## A Omitted Related Work

Neural Kernel Bandits. [66] initiated the study of kernelized linear bandits, showing regret dependent on the information gain. The work of [76, 73] specialized this to the Neural Tangent Kernel (NTK) [54, 24, 40, 27, 10], where the algorithm utilizes gradient descent but remains close to initialization and thus remains a kernel class. Furthermore, NTK methods require $d^{p}$ samples to express a degree $p$ polynomial in $d$ dimensions [31], similar to eluder dimension of polynomials, and so lack the inductive biases necessary for real-world applications of decision-making problems [60].

Neural Bandits. For bandits with practical neural networks (instead of overparameterized NTKs) as the function approximator, we are not aware of any previous paper that gives provably efficient bandit algorithm for this case. Our paper gives the first provably efficient algorithm for neural bandits with noiseless reward and deterministic activation. We note that however, previous paper has already solved neural bandits when the neural network happens to be convex [21].

Concave Bandits. There has been a rich line of work on concave bandits starting with [26, 46]. [4] attained the first $\sqrt{T}$ regret algorithm for concave bandits though with a large poly $(d)$ dependence. In the adversarial setting, a line of work [38, 14, 49] have attained polynomial-time algorithms with $\sqrt{T}$ regret with increasingly improved dimension dependence. The sharp dimension dependence remains unknown.

Noiseless Bandits. In the noiseless setting, there is some investigation in phase retrieval borrowing the tools from algebraic geometry (see e.g. [69]). In this paper, we will study the bandit problem with more general reward functions: neural nets with polynomial activation (structured polynomials) including phase retrieval. [45] study similar structured polynomials, also using tools from algebraic geometry, but they only study the expressivity of those polynomials and do not consider the learning problems. [21] study noiseless bandits with bounded Sequential Rademacher Complexity, but focus on attaining local optimality.

Concurrent work. [50] address the phase retrieval bandit problem which is equivalent to a symmetric rank 1 variant of the bilinear bandit of [43] and attain $\widetilde{O}\left(\sqrt{d^{2} T}\right)$ regret. Our work in Section 3.1 specialized to the rank 1 case attains the same regret.

Matrix/Tensor Power Method. Our analysis stems from noisy power methods for matrix/tensor decomposition problems. Robust power method, subspace iteration, and tensor decomposition that tolerate noise first appeared in [37, 8]. Follow-up work attained the optimal rate for both gapdependence and gap-free settings for matrix decomposition [58, 6]. An improvement on the problem dimension for tensor power method is established in [70]. [63] considers the convergence of tensor power method in the non-orthogonal case.

## B Additional Preliminaries

In this section we show that adapting the eluder UCB algorithms from [62] would yield the sample complexity in Theorem 2.1. Especially we give the rates in Table 1 for our stochastic settings.

The algorithm [62] consider Algorithm 2 for the stochastic generalized linear bandit problem. Assume that $\boldsymbol{\theta}^{*}$ is the true parameter of the reward model. The reward is $r_{t}=f_{\boldsymbol{\theta}^{*}}\left(\boldsymbol{a}_{t}\right)+\eta_{t}$ for $f_{\boldsymbol{\theta}^{*}} \in \mathcal{F}$. Let $N$ be the $\alpha$-covering-number (under $\|\cdot\|_{\infty}$ ) of $\mathcal{F}, d_{E}$ be the $\alpha$-eluder-dimension of $\mathcal{F}$ (see Definition 3,4 in [62]). Let $C=\sup _{f \in \mathcal{F}, a \in \mathcal{A}}|f(a)|$. We set $\alpha=\frac{1}{T^{2}}$ in the algorithm.

The regret analysis Choosing $\alpha=1 / T^{2}$, proposition 4 in [62] state that with probability $1-\delta$, for some universal constant $C$, the total regret $\mathfrak{R}(T) \leq \frac{1}{T}+C \min \left\{d_{E}, T\right\}+4 \sqrt{d_{E} \beta_{T} T} \leq 1+$ $C \sqrt{d_{E} T}+4 \sqrt{d_{E} \beta_{T} T}=O\left(\sqrt{d_{E}\left(1+\beta_{T}\right) T}\right)$. In our settings with $\alpha=1 / T^{2}, \beta_{T}=8 \log (N / \delta)+$ $2\left(8 C+\sqrt{8 \ln \left(4 T^{2} / \delta\right)}\right) / T=O(\log (N / \delta))$ where $\log (N)=\Omega(1)$ for our action sets, and thus

$$
\mathfrak{R}(T)=\widetilde{O}\left(\sqrt{d_{E} T \log N}\right)
$$

```
Algorithm 2 Eluder UCB
    Input: Function class \(\mathcal{F}\), failure probability \(\delta\), parameters \(\alpha, N, C\).
    Initialization: \(\mathcal{F}_{0} \leftarrow \mathcal{F}\).
    for \(t\) from 1 to \(T\) do
        Select Action:
        \(\boldsymbol{a}_{t} \in \arg \max _{\boldsymbol{a} \in \mathcal{A}} \sup _{f_{\boldsymbol{\theta}} \in \mathcal{F}_{t-1}} f_{\boldsymbol{\theta}}(\boldsymbol{a})\)
        Play action \(\boldsymbol{a}_{t}\) and observe reward \(r_{t}\)
        Update Statistics:
        \(\widehat{\boldsymbol{\theta}}_{t} \in \arg \min _{\boldsymbol{\theta}} \sum_{s=1}^{t}\left(f_{\boldsymbol{\theta}}\left(\boldsymbol{a}_{s}\right)-r_{s}\right)^{2}\)
        \(\beta_{t} \leftarrow 8 \log (N / \delta)+2 \alpha t\left(8 C+\sqrt{8 \ln \left(4 t^{2} / \delta\right)}\right)\)
        \(\mathcal{F}_{t} \leftarrow\left\{f_{\boldsymbol{\theta}}: \sum_{s=1}^{t}\left(f_{\boldsymbol{\theta}}-f_{\widehat{\boldsymbol{\theta}}_{t}}\right)^{2}\left(\boldsymbol{a}_{s}\right) \leq \beta_{t}\right\}\)
```

Applications in our settings We show that in our settings Theorem 2.1 will obtain the rates listed in Table 1

## The covering numbers

Lemma B.1. The log-covering-number (of radius $\alpha$ with $\alpha \ll 1$, under $\|\cdot\|_{\infty}$ ) of the function classes are: $\log N\left(\mathcal{F}_{S Y M}\right)=O\left(d k \log \frac{k}{\alpha}\right), \log N\left(\mathcal{F}_{A S Y M}\right)=O\left(d k \log \frac{k}{\alpha}\right), \log N\left(\mathcal{F}_{E V}\right)=O\left(d k \log \frac{k}{\alpha}\right)$, and $\log N\left(\mathcal{F}_{L R}\right)=O\left(d k \log \frac{k}{\alpha}\right)$.

Proof. Let $S_{\xi}^{d}$ denote a minimal $\xi$-covering of $\mathbb{S}^{d-1}$ (under $\|\cdot\|_{2}$ ) for $0<\xi<\frac{1}{10}$, and $\left|S_{\xi}^{d}\right|=$ $O(d \log 1 / \xi)$ (see for example [62]). Then we can construct the coverings in our settings from $S_{\xi}^{d}$ :

- $\mathcal{F}_{\text {SYM }}$ : let $\xi=\frac{\alpha}{k p}$, and for $k$ copies of $S_{\xi}^{d}$, we can construct a covering of $\mathcal{F}_{\text {SYM }}$ with size $\left|S_{\xi}^{d}\right|^{k}$. Specifically, let the covering be $S_{\text {SYM }}=\left\{g(\boldsymbol{a})=\sum_{j=1}^{k} \lambda_{j}\left(\boldsymbol{u}_{j}^{\top} \boldsymbol{a}\right)^{p}\right.$ : $\left.\left(\boldsymbol{u}_{1}, \boldsymbol{u}_{2}, \cdots, \boldsymbol{u}_{k}\right) \in S_{\xi}^{d} \times S_{\xi}^{d} \times \cdots \times S_{\xi}^{d}\right\}$, then for each $f(\boldsymbol{a})=\sum_{j=1}^{k} \lambda_{j}\left(\boldsymbol{v}_{j}^{\top} \boldsymbol{a}\right)^{p} \in \mathcal{F}_{\mathrm{SYM}}$, as we can find $\boldsymbol{u}_{j} \in S_{\xi}^{d}$ that $\left\|\boldsymbol{u}_{j}-\boldsymbol{v}_{j}\right\|_{2} \leq \xi$,

$$
\sup _{\boldsymbol{a}}[f(\boldsymbol{a})-g(\boldsymbol{a})] \leq \sup _{\boldsymbol{a}}\left[\sum_{j=1}^{k}\left|\lambda_{j}\left\|\boldsymbol{u}_{j}^{\top} \boldsymbol{a}-\boldsymbol{v}_{j}^{\top} \boldsymbol{a}\right\| \sum_{q=0}^{p-1}\left(\boldsymbol{u}_{j}^{\top} \boldsymbol{a}\right)^{q}\left(\boldsymbol{v}_{j}^{\top} \boldsymbol{a}\right)^{p-q-1}\right|\right] \leq p k \xi=\alpha
$$

- $\mathcal{F}_{\mathrm{ASYM}}:$ let $\xi=\frac{\alpha}{k p}$, and for $k p$ copies of $S_{\xi}^{d}$, let the covering be $S_{\mathrm{ASYM}}=\{g(\boldsymbol{a})=$ $\sum_{j=1}^{k} \lambda_{j} \prod_{q=1}^{p}\left(\boldsymbol{u}_{j}(q)^{\top} \boldsymbol{a}(q)\right):\left(\boldsymbol{u}_{1}(1), \boldsymbol{u}_{1}(2), \cdots, \boldsymbol{u}_{1}(p), \boldsymbol{u}_{2}(1), \cdots, \boldsymbol{u}_{k}(p)\right) \in S_{\xi}^{d} \times$ $\left.S_{\xi}^{d} \times \cdots \times S_{\xi}^{d}\right\}$ with size $\left|S_{\xi}^{d}\right|^{k p}$. Then for each $f(\boldsymbol{a})=\sum_{j=1}^{k} \lambda_{j} \prod_{q=1}^{p}\left(\boldsymbol{v}_{j}(q)^{\top} \boldsymbol{a}(q)\right) \in$ $\mathcal{F}_{\text {ASYM }}$, as we can find $\boldsymbol{u}_{j}(q) \in S_{\xi}^{d}$ that $\left\|\boldsymbol{u}_{j}(q)-\boldsymbol{v}_{j}(q)\right\|_{2} \leq \xi$,

$$
\begin{aligned}
\sup _{\boldsymbol{a}}[f(\boldsymbol{a})-g(\boldsymbol{a})] & \leq \sup _{\boldsymbol{a}}\left[\sum_{j=1}^{k}\left|\lambda_{j}\right| \sum_{q=1}^{p}\left|\boldsymbol{u}_{j}(q)^{\top} \boldsymbol{a}-\boldsymbol{v}_{j}(q)^{\top} \boldsymbol{a}\right| \cdot\right. \\
& \left.\leq \prod_{r<q}\left(\boldsymbol{u}_{j}(r)^{\top} \boldsymbol{a}\right) \prod_{r>q}\left(\boldsymbol{v}_{j}(r)^{\top} \boldsymbol{a}\right) \mid\right] \\
& \leq \operatorname{pk}=\alpha ;
\end{aligned}
$$

- $\mathcal{F}_{\mathrm{EV}}$ : the construction follows that of $\mathcal{F}_{\mathrm{SYM}}$ by taking $p=2$;
- $\mathcal{F}_{\mathrm{LR}}$ : taking the construction of $\mathcal{F}_{\mathrm{SYM}}$ with $p=2$ and $\xi=\frac{\alpha}{2 k}$, for $\boldsymbol{N}=$ $\sum_{j=1}^{k} \lambda_{j} \boldsymbol{u}_{j} \boldsymbol{u}_{j}^{\top}$ and $\boldsymbol{M}=\sum_{j=1}^{k} \lambda_{j} \boldsymbol{v}_{j} \boldsymbol{v}_{j}^{\top}$ with $\left\|\boldsymbol{u}_{j}-\boldsymbol{v}_{j}\right\|_{2} \leq \xi$, we know $\|\boldsymbol{N}-\boldsymbol{M}\|_{\mathrm{F}} \leq$ $\left\|\boldsymbol{N}-\sum_{j=1}^{k} \lambda_{j} \boldsymbol{u}_{j} \boldsymbol{v}_{j}^{\top}\right\|_{\mathrm{F}}+\left\|\sum_{j=1}^{k} \lambda_{j} \boldsymbol{u}_{j} \boldsymbol{v}_{j}^{\top}-\boldsymbol{M}\right\|_{\mathrm{F}} \leq \sum_{j=1}^{k} 2\left|\lambda_{j}\right| \xi \leq \alpha$. Then $\sup _{\boldsymbol{A}}\left[f_{\boldsymbol{M}}(\boldsymbol{A})-f_{\boldsymbol{N}}(\boldsymbol{A})\right] \leq \sup _{\boldsymbol{A}}\|\boldsymbol{M}-\boldsymbol{N}\|_{\mathrm{F}} \cdot\|\boldsymbol{A}\|_{\mathrm{F}} \leq \alpha$.

Then we can bound the covering numbers in Theorem 2.1 Notice that in the settings the log-covering numbers are only different by constant factors.

## The eluder dimensions

Lemma B.2. The $\epsilon$-eluder-dimension $(\epsilon<1) d_{E}$ of the function classes are: $d_{E}\left(\mathcal{F}_{S Y M}\right)=\widetilde{\Theta}\left(d^{p}\right)$ $($ for $k \geq p), d_{E}\left(\mathcal{F}_{A S Y M}\right)=\widetilde{\Theta}\left(d^{p}\right), d_{E}\left(\mathcal{F}_{E V}\right)=\widetilde{\Theta}\left(d^{2}\right)$, and $d_{E}\left(\mathcal{F}_{L R}\right)=\widetilde{\Theta}\left(d^{2}\right)$. In the settings WLOG we assume the top eigenvalue is $r^{*}=\lambda_{1}=1$ as we are mostly interested in the cases where $r^{*}>\epsilon$.

Proof. The upper bounds for the eluder dimension can be given by the linear argument. [62] show that the $d$-dimension linear model $\left\{f_{\boldsymbol{\theta}}(\boldsymbol{a})=\boldsymbol{\theta}^{\top} \boldsymbol{a}\right\}$ has $\epsilon$-eluder-dimension $O\left(d \log \frac{1}{\epsilon}\right)$. In all of these settings, we can find feature maps $\phi$ and $\psi$ so that $\mathcal{F}=\left\{f_{\boldsymbol{\theta}}(\boldsymbol{a}), f_{\boldsymbol{\theta}}(\boldsymbol{a})=\phi(\boldsymbol{\theta})^{\top} \psi(\boldsymbol{a}),\|\phi(\boldsymbol{\theta})\|_{2} \leq\right.$ $\left.k,\|\psi(\boldsymbol{a})\|_{2} \leq k\right\}$. Then the eluder dimensions will be bounded by the corresponding linear dimension as an original $\epsilon$-independent sequence $\left\{\boldsymbol{a}_{i}\right\}$ will induce an $\epsilon$-independent sequence $\left\{\psi\left(\boldsymbol{a}_{i}\right)\right\}$ in the linear model. Therefore for matrices $\left(\mathcal{F}_{\mathrm{LR}}\right.$ and $\left.\mathcal{F}_{\mathrm{EV}}\right)$ the eluder dimension is $O\left(d^{2} \log \frac{k}{\epsilon}\right)$ and for the tensors $\left(\mathcal{F}_{\text {SYM }}\right.$ and $\left.\mathcal{F}_{\mathrm{ASYM}}\right)$ it is $O\left(d^{p} \log \frac{k}{\epsilon}\right)$.
Then we consider the lower bounds. We provide the following example of $O(1)$-independent sequences to bound the eluder dimension in our settings up to a log factor.

- $\mathcal{F}_{\mathrm{SYM}}$ : the sequence is $\left\{\boldsymbol{a}_{i}=\left(\boldsymbol{e}_{i_{1}}, \boldsymbol{e}_{i_{2}}, \cdots, \boldsymbol{e}_{i_{p}}\right): i=\left(i_{1}, i_{2}, \cdots, i_{p}\right) \in[d]^{p}\right\}$. For $f_{j}(\boldsymbol{a})=\prod_{q=1}^{p} \boldsymbol{e}_{j_{q}}^{\top} \boldsymbol{a}(q), f_{j}\left(\boldsymbol{a}_{i}\right)$ is only 1 when $i=j$ and 0 otherwise. Then each $\boldsymbol{a}_{i}$ is 1-independent to the predecessors on $f_{i}$ and zero, and thus the eluder dimension is lower bounded by $d^{p}$.
- $\mathcal{F}_{\mathrm{ASYM}}$ : for $p \leq d$ and $k \geq p$, the sequence is $\left\{\boldsymbol{a}_{i}=\frac{1}{\sqrt{p}}\left(\boldsymbol{e}_{i_{1}}+\boldsymbol{e}_{i_{2}}+\cdots+\boldsymbol{e}_{i_{p}}\right): i=\right.$ $\left.\left(i_{1}, i_{2}, \cdots, i_{p}\right) \in[d]^{p}, i_{1}<i_{2}<\cdots<i_{p}\right\}$. There are tensors $f_{j}$ and $g_{j}$ of CP-rank $k$ that $\left(f_{j}-g_{j}\right)(\boldsymbol{a})=\prod_{q=1}^{p}\left(\boldsymbol{e}_{j_{q}}^{\top} \boldsymbol{a}\right)$ where $j_{1}<j_{2}<\cdots<j_{p},\left(f_{j}-g_{j}\right)\left(\boldsymbol{a}_{i}\right)$ is only 1 when $i=j$ and 0 otherwise. Then each $\boldsymbol{a}_{i}$ is 1-independent to the predecessors on $f_{i}$ and $g_{i}$, and thus the eluder dimension is lower bounded by $\binom{d}{p}$.
- $\mathcal{F}_{\mathrm{EV}}$ : the sequence is $\left\{\boldsymbol{a}_{i}=\frac{1}{\sqrt{2}}\left(\boldsymbol{e}_{i_{1}}+\boldsymbol{e}_{i_{2}}\right): i=\left(i_{1}, i_{2}\right) \in[d]^{2}, i_{1} \leq i_{2}\right\}$. For $f_{j}(\boldsymbol{a})=$ $\frac{1}{2} \boldsymbol{a}^{\top}\left(\boldsymbol{e}_{j_{1}}+\boldsymbol{e}_{j_{2}}\right)\left(\boldsymbol{e}_{j_{1}}+\boldsymbol{e}_{j_{2}}\right)^{\top} \boldsymbol{a}$ and $g_{j}(\boldsymbol{a})=\frac{1}{2} \boldsymbol{a}^{\top}\left(\boldsymbol{e}_{j_{1}}-\boldsymbol{e}_{j_{2}}\right)\left(\boldsymbol{e}_{j_{1}}-\boldsymbol{e}_{j_{2}}\right)^{\top} \boldsymbol{a}$ with $j_{1} \leq j_{2}$, $\left(f_{j}-g_{j}\right)\left(\boldsymbol{a}_{i}\right)$ is only 1 when $i=j$ and 0 otherwise. Then each $\boldsymbol{a}_{i}$ is 1-independent to the predecessors on $f_{i}$ and $g_{i}$, and thus the eluder dimension is lower bounded by $\binom{d}{2}$.
- $\mathcal{F}_{\mathrm{LR}}$ : the sequence is $\left\{\boldsymbol{A}_{i}=\frac{1}{2} \boldsymbol{e}_{i_{1}} \boldsymbol{e}_{i_{2}}^{T}+\boldsymbol{e}_{i_{1}} \boldsymbol{e}_{i_{2}}^{T}: i=\left(i_{1}, i_{2}\right) \in[d]^{2}, i_{1} \leq i_{2}\right\}$. For $f_{j}(\boldsymbol{A})=\left\langle\frac{1}{2}\left(\boldsymbol{e}_{j_{1}} \boldsymbol{e}_{j_{2}}^{T}+\boldsymbol{e}_{j_{2}} \boldsymbol{e}_{j_{1}}^{T}\right), \boldsymbol{A}\right\rangle$ with $j_{1} \leq j_{2}, f_{j}\left(\boldsymbol{A}_{i}\right)$ is only 1 when $i=j$ and 0 otherwise. Then each $\boldsymbol{A}_{i}$ is 1-independent to the predecessors on $f_{i}$ and zero, and thus the eluder dimension is lower bounded by $\binom{d}{2}$.

Then we are all set for the results in the first line of 1 Notice that when we choose $\alpha=O\left(1 / T^{2}\right)$ and $\epsilon=O\left(1 / T^{2}\right)$ in our analysis of Algorithm 2, the regret upper bound would only expand by $\log (T)$ factors.

## C Omitted Proofs for Quadratic Reward

In this section we include all the omitted proof of the theorems presented in the main paper.

## C. 1 Omitted Proofs of Main Results for Stochastic Bandit Eigenvector Problem

Proof of Theorem 3.3. Notice in Algorithm 1. for each iterate $\boldsymbol{a}$, its next iterate $\boldsymbol{y}$ satisfies

$$
\begin{aligned}
\boldsymbol{y} & =\frac{1}{n_{s}} \sum_{i=1}^{n_{s}}\left(\boldsymbol{a} / 2+\boldsymbol{z}_{i} / 2\right)^{\top} \boldsymbol{M}\left(\boldsymbol{a} / 2+\boldsymbol{z}_{i} / 2\right) \boldsymbol{z}_{i}+\eta_{i} \boldsymbol{z}_{i} \\
& =\frac{m_{s}}{n_{s}} \sum_{i=1}^{n_{s}}\left(\frac{1}{4} \boldsymbol{a}^{\top} \boldsymbol{M} \boldsymbol{a}+\frac{1}{2} \boldsymbol{a}^{\top} \boldsymbol{M} \boldsymbol{z}_{i}+\eta_{i}\right) \boldsymbol{z}_{i} .
\end{aligned}
$$

Therefore $\mathbb{E}[\boldsymbol{y}]=\frac{1}{2} \boldsymbol{M a}$. We can write $2 \boldsymbol{y}=\boldsymbol{M a}+\boldsymbol{g}$ where $\boldsymbol{g}:=\frac{m_{s}}{n_{s}} \sum_{i=1}^{n_{s}}\left(\frac{1}{2} \boldsymbol{a}^{\top} \boldsymbol{M a}+2 \eta_{i}\right) \boldsymbol{z}_{i}$. With Claim D. 12 and Claim D. 11 we get that $\|\boldsymbol{g}\| \leq C \sqrt{\frac{m_{s} \log ^{2}(n / \delta) \log (d / \delta) d}{n_{s}}}$. Therefore with our choice of $n_{s} \geq \widetilde{\Theta}\left(\frac{d^{2}}{\varepsilon_{s}^{2}\left(\lambda_{1}-\left|\lambda_{2}\right|\right)^{2}}\right)$ we guarantee $\|\boldsymbol{g}\| \leq \varepsilon_{s}\left(\lambda_{1}-\left|\lambda_{2}\right|\right)$. Therefore it satisfies the requirements for noisy power method, and by applying Corollary C.4 we have with $L=$ $O(\kappa \log (d / \varepsilon))$ iterations we will be able to find $\left\|\widehat{\boldsymbol{a}}-\boldsymbol{a}^{*}\right\| \leq \varepsilon$. By setting $\delta<0.1 / L$ in the algorithm we can guarantee the whole process succeed with high probability. Altogether it is sufficient to take $L n_{s}=\widetilde{O}\left(\kappa d^{2} /(\varepsilon \Delta)^{2}\right)$ actions to get an $\varepsilon$-optimal arm.
Finally to get the cumulative regret bound, we apply Claim D.7 with $A=\frac{d^{2} \kappa}{\Delta^{2}}$ and $a=2$. Therefore we set $\varepsilon=A^{1 / 4} T^{-1 / 4}=\frac{d^{1 / 2} \kappa^{1 / 4}}{\Delta^{1 / 2} T^{1 / 4}}$ and get:

$$
\operatorname{Reg}(T) \lesssim T^{1 / 2} A^{1 / 2} r^{*}=\sqrt{\frac{d^{2} \kappa}{\Delta^{2}} T r^{*}}=\sqrt{d^{2} \kappa^{3} T}
$$

Corollary C. 1 (Formal statement for Corollary 3.6. In Algorithm 1. by setting $\alpha=1-\varepsilon^{2} / 2$, one can get $\varepsilon$-optimal reward with a total of $\widetilde{O}\left(d^{2} \lambda_{1}^{2} / \varepsilon^{4}\right)$ total samples to get a such that $r^{*}-f(\boldsymbol{a}) \leq \varepsilon$. Therefore one can get an accumulative regret of $\widetilde{O}\left(\lambda_{1}^{3 / 5} d^{2 / 5} T^{4 / 5}\right)$.

