
Learning in Observable POMDPs, without Computationally Intractable Oracles

Noah Golowich
MIT
nzg@mit.edu

Ankur Moitra
MIT
moitra@mit.edu

Dhruv Rohatgi
MIT
drohatgi@mit.edu

Abstract

Much of reinforcement learning theory is built on top of oracles that are computationally hard to implement. Specifically for learning near-optimal policies in Partially Observable Markov Decision Processes (POMDPs), existing algorithms either need to make strong assumptions about the model dynamics (e.g. deterministic transitions) or assume access to an oracle for solving a hard optimistic planning or estimation problem as a subroutine. In this work we develop the first oracle-free learning algorithm for POMDPs under reasonable assumptions. Specifically, we give a quasipolynomial-time end-to-end algorithm for learning in “observable” POMDPs, where observability is the assumption that well-separated distributions over states induce well-separated distributions over observations. Our techniques circumvent the more traditional approach of using the principle of optimism under uncertainty to promote exploration, and instead give a novel application of barycentric spanners to constructing policy covers.

1 Introduction

Markov Decision Processes (MDPs) are a ubiquitous model in reinforcement learning that aim to capture sequential decision-making problems in a variety of applications spanning robotics to healthcare. However, modelling a problem with an MDP makes the often-unrealistic assumption that the agent has perfect knowledge about the state of the world. Partially Observable Markov Decision Processes (POMDPs) are a broad generalization of MDPs which capture an agent’s inherent uncertainty about the state: while there is still an underlying state that updates according to the agent’s actions, the agent never directly observes the state, but instead receives samples from a state-dependent observation distribution. The greater generality afforded by partial observability is crucial to applications in game theory [BS18], healthcare [Hau00, HF00b], market design [WME⁺22], and robotics [CKK96].

Unfortunately, this greater generality comes with steep statistical and computational costs. There are well-known statistical lower bounds [JKKL20, KAL16], which show that in the worst case, it is statistically intractable to find a near-optimal policy for a POMDP given the ability to play policies on it (the *learning* problem), even given unlimited computation. Furthermore, there are worst-case computational lower bounds [PT87, Lit94, BDRS96, LGM01, VLB12], which establish that it is computationally intractable to find a near-optimal policy even when given the exact parameters of the model (the substantially simpler *planning* problem).

Nevertheless there is a sizeable literature devoted to overcoming the *statistical* intractability of the learning problem by restricting to natural subclasses of POMDPs [KAL16, GDB16, ALA16, JKKL20, XCGZ21, KECM21a, KECM21b, LCSJ22]. There are far fewer works attempting to overcome *computational* intractability, and all make severe restrictions on either the model dynamics [JKKL20, KAL16] or the structure of the uncertainty [BDRS96, KECM21a]. The standard practice is to simply sidestep computational issues by assuming access to strong oracles such as ones that

solve *Optimistic Planning* (given a constrained, non-convex set of POMDPs, find the maximum value achievable by any policy on any POMDP in the set) [JKKL20] or *Optimistic Maximum Likelihood Estimation* (given a set of action/observation sample trajectories, find a POMDP which obtains maximum value conditioned on approximately maximizing the likelihood of seeing those trajectories) [LCSJ22]. Unsurprisingly, these oracles are computationally intractable to implement. Is there any hope for giving computationally efficient, oracle-free learning algorithms for POMDPs under reasonable assumptions? The naïve approach would require exponential time, and thus even a quasi-polynomial time algorithm would represent a dramatic improvement.

A necessary first step towards solving the learning problem is having a computationally efficient planning algorithm. Few such algorithms have provable guarantees under reasonable model assumptions, but recently it was shown [GMR22] that there is a quasipolynomial-time planning algorithm for POMDPs which satisfy an *observability* assumption. Let $H \in \mathbb{N}$ be the horizon length of the POMDP, and for each state s and step $h \in [H]$, let $\mathbb{O}_h(\cdot|s)$ denote the observation distribution at state s and step h . Then observability is defined as follows:

Assumption 1.1 ([EDKM07, GMR22]). *Let $\gamma > 0$. For $h \in [H]$, let \mathbb{O}_h be the matrix with columns $\mathbb{O}_h(\cdot|s)$, indexed by states s . We say that the matrix \mathbb{O}_h satisfies γ -observability if for each h , for any distributions b, b' over states, $\|\mathbb{O}_h b - \mathbb{O}_h b'\|_1 \geq \gamma \|b - b'\|_1$. A POMDP satisfies (one-step) γ -observability if all H of its observation matrices do.*

Compared to previous assumptions enabling computationally efficient planning, observability is much milder because it makes no assumptions about the dynamics of the POMDP, and it allows for natural observation models such as noisy or lossy sensors [GMR22]. It is known that statistically efficient learning is possible under somewhat weaker assumptions than observability [JKKL20], however these works rely on solving a planning problem that is computationally intractable. This raises the question: *Is observability enough to remedy both the computational and statistical woes of learning POMDPs?* In particular, can we get not only efficient planning but efficient learning too?

Overview of results. This work provides an affirmative answer to the questions above: we give an algorithm (BaSeCAMP, Algorithm 3) with quasi-polynomial time (and sample) complexity for learning near-optimal policies in observable POMDPs – see Theorem 3.1. While this falls just short of polynomial time, it turns out to be optimal in the sense that even for observable POMDPs there is a quasi-polynomial time lower bound for the (simpler) planning problem under standard complexity assumptions [GMR22]. A key innovation of our approach is an alternative technique to encourage exploration: whereas essentially all previous approaches for partially observable RL used the principle of *optimism under uncertainty* to encourage the algorithm to visit states [JKKL20, LCSJ22], we introduce a new framework based on the use of *barycentric spanners* [AK08] and *policy covers* [DKJ⁺19]. While each of these tools has previously been used in the broader RL literature to promote exploration (e.g. [LS17, FKQR21, DKJ⁺19, AHKS20, DGZ22, JKSJ20, MHKL20]), they have not been used specifically in the study of POMDPs with imperfect observations,¹ and indeed our usage of them differs substantially from past instances.

The starting point for our approach is a result of [GMR22] (restated as Theorem 2.1) which implies that the dynamics of an observable POMDP \mathcal{P} may be approximated by those of a *Markov decision process* \mathcal{M} with a quasi-polynomial number of states. If we knew the transitions of \mathcal{M} , then we could simply use dynamic programming to find an optimal policy for \mathcal{M} , which would be guaranteed to be a near-optimal policy for \mathcal{P} . Instead, we must *learn* the transitions of \mathcal{M} , for which it is necessary to explore all (reachable) states of the underlying POMDP \mathcal{P} . A naïve approach to encourage exploration is to learn the transitions of \mathcal{P} via forwards induction on the layer h , using, at each step h , our knowledge of the learned transitions at steps prior to h to find a policy which visits each reachable state at step h . Such an approach would lead to a *policy cover*, namely a collection of policies which visits all reachable latent states.

A major problem with this approach is that the latent states are not observed: instead, we only see observations. Hence a natural approach might be to choose policies which lead to all possible *observations* at each step h . This approach is clearly insufficient, since, e.g., a single state could emit a uniform distribution over observations. Thus we instead compute the following stronger concept: for each step h , we consider the set \mathcal{X} of all possible distributions over observations at

¹Policy covers have been used in the special case of *block MDPs* [DKJ⁺19, MHKL20], namely where different states produce disjoint observations.

step h under any general policy, and attempt to find a *barycentric spanner* of \mathcal{X} , namely a small subset $\mathcal{X}' \subset \mathcal{X}$ so that all other distributions in \mathcal{X} can be expressed as a linear combination of elements of \mathcal{X}' with bounded coefficients. By playing a policy which realizes each distribution in such a barycentric spanner \mathcal{X}' , we are able to explore all reachable latent states, despite having no knowledge about which states we are exploring. This discussion omits a key technical aspect of the proof, which is the fact that we can only compute a barycentric spanner for a set \mathcal{X} corresponding to an empirical estimate $\widehat{\mathcal{M}}$ of \mathcal{M} . A key innovation in our proof is a technique to dynamically use such barycentric spanners, even when $\widehat{\mathcal{M}}$ is inaccurate, to improve the quality of the estimate $\widehat{\mathcal{M}}$. We remark that a similar dynamic usage of barycentric spanners appeared in [FKQR21]; we discuss in the appendix why the approach of [FKQR21], as well as related approaches involving nonstationary MDPs [WL21, WZ22, WYDW21], cannot be applied here.

Taking a step back, few models in reinforcement learning (beyond tabular or linear MDPs) admit computationally efficient end-to-end learning algorithms – indeed, our main contribution is a way to circumvent the daunting task of implementing any of the various constrained optimistic planning oracles assumed in previous optimism-based approaches. We hope that our techniques may be useful in other contexts for avoiding computational intractability without resorting to oracles.

2 Preliminaries

For sets \mathcal{T}, \mathcal{Q} , let $\mathcal{Q}^{\mathcal{T}}$ denote the set of mappings from $\mathcal{T} \rightarrow \mathcal{Q}$. Accordingly, we will identify $\mathbb{R}^{\mathcal{T}}$ with $|\mathcal{T}|$ -dimensional Euclidean space, and let $\Delta(\mathcal{T}) \subset \mathbb{R}^{\mathcal{T}}$ consist of distributions on \mathcal{T} . For $d \in \mathbb{N}$ and a vector $\mathbf{v} \in \mathbb{R}^d$, we denote its components by $\mathbf{v}(1), \dots, \mathbf{v}(d)$. For integers $a \leq b$, we abbreviate a sequence $(x_a, x_{a+1}, \dots, x_b)$ by $x_{a:b}$. If $a > b$, then we let $x_{a:b}$ denote the empty sequence. Sometimes we refer to negative indices of a sequence $x_{1:n}$: in such cases the elements with negative indices may be taken to be arbitrary, as they will never affect the value of the expression. See Appendix B.1 for clarification. For $x \in \mathbb{R}$, we write $[x]_+ = \max\{x, 0\}$, and $[x]_- = -\min\{x, 0\}$. For sets \mathcal{S}, \mathcal{T} , the notation $\mathcal{S} \subset \mathcal{T}$ allows for the possibility that $\mathcal{S} = \mathcal{T}$.

2.1 Background on POMDPs

In this paper we address the problem of learning finite-horizon *partially observable Markov decision processes (POMDPs)*. Formally, a POMDP \mathcal{P} is a tuple $\mathcal{P} = (H, \mathcal{S}, \mathcal{A}, \mathcal{O}, b_1, R, \mathbb{T}, \mathbb{O})$, where: $H \in \mathbb{N}$ is a positive integer denoting the *horizon length*; \mathcal{S} is a finite set of *states* of size $S := |\mathcal{S}|$; \mathcal{A} is a finite set of *actions* of size $A := |\mathcal{A}|$; \mathcal{O} is a finite set of *observations* of size $O := |\mathcal{O}|$; b_1 is the initial distribution over states; and $R, \mathbb{T}, \mathbb{O}$ are given as follows. First, $R = (R_1, \dots, R_H)$ denotes a tuple of *reward functions*, where, for $h \in [H]$, $R_h : \mathcal{O} \rightarrow [0, 1]$ gives the reward received as a function of the observation at step h . (It is customary in the literature [JKKL20, LCSJ22] to define the rewards as being a function of the observations as opposed to being observed by the algorithm as separate information.) Second, $\mathbb{T} = (\mathbb{T}_1, \dots, \mathbb{T}_H)$ is a tuple of *transition kernels*, where, for $h \in [H]$, $s, s' \in \mathcal{S}, a \in \mathcal{A}$, $\mathbb{T}_h(s'|s, a)$ denotes the probability of transitioning from s to s' at step h when action a is taken. For each $a \in \mathcal{A}$, we will write $\mathbb{T}_h(a) \in \mathbb{R}^{S \times S}$ to denote the matrix with $\mathbb{T}_h(a)_{s,s'} = \mathbb{T}_h(s'|s, a)$. Third, $\mathbb{O} = (\mathbb{O}_1, \dots, \mathbb{O}_H)$ is a tuple of *observation matrices*, where for $h \in [H]$, $s \in \mathcal{S}, o \in \mathcal{O}$, $(\mathbb{O}_h)_{o,s}$, also written as $\mathbb{O}_h(o|s)$, denotes the probability observing o while in state s at step h . Thus $\mathbb{O}_h \in \mathbb{R}^{O \times S}$ for each h . Sometimes, for disambiguation, we will refer to the states \mathcal{S} as the *latent states* of the POMDP \mathcal{P} .

The interaction (namely, a single *episode*) with \mathcal{P} proceeds as follows: initially a state $s_1 \sim b_1$ is drawn from the initial state distribution. At each step $1 \leq h < H$, an action $a_h \in \mathcal{A}$ is chosen (as a function of previous observations and actions taken), \mathcal{P} transitions to a new state $s_{h+1} \sim \mathbb{T}_h(\cdot|s_h, a_h)$, a new observation is observed, $o_{h+1} \sim \mathbb{O}_{h+1}(\cdot|s_{h+1})$, and a reward of $R_{h+1}(o_{h+1})$ is received (and observed). We emphasize that the underlying states $s_{1:H}$ are never observed directly. As a matter of convention, we assume that no observation is observed at step $h = 1$; thus the first observation is o_2 .

2.2 Policies, value functions

A *deterministic policy* σ is a tuple $\sigma = (\sigma_1, \dots, \sigma_H)$, where $\sigma_h : \mathcal{A}^{h-1} \times \mathcal{O}^{h-1} \rightarrow \mathcal{A}$ is a mapping from *histories* up to step h , namely tuples $(a_{1:h-1}, o_{2:h})$, to actions. We will denote the collection of histories up to step h by $\mathcal{H}_h := \mathcal{A}^{h-1} \times \mathcal{O}^{h-1}$ and the set of deterministic policies by Π^{det} ,

meaning that $\Pi^{\text{det}} = \prod_{h=1}^H \mathcal{A}^{\mathcal{H}_h}$. A *general policy* π is a distribution over deterministic policies; the set of general policies is denoted by $\Pi^{\text{gen}} := \Delta(\prod_{h=1}^H \mathcal{A}^{\mathcal{H}_h})$. Given a general policy $\pi \in \Pi^{\text{gen}}$, we denote by $\sigma \sim \pi$ the draw of a deterministic policy from the distribution π ; to execute a general policy π , a sample $\sigma \sim \pi$ is first drawn and then followed for an episode of the POMDP. For a general policy π and some event \mathcal{E} , write $\mathbb{P}_{a_{1:H-1}, o_{2:H}, s_{1:H} \sim \pi}^{\mathcal{P}}(\mathcal{E})$ to denote the probability of \mathcal{E} when $s_{1:H}, a_{1:H-1}, o_{2:H}$ is drawn from a trajectory following policy π for the POMDP \mathcal{P} . At times we will compress notation in the subscript, e.g., write $\mathbb{P}_{\pi}^{\mathcal{P}}(\mathcal{E})$ if the definition of $s_{1:H}, a_{1:H-1}, o_{2:H}$ is evident. In similar spirit, we will write $\mathbb{E}_{a_{1:H-1}, o_{2:H}, s_{1:H} \sim \pi}^{\mathcal{P}}[\cdot]$ to denote expectations.

Given a general policy $\pi \in \Pi^{\text{gen}}$, define the *value function* for π at step 1 by $V_1^{\pi, \mathcal{P}}(\emptyset) = \mathbb{E}_{o_{1:H} \sim \pi}^{\mathcal{P}} \left[\sum_{h=2}^H R_h(o_h) \right]$, namely as the expected reward received by following π .

Our objective is to find a policy π which maximizes $V_1^{\pi, \mathcal{P}}(\emptyset)$, in the *PAC-RL model* [KS02]: in particular, the algorithm does not have access to the transition kernel, reward function, or observation matrices of \mathcal{P} , but can repeatedly choose a general policy π and observe the following data from a single trajectory drawn according to π : $(a_1, o_2, R_2(o_2), a_2, \dots, a_{H-1}, o_H, R_H(o_H))$. The challenge is to choose such policies π which can sufficiently explore the environment.

Finally, we remark that *Markov decision processes* (MDPs) are a special case of POMDPs where $\mathcal{O} = \mathcal{S}$ and $\mathbb{O}_h(o|s) = \mathbb{1}[o = s]$ for all $h \in [H], o, s \in \mathcal{S}$. For the MDPs we will consider, the initial state distribution will be left unspecified (indeed, the optimal policy of an MDP does not depend on the initial state distribution). Thus, we consider MDPs \mathcal{M} described by a tuple $\mathcal{M} = (H, \mathcal{S}, \mathcal{A}, R, \mathbb{T})$.

2.3 Belief contraction

A prerequisite for a computationally efficient *learning* algorithm in observable POMDPs is a computationally efficient *planning* algorithm, i.e. an algorithm to find an approximately optimal policy when the POMDP is known. Recent work [GMR22] obtains such a planning algorithm taking quasipolynomial time; we now introduce the key tools used in [GMR22], which are used in our algorithm as well.

Consider a POMDP $\mathcal{P} = (H, \mathcal{S}, \mathcal{A}, \mathcal{O}, b_1, R, \mathbb{T}, \mathbb{O})$. Given some $h \in [H]$ and a history $(a_{1:h-1}, o_{2:h}) \in \mathcal{H}_h$, the *belief state* $\mathbf{b}_h^{\mathcal{P}}(a_{1:h-1}, o_{2:h}) \in \Delta(\mathcal{S})$ is given by the distribution of the state s_h conditioned on taking actions $a_{1:h-1}$ and observing $o_{2:h}$ in the first h steps. Formally, the belief state is defined inductively as follows: $\mathbf{b}_1^{\mathcal{P}}(\emptyset) = b_1$, and for $2 \leq h \leq H$ and any $(a_{1:h-1}, o_{2:h}) \in \mathcal{H}_h$,

$$\mathbf{b}_h^{\mathcal{P}}(a_{1:h-1}, o_{2:h}) = U_{h-1}^{\mathcal{P}}(\mathbf{b}_{h-1}^{\mathcal{P}}(a_{1:h-2}, o_{2:h-1}); a_{h-1}, o_h),$$

where for $\mathbf{b} \in \Delta(\mathcal{S}), a \in \mathcal{A}, o \in \mathcal{O}$, $U_h^{\mathcal{P}}(\mathbf{b}; a, o) \in \Delta(\mathcal{S})$ is the distribution defined by

$$U_h^{\mathcal{P}}(\mathbf{b}; a, o)(s) := \frac{\mathbb{O}_{h+1}(o|s) \cdot \sum_{s' \in \mathcal{S}} \mathbf{b}(s') \cdot \mathbb{T}_h(s|s', a)}{\sum_{x \in \mathcal{S}} \mathbb{O}_{h+1}(o|x) \sum_{s' \in \mathcal{S}} \mathbf{b}(s') \cdot \mathbb{T}_h(x|s', a)}.$$

We call $U_h^{\mathcal{P}}$ the *belief update operator*. The belief state $\mathbf{b}_h^{\mathcal{P}}(a_{1:h-1}, o_{2:h})$ is a sufficient statistic for the sequence of future actions and observations under any deterministic policy. In particular, the optimal policy can be expressed as a function of the belief state, rather than the entire history. Thus, a natural approach to plan a near-optimal policy is to find a small set $\mathcal{B} \subset \Delta(\mathcal{S})$ of distributions over states such that each possible belief state $\mathbf{b}_h^{\mathcal{P}}(a_{1:h-1}, o_{2:h})$ is close to some element of \mathcal{B} . Unfortunately, this is not possible, even in observable POMDPs [GMR22, Example D.2]. The main result of [GMR22] circumvents this issue by showing that there is a subset $\mathcal{B} \subset \Delta(\mathcal{S})$ of quasipolynomial size (depending on \mathcal{P}) so that $\mathbf{b}_h^{\mathcal{P}}(a_{1:h-1}, o_{2:h})$ is close to some element of \mathcal{B} *in expectation* under any given policy. To state the result of [GMR22], we need to introduce *approximate belief states*:

Definition 2.1 (Approximate belief state). Fix a POMDP $\mathcal{P} = (H, \mathcal{S}, \mathcal{A}, \mathcal{O}, b_1, R, \mathbb{T}, \mathbb{O})$. For any distribution $\mathcal{D} \in \Delta(\mathcal{S})$, as well as any choices of $1 \leq h \leq H$ and $L \geq 0$, the approximate belief state $\mathbf{b}_h^{\text{apx}, \mathcal{P}}(a_{h-L:h-1}, o_{h-L+1:h}; \mathcal{D})$ is defined as follows, via induction on L : in the case that $L = 0$, then we define

$$\mathbf{b}_h^{\text{apx}, \mathcal{P}}(\emptyset; \mathcal{D}) := \begin{cases} b_1 & : h = 1 \\ \mathcal{D} & : h > 1, \end{cases}$$

and for the case that $L > 0$, define, for $h > L$,

$$\mathbf{b}_h^{\text{apx},\mathcal{P}}(a_{h-L:h-1}, o_{h-L+1:h}; \mathcal{D}) := U_{h-1}^{\mathcal{P}}(\mathbf{b}_{h-1}^{\text{apx},\mathcal{P}}(a_{h-L:h-2}, o_{h-L+1:h-1}; \mathcal{D}); a_{h-1}, o_h).$$

We extend the above definition to the case that $h \leq L$ by defining $\mathbf{b}_h^{\text{apx},\mathcal{P}}(a_{h-L:h-1}, o_{h-L+1:h}; \mathcal{D}) := \mathbf{b}_h^{\text{apx},\mathcal{P}}(a_{\max\{1, h-L\}:h-1}, o_{\max\{2, h-L+1\}:h}; \mathcal{D})$. In words, the approximate belief state $\mathbf{b}_h^{\text{apx},\mathcal{P}}(a_{h-L:h-1}, o_{h-L+1:h}; \mathcal{D})$ is obtained by applying the belief update operator starting from the distribution \mathcal{D} at step $h-L$, if $h-L > 1$ (and otherwise, starting from b_1 , at step 1). At times, we will drop the superscript \mathcal{P} from the above definitions and write $\mathbf{b}_h, \mathbf{b}_h^{\text{apx}}$.

The main technical result of [GMR22], stated as Theorem 2.1 below (with slight differences, see Appendix B), proves that if the POMDP \mathcal{P} is γ -observable for some $\gamma > 0$, then for a wide range of distributions \mathcal{D} , for sufficiently large L , the approximate belief state $\mathbf{b}_h^{\text{apx},\mathcal{P}}(a_{h-L:h-1}, o_{h-L+1:h}; \mathcal{D})$ will be close to (i.e., “contract to”) the true belief state $\mathbf{b}_h^{\mathcal{P}}(a_{1:h-1}, o_{2:h})$.

Theorem 2.1 (Belief contraction; Theorems 4.1 and 4.7 of [GMR22]). *Consider any γ -observable POMDP \mathcal{P} , any $\epsilon > 0$ and $L \in \mathbb{N}$ so that $L \geq C \cdot \min \left\{ \frac{\log(1/(\epsilon\phi)) \log(\log(1/\phi)/\epsilon)}{\gamma^2}, \frac{\log(1/(\epsilon\phi))}{\gamma^4} \right\}$. Fix any $\pi \in \Pi^{\text{gen}}$, and suppose that $\mathcal{D} \in \Delta^{\mathcal{S}}$ satisfies $\frac{\mathbf{b}_h^{\mathcal{P}}(a_{1:h-L-1}, o_{2:h-L})(s)}{\mathcal{D}(s)} \leq \frac{1}{\phi}$ for all $(a_{h-1}, o_{2:h})$. Then*

$$\mathbb{E}_{(a_{1:h-1}, o_{2:h}) \sim \pi}^{\mathcal{P}} \left\| \mathbf{b}_h^{\mathcal{P}}(a_{1:h-1}, o_{2:h}) - \mathbf{b}_h^{\text{apx},\mathcal{P}}(a_{h-L:h-1}, o_{h-L+1:h}; \mathcal{D}) \right\|_1 \leq \epsilon.$$

2.4 Visitation distributions

For a POMDP $\mathcal{P} = (H, \mathcal{S}, \mathcal{A}, \mathcal{O}, b_1, R, \mathbb{T}, \mathbb{O})$, policy $\pi \in \Pi^{\text{gen}}$, and step $h \in [H]$, the (latent) *state visitation distribution* at step h is $d_{\mathcal{S},h}^{\mathcal{P},\pi} \in \Delta(\mathcal{S})$ defined by $d_{\mathcal{S},h}^{\mathcal{P},\pi}(s) := \mathbb{P}_{s_{1:H} \sim \pi}^{\mathcal{P}}(s_h = s)$, and the *observation visitation distribution* at step h is $d_{\mathcal{O},h}^{\mathcal{P},\pi} \in \Delta(\mathcal{O})$ defined by $d_{\mathcal{O},h}^{\mathcal{P},\pi} := \mathbb{P}_{o_{1:H} \sim \pi}^{\mathcal{P}}(o_h = \cdot) = \mathbb{O}_h \cdot d_{\mathcal{S},h}^{\mathcal{P},\pi}$. As will be discussed in Section 4.1, Theorem 2.1 implies that the transitions of the POMDP \mathcal{P} can be approximated by those of an MDP \mathcal{M} whose states consist of L -tuples of actions and observations. Thus we will often deal with such MDPs \mathcal{M} of the form $\mathcal{M} = (\bar{H}, \bar{\mathcal{Z}}, \mathcal{A}, R, \mathbb{T})$ where the set of states has the product structure $\bar{\mathcal{Z}} = \mathcal{A}^L \times \mathcal{O}^L$. We then define $\circ : \bar{\mathcal{Z}} \rightarrow \mathcal{O}$ by $\circ(a_{1:L}, o_{1:L}) = o_L$. For such MDPs, we define the observation visitation distributions by $d_{\mathcal{O},h}^{\mathcal{M},\pi}(o) := \mathbb{P}_{z_{1:H} \sim \pi}^{\mathcal{M}}(\circ(z_h) = o)$. Finally, for $o \in \mathcal{O}$, we let $e_o \in \mathbb{R}^{\mathcal{O}}$ denote the o th unit vector; thus, for instance, we have $d_{\mathcal{O},h}^{\mathcal{P},\pi}(o) = \langle e_o, d_{\mathcal{O},h}^{\mathcal{P},\pi} \rangle$.

3 Main result: learning observable POMDPs in quasipolynomial time

Theorem 3.1 below states our main guarantee for BaSeCAMP (**B**arycentric **S**panner policy **C**over with **A**pproximate **M**DP; Algorithm 3).

Theorem 3.1. *Given any $\alpha, \beta, \gamma > 0$ and γ -observable POMDP \mathcal{P} , BaSeCAMP with parameter settings as described in Section C.1 outputs a policy which is α -suboptimal with probability at least $1 - \beta$, using time and sample complexity bounded by $(OA)^{CL} \log(1/\beta)$, where $C > 1$ is a constant and $L = \min \left\{ \frac{\log(HSO/(\alpha\gamma))}{\gamma^4}, \frac{\log^2(HSO/(\alpha\gamma))}{\gamma^2} \right\}$.*

It is natural to ask whether the complexity guarantee of Theorem 3.1 can be improved further. [GMR22, Theorem 6.4] shows that under the Exponential Time Hypothesis, there is no algorithm for planning in γ -observable POMDPs which runs in time $(SAHO)^{o(\log(SAHO/\alpha)/\gamma)}$ and produces α -suboptimal policies, *even if the POMDP is known*. Thus, up to polynomial factors in the exponent, Theorem 3.1 is optimal. It is plausible, however, that there could be an algorithm which runs in quasipolynomial time yet only needs *polynomially* many samples; we leave this question for future work.

4 Algorithm description

4.1 Approximating \mathcal{P} with an MDP.

A key consequence of observability is that, by the belief contraction result of Theorem 2.1, the POMDP \mathcal{P} is well-approximated by an MDP \mathcal{M} of quasi-polynomial size. In more detail, we will apply Theorem 2.1 with $\phi = 1/S$, $\mathcal{D} = \text{Unif}(\mathcal{S})$, and some $L = \text{poly}(\log(S/\epsilon)/\gamma)$ sufficiently large so as to satisfy the requirement of the theorem statement. The MDP \mathcal{M} has state space $\mathcal{Z} := \mathcal{A}^L \times \mathcal{O}^L$, horizon H , action space \mathcal{A} , and transitions $\mathbb{P}_h^{\mathcal{M}}(\cdot|z_h, a_h)$ which are defined via a belief update on the approximate belief state $\mathbf{b}_h^{\text{apx}, \mathcal{P}}(z_h; \text{Unif}(\mathcal{S}))$ ²: in particular, for a state $z_h = (a_{h-L:h-1}, o_{h-L+1:h}) \in \mathcal{Z}$ of \mathcal{M} , action a_h , and subsequent observation $o_{h+1} \in \mathcal{O}$, define

$$\mathbb{P}_h^{\mathcal{M}}((a_{h-L+1:h}, o_{h-L+2:h+1})|z_h, a_h) := e_{o_{h+1}}^\top \cdot \mathbb{O}_{h+1}^{\mathcal{P}} \cdot \mathbb{T}_h^{\mathcal{P}}(a_h) \cdot \mathbf{b}_h^{\text{apx}, \mathcal{P}}(z_h; \text{Unif}(\mathcal{S})). \quad (1)$$

The above definition should be compared with the probability of observing o_{h+1} given history $(a_{1:h}, o_{2:h})$ and policy π when interacting with the POMDP \mathcal{P} , which is

$$\mathbb{P}_{o_{h+1} \sim \pi}^{\mathcal{P}}(o_{h+1}|a_{1:h}, o_{2:h}) = e_{o_{h+1}}^\top \cdot \mathbb{O}_{h+1}^{\mathcal{P}} \cdot \mathbb{T}_h^{\mathcal{P}}(a_h) \cdot \mathbf{b}_h^{\mathcal{P}}(a_{1:h-1}, o_{2:h}). \quad (2)$$

Theorem 2.1 gives that $\left\| \mathbf{b}_h^{\mathcal{P}}(a_{1:h-1}, o_{2:h}) - \mathbf{b}_h^{\text{apx}, \mathcal{P}}(a_{h-L:h-1}, o_{h-L+1:h}; \text{Unif}(\mathcal{S})) \right\|_1$ is small in expectation under any general policy π , which, using (1) and (2), gives that, for all $\pi \in \Pi^{\text{gen}}$ and $h \in [H]$,

$$\mathbb{E}_{a_{1:h}, o_{2:h} \sim \pi} \sum_{o_{h+1} \in \mathcal{O}} \left| \mathbb{P}_h^{\mathcal{M}}(o_{h+1}|z_h, a_h) - \mathbb{P}_h^{\mathcal{P}}(o_{h+1}|a_{1:h}, o_{2:h}) \right| \leq \epsilon. \quad (3)$$

(Above we have written $z_h = (a_{h-L:h-1}, o_{h-L+1:h})$ and, via abuse of notation, $\mathbb{P}_h^{\mathcal{M}}(o_{h+1}|z_h, a_h)$ in place of $\mathbb{P}_h^{\mathcal{M}}((a_{h-L+1:h}, o_{h-L+2:h+1})|z_h, a_h)$.) The inequality (3) establishes that the dynamics of \mathcal{P} under any policy may be approximated by those of the MDP \mathcal{M} . Crucially, this implies that there exists a deterministic Markov policy for \mathcal{M} which is near-optimal among general policies for \mathcal{P} ; the set of such Markov policies for \mathcal{M} is denoted by $\Pi_{\mathcal{Z}}^{\text{markov}}$. Because of the Markovian structure of \mathcal{M} , such a policy can be found in time polynomial in the size of \mathcal{M} (which is quasi-polynomial in the underlying problem parameters), if \mathcal{M} is known. Of course, \mathcal{M} is not known.

Approximately learning the MDP \mathcal{M} . These observations suggest the following model-based approach of trying to learn the transitions of \mathcal{M} . Suppose that we know a sequence of general policies π^1, \dots, π^H (abbreviated as $\pi^{1:H}$) so that for each h , π^h visits a uniformly random state of \mathcal{P} at step $h-L$ (i.e. $d_{\mathcal{S}, h-L}^{\mathcal{P}, \pi^h} = \text{Unif}(\mathcal{S})$). Then we can estimate the transitions of \mathcal{M} as follows: play π^h for $h-L-1$ steps and then play $L+1$ random actions, generating a trajectory $(a_{1:h}, o_{2:h+1})$. Conditioned on $z_h = (a_{h-L:h-1}, o_{h-L+1:h})$ and final action a_h , the last observation of this sample trajectory, o_{h+1} , would give an unbiased draw from the transition distribution $\mathbb{P}_h^{\mathcal{M}}(\cdot|z_h, a_h)$. Repeating this procedure would allow estimation of $\mathbb{P}_h^{\mathcal{M}}$ (see Lemma E.1).

This idea is formalized in the procedure `ConstructMDP` (Algorithm 1): given a sequence of general policies π^1, \dots, π^H (abbreviated $\pi^{1:H}$), `ConstructMDP` constructs an MDP, denoted $\widehat{\mathcal{M}} = \widehat{\mathcal{M}}(\pi^{1:H})$, which empirically approximates \mathcal{M} using the sampling procedure described above. For technical reasons, $\widehat{\mathcal{M}}$ actually has state space $\overline{\mathcal{Z}} := \mathcal{A}^L \cdot \overline{\mathcal{O}}^L$, where $\overline{\mathcal{O}} := \mathcal{O} \cup \{o^{\text{sink}}\}$ and o^{sink} is a special ‘‘sink observation’’ so that, after o^{sink} is observed, all future observations are also o^{sink} . Furthermore, we remark that $d_{\mathcal{S}, h-L}^{\mathcal{P}, \pi^h}$ does not have to be exactly uniform – it suffices if π^h visits all states of \mathcal{P} at step $h-L$ with non-negligible probability.

4.2 Exploration via barycentric spanners

The above procedure for approximating \mathcal{M} with $\widehat{\mathcal{M}}$ omits a crucial detail: *how can we find the ‘‘exploratory policies’’ $\pi^{1:H}$?* Indeed a major obstacle to finding such policies is that *we never directly*

²For simplicity, descriptions of the reward function of \mathcal{M} are omitted; we refer the reader to the appendix for the full details of the proof

observe the states of \mathcal{P} . By repeatedly playing a policy π on \mathcal{P} , we can estimate the induced observation visitation distribution $d_{\mathcal{O},h-L}^{\mathcal{P},\pi}$, which is related to the state visitation distribution via the equality $\mathbb{O}_{h-L}^\dagger \cdot d_{\mathcal{O},h-L}^{\mathcal{P},\pi} = d_{\mathcal{S},h-L}^{\mathcal{P},\pi}$. Unfortunately, the matrix \mathbb{O}_{h-L} is still unknown, and in general *unidentifiable*.

On the positive side, we can attempt to learn \mathcal{M} layer by layer – in particular, when learning the h th layer, we can assume that we have learned previous layers, i.e., $d_{\mathcal{O},h-L}^{\widehat{\mathcal{M}},\pi}$ approximates $d_{\mathcal{O},h-L}^{\mathcal{M},\pi}$, and therefore $d_{\mathcal{O},h-L}^{\mathcal{P},\pi}$. Even though $\widehat{\mathcal{M}}$ does not have underlying latent states, we can define “formal” latent state distributions on $\widehat{\mathcal{M}}$ in analogy with \mathcal{P} , i.e. $d_{\mathcal{S},h-L}^{\widehat{\mathcal{M}},\pi} := \mathbb{O}_{h-L}^\dagger \cdot d_{\mathcal{O},h-L}^{\widehat{\mathcal{M}},\pi}$. But this does not seem helpful, again because \mathbb{O}_{h-L} is unknown. Our first key insight is that a policy π^h , for which $d_{\mathcal{S},h-L}^{\widehat{\mathcal{M}},\pi^h}$ puts non-negligible mass on all states, can be found (when it exists) via knowledge of $\widehat{\mathcal{M}}$ and the technique of *Barycentric Spanners* – all without ever explicitly computing $d_{\mathcal{S},h-L}^{\widehat{\mathcal{M}},\pi^h}$.

Barycentric spanners. Suppose we knew that the transitions of our empirical estimate $\widehat{\mathcal{M}}$ approximate those of \mathcal{M} up to step h , and we want to find a policy π^h for which the (formal) latent state distribution $d_{\mathcal{S},h-L}^{\widehat{\mathcal{M}},\pi^h}$ is non-negligible on all states. Unfortunately, the set of achievable latent state distributions $\{\mathbb{O}_{h-L}^\dagger \cdot d_{\mathcal{O},h-L}^{\widehat{\mathcal{M}},\pi} : \pi \in \Pi^{\text{gen}}\} \subset \mathbb{R}^{\mathcal{S}}$ is defined implicitly, via the unknown observation matrix \mathbb{O}_{h-L} . But we do have access to $\mathcal{X}_{\widehat{\mathcal{M}},h-L} := \{d_{\mathcal{O},h-L}^{\widehat{\mathcal{M}},\pi} : \pi \in \Pi^{\text{gen}}\} \subset \mathbb{R}^{\mathcal{O}}$, the set of achievable distributions over the observation at step $h-L$. In particular, for any reward function on observations at step $h-L$, we can efficiently (by dynamic programming) find a policy π that maximizes reward on $\widehat{\mathcal{M}}$ over all policies. In other words, we can solve *linear optimization* problems over \mathcal{X} . By a classic result [AK08], we can thus efficiently find a *barycentric spanner* for $\mathcal{X}_{\widehat{\mathcal{M}},h-L}$:

Definition 4.1 (Barycentric spanner). Consider a subset $\mathcal{X} \subset \mathbb{R}^d$. For $B \geq 1$, a set $\mathcal{X}' \subset \mathcal{X}$ of size d is a B -approximate barycentric spanner of \mathcal{X} if each $x \in \mathcal{X}$ can be expressed as a linear combination of elements in \mathcal{X}' with coefficients in $[-B, B]$.

