# Multi-User Reinforcement Learning with Low Rank Rewards 

Dheeraj Nagaraj ${ }^{1}$ Suhas S Kowshik ${ }^{23}$ Naman Agarwal ${ }^{4}$ Praneeth Netrapalli ${ }^{1}$ Prateek Jain ${ }^{1}$


#### Abstract

We consider collaborative multi-user reinforcement learning, where multiple users have the same state-action space and transition probabilities but different rewards. Under the assumption that the reward matrix of the $N$ users has a low-rank structure - a standard and practically successful assumption in the collaborative filtering setting - we design algorithms with significantly lower sample complexity compared to the ones that learn the MDP individually for each user. Our main contribution is an algorithm which explores rewards collaboratively with $N$ user-specific MDPs and can learn rewards efficiently in two key settings: tabular MDPs and linear MDPs. When $N$ is large and the rank is constant, the sample complexity per MDP depends logarithmically over the size of the state-space, which represents an exponential reduction (in the state-space size) when compared to the standard "non-collaborative" algorithms. Our main technical contribution is a method to construct policies which obtain data such that low rank matrix completion is possible (without a generative model). This goes beyond the regular RL framework and is closely related to mean field limits of multi-agent RL.


## 1. Introduction

Reinforcement learning (RL) has recently seen tremendous empirical and theoretical success (Mnih et al., 2015; Sutton et al., 1992; Jin et al. 2020b, Gheshlaghi Azar et al., 2013; Dann \& Brunskill, 2015). Near optimal algorithms have been proposed to explore and learn a given MDP with sample access to trajectories. Multi-agent RL, where multiple agents interact among themselves and the environment to collect rewards, has gained a lot of interest due to immense

[^0]practical applications. However, even simple instances of multi-agent RL, like restless bandits, can be provably hard (Papadimitriou \& Tsitsiklis, 1999).

In this work, we consider the problem of learning optimal policies for multiple MDPs collaboratively so that the total number of trajectories sampled per MDP is smaller than the number of trajectories required to learn them individually. We assume that the various users have the same transition matrices, but different rewards and the rewards have a low rank structure. This is closely related to mean-field limits of certain instances of multi-agent RL as described in Section 3 .

Motivation From the point of view of RL, this is an instance of multi-task reinforcement learning (MTRL), various versions of which have been considered in the literature (Brunskill \& Li, 2013, D’Eramo et al., 2020; Teh et al., 2017, Hessel et al., 2019, Lazaric, 2012). Here, an agent learns different MDPs together with certain common structures. This shared structure could be a common domain (such as moving towards different target points, but in the same environment) or similarity in the task to be performed (such as picking up and moving different kinds of objects). Our problem setup falls in the former category.

Recently, collaborative filtering has been studied in the online learning setting (Bresler \& Karzand, 2021; Jain \& Pal, 2022; Ariu et al. 2020, Huleihel et al. 2021; Nguyen-Thanh et al., 2019), where multiple bandit instances are simultaneously explored under low rank assumptions in order to learn the preferences of multiple users simultaneously. From this point of view, our work adds temporal dynamics, such as change in preferences over time, based on past actions. That is, we consider non-stationary environments via Markov Decision Processes.

To motivate our setup, we consider the example of recommendation systems in e-commerce or video streaming. In this context, classical collaborative filtering setup assumes that the preferences are static and not influenced by the recommendation system itself. However, this is untrue in the real world. Buying an item changes the preferences of the customer. In case the system recommends a TV and the customer buys a TV, a TV stand might be the most relevant recommendation. In case they watch a recommended video about astrophysics, they might want to watch other astro-
physics videos because the first video piqued their interest in the topic.

Our work captures such scenarios by formulating this as an RL problem where the state is affected by the actions of the recommendation system and the state in turn affects the rewards (i.e, the user preferences). This idea has gained traction recently, as shown by the survey paper (Afsar et al., 2022). However, in the papers which have been discussed in the survey, the user information is apriori encoded into the state embedding and the resulting system is treated as a single agent MDP. In this work, we bring forth the multiagent aspect by viewing the RL approach as an enhancement of the classical collaborative filtering allowing us discover similarity among users. We believe the theoretical insights gained in such a setting can lead to Deep RL algorithms which utilize this multi-agent structure effectively and explicitly. The assumption of a common transition matrix can be relaxed in practice by clustering users based on side information and modeling each cluster to have a common transition matrix (see (Mate et al., 2022) and references therein).

Our Contributions We introduce the setting of multi-user collaborative reinforcement learning in the case of tabular and linear MDPs. In our study, we isolate and overcome several technical and conceptual challenges in order achieve sample efficient learning. The main technical challenge we encounter is obtaining the right distribution of stateaction pairs from users such that we can successfully run low-rank matrix estimation algorithms, without access to a generative model (i.e, we can only deploy policies and query trajectories corresponding to this policy). This requires clever algorithm design since some states can be hard to even reach. In fact, this endeavor goes beyond standard RL methods and is related to functional reward maximization and mean field limits of multi-agent RL as explained in Section 3. To summarize our contributions:
a) Improved Sample Complexity: We provide sample efficient algorithms for both these scenarios without access to a generative model. Under the low rank assumption on the reward matrix, the total sample complexity required to learn the near-optimal policies for every user scales as $\tilde{O}\left(N_{\tilde{O}}+|\mathcal{S} \| \mathcal{A}|\right)$ instead of $O(N|\mathcal{S}||\mathcal{A}|)$ for tabular MDPs and $\tilde{O}(N+d)$ instead of $O\left(N d^{2}\right)$ for linear MDPs.
b) Collaborative Exploration: In order to learn the rewards of all the users efficiently under the low-rank assumption, we need to deploy standard low rank matrix estimation algorithms. These require specific kinds of linear measurements (See Section 1.1). Without access to a generative model, the main challenge in this setting is to obtain these linear measurements by querying trajectories of carefully designed policies. We design such algorithms in Section 4
c) Functional Reward Maximization: In the case of linear MDPs, matrix completion is more challenging since we observe measurements of the form $e_{i}^{\top} \Theta^{*} \psi$ where $\Theta^{*} \in$ $\mathbb{R}^{N \times d}$, corresponding to the reward obtained by user $i$, with respect to an embedding $\psi$. Estimating $\Theta^{*}$ under low rank assumptions requires the distribution of $\psi$ to have certain isotropy properties (see Section 7). We design a procedure which can sample-efficiently estimate policies which lead to these isotropic measurements (Section 6).
d) Matrix Completion With Row-Wise Linear Measurements: For the linear MDP setting, the low rank matrix estimation problem lies somewhere in between the matrix completion (Recht, 2011; Jain et al., 2013) and matrix estimation with restricted strong convexity (Negahban et al., 2009). We give a novel active learning based algorithm where we estimate $\Theta^{*}$ row by row without any assumptions like incoherence. This algorithm maybe of independent interest. This is described in Section 7 .

### 1.1. Related Works

Related Settings: Multi-task Reinforcement learning has been studied empirically and theoretically (Brunskill \& Li 2013, Taylor \& Stone, 2009, D’Eramo et al., 2020, Teh et al. 2017, Hessel et al. 2019, Sodhani et al., 2021). (Modi et al. 2017) considers learning a sequence of MDPs with side information, where the parameters of the MDP varies smoothly with the context. (Shah et al. 2020) assumes the optimal Q function $Q^{*}(s, a)$, when represented as a $\mathcal{S} \times \mathcal{A}$ matrix, has low rank. With a generative model, they obtain algorithms which makes use of this structure to obtain a smaller sample complexity whenever the discount factor is bounded by a constant. (Sam et al. 2022) improves the results in this setting with additional assumptions on the transition matrices. Our setting is different in that we consider multiple users, and do not assume a generative model. Our main contribution is to efficiently obtain measurements conducive to matrix completion. (Hu et al., 2021) considers a multi-task RL problem with linear function approximation similar to our setting, but with the assumption of low-rank Bellman closure, where the application of the Bellman operator retains the low rank structure. They obtain a bound depending on the quantity $N \sqrt{d}$ instead of $(N+d)$ like in our work. (Lei \& Li, 2019) RL with low rank assumptions in an experimental context.
Low Rank Matrix Estimation:The low rank assumption is popular in the collaborative filtering literature and has been deployed successfully in a variety of tasks (Bell \& Koren, 2007, Gleich \& Lim, 2011; Hsieh et al., 2012). Low rank matrix estimation has been extensively studied in the statistics and ML community for decades in the context of supervised learning (Candès \& Tao, 2010; Negahban \& Wainwright, 2011; Fazel, 2002; Chen et al., 2019; Jain
et al. 2013, 2017; Recht, 2011; Chen et al., 2020, Chi et al., 2019) in multi-user collaborative filtering settings. The basic question is to estimate a $d_{1} \times d_{2}$ matrix $M$ given linear measurements $\left(x_{i}^{\top} M y_{i}\right)_{i=1}^{n}$ when the number of samples is much smaller than $d_{1} \times d_{2}$ using the assumption that $M$ has low rank.
a) Matrix Completion: $x_{i}$ and $y_{i}$ are standard basis vectors. Typically $x_{i}$ and $y_{i}$ are picked uniformly at random and recovery guarantees are given whenever the matrix $M$ is incoherent (Recht, 2011).
b) Matrix Estimation: $x_{i}$ and $y_{i}$ are not restricted to be standard basis vectors. Typically, they are chosen i.i.d such that the restricted strong convexity holds Negahban et al., 2009).

For the case of tabular MDPs, we use the matrix completion setting and for the case of linear MDPs, our setting lies some where in between settings $a$ ) and $b$ ) as explained above.

### 1.2. Notation

By $\|\cdot\|$ we denote the Euclidean norm and by $e_{1}, \ldots, e_{m}$ the standard basis vectors of the space $\mathbb{R}^{m}$ for some $m \in \mathbb{N}$. Let $\mathcal{S}^{d-1}:=\left\{x \in \mathbb{R}^{d}:\|x\|=1\right\}, \mathcal{B}_{d}(r):=\left\{x \in \mathbb{R}^{d}:\right.$ $\|x\| \leq r\}$. For any $m \times n$ matrix $A$ and a set $\Omega \subseteq[n]$ by $A^{\Omega}$, we denote the sub-matrix of $A$ where the columns corresponding to $\Omega^{\complement}$ are deleted. By $\Delta(\mathcal{A})$, we denote the set of all Borel probability measures over the set $\mathcal{A}$. In the sequel,

## 2. Problem Setting

We consider $N$ users indexed by $[N]$, each of them associated with an MDP with the same state-space $\mathcal{S}$, action space $\mathcal{A}$, horizon $H$ and transition matrices $\mathcal{P}=\left(P_{1}, \ldots, P_{H-1}\right)$. Here $P_{h}\left(\cdot \mid s_{h}, a_{h}\right)$ is a probability measure over $\mathcal{S}$, which gives the distribution of the state at time $h+1$ given the action $a_{h} \in \mathcal{A}$ was taken in state $s_{h} \in \mathcal{S}$ at time $h$. Each user has a different reward denoted by $\mathcal{R}_{u}=\left(R_{1 u}, \ldots, R_{H u}\right)$ where $R_{h u}: \mathcal{S} \times \mathcal{A} \rightarrow[0,1]$. Denote the MDP associated with the user $u$ by $\mathcal{M}_{u}:=\left(\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}_{u}\right)$. For the sake of simplicity, we will assume that the rewards are deterministic.

Assume that all the MDPs start at a random state $S_{1}$ with the same distribution. Consider a policy $\Pi:=\left(\pi_{1}, \ldots, \pi_{H}\right)$ where $\pi_{h}: \Delta(\mathcal{A}) \times \mathcal{S} \rightarrow \mathbb{R}^{+}$is a kernel - i.e, $\pi_{h}(\cdot \mid s)$ gives the probability distribution over actions given a state $s$ at time $h$. By $\left(S_{1: H}, A_{1: H}\right)$ we denote the trajectory $\left(S_{1}, A_{1}\right),\left(S_{2}, A_{2}\right), \ldots,\left(S_{H}, A_{H}\right) \in \mathcal{S} \times \mathcal{A}$. By $\left(S_{1: H}, A_{1: H}\right) \sim \mathcal{M}(\Pi)$ we mean the random trajectory under the policy $\Pi$ - where $A_{h} \sim \pi_{h}\left(\cdot \mid S_{h}\right)$ and $S_{h+1} \sim$ $P_{h}\left(\cdot \mid S_{h}, A_{h}\right)$. That is, it is the trajectory of the MDP under the policy $\Pi$. Define the value function of $\mathcal{M}_{u}$ under policy
$\Pi$ as: $V\left(\Pi, \mathcal{M}_{u}\right):=\mathbb{E}_{\left(S_{1: H}, A_{1: H}\right) \sim \Pi} \sum_{h=1}^{H} R_{h u}\left(S_{h}, A_{h}\right)$. We will call a policy $\hat{\Pi}_{u}$ to be $\epsilon$ optimal for $\mathcal{M}_{u}$ if $V\left(\hat{\Pi}_{u}, \mathcal{M}_{u}\right) \geq \sup _{\Pi} V\left(\Pi, \mathcal{M}_{u}\right)-\epsilon$. Our goal is to find $\epsilon$ optimal policies for every $u \in[N]$ under low rank assumptions on the rewards $R_{u h}$. We assume that we are allowed to pick any user $u$ and query a trajectory corresponding to any policy $\Pi$.

Reward Free Exploration: The objective of reward free RL is to explore an MDP (without looking at the rewards) such that we can obtain the optimal policy for every possible reward. After collecting $K$ trajectories from the MDP sequentially (denoted by $\mathcal{D}_{K}$ ), the algorithm outputs functions $\hat{\Pi}$ and $\hat{V}$ whose input is a reward function $\mathcal{R}=\left(R_{h}(\cdot, \cdot)\right)_{h=1}^{H}$ (bounded between $[0,1]$ ) and the output is a nearly-optimal policy $\hat{\Pi}(\mathcal{R})$ and its estimated value $\hat{V}(\hat{\Pi}(\mathcal{R}))$ for this reward function. Denote the MDP with this reward function by $\mathcal{M}_{\mathcal{R}}$. Given $\epsilon>0$ and $\delta \in[0,1]$, we let $K^{\text {rf }}(\epsilon, \delta)$ to be such that whenever $K \geq K^{\text {rf }}(\epsilon, \delta)$, with probability at-least $1-\delta$ we have:
a) $\sup _{\mathcal{R}}\left|V\left(\hat{\Pi}, \mathcal{M}_{\mathcal{R}}\right)-\hat{V}(\hat{\Pi}(\mathcal{R}))\right| \leq \epsilon$ and b) $\hat{\Pi}$ is an $\epsilon$ optimal policy for $\mathcal{M}_{\mathcal{R}}$ for every $\mathcal{R}$.

This setting was introduced in (Jin et al., 2020a). In this work, we will use the reward free exploration algorithms in (Zhang et al., 2020) for tabular MDPs and (Wagenmaker et al. 2022) for linear MDPs.

Tabular MDP Setting $\mathcal{S}$ and $\mathcal{A}$ are finite sets. Denote the reward $R_{h u}(s, a)$ by the $N \times|\mathcal{S}||\mathcal{A}|$ matrix $R_{h}$ where $R_{h}(u,(s, a))=R_{h u}(s, a)$. We have the following low-rank assumption:
Assumption (Tab) 2.1. The matrix $R_{h}$ has rank $r$ for some $r \leq \frac{1}{2} \min (N,|\mathcal{S} \| \mathcal{A}|)$.

Linear MDP Setting Our definition is slightly different from the one in (Jin et al., 2020b): a) we use two different embeddings for rewards and transitions and $b$ ) we impose an $\ell_{1}$ constraint on the transition embedding and an $\ell_{2}$ constraint on the reward embedding (instead of $\ell_{2}$ on both). This is a natural choice since transition embeddings describe a mixture of probability measures as the law of the next state. On a technical level, $\ell_{1}$ norm is natural when controlling the statistical error due to policy search described in Section 6 , which is based on the structural result in Lemma 6.1
We consider embeddings $\phi: \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^{d}, \psi: \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^{d}$ such that $\|\phi(s, a)\|_{1} \leq 1,\|\psi(s, a)\|_{2} \leq 1$. We make the following assumptions:

1. There exists $\theta_{h u} \in \mathbb{R}^{d},\left\|\theta_{h u}\right\|_{2} \leq \sqrt{d}$ such that $R_{h u}(s, a)=\left\langle\theta_{h u}, \psi(s, a)\right\rangle$ and $R_{h}(s, a) \in[0,1]$.
2. There exist signed measures $\mu_{1 h}, \ldots, \mu_{d h}$
over the space $\mathcal{S}$ such that: $P_{h}(\cdot \mid s, a)=$ $\sum_{i=1}^{d} \mu_{i h}(\cdot)\left\langle\phi(s, a), e_{i}\right\rangle$

We will assume that $\mu_{i}$ are such that $\left\|\int \mu_{i h}(d s) \phi(s, a) \pi(d a \mid s)\right\|_{1} \quad \leq \quad 1 \quad$ and $\sup _{i, h} \int\left|\mu_{i h}(d s)\right| \leq C_{\mu}$. This is true whenever $\mu_{i h}$ are probability measures. We consider different embeddings for transition $(\phi)$ and reward $(\psi)$ as the transition embeddings have a natural $\|\cdot\|_{1}$ structure since they give linear combinations of measure which make up $P_{h}(\cdot \mid s, a)$. We denote the $N \times d$ matrix whose $u$-th row is $\theta_{u h}^{\top}$ to be $\Theta_{h}$. The low-rank assumption in this setting takes the following form:
Assumption (Lin) 2.2. . The $N \times d$ matrix $\Theta_{h}$ has rank $r \leq \frac{1}{2} \min (N, d)$.

We restrict our attention to policies given by some fixed policy space $\mathbb{Q}$. As explained below in Section 3, this is necessitated by the fact that our techniques are required to go beyond the standard RL setup and might necessarily require non-deterministic policies. However, the space of all possible policies can be very large and intractable. We refer to Section $A$ for the construction of randomized policy class $\mathbb{Q}$ such that it contains all $\epsilon$-optimal policies for every possible linear reward. With some abuse of notation, we define the total variation distance between two kernels as: $\operatorname{TV}\left(\pi_{h}, \pi_{h}^{\prime}\right):=$ $\sup _{s \in \mathcal{S}} \operatorname{TV}\left(\pi_{h}(\cdot \mid s), \pi_{h}^{\prime}(\cdot \mid s)\right)$. We define a distance over $\mathbb{Q}$ by $D_{\mathbb{Q}}\left(\Pi^{(1)}, \Pi^{(2)}\right)=\sup _{h \in[H]} \operatorname{TV}\left(\pi_{h}^{(1)}, \pi_{h}^{(2)}\right)$, where $\Pi^{(i)}=\left(\pi_{h}^{(i)}: h \in[H]\right)$.

## 3. Connection to Multi-Agent RL

We now connect our results to multi-agent reinforcement learning in order to demonstrate why the problem of collaborative RL as described above can be hard. Low rank matrix estimation requires random measurements with specific isotropy properties. For instance, matrix completion results are derived when we observe entries from uniformly random indices $((\widehat{\text { Recht }}, 2011))$. In our context, this translates to sampling from specific distribution of state and action $\left(S_{h}, A_{h}\right)$ at time $h$ by depolying a policy $\Pi$ over a uniformly random user $u \in[N]$. Rather than maximizing a scalar reward, this requires us to sample from a distribution with certain properties, going beyond the framework of standard RL. In fact, Section A. we show that this requires randomized policies even for simple MDPs. In the tabular MDP case, we sidestep these issues with clever algorithm design. However, in the linear MDP case, this does not seem to be feasible. We use the following connection to mean-field limits of multi-agent RL as sketched below in order to solve the sampling question.

Suppose we pick users uniformly at random $\left(U_{1}, \ldots, U_{T} \in\right.$
[ $N]$ ) and deploy a policy $\Pi$ for each of them with corresponding trajectories $S_{1: H}^{(t)}, A_{1: H}^{(t)} \sim \mathcal{M}(\Pi)$. We observe 'linear measurements' of $\Theta_{h}$ of the form $\left(e_{U_{t}}, \psi\left(S_{h}^{(t)}, A_{h}^{(t)}\right), e_{U_{t}}^{\top} \Theta_{h} \psi\left(S_{h}^{(t)}, A_{h}^{(t)}\right)\right)$. To achieve matrix estimation, we need to query $\left(S_{h}^{(t)}, A_{h}^{(t)}\right)$ such that the distribution of $\psi\left(S_{h}^{(t)}, A_{h}^{(t)}\right)$ is 'nearly isotropic' ( See (4) in Section 7. Let $\Gamma_{h}(\Pi)$ denote the distribution of $S_{h}^{(1)}, A_{h}^{(1)}$. In Theorem6.2 we show that the conditions in (4) are satisfied whenever $J\left(\Gamma_{h}(\Pi)\right)>0$ for $J: \Delta(\mathcal{S} \times \mathcal{A}) \rightarrow \mathbb{R}$ given by:

$$
\begin{aligned}
& J\left(\Gamma_{h}(\Pi)\right):= \\
& \quad \inf _{x \in \mathcal{S}^{d-1}} \mathbb{E}\left|\left\langle\psi\left(S_{h}^{(1)}, A_{h}^{(1)}\right), x\right\rangle\right| \sqrt{d}-\xi d\left\langle\psi\left(S_{h}^{(1)}, A_{h}^{(1)}\right), x\right\rangle^{2}
\end{aligned}
$$

Our objective now is to find a policy by solving the following optimization problem

$$
\begin{equation*}
\arg \sup _{\Pi \in \mathbb{Q}} J\left(\Gamma_{h}(\Pi)\right) \tag{1}
\end{equation*}
$$

This is similar to the mean field multi-agent control problem presented in (Cammardella et al. 2020). To demonstrate the connection to multi-agent systems, consider $n$ agents with the same MDP $\mathcal{M}$ and embeddings $\phi, \psi$. A trajectory here corresponds to jointly and independently running MDP associated with each agent with the same policy. The collective reward of the system is given by $J\left(\hat{\Gamma}_{h}\right)$, where $\hat{\Gamma}_{h}$ denotes the empirical distribution of state-actions of the $n$ agents at time $h$. Note that, picking a policy $\Pi_{n}$ to maximize this reward is a reward maximization problem on the joint multi-agent system. And, for any fixed policy $\Pi$, as $n \rightarrow \infty, \hat{\Gamma}_{h}(\Pi) \rightarrow \Gamma_{h}(\Pi)$ under reasonable assumptions on the state space via the law of large numbers. Hence $J\left(\hat{\Gamma}_{h}(\Pi)\right) \rightarrow J\left(\Gamma_{h}(\Pi)\right)$ under continuity (in some appropriate distance between probability measures). Therefore the planning problem in (1) is the same as the multi-agent RL problem described above in the limit $n \rightarrow \infty$.

## 4. The Algorithm

Our algorithm proceeds in 4 phases. In phase 1, we run reward free RL, which selects trajectories from uniformly random users since all the users share the same MDP. Thus, this does not incur a large per-user sample complexity. This step allows us to find near-optimal for any reward function of our choice. At the end of phase 1, the main unknown will the reward matrix of the users.

Phase 2 is the main technical contribution of work. In phase 2 , we use the reward free RL output from phase 1 in order to design collaborative exploration policies. These policies obtain the right linear measurements of the low-rank reward
matrix so that we can successfully apply matrix estimation algorithms in phase 3. Phase 4 uses the reward estimate from phase 3 and the reward free RL output from phase 1 in order to learn the optimal policy for every user.

### 4.1. Tabular MDP Case:

Phase 1: Reward Free Exploration We run the reward free RL algorithm in (Zhang et al. 2020) for $K^{\text {rf }}\left(\frac{\epsilon}{8}, \frac{\delta}{2}\right)=$ $C \frac{|\mathcal{S} \| \mathcal{A}| H^{2}\left(|\mathcal{S}|+\log \left(\frac{1}{\delta}\right)\right)}{\epsilon^{2}}$ polylog $\left(\frac{|\mathcal{S}||\mathcal{A}| H}{\epsilon}\right)$ time steps by picking the MDP corresponding to a uniformly random user whenever the reward free RL algorithm queries a trajectory. Let the output of the reward free RL algorithm be $\hat{\Pi}$ and $\hat{V}$.

