## Checklist

1. For all authors...
(a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? Yes.
(b) Did you describe the limitations of your work? Yes.
(c) Did you discuss any potential negative societal impacts of your work? No, this paper is purely theoretical.
(d) Have you read the ethics review guidelines and ensured that your paper conforms to them? Yes.
2. If you are including theoretical results...
(a) Did you state the full set of assumptions of all theoretical results? Yes, they are presented in Section 2.1.
(b) Did you include complete proofs of all theoretical results? Yes, they are in the appendix.
3. If you ran experiments...
(a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? N/A.
(b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? N/A.
(c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? N/A.
(d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? N/A.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
(a) If your work uses existing assets, did you cite the creators? N/A.
(b) Did you mention the license of the assets? N/A.
(c) Did you include any new assets either in the supplemental material or as a URL? N/A.
(d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? N/A.
(e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? N/A.
5. If you used crowdsourcing or conducted research with human subjects...
(a) Did you include the full text of instructions given to participants and screenshots, if applicable? N/A.
(b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? N/A.
(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? N/A.
```
Algorithm \(2 \operatorname{OLIVE}\left(\mathcal{F}, \zeta_{\text {act }}, \zeta_{\text {elim }}, n_{\text {act }}, n_{\text {elim }}\right)\)
    Initialize: \(\mathcal{B}^{0} \leftarrow \mathcal{F}, \mathcal{D}_{h} \leftarrow \emptyset\) for all \(h, k\).
    for phase \(k=1,2, \ldots\) do
        Choose policy \(\pi^{k}=\pi_{f^{k}}\), where \(f^{k}=\operatorname{argmax}_{f \in \mathcal{B}^{k-1}} f\left(s_{1}, \pi_{f}\left(s_{1}\right)\right)\).
        Execute \(\pi^{k}\) for \(n_{\text {act }}\) episodes and refresh \(\mathcal{D}_{h}\) to include the fresh \(\left(s_{h}, a_{h}, r_{h}, s_{h+1}\right)\) tuples.
        Estimate \(\hat{\mathcal{E}}\left(f^{k}, \pi^{k}, h\right)\) for all \(h \in[H]\), where
            \(\hat{\mathcal{E}}\left(g, \pi^{k}, h\right)=\frac{1}{\left|\mathcal{D}_{h}\right|} \sum_{\left(s, a, r, s^{\prime}\right) \in \mathcal{D}_{h}}\left(g_{h}(s, a)-r-\max _{a^{\prime} \in \mathcal{A}} g_{h+1}\left(s^{\prime}, a^{\prime}\right)\right)\).
        if \(\sum_{h=1}^{H} \hat{\mathcal{E}}\left(f^{k}, \pi^{k}, h\right) \leq H \zeta_{\text {act }}\) then
            Terminate and output \(\pi^{k}\).
        Pick any \(t \in[H]\) for which \(\hat{\mathcal{E}}\left(f^{k}, \pi^{k}, t\right) \geq \zeta_{\text {act }}\).
        Execute \(\pi^{k}\) for \(n_{\text {elim }}\) episodes and refresh \(\mathcal{D}_{h}\) to include the fresh \(\left(s_{h}, a_{h}, r_{h}, s_{h+1}\right)\) tuples.
        Estimate \(\hat{\mathcal{E}}\left(f, \pi^{k}, t\right)\) for all \(f \in \mathcal{F}\).
        Update \(\mathcal{B}^{k}=\left\{f \in \mathcal{B}^{k-1}:\left|\hat{\mathcal{E}}\left(f, \pi^{k}, t\right)\right| \leq \zeta_{\text {elim }}\right\}\).
```


## A Algorithm Olive

In this section, we analyze algorithm Olive proposed in [1], which is based on hypothesis elimination. We prove that, despite Olive was originally designed for solving low Bellman rank problems, it naturally learns RL problems with low BE dimension as well.

The main advantage of Olive comparing to Golf is that Olive does not require the completeness assumption. In return, Olive has several disadvantages including worse sample complexity, and no sublinear regret.
The pseudocode of Olive is presented in Algorithm 2, where in each phase the algorithm contains the following three main components:

- Line 3 (Optimistic planning): compute the most optimistic value function $f^{k}$ from the candidate set $\mathcal{B}^{k-1}$, and choose $\pi^{k}$ to be its greedy policy.
- Line 4-7 (Estimate Bellman error): estimate the Bellman error of $f^{k}$ under $\pi^{k}$; output $\pi^{k}$ if the estimated error is small, and otherwise activate the elimination procedure.
- Line 8-11 (Eliminate functions with large Bellman error): pick a step $t \in[H]$ where the estimated Bellman error exceeds the activation threshold $\zeta_{\text {act }}$; eliminate all functions in the candidate set whose Bellman error at step $t$ exceeds the elimination threshold $\zeta_{\text {elim }}$.

We comment that Olive is computationally inefficient in general because implementing the optimistic planning part requires solving an NP-hard problem in the worst case [Theorem 4, 48].

## A. 1 Theoretical guarantees

Now, we are ready to present the theoretical guarantee for Olive.
Theorem 18 (OLIVE). Under Assumption 1, there exists absolute constant c such that if we choose

$$
\zeta_{\text {act }}=\frac{2 \epsilon}{H}, \zeta_{\text {elim }}=\frac{\epsilon}{2 H \sqrt{d}}, n_{\text {act }}=\frac{H^{2} \iota}{\epsilon^{2}}, \text { and } n_{\text {elim }}=\frac{H^{2} d \log \left(\mathcal{N}_{\mathcal{F}}\left(\zeta_{\text {elim }} / 8\right)\right) \cdot \iota}{\epsilon^{2}}
$$

where $d=\operatorname{dim}_{\mathrm{BE}}\left(\mathcal{F}, \mathcal{D}_{\mathcal{F}}, \epsilon / H\right)$ and $\iota=c \log (H d / \delta \epsilon)$, then with probability at least $1-\delta$, Algorithm 2 will output an $\mathcal{O}(\epsilon)$-optimal policy using at most $\mathcal{O}\left(H^{3} d^{2} \log \left[\mathcal{N}_{\mathcal{F}}\left(\zeta_{\text {elim }} / 8\right)\right] \cdot \iota / \epsilon^{2}\right)$ episodes.

Theorem 18 claims that Olive learns an $\epsilon$-optimal policy of an MDP with BE dimension $d$ within $\tilde{\mathcal{O}}\left(H^{3} d^{2} \log \left(\mathcal{N}_{\mathcal{F}}\right) / \epsilon^{2}\right)$ episodes. When specialized to low Bellman rank problems, our sample complexity has the same quadratic dependence on Bellman rank $d$ as in [1].
Comparing to Golf, the major advantage of Olive is that Olive does not require completeness assumption (Assumption 2) to work. Nevertheless, Olive only learns the RL problems that have
low BE dimension with respect to distribution family $\mathcal{D}_{\mathcal{F}}$, not $\mathcal{D}_{\Delta}$. The sample complexity of OLIVE is also worse than the sample complexity Golf (as presented in Corollary 16).
Finally, we comment that interpreting OLIVE through the lens of BE dimension, makes the proof of Theorem 18 surprisingly natural, which follows from the definition of BE dimension along with some standard concentration arguments.

## A. 2 Interpret Olive with BE dimension

In this subsection, we explain the key idea behind OlIVE through the lens of BE dimension.
To provide a clean high-level view, let us assume all estimates are accurate for now, and the activation threshold $\zeta_{\text {act }}$ and the elimination threshold $\zeta_{\text {elim }}$ satisfy $\zeta_{\text {elim }} \sqrt{d} \leq \zeta_{\text {act }}$, where $d=$ $\operatorname{dim}_{\text {BE }}\left(\mathcal{F}, \mathcal{D}_{\mathcal{F}}, \zeta_{\text {act }}\right)$. Since $\mathcal{E}\left(Q^{\star}, \pi, h\right) \equiv 0$ for any $(\pi, h), Q^{\star}$ is always in the candidate set. Therefore, the optimistic planning (Line 3) guarantees $\max _{a} f_{1}^{k}\left(s_{1}, a\right) \geq V_{1}^{\star}\left(s_{1}\right)$.
If the Bellman error summation is small (Line 6) i.e., $\sum_{h=1}^{H} \mathcal{E}\left(f^{k}, \pi^{k}, h\right) \leq H \zeta_{\text {act }}$, then by simple policy loss decomposition (e.g., Lemma 1 in [1]) and the optimism of $f^{k}$, $\pi^{k}$ is $H \zeta_{\text {act }}$-optimal. Otherwise, the elimination procedure is activated at some step $t$ satisfying $\mathcal{E}\left(f^{k}, \pi^{k}, t\right) \geq \zeta_{\text {act }}$ and all $f$ with $\mathcal{E}\left(f, \pi^{k}, t\right) \geq \zeta_{\text {elim }}$ get eliminated. The key observation here is:

If the elimination procedure is activated at step $h$ in phase $k_{1}<\ldots<k_{m}$, then the roll-in distribution of $\pi^{k_{1}}, \ldots, \pi^{k_{m}}$ at step $h$ is an $\zeta_{\text {act }}$-independent sequence with respect to the class of Bellman residuals $\left(I-\mathcal{T}_{h}\right) \mathcal{F}$ at step $h$. Therefore, we should have $m \leq d$.

For the sake of contradiction, assume $m \geq d+1$. Let us prove $\pi^{k_{1}}, \ldots, \pi^{k_{d+1}}$ is a $\zeta_{\text {act }}$-independent sequence. Firstly, for any $j \in[d+1]$, since $f^{k_{j}}$ is not eliminated in phase $k_{1}, \ldots, k_{j-1}$, we have

$$
\sqrt{\sum_{i=1}^{j-1}\left(\mathcal{E}\left(f^{k_{j}}, \pi^{k_{i}}, h\right)\right)^{2}} \leq \sqrt{d} \times \zeta_{\text {elim }} \leq \zeta_{\text {act }}
$$

Besides, because the elimination procedure is activated at step $h$ in phase $k_{j}$, we have $\mathcal{E}\left(f^{k_{j}}, \pi^{k_{j}}, h\right) \geq \zeta_{\text {act }}$. By Definition 6, we obtain that the roll-in distribution of $\pi^{k_{j}}$ at step $h$ is $\zeta_{\text {act }}$-independent of those of $\pi^{k_{1}}, \ldots, \pi^{k_{j-1}}$ for $j \in[d+1]$, which contradicts the definition $d=\operatorname{dim}_{\mathrm{BE}}\left(\mathcal{F}, \mathcal{D}_{\mathcal{F}}, \zeta_{\text {act }}\right)$. As a result, the elimination procedure can happen at most $d$ times for each $h \in[H]$, which means the algorithm should terminate within $d H+1$ phases and output an $H \zeta_{\text {act-optimal policy. }}$

## B V-type BE Dimension and Algorithms

The definition of Bellman rank, mentioned in Definition 10 and Proposition 11, is slightly different from the original definition in [1]. We denote the former by Q-type and the latter (the original definition) by V-type. In this section we introduce V-type BE Dimension as well as V-type variants of Golf and Olive. We show that similar results also hold for the V-type variants.
Definition 19 (V-type Bellman rank). The V-type Bellman rank is the minimum integer $d$ so that there exists $\phi_{h}: \mathcal{F} \rightarrow \mathbb{R}^{d}$ and $\psi_{h}: \mathcal{F} \rightarrow \mathbb{R}^{d}$ for each $h \in[H]$, such that for any $f, f^{\prime} \in \mathcal{F}$, the average V-type Bellman error

$$
\mathcal{E}_{V}\left(f, \pi_{f^{\prime}}, h\right):=\mathbb{E}\left[\left(f_{h}-\mathcal{T}_{h} f_{h+1}\right)\left(s_{h}, a_{h}\right) \mid s_{h} \sim \pi_{f^{\prime}}, a_{h} \sim \pi_{f}\right]=\left\langle\phi_{h}(f), \psi_{h}\left(f^{\prime}\right)\right\rangle,
$$

where $\left\|\phi_{h}(f)\right\|_{2} \cdot\left\|\psi_{h}\left(f^{\prime}\right)\right\|_{2} \leq \zeta$, and $\zeta$ is the normalization parameter.
The only difference between these two definitions is how we sample $a_{h}$. In the Q-type definition we have $a_{h} \sim \pi_{f^{\prime}}$ (the roll-in policy), however in the V-type definition we have $a_{h} \sim \pi_{f}$ (the greedy policy of the function evaluated in the Bellman error) instead. It is worth mentioning that the Q-type and V-type bellman error coincide whenever $f=f^{\prime}$; namely, $\mathcal{E}\left(f, \pi_{f}, h\right)=\mathcal{E}_{\mathrm{V}}\left(f, \pi_{f}, h\right)$ for all $f \in \mathcal{F}$.
We can similarly define the V-type variant of BE Dimension. At a high level, V-type BE dimension $\operatorname{dim}_{\mathrm{VBE}}(\mathcal{F}, \Pi, \epsilon)$ measures the complexity of finding a function in $\mathcal{F}$ such that its expected Bellman error under any state distribution in $\Pi$ is smaller than $\epsilon$.

```
Algorithm 3 V-type \(\operatorname{Golf}(\mathcal{F}, K, \beta)\)
    Initialize: \(\mathcal{D}_{1}, \ldots, \mathcal{D}_{H} \leftarrow \emptyset, \mathcal{B}^{0} \leftarrow \mathcal{F}\).
    for epoch \(k\) from 1 to \(K\) do
        Choose policy \(\pi^{k}=\pi_{f^{k}}\), where \(f^{k}=\operatorname{argmax}_{f \in \mathcal{B}^{k-1}} f\left(s_{1}, \pi_{f}\left(s_{1}\right)\right)\).
        for step \(h\) from 1 to \(H\) do
            Collect a tuple \(\left(s_{h}, a_{h}, r_{h}, s_{h+1}\right)\) by executing \(\pi^{k}\) at step \(1, \ldots, h-1\) and taking action
            uniformly at random at step \(h\).
            Augment \(\mathcal{D}_{h}=\mathcal{D}_{h} \cup\left\{\left(s_{h}, a_{h}, r_{h}, s_{h+1}\right)\right\}\) for all \(h \in[H]\).
        Update
            \(\mathcal{B}^{k}=\left\{f \in \mathcal{F}: \mathcal{L}_{\mathcal{D}_{h}}\left(f_{h}, f_{h+1}\right) \leq \inf _{g \in \mathcal{G}_{h}} \mathcal{L}_{\mathcal{D}_{h}}\left(g, f_{h+1}\right)+\beta\right.\) for all \(\left.h \in[H]\right\}\),
                        where \(\mathcal{L}_{\mathcal{D}_{h}}\left(\xi_{h}, \zeta_{h+1}\right)=\sum_{\left(s, a, r, s^{\prime}\right) \in \mathcal{D}_{h}}\left[\xi_{h}(s, a)-r-\max _{a^{\prime} \in \mathcal{A}} \zeta_{h+1}\left(s^{\prime}, a^{\prime}\right)\right]^{2}\).
    Output \(\pi^{\text {out }}\) sampled uniformly at random from \(\left\{\pi^{k}\right\}_{k=1}^{K}\).
```

Definition 20 (V-type BE dimension). Let $\left(I-\mathcal{T}_{h}\right) V_{\mathcal{F}} \subseteq(\mathcal{S} \rightarrow \mathbb{R})$ be the state-wise Bellman residual class of $\mathcal{F}$ at step $h$ which is defined as

$$
\left(I-\mathcal{T}_{h}\right) V_{\mathcal{F}}:=\left\{s \mapsto\left(f_{h}-\mathcal{T}_{h} f_{h+1}\right)\left(s, \pi_{f_{h}}(s)\right): f \in \mathcal{F}\right\} .
$$

Let $\Pi=\left\{\Pi_{h}\right\}_{h=1}^{H}$ be a collection of $H$ probability measure families over $\mathcal{S}$. The V-type $\epsilon$-BE dimension of $\mathcal{F}$ with respect to $\Pi$ is defined as

$$
\operatorname{dim}_{\mathrm{VBE}}(\mathcal{F}, \Pi, \epsilon):=\max _{h \in[H]} \operatorname{dim}_{\mathrm{DE}}\left(\left(I-\mathcal{T}_{h}\right) V_{\mathcal{F}}, \Pi_{h}, \epsilon\right)
$$

Relation with low V-type Bellman rank With slight abuse of notation, denote by $\mathcal{D}_{\mathcal{F}, h}$ the collection of all probability measures over $\mathcal{S}$ at the $h^{\text {th }}$ step, which can be generated by rolling in with a greedy policy $\pi_{f}$ with $f \in \mathcal{F}$. Similar to Proposition 11 , the following proposition claims that the V-type BE dimension of $\mathcal{F}$ with respect to $\mathcal{D}_{\mathcal{F}}:=\left\{\mathcal{D}_{\mathcal{F}, h}\right\}_{h \in[H]}$ is always upper bounded by its V-type Bellman rank up to some logarithmic factor.
Proposition 21 (low V-type Bellman rank $\subset$ low V-type BE dimension). If an MDP with function class $\mathcal{F}$ has V-type Bellman rank d with normalization parameter $\zeta$, then

$$
\operatorname{dim}_{\mathrm{VBE}}\left(\mathcal{F}, \mathcal{D}_{\mathcal{F}}, \epsilon\right) \leq \mathcal{O}(1+d \log (1+\zeta / \epsilon))
$$

The proof of Proposition 21 is almost the same as that of Proposition 11 in Appendix D.1. We omit it here since the only modification is to replace Q-type Bellman rank with its V-type variant wherever it is used.