Using the guarantee of [AK08] (restated as Lemma D.1) applied to the set $\mathcal{X}_{\widehat{\mathcal{M}},h-L}$, we can find, in time polynomial in the size of $\widehat{\mathcal{M}}$, a 2-approximate barycentric spanner $\tilde{\pi}^1, \dots, \tilde{\pi}^O$ of $\mathcal{X}_{\widehat{\mathcal{M}},h-L}$; this procedure is formalized in `BarySpannerPolicy` (Algorithm 2). Thus, for any policy π , the observation distribution $d_{\mathcal{O},h-L}^{\widehat{\mathcal{M}},\pi}$ induced by π is a linear combination of the distributions $\{d_{\mathcal{O},h-L}^{\widehat{\mathcal{M}},\tilde{\pi}^i} : i \in [O]\}$ with coefficients in $[-2, 2]$. Since $d_{\mathcal{S},h-L}^{\widehat{\mathcal{M}},\pi} = \mathbb{O}_{h-L}^\dagger \cdot d_{\mathcal{O},h-L}^{\widehat{\mathcal{M}},\pi}$ for all π , it follows that the formal latent state distribution $d_{\mathcal{S},h-L}^{\widehat{\mathcal{M}},\pi}$ induced by π is a linear combination of the formal latent state distributions $\{d_{\mathcal{S},h-L}^{\widehat{\mathcal{M}},\tilde{\pi}^i} : i \in [O]\}$ with the same coefficients. Now, intuitively, the randomized mixture policy $\pi_{\text{mix}} = \frac{1}{O}(\tilde{\pi}^1 + \dots + \tilde{\pi}^O)$ should explore every reachable state; indeed, we show the following guarantee for `BarySpannerPolicy`:

Lemma 4.1 (Informal version & special case of Lemma D.2). *In the above setting, under some technical conditions, for all $s \in \mathcal{S}$ and $\pi \in \Pi^{\text{gen}}$, it holds that $d_{\mathcal{S},h-L}^{\widehat{\mathcal{M}},\pi_{\text{mix}}}(s) \geq \frac{1}{4O^2} \cdot d_{\mathcal{S},h-L}^{\widehat{\mathcal{M}},\pi}(s)$.*

But recall the original goal: a policy π^h which explores \mathcal{P} – not $\widehat{\mathcal{M}}$. Unfortunately, those states in \mathcal{P} which can only be reached with probability that is small compared to the distance between \mathcal{P} and $\widehat{\mathcal{M}}$ may be missed by π_{mix} . When we use π_{mix} to compute the next-step transitions of $\widehat{\mathcal{M}}$, this will lead to additional error when applying belief contraction (Theorem 2.1), and therefore additional error between \mathcal{P} and $\widehat{\mathcal{M}}$. If not handled carefully, this error will compound exponentially over layers.

4.3 The full algorithm via iterative discovery

The solution to the dilemma discussed above is to not try to construct our estimate $\widehat{\mathcal{M}}$ of \mathcal{M} layer by layer, hoping that at each layer we can explore all reachable states of \mathcal{P} despite making errors in

Algorithm 1 ConstructMDP

```
1: procedure CONSTRUCTMDP( $L, N_0, N_1, \pi^1, \dots, \pi^H$ )
2:   for  $1 \leq h \leq H$  do
3:     Let  $\widehat{\pi}^h$  be the policy which follows  $\pi^h$  for the first  $\max\{h-L-1, 0\}$  steps and thereafter
     chooses uniformly random actions.
4:     Draw  $N_0$  independent trajectories from the policy  $\widehat{\pi}^h$ :
5:     Denote the data from the  $i$ th trajectory by  $a_{1:H-1}^i, o_{2:H}^i$ , for  $i \in [N_0]$ .
6:     Set  $z_h^i = (a_{h-L:h-1}^i, o_{h-L+1:h}^i)$  for all  $i \in [N_0], h \in [H]$ .
7:     // Construct the transitions  $\mathbb{P}_h^{\widehat{\mathcal{M}}}(z_{h+1}|z_h, a_h)$  as follows:
8:     for each  $z_h = (a_{h-L:h-1}, o_{h-L+1:h}) \in \mathcal{Z}, a_h \in \mathcal{A}$  do
9:       // Define  $\mathbb{P}_h^{\widehat{\mathcal{M}}}(\cdot|z_h, a_h)$  to be the empirical distribution of  $z_{h+1}^i|z_h^i, a_h^i$ , as follows:
10:      For  $o_{h+1} \in \mathcal{O}$ , define  $\varphi(o_{h+1}) := |\{i : (a_{\max\{1, h-L\}:h}^i, o_{\max\{2, h-L+1\}:h+1}^i) =$ 
      ( $a_{\max\{1, h-L\}:h}, o_{\max\{2, h-L+1\}:h+1}$ ) $\}|$ .
11:      if  $\sum_{o_{h+1}} \varphi(o_{h+1}) \geq N_1$  then
12:        Set  $\mathbb{P}_h^{\widehat{\mathcal{M}}}((a_{h-L+1:h}, o_{h-L+2:h+1})|z_h, a_h) := \frac{\varphi(o_{h+1})}{\sum_{o_{h+1}} \varphi(o_{h+1})}$  for all  $o_{h+1} \in \mathcal{O}$ .
13:        Set  $R_h^{\widehat{\mathcal{M}}}(z_h, a_h) := R_h^{\mathcal{P}}(o_h^i)$  for some  $i$  with  $o_h^i = o_h$ .  $\triangleright R_h^{\mathcal{P}}(o_h^i)$  is observed.
14:      else
15:        Let  $\mathbb{P}_h^{\widehat{\mathcal{M}}}(\cdot|z_h, a_h)$  put all its mass on  $(a_{h-L+1:h}, (o_{h-L+2:h}, o^{\text{sink}}))$ .
16:      for each  $z_h = (a_{h-L:h-1}, o_{h-L+1:h}) \in \mathcal{Z} \setminus \mathcal{Z}$  and  $a_h \in \mathcal{A}$  do
17:        Let  $\mathbb{P}_h^{\widehat{\mathcal{M}}}(\cdot|z_h, a_h)$  put all its mass on  $(a_{h-L+1:h}, (o_{h-L+2:h}, o^{\text{sink}}))$ .
18:    Let  $\widehat{\mathcal{M}}$  denote the MDP  $(\mathcal{Z}, H, \mathcal{A}, R^{\widehat{\mathcal{M}}}, \mathbb{P}^{\widehat{\mathcal{M}}})$ .
19:    return the MDP  $\widehat{\mathcal{M}}$ , which we denote by  $\widehat{\mathcal{M}}(\pi^{1:H})$ .
```

Algorithm 2 BarySpannerPolicy

```
1: procedure BARYSPANNERPOLICY( $\widehat{\mathcal{M}}, h$ )  $\triangleright \widehat{\mathcal{M}}$  is MDP on state space  $\mathcal{Z}$ , horizon  $H$ ;  $h \in [H]$ 
2:   if  $h \leq L$  then return an arbitrary general policy.
3:   Let  $\mathcal{O}$  be the linear optimization oracle which given  $r \in \mathbb{R}^{\mathcal{O}}$ , returns
    $\arg \max_{\pi \in \Pi_{\mathcal{Z}}^{\text{markov}}} \langle r, d_{\mathcal{O}, h-L}^{\pi, \widehat{\mathcal{M}}} \rangle$  and  $\max_{\pi \in \Pi_{\mathcal{Z}}^{\text{markov}}} \langle r, d_{\mathcal{O}, h-L}^{\pi, \widehat{\mathcal{M}}} \rangle$   $\triangleright$  Note that  $\mathcal{O}$  can be implemented
   in time  $\tilde{O}(|\mathcal{Z}| \cdot HO)$  using dynamic programming
4:   Using the algorithm of [AK08] with oracle  $\mathcal{O}$ , compute policies  $\{\pi^1, \dots, \pi^{\mathcal{O}}\}$  so that
    $\{d_{\mathcal{O}, h-L}^{\pi^i, \widehat{\mathcal{M}}} : i \in [\mathcal{O}]\}$  is a 2-approximate barycentric spanner of  $\{d_{\mathcal{O}, h-L}^{\pi, \widehat{\mathcal{M}}} : \pi \in \Pi_{\mathcal{Z}}^{\text{markov}}\}$ .
    $\triangleright$  This algorithm requires only  $O(\mathcal{O}^2 \log \mathcal{O})$  calls to  $\mathcal{O}$ 
5:   return the general policy  $\frac{1}{\mathcal{O}} \cdot \sum_{i=1}^{\mathcal{O}} \pi^i$ .
```

earlier layers of $\widehat{\mathcal{M}}$. Instead, we have to be able to go back and “fix” errors in our empirical estimates at earlier layers. This task is performed in our main algorithm, BaSeCAMP (Algorithm 3). For some $K \in \mathbb{N}$, BaSeCAMP runs for a total of K iterations: at each iteration $k \in [K]$, BaSeCAMP defines H general policies, $\bar{\pi}^{k,1}, \dots, \bar{\pi}^{k,H} \in \Pi^{\text{gen}}$ (abbreviated $\bar{\pi}^{k,1:H}$; step 4). The algorithm’s overall goal is that, for some k , for each $h \in [H]$, $\bar{\pi}^{k,h}$ explores all latent states at step $h-L$ that are reachable by any policy.

To ensure that this condition holds at some iteration, BaSeCAMP performs the following two main steps for each iteration k : first, it calls Algorithm 1 to construct an MDP, denoted $\widehat{\mathcal{M}}^{(k)}$, using the policies $\bar{\pi}^{k,1:H}$. Then, for each $h \in [H]$, it passes the tuple $(\widehat{\mathcal{M}}^{(k)}, h)$ to BarySpannerPolicy, which returns as output a general policy, $\pi^{k+1,h,0}$. Then, policies $\pi^{k+1,h}$ are produced (step 7) by averaging $\pi^{k+1,h',0}$, for all $h' \geq h$. The policies $\pi^{k+1,h}$ are then mixed into the policies $\bar{\pi}^{k+1,h}$ for the next iteration $k+1$. It follows from properties of BarySpannerPolicy that, in the event that the policies $\bar{\pi}^{k,h}$ are not sufficiently exploratory, one of the new policies $\pi^{k+1,h}$ visits *some* latent state $(s, h') \in \mathcal{S} \times [H]$, which was not previously visited by $\bar{\pi}^{k,1:H}$ with significant probability. Thus, after a total of $K = O(SH)$ iterations k , it follows that we must have visited all (reachable) latent

states of the POMDP. At the end of these K iterations, BaSeCAMP computes an optimal policy for each $\widehat{\mathcal{M}}^{(k)}$ and returns the best of them (as evaluated on fresh trajectories drawn from \mathcal{P} ; step 12).

Algorithm 3 BaSeCAMP (Barycentric Spanner policy Cover with Approximate MDP)

```

1: procedure BaSeCAMP( $L, N_0, N_1, \alpha, \beta, K$ )
2:   Initialize  $\pi^{1,1}, \dots, \pi^{1,H}$  to be arbitrary policies.
3:   for  $k \in [K]$  do
4:     For each  $h \in [H]$ , set  $\bar{\pi}^{k,h} = \frac{1}{k} \sum_{k'=1}^k \pi^{k',h}$ .
5:     Run ConstructMDP( $L, N_0, N_1, \bar{\pi}^{k,1:H}$ ) and let its output be  $\widehat{\mathcal{M}}^{(k)}$ .
6:     For each  $h \in [H]$ , let  $\pi^{k+1,h,0}$  be the output of BarySpannerPolicy( $\widehat{\mathcal{M}}^{(k)}, h$ ).
7:     For each  $h \in [H]$ , define  $\pi^{k+1,h} := \frac{1}{H-h+1} \sum_{h'=h}^H \pi^{k+1,h',0}$ .
8:     // Choose the best optimal policy amongst all  $\widehat{\mathcal{M}}^{(k)}$ 
9:     for  $k \in [K]$  do
10:      Let  $\pi_*^k$  denote an optimal policy of  $\widehat{\mathcal{M}}^{(k)}$ .
11:      Execute  $\pi_*^k$  for  $\frac{100H^2 \log K/\beta}{\alpha^2}$  trajectories and let the mean reward across them be  $\hat{r}^k$ .
12:      Let  $k^* = \arg \max_{k \in [K]} \hat{r}^k$ .
13:   return  $\pi_*^{k^*}$ .

```

5 Proof Outline

We now briefly outline the proof of Theorem 3.1; further details may be found in the appendix. The high-level idea of the proof is to show that the algorithm BaSeCAMP makes a given amount of progress for each iteration k , as specified in the following lemma:

Lemma 5.1 (“Progress lemma”: informal version of Lemma I.2). *Fix any iteration k in Algorithm 3, step 3. Then, for some parameters δ, ϕ with $\alpha \gg \delta \gg \phi > 0$, one of the following statements holds:*

1. Any (s, h) with $d_{\mathcal{S}, h-L}^{\mathcal{P}, \bar{\pi}^{k,h}}(s) < \phi$ satisfies $d_{\mathcal{S}, h-L}^{\mathcal{P}, \pi}(s) \leq \delta$ for all general policies π .
2. There is some $(h, s) \in [H] \times \mathcal{S}$ so that:

$$d_{\mathcal{S}, h-L}^{\mathcal{P}, \bar{\pi}^{k,h}}(s) < \phi \cdot H^2 S, \quad \text{yet, for all } k' > k: \quad d_{\mathcal{S}, h-L}^{\mathcal{P}, \bar{\pi}^{k',h}}(s) \geq \phi \cdot H^2 S.$$

Given Lemma 5.1, the proof of Theorem 3.1 is fairly straightforward. In particular, each $(h, s) \in [H] \times \mathcal{S}$ can only appear as the specified pair in item 2 of the lemma for a single iteration k . Thus, as long as $K > HS$, item 1 must hold for some value of k , say $k^* \in [K]$. In turn, it is not difficult to show from this that the MDP $\widehat{\mathcal{M}}^{(k^*)}$ is a good approximation of \mathcal{P} , in the sense that for any general policy π , the values of π in $\widehat{\mathcal{M}}^{(k^*)}$ and in \mathcal{P} are close (Lemma H.3). Thus, the optimal policy $\pi_*^{k^*}$ of $\widehat{\mathcal{M}}^{(k^*)}$ will be a near-optimal policy of \mathcal{P} , and steps 9 through 12 of BaSeCAMP will identify either the policy $\pi_*^{k^*}$ or some other policy $\pi_*^{k'}$ which has even higher reward on \mathcal{P} .

Proof of the progress lemma. The bulk of the proof of Theorem 3.1 consists of the proof of Lemma 5.1, which we proceed to outline. Suppose that item 1 of the lemma does not hold, meaning that there is some $\pi \in \Pi^{\text{gen}}$ and some $(h, s) \in [H] \times \mathcal{S}$ so that $d_{\mathcal{S}, h-L}^{\mathcal{P}, \bar{\pi}^{k,h}}(s) < \phi$ yet $d_{\mathcal{S}, h-L}^{\mathcal{P}, \pi}(s) > \delta$; i.e., $\bar{\pi}^{k,h}$ does *not* explore $(s, h-L)$, but the policy π does. Roughly speaking, BaSeCAMP ensures that item 2 holds in this case, using the following two steps:

(A) First, we show that $\langle e_s, \mathbb{O}_{h-L}^\dagger \cdot d_{\mathcal{O}, h-L}^{\widehat{\mathcal{M}}^{(k)}, \pi} \rangle \geq \delta'$, where δ' is some parameter satisfying $\delta \gg \delta' \gg \phi$.

In words, the estimate of the underlying state distribution provided by $\widehat{\mathcal{M}}^{(k)}$ also has the property that some policy π visits $(s, h-L)$ with non-negligible probability (namely, δ'). While this statement would be straightforward if $\widehat{\mathcal{M}}^{(k)}$ were a close approximation \mathcal{P} , this is *not* necessarily the case (indeed, if it were the case, then item 1 of Lemma 5.1 would hold). To circumvent this issue, we introduce a family of intermediate POMDPs indexed by $H' \in [H]$ and denoted $\overline{\mathcal{P}}_{\phi, H'}(\bar{\pi}^{k,1:H})$,

which we call *truncated POMDPs*. Roughly speaking, the truncated POMDP $\overline{\mathcal{P}}_{\phi, H'}(\overline{\pi}^{k, 1:H})$ diverts transitions away from all states at step $H' - L$ which $\overline{\pi}^{k, H'}$ does not visit with probability at least ϕ . This modification is made to allow Theorem 2.1 to be applied to $\overline{\mathcal{P}}_{\phi, H'}(\overline{\pi}^{k, 1:H})$ and any general policy π . By doing so, we may show a *one-sided* error bound between $\overline{\mathcal{P}}_{\phi, H'}(\overline{\pi}^{k, 1:H})$ and $\widehat{\mathcal{M}}^{(k)}$ (Lemma G.4) which, importantly, holds even when the policies $\overline{\pi}^{k, 1:H}$ may fail to explore some states. It is this one-sided error bound which implies the lower bound on $\langle e_s, \mathbb{O}_{h-L}^\dagger \cdot d_{\mathbb{O}, h-L}^{\widehat{\mathcal{M}}^{(k)}, \pi} \rangle$.

(B) Second, we show that the policy $\pi^{k+1, h, 0}$ produced by `BarySpannerPolicy` in step 6 explores $(s, h - L)$ with sufficient probability. To do so, we first use Lemma 4.1 to conclude that $\langle e_s, \mathbb{O}_{h-L}^\dagger \cdot d_{\mathbb{O}, h-L}^{\widehat{\mathcal{M}}^{(k)}, \pi^{k+1, h, 0}} \rangle \geq \frac{\delta'}{4O^2}$. The more challenging step is to use this fact to conclude a lower bound on $d_{\mathbb{S}, h-L}^{\mathcal{P}, \pi^{k+1, h, 0}}(s)$; unfortunately, the one-sided error bound between \mathcal{P} and $\widehat{\mathcal{M}}^{(k)}$ that we used in the previous paragraph goes in the wrong direction here. The solution is to use Lemma H.3, which has the following consequence: either $d_{\mathbb{S}, h-L}^{\mathcal{P}, \pi^{k+1, h, 0}}(s)$ is not too small, or else the policy $\pi^{k+1, h, 0}$ visits some state at a step *prior to* $h - L$ which was not sufficiently explored by any of the policies $\overline{\pi}^{k, 1:H}$ (see Section C for further details). In either case $\pi^{k+1, h, 0}$ visits a state that was not previously explored, and the fact that $\pi^{k+1, h, 0}$ is mixed in to $\overline{\pi}^{k', h'}$ for $k' > k, h' \leq h$ (steps 4, 7 of `BaSeCAMP`) allows us to conclude item 2 of Lemma 5.1.

6 Conclusion

In this paper we have demonstrated the first quasipolynomial-time (and quasipolynomial-sample) algorithm for learning observable POMDPs. Several interesting directions for future work remain:

- It is straightforward to show that a γ -observable POMDP is $\Omega(\gamma/\sqrt{S})$ weakly-revealing in the sense of [LCSJ22, Assumption 1]. Thus, the results of [JKKL20, LCSJ22] imply that γ -observable POMDPs can be learned with polynomially many samples, albeit by computationally inefficient algorithms. Thus, as discussed following Theorem 3.1, it is natural to wonder whether we can achieve the best of both worlds: is there a quasipolynomial-time algorithm that only needs polynomially many samples, or can one show a computational-statistical tradeoff?
- It is also natural to ask whether an analogue of Assumption 1.1 for the ℓ_2 norm, namely that $\|\mathbb{O}_h b - \mathbb{O}_h b'\|_2 \geq \gamma \|b - b'\|_2$ for all $h \in [H]$, is sufficient for computationally efficient learnability. Even the planning version of this question (where the parameters of the POMDP are known and the problem is to find a near-optimal policy) is open.

Acknowledgments and Disclosure of Funding

N.G. is supported by a Fannie & John Hertz Foundation Fellowship and an NSF Graduate Fellowship. A.M. is supported by a Microsoft Trustworthy AI Grant, NSF Large CCF-1565235, a David and Lucile Packard Fellowship and an ONR Young Investigator Award. D.R. is supported by an Akamai Presidential Fellowship and a U.S. DoD NDSEG Fellowship.

References

- [AGH⁺14] Animashree Anandkumar, Rong Ge, Daniel J Hsu, Sham M Kakade, and Matus Telgarsky. Tensor decompositions for learning latent variable models. *Journal of Machine Learning Research*, 15(1):2773–2832, 2014.
- [AHKS20] Alekh Agarwal, Mikael Henaff, Sham Kakade, and Wen Sun. Pc-pg: Policy cover directed exploration for provable policy gradient learning. *Advances in Neural Information Processing Systems*, 33:13399–13412, 2020.
- [AK08] Baruch Awerbuch and Robert Kleinberg. Online linear optimization and adaptive routing. *Journal of Computer and System Sciences*, 74(1):97–114, 2008.

- [ALA16] Kamyar Azizzadenesheli, Alessandro Lazaric, and Animashree Anandkumar. Open problem: Approximate planning of pomdps in the class of memoryless policies. In *Conference on Learning Theory*, pages 1639–1642. PMLR, 2016.
- [AYA18] Kamyar Azizzadenesheli, Yisong Yue, and Animashree Anandkumar. Policy gradient in partially observable environments: Approximation and convergence. *arXiv preprint arXiv:1810.07900*, 2018.
- [BDRS96] Dima Burago, Michel De Rougemont, and Anatol Slissenko. On the complexity of partially observed markov decision processes. *Theoretical Computer Science*, 157(2):161–183, 1996.
- [BS18] Noam Brown and Tuomas Sandholm. Superhuman ai for heads-up no-limit poker: Libratus beats top professionals. *Science*, 359(6374):418–424, 2018.
- [Can20] Clément L Canonne. A short note on learning discrete distributions. *arXiv preprint arXiv:2002.11457*, 2020.
- [CJJ⁺21] Murat Cubuktepe, Nils Jansen, Sebastian Junges, Ahmadreza Marandi, Marnix Suilen, and Ufuk Topcu. Robust finite-state controllers for uncertain pomdps. In *AAAI*, pages 11792–11800. AAAI Press, 2021.
- [CJJT22] Steven Carr, Nils Jansen, Sebastian Junges, and Ufuk Topcu. Safe reinforcement learning via shielding for pomdps. *CoRR*, abs/2204.00755, 2022.
- [CKK96] Anthony R Cassandra, Leslie Pack Kaelbling, and James A Kurien. Acting under uncertainty: Discrete bayesian models for mobile-robot navigation. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems. IROS'96*, volume 2, pages 963–972. IEEE, 1996.
- [CLC09] Chenghui Cai, Xuejun Liao, and Lawrence Carin. Learning to explore and exploit in pomdps. In Y. Bengio, D. Schuurmans, J. Lafferty, C. Williams, and A. Culotta, editors, *Advances in Neural Information Processing Systems*, volume 22. Curran Associates, Inc., 2009.
- [CLZ97] Anthony Cassandra, Michael L. Littman, and Nevin L. Zhang. Incremental pruning: A simple, fast, exact method for partially observable markov decision processes. In *Proceedings of the Thirteenth Conference on Uncertainty in Artificial Intelligence, UAI'97*, page 54–61, San Francisco, CA, USA, 1997. Morgan Kaufmann Publishers Inc.
- [DGZ22] Constantinos Daskalakis, Noah Golowich, and Kaiqing Zhang. The complexity of markov equilibrium in stochastic games. *arXiv preprint arXiv:2204.03991*, 2022.
- [DKJ⁺19] 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.
- [DLB17] Christoph Dann, Tor Lattimore, and Emma Brunskill. Unifying pac and regret: Uniform pac bounds for episodic reinforcement learning. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- [DVPWR15] Finale Doshi-Velez, David Pfau, Frank Wood, and Nicholas Roy. Bayesian nonparametric methods for partially-observable reinforcement learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37:394–407, 2015.
- [EDKM05] Eyal Even-Dar, Sham M Kakade, and Yishay Mansour. Reinforcement learning in pomdps without resets. 2005.
- [EDKM07] Eyal Even-Dar, Sham M Kakade, and Yishay Mansour. The value of observation for monitoring dynamic systems. In *IJCAI*, pages 2474–2479, 2007.

- [EJKM22] Yonathan Efroni, Chi Jin, Akshay Krishnamurthy, and Sobhan Miryoosefi. Provable reinforcement learning with a short-term memory. *CoRR*, abs/2202.03983, 2022.
- [FKQR21] Dylan J Foster, Sham M Kakade, Jian Qian, and Alexander Rakhlin. The statistical complexity of interactive decision making. *arXiv preprint arXiv:2112.13487*, 2021.
- [GDB16] Zhaohan Daniel Guo, Shayan Doroudi, and Emma Brunskill. A pac rl algorithm for episodic pomdps. In *Artificial Intelligence and Statistics*, pages 510–518. PMLR, 2016.
- [GHL19] Neha Priyadarshini Garg, David Hsu, and Wee Sun Lee. Despot-alpha: Online pomdp planning with large state and observation spaces. In *Robotics: Science and Systems*, 2019.
- [GMR22] Noah Golowich, Ankur Moitra, and Dhruv Rohatgi. Planning in observable pomdps in quasipolynomial time. *arXiv preprint arXiv:2201.04735*, 2022.
- [Han98] Eric A Hansen. Solving pomdps by searching in policy space. In *Proceedings of the Fourteenth conference on Uncertainty in artificial intelligence*, pages 211–219, 1998.
- [Hau00] Milos Hauskrecht. Value-function approximations for partially observable markov decision processes. *Journal of artificial intelligence research*, 13:33–94, 2000.
- [HDT20] Dongqi Han, Kenji Doya, and Jun Tani. Variational recurrent models for solving partially observable control tasks. In *ICLR*, 2020.
- [HF00a] Eric A. Hansen and Zhengzhu Feng. Dynamic programming for pomdps using a factored state representation. In *In Proceedings of the Fifth International Conference on AI Planning Systems*, pages 130–139, 2000.
- [HF00b] Milos Hauskrecht and Hamish Fraser. Planning treatment of ischemic heart disease with partially observable markov decision processes. *Artificial intelligence in medicine*, 18(3):221–244, 2000.
- [HKZ12] Daniel Hsu, Sham M. Kakade, and Tong Zhang. A spectral algorithm for learning hidden markov models. *Journal of Computer and System Sciences*, 78(5):1460–1480, 2012. JCSS Special Issue: Cloud Computing 2011.
- [JN21] Mehdi Jafarnia-Jahromi, Rahul Jain, and Ashutosh Nayyar. Online learning for unknown partially observable mdps. *CoRR*, abs/2102.12661, 2021.
- [JJW⁺17] Sebastian Junges, Nils Jansen, Ralf Wimmer, Tim Quatmann, Leonore Winterer, Joost-Pieter Katoen, and Bernd Becker. Permissive finite-state controllers of pomdps using parameter synthesis. *CoRR*, abs/1710.10294, 2017.
- [JKA⁺17] Nan Jiang, Akshay Krishnamurthy, Alekh Agarwal, John Langford, and Robert E. Schapire. Contextual decision processes with low Bellman rank are PAC-learnable. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1704–1713. PMLR, 06–11 Aug 2017.
- [JKKL20] Chi Jin, Sham M Kakade, Akshay Krishnamurthy, and Qinghua Liu. Sample-efficient reinforcement learning of undercomplete pomdps. *arXiv preprint arXiv:2006.12484*, 2020.
- [JKSY20] Chi Jin, Akshay Krishnamurthy, Max Simchowitz, and Tiancheng Yu. Reward-free exploration for reinforcement learning. In *International Conference on Machine Learning (ICML)*, pages 4870–4879. PMLR, 2020.
- [KAL16] Akshay Krishnamurthy, Alekh Agarwal, and John Langford. Pac reinforcement learning with rich observations. *arXiv preprint arXiv:1602.02722*, 2016.
- [KECM21a] Jeongyeol Kwon, Yonathan Efroni, Constantine Caramanis, and Shie Mannor. Reinforcement learning in reward-mixing mdps. *arXiv preprint arXiv:2110.03743*, 2021.

- [KECM21b] Jeongyeol Kwon, Yonathan Efroni, Constantine Caramanis, and Shie Mannor. RI for latent mdps: Regret guarantees and a lower bound. *arXiv preprint arXiv:2102.04939*, 2021.
- [KLLM22] Daniel Kane, Sihan Liu, Shachar Lovett, and Gaurav Mahajan. Computational-statistical gaps in reinforcement learning. *arXiv preprint arXiv:2202.05444*, 2022.
- [KMN99] Michael J Kearns, Yishay Mansour, and Andrew Y Ng. Approximate planning in large pomdps via reusable trajectories. In *NIPS*, pages 1001–1007. Citeseer, 1999.
- [KOA19] Sammie Katt, Frans A. Oliehoek, and Christopher Amato. Bayesian reinforcement learning in factored pomdps. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS '19*, page 7–15, Richland, SC, 2019. International Foundation for Autonomous Agents and Multiagent Systems.
- [KS02] Michael Kearns and Satinder Singh. Near-optimal reinforcement learning in polynomial time. *Machine learning*, 49(2):209–232, 2002.
- [LCSJ22] Qinghua Liu, Alan Chung, Csaba Szepesvári, and Chi Jin. When is partially observable reinforcement learning not scary? *arXiv preprint arXiv:2204.08967*, 2022.
- [LGM01] Christopher Lusena, Judy Goldsmith, and Martin Mundhenk. Nonapproximability results for partially observable markov decision processes. *Journal of artificial intelligence research*, 14:83–103, 2001.
- [Lit94] Michael L Littman. Memoryless policies: Theoretical limitations and practical results. *From animals to animats*, 3:238–245, 1994.
- [LLC09] Hui Li, Xuejun Liao, and Lawrence Carin. Multi-task reinforcement learning in partially observable stochastic environments. *J. Mach. Learn. Res.*, 10:1131–1186, jun 2009.
- [LLC11] Miao Liu, Xuejun Liao, and Lawrence Carin. The Infinite Regionalized Policy Representation. In *Proceedings of the 28th International Conference on Machine Learning*, pages 769–776. Omnipress, 2011.
- [LLC13] Miao Liu, Xuejun Liao, and Lawrence Carin. Online expectation maximization for reinforcement learning in pomdps. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence, IJCAI '13*, page 1501–1507. AAAI Press, 2013.
- [LS17] Tor Lattimore and Csaba Szepesvari. The end of optimism? an asymptotic analysis of finite-armed linear bandits. In *Artificial Intelligence and Statistics*, pages 728–737. PMLR, 2017.
- [LYX11] Yanjie Li, Baoqun Yin, and Hongsheng Xi. Finding optimal memoryless policies of pomdps under the expected average reward criterion. *European Journal of Operational Research*, 211(3):556–567, 2011.
- [MGK21] Lingheng Meng, Rob Gorbet, and Dana Kuli’c. Memory-based deep reinforcement learning for pomdp. *ArXiv*, abs/2102.12344, 2021.
- [MHKL20] Dipendra Misra, Mikael Henaff, Akshay Krishnamurthy, and John Langford. Kinematic state abstraction and provably efficient rich-observation reinforcement learning. 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 6961–6971. PMLR, 13–18 Jul 2020.
- [MKKC99] Nicolas Meuleau, Kee-Eung Kim, Leslie Pack Kaelbling, and Anthony R. Cassandra. Solving pomdps by searching the space of finite policies. In *Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence, UAI'99*, page 417–426, San Francisco, CA, USA, 1999. Morgan Kaufmann Publishers Inc.

- [Mon82] George E. Monahan. State of the art-a survey of partially observable markov decision processes: Theory, models, and algorithms. *Manage. Sci.*, 28(1):1–16, jan 1982.
- [MR05] Elchanan Mossel and Sébastien Roch. Learning nonsingular phylogenies and hidden markov models. *CoRR*, abs/cs/0502076, 2005.
- [NES22] Tianwei Ni, Benjamin Eysenbach, and Ruslan Salakhutdinov. Recurrent model-free rl can be a strong baseline for many pomdps, 2022.
- [PB04] Pascal Poupart and Craig Boutilier. Vdcbpi: an approximate scalable algorithm for large pomdps. In *NIPS*, pages 1081–1088. Citeseer, 2004.
- [PGT06] Joelle Pineau, Geoffrey Gordon, and Sebastian Thrun. Anytime point-based approximations for large pomdps. *Journal of Artificial Intelligence Research*, 27:335–380, 2006.
- [PT87] Christos H Papadimitriou and John N Tsitsiklis. The complexity of markov decision processes. *Mathematics of operations research*, 12(3):441–450, 1987.
- [PV08] Pascal Poupart and Nikos Vlassis. Model-based bayesian reinforcement learning in partially observable domains. In *Proc Int. Symp. on Artificial Intelligence and Mathematics*,, pages 1–2, 2008.
- [RCdP07] Stephane Ross, Brahim Chaib-draa, and Joelle Pineau. Bayes-adaptive pomdps. In J. Platt, D. Koller, Y. Singer, and S. Roweis, editors, *Advances in Neural Information Processing Systems*, volume 20. Curran Associates, Inc., 2007.
- [RG02] Nicholas Roy and Geoffrey J Gordon. Exponential family pca for belief compression in pomdps. *Advances in Neural Information Processing Systems*, 15:1667–1674, 2002.
- [RPCdK11] Stéphane Ross, Joelle Pineau, Brahim Chaib-draa, and Pierre Kreitmann. A bayesian approach for learning and planning in partially observable markov decision processes. *J. Mach. Learn. Res.*, 12(null):1729–1770, jul 2011.
- [RPPCD08] Stéphane Ross, Joelle Pineau, Sébastien Paquet, and Brahim Chaib-Draa. Online planning algorithms for pomdps. *Journal of Artificial Intelligence Research*, 32:663–704, 2008.
- [SBS05] Guy Shani, Ronen I. Brafman, and Solomon E. Shimony. Model-based online learning of pomdps. In João Gama, Rui Camacho, Pavel B. Brazdil, Alípio Mário Jorge, and Luís Torgo, editors, *Machine Learning: ECML 2005*, pages 353–364, Berlin, Heidelberg, 2005. Springer Berlin Heidelberg.
- [Sch90] Jürgen Schmidhuber. Reinforcement learning in markovian and non-markovian environments. In *Proceedings of the 3rd International Conference on Neural Information Processing Systems*, NIPS’90, page 500–506, San Francisco, CA, USA, 1990. Morgan Kaufmann Publishers Inc.
- [SS12] Trey Smith and Reid Simmons. Heuristic search value iteration for pomdps. *arXiv preprint arXiv:1207.4166*, 2012.
- [SV05] Matthijs TJ Spaan and Nikos Vlassis. Perseus: Randomized point-based value iteration for pomdps. *Journal of artificial intelligence research*, 24:195–220, 2005.
- [SV10] David Silver and Joel Veness. Monte-carlo planning in large pomdps. *Neural Information Processing Systems*, 2010.
- [SYHL13] Adhiraj Somani, Nan Ye, David Hsu, and Wee Sun Lee. Despot: Online pomdp planning with regularization. *Advances in neural information processing systems*, 26:1772–1780, 2013.
- [TK03] Georgios Theodorou and Leslie Kaelbling. Approximate planning in pomdps with macro-actions. *Advances in Neural Information Processing Systems*, 16:775–782, 2003.

- [VLB12] Nikos Vlassis, Michael L. Littman, and David Barber. On the computational complexity of stochastic controller optimization in pomdps. *ACM Trans. Comput. Theory*, 4(4), nov 2012.
- [WAJ⁺21] Gellert Weisz, Philip Amortila, Barnabás Janzer, Yasin Abbasi-Yadkori, Nan Jiang, and Csaba Szepesvári. On query-efficient planning in mdps under linear realizability of the optimal state-value function. In *Conference on Learning Theory*, pages 4355–4385. PMLR, 2021.
- [WDZ22] Chen-Yu Wei, Christoph Dann, and Julian Zimmert. A model selection approach for corruption robust reinforcement learning. In Sanjoy Dasgupta and Nika Haghtalab, editors, *Proceedings of The 33rd International Conference on Algorithmic Learning Theory*, volume 167 of *Proceedings of Machine Learning Research*, pages 1043–1096. PMLR, 29 Mar–01 Apr 2022.
- [WL21] Chen-Yu Wei and Haipeng Luo. Non-stationary reinforcement learning without prior knowledge: An optimal black-box approach. *CoRR*, abs/2102.05406, 2021.
- [WME⁺22] Xintong Wang, Gary Qiurui Ma, Alon Eden, Clara Li, Alexander Trott, Stephan Zheng, and David C. Parkes. Using reinforcement learning to study platform economies under market shocks, 2022.
- [WYDW21] Tianhao Wu, Yunchang Yang, Simon Du, and Liwei Wang. On reinforcement learning with adversarial corruption and its application to block mdp. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 11296–11306. PMLR, 18–24 Jul 2021.
- [XCGZ21] Yi Xiong, Ningyuan Chen, Xuefeng Gao, and Xiang Zhou. Sublinear regret for learning pomdps. *arXiv preprint arXiv:2107.03635*, 2021.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [Yes]
 - (b) Did you describe the limitations of your work? [Yes]
 - (c) Did you discuss any potential negative societal impacts of your work? [N/A]
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [Yes]
 - (b) Did you include complete proofs of all theoretical results? [Yes]
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [N/A]
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [N/A]
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A]
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [N/A]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [N/A]
 - (b) Did you mention the license of the assets? [N/A]

- (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
- (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

A Additional Related Work

Learning POMDPs (without computational efficiency). Over the past several decades several algorithmic paradigms have been developed for the problem of learning the optimal policy in POMDPs. [LLC09] introduced the *regionalized policy representation* (later developed in [LLC11, LLC13, CLC09]), a model-free Bayesian method which attempts to learn the optimal policy directly by learning a model which aggregates belief states into regions and then marginalizing out over the regions. Several papers [PV08, RCdP07, RPCdK11, KOA19] have also applied model-based Bayesian methods to learning POMDPs, by augmenting the state space with empirical estimates of transition and observation probabilities and redefining the belief state updates in the natural way. See also [DVPWR15] for related Bayesian methods. [AYA18] adapts policy gradient methods to the problem of policy optimization in POMDPs. Various other heuristics have been considered as well (e.g., [SBS05, CJJ⁺21, CJJT22, JJW⁺17]). More recently, several works have used recurrent neural networks to parametrize history-dependent policies and value functions for learning POMDPs [Sch90, NES22, MGK21, HDT20]. Essentially all of these results come with no formal guarantees (i.e., achieving PAC bounds computationally efficiently) regarding optimality of the learned policies.