Phase 2: Querying the Reward Matrix In this phase we query a 'uniform mask' with the parameter $p$ for the reward matrix $R_{h}$ using Algorithm 1. For each $(s, a) \in \mathcal{S} \times \mathcal{A}$ and $h \in[H]$, maintain a counter $T_{h,(s, a)}$ for $(s, a) \in \mathcal{S} \times \mathcal{A}$ and $h \in[H]$, initialized at 0 . Given the 'active sets' $\mathcal{G}=$ $\left(\mathcal{G}_{h}\right)_{h \in[H]} \subseteq \mathcal{S} \times \mathcal{A}$ and $h \in[H]$, we define the reward $\mathcal{J}(; \mathcal{G})=\left(J_{1}, \ldots, J_{h}\right)$ by

$$
\begin{equation*}
J_{h}(s, a ; \mathcal{G})=\mathbb{1}\left((s, a) \in \mathcal{G}_{h}\right) \tag{2}
\end{equation*}
$$

We will denote this reward by $\mathcal{J}(\cdot ; \mathcal{G})$. Initialize active set $\mathcal{G}=\left(\mathcal{G}_{h}\right)_{h \in[H]}$ such that $\mathcal{G}_{h}=\mathcal{S} \times \mathcal{A}$. We initialize the reward matrix $\hat{R}_{h}(u,(s, a))=*$, where $*$ denotes unknown entry. This algorithm terminates when it detects that sufficient number of samples have been collected for matrix completion.

Phase 3: Reward Matrix Completion We receive $\mathcal{G}_{h}$ and the partially observed matrix $\hat{R}_{h}$ for each $h$ as the output of Algorithm 1. By $\hat{R}_{h}^{\mathcal{G}_{h}^{\mathfrak{C}}}$, we denote the sub-matrix where the columns corresponding to $\mathcal{G}_{h}$ are deleted. We use the nuclear norm minimization algorithm given in (Recht, 2011) to recover $R_{h}^{\mathcal{G}_{h}^{\mathcal{C}}}$ from $\hat{R}_{h}^{\mathcal{G}_{h}^{\complement}}$ for every $h \in[H]$.

Phase 4: Computing the Optimal Policy Phase 3 outputs the completed sub-matrix $R_{h}^{\mathcal{G}_{h}^{\mathrm{C}}}$, where only the columns corresponding to $\left|\mathcal{G}_{h}^{\mathrm{C}}\right|$ are recovered. We construct the recovered matrix $\bar{R}_{h}$ by setting $\bar{R}_{h}^{\mathcal{G}_{h}^{\text {© }}}=R_{h}^{\mathcal{G}_{h}^{\text {© }}}$ and $\bar{R}_{h}^{\mathcal{G}_{h}}=0$. We compute the optimal policy for each user using the rewards from $\bar{R}_{h}$ via the output of the reward free RL, $\hat{\Pi}$, from Phase 1.

### 4.2. Linear MDP Case:

Phase 1 : Reward Free RL We run the reward free RL algorithm for Linear MDPs from (Wagenmaker et al., 2022), with error $\epsilon$ and probability of failure $\frac{\partial}{4}$. We use trajectories from random users whenever a trajectory is queried. Here, $K^{\text {rf }}(\epsilon, \delta / 4)=\frac{C d H^{5}\left(d+\log \left(\frac{1}{\delta}\right)\right)}{\epsilon^{2}}+\frac{C d^{9 / 2} H^{6} \log ^{4}\left(\frac{1}{\delta}\right)}{\epsilon}$.
Phase 2: Querying Linear Measurements of the Reward

```
Algorithm 1 Uniform Mask Sampler for Tabular MDPs
    Output:Active sets \(\mathcal{G}=\left(\mathcal{G}_{h}\right)_{h \in[H]}\), Partially complete
    matrix \(\hat{R}_{h}\)
    \(t \leftarrow 0 \quad \hat{P}^{\mathcal{G}} \leftarrow \hat{V}(\mathcal{J}(\cdot ; \mathcal{G}))\)
    \(\hat{\Pi}^{\mathcal{G}} \leftarrow \hat{\Pi}(\mathcal{J}(\cdot ; \mathcal{G}))\)
    while \(\hat{P}^{\mathcal{G}}>\frac{\epsilon}{2}\) do
        \(U_{t} \leftarrow \operatorname{Unif}([N])\{\) Pick a user uniformly at random \(\}\)
        \(S_{1: H}^{(t)}, A_{1: H}^{(t)}, R_{1: H}^{(t)} \sim \mathcal{M}_{U_{t}}\left(\hat{\Pi}^{\mathcal{G}}\right)\{\) Query trajectory \(\}\)
        for \(h \in[H]\) do
            if \(\left(S_{h}^{(t)}, A_{h}^{(t)}\right) \in \mathcal{G}_{h}\) and \(R_{h}\left(U_{t},\left(S_{h}^{(t)}, A_{h}^{(t)}\right)\right)=*\)
            then
                \(T_{h,\left(S_{h}^{(t)}, A_{h}^{(t)}\right)} \leftarrow T_{h,\left(S_{h}^{(t)}, A_{h}^{(t)}\right)}+1\{\) Increment
                count \(\}\)
                \(\hat{R}_{h}\left(U_{t},\left(S_{h}^{(t)}, A_{h}^{(t)}\right)\right) \leftarrow R_{h}^{(t)}\) \{Fill Missing En-
                try \(\}\)
            end if
            if \(T_{h,\left(S_{h}^{(t)}, A_{h}^{(t)}\right)}=N p\) then
                \(\mathcal{G}_{h} \leftarrow \mathcal{G}_{h} \backslash\left\{\left(S_{h}^{(t)}, A_{h}^{(t)}\right)\right\}\{\) Remove element from
                Active Set \(\}\)
            end if
        end for
        \(t \leftarrow t+1\)
        \(\hat{P}^{\mathcal{G}} \leftarrow \hat{V}(\mathcal{J}(\cdot ; \mathcal{G}))\)
        \(\hat{\Pi}^{\mathcal{G}} \leftarrow \hat{\Pi}(\mathcal{J}(\cdot ; \mathcal{G}))\)
    end while
```

Matrix We obtain policies whose trajectory data allows low rank matrix estimation of the reward matrix.

Step 1: For each time step $h \in[H]$, we want to query obtain samples $\left(s_{h}^{(t)}, a_{h}^{(t)}\right)$ such that $\sum_{t=1}^{T} \phi\left(s_{h}^{(t)}, a_{h}^{(t)}\right) \phi^{\top}\left(s_{h}^{(t)}, a_{h}^{(t)}\right) \succeq \kappa^{2} I$. This can be done by Algorithm 2 Given a projector $Q$ to some subspace of $\mathbb{R}^{d}$, by $\mathcal{Q}_{h, Q}$ denote the reward $\|Q \phi(s, a)\|^{2}$ at time $h$ and 0 otherwise. The termination condition ensures that we see enough data in all directions $\phi$, which allows us to find collaborative exploration policy below.

Step 2: Using the observations given in Step 1, we compute the policy $\hat{\Pi}^{f, h}$ which approximately satisfies the property given in Assumption 5.6. This procedure is described in Section 6.

Phase 3: Estimating Low Rank Reward Matrix For this, we use the active learning procedure given in Section 7 via row-wise linear measurements along with the policy $\Pi^{\mathrm{MC}, h}=\hat{\Pi}^{f, h}$, which was computed in Phase 2.

Phase 4: Computing the Optimal Policy Once the reward matrix $\Theta_{h}$ have been reconstructed for every $h$ in Phase 3, we use the output of reward free RL in order to compute the $\epsilon$ optimal policy for each user.

```
Algorithm 2 Well Conditioned Matrix Sampler
    Input:Total time \(T\); Tolerance \(\gamma\); lower isometry \(\kappa\)
    Output: \(\phi_{h t}, S_{(h+1) t}\) for \(h \in[H-1], t \in[T]\)
    for \(h=1\) to \(h=H-1\) do
        \(Q \leftarrow I \quad \hat{\Pi}^{Q} \leftarrow \hat{\Pi}\left(\mathcal{Q}_{h, Q}\right)\)
        \(G_{\phi, h} \leftarrow 0\) \{Grammian initialized to 0\(\}\)
        for \(t=1\) : to \(t=T\) do
            \(U_{t} \sim \operatorname{Unif}([N])\{\) Pick a uniformly random user \(\}\)
            \(S_{1: H}, A_{1: H} \sim \mathcal{M}_{U_{t}}\left(\hat{\Pi}^{Q}\right)\{\) Obtain Trajectory \(\}\)
            \(\phi_{h t} \leftarrow \phi\left(S_{h}, A_{h}\right) \quad S_{(h+1) t} \leftarrow S_{h+1}\)
            \(G_{\phi, h} \leftarrow G_{\phi, h}+\phi_{h t} \phi_{h t}^{\top}\) \{Update Grammian \(\}\)
            if \(\exists y \in \mathbb{S}^{d-1}: y^{\boldsymbol{\top}} G_{\phi, h} y<\kappa^{2}\) then
                \(Q \leftarrow\) eigenspace of \(G_{\phi, h}\) with eigenvalues \(<\kappa^{2}\)
                \(\hat{\Pi}^{Q} \leftarrow \hat{\Pi}\left(\mathcal{Q}_{h, Q}\right)\)
            end if
        end for
    end for
```


## 5. Main Results

### 5.1. Tabular MDP:

Incoherence is a standard assumption for low rank matrix completion. This ensures that the matrix is not too sparse so that sparse measurements are sufficient to learn it. The following definition is used in (Recht, 2011).
Definition 5.1. Given a $r$ dimensional sub-space $U$ of $\mathbb{R}^{n}$, we define the coherence of $U$ as:

$$
\mu(U):=\frac{n}{r} \sup _{1 \leq i \leq n}\left\|P_{U} e_{i}\right\|^{2}
$$

A $n_{1} \times n_{2}$ matrix $M$ with singular value decomposition $U \Sigma V^{\top}$ is called $\left(\mu_{0}, \mu_{1}\right)$ coherent if:
a) The coherence of the row and column spaces of $M$ are at-most $\mu_{0} \mathrm{~b}$ ) The absolute value of every entry of $U V^{\top}$ is bounded above by $\mu_{1} \sqrt{\frac{r}{n_{1} n_{2}}}$.

Given a policy $\Pi$, and $\Omega \subseteq \mathcal{S} \times \mathcal{A}$, by $P_{h}^{\Pi}(\Omega)$ we denote the probability that at time $h$ we have $\left(S_{h}, A_{h}\right) \in \Omega$ under the policy $\Pi$.
Assumption (Tab) 5.2. Given the reward matrix $R_{h}$ and $\Omega_{h} \subset \mathcal{S} \times \mathcal{A}$, recall the notation for the sub-matrix $R_{h}^{\Omega_{h}}$ of $R_{h}$. If $\sup _{\Pi} P^{\Pi}\left(\Omega_{h}^{\complement}\right)<\epsilon$ have:

1) $R_{h}^{\Omega_{h}}$ is $\left(\mu_{0}, \mu_{1}\right)$ incoherent 2) $\left|\Omega_{h}^{\complement}\right| \leq \frac{|\mathcal{S}|}{2}$

The incoherence assumption for $R_{h}^{\Omega}$ makes sense since the set $\Omega^{\complement}$ cannot be easily reached with any policy with a probability larger than $\epsilon$. In fact we can arrive at an $\epsilon$ optimal policy for the original reward by just setting the rewards at $\Omega^{\text {C }}$ to be 0 . These can be thought of as redundant states which do not matter for our RL model with any reward.

Theorem 5.3. Suppose Assumption (Tab) 2.1 5.2 hold. Let the parameter $p=$ $C \frac{\max \left(\mu_{1}^{2}, \mu_{0}\right) r(N+|\mathcal{S} \| \mathcal{A}|) \log ^{2}|\mathcal{S}||\mathcal{A}| \log \left(\frac{H}{\delta}\right)}{N|\mathcal{S}||\mathcal{A}|}$ for some large enough constant $C$. Assume that $|\mathcal{S}||\mathcal{A}|$ and $N$ are large enough such that $p<1 / 2$. Then, with probability at-least $1-\delta$, we can find an $\epsilon$ optimal policy $\hat{\Pi}_{u}$ for every user $u \in[N]$ whenever the total number of trajectories queried is:

$$
\begin{aligned}
C & \frac{|\mathcal{S}||\mathcal{A}| H^{2}\left(|\mathcal{S}|+\log \left(\frac{1}{\delta}\right)\right)}{\epsilon^{2}} \operatorname{poly} \log \left(\frac{|\mathcal{S}||\mathcal{A}| H}{\epsilon}\right) \\
& +\frac{C \max \left(\mu_{1}^{2}, \mu_{0}\right) r(N+|\mathcal{S}||\mathcal{A}|) H \log ^{2}|\mathcal{S}||\mathcal{A}| \log \left(\frac{H}{\delta}\right)}{\epsilon}
\end{aligned}
$$

Remark 5.4. For large $N$, the number of trajectories queried per user is $\tilde{O}\left(\frac{r H \log ^{2}(|\mathcal{S} \| \mathcal{A}|)}{\epsilon}\right)$, which is an exponential improvement in the state-space size dependence when compared to the minimax rate of $\frac{|\mathcal{S} \| \mathcal{A}| H^{2}}{\epsilon^{2}}$ Dann \& Brunskill, 2015) for learning a single MDP. Every phase in the algorithm has polynomial computational complexity in $N,|\mathcal{S}||\mathcal{A}|$ and $\frac{1}{\epsilon}$. The probability $p$ is chosen such that $p|\mathcal{S} \| \mathcal{A}| N=\tilde{O}(r(|\mathcal{S}||\mathcal{A}|+N))$, which is the number of free parameters required to describe a rank $r$ matrix.

### 5.2. Linear MDP

Assumption (Lin) 5.5. There exists a $\gamma>0$ such that for every $x \in \mathcal{S}^{d-1}$, and every $h \in[H]$ there exists a policy $\Pi_{x, h}$ such that whenever $S_{1: H}, A_{1: H} \sim \Pi_{x, h}$, $\mathbb{E}\left\langle\phi\left(S_{h}, A_{h}\right), x\right\rangle^{2} \geq \gamma$

The assumption above shows that we can obtain information about all directions. If this does not hold for any $\gamma$, then $\phi\left(S_{h}, A_{h}\right)$ does not have any component in some direction $x_{0}$ with any policy. Thus, we can remove the sub-space spanned by $x_{0}$ and make the embedding space $\mathbb{R}^{d-1}$ at time $h$.

Assumption (Lin) 5.6. There exist $\zeta, \xi>0$ such that for every $h \in[H]$, there exists a policy $\Pi^{h, \zeta, \xi} \in \mathbb{Q}$ such that whenever $S_{1: H}, A_{1: H} \sim \mathcal{M}\left(\Pi^{h, \zeta, \xi}\right)$, we have:

$$
\inf _{x \in \mathcal{S}^{d-1}} \mathbb{E}\left[\left|\left\langle x, \psi\left(S_{h}, A_{h}\right)\right\rangle\right| \sqrt{d}-\xi d\left\langle x, \psi\left(S_{h}, A_{h}\right)\right\rangle^{2}\right] \geq \zeta
$$

The assumption above ensures that there exist measurements $\psi\left(S_{h}, A_{h}\right)$ which are conducive to low rank matrix estimation as considered in Section 7. This means that $\left|\left\langle\psi\left(S_{h}, A_{h}\right), x\right\rangle\right|=\Omega(1 / \sqrt{d})$ just like an isotropic random vector, which gives us information about all directions. However, this condition is much looser than the assumption of uniform distribution on the sphere.
Assumption (Lin) 5.7. For any $1 \geq \eta>0$, there exists an $\eta$ net for $\mathbb{Q}$, denoted by $\hat{\mathbb{Q}}_{\eta}$ such that $\log \left|\hat{\mathbb{Q}}_{\eta}\right| \leq D \log \left(\frac{1}{\eta}\right)$.

We refer to Section A, where we justify this assumption. We first demonstrate that deterministic policies which are sufficient for reward maximization (as used in (Jin et al., 2020b) cannot be used in this context, so a set of stochastic policies is required. We then construct such policy classes with $D=O(d H \log d H \log \log (|\mathcal{A}|))$.
Theorem 5.8. Suppose Assumptions (Lin) 2.25 .55 .6 and 5.7 hold and suppose $\epsilon<\frac{\gamma}{2}$. In Algorithm 2, we set $\kappa=\frac{C C_{\mu} d H \sqrt{d H+D}(\sqrt{d}+\xi d)}{\zeta} \sqrt{\log \left(\frac{C_{\mu} H(d+D)}{\zeta \gamma \delta}\right)}$ and $T=C \frac{\kappa^{2} d}{\gamma^{2}} \log \frac{d \kappa}{\gamma}$.
Then, with probability at least $1-\delta$, our algorithm finds $\epsilon$ optimal policy for every user $u \in[N]$ with the total number of trajectories being bounded by: $T_{\mathrm{rf}}+T_{\mathrm{pol}}+T_{\mathrm{mat}-c o m p}$, where:

$$
\begin{gather*}
T_{\mathrm{rf}}=\frac{d H^{5}\left(d+\log \left(\frac{1}{\delta}\right)\right)}{\epsilon^{2}}+\frac{d^{9 / 2} H^{6} \log ^{4}\left(\frac{1}{\delta}\right)}{\epsilon} \\
T_{\mathrm{pol}}=\frac{C_{\mu}^{2} d^{5} H^{3}(d H+D) \log ^{2}\left(\frac{C_{\mu} H(d+D)}{\zeta \gamma \delta}\right)}{\zeta^{2} \gamma^{2}} \\
T_{\text {mat-comp }}=C \frac{H r(N+d \log N) \log \frac{d}{\zeta \xi}}{\zeta^{2} \xi^{2}} \\
+\frac{H \log N \log \left(\frac{\log N}{\delta}\right)}{\zeta^{2} \xi^{2}} \tag{3}
\end{gather*}
$$

Remark 5.9. When $N$ is very large, the per user sample complexity is $O(H r)$, which is much better than the mini-max optimal complexity of $\Omega\left(d^{2} H^{2}\right)$ (Wagenmaker et al. 2022). While Phases 1 and 2 of the algorithm have a computational complexity which is polynomial in $d$ and $\frac{1}{\epsilon}$, the optimization problems posed in Phase 3 and 4 are not necessarily polynomial time. We leave the computational aspects to future work. The sample complexity $\tilde{O}(r(N+d))$ corresponds to the number of free parameters required to describe a rank $r$ matrix.

## 6. Obtaining Policies With Given Statistics

In this section, we consider the Linear MDP setting and describe the sub-routine described in Step 2 of Phase 2 of the algorithm where we compute a policy $\Pi^{f, h}$ such that the law of $\phi\left(S_{h}, A_{h}\right)$ under this policy approximately satisfies the property given in Assumption 5.6 This is required in order to use the guarantees for low matrix estimation in Phase 3, which is described in Section 7. We first state a structural lemma which characterizes the law of $S_{h+1}, A_{h+1}$ under any policy $\Pi$.
Lemma 6.1. Consider any policy $\Pi=\left(\pi_{1}, \ldots, \pi_{H-1}, \pi_{H}\right)$ to the MDP $\mathcal{M}$. Let $S_{1: H}, A_{1: H} \sim \mathcal{M}(\Pi)$. Then for any
bounded, measurable function $g: \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$, we have:
$\mathbb{E} g\left(S_{h+1}, A_{h+1}\right)=\sum_{i=1}^{d} \nu_{i} \int g(s, a) d \mu_{i h}(d s) \pi_{h+1}(d a \mid s)$

Where $\nu_{i}:=\left\langle\mathbb{E} \phi\left(S_{h}, A_{h}\right), e_{i}\right\rangle$
We now want to estimate certain statistics under any policy using available data, obtained from the output of Algorithm 2. Notice that the output of Algorithm 2 gives a sequence of random variables $\left(\phi_{h 1}, s_{(h+1) 1}\right), \ldots,\left(\phi_{h T}, s_{(h+1) T}\right) \in \mathbb{R}^{d} \times \mathcal{S}$ such that $\left(s_{(h+1) l}\right)_{l=1}^{T} \mid\left(\phi_{h l}\right)_{l=1}^{T} \sim \prod_{l=1}^{T}\left(\sum_{i=1}^{d}\left\langle\phi_{h l}, e_{i}\right\rangle \mu_{h i}(\cdot)\right)$ and $G_{\phi, h}:=\sum_{t=1}^{T} \phi_{h t} \phi_{h t}^{\top}$. For any measurable function $g: \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^{K}, \nu \in \mathbb{R}^{d}$ such that $\|\nu\| \leq 1$ and any randomized policy $\Pi=\left(\pi_{1}, \ldots, \pi_{H}\right)$ we define:

1. $\mathcal{T}_{1}\left(g, \pi_{1}\right)=\mathbb{E} \int g\left(S_{1}, a\right) \pi_{1}\left(d a \mid S_{1}\right)$
2. $\mathcal{T}_{h+1}\left(g ; \nu, \pi_{h+1}\right)=$ $\sum_{i=1}^{d}\left\langle\nu, e_{i}\right\rangle \int \mu_{i h}(d s) \pi_{h+1}(d a \mid s) g(s, a)$ when $h \leq$
3. $E_{1}^{\nu}(\Pi):=\left\|\mathcal{T}_{1}\left(\phi, \pi_{1}\right)-\nu\right\|_{1}$
4. $E_{h}^{\nu}(\Pi)=\inf _{\nu_{1}, \ldots, \nu_{h-1} \in \mathcal{B}_{d}(1)} F\left(\Pi, \nu_{1}, \ldots, \nu_{h-1}, \nu\right)$ whenever $h>1$

Define $\alpha_{h t, \nu}=\phi_{h t}^{\top} G_{\phi, h}^{-1} \nu$. We estimate these operators from data as follows:

1. $\hat{\mathcal{T}}_{1}\left(g, \pi_{1}\right)=\frac{1}{T} \sum_{t=1}^{T} \int g\left(s_{1 t}, a\right) \pi_{1}\left(d a \mid s_{1 t}\right)$
2. $\hat{\mathcal{T}}_{h+1}\left(g, \nu, \pi_{h+1}\right)=$
$\sum_{t=1}^{T} \alpha_{h t, \nu} \int g\left(s_{(h+1) t}, a\right) \pi_{h+1}\left(d a \mid s_{(h+1) t}\right)$
3. $\hat{E}_{1}^{\nu}(\Pi):=\left\|\hat{\mathcal{T}}_{1}\left(\phi, \pi_{1}\right)-\nu\right\|_{1}$
4. $\hat{E}_{h}^{\nu}(\Pi)=\inf _{\nu_{1}, \ldots, \nu_{h-1} \in \mathcal{B}_{d}(1)} \hat{F}\left(\Pi, \nu_{1}, \ldots, \nu_{h-1}, \nu\right)$ whenever $h>1$

Where, for $h>1$ and $\nu_{1}, \ldots, \nu_{h-1} \in \mathcal{B}_{d}(1)$, we have defined:

1. $F\left(\Pi, \nu_{1}, \ldots, \nu_{h-1}, \nu_{h}\right) \quad:=\quad E_{1}^{\nu_{1}}(\Pi) \quad+$ $\sum_{j=2}^{h}\left\|\mathcal{T}_{j}\left(\phi, \nu_{j-1}, \pi_{j}\right)-\nu_{j}\right\|_{1}$
2. $\hat{F}\left(\Pi, \nu_{1}, \ldots, \nu_{h-1}, \nu_{h}\right) \quad:=\quad \hat{E}_{1}^{\nu_{1}}(\Pi) \quad+$ $\sum_{j=2}^{h}\left\|\hat{\mathcal{T}}_{j}\left(\phi, \nu_{j-1}, \pi_{j}\right)-\nu_{j}\right\|_{1}$

Define $f(s, a ; x):=|\langle x, \psi(s, a)\rangle| \sqrt{d}-\xi d\langle x, \psi(s, a)\rangle^{2}$. The output of our method is:
1.

$$
\begin{aligned}
& \hat{\Pi}^{f, 1}= \\
& \arg \sup _{\Pi=\left(\pi_{1}, \ldots, \pi_{H}\right) \in \mathbb{Q}} \inf _{x \in \mathcal{S}^{d-1}} \hat{\mathcal{T}}_{1}\left(f(\cdot ; x), \pi_{1}\right)
\end{aligned}
$$

2. 

$$
\begin{aligned}
& \left(\hat{\Pi}^{f, h}, \hat{\nu}\right)= \\
& \quad \arg \sup _{\substack{\nu \in \mathcal{B}(1) \\
\Pi=\left(\pi_{1}, \ldots, \pi_{H}\right) \in \mathbb{Q}}} \inf _{x \in \mathcal{S}^{d-1}} \hat{\mathcal{T}}_{h}\left(f(\cdot ; x) ; \nu, \pi_{h}\right) \\
&
\end{aligned}
$$

whenever $h>1$, subject to $\hat{E}^{\hat{\nu}, h-1}\left(\hat{\Pi}^{f, h}\right) \leq \eta_{0}$
3. Assign output: $\hat{\Pi}^{\zeta, \xi, h}=\hat{\Pi}^{f, h}$