Proof of Lemma 3.6. In order to find an arm with $\lambda_{1} \varepsilon^{2}$-optimal reward, one will want to recover an arm that is $\varepsilon / 2$-close (meaning to find an $\boldsymbol{a}$ such that $\tan \theta\left(\boldsymbol{V}_{l}, \boldsymbol{a}\right) \leq \varepsilon / 2$ ) to the top eigenspace $\operatorname{span}\left(\boldsymbol{v}_{1}, \cdots \boldsymbol{v}_{l}\right)$, where $l$ satisfies $\lambda_{l} \geq \lambda_{1}-\widetilde{\varepsilon}$ and $\lambda_{l+1} \leq \lambda_{1}-\widetilde{\varepsilon}$. Here we set $\widetilde{\varepsilon}:=\lambda_{1} \varepsilon^{2} / 2$. We first show 1) this is sufficient to get an $\lambda_{1} \varepsilon$-optimal reward, and next show 2 ) how to set parameter to achieve this.

To get 1), we write $\boldsymbol{V}_{l}=\left[\boldsymbol{v}_{1}, \cdots \boldsymbol{v}_{l}\right] \in \mathbb{R}^{d \times l}$ and $\boldsymbol{V}_{l}^{\perp}=\left[\boldsymbol{v}_{l+1}, \cdots \boldsymbol{v}_{k}\right]$. When $\tan \theta\left(\boldsymbol{V}_{l}, \boldsymbol{a}_{T}\right)=$ $\left\|\boldsymbol{V}^{\perp} \boldsymbol{a}\right\| /\|\boldsymbol{V} \boldsymbol{a}\| \leq \varepsilon / 2$, from the proof of ClaimD.6, we get $r^{*}-f(\boldsymbol{a}) \leq \min \left\{\lambda_{1}, \lambda_{1} 2(\varepsilon / 2)^{2}+\widetilde{\varepsilon}\right\}=$ $\lambda_{1} \varepsilon^{2}$.
Now to get 2), we note that in each iteration we try to conduct the power iteration to find an action $\tan \theta\left(\boldsymbol{V}_{l}, \widehat{\boldsymbol{a}}\right) \leq \varepsilon / 2$ and with eigengap $\geq \widetilde{\varepsilon}:=\lambda_{1} \varepsilon^{2} / 2$. Therefore it is sufficient to let $\|\boldsymbol{g}\| \leq 0.1 \widetilde{\varepsilon} \varepsilon$ and $\left|\boldsymbol{v}_{1}^{\top} \boldsymbol{g}\right| \leq 0.1 \widetilde{\epsilon} \frac{1}{\sqrt{d}}$, and thus $n_{s} \geq \widetilde{\Theta}\left(\frac{d^{2}}{\varepsilon^{2} \widetilde{\varepsilon}^{2}}\right) \leq \widetilde{\Theta}\left(d^{2} / \lambda_{1}^{2} \varepsilon^{6}\right)$. Together we need $\lambda_{1} / \widetilde{\epsilon} \log (2 d / \varepsilon) n_{s}=\widetilde{\Theta}\left(d^{2} / \lambda_{1}^{2} \varepsilon^{8}\right)$ samples to get an $\lambda_{1} \varepsilon^{2}$-optimal reward. Namely we get $\widetilde{\varepsilon}$-optimal reward with $\widetilde{O}\left(d^{2} \lambda_{1}^{2} / \widetilde{\varepsilon}^{4}\right)$ samples.
Finally by applying Claim D.7 we get:

$$
\mathfrak{R}(T) \lesssim\left(d^{2} \lambda_{1}^{2}\right)^{\frac{1}{5}} T^{\frac{4}{5}} \lambda_{1}^{\frac{1}{5}} \leq \widetilde{O}\left(\lambda_{1}^{3 / 5} d^{2 / 5} T^{4 / 5}\right)
$$

Theorem C. 2 (Formal statement of Theorem 3.7. In Algorithm 3. if we set $n=\widetilde{\Theta}\left(\frac{d^{2} \lambda_{1}^{2}}{\varepsilon^{2} \lambda_{k}^{2}}\right), m=$ $d \log (n / \delta), L=\Theta(\log (d / \varepsilon)), \delta=0.1 / L$, we will be able to identify an action $\widehat{\boldsymbol{a}}$ that yield at most $\varepsilon$-regret with probability 0.9. Therefore by applying the standard PAC to regret conversion as discussed in Claim $D .7$ we get a cumulative regret of $\widetilde{O}\left(\lambda_{1}^{1 / 3} k^{1 / 3}(\widetilde{\kappa} d T)^{2 / 3}\right)$ for large enough $T$, where $\widetilde{\kappa}=\lambda_{1} /\left|\lambda_{k}\right|$.
On the other hand, we set $n=\widetilde{\Theta}\left(\frac{d^{2} k^{2}}{\varepsilon^{2}}\right)$ and keep the other parameters. If we play Algorithm $3 k$ times by setting $k^{\prime}=2,4,6, \cdots 2 k$ and select the best output among them, we can get a gap-free cumulative regret of $\widetilde{O}\left(\lambda_{1}^{1 / 3} k^{4 / 3}(d T)^{2 / 3}\right)$ for large enough $T$ with high probability.

Proof of Theorem 3.7. First we show the first setting identify an $\varepsilon$-optimal reward with $\widetilde{O}\left(\widetilde{\kappa}^{2} d^{2} k \epsilon^{-2}\right)$ samples.
Similarly as Theorem 3.8, when setting $n \geq \widetilde{\Theta}\left(d^{2} /\left(\sigma_{k}^{2} \widetilde{\epsilon}^{2}\right)\right)$, we can find $\boldsymbol{X}_{L}$ that satisfies $\left\|\left(\boldsymbol{X}_{L} \boldsymbol{X}_{L}^{\top}-I\right) \boldsymbol{U}\right\| \leq \widetilde{\epsilon}$, and therefore we recover an $\boldsymbol{Y}_{L}=\boldsymbol{M} \boldsymbol{X}_{L-1}+\boldsymbol{G}_{L}$ with $\left\|\boldsymbol{G}_{L}\right\| \leq \sigma_{k} \widetilde{\varepsilon}$ and

```
Algorithm 3 Gap-free Subspace Iteration for Bilinear Bandit
    Input: Quadratic reward \(f: \mathcal{X} \rightarrow \mathbb{R}\) generating noisy reward, failure probability \(\delta\), error \(\varepsilon\).
    Initialization: Set \(k^{\prime}=2 k\). Initial candidate matrix \(\boldsymbol{X}_{0} \in \mathbb{R}^{d \times k^{\prime}}, \boldsymbol{X}_{0}(j) \in \mathbb{R}^{d}, j=1,2, \cdots k^{\prime}\)
    is the \(j\)-th column of \(X_{0}\) and are i.i.d sampled on the unit sphere \(\mathbb{S}^{d-1}\) uniformly. Sample
    variance \(m\), \# sample per iteration \(n\), total iteration \(L\).
    for Iteration \(l\) from 1 to \(L\) do
        for \(s\) from 1 to \(k^{\prime}\) do
            Noisy subspace iteration:
            \(\overline{\text { Sample } \boldsymbol{z}_{i} \sim \mathcal{N}\left(0,1 / m I_{d}\right)}, i=1,2, \cdots n_{s}\).
            Calculate tentative rank-1 arms \(\widetilde{\boldsymbol{a}}_{i}=\frac{1}{2}\left(\boldsymbol{X}_{l-1}(s)+\boldsymbol{z}_{i}\right)\).
            Conduct estimation \(\boldsymbol{Y}_{l}(s) \leftarrow 4 m / n \sum_{i=1}^{n}\left(f\left(\widetilde{\boldsymbol{a}}_{i}\right)+\eta_{i}\right) \boldsymbol{z}_{i} .\left(\boldsymbol{Y}_{l} \in \mathbb{R}^{d \times k^{\prime}}\right)\)
        Let \(\boldsymbol{Y}_{l}=\boldsymbol{X}_{l} \boldsymbol{R}_{l}\) be a QR-factorization of \(\boldsymbol{Y}_{l}\)
        Update target arm \(\boldsymbol{a}_{l} \leftarrow \arg \max _{\|\boldsymbol{a}\|=1} \boldsymbol{a}^{\top} \boldsymbol{Y}_{l} \boldsymbol{X}_{l-1}^{\top} \boldsymbol{a}\).
    Output: \(\boldsymbol{a}_{L}\).
```

$\left\|\boldsymbol{Y}_{L} \boldsymbol{X}_{L-1}^{\top}-\boldsymbol{M}\right\|_{2}=\left\|\boldsymbol{M} \boldsymbol{X}_{L-1} \boldsymbol{X}_{L-1}^{\top}-\boldsymbol{M}+\boldsymbol{G}_{L} \boldsymbol{X}_{L-1}^{\top}\right\|_{2} \leq\left(\lambda_{1}+\left|\lambda_{k}\right|\right) \widetilde{\varepsilon}$. Therefore by definition of $\boldsymbol{a}_{L}, \boldsymbol{a}_{L}^{\top} \boldsymbol{Y}_{L} \boldsymbol{X}_{L-1}^{\top} \boldsymbol{a}_{L}=\max _{\|\boldsymbol{a}\|=1} \boldsymbol{a}^{\top}\left(\boldsymbol{M} \boldsymbol{X}_{L-1} \boldsymbol{X}_{L-1}^{\top}+\boldsymbol{G}_{L} \boldsymbol{X}_{L-1}^{\top}\right) \boldsymbol{a} \geq \lambda_{1}-\left(\lambda_{1}+\left|\lambda_{k}\right|\right) \widetilde{\text {. }}$. Therefore $\boldsymbol{a}_{L}^{\top} \boldsymbol{M} \boldsymbol{a}_{L} \geq \lambda_{1}-2\left(\lambda_{1}+\left|\lambda_{k}\right|\right) \widetilde{\varepsilon}$. Therefore we set $2\left(\lambda_{1}+\left|\lambda_{k}\right|\right) \widetilde{\varepsilon}=\epsilon$, i.e., $\widetilde{\varepsilon}=0.5 \epsilon /\left(\lambda_{1}+\left|\lambda_{k}\right|\right)$ which will get a total sample of $T=\widetilde{\Theta}(k n)=\widetilde{\Theta}\left(d^{2} \widetilde{\kappa}^{2} k \varepsilon^{-2}\right)$. Then by applying Claim D. 7 we get the cumulative regret bound.

Next we show how to estimate the action with $\widetilde{O}\left(d^{2} k^{4} \varepsilon^{-2}\right)$ samples. To achieve this result, we need to slightly alter Algorithm 3 where we respectively set $k^{\prime}=2,4,6, \cdots 2 k$ and keep the best arm among the $k$ outputs. We argue that among all the choices of $k^{\prime}$, at least for one $l \in[k], k^{\prime}=2 l$, we have $\left|\lambda_{l}\right|-\left|\lambda_{l+1}\right| \geq \lambda_{1} / k$. Notice with similar argument as above, when we set $n=\widetilde{\Theta}\left(d^{2} \lambda_{l}^{-2} \widetilde{\varepsilon}^{-2}\right) \leq$ $\widetilde{\Theta}\left(d^{2} k^{2} \lambda_{1}^{-2} \widetilde{\varepsilon^{-2}}\right)$ we can get $\|\boldsymbol{G}\| \leq \widetilde{\varepsilon} \lambda_{l}$ as required by Corollary C.4 the total number of iterations $L=O\left(\sigma_{l} /\left(\sigma_{l}-\sigma_{l+1}\right) \log (2 d / \epsilon)=\widetilde{O}(k)\right.$. Finally by setting $\widetilde{\varepsilon}=\epsilon /\left(4 \lambda_{1}\right)$ we get the overall samples we required is $\widetilde{O}\left(k^{2} n\right)=\widetilde{O}\left(d^{2} k^{4} \epsilon^{-2}\right)$.
For both settings, directly applying our arguments in the PAC to regret conversion: Claim D.7 will finish the proof.

## C. 2 Omitted Details of Main Results of Low-Rank Linear Reward

```
Algorithm 4 Subspace Iteration Exploration for Low-rank Linear Reward.
    Input: Quadratic function \(f: \mathcal{A} \rightarrow \mathbb{R}\) with noisy reward, failure probability \(\delta\), error \(\varepsilon\).
    Initialization: Set \(k^{\prime}=2 k\). Initial candidate matrix \(\boldsymbol{X}_{0} \in \mathbb{R}^{d \times k^{\prime}}, \boldsymbol{X}_{0}(j) \in \mathbb{R}^{d}, j=1,2, \cdots k^{\prime}\)
    is the \(j\)-th column of \(\boldsymbol{X}_{0}\) and are i.i.d sampled on the unit sphere \(\mathbb{S}^{d-1}\) uniformly. Sample
    variance \(m\), \# sample per iteration \(n\), total iteration \(L\).
    for Iteration \(l\) from 1 to \(L\) do
        Sample \(\boldsymbol{z}_{i} \sim \mathcal{N}\left(0,1 / m I_{d}\right), i=1,2, \cdots n\).
        for \(s\) from 1 to \(k^{\prime}\) do
            Noisy subspace iteration:
            Calculate tentative rank-1 actions \(\widetilde{\boldsymbol{A}}_{i}=\boldsymbol{X}_{l-1}(s) \boldsymbol{z}_{i}^{\top}\).
            Conduct estimation \(\boldsymbol{Y}_{l}(s) \leftarrow m / n \sum_{i=1}^{n}\left(\left\langle\boldsymbol{M}, \widetilde{\boldsymbol{A}}_{i}\right\rangle+\eta_{i, s}\right) \boldsymbol{z}_{i} .\left(\boldsymbol{Y}_{l} \in \mathbb{R}^{d \times k^{\prime}}\right)\)
        Let \(\boldsymbol{Y}_{l}=\boldsymbol{X}_{l} \boldsymbol{R}_{l}\) be a QR-factorization of \(\boldsymbol{Y}_{l}\)
        Update target action \(\boldsymbol{A}_{l} \leftarrow \boldsymbol{Y}_{l} \boldsymbol{X}_{l}^{\top}\).
    Output: \(\widehat{\boldsymbol{A}}=\boldsymbol{A}_{L} /\left\|\boldsymbol{A}_{L}\right\|_{F}\)
```

Theorem C. 3 (Formal statement of Theorem 3.8). In Algorithm 4 for large enough constants $C_{n}, C_{L}, C_{m}$, let $n=C_{n} d^{2} \log ^{2}(d / \delta) \sigma_{k}^{-2} \varepsilon^{-2}, m=C_{m} d \log (n / \delta)$, and $L=C_{L} \log (d / \varepsilon), \boldsymbol{X}_{L}$
satisfies $\left\|\left(\boldsymbol{I}-\boldsymbol{X}_{L} \boldsymbol{X}_{L}^{\top}\right) \boldsymbol{V}\right\| \leq \varepsilon / 4$, and the output $\widehat{\boldsymbol{A}}$ satisfies $\left\|\widehat{\boldsymbol{A}}-\boldsymbol{A}^{*}\right\|_{F} \leq\left\|\boldsymbol{M}_{\sim}\right\|_{F} \varepsilon$. Altogether to get an $\varepsilon$-optimal action, it is sufficient to have total sample complexity of $T \leq \widetilde{O}\left(d^{2} k \lambda_{k}^{-2} \varepsilon^{-2}\right)$.

Proof of Theorem 3.8. Let $\boldsymbol{M}=\boldsymbol{V} \boldsymbol{\Sigma} \boldsymbol{V}^{\top}$. From Claim C.6 we get that for each noisy subspace iteration step we get $\boldsymbol{Y}_{l}=\boldsymbol{M} \boldsymbol{X}_{l}+\boldsymbol{G}_{l}$ with $5\left\|\boldsymbol{G}_{l}\right\| \leq \varepsilon \sigma_{k}$ and $\left\|\boldsymbol{V}^{\top} \boldsymbol{G}\right\| \leq \sigma_{k} \sqrt{k} / 3 \sqrt{d} \leq$ $\sigma_{k}(\sqrt{2 k}-\sqrt{k}) / 2 \sqrt{d}$. Therefore we can apply Corollary C.4. and get $\left\|\boldsymbol{V}\left(\boldsymbol{X}_{L} \boldsymbol{X}_{L}^{\top}-\boldsymbol{I}\right)\right\| \leq \varepsilon / 4$ with $O(\log 2 d / \epsilon)$ steps. Therefore we have:

$$
\begin{aligned}
\left\|\boldsymbol{A}_{L}-\boldsymbol{M}\right\|_{F} & \left.=\left\|\left(\boldsymbol{M} \boldsymbol{X}_{L}+\boldsymbol{G}_{L}\right) \boldsymbol{X}_{L}^{\top}-\boldsymbol{M}\right\|_{F}=\| \boldsymbol{V}^{\top} \boldsymbol{\Sigma} \boldsymbol{V}\left(\boldsymbol{X}_{L} \boldsymbol{X}_{L}^{\top}-\boldsymbol{I}\right)\right)+\boldsymbol{G}_{L} \boldsymbol{X}_{L}^{\top} \|_{F} \\
& \leq\|\boldsymbol{M}\|_{F}\left\|\boldsymbol{V}\left(\boldsymbol{X}_{L} \boldsymbol{X}_{L}^{\top}-\boldsymbol{I}\right)\right\|+\left\|\boldsymbol{G}_{L}\right\|\left\|\boldsymbol{X}_{L}\right\|_{F} \\
& \leq\left(\|\boldsymbol{M}\|_{F}+\sigma_{k}\right) \varepsilon / 4<\|\boldsymbol{M}\|_{F} \varepsilon / 2
\end{aligned}
$$

Meanwhile, notice $\left\|\boldsymbol{A}^{*}\right\|_{F}=1,\|\boldsymbol{M}\|_{F}=r^{*}$ and $\|\widehat{\boldsymbol{A}}\|_{F}=1$. $\left\|\boldsymbol{A}_{L} / r^{*}-\boldsymbol{A}^{*}\right\|_{F} \leq \varepsilon / 2$. $\| \widehat{\boldsymbol{A}}-$ $\boldsymbol{A}^{*}\left\|_{F}=\right\| \boldsymbol{A}_{L} /\left\|\boldsymbol{A}_{L}\right\|_{F}-\boldsymbol{A}^{*}\left\|_{F}=\right\| \operatorname{vec}\left(\boldsymbol{A}_{L}\right) /\left\|\operatorname{vec}\left(\boldsymbol{A}_{L}\right)\right\|_{2}-\operatorname{vec}\left(\boldsymbol{A}^{*}\right) \|_{2}$.

Write $\theta_{A}:=\theta\left(\operatorname{vec}\left(\boldsymbol{A}_{L}\right), \operatorname{vec}\left(\boldsymbol{A}^{*}\right)\right.$. The worst case that makes $\left\|\operatorname{vec}(\widehat{\boldsymbol{A}})-\operatorname{vec}\left(\boldsymbol{A}^{*}\right)\right\|$ to be larger than $\left\|\operatorname{vec}\left(\boldsymbol{A}_{L} / r^{*}\right)-\operatorname{vec}\left(\boldsymbol{A}^{*}\right)\right\|$ is when $\left\|\operatorname{vec}\left(\boldsymbol{A}_{L} / r^{*}\right)-\operatorname{vec}\left(\boldsymbol{A}^{*}\right)\right\|=\sin \theta_{A}$ and $\left\|\operatorname{vec}(\widehat{\boldsymbol{A}})-\operatorname{vec}\left(\boldsymbol{A}^{*}\right)\right\|$ is always $2 \sin \left(\theta_{A} / 2\right)$. Notice trivially $2 \sin \left(\theta_{A} / 2\right) \leq 2 \sin \left(\theta_{A}\right)$ Therefore we could get $\left\|\widehat{\boldsymbol{A}}-\boldsymbol{A}^{*}\right\|_{F} \leq$ $2\left\|\boldsymbol{A}_{L} / r^{*}-\boldsymbol{A}^{*}\right\|_{F} \leq \varepsilon$.

Proof of Corollary 3.9. The corollary uses a special property of the strongly convex action set that ensures: $\boldsymbol{A}^{*}=\boldsymbol{M} / r^{*}$. With $\widehat{\boldsymbol{A}}$ that satisfies $\|\widehat{\boldsymbol{A}}\|_{F}=1$, we have

$$
\begin{align*}
r^{*}-f_{\boldsymbol{M}}(\boldsymbol{A}) & =r^{*}-\langle\widehat{\boldsymbol{A}}, \boldsymbol{M}\rangle=r^{*}-\left\langle\widehat{\boldsymbol{A}}, r^{*} \boldsymbol{A}^{*}\right\rangle \\
& =\frac{r^{*}}{2}\left(2-2\left\langle\widehat{\boldsymbol{A}}, \boldsymbol{A}^{*}\right\rangle\right)=\frac{r^{*}}{2}\left(\|\widehat{\boldsymbol{A}}\|_{F}^{2}+\left\|\boldsymbol{A}^{*}\right\|_{F}^{2}-\left\langle\widehat{\boldsymbol{A}}, \boldsymbol{A}^{*}\right\rangle\right) \\
& =\frac{r^{*}}{2}\left\|\widehat{\boldsymbol{A}}-\boldsymbol{A}^{*}\right\|_{F}^{2} \leq \frac{r^{*} \varepsilon^{2}}{2} \tag{2}
\end{align*}
$$

Therefore, with first $T_{1}=\widetilde{O}\left(d^{2} k \lambda_{k}^{-2} \varepsilon^{-2}\right)$ exploratory samples we get $r^{*}-f(\widehat{\boldsymbol{A}}) \leq r^{*} \varepsilon^{2} / 2=$ $r^{*} \sqrt{\frac{d^{2} k}{\lambda_{k}^{2} T}}=\sqrt{\frac{\left(r^{*}\right)^{2} d^{2} k}{\lambda_{k}^{2} T}}$. Together we have:

$$
\begin{aligned}
\mathfrak{R}(T) & =\sum_{t=1}^{T_{1}} r^{*}-f\left(\boldsymbol{A}_{t}\right)+\sum_{t=T_{1}+1}^{T} r^{*}-f(\widehat{\boldsymbol{A}}) \\
& <r^{*} T_{1}+T r^{*} \varepsilon^{2} \\
& \leq \widetilde{O}\left(\sqrt{d^{2} k\left(r^{*}\right)^{2} \lambda_{k}^{-2} T}\right) .
\end{aligned}
$$

Proof of Theorem 3.10 We find an $l$ to be the smallest integer such that $\sum_{i=l+1}^{k} \sigma_{i}^{2} \leq \epsilon^{2}\|\boldsymbol{M}\|_{F}^{2}$. Then we have $\sigma_{l} \geq \epsilon / \sqrt{k-l}>\epsilon / \sqrt{k}$.
Notice that in Algorithm 4, we set $n \geq \widetilde{\Theta}\left(\frac{d^{2} k}{\left(r^{*}\right)^{2} \varepsilon^{4}}\right)$ large enough such that $\|\boldsymbol{G}\|_{2} \leq$ $O\left(\|\boldsymbol{M}\|_{F} \epsilon^{2} / \sqrt{k}\right) \lesssim \epsilon\left(\sigma_{l}-0\right)$ and $\left\|\boldsymbol{U}^{\top} \boldsymbol{G}\right\|_{2} \leq\|\boldsymbol{M}\|_{F} \epsilon / \sqrt{k} \frac{\sqrt{k^{\prime}}-\sqrt{k-1}}{2 \sqrt{d}}$. (This comes from the argument proved in Claim C.6.)
Therefore by conducting noisy power method we get with $O(n k)=\widetilde{O}\left(\frac{d^{2} k^{2}}{\left(r^{*}\right)^{2} \varepsilon^{4}}\right)$ samples we can get an action $\widehat{\boldsymbol{A}}$ that satisfies:

$$
\left\|\boldsymbol{M}-\boldsymbol{X}_{L} \boldsymbol{X}_{L}^{\top} \boldsymbol{M}\right\|_{F}^{2} \leq \sum_{i=l+1}^{k} \sigma_{i}^{2}+l \epsilon^{2} \sigma_{l}^{2} \leq 2\|\boldsymbol{M}\|_{F}^{2} \epsilon^{2}
$$

Therefore we could get $\left\|\boldsymbol{A}^{*}-\widehat{\boldsymbol{A}}\right\| \leq 2 \epsilon$, and with similar argument as (2) we have $r^{*}-f(\widehat{\boldsymbol{A}}) \leq$ $\|\boldsymbol{M}\|_{F} \epsilon^{2}$.
Therefore if we want to take a total of $T$ actions, we will set $\epsilon^{6}=\widetilde{\Theta}\left(\frac{d^{2} k^{2}}{\left(r^{*}\right)^{2} T}\right)$ and we get:

$$
\begin{aligned}
\mathfrak{R}(T) & =\sum_{t=1}^{T_{1}} r^{*}-f\left(\boldsymbol{A}_{t}\right)+\sum_{t=T_{1}+1}^{T} r^{*}-f(\widehat{\boldsymbol{A}}) \\
& <r^{*} T_{1}+T r^{*} \varepsilon^{2} \\
& \leq \widetilde{O}\left(d^{2 / 3} k^{2 / 3}\left(r^{*}\right)^{1 / 3} T^{2 / 3}\right) .
\end{aligned}
$$

## C. 3 Technical Details for Quadratic Reward

## Noisy Power Method.

Corollary C. 4 (Adapted from Corollary 1.1 from [37]). Let $k^{\prime} \geq l$ l. Let $\boldsymbol{U} \in \mathbb{R}^{d \times l}$ represent the top $l$ singular vectors of $M$ and let $\sigma_{1} \geq \cdots \geq \sigma_{k}>0$ denote its singular values. Suppose $X_{0}$ is an orthonormal basis of a random $k^{\prime}$-dimensional subspace. Further suppose that at every step of NPM we have

$$
\begin{aligned}
5\|\boldsymbol{G}\| & \leq \epsilon\left(\sigma_{l}-\sigma_{l+1}\right), \\
\text { and } 5\left\|\boldsymbol{U}^{\top} \boldsymbol{G}\right\| & \leq\left(\sigma_{l}-\sigma_{l+1}\right) \frac{\sqrt{k^{\prime}}-\sqrt{l-1}}{2 \sqrt{d}}
\end{aligned}
$$

for some fixed parameter $\epsilon<1 / 2$. Then with all but $2^{-\Omega\left(k^{\prime}+1-l\right)}+e^{\Omega(d)}$ probability, there exists an $L=O\left(\frac{\sigma_{l}}{\sigma_{l}-\sigma_{l+1}} \log (2 d / \epsilon)\right)$ so that after $L$ steps we have that $\left\|\left(\boldsymbol{I}-\boldsymbol{X}_{L} \boldsymbol{X}_{L}^{\top}\right) \boldsymbol{U}\right\| \leq \epsilon$.
Theorem C. 5 (Adapted from Theorem 2.2 from [11]). Let $\boldsymbol{U}_{l} \in \mathbb{R}^{d \times l}$ represent the top $l$ singular vectors of $M$ and let $\sigma_{1} \geq \cdots \geq \sigma_{k}>0$ denote its singular values. Naturally $l \leq k$. Suppose $X_{0}$ is an orthonormal basis of a random $k^{\prime}$-dimensional subspace where $k^{\prime} \geq k$. Further suppose that at every step of NPM we have

$$
\begin{aligned}
\|\boldsymbol{G}\| & \leq O\left(\epsilon \sigma_{l}\right) \\
\text { and }\left\|\boldsymbol{U}_{k}^{\top} \boldsymbol{G}\right\|_{2} & \leq O\left(\sigma_{l} \frac{\sqrt{k^{\prime}}-\sqrt{k-1}}{2 \sqrt{d}}\right)
\end{aligned}
$$

for small enough $\epsilon$. Then with all but $2^{-\Omega\left(k^{\prime}+1-k\right)}+e^{\Omega(d)}$ probability, there exists an $L=$ $O(\log (2 d / \epsilon))$ so that after $L$ steps we have that $\left\|\left(\boldsymbol{I}-\boldsymbol{X}_{L} \boldsymbol{X}_{L}^{\top}\right) \boldsymbol{U}_{l}\right\| \leq \epsilon$. Furthermore:

$$
\left\|\boldsymbol{M}-\boldsymbol{X}_{L} \boldsymbol{X}_{L}^{\top} \boldsymbol{M}\right\|_{F}^{2} \leq \sum_{i=l+1}^{k} \sigma_{i}^{2}+l \sigma^{2} \sigma_{l}^{2} .
$$

## Concentration Bounds.

Claim C.6. Write the eigendecomposition for $\boldsymbol{M}$ as $\boldsymbol{M}=\boldsymbol{U} \boldsymbol{\Sigma} \boldsymbol{U}^{\top}$. In Algorithm 4 when $n \geq$ $\widetilde{\Theta}\left(d^{2} /\left(\lambda_{k}^{2} \varepsilon^{2}\right)\right)$, the noisy subspace iteration step can be written as: $\boldsymbol{Y}_{l}=\boldsymbol{M} \boldsymbol{X}_{l-1}+\boldsymbol{G}_{l}$, where the noise term satisfies:

$$
\begin{aligned}
5\left\|\boldsymbol{G}_{l}\right\| & \leq \varepsilon\left|\lambda_{k}\right| \\
5\left\|\boldsymbol{U}^{\top} \boldsymbol{G}_{l}\right\| & \leq \varepsilon\left|\lambda_{k}\right| \frac{\sqrt{k}}{3 \sqrt{d}} .
\end{aligned}
$$

with high probability for our choice of $n$.
Proof. For compact notation, write vector $\boldsymbol{\eta}_{i}:=\left[\eta_{i, 1}, \eta_{i, 2}, \cdots \eta_{i, k^{\prime}}\right]^{\top} \in \mathbb{R}^{k^{\prime}}$. We have:

$$
\begin{aligned}
\boldsymbol{G}_{\boldsymbol{l}}(s) & =\frac{m}{n} \sum_{i=1}^{n}\left(\boldsymbol{z}_{i}^{\top} \boldsymbol{M} \boldsymbol{X}_{l}(s)\right) \boldsymbol{z}_{i}+\frac{m}{n} \sum_{i=1}^{n} \eta_{i, s} \boldsymbol{z}_{i}-\boldsymbol{M} \boldsymbol{X}_{l}(s), \text { therefore } \\
\boldsymbol{G}_{l} & =\left(\frac{m}{n} \sum_{i=1}^{n}\left[\boldsymbol{z}_{i} \boldsymbol{z}_{i}^{\top}\right]-I\right) \boldsymbol{M} \boldsymbol{X}_{l}+\frac{m}{n} \sum_{i=1}^{n} \boldsymbol{z}_{i} \boldsymbol{\eta}_{i}^{\top} .
\end{aligned}
$$

First note that for orthogonal matrix $\boldsymbol{X}_{l},\left\|\boldsymbol{M} \boldsymbol{X}_{l}\right\| \leq \lambda_{1}$, and $\left\|\frac{m}{n} \sum_{i=1}^{n}\left[\boldsymbol{z}_{i} \boldsymbol{z}_{i}^{\top}\right]-I\right\| \leq$ $O\left(\sqrt{\frac{d+\log (1 / \delta)}{n}}\right)$. The bottleneck is from the second term and we will use Matrix Bernstein to concentrate it. Write $\boldsymbol{S}_{i}=\frac{m}{n} \boldsymbol{z}_{i} \boldsymbol{\eta}_{i}^{\top}$. We have $\left\|\boldsymbol{S}_{i}\right\| \leq O\left(\frac{\sqrt{m k^{\prime}} \log (n / \delta)}{n}\right)$ with probability $1-\delta$ and $\mathbb{E}\left[\sum_{i} \boldsymbol{S}_{i} \boldsymbol{S}_{i}^{\top}\right]=\frac{m k^{\prime}}{n} I_{d}$ and $\mathbb{E}\left[\sum_{i} \boldsymbol{S}_{i}^{\top} \boldsymbol{S}_{i}\right]=\frac{m d}{n} I_{k^{\prime}}$. Therefore with matrix Bernstein we can get that $\left\|\sum_{i} \boldsymbol{S}_{i}\right\|_{i} \leq O\left(\sqrt{\frac{m d}{n}} \log (d / \delta)\right)$ with probability $1-\delta$.
Therefore for $n \geq \widetilde{\Omega}\left(d^{2} /\left(\lambda_{k}^{2} \varepsilon^{2}\right)\right.$, we can get that $5\left\|\boldsymbol{G}_{l}\right\| \leq \varepsilon\left|\lambda_{k}\right|$.
Similarly since $\boldsymbol{U}^{\top} \boldsymbol{z}_{i} \sim \mathcal{N}\left(0, \frac{1}{m} I_{k^{\prime}}\right)$, with the same argument one can easily get that $\left\|\boldsymbol{U}^{\top} \boldsymbol{G}_{l}\right\| \leq$ $O\left(\sqrt{\frac{m k^{\prime}}{n}} \log (d / \delta)\right)$. Therefore with the same lower bound for $n$ one can get $15\left\|\boldsymbol{U}^{\top} \boldsymbol{G}_{l}\right\| \leq$ $\varepsilon\left|\lambda_{k}\right| \sqrt{\frac{k}{d}}$.

## C. 4 Omitted Proof for RL with Quadratic Q function

```
Algorithm 5 Learn policy complete polynomial with simulator.
    Initialize: Set \(n=\widetilde{\Theta}\left(\widetilde{\kappa}^{2} d^{2} H^{3} / \varepsilon^{2}\right)\), Oracle to estimate \(\widehat{T}_{h}\) from noisy observations.
    for \(h=H, \ldots 1\) do
        Sample \(\phi\left(s_{h}^{i}, a_{h}^{i}\right), i \in[n]\) from standard Gaussian \(N\left(0, I_{d}\right)\)
        for \(i \in[n]\) do
            Query \(\left(s_{h}^{i}, a_{h}^{i}\right)\) and use \(\pi_{h+1}, \ldots, \pi_{H}\) as the roll-out to get estimation
    \(\widehat{Q}_{h}^{\pi_{h+1}, \ldots, \pi_{H}}\left(s_{h}^{i}, a_{h}^{i}\right)\)
        Retrieve \(\widehat{\boldsymbol{M}}_{h}\) from estimation \(\widehat{Q}_{h}^{\pi_{h+1}, \ldots, \pi_{H}}\left(s_{h}^{i}, a_{h}^{i}\right), i \in[n]\)
        Set \(\widehat{Q}_{h}(s, a) \leftarrow f_{\widehat{T}_{h}}\)
        Set \(\pi_{h}(s) \leftarrow \arg \max _{a \in \mathcal{S}} \widehat{Q}_{h}(s, a)\)
    Return \(\pi_{1}, \ldots, \pi_{H}\)
```

Proof of Theorem 3.13. With the oracle, at horizon $H$, we can estimate $\widehat{\boldsymbol{M}}_{H}$ that is $\epsilon / H$ close to $M_{H}^{*}$ in spectral norm through noisy observations from the reward function with $\widetilde{O}\left(\widetilde{\kappa}^{2} d^{2} H^{2} / \varepsilon^{2}\right)$ samples. Next, for each horizon $h=H-1, H-1, \cdots, 1$, sample $s_{i}^{\prime} \sim \mathbb{P}(\cdot \mid s, a)$, we define $\eta_{i}=$ $\max _{a^{\prime}} f_{\widehat{M}_{h+1}}\left(s_{i}^{\prime}, a^{\prime}\right)-\mathbb{E}_{s^{\prime} \sim \mathbb{P}(\cdot \mid s, a)} \max _{a^{\prime}} f_{\widehat{M}_{h+1}}\left(s^{\prime}, a^{\prime}\right) . \eta_{i}$ is mean-zero and $O(1)$-sub-gaussian since it is bounded. Denote $\boldsymbol{M}_{h}$ as the matrix that satisfies $f_{M_{h}}:=\mathcal{T} f_{\widehat{\boldsymbol{M}}_{h+1}}$, which is well-defined due to Bellman completeness. We estimate $\widehat{\boldsymbol{M}}_{h}$ from the noisy observations $y_{i}=r_{h}(s, a)+$ $\max _{a^{\prime}} f_{\widehat{\boldsymbol{M}}_{h+1}}\left(s_{i}^{\prime}, a^{\prime}\right)=\mathcal{T} f_{\widehat{\boldsymbol{M}}_{h+1}}+\eta_{i}=: f_{\boldsymbol{M}_{h}}+\eta_{i}$. Therefore with the oracle, we can estimate $\widehat{\boldsymbol{M}}_{h}$ such that $\left\|\widehat{\boldsymbol{M}}_{h}-\boldsymbol{M}_{h}\right\|_{2} \leq \epsilon / H$ with $\Theta\left(\widetilde{\kappa}^{2} d^{2} k^{2} H^{2} / \epsilon^{2}\right)$ bandits. Together we have:

$$
\begin{aligned}
\left\|f_{\widehat{\boldsymbol{M}}_{h}}-f_{\widehat{\boldsymbol{M}}_{h}^{*}}\right\|_{\infty} & =\left\|\widehat{\boldsymbol{M}}_{h}-\boldsymbol{M}_{h}^{*}\right\| \\
& \leq\left\|\widehat{\boldsymbol{M}}_{h}-\boldsymbol{M}_{h}\right\|+\left\|\boldsymbol{M}_{h}-\boldsymbol{M}_{h}^{*}\right\| \\
& \leq \epsilon / H+\left\|\mathcal{T} f_{\widehat{\boldsymbol{M}}_{h+1}}-\mathcal{T} f_{\boldsymbol{M}_{h+1}^{*}}\right\|_{\infty} \\
& \leq \epsilon / H+\left\|f_{\widehat{\boldsymbol{M}}_{h+1}}-f_{\boldsymbol{M}_{h+1}^{*}}\right\|_{\infty} \\
& \leq 2 \epsilon / H+\left\|\widehat{\widehat{\boldsymbol{M}}}_{h+2}-f_{\boldsymbol{M}_{h+2}^{*}}\right\|_{\infty} \\
& \leq \cdots \\
& \leq(H-h) \epsilon / H .
\end{aligned}
$$

Finally for $h=1$ we have $\left\|\widehat{\boldsymbol{M}}_{1}-\boldsymbol{M}^{*}\right\| \leq \epsilon$ if we sample $n=\widetilde{\Theta}\left(\widetilde{\kappa}^{2} d^{2} k^{2} H^{2} / \epsilon^{2}\right)$ for each $h \in[H]$. Therefore for all the $H$ timesteps we need $\Theta\left(\widetilde{\kappa}^{2} d^{2} k^{2} H^{3} / \epsilon^{2}\right)$.

## D Technical details for General Tensor Reward

```
Algorithm 6 Phased elimination with zeroth order exploration.
    Input: Function \(f: \mathcal{A} \rightarrow \mathbb{R}\) of polynomial degree \(p\) generating noisy reward, failure probability
    \(\overline{\delta \text {, error }} \varepsilon\).
    Initialization: \(L_{0}=C_{L} k \log (1 / \delta)\); Total number of stages \(S=C_{S}\lceil\log (1 / \varepsilon)\rceil+1, \mathcal{A}_{0}=\)
    \(\left\{\boldsymbol{a}_{0}^{(1)}, \boldsymbol{a}_{0}^{(2)}, \cdots \boldsymbol{a}_{0}^{\left(L_{0}\right)}\right\}\) where each \(\boldsymbol{a}_{0}^{(l)}\) is uniformly sampled on the unit sphere \(\mathbb{S}^{d-1} . \widetilde{\varepsilon}_{0}=1\).
    for \(s\) from 1 to \(S\) do
        \(\left.\widetilde{\varepsilon}_{s} \leftarrow \widetilde{\varepsilon}_{s-1} / 2, n_{s} \leftarrow C_{n} d^{p} \log (d / \delta) / \lambda_{1}^{2} \widetilde{\varepsilon}_{s}^{2}\right), n_{s} \leftarrow n_{s} \cdot \log ^{3}\left(n_{s} / \delta\right), m_{s} \leftarrow C_{m} d \log \left(n_{s} / \delta\right)\),
    \(\mathcal{A}_{s}=\varnothing\).
        for \(l\) from 1 to \(L_{s-1}\) do
            Zeroth-order optimization:
            Locate current action \(\widetilde{\boldsymbol{a}}=\boldsymbol{a}_{s-1}^{(l)}\).
            for \(\lceil(1 /(1-\alpha)) \log (2 d)\rceil\) times do
                Sample \(\boldsymbol{z}_{i} \sim \mathcal{N}\left(0,1 / m_{s} I_{d}\right), i=1,2, \cdots n_{s}\).
                    Take actions \(\boldsymbol{a}_{i}=\left(1-\frac{1}{2 p}\right) \widetilde{\boldsymbol{a}}+\frac{1}{2 p} \boldsymbol{z}_{i}\) and observe \(r_{i}=\boldsymbol{T}\left(\boldsymbol{a}_{i}\right)+\eta_{i}, i \in\left[n_{s}\right]\); Take
    actions \(\frac{1}{2 p} \boldsymbol{z}_{i}\) and observe \(r_{i}^{\prime}=\boldsymbol{T}\left(\frac{1}{2 p} \boldsymbol{z}_{i}\right)+\eta_{i}^{\prime}, i \in\left[n_{s}\right]\).
                    Conduct estimation \(\boldsymbol{y} \leftarrow 1 / n_{s} \sum_{i=1}^{n_{s}}\left(r_{i}-r_{i}^{\prime}\right) \boldsymbol{z}_{i}\).
                    Update the current action \(\widetilde{\boldsymbol{a}} \leftarrow \boldsymbol{y} /\|\boldsymbol{y}\|\).
            Estimate the expected reward for \(\widetilde{\boldsymbol{a}}\) through \(n_{s}\) samples: \(r_{n}(\widetilde{\boldsymbol{a}})=1 / n_{s} \sum_{i=1}^{n_{s}}\left(\boldsymbol{T}(\widetilde{\boldsymbol{a}})+\eta_{i}\right)\).
            Candidate Elimination:
            if \(r_{n} \geq \lambda_{1}\left(1-p \widetilde{\varepsilon}_{s}^{2}\right)\) then
                    Keep the action \(\mathcal{A}_{s} \leftarrow \mathcal{A}_{s} \cup\{\widetilde{\boldsymbol{a}}\}\)
            Label the actions: \(L_{s}=\left|\mathcal{A}_{s}\right|, \mathcal{A}_{s}=:\left\{\boldsymbol{a}_{s}^{(1)}, \cdots \boldsymbol{a}_{s}^{\left(L_{s}\right)}\right\}\).
    Run UCB (Algorithm7) with the candidate set \(\mathcal{A}_{S}\).
```


## D. 1 Technical Details for Symmetric Setting

Lemma D. 1 (Zeroth order optimization for noiseless setting). For $p \geq 3$, suppose $0.5 \boldsymbol{a}^{\top} \boldsymbol{v}_{1}>\left|\boldsymbol{a}^{\top} \boldsymbol{v}_{j}\right|$ for all $j \geq 2$, we have:

$$
\tan \theta\left(G(\boldsymbol{a}), \boldsymbol{v}_{1}\right) \leq \frac{1}{2} \tan \theta\left(\boldsymbol{a}, \boldsymbol{v}_{1}\right)
$$

Proof. We first simplify $G(\boldsymbol{a})=\sum_{j=1}^{r} \lambda_{j} \boldsymbol{v}_{j} \cdot S_{j}$, where

$$
\begin{aligned}
G(\boldsymbol{a}) & =\sum_{s=0}^{\lfloor(p-3) / 2\rfloor} \frac{\left(1-\frac{1}{2 p}\right)^{p-2 s-1}\left(\frac{1}{2 p}\right)^{2 s+1}}{m^{s}}\binom{p}{2 s+1} T\left(I^{\otimes s+1} \otimes \boldsymbol{a}^{\otimes p-2 s-1}\right) \\
& =\sum_{s=0}^{\lfloor(p-3) / 2\rfloor} \frac{\left(1-\frac{1}{2 p}\right)^{p-2 s-1}\left(\frac{1}{2 p}\right)^{2 s+1}}{m^{s}}\binom{p}{2 s+1} \sum_{j=1}^{k} \lambda_{j}\left(\boldsymbol{v}_{j}^{\top} \boldsymbol{a}\right)^{p-2 j-1} \boldsymbol{v}_{j} \\
& =\sum_{j=1}^{k} v_{j} \cdot \overbrace{\lambda_{j}}^{\sum_{s=0}^{\lfloor(p-3) / 2\rfloor} \frac{\left(1-\frac{1}{2 p}\right)^{p-2 s-1}\left(\frac{1}{2 p}\right)^{2 s+1}}{m^{s}}\binom{p}{2 s+1}\left(\boldsymbol{v}_{j}^{\top} \boldsymbol{a}\right)^{p-2 s-1}} \\
& =\sum_{j=1}^{k} S_{j} \boldsymbol{v}_{j} .
\end{aligned}
$$

Notice for even $p$,

$$
\begin{aligned}
& S_{j}=\lambda_{j}\left(\boldsymbol{v}_{j}^{\top} \boldsymbol{a}\right)^{3} \cdot \sum_{s=0}^{p / 2-2} \frac{\left(1-\frac{1}{2 p}\right)^{p-2 s-1}\left(\frac{1}{2 p}\right)^{2 s+1}}{m^{s}}\binom{p}{2 s+1}\left(\boldsymbol{v}_{j}^{\top} \boldsymbol{a}\right)^{p-2 s-4} \\
&=\lambda_{j}\left(\boldsymbol{v}_{j}^{\top} \boldsymbol{a}\right)^{3} \cdot \sum_{r=0}^{p / 2-2} \frac{\left(1-\frac{1}{2 p}\right)^{2 r+3}\left(\frac{1}{2 p}\right)^{p-3-2 r}}{m^{p / 2-2-r}}\binom{p}{p-2 r-3}\left(\boldsymbol{v}_{j}^{\top} \boldsymbol{a}\right)^{2 r} . \\
& \quad(\text { let } 2 r=p-4-2 s) \\
& \frac{S_{j}}{\lambda_{j}\left(\boldsymbol{v}_{j}^{\top} \boldsymbol{a}\right)^{3}}=\sum_{r=0}^{p / 2-2} \frac{\left(1-\frac{1}{2 p}\right)^{2 r+3}\left(\frac{1}{2 p}\right)^{p-3-2 r}}{m^{p / 2-2-r}}\binom{p}{p-2 r-3}\left(\boldsymbol{v}_{j}^{\top} \boldsymbol{a}\right)^{2 r}
\end{aligned}
$$

(Divide both sides by $\lambda_{j}\left(\boldsymbol{v}_{j}^{\top} \boldsymbol{a}\right)^{3}$ )
$\leq \sum_{r=0}^{p / 2-2} \frac{\left(1-\frac{1}{2 p}\right)^{2 r+3}\left(\frac{1}{2 p}\right)^{p-3-2 r}}{m^{p / 2-2-r}}\binom{p}{p-2 r-3}\left(\boldsymbol{v}_{1}^{\top} \boldsymbol{a}\right)^{2 r}$.
(Since the first term is constant and $\left|\boldsymbol{v}_{j}^{\top} \boldsymbol{a}\right| \leq \boldsymbol{v}_{1}^{\top} \boldsymbol{a}$ for $r \geq 1$ )

$$
=\frac{S_{1}}{\lambda_{1}\left(\boldsymbol{v}_{1}^{\top} \boldsymbol{a}\right)^{3}}
$$

Therefore for even $p \geq 4$ :

$$
\begin{equation*}
\left|S_{j}\right| \leq \frac{\left|\lambda_{j}\right|}{\lambda_{1}} \frac{\left|\boldsymbol{v}_{j}^{\top} \boldsymbol{a}\right|^{3}}{\left|\boldsymbol{v}_{1}^{\top} \boldsymbol{a}\right|^{3}} S_{1} \leq \frac{1}{4} \frac{\left|\boldsymbol{v}_{j}^{\top} \boldsymbol{a}\right|}{\left|\boldsymbol{v}_{1}^{\top} \boldsymbol{a}\right|} S_{1}, \forall j \geq 2 \tag{3}
\end{equation*}
$$

Similarly for odd $p$, we have:

$$
\begin{aligned}
& S_{j}= \lambda_{j}\left(\boldsymbol{v}_{j}^{\top} \boldsymbol{a}\right)^{2} \cdot \sum_{s=0}^{(p-3) / 2} \frac{\left(1-\frac{1}{2 p}\right)^{p-2 s-1}\left(\frac{1}{2 p}\right)^{2 s+1}}{m^{s}}\binom{p}{2 s+1}\left(v_{j}^{\top} a\right)^{p-2 s-3} \\
&= \lambda_{j}\left(\boldsymbol{v}_{j}^{\top} \boldsymbol{a}\right)^{2} \cdot \sum_{r=0}^{(p-3) / 2} \frac{\left(1-\frac{1}{2 p}\right)^{2 r+2}\left(\frac{1}{2 p}\right)^{p-2-2 r}}{m^{(p-3) / 2-r}}\binom{p}{p-2-2 r}\left(\boldsymbol{v}_{j}^{\top} \boldsymbol{a}\right)^{2 r}, \\
& \quad(\text { Let } r=(p-3) / 2-s) \\
& \frac{S_{j}}{\lambda_{j}\left(\boldsymbol{v}_{j}^{\top} \boldsymbol{a}\right)^{2}}=\sum_{r=0}^{(p-3) / 2} \frac{\left(1-\frac{1}{2 p}\right)^{2 r+2}\left(\frac{1}{2 p}\right)^{p-2-2 r}}{m^{(p-3) / 2-r}}\binom{p}{p-2-2 r}\left(\boldsymbol{v}_{j}^{\top} \boldsymbol{a}\right)^{2 r}
\end{aligned}
$$

(Divide both sides by $\lambda_{j}\left(\boldsymbol{v}_{j}^{\top} \boldsymbol{a}\right)^{2}$ )

$$
\leq \sum_{r=0}^{(p-3) / 2} \frac{\left(1-\frac{1}{2 p}\right)^{2 r+2}\left(\frac{1}{2 p}\right)^{p-2-2 r}}{m^{(p-3) / 2-r}}\binom{p}{p-2-2 r}\left(\boldsymbol{v}_{1}^{\top} \boldsymbol{a}\right)^{2 r}
$$

(Since the first term is constant and $\left|\boldsymbol{v}_{j}^{\top} \boldsymbol{a}\right| \leq \boldsymbol{v}_{1}^{\top} \boldsymbol{a}$ for $r \geq 1$ )

$$
=\frac{S_{1}}{\lambda_{1}\left(\boldsymbol{v}_{1}^{\top} \boldsymbol{a}\right)^{2}} .
$$

Therefore for odd $p$ we have:

$$
\begin{equation*}
\left|S_{j}\right| \leq \frac{\left|\lambda_{j}\right|}{\lambda_{1}} \frac{\left|\boldsymbol{v}_{j}^{\top} \boldsymbol{a}\right|^{2}}{\left|\boldsymbol{v}_{1}^{\top} \boldsymbol{a}\right|^{2}} S_{1} \leq \frac{1}{2} \frac{\left|\boldsymbol{v}_{j}^{\top} \boldsymbol{a}\right|}{\left|\boldsymbol{v}_{1}^{\top} \boldsymbol{a}\right|} S_{1}, \forall j \geq 2 \tag{4}
\end{equation*}
$$

Write $\boldsymbol{V}=\left[\boldsymbol{v}_{2}, \boldsymbol{v}_{3}, \cdots, \boldsymbol{v}_{k}\right] \in \mathbb{R}^{d \times k}$ be the complement for $\boldsymbol{v}_{1}$. Therefore for any $\boldsymbol{x}$ without normalization, one can conveniently represent $\left|\tan \theta\left(\boldsymbol{x}, \boldsymbol{v}_{1}\right)\right|$ as $\left\|\boldsymbol{V}^{\top} \boldsymbol{x}\right\|_{2} /\left|\boldsymbol{v}_{1}^{\top} x\right|$.

$$
\begin{align*}
\left\|\boldsymbol{V}^{\top} G(\boldsymbol{a})\right\|^{2} & =\sum_{j=2}^{k} S_{j}^{2}  \tag{5}\\
& \leq \sum_{j=2}^{k} \frac{\left|\boldsymbol{v}_{j}^{\top} \boldsymbol{a}\right|^{2}}{4\left|\boldsymbol{v}_{1}^{\top} \boldsymbol{a}\right|^{2}} S_{1}^{2}  \tag{4}\\
& =\frac{1}{4} \tan ^{2} \theta\left(\boldsymbol{v}_{1}, \boldsymbol{a}\right)\left(\boldsymbol{v}_{1}^{\top} G(\boldsymbol{a})\right)^{2} \tag{6}
\end{align*}
$$

Therefore for $p \geq 3, \tan \theta\left(G(\boldsymbol{a}), \boldsymbol{v}_{1}\right) \leq \frac{1}{2} \tan \theta\left(\boldsymbol{a}, \boldsymbol{v}_{1}\right)$.