The lack of provable guarantees for learning POMDPs without additional assumptions is unavoidable: [JKKL20, KAL16] show that exponentially many samples are needed to learn an approximately optimal policy in general POMDPs. Furthermore, [MR05] proved the related result that learning the parameters of an HMM is as hard as learning parity with noise. If given exponentially many samples, however, it is possible to prove positive results: in the infinite-horizon nonepisodic setting, [EDKM05] introduces an algorithm which provably outputs a near-optimal policy in time (and sample complexity) scaling exponentially in a certain parameter which depends on a horizon-like quantity of the POMDP.

Closer to our work, by placing appropriate assumptions on the POMDP, many works have proven polynomial sample complexity bounds. Several works [GDB16, ALA16, XCGZ21] have used *spectral methods* (see [HKZ12, AGH⁺14]) to learn POMDPs with polynomial sample complexity; while, in a sense, these papers do address the question of exploration, they do so by making strong assumptions on the underlying *transitions*, namely that the transitions must be full-rank and furthermore must satisfy a mixing-type condition which implies that the Markov chain induced by any policy mixes quickly. Such assumptions on the transitions were removed in the recent work [JKKL20], which modified the technique of spectral methods to obtain polynomial sample complexity for learning POMDPs that satisfy an ℓ_2 version of Assumption 1.1: in particular, their sample complexity bound depends polynomially on the inverse of the smallest singular value of any of the observation matrices. Later work [LCSJ22] presented a simpler algorithm for learning such POMDPs, based on maximum likelihood estimation; further, the result of [LCSJ22] applies also to the more general class of *weakly revealing POMDPs*, which includes the *overcomplete* case, namely where there are more states than observations. [JJN21] applies the technique of posterior sampling to learn POMDPs in the Bayesian setting, though with the strong assumption that either different models to be learned are separated in KL divergence, or that the models sampled in the course of the algorithm are good estimates of the true model. [EJKM22] establishes polynomial sample complexity bounds in the case when the latent state is *decodable* from a small suffix of the history. [KECM21b] shows that latent MDPs, a special case of overcomplete POMDPs, are learnable when the individual MDPs have transitions that are well-separated. We emphasize that all of the above works do *not* provide end-to-end computational guarantees on learning POMDPs.

Learning POMDPs (with computational efficiency). Prior to our work, the only computationally efficient learning algorithms for POMDPs applied to very special cases. [JKKL20, KAL16] proved polynomial-time learning results assuming deterministic transitions of the POMDP. [KECM21a] obtains results for learning latent MDPs which are a mixture of 2 underlying models; since the uncertainty in the system can be modeled by a single-dimensional parameter, their results are computationally efficient.

Planning algorithms in RL. The problem of *planning* in POMDPs (namely, finding the optimal policy when the model is known) has been extensively studied, with various proposed heuristics (e.g., [Mon82, CLZ97, HF00a, Hau00, RG02, TK03, PB04, SV05, PGT06, RPPCD08, SV10, SS12, SYHL13, GHL19, Han98, MKKC99, KMN99, LYX11, AYA18]), and a few provably efficient

algorithms [BDRS96, KECM21a, GMR22]. Most closely related to our work is [GMR22], which shows a quasipolynomial-time planning algorithm for observable POMDPs.

Computational-statistical considerations in RL. Several recent works have shown that certain learning problems in RL can be solved with few (i.e., polynomial) samples yet have superpolynomial *computational* complexity, under standard assumptions. This is the case for finding a near-optimal policy in γ -observable POMDPs with $\gamma = 1/\text{poly}(O)$, which requires exponential time under ETH [GMR22] yet can be learned with polynomially many samples [JKKL20]; finding stationary coarse correlated equilibria in general-sum Markov games [DGZ22]; and finding the optimal policy in MDPs which have linear Q^* and V^* functions [KLLM22]. The latter result establishes a computational-statistical gap of stronger nature in the sense that (unlike the previous two examples) the planning version of the problem can be solved efficiently. In contrast to the above lower bounds for computationally efficient learning, Theorem 3.1 shows that learning γ -observable POMDPs with constant γ can be done computationally (quasi)-efficiently.

As a whole, the issue of developing computationally efficient algorithms for RL problems is wide open, with a plethora of intriguing open questions (see, e.g., [WAJ⁺21, JKA⁺17, FKQR21]).

Nonstationary MDPs and related approaches. In light of the fact that the dynamics of a γ -observable POMDP \mathcal{P} can be approximated by those of an MDP \mathcal{M} with quasipolynomially-sized state space (see Section 4.1), a natural approach to learn an optimal policy of \mathcal{P} is to use one of many recent results (e.g., [WL21, WYDW21, WDZ22]) which establishes regret bounds for learning in MDPs when an adversary is allowed to corrupt the trajectories for some episodes. Such an adversary may be used to model the fact that we can simulate trajectory draws from \mathcal{P} by viewing them as trajectory draws from \mathcal{M} with a certain fraction of episodes being corrupted. In particular, [WYDW21] shows that if C episodes (out of T total episodes) are corrupted, then there is an efficient algorithm with regret upper bounded by $\text{poly}(H) \cdot \tilde{O}(\sqrt{ZAT} + Z^2A + CZA)$, where Z denotes the number of states of the MDP \mathcal{M} . In our application, we would have $Z = (OA)^L$ (with L defined as in Theorem 3.1) and $C \geq \alpha \cdot T$ (as Theorem 3.1 gives an upper bound on the probability of an episode being corrupted which is no smaller than α). Thus $CZA \geq \alpha T \cdot (OA)^{\log(1/\alpha)} \gg T$, meaning that the bound from [WYDW21] is vacuous. Furthermore, this is not just an artifact of their analysis: [WYDW21] showed a lower bound of $\Omega(CZA)$ on the regret when an adversary is allowed to corrupt C episodes. Thus, in a sense, the adversarial model studied by [WL21, WYDW21, WDZ22] is too strong to prove nonvacuous results in our setting.

Another related approach is the PC-IGW.Bilinear algorithm (as well as its IGW-Spanner subroutine) of [FKQR21], which use similar techniques (namely, a policy cover and a barycentric spanner) to ours. Moreover, since POMDPs fit into the framework of *decision-making with structured observations* of [FKQR21], one can attempt to use the PC-IGW.Bilinear algorithm together with E2D [FKQR21] to attempt to efficiently learn the optimal policy in POMDPs. However, there is a significant issue: the full E2D meta-algorithm of [FKQR21] requires a *model-estimation oracle* (the output of which is passed to PC-IGW.Bilinear). One might hope that we can implement this oracle efficiently by using subtrajectories of length L to estimate the transitions of the MDP \mathcal{M} that approximates \mathcal{P} . However (using the notation of the previous paragraph) the estimation error would necessarily scale as αT . Furthermore, in order to ensure that the output of the estimation oracle belongs to the class of models, the model class must be augmented to contain all MDPs on the state space \mathcal{Z} ; doing so leads to a bound on the decision estimation coefficient which is at least $\frac{Z}{\gamma}$. The resulting regret bound of [FKQR21, Theorem 4.1] scales as $\inf_{\gamma>0} \frac{TZ}{\gamma} + \gamma \cdot \mathbf{Est}$, where $\mathbf{Est} \geq \alpha T$ denotes the estimation error. This expression is minimized when $\gamma = \sqrt{TZ/\mathbf{Est}}$, and yields a regret bound which is at least $\sqrt{TZ \cdot \mathbf{Est}} \geq T \cdot \sqrt{\alpha \cdot (OA)^{\log 1/\alpha}} \gg T$, which is again vacuous. Finding a general model of learning in “approximate MDPs” which yields nonvacuous regret bounds in our setting is an interesting open problem.

We also remark that the use of barycentric spanners in [FKQR21] differs significantly from ours. In [FKQR21], a weaker notion, known as *approximate G-optimal design*, suffices for their application, yet in order to compute such a G-optimal design efficiently in their setting, it turns out to be more convenient to search for the stronger notion of barycentric spanner. Replacing the barycentric spanner in our algorithm (Algorithm 2) with a G-optimal design does not seem to suffice for our needs.

B Additional preliminaries

In this section we state some additional preliminaries which are useful in the proof and formalize some notions which were introduced informally in the main body.

Extending the state and observation spaces. Fix a POMDP $\mathcal{P} = (\mathcal{S}, \mathcal{A}, \mathcal{O}, H, b_1, R, \mathbb{T}, \mathbb{O})$, as in Theorem 3.1. For technical reasons, we will augment the state space \mathcal{S} with a special “sink state” s^{sink} , and augment the observation space \mathcal{O} with a special “sink observation” o^{sink} . We write $\bar{\mathcal{S}} := \mathcal{S} \cup \{s^{\text{sink}}\}$ and $\bar{\mathcal{O}} := \mathcal{O} \cup \{o^{\text{sink}}\}$, and call the spaces $\bar{\mathcal{S}}, \bar{\mathcal{O}}$, the *extended state and observation spaces*, respectively. We extend the various components of \mathcal{P} to the extended state and observation sequences as follows: b_1 remains unchanged, $R_h(o^{\text{sink}}) = 0$ for all h , $\mathbb{T}_h^{\mathcal{P}}(s^{\text{sink}}|s^{\text{sink}}, a) = 1$ for all $a \in \mathcal{A}, h \in [H]$ (and all transitions from states $s \in \mathcal{S}$ remain unchanged), and $\mathbb{O}_h^{\mathcal{P}}(o^{\text{sink}}|s^{\text{sink}}) = 1$ for all $h \in [H]$ (and all observation distributions at states $s \in \mathcal{S}$ remain unchanged). It is clear that the distribution over trajectories given any fixed policy are the same in the original and extended POMDP, and the extended POMDP satisfies γ -observability if the original POMDP does. Thus, it is without loss of generality to replace \mathcal{P} with the extended POMDP $\mathcal{P} = (\bar{\mathcal{S}}, \mathcal{A}, \bar{\mathcal{O}}, H, b_1, R, \mathbb{T}, \mathbb{O})$, as we shall do in the entirety of the proof (we do so since we will introduce other POMDPs which, unlike \mathcal{P} , will at times transition into s^{sink} , unlike the original POMDP \mathcal{P} ; see Section E.3). Thus, we have that the belief states $\mathbf{b}_h^{\mathcal{P}}(a_{1:h-1}, o_{2:h})$, as well as the approximate belief states $\mathbf{b}_h^{\text{apx}, \mathcal{P}}(a_{h-L:h-1}, o_{h-L+1:h}; \mathcal{D})$, belong to $\mathbb{R}^{\bar{\mathcal{S}}}$. Moreover, it is immediate from the block structure of $\mathbb{O}_h^{\mathcal{P}} \in \mathbb{R}^{\bar{\mathcal{O}} \times \bar{\mathcal{S}}}$ that its pseudoinverse $(\mathbb{O}_h^{\mathcal{P}})^\dagger \in \mathbb{R}^{\bar{\mathcal{S}} \times \bar{\mathcal{O}}}$ has the same block structure, i.e., $(\mathbb{O}_h^{\mathcal{P}})^\dagger_{s^{\text{sink}}, o} = (\mathbb{O}_h^{\mathcal{P}})^\dagger_{s, o^{\text{sink}}} = 0$ for all $s \in \mathcal{S}, o \in \mathcal{O}$.

Finally, we remark that for some choices of $a_{1:h-1}, o_{2:h}$, the belief state $\mathbf{b}_h^{\mathcal{P}}(a_{1:h-1}, o_{2:h})$ or the approximate belief state $\mathbf{b}_h^{\text{apx}, \mathcal{P}}(a_{h-L:h-1}, o_{h-L+1:h})$ may be undefined (for instance, if $o_{h-L+1:h} \notin \bar{\mathcal{O}}^L$ and $\mathcal{D}(s^{\text{sink}}) = 0$, then $\mathbf{b}_h^{\text{apx}, \mathcal{P}}(a_{h-L:h-1}, o_{h-L+1:h}; \mathcal{D})$ is undefined since the transitions of \mathcal{P} never lead to s^{sink}). In such a case, as a matter of convention, we will define the belief state to put all its mass on s^{sink} .

MDPs with Z-structure. At times, we will write $\mathbb{T}^{\mathcal{P}}, \mathbb{O}^{\mathcal{P}}, R^{\mathcal{P}}$ in place of $\mathbb{T}, \mathbb{O}, R$, respectively, to clarify that the corresponding transition matrices, observation matrices, and reward function are for the POMDP \mathcal{P} . Furthermore, we will consider a fixed value $L \in \mathbb{N}$ (the particular value of L , as a function of the parameters of the problem, is given in Section C.1; it is chosen so as to satisfy the requirements of the belief contraction theorem (Theorem B.2)). Given \mathcal{P} , write $\mathcal{Z} := \mathcal{A}^L \times \mathcal{O}^L$ and $\bar{\mathcal{Z}} := \mathcal{A} \times \bar{\mathcal{O}}^L$; the spaces $\mathcal{Z}, \bar{\mathcal{Z}}$ should be interpreted as representing a sequence of L consecutive action-observation pairs for a trajectory drawn from \mathcal{P} . Throughout the course of the proof, we will consider several *Markov Decision Processes* (MDPs) whose state space is $\bar{\mathcal{Z}}$. In particular, for various choices of the probability transition matrices $\mathbb{P}_h(z_{h+1}|z_h, a_h)$ (for $z_h, z_{h+1} \in \bar{\mathcal{Z}}$) and reward functions $R = (R_1, \dots, R_H), R_h : \bar{\mathcal{Z}} \rightarrow \mathbb{R}$, we will consider MDPs of the form $\mathcal{M} = (\bar{\mathcal{Z}}, \mathcal{A}, H, R, \mathbb{P})$ (we will typically leave the initial state distribution unspecified, as it does not affect the optimal policy of an MDP).

For such an \mathcal{M} , we say that \mathcal{M} has *Z-structure* if: (a) $R_h(z)$ is only a function of $o(z)$ for all $h \in [H]$, and (b) the transitions $\mathbb{P}_h(\cdot|z_h, a_h)$ have the following property, for $z_h \in \bar{\mathcal{Z}}, a_h \in \mathcal{A}$: writing $z_h = (a_{h-L:h-1}, o_{h-L+1:h})$, $\mathbb{P}_h(z_{h+1}|z_h, a_h)$ is nonzero only for those z_{h+1} of the form $z_{h+1} = (a_{h-L+1:h}, o_{h-L+2:h+1})$, where $o_{h+1} \in \bar{\mathcal{O}}$. Throughout this section, the MDPs with state space $\bar{\mathcal{Z}}$ that we consider all will have Z-structure.

Throughout, we will use the following notational convention: given a sequence of L action-observation pairs for the POMDP \mathcal{P} , namely $(a_{h-L:h-1}, o_{h-L+1:h}) \in (\mathcal{A} \times \bar{\mathcal{O}})^L$, we will denote this sequence by $z_h = (a_{h-L:h-1}, o_{h-L+1:h})$. When the identity of the action-observation sequence is clear, we will use z_h and $(a_{h-L:h-1}, o_{h-L+1:h})$ interchangeably. Furthermore, for a sequence $z_h = (a_{h-L:h-1}, o_{h-L+1:h})$, write $o(z_h) := o_h$ to denote the final observation of this sequence.

We define $\Pi_{\mathcal{Z}}^{\text{markov}}$ to be the set of sequences $\pi = (\pi_1, \dots, \pi_H)$, where each $\pi_h : \bar{\mathcal{Z}} \rightarrow \mathcal{A}$. $\Pi_{\mathcal{Z}}^{\text{markov}}$ may be identified as a subset of Π^{det} : a policy $\pi = (\pi_1, \dots, \pi_H) \in \Pi_{\mathcal{Z}}^{\text{markov}}$ may be seen as a history-dependent deterministic policy via the mapping $(a_{1:h-1}, o_{2:h}) \mapsto \pi_h(a_{h-L:h-1}, o_{h-L+1:h})$

(i.e., by ignoring the portion of the history prior to step $h - L$). Given an MDP \mathcal{M} with \mathbf{Z} -structure and a policy $\pi \in \Pi_{\mathbf{Z}}^{\text{markov}}$, the definition of a trajectory drawn from \mathcal{M} by following policy π is standard (see Section 2.1). We next extend this definition to the case of a *general policy* $\pi \in \Pi^{\text{gen}}$, for which a trajectory $z_{1:H}$ (which is equivalent to the data of $a_{1:H-1}, o_{2:H}$) is drawn as follows: first, a deterministic policy $\sigma \in \Pi^{\text{det}}$ is drawn, $\sigma \sim \pi$, and then, at each step $h \in [H]$, given history $z_{1:h}$ (which is equivalent to the data of $a_{1:h-1}, o_{2:h}$, since \mathcal{M} has \mathbf{Z} -structure), we take action $\sigma_h(a_{1:h-1}, o_{2:h})$. At times, we will abuse notation and write $\sigma_h(z_{1:h}) = \sigma_h(a_{1:h-1}, o_{2:h})$. Additionally, we write $\mathbb{P}_{z_{1:H} \sim \pi}^{\mathcal{M}}(\mathcal{E})$ to denote the probability of \mathcal{E} when $z_{1:H}$ (which is equivalent to the data of $a_{1:H-1}, o_{2:H}$) is drawn from a trajectory following policy π for the MDP \mathcal{M} . In similar spirit, we will write $\mathbb{E}_{z_{1:H} \sim \pi}^{\mathcal{M}}[\cdot]$ to denote expectations.

Visitation distributions. Next, for an element $z = (a_{1:L}, o_{1:L}) \in \bar{\mathcal{Z}}$, as well as $0 \leq h$, write $\text{suff}_h(z) := (a_{\max\{1, h-L+1\}:L}, o_{\max\{1, h-L+1\}:L}) \in \mathcal{A}^{\min\{h, L\}} \times \bar{\mathcal{O}}^{\min\{h, L\}}$ to denote the suffix of z that consists of the last h actions and observations in z , or all of z if $h > L$. Recall from Section 2.4 the following notation for visitation distributions for the POMDP \mathcal{P} and an MDP \mathcal{M} with \mathbf{Z} -structure:

- Given a general policy $\pi \in \Pi^{\text{gen}}$, $h \in [H]$, and $s \in \bar{\mathcal{S}}$, define $d_{\mathcal{S}, h}^{\mathcal{P}, \pi}(s) = \mathbb{P}_{s_{1:H} \sim \pi}^{\mathcal{P}}(s_h = s)$.
- For a general policy $\pi \in \Pi^{\text{gen}}$, $h \in [H]$, and $o \in \bar{\mathcal{O}}$, define $d_{\mathcal{O}, h}^{\mathcal{P}, \pi}(o) = \mathbb{P}_{o_{1:H} \sim \pi}^{\mathcal{P}}(o_h = o)$.
- For a general policy $\pi \in \Pi^{\text{gen}}$, $h \in [H]$, and $z \in \bar{\mathcal{Z}}$, define $d_{\mathcal{Z}, h}^{\mathcal{P}, \pi}(z) = \mathbb{P}_{a_{1:H}, o_{2:H} \sim \pi}^{\mathcal{P}}(\text{suff}_{h-1}(a_{h-L:h-1}, o_{h-L+1:h}) = \text{suff}_{h-1}(z))$.
- For a general policy $\pi \in \Pi^{\text{gen}}$, $h \in [H]$, and $z \in \bar{\mathcal{Z}}$, define $d_{\mathcal{Z}, h}^{\mathcal{M}, \pi}(z) = \mathbb{P}_{z_{1:H} \sim \pi}^{\mathcal{M}}(\text{suff}_{h-1}(z_h) = \text{suff}_{h-1}(z))$.
- For a general policy $\pi \in \Pi^{\text{gen}}$, $h \in [H]$, and $o \in \bar{\mathcal{O}}$, define $d_{\mathcal{O}, h}^{\mathcal{M}, \pi}(o) = \mathbb{P}_{z_{1:H} \sim \pi}^{\mathcal{M}}(o(z_h) = o)$.

Notice that, for $h \leq L$, $d_{\mathcal{Z}, h}^{\mathcal{P}, \pi}$ and $d_{\mathcal{Z}, h}^{\mathcal{M}, \pi}$ are not distributions, i.e., we will have $\sum_{z \in \bar{\mathcal{Z}}} d_{\mathcal{Z}, h}^{\mathcal{P}, \pi}(z) > 1$, since we consider the suffix $\text{suff}_{h-1}(z)$ in the above definitions. This choice is made to ensure that, for $z = (a_{h-L:h-1}, o_{h-L+1:h})$, the values $d_{\mathcal{Z}, h}^{\mathcal{P}, \pi}(z), d_{\mathcal{Z}, h}^{\mathcal{M}, \pi}(z)$ do not depend on elements of z with nonpositive indices when $h \leq L$.

Given the definitions above, we will view the corresponding visitations as vectors in their respective spaces: in particular, $d_{\mathcal{S}, h}^{\mathcal{P}, \pi} \in \mathbb{R}^{\bar{\mathcal{S}}}$, $d_{\mathcal{O}, h}^{\mathcal{P}, \pi}, d_{\mathcal{O}, h}^{\mathcal{M}, \pi} \in \mathbb{R}^{\bar{\mathcal{O}}}$, and $d_{\mathcal{Z}, h}^{\mathcal{P}, \pi}, d_{\mathcal{Z}, h}^{\mathcal{M}, \pi} \in \mathbb{R}^{\bar{\mathcal{Z}}}$. Note that the vectors $d_{\mathcal{O}, h}^{\mathcal{P}, \pi} \in \mathbb{R}^{\bar{\mathcal{O}}}$ and $d_{\mathcal{S}, h}^{\mathcal{P}, \pi} \in \mathbb{R}^{\bar{\mathcal{S}}}$ satisfy $d_{\mathcal{O}, h}^{\mathcal{P}, \pi} = \mathbb{O}_h^{\mathcal{P}} \cdot d_{\mathcal{S}, h}^{\mathcal{P}, \pi}$. Furthermore, we will slightly overload notation as follows: for a state visitation vector d in $\mathbb{R}^{\mathcal{T}}$ (for some set \mathcal{T}), and $\mathcal{T}' \subset \mathcal{T}$, we write $d(\mathcal{T}') := \sum_{t \in \mathcal{T}'} d(t)$.

B.1 Negative indices.

We remark that our expressions will, at times, refer to state visitation distributions of the form $d_{\mathcal{S}, h-L}^{\mathcal{P}, \pi}$ (or $d_{\mathcal{O}, h-L}^{\mathcal{P}, \pi}$, etc.), for some fixed $L \in \mathbb{N}$ and $h \in [H]$. Of course, if $h \leq L$, then this distribution is not defined. In all such cases, though, it will be the case that, for $h \leq L$, the expression in question does not depend on the distribution $d_{\mathcal{S}, h-L}^{\mathcal{P}, \pi}$, and thus the expression is well-defined. In a similar manner, for a trajectory $(a_{1:H-1}, o_{2:H})$ drawn from a POMDP \mathcal{P} , we will at times consider expressions of the form $(a_{h-L:h-1}, o_{h-L+1:h})$, for $h \in [H]$. In the case that $h \leq L$, we consider that the actions a_{h-L}, \dots, a_0 and observations o_{h-L+1}, \dots, o_1 to be chosen arbitrarily (to be concrete, one may consider that all $a_{h'} = a^*$ and $o_{h'+1} = o^*$ for $h' \leq 0$ and some universal $a^* \in \mathcal{A}, o^* \in \mathcal{O}$) – they will not affect the value of the relevant expression, and we use this convention purely to simplify notation.

B.2 Contextual decision processes

Both POMDPs and MDPs with \mathbf{Z} -structure are special cases of *contextual decision processes* [JKA⁺17]: a contextual decision process is a tuple $\mathcal{P} = (\mathcal{A}, \bar{\mathcal{O}}, H, R, \mathbb{P})$, where: \mathcal{A} is the ac-

tion space, $\bar{\mathcal{O}}$ is the *observation space* (also known as the *context space*), H is the horizon, $R = (R_1, \dots, R_H)$ is a collection of reward functions, $R_h : \bar{\mathcal{O}} \rightarrow \mathbb{R}$, and \mathbb{P} is the *transition kernel*. Formally, we have $\mathbb{P} = (\mathbb{P}_1, \dots, \mathbb{P}_H)$, where $\mathbb{P}_h : \mathcal{H}_h \times \mathcal{A} \rightarrow \Delta(\bar{\mathcal{O}})$ is a mapping from the space of histories \mathcal{H}_h (where $\mathcal{H}_h = \mathcal{A}^{h-1} \times \bar{\mathcal{O}}^{h-1}$) and actions to the space of distributions over the next observation (context). Given $h \geq 1$, a history $(a_{1:h-1}, o_{2:h})$ and an action $a_h \in \mathcal{A}$, we write $\mathbb{P}_h(o_{h+1} | a_{1:h}, o_{2:h})$ to denote the probability of observing o_{h+1} given the history $(a_{1:h-1}, o_{2:h})$ and that we take action a_h . We will consider the space of general policies Π^{gen} (defined in Section 2 as the convex hull of Π^{det} , which is the space of tuples $\pi = (\pi_1, \dots, \pi_H)$, where $\pi_h : \mathcal{H}_h \rightarrow \mathcal{A}$) acting on a contextual decision process.

POMDPs and MDPs with Z-structure may seem to be special case of contextual decision processes, as follows:

- Given a POMDP $\mathcal{P} = (\bar{\mathcal{S}}, \mathcal{A}, \bar{\mathcal{O}}, H, b_1, R, \mathbb{T}, \mathbb{O})$, we view it as the contextual decision process $\mathcal{P} = (\mathcal{A}, \bar{\mathcal{O}}, H, R, \mathbb{P})$, where \mathbb{P} is defined as follows: given a history $(a_{1:h-1}, o_{2:h})$ and an action $a_h \in \mathcal{A}$, define

$$\mathbb{P}_h^{\mathcal{P}}(o_{h+1} | a_{1:h}, o_{2:h}) := e_{o_{h+1}}^\top \mathbb{O}_{h+1}^{\mathcal{P}} \cdot \mathbb{T}_h^{\mathcal{P}}(a_h) \cdot \mathbf{b}_h^{\mathcal{P}}(a_{1:h-1}, o_{2:h}),$$

which is exactly the probability (in the POMDP \mathcal{P}) of next observing o_{h+1} given the history $(a_{1:h-1}, o_{2:h})$ and action a_h .

- Given an MDP with Z-structure, $\mathcal{M} = (\bar{\mathcal{Z}}, \mathcal{A}, H, R, \mathbb{P}^{\mathcal{M}})$, we view it as the contextual decision process $\mathcal{P} = (\mathcal{A}, \bar{\mathcal{O}}, H, R, \mathbb{P}^{\mathcal{P}})$, where we use the fact that (by Z-structure) $R = (R_1, \dots, R_H)$ has the property that each $R_h(z)$ is only a function of $o(z)$, and $\mathbb{P}^{\mathcal{P}}$ is defined as follows. We have $\mathbb{P}^{\mathcal{P}} = (\mathbb{P}_1^{\mathcal{P}}, \dots, \mathbb{P}_H^{\mathcal{P}})$, where

$$\mathbb{P}_h^{\mathcal{P}}(o_{h+1} | a_{1:h}, o_{2:h}) := \mathbb{P}_h^{\mathcal{M}}(z_{h+1} | z_h, a_h),$$

where $z_h = (a_{h-L:h-1}, o_{h-L+1:h})$ and $z_{h+1} = (a_{h-L+1:h}, o_{h-L+2:h+1})$, and we use the fact that \mathcal{M} has Z-structure to guarantee that $\mathbb{P}_h^{\mathcal{P}}(\cdot | a_{1:h}, o_{2:h})$ is a probability distribution.

Value functions. Given a deterministic policy $\sigma \in \Pi^{\text{det}}$ and a contextual decision process $\mathcal{P} = (\mathcal{A}, \bar{\mathcal{O}}, H, R, \mathbb{P})$, the value function $V_h^{\sigma, \mathcal{P}}(a_{1:h-1}, o_{2:h})$ is defined in an inductive manner:

$$V_h^{\sigma, \mathcal{P}}(a_{1:h-1}, o_{2:h}) := \mathbb{E}_{o_{h+1} \sim \mathbb{P}_h(\cdot | a_{1:h}, o_{2:h})} \left[R_{h+1}(o_{h+1}) + V_{h+1}^{\sigma, \mathcal{P}}(a_{1:h}, o_{2:h+1}) \right], \quad (4)$$

where we have written $a_h := \pi_h(a_{1:h-1}, o_{2:h})$ above. Then, for a general policy $\pi \in \Pi^{\text{gen}}$, we define $V_h^{\pi, \mathcal{P}}(a_{1:h-1}, o_{2:h}) := \mathbb{E}_{\sigma \sim \pi} \left[V_h^{\sigma, \mathcal{P}}(a_{1:h-1}, o_{2:h}) \right]$. We will often modify contextual decision processes by considering different choices of their reward function. Given a contextual decision process \mathcal{P} and a reward function $R = (R_1, \dots, R_H)$ (which may not be the reward function of \mathcal{P}), where $R_h : \bar{\mathcal{O}} \rightarrow \mathbb{R}$ for each $h \in [H]$, we write $V_h^{\pi, \mathcal{P}, R}(\cdot)$ to denote the value under π of \mathcal{P} where its reward function is replaced by R . We will often consider the value function of a contextual decision process \mathcal{P} which is actually either a POMDP or a MDP with Z-structure, in which case the following remarks apply:

- If \mathcal{P} is a POMDP, then $V_1^{\pi, \mathcal{P}}(\emptyset) = \mathbb{E}_{(a_{1:H-1}, o_{2:H}) \sim \pi} \left[\sum_{h=2}^H R_h(o_h) \right]$, i.e., our definition of the value function $V_1^{\pi, \mathcal{P}}(\emptyset)$ given here by reverse induction corresponds with that given in Section 2.
- If \mathcal{P} is a MDP with Z-structure, then, it is straightforward to see that $V_h^{\pi, \mathcal{P}}(a_{1:h-1}, o_{2:h})$ depends only on $z_h = (a_{h-L:h-1}, o_{h-L+1:h})$ (by the Markovian structure of \mathcal{P}), so we will often write $V_h^{\pi, \mathcal{P}}(z_h)$ in place of $V_h^{\pi, \mathcal{P}}(a_{1:h-1}, o_{2:h})$.

Viewing \mathbb{P}_h as an operator. Consider a contextual decision process $\mathcal{P} = (\mathcal{A}, \bar{\mathcal{O}}, H, R, \mathbb{P})$. For $h \in [H]$, we will view \mathbb{P}_h as the following operator: given a function $f : \mathcal{H}_{h+1} \rightarrow \mathbb{R}$, define $\mathbb{P}_h f : \mathcal{H}_h \times \mathcal{A} \rightarrow \mathbb{R}$ as:

$$(\mathbb{P}_h f)(a_{1:h}, o_{2:h}) := \mathbb{E}_{o_{h+1} \sim \mathbb{P}_h(\cdot | a_{1:h}, o_{2:h})} [f(a_{1:h}, o_{2:h+1})].$$

Thus we may alternatively write (4) as $V_h^{\sigma, \mathcal{P}}(a_{1:h-1}, o_{2:h}) = (\mathbb{P}_h(R_{h+1} + V_{h+1}^{\sigma, \mathcal{P}}))(a_{1:h}, o_{2:h})$, where we view R_{h+1} as a function on \mathcal{H}_{h+1} in the natural way, i.e., with abuse of notation, $R_{h+1}(a_{1:h}, o_{2:h+1}) := R_{h+1}(o_{h+1})$.

If \mathcal{P} is an MDP with Z-structure and $f(a_{1:h}, o_{2:h+1})$ only depends on the most recent L actions and observations (i.e., $(a_{h-L+1:h}, o_{h-L+2:h+1})$), then $(\mathbb{P}_h f)(a_{1:h}, o_{2:h})$ only depends on $z_h := (a_{h-L:h-1}, o_{h-L+1:h})$ and a_h . Thus, we will often write $(\mathbb{P}_h f)(z_h, a_h) := (\mathbb{P}_h f)(a_{1:h}, o_{2:h})$.

Lemma B.1 ([DLB17], Lemma E.15). *Consider any two CDPs $\mathcal{P}, \mathcal{P}'$, with reward functions $R^{\mathcal{P}}, R^{\mathcal{P}'}$ and transition kernels $\mathbb{P}^{\mathcal{P}}, \mathbb{P}^{\mathcal{P}'}$. Then for any deterministic policy π and all $h \in [H]$, it holds that, for all $a_{1:h-1}, o_{2:h}$,*

$$\begin{aligned} & V_h^{\pi, \mathcal{P}}(a_{1:h-1}, o_{2:h}) - V_h^{\pi, \mathcal{P}'}(a_{1:h-1}, o_{2:h}) \\ &= \mathbb{E}_{(a_{h:H-1}, o_{h+1:H}) \sim \pi} \left[\sum_{t=h}^H \left((\mathbb{P}_t^{\mathcal{P}} R_{t+1}^{\mathcal{P}})(a_{1:t}, o_{2:t}) - (\mathbb{P}_t^{\mathcal{P}'} R_{t+1}^{\mathcal{P}'})(a_{1:t}, o_{2:t}) \right. \right. \\ & \quad \left. \left. + ((\mathbb{P}_t^{\mathcal{P}} - \mathbb{P}_t^{\mathcal{P}'}) V_{t+1}^{\pi, \mathcal{P}'})(a_{1:t}, o_{2:t}) \right) \mid a_{1:h-1}, o_{2:h} \right]. \end{aligned}$$

Proof. The lemma is essentially an immediate consequence of [DLB17, Lemma E.15], since we may view $\mathcal{P}, \mathcal{P}'$ as MDPs on the state space of histories. For completeness, we write out the full proof in our setting.

We use reverse induction on h , noting that the lemma statement is immediate for $h = H + 1$ since $V_{H+1}^{\pi, \mathcal{P}}$ and $V_{H+1}^{\pi, \mathcal{P}'}$ are identically 0. Now assume that the statement of the lemma holds at time step $h + 1$. Fix any $(a_{1:h-1}, o_{2:h}) \in \mathcal{H}_h$, and let $a_h := \pi_h(a_{1:h-1}, o_{2:h})$. Then

$$\begin{aligned} V_h^{\pi, \mathcal{P}}(a_{1:h-1}, o_{2:h}) &= (\mathbb{P}_h^{\mathcal{P}} R_{h+1}^{\mathcal{P}})(a_{1:h}, o_{2:h}) + (\mathbb{P}_h^{\mathcal{P}} V_{h+1}^{\pi, \mathcal{P}})(a_{1:h}, o_{2:h}) \\ V_h^{\pi, \mathcal{P}'}(a_{1:h-1}, o_{2:h}) &= (\mathbb{P}_h^{\mathcal{P}'} R_{h+1}^{\mathcal{P}'})(a_{1:h}, o_{2:h}) + (\mathbb{P}_h^{\mathcal{P}'} V_{h+1}^{\pi, \mathcal{P}'})(a_{1:h}, o_{2:h}), \end{aligned}$$

and an analogous equality holds for \mathcal{P}' . Then we may compute, for any fixed $(a_{1:h-1}, o_{2:h}) \in \mathcal{H}_h$ (again letting $a_h := \pi_h(a_{1:h-1}, o_{2:h})$),

$$\begin{aligned} & V_h^{\pi, \mathcal{P}}(a_{1:h-1}, o_{2:h}) - V_h^{\pi, \mathcal{P}'}(a_{1:h-1}, o_{2:h}) \\ &= (\mathbb{P}_h^{\mathcal{P}} R_{h+1}^{\mathcal{P}})(a_{1:h}, o_{2:h}) - (\mathbb{P}_h^{\mathcal{P}'} R_{h+1}^{\mathcal{P}'})(a_{1:h}, o_{2:h}) \\ & \quad + ((\mathbb{P}_h^{\mathcal{P}} - \mathbb{P}_h^{\mathcal{P}'}) V_{h+1}^{\pi, \mathcal{P}'})(a_{1:h}, o_{2:h}) + (\mathbb{P}_h^{\mathcal{P}} (V_{h+1}^{\pi, \mathcal{P}} - V_{h+1}^{\pi, \mathcal{P}'}))(a_{1:h}, o_{2:h}) \\ &= (\mathbb{P}_h^{\mathcal{P}} R_{h+1}^{\mathcal{P}})(a_{1:h}, o_{2:h}) - (\mathbb{P}_h^{\mathcal{P}'} R_{h+1}^{\mathcal{P}'})(a_{1:h}, o_{2:h}) + ((\mathbb{P}_h^{\mathcal{P}} - \mathbb{P}_h^{\mathcal{P}'}) V_{h+1}^{\pi, \mathcal{P}'})(a_{1:h}, o_{2:h}) \\ & \quad + \mathbb{E}_{(a_h, o_{h+1}) \sim \pi} \left[(V_{h+1}^{\pi, \mathcal{P}} - V_{h+1}^{\pi, \mathcal{P}'})(a_{1:h}, o_{2:h+1}) \mid a_{1:h-1}, o_{2:h} \right] \\ &= (\mathbb{P}_h^{\mathcal{P}} R_{h+1}^{\mathcal{P}})(a_{1:h}, o_{2:h}) - (\mathbb{P}_h^{\mathcal{P}'} R_{h+1}^{\mathcal{P}'})(a_{1:h}, o_{2:h}) + ((\mathbb{P}_h^{\mathcal{P}} - \mathbb{P}_h^{\mathcal{P}'}) V_{h+1}^{\pi, \mathcal{P}'})(a_{1:h}, o_{2:h}) \\ & \quad + \mathbb{E}_{(a_h, o_{h+1}) \sim \pi} \left[\mathbb{E}_{(a_{h+1:H-1}, o_{h+2:H}) \sim \pi} \left[\sum_{t=h+1}^H (\mathbb{P}_t^{\mathcal{P}} R_{t+1}^{\mathcal{P}})(a_{1:t}, o_{2:t}) - (\mathbb{P}_t^{\mathcal{P}'} R_{t+1}^{\mathcal{P}'})(a_{1:t}, o_{2:t}) \right. \right. \\ & \quad \left. \left. + \sum_{t=h+1}^H ((\mathbb{P}_t^{\mathcal{P}} - \mathbb{P}_t^{\mathcal{P}'}) V_{t+1}^{\pi, \mathcal{P}'})(a_{1:t}, o_{2:t}) \mid a_{1:h}, o_{2:h+1} \right] \mid a_{1:h-1}, o_{2:h} \right] \\ &= \mathbb{E}_{(a_{h:H-1}, o_{h+1:H}) \sim \pi} \left[\sum_{t=h}^H (\mathbb{P}_t^{\mathcal{P}} R_{t+1}^{\mathcal{P}})(a_{1:t}, o_{2:t}) - (\mathbb{P}_t^{\mathcal{P}'} R_{t+1}^{\mathcal{P}'})(a_{1:t}, o_{2:t}) \right. \\ & \quad \left. + \sum_{t=h}^H ((\mathbb{P}_t^{\mathcal{P}} - \mathbb{P}_t^{\mathcal{P}'}) V_{t+1}^{\pi, \mathcal{P}'})(a_{1:t}, o_{2:t}) \mid a_{1:h-1}, o_{2:h} \right], \end{aligned}$$

which completes the inductive step. \square

B.3 Belief contraction, expanded

Below we state a slight strengthening of Theorem 2.1 which will be needed in our proofs.

