## A Proof of Theorem 3.1

Hyperparameters for the algorithm We define

$$
\begin{align*}
E_{d} & =3 d \frac{e}{e-1} \ln \left\{3+3(2 B / \varepsilon)^{2}\right\}+1  \tag{9}\\
n_{\text {rollout }} & =\frac{2 H^{2}(1+2 B)^{2} \log \left(\frac{2\left(E_{d}+1\right)}{\delta}\right)}{\varepsilon_{\text {target }}^{2}}  \tag{10}\\
N & =\left(E_{d}+1\right) n_{\text {rollout }} H A  \tag{11}\\
n_{\text {eval }} & =\frac{50 H^{2}\left(1+B^{2}\right)(d+1) \log \left(\frac{2(d+1) N}{\delta}\right)}{\varepsilon_{\text {target }}^{2}}  \tag{12}\\
\varepsilon_{\text {eval }} & =\sqrt{\frac{\log \left(\frac{2(d+1) N}{\delta}\right)}{2 n_{\text {eval }}}}  \tag{13}\\
\bar{\varepsilon}_{\text {eval }} & =\sqrt{1+B^{2}} \sqrt{d+1} \varepsilon_{\text {eval }}  \tag{14}\\
\varepsilon_{\text {tol }} & =4 \bar{\varepsilon}_{\text {eval }}  \tag{15}\\
\varepsilon_{\text {roll }} & =H(1+2 B) \sqrt{\frac{\log \left(2\left(E_{d}+1\right) / \delta\right)}{2 n_{\text {rollout }}}} \tag{16}
\end{align*}
$$

As we will see, $E_{d}$ is the upper bound on the number of parameters observed before global consistency holds, $n_{\text {eval }}$ is the number of samples taken to estimate each TD vector, $n_{\text {rollout }}$ is the number of rollouts executed for each parameter, $N$ is the maximum number of state-action pairs which will be seen by DELPHI, $\varepsilon_{\text {eval }}$ is the error between the estimated TD vector and the true TD vector, and $\varepsilon_{\text {roll }}$ is the error of replacing the expected total rewards with the average over several rollouts.

## A. 1 Part 1: Concentration Inequalities

Recall that we write

$$
\Delta_{s, a}=r(s, a) \oplus\left(\mathbb{E}_{P(s, a)}\left[\varphi\left(s^{\prime}\right)\right]-\varphi(s)\right)
$$

for the true TD vector for state-action $(s, a)$. We write $\hat{\Delta}_{s, a}$ for any estimated TD vector resulting from Line 11 , and $\tilde{\Delta}_{s, a}$ for any doubly-estimated TD vector resulting from Line 16
Lemma 3.2 (Concentration of $\hat{\Delta}_{s, a}$ (Line 11). For any $s, a \in \mathcal{S} \times \mathcal{A}$ that is observed throughout the execution DELPHI, with $n_{\text {eval }}$ samples in Line 11 we have that with probability $\geq 1-\delta$, $\left\|\hat{\Delta}_{s, a}-\Delta_{s, a}\right\|_{\infty} \leq \varepsilon_{e v a l}$ and thus that $\left\langle 1 \oplus \theta, \hat{\Delta}_{s, a}-\Delta_{s, a}\right\rangle \leq \bar{\varepsilon}_{\text {eval }}$.

Proof. Pick a single $s, a$, and omit the dependence on $s, a$ for cleanliness. Starting with the reward concentration: at every step we collect $n_{\text {eval }}$ iid samples from $R\left(s_{i}, a_{i}\right)$. By Hoeffding, since rewards are bounded in $[0,1]$, the empirical average satisfies

$$
\left|\frac{1}{n} \sum r_{i}-\mathbb{E}[r]\right|_{\substack{\text { w.p. } 1-\delta /((d+1) N)}}^{\frac{\log (2(d+1) N / \delta)}{2 n_{\text {eval }}}}
$$

Onto the transition probabilities. Since the features are such that $\|\varphi\|_{2} \leq 1$, we also have that $\|\varphi\|_{\infty} \leq 1$. Thus, each coordinate of the feature map is bounded by 1. By Hoeffding, for each coordinate $i$ of the vector, we have that

$$
\left|\frac{1}{n} \sum_{j=1}^{n} \varphi_{i}^{(j)}-\mathbb{E} \varphi_{i}^{(j)}\right|^{\text {w.p. } 1-\delta /((d+1) N)} \leq \sqrt{\frac{\log (2(d+1) / \delta)}{2 n_{\mathrm{eval}}}}
$$

Thus, with a union bound over these events for each coordinate $i$ and the event for the reward concentration, we have that w.p. $\geq 1-\delta / N$

$$
\|\hat{\Delta}-\Delta\|_{\infty} \leq \sqrt{\frac{\log \left(\frac{2(d+1) N}{\delta}\right)}{2 n_{\text {eval }}}}=\varepsilon_{\text {eval }}
$$

Finally, by a union bound over all $N$ state-action pairs observed by DELPHI, we have that for every $s, a$ encountered:

$$
\left\|\hat{\Delta}_{s, a}-\Delta_{s, a}\right\|_{\infty} \leq \sqrt{\frac{\log \left(\frac{2(d+1) N}{\delta}\right)}{2 n_{\text {eval }}}}=\varepsilon_{\text {eval }} \text { with probability } \geq 1-\delta
$$

We further find that: $\|\hat{\Delta}-\Delta\|_{2} \leq \sqrt{d+1} \varepsilon_{\text {eval }}$ (since $\Delta$ and $\hat{\Delta}$ are $(d+1)$-dimensional). By Cauchy-Schwartz this gives, $\forall \theta \in \operatorname{Ball}_{\ell_{2}}(B)$ :

$$
|\langle 1 \oplus \theta, \hat{\Delta}-\Delta\rangle| \leq\|1 \oplus \theta\|_{2}\|\hat{\Delta}-\Delta\|_{2}=\sqrt{1+\|\theta\|_{2}^{2}}\|\hat{\Delta}-\Delta\|_{2}=\sqrt{1+B^{2}} \sqrt{d+1} \varepsilon_{\mathrm{eval}}=\bar{\varepsilon}_{\mathrm{eval}}
$$

Lemma 3.3 ( $\tilde{\Delta}_{s, a}$ concentrates even more (Line 16). Similarly, for all $s$, a where we call the oracle, with probability $1-\delta$, we have $\left\|\tilde{\Delta}_{s, a}-\Delta_{s, a}\right\|_{\infty} \leq \varepsilon_{\text {eval }} /\left(2 \sqrt{E_{d}}\right)$, and thus, $\forall \theta \in \operatorname{Ball}_{\ell_{2}}(B)$, $\left|\left\langle 1 \oplus \theta, \tilde{\Delta}_{s, a}-\Delta_{s, a}\right\rangle\right| \leq \bar{\varepsilon}_{\text {eval }} /\left(2 \sqrt{E_{d}}\right)$.

Proof. Follows from the same proof as Lemma3.2, just replace $n_{\text {eval }}$ by $4 E_{d} n_{\text {eval }}$.

## A. 2 Part 2: Optimism

Lemma 3.4 ( $\theta^{\circ}$ not eliminated). With probability $\geq 1-\delta, \theta^{\circ} \in \Theta_{t}$ for all iterations $t \in\left[E_{d}+1\right]$.
Proof. Let $\tilde{\Delta}_{t}$ be the TD vector which is added to $\Theta_{t}$ at time $t \in\left[E_{d}\right]$. Note that, by Lemma 3.3, for any given $t$, we have

$$
\begin{equation*}
\left|\left\langle 1 \oplus \theta^{\circ}, \tilde{\Delta}_{t}-\Delta_{t}\right\rangle\right| \leq \bar{\varepsilon}_{\text {eval }} /\left(2 \sqrt{E_{d}}\right) \tag{17}
\end{equation*}
$$

for all $t$, w.p. $\geq 1-\delta$. The parameter set is defined as

$$
\Theta_{t}=\left\{\theta:\left|\left\langle 1 \oplus \theta, \tilde{\Delta}_{i}\right\rangle\right| \leq \frac{\bar{\varepsilon}_{\text {eval }}}{2 \sqrt{E_{d}}} \forall i \in[t]\right\} .
$$

Note that, for $\theta^{\circ}$, this is equivalent to requiring

$$
\left|\left\langle 1 \oplus \theta^{\circ}, \tilde{\Delta}_{t}-\Delta_{t}\right\rangle\right| \leq \frac{\bar{\varepsilon}_{\mathrm{eval}}}{2 \sqrt{E_{d}}}
$$

since $\left\langle 1 \oplus \theta^{\circ}, \Delta_{t}\right\rangle=0$ for all $\Delta_{t}$ (recall Equation 55. This holds by Eq. 17).
Lemma 3.5 (Optimism). Under the event of Lemma 3.4 we have $v_{t}\left(s_{0}\right) \geq v^{\circ}\left(s_{0}\right), \forall t \in\left[E_{d}\right]$.
Proof (of Lemma 3.5). Under the event of Lemma3.4, $\theta^{\circ}$ is never eliminated from $\Theta_{t}$. Thus, by the optimistic update rule, we have $\theta_{t+1}=\operatorname{argmax}_{\theta \in \Theta_{t}} \theta^{\top} \varphi\left(s_{0}\right)$, and since $\theta^{\circ} \in \Theta_{t}$, then

$$
\theta_{t+1}^{\top} \varphi\left(s_{0}\right)=v_{\theta_{t+1}}\left(s_{0}\right) \geq\left(\theta^{\circ}\right)^{\top} \varphi\left(s_{0}\right)=v_{\theta^{\circ}}\left(s_{0}\right)=v^{\circ}\left(s_{0}\right)
$$

## A. 3 Part 3: Iteration Complexity

To bound the iteration complexity of our algorithm, we will need to introduce the notion of Eluder dimension. We use a simplified form introduced by [LKFS21], although the first definition is due to [RVR13].
Definition A. 1 (Eluder dimension [LKFS21]). Let $\mathcal{F}$ be a real-valued function class on domain $\mathcal{X}$. Fix a reference function $f^{\star} \in \mathcal{F}$, and a scale $\varepsilon$. The Eluder dimension of $\mathcal{F}$ at a scale $\varepsilon$, w.r.t. $f^{\star}$, is the length $\tau \in \mathbb{N}$ of the longest sequence of points $\left(\left(x_{1}, f_{1}\right), \ldots\left(x_{\tau}, f_{\tau}\right)\right)$ such that

$$
\begin{equation*}
\forall i \in[\tau]: \quad\left|f_{i}\left(x_{i}\right)-f^{\star}\left(x_{i}\right)\right|>\varepsilon, \quad \text { and } \quad \sum_{j<i}\left(f_{i}\left(x_{j}\right)-f^{\star}\left(x_{j}\right)\right)^{2} \leq \varepsilon^{2} \tag{18}
\end{equation*}
$$

An Eluder sequence of length $\tau$ (with respect to $f^{\star}$ ) is any sequence $\left(x_{i}, f_{i}\right)_{i=1}^{\tau}$ which satisfies Eq (18) for each $i$.