The idea behind this method is as follows. First, using the output of algorithm 2 we construct $\hat{\mathcal{T}}_{h}\left(g, \nu, \pi_{h}\right)$, which approximates the functional $\mathcal{T}_{h}\left(g, \nu, \pi_{h}\right)$ uniformly for every $\nu, \pi_{h}$. This is shown in Lemma E. 1 in the appendix. We will show in Theorem 6.2 that obtaining policies which can be used with the matrix completion routine reduces to picking a policy $\Pi^{f, h}=\left(\pi_{1}, \ldots, \pi_{H}\right)$ such that whenever $S_{1: H}, A_{1: H} \sim \mathcal{M}\left(\Pi^{f, h}\right)$, we must have: $\inf _{x \in \mathcal{S}^{d-1}} \mathbb{E} f\left(S_{h}, A_{h} ; x\right) \geq \zeta$. Now, we use Lemma 6.1 to conclude that if such a policy exists, then there exist $\nu_{1}, \ldots, \nu_{h-1}$ such that $\mathbb{E} \phi\left(S_{j}, A_{j}\right)=\nu_{j}$ and $\inf _{x \in \mathcal{S}^{d-1}} \mathcal{T}_{h}\left(f(; x) ; \nu_{h}, \pi_{h}\right) \geq \zeta$. Since we only have sample access, we find such a policy approximately by optimizing using the estimates $\hat{\mathcal{T}}$ instead of the exact functional $\mathcal{T}$ as described above.
Theorem 6.2. We condition on the event $G_{\phi, h} \geq \kappa^{2} I$ for every $h \in[H]$. Let $\kappa, \eta, \eta_{0}$ be such that for some small enough constants $c_{0}, c>0$ and a large enough constant $C>0$ :

$$
\begin{aligned}
& \text { 1. } \eta \leq c \frac{\zeta}{C_{\mu} d H(\sqrt{d}+\xi d) H} \sqrt{\frac{\kappa^{2}}{T}} ; \quad \eta_{0}=c_{0} \frac{\zeta}{C_{\mu}(\sqrt{d}+d \xi)} \\
& \text { 2. } \kappa \geq C \frac{C_{\mu}(\sqrt{d}+\xi d) d H}{\zeta} \sqrt{\log \left(\frac{d H\left|\hat{\mathbb{Q}}_{\eta}\right|}{\delta}\right)+d H \log \left(\frac{d}{\eta}\right)}
\end{aligned}
$$

Recall the policy $\hat{\Pi}^{f, h}$. Suppose the Assumption 5.6 holds. Then, with probability at-least $1-\frac{\delta}{4}$ we obtain the policy $\hat{\Pi}^{f, h}$ is such that whenever $S_{1: H}, A_{1: H} \sim \mathcal{M}\left(\hat{\Pi}^{f, h}\right)$, we have:

$$
\inf _{x \in \mathcal{S}^{d-1}} \mathbb{E} f\left(S_{h}, A_{h} ; x\right) \geq \frac{\zeta}{2}
$$

This implies that $\psi\left(S_{h}, A_{h}\right)$ satisfies $\mathbb{E}\left|\left\langle\psi\left(S_{h}, A_{h}\right), x\right\rangle\right| \geq$ $\frac{\zeta}{2 \sqrt{d}} ; \quad \mathbb{E} \psi\left(S_{h}, A_{h}\right) \psi_{i k}^{\top} \leq \frac{1}{d \xi^{2}}$

## 7. Matrix Estimation with Row-wise Linear Measurements

We now describe the active learning based low rank matrix estimation procedure. For an unknown rank $r$ matrix $\Theta^{*}$ (corresponding to $\Theta_{h}^{*}$ in the definition of Linear MDPs) of dimensions $N \times d$, we are allowed to query samples of the form $\left(e_{i}, \psi, e_{i}^{\top} \Theta^{*} \psi\right)$ for any $i \in[N]$ of our choice and $\psi=\psi\left(S_{h}, A_{h}\right)$ where $S_{1: H}, A_{1: H} \sim \mathcal{M}\left(\Pi^{\mathrm{MC}, h}\right)$, for some input policy $\Pi^{\mathrm{MC}, h}$. This corresponds to running the MDP of user $i$, with the policy $\Pi^{\mathrm{MC}, h}$ and observing the reward at time $h$, given by $\left\langle e_{i}, \Theta_{h}^{*} \Psi\left(S_{h}, A_{h}\right)\right\rangle$. We want to estimate the matrix $\Theta^{*}$ from these samples with high-probability.

### 7.1. The Estimator

Given any $N \times d$ matrix $\Delta$, by $\Delta_{i}^{\top}$, we denote its $i$-th row. Given $K \in \mathbb{N}$, and a sequence of vectors $\Psi=\left(\psi_{i k} \in\right.$ $\left.\mathbb{R}^{d}\right)_{i \in[N], k \in[K]}$.

$$
L(\Delta, \Psi):=\frac{1}{N K} \sum_{i=1}^{N} \sum_{k=1}^{K}\left|\left\langle\Delta_{i}, \psi_{i k}\right\rangle\right|^{2}
$$

We estimate $\Theta^{*}$ row-wise using the following iterative procedure, where recover some rows of $\Theta^{*}$ into $\hat{\Theta}$ in each iteration and obtain the corresponding linear measurements of $\Theta^{*}$. Letting the set of unknown rows at iteration $t$ to be $\bar{I}_{t-1}$ (with $\bar{I}_{0}=[N]$ ). We draw a fresh sequence of vectors $\Psi^{(t)}$ from some distribution, we then recover some rows $\bar{I}_{t}^{\mathrm{C}} \subseteq \bar{I}_{t-1}$ of $\Theta^{*}$ and store them in $\hat{\Theta}$.

1. Draw $\Psi^{(t)}=\left(\psi_{i k}^{(t)}\right)_{k \in\left[K_{t}\right], i \in \bar{I}_{t-1}}$, we obtain $\theta_{i k}^{*}=$ $e_{i}^{\top} \Theta^{*} \psi_{i k}^{(t)}$.
2. Consider the loss function

$$
\begin{aligned}
& L\left(\Theta-\Theta^{*}, \Psi^{(t)}\right) \\
& :=\frac{1}{K_{t}\left|\bar{I}_{t-1}\right|} \sum_{i \in \bar{I}_{t-1}} \sum_{k=1}^{K}\left|\left\langle\Theta_{i}, \psi_{i k}^{(t)}\right\rangle-\theta_{i k}^{*}\right|^{2} .
\end{aligned}
$$

3. Find a matrix $\Theta$ with rank $\leq r$ such that $L(\Theta-$ $\left.\Theta^{*}, \Psi^{(t)}\right)=0$.
4. Initialize $\bar{I}_{t} \leftarrow \emptyset$.
5. For every $i \in \bar{I}_{t-1}$, draw $K$ fresh samples using $\tilde{\psi}_{i 1}^{(t)}, \cdots, \tilde{\psi}_{i K}^{(t)}$ and compute $\sum_{k=1}^{K}\left|\left\langle\Theta_{i}, \tilde{\psi}_{i k}^{(t)}\right\rangle-\theta_{i k}^{*}\right|^{2}$. If $\sum_{k=1}^{K}\left|\left\langle\Theta_{i}, \tilde{\psi}_{i k}^{(t)}\right\rangle-\theta_{i k}^{*}\right|^{2}>0$ then add $i$ to $\bar{I}_{t}$ i.e., $\bar{I}_{t} \leftarrow \bar{I}_{t} \cup\{i\}$.
6. End routine when $\bar{I}_{t}=\emptyset$.

Suppose $\psi_{i k}$ are i.i.d random vectors such that there exist $\zeta, \xi>0$ such that for any $x \in \mathbb{R}^{d},\|x\|=1$ we have:

$$
\begin{align*}
& \left\|\psi_{i k}\right\| \leq 1 \text { almost surely; } \quad \mathbb{E}\left|\left\langle\psi_{i k}, x\right\rangle\right| \geq \frac{\zeta}{\sqrt{d}} \\
& \mathbb{E} \psi_{i k} \psi_{i k}^{\top} \leq \frac{1}{d \xi^{2}} \tag{4}
\end{align*}
$$

To give some intuition, the second condition above means that given any vector $x$, there is some overlap between the random vector $\psi$ and $x$, ensuring that every measurement gives us some information helping us to complete the matrix. The third assumption is a standard bound on the covariance matrix. Then we have the following theorem whose proof is presented in Section F
Theorem 7.1. Assume that $\sup _{i}\left\|\Theta_{i}^{*}\right\| \leq C_{\theta}$ and that the distribution of $\psi_{i k}^{(t)}$ satisfies (4). Suppose $K_{t}\left|\bar{I}_{t-1}\right|=$ $C \frac{r\left|\bar{I}_{t-1}\right|+d r}{\zeta^{2} \xi^{2}} \log \frac{d}{\zeta \xi}+C \frac{\log \left(\frac{\log N}{\delta}\right)}{\zeta^{2} \xi^{2}}$. With probability at-least $1-\delta$, the algorithm terminates after $\log N$ iterations and the output $\hat{\Theta}$ satisfies $\hat{\Theta}=\Theta^{*}$. Therefore, with probability at-least $1-\delta$, the sample complexity for estimation of $\Theta^{*}$ is:

$$
C \frac{r(N+d \log N)}{\zeta^{2} \xi^{2}} \log \frac{d}{\zeta \xi}+C \frac{\log N \log \left(\frac{\log N}{\delta}\right)}{\zeta^{2} \xi^{2}}
$$

## 8. Discussion

In this work, we designed methods to perform collaborative exploration of a number of MDPs with near optimal sample complexity. In particular, we encountered and solved the important problem of exploring such that the data can be used down-stream to learn the optimal policy for every one of the MDPs. We also established connections to mean-field limits of multi-agent reinforcement learning problems. In future work, we hope to use the observations in the current work in order to design collaborative RL algorithms based on practically deployed RL algorithms like PPO, DQN and TD3.

## References

Afsar, M. M., Crump, T., and Far, B. Reinforcement learning based recommender systems: A survey. ACM Computing Surveys, 55(7):1-38, 2022.

Agarwal, A., Jiang, N., Kakade, S. M., and Sun, W. Reinforcement learning: Theory and algorithms. CS Dept., UW Seattle, Seattle, WA, USA, Tech. Rep, pp. 10-4, 2019.

Ariu, K., Ryu, N., Yun, S.-Y., and Proutière, A. Regret in online recommendation systems. Advances in Neural Information Processing Systems, 33:21141-21150, 2020.

Bell, R. M. and Koren, Y. Scalable collaborative filtering with jointly derived neighborhood interpolation weights. In Seventh IEEE international conference on data mining (ICDM 2007), pp. 43-52. IEEE, 2007.

Boucheron, S., Lugosi, G., and Massart, P. Concentration inequalities: A nonasymptotic theory of independence. Oxford university press, 2013.

Bresler, G. and Karzand, M. Regret bounds and regimes of optimality for user-user and item-item collaborative filtering. IEEE Transactions on Information Theory, 67 (6):4197-4222, 2021.

Brunskill, E. and Li, L. Sample complexity of multi-task reinforcement learning. arXiv preprint arXiv:1309.6821, 2013.

Cammardella, N., Bušić, A., and Meyn, S. Kullback-leibler-quadratic optimal control. arXiv preprint arXiv:2004.01798, 2020.

Candès, E. J. and Tao, T. The power of convex relaxation: Near-optimal matrix completion. IEEE Transactions on Information Theory, 56(5):2053-2080, 2010.

Chen, Y., Fan, J., Ma, C., and Yan, Y. Inference and uncertainty quantification for noisy matrix completion. Proceedings of the National Academy of Sciences, 116(46): 22931-22937, 2019.

Chen, Y., Chi, Y., Fan, J., Ma, C., and Yan, Y. Noisy matrix completion: Understanding statistical guarantees for convex relaxation via nonconvex optimization. SIAM journal on optimization, 30(4):3098-3121, 2020.

Chi, Y., Lu, Y. M., and Chen, Y. Nonconvex optimization meets low-rank matrix factorization: An overview. IEEE Transactions on Signal Processing, 67(20):5239-5269, 2019.

Dann, C. and Brunskill, E. Sample complexity of episodic fixed-horizon reinforcement learning. Advances in Neural Information Processing Systems, 28, 2015.

D’Eramo, C., Tateo, D., Bonarini, A., Restelli, M., Peters, J., et al. Sharing knowledge in multi-task deep reinforcement learning. In 8th International Conference on Learning Representations, pp. 1-11. OpenReview. net, 2020.

Dubhashi, D. P. and Ranjan, D. Balls and bins: A study in negative dependence. BRICS Report Series, 3(25), 1996.
Fazel, M. Matrix rank minimization with applications. PhD thesis, PhD thesis, Stanford University, 2002.

Gheshlaghi Azar, M., Munos, R., and Kappen, H. J. Minimax pac bounds on the sample complexity of reinforcement learning with a generative model. Machine learning, 91(3):325-349, 2013.

Gleich, D. F. and Lim, L.-h. Rank aggregation via nuclear norm minimization. In Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining, pp. 60-68, 2011.

Hessel, M., Soyer, H., Espeholt, L., Czarnecki, W., Schmitt, S., and van Hasselt, H. Multi-task deep reinforcement learning with popart. In Proceedings of the AAAI Conference on Artificial Intelligence, volume 33, pp. 3796-3803, 2019.

Hsieh, C.-J., Chiang, K.-Y., and Dhillon, I. S. Low rank modeling of signed networks. In Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining, pp. 507-515, 2012.

Hu, J., Chen, X., Jin, C., Li, L., and Wang, L. Near-optimal representation learning for linear bandits and linear rl. In International Conference on Machine Learning, pp. 4349-4358. PMLR, 2021.

Huleihel, W., Pal, S., and Shayevitz, O. Learning user preferences in non-stationary environments. In International Conference on Artificial Intelligence and Statistics, pp. 1432-1440. PMLR, 2021.

Jain, P. and Pal, S. Online low rank matrix completion. arXiv preprint arXiv:2209.03997, 2022.

Jain, P., Netrapalli, P., and Sanghavi, S. Low-rank matrix completion using alternating minimization. In Proceedings of the forty-fifth annual ACM symposium on Theory of computing, pp. 665-674, 2013.

Jain, P., Kar, P., et al. Non-convex optimization for machine learning. Foundations and Trends® in Machine Learning, 10(3-4):142-363, 2017.

Jin, C., Krishnamurthy, A., Simchowitz, M., and Yu, T. Reward-free exploration for reinforcement learning. In International Conference on Machine Learning, pp. 48704879. PMLR, 2020a.

Jin, C., Yang, Z., Wang, Z., and Jordan, M. I. Provably efficient reinforcement learning with linear function approximation. In Conference on Learning Theory, pp. 2137-2143. PMLR, 2020b.

Lazaric, A. Transfer in reinforcement learning: a framework and a survey. In Reinforcement Learning, pp. 143-173. Springer, 2012.

Lei, Y. and Li, W. When collaborative filtering meets reinforcement learning. arXiv preprint arXiv:1902.00715, 2019.

Mate, A., Madaan, L., Taneja, A., Madhiwalla, N., Verma, S., Singh, G., Hegde, A., Varakantham, P., and Tambe,
M. Field study in deploying restless multi-armed bandits: Assisting non-profits in improving maternal and child health. In Proceedings of the AAAI Conference on Artificial Intelligence, volume 36, pp. 12017-12025, 2022.

Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al. Human-level control through deep reinforcement learning. nature, 518(7540): 529-533, 2015.

Modi, A., Jiang, N., Singh, S., and Tewari, A. Markov decision processes with continuous side information. arXiv preprint arXiv:1711.05726, 2017.

Negahban, S. and Wainwright, M. J. Estimation of (near) low-rank matrices with noise and high-dimensional scaling. The Annals of Statistics, 39(2):1069-1097, 2011.

Negahban, S., Yu, B., Wainwright, M. J., and Ravikumar, P. A unified framework for high-dimensional analysis of $m$-estimators with decomposable regularizers. Advances in neural information processing systems, 22, 2009.

Nguyen-Thanh, N., Marinca, D., Khawam, K., Rohde, D., Vasile, F., Lohan, E. S., Martin, S., and Quadri, D. Recommendation system-based upper confidence bound for online advertising. arXiv preprint arXiv:1909.04190, 2019.

Papadimitriou, C. H. and Tsitsiklis, J. N. The complexity of optimal queuing network control. Mathematics of Operations Research, 24(2), 1999.

Recht, B. A simpler approach to matrix completion. Journal of Machine Learning Research, 12(12), 2011.

Sam, T., Chen, Y., and Yu, C. L. Overcoming the long horizon barrier for sample-efficient reinforcement learning with latent low-rank structure. arXiv preprint arXiv:2206.03569, 2022.

Shah, D., Song, D., Xu, Z., and Yang, Y. Sample efficient reinforcement learning via low-rank matrix estimation. Advances in Neural Information Processing Systems, 33: 12092-12103, 2020.

Sodhani, S., Zhang, A., and Pineau, J. Multi-task reinforcement learning with context-based representations. In International Conference on Machine Learning, pp. 9767-9779. PMLR, 2021.

Sutton, R. S., Barto, A. G., and Williams, R. J. Reinforcement learning is direct adaptive optimal control. IEEE control systems magazine, 12(2):19-22, 1992.

Taylor, M. E. and Stone, P. Transfer learning for reinforcement learning domains: A survey. Journal of Machine Learning Research, 10(7), 2009.

Teh, Y., Bapst, V., Czarnecki, W. M., Quan, J., Kirkpatrick, J., Hadsell, R., Heess, N., and Pascanu, R. Distral: Robust multitask reinforcement learning. Advances in neural information processing systems, 30, 2017.

Vershynin, R. High-dimensional probability: An introduction with applications in data science, volume 47. Cambridge university press, 2018.

Wagenmaker, A. J., Chen, Y., Simchowitz, M., Du, S., and Jamieson, K. Reward-free rl is no harder than rewardaware rl in linear markov decision processes. In International Conference on Machine Learning, pp. 2243022456. PMLR, 2022.

Zhang, Z., Du, S. S., and Ji, X. Nearly minimax optimal reward-free reinforcement learning. arXiv preprint arXiv:2010.05901, 2020.

## A. More Discussion Regarding Policy Space

## A.1. Necessity of Randomized Policies

We will first show that randomized policies might be necessary in such contexts with a simple example and show that obtaining states which satisfy conditions like (4) goes beyond simple reward maximization. Suppose $H=1, \mathcal{S}=\{1\}$ and $\mathcal{A}=\{1, \ldots, d\}$. We consider the embedding $\psi(s, a)=e_{a}$. Suppose we want to obtain a policy $\pi$ such that whenever $S_{1}, A_{1} \sim \pi$, $\lambda_{\text {min }}\left(\mathbb{E} \psi\left(S_{1}, A_{1}\right) \psi\left(S_{1}, A_{1}\right)^{\top}\right)$ is maximized (where $\lambda_{\text {min }}$ denotes the minimum eigenvalue). This is maximized when $\pi(d a \mid s)$ is chosen to be the uniform distribution over $\mathcal{A}$ and the corresponding value is $1 / d$. Note that whenever $\pi$ is a deterministic policy we will have $\lambda_{\min }=0$ whenever $d>1$. This is in contrast to reward maximization problems where, under general conditions, a deterministic optimal policy exists (See Theorem 1.7 in (Agarwal et al., 2019)).
If fact, we can also show that the policy which minimizes $\left\|\mathbb{E} \psi\left(S_{1}, A_{1}\right)-\frac{1}{d} \sum_{a=1}^{d} e_{a}\right\|$ must also necessarily be random.
In the case of linear MDPs, we can find such a deterministic optimal policy $\Pi=\left(\pi_{1}, \ldots, \pi_{H}\right)$ as $\pi_{h}(s)=$ $\arg \sup _{a}\left\langle\psi(s, a), u_{h}^{*}\right\rangle+\left\langle\phi(s, a), v_{h}^{*}\right\rangle$ (Jin et al., 2020b). This reduces the problem to estimating the parameters $u_{h}^{*}, v_{h}^{*}$ even when the state-action space is an infinite set. However, when such policies are not guaranteed to exist, as in case of functional maximization required in Section 7 , the set of all policies can be intractably large. This is the justification for picking a nice enough policy space denoted by $\mathbb{Q}$.

## A.2. Constructing Policy Spaces

We consider any linear MDP satisfying the definition given in Section 2 and suppose $\mathcal{A}$ is finite. We consider the set of all probability distributions $\pi_{h}(a \mid s ; u, v) \propto \exp (\langle\phi(s, a), u\rangle+\langle\psi(s, a), v\rangle)$. We consider $\mathbb{Q}_{h}=\left\{\pi_{h}(a \mid s, u, v): u, v \in\right.$ $\left.\mathcal{B}_{d}(R)\right\}$, We let our policy space be $\mathbb{Q}=\left\{\Pi=\left(\pi_{1}, \ldots, \pi_{H}\right): \pi_{h} \in \mathbb{Q}_{h}\right\}$.
Lemma A.1. Consider the probability distribution over a finite set $[|\mathcal{A}|]$ give by $p_{\beta}(a) \propto \exp \left(\beta x_{a}\right)$ for every $a \in[|\mathcal{A}|]$ some $x_{a} \in \mathbb{R}^{+}$and $\beta \in \mathbb{R}^{+}$. For any $\epsilon>0$ and random variable $A \sim p_{\beta}$, we must have:

$$
\mathbb{P}\left(x_{A}<\sup _{a} x_{a}-\epsilon\right) \leq|\mathcal{A}| \exp (-\beta \epsilon)
$$

And

$$
\mathbb{E} x_{A} \geq\left(\sup _{a} x_{a}-\epsilon\right)(1-|\mathcal{A}| \exp (-\beta \epsilon))
$$

Lemma A.2. Let $Q_{h}^{*}(s, a)$ be the optimal action-value function for the MDP. Then the policy $\Pi=\left(\pi_{1}, \ldots, \pi_{h}\right)$ given by $\pi_{h}(a \mid s) \propto \exp \left(\beta Q_{h}^{*}(s, a)\right)$ is $\epsilon H+H^{2}|\mathcal{A}| \exp (-\beta \epsilon)$ sub-optimal for any $\epsilon>0$

Proof. Consider the optimal value function defined by $V_{h}^{*}(s)=\sup _{a} Q_{h}^{*}(s, a)$. Let $\bar{Q}_{h}(s, a)$ denote the optimal action value function under the policy $\Pi$ and let $\bar{V}(s)=\int \bar{Q}_{h}(s, a) \pi_{h}(d a \mid s)$ denote the value at state $s$ with the policy $\Pi$. Clearly, we have: $\bar{Q}_{H}(s, a)=Q_{H}^{*}(s, a)=R(s, a) . \bar{Q}_{h}(s, a) \geq Q_{h}^{*}(s, a)-\eta$ uniformly. Then we have

$$
\begin{align*}
\bar{V}_{h}(s) & \geq \int Q_{h}^{*}(s, a) \pi_{h}(d a \mid s)-\eta \\
& \left.\geq \sup _{a} Q_{h}^{*}(s, a)-\epsilon\right)(1-|\mathcal{A}| \exp (-\beta \epsilon))-\eta \\
& \geq \sup _{a} Q_{h}^{*}(s, a)-\epsilon-H|\mathcal{A}| \exp (-\beta \epsilon)-\eta=V_{h}^{*}(s)-\epsilon-H|\mathcal{A}| \exp (-\beta \epsilon)-\eta \tag{5}
\end{align*}
$$

In the second step, we have invoked Lemma A.1. In the last step, we have used the fact that $Q_{h}^{*} \in[0, H]$ uniformly. Now, by the Bellman iteration, we have:

$$
\begin{align*}
Q_{h-1}^{*}(s, a) & =R_{h-1}(s, a)+\mathbb{E}_{s^{\prime} \sim P_{h-1}(\mid s, a)} V_{h}^{*}\left(s^{\prime}\right) \\
& \geq R_{h-1}(s, a)+\mathbb{E}_{s^{\prime} \sim P_{h-1}(\mid s, a)} \bar{V}_{h}\left(s^{\prime}\right)-\epsilon-H|\mathcal{A}| \exp (-\beta \epsilon)-\eta \\
& =\bar{Q}_{h-1}(s, a)-\epsilon-H|\mathcal{A}| \exp (-\beta \epsilon)-\eta \tag{6}
\end{align*}
$$

Therefore, by induction, we conclude that $\bar{V}_{1}(s) \geq V_{1}^{*}(s)-\epsilon H-H^{2}|\mathcal{A}| \exp (-\beta \epsilon)$

Therefore, by the definition of the value function, we conclude the claim.

