# Bayesian Learning of Optimal Policies in Markov Decision Processes with Countably Infinite State-Space 

Saghar Adler<br>University of Michigan

Vijay Subramanian<br>University of Michigan


#### Abstract

Models of many real-life applications, such as queueing models of communication networks or computing systems, have a countably infinite state-space. Algorithmic and learning procedures that have been developed to produce optimal policies mainly focus on finite state settings, and do not directly apply to these models. To overcome this lacuna, in this work we study the problem of optimal control of a family of discrete-time countable state-space Markov Decision Processes (MDPs) governed by an unknown parameter $\theta \in \Theta$, and defined on a countably-infinite state-space $\mathcal{X}=\mathbb{Z}_{+}^{d}$, with finite action space $\mathcal{A}$, and an unbounded cost function. We take a Bayesian perspective with the random unknown parameter $\boldsymbol{\theta}^{*}$ generated via a given fixed prior distribution on $\Theta$. To optimally control the unknown MDP, we propose an algorithm based on Thompson sampling with dynamically-sized episodes: at the beginning of each episode, the posterior distribution formed via Bayes' rule is used to produce a parameter estimate, which then decides the policy applied during the episode. To ensure the stability of the Markov chain obtained by following the policy chosen for each parameter, we impose ergodicity assumptions. From this condition and using the solution of the average cost Bellman equation, we establish an $\tilde{O}\left(d h^{d} \sqrt{|\mathcal{A}| T}\right)$ upper bound on the Bayesian regret of our algorithm, where $T$ is the time-horizon. Finally, to elucidate the applicability of our algorithm, we consider two different queueing models with unknown dynamics, and show that our algorithm can be applied to develop approximately optimal control algorithms.


## 1 Introduction

Many real-life applications, such as communication networks, supply chains, and computing systems, are modeled using queueing models with countably infinite state-space. In the existing analysis of these systems, the models are assumed to be known, but despite this, developing optimal control schemes is hard, with only a few examples worked out [35, 9, [54]. However, knowing the model, algorithmic procedures exist to produce approximately optimal policies [35] (such as value iteration and linear programming). Given the success of data-driven optimal control design, in particular Reinforcement Learning (RL), we explore the use of such methods for the countable state-space controlled Markov processes. However, current RL methods that focus on finite-state settings do not apply to the mentioned queueing models. With the model unknown, our goal is to develop a meta-learning scheme that is RL-based but obtains good performance by utilizing algorithms developed when models are known. Specifically, we study the problem of optimal control of a family of discrete-time countable state-space MDPs governed by an unknown parameter $\theta$ from a general space $\Theta$ with each MDP evolving on the countable state-space $\mathcal{X}=\mathbb{Z}_{+}^{d}$ and finite action space $\mathcal{A}$. The cost function is unbounded and polynomially dependent on the state, following the examples of minimizing waiting times in queueing systems. Taking a Bayesian view, we assume the model is governed by an unknown parameter $\boldsymbol{\theta}^{*} \in \Theta$ generated from a fixed and known prior distribution. We aim to learn a policy $\pi$ that minimizes the optimal infinite-horizon average cost over a given class of policies $\Pi$ with low Bayesian regret with respect to the (parameter-dependent) optimal policy in $\Pi$.

To avoid many technical difficulties in countably infinite state-space settings, it is crucial to establish certain assumptions regarding the class of models from which the unknown system is drawn; some examples are: i) the number of deterministic stationary policies is not finite; and ii) in average cost optimal control problems, without stability/ergodicity assumptions, an optimal policy may not exist [40], and when it exists, it may not be stationary or deterministic [20]. With these in mind, we assume that for any state-action pair, the transition kernels in the model class are categorical and skip-free to the right, i.e., with finite support with a bound depending on the state only in an additive manner; both are common features of queueing models where an increase in state is due to arrivals. A second set of assumptions ensure stability by assuming that the Markov chains obtained by using different policies in $\Pi$ are geometrically ergodic with uniformity across $\Theta$. From these assumptions, moments on hitting times are derived in terms of Lyapunov functions for polynomial ergodicity. These assumptions also yield a solution to the average cost optimality equation (ACOE) [9].

Contributions: To optimally control the unknown MDP, we propose an algorithm based on Thompson sampling with dynamically-sized episodes; posterior sampling is used based on its broad applicability and computational efficiency [46, 47]. At the beginning of each episode, a posterior distribution is formed using Bayes' rule, and an estimate is realized from this distribution which then decides the policy used throughout the episode. To evaluate the performance of our proposed algorithm, we use the metric of Bayesian regret, which compares the expected total cost achieved by a learning policy $\pi_{L}$ until time horizon $T$ with the policy achieving the optimal infinite-horizon average cost in the policy class $\Pi$. We consider regret guarantees in three different settings as follows:

1. In Theorem 1 , for $\Pi$ being the set of all policies and assuming that we have oracle access to the optimal policy for each parameter, we establish an $\tilde{O}\left(d h^{d} \sqrt{|\mathcal{A}| T}\right)$ upper bound on the Bayesian regret of this algorithm compared to the optimal policy.
2. In Corollary 1, where class $\Pi$ is a subset of all stationary policies and where we know the best policy within this subset for each parameter via an oracle, we prove an $\tilde{O}\left(d h^{d} \sqrt{|\mathcal{A}| T}\right)$ upper bound on the Bayesian regret of our proposed algorithm, relative to the best-in-class policy.
3. In Theorem 2, we explore a scenario where we have access to an approximately optimal policy, rather than the optimal policy in set $\Pi$ (which are all assumed to be stationary policies). When the approximately optimal policies satisfy Assumptions 3. 4 , we prove an $\tilde{O}\left(d h^{d} \sqrt{|\mathcal{A}| T}\right)$ regret bound, relative to the optimal policy in set $\Pi$.

Finally, to provide examples of our framework for developing approximately optimal control algorithms for stochastic systems with unknown dynamics, we study two different queueing models that meet our technical conditions. The first example is a continuous-time queueing system with two heterogeneous servers with unknown service rates and a common infinite buffer with the decision being the use of the slower server. Here, the optimal policy that minimizes the average waiting time is a threshold policy [38] which yields a queue-length after which the slower server is always used. The second model is a two-server queueing system, each with separate infinite buffers, to one of which a dispatcher routes an incoming arrival. Here, the optimal policy minimizing the waiting time is a switching-curve [26] with the specifics unknown for general parameter values, so we find the best policy within a commonly used set of switching-curve policies (Max-Weight policies [58, 59]), and assign the arrival to the queue with minimum weighted queue-length. For both models, we verify our assumptions for the class of optimal/best-in-class policies corresponding to different service rates and conclude that our proposed algorithm can be used to learn the optimal/best-in-class policy.
Related Work: Thompson sampling [62], or posterior sampling, has been applied to RL in many contexts of unknown MDPs [55, 45] and partially observed MDPs [28]; see tutorials [22, 50] for a comprehensive survey. It has been used in the parametric learning context [6] to minimize either Bayesian [46, 47, 49, 1, 60, 61] or frequentist [5, 23] regret. The bulk of the literature, including [5, 23, 49], analyzes finite-state and finite-action models but with different parameterizations such that a general dependence of the models on the parameters is allowed. The work in [61] studies general state-space MDPs but with a scalar parameterization with a Lipschitz dependence of the underlying models. Our problem formulation specifically considers countable state-space models with the models related via ergodicity, which we believe is a natural choice. Our focus on parametric learning is also connected to older work in adaptive control [3, 24] which studies asymptotically optimal learning for general parameter settings but with either a finite or countably infinite number of policies. Learning-based asymptotically optimal control in queues has a long history [36, 35] but recently there is increased work that also characterizes finite-time regret performance with respect to a well-known good policy or the optimal policy; see [63] for a survey. A series of work has studied
learning with Max-Weight policies to get stability and linear regret [44, 30] or just stability [65]. A recent related work [18] considers learning optimal paramterized policies in queueing networks when the MDP is known. In a finite or countable state-space setting of specific queueing models where the parameters can be estimated, many works [2, 17, 53, 32, 31, 14, 21, 16] have used forced exploration type schemes to obtain either regret that is constant or scaling logarithmically in the time-horizon.
Another line of work studies the problem of learning the optimal policy in an undiscounted finitehorizon MDP with a bounded reward function. Reference [66] uses a Thompson sampling-based learning algorithm with linear value function approximation to study an MDP with a bounded reward function in a finite-horizon setting. Reference [15] considers an episodic finite-horizon MDP with known bounded rewards but unknown transition kernels modeled using linearly parameterized exponential families. A maximum likelihood (ML) based algorithm coupled with exploration done by constructing high probability confidence sets around the ML estimate is used to learn the unknown parameters. In another work, [48] extends the problem setting of [15] to an episodic finite-horizon MDP with unknown rewards and transitions modeled using parametric bilinear exponential families. To learn the unknown parameters, they use a ML based algorithm with exploration done with explicit perturbation. We note that all mentioned works consider a finite-horizon problem. In contrast, our work considers an average cost problem, an infinite-horizon setting, and provides finite-time performance guarantees. In addition, these works focus on an MDP with a bounded reward function. Our focus, however, is learning in MDPs with unbounded rewards with the goal of covering practical queueing examples. We note that the parameterization of transitions used in [48, 15] can be used within our framework. However, similar to our work, additional stability assumptions are necessary to guarantee asymptotic learning and sub-linear regret. Another issue with exponential transition families is that they do not allow for 0 entries, which limits their applicability in queueing models.

In another work, [51] studies discounted MDPs with unknown dynamics, and unbounded state-space, but with bounded rewards, and learns an online policy that satisfies a specific notion of stability. It is also assumed that a Lyapunov function ensuring stability for the optimal policy exists. We note that [51] ignores optimality and focuses on finding a stable policy, which contrasts with our work that evaluates performance relative to the optimal policy. Secondly, [51] considers a discounted reward problem, essentially a finite-time horizon problem. Average cost problems, such as ours, are infinite-time horizon problems, so connections to discounted problems can only be made in the limit of the discount parameter going to 1 . Moreover, [51] considers a bounded reward function, simplifying their analysis but not practical for many queueing examples. Further, the assumption of a stable optimal policy with a Lyapunov function (as in [51]) is highly restrictive for bounded reward settings with discounting. Additionally, average cost problems with bounded costs need strong state-independent recurrence conditions for the existence of (stationary) optimal solutions, which many queueing examples don't satisfy; see [12]. Further complications can also arise with bounded costs: e.g., [20] shows that a stationary average cost optimal policy may not exist.

## 2 Problem formulation

We consider a family of discrete-time Markov Decision Processes (MDPs) governed by parameter $\theta \in \Theta$ with the MDP for parameter $\theta$ described by $\left(\mathcal{X}, \mathcal{A}, c, P_{\theta}\right)$. For exposition purposes, we assume that all the MDPs are on (a common) countably infinite state-space $\mathcal{X}=\mathbb{Z}_{+}^{d}$. We denote the finite action space by $\mathcal{A}$, the transition kernel by $P_{\theta}: \mathcal{X} \times \mathcal{A} \rightarrow \Delta(\mathcal{X})$, and the cost function by $c: \mathcal{X} \times \mathcal{A} \rightarrow \mathbb{R}_{+}$. As mentioned earlier, we will take a Bayesian view of the problem and assume that the model is generated using an unknown parameter $\boldsymbol{\theta}^{*} \in \Theta$, which is generated from a given fixed prior distribution $\nu(\cdot)$ on $\Theta$. Our goal is to find a policy $\pi: \mathcal{X} \rightarrow \mathcal{A}$ that tries to achieve Bayesian optimal performance in policy class $\Pi$, i.e., minimizes the expected regret with $\boldsymbol{\theta}^{*}$ chosen from the prior distribution $\nu(\cdot)$. For each value $\theta \in \Theta$, the minimum infinite-horizon average cost is defined as

$$
\begin{equation*}
J(\theta)=\inf _{\pi \in \Pi} \limsup _{T \rightarrow \infty} \frac{1}{T} \mathbb{E}\left[\sum_{t=1}^{T} c(\boldsymbol{X}(t), A(t))\right] \tag{1}
\end{equation*}
$$

where we optimize over a given class of policies $\Pi$ and $\boldsymbol{X}(t)=\left(X_{1}(t), \ldots, X_{d}(t)\right) \in \mathcal{X}$ and $A(t) \in \mathcal{A}$ are the state and action at $t \in \mathbb{N}$. Typically, we set this class to be all (causal) policies, but it is also possible to consider $\Pi$ to be a proper subset of all policies as we will explore in our results. For a learning policy $\pi_{L}$ that aims to select the optimal control without model knowledge but with
knowledge of $\Theta$ and the prior $\nu$, the Bayesian regret until time horizon $T \geq 2$ is defined as

$$
\begin{equation*}
R\left(T, \pi_{L}\right)=\mathbb{E}\left[\sum_{t=1}^{T}\left[c(\boldsymbol{X}(t), A(t))-J\left(\boldsymbol{\theta}^{*}\right)\right]\right] \tag{2}
\end{equation*}
$$

where the expectation is taken over $\boldsymbol{\theta}^{*} \sim \nu$ and the dynamics induced by $\pi_{L}$. Owing to underlying challenges in countable state-space MDPs, we require the below assumptions on the cost function.
Assumption 1. The cost function $c: \mathcal{X} \times \mathcal{A} \rightarrow \mathbb{R}_{+}$is assumed to satisfy the following two conditions:

1. For every number $z \geq 0$ and action $a \in \mathcal{A}, c(\boldsymbol{x}, a) \geq z$ outside a finite subset of $\mathcal{X}$.
2. The cost function is upper-bounded by a multivariate polynomial $f_{c}: \mathbb{Z}_{+}^{d} \rightarrow \mathbb{R}_{+}$which is increasing in every component on $\boldsymbol{x} \in \mathbb{Z}_{+}^{d}$ and has maximum degree of $r(\geq 1)$ in any dimension. We can assume that $f_{c}(\boldsymbol{x})=K \sum_{i=1}^{d}\left(x_{i}\right)^{r}$ for some $K>0$, where $\boldsymbol{x}=\left(x_{1}, \ldots, x_{d}\right)$.

Thus, the cost function increases without bound (in the state) at a polynomial rate. This assumption is common in practice-holding costs in queueing models are polynomial in the state components. To avoid technical issues the infinite state-space setting also necessitates some assumptions on the class from which the unknown model is drawn. For instance, irreducibility of Markov chains on such state-spaces does not ensure positive recurrence (and ergodicity). Moreover, for average cost optimal control problems, without stability even the existence of an optimal policy is not guaranteed, and we need more conditions. The following assumption ensures a skip-free behaviour for transitions, which holds in many queueing models, where an increase in state corresponds to (new) arrivals.
Assumption 2. From any state-action pair $(\boldsymbol{x}, a)$, the transition is to a finite number of states. We also assume that all transition kernels are skip-free to the right: for some $h \geq 1$ which is independent of $\theta \in \Theta$ and $(\boldsymbol{x}, a) \in \mathcal{X} \times \mathcal{A}$, we have $P_{\theta}\left(\boldsymbol{x}^{\prime} ; \boldsymbol{x}, a\right)=0$ for all $\boldsymbol{x}^{\prime} \in\left\{\tilde{\boldsymbol{x}} \in \mathbb{Z}_{+}^{d}:\|\tilde{\boldsymbol{x}}\|_{1}>\|\boldsymbol{x}\|_{1}+h\right\}$.