In other words, the Eluder dimension is the length of the longest sequence of points such that, for each $i$, we can find a new function $f_{i}$ which is large with respect to $f^{\star}$ on $x_{i}$ but correctly fits $f^{\star}$ on historical data $x_{1}, \ldots, x_{j}$, for $j<i$. We will use the folllowing bound for the Eluder dimension of linear functions with $d$-dimensional parameters.
Lemma A. 2 ([RVR13]). For any $f^{\star}$, the Eluder dimension of the function class $\mathcal{F}=\left\{f_{\theta}(x)=\theta^{\top} x\right\}$, assuming that $\|\theta\|_{2} \leq B$ and $\|x\|_{2} \leq \gamma$, is

$$
\operatorname{dim}_{E}(\mathcal{F}, \varepsilon) \leq 3 d \frac{e}{e-1} \ln \left\{3+3(2 B / \varepsilon)^{2}\right\}+1=\mathcal{O}(d \ln (B / \varepsilon))
$$

This gives us enough to prove our iteration bound - we will show that the sequence of linear parameters chosen by our algorithm together with each new TD vector obtained from the oracle forms an Eluder sequence with respect to $\theta^{\star}$

Lemma 3.6 (Iteration Complexity). With probability $\geq 1-2 \delta$, the iteration complexity of the algorithm is at most the Eluder dimension at scale $\bar{\varepsilon}_{\text {eval }}$, i.e. $E_{d}=\mathcal{O}\left(d \ln \left(B / \bar{\varepsilon}_{\text {eval }}\right)\right)$.

Proof (of Lemma 3.6). Assume the events of Lemma 3.2 and Lemma 3.3, which happen together with probability $\geq 1-2 \delta$. Our function class is $\mathcal{F}=\left\{f_{\theta}(x)=\langle 1 \oplus \theta, x\rangle\right\}$. This is a subset of all linear functions on $d+1$ dimensions, so by Lemma A. $2^{7}$ it has Eluder dimension at least $\mathcal{O}((d+1) \ln (B / \varepsilon))=\mathcal{O}(d \ln (B / \varepsilon))$. We pick $f^{\star}=f_{\theta^{\circ}}$. We will show that $\forall i$, the sequence $\left(\tilde{\Delta}_{i}, f_{\theta_{i}}\right)$ forms an Eluder sequence with respect to $f^{\star}$, from which it will follow that its length is bounded by $E_{d}$. We do this by induction. The base case is obvious. By the constraint set definition (Line 17, we have $\left|f_{\theta_{i}}\left(\tilde{\Delta}_{j}\right)\right|=\left|\left\langle 1 \oplus \theta_{i}, \tilde{\Delta}_{j}\right\rangle\right| \leq \frac{\bar{\varepsilon}_{\text {eval }}}{2 \sqrt{E_{d}}} \forall j<i$. Assuming the event of Lemma 3.4 we also have $\left|f_{\theta^{\circ}}\left(\tilde{\Delta}_{j}\right)\right| \leq \frac{\bar{\varepsilon}_{\text {eval }}}{2 \sqrt{E_{d}}}$, since $\theta^{\circ} \in \Theta_{i-1}$. Thus

$$
\left|f_{\theta_{i}}\left(\tilde{\Delta}_{j}\right)-f_{\theta^{\circ}}\left(\tilde{\Delta}_{j}\right)\right| \leq 2 \frac{\bar{\varepsilon}_{\text {eval }}}{2 \sqrt{E_{d}}} \forall i \Longrightarrow \sum_{j<i}\left(f_{\theta_{i}}\left(\tilde{\Delta}_{j}\right)-f_{\theta^{\circ}}\left(\tilde{\Delta}_{j}\right)\right)^{2} \leq \bar{\varepsilon}_{\mathrm{eval}}^{2} \frac{i-1}{E_{d}} \leq \bar{\varepsilon}_{\mathrm{eval}}^{2}
$$

where in the last inequality we have used $i-1 \leq E_{d}$ by the induction hypothesis. Thus, the second condition of the Eluder dimension is satisfied. For the first condition, we want to show that $\left|f_{\theta_{i}}\left(\tilde{\Delta}_{i}\right)-f_{\theta^{\circ}}\left(\tilde{\Delta}_{i}\right)\right|>\bar{\varepsilon}_{\text {eval }}$. Note that, by Lemma 3.4 we have $\left|f_{\theta^{\circ}}\left(\tilde{\Delta}_{i}\right)\right| \leq \bar{\varepsilon}_{\text {eval }} /\left(2 \sqrt{E_{d}}\right) \leq \bar{\varepsilon}_{\text {eval }}$. Recall from Line 13 that $\left|\left\langle\hat{\Delta}_{i}, 1 \oplus \theta_{i}\right\rangle\right|=\left|f_{\theta_{i}}\left(\hat{\Delta}_{i}\right)\right|>\varepsilon_{\text {tol }}$. Using concentration and linearity, we have $\left|f_{\theta_{i}}\left(\hat{\Delta}_{i}\right)\right|-\left|f_{\theta_{i}}\left(\tilde{\Delta}_{i}\right)\right| \leq 2 \bar{\varepsilon}_{\text {eval }}$. Thus, $\left|f_{\theta_{i}}\left(\tilde{\Delta}_{i}\right)\right| \geq \varepsilon_{\text {tol }}-2 \bar{\varepsilon}_{\text {eval }}=2 \bar{\varepsilon}_{\text {eval }}$ since $\varepsilon_{\text {tol }}=4 \bar{\varepsilon}_{\text {eval }}$. Putting this together gives

$$
\left|f_{\theta_{i}}\left(\tilde{\Delta}_{i}\right)-f_{\theta^{\circ}}\left(\tilde{\Delta}_{i}\right)\right| \geq f_{\theta_{i}}\left(\tilde{\Delta}_{i}\right)-f_{\theta^{\circ}}\left(\tilde{\Delta}_{i}\right) \geq 2 \bar{\varepsilon}_{\text {eval }}-\bar{\varepsilon}_{\text {eval }}=\bar{\varepsilon}_{\text {eval }}
$$

And we are done.

## A. 4 Part 4: Consistency, and Putting Everything Together

Lemma 3.7 (Consistency $\Longrightarrow$ accurate prediction). If $m$ rollouts have occured without any inconsistencies (i.e., the if statement of Line 13 never gets triggered), then $v^{\pi_{\theta}}\left(s_{0}\right)>v_{\theta}\left(s_{0}\right)-$ $5 H \bar{\varepsilon}_{\text {eval }}-\varepsilon_{\text {roll }}$ with probability $\geq 1-3 \delta$.

Proof (of Lemma 3.7). Assume the event of Lemmas 3.2 and Lemmas 3.3, which together happen with probability $\geq 1-2 \delta$. The third event which we assume will be introduced shortly.

For cleanliness, let us write $\theta$ for the final parameter which observes $m$ rollouts without consistency break. Observe the following calculation:

[^0]\[

$$
\begin{array}{rlr}
v^{\pi_{\theta}}\left(s_{0}\right) & =\mathbb{E}_{\pi_{\theta}}\left[\sum_{j \in[H]} r_{S_{j}, A_{j}}\right] \\
& =\mathbb{E}\left[\left\langle\left(\sum_{j \in[H]} r_{S_{j}, A_{j}}\right) \oplus \varphi\left(s_{H+1}\right), 1 \oplus \theta\right\rangle\right] \quad\left(\varphi\left(s_{H+1}\right)=0\right) \\
& =\mathbb{E}\left[\left\langle\varphi\left(s_{0}\right), \theta\right\rangle+\sum_{j \in[H]}\left\langle r_{S_{j}, A_{j}} \oplus\left(\varphi\left(S_{j+1}\right)-\varphi\left(S_{j}\right)\right), 1 \oplus \theta\right\rangle\right] \quad \text { (telescoping sum) } \\
& =\left\langle\varphi\left(s_{0}\right), \theta\right\rangle+\mathbb{E}\left[\sum_{j \in[H]}\left\langle r_{S_{j}, A_{j}} \oplus\left(P_{S_{j}, A_{j}} \varphi(\cdot)-\varphi\left(S_{j}\right)\right), 1 \oplus \theta\right\rangle\right],
\end{array}
$$
\]