## B. 1 Algorithm V-type GoLF

In this section we describe the V-type variant of Golf. The pseudocode is provided in Algorithm 3. Its only difference from the Q-type analogue is in Line 5: for each $h \in[H]$, we roll in with policy $\pi^{k}$ to sample $s_{h}$, and then instead of continuing following $\pi^{k}$ we take random action at step $h$.
Now we present the theoretical guarantee for Algorithm 3. Its proof is almost the same as that of Corollary 16 and can be found in appendix G.2.
Theorem 22 (V-type Golf). Under Assumption 1, 14, there exists an absolute constant c such that for any given $\epsilon>0$, if we choose $\beta=c \log \left[K H \mathcal{N}_{\mathcal{F} \cup \mathcal{G}}\left(\epsilon^{2} /\left(d|\mathcal{A}| H^{2}\right)\right)\right]$, then with probability at least $0.99, \pi^{\text {out }}$ is $\mathcal{O}(\epsilon)$-optimal, if

$$
K \geq \Omega\left(\frac{H^{2} d|\mathcal{A}|}{\epsilon^{2}} \cdot \log \left[\mathcal{N}_{\mathcal{F} \cup \mathcal{G}}\left(\frac{\epsilon^{2}}{H^{2} d|\mathcal{A}|}\right) \cdot \frac{H d|\mathcal{A}|}{\epsilon}\right]\right)
$$

where $d=\min _{\Pi \in\left\{\mathcal{D}_{\Delta}, \mathcal{D}_{\mathcal{F}}\right\}} \operatorname{dim}_{\mathrm{VBE}}(\mathcal{F}, \Pi, \epsilon / H)$.
Compared with Theorem 23 (V-type Olive), Theorem 22 (V-type Golf) has the following two advantages.

```
Algorithm 4 V-type \(\operatorname{OLIVE}\left(\mathcal{F}, \zeta_{\text {act }}, \zeta_{\text {elim }}, n_{\text {act }}, n_{\text {elim }}\right)\)
    Initialize: \(\mathcal{B}^{0} \leftarrow \mathcal{F}, \mathcal{D}_{h} \leftarrow \emptyset\) for all \(h, k\).
    for phase \(k=1,2, \ldots\) do
        Choose policy \(\pi^{k}=\pi_{f^{k}}\), where \(f^{k}=\operatorname{argmax}_{f \in \mathcal{B}^{k-1}} f\left(s_{1}, \pi_{f}\left(s_{1}\right)\right)\).
        Execute \(\pi^{k}\) for \(n_{\text {act }}\) episodes and refresh \(\mathcal{D}_{h}\) to include the fresh \(\left(s_{h}, a_{h}, r_{h}, s_{h+1}\right)\) tuples.
        Estimate \(\tilde{\mathcal{E}}_{\mathrm{V}}\left(f^{k}, \pi^{k}, h\right)\) for all \(h \in[H]\), where
                \(\tilde{\mathcal{E}}_{\mathrm{V}}\left(f^{k}, \pi^{k}, h\right)=\frac{1}{\left|\mathcal{D}_{h}\right|} \sum_{\left(s, a, r, s^{\prime}\right) \in \mathcal{D}_{h}}\left(f_{h}^{k}(s, a)-r-\max _{a^{\prime} \in \mathcal{A}} f_{h+1}^{k}\left(s^{\prime}, a^{\prime}\right)\right)\).
        if \(\sum_{h=1}^{H} \tilde{\mathcal{E}}_{\mathrm{V}}\left(f^{k}, \pi^{k}, h\right) \leq H \zeta_{\text {act }}\) then
            Terminate and output \(\pi^{k}\).
        Pick any \(t \in[H]\) for which \(\tilde{\mathcal{E}}_{\mathrm{V}}\left(f^{k}, \pi^{k}, t\right)>\zeta_{\text {act }}\).
        Collect \(n_{\text {elim }}\) episodes by executing \(\pi^{k}\) for step \(1, \ldots, t-1\) and picking action uniform at
        random for step \(t\). Refresh \(\mathcal{D}_{h}\) to include the fresh \(\left(s_{h}, a_{h}, r_{h}, s_{h+1}\right)\) tuples.
10: \(\quad\) Estimate \(\hat{\mathcal{E}}_{\mathrm{V}}\left(f, \pi^{k}, t\right)\) for all \(f \in \mathcal{F}\), where
            \(\hat{\mathcal{E}}_{\mathrm{V}}\left(f, \pi^{k}, h\right)=\frac{1}{\left|\mathcal{D}_{h}\right|} \sum_{\left(s, a, r, s^{\prime}\right) \in \mathcal{D}_{h}} \frac{\mathbf{1}\left[a=\pi_{f}(s)\right]}{1 /|\mathcal{A}|}\left(f_{h}(s, a)-r-\max _{a^{\prime} \in \mathcal{A}} f_{h+1}\left(s^{\prime}, a^{\prime}\right)\right)\).
11: \(\quad\) Update \(\mathcal{B}^{k}=\left\{f \in \mathcal{B}^{k-1}:\left|\hat{\mathcal{E}}_{\mathrm{V}}\left(f, \pi^{k}, t\right)\right| \leq \zeta_{\text {elim }}\right\}\).
```

- The sample complexity in Theorem 22 depends linearly on the V-type BE-dimension while the dependence in Theorem 23 is quadratic.
- Theorem 22 applies to RL problems of finite V-type BE dimension with respect to either $\mathcal{D}_{\mathcal{F}}$ or $\mathcal{D}_{\Delta}$. In comparison, Theorem 23 provides no guarantee for the $\mathcal{D}_{\Delta}$ case.

Finally, we comment that for the low Q-type BE dimension family, we provide both regret and sample complexity guarantees while for the low V-type counterpart, we only derive sample complexity result due to the need of taking actions uniformly at random in Algorithm 4 and Algorithm 3. [38] propose an algorithm that can achieve $\sqrt{T}$-regret for problems of low V-type Bellman rank. It is an interesting open problem to study whether similar techniques can be adapted to the low V-type BE dimension setting so that we can also obtain $\sqrt{T}$-regret.

## B. 2 Algorithm V-type Olive

In this section, we describe the original Olive (i.e., V-type Olive) proposed by [1], and its theoretical guarantee in terms of V-type BE dimension.

The pseudocode is provided in Algorithm 4. Its only difference from Algorithm 2 is Line 9-10: note that V-type Bellman rank needs the action at step $t$ to be greedy with respect to the function $f$ instead of being picked by the roll-in policy $\pi^{k}$, so we choose action $a_{t}$ uniformly at random and use the importance-weighted estimator to estimate the Bellman error for each $f$.

We have the following similar theoretical guarantee for Algorithm 4. Its proof is almost the same as that of Theorem 18 and can be found in Appendix G.1.
Theorem 23 (V-type OlIVE). Assume realizability (Assumption 1) holds and $\mathcal{F}$ is finite. There exists absolute constant $c$ such that if we choose

$$
\zeta_{\text {act }}=\frac{2 \epsilon}{H}, \zeta_{\text {elim }}=\frac{\epsilon}{2 H \sqrt{d}}, n_{\text {act }}=\frac{H^{2} \iota}{\epsilon^{2}}, \text { and } n_{\text {elim }}=\frac{H^{2} d|\mathcal{A}| \log (|\mathcal{F}|) \cdot \iota}{\epsilon^{2}}
$$

where $d=\operatorname{dim}_{\mathrm{VBE}}\left(\mathcal{F}, \mathcal{D}_{\mathcal{F}}, \epsilon / H\right)$ and $\iota=c \log [H d|\mathcal{A}| / \delta \epsilon]$, then with probability at least $1-\delta$, Algorithm 4 will output an $\mathcal{O}(\epsilon)$-optimal policy using at most $\mathcal{O}\left(H^{3} d^{2}|\mathcal{A}| \log (|\mathcal{F}|) \cdot \iota / \epsilon^{2}\right)$ episodes.

For problems with Bellman rank $d$ and finite function class $\mathcal{F}$, Theorem 23 together with Proposition 21 guarantees $\tilde{\mathcal{O}}\left(H^{3} d^{2}|\mathcal{A}| \log (|\mathcal{F}|) / \epsilon^{2}\right)$ samples suffice for finding an $\epsilon$-optimal policy, which matches the result in [1]. For function class $\mathcal{F}$ of infinite cardinality but with finite covering number, we can first compute an $\mathcal{O}\left(\zeta_{\text {elim }}\right)$-cover of $\mathcal{F}$, which we denote as $\mathcal{Z}_{\rho}$, and then run Algorithm 4 on
$\mathcal{Z}_{\rho}$. By following almost the same arguments in the proof of Theorem 23 (the only difference is to replace $Q^{\star}$ by its proxy in $\mathcal{Z}_{\rho}$ ), we can show Algorithm 4 will output an $\mathcal{O}(\epsilon)$-optimal policy using at most $\tilde{\Omega}\left(H^{3} d^{2}|\mathcal{A}| \log (N) / \epsilon^{2}\right)$ episodes where $N=\mathcal{N}_{\mathcal{F}}\left(\mathcal{O}\left(\zeta_{\text {elim }}\right)\right)$.

## B. 3 Discussions on Q-type versus V-type

In this paper, we have introduced two complementary definitions of Bellman rank: Q-type Bellman rank and V-type Bellman rank. And we prove they are upper bounds for Q-type and V-type BE dimension, respectively. Here, we want to emphasize that both Q-type and V-type Bellman rank have their own advantages. Specifically, the Q-type version has the following strengths.

1. There are natural RL problems whose Q-type Bellman rank is small, while their V-type Bellman rank is very large, e.g., the linear function approximation setting studied in in [8].
2. All the existing sample complexity results for the V-type cases scale linearly with respect to the number of actions, while those for the Q-type cases are independent of the number of actions. Therefore, for control problems such as Linear Quadratic Regulator (LQR), which has both small Q-type and V-type Bellman rank but infinite number of actions, the notion of Q-type is more suitable.

On the other hand, there are problems that naturally induce low V-type Bellman rank but have large Q-type Bellman rank, e.g., reactive POMDPs.

## C Examples

In this section, we introduce examples with low BE dimension. We will start with linear models and their variants, then introduce kernel MDPs, and finally present kernel reactive POMDPs which have low BE dimension, but possibly large Bellman rank and large Eluder dimension. All the proofs for this section are deferred to Appendix H.

## C. 1 Linear models and their variants

In this subsection, we review problems with linear structure in ascending order of generality. We start with the definition of linear MDPs [e.g., 7].
Definition 24 (Linear MDPs). We say an MDP is linear of dimension d if for each $h \in[H]$, there exists feature mappings $\phi_{h}: \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^{d}$, and d unknown signed measures $\psi_{h}=\left(\psi_{h}^{(1)}, \ldots, \psi_{h}^{(d)}\right)$ over $\mathcal{S}$, and an unknown vector $\theta_{h}^{r} \in \mathbb{R}^{d}$, such that $\mathbb{P}_{h}(\cdot \mid s, a)=\phi_{h}(s, a)^{\top} \psi_{h}(\cdot)$ and $r_{h}(s, a)=$ $\phi_{h}(s, a)^{\top} \theta_{h}^{r}$ for all $(s, a) \in \mathcal{S} \times \mathcal{A}$.

We remark that existing works [e.g., 7] usually assumxe $\phi$ is known to the learner. Next, we review a more general setting-the linear completeness setting [e.g., 8].
Definition 25 (Linear completeness setting). We say an MDP is in the linear completeness setting of dimension d, if there exists a feature mapping $\phi_{h}: \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^{d}$, such that for the linear function class $\mathcal{F}_{h}=\left\{\phi_{h}(\cdot)^{\top} \theta \mid \theta \in \mathbb{R}^{d}\right\}$, both Assumption 1 and 2 are satisfied.

We make three comments here. Firstly, we note that linear MDPs automatically satisfy both linear realizability and linear completeness assumptions, therefore are special cases of the linear completeness setting with the same ambient dimension. Secondly, only assuming linear realizability but without completeness is insufficient for sample-efficient learning (see exponential lower bounds in [44]). Finally, as mentioned in Appendix B.3, though MDPs in the linear completeness setting have low Q-type Bellman rank, their V-type Bellman rank can be arbitrarily large.

Finally, we review the generalized linear completeness setting [6], which generalizes the linear completeness setting by adding nonlinearity.
Definition 26 (Generalized linear completeness setting). We say an MDP is in the generalized linear completeness setting of dimension $d$, if there exists a feature mapping $\phi_{h}: \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^{d}$, and a link function $\sigma$, such that for the generalized linear function class $\mathcal{F}_{h}=\left\{\sigma\left(\phi_{h}(\cdot)^{\top} \theta\right) \mid \theta \in \mathbb{R}^{d}\right\}$, both Assumption 1 and 2 are satisfied, and the link function is strictly monotone, i.e., there exist $0<c_{1}<c_{2}<\infty$ such that $\sigma^{\prime}(x) \in\left[c_{1}, c_{2}\right]$ for all $x$.

One can directly verify by definition that when we choose link function $\sigma(x)=x$ in the generalized linear completeness setting, it will reduce to the standard linear version. Besides, it is known [18] the generalized linear completeness setting is a special case of low Eluder dimension, thus belonging to the low BE dimension family. Finally, we comment that despite the linear completeness setting belongs to the low Bellman rank family, the generalized version does not because of the possible nonlinearity of the link function.

## C. 2 Effective dimension and kernel MDPs

In this subsection, we introduce the notion of effective dimension. With this notion, we prove a useful proposition that any linear kernel function class with low effective dimension also has low Eluder dimension. This proposition directly implies that kernel MDPs are special cases of low Eluder dimension, which are also special cases of low BE dimension.

Effective dimension We start with the definition of effective dimension for a set, which is also known as critical information gain in Du et al. [41].
Definition 27 ( $\epsilon$-effective dimension of a set). The $\epsilon$-effective dimension of a set $\mathcal{X}$ is the minimum integer $d_{\mathrm{eff}}(\mathcal{X}, \epsilon)=n$ such that

$$
\begin{equation*}
\sup _{x_{1}, \ldots, x_{n} \in \mathcal{X}} \frac{1}{n} \log \operatorname{det}\left(\mathrm{I}+\frac{1}{\epsilon^{2}} \sum_{i=1}^{n} x_{i} x_{i}^{\top}\right) \leq e^{-1} \tag{6}
\end{equation*}
$$

Based on this definition, we can also define the effective dimension of a function class.
Definition 28 ( $\epsilon$-effective dimension of a function class). Given a function class $\mathcal{F}$ defined on $\mathcal{X}$, its $\epsilon$-effective dimension $d_{\text {eff }}(\mathcal{F}, \epsilon)=n$ is the minimum integer $n$ such that there exists a separable Hilbert space $\mathcal{H}$ and a mapping $\phi: \mathcal{X} \rightarrow \mathcal{H}$ so that

- for every $f \in \mathcal{F}$ there exists $\theta_{f} \in B_{\mathcal{H}}(1)$ satisfying $f(x)=\left\langle\theta_{f}, \phi(x)\right\rangle_{\mathcal{H}}$ for all $x \in \mathcal{X}$,
- $d_{\mathrm{eff}}(\phi(\mathcal{X}), \epsilon)=n$ where $\phi(\mathcal{X})=\{\phi(x): x \in \mathcal{X}\}$.

The following proposition shows that the Eluder dimension of any function class is always upper bounded by its effective dimension.
Proposition 29 (low effective dimension $\subset$ low Eluder dimension). For any function class $\mathcal{F}$ and domain $\mathcal{X}$, we have

$$
\operatorname{dim}_{\mathrm{E}}(\mathcal{F}, \epsilon) \leq \operatorname{dim}_{\mathrm{eff}}(\mathcal{F}, \epsilon / 2)
$$

On the other hand, we remark that effective dimension requires the existence of a benign linear structure in certain Hilbert spaces. In constrast, Eluder dimension does not require such conditions. Therefore, the function class of low Eluder dimension is more general than the function class of low effective dimension.

Kernel MDPs Now, we are ready to define kernel MDPs and prove it is a subclass of low Eluder dimension.

Definition 30 (Kernel MDPs). In a kernel MDP of effective dimension $d(\epsilon)$, for each step $h \in$ $[H]$, there exist feature mappings $\phi_{h}: \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{H}$ and $\psi_{h}: \mathcal{S} \rightarrow \mathcal{H}$ where $\mathcal{H}$ is a separable Hilbert space, so that the transition measure can be represented as the inner product of features, i.e., $\mathbb{P}_{h}\left(s^{\prime} \mid s, a\right)=\left\langle\phi_{h}(s, a), \psi_{h}\left(s^{\prime}\right)\right\rangle_{\mathcal{H}}$. Besides, the reward function is linear in $\phi$, i.e., $r_{h}(s, a)=$ $\left\langle\phi_{h}(s, a), \theta_{h}^{r}\right\rangle_{\mathcal{H}}$ for some $\theta_{h}^{r} \in \mathcal{H}$. Here, $\phi$ is known to the learner while $\psi$ and $\theta^{r}$ are unknown. Moreover, a kernel MDP satisfies the following regularization conditions: for all $h$

- $\left\|\theta_{h}^{r}\right\|_{\mathcal{H}} \leq 1$ and $\left\|\phi_{h}(s, a)\right\|_{\mathcal{H}} \leq 1$ for all $s, a$.
- $\left\|\sum_{s \in \mathcal{S}} \mathcal{V}(s) \psi_{h}(s)\right\|_{\mathcal{H}} \leq 1$ for any function $\mathcal{V}: \mathcal{S} \rightarrow[0,1]$.
- $\operatorname{dim}_{\text {eff }}\left(\mathcal{X}_{h}, \epsilon\right) \leq d(\epsilon)$ for all $h$ and $\epsilon$, where $\mathcal{X}_{h}=\left\{\phi_{h}(s, a):(s, a) \in \mathcal{S} \times \mathcal{A}\right\}$.