Learning necessitates some commonalities within the class of models so that using a policy well-suited to one model provides information on other models too. For us, these are in the form of constraints on the transition kernels of the models and stability assumptions. As simple union bound arguments don't work in the countably infinite state-space setting, we will use the stability assumptions instead. In our setting, we consider a class of models, each with a policy being well-suited to at least one model in the class, and use the set of policies to search within. Using a reduced set of policies is necessary as the number of deterministic stationary policies is infinite. To learn correctly while restricting attention to this subset policy class, requires some regularity assumptions when a policy well-suited to one model is tried on a different model. Our ergodicity assumptions are one convenient choice; see Appendix A. 1 for details. These assumptions let us characterize the distributions of the first passage times of the Markov processes via stability conditions; see Lemmas 10 and 11
Assumption 3. For any $\operatorname{MDP}\left(\mathcal{X}, \mathcal{A}, c, P_{\theta}\right)$ with parameter $\theta \in \Theta$, there exists a unique optimal policy $\pi_{\theta}^{*}$ that minimizes the infinite-horizon average cost within the class of policies $\Pi$. Furthermore, for any $\theta_{1}, \theta_{2} \in \Theta$, the Markov process with transition kernel $P_{\theta_{1}}^{\pi_{\theta_{2}}^{*}}$ obtained from the MDP $\left(\mathcal{X}, \mathcal{A}, c, P_{\theta_{1}}\right)$ by following policy $\pi_{\theta_{2}}^{*}$ is irreducible, aperiodic, and geometrically ergodic with geometric ergodicity coefficient $\gamma_{\theta_{1}, \theta_{2}}^{g} \in(0,1)$ and stationary distribution $\mu_{\theta_{1}, \theta_{2}}$. This is equivalent to the existence of finite set $C_{\theta_{1}, \theta_{2}}^{g}$ and Lyapunov function $V_{\theta_{1}, \theta_{2}}^{g}: \mathcal{X} \rightarrow[1,+\infty)$ satisfying

$$
\Delta V_{\theta_{1}, \theta_{2}}^{g}(\boldsymbol{x}) \leq-\left(1-\gamma_{\theta_{1}, \theta_{2}}^{g}\right) V_{\theta_{1}, \theta_{2}}^{g}(\boldsymbol{x}), \boldsymbol{x} \in \mathcal{X} \backslash C_{\theta_{1}, \theta_{2}}^{g} \text { and } P_{\theta_{1}}^{\pi_{\theta_{2}}^{*}} V_{\theta_{1}, \theta_{2}}^{g}(\boldsymbol{x})<+\infty, \boldsymbol{x} \in C_{\theta_{1}, \theta_{2}}^{g}
$$

where $\Delta V_{\theta_{1}, \theta_{2}}^{g}(\boldsymbol{x}):=P_{\theta_{1}}^{\pi_{\theta_{2}}^{*}} V_{\theta_{1}, \theta_{2}}^{g}(\boldsymbol{x})-V_{\theta_{1}, \theta_{2}}^{g}(\boldsymbol{x})$. Setting $b_{\theta_{1}, \theta_{2}}^{g}:=\max _{\boldsymbol{x} \in C_{\theta_{1}, \theta_{2}}^{g}} P_{\theta_{1}}^{\pi_{\theta_{2}}^{*}} V_{\theta_{1}, \theta_{2}}^{g}(\boldsymbol{x})+$ $V_{\theta_{1}, \theta_{2}}^{g}(\boldsymbol{x})$ yields

$$
\begin{equation*}
\Delta V_{\theta_{1}, \theta_{2}}^{g}(\boldsymbol{x}) \leq-\left(1-\gamma_{\theta_{1}, \theta_{2}}^{g}\right) V_{\theta_{1}, \theta_{2}}^{g}(\boldsymbol{x})+b_{\theta_{1}, \theta_{2}}^{g} \mathbb{I}_{C_{\theta_{1}, \theta_{2}}^{g}}^{g}(\boldsymbol{x}), \quad \boldsymbol{x} \in \mathcal{X} \tag{3}
\end{equation*}
$$

Then, we have the following assumptions relating all the models in $\Theta$ :

1. The geometric ergodicity coefficient is uniformly bounded below 1 : $\gamma_{*}^{g}:=\sup _{\theta_{1}, \theta_{2} \in \Theta} \gamma_{\theta_{1}, \theta_{2}}^{g}<1$.
2. We assume that $\left\{0^{d}\right\} \subseteq \cap_{\theta_{1}, \theta_{2} \in \Theta} C_{\theta_{1}, \theta_{2}}^{g}$ and $C_{*}^{g}=\cup_{\theta_{1}, \theta_{2} \in \Theta} C_{\theta_{1}, \theta_{2}}^{g}$ is a finite set. We further assume that $b_{*}^{g}:=\sup _{\theta_{1}, \theta_{2}} b_{\theta_{1}, \theta_{2}}^{g}<+\infty$.

Remark 1. The uniqueness of the optimal policy is not essential for the validity of our results, provided that all optimal policies satisfy our assumptions. When this condition is not met, we need to select an optimal policy that is geometrically ergodic for all $\theta \in \Theta$. This issue can be avoided by using a smaller subset of policies for which ergodicity can be shown, such as Max-Weight policies.

Geometric ergodicity implies that all moments of the hitting time of state $0^{d}$, say $\tau_{0^{d}}$, from any initial state $\boldsymbol{x} \neq 0^{d}$ are finite as $\mathbb{E}_{\boldsymbol{x}}\left[\kappa^{\tau_{0} d}\right] \leq c_{1} V^{g}(\boldsymbol{x})$ (for specific $\kappa>1$ and $c_{1}$ ), and so, $\mathbb{E}_{\boldsymbol{x}}\left[\tau_{0^{d}}^{k}\right] \leq$ $c_{1} V^{g}(\boldsymbol{x}) k!/ \log ^{k}(\kappa)<+\infty$ for all $k \in \mathbb{N}$; see Appendix A. 2 Function $V^{g}$ is typically exponential in some norm of the state and yields an exponential bound for moments of hitting times, and a poor regret bound. To improve the regret bound, we need a different drift equation with function $V^{p}$ with polynomial dependence on a norm of the state that bounds certain polynomial moments of $\tau_{0^{d}}$.
Assumption 4. Given $\theta_{1}, \theta_{2} \in \Theta$, Markov process obtained from $\operatorname{MDP}\left(\mathcal{X}, \mathcal{A}, c, P_{\theta_{1}}\right)$ by following policy $\pi_{\theta_{2}}^{*}$ is polynomially ergodic through the Foster-Lyapunov criteria: there exists a finite set $C_{\theta_{1}, \theta_{2}}^{p}$, constants $\beta_{\theta_{1}, \theta_{2}}^{p}, b_{\theta_{1}, \theta_{2}}^{p}>0, \alpha_{\theta_{1}, \theta_{2}}^{p} \in\left[\frac{r}{r+1}, 1\right)$, and function $V_{\theta_{1}, \theta_{2}}^{p}: \mathcal{X} \rightarrow[1,+\infty)$ satisfying

$$
\begin{equation*}
\Delta V_{\theta_{1}, \theta_{2}}^{p}(\boldsymbol{x}) \leq-\beta_{\theta_{1}, \theta_{2}}^{p}\left(V_{\theta_{1}, \theta_{2}}^{p}(\boldsymbol{x})\right)^{\alpha_{\theta_{1}, \theta_{2}}^{p}}+b_{\theta_{1}, \theta_{2}}^{p} \mathbb{I}_{C_{\theta_{1}, \theta_{2}}^{p}}(\boldsymbol{x}), \quad \boldsymbol{x} \in \mathcal{X} \tag{4}
\end{equation*}
$$

Then, we have the following assumptions relating all the models in $\Theta$ :

1. $V_{\theta_{1}, \theta_{2}}^{p}$ is a polynomial with positive coefficients, maximum degree (in any dimension) $r_{\theta_{1}, \theta_{2}}^{p}$, and sum of coefficients $s_{\theta_{1}, \theta_{2}}^{p}$. We assume $r_{*}^{p}=\sup _{\theta_{1}, \theta_{2}} r_{\theta_{1}, \theta_{2}}^{p}<\infty$ and $s_{*}^{p}=\sup _{\theta_{1}, \theta_{2}} s_{\theta_{1}, \theta_{2}}^{p}<\infty$.
2. We assume that $\left\{0^{d}\right\} \subseteq \cap_{\theta_{1}, \theta_{2} \in \Theta} C_{\theta_{1}, \theta_{2}}^{p}$ and $C_{*}^{p}=\cup_{\theta_{1}, \theta_{2} \in \Theta} C_{\theta_{1}, \theta_{2}}^{p}$ is a finite set. We further assume that $\beta_{*}^{p}:=\inf _{\theta_{1}, \theta_{2}} \beta_{\theta_{1}, \theta_{2}}^{p}>0$ and $b_{*}^{p}:=\sup _{\theta_{1}, \theta_{2}} b_{\theta_{1}, \theta_{2}}^{p}<\infty$.
3. Let $K_{\theta_{1}, \theta_{2}}(\boldsymbol{x}):=\sum_{n=0}^{\infty} 2^{-n-2}\left(P_{\theta_{1}}^{\pi_{\theta_{2}}^{*}}\right)^{n}\left(\boldsymbol{x}, 0^{d}\right)$, which is positive for any pair $\theta_{1}, \theta_{2} \in \Theta$ by irreducibility. We assume that it is strictly positive in $\Theta: K_{*}:=\inf _{\theta_{1}, \theta_{2}} \min _{\boldsymbol{x} \in C_{*}^{p}} K_{\theta_{1}, \theta_{2}}(\boldsymbol{x})>0$.

Assumptions 3 4hold in many models of interest; see Appendix E. As average cost optimality is our design criterion, we need to ensure the existence of solutions to ACOE when $\Pi$ is the set of all policies, or Poisson equation when $\Pi$ is a subset of all policies. We discuss these two cases separately. Case 1: $\Pi$ is the set of all policies. For any parameter $\theta \in \Theta$, the $\operatorname{MDP}\left(\mathcal{X}, \mathcal{A}, c, P_{\theta}\right)$ is said to satisfy the ACOE if there exists a constant $J(\theta)$ and a unique function $v(\cdot ; \theta): \mathcal{X} \rightarrow \mathbb{R}$ such that

$$
J(\theta)+v(\boldsymbol{x} ; \theta)=\min _{a \in \mathcal{A}}\left\{c(\boldsymbol{x}, a)+\sum_{\boldsymbol{y} \in \mathcal{X}} P_{\theta}(\boldsymbol{y} \mid \boldsymbol{x}, a) v(\boldsymbol{y} ; \theta)\right\} \text { with } v\left(0^{d} ; \theta\right)=0 .
$$

From [13] if the following conditions hold, ACOE has a solution, $J_{\theta}$ is the optimal infinite-horizon average cost, and there is an optimal stationary policy with ACOE becoming (5): (i) for every $(\boldsymbol{x}, a)$ and $z \geq 0$, cost function $c(\boldsymbol{x}, a) \geq z$ outside a finite subset of $\mathcal{X}$; (ii) there is a stationary policy with an irreducible and aperiodic Markov process with finite average cost; and (iii) from every ( $\boldsymbol{x}, a$ ) transition to a finite number of states is possible. From Assumptions 1/3, the above conditions hold. Case 2: $\Pi$ is a proper subset of all policies. Here, we posit that for every $\theta \in \Theta$ and its best in-class policy $\pi_{\theta}^{*}$, there exists a constant $J(\theta)$, the average cost of $\pi_{\theta}^{*}$, and a function $v(\cdot ; \theta): \mathcal{X} \rightarrow \mathbb{R}$ with

$$
\begin{equation*}
J(\theta)+v(\boldsymbol{x} ; \theta)=c\left(\boldsymbol{x}, \pi_{\theta}^{*}(\boldsymbol{x})\right)+\sum_{\boldsymbol{y} \in \mathcal{X}} P_{\theta}\left(\boldsymbol{y} \mid \boldsymbol{x}, \pi_{\theta}^{*}(\boldsymbol{x})\right) v(\boldsymbol{y} ; \theta) . \tag{5}
\end{equation*}
$$

This holds by the solution of the Poisson equation with the appropriate forcing function. For a Markov process $\boldsymbol{X}$ on the space $\mathcal{X}$ with transition kernel $P$ and cost function $\bar{c}(\cdot)$, a solution to the Poisson equation [41] is a scalar $J$ and function $v(\cdot): \mathcal{X} \mapsto \mathbb{R}$ such that $J+v=\bar{c}+P v$, where $v(\boldsymbol{z})=0$ for some $z \in \mathcal{X}$. In our setting using [41, Sections 9.6-9.8], for a model governed by $\theta \in \Theta$ following policy $\pi_{\theta}^{*}$, we show a solution to the Poisson equation exists and is given by $v^{\pi_{\theta}^{*}}\left(0^{d}\right)=0$, and

$$
\begin{equation*}
J(\theta)=\bar{C}^{\pi_{\theta}^{*}}\left(0^{d}\right) / \mathbb{E}_{0^{d}}^{\pi_{\theta}^{*}}\left[\tau_{0^{d}}\right] \text { and } v^{\pi_{\theta}^{*}}(\boldsymbol{x})=\bar{C}^{\pi_{\theta}^{*}}(\boldsymbol{x})-J(\theta) \mathbb{E}_{\boldsymbol{x}}^{\pi_{\theta}^{*}}\left[\tau_{0^{d}}\right], \quad \forall \boldsymbol{x} \in \mathcal{X} \tag{6}
\end{equation*}
$$

where $\bar{C}^{\pi_{\theta}^{*}}(\boldsymbol{x})=\mathbb{E}_{\boldsymbol{x}}^{\pi_{\theta}^{*}}\left[\sum_{i=0}^{\tau_{0 d}-1} c\left(\boldsymbol{X}(i), \pi_{\theta}^{*}(\boldsymbol{X}(i))\right)\right]$, and expectation is over trajectories of Markov chain $\boldsymbol{X}$ with transition kernel $P_{\theta}^{\pi_{\theta}^{*}}$ starting in state $\boldsymbol{x}$. In Appendix A.3 we present related definitions and show that from Assumptions $3 \sqrt{4}$, the requirements for the existence and finiteness of the solutions to Poisson equation are satisfied. Finally, we assume $\sup _{\theta \in \Theta} J(\theta)$ is finite, which typically holds as a result of the boundedness assumptions stated in Asumptions 3 or 4 along with Assumption 1

```
Algorithm 1 Thompson Sampling with Dynamically-sized Episodes (TSDE)
    Input: \(\nu_{0}\)
    Initialization: \(\boldsymbol{X}(1)=0^{d}, t \leftarrow 1\)
    for episodes \(k=1,2, \ldots\) do
        \(t_{k} \leftarrow t\)
        Generate \(\theta_{k} \sim \nu_{t_{k}}\)
        while \(t \leq t_{k}+\tilde{T}_{k-1}\) and \(N_{t}(\boldsymbol{x}, a) \leq 2 N_{t_{k}}(\boldsymbol{x}, a)\) for all \((\boldsymbol{x}, a) \in \mathcal{X} \times \mathcal{A}\) do
            Apply action \(A(t)=\pi_{\theta_{k}}^{*}(\boldsymbol{X}(t))\)
            \(N_{t}(\boldsymbol{X}(t), A(t)) \leftarrow N_{t}(\boldsymbol{X}(t), A(t))+1\)
            Observe new state \(\boldsymbol{X}(t+1)\)
            Update \(\nu_{t+1}\) according to (7)
            \(t \leftarrow t+1\)
        end while
        \(\tilde{T}_{k} \leftarrow t-t_{k}\)
        while \(\boldsymbol{X}(t) \neq 0^{d}\) do
            Apply action \(A(t)=\pi_{\theta_{k}}^{*}(\boldsymbol{X}(t))\)
            Observe new state \(\boldsymbol{X}(t+1)\)
        end while
        \(T_{k} \leftarrow t-t_{k}\)
    end for
```

Remark 2. In Assumption 4 we can use any other policy $\pi_{\theta_{2}}$ such that the Markov process obtained from $\operatorname{MDP}\left(\mathcal{X}, \mathcal{A}, c, P_{\theta_{1}}\right)$ by following policy $\pi_{\theta_{2}}$ is irreducible and polynomially ergodic via the Foster-Lyapunov criteria with the uniformity discussed. Irreducibility is important as the policy will be used at times when the state is not known in advance, specifically at Steps 14-17 in Algorithm 11
Assumption 5. We assume that $J^{*}:=\sup _{\theta \in \Theta} J(\theta)<+\infty$.

## 3 Thompson sampling based learning algorithm

We will use the learning algorithm Thompson sampling with dynamically-sized episodes from [49] to learn the unknown parameter $\boldsymbol{\theta}^{*} \in \Theta$ and the corresponding policy, $\pi_{\boldsymbol{\theta}^{*}}^{*}$, but suitably modify it for our countable state-space setting. Consider the prior distribution $\nu_{0}=\nu$ defined on $\Theta$ from which $\boldsymbol{\theta}^{*}$ is sampled. At each time $t \in \mathbb{N}$, the posterior distribution $\nu_{t}$ is updated according to Bayes' rule as

$$
\begin{equation*}
\nu_{t+1}(d \theta)=\frac{\mathbb{P}_{\theta}(\boldsymbol{X}(t+1) \mid \boldsymbol{X}(t), A(t)) \nu_{t}(d \theta)}{\int_{\theta^{\prime} \in \Theta} \mathbb{P}_{\theta^{\prime}}(\boldsymbol{X}(t+1) \mid \boldsymbol{X}(t), A(t)) \nu_{t}\left(d \theta^{\prime}\right)}, \tag{7}
\end{equation*}
$$

and the posterior estimate $\theta_{t+1}$, if generated, is from the posterior distribution $\nu_{t+1}$. The modified Thompson-sampling with dynamically-sized episodes algorithm (TSDE) is presented in Algorithm 1 The TSDE algorithm operates in episodes: at the beginning of episode $k$, parameter $\theta_{k}$ is sampled from the posterior distribution $\nu_{t_{k}}$ and during episode $k$, actions are generated from the stationary policy according to $\theta_{k}$, i.e., $\pi_{\theta_{k}}^{*}$. Let $t_{k}$ be the time the $k$-th episode begins. Define $\tilde{t}_{k+1}$ as the first time after $t_{k}$ that the conditions of Line 6 of Algorithm 1 is triggered and $t_{k+1}$ as the first time at or after $\tilde{t}_{k+1}$ where state $0^{d}$ is visited; for the last episode started before or at $T$, we ensure that $t_{k}$ and $\tilde{t}_{k}$ are less than or equal $T+1$. Explicitly, $t_{1}=1$ and for $k>1, t_{k}=\min _{\tilde{T_{2}}}\left\{t \geq \tilde{t}_{k}: \boldsymbol{X}(t)=0^{d}\right.$ or $\left.t>T\right\}$. Let $T_{k}=t_{k+1}-t_{k}$ be the length of the $k$-th episode and set $\tilde{T}_{k}=\tilde{t}_{k+1}-t_{k}$ with the convention $\tilde{T}_{0}=1$. For any state-action pair $(\boldsymbol{x}, a)$, we define $N_{1}(\boldsymbol{x}, a)=0$ and for $t>1$,

$$
N_{t}(\boldsymbol{x}, a)=\mid\left\{t_{k} \leq i<\tilde{t}_{k+1} \leq t \text { for some } k \geq 1:(\boldsymbol{X}(i), A(i))=(\boldsymbol{x}, a)\right\} \mid
$$

Notice that for all state-action pairs $(\boldsymbol{x}, a)$ and $\tilde{t}_{k+1} \leq t \leq t_{k+1}$, we have $N_{t}(\boldsymbol{x}, a)=N_{\tilde{t}_{k+1}}(\boldsymbol{x}, a)$. We denote $K_{T}$ as the number of episodes started by or at time $T$, or $K_{T}=\max \left\{k: t_{k} \leq T\right\}$. The length of episode $k<K_{T}$ is not fixed and is determined according to two stopping criteria: (1) $t>t_{k}+\tilde{T}_{k-1}$, (2) $N_{t}(\boldsymbol{x}, a)>2 N_{t_{k}}(\boldsymbol{x}, a)$ for some state-action pair $(\boldsymbol{x}, a)$. After either criterion is met, the system will still follow policy $\pi_{\theta_{k}}^{*}$ until the first time at which state $0^{d}$ is visited; see Line 14 and Figure 1 . We use this settling period to $0^{d}$ because the system state can be arbitrary when the


Figure 1: MDP evolution in episode $k<K_{T}$.
first stopping criterion is met. As the countable state-space setting precludes a simple union-bound argument to overcome this uncertainty (as in the literature for finite state settings), we let the system reach the special state $0^{d}$. Another (essentially equivalent) option is to wait until the state hits the finite set $C_{*}^{g}$ or $C_{*}^{p}$ and then use a union bound argument for all states in either set. For analytical convenience, we only use the state samples observed before arrival $\tilde{t}_{k+1}$ to update the posterior distribution. The posterior update is halted during the settling period to $0^{d}$ as we have no control on the states visited during it, despite it being finite in duration (by our assumptions).

## 4 Regret analysis of Algorithm 1

The performance of any learning policy $\pi_{L}$ is evaluated using the metric of expected regret compared to the optimal expected average cost of true parameter $\boldsymbol{\theta}^{*}$, namely, $J\left(\boldsymbol{\theta}^{*}\right)$. In this section, we evaluate the performance of Algorithm 1 and derive an upper bound for $R\left(T, \pi_{T S D E}\right)$, its expected regret up to time $T$. In Section 2, we argued that at time $t$ in episode $k\left(t_{k} \leq t<t_{k+1}\right)$, there exist a constant $J\left(\theta_{k}\right)$ and a unique function $v\left(\cdot ; \theta_{k}\right): \mathcal{X} \rightarrow \mathbb{R}$ such that $v\left(0^{d} ; \theta_{k}\right)=0$ and

$$
\begin{equation*}
J\left(\theta_{k}\right)+v\left(\boldsymbol{X}(t) ; \theta_{k}\right)=c\left(\boldsymbol{X}(t), \pi_{\theta_{k}}^{*}(\boldsymbol{X}(t))\right)+\sum_{\boldsymbol{y} \in \mathcal{X}} P_{\theta_{k}}\left(\boldsymbol{y} \mid \boldsymbol{X}(t), \pi_{\theta_{k}}^{*}(\boldsymbol{X}(t))\right) v\left(\boldsymbol{y} ; \theta_{k}\right), \tag{8}
\end{equation*}
$$

in which $\pi_{\theta_{k}}^{*}$ is the optimal or best-in-class policy (depending on the context) according to parameter $\theta_{k}$ and $J\left(\theta_{k}\right)$ is the average cost for the Markov process obtained from $\operatorname{MDP}\left(\mathcal{X}, \mathcal{A}, c, P_{\theta_{k}}\right)$ by following $\pi_{\theta_{k}}^{*}$. We derive a bound for the expected regret $R\left(T, \pi_{T S D E}\right)$ following the proof steps of [49] while extending it to the countable state-space setting of our problem. Using (8), the regret is decomposed into three terms and each term is bounded separately:

$$
\begin{align*}
& R\left(T, \pi_{T S D E}\right)=\mathbb{E}\left[\sum_{k=1}^{K_{T}} \sum_{t=t_{k}}^{t_{k+1}-1} c\left(\boldsymbol{X}(t), \pi_{\theta_{k}}^{*}(\boldsymbol{X}(t))\right)\right]-T \mathbb{E}\left[J\left(\boldsymbol{\theta}^{*}\right)\right]=R_{0}+R_{1}+R_{2}  \tag{9}\\
& \text { with } R_{0}=\mathbb{E}\left[\sum_{k=1}^{K_{T}} T_{k} J\left(\theta_{k}\right)\right]-T \mathbb{E}\left[J\left(\boldsymbol{\theta}^{*}\right)\right],  \tag{10}\\
& R_{1}=\mathbb{E}\left[\sum_{k=1}^{K_{T}} \sum_{t=t_{k}}^{t_{k+1}-1}\left[v\left(\boldsymbol{X}(t) ; \theta_{k}\right)-v\left(\boldsymbol{X}(t+1) ; \theta_{k}\right)\right]\right]  \tag{11}\\
& R_{2}=\mathbb{E}\left[\sum_{k=1}^{K_{T}} \sum_{t=t_{k}}^{t_{k+1}-1}\left[v\left(\boldsymbol{X}(t+1) ; \theta_{k}\right)-\sum_{\boldsymbol{y} \in \mathcal{X}} P_{\theta_{k}}\left(\boldsymbol{y} \mid \boldsymbol{X}(t), \pi_{\theta_{k}}^{*}(\boldsymbol{X}(t))\right) v\left(\boldsymbol{y} ; \theta_{k}\right)\right]\right] . \tag{12}
\end{align*}
$$

Before bounding the above regret terms, we address the complexities arising from the countable state-space setting. Firstly, we need to study the maximum state (with respect to the $\ell_{\infty}$-norm) visited up to time $T$ in the $\operatorname{MDP}\left(\mathcal{X}, \mathcal{A}, c, P_{\boldsymbol{\theta}^{*}}\right)$ following Algorithm 1 , we denote this maximum state by $M_{\boldsymbol{\theta}^{*}}^{T}$. In Appendix C, we derive upper bounds on the moments of hitting times of state $0^{d}$ and utilize this to bound the moments of random variable $M_{\boldsymbol{\theta}^{*}}^{T}$, which then lets us study the number of episodes $K_{T}$ by time $T$. Another challenge in analyzing the regret is that the relative value function $v(\boldsymbol{x} ; \theta)$ is unlikely to be bounded in the countable state-space setting. Hence, in (13) and (14), we find bounds for the relative value function in terms of hitting time $\tau_{0^{d}}$ from the initial state $\boldsymbol{x}$. Based on these results, we provide an upper bound for the regret of Algorithm 1 in Theorem 1.
Maximum state norm under polynomial and geometric ergodicity. Here we state the results that characterize the maximum $l_{\infty}$-norm of the state vector achieved up until and including time $T$, and
the resulting bounds on the number of episodes executed until time $T$. Owing to space constraints the details (including formal statements) are presented in Appendix B The results are listed as below: (i) In Lemma 6, we bound the moments of the maximum length of recurrence times of state $0^{d}$, using the ergodicity assumptions 3 and 4 . This, along with the skip-free property, allows us to prove that the $p$-th moment of $\max _{1 \leq i \leq T} \tau_{0^{d}}^{(i)}$ and $M_{\boldsymbol{\theta}^{*}}^{T}$ are both of order $O\left(\log ^{p} T\right)$.
(ii) In Lemma 7 , we find an upper bound for the number of episodes in which the second stopping criterion is met or there exists a state-action pair for which $N_{t}(\boldsymbol{x}, a)$ has increased more than twice. (iii) In Lemma 8, we bound the total number of episodes $K_{T}$ by time $T$ by bounding the number of episodes triggered by the first stopping criterion, using the fact that in such episodes, $\tilde{T}_{k}=\tilde{T}_{k-1}+1$. Moreover, to account for the settling time of each episode, we use geometric ergodicity and Lemma 6 It follows that the expected value of the number of episodes $K_{T}$ is of the order $\tilde{O}\left(\sqrt{h^{d}|\mathcal{A}| T}\right)$.
Regret analysis. Next, we bound regret terms $R_{0}, R_{1}$ and $R_{2}$ using the approach of [49] along with additional arguments to extend their result to a countably infinite state-space. We consider the relative value function $v(\boldsymbol{x} ; \theta)$ of policy $\pi_{\theta}^{*}$ introduced for the optimal policy in ACOE or for the best in-class policy in the Poisson equation. In either of these cases, policy $\pi_{\theta}^{*}$ satisfies (5), which is the corresponding Poisson equation with forcing function $c\left(\boldsymbol{x}, \pi_{\theta}^{*}(\boldsymbol{x})\right)$ in a Markov chain with transition matrix $P_{\theta}^{\pi_{\theta}^{*}}$. In (6), we presented the solution $(J, v)$ to the Poisson equation, which yields the following upper bound for the relative value function, as argued in Appendix A. 3 .

$$
\begin{equation*}
v(\boldsymbol{x} ; \theta) \leq \bar{C}^{\pi_{\theta}^{*}}(\boldsymbol{x}) \leq \mathbb{E}_{\boldsymbol{x}}^{\pi_{\theta}^{*}}\left[K d\left(\|\boldsymbol{x}\|_{\infty}+h \tau_{0^{d}}\right)^{r} \tau_{0^{d}}\right] \tag{13}
\end{equation*}
$$

We can similarly lower bound the relative value function using Assumption 5 as

$$
\begin{equation*}
v(\boldsymbol{x} ; \theta) \geq-J(\theta) \mathbb{E}_{\boldsymbol{x}}^{\pi_{\theta}^{*}}\left[\tau_{0^{d}}\right] \geq-J^{*} \mathbb{E}_{\boldsymbol{x}}^{\pi_{\theta}^{*}}\left[\tau_{0^{d}}\right] \tag{14}
\end{equation*}
$$

From Assumption 3 all moments of $\tau_{0^{d}}$ and thus, the derived bounds are finite. Also, in Lemma 10 we bound the moments of $\tau_{0^{d}}$ of order $i \leq r+1$ using the polynomial Lyapunov function $V_{\theta_{1}, \theta_{2}}^{p}$, which is then used to bound the expected regret. We next bound the first regret term $R_{0}$ from the first stopping criterion in terms of the number of episodes $K_{T}$ and the settling time of each episode $k$.
Lemma 1. The first regret term $R_{0}$ satisfies $R_{0} \leq J^{*} \mathbb{E}\left[K_{T}\left(\max _{1 \leq i \leq T} \tau_{0^{d}}^{(i)}+1\right)\right]$.
Proof of Lemma 1 is given in Appendix B. 4 . From Lemma 6, all moments of $\max _{1 \leq i \leq T} \tau_{0^{d}}^{(i)}$ are bounded by a polylogarithmic function. Futhermore, as a result of Lemma 8 , expected value of the number of episodes $K_{T}$ is of the order $\tilde{O}\left(\sqrt{h^{d}|\mathcal{A}| T}\right)$, which leads to a $\tilde{O}\left(\sqrt{h^{d}|\mathcal{A}| T}\right)$ regret term $R_{0}$. Next, an upper bound on $R_{1}$ defined in (11) is derived. In the proof of Lemma 2 we argue that as the relative value function is equal to 0 at all time instances $t_{k}$ for $k \leq K_{T}$, the only term that contributes to the regret is the value function at the end of time horizon $T$. We use the lower bound derived in (14) to show that the second regret term $R_{1}$ is $\tilde{O}(1)$; the proof is given in Appendix B. 5

Lemma 2. The second regret term $R_{1}$ satisfies $R_{1} \leq c_{2} \mathbb{E}\left[\left(M_{\boldsymbol{\theta}^{*}}^{T}\right)^{r_{*}^{p}}\right]+c_{3}$, where $c_{2}=J^{*} 2^{r_{*}^{p}} s_{*}^{p}\left(\beta_{*}^{p}\right)^{-1}$ and $c_{3}=J^{*}\left(\beta_{*}^{p}\right)^{-1}\left(s_{*}^{p}(2 h)^{r_{*}^{p}}+b_{*}^{p}\left(K_{*}\right)^{-1}\right)$.

