisy expert setting. We provide ablation studies of these hyperparameters in Appendix H.2.2.
All methods were optimized using the Adam optimizer with a learning rate of 0.00025 and a batch size of 256 . The training objective is to maximize the log-likelihood. The training process consisted of 200 K gradient steps. Every 2 K gradient steps, the algorithms were evaluated by running 10 episodes and computing the raw game scores. The normalized score in the last column of Table 3 is computed by Eq. (21).

## H.1.3 Object Recognition

We utilize the publicly available DomainNet dataset [29] for our experiments, which can be accessed at http://csr.bu.edu/ftp/visda/2019/multi-source. This dataset comprises six sub-datasets: clipart, infograph, painting, quickdraw, real, and sketch, with 2103, 2626, 2472, 4000, 4864, and 2213 images, respectively. Our task involves recognizing objects from 10 different classes: bird, feather, headphones, ice_cream, teapot, tiger, whale, windmill, wine_glass, and zebra. We divided the images into training and test sets, with $80 \%$ for training and $20 \%$ for testing.
We employ a 2-hidden-layer neural network with a hidden size of 512 and ReLU activation as the classifier. To extract features from images, we utilize the pretrained ResNet-18 model (trained on ImageNet), which has a feature dimension of 512. The ResNet-18 model can be accessed at https://pytorch.org/vision/main/models/generated/torchvision.models. resnet18.html. We opted for this approach as training such a large convolutional neural network directly on the DomainNet dataset proved to be ineffective. The training objective is to minimize the cross-entropy loss. To optimize the network parameters, we use the stochastic gradient descent (SGD) optimizer with a learning rate of 0.01 and momentum of 0.9 . Additionally, we apply weight decay with a coefficient of 0.0005 . The models are trained for 100 epochs with a batch size of 100 , following the standard practice.

The discriminators used in ISW-BC and DWBC are implemented as 2-hidden-layer neural networks with ReLU activation. It's important to note that these discriminators take both the image and label as inputs. The image input is processed by the pre-trained and fixed ResNet-18, while the label input is projected to the same dimension (512) by a random projection matrix. The hidden size for the discriminator is set to 1024 for ISW-BC and 1025 for DWBC, as the discriminator in DWBC also takes the log-likelihood as an input. For ISW-BC, the discriminator is trained independently for 100 epochs with the same optimization configuration as the classifier. Afterward, the discriminator is fixed, and its output is used to compute the importance sampling ratio, which is then used to train the classifier.

## H. 2 Additional Results

## H.2.1 Training Curves

Training curves. The training curves on the MuJoco locomotion control tasks are displayed in Figure 6 and Figure 7. The training curves on Atari games are displayed in Figure 8 and Figure 9. The training curves on the object recognition task are displayed in Figure 10.


Figure 6: Training curves of algorithms on the locomotion control task in the full replay setting. Solid lines correspond to the mean performance and shaded regions correspond to the $95 \%$ confidence interval. Same as other figures.


Figure 7: Training curves of algorithms on the locomotion control task in the noisy expert setting.

## H.2.2 Ablation Study

In this section, we present ablation studies conducted on Atari games, aiming to provide insights into the underlying working scheme of our method. We specifically emphasize Atari games due to their high-dimensional image inputs, making these tasks particularly challenging. In contrast, the other


Figure 8: Training curves of algorithms on the Atari games in the full replay setting.


Figure 9: Training curves of algorithms on the Atari games in the noisy expert setting.


Figure 10: Training curves of algorithms on the object recognition task using the DomainNet dataset.

Figure 11: Training curves of ISW-BC on the Atari games in the full replay setting. We test the performance with different feature extractors of the discriminator.
two tasks, locomotion control and object recognition, involve informative vector inputs, setting them apart from the unique characteristics of Atari games.

Ablation Study on Feature Representations of Discriminator Network. Our study reveals that employing a separate CNN for the discriminator yields inferior results compared to utilizing the feature extractor of the policy network. Please refer to Figure 11. Our conjecture is that training the discriminator independently may cause it to fit noise information (e.g., background). In contrast, the policy CNN network is capable of learning decision-centric information, enabling an effective approach to building the discriminator network through the feature extractor of the policy network.


Ablation Study on Depth of Discriminator Network. We have discovered that the number of discriminator layers plays a crucial role in the performance of Atari games. The training curves, depicted in both Figure 12 and Figure 13, illustrate the performance variation based on the number of layers in the discriminator network. Notably, a 1-hidden-layer neural network yields the best results for the full replay setting, while a 3-hidden-layer neural network performs optimally in the noisy expert setting. It is important to note that this phenomenon is specific to Atari games. We do not have a good explanation yet. We believe this deserves further investigation in the future work.


Figure 12: Training curves of ISW-BC on the Atari games in the full replay setting. We test the performance with different number of layers for the discriminator network.


Figure 13: Training curves of ISW-BC on the Atari games with the noisy expert setting. We test the performance with different number of layers for the discriminator network.


[^0]:    ${ }^{2}$ For a moment, we use the notations in [22] and present their results under the stationary and infinite-horizon MDPs. Same as the discussion of DWBC [47].

[^1]:    ${ }^{3}$ Note that Theorem 4 holds for the case where the expert policy is stochastic.

[^2]:    ${ }^{4}$ Loading 200 M frames requires over 500 GB memory.