Now observe that after $n_{\text {rollout }}$ number of rollouts, we have $n_{\text {rollout }}$ unbiased estimates of the expected trajectories. Thus we can approximate $\mathbb{E}\left[\sum_{j \in[H]}\left\langle r_{S_{j}, A_{j}} \oplus\left(P_{S_{j}, A_{j}} \varphi(\cdot)-\varphi\left(S_{j}\right)\right), 1 \oplus \theta\right\rangle\right] \quad \approx \quad \frac{1}{n_{\text {rollout }}} \sum_{i=1}^{n_{\text {rollout }}} \mathbb{E}\left[r_{S_{j}^{i}, A_{j}^{i}}\right] \oplus$ $\left(P_{S_{j}^{i}, A_{j}^{i}} \varphi(\cdot)-\varphi\left(S_{j}^{i}\right)\right)$, where $S_{j}^{i}$ and $A_{j}^{i}$ are the states and actions in horizon $j \in[H]$ of rollout $i \in\left[n_{\text {rollout }}\right]$. More precisely, using Hoeffding's and that $\left.\left\langle\mathbb{E}\left[r_{S_{j}^{i}, A_{j}^{i}}\right] \oplus\left(P_{S_{j}^{i}, A_{j}^{i}} \varphi(\cdot)-\varphi\left(S_{j}\right)\right)\right], 1 \oplus \theta\right\rangle \leq 1+2 B$, we can get $\left|\frac{1}{n_{\text {rollout }}} \sum_{i \in[m]}\left(\sum_{j \in[H]}\left\langle\Delta_{j}^{i}, 1 \oplus \theta\right\rangle\right)-\left(\mathbb{E} \sum_{j \in[H]}\left\langle\Delta_{j}^{i}, 1 \oplus \theta\right\rangle\right)\right| \leq H(1+2 B) \sqrt{\frac{\log \left(2\left(E_{d}+1\right) / \delta\right)}{2 m}}:=\varepsilon_{\text {roll }}$, w.p. $\geq 1-\frac{\delta}{E_{d}+1}$, where we wrote $\Delta_{j}^{i}=\mathbb{E}\left[r_{S_{j}^{i}, A_{j}^{i}}\right] \oplus\left(P_{S_{j}^{i}, A_{j}^{i}} \varphi(\cdot)-\varphi\left(S_{j}\right)\right)$. By a union bound this happens for all $t \in\left[E_{d}+1\right]$ with probability $\geq 1-\delta$. Picking up where we left off:

$$
\left.\begin{array}{rl}
v^{\pi_{\theta}}\left(s_{0}\right) & \left.\geq\left\langle\varphi\left(s_{0}\right), \theta\right\rangle+\frac{1}{n_{\text {rollout }}} \sum_{i=1}^{n_{\text {rollout }}} \sum_{j \in[H]}\left\langle\mathbb{E}\left[r_{S_{j}^{i}, A_{j}^{i}}\right] \oplus\left(P_{S_{j}^{i}, A_{j}^{i}} \varphi(\cdot)-\varphi\left(S_{j}\right)\right)\right], 1 \oplus \theta\right\rangle-\varepsilon_{\text {roll }} \\
& \geq\left\langle\varphi\left(s_{0}\right), \theta\right\rangle+\frac{1}{m} \sum_{i=1}^{m} \sum_{j \in[H]}\left\langle\hat{\Delta}_{s_{j}^{i}, a_{j}^{i}}, 1 \oplus \theta\right\rangle-H \bar{\varepsilon}_{\text {eval }}-\varepsilon_{\text {roll }} \quad \quad \text { (evaluation error) } \\
& \geq v_{\theta}\left(s_{0}\right)-\frac{1}{m} \sum_{i=1}^{m} \sum_{j \in[H]} 4 \bar{\varepsilon}_{\text {eval }}-H \bar{\varepsilon}_{\text {eval }}-\varepsilon_{\text {roll }} \\
& =v_{\theta}\left(s_{0}\right)-5 H \bar{\varepsilon}_{\text {eval }}-\varepsilon_{\text {roll }}
\end{array} \quad \text { (consistency holds) }\right)
$$

and we are done. The second inequality follows from Lemma 3.2, which holds for all steps in all $m$ trajectories in all $E_{d}+1$ iterations. The third inequality holds since, for every $S_{j}^{i}$, there exists a consistent action, i.e. an action such that $\left|\left\langle\hat{\Delta}_{j}^{i}, 1 \oplus \theta\right\rangle\right| \leq \varepsilon_{\text {tol }}=4 \bar{\varepsilon}_{\text {eval }}$.

## A. 5 Simulator inaccuracy and misspecification

We start with the case of inaccurate simulators. Recall that we say that a simulator is $\lambda$-innacurate if the samples obtained are of the form $\left(\Pi\left(r+\lambda_{s, a}\right), s^{\prime}\right)$, for any $(s, a)$ and for some constant $\lambda_{s, a}$ that satisfies $\left|\lambda_{s, a}\right| \leq \lambda$.
Theorem 3.8 (DELPHI with misspecification). Redefine $n_{\text {eval }}^{\prime}=4 n_{\text {eval }}$ (Eq. 12) and all subsequent hyperparameters which depend on $n_{\text {eval }}^{\prime}$. Then we have that, for all MDPs that are at most $\frac{\varepsilon_{\text {eval }}}{8 \sqrt{E_{d}}}-$ misspecified or for all simulators that are at most $\frac{\bar{\varepsilon}_{\text {eval }}}{4 \sqrt{E_{d}}}$-innaccurate, the conclusions of Theorem 3.1 continue to hold when running DELPHI with the new hyperparameters.

We can repeat the proof of DELPHI, and in fact the only difference will be in Part 1 of the proof (Lemmas 3.2 and 3.3 ) will continue to hold, which also implies that the rest of the proof will continue to hold.

Lemma A. 3 (Concentration of $\hat{\Delta}_{s, a}$ (Line 11). For any $s, a \in \mathcal{S} \times \mathcal{A}$ that is observed throughout the execution DELPHI, with $n_{\text {eval }}$ samples in Line 11 we have that with probability $\geq 1-\delta$, $\left\|\hat{\Delta}_{s, a}-\Delta_{s, a}\right\|_{\infty} \leq \varepsilon_{e v a l}$ and thus that $\left\langle 1 \oplus \theta, \hat{\Delta}_{s, a}-\Delta_{s, a}\right\rangle \leq \bar{\varepsilon}_{\text {eval }}$.
Lemma A. 4 ( $\tilde{\Delta}_{s, a}$ concentrates even more (Line 16). Similarly, for all $s$, a where we call the oracle, with probability $1-\delta$, we have $\left\|\tilde{\Delta}_{s, a}-\Delta_{s, a}\right\|_{\infty} \leq \varepsilon_{\text {eval }} /\left(2 \sqrt{E_{d}}\right)$, and thus, $\forall \theta \in B a l l_{\ell_{2}}(B)$, $\left|\left\langle 1 \oplus \theta, \tilde{\Delta}_{s, a}-\Delta_{s, a}\right\rangle\right| \leq \bar{\varepsilon}_{\text {eval }} /\left(2 \sqrt{E_{d}}\right)$.

Proof (of Lemmas A.3 and A.4. We repeat the proof of Lemmas 3.2 and 3.3 . Note that $\hat{\Delta}_{s, a}$ is an average of i.i.d. random variables with mean $\Delta_{s, a}^{\prime}=\left(r(s, a)+\lambda_{s, a}\right) \oplus P_{s, a} \varphi-\varphi(s)$. As before, by Hoeffdings, we have $\left|\left\langle 1 \oplus \theta, \hat{\Delta}-\Delta^{\prime}\right\rangle\right| \leq \bar{\varepsilon}_{\text {eval }} / 2$, where the factor of $1 / 2$ is due to the definition of $n_{\text {eval }}^{\prime}$. This gives that
$|\langle 1 \oplus \theta, \hat{\Delta}-\Delta\rangle| \leq\left|\left\langle 1 \oplus \theta, \hat{\Delta}-\Delta^{\prime}\right\rangle\right|+\left|\left\langle 1 \oplus \theta, \Delta^{\prime}-\Delta\right\rangle\right| \leq \bar{\varepsilon}_{\text {eval }} / 2+\lambda \leq \bar{\varepsilon}_{\text {eval }} / 2+\bar{\varepsilon}_{\text {eval }} / 4 \sqrt{E_{d}} \leq \bar{\varepsilon}_{\text {eval }}$
The proof of the second part follows similarly, since we have

$$
|\langle 1 \oplus \theta, \tilde{\Delta}-\Delta\rangle| \leq\left|\left\langle 1 \oplus \theta, \tilde{\Delta}-\Delta^{\prime}\right\rangle\right|+\left|\left\langle 1 \oplus \theta, \Delta^{\prime}-\Delta\right\rangle\right| \leq \bar{\varepsilon}_{\mathrm{eval}} / 4 \sqrt{E_{d}}+\lambda \leq \bar{\varepsilon}_{\mathrm{eval}} / 2 \sqrt{E_{d}}
$$

Next we handle the value misspecification case. Recall that the MDP is $\eta$-misspecified for the expert policy $\pi^{\circ}$ and the features $\varphi$ if there exists $\theta^{\circ}$ such that $\sup _{s}\left|v^{\circ}(s)-\left\langle\varphi(s), \theta^{\circ}\right\rangle\right| \leq \eta$. Here, we use a reduction argument to show that an $\eta$-misspecified MDP can be reduced to a $2 \eta$-innacurate simulator of a realizable MDP. Namely, using the same reduction as Appendix D of [WAJAYJS21], we can construct an alternative MDP $\mathcal{M}^{\prime}$ such that $\mathcal{M}^{\prime}$ is realizable but is a $2 \eta$-inaccurate simulator of $\mathcal{M}$. The result then follows from the first part of the proof.