## D.1.1 Proof Sketch of Theorem $\mathbf{3 . 1 4}$

Definition D. 2 (Zeroth order gradient function). For some scalar $m$, we define an empirical operator $G_{n}: \mathcal{A} \rightarrow \mathcal{A}$ that is similar to the zeroth-order gradient of $f$ through $n$ samples:

$$
G_{n}(\boldsymbol{a}):=\frac{m}{n} \sum_{i=1}^{n}\left(T\left(\left(\left(1-\frac{1}{2 p}\right) \boldsymbol{a}+\frac{1}{2 p} \boldsymbol{z}_{i}\right)^{\otimes p}\right)-\boldsymbol{T}\left(\frac{1}{2 p} \boldsymbol{z}_{i}\right)\right) \boldsymbol{z}_{i}+\left(\eta_{i}-\eta_{i}^{\prime}\right) \boldsymbol{z}_{i}
$$

where $\boldsymbol{z}_{i} \sim \mathcal{N}\left(0, \frac{1}{m} \boldsymbol{I}\right)$ and $\eta_{i}, \eta_{i}^{\prime}$ are independent zero-mean 1-sub-Gaussian noise. Therefore we have:

$$
\mathbb{E}\left[G_{n}(\boldsymbol{a})\right]=m \mathbb{E}\left[\sum_{l=0}^{p-1}\binom{p}{l} \boldsymbol{T}\left(\left(1-\frac{1}{2 p}\right)^{p-l} \boldsymbol{a}^{\otimes(p-l)} \otimes\left(\frac{1}{2 p}\right)^{l} \boldsymbol{z}^{\otimes l}\right) \boldsymbol{z}\right]
$$

(Due to symmetry of Gaussian only for odd $l=: 2 s+1$ expectation is nonzero)

$$
=\left(1-\frac{1}{2 p}\right)^{p-2 s-1}\left(\frac{1}{2 p}\right)^{2 s+1}\left[\sum_{s=0}^{\lfloor p / 2-1\rfloor} m^{-s}\binom{p}{2 s+1} \boldsymbol{T}\left(\boldsymbol{a}^{\otimes(p-2 s-1)} \otimes \boldsymbol{I}^{\otimes s+1}\right)\right]
$$

Note that for even $p$ the last term (when $s=p / 2-1)$ is $\boldsymbol{T}\left(\boldsymbol{a} \otimes \boldsymbol{I}^{\otimes p / 2}\right)=\sum_{j=1}^{k} \lambda_{j}\left(\boldsymbol{a}^{\top} \boldsymbol{v}_{j}\right) \boldsymbol{v}_{j}$. While all other terms will push the iterate towards the optimal action at a superlinear speed, the last term perform a matrix multiplication and the convergence speed will depend on the eigengap. Therefore for $p \geq 4$ we will remove the extra bias in the last term that is orthogonal to $\boldsymbol{v}_{1}$ and will treat it as noise. (Notice for quadratic function $s=0=p / 2-1$ is the only term in $\mathbb{E}\left[G_{n}(\boldsymbol{a})\right]$. This is the distinction between $p=2$ and larger $p$, and why its convergence depends on eigengap.)

We further define $G(\boldsymbol{a})$ as the population version of $G_{n}(\boldsymbol{a})$ by removing this undesirable bias term that will be treated as noise:

$$
\begin{aligned}
G(\boldsymbol{a}) & = \begin{cases}\mathbb{E}\left[G_{n}\right]-\frac{\left(\frac{1}{2 p}\right)^{p-1}\left(1-\frac{1}{2 p}\right) p}{m^{p / 2-1}} \sum_{j=2}^{k} \lambda_{j}\left(\boldsymbol{v}_{j}^{\top} \boldsymbol{a}\right) \boldsymbol{v}_{j}, & \begin{array}{l}
\text { when } p \text { is even } \\
\text { when } p \text { is odd. }
\end{array} \\
& =\sum_{s=0}^{\lfloor(p-3) / 2\rfloor} \frac{\left(\frac{1}{2 p}\right)^{2 s+1}}{m^{s}}\binom{p}{2 s+1} \boldsymbol{T}\left(\boldsymbol{I}^{\otimes s+1} \otimes\left(\left(1-\frac{1}{2 p}\right) \boldsymbol{a}\right)^{\otimes p-2 s-1}\right) \\
& =\frac{1}{2}\left(1-\frac{1}{2 p}\right)^{p-1} \boldsymbol{T}\left(\boldsymbol{I}, \boldsymbol{a}^{\otimes p-1}\right)+O(1 / m) .\end{cases}
\end{aligned}
$$

We define $G(\boldsymbol{a})$ to push the action $\boldsymbol{a}$ towards the $\boldsymbol{v}_{1}$ direction with at least linear convergence rate. More precisely, their angle $\tan \theta\left(G(\boldsymbol{a}), \boldsymbol{v}_{1}\right)$ will converge linearly to 0 for proper initialization with the dynamics $\boldsymbol{a} \rightarrow G(\boldsymbol{a})$. An easy way to see that is when $p=2$ or $3, G$ is conducting ( 3 -order tensor) power iteration. For higher-order problems, this operation $G$ is equivalent to the summation of $p, p-2, p-4, \cdots$-th order tensor product and hence the linear convergence.
The estimation error $G_{n}(\boldsymbol{a})-G(\boldsymbol{a})$ will be treated as noise (which is not mean zero when $p$ is even but will be small enough: $O\left((2 p)^{-p} m^{-(p-1) / 2}\right)$ ). Therefore the iterative algorithm with $\boldsymbol{a} \rightarrow G_{n}(\boldsymbol{a})$ will converge to a small neighborhood of $\boldsymbol{v}_{1}$ depending on the estimation error. This estimation error is controlled by the choice of sample size $n$ in each iteration. We now provide the proof sketch:

Lemma D. 3 (Initialization for $p \geq 3$; Corollary C. 1 from [70] ). For any $\eta \in(0,1 / 2)$, with $L=\Theta(k \log (1 / \eta))$ samples $\mathcal{A}=\left\{\overline{\boldsymbol{a}}^{(1)}, \boldsymbol{a}^{(2)}, \cdots \boldsymbol{a}^{(L)}\right\}$ where each $\boldsymbol{a}^{(l)}$ is sampled uniformly on the sphere $\mathbb{S}^{d-1}$. At least one sample $a \in \mathcal{A}$ satisfies

$$
\begin{equation*}
\max _{j \neq 1}\left|\boldsymbol{v}_{j}^{\top} \boldsymbol{a}\right| \leq 0.5\left|\boldsymbol{v}_{1}^{\top} \boldsymbol{a}\right|, \text { and }\left|\boldsymbol{v}_{1}^{\top} \boldsymbol{a}\right| \geq 1 / \sqrt{d} \tag{7}
\end{equation*}
$$

with probability at least $1-\eta$.
Lemma D. 4 (Iterative progress). Let $\alpha=1 / 2$ for $p \geq 3$ in Algorithm 6 Consider noisy operation $\boldsymbol{a}^{+} \rightarrow G(\boldsymbol{a})+\boldsymbol{g}$. If the error term $\boldsymbol{g}$ satisfies:

$$
\begin{aligned}
\|\boldsymbol{g}\| \leq & \min \left\{\frac{0.025}{p} \lambda_{1}\left(\boldsymbol{v}_{1}^{\top} \boldsymbol{a}\right)^{p-2}, 0.1 \lambda_{1} \widetilde{\varepsilon}\right\} \\
& +0.03 \lambda_{1}\left|\sin \theta\left(\boldsymbol{v}_{1}, \boldsymbol{a}\right)\right|\left(\boldsymbol{v}_{1}^{\top} \boldsymbol{a}\right)^{p-2} \\
\left|\boldsymbol{v}_{1}^{\top} \boldsymbol{g}\right| \leq & 0.05 \lambda_{1}\left(\boldsymbol{v}_{1}^{\top} \boldsymbol{a}\right)^{p-1}
\end{aligned}
$$

Suppose $\boldsymbol{a}$ satisfies $0.5\left|\boldsymbol{v}_{1}^{\top} \boldsymbol{a}\right| \geq \max _{j \geq 2}\left|\boldsymbol{v}_{j}^{\top} \boldsymbol{a}\right|$, we have:

$$
\tan \theta\left(\boldsymbol{a}^{+}, \boldsymbol{v}_{1}\right) \leq 0.8 \tan \theta\left(\boldsymbol{a}, \boldsymbol{v}_{1}\right)+\widetilde{\varepsilon}
$$

We can also bound $\boldsymbol{g}$ by standard concentration plus an additional small bias term.
Lemma D. 5 (Estimation error bound for $G$ ). For fixed value $\delta \in(0,1)$ and large enough universal constant $c_{1}, c_{2}, c_{m}, c_{n}$, when $m=c_{m} d \log (n / \delta), n \geq c_{n} d \log (d / \delta)$, we have

$$
\begin{array}{r}
\|\boldsymbol{g}\| \equiv\left\|G_{n}(\boldsymbol{a})-G(\boldsymbol{a})\right\| \leq c_{1} \sqrt{\frac{d^{2} \log ^{3}(n / \delta) \log (d / \delta)}{n}}+e \lambda_{2}\left|\sin \theta\left(\boldsymbol{a}, \boldsymbol{v}_{1}\right)\right|, \\
\left|\boldsymbol{v}_{1}^{\top} \boldsymbol{g}\right| \equiv\left|\boldsymbol{v}_{1}^{\top} G_{n}(\boldsymbol{a})-\boldsymbol{v}_{1}^{\top} G(\boldsymbol{a})\right| \leq c_{2} \sqrt{\frac{d \log ^{3}(n / \delta) \log (d / \delta)}{n}} .
\end{array}
$$

with probability $1-\delta . e=0$ for odd $p$ and $e=(2 p)^{-(p-1)} m^{-(p / 2-1)}$ for even $p$.
Together we are able to prove Theorem 3.14 .
Proof of Theorem 3.14. Initially with high probability there exists an $a_{0} \in \mathcal{A}_{0}$ such that Eqn. (7) holds, i.e., $\boldsymbol{v}_{1}^{\top} \boldsymbol{a}_{0} \geq 1 / \sqrt{d}$ and $\boldsymbol{v}_{1}^{\top} \boldsymbol{a}_{0} \geq 2\left|\boldsymbol{v}_{j}^{\top} \boldsymbol{a}_{0}\right|, \forall j \geq 2$.

Next, from Lemma D.5, the extra bias term is bounded by $e \lambda_{2}\left|\sin \theta\left(\boldsymbol{a}, \boldsymbol{v}_{1}\right)\right|$
$\leq 0.03 \lambda_{1}\left(\boldsymbol{v}_{1}^{\top} \boldsymbol{a}\right)^{p-2}\left|\sin \theta\left(\boldsymbol{a}, \boldsymbol{v}_{1}\right)\right|$ since $e=(2 p)^{-p+1} m_{s}^{-p / 2+1}$ and with our choice of variance $m_{s} \geq d \geq\left(\boldsymbol{v}_{1}^{\top} \boldsymbol{a}\right)^{-2}$, plus $p \geq 3$. Next with our setting of $n_{s}=\widetilde{\Theta}\left(d^{p} /\left(\lambda_{1}^{2} \widetilde{\varepsilon}_{t}^{2}\right)\right)$, the error term $\left\|\mathbb{E}[G(\boldsymbol{a})]-G_{n}(\boldsymbol{a})\right\|$ is upper bounded by $\widetilde{O}\left(\sqrt{\frac{d^{2}}{n}}\right) \leq 0.025 \lambda_{1} d^{-(p-2) / 2} \widetilde{\varepsilon}_{s} / p+0.1 \lambda_{1} \widetilde{\varepsilon}_{s}$. Meanwhile $\left|\boldsymbol{v}_{1}^{\top} \boldsymbol{g}\right| \leq \widetilde{O}\left(\sqrt{\frac{d}{n_{s}}}\right) \leq 0.05 \lambda_{1}\left(\boldsymbol{v}_{1}^{\top} \boldsymbol{a}\right)^{p-1}$.
This meets the requirements for TheoremD. 4 and therefore $\tan \theta\left(G_{n}\left(\boldsymbol{a}_{0}\right), \boldsymbol{v}_{1}\right) \leq 0.8 \tan \theta\left(\boldsymbol{a}_{0}, \boldsymbol{v}_{1}\right)+$ $0.1 \lambda_{1} \widetilde{\varepsilon}_{s}$. Therefore after $l$ steps will have

$$
\begin{aligned}
\tan \theta\left(G_{n}^{l}\left(\boldsymbol{a}_{0}\right), \boldsymbol{v}_{1}\right) & \leq 0.8^{l} \tan \theta\left(\boldsymbol{a}_{0}, \boldsymbol{v}_{1}\right)+\sum_{i=1}^{l} 0.8^{i} \cdot 0.1 \widetilde{\varepsilon}_{s} \\
& \leq 0.8^{l} \tan \theta\left(\boldsymbol{a}_{0}, \boldsymbol{v}_{1}\right)+0.5 \widetilde{\varepsilon}_{s}
\end{aligned}
$$

Notice initially $\tan \theta\left(\boldsymbol{a}_{0}, \boldsymbol{v}_{1}\right) \leq 1 /\left(\boldsymbol{v}_{1}^{\top} \boldsymbol{a}_{0}\right) \leq \sqrt{d}$. Therefore after at most
$l=O\left(\log _{2}\left(\tan \theta\left(\boldsymbol{a}_{0}, \boldsymbol{v}_{1}\right)\right)\right) \leq O\left(\log _{2}(d)\right)$ steps, we will have $\tan \left(G_{n}^{l}\left(\boldsymbol{a}_{0}\right), \boldsymbol{v}_{1}\right) \leq \widetilde{\varepsilon}_{0} / 2=\widetilde{\varepsilon}_{1}$. With the same argument, the progress also holds for $s>0$ with even smaller $l$.

Proof of Lemma D.5. We first estimate $G_{n}(\boldsymbol{a})-\mathbb{E}\left[G_{n}(\boldsymbol{a})\right]$, which is want we want for even $p$. For odd $p$ we will need to analyze an extra bias term that is orthogonal to $\boldsymbol{v}_{1}, \boldsymbol{e}:=$ $\frac{\left(\frac{1}{2 p}\right)^{p-1}\left(1-\frac{1}{2 p}\right) p}{m^{p / 2-1}} \sum_{j=2}^{k} \lambda_{j}\left(\boldsymbol{v}_{j}^{\top} \boldsymbol{a}\right) \boldsymbol{v}_{j} ;$ and we have $G_{n}(\boldsymbol{a})-\mathbb{E}\left[G_{n}(\boldsymbol{a})\right]=G_{n}(\boldsymbol{a})-G(\boldsymbol{a})+\boldsymbol{e}$.

We decompose $G_{n}(\boldsymbol{a})$ as $G_{n}(\boldsymbol{a})=\sum_{s=1}^{k} G_{n}^{(s)}+N$, where $G_{n}^{(s)}:=\frac{m}{n} \sum_{i=1}^{n}\binom{p}{s} T(((1-$ $\left.0.5 / p) \boldsymbol{a})^{\otimes p-s} \otimes\left(\boldsymbol{z}_{i} /(2 p)\right)^{\otimes s}\right) \boldsymbol{z}_{i}$. The noise term $N:=\frac{m}{n} \sum \epsilon_{i} \boldsymbol{z}_{i}$.

$$
\begin{aligned}
G_{n}^{(s)} & :=\frac{m}{n} \sum_{i=1}^{n}\binom{p}{s} T\left(((1-0.5 / p) \boldsymbol{a})^{\otimes p-s} \otimes\left(\boldsymbol{z}_{i} /(2 p)\right)^{\otimes s}\right) \boldsymbol{z}_{i} \\
& =\frac{m}{n}\left(1-\frac{1}{2 p}\right)^{p-s}\left(\frac{1}{2 p}\right)^{s}\binom{p}{s} \sum_{i=1}^{n} \sum_{j=1}^{k} \lambda_{j}\left(\boldsymbol{a}^{\top} \boldsymbol{v}_{j}\right)^{p-s}\left(\boldsymbol{z}_{i}^{\top} \boldsymbol{v}_{j}\right)^{s} \boldsymbol{z}_{i} . \\
\mathbb{E}\left[G_{n}^{(s)}\right] & =m\left(1-\frac{1}{2 p}\right)^{p-s}\left(\frac{1}{2 p}\right)^{s}\binom{p}{s} \sum_{j=1}^{k} \lambda_{j}\left(\boldsymbol{a}^{\top} \boldsymbol{v}_{j}\right)^{p-s} \mathbb{E}\left[\left(\boldsymbol{z}^{\top} \boldsymbol{v}_{j}\right)^{s} \boldsymbol{z}\right] \\
& = \begin{cases}\left(1-\frac{1}{2 p}\right)^{p-s}\left(\frac{1}{2 p}\right)^{s} m\binom{p}{s} \sum_{j=1}^{k} \lambda_{j}\left(\boldsymbol{a}^{\top} \boldsymbol{v}_{j}\right)^{p-s} \frac{1}{m^{(s+1) / 2}}(s)!!\boldsymbol{v}_{j}, & \text { for odd } s, \\
0, & \text { for even } s\end{cases} \\
& = \begin{cases}\left(1-\frac{1}{2 p}\right)^{p-s}\left(\frac{1}{2 p}\right)^{s} \frac{s!!}{m^{(s-1) / 2}}\binom{p}{s} \sum_{j=1}^{k} \lambda_{j}\left(\boldsymbol{a}^{\top} \boldsymbol{v}_{j}\right)^{p-s} \boldsymbol{v}_{j}, & \text { for odd } s, \\
0, & \text { for even } s\end{cases} \\
& G_{n}^{(s)}-\mathbb{E}\left[G_{n}^{(s)}\right]=m\left(1-\frac{1}{2 p}\right)^{p-s}\left(\frac{1}{2 p}\right)^{s}\binom{p}{s} \sum_{j=1}^{k} \lambda_{j}\left(\boldsymbol{a}^{\top} \boldsymbol{v}_{j}\right)^{p-s} \boldsymbol{g}_{n, s}(j),
\end{aligned}
$$

where $\boldsymbol{g}_{n, s}(j):=\frac{1}{n} \sum_{i=1}^{n}\left(\boldsymbol{z}_{i}^{\top} \boldsymbol{v}_{j}\right)^{s} \boldsymbol{z}_{i}-\mathbb{E}\left[\left(\boldsymbol{z}^{\top} \boldsymbol{v}_{j}\right)^{s} \boldsymbol{z}\right]$.
Notice the scaling in each $G_{n}^{(s)}$ is $\left(1-\frac{1}{2 p}\right)^{p-s}\left(\frac{1}{2 p}\right)^{s}\binom{p}{s} \leq\left(\frac{1}{2 p}\right)^{s} p^{s} /(s!)<2^{-s}$ decays exponentially. In Claim D.12 we give bounds for $g_{n, s}(j)$. We note the bound for each $g_{n, s}$ also decays with $s$.
Therefore the bottleneck of the upper bound mostly depend on $\boldsymbol{g}_{n, 0}$ and $\boldsymbol{v}_{1}^{\top} \boldsymbol{g}_{n, 0}$, and we get:

$$
\begin{aligned}
\left\|G_{n}-\mathbb{E}\left[G_{n}\right]\right\| \leq C_{1} \lambda_{1} \sqrt{\frac{(d+\log (1 / \delta)) d \log (n / \delta)}{n}}+N, \\
\left|\boldsymbol{v}_{1}^{\top} G_{n}-\mathbb{E}\left[\boldsymbol{v}_{1}^{\top} G_{n}\right]\right| \leq C_{2} \lambda_{1} \sqrt{\frac{d \log (n / \delta)(1+\log (1 / \delta))}{n}}+\boldsymbol{v}_{1}^{\top} N .
\end{aligned}
$$

Next from Claim D.11, the noise term

$$
\begin{aligned}
N \leq & C_{3} \sqrt{\frac{m \log (n / \delta)(d+\log (n / \delta)) \log (d / \delta)}{n}}, \\
& \left|v_{1}^{\top} N\right| \leq C_{4} \sqrt{m \frac{\log ^{2}(n / \delta) \log (d / \delta)}{n}},
\end{aligned}
$$

Finally $\boldsymbol{e}$ is very small: $\|\boldsymbol{e}\| \leq \frac{1}{m^{p / 2-1}(2 p)^{(p-1)}} \lambda_{2}\left\|\boldsymbol{V}^{\top} \boldsymbol{a}\right\|=\lambda_{2} \frac{1}{m^{p / 2-1}(2 p)(p-1)} \sin \theta\left(\boldsymbol{a}, \boldsymbol{v}_{1}\right)$. $\left|\boldsymbol{e}^{\top} \boldsymbol{v}_{1}\right|=0$.
Together we can bound $G_{n}(\boldsymbol{a})-G(\boldsymbol{a})$ and finish the proof.

Proof of Lemma D.4 From Lemma D.1 we have: $\left|\tan \theta\left(G(\boldsymbol{a}), \boldsymbol{v}_{1}\right)\right| \leq 1 / 2\left|\tan \theta\left(\boldsymbol{a}, \boldsymbol{v}_{1}\right)\right|$. Let $\boldsymbol{V}=\left[\boldsymbol{v}_{2}, \cdots \boldsymbol{v}_{k}\right]$. For any $p \geq 2$, we have:

$$
\begin{aligned}
\left|\tan \theta\left(\boldsymbol{a}^{+}, \boldsymbol{v}_{1}\right)\right| & =\frac{\left\|\boldsymbol{V}^{\top} \boldsymbol{a}^{+}\right\|_{2}}{\left|\boldsymbol{v}_{1}^{\top} \boldsymbol{a}^{+}\right|} \\
& =\frac{\left\|\boldsymbol{V}^{\top}(G(\boldsymbol{a})+\boldsymbol{g})\right\|}{\left|\boldsymbol{v}_{1}^{\top}(G(\boldsymbol{a})+\boldsymbol{g})\right|} \\
& \leq \frac{\left\|\boldsymbol{V}^{\top} G(\boldsymbol{a})\right\|+\left\|\boldsymbol{V}^{\top} \boldsymbol{g}\right\|}{\left|\boldsymbol{v}_{1}^{\top} G(\boldsymbol{a})\right|-\left|\boldsymbol{v}_{1}^{\top} \boldsymbol{g}\right|} \\
& \leq \frac{1 / 2\left|\tan \theta\left(\boldsymbol{a}, \boldsymbol{v}_{1}\right)\right|\left\|\boldsymbol{v}_{1}^{\top} G(\boldsymbol{a}) \mid+\right\| \boldsymbol{g} \|}{\left|\boldsymbol{v}_{1}^{\top} G(\boldsymbol{a})\right|-\left|\boldsymbol{v}_{1}^{\top} \boldsymbol{g}\right|} \\
& =\alpha\left|\tan \theta\left(\boldsymbol{a}, \boldsymbol{v}_{1}\right)\right| \frac{S_{1}}{S_{1}-\left\|\boldsymbol{v}_{1}^{\top} \boldsymbol{g}\right\|}+\frac{\|\boldsymbol{g}\|}{S_{1}-\left\|\boldsymbol{v}_{1}^{\top} \boldsymbol{g}\right\|},
\end{aligned}
$$

where $S_{1}:=\boldsymbol{v}_{1}^{\top} G(\boldsymbol{a})=\boldsymbol{v}_{1}^{\top} G(\boldsymbol{a})=\lambda_{1} \sum_{s=0}^{\lfloor(p-3) / 2\rfloor} \frac{\left(1-\frac{1}{2 p}\right)^{p-2 s-1}\left(\frac{1}{2 p}\right)^{2 s+1}}{m^{s}}\binom{p}{2 s+1}\left(\boldsymbol{v}_{1}^{\top} \boldsymbol{a}\right)^{p-2 s-1} \geq$ $\lambda_{1}\left(1-\frac{1}{2 p}\right)^{p-1}\left(\frac{1}{2 p}\right) p\left(\boldsymbol{v}_{1}^{\top} \boldsymbol{a}\right)^{p-1} \geq \frac{\lambda_{1}}{4}\left(\boldsymbol{v}_{1}^{\top} \boldsymbol{a}\right)^{p-1}$. The inequality comes from keeping only the first term where $s=0$. With the assumption that $\left|\boldsymbol{v}_{1}^{\top} \boldsymbol{g}\right| \leq 0.05 \lambda_{1}\left(\boldsymbol{v}_{1}^{\top} \boldsymbol{a}\right)^{p-1}$, we have $\left|\boldsymbol{v}_{1}^{\top} \boldsymbol{g}\right| \leq 0.2 S_{1}$. Therefore

$$
\begin{aligned}
\left|\tan \theta\left(\boldsymbol{a}^{+}, \boldsymbol{v}_{1}\right)\right| & \leq 1.25 / 2\left|\tan \theta\left(\boldsymbol{a}, \boldsymbol{v}_{1}\right)\right|+\frac{\|\boldsymbol{g}\|}{S_{1}-\left|\boldsymbol{v}_{1}^{\top} \boldsymbol{g}\right|} \\
& \leq 1.25 / 2\left|\tan \theta\left(\boldsymbol{a}, \boldsymbol{v}_{1}\right)\right|+5 / 4 \frac{\|\boldsymbol{g}\|}{S_{1}} \\
& \leq 1.25 / 2\left|\tan \theta\left(\boldsymbol{a}, \boldsymbol{v}_{1}\right)\right|+5 \frac{\|\boldsymbol{g}\|}{\lambda_{1}\left(\boldsymbol{v}_{1}^{\top} \boldsymbol{a}\right)^{p-1}} .
\end{aligned}
$$

Notice when $5 \frac{\|\boldsymbol{g}\|}{\lambda_{1}\left(\boldsymbol{v}_{1}^{\top} \boldsymbol{a}\right)^{p-1}} \leq \max \left\{0.125\left|\tan \theta\left(\boldsymbol{a}, \boldsymbol{v}_{1}\right)\right|, \widetilde{\epsilon}\right\}$, which will ensure $\left|\tan \theta\left(G(\boldsymbol{a}), \boldsymbol{v}_{1}\right)\right| \leq$ $(1.25 / 2+0.125)\left|\tan \theta\left(G(\boldsymbol{a}), \boldsymbol{v}_{1}\right)\right|+\widetilde{\epsilon}$. (We will prove this condition is satisfied when $\|\boldsymbol{g}\| \leq$ $\min \left\{\frac{0.025}{p} \lambda_{1}\left(\boldsymbol{v}_{1}^{\top} \boldsymbol{a}\right)^{p-2}, 0.1 \lambda_{1} \widetilde{\varepsilon}\right\}$. We will handle the additional term in the upper bound of $\|\boldsymbol{g}\|$ later.) We divide this requirement into the following two cases. On one hand, when $\left|\boldsymbol{v}_{1}^{\top} \boldsymbol{a}\right| \leq 1-1 /(p-1)$, $\left\|\boldsymbol{V}^{\top} \boldsymbol{a}\right\| \geq \sqrt{1-(1-1 /(p-1))^{2}}>1 / p$, therefore $\left|\tan \theta\left(\boldsymbol{a}, \boldsymbol{v}_{1}\right)\right| \geq 1 / p\left|\boldsymbol{v}_{1}^{\top} \boldsymbol{a}\right|$. Therefore

$$
\begin{aligned}
& 5 \frac{\|\boldsymbol{g}\|}{\lambda_{1}\left(\boldsymbol{v}_{1}^{\top} \boldsymbol{a}\right)^{p-1}} \leq 0.125\left|\tan \theta\left(\boldsymbol{a}, \boldsymbol{v}_{1}\right)\right| \\
\Leftarrow & 5 \frac{\|\boldsymbol{g}\|}{\lambda_{1}\left(\boldsymbol{v}_{1}^{\top} \boldsymbol{a}\right)^{p-1}} \leq 0.125 /\left(p\left|\boldsymbol{v}_{1}^{\top} \boldsymbol{a}\right|\right) \\
\Leftrightarrow & \|\boldsymbol{g}\| \leq 0.025 \lambda_{1}\left(\boldsymbol{v}_{1}^{\top} \boldsymbol{a}\right)^{p-2} / p .
\end{aligned}
$$

On the other hand, when $\left|\boldsymbol{v}_{1}^{\top} \boldsymbol{a}\right| \geq 1-1 /(p-1),\left|\boldsymbol{v}_{1}^{\boldsymbol{a}}\right|^{p-1} \geq 1 / 4$ when $p=3$. Therefore $5 \frac{\|\boldsymbol{g}\|}{\lambda_{1}\left(\boldsymbol{v}_{1}^{\top} \boldsymbol{a}\right)^{p-1}} \leq 20\|\boldsymbol{g}\| / \lambda_{1}$. Therefore we will need $\|\boldsymbol{g}\| \leq 0.05 \lambda_{1} \tilde{\epsilon}$, and then the requirement that $5 \frac{\|\boldsymbol{g}\|}{\lambda_{1}\left(\boldsymbol{v}_{1}^{\top} \boldsymbol{a}\right)^{p-1}} \leq \tilde{\epsilon}$ is satisfied.
Altogether in both cases we have: $\left|\tan \theta\left(G(\boldsymbol{a}), \boldsymbol{v}_{1}\right)\right| \leq 0.75\left|\tan \theta\left(G(\boldsymbol{a}), \boldsymbol{v}_{1}\right)\right|+\widetilde{\epsilon}$. Finally if we additionally increase $\|\boldsymbol{g}\|$ by $0.05 \lambda_{1}\left(\boldsymbol{v}_{1}^{\top} \boldsymbol{a}\right)^{p-1}$ we will have: $\left|\tan \theta\left(G(\boldsymbol{a}), \boldsymbol{v}_{1}\right)\right| \leq$ $0.8\left|\tan \theta\left(G(\boldsymbol{a}), \boldsymbol{v}_{1}\right)\right|+\widetilde{\epsilon}$.

Proof of Corollary 3.16. As shown in Theorem 3.14 at least one action $\boldsymbol{a}$ in $\mathcal{A}_{S},\left|\mathcal{A}_{S}\right| \leq \widetilde{O}(k)$ satisfies $\tan \theta\left(\boldsymbol{a}, \boldsymbol{a}^{*}\right) \leq \varepsilon$ with a total of $\widetilde{O}\left(\frac{d^{p} k}{\lambda_{1}^{2} \varepsilon^{2}}\right)$ steps. Therefore with Claim D. 6 we have to get $\widetilde{\varepsilon}$-optimal reward we need $\widetilde{O}\left(\frac{d^{p} k}{\lambda_{1} \tilde{\varepsilon}}\right)$ steps. Notice the eluder dimension for symmetric polynomials is $d^{p}$ and the size of $\mathcal{A}_{S}$ is at most $\widetilde{O}(k)$. Then by applying Corollary D.9 we get that the total regret is at most $\widetilde{O}\left(\sqrt{d^{p} k T}+\sqrt{\left|\mathcal{A}_{S}\right| T}\right)=\widetilde{O}\left(\sqrt{d^{p} k T}\right)$.

## D. 2 PAC to Regret Bound Relation.

Claim D. 6 (Connecting angle to regret). When $0<\tan \theta\left(\boldsymbol{a}, \boldsymbol{v}_{1}\right) \leq \zeta$, we have regret $r^{*}-r(\boldsymbol{a}) \leq$ $r^{*} \min \left\{2, p \zeta^{2}\right\}$.

Proof.

$$
\begin{aligned}
\left|\cos \theta\left(\boldsymbol{a}, \boldsymbol{v}_{1}\right)\right| & =\left|\boldsymbol{a}^{\top} \boldsymbol{v}_{1}\right|=: b \\
\left|\tan \theta\left(\boldsymbol{a}, \boldsymbol{v}_{1}\right)\right| & =\frac{\sqrt{1-b^{2}}}{b} \leq \zeta \Leftrightarrow b \geq \frac{1}{\sqrt{\zeta^{2}+1}} . \\
\Rightarrow r^{*}-r(\boldsymbol{a}) & \leq \lambda_{1}-\lambda_{1} b^{p} \\
& \leq \lambda_{1}\left(1-\left(\zeta^{2}+1\right)^{-p / 2}\right) \\
& =\lambda_{1} \frac{\left(\zeta^{2}+1\right)^{p / 2}-1}{\left(\zeta^{2}+1\right)^{p / 2}} \\
& \leq \lambda_{1}\left(\left(\zeta^{2}+1\right)^{p / 2}-1\right) \\
& \leq \lambda_{1} p \zeta^{2}, \text { when } \zeta^{2} \leq 1 / p
\end{aligned}
$$

$$
\leq \lambda_{1}\left(\left(\zeta^{2}+1\right)^{p / 2}-1\right) \quad\left(\text { since denominator }\left(\zeta^{2}+1\right)^{p / 2} \geq 1\right)
$$

Additionally by definition $r^{*}-r(\boldsymbol{a}) \leq \lambda_{1}-\left(-\lambda_{1}\right)=2 \lambda_{1}$ and thus $r^{*}-r(a) \leq \lambda_{1} \min \left\{2, p \zeta^{2}\right\}$. We now derive the last inequality. When $\zeta \geq 1 / p$ it is trivially true. When $\zeta \leq 1 / p$, we have $\left(1+\zeta^{2}\right)^{p / 2} \leq 1+p \zeta^{2}$ for any $p \geq 2$. Since the LHS is a convex function for $\zeta$ when $p \geq 2$ and when $\zeta=0$ LHS $=$ RHS and when $\bar{\zeta}^{2}=1 / p$ LHS is always smaller than RHS (=2).
Notice the argument is straightforward to extend to the setting where the angle is between $\boldsymbol{a}$ and subspace $V_{1}$ that satisfies $\forall \boldsymbol{v} \in \boldsymbol{V}_{1}, T(\boldsymbol{v}) \geq \lambda_{1}-\epsilon$, then one also get $r^{*}-r(\boldsymbol{a}) \leq \lambda_{1}-\left(\lambda_{1}-\epsilon\right) b^{p} \leq$ $\min \left\{\lambda_{1}, \lambda_{1} p \zeta^{2}+\epsilon b^{p}\right\} \leq \min \left\{\lambda_{1}, \lambda_{1} p \zeta^{2}+\epsilon\right\}$.

Claim D. 7 (Connecting PAC to Cumulative Regret). Suppose we have an algorithm alg( $\zeta$ ) that finds $\zeta$-optimal action $\widehat{\boldsymbol{a}}$ that satisfies $0<\tan \theta\left(\boldsymbol{a}, \boldsymbol{v}_{1}\right) \leq \zeta$ by taking $A \zeta^{-a}$ actions. Here $A$ can depend on any parameters such as $d, \lambda_{1}$, probability error $\delta$, etc., that are not $\zeta$. Then for large enough $T$, by calling alg with $\zeta=A^{\frac{1}{a+2}} T^{-\frac{1}{a+2}} p^{-\frac{1}{a+2}}$ and playing its output action $\widehat{\boldsymbol{a}}$ for the remaining actions, one can get a cumulative regret of:

$$
\mathfrak{R}(T) \lesssim T^{\frac{a}{a+2}} p^{\frac{a}{a+2}} A^{\frac{2}{a+2}} r^{*} .
$$

Similarly, if an oracle finds $\varepsilon$-optimal action $\widehat{\boldsymbol{a}}$ that satisfies $r^{*}-r(\boldsymbol{a}) \leq \varepsilon$ with $B \varepsilon^{-b}$ samples, then by setting $\varepsilon=\left(B r^{*} / T\right)^{\frac{1}{1+b}}$, and playing the output arm for the remaining actions, one can get cumulative regret of:

$$
\mathfrak{R}(T) \lesssim B^{\frac{1}{1+b}} T^{\frac{b}{1+b}} r^{\frac{1}{1+b}}
$$

Proof. For the chosen $\zeta$, write $T_{1}=A \zeta^{-a}$ be the number of actions that finds $\zeta$-optimal action. Therefore $T_{1}=A^{\frac{2}{a+2}} T^{\frac{a}{a+2}} p^{\frac{a}{a+2}}$. First, when $T \geq A p^{a / 2}, \zeta^{2} \leq 1 / p$, namely $r^{*}-r(\boldsymbol{a}) \leq r^{*} p \zeta^{2}$. We have:

$$
\begin{aligned}
\mathfrak{R}(T) & \leq \sum_{t=1}^{T_{1}} 2 r^{*}+\sum_{t=T_{1}+1}^{T} r^{*} p \zeta^{2} \\
& \leq 2 r^{*} T_{1}+T r^{*} p \zeta^{2} \\
& \leq 3 T^{\frac{a}{a+2}} p^{\frac{a}{a+2}} A^{\frac{2}{a+2}} r^{*} .
\end{aligned}
$$

When $T<A p^{a / 2}$, it trivially holds that $\mathfrak{R}(T) \leq 2 r^{*} T<2 T^{\frac{a}{a+2}} p^{\frac{a}{a+2}} A^{\frac{2}{a+2}} r^{*}$.
Theorem D. 8 (Theorem 5.1 from [5]). With UCB algorithm on action set with size $K$, we have with probability $1-\delta$,

$$
\mathfrak{R}(T)=\widetilde{O}(\min \{\sqrt{K T}\}+K)
$$

```
Algorithm 7 UCB (Algorithm 1 in Section 5 of [5])
    Input: Stochastic reward function \(f\), failure probability \(\delta\), action set \(\mathcal{A}\) with finite size \(K\).
    for \(t\) from 1 to \(T-1-K\) do
        Execute \(\operatorname{arm} I_{t}=\arg \max _{i \in[K]}\left(\widehat{\mu}^{t}(i)+\sqrt{\frac{\log (T K / \delta)}{N^{t}(i)}}\right)\). Here \(N^{t}(\boldsymbol{a})=1+\sum_{i=1}^{t} \mathbf{1}\left\{I_{i}=\right.\)
        \(\boldsymbol{a}\}\); and \(\widehat{\mu}^{t}(\boldsymbol{a})=\frac{1}{N^{t}(\boldsymbol{a})}\left(r_{a}+\sum_{i=1}^{t} \mathbf{1}\left\{I_{i}=\boldsymbol{a}\right\} r_{i}\right)\).
            Observe \(r_{I_{t}}\)
```

Corollary D.9. With the same setting of Claim D.7 except that now the algorithm alg $(\varepsilon)$ finds a set $\mathcal{A}$ of size $S$ where at least one action $\boldsymbol{a} \in \mathcal{A}$ satisfies $r^{*}-f(\boldsymbol{a}) \leq \varepsilon$. Then all argument in Claim D. 7 still hold by adding $\widetilde{O}(\sqrt{S T})$ on the RHS of each regret bound.

Proof. Suppose we run alg for $T_{1}$ steps and achieve $\varepsilon$-optimal reward.
Let $r_{\varepsilon}:=\max _{\boldsymbol{a} \in \mathcal{A}} f(\boldsymbol{a})$. Therefore with UCB on mutiarm bandit we have: $\sum_{t=T_{1}+1}^{T} r_{\varepsilon}-f\left(\boldsymbol{a}_{t}\right) \leq$ $\widetilde{O}(\sqrt{S T})$ by Theorem D.8
From the statement $r_{\varepsilon} \geq r^{*}-\zeta$. Therefore $\sum_{t=T_{1}+1}^{T} r^{*}-f\left(\boldsymbol{a}_{t}\right) \leq \widetilde{O}(\sqrt{S T})+\varepsilon\left(T-T_{1}\right)$. Therefore

$$
\mathfrak{R}(T) \leq \sum_{t=1}^{T_{1}} 2 r^{*}+\varepsilon\left(T-T_{1}\right)+\widetilde{O}(\sqrt{S T})
$$

With the same choices of $T_{1}$ in Claim D.7, the same conclusion still holds with an additional term of $\widetilde{O}(\sqrt{S T})$.

For symmetric tensor problems the set size is $\widetilde{O}(k)$ and therefore we will have an additional $\sqrt{k T}$ term which will be subsumed in our regret bound.

## D. 3 Variance and Noise Concentration

Lemma D. 10 (Vector Bernstein; adapted from Theorem 7.3.1 in [65]). Consider a finite sequence $\left\{\boldsymbol{x}_{k}\right\}_{k=1}^{n}$ be i.i.d randomly generated samples, $x_{k} \in \mathbb{R}^{d}$, and assume that $\mathbb{E}\left[\boldsymbol{x}_{k}\right]=0,\left\|\boldsymbol{x}_{k}\right\| \leq L$, and covariance matrix of $x_{k}$ is $\Sigma$. Then it satisfies that when $n \geq \log d / \delta$, we have:

$$
\left\|\frac{\sum_{i=1}^{n} \boldsymbol{x}_{i}}{n}\right\| \leq C \sqrt{\frac{\left(\|\Sigma\|+L^{2}\right) \log d / \delta}{n}}
$$

with probability $1-\delta$.
Claim D. 11 (Noise concentration). Let independent samples $\boldsymbol{z}_{i} \sim \mathcal{N}\left(0,1 / m I_{d}\right)$ and $\epsilon_{i} \sim \mathcal{N}(0,1)$. With probability $1-\delta, \delta \in(0,1)$ :

$$
\begin{aligned}
\left\|\frac{m}{n} \sum_{i=1}^{n} \epsilon_{i} \boldsymbol{z}_{i}\right\| & \leq C \sqrt{\frac{m \log (n / \delta)(d+\log (n / \delta)) \log (d / \delta)}{n}} \\
\left|\frac{m}{n} \sum_{i=1}^{n} \epsilon_{i} \boldsymbol{z}_{i}^{\top} \boldsymbol{v}_{1}\right| & \leq C^{\prime} \sqrt{\frac{m \log ^{2}(n / \delta) \log (d / \delta)}{n}}
\end{aligned}
$$

Proof. We use the Vector Bernstein Lemma D.10. The covariance matrix for $\boldsymbol{x}_{i}=\epsilon_{i} \boldsymbol{z}_{i}$ satisfies $\mathbb{E}\left[\boldsymbol{x}_{i} \boldsymbol{x}_{i}^{\top}\right]=1 / m I_{d} . \boldsymbol{x}_{i}$ is mean zero. $\left\|\epsilon_{i} \boldsymbol{z}_{i}\right\|^{2}=\epsilon_{i}^{2}\left\|\boldsymbol{z}_{i}\right\|^{2}$. Notice $\epsilon_{i}^{2} \sim \chi(1) \lesssim 1+\log (1 / \delta)$ and $m \boldsymbol{z}_{i}^{\top} \boldsymbol{z}_{i} \sim \chi(d) \lesssim d+\log (1 / \delta)$. Therefore by directly applying Vector Bernstein $\left\|\epsilon_{i} \boldsymbol{z}_{i}\right\| \leq$ $c \sqrt{\frac{(1+\log (1 / \delta))(d+\log (1 / \delta))}{m}}$ with probability $1-\delta$. By union bound we have: for all $i,\left\|\epsilon_{i} \boldsymbol{z}_{i}\right\| \leq$ $c \sqrt{\frac{\log (n / \delta)(d+\log (n / \delta))}{m}}$ with probability $1-\delta$. Therefore

$$
\left\|\frac{1}{n} \sum_{i=1}^{n} \epsilon_{i} \boldsymbol{z}_{i}\right\| \leq C \sqrt{\frac{\log (n / \delta)(d+\log (n / \delta)) \log (d / \delta)}{m n}},
$$

with probability $1-\delta$. Similarly

$$
\begin{aligned}
\left|\frac{1}{n} \sum_{i=1}^{n} \epsilon_{i} \boldsymbol{z}_{i}^{\top} \boldsymbol{v}_{1}\right| & \leq C \sqrt{\frac{\log (n / \delta)(1+\log (n / \delta)) \log (d / \delta)}{m n}} \\
& =C^{\prime} \sqrt{\frac{\log ^{2}(n / \delta) \log (d / \delta)}{m n}}
\end{aligned}
$$

Claim D.12. Let $\left\{\boldsymbol{z}_{i}\right\}_{i=1}^{n}$ be i.i.d samples from $\mathcal{N}\left(0,1 / m I_{d}\right)$. Let $g_{n, s}(j):=\frac{1}{n} \sum_{i=1}^{n}\left(\boldsymbol{z}_{i}^{\top} \boldsymbol{v}_{j}\right)^{s} \boldsymbol{z}_{i}-$ $\mathbb{E}\left[\left(\boldsymbol{z}^{\top} \boldsymbol{v}_{j}\right)^{s} \boldsymbol{z}\right]$. We have:

$$
\begin{gathered}
\left\|g_{n, 0}(j)\right\| \lesssim \sqrt{\frac{d+\log (1 / \delta)}{n m}}, \\
\left|\boldsymbol{v}_{1}^{\top} g_{n, 0}(j)\right| \lesssim \sqrt{\frac{1+\log (1 / \delta)}{n m}}, \\
\left|\boldsymbol{v}_{1}^{\top} g_{n, 1}(j)\right| \leq\left\|g_{n, 1}(j)\right\| \lesssim \sqrt{\frac{d+\log (1 / \delta)}{m^{2} n}}, \text { when } n \geq d \log (1 / \delta) \\
\left|\boldsymbol{v}_{1}^{\top} g_{n, s}(j)\right| \leq\left\|g_{n, s}(j)\right\| \lesssim \sqrt{\frac{\log (d / \delta)}{d^{s} n}}, \text { when } n \geq \log (d / \delta), m \geq c_{0} d \log (n / \delta), s \geq 2
\end{gathered}
$$

For any $j \in[k]$.
We mostly care about the correct concentration for smaller $s$. For larger $s$ a very loose bound will already suffice our requirement.

Proof of Claim D.12. For $s=0, n m\left\|\frac{1}{n} \sum_{i=1}^{n} \boldsymbol{z}_{i}\right\|^{2} \sim \chi(d)$, therefore $\left\|\frac{1}{n} \sum_{i=1}^{n} \boldsymbol{z}_{i}\right\| \lesssim$ $\sqrt{\frac{d+\log (1 / \delta)}{n m}} . n m\left(\frac{1}{n} \sum_{i=1}^{n} \boldsymbol{z}_{i}^{\top} \boldsymbol{v}_{1}\right)^{2} \sim \chi(1)$. Therefore $\left|\frac{1}{n} \sum_{i=1}^{n} \boldsymbol{z}_{i}^{\top} \boldsymbol{v}_{1}\right| \lesssim \sqrt{\frac{1+\log (1 / \delta)}{n m}}$.
For $s=1$, due to standard concentration for covariance matrices (see e.g. [65, 22] ), we have:

$$
m\left\|\left(\frac{1}{n} \sum_{i=1}^{n} \boldsymbol{z}_{i} \boldsymbol{z}_{i}^{\top}-\mathbb{E}\left[\boldsymbol{z} \boldsymbol{z}^{\top}\right]\right)\right\| \leq \max \left\{\sqrt{\frac{d+\log (2 / \delta)}{n}}, \frac{d+\log (2 / \delta)}{n}\right\}
$$

Therefore when $n \geq d \log (1 / \delta)$, both results

$$
\begin{aligned}
\left\|g_{n, 1}(j)\right\| & \lesssim \sqrt{\frac{d+\log (1 / \delta)}{m^{2} n}}\left\|\boldsymbol{v}_{j}\right\| \\
& =\sqrt{\frac{d+\log (1 / \delta)}{m^{2} n}}, \text { and } \\
\left\|\boldsymbol{v}_{1}^{\top} g_{n, 1}(j)\right\| & \lesssim \sqrt{\frac{d+\log (1 / \delta)}{m^{2} n}}\left\|\boldsymbol{v}_{1}\right\|\left\|\boldsymbol{v}_{j}\right\| \\
& =\sqrt{\frac{d+\log (1 / \delta)}{m^{2} n}}
\end{aligned}
$$

hold.
For larger $s \geq 2$, with probability $1-\delta,\left|\boldsymbol{z}_{i}^{\top} \boldsymbol{v}_{j}\right| \leq C \sqrt{\log (n / \delta) / m}=C c_{0} / \sqrt{d} \leq 1 / \sqrt{d}$. When $m \geq c_{0} d \log (n / \delta)$, for small enough $c_{0}$ we have $\left|\boldsymbol{z}_{i}^{\top} \boldsymbol{v}_{j}\right| \leq 1 / \sqrt{d}$ and $\left\|\boldsymbol{z}_{i}\right\| \leq 1$ for all $i \in[n]$. Therefore $\left\|\left(\boldsymbol{z}_{i}^{\top} \boldsymbol{v}_{j}\right)^{s} \boldsymbol{z}_{i}\right\| \leq d^{-s / 2}$ We can use vector Bernstein, i.e., Lemma D. 10 to get:

$$
\left\|g_{n, s}(j)\right\| \leq C_{1} \sqrt{\frac{\log (d / \delta)}{d^{s} n}}
$$

Therefore we have:

$$
\left|g_{n, s}(j)^{\top} \boldsymbol{v}_{1}\right| \leq C_{1} \sqrt{\frac{\log (d / \delta)}{d^{s} n}}
$$

## D.3.1 The asymmetric setting

Now we consider the asymmetric tensor problem with reward $f: \mathcal{A} \rightarrow \mathbb{R}$. The input space $\mathcal{A}$ consists of $p$ vectors in a unit ball: $\overrightarrow{\boldsymbol{a}}=(\boldsymbol{a}(1), \boldsymbol{a}(2), \cdots \boldsymbol{a}(p)) \in \mathcal{A},\|\boldsymbol{a}(s)\| \leq 1, \forall s \in[p]$. $f(\overrightarrow{\boldsymbol{a}})=\boldsymbol{T}\left(\otimes_{s=1}^{p} \boldsymbol{a}(s)\right)+\eta$. Tensor $\boldsymbol{T}=\sum_{j=1}^{k} \lambda_{j} \boldsymbol{v}_{j}(1) \otimes \boldsymbol{v}_{j}(2) \cdots \otimes \boldsymbol{v}_{j}(p)$. For each $s \in[p]$, $\left\{\boldsymbol{v}_{1}(s), \boldsymbol{v}_{2}(s), \cdots \boldsymbol{v}_{k}(s)\right\}$ are orthonormal vectors. We order the eigenvalues such that $\lambda_{1} \geq$ $\left|\lambda_{2}\right| \cdots \geq\left|\lambda_{k}\right|$. Therefore the optimal reward is $\lambda_{1}$ and can be achieved by $\boldsymbol{a}^{*}(s)=\boldsymbol{v}_{1}(s), s \in[p]$. In this section we only consider $p \geq 3$ and leave the quadratic and low-rank matrix setting to the next section.
Theorem D.13. For $p \geq 3$, by conducting alternating power iteration, one can get a $\varepsilon$-optimal reward with a total $\widetilde{O}\left((2 k)^{p} \log ^{p}(p / \delta) d^{p} \lambda_{1}^{-1} \varepsilon^{-1}\right)$ actions; therefore the regret bound is at most $\widetilde{O}\left(\sqrt{k^{p} d^{p} T}\right)$.

This setting is actually much easier than the symmetric setting. Notice by replacing one slice of $\overrightarrow{\boldsymbol{a}}$ by random Gaussian $\boldsymbol{z}_{i} \sim \mathcal{N}(0,2 / d \log (d / \delta))$, one directly gets $\boldsymbol{T}(\boldsymbol{a}(1), \cdots \boldsymbol{a}(s-1), \boldsymbol{I}, \boldsymbol{a}(s+$ $1), \cdots \boldsymbol{a}(p))$ on each slice with $1 / n \sum_{i} f\left(\boldsymbol{a}(1), \cdots \boldsymbol{a}(s-1), \boldsymbol{z}_{i}, \boldsymbol{a}(s+1), \cdots \boldsymbol{a}(p)\right) \boldsymbol{z}_{i}$ which is tensor product. We defer the proof to Appendix D. 4

## D. 4 Omitted Details for Asymmetric Tensors

```
Algorithm 8 Phased elimination with alternating tensor product.
    Input: Stochastic reward \(r:\left(B_{1}^{d}\right)^{\otimes p} \rightarrow \mathbb{R}\) of polynomial degree \(p\), failure probability \(\delta\), error \(\varepsilon\).
    Initialization: \(L_{0}=C_{L} k \log (1 / \delta)\); Total number of stages \(S=C_{S}\lceil\log (1 / \varepsilon)\rceil+1, \mathcal{A}_{0}=\)
    \(\left\{\boldsymbol{a}_{0}^{(1)}, \boldsymbol{a}_{0}^{(2)}, \cdots \boldsymbol{a}_{0}^{\left(L_{0}\right)}\right\} \subset\left(B_{1}^{d}\right)^{\otimes p}\) where each \(\boldsymbol{a}_{0}^{(l)}(j), j \in[p]\) is uniformly sampled on the unit
    sphere \(\mathbb{S}^{d-1} . \widetilde{\varepsilon}_{0}=1\).
    for \(s\) from 1 to \(S\) do
        \(\widetilde{\varepsilon}_{s} \leftarrow \widetilde{\varepsilon}_{s-1} / 2, n_{s} \leftarrow C_{n} d^{p} \log (d / \delta) / \widetilde{\varepsilon}_{s}^{2}, n_{s} \leftarrow n_{s} \cdot \log ^{3}\left(n_{s} / \delta\right), m_{s} \leftarrow C_{m} d \log (n / \delta)\),
    \(\mathcal{A}_{s}=\varnothing\).
        for \(l\) from 1 to \(L_{s-1}\) do
            Tensor product update:
            Locate current arm \(\widetilde{\boldsymbol{a}}=\boldsymbol{a}_{s-1}^{(l)}\).
            for \(\left\lceil\left(\lambda_{1} / \Delta\right) \log (2 d)\right\rceil\) times do
                    for \(j\) from 1 to \(p\) do
                    Sample \(\boldsymbol{z}_{i} \sim \mathcal{N}\left(0,1 / m_{s} I_{d}\right), i=1,2, \cdots n_{s}\).
                    Calculate tentative arm \(\boldsymbol{a}_{i} \leftarrow \widetilde{\boldsymbol{a}}, \boldsymbol{a}_{i}(j)=\left(1-\widetilde{\varepsilon_{s}}\right) \widetilde{\boldsymbol{a}}(j)+\widetilde{\varepsilon}_{s} \boldsymbol{z}_{i}\)
                    Conduct estimation \(\boldsymbol{y} \leftarrow 1 / n_{s} \sum_{i=1}^{n_{s}} r_{\epsilon_{i}}\left(\boldsymbol{a}_{i}\right) \boldsymbol{z}_{i}\).
                    Update the current arm \(\widetilde{\boldsymbol{a}}(j) \leftarrow \boldsymbol{y} /\|\boldsymbol{y}\|\).
                Estimate the expected reward for \(\widetilde{\boldsymbol{a}}\) through \(n_{s}\) samples: \(r_{n}=1 / n_{s} \sum_{i=1}^{n_{s}} r_{\epsilon_{i}}(\widetilde{\boldsymbol{a}})\).
                Candidate Elimination:
                if \(r_{n} \geq \lambda_{1}\left(1-p \widetilde{\varepsilon}_{s}\right)\) then
                    Keep the arm \(\mathcal{A}_{s} \leftarrow \mathcal{A}_{s} \cup\{\widetilde{\boldsymbol{a}}\}\)
        Label the arms: \(L_{s}=\left|\mathcal{A}_{s}\right|, \mathcal{A}_{t}=:\left\{\boldsymbol{a}_{s}^{(1)}, \cdots \boldsymbol{a}_{s}^{\left(L_{s}\right)}\right\}\).
    Run Algorithm 7 with \(\mathcal{A}_{S}\).
```

Lemma D. 14 (Asymmetric Tensor Initialization). With probability $1-\delta$, with $L=$ $\widetilde{\Theta}\left((2 k)^{p} \log ^{p}(p / \delta)\right)$ random initializations $\mathcal{A}_{0}=\left\{\boldsymbol{a}_{0}^{(0)}, \boldsymbol{a}_{0}^{(1)}, \cdots \boldsymbol{a}_{0}^{(L)}\right\}$, there exists an initialization $\boldsymbol{a}_{0} \in \mathcal{A}_{0}$ that satisfies:

$$
\begin{array}{r}
\alpha \boldsymbol{a}_{0}(s)^{\top} \boldsymbol{v}_{1}^{(s)} \geq\left|\boldsymbol{a}_{0}(s)^{\top} \boldsymbol{v}_{j}^{(s)}\right|, \forall j \geq 2 \& j \in[k], \forall s \in[p],  \tag{8}\\
\boldsymbol{a}_{0}(s)^{\top} \boldsymbol{v}_{1}^{(s)} \geq 1 / \sqrt{d .}
\end{array}
$$

with some constant $\alpha<1$.
Proof. This lemma simply comes from applying Lemma D. 3 for $p$ times and we need $\geq 2 k \log _{2}(p \delta)$ to ensure the condition for each $\boldsymbol{a}_{0}(s), s \in[p]$ holds. Therefore together we will need $\left(2 k \log _{2}(p / \delta)\right)^{p}$ samples.

Lemma D. 15 (Asymmetric tensor progress). For each a that satisfies Eqn. 8) with constant $\alpha<1$, we have:

$$
\tan \theta\left(\boldsymbol{T}(\boldsymbol{a}(1), \cdots \boldsymbol{a}(s-1), \boldsymbol{I}, \boldsymbol{a}(s+1), \cdots \boldsymbol{a}(p)), \boldsymbol{v}_{1}^{(s)}\right) \leq \alpha \tan \theta\left(\boldsymbol{a}_{j}, \boldsymbol{v}_{1}^{(j)}\right)
$$

for any $j$ that is in $[p]$ but is not $s$. When $n \geq \Theta\left(d^{p} \log (d / \delta) \log ^{3}(n / \delta) / \widetilde{\varepsilon}^{2}\right)$ and $m=\Theta(d \log (n / \delta))$, we have:

$$
\begin{aligned}
& \tan \theta\left(\boldsymbol{T}(\boldsymbol{a}(1), \cdots \boldsymbol{a}(s-1), \boldsymbol{I}, \boldsymbol{a}(s+1), \cdots \boldsymbol{a}(p)), \boldsymbol{v}_{1}^{(s)}\right) \\
\leq & (1+\alpha) / 2 \tan \theta\left(\boldsymbol{a}_{j}, \boldsymbol{v}_{1}^{(j)}\right)+\widetilde{\varepsilon}, \forall j \in[p] \& j \neq s
\end{aligned}
$$

The remaining proof is a simpler version for the symmetric tensor setting on conducting noisy power method with the good initialization and iterative progress.
Finally due to the good initialization that satisfies (8) and together with Lemma D.15 we can finish the proof for Theorem D. 13 .

## E Proof of Theorem 3.21

## E. 1 Additional Notations

Here, we briefly introduce complex and real algebraic geometry. This section is based on [56, 64, 12 , 69].
An (affine) algebraic variety is the common zero loci of a set of polynomials, defined as $V=$ $Z(S)=\left\{\boldsymbol{x} \in \mathbb{C}^{n}: f(\boldsymbol{x})=0, \forall f \in S\right\} \subseteq \mathbb{A}^{n}=\mathbb{C}^{n}$ for some $S \subseteq \mathbb{C}\left[x_{1}, \cdots, x_{n}\right]$. A projective variety $U$ is a subset of $\mathbb{P}^{n}=\left(\mathbb{C}^{n+1} \backslash\{0\}\right) / \sim$, where $\left(x_{0}, \cdots, x_{n}\right) \sim k\left(x_{0}, \cdots, x_{n}\right)$ for $k \neq 0$ and $S$ is a set of homogeneous polynomials of $(n+1)$ variables.
For an affine variety $V$, its projectivization is the variety $\mathbb{P}(V)=\{[\boldsymbol{x}]: x \in V\} \subseteq \mathbb{P}^{n-1}$, where $[\boldsymbol{x}]$ is the line corresponding to $\boldsymbol{x}$.
The Zariski topology is the topology generated by taking all varieties to be the closed sets.
A set is irreducible if it is not the union of two proper closed subsets.
A variety is irreducible if and only if it is irreducible under the Zariski topology.
The algebraic dimension $d=\operatorname{dim} V$ of a variety $V$ is defined as the length of the longest chain $V_{0} \subset V_{1} \subset \cdots \subset V_{d}=V$, such that each $V_{i}$ is irreducible.
A variety $V$ is said to be admissible to a set of linear functions $\left.\left\{\ell_{\alpha}: \mathbb{C}^{d} \rightarrow \mathbb{C}\right\}_{\alpha \in I}\right\}$, if for every $\ell_{\alpha}$, we have $\operatorname{dim}\left(V \cap\left\{\boldsymbol{x} \in \mathbb{C}^{d}: \ell_{\alpha}(\boldsymbol{x})=0\right\}\right)<\operatorname{dim} V$.
A map $f=\left(f_{1}, \cdots, f_{m}\right): \mathbb{A}^{n} \rightarrow \mathbb{A}^{m}$ is regular if each $f_{i}$ is a polynomial.
A algebraic set is the common real zero loci of a set of polynomials.
For a complex variety $V \subseteq \mathbb{A}^{n}$, its real points form a algebraic set $V_{\mathbb{R}}$.
For an algebraic set $V_{\mathbb{R}}$, its real dimension $d=\operatorname{dim}_{\mathbb{R}} V_{\mathbb{R}}$ is the maximum number $d$ such that $V_{\mathbb{R}}$ is locally semi-algebraically homeomorphic to the unit cube $(0,1)^{d}$, details can be found in [12].

## E. 2 Proof of Sample Complexity

Lemma E. 1 ([69], Theorem 3.2). For $i=1, \ldots, T$, let $L_{i}: \mathbb{C}^{n} \times \mathbb{C}^{m} \rightarrow \mathbb{C}$ be bilinear functions and $V_{i}$ be varieties given by homogeneous polynomials in $\mathbb{C}^{n}$. Let $V=V_{1} \times \cdots \times V_{T} \subseteq\left(\mathbb{C}^{n}\right)^{N}$. Let $W \subseteq \mathbb{C}^{m}$ be a variety given by homogeneous polynomials. In addition, we assume $V_{i}$ is admissible with respect to the linear functions $\left\{f^{\boldsymbol{w}}(\cdot)=L_{i}(\cdot, \boldsymbol{w}): \boldsymbol{w} \in W \backslash\{0\}\right\}$. When $T \geq \operatorname{dim} W$, let $\delta=T-\operatorname{dim} W+1 \geq 1$. Then there exists a subvariety $Z \subseteq V$ with $\operatorname{dim} Z \leq \operatorname{dim} V-\delta$ such that for any $\left(\boldsymbol{x}_{1}, \ldots, \boldsymbol{x}_{T}\right) \in V \backslash Z$ and $\boldsymbol{w} \in W$, if $L_{1}\left(\boldsymbol{x}_{1}, \boldsymbol{w}\right)=\cdots=L_{T}\left(\boldsymbol{x}_{T}, \boldsymbol{w}\right)=0$, then $\boldsymbol{w}=0$.
Lemma E. 2 ([69], Lemma 3.1). Let $V$ be an algebraic variety in $\mathbb{C}^{d}$. Then $\operatorname{dim}_{\mathbb{R}} V_{\mathbb{R}} \leq \operatorname{dim} V$.
Lemma E.3. Let $W$ be a vector space. For vectors $\boldsymbol{x}_{1}, \cdots, \boldsymbol{x}_{T}$, if the map $f: \boldsymbol{w} \mapsto$ $\left(\left\langle\boldsymbol{x}_{1}, w\right\rangle, \ldots,\left\langle\boldsymbol{x}_{T}, \boldsymbol{w}\right\rangle\right)$ is not injective over $W-W:=\left\{\boldsymbol{w}_{1}-\boldsymbol{w}_{2}: \boldsymbol{w}_{1}, \boldsymbol{w}_{2} \in W\right\}$, then there exists $\boldsymbol{v} \in W$ such that $f(\boldsymbol{v})=0$.

Proof. Suppose $f\left(\boldsymbol{w}_{1}\right)=f\left(\boldsymbol{w}_{2}\right)$. Let $\boldsymbol{v}=\boldsymbol{w}_{1}-\boldsymbol{w}_{2}$. Then $\boldsymbol{v} \in W-W$ and $f(\boldsymbol{v})=f\left(\boldsymbol{w}_{1}\right)-$ $f\left(\boldsymbol{w}_{2}\right)=0$.
Definition E. 4 (Tensorization). Let $f$ be a polynomial of $x_{1}, \cdots, x_{d}$ with degree $\operatorname{deg} f \leq p$. Then every $p$-tensor $\mathbf{W}_{f}$ satisfying $\left\langle\mathbf{W}_{f}, \mathbf{X}_{\boldsymbol{x}}\right\rangle=f(\boldsymbol{x})$ is said to be a tensorization of the polynomial $f$, where $\mathbf{X}_{\boldsymbol{x}}$ is the tensorization of $\boldsymbol{x}$ itself:

$$
\begin{equation*}
\mathbf{X}_{\boldsymbol{x}}=\binom{1}{\boldsymbol{x}}^{\otimes p} \tag{9}
\end{equation*}
$$

Let $\mathcal{F}$ be a class of polynomials. A variety of tensorization of $\mathcal{F}$ is defined to be an irredicuble closed variety defined by homogeneous polynomials $W$, such that for every $f \in \mathcal{F}$, there is a tensorization $\mathbf{W}_{f}$ of $f$, such that $W \ni \mathbf{W}_{f}$ contains its tensorization. Note that neither tensorization of $f$ nor variety of tensorization of $\mathcal{F}$ is unique.
We define the variety of tensorization of $\boldsymbol{x}$ as follows. (Note that this is uniquely defined.) Consider the regular map

$$
\begin{equation*}
\varphi_{1}: \mathbb{C}^{d} \rightarrow \mathbb{C}^{(d+1)^{p}}, \quad \boldsymbol{x} \mapsto\binom{1}{\boldsymbol{x}}^{\otimes p} \tag{10}
\end{equation*}
$$

the tensorization of $\boldsymbol{x}$ is defined as $V_{i}=\mathbb{P}\left(\overline{\operatorname{Im} \varphi_{1}}\right)$.
Note that $V_{i}$ is irreducible because $\varphi_{1}$ is regular and $\mathbb{C}^{d}$ is irreducible. By [56, Theorem 9.9], its dimension is given by

$$
\begin{equation*}
\operatorname{dim} V_{i} \leq \operatorname{dim} \overline{\operatorname{Im} \varphi_{1}}+1 \leq \operatorname{dim} \mathbb{C}^{d}+1=d+1 \tag{11}
\end{equation*}
$$

Lemma E.5. For any non-zero polynomial $f \neq 0$ with $\operatorname{deg} f \leq p$. Let $W_{f}$ be a tensorization of $f$. Then $V_{i}$ is admissible with respect to $\left\{L_{i}(\cdot)=\left\langle\cdot, W_{f}\right\rangle\right\}$.

Proof. Since $V_{i}$ is irredicuble and $L_{i}$ is a linear function, it suffices to verify that $\left\langle\mathbf{X}_{\boldsymbol{x}}, W_{f}\right\rangle \neq 0$ [69]. But according to Definition E. $4,\left\langle\mathbf{X}_{\boldsymbol{x}}, W_{f}\right\rangle \neq 0$ is equivalent to

$$
\begin{equation*}
f(\boldsymbol{x})=\left\langle\mathbf{W}_{f},\binom{1}{\boldsymbol{x}}^{\otimes p}\right\rangle \neq 0 \tag{12}
\end{equation*}
$$

Since $f \neq 0$, we must have $f(\boldsymbol{x}) \neq 0$ for some $\boldsymbol{x}$, which gives a non-zero $\mathbf{X}_{\boldsymbol{x}} \neq 0$ for the above equation: $\left\langle\mathbf{X}_{x}, W_{f}\right\rangle \neq 0$, and we conclude that $V_{i}$ is admissible.
Lemma E.6. Let $V \subset \mathbb{C}^{n}$ be a (Zariski) closed proper subset, $V \neq \mathbb{C}^{n}$. Then $V$ is a null set, i.e. it has (Lebesgue) measure zero.

Proof. Suppose $V=Z(S)$ is the vanishing set for some $S \subseteq \mathbb{C}\left[x_{1}, \cdots, x_{n}\right]$. Since $V \neq \mathbb{C}^{n}$, let $f \in S$, we have $V \subseteq Z(f)$, so it suffices to show $\operatorname{Leb}(Z(f))=0$, which is because $Z(f)=f^{-1}(0), \operatorname{Leb}(\{0\})=0, f$ is a continuous function (under Euclidean topology), and $\operatorname{Leb}(\{\boldsymbol{x}: \nabla f(\boldsymbol{x})=0\})=0$.

Theorem E.7. Assume that the reward function class is a class of polynomials $\mathcal{F}$. Let $W$ be (one of) its variety of tensorization. If we sample $T \geq \operatorname{dim} W$ times, and the sample points satisfying $\left(\boldsymbol{x}_{1}, \cdots, \boldsymbol{x}_{T}\right) \in\left(\mathbb{C}^{d}\right)^{T} \backslash Z$ for some null set $Z$. Then we can uniquely determine the reward function $f$ from the observed rewards $\left(f\left(\boldsymbol{x}_{1}\right), \cdots, f\left(\boldsymbol{x}_{T}\right)\right)$.

Proof. Let $n=m=(d+1)^{p}, L_{i}(\boldsymbol{x}, \boldsymbol{w})=\langle\boldsymbol{x}, \boldsymbol{w}\rangle, V=V_{1} \times \cdots \times V_{T}$, where $V_{i}$ is as in Definition E. 4 By [64, Example 1.33], we have $\operatorname{dim} V \leq(d+1) T$. Since $W$ is a vareity of tensorization, by LemmaE.5. $V_{i}$ is admissible with respect to $\left\{L_{i}(\cdot, \mathbf{W}): \mathbf{W} \in W\right\}$.
We are now ready to apply Lemma E. 1 , which gives that when $T \geq \operatorname{dim} W$, there exists subvariety $Z \subset V$ with $\operatorname{dim} Z<\operatorname{dim} V \leq r T$, and for any $\left(\mathbf{X}_{1}, \cdots, \mathbf{X}_{T}\right) \in V \backslash Z$ and any $\mathbf{W} \in W$, if $\left\langle\mathbf{X}_{1}, \mathbf{W}\right\rangle=\cdots=\left\langle\mathbf{X}_{T}, \mathbf{W}\right\rangle=0$, then $\mathbf{W}=0$. By Lemma E.3, we have for every $\left(\mathbf{X}_{1}, \cdots, \mathbf{X}_{T}\right) \in$ $V \backslash Z$, the map $\mathbf{W} \mapsto\left(\left\langle\mathbf{X}_{1}, \mathbf{W}\right\rangle, \cdots,\left\langle\mathbf{X}_{T}, \mathbf{W}\right\rangle\right)$ is injective, so $\mathbf{W}_{f}$ and thus $f$ can be uniquely recovered from the observed rewards.

Finally, we show that $\left(\varphi_{1}^{-1} \times \cdots \times \varphi_{1}^{-1}\right)(Z)$ is a null set, where $\varphi_{1}$ is as in 10 . According to the proof of Lemma E. 1 by [69], we find that $Z$ is also defined by homogeneous polynomials. We take the slice $Z^{\prime}=\left\{\boldsymbol{x} \in Z: x_{11}=\cdots=x_{T 1}=1\right\}, V^{\prime}=\left\{\boldsymbol{x} \in V: x_{11}=\cdots=x_{T 1}=1\right\}$, (here $x_{i j}$ is the $j$-th coordinate of $\boldsymbol{x}_{i}$ ), then $Z^{\prime}, V^{\prime}$ are varieties. Since $\operatorname{dim} Z<\operatorname{dim} V$, we have $\operatorname{dim} Z^{\prime}=\operatorname{dim} Z-T<\operatorname{dim} V-T=\operatorname{dim} V^{\prime}$ and $Z^{\prime} \subset V^{\prime}$.
Now consider the regular map $\varphi_{1}^{\prime}: V^{\prime} \rightarrow\left(\mathbb{C}^{d}\right)^{T}$,

$$
\begin{equation*}
\left(\binom{1}{\boldsymbol{x}_{1}}^{\otimes p}, \cdots,\binom{1}{\boldsymbol{x}_{1}}^{\otimes p}\right) \mapsto\left(\boldsymbol{x}_{1}, \cdots, \boldsymbol{x}_{T}\right) \tag{13}
\end{equation*}
$$

Then $\varphi_{1}^{\prime}\left(Z^{\prime}\right), \varphi_{1}^{\prime}\left(V^{\prime}\right)$ are both varieties. By [56, Lemma 9.9], we have $\operatorname{dim} \overline{\varphi_{1}^{\prime}\left(Z^{\prime}\right)} \leq \operatorname{dim} Z$. Since $\varphi_{1}^{\prime}(V)=\left(\mathbb{C}^{d}\right)^{T}$ and $\operatorname{dim} \varphi_{1}^{\prime}\left(V^{\prime}\right) \leq \operatorname{dim} \underline{V^{\prime}=\operatorname{dim} V-T \leq(d+1) T-T \text {, we have }{ }^{\prime} \text {, }{ }^{\prime}(Z)}$ $\operatorname{dim} V=\operatorname{dim} \varphi_{1}^{\prime}(V)=d T$ and as a result, $\operatorname{dim} \overline{\varphi_{1}^{\prime}(Z)} \leq \operatorname{dim} Z<\operatorname{dim} V=d T$. By Lemma E. $6 \overline{\varphi_{1}^{\prime}(Z)}$ is a null set. Since $\left(\boldsymbol{x}_{1}, \cdots, \boldsymbol{x}_{T}\right) \notin \overline{\varphi_{4}^{\prime}(Z)}$ implies that $\left(\varphi_{1}\left(\boldsymbol{x}_{1}\right), \cdots, \varphi_{1}\left(\boldsymbol{x}_{T}\right)\right) \notin Z$, we conclude the proof.

TheoremE.7 is stated for complex sample points. Next we extend it to the real case.
Lemma E.8. In Lemma E.1] if we assume in addition that $\operatorname{dim}_{\mathbb{R}} V_{\mathbb{R}}=\operatorname{dim} V$, then the conclusion can be enhanced to ensure that $Z$ is a real subvariety and $\operatorname{dim}_{\mathbb{R}} Z<\operatorname{dim}_{\mathbb{R}} V_{\mathbb{R}}$.
Lemma E.9. Let $V \subset \mathbb{R}^{n}$ be a (Zariski) closed proper subset, $V \neq \mathbb{R}^{n}$. Then $V$ is a null set.
The proof of LemmaE. 9 is the same as that of LemmaE. 6
Theorem E.10. We can additionally assume $\boldsymbol{x}_{i} \in \mathbb{R}^{d}$ in Theorem E. 7
Proof. We verify that $\operatorname{dim} V=\operatorname{dim}_{\mathbb{R}} V_{\mathbb{R}}$, where $V$ is defined in the proof of Theorem E.7, but this follows clearly by [12, Corollary 2.8.2]. We conclude the proof by applying LemmaE. 8 .

Finally, we apply Theorem E. 10 to two concrete classes of polynomials, namely Examples 3.22 and 3.23 For Example 3.22, we construct its variety of tensorization of $\mathcal{R}_{\mathcal{V}}$ as follows. We first construct the tensorization of each polynomial. We define

$$
\begin{equation*}
\mathbf{W}_{f}=\sum_{i=1}^{r} a_{i}\binom{1}{\boldsymbol{w}_{i}}^{\otimes p_{i}} \otimes\binom{1}{0}^{\otimes\left(p-p_{i}\right)} \tag{14}
\end{equation*}
$$

Next we construct the variety of tensorization $W$. Consider the map $\varphi_{2}:\left(\mathbb{C}^{d}\right)^{r} \rightarrow \mathbb{C}^{(d+1)^{p}}$,

$$
\begin{equation*}
\varphi_{2}\left(\boldsymbol{w}_{1}, \cdots, \boldsymbol{w}_{r}\right)=\sum_{i=1}^{r}\binom{1}{\boldsymbol{w}_{i}}^{\otimes p_{i}} \otimes\binom{1}{0}^{\otimes\left(p-p_{i}\right)} \tag{15}
\end{equation*}
$$

and let $Y=\mathbb{P}\left(\overline{\operatorname{Im} \varphi_{2}}\right)$. Similar to $V_{i}$, we can prove that $Y$ is an irredicuble closed variety defined by homogeneous polynomials with $\operatorname{dim} Y \leq d r+1$. Next consider the map $\varphi_{2}^{\prime}:\left(\mathbb{C}^{d}\right)^{2 r} \rightarrow \mathbb{C}^{(d+1)^{p}}$,

$$
\begin{equation*}
\varphi_{2}^{\prime}\left(\boldsymbol{w}_{1}, \cdots, \boldsymbol{w}_{2 r}\right)=\varphi_{2}\left(\boldsymbol{w}_{1}, \cdots, \boldsymbol{w}_{r}\right)-\varphi_{2}\left(\boldsymbol{w}_{r+1}, \cdots, \boldsymbol{w}_{2 r}\right) \tag{16}
\end{equation*}
$$

and let $W=\mathbb{P}\left(\overline{\operatorname{Im} \varphi_{2}^{\prime}}\right)$. Similar to $Y$, we can prove that $W$ is an irredicuble closed variety defined by homogeneous polynomials with $\operatorname{dim} W \leq 2 d r+1$. Together with Theorem E.10, we can conlude that the optimal action for Example 3.22 can be uniquely determined using at most $2 d r+1$ samples.
For Example 3.23, we construct $W$ as follows. Let

$$
\boldsymbol{U}=\left(\begin{array}{lll}
\boldsymbol{w}_{1} & \cdots & \boldsymbol{w}_{k}
\end{array}\right), \quad q=\sum_{I \subseteq[k]:|I| \leq p} a_{I} x^{I},
$$

then we construct the tensorization of each polynomial by

$$
\begin{equation*}
\mathbf{W}_{f}=\sum_{I \subseteq[k]:|I| \leq p} a_{I} \bigotimes_{i \in I}\binom{1}{\boldsymbol{w}_{i}} \otimes\binom{1}{0}^{\otimes(p-|I|)} \tag{17}
\end{equation*}
$$

Then we have $f(\boldsymbol{x})=\left\langle\mathbf{W}_{f}, \mathbf{X}_{\boldsymbol{x}}\right\rangle$. To reduce the dimension of $W$ and get better sample complexity bound, we construct in a manner slightly different from what we did for Example 3.22. Consider the $\operatorname{map} \varphi_{3}:\left(\mathbb{C}^{d}\right)^{k} \times \mathbb{C}^{(k+1)^{p}} \rightarrow \mathbb{C}^{(d+1)^{p}}$,

$$
\begin{equation*}
\left(\boldsymbol{w}_{1}, \cdots, \boldsymbol{w}_{k}\right) \times\left(a_{I}: I \subseteq[k],|I| \leq p\right) \mapsto \mathbf{W}_{f} \tag{18}
\end{equation*}
$$

where $\mathbf{W}_{f}$ is as defined in 17). Let $Y=\mathbb{P}\left(\overline{\operatorname{Im} \varphi_{3}}\right)$ and $W=\mathbb{P}\left(\overline{\operatorname{Im} \varphi_{3}-\operatorname{Im} \varphi_{3}}\right)$. We end up with $\operatorname{dim} Y=\leq d k+(k+1)^{p}+1, \operatorname{dim} W \leq 2\left(d k+(k+1)^{p}\right)+1$. So we conlude that the optimal action for Example 3.23 can be uniquely determined using at most $2 d k+2(k+1)^{p}+1$ samples.

## F Omitted Proof for Lower Bounds with UCB Algorithms

In this section, we provide the proof for the lower bounds for learning with UCB algorithms in Subsection G.0.1

Notation Recall that we use $\Lambda$ to denote the subset of the $p$-th multi-indices $\Lambda=\left\{\left(\alpha_{1}, \ldots, \alpha_{p}\right) \mid 1 \leq\right.$ $\left.\alpha_{1}<\cdots<\alpha_{p} \leq d\right\}$. For an $\alpha=\left(\alpha_{1}, \ldots, \alpha_{p}\right) \in \Lambda$, denote $\boldsymbol{M}_{\alpha}=\boldsymbol{e}_{\alpha_{1}} \otimes \cdots \otimes \boldsymbol{e}_{\alpha_{p}}, \boldsymbol{A}_{\alpha}=$ $\left(\boldsymbol{e}_{\alpha_{1}}+\cdots+\boldsymbol{e}_{\alpha_{p}}\right)^{\otimes p}$. The model space $\mathcal{M}$ is a subset of rank-1 $p$-th order tensors, which is defined as $\mathcal{M}=\left\{\boldsymbol{M}_{\alpha} \mid \alpha \in \Lambda\right\}$. We define the core action set $\mathcal{A}_{0}$ as $\mathcal{A}_{0}=\left\{\boldsymbol{e}_{\alpha_{1}}+\cdots+\boldsymbol{e}_{\alpha_{p}} \mid \alpha \in \Lambda\right\}$. The action set $\mathcal{A}$ is the convex hull of $\mathcal{A}_{0}: \mathcal{A}=\operatorname{conv}\left(\mathcal{A}_{0}\right)$. Assume that the ground-truth parameter is $\boldsymbol{M}^{*}=\boldsymbol{M}_{\alpha^{*}} \in \mathcal{M}$. At round $t$, the algorithm chooses an action $\boldsymbol{a}_{t} \in \mathcal{A}$, and gets the noiseless reward $r_{t}=r\left(\boldsymbol{M}^{*}, \boldsymbol{a}_{t}\right)=\left\langle\boldsymbol{M}^{*},\left(\boldsymbol{a}_{t}\right)^{\otimes p}\right\rangle=\prod_{i=1}^{p}\left\langle\boldsymbol{e}_{\alpha_{i}^{*}}, \boldsymbol{a}_{t}\right\rangle$.

## F. 1 Proof for Theorem G. 2

We introduce a lemma showing that if the action set is restricted to the core action set $\mathcal{A}_{0}$, then at least $\left|\mathcal{A}_{0}\right|-1=\binom{d}{p}-1$ actions are needed to identify the ground-truth.
Lemma F.1. If the actions are restricted to $\mathcal{A}_{0}$, then for the noiseless degree-p polynomial bandits, any algorithm needs to play at least $\binom{d}{p}-1$ actions to determine $\boldsymbol{M}^{*}$ in the worst case. Furthermore, the worst-case cumulative regret at round $T$ can be lower bounded by

$$
\mathfrak{R}(T) \geq \min \left\{T,\binom{d}{p}-1\right\}
$$

proof of Lemma F. 1 For any $\alpha$ and $\alpha^{\prime}$, the reward of playing $\boldsymbol{e}_{\alpha_{1}}+\cdots+\boldsymbol{e}_{\alpha_{p}}$ when the ground-truth model is $\boldsymbol{M}_{\alpha}^{\prime}$ is

$$
\begin{aligned}
\left\langle\boldsymbol{M}_{\alpha}^{\prime},\left(\boldsymbol{e}_{\alpha_{1}}+\cdots+\boldsymbol{e}_{\alpha_{p}}\right)^{\otimes p}\right\rangle & =\prod_{i=1}^{p}\left\langle\boldsymbol{e}_{\alpha_{i}^{\prime}}, \boldsymbol{e}_{\alpha_{1}}+\cdots+\boldsymbol{e}_{\alpha_{p}}\right\rangle \\
& =\prod_{i=1}^{p} \mathbb{I}\left\{\alpha_{i}^{\prime} \in \alpha\right\} \\
& = \begin{cases}1, & \text { if } \alpha=\alpha^{\prime} \\
0, & \text { otherwise } .\end{cases}
\end{aligned}
$$

Hence, no matter how the algorithm adaptively chooses the actions, in the worst case $\binom{d}{p}-1$ actions are needed to determine $\boldsymbol{M}^{*}$. Also notice that the reward for $\boldsymbol{e}_{\alpha_{1}}+\cdots+\boldsymbol{e}_{\alpha_{p}}$ is zero if $\alpha \neq \alpha^{*}$. Therefore the regret lower bound follows.

Next, we show that even when the action set is unrestricted, any UCB algorithm fails to explore in an unrestricted way. This is because the optimistic mechanism forbids the algorithm to play an informative action that is known to be low reward for all models in the confidence set. We first recall the definition of UCB algorithms.

UCB Algorithms The UCB algorithms sequentially maintain a confidence set $\mathcal{C}_{t}$ after playing actions $\boldsymbol{a}_{1}, \ldots, \boldsymbol{a}_{t}$. Then UCB algorithms play

$$
\boldsymbol{a}_{t+1} \in \underset{\boldsymbol{a} \in \mathcal{A}}{\arg \max } \mathrm{UCB}_{t}(\boldsymbol{a}),
$$

where

$$
\mathrm{UCB}_{t}(\boldsymbol{a})=\max _{\boldsymbol{M} \in \mathcal{C}_{t}}\left\langle\boldsymbol{M},(\boldsymbol{a})^{\otimes p}\right\rangle
$$

proof of Theorem G.2. We prove that even if the action set is unrestrcited, the optimistic mechanism in the UCB algorithm above forces it to choose actions in the restricted action set $\mathcal{A}_{0}$.
Assume $\boldsymbol{M}^{*}=\boldsymbol{M}_{\alpha^{*}}$. Next we show that for all $\boldsymbol{a} \in \mathcal{A}-\mathcal{A}_{0}$ (where the minus sign should be understood as set difference), we have

$$
\mathrm{UCB}_{t}(\boldsymbol{a})<1
$$

For all $\boldsymbol{a} \in \mathcal{A}$, since $\mathcal{A}=\operatorname{conv}\left(\mathcal{A}_{0}\right)$, we can write

$$
\boldsymbol{a}=\sum_{\alpha \in \Lambda} p_{\alpha}\left(\boldsymbol{e}_{\alpha_{1}}+\cdots+\boldsymbol{e}_{\alpha_{p}}\right)
$$

where $\sum_{\alpha \in \Lambda} p_{\alpha}=1$ and $p_{\alpha} \geq 0$. Therefore,

$$
\begin{aligned}
\mathrm{UCB}_{t}(\boldsymbol{a}) & =\max _{\boldsymbol{M} \in \mathcal{C}_{t}}\left\langle\boldsymbol{M},(\boldsymbol{a})^{\otimes p}\right\rangle \\
& \leq \max _{\boldsymbol{M} \in \mathcal{M}}\left\langle\boldsymbol{M},(\boldsymbol{a})^{\otimes p}\right\rangle \\
& =\max _{\alpha^{\prime}}\left\langle\boldsymbol{M}_{\alpha^{\prime}},(\boldsymbol{a})^{\otimes p}\right\rangle \\
& =\max _{\alpha^{\prime}} \prod_{i=1}^{p}\left\langle\boldsymbol{e}_{\alpha_{i}^{\prime}}, \boldsymbol{a}\right\rangle .
\end{aligned}
$$

Plug in the expression of $\boldsymbol{a}$, we have

$$
\begin{aligned}
\left\langle\boldsymbol{e}_{\alpha_{i}^{\prime}}, \boldsymbol{a}\right\rangle & =\sum_{\alpha} p_{\alpha}\left\langle\boldsymbol{e}_{\alpha_{i}^{\prime}}, \boldsymbol{e}_{\alpha_{1}}+\cdots+\boldsymbol{e}_{\alpha_{p}}\right\rangle \\
& =\sum_{\alpha} p_{\alpha} \mathbb{I}\left\{\alpha_{i}^{\prime} \in \alpha\right\} \\
& \leq \sum_{\alpha} p_{\alpha}=1 .
\end{aligned}
$$

Therefore, for any fixed $\alpha^{\prime}=\left(\alpha_{1}^{\prime}, \ldots, \alpha_{p}^{\prime}\right)$,

$$
\begin{aligned}
\prod_{i=1}^{p}\left\langle\boldsymbol{e}_{\alpha_{i}^{\prime}}, \boldsymbol{a}\right\rangle & =\left(\sum_{\alpha} p_{\alpha} \mathbb{I}\left\{\alpha_{1}^{\prime} \in \alpha\right\}\right) \cdots\left(\sum_{\alpha} p_{\alpha} \mathbb{I}\left\{\alpha_{p}^{\prime} \in \alpha\right\}\right) \\
& \leq 1,
\end{aligned}
$$

where the equality holds if and only if for any $p_{\alpha}>0, \alpha=\alpha^{\prime}$, which is equivalent to $\boldsymbol{a}=$ $\boldsymbol{e}_{\alpha_{1}^{\prime}}+\cdots+\boldsymbol{e}_{\alpha_{p}^{\prime}}$. Therefore, if $\boldsymbol{a} \in \mathcal{A}-\mathcal{A}_{0}$, for any $\alpha^{\prime} \in \Lambda$, we have $\prod_{i=1}^{p}\left\langle\boldsymbol{e}_{\alpha_{i}^{\prime}}, \boldsymbol{a}\right\rangle<1$. This means

$$
\mathrm{UCB}_{t}(\boldsymbol{a})<1
$$

Meanwhile, we can see that for the action $\boldsymbol{a}^{*}=\boldsymbol{e}_{\alpha_{1}^{*}}+\cdots+\boldsymbol{e}_{\alpha_{p}^{*}} \in \mathcal{A}_{0}$,

$$
\begin{aligned}
\mathrm{UCB}_{t}\left(\boldsymbol{a}^{*}\right) & =\max _{\boldsymbol{M} \in \mathcal{C}_{t}}\left\langle\boldsymbol{M},\left(\boldsymbol{a}^{*}\right)^{\otimes p}\right\rangle \\
& \geq\left\langle\boldsymbol{M}^{*},\left(\boldsymbol{a}^{*}\right)^{\otimes p}\right\rangle \\
& =\left\langle\boldsymbol{M}^{*}, \boldsymbol{A}_{\alpha^{*}}\right\rangle=1 .
\end{aligned} \quad\left(\boldsymbol{M}^{*} \in \mathcal{C}_{t}\right)
$$

Therefore, we see that $\left(\mathcal{A}-\mathcal{A}_{0}\right) \cap \arg \max _{\boldsymbol{a} \in \mathcal{A}} \mathrm{UCB}_{t}(\boldsymbol{a})=\varnothing$, which means $\boldsymbol{a}_{t+1} \in \mathcal{A}_{0}$ for all $t \geq 0$. Therefore, by Lemma F.1, the theorem holds.

## F. $2 O(d)$ Actions via Solving Polynomial Equations

Firstly, we verify that the model falls into the category of Example 3.23 with $k=p$. For every $\alpha \in \Lambda$, the reward of playing $\boldsymbol{a}$ when the ground-truth model is $\boldsymbol{M}_{\alpha}$ is

$$
\left\langle\boldsymbol{M}_{\alpha},(\boldsymbol{a})^{\otimes p}\right\rangle=\prod_{i=1}^{p}\left\langle\boldsymbol{e}_{\alpha_{i}}, \boldsymbol{a}\right\rangle
$$

which can be written as $q_{0}\left(\boldsymbol{U}_{\alpha} \boldsymbol{a}\right)$, where $q_{0}\left(x_{1}, \ldots, x_{p}\right)=x_{1} x_{2} \cdots x_{p}$ and $\boldsymbol{U}_{\alpha} \in \mathbb{R}^{p \times d}$ is a matrix with $\boldsymbol{e}_{\alpha_{i}}$ as the $i$-th row.
Secondly, we show that since the ground-truth model is $p$-homogenous, we can extend the action set to $\operatorname{conv}(\mathcal{A}, \mathbf{0})$. This is because for every action of the form $c \boldsymbol{a}$, where $0 \leq c \leq 1$ and $\boldsymbol{a} \in \mathcal{A}$, the reward is $c^{p}$ times the reward at $\boldsymbol{a}$. Therefore, to get the reward at $c \boldsymbol{a}$, we only need to play at $\boldsymbol{a}$ and multiply the reward by $c^{p}$.
Notice that $\operatorname{conv}(\mathcal{A}, \mathbf{0})$ is of positive Lebesgue measure. By Theorem 3.21, we know that only $2\left(d k+(p+1)^{p}\right)=O(d)$ actions are needed to determine the optimal action almost surely.

## G Proof of Section 3.3.2

We present the proof of Theorem 3.18 in the following.

Proof. We overload the notation and use $[d]$ to denote the set $\left\{e_{1}, e_{2}, \ldots, e_{d}\right\}$. The hard instances are chosen in $\Delta \cdot[d]^{p}$, i.e. $\left(\boldsymbol{\theta}_{1}, \ldots, \boldsymbol{\theta}_{p}\right)=\Delta \cdot\left(\widehat{\boldsymbol{\theta}}_{1}, \ldots, \widehat{\boldsymbol{\theta}}_{p}\right)$ where $\left(\widehat{\boldsymbol{\theta}}_{1}, \ldots, \widehat{\boldsymbol{\theta}}_{p}\right) \in[d]^{p}$. For a group of vectors $\boldsymbol{\theta}_{1}, \ldots, \boldsymbol{\theta}_{p} \in[d]$, we use

$$
\operatorname{supp}\left(\boldsymbol{\theta}_{1}, \ldots, \boldsymbol{\theta}_{p}\right):=\left(\max _{i \in[p]}\left(\boldsymbol{\theta}_{i}\right)_{1}, \ldots, \max _{i \in[p]}\left(\boldsymbol{\theta}_{i}\right)_{d}\right) \in\{0,1\}^{d}
$$

to denote the support of these vectors. We use $\boldsymbol{a}^{(t)} \in \mathbb{R}^{d}$ to denote the action in $t$-th episode.
We use $\mathbb{P}_{\left(\boldsymbol{\theta}_{1}, \ldots, \boldsymbol{\theta}_{p}\right)}$ to denote the measure on outcomes induced by the interaction of the fixed policy and the bandit paramterised by $r=\prod_{i=1}^{p}\left(\boldsymbol{\theta}_{i}^{\top} \boldsymbol{a}\right)+\epsilon$. Specifically, We use $\mathbb{P}_{0}$ to denote the measure on outcomes induced by the interaction of the fixed policy and the pure noise bandit $r=\epsilon$.

$$
\begin{aligned}
& \mathfrak{R}(d, p, T) \\
& \geq \frac{1}{d^{p}} \sum_{\left(\boldsymbol{\theta}_{1}, \ldots, \boldsymbol{\theta}_{p}\right) \in \Delta \cdot[d]^{p}} \mathbb{E}_{\left(\boldsymbol{\theta}_{1}, \ldots, \boldsymbol{\theta}_{p}\right)}\left[T \Delta^{p} / p^{p / 2}-\sum_{t=1}^{T} \prod_{i=1}^{p}\left(\boldsymbol{\theta}_{i}^{\top} \boldsymbol{a}^{(t)}\right)\right] \\
& =\frac{\Delta^{p}}{d^{p}} \sum_{\left(\boldsymbol{\theta}_{1}, \ldots, \boldsymbol{\theta}_{p}\right) \in \Delta \cdot[d]^{p}}\left(T / p^{p / 2}-\mathbb{E}_{\left(\boldsymbol{\theta}_{1}, \ldots, \boldsymbol{\theta}_{p}\right)}\left[\sum_{t=1}^{T} \prod_{i=1}^{p}\left(\widehat{\boldsymbol{\theta}}_{i}^{\top} \boldsymbol{a}^{(t)}\right)\right]\right) \\
& \geq \frac{\Delta^{p}}{d^{p}} \sum_{\left(\boldsymbol{\theta}_{1}, \ldots, \boldsymbol{\theta}_{p}\right) \in \Delta \cdot[d]^{p}}\left(T / p^{p / 2}-\mathbb{E}_{0}\left[\sum_{t=1}^{T} \prod_{i=1}^{p}\left(\widehat{\boldsymbol{\theta}}_{i}^{\top} \boldsymbol{a}^{(t)}\right)\right]-T\left\|\mathbb{P}_{0}-\mathbb{P}_{\left(\boldsymbol{\theta}_{1}, \ldots, \boldsymbol{\theta}_{p}\right)}\right\|_{\mathrm{TV}}\right) \\
& \geq \frac{\Delta^{p}}{d^{p}} \sum_{\left(\boldsymbol{\theta}_{1}, \ldots, \boldsymbol{\theta}_{p}\right) \in \Delta \cdot[d]^{p}}\left(T / p^{p / 2}-\mathbb{E}_{0}\left[\sum_{t=1}^{T} \prod_{i=1}^{p}\left(\hat{\boldsymbol{\theta}}_{i}^{\top} \boldsymbol{a}^{(t)}\right)\right]-T \sqrt{D_{\mathrm{KL}}\left(\mathbb{P}_{0} \| \mathbb{P}_{\left(\boldsymbol{\theta}_{1}, \ldots, \boldsymbol{\theta}_{p}\right)}\right)}\right) \\
& \geq \frac{\Delta^{p}}{d^{p}} \sum_{\left(\boldsymbol{\theta}_{1}, \ldots, \boldsymbol{\theta}_{p}\right) \in \Delta \cdot[d]^{p}}\left(T / p^{p / 2}-\mathbb{E}_{0}\left[\sum_{t=1}^{T} \prod_{i=1}^{p}\left(\hat{\boldsymbol{\theta}}_{i}^{\top} \boldsymbol{a}^{(t)}\right)\right]-T \sqrt{\Delta^{2 p} \mathbb{E}_{0}\left[\sum_{t=1}^{T} \prod_{i=1}^{p}\left(\hat{\boldsymbol{\theta}}_{i}^{\top} \boldsymbol{a}^{(t)}\right)^{2}\right]}\right) \\
& \geq \frac{\Delta^{p}}{d^{p}}\left(\frac{d^{p} T}{p^{p}}-\mathbb{E}_{0}\left[\sum_{\left(\widehat{\boldsymbol{\theta}}_{1}, \ldots, \widehat{\boldsymbol{\theta}}_{p}\right) \in[d]^{p}} \sum_{t=1}^{T} \prod_{i=1}^{p}\left(\widehat{\boldsymbol{\theta}}_{i}^{\top} \boldsymbol{a}^{(t)}\right)\right]-T d^{\frac{p}{2}} \Delta^{p} \sqrt{\mathbb{E}_{0}\left[\sum_{\left.\left(\widehat{\boldsymbol{\theta}}_{1}, \ldots, \widehat{\boldsymbol{\theta}}_{p}\right) \in[d]\right]^{p}} \sum_{t=1}^{T} \prod_{i=1}^{p}\left(\widehat{\boldsymbol{\theta}}_{i}^{\top} \boldsymbol{a}^{(t)}\right)^{2}\right]}\right)
\end{aligned}
$$

where the first step comes from

$$
\text { Regret } \geq \mathbb{E}_{\left(\boldsymbol{\theta}_{1}, \ldots, \boldsymbol{\theta}_{p}\right)}\left[T \Delta^{p} / p^{p / 2}-\sum_{t=1}^{T} \prod_{i=1}^{p}\left(\boldsymbol{\theta}_{i}^{\top} \boldsymbol{a}^{(t)}\right)\right]
$$

(the optimal action in hindsight is $\boldsymbol{a}=\operatorname{supp}\left(\boldsymbol{\theta}_{1}, \ldots, \boldsymbol{\theta}_{p}\right) / \sqrt{p}$ ); the second step comes from $\left(\boldsymbol{\theta}_{1}, \ldots, \boldsymbol{\theta}_{p}\right)=\Delta \cdot\left(\widehat{\boldsymbol{\theta}}_{1}, \ldots, \widehat{\boldsymbol{\theta}}_{p}\right)$ and algebra; the third step comes from $\left|\sum_{t=1}^{T} \prod_{i=1}^{p}\left(\widehat{\boldsymbol{\theta}}_{i}^{\top} \boldsymbol{a}^{(t)}\right)\right| \leq T ;$ the fourth step comes from Pinsker's inequality; the fifth step comes from

$$
\begin{aligned}
D_{\mathrm{KL}}\left(\mathbb{P}_{0} \| \mathbb{P}_{\boldsymbol{\theta}_{1}, \ldots, \boldsymbol{\theta}_{p}}\right) & =\mathbb{E}_{0}\left[\sum_{t=1}^{T} D_{\mathrm{KL}}\left(N(0,1) \| N\left(\prod_{i=1}^{p}\left(\boldsymbol{\theta}_{i}^{\top} \boldsymbol{a}^{(t)}\right), 1\right)\right)\right] \\
& =\Delta^{2 p} \mathbb{E}_{0}\left[\sum_{t=1}^{T} \prod_{i=1}^{p}\left(\hat{\boldsymbol{\theta}}_{i}^{\top} \boldsymbol{a}^{(t)}\right)^{2}\right]
\end{aligned}
$$

and the final step comes from Jensen's inequality and algebra.
Notice that

$$
\begin{aligned}
\mathbb{E}_{0}\left[\sum_{\left(\widehat{\boldsymbol{\theta}}_{1}, \ldots, \widehat{\boldsymbol{\theta}}_{p}\right) \in[d]^{p}} \sum_{t=1}^{T} \prod_{i=1}^{p}\left(\widehat{\boldsymbol{\theta}}_{i}^{\top} \boldsymbol{a}^{(t)}\right)\right] & =\mathbb{E}_{0}\left[\sum_{\left(j_{1}, \ldots, j_{p}\right) \in[d]^{p}} \sum_{t=1}^{T} \prod_{i=1}^{p}\left(\boldsymbol{a}_{j_{i}}^{(t)}\right)\right] \\
& =\mathbb{E}_{0}\left[\sum_{t=1}^{T} \prod_{i=1}^{p}\left(\sum_{j=1}^{d} \boldsymbol{a}_{j}^{(t)}\right)\right] \\
& \leq \mathbb{E}_{0}\left[\sum_{t=1}^{T} \prod_{i=1}^{p}\left\|\boldsymbol{a}^{(t)}\right\|_{1}\right] \\
& \leq d^{p / 2} T
\end{aligned}
$$

and

$$
\begin{aligned}
\mathbb{E}_{0}\left[\sum_{\left(\widehat{\boldsymbol{\theta}}_{1}, \ldots, \widehat{\boldsymbol{\theta}}_{p}\right) \in[d]^{p}} \sum_{t=1}^{T} \prod_{i=1}^{p}\left(\widehat{\boldsymbol{\theta}}_{i}^{\top} \boldsymbol{a}^{(t)}\right)^{2}\right] & =\mathbb{E}_{0}\left[\sum_{\left(j_{1}, \ldots, j_{p}\right) \in[d]^{p}} \sum_{t=1}^{T} \prod_{i=1}^{p}\left(\boldsymbol{a}_{j_{i}}^{(t)}\right)^{2}\right] \\
& =\mathbb{E}_{0}\left[\sum_{t=1}^{T} \prod_{i=1}^{p}\left\|\boldsymbol{a}^{(t)}\right\|_{2}^{2}\right] \\
& \leq T
\end{aligned}
$$

where we used $\left\|\boldsymbol{a}^{(t)}\right\|_{2} \leq 1, \forall t \in[T]$. Therefore plugging back we have

$$
\mathfrak{R}(d, p, T) \geq \frac{\Delta^{p}}{d^{p}}\left(\frac{d^{p} T}{p^{\frac{p}{2}}}-d^{\frac{p}{2}} T-T d^{\frac{p}{2}} \Delta^{p} \sqrt{T}\right)
$$

and finally letting $\Delta^{p}=\sqrt{\frac{d^{p}}{4 T p^{p}}}$ leads to

$$
\mathfrak{R}(d, p, T) \geq O\left(\sqrt{d^{p} T} / p^{p}\right)
$$

Remark G.1. Better result $O\left(\sqrt{d^{p} T}\right)$ holds for bandits $r=\prod_{i=1}^{p}\left(\boldsymbol{\theta}_{i}^{\top} \boldsymbol{a}_{i}\right)+\epsilon$ where $\boldsymbol{a}_{i} \in$ $\mathbb{R}^{d},\left\|\boldsymbol{a}_{i}\right\|_{2} \leq 1$.

For completeness, we show the proof of the above remark.
Proof. We overload the notation and use $[d]$ to denote the set $\left\{e_{1}, e_{2}, \ldots, e_{d}\right\}$. The hard instances are chosen in $\Delta \cdot[d]^{p}$, i.e. $\left(\boldsymbol{\theta}_{1}, \ldots, \boldsymbol{\theta}_{p}\right)=\Delta \cdot\left(\widehat{\boldsymbol{\theta}}_{1}, \ldots, \widehat{\boldsymbol{\theta}}_{p}\right)$ where $\left(\widehat{\boldsymbol{\theta}}_{1}, \ldots, \widehat{\boldsymbol{\theta}}_{p}\right) \in[d]^{p}$. We use $\boldsymbol{a}_{i}^{(t)} \in \mathbb{R}^{d}$ to denote the $i$-th action in $t$-th episode, where $i \in[p], t \in[T]$.

We use $\mathbb{P}_{\left(\boldsymbol{\theta}_{1}, \ldots, \boldsymbol{\theta}_{p}\right)}$ to indicate the measure on outcomes induced by the interaction of the fixed policy and the bandit paramterised by $r=\prod_{i=1}^{p}\left(\boldsymbol{\theta}_{i}^{\top} \boldsymbol{a}_{i}\right)+\epsilon$. Specifically, We use $\mathbb{P}_{0}$ to indicate the measure on outcomes induced by the interaction of the fixed policy and the pure noise bandit $r=\epsilon$.

$$
\begin{aligned}
& \mathfrak{R}(d, p, T) \\
& \geq \frac{1}{d^{p}} \sum_{\left(\boldsymbol{\theta}_{1}, \ldots, \boldsymbol{\theta}_{p}\right) \in \Delta \cdot[d]^{p}} \mathbb{E}_{\left(\boldsymbol{\theta}_{1}, \ldots, \boldsymbol{\theta}_{p}\right)}\left[T \Delta^{p}-\sum_{t=1}^{T} \prod_{i=1}^{p}\left(\boldsymbol{\theta}_{i}^{\top} \boldsymbol{a}_{i}^{(t)}\right)\right] \\
& =\frac{\Delta^{p}}{d^{p}} \sum_{\left(\boldsymbol{\theta}_{1}, \ldots, \boldsymbol{\theta}_{p}\right) \in \Delta \cdot[d]^{p}}\left(T-\mathbb{E}_{\left(\boldsymbol{\theta}_{1}, \ldots, \boldsymbol{\theta}_{p}\right)}\left[\sum_{t=1}^{T} \prod_{i=1}^{p}\left(\widehat{\boldsymbol{\theta}}_{i}^{\top} \boldsymbol{a}_{i}^{(t)}\right)\right]\right) \\
& \geq \frac{\Delta^{p}}{d^{p}} \sum_{\left(\boldsymbol{\theta}_{1}, \ldots, \boldsymbol{\theta}_{p}\right) \in \Delta \cdot[d]^{p}}\left(T-\mathbb{E}_{0}\left[\sum_{t=1}^{T} \prod_{i=1}^{p}\left(\widehat{\boldsymbol{\theta}}_{i}^{\top} \boldsymbol{a}_{i}^{(t)}\right)\right]-T\left\|\mathbb{P}_{0}-\mathbb{P}_{\left(\boldsymbol{\theta}_{1}, \ldots, \boldsymbol{\theta}_{p}\right)}\right\|_{\mathrm{TV}}\right) \\
& \geq \frac{\Delta^{p}}{d^{p}} \sum_{\left(\boldsymbol{\theta}_{1}, \ldots, \boldsymbol{\theta}_{p}\right) \in \Delta \cdot[d]^{p}}\left(T-\mathbb{E}_{0}\left[\sum_{t=1}^{T} \prod_{i=1}^{p}\left(\widehat{\boldsymbol{\theta}}_{i}^{\top} \boldsymbol{a}_{i}^{(t)}\right)\right]-T \sqrt{D_{\mathrm{KL}}\left(\mathbb{P}_{0} \| \mathbb{P}_{\left(\boldsymbol{\theta}_{1}, \ldots, \boldsymbol{\theta}_{p}\right)}\right)}\right) \\
& \geq \frac{\Delta^{p}}{d^{p}} \sum_{\left(\boldsymbol{\theta}_{1}, \ldots, \boldsymbol{\theta}_{p}\right) \in \Delta \cdot[d]^{p}}\left(T-\mathbb{E}_{0}\left[\sum_{t=1}^{T} \prod_{i=1}^{p}\left(\hat{\boldsymbol{\theta}}_{i}^{\top} \boldsymbol{a}_{i}^{(t)}\right)\right]-T \sqrt{\Delta^{2 p} \mathbb{E}_{0}\left[\sum_{t=1}^{T} \prod_{i=1}^{p}\left(\hat{\boldsymbol{\theta}}_{i}^{\top} \boldsymbol{a}_{i}^{(t)}\right)^{2}\right]}\right) \\
& \geq \frac{\Delta^{p}}{d^{p}}\left(d^{p} T-\mathbb{E}_{0}\left[\sum_{\left(\widehat{\boldsymbol{\theta}}_{1}, \ldots, \widehat{\boldsymbol{\theta}}_{p}\right) \in[d]^{p}} \sum_{t=1}^{T} \prod_{i=1}^{p}\left(\widehat{\boldsymbol{\theta}}_{i}^{\top} \boldsymbol{a}_{i}^{(t)}\right)\right]-T d^{\frac{p}{2}} \Delta^{p} \sqrt{\mathbb{E}_{0}\left[\sum_{\left(\widehat{\boldsymbol{\theta}}_{1}, \ldots, \widehat{\boldsymbol{\theta}}_{p}\right) \in[d]^{p}} \sum_{t=1}^{T} \prod_{i=1}^{p}\left(\widehat{\boldsymbol{\theta}}_{i}^{\top} \boldsymbol{a}_{i}^{(t)}\right)^{2}\right]}\right)
\end{aligned}
$$

where the first step comes from

$$
\text { Regret } \geq \mathbb{E}_{\left(\boldsymbol{\theta}_{1}, \ldots, \boldsymbol{\theta}_{p}\right)}\left[T \Delta^{p}-\sum_{t=1}^{T} \prod_{i=1}^{p}\left(\boldsymbol{\theta}_{i}^{\top} \boldsymbol{a}_{i}^{(t)}\right)\right]
$$

(the optimal action in hindsight is $\boldsymbol{a}_{i}=\widehat{\boldsymbol{\theta}}_{i}$ ); the second step comes from $\left(\boldsymbol{\theta}_{1}, \ldots, \boldsymbol{\theta}_{p}\right)=\Delta$. $\left(\widehat{\boldsymbol{\theta}}_{1}, \ldots, \widehat{\boldsymbol{\theta}}_{p}\right)$ and algebra; the third step comes from $\left|\sum_{t=1}^{T} \prod_{i=1}^{p}\left(\widehat{\boldsymbol{\theta}}_{i}^{\top} \boldsymbol{a}_{i}^{(t)}\right)\right| \leq T$; the fourth step comes from Pinsker's inequality; the fifth step comes from

$$
\begin{aligned}
D_{\mathrm{KL}}\left(\mathbb{P}_{0} \| \mathbb{P}_{\boldsymbol{\theta}_{1}, \ldots, \boldsymbol{\theta}_{p}}\right) & =\mathbb{E}_{0}\left[\sum_{t=1}^{T} D_{\mathrm{KL}}\left(N(0,1) \| N\left(\prod_{i=1}^{p}\left(\boldsymbol{\theta}_{i}^{\top} \boldsymbol{a}_{i}^{(t)}\right), 1\right)\right)\right] \\
& =\Delta^{2 p} \mathbb{E}_{0}\left[\sum_{t=1}^{T} \prod_{i=1}^{p}\left(\widehat{\boldsymbol{\theta}}_{i}^{\top} \boldsymbol{a}_{i}^{(t)}\right)^{2}\right]
\end{aligned}
$$

and the final step comes from Jensen's inequality and algebra.
Notice that

$$
\begin{aligned}
\mathbb{E}_{0}\left[\sum_{\left(\widehat{\boldsymbol{\theta}}_{1}, \ldots, \widehat{\boldsymbol{\theta}}_{p}\right) \in[d]^{p}} \sum_{t=1}^{T} \prod_{i=1}^{p}\left(\hat{\boldsymbol{\theta}}_{i}^{\top} \boldsymbol{a}_{i}^{(t)}\right)\right] & =\mathbb{E}_{0}\left[\sum_{\left(j_{1}, \ldots, j_{p}\right) \in[d]^{p}} \sum_{t=1}^{T} \prod_{i=1}^{p}\left(\left(\boldsymbol{a}_{i}^{(t)}\right)_{j_{i}}\right)\right] \\
& =\mathbb{E}_{0}\left[\sum_{t=1}^{T} \prod_{i=1}^{p}\left(\sum_{j=1}^{d}\left(\boldsymbol{a}_{i}^{(t)}\right)_{j}\right)\right] \\
& \leq \mathbb{E}_{0}\left[\sum_{t=1}^{T} \prod_{i=1}^{p}\left\|\boldsymbol{a}_{i}^{(t)}\right\|_{1}\right] \\
& \leq d^{p / 2} T
\end{aligned}
$$

and

$$
\begin{aligned}
\mathbb{E}_{0}\left[\sum_{\left(\hat{\boldsymbol{\theta}}_{1}, \ldots, \widehat{\boldsymbol{\theta}}_{p}\right) \in[d]^{p}} \sum_{t=1}^{T} \prod_{i=1}^{p}\left(\hat{\boldsymbol{\theta}}_{i}^{\top} \boldsymbol{a}_{i}^{(t)}\right)^{2}\right] & =\mathbb{E}_{0}\left[\sum_{\left(j_{1}, \ldots, j_{p}\right) \in[d]^{p}} \sum_{t=1}^{T} \prod_{i=1}^{p}\left(\left(\boldsymbol{a}_{i}^{(t)}\right)_{j_{i}}\right)^{2}\right] \\
& =\mathbb{E}_{0}\left[\sum_{t=1}^{T} \prod_{i=1}^{p}\left\|\boldsymbol{a}_{i}^{(t)}\right\|_{2}^{2}\right] \\
& \leq T
\end{aligned}
$$

where we used $\left\|\boldsymbol{a}_{i}^{(t)}\right\|_{2} \leq 1, \forall t \in[T]$. Therefore plugging back we have

$$
\mathfrak{R}(d, p, T) \geq \frac{\Delta^{p}}{d^{p}}\left(d^{p} T-d^{\frac{p}{2}} T-T d^{\frac{p}{2}} \Delta^{p} \sqrt{T}\right)
$$

and finally letting $\Delta^{p}=\sqrt{\frac{d^{p}}{4 T}}$ leads to

$$
\mathfrak{R}(d, p, T) \geq O\left(\sqrt{d^{p} T}\right)
$$

We present the proof of Theorem 3.19 in the following.
Proof. Denote the optimal action in hindsight as $\boldsymbol{a}^{*}=\operatorname{supp}\left(\boldsymbol{\theta}_{1}, \ldots, \boldsymbol{\theta}_{p}\right) / \sqrt{p}$. From the proof of Theorem 3.18 we know that if $T \leq \frac{1}{4 p^{p}} \cdot \frac{d^{p}}{\Delta^{2 p}}$, then

$$
\begin{aligned}
& \frac{1}{d^{p}} \sum_{\left(\boldsymbol{\theta}_{1}, \ldots, \boldsymbol{\theta}_{p}\right) \in \Delta \cdot[d]^{p}} \mathbb{E}_{\left(\boldsymbol{\theta}_{1}, \ldots, \boldsymbol{\theta}_{p}\right)}\left[\prod_{i=1}^{p}\left(\boldsymbol{\theta}_{i}^{\top} \boldsymbol{a}^{*}\right)-\prod_{i=1}^{p}\left(\boldsymbol{\theta}_{i}^{\top} \boldsymbol{a}^{(t)}\right)\right] \\
\geq & \frac{\Delta^{p}}{d^{p}}\left(\frac{d^{p}}{p^{\frac{p}{2}}}-d^{\frac{p}{2}}-d^{\frac{p}{2}} \Delta^{p} \sqrt{T}\right) \\
\geq & \frac{\Delta^{p}}{4 p^{\frac{p}{2}}} \\
\geq & \frac{1}{4} \cdot \frac{1}{d^{p}} \sum_{\left(\boldsymbol{\theta}_{1}, \ldots, \boldsymbol{\theta}_{p}\right) \in \Delta \cdot[d]^{p}} \mathbb{E}_{\left(\boldsymbol{\theta}_{\boldsymbol{1}}, \ldots, \boldsymbol{\theta}_{p}\right)}\left[\prod_{i=1}^{p}\left(\boldsymbol{\theta}_{i}^{\top} \boldsymbol{a}^{*}\right)\right]
\end{aligned}
$$

which indicates the following

$$
\inf _{\pi} \sup _{\left(\boldsymbol{\theta}_{1}, \ldots, \boldsymbol{\theta}_{p}\right)} \mathbb{E}_{\left(\boldsymbol{\theta}_{1}, \ldots, \boldsymbol{\theta}_{p}\right)}\left[\frac{3}{4} \cdot \prod_{i=1}^{p}\left(\boldsymbol{\theta}_{i}^{\top} \boldsymbol{a}^{*}\right)-\prod_{i=1}^{p}\left(\boldsymbol{\theta}_{i}^{\top} \boldsymbol{a}^{(t)}\right)\right] \geq 0 .
$$

## G.0.1 Lower Bounds with UCB Algorithms

In this subsection, we construct a hard bandit problem where the rewards are noiseless degree-p polynomial, and show that any UCB algorithm needs at least $\Omega\left(d^{p}\right)$ actions to learn the optimal action. On the contrary, Theorem 3.21 shows that by playing actions randomly, we only need $2\left(d k+(p+1)^{p}\right)=O(d)$ actions.

Hard Case Construction Let $\boldsymbol{e}_{i}$ denotes the $i$-th standard orthonormal basis of $\mathbb{R}^{d}$, i.e., $\boldsymbol{e}_{i}$ has only one 1 at the $i$-th entry and 0 's for other entries. We define a $p$-th multi-indices set $\Lambda$ as $\Lambda=\left\{\left(\alpha_{1}, \ldots, \alpha_{p}\right) \mid 1 \leq \alpha_{1}<\cdots<\alpha_{p} \leq d\right\}$. For an $\alpha=\left(\alpha_{1}, \ldots, \alpha_{p}\right) \in \Lambda$, denote $\boldsymbol{M}_{\alpha}=$ $\boldsymbol{e}_{\alpha_{1}} \otimes \cdots \otimes \boldsymbol{e}_{\alpha_{p}}$. Then the model space $\mathcal{M}$ is defined as $\mathcal{M}=\left\{\boldsymbol{M}_{\alpha} \mid \alpha \in \Lambda\right\}$, which is a subset of rank$1 p$-th order tensors. The action set $\mathcal{A}$ is defined as $\mathcal{A}=\operatorname{conv}\left(\left\{\boldsymbol{e}_{\alpha_{1}}+\cdots+\boldsymbol{e}_{\alpha_{p}} \mid \alpha \in \Lambda\right\}\right)$. Assume that the ground-truth parameter is $\boldsymbol{M}^{*}=\boldsymbol{M}_{\alpha^{*}} \in \mathcal{M}$. The noiseless reward $r_{t}=r\left(\boldsymbol{M}^{*}, \boldsymbol{a}_{t}\right)=$ $\left\langle\boldsymbol{M}^{*},\left(\boldsymbol{a}_{t}\right)^{\otimes p}\right\rangle=\prod_{i=1}^{p}\left\langle\boldsymbol{e}_{\alpha_{i}^{*}}, \boldsymbol{a}_{t}\right\rangle$ is a polynomial of $\boldsymbol{a}_{t}$ and falls into the case of Example 3.23

UCB Algorithms The UCB algorithms sequentially maintain a confidence set $\mathcal{C}_{t}$ after playing actions $\boldsymbol{a}_{1}, \ldots, \boldsymbol{a}_{t}$. Then UCB algorithms play $\boldsymbol{a}_{t+1} \in \arg \max _{\boldsymbol{a} \in \mathcal{A}} \mathrm{UCB}_{t}(\boldsymbol{a})$, where $\mathrm{UCB}_{t}(\boldsymbol{a})=$ $\max _{\boldsymbol{M} \in \mathcal{C}_{t}}\left\langle\boldsymbol{M},(\boldsymbol{a})^{\otimes p}\right\rangle$.
Theorem G.2. Assume that for each $t \geq 0$, the confidence set $\mathcal{C}_{t}$ contains the ground-truth model, i.e., $\boldsymbol{M}^{*} \in \mathcal{C}_{t}$. Then for the noiseless degree-p polynomial bandits, any UCB algorithm needs to play at least $\binom{d}{p}-1$ actions to distinguish models in $\mathcal{M}$. Furthermore, the worst-case cumulative regret at round $T$ can be lower bounded by

$$
\mathfrak{R}(T) \geq \min \left\{T,\binom{d}{p}-1\right\}
$$

Theorem G. 2 shows the failure of the optimistic mechanism, which forbids the algorithm to play an informative action that is known to be of low reward for all models in the confidence set. On the contrary, the reward function class falls into the form of $q(\boldsymbol{U a})$, therefore, by playing actions randomly ${ }^{6}$ we only need $O(d)$ actions as Theorem 3.21 suggests.

[^0]
[^0]:    ${ }^{6}$ Careful readers may notice that $\mathcal{A}$ is of measure zero in this setting. However, since the reward function is a homogenous polynomial of degree $p$, we can actually obtain the rewards on $\operatorname{conv}(\mathcal{A}, \mathbf{0})$, which is of positive measure.