From Lemma 6, $\mathbb{E}\left[\left(M_{\theta^{*}}^{T}\right)^{r_{*}^{p}}\right]$ is $O\left(\log ^{r_{*}^{p}} T\right)$; hence, $R_{1}$ is upper bounded by a polylogarithmic function of the order $r_{*}^{p}$. Finally, in Lemma 3, we derive an upper bound for the third regret term $R_{2}$ defined in (12) using the bound derived for the relative value function in (13). To bound $R_{2}$, we characterize it in terms of the difference between the empirical and true unknown transition kernel and following the concentration method used in [64, 10, 49, 7, 8], we argue that with high probability the total variation distance between the two distributions is small; for proof, see Appendix B. 6 .
Lemma 3. For problem-dependent constant $c_{p_{3}}$ and polynomial $Q(T)=c_{p_{3}}(T h)^{r+r_{*}^{p}} / 48$, we have

$$
R_{2} \leq(\log (h T+h)+1)^{d}+c_{p_{3}} \sqrt{|\mathcal{A}| T} \log _{2}\left(2|\mathcal{A}| T^{2} Q(T)\right) \mathbb{E}\left[\left(M_{\boldsymbol{\theta}^{*}}^{T}+h\right)^{d+r+r_{*}^{p}}\left(\max _{1 \leq i \leq T} \tau_{0^{d}}^{(i)}\right)\right]
$$

The above Lemma results in a $\tilde{O}\left(K r d J^{*} h^{d+2 r+r_{*}^{p}} \sqrt{|\mathcal{A}| T}\right)$ regret term as a result of Lemma 6 where $h$ is the skip-free parameter defined in Assumption 2, $d$ is the dimension of the state-space, $K$ and $r$ are the cost function parameters defined in Assumption 1, $J^{*}$ is the supremum on the optimal cost, $r_{*}^{p}$ is defined in Assumption 4 and where $\tilde{O}$ hides logarithmic factors in problem parameters one of which is $\log ^{d+r+r_{*}^{p}+2}(T)$. For simplicity, we have not included the Lyapunov
functions related parameters in the regret. Finally, from Lemmas 1, 2, 3, along with the CauchySchwarz inequality, we conclude that the regret of Algorithm $1 R\left(T, \pi_{T S D E}\right)\left(=R_{0}+R_{1}+R_{2}\right)$ is $\tilde{O}\left(K r d J^{*} h^{d+2 r+r_{*}^{p}} \sqrt{|\mathcal{A}| T}\right)$; for brevity, we will state that regret is of the order $\tilde{O}\left(d h^{d} \sqrt{|\mathcal{A}| T}\right)$.
Theorem 1. Under Assumptions 1.5 the regret of Algorithm $1 \quad R\left(T, \pi_{T S D E}\right)$, is $\tilde{O}\left(d h^{d} \sqrt{|\mathcal{A}| T}\right)$.
Theorem 1 can be extended to the problem of finding the best policy within a sub-class of policies in set $\Pi$, which may or may not contain the optimal policy. In Section 2 , we stated that Assumptions 3 and 4 hold for policies in $\Pi$ and we used this to argue that the Poisson equation has a solution given in (6). As a result, repeating the same arguments as in Theorem 1 with the modification that $\pi_{\theta}^{*}$ is the best in-class policy of the MDP governed by parameter $\theta$, yields the following corollary.
Corollary 1. Under Assumptions 1 through $\sqrt{5}$ the regret of Algorithm 1 when using the best in-class policy is $\tilde{O}\left(d h^{d} \sqrt{|\mathcal{A}| T}\right)$.

Requirement of an optimal policy oracle. To implement our algorithm, we need to find the optimal policy for each model sampled by the algorithm—optimal policy for Theorem 1 and optimal policy within policy class $\Pi$ for Corollary 1 . In the finite state-space setting, [49] provides a schedule of $\epsilon$ values and selects $\epsilon$-optimal policies to obtain $\tilde{O}(\sqrt{T})$ regret guarantees. The issue with extending the analysis of [49] to the countable state-space setting is that we need to ensure (uniform) ergodicity for the chosen $\epsilon$-optimal policies. Another issue is that, to the best of our knowledge, there isn't a general structural characterization of all $\epsilon$-optimal stationary policies for countable state-space MDPs or even a characterization of the policy within this set that is selected by any computational procedure in the literature; current results only discuss characterization of the stationary optimal policy. In the absence of such results, stability assumptions with the same uniformity across models as in our submission will be needed, which are likely too strong to be useful. However, if we could verify the stability requirements of Assumptions 3 and 4 for a subset of policies, the optimal oracle is not needed, and instead, by choosing approximately optimal policies within this subset, we can follow the same proof steps as [49] to guarantee regret performance similar to Corollary 1 (without knowledge of model parameters). Thus, in Theorem 2 we extend the previous regret guarantees to the algorithm employing $\epsilon$-optimal policy; proof is given in Appendix B. 8
Theorem 2. Consider a non-negative sequence $\left\{\epsilon_{k}\right\}_{k=1}^{\infty}$ such that for every $k \in \mathbb{N}, \epsilon_{k}$ is bounded above by $\frac{1}{k+1}$ and an $\epsilon_{k}$-optimal policy satisfying Assumptions 3 and 4 is given. The regret incurred by Algorithm $\left[1\right.$ while using the $\epsilon_{k}$-optimal policy during any episode $k$ is $\tilde{O}\left(d h^{d} \sqrt{|\mathcal{A}| T}\right)$.

## 5 Evaluation and Conclusion: Application of Algorithm 1 to queueing models

Next, we present an evaluation of our algorithm. We study two different queueing models shown in Figure 2, each with Poisson arrivals at rate $\lambda$, and two heterogeneous servers with exponentially distributed services times with unknown service rate vector $\boldsymbol{\theta}^{*}=\left(\theta_{1}^{*}, \theta_{2}^{*}\right)$. Vector $\boldsymbol{\theta}^{*}$ is sampled from the prior distribution $\nu$ defined on the space $\Theta$ given as $\Theta=\left\{\left(\theta_{1}, \theta_{2}\right) \in \mathbb{R}_{+}^{2}: \frac{\lambda}{\theta_{1}+\theta_{2}} \leq \frac{1-\delta}{1+\delta}, 1 \leq\right.$ $\left.\frac{\theta_{1}}{\theta_{2}} \leq R\right\}$, for fixed $R \geq 1$ and $\delta \in(0,0.5)$. The first condition ensures the stability of the queueing models, while the second guarantees the compactness of the parameter space of the parameterized policies. In both systems, the goal of the dispatcher is to minimize the expected sojourn time of jobs, which by Little's law [52] is equivalent to minimizing the average number of jobs in the system. After verifying Assumptions 1 | 5 in Appendix Efor the cost function $c(\boldsymbol{x})=\|\boldsymbol{x}\|_{1}$, Theorem 1 yields a Bayesian regret of order $O(\sqrt{|\mathcal{A}| T})$ for Algorithm 1 .

Model 1. Two-server queueing system with a common buffer. We consider the continuous-time queueing system of Figure 2a, where the countable state space is $\mathcal{X}=\left\{\boldsymbol{x}=\left(x_{0}, x_{1}, x_{2}\right) \in\right.$ $\left.\mathbb{Z}_{+} \times\{0,1\}^{2}\right\}$, where $x_{0}$ is the queue length, and $x_{i}, i=1,2$ equal 1 if server $i$ is busy. The action space is $\mathcal{A}=\{h, b, 1,2\}$, where $h$ means no action, $b$ sends a job to both servers, and $i=1,2$ assigns a job to server $i$. In [38], it is shown that by uniformization [39] and sampling the continuous-time Markov process at rate $\lambda+\theta_{1}^{*}+\theta_{2}^{*}$, a discrete-time Markov chain is obtained, which converts the original continuous-time problem to an equivalent discrete-time problem where we need to minimize $\limsup _{T \rightarrow \infty} T^{-1} \sum_{t=0}^{T-1}\|\boldsymbol{X}(t)\|_{1}$. Further, [38] shows that the optimal policy is a threshold policy $\pi_{t\left(\boldsymbol{\theta}^{*}\right)}$ with optimal finite threshold $t\left(\boldsymbol{\theta}^{*}\right) \in \mathbb{N}$ : always assign a job to the faster (first) server when free, and to the second server if it is free and $\|\boldsymbol{x}\|_{1}>t\left(\boldsymbol{\theta}^{*}\right)$, and take no action otherwise. In


Figure 2: Two-server queueing systems with heterogeneous service rates.


Figure 3: Regret performance for $\lambda=0.3,0.5,0.7$. Shaded region shows the $\pm \sigma$ area of mean regret.

Appendix E.1, we argue that the discrete-time Markov process governed by $\theta \in \Theta$ and following threshold policy $\pi_{t}$ for any threshold $t$ belonging to a compact set satisfies Assumptions 11.5
Model 2. Two heterogeneous parallel queues. We consider the continuous-time queueing system of Figure 2b with countable state space $\mathcal{X}=\left\{\boldsymbol{x}=\left(x_{1}, x_{2}\right) \in \mathbb{Z}_{+}^{2}\right\}$, where $x_{i}$ is the number of jobs in the server-queue pair $i$. The action space is $\mathcal{A}=\{1,2\}$, where action $i$ sends the arrival to queue $i$. We obtain the discrete-time MDP by sampling the queueing system at the arrivals, and then aim to find the average cost minimizing policy within the class $\Pi=\left\{\pi_{\omega} ; \omega \in\left[\left(c_{R} R\right)^{-1}, c_{R} R\right]\right\}$, $c_{R} \geq 1$. Policy $\pi_{\omega}: \mathcal{X} \rightarrow \mathcal{A}$ routes arrivals based on the weighted queue lengths: $\pi_{\omega}(\boldsymbol{x})=$ $\arg \min \left(1+x_{1}, \omega\left(1+x_{2}\right)\right)$ with ties broken for 1 . Even with the transition kernel fully specified (by the values of arrival and service rates), the optimal policy in $\Pi$ is not known except when $\theta_{1}=\theta_{2}$ where the optimal value is $\omega=1$, and so, to learn it, we will use Proximal Policy Optimization for countable state-space MDPs [18]. Note that [18] requires full model knowledge, which holds in our scheme as we use parameters sampled from the posterior for choosing the policy at the beginning of each episode. In Appendix E.2, we argue that the discrete-time Markov process governed by parameter $\theta \in \Theta$ and following policy $\pi_{\omega}$ for $\omega \in\left[\left(c_{R} R\right)^{-1}, c_{R} R\right]$ satisfies Assumptions $1 \mid 5$
Next, we report the numerical results of Algorithm 1 in the two queueing models of Figure 2 and calculate regret using (2). The regret is averaged over 2000 simulation runs and plotted against the number of transitions in the sampled discrete-time Markov process. Figure 3 shows the behavior of the regret of the two queuing models for three different arrival rates and service rates distributed according to a Dirichlet prior over $[0.5,1.9]^{2}$. We observe that the regret is sub-linear in time and grows as the arrival rate increases. For the queueing model of Figure 2a the minimum average cost $J(\theta)$ and optimal policy $\pi_{\theta}^{*}$ are known explicitly [38] for every $\theta \in \Theta$, which are used in Algorithm 1 and for regret calculation. Conversely, for the second queueing model, $J(\theta)$ and $\pi_{\theta}^{*}$ are not known. The PPO algorithm [18] is used to empirically find both the optimal weight and the policy's average cost. Additional details of the simulations and more plots are presented in Appendix $G$
Conclusions and future work. We studied the problem of learning optimal policies in countable state-space MDPs governed by unknown parameters. We proposed a learning policy based on Thompson sampling and established finite-time performance guarantees on the Bayesian regret. We highlighted the practicality of our proposed algorithm by considering two different queuing models and showing that our algorithm can be applied to develop optimal control policies. For future work we plan two directions to explore: to generalize our algorithm to consider polices that might not all be stabilizing, and also to simplify the algorithm using ideas from [61, 57].

## Disclosure of Funding

SA's research was supported by NSF via grants ECCS2038416, CCF2008130, and CNS1955777, and a grant from General Dynamics via MIDAS at the University of Michigan, Ann Arbor. VS's research is supported in part by NSF via grants ECCS2038416, CCF2008130, CNS1955777, and CMMI2240981.

## References

[1] Yasin Abbasi-Yadkori and Csaba Szepesvari. Bayesian optimal control of smoothly parameterized systems: The lazy posterior sampling algorithm. arXiv preprint arXiv:1406.3926, 2014.
[2] Saghar Adler, Mehrdad Moharrami, and Vijay Subramanian. Learning a discrete set of optimal allocation rules in queueing systems with unknown service rates. arXiv preprint arXiv:2202.02419, 2022.
[3] R. Agrawal, D. Teneketzis, and V. Anantharam. Asymptotically efficient adaptive allocation schemes for controlled Markov chains: Finite parameter space. IEEE Transactions on Automatic Control, 34(12):1249-1259, 1989.
[4] Rajeev Agrawal and Demosthenis Teneketzis. Certainty equivalence control with forcing: Revisited. Systems \& Control Letters, 13(5):405-412, 1989.
[5] Shipra Agrawal and Randy Jia. Optimistic posterior sampling for reinforcement learning: Worst-case regret bounds. Advances in Neural Information Processing Systems, 30, 2017.
[6] Shipra Agrawal and Randy Jia. Learning in structured MDPs with convex cost functions: Improved regret bounds for inventory management. In Proceedings of the 2019 ACM Conference on Economics and Computation, pages 743-744, 2019.
[7] Nima Akbarzadeh and Aditya Mahajan. On learning Whittle index policy for restless bandits with scalable regret. arXiv preprint arXiv:2202.03463, 2022.
[8] Nima Akbarzadeh and Aditya Mahajan. On learning Whittle index policy for restless bandits with scalable regret. IEEE Transactions on Control of Network Systems, 2023.
[9] Aristotle Arapostathis, Vivek S Borkar, Emmanuel Fernández-Gaucherand, Mrinal K Ghosh, and Steven I Marcus. Discrete-time controlled Markov processes with average cost criterion: A survey. SIAM Journal on Control and Optimization, 31(2):282-344, 1993.
[10] Peter Auer, Thomas Jaksch, and Ronald Ortner. Near-optimal regret bounds for reinforcement learning. Advances in Neural Information Processing Systems, 21, 2008.
[11] V S Borkar. The Kumar-Becker-Lin scheme revisited. Journal of Optimization Theory and Applications, 66:289-309, 1990.
[12] Rolando Cavazos-Cadena. Necessary conditions for the optimality equation in average-reward Markov decision processes. Applied Mathematics and Optimization, 19(1):97-112, 1989.
[13] Rolando Cavazos-Cadena. Weak conditions for the existence of optimal stationary policies in average Markov decision chains with unbounded costs. Kybernetika, 25(3):145-156, 1989.
[14] Tuhinangshu Choudhury, Gauri Joshi, Weina Wang, and Sanjay Shakkottai. Job dispatching policies for queueing systems with unknown service rates. In Proceedings of the Twenty-Second International Symposium on Theory, Algorithmic Foundations, and Protocol Design for Mobile Networks and Mobile Computing, MobiHoc '21, page 181-190, New York, NY, USA, 2021. Association for Computing Machinery.
[15] Sayak Ray Chowdhury, Aditya Gopalan, and Odalric-Ambrym Maillard. Reinforcement learning in parametric MDPs with exponential families. In International Conference on Artificial Intelligence and Statistics, pages 1855-1863. PMLR, 2021.
[16] Asaf Cohen, Vijay Subramanian, and Yili Zhang. Learning-based optimal admission control in a single-server queuing system. Stochastic Systems, 2024.
[17] Asaf Cohen, Vijay G. Subramanian, and Yili Zhang. Learning-based optimal admission control in a single server queuing system. https://arxiv.org/abs/2212.11316, 2022.
[18] Jim G Dai and Mark Gluzman. Queueing network controls via deep reinforcement learning. Stochastic Systems, 12(1):30-67, 2022.
[19] Anthony Ephremides, Pravin Varaiya, and Jean Walrand. A simple dynamic routing problem. IEEE Transactions on Automatic Control, 25(4):690-693, 1980.
[20] Lloyd Fisher and Sheldon M Ross. An example in denumerable decision processes. The Annals of Mathematical Statistics, 39(2):674-675, 1968.
[21] Daniel Freund, Thodoris Lykouris, and Wentao Weng. Efficient decentralized multi-agent learning in asymmetric queuing systems. In Conference on Learning Theory, pages 4080-4084. PMLR, 2022.
[22] Mohammad Ghavamzadeh, Shie Mannor, Joelle Pineau, and Aviv Tamar. Bayesian reinforcement learning: A survey. Foundations and Trends® in Machine Learning, 8(5-6):359-483, 2015.
[23] Aditya Gopalan and Shie Mannor. Thompson sampling for learning parameterized Markov decision processes. In Conference on Learning Theory, pages 861-898. PMLR, 2015.
[24] Todd L Graves and Tze-Leung Lai. Asymptotically efficient adaptive choice of control laws incontrolled Markov chains. SIAM Journal on Control and Optimization, 35(3):715-743, 1997.
[25] Bruce Hajek. Hitting-time and occupation-time bounds implied by drift analysis with applications. Advances in Applied Probability, 14(3):502-525, 1982.
[26] Bruce Hajek. Optimal control of two interacting service stations. IEEE Transactions on Automatic Control, 29(6):491-499, 1984.
[27] Arie Hordijk and Flora Spieksma. On ergodicity and recurrence properties of a Markov chain by an application to an open Jackson network. Advances in Applied Probability, 24(2):343-376, 1992.
[28] Mehdi Jafarnia Jahromi, Rahul Jain, and Ashutosh Nayyar. Online learning for unknown partially observable MDPs. In International Conference on Artificial Intelligence and Statistics, pages 1712-1732. PMLR, 2022.
[29] Soren F Jarner and Gareth O Roberts. Polynomial convergence rates of Markov chains. The Annals of Applied Probability, 12(1):224-247, 2002.
[30] Subhashini Krishnasamy, PT Akhil, Ari Arapostathis, Rajesh Sundaresan, and Sanjay Shakkottai. Augmenting Max-Weight with explicit learning for wireless scheduling with switching costs. IEEE/ACM Transactions on Networking, 26(6):2501-2514, 2018.
[31] Subhashini Krishnasamy, Ari Arapostathis, Ramesh Johari, and Sanjay Shakkottai. On learning the $\mathrm{c} \mu$ rule in single and parallel server networks. https://arxiv.org/abs/1802.06723, 2018.
[32] Subhashini Krishnasamy, Rajat Sen, Ramesh Johari, and Sanjay Shakkottai. Learning unknown service rates in queues: A multiarmed bandit approach. Operations Research, 69(1):315-330, 2021.
[33] P R Kumar and A Becker. A new family of optimal adaptive controllers for Markov chains. IEEE Transactions on Automatic Control, 27(1):137-146, 1982.
[34] P R Kumar and Woei Lin. Optimal adaptive controllers for unknown Markov chains. IEEE Transactions on Automatic Control, 27(4):765-774, 1982.
[35] P R Kumar and Pravin Varaiya. Stochastic systems: Estimation, identification, and adaptive control. SIAM, 2015.
[36] Tze-Leung Lai and Sidney Yakowitz. Machine learning and nonparametric bandit theory. IEEE Transactions on Automatic Control, 40(7):1199-1209, 1995.
[37] Ronald Larsen. Control of multiple exponential servers with application to computer systems. PhD thesis, University of Maryland, 1981.
[38] Woei Lin and P R Kumar. Optimal control of a queueing system with two heterogeneous servers. IEEE Transactions on Automatic Control, 29(8):696-703, 1984.
[39] Steven A Lippman. Applying a new device in the optimization of exponential queuing systems. Operations Research, 23(4):687-710, 1975.
[40] Ashok P. Maitra. Dynamic programming for countable state systems. PhD thesis, University of California, Berkeley, 1963.
[41] Armand M Makowski and Adam Shwartz. The Poisson equation for countable Markov chains: Probabilistic methods and interpretations. Handbook of Markov Decision Processes: Methods and Applications, pages 269-303, 2002.
[42] Akshay Mete, Rahul Singh, Xi Liu, and P R Kumar. Reward biased maximum likelihood estimation for reinforcement learning. In Learning for Dynamics and Control, pages 815-827. PMLR, 2021.
[43] Sean P Meyn and Richard L Tweedie. Markov chains and stochastic stability. Springer Science \& Business Media, 2012.
[44] Michael J Neely, Scott T Rager, and Thomas F La Porta. Max-Weight learning algorithms for scheduling in unknown environments. IEEE Transactions on Automatic Control, 57(5):11791191, 2012.
[45] Pedro A Ortega and Daniel A Braun. A minimum relative entropy principle for learning and acting. Journal of Artificial Intelligence Research, 38:475-511, 2010.
[46] Ian Osband, Daniel Russo, and Benjamin Van Roy. (More) Efficient reinforcement learning via posterior sampling. Advances in Neural Information Processing Systems, 26, 2013.
[47] Ian Osband and Benjamin Van Roy. Why is posterior sampling better than optimism for reinforcement learning? In International Conference on Machine Learning, pages 2701-2710. PMLR, 2017.
[48] Reda Ouhamma, Debabrota Basu, and Odalric Maillard. Bilinear exponential family of MDPs: Frequentist regret bound with tractable exploration and planning. In Proceedings of the AAAI Conference on Artificial Intelligence, volume 37, pages 9336-9344, 2023.
[49] Yi Ouyang, Mukul Gagrani, Ashutosh Nayyar, and Rahul Jain. Learning unknown Markov decision processes: A Thompson sampling approach. Advances in Neural Information Processing Systems, 30, 2017.
[50] Daniel J Russo, Benjamin Van Roy, Abbas Kazerouni, Ian Osband, and Zheng Wen. A tutorial on Thompson sampling. Foundations and Trends ${ }^{\circledR}$ in Machine Learning, 11(1):1-96, 2018.
[51] Devavrat Shah, Qiaomin Xie, and Zhi Xu. Stable reinforcement learning with unbounded state space. arXiv preprint arXiv:2006.04353, 2020.
[52] R. Srikant and Lei Ying. Communication networks: An optimization, control, and stochastic networks perspective. Cambridge University Press, 2013.
[53] Thomas Stahlbuhk, Brooke Shrader, and Eytan Modiano. Learning algorithms for minimizing queue length regret. IEEE Transactions on Information Theory, 67(3):1759-1781, 2021.
[54] Shaler Stidham and Richard Weber. A survey of Markov decision models for control of networks of queues. Queueing Systems, 13:291-314, 1993.
[55] Malcolm Strens. A Bayesian framework for reinforcement learning. In ICML, volume 2000, pages 943-950, 2000.
[56] Wojciech Szpankowski and Vernon Rego. Yet another application of a binomial recurrence order statistics. Computing, 43(4):401-410, 1990.
[57] Dengwang Tang, Rahul Jain, Botao Hao, and Zheng Wen. Efficient online learning with offline datasets for infinite horizon MDPs: A Bayesian approach. arXiv preprint arXiv:2310.11531, 2023.
[58] Leandros Tassiulas and Anthony Ephremides. Jointly optimal routing and scheduling in packet ratio networks. IEEE Transactions on Information Theory, 38(1):165-168, 1992.
[59] Leandros Tassiulas and Anthony Ephremides. Dynamic server allocation to parallel queues with randomly varying connectivity. IEEE Transactions on Information Theory, 39(2):466-478, 1993.
[60] Georgios Theocharous, Zheng Wen, Yasin Abbasi-Yadkori, and Nikos Vlassis. Posterior sampling for large scale reinforcement learning. arXiv preprint arXiv:1711.07979, 2017.
[61] Georgios Theocharous, Zheng Wen, Yasin Abbasi Yadkori, and Nikos Vlassis. Scalar posterior sampling with applications. Advances in Neural Information Processing Systems, 31, 2018.
[62] William R Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. Biometrika, 25(3-4):285-294, 1933.
[63] Neil Walton and Kuang Xu. Learning and information in stochastic networks and queues. In Tutorials in Operations Research: Emerging Optimization Methods and Modeling Techniques with Applications, pages 161-198. INFORMS, 2021.
[64] Tsachy Weissman, Erik Ordentlich, Gadiel Seroussi, Sergio Verdu, and Marcelo J Weinberger. Inequalities for the $L_{1}$ deviation of the empirical distribution. Hewlett-Packard Labs, Tech. Rep, 2003.
[65] Zixian Yang, R Srikant, and Lei Ying. Learning while scheduling in multi-server systems with unknown statistics: Maxweight with discounted UCB. In International Conference on Artificial Intelligence and Statistics, pages 4275-4312. PMLR, 2023.
[66] Andrea Zanette, David Brandfonbrener, Emma Brunskill, Matteo Pirotta, and Alessandro Lazaric. Frequentist regret bounds for randomized least-squares value iteration. In International Conference on Artificial Intelligence and Statistics, pages 1954-1964. PMLR, 2020.

## A Proofs related to problem formulation

## A. 1 Ergodicity definitions

Suppose that Markov process $\boldsymbol{X}$ on $\mathcal{X}$ with transition kernel $P$ is irreducible, aperiodic and positive recurrent with stationary distribution $\mu$ and let $f: \mathcal{X} \mapsto[1, \infty)$ be a measurable function such that $\mu(f):=\mathbb{E}_{\mu}[f(Y)]<+\infty$ with $Y \sim \mu$. We are interested in conditions under which for a sequence of positive numbers $\rho:=(\rho(n))_{n \geq 0}$,

$$
\begin{equation*}
\lim _{n \rightarrow \infty} \rho(n)\left\|P^{n}(\boldsymbol{x}, \cdot)-\mu(\cdot)\right\|_{f}=0, \quad \forall \boldsymbol{x} \in \mathcal{X} \tag{15}
\end{equation*}
$$

where for a signed measure $\tilde{\mu}$ on $\mathcal{X},\|\tilde{\mu}\|_{f}:=\sup _{|g| \leq f}|\tilde{\mu}(g)|$. The sequence $\rho$ is interpreted as the rate function, and three different notions of ergodicity are distinguished based on the following rate functions: $\rho(n) \equiv 1, \rho(n)=\zeta^{n}$ for $\zeta>1$, and $\rho(n)=n^{\zeta-1}$ for $\zeta \geq 1$. Further, for each rate function $\rho$, we state the Foster-Lyapunov characterization of ergodicity of the Markov process $\boldsymbol{X}$, which provides sufficient conditions for (15) to hold.

1. If $\rho(n) \equiv 1$ for all $n \geq 0$, the Markov process $\boldsymbol{X}$ satisfying (15) is said to be $f$-ergodic. From [43], for an irreducible and aperiodic chain, $f$-ergodicity is equivalent to the existence of a function $V: \mathcal{X} \mapsto[0, \infty)$, a finite set $C$, and positive constant $b$ such that

$$
\begin{equation*}
\Delta V \leq-f+b \mathbb{I}_{C} \tag{16}
\end{equation*}
$$

where $\Delta V:=P V-V$ with $P V(\boldsymbol{x}):=\sum_{\boldsymbol{x}^{\prime} \in \mathcal{X}} P\left(\boldsymbol{x}, \boldsymbol{x}^{\prime}\right) V\left(\boldsymbol{x}^{\prime}\right)$. The drift condition (16) implies positive recurrence of the Markov process, existence of a unique stationary distribution $\mu$, and $\mu(f) \leq b<+\infty$ ([43], Theorem 14.3.7).
2. If $\rho(n)=\zeta^{n}$ for some $\zeta>1$, the Markov process $\boldsymbol{X}$ satisfying 15 is said to be $f$ geometrically ergodic. From [43], for an irreducible and aperiodic chain, $f$-geometric ergodicity is equivalent to the existence of a function $V: \mathcal{X} \mapsto[1, \infty)$, a finite set $C$, a constant $\gamma \in(0,1)$ and positive constant $b$ such that

$$
\begin{equation*}
\Delta V \leq-(1-\gamma) V+b \mathbb{I}_{C} \tag{17}
\end{equation*}
$$

The drift condition 17 implies positive recurrence of the Markov process, existence of a unique stationary distribution $\mu$, and $\mu(V) \leq \frac{b}{1-\gamma}<+\infty$ ([43], Theorem 14.3.7). Moreover, if $f(\cdot) \equiv 1$ in (15), then the Markov process $\boldsymbol{X}$ is called geometrically ergodic.
3. If $\rho(n)=n^{\zeta-1}$ for some $\zeta \geq 1$, the Markov process $\boldsymbol{X}$ satisfying (15) is said to be $f$ polynomially ergodic. From [43, 29], for an irreducible and aperiodic chain, the existence of a function $V: \mathcal{X} \mapsto[1, \infty)$, a finite set $C$, a constant $\alpha \in[0,1)$, and positive constants $c$ and $b$ such that

$$
\begin{equation*}
\Delta V \leq-c V^{\alpha}+b \mathbb{I}_{C} \tag{18}
\end{equation*}
$$

implies $V_{\zeta}$-polynomial ergodicity of $\boldsymbol{X}$ at rate $\rho(n)=n^{\zeta-1}$ for all $\zeta \in[1,1 /(1-\alpha)]$ with $V_{\zeta}=V^{1-\zeta(1-\alpha)}$. The drift condition (18) implies positive recurrence of the Markov process, existence of a unique stationary distribution $\mu$, and $\mu\left(V^{\alpha}\right) \leq \frac{b}{c}<+\infty$.

## A. 2 Lemma 4

Lemma 4. For any state $\boldsymbol{x} \neq 0^{d}$, there exists constants $\kappa>1$ and $c_{1}$ such that the following holds for the hitting time of state $0^{d}, \tau_{0^{d}}$,

$$
\mathbb{E}_{\boldsymbol{x}}\left[\kappa^{\tau_{0^{d}}}\right] \leq c_{1} V^{g}(\boldsymbol{x}) .
$$

Proof. We define $\tilde{V}:=\sum_{n=0}^{\infty}{ }_{0^{d}} P^{n} V^{g}$ where ${ }_{0^{d}} P^{n}$ is the $n$-step taboo probability [43] defined as

$$
{ }_{A} P_{\boldsymbol{x} B}^{n}=\mathbb{P}_{\boldsymbol{x}}\left(\boldsymbol{X}_{n} \in B, \tau_{A}>n\right),
$$

for $A, B \subseteq \mathcal{X}$, and $\tau_{A}$ is the first hitting time of set $A$. We also let ${ }_{A} P_{\boldsymbol{x} B}^{0}=\mathbb{I}_{B}(\boldsymbol{x})$. We have

$$
\begin{aligned}
{ }_{0^{d}} P \tilde{V}(\boldsymbol{x}) & =\sum_{\boldsymbol{y} \neq 0^{d}} P_{\boldsymbol{x} \boldsymbol{y}} \tilde{V}(\boldsymbol{y})=\sum_{n=0}^{\infty} \sum_{\boldsymbol{y}, \boldsymbol{z} \neq 0^{d}} P_{\boldsymbol{x} \boldsymbol{y}{ }_{0^{d}}} P_{\boldsymbol{y} \boldsymbol{z}}^{n} V^{g}(\boldsymbol{z}) \\
& =\sum_{n=0}^{\infty} \sum_{\boldsymbol{z} \neq 0^{d}}{ }_{0}{ }^{d} P_{\boldsymbol{x} \boldsymbol{z}}^{n+1} V^{g}(\boldsymbol{z})=\tilde{V}(\boldsymbol{x})-V^{g}(\boldsymbol{x}) .
\end{aligned}
$$

In Appendix D.3, we argue that there exists $\tilde{b}^{g}>1$ such that $\tilde{V}(\boldsymbol{y}) \leq \tilde{b}^{g} V^{g}(\boldsymbol{y})$ for all $\boldsymbol{y} \in \mathcal{X}$, which leads to

$$
\begin{equation*}
{ }_{0^{d}} P \tilde{V}=\tilde{V}-V^{g} \leq \tilde{V}-\frac{1}{\tilde{b}^{g}} \tilde{V}=\left(1-\frac{1}{\tilde{b}^{g}}\right) \tilde{V} \tag{19}
\end{equation*}
$$

Define Lyapunov function

$$
\tilde{V}^{g}(\boldsymbol{x})= \begin{cases}\left(1+2 \tilde{b}^{g}\right) \tilde{V}(\boldsymbol{x}), & \text { if } \boldsymbol{x} \neq 0^{d} \\ 1+\left(2 \tilde{b}^{g}\right)^{-1}, & \text { if } \boldsymbol{x}=0^{d}\end{cases}
$$

From the above equation and (19), we get

$$
\begin{aligned}
P \tilde{V}^{g}(\boldsymbol{x}) & =\sum_{\boldsymbol{y} \neq 0^{d}} P_{\boldsymbol{x} \boldsymbol{y}} \tilde{V}^{g}(\boldsymbol{y})+P_{\boldsymbol{x} 0^{d}} \tilde{V}^{g}\left(0^{d}\right) \\
& =\sum_{\boldsymbol{y} \neq 0^{d}} P_{\boldsymbol{x} \boldsymbol{y}}\left(1+2 \tilde{b}^{g}\right) \tilde{V}(\boldsymbol{y})+P_{\boldsymbol{x} 0^{d}}\left(1+\frac{1}{2 \tilde{b}^{g}}\right) \\
& \leq\left(1-\frac{1}{\tilde{b}^{g}}\right)\left(1+2 \tilde{b}^{g}\right) \tilde{V}(\boldsymbol{x})+1+\frac{1}{2 \tilde{b}^{g}} \\
& \leq\left(1-\frac{1}{\tilde{b}^{g}}\right)\left(1+2 \tilde{b}^{g}\right) \tilde{V}(\boldsymbol{x})+\left(1+\frac{1}{2 \tilde{b}^{g}}\right) \tilde{V}(\boldsymbol{x}) \\
& =\left(1-\frac{1}{2 \tilde{b}^{g}}\right)\left(1+2 \tilde{b}^{g}\right) \tilde{V}(\boldsymbol{x}) .
\end{aligned}
$$

Thus,

$$
P \tilde{V}^{g}(\boldsymbol{x}) \leq\left(1-\frac{1}{2 \tilde{b}^{g}}\right) \tilde{V}^{g}(\boldsymbol{x})+\left(1-\frac{1}{2 \tilde{b}^{g}}\right)\left(1+2 \tilde{b}^{g}\right) \tilde{V}\left(0^{d}\right) \mathbb{I}_{0^{d}}(\boldsymbol{x}), \quad \boldsymbol{x} \in \mathcal{X}
$$

To find an upper bound for $\mathbb{E}_{\boldsymbol{x}}\left[\kappa^{\tau_{0} d}\right]$, we apply [43, Theorem 15.2.5], which is a generalization of Lemma 12 For any $1 \leq \kappa \leq \frac{2 \tilde{b}^{g}}{2 \tilde{b}^{g}-1}$, there exists $\epsilon>0$ such that

$$
\mathbb{E}_{\boldsymbol{x}}\left[\sum_{i=0}^{\tau_{0}-1} \tilde{V}^{g}\left(\boldsymbol{X}_{i}\right) \kappa^{i}\right] \leq \epsilon^{-1} \kappa^{-1} \tilde{V}^{g}(\boldsymbol{x})
$$

As $\tilde{V}^{g}(\boldsymbol{y}) \geq 1$ for all $\boldsymbol{y} \in \mathcal{X}$, we have

$$
\begin{aligned}
\mathbb{E}_{\boldsymbol{x}}\left[\kappa^{\tau_{0 d}}\right] & \leq \kappa \mathbb{E}_{\boldsymbol{x}}\left[\sum_{i=0}^{\tau_{0 d}-1} \tilde{V}^{g}\left(\boldsymbol{X}_{i}\right) \kappa^{i}\right] \leq \epsilon^{-1} \tilde{V}^{g}(\boldsymbol{x}) \\
& =\epsilon^{-1}\left(1+2 \tilde{b}^{g}\right) \tilde{V}(\boldsymbol{x}) \leq \tilde{b}^{g} \epsilon^{-1}\left(1+2 \tilde{b}^{g}\right) V^{g}(\boldsymbol{x})
\end{aligned}
$$

and the claim holds for any $\kappa \in\left[1, \frac{2 \tilde{b}^{g}}{2 \tilde{b}^{g}-1}\right]$ and $c_{1}=\tilde{b}^{g} \epsilon^{-1}\left(1+2 \tilde{b}^{g}\right)$.

## A. 3 Poisson equation

For an irreducible Markov process on the countably-infinite space $\mathcal{X}$ with time-homogeneous transition kernel $P$ and cost function $\bar{c}(\cdot)$, a solution pair to the Poisson equation [41] is a scalar $J$ and function $v(\cdot): \mathcal{X} \mapsto \mathbb{R}$ such that $J+v=\bar{c}+P v$, where $v(\boldsymbol{z})=0$ for some $\boldsymbol{z} \in \mathcal{X}$. If the Markov process is also positive recurrent and $\mathbb{E}_{\boldsymbol{x}}\left[\sum_{i=0}^{\tau_{\boldsymbol{y}}-1}|\bar{c}(\boldsymbol{X}(i))|\right]<\infty$, where $\tau_{\boldsymbol{y}}$ is the first hitting time of some state $\boldsymbol{y} \in \mathcal{X}$, then solution pair $(J, v)$ given as

$$
J=\frac{\mathbb{E}_{\boldsymbol{y}}\left[\sum_{i=0}^{\tau_{y}-1}|\bar{c}(\boldsymbol{X}(i))|\right]}{\mathbb{E}_{\boldsymbol{y}}\left[\tau_{\boldsymbol{y}}\right]} \text { and } v(\boldsymbol{x})=\mathbb{E}_{\boldsymbol{y}}\left[\sum_{i=0}^{\tau_{\boldsymbol{x}}-1}|\bar{c}(\boldsymbol{X}(i))|\right]-J \mathbb{E}_{\boldsymbol{x}}\left[\tau_{\boldsymbol{y}}\right], \quad \forall \boldsymbol{x} \in \mathcal{X}
$$

is a solution to the Poisson equation $J+v=\bar{c}+P v$ with $v(\boldsymbol{z})=0$ [41, Theorem 9.5].

Lemma 5. Consider Markov Decision Processes ( $\left.\mathcal{X}, \mathcal{A}, c, P_{\theta}\right)$ governed by parameter $\theta \in \Theta$ following the best-in-class policy $\pi_{\theta}^{*}$. Then the pair $\left.(J)(\theta), v^{\pi_{\theta}^{*}}\right)$ given as

$$
J(\theta):=\frac{\bar{C}^{\pi_{\theta}^{*}}\left(0^{d}\right)}{\mathbb{E}_{0^{d}}^{\pi_{\theta}^{*}}\left[\tau_{0^{d}}\right]} \text { and } v^{\pi_{\theta}^{*}}(\boldsymbol{x})=\bar{C}^{\pi_{\theta}^{*}}(\boldsymbol{x})-J(\theta) \mathbb{E}_{x}^{\pi_{\theta}^{*}}\left[\tau_{0^{d}}\right], \quad \forall \boldsymbol{x} \in \mathcal{X},
$$

is a solution to the Poisson equation $v+J=c+P_{\theta}^{\pi_{\theta}^{*}} v$, where $v^{\pi_{\theta}^{*}}\left(0^{d}\right)=0$ and $\bar{C}^{\pi_{\theta}^{*}}(\boldsymbol{x})=$ $\mathbb{E}_{\boldsymbol{x}}^{\pi_{\boldsymbol{\theta}}^{*}}\left[\sum_{i=0}^{\tau_{\sigma_{d}}-1} c\left(\boldsymbol{X}(i), \pi_{\theta}^{*}(\boldsymbol{X}(i))\right)\right]$.

Proof. From [41 Theorem 9.5], a solution pair to the Poisson equation exists if $\mathbb{E}_{\boldsymbol{x}}^{\pi_{\theta}^{*}}\left[\tau_{0^{d}}\right]$ and $\bar{C}^{\pi_{\theta}^{*}}(\boldsymbol{x})$ are finite for all $\boldsymbol{x} \in \mathcal{X}$. The former follows from positive recurrence assumed in Assumption 3 and for the latter, from Assumptions 1 and 2 ,

$$
\begin{aligned}
\bar{C}^{\pi_{\theta}^{*}}(\boldsymbol{x}) & =\mathbb{E}_{\boldsymbol{x}}^{\pi_{\theta}^{*}}\left[\sum_{i=0}^{\tau_{0} d-1} c\left(\boldsymbol{X}(i), \pi_{\theta}^{*}(\boldsymbol{X}(i))\right)\right] \leq \mathbb{E}_{\boldsymbol{x}}^{\pi_{\boldsymbol{\pi}}^{*}}\left[\sum_{i=0}^{\tau_{0} d-1} \sum_{j=1}^{d} K\left(X_{j}(i)\right)^{r}\right] \\
& \leq \mathbb{E}_{\boldsymbol{x}}^{\pi_{\theta}^{*}}\left[\sum_{i=0}^{\tau_{0}-1} K d\left(\|\boldsymbol{x}\|_{\infty}+h i\right)^{r}\right] \leq \mathbb{E}_{\boldsymbol{x}}^{\pi_{\boldsymbol{x}}^{*}}\left[K d\left(\|\boldsymbol{x}\|_{\infty}+h \tau_{0^{d}}\right)^{r} \tau_{0^{d} d},\right.
\end{aligned}
$$

which is finite from geometric ergodicity (Assumption 3) and the discussion following that.

## B Proofs of regret analysis

In this section, we state the proofs related to regret analysis of Section 4 We first note a key property of Thompson sampling from [49], which states that for any episode $k$, measurable function $f$, and $\mathcal{H}_{t_{k}}$-measurable random variable $Y$, we have

$$
\begin{equation*}
\mathbb{E}\left[f\left(\theta_{k}, Y\right)\right]=\mathbb{E}\left[f\left(\boldsymbol{\theta}^{*}, Y\right)\right], \tag{20}
\end{equation*}
$$

where $\mathcal{H}_{t}:=\sigma(\boldsymbol{X}(1), \ldots, \boldsymbol{X}(t), A(1), \ldots, A(t-1))$ for all $t \in \mathbb{N}$. We start with deriving upper bounds on the hitting times of state $0^{d}$ using the ergodicity conditions of Assumptions 3 and 4 Previous works [25, 27, 29] have already established bounds on hitting times in geometrically and polynomially ergodic chains in terms of their corresponding Lyapunov function. However, our objective is to provide a precise characterization of all constants included in these bounds in terms of the constants of the drift equations 3 and 4 . This characterization allows us to derive uniform bounds across the model class. In Appendix C.1 , using the polynomial Lyapunov function provided in Assumption 4 , we establish upper bounds on the $i$-th moment of hitting time of state $0^{d}$ from any state $\boldsymbol{x} \in \mathcal{X}$ and for $1 \leq i \leq r+1$. Importantly, the derived bound is polynomial in terms of any component of the state $x_{i}$. Additionally, in Appendix C.2. we characterize the tail probabilities of the return time to state $0^{d}$ starting from $0^{d}$ in terms of the geometric Lyapunov function of Assumption 3 The derived tail bounds will be used in Lemma6 6 to derive upper bounds for all moments of hitting times in the model class. These bounds, along with the skip-free behavior of the model, allow us to study the maximum state (with respect to $\ell_{\infty}$-norm) achieved up to time $T$ in $\operatorname{MDP}\left(\mathcal{X}, \mathcal{A}, c, P_{\boldsymbol{\theta}^{*}}\right)$ following Algorithm 1 as follows.
Lemma 6. For $p \in \mathbb{N}$, the $p$-th moment of $\max _{1 \leq i \leq T} \tau_{0^{d}}^{(i)}$ and $M_{\theta^{*}}^{T}$, that is the maximum $\ell_{\infty}$-norm of the state vector achieved up until and including time $T$ is $O\left(\log ^{p} T\right)$.

In the proof of Lemma 6 given in Appendix B.1 , we make use of geometric ergodicity of the chain and the fact that hitting times have geometric tails to find an upper bound for moments of $M_{\theta^{*}}^{T}$. Using this, we aim to bound the number of episodes started before or at $T$, denoted by $K_{T}$. We first find an upper bound for the number of episodes in which the second stopping criterion is met or there exists a state-action pair for which $N_{t}(\boldsymbol{x}, a)$ has increased more than twice. In the following lemma, we bound the number of such episodes, which we denote by $K_{M}$, in terms of random variable $M_{\theta^{*}}^{T}$ and other problem-dependent constants. Proof of Lemma 7 is given in Appendix B. 2.
Lemma 7. The number of episodes triggered by the second stopping criterion and started before or at time $T$, denoted by $K_{M}$, satisfies $K_{M} \leq 2|\mathcal{A}|\left(M_{\theta^{*}}^{T}+1\right)^{d} \log _{2} T$ a.s.

We next bound the total number of episodes $K_{T}$ by bounding the number of episodes triggered by the first stopping criterion, using the fact that in such episodes, $\tilde{T}_{k}=\tilde{T}_{k-1}+1$. Moreover, to address the settling time of each episode $k$, shown by $E_{k}=T_{k}-\tilde{T}_{k}$, we use the geometric ergodicity property and Lemma6. Finally, the proof of Lemma 8 is given in Appendix B. 3
Lemma 8. The number of episodes started by $T$ satisfies $K_{T} \leq 2 \sqrt{|\mathcal{A}|\left(M_{\boldsymbol{\theta}^{*}}^{T}+1\right)^{d} T \log _{2} T}$ a.s.
From Lemma 8 , the upper bound given in Lemma 6 for moments of $M_{\boldsymbol{\theta}^{*}}^{T}$, and Cauchy-Schwarz inequality, it follows that the expected value of the number of episodes $K_{T}$ is of the order $\tilde{O}\left(\sqrt{h^{d}|\mathcal{A}| T}\right)$. This term has a crucial role in determining the overall order of the total regret up to time $T$. In the rest of this section, we present a detailed proof of the lemmas and other results used to prove Theorem 1
Remark 3. The skip-free to the right property in Assumption 2 yields a polynomially-sized subset of the underlying state-space that can be explored as a function of T. This polynomially-sized subset can be viewed as the effective finite-size of the system in the worst-case, and then, directly applying finite-state problem bounds [49] would result in a regret of order $\tilde{O}\left(T^{d+0.5}\right)$; since $d \geq 1$, such a coarse bound is not helpful even for asserting asymptotic optimality! However, to achieve a regret of $\tilde{O}(\sqrt{T})$, it is essential to carefully understand and characterize the distribution of $M_{\boldsymbol{\theta}^{*}}^{T}$ and then its moments, as demonstrated in Lemma 6
Remark 4. The derived regret bound can be extended to a larger class of MDPs which consist of transient states in addition to the single irreducible class. Specifically, for any $\theta_{1}, \theta_{2} \in \Theta$, the Markov process with transition kernel $P_{\theta_{1}}^{\pi_{\theta_{2}}^{*}}$ obtained from the $\operatorname{MDP}\left(\mathcal{X}, \mathcal{A}, c, P_{\theta_{1}}\right)$ by following policy $\pi_{\theta_{2}}^{*}$ has a single irreducible class $I_{\theta_{1}, \theta_{2}}$ and a set of transient states $T_{\theta_{1}, \theta_{2}}$. Furthermore, Assumptions 3 and 4 hold for the single irreducible class. The reasoning behind the proof remains true in this case using the following argument: each episode $k$ starts at $0^{d}$ which is in the irreducible set for the chosen policy $\pi_{\theta_{k}}^{*}$, hence, throughout the episode the algorithm remains in the irreducible set that is positive recurrent and never visits any transient states. In other words, episodes starting and ending at $0^{d}$ with a fixed episode dependent policy implies that reachable set of $0^{d}$ is all that can be explored, which is positive recurrent by our assumptions. As a result, we can restrict our proof derivations to the subset that is reachable from $0^{d}$ in each episode and follow the same analysis. The Lyapunov function based bounds apply to the positive recurrent states, and hence, restricting attention to states reachable from $0^{d}$ within each episode, we can use these bounds for our assessment of regret using norms of the state. Thereafter, the coarse bounds on the norms of the state can be applied as carried out in our proof.
Remark 5. By problem-dependent parameters, we refer to the parameters that characterize the complexity or size of the model class $\Theta$. These parameters are not just a function of the size of the statespace and diameter of the MDP (as mentioned in the literature on finite-size problems [5] 23, 49]), as stability needs to be accounted for in the countable state-space setting. The dependence is, thus, more complex and requires the inclusion of stability parameters, such as Lyapunov functions, petite sets, and ergodicity coefficients that are discussed in Assumptions 1.4
Remark 6. In the subsequent sections, several equalities and inequalities in the proofs are between random variables and hold almost surely (a.s.). Throughout the remainder, we will omit the explicit mention of a.s., but any such statement should be interpreted in this context.

## B. 1 Proof of Lemma 6

Proof. Let $\left\{\alpha_{i}\right\}_{i \geq 0}$ be the sequence of hitting times of state $0^{d}$ starting from $0^{d}$ (set $\alpha_{0}=0$ ). Define $\tau_{0^{d}}^{(i)}$ as the length of the $i$-th recurrence time of state $0^{d}$ for $i \in \mathbb{N}$, i.e., $\tau_{0^{d}}^{(i)}=\alpha_{i}-\alpha_{i-1}$. For simplicity, we take $\tau_{0^{d}}=\tau_{0^{d}}^{(1)}$. Each such recurrence time is generated using policy $\pi_{\theta_{i}}^{*}$ that is determined using the algorithm in operation in an MDP governed by parameter $\boldsymbol{\theta}^{*}$. Furthermore, $\left\{\tau_{0^{d}}^{(i)}\right\}_{i \in \mathbb{N}}$ are independent with length at least 1 , but they need not be identically distributed. The time $T$ can be in the middle of one of these recurrence times, hence the current recurrence interval count is $N(T)=\inf \left\{n: \sum_{i=1}^{n} \tau_{0^{d}}^{(i)} \geq T\right\}$. Note that the lower bound of 1 on every $\tau_{0^{d}}^{(i)}$ says that $N(T) \leq T$ a.s. Further, from the skip-free to the right property, the most any component of state can increase in during recurrence time $\tau_{0^{d}}^{(i)}$ is $h \tau_{0^{d}}^{(i)}$. Hence, the most any component of the state (and also the $\|\cdot\|_{\infty}$ norm of the state) can increase is given by $h \max _{i=1, \ldots, T} \tau_{0^{d}}^{(i)}$ where the random variables
are independent with geometrically decaying tails with a worst case rate of

$$
\sup _{\theta_{1}, \theta_{2} \in \Theta} \tilde{\gamma}_{\theta_{1}, \theta_{2}}^{g}=1-\left(\sup _{\theta_{1}, \theta_{2} \in \Theta} \tilde{b}_{\theta_{1}, \theta_{2}}^{g}\right)^{-1}
$$

see Lemma 11 From Lemma 10, we have

$$
\begin{align*}
\tilde{b}_{\theta_{1}, \theta_{2}}^{g} & =\frac{3 b_{\theta_{1}, \theta_{2}}^{g}+1}{1-\gamma_{\theta_{1}, \theta_{2}}^{g}}\left(\left|C_{\theta_{1}, \theta_{2}}^{g}\right|^{2} \max \left(1, \underset{\boldsymbol{u} \in C_{\theta_{1}, \theta_{2}}^{g} \backslash\left\{0^{d}\right\}}{ } \mathbb{E}_{\boldsymbol{u}}^{\pi_{\theta_{2}}^{*}}\left[\tau_{\left.0^{d}\right]}\right]\right)\right) \\
& \leq \frac{3 b_{*}^{g}+1}{1-\gamma_{*}^{g}}\left(\left|C_{*}^{g}\right|^{2} \max \left(1, \sup _{\substack{\boldsymbol{u} \in C_{1}^{g} \backslash\left\{0^{d}\right\} \\
\theta_{1}, \theta_{2} \in \Theta}} \phi_{\theta_{1}, \theta_{2}}^{p}(1)\left(V_{\theta_{1}, \theta_{2}}^{p}(\boldsymbol{u})+b_{\theta_{1}, \theta_{2}}^{p} \alpha_{C_{\theta_{1}, \theta_{2}}^{p}}^{p}\right)\right)\right) \\
& \leq \frac{3 b_{*}^{g}+1}{1-\gamma_{*}^{g}}\left(\left|C_{*}^{g}\right|^{2} \max \left(1, \sup _{\substack{\boldsymbol{u} \in C_{1}^{g} \backslash\left\{0^{d}\right\} \\
\theta_{1}, \theta_{2} \in \Theta}} \frac{1}{\beta_{\theta_{1}, \theta_{2}}^{p}}\left(s_{\theta_{1}, \theta_{2}}^{p}\|\boldsymbol{u}\|_{\infty}^{r_{\theta_{1}, \theta_{2}}^{p}}+\frac{b_{\theta_{1}, \theta_{2}}^{p}}{\min _{\boldsymbol{y} \in C_{\theta_{1}, \theta_{2}}^{p}}^{p} K_{\theta_{1}, \theta_{2}}(\boldsymbol{y})}\right)\right)\right) \\
& \leq \frac{3 b_{*}^{g}+1}{1-\gamma_{*}^{g}}\left(\left|C_{*}^{g}\right|^{2} \max \left(1, \sup _{\boldsymbol{u} \in C_{*}^{g} \backslash\left\{0^{d}\right\}} \frac{1}{\beta_{*}^{p}}\left(s_{*}^{p}\|\boldsymbol{u}\|_{\infty}^{r_{*}^{p}}+\frac{b_{*}^{p}}{K_{*}}\right)\right)\right)  \tag{21}\\
& :=\tilde{b}_{*}^{g},
\end{align*}
$$

and we define $\tilde{\gamma}_{*}^{g}:=1-\left(\tilde{b}_{*}^{g}\right)^{-1}$. From the definition of $b_{\theta_{1}, \theta_{2}}^{g}$ in Assumption 3, $b_{\theta_{1}, \theta_{2}}^{g}$ is greater than or equal to 2 . Thus, $\tilde{b}_{\theta_{1}, \theta_{2}}^{g} \geq 7$ and we have

$$
\sup _{\theta_{1}, \theta_{2} \in \Theta} c_{\theta_{1}, \theta_{2}}^{g}=\sup _{\theta_{1}, \theta_{2} \in \Theta} \frac{b_{\theta_{1}, \theta_{2}}^{g}\left(\tilde{b}_{\theta_{1}, \theta_{2}}^{g}\right)^{2}}{\tilde{b}_{\theta_{1}, \theta_{2}}^{g}-1} \leq \frac{b_{*}^{g}\left(\tilde{b}_{*}^{g}\right)^{2}}{6}:=c_{*}^{g},
$$

and as a result of Lemma 11 ,

$$
\begin{equation*}
\mathbb{P}_{0^{d}}\left(\tau_{0^{d}}^{(i)}>n\right) \leq c_{*}^{g}\left(\gamma_{*}^{g}\right)^{n}, \quad 1 \leq i \leq T \tag{22}
\end{equation*}
$$

We upper bound $\mathbb{E}\left[M_{\boldsymbol{\theta}^{*}}^{T}\right]$ using the independence of $\left\{\tau_{0^{d}}^{(i)}\right\}_{i \in \mathbb{N}}$ and the above equation,

$$
\begin{aligned}
\mathbb{E}\left[M_{\boldsymbol{\theta}^{*}}^{T}\right] & \leq h \mathbb{E}\left[\max _{1 \leq i \leq T} \tau_{0^{d}}^{(i)}\right]=h \sum_{n=0}^{\infty} \mathbb{P}\left(\max _{1 \leq i \leq T} \tau_{0^{d}}^{(i)}>n\right) \\
& =h \sum_{n=0}^{\infty}\left(1-\mathbb{P}\left(\max _{1 \leq i \leq T} \tau_{0^{d}}^{(i)} \leq n\right)\right)=h \sum_{n=0}^{\infty}\left(1-\prod_{i=1}^{T} \mathbb{P}\left(\tau_{0^{d}}^{(i)} \leq n\right)\right) \\
& \leq h n_{0}+h \sum_{n=n_{0}}^{\infty} 1-\left(1-c_{*}^{g}\left(\gamma_{*}^{g}\right)^{n_{0}}\left(\gamma_{*}^{g}\right)^{n-n_{0}}\right)^{T} \\
& \leq h\left(n_{0}+1\right)+h \sum_{n=n_{0}+1}^{\infty} 1-\left(1-\left(\gamma_{*}^{g}\right)^{n-n_{0}}\right)^{T},
\end{aligned}
$$

where $n_{0}$ is the smallest $n \geq 0$ such that $c_{*}^{g}\left(\gamma_{*}^{g}\right)^{n}<1$. By Reimann sum approximation, we get

$$
\begin{aligned}
\mathbb{E}\left[M_{\boldsymbol{\theta}^{*}}^{T}\right] & \leq h\left(n_{0}+1\right)+h \sum_{n=1}^{\infty} 1-\left(1-\left(\gamma_{*}^{g}\right)^{n}\right)^{T} \\
& <h\left(n_{0}+1\right)+h \int_{0}^{\infty} 1-\left(1-\left(\gamma_{*}^{g}\right)^{u}\right)^{T} d u \\
& =h\left(n_{0}+1\right)+\frac{h}{\log \gamma_{*}^{g}} \int_{0}^{1} \frac{1-u^{T}}{1-u} d u \\
& \leq h\left(n_{0}+1\right)+\frac{h}{\log \gamma_{*}^{g}}(\log T+1),
\end{aligned}
$$

where the last inequality follows from $\sum_{n=1}^{T} n^{-1} \leq \log T+1$ and thus $\mathbb{E}\left[M_{\boldsymbol{\theta}^{*}}^{T}\right]$ is $O(h \log T)$. We now extend the result to moments of order greater than one. From (22), for $1 \leq i \leq T$,

$$
\mathbb{P}_{0^{d}}\left(\tau_{0^{d}}^{(i)}>n\right) \leq c_{*}^{g}\left(\gamma_{*}^{g}\right)^{n}=c_{*}^{g}\left(\gamma_{*}^{g}\right)^{n_{0}}\left(\gamma_{*}^{g}\right)^{n-n_{0}}<\left(\gamma_{*}^{g}\right)^{n-n_{0}}
$$

For $n \geq n_{0}$, let $t=n-n_{0} \geq 0$ and $Y_{i}=\max \left(\tau_{0^{d}}^{(i)}-n_{0}, 0\right)$ to get

$$
\mathbb{P}_{0^{d}}\left(Y_{i}>t\right)=\mathbb{P}_{0^{d}}\left(\tau_{0^{d}}^{(i)}-n_{0}>t\right)<\left(\gamma_{*}^{g}\right)^{t}
$$

which means random variables $\left\{Y_{i}\right\}_{i=1}^{T}$ are stochastically dominated by independent and identically distributed geometric random variables with parameter $1-\gamma_{*}^{g}$. Furthermore, [56] argues that the $p$-th moment of the maximum of $T$ independent and identically distributed geometric random variables is $O\left(\log ^{p} T\right)$. Thus, the $p$-th moment of $\max _{1 \leq i \leq T} Y_{i}$ is $O\left(\log ^{p} T\right)$ and

$$
\begin{aligned}
\max _{1 \leq i \leq T} Y_{i} & =\max \left(\tau_{0^{d}}^{(1)}-n_{0}, \ldots, \tau_{0^{d}}^{(T)}-n_{0}, 0\right)=\max \left(\tau_{0^{d}}^{(1)}, \ldots, \tau_{0^{d}}^{(T)}, n_{0}\right)-n_{0} \\
& \geq \max \left(\tau_{0^{d}}^{(1)}, \ldots, \tau_{0^{d}}^{(T)}\right)-n_{0} \geq h^{-1} M_{\boldsymbol{\theta}^{*}}^{T}-n_{0}
\end{aligned}
$$

which gives

$$
\mathbb{E}\left[\left(M_{\boldsymbol{\theta}^{*}}^{T}\right)^{p}\right] \leq h^{p} \mathbb{E}\left[\left(\max _{1 \leq i \leq T} \tau_{0^{d}}^{(i)}\right)^{p}\right] \leq h^{p} \mathbb{E}\left[\left(\max _{1 \leq i \leq T} Y_{i}+n_{0}\right)^{p}\right]
$$

Since the right-hand side of the above equation is $O\left(h^{p} \log ^{p} T\right)$, the claim is proved.

## B. 2 Proof of Lemma 7

Proof. Let $K_{M}(\boldsymbol{x}, a)$ be the number of episodes $k$ such that $1 \leq k \leq K_{T}$ and in which the number of visits to the state-action pair $(\boldsymbol{x}, a)$ is increased more than twice at episode $k$, or

$$
K_{M}(\boldsymbol{x}, a)=\left|\left\{k \leq K_{T}: N_{\tilde{t}_{k+1}}(\boldsymbol{x}, a)>2 N_{t_{k}}(\boldsymbol{x}, a)\right\}\right|
$$

As for every episode in the above set the number of visits to $(x, a)$ doubles,

$$
K_{M}(\boldsymbol{x}, a) \leq \log _{2}\left(N_{T+1}(\boldsymbol{x}, a)\right)+1
$$

and we can upper bound $K_{M}$ as follows

$$
\begin{aligned}
K_{M} & =\sum_{\boldsymbol{x} \in \mathcal{X}, a \in \mathcal{A}} K_{M}(\boldsymbol{x}, a)=\sum_{\substack{\|\boldsymbol{x}\|_{\infty} \leq M_{\boldsymbol{\theta}^{*}}^{T} \\
a \in \mathcal{A}}} K_{M}(\boldsymbol{x}, a) \\
& \leq \sum_{\substack{\|\boldsymbol{x}\|_{\infty} \leq M_{\boldsymbol{\theta}^{*}}^{T} \\
a \in \mathcal{A}}}\left(1+\log _{2} N_{T+1}(\boldsymbol{x}, a)\right) \leq|\mathcal{A}|\left(M_{\boldsymbol{\theta}^{*}}^{T}+1\right)^{d}\left(1+\log _{2} T\right) .
\end{aligned}
$$

This completes the proof.