Now, by a simple extension of Proposition 2.3 in (Jin et al. 2020b), we conclude that the optimal $Q_{h}^{*}$ function for any linear MDP can be written as:

$$
Q_{h}^{*}(s, a)=\left\langle\psi(s, a), u_{h}^{*}\right\rangle+\left\langle\phi(s, a), v_{h}^{*}\right\rangle
$$

Where $\left\|u_{h}^{*}\right\|_{2} \leq \sqrt{d}$ and $\left\|v_{h}^{*}\right\|_{\infty} \leq H C_{\mu}$. Observe that choosing $\epsilon=\frac{\eta}{2 H}$ and $\beta=2 \frac{\log (2 H|\mathcal{A}| / \eta)}{\eta}$ will ensure that the randomized policy $\Pi$ in the statement of Lemma A.2 is $\eta$ optimal. Therefore, we can take $R=2 d H C_{\mu} \frac{\log (2 H|\mathcal{A}| / \eta)}{\eta}$ in the definition of $\mathbb{Q}_{h}$ above and conclude that this includes every $\eta$ optimal policy for every MDP with embedding functions $\phi, \psi$. We will now bound the covering number. Recall the definition of the distance $D_{\mathbb{Q}}\left(\Pi_{1}, \Pi_{2}\right)=\sup _{h \in[H]} \mathrm{TV}\left(\pi_{h}^{(1)}, \pi_{h}^{(2)}\right)$. Therefore it is sufficient to obtain an $\eta$ cover for $\mathbb{Q}_{h}$ (denoted by $\hat{\mathbb{Q}}_{h, \eta}$ ) and then construct $\hat{\mathbb{Q}}_{\eta}=\left\{\Pi=\left(\pi_{1}, \ldots, \pi_{H}\right): \pi_{h} \in\right.$ $\left.\hat{\mathbb{Q}}_{h, \eta} \forall h \in[H]\right\}=\prod_{h=1}^{H} \hat{\mathbb{Q}}_{h, \eta}$.
Lemma A.3. $\pi(\mid s ; u, v)$ be as defined in the beginning of this Subsection.

$$
\operatorname{TV}\left(\pi(\cdot \mid s ; u, v), \pi\left(\cdot \mid s ; u^{\prime}, v^{\prime}\right)\right) \leq \frac{1}{2}\left(\exp \left(2\left\|u-u^{\prime}\right\|_{2}+2\left\|v-v^{\prime}\right\|_{\infty}\right)-1\right)
$$

Proof. Denote $\pi(a \mid s ; u, v)$ by $\pi(a)$ and $\pi\left(a \mid s ; u^{\prime}, v^{\prime}\right)$ by $\pi^{\prime}(a)$. Consider the corresponding partition functions denoted by $Z:=\sum_{a \in \mathcal{A}} \exp (\langle\psi(s, a), u\rangle+\langle\phi(s, a), v\rangle)$ and $Z^{\prime}:=\sum_{a \in \mathcal{A}} \exp \left(\left\langle\psi(s, a), u^{\prime}\right\rangle+\left\langle\phi(s, a), v^{\prime}\right\rangle\right)$. We conclude that using Hölder's inequality for $\left\langle u-u^{\prime}, \psi\right\rangle$ and $\left\langle v-v^{\prime}, \phi\right\rangle$ that:

$$
\begin{align*}
& \exp \left(-\left\|u-u^{\prime}\right\|_{2}-\left\|v-v^{\prime}\right\|_{\infty}\right) \leq \frac{Z^{\prime}}{Z} \leq \exp \left(\left\|u-u^{\prime}\right\|_{2}+\left\|v-v^{\prime}\right\|_{\infty}\right) \\
& \exp \left(-2\left\|u-u^{\prime}\right\|_{2}-2\left\|v-v^{\prime}\right\|_{\infty}\right) \leq \frac{\pi^{\prime}(a)}{\pi(a)} \leq \exp \left(2\left\|u-u^{\prime}\right\|_{2}+2\left\|v-v^{\prime}\right\|_{\infty}\right)  \tag{7}\\
& \operatorname{TV}\left(\pi, \pi^{\prime}\right)=\frac{1}{2} \sum_{a \in \mathcal{A}}\left|\pi(a)-\pi^{\prime}(a)\right| \\
&=\frac{1}{2} \sum_{a \in \mathcal{A}} \pi(a)\left|1-\frac{\pi^{\prime}(a)}{\pi(a)}\right| \\
& \leq \frac{1}{2}\left(\exp \left(2\left\|u-u^{\prime}\right\|_{2}+2\left\|v-v^{\prime}\right\|_{\infty}\right)-1\right) \tag{8}
\end{align*}
$$

Using the lemma above, we conclude that $\hat{\mathbb{Q}}_{h, \eta}=\left\{\pi_{h}(\mid s ; u, v): u, v \in \hat{\mathcal{B}}_{d, \eta / 4}(R)\right\}$ whenever $\eta \leq 1$. Here $\hat{\mathcal{B}}_{d, \eta / 4}(R)$ an $\eta / 4$ net over $\mathcal{B}_{d}(R)$ with respect to the norm $\|\cdot\|_{2}$. From the results in (Vershynin, 2018), we can therefore take:

$$
\begin{equation*}
\left|\hat{\mathbb{Q}}_{h, \eta}\right| \leq\left|\hat{\mathcal{B}}_{d, \eta / 4}(R)\right|^{2} \leq \exp \left(C d \log \left(\frac{C \eta}{R}\right)\right) \tag{9}
\end{equation*}
$$

Since we had $\hat{\mathbb{Q}}_{\eta}=\prod_{h=1}^{H}$, we conclude that:

$$
\log \left(\left|\hat{\mathbb{Q}}_{\eta}\right|\right) \leq c d H\left(\log \frac{d H}{\eta}+\log \log (2 H|\mathcal{A}| / \eta)\right)
$$

## B. Analysis - Tabular MDPs

We will call the reward free RL procedure in Phase 1 to be successful if it outputs the $\epsilon$ optimal policy. This has probability atleast $1-\frac{\delta}{2}$.

## B.1. Analysis of Algorithm 1

Lemma B.1. Suppose $p \leq \frac{1}{2}$, conditioned on the success of Phase 1 , with probability at-least $1-\exp (-c N p|\mathcal{S} \| \mathcal{A}| H)$, Algorithm 1 terminates after querying $\frac{C|\mathcal{S} \| \mathcal{A}| N H p}{\epsilon}$ trajectories. $\left(\mathcal{G}_{h}\right)_{h \in[H]}$, the active sets at the termination of the algorithm. They satisfy:

$$
\begin{equation*}
\sup _{\pi} \sum_{h=1}^{H} P_{h}^{\pi}\left(\mathcal{G}_{h}\right) \leq \frac{5 \epsilon}{8} \tag{10}
\end{equation*}
$$

For any $a \times b$ matrix $R$, let $\hat{R}$ be its partially observed version (that is, there exists a set of indices $I \subseteq[a] \times[b]$ such that $\hat{R}_{i j}=R_{i j}$ if $(i, j) \in I$ and $\hat{R}_{i j}=*$ otherwise $)$. We call a random set of indices $J$ to have the distribution Unif $(m,[a],[b])$ if $J$ is drawn uniformly at random such that $|J|=m$.
Lemma B. 2 (Modification: Mod1). Suppose we run, independently, a modification of algorithm 1 where on the "Query trajectory" step the trajectories are sampled from a fixed MDP $\mathcal{M}_{1}$ (but rewards are from the reward function corresponding to $\left.U_{t}\right)$. Consider all the random variables that determine the trajectory of this algorithm: $\left(\hat{V}, \hat{\Pi},\left(S_{1: H}^{(t)}, A_{1: H}^{(t)}, R_{1: H}^{(t)}\right)_{t},\left(U_{t}\right)_{t}\right)$. Then the joint distribution of this collection of random variables is unchanged under the modification.

Proof. The proof follows from an induction argument on the time index $t$. We describe the key steps here. For the ease of notation, let $\mathcal{T}_{t}=\left(\left(S_{1: H}^{(t)}, A_{1: H}^{(t)}, R_{1: H}^{(t)}\right), U_{t}\right)$ Let $X_{T}=\left(\hat{V}, \hat{\Pi},\left(\mathcal{T}_{t}\right)_{t \leq T}\right)$. Let $\tilde{X}_{T}$ and $\tilde{\mathcal{T}}_{t}$ denote the corresponding quantities under the modification. It is enough to show that finite dimensional marginals have the same joint distribution under the modification. In particular, we will show:

## 1. $\mathcal{T}_{0} \stackrel{d}{=} \tilde{\mathcal{T}}_{0}$

2. Suppose $X_{T} \stackrel{d}{=} \tilde{X}_{T}$. Then the Markov kernel $k_{\mathcal{T}_{T+1} \mid X_{T}}$ is almost surely (under the common distribution of $X_{T}, \tilde{X}_{T}$ ) equal to $k_{\tilde{\mathcal{T}}_{T+1} \mid \tilde{X}_{T}}$. Thus $X_{T+1} \stackrel{d}{=} \tilde{X}_{T+1}$

The first statement is straightforward since, in the zeroth step, the distribution of $\hat{\Pi}^{\mathcal{G}}$ not affected by the modification, and thus due to identical MDP transitions across users, the distribution of $\mathcal{T}_{0}$ is preserved under modification. A similar argument proves the second statement. Roughly, given a realization of $X_{T}$ the distribution of $\mathcal{T}_{T+1}$ is same as the distribution of $\tilde{\mathcal{T}}_{T+1}$ given the same realization of $\tilde{X}_{T}$, due to the exact same reason presented for the first statement. A fully formal proof requires setting up appropriate proability spaces, so we omit it here. Furthermore, since the random variables considered are all discrete, one can argue via PMFs as well.
Lemma B.3. Suppose $p \leq \frac{1}{2}$. conditioned on the success of Phase 1 and termination of Algorithm 1$]$ for every $h \in[H]$, the Algorithm $\left\lfloor\right.$ returns partially filled reward matrices $\hat{R}_{h}$. Consider the sub-matrix $\hat{R}_{h}^{\mathcal{G}_{h}^{\mathbb{C}}}$. Let $I_{h} \subseteq[N] \times \mathcal{G}_{h}^{\mathbb{C}}$ be the sub-set of observed indices for $\hat{R}_{h}$. Let $J_{h} \left\lvert\, \mathcal{G}_{h} \sim \operatorname{Unif}\left(\frac{N p\left|\mathcal{G}_{h}^{\mathrm{C}}\right|}{2},[N], \mathcal{G}_{h}^{\mathrm{C}}\right)\right.$. There exists a coupling between $J_{h}$ and $I_{h}$ such that

$$
\mathbb{P}\left(J_{h} \subseteq I_{h} \mid \mathcal{G}_{h}\right) \geq 1-|\mathcal{S}||\mathcal{A}| \exp (-c N p)
$$

Proof. Let us fix $\mathcal{G}_{h}$ and construct a coupling between $I_{h}$ and $J_{h}$.
Consider any fixed, arbitrary permutations $\sigma_{g}$ over $[N]$, for $g \in \mathcal{G}_{h}^{C}$. By $\sigma\left(I_{h}\right)$, we denote $\left\{\left(\sigma_{g}(i), g\right):(i, g) \in I_{h}\right\}$.
Claim B.4. Conditioned on $\mathcal{G}_{h}, \sigma\left(I_{h}\right)$ has the same distribution as $I_{h}$.
Proof. Let $\left\{\sigma_{(s, a):[N] \rightarrow[N]} \mid(s, a) \in \mathcal{S} \times \mathcal{A}\right\}$ be a set of arbitrary permutations on $[N]$. From lemma B.2 it is enough to prove the statement for the random variables under the modification described in that lemma (call this Mod1). Now consider a further modification (call it Mod2) where in every iteration $t$, we sample $U_{t} \sim \operatorname{Unif}([N])$, for each horizon $h$ we set $\tilde{U}_{h}^{(t)}=\sigma_{\left(S_{h}^{(t)}, A_{h}^{(t)}\right)}\left(U_{t}\right)$, and then update the entries of $R_{h}\left(\tilde{U}_{h}^{(t)},\left(S_{h}^{(t)}, A_{h}^{(t)}\right)\right)\left(\right.$ instead of $R_{h}\left(U_{t},\left(S_{h}^{(t)}, A_{h}^{(t)}\right)\right)$ ). Next, we couple these two modifications by using same $(\hat{V}, \hat{\Pi})$ and the same set of $U_{t}$ 's for both the modifications. Further, we couple the MDP used in these modifications to be the same, single MDP.

Now an induction argument shows that the sequence of active sets $\mathcal{G}$ obtained in these modifications are also identical for every time $t$; only the rows of $R_{h}$ where entries are filled change according to the set of permutations chosen.Thus, under the described coupling, Mod1 and Mod2 produce identical trajectories (i.e., $\left(S_{1: H}^{(t)}, A_{1: H}^{(t)}\right)$ ), the columns of reward matrices are just permutations of each other described by the chosen set of permutations, and algorithm 1 terminate at the same time in both these cases. However, the same induction argument also shows that for each $t$ and $h$, conditioned on $\mathcal{G}_{h}$, trajectories (which is same in Mod1 and Mod2) until at beginning of iteration $t$, we have $\left(U_{t},\left(S_{h}^{(t)}, A_{h}^{(t)}\right)\right) \stackrel{d}{=}\left(\tilde{U}_{h}^{(t)},\left(S_{h}^{(t)}, A_{h}^{(t)}\right)\right)$.
Therefore if $I_{h}, \tilde{I}_{h} \subset[N] \times \mathcal{G}_{h}^{C}$ denotes the subset of observed indices at termination (outside the active set), then $\tilde{I}_{h}=\sigma\left(I_{h}\right) \equiv\left\{\left(\sigma_{(s, a)}(i),(s, a)\right):\left(i,(s, a) \in I_{h}\right)\right\}$ and, conditioned on $\mathcal{G}_{h}, \tilde{I}_{h} \stackrel{d}{=} I_{h}$

Claim B.5. At termination, conditioned on $\mathcal{G}_{h}$, random sets $I_{h}^{g}=\left\{(i, g):(i, g) \in I_{h}\right\}$ are jointly independent.

Proof. Again we work with the modification Mod1 described in lemma B. 2 . For each $(s, a)$ consider the collection of $U_{t}$ 's that are used populate the column $(s, a)$ of matrix $R_{h}$ in algorithm 1 . Call this collection $\mathcal{U}_{(s, a)}$.

Remark B.6. Since the columns of $I_{h}$ have exactly $N p$ entries, permutation invariance proved in the above claim implies that

For any set $\bar{J} \subseteq[N] \times \mathcal{G}_{h}^{\complement}$, define the count function $\left(N_{g}^{\bar{J}}\right)_{g \in \mathcal{G}_{h}^{\text {© }}}$ such that $N_{g}^{\bar{J}}=|\{i \in[N]:(i, g) \in \bar{J}\}|$.
We are now ready to give the coupling: given $\mathcal{G}_{h}$, draw uniformly random, independent permutations $\sigma_{g}$ for $g \in \mathcal{G}_{h}^{\complement}$. Draw $\left(N_{g}\right)_{g \in \mathcal{G}_{h}^{\mathrm{C}}}$ independent of $\sigma_{g}$ and to have the joint law of $\left(N_{g}^{J_{h}}\right)_{g \in \mathcal{G}_{h}^{\mathrm{c}}}$. Define:

$$
\begin{aligned}
& \tilde{J}_{h}=\left\{\left(\sigma_{g}(i), g\right): i \leq N_{g}, g \in \mathcal{G}_{h}^{\complement}\right\} \\
& \tilde{I}_{h}=\left\{\left(\sigma_{g}(i), g\right): i \leq N p, g \in \mathcal{G}_{h}^{\complement}\right\}
\end{aligned}
$$

Claim B.7. The marginal distributions of $\tilde{J}_{h}$ and $\tilde{I}_{h}$ are respectively the distributions of $J_{h}$ and $I_{h}$.

Proof. First we will prove a general statement about $J \sim \operatorname{Unif}(r,[N],[M])$. Let $X \in\{0,1\}^{N \times M}$ with $X_{i, m}=1$ iff $(i, m) \in J$. Let $\left(N_{m}\right)_{m \in[M]}$ be the count functions corresponding to $J$ i.e., $N_{m}=\sum_{i} X_{i, m}$. Let $Y_{m}=\left(X_{1, m}, \cdots, X_{N, m}\right)$. We will argue that conditional on $\left\{N_{m}: m \in[M]\right\}$, the random vectors $Y_{m}$ are jointly independent. Indeed, pick any $x \in\{0,1\}^{N \times M}$ and $\left(n_{1}, \cdots, n_{m}\right)$. Let $y_{m}$ be the $m$ 'th column of $x$. Then

$$
\mathbb{P}\left[X=x, \cap_{m}\left\{N_{m}=n_{m}\right\}\right]=\left(\prod_{m} 1\left[\sum_{i} x_{i, m}=n_{m}\right]\right) 1\left[\sum_{i, m} x_{i, m}=r\right] \frac{1}{\binom{M_{r}}{r}}
$$

The above can also be written as

$$
\mathbb{P}\left[\cap_{m}\left\{Y_{m}=y_{m}\right\}, \cap_{m}\left\{N_{m}=n_{m}\right\}\right]=\left(\prod_{m} 1\left[\sum_{i} x_{i, m}=n_{m}\right]\right) 1\left[\sum_{m} n_{m}=r\right] \frac{1}{\binom{M N}{r}}
$$

Let 1 denote the all 1 vector in $\mathbb{R}^{N}$. Note that $y_{m}=\left(x_{1, m}, \cdots, x_{N, m}\right)^{\top}$. Marginalizing the above, we see that

$$
\mathbb{P}\left[\cap_{m}\left\{N_{m}=n_{m}\right\}\right]=\left(\prod_{m} \frac{1}{\binom{N}{n_{m}}}\right) 1\left[\sum_{m} n_{m}=r\right] \frac{1}{\binom{M N}{r}}
$$

Thus the conditional distribution can be expressed as

$$
\mathbb{P}\left[\cap_{m}\left\{Y_{m}=y_{m}\right\} \mid \cap_{m}\left\{N_{m}=n_{m}\right\}\right]=\left(\prod_{m} \frac{1\left[\mathbf{1}^{\top} y_{m}=n_{m}\right]}{\binom{N}{n_{m}}}\right) \frac{1\left[\sum_{m} n_{m}=r\right]}{\binom{M N}{r}}
$$

Since the (conditional) joint PMF factors, it is an easy calculation to show the conditional independence i.e.,

$$
\mathbb{P}\left[\cap_{m}\left\{Y_{m}=y_{m}\right\} \mid \cap_{m}\left\{N_{m}=n_{m}\right\}\right]=\prod_{m} \mathbb{P}\left[Y_{m}=y_{m} \mid \cap_{m}\left\{N_{m}=n_{m}\right\}\right]
$$

Furthermore, for any $n_{1}, \cdots n_{m}$ such that $\sum_{m} n_{m}=r$, marginalization shows

$$
\mathbb{P}\left[Y_{m}=y_{m} \mid \cap_{m}\left\{N_{m}=n_{m}\right\}\right]=\frac{1\left[\mathbf{1}^{\top} y_{m}=n_{m}\right]}{\binom{N}{n_{m}}}
$$

Let $N_{-m}=\left(N_{1}, \cdots, N_{m-1}, N_{m+1}, \cdots, N_{M}\right)$, and similarly for $n_{-m}$. Then

$$
\begin{aligned}
& \mathbb{P}\left[Y_{m}=y_{m}, N_{-m}=n_{-m} \mid N_{m}=n_{m}\right] \\
& =\mathbb{P}\left[Y_{m}=y_{m} \mid \cap_{m}\left\{N_{m}=n_{m}\right\}\right] \mathbb{P}\left[N_{-m}=n_{-m} \mid N_{m}=n_{m}\right] \\
& = \begin{cases}0, & \sum_{m} n_{m} \neq r \\
\frac{1\left[\mathbf{1}^{\top} y_{m}=n_{m}\right]}{\binom{N}{n_{m}}} \mathbb{P}\left[N_{-m}=n_{-m} \mid N_{m}=n_{m}\right], & \text { otherwise }\end{cases}
\end{aligned}
$$

The above factorization directly implies that $Y_{m}$, conditioned on $N_{m}$ is uniformly distributed on its support $\left\{y: \mathbf{1}^{\top} y=N_{m}\right\}$ and is independent of $N_{-m}$. Thus

$$
\begin{aligned}
& \mathbb{P}\left[\cap_{m}\left\{Y_{m}=y_{m}\right\} \mid \cap_{m}\left\{N_{m}=n_{m}\right\}\right] \\
& =\prod_{m} \mathbb{P}\left[Y_{m}=y_{m} \mid N_{m}=n_{m}\right] \\
& =\prod_{m} \frac{1\left[\mathbf{1}^{\top} y_{m}=n_{m}\right]}{\binom{N}{n_{m}}}
\end{aligned}
$$

Observation: The above calculations give another way to generate $Y$ : first generate $N_{1}, \cdots, N_{M}$ from the right distribution, and then conditioned on $N_{m}$ generate each $Y_{m}$ uniformly such that $\mathbf{1}^{\top} Y_{m}=N_{m}$.
Next we apply the above calculations and observation to $J=J_{h} \left\lvert\, \mathcal{G}_{h} \sim \operatorname{Unif}\left(\frac{N p\left|\mathcal{G}_{h}^{\mathrm{C}}\right|}{2},[N], \mathcal{G}_{h}^{\mathrm{C}}\right)\right.$. For a uniformly random permutation $\sigma$ on $[N]$, the set $\{\sigma(i): 1 \leq i \leq k\}$ is uniformly distributed on all $k$-sized subsets of $[N]$. In the statement of the claim the permutations are chosen independently for each $g \in \mathcal{G}_{c}^{\complement}$. Thus from the above observation, we have $J_{h} \stackrel{d}{=} \tilde{J}_{h}$ conditioned on $\mathcal{G}_{h}$.
The claim about $\tilde{I}_{h}$ follows directly from permutation invariance proved by claim B. 4 .

Claim B.8. $\mathbb{P}\left(N_{g}>N p \mid \mathcal{G}_{h}\right) \leq \exp \left(-c_{0} N p\right)$ for every $g \in \mathcal{G}_{h}^{\complement}$
Proof. Throughout this proof, we will condition on the terminal active set $\mathcal{G}_{h}$. We will show this using the results on concentration with negative regression property as established in Proposition 29 in (Dubhashi \& Ranjan, 1996). $N_{g}=$ $\sum_{i=1}^{N} \mathbb{1}\left((i, g) \in \tilde{J}_{h}\right)$. Now we will show that the collection $X_{i g}:=\mathbb{1}\left((i, g) \in \tilde{J}_{h}\right)$ for $i \in[N], g \in \mathcal{G}_{h}^{C}$ satisfy the negative regression property. By the definition of negative regression, we can conclude that the sub-collection $\left(X_{i g}\right)_{i \in[N]}$ also satisfies this property for every $g \in \mathcal{G}_{h}^{\mathrm{C}}$.
Consider the partial order over binary vectors $X \succeq Y$ iff $X_{l} \geq Y_{l}$ for every $l$. The negative regression property is satisfied iff for every $K_{1}, K_{2} \subseteq[N] \times \mathcal{G}_{h}^{\complement}$ such that $K_{1} \cap K_{2}=\emptyset$, and a real valued function $f\left(X_{m}: m \in K_{1}\right)$ which is non-decreasing with respect to the partial order, we must have:

$$
g\left(t_{l}: l \in K_{2}\right):=\mathbb{E}\left[f\left(X_{m}: m \in K_{1}\right) \mid X_{l}=t_{l}, \forall l \in K_{2}\right]
$$

be such that $g$ is a non-increasing function in $t_{l}$ with respect to the partial order. Note that in the case of uniform distribution as in $\tilde{J}_{h}$, the distribution $\left(X_{m}\right)_{m \in K_{1}}$ is the uniform, permutation invariant distribution with constant sum almost surely. The sum being $\frac{N p\left|\mathcal{G}_{h}^{\mathrm{C}}\right|}{2}-\sum_{l \in K_{2}} t_{l}$. Therefore, whenever $t_{l}^{\prime} \geq t_{l}$ for every $l \in K_{2}$, we have the following stochastic dominance:

$$
\left[\left(X_{m}\right)_{m \in K_{1}} \mid X_{l}=t_{l}^{\prime} \forall l \in K_{2}\right] \preceq\left[\left(X_{m}\right)_{m \in K_{1}} \mid X_{l}=t_{l} \forall l \in K_{2}\right]
$$

Therefore, this coupling leads us to conclude that:

$$
\begin{align*}
g\left(t_{l}: l \in K_{2}\right) & :=\mathbb{E}\left[f\left(X_{m}: m \in K_{1}\right) \mid X_{l}=t_{l} \forall l \in K_{2}\right] \\
& \geq \mathbb{E}\left[f\left(X_{m}: m \in K_{1}\right) \mid X_{l}=t_{l}^{\prime} \forall l \in K_{2}\right] \\
& =g\left(t_{l}^{\prime}: l \in K_{2}\right) \tag{11}
\end{align*}
$$

The second step follows from stochastic dominance. This implies that the function $g$ is non-increasing which establishes the negative regression property. Now, we consult Proposition 29 in (Dubhashi \& Ranjan, 1996) to show that we can take Chernoff bounds on $N_{g}=\sum_{i \in[N]} X_{i g}$ as though $X_{i g}$ were i.i.d $\operatorname{Ber}(p)$. Therefore, from an application of Bernstein's inequality (Boucheron et al. 2013), we conclude the statement of the claim.

Now, $J_{h} \subseteq I_{h}$ if and only if $N_{g} \leq N p$ for every $g \in \mathcal{G}_{h}^{\complement}$. Therefore, from the claim above, we have $\mathbb{P}\left(J_{h} \subseteq I_{h} \mid \mathcal{G}_{h}^{\complement}\right) \geq$ $1-|\mathcal{S}||\mathcal{A}| \exp \left(-c_{0} N p\right)$.

We are now ready to prove Theorem 5.3
Proof of Theorem 5.3 In order to establish the result, we need to show that with $p$ as set in the statement, the algorithm returns $\epsilon$ optimal policies $\hat{\Pi}_{u}$ for every user $u \in[N]$ with probability at-least $1-\delta$.

The total sample complexity is the number of trajectories queried in Phase 1 plus the number of trajectories queried in Phase 2. Phase 1 queries $K^{\text {rf }}\left(\frac{\epsilon}{8}, \frac{\delta}{2}\right)$ trajectories, which is $C \frac{|\mathcal{S} \| \mathcal{A}| H^{2}\left(|\mathcal{S}|+\log \left(\frac{1}{\delta}\right)\right)}{\epsilon^{2}} \operatorname{polylog}\left(\frac{|\mathcal{S} \| \mathcal{A}| H}{\epsilon}\right)$ by the results of (Zhang et al. 2020). By Lemma B.1. we conclude that the sample complexity of phase 2 is $\frac{C|\mathcal{S} \| \mathcal{A}| N H p}{\epsilon}$ and with the value of $p$ given in the statement of the theorem, this succeeds with probability at-least $1-\frac{\delta}{4}$ when conditioned on the success of Phase 1.
We will show that conditioned on the success of Phase 2 , with probability at-least $1-\frac{\delta}{4}$, the nuclear norm minimization algorithm of (Recht, 2011) successfully obtains $R_{h}^{\mathcal{G}_{h}^{\mathcal{C}}}$. Indeed by Theorem 1 in (Recht, 2011, we see that whenever co-ordinates of $R_{h}^{\mathcal{G}_{h}^{\mathrm{L}}}$ corresponding to random indices drawn from Unif $\left(m,[N], \mathcal{G}_{h}^{\mathrm{C}}\right)$ are observed with $m=C_{1} \max \left(\mu_{1}^{2}, \mu_{0}\right) r\left(N+\left|\mathcal{G}_{h}^{\mathrm{C}}\right|\right) \log ^{2}\left|\mathcal{G}_{h}^{\mathrm{C}}\right| \log \left(\frac{H}{\delta}\right)$, the algorithm succeeds at recovering $R_{h}^{\mathcal{G}_{h}^{\mathrm{C}}}$ with probability at-least $1-\frac{\delta}{8 H}$. The number of co-ordinates we observe is

$$
N p\left|\mathcal{G}_{h}^{\mathrm{C}}\right| \geq \frac{N p|\mathcal{S} \mathcal{A}|}{2} \geq 2 C_{1} \max \left(\mu_{1}^{2}, \mu_{0}\right) r\left(N+\left|\mathcal{G}_{h}^{\mathrm{C}}\right|\right) \log ^{2}\left|\mathcal{G}_{h}^{\mathrm{C}}\right| \log \left(\frac{H}{\delta}\right)
$$

In the last step, we have used Assumption 5.2 to conclude that $\left|\mathcal{G}_{h}^{\mathcal{C}}\right| \geq \frac{|\mathcal{S}||\mathcal{A}|}{2}$. For the constant $C$ in the definition of $p$ large enough, we must have:

$$
m \leq \frac{N p\left|\mathcal{G}_{h}^{\mathrm{C}}\right|}{2}
$$

Note that the results of (Recht, 2011) requires at-least $m$ observations to be chosen uniformly at random co-ordinates, but we do not obtain observations which are uniformly at uniformly random co-ordinates. Here, we will use the results of Lemma B.3. Let $J_{h}$ be a fictitious subset of co-ordinates distributed as $\operatorname{Unif}\left(m,[N], \mathcal{G}_{h}^{C}\right)$ when conditioned on $\mathcal{G}_{h}^{\complement}$. If the observed co-ordinates are $J_{h}$, then we can successfully estimate the reward matrix $R_{h}$ with proability at-least $1-\frac{\delta}{8 H}$ in this case. Now, suppose that the actually observed co-ordinates are $I_{h}$, which is a strict super-set of $J_{h}$. Then we check that the matrix completion algorithm, which is based on constrained nuclear-norm minimization, still succeeds with observed co-ordinates corresponding to $I_{h}$ whenever it succeeds with the observed co-ordinates correspond to $J_{h}$.

We now refer to the coupling in Lemma B.3. which shows that we can couple $J_{h}$ to the real distribution $I_{h}$ such that $J_{h} \subseteq I_{h}$ with probability at-least $1-\frac{\delta}{8 H}$ When the constant $C_{1}$ in the definition of $p$ is large enough, we conclude by invoking Lemma B. 3 that: $J_{h} \subseteq I_{h}$ with probability at-least $1-\frac{\delta}{8 H}$. Applying union bound for $h \in[H]$, we conclude that Phase 3 succeeds with probability at-least $1-\frac{\delta}{4}$ when conditioned on the success of Phases 1 and 2.

Therefore, from the arguments above, we conclude that Phases 1,2 and 3 succeed with probability at-least $1-\delta$ and give us the reward matrices $R_{h}^{\mathcal{G}_{h}^{\mathrm{C}}}$ where the sets satisfy the following equation from Lemma B. 1 .

$$
\begin{equation*}
\sup _{\pi} \sum_{h=1}^{H} P_{h}^{\pi}\left(\mathcal{G}_{h}\right) \leq \frac{5 \epsilon}{8} \tag{12}
\end{equation*}
$$

It now remains to show that we obtain $\epsilon$ optimal policies for each user after Phase 4. Note that whenever Phase 1 succeeds, we can compute $\epsilon / 4$ optimal policies for every possible reward function bounded in $[0,1]$. Since we do not know the rewards over the set $\mathcal{G}_{h}$, we set it to zero as described in the algorithm to obtain $\bar{R}_{h}$. It remains to show that planning with $\bar{R}_{h}$ and using it with the reward free RL algorithm gives us an $\epsilon$ optimal policy. Suppose $\Pi_{u}^{*}$ is the optimal policy for user $u$ and suppose $\bar{\Pi}_{u}$ be the optimal policy for user $u$ under rewards $\bar{R}_{h}(u,(s, a))$. Note that combined with the guarantees for the reward free RL, in order to complete the proof of the theorem, it is sufficient to show that the policy $\bar{\Pi}_{u}$ is $3 \epsilon / 4$ optimal with respect to the actual rewards $R_{h}(u,(s, a))$. Let $S_{1: H}^{*}, A_{1: H}^{*} \sim \mathcal{M}\left(\Pi_{u}^{*}\right)$ and $\bar{S}_{1: H}, \bar{A}_{1: H} \sim \mathcal{M}\left(\bar{\Pi}_{u}\right)$.

$$
\begin{align*}
& \mathbb{E} \sum_{h=1}^{H} R_{h}\left(u,\left(S_{h}^{*}, A_{h}^{*}\right)\right) \leq \mathbb{E} \sum_{h=1}^{H} R_{h}\left(u,\left(S_{h}^{*}, A_{h}^{*}\right)\right) \mathbb{1}\left(\left(S_{h}^{*}, A_{h}^{*}\right) \in \mathcal{G}_{h}^{\text {C }}\right)+\mathbb{1}\left(\left(S_{h}^{*}, A_{h}^{*}\right) \in \mathcal{G}_{h}\right) \\
& =\mathbb{E} \sum_{h=1}^{H} \bar{R}_{h}\left(u,\left(S_{h}^{*}, A_{h}^{*}\right)\right)+\mathbb{1}\left(\left(S_{h}^{*}, A_{h}^{*}\right) \in \mathcal{G}_{h}\right)=\mathbb{E} \sum_{h=1}^{H} \bar{R}_{h}\left(u,\left(S_{h}^{*}, A_{h}^{*}\right)\right)+\sum_{h=1}^{H} P_{h}^{\Pi^{*}}\left(\mathcal{G}_{h}\right) \\
& \leq \mathbb{E} \sum_{h=1}^{H} \bar{R}_{h}\left(u,\left(S_{h}^{*}, A_{h}^{*}\right)+\frac{5 \epsilon}{8}\right. \\
& \leq \mathbb{E} \sum_{h=1}^{H} \bar{R}_{h}\left(u,\left(\bar{S}_{h}, \bar{A}_{h}\right)+\frac{5 \epsilon}{8}\right. \\
& \leq \mathbb{E} \sum_{h=1}^{H} R_{h}\left(u,\left(\bar{S}_{h}, \bar{A}_{h}\right)+\frac{5 \epsilon}{8}\right. \tag{13}
\end{align*}
$$

In the first step we have used the fact that the rewards are uniformly bounded in $[0,1]$. In the second step, we have used the definition of $\bar{R}_{h}(u,(s, a)):=R_{h}(u,(s, a)) \mathbb{1}\left((s, a) \in \mathcal{G}_{h}^{\mathrm{C}}\right)$. In the third step, we have used the guarantee in (12). In the fourth step, we have used the fact that $\bar{P} i$ maximizes the reward $\bar{R}_{h}$. In the fifth step, we have used the fact that $R_{h}(u,(s, a)) \geq \bar{R}_{h}(u,(s, a))$ uniformly. From the discussion above, this concludes the proof of the theorem.

## C. Analysis - Linear MDPs

Lemma C.1. Suppose Assumption 5.5 holds. Let $\kappa>1$ and $T \geq C \frac{d \kappa^{2}}{(\gamma-\epsilon)^{2}} \log \frac{d \kappa}{\gamma-\epsilon}$. With probability at-least $1-$ $H \exp (-c(\gamma-\epsilon) T)$, the output of Algorithm 2 returns $\phi_{t h}$ such that $\sum_{t=1}^{T} \phi_{t h} \phi_{t h}^{\top} \succeq \kappa^{2} I$ for every $h \in[H]$

Proof of Theorem 5.8 By Theorem 1 in Wagenmaker et al., 2022), we take $K^{\text {rf }}(\epsilon, \delta / 4)=\frac{C d H^{5}\left(d+\log \left(\frac{1}{\delta}\right)\right)}{\epsilon^{2}}+$ $\frac{C d^{9 / 2} H^{6} \log ^{4}\left(\frac{1}{\delta}\right)}{\epsilon}$. Phase 1 succeeds with probability $1-\frac{\delta}{4}$.
Note that this is the quantity $T_{\mathrm{rf}}$ in the statement of the theorem. We now condition on the success of Phase 1 . The number of trajectories queried by Algorithm 2 which is given by $H T=T_{\text {pol }}$. By Lemma C.1, we conclude that for the given values of $T$ and $\kappa$, this algorithm successfully outputs $\phi_{h t}$ such that $G_{\phi, h} \succeq \kappa^{2} I$ for every $h \in[H]$, with probability at-least $1-\frac{\delta}{4}$.

Now, condition on the success of Algorithm 2 By theorem 6.2, we conclude that with these conclude that with proabability at-least $1-\frac{\delta}{4}$, with the values of the given parameters, for every $h \in[H]$, the procedure in Step 2 of Phase 2 outputs a policy $\hat{\Pi}^{f, h}$ such that whenever $S_{1: H}, A_{1: H} \sim \mathcal{M}\left(\hat{\Pi}^{f, h}\right)$, the conditions in 4] is satisfied for the random vector $\psi\left(S_{h}, A_{h}\right)$ with $\zeta$ replaced by $\zeta / 2$. We then use the active learning based matrix completion procedure given in Section 7 , where the vectors $\psi_{j k}$ are sample using the policy $\hat{\Pi}^{f, h}$ on the given user. By theorem 7.1, we conclude that conditioned on the success of all the steps above, with probability $1-\frac{\delta}{4}$, we can exactly estimate each of the matrices $\Theta_{h}^{*}$ for $h \in[H]$ with $T_{\text {mat-comp }}$ number of samples.
Upon the success of Phases 1,2,3 (which occurs with probability at-least $1-\delta$ by union bound), we conclude that Phase 4 gives the $\epsilon$ optimal policy for each user $u \in[N]$ because of the guarantees of reward free RL.

## D. Deferred Proofs

## D.1. Proof of Lemma B. 1

Proof. We suppose that the reward free RL in Phase 1 succeeds and returns the $\frac{\epsilon}{8}$ optimal policy for every choice of rewards bounded in $[0,1]$. The algorithm terminates whenever the active sets are such that

$$
\begin{equation*}
\hat{V}(\mathcal{J}(; \mathcal{G})) \leq \frac{\epsilon}{2} \tag{14}
\end{equation*}
$$

Note that by the definition of $\mathcal{J}(; \mathcal{G})$, the maximum value for the MDP with reward $\mathcal{J}(; \mathcal{G})$ is $\sup _{\Pi} \sum_{h=1}^{H} P^{\Pi}\left(\mathcal{G}_{h}\right)$. Since $\hat{V}$ is the output of the reward free RL algorithm, we conclude that we have:

$$
\begin{equation*}
\left|\hat{V}(\mathcal{J}(; \mathcal{G}))-\sup _{\Pi} \sum_{h=1}^{H} P_{h}^{\Pi}\left(\mathcal{G}_{h}\right)\right| \leq \frac{\epsilon}{8} \tag{15}
\end{equation*}
$$

We conclude via (14) and (15) that (10) holds, which establishes the second part of the theorem. We now consider the termination time.
Suppose $\mathcal{G}^{(t)}$ is the sequence of active sets before termination at step $t$ (i.e, it satisfies $\left.\hat{V}\left(\mathcal{J}\left(; \mathcal{G}^{(t)}\right)\right)>\frac{\epsilon}{2}\right)$. Recall $\hat{\Pi}$, the output of the reward free RL algorithm. It follows from the guarantees for reward free RL that:

$$
\left|\sum_{h=1}^{H} P_{h}^{\hat{\Pi}^{\mathcal{G}}}\left(\mathcal{G}_{h}^{(t)}\right)-\sup _{\Pi} \sum_{h=1}^{H} P_{h}^{\Pi}\left(\mathcal{G}_{h}^{(t)}\right)\right| \leq \frac{\epsilon}{8}
$$

Combining this with Equation (15) and the fact that $\hat{V}\left(\mathcal{J}\left(; \mathcal{G}^{(t)}\right)\right)>\frac{\epsilon}{2}$, we conclude:

$$
\begin{equation*}
\sum_{h=1}^{H} P_{h}^{\hat{\Pi}^{\mathcal{G}}}\left(\mathcal{G}_{h}^{(t)}\right) \geq \frac{\epsilon}{4} \tag{16}
\end{equation*}
$$

We consider the potential function with $\varphi(0)=0$ and $\varphi(t)=\sum_{h \in[H]} \sum_{(s, a) \in \mathcal{S} \times \mathcal{A}} T_{h,(s, a)}^{(t)}$, where $T_{h,(s, a)}^{(t)}$ is the counter $T_{h,(s, a)}$ inside Algorithm 1 at the beginning of the step $t$.
Whenever $\mathcal{G}^{(t)}$ is such that $\hat{V}\left(\mathcal{J}\left(; \mathcal{G}^{(t)}\right)\right)>\frac{\epsilon}{2}$, we define $N_{t}:=\varphi(t+1)-\varphi(t)$ (i.e, before termination). Just for the sake of theoretical arguments, we define the fictious random variables $N_{t}=\operatorname{Ber}\left(\frac{\epsilon}{8}\right)$ i.i.d after termination. Let $\mathcal{F}_{t}=\sigma\left(\mathcal{G}^{(s)}, S_{1: H}^{(s)}, A_{1: H}^{(s)}, R_{1: H}^{(s)}, U^{(s)}: s \leq t\right)$
Claim D.1. The following relations hold:

1. $\mathbb{E}\left[N_{t} \mid \mathcal{F}_{t}\right] \geq \frac{\epsilon}{8}$
2. $\mathbb{E}\left[N_{t}^{2} \mid \mathcal{F}_{t}\right] \leq \frac{\mathbb{E}\left[N_{t} \mid \mathcal{F}_{t}\right] H}{4}$
3. $\left|N_{t}\right| \leq H$ almost surely.

Proof. The inequalities are clear when $\mathcal{G}^{(t)}$ is such that $\hat{V}\left(\mathcal{J}\left(; \mathcal{G}^{(t)}\right)\right) \leq \frac{\epsilon}{2}$. Now consider the case $\hat{V}\left(\mathcal{J}\left(; \mathcal{G}^{(t)}\right)\right)>\frac{\epsilon}{2}$. By definition, conditioned on this event, we have almost surely:

$$
N_{t}=\sum_{h=1}^{H} \mathbb{1}\left(\left(S_{h}^{(t)}, A_{h}^{(t)}\right) \in \mathcal{G}_{h}^{(t)}\right) \cdot \mathbb{1}\left(\hat{R}_{h}^{(t)}\left(U_{t},\left(S_{h}^{(t)}, A_{h}^{(t)}\right)\right)=*\right)
$$

That is, we increment the $T_{h,(s, a)}$ only when we encounter an element of the active set such that the entry for this user has not been observed before. Observe that for any arbitrary $(s, a) \in \mathcal{S} \times \mathcal{A}$

$$
\begin{equation*}
\mathbb{P}\left(\hat{R}_{h}^{(t)}\left(U_{t},(s, a)\right)=* \mid \mathcal{F}_{t},\left(S_{h}^{(t)}, A_{h}^{(t)}\right)=(s, a)\right)=\frac{\left|\left\{u: \hat{R}_{h}^{(t)}(s, a)=*\right\}\right|}{N} \tag{17}
\end{equation*}
$$

This is true since the law of $S_{h}^{(t)}, A_{h}^{(t)}$ is independent of $U_{t}$ (since all users share the same MDP), when conditioned on $\mathcal{F}_{t}$. Now, the algorithm only fills the column corresponding to $(s, a)$ until the number of entries is smaller than $N p \leq \frac{N}{2}$. We conclude that:

$$
\left|\left\{u: \hat{R}^{(t)}(h,(s, a))=*\right\}\right| \geq N-N p \geq \frac{N}{2}
$$

This allows us to conclude $\mathbb{P}\left(\hat{R}_{h}^{(t)}\left(U_{t},(s, a)\right)=* \mid \mathcal{F}_{t},\left(S_{h}^{(t)}, A_{h}^{(t)}\right)=(s, a)\right) \geq \frac{1}{2}$ and hence:

$$
\begin{align*}
\mathbb{E} N_{t} & =\sum_{h=1}^{H} \mathbb{E} \mathbb{1}\left(\left(S_{h}^{(t)}, A_{h}^{(t)}\right) \in \mathcal{G}_{h}^{(t)}\right) \cdot \mathbb{1}\left(R_{h}^{(t)}\left(U_{t},\left(S_{h}^{(t)}, A_{h}^{(t)}\right)\right)=*\right) \\
& \geq \frac{1}{2} \sum_{h=1}^{H} \mathbb{E} \mathbb{1}\left(\left(S_{h}^{(t)}, A_{h}^{(t)}\right) \in \mathcal{G}_{h}^{(t)}\right) \\
& =\frac{1}{2} \sum_{h=1}^{H} P_{h}^{\hat{\Pi}^{G}}\left(\mathcal{G}_{h}^{(t)}\right) \geq \frac{\epsilon}{8} \tag{18}
\end{align*}
$$

In the last step we have used 16 . The bound $\left|N_{t}\right| \leq H$ almost surely follows from definition. Now note that $\mathbb{E}\left[N_{t}^{2} \mid \mathcal{F}_{t}\right] \leq$ $H \mathbb{E}\left[N_{t} \mid \mathcal{F}_{t}\right]$.

Claim D.2. For any $\tau \in \mathbb{N}$ and some $c_{0}>0$ small enough, we have:

$$
\mathbb{P}\left(\sum_{t=0}^{\tau-1} N_{t}<\frac{\epsilon \tau}{16}\right) \leq \exp \left(-c_{0} \frac{\epsilon \tau}{H}\right)
$$

Proof. For $\frac{3}{4 H}>\lambda>0$, consider $M_{t}=-\frac{\lambda^{2} \mathbb{E}\left[N_{t}^{2} \mid \mathcal{F}_{t}\right]}{1-\frac{\lambda H}{3}}+\lambda\left(\mathbb{E}\left[N_{t} \mid \mathcal{F}_{t}\right]-N_{t}\right)$. Now consider:

$$
\begin{align*}
\mathbb{E} \exp \left(\sum_{t=0}^{\tau-1} M_{t}\right) & =\mathbb{E}\left[\mathbb{E}\left[\exp \left(M_{\tau-1}\right) \mid \mathcal{F}_{\tau-1}\right] \exp \left(\sum_{t=0}^{\tau-1} M_{t}\right)\right] \\
& =\mathbb{E}\left[\exp \left(\lambda \mathbb{E}\left[N_{\tau-1} \mid \mathcal{F}_{\tau-1}\right]-\lambda N_{t}\right) \mid \mathcal{F}_{\tau-1}\right] \exp \left(\sum_{t=0}^{\tau-2} M_{t}\right) \exp \left(-\frac{\lambda^{2} \mathbb{E}\left[N_{\tau-1}^{2} \mid \mathcal{F}_{\tau-1}\right]}{1-\frac{\lambda H}{3}}\right) \\
& \leq \mathbb{E} \exp \left(-\frac{\lambda^{2} \mathbb{E}\left[N_{\tau-1}^{2} \mid \mathcal{F}_{\tau-1}\right]}{1-\frac{\lambda H}{3}}\right) \exp \left(\sum_{t=0}^{\tau-2} M_{t}\right) \exp \left(-\frac{\lambda^{2} \mathbb{E}\left[N_{\tau-1}^{2} \mid \mathcal{F}_{\tau-1}\right]}{1-\frac{\lambda H}{3}}\right) \\
& =\mathbb{E} \exp \left(\sum_{t=0}^{\tau-2} M_{t}\right) \tag{19}
\end{align*}
$$

In the first step we have used the fact that $\sum_{t=0}^{\tau-2} M_{t}$ is $\mathcal{F}_{\tau-1}$ measurable and the towering property of conditional expectation. In the third step, we have used the exponential moment bound given in Exercise 2.8.5 in (Vershynin, 2018), as applied to $N_{\tau}-\mathbb{E}\left[N_{\tau-1} \mid \mathcal{F}_{\tau-1}\right]$ along with the fact that $N_{t} \in[0, H]$ almost surely. From $\sqrt{19}$, we conclude that $\mathbb{E} \exp \left(\sum_{t=0}^{\tau} M_{t}\right) \leq 1$ and thus applying the Chernoff bound, we conclude that for any $\beta>0$

$$
\mathbb{P}\left(\sum_{t=0}^{\tau-1}-\frac{\lambda \mathbb{E}\left[N_{t}^{2} \mid \mathcal{F}_{t}\right]}{1-\frac{\lambda H}{3}}+\left(\mathbb{E}\left[N_{t} \mid \mathcal{F}_{t}\right]-N_{t}\right)>\beta\right) \leq \exp (-\lambda \beta)
$$

Now, using item 2 from Claim D.1 we conclude that

$$
\mathbb{P}\left(\sum_{t=0}^{\tau-1} N_{t}<-\beta+\frac{3-4 \lambda H}{3-\lambda H} \sum_{t=0}^{\tau-1} \mathbb{E}\left[N_{t} \mid \mathcal{F}_{t}\right]\right) \leq \exp (-\lambda \beta)
$$

Now, using item 1 from Claim D.1, we note that $\sum_{t=0}^{\tau-1} \mathbb{E}\left[N_{t} \mid \mathcal{F}_{t}\right] \geq \frac{\epsilon \tau}{8}$ almost surely. Setting $\lambda=\frac{1}{4 H}$ and $\beta=c_{0} \epsilon \tau$ for some small enough constant $\epsilon$, we conclude:

$$
\mathbb{P}\left(\sum_{t=0}^{\tau-1} N_{t}<\frac{\epsilon \tau}{16}\right) \leq \exp \left(-c_{0} \frac{\epsilon \tau}{H}\right)
$$

Let $\tau^{\text {term }}$ denote the termination time for the algorithm. This is true since $\varphi(t)$ is increasing in $t, \varphi(t) \leq N p H|\mathcal{S} \| \mathcal{A}|$, and strict inequality holds when $t<\tau^{\text {term }}$. For every $\tau<\tau^{\text {term }}$ we have $\varphi(\tau)=\sum_{t=0}^{\tau-1} N_{\tau}<N p H|\mathcal{S}||\mathcal{A}|$. Therefore, we have the following relationship between the events:

$$
\left\{\tau^{\mathrm{term}}>\tau\right\} \subseteq\left\{\sum_{t=1}^{\tau} N_{\tau}<N p|\mathcal{S} \| \mathcal{A}| H\right\}
$$

Setting $\tau=\frac{16 N p|\mathcal{S}||\mathcal{A}| H}{\epsilon}$, we have:

$$
\mathbb{P}\left(\tau^{\text {term }}>\tau\right) \leq \mathbb{P}\left(\sum_{t=1}^{\tau} N_{\tau}<N p|\mathcal{S} \| \mathcal{A}| H\right) \leq \exp (-c N p|\mathcal{S} \| \mathcal{A}|)
$$

## D.2. Proof of Lemma C. 1

Proof. By Remark 4.3 in (Wagenmaker et al. 2022), we show that non-linear rewards can be handled by the reward free RL algorithm in Phase 1 as long all the reward are uniformly bounded in $[0,1]$. Let $B_{t h}$ be the matrix $I+A_{\phi}$ in Algorithm 2 at step $t$ for horizon $h$. Similarly, let the corresponding projection $Q$ be $Q_{t h}$. Recall that $Q_{t h}$ is the projection onto an eigenspace of $B_{t h}$. Now, suppose $S_{1: H}, A_{1: H} \sim \mathcal{M}_{U_{t}}\left(\hat{\Pi}^{Q_{t, h}}\right)$ as in the algorithm. Let $\phi_{t h}:=\phi\left(S_{h}, A_{h}\right)$. Now, if $Q_{t h} \neq 0$, then:

$$
\begin{align*}
\phi_{t h}^{\top} B_{t h}^{-1} \phi_{t h} & \geq \phi_{t h}^{\top} Q_{t h} B_{t h}^{-1} Q_{t h} \phi_{t h} \\
& \geq \phi_{t h}^{\top} Q_{t h} \frac{I}{1+\kappa^{2}} Q_{t h} \phi_{t h} \\
& =\frac{\left\|Q_{t h} \phi_{t h}\right\|^{2}}{1+\kappa^{2}} \tag{20}
\end{align*}
$$

In the first step, we have used the fact that $Q_{t h}$ is the projector to the eigenspace of $B_{t h}^{-1}$. In the second step, we have used the fact that over the eigenspace corresponding to $Q_{t h}$, the eigenvalues of $B_{t h}^{-1}$ are at-least $\frac{1}{1+\kappa^{2}}$. We now invoke Assumption 5.5 in order to show that, along with the guarantees of reward free RL in phase 1 , we conclude that:

$$
\begin{equation*}
\mathbb{E}\left[\left\|Q_{t h} \phi_{t h}\right\|^{2} \mid Q_{t h} \neq 0, B_{t h}\right] \geq \gamma-\epsilon \tag{21}
\end{equation*}
$$

Now, note by the fact that $Q_{t h}$ is a projector and that $\left\|\phi_{t h}\right\| \leq 1$, we have:

$$
\begin{equation*}
\mathbb{E}\left[\left\|Q_{t h} \phi_{t h}\right\|^{4} \mid Q_{t h} \neq 0, B_{t h}\right] \leq \mathbb{E}\left[\left\|Q_{t h} \phi_{t h}\right\|^{2} \mid Q_{t h} \neq 0, B_{t h}\right] \tag{22}
\end{equation*}
$$

Recall the Paley-Zygmund inequality which states that for any positive random variable $Z$, we must have: $\mathbb{P}\left(Z>\frac{\mathbb{E} Z}{2}\right) \geq$ $\frac{1}{4} \frac{(\mathbb{E} Z)^{2}}{\mathbb{E} Z^{2}}$. Therefore,

$$
\begin{align*}
& \mathbb{P}\left[\left.\phi_{t h}^{\top} B_{t h}^{-1} \phi_{t h}>\frac{\gamma-\epsilon}{2\left(1+\kappa^{2}\right)} \right\rvert\, Q_{t h} \neq 0, B_{t h}\right] \geq \mathbb{P}\left[\left.\left\|Q_{t h} \phi_{t h}\right\|^{2}>\frac{\gamma-\epsilon}{2} \right\rvert\, Q_{t h} \neq 0, B_{t h}\right] \\
& \geq \mathbb{P}\left[\left.\left\|Q_{t h} \phi_{t h}\right\|^{2}>\frac{1}{2} \mathbb{E}\left[\left\|Q_{t h} \phi_{t h}\right\|^{2} \mid Q_{t h} \neq 0, B_{t h}\right] \right\rvert\, Q_{t h} \neq 0, B_{t h}\right] \\
& \geq \frac{1}{4} \frac{\mathbb{E}\left[\left\|Q_{t h} \phi_{t h}\right\|^{2} \mid Q_{t h} \neq 0, B_{t h}\right]^{2}}{\mathbb{E}\left[\left\|Q_{t h} \phi_{t h}\right\|^{4} \mid Q_{t h} \neq 0, B_{t h}\right]} \geq \frac{1}{4} \mathbb{E}\left[\left\|Q_{t h} \phi_{t h}\right\|^{2} \mid Q_{t h} \neq 0, B_{t h}\right] \\
& \geq \frac{\gamma-\epsilon}{4} \tag{23}
\end{align*}
$$

In the first step, we have used 20. In the second step, we have used 21). In the third step, we have used the Paley-Zygmund inequality and the moment bound in (22).
Define the stopping time $\tau=\inf \left\{t \leq T: Q_{t h}=0\right\}$ and $\tau=\infty$ if the set in the RHS is empty. Let $\Xi_{t}^{0}$ for $t \in\{0\} \cup \mathbb{N}$ be a sequence of i.i.d random variables with the law $\frac{\gamma-\epsilon}{2\left(1+\kappa^{2}\right)} \operatorname{Ber}\left(\frac{\gamma-\epsilon}{4}\right)$. We consider the sequence of random variables $\Xi_{t}=\phi_{t h}^{\top} B_{t h}^{-1} \phi_{t h}$ for $t<\tau$ and $\Xi_{t}=\Xi_{t}^{0}$ for $t \geq \tau$
Now, we apply the matrix determinant lemma which states that $\operatorname{det}\left(B+u u^{\boldsymbol{\top}}\right)=\operatorname{det}(B)\left(1+u^{\boldsymbol{\top}} B^{-1} u\right)$. We note that $B_{(t+1) h}=B_{t h}+\phi_{t h} \phi_{t h}^{\top}$. Therefore, whenever $t<\tau$, we must have:

$$
\begin{equation*}
\operatorname{det}\left(B_{(t+1) h}\right)=\operatorname{det}\left(B_{t h}\right)\left(1+\Xi_{t}\right) \tag{24}
\end{equation*}
$$

Since $\left\|\phi_{t h}\right\| \leq 1$ almost surely, we must have

$$
\operatorname{Tr}\left(B_{t h}\right)=\sum_{i=1}^{d}\left\langle e_{i}, B_{t h} e_{i}\right\rangle \leq d+t
$$

It is easy to show that for any PSD matrix, $A$, if $\operatorname{Tr}(A) \leq \alpha$, then $\operatorname{det}(A) \leq\left(\frac{\alpha}{d}\right)^{d}$ (since trace is the sum of the eigenvalues and the determinant is the product). Combining the equations above, we conclude that whenever $t<\tau$, we must have:

$$
\left(\frac{t+1+d}{d}\right)^{d} \geq \prod_{s=0}^{t}\left(1+\Xi_{t}\right)
$$

Therefore, the event

$$
\begin{equation*}
\{\tau>T\} \subseteq\left\{\left(\frac{T+1+d}{d}\right)^{d} \geq \prod_{s=0}^{T}\left(1+\Xi_{t}\right)\right\} \tag{25}
\end{equation*}
$$

Claim D.3.

$$
\mathbb{P}\left[\prod_{s=0}^{T}\left(1+\Xi_{t}\right) \geq\left(1+\frac{\gamma-\epsilon}{2\left(1+\kappa^{2}\right)}\right)^{\frac{(\gamma-\epsilon) T}{8}}\right] \geq 1-\exp \left(-c_{0} T(\gamma-\epsilon)\right)
$$

Let $\kappa>1$ and $T \geq C \frac{d \kappa^{2}}{(\gamma-\epsilon)^{2}} \log \frac{d \kappa}{\gamma-\epsilon}$, we have:

$$
\mathbb{P}\left[\prod_{s=0}^{T}\left(1+\Xi_{t}\right) \geq\left(\frac{T+1+d}{d}\right)^{d}\right] \geq 1-\exp \left(-c_{0} T(\gamma-\epsilon)\right)
$$

Proof. Let $N_{T}$ be the number of variables $\left(\Xi_{t}\right)_{t=0}^{T}$ such that $\Xi_{t} \geq \frac{\gamma-\epsilon}{2 \kappa^{2}}$. Then, it is clear that $\prod_{s=0}^{T}\left(1+\Xi_{t}\right) \geq\left(1+\frac{\gamma-\epsilon}{2 \kappa^{2}}\right)^{N_{T}}$. Therefore,

$$
\begin{align*}
\mathbb{P}\left[\prod_{s=0}^{T}\left(1+\Xi_{t}\right) \geq\left(1+\frac{\gamma-\epsilon}{2\left(1+\kappa^{2}\right)}\right)^{\frac{(\gamma-\epsilon) T}{8}}\right] & \geq \mathbb{P}\left(N_{T} \geq \frac{(\gamma-\epsilon) T}{8}\right) \\
& \geq \mathbb{P}\left(\operatorname{Bin}\left(T, \frac{\gamma-\epsilon}{4}\right) \geq \frac{(\gamma-\epsilon) T}{8}\right) \\
& \geq 1-\exp \left(-c_{0} T(\gamma-\epsilon)\right) \tag{26}
\end{align*}
$$

Here Bin refers to the law of a binomial random variable. The first step follows from the fact that $\prod_{s=0}^{T}\left(1+\Xi_{t}\right) \geq$ $\left(1+\frac{\gamma-\epsilon}{2\left(1+\kappa^{2}\right)}\right)^{N_{T}}$ almost surely. The second step follows from (23), which shows that conditioned on $Q_{t h}, B_{t h}$, the random variable $\mathbb{1}\left(\Xi_{t} \geq \frac{\gamma-\epsilon}{2\left(1+\kappa^{2}\right)}\right)$ stochastically dominates $\operatorname{Ber}\left(\frac{\gamma-\epsilon}{4}\right)$. The last step follows from an application of Bernstein's inequality for binomial random variables.

Now, using (25) along with Claim D.3, we conclude:

$$
\mathbb{P}(\tau>T) \leq \mathbb{P}\left(\left(\frac{T+1+d}{d}\right)^{d} \geq \prod_{s=0}^{T}\left(1+\Xi_{t}\right)\right) \leq \exp \left(-c_{0} T(\gamma-\epsilon)\right)
$$

## D.3. Proof of Lemma 6.1

Proof. By the definition of Linear MDP, we must have $S_{h+1} \mid S_{h}, A_{h} \sim \sum_{i=1}^{d}\left\langle\phi\left(S_{h}, A_{h}\right), e_{i}\right\rangle \mu_{i h}(\cdot)$ and $A_{h+1} \mid S_{h+1} \sim$ $\pi_{h+1}\left(\cdot \mid S_{h}\right)$. Therefore, for any bounded, measurable function $g: \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$, we must have:

$$
\begin{align*}
\mathbb{E} g\left(S_{h+1}, A_{h+1}\right) & =\mathbb{E}\left[\mathbb{E}\left[g\left(S_{h+1}, A_{h+1}\right) \mid S_{h}, A_{h}\right]\right] \\
& =\mathbb{E} \sum_{i=1}^{d}\left\langle\phi\left(S_{h}, A_{h}\right), e_{i}\right\rangle \int \mu_{i(h-1)}(d s) \pi_{h+1}(d a \mid s) g(s, a) \\
& =\sum_{i=1}^{d} \nu_{i h} \int \mu_{i h}(d s) \pi_{h+1}(d a \mid s) g(s, a) \tag{27}
\end{align*}
$$

## D.4. Proof of Lemma E. 1

Proof. It is clear from the assumption that $\mathbb{E}\left[\int g\left(s_{(h+1) t}, a\right) \pi_{h+1}\left(d a \mid s_{t}\right) \mid\left(\phi_{h t}\right)_{t \leq T}\right] \quad=$ $\sum_{i=1}^{d}\left\langle\phi_{h t}, e_{i}\right\rangle \int \mu_{i h}(d s) \pi_{h+1}(d a \mid s) g(s, a)$.
Note that

$$
\begin{align*}
\sum_{t=1}^{T} \alpha_{h t, \nu} \phi_{h t} & =\sum_{t=1}^{T} \phi_{h t}^{\top} G_{\phi, h}^{-1} \nu \phi_{h t} \\
& =\sum_{t=1}^{T} \phi_{h t} \phi_{h t}^{\top} G_{\phi, h}^{-1} \nu=\left(\sum_{t=1}^{T} \phi_{h t} \phi_{h t}^{\top}\right) G_{\phi, h}^{-1} \nu \\
& =G_{\phi, h} G_{\phi, h}^{-1} \nu=\nu \tag{28}
\end{align*}
$$

Therefore,

$$
\begin{align*}
\mathbb{E}\left[\hat{\mathcal{T}}\left(g ; \nu, \pi_{h}\right) \mid\left(\phi_{l}\right)_{t \in[T]}\right] & =\sum_{i=1}^{d}\left\langle\sum_{t=1}^{T} \alpha_{t, \nu} \phi_{t}, e_{i}\right\rangle \int \mu_{i(h-1)}(d s) \pi_{h}(d a \mid s) g(s, a) \\
& =\sum_{i=1}^{d}\left\langle\nu, e_{i}\right\rangle \int \mu_{i(h-1)}(d s) \pi_{h}(d a \mid s) g(s, a)=\mathcal{T}\left(g ; \nu, \pi_{h}\right) \tag{29}
\end{align*}
$$

Note that, conditioned on $\left(\phi_{t}\right)_{t \in[T]}, \alpha_{h t, \nu} \int g\left(s_{(h+1) t}, a\right) \pi_{h+1}\left(d a \mid s_{(h+1) t}\right)$ are independent random variables bounded above by $\alpha_{h t, \nu} B$. Therefore, applying the Azuma-Hoeffding inequality, we conclude:

$$
\mathbb{P}\left(\left|\hat{\mathcal{T}}\left(g ; \nu, \pi_{h}\right)-\mathcal{T}\left(g ; \nu, \pi_{h}\right)\right|>\beta \mid\left(\phi_{t}\right)_{t \in[T]}\right) \leq 2 \exp \left(-\frac{t^{2}}{2 B^{2} \sum_{t} \alpha_{t, \nu}^{2}}\right)
$$

Now, observe that $\sum_{t} \alpha_{h t, \nu}^{2}=\sum_{t} \nu^{\top} G_{\phi, h}^{-1} \nu \leq \frac{1}{\kappa^{2}}$ whenever $G_{\phi, h} \succeq \kappa^{2} I$ This concludes the proof.

## D.5. Proof of Lemma E. 2

Proof. Notice that:

$$
\begin{align*}
& \left|E_{1}^{\nu_{1}}(\Pi)-E_{1}^{\nu_{1}^{\prime}}\left(\Pi^{\prime}\right)\right| \leq\left|E_{1}^{\nu_{1}}(\Pi)-E_{1}^{\nu_{1}}\left(\Pi^{\prime}\right)\right|+\left|E_{1}^{\nu_{1}}\left(\Pi^{\prime}\right)-E_{1}^{\nu_{1}^{\prime}}\left(\Pi^{\prime}\right)\right| \\
& \leq\left|E_{1}^{\nu_{1}}(\Pi)-E_{1}^{\nu_{1}}\left(\Pi^{\prime}\right)\right|+\left\|\nu_{1}-\nu_{1}^{\prime}\right\|_{1} \\
& \leq\left\|\mathbb{E} \int \phi\left(S_{1}, a\right) \pi_{1}\left(d a \mid S_{1}\right)-\mathbb{E} \int \phi\left(S_{1}, a\right) \pi_{1}^{\prime}\left(d a \mid S_{1}\right)\right\|_{1}+\left\|\nu_{1}-\nu_{1}^{\prime}\right\|_{1} \\
& \leq \sup _{(s, a)}\|\phi(s, a)\|_{1} \operatorname{TV}\left(\pi_{1}, \pi_{1}^{\prime}\right)+\left\|\nu_{1}-\nu_{1}^{\prime}\right\|_{1} \leq \operatorname{TV}\left(\pi_{1}, \pi_{1}^{\prime}\right)+\left\|\nu_{1}-\nu_{1}^{\prime}\right\|_{1} \tag{30}
\end{align*}
$$

In the first, second and third steps we have used the triangle inequality. In the last step, we have used the fact that for any bounded function, and any probability measures $\mu, \nu$, we have $\left|\int f(x) \mu(d x)-\int f(x) \nu(d x)\right| \leq \sup _{x}|f(x)| \mathrm{TV}(\nu, \mu)$.
Now consider:

$$
\begin{align*}
& \left|\left\|\mathcal{T}_{j}\left(\phi, \nu_{j-1}, \pi_{j}\right)-\nu_{j}\right\|_{1}-\left\|\mathcal{T}_{j}\left(\phi, \nu_{j-1}^{\prime}, \pi_{j}^{\prime}\right)-\nu_{j}^{\prime}\right\|_{1}\right|^{\prime} \\
& \leq\left\|\nu_{j}-\nu_{j}^{\prime}\right\|_{1}+\left\|\mathcal{T}_{j}\left(\phi, \nu_{j-1}, \pi_{j}\right)-\mathcal{T}_{j}\left(\phi, \nu_{j-1}^{\prime}, \pi_{j}^{\prime}\right)\right\|_{1} \\
& \leq\left\|\nu_{j}-\nu_{j}^{\prime}\right\|_{1}+\left\|\mathcal{T}_{j}\left(\phi, \nu_{j-1}, \pi_{j}\right)-\mathcal{T}_{j}\left(\phi, \nu_{j-1}^{\prime}, \pi_{j}\right)\right\|_{1}+\left\|\mathcal{T}_{j}\left(\phi, \nu_{j-1}^{\prime}, \pi_{j}\right)-\mathcal{T}_{j}\left(\phi, \nu_{j-1}^{\prime}, \pi_{j}^{\prime}\right)\right\|_{1} \tag{31}
\end{align*}
$$

Now, observe that:

$$
\begin{align*}
& \left\|\mathcal{T}_{j}\left(\phi, \nu_{j-1}, \pi_{j}\right)-\mathcal{T}_{j}\left(\phi, \nu_{j-1}^{\prime}, \pi_{j}\right)\right\|_{1} \leq \sum_{i=1}^{d}\left|\left\langle\nu_{j-1}-\nu_{j-1}^{\prime}, e_{i}\right\rangle\right|\left\|\int \phi(s, a) \mu_{i(j-1)}(d s) \pi_{j}(d a \mid s)\right\|_{1} \\
& \leq \sum_{i=1}^{d}\left|\left\langle\nu_{j-1}-\nu_{j-1}^{\prime}, e_{i}\right\rangle\right|=\left\|\nu_{j-1}-\nu_{j-1}^{\prime}\right\|_{1} \tag{32}
\end{align*}
$$

Where we recall $\sup _{i, h, \pi}\left\|\int \phi(s, a) \mu_{i h}(d s) \pi(d a \mid s)\right\|_{1} \leq 1$ as given in the definition of Linear MDP. Using the Hahn-Jordan decomposition of a signed measure, we conclude:

$$
\begin{align*}
& \left\|\mathcal{T}_{j}\left(\phi, \nu_{j-1}^{\prime}, \pi_{j}\right)-\mathcal{T}_{j}\left(\phi, \nu_{j-1}^{\prime}, \pi_{j}\right)\right\|_{1} \\
& \leq \sum_{i=1}^{d}\left|\left\langle\nu_{j-1}^{\prime}, e_{i}\right\rangle\right|\left\|\int \phi(s, a) \mu_{i(j-1)}(d s)\left(\pi_{j}(d a \mid s)-\pi_{j}^{\prime}(d a \mid s)\right)\right\|_{1} \\
& \leq \sum_{i=1}^{d} C_{\mu}\left|\left\langle\nu_{j-1}^{\prime}, e_{i}\right\rangle\right| \operatorname{TV}\left(\pi_{j}, \pi_{j}^{\prime}\right) \leq C_{\mu}\left\|\nu_{j-1}^{\prime}\right\|_{1} \operatorname{TV}\left(\pi_{j}, \pi_{j}^{\prime}\right) \tag{33}
\end{align*}
$$

Combining (31), (32) and (33) we conclude:

$$
\begin{align*}
& \left|\left\|\mathcal{T}_{j}\left(\phi, \nu_{j-1}, \pi_{j}\right)-\nu_{j}\right\|_{1}-\left\|\mathcal{T}_{j}\left(\phi, \nu_{j-1}^{\prime}, \pi_{j}^{\prime}\right)-\nu_{j}^{\prime}\right\|_{1}\right| \\
& \quad \leq\left\|\nu_{j}-\nu_{j}^{\prime}\right\|_{1}+C_{\mu}\left\|\nu_{j-1}^{\prime}\right\|_{1} \operatorname{TV}\left(\pi_{j}, \pi_{j}^{\prime}\right)+\left\|\nu_{j-1}-\nu_{j-1}^{\prime}\right\|_{1} \tag{34}
\end{align*}
$$

Combining (30) and (34), we conclude the first inequality in the statement of the lemma.
With a reasoning very similar to that in 30, we have:

$$
\begin{equation*}
\left|\hat{E}_{1}^{\nu_{1}}(\Pi)-\hat{E}_{1}^{\nu_{1}^{\prime}}\left(\Pi^{\prime}\right)\right| \leq \operatorname{TV}\left(\pi_{1}, \pi_{1}^{\prime}\right)+\left\|\nu_{1}-\nu_{1}^{\prime}\right\|_{1} \tag{35}
\end{equation*}
$$

Using similar reasoning as in (34):

$$
\begin{align*}
& \left|\left\|\mathcal{T}_{j}\left(\phi, \nu_{j-1}, \pi_{j}\right)-\nu_{j}\right\|_{1}-\left\|\mathcal{T}_{j}\left(\phi, \nu_{j-1}^{\prime}, \pi_{j}^{\prime}\right)-\nu_{j}^{\prime}\right\|_{1}\right| \\
& \leq\left\|\nu_{j}-\nu_{j}^{\prime}\right\|_{1}+\left(\sum_{t=1}^{T}\left|\left(\nu_{j-1}-\nu_{j-1}^{\prime}\right)^{\top} G_{\phi, j-1}^{-1} \phi_{(j-1) t}\right|\right)+\left(\sum_{t=1}^{T}\left|\left(\nu_{j-1}^{\prime}\right)^{\top} G_{\phi, j-1}^{-1} \phi_{(j-1) t}\right|\right) \operatorname{TV}\left(\pi_{j}, \pi_{j}^{\prime}\right) \tag{36}
\end{align*}
$$

Now note that for any $\nu \in \mathbb{R}^{d}$, we have:

$$
\begin{align*}
\sum_{t=1}^{T}\left|\nu^{\top} G_{\phi, j-1}^{-1} \phi_{(j-1) t}\right| & \leq \sqrt{T \sum_{t}\left|\nu^{\top} G_{\phi, j-1}^{-1} \phi_{(j-1) t}\right|^{2}} \\
& =\sqrt{T \sum_{t=1}^{T} \nu^{\top} G_{\phi, j-1}^{-1} \phi_{(j-1) t} \phi_{(j-1) t}^{\top} G_{\phi, j-1}^{-1} \nu} \\
& =\sqrt{T \nu^{\top} G_{\phi, j-1}^{-1} \nu} \\
& \leq \sqrt{\frac{T}{\kappa^{2}}}\|\nu\|_{2} \tag{37}
\end{align*}
$$

Here, in the first step we have used the fact that whenever $x \in \mathbb{R}^{K}$, we must have $\|x\|_{1} \leq \sqrt{K}\|x\|_{2}$. In the third step, we have used the fact that $\sum_{t=1}^{T} \phi_{(j-1) t} \phi_{(j-1) t}^{\top}=G_{\phi, j-1}$ by definition. In the last step, we have used the fact that $\hat{G}_{\phi, j-1} \succeq \kappa^{2} I$. Plugging this into (36), we conclude:

$$
\begin{align*}
& \left|\left\|\mathcal{T}_{j}\left(\phi, \nu_{j-1}, \pi_{j}\right)-\nu_{j}\right\|_{1}-\left\|\mathcal{T}_{j}\left(\phi, \nu_{j-1}^{\prime}, \pi_{j}^{\prime}\right)-\nu_{j}^{\prime}\right\|_{1}\right| \\
& \quad \leq\left\|\nu_{j}-\nu_{j}^{\prime}\right\|_{1}+\sqrt{\frac{T}{\kappa^{2}}}\left\|\nu_{j-1}-\nu_{j-1}^{\prime}\right\|_{2}+\sqrt{\frac{T}{\kappa^{2}}}\left\|\nu_{j-1}^{\prime}\right\|_{2} \operatorname{TV}\left(\pi_{j}, \pi_{j}^{\prime}\right) \tag{38}
\end{align*}
$$

Using this and the definition of $\hat{F}$ we conclude the second inequality in the statement of the lemma. (45) and (46) follow from a similar reasoning.

## D.6. Proof of Lemma E. 3

Proof. First consider the case $h=1$. Let $g(s, a):=\phi(s, a)$. In this case, $\sup _{\nu \in \mathcal{B}_{d}(1)}\left|\hat{E}^{\nu, 1}(\Pi)-E^{\nu, 1}(\Pi)\right| \leq \| \mathcal{T}_{0}\left(g ; \pi_{1}\right)-$ $\hat{\mathcal{T}}_{0}\left(g ; \pi_{1}\right) \|_{1}$. By (30) and (35), we conclude that $\pi_{1} \rightarrow \mathcal{T}_{0}\left(\phi ; \pi_{1}\right)$ and $\pi_{1} \rightarrow \hat{\mathcal{T}}_{0}\left(\phi ; \pi_{1}\right)$ are 1-Lipschitz with respect to $\operatorname{TV}()$ and $\|\cdot\|_{1}$.

$$
\sup _{\Pi=\pi_{1}, \ldots, \pi_{H} \in \mathbb{Q}}\left\|\mathcal{T}_{0}\left(g ; \pi_{1}\right)-\hat{\mathcal{T}}_{0}\left(g ; \pi_{1}\right)\right\|_{1} \leq \sup _{\Pi=\pi_{1}, \ldots, \pi_{H} \in \hat{\mathbb{Q}}_{\eta}}\left\|\mathcal{T}_{0}\left(g ; \pi_{1}\right)-\hat{\mathcal{T}}_{0}\left(g ; \pi_{1}\right)\right\|_{1}+2 \eta
$$

We apply LemmaE.1 co-ordinate wise to the co-ordinates of $\phi$ and union bound over $\hat{\mathbb{Q}}_{\eta}$. We have:

$$
\begin{equation*}
\mathbb{P}\left(\sup _{\Pi=\pi_{1}, \ldots, \pi_{H} \in \mathbb{Q}}\left\|\mathcal{T}_{0}\left(g ; \pi_{1}\right)-\hat{\mathcal{T}}_{0}\left(g ; \pi_{1}\right)\right\|_{1}>2 \eta+d \beta\right) \leq d\left|\hat{\mathbb{Q}}_{\eta}\right| \exp \left(-\frac{\beta^{2} \kappa^{2}}{2}\right) \tag{39}
\end{equation*}
$$

Now, consider $h>1$. Consider any $\eta$ net over $\mathcal{B}_{d}(1)$, denoted by $\hat{\mathcal{B}}_{d, \eta}$ with respect to the norm $\|\cdot\|_{1}$. We can take $\left|\hat{\mathcal{B}}_{d, \eta}\right| \leq \exp (C d \log (d / \eta))$ Vershynin 2018). Invoking Lemma E. 2 , we conclude:

$$
\begin{align*}
& \sup _{\Pi \in \mathbb{Q}} \sup _{\nu \in \mathcal{B}_{d}(1)}\left|E^{\nu, h}(\Pi)-\hat{E}^{\nu, h}(\Pi)\right| \leq \sup _{\substack{\nu_{1}, \ldots, \nu_{h} \in \mathcal{B}_{d}(1) \\
\Pi \in \mathbb{Q}}}\left|\hat{F}\left(\Pi, \nu_{1}, \ldots, \nu_{h}\right)-F\left(\Pi, \nu_{1}, \ldots, \nu_{h}\right)\right| \\
& \leq \sup _{\substack{\nu_{1}, \ldots, \nu_{h} \in \hat{\mathcal{B}}_{d, \eta} \\
\Pi \in \hat{\mathbb{Q}}_{\eta}}}\left|\hat{F}\left(\Pi, \nu_{1}, \ldots, \nu_{h}\right)-F\left(\Pi, \nu_{1}, \ldots, \nu_{h}\right)\right|+2\left(1+C_{\mu}+\sqrt{\frac{T}{\kappa^{2}}}\right) \eta h \tag{40}
\end{align*}
$$

Now, by the triangle inequality, we have:

$$
\begin{align*}
& \left|\hat{F}\left(\Pi, \nu_{1}, \ldots, \nu_{h}\right)-F\left(\Pi, \nu_{1}, \ldots, \nu_{h}\right)\right| \\
& \quad \leq\left\|\mathcal{T}_{0}\left(\phi, \pi_{1}\right)-\hat{\mathcal{T}}_{0}\left(\phi, \pi_{1}\right)\right\|_{1}+\sum_{j=1}^{h-1}\left\|\mathcal{T}_{j}\left(\phi, \nu_{j}, \pi_{j+1}\right)-\hat{\mathcal{T}}_{j}\left(\phi, \nu_{j}, \pi_{j+1}\right)\right\|_{1} \tag{41}
\end{align*}
$$

Therefore, by invoking Lemma E. 1 , along with union bound over every component in the sum in 41) and over the net in (40) we conclude that:

$$
\begin{align*}
& \mathbb{P}\left[\sup _{\substack{\nu_{1}, \ldots, \nu_{h} \in \hat{\mathcal{B}}_{d, \eta} \\
\Pi \in \hat{\mathbb{Q}}_{\eta}}}\left|\hat{F}\left(\Pi, \nu_{1}, \ldots, \nu_{h}\right)-F\left(\Pi, \nu_{1}, \ldots, \nu_{h}\right)\right|>\beta d h\right] \\
& \leq 2 d h\left|\hat{\mathbb{Q}}_{\eta}\right|\left|\hat{\mathcal{B}}_{d, \eta}\right|^{h} \exp \left(-\frac{\beta^{2} \kappa^{2}}{2}\right) \tag{42}
\end{align*}
$$

Combining (40) and (42), we conclude the second item in the statement of the lemma.
The concentration of $X_{1}$ and $X_{h}$ follow in a similar fashion, but here we consider an $\eta$ net even over $x$ and use the Lipschitzness results given in Lemma E. 2 and the fact that $x \rightarrow f(\phi ; x)$ is 1 Lipschitz.

## E. Proof of Theorem 6.2

Lemma E.1. Suppose $h \in[H-1]$, and $g: \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ be such that $|g(s, a)| \leq B$ for every $(s, a)$. For any policy $\pi_{h}$ and any $\nu$ such that $\|\nu\|_{2} \leq 1$, we must have:

$$
\mathbb{P}\left(\left|\hat{\mathcal{T}}_{h}\left(g ; \nu, \pi_{h}\right)-\mathcal{T}_{h}\left(g ; \nu, \pi_{h}\right)\right|>\beta \mid\left(\phi_{h t}\right)_{t \in[T]}, G_{\phi, h} \succeq \kappa^{2} I\right) \leq 2 \exp \left(-\frac{\beta^{2} \kappa^{2}}{2 B^{2}}\right)
$$

Lemma E.2. Let $\Pi=\left(\pi_{1}, \ldots, \pi_{H}\right), \Pi^{\prime}=\left(\pi_{1}^{\prime}, \ldots, \pi_{H}^{\prime}\right)$ be policies in $\mathbb{Q}$. Conditioned on the event $G_{\phi, h} \succeq \kappa^{2} I$, the following hold:

$$
\begin{align*}
& \quad\left|F\left(\Pi, \nu_{1}, \ldots, \nu_{h-1}, \nu_{h}\right)-F\left(\Pi^{\prime}, \nu_{1}^{\prime}, \ldots, \nu_{h-1}^{\prime}, \nu_{h}^{\prime}\right)\right| \\
& \quad \leq\left(\sum_{j=2}^{h} C_{\mu} \operatorname{TV}\left(\pi_{j}, \pi_{j}^{\prime}\right)\left\|\nu_{j}\right\|_{1}+2\left\|\nu_{j}-\nu_{j}^{\prime}\right\|_{1}\right)+\operatorname{TV}\left(\pi_{1}, \pi_{1}^{\prime}\right)+2\left\|\nu_{1}-\nu_{1}^{\prime}\right\|_{1}  \tag{43}\\
& \left|\hat{F}\left(\Pi, \nu_{1}, \ldots, \nu_{h-1}, \nu_{h}\right)-\hat{F}\left(\Pi^{\prime}, \nu_{1}^{\prime}, \ldots, \nu_{h-1}^{\prime}, \nu_{h}^{\prime}\right)\right| \\
& \leq  \tag{44}\\
& \leq \sqrt{\frac{T}{\kappa^{2}}}\left(\sum_{j=2}^{h} \operatorname{TV}\left(\pi_{j}, \pi_{j}^{\prime}\right)\left\|\nu_{j}\right\|_{2}+\left\|\nu_{j}-\nu_{j}^{\prime}\right\|_{2}\right)+\operatorname{TV}\left(\pi_{1}, \pi_{1}^{\prime}\right)+\sum_{j=1}^{h}\left\|\nu_{j}-\nu_{j}^{\prime}\right\|_{1}
\end{align*}
$$

Suppose $x \in \mathcal{S}^{d-1}$

$$
\begin{align*}
& \left|\mathcal{T}_{h}\left(f(\cdot ; x), \nu, \pi_{h}\right)-\mathcal{T}_{h}\left(f\left(\cdot ; x^{\prime}\right), \nu^{\prime}, \pi_{h}^{\prime}\right)\right| \\
& \quad \leq 2 C_{\mu}(\sqrt{d}+\xi d)\left(\left\|\nu-\nu^{\prime}\right\|_{1}+\operatorname{TV}\left(\pi_{h}, \pi_{h}^{\prime}\right)+\left\|x-x^{\prime}\right\|_{2}\|\nu\|_{1}\right)  \tag{45}\\
& \left|\hat{\mathcal{T}}_{h}\left(f(\cdot ; x), \nu, \pi_{h}\right)-\hat{\mathcal{T}}_{h}\left(f\left(\cdot ; x^{\prime}\right), \nu^{\prime}, \pi_{h}^{\prime}\right)\right| \\
& \leq 2 \sqrt{\frac{T}{\kappa^{2}}}(\sqrt{d}+\xi d)\left(\left\|\nu-\nu^{\prime}\right\|_{2}+\operatorname{TV}\left(\pi_{h}, \pi_{h}^{\prime}\right)+\left\|x-x^{\prime}\right\|_{2}\|\nu\|_{2}\right) \tag{46}
\end{align*}
$$

Lemma E.3. Condition on the event $G_{\phi, h} \succeq \kappa^{2}$ I for every $h \in[H]$. Fix some $\eta>0$ and let $\hat{\mathbb{Q}}_{\eta}$ denote any $\eta$-net over $\mathbb{Q}$. With probability at-least $1-\delta / 4$, the following hold simultaneously:
1.

$$
\sup _{\nu} \sup _{\Pi \in \mathbb{Q}}\left|E_{1}^{\nu}(\Pi)-\hat{E}_{1}^{\nu}(\Pi)\right| \leq C \frac{d}{\kappa} \sqrt{\log \left(\frac{d\left|\hat{\mathbb{Q}}_{\eta}\right|}{\delta}\right)}+C \eta
$$

2. For $h>1$ :

$$
\sup _{\Pi \in \mathbb{Q}} \sup _{\nu \in \mathcal{B}_{d}(1)}\left|E_{h}^{\nu}(\Pi)-\hat{E}_{h}^{\nu}(\Pi)\right| \leq \frac{C d H}{\kappa} \sqrt{\log \left(\frac{d H\left|\hat{\mathbb{Q}}_{\eta}\right|}{\delta}\right)+H d \log \left(\frac{d}{\eta}\right)}+C\left(\sqrt{\frac{T}{\kappa^{2}}}\right) \eta H
$$

3. $X_{1}:=\sup _{\Pi=\left(\pi_{1}, \ldots, \pi_{H}\right) \in \mathbb{Q}}\left|\inf _{x \in \mathcal{S}^{d-1}} \hat{\mathcal{T}}_{1}\left(f(\cdot ; x), \pi_{1}\right)-\inf _{x \in \mathcal{S}^{d-1}} \mathcal{T}_{1}\left(f(\cdot ; x), \pi_{1}\right)\right|$

$$
X_{1} \leq \frac{C(\sqrt{d}+\xi d)}{\kappa} \sqrt{\log \left(\frac{\left|\hat{\mathbb{Q}}_{n}\right|}{\delta}\right)+d \log \left(\frac{d}{\eta}\right)}+C \eta(\sqrt{d}+\xi d)
$$

4. $X_{h}:=\sup _{\Pi=\left(\pi_{1}, \ldots, \pi_{H}\right) \in \mathbb{Q}}^{\nu \in \mathcal{B}(1)}\left|\inf _{x \in \mathcal{S}^{d-1}} \hat{\mathcal{T}}_{h}\left(f(\cdot ; x) ; \nu, \pi_{h}\right)-\inf _{x \in \mathcal{S}^{d-1}} \mathcal{T}_{h}\left(f(\cdot ; x) ; \nu, \pi_{h}\right)\right|$

$$
X_{h} \leq \frac{C(\sqrt{d}+\xi d)}{\kappa} \sqrt{\log \left(\frac{\left|\hat{\mathbb{Q}}_{\eta}\right| H}{\delta}\right)+d \log \left(\frac{d}{\eta}\right)}+C \eta(\sqrt{d}+\xi d)\left(C_{\mu}+\sqrt{\frac{T}{\kappa^{2}}}\right)
$$

Lemma E.4. $\Pi=\left(\pi_{1}, \ldots, \pi_{H}\right)$. For any $\eta \geq 0$, and $h \in[H]$, suppose $E_{h}^{\nu}(\Pi) \leq \eta$. Then, we have:

$$
\left\|\mathbb{E} \phi\left(S_{h}, A_{h}\right)-\nu\right\|_{1} \leq \eta
$$

Proof. Let $S_{1: H}, A_{1: H} \sim \mathcal{M}(\Pi)$. By Lemma 6.1. we conclude that: $\mathbb{E} \phi\left(S_{1}, A_{1}\right)=\mathcal{T}_{1}\left(\phi, \pi_{1}\right)$. Therefore we conclude the lemma for the case $h=1$. Now let $h>1$.
Now, note that for $j>1$, we have: $\mathbb{E} \phi\left(S_{j}, A_{j}\right)=\mathcal{T}_{j}\left(\phi, \mathbb{E} \phi\left(S_{j-1}, A_{j-1}\right), \pi_{j}\right)$. There exists a sequence $\nu_{1}, \ldots, \nu_{h-1}$ such that

$$
E_{1}^{\nu_{1}}(\Pi)+\sum_{j=2}^{h}\left\|\mathcal{T}_{j}\left(\phi, \nu_{j-1}, \pi_{j}\right)-\nu_{j}\right\|_{1} \leq \eta_{0}
$$

Letting $E_{1}^{\nu_{1}}(\Pi)=: \eta_{1},\left\|\mathcal{T}_{j}\left(\phi, \nu_{j-1}, \pi_{j}\right)-\nu_{j}\right\|_{1}=: \eta_{j}$, we have from the case $h=1:\left\|\mathbb{E} \phi\left(S_{1}, A_{1}\right)-\nu_{1}\right\|_{1} \leq \eta_{1}$.

$$
\begin{align*}
& \left\|\mathbb{E} \phi\left(S_{j}, A_{j}\right)-\nu_{j}\right\|_{1}=\left\|\mathcal{T}_{j}\left(\phi, \mathbb{E} \phi\left(S_{j-1}, A_{j-1}\right), \pi_{j}\right)-\nu_{j}\right\|_{1} \\
& \leq\left\|\mathcal{T}_{j}\left(\phi, \mathbb{E} \phi\left(S_{j-1}, A_{j-1}\right), \pi_{j}\right)-\mathcal{T}_{j}\left(\phi, \nu_{j-1}, \pi_{j}\right)\right\|_{1}+\left\|\mathcal{T}_{j}\left(\phi, \nu_{j-1}, \pi_{j}\right)-\nu_{j}\right\|_{1} \\
& \leq\left\|\nu_{j-1}-\mathbb{E} \phi\left(S_{j-1}, A_{j-1}\right)\right\|_{1}+\eta_{j} \tag{47}
\end{align*}
$$

We have used $\sqrt{32}$ in the last step. Continuing recursively, we conclude the result

Proof of Theorem 6.2 We condition on the event described in LemmaE.3. We suppose that $\kappa, \eta$ and $\eta_{0}$ are related as in the statement of the theorem. We will apply these values whenever we invoke the concentration bounds obtained from LemmaE. 3 in all the inequalities below. First consider $h=1$. Let $\hat{\Pi}^{f, 1}=\left(\pi_{H}^{f, 1}, \ldots, \pi_{H}^{f, H}\right)$. By item 3 in Lemma E. 3 we have (with $X_{1}$ as defined in the lemma):

$$
\inf _{x \in \mathcal{S}^{d-1}} \hat{\mathcal{T}}_{1}\left(f(; x), \pi_{1}^{f, 1}\right) \geq \sup _{\Pi \in \mathbb{Q}} \inf _{x \in \mathcal{S}^{d-1}} \mathcal{T}_{1}\left(f(; x), \pi_{1}\right)-X_{1} \geq \zeta-X_{1} \geq \frac{3 \zeta}{4}
$$

Similarly, we have:

$$
\inf _{x \in \mathcal{S}^{d-1}} \mathcal{T}_{1}\left(f(; x), \pi_{1}^{f, 1}\right) \geq \inf _{x \in \mathcal{S}^{d-1}} \hat{\mathcal{T}}_{1}\left(f(; x), \pi_{1}^{f, 1}\right)-\frac{\zeta}{4}
$$

Combining the two displays above, we conclude the theorem for $h=1$. Now consider $h>1$. We will first show that the constraint $\hat{E}^{\nu, h-1}(\Pi) \leq \eta_{0}$ is feasible for some $\Pi \in \mathbb{Q}$ and some $\nu$. Note that, for any policy $\Pi$ there exists a $\nu_{1}, \ldots, \nu_{h-1} \in \mathbb{R}^{d}$ such that $\mathbb{E} \phi\left(S_{j}, A_{j}\right)=\nu_{j}$ whenever $S_{1: H}, A_{1: H} \sim \mathcal{M}(\Pi)$. For the choice $\nu=\nu_{h-1}$, we must have $E^{\nu, h-1}(\Pi)=0$. Now, by item 1 and 2 of Lemma E.3. we conclude that $\hat{E}^{\nu, h-1} \leq \eta_{0}$. Therefore this optimization is feasible.
Consider the solutions to the optimization problem given by $\hat{\nu}$ and $\hat{\Pi}^{f, h}$. Note again from LemmaE. 3 that $E^{\hat{\nu}, h-1}\left(\hat{\Pi}^{f, h}\right) \leq$ $\hat{E}^{\hat{\nu}, h-1}\left(\hat{\Pi}^{f, h}\right)+\eta_{0} \leq 2 \eta_{0}$. Now, applying Lemma E.4, we conclude that whenever $S_{1: H}, A_{1: H} \sim \mathcal{M}\left(\hat{\Pi}^{f, h}\right)$

$$
\left\|\mathbb{E} \phi\left(S_{h-1}, A_{h-1}\right)-\hat{\nu}\right\|_{1} \leq 2 \eta_{0}
$$

By a similar reasoning as the case $h=1$, we conclude:

$$
\inf _{x \in \mathcal{S}^{d-1}} \mathcal{T}_{h}\left(f(; x), \hat{\nu}, \pi_{h}^{f, h}\right) \geq 3 \frac{\zeta}{4}
$$

Now, applying (45), we conclude:

$$
\begin{align*}
& \inf _{x \in \mathcal{S}^{d-1}} \mathbb{E} f\left(S_{h}, A_{h} ; x\right)=\inf _{x \in \mathcal{S}^{d-1}} \mathcal{T}_{h}\left(f(; x), \mathbb{E} \phi\left(S_{h-1}, A_{h-1}\right), \pi_{h}^{f, h}\right) \\
& \geq \inf _{x \in \mathcal{S}^{d-1}} \mathcal{T}_{h}\left(f(; x), \hat{\nu}, \pi_{h}^{f, h}\right)-\sup _{x \in \mathcal{S}^{d-1}}\left|\mathcal{T}_{h}\left(f(; x), \mathbb{E} \phi\left(S_{h-1}, A_{h-1}\right), \pi_{h}^{f, h}\right)-\mathcal{T}_{h}\left(f(; x), \hat{\nu}, \pi_{h}^{f, h}\right)\right| \\
& \geq \frac{3 \zeta}{4}-2 C_{\mu}(\sqrt{d}+\xi d)\left(\left\|\hat{\nu}-\mathbb{E} \phi\left(S_{h-1}, A_{h-1}\right)\right\|_{1}\right) \geq \frac{\zeta}{2} \tag{48}
\end{align*}
$$

In the last step, we have used the lipschitzness bound for $\mathcal{T}_{h}$ given in Lemma E. 2 We will show that the conditions given in (4) are satisfied for $\psi\left(S_{h}, A_{h}\right)$ with parameters $\zeta / 2$ instead of $\zeta$.
$\left\|\psi\left(S_{h}, A_{h}\right)\right\|_{2} \leq 1$ almost surely follows from the definition of $\psi$. Now, $\mathbb{E} f\left(S_{h}, A_{h}, x\right) \geq \frac{\zeta}{2}$ for every $x \in \mathcal{S}^{d-1}$ implies $\mathbb{E}\left|\left\langle x, \psi\left(S_{h}, A_{h}\right)\right\rangle\right| \geq \frac{\zeta}{2 \sqrt{d}}$. Using the definition of $f\left(S_{h}, A_{h}, x\right)$ (see Section 6) and the fact that $\mathbb{E} f\left(S_{h}, A_{h}, x\right) \geq \frac{\zeta}{2}$ as established above, we conclude that for every $x \in \mathcal{S}^{d-1}$, we also have:

$$
\begin{align*}
d \xi \mathbb{E}\left\langle x, \psi\left(S_{h}, A_{h}\right)\right\rangle^{2} & \leq \sqrt{d} \mathbb{E}\left|\left\langle x, \psi\left(S_{h}, A_{h}\right)\right\rangle\right|-\frac{\zeta}{2} \\
& \leq \sqrt{d} \mathbb{E}\left|\left\langle x, \psi\left(S_{h}, A_{h}\right)\right\rangle\right| \leq \sqrt{d} \sqrt{\mathbb{E}\left\langle x, \psi\left(S_{h}, A_{h}\right)\right\rangle^{2}} \tag{49}
\end{align*}
$$

In the second step, we have used Jensen's inequality. From this, we conclude $\mathbb{E}\left\langle x, \psi\left(S_{h}, A_{h}\right)\right\rangle^{2} \leq \frac{1}{d \xi^{2}}$ for every $x \in \mathcal{S}^{d-1}$ and thence $\mathbb{E} \psi\left(S_{h}, A_{h}\right) \psi\left(S_{h}, A_{h}\right)^{\top} \preceq \frac{1}{d \xi^{2}}$.

## F. Proof of Theorem 7.1

Let the unknown row set in the iteration $t$ in the matrix estimation procedure of Section 7.1 be denoted by $\bar{I}_{t-1}$. For the analysis, we will use the convention that $\bar{I}_{t}=\emptyset$ if the procedure terminates before the $t$-th iteration. Suppose $K_{t}$ is such that for every $t \leq \log N$, we have: $K_{t}\left|\bar{I}_{t-1}\right| \geq C \frac{r\left|\bar{I}_{t-1}\right|+d r}{\zeta^{2} \xi^{2}} \log \frac{d}{\zeta \xi}+C \frac{\log \left(\frac{\log N}{\delta}\right)}{\zeta^{2} \xi^{2}}$. We will then show that the event $\left\{\left|\bar{I}_{t}\right| \leq \frac{1}{10}\left|\bar{I}_{t-1}\right| \forall t \leq \log N\right\} \cap\left\{\hat{\Theta}_{i}=\Theta_{i}^{*}, \forall i \in \bar{I}_{\log N}^{\complement}\right\}$ has probability at-least $1-\delta$. To show this, it is sufficient to
consider the step $t=1$ with $\bar{I}_{0}=[N], K_{1}=K, \Psi^{(1)}=\Psi$ and show that with probability $1-\frac{\delta}{\log N}, \bar{I}_{1} \leq \frac{9 N}{10}$ and $\hat{\Theta}_{i}=\Theta_{i}^{*}$ for every $i \in \bar{I}_{1}^{\complement}$. The result then follows from a union bound. We will therefore establish the following structural lemma and prove the Theorem 7.1. The rest of the section is then dedicated to proving Lemma F.1.
Lemma F.1. Suppose the distribution of $\psi_{i k}$ satisfies 44. Let $K \geq \frac{C\left(r+\frac{d r}{N}\right)}{\zeta^{2} \xi^{2}} \log \frac{d}{\zeta \xi}$. Let $\mathcal{Y}(\Psi)$ denote the set of all matrices $\Delta$ with rank at most $2 r$ such that $L(\Delta, \Psi)=0$. Let $I_{\mathcal{Z}}(\Delta)=\left\{i \in[N]: \Delta_{i} \neq 0\right\}$. With probability $1-\exp \left(-c \zeta^{2} \xi^{2} N K\right)$ we must have:

$$
\mathcal{Y}(\Psi) \cap\left\{\Delta:\left|I_{\mathcal{Z}}(\Delta)\right|>\frac{N}{10}\right\}=\emptyset
$$

Proof of Theorem 7.1 Let $\bar{\Theta}$ be the rank $\leq r$ matrix found satisfying $L\left(\bar{\Theta}-\Theta^{*}, \Psi^{(t)}\right)=0$. By Lemma F.1, we have that $\left|I_{\mathcal{Z}}\left(\bar{\Theta}-\Theta^{*}\right)\right| \leq \frac{N}{10}$ with probability at-least $1-\frac{\delta}{\log N}$ (by setting $K=K_{1}$ as in the statement of Theorem 7.1 . By Lemma F.2. the probability that there exists an $i \in[N]$ such that $\bar{\Theta}_{i} \neq \Theta_{i}^{*}$, and $\sum_{k=1}^{K}\left|\left\langle\bar{\Theta}_{i}, \tilde{\psi}_{i k}\right\rangle-\theta_{i k}^{*}\right|^{2}=0$ is at most $|I| \cdot \exp \left(-c \zeta^{2} \xi^{2} K\right) \leq \delta \cdot N^{-c}$ for some large constant $c$. From this we conclude that $\bar{\Theta}_{i}=\Theta_{i}^{*}$ for every $i \in \bar{I}_{1}^{\complement}$.

Lemma F.2. Fix any $\hat{\Theta}$. Suppose the distribution of $\left(\psi_{i k}\right)_{i \in[N], k \in[K]}$ satisfies 4]. Then, there exists a small enough constant $c$ such that:

$$
\mathbb{P}\left(\exists i \text { s.t. } \sum_{k=1}^{K}\left|\left\langle\hat{\Theta}_{i}, \psi_{i k}\right\rangle-\theta_{i k}^{*}\right|^{2}<\frac{K \zeta^{4} \xi^{2}\left\|\hat{\Theta}_{i}-\Theta_{i}^{*}\right\|^{2}}{32 d}\right) \leq|I| \cdot \exp \left(-c \zeta^{2} \xi^{2} K\right)
$$

Proof. Consider the Paley-Zygmund inequality, which states that for any positive random variable $Z$,

$$
\mathbb{P}\left(Z \geq \frac{\mathbb{E} Z}{2}\right) \geq \frac{(\mathbb{E} Z)^{2}}{4 \mathbb{E} Z^{2}}
$$

Suppose $i \in I$ and denote $\Gamma_{i}:=\hat{\Theta}_{i}-\Theta_{i}^{*}$. By the properties of $\psi_{i k}$, we have that $\mathbb{E}\left|\left\langle\Gamma_{i}, \psi_{i k}\right\rangle\right| \geq \frac{\zeta\left\|\Gamma_{i}\right\|}{\sqrt{d}}$ and $\mathbb{E}\left|\left\langle\Gamma_{i}, \psi_{i k}\right\rangle\right|^{2} \leq$ $\frac{\left\|\Gamma_{i}\right\|^{2}}{\xi^{2} d}$.
Applying the Paley-Zygmund inequality to the random variable $\left|\left\langle\Gamma_{i}, \psi_{i k}\right\rangle\right|$, we conclude the result in 53 ):

$$
\begin{equation*}
\mathbb{P}\left(\left|\left\langle\psi_{i k}, \Gamma_{i}\right\rangle\right| \geq \frac{\zeta\left\|\Gamma_{i}\right\|}{2 \sqrt{d}}\right) \geq \frac{\zeta^{2} \xi^{2}}{4} \tag{50}
\end{equation*}
$$

Let $p_{0}:=\frac{\zeta^{2} \xi^{2}}{4}$. Let $N\left(\Gamma_{i}, \Psi_{i}\right):=\sum_{k=1}^{K} \mathbb{1}\left(\left|\left\langle\psi_{i k}, \Gamma_{i}\right\rangle\right|>\frac{\zeta}{2 \sqrt{d}}\right)$. Clearly, $\sum_{k=1}^{K}\left|\left\langle\hat{\Theta}_{i}, \psi_{i k}\right\rangle-\theta_{i k}^{*}\right|^{2} \geq \frac{\zeta^{2}\left\|\Gamma_{i}\right\|^{2}}{4 d} \cdot N\left(\Gamma_{i}, \Psi_{i}\right)$. Therefore, we have:

$$
\begin{align*}
\mathbb{P}\left(\sum_{k=1}^{K}\left|\left\langle\hat{\Theta}_{i}, \psi_{i k}\right\rangle-\theta_{i k}^{*}\right|^{2}<\frac{K \zeta^{4} \xi^{2}\left\|\Gamma_{i}\right\|^{2}}{32 d}\right) & \leq \mathbb{P}\left(N\left(\Gamma_{i}, \Psi_{i}\right)<\frac{K \zeta^{2} \xi^{2}}{8}\right) \\
& \leq \mathbb{P}\left(\operatorname{Bin}\left(K, p_{0}\right) \leq \frac{K p_{0}}{2}\right) \\
& \leq \exp \left(-c p_{0} K\right) \tag{51}
\end{align*}
$$

Here $\operatorname{Bin}\left(K, p_{0}\right)$ denotes the binomial random variable. In the second step we have used the fact that $N\left(\Gamma_{i}, \Psi_{i}\right)$ is a sum of $K$ independent Bernoulli random variables with probability of being 1 for each of them being at-least $p_{0}=\frac{\zeta^{2} \xi^{2}}{4}$. In the last step, we have used Sanov's theorem for large deviations. In the last step we have used Bernstein's inequality for concentration of sums of Bernoulli random variables (see (Boucheron et al. 2013)). The statement of the result then follows from a union bound argument over $i \in I$.

## F.1. Proof of Lemma $\mathbf{F . 1}$

Suppose $I \neq \emptyset, I \subseteq[N]$ be any fixed subset. By $\mathcal{M}(N, d, I, 2 r)$, we denote the set of all $N \times d$ matrices $\Delta$ with rank at-most $2 r$ such that $\left\|\Delta_{i}\right\|>0$ for all $i \in I$. By $\mathcal{B}(N, d, I, 2 r)$ we denote the set of all $N \times d$ matrices with rank at-most $2 r$ such that $\left\|\Gamma_{i}\right\|=1$ whenever $i \in I$.
Lemma F.3. Suppose $\inf _{\Gamma \in \mathcal{B}(N, d, I, 2 r)} L(\Gamma, \Psi)>0$. Then, $L(\Delta, \Psi)>0$ for every $\Delta \in \mathcal{M}(N, d, I, 2 r)$

Proof. For every $\Delta \in \mathcal{M}(N, d, I, 2 r)$, we construct $\Gamma$ such that:

$$
\Gamma_{i}= \begin{cases}\frac{\Delta_{i}}{\left\|\Delta_{i}\right\|} & \text { whenever } i \in I  \tag{52}\\ 0 & \text { otherwise }\end{cases}
$$

Now, by hypothesis, $L(\Gamma, \Psi)>0$. This implies, there exists an $i \in I$ and $k \in K$ such that $\left|\left\langle\psi_{i k}, \Gamma_{i}\right\rangle\right|>0$. This implies $\left|\left\langle\psi_{i k}, \Delta_{i}\right\rangle\right|>0$ and thence we conclude that $L(\Delta, \Psi)>0$.

Lemma F.4. Suppose $\Gamma$ is such that $\left\|\Gamma_{i}\right\|=1$ for every $i \in I$. Suppose the distribution of $\left(\psi_{i k}\right)_{i \in[N], k \in[K]}$ satisfy (4). Then, there exists a small enough constant $c$ such that:

$$
\mathbb{P}\left(L(\Gamma, \Psi)<\frac{|I| \zeta^{4} \xi^{2}}{32 N d}\right) \leq|I|^{2} K^{2} \exp \left(-c \zeta^{2} \xi^{2}|I| K\right)
$$

Proof. Consider the Paley-Zygmund inequality, which states that for any positive random variable $Z$,

$$
\mathbb{P}\left(Z \geq \frac{\mathbb{E} Z}{2}\right) \geq \frac{(\mathbb{E} Z)^{2}}{4 \mathbb{E} Z^{2}}
$$

Suppose $i \in I$. By the properties of $\psi_{i k}$, we have that $\mathbb{E}\left|\left\langle\Gamma_{i}, \psi_{i k}\right\rangle\right| \geq \frac{\zeta}{\sqrt{d}}$ and $\mathbb{E}\left|\left\langle\Gamma_{i}, \psi_{i k}\right\rangle\right|^{2} \leq \frac{1}{\xi^{2} d}$
Applying the Paley-Zygmund inequality to the random variable $\left|\left\langle\Gamma_{i}, \psi_{i k}\right\rangle\right|$, we conclude the result in 53):

$$
\begin{equation*}
\mathbb{P}\left(\left|\left\langle\psi_{i k}, \Gamma_{i}\right\rangle\right| \geq \frac{\zeta}{2 \sqrt{d}}\right) \geq \frac{\zeta^{2} \xi^{2}}{4} \tag{53}
\end{equation*}
$$

Let $p_{0}:=\frac{\zeta^{2} \xi^{2}}{4}$. Let $N(\Gamma, \Psi):=\sum_{i \in I} \sum_{k=1}^{K} \mathbb{1}\left(\left|\left\langle\psi_{i k}, \Gamma_{i}\right\rangle\right|>\frac{\zeta}{2 \sqrt{d}}\right)$. Clearly, $L(\Gamma, \Psi) \geq \frac{\zeta^{4}}{4 d N K} N(\Gamma, \Psi)$ almost surely. Therefore, we have:

$$
\begin{align*}
\mathbb{P}\left(L(\Gamma, \Psi)<\frac{|I| \zeta^{4} \xi^{2}}{32 N \sqrt{d}}\right) & \leq \mathbb{P}\left(N(\Gamma, \Psi)<\frac{|I| K \zeta^{2} \xi^{2}}{8}\right) \\
& \leq \mathbb{P}\left(\operatorname{Bin}\left(|I| K, p_{0}\right) \leq \frac{|I| K p_{0}}{2}\right) \\
& \leq \exp \left(-c p_{0}|I| K\right) \tag{54}
\end{align*}
$$

Here $\operatorname{Bin}\left(|I| K, p_{0}\right)$ denotes the binomial random variable. In the second step we have used the fact that $N(\Gamma, \Psi)$ is a sum of $|I| K$ independent Bernoulli random variables with probability of being 1 for each of them being at-least $p_{0}=\frac{\zeta^{2} \xi^{2}}{4}$. In the last step, we have used Sanov's theorem for large deviations. In the last step we have used Bernstein's inequality for concentration of sums of Bernoulli random variables (see (Boucheron et al., 2013))

Lemma F.5. Suppose the distribution of $\left(\psi_{i k}\right)_{i \in[N], k \in[K]}$ satisfy (4]. Let $|I| \geq \frac{N}{10}$. There exist positive constants $c_{0}, c, C$ such that whenever $K N \geq \frac{C r(N+d)}{\zeta^{2} \xi^{2}} \log \frac{d}{\zeta \xi}$, we have:

$$
\mathbb{P}\left(\inf _{\Gamma \in \mathcal{B}(N, d, I, 2 r)} L(\Gamma, \Psi)<c_{0} \frac{\zeta^{4} \xi^{2}}{d}\right) \leq \exp \left(-c \zeta^{2} \xi^{2} N K\right)
$$

Proof. It is sufficient to prove this result for $\Gamma \in \mathcal{B}_{0}(N, d, I, 2 r) \subseteq \mathcal{B}(N, d, I, 2 r)$, which is the set of all matrices such that $\left\|\Gamma_{i}\right\|=1$ for every $i \in I$ and 0 otherwise. Define $\|\Gamma\|_{1,2, \boldsymbol{T}}:=\frac{1}{N} \sum_{i=1}^{N}\left\|\Gamma_{i}\right\|$. Suppose $\hat{\Gamma} \in \mathcal{B}_{0}(N, d, I, 2 r)$ is such that $\|\Gamma-\hat{\Gamma}\|_{1,2, \boldsymbol{\top}}<\eta$. Then,

$$
\begin{align*}
L(\Gamma, \Psi) & =\frac{1}{N K} \sum_{i=1}^{N} \sum_{k=1}^{N}\left|\left\langle\Gamma_{i}, \psi_{i k}\right\rangle\right|^{2} \\
& \geq \frac{1}{N K} \sum_{i=1}^{N} \sum_{k=1}^{N}\left|\left\langle\hat{\Gamma}_{i}, \psi_{i k}\right\rangle\right|^{2}-2\left|\left\langle\hat{\Gamma}_{i}-\Gamma_{i}, \psi_{i k}\right\rangle\right|\left|\left\langle\hat{\Gamma}_{i}, \psi_{i k}\right\rangle\right|  \tag{55}\\
& =L(\hat{\Gamma}, \Psi)-\|\Gamma-\hat{\Gamma}\|_{1,2, \mathrm{\top}} \geq L(\hat{\Gamma}, \Psi)-2 \eta \tag{56}
\end{align*}
$$

In the third step, we have used the fact that $\left\|\psi_{i k}\right\| \leq 1$ and the Cauchy-Schwarz inequality to imply $\left|\left\langle\hat{\Gamma}_{i}-\Gamma_{i}, \psi_{i k}\right\rangle\right| \leq$ $\left\|\hat{\Gamma}_{i}-\Gamma_{i}\right\|$. Therefore, given any $\eta$ net of $\mathcal{B}_{0}(N, d, I, 2 r)$, denoted by $\hat{\mathcal{B}}_{0, \eta}$, we must have:

$$
\begin{equation*}
\inf _{\Gamma \in \mathcal{B}_{0}(N, d, I, 2 r)} L(\Gamma, \Psi) \geq \inf _{\Gamma \in \hat{\mathcal{B}}_{0, \eta}} L(\hat{\Gamma}, \Psi)-2 \eta \tag{57}
\end{equation*}
$$

We will now parametrize $\mathcal{B}_{0}(N, d, I, 2 r)$ as follows:
Claim F.6. Every $\Gamma \in \mathcal{B}_{0}(N, d, I, 2 r)$ can be written as

$$
\Gamma_{i}=\left\{\begin{array}{l}
\sum_{k=1}^{2 r} u_{i k} v_{k} \text { if } i \in I  \tag{58}\\
0 \text { otherwise }
\end{array}\right.
$$

Where $v_{1}, \ldots, v_{2 r}$ are orthonormal vectors in $\mathbb{R}^{d}$ and $u_{i}=\left(u_{i k}\right)_{k=1}^{2 r} \in \mathbb{R}^{2 r}$ are such that $\left\|u_{i}\right\|=1$.
Proof. By the singular value decomposition, we have: $\Gamma=W \Sigma V^{\top}$ for orthogonal matrices $W, V$ and the singular value matrix $\Sigma$. Therefore, $\Gamma_{i j}=\sum_{k=1}^{2 r} w_{i k} \sigma_{k} v_{k j}$ Denoting $u_{i k}:=w_{i k} \sigma_{k}$, we note that $\Gamma_{i}=\sum_{k=1}^{2 r} u_{i k} v_{k}$, where $v_{k}$ is the $k$-th column of $V$.
Now, it remains to show that $\left\|u_{i}\right\|=1$. By ortho-normality of $v_{1}, \ldots, v_{2 r}$ and the definition of $\Gamma$, we have: $1=\left\|\Gamma_{i}\right\|^{2}=$ $\sum_{k=1}^{2 r}\left|u_{i k}\right|^{2}=\left\|u_{i}\right\|^{2}$

Therefore, we construct an $\eta$-net for $\mathcal{B}_{0}(N, d, I, 2 r)$ as follows: consider any $\eta / 2$-net over the sphere $\mathcal{S}^{2 r-1}$, denoted by $\hat{\mathcal{S}}_{\frac{\eta}{2}}(2 r)$ with respect to the Euclidean norm. Similarly, consider any $\frac{\eta}{2 \sqrt{2 r}}$-net over the sphere $\mathcal{S}^{d-1}$, denoted by $\hat{\mathcal{S}}_{\frac{\eta}{2 \sqrt{2 r}}}(d)$. We draw $\left(u_{i}\right)_{i \in I},\left(v_{k}\right)_{k \in[2 r]}$ from the set $\prod_{i \in I} \hat{\mathcal{S}}_{\frac{\eta}{2}}(2 r) \prod_{k \in[2 r]} \hat{\mathcal{S}}_{\frac{\eta}{2 \sqrt{2 r}}}(d)$ and take $\hat{\mathcal{B}}_{0, \eta}$ to be the set of all $\hat{\Gamma}(u, v)$ of the form given in Claim F. 6
Claim F.7. $\hat{\mathcal{B}}_{0, \eta}$ is an $\eta$ net for $\mathcal{B}_{0}(N, d, I, 2 r)$ with respect to the norm $\|\cdot\|_{1,2, \mathrm{~T}}$.

$$
\left|\hat{\mathcal{B}}_{0, \eta}\right| \leq \exp \left(2 d r \log \left(\frac{4 \sqrt{2 r}}{\eta}+1\right)+2|I| r \log \left(\frac{4}{\eta}+1\right)\right)
$$

Proof of ClaimF. 7 Let $\Gamma \in \mathcal{B}_{0}(N, d, I, 2 r)$. Let $\left(u_{i}\right),\left(v_{k}\right)$ be such that: ClaimF.6, $\Gamma_{i}=\sum_{k=1}^{2 r} u_{i k} v_{k}$. By construction, there exists $\hat{\Gamma} \in \hat{\mathcal{B}}_{0, \eta}$ such that:

$$
\hat{\Gamma}_{i}=\sum_{k=1}^{2 r} \hat{u}_{i k} \hat{v}_{k}
$$

with $\left\|u_{i}-\hat{u_{i}}\right\| \leq \frac{\eta}{2}$ and $\left\|v_{k}-\hat{v_{k}}\right\| \leq \frac{\eta}{2 \sqrt{2 r}}$ for every $i \in I$ and $k \in[2 r]$.
In order to show that $\|\Gamma-\hat{\Gamma}\|_{1,2, \mathrm{~T}} \leq \eta$, it is sufficient to show that $\left\|\hat{\Gamma}_{i}-\hat{\Gamma}_{i}\right\| \leq \eta$ for every $i \in[I]$.

$$
\begin{align*}
\left\|\hat{\Gamma}_{i}-\Gamma_{i}\right\| & =\left\|\sum_{k=1}^{2 r}\left(\hat{u}_{i k}-u_{i k}\right) v_{k}+\sum_{k=1}^{2 r} u_{i k}\left(v_{k}-\hat{v}_{k}\right)\right\| \\
& \leq\left\|\sum_{k=1}^{2 r}\left(\hat{u}_{i k}-u_{i k}\right) v_{k}\right\|+\left\|\sum_{k=1}^{2 r} u_{i k}\left(v_{k}-\hat{v}_{k}\right)\right\| \\
& =\sqrt{\sum_{k=1}^{2 r}\left(\hat{u}_{i k}-u_{i k}\right)^{2}}+\left\|\sum_{k=1}^{2 r} u_{i k}\left(v_{k}-\hat{v}_{k}\right)\right\| \leq \frac{\eta}{2}+\left\|\sum_{k=1}^{2 r} u_{i k}\left(v_{k}-\hat{v}_{k}\right)\right\| \\
& \leq \frac{\eta}{2}+\sqrt{\sum_{k=1}^{2 r} u_{i k}^{2}} \sqrt{\sum_{k=1}^{2 r}\left\|v_{k}-\hat{v}_{k}\right\|^{2}} \leq \eta \tag{59}
\end{align*}
$$

Therefore $\hat{\mathcal{B}}_{0, \eta}$ is an $\eta$ net with respect to $\|\cdot\|_{1,2, \mathrm{~T}}$. By Corollary 4.2.13 in Vershynin 2018, we can pick: $\left|\hat{\mathcal{S}} \frac{\eta}{2 \sqrt{2 r}}(d)\right| \leq$ $\left(\frac{4 \sqrt{2 r}}{\eta}+1\right)^{d}$ and $\left|\hat{\mathcal{S}}_{\frac{\eta}{2}}(2 r)\right| \leq\left(\frac{4}{\eta}+1\right)^{2 r}$ and conclude the bound on the cardinality of $\hat{\mathcal{B}}_{0, \eta}$.

By Lemma F. 4 and a union bound,

$$
\begin{align*}
& \mathbb{P}\left(\inf _{\hat{\Gamma} \in \hat{B}_{0, \eta}} L(\hat{\Gamma}, \Psi)<\frac{\zeta^{4} \xi^{2}|I|}{32 N d}\right) \leq\left|\hat{\mathcal{B}}_{0, \eta}\right| \exp \left(-c \zeta^{2} \xi^{2}|I| K\right) \\
& \leq \exp \left(2 d r \log \left(\frac{4 \sqrt{2 r}}{\eta}+1\right)+2|I| r \log \left(\frac{4}{\eta}+1\right)-c \zeta^{2} \xi^{2}|I| K\right) \tag{60}
\end{align*}
$$

Therefore, whenever taking $|I| \geq \frac{N}{10}$ and $\eta=c_{1} \frac{\zeta^{4} \xi^{2}}{d}$ for some constant $c_{1}$ small enough, and combining (60) with (56), we conclude that whenever $K \geq \frac{C\left(r+\frac{d r}{N}\right)}{\zeta^{2} \xi^{2}} \log \frac{d}{\zeta \xi}$ for a large enough constant $C$, we have:

$$
\mathbb{P}\left(\inf _{\Gamma \in \mathcal{B}_{0}(N, d, I, 2 r)} L(\Gamma, \Psi)<c_{0} \frac{\zeta^{4} \xi^{2}}{d}\right) \leq \exp \left(-c \zeta^{2} \xi^{2} N K\right)
$$

Now, consider $|I| \geq \frac{N}{10}$. The number of such sets $I$ is at-most $\exp \left(c_{1} N\right)$ for some constant $c_{1}>0$. Therefore, applying Lemma F. 5 along with the union bound over all $I$ such that $|I| \geq \frac{N}{10}$ we have:
Corollary F.8. Under the conditions of Lemma F.5, we have:

$$
\inf _{\substack{I \subseteq N \\|I| \geq \frac{N}{10}}} \inf _{\mathcal{B}(N, d, I, 2 r)} L(\Gamma, \Psi)>c_{0} \frac{\zeta^{4} \xi^{2}}{d}
$$

with probability at-least $1-\exp \left(-c \zeta^{2} \xi^{2} N K\right)$
We are now ready to prove Lemma F.1.
Proof of Lemma F.1. Combining Lemma F. 3 and Corollary F.8, we conclude that with probability at-least 1 $\exp \left(-c \zeta^{2} \xi^{2} N K\right)$,

$$
\mathcal{Y}(\Psi) \bigcap\left(\bigcup_{\substack{I \subseteq[N] \\|I| \geq N \\ 10}} \mathcal{N}(N, d, I, 2 r)\right)=\emptyset
$$

Note that if $\Delta \in \mathcal{Y}(\Psi)$ such that $\left|I_{\mathcal{Z}}(\Delta)\right|>\frac{N}{10}$ implies $\Delta \in \mathcal{M}(N, d, I, 2 r)$ for some $|I|>\frac{N}{10}$. This allows us to conclude the statement of the lemma.


[^0]:    *Equal contribution ${ }^{1}$ Google Research, Bangalore ${ }^{2}$ Amazon India ${ }^{3}$ Work was done prior to joining Amazon ${ }^{4}$ Google Research, Princeton. Correspondence to: Dheeraj Nagaraj [dheerajnagaraj@google.com](mailto:dheerajnagaraj@google.com).

    Proceedings of the $40^{\text {th }}$ International Conference on Machine Learning, Honolulu, Hawaii, USA. PMLR 202, 2023. Copyright 2023 by the author(s).