Theorem B.2 (Belief contraction; Theorems 4.1 and 4.7 of [GMR22]). *Consider any γ -observable POMDP \mathcal{P} , any $\epsilon > 0$ and $L \in \mathbb{N}$ so that $L \geq C \cdot \min \left\{ \frac{\log(1/(\epsilon\phi)) \log(\log(1/\phi)/\epsilon)}{\gamma^2}, \frac{\log(1/(\epsilon\phi))}{\gamma^4} \right\}$. Fix any $\pi \in \Pi^{\text{gen}}$, and suppose that $\mathcal{D}', \mathcal{D} \in \Delta^{\mathcal{S}}$ satisfy $\frac{\mathcal{D}'(s)}{\mathcal{D}(s)} \leq \frac{1}{\phi}$ (where we take $0/0 = 0$). Then*

$$\mathbb{E}_{\substack{\mathcal{P} \\ s_{h-L} \sim \mathcal{D}' \\ (a_{h-L:h-1}, o_{h-L+1:h}) \sim \pi}} \left\| \mathbf{b}_h^{\text{apx}, \mathcal{P}}(a_{h-L:h-1}, o_{h-L+1:h}; \mathcal{D}') - \mathbf{b}_h^{\text{apx}, \mathcal{P}}(a_{h-L:h-1}, o_{h-L+1:h}; \mathcal{D}) \right\|_1 \leq \epsilon. \quad (5)$$

Note that the theorems in [GMR22] only state the special case $\mathcal{D} = \text{Unif}(\mathcal{S})$, in which case the precondition of the theorem is always satisfied with $\phi = 1/S$. However, the proof extends unchanged to this generality.

B.4 Pseudoinverses

As previously stated, for a matrix $\mathbb{O} \in \mathbb{R}^{O \times S}$, we let \mathbb{O}^\dagger denote the Moore-Penrose pseudoinverse of \mathbb{O} . The following lemma is straightforward:

Lemma B.3. *If $\mathbb{O} \in \mathbb{R}^{O \times S}$ is γ -observable, then for all $x \in \mathbb{R}^O$, $\|\mathbb{O}^\dagger x\|_1 \leq \frac{\sqrt{S}}{\gamma} \cdot \|x\|_1$. In particular, all entries of \mathbb{O}^\dagger are bounded in magnitude by $\frac{\sqrt{S}}{\gamma}$.*

Proof. Write $y = \mathbb{O}^\dagger x$. Then

$$\|y\|_1 \leq \frac{1}{\gamma} \cdot \|\mathbb{O}y\|_1 \leq \frac{\sqrt{S}}{\gamma} \cdot \|\mathbb{O}y\|_2 = \frac{\sqrt{S}}{\gamma} \cdot \|\mathbb{O}\mathbb{O}^\dagger x\|_2 \leq \frac{\sqrt{S}}{\gamma} \cdot \|x\|_2 \leq \frac{\sqrt{S}}{\gamma} \cdot \|x\|_1.$$

where the first inequality follows from observability and the second follows from Cauchy-Schwartz. To establish the second statement, note that, for each $o \in [O]$, we have that $\|\mathbb{O}^\dagger \cdot e_o\|_1 \leq \frac{\sqrt{S}}{\gamma} \cdot \|e_o\|_1 = \frac{\sqrt{S}}{\gamma}$. \square

C Proof Overview and Organization

Below we describe the organization of the proof of Theorem 3.1 which obtains a guarantee for BaSeCAMP (Algorithm 3). At a high level, our task is to show that the true POMDP \mathcal{P} and the MDPs with Z-structure $\widehat{\mathcal{M}}^{(k)}$ constructed in the course of BaSeCAMP are close in some sense; doing so is crucial to proving the main ‘‘progress lemma’’ (Lemma I.2, the formal version of Lemma 5.1) as discussed in Section 5:

- In Section D, we introduce the concept of barycentric spanners (Definition D.1) and prove our main lemma (Lemma D.2, the formal version of Lemma 4.1) which establishes a guarantee on the output of Algorithm 2 (BarySpannerPolicy).
- In Section E we introduce several intermediary contextual decision processes which we will use to relate the true POMDP \mathcal{P} and the empirical MDPs $\widehat{\mathcal{M}}^{(k)}$ constructed in BaSeCAMP: these include an MDP $\widehat{\mathcal{M}}(\pi^{1:H})$ with Z-structure (which depends on a sequence of general policies $\pi^{1:H}$; see Section E.1) and a collection of truncated POMDPs, $\overline{\mathcal{P}}_{\phi, h}(\pi^{1:H})$, $h \in [H]$ (which depend on a parameter $\phi > 0$ and a sequence of general policies $\pi^{1:H}$; see Section E.3). Omitting some details, the way we relate the various contextual decision processes is summarized as follows:

$$\mathcal{P} \stackrel{\text{Lems. F.1 \& F.2}}{\sim} \overline{\mathcal{P}}_{\phi, h}(\pi^{1:H}) \stackrel{\text{Lems. F.3 \& G.4}}{\sim} \widehat{\mathcal{M}}(\pi^{1:H}) \stackrel{\text{Lems. E.1, F.4 \& G.4}}{\sim} \widehat{\mathcal{M}}^{(k)}. \quad (6)$$

- The truncated POMDPs $\overline{\mathcal{P}}_{\phi, h}(\pi^{1:H})$ constructed in Section E.3 are a special case of a more general notion of *truncation* of the given POMDP \mathcal{P} (Definition F.1). In Section F we prove several properties of truncations of POMDPs, which are useful for establishing several steps in the relations (6).

- In Section G we use the results of Section F to establish a certain type of “one-sided” relation between $\overline{\mathcal{P}}_{\phi,h}(\pi^{1:H})$ and $\widehat{\mathcal{M}}^{(k)}$ (for each $k \in [K]$). In particular, the main result of Section G is Lemma G.4, which establishes the following: for any reward function R which is non-negative in a certain sense, the value of any general policy π in $\overline{\mathcal{P}}_{\phi,h}(\pi^{1:H})$ with the reward function R is approximately bounded above by the value of π in $\widehat{\mathcal{M}}^{(k)}$ with the reward function R ; i.e., we have an upper bound on $V_1^{\pi, \overline{\mathcal{P}}_{\phi,h}(\pi^{1:H}), R}(\emptyset) - V_1^{\pi, \widehat{\mathcal{M}}^{(k)}, R}(\emptyset)$. The inequality does not hold, though, in the opposite direction: this is because $\overline{\mathcal{P}}_{\phi,h}(\pi^{1:H})$ is obtained by redirecting some of the transitions of \mathcal{P} into the sink state s^{sink} (which will always lead to 0 reward), and lacking a bound on the probability of ending up in s^{sink} , it is impossible to show an inequality in the opposite direction.
- Our proof also requires a two-sided inequality that upper bounds $\left| V_1^{\pi, \mathcal{P}, R}(\emptyset) - V_1^{\pi, \widehat{\mathcal{M}}^{(k)}, R}(\emptyset) \right|$ for general policies π . Such an inequality is established in Section H, namely in Lemma H.3. The proof of Lemma H.3 proceeds somewhat differently than that of the one-sided inequalities, in that it does not make use of the truncated POMDPs $\overline{\mathcal{P}}_{\phi,h}(\pi^{1:H})$; in particular, as opposed to (6), the rough outline is as follows:

$$\mathcal{P} \stackrel{\text{Lems. H.2 \& H.3}}{\sim} \widetilde{\mathcal{M}}(\pi^{1:H}) \stackrel{\text{Lems. E.1, F.4, \& H.3}}{\sim} \widehat{\mathcal{M}}^{(k)}. \quad (7)$$

Unsurprisingly, there is a cost to the fact that this inequality is two-sided (which is of course stronger than being one-sided): the upper bound on $\left| V_1^{\pi, \mathcal{P}, R}(\emptyset) - V_1^{\pi, \widehat{\mathcal{M}}^{(k)}, R}(\emptyset) \right|$ depends on the probability, say $\rho(\pi)$, that π visits the set of states which are “underexplored” (see Definition E.2) with respect to a certain set of exploratory policies used by BaSeCAMP. This additional term of $\rho(\pi)$ is dealt as follows when proving the “progress lemma”, Lemma I.2: we will be applying Lemma H.3 with $\pi = \pi^{k+1, h, 0}$ for some $h \in [H]$, namely one of the “exploratory policies” constructed in BaSeCAMP with the purpose of exploring states at step $h - L$. If the probability $\rho(\pi^{k+1, h, 0})$ is too large, then, by the precise definition of $\rho(\pi^{k+1, h, 0})$, it turns out that $\pi^{k+1, h, 0}$ must explore some state *prior to* step $h - L$ (even though $\pi^{k+1, h, 0}$ was constructed with the “purpose” of exploring states at step $h - L$!). We can then use this fact to conclude that item 2 of Lemma I.2 holds in this case.

- Finally, Section I analyzes BaSeCAMP, proving Theorem 3.1. As outlined in Section 5, the main step is to prove Lemma I.2, which establishes that BaSeCAMP makes a quantifiable amount of progress at each iteration k . Using Lemma I.2, Lemma I.3 establishes that, in fact, at some iteration \bar{k} in BaSeCAMP, the item 1 of Lemma I.2 must hold, which, roughly speaking, means the POMDP \mathcal{P} has been sufficiently explored. Lemma I.4 uses this fact together with Lemma H.3 to establish a *two-sided* upper bound between \mathcal{P} and $\widehat{\mathcal{M}}^{(\bar{k})}$ (without the additional term involving underexplored states referred to above). From this result, the proof of Theorem 3.1 is straightforward.

C.1 Definitions of hyperparameters

Below we give the values for the various hyperparameters employed by our algorithm and used in the proof, given a γ -observable POMDP with S states, A actions, O observations, horizon H , and a desired accuracy level α and probability of failure β (in the context of Theorem 3.1). C^* refers to a sufficiently large constant. Throughout the remainder of the appendix, when we will refer to the values of the constants below with any qualifiers, they will be assumed to take the below values.

- $\epsilon = \alpha \cdot \frac{\gamma}{O^2 H^5 S^{3/2} (C^*)^2}$.
- $\phi = \frac{1}{C^*} \cdot \frac{\gamma}{H^5 S^{7/2} O^2} \cdot \epsilon$.
- $L = C^* \cdot \min \left\{ \frac{\log(1/(\epsilon\phi)) \log(\log(1/\phi)/\epsilon)}{\gamma^2}, \frac{\log(1/(\epsilon\phi))}{\gamma^4} \right\}$.
- $\theta = \epsilon$ (see Lemma E.1).
- $\zeta = \frac{\epsilon\phi}{A^{2L} O^L}$ (see Lemma E.1).
- $\delta = C^* \cdot \frac{O^2 H^3 \sqrt{S}}{\gamma} \cdot \epsilon$.

- $\delta' = \delta/2$.
- $K = 2HS$.
- $p = \beta/(2K)$.
- $N_1 = C^\star \cdot \frac{LA^{L+1}O^L \log(AO/p)}{\theta^2}$ (used in Algorithm 1).
- $N_0 = C^\star \cdot \frac{N_1 AL \log(OA/p)}{\zeta}$ (used in Algorithm 1).

We assume that H is chosen sufficiently large so that $H > L$ (if not, then we can increase H to be $L + 1$, which only affects the time and sample complexities by a constant factor in the exponent).

D Barycentric spanners

Definition D.1 below formally defines the notion of *barycentric spanner*.

Definition D.1. Consider a subset $\mathcal{X} \subset \mathbb{R}^d$. For $C \geq 1$, a set $\mathcal{X}' \subset \mathcal{X}$ of size d is a C -approximate barycentric spanner of \mathcal{X} if each $x \in \mathcal{X}$ can be expressed as a linear combination of elements in \mathcal{X}' with coefficients in $[-C, C]$.

It is easy to show that (exact) barycentric spanners exist: it is without loss of generality to assume that \mathcal{X} is full-rank, since otherwise we may apply a linear transformation to it so that the transformed set is full-rank (in a smaller-dimensional space). Then, given that \mathcal{X} is full-rank, it may be verified that a collection of d vectors $x_1, \dots, x_d \in \mathcal{X}$ maximizing $|\det(x_1, \dots, x_d)|$ is a (1-approximate) barycentric spanner. We will use the following result of Awerbuch & Kleinberg, showing that approximate barycentric spanners can be computed efficiently:

Lemma D.1 ([AK08], Proposition 2.5). *Suppose $\mathcal{X} \subset \mathbb{R}^d$ and $\lambda > 1$. Given an oracle for optimizing linear functions over \mathcal{X} , a λ -approximate barycentric spanner for \mathcal{X} may be computed in polynomial time using $O(d^2 \log_\lambda(d))$ calls to the optimization oracle.*

Further, each element x_i of the resulting approximate barycentric spanner is equal to the output of the optimization oracle for some input cost vector.

We remark that the algorithm of [AK08] technically applies to the setting that \mathcal{X} is full-dimensional. In general this will not be the case for sets \mathcal{X} to which we will apply Lemma D.1. We may deal with sets $\mathcal{X} \subset \mathbb{R}^d$ which are not full-dimensional as follows: using at most d oracle calls to the optimization oracle, we may compute a set of $k \leq d$ linearly independent vectors which spans \mathcal{X} . Using these k vectors, we may compute a one-to-one linear mapping T from \mathcal{X} into \mathbb{R}^k , so that the image of \mathcal{X} in \mathbb{R}^k , say \mathcal{Y} , is full-dimensional. We may then use Lemma D.1 to compute a λ -approximate barycentric spanner of \mathcal{Y} , and then map this back into \mathcal{X} via T^{-1} , yielding a λ -approximate barycentric spanner of \mathcal{X} .

In Lemma D.2 below, we consider positive integers $S, O \in \mathbb{N}$, a matrix $\mathbb{M} \in \mathbb{R}^{S \times O}$, as well as a set $\mathcal{X} \subset \mathbb{R}^O$ indexed by some given set Π . In particular, there is a bijection between Π and \mathcal{X} , which we denote as $\pi \mapsto x(\pi)$. In our usage of Lemma D.2, the set \mathcal{X} will be the set of all possible observation visitation distributions (indexed by policies π) at a particular layer of the current empirical MDP with Z-structure $\widehat{\mathcal{M}}^{(k)}$. The matrix \mathbb{M} will be equal to \mathbb{O}_h^\dagger for some layer h of the POMDP. Thus, for $\pi \in \Pi$, $\mathbb{M} \cdot x(\pi) = \mathbb{O}_h^\dagger \cdot x(\pi)$ represents an estimate of the state visitation distribution of \mathcal{P} under policy π at step h . In this context, the conclusion of Lemma D.2 should be interpreted as saying that the mixture policy defined as $\frac{\pi^1 + \dots + \pi^O}{O}$ visits all states which can be visited with non-negligible probability.

Lemma D.2. *Fix $\eta > 0$, $S, O \in \mathbb{N}$, a matrix $\mathbb{M} \in \mathbb{R}^{S \times O}$, and a set $\mathcal{X} \subset \mathbb{R}^O$, indexed by some class Π . Assume that for each $x \in \mathcal{X}$, we have*

$$\sum_{s \in S} [(\mathbb{M} \cdot x)(s)]_- \leq \frac{\eta}{4O^2}. \quad (8)$$

Consider an oracle \mathcal{O} which returns, given any $r \in \mathbb{R}^O$, some $\pi^(r) \in \arg \max_{\pi \in \Pi} \langle r, x(\pi) \rangle$ as well as $\langle r, x(\pi^*(r)) \rangle$. Then there is a polynomial-time algorithm which only has access to the oracle \mathcal{O}*

(i.e., does not know \mathbb{M}), makes $O(O^2 \log O)$ oracle calls, and returns $\pi^1, \dots, \pi^O \in \Pi$ satisfying

$$\forall \mathcal{S}' \subset [S] \text{ s.t. } \max_{\pi \in \Pi} \sum_{s \in \mathcal{S}'} \langle e_s, \mathbb{M} \cdot x(\pi) \rangle \geq \eta, \quad \sum_{s \in \mathcal{S}'} \left\langle e_s, \frac{\mathbb{M} \cdot (x(\pi^1) + \dots + x(\pi^O))}{O} \right\rangle \geq \frac{\eta}{4O^2}. \quad (9)$$

Proof. By Lemma D.1, with $O(O^2 \log O)$ calls to the oracle \mathcal{O} , we may compute a set $\{\pi^1, \dots, \pi^O\}$ so that $\{x(\pi^1), \dots, x(\pi^O)\}$ is a 2-approximate barycentric spanner of \mathcal{X} .

Consider any $\mathcal{S}' \subset [S]$ so that, for some $\pi \in \Pi$, $\sum_{s \in \mathcal{S}'} \langle e_s, \mathbb{M} \cdot x(\pi) \rangle \geq \eta$. By the definition of barycentric spanner, we can write $x(\pi) = \sum_{o=1}^O \alpha_o \cdot x(\pi^o)$, for some reals $\alpha_o \in [-2, 2]$, $o \in [O]$. Then we have

$$\eta \leq \sum_{s \in \mathcal{S}'} \langle e_s, \mathbb{M} \cdot x(\pi) \rangle = \sum_{o=1}^O \alpha_o \cdot \left\langle \sum_{s \in \mathcal{S}'} e_s, \mathbb{M} \cdot x(\pi^o) \right\rangle.$$

Thus there must be some $o \in [O]$ so that $|\langle \sum_{s \in \mathcal{S}'} e_s, \mathbb{M} \cdot x(\pi^o) \rangle| \geq \frac{\eta}{2O}$, which implies, by the assumption (8), that $\langle \sum_{s \in \mathcal{S}'} e_s, \mathbb{M} \cdot x(\pi^o) \rangle \geq \frac{\eta}{2O}$. Since, for all $o' \in [O]$, we similarly have, by (8),

$$\left\langle \sum_{s \in \mathcal{S}'} e_s, \mathbb{M} \cdot x(\pi^{o'}) \right\rangle \geq -\sum_{s \in \mathcal{S}} \left[(\mathbb{M} \cdot x(\pi^{o'})) (s) \right]_- \geq -\frac{\eta}{4O^2},$$

it follows that

$$\left\langle \sum_{s \in \mathcal{S}'} e_s, \mathbb{M} \cdot \frac{x(\pi^1) + \dots + x(\pi^O)}{O} \right\rangle \geq \frac{\eta}{2O^2} - \frac{\eta}{4O^2} = \frac{\eta}{4O^2}.$$

□

E Intermediary Models for the Analysis

In Section 4.1, we gave an overview of the notion of an MDP \mathcal{M} which approximates the POMDP \mathcal{P} , using approximate belief states initialized with the uniform distribution. We then suggested that, if appropriate policies $\pi^{1:H}$ can be found, then (using `ConstructMDP`, Algorithm 1), we can construct an MDP $\widehat{\mathcal{M}}(\pi^{1:H})$ which can approximate \mathcal{M} , and therefore \mathcal{P} . To make this argument more formal, we use an MDP which differs slightly from the “intermediate” MDP \mathcal{M} introduced in Section 4.1 to relate $\widehat{\mathcal{M}}(\pi^{1:H})$ and \mathcal{P} . In particular, in Section E.1 we introduce an MDP $\widetilde{\mathcal{M}}(\pi^{1:H})$ for any set of policies $\pi^{1:H}$. At a high level, $\widetilde{\mathcal{M}}(\pi^{1:H})$ uses approximate belief states initialized with the visitation distributions induced by $\pi^{1:H}$. In Section E.2 we show by concentration arguments that $\widetilde{\mathcal{M}}(\pi^{1:H})$ is close to the “empirical MDP” $\widehat{\mathcal{M}}(\pi^{1:H})$ (constructed in Algorithm 1) for arbitrary $\pi^{1:H}$.

Subsequently, the core technical challenge will be to relate $\widetilde{\mathcal{M}}(\pi^{1:H})$ to \mathcal{P} under various assumptions on $\pi^{1:H}$. As described in the proof overview, we will crucially use *truncated POMDPs* as an intermediary model. These are introduced in Section E.3.

E.1 Construction of Approximate MDP $\widetilde{\mathcal{M}}(\pi^{1:H})$

Fix a POMDP \mathcal{P} , and consider any collection of H general policies π^1, \dots, π^H , which we abbreviate as $\pi^{1:H}$. For each $h \in [H]$, let $\widehat{\pi}^h$ denote the policy which follows π^h for the first $\max\{h-L-1, 0\}$ steps and then chooses a uniformly random action for steps $h-L$ through h . Write $\widehat{\pi} := \frac{1}{H} \cdot \sum_{h=1}^H \widehat{\pi}^h$.

Given the policies $\pi^{1:H}$ as above, we now define an MDP with Z-structure, denoted $\widetilde{\mathcal{M}}(\pi^{1:H}) = (\overline{\mathcal{Z}}, H, \mathcal{A}, \mathbb{P}^{\widetilde{\mathcal{M}}(\pi^{1:H})})$, as follows. To define the transitions $\mathbb{P}_h^{\widetilde{\mathcal{M}}(\pi^{1:H})}$, first define, for $\pi' \in \Pi^{\text{gen}}$,

$$\widetilde{\mathbf{b}}_h^{\pi'}(a_{h-L:h-1}, o_{h-L+1:h}) := \mathbf{b}_h^{\text{apx}, \mathcal{P}}(a_{h-L:h-1}, o_{h-L+1:h}; d_{\mathcal{S}, h-L}^{\pi'}).$$

(Recall the definition of $\mathbf{b}_h^{\text{apx}, \mathcal{P}}(\cdot)$ in Section 2.3; furthermore, recall that, if $h \leq L$, $\mathbf{b}_h^{\text{apx}, \mathcal{P}}(a_{h-L:h-1}, o_{h-L+1:h}; \mathcal{D})$ is defined as $\mathbf{b}_h^{\text{apx}, \mathcal{P}}(a_{1:h-1}, o_{2:h}; \mathcal{D})$, which does not depend on \mathcal{D} , as it is $\mathbf{b}_h^{\mathcal{P}}(a_{1:h-1}, o_{2:h})$. Thus, even for $h \leq L$, $\tilde{\mathbf{b}}_h^{\pi^h}$ is well-defined despite the fact that $d_{\mathcal{S}, h-L}^{\pi^h, \mathcal{P}}$ is not.)

Given a sequence $a_{1:h}, o_{2:h+1}$, writing $z_h = (a_{h-L:h-1}, o_{h-L+1:h})$, $z_{h+1} = (a_{h-L+1:h}, o_{h-L+2:h+1})$, define, in the case that $z_h \in \mathcal{Z}$,

$$\mathbb{P}_h^{\tilde{\mathcal{M}}(\pi^{1:H})}(z_{h+1}|z_h, a_h) = e_{o_{h+1}}^\top \cdot \mathbb{O}_{h+1} \cdot \mathbb{T}_h(a_h) \cdot \tilde{\mathbf{b}}_h^{\pi^h}(a_{h-L:h-1}, o_{h-L+1:h}). \quad (10)$$

Note that $\tilde{\mathcal{M}}(\pi^{1:H})$ has \mathbf{Z} -structure: given $z_h = (a_{h-L:h-1}, o_{h-L+1:h}), a_h$, under the transitions of $\tilde{\mathcal{M}}(\pi^{1:H})$, all coordinates of z_{h+1} except the last observation are full determined (namely, they are given by $(a_{h-L+1:h}, o_{h-L+2:h})$); thus, with slight abuse of notation, we will at times write $\mathbb{P}_h^{\tilde{\mathcal{M}}(\pi^{1:H})}(o_{h+1}|z_h, a_h) = \mathbb{P}_h^{\tilde{\mathcal{M}}(\pi^{1:H})}(z_{h+1}|z_h, a_h)$. If $z_h \notin \mathcal{Z}$, then we define $\mathbb{P}_h^{\tilde{\mathcal{M}}(\pi^{1:H})}(o^{\text{sink}}|z_h, a_h) = 1$. Finally, we remark that we have not specified a reward function or initial state distribution for $\tilde{\mathcal{M}}(\pi^{1:H})$; such choices turn out not to affect the analysis.

E.2 Concentration of Empirical Approximate MDP $\widehat{\mathcal{M}}(\pi^{1:H})$.

While we do not have algorithmic access to $\tilde{\mathcal{M}}(\pi^{1:H})$, we can algorithmically construct an empirical approximation of it. This is precisely the (random) MDP $\widehat{\mathcal{M}}(\pi^{1:H}) = (\mathcal{Z}, H, \mathcal{A}, R^{\widehat{\mathcal{M}}(\pi^{1:H})}, \mathbb{P}^{\widehat{\mathcal{M}}(\pi^{1:H})})$, defined as the output of Algorithm 1 when given input policies $\pi^{1:H}$. Notice that the MDP $\widehat{\mathcal{M}}(\pi^{1:H})$ does come with a reward function $R^{\widehat{\mathcal{M}}(\pi^{1:H})}$; we remark that, at some points in the proof, we will consider an alternative reward function (generally denoted by R , without the superscript) attached to the remaining data of $\widehat{\mathcal{M}}(\pi^{1:H})$, purely for purposes of analysis.

Lemma E.1 shows by Chernoff bounds that outside the set of states which are rarely visited by $\hat{\pi}^{1:H}$ (defined in Definition E.1 below), we have that $\widehat{\mathcal{M}}(\pi^{1:H})$ closely approximates $\tilde{\mathcal{M}}(\pi^{1:H})$.

To state Lemma E.1, we need to introduce some notation.

Definition E.1. For each $h \in [H]$, $\zeta > 0$, and a general policy π' , define the set $\mathcal{Z}_{h,\zeta}^{\text{low}}(\pi') \subset \mathcal{Z}$, as follows:

$$\mathcal{Z}_{h,\zeta}^{\text{low}}(\pi') := \left\{ z \in \mathcal{Z} : d_{\mathcal{Z}, h}^{\mathcal{P}, \pi'}(z) \leq \zeta \right\}.$$

Lemma E.1. Fix any sequence of policies $\pi^{1:H}$ and the resulting $\hat{\pi}^{1:H}$ as defined in Section E.1. There is an event \mathcal{E}^{low} that occurs with probability $1 - p$, so that under the event \mathcal{E}^{low} , we have:

1. The output $\widehat{\mathcal{M}}(\pi^{1:H})$ of `ConstructMDP` with parameters N_1, N_0 satisfying

$$N_1 \geq C_0 \cdot \left(\frac{LA^{L+1}O^L \log(AO/p)}{\theta^2} \right), \quad N_0 \geq C_0 \cdot \frac{N_1 AL \log(OA/p)}{\zeta} \quad (11)$$

(where C_0 is a sufficiently large constant) satisfies the following: for all $z_h \notin \mathcal{Z}_{h,\zeta}^{\text{low}}(\hat{\pi}^h)$ and all $a_h \in \mathcal{A}$, we have that $R_h^{\widehat{\mathcal{M}}(\pi^{1:H})}(o(z_h)) = R_h^{\mathcal{P}}(o(z_h))$ and

$$\left\| \mathbb{P}_h^{\widehat{\mathcal{M}}(\pi^{1:H})}(\cdot|z_h, a_h) - \mathbb{P}_h^{\tilde{\mathcal{M}}(\pi^{1:H})}(\cdot|z_h, a_h) \right\|_1 \leq \theta. \quad (12)$$

2. For all $z_h \in \mathcal{Z}_{h,\zeta}^{\text{low}}(\hat{\pi}^h)$, if (12) does not hold, then under the transitions of $\widehat{\mathcal{M}}(\pi^{1:H})$, for all $a_h \in \mathcal{A}$, $z_h = (a_{h-L:h-1}, o_{h-L+1:h})$ transitions to $(a_{h-L+1:h}, (o_{h-L+2:h}, o^{\text{sink}}))$ with probability 1.

Notice that, for $h \leq L$, $z_h = (a_{h-L:h-1}, o_{h-L+1:h})$ includes actions (and observations) with nonpositive indices: importantly, such actions and observations do not affect the distribution of the next observation, namely $\mathbb{P}_h^{\widehat{\mathcal{M}}(\pi^{1:H})}(o_{h+1}|z_h, a_h)$ or $\mathbb{P}_h^{\tilde{\mathcal{M}}(\pi^{1:H})}(o_{h+1}|z_h, a_h)$, which intuitively explains why (12) can hold even for such h .

Proof of Lemma E.1. First, note that (12) is immediate if $z_h \notin \mathcal{Z}$, since then

$$\mathbb{P}_h^{\widehat{\mathcal{M}}(\pi^{1:H})}(\cdot | z_h, a_h) = \mathbb{P}_h^{\widehat{\mathcal{M}}(\pi^{1:H})}(\cdot | z_h, a_h) = o^{\text{sink}}$$

with probability 1. Thus for the remainder of the proof we may consider $z_h \in \mathcal{Z}$.

Consider some $p' \in (0, 1)$, to be specified below. Fix any $h \in [H]$: we will argue about the h th step in the for loop on step 2 of Algorithm 1. We will use the following key fact: for any $(z_h, a_h) \in \mathcal{Z} \times \mathcal{A}$, and for any iteration i of ConstructMDP, the distribution of z_{h+1}^i conditioned on $a_h^i = a_h$ and $\text{suff}_{h-1}(z_h^i) = \text{suff}_{h-1}(z_h)$ is the same as $\mathbb{P}_h^{\widehat{\mathcal{M}}(\pi^{1:H})}(\cdot | z_h, a_h)$. To see that this fact holds, it suffices to note that, since (z_h^i, a_h^i) is drawn according to the policy $\widehat{\pi}^h$, the distribution of s_h , conditioned on the observation-action sequence $(a_{\max\{1, h-L\}:h-1}, o_{\max\{2, h-L+1\}:h})$, is exactly $\widetilde{\mathbf{b}}_h^{\pi^h}(a_{h-L:h-1}, o_{h-L+1:h})$. Thus, given that action a_h is taken at step h , the probability that the next observation is o_{h+1} is given by $e_{o_{h+1}}^\top \cdot \mathbb{O}_{h+1} \cdot \mathbb{T}_h(a_h) \cdot \widetilde{\mathbf{b}}_h^{\pi^h}(a_{h-L:h-1}, o_{h-L+1:h})$, which, by definition, is exactly $\mathbb{P}_h^{\widehat{\mathcal{M}}(\pi^{1:H})}((a_{h-L+1:h}, o_{h-L+2:h+1}) | z_h, a_h)$.

Consider any $z_h \in \mathcal{Z}$ with $z_h \notin \mathcal{Z}_{h, \zeta}^{\text{low}}(\widehat{\pi}^h)$ and any $a_h \in \mathcal{A}$. By the Chernoff bound, with probability at least $1 - \exp(-\zeta N_0 / (8A))$, there are at least $\frac{\zeta N_0}{2A} \geq N_1$ iterations $i \in [N_0]$ (in the for loop in step 4 of Algorithm 1) so that $a_h^i = a_h$ and $\text{suff}_{h-1}(z_h^i) = \text{suff}_{h-1}(z_h)$. Under such an event, since the reward function $R^{\mathcal{P}}$ of \mathcal{P} is deterministic, it must be the case that $R_h^{\widehat{\mathcal{M}}(\pi^{1:H})}(o(z_h)) = R_h^{\mathcal{P}}(o(z_h))$ (step 13 of Algorithm 1).

Next, consider any $(z_h, a_h) \in \mathcal{Z} \times \mathcal{A}$ so that there are at least N_1 values of i so that $a_h^i = a_h$ and $\text{suff}_{h-1}(z_h^i) = \text{suff}_{h-1}(z_h)$ (i.e., for the chosen value of (z_h, a_h) , the if statement at step 11 holds). By the fact established in the previous paragraph and [Can20, Theorem 1], with probability at least $1 - p'$, (12) holds as long as

$$N_1 \geq C \cdot \frac{|\mathcal{Z}| + \log(1/p')}{\theta^2} \quad (13)$$

for some constant $C > 1$. Thus, by a union bound over all $h \in [H]$, and all $(z_h, a_h) \in \mathcal{Z} \times \mathcal{A}$, under some event \mathcal{E}^{low} that occurs with probability at least

$$1 - |\mathcal{Z}|A \cdot \exp(-\zeta N_0 / (8A)) - |\mathcal{Z}|A \cdot p', \quad (14)$$

we have that both:

1. For all $h \in [H]$, $a_h \in \mathcal{A}$, and $z_h \notin \mathcal{Z}_{h, \zeta}^{\text{low}}(\widehat{\pi}^h)$, there are at least N_1 values of i (in the h th step of the loop in step 4) so that $(a_h^i = a_h)$ and $\text{suff}_{h-1}(z_h^i) = \text{suff}_{h-1}(z_h)$;
2. For all $h \in [H]$, $a_h \in \mathcal{Z}$, $z_h \in \mathcal{Z}$ for which there are at least N_1 values of i so that $(a_h^i = a_h)$ and $\text{suff}_{h-1}(z_h^i) = \text{suff}_{h-1}(z_h)$, (12) holds; and
3. Thus, if (12) does not hold, we must have that there are fewer than N_1 values of i so that $(z_h^i, a_h^i) = (z_h, a_h)$, which means, by step 15, that $z_h = (a_{h-L:h-1}, o_{h-L+1:h})$ transitions to $(a_{h-L+1:h}, (o_{h-L+2:o^{\text{sink}}}))$ with probability 1.

The above points verify the two points in the lemma statement.

Finally we must verify that our settings N_0, N_1 satisfy the above constraints and lead to the probability in (14) being at least $1 - p$. In particular, we choose $p' = p / (2|\mathcal{Z}|A)$, so that (13) holds as long as we set

$$N_1 = \frac{C_0 L A^{L+1} O^L \log(AO/p)}{\theta^2},$$

for sufficiently large constant C_0 . Finally, note that $|\mathcal{Z}|A \cdot \exp(-\zeta N_0 / (8A)) \leq p/2$ if

$$N_0 > \log\left(\frac{2|\mathcal{Z}|A}{p}\right) \cdot \frac{8A}{\zeta},$$

which holds as long as the constant C_0 in (11) is sufficiently large. Thus, the probability of \mathcal{E}^{low} in (14) is at least $1 - p$, as desired. \square

E.3 Construction of Truncated POMDPs $\overline{\mathcal{P}}_{\phi,h}(\pi^{1:H})$.

One of the main technical challenges is accounting for states of \mathcal{P} which the current policy set *underexplores*, i.e. fails to explore with non-trivial probability. To deal with this issue, given the POMDP \mathcal{P} and a sequence of policies $\pi^{1:H}$, we introduce the notion of *truncated POMDPs*. We start with a description of how to construct a sequence of truncated POMDPs $\overline{\mathcal{P}}_{\phi,h}(\pi^{1:H})$ for $h \in [H]$, which will aid in our analysis relating \mathcal{P} and $\widehat{\mathcal{M}}^{(k)}$. Note that (like the MDPs $\widehat{\mathcal{M}}(\pi^{1:H})$) truncated POMDPs do not appear in our algorithms; they are solely an analytic tool.

We begin by making the following definition of underexplored sets of states.

Definition E.2. Given a real number $\phi \in (0, 1)$, a POMDP \mathcal{P} , and a general policy π' , define the ϕ -underexplored set for π' , denoted $\mathcal{U}_{\phi}^{\mathcal{P}}(\pi')$, to be

$$\mathcal{U}_{\phi}^{\mathcal{P}}(\pi') := \left\{ (h, s) \in [H] \times \mathcal{S} : d_{\mathcal{S},h}^{\mathcal{P},\pi'}(s) < \phi \right\}.$$

For some $h \in [H]$, we also define the ϕ -underexplored set at step h , as:

$$\mathcal{U}_{\phi,h}^{\mathcal{P}}(\pi') := \{s \in \mathcal{S} : (h, s) \in \mathcal{U}_{\phi}^{\mathcal{P}}(\pi')\}.$$

In words, $\mathcal{U}_{\phi}^{\mathcal{P}}(\pi')$ is the set of pairs $(h, s) \in [H] \times \mathcal{S}$ so that π' does not put sufficient mass on state s at step h .