## B Proof of Section 4

## B. 1 The MDP construction

Theorems 4.1 and 4.3 use the same MDP construction, which is inspired by the recent and remarkable lower bound of [WSG21] (itself an extension of the lower bound of [WAS21]). We give an overview of the MDP and its features, and describe the changes from the original construction of [WSG21]. The state space consists of a hypercube in $p$ dimensions, $\mathcal{S}=\{ \pm 1\}^{p}$, for some $p \in \mathbb{N}$, and assume for simplicity that $p$ is divisible by 4 . More specifically, the state space at stage $h$ also contains the history of vectors encountered so far $\left(s_{1}, \ldots, s_{h}\right)$, which is uniquely defined as the transitions are deterministic. The dimension $p$ will end up being $p \approx \sqrt{d}$ for the value-based lower bound and $p \approx d$ for the policy-based lower bound. We write $\rho(\cdot, \cdot)$ for the Hamming distance on $\mathcal{S}$, recalling that is a bilinear function of its arguments, i.e. $\rho\left(s_{1}, s_{2}\right)=\frac{1}{2}\left(p-\left\langle s_{1}, s_{2}\right\rangle\right)$. The action set is $\mathcal{A}=[p]$, and each action $a \in[p]$ will correspond to flipping the $a^{\text {th }}$ bit of the current state. Each trajectory of horizon $H$ has $K$ phases $(K \in \mathbb{N}$ ), and each phase consists of $p$ steps (thus, $H=K p$ ). We write $s_{0}, s_{1}, \ldots, s_{k}, \ldots s_{K}$ for the states reached at the end of each phase, and when we need to we will use $s_{k, i}$ for a state at step $i$ of phase $k$. The start state is $s_{0}=\overrightarrow{1}$, the all-ones vector.

There is a special "goal state" $s^{\star}$, and reward is given only if a) at the end of any phase, $\rho\left(s_{k}, s^{\star}\right) \leq$ $p / 4$, or b) the learner reaches horizon $H$, i.e. the end of phase $K$. The reward function decays geometrically, and is defined according to the sequence of states $s_{0}, \ldots s_{k}$ obtained at the end of each phase. Letting $g\left(s_{1}, s_{2}\right):=1-\rho\left(s_{1}, s_{2}\right) / p$ denote one minus the proportion of bits where two states differ, the reward function for reaching a $p / 4$-neighbourhood of $s^{\star}$ at stage $k$ (condition a) just described) is deterministic and has value

$$
\begin{equation*}
r_{w^{\star}}\left(\left(s_{i}\right)_{i=1}^{k}\right)=\left(\prod_{i=1}^{k} g\left(s_{i-1}, s_{i}\right)\right) g\left(s_{i}, s^{\star}\right) \tag{19}
\end{equation*}
$$

The reward function at stage $H$ (condition b) above) is always given and has the same expectation as Eq. (19), but will be a Bernoulli random variable.
Modulo some exceptions (to be described shortly), the transition function is deterministic and is defined via $\mathcal{P}(\tau(s, a) \mid s, a)=1$, where $\tau(s, a)=\left(s_{1}, \ldots,-s_{a}, \ldots s_{p}\right)$ corresponds to the new vector obtained from flipping the bit at index $a \in[p]$. The exceptions to these transition dynamics are if 1) a state within a $p / 4$-neighbourhood of $s^{\star}$ is reached, or 2 ) a move is repeated. In the first case, the MDP transitions to a game-over state $\perp$ (after which nothing else is possible and no reward is given). In the second case, if the move is repeated within the first $p / 4$ steps of a phase then the MDP similarly transitions to $\perp$, and if the move is repeated in the second $3 p / 4$ turns then the current state becomes frozen until the end of the current phase (i.e. no further bit flips are allowed). This implies that each bit can only be flipped once in each phase, and further that $g\left(s_{i-1}, s_{i}\right) \leq 3 / 4$ for each "legal" trajectory in the MDP (thus, the reward decays geometrically, cf. Equation (19)).
So far, the main modification in our construction from that of [WSG21] is the reward function. Their $g$ function is chosen to be $2^{\text {nd }}$ order in $\rho\left(s_{1}, s_{2}\right)$, while ours is linea in $\rho$. We now introduce the definition of the expert policy: we simply choose it to be the one which will flip the earliest index such that $s_{k, i}$ differs from $s^{\star}$ and such that its index has not yet been played in the round. If no such index exists, or if $s_{k, i}=s^{\star}$, the expert policy will simply freeze the current round by repeating the earliest index that has already been played. (Note that there is always a repeated index in this case, and that repeating an index will lead to freezing instead of termination, since at the start of any phase $k$ the trajectory must satisfy $\rho\left(s_{k}, s^{\star}\right)>p / 4$, otherwise the trajectory would have terminated). It will turn out that the expert's trajectory from the start state will be identical to that of the optimal policy (and, thus, will be equally difficult to compete with, in the sense of Eq. (2)). While our expert policy might seem optimal, this need not be the case for arbitrary states since it might sometimes be more advantageous for a state in phase $k$ which can no longer reach the $p / 4$-neighbourhood of $s^{\star}$ to simply aim at minimizing the inevitable factor of $g\left(s_{k-1}, s_{k}\right)$ that it will incur. The following lemma (proved in Appendix B.2) gives an expression for the value function corresponding to our expert policy, and shows that the value function satisfies Assumption 2.2

Lemma B.1. Let $s_{k, i}$ be a state in round $i$ of phase $k$. Note that, from $s_{k, i}$, the state $s_{k+1}$ that $\pi^{\circ}$ will reach at the end of the current phase is deterministic. The value function of $\pi^{\circ}$ is

$$
v^{\circ}\left(s_{k, i}\right)=\left(\prod_{k^{\prime} \in[k]} g\left(s_{k^{\prime}-1}, s_{k^{\prime}}\right)\right) g\left(s_{k}, s_{k+1}\right) g\left(s_{k+1}, s^{\star}\right),
$$

and furthermore this is linear with features $\varphi_{v}$ of dimension $d=\Theta\left(p^{2}\right)$.
The next lemma gives an expression showing that this same expert policy is linear in a different set of features $\varphi_{\pi}$, and thus that the expert policy satisfies Assumption 4.2
Lemma B.2. There exists a feature map $\varphi_{\pi}$ and parameter $\theta_{\pi}$ of dimension $d=\Theta(p)$ such that

$$
\pi^{\circ}\left(s_{k, i}\right)=\operatorname{argmax}_{a}\left\{\left\langle\varphi_{\pi}(s, a), \theta_{\pi}\right\rangle\right\}
$$

This lemma is also proved in Appendix B. 2
The essence of our lower bound is that each oracle call will reveal one bit of the secret state $s^{\star}$, and thus without $\tilde{\Omega}(p)$ calls and learner will need exponentially-many exploratory samples to solve this MDP. Thus, for the value-based lower bound, an agent with a $d$-dimensional feature map must be given an MDP which has a $p=\Theta(\sqrt{d})$-dimensional state space. However, for the policy-based lower bound, we can give the agent a $p=\Theta(d)$-dimensional MDP.

Information-theoretic lower bound Information-theoretically, a learning algorithm can solve this MDP only if they recover the secret state $s^{\star}$. We let $\operatorname{MDP}(p, K)$ denote an instance of the above MDP with dimension $p$ and parameters $K$. Following the approach of [WSG21], we prove the sample complexity hardness by reducing each $\operatorname{MDP}(p, K)$ to an instance of an abstract game called $\operatorname{CubeGame}(p, K)$. Details on CubeGame, and the proof of the following theorem, are deferred to Appendix B.3. The main result is the following:
Theorem B.3. Any learner which solves $\operatorname{MDP}(p, K)$ can be used to solve $\operatorname{CubEGAME}(p, K)$. Unless the number of oracle calls is $\Omega\left(\frac{p}{\log p}\right)$, any learner which is 0.01 -optimal on $\operatorname{CUBEGAME}(p, K)$, with a sample complexity of $N$, will need

$$
N \geq 2^{\Omega(p \wedge K)}
$$

Combined with the fact that a learner provided with $d$-dimensional features is given an MDP with parameter $p \approx \sqrt{d}$, this gives the result of Theorem 4.1

## B. 2 Proofs of Lemma B. 1, B. 2

We need some more notation in order to linearize the value function. Let $s_{k, i} \neq \perp$ be a state in step $i$ of phase $k$. We define the variable $\mathrm{ct}_{k, i}^{\text {flip }}=\rho\left(s_{k, 0}, s_{k, i}\right)$, which simply measures the number of components flipped so far in round $k$ of step $i$, and fix $k, i \in\{0,1\}^{p}$ which is a vector with 1 at a component $j$ if said component is currently frozen (i.e. if it has been played or if the entire state has been frozen), and 0 otherwise. Similarly, there are two scalars $e_{k, i}^{\text {fix }}$ and $e_{k, i}^{\neg \mathrm{fix}}$ which simply count the number of components which disagree with $s^{\star}$ that are currently frozen (for $e_{k, i}^{\mathrm{fix}}$ ) or not frozen (for $\left.e_{k, i}^{-\mathrm{fix}}\right)$. Note that only $e_{k, i}^{\mathrm{fix}}$ and $e_{k, i}^{\sim \mathrm{fix}}$ depend on $s^{\star}$, and in fact we have:

$$
\begin{align*}
e_{k, i}^{\mathrm{fix}} & =\frac{1}{2}\left(\left\langle\overrightarrow{\mathrm{1}}, \mathrm{fix}_{k, i}\right\rangle-\left\langle\mathrm{fix}_{k, i} \cdot s_{k, i}, s^{\star}\right\rangle\right)  \tag{20}\\
e_{k, i}^{\neg \mathrm{fix}} & =\frac{1}{2}\left(\left\langle\overrightarrow{\mathrm{1}}, \neg \mathrm{fix}_{k, i}\right\rangle-\left\langle\neg \mathrm{fix}_{k, i} \cdot s_{k, i}, s^{\star}\right\rangle\right), \tag{21}
\end{align*}
$$

where $\mathrm{fix}_{k, i} \cdot s_{k, i}$ denotes component-wise multiplication, i.e. $\left(\mathrm{fix}_{k, i} \cdot s_{k, i}\right)_{j}=\left(\mathrm{fix}_{k, i}\right)_{j} \cdot\left(s_{k, i}\right)_{j}$. Now, we have that:
Lemma B. 4 (Value of $v^{\circ}$ ). Let $s_{k, i}$ be a state in round $i$ of phase $k$. Let $s_{k+1}$ denote the state that $\pi^{\circ}$ will reach at the end of the current phase when starting from $s_{k, i}$ (and note that this choice is deterministic given $s_{k, i}$, and that it may not be in the p/4-neighbourhood of $\left.s^{\star}\right)$. Then we have:

$$
v^{\circ}\left(s_{k, i}\right)=\left(\prod_{k^{\prime} \in[k]} g\left(s_{k^{\prime}-1}, s_{k^{\prime}}\right)\right) g\left(s_{k}, s_{k+1}\right) g\left(s_{k+1}, s^{\star}\right),
$$

or, overloading notation and letting $g(x)=1-x / p$, we have

$$
\begin{equation*}
v^{\circ}\left(s_{k, i}\right)=\left(\prod_{k^{\prime} \in[k]} g\left(s_{k^{\prime}-1}, s_{k^{\prime}}\right)\right) g\left(c t_{k, i}^{f i p}+e_{k, i}^{\neg f i x}\right) g\left(e_{k, i}^{f i x}\right) . \tag{22}
\end{equation*}
$$

Proof. Identical to WSG21, Lemma 4.9]
Lemma B.5. The value function $v^{\circ}$ is linear in features $\varphi_{v}$ with dimension $d=\Theta\left(p^{2}\right)$.
Proof. Starting from Equation 22, we observe that only $e_{k, i}^{\text {fix }}$ and $e_{k, i}^{\neg \text { fix }}$ depend on $s^{\star}$. The first term in parentheses in simply a scalar which multiplies the features. We thus calculate linear expressions for $x=\mathrm{ct}_{k, i}^{\mathrm{ftip}}+e_{k, i}^{\neg \mathrm{fix}}$ and $y=e_{k, i}^{\mathrm{fix}}$. Starting with $y=e_{k, i}^{\mathrm{fix}}$, we have:

$$
y=\frac{1}{2}\left(\left\langle\overrightarrow{1}, \mathrm{fix}_{k, i}\right\rangle-\left\langle\mathrm{fix}_{k, i} \cdot s_{k, i}, s^{\star}\right\rangle\right)=a+\left\langle b, s^{\star}\right\rangle,
$$

where $a=\frac{1}{2}\left\langle\overrightarrow{1}, \mathrm{fix}_{k, i}\right\rangle$ and $b=-\frac{1}{2} \mathrm{fix}_{k, i} \cdot s_{k, i}$. Similarly:

$$
x=\mathrm{cf}_{k, i}^{\mathrm{flip}}+\frac{1}{2}\left(\left\langle\overrightarrow{1}, \neg \mathrm{fix}_{k, i}\right\rangle-\left\langle\neg \mathrm{fix}_{k, i} \cdot s_{k, i}, w^{\star}\right\rangle\right)=c+\left\langle d, s^{\star}\right\rangle,
$$

where $c=\operatorname{ct}_{k, i}^{\text {fiip }}+\frac{1}{2}\left(\left\langle\overrightarrow{1}, \neg \mathrm{fix}_{k, i}\right\rangle\right)$ and $d=-\frac{1}{2} \neg \mathrm{fix}_{k, i} \cdot s_{k, i}$. Thus we have: $g(y)=1-y / p=$ $1-\left(a+\left\langle d, s^{\star}\right\rangle\right) / p=(1-a / p)-\left\langle\bar{b}, \bar{s}^{\star}\right\rangle=a^{\prime}+\left\langle\bar{b}, \bar{s}^{\star}\right\rangle$, where $a^{\prime}=(1-a / p), \bar{b}=-b / \sqrt{p}$ and $\overline{s^{\star}}=s^{\star} / \sqrt{p}$. Similarly, $g(x)=1-x / p=1-\left(c+\left\langle d, s^{\star}\right\rangle\right) / p=c^{\prime}-\left\langle\bar{d}, \bar{s}^{\star}\right\rangle$, where $c^{\prime}=1-c / p$ and $\bar{d}=d / \sqrt{p}$. Putting this together we have that

$$
\begin{aligned}
v^{\circ}\left(s_{k, i}\right) & =\left(\prod_{k^{\prime} \in[k]} g\left(s_{k^{\prime}-1}, s_{k^{\prime}}\right)\right) g(x) g(y) \\
& =\left(\prod_{k^{\prime} \in[k]} g\left(s_{k^{\prime}-1}, s_{k^{\prime}}\right)\right)\left(a^{\prime}+\left\langle\bar{b}, \bar{s}^{\star}\right\rangle\right)\left(c^{\prime}+\left\langle\bar{d}, \bar{s}^{\star}\right\rangle\right) \\
& =\left(\prod_{k^{\prime} \in[k]} g\left(s_{k^{\prime}-1}, s_{k^{\prime}}\right)\right)\left(a^{\prime} c^{\prime}+c^{\prime}\left\langle\bar{b}, \bar{s}^{\star}\right\rangle+a^{\prime}\left\langle\bar{d}, \bar{s}^{\star}\right\rangle+\left\langle\bar{b}, \bar{s}^{\star}\right\rangle\left\langle\bar{d}, \bar{s}^{\star}\right\rangle\right) \\
& =\left(\prod_{k^{\prime} \in[k]} g\left(s_{k^{\prime}-1}, s_{k^{\prime}}\right)\right)\left(a^{\prime} c^{\prime}+\left\langle c^{\prime} \bar{b}+a^{\prime} \bar{d}, \bar{s}^{\star}\right\rangle+\left\langle\bar{b} \otimes \bar{d}, \bar{s}^{\star} \otimes \bar{s}^{\star}\right\rangle\right)
\end{aligned}
$$

where we have use a property of the tensor product that $\left\langle a_{1}, b_{1}\right\rangle\left\langle a_{2}, b_{2}\right\rangle=\left\langle a_{1} \otimes a_{2}, b_{1} \otimes b_{2}\right\rangle$, where $a_{1} \otimes a_{2}, b_{1} \otimes b_{2} \in \mathbb{R}^{p \times p}$ and their inner product is interpreted as the inner product between the vectorized matrices. Thus, if we take

$$
\theta_{v}=1 \oplus \bar{s}^{\star} \oplus\left(\bar{s}^{\star} \otimes \bar{s}^{\star}\right) \in \mathbb{R}^{1+p+p^{2}}
$$

and

$$
\varphi\left(s_{k, i}\right)=\left(\prod_{k^{\prime} \in[k]} g\left(s_{k^{\prime}-1}, s_{k^{\prime}}\right)\right)\left(a^{\prime} c^{\prime} \oplus\left(c^{\prime} \bar{b}+a^{\prime} \bar{d}\right) \oplus(\bar{b} \otimes \bar{d})\right) \in \mathbb{R}^{1+p+p^{2}}
$$

then we have $v^{\circ}\left(s_{k, i}\right)=\left\langle\varphi\left(s_{k, i}\right), \theta_{v}\right\rangle$ as desired. Thus $v^{\circ}$ is linear with features in dimension $1+p+p^{2}$. Note that the norm of the features and the parameter $\theta_{v}$ is also bounded by constants.

This completes the proof for $v^{\circ}$-linearity. We next tackle the analogous statement for $\pi^{\circ}$-linearity.
Lemma B.6. There exists a feature map $\varphi_{\pi}$ and parameter $\theta_{\pi}$ of dimension $d \approx p$ such that

$$
\pi^{\circ}\left(s_{k, i}\right)=\operatorname{argmax}_{a}\left\{\left\langle\varphi_{\pi}(s, a), \theta_{\pi}\right\rangle\right\}
$$

Proof. Let $s_{k, i}$ be a state of interest. Recall that $\pi^{\circ}$ will either 1) flip the earliest index which has not been fixed such that the value of $s_{k, i}$ at that index disagrees with $s^{\star}$, or 2) if no such index exists, freeze the current round by playing a frozen index. First consider a state $s_{k, i}$ such that $\pi^{\circ}\left(s_{k, i}\right)$ will flip an index. Since the index flipped was previously incorrect and will thereafter agree with $s^{\star}$ on that bit, this corresponds to minimizing the distance between $s^{\prime}$ and $s^{\star}$ amongst all possible $s^{\prime}$ which can be reached in one step from $s_{k, i}$ (i.e. amongst all possible $\left\{\tau\left(s_{k, i}, a\right)\right\}_{a}$, recalling that $\tau\left(s_{k, i}, a\right)$ is our notation for the transition function of the MDP). Thus, $\pi^{\circ}\left(s_{k, i}\right) \in \operatorname{argmin}_{a}\left\{\rho\left(\tau\left(s_{k, i}, a\right), s^{\star}\right)\right\}$. This can be written linearly as $\operatorname{argmin}_{a}\left\{\frac{1}{2}\left(p-\left\langle\tau\left(s_{k, i}, a\right), s^{\star}\right\rangle\right)\right\}=\operatorname{argmax}_{a}\left\{\frac{1}{2}\left(\left\langle\tau\left(s_{k, i}, a\right), s^{\star}\right\rangle-p\right)\right\}=$ $\operatorname{argmax}_{a}\left\{\frac{1}{2}\left(\left\langle p \oplus\left\langle\tau\left(s_{k, i}, a\right), s^{\star}\right\rangle, 1 \oplus s^{\star}\right\rangle\right)\right\}$. The second case is that $\pi^{\circ}$ will freeze the round at the state $s_{k, i}$. This means that there are no indices which are incorrect that have not been frozen in this round. Again, this corresponds to minimizing the distance between $\tau\left(s_{k, i}, a\right)$ and $s^{\star}$ : all other choices will either result in a game over (which has 0 value) or will flip an incorrect bit (which increases the distance). Thus, again we have $\pi^{\circ}\left(s_{k, i}\right) \in \operatorname{argmin}_{a}\left\{\rho\left(\tau\left(s_{k, i}, a\right), s^{\star}\right)\right\}=$ $\operatorname{argmax}_{a}\left\{\frac{1}{2}\left(\left\langle p \oplus\left\langle\tau\left(s_{k, i}, a\right), s^{\star}\right\rangle, 1 \oplus s^{\star}\right\rangle\right)\right\}$.
Thus, in either case, we have that $\pi^{\circ}$ is linear with features $\varphi_{\pi}(s, a)=p \oplus \tau\left(s_{k, i}, a\right)$ and $\theta^{\circ}=1 \oplus s^{\star}$. Since the definition of the argmax is scale-insensitive, we can further normalize to obtain that the features are bounded in magnitude by a constant.

## B. 3 CUbeGame with expert advice, and Proof of Theorem B. 3

Following the approach of [WSG21], we give our lower bound by providing a reduction to an abstract game called CubeGame. Any learning algorithm which can solve the MDPs in our construction can also be used to solve CUBEGAME, and thus it follows that the learner will be subject to the same lower bound. For our setting, we modify the reward function CUBEGAME and augment the learner with the ability to query an expert, which will behave identically to the expert policy which we have defined in our MDPs (that is, it will flip the first bit which is incorrect or give a special actions to indicate if all of the bits are correct). For the rest of this section, when referring to CUBEGAME we are referring to our modified version.

In [WSG21] it is shown that any algorithm which outputs a 0.01 -optimal answer for the (expert-less) CUBEGAME will need a query complexity of $N \geq 2^{\Omega(p \wedge K)}$. In what follows, we will provide the analogous proof of this for our modified game. Thus, the main result of this section is the following, which states that if the learner is not given a budget of $\Omega(p / \log p)$ expert queries, the sample complexity remains exponential. A learning algorithm for CUBEGAME will be called a planner, and a planner which returns a 0.01 -optimal answer at the end of CUBEGAME will be called sound.

Rules of CubeGame CubeGame is a bandit-like game which is defined by two parameters: a length $K \in \mathbb{N}_{+}$and a dimension $p \in \mathbb{N}_{+}$. The "action space" is $W=\{ \pm 1\}^{p}$. Recall that $\rho(x, y)=\frac{1}{2}\left(p-x^{\top} y\right)$ is the Hamming distance between two vectors in $W$. The secret parameter which solves the game is housed in the set $W^{\star}=\left\{w \in W \left\lvert\, \frac{p}{4} \leq \rho(\overrightarrow{1}, w) \leq \frac{3 p}{4}\right.\right\}$. The planner can only input sequences of vectors where each vector is sufficiently far from the previous one. Formally, for any $k \in[K]$, we let $W^{\circ k}=\left\{\left(w_{i}\right)_{i \in[k]} \in W^{k} \mid \forall i \in[k]: \rho\left(w_{i-1}, w_{i}\right) \geq p / 4\right\}$, with $w_{0}:=\overrightarrow{1}$. The action space is: $\mathcal{A}=\cup_{k \in[K]} W^{\circ k}$, thus the planner can input any sequence of length $k \leq K$ satisfying that $\left(w_{i}\right) \in W^{\circ k}$.

The reward function is defined by

$$
\begin{aligned}
f_{w^{\star}}\left(\left(w_{i}\right)_{i \in[k]}\right) & =\left(\prod_{i \in[k]} g\left(\rho\left(w_{i-1}, w_{i}\right)\right)\right) g\left(\rho\left(w_{k}, w^{\star}\right)\right) \\
\text { where: } g(x) & =1-\frac{x}{p}
\end{aligned}
$$

with base case $f(())=g\left(\rho\left(w_{0}, w^{\star}\right)\right)$.
While the planner plays, it chooses sequence lengths $L_{t} \in[K]$ and input sequences $S_{t}=\left(w_{i}^{t}\right)_{i \in\left[L_{t}\right]} \in$ $W^{\circ L_{t}}$. If it chooses to stop playing, it chooses an output $S_{t}=\left(w_{i}^{t}\right)_{i \in[8]} \in W^{\circ 8}$ (we distinguish this
case by letting $L_{t}=0$ denote that the planner has chosen to terminate). Thus the number of actions taken is $N=\min \left\{t \in \mathbb{N}_{+} \mid L_{t}=0\right\}$.
After any action $S_{t}$, the reward is given if either $\rho\left(w_{L_{t}}^{t}, w^{\star}\right)<p / 4$ or if $L_{t}=K$. In both cases the reward is $\operatorname{Ber}\left(f_{w^{\star}}\left(S_{t}\right)\right.$ ), a Bernoulli random variable with mean $f_{w^{\star}}(\cdot)$. Similarly, if the planner is done (i.e. the input is $\left.S_{N} \in W^{\circ 8}\right)$, then the reward given is $R=f_{w^{\star}}\left(\left(w_{i}^{N}\right)_{i \in\left[k^{\star}\right]}\right)$, where $k^{\star}=\min \left\{8, \min \left\{k \mid \rho\left(w_{k}^{N}, w^{\star}\right)<p / 4\right\}\right\}$.
The last thing to specify is the oracle. Here, we allow the planner to query the oracle part-way through a sequence. Namely, if the planner chooses to input a sequence of length $L_{t}<K$, then the oracle can be queried at the end of the sequence, and a second sequence of length $L_{t}^{2} \leq K-L_{t}$ can be inputted. This can be repeated as many times as desired, given that the total sequence length remains $\leq K$. The oracle will simulate the expert policy from before: upon being called at a vector $w_{k, i}$, it will either return the index of the first bit which does not agree with $w^{\star}$, or it will return a special action indicating that all bits are correct.

We are now ready for the main theorem.
Theorem B.7. Any sound planner for CUBEGAME which has less than $\Omega(p / \log p)$ oracle queries, will have a sample complexity

$$
N \geq 2^{\Omega(p \wedge K)}
$$

The proof comes in 5 lemmas, two of which (Lemmas B.8 and B.10) are analogous to properties from the expert-less CubeGame. The other 3 lemmas are information-theoretic and are specific to the oracle setting.

First, some properties about the reward function of CUBEGAME.
Lemma B. 8 (Properties of $f_{w^{\star}}$ ). For any $w^{\star} \in W^{\star}, k \in \mathbb{N}, s=\left(w_{k^{\prime}}\right)_{k^{\prime} \in[k]} \in W^{\circ k}$, we have

$$
\begin{aligned}
& \frac{1}{4} \leq f_{w^{\star}}(()) \leq \frac{3}{4} \\
& 0<f_{w^{\star}}(s) \leq\left(\frac{3}{4}\right)^{k+\mathbb{1}\left[\rho\left(w_{k}, w^{\star}\right) \geq p / 4\right]}
\end{aligned}
$$

Proof. The proof is analogous to [WSG21, Lemma 4.2], substituting our first-order $g$ function.

The following parameters will control our sample complexity:

$$
\begin{align*}
& n=\min \left\{\exp (p / 8) p^{-x} / 20-5,\left(\frac{1}{\varepsilon}-1\right) / 9.9\right\}, \quad \text { where }  \tag{23}\\
& x=\frac{p}{16} \log p, \quad \text { and }  \tag{24}\\
& \varepsilon=\left(\frac{3}{4}\right)^{K+1} \tag{25}
\end{align*}
$$

We will see that $N=\Omega(n)$ for any sound planner.
In the original game of [WSG21], the interaction protocol is captured by $\left(X_{t}, Y_{t}\right)_{t \in[N]}$, where

- $N=\min \left\{t \in \mathbb{N}_{+} \mid L_{t}=0\right\}$ is the interaction length
- $L_{t} \in[K]$ is the input length chosen,
- $S_{t}$ is the sequence $\left(w_{i}\right)_{i \in\left[L_{t}\right]}$ inputted, satisfying $\rho\left(w_{i-1}, w_{i}\right) \geq \frac{p}{4}$
- $X_{t}=\left(L_{t}, S_{t}\right)$,
- $U_{t}=\mathbb{1}\left\{\rho\left(w_{L_{t}-1}, w^{\star}\right)<p / 4\right\}$,
- $V_{t}=\mathbb{1}\left\{\rho\left(w_{L_{t}}, w^{\star}\right)<p / 4\right\}$, and
- $Z_{t}=0$ unless $V_{t}=1$ or $L_{t}=K$, in which case $Z_{t}=\operatorname{Ber}\left(f_{w^{\star}}\left(S_{t}\right)\right)$,
- $Y_{t}=\left(U_{t}, V_{t}, Z_{t}\right)$.

In our case, we need the sequence $S_{t}$ to include every bit flip, thus $S_{t}=\left(w_{k, i}\right)_{k \in\left[L_{t}\right], i \in[p]}$ We also have two new variables, namely $O_{t}=\left(o_{k, i}\right)$ which is an indicator that the oracle was called at step $i$ of phrase $k$ and $E_{t}=\left(e_{k, i}\right)$ which is the answer returned by the oracle. Thus our new interaction protocol is defined by $X_{t}=\left(L_{t}, S_{t}, O_{t}, E_{t}\right)$ and $Y_{t}=\left(U_{t}, V_{t}, Z_{t}\right)$, where $Y_{t}$ remains unchanged.

The planner $A$, with $n$ interactions, in the environment defined by $w^{\star}$, defines a distribution over the environment:
$P_{w^{\star}}^{A, n}\left(\left(X_{t}, O_{t}, E_{t}, Y_{t}\right)_{t}\right)=\prod_{i=1}^{n} p\left(x_{i} \mid x_{1: i-1}, o_{1: i-1}, e_{1: i-1}, y_{1: i-1}\right) p\left(o_{i} \mid x_{1: i}, o_{1: i-1}, e_{1: i-1}, y_{1: i-1}\right) p\left(e_{i} \mid x_{i}, o_{i}\right) p\left(y_{i} \mid x_{i}\right)$.
Note that $p\left(x_{i} \mid x_{1: i-1}, o_{1: i-1}, e_{1: i-1}, y_{1: i-1}\right)$ and $p\left(o_{i} \mid x_{1: i}, o_{1: i-1}, e_{1: i-1}, y_{1: i-1}\right)$ are decisions made by the planner and $p\left(e_{i} \mid x_{i}, o_{i}\right) p\left(y_{i} \mid x_{i}\right)$ and $p\left(y_{i} \mid x_{i}\right)$ are obtained by querying the environment. We define the "abstract game $\left(0, w^{\star}\right)$ " to always yield reward 0 , and which has the same oracle as environment $w^{\star}$. It's distribution will be written as $P_{\left(0, w^{\star}\right)}^{A}$. Let $E_{n}^{w^{\star}}$ be the event that in the first $n$ steps the planner does not hit on any vector that is close to $w^{\star}$ :
$E_{n}^{w^{\star}}=\cap_{t \in[n]}\left\{t>N\right.$ or $\left(t=N\right.$ and $\left.\min _{i \in[8]} \rho\left(w_{i}^{N}, w^{\star}\right) \geq \frac{p}{4}\right)$ or $\left(t<N\right.$ and $\rho\left(w_{L_{t}-1}, w^{\star}\right) \geq \frac{p}{4}$ and $\left.\left.\rho\left(w_{L_{t}}, w^{\star}\right) \geq \frac{p}{4}\right)\right\}$
Lemma B. 9 (A first change of measure). For any planner $A$ and any $w^{\star} \in W$, we have

$$
P_{w^{\star}}^{A}\left(E_{n}^{w^{\star}}\right) \geq \frac{9}{10} P_{\left(0, w^{\star}\right)}^{A}\left(E_{n}^{w^{\star}}\right)
$$

Proof. It will be shown that

$$
\begin{equation*}
P_{w^{\star}}^{A}\left(E_{n}^{w^{\star}}\right) \geq(1-\varepsilon)^{n} P_{\left(0, w^{\star}\right)}^{A}\left(E_{n}^{w^{\star}}\right) . \tag{27}
\end{equation*}
$$

Since $1-\varepsilon \geq 1-\frac{1}{1+9.9 n}$ by our definition of $n$ we have that $(1-\varepsilon)^{n} \geq\left(1-\frac{1}{1+9.9 n}\right)^{n} \geq$ $\lim _{n \rightarrow \infty}\left(1-\frac{1}{1+9.9 . n}\right)^{n}>9 / 10$, and thus Equation (27) implies our result.

Let $\mathcal{H}$ be the set of all possible histories $\left(x_{t}, o_{t}, e_{t}, y_{t}\right)$ of length $n$, and $E_{h}=E_{n}^{w^{\star}} \cap\{H=h\}$. Note that $E_{n}^{w^{\star}}$ is the disjoint union of $E_{h}$, so it is enough to show that for each $h$ we have:

$$
\rho=\frac{P_{w^{\star}}^{A}\left[E_{h}\right]}{P_{\left(0, w^{\star}\right)}^{A}\left[E_{h}\right]} \geq(1-\varepsilon)^{n}
$$

for each $h$ such that $P_{\left(0, w^{\star}\right)}^{A}\left(E_{h}\right)>0$. So, let $h=\left(x_{t}, o_{t}, e_{t}, y_{t}\right)$ be such that $P_{\left(0, w^{\star}\right)}^{A}\left(E_{h}\right)>0$. This implies in particular that $y_{t}=(0,0,0) \forall t$. Now, both $P_{w^{\star}}^{A}$ and $P_{\left(0, w^{\star}\right)}^{A}$ factorize according to Eq. (26), giving:

$$
\frac{P_{w^{\star}}^{A}\left[E_{h}\right]}{P_{\left(0, w^{\star}\right)}^{A}\left[E_{h}\right]}=\prod_{i} \frac{p\left(x_{i} \mid x_{1: i-1}, o_{1: i-1}, e_{1: i-1}, y_{1: i-1}\right) p\left(o_{i} \mid x_{1: i}, o_{1: i-1}, e_{1: i-1}, y_{1: i-1}\right) p^{w^{\star}}\left(e_{i} \mid x_{i}, o_{i}\right) p^{w^{\star}}\left(y_{i} \mid x_{i}\right)}{p^{A}\left(x_{i} \mid x_{1: i-1}, o_{1: i-1}, e_{1: i-1}, y_{1: i-1}\right) p^{A}\left(o_{i} \mid x_{1: i}, o_{1: i-1}, e_{1: i-1}, y_{1: i-1}\right) p^{\left(0, w^{\star}\right)}\left(e_{i} \mid x_{i}, o_{i}\right) p^{\left(0, w^{\star}\right)}\left(y_{i} \mid x_{i}\right)} .
$$

Since we are conditioning on the same fixed history, the terms involving decisions made by the planner will cancel, and similarly the variable $E_{i}$ also behaves the same in both environments (since the oracle for $w^{\star}$ is the same). We are left with:

$$
\rho=\prod_{i} \frac{p^{w^{\star}}\left(y_{i}=(0,0,0) \mid x_{i}\right)}{p^{\left(0, w^{\star}\right)}\left(y_{i}=(0,0,0) \mid x_{i}\right)}
$$

The denominator always has probability 1 in the environment $\left(0, w^{\star}\right)$, and since $U_{t}=V_{t}=0$ under the set $E_{n}^{w^{\star}}$ (the planner is never close to $w^{\star}$ ), we have $Y_{t}=0 \Longleftrightarrow Z_{t}=0$, so it remains to control

$$
\rho=\prod_{i=1}^{n} P_{w^{\star}}^{A}\left(Z_{t}=0 \mid x_{t}\right)
$$

Again since the planner is never close to $w^{\star}, Z_{t}=1$ only if $l_{t}=K$, in which case we have $P_{w^{\star}}^{A}\left(Z_{t}=1 \mid x_{t}\right) \geq(3 / 4)^{K+1}=\varepsilon$ by definition of the reward $f_{w^{\star}}$ obtained from reaching level $K$.

The next lemma simply bounds the number of vectors in $W$ which are close to any fixed vector.
Lemma B. 10 (Hypercube counting). For $\tilde{w} \in W$, let $\mathrm{W}_{\text {close }}(\tilde{w})=\{w \in W \mid \rho(w, \tilde{w})<p / 4\}$.
Then $\left|\mathrm{W}_{\text {close }}(\tilde{w})\right| \leq 2^{p} \exp (-p / 8)$
Proof. Identical to [WSG21, Lemma 4.4].
Recall that $E_{n}^{w^{\star}}$ is the "bad event" for the planner. We study its complement, $\left(E_{n}^{w^{\star}}\right)^{c}$, which satisfies $\left(E_{n}^{w^{\star}}\right)^{c} \subset\left\{w^{\star} \in Z\right\}$, where

$$
Z:=\cup_{t \in[n \wedge(N-1)]}\left(\mathrm{W}_{\text {close }}\left(w_{L_{t}-1}^{t}\right) \cup \mathrm{W}_{\text {close }}\left(w_{L_{t}}^{t}\right) \cup\left(\cup_{i \in[8]} \mathrm{W}_{\text {close }}\left(w_{i}^{N}\right)\right),\right.
$$

i.e. the event that for some $t$ we have $\rho\left(w_{L_{t}-1}^{t}, w^{\star}\right)<p / 4$ or $\rho\left(w_{L_{t}}^{t}, w^{\star}\right)<p / 4$ or that for some $i \in[8]$ we have $\rho\left(w_{i}^{N}, w^{\star}\right)<p / 4$. We define the "abstract game" $(0,0)$, where the planner has access to an oracle but, when queried, rather than giving information about the "true" $w^{\star}$, the oracle will simply return a uniformly random bit in $[p]$.
Lemma B. 11 (A second change of measure). For any planner $A$ with an oracle budget of $x$, we have that:

$$
P_{0, \hat{w}}^{A}(\hat{w} \in Z)=p^{x} P_{0,0}^{A}(\hat{w} \in Z)
$$

Proof. As before, consider the set of histories $\mathcal{H}$, and let $Z_{h}=Z \cap\{H=h\}$. Writing out the importance ratio gives:
$\rho=\frac{P_{0, w^{\star}}^{A}\left[Z_{h}\right]}{P_{(0,0)}^{A}\left[Z_{h}\right]}=\prod_{i} \frac{p\left(x_{i} \mid x_{1: i-1}, o_{1: i-1}, e_{1: i-1}, y_{1: i-1}\right) p\left(o_{i} \mid x_{1: i}, o_{1: i-1}, e_{1: i-1}, y_{1: i-1}\right) p^{0, w^{\star}}\left(e_{i} \mid x_{i}, o_{i}\right) p^{\left(0, w^{\star}\right)}\left(y_{i} \mid x_{i}\right)}{p\left(x_{i} \mid x_{1: i-1}, o_{1: i-1}, e_{1: i-1}, y_{1: i-1}\right) p\left(o_{i} \mid x_{1: i}, o_{1: i-1}, e_{1: i-1}, y_{1: i-1}\right) p^{(0,0)}\left(e_{i} \mid x_{i}, o_{i}\right) p^{(0,0)}\left(y_{i} \mid x_{i}\right)}$
Again, as before, all the terms involving the planner will cancel, since they are conditioned on the same history and thus the planner will make the same decisions. Similarly, in both games the reward is deterministically 0 thus the $p^{\left(0, w^{\star}\right)}\left(y_{i} \mid x_{i}\right)=p^{(0,0)}\left(y_{i} \mid x_{i}\right)$. We are left with

$$
\rho=\prod_{i=1}^{n} \frac{p^{\left(0, w^{\star}\right)}\left(E_{i}=e_{i} \mid x_{i}, o_{i}\right)}{p^{(0,0)}\left(E_{i}=e_{i} \mid x_{i}, o_{i}\right)}
$$

Note that the top probability is deterministic (since the true expert is) and is only equal to 1 at most $x$ times (recalling that $x$ is the total number of oracle calls allowed). We are simply left with the (inverse of the) probability that the random oracle returns any given answer, which is $1 / p$. Thus we end up with $\rho=p^{x}$.
Lemma B. 12 (Finding a bad $w^{\star}$ for planner $A$ ). For any abstract planner $A$ there exists $w^{\star} \in W^{\star}$ such that

$$
P_{w^{\star}}^{A}\left(E_{n}^{w^{\star}}\right) \geq(9 / 10)^{2}
$$

Proof. Note that, by Lemma B. 9 it is sufficient to show that $P_{0, w^{\star}}^{A}\left(\left(E_{n}^{w^{\star}}\right)^{c}\right) \leq \frac{1}{10}$. Recall that for any $\hat{w} \in W^{\star}$ we have $\left(E_{n}^{\hat{w}}\right)^{c} \subseteq\{\hat{w} \in Z\}$, where

$$
Z:=\cup_{t \in[n \wedge(N-1)]}\left(\mathrm{W}_{\text {close }}\left(w_{L_{t}-1}^{t}\right) \cup \mathrm{W}_{\text {close }}\left(w_{L_{t}}^{t}\right) \cup\left(\cup_{i \in[8]} \mathrm{W}_{\text {close }}\left(w_{i}^{N}\right)\right)\right.
$$

By a union bound and Lemma B. 10 we have that $|Z| \leq(2 n+8) 2^{p} \exp (-p / 8)$. Since $W^{\star}=$ $W \backslash \mathrm{~W}_{\text {close }}(1) \backslash \mathrm{W}_{\text {close }}(-1)$, this also gives that $\left|W^{\star}\right| \geq 2^{p}(1-2 \exp (-p / 8))$. We pick $w^{\star}$ according to:

$$
w^{\star}=\operatorname{argmin}_{\hat{w} \in W^{\star}} P_{0, w^{\star}}^{A}(\hat{w} \in Z)
$$

Putting things together and using Lemma B. 11 gives:

$$
\begin{aligned}
& 2^{p}(1-2 \exp (-p / 8)) P_{0, w^{\star}}^{A}\left(w^{\star} \in Z\right) \leq\left|W^{\star}\right| P_{0, \hat{w}}^{A}\left(w^{\star} \in Z\right) \\
& \leq \sum_{\hat{w} \in W^{\star}} P_{0, \hat{w}}^{A}(\hat{w} \in Z) \leq \sum_{\hat{w} \in W} P_{0, \hat{w}}^{A}(\hat{w} \in Z) \leq \sum_{\hat{w} \in W} p^{x} P_{0,0}^{A}(\hat{w} \in Z) \\
& =p^{x} \sum_{\hat{w} \in W} P_{0,0}^{A}(\hat{w} \in Z)=p^{x} \sum_{\hat{w} \in W} \mathbb{E}_{0,0}^{A}[\mathbb{1}[\hat{w} \in Z]]=p^{x} \mathbb{E}_{0,0}^{A}\left[\sum_{\hat{w} \in W} \mathbb{1}[\hat{w} \in Z]\right] \\
& =p^{x} \mathbb{E}_{0,0}^{A}[|Z|] \leq p^{x}(2 n+8) 2^{p} \exp (-p / 8)
\end{aligned}
$$

Rearranging gives that

$$
P_{0, w^{\star}}^{A}\left(\left(E_{n}^{w^{\star}}\right)^{c}\right) \leq P_{0, w^{\star}}^{A}\left(w^{\star} \in Z\right) \leq \frac{(2 n+8) 2^{p} \exp (-p / 8) p^{x}}{2^{p}(1-2 \exp (-p / 8)} \leq 2(n+5) p^{x} \exp (-p / 8) \leq \frac{1}{10}
$$

where the last line followed from our bound on $n$ (Eq. (23)).
We are now ready to prove Theorem B. 7 In fact, there is not much left to do.
Proof (of Theorem B.7). Let the planner $A$ be sound and have an expected query cost $\bar{N}$, and $w^{\star}$ the vector from the previous lemma. Then by Markov's inequality we have:

$$
P_{\left(0, w^{\star}\right)}^{A}[N-1 \geq n] \leq \frac{\bar{N}}{n}
$$

Letting $E^{\prime}=E_{n}^{w^{\star}} \cap\{N-1<n\}$ we have

$$
P_{0, w^{\star}}^{A}\left[E^{\prime}\right] \geq(9 / 10)^{2}-\frac{\bar{N}}{n}
$$

Under event $E^{\prime}$, the output of the planner satisfies $\rho\left(w_{i}^{N}, w^{\star}\right) \geq p / 4$ for all $i \in[8]$, so the reward at the end of the game is $R<(3 / 4)^{9}$. Combined with soundness this gives

$$
\begin{aligned}
\frac{1}{4}-0.01 \leq f_{w^{\star}}(())-0.01 \leq \mathbb{E}_{w^{\star}}^{A}[R] & \leq\left(\frac{3}{4}\right)^{9}+\left(1-P_{0, w^{\star}}^{A}\left[E^{\prime}\right]\right) \frac{3}{4} \\
& \leq\left(\frac{3}{4}\right)^{9}+\left(1-(9 / 10)^{2}\right) \frac{3}{4}+\frac{\bar{N}}{n} \frac{3}{4}
\end{aligned}
$$

which requires $\bar{N}>0.02 n$, namely

$$
N>0.02 \min \left\{\exp (p / 8) p^{-x} / 16-5,\left(\frac{1}{\varepsilon}-1\right) / 7.5\right\}
$$

Lastly, note that when $x \leq \frac{p}{16 \log p}$ we have

$$
\log n=\log (\exp (p / 8))-\log \left(p^{x}\right)=\frac{p}{8}-x \log p \geq \frac{p}{8}-\frac{p}{16 \log p} \log p=\frac{p}{16},
$$

thus $n=\Omega(\exp (p / 16))$ and in particular $N=\Omega\left\{2^{p \wedge K}\right\}$.
To prove Theorem B.3 the last thing to show is that a learner which solves the MDP can be used to solve CubeGame. This reduction follows exactly as in Section 4.8 of [WSG21]

## C On $\pi^{\circ}$ linearity

This section shows that, when $\pi^{\circ} \neq \pi^{\star}, q^{\circ}$ can be linear with $d$-dimensional features yet these features do not realize $\pi^{\circ}$-linearity.

The MDP is as follows: the states are arranged in a binary tree of length $H$. The action space is $\{\ell, r\}$, corresponding to the $\ell$ eft and right actions. Transitions are deterministic. The reward for every $\ell$ eft action is -1 , the reward for every right action is +1 . See Figure 1


Figure 1: $q^{\circ}$-linearity does not imply $\pi^{\circ}$-linearity
Note that we can identify every state with the action sequence that led to it (with the starting state corresponding to the empty sequence). The policy $\pi^{\circ}$ is defined such that, if $s=\left(a_{0}, \ldots, \ell\right)$ then $\pi^{\circ}(s)=\mathrm{r}$ and otherwise if $s=\left(a_{0}, \ldots, \mathrm{r}\right)$ then $\pi^{\circ}(s)=\ell$. Thus the policy will alternate the action taken at every step. This defines the $q^{\circ}$ function:

$$
q^{\circ}(s, a)= \begin{cases}0, & \text { if } h=0 \bmod 2 \\ -1, & \text { if } h=1 \bmod 2 \text { and } a=\ell \\ +1, & \text { if } h=1 \bmod 2 \text { and } a=\mathrm{r}\end{cases}
$$

Note that we can linearize this in one dimension via the features $\varphi^{\circ}(s, a)=q^{\circ}(s, a)$ and $\theta=1$. However, these features do not linearly-realize $\pi^{\circ}$ : since $\theta>1$ then the argmax at every odd horizon will always be the right action (since $\varphi^{\circ}(s, r)=1$ and $\varphi^{\circ}(s, \ell)=-1$ ).
Note that $\pi^{\prime}=\operatorname{argmax}_{a}\left\{q^{\circ}(s, a)\right\}$, the greedy policy derived from $q^{\circ}$ is by definition linear with those features. For the special case where $\pi^{\circ}=\pi^{\star}$ then the greedy policy $\pi^{\prime}$ lines up with the policy $\pi^{\circ}$, so we get linearity for free in that case.


[^0]:    ${ }^{7}$ (and using that $\mathcal{F} \subseteq \mathcal{F}^{\prime} \Longrightarrow \operatorname{dim}_{E}(\mathcal{F}, \varepsilon) \leq \operatorname{dim}_{E}\left(\mathcal{F}^{\prime}, \varepsilon\right)$ )