## B. 3 Proof of Lemma 8

Proof. We define macro episodes with start times $t_{n_{k}}, k=1,2, \ldots, K_{M}+1$ where $t_{n_{1}}=t_{1}$, $t_{n_{K_{M}+1}}=T+1$ (which is equivalent to $n_{K_{M}+1}=K_{T}+1$ ), and for $1<k<K_{M}+1$

$$
t_{n_{k+1}}=\min \left\{t_{j}>t_{n_{k}}: \quad N_{t_{j}}(\boldsymbol{x}, a)>2 N_{t_{j-1}}(\boldsymbol{x}, a) \text { for some }(\boldsymbol{x}, a)\right\}
$$

which are episodes wherein the second stopping criterion is triggered. Any episode (except for the last episode) in a macro episode must be triggered by the first stopping criterion; equivalently, $\tilde{T}_{j}=\tilde{T}_{j-1}+1$ for all $j=n_{k}, n_{k}+1, \ldots, n_{k+1}-2$. For $1 \leq k \leq K_{M}$, let $T_{k}^{M}=\sum_{j=n_{k}}^{n_{k+1}-1} T_{j}$ be the length of the $k$-th macro episode. We have

$$
T_{k}^{M}=\sum_{j=n_{k}}^{n_{k+1}-1} T_{j} \geq \sum_{j=n_{k}}^{n_{k+1}-1} \tilde{T}_{j} \geq 1+\sum_{j=n_{k}}^{n_{k+1}-2}\left(j-n_{k}+2\right)=0.5\left(n_{k+1}-n_{k}\right)\left(n_{k+1}-n_{k}+1\right)
$$

Consequently, $n_{k+1}-n_{k} \leq \sqrt{2 T_{k}^{M}}$ for all $1 \leq k \leq K_{M}$. From this, we obtain

$$
K_{T}=n_{K_{M}+1}-1=\sum_{k=1}^{K_{M}}\left(n_{k+1}-n_{k}\right) \leq \sum_{k=1}^{K_{M}} \sqrt{2 T_{k}^{M}} .
$$

Using the above equation and the fact that $\sum_{k=1}^{K_{M}} T_{k}^{M}=T$ we get

$$
K_{T} \leq \sum_{k=1}^{K_{M}} \sqrt{2 T_{k}^{M}} \leq \sqrt{K_{M} \sum_{k=1}^{K_{M}} 2 T_{k}^{M}}=\sqrt{2 K_{M} T}
$$

Finally, from Lemma 7 we get

$$
K_{T} \leq \sqrt{2 K_{M} T} \leq 2 \sqrt{|\mathcal{A}|\left(M_{\boldsymbol{\theta}^{*}}^{T}+1\right)^{d} T \log _{2} T}
$$

This completes the proof.

## B. 4 Proof of Lemma 1

Proof. Let $E_{k}=T_{k}-\tilde{T}_{k} \geq 0$ be the settling time needed to return to state $0^{d}$ after a stopping criterion is realized in episode $k$. We have

$$
\begin{align*}
R_{0} & =\mathbb{E}\left[\sum_{k=1}^{K_{T}} T_{k} J\left(\theta_{k}\right)\right]-T \mathbb{E}\left[J\left(\boldsymbol{\theta}^{*}\right)\right] \\
& =\mathbb{E}\left[\sum_{k=1}^{K_{T}} \tilde{T}_{k} J\left(\theta_{k}\right)\right]+\mathbb{E}\left[\sum_{k=1}^{K_{T}} E_{k} J\left(\theta_{k}\right)\right]-T \mathbb{E}\left[J\left(\boldsymbol{\theta}^{*}\right)\right] \tag{23}
\end{align*}
$$

We first simplify the first term in the above summation. From the monotone convergence theorem,

$$
\mathbb{E}\left[\sum_{k=1}^{K_{T}} \tilde{T}_{k} J\left(\theta_{k}\right)\right]=\sum_{k=1}^{\infty} \mathbb{E}\left[\mathbb{I}_{\left\{t_{k} \leq T\right\}} \tilde{T}_{k} J\left(\theta_{k}\right)\right]
$$

Note that the first stopping criterion of Algorithm 1 ensures that $\tilde{T}_{k} \leq \tilde{T}_{k-1}+1$ at all episodes $k \geq 1$. Hence

$$
\mathbb{E}\left[\mathbb{I}_{\left\{t_{k} \leq T\right\}} \tilde{T}_{k} J\left(\theta_{k}\right)\right] \leq \mathbb{E}\left[\mathbb{I}_{\left\{t_{k} \leq T\right\}}\left(\tilde{T}_{k-1}+1\right) J\left(\theta_{k}\right)\right]
$$

Since $\mathbb{I}_{\left\{t_{k} \leq T\right\}}\left(\tilde{T}_{k-1}+1\right)$ is measurable with respect to $\mathcal{H}_{t_{k}}$, by we get

$$
\mathbb{E}\left[\mathbb{I}_{\left\{t_{k} \leq T\right\}}\left(\tilde{T}_{k-1}+1\right) J\left(\theta_{k}\right)\right]=\mathbb{E}\left[\mathbb{I}_{\left\{t_{k} \leq T\right\}}\left(\tilde{T}_{k-1}+1\right) J\left(\boldsymbol{\theta}^{*}\right)\right]
$$

Therefore,

$$
\mathbb{E}\left[\sum_{k=1}^{K_{T}} \tilde{T}_{k} J\left(\theta_{k}\right)\right] \leq \sum_{k=1}^{\infty} \mathbb{E}\left[\mathbb{I}_{\left\{t_{k} \leq T\right\}}\left(\tilde{T}_{k-1}+1\right) J\left(\boldsymbol{\theta}^{*}\right)\right]=\mathbb{E}\left[\sum_{k=1}^{K_{T}}\left(\tilde{T}_{k-1}+1\right) J\left(\boldsymbol{\theta}^{*}\right)\right]
$$

Thus,

$$
\begin{align*}
\mathbb{E}\left[\sum_{k=1}^{K_{T}} \tilde{T}_{k} J\left(\theta_{k}\right)\right]-T \mathbb{E}\left[J\left(\boldsymbol{\theta}^{*}\right)\right] & \leq \mathbb{E}\left[J\left(\boldsymbol{\theta}^{*}\right) \sum_{k=1}^{K_{T}}\left(\tilde{T}_{k-1}+1\right)\right]-\mathbb{E}\left[J\left(\boldsymbol{\theta}^{*}\right) \sum_{k=1}^{K_{T}} T_{k}\right] \\
& =\mathbb{E}\left[J\left(\boldsymbol{\theta}^{*}\right)\left(K_{T}+1-T_{K_{T}}-\sum_{k=1}^{K_{T}-1} E_{k}\right)\right] \\
& \leq \mathbb{E}\left[J\left(\boldsymbol{\theta}^{*}\right) K_{T}\right] . \tag{24}
\end{align*}
$$

For the second term in 23), from Assumption 5

$$
\begin{equation*}
\mathbb{E}\left[\sum_{k=1}^{K_{T}} E_{k} J\left(\theta_{k}\right)\right] \leq J^{*} \mathbb{E}\left[\sum_{k=1}^{K_{T}} E_{k}\right] \leq J^{*} \mathbb{E}\left[K_{T} \max _{1 \leq i \leq T} \tau_{0^{d}}^{(i)}\right] \tag{25}
\end{equation*}
$$

Substitutinh 24) and (25) in 23), we get

$$
\begin{aligned}
R_{0} & \leq \mathbb{E}\left[K_{T} J\left(\boldsymbol{\theta}^{*}\right)\right]+J^{*} \mathbb{E}\left[K_{T} \max _{1 \leq i \leq T} \tau_{0^{d}}^{(i)}\right] \\
& \leq J^{*} \mathbb{E}\left[K_{T}\right]+J^{*} \mathbb{E}\left[K_{T} \max _{1 \leq i \leq T} \tau_{0^{d}}^{(i)}\right] \\
& =J^{*} \mathbb{E}\left[K_{T}\left(\max _{1 \leq i \leq T} \tau_{0^{d}}^{(i)}+1\right)\right] .
\end{aligned}
$$

## B. 5 Proof of Lemma 2

Proof. We note that the state of the MDP is equal to $0^{d}$ at the beginning of all episodes and the relative value function $v(\boldsymbol{x} ; \theta)$ is equal to 0 at $\boldsymbol{x}=0^{d}$ for all $\theta$. Thus,

$$
\begin{aligned}
R_{1} & =\mathbb{E}\left[\sum_{k=1}^{K_{T}} \sum_{t=t_{k}}^{t_{k+1}-1}\left[v\left(\boldsymbol{X}(t) ; \theta_{k}\right)-v\left(\boldsymbol{X}(t+1) ; \theta_{k}\right)\right]\right] \\
& =\mathbb{E}\left[\sum_{k=1}^{K_{T}}\left[v\left(\boldsymbol{X}\left(t_{k}\right) ; \theta_{k}\right)-v\left(\boldsymbol{X}\left(t_{k+1}\right) ; \theta_{k}\right)\right]\right] \\
& =\mathbb{E}\left[\sum_{k=1}^{K_{T}-1}\left[v\left(0^{d} ; \theta_{k}\right)-v\left(0^{d} ; \theta_{k}\right)\right]+v\left(0^{d} ; \theta_{K_{T}}\right)-v\left(\boldsymbol{X}(T+1) ; \theta_{K_{T}}\right)\right] \\
& =-\mathbb{E}\left[v\left(\boldsymbol{X}(T+1) ; \theta_{K_{T}}\right)\right] .
\end{aligned}
$$

From the lower bound derived for the relative value function in 14 ,

$$
-v(\boldsymbol{x} ; \theta) \leq J^{*} \mathbb{E}_{\boldsymbol{x}}^{\pi_{\theta}^{*}}\left[\tau_{0^{d}}\right] \leq \frac{J^{*}}{\beta_{*}^{p}}\left(s_{*}^{p}\|\boldsymbol{x}\|_{\infty}^{r_{*}^{p}}+\frac{b_{*}^{p}}{K_{*}}\right),
$$

where the second inequality follows from 21 in the proof of Lemma6. We also note that $\| \boldsymbol{X}(T+$ 1) $\|_{\infty} \leq M_{\boldsymbol{\theta}^{*}}^{T}+h$. Thus,

$$
R_{1}=-\mathbb{E}\left[v\left(\boldsymbol{X}(T+1) ; \theta_{K_{T}}\right)\right] \leq \mathbb{E}\left[\frac{J^{*}}{\beta_{*}^{p}}\left(s_{*}^{p}\left(M_{\boldsymbol{\theta}^{*}}^{T}+h\right)^{r_{*}^{p}}+\frac{b_{*}^{p}}{K_{*}}\right)\right]
$$

From the inequality $(a+b)^{r} \leq 2^{r}\left(a^{r}+b^{r}\right)$, we have

$$
R_{1} \leq \frac{J^{*} 2^{r_{*}^{p}} s_{*}^{p}}{\beta_{*}^{p}} \mathbb{E}\left[\left(M_{\boldsymbol{\theta}^{*}}^{T}\right)^{r_{*}^{p}}\right]+\frac{J^{*}}{\beta_{*}^{p}}\left(s_{*}^{p}(2 h)^{r_{*}^{p}}+\frac{b_{*}^{p}}{K_{*}}\right) .
$$

## B. 6 Proof of Lemma 3

Proof. Let $\boldsymbol{Z}(t)=\left(\boldsymbol{X}(t), \pi_{\theta_{k}}^{*}(\boldsymbol{X}(t))\right)$ be the state-action pair at $t_{k} \leq t<t_{k+1} . R_{2}$ can be upper bounded as

$$
\begin{align*}
R_{2} & =\mathbb{E}\left[\sum_{k=1}^{K_{T}} \sum_{t=t_{k}}^{t_{k+1}^{-1}}\left[v\left(\boldsymbol{X}(t+1) ; \theta_{k}\right)-\sum_{\boldsymbol{y} \in \mathcal{X}} P_{\theta_{k}}\left(\boldsymbol{y} \mid \boldsymbol{X}(t), \pi_{\theta_{k}}^{*}(\boldsymbol{X}(t))\right) v\left(\boldsymbol{y} ; \theta_{k}\right)\right]\right] \\
& \leq \mathbb{E}\left[\sum_{k=1}^{K_{T}} \sum_{t=t_{k}}^{t_{k+1}-1}\left[\sum_{\boldsymbol{y} \in \mathcal{X}}\left|P_{\boldsymbol{\theta}^{*}}(\boldsymbol{y} \mid \boldsymbol{Z}(t))-P_{\theta_{k}}(\boldsymbol{y} \mid \boldsymbol{Z}(t))\right| v\left(\boldsymbol{y} ; \theta_{k}\right) \mid\right]\right] \\
& \leq \sum_{t=1}^{T} \mathbb{E}\left[\left(\max _{\substack{1 \leq k \leq K_{T} \\
\|\boldsymbol{x}\|_{\infty} \leq M_{\boldsymbol{\theta}^{*}}^{T}}}\left|v\left(\boldsymbol{x} ; \theta_{k}\right)\right|\right)\left\|P_{\boldsymbol{\theta}^{*}}(\cdot \mid \boldsymbol{Z}(t))-P_{\theta_{k}}(\cdot \mid \boldsymbol{Z}(t))\right\|_{1}\right] . \tag{26}
\end{align*}
$$

We have
$\left\|P_{\boldsymbol{\theta}^{*}}(\cdot \mid \boldsymbol{Z}(t))-P_{\theta_{k}}(\cdot \mid \boldsymbol{Z}(t))\right\|_{1} \leq\left\|P_{\boldsymbol{\theta}^{*}}(\cdot \mid \boldsymbol{Z}(t))-P_{\hat{\theta}_{k}}(\cdot \mid \boldsymbol{Z}(t))\right\|_{1}+\left\|P_{\theta_{k}}(\cdot \mid \boldsymbol{Z}(t))-P_{\hat{\theta}_{k}}(\cdot \mid \boldsymbol{Z}(t))\right\|_{1}$, where $P_{\hat{\theta}_{k}}(\boldsymbol{y} \mid \boldsymbol{Z}(t))$ is the empirical transition probability defined as

$$
P_{\hat{\theta}_{k}}(\boldsymbol{y} \mid \boldsymbol{Z}(t))=\frac{N_{t_{k}}(\boldsymbol{Z}(t), \boldsymbol{y})}{\max \left(1, N_{t_{k}}(\boldsymbol{Z}(t))\right)}
$$

and for any tuple $(\boldsymbol{x}, a, \boldsymbol{y})$, we define $N_{1}(\boldsymbol{x}, a, \boldsymbol{y})=0$ and for $t>1$,

$$
N_{t}(\boldsymbol{x}, a, \boldsymbol{y})=\mid\left\{t_{k} \leq i<\tilde{t}_{k+1} \leq t \text { for some } k \geq 1:(\boldsymbol{X}(i), A(i), \boldsymbol{X}(i+1))=(\boldsymbol{x}, a, \boldsymbol{y})\right\} \mid
$$

Thus, from (26) and defining random variable $v_{M}=\max { }_{1 \leq k \leq K_{T}}\left|v\left(\boldsymbol{x} ; \theta_{k}\right)\right|$,

$$
\begin{aligned}
& \|\boldsymbol{x}\|_{\infty} \leq M_{\boldsymbol{\theta}^{*}}^{T}
\end{aligned}
$$

$$
\begin{equation*}
R_{2} \leq \sum_{t=1}^{T} \mathbb{E}\left[v_{M}\left\|P_{\boldsymbol{\theta}^{*}}(\cdot \mid \boldsymbol{Z}(t))-P_{\hat{\theta}_{k}}(\cdot \mid \boldsymbol{Z}(t))\right\|_{1}\right]+\sum_{t=1}^{T} \mathbb{E}\left[v_{M}\left\|P_{\theta_{k}}(\cdot \mid \boldsymbol{Z}(t))-P_{\hat{\theta}_{k}}(\cdot \mid \boldsