

# Linear Bellman Completeness Suffices for Efficient Online Reinforcement Learning with Few Actions

Noah Golowich

Ankur Moitra

*Massachusetts Institute of Technology*

NZG@MIT.EDU

MOITRA@MIT.EDU

**Editors:** Shipra Agrawal and Aaron Roth

## Abstract

One of the most natural approaches to reinforcement learning (RL) with function approximation is *value iteration*, which inductively generates approximations to the optimal value function by solving a sequence of regression problems. To ensure the success of value iteration, it is typically assumed that *Bellman completeness* holds, which ensures that these regression problems are well-specified. We study the problem of learning an optimal policy under Bellman completeness in the *online* model of RL with linear function approximation. In the linear setting, while statistically efficient algorithms are known under Bellman completeness (e.g., (Jiang et al., 2017; Zanette et al., 2020a)), these algorithms all rely on the principle of *global optimism* which requires solving a nonconvex optimization problem. In particular, it has remained open as to whether *computationally efficient* algorithms exist. In this paper we give the first polynomial-time algorithm for RL under linear Bellman completeness when the number of actions is any constant.

**Keywords:** Reinforcement learning, Linear Bellman completeness, Optimism

## 1. Introduction

*Reinforcement learning (RL)* describes the problem of solving sequential decision-making problems in a stochastically changing environment, and is typically studied using the formalism of *Markov Decision Processes (MDPs)*. In an MDP, a learning agent must choose a sequence of *actions* over some number of time steps, each of which affects the *state* of the environment and potentially yields some *reward* to the agent. The agent aims to find a *policy*, or a mapping that describes which action to take at each state, that maximizes its expected total reward. In order for RL to yield effective learning strategies in its various application domains, including robotics (Gu et al., 2017), economics (Zheng et al., 2022), and healthcare (Yu et al., 2021), it is necessary to come up with efficient strategies for exploring complex state spaces, which may be infinite or exponentially large.

A general approach to RL, which dates back decades (Bradtke and Barto, 1996; Melo and Ribeiro, 2007; Sutton and Barto, 2018) and yet still forms the basis for many current empirical approaches (Hasselt et al., 2016; Schulman et al., 2017), involves the use of value function approximation. Recall that the optimal *value function* maps a state-action pair to the agent’s expected reward under the optimal policy starting from that state-action pair. Then this approach posits that the optimal value function belongs to some known function class  $\mathcal{F}$ , such as a class of linear functions or neural networks. A key question is: *under what assumptions on the class  $\mathcal{F}$  can we efficiently learn a near-optimal policy, i.e., one with near-maximal expected reward?*

**Value iteration and Bellman completeness.** A popular and time-tested approach to finding a near-optimal policy with value function approximation is *value iteration*. To explain this procedure, we consider the *finite-horizon* setting, whereby interactions with the environment occur in *episodes* lasting  $H$  time steps. Letting  $\mathcal{X}$  denote the state space and  $\mathcal{A}$  denote the action space, value iteration computes mappings  $\hat{Q}_h : \mathcal{X} \times \mathcal{A} \rightarrow \mathbb{R}$  in a backwards-inductive manner, i.e., for  $h = H, H - 1, \dots, 1$ . The values  $\hat{Q}_h(x, a)$  should be interpreted as estimates of the optimal value<sup>1</sup> at step  $h$  given that the state-action pair taken at step  $h$  is  $(x, a)$ . Since the environment’s reward functions and transitions are unknown,  $\hat{Q}_h$  must be estimated empirically. Given a dataset  $\mathcal{D}$  consisting of tuples  $(x_h, a_h, r_h, x_{h+1})$  of states, actions, and rewards drawn from the environment at step  $h$ , together with the subsequent state at step  $h + 1$ ,  $\hat{Q}_h$  is typically chosen to be the function in  $\mathcal{F}$  which minimizes the following square-loss objective whose labels are defined in terms of  $\hat{Q}_{h+1}$ :

$$\hat{Q}_h := \arg \min_{Q_h \in \mathcal{F}} \sum_{(x_h, a_h, r_h, x_{h+1}) \in \mathcal{D}} \left( Q_h(x_h, a_h) - \left( r_h + \max_{a'} \hat{Q}_{h+1}(x_{h+1}, a') \right) \right)^2. \quad (1)$$

The procedure defined by (1) is often known as *least-squares value iteration* (LSVI) (Bradtke and Barto, 1996; Osband et al., 2016). A natural condition under which LSVI might yield good estimates of the optimal value function, and thereby a near-optimal policy, is that the least-squares problem (1) be *well-specified*. This requires that for any  $\hat{Q}_{h+1} \in \mathcal{F}$ , there is some  $Q'_h \in \mathcal{F}$  so that, for all  $x_h \in \mathcal{X}, a_h \in \mathcal{A}$ ,

$$Q'_h(x_h, a_h) = \mathbb{E} \left[ r_h + \max_{a'} \hat{Q}_{h+1}(x_{h+1}, a') \mid (x_h, a_h) \right]. \quad (2)$$

(2) is known to be necessary for LSVI to succeed, in the sense that without it, the value functions computed by LSVI may be wildly divergent (Tsitsiklis and van Roy, 1996). Moreover, classical results (Munos, 2005; Munos and Szepesvári, 2008) showed that if the dataset  $\mathcal{D}$  is sufficiently exploratory, then (2) is sufficient for LSVI to succeed. The requirement that (2) holds for any  $\hat{Q}_{h+1} \in \mathcal{F}$  is often known as *Bellman completeness*; it is a property of the MDP and the value function class  $\mathcal{F}$ .

In this paper, our focus is on finding computationally efficient algorithms. Since regression problems such as (1) can be computationally intractable for even relatively simple nonlinear classes  $\mathcal{F}$  such as shallow neural networks (Dey et al., 2020; Bakshi et al., 2019; Goel et al., 2020), we focus on the case where  $\mathcal{F}$  is simply the class of linear functions in  $\phi_h(x, a)$ , for some known *feature mappings*  $\phi_h : \mathcal{X} \times \mathcal{A} \rightarrow \mathbb{R}^d$ . In this setting, the requirement of Bellman completeness in (2) is known as *linear Bellman completeness* (see Definition 2.1 for a formal definition). In the remainder of this section we discuss our results on computationally efficient learning of MDPs satisfying linear Bellman completeness.

### 1.1. Exploration and linear Bellman completeness

The results of (Munos, 2005; Munos and Szepesvári, 2008) referenced above showing sufficiency of Bellman completeness do not address the following central problem in RL: *how can we find a dataset  $\mathcal{D}$  which is sufficiently exploratory to compute a near-optimal policy?* More precisely, we would like that for tuples  $(x_h, a_h, r_h, x_{h+1}) \in \mathcal{D}$ , the feature vectors  $\phi_h(x_h, a_h)$  span a “sufficient

---

1. I.e., the value of the optimal policy.

number” of distinct directions in  $\mathbb{R}^d$ . We adopt the standard online setting in RL, allowing the learning algorithm to form  $\mathcal{D}$  by repeatedly sampling trajectories from the MDP using adaptively chosen policies.<sup>2</sup> We first note that this problem is known to be statistically tractable using a technique known as *global optimism* (see Section 3.1). This technique has various instantiations as specific algorithms in the setting of linear Bellman completeness, including ELEANOR (Zanette et al., 2020a), GOLF (Jin et al., 2021), OLIVE (Jiang et al., 2017), and BilinUCB (Du et al., 2021). While these algorithms require only polynomially many samples (i.e., rounds of interaction) to output a near-optimal policy, all are computationally inefficient, even when specialized to the linear case, since they require solving nonconvex optimization problems. Our goal is to find a computationally efficient algorithm which achieves the same guarantee:

**Question 1.1.** Is there an algorithm which learns an  $\epsilon$ -optimal policy in an unknown linear Bellman complete MDP using  $\text{poly}(H, d, |\mathcal{A}|, \epsilon^{-1})$  samples and time?

A sizeable portion of the work on computationally efficient RL in the last several years has been focused on answering Question 1.1 for settings which are strict special cases of linear Bellman completeness. The simplest such setting is the *tabular setting*, which describes the case that  $|\mathcal{X}|, |\mathcal{A}|$  are finite and the goal is to obtain sample and computational complexities scaling as  $\text{poly}(H, |\mathcal{X}|, |\mathcal{A}|)$ . In this setting, there are several computationally efficient algorithms which can be viewed as variants of value iteration that are *optimistic* in the sense that they add bonuses to the rewards to induce exploration: these include UCBVI (Azar et al., 2017) and Q-learning-UCB (Jin et al., 2018; Zhang et al., 2020), which are known to obtain near-optimal rates. Tabular MDPs are generalized by the *linear MDP* setting, in which feature vectors  $\phi_h(x, a) \in \mathbb{R}^d$  are given, and the state-action transition probabilities are assumed to be linear in  $\phi_h$ . Here too there are computationally efficient algorithms, namely LSVI-UCB (Jin et al., 2020), an optimistic version of LSVI, as well as more recent rate-optimal variants (Agarwal et al., 2022; He et al., 2023). Finally, in the setting of linear Bellman completeness, which is a strict generalization of linear MDPs (Zanette et al., 2020a, Proposition 3), the algorithm FRANCIS of (Zanette et al., 2020b) is computationally efficient and learns a near-optimal policy in the special case that the MDP is *reachable*, meaning that any direction in  $\mathbb{R}^d$  can be reached under some policy.

Despite the above line of work, Question 1.1 in its full generality has remained open. Part of the reason for this is that the two principal techniques to perform computationally efficient exploration in RL both break down in the general setting of linear Bellman completeness:

1. The first technique is *local optimism*,<sup>3</sup> which adds exploration bonuses to the reward at each state which scale inversely with how often the state has been visited. It includes the UCBVI and Q-learning-UCB algorithms for the tabular setting, and the LSVI-UCB algorithm for the linear MDP setting, among others. Local optimism requires that the value function class be complete with respect to the exploration bonuses, which is satisfied for linear MDPs but which fails more generally (see Section 3).
2. The second technique is to construct a *policy cover*, which is a small set of policies which, roughly speaking, covers all states to the maximum extent possible. This technique includes the FRANCIS algorithm, as well as computationally efficient and oracle efficient algorithms for related

2. See Section 2.1 for a formal definition.

3. “Local” is used to distinguish this technique from *global optimism*, which, as discussed above, works in the setting of linear Bellman completeness but is computationally inefficient.

tasks in RL such as representation learning (Du et al., 2019; Misra et al., 2020; Mhammedi et al., 2023b,a; Golowich et al., 2023) and learning in POMDPs (Golowich et al., 2022). In order for this approach to work in the absence of reachability, it is necessary to analyze a *truncated* version of the true MDP. But doing so seems impossible in our setting, since truncating the MDP breaks the property of linear Bellman completeness.

Thus, in addition to generalizing a long line of work on computationally efficient learning of MDPs, [Question 1.1](#) captures exactly the point where known exploration paradigms in RL break down.

### 1.1.1. MAIN RESULT FOR THE ONLINE SETTING

Our main result is a positive answer to [Question 1.1](#) in the case that  $|\mathcal{A}|$  is any constant:

**Theorem 1.1** (Informal version of [Theorem D.10](#)). *Suppose the ground-truth MDP satisfies linear Bellman completeness. Then for any  $\epsilon > 0$ , there is an algorithm (PSDP-UCB; [Algorithm 1](#)) which with high probability learns an  $\epsilon$ -optimal policy using  $(Hd|\mathcal{A}|\epsilon^{-1})^{O(|\mathcal{A}|)}$  samples and time.*

We remark that the exponential dependence of the sample and computational complexities on  $|\mathcal{A}|$  is somewhat unusual in RL. We are not aware of any prior work on an RL problem for which  $\text{poly}(|\mathcal{A}|)$  dependence is possible but for which a “natural” algorithm for even the  $|\mathcal{A}| = 2$  case does not extend to the general case. [Theorem 1.1](#) thus presents an intriguing challenge: to either improve the guarantee to fully answer [Question 1.1](#) in the affirmative, or to find a lower bound.

**Organization of the paper.** We discuss preliminaries in [Section 2](#), and overview the proof of [Theorem 1.1](#) in [Section 3](#), which is proved formally in [Appendices C](#) and [D](#).

## 2. Preliminaries

A finite-horizon *Markov Decision Process (MDP)* is given by a tuple  $M = (H, \mathcal{X}, \mathcal{A}, (P_h^M)_{h=1}^H, (r_h^M)_{h=1}^H, d_1^M)$ , where  $H \in \mathbb{N}$ ,  $\mathcal{X}$  is a measure space denoting the state set,  $\mathcal{A}$  denotes the action set,  $P_h^M(\cdot|x, a) \in \Delta(\mathcal{X})$  (for  $h \in [H]$ ) denotes the probability transition kernels,  $r_h^M : \mathcal{X} \times \mathcal{A} \rightarrow [0, 1]$  (for  $h \in [H]$ ) denotes the reward functions, and  $d_1^M \in \Delta(\mathcal{X})$  denotes the initial state distribution. When the MDP  $M$  is clear from context, we will drop the superscript  $M$  in these notations. We let  $A := |\mathcal{A}|$  denote the number of actions in the MDP.

A *policy*  $\pi$  consists of a tuple  $\pi = (\pi_1, \dots, \pi_H)$ , where each  $\pi_h : \mathcal{X} \rightarrow \Delta(\mathcal{A})$  is a mapping from states to distributions over actions. Note that we allow policies to be nonstationary and randomized;  $\Pi$  denotes the set of all such policies. A policy  $\pi \in \Pi$  defines a distribution over *trajectories*  $(x_1, a_1, r_1, \dots, x_H, a_H, r_H) \in (\mathcal{X} \times \mathcal{A} \times [0, 1])^H$ , as follows: first,  $x_1 \sim d_1$ , and then for each  $h \in [H]$ , we draw  $a_h \sim \pi_h(x_h)$ , observe  $r_h(x_h, a_h)$ , and transition to  $x_{h+1} \sim P_h(\cdot|x_h, a_h)$ . Let  $\mathcal{H}_H$  denote the set of trajectories. For a function  $f : \mathcal{H}_H \rightarrow \mathbb{R}^k$ , we often write  $\mathbb{E}^{M, \pi}[f(x_1, a_1, r_1, \dots, x_H, a_H, r_H)]$  to denote the expectation of  $f$  over trajectories drawn from  $M$  under policy  $\pi$ . If  $M$  is clear from context, we will simply write  $\mathbb{E}^\pi[f(x_1, a_1, r_1, \dots, x_H, a_H, r_H)]$ . Given policies  $\pi = (\pi_1, \dots, \pi_H), \pi' = (\pi'_1, \dots, \pi'_H) \in \Pi$  and a step  $h \in [H]$ , we let  $\pi \circ_h \pi' \in \Pi$  denote the policy  $(\pi_1, \dots, \pi_{h-1}, \pi'_h, \dots, \pi'_H)$ , i.e., which acts according to  $\pi$  during the first  $h - 1$  steps and thereafter acts according to  $\pi'$ .

For a function  $f : \mathcal{X} \times \mathcal{A} \rightarrow \mathbb{R}$ , a (randomized) policy  $\pi$  and  $h \in [H]$ , we write  $f(x, \pi_h(x)) := \mathbb{E}_{a \sim \pi_h(x)}[f(x, a)]$ . The state-action value function (or *Q-function*) and state-value function (or *V-*

function) of a policy  $\pi \in \Pi$  are then defined as follows: for  $h \in [H], x \in \mathcal{X}, a \in \mathcal{A}$ ,

$$Q_h^\pi(x, a) := r_h(x, a) + \mathbb{E}^\pi \left[ \sum_{g=h+1}^H r_g(x_g, a_g) \mid (x_h, a_h) = (x, a) \right], \quad V_h^\pi(x) := Q_h^\pi(x, \pi_h(x)).$$

The *optimal policy* is defined as  $\pi^* := \arg \max_{\pi \in \Pi} \mathbb{E}[V_1^\pi(x_1)]$  (where expectation is over  $x_1 \sim d_1$ ). It is known that there is always a deterministic optimal policy  $\pi^*$  (i.e., so that  $\pi_h^*(x)$  is a singleton for all  $x, h$ ). We often abbreviate  $Q_h^*(x, a) := Q_h^{\pi^*}(x, a)$  and  $V_h^*(x) := V_h^{\pi^*}(x)$ .

## 2.1. The online learning problem

We assume that the transitions, rewards, and initial state distribution of the ground-truth MDP  $M$  are *unknown* to the algorithm. To learn information about the MDP, the algorithm interacts with  $M$  via the *episodic online learning model*, as follows. The interaction proceeds over a series of  $T$  episodes. In each step  $h \in [H]$  of each episode, the algorithm observes the current state  $x_h$ , specifies an action  $a_h \in \mathcal{A}$  to take, and observes a reward of  $r_h = r_h(x_h, a_h)$ . Then the environment transitions to a new state  $x_{h+1} \sim P_h(\cdot | x_h, a_h)$ . We assume that the algorithm can query  $\phi_h(x, a) \in \mathbb{R}^d$  for any  $x, a, h$ . We remark that our algorithm only needs to query  $\phi_h(x, a)$  for states  $x$  which are visited at some point in some episode.

The goal is as follows: for  $\epsilon, \delta \in (0, 1)$ , to give an algorithm which interacts with the environment for  $T = T(\epsilon, \delta)$  episodes in the manner described above, and then to output a policy  $\hat{\pi}$  so that, with probability  $1 - \delta$ ,  $\mathbb{E}[V_1^*(x_1) - V_1^{\hat{\pi}}(x_1)] \leq \epsilon$ . Moreover, we wish the algorithm to be computationally efficient (both in terms of the computation required when interacting with the environment over the course of  $T$  episodes and when outputting  $\hat{\pi}$  at termination).

## 2.2. Function approximation

In order for the above learning problem to be tractable, it is necessary to make some assumptions on the ground-truth MDP  $M$  being learned. While it is well-known that boundedness of  $|\mathcal{X}|, |\mathcal{A}|$  implies efficient learning algorithms whose computational and statistical costs scale with  $\text{poly}(|\mathcal{X}|, |\mathcal{A}|)$  (Azar et al., 2017; Jin et al., 2018), realistic RL environments typically have enormous state spaces. Thus we aim for weaker function approximation assumptions: a common and longstanding such assumption (Bradtke and Barto, 1996; Melo and Ribeiro, 2007; Sutton and Barto, 2018; Yang and Wang, 2020; Jin et al., 2020) is *linearity* of the value functions, with respect to some known features. In particular, given  $d \in \mathbb{N}$ , we assume that functions  $\phi_h : \mathcal{X} \times \mathcal{A} \rightarrow \mathbb{R}^d$  are given for all  $h \in [H]$ , mapping each state-action pair  $(x, a)$  to a collection of  $d$  features which should be interpreted as capturing all relevant information about  $(x, a)$ .

The weakest assumption on value function linearity is simply that  $Q_h^*$  is linear, i.e., for some  $w_h^* \in \mathbb{R}^d$ , we have  $Q_h^*(x, a) = \langle w_h^*, \phi_h(x, a) \rangle$  for all  $(x, a) \in \mathcal{X} \times \mathcal{A}$ . Unfortunately, unless  $\text{NP} = \text{RP}$ , it is not possible to computationally efficiently learn a near-optimal policy under this assumption (Kane et al., 2022; Liu et al., 2023). Accordingly, we make a stronger assumption known as *linear Bellman completeness*, which states, roughly speaking, that the Bellman backup of all linear functions is linear. To formally state this assumption (in Definition 2.1 below), we need to introduce the following notation: for  $h \in [H]$ , define  $\mathcal{B}_h := \{\theta \in \mathbb{R}^d : |\langle \phi_h(x, a), \theta \rangle| \leq 1 \ \forall (x, a) \in \mathcal{X} \times \mathcal{A}\}$ . In words,  $\mathcal{B}_h$  denotes the set of coefficient vectors inducing bounded (linear) functions on  $\mathcal{X} \times \mathcal{A}$ .

**Definition 2.1** (Linear Bellman Completeness). The MDP  $M$  is defined to be *linear Bellman complete* with respect to the feature mappings  $(\phi_h)_{h \in [H]}$  if, for each  $h \in [H]$ , there is a mapping  $\mathcal{T}_h : \mathcal{B}_{h+1} \rightarrow \mathcal{B}_h$  so that, for all  $\theta \in \mathcal{B}_{h+1}$  and all  $(x, a) \in \mathcal{X} \times \mathcal{A}$ ,  $\langle \phi_h(x, a), \mathcal{T}_h \theta \rangle = \mathbb{E}_{x' \sim P_h(x, a)} [\max_{a' \in \mathcal{A}} \langle \phi_{h+1}(x', a'), \theta \rangle]$ . Moreover, we require that for all  $h \in [H]$ ,  $(x, a) \in \mathcal{X} \times \mathcal{A}$ , the reward at  $(x, a)$  at step  $h$  is given by:  $r_h(x, a) := \langle \phi_h(x, a), \theta_h^r \rangle$  for some vectors  $\theta_h^r \in \mathcal{B}_h$ .

It is immediate from [Definition 2.1](#) that if  $M$  is linear Bellman complete, then there are vectors  $w_h^* \in H \cdot \mathcal{B}_h$  so that  $Q_h^*(x, a) = \langle w_h^*, \phi_h(x, a) \rangle$  for all  $x, a$ . In addition to linear Bellman completeness, we make the following standard boundedness assumptions on the coefficient and feature vectors:

**Assumption 2.1** (Boundedness). *We assume the following:*

1. For all  $h \in [H]$ ,  $x \in \mathcal{X}$ ,  $a \in \mathcal{A}$ , we have  $\|\phi_h(x, a)\|_2 \leq 1$ .
2. For some parameter  $B \in \mathbb{R}_+$ : for all  $w_h \in \mathcal{B}_h$ , it holds that  $\|w_h\|_2 \leq B$ .
3. For all  $h \in [H]$ ,  $\|\theta_h^r\|_2 \leq 1$  (and hence  $\sup_{x, a, h} |r_h(x, a)| \leq 1$ ).

We present additional preliminaries in [Appendix A](#). Here we highlight the definition of a linear policy, which is defined as follows: for  $w \in \mathbb{R}^d$  and  $h \in [H]$ , the associated linear policy at step  $h$  is  $\pi_{h, w}(x) := \arg \max_{a \in \mathcal{A}} \langle w, \phi_h(x, a) \rangle$  if the argmax is unique. We discuss in [Appendix A](#) how to deal with the (typically measure-zero) situation that the argmax is not unique.

### 3. Technical overview

In this section, we overview the proof of [Theorem 1.1](#) (stated formally in [Theorem D.10](#)), which shows how, in the online setting ([Section 2.1](#)), [Algorithm 1](#) can efficiently learn a near-optimal policy for an unknown MDP which is linear Bellman complete. To simplify our notation, we assume in this section that the parameter  $B$  in [Assumption 2.1](#) is bounded by  $B \leq O(1)$ . The bulk of the challenge is to perform the task of *exploration*: how can we interact with the environment so as to reach state-action pairs  $(x, a)$  for which  $\phi_h(x, a)$  points in a new direction?

#### 3.1. Prior work: exploration via optimism

A popular approach to exploration in RL involves the use of *optimism*, which describes, loosely speaking, the technique of perturbing the algorithm’s estimates of the MDP’s value function, typically to *increase* the estimated values, so as to induce the algorithm to visit new directions in feature space. Two distinct flavors of optimism have emerged in the literature: the first, *global optimism*,<sup>4</sup> constructs a confidence set consisting of all possible vectors  $w = (w_1, \dots, w_H)$  which could be consistent with the optimal value function given the data observed so far. The global optimism procedure then chooses some  $\bar{w}$  in this confidence set which maximizes  $\mathbb{E}[\max_a \langle \bar{w}_1, \phi_1(x_1, a) \rangle]$ . It then executes the linear policy defined by  $\bar{w}$ , namely the policy  $(\pi_{1, \bar{w}_1}, \dots, \pi_{H, \bar{w}_H})$ , uses the resulting data to update the confidence sets, and repeats. Unfortunately, the optimization problem of finding a maximizing  $\bar{w}$  is nonconvex, and seems unlikely to have an efficient algorithm.

The second type of optimism-based exploration technique in RL is a more *local* approach: at each episode  $t$ , for some function  $B_h^{(t)} : \mathcal{X} \times \mathcal{A} \rightarrow \mathbb{R}_{\geq 0}$ , an *exploration bonus* of  $B_h^{(t)}(x, a)$  is

4. See, e.g., ([Zanette et al., 2020a](#)), which used global optimism to computationally inefficiently learn a near-optimal policy under linear Bellman completeness; many other papers, including ([Jin et al., 2021](#); [Jiang et al., 2017](#); [Du et al., 2021](#)), use global optimism in settings with more general nonlinear function approximation.



added to the reward  $r_h(x, a)$  for the pair  $(x, a)$  at step  $h$ . Then, the algorithm uses the data gathered prior to episode  $t$  to estimate a policy  $\pi^{(t)}$  which, roughly speaking, maximizes the expected sum of rewards  $r_h(x_h, a_h)$  and bonuses  $B_h^{(t)}(x_h, a_h)$  over a trajectory. This technique has proved successful for computationally efficient learning of tabular MDPs (i.e., the setting where  $|\mathcal{X}|, |\mathcal{A}|$  are finite) (Azar et al., 2017; Jin et al., 2018), as well as *linear MDPs* (Jin et al., 2020), which constitute a generalization of tabular MDPs and a strict subclass of linear Bellman complete MDPs.

Unfortunately, the local optimism approach fails in the more general setting of linear Bellman completeness. At a high level, this failure of local optimism results from the fact that the exploration bonuses  $B_h^{(t)}(x, a)$  may not be linear. To illustrate, we consider the LSVI-UCB approach of (Jin et al., 2020). This approach computes optimistic  $Q$ -function estimates,  $\hat{Q}_h^{(t)}$ , in a backwards-inductive manner. In particular, it defines:

$$\hat{Q}_h^{(t)}(x, a) := \min\{\langle \hat{w}_h^{(t)}, \phi_h(x, a) \rangle + B_h^{(t)}(x, a), H\}, \quad (3)$$

where  $\hat{w}_h^{(t)}$  is the solution to a least-squares objective function whose labels are given by  $\max_{a \in \mathcal{A}} \hat{Q}_{h+1}^{(t)}(x^i, a)$ , for various states  $x^i$ .<sup>5</sup> The bonus  $B_h^{(t)}(x, a)$  is defined as follows: let  $\Sigma_h^{(t)}$  be the covariance matrix of features at step  $h$  observed prior to episode  $t$ . Then  $B_h^{(t)}$  is given by a scaling of the *quadratic bonus*, i.e., for some scalar  $\beta_h$ ,  $B_h^{(t)}(x, a) := \beta_h \cdot B_h^{\text{quad}}(x, a; (\Sigma_h^{(t)})^{-1})$ , where the quadratic bonus,  $B_h^{\text{quad}}(x, a; \Sigma)$ , is defined for a general PSD matrix  $\Sigma$ , by

$$B_h^{\text{quad}}(x, a; \Sigma) := (\phi_h(x, a)^\top \cdot \Sigma \cdot \phi_h(x, a))^{1/2}. \quad (4)$$

The intuition behind the quadratic bonus is as follows:  $B_h^{\text{quad}}(x, a; (\Sigma_h^{(t)})^{-1})$  will be particularly large if  $\phi_h(x, a)$  points in the direction of eigenvectors of  $\Sigma_h^{(t)}$  with small eigenvalues, i.e., directions which have not been explored in prior episodes. Thus such *unexplored* states receive larger bonuses, and should be explored more during later episodes. More formally,  $B_h^{\text{quad}}(\cdot)$  takes the same form as the standard error bound from least-squares regression (see Lemma D.7), meaning that by adding it to the reward function, one can “cancel out” regression errors and thus show that  $\hat{Q}_h^{(t)}$  in (3) is optimistic (see (Jin et al., 2020; Agarwal et al., 2022; Zhang et al., 2022; He et al., 2023)).

To ensure that  $\hat{w}_h^{(t)}$ , used in the definition of  $\hat{Q}_h^{(t)}$ , is a good estimator of the future rewards and bonuses, we certainly need the regression problem to be *well-specified*, i.e., the expectation of the regression labels is a linear function in the features  $\phi_h(x, a)$ . In particular, this approach crucially relies on the fact that, in a linear MDP,  $\mathbb{E}_{x' \sim P_h(x, a)}[\max_{a'} \hat{Q}_{h+1}^{(t)}(x', a')]$ , is a linear function of  $\phi_h(x, a)$ . In fact, an even stronger statement holds for linear MDPs: for *any* function  $F : \mathcal{X} \rightarrow \mathbb{R}$ , its *Bellman backup*, namely  $\mathbb{E}_{x' \sim P_h(x, a)}[F(x')]$ , is linear in  $\phi_h(x, a)$ . Unfortunately, this fact fails to hold under the weaker assumption of linear Bellman completeness. We provide an example in Proposition F.1 for which  $\mathbb{E}_{x' \sim P_h(x, a)}[\max_{a'} \hat{Q}_{h+1}^{(t)}(x', a')]$  is *not linear*, for some  $\hat{Q}_{h+1}^{(t)}$  as in (3).

Our main innovation is to show that if the bonus  $B_h^{(t)}(x, a)$  is carefully defined to be somewhat different from a quadratic bonus, then a variant of (3) *does* have a linear Bellman backup. In the remainder of this section, we discuss in detail how to execute this strategy: in Section 3.2, we first discuss the overall structure of our algorithm, PSDP-UCB, which is a variant of the LSVI-UCB

5. The choice of the scalar  $H$  in the minimum in (3) results from the fact that the  $Q$ -function of any policy is bounded above in absolute value by  $H$ , as the reward at each step has absolute value at most 1.

algorithm discussed above but which lends itself to a simpler analysis for the setting of linear Bellman completeness. Then, in [Sections 3.3](#) and [3.4](#), we discuss how to construct exploration bonuses for use in PSDP-UCB which do have a linear Bellman backup.

### 3.2. Overview of PSDP-UCB

---

**Algorithm 1** PSDP-UCB( $T, H, \lambda, \beta, \lambda_1$ ): PSDP with upper confidence bounds

---

**Require:** Number of episodes  $T$ , horizon  $H$ , non-negative parameters  $\lambda, \beta, \lambda_1$ .

- 1: **for** Round  $t = 1, \dots, T$  **do**
- 2:     **for** Step  $h = H, \dots, 1$  **do**
- 3:         For  $i \in [n]$ , draw  $\{(x_k^{(t,i,h)}, a_k^{(t,i,h)}, r_k^{(t,i,h)})\}_{k=1}^H$  from  $\text{unif}(\{\hat{\pi}^{(s)} \circ_h \tilde{\pi}^{(s)} \circ_{h+1} \hat{\pi}^{(t)}\}_{s=1}^{t-1})$ .
- 4:         Set  $\Sigma_h^{(t)} \leftarrow \lambda I + \sum_{i=1}^n \phi_h(x_h^{(t,i,h)}, a_h^{(t,i,h)}) \phi_h(x_h^{(t,i,h)}, a_h^{(t,i,h)})^\top$ .
- 5:         For each  $g > h$  and  $i \in [n]$ , define  $\triangleright F_g^{(t)}$  defined in (33).

$$\hat{r}_g^{(t,i,h)} := r_g^{(t,i,h)} + F_g^{(t)}(x_g^{(t,i,h)}). \quad (5)$$

- 6:         Set  $\hat{w}_h^{(t)} \leftarrow (\Sigma_h^{(t)})^{-1} \cdot \sum_{i=1}^n \phi_h(x_h^{(t,i,h)}, a_h^{(t,i,h)}) \cdot (r_h^{(t,i,h)} + \sum_{g=h+1}^H \hat{r}_g^{(t,i,h)})$ .
  - 7:         Define  $\hat{\pi}_h^{(t)}$  by  $\hat{\pi}_h^{(t)}(x) := \arg \max_{a \in \mathcal{A}} \langle \phi_h(x, a), \hat{w}_h^{(t)} \rangle$ . (If  $t = 1$ , let  $\hat{\pi}^{(t)}$  be arbitrary.)
  - 8:         Define  $(\Sigma', \Lambda') := \text{trunc}_{\sigma_{\text{tr}}}(\frac{\beta}{\lambda_1} \cdot (\Sigma_h^{(t)})^{-1/2})$  (per [Definition D.2](#)).
  - 9:         Define  $\tilde{\pi}_h^{(t)} : \mathcal{X} \rightarrow \Delta(\mathcal{A})$  by, for all  $(x, a) \in \mathcal{X} \times \mathcal{A}$ ,  $\tilde{\pi}_h^{(t)}(x)(a) := \Pr_{w \sim \mathcal{N}(0, \Sigma')} (a = \arg \max_{a' \in \mathcal{A}} \langle w, \phi_h(x, a') \rangle)$ .
- 

Our algorithm, PSDP-UCB, is presented in [Algorithm 1](#). For some  $T \in \mathbb{N}$ , the algorithm proceeds over  $T$  rounds. In each round  $t \in [T]$ , PSDP-UCB constructs a policy  $\hat{\pi}^{(t)} = (\hat{\pi}_1^{(t)}, \dots, \hat{\pi}_H^{(t)})$ , which may be interpreted as an estimate of the optimal policy. We will show that, for  $T$  sufficiently large, with high probability there is some  $t \in [T]$  so that  $\mathbb{E}[V_1^*(x_1) - V_1^{\hat{\pi}^{(t)}}(x_1)]$  is small.

In each round  $t \in [T]$ , the algorithm iterates through steps  $h = H, H-1, \dots, 1$ : for each such  $h$ , it will define the mapping  $\hat{\pi}_h^{(t)} : \mathcal{X} \rightarrow \Delta(\mathcal{A})$ . Thus, at step  $h$ ,  $\hat{\pi}_{h+1}^{(t)}, \dots, \hat{\pi}_H^{(t)}$  have already been defined. At each step  $h$ , in [Line 3](#), the algorithm samples some number  $n$  of trajectories from the MDP according to a uniformly random policy from (roughly) the set  $\{\hat{\pi}^{(s)} \circ_{h+1} \hat{\pi}^{(t)}\}_{s \in [t-1]}$ .<sup>6</sup> Next, in [Line 4](#), the algorithm sets  $\Sigma_h^{(t)}$  to be the empirical covariance matrix of features at step  $h$  under the  $n$  trajectories just sampled.

PSDP-UCB then modifies the rewards in each of these  $n$  trajectories by adding to the true rewards an exploration bonus, denoted by  $F_g^{(t)}(\cdot)$  in [Line 5](#). We call these modified rewards, denoted by  $\hat{r}_h^{(t,i,h)}$  in [Algorithm 1](#), *optimistic rewards*. The remainder of this section will be focused on defining and explaining the intuition behind the bonus function  $F_g^{(t)}$  (the formal definition is given in (33)). Then, in [Line 6](#), PSDP-UCB sets  $\hat{w}_h^{(t)}$  to be the solution to a least-squares regression

---

6. Technically, this statement is slightly inaccurate: [Algorithm 1](#) actually samples a uniform policy from the set  $\{\hat{\pi}^{(s)} \circ_h \tilde{\pi}_h^{(s)} \circ_{h+1} \hat{\pi}^{(t)}\}_{s \in [t-1]}$ , where  $\tilde{\pi}_h^{(s)}$  is defined for each  $s$  to be a policy which performs a sort of “uniform” exploration at each step  $h$  ([Line 9](#)). This choice is made to simplify the analysis and is immaterial for our present discussion.



problem where the covariance matrix is  $\Sigma_h^{(t)}$  and the labels are the cumulative optimistic rewards from step  $h$  to  $H$ . The algorithm then repeats this procedure for step  $h - 1$ , and so on.

We have omitted a description of [Lines 8 and 9](#), in which a policy at step  $h$ ,  $\hat{\pi}_h^{(t)} : \mathcal{X} \rightarrow \Delta(\mathcal{A})$ , is defined. This policy plays a relatively minor role (see [Footnote 6](#)) and may be ignored for the purpose of the present discussion.

**Overview of the analysis of PSDP-UCB.** To analyze PSDP-UCB, we define

$$Q_h^{(t)}(x, a) := r_h(x, a) + \mathbb{E}^{\hat{\pi}^{(t)}} \left[ \sum_{g=h+1}^H r_g(x_g, a_g) + F_g^{(t)}(x_g) \mid (x_h, a_h) = (x, a) \right], \quad (6)$$

and  $V_h^{(t)}(x) := Q_h^{(t)}(x, \hat{\pi}_h^{(t)}(x))$ .  $Q_h^{(t)}$  represents the  $Q$ -value function corresponding to the policy  $\hat{\pi}^{(t)}$  and the optimistic rewards discussed above. Our main technical lemma is the following statement showing that adding the exploration bonuses  $F_h^{(t)}(\cdot)$  to the rewards ensures that  $Q_h^{(t)}$  is an upper bound on  $Q_h^*$ :

**Lemma 3.1** (Informal version of [Lemma D.9](#)). *With high probability over the execution of PSDP-UCB ([Algorithm 1](#)), for all  $x, a, h, t$ , we have  $Q_h^{(t)}(x, a) \geq Q_h^*(x, a) - o(1)$  and  $V_h^{(t)}(x) + F_h^{(t)}(x) \geq V_h^*(x) - o(1)$ .*

In light of [Lemma 3.1](#), we can bound the average suboptimality of the policies  $\hat{\pi}^{(t)}$  as follows:

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^T \left( \mathbb{E}[V_1^*(x_1) - V_1^{\hat{\pi}^{(t)}}(x_1)] \right) &\leq \frac{1}{T} \sum_{t=1}^T \left( \mathbb{E}[V_1^{(t)}(x_1) + F_1^{(t)}(x_1) - V_1^{\hat{\pi}^{(t)}}(x_1)] \right) - o(1) \\ &= \frac{1}{T} \sum_{t=1}^T \mathbb{E}^{\hat{\pi}^{(t)}} \left[ \sum_{h=1}^H F_h^{(t)}(x_h) \right] - o(1), \end{aligned} \quad (7)$$

where the inequality uses [Lemma 3.1](#), and the equality uses the definition of  $V_h^{(t)}$ . To upper bound [\(7\)](#), we need to show, roughly speaking, the following: for states  $x$  which are likely to be visited by many of the policies  $\hat{\pi}^{(t)}$ , the sum of exploration bonuses at  $x$  is sublinear in  $T$ , i.e.,  $\frac{1}{T} \sum_{t=1}^T F_h^{(t)}(x) = o(T)$ . We show a version of this statement “in feature space”: in particular, we first upper bound  $|F_h^{(t)}(x)|$ , for all states  $x$ , as follows.

**Lemma 3.2** (Informal version of [Lemmas D.2, D.3 and D.5](#)). *For all  $x, h, t$  in the execution of PSDP-UCB, it holds that  $|F_h^{(t)}(x)| \leq \text{poly}(d, H)^A \cdot \max_{a \in \mathcal{A}} \left( \phi_h(x, a)^\top \cdot (\Sigma_h^{(t)})^{-1} \cdot \phi_h(x, a) \right)^{1/2}$ .*

[Lemma 3.2](#) shows that for states  $x$  for which all actions  $a \in \mathcal{A}$  point in directions of  $\Sigma_h^{(t)}$  corresponding to eigenvectors with large eigenvalues (i.e., “well-explored” directions),  $|F_h^{(t)}(x)|$  must be small. Using [Lemma 3.2](#), the quantity in [\(7\)](#) may be bounded by  $o(1)$  using a variant of the elliptic potential lemma ([Lemma E.4](#)). The elliptic potential lemma formalizes the notion that there are finitely many “different directions” in  $\mathbb{R}^d$ , so if  $T$  is sufficiently large, all directions  $\mathbb{R}^d$  frequently explored by policies  $\hat{\pi}^{(t)}$  must correspond to large-eigenvalue directions for  $\Sigma_h^{(t)}$ .

Thus, it remains to discuss the proofs of [Lemmas 3.1 and 3.2](#). Notice that the statements of these two lemmas are in tension: [Lemma 3.1](#) requires the bonuses  $F_h^{(t)}$  be sufficiently large so that  $Q_h^{(t)}$

is an upper bound on  $Q_h^*$ , while [Lemma 3.2](#) requires that  $F_h^{(t)}$  be sufficiently small, as controlled by  $\Sigma_h^{(t)}$ . As discussed in the following subsections, by carefully defining  $F_h^{(t)}$ , we can ensure that both of these criteria are met.

### 3.3. Implementing local optimism via Bellman-linearity

Essentially any approach to proving [Lemma 3.1](#) requires that  $Q_h^{(t)}(x, a)$  be a linear function of  $\phi_h(x, a)$ .<sup>7</sup> In turn, as discussed in [Section 3.1](#), this requires that the bonuses  $F_{h+1}^{(t)}(\cdot)$  have a Bellman backup which is linear in  $\phi_h(x, a)$ . It will be useful to formalize the following notion of *Bellman-linearity*, which describes those functions whose Bellman backup is linear in the features:

**Definition 3.1.** Fix  $h \in [H]$  with  $h > 1$  and  $k \in \mathbb{N}$ . We say that a function  $G : \mathcal{X} \rightarrow \mathbb{R}^k$  is *Bellman-linear at step  $h$*  if there exists  $W_G \in \mathbb{R}^{k \times d}$  so that, for all  $(x, a) \in \mathcal{X} \times \mathcal{A}$ , it holds that  $W_G \cdot \phi_{h-1}(x, a) = \mathbb{E}_{x' \sim P_{h-1}(x, a)}[G(x')]$ . In particular, in the case that  $k = 1$ ,  $G : \mathcal{X} \rightarrow \mathbb{R}$  is Bellman-linear at step  $h$  if there is  $w_G \in \mathbb{R}^d$  so that  $\langle w_G, \phi_{h-1}(x, a) \rangle = \mathbb{E}_{x' \sim P_{h-1}(x, a)}[G(x')]$ .

If the value of  $h$  is clear, we will simply say the function  $G$  is Bellman-linear. The canonical example of a Bellman-linear function is the mapping  $x \mapsto \max_a \langle w, \phi_h(x, a) \rangle$ , for any  $w \in \mathbb{R}^d$ . Given that such functions are Bellman-linear, we can construct additional Bellman-linear functions by averaging, taking limits ([Lemma C.1](#)), and differentiating ([Lemma C.2](#)). The following lemma shows that Bellman-linearity of the bonuses  $F_h^{(t)}$  is indeed a sufficient condition for linearity of  $Q_h^{(t)}$ :

**Lemma 3.3** (Variant of [Lemma D.1](#)). *Suppose that, in PSDP-UCB, for each  $h \in [H], t \in [T]$ ,  $F_h^{(t)} : \mathcal{X} \rightarrow \mathbb{R}$  is Bellman-linear at step  $h$ . Then for all  $h, t$ , there is some  $w_h^{(t)} \in \mathbb{R}^d$  so that  $Q_h^{(t)}(x, a) = \langle \phi_h(x, a), w_h^{(t)} \rangle$  for all  $(x, a) \in \mathcal{X} \times \mathcal{A}$ .*

The proof of [Lemma 3.3](#) requires some work to show that the expectation of the bonuses  $F_g^{(t)}$  for  $g > h + 1$  is linear at step  $h$ : this fact crucially relies on the fact that the policies  $\hat{\pi}^{(t)}$  are linear policies (see [Lemma C.3](#)).

**A Bellman-linear approximation of  $B_h^{\text{quad}}$ .** In light of our discussion pertaining to the quadratic bonus  $B_h^{\text{quad}}$  around (4) above, a natural strategy is to show that  $\max_a B_h^{\text{quad}}(\cdot, a; \Sigma)$  is Bellman-linear at step  $h$  for any  $\Sigma$ . Though this is not the case ([Proposition F.2](#)), [Lemma 3.4](#) below suggests a workaround: the quadratic bonus  $B_h^{\text{quad}}(x, a; \Sigma)$  can be *approximated* by a Bellman-linear function.

**Lemma 3.4** (Simplified version of [Lemma D.3](#)). *Fix  $h \in [H]$  and a PSD matrix  $\Sigma$ , and define*

$$F_h^{\text{normal}}(x; \Sigma) := \mathbb{E}_{w \sim \mathcal{N}(0, \Sigma)} \left[ \max_{a \in \mathcal{A}} \langle w, \phi_h(x, a) \rangle \right]. \quad (8)$$

*Then  $F_h^{\text{normal}}(\cdot; \Sigma)$  is Bellman-linear at step  $h$ , and for any  $x \in \mathcal{X}$ ,*

$$\frac{1}{\sqrt{2\pi}} \max_{a, a' \in \mathcal{A}} ((\phi_h(x, a) - \phi_h(x, a'))^\top \Sigma (\phi_h(x, a) - \phi_h(x, a')))^{1/2} \leq F_h^{\text{normal}}(x; \Sigma) \leq \sqrt{d} \max_{a \in \mathcal{A}} (\phi_h(x, a)^\top \Sigma \phi_h(x, a))^{1/2}.$$

7. In particular, linearity of  $Q_h^{(t)}$  establishes that the regression problem solved in [Line 6](#) of [Algorithm 1](#) is well-specified, which allows us to show upper bounds on the prediction error  $|\langle \hat{w}_h^{(t)}, \phi_h(x, a) \rangle - Q_h^{(t)}(x, a)|$  and is in turn used to establish optimism.

The proof of [Lemma 3.4](#) is straightforward: that  $x \mapsto F_h^{\text{normal}}(x; \Sigma)$  is Bellman-linear at step  $h$  follows from the fact that  $F_h^{\text{normal}}(x; \Sigma)$  is defined as an average of Bellman-linear functions. The upper and lower bounds on  $F_h^{\text{normal}}(x; \Sigma)$  follow from a direct computation.

[Lemma 3.4](#) suggests the following approach: for some scalar  $\beta_h$ , suppose we define

$$F_h^{(t)}(x) := \beta_h \cdot F_h^{\text{normal}}(x; (\Sigma_h^{(t)})^{-1}). \quad (9)$$

Since  $F_h^{\text{normal}}(\cdot; \Sigma)$  is Bellman-linear and approximates the quadratic bonus ([Lemma 3.4](#)), we might hope that it acts similarly to the quadratic bonus in `LSVI-UCB` to induce an optimistic property of the value functions  $Q_h^{(t)}(x, a)$  in `PSDP-UCB`, thereby allowing us to establish [Lemma 3.1](#). Unfortunately, this argument encounters a few snags:

1. The lower bound on  $F_h^{\text{normal}}(x; \Sigma)$  in [Lemma 3.4](#) is not exactly a constant times  $B_h^{\text{quad}}(x; \Sigma)$ , but instead involves a *difference* between feature vectors at the state  $x$ . Thus, strictly speaking,  $F_h^{\text{normal}}(x; \Sigma)$  does not give a multiplicative approximation of the quadratic bonus  $B_h^{\text{quad}}(x; \Sigma)$ .
2. Even if  $F_h^{\text{normal}}(x; \Sigma)$  did give a multiplicative approximation of  $B_h^{\text{quad}}(x; \Sigma)$ , we would need to ensure that this multiplicative approximation is good enough for all arguments to go through.
3. Ignoring the previous two items, we still suffer from the following exponential blowup. Roughly speaking, the proof of [Lemma 3.1](#) proceeds by showing that the loss in value from any sub-optimal action taken by  $\hat{\pi}^{(t)}$  at step  $h$  is canceled out by the bonus  $F_h^{(t)}(x_h)$ . To show this, we need to show that the exploration bonus  $\beta_h \cdot F_h^{\text{normal}}(x_h; (\Sigma_h^{(t)})^{-1})$  is at least as large as the regression error in our estimate of  $\hat{w}_h^{(t)}$  from [Line 6](#). But the labels in this linear regression involve the sum of bonuses  $F_g^{(t)}(\cdot)$  for  $h+1 \leq g \leq H$ , and thus scale with  $\beta_{h+1} + \dots + \beta_H$ . Working through standard high-confidence bounds for linear regression (see [Lemma D.7](#) for details) establishes that we would need  $\beta_h \geq \text{poly}(d) \cdot (\beta_{h+1} + \dots + \beta_H)$ . This recursion leads to  $\beta_1 \geq d^{O(H)}$ , which would require exponentially many samples to cancel out exponentially large error terms.

[Items 1](#) and [2](#) above turn out not to be major issues, as discussed in the proof of [Lemma D.9](#).

The main technical challenge is dealing with [Item 3](#) above, namely the exponential blowup of  $\beta_h$ . Typically, the solution to the above dilemma is to *truncate* the bonus, i.e., define  $F_h^{(t)}(x) := \min\{\beta \cdot F_h^{\text{normal}}(x; (\Sigma_h^{(t)})^{-1}), \bar{F}_h(x)\}$ , for some  $\bar{F}_h(x) \ll \beta$  which is a uniform upper bound on  $V_h^*(x)$ . (Note that, when using such an approach, the parameter  $\beta$  does not need any dependence on  $h$ .) Indeed, this truncation is used in `LSVI-UCB`, where the  $\hat{Q}_h^{(t)}$  function in [\(3\)](#) is truncated by the constant  $\bar{F}_h(x) \equiv H$ . Unfortunately, in the setting of linear Bellman completeness, truncation, in general, breaks Bellman-linearity of  $F_h^{(t)}(x)$ . In the following section, we discuss how to implement this truncation in a way that preserves Bellman-linearity.

### 3.4. Constructing truncated bonuses

**Approximating  $(\Sigma_h^{(t)})^{-1}$ .** For each  $h \in [H]$  and  $t \in [T]$ , the inverse covariance matrix  $(\Sigma_h^{(t)})^{-1}$  captures which directions of feature space are explored by policies  $\hat{\pi}^{(s)} \circ_h \tilde{\pi}^{(s)}$ , for  $s \leq t-1$ . We split  $(\Sigma_h^{(t)})^{-1}$  into matrices representing “well-explored directions” and “unexplored directions”, as follows: for an appropriate threshold value  $\sigma$ , we let  $\Sigma'$  denote the projection matrix onto the subspace spanned by eigenvectors of  $(\Sigma_h^{(t)})^{-1/2}$  with eigenvalues greater than  $\sigma$ , and  $\Lambda'$  denote the

projection matrix onto the subspace spanned by eigenvectors of  $(\Sigma_h^{(t)})^{-1/2}$  with eigenvalues less than  $\sigma$ . Formally, as defined in [Definition D.2](#), we write  $(\Sigma', \Lambda') = \text{trunc}_\sigma((\Sigma_h^{(t)})^{-1/2})$ . We denote subspaces onto which  $\Sigma', \Lambda'$  project by  $\mathcal{S}_{\Sigma'}, \mathcal{S}_{\Lambda'}$ , respectively. Note that  $\Sigma' \Lambda' = \Lambda' \Sigma' = 0$  and  $\Sigma' + \Lambda' = I_d$ ; we call the pair  $(\Sigma', \Lambda')$  an *orthogonal pair* ([Definition C.1](#)).

If  $\sigma$  is chosen sufficiently small to be less than our desired error threshold, then the estimations  $\hat{w}_h^{(t)}$  produced in [Line 6](#) of [Algorithm 1](#) are sufficiently accurate in directions spanned by  $\mathcal{S}_{\Lambda'}$ . Thus, we only have to add an exploration bonus at a state  $x$  which scales in proportion to the width of the features  $\phi_h(x, a)$ ,  $a \in \mathcal{A}$  in directions spanned by the orthogonal complement of  $\mathcal{S}_{\Lambda'}$ , namely  $\mathcal{S}_{\Sigma'}$ . This leads us to aim to construct an exploration bonus which scales roughly as  $x \mapsto \min\{\beta \cdot F_h^{\text{normal}}(x; \Sigma'), \bar{F}_h(x)\}$ , for some parameter  $\beta > 1$  and an appropriate choice of  $\bar{F}_h$  which should (a) be uniformly upper bounded by a parameter which does not scale with  $\beta$ , and (b) be an upper bound on  $V_h^*$ . It turns out that condition (b) can be relaxed to require only that  $\bar{F}_h(x)$  be an upper bound on  $\max_{a \in \mathcal{A}} Q_h^*(x, a) - \min_{a \in \mathcal{A}} Q_h^*(x, a) = \max_{a, a' \in \mathcal{A}} \langle w_h^*, \phi_h(x, a) - \phi_h(x, a') \rangle$ . Thus, to summarize, we aim to find a bonus function which satisfies the following properties:

**Problem 3.2** (Bonus function). Fix  $h, t$  (which determine  $(\Sigma', \Lambda') = \text{trunc}_\sigma((\Sigma_h^{(t)})^{-1/2})$  as above), and some  $\beta > 1$ . Can we find a Bellman-linear function  $F_h^{(t)} : \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$  which satisfies:

1. For all  $x \in \mathcal{X}$ ,  $F_h^{(t)}(x)$  is lower bounded as

$$F_h^{(t)}(x) \geq BH \cdot \max_{a, a' \in \mathcal{A}} \min \{ \beta \cdot \|\Sigma' \cdot (\phi_h(x, a) - \phi_h(x, a'))\|_2, \|\Lambda' \cdot (\phi_h(x, a) - \phi_h(x, a'))\|_2 \}. \quad (10)$$

2. For all  $x \in \mathcal{X}$ ,  $F_h^{(t)}(x)$  is upper bounded as  $F_h^{(t)}(x) \leq \text{poly}(d, H, B) \cdot \max_{a \in \mathcal{A}} \min \{ \beta \|\Sigma' \cdot \phi_h(x, a)\|_2, \|\phi_h(x, a)\|_2 \}$ .

**The case  $d = A = 2$ .** To proceed, we consider the special case in which  $d = 2$ . Let us suppose that  $\mathcal{S}_{\Sigma'}, \mathcal{S}_{\Lambda'}$  are each one dimensional, and are spanned by unit vectors  $u, v \in \mathbb{R}^2$ , respectively, which are orthogonal. Let us consider the function

$$F_h^{\text{tl}}(x; u, v) := \max_{a \in \mathcal{A}} \langle u, \phi_h(x, a) \rangle + \max_{a \in \mathcal{A}} \langle v, \phi_h(x, a) \rangle - \max_{a \in \mathcal{A}} \langle u + v, \phi_h(x, a) \rangle,$$

for some parameter  $\beta > 1$ . If we suppose that there are only 2 actions (i.e.,  $|\mathcal{A}| = 2$ ), then  $x \mapsto \sum_{\epsilon_u, \epsilon_v \in \{\pm 1\}} F_h^{\text{tl}}(x; \beta \epsilon_u \cdot u, \epsilon_v \cdot v)$  (scaled by a  $\text{poly}(B, H)$  factor) satisfies the conditions of [Problem 3.2](#). To see this, let us write  $\mathcal{A} = \{0, 1\}$ , and set  $\phi_h^\Delta(x) := \phi_h(x, 1) - \phi_h(x, 0)$ . Fix  $x \in \mathcal{X}$ , and choose  $\epsilon_u, \epsilon_v \in \{\pm 1\}$  so that  $\langle \epsilon_u \cdot u, \phi_h^\Delta(x) \rangle$  and  $\langle \epsilon_v \cdot v, \phi_h^\Delta(x) \rangle$  have opposite signs. Then

$$\begin{aligned} F_h^{\text{tl}}(x; \beta \epsilon_u u, \epsilon_v v) &= |\langle \beta \epsilon_u u, \phi_h^\Delta(x) \rangle| + |\langle \epsilon_v v, \phi_h^\Delta(x) \rangle| - |\langle \beta \epsilon_u u + \epsilon_v v, \phi_h^\Delta(x) \rangle| \\ &\geq \min\{|\langle \beta u, \phi_h^\Delta(x) \rangle|, |\langle v, \phi_h^\Delta(x) \rangle|\}, \end{aligned} \quad (11)$$

where the inequality uses the property that for real numbers  $a, b$  with opposite signs, we have  $|a| + |b| - |a + b| \geq \min\{|a|, |b|\}$ . By combining [\(11\)](#) with non-negativity of  $F_h^{\text{tl}}(\cdot)$ , we obtain the first property of [Problem 3.2](#). The second property may be verified in a similar manner (see [Lemma C.8](#)).

**The case of general  $d$  and  $A$ .** A direct extension of the above argument for  $d = A = 2$  runs into some snags even in the case that  $d = 2$  and  $A = 4$ . We discuss these issues and additional techniques we introduce to overcome them in [Appendix B](#).

## Acknowledgements

NG is supported by a Fannie & John Hertz Foundation Fellowship and an NSF Graduate Fellowship. AM is supported in part by a Microsoft Trustworthy AI Grant, an ONR grant and a David and Lucile Packard Fellowship.

## References

- Alekh Agarwal, Yujia Jin, and Tong Zhang. Voql: Towards optimal regret in model-free rl with nonlinear function approximation. 2022.
- Mohammad Gheshlaghi Azar, Ian Osband, and Rémi Munos. Minimax regret bounds for reinforcement learning. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70*, ICML'17, page 263–272. JMLR.org, 2017.
- Ainesh Bakshi, Rajesh Jayaram, and David P Woodruff. Learning two layer rectified neural networks in polynomial time. In Alina Beygelzimer and Daniel Hsu, editors, *Proceedings of the Thirty-Second Conference on Learning Theory*, volume 99 of *Proceedings of Machine Learning Research*, pages 195–268. PMLR, 25–28 Jun 2019.
- Steven J. Bradtke and Andrew G. Barto. Linear least-squares algorithms for temporal difference learning. *Mach. Learn.*, 22(1–3):33–57, jan 1996. ISSN 0885-6125.
- Santanu Dey, Guanyi Wang, and Yao Xie. Approximation algorithms for training one-node relu neural networks. *IEEE Transactions on Signal Processing*, PP:1–1, 11 2020.
- Simon Du, Akshay Krishnamurthy, Nan Jiang, Alekh Agarwal, Miroslav Dudik, and John Langford. Provably efficient rl with rich observations via latent state decoding. In *International Conference on Machine Learning*, pages 1665–1674. PMLR, 2019.
- Simon Du, Sham Kakade, Jason Lee, Shachar Lovett, Gaurav Mahajan, Wen Sun, and Ruosong Wang. Bilinear classes: A structural framework for provable generalization in rl. In *International Conference on Machine Learning*, pages 2826–2836. PMLR, 2021.
- Surbhi Goel, Adam R. Klivans, Pasin Manurangsi, and Daniel Reichman. Tight hardness results for training depth-2 relu networks. In *Information Technology Convergence and Services*, 2020.
- Noah Golowich, Ankur Moitra, and Dhruv Rohatgi. Learning in observable pomdps, without computationally intractable oracles. *Advances in Neural Information Processing Systems*, 35:1458–1473, 2022.
- Noah Golowich, Ankur Moitra, and Dhruv Rohatgi. Exploring and learning in sparse linear mdps without computationally intractable oracles. 2023.
- Shixiang Gu, Ethan Holly, Timothy Lillicrap, and Sergey Levine. Deep reinforcement learning for robotic manipulation with asynchronous off-policy updates. In *2017 IEEE international conference on robotics and automation (ICRA)*, pages 3389–3396. IEEE, 2017.
- Hado van Hasselt, Arthur Guez, and David Silver. Deep reinforcement learning with double q-learning. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, AAAI'16, page 2094–2100. AAAI Press, 2016.

- Jiafan He, Heyang Zhao, Dongruo Zhou, and Quanquan Gu. Nearly minimax optimal reinforcement learning for linear Markov decision processes. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett, editors, *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 12790–12822. PMLR, 23–29 Jul 2023.
- Nan Jiang, Akshay Krishnamurthy, Alekh Agarwal, John Langford, and Robert E Schapire. Contextual decision processes with low bellman rank are pac-learnable. In *International Conference on Machine Learning*, pages 1704–1713. PMLR, 2017.
- Chi Jin, Zeyuan Allen-Zhu, Sebastien Bubeck, and Michael I Jordan. Is q-learning provably efficient? *Advances in neural information processing systems*, 31, 2018.
- Chi Jin, Zhuoran Yang, Zhaoran Wang, and Michael I Jordan. Provably efficient reinforcement learning with linear function approximation. In *Conference on Learning Theory*, pages 2137–2143. PMLR, 2020.
- Chi Jin, Qinghua Liu, and Sobhan Miryoosefi. Bellman eluder dimension: New rich classes of rl problems, and sample-efficient algorithms. *Advances in neural information processing systems*, 34:13406–13418, 2021.
- Sham Kakade and John Langford. Approximately optimal approximate reinforcement learning. In *Proceedings of the Nineteenth International Conference on Machine Learning*, pages 267–274, 2002.
- Daniel Kane, Sihan Liu, Shachar Lovett, and Gaurav Mahajan. Computational-statistical gap in reinforcement learning. In *Conference on Learning Theory*, pages 1282–1302. PMLR, 2022.
- Sihan Liu, Gaurav Mahajan, Daniel Kane, Shachar Lovett, Gellért Weisz, and Csaba Szepesvári. Exponential hardness of reinforcement learning with linear function approximation. In *The Thirty Sixth Annual Conference on Learning Theory*, pages 1588–1617. PMLR, 2023.
- Francisco S. Melo and M. Isabel Ribeiro. Convergence of q-learning with linear function approximation. In *2007 European Control Conference (ECC)*, pages 2671–2678, 2007.
- Zakaria Mhammedi, Adam Block, Dylan J Foster, and Alexander Rakhlin. Efficient model-free exploration in low-rank mdps. *arXiv preprint arXiv:2307.03997*, 2023a.
- Zakaria Mhammedi, Dylan J Foster, and Alexander Rakhlin. Representation learning with multi-step inverse kinematics: An efficient and optimal approach to rich-observation rl. *arXiv preprint arXiv:2304.05889*, 2023b.
- Dipendra Misra, Mikael Henaff, Akshay Krishnamurthy, and John Langford. Kinematic state abstraction and provably efficient rich-observation reinforcement learning. In *International conference on machine learning*, pages 6961–6971. PMLR, 2020.
- Rémi Munos. Error bounds for approximate value iteration. In *Proceedings of the 20th National Conference on Artificial Intelligence - Volume 2, AAAI’05*, page 1006–1011. AAAI Press, 2005. ISBN 157735236x.



- Rémi Munos and Csaba Szepesvári. Finite-time bounds for fitted value iteration. *Journal of Machine Learning Research*, 9(27):815–857, 2008.
- Ian Osband, Benjamin Van Roy, and Zheng Wen. Generalization and exploration via randomized value functions. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48, ICML'16*, page 2377–2386. JMLR.org, 2016.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *CoRR*, abs/1707.06347, 2017.
- Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. The MIT Press, second edition, 2018.
- John N. Tsitsiklis and Benjamin van Roy. Feature-based methods for large scale dynamic programming. *Mach. Learn.*, 22(1–3):59–94, jan 1996. ISSN 0885-6125.
- Lin Yang and Mengdi Wang. Reinforcement learning in feature space: Matrix bandit, kernels, and regret bound. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 10746–10756. PMLR, 13–18 Jul 2020.
- Chao Yu, Jiming Liu, Shamim Nemati, and Guosheng Yin. Reinforcement learning in healthcare: A survey. *ACM Computing Surveys (CSUR)*, 55(1):1–36, 2021.
- Andrea Zanette, Alessandro Lazaric, Mykel Kochenderfer, and Emma Brunskill. Learning near optimal policies with low inherent Bellman error. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 10978–10989. PMLR, 13–18 Jul 2020a.
- Andrea Zanette, Alessandro Lazaric, Mykel J Kochenderfer, and Emma Brunskill. Provably efficient reward-agnostic navigation with linear value iteration. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 11756–11766. Curran Associates, Inc., 2020b.
- Andrea Zanette, Ching-An Cheng, and Alekh Agarwal. Cautiously optimistic policy optimization and exploration with linear function approximation. In Mikhail Belkin and Samory Kpotufe, editors, *Proceedings of Thirty Fourth Conference on Learning Theory*, volume 134 of *Proceedings of Machine Learning Research*, pages 4473–4525. PMLR, 15–19 Aug 2021.
- Xuezhou Zhang, Yuda Song, Masatoshi Uehara, Mengdi Wang, Alekh Agarwal, and Wen Sun. Efficient reinforcement learning in block mdps: A model-free representation learning approach. In *International Conference on Machine Learning*, pages 26517–26547. PMLR, 2022.
- Zihan Zhang, Yuan Zhou, and Xiangyang Ji. Almost optimal model-free reinforcement learning via reference-advantage decomposition. In *Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS'20*, Red Hook, NY, USA, 2020. Curran Associates Inc. ISBN 9781713829546.
- Stephan Zheng, Alexander Trott, Sunil Srinivasa, David C. Parkes, and Richard Socher. The ai economist: Taxation policy design via two-level deep multiagent reinforcement learning. *Science Advances*, 8(18):eabk2607, 2022.

## Appendix A. Additional preliminaries

In this section, we present several additional preliminaries that will be useful in our proofs.

**Polytope of actions.** We assume that  $\mathcal{A}$  is finite and write  $A := |\mathcal{A}|$ . Let  $\mathcal{P}^d$  denote the space of  $d$ -dimensional (bounded) polyhedra. For  $A \in \mathbb{N}$ , let  $\mathcal{P}_A^d \subset \mathcal{P}^d$  denote the space of  $d$ -dimensional polyhedra  $\Phi$  so that  $\dim(\text{Span}\Phi) \leq A$ . For each  $x \in \mathcal{X}$ , we define  $\Phi_h(x) := \{\phi_h(x, a) : a \in \mathcal{A}\} \subset \mathbb{R}^d$  and  $\bar{\Phi}_h(x) := \text{co}(\Phi_h(x)) \in \mathcal{P}_A^d \subset \mathcal{P}^d$ .

**Gaussian smoothing.** For  $\theta \in \mathbb{R}^d$  and  $\sigma > 0$ , we write

$$\mathcal{N}_\sigma(\theta) := \mathcal{N}(0, \sigma^2 \cdot I_d)(\theta) = \frac{1}{(2\pi)^{d/2} \sigma^d} \cdot \exp\left(-\frac{1}{2\sigma^2} \|\theta\|_2^2\right)$$

to denote the probability density function of the standard normal distribution with covariance  $\sigma^2 I_d$ . Furthermore, for  $f : \mathbb{R}^d \rightarrow \mathbb{R}$ , we write  $S_\sigma f(\theta)$  to denote the convolution of  $f$  with  $\mathcal{N}_\sigma$ , namely

$$S_\sigma f(\theta) := \int_{\mathbb{R}^d} f(z) \mathcal{N}_\sigma(\theta - z) dz = \int_{\mathbb{R}^d} f(\theta - z) \mathcal{N}_\sigma(z) dz = \mathbb{E}_{z \sim \mathcal{N}(0, \sigma^2 I_d)}[f(\theta - z)].$$

**Miscellaneous notation.** Given  $d \in \mathbb{N}$ , we let  $\mathbb{S}^d$  denote the space of symmetric  $d \times d$  matrices,  $\mathbb{S}_+^d \subset \mathbb{S}^d$  denote the space of positive semidefinite (PSD) matrices, and  $\mathbb{S}_{++}^d \subset \mathbb{S}_+^d$  denote the space of positive definite matrices. For a PSD matrix  $T \in \mathbb{S}_+^d$ , we let  $T^{1/2}$  denote the unique PSD matrix whose square is  $T$ . Given a subset  $\mathcal{S} \subset \mathbb{R}^d$  and a matrix  $T \in \mathbb{R}^{d \times d}$ , we let  $T \cdot \mathcal{S} := \{Tv : v \in \mathcal{S}\}$ . For vectors  $v, v' \in \mathbb{R}^d$ , we let  $[v, v']$  denote the segment  $[v, v'] := \{\lambda v + (1 - \lambda)v' : \lambda \in [0, 1]\}$ . Let  $S^{d-1} = \{u \in \mathbb{R}^d : \|u\|_2 = 1\}$  denote the  $d$ -dimensional unit sphere. For a square matrix  $T$ , we let  $\|T\|$  denote its spectral norm.

### A.1. Linear policies

Under the assumption of linear Bellman completeness, there is always an optimal policy with the additional structure of being a *linear policy*. Roughly speaking, linear policies take an action defined by the argmax with respect to some fixed coefficient vector: in particular, given  $w \in \mathbb{R}^d$  and  $h \in [H]$ , the associated linear policy at step  $h$  is  $\pi_{h,w}(x) := \arg \max_{a \in \mathcal{A}} \langle w, \phi_h(x, a) \rangle$  if the argmax is unique. It requires some care to appropriately break ties for states  $x$  at which the argmax is not unique.<sup>8</sup> To do so, given  $w \in \mathbb{R}^d$ ,  $h \in [H]$ ,  $x \in \mathcal{X}$ , define  $\mathcal{A}_{h,w}(x) := \arg \max_{a \in \mathcal{A}} \langle w, \phi_h(x, a) \rangle \subset \mathcal{A}$ , where argmax is interpreted as the set of all actions maximizing  $\langle w, \phi_h(x, a) \rangle$ . Then we set, for each  $a \in \mathcal{A}_{h,w}(x)$ ,

$$\mathcal{G}_{h,w}(x, a) := \left\{ \theta \in \mathbb{R}^d : \|\theta\|_2 = 1, \langle \theta, \phi_h(x, a) \rangle > \max_{a' \in \mathcal{A}_{h,w}(x) \setminus \{a\}} \langle \theta, \phi_h(x, a') \rangle \right\}.$$

Let  $\nu_d$  denote the spherical measure on  $S^{d-1}$ . It is straightforward to see that, for all  $w$  and  $x$ ,  $S^{d-1} \setminus \bigcup_{a \in \mathcal{A}_{h,w}(x)} \mathcal{G}_{h,w}(x, a)$  has measure 0 with respect to  $\nu_d$ . We now define,  $\pi_{h,w} : x \rightarrow \Delta(\mathcal{A})$  to be the following randomized policy: for all  $w \in \mathbb{R}^d$ ,  $h \in [H]$ ,  $x \in \mathcal{X}$ ,  $a \in \mathcal{A}$ ,

$$\pi_{h,w}(a|x) := \mathbb{1}\{a \in \mathcal{A}_{h,w}(x)\} \cdot \nu_d(\mathcal{G}_{h,w}(x, a)). \quad (12)$$

It is straightforward to see that  $\pi_{h,w}(\cdot|x) \in \Delta(\mathcal{A})$ . We say that a policy  $\pi$  is a *linear policy* if it is of the form  $\pi = (\pi_{1,w_1}, \dots, \pi_{H,w_H})$ , for  $w_1, \dots, w_H \in \mathbb{R}^d$ . We let  $\Pi_h^{\text{lin}} := \{\pi_{h,w} : w \in \mathbb{R}^d\}$  denote the space of linear policies at step  $h$ , and  $\Pi^{\text{lin}} := \{(\pi_1, \dots, \pi_H) : \pi_h \in \Pi_h^{\text{lin}}\}$  denote the space of linear policies.

<sup>8</sup> In particular, an appropriate tie-breaking procedure is necessary for [Lemma C.3](#) to hold.

## Appendix B. Additional technical overview

Recall that in [Section 3.4](#) we gave an overview of the construction of truncated confidence bonuses (namely, those satisfying the conditions of [Problem 3.2](#)) in the special case that  $d = A = 2$ . In this section, we extend this argument to the case of general values of  $d, A$ .

**The case  $d = 2$  and general  $A$ .** Let us try to extend the above argument (i.e., in [Section 3.4](#)) for the case  $d = 2$  and  $A = 2$  to the case of general  $A$  (still keeping  $d = 2$  fixed). Unfortunately, the function  $x \mapsto \sum_{\epsilon_u, \epsilon_v \in \{\pm 1\}} F_h^{\text{tl}}(x; \beta \epsilon_u \cdot u, \epsilon_v \cdot v)$  we used above no longer satisfies the first condition of [Problem 3.2](#). To see this, suppose that  $\mathcal{A} = \{0, 1, 2, 3\}$  and the feature vectors at some state  $x$  form a rectangle, i.e.,  $\phi_h(x, 0) = (0, 0)$ ,  $\phi_h(x, 1) = (u, 0)$ ,  $\phi_h(x, 2) = (u + v, u + v)$ ,  $\phi_h(x, 3) = (0, v)$ . It is straightforward to see that  $F_h^{\text{tl}}(x; u, v) = 0$ , whereas  $\max_{a, a' \in \mathcal{A}} \min\{|\langle \beta u, \phi_h(x, a) - \phi_h(x, a') \rangle|, |\langle v, \phi_h(x, a) - \phi_h(x, a') \rangle|\} \geq 1$  for any  $\beta \geq 1$  (e.g., take  $a = 0, a' = 2$ ). Thus, [\(11\)](#) cannot hold, even if we only wish for it to hold up to a  $\text{poly}(H, B, d)$  factor.

The above issue runs fairly deep: it can be shown that, under some mild conditions on  $F_h^{\text{tl}}$  which ensure that it be Bellman-linear,<sup>9</sup> there is *no function*  $F_h^{(t)}$  satisfying both conditions of [Problem 3.2](#). To rectify this issue, we will relax the first condition of [Problem 3.2](#), as follows:

**Problem B.1.** In the setting of [Problem 3.2](#), can we find a Bellman-linear function  $F_h^{(t)} : \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$  which satisfies the constraints of [Problem 3.2](#) where [Item 1](#) is replaced by the following constraint:

1. For all  $x \in \mathcal{X}$ ,  $F_h^{(t)}(x)$  is lower bounded as

$$F_h^{(t)}(x) \geq BH \cdot \max_{a, a' \in \mathcal{A}} \min_{\psi \in \bar{\Phi}(x)} \{ \|\beta \Sigma' \cdot (\phi_h(x, a) - \psi)\|_2 + \|\Lambda' \cdot (\psi - \phi_h(x, a'))\|_2 \}, \quad (13)$$

where  $\bar{\Phi}_h(x) := \text{co}\{\phi_h(x, a) : a \in \mathcal{A}\}$  denotes the convex hull of the feature vectors  $\phi_h(x, a)$ ,  $a \in \mathcal{A}$ .

Note that [\(13\)](#) is weaker than [\(11\)](#), as can be verified by considering the choices  $\psi = \phi_h(x, a)$  and  $\psi = \phi_h(x, a')$  for given  $a, a'$  in the outer maximum. As we discuss in [Appendix D.3](#), this weaker lower bound on  $F_h^{(t)}(x)$  is still sufficient to ensure optimism. Moreover, for the example discussed above in which the feature vectors at a state  $x$  form a rectangle, the right-hand side of [\(13\)](#) is 0: if  $a, a'$  in the outer maximum are chosen to represent opposing corners of the rectangle, then taking  $\psi$  to be one of the other two corners yields  $|\langle \beta u, \phi_h(x, a) - \psi \rangle| = |\langle v, \psi - \phi_h(x, a') \rangle| = 0$ . For this example, we therefore have that  $x \mapsto \sum_{\epsilon_u, \epsilon_v \in \{\pm 1\}} F_h^{\text{tl}}(x; \beta \epsilon_u \cdot u, \epsilon_v \cdot v)$  satisfies the properties of [Problem B.1](#).

In fact, this holds more generally in the  $d = 2$  case, as can be seen by the following argument: fix any  $a, a' \in \mathcal{A}$ , and choose

$$\psi^* := \arg \min_{\psi \in \bar{\Phi}_h(x)} \{ |\langle \beta u, \phi_h(x, a) - \psi \rangle| + |\langle v, \psi - \phi_h(x, a') \rangle| \}. \quad (14)$$

Moreover choose  $\epsilon_u, \epsilon_v \in \{\pm 1\}$  so that  $\langle \phi_h(x, a) - \psi^*, \epsilon_u u \rangle \geq 0$  and  $\langle \phi_h(x, a') - \psi^*, \epsilon_v v \rangle \geq 0$ . Note that [\(14\)](#) is a convex optimization problem; its KKT conditions tell us that  $\langle \psi^* - \psi, \beta \epsilon_u u +$

9. Namely, that  $F_h^{\text{tl}}(x; u, v) = \int_{w \in \mathbb{R}^d} f(w) \cdot \max_{a \in \mathcal{A}} \langle \phi_h(x, a), w \rangle d\mu(w)$ , for some Borel measure  $\mu$  and function  $f : \mathbb{R}^d \rightarrow \mathbb{R}$ .

$\epsilon_v v \rangle \geq 0$  for all  $\psi \in \bar{\Phi}_h(x)$ . Then, recalling that  $a, a' \in \mathcal{A}$  were fixed, we have

$$\begin{aligned}
 F_h^{\text{tl}}(x; \beta \epsilon_u u, \epsilon_v v) &\geq \langle \beta \epsilon_u u, \phi_h(x, a) \rangle + \langle \epsilon_v v, \phi_h(x, a') \rangle - \max_{\bar{a} \in \mathcal{A}} \langle \beta \epsilon_u u + \epsilon_v v, \phi_h(x, \bar{a}) \rangle \\
 &\geq \langle \beta \epsilon_u u, \phi_h(x, a) \rangle + \langle \epsilon_v v, \phi_h(x, a') \rangle - \langle \beta \epsilon_u u + \epsilon_v v, \psi^* \rangle \\
 &= \langle \beta \epsilon_u u, \phi_h(x, a) - \psi^* \rangle + \langle \epsilon_v v, \phi_h(x, a') - \psi^* \rangle \\
 &= |\langle \beta u, \phi_h(x, a) - \psi^* \rangle| + |\langle v, \phi_h(x, a') - \psi^* \rangle|, \tag{15}
 \end{aligned}$$

where the second inequality uses the KKT conditions of (14) and the final equality uses the choices of  $\epsilon_u, \epsilon_v$ . (15), combined with non-negativity of  $F_h^{\text{tl}}(\cdot)$ , verifies the first property of Problem B.1, as desired; the second property (namely, Item 2 of Problem 3.2) can be verified by a direct computation (Lemma C.8).

**The case of general  $d$  and general  $A$ .** The argument presented above for  $d = 2$  generalizes to the case of arbitrary  $d \in \mathbb{N}$ . In particular, we will aim to construct a bonus of the form  $G_h^{(t)}(x) := O(BH) \cdot \mathbb{E}_{u \sim \mathcal{D}_{\Sigma'}, v \sim \mathcal{D}_{\Lambda'}} [F_h^{\text{tl}}(x; u, v)]$ , where  $\mathcal{D}_{\Sigma'}, \mathcal{D}_{\Lambda'}$  are distributions supported on  $\mathcal{S}_{\Sigma'}, \mathcal{S}_{\Lambda'}$ , respectively. A key consideration is: *how do we choose the distributions  $\mathcal{D}_{\Sigma'}, \mathcal{D}_{\Lambda'}$ ?* To answer this question, it is instructive to consider how the above argument verifying Item 1 of Problem B.1 generalizes to higher dimensions. In particular, given arbitrary  $a, a' \in \mathcal{A}$ , in lieu of (14), we define

$$\psi^* := \arg \min_{\psi \in \bar{\Phi}_h(x)} \{ \|\beta \Sigma' \cdot (\phi_h(x, a) - \psi)\|_2 + \|\Lambda' \cdot (\psi - \phi_h(x, a'))\|_2 \}. \tag{16}$$

Using the KKT conditions of (16), it can be verified that, for  $u = \frac{\Sigma' \cdot (\phi_h(x, a) - \psi^*)}{\|\Sigma' \cdot (\phi_h(x, a) - \psi^*)\|_2}$  and  $v = \frac{\Lambda' \cdot (\phi_h(x, a') - \psi^*)}{\|\Lambda' \cdot (\phi_h(x, a') - \psi^*)\|_2}$ , we have that

$$F_h^{\text{tl}}(x; \beta u, v) \geq \|\beta \Sigma' \cdot (\phi_h(x, a) - \psi^*)\|_2 + \|\Lambda' \cdot (\phi_h(x, a') - \psi^*)\|_2. \tag{17}$$

Unfortunately these choices of  $u, v$  depend on the state  $x$ , and  $\mathcal{D}_{\Sigma'}, \mathcal{D}_{\Lambda'}$  must be independent of  $x$  in order to ensure Bellman-linearity of the bonus function.

Fortunately, as shown in Lemma C.9, under some additional technical conditions, (17) holds up to an additive term of  $O(\varepsilon)$  when the left-hand side is replaced by  $F_h^{\text{tl}}(x; \beta u', v')$ , for any vectors  $u', v'$  which are  $\varepsilon$ -close in  $\ell_2$  norm to  $u, v$ , respectively. If we take  $\mathcal{D}_{\Sigma'}, \mathcal{D}_{\Lambda'}$  to be the uniform distributions on the unit sphere in the subspaces  $\mathcal{S}_{\Sigma'}, \mathcal{S}_{\Lambda'}$ , respectively, then  $u' \sim \mathcal{D}_{\Sigma'}, v' \sim \mathcal{D}_{\Lambda'}$  satisfy  $\|u' - u\|_2 \leq \varepsilon, \|v' - v\|_2 \leq \varepsilon$  with probability  $\varepsilon^{O(d)}$ . With these choices of  $\mathcal{D}_{\Sigma'}, \mathcal{D}_{\Lambda'}$ , the bonus  $G_h^{(t)}$  satisfies the constraint (13) up to a factor of  $\varepsilon^{-O(d)}$ . In fact, this can easily be improved to a factor of  $\varepsilon^{-O(A)}$  by using the fact that closeness of  $u, u'$  and of  $v, v'$  only needs to hold in the subspace spanned by  $\phi_h(x, a)$  (for  $a \in \mathcal{A}$ ), which has at most  $A$  dimensions – see Lemma C.10.

It turns out to simplify the analysis to take  $\mathcal{D}_{\Sigma'}, \mathcal{D}_{\Lambda'}$  to be the distributions  $\mathcal{N}(0, \Sigma'), \mathcal{N}(0, \Lambda')$ , respectively, which permit essentially the same analysis as described above. Moreover, to satisfy certain technical conditions in Lemma C.9, we add a term to the bonus function which is proportional to  $F_h^{\text{normal}}(x; \Sigma') = \mathbb{E}_{w \sim \mathcal{N}(0, \Sigma')} [\max_{a \in \mathcal{A}} \langle w, \phi_h(x, a) \rangle]$ . To sum up, we define the bonus function as

$$F_h^{(t)}(x) := \beta_1 \cdot \mathbb{E}_{u' \sim \mathcal{N}(0, \Sigma')} \mathbb{E}_{v' \sim \mathcal{N}(0, \Lambda')} \left[ F_h^{\text{tl}}(x; \beta u', v') \right] + \beta_2 \cdot \mathbb{E}_{w \sim \mathcal{N}(0, \Sigma')} \left[ \max_{a \in \mathcal{A}} \langle w, \phi_h(x, a) \rangle \right], \tag{18}$$

where  $\beta_1, \beta_2, \beta$  are carefully chosen parameters bounded above by  $(BHD/\epsilon)^{O(A)}$ , where  $\epsilon$  is the desired error level of the algorithm's output policy. The precise values of these parameters appear in the formal definition of  $F_h^{(t)}$  in (33). We remark that, crucially,  $\beta_2 \ll \beta_1 \cdot \beta$  – the entire point of the truncation procedure we have just described is to ensure that the size of  $F_h^{(t)}(x)$  (which may be as large as  $O(\beta_2)$ ) does *not* scale with  $\beta$ .

In Appendices C.2 and C.3, we formally verify, per the arguments outlined above, that the bonus  $F_h^{(t)}$  in (18) satisfies the properties laid out in Problem B.1. Then, in Appendices D.1 to D.3, we show how this implies the statements of Lemmas 3.1 and 3.2, which suffices to prove our main theorem (Theorem 1.1).

### Appendix C. Bellman linear functions

Recall the definition of Bellman linearity in Definition 3.1. In this section, we discuss several properties of Bellman linear functions.

**Lemma C.1** (Limiting preserves Bellman-linearity). *Fix  $h \in [H]$ . Suppose that for each  $i \in \mathbb{N}$ , we are given a function  $G_i : \mathcal{X} \rightarrow \mathbb{R}$  which is Bellman-linear at step  $h$  and so that  $\sup_{x \in \mathcal{X}, i \in \mathbb{N}} |G_i(x)| \leq C$ , for some  $C > 0$ . Furthermore suppose that, for some  $G : \mathcal{X} \rightarrow \mathbb{R}$ ,  $\lim_{i \rightarrow \infty} G_i(x) = G(x)$  for all  $x \in \mathcal{X}$ . Then  $G$  is Bellman-linear at step  $h$ .*

*Proof.* For each  $i \in \mathbb{N}$ , we know that there is some  $w_i \in \mathbb{R}^d$  so that, for all  $(x, a) \in \mathcal{X} \times \mathcal{A}$ ,  $\langle \phi_{h-1}(x, a), w_i \rangle = \mathbb{E}_{x' \sim P_{h-1}(x, a)} [G_i(x')]$ . For each  $i, a, x$ , we know that  $|\mathbb{E}_{x' \sim P_{h-1}(x, a)} [G_i(x')]| \leq C$  by uniform boundedness of  $G_i$ , so  $|\langle \phi_{h-1}(x, a), w_i \rangle| \leq C$  for all  $i, a, x$ . By definition of  $\mathcal{B}_{h-1}$ , it follows that  $w_i \in C \cdot \mathcal{B}_{h-1}$ , which gives  $\|w_i\|_2 \leq CB$  by Assumption 2.1. Hence we may find a subsequence  $(w_{i_j})_{j \in \mathbb{N}}$  so that  $\lim_{j \rightarrow \infty} w_{i_j} = w^*$ , for some  $w^* \in C \cdot \mathcal{B}_{h-1}$ . Thus, for any  $(x, a) \in \mathcal{X} \times \mathcal{A}$ , we have

$$\begin{aligned} \langle \phi_{h-1}(x, a), w^* \rangle &= \lim_{j \rightarrow \infty} \langle \phi_{h-1}(x, a), w_{i_j} \rangle = \lim_{j \rightarrow \infty} \mathbb{E}_{x' \sim P_{h-1}(x, a)} [G_{i_j}(x')] \\ &= \mathbb{E}_{x' \sim P_{h-1}(x, a)} \left[ \lim_{j \rightarrow \infty} G_{i_j}(x') \right] = \mathbb{E}_{x' \sim P_{h-1}(x, a)} [G(x')], \end{aligned}$$

where the second-to-last equality uses the dominated convergence theorem and uniform boundedness of the functions  $G_i$ .  $\square$

**Lemma C.2** (Differentiation preserves Bellman-linearity). *Fix any  $h, k \in \mathbb{N}$  and consider an open set  $U \in \mathbb{R}^k$  as well as a function  $F : U \times \mathcal{X} \rightarrow \mathbb{R}$  so that  $F(u, x)$  is continuously differentiable in  $u$  for all  $x \in \mathcal{X}$ , and that  $\max_{x \in \mathcal{X}, u \in U} \max\{F(u, x), \|\nabla_u F(u, x)\|_2\} \leq C$ . Suppose that, for all  $u \in U$ , the function  $F(u, \cdot)$  is Bellman-linear at step  $h$ . Then the function  $\nabla_u F(u, \cdot)$  is Bellman-linear at step  $h$ , for all  $u \in U$ .*

*Proof.* By Bellman-linearity, we know that for all  $u \in U$ , there is some  $w_u \in \mathbb{R}^d$  so that, for all  $(x, a) \in \mathcal{X} \times \mathcal{A}$ , it holds that

$$\langle \phi_h(x, a), w_u \rangle = \mathbb{E}_{x' \sim P_{h-1}(x, a)} [F(u, x')].$$

By the dominated convergence theorem (which may be applied as a consequence of uniform boundedness of  $\|\nabla_u F(u, x)\|_2$  over  $u \in U$ ), it holds that  $u \mapsto \langle \phi_h(x, a), w_u \rangle$  is continuously differentiable for all  $x, a$ , and moreover,

$$\nabla_u \langle \phi_h(x, a), w_u \rangle = \mathbb{E}_{x' \sim P_{h-1}(x, a)} [\nabla_u F(u, x')]. \quad (19)$$

Choose  $(x_1, a_1), \dots, (x_\ell, a_\ell) \in \mathcal{X} \times \mathcal{A}$  so that  $\{\phi_h(x_i, a_i)\}_{i=1}^\ell$  is linearly independent and spans  $\{\phi_h(x, a)\}_{(x,a) \in \mathcal{X} \times \mathcal{A}}$ ; then  $\ell \leq d$ . Let  $W_u \in \mathbb{R}^{d \times d}$  be any matrix satisfying

$$W_u \cdot \phi_h(x_i, a_i) = \nabla_u \langle \phi_h(x_i, a_i), w_u \rangle \quad \forall i \in [\ell],$$

which is possible by linear independence of  $\phi_h(x_i, a_i)$ ,  $i \in [\ell]$ . Now consider any  $(x, a) \in \mathcal{X} \times \mathcal{A}$ ; we can find  $\alpha_1, \dots, \alpha_\ell \in \mathbb{R}$  so that  $\phi_h(x, a) = \sum_{i=1}^\ell \alpha_i \cdot \phi_h(x_i, a_i)$ . Then

$$W_u \cdot \phi_h(x, a) = \sum_{i=1}^\ell \alpha_i \cdot \nabla_u \langle \phi_h(x_i, a_i), w_u \rangle = \nabla_u \sum_{i=1}^\ell \alpha_i \cdot \langle \phi_h(x_i, a_i), w_u \rangle = \nabla_u \langle \phi_h(x, a), w_u \rangle,$$

which completes the proof by (19).  $\square$

**Lemma C.3.** *Suppose that  $M$  is linear Bellman complete. Then for each  $h \in [H]$  and each  $w \in \mathbb{R}^d$ , the mapping  $x' \mapsto \phi_{h+1}(x', \pi_{h+1, w}(x'))$  is Bellman-linear at step  $h+1$ , i.e., there is a linear map  $L_h(w) : \mathbb{R}^d \rightarrow \mathbb{R}^d$  so that, for all  $(x, a) \in \mathcal{X} \times \mathcal{A}$ ,*

$$L_h(w)^\top \cdot \phi_h(x, a) = \mathbb{E}_{x' \sim P_h(x, a)}[\phi_{h+1}(x', \pi_{h+1, w}(x'))].$$

*Proof.* Fix  $h \in [H]$ . Define  $V(x', \theta) := \max_{a' \in \mathcal{A}} \langle \phi_{h+1}(x', a'), \theta \rangle$  and  $Q(x, a, \theta) = \langle \phi_h(x, a), \mathcal{T}_h \theta \rangle$ , so that  $Q(x, a, \theta) = \langle \phi_h(x, a), \mathcal{T}_h \theta \rangle = \mathbb{E}_{x' \sim P_h(x, a)}[V(x', \theta)]$ . By Fubini's theorem, for any  $\sigma > 0$ , we have

$$\mathbb{S}_\sigma Q(x, a, \theta) = \mathbb{E}_{z \sim \mathcal{N}(0, \sigma^2 I_d)}[Q(x, a, \theta - z)] = \mathbb{E}_{x' \sim P_h(x, a)} \mathbb{E}_{z \sim \mathcal{N}(0, \sigma^2 I_d)}[V(x', \theta - z)] = \mathbb{E}_{x' \sim P_h(x, a)}[\mathbb{S}_\sigma V(x', \theta)].$$

Linear Bellman completeness of  $M$  gives that for each  $\theta \in \mathbb{R}^d$ , the function  $x' \mapsto V(x', \theta)$ , and thus  $x' \mapsto \mathbb{S}_\sigma V(x', \theta)$  is Bellman-linear at step  $h+1$ . The function  $\mathbb{S}_\sigma V(x', \theta)$  is infinitely differentiable as a function of  $\theta$ , for all  $x'$ ; moreover, since

$$\nabla \mathbb{S}_\sigma V(x', \theta) = \mathbb{E}_{z \sim \mathcal{N}(0, \sigma^2 I_d)}[\nabla V(x', \theta - z)] = \mathbb{E}_{z \sim \mathcal{N}(0, \sigma^2 I_d)}[\phi_{h+1}(x', \pi_{h+1, \theta - z}(x'))], \quad (20)$$

we have that  $\max\{\mathbb{S}_\sigma V(x', \theta), \|\nabla \mathbb{S}_\sigma V(x', \theta)\|_2\}$  is bounded uniformly over  $x'$  in the neighborhood of any  $\theta$ . We remark that the first equality in (20) uses the dominated convergence theorem, and the second equality uses that the  $V(x', \theta - z)$  is differentiable almost surely over  $z \sim \mathcal{N}(0, \sigma^2 I_d)$ . Thus, by Lemma C.2, the function  $x' \mapsto \nabla \mathbb{S}_\sigma V(x', \theta)$  is Bellman-linear at step  $h$ , for all  $\theta \in \mathbb{R}^d$ . But for  $x' \in \mathcal{X}$  and  $w \in \mathbb{R}^d$ ,

$$\nabla \mathbb{S}_\sigma V(x', w) = \phi_{h+1}(x', \pi_{h+1, w, \sigma}(x')),$$

where the equality above uses (20) together with Definition F.1. Next, (72) ensures that for each  $x' \in \mathcal{X}$ ,  $\lim_{\sigma \rightarrow 0^+} \phi_{h+1}(x', \pi_{h+1, w, \sigma}(x')) = \phi_{h+1}(x', \pi_{h+1, w}(x'))$ . Then Lemma C.1 (applied to each coordinate of the function  $x' \mapsto \phi_{h+1}(x', \pi_{h+1, w, \sigma}(x'))$ , for  $\sigma > 0$ ) yields that  $x' \mapsto \phi_{h+1}(x', \pi_{h+1, w}(x'))$  is Bellman-linear at step  $h+1$ , for any  $w \in \mathbb{R}^d$ .  $\square$

Given any  $\pi_{h+1} \in \Pi_{h+1}^{\text{lin}}$  (so that  $\pi_{h+1} = \pi_{h+1, v}$  for some  $v \in \mathbb{R}^d$ ) and  $w \in \mathbb{R}^d$ , we define  $\mathcal{T}_h^\pi w := L_h(v) \cdot w$ . Then by Lemma C.3, for all  $(x, a) \in \mathcal{X} \times \mathcal{A}$ ,

$$\langle \phi_h(x, a), \mathcal{T}_h^\pi w \rangle = \mathbb{E}_{x' \sim P_h(x, a)}[\langle \phi_{h+1}(x', \pi_{h+1}(x')), w \rangle]. \quad (21)$$



Note that  $\mathcal{T}_h^\pi$  depends only on  $\pi_{h+1}$ ; thus with slight abuse of notation, we will sometimes write  $\mathcal{T}_h^{\pi_{h+1}}w$  in place of  $\mathcal{T}_h^\pi w$ . We also remark that by linearity of expectation, for any  $\pi \in \Pi^{\text{Plin}}$  and  $w \in \mathbb{R}^d$ , there is a vector  $\mathcal{T}_h^\pi w \in \mathbb{R}^d$  so that (21) holds for the perturbed linear policy  $\pi$ .

As an immediate consequence of Lemma C.3, we can show that the  $Q$ -function for a linear policy is linear.

**Corollary C.4.** *Suppose that  $M$  is linear Bellman complete, and that  $\pi \in \Pi^{\text{Plin}}$ . Then for each  $h \in [H]$ , there is a vector  $w_h^\pi \in H \cdot \mathcal{B}_h \subset \mathbb{R}^d$  so that for all  $(x, a) \in \mathcal{X} \times \mathcal{A}$ ,  $Q_h^\pi(x, a) = \langle w_h^\pi, \phi_h(x, a) \rangle$ . Moreover,  $\|w_h^\pi\|_2 \leq HB$ .*

*Proof.* Since  $Q_h^\pi(x, a) = r_h(x, a) + \mathbb{E}_{x' \sim P_h(x, a)}[Q_{h+1}^\pi(x, \pi_{h+1}(x))]$ , by induction and the fact that  $r_h(x, a)$  is linear, it suffices to show that for all  $v \in \mathbb{R}^d$ , there is some  $w \in \mathbb{R}^d$  so that  $\langle \phi_h(x, a), w \rangle = \mathbb{E}_{x' \sim P_h(x, a)}[\langle \phi_h(x, \pi_{h+1}(x)), v \rangle]$ . We may write  $\pi_{h+1} = \pi_{h+1, y, \sigma}$  for some  $y \in \mathbb{R}^d$  and  $\sigma \geq 0$ ; then by Lemma C.3, we may take  $w = \mathbb{E}_{z \sim \mathcal{N}(y, \sigma^2)}[L_h(z)] \cdot v$ .

To see the upper bound on  $\|w_h^\pi\|_2$ , note that, by definition of  $Q_h^\pi$ , we have  $|\langle w_h^\pi, \phi_h(x, a) \rangle| \leq H$  for all  $x, a, h$ . Then it follows that  $\|w_h^\pi\|_2 \leq HB$  by Assumption 2.1.  $\square$

### C.1. Construction of Bellman-linear functions for bonuses

A central component of the proof consists of the construction of functions, to be used as exploration bonuses, which (a) are Bellman-linear and (b) approximate a truncated version of arbitrary norms induced by PSD matrices. In this section, we establish Bellman linearity of certain functions which form the building blocks of these bonuses.

Given orthogonal vectors  $u, v \in \mathbb{R}^d$ , we let  $\Pi_{u, v} := \frac{uu^\top}{\|u\|_2} + \frac{vv^\top}{\|v\|_2}$ . Note that if  $u, v$  are unit vectors, then  $\Pi_{u, v}$  is the projection matrix onto  $\text{Span}\{u, v\}$ .

**Definition C.1** (Orthogonal pair). We define a tuple of  $d \times d$  PSD matrices  $(\Sigma, \Lambda)$  to be an *orthogonal pair* if the following equalities hold:

$$\Sigma^2 = \Sigma, \quad \Lambda^2 = \Lambda, \quad \Sigma\Lambda = \Lambda\Sigma = 0, \quad \Sigma + \Lambda = I_d.$$

For an orthogonal pair  $(\Sigma, \Lambda)$ ,  $\Sigma, \Lambda$  are projections onto subspaces of  $\mathbb{R}^d$ , which we denote by  $\mathcal{S}_\Sigma := \{\Sigma v : v \in \mathbb{R}^d\}$ ,  $\mathcal{S}_\Lambda := \{\Lambda v : v \in \mathbb{R}^d\}$ , respectively. Since  $\Sigma\Lambda = \Lambda\Sigma = 0$ , the subspaces  $\mathcal{S}_\Sigma, \mathcal{S}_\Lambda$  are orthogonal. Since  $\Sigma + \Lambda = I_d$ , we have that  $\dim(\mathcal{S}_\Sigma) + \dim(\mathcal{S}_\Lambda) = d$ .

**Definition C.2** (Truncated linear bonus). Consider orthogonal vectors  $u, v \in \mathbb{R}^d$ . For  $\Phi \in \mathcal{P}^d$ , we define

$$F^{\text{tl}}(\Phi; u, v) = \max_{\phi \in \Phi} \langle u, \phi \rangle + \max_{\phi \in \Phi} \langle v, \phi \rangle - \max_{\phi \in \Phi} \langle u + v, \phi \rangle.$$

For  $x \in \mathcal{X}$  and  $h \in [H]$ , we then define  $F_h^{\text{tl}}(x; u, v) := F^{\text{tl}}(\bar{\Phi}_h(x); u, v)$ .

The following lemma is immediate from the definition of  $F^{\text{tl}}(\cdot)$ .

**Lemma C.5** (Bellman-linearity of truncated linear bonus). *Given any  $h > 1$  and vectors  $u, v \in \mathbb{R}^d$ , the mapping  $x \mapsto F_h^{\text{tl}}(x; u, v)$  is Bellman-linear at step  $h$ .*

We next record the following lemmas for future use, which follow as an immediate consequence of the definition of  $F^{\text{tl}}$ .

**Lemma C.6.** Let  $\Phi \in \mathcal{P}^d$ . Suppose that  $u, u', v, v' \in \mathbb{R}^d$  satisfy  $\langle u, \phi \rangle = \langle u', \phi \rangle$  and  $\langle v, \phi \rangle = \langle v', \phi \rangle$  for all  $\phi \in \Phi$ . Then

$$F^{\text{tl}}(\Phi; u, v) = F^{\text{tl}}(\Phi; u', v').$$

*Proof.* We have that  $\max_{\phi \in \Phi} \langle w, \phi \rangle = \max_{\phi \in \Phi} \langle w', \phi \rangle$  for each pair  $(w, w') \in \{(u, u'), (v, v'), (u+v, u'+v')\}$ .  $\square$

**Lemma C.7.** Let  $\Phi \in \mathcal{P}^d$ . Suppose that  $u, v \in \mathbb{R}^d$  are given and that  $\alpha_u, \alpha_v \geq 0$  are real numbers. Then

$$F^{\text{tl}}(\Phi; \alpha_u \cdot u, \alpha_v \cdot v) \geq \min\{\alpha_u, \alpha_v\} \cdot F^{\text{tl}}(\Phi; u, v).$$

*Proof.* Since  $F^{\text{tl}}(\Phi; \alpha u, \alpha v) = \alpha \cdot F^{\text{tl}}(\Phi; u, v)$  for any  $\alpha \geq 0$ , it suffices to consider the case that  $\min\{\alpha_u, \alpha_v\} = 1$ . Without loss of generality let us suppose that  $\alpha_u = 1 \leq \alpha_v$ . Then

$$F^{\text{tl}}(\Phi; u, \alpha_v \cdot v) - F^{\text{tl}}(\Phi; u, v) = (\alpha_v - 1) \cdot \max_{\phi \in \Phi} \langle v, \phi \rangle - \max_{\phi \in \Phi} \langle u + \alpha_v \cdot v, \phi \rangle + \max_{\phi \in \Phi} \langle u + v, \phi \rangle \geq 0,$$

where the final inequality follows since  $u + \alpha_v \cdot v = (u + v) + (\alpha_v - 1) \cdot v$ .  $\square$

## C.2. An upper bound on the truncated linear bonus

[Lemma C.8](#) below shows that the truncated linear bonus  $F^{\text{tl}}(\Phi; u, v)$  can be bounded above by the length of the projection of  $\Phi$  onto each of  $u$  and  $v$ .

**Lemma C.8.** Suppose  $u, v \in \mathbb{R}^d$  and  $\Phi \in \mathcal{P}^d$ . Then

$$0 \leq F^{\text{tl}}(\Phi; u, v) \leq 2 \cdot \min \left\{ \max_{\phi, \phi' \in \Phi} \langle u, \phi - \phi' \rangle, \max_{\phi, \phi' \in \Phi} \langle v, \phi - \phi' \rangle \right\}.$$

*Proof.* For any fixed  $\phi$ , we have

$$\langle u + v, \phi \rangle \leq \max_{\phi' \in \Phi} \langle u, \phi' \rangle + \max_{\phi' \in \Phi} \langle v, \phi' \rangle,$$

which implies  $F^{\text{tl}}(\Phi; u, v) \geq 0$ .

To upper bound  $F^{\text{tl}}(\Phi; u, v)$ , we first note that

$$\begin{aligned} F^{\text{tl}}(\Phi; u, v) &\leq F^{\text{tl}}(\Phi; u, v) + F^{\text{tl}}(\Phi; -u, -v) \\ &= \max_{\phi, \phi' \in \Phi} \langle u, \phi - \phi' \rangle + \max_{\phi, \phi' \in \Phi} \langle v, \phi - \phi' \rangle - \max_{\phi, \phi' \in \Phi} \langle u + v, \phi - \phi' \rangle. \end{aligned}$$

We will now upper bound  $F^{\text{tl}}(\Phi; u, v) + F^{\text{tl}}(\Phi; -u, -v)$  as follows. By symmetry, it is without loss of generality to assume that  $\max_{\phi, \phi' \in \Phi} \langle u, \phi - \phi' \rangle \leq \max_{\phi, \phi' \in \Phi} \langle v, \phi - \phi' \rangle$ . Now we have

$$\begin{aligned} \max_{\phi, \phi' \in \Phi} \langle u + v, \phi - \phi' \rangle &\geq \max_{\phi, \phi' \in \Phi} \langle v, \phi - \phi' \rangle - \max_{\phi, \phi' \in \Phi} \langle -u, \phi - \phi' \rangle \\ &= \max_{\phi, \phi' \in \Phi} \langle v, \phi - \phi' \rangle - \max_{\phi, \phi' \in \Phi} \langle u, \phi - \phi' \rangle. \end{aligned}$$

Rearranging, we see that  $F^{\text{tl}}(\Phi; u, v) + F^{\text{tl}}(\Phi; -u, -v) \leq 2 \max_{\phi, \phi' \in \Phi} \langle u, \phi - \phi' \rangle$ , as desired.  $\square$

### C.3. A lower bound on the modified truncated linear bonus

Given  $\Phi \in \mathcal{P}^d$  and a matrix  $\Gamma \in \mathbb{R}^{d \times d}$ , let  $\mathcal{S}_\Phi \subset \mathbb{R}^d$  denote the subspace  $\mathcal{S}_\Phi := \text{Span}\{\phi : \phi \in \Phi\}$ , and  $\mathcal{S}_{\Phi, \Gamma} \subset \mathbb{R}^d$  denote the subspace  $\mathcal{S}_{\Phi, \Gamma} := \text{Span}\{\Gamma \cdot \phi : \phi \in \Phi\} = \Gamma \cdot \mathcal{S}_\Phi$ . Let  $\Pi_{\Phi, \Gamma} \in \mathbb{R}^{d \times d}$  denote the matrix which projects onto  $\mathcal{S}_{\Phi, \Gamma}$ . Recall that for an orthogonal pair  $(\Sigma, \Lambda)$  (Definition C.1), we have  $\mathcal{S}_\Sigma = \Sigma \cdot \mathbb{R}^d$ ,  $\mathcal{S}_\Lambda = \Lambda \cdot \mathbb{R}^d$ . Then for any  $\Phi \in \mathcal{P}^d$ ,  $\mathcal{S}_{\Phi, \Sigma} \subset \mathcal{S}_\Sigma$  and  $\mathcal{S}_{\Phi, \Lambda} \subset \mathcal{S}_\Lambda$ .

Moreover, given  $\Phi \in \mathcal{P}^d$  and  $\phi_1, \phi_2 \in \Phi$ , we define

$$\mu_{\Phi, \phi_1, \phi_2}^* := \arg \min_{\xi \in \Phi} \{\|\beta \Sigma \cdot (\phi_1 - \xi)\|_2 + \|\Lambda \cdot (\xi - \phi_2)\|_2\}. \quad (22)$$

The point  $\mu_{\Phi, \phi_1, \phi_2}^*$  can be thought of as a sort of ‘‘midpoint’’ between  $\phi_1$  and  $\phi_2$  in  $\Phi$ , where the distance between  $\phi_1$  and  $\mu_{\Phi, \phi_1, \phi_2}^*$  is measured with respect to  $\beta \Sigma$  and the distance between  $\mu_{\Phi, \phi_1, \phi_2}^*$  and  $\phi_2$  is measured with respect to  $\Lambda$ . As it turns out, this particular notion of midpoint will be useful in proving that our exploration bonuses induce optimistic value functions.

Finally, we introduce some notation regarding the normal cone of a convex body. For a convex subset  $\mathcal{C} \subset \mathbb{R}^d$  and  $z \in \mathcal{C}$ , the normal cone of  $\mathcal{C}$  at  $z$ , denoted  $N_{\mathcal{C}}(z)$  is defined as  $N_{\mathcal{C}}(z) := \{y \in \mathbb{R}^d : \langle y, z - z' \rangle \geq 0 \forall z' \in \mathcal{C}\}$ .

**Lemma C.9.** *Consider an orthogonal pair  $(\Sigma, \Lambda)$  and  $\beta \geq 1, \zeta > 0$ . Suppose that  $\Phi \in \mathcal{P}^d$  satisfies*

$$\max \left\{ \max_{\phi_1, \phi_2 \in \Phi} \|\beta \Sigma \cdot (\phi_1 - \phi_2)\|_2, \max_{\phi_1, \phi_2 \in \Phi} \|\Lambda \cdot (\phi_1 - \phi_2)\|_2 \right\} \leq \zeta. \quad (23)$$

Fix any  $\phi_1, \phi_2 \in \Phi$ . Then there are unit vectors  $u \in \mathcal{S}_{\Phi, \Sigma}$  and  $v \in \mathcal{S}_{\Phi, \Lambda}$  (depending on  $\Phi, \phi_1, \phi_2$ ) so that for all unit vectors  $u' \in \mathcal{S}_\Sigma, v' \in \mathcal{S}_\Lambda$  satisfying

$$\|u - u'\|_2 \leq \varepsilon, \quad \|v - v'\|_2 \leq \varepsilon, \quad \langle u', u \rangle \geq \eta, \quad \langle v', v \rangle \geq \eta, \quad (24)$$

it holds that

$$F^{\text{tl}}(\Phi; \beta u', v') \geq \eta \|\beta \Sigma \cdot (\phi_1 - \mu_{\Phi, \phi_1, \phi_2}^*)\|_2 + \eta \|\Lambda \cdot (\mu_{\Phi, \phi_1, \phi_2}^* - \phi_2)\|_2 - 2\varepsilon \zeta.$$

*Proof.* Fix  $\Phi \in \mathcal{P}^d$  satisfying (23) and  $\phi_1, \phi_2 \in \Phi$ .

**Step 1: Choosing  $u, v$ .** The KKT conditions for optimality of  $\mu_{\Phi, \phi_1, \phi_2}^*$  (as defined in (22)) give that, for all  $z \in \bar{\Phi}$ ,

$$\frac{\beta^2 \Sigma \cdot (\phi_1 - \mu_{\Phi, \phi_1, \phi_2}^*)}{\|\beta \Sigma \cdot (\phi_1 - \mu_{\Phi, \phi_1, \phi_2}^*)\|_2} + \frac{\Lambda \cdot (\phi_2 - \mu_{\Phi, \phi_1, \phi_2}^*)}{\|\Lambda \cdot (\phi_2 - \mu_{\Phi, \phi_1, \phi_2}^*)\|_2} \in N_{\Phi}(\mu_{\Phi, \phi_1, \phi_2}^*), \quad (25)$$

where  $N_{\Phi}(\xi)$  denotes the normal cone of  $\Phi$  at  $\xi$ . We now set

$$u := \frac{\Sigma \cdot (\phi_1 - \mu_{\Phi, \phi_1, \phi_2}^*)}{\|\Sigma \cdot (\phi_1 - \mu_{\Phi, \phi_1, \phi_2}^*)\|_2}, \quad v := \frac{\Lambda \cdot (\phi_2 - \mu_{\Phi, \phi_1, \phi_2}^*)}{\|\Lambda \cdot (\phi_2 - \mu_{\Phi, \phi_1, \phi_2}^*)\|_2}.$$

Then (25) states that for all  $z \in \Phi$ ,  $\langle \beta u + v, \mu_{\Phi, \phi_1, \phi_2}^* - z \rangle \geq 0$ . Moreover, since  $\mu_{\Phi, \phi_1, \phi_2}^* \in \Phi$ , we certainly have that  $u \in \mathcal{S}_{\Phi, \Sigma} \subset \mathcal{S}_\Sigma$ ,  $v \in \mathcal{S}_{\Phi, \Lambda} \subset \mathcal{S}_\Lambda$ .

**Step 2: Relating  $u, v$  to  $u', v'$  satisfying (24).** Now consider any unit vectors  $u' \in \mathcal{S}_\Sigma, v' \in \mathcal{S}_\Lambda$  satisfying (24). Choose

$$\xi' := \arg \max_{\xi \in \Phi} \langle \beta u' + v', \xi \rangle. \quad (26)$$

Next, we compute

$$\begin{aligned} \langle \beta u' + v', \xi' - \mu_{\Phi, \phi_1, \phi_2}^* \rangle &= \langle (\beta u' + v') - (\beta u + v), \xi' - \mu_{\Phi, \phi_1, \phi_2}^* \rangle + \langle \beta u + v, \xi' - \mu_{\Phi, \phi_1, \phi_2}^* \rangle \\ &\leq \beta |\langle u - u', \xi' - \mu_{\Phi, \phi_1, \phi_2}^* \rangle| + |\langle v - v', \xi' - \mu_{\Phi, \phi_1, \phi_2}^* \rangle| \\ &= \beta |\langle u - u', \Sigma \cdot (\xi' - \mu_{\Phi, \phi_1, \phi_2}^*) \rangle| + |\langle v - v', \Lambda \cdot (\xi' - \mu_{\Phi, \phi_1, \phi_2}^*) \rangle| \\ &= \beta |\langle \Pi_{\Phi, \Sigma} \cdot (u - u'), \Sigma \cdot (\xi' - \mu_{\Phi, \phi_1, \phi_2}^*) \rangle| + |\langle \Pi_{\Phi, \Lambda} \cdot (v - v'), \Lambda \cdot (\xi' - \mu_{\Phi, \phi_1, \phi_2}^*) \rangle| \\ &\leq \beta \|\Pi_{\Phi, \Sigma} \cdot (u - u')\|_2 \cdot \|\Sigma \cdot (\xi' - \mu_{\Phi, \phi_1, \phi_2}^*)\|_2 + \|\Pi_{\Phi, \Lambda} \cdot (v - v')\|_2 \cdot \|\Lambda \cdot (\xi' - \mu_{\Phi, \phi_1, \phi_2}^*)\|_2 \\ &\leq 2\varepsilon\zeta, \end{aligned} \quad (27)$$

where the first inequality uses the triangle inequality and the fact that  $\langle \beta u + v, \xi' - \mu_{\Phi, \phi_1, \phi_2}^* \rangle \leq 0$  since  $\xi' \in \Phi$ ; the second equality uses the fact that  $u - u' \in \mathcal{S}_\Sigma$  and  $v - v' \in \mathcal{S}_\Lambda$ ; the third equality uses the fact that  $\Sigma \cdot (\xi' - \mu_{\Phi, \phi_1, \phi_2}^*) \in \mathcal{S}_{\Phi, \Sigma}$  and  $\Lambda \cdot (\xi' - \mu_{\Phi, \phi_1, \phi_2}^*) \in \mathcal{S}_{\Phi, \Lambda}$ ; the second inequality uses Cauchy-Schwarz; and the third inequality uses (24) and (23).

We then have

$$\begin{aligned} &\langle \phi_1 - \xi', \beta u' \rangle + \langle \phi_2 - \xi', v' \rangle \\ &= \langle \Sigma \cdot (\phi_1 - \xi'), \beta u' + v' \rangle + \langle \Lambda \cdot (\phi_2 - \xi'), v' + \beta u' \rangle \\ &= \langle \Sigma \cdot \phi_1 + \Lambda \cdot \phi_2 - \xi', \beta u' + v' \rangle \\ &\geq \langle \Sigma \cdot \phi_1 + \Lambda \cdot \phi_2 - \mu_{\Phi, \phi_1, \phi_2}^*, \beta u' + v' \rangle - 2\varepsilon\zeta \\ &= \langle \Sigma \cdot (\phi_1 - \mu_{\Phi, \phi_1, \phi_2}^*), \beta u' \rangle + \langle \Lambda \cdot (\phi_2 - \mu_{\Phi, \phi_1, \phi_2}^*), v' \rangle - 2\varepsilon\zeta \\ &\geq \eta\beta \|\Sigma \cdot (\phi_1 - \mu_{\Phi, \phi_1, \phi_2}^*)\|_2 + \eta \|\Lambda \cdot (\phi_2 - \mu_{\Phi, \phi_1, \phi_2}^*)\|_2 - 2\varepsilon\zeta, \end{aligned} \quad (28)$$

where the first equality uses that  $\Sigma u' = u', \Sigma v' = v',$  and  $\Sigma v' = \Lambda u' = 0$  since  $(\Sigma, \Lambda)$  is an orthogonal pair, the second equality uses that  $\Sigma + \Lambda = I_d,$  the first inequality uses (27), the third equality again uses  $\Sigma + \Lambda = I_d$  and the fact that  $\Sigma v' = \Lambda u' = 0,$  and the second inequality uses (24) and the definition of  $u, v.$

**Step 3: wrapping up.** Finally, we may write

$$\begin{aligned} F^{\text{tl}}(\Phi; \beta u', v') &= \max_{\phi \in \Phi} \langle \phi, \beta u' \rangle + \max_{\phi \in \Phi} \langle \phi, v' \rangle - \max_{\phi \in \Phi} \langle \phi, \beta u' + v' \rangle \\ &= \max_{\phi \in \Phi} \langle \phi, \beta u' \rangle + \max_{\phi \in \Phi} \langle \phi, v' \rangle - \langle \xi, \beta u' + v' \rangle \\ &\geq \langle \phi_1, \beta u' \rangle + \langle \phi_2, v' \rangle - \langle \xi, \beta u' + v' \rangle \\ &\geq \eta\beta \|\Sigma \cdot (\phi_1 - \mu_{\Phi, \phi_1, \phi_2}^*)\|_2 + \eta \|\Lambda \cdot (\phi_2 - \mu_{\Phi, \phi_1, \phi_2}^*)\|_2 - 2\varepsilon\zeta, \end{aligned}$$

where the second equality uses the definition of  $\xi$  in (26), and the second inequality uses (28).  $\square$

By averaging over  $u', v'$  drawn randomly from appropriate subspaces, we have the following consequence of Lemma C.9.

**Lemma C.10.** *There is a constant  $C_{C.10}$  so that the following holds. Fix  $A \in \mathbb{N}$ , an orthogonal pair  $(\Sigma, \Lambda)$ ,  $\beta \geq 1$ ,  $\zeta > 0$ ,  $\Phi \in \mathcal{P}_A^d$ ,  $\phi_1, \phi_2 \in \Phi$ , and suppose that  $\Phi$  satisfies (23). Then for any  $\varepsilon > 0$ ,*

$$\begin{aligned} & \left( \frac{C_{C.10}}{\varepsilon} \right)^{2A} \cdot \mathbb{E}_{u' \sim \mathcal{N}(0, \Sigma)} \mathbb{E}_{v' \sim \mathcal{N}(0, \Lambda)} \left[ F^{\text{tl}}(\Phi; \beta u', v') \right] \\ & \geq \|\beta \Sigma \cdot (\phi_1 - \mu_{\Phi, \phi_1, \phi_2}^*)\|_2 + \|\Lambda \cdot (\mu_{\Phi, \phi_1, \phi_2}^* - \phi_2)\|_2 - 4\varepsilon \zeta. \end{aligned}$$

*Proof.* Note that if  $\dim(\mathcal{S}_{\Phi, \Sigma}) = 0$ , then we could choose  $\mu_{\Phi, \phi_1, \phi_2}^* = \phi_2$  in (22), and thus the lemma statement becomes immediate. A symmetric argument applies to the case that  $\dim(\mathcal{S}_{\Phi, \Lambda}) = 0$ . Thus we may assume for the remainder of the proof that  $\dim(\mathcal{S}_{\Phi, \Sigma}), \dim(\mathcal{S}_{\Phi, \Lambda}) \geq 1$ .

Note that the subspace  $\mathcal{S}_{\Phi, \Sigma} \subseteq \mathcal{S}_{\Sigma}$  satisfies  $\dim(\mathcal{S}_{\Phi, \Sigma}) \leq A$ , and similarly,  $\dim(\mathcal{S}_{\Phi, \Lambda}) \leq A$ . Note that  $\mathbb{E}_{u' \sim \mathcal{N}(0, \Sigma)} [(\Pi_{\Phi, \Sigma} u') \cdot (\Pi_{\Phi, \Sigma} u')^\top] = \Pi_{\Phi, \Sigma} \cdot \Sigma \cdot \Pi_{\Phi, \Sigma} = \Pi_{\Phi, \Sigma}$ , meaning that  $\Pi_{\Phi, \Sigma} u' \sim \mathcal{N}(0, \Pi_{\Phi, \Sigma})$ . Similarly, we have that  $\Pi_{\Phi, \Lambda} v' \sim \mathcal{N}(0, \Pi_{\Phi, \Lambda})$ .

Therefore, we have that  $\frac{\Pi_{\Phi, \Sigma} u'}{\|\Pi_{\Phi, \Sigma} u'\|_2} \sim \text{unif}(S^{d-1} \cap \mathcal{S}_{\Phi, \Sigma})$ . Since  $S^{d-1} \cap \mathcal{S}_{\Phi, \Sigma}$  may be covered with  $(3/\varepsilon)^A$  balls of radius  $\varepsilon$ , it follows that, for any  $u \in S^{d-1} \cap \mathcal{S}_{\Phi, \Sigma}$ , we have

$$\Pr_{u' \sim \mathcal{N}(0, \Sigma)} \left( \left\| \frac{\Pi_{\Phi, \Sigma} u'}{\|\Pi_{\Phi, \Sigma} u'\|_2} - u \right\|_2 \leq \varepsilon \right) \geq (\varepsilon/3)^A.$$

Moreover, we have that  $\Pr_{u' \sim \mathcal{N}(0, \Sigma)} (\|\Pi_{\Phi, \Sigma} u'\|_2 \geq 1/2) \geq 1/2$ : this holds because  $\|\Pi_{\Phi, \Sigma} u'\|_2^2$  is a chi-squared random variable with  $\dim(\mathcal{S}_{\Phi, \Sigma}) \geq 1$  degrees of freedom. Using the fact that  $\|\Pi_{\Phi, \Sigma} u'\|_2$  and  $\frac{\Pi_{\Phi, \Sigma} u'}{\|\Pi_{\Phi, \Sigma} u'\|_2}$  are independent random variables, it follows that

$$\frac{1}{2} \cdot (\varepsilon/3)^A \leq \Pr_{u' \sim \mathcal{N}(0, \Sigma)} \left( \left\| \frac{\Pi_{\Phi, \Sigma} u'}{\|\Pi_{\Phi, \Sigma} u'\|_2} - u \right\|_2 \leq \varepsilon \text{ and } \|\Pi_{\Phi, \Sigma} u'\|_2 \geq 1/2 \right).$$

In a symmetric manner, we obtain that

$$\frac{1}{2} \cdot (\varepsilon/3)^A \leq \Pr_{v' \sim \mathcal{N}(0, \Lambda)} \left( \left\| \frac{\Pi_{\Phi, \Lambda} v'}{\|\Pi_{\Phi, \Lambda} v'\|_2} - v \right\|_2 \leq \varepsilon \text{ and } \|\Pi_{\Phi, \Lambda} v'\|_2 \geq 1/2 \right).$$

Let us define the random variables  $u'' := \frac{\Pi_{\Phi, \Sigma} u'}{\|\Pi_{\Phi, \Sigma} u'\|_2}$  and  $v'' := \frac{\Pi_{\Phi, \Lambda} v'}{\|\Pi_{\Phi, \Lambda} v'\|_2}$ . Let  $u \in \mathcal{S}_{\Phi, \Sigma}$  and  $v \in \mathcal{S}_{\Phi, \Lambda}$  be chosen according to the statement of Lemma C.9 given  $\Phi, \phi_1, \phi_2$ . Then with probability at least  $(1/4) \cdot (\varepsilon/3)^{2A}$  over the independent draws of  $u'$  and  $v'$  (which induce values of  $u'', v''$ ), we have that:

$$\|u'' - u\|_2 \leq \varepsilon, \quad \|v'' - v\|_2 \leq \varepsilon, \quad \langle u'', u \rangle \geq 1/2, \quad \langle v'', v \rangle \geq 1/2, \quad \|\Pi_{\Phi, \Sigma} u'\|_2 \geq 1/2, \quad \|\Pi_{\Phi, \Lambda} v'\|_2 \geq 1/2, \quad (29)$$

where we have used that  $\varepsilon \leq 1$  and that for unit vectors  $y, y'$ , we have  $\langle y, y' \rangle = \frac{2 - \|y - y'\|_2^2}{2}$ .

For any  $\phi \in \Phi \subset \mathcal{S}_{\Phi}$ , we have

$$\langle u', \phi \rangle = \langle u', \Sigma \phi \rangle = \langle \Pi_{\Phi, \Sigma} u', \Sigma \phi \rangle = \langle \Pi_{\Phi, \Sigma} u', \phi \rangle,$$

where the first equality uses that  $u' \in \mathcal{S}_\Sigma$ , the second equality uses that  $\Sigma\phi \in \mathcal{S}_{\Phi,\Sigma}$ , and the third equality uses that  $\Pi_{\Phi,\Sigma}u' \in \mathcal{S}_{\Phi,\Sigma} \subseteq \mathcal{S}_\Sigma$ . In a similar manner, we have that, for all  $\phi \in \Phi$ ,  $\langle v', \phi \rangle = \langle \Pi_{\Phi,\Lambda}v', \phi \rangle$ . It then follows from [Lemma C.6](#) that

$$F^{\text{tl}}(\Phi; \beta u', v') = F^{\text{tl}}(\Phi; \beta \cdot \Pi_{\Phi,\Sigma}u', \Pi_{\Phi,\Lambda}v'). \quad (30)$$

Under the event that  $\|\Pi_{\Phi,\Sigma}u'\|_2 \geq 1/2$  and  $\|\Pi_{\Phi,\Lambda}v'\|_2 \geq 1/2$ , we have by [Lemma C.7](#) that

$$F^{\text{tl}}(\Phi; \beta \Pi_{\Phi,\Sigma}u', \Pi_{\Phi,\Lambda}v') \geq \frac{1}{2} \cdot F^{\text{tl}}(\Phi; \beta u'', v''). \quad (31)$$

Combining [Eqs. \(30\) and \(31\)](#), we see that

$$\begin{aligned} & \mathbb{E}_{u' \sim \mathcal{N}(0,\Sigma)} \mathbb{E}_{v' \sim \mathcal{N}(0,\Lambda)} \left[ F^{\text{tl}}(\Phi; \beta u', v') \right] \\ & \geq \frac{1}{2} \mathbb{E}_{u' \sim \mathcal{N}(0,\Sigma)} \mathbb{E}_{v' \sim \mathcal{N}(0,\Lambda)} \left[ \tilde{F}^{\text{tl}}(\Phi; \beta u'', v'') \right] \\ & \geq \frac{1}{4} \cdot \frac{1}{4} \cdot \left( \frac{\varepsilon}{3} \right)^{2A} \cdot (\|\beta \Sigma \cdot (\phi_1 - \mu_{\Phi,\phi_1,\phi_2}^* )\|_2 + \|\Lambda \cdot (\mu_{\Phi,\phi_1,\phi_2}^* - \phi_2)\|_2 - 4\varepsilon\zeta), \end{aligned}$$

where the second inequality uses the fact that [\(29\)](#) holds with probability at least  $(1/4) \cdot (\varepsilon/3)^{2A}$  together with [Lemma C.9](#) with  $\eta = 1/2$ ,  $u'$  set to  $u''$ , and  $v'$  set to  $v''$ . Rearranging and using that  $A \geq 2$ , we obtain that

$$\begin{aligned} & \left( \frac{6}{\varepsilon} \right)^{2A} \cdot \mathbb{E}_{u' \sim \mathcal{N}(0,\Sigma)} \mathbb{E}_{v' \sim \mathcal{N}(0,\Lambda)} \left[ \tilde{F}^{\text{tl}}(\Phi; \beta u', v') \right] \\ & \geq \|\beta \Sigma \cdot (\phi_1 - \mu_{\Phi,\phi_1,\phi_2}^* )\|_2 + \|\Lambda \cdot (\mu_{\Phi,\phi_1,\phi_2}^* - \phi_2)\|_2 - 4\varepsilon\zeta, \end{aligned}$$

as desired. (In particular, we may take  $C_{C.10} = 6$ .)  $\square$

## Appendix D. The algorithm

**The algorithm.** [Algorithm 1](#) presents our algorithm for learning linear Bellman complete MDPs. For some  $T \in \mathbb{N}$ , the algorithm proceeds for  $T$  rounds. In each round  $t \in [T]$ , the algorithm computes two policies  $\hat{\pi}_h^{(t)}, \tilde{\pi}_h^{(t)} : \mathcal{X} \rightarrow \mathcal{A}$  at step  $h$ , in order of decreasing  $h$ . For each step  $h$ , trajectories are gathered from the uniform mixture over  $\hat{\pi}^{(s)} \circ_h \tilde{\pi}^{(s)} \circ_{h+1} \hat{\pi}^{(t)}$ , for  $s \in [t-1]$  ([Line 3](#)). These trajectories are then used, together with a procedure to construct an optimism-based upper confidence bonus, to define a policy  $\hat{\pi}_h^{(t)}$  ([Lines 6 and 7](#)), which represents the algorithm's current optimistic estimate of the optimal policy. Finally,  $\tilde{\pi}_h^{(t)}$  is defined to be a sort of uniform policy ([Line 9](#)), with respect to the current covariance matrix of features  $\Sigma_h^{(t)}$ .

**Parameter definitions.** Below, we specify the parameters  $\lambda_1, \beta$  used in the definition of the bonuses  $F_h^{(t)}$  (in [\(33\)](#) below), the parameters  $\lambda, n, T$  used in [Algorithm 1](#), as well as several other parameters used in the analysis of this section:

**Definition D.1** (Parameter definitions). Fix  $\delta \in (0, 1)$ ,  $\varepsilon_{\text{final}} \in (0, 1)$ , as well as parameters  $A, H, B, d$  of the MDP. We define the following parameters, which are used in [Algorithm 1](#) as well in the remainder of this section:

- $\lambda = C_{E.3} d \log(2THn/\delta)$ , where  $C_{E.3}$  is a constant chosen sufficiently large in [Lemma E.3](#).



- $T = d \cdot \left( \frac{C_{D.10} H^4 B^3 d A^{1/2} \log^{1/2}(HABd/(\varepsilon_{\text{final}}\delta))}{\varepsilon_{\text{final}}} \right)^{6A+2}$ , where  $C_{D.10}$  is a constant chosen sufficiently large in the proof of [Theorem D.10](#).
- $n = 3T$ .
- $\iota = \log \left( \frac{TH(\lambda d+n)}{\delta} \right)$ .
- $\lambda_1 = BH$ .
- $\varepsilon_{\text{bkup}} = \frac{\varepsilon_{\text{final}}}{2H}$ .
- $\sigma_{\text{tr}} = \frac{\varepsilon_{\text{bkup}}}{4\lambda_1}$ .
- $\varepsilon_{\text{apx}} = \frac{\varepsilon_{\text{bkup}}}{128\sqrt{2\pi}C_{D.7}\lambda_1^2HBd\sqrt{\iota}}$ , where  $C_{D.7}$  is chosen sufficiently large in the proof of [Lemma D.7](#) holds.
- $\beta = 4C_{D.7}HB\sqrt{d\iota} \cdot 5\lambda_1\sqrt{d} \cdot \left( \frac{C_{C.10}}{\varepsilon_{\text{apx}}} \right)^{2A}$ .
- $\xi = \frac{\beta}{4C_{D.7}HB\sqrt{d\iota} \cdot 2\sqrt{2\pi}\lambda_1\sqrt{d}}$ . (Note that  $\xi \geq 1$ .)

**Definition of bonuses.** Next we define the exploration bonus function  $F_g^{(t)} : \mathcal{X} \rightarrow \mathbb{R}$  used in the construction of the rewards  $\hat{r}_h^{(t,i,h)}$  in [Line 5](#). The key component in  $F_g^{(t)}$  is a linear combination of truncated linear bonuses  $F_h^{\text{tl}}(x; u, v)$ , over various values of  $u, v$ . The vectors  $u, v$  are chosen to belong to the subspaces  $\mathcal{S}_\Sigma, \mathcal{S}_\Lambda$ , for matrices  $\Sigma, \Lambda$  forming an orthogonal pair which are obtained from the covariance matrix  $\Sigma_h^{(t)}$  via a certain truncation procedure, defined below:

**Definition D.2.** Let  $\Gamma \in \mathbb{R}^{d \times d}$  be a PSD matrix, and suppose that  $\Gamma$  may be diagonalized as  $\Gamma = UDU^\top$ , for a diagonal matrix  $D$  and an orthogonal matrix  $U$ . Given  $\sigma > 0$ , we define the  $\sigma$ -truncated pair of  $\Gamma$ , denoted  $(\Sigma', \Lambda') := \text{trunc}_\sigma(\Gamma)$ , as  $\Sigma' = UD_+U^\top, \Lambda' = UD_-U^\top$ , where  $D_+, D_-$  are diagonal matrices whose entries are given by:

$$(D_+)_{ii} = \begin{cases} 1 & : D_{ii} \geq \sigma \\ 0 & : D_{ii} < \sigma \end{cases}, \quad (D_-)_{ii} = \begin{cases} 0 & : D_{ii} \geq \sigma \\ 1 & : D_{ii} < \sigma \end{cases}. \quad (32)$$

In words, for  $(\Sigma', \Lambda') = \text{trunc}_\sigma(\Gamma)$ ,  $\Sigma'$  contains the large components of  $\Gamma$  (as parametrized by  $\sigma$ ), and  $\Lambda'$  contains the small components of  $\Gamma$ . Note that the definition of  $\text{trunc}_\sigma(\Gamma)$  does not depend on the choice of diagonalization. Moreover, note that  $(\Sigma', \Lambda') = \text{trunc}_\sigma(\Gamma)$  is an orthogonal pair.

The bonus function  $F_h^{(t)}$  used in [Algorithm 1](#) is defined as follows: for each  $t \in [T], h \in [H]$ ,

$$F_h^{(t)}(x) := \lambda_1 \cdot \left( \frac{C_{C.10}}{\varepsilon_{\text{apx}}} \right)^{2A} \cdot \mathbb{E}_{u' \sim \mathcal{N}(0, \Sigma')} \mathbb{E}_{v' \sim \mathcal{N}(0, \Lambda')} \left[ F_h^{\text{tl}}(x; \beta u', v') \right] \\ + 2\sqrt{2\pi}\lambda_1\xi \cdot \mathbb{E}_{w \sim \mathcal{N}(0, \Sigma')} \left[ \max_{a \in \mathcal{A}} \langle w, \phi_h(x, a) \rangle \right] \quad (33)$$

$$\text{for } (\Sigma', \Lambda') := \text{trunc}_{\sigma_{\text{tr}}}((\beta/\lambda_1) \cdot (\Sigma_h^{(t)})^{-1/2}).$$

**Q- and V-function definitions.** For each  $h \in [H], t \in [T]$ , we make the following definitions:

$$Q_h^{(t)}(x, a) := r_h(x, a) + \mathbb{E}^{\hat{\pi}^{(t)}} \left[ \sum_{g=h+1}^H r_g(x_g, a_g) + F_g^{(t)}(x_g) \mid (x_h, a_h) = (x, a) \right] \quad (34)$$

$$V_h^{(t)}(x) := Q_h^{(t)}(x, \hat{\pi}_h^{(t)}(x)) \quad (35)$$

$$\hat{Q}_h^{(t)}(x, a) := \langle \hat{w}_h^{(t)}, \phi_h(x, a) \rangle$$

$$\hat{V}_h^{(t)}(x) := \hat{Q}_h^{(t)}(x, \hat{\pi}_h^{(t)}(x)) = \max_{a \in \mathcal{A}} \hat{Q}_h^{(t)}(x, a).$$

The functions  $Q_h^{(t)}(\cdot)$  and  $V_h^{(t)}(\cdot)$  represent the Q- and V-value functions for the policy  $\hat{\pi}^{(t)}$ , with respect to the rewards  $\hat{r}_g^{(t)}$  defined in (5). In Lemma D.1 below, we show that  $Q_h^{(t)}$  is a linear function, i.e., for some  $w_h^{(t)} \in \mathbb{R}^d$ ,  $Q_h^{(t)}(x, a) = \langle \phi_h(x, a), w_h^{(t)} \rangle$  for all  $(x, a) \in \mathcal{X} \times \mathcal{A}$ . The vector  $\hat{w}_h^{(t)}$  computed in Algorithm 1 may be viewed as an empirical approximation to  $w_h^{(t)}$ , so that  $\hat{Q}_h^{(t)}(\cdot)$  and  $\hat{V}_h^{(t)}(\cdot)$  represent empirical approximations to  $Q_h^{(t)}, V_h^{(t)}$ , respectively.

**Lemma D.1.** *For all  $t \in [T], h \in [h]$ , there is some  $w_h^{(t)} \in \mathbb{R}^d$  so that  $Q_h^{(t)}(x, a) = \langle \phi_h(x, a), w_h^{(t)} \rangle$  for all  $(x, a) \in \mathcal{X} \times \mathcal{A}$ .*

*Proof.* For all  $x, a, h, t$ , we have

$$Q_h^{(t)}(x, a) = r_h(x, a) + \mathbb{E}_{x' \sim \mathbb{P}_h(x, a)} \left[ F_{h+1}^{(t)}(x') + Q_{h+1}^{(t)}(x', \hat{\pi}_{h+1}^{(t)}(x')) \right].$$

For fixed  $t \in [T]$ , we will use reverse induction on  $h$  to establish the stated claim. The base case  $h = H$  follows since  $Q_H^{(t)}(x, a) = r_H(x, a) = \langle w_H^{(t)}, \phi_H(x, a) \rangle$  for some  $w_H \in \mathbb{R}^d$ , by assumption.

Now, assume for the inductive hypothesis that for all  $(x, a) \in \mathcal{X} \times \mathcal{A}$ ,  $Q_{h+1}^{(t)}(x, a) = \langle w_{h+1}^{(t)}, \phi_{h+1}(x, a) \rangle$  for some  $w_{h+1}^{(t)} \in \mathbb{R}^d$ . Then Lemma C.3 gives that there is some  $\theta^1 \in \mathbb{R}^d$  so that for all  $(x, a) \in \mathcal{X} \times \mathcal{A}$ ,

$$\langle \theta^1, \phi_h(x, a) \rangle = \mathbb{E}_{x' \sim \mathbb{P}_h(x, a)} \left[ \langle \phi_{h+1}(x', \hat{\pi}_{h+1}^{(t)}(x')), w_{h+1}^{(t)} \rangle \right] = \mathbb{E}_{x' \sim \mathbb{P}_h(x, a)} \left[ Q_{h+1}^{(t)}(x', \hat{\pi}_{h+1}^{(t)}(x')) \right].$$

Next, note that it is immediate from the definition (33) that, with probability 1 over  $u' \sim \mathcal{N}(0, \Sigma')$  and  $v' \sim \mathcal{N}(0, \Lambda')$ , the function  $F_{h+1}^{\text{tl}}(\cdot; u', v', \beta \Sigma', \Lambda)$  is Bellman-linear at step  $h + 1$ . Note that the function  $x \mapsto \max_{a \in \mathcal{A}} \langle w, \phi_{h+1}(x, a) \rangle$  is also Bellman-linear at step  $h + 1$ , for all  $w \in \mathbb{R}^d$ . Since the sum of Bellman-linear functions is Bellman-linear, the definition of  $F_{h+1}^{(t)}$  in (33) gives that  $F_{h+1}^{(t)}$  is Bellman-linear at step  $h + 1$ . Thus, there is some  $\theta^2 \in \mathbb{R}^2$  so that

$$\langle \phi_h(x, a), \theta^2 \rangle = \mathbb{E}_{x' \sim \mathbb{P}_h(x, a)} \left[ F_{h+1}^{(t)}(x') \right]$$

for all  $(x, a) \in \mathcal{X} \times \mathcal{A}$ .

Finally, by assumption, there is some  $\theta^3 \in \mathbb{R}^d$  so that  $r_h(x, a) = \langle \theta^3, \phi_h(x, a) \rangle$  for all  $(x, a) \in \mathcal{X} \times \mathcal{A}$ . Hence, writing  $w_h^{(t)} = \theta^1 + \theta^2 + \theta^3$ , we have that  $Q_h^{(t)}(x, a) = \langle w_h^{(t)}, \phi_h(x, a) \rangle$  for all  $(x, a) \in \mathcal{X} \times \mathcal{A}$ .  $\square$

### D.1. Boundedness of the bonuses

Next, we prove some lemmas which establish bounds on the magnitude of the bonus functions  $F_h^{(t)}(\cdot)$ , and thereby on the magnitude of the optimistic  $Q$ -functions  $Q_h^{(t)}(\cdot)$ .

**Lemma D.2.** *For a PSD matrix  $\Gamma$  and  $\sigma > 0$ , writing  $(\Sigma', \Lambda') := \text{trunc}_\sigma(\Gamma)$ , we have that  $\Sigma' \preceq \frac{1}{\sigma} \cdot \Gamma$ .*

*Proof.* We write  $\Gamma = UDU^\top$  and  $\Sigma' = UD_+U^\top$ , where  $D_+$  is as defined in (32). Certainly  $0 = (D_+)_{ii} \leq D_{ii}$  if  $D_{ii} < \sigma$ ; otherwise, we have  $1 = (D_+)_{ii} \leq \frac{1}{\sigma} \cdot D_{ii}$ , which implies that  $D_+ \preceq \frac{D}{\sigma}$ , as desired.  $\square$

**Lemma D.3.** *Let  $\Phi \in \mathcal{P}^d$  be a given polyhedron. Consider any PSD matrix  $\Sigma \in \mathbb{R}^{d \times d}$ . Then*

$$\frac{1}{\sqrt{2\pi}} \cdot \max_{\phi, \phi' \in \Phi} ((\phi - \phi')^\top \Sigma (\phi - \phi'))^{1/2} \leq \mathbb{E}_{w \sim \mathcal{N}(0, \Sigma)} \left[ \max_{\phi \in \Phi} \langle w, \phi \rangle \right] \leq \sqrt{d} \cdot \mathbb{E}_{w \sim \mathcal{N}(0, \Sigma)} \left[ \phi_w^\top \Sigma \phi_w \right]^{1/2},$$

where  $\phi_w := \arg \max_{\phi \in \Phi} \langle w, \phi \rangle$ .

*Proof.* Define  $\Psi = \{\Sigma^{1/2} \cdot \phi : \phi \in \Phi\}$ , and, for  $v \in \mathbb{R}^d$ , write  $\psi_v := \arg \max_{\psi \in \Psi} \langle v, \psi \rangle$ . Then

$$\Sigma^{1/2} \cdot \phi_{\Sigma^{1/2} v} = \Sigma^{1/2} \cdot \arg \max_{\phi \in \Phi} \langle v, \Sigma^{1/2} \phi \rangle = \arg \max_{\psi \in \Psi} \langle v, \psi \rangle = \psi_v.$$

We then have the following equalities:

$$\max_{\phi, \phi' \in \Phi} ((\phi - \phi')^\top \Sigma (\phi - \phi'))^{1/2} = \max_{\psi, \psi' \in \Psi} \|\psi - \psi'\|_2 \quad (36)$$

$$\mathbb{E}_{w \sim \mathcal{N}(0, \Sigma)} \left[ \phi_w^\top \Sigma \phi_w \right]^{1/2} = \mathbb{E}_{v \sim \mathcal{N}(0, I_d)} \left[ \phi_{\Sigma^{1/2} v}^\top \Sigma \phi_{\Sigma^{1/2} v} \right]^{1/2} = \mathbb{E}_{v \sim \mathcal{N}(0, I_d)} \left[ \|\psi_v\|_2^2 \right]^{1/2} \quad (37)$$

$$\mathbb{E}_{w \sim \mathcal{N}(0, \Sigma)} \left[ \max_{\phi \in \Phi} \langle w, \phi \rangle \right] = \mathbb{E}_{v \sim \mathcal{N}(0, I_d)} \left[ \max_{\phi \in \Phi} \langle \Sigma^{1/2} v, \phi \rangle \right] = \mathbb{E}_{v \sim \mathcal{N}(0, I_d)} \left[ \max_{\psi \in \Psi} \langle v, \psi \rangle \right]. \quad (38)$$

It therefore suffices to show that

$$\frac{1}{\sqrt{2\pi}} \cdot \max_{\psi, \psi' \in \Psi} \|\psi - \psi'\|_2 \leq \mathbb{E}_{v \sim \mathcal{N}(0, I_d)} \left[ \max_{\psi \in \Psi} \langle v, \psi \rangle \right] \leq \sqrt{d} \cdot \mathbb{E}_{v \sim \mathcal{N}(0, I_d)} \left[ \|\psi_v\|_2^2 \right]^{1/2}. \quad (39)$$

To show the first inequality in (39), pick  $\psi_0, \psi_1 \in \Psi$  maximizing  $\|\psi_0 - \psi_1\|_2$ . Then

$$\begin{aligned} \mathbb{E}_{v \sim \mathcal{N}(0, I_d)} \left[ \max_{\psi \in \Psi} \langle v, \psi \rangle \right] &\geq \mathbb{E}_{v \sim \mathcal{N}(0, I_d)} \left[ \max\{\langle v, \psi_0 \rangle, \langle v, \psi_1 \rangle\} \right] \\ &= \mathbb{E}_{v \sim \mathcal{N}(0, I_d)} \left[ \langle v, (\psi_0 + \psi_1)/2 \rangle + \max\{\langle v, (\psi_0 - \psi_1)/2 \rangle, \langle v, (\psi_1 - \psi_0)/2 \rangle\} \right] \\ &= \mathbb{E}_{v \sim \mathcal{N}(0, I_d)} \left[ \langle v, (\psi_0 - \psi_1)/2 \rangle \right] \\ &= \frac{\sqrt{2}}{2\sqrt{\pi}} \|\psi_0 - \psi_1\|_2, \end{aligned}$$

where the second equality uses that  $\mathbb{E}_{v \sim \mathcal{N}(0, I_d)} [\langle v, (\psi_0 + \psi_1)/2 \rangle] = 0$ , and the final inequality uses that  $\langle v, (\psi_0 - \psi_1)/2 \rangle \sim \mathcal{N}(0, \|(\psi_0 - \psi_1)/2\|_2^2)$ .

To show the second inequality in (39), note that

$$\begin{aligned} \mathbb{E}_{v \sim \mathcal{N}(0, I_d)} \left[ \max_{\psi \in \Psi} \langle v, \psi \rangle \right] &= \mathbb{E}_{v \sim \mathcal{N}(0, I_d)} [\langle v, \psi_v \rangle] \\ &\leq \mathbb{E}_{v \sim \mathcal{N}(0, I_d)} [\|v\|_2^2]^{1/2} \cdot \mathbb{E}_{v \sim \mathcal{N}(0, I_d)} [\|\psi_v\|_2^2]^{1/2} \\ &= \sqrt{d} \cdot \mathbb{E}_{v \sim \mathcal{N}(0, I_d)} [\|\psi_v\|_2^2]^{1/2}, \end{aligned}$$

where the inequality uses Cauchy-Schwarz.  $\square$

**Lemma D.4.** *Suppose that  $\Gamma \in \mathbb{R}^d$  is a PSD matrix, and let  $\sigma > 0$  be given. Let  $(\Sigma', \Lambda') := \text{trunc}_\sigma(\Gamma)$ . Then for all  $h \in [H], x \in \mathcal{X}, a, a' \in \mathcal{A}$ , and  $v \in \bar{\Phi}_h(x)$ , it holds that*

$$\begin{aligned} &\|\Gamma \cdot (\phi_h(x, a) - v)\|_2 + \|v - \phi_h(x, a')\|_2 \\ &\leq \|\Gamma\| \cdot \|\Sigma' \cdot (\phi_h(x, a) - v)\|_2 + \|\Lambda' \cdot (v - \phi_h(x, a'))\|_2 + \sqrt{2\pi} \cdot \mathbb{E}_{w \sim \mathcal{N}(0, \Sigma')} \left[ \max_{\bar{a} \in \mathcal{A}} \langle w, \phi_h(x, \bar{a}) \rangle \right] + 2\sigma. \end{aligned}$$

*Proof.* Note that  $\Gamma^2 \preceq \|\Gamma\|^2 \cdot \Sigma' + \sigma^2 \cdot I_d$ , so that

$$\|\Gamma \cdot (\phi_h(x, a) - v)\|_2 \leq \|\Gamma\| \cdot \|\Sigma' \cdot (\phi_h(x, a) - v)\|_2 + \sigma \cdot \|\phi_h(x, a) - v\|_2 \leq \|\Sigma' \cdot (\phi_h(x, a) - v)\|_2 + 2\sigma \quad (40)$$

where the final inequality uses that  $\max_{\phi \in \bar{\Phi}_h(x)} \|\phi\|_2 \leq 1$  (Assumption 2.1). Next, the fact that  $\Lambda' + \Sigma' = I_d$  gives that

$$\begin{aligned} \|v - \phi_h(x, a')\|_2 &\leq \|\Lambda' \cdot (v - \phi_h(x, a'))\|_2 + \max_{a_1, a_2 \in \mathcal{A}} \|\Sigma' \cdot (\phi_h(x, a_1) - \phi_h(x, a_2))\|_2 \\ &\leq \|\Lambda' \cdot (v - \phi_h(x, a'))\|_2 + \sqrt{2\pi} \cdot \mathbb{E}_{w \sim \mathcal{N}(0, \Sigma')} \left[ \max_{\bar{a} \in \mathcal{A}} \langle w, \phi_h(x, \bar{a}) \rangle \right], \quad (41) \end{aligned}$$

where the second inequality uses Lemma D.3. The conclusion of the lemma statement follows by adding Eqs. (40) and (41).  $\square$

**Lemma D.5** (Bound on bonuses in terms of  $\Sigma'$ ). *There is a sufficiently large constant  $C_{D.5}$  so that the following holds. Fix  $t \in [T], h \in [H]$ , and let  $(\Sigma', \Lambda') := \text{trunc}_{\sigma_{\text{tr}}}((\beta/\lambda_1) \cdot (\Sigma_h^{(t)})^{-1/2})$ . Then for all  $x \in \mathcal{X}$ ,*

$$|F_h^{(t)}(x)| \leq \sqrt{d} \beta \lambda_1 \cdot \left( \frac{C_{D.5}}{\varepsilon_{\text{apx}}} \right)^{2A} \cdot \mathbb{E}_{w \sim \mathcal{N}(0, \Sigma')} \left[ \max_{a \in \mathcal{A}} \langle w, \phi_h(x, a) \rangle \right].$$

*Proof.* Fix any  $u' \in \mathcal{S}_{\Sigma'}, v' \in \mathcal{S}_{\Lambda'}$ . Then

$$\begin{aligned} F_h^{\text{tl}}(x; \beta u', v') &\leq 2 \|u'\|_2 \cdot \max_{a, a' \in \mathcal{A}} \|\beta \Sigma' \cdot (\phi_h(x, a) - \phi_h(x, a'))\|_2 \\ &\leq 2 \sqrt{2\pi} \beta \|u'\|_2 \cdot \mathbb{E}_{w \sim \mathcal{N}(0, \Sigma')} \left[ \max_{a \in \mathcal{A}} \langle w, \phi_h(x, a) \rangle \right], \end{aligned}$$

where the first inequality uses [Lemma C.8](#) with  $u = \beta u'$ ,  $v = v'$  as well as the fact that  $\Sigma u' = u'$ , and the second inequality uses [Lemma D.3](#). Using the definition of  $F_h^{(t)}(\cdot)$  in [\(33\)](#), it follows that, for a sufficiently large constant  $C > 0$ ,

$$\begin{aligned} |F_h^{(t)}(x)| &\leq \lambda_1 \cdot \left( \frac{C_{C.10}}{\varepsilon_{\text{apx}}} \right)^{2A} \cdot \mathbb{E}_{u' \sim \mathcal{N}(0, \Sigma')} \left[ 4\sqrt{2\pi}\beta \|u'\|_2 \cdot \mathbb{E}_{w \sim \mathcal{N}(0, \Sigma')} \left[ \max_{a \in \mathcal{A}} \langle w, \phi_h(x, a) \rangle \right] \right] \\ &\quad + 2\sqrt{2\pi}\lambda_1 \xi \cdot \mathbb{E}_{w \sim \mathcal{N}(0, \Sigma')} \left[ \max_{a \in \mathcal{A}} \langle w, \phi_h(x, a) \rangle \right] \\ &\leq \left( 4\sqrt{2\pi}\sqrt{d}\beta\lambda_1 \cdot \left( \frac{C_{C.10}}{\varepsilon_{\text{apx}}} \right)^{2A} + 2\sqrt{2\pi}\lambda_1 \xi \right) \cdot \mathbb{E}_{w \sim \mathcal{N}(0, \Sigma')} \left[ \max_{a \in \mathcal{A}} \langle w, \phi_h(x, a) \rangle \right] \\ &\leq \sqrt{d}\beta\lambda_1 \cdot \left( \frac{C}{\varepsilon_{\text{apx}}} \right)^{2A} \cdot \mathbb{E}_{w \sim \mathcal{N}(0, \Sigma')} \left[ \max_{a \in \mathcal{A}} \langle w, \phi_h(x, a) \rangle \right], \end{aligned}$$

where the final inequality uses that  $\xi \leq \beta$  ([Definition D.1](#)).  $\square$

**Lemma D.6** (Absolute bound on bonuses). *For all  $h, t, x$ , it holds that  $|F_h^{(t)}(x)| \leq \frac{\beta}{2C_{D.7}HB\sqrt{d}}$ , and that  $\|w_h^{(t)}\|_2 \leq \frac{\beta}{C_{D.7}\sqrt{d}}$ .*

*Proof.* Fix  $t \in [T]$ ,  $h \in [H]$ , and set  $(\Sigma', \Lambda') := \text{trunc}_{\sigma_{\text{tr}}}((\beta/\lambda_1) \cdot (\Sigma_h^{(t)})^{-1/2})$ , as per [\(33\)](#). For any fixed  $u', v' \in \mathbb{R}^d$  with  $u' \in \mathcal{S}_{\Sigma'}$ ,  $v' \in \mathcal{S}_{\Lambda'}$ , we have

$$F_h^{\text{tl}}(x; \beta u', v') \leq 2\|v'\|_2 \cdot \max_{\phi, \phi' \in \Phi_h(x)} \|\Lambda' \cdot (\phi - \phi')\|_2 \leq 4\|v'\|_2,$$

where the first inequality uses [Lemma C.8](#) and the fact that  $\Lambda v' = v'$ , and the second inequality uses that  $\max_{h,x,a} \|\phi_h(x, a)\|_2 \leq 1$ . Then, by the definition of  $F_h^{(t)}$  in [\(33\)](#), we have

$$\begin{aligned} |F_h^{(t)}(x)| &\leq \lambda_1 \cdot \left( \frac{C_{C.10}}{\varepsilon_{\text{apx}}} \right)^{2A} \cdot \mathbb{E}_{v' \sim \mathcal{N}(0, \Lambda)} [4\|v'\|_2] + (2\sqrt{2\pi}\lambda_1 \xi) \cdot \mathbb{E}_{w \sim \mathcal{N}(0, \Sigma')} \left[ \max_{a \in \mathcal{A}} \langle w, \phi_h(x, a) \rangle \right] \\ &\leq 4\lambda_1 \sqrt{d} \cdot \left( \frac{C_{C.10}}{\varepsilon_{\text{apx}}} \right)^{2A} + 2\sqrt{2\pi}\lambda_1 \xi \cdot \sqrt{d} \leq \frac{\beta}{2C_{D.7}HB\sqrt{d}}, \end{aligned}$$

where the second inequality uses [Lemma D.3](#), and the final inequality is derived as follows:

$$4\lambda_1 \sqrt{d} \cdot \left( \frac{C_{C.10}}{\varepsilon_{\text{apx}}} \right)^{2A} + (2\sqrt{2\pi}\lambda_1 \xi) \cdot \sqrt{d} \leq 4\lambda_1 \sqrt{d} \cdot \left( \frac{C_{C.10}}{\varepsilon_{\text{apx}}} \right)^{2A} + \frac{\beta}{4C_{D.7}HB\sqrt{d}} \leq \frac{\beta}{2C_{D.7}HB\sqrt{d}},$$

where we have used the fact that  $2\sqrt{2\pi}\lambda_1 \sqrt{d} \cdot \xi \leq \frac{\beta}{4C_{D.7}HB\sqrt{d}}$  and that  $\frac{\beta}{4C_{D.7}HB\sqrt{d}} \geq 4\lambda_1 \sqrt{d} \cdot \left( \frac{C_{C.10}}{\varepsilon_{\text{apx}}} \right)^{2A}$  (by [Definition D.1](#)).

The definition of  $w_h^{(t)}$  (per [\(34\)](#) and [Lemma D.1](#)), combined with the fact that  $|r_h(x, a)| \leq 1$  for all  $h, x, a$  ([Assumption 2.1](#)), gives that  $|\langle w_h^{(t)}, \phi_h(x, a) \rangle| \leq H + H \cdot \max_{h,x} \{|F_h^{(t)}(x)|\}$ , from which it follows that  $w_h^{(t)} \in (H + H \cdot \max_{h,x} \{|F_h^{(t)}(x)|\}) \cdot \mathcal{B}_h$ . Since  $\max_{w \in \mathcal{B}_h} \|w\|_2 \leq B$  (again by [Assumption 2.1](#)), we have as a consequence that  $\|w_h^{(t)}\|_2 \leq \frac{\beta}{C_{D.7}\sqrt{d}}$ .  $\square$

## D.2. Linear regression lemma

We will need the following lemma which establishes a standard concentration statement regarding closeness of  $\hat{w}_h^{(t)}$  to  $w_h^{(t)}$ .

**Lemma D.7.** *There is a sufficiently large constant  $C_{D.7}$  (which is used in the definitions of  $\varepsilon_{\text{apx}}, \beta, \xi$  in [Definition D.1](#)) so that the following holds. There is an event  $\mathcal{E}$  which occurs with probability at least  $1 - \delta/2$  so that, for all  $h \in [H], t \in [T], \phi \in \mathbb{R}^d$ , it holds that*

$$|\langle \phi, \hat{w}_h^{(t)} - w_h^{(t)} \rangle| \leq \beta \cdot \left( \phi^\top (\Sigma_h^{(t)})^{-1} \phi \right)^{1/2}. \quad (42)$$

*Proof of Lemma D.7.* Fix any  $h \in [H], t \in [T]$ . Note that, by definition of  $\Sigma_h^{(t)}$  in [Line 4](#) of [Algorithm 1](#),

$$w_h^{(t)} = (\Sigma_h^{(t)})^{-1} \Sigma_h^{(t)} w_h^{(t)} = (\Sigma_h^{(t)})^{-1} \cdot \left( \lambda w_h^{(t)} + \sum_{i=1}^n \phi_h(x_h^{(t,i,h)}, a_h^{(t,i,h)}) \cdot \langle \phi_h(x_h^{(t,i,h)}, a_h^{(t,i,h)}), w_h^{(t)} \rangle \right).$$

Also recall that the definition of  $\hat{w}_h^{(t)}$  in [Line 6](#) gives

$$\hat{w}_h^{(t)} = (\Sigma_h^{(t)})^{-1} \cdot \sum_{i=1}^n \phi_h(x_h^{(t,i,h)}, a_h^{(t,i,h)}) \cdot \left( r_h^{(t,i,h)} + \sum_{g=h+1}^H \hat{r}_g^{(t,i,h)} \right).$$

Thus,

$$\begin{aligned} w_h^{(t)} - \hat{w}_h^{(t)} &= (\Sigma_h^{(t)})^{-1} \cdot \left( \lambda w_h^{(t)} + \sum_{i=1}^n \phi_h(x_h^{(t,i,h)}, a_h^{(t,i,h)}) \right. \\ &\quad \left. \cdot \left( \langle \phi_h(x_h^{(t,i,h)}, a_h^{(t,i,h)}), w_h^{(t)} \rangle - r_h^{(t,i,h)} - \sum_{g=h+1}^H \hat{r}_g^{(t,i,h)} \right) \right). \end{aligned} \quad (43)$$

For each  $i \in [n]$ , we have, by [\(34\)](#), that

$$\mathbb{E}^{\hat{\pi}^{(t)}} \left[ r_h^{(t,i,h)} + \sum_{g=h+1}^H \hat{r}_g^{(t,i,h)} \mid (x_h, a_h) = (x_h^{(t,i,h)}, a_h^{(t,i,h)}) \right] = Q_h^{(t)}(x_h^{(t,i,h)}, a_h^{(t,i,h)}) = \langle \phi_h(x_h^{(t,i,h)}, a_h^{(t,i,h)}), w_h^{(t)} \rangle.$$

(In particular, note that by definition of the sampling procedure in [Line 3](#), the trajectory  $(x_h, a_h, \dots, x_H, a_H)$  is indeed drawn from the MDP with policy  $\hat{\pi}^{(t)}$ .)

Let us write  $Y := \frac{\beta}{C_{D.7} \sqrt{d_t}}$ . The definition of  $\beta$  in [Definition D.1](#) together with [Lemma D.6](#) ensures that  $Y \geq 2H \cdot \max \left\{ 1, \max_{h,x} |F_h^{(t)}(x)| \right\} \geq \left| r_h^{(t,i,h)} + \sum_{g=h+1}^H \hat{r}_g^{(t,i,h)} \right|$  with probability 1. We also have from [Lemma D.6](#) that  $Y \geq \|w_h^{(t)}\|_2$  with probability 1.

Thus, by [Corollary E.2](#) with  $\Gamma_0 = \lambda I$ ,  $\phi_i = \phi_h(x_h^{(t,i,h)}, a_h^{(t,i,h)})$  and  $\varepsilon_i = \langle \phi_h(x_h^{(t,i,h)}, a_h^{(t,i,h)}), w_h^{(t)} \rangle - r_h^{(t,i,h)} - \sum_{g=h+1}^H \hat{r}_g^{(t,i,h)}$  (so that we may take  $\sigma = 2Y$ ), under some event  $\mathcal{E}^{(t,h)}$  which holds with



probability at least  $1 - \delta/(2TH)$ , we have

$$\begin{aligned}
 & \left\| \sum_{i=1}^n \phi_h(x_h^{(t,i,h)}, a_h^{(t,i,h)}) \cdot \left( \langle \phi_h(x_h^{(t,i,h)}, a_h^{(t,i,h)}), w_h^{(t)} \rangle - r_h^{(t,i,h)} - \sum_{g=h+1}^H \hat{r}_g^{(t,i,h)} \right) \right\|_{(\Sigma_h^{(t)})^{-1}} \\
 & \leq 8Y \cdot \sqrt{\log \left( \frac{TH \cdot \det(\Sigma_h^{(t)})^{1/2} \det(\lambda I)^{-1/2}}{\delta/2} \right)} \\
 & \leq 8Y \sqrt{d} \cdot \sqrt{\log \left( TH \cdot \frac{\lambda d + n}{\lambda d \cdot \delta/2} \right)}, \tag{44}
 \end{aligned}$$

where the final inequality follows since

$$\det(\Sigma_h^{(t)}) \leq \left( \frac{1}{d} \text{Tr} \Sigma_h^{(t)} \right)^d \leq \left( \frac{1}{d} \cdot \left( \lambda d + \sum_{i=1}^n \|\phi_h(x_h^{(t,i,h)}, a_h^{(t,i,h)})\|_2^2 \right) \right)^d \leq \left( \frac{\lambda d + n}{d} \right)^d.$$

Let us define  $\omega_h^{(t)} := w_h^{(t)} - \hat{w}_h^{(t)} - (\Sigma_h^{(t)})^{-1} \cdot \lambda w_h^{(t)}$ , so that (43) and (44) give that, under  $\mathcal{E}^{(t,h)}$ ,

$$\left\| (\Sigma_h^{(t)})^{1/2} \omega_h^{(t)} \right\|_2 \leq 8Y \sqrt{d} \cdot \sqrt{\log \left( TH \cdot \frac{\lambda d + n}{\lambda d \delta/2} \right)}. \tag{45}$$

It follows that, under  $\mathcal{E}^{(t,h)}$ , for all  $\phi \in \mathbb{R}^d$ ,

$$\begin{aligned}
 |\langle \hat{w}_h^{(t)} - w_h^{(t)}, \phi \rangle| &= |\lambda \langle (\Sigma_h^{(t)})^{-1} w_h^{(t)}, \phi \rangle| + |\langle \omega_h^{(t)}, \phi \rangle| \\
 &= |\lambda \langle (\Sigma_h^{(t)})^{-1} w_h^{(t)}, \phi \rangle| + |\langle (\Sigma_h^{(t)})^{1/2} \cdot \omega_h^{(t)}, (\Sigma_h^{(t)})^{-1/2} \cdot \phi \rangle| \\
 &\leq \left( \sqrt{\lambda} \cdot Y + 8Y \sqrt{d} \cdot \sqrt{\log \left( TH \cdot \frac{\lambda d + n}{\lambda d \delta/2} \right)} \right) \cdot \sqrt{\phi^\top \cdot (\Sigma_h^{(t)})^{-1} \cdot \phi}, \tag{46}
 \end{aligned}$$

where the final inequality uses (45), Cauchy-Schwarz, and

$$\begin{aligned}
 |\lambda \langle (\Sigma_h^{(t)})^{-1} w_h^{(t)}, \phi \rangle| &= |\langle w_h^{(t)}, \lambda (\Sigma_h^{(t)})^{-1} \phi \rangle| \\
 &\leq \|w_h^{(t)}\|_2 \cdot \sqrt{\phi^\top \cdot \lambda^2 (\Sigma_h^{(t)})^{-2} \cdot \phi} \\
 &\leq \sqrt{\lambda} \cdot Y \cdot \sqrt{\phi^\top \cdot (\Sigma_h^{(t)})^{-1} \cdot \phi},
 \end{aligned}$$

where the final inequality above uses the fact that  $\lambda \cdot (\Sigma_h^{(t)})^{-1} \preceq I_d$ . Finally, we note that (46) gives that (42) holds under  $\mathcal{E}^{(h,t)}$  since we have

$$\begin{aligned}
 \beta &= Y \cdot C_{D,7} \sqrt{dt} \geq Y \cdot \left( \sqrt{C_{E,3} d \log(2THn/\delta)} + 8 \sqrt{d \log \left( \frac{TH(\lambda d + n)}{\lambda d \delta/2} \right)} \right) \\
 &= Y \cdot \left( \sqrt{\lambda} + 8 \sqrt{d} \sqrt{\log \left( \frac{TH(\lambda d + n)}{\lambda d \delta/2} \right)} \right), \tag{47}
 \end{aligned}$$

where the above follows by the choices of  $\lambda = C_{E.3} d \log(2THn/\delta)$ ,  $\iota = \log\left(\frac{TH(\lambda d+n)}{\delta/2}\right) \geq \log(2THn/\delta)$ , and as long as  $C_{D.7}$  is chosen sufficiently large. The conclusion of the lemma follows by taking  $\mathcal{E} = \bigcap_{h \in [H], t \in [T]} \mathcal{E}^{(h,t)}$  and noting that by a union bound,  $\mathcal{E}$  occurs with probability at least  $1 - \delta/2$ .  $\square$

### D.3. Establishing optimism

In this section, we establish that that  $Q_h^{(t)}(\cdot)$  is (approximately) an upper bound on  $Q_h^*(\cdot)$ , in [Lemma D.9](#) below.

First, we prove [Lemma D.8](#) below, which shows that for states  $x$  which satisfy a certain inequality (namely, [\(48\)](#)),  $V_h^\pi(x)$  is approximately upper bounded by  $V_h^{(t)}(x)$  for any linear policy  $\pi$ . It is used to deal with one particular case in the proof of [Lemma D.9](#). Recall that for any  $\pi \in \Pi^{\text{lin}}$ , [Corollary C.4](#) guarantees the existence of a vector  $w_h^\pi \in \mathbb{R}^d$  so that for all  $(x, a) \in \mathcal{X} \times \mathcal{A}$ ,  $Q_h^\pi(x, a) = \langle w_h^\pi, \phi_h(x, a) \rangle$ .

**Lemma D.8.** *Consider an orthogonal pair of PSD matrices  $(\Sigma, \Lambda)$ , and fix  $x \in \mathcal{X}$ . For some  $\xi \geq 1$ , suppose that*

$$\max_{a_1, a_2 \in \mathcal{A}} \|\Lambda \cdot (\phi_h(x, a_1) - \phi_h(x, a_2))\|_2 \leq \xi \cdot \max_{a_1, a_2 \in \mathcal{A}} \|\Sigma \cdot (\phi_h(x, a_1) - \phi_h(x, a_2))\|_2. \quad (48)$$

Then for any  $\pi \in \Pi^{\text{lin}}$ ,

$$V_h^{(t)}(x) + 2\xi \|w_h^\pi\|_2 \cdot \max_{a_1, a_2 \in \mathcal{A}} \|\Sigma \cdot (\phi_h(x, a_1) - \phi_h(x, a_2))\|_2 \geq V_h^\pi(x) + \left( Q_h^{(t)}(x, \hat{\pi}_h^{(t)}(x)) - Q_h^\pi(x, \hat{\pi}_h^{(t)}(x)) \right).$$

*Proof.* Note that, for any actions  $a, a' \in \mathcal{A}$ , we have

$$\begin{aligned} \|\phi_h(x, a) - \phi_h(x, a')\|_2 &\leq \|\Lambda \cdot (\phi_h(x, a) - \phi_h(x, a'))\|_2 + \|\Sigma \cdot (\phi_h(x, a) - \phi_h(x, a'))\|_2 \\ &\leq 2\xi \cdot \max_{a, a' \in \mathcal{A}} \|\Sigma \cdot (\phi_h(x, a) - \phi_h(x, a'))\|_2, \end{aligned} \quad (49)$$

where the second inequality uses [\(48\)](#) and the fact that  $\xi \geq 1$ .

We compute

$$\begin{aligned} &V_h^{(t)}(x) - V_h^\pi(x) \\ &= Q_h^{(t)}(x, \hat{\pi}_h^{(t)}(x)) - Q_h^\pi(x, \pi_h(x)) \\ &= \langle w_h^\pi, \phi_h(x, \hat{\pi}_h^{(t)}(x)) - \phi_h(x, \pi_h(x)) \rangle + \left( Q_h^{(t)}(x, \hat{\pi}_h^{(t)}(x)) - Q_h^\pi(x, \hat{\pi}_h^{(t)}(x)) \right) \\ &\geq -\|w_h^\pi\|_2 \cdot \|\phi_h(x, \hat{\pi}_h^{(t)}(x)) - \phi_h(x, \pi_h(x))\|_2 + \left( Q_h^{(t)}(x, \hat{\pi}_h^{(t)}(x)) - Q_h^\pi(x, \hat{\pi}_h^{(t)}(x)) \right) \\ &\geq -2\xi \|w_h^\pi\|_2 \cdot \max_{a, a' \in \mathcal{A}} \|\Sigma \cdot (\phi_h(x, a) - \phi_h(x, a'))\|_2 + \left( Q_h^{(t)}(x, \hat{\pi}_h^{(t)}(x)) - Q_h^\pi(x, \hat{\pi}_h^{(t)}(x)) \right), \end{aligned}$$

where the final inequality uses [\(49\)](#).  $\square$

[Lemma D.9](#) establishes that the functions  $Q_h^{(t)}$  are an approximate upper bound on  $Q_h^*$ .

**Lemma D.9 (Optimism).** *Under the event  $\mathcal{E}$  of [Lemma D.7](#), it holds that, for all  $x, a, h, t$ ,*

$$\varepsilon_{\text{bkup}} \cdot (H - h) + Q_h^{(t)}(x, a) \geq Q_h^*(x, a) \quad (50)$$

$$\varepsilon_{\text{bkup}} \cdot (H + 1 - h) + V_h^{(t)}(x) + F_h^{(t)}(x) \geq V_h^*(x). \quad (51)$$

*Proof.* Condition on the event  $\mathcal{E}$ , and fix  $t \in [T]$ . We prove the statement using reverse induction on  $h$ . The base case  $h = H + 1$  is immediate since  $Q_{H+1}^{(t)}(x, a) = Q_{H+1}^*(x, a) = V_{H+1}^{(t)}(x) = V_{H+1}^*(x) = 0$  for all  $x, a, t$ .

Now suppose that the statement of the lemma holds at step  $h + 1$ . For all  $x, a$ , we have

$$Q_h^{(t)}(x, a) = r_h(x, a) + \mathbb{E}_{x' \sim \mathbb{P}_h(x, a)} \left[ F_{h+1}^{(t)}(x') + V_{h+1}^{(t)}(x') \right] \quad (52)$$

$$Q_h^*(x, a) = r_h(x, a) + \mathbb{E}_{x' \sim \mathbb{P}_h(x, a)} \left[ V_{h+1}^*(x') \right], \quad (53)$$

where (52) uses the definition of  $Q_h^{(t)}, V_h^{(t)}$  in Eqs. (34) and (35). The inductive hypothesis gives that

$$\varepsilon_{\text{bkup}} \cdot (H + 1 - h - 1) + V_{h+1}^{(t)}(x') + F_{h+1}^{(t)}(x') \geq V_{h+1}^*(x') \quad (54)$$

for all  $x' \in \mathcal{X}$ . Combining Eqs. (52) and (53), we see that

$$Q_h^{(t)}(x, a) - Q_h^*(x, a) = \mathbb{E}_{x' \sim \mathbb{P}_h(x, a)} \left[ F_{h+1}^{(t)}(x') + V_{h+1}^{(t)}(x') - V_{h+1}^*(x') \right] \geq -\varepsilon_{\text{bkup}} \cdot (H - h), \quad (55)$$

where the inequality uses (54). This verifies (50) at step  $h$ .

To establish (51) at step  $h$ , set  $(\Sigma', \Lambda') := \text{trunc}_{\sigma_{\text{tr}}}((\beta/\lambda_1) \cdot (\Sigma_h^{(t)})^{-1/2})$  (per Definition D.2), so that  $(\Sigma', \Lambda')$  is an orthogonal pair (note that this is the same choice as is made in the definition of  $F_h^{(t)}(\cdot)$  in (33)).

Let  $\xi$  be defined as per Definition D.1, so that  $\xi \geq 1$ . Now consider any  $x \in \mathcal{X}$  and  $\pi \in \Pi$ . We consider two cases based on the value of  $x$ .

**Case 1.** In the first case, we assume that

$$\max_{a, a' \in \mathcal{A}} \|\Lambda' \cdot (\phi_h(x, a) - \phi_h(x, a'))\|_2 \leq \xi \cdot \max_{a, a' \in \mathcal{A}} \|\Sigma' \cdot (\phi_h(x, a) - \phi_h(x, a'))\|_2. \quad (56)$$

Then Lemma D.8 with  $(\Sigma, \Lambda) = (\Sigma', \Lambda')$  and  $\pi = \pi^*$  gives that

$$\begin{aligned} V_h^{(t)}(x) + 2\xi \|w_h^*\|_2 \cdot \max_{a, a' \in \mathcal{A}} \|\Sigma' \cdot (\phi_h(x, a) - \phi_h(x, a'))\|_2 &\geq V_h^*(x) + \left( Q_h^{(t)}(x, \hat{\pi}_h^{(t)}(x)) - Q_h^*(x, \hat{\pi}_h^{(t)}(x)) \right) \\ &\geq V_h^*(x) - \varepsilon_{\text{bkup}} \cdot (H - h), \end{aligned}$$

where the second inequality uses (55). We may compute

$$2\xi \|w_h^*\|_2 \cdot \max_{a, a' \in \mathcal{A}} \|\Sigma' \cdot (\phi_h(x, a) - \phi_h(x, a'))\|_2 \leq 2\xi \lambda_1 \cdot \sqrt{2\pi} \mathbb{E}_{w \sim \mathcal{N}(0, \Sigma')} \left[ \max_{a \in \mathcal{A}} \langle w, \phi_h(x, a) \rangle \right] \leq F_h^{(t)}(x),$$

where the first inequality uses Lemma D.3 and Corollary C.4 as well as the fact that  $(\Sigma')^2 = \Sigma'$ , and the second inequality uses the definition of  $F_h^{(t)}(\cdot)$  in (33). It then follows that (51) holds at step  $h$ .

**Case 2.** In the second case, we assume that (56) does not hold, which implies that

$$\begin{aligned} & \max \left\{ \xi \cdot \max_{a,a' \in \mathcal{A}} \|\Sigma' \cdot (\phi_h(x, a) - \phi_h(x, a'))\|_2, \max_{a,a' \in \mathcal{A}} \|\Lambda' \cdot (\phi_h(x, a) - \phi_h(x, a'))\|_2 \right\} \\ & \leq \max_{a,a' \in \mathcal{A}} \|\Lambda' \cdot (\phi_h(x, a) - \phi_h(x, a'))\|_2 \leq \max_{a,a' \in \mathcal{A}} \|\phi_h(x, a) - \phi_h(x, a')\|_2 \leq 2. \end{aligned} \quad (57)$$

We will apply Lemma C.10 with  $(\Sigma, \Lambda) = (\Sigma', \Lambda')$ , the value of  $\beta$  set to the present value of  $\beta$  (defined in Definition D.1),  $\varepsilon = \varepsilon_{\text{apx}}$ , and  $\Phi = \bar{\Phi}_h(x)$ . We need to check that (23) holds; indeed, we have

$$\begin{aligned} & \max \left\{ \beta \cdot \max_{a,a' \in \mathcal{A}} \|\Sigma' \cdot (\phi_h(x, a) - \phi_h(x, a'))\|_2, \max_{a,a' \in \mathcal{A}} \|\Lambda' \cdot (\phi_h(x, a) - \phi_h(x, a'))\|_2 \right\} \\ & \leq 8\sqrt{2\pi}C_{D.7}\lambda_1 H B d \sqrt{l} \cdot \max \left\{ \xi \cdot \max_{a,a' \in \mathcal{A}} \|\Sigma' \cdot (\phi_h(x, a) - \phi_h(x, a'))\|_2, \max_{a,a' \in \mathcal{A}} \|\Lambda' \cdot (\phi_h(x, a) - \phi_h(x, a'))\|_2 \right\} \\ & \leq 16\sqrt{2\pi}C_{D.7}\lambda_1 H B d \sqrt{l}, \end{aligned}$$

where the first inequality uses the definition of  $\xi$  in Definition D.1 and the second inequality uses (57). Thus we may apply Lemma C.10 with  $\zeta = 16\sqrt{2\pi}C_{D.7}\lambda_1 H B d \sqrt{l}$ . For  $a, a' \in \mathcal{A}$ , let us write  $\mu_{x,h,a,a'}^*$  as shorthand for  $\mu_{\bar{\Phi}_h(x), \phi_h(x,a), \phi_h(x,a')}^*$ , which was defined in (22). Fix any  $a, a' \in \mathcal{A}$ ; then applying Lemma C.10 with  $\phi = \phi_h(x, a)$ ,  $\phi' = \phi_h(x, a')$ , we have

$$\begin{aligned} \frac{1}{\lambda_1} \cdot F_h^{(t)}(x) & \geq \left( \frac{C_{C.10}}{\varepsilon_{\text{apx}}} \right)^{2A} \cdot \mathbb{E}_{u' \sim \mathcal{N}(0, \Sigma')} \mathbb{E}_{v' \sim \mathcal{N}(0, \Lambda')} \left[ F_h^{\text{tl}}(x; \beta u', v') \right] \\ & \quad + \sqrt{2\pi} \cdot \mathbb{E}_{w \sim \mathcal{N}(0, \Sigma')} \left[ \max_{\bar{a} \in \mathcal{A}} \langle w, \phi_h(x, \bar{a}) \rangle \right] \\ & \geq \|\beta \Sigma' \cdot (\phi_h(x, a) - \mu_{x,h,a,a'}^*)\|_2 + \|\Lambda' \cdot (\mu_{x,h,a,a'}^* - \phi_h(x, a'))\|_2 - 64\sqrt{2\pi}C_{D.7}\lambda_1 H B d \sqrt{l} \cdot \varepsilon_{\text{apx}} \\ & \quad + \sqrt{2\pi} \cdot \mathbb{E}_{w \sim \mathcal{N}(0, \Sigma')} \left[ \max_{\bar{a} \in \mathcal{A}} \langle w, \phi_h(x, \bar{a}) \rangle \right] \\ & \geq \|\beta \Sigma' \cdot (\phi_h(x, a) - \mu_{x,h,a,a'}^*)\|_2 + \|\Lambda' \cdot (\mu_{x,h,a,a'}^* - \phi_h(x, a'))\|_2 - \varepsilon_{\text{bkup}} / (2\lambda_1) \\ & \quad + \sqrt{2\pi} \cdot \mathbb{E}_{w \sim \mathcal{N}(0, \Sigma')} \left[ \max_{\bar{a} \in \mathcal{A}} \langle w, \phi_h(x, \bar{a}) \rangle \right], \end{aligned} \quad (58)$$

where the first inequality follows from the definition of  $F_h^{(t)}(x)$  in (33) and the fact that  $\xi \geq 1$ , the second inequality uses Lemma C.10 with the above parameter settings, and the final inequality uses the definition of  $\varepsilon_{\text{apx}}$  in Definition D.1.

Write  $\mu^* := \mu_{x,h,\hat{\pi}_h^{(t)}(x),\pi_h^*(x)}^*$ . We now apply Lemma D.7 with  $\phi = \phi_h(x, \hat{\pi}_h^{(t)}(x)) - \mu^*$ . Then Lemma D.7 ensures that

$$\left| \langle w_h^{(t)} - \hat{w}_h^{(t)}, \phi_h(x, \hat{\pi}_h^{(t)}(x)) - \mu^* \rangle \right| \leq \beta \cdot \left\| (\Sigma_h^{(t)})^{-1/2} \cdot (\phi_h(x, \hat{\pi}_h^{(t)}(x)) - \mu^*) \right\|_2. \quad (59)$$

Moreover, since  $\mu^* \in \bar{\Phi}_h(x)$ , we have from the fact that (50) holds at step  $h$  that

$$\langle w_h^{(t)} - w_h^*, \mu^* \rangle \geq -\varepsilon_{\text{bkup}} \cdot (H - h). \quad (60)$$

Finally, we have

$$\langle w_h^*, \mu^* - \phi_h(x, \pi_h^*(x)) \rangle \geq - \|w_h^*\|_2 \cdot \|\mu^* - \phi_h(x, \pi_h^*(x))\|_2. \quad (61)$$

We may now compute

$$\begin{aligned} V_h^{(t)}(x) - V_h^*(x) &= Q_h^{(t)}(x, \hat{\pi}_h^{(t)}(x)) - Q_h^*(x, \pi_h^*(x)) \\ &\geq Q_h^{(t)}(x, \hat{\pi}_h^{(t)}(x)) - Q_h^*(x, \pi_h^*(x)) + \langle \hat{w}_h^{(t)}, \mu^* - \phi_h(x, \hat{\pi}_h^{(t)}(x)) \rangle \\ &= \langle w_h^{(t)} - w_h^*, \mu^* \rangle + \langle \hat{w}_h^{(t)} - w_h^{(t)}, \mu^* - \phi_h(x, \hat{\pi}_h^{(t)}(x)) \rangle + \langle w_h^*, \mu^* - \phi_h(x, \pi_h^*(x)) \rangle \\ &\geq -\varepsilon_{\text{bkup}} \cdot (H - h) - \beta \cdot \left\| (\Sigma_h^{(t)})^{-1/2} \cdot (\phi_h(x, \hat{\pi}_h^{(t)}(x)) - \mu^*) \right\|_2 - \|w_h^*\|_2 \cdot \|\mu^* - \phi_h(x, \pi_h^*(x))\|_2, \end{aligned} \quad (62)$$

where the first inequality uses that  $\hat{\pi}_h^{(t)}(x) = \arg \max_{a' \in \mathcal{A}} \hat{Q}_h^{(t)}(x, a')$ , and the second inequality uses [Eqs. \(59\) to \(61\)](#). Moreover,

$$\begin{aligned} &\beta \cdot \left\| (\Sigma_h^{(t)})^{-1/2} \cdot (\phi_h(x, \hat{\pi}_h^{(t)}(x)) - \mu^*) \right\|_2 + \|w_h^*\|_2 \cdot \|\mu^* - \phi_h(x, \pi_h^*(x))\|_2 \\ &\leq \beta \cdot \left\| (\Sigma_h^{(t)})^{-1/2} \cdot (\phi_h(x, \hat{\pi}_h^{(t)}(x)) - \mu^*) \right\|_2 + \lambda_1 \cdot \|\mu^* - \phi_h(x, \pi_h^*(x))\|_2 \\ &\leq \lambda_1 \cdot \left( \|(\beta/\lambda_1) \cdot \Sigma' \cdot (\phi_h(x, \hat{\pi}_h^{(t)}(x)) - \mu^*)\|_2 + \|\Lambda' \cdot (\mu^* - \phi_h(x, \pi_h^*(x)))\|_2 \right) \\ &\quad + \sqrt{2\pi} \cdot \mathbb{E}_{w \sim \mathcal{N}(0, \Sigma')} \left[ \max_{a \in \mathcal{A}} \langle w, \phi_h(x, a) \rangle \right] + 2\sigma_{\text{tr}}, \end{aligned} \quad (63)$$

where the first inequality uses that  $\|w_h^*\|_2 \leq \lambda_1$  ([Corollary C.4](#)) and the second inequality uses [Lemma D.4](#) with  $\Gamma = (\beta/\lambda_1) \cdot (\Sigma_h^{(t)})^{-1/2}$  and  $\sigma = \sigma_{\text{tr}}$ , as well as the fact that  $\|(\Sigma_h^{(t)})^{-1/2}\| \leq 1$  (which follows by the choice of  $\lambda = 1$  in [Definition D.1](#)), so that  $\|\Gamma\| \leq \beta/\lambda_1$ .

Combining [Eqs. \(58\), \(62\) and \(63\)](#) and the fact that  $\lambda_1 \geq 1$ , it follows that

$$\begin{aligned} &V_h^{(t)}(x) - V_h^*(x) \\ &\geq -\varepsilon_{\text{bkup}} \cdot (H - h) - \|\beta \Sigma' \cdot (\phi_h(x, \hat{\pi}_h^{(t)}(x)) - \mu^*)\|_2 - \lambda_1 \cdot \|\Lambda' \cdot (\mu^* - \phi_h(x, \pi_h^*(x)))\|_2 \\ &\quad - \sqrt{2\pi} \lambda_1 \cdot \mathbb{E}_{w \sim \mathcal{N}(0, \Sigma')} \left[ \max_{a \in \mathcal{A}} \langle w, \phi_h(x, a) \rangle \right] - 2\lambda_1 \sigma_{\text{tr}} \\ &\geq -\varepsilon_{\text{bkup}} \cdot (H - h) - F_h^{(t)}(x) - \varepsilon_{\text{bkup}}/2 - 2\lambda_1 \sigma_{\text{tr}} \end{aligned} \quad (64)$$

$$\geq -\varepsilon_{\text{bkup}} \cdot (H + 1 - h) - F_h^{(t)}(x), \quad (65)$$

where the final inequality follows since  $\varepsilon_{\text{bkup}}/2 + 2\lambda_1 \sigma_{\text{tr}} \leq \varepsilon_{\text{bkup}}$  by the choice of  $\sigma_{\text{tr}}$  in [Definition D.1](#). This verifies that [\(51\)](#) holds at step  $h$ , completing the proof of the lemma.  $\square$

**Theorem D.10.** *For any  $\varepsilon_{\text{final}}, \delta \in (0, 1)$ , the policy  $\hat{\pi} := \frac{1}{T} \sum_{t=1}^T \hat{\pi}^{(t)}$  produced by [Algorithm 1](#) has suboptimality bounded by  $V_1^*(x_1) - V_1^{\hat{\pi}}(x_1) \leq \varepsilon_{\text{final}}$  after the algorithm has gathered*

$$O \left( H d^2 \cdot \left( \frac{H^4 B^3 d A^{1/2} \log^{1/2}(H A B d / (\varepsilon_{\text{final}} \delta))}{\varepsilon_{\text{final}}} \right)^{12A+4} \right)$$

*samples and used  $(H B d A \log(1/(\varepsilon_{\text{final}} \delta)))^{O(A)}$  time.*

*Proof.* Using the fact that  $V_1^{\hat{\pi}^{(t)}}(x_1) = \mathbb{E}^{\hat{\pi}^{(t)}} \left[ \sum_{h=1}^H r_h(x_h, a_h) \right]$  as well as the definition of  $V_h^{(t)}$  in Eqs. (34) and (35), we see that

$$V_1^{(t)}(x_1) - V_1^{\hat{\pi}^{(t)}}(x_1) = \sum_{h=2}^H \mathbb{E}^{\hat{\pi}^{(t)}} \left[ F_h^{(t)}(x_h) \right]. \quad (66)$$

Set  $\gamma = \sqrt{d}\beta\lambda_1 \cdot \left( \frac{C_{D.5}}{\varepsilon_{\text{apx}}} \right)^{2A}$ . For each  $h \in [H], t \in [T]$ , let us write  $((\Sigma'_h)^{(t)}, (\Lambda'_h)^{(t)}) := \text{trunc}_{\sigma_{\text{tr}}}((\beta/\lambda_1) \cdot (\Sigma_h^{(t)})^{-1/2})$ . Then under the event  $\mathcal{E}$  of Lemma D.9 (which satisfies  $\Pr(\mathcal{E}) \geq 1 - \delta/2$ ),

$$\begin{aligned} V_1^*(x_1) - V_1^{\hat{\pi}^{(t)}}(x_1) &\leq H \cdot \varepsilon_{\text{bkup}} + V_1^{(t)}(x_1) + F_1^{(t)}(x_1) - V_1^{\hat{\pi}^{(t)}}(x_1) \\ &= H \cdot \varepsilon_{\text{bkup}} + \sum_{h=1}^H \mathbb{E}^{\hat{\pi}^{(t)}} \left[ F_h^{(t)}(x_h) \right] \\ &\leq H \cdot \varepsilon_{\text{bkup}} + \gamma \sum_{h=1}^H \mathbb{E}^{\hat{\pi}^{(t)}} \left[ \mathbb{E}_{w \sim \mathcal{N}(0, (\Sigma'_h)^{(t)})} \left[ \max_{a \in \mathcal{A}} \langle w, \phi_h(x_h, a) \rangle \right] \right] \\ &\leq H \cdot \varepsilon_{\text{bkup}} + \gamma \sum_{h=1}^H \mathbb{E}^{\hat{\pi}^{(t)}} \left[ \sqrt{d} \cdot \mathbb{E}_{w \sim \mathcal{N}(0, (\Sigma'_h)^{(t)})} \left[ \phi_h(x_h, \pi_{h,w}(x_h))^\top (\Sigma'_h)^{(t)} \phi_h(x_h, \pi_{h,w}(x_h)) \right]^{1/2} \right] \\ &\leq H \cdot \varepsilon_{\text{bkup}} + \gamma \sqrt{d} \sum_{h=1}^H \mathbb{E}^{\hat{\pi}^{(t)} \circ_h \tilde{\pi}^{(t)}} \left[ \phi_h(x_h, a_h)^\top (\Sigma'_h)^{(t)} \phi_h(x_h, a_h) \right]^{1/2} \\ &\leq H \cdot \varepsilon_{\text{bkup}} + \gamma \sqrt{d} \cdot \frac{\beta}{\sigma_{\text{tr}} \lambda_1} \sum_{h=1}^H \mathbb{E}^{\hat{\pi}^{(t)} \circ_h \tilde{\pi}^{(t)}} \left[ \phi_h(x_h, a_h)^\top (\Sigma_h^{(t)})^{-1} \phi_h(x_h, a_h) \right]^{1/2}, \end{aligned} \quad (67)$$

where the first inequality uses Lemma D.9, the second inequality uses Lemma D.5 together with the definition of  $\gamma$ , the third inequality uses Lemma D.3, the fourth inequality uses Jensen's inequality as well as the definition of  $\tilde{\pi}_h^{(t)}$  in Line 9 of Algorithm 1, and the final inequality uses the fact that  $((\Sigma'_h)^{(t)}, (\Lambda'_h)^{(t)}) = \text{trunc}_{\sigma_{\text{tr}}}(\frac{\beta}{\lambda_1} \cdot (\Sigma_h^{(t)})^{-1/2})$  together with Lemma D.2, which implies that  $((\Sigma'_h)^{(t)})^2 = (\Sigma'_h)^{(t)} \preceq \frac{\beta}{\sigma_{\text{tr}} \lambda_1} \cdot (\Sigma_h^{(t)})^{-1/2}$ .

For each  $h \in [H], t \in [T]$ , let us define

$$\Gamma_h^{(t)} := \mathbb{E}^{\hat{\pi}^{(t)} \circ_h \tilde{\pi}^{(t)}} \left[ \phi_h(x_h, a_h) \phi_h(x_h, a_h)^\top \right].$$

Next, for each  $t \in [T]$  and  $h \in [H]$ , we will now apply Lemma E.3 with  $P$  given by the distribution of  $\phi_h(x_h, a_h)$  for  $(x_h, a_h) \sim \mathbb{P}^{\text{unif}(\{\hat{\pi}^s \circ_h \tilde{\pi}^s\}_{s=1}^{t-1})}$ . Write

$$\Lambda_h^{(t)} = \mathbb{E}^{\text{unif}(\{\hat{\pi}^s \circ_h \tilde{\pi}^s\}_{s=1}^{t-1})} \left[ \phi_h(x_h, a_h) \phi_h(x_h, a_h)^\top \right],$$

so that for each  $i \in [n]$ ,  $\mathbb{E}[\phi_h(x_h^{(t,i,h)}, a_h^{(t,i,h)}) \phi_h(x_h^{(t,i,h)}, a_h^{(t,i,h)})^\top] = \Lambda_h^{(t)}$ . Then Lemma E.3 together with the definition of  $\Sigma_h^{(t)}$  in Line 4 of Algorithm 1 gives that, since  $\lambda \geq C_{E.3} d \log(2THn/\delta)$

(per [Definition D.1](#)),

$$\Pr\left(\frac{1}{3}(\lambda I + n\Lambda_h^{(t)}) \preceq \Sigma_h^{(t)} \preceq \frac{5}{3}(\lambda I + n\Lambda_h^{(t)})\right) \geq 1 - \delta/(2TH). \quad (68)$$

Let  $\mathcal{E}'$  denote the event that (68) holds for all  $h \in [H], t \in [T]$ , so that  $\mathcal{E}'$  occurs with probability at least  $1 - \delta/2$ . Since  $n \geq 3T$ , we have that  $\frac{n}{3}\Lambda_h^{(t)} \succeq \sum_{s=1}^t \Gamma_h^{(s)}$ , and therefore, under  $\mathcal{E}'$ , we have  $\tilde{\Sigma}_h^{(t)} := \frac{\lambda}{3}I_d + \sum_{s=1}^t \Gamma_h^{(s)} \preceq \Sigma_h^{(t)}$ . Next, we may compute, for each  $h \in [H]$ ,

$$\sum_{t=1}^T \mathbb{E}^{\hat{\pi}^{(t)} \circ_h \tilde{\pi}^{(t)}}[\phi_h(x_h, a_h)^\top (\Sigma_h^{(t)})^{-1} \phi_h(x_h, a_h)] = \sum_{t=1}^T \langle \Gamma_h^{(t)}, (\Sigma_h^{(t)})^{-1} \rangle \leq \sum_{t=1}^T \langle \Gamma_h^{(t)}, (\tilde{\Sigma}_h^{(t)})^{-1} \rangle \leq 2d \log(2T), \quad (69)$$

where the final inequality uses [Lemma E.4](#) together with the fact that  $\lambda/3 \geq 1$ . Combining (67) and (69), and using Jensen's inequality, we obtain that under  $\mathcal{E} \cap \mathcal{E}'$  (which occurs with probability at least  $1 - \delta$ ), for sufficiently large constants  $C, C'$ ,

$$\begin{aligned} \sum_{t=1}^T (V_1^*(x_1) - V_1^{\hat{\pi}^{(t)}}(x_1)) &\leq TH\varepsilon_{\text{bkup}} + \frac{H\gamma\beta\sqrt{d}}{\sigma_{\text{tr}}\lambda_1} \cdot \sqrt{2Td \log(2T)} \\ &\leq TH\varepsilon_{\text{bkup}} + \frac{4Hd\beta^2\lambda_1}{\varepsilon_{\text{bkup}}} \cdot \left(\frac{C_{D.5}}{\varepsilon_{\text{apx}}}\right)^{2A} \cdot \sqrt{2Td \log(2T)} \\ &\leq TH\varepsilon_{\text{bkup}} + \frac{CH^6B^5d^3\iota}{\varepsilon_{\text{bkup}}} \cdot \left(\frac{C_{C.10}}{\varepsilon_{\text{apx}}}\right)^{4A} \cdot \left(\frac{C_{D.5}}{\varepsilon_{\text{apx}}}\right)^{2A} \cdot \sqrt{2Td \log(2T)} \\ &\leq TH\varepsilon_{\text{bkup}} + \frac{CH^6B^5d^3\iota}{\varepsilon_{\text{bkup}}} \cdot \left(\frac{C'H^3B^3d\sqrt{\iota}}{\varepsilon_{\text{bkup}}}\right)^{6A} \\ &\leq T\varepsilon_{\text{final}}/2 + CH^3B^2d^2\sqrt{\iota} \left(\frac{2C'H^4B^3d\sqrt{\iota}}{\varepsilon_{\text{final}}}\right)^{6A+1}, \end{aligned} \quad (70)$$

where the second inequality uses the definition of  $\gamma, \sigma_{\text{tr}}$ , and the remaining inequalities use the definitions of  $\varepsilon_{\text{apx}}, \beta, \lambda_1, \varepsilon_{\text{final}}$  in [Definition D.1](#). Next, we have the following claim:

**Claim D.11.** *The choice of  $T, \iota$  (in [Definition D.1](#)) ensures that  $T \geq d \cdot \left(\frac{C_{D.10}^{1/2}H^4B^3d\sqrt{\iota}}{\varepsilon_{\text{final}}}\right)^{6A+2}$ .*

[Claim D.11](#) is proved at the end of the section. By [Claim D.11](#), as long as the constant  $C_{D.10}$  is chosen sufficiently large, the expression in (70) is bounded above by  $T\varepsilon_{\text{final}}$ . Hence the policy  $\hat{\pi} := \frac{1}{T} \sum_{t=1}^T \hat{\pi}^{(t)}$  satisfies  $V_1^*(x_1) - V_1^{\hat{\pi}}(x_1) \leq \varepsilon_{\text{final}}$ .

Altogether, [Algorithm 1](#) collects  $THn \leq O(HT^2) \leq O\left(Hd^2 \cdot \left(\frac{H^4B^3dA^{1/2} \log^{1/2}(HABd/(\varepsilon_{\text{final}}\delta))}{\varepsilon_{\text{final}}}\right)^{12A+4}\right)$  trajectories.

Finally, we analyze the computation time of [Algorithm 1](#). It is straightforward to see that each step of the algorithm can be implemented in  $\text{poly}(T)$  time, with the exception of the computation of  $F_g^{(t)}(x_g^{(t,i,h)})$  in (5), which requires integrating a nonlinear function over a Gaussian (see (33)). To handle this difficulty, we discuss a slight modification of our algorithm and its analysis which



proceeds by approximating this integral. To efficiently approximate the result of each such integration, we use  $O(\log(T/\delta)/\varepsilon_{\text{apx}}^2)$  draws from the appropriate Gaussian to approximate each integral to within  $\varepsilon_{\text{apx}}$ . The error term of  $\varepsilon_{\text{apx}}$  that arises in this approximation leads to an additional error term of  $\varepsilon_{\text{apx}}$  in (42) (Lemma D.7), which in turn leads to an additional term of  $-\varepsilon_{\text{apx}}$  in (64). By an appropriate choice of the constants in Definition D.1, we can ensure that  $\varepsilon_{\text{apx}} + \varepsilon_{\text{bkup}}/2 + 2\lambda_1\sigma_{\text{tr}} \leq \varepsilon_{\text{bkup}}$ , which ensures that (65) still holds. The remaining details of the proof remain unchanged.  $\square$

*Proof of Claim D.11.* Note that, since  $n = 3T$  and by the definition of  $T$  in Definition D.1, for sufficiently large constants  $C, C'$  (independent of  $C_{D.10}$ ), we have  $\iota \leq C \cdot (\log(HBd/\delta) + \log(T)) \leq C' \cdot A \log\left(\frac{C_{D.10}HBdA}{\varepsilon_{\text{final}}\delta}\right)$ . Then

$$d \cdot \left( \frac{C_{D.10}^{1/2} H^4 B^3 d \sqrt{\iota}}{\varepsilon_{\text{final}}} \right)^{6A+2} \leq d \cdot \left( \frac{\sqrt{C'} \log^{1/2}(C_{D.10}) \cdot C_{D.10}^{1/2} H^4 B^3 d A^{1/2} \log^{1/2}(HBdA/(\varepsilon_{\text{final}}\delta))}{\varepsilon_{\text{final}}} \right)^{6A+2}.$$

As long as  $C_{D.10}$  is chosen sufficiently large so that  $\sqrt{C'} \log^{1/2}(C_{D.10}) C_{D.10}^{1/2} \leq C_{D.10}$ , the right-hand side of the above expression is bounded above by  $T$ , as desired.  $\square$

## Appendix E. Useful lemmas

### E.1. Concentration

**Lemma E.1** (Concentration for self-normalized process; e.g., Theorem D.3 of (Jin et al., 2020)). *Fix  $n \in \mathbb{N}$  and let  $\varepsilon_1, \dots, \varepsilon_n$  be random variables which are adapted to a filtration  $(\mathcal{F}_i)_{0 \leq i \leq n}$ . Suppose that for each  $i \in [n]$ ,  $\mathbb{E}[\varepsilon_i | \mathcal{F}_{i-1}] = 0$  and  $\mathbb{E}[e^{\lambda \varepsilon_i} | \mathcal{F}_{i-1}] \leq e^{\lambda^2 \sigma^2 / 2}$ . Suppose that  $\phi_1, \dots, \phi_n$  is a sequence which is predictable with respect to  $(\mathcal{F}_i)_{0 \leq i \leq n}$ , i.e.,  $\phi_i$  is measurable with respect to  $\mathcal{F}_{i-1}$  for all  $i \in [n]$ . Suppose that  $\Gamma_0 \in \mathbb{R}^{d \times d}$  is positive definite, and let  $\Gamma_i = \Gamma_0 + \sum_{j=1}^i \phi_j \phi_j^\top$ . Then for any  $\delta > 0$ , with probability at least  $1 - \delta$ ,*

$$\left\| \sum_{i=1}^n \phi_i \varepsilon_i \right\|_{\Gamma_i^{-1}}^2 \leq 2\sigma^2 \log \left( \frac{\det(\Gamma_i)^{1/2} \det(\Gamma_0)^{-1/2}}{\delta} \right).$$

In the special case that the random variables  $\varepsilon_i$  are i.i.d., we obtain the following:

**Corollary E.2** (Concentration for self-normalized process; i.i.d. data). *Fix  $n \in \mathbb{N}$ , and let  $\varepsilon_1, \dots, \varepsilon_n$  be independent real-valued random variables, so that for each  $i \in [n]$ ,  $\mathbb{E}[\varepsilon_i] = 0$  and  $\mathbb{E}[e^{\lambda \varepsilon_i}] \leq e^{\lambda^2 \sigma^2 / 2}$  for all  $\lambda \in \mathbb{R}$ . Let  $\phi_1, \dots, \phi_n \in \mathbb{R}^d$  be given. Suppose that  $\Gamma_0 \in \mathbb{R}^{d \times d}$  is positive definite, and let  $\Gamma_i = \Gamma_0 + \sum_{j=1}^i \phi_j \phi_j^\top$ . Then for any  $\delta > 0$ , with probability at least  $1 - \delta$ ,*

$$\left\| \sum_{i=1}^n \phi_i \varepsilon_i \right\|_{\Gamma_i^{-1}}^2 \leq 2\sigma^2 \log \left( \frac{\det(\Gamma_i)^{1/2} \det(\Gamma_0)^{-1/2}}{\delta} \right).$$

**Lemma E.3** (Lemma 39 of (Zanette et al., 2021)). *There is a constant  $C_{E.3} > 0$  so that the following holds. Suppose  $P$  is a distribution supported on the unit Euclidean ball in  $\mathbb{R}^d$ . Write  $\Sigma = \mathbb{E}_{\phi \sim P}[\phi \phi^\top]$ . Suppose  $n \in \mathbb{N}, \delta > 0, \lambda$  are given so that  $\lambda \geq C_{E.3} d \log(n/\delta)$ . Suppose  $\phi_1, \dots, \phi_n \sim P$  are i.i.d. Then*

$$\Pr \left( \frac{1}{3}(\lambda I + n\Sigma) \preceq \lambda I + \sum_{i=1}^n \phi_i \phi_i^\top \preceq \frac{5}{3}(\lambda I + n\Sigma) \right) \geq 1 - \delta.$$

## E.2. Elliptical potential

The below lemma generalizes the elliptical potential lemma.

**Lemma E.4.** *Consider any sequence  $\Gamma_1, \dots, \Gamma_T \in \mathbb{R}^{d \times d}$  of PSD matrices, and suppose that  $\text{Tr} \Gamma_t \leq 1$  for all  $t \in [T]$ . Define  $\Sigma_t = \lambda I + \sum_{s=1}^t \Gamma_s$ , for some  $\lambda \geq 1$ . Then*

$$\sum_{t=1}^T \langle \Gamma_t, \Sigma_{t-1}^{-1} \rangle \leq 2d \log(2T).$$

*Proof.* Write  $\Sigma_0 = \lambda I$ . Since  $\lambda \geq 1$ , we have  $\Sigma_{t-1}^{-1} \preceq I_d$  and therefore  $\langle \Gamma_t, \Sigma_{t-1}^{-1} \rangle \leq \text{Tr} \Gamma_t \leq 1$  for all  $t \in [T]$ . Thus  $\langle \Gamma_t, \Sigma_{t-1}^{-1} \rangle \leq 2 \log(1 + \langle \Gamma_t, \Sigma_{t-1}^{-1} \rangle)$ . Moreover, for each  $t \in [T]$ ,

$$\begin{aligned} \det(\Sigma_t) &= \det(\Sigma_{t-1} + \Gamma_t) \\ &= \det(\Sigma_{t-1}) \cdot \det\left(I + \Sigma_{t-1}^{-1/2} \Gamma_t \Sigma_{t-1}^{-1/2}\right) \\ &\geq \det(\Sigma_{t-1}) \cdot (1 + \text{Tr}(\Sigma_{t-1}^{-1/2} \Gamma_t \Sigma_{t-1}^{-1/2})) = \det(\Sigma_{t-1}) \cdot (1 + \langle \Gamma_t, \Sigma_{t-1}^{-1} \rangle), \end{aligned}$$

where the inequality uses that for any real numbers  $\sigma_1, \dots, \sigma_d \geq 0$ ,  $\prod_{i=1}^d (1 + \sigma_i) \geq 1 + \sum_{i=1}^d \sigma_i$ , which implies that  $\det(I + A) \geq 1 + \text{Tr}(A)$  for any PSD matrix  $A$  (by taking  $\sigma_1, \dots, \sigma_d$  to be the eigenvalues of  $A$ ). Telescoping the above display, we obtain that

$$\frac{1}{2} \sum_{t=1}^T \langle \Gamma_t, \Sigma_{t-1}^{-1} \rangle \leq \sum_{t=1}^T \log(1 + \langle \Gamma_t, \Sigma_{t-1}^{-1} \rangle) \leq \log\left(\frac{\det(\Sigma_T)}{\det(\lambda I)}\right) \leq \log\left(\frac{(\frac{1}{d} \text{Tr}(\Sigma_T))^d}{\det(\lambda I)}\right) \leq d \log\left(\frac{\lambda d + T}{\lambda d}\right) \leq d \log(2T). \quad \square$$

## E.3. Performance difference lemma

**Lemma E.5** (Performance difference lemma (Kakade and Langford, 2002)). *For any MDP  $M$ , policies  $\pi, \pi' \in \Pi$ , it holds that*

$$\mathbb{E}^{M, \pi} \left[ \sum_{h=1}^H r_h(x_h, a_h) \right] - \mathbb{E}^{M, \pi'} \left[ \sum_{h=1}^H r_h(x_h, a_h) \right] = \sum_{h=1}^H \mathbb{E}^{M, \pi'} [V_h^{M, \pi}(x_h) - Q_h^{M, \pi}(x_h, a_h)].$$

## E.4. Linear algebraic lemmas

**Lemma E.6.** *Two symmetric matrices  $\Sigma, \Gamma \in \mathbb{R}^{d \times d}$  commute if and only if their eigenspaces coincide.*

## Acknowledgements

We thank Dylan Foster and Sham Kakade for bringing this problem to our attention and for helpful discussions.

## Appendix F. Proofs for Appendix C

### F.1. Functions which are not Bellman-linear

**Proposition F.1.** *There is an MDP  $M = (H, \mathcal{X}, \mathcal{A}, (P_h)_h, (r_h)_h, d_1)$  with horizon  $H = 2$  satisfying linear Bellman completeness with respect to some feature mappings  $\phi_h$  in dimension  $d = 1$ , together with a vector  $w^* \in \mathbb{R}^d$  so that the function*

$$V_h(x) := \max_{a \in \mathcal{A}} \min \{ \langle w^*, \phi_h(x, a) \rangle, H \}$$

*is not Bellman-linear at step 2 (per Definition 3.1).*

*Proof.* We set  $H = 2, d = 1, \mathcal{X} = \{\mathfrak{s}_1, \mathfrak{s}_{2,0}, \mathfrak{s}_{2,1}\}$ , and  $\mathcal{A} = \{0, 1\}$ , and define:

$$\begin{aligned} \phi_1(\mathfrak{s}_1, 0) &= 1, \phi_1(\mathfrak{s}_1, 1) = 2, \\ \phi_2(\mathfrak{s}_{2,0}, 0) &= H, \phi_2(\mathfrak{s}_{2,0}, 1) = -H, \quad \phi_2(\mathfrak{s}_{2,1}, 0) = 2H, \phi_2(\mathfrak{s}_{2,1}, 1) = -2H. \end{aligned}$$

The transitions are defined as follows:  $(\mathfrak{s}_1, 0)$  transitions to  $\mathfrak{s}_{2,0}$ , and  $(\mathfrak{s}_1, 1)$  transition to  $\mathfrak{s}_{2,1}$ . All rewards are 0. By defining  $\mathcal{T}_1 w := H \cdot |w|$ , it is readily seen that  $M$  is linear Bellman complete with respect to the above feature mappings.

Let us choose  $w^* = 1$ . Then Bellman-linearity of  $V_2(x)$  would require that there is some  $\bar{w} \in \mathbb{R}$  so that

$$\phi_1(\mathfrak{s}_1, 0) \cdot \bar{w} = H, \quad \phi_1(\mathfrak{s}_1, 1) \cdot \bar{w} = H.$$

But since  $\phi_1(\mathfrak{s}_1, 1) = 2$  and  $\phi_1(\mathfrak{s}_1, 0) = 1$ , there is no such  $\bar{w}$ .  $\square$

**Proposition F.1** provides an example for which a function  $\hat{Q}_h^{(t)}$  as defined in (3), with  $B_h^{(t)} \equiv 0$ , is so that  $\max_{a \in \mathcal{A}} \hat{Q}_h^{(t)}(\cdot, a)$  is not Bellman-linear. We have considered that case that  $B_h^{(t)}$  is 0 for simplicity. Note that in an instantiation of LSVI-UCB,  $B_h^{(t)}$  would not be identically 0, as it is defined by (4). Nevertheless, the example in **Proposition F.1** can be readily modified to account for nonzero quadratic bonus functions  $B_h^{(t)}$ . In a similar manner, by rescaling the features we may modify the example to ensure that  $\|\phi_h(x, a)\|_2 \leq 1$  for all  $h, x, a$  (so as to satisfy **Assumption 2.1**).

**Proposition F.2.** *There is an MDP  $M = (H, \mathcal{X}, \mathcal{A}, (P_h)_h, (r_h)_h, d_1)$  with horizon  $H = 2$  satisfying linear Bellman completeness with respect to some feature mappings  $\phi_h$  in dimension  $d = 1$ , so that the function*

$$F_h(x) := \max_{a \in \mathcal{A}} \|\phi_h(x, a)\|_2$$

*is not Bellman-linear at step 2.*

*Proof.* We set  $H = 2, d = 1, \mathcal{X} = \{\mathfrak{s}_1, \mathfrak{s}_{2,0}, \mathfrak{s}_{2,1}, \mathfrak{s}_{2,2}\}$ ,  $\mathcal{A} = \{0, 1\}$ , and define:

$$\begin{aligned} \phi_1(\mathfrak{s}_1, 0) &= 1, \phi_1(\mathfrak{s}_1, 1) = 1, \\ \phi_2(\mathfrak{s}_{2,0}, 0) &= 1, \phi_2(\mathfrak{s}_{2,0}, 1) = -1, \quad \phi_2(\mathfrak{s}_{2,1}, 0) = 2, \phi_2(\mathfrak{s}_{2,1}, 1) = 0, \quad \phi_2(\mathfrak{s}_{2,2}, 0) = -2, \phi_2(\mathfrak{s}_{2,2}, 1) = 0. \end{aligned}$$

The transition are defined as follows:  $(\mathfrak{s}_1, 0)$  transitions to  $\mathfrak{s}_{2,0}$ , and  $(\mathfrak{s}_1, 1)$  transitions to each of  $\mathfrak{s}_{2,1}, \mathfrak{s}_{2,2}$  with probability  $1/2$ . All rewards are 0. By defining  $\mathcal{T}_1 w := |w|$ , it is readily seen that  $M$  is linear Bellman complete with respect to the above feature mappings.

But Bellman-linearity of  $F_2(x)$  would require that there is some  $\bar{w} \in \mathbb{R}$  so that

$$\phi_1(\mathfrak{s}_1, 0) \cdot \bar{w} = 1, \quad \phi_1(\mathfrak{s}_1, 1) \cdot \bar{w} = 2.$$

No such  $\bar{w}$  exists since  $\phi_1(\mathfrak{s}_1, 0) = \phi_1(\mathfrak{s}_1, 1)$ .  $\square$

**Proposition F.2** establishes that, for  $B_h^{\text{quad}}(x, a; \Sigma)$  defined as in (4), the mapping  $\max_{a \in \mathcal{A}} B_h^{\text{quad}}(\cdot, a; I_d)$  is, in general, not Bellman-linear.

## F.2. Perturbed linear policies

We begin by defining a perturbed version of linear policies.

**Definition F.1** (Perturbed linear policies). For  $\sigma > 0$ ,  $h \in [H]$  and  $w \in \mathbb{R}^d$ , define  $\pi_{h,w,\sigma} : \mathcal{X} \rightarrow \Delta(\mathcal{A})$  by

$$\pi_{h,w,\sigma}(x)(a) = \Pr_{\theta \sim \mathcal{N}(w, \sigma^2 \cdot I_d)}(a \in \mathcal{A}_{h,\theta}(x)).$$

In words, to draw an action  $a \sim \pi_{h,w,\sigma}(x)$ , we draw  $\theta \sim \mathcal{N}(w, \sigma^2 \cdot I_d)$  and then play  $\arg \max_{a' \in \mathcal{A}} \langle \phi_h(x, a'), \theta \rangle$ . Given  $\sigma > 0$ , we denote the set of all  $\pi_{h,w,\sigma'}$ , where  $w \in \mathbb{R}^d$ ,  $\sigma' > 0$  satisfy  $\sigma'/\|w\|_2 \geq \sigma$ , by  $\Pi_h^{\text{Plin},\sigma}$ , and  $\Pi^{\text{Plin},\sigma} := \prod_{h=1}^H \Pi_h^{\text{Plin},\sigma}$ . As a matter of convention we write  $\Pi_h^{\text{Plin},0} := \Pi_h^{\text{lin}}$  and  $\Pi^{\text{Plin},0} = \Pi^{\text{lin}}$ . Moreover, we write  $\Pi_h^{\text{Plin}} := \bigcup_{\sigma \geq 0} \Pi_h^{\text{Plin},\sigma}$  and  $\Pi^{\text{Plin}} = \bigcup_{\sigma \geq 0} \Pi^{\text{Plin},\sigma}$ .

Note that, for any  $c > 0$ ,  $\pi_{h,cw,\sigma} = \pi_{h,w,\sigma/c}$ . Moreover, note that tie-breaking is not an issue for perturbed linear policies, since the measure of all  $\theta$  so that  $|\mathcal{A}_{h,\theta}(x)| > 1$  is 0. The following lemma provides another interpretation of linear policies, as a limit of perturbed linear policies:

**Lemma F.3.** For any  $w \in \mathbb{R}^d$ ,  $h \in [H]$ ,  $x \in \mathcal{X}$ , we have

$$\pi_{h,w}(a|x) = \lim_{\sigma \rightarrow 0^+} \pi_{h,w,\sigma}(a|x) = \lim_{\sigma \rightarrow 0^+} \Pr_{\theta \sim \mathcal{N}(w, \sigma^2 \cdot I_d)}(a \in \mathcal{A}_{h,\theta}(x)). \quad (71)$$

*Proof.* It is straightforward to see that for any  $a \notin \mathcal{A}_{h,w}(x)$ ,  $\lim_{\sigma \rightarrow 0^+} \Pr_{\theta \sim \mathcal{N}(w, \sigma^2 \cdot I_d)}(a \in \mathcal{A}_{h,\theta}(x)) = 0$ . This verifies (71) for  $a \notin \mathcal{A}_{h,w}(x)$ . Moreover, it also implies that for each  $a \in \mathcal{A}_{h,w}(x)$ , we have

$$\lim_{\sigma \rightarrow 0^+} \left| \Pr_{\theta \sim \mathcal{N}(w, \sigma^2 \cdot I_d)}(a \in \mathcal{A}_{h,\theta}(x)) - \Pr_{\theta \sim \mathcal{N}(w, \sigma^2 \cdot I_d)} \left( \langle \theta, \phi_h(x, a) \rangle > \max_{a' \in \mathcal{A}_{h,w}(x) \setminus \{a\}} \langle \theta, \phi_h(x, a') \rangle \right) \right| = 0.$$

Since  $\langle w, \phi_h(x, a) \rangle = \langle w, \phi_h(x, a') \rangle$  for all  $a, a' \in \mathcal{A}_{h,w}(x)$ , we have that

$$\begin{aligned} & \Pr_{\theta \sim \mathcal{N}(w, \sigma^2 \cdot I_d)} \left( \langle \theta, \phi_h(x, a) \rangle > \max_{a' \in \mathcal{A}_{h,w}(x) \setminus \{a\}} \langle \theta, \phi_h(x, a') \rangle \right) \\ &= \Pr_{\theta \sim \mathcal{N}(0, \sigma^2 \cdot I_d)} \left( \langle \theta, \phi_h(x, a) \rangle > \max_{a' \in \mathcal{A}_{h,w}(x) \setminus \{a\}} \langle \theta, \phi_h(x, a') \rangle \right) \\ &= \Pr_{\theta \sim \nu_d} \left( \langle \theta, \phi_h(x, a) \rangle > \max_{a' \in \mathcal{A}_{h,w}(x) \setminus \{a\}} \langle \theta, \phi_h(x, a') \rangle \right) = \nu_d(\mathcal{G}_{h,w}(x, a)), \end{aligned}$$

where the second equality follows by rescaling. Combining the two displays above gives the result.  $\square$

For a randomized policy  $\pi$  and  $h \in [H]$ , we use the shorthand

$$\phi_h(x, \pi_h(x)) := \mathbb{E}_{a \sim \pi_h(x)}[\phi_h(x, a)].$$

Then **Lemma F.3** yields that, for any  $h \in [H]$ ,  $w \in \mathbb{R}^d$ ,  $f : \mathcal{X} \times \mathcal{A} \rightarrow \mathbb{R}$ , we have

$$\phi_h(x, \pi_{h,w}(x)) = \lim_{\sigma \rightarrow 0^+} \phi_h(x, \pi_{h,w,\sigma}(x)), \quad f_h(x, \pi_{h,w}(x)) = \lim_{\sigma \rightarrow 0^+} f(x, \pi_{h,w,\sigma}(x)). \quad (72)$$