Given a sequence of general policies $\pi^{1:H}$ as above together with a parameter $\phi > 0$ representing a threshold visitation probability and some $H' \in [H]$, we proceed to describe how to construct a POMDP $\overline{\mathcal{P}}_{\phi,H'}(\pi^{1:H}) = (H, \overline{\mathcal{S}}, \mathcal{A}, \overline{\mathcal{O}}, b_1, \overline{\mathbb{T}}^{H'}, \overline{\mathbb{O}}^{H'})$, whose observation space is $\overline{\mathcal{O}} = \mathcal{O} \cup \{o^{\text{sink}}\}$, and whose state space is $\overline{\mathcal{S}} = \mathcal{S} \cup \{s^{\text{sink}}\}$. The observation matrices of $\overline{\mathcal{P}}_{\phi,H'}(\pi^{1:H})$, denoted by $\overline{\mathbb{O}}_h^{H'}$, are equal to those of \mathcal{P} , i.e., we have $\overline{\mathbb{O}}_h^{H'} = \mathbb{O}_h^{\mathcal{P}}$ for all $h \in [H]$. The state transitions of $\overline{\mathcal{P}}_{\phi,H'}(\pi^{1:H})$, denoted by $\overline{\mathbb{T}}^{H'}$, are constructed inductively (with respect to H') according to the following procedure:

1. First set $\overline{\mathcal{P}}_{\phi,1}(\pi^{1:H}), \dots, \overline{\mathcal{P}}_{\phi,L}(\pi^{1:H})$ to have identical transitions $\overline{\mathbb{T}}^1, \dots, \overline{\mathbb{T}}^L$ (respectively) to those of \mathcal{P} out of any state $s \in \overline{\mathcal{S}}$ at each step $h \in [H]$.
2. For $H' = L + 1, L + 2, \dots, H$:
 - (a) Define the transition matrices $\overline{\mathbb{T}}^{H'}$ of $\overline{\mathcal{P}}_{\phi,H'}(\pi^{1:H})$ as follows:
 - (b) Initially set $\overline{\mathbb{T}}^{H'} \leftarrow \overline{\mathbb{T}}^{H'-1}$.
 - (c) Then, for each state $s \in \mathcal{U}_{\phi,H'-L}^{\overline{\mathcal{P}}_{\phi,H'-1}(\pi^{1:H})}(\pi^{H'})$:
 - For all $s' \in \mathcal{S}$ and $a \in \mathcal{A}$, add $\mathbb{T}_{H'-L-1}^{H'}(s|s', a)$ to $\overline{\mathbb{T}}_{H'-L-1}^{H'}(s^{\text{sink}}|s', a)$ and set $\overline{\mathbb{T}}_{H'-L-1}^{H'}(s|s', a)$ to 0.

Note that for all H' , the above construction ensures that (at the end of the above procedure) $\overline{\mathbb{T}}_h^{H'}(s^{\text{sink}}|s^{\text{sink}}, a) = 1$ for all $h \in [H], a \in \mathcal{A}$.

The main point of the construction of the truncated POMDPs $\overline{\mathcal{P}}_{\phi,H'}(\pi^{1:H})$ is to ensure that the following lemma holds:

Lemma E.2. Consider a sequence of general policies $\pi^{1:H}$ and $\phi > 0$. For each $H' \in [H]$, the truncated POMDP $\overline{\mathcal{P}}_{\phi,H'}(\pi^{1:H})$ satisfies the following: for all $(s, h) \in \mathcal{S} \times [H]$ with $L < h \leq H'$, if there is some policy $\pi \in \Pi^{\text{gen}}$ with $d_{\mathcal{S},h-L}^{\overline{\mathcal{P}}_{\phi,H'}(\pi^{1:H}),\pi}(s) > 0$, then $d_{\mathcal{S},h-L}^{\overline{\mathcal{P}}_{\phi,H'}(\pi^{1:H}),\pi^h}(s) \geq \phi$, i.e., $s \notin \mathcal{U}_{\phi,h-L}^{\overline{\mathcal{P}}_{\phi,H'}(\pi^{1:H})}(\pi^h)$.

Proof. For each $h \in [H]$ let $\overline{\mathbb{T}}^h$ denote the state transitions of $\overline{\mathcal{P}}_{\phi,h}(\pi^{1:H})$.

Since the transitions of $\overline{\mathcal{P}}_{\phi,H'}(\pi^{1:H})$ and $\overline{\mathcal{P}}_{\phi,h}(\pi^{1:H})$ do not differ prior to step h , it suffices to consider the case that $H' = h$. So suppose that $d_{\mathcal{S},h-L}^{\overline{\mathcal{P}}_{\phi,h}(\pi^{1:H}),\pi}(s) > 0$ for some $\pi \in \Pi^{\text{gen}}$. Then

by construction of $\overline{\mathcal{P}}_{\phi,h}(\pi^{1:H})$ it must be the case that $s \notin \mathcal{U}_{\phi,h-L}^{\overline{\mathcal{P}}_{\phi,h-1}(\pi^{1:H})}(\pi^h)$. Since \mathbb{T}^h and \mathbb{T}^{h-1} only differ in the transitions at step $h-L-1$ and furthermore for all $s' \in \overline{\mathcal{S}}$ and $a \in \mathcal{A}$, $\overline{\mathbb{T}}_{h-L-1}^{h-1}(s|s', a) = \overline{\mathbb{T}}_{h-L-1}^h(s|s', a)$, we have that $d_{\mathcal{S},h-L}^{\overline{\mathcal{P}}_{\phi,h-1}(\pi^{1:H}),\pi^h}(s) = d_{\mathcal{S},h-L}^{\overline{\mathcal{P}}_{\phi,h}(\pi^{1:H}),\pi^h}(s)$. Thus $s \notin \mathcal{U}_{\phi,h-L}^{\overline{\mathcal{P}}_{\phi,h}(\pi^{1:H})}(\pi^h)$, as desired. \square

F Generic truncated POMDPs

We next define a generic notion of truncation (of a POMDP \mathcal{P}), which is satisfied by the truncated POMDPs $\overline{\mathcal{P}}_{\phi,h}(\pi^{1:H})$ as constructed in the previous section (Lemma F.1 below).

Definition F.1. We say that a POMDP $\overline{\mathcal{P}}$ with extended state and observation spaces $\overline{\mathcal{S}}, \overline{\mathcal{O}}$ is a *truncation* of \mathcal{P} if for all general policies $\pi \in \Pi^{\text{gen}}$, $h \in [H]$, $a'_{1:h-1} \in \mathcal{A}^{h-1}$, $o'_{2:h} \in \mathcal{O}^h$, $s'_{1:h} \in \mathcal{S}^h$, it holds that

$$\begin{aligned} & \mathbb{P}_{(a_{1:h-1}, s_{1:h}, o_{2:h}) \sim \pi}^{\overline{\mathcal{P}}}((a_{1:h-1}, s_{1:h}, o_{2:h}) = (a'_{1:h-1}, s'_{1:h}, o'_{2:h})) \\ & \leq \mathbb{P}_{(a_{1:h-1}, s_{1:h}, o_{2:h}) \sim \pi}^{\mathcal{P}}((a_{1:h-1}, s_{1:h}, o_{2:h}) = (a'_{1:h-1}, s'_{1:h}, o'_{2:h})). \end{aligned} \quad (15)$$

Furthermore, we require that $\mathbb{T}_h^{\overline{\mathcal{P}}}(s^{\text{sink}}|s^{\text{sink}}, a) = 1$ for all $h \in [H]$ and $a \in \mathcal{A}$.

Note that an immediate consequence of (15) is that for all $(h, s) \in [H] \times \mathcal{S}$, $d_{\mathcal{S},h}^{\overline{\mathcal{P}},\pi}(s) \geq d_{\mathcal{S},h}^{\mathcal{P},\pi}(s)$.

Lemma F.1. Fix any sequence of general policies $\pi^{1:H}$ and $\phi > 0$. Then for each $H', H'' \in [H]$ satisfying $H'' \leq H'$, the POMDP $\overline{\mathcal{P}}_{\phi,H'}(\pi^{1:H})$ as constructed above is a truncation of $\overline{\mathcal{P}}_{\phi,H''}(\pi^{1:H})$ (per Definition F.1).

Since $\overline{\mathcal{P}}_{\phi,1}(\pi^{1:H}) = \overline{\mathcal{P}}$, an immediate consequence of Lemma F.1 is that for any $H' \in [H]$, $\overline{\mathcal{P}}_{\phi,H'}(\pi^{1:H})$ is a truncation of \mathcal{P} .

Proof of Lemma F.1. Fix any $H'' \leq H'$. Write $\overline{\mathcal{P}}_{H'} = \overline{\mathcal{P}}_{\phi,H'}(\pi^{1:H})$ and $\overline{\mathcal{P}}_{H''} = \overline{\mathcal{P}}_{\phi,H''}(\pi^{1:H})$. By convexity it suffices to consider the case that π is a deterministic policy, i.e., $\pi \in \Pi^{\text{det}}$. The proof uses induction on h to verify (15), noting that the base case $h = 1$ is immediate since the initial state distributions of $\overline{\mathcal{P}}_{H''}$ and $\overline{\mathcal{P}}_{H'}$ are identical.

Fix any $h \in [H]$, and suppose that (15) holds at all steps up to (and including) step h . Let $\tau = (s_1, a_1, s_2, o_2, a_2, \dots, s_{h-1}, o_{h-1}, a_{h-1}, s_h, o_h)$ denote a trajectory drawn according to the policy π up to step h . Our inductive assumption gives that $\mathbb{P}_{\pi}^{\overline{\mathcal{P}}_{H'}}(\tau) \leq \mathbb{P}_{\pi}^{\overline{\mathcal{P}}_{H''}}(\tau)$. Since π is a deterministic policy, the action chosen by π at step h is given by $\pi_h(a_{1:h-1}, o_{2:h})$, i.e., it is a function of τ . Furthermore, for any $a \in \mathcal{A}$ and $s' \in \mathcal{S}$, we have that $\mathbb{T}_h^{\overline{\mathcal{P}}_{H''}}(s'|s_h, a) \geq \mathbb{T}_h^{\overline{\mathcal{P}}_{H'}}(s'|s_h, a)$ by definition of $\overline{\mathcal{P}}_{H'}, \overline{\mathcal{P}}_{H''}$. Since furthermore $\mathbb{O}_{h+1}^{\overline{\mathcal{P}}_{H''}}(o|s') = \mathbb{O}_{h+1}^{\overline{\mathcal{P}}_{H'}}(o|s')$ for all $o \in \overline{\mathcal{O}}$, $s' \in \overline{\mathcal{S}}$, we get that (15) holds at step $h+1$, as desired. \square

The below lemma establishes some additional properties of the truncated POMDPs $\overline{\mathcal{P}}_{\phi,h}(\pi^{1:H})$.

Lemma F.2. Fix any sequence of general policies $\pi^{1:H}$ and $\phi > 0$. Then for each $h \in [H]$, we have the following:

1. For any general policy $\pi \in \Pi^{\text{gen}}$ and all $(h, s) \in [H] \times \mathcal{S}$ with $h > L$, it holds that

$$d_{\mathcal{S},h-L}^{\pi,\mathcal{P}}(s) \leq d_{\mathcal{S},h-L}^{\pi,\overline{\mathcal{P}}_{\phi,h}(\pi^{1:H})}(s) + d_{\mathcal{S},h-L}^{\pi,\overline{\mathcal{P}}_{\phi,h}(\pi^{1:H})}(s^{\text{sink}}).$$

2. For any $H', H'' \in [H]$ satisfying $H'' \leq H'$, and any $(s, h) \in [H] \times \mathcal{S}$, it holds that $d_{\mathcal{S},h}^{\pi,\overline{\mathcal{P}}_{\phi,H'}(\pi^{1:H})}(s) \leq d_{\mathcal{S},h}^{\overline{\mathcal{P}}_{\phi,H''}(\pi^{1:H}),\pi}(s)$, and thus $d_{\mathcal{S},h}^{\overline{\mathcal{P}}_{\phi,H''}(\pi^{1:H}),\pi}(s^{\text{sink}}) \leq d_{\mathcal{S},h}^{\overline{\mathcal{P}}_{\phi,H'}(\pi^{1:H}),\pi}(s^{\text{sink}})$.

3. For any $h, h' \in [H]$ satisfying $h' \leq h$, as well as any $H' \in [H]$, it holds that $d_{\mathcal{S},h'}^{\overline{\mathcal{P}}_{\phi,H'}(\pi^{1:H}),\pi}(s^{\text{sink}}) \leq d_{\mathcal{S},h}^{\overline{\mathcal{P}}_{\phi,H'}(\pi^{1:H}),\pi}(s^{\text{sink}})$.

Proof. For $h \in [H]$, write $\bar{\mathcal{P}}_h := \bar{\mathcal{P}}_{\phi, h}(\pi^{1:H})$. We begin with a proof of the first item. By linearity of expectation it suffices to consider the case that π is a deterministic policy, i.e., we have $\pi = (\pi_1, \dots, \pi_H)$, where each $\pi_h : \mathcal{A}^{h-1} \times \mathcal{O}^{h-1} \rightarrow \mathcal{A}$. Consider any trajectory up to step $h-L$, denoted $\tau = (s_1, a_1, s_2, o_2, a_2, \dots, s_{h-L}, o_{h-L})$, which has positive probability under \mathcal{P} when policy π is used. For such a trajectory τ and $h' \leq h-L$, let $s_{h'}(\tau)$ denote the state in τ at step h' (i.e., the state $s_{h'}$). By the construction of $\bar{\mathcal{P}}_h$, if τ also has positive probability under $\bar{\mathcal{P}}_h$ for the policy π , then $\mathbb{P}_\pi^{\mathcal{P}}(\tau) = \mathbb{P}_\pi^{\bar{\mathcal{P}}_h}(\tau)$. On the other hand, it is evident from construction of $\bar{\mathcal{P}}_h = \bar{\mathcal{P}}_{\phi, h}(\pi^{1:H})$ that

$$d_{\mathcal{S}, h-L}^{\bar{\mathcal{P}}_h, \pi}(s^{\text{sink}}) = \sum_{\tau: \mathbb{P}_\pi^{\bar{\mathcal{P}}_h}(\tau)=0, s_{h-L}(\tau) \in \mathcal{S}} \mathbb{P}_\pi^{\mathcal{P}}(\tau). \quad (16)$$

To see that the above equality holds, consider any trajectory $\tau = (s_1, a_1, s_2, o_2, a_2, \dots, s_{h-L}, o_{h-L})$ with $s_{h-L} = s^{\text{sink}}$ and which occurs with positive probability under $\bar{\mathcal{P}}_h, \pi$. Choose h' as small as possible so that $s_{h'} = s^{\text{sink}}$. We associate τ with the set $\mathcal{T}(\tau)$ of trajectories that agree with τ up to step $h'-1$, and then transition to a state $\tilde{s}_{h'}$ from which mass was diverted away in the construction of $\bar{\mathcal{P}}_{h'+L}$ (to be precise, the set of such states is $\mathcal{U}_{\phi, h'+L}^{\bar{\mathcal{P}}_{h'+L-1}}(\pi^{h'+L})$). It is clear that all trajectories $\tau' \in \mathcal{T}(\tau)$ have $\mathbb{P}_\pi^{\bar{\mathcal{P}}_h}(\tau') = 0$, and moreover, $\mathbb{P}_\pi^{\bar{\mathcal{P}}_h}(\tau) = \sum_{\tau' \in \mathcal{T}(\tau)} \mathbb{P}_\pi^{\mathcal{P}}(\tau')$. Furthermore, any trajectory τ' with $s_{h-L}(\tau') \in \mathcal{S}$, $\mathbb{P}_\pi^{\mathcal{P}}(\tau') > 0$, and $\mathbb{P}_\pi^{\bar{\mathcal{P}}_h}(\tau') = 0$ must be in some (unique) $\mathcal{T}(\tau)$, for some τ with $s_{h-L}(\tau) = s^{\text{sink}}$. Combining these facts verifies (16).

Then we have

$$\begin{aligned} d_{\mathcal{S}, h-L}^{\mathcal{P}, \pi}(s) &= \sum_{\tau: s_{h-L}(\tau)=s} \mathbb{P}_\pi^{\mathcal{P}}(\tau) \leq \sum_{\tau: s_{h-L}(\tau)=s} \mathbb{P}_\pi^{\bar{\mathcal{P}}_h}(\tau) + d_{\mathcal{S}, h-L}^{\bar{\mathcal{P}}_h, \pi}(s^{\text{sink}}) \\ &= d_{\mathcal{S}, h-L}^{\bar{\mathcal{P}}_h, \pi}(s) + d_{\mathcal{S}, h-L}^{\bar{\mathcal{P}}_h, \pi}(s^{\text{sink}}), \end{aligned} \quad (17)$$

which completes the proof of the first item.

The second item of the lemma statement is a direct consequence of the fact that $\bar{\mathcal{P}}_{\phi, H'}(\pi^{1:H})$ is a truncation of $\bar{\mathcal{P}}_{\phi, H'}(\pi^{1:H})$ (Lemma F.1).

Finally, the third item of the lemma statement follows from the fact that for all $h \in [H]$ and $a \in \mathcal{A}$, $\mathbb{T}_h^{\bar{\mathcal{P}}_{\phi, H'}(\pi^{1:H})}(s^{\text{sink}} | s^{\text{sink}}, a) = 1$. \square

For a non-negative vector $\mathbf{v} \in \mathbb{R}_{\geq 0}^d$, we define $\mathbf{n}(\mathbf{v}) \in \Delta([d])$ by, for $i \in [d]$, $\mathbf{n}(\mathbf{v})(i) = \frac{\mathbf{v}(i)}{\sum_{j=1}^d \mathbf{v}(j)}$. Lemma F.3 below can be viewed as a ‘‘one-sided’’ corollary of the belief contraction result in Theorem B.2, applied to a truncated POMDP $\bar{\mathcal{P}}$. It makes a somewhat weaker assumption on the distribution \mathcal{D} in Theorem B.2 (for a particular type of choice for \mathcal{D}), at the cost of only obtaining a one-sided inequality on the difference between the true beliefs $\mathbf{b}_h^{\bar{\mathcal{P}}}$ and the approximate beliefs $\tilde{\mathbf{b}}_h^{\pi'}$.

Lemma F.3. Fix any $\phi > 0$, $h \in [H]$, POMDP $\bar{\mathcal{P}}$ with extended state and observation spaces $\bar{\mathcal{S}}, \bar{\mathcal{O}}$ which is a truncation of \mathcal{P} , and general policy $\pi' \in \Pi^{\text{gen}}$. Suppose that $\mathbb{T}_{h'}^{\bar{\mathcal{P}}} = \mathbb{T}_{h'}^{\mathcal{P}}$, for $h' \geq h-L$ and $\mathbb{O}_{h'}^{\bar{\mathcal{P}}} = \mathbb{O}_{h'}^{\mathcal{P}}$, for all $h' \in [H]$. Suppose further that π' satisfies the following: for each $s \in \mathcal{S}$, if $h > L$ and there is some policy $\pi \in \Pi^{\text{gen}}$ so that $d_{\mathcal{S}, h-L}^{\bar{\mathcal{P}}, \pi}(s) > 0$, then $s \notin \mathcal{U}_{\phi, h-L}^{\bar{\mathcal{P}}}(\pi')$. Then for general policies π , it holds that

$$\mathbb{E}_{a_{1:h-1}, o_{2:h} \sim \pi} \left[\sum_{s \in \mathcal{S}} \left| \mathbf{b}_h^{\bar{\mathcal{P}}}(a_{1:h-1}, o_{2:h})(s) - \tilde{\mathbf{b}}_h^{\pi'}(a_{h-L:h-1}, o_{h-L+1:h})(s) \right| \right] \leq \epsilon.$$

Proof of Lemma F.3. Set $\underline{h} := \max\{h-L, 1\}$. Furthermore, define $\mathcal{D} = d_{\mathcal{S}, \underline{h}}^{\mathcal{P}, \pi'}$. Note that, if $\underline{h} = 1$, then we have $\mathcal{D} = b_1$, the initial state distribution of \mathcal{P} . By definition of $\tilde{\mathbf{b}}_h^{\pi'}$ and by our assumption that $\mathbb{T}_{h'}^{\bar{\mathcal{P}}} = \mathbb{T}_{h'}^{\mathcal{P}}$ and $\mathbb{O}_{h'+1}^{\bar{\mathcal{P}}} = \mathbb{O}_{h'+1}^{\mathcal{P}}$ for all $h' \geq h-L$, we have that

$$\tilde{\mathbf{b}}_h^{\pi'}(a_{\underline{h}:h-1}, o_{\underline{h}+1:h}) = \mathbf{b}_h^{\text{apx}, \mathcal{P}}(a_{\underline{h}:h-1}, o_{\underline{h}+1:h}; \mathcal{D}) = \mathbf{b}^{\text{apx}, \bar{\mathcal{P}}}(a_{\underline{h}:h-1}, o_{\underline{h}+1:h}; \mathcal{D}).$$

Consider any general policy π , and fix any sequence $(a_{1:\underline{h}-1}, o_{2:\underline{h}}) \in (\mathcal{A} \times \mathcal{O})^{\underline{h}-1}$ that occurs with positive probability under $(\bar{\mathcal{P}}, \pi)$. If $o_{\underline{h}} = o^{\text{sink}}$, then for all sequences $(a_{\underline{h}:h-1}, o_{\underline{h}+1:h})$ that occur with positive probability under $(\bar{\mathcal{P}}, \pi)$ conditioned on $(a_{1:\underline{h}-1}, o_{2:\underline{h}})$ (which in particular implies that $o_{h'} = o^{\text{sink}}$ for all $\underline{h} + 1 \leq h' \leq h$), we have

$$\mathbf{b}_h^{\bar{\mathcal{P}}}(a_{1:h-1}, o_{2:h})(s) = \tilde{\mathbf{b}}_h^{\pi'}(a_{\underline{h}:h-1}, o_{\underline{h}+1:h})(s) = 0$$

for all $s \in \mathcal{S}$. This uses our convention to define $\tilde{\mathbf{b}}_h^{\pi'}(a_{\underline{h}:h-1}, o_{\underline{h}+1:h})(s^{\text{sink}}) = 1$ (see Section B) when any of $o_{\underline{h}+1}, \dots, o_h$ are equal to o^{sink} .

Next suppose that $o_{\underline{h}} \neq o^{\text{sink}}$; then $\mathbf{b}_h^{\bar{\mathcal{P}}}(a_{1:\underline{h}-1}, o_{2:\underline{h}})(s^{\text{sink}}) = 0$. If $\underline{h} > 1$, then for any $s \in \bar{\mathcal{S}}$ so that $\mathbf{b}_h^{\bar{\mathcal{P}}}(a_{1:\underline{h}-1}, o_{2:\underline{h}})(s) > 0$, we must have that $d_{\mathcal{S}, \underline{h}}^{\bar{\mathcal{P}}, \pi}(s) > 0$ and $s \in \mathcal{S}$, and therefore, $s \notin \mathcal{U}_{\phi, \underline{h}}^{\bar{\mathcal{P}}}(\pi')$ (by assumption), i.e., $\mathcal{D}(s) = d_{\mathcal{S}, \underline{h}}^{\bar{\mathcal{P}}, \pi'}(s) \geq d_{\mathcal{S}, \underline{h}}^{\bar{\mathcal{P}}, \pi}(s) \geq \phi$. On the other hand, if $\underline{h} = 1$, it holds that $\mathcal{D} = b_1 = \mathbf{b}_1^{\bar{\mathcal{P}}}(\emptyset)$. Therefore, by Theorem B.2 applied to the POMDP $\bar{\mathcal{P}}$, we have

$$\mathbb{E}_{\substack{s_{\underline{h}} \sim \mathbf{b}_h^{\bar{\mathcal{P}}}(a_{1:\underline{h}-1}, o_{2:\underline{h}}), \\ a_{\underline{h}:h-1}, o_{\underline{h}+1:h} \sim \pi}}^{\bar{\mathcal{P}}} \left[\left\| \mathbf{b}_h^{\bar{\mathcal{P}}}(a_{1:h-1}, o_{2:h}) - \tilde{\mathbf{b}}_h^{\pi'}(a_{\underline{h}:h-1}, o_{\underline{h}+1:h}) \right\|_1 \right] \leq \epsilon.$$

Taking expectation over $(a_{1:\underline{h}-1}, o_{2:\underline{h}}) \sim \pi$ under the POMDP $\bar{\mathcal{P}}$, and using the fact that the conditional distribution of $s_{\underline{h}}$ given $a_{1:\underline{h}-1}, o_{2:\underline{h}}$ is exactly $\mathbf{b}_h^{\bar{\mathcal{P}}}(a_{1:\underline{h}-1}, o_{2:\underline{h}})$, we get

$$\mathbb{E}_{a_{1:h-1}, o_{2:h} \sim \pi}^{\bar{\mathcal{P}}} \left[\sum_{s \in \mathcal{S}} \left| \mathbf{b}_h^{\bar{\mathcal{P}}}(a_{1:h-1}, o_{2:h})(s) - \tilde{\mathbf{b}}_h^{\pi'}(a_{\underline{h}:h-1}, o_{\underline{h}+1:h})(s) \right| \right] \leq \epsilon. \quad (18)$$

as desired. \square

Recall that Lemma E.1 shows that the empirical MDP $\widehat{\mathcal{M}}(\pi^{1:H})$ is close to $\widetilde{\mathcal{M}}(\pi^{1:H})$ at \mathcal{Z} -states which are not in $\mathcal{Z}_{h, \zeta}^{\text{low}}(\pi')$, i.e., those which are not visited too infrequently under the dynamics of \mathcal{P} (recall Definition E.1). Thus, to effectively apply Lemma E.1, we need to bound the probability of reaching such states; Lemma F.4 does so, with respect to the dynamics of any truncated POMDP $\bar{\mathcal{P}}$.

Lemma F.4. *Fix any $\zeta > 0$, $\delta \geq 0$, $h \leq H$ and consider any policies $\pi', \pi \in \Pi^{\text{gen}}$, as well as any POMDP $\bar{\mathcal{P}}$ with extended state and observation spaces $\bar{\mathcal{S}}, \bar{\mathcal{O}}$ which is a truncation of \mathcal{P} . Suppose that π' takes uniformly random actions at each step $h' \in \{\max\{h-L, 1\}, \dots, h\}$, each chosen independently of all previous states, actions, and observations $(s_{1:h'}, a_{1:h'-1}, o_{2:h'})$. Then*

$$d_{\mathcal{Z}, h}^{\bar{\mathcal{P}}, \pi}(\mathcal{Z}_{h, \zeta}^{\text{low}}(\pi')) \leq \frac{A^{2L} O^L \zeta}{\phi} + \mathbb{1}[h > L] \cdot d_{\mathcal{S}, h-L}^{\bar{\mathcal{P}}, \pi}(\mathcal{U}_{\phi, h-L}^{\bar{\mathcal{P}}}(\pi')).$$

Proof. By linearity of expectation, we may assume that π is a deterministic policy, i.e., $\pi \in \Pi^{\text{det}}$.

Fix any h as in the lemma statement, and write $\underline{h} := \max\{1, h-L\}$. By definition of $\mathcal{U}_{\phi, \underline{h}}^{\bar{\mathcal{P}}}(\pi')$, for any $s \notin \mathcal{U}_{\phi, \underline{h}}^{\bar{\mathcal{P}}}(\pi')$, it holds that $d_{\mathcal{S}, \underline{h}}^{\bar{\mathcal{P}}, \pi'}(s) \geq \phi$. For convenience we write

$$\mathcal{U} := \begin{cases} \mathcal{U}_{\phi, \underline{h}}^{\bar{\mathcal{P}}}(\pi') & : \underline{h} > 1 \\ \emptyset & : \underline{h} = 1. \end{cases}$$

throughout the proof of the lemma.

Suppose $\underline{h} = 1$: since $d_{\mathcal{S}, \underline{h}}^{\bar{\mathcal{P}}, \pi}(s) \leq 1$ for all $s \in \mathcal{S}$, it then holds that $\frac{d_{\mathcal{S}, \underline{h}}^{\bar{\mathcal{P}}, \pi}(s)}{d_{\mathcal{S}, \underline{h}}^{\bar{\mathcal{P}}, \pi'}(s)} \leq \frac{1}{\phi}$ for all $s \in \mathcal{S} \setminus \mathcal{U}$. On

the other hand, when $\underline{h} = 1$, it holds that $\frac{d_{\mathcal{S}, \underline{h}}^{\bar{\mathcal{P}}, \pi}(s)}{d_{\mathcal{S}, \underline{h}}^{\bar{\mathcal{P}}, \pi'}(s)} = 1 \leq \frac{1}{\phi}$ for all $s \in \mathcal{S} = \mathcal{S} \setminus \mathcal{U}$.

We next prove by induction that for all h' satisfying $\underline{h} \leq h' \leq h$, it holds that

$$\forall a'_{\underline{h}:h'-1}, o'_{\underline{h}+1:h'}, s'_{\underline{h}:h'} \text{ such that } s'_{\underline{h}} \notin \mathcal{U}, \quad \frac{\mathbb{P}_{\pi}^{\bar{\mathcal{P}}}((a_{\underline{h}:h'-1}, o_{\underline{h}+1:h'}, s_{\underline{h}:h'}) = (a'_{\underline{h}:h'-1}, o'_{\underline{h}+1:h'}, s'_{\underline{h}:h'}))}{\mathbb{P}_{\pi'}^{\bar{\mathcal{P}}}((a_{\underline{h}:h'-1}, o_{\underline{h}+1:h'}, s_{\underline{h}:h'}) = (a'_{\underline{h}:h'-1}, o'_{\underline{h}+1:h'}, s'_{\underline{h}:h'}))} \leq \frac{A^{h'-\underline{h}}}{\phi}. \quad (19)$$

We first note that (19) holds for $h' = \underline{h}$ because $\frac{d_{\mathcal{S}, \underline{h}}^{\pi, \overline{\mathcal{P}}}(s)}{d_{\mathcal{S}, \underline{h}}^{\pi', \overline{\mathcal{P}}}(s)} \leq \frac{1}{\phi}$ for all $s \in \mathcal{S} \setminus \mathcal{U}$.

If (19) holds at some step h' , then we proceed to show that (19) holds at step $h' + 1$: conditioned on any history $a_{1:h'-1}, o_{2:h'}, s_{\underline{h}:h'}$, we have by assumption that π' chooses a uniformly random action at step h' , meaning that with probability $1/A$, π' chooses the action $\pi_{h'}(a_{a:h'-1}, o_{2:h'})$ at step h' . (Recall that we have assumed that π is a deterministic policy.) Thus, we have that for all $a' \in \mathcal{A}$, $s' \in \mathcal{S}$, $o' \in \mathcal{O}$,

$$\frac{\mathbb{P}_{\pi}^{\overline{\mathcal{P}}}(a_{h'} = a', s_{h'+1} = s', o_{h'+1} = o' | a_{\underline{h}:h'-1}, o_{\underline{h}+1:h'}, s_{\underline{h}:h'})}{\mathbb{P}_{\pi'}^{\overline{\mathcal{P}}}(a_{h'} = a', s_{h'+1} = s', o_{h'+1} = o' | a_{\underline{h}:h'-1}, o_{\underline{h}+1:h'}, s_{\underline{h}:h'})} \leq A.$$

Then using (19) at step h' gives (19) at step $h' + 1$.

Given the validity of (19) at step $h' = h$, we sum over all possible values of $s'_{\underline{h}:h} \in \mathcal{S}^{L+1}$ and $(a'_{h-L:h-1}, o'_{h-L+1:h}) \in \mathcal{Z}_{h,\zeta}^{\text{low}}(\pi')$ to obtain that

$$\begin{aligned} & \sum_{(a'_{h-L:h-1}, o'_{h-L+1:h}) \in \mathcal{Z}_{h,\zeta}^{\text{low}}(\pi')} \mathbb{P}_{\pi}^{\overline{\mathcal{P}}}((a_{\underline{h}:h-1}, o_{\underline{h}+1:h}) = (a'_{\underline{h}:h-1}, o'_{\underline{h}+1:h})) \\ & \leq \sum_{(a'_{h-L:h-1}, o'_{h-L+1:h}) \in \mathcal{Z}_{h,\zeta}^{\text{low}}(\pi')} \mathbb{P}_{\pi'}^{\overline{\mathcal{P}}}((a_{\underline{h}:h-1}, o_{\underline{h}+1:h}) = (a'_{\underline{h}:h-1}, o'_{\underline{h}+1:h})) \cdot \frac{A^L}{\phi} + \mathbb{P}_{\pi}^{\overline{\mathcal{P}}}(s_{\underline{h}} \in \mathcal{U}). \end{aligned} \quad (20)$$

(Note that in the case that $h - L \leq 0$, we are multi-counting each value of $(a'_{\underline{h}:h-1}, o'_{\underline{h}+1:h})$ on both sides of the equation.) Consider any $z = (a'_{h-L:h-1}, o'_{h-L+1:h}) \in \mathcal{Z}_{h,\zeta}^{\text{low}}(\pi')$. Since $\overline{\mathcal{P}}$ is a truncation of \mathcal{P} (Definition F.1), we have that

$$\begin{aligned} \mathbb{P}_{\pi'}^{\overline{\mathcal{P}}}((a_{\underline{h}:h-1}, o_{\underline{h}+1:h}) = (a'_{\underline{h}:h-1}, o'_{\underline{h}+1:h})) & \leq \mathbb{P}_{\pi'}^{\mathcal{P}}((a_{\underline{h}:h-1}, o_{\underline{h}+1:h}) = (a'_{\underline{h}:h-1}, o'_{\underline{h}+1:h})) \\ & = d_{\mathcal{Z}, h}^{\mathcal{P}, \pi'}(z) \leq \zeta. \end{aligned}$$

Since $|\mathcal{Z}_{h,\zeta}^{\text{low}}(\pi')| \leq (OA)^L$, it follows from (20) that $d_{\mathcal{Z}, h}^{\overline{\mathcal{P}}, \pi}(\mathcal{Z}_{h,\zeta}^{\text{low}}(\pi')) \leq \frac{A^{2L} O^L \zeta}{\phi} + \sum_{s \in \mathcal{U}} d_{\mathcal{S}, \underline{h}}^{\overline{\mathcal{P}}, \pi}(s) \leq \frac{A^{2L} O^L \zeta}{\phi} + d_{\mathcal{S}, \underline{h}}^{\overline{\mathcal{P}}, \pi}(\mathcal{U})$, as desired (recall that for the case $\underline{h} = 1$, we have $\mathcal{U} = \emptyset$, so that $d_{\mathcal{S}, \underline{h}}^{\overline{\mathcal{P}}, \pi}(\mathcal{U}) = 0$). \square

G One-sided Comparison: Relating $\overline{\mathcal{P}}_{\phi, H'}(\pi^{1:H})$ and $\widehat{\mathcal{M}}(\pi^{1:H})$.

In this section we develop several bounds on the distance between the distributions of action-observation sequences under the truncated POMDPs $\overline{\mathcal{P}}_{\phi, H'}(\pi^{1:H})$ and the MDPs $\widehat{\mathcal{M}}^{(k)}$ constructed in the course of Algorithm 3. We will do this by relating both to the MDPs $\widehat{\mathcal{M}}(\pi^{1:H})$ defined in Section E.1. Roughly speaking, we show that, if we attach a non-negative reward function to $\overline{\mathcal{P}}_{\phi, H'}(\pi^{1:H})$ and $\widehat{\mathcal{M}}(\pi^{1:H})$, then the value of any policy π in $\overline{\mathcal{P}}_{\phi, H'}(\pi^{1:H})$ cannot be much greater than the value of π in $\widehat{\mathcal{M}}(\pi^{1:H})$ (but the inequality in the other direction need not necessarily hold).

Our first lemma, Lemma G.1, can be viewed as a modification of Lemma B.1, which expresses the difference in value functions between two contextual decision processes in terms of the differences between their transitions and rewards. Lemma G.1 in particular applies to the case where one of the contextual decision processes is one of the truncated POMDPs $\overline{\mathcal{P}}_{\phi, H'-1}(\pi^{1:H})$, for some choice of $H', \phi, \pi^{1:H}$ (see Section E.3), the rewards are identical, and the value function is always approximately non-negative. In particular, Lemma G.1 derives an *upper bound* on the value function of this truncated POMDP; the one-sided nature of this upper bound should not be surprising in light of the fact that the truncated POMDPs divert some mass to the sink state, which has 0 value, while the value elsewhere is (approximately) non-negative. Importantly, the upper bound is phrased in terms of the difference between the transition functions of $\overline{\mathcal{P}}_{\phi, h}(\pi^{1:H})$ and \mathcal{P}' , for values of $h \leq H' - 1$; this particular detail is needed to be able to apply Lemma F.3 to further upper bound this expression.

Lemma G.1. Fix the POMDP \mathcal{P} , and consider any reward-free contextual decision process $\mathcal{P}' = (\mathcal{A}, \bar{\mathcal{O}}, H, \mathbb{P})$. Fix any $H' \in [H]$, and consider any collection of reward functions $R = (R_1, \dots, R_H)$, $R_h : \bar{\mathcal{O}} \rightarrow \mathbb{R}$ so that:

- $R_{h'} \equiv 0$ for all $h' \neq H'$;
- $R_{h'}(o^{\text{sink}}) = 0$ for all $h' \in [H]$;
- For some $\psi > 0$, $V_h^{\pi, \mathcal{P}', R}(a_{1:h-1}, o_{2:h}) \geq -\psi$ for all $h \in [H]$ and $(a_{1:h-1}, o_{2:h}) \in \mathcal{H}_h$;
- For any policy $\pi \in \Pi^{\text{gen}}$ and any trajectory $(a_{1:H-1}, o_{2:H})$ with positive probability under (π, \mathcal{P}') , if $o_h = o^{\text{sink}}$ for some h , then $o_{h'} = o^{\text{sink}}$ for all $h' > h$.

Consider any sequence of general policies $\pi^1, \dots, \pi^H \in \Pi^{\text{gen}}$ and $\phi > 0$. Then

$$\begin{aligned} V_1^{\pi, \bar{\mathcal{P}}_{H'-1}, R}(\emptyset) - V_1^{\pi, \mathcal{P}', R}(\emptyset) &\leq \psi H + \sum_{h=1}^{H'-2} \mathbb{E}_{(a_{1:h}, o_{2:h}) \sim \pi}^{\bar{\mathcal{P}}_h} \left[\left(\left(\mathbb{P}_h^{\bar{\mathcal{P}}_h} - \mathbb{P}_h^{\mathcal{P}'} \right) V_{h+1}^{\pi, \mathcal{P}', R} \right) (a_{1:h}, o_{2:h}) \right] \\ &\quad + \mathbb{E}_{(a_{1:H'-1}, o_{2:H'-1}) \sim \pi}^{\bar{\mathcal{P}}_{H'-1}} \left[\left(\left(\mathbb{P}_{H'-1}^{\bar{\mathcal{P}}_{H'-1}} - \mathbb{P}_{H'-1}^{\mathcal{P}'} \right) R_{H'} \right) (a_{1:H'-1}, o_{2:H'-1}) \right], \end{aligned}$$

where we have written $\bar{\mathcal{P}}_h = \bar{\mathcal{P}}_{\phi, h}(\pi^{1:H})$ in the above expression.

Proof. Fix a deterministic policy $\pi \in \Pi^{\text{det}}$. For $h \in [H]$, write $\bar{\mathcal{P}}_h := \bar{\mathcal{P}}_{\phi, h}(\pi^{1:H})$. For $h \in [H]$ satisfying $h > L$, write $\mathcal{U}_{h-L} := \mathcal{U}_{\phi, h-L}^{\bar{\mathcal{P}}_h}(\pi^h)$.

The definition of $\bar{\mathcal{P}}_h$, for each $h \in [H]$, ensures that for any $g \geq h$, and any $a_{1:g-1} \in \mathcal{A}^{g-1}$, $o_{2:g} \in \mathcal{O}^{g-1}$, $s_{1:g} \in \mathcal{S}^g$ so that for all $h' \leq h - L$, $s_{h'} \notin \mathcal{U}_{h'}$, it holds that $\mathbb{P}_\pi^{\mathcal{P}}(s_{1:g}, a_{1:g-1}, o_{2:g}) = \mathbb{P}_\pi^{\bar{\mathcal{P}}_h}(s_{1:g}, a_{1:g-1}, o_{2:g})$. Therefore, for any $h, g \in [H]$ with $g \geq h$, we have, for all $(a_{1:g-1}, o_{2:g}) \in (\mathcal{A} \times \mathcal{O})^{g-1}$,

$$\mathbb{P}_\pi^{\mathcal{P}}(a_{1:g-1}, o_{2:g} \text{ and } \forall h' \leq h - L, s_{h'} \notin \mathcal{U}_{h'}) = \mathbb{P}_\pi^{\bar{\mathcal{P}}_h}(a_{1:g-1}, o_{2:g}). \quad (21)$$

Thus, letting $a_g = \pi_g(a_{1:g-1}, o_{2:g})$, we have

$$\mathbb{P}_g^{\bar{\mathcal{P}}_h}(o_{g+1} | a_{1:g}, o_{2:g}) = \mathbb{P}_\pi^{\mathcal{P}}(o_{g+1} | a_{1:g}, o_{2:g} \text{ and } \forall h' \leq h - L, s_{h'} \notin \mathcal{U}_{h'}). \quad (22)$$

For $h \leq H'$ and $(a_{1:h-1}, o_{2:h}) \in (\mathcal{A} \times \mathcal{O})^{h-1}$, write

$$\begin{aligned} U_h^\pi(a_{1:h-1}, o_{2:h}) &:= \mathbb{E}_\pi^{\mathcal{P}} [R_{H'}(o_{H'}) \cdot \mathbb{1}[\forall h' \leq H' - L - 1, s_{h'} \notin \mathcal{U}_{h'}] \\ &\quad | a_{1:h-1}, o_{2:h}, \forall h' \leq h - L - 1, s_{h'} \notin \mathcal{U}_{h'}]. \end{aligned}$$

From (21) and the fact that $R_{H'}(o^{\text{sink}}) = 0$, we have that $U_1^\pi(\emptyset) = V_1^{\pi, \bar{\mathcal{P}}_{H'-1}, R}(\emptyset)$. Furthermore, for all $a_{1:H'-1}, o_{2:H'}$ so that $\mathbb{P}_\pi^{\mathcal{P}}(a_{1:H'-1}, o_{2:H'}, \forall h' \leq H' - L - 1, s_{h'} \notin \mathcal{U}_{h'}) > 0$, we have $U_{H'}^\pi(a_{1:H'-1}, o_{2:H'}) = R_{H'}(o_{H'})$.

Next for any $h \leq H' - 1$ and $(a_{1:h-1}, o_{2:h}) \in (\mathcal{A} \times \mathcal{O})^{h-1}$, letting $a_h = \pi_h(a_{1:h-1}, o_{2:h})$, we have

$$\begin{aligned} &U_h^\pi(a_{1:h-1}, o_{2:h}) \\ &= \sum_{o_{h+1} \in \mathcal{O}} \mathbb{P}_\pi^{\mathcal{P}} \left(\begin{array}{c} (a_{1:h}, o_{2:h+1}) \text{ and} \\ s_{h-L} \notin \mathcal{U}_{h-L} \end{array} \middle| \begin{array}{c} a_{1:h}, o_{2:h} \text{ and} \\ \forall h' \leq h-L-1, s_{h'} \notin \mathcal{U}_{h'} \end{array} \right) \cdot U_{h+1}^\pi(a_{1:h}, o_{2:h+1}) \\ &= \mathbb{P}_\pi^{\mathcal{P}} \left(s_{h-L} \notin \mathcal{U}_{h-L} \middle| \begin{array}{c} a_{1:h}, o_{2:h} \text{ and} \\ \forall h' \leq h-L-1, s_{h'} \notin \mathcal{U}_{h'} \end{array} \right) \cdot \sum_{o_{h+1} \in \mathcal{O}} \mathbb{P}_\pi^{\mathcal{P}} \left(o_{h+1} \middle| \begin{array}{c} a_{1:h}, o_{2:h} \text{ and} \\ \forall h' \leq h-L, s_{h'} \notin \mathcal{U}_{h'} \end{array} \right) \cdot U_{h+1}^\pi(a_{1:h}, o_{2:h+1}) \\ &= \mathbb{P}_\pi^{\mathcal{P}} \left(s_{h-L} \notin \mathcal{U}_{h-L} \middle| \begin{array}{c} a_{1:h}, o_{2:h} \text{ and} \\ \forall h' \leq h-L-1, s_{h'} \notin \mathcal{U}_{h'} \end{array} \right) \cdot \sum_{o_{h+1} \in \mathcal{O}} \mathbb{P}_\pi^{\bar{\mathcal{P}}_h} (o_{h+1} | a_{1:h}, o_{2:h}) \cdot U_{h+1}^\pi(a_{1:h}, o_{2:h+1}), \end{aligned} \quad (23)$$

where the first equality follows from the definition of U_h^π , the second equality follows from rearranging, and the third equality follows from (22).

For each $h \leq H'$ and $(a_{1:h-1}, o_{2:h}) \in (\mathcal{A} \times \mathcal{O})^{h-1}$, define

$$\tilde{V}_h^\pi(a_{1:h-1}, o_{2:h}) := \begin{cases} V_h^{\pi, \mathcal{P}', R}(a_{1:h-1}, o_{2:h}) & : h \leq H' - 1 \\ R_{H'}(o_{H'}) & : h = H'. \end{cases}$$

Moreover, from the definition of $V_h^{\pi, \mathcal{P}'}(\cdot)$, we have (again for $h \leq H' - 1$)

$$\tilde{V}_h^\pi(a_{1:h-1}, o_{2:h}) = \sum_{o_{h+1} \in \mathcal{O}} \mathbb{P}_h^{\mathcal{P}'}(o_{h+1} | a_{1:h}, o_{2:h}) \cdot \tilde{V}_{h+1}^\pi(a_{1:h}, o_{2:h+1}). \quad (24)$$

The above is immediate for $h < H' - 1$ by definition of $V_h^{\pi, \mathcal{P}', R}(\cdot)$ and the fact that $R_{h'} \equiv 0$ for $h' \neq H'$, whereas for $h = H' - 1$, the fact that $V_{H'}^{\pi, \mathcal{P}', R} \equiv 0$ means that $V_h^{\pi, \mathcal{P}', R}(a_{1:h-1}, o_{2:h}) = (\mathbb{P}_h^{\mathcal{P}'}(R_{h+1} + V_{h+1}^{\pi, \mathcal{P}', R}))(a_{1:h}, o_{2:h}) = (\mathbb{P}_h^{\mathcal{P}'} R_{h+1})(a_{1:h}, o_{2:h}) = (\mathbb{P}_h^{\mathcal{P}'} \tilde{V}_{h+1}^\pi)(a_{1:h}, o_{2:h})$, which is exactly the content of (24). Furthermore, in both cases it is permissible to take the sum over \mathcal{O} (as opposed to $\bar{\mathcal{O}}$) since if $o_{h+1} = o^{\text{sink}}$, then $\tilde{V}_{h+1}^\pi(a_{1:h}, o_{2:h+1}) = 0$ by our assumptions in the lemma statement.

Now, for any $h \in [H' - 1]$, we may bound

$$\begin{aligned} & \sum_{a_{1:h}, o_{2:h}} \mathbb{P}_\pi^{\mathcal{P}} \left(\begin{array}{c} (a_{1:h}, o_{2:h}) \text{ and} \\ \forall h' \leq h-L-1, s_{h'} \notin \mathcal{U}_{h'} \end{array} \right) \cdot \left(U_h^\pi(a_{1:h-1}, o_{2:h}) - \tilde{V}_h^\pi(a_{1:h-1}, o_{2:h}) \right) \\ \leq & \psi + \sum_{a_{1:h}, o_{2:h}} \mathbb{P}_\pi^{\mathcal{P}} \left(\begin{array}{c} (a_{1:h}, o_{2:h}) \text{ and} \\ \forall h' \leq h-L-1, s_{h'} \notin \mathcal{U}_{h'} \end{array} \right) \cdot \mathbb{P}_\pi^{\mathcal{P}} \left(s_{h-L} \notin \mathcal{U}_{h-L} \mid \begin{array}{c} a_{1:h}, o_{2:h} \text{ and} \\ \forall h' \leq h-L-1, s_{h'} \notin \mathcal{U}_{h'} \end{array} \right) \\ & \cdot \left(\sum_{o_{h+1} \in \mathcal{O}} \mathbb{P}_h^{\bar{\mathcal{P}}}(o_{h+1} | a_{1:h}, o_{2:h}) \cdot U_{h+1}^\pi(a_{1:h}, o_{2:h+1}) - \sum_{o_{h+1} \in \mathcal{O}} \mathbb{P}_h^{\mathcal{P}'}(o_{h+1} | a_{1:h}, o_{2:h}) \cdot \tilde{V}_{h+1}^\pi(a_{1:h}, o_{2:h+1}) \right) \quad (25) \\ = & \psi + \sum_{a_{1:h}, o_{2:h}} \mathbb{P}_\pi^{\mathcal{P}} \left(\begin{array}{c} (a_{1:h}, o_{2:h}) \text{ and} \\ \forall h' \leq h-L, s_{h'} \notin \mathcal{U}_{h'} \end{array} \right) \cdot \left(\sum_{o_{h+1} \in \mathcal{O}} \left(\mathbb{P}_h^{\bar{\mathcal{P}}}(o_{h+1} | a_{1:h}, o_{2:h}) - \mathbb{P}_h^{\mathcal{P}'}(o_{h+1} | a_{1:h}, o_{2:h}) \right) \cdot \tilde{V}_{h+1}^\pi(a_{1:h}, o_{2:h+1}) \right) \\ & + \sum_{a_{1:h}, o_{2:h}} \mathbb{P}_\pi^{\mathcal{P}} \left(\begin{array}{c} (a_{1:h}, o_{2:h}) \text{ and} \\ \forall h' \leq h-L, s_{h'} \notin \mathcal{U}_{h'} \end{array} \right) \cdot \sum_{o_{h+1} \in \mathcal{O}} \mathbb{P}_h^{\bar{\mathcal{P}}}(o_{h+1} | a_{1:h}, o_{2:h}) \cdot \left(U_{h+1}^\pi(a_{1:h}, o_{2:h+1}) - \tilde{V}_{h+1}^\pi(a_{1:h}, o_{2:h+1}) \right) \\ = & \psi + \sum_{a_{1:h}, o_{2:h}} \mathbb{P}_\pi^{\mathcal{P}} \left(\begin{array}{c} (a_{1:h}, o_{2:h}) \text{ and} \\ \forall h' \leq h-L, s_{h'} \notin \mathcal{U}_{h'} \end{array} \right) \cdot \left(\sum_{o_{h+1} \in \mathcal{O}} \left(\mathbb{P}_h^{\bar{\mathcal{P}}}(o_{h+1} | a_{1:h}, o_{2:h}) - \mathbb{P}_h^{\mathcal{P}'}(o_{h+1} | a_{1:h}, o_{2:h}) \right) \cdot \tilde{V}_{h+1}^\pi(a_{1:h}, o_{2:h+1}) \right) \\ & + \sum_{a_{1:h+1}, o_{2:h+1}} \mathbb{P}_\pi^{\mathcal{P}} \left(\begin{array}{c} (a_{1:h+1}, o_{2:h+1}) \text{ and} \\ \forall h' \leq h-L, s_{h'} \notin \mathcal{U}_{h'} \end{array} \right) \cdot \left(U_{h+1}^\pi(a_{1:h}, o_{2:h+1}) - \tilde{V}_{h+1}^\pi(a_{1:h}, o_{2:h+1}) \right), \quad (26) \end{aligned}$$

where (25) uses (23), (24), and the fact that $\tilde{V}_h^\pi(a_{1:h-1}, o_{2:h}) \geq -\psi$ for all $h \leq H' - 1$ and all $(a_{1:h-1}, o_{2:h}) \in \mathcal{H}_h$. Furthermore, (26) uses (22) and the fact that $\mathbb{P}_\pi^{\mathcal{P}} \left(\begin{array}{c} (a_{1:h}, o_{2:h+1}) \text{ and} \\ \forall h' \leq h-L, s_{h'} \notin \mathcal{U}_{h'} \end{array} \right) = \mathbb{P}_\pi^{\mathcal{P}} \left(\begin{array}{c} (a_{1:h+1}, o_{2:h+1}) \text{ and} \\ \forall h' \leq h-L, s_{h'} \notin \mathcal{U}_{h'} \end{array} \right)$ for $a_{h+1} = \pi_{h+1}(a_{1:h}, o_{2:h+1})$ and $\mathbb{P}_\pi^{\mathcal{P}} \left(\begin{array}{c} (a_{1:h+1}, o_{2:h+1}) \text{ and} \\ \forall h' \leq h-L, s_{h'} \notin \mathcal{U}_{h'} \end{array} \right) = 0$ whenever $a_{h+1} \neq \pi_{h+1}(a_{1:h}, o_{2:h+1})$.

Thus, by adding the display (26) for $1 \leq h \leq H' - 1$, we see that

$$\begin{aligned}
& U_1^\pi(\emptyset) - V_1^{\pi, \mathcal{P}', R}(\emptyset) \\
&= U_1^\pi(\emptyset) - \tilde{V}_1^\pi(\emptyset) \\
&\leq \sum_{h=1}^{H'-1} \sum_{a_{1:h}, o_{2:h}} \mathbb{P}_\pi^{\mathcal{P}} \left(\begin{array}{c} (a_{1:h}, o_{2:h}) \text{ and} \\ \forall h' \leq h-L, s_{h'} \notin \mathcal{U}_{h'} \end{array} \right) \cdot \left(\sum_{o_{h+1} \in \mathcal{O}} \left(\mathbb{P}_h^{\bar{\mathcal{P}}^h}(o_{h+1}|a_{1:h}, o_{2:h}) - \mathbb{P}_h^{\mathcal{P}'}(o_{h+1}|a_{1:h}, o_{2:h}) \right) \cdot \tilde{V}_{h+1}^\pi(a_{1:h}, o_{2:h+1}) \right) \\
&\quad + \sum_{a_{1:H'}, o_{2:H'}} \mathbb{P}_\pi^{\mathcal{P}} \left(\begin{array}{c} (a_{1:H'}, o_{2:H'}) \text{ and} \\ \forall h' \leq H'-L-1, s_{h'} \notin \mathcal{U}_{h'} \end{array} \right) \cdot \left(U_{H'}^\pi(a_{1:H'-1}, o_{2:H'}) - \tilde{V}_{H'}^\pi(a_{1:H'-1}, o_{2:H'}) \right) + \psi H \\
&= \sum_{h=1}^{H'-2} \sum_{a_{1:h}, o_{2:h}} \mathbb{P}_\pi^{\bar{\mathcal{P}}^h}(a_{1:h}, o_{2:h}) \cdot \left(\sum_{o_{h+1} \in \mathcal{O}} \left(\mathbb{P}_h^{\bar{\mathcal{P}}^h}(o_{h+1}|a_{1:h}, o_{2:h}) - \mathbb{P}_h^{\mathcal{P}'}(o_{h+1}|a_{1:h}, o_{2:h}) \right) \cdot \tilde{V}_{h+1}^\pi(a_{1:h}, o_{2:h+1}) \right) \\
&\quad + \sum_{a_{1:H'}, o_{2:H'}} \mathbb{P}_\pi^{\mathcal{P}} \left(\begin{array}{c} (a_{1:H'}, o_{2:H'}) \text{ and} \\ \forall h' \leq H'-L-1, s_{h'} \notin \mathcal{U}_{h'} \end{array} \right) \cdot (R_{H'}(o_{H'}) - R_{H'}(o_{H'})) + \psi H \\
&= \psi H + \sum_{h=1}^{H'-2} \mathbb{E}_{(a_{1:h}, o_{2:h}) \sim \pi}^{\bar{\mathcal{P}}^h} \left[\left(\left(\mathbb{P}_h^{\bar{\mathcal{P}}^h} - \mathbb{P}_h^{\mathcal{P}'} \right) V_{h+1}^{\pi, \mathcal{P}', R} \right) (a_{1:h}, o_{2:h}) \right] \\
&\quad + \mathbb{E}_{(a_{1:H'-1}, o_{2:H'-1}) \sim \pi}^{\bar{\mathcal{P}}^{H'-1}} \left[\left(\left(\mathbb{P}_{H'-1}^{\bar{\mathcal{P}}^{H'-1}} - \mathbb{P}_{H'-1}^{\mathcal{P}'} \right) R_{H'} \right) (a_{1:H'-1}, o_{2:H'-1}) \right].
\end{aligned}$$

where the final equality follows from the fact that for $(a_{1:h-1}, o_{2:h}) \in \mathcal{H}_h \setminus (\mathcal{A} \times \mathcal{O})^{h-1}$, we have $V_{h+1}^{\pi, \mathcal{P}', R}(a_{1:h}, o_{2:h+1}) = 0$ for all $o_{h+1} \in \bar{\mathcal{O}}$ by our assumptions in the lemma statement. \square

Lemma G.2 shows that the transition function of the MDP $\widehat{\mathcal{M}}(\pi^{1:H})$ approximates that of the truncated POMDP $\bar{\mathcal{P}}_{\phi, h}(\pi^{1:H})$ at step h . In light of the fact that $\bar{\mathcal{P}}_{\phi, h}(\pi^{1:H})$ diverts a potentially large amount of mass to the sink state s^{sink} , it may be somewhat surprising that the result below is two-sided in nature, i.e., we get an upper bound on the total variation distance between $\mathbb{P}_h^{\bar{\mathcal{P}}_{\phi, h}(\pi^{1:H})}(\cdot|z_{1:h}, a_h)$ and $\mathbb{P}_h^{\widehat{\mathcal{M}}(\pi^{1:H})}(\cdot|z_h, a_h)$. This may be explained by the fact that the roll-in trajectory $(z_{1:h}, a_h)$ is drawn from the truncated POMDP $\bar{\mathcal{P}}_{\phi, h}(\pi^{1:H})$, and the (latent state) transition functions of $\bar{\mathcal{P}}_{\phi, h}(\pi^{1:H})$ and \mathcal{P} are identical at step $h-L$ and thereafter. Thus, trajectories $(z_{1:h}, a_h)$ which $\bar{\mathcal{P}}_{\phi, h}(\pi^{1:H})$ does not reach do not contribute to the left-hand side of (27), and trajectories which are reached are not truncated.

Lemma G.2. *Consider any sequence of general policies $\pi^1, \dots, \pi^H \in \Pi^{\text{gen}}$ and $\phi > 0$. Suppose that, in the construction of $\widehat{\mathcal{M}}(\pi^{1:H})$, the event \mathcal{E}^{low} of Lemma E.1 holds. Then for general policies π and all $h \in [H]$,*

$$\mathbb{E}_{(z_{1:h}, a_h) \sim \pi}^{\bar{\mathcal{P}}_{\phi, h}(\pi^{1:H})} \sum_{z_{h+1} \in \bar{\mathcal{Z}}} \left| \mathbb{P}_h^{\bar{\mathcal{P}}_{\phi, h}(\pi^{1:H})}(z_{h+1}|z_{1:h}, a_h) - \mathbb{P}_h^{\widehat{\mathcal{M}}(\pi^{1:H})}(z_{h+1}|z_h, a_h) \right| \leq \epsilon + \theta + \frac{A^{2L} O^L \zeta}{\phi}. \tag{27}$$

Proof. We write $\bar{\mathcal{P}} := \bar{\mathcal{P}}_{\phi, h}(\pi^{1:H})$, $\widetilde{\mathcal{M}} := \widetilde{\mathcal{M}}(\pi^{1:H})$, $\widehat{\mathcal{M}} := \widehat{\mathcal{M}}(\pi^{1:H})$. Throughout the proof of the lemma we write $z_{h'} = (a_{h'-L:h'-1}, o_{h'-L+1:h'})$ for $h' \in [h]$.

First note that

$$\begin{aligned}
& \mathbb{E}_{(z_{1:h}, a_h) \sim \pi}^{\bar{\mathcal{P}}} \sum_{z_{h+1} \in \bar{\mathcal{Z}}} \left| \mathbb{P}_h^{\bar{\mathcal{P}}}(z_{h+1}|z_{1:h}, a_h) - \mathbb{P}_h^{\widehat{\mathcal{M}}}(z_{h+1}|z_h, a_h) \right| \\
&\leq \mathbb{E}_{(z_{1:h}, a_h) \sim \pi}^{\bar{\mathcal{P}}} \sum_{z_{h+1} \in \bar{\mathcal{Z}}} \left| \mathbb{P}_h^{\bar{\mathcal{P}}}(z_{h+1}|z_{1:h}, a_h) - \mathbb{P}_h^{\widetilde{\mathcal{M}}}(z_{h+1}|z_h, a_h) \right| \\
&\quad + \mathbb{E}_{(z_{1:h}, a_h) \sim \pi}^{\bar{\mathcal{P}}} \sum_{z_{h+1} \in \bar{\mathcal{Z}}} \left| \mathbb{P}_h^{\widetilde{\mathcal{M}}}(z_{h+1}|z_{1:h}, a_h) - \mathbb{P}_h^{\widehat{\mathcal{M}}}(z_{h+1}|z_h, a_h) \right|.
\end{aligned}$$

Therefore, to prove the statement of the lemma, it suffices to show that

$$\mathbb{E}_{(z_{1:h}, a_h) \sim \pi} \sum_{z_{h+1} \in \bar{\mathcal{Z}}} \left| \mathbb{P}_h^{\bar{\mathcal{P}}}(z_{h+1}|z_{1:h}, a_h) - \mathbb{P}_h^{\tilde{\mathcal{M}}}(z_{h+1}|z_h, a_h) \right| \leq \epsilon \quad (28)$$

$$\mathbb{E}_{(z_{1:h}, a_h) \sim \pi} \sum_{z_{h+1} \in \bar{\mathcal{Z}}} \left| \mathbb{P}_h^{\tilde{\mathcal{M}}}(z_{h+1}|z_h, a_h) - \mathbb{P}_h^{\widehat{\mathcal{M}}}(z_{h+1}|z_h, a_h) \right| \leq \theta + \frac{A^{2L} O^L \zeta}{\phi}. \quad (29)$$

Note that for any $s \in \mathcal{S}$, if $h > L$ and there is some policy $\pi \in \Pi^{\text{gen}}$ so that $d_{\mathcal{S}, h-L}^{\bar{\mathcal{P}}, \pi}(s) > 0$, then the fact that $\bar{\mathcal{P}} = \bar{\mathcal{P}}_{\phi, h}(\pi^{1:H})$ and Lemma E.2 ensures that $s \notin \mathcal{U}_{\phi, h-L}^{\bar{\mathcal{P}}}(\pi^h)$. Furthermore, we have $\mathbb{T}_{h'}^{\bar{\mathcal{P}}} = \mathbb{T}_{h'}^{\mathcal{P}}$ for all $h' \geq h-L$ and $\mathbb{O}_{h'}^{\bar{\mathcal{P}}} = \mathbb{O}_{h'}^{\mathcal{P}}$ for all $h' \in [H]$. Thus we may apply Lemma F.3 to the POMDP $\bar{\mathcal{P}}$ with $\pi' = \pi^h$, which gives that for general policies π ,

$$\mathbb{E}_{a_{1:h-1}, o_{2:h} \sim \pi} \left[\sum_{s \in \mathcal{S}} \left| \mathbf{b}_h^{\bar{\mathcal{P}}}(a_{1:h-1}, o_{2:h})(s) - \tilde{\mathbf{b}}_h^{\pi^h}(a_{h-L:h-1}, o_{h-L+1:h})(s) \right| \right] \leq \epsilon. \quad (30)$$

Now, for any choice of $(z_{1:h}, a_h) \in \mathcal{Z}^h \times \mathcal{A}$ (which is equivalent to the data of $((a_{1:h-1}, o_{2:h}), a_h)$) we have that, for $o_{h+1} \in \mathcal{O}$, letting $z_{h+1} = (a_{h-L+1:h}, o_{h-L+2:h+1})$,

$$\mathbb{P}_h^{\bar{\mathcal{P}}}(z_{h+1}|z_{1:h}, a_h) = e_{o_{h+1}}^\top \cdot \mathbb{O}_{h+1}^{\mathcal{P}} \cdot \mathbb{T}_h^{\mathcal{P}}(a_h) \cdot \mathbf{b}_h^{\bar{\mathcal{P}}}(a_{1:h-1}, o_{2:h}).$$

On the other hand, (10) gives

$$\mathbb{P}_h^{\tilde{\mathcal{M}}}(z_{h+1}|z_h, a_h) := e_{o_{h+1}}^\top \cdot \mathbb{O}_{h+1}^{\mathcal{P}} \cdot \mathbb{T}_h^{\mathcal{P}}(a_h) \cdot \tilde{\mathbf{b}}_h^{\pi^h}(a_{h-L:h-1}, o_{h-L+1:h}).$$

Moreover, for $z_{h+1} \in \bar{\mathcal{Z}} \setminus \mathcal{Z}$, we have $\mathbb{P}_h^{\bar{\mathcal{P}}}(z_{h+1}|z_{1:h}, a_h) = \mathbb{P}_h^{\tilde{\mathcal{M}}}(z_{h+1}|z_h, a_h) = 0$. Thus, we get that

$$\begin{aligned} & \sum_{z_{h+1} \in \bar{\mathcal{Z}}} \left| \mathbb{P}_h^{\bar{\mathcal{P}}}(z_{h+1}|z_{1:h}, a_{1:h}) - \mathbb{P}_h^{\tilde{\mathcal{M}}}(z_{h+1}|z_h, a_h) \right| \\ & \leq \sum_{o_{h+1} \in \mathcal{O}} \left| e_{o_{h+1}}^\top \cdot \mathbb{O}_{h+1}^{\mathcal{P}} \cdot \mathbb{T}_h^{\mathcal{P}}(a_h) \cdot \left(\mathbf{b}_h^{\bar{\mathcal{P}}}(a_{1:h-1}, o_{2:h}) - \tilde{\mathbf{b}}_h^{\pi^h}(a_{h-L:h-1}, o_{h-L+1:h}) \right) \right| \end{aligned} \quad (31)$$

$$\leq \sum_{s \in \mathcal{S}} \left| \mathbf{b}_h^{\bar{\mathcal{P}}}(a_{1:h-1}, o_{2:h})(s) - \tilde{\mathbf{b}}_h^{\pi^h}(a_{h-L:h-1}, o_{h-L+1:h})(s) \right|, \quad (32)$$

where (32) uses the data processing inequality for total variation distance (together with the fact that, for $(z_{1:h}, a_h) \in \mathcal{Z}^h \times \mathcal{A}$, we have $\mathbf{b}_h^{\bar{\mathcal{P}}}(a_{1:h-1}, o_{2:h})(s^{\text{sink}}) = \tilde{\mathbf{b}}_h^{\pi^h}(a_{h-L:h-1}, o_{h-L+1:h})(s^{\text{sink}}) = 0$). The above display, together with (30) and the fact that for any $z_h \in \bar{\mathcal{Z}} \setminus \mathcal{Z}$, $\mathbb{P}_h^{\bar{\mathcal{P}}}(o^{\text{sink}}|z_{1:h}, a_h) = \mathbb{P}_h^{\tilde{\mathcal{M}}}(o^{\text{sink}}|z_h, a_h) = 1$, gives (28).

Next, we remark that for all $z_h \in \mathcal{Z}$, $a_h \in \mathcal{A}$, it holds that

$$\sum_{z_{h+1} \in \bar{\mathcal{Z}}} \left| \mathbb{P}_h^{\tilde{\mathcal{M}}}(z_{h+1}|z_h, a_h) - \mathbb{P}_h^{\widehat{\mathcal{M}}}(z_{h+1}|z_h, a_h) \right| = \frac{1}{2} \left\| \mathbb{P}_h^{\tilde{\mathcal{M}}}(\cdot|z_h, a_h) - \mathbb{P}_h^{\widehat{\mathcal{M}}}(\cdot|z_h, a_h) \right\|_1 \leq \theta + \mathbb{1}[z_h \in \mathcal{Z}_{h, \zeta}^{\text{low}}(\hat{\pi}^h)],$$

where the inequality follows from Lemma E.1 (and assumption that \mathcal{E}^{low} holds). Taking expectation over $(z_{1:h}, a_{1:h}) \sim (\bar{\mathcal{P}}, \pi)$, we obtain that

$$\begin{aligned} & \mathbb{E}_{(z_{1:h}, a_{1:h}) \sim \pi} \sum_{z_{h+1} \in \bar{\mathcal{Z}}} \left| \mathbb{P}_h^{\tilde{\mathcal{M}}}(z_{h+1}|z_h, a_h) - \mathbb{P}_h^{\widehat{\mathcal{M}}}(z_{h+1}|z_h, a_h) \right| \\ & \leq \theta + \mathbb{P}_{(z_{1:h}, a_{1:h}) \sim \pi}^{\bar{\mathcal{P}}}(z_h \in \mathcal{Z}_{h, \zeta}^{\text{low}}(\hat{\pi}^h)) \\ & = \theta + d_{\mathcal{Z}, h}^{\bar{\mathcal{P}}, \pi}(\mathcal{Z}_{h, \zeta}^{\text{low}}(\hat{\pi}^h)) \\ & \leq \theta + \frac{A^{2L} O^L \zeta}{\phi}, \end{aligned}$$

where the final inequality follows by Lemma F.4 with $\pi' = \pi^h$ and the fact that if $h > L$, for all s with $d_{\mathcal{S}, h-L}^{\bar{\mathcal{P}}, \pi}(s) > 0$, it holds that $s \notin \mathcal{U}_{\phi, h-L}^{\bar{\mathcal{P}}}(\pi^h)$ (Lemma E.2), meaning that $d_{\mathcal{S}, h-L}^{\bar{\mathcal{P}}, \pi}(\mathcal{U}_{\phi, h-L}^{\bar{\mathcal{P}}}(\pi^h)) = 0$. \square

Lemma G.3 shows that for any reward function of a particular form, the associated value function of $\widehat{\mathcal{M}}(\pi^{1:H})$, under any general policy, is nonnegative. Roughly speaking, this holds because the reward functions considered in Lemma G.3 belong to the convex hull of indicators of subsets of (latent) states. Since $\widehat{\mathcal{M}}(\pi^{1:H})$ is an approximation of \mathcal{P} and the value function for such rewards would be non-negative in \mathcal{P} , it is too in $\widehat{\mathcal{M}}(\pi^{1:H})$.

Lemma G.3. *Consider any sequence of general policies $\pi^1, \dots, \pi^H \in \Pi^{\text{gen}}$ and $\phi > 0$. Fix any $H' \in [H]$. Consider any collection of reward functions $R = (R_1, \dots, R_H)$, $R_h : \overline{\mathcal{O}} \rightarrow \mathbb{R}$ so that $R_{h'} \equiv 0$ for all $h' > H'$, and so that for each $h' \leq H'$, there are weights $\alpha_{h',s} \in [0, 1]$, $s \in \mathcal{S}$ so that $R_{h'}(o) = \sum_{s \in \mathcal{S}} \alpha_{h',s} \cdot (\mathbb{O}_{h'}^{\mathcal{P}})_{s,o}^\dagger$ for all $o \in \mathcal{O}$. Suppose that in the construction of $\widehat{\mathcal{M}}(\pi^{1:H})$ the event \mathcal{E}^{low} of Lemma E.1 holds. Then for all $\pi \in \Pi^{\text{gen}}$, $h \in [H]$, and $z_h \in \mathcal{Z}$,*

$$V_h^{\pi, \widehat{\mathcal{M}}(\pi^{1:H}), R}(z_h) \geq -\frac{H(H-h)\sqrt{S}}{\gamma} \cdot \theta.$$

Proof. Write $\widehat{\mathcal{M}} = \widehat{\mathcal{M}}(\pi^{1:H})$. Note that by our assumption on R and Lemma B.3, we have that $|R_h(o)| \leq \frac{\sqrt{S}}{\gamma}$ for all h, o , meaning that $|V_h^{\pi, \widehat{\mathcal{M}}, R}(z_h)| \leq \frac{H\sqrt{S}}{\gamma}$ for all h, z_h .

Note that for all $h \in [H]$ and $z_h \in \overline{\mathcal{Z}} \setminus \mathcal{Z}$, we have $V_h^{\pi, \widehat{\mathcal{M}}, R}(z_h) = 0$, so it suffices to consider the case that $z_h \in \mathcal{Z}$; we will do so via reverse induction on h . For the base case, we have that $V_{H'}^{\pi, \widehat{\mathcal{M}}, R} \equiv 0$ (which follows from $R_{h'} \equiv 0$ for all $h' > H'$). Now fix some $h < H'$, assume that $V_{h+1}^{\pi, \widehat{\mathcal{M}}, R}(z_{h+1}) \geq -\frac{H(H-h-1)\sqrt{S}}{\gamma} \cdot \theta$ for all $z_{h+1} \in \overline{\mathcal{Z}}$, and consider any $z_h = (a_{h-L:h-1}, o_{h-L+1:h}) \in \overline{\mathcal{Z}}$. Then we have, for any $z_h \in \mathcal{Z}$:

$$\begin{aligned} V_h^{\pi, \widehat{\mathcal{M}}, R}(z_h) &= \mathbb{E}_{z_{h+1} \sim \mathbb{P}_h^{\widehat{\mathcal{M}}}(\cdot | z_h, a_h)} \left[R_{h+1}(o(z_{h+1})) + V_{h+1}^{\pi, \widehat{\mathcal{M}}, R}(z_{h+1}) \right] \\ &\geq \mathbb{E}_{z_{h+1} \sim \mathbb{P}_h^{\widehat{\mathcal{M}}}(\cdot | z_h, a_h)} \left[R_{h+1}(o(z_{h+1})) + V_{h+1}^{\pi, \widehat{\mathcal{M}}, R}(z_{h+1}) \right] - \frac{H\sqrt{S}}{\gamma} \cdot \theta \end{aligned} \quad (33)$$

$$\geq \mathbb{E}_{z_{h+1} \sim \mathbb{P}_h^{\widehat{\mathcal{M}}}(\cdot | z_h, a_h)} [R_{h+1}(o(z_{h+1}))] - \frac{H(H-h)\sqrt{S}}{\gamma} \cdot \theta \quad (34)$$

$$\begin{aligned} &= \sum_{o_{h+1} \in \mathcal{O}} \mathbb{P}_h^{\widehat{\mathcal{M}}}((a_{h-L+1:h}, o_{h-L+2:h+1}) | z_h, a_h) \cdot R_{h+1}(o_{h+1}) - \frac{H(H-h)\sqrt{S}}{\gamma} \cdot \theta \\ &= \sum_{s \in \mathcal{S}} \alpha_{h+1,s} \cdot \sum_{o_{h+1} \in \mathcal{O}} (\mathbb{O}_{h+1}^{\mathcal{P}})_{s, o_{h+1}}^\dagger \cdot e_{o_{h+1}}^\top \cdot \mathbb{O}_{h+1}^{\mathcal{P}} \cdot \mathbb{T}_h^{\mathcal{P}}(a_h) \cdot \widetilde{\mathbf{b}}_h^{\pi^h}(a_{h-L:h-1}, o_{h-L+1:h}) \\ &\quad - \frac{H(H-h)\sqrt{S}}{\gamma} \cdot \theta \end{aligned} \quad (35)$$

$$\begin{aligned} &= \sum_{s \in \mathcal{S}} \alpha_{h+1,s} \cdot e_s^\top \cdot \mathbb{T}_h^{\mathcal{P}}(a_h) \cdot \widetilde{\mathbf{b}}_h^{\pi^h}(a_{h-L:h-1}, o_{h-L+1:h}) \\ &\quad - \frac{H(H-h)\sqrt{S}}{\gamma} \cdot \theta \end{aligned} \quad (36)$$

$$\geq -\frac{H(H-h)\sqrt{S}}{\gamma} \cdot \theta, \quad (37)$$

where:

- (33) follows because if z_h, a_h are so that $\left\| \mathbb{P}_h^{\widehat{\mathcal{M}}}(\cdot | z_h, a_h) - \mathbb{P}_h^{\widehat{\mathcal{M}}}(\cdot | z_h, a_h) \right\|_1 > \theta$, then by Lemma E.1 (and assumption that \mathcal{E}^{low} holds) we have that with probability 1 over $z_{h+1} \sim \mathbb{P}_h^{\widehat{\mathcal{M}}}(\cdot | z_h, a_h)$, it holds that $z_{h+1} \notin \mathcal{Z}$, meaning that $R_{h+1}(z_{h+1}) + V_{h+1}^{\pi, \widehat{\mathcal{M}}, R}(z_{h+1}) = 0$ with probability 1.

On the other hand, if $\left\| \mathbb{P}_h^{\widehat{\mathcal{M}}}(\cdot | z_h, a_h) - \mathbb{P}_h^{\widetilde{\mathcal{M}}}(\cdot | z_h, a_h) \right\|_1 \leq \theta$, then the fact that $\left| R_{h+1}(o(z_{h+1})) + V_{h+1}^{\pi, \widehat{\mathcal{M}}, R}(z_{h+1}) \right| \leq \frac{H\sqrt{S}}{\gamma}$ for all $z_{h+1} \in \mathcal{Z}$ establishes (33).

- (34) follows from the inductive hypothesis;
- (35) follows from (10);
- the final inequality in (37) follows since all $\alpha_{h+1,s} \geq 0$ and the entries of $\mathbb{T}_h^{\mathcal{P}}(a_h)$, $\widetilde{\mathbf{b}}_h^{\pi^h}(a_{h-L:h-1}, o_{h-L+1:h})$ are all non-negative.

(37) completes the inductive step, and thus the proof of the lemma. \square

Using Lemmas G.2 and G.3, we next establish Lemma G.4, the main result of the section, which shows that for any reward function which is a special case of the type used in Lemma G.3, the value function of $\overline{\mathcal{P}}_{\phi, H''}(\pi^{1:H})$ is approximately upper bounded by that of $\widehat{\mathcal{M}}(\pi^{1:H})$ (for appropriate choices of H''). Intuitively, the lemma holds since, as established in Lemma G.2, the transitions of $\overline{\mathcal{P}}_{\phi, h}(\pi^{1:H})$ and $\widehat{\mathcal{M}}(\pi^{1:H})$ are close at step h . Furthermore, Lemma G.1 ensures that the difference in value functions of $\overline{\mathcal{P}}_{\phi, H''}(\pi^{1:H})$ and $\widehat{\mathcal{M}}(\pi^{1:H})$ may be bounded above by that difference in transitions at step h .

Lemma G.4. *Consider any sequence of general policies $\pi^1, \dots, \pi^H \in \Pi^{\text{gen}}$ and $\phi > 0$. Fix any $H' \in [H]$. Consider any collection of reward functions $R = (R_1, \dots, R_H)$, $R_h : \overline{\mathcal{O}} \rightarrow \mathbb{R}$ so that $R_h \equiv 0$ for all $h \neq H'$, and so that there are weights $\alpha_s \in [0, 1]$, $s \in \mathcal{S}$ so that $R_{H'}(o) = \sum_{s \in \mathcal{S}} \alpha_s \cdot (\mathbb{O}_{H'}^{\mathcal{P}})_{s,o}^\dagger$ for all $o \in \mathcal{O}$. Suppose that, in the construction of $\widehat{\mathcal{M}}(\pi^{1:H})$, the event \mathcal{E}^{low} of Lemma E.1 holds. Then for all general policies π and all $H'' \geq H' - 1$,*

$$V_1^{\pi, \overline{\mathcal{P}}_{\phi, H''}(\pi^{1:H}), R}(\emptyset) - V_1^{\pi, \widehat{\mathcal{M}}(\pi^{1:H}), R}(\emptyset) \leq \frac{H^2 \sqrt{S} \epsilon}{\gamma} + \frac{2H^3 \theta \sqrt{S}}{\gamma} + \frac{H^2 \sqrt{S} A^{2L} O^L \zeta}{\phi \gamma} \quad (38)$$

In the above lemma statement we consider R to be a reward function for the MDP $\widehat{\mathcal{M}}(\pi^{1:H})$ in the natural way, i.e., via abuse of notation, we use the induced mapping $R_h : \overline{\mathcal{Z}} \rightarrow \mathbb{R}$ defined by $R_h(z) = R_h(o(z))$. Note also that for all $h' \in [H]$ we have $R_{h'}(o^{\text{sink}}) = 0$, since $(\mathbb{O}_{h'}^{\mathcal{P}})_{s, o^{\text{sink}}}^\dagger = 0$ for all $s \in \mathcal{S}$.

Proof. For $h \in [H]$, we write $\overline{\mathcal{P}}_h := \overline{\mathcal{P}}_{\phi, h}(\pi^{1:H})$, and furthermore write $\widetilde{\mathcal{M}} := \widetilde{\mathcal{M}}(\pi^{1:H})$, and $\widehat{\mathcal{M}} := \widehat{\mathcal{M}}(\pi^{1:H})$. By Jensen's inequality, it suffices to consider the case that π is a deterministic policy.

Since $\overline{\mathcal{P}}_{H''}$ is a truncation of $\overline{\mathcal{P}}_{H'-1}$ (Lemma F.1), we have that, for any $s \in \mathcal{S}$, $\mathbb{P}_\pi^{\overline{\mathcal{P}}_{H''}}[s_h = s] \leq \mathbb{P}_\pi^{\overline{\mathcal{P}}_{H'-1}}[s_h = s]$. In light of the fact that $V_1^{\pi, \overline{\mathcal{P}}_h, R}(\emptyset) = \mathbb{E}_\pi^{\overline{\mathcal{P}}_h}[\sum_{s \in \mathcal{S}} \alpha_s \cdot \mathbb{1}[s_h = s]]$ for each $h \in [H]$, it follows that $V_1^{\pi, \overline{\mathcal{P}}_{H''}, R}(\emptyset) \leq V_1^{\pi, \overline{\mathcal{P}}_{H'-1}, R}(\emptyset)$. Thus it suffices to upper bound $V_1^{\pi, \overline{\mathcal{P}}_{H'-1}, R}(\emptyset) - V_1^{\pi, \widehat{\mathcal{M}}, R}(\emptyset)$.

Write $\psi := \frac{H^2 \theta \sqrt{S}}{\gamma}$. We will now apply Lemma G.1 to the POMDP \mathcal{P} with $\mathcal{P}' = \widehat{\mathcal{M}}$. Its preconditions are verified as follows: certainly $R_{h'} \equiv 0$ for $h' \neq H'$, and moreover $R_{H'}(o^{\text{sink}}) = 0$ since $(\mathbb{O}_{H'}^{\mathcal{P}})_{s, o^{\text{sink}}}^\dagger = 0$ for all $s \in \mathcal{S}$. By Lemma G.3, $V_{h+1}^{\pi, \widehat{\mathcal{M}}, R}(a_{1:h-1}, o_{2:h}) \geq -\psi$ for all h and $a_{1:h-1}, o_{2:h} \in (\mathcal{A} \times \mathcal{O})^{h-1}$. Furthermore, for all $(a_{1:H-1}, o_{2:H}) \in \mathcal{H}_H$ occurring with positive probability under $(\sigma, \widehat{\mathcal{M}})$ for any policy σ , the definition of $\widehat{\mathcal{M}}(\pi^{1:H})$ ensures that if $o_h = o^{\text{sink}}$ for

some $h \in [H]$, then $o_{h'} = o^{\text{sink}}$ for all $h' > h$. Thus, we may apply Lemma G.1, giving that

$$\begin{aligned}
& V_1^{\pi, \bar{\mathcal{P}}_{H'-1}, R}(\emptyset) - V_1^{\pi, \widehat{\mathcal{M}}, R}(\emptyset) \\
& \leq \sum_{h=1}^{H'-2} \mathbb{E}_{(a_{1:h}, o_{2:h}) \sim \pi} \left[\left((\mathbb{P}_h^{\bar{\mathcal{P}}_h} - \mathbb{P}_h^{\widehat{\mathcal{M}}})(V_{h+1}^{\pi, \widehat{\mathcal{M}}, R}) \right) (a_{1:h}, o_{2:h}) \right] \\
& \quad + \mathbb{E}_{(a_{1:H'-1}, o_{2:H'-1}) \sim \pi} \left[\left((\mathbb{P}_{H'-1}^{\bar{\mathcal{P}}_{H'-1}} - \mathbb{P}_{H'-1}^{\widehat{\mathcal{M}}}) R_{H'} \right) (a_{1:H'-1}, o_{2:H'-1}) \right] + \psi H \\
& \leq \psi H + \frac{H\sqrt{S}}{\gamma} \cdot \sum_{h=1}^{H'-1} \mathbb{E}_{(a_{1:h}, o_{2:h}) \sim \pi} \left[\left\| \mathbb{P}_h^{\bar{\mathcal{P}}_h}(\cdot | a_{1:h}, o_{2:h}) - \mathbb{P}_h^{\widehat{\mathcal{M}}}(\cdot | a_{1:h}, o_{2:h}) \right\|_1 \right] \\
& \leq \frac{H^3\theta\sqrt{S}}{\gamma} + \frac{H^2\sqrt{S}}{\gamma} \cdot \left(\epsilon + \theta + \frac{A^{2L}O^L\zeta}{\phi} \right) \\
& \leq \frac{H^2\sqrt{S}\epsilon}{\gamma} + \frac{2H^3\theta\sqrt{S}}{\gamma} + \frac{H^2\sqrt{S}A^{2L}O^L\zeta}{\phi\gamma}, \tag{39}
\end{aligned}$$

where the third-to-last inequality follows from the fact that $\left| V_h^{\pi, \widehat{\mathcal{M}}, R}(a_{1:h-1}, o_{2:h}) \right| \leq \frac{H\sqrt{S}}{\gamma}$ for all $h \in [H]$ and $(a_{1:h-1}, o_{2:h}) \in \mathcal{H}_h$, and the second-to-last inequality follows from Lemma G.2, applied for each $h \in [H' - 1]$. \square

H Two-sided comparison inequalities

In this section, we prove several inequalities which relate \mathcal{P} and $\widehat{\mathcal{M}}(\pi^{1:H})$, which are similar in spirit to those in Section G (which related $\bar{\mathcal{P}}_{\phi, H'}(\pi^{1:H})$ and $\widehat{\mathcal{M}}(\pi^{1:H})$), but are of a *two-sided* nature. Accordingly, the upper bounds we show on the difference in the value functions for any policy π depend on the probability that π visits the underexplored sets (Definition E.2) associated to the policies $\pi^{1:H}$.

Given $S \in \mathbb{N}$ and a distribution $Q \in \Delta([S])$, as well as a subset $S' \subset [S]$, let $\mathcal{R}_{S'}(Q) \in \mathbb{R}_{\geq 0}^S$ be the vector which is identical to Q except that for $s \in S'$, $\mathcal{R}_{S'}(Q)(s) = 0$. Furthermore, recall that for a non-negative vector $\mathbf{v} \in \mathbb{R}_{\geq 0}^d$, we define $\mathbf{n}(\mathbf{v}) \in \Delta([d])$ by, for $i \in [d]$, $\mathbf{n}(\mathbf{v})(i) = \frac{\mathbf{v}(i)}{\sum_{j=1}^d \mathbf{v}(j)}$.

Lemma H.1. *If $Q \in \Delta^S$, $S' \subset [S]$, then $\|Q - \mathbf{n}(\mathcal{R}_{S'}(Q))\|_1 = 2 \cdot Q(S')$.*

Proof. We compute

$$\begin{aligned}
\|Q - \mathbf{n}(\mathcal{R}_{S'}(Q))\|_1 &= Q(S') + \sum_{s \notin S'} \left| Q(s) - \frac{Q(s)}{1 - Q(S')} \right| \\
&= Q(S') + \sum_{s \notin S'} Q(S') \cdot \left| \frac{Q(s)}{1 - Q(S')} \right| \\
&= 2 \cdot Q(S').
\end{aligned}$$

\square

Lemma H.2 is a corollary of the belief contraction theorem (Theorem B.2), which relaxes the constraint of Theorem B.2 that $\frac{\mathcal{D}'(s)}{\mathcal{D}(s)} \leq \frac{1}{\phi}$ for all $s \in S$ (for particular types of choices for \mathcal{D}' , \mathcal{D}). The cost of doing so is that the upper bound contains a term denoting the probability that π visits a certain set of underexplored states.

Lemma H.2. *There is a constant $C \geq 1$ so that the following holds. For any $\epsilon > 0$, $L \in \mathbb{N}$ so that $L \geq C \cdot \min \left\{ \frac{\log(1/(\epsilon\phi)) \log(\log(1/\phi)/\epsilon)}{\gamma^2}, \frac{\log(1/(\epsilon\phi))}{\gamma^4} \right\}$, it holds that, for all general policies*

$\pi, \pi' \in \Pi^{\text{gen}}$,

$$\begin{aligned} & \mathbb{E}_{a_{1:h-1}, o_{2:h} \sim \pi}^{\mathcal{P}} \left\| \mathbf{b}_h(a_{1:h-1}, o_{2:h}) - \tilde{\mathbf{b}}_h^{\pi'}(a_{h-L:h-1}, o_{h-L+1:h}) \right\|_1 \\ & \leq \epsilon + \mathbb{1}[h > L] \cdot 6 \cdot d_{\mathcal{S}, h-L}^{\mathcal{P}, \pi}(\mathcal{U}_{\phi, h-L}^{\mathcal{P}}(\pi')). \end{aligned}$$

Proof. By linearity of expectation it suffices to prove the result for the case that π is a deterministic policy, i.e., we have $\pi = (\pi_1, \dots, \pi_H)$ where $\pi_h : \mathcal{A}^{h-1} \times \mathcal{O}^{h-1} \rightarrow \mathcal{A}$. Fix any such deterministic policy π and a general policy $\pi' \in \Pi^{\text{gen}}$.

Note that the lemma statement is immediate in the case that $h \leq L$, as we have $\mathbf{b}_h(a_{1:h-1}, o_{2:h}) = \tilde{\mathbf{b}}_h^{\pi'}(a_{h-L:h-1}, o_{h-L+1:h})$ (see Section E.1). Thus, for the remainder of the proof, we may assume $h > L$.

Set $\underline{h} := h - L$. Fix any choice of $\tau := (a_{1:\underline{h}-1}, o_{2:\underline{h}}) \in \mathcal{A}^{\underline{h}-1} \times \mathcal{O}^{\underline{h}-1}$. Let π^τ be the policy defined at steps \underline{h} onward, given the history τ ; in particular, for $h' \geq \underline{h}$, define $\pi_{h'}^\tau : \mathcal{A}^{h'-h+L} \times \mathcal{O}^{h'-h+L} \rightarrow \mathcal{A}$ by $\pi_{h'}^\tau(a_{\underline{h}:h'-1}, o_{\underline{h}+1:h'}) = \pi_{h'}(a_{1:h'-1}, o_{2:h'})$.

Theorem B.2 gives that for any distribution $Q \in \Delta(\mathcal{S})$ so that $\frac{Q(s)}{d_{\mathcal{S}, \underline{h}}^{\mathcal{P}, \pi'}(s)} \leq \frac{1}{\phi}$ for all $s \in \mathcal{S}$ (with $\frac{0}{0}$ interpreted as 0), it holds that

$$\mathbb{E}_{s_{\underline{h}} \sim Q, (a_{\underline{h}:h-1}, o_{\underline{h}+1:h}) \sim \pi^\tau}^{\mathcal{P}} \left\| \mathbf{b}_h^{\text{apx}}(a_{\underline{h}:h-1}, o_{\underline{h}+1:h}; Q) - \tilde{\mathbf{b}}_h^{\pi'}(a_{\underline{h}:h-1}, o_{\underline{h}+1:h}) \right\|_1 \leq \epsilon. \quad (40)$$

Now set $Q^\tau := \mathbf{n} \left(\mathcal{R}_{\mathcal{U}_{\phi, \underline{h}}^{\mathcal{P}}(\pi')}(\mathbf{b}_{\underline{h}}(a_{1:\underline{h}-1}, o_{2:\underline{h}})) \right)$. By Lemma H.1, $\frac{1}{2} \left\| \mathbf{b}_{\underline{h}}(a_{1:\underline{h}-1}, o_{2:\underline{h}}) - Q^\tau \right\|_1 \leq \epsilon^\tau$, where $\epsilon^\tau = \sum_{s \in \mathcal{U}_{\phi, \underline{h}}^{\mathcal{P}}(\pi')} \mathbf{b}_{\underline{h}}(a_{1:\underline{h}-1}, o_{2:\underline{h}})(s)$. We now see that

$$\begin{aligned} & \mathbb{E}_{s_{\underline{h}} \sim \mathbf{b}_{\underline{h}}(a_{1:\underline{h}-1}, o_{2:\underline{h}}), (a_{\underline{h}:h-1}, o_{\underline{h}+1:h}) \sim \pi^\tau}^{\mathcal{P}} \left\| \mathbf{b}_h(a_{1:h-1}, o_{2:h}) - \tilde{\mathbf{b}}_h^{\pi'}(a_{\underline{h}:h-1}, o_{\underline{h}+1:h}) \right\|_1 \\ & \leq 2\epsilon^\tau + \mathbb{E}_{s_{\underline{h}} \sim Q^\tau, (a_{\underline{h}:h-1}, o_{\underline{h}+1:h}) \sim \pi^\tau}^{\mathcal{P}} \left\| \mathbf{b}_h(a_{1:h-1}, o_{2:h}) - \tilde{\mathbf{b}}_h^{\pi'}(a_{\underline{h}:h-1}, o_{\underline{h}+1:h}) \right\|_1 \\ & \leq 2\epsilon^\tau + \mathbb{E}_{s_{\underline{h}} \sim Q^\tau, (a_{\underline{h}:h-1}, o_{\underline{h}+1:h}) \sim \pi^\tau}^{\mathcal{P}} \left\| \mathbf{b}_h^{\text{apx}}(a_{\underline{h}:h-1}, o_{\underline{h}+1:h}; Q^\tau) - \tilde{\mathbf{b}}_h^{\pi'}(a_{\underline{h}:h-1}, o_{\underline{h}+1:h}) \right\|_1 \\ & \quad + \mathbb{E}_{s_{\underline{h}} \sim Q^\tau, (a_{\underline{h}:h-1}, o_{\underline{h}+1:h}) \sim \pi^\tau}^{\mathcal{P}} \left\| \mathbf{b}_h(a_{1:h-1}, o_{2:h}) - \mathbf{b}_h^{\text{apx}}(a_{\underline{h}:h-1}, o_{\underline{h}+1:h}; Q^\tau) \right\|_1 \\ & \leq 6\epsilon^\tau + \mathbb{E}_{s_{\underline{h}} \sim Q^\tau, (a_{\underline{h}:h-1}, o_{\underline{h}+1:h}) \sim \pi^\tau}^{\mathcal{P}} \left\| \mathbf{b}_h^{\text{apx}}(a_{\underline{h}:h-1}, o_{\underline{h}+1:h}; Q^\tau) - \tilde{\mathbf{b}}_h^{\pi'}(a_{\underline{h}:h-1}, o_{\underline{h}+1:h}) \right\|_1 \quad (41) \\ & \leq 6\epsilon^\tau + \epsilon, \quad (42) \end{aligned}$$

where (42) uses (40) together with the fact that $\frac{Q^\tau(s)}{d_{\mathcal{S}, \underline{h}}^{\mathcal{P}, \pi'}(s)} \leq \frac{1}{\phi}$ for all $s \in \mathcal{S}$ (as those $s \in \mathcal{S}$ with $d_{\mathcal{S}, \underline{h}}^{\mathcal{P}, \pi'}(s) < \phi$ have been “zeroed out” from $\mathbf{b}_{\underline{h}}(a_{1:\underline{h}-1}, o_{2:\underline{h}})$), and (41) uses the fact that

$$\begin{aligned} & \mathbb{E}_{s_{\underline{h}} \sim Q^\tau, (a_{\underline{h}:h-1}, o_{\underline{h}+1:h}) \sim \pi^\tau}^{\mathcal{P}} \left\| \mathbf{b}_h(a_{1:h-1}, o_{2:h}) - \mathbf{b}_h^{\text{apx}}(a_{\underline{h}:h-1}, o_{\underline{h}+1:h}; Q^\tau) \right\|_1 \\ & \leq 2 \cdot \min \left\{ \mathbb{E}_{s_{\underline{h}} \sim Q^\tau, (a_{\underline{h}:h-1}, o_{\underline{h}+1:h}) \sim \pi^\tau}^{\mathcal{P}} \left(\left\| \frac{\mathbf{b}_h^{\text{apx}}(a_{\underline{h}:h-1}, o_{\underline{h}+1:h}; Q^\tau)}{\mathbf{b}_h(a_{1:h-1}, o_{2:h})} \right\|_\infty - 1 \right), 1 \right\} \quad (43) \end{aligned}$$

$$\begin{aligned} & \leq 2 \cdot \min \left\{ \left(\left\| \frac{Q^\tau}{\mathbf{b}_h(a_{1:\underline{h}-1}, o_{2:\underline{h}})} \right\|_\infty - 1 \right), 1 \right\} \quad (44) \\ & \leq 2 \cdot \min \left\{ \frac{\epsilon^\tau}{1 - \epsilon^\tau}, 1 \right\} \leq 4\epsilon^\tau, \end{aligned}$$

where (43) uses the fact that for $P, Q \in \Delta^{\mathcal{S}}$, $\|P - Q\|_1 \leq 2 \cdot \left(\left\| \frac{P}{Q} \right\|_\infty - 1 \right)$, and (44) uses [GMR22, Lemma 4.6].

Note that $\mathbb{E}_{(a_{1:\underline{h}-1}, o_{2:\underline{h}}) \sim \pi}^{\mathcal{P}} [\mathbf{b}_{\underline{h}}(a_{1:\underline{h}-1}, o_{2:\underline{h}})] = d_{\mathcal{S}, \underline{h}}^{\mathcal{P}, \pi}$, meaning that $\mathbb{E}_\pi^{\mathcal{P}}[\epsilon^\tau] = d_{\mathcal{S}, \underline{h}}^{\mathcal{P}, \pi}(\mathcal{U}_{\phi, \underline{h}}^{\mathcal{P}}(\pi'))$. Furthermore, note that the conditional distribution of $s_{\underline{h}}$ given $(a_{1:\underline{h}-1}, o_{2:\underline{h}})$ is exactly $\mathbf{b}_{\underline{h}}(a_{1:\underline{h}-1}, o_{2:\underline{h}})$.

Thus, taking an expectation of the display (42) over $(a_{1:h-1}, o_{2:h}) \sim \pi$, we get that

$$\begin{aligned} & \mathbb{E}_{(a_{1:h-1}, o_{2:h}) \sim \pi}^{\mathcal{P}} \left\| \mathbf{b}_h(a_{1:h-1}, o_{2:h}) - \tilde{\mathbf{b}}_h^{\pi'}(a_{h:h-1}, o_{h+1:h}) \right\|_1 \\ & \leq \epsilon + \mathbb{E}_{\pi}^{\mathcal{P}} [6\epsilon^{\tau}] = \epsilon + 6 \cdot d_{\mathcal{S}, \underline{h}}^{\mathcal{P}, \pi}(\mathcal{U}_{\phi, \underline{h}}^{\mathcal{P}}(\pi')), \end{aligned}$$

as desired. \square

Lemma H.3 is the main result of this section, which uses Lemma H.2 to show an upper bound on the absolute value of the difference between the value functions of \mathcal{P} and $\widehat{\mathcal{M}}(\pi^{1:H})$, for any policy π .

Lemma H.3. *Consider any sequence of general policies $\pi^1, \dots, \pi^H \in \Pi^{\text{gen}}$ and $\phi \geq 0$. Fix any $H' \in [H]$. Consider any collection of reward functions $R = (R_1, \dots, R_H)$, $R_h : \mathcal{O} \rightarrow [-1, 1]$ so that $R_{h'} \equiv 0$ for all $h' > H'$. Suppose that in the construction of $\widehat{\mathcal{M}}(\pi^{1:H})$, the event \mathcal{E}^{low} of Lemma E.1 holds. Then for all general policies $\pi \in \Pi^{\text{gen}}$,*

$$\left| V_1^{\pi, \mathcal{P}, R}(\emptyset) - V_1^{\pi, \widehat{\mathcal{M}}(\pi^{1:H}), R}(\emptyset) \right| \leq H^2 \epsilon + H^2 \theta + \frac{H^2 A^{2L} O^L \zeta}{\phi} + 4H \cdot \sum_{h=L+1}^{H'-1} d_{\mathcal{S}, h-L}^{\mathcal{P}, \pi}(\mathcal{U}_{\phi, h-L}^{\mathcal{P}}(\pi^h)).$$

Proof. Fix general policies $\pi^{1:H}$, and for $h \in [H]$, let $\hat{\pi}^h$ be the policy which follows π^h up to step $\max\{h-L-1, 0\}$ and thereafter chooses actions uniformly at random. Write $\widehat{\mathcal{M}} = \widehat{\mathcal{M}}(\pi^{1:H})$ and $\widetilde{\mathcal{M}} = \widetilde{\mathcal{M}}(\pi^{1:H})$. By Jensen's inequality, it suffices to consider the case that π is a deterministic policy. Using the assumption that $R_{h'}$ is identically 0 for $h' > H'$, we have from Lemma B.1 that for all general policies π , writing $z_h = (a_{h-L:h-1}, o_{h-L+1:h})$,

$$\begin{aligned} \left| V_1^{\pi, \mathcal{P}, R}(\emptyset) - V_1^{\pi, \widehat{\mathcal{M}}, R}(\emptyset) \right| &= \left| \mathbb{E}_{(a_{1:H-1}, o_{2:H}) \sim \pi}^{\mathcal{P}} \left[\sum_{h=1}^H ((\mathbb{P}_h^{\mathcal{P}} - \mathbb{P}_h^{\widehat{\mathcal{M}}})(V_{h+1}^{\pi, \widehat{\mathcal{M}}, R} + R_{h+1}))(a_{1:h}, o_{2:h}) \right] \right| \\ &= \left| \mathbb{E}_{(a_{1:H-1}, o_{2:H}) \sim \pi}^{\mathcal{P}} \left[\sum_{h=1}^{H'-1} ((\mathbb{P}_h^{\mathcal{P}} - \mathbb{P}_h^{\widehat{\mathcal{M}}})(V_{h+1}^{\pi, \widehat{\mathcal{M}}, R} + R_{h+1}))(a_{1:h}, o_{2:h}) \right] \right| \\ &\leq \frac{H}{2} \cdot \mathbb{E}_{(a_{1:H-1}, o_{2:H}) \sim \pi}^{\mathcal{P}} \left[\sum_{h=1}^{H'-1} \left\| \mathbb{P}_h^{\mathcal{P}}(\cdot | a_{1:h}, o_{2:h}) - \mathbb{P}_h^{\widehat{\mathcal{M}}}(\cdot | z_h, a_h) \right\|_1 \right] \end{aligned} \quad (45)$$

$$\begin{aligned} &\leq \frac{H}{2} \cdot \mathbb{E}_{(a_{1:H-1}, o_{2:H}) \sim \pi}^{\mathcal{P}} \left[\sum_{h=1}^{H'-1} \left\| \mathbb{P}_h^{\mathcal{P}}(\cdot | a_{1:h}, o_{2:h}) - \mathbb{P}_h^{\widetilde{\mathcal{M}}}(\cdot | z_h, a_h) \right\|_1 \right] \\ &\quad + \frac{H}{2} \cdot \mathbb{E}_{(a_{1:H-1}, o_{2:H}) \sim \pi}^{\mathcal{P}} \left[\sum_{h=1}^{H'-1} \left\| \mathbb{P}_h^{\widetilde{\mathcal{M}}}(\cdot | z_h, a_h) - \mathbb{P}_h^{\widehat{\mathcal{M}}}(\cdot | z_h, a_h) \right\|_1 \right], \end{aligned} \quad (46)$$

where (45) uses the fact that $\left| \left(V_{h+1}^{\pi, \widehat{\mathcal{M}}, R} + R_{h+1} \right) (a_{1:h}, o_{2:h+1}) \right| \leq H$ for all $a_{1:h}, o_{2:h+1}$. For all $h \in [H]$, $a_{1:h} \in \mathcal{A}^h$ and $o_{2:h} \in \mathcal{O}^{h-1}$, we have (with $z_h = (a_{h-L:h-1}, o_{h-L+1:h})$)

$$\begin{aligned} \mathbb{P}_h^{\mathcal{P}}(o_{h+1} | a_{1:h}, o_{2:h}) &= e_{o_{h+1}}^{\top} \cdot \mathbb{O}_{h+1}^{\mathcal{P}} \cdot \mathbb{T}_h^{\mathcal{P}}(a_h) \cdot \mathbf{b}_h^{\mathcal{P}}(a_{1:h-1}, o_{2:h}) \\ \mathbb{P}_h^{\widetilde{\mathcal{M}}}(o_{h+1} | z_h, a_h) &= e_{o_{h+1}}^{\top} \cdot \mathbb{O}_{h+1}^{\mathcal{P}} \cdot \mathbb{T}_h^{\mathcal{P}}(a_h) \cdot \tilde{\mathbf{b}}_h^{\pi^h}(a_{h-L:h-1}, o_{h-L+1:h}), \end{aligned} \quad (47)$$

where (47) follows from (10). Thus, by the data processing inequality, for all $a_{1:h}, o_{2:h}$, (again writing $z_h = (a_{h-L:h-1}, o_{h-L+1:h})$)

$$\left\| \mathbb{P}_h^{\mathcal{P}}(\cdot | a_{1:h}, o_{2:h}) - \mathbb{P}_h^{\widetilde{\mathcal{M}}}(\cdot | z_h, a_h) \right\|_1 \leq \left\| \mathbf{b}_h^{\mathcal{P}}(a_{1:h-1}, o_{2:h}) - \tilde{\mathbf{b}}_h^{\pi^h}(a_{h-L:h-1}, o_{h-L+1:h}) \right\|_1. \quad (48)$$

Next, by Lemma E.1 and the assumption that \mathcal{E}^{low} holds, we have that for all $h \in [H]$, $z_h \in \mathcal{Z}$, $a_h \in \mathcal{A}$,

$$\left\| \mathbb{P}_h^{\widehat{\mathcal{M}}}(\cdot | z_h, a_h) - \mathbb{P}_h^{\widetilde{\mathcal{M}}}(\cdot | z_h, a_h) \right\|_1 \leq \theta + 2\mathbb{1}[z_h \in \mathcal{Z}_{h,\zeta}^{\text{low}}(\widehat{\pi}^h)]. \quad (49)$$

Moreover, by (48), we have

$$\begin{aligned} & \mathbb{E}_{(a_{1:H-1}, o_{2:H}) \sim \pi} \left[\sum_{h=1}^{H'-1} \left\| \mathbb{P}_h^{\mathcal{P}}(\cdot | a_{1:h}, o_{2:h}) - \mathbb{P}_h^{\widetilde{\mathcal{M}}}(\cdot | a_{1:h}, o_{2:h}) \right\|_1 \right] \\ & \leq \mathbb{E}_{(a_{1:H-1}, o_{2:H}) \sim \pi} \left[\sum_{h=1}^{H'-1} \left\| \mathbf{b}_h^{\mathcal{P}}(a_{1:h-1}, o_{2:h}) - \widetilde{\mathbf{b}}_h^{\pi^h}(a_{h-L:h-1}, o_{h-L+1:h}) \right\|_1 \right] \\ & \leq H\epsilon + 6 \cdot \sum_{h=L+1}^{H'-1} d_{\mathcal{S}, h-L}^{\mathcal{P}, \pi}(\mathcal{U}_{\phi, h-L}^{\mathcal{P}}(\pi^h)), \end{aligned} \quad (50)$$

where (50) uses Lemma H.2 for each h satisfying $L+1 \leq h \leq H'-1$ (in particular, for each such value of h , the policy π' in Lemma H.2 is set to π^h).

Next, using (49), we compute

$$\begin{aligned} & \mathbb{E}_{(a_{1:H-1}, o_{2:H}) \sim \pi} \left[\sum_{h=1}^{H'-1} \left\| \mathbb{P}_h^{\widehat{\mathcal{M}}}(\cdot | a_{1:h}, o_{2:h}) - \mathbb{P}_h^{\widetilde{\mathcal{M}}}(\cdot | a_{1:h}, o_{2:h}) \right\|_1 \right] \\ & \leq \theta H + 2 \cdot \sum_{h=1}^{H'-1} \mathbb{P}_{z_h \sim \pi}^{\mathcal{P}}(z_h \in \mathcal{Z}_{h,\zeta}^{\text{low}}(\widehat{\pi}^h)) \\ & \leq \theta H + 2 \cdot \frac{HA^2LO^L\zeta}{\phi} + 2 \cdot \sum_{h=L+1}^{H'-1} d_{\mathcal{S}, h-L}^{\mathcal{P}, \pi}(\mathcal{U}_{\phi, h-L}^{\mathcal{P}}(\pi^h)), \end{aligned} \quad (51)$$

where (51) uses Lemma F.4 for each $1 \leq h \leq H'-1$: in particular, for each value of h , the lemma is applied with $\overline{\mathcal{P}} = \mathcal{P}$ and $\pi' = \widehat{\pi}^h$.

Combining (46), (50), and (51), we obtain that

$$\left| V_1^{\pi, \mathcal{P}, R}(\emptyset) - V_1^{\pi, \widehat{\mathcal{M}}, R}(\emptyset) \right| \leq H^2\epsilon + H^2\theta + \frac{H^2A^2LO^L\zeta}{\phi} + 4H \cdot \sum_{h=L+1}^{H'-1} d_{\mathcal{S}, h-L}^{\mathcal{P}, \pi}(\mathcal{U}_{\phi, h-L}^{\mathcal{P}}(\pi^h)),$$

as desired. \square

I Analysis of BaSeCAMP (Algorithm 3)

In this section we prove Theorem 3.1. We first show the following technical lemma, which relates the underexplored sets (Definition E.2) of the truncated POMDPs $\overline{\mathcal{P}}_{\phi, h}(\pi^{1:H})$ and \mathcal{P} .

Lemma I.1. *Consider any sequence of general policies $\pi^{1:H,0}$, and define $\pi^h := \frac{1}{H-h+1} \sum_{h'=h}^H \pi^{h',0}$ for each $h \in [H]$. Then, for each h satisfying $L < h \leq H$ and any $\phi > 0$,*

$$\mathcal{U}_{\phi, h-L}^{\overline{\mathcal{P}}_{\phi, h}(\pi^{1:H})}(\pi^h) \subset \mathcal{U}_{\phi, H^2S, h-L}^{\mathcal{P}}(\pi^h).$$

Proof. Fix ϕ and $\pi^{1:H}$, and write $\overline{\mathcal{P}}_h := \overline{\mathcal{P}}_{\phi, h}(\pi^{1:H})$.

Consider any $h > L$, and any $s \in \mathcal{U}_{\phi, h-L}^{\overline{\mathcal{P}}_h}(\pi^h)$; in particular, we have $d_{\mathcal{S}, h-L}^{\overline{\mathcal{P}}_h, \pi^h}(s) \leq \phi$. By item 1 of Lemma F.2, it holds that

$$d_{\mathcal{S}, h-L}^{\mathcal{P}, \pi^h}(s) \leq d_{\mathcal{S}, h-L}^{\overline{\mathcal{P}}_h, \pi^h}(s) + d_{\mathcal{S}, h-L}^{\overline{\mathcal{P}}_h, \pi^h}(s^{\text{sink}}). \quad (52)$$

For each $H' > L$, write $\mathcal{U}_{H'-L} := \mathcal{U}_{\phi, H'-L}^{\bar{\mathcal{P}}_{H'-1}(\pi^{H'})}$; in words, $\mathcal{U}_{H'-L}$ is the set of underexplored states in $\bar{\mathcal{P}}_{H'-1}$ at step $H' - L$ which all mass is diverted away from in the construction of $\bar{\mathcal{P}}_{H'}$ (see Section E.3). It holds that

$$\begin{aligned} d_{\mathcal{S}, h-L}^{\bar{\mathcal{P}}_h, \pi^h}(s^{\text{sink}}) &= \sum_{H'=L+1}^h d_{\mathcal{S}, H'-L}^{\bar{\mathcal{P}}_{H'-1}, \pi^{H'}}(\mathcal{U}_{H'-L}) \\ &\leq H \cdot \sum_{H'=L+1}^h d_{\mathcal{S}, H'-L}^{\bar{\mathcal{P}}_{H'-1}, \pi^{H'}}(\mathcal{U}_{H'-L}) \end{aligned} \quad (53)$$

$$\leq H \cdot \sum_{H'=L+1}^h S\phi \leq H(H-1)S\phi, \quad (54)$$

where (53) follows from the fact that for all $s \in \mathcal{S}$ and any $H' \leq h$,

$$\begin{aligned} d_{\mathcal{S}, H'-L}^{\bar{\mathcal{P}}_{H'-1}, \pi^h}(s) &= \frac{1}{H-h+1} \sum_{h'=h}^H d_{\mathcal{S}, H'-L}^{\bar{\mathcal{P}}_{H'-1}, \pi^{h',0}}(s) \\ &\leq H \cdot \frac{1}{H-H'+1} \sum_{h'=H'}^H d_{\mathcal{S}, H'-L}^{\bar{\mathcal{P}}_{H'-1}, \pi^{h',0}}(s) \\ &= H \cdot d_{\mathcal{S}, H'-L}^{\bar{\mathcal{P}}_{H'-1}, \pi^{H'}}(s), \end{aligned}$$

and (54) follows by definition of $\mathcal{U}_{H'-L}$. From (52) it follows that $d_{\mathcal{S}, h-L}^{\bar{\mathcal{P}}_h, \pi^h}(s) \leq \phi + H(H-1)S\phi \leq H^2S\phi$, i.e., $s \in \mathcal{U}_{\phi, H^2S, h-L}^{\bar{\mathcal{P}}_h, \pi^h}$, as desired. \square

Lemma I.2 is the main technical lemma in the proof of Theorem 3.1, and shows that for each iteration k of BaSeCAMP, one of two options must hold: option 1 implies that the policies $\bar{\pi}^{k,1:H}$ of BaSeCAMP have sufficiently explored \mathcal{P} and will be used later to show that the optimal policy for $\widehat{\mathcal{M}}^{(k)}$ is an approximately optimal policy for \mathcal{P} . Option 2 states that progress is made by BaSeCAMP at iteration k in sense that some new latent state (s, h) of \mathcal{P} must be explored. The proof of the lemma proceeds by combining the results of Section G and H which show that the true POMDP \mathcal{P} is “sufficiently close” to the approximating MDPs $\widehat{\mathcal{M}}^{(k)}$, with the result of Section D which states that the policy obtained from a barycentric spanner (as in Algorithm 2) has desirable exploration properties.

Lemma I.2 (“Progress Lemma”; formal version of Lemma 5.1). *Consider the execution of Algorithm 3 given the values of the parameters defined in Section C.1. Fix any $k \in [K]$ in step 3 of Algorithm 3 and suppose that, in the construction of $\widehat{\mathcal{M}}^{(k)}$ in step 5, the event \mathcal{E}^{low} of Lemma E.1 holds. Then one of the two following statements holds:*

1. For all general policies π , it holds that $d_{\mathcal{S}, H-L}^{\bar{\mathcal{P}}_H(\bar{\pi}^{k,1:H}), \pi}(s^{\text{sink}}) \leq \delta H$.
2. There is $H' \in [H]$ and $(h, s) \in [H'] \times \mathcal{S}$, so that $h > L$ and the following holds:

$$(h-L, s) \in \mathcal{U}_{\phi, H^2S, h-L}^{\bar{\mathcal{P}}_h, \pi^h}(\bar{\pi}^{k,h}) \quad \text{and} \quad d_{\mathcal{S}, h-L}^{\bar{\mathcal{P}}_h, \pi^{k+1, H', 0}}(s) \geq \frac{\delta' \cdot \gamma}{56HS^{3/2}O^2}.$$

Proof. Fix $k \in [K]$. For $h \in [H]$, we will write $\pi^h := \bar{\pi}^{k,h}$ for the course of the proof. Furthermore, according to our convention, let $\hat{\pi}^h$ denote the policy which follows π^h for the first $\max\{h-L-1, 0\}$ steps and thereafter chooses uniformly random actions. For each $H' \in [H]$, define $\bar{\mathcal{P}}_{H'} := \bar{\mathcal{P}}_{\phi, H'}(\pi^{1:H})$ and $\widehat{\mathcal{M}} = \widehat{\mathcal{M}}(\pi^{1:H}) = \widehat{\mathcal{M}}^{(k)}$.

We consider two cases:

Case 1. First, suppose that for all general policies $\pi \in \Pi^{\text{gen}}$ and all $H' \in [H]$ satisfying $H' > L$, it holds that

$$d_{\mathcal{S}, H'-L}^{\bar{\mathcal{P}}_{H'-1}, \pi} \left(\mathcal{U}_{\phi, H'-L}^{\bar{\mathcal{P}}_{H'-1}, \pi}(\pi^{H'}) \right) \leq \delta. \quad (55)$$

We now claim that

$$d_{\mathcal{S}, H-L}^{\overline{\mathcal{P}}_{H,\pi}}(s^{\text{sink}}) \leq \delta \cdot H. \quad (56)$$

To see that (56) holds, we argue by induction, noting that $d_{\mathcal{S}, L}^{\overline{\mathcal{P}}_{L,\pi}}(s^{\text{sink}}) = 0$ since $\overline{\mathcal{P}}_L = \mathcal{P}$. Next, we note that for each $H' \geq L + 1$,

$$\begin{aligned} d_{\mathcal{S}, H'-L}^{\overline{\mathcal{P}}_{H',\pi}}(s^{\text{sink}}) &= d_{\mathcal{S}, H'-L-1}^{\overline{\mathcal{P}}_{H',\pi}}(s^{\text{sink}}) + d_{\mathcal{S}, H'-L}^{\overline{\mathcal{P}}_{H'-1,\pi}} \left(\mathcal{U}_{\phi, H'-L}^{\overline{\mathcal{P}}_{H'-1}}(\pi^{H'}) \right) \\ &= d_{\mathcal{S}, H'-L-1}^{\overline{\mathcal{P}}_{H'-1,\pi}}(s^{\text{sink}}) + d_{\mathcal{S}, H'-L}^{\overline{\mathcal{P}}_{H'-1,\pi}} \left(\mathcal{U}_{\phi, H'-L}^{\overline{\mathcal{P}}_{H'-1}}(\pi^{H'}) \right) \\ &\leq \delta \cdot H' + \delta = \delta \cdot (H' + 1), \end{aligned}$$

where the first equality holds by definition of $\overline{\mathcal{P}}_{H'+1}$ (see Section E.3), the second equality holds since the transitions at steps $1, 2, \dots, H' - L - 2$ of $\overline{\mathcal{P}}_{H'-1}, \overline{\mathcal{P}}_{H'}$ are identical, and the inequality follows by assumption (55) and the inductive hypothesis. Thus, in this case, we have established that (56) holds (i.e., item 1 in the lemma statement holds).

Case 2. Next suppose that the previous case does not hold, and choose $H' > L$ as small as possible so that, for some $\pi^* \in \Pi^{\text{gen}}$,

$$d_{\mathcal{S}, H'-L}^{\overline{\mathcal{P}}_{H'-1,\pi^*}} \left(\mathcal{U}_{\phi, H'-L}^{\overline{\mathcal{P}}_{H'-1}}(\pi^{H'}) \right) > \delta. \quad (57)$$

We will now apply Lemma G.4 with the value of H' in Lemma G.4 set to $H' - L$, the value of H'' in Lemma G.4 set to $H' - 1$, the value of ϕ in Lemma G.4 set to the value of ϕ defined in Section C.1, $\pi^{1:H}$ set to the present value of $\pi^{1:H}$, and $\widehat{\mathcal{M}}(\pi^{1:H})$ set to $\widehat{\mathcal{M}} = \widehat{\mathcal{M}}^{(k)}$. We also need to define the rewards $R = (R_1, \dots, R_H)$, where each $R_h : \overline{\mathcal{O}} \rightarrow \mathbb{R}$. We set each R_h , $h \neq H' - L$, to be identically 0. Furthermore, for $o \in \overline{\mathcal{O}}$, we define

$$R_{H'-L}(o) := \sum_{s \in \mathcal{S}} \mathbb{1} \left[s \in \mathcal{U}_{\phi, H'-L}^{\overline{\mathcal{P}}_{H'-1}}(\pi^{H'}) \right] \cdot (\mathbb{O}_{H'-L}^{\mathcal{P}})_{s,o}^\dagger.$$

Note that the rewards R satisfy the requirements of Lemma G.4. The definition of R ensures that, for all $\pi \in \Pi^{\text{gen}}$,

$$\begin{aligned} \sum_{s \in \mathcal{U}_{\phi, H'-L}^{\overline{\mathcal{P}}_{H'-1}}(\pi^{H'})} \left\langle e_s, (\mathbb{O}_{H'-L}^{\mathcal{P}})^\dagger \cdot d_{\mathcal{O}, H'-L}^{\widehat{\mathcal{M}}, \pi} \right\rangle &= V_1^{\pi, \widehat{\mathcal{M}}, R}(\emptyset) \\ d_{\mathcal{S}, H'-L}^{\overline{\mathcal{P}}_{H'-1,\pi}} \left(\mathcal{U}_{\phi, H'-L}^{\overline{\mathcal{P}}_{H'-1}}(\pi^{H'}) \right) &= \sum_{s \in \mathcal{U}_{\phi, H'-L}^{\overline{\mathcal{P}}_{H'-1}}(\pi^{H'})} \left\langle e_s, (\mathbb{O}_{H'-L}^{\mathcal{P}})^\dagger \cdot d_{\mathcal{O}, H'-L}^{\overline{\mathcal{P}}_{H'-1,\pi}} \right\rangle = V_1^{\pi, \overline{\mathcal{P}}_{H'-1}, R}(\emptyset). \end{aligned}$$

Thus, by Lemma G.4, the fact that \mathcal{E}^{low} holds for the construction of $\widehat{\mathcal{M}}$, and (57), it follows that

$$\begin{aligned} &\sum_{s \in \mathcal{U}_{\phi, H'-L}^{\overline{\mathcal{P}}_{H'-1}}(\pi^{H'})} \left\langle e_s, (\mathbb{O}_{H'-L}^{\mathcal{P}})^\dagger \cdot d_{\mathcal{O}, H'-L}^{\widehat{\mathcal{M}}, \pi^*} \right\rangle \\ &= V_1^{\pi^*, \widehat{\mathcal{M}}, R}(\emptyset) \\ &\geq V_1^{\pi^*, \overline{\mathcal{P}}_{H'-1}, R}(\emptyset) - \left(\frac{H^2 \sqrt{S} \epsilon}{\gamma} + \frac{2H^3 \sqrt{S} \theta}{\gamma} + \frac{H^2 A^{2L} O^L \sqrt{S} \zeta}{\phi \gamma} \right) \\ &\geq \delta - \left(\frac{H^2 \sqrt{S} \epsilon}{\gamma} + \frac{2H^3 \sqrt{S} \theta}{\gamma} + \frac{H^2 A^{2L} O^L \sqrt{S} \zeta}{\phi \gamma} \right) \geq \delta', \end{aligned}$$

where we have used the definitions of the parameters given in Section C.1: in particular, since $\theta = \epsilon$ and $\frac{\zeta \cdot A^{2L} O^L}{\phi} \leq \epsilon$, we need that $\delta - \frac{H^3 \sqrt{S}}{\gamma} \cdot 4\epsilon \geq \delta'$, which holds since $\delta = 2\delta' = C^* \cdot \frac{O^4 H^3 \sqrt{S}}{\gamma} \cdot \epsilon$ for some constant $C^* > 1$ and C^* is sufficiently large.

Since $\widehat{\mathcal{M}}$ is an MDP on the state space $\overline{\mathcal{Z}}$, and any MDP has a deterministic Markov policy achieving its optimal value, there must be some $\sigma^* \in \Pi_{\mathcal{Z}}^{\text{markov}}$ so that

$$\sum_{s \in \mathcal{U}_{\phi, H'-L}^{\overline{\mathcal{P}}_{H'-1}}(\pi^{H'})} \left\langle e_s, (\mathbb{O}_{H'-L}^{\mathcal{P}})^{\dagger} \cdot d_{\mathcal{O}, H'-L}^{\widehat{\mathcal{M}}, \sigma^*} \right\rangle = V_1^{\sigma^*, \widehat{\mathcal{M}}, R}(\emptyset) \geq \delta'. \quad (58)$$

We will now apply Lemma D.2 with the following settings (importantly, the values of \mathcal{X}, \mathbb{M} below are the same as the settings used in step 4 of Algorithm 2 in the construction of the policies $\pi^{k+1, H', 0}$):

$$\eta = \delta', \quad \mathcal{X} = \left\{ d_{\mathcal{O}, H'-L}^{\widehat{\mathcal{M}}, \pi} : \pi \in \Pi_{\mathcal{Z}}^{\text{markov}} \right\}, \quad \mathbb{M} = (\mathbb{O}_{H'-L}^{\mathcal{P}})^{\dagger}.$$

We must first verify that the preconditions of the lemma hold: in particular, it suffices to show that for any subset $\mathcal{S}' \subset \mathcal{S}$, and any $\pi \in \Pi_{\mathcal{Z}}^{\text{markov}} \subset \Pi^{\text{gen}}$,

$$\left\langle \sum_{s \in \mathcal{S}'} e_s, (\mathbb{O}_{H'-L}^{\mathcal{P}})^{\dagger} \cdot d_{\mathcal{O}, H'-L}^{\widehat{\mathcal{M}}, \pi} \right\rangle \geq -\frac{\delta'}{4O^2}. \quad (59)$$

To do so, we will again apply Lemma G.4, with the same parameter settings as previously, except with different reward functions: in particular, we define the rewards $R = (R_1, \dots, R_H)$, for $R_h : \overline{\mathcal{O}} \rightarrow \mathbb{R}$, as follows. We set each R_h , $h \neq H' - L$, to be identically 0. Furthermore, for $o \in \overline{\mathcal{O}}$, we define

$$R_{H'-L}(o) := \sum_{s \in \mathcal{S}} \mathbb{1}[s \in \mathcal{S}'] \cdot (\mathbb{O}_{H'-L}^{\mathcal{P}})^{\dagger}_{s,o}.$$

Then the following equalities hold true:

$$\begin{aligned} V_1^{\pi, \widehat{\mathcal{M}}, R}(\emptyset) &= \left\langle \sum_{s \in \mathcal{S}'} e_s, (\mathbb{O}_{H'-L}^{\mathcal{P}})^{\dagger} \cdot d_{\mathcal{O}, H'-L}^{\widehat{\mathcal{M}}, \pi} \right\rangle \\ V_1^{\pi, \overline{\mathcal{P}}_{H'-1}, R}(\emptyset) &= \left\langle \sum_{s \in \mathcal{S}'} e_s, (\mathbb{O}_{H'-L}^{\mathcal{P}})^{\dagger} \cdot d_{\mathcal{O}, H'-L}^{\overline{\mathcal{P}}_{H'-1}, \pi} \right\rangle. \end{aligned} \quad (60)$$

Since, for each $s \in \mathcal{S}' \subset \mathcal{S}$, it holds that

$$\left\langle e_s, (\mathbb{O}_{H'-L}^{\mathcal{P}})^{\dagger} \cdot d_{\mathcal{O}, H'-L}^{\overline{\mathcal{P}}_{H'-1}, \pi} \right\rangle = \left\langle e_s, d_{\mathcal{S}, H'-L}^{\overline{\mathcal{P}}_{H'-1}, \pi} \right\rangle = d_{\mathcal{S}, H'-L}^{\overline{\mathcal{P}}_{H'-1}, \pi}(s) \geq 0,$$

we have that $V_1^{\pi, \overline{\mathcal{P}}_{H'-1}, R}(\emptyset) \geq 0$, which means, by Lemma G.4 and (60), that (59) holds as long as

$$\frac{\delta'}{4O^2} \geq \frac{H^2 \sqrt{S} \epsilon}{\gamma} + \frac{2H^3 \sqrt{S} \theta}{\gamma} + \frac{H^2 A^{2L} O^L \sqrt{S} \zeta}{\phi \gamma},$$

which is true by our parameter settings in Section C.1 (in particular, since $\theta = \epsilon$ and $\frac{\zeta A^{2L} O^L}{\phi} = \epsilon$, we need that $\frac{\delta'}{4O^2} \geq \frac{H^3 \sqrt{S}}{\gamma} \cdot 4\epsilon$, which holds as long as C^* is sufficiently large).

By Lemma D.2 and (58), it holds that the policy $\pi^{k+1, H', 0}$ output by Algorithm 2 satisfies

$$\sum_{s \in \mathcal{U}_{\phi, H'-L}^{\overline{\mathcal{P}}_{H'-1}}(\pi^{H'})} \left\langle e_s, (\mathbb{O}_{H'-L}^{\mathcal{P}})^{\dagger} \cdot d_{\mathcal{O}, H'-L}^{\widehat{\mathcal{M}}, \pi^{k+1, H', 0}} \right\rangle \geq \frac{\delta'}{4O^2}. \quad (61)$$

We next apply Lemma H.3, with the present values of H' , $\pi^{1:H}$ and ϕ (namely, the ϕ defined in Section C.1), and the following setting of $R = (R_1, \dots, R_H)$: $R_{h'} \equiv 0$ for all $h' \neq H' - L$, and, for $o \in \overline{\mathcal{O}}$,

$$R_{H'-L}(o) := \sum_{s \in \mathcal{U}_{\phi, H'-L}^{\overline{\mathcal{P}}_{H'-1}}(\pi^{H'})} (\mathbb{O}_{H'-L}^{\mathcal{P}})^{\dagger}_{s,o}.$$

Since $|R_{H'-L}(o)| \leq \frac{\sqrt{S}}{\gamma}$ for all $o \in \overline{\mathcal{O}}$ (which follows from Lemma B.3), we have from Lemma H.3 (and the fact that \mathcal{E}^{low} holds in the construction of $\widehat{\mathcal{M}}$) that

$$\begin{aligned} & \left| d_{\mathbb{S}, H'-L}^{\mathcal{P}, \pi^{k+1, H', 0}}(\mathcal{U}_{\phi, H'-L}^{\overline{\mathcal{P}}_{H'-1}}(\pi^{H'})) - \sum_{s \in \mathcal{U}_{\phi, H'-L}^{\overline{\mathcal{P}}_{H'-1}}(\pi^{H'})} \left\langle e_s, (\mathbb{O}_{H'-L}^{\mathcal{P}})^\dagger \cdot d_{\mathbb{O}, H'-L}^{\widehat{\mathcal{M}}, \pi^{k+1, H', 0}} \right\rangle \right| \\ &= \left| V_1^{\pi^{k+1, H', 0}, \mathcal{P}, R}(\emptyset) - V_1^{\pi^{k+1, H', 0}, \widehat{\mathcal{M}}, R}(\emptyset) \right| \\ &\leq \frac{\sqrt{S}}{\gamma} \cdot \left(H^2 \epsilon + H^2 \theta + \frac{H^2 A^{2L} O^L \zeta}{\phi} + 4H \cdot \sum_{h=L+1}^{H'-1} d_{\mathbb{S}, h-L}^{\pi^{k+1, H', 0}, \mathcal{P}}(\mathcal{U}_{\phi, h-L}^{\mathcal{P}}(\pi^h)) \right) \\ &\leq \frac{\delta'}{8O^2} + \frac{4H\sqrt{S}}{\gamma} \cdot \sum_{h=L+1}^{H'-1} d_{\mathbb{S}, h-L}^{\mathcal{P}, \pi^{k+1, H', 0}}(\mathcal{U}_{\phi, h-L}^{\mathcal{P}}(\pi^h)), \end{aligned}$$

where the final inequality follows since

$$\frac{\sqrt{S}}{\gamma} \cdot \left(H^2 \epsilon + H^2 \theta + \frac{H^2 A^{2L} O^L \zeta}{\phi} \right) \leq \frac{\delta'}{8O^2},$$

which holds by our parameter settings in Section C.1. In particular, since we have $\theta = \epsilon$ and $\frac{A^{2L} O^L \zeta}{\phi} \leq \epsilon$, it suffices to have $\frac{\delta'}{8O^2} \geq \frac{\sqrt{S} H^2}{\gamma} \cdot 3\epsilon$, which holds as long as the constant C^* is sufficiently large. Thus, using (61), we get that

$$d_{\mathbb{S}, H'-L}^{\mathcal{P}, \pi^{k+1, H', 0}}(\mathcal{U}_{\phi, H'-L}^{\overline{\mathcal{P}}_{H'-1}}(\pi^{H'})) \geq \frac{\delta'}{8O^2} - \frac{4H\sqrt{S}}{\gamma} \cdot \sum_{h=L+1}^{H'-1} d_{\mathbb{S}, h-L}^{\mathcal{P}, \pi^{k+1, H', 0}}(\mathcal{U}_{\phi, h-L}^{\mathcal{P}}(\pi^h)). \quad (62)$$

We next apply Lemma I.1, with $\pi^{h,0}$ in the lemma statement set to $\overline{\pi}^{k, h, 0} := \frac{1}{k} \sum_{k'=1}^k \pi^{k', h, 0}$. Then we have

$$\pi^h = \overline{\pi}^{k, h} = \frac{1}{k} \sum_{k'=1}^k \frac{1}{H-h+1} \sum_{h'=h}^H \pi^{k', h', 0} = \frac{1}{H-h+1} \sum_{h'=h}^H \overline{\pi}^{k, h', 0},$$

so Lemma I.1 gives that $\mathcal{U}_{\phi, H'-L}^{\overline{\mathcal{P}}_{H'-1}}(\pi^{H'}) \subset \mathcal{U}_{\phi, H^2 S, H'-L}^{\mathcal{P}}(\pi^{H'})$, so rearranging (62), we get

$$d_{\mathbb{S}, H'-L}^{\mathcal{P}, \pi^{k+1, H', 0}}(\mathcal{U}_{\phi, H^2 S, H'-L}^{\mathcal{P}}(\pi^{H'})) + \frac{4H\sqrt{S}}{\gamma} \cdot \sum_{h=L+1}^{H'-1} d_{\mathbb{S}, h-L}^{\mathcal{P}, \pi^{k+1, H', 0}}(\mathcal{U}_{\phi, h-L}^{\mathcal{P}}(\pi^h)) \geq \frac{\delta'}{8O^2}.$$

Using that $\mathcal{U}_{\phi, h-L}^{\mathcal{P}}(\pi^h) \subset \mathcal{U}_{\phi, H^2 S, h-L}^{\mathcal{P}}(\pi^h)$ for all h and simplifying gives

$$\frac{4H\sqrt{S}}{\gamma} \cdot \sum_{h=L+1}^{H'} d_{\mathbb{S}, h-L}^{\mathcal{P}, \pi^{k+1, H', 0}}(\mathcal{U}_{\phi, H^2 S, h-L}^{\mathcal{P}}(\pi^h)) \geq \frac{\delta'}{8O^2}.$$

It then follows that there is some $(h, s) \in [H'] \times \mathcal{S}$ so that $h > L$, $(h-L, s) \in \mathcal{U}_{\phi, H^2 S}^{\mathcal{P}}(\pi^h)$ and

$$d_{\mathbb{S}, h-L}^{\mathcal{P}, \pi^{k+1, H', 0}}(s) \geq \frac{\delta' \cdot \gamma}{32H^2 S^{3/2} O^2},$$

which means that item (2) of the lemma statement holds. \square

Recall that in Section C.1 we set $K = 2HS$. Lemma I.3 formalizes the intuition that item 2 of Lemma I.2 can only hold for HS iterations, showing that indeed item 1 must hold for some iteration $k \in [K]$ (using that $K > HS$).

Lemma I.3. *Suppose that the event \mathcal{E}^{low} (of Lemma E.1) holds for the construction of $\widehat{\mathcal{M}}^{(k)}$, for all rounds k in Algorithm 3. Then at least $K/2$ iterations over k in Algorithm 3 satisfy the following: for all $\pi \in \Pi^{\text{gen}}$, $d_{\mathbb{S}, H-L}^{\overline{\mathcal{P}}_{\phi, H}(\overline{\pi}^{k, 1: H}), \pi}(s^{\text{sink}}) \leq \delta H$.*

Proof. Let the set of $k \in [K]$ so that there is some $\pi \in \Pi^{\text{gen}}$ with $d_{\mathcal{S}, H-L}^{\overline{\mathcal{P}}_{\phi, H}(\overline{\pi}^{k, 1:H}), \pi}(s^{\text{sink}}) > \delta H$ be denoted by $\mathcal{K}_0 \subset [K]$. Lemma I.2 gives that for each $k \in \mathcal{K}_0$, there is $H' \in [H]$, $(h, s) \in [H'] \times \mathcal{S}$ so that $h > L$ and $(h-L, s) \in \mathcal{U}_{\phi, H^2 S, h-L}^{\mathcal{P}}(\overline{\pi}^{k, h})$ so that $d_{\mathcal{S}, h-L}^{\mathcal{P}, \overline{\pi}^{k+1, H'}, 0}(s) \geq \frac{\delta' \cdot \gamma}{56 H S^{3/2} O^2}$. For all $k' \in [K]$ with $k' \geq k+1$, $\overline{\pi}^{k', h}$ is a mixture policy which puts weight at least $\frac{1}{K} \cdot \frac{1}{H} = \frac{1}{2H^2 S}$ on $\overline{\pi}^{k+1, H'}, 0$, meaning that for all such k' , it holds that

$$d_{\mathcal{S}, h-L}^{\mathcal{P}, \overline{\pi}^{k', h}}(s) \geq \frac{1}{2H^2 S} \cdot \frac{\delta' \cdot \gamma}{56 H S^{3/2} O^2} > \phi \cdot H^2 S,$$

which means that $(h-L, s) \notin \mathcal{U}_{\phi, H^2 S, h-L}^{\mathcal{P}}(\overline{\pi}^{k', h})$. The second inequality above holds since we have, in Section C.1, set $\phi = \frac{\epsilon \cdot \gamma}{C^* \cdot H^5 S^{7/2} O^2} < \frac{\delta' \cdot \gamma}{C^* \cdot H^5 S^{7/2} O^2}$ for a sufficiently large constant C^* .

Thus, if $|\mathcal{K}_0| > HS$, then letting k^* be the largest element of \mathcal{K}_0 , we get that for all $h \in [H]$ with $h > L$, $\mathcal{U}_{\phi, H^2 S, h-L}^{\mathcal{P}}(\overline{\pi}^{k^*, h})$ must be empty, which contradicts Lemma I.2. Hence $|\mathcal{K}_0| \leq HS = K/2$, as desired. \square

Lemma I.4 shows that, if the property in item 1 of Lemma I.2 is satisfied, then the value functions of \mathcal{P} and $\widehat{\mathcal{M}}(\pi^{1:H})$ are close under any policy. This result is similar to Lemma H.3 (which is used in its proof), but unlike Lemma H.3, the upper bound does not depend on the probability of reaching any underexplored set. Instead, such a probability is upper bounded by the assumption that item 1 of Lemma I.2 holds.

Lemma I.4. *Let $\phi > 0$ and suppose $\pi^{1:H}$ are general policies so that $\overline{\mathcal{P}}_{\phi, H}(\pi^{1:H})$ satisfies the following, for some $\delta > 0$: for all general policies π , $d_{\mathcal{S}, H-L}^{\overline{\mathcal{P}}_{\phi, H}(\pi^{1:H}), \pi}(s^{\text{sink}}) \leq \delta$. Suppose that, in the construction of $\widehat{\mathcal{M}}(\pi^{1:H})$, the event \mathcal{E}^{low} of Lemma E.1 holds. Then, for all $h \in [H]$ with $h > L$, we have*

$$d_{\mathcal{S}, h-L}^{\mathcal{P}, \pi}(\mathcal{U}_{\phi, h-L}^{\mathcal{P}}(\pi^h)) \leq S\delta. \quad (63)$$

Furthermore, for any reward functions $R = (R_1, \dots, R_H)$, $R_h : \overline{\mathcal{O}} \rightarrow [0, 1]$, it holds that

$$\left| V_1^{\pi, \mathcal{P}, R}(\emptyset) - V_1^{\pi, \widehat{\mathcal{M}}(\pi^{1:H}), R}(\emptyset) \right| \leq H^2 \epsilon + H^2 \theta + \frac{H^2 A^{2L} O^L \zeta}{\phi} + 4H^2 S \delta.$$

Proof. Since \mathcal{E}^{low} holds, Lemma H.3 with the given values of ϕ , $\pi^{1:H}$ and $H' = H$ gives that, for all general policies π ,

$$\left| V_1^{\pi, \mathcal{P}, R}(\emptyset) - V_1^{\pi, \widehat{\mathcal{M}}(\pi^{1:H}), R}(\emptyset) \right| \leq H^2 \epsilon + H^2 \theta + \frac{H^2 A^{2L} O^L \zeta}{\phi} + 4H \cdot \sum_{h=L+1}^H d_{\mathcal{S}, h-L}^{\mathcal{P}, \pi}(\mathcal{U}_{\phi, h-L}^{\mathcal{P}}(\pi^h)).$$

Thus it suffices to bound $d_{\mathcal{S}, h-L}^{\mathcal{P}, \pi}(\mathcal{U}_{\phi, h-L}^{\mathcal{P}}(\pi^h))$ for each $h > L$. For this purpose, we first use item 1 of Lemma F.2 to get that for all $s \in \mathcal{S}$ and $h > L$, and all $\pi \in \Pi^{\text{gen}}$,

$$d_{\mathcal{S}, h-L}^{\mathcal{P}, \pi}(s) \leq d_{\mathcal{S}, h-L}^{\overline{\mathcal{P}}_{\phi, h}(\pi^{1:H}), \pi}(s) + d_{\mathcal{S}, h-L}^{\overline{\mathcal{P}}_{\phi, h}(\pi^{1:H}), \pi}(s^{\text{sink}}). \quad (64)$$

By items 2 and 3 of Lemma F.2 and our assumption in the statement of the present lemma that $d_{\mathcal{S}, H-L}^{\overline{\mathcal{P}}_{\phi, H}(\pi^{1:H}), \pi}(s^{\text{sink}}) \leq \delta$, it holds that

$$d_{\mathcal{S}, h-L}^{\overline{\mathcal{P}}_{\phi, h}(\pi^{1:H}), \pi}(s^{\text{sink}}) \leq d_{\mathcal{S}, H-L}^{\overline{\mathcal{P}}_{\phi, H}(\pi^{1:H}), \pi}(s^{\text{sink}}) \leq \delta.$$

Lemma F.1 gives that for all $h \in [H]$, $\mathcal{U}_{\phi, h-L}^{\mathcal{P}}(\pi^h) \subset \mathcal{U}_{\phi, h-L}^{\overline{\mathcal{P}}_{\phi, h}(\pi^{1:H})}(\pi^h)$. But Lemma E.2 with $H' = h$ ensures that for any $s \in \mathcal{U}_{\phi, h-L}^{\overline{\mathcal{P}}_{\phi, h}(\pi^{1:H})}(\pi^h)$, we have $d_{\mathcal{S}, h-L}^{\overline{\mathcal{P}}_{\phi, h}(\pi^{1:H}), \pi}(s) = 0$. Thus, from (64) we have

$$d_{\mathcal{S}, h-L}^{\mathcal{P}, \pi}(\mathcal{U}_{\phi, h-L}^{\mathcal{P}}(\pi^h)) \leq S\delta,$$

which gives the desired result. \square

Finally, using Lemma I.4, the proof of Theorem 3.1 is straightforward.

Proof of Theorem 3.1. Let \mathcal{E}^1 be the event that the event \mathcal{E}^{low} of Lemma E.1 holds in the construction of $\widehat{\mathcal{M}}^{(k)}$ in Algorithm 3 for all $k \in [K]$. Then, by Lemma E.1, the probability of \mathcal{E}^1 is at least $1 - pK \geq 1 - \beta/2$ (where we have used $p = \beta/(2K)$ as defined in Section C.1).

By Lemma I.3, under the event \mathcal{E}^1 , there must be some iteration k in Algorithm 3 so that, for all $\pi \in \Pi^{\text{gen}}$, $d_{\mathcal{S}, H-L}^{\mathcal{P}, H}(\bar{\pi}^{k, 1:H}, \pi)(s^{\text{sink}}) \leq \delta H$. Then by Lemma I.4, for any reward function R with values in $[0, 1]$, it holds that, under \mathcal{E}^1 ,

$$\left| V_1^{\pi, \mathcal{P}, R}(\emptyset) - V_1^{\pi, \widehat{\mathcal{M}}^{(k)}, R}(\emptyset) \right| \leq H^2 \epsilon + H^2 \theta + \frac{H^2 A^{2L} O^L \zeta}{\phi} + 4H^3 S \delta \leq \frac{\alpha}{8}, \quad (65)$$

where the final inequality follows by definition of α in Section C.1, assuming the constant C^* is large enough.

For each $h \in [H]$, let $\pi^h := \bar{\pi}^{k, h}$, and $\hat{\pi}^h$ denote the policy which follows $\bar{\pi}^{k, h}$ up to step $\max\{h - L - 1, 0\}$ and thereafter chooses uniformly random actions. Denote the reward function of $\widehat{\mathcal{M}}^{(k)}$ as constructed in Algorithm 1 by $\widehat{R}^{(k)} := R^{\widehat{\mathcal{M}}^{(k)}}$, and denote the true reward function of \mathcal{P} by $R^{\mathcal{P}}$. By Lemma E.1, under the event \mathcal{E}^1 , for any $h \in [H]$ and any trajectory $o_{1:h-1}, o_{2:h}$ for which $(o_{h-L:h-1}, o_{h-L+1:h}) \notin \mathcal{Z}_{h, \zeta}^{\text{low}}(\hat{\pi}^h)$, it holds that $\widehat{R}^{(k)}(o_h) = R_h^{\mathcal{P}}(o_h)$. Thus, by the definition of the value function,

$$\begin{aligned} & \left| V_1^{\pi, \mathcal{P}, R^{\mathcal{P}}}(\emptyset) - V_1^{\pi, \mathcal{P}, \widehat{R}^{(k)}}(\emptyset) \right| \\ &= \left| \mathbb{E}_{(a_{1:H-1}, o_{2:H}) \sim \pi} \left[\sum_{h=2}^H R_h^{\mathcal{P}}(o_h) - \widehat{R}^{(k)}(o_h) \right] \right| \\ &\leq \mathbb{E}_{(a_{1:H-1}, o_{2:H}) \sim \pi} \left[\sum_{h=2}^H \mathbb{1}[(a_{h-L:h-1}, o_{h-L+1:h}) \notin \mathcal{Z}_{h, \zeta}^{\text{low}}(\hat{\pi}^h)] \right] \\ &\leq \frac{H A^{2L} O^L \zeta}{\phi} + \sum_{h=2}^H d_{\mathcal{S}, h-L}^{\mathcal{P}, \pi}(\mathcal{U}_{\phi, h-L}^{\mathcal{P}}(\pi^h)) \end{aligned} \quad (66)$$

$$\leq \frac{H A^{2L} O^L \zeta}{\phi} + H S \delta \leq \alpha/8, \quad (67)$$

where (66) uses Lemma F.4 for each $2 \leq h \leq H$, applied to the POMDP \mathcal{P} with $\pi' = \hat{\pi}^h$, and (67) uses (63) of Lemma I.4.

Combining (65) and (67) gives that, for all general policies π ,

$$\left| V_1^{\pi, \mathcal{P}, R^{\mathcal{P}}}(\emptyset) - V_1^{\pi, \widehat{\mathcal{M}}^{(k)}, \widehat{R}^{(k)}}(\emptyset) \right| \leq \frac{\alpha}{4}. \quad (68)$$

The above inequality holds in particular for the optimal policy π^* of \mathcal{P} . Since π_*^k (as computed in step 10 of BaSeCAMP) is the optimal policy of $\widehat{\mathcal{M}}^{(k)}$, we have (using (68))

$$\begin{aligned} V_1^{\pi_*^k, \mathcal{P}, R^{\mathcal{P}}}(\emptyset) &\geq V_1^{\pi_*^k, \widehat{\mathcal{M}}^{(k)}, \widehat{R}^{(k)}}(\emptyset) - \frac{\alpha}{4} \\ &\geq V_1^{\pi^*, \widehat{\mathcal{M}}^{(k)}, \widehat{R}^{(k)}}(\emptyset) - \frac{\alpha}{4} \\ &\geq V_1^{\pi^*, \mathcal{P}, R^{\mathcal{P}}}(\emptyset) - \frac{\alpha}{2}. \end{aligned}$$

Denote by \mathcal{E}^2 the event that, for all $k \in [K]$, the estimate \hat{r}^k produced in step 11 of Algorithm 3 is within $\alpha/4$ of the true value of π_*^k , namely $V_1^{\pi_*^k, \mathcal{P}, R^{\mathcal{P}}}(\emptyset)$. By Hoeffding's inequality and the union bound, \mathcal{E}^2 occurs with probability at least $1 - \beta/2$. Thus, under the event $\mathcal{E}^1 \cap \mathcal{E}^2$ (which occurs with probability at least $1 - \beta$), the policy $\pi_*^{k^*}$ output by BaSeCAMP (Algorithm 3) satisfies

$$V_1^{\pi_*^{k^*}, \mathcal{P}, R^{\mathcal{P}}}(\emptyset) \geq V_1^{\pi^*, \mathcal{P}, R^{\mathcal{P}}}(\emptyset) - \alpha,$$

as desired.

Finally, we analyze the runtime of Algorithm 3. Let T_1 be the running time for a single call of $\text{ConstructMDP}(L, N_0, N_1, \bar{\pi}^{k, 1:H})$, and let T_2 be the running time for a single call of $\text{BarySpannerPolicy}(\widehat{\mathcal{M}}^{(k)}, h)$. Since it takes time $\tilde{O}(|\mathcal{Z}|^2 AH)$ to compute an optimal policy of $\widehat{\mathcal{M}}^{(k)}$ for each $k \in [K]$ and $K = 2SH \leq 2OH$, the total running time is bounded above by

$$2OH \cdot \left(T_1 + T_2 + \tilde{O}(|\mathcal{Z}|^2 AH) + \tilde{O}\left(\frac{H^2}{\alpha^2}\right) \right),$$

where $\tilde{O}(\cdot)$ denotes that logarithmic factors (in $H, O, A, 1/\alpha$) are suppressed (above the $\tilde{O}(\cdot)$ accounts for bit arithmetic in the algorithm's operations).

Next, from the definition of ConstructMDP (Algorithm 1), it is straightforward to see that

$$T_1 = \tilde{O}(H^2 N_0 |\mathcal{Z}| AO).$$

Next we bound T_2 . The linear optimization oracle \mathcal{O} defined in step 3 of Algorithm 2 can be implemented in time $\tilde{O}(|\mathcal{Z}| \cdot HO)$ using dynamic programming, since the problem $\arg \max_{\pi \in \Pi_{\mathcal{Z}}^{\text{markov}}} \langle r, d_{\mathcal{O}, h-L}^{\pi, \widehat{\mathcal{M}}} \rangle$ is equivalent to finding an optimal policy in $\widehat{\mathcal{M}}$ when the reward function is given by r at step $h - L$ and is 0 otherwise. Step 4 requires $O(O^2 \log O)$ calls to \mathcal{O} , meaning that a single run of $\text{BarySpannerPolicy}(\widehat{\mathcal{M}}, h)$ (Algorithm 2) takes time

$$T_2 = \tilde{O}(O^3 H \cdot |\mathcal{Z}|). \tag{69}$$

Altogether, since $|\mathcal{Z}| = (AO)^L$, and using the definition of N_0 in Section C.1, the running time (and number of samples) taken by BaSeCAMP is $(AO)^{CL} \cdot \log(1/\beta)$, for some constant $C > 1$. \square