In order to learn kernel MDPs, we need to construct a proper function class $\mathcal{F}$. Formally, for each $h \in[H]$, we choose $\mathcal{F}_{h}=\left\{\phi_{h}(\cdot, \cdot)^{\top} \theta \mid \theta \in B_{\mathcal{H}}(H+1-h)\right\}$. One can easily verify $\mathcal{F}$ satisfies
both realizability and completeness by following the same arguments as in linear MDPs [7]. In order to apply Golf or Olive, we also need to show it has low BE dimension and bounded log-covering number. Below, we prove in sequence that $\mathcal{F}$ has low Eluder dimension and low log-covering number. Therefore, kernel MDPs fall into our low BE dimension framework.
Proposition 31 (kernel MDPs $\subset$ low Eluder dimension). Let $\mathcal{M}$ be a kernel MDP of effective dimension $d(\epsilon)$, then

$$
\operatorname{dim}_{\mathrm{E}}(\mathcal{F}, \epsilon) \leq d(\epsilon / 2 H)
$$

Proposition 31 follows directly from Proposition 29 by rescaling the parameters. Utilizing Proposition 31, we can further prove the log-covering number of $\mathcal{F}$ is also upper bounded by the effective dimension of the kernel MDP up to some logarithmic factor.
Proposition 32 (bounded covering number). Let $\mathcal{M}$ be a kernel MDP of effective dimension $d(\epsilon)$, then

$$
\log \mathcal{N}_{\mathcal{F}}(\epsilon) \leq \mathcal{O}(H d(\epsilon) \cdot \log (1+d(\epsilon) H / \epsilon))
$$

## C. 3 Effective Bellman rank and kernel reactive POMDPs

To begin with, we introduce the definition of effective Bellman rank and prove that it is always an upper bound for BE dimension. We will see effective Bellman rank serves as a useful tool for controlling the BE dimension of the example discussed in this section-kernel reactive POMDPs.

Q-type effective Bellman rank We start with Q-type $\epsilon$-effective Bellman rank which is simply the $\epsilon$-effective dimension of a special feature set.
Definition 33 (Q-type $\epsilon$-effective Bellman rank). The Q-type $\epsilon$-effective Bellman rank is the minimum integer $d$ so that

- There exists $\phi_{h}: \mathcal{F} \rightarrow \mathcal{H}$ and $\psi_{h}: \mathcal{F} \rightarrow \mathcal{H}$ for each $h \in[H]$ where $\mathcal{H}$ is a separable Hilbert space, such that for any $f, f^{\prime} \in \mathcal{F}$, the average Bellman error

$$
\mathcal{E}\left(f, \pi_{f^{\prime}}, h\right):=\mathbb{E}_{\pi_{f^{\prime}}}\left[\left(f_{h}-\mathcal{T}_{h} f_{h+1}\right)\left(s_{h}, a_{h}\right)\right]=\left\langle\phi_{h}(f), \psi_{h}\left(f^{\prime}\right)\right\rangle_{\mathcal{H}}
$$

where $\left\|\phi_{h}(f)\right\|_{\mathcal{H}} \leq \zeta$, and $\zeta$ is the normalization parameter.

- $d=\max _{h \in[H]} d_{\text {eff }}\left(\mathcal{X}_{h}(\psi, \mathcal{F}), \epsilon / \zeta\right)$ where $\mathcal{X}_{h}(\psi, \mathcal{F})=\left\{\psi_{h}\left(f_{h}\right): f_{h} \in \mathcal{F}_{h}\right\}$.

One can easily verify that when $\mathcal{H}$ is a finite-dimensional Euclidean space, the $\epsilon$-effective Bellman rank is always upper bounded by the original Bellman rank up to a logarithmic factor in $\zeta$ and $\epsilon^{-1}$. Moreover, the effective Bellman rank can be much smaller than the original Bellman rank if the induced feature set $\left\{\mathcal{X}_{h}(\psi, \mathcal{F})\right\}_{h \in[H]}$ approximately lies in a low-dimensional linear subspace. Therefore, effective Bellman rank can be viewed as a strict generalization of the original version.
Proposition 34 (low Q-type effective Bellman rank $\subset$ low Q-type BE dimension). Suppose function class $\mathcal{F}$ has $Q$-type $\epsilon$-effective Bellman rank d, then

$$
\operatorname{dim}_{\mathrm{BE}}\left(\mathcal{F}, \mathcal{D}_{\mathcal{F}}, \epsilon\right) \leq d
$$

Proposition 34 claims that problems with low Q-type effective Bellman rank also have low Q-type BE dimension.

V-type effective Bellman rank We can similarly define the V-type variant of effective Bellman rank, and prove it is always an upper bound for V-type BE dimension.
Definition 35 (V-type $\epsilon$-effective Bellman rank). The V-type $\epsilon$-effective Bellman rank is the minimum integer $d$ so that

- There exists $\phi_{h}: \mathcal{F} \rightarrow \mathcal{H}$ and $\psi_{h}: \mathcal{F} \rightarrow \mathcal{H}$ for each $h \in[H]$ where $\mathcal{H}$ is a separable Hilbert space, such that for any $f, f^{\prime} \in \mathcal{F}$, the average Bellman error

$$
\mathcal{E}_{V}\left(f, \pi_{f^{\prime}}, h\right):=\mathbb{E}\left[\left(f_{h}-\mathcal{T}_{h} f_{h+1}\right)\left(s_{h}, a_{h}\right) \mid s_{h} \sim \pi_{f^{\prime}}, a_{h} \sim \pi_{f}\right]=\left\langle\phi_{h}(f), \psi_{h}\left(f^{\prime}\right)\right\rangle_{\mathcal{H}}
$$

where $\left\|\phi_{h}(f)\right\|_{\mathcal{H}} \leq \zeta$, and $\zeta$ is the normalization parameter.

- $d=\max _{h \in[H]} d_{\text {eff }}\left(\mathcal{X}_{h}(\psi, \mathcal{F}), \epsilon / \zeta\right)$ where $\mathcal{X}_{h}(\psi, \mathcal{F})=\left\{\psi_{h}\left(f_{h}\right): f_{h} \in \mathcal{F}_{h}\right\}$.

Proposition 36 (low V-type effective Bellman rank $\subset$ low V-type BE dimension). Suppose function class $\mathcal{F}$ has $V$-type $\epsilon$-effective Bellman rank d, then

$$
\operatorname{dim}_{\mathrm{VBE}}\left(\mathcal{F}, \mathcal{D}_{\mathcal{F}}, \epsilon\right) \leq d
$$

The proof of Proposition 36 is almost the same as that of Proposition 34. We omit it since the only modification is to replace Q-type effective Bellman rank with its V-type variant wherever it is used.

We want to briefly comment that the majority of examples introduced in [41] have low effective Bellman rank. For example, low occupancy complexity, linear $Q^{*} / V^{*}$, linear Bellman complete and $Q^{*}$ state aggregation have low Q-type effective Bellman rank. And the feature selection problem has low V-type Bellman rank.

Kernel reactive POMDPs We start with the definition of POMDPs. A POMDP is defined by a tuple $(\mathcal{S}, \mathcal{A}, \mathcal{O}, \mathbb{T}, \mathbb{O}, r, H)$ where $\mathcal{S}$ denotes the set of hidden states, $\mathcal{A}$ denotes the set of actions, $\mathcal{O}$ denotes the set of observations, $\mathbb{T}$ denotes the transition measure, $\mathbb{O}$ denotes the emission measure, $r=\left\{r_{h}\right\}_{h=1}^{H}$ denotes the collections of reward functions, and $H$ denotes the length of each episode. At the beginning of each episode, the agent always starts from a fixed initial state. At each step $h \in[H]$, after reaching $s_{h}$, the agent will observe $o_{h} \sim \mathbb{O}_{h}\left(\cdot \mid s_{h}\right)$. Then the agent picks action $a_{h}$, receives $r_{h}\left(o_{h}, a_{h}\right)$ and transits to $s_{h+1} \sim \mathbb{T}_{h}\left(\cdot \mid s_{h}, a_{h}\right)$. In POMDPs, the agent can never directly observe the states $s_{1: H}$. It can only observe $o_{1: H}$ and $r_{1: H}$. Now we are ready to formally define kernel reactive POMDPs.
Definition 37 (Kernel reactive POMDPs). A kernel reactive POMDP is a POMDP that additionally satisfies the following two conditions

- For each $h \in[H]$, there exist mappings $\phi_{h}: \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{H}$ and $\psi_{h}: \mathcal{S} \rightarrow \mathcal{H}$ where $\mathcal{H}$ is a separable Hilbert space, such that $\mathbb{T}_{h}\left(s^{\prime} \mid s, a\right)=\left\langle\phi_{h}(s, a), \psi_{h}\left(s^{\prime}\right)\right\rangle_{\mathcal{H}}$ for all $s^{\prime}, a, s$. Moreover, for any function $\mathcal{V}: \mathcal{S} \rightarrow[0,1],\left\|\sum_{s^{\prime} \in \mathcal{S}} \mathcal{V}\left(s^{\prime}\right) \psi_{h}\left(s^{\prime}\right)\right\|_{\mathcal{H}} \leq 1$.
- (Reactiveness) The optimal action-value function $Q^{*}$ only depends on the current observation and action, i.e., for each $h \in[H]$, there exists function $f_{h}^{*}: \mathcal{O} \times \mathcal{A} \rightarrow[0,1]$ such that for all $\tau_{h}=\left[o_{1}, a_{1}, r_{1}, \ldots, o_{h}\right]$ and $a_{h}$

$$
Q_{h}^{*}\left(\tau_{h}, a_{h}\right)=f_{h}^{*}\left(o_{h}, a_{h}\right)
$$

The following proposition shows that when a kernel reactive POMDP has low effective dimension, it also has low V-type BE dimension.

Proposition 38 (kernel reactive POMDPs $\subset$ low V-type BE dimension). Any kernel reactive POMDP and function class $\mathcal{F} \subseteq(\mathcal{O} \times \mathcal{A} \rightarrow[0,1])$ satisfy

$$
\operatorname{dim}_{\mathrm{VBE}}\left(\mathcal{F}, \mathcal{D}_{\mathcal{F}}, \epsilon\right) \leq \max _{h \in[H]} d_{\mathrm{eff}}\left(\mathcal{X}_{h}, \epsilon / 2\right)
$$

where $\mathcal{X}_{h}=\left\{\mathbb{E}_{\pi_{f}}\left[\phi_{h}\left(s_{h}, a_{h}\right)\right]: f \in \mathcal{F}\right\}$.
We comment that when $\mathcal{H}$ approximately aligns with a low-dimensional linear subspace, the V-type effective Bellman rank in Proposition 38 will also be low. However, the Eluder dimension of $\mathcal{F}$ can be arbitrarily large because we basically pose no structural assumption on $\mathcal{F}$. Besides, its V/Q-type original Bellman rank can also be arbitrarily large, because $\mathcal{H}$ may be infinite-dimensional and the observation set $\mathcal{O}$ may be exponentially large. If we additionally assume $\mathcal{F}$ satisfies realizability $\left(f^{*} \in \mathcal{F}\right)$, then we can apply V-type OLIVE and obtain polynomial sample-complexity guarantee.

## D Proofs for BE Dimension

In this section, we provide formal proofs for the results stated in Section 3.

## D. 1 Proof of Proposition 11

The proof is basically the same as that of Example 3 in [18] with minor modification.

Proof. Without loss of generality, assume $\max \left\{\left\|\phi_{h}(f)\right\|_{2},\left\|\psi_{h}(f)\right\|_{2}\right\} \leq \sqrt{\zeta}$, otherwise we can satisfy this assumption by rescaling the feature mappings. Assume there exists $h \in[H]$ such that $\operatorname{dim}_{\text {DE }}\left(\left(I-\mathcal{T}_{h}\right) \mathcal{F}, \mathcal{D}_{\mathcal{F}, h}, \epsilon\right) \geq m$. Let $\mu_{1}, \ldots, \mu_{m} \in \mathcal{D}_{\mathcal{F}, h}$ be a an $\epsilon$-independent sequence with respect to $\left(I-\mathcal{T}_{h}\right) \mathcal{F}$. By Definition 6, there exists $f^{1}, \ldots, f^{m}$ such that for all $i \in[m]$, $\sqrt{\sum_{t=1}^{i-1}\left(\mathbb{E}_{\mu_{t}}\left[f_{h}^{i}-\mathcal{T}_{h} f_{h+1}^{i}\right]\right)^{2}} \leq \epsilon$ and $\left|\mathbb{E}_{\mu_{i}}\left[f_{h}^{i}-\mathcal{T}_{h} f_{h+1}^{i}\right]\right|>\epsilon$. Since $\mu_{1}, \ldots, \mu_{n} \in \mathcal{D}_{\mathcal{F}, h}$, there exist $g^{1}, \ldots, g^{n} \in \mathcal{F}$ so that $\mu_{i}$ is generated by executing $\pi_{g^{i}}$ for all $i \in[n]$.
By the definition of Bellman rank, this is equivalent to: for all $i \in[m]$, $\sqrt{\sum_{t=1}^{i-1}\left(\left\langle\phi_{h}\left(g^{i}\right), \psi_{h}\left(f^{t}\right)\right\rangle\right)^{2}} \leq \epsilon$ and $\left|\left\langle\phi_{h}\left(g^{i}\right), \psi_{h}\left(f^{i}\right)\right\rangle\right|>\epsilon$.
For notational simplicity, define $\mathbf{x}_{i}=\phi_{h}\left(g^{i}\right), \mathbf{z}_{i}=\psi_{h}\left(f^{i}\right)$ and $\mathbf{V}_{i}=\sum_{t=1}^{i-1} \mathbf{z}_{t} \mathbf{z}_{t}^{\top}+\frac{\epsilon^{2}}{\zeta} \cdot \mathbf{I}$. The previous argument directly implies: for all $i \in[m],\left\|\mathbf{x}_{i}\right\|_{\mathbf{V}_{i}} \leq \sqrt{2} \epsilon$ and $\left\|\mathbf{x}_{i}\right\|_{\mathbf{V}_{i}} \cdot\left\|\mathbf{z}_{i}\right\|_{\mathbf{V}_{i}^{-1}}>\epsilon$. Therefore, we have $\left\|\mathbf{z}_{i}\right\|_{\mathbf{V}_{i}^{-1}} \geq \frac{1}{\sqrt{2}}$.
By the matrix determinant lemma,
$\operatorname{det}\left[\mathbf{V}_{m}\right]=\operatorname{det}\left[\mathbf{V}_{m-1}\right]\left(1+\left\|\mathbf{z}_{m}\right\|_{\mathbf{V}_{m}^{-1}}^{2}\right) \geq \frac{3}{2} \operatorname{det}\left[\mathbf{V}_{m-1}\right] \geq \ldots \geq \operatorname{det}\left[\frac{\epsilon^{2}}{\zeta} \cdot \mathbf{I}\right]\left(\frac{3}{2}\right)^{m-1}=\left(\frac{\epsilon^{2}}{\zeta}\right)^{d}\left(\frac{3}{2}\right)^{m-1}$.
On the other hand,

$$
\operatorname{det}\left[\mathbf{V}_{m}\right] \leq\left(\frac{\operatorname{trace}\left[\mathbf{V}_{m}\right]}{d}\right)^{d} \leq\left(\frac{\zeta(m-1)}{d}+\frac{\epsilon^{2}}{\zeta}\right)^{d}
$$

Therefore, we obtain

$$
\left(\frac{3}{2}\right)^{m-1} \leq\left(\frac{\zeta^{2}(m-1)}{d \epsilon^{2}}+1\right)^{d}
$$

Take logarithm on both sides,

$$
m \leq 4\left[1+d \log \left(\frac{\zeta^{2}(m-1)}{d \epsilon^{2}}+1\right)\right]
$$

which, by simple calculation, implies

$$
m \leq \mathcal{O}\left(1+d \log \left(\frac{\zeta^{2}}{\epsilon^{2}}+1\right)\right)
$$

## D. 2 Proof of Proposition 12

Proof. Assume $\delta_{z_{1}}, \ldots, \delta_{z_{m}}$ is an $\epsilon$-independent sequence of distributions with respect to $\left(I-\mathcal{T}_{h}\right) \mathcal{F}$, where $\delta_{z_{i}} \in \mathcal{D}_{\Delta}$. By Definition 6, there exist functions $f^{1}, \ldots, f^{m} \in \mathcal{F}$ such that for all $i \in[m]$, we have $\left|\left(f_{h}^{i}-\mathcal{T}_{h} f_{h+1}^{i}\right)\left(z_{i}\right)\right|>\epsilon$ and $\sqrt{\sum_{t=1}^{i-1}\left|\left(f_{h}^{i}-\mathcal{T}_{h} f_{h+1}^{i}\right)\left(z_{t}\right)\right|^{2}} \leq \epsilon$. Define $g_{h}^{i}=\mathcal{T}_{h} f_{h+1}^{i}$. Note that $g_{h}^{i} \in \mathcal{F}_{h}$ because $\mathcal{T}_{h} \mathcal{F}_{h+1} \subset \mathcal{F}_{h}$. Therefore, we have for all $i \in[m],\left|\left(f_{h}^{i}-g_{h}^{i}\right)\left(z_{i}\right)\right|>\epsilon$ and $\sqrt{\sum_{t=1}^{i-1}\left|\left(f_{h}^{i}-g_{h}^{i}\right)\left(z_{t}\right)\right|^{2}} \leq \epsilon$ with $f_{h}^{i}, g_{h}^{i} \in \mathcal{F}_{h}$. By Definition 4 and 5, this implies $\operatorname{dim}_{\mathrm{E}}\left(\mathcal{F}_{h}, \epsilon\right) \geq$ $m$, which completes the proof.

## D. 3 Proof of Proposition 13

Proof. For any $m \in \mathbb{N}^{+}$, denote by $e_{1}, \ldots, e_{m}$ the basis vectors in $\mathbb{R}^{m}$, and consider the following linear bandits $(|\mathcal{S}|=H=1)$ problem.

- The action set $\mathcal{A}=\left\{a_{i}=\left(1 ; e_{i}\right) \in \mathbb{R}^{m+1}: i \in[m]\right\}$.
- The function set $\mathcal{F}_{1}=\left\{f_{\theta_{i}}(a)=a^{\top} \theta_{i}: \theta_{i}=\left(1 ; e_{i}\right), i \in[m]\right\}$.
- The reward function is always zero, i.e., $r \equiv 0$.

Eluder dimension For any $\epsilon \in(0,1], a_{1}, \ldots, a_{m-1}$ is an $\epsilon$-independent sequence of points because: (a) for any $t \in[m-1], \sum_{i=1}^{t-1}\left(f_{\theta_{t}}\left(a_{i}\right)-f_{\theta_{t+1}}\left(a_{i}\right)\right)^{2}=0$; (b) for any $t \in[m-1]$, $f_{\theta_{t}}\left(a_{t}\right)-f_{\theta_{t+1}}\left(a_{t}\right)=1 \geq \epsilon$. Therefore, $\min _{h \in[H]} \operatorname{dim}_{\mathrm{E}}\left(\mathcal{F}_{h}, \epsilon\right)=\operatorname{dim}_{\mathrm{E}}\left(\mathcal{F}_{1}, \epsilon\right) \geq m-1$.

Bellman rank It is direct to see the Bellman residual matrix is $\mathcal{E}:=\Theta^{\top} \Theta \in \mathbb{R}^{m \times m}$ with rank $m$, where $\Theta=\left[\theta_{1}, \theta_{2}, \ldots, \theta_{m}\right]$. As a result, the Bellman rank is at least $m$.

BE dimension First, note in this setting $\left(I-\mathcal{T}_{1}\right) \mathcal{F}$ is simply $\mathcal{F}_{1}$ (because $\mathcal{F}_{2}=\{0\}$ and $r \equiv 0$ ), and $\mathcal{D}_{\mathcal{F}}$ coincides with $\mathcal{D}_{\Delta}$, so it suffices to show $\operatorname{dim}_{\text {DE }}\left(\mathcal{F}_{1}, \mathcal{D}_{\Delta}, \epsilon\right) \leq 5$.
Assume $\operatorname{dim}_{\mathrm{DE}}\left(\mathcal{F}_{1}, \mathcal{D}_{\Delta}, \epsilon\right)=k$. Then there exist $q_{1}, \ldots, q_{k} \in \mathcal{A}$ and $w_{1}, \ldots, w_{k} \in \mathcal{A}$ such that for all $t \in[k], \sqrt{\sum_{i=1}^{t-1}\left(\left\langle q_{t}, w_{i}\right\rangle\right)^{2}} \leq \epsilon$ and $\left|\left\langle q_{t}, w_{t}\right\rangle\right|>\epsilon$. By simple calculation, we have $q_{i}^{\top} w_{j} \in[1,2]$ for all $i, j \in[k]$. Therefore, if $\epsilon>2$, then $k=0$ because $\left|\left\langle q_{t}, w_{t}\right\rangle\right| \leq 2$; if $\epsilon \leq 2$, then $k \leq 5$ because $\sqrt{k-1} \leq \sqrt{\sum_{i=1}^{k-1}\left(\left\langle q_{k}, w_{i}\right\rangle\right)^{2}} \leq \epsilon$.

## E Proofs for GOLF

In this section, we provide formal proofs for the results stated in Section 4.

## E. 1 Proof of Theorem 15

We start the proof with the following two lemmas. The first lemma shows that with high probability any function in the confidence set has low Bellman-error over the collected datasets $\mathcal{D}_{1}, \ldots, \mathcal{D}_{H}$ as well as the distributions from which $\mathcal{D}_{1}, \ldots, \mathcal{D}_{H}$ are sampled.
Lemma 39. Let $\rho>0$ be an arbitrary fixed number. If we choose $\beta=c\left(\log \left[K H \mathcal{N}_{\mathcal{F} \cup \mathcal{G}}(\rho) / \delta\right]+\right.$ $K \rho$ ) with some large absolute constant $c$ in Algorithm 1, then with probability at least $1-\delta$, for all $(k, h) \in[K] \times[H]$, we have
(a) $\sum_{i=1}^{k-1} \mathbb{E}\left[\left(f_{h}^{k}\left(s_{h}, a_{h}\right)-\left(\mathcal{T} f_{h+1}^{k}\right)\left(s_{h}, a_{h}\right)\right)^{2} \mid s_{h}, a_{h} \sim \pi^{i}\right] \leq \mathcal{O}(\beta)$.
(b) $\sum_{i=1}^{k-1}\left(f_{h}^{k}\left(s_{h}^{i}, a_{h}^{i}\right)-\left(\mathcal{T} f_{h+1}^{k}\right)\left(s_{h}^{i}, a_{h}^{i}\right)\right)^{2} \leq \mathcal{O}(\beta)$,
where $\left(s_{1}^{i}, a_{1}^{i}, \ldots, s_{H}^{i}, a_{H}^{i}, s_{H+1}^{i}\right)$ denotes the trajectory sampled by following $\pi^{i}$ in the $i^{\text {th }}$ episode.
The second lemma guarantees that the optimal value function is inside the confidence with high probability. As a result, the selected value function $f^{k}$ in each iteration shall be an upper bound of $Q^{\star}$ with high probability.
Lemma 40. Under the same condition of Lemma 39, with probability at least $1-\delta$, we have $Q^{\star} \in \mathcal{B}^{k}$ for all $k \in[K]$.

The proof of Lemma 39 and 40 relies on standard martingale concentration (e.g. Freedman's inequality) and can be found in Appendix E.3.

Step 1. Bounding the regret by Bellman error By Lemma 40, we can upper bound the cumulative regret by the summation of Bellman error with probability at least $1-\delta$ :

$$
\begin{equation*}
\sum_{k=1}^{K}\left(V_{1}^{\star}\left(s_{1}\right)-V_{1}^{\pi^{k}}\left(s_{1}\right)\right) \leq \sum_{k=1}^{K}\left(\max _{a} f_{1}^{k}\left(s_{1}, a\right)-V_{1}^{\pi^{k}}\left(s_{1}\right)\right) \stackrel{(i)}{=} \sum_{k=1}^{K} \sum_{h=1}^{H} \mathcal{E}\left(f^{k}, \pi^{k}, h\right) \tag{7}
\end{equation*}
$$

where $(i)$ follows from standard policy loss decomposition (e.g. Lemma 1 in [1]).
Step 2. Bounding cumulative Bellman error using DE dimension Next, we focus on a fixed step $h$ and bound the cumulative Bellman error $\sum_{k=1}^{K} \mathcal{E}\left(f^{k}, \pi^{k}, h\right)$ using Lemma 39. To proceed, we need the following lemma to control the accumulating rate of Bellman error.
Lemma 41. Given a function class $\Phi$ defined on $\mathcal{X}$ with $|\phi(x)| \leq C$ for all $(g, x) \in \Phi \times \mathcal{X}$, and a family of probability measures $\Pi$ over $\mathcal{X}$. Suppose sequence $\left\{\phi_{k}\right\}_{k=1}^{K} \subset \Phi$ and $\left\{\mu_{k}\right\}_{k=1}^{K} \subset \Pi$ satisfy that for all $k \in[K], \sum_{t=1}^{k-1}\left(\mathbb{E}_{\mu_{t}}\left[\phi_{k}\right]\right)^{2} \leq \beta$. Then for all $k \in[K]$ and $\omega>0$,

$$
\sum_{t=1}^{k}\left|\mathbb{E}_{\mu_{t}}\left[\phi_{t}\right]\right| \leq \mathcal{O}\left(\sqrt{\operatorname{dim}_{\mathrm{DE}}(\Phi, \Pi, \omega) \beta k}+\min \left\{k, \operatorname{dim}_{\mathrm{DE}}(\Phi, \Pi, \omega)\right\} C+k \omega\right)
$$

Lemma 41 is a simple modification of Lemma 2 in [18] and its proof can be found in Appendix E.4. We provide two ways to apply Lemma 41, which can produce regret bounds in term of two different complexity measures. If we invoke Lemma 39 (a) and Lemma 41 with

$$
\left\{\begin{array}{l}
\rho=\frac{1}{K}, \omega=\sqrt{\frac{1}{K}}, C=1, \\
\mathcal{X}=\mathcal{S} \times \mathcal{A}, \Phi=\left(I-\mathcal{T}_{h}\right) \mathcal{F}, \Pi=\mathcal{D}_{\mathcal{F}, h}, \\
\phi_{k}=f_{h}^{k}-\mathcal{T}_{h} f_{h+1}^{k} \text { and } \mu_{k}=\mathbb{P}^{\pi^{k}}\left(s_{h}=\cdot, a_{h}=\cdot\right),
\end{array}\right.
$$

we obtain

$$
\begin{equation*}
\sum_{t=1}^{k} \mathcal{E}\left(f^{t}, \pi^{t}, h\right) \leq \mathcal{O}\left(\sqrt{k \cdot \operatorname{dim}_{\mathrm{BE}}\left(\mathcal{F}, \mathcal{D}_{\mathcal{F}}, \sqrt{1 / K}\right) \log \left[K H \mathcal{N}_{\mathcal{F} \cup \mathcal{G}}(1 / K) / \delta\right]}\right) \tag{8}
\end{equation*}
$$

We can also invoke Lemma 39 (b) and Lemma 41 with

$$
\left\{\begin{array}{l}
\rho=\frac{1}{K}, \omega=\sqrt{\frac{1}{K}}, C=1 \\
\mathcal{X}=\mathcal{S} \times \mathcal{A}, \Phi=\left(I-\mathcal{T}_{h}\right) \mathcal{F}, \text { and } \Pi=\mathcal{D}_{\Delta, h} \\
\phi_{k}=f_{h}^{k}-\mathcal{T}_{h} f_{h+1}^{k} \text { and } \mu_{k}=\mathbf{1}\left\{\cdot=\left(s_{h}^{k}, a_{h}^{k}\right)\right\}
\end{array}\right.
$$

and obtain

$$
\begin{align*}
\sum_{t=1}^{k} \mathcal{E}\left(f^{t}, \pi^{t}, h\right) & \leq \sum_{t=1}^{k}\left(f_{h}^{t}-\mathcal{T} f_{h+1}^{t}\right)\left(s_{h}^{t}, a_{h}^{t}\right)+\mathcal{O}(\sqrt{k \log (k)})  \tag{9}\\
& \leq \mathcal{O}\left(\sqrt{k \cdot \operatorname{dim}_{\mathrm{BE}}\left(\mathcal{F}, \mathcal{D}_{\Delta}, \sqrt{1 / K}\right) \log \left[K H \mathcal{N}_{\mathcal{F} \cup \mathcal{G}}(1 / K) / \delta\right]}\right)
\end{align*}
$$

where the first inequality follows from standard martingale concentration.
Plugging either equation (8) or (9) back into equation (7) completes the proof.

## E. 2 Proof of Corollary 16

Step 1. Bounding the regret by Bellman error By Lemma 40, we can upper bound the cumulative regret by the summation of Bellman error with probability at least $1-\delta$ :

$$
\begin{equation*}
\sum_{k=1}^{K}\left(V_{1}^{\star}\left(s_{1}\right)-V_{1}^{\pi^{k}}\left(s_{1}\right)\right) \leq \sum_{k=1}^{K}\left(\max _{a} f_{1}^{k}\left(s_{1}, a\right)-V_{1}^{\pi^{k}}\left(s_{1}\right)\right) \stackrel{(i)}{=} \sum_{k=1}^{K} \sum_{h=1}^{H} \mathcal{E}\left(f^{k}, \pi^{k}, h\right) \tag{10}
\end{equation*}
$$

where $(i)$ follows from standard policy loss decomposition (e.g. Lemma 1 in [1]).
Step 2. Bounding cumulative Bellman error using DE dimension Next, we focus on a fixed step $h$ and bound the cumulative Bellman error $\sum_{k=1}^{K} \mathcal{E}\left(f^{k}, \pi^{k}, h\right)$ using Lemma 39.
If we invoke Lemma 39 (a) with

$$
\rho=\frac{\epsilon^{2}}{H^{2} \cdot \operatorname{dim}_{\mathrm{BE}}\left(\mathcal{F}, \mathcal{D}_{\mathcal{F}}, \epsilon / H\right)}
$$

and Lemma 41 with

$$
\left\{\begin{array}{l}
\omega=\frac{\epsilon}{H}, C=1 \\
\mathcal{X}=\mathcal{S} \times \mathcal{A}, \Phi=\left(I-\mathcal{T}_{h}\right) \mathcal{F}, \Pi=\mathcal{D}_{\mathcal{F}, h} \\
\phi_{k}=f_{h}^{k}-\mathcal{T}_{h} f_{h+1}^{k} \text { and } \mu_{k}=\mathbb{P}^{\pi^{k}}\left(s_{h}=\cdot, a_{h}=\cdot\right)
\end{array}\right.
$$

we obtain with probability at least $1-10^{-3}$,

$$
\begin{align*}
\frac{1}{K} \sum_{k=1}^{K} \mathcal{E}\left(f^{k}, \pi^{k}, h\right) & \leq \mathcal{O}\left(\sqrt{\operatorname{dim}_{\mathrm{BE}}\left(\mathcal{F}, \mathcal{D}_{\mathcal{F}}, \epsilon / H\right)\left[\frac{\log \left[K H \mathcal{N}_{\mathcal{F} \cup \mathcal{G}}(\rho)\right]}{K}+\rho\right]}+\frac{\epsilon}{H}\right)  \tag{11}\\
& \leq \mathcal{O}\left(\frac{\epsilon}{H}+\sqrt{\frac{d \log \left[K H \mathcal{N}_{\mathcal{F} \cup \mathcal{G}}(\rho)\right]}{K}}\right)
\end{align*}
$$

where the second inequality follows from the choice of $\rho$ and $d:=\operatorname{dim}_{B E}\left(\mathcal{F}, \mathcal{D}_{\mathcal{F}}, \epsilon / H\right)$. Now we need to choose $K$ such that

$$
\begin{equation*}
\sqrt{\frac{d \log \left[K H \mathcal{N}_{\mathcal{F} \cup \mathcal{G}}(\rho)\right]}{K}} \leq \frac{\epsilon}{H} \tag{12}
\end{equation*}
$$

By simple calculation, one can verify it suffices to choose

$$
\begin{equation*}
K=\frac{H^{2} d \log \left(H d \mathcal{N}_{\mathcal{F} \cup \mathcal{G}}(\rho) / \epsilon\right)}{\epsilon^{2}} \tag{13}
\end{equation*}
$$

Plugging equation (11) back into equation (10) completes the proof. We can similarly prove the bound in terms of the BE dimension with respect to $\mathcal{D}_{\Delta}$.

## E. 3 Proofs of concentration lemmas

To begin with, recall the Freedman's inequality that controls the sum of martingale difference by the sum of their predicted variance.
Lemma 42 (Freedman's inequality [e.g., 49]). Let $\left(Z_{t}\right)_{t \leq T}$ be a real-valued martingale difference sequence adapted to filtration $\mathfrak{F}_{t}$, and let $\mathbb{E}_{t}[\cdot]=\mathbb{E}\left[\cdot \mid \mathfrak{F}_{t}\right]$. If $\left|Z_{t}\right| \leq R$ almost surely, then for any $\eta \in\left(0, \frac{1}{R}\right)$ it holds that with probability at least $1-\delta$,

$$
\sum_{t=1}^{T} Z_{t} \leq \mathcal{O}\left(\eta \sum_{t=1}^{T} \mathbb{E}_{t-1}\left[Z_{t}^{2}\right]+\frac{\log \left(\delta^{-1}\right)}{\eta}\right)
$$

## E.3.1 Proof of Lemma 39

Proof. We prove inequality (b) first.
Consider a fixed $(k, h, f)$ tuple. Let
$X_{t}(h, f):=\left(f_{h}\left(s_{h}^{t}, a_{h}^{t}\right)-r_{h}^{t}-f_{h+1}\left(s_{h+1}^{t}, \pi_{f}\left(s_{h+1}^{t}\right)\right)\right)^{2}-\left(\left(\mathcal{T} f_{h+1}\right)\left(s_{h}^{t}, a_{h}^{t}\right)-r_{h}^{t}-f_{h+1}\left(s_{h+1}^{t}, \pi_{f}\left(s_{h+1}^{t}\right)\right)\right)^{2}$
and $\mathfrak{F}_{t, h}$ be the filtration induced by $\left\{s_{1}^{i}, a_{1}^{i}, r_{1}^{i}, \ldots, s_{H}^{i}\right\}_{i=1}^{t-1} \bigcup\left\{s_{1}^{t}, a_{1}^{t}, r_{1}^{t}, \ldots, s_{h}^{t}, a_{h}^{t}\right\}$. We have

$$
\mathbb{E}\left[X_{t}(h, f) \mid \mathfrak{F}_{t, h}\right]=\left[\left(f_{h}-\mathcal{T} f_{h+1}\right)\left(s_{h}^{t}, a_{h}^{t}\right)\right]^{2}
$$

and
$\operatorname{Var}\left[X_{t}(h, f) \mid \mathfrak{F}_{t, h}\right] \leq \mathbb{E}\left[\left(X_{t}(h, f)\right)^{2} \mid \mathfrak{F}_{t, h}\right] \leq 36\left[\left(f_{h}-\mathcal{T} f_{h+1}\right)\left(s_{h}^{t}, a_{h}^{t}\right)\right]^{2}=36 \mathbb{E}\left[X_{t}(h, f) \mid \mathfrak{F}_{t, h}\right]$.
By Freedman's inequality, we have, with probability at least $1-\delta$,

$$
\left|\sum_{t=1}^{k} X_{t}(h, f)-\sum_{t=1}^{k} \mathbb{E}\left[X_{t}(h, f) \mid \mathfrak{F}_{t, h}\right]\right| \leq \mathcal{O}\left(\sqrt{\log (1 / \delta) \sum_{t=1}^{k} \mathbb{E}\left[X_{t} \mid \mathfrak{F}_{t, h}\right]}+\log (1 / \delta)\right)
$$

Let $\mathcal{Z}_{\rho}$ be a $\rho$-cover of $\mathcal{F}$. Now taking a union bound for all $(k, h, \phi) \in[K] \times[H] \times \mathcal{Z}_{\rho}$, we obtain that with probability at least $1-\delta$, for all $(k, h, \phi) \in[K] \times[H] \times \mathcal{Z}_{\rho}$

$$
\begin{equation*}
\left|\sum_{t=1}^{k} X_{t}(h, \phi)-\sum_{t=1}^{k}\left[\left(\phi_{h}-\mathcal{T} \phi_{h+1}\right)\left(s_{h}^{t}, a_{h}^{t}\right)\right]^{2}\right| \leq \mathcal{O}\left(\sqrt{\iota \sum_{t=1}^{k}\left[\left(\phi_{h}-\mathcal{T} \phi_{h+1}\right)\left(s_{h}^{t}, a_{h}^{t}\right)\right]^{2}}+\iota\right) \tag{14}
\end{equation*}
$$

where $\iota=\log \left(H K\left|\mathcal{Z}_{\rho}\right| / \delta\right)$. From now on, we will do all the analysis conditioning on this event being true.

Consider an arbitrary $(h, k) \in[H] \times[K]$ pair. By the definition of $\mathcal{B}^{k}$ and Assumption 14

$$
\begin{aligned}
\sum_{t=1}^{k-1} X_{t}\left(h, f^{k}\right)= & \sum_{t=1}^{k-1}\left[f_{h}^{k}\left(s_{h}^{t}, a_{h}^{t}\right)-r_{h}^{t}-f_{h+1}^{k}\left(s_{h+1}^{t}, \pi_{f^{k}}\left(s_{h+1}^{t}\right)\right)\right]^{2} \\
& -\sum_{t=1}^{k-1}\left[\left(\mathcal{T} f_{h+1}^{k}\right)\left(s_{h}^{t}, a_{h}^{t}\right)-r_{h}^{t}-f_{h+1}^{k}\left(s_{h+1}^{t}, \pi_{f^{k}}\left(s_{h+1}^{t}\right)\right)\right]^{2} \\
\leq & \sum_{t=1}^{k-1}\left[f_{h}^{k}\left(s_{h}^{t}, a_{h}^{t}\right)-r_{h}^{t}-f_{h+1}^{k}\left(s_{h+1}^{t}, \pi_{f^{k}}\left(s_{h+1}^{t}\right)\right)\right]^{2} \\
& -\inf _{g \in \mathcal{G}} \sum_{t=1}^{k-1}\left[g_{h}\left(s_{h}^{t}, a_{h}^{t}\right)-r_{h}^{t}-f_{h+1}^{k}\left(s_{h+1}^{t}, \pi_{f^{k}}\left(s_{h+1}^{t}\right)\right)\right]^{2} \leq \beta
\end{aligned}
$$

Define $\phi^{k}=\operatorname{argmin}_{\phi \in \mathcal{Z}_{\rho}} \max _{h \in[H]}\left\|f_{h}^{k}-\phi_{h}^{k}\right\|_{\infty}$. By the definition of $\mathcal{Z}_{\rho}$, we have

$$
\left|\sum_{t=1}^{k-1} X_{t}\left(h, f^{k}\right)-\sum_{t=1}^{k-1} X_{t}\left(h, \phi^{k}\right)\right| \leq \mathcal{O}(k \rho)
$$

Therefore,

$$
\begin{equation*}
\sum_{t=1}^{k-1} X_{t}\left(h, \phi^{k}\right) \leq \mathcal{O}(k \rho)+\beta \tag{15}
\end{equation*}
$$

Recall inequality (14) implies

$$
\begin{equation*}
\left|\sum_{t=1}^{k-1} X_{t}\left(h, \phi^{k}\right)-\sum_{t=1}^{k-1}\left[\left(\phi_{h}^{k}-\mathcal{T} \phi_{h+1}^{k}\right)\left(s_{h}^{t}, a_{h}^{t}\right)\right]^{2}\right| \leq \mathcal{O}\left(\sqrt{\iota \sum_{t=1}^{k-1}\left[\left(\phi_{h}^{k}-\mathcal{T} \phi_{h+1}^{k}\right)\left(s_{h}^{t}, a_{h}^{t}\right)\right]^{2}}+\iota\right) \tag{16}
\end{equation*}
$$

Putting (15) and (16) together, we obtain

$$
\sum_{t=1}^{k-1}\left[\left(\phi_{h}^{k}-\mathcal{T} \phi_{h+1}^{k}\right)\left(s_{h}^{t}, a_{h}^{t}\right)\right]^{2} \leq \mathcal{O}(\iota+k \rho+\beta)
$$

Because $\phi^{k}$ is an $\rho$-approximation to $f^{k}$, we conclude

$$
\sum_{t=1}^{k-1}\left[\left(f_{h}^{k}-\mathcal{T} f_{h+1}^{k}\right)\left(s_{h}^{t}, a_{h}^{t}\right)\right]^{2} \leq \mathcal{O}(\iota+k \rho+\beta)
$$

Therefore, we prove inequality $(b)$ in Lemma 39.
To prove inequality $(a)$, we only need to redefine $\mathfrak{F}_{t, h}$ to be the filtration induced by
$\left\{s_{1}^{i}, a_{1}^{i}, r_{1}^{i}, \ldots, s_{H}^{i}\right\}_{i=1}^{t-1}$ and then repeat the arguments above verbatim.

## E.3.2 Proof of Lemma 40

Proof. Let $\mathcal{V}_{\rho}$ be a $\rho$-cover of $\mathcal{G}$.
Consider an arbitrary fixed tuple $(k, h, g) \in[K] \times[H] \times \mathcal{G}$. Let
$W_{t}(h, g):=\left(g_{h}\left(s_{h}^{t}, a_{h}^{t}\right)-r_{h}^{t}-Q_{h+1}^{\star}\left(s_{h+1}^{t}, \pi_{Q^{\star}}\left(s_{h+1}^{t}\right)\right)\right)^{2}-\left(Q_{h}^{\star}\left(s_{h}^{t}, a_{h}^{t}\right)-r_{h}^{t}-Q_{h+1}^{\star}\left(s_{h+1}^{t}, \pi_{Q^{\star}}\left(s_{h+1}^{t}\right)\right)\right)^{2}$ and $\mathfrak{F}_{t, h}$ be the filtration induced by $\left\{s_{1}^{i}, a_{1}^{i}, r_{1}^{i}, \ldots, s_{H}^{i}\right\}_{i=1}^{t-1} \bigcup\left\{s_{1}^{t}, a_{1}^{t}, r_{1}^{t}, \ldots, s_{h}^{t}, a_{h}^{t}\right\}$. We have

$$
\mathbb{E}\left[W_{t}(h, g) \mid \mathfrak{F}_{t, h}\right]=\left[\left(g_{h}-Q_{h}^{\star}\right)\left(s_{h}^{t}, a_{h}^{t}\right)\right]^{2}
$$

and

$$
\operatorname{Var}\left[W_{t}(h, g) \mid \mathfrak{F}_{t, h}\right] \leq \mathbb{E}\left[\left(W_{t}(h, g)\right)^{2} \mid \mathfrak{F}_{t, h}\right] \leq 36\left(\left(g_{h}-Q_{h}^{\star}\right)\left(s_{h}^{t}, a_{h}^{t}\right)\right)^{2}=36 \mathbb{E}\left[W_{t}(h, g) \mid \mathfrak{F}_{t, h}\right] .
$$

By Freedman's inequality, with probability at least $1-\delta$,

$$
\left|\sum_{t=1}^{k} W_{t}(h, g)-\sum_{t=1}^{k}\left[\left(g_{h}-Q_{h}^{\star}\right)\left(s_{h}^{t}, a_{h}^{t}\right)\right]^{2}\right| \leq \mathcal{O}\left(\sqrt{\log (1 / \delta) \sum_{t=1}^{k}\left[\left(g_{h}-Q_{h}^{\star}\right)\left(s_{h}^{t}, a_{h}^{t}\right)\right]^{2}}+\log (1 / \delta)\right)
$$

By taking a union bound over $[K] \times[H] \times \mathcal{V}_{\rho}$ and the non-negativity of $\sum_{t=1}^{k}\left[\left(g_{h}-Q_{h}^{\star}\right)\left(s_{h}^{t}, a_{h}^{t}\right)\right]^{2}$, we obtain that with probability at least $1-\delta$, for all $(k, h, \psi) \in[K] \times[H] \times \mathcal{V}_{\rho}$

$$
-\sum_{t=1}^{k} W_{t}(h, \psi) \leq \mathcal{O}(\iota)
$$

where $\iota=\log \left(H K\left|\mathcal{V}_{\rho}\right| / \delta\right)$. This directly implies for all $(k, h, g) \in[K] \times[H] \times \mathcal{G}$

$$
\begin{aligned}
& \sum_{t=1}^{k-1}\left[Q_{h}^{\star}\left(s_{h}^{t}, a_{h}^{t}\right)-r_{h}^{t}-Q_{h+1}^{\star}\left(s_{h+1}^{t}, \pi_{Q^{\star}}\left(s_{h+1}^{t}\right)\right)\right]^{2} \\
\leq & \sum_{t=1}^{k-1}\left[g_{h}\left(s_{h}^{t}, a_{h}^{t}\right)-r_{h}^{t}-Q_{h+1}^{\star}\left(s_{h+1}^{t}, \pi_{Q^{\star}}\left(s_{h+1}^{t}\right)\right)\right]^{2}+\mathcal{O}(\iota+k \rho) .
\end{aligned}
$$

Finally, by recalling the definition of $\mathcal{B}^{k}$, we conclude that with probability at least $1-\delta, Q^{\star} \in \mathcal{B}^{k}$ for all $k \in[K]$.

## E. 4 Proof of Lemma 41

The proof in this subsection basically follows the same arguments as in Appendix C of [18]. We firstly prove the following proposition which bounds the number of times $\left|\mathbb{E}_{\mu_{t}}\left[\phi_{t}\right]\right|$ can exceed a certain threshold.
Proposition 43. Given a function class $\Phi$ defined on $\mathcal{X}$, and a family of probability measures $\Pi$ over $\mathcal{X}$. Suppose sequence $\left\{\phi_{k}\right\}_{k=1}^{K} \subset \Phi$ and $\left\{\mu_{k}\right\}_{k=1}^{K} \subset \Pi$ satisfy that for all $k \in[K]$, $\sum_{t=1}^{k-1}\left(\mathbb{E}_{\mu_{t}}\left[\phi_{k}\right]\right)^{2} \leq \beta$. Then for all $k \in[K]$,

$$
\sum_{t=1}^{k} \mathbf{1}\left\{\left|\mathbb{E}_{\mu_{t}}\left[\phi_{t}\right]\right|>\epsilon\right\} \leq\left(\frac{\beta}{\epsilon^{2}}+1\right) \operatorname{dim}_{\mathrm{DE}}(\Phi, \Pi, \epsilon)
$$

Proof of Proposition 43. We first show that if for some $k$ we have $\left|\mathbb{E}_{\mu_{k}}\left[\phi_{k}\right]\right|>\epsilon$, then $\mu_{k}$ is $\epsilon$ dependent on at most $\beta / \epsilon^{2}$ disjoint subsequences in $\left\{\mu_{1}, \ldots, \mu_{k-1}\right\}$. By definition of DE dimension, if $\left|\mathbb{E}_{\mu_{k}}\left[\phi_{k}\right]\right|>\epsilon$ and $\mu_{k}$ is $\epsilon$-dependent on a subsequence $\left\{\nu_{1}, \ldots, \nu_{\ell}\right\}$ of $\left\{\mu_{1}, \ldots, \mu_{k-1}\right\}$, then we should have $\sum_{t=1}^{\ell}\left(\mathbb{E}_{\nu_{t}}\left[\phi_{k}\right]\right)^{2} \geq \epsilon^{2}$. It implies that if $\mu_{k}$ is $\epsilon$-dependent on $L$ disjoint subsequences in $\left\{\mu_{1}, \ldots, \mu_{k-1}\right\}$, we have

$$
\beta \geq \sum_{t=1}^{k-1}\left(\mathbb{E}_{\mu_{t}}\left[\phi_{k}\right]\right)^{2} \geq L \epsilon^{2}
$$

resulting in $L \leq \beta / \epsilon^{2}$.
Now we want to show that for any sequence $\left\{\nu_{1}, \ldots, \nu_{\kappa}\right\} \subseteq \Pi$, there exists $j \in[\kappa]$ such that $\nu_{j}$ is $\epsilon$-dependent on at least $L=\left\lceil(\kappa-1) / \operatorname{dim}_{\mathrm{DE}}(\Phi, \Pi, \epsilon)\right\rceil$ disjoint subsequences in $\left\{\nu_{1}, \ldots, \nu_{j-1}\right\}$. We argue by the following mental procedure: we start with singleton sequences $B_{1}=\left\{\nu_{1}\right\}, \ldots, B_{L}$ $=\left\{\nu_{L}\right\}$ and $j=L+1$. For each $j$, if $\nu_{j}$ is $\epsilon$-dependent on $B_{1}, \ldots, B_{L}$ we already achieved our goal so we stop; otherwise, we pick an $i \in[L]$ such that $\nu_{j}$ is $\epsilon$-independent of $B_{i}$ and update $B_{i}=B_{i} \cup\left\{\nu_{j}\right\}$. Then we increment $j$ by 1 and continue this process. By the definition of DE dimension, the size of each $B_{1}, \ldots, B_{L}$ cannot get bigger than $\operatorname{dim}_{\mathrm{DE}}(\Phi, \Pi, \epsilon)$ at any point in this process. Therefore, the process stops before or on $j=L \operatorname{dim}_{\mathrm{DE}}(\Phi, \Pi, \epsilon)+1 \leq \kappa$.
Fix $k \in[K]$ and let $\left\{\nu_{1}, \ldots, \nu_{\kappa}\right\}$ be subsequence of $\left\{\mu_{1}, \ldots, \mu_{k}\right\}$, consisting of elements for which $\left|\mathbb{E}_{\mu_{t}}\left[\phi_{t}\right]\right|>\epsilon$. Using the first claim, we know that each $\nu_{j}$ is $\epsilon$-dependent on at most $\beta / \epsilon^{2}$ disjoint
subsequences of $\left\{\nu_{1}, \ldots, \nu_{j-1}\right\}$. Using the second claim, we know there exists $j \in[\kappa]$ such that $\nu_{j}$ is $\epsilon$-dependent on at least $\left(\kappa / \operatorname{dim}_{\mathrm{DE}}(\Phi, \Pi, \epsilon)\right)-1$ disjoint subsequences of $\left\{\nu_{1}, \ldots, \nu_{j-1}\right\}$. Therefore, we have $\kappa / \operatorname{dim}_{\mathrm{DE}}(\Phi, \Pi, \epsilon)-1 \leq \beta / \epsilon^{2}$ which results in

$$
\kappa \leq\left(\frac{\beta}{\epsilon^{2}}+1\right) \operatorname{dim}_{\mathrm{DE}}(\Phi, \Pi, \epsilon)
$$

and completes the proof.
Proof of Lemma 41. Fix $k \in[K]$; let $d=\operatorname{dim}_{D E}(\Phi, \Pi, \omega)$. Sort the sequence $\left\{\left|\mathbb{E}_{\phi_{1}}\left[\phi_{1}\right]\right|, \ldots\right.$, $\left.\left|\mathbb{E}_{\mu_{k}}\left[\phi_{k}\right]\right|\right\}$ in a decreasing order and denote it by $\left\{e_{1}, \ldots, e_{k}\right\}\left(e_{1} \geq e_{2} \geq \cdots \geq e_{k}\right)$.

$$
\sum_{t=1}^{k}\left|\mathbb{E}_{\mu_{t}}\left[\phi_{t}\right]\right|=\sum_{t=1}^{k} e_{t}=\sum_{t=1}^{k} e_{t} \mathbf{1}\left\{e_{t} \leq \omega\right\}+\sum_{t=1}^{k} e_{t} \mathbf{1}\left\{e_{t}>\omega\right\} \leq k \omega+\sum_{t=1}^{k} e_{t} \mathbf{1}\left\{e_{t}>\omega\right\}
$$

For $t \in[k]$, we want to prove that if $e_{t}>\omega$, then we have $e_{t} \leq \min \left\{\sqrt{\frac{d \beta}{t-d}}, C\right\}$. Assume $t \in[k]$ satisfies $e_{t}>\omega$. Then there exists $\alpha$ such that $e_{t}>\alpha \geq \omega$. By Proposition 43, we have

$$
t \leq \sum_{i=1}^{k} \mathbf{1}\left\{e_{i}>\alpha\right\} \leq\left(\frac{\beta}{\alpha^{2}}+1\right) \operatorname{dim}_{\mathrm{DE}}(\Phi, \Pi, \alpha) \leq\left(\frac{\beta}{\alpha^{2}}+1\right) \operatorname{dim}_{\mathrm{DE}}(\Phi, \Pi, \omega)
$$

which implies $\alpha \leq \sqrt{\frac{d \beta}{t-d}}$. Besides, recall $e_{t} \leq C$, so we have $e_{t} \leq \min \left\{\sqrt{\frac{d \beta}{t-d}}, C\right\}$.
Finally, we have

$$
\begin{aligned}
\sum_{t=1}^{k} e_{t} \mathbf{1}\left\{e_{t}>\omega\right\} \leq \min \{d, k\} C+\sum_{t=d+1}^{k} \sqrt{\frac{d \beta}{t-d}} & \leq \min \{d, k\} C+\sqrt{d \beta} \int_{0}^{k} \frac{1}{\sqrt{t}} d t \\
& \leq \min \{d, k\} C+2 \sqrt{d \beta k}
\end{aligned}
$$

which completes the proof.

## F Proofs for OLIVE

In this section, we provide the formal proof for the results stated in Appendix A.

## F. 1 Full proof of Theorem 18

Proof of Theorem 18. By standard concentration arguments (Hoeffding's inequality plus union bound argument), with probability at least $1-\delta$, the following events hold for the first $d H+1$ phases (please refer to Appendix F. 2 for the proof)

1. If the elimination procedure is activated at the $h^{\text {th }}$ step in the $k^{\text {th }}$ phase, then $\mathcal{E}\left(f^{k}, \pi^{k}, h\right)>\zeta_{\text {act }} / 2$ and all $f \in \mathcal{F}$ satisfying $\left|\mathcal{E}\left(f, \pi^{k}, h\right)\right| \geq 2 \zeta_{\text {elim }}$ get eliminated.
2. If the elimination procedure is not activated in the $k^{\text {th }}$ phase, then, $\sum_{h=1}^{H} \mathcal{E}\left(f^{k}, \pi^{k}, h\right)<$ $2 H \zeta_{\text {act }}=4 \epsilon$.
3. $Q^{\star}$ is not eliminated.

Therefore, if we can show OLIVE terminates within $d H+1$ phases, then with high probability the output policy is $4 \epsilon$-optimal by the optimism of $f^{k}$ and simple policy loss decomposition (e.g. Lemma 1 in [1]):

$$
\begin{equation*}
\left(V_{1}^{\star}\left(s_{1}\right)-V_{1}^{\pi^{k}}\left(s_{1}\right)\right) \leq \max _{a} f^{k}\left(s_{1}, a\right)-V^{\pi^{k}}\left(s_{1}\right)=\sum_{h=1}^{H} \mathcal{E}\left(f^{k}, \pi^{k}, h\right) \leq 4 \epsilon \tag{17}
\end{equation*}
$$

In order to prove that OLIVE terminates within $d H+1$ phases, it suffices to show that for each $h \in[H]$, we can activate the elimination procedure at the $h^{\text {th }}$ step for at most $d$ times.

For the sake of contradiction, assume that Olive does not terminate in $d H+1$ phases. Within these $d H+1$ phases, there exists some $h \in[H]$ for which the activation process has been activated for at least $d+1$ times. Denote by $k_{1}<\cdots<k_{d+1} \leq d H+1$ the indices of the phases where the elimination is activated at the $h^{\text {th }}$ step. By the high-probability events, for all $i<j \leq d+1$, we have $\left|\mathcal{E}\left(f^{k_{j}}, \pi^{k_{i}}, h\right)\right|<2 \zeta_{\text {elim }}$ and for all $l \leq d+1$, we have $\mathcal{E}\left(f^{k_{l}}, \pi^{k_{l}}, h\right)>\zeta_{\text {act }} / 2$. This means for all $l \leq d+1$, we have both $\sqrt{\sum_{i=1}^{l-1}\left(\mathcal{E}\left(f^{k_{l}}, \pi^{k_{i}}, h\right)\right)^{2}}<\sqrt{d} \times 2 \zeta_{\text {elim }}=\epsilon / H$ and $\mathcal{E}\left(f^{k_{l}}, \pi^{k_{l}}, h\right)>\zeta_{\text {act }} / 2=\epsilon / H$. Therefore, the roll-in distribution of $\pi^{k_{1}}, \ldots, \pi^{k_{d+1}}$ at step $h$ is an $\epsilon / H$-independent sequence of length $d+1$, which contradicts with the definition of BE dimension. So Olive should terminate within $d H+1$ phases.

In sum, with probability at least $1-\delta$, Algorithm 2 will terminate and output a $4 \epsilon$-optimal policy using at most

$$
(d H+1)\left(n_{\mathrm{act}}+n_{\mathrm{elim}}\right) \leq \frac{3 c H^{3} d^{2} \log \left(\mathcal{N}\left(\mathcal{F}, \zeta_{\mathrm{elim}} / 8\right)\right) \cdot \iota}{\epsilon^{2}}
$$

episodes.

## F. 2 Concentration arguments for Theorem 18

Recall in Algorithm 2 we choose

$$
\zeta_{\text {act }}=\frac{2 \epsilon}{H}, \zeta_{\mathrm{elim}}=\frac{\epsilon}{2 H \sqrt{d}}, n_{\mathrm{act}}=\frac{c H^{2} \iota}{\epsilon^{2}}, \text { and } n_{\mathrm{elim}}=\frac{c H^{2} d \log \left(\mathcal{N}\left(\mathcal{F}, \zeta_{\mathrm{elim}} / 8\right)\right) \cdot \iota}{\epsilon^{2}}
$$

where $d=\max _{h \in[H]} \operatorname{dim}_{\mathrm{BE}}\left(\mathcal{F}, \mathcal{D}_{\mathcal{F}, h}, \epsilon / H\right), \iota=\log [H d / \delta \epsilon]$ and $c$ is a large absolute constant.Our goal is to prove with probability at least $1-\delta$, the following events hold for the first $d H+1$ phases

1. If the elimination procedure is activated at the $h^{\text {th }}$ step in the $k^{\text {th }}$ phase, then $\mathcal{E}\left(f^{k}, \pi^{k}, h\right)>\zeta_{\text {act }} / 2$ and all $f \in \mathcal{F}$ satisfying $\left|\mathcal{E}\left(f, \pi^{k}, h\right)\right| \geq 2 \zeta_{\text {elim }}$ get eliminated.
2. If the elimination procedure is not activated in the $k^{\text {th }}$ phase, then, $\sum_{h=1}^{H} \mathcal{E}\left(f^{k}, \pi^{k}, h\right)<$ $2 H \zeta_{\text {act }}=4 \epsilon$.
3. $Q^{\star}$ is not eliminated.

We begin with the activation procedure.
Concentration in the activation procedure Consider a fixed $(k, h) \in[d H+1] \times[H]$ pair. By Azuma-Hoefdding's inequality, with probability at least $1-\frac{\delta}{8 H\left(d H^{2}+1\right)}$, we have

$$
\left|\hat{\mathcal{E}}\left(f^{k}, \pi^{k}, h\right)-\mathcal{E}\left(f^{k}, \pi^{k}, h\right)\right| \leq \mathcal{O}\left(\sqrt{\frac{\iota}{n_{\mathrm{act}}}}\right) \leq \frac{\epsilon}{2 H} \leq \zeta_{\mathrm{act}} / 4
$$

where the second inequality follows from $n_{\text {act }}=C \frac{H^{2} \iota}{\epsilon^{2}}$ with $C$ being chosen large enough.
Take a union bound for all $(k, h) \in[d H+1] \times[H]$, we have with probability at least $1-\delta / 4$, the following holds for all $(k, h) \in[d H+1] \times[H]$

$$
\left|\hat{\mathcal{E}}\left(f^{k}, \pi^{k}, h\right)-\mathcal{E}\left(f^{k}, \pi^{k}, h\right)\right| \leq \zeta_{\text {act }} / 4
$$

By Algorithm 2, if the elimination procedure is not activated in the $k^{\text {th }}$ phase, we have $\sum_{h=1}^{H} \hat{\mathcal{E}}\left(f^{k}, \pi^{k}, h\right) \leq H \zeta_{\text {act }}$. Combine it with the concentration argument we just proved,

$$
\sum_{h=1}^{H} \mathcal{E}\left(f^{k}, \pi^{k}, h\right) \leq \sum_{h=1}^{H} \hat{\mathcal{E}}\left(f^{k}, \pi^{k}, h\right)+\frac{H \zeta_{\mathrm{act}}}{4}<\frac{5 H \zeta_{\mathrm{act}}}{4}
$$

On the other hand, if the elimination procedure is activated at the $h^{\text {th }}$ step in the $k^{\text {th }}$ phase, then $\hat{\mathcal{E}}\left(f^{k}, \pi^{k}, h\right)>\zeta_{\text {act }}$. Again combine it with the concentration argument we just proved,

$$
\mathcal{E}\left(f^{k}, \pi^{k}, h\right) \geq \hat{\mathcal{E}}\left(f^{k}, \pi^{k}, h\right)-\frac{\zeta_{\text {act }}}{4}>\frac{3 \zeta_{\mathrm{act}}}{4}
$$

Concentration in the elimination procedure Now, let us turn to the elimination procedure. First, let $\mathcal{Z}$ be an $\zeta_{\text {elim }} / 8$-cover of $\mathcal{F}$ with cardinality $\mathcal{N}\left(\mathcal{F}, \zeta_{\text {elim }} / 8\right)$. With a little abuse of notation, for every $f \in \mathcal{F}$, define $\hat{f}=\operatorname{argmin}_{g \in \mathcal{Z}} \max _{h \in[H]}\left\|f_{h}-g_{h}\right\|_{\infty}$. By applying Azuma-Hoeffding's inequality to all $(k, g) \in[d H+1] \times \mathcal{Z}$ and taking a union bound, we have with probability at least $1-\delta / 4$, the following holds for all $(k, g) \in[d H+1] \times \mathcal{Z}$

$$
\left|\hat{\mathcal{E}}\left(g, \pi^{k}, h_{k}\right)-\mathcal{E}\left(g, \pi^{k}, h_{k}\right)\right| \leq \zeta_{\mathrm{elim}} / 4
$$

Recall that Algorithm 2 eliminates all $f$ satisfying $\left|\hat{\mathcal{E}}\left(f, \pi^{k}, h_{k}\right)\right|>\zeta_{\text {elim }}$ when the elimination procedure is activated at the $h_{k}^{\text {th }}$ step in the $k^{\text {th }}$ phase. Therefore, if $\left|\mathcal{E}\left(f, \pi^{k}, h_{k}\right)\right| \geq 2 \zeta_{\text {elim }}$, $f$ will be eliminated because

$$
\begin{aligned}
\left|\hat{\mathcal{E}}\left(f, \pi^{k}, h_{k}\right)\right| & \geq\left|\hat{\mathcal{E}}\left(\hat{f}, \pi^{k}, h_{k}\right)\right|-2 \times \frac{\zeta_{\text {elim }}}{8} \\
& \geq\left|\mathcal{E}\left(\hat{f}, \pi^{k}, h_{k}\right)\right|-\frac{\zeta_{\text {elim }}}{2} \\
& \geq\left|\mathcal{E}\left(f, \pi^{k}, h_{k}\right)\right|-\frac{\zeta_{\text {elim }}}{2}-2 \times \frac{\zeta_{\text {elim }}}{8}>\zeta_{\text {elim }}
\end{aligned}
$$

Finally, note that $\mathcal{E}\left(Q^{\star}, \pi, h\right) \equiv 0$ for any $\pi$ and $h$. As a result, it will never be eliminated within the first $d H+1$ phases because we can similarly prove

$$
\left|\hat{\mathcal{E}}\left(Q^{\star}, \pi^{k}, h_{k}\right)\right| \leq\left|\mathcal{E}\left(Q^{\star}, \pi^{k}, h_{k}\right)\right|+\frac{3 \zeta_{\text {elim }}}{4}<\zeta_{\text {elim }} .
$$

Wrapping up: take a union bound for the activation and elimination procedure, and conclude that the three events, listed at the beginning of this section, hold for the the first $d H+1$ phases with probability at least $1-\delta / 2$.

## G Proofs for V-type Variants

In this section, we provide formal proofs for the results stated in Section B.

## G. 1 Proof of Theorem 23

The proof is similar to that in Appendix F.

Proof of Theorem 23. By standard concentration arguments (Hoeffding's inequality, Bernstein's inequality, and union bound argument), with probability at least $1-\delta$, the following events hold for the first $d H+1$ phases (please refer to Appendix G.1. 1 for the proof)

1. If the elimination procedure is activated at the $h^{\text {th }}$ step in the $k^{\text {th }}$ phase, then $\mathcal{E}_{\mathrm{V}}\left(f^{k}, \pi^{k}, h\right)>\zeta_{\text {act }} / 2$ and all $f \in \mathcal{F}$ satisfying $\left|\mathcal{E}_{\mathrm{V}}\left(f, \pi^{k}, h\right)\right| \geq 2 \zeta_{\text {elim }}$ get eliminated.
2. If the elimination procedure is not activated in the $k^{\text {th }}$ phase, then, $\sum_{h=1}^{H} \mathcal{E}_{\mathrm{V}}\left(f^{k}, \pi^{k}, h\right)<$ $2 H \zeta_{\text {act }}=4 \epsilon$.
3. $Q^{\star}$ is not eliminated.

Therefore, if we can show OLIVE terminates within $d H+1$ phases, then with high probability the output policy is $4 \epsilon$-optimal by the optimism of $f^{k}$ and simple policy loss decomposition (e.g., Lemma 1 in [1]):

$$
\begin{equation*}
\left(V_{1}^{\star}\left(s_{1}\right)-V_{1}^{\pi^{k}}\left(s_{1}\right)\right) \leq \max _{a} f^{k}\left(s_{1}, a\right)-V^{\pi^{k}}\left(s_{1}\right)=\sum_{h=1}^{H} \mathcal{E}_{\mathrm{V}}\left(f^{k}, \pi^{k}, h\right) \leq 4 \epsilon \tag{18}
\end{equation*}
$$

In order to prove that Olive terminates within $d H+1$ phases, it suffices to show that for each $h \in[H]$, we can activate the elimination procedure at the $h^{\text {th }}$ step for at most $d$ times.

For the sake of contradiction, assume that Olive does not terminate in $d H+1$ phases. Within these $d H+1$ phases, there exists some $h \in[H]$ for which the activation process has been activated for at least $d+1$ times. Denote by $k_{1}<\cdots<k_{d+1} \leq d H+1$ the indices of the phases where the elimination is activated at the $h^{\text {th }}$ step. By the high-probability events, for all $i<j \leq d+1$, we have $\left|\mathcal{E}_{\mathrm{V}}\left(f^{k_{j}}, \pi^{k_{i}}, h\right)\right|<2 \zeta_{\text {elim }}$ and for all $l \leq d+1$, we have $\mathcal{E}_{\mathrm{V}}\left(f^{k_{l}}, \pi^{k_{l}}, h\right)>\zeta_{\text {act }} / 2$. This means for all $l \leq d+1$, we have both $\sqrt{\sum_{i=1}^{l-1}\left(\mathcal{E}_{\mathrm{V}}\left(f^{k_{l}}, \pi^{k_{i}}, h\right)\right)^{2}}<\sqrt{d} \times 2 \zeta_{\text {elim }}=\epsilon / H$ and $\mathcal{E}_{\mathrm{V}}\left(f^{k_{l}}, \pi^{k_{l}}, h\right)>\zeta_{\text {act }} / 2=\epsilon / H$. Therefore, the roll-in distribution of $\pi^{k_{1}}, \ldots, \pi^{k_{d+1}}$ at step $h$ is an $\epsilon / H$-independent sequence of length $d+1$ with respect to $\left(I-\mathcal{T}_{h}\right) V_{\mathcal{F}}$, which contradicts with the definition of BE dimension. So OLIVE should terminate within $d H+1$ phases.

In sum, with probability at least $1-\delta$, Algorithm 2 will terminate and output a $4 \epsilon$-optimal policy using at most

$$
(d H+1)\left(n_{\mathrm{act}}+n_{\mathrm{elim}}\right) \leq \frac{3 c H^{3} d^{2}|\mathcal{A}| \log (|\mathcal{F}|) \cdot \iota}{\epsilon^{2}}
$$

episodes.

## G.1.1 Concentration arguments for Theorem 23

Recall in Algorithm 4 we choose

$$
\zeta_{\text {act }}=\frac{2 \epsilon}{H}, \zeta_{\mathrm{elim}}=\frac{\epsilon}{2 H \sqrt{d}}, n_{\text {act }}=\frac{c H^{2} \iota}{\epsilon^{2}}, \text { and } n_{\mathrm{elim}}=\frac{c|\mathcal{A}| H^{2} d \log \left(\mathcal{N}\left(\mathcal{F}, \zeta_{\mathrm{elim}} / 8\right)\right) \cdot \iota}{\epsilon^{2}}
$$

where $d=\max _{h \in[H]} \operatorname{dim}_{\operatorname{VBE}}\left(\mathcal{F}, \mathcal{D}_{\mathcal{F}, h}, \epsilon / H\right), \iota=\log [H d / \delta \epsilon]$ and $c$ is a large absolute constant. Our goal is to prove with probability at least $1-\delta$, the following events hold for the first $d H+1$ phases

1. If the elimination procedure is activated at the $h^{\text {th }}$ step in the $k^{\text {th }}$ phase, then $\mathcal{E}_{\mathrm{V}}\left(f^{k}, \pi^{k}, h\right)>\zeta_{\text {act }} / 2$ and all $f \in \mathcal{F}$ satisfying $\left|\mathcal{E}_{\mathrm{V}}\left(f, \pi^{k}, h\right)\right| \geq 2 \zeta_{\text {elim }}$ get eliminated.
2. If the elimination procedure is not activated in the $k^{\text {th }}$ phase, then, $\sum_{h=1}^{H} \mathcal{E}_{\mathrm{V}}\left(f^{k}, \pi^{k}, h\right)<$ $2 H \zeta_{\text {act }}=4 \epsilon$.
3. $Q^{\star}$ is not eliminated.

We begin with the activation procedure.
Concentration in the activation procedure Consider a fixed $(k, h) \in[d H+1] \times[H]$ pair. By Azuma-Hoefdding's inequality, with probability at least $1-\frac{\delta}{8 H(d H+1)}$, we have

$$
\left|\tilde{\mathcal{E}}_{\mathrm{V}}\left(f^{k}, \pi^{k}, h\right)-\mathcal{E}_{\mathrm{V}}\left(f^{k}, \pi^{k}, h\right)\right| \leq \mathcal{O}\left(\sqrt{\frac{\iota}{n_{\mathrm{act}}}}\right) \leq \frac{\epsilon}{2 H} \leq \zeta_{\mathrm{act}} / 4
$$

where the second inequality follows from $n_{\text {act }}=C \frac{H^{2} \iota}{\epsilon^{2}}$ with $C$ being chosen large enough.
Take a union bound for all $(k, h) \in[d H+1] \times[H]$, we have with probability at least $1-\delta / 4$, the following holds for all $(k, h) \in[d H+1] \times[H]$

$$
\left|\tilde{\mathcal{E}}_{\mathrm{V}}\left(f^{k}, \pi^{k}, h\right)-\mathcal{E}_{\mathrm{V}}\left(f^{k}, \pi^{k}, h\right)\right| \leq \zeta_{\mathrm{act}} / 4
$$

By Algorithm 4, if the elimination procedure is not activated in the $k^{\text {th }}$ phase, we have $\sum_{h=1}^{H} \tilde{\mathcal{E}}_{\mathrm{V}}\left(f^{k}, \pi^{k}, h\right) \leq H \zeta_{\text {act }}$. Combine it with the concentration argument we just proved,

$$
\sum_{h=1}^{H} \mathcal{E}_{\mathrm{V}}\left(f^{k}, \pi^{k}, h\right) \leq \sum_{h=1}^{H} \tilde{\mathcal{E}}_{\mathrm{V}}\left(f^{k}, \pi^{k}, h\right)+\frac{H \zeta_{\mathrm{act}}}{4} \leq \frac{5 H \zeta_{\mathrm{act}}}{4}
$$

On the other hand, if the elimination procedure is activated at the $h^{\text {th }}$ step in the $k^{\text {th }}$ phase, then $\tilde{\mathcal{E}}_{\mathrm{V}}\left(f^{k}, \pi^{k}, h\right)>\zeta_{\text {act }}$. Again combine it with the concentration argument we just proved,

$$
\mathcal{E}_{\mathrm{V}}\left(f^{k}, \pi^{k}, h\right) \geq \tilde{\mathcal{E}}_{\mathrm{V}}\left(f^{k}, \pi^{k}, h\right)-\frac{\zeta_{\mathrm{act}}}{4}>\frac{3 \zeta_{\mathrm{act}}}{4}
$$

Concentration in the elimination procedure Now, let us turn to the elimination procedure. We start by bounding the the second moment of

$$
\frac{\mathbf{1}\left[\pi_{f}\left(s_{h}\right)=a_{h}\right]}{1 /|\mathcal{A}|}\left(f_{h}\left(s_{h}, a_{h}\right)-r_{h}-\max _{a^{\prime} \in \mathcal{A}} f_{h+1}\left(s_{h+1}, a^{\prime}\right)\right)
$$

for all $f \in \mathcal{F}$. Let $y\left(s_{h}, a_{h}, r_{h}, s_{h+1}\right)=f_{h}\left(s_{h}, a_{h}\right)-r_{h}-\max _{a^{\prime} \in \mathcal{A}} f_{h+1}\left(s_{h+1}, a^{\prime}\right) \in[-2,1]$, then we have

$$
\begin{aligned}
& \mathbb{E}\left[\left(|\mathcal{A}| \mathbf{1}\left[\pi_{f}\left(s_{h}\right)=a_{h}\right] y\left(s_{h}, a_{h}, r_{h}, s_{h+1}\right)\right)^{2} \mid s_{h} \sim \pi^{k}, a_{h} \sim \operatorname{Uniform}(\mathcal{A})\right] \\
\leq & 4|\mathcal{A}|^{2} \mathbb{E}\left[\mathbf{1}\left[\pi_{f}\left(s_{h}\right)=a_{h}\right] \mid s_{h} \sim \pi^{k}, a_{h} \sim \operatorname{Uniform}(\mathcal{A})\right]=4|\mathcal{A}| .
\end{aligned}
$$

For a fixed $(k, f) \in[d H+1] \times \mathcal{F}$, by applying Azuma-Bernstein's inequality, with probability at least $1-\frac{\delta}{8(d H+1)|\mathcal{F}|}$ we have

$$
\left|\hat{\mathcal{E}}_{\mathrm{V}}\left(f, \pi^{k}, h_{k}\right)-\mathcal{E}_{\mathrm{V}}\left(f, \pi^{k}, h_{k}\right)\right| \leq \mathcal{O}\left(\sqrt{\frac{|\mathcal{A}| \iota^{\prime}}{n_{\mathrm{elim}}}}+\frac{|\mathcal{A}| \iota^{\prime}}{n_{\mathrm{elim}}}\right) \leq \mathcal{O}\left(\sqrt{\frac{|\mathcal{A}| \iota^{\prime}}{n_{\mathrm{elim}}}}\right) \leq \zeta_{\mathrm{elim}} / 2
$$

where $\iota^{\prime}=\log [8(d H+1)|\mathcal{F}| / \delta]$, and the third inequality follows from $n_{\text {elim }}=C|\mathcal{A}| \iota / \zeta_{\text {elim }}^{2}$ with $C$ being chosen large enough.

Taking a union bound over $[d H+1] \times \mathcal{F}$, we have with probability at least $1-\delta / 4$, the following holds for all $(k, f) \in[d H+1] \times \mathcal{F}$

$$
\left|\hat{\mathcal{E}}_{\mathrm{V}}\left(f, \pi^{k}, h_{k}\right)-\mathcal{E}_{\mathrm{V}}\left(f, \pi^{k}, h_{k}\right)\right| \leq \zeta_{\text {elim }} / 2
$$

Recall that Algorithm 4 eliminates all $f$ satisfying $\left|\hat{\mathcal{E}}_{\mathrm{V}}\left(f, \pi^{k}, h_{k}\right)\right|>\zeta_{\text {elim }}$ when the elimination procedure is activated at the $h_{k}^{\text {th }}$ step in the $k^{\text {th }}$ phase. Therefore, if $\left|\mathcal{E}_{\mathrm{V}}\left(f, \pi^{k}, h_{k}\right)\right| \geq 2 \zeta_{\text {elim }}$, $f$ will be eliminated because

$$
\left|\hat{\mathcal{E}}_{\mathrm{V}}\left(f, \pi^{k}, h_{k}\right)\right| \geq\left|\mathcal{E}_{\mathrm{V}}\left(f, \pi^{k}, h_{k}\right)\right|-\frac{\zeta_{\mathrm{elim}}}{2}>\zeta_{\mathrm{elim}}
$$

Finally, note that $\mathcal{E}_{\mathrm{V}}\left(Q^{\star}, \pi, h\right) \equiv 0$ for any $\pi$ and $h$. As a result, it will never be eliminated within the first $d H+1$ phases because we can similarly prove

$$
\left|\hat{\mathcal{E}}_{\mathrm{V}}\left(Q^{\star}, \pi^{k}, h_{k}\right)\right| \leq\left|\mathcal{E}_{\mathrm{V}}\left(Q^{\star}, \pi^{k}, h_{k}\right)\right|+\frac{\zeta_{\text {elim }}}{2}<\zeta_{\text {elim }}
$$

Wrapping up: take a union bound for the activation and elimination procedure, and conclude that the three events, listed at the beginning of this section, hold for the the first $d H+1$ phases with probability at least $1-\delta / 2$.

## G. 2 Proof of Theorem 22

The proof is basically the same as that of Theorem 15 in Appendix E.
To begin with, we have the following lemma (akin to Lemma 39 and 40) showing that with high probability: $(i)$ any function in the confidence set has low Bellman-error over the collected Datasets $\mathcal{D}_{1}, \ldots, \mathcal{D}_{H}$ as well as the distributions from which $\mathcal{D}_{1}, \ldots, \mathcal{D}_{H}$ are sampled; (ii) the optimal value function is inside the confidence set. Its proof is almost identical to that of Lemma 39 and 40 which can be found in Appendix E.3.
Lemma 44 (Akin to Lemma 39 and 40). Let $\rho>0$ be an arbitrary fixed number. If we choose $\beta=c\left(\log \left[K H \mathcal{N}_{\mathcal{F} \cup \mathcal{G}}(\rho) / \delta\right]+K \rho\right)$ with some large absolute constant $c$ in Algorithm 3, then with probability at least $1-\delta$, for all $(k, h) \in[K] \times[H]$, we have
(a) $\sum_{i=1}^{k-1} \mathbb{E}\left[\left(f_{h}^{k}\left(s_{h}, a_{h}\right)-\left(\mathcal{T} f_{h+1}^{k}\right)\left(s_{h}, a_{h}\right)\right)^{2} \mid s_{h} \sim \pi^{i}, a_{h} \sim \operatorname{Uniform}(\mathcal{A})\right] \leq \mathcal{O}(\beta)$,
(b) $\frac{1}{|\mathcal{A}|} \sum_{i=1}^{k-1} \sum_{a \in \mathcal{A}}\left(f_{h}^{k}\left(s_{h}^{i}, a\right)-\left(\mathcal{T} f_{h+1}^{k}\right)\left(s_{h}^{i}, a\right)\right)^{2} \leq \mathcal{O}(\beta)$,
(c) $Q^{\star} \in \mathcal{B}^{k}$,
where $s_{h}^{i}$ denotes the state at step $h$ collected according to Line 5 in Algorithm 3 following $\pi^{i}$.
Proof of Lemma 44. To prove inequality $(a)$, we only need to redefine the filtration $\mathfrak{F}_{t, h}$ in Appendix E.3.1 to be the filtration induced by $\left\{s_{1}^{i}, a_{1}^{i}, r_{1}^{i}, \ldots, s_{H}^{i}\right\}_{i=1}^{t-1}$ and repeat the arguments there verbatim.

To prove inequality $(b)$, we only need to redefine the filtration $\mathfrak{F}_{t, h}$ in Appendix E.3.1 to be the filtration induced by $\left\{s_{1}^{i}, a_{1}^{i}, r_{1}^{i}, \ldots, s_{H}^{i}\right\}_{i=1}^{t-1} \bigcup\left\{s_{1}^{t}, a_{1}^{t}, r_{1}^{t}, \ldots, s_{h}^{t}\right\}$ and repeat the arguments there verbatim.
The proof of $(c)$ is the same as that of Lemma 40 in Appendix E.3.2.
Step 1. Bounding the regret by Bellman error By Lemma 44 (c), we can upper bound the cumulative regret by the summation of Bellman error with probability at least $1-\delta$ :

$$
\begin{equation*}
\sum_{k=1}^{K}\left(V_{1}^{\star}\left(s_{1}\right)-V_{1}^{\pi^{k}}\left(s_{1}\right)\right) \leq \sum_{k=1}^{K}\left(\max _{a} f_{1}^{k}\left(s_{1}, a\right)-V_{1}^{\pi^{k}}\left(s_{1}\right)\right) \stackrel{(i)}{=} \sum_{k=1}^{K} \sum_{h=1}^{H} \mathcal{E}_{\mathrm{V}}\left(f^{k}, \pi^{k}, h\right) \tag{19}
\end{equation*}
$$

where $(i)$ follows from standard policy loss decomposition (e.g. Lemma 1 in [1]).
Step 2. Bounding cumulative Bellman error using DE dimension Next, we focus on a fixed step $h$ and bound the cumulative Bellman error $\sum_{k=1}^{K} \mathcal{E}_{\mathrm{V}}\left(f^{k}, \pi^{k}, h\right)$ using Lemma 44.
Invoking Lemma 44 (a) with

$$
\rho=\frac{\epsilon^{2}}{H^{2} \cdot \operatorname{dim}_{\mathrm{VBE}}\left(\mathcal{F}, \mathcal{D}_{\mathcal{F}}, \epsilon / H\right) \cdot|\mathcal{A}|}
$$

implies that with probability at least $1-\delta$, for all $(k, h) \in[K] \times[H]$, we have

$$
\sum_{i=1}^{k-1} \mathbb{E}\left[\left(f_{h}^{k}\left(s_{h}, \pi_{f_{h}^{k}}\left(s_{h}\right)\right)-\left(\mathcal{T} f_{h+1}^{k}\right)\left(s_{h}, \pi_{f_{h}^{k}}\left(s_{h}\right)\right)\right)^{2} \mid s_{h} \sim \pi^{i}\right] \leq \mathcal{O}(|\mathcal{A}| \beta)
$$

Further invoking Lemma 41 with

$$
\left\{\begin{array}{l}
\omega=\frac{\epsilon}{H}, C=1 \\
\mathcal{X}=\mathcal{S}, \Phi=\left(I-\mathcal{T}_{h}\right) V_{\mathcal{F}}, \Pi=\mathcal{D}_{\mathcal{F}, h} \\
\phi_{k}(s):=\left(f_{h}^{k}-\mathcal{T}_{h} f_{h+1}^{k}\right)\left(s, \pi_{f_{h}^{k}}(s)\right) \text { and } \mu_{k}=\mathbb{P}^{\pi^{k}}\left(s_{h}=\cdot\right)
\end{array}\right.
$$

we obtain

$$
\frac{1}{K} \sum_{t=1}^{K} \mathcal{E}_{\mathrm{V}}\left(f^{t}, \pi^{t}, h\right) \leq \mathcal{O}\left(\sqrt{\frac{\operatorname{dim}_{\mathrm{VBE}}\left(\mathcal{F}, \mathcal{D}_{\mathcal{F}}, \epsilon / H\right)|\mathcal{A}| \log \left[K H \mathcal{N}_{\mathcal{F} \cup \mathcal{G}}(\rho) / \delta\right]}{K}}+\frac{\epsilon}{H}\right)
$$

Plugging in the choice of $K$ completes the proof.
Similarly, for $\mathcal{D}_{\Delta}$, we can invoke Lemma 44 (b) witht

$$
\rho=\frac{\epsilon^{2}}{H^{2} \cdot \operatorname{dim}_{\mathrm{VBE}}\left(\mathcal{F}, \mathcal{D}_{\Delta}, \epsilon / H\right) \cdot|\mathcal{A}|}
$$

and Lemma 41 with

$$
\left\{\begin{array}{l}
\omega=\frac{\epsilon}{H}, C=1 \\
\mathcal{X}=\mathcal{S}, \Phi=\left(I-\mathcal{T}_{h}\right) V_{\mathcal{F}}, \Pi=\mathcal{D}_{\Delta, h} \\
\phi_{k}(s):=\left(f_{h}^{k}-\mathcal{T}_{h} f_{h+1}^{k}\right)\left(s, \pi_{f_{h}^{k}}(s)\right) \text { and } \mu_{k}=\mathbf{1}\left\{\cdot=s_{h}^{k}\right\}
\end{array}\right.
$$

and obtain

$$
\begin{aligned}
\frac{1}{K} \sum_{t=1}^{K} \mathcal{E}_{\mathrm{V}}\left(f^{t}, \pi^{t}, h\right) & \leq \frac{1}{K} \sum_{t=1}^{K}\left(f_{h}^{t}-\mathcal{T} f_{h+1}^{t}\right)\left(s_{h}^{t}, \pi_{f_{h}^{t}}\left(s_{h}^{t}\right)\right)+\mathcal{O}\left(\sqrt{\frac{\log K}{K}}\right) \\
& \leq \mathcal{O}\left(\sqrt{\frac{\operatorname{dim}_{\mathrm{VBE}}\left(\mathcal{F}, \mathcal{D}_{\Delta}, \epsilon / H\right)|\mathcal{A}| \log \left[K H \mathcal{N}_{\mathcal{F} \cup \mathcal{G}}(\rho) / \delta\right]}{K}}+\frac{\epsilon}{H}+\sqrt{\frac{\log K}{K}}\right)
\end{aligned}
$$

where the first inequality follows from standard martingale concentration.
Plugging in the choice of $K$ completes the proof.

## H Proofs for Examples

## H. 1 Proof of Proposition 29

Proof. Suppose $\mathcal{F}$ has finite $\epsilon$-effective dimension and denote the corresponding mapping by $\phi$. Then we can rewrite $\mathcal{F}$ in the form of $\mathcal{F}=\left\{f_{\theta}(\cdot)=\langle\phi(\cdot), \theta\rangle_{\mathcal{H}} \mid \theta \in \Theta\right\}$, where $\Theta \subset B_{\mathcal{H}}(1)$.

Suppose there exists an $\epsilon^{\prime}$-independent sequence $x_{1}^{\prime}, \ldots, x_{n}^{\prime} \in \mathcal{X}$ with respect to $\mathcal{F}$ where $\epsilon^{\prime} \geq \epsilon$. By the definition of independent sequence, this is equivalent to the existence of $\theta_{1}, \ldots, \theta_{n} \in(\Theta-\Theta)$ and $x_{1}, \ldots, x_{n} \in \phi(\mathcal{X})$ such that

$$
\left\{\begin{array}{l}
\sum_{i=1}^{t-1}\left(x_{i}^{\top} \theta_{t}\right)^{2} \leq \epsilon^{\prime 2}, \quad t \in[n]  \tag{20}\\
\left|x_{t}^{\top} \theta_{t}\right| \geq \epsilon^{\prime}, \quad t \in[n]
\end{array}\right.
$$

Define $\Sigma_{t}=\sum_{i=1}^{t-1} x_{i} x_{i}^{\top}+\frac{\epsilon^{\prime 2}}{4} \cdot$ I. We have

$$
\begin{equation*}
\left\|\theta_{t}\right\|_{\Sigma_{t}} \leq \sqrt{2} \epsilon^{\prime} \quad \Longrightarrow \quad \epsilon^{\prime} \leq\left|x_{t}^{\top} \theta_{t}\right| \leq\left\|\theta_{t}\right\|_{\Sigma_{t}} \cdot\left\|x_{t}\right\|_{\Sigma_{t}^{-1}} \leq \sqrt{2} \epsilon^{\prime}\left\|x_{t}\right\|_{\Sigma_{t}^{-1}}, \quad t \in[n] \tag{21}
\end{equation*}
$$

As a result, we should have $\left\|x_{t}\right\|_{\Sigma_{t}^{-1}}^{2} \geq 1 / 2$ for all $t \in[n]$. Now we can apply the standard logdeterminant argument,

$$
\sum_{t=1}^{n} \log \left(1+\left\|x_{t}\right\|_{\Sigma_{t}^{-1}}^{2}\right)=\log \left(\frac{\operatorname{det}\left(\Sigma_{n+1}\right)}{\operatorname{det}\left(\Sigma_{1}\right)}\right)=\log \operatorname{det}\left(\mathrm{I}+\frac{4}{\epsilon^{\prime 2}} \sum_{i=1}^{n} x_{i} x_{i}^{\top}\right)
$$

which implies

$$
\begin{equation*}
0.5 \leq \min _{t \in[n]}\left\|x_{t}\right\|_{\Sigma_{t}^{-1}}^{2} \leq \exp \left(\frac{1}{n} \log \operatorname{det}\left(\mathrm{I}+\frac{4}{\epsilon^{\prime 2}} \sum_{i=1}^{n} x_{i} x_{i}^{\top}\right)\right)-1 \tag{22}
\end{equation*}
$$

Choose $n=d_{\mathrm{eff}}(\mathcal{F}, \epsilon / 2)$ that is the minimum positive integer satisfying

$$
\begin{equation*}
\sup _{x_{1}, \ldots, x_{n} \in \phi(\mathcal{X})} \frac{1}{n} \log \operatorname{det}\left(\mathrm{I}+\frac{4}{\epsilon^{2}} \sum_{i=1}^{n} x_{i} x_{i}^{\top}\right) \leq e^{-1} \tag{23}
\end{equation*}
$$

This leads to a contradiction because $\epsilon^{\prime} \geq \epsilon$ and $0.5>e^{e^{-1}}-1$. So we must have

$$
\operatorname{dim}_{\mathrm{E}}(\mathcal{F}, \epsilon) \leq d_{\mathrm{eff}}(\mathcal{F}, \epsilon / 2)
$$

## H. 2 Proof of Proposition 32

Proof. Consider fixed $\epsilon \in \mathbb{R}^{+}$and $h \in[H]$, and denote $n=\operatorname{dim}_{E}(\mathcal{F}, \epsilon)$. Then by the definition of Eluder dimension, there must exist $x_{1}, \ldots, x_{n} \in \mathcal{X}_{h}$ where $\mathcal{X}_{h}=\left\{\phi_{h}(s, a):(s, a) \in \mathcal{S} \times \mathcal{A}\right\}$ so that for any $\theta, \theta^{\prime} \in B_{\mathcal{H}}(H+1-h)$, if $\sum_{i=1}^{n}\left(\left\langle x_{i}, \theta-\theta^{\prime}\right\rangle_{\mathcal{H}}\right)^{2} \leq \epsilon^{2}$, then $\left|\left\langle z, \theta-\theta^{\prime}\right\rangle_{\mathcal{H}}\right| \leq \epsilon$ for any $z \in \mathcal{X}_{h}$. In other words, $x_{1}, \ldots, x_{n}$ is one of the longest independent subsequences. Therefore, in order to cover $\mathcal{F}_{h}$, we only need cover the projection of $B_{\mathcal{H}}(H+1-h)$ onto the linear subspace spanned by $x_{1}, \ldots, x_{n}$, which is at most $n$ dimensional.
By standard $\epsilon$-net argument, there exists $\mathcal{C} \subset B_{\mathcal{H}}(H+1-h)$ such that: (a) $\log |\mathcal{C}| \leq \mathcal{O}(n \cdot \log (1+$ $n H / \epsilon)$ ), (b) for any $\theta \in B_{\mathcal{H}}(H+1-h)$, there exists $\hat{\theta} \in \mathcal{C}$ satisfying $\sum_{i=1}^{n}\left(\left\langle x_{i}, \theta-\hat{\theta}\right\rangle_{\mathcal{H}}\right)^{2} \leq \epsilon^{2}$. By the property of $x_{1}, \ldots, x_{n},\left\{\phi_{h}(\cdot, \cdot)^{\top} \hat{\theta} \mid \hat{\theta} \in \mathcal{C}\right\}$ is an $\epsilon$-cover of $\mathcal{F}_{h}$. Since $\mathcal{F}=\mathcal{F}_{1} \times \cdots \times \mathcal{F}_{H}$, we obtain $\log \mathcal{N}_{\mathcal{F}}(\epsilon) \leq \mathcal{O}(H n \cdot \log (1+n H / \epsilon))$. Finally, by Proposition 31, $n \leq d(\epsilon)$, which concludes the proof.

## H. 3 Proof of Proposition 34

Proof. Assume there exists $h \in[H]$ such that $\operatorname{dim}_{\text {DE }}\left(\left(I-\mathcal{T}_{h}\right) \mathcal{F}, \mathcal{D}_{\mathcal{F}, h}, \epsilon\right) \geq m$. Let $\mu_{1}, \ldots, \mu_{n} \in$ $\mathcal{D}_{\mathcal{F}, h}$ be a an $\epsilon$-independent sequence with respect to $\left(I-\mathcal{T}_{h}\right) \mathcal{F}$. By Definition 6 , there exist $f^{1}, \ldots, f^{n}$ such that for all $t \in[n], \sqrt{\sum_{i=1}^{t-1}\left(\mathbb{E}_{\mu_{i}}\left[f_{h}^{t}-\mathcal{T}_{h} f_{h+1}^{t}\right]\right)^{2}} \leq \epsilon$ and $\left|\mathbb{E}_{\mu_{t}}\left[f_{h}^{t}-\mathcal{T}_{h} f_{h+1}^{t}\right]\right|>\epsilon$. Since $\mu_{1}, \ldots, \mu_{n} \in \mathcal{D}_{\mathcal{F}, h}$, there exist $g^{1}, \ldots, g^{n} \in \mathcal{F}$ so that $\mu_{i}$ is generated by executing $\pi_{g^{i}}$, for all $i \in[n]$.
By the definition of effective Bellman rank, this is equivalent to: $\sqrt{\sum_{i=1}^{t-1}\left(\left\langle\phi_{h}\left(f^{t}\right), \psi_{h}\left(g^{i}\right)\right\rangle\right)^{2}} \leq \epsilon$ and $\left|\left\langle\phi_{h}\left(f^{t}\right), \psi_{h}\left(g^{t}\right)\right\rangle\right|>\epsilon$ for all $t \in[n]$. For notational simplicity, define $x_{i}=\psi_{h}\left(g^{i}\right)$ and $\theta_{i}=\phi_{h}\left(f^{i}\right)$. Then

$$
\left\{\begin{array}{l}
\sum_{i=1}^{t-1}\left(x_{i}^{\top} \theta_{t}\right)^{2} \leq \epsilon^{2}, \quad t \in[n]  \tag{24}\\
\left|x_{t}^{\top} \theta_{t}\right| \geq \epsilon, \quad t \in[n] .
\end{array}\right.
$$

The remaining arguments follow the same as in the proof of Proposition 29 except that we replace $\epsilon$ by $\epsilon / \zeta$.

## H. 4 Proof of Proposition 38

Proof. Note that the case $h=1$ is trivial because each episode always starts from a fixed initial state independent of the policy. For any policy $\pi$, function $f \in \mathcal{F}$, and step $h \geq 2$

$$
\begin{aligned}
\mathcal{E}_{\mathrm{V}}(f, \pi, h) & =\mathbb{E}\left[f_{h}\left(o_{h}, a_{h}\right)-r_{h}\left(o_{h}, a_{h}\right)-f_{h+1}\left(o_{h+1}, a_{h+1}\right) \mid s_{h} \sim \pi, a_{h: h+1} \sim \pi_{f}\right] \\
& =\mathbb{E}\left[f_{h}\left(o_{h}, a_{h}\right)-r_{h}\left(o_{h}, a_{h}\right)-f_{h+1}\left(o_{h+1}, a_{h+1}\right) \mid\left(s_{h-1}, a_{h-1}\right) \sim \pi, a_{h: h+1} \sim \pi_{f}\right] \\
& =\sum_{s, a \in \mathcal{S}} \sum_{s^{\prime} \in \mathcal{S}} \mathbb{P}^{\pi}\left(s_{h-1}=s, a_{h-1}=a\right) \cdot\left\langle\phi_{h-1}(s, a), \psi_{h-1}\left(s^{\prime}\right)\right\rangle_{\mathcal{H}} \cdot \mathcal{V}\left(s^{\prime}\right),
\end{aligned}
$$

where

$$
\mathcal{V}\left(s^{\prime}\right)=\mathbb{E}\left[f_{h}\left(o_{h}, a_{h}\right)-r_{h}\left(o_{h}, a_{h}\right)-f_{h+1}\left(o_{h+1}, a_{h+1}\right) \mid s_{h}=s^{\prime}, a_{h: h+1} \sim \pi_{f}\right] .
$$

As a result, we obtain

$$
\begin{aligned}
& \mathbb{E}\left[f_{h}\left(o_{h}, a_{h}\right)-r_{h}\left(o_{h}, a_{h}\right)-f_{h+1}\left(o_{h+1}, a_{h+1}\right) \mid s_{h} \sim \pi, a_{h: h+1} \sim \pi_{f}\right] \\
= & \left\langle\mathbb{E}_{\pi}\left[\phi_{h-1}\left(s_{h-1}, a_{h-1}\right)\right], \sum_{s^{\prime} \in \mathcal{S}} \psi_{h-1}\left(s^{\prime}\right) \mathcal{V}\left(s^{\prime}\right)\right\rangle_{\mathcal{H}} .
\end{aligned}
$$

Notice that the left hand side of the inner product only depends on $\pi$ while the right hand side only depends on $f$. Moreover, by the definition of kernel reactive POMDPs, the RHS has norm at most 2. Therefore, we conclude the proof by revoking Proposition 36 with $\zeta=2$.

## I Discussions on $\mathcal{D}_{\mathcal{F}}$ versus $\mathcal{D}_{\Delta}$ in BE Dimension

In this paper, we have mainly focused on the BE dimension induced by two special distribution families: $(a) \mathcal{D}_{\mathcal{F}}$ - the roll-in distributions produced by executing the greedy policies induced by the functions in $\mathcal{F}$, (b) $\mathcal{D}_{\Delta}$ - the collection of all Dirac distributions. And we prove that both low $\operatorname{dim}_{\mathrm{BE}}\left(\mathcal{F}, \mathcal{D}_{\mathcal{F}}, \epsilon\right)$ and low $\operatorname{dim}_{\mathrm{BE}}\left(\mathcal{F}, \mathcal{D}_{\Delta}, \epsilon\right)$ can imply sample-efficient learning. As a result, it is natural to ask what is the relation between $\operatorname{dim}_{\mathrm{BE}}\left(\mathcal{F}, \mathcal{D}_{\mathcal{F}}, \epsilon\right)$ and $\operatorname{dim}_{\mathrm{BE}}\left(\mathcal{F}, \mathcal{D}_{\Delta}, \epsilon\right)$ ? Is it possible that one of them is always no larger than the other so that we only need to use the smaller one? We answer this question with the following proposition, showing that either of them can be arbitrarily larger than the other.
Proposition 45. There exists absolute constant $c$ such that for any $m \in \mathbb{N}^{+}$,
(a) there exist an MDP and a function class $\mathcal{F}$ satisfying for all $\epsilon \in(0,1 / 2]$, $\operatorname{dim}_{\mathrm{BE}}\left(\mathcal{F}, \mathcal{D}_{\mathcal{F}}, \epsilon\right) \leq c$ while $\operatorname{dim}_{\mathrm{BE}}\left(\mathcal{F}, \mathcal{D}_{\Delta}, \epsilon\right) \geq m$.
(b) there exist an MDP and a function class $\mathcal{F}$ satisfying for all $\epsilon \in(0,1 / 2]$, $\operatorname{dim}_{\mathrm{BE}}\left(\mathcal{F}, \mathcal{D}_{\Delta}, \epsilon\right) \leq c$ while $\operatorname{dim}_{\mathrm{BE}}\left(\mathcal{F}, \mathcal{D}_{\mathcal{F}}, \epsilon\right) \geq m$.

Proof. We prove (a) first. Consider the following contextual bandits problem $(H=1)$.

- There are $m$ states $s_{1}, \ldots, s_{m}$ but the agent always starts at $s_{1}$. This means the agent can never visit other states because each episode contains only one step ( $H=1$ ).
- There are two actions $a_{1}$ and $a_{2}$. The reward function is zero for any state-action pair.
- The function class $\mathcal{F}_{1}=\left\{f_{i}(s, a)=\mathbf{1}\left(s=s_{i}\right)+\mathbf{1}\left(a=a_{1}\right): i \in[m]\right\}$.

First of all, note in this setting $\mathcal{D}_{\Delta}$ is the collection of all Dirac distributions over $\mathcal{S} \times \mathcal{A}, \mathcal{D}_{\mathcal{F}, 1}$ is a singleton containing only $\delta_{\left(s_{1}, a_{1}\right)}$, and $\left(I-\mathcal{T}_{1}\right) \mathcal{F}$ is simply $\mathcal{F}_{1}$ because $H=1$ and $r \equiv 0$. Since $\mathcal{D}_{\mathcal{F}, 1}$ has cardinality one, it follows directly from definition that $\operatorname{dim}_{\mathrm{BE}}\left(\mathcal{F}, \mathcal{D}_{\Delta}, \epsilon\right)$ is at most 1. Moreover, it is easy to verify that $\left(s_{1}, a_{2}\right),\left(s_{2}, a_{2}\right), \ldots,\left(s_{m}, a_{m}\right)$ is a 1-independent sequence with respect to $\mathcal{F}$ because we have $f_{i}\left(s_{j}, a_{2}\right)=\mathbf{1}(i=j)$ for all $i, j \in[m]$. As a result, we have $\operatorname{dim}_{\mathrm{BE}}\left(\mathcal{F}, \mathcal{D}_{\Delta}, \epsilon\right) \geq m$ for all $\epsilon \in(0,1]$.
Now we come to the proof of $(b)$. Consider the following contextual bandits problem $(H=1)$.

- There are 2 states $s_{1}$ and $s_{2}$. In each episode, the agent starts at $s_{1}$ or $s_{2}$ uniformly at random.
- There are $m$ actions $a_{1}, \ldots, a_{m}$. The reward function is zero for any state-action pair.
- The function class $\mathcal{F}_{1}=\left\{f_{i}(s, a)=\left(2 \cdot \mathbf{1}\left(s=s_{1}\right)-1\right)+0.5 \cdot \mathbf{1}\left(a=a_{i}\right): i \in[m]\right\}$.

First of all, note in this setting $\left(I-\mathcal{T}_{1}\right) \mathcal{F}$ is simply $\mathcal{F}_{1}$ and the roll-in distribution induced by the greedy policy of $f_{i}$ is the uniform distribution over $\left(s_{1}, a_{i}\right)$ and $\left(s_{2}, a_{i}\right)$, which we denote as $\mu_{i}$. It is easy to verify that $\mu_{1}, \ldots, \mu_{m}$ is a 0.5 -independent sequence with respect to $\mathcal{F}$ because $\mathbb{E}_{(s, a) \sim \mu_{i}}\left[f_{j}(s, a)\right]=0.5 \cdot \mathbf{1}(i=j)$. Therefore, for all $\epsilon \in(0,0.5], \operatorname{dim}_{\mathrm{BE}}\left(\mathcal{F}, \mathcal{D}_{\mathcal{F}}, \epsilon\right) \geq m$.
Next, we upper bound $\operatorname{dim}_{\mathrm{BE}}\left(\mathcal{F}, \mathcal{D}_{\Delta}, \epsilon\right)$ which is equivalent to $\operatorname{dim}_{\mathrm{DE}}\left(\mathcal{F}_{1}, \mathcal{D}_{\Delta}, \epsilon\right)$ in this problem. Assume $\operatorname{dim}_{\mathrm{DE}}\left(\mathcal{F}_{1}, \mathcal{D}_{\Delta}, \epsilon\right)=k$. Then there exist $g_{1}, \ldots, g_{k} \in \mathcal{F}_{1}$ and $w_{1}, \ldots, w_{k} \in \mathcal{S} \times \mathcal{A}$ such that for all $i \in[k], \sqrt{\sum_{t=1}^{i-1}\left(g_{i}\left(w_{i}\right)\right)^{2}} \leq \epsilon$ and $\left|g_{i}\left(w_{i}\right)\right|>\epsilon$. Note that we have $|f(s, a)| \in[0.5,1.5]$ for all $(s, a, f) \in \mathcal{S} \times \mathcal{A} \times \mathcal{F}_{1}$. Therefore, if $\epsilon>1.5$, then $k=0$; if $\epsilon \leq 1.5$, then $k \leq 10$ because $0.5 \times \sqrt{k-1} \leq \sqrt{\sum_{t=1}^{k-1}\left(g_{k}\left(w_{t}\right)\right)^{2}} \leq \epsilon \leq 1.5$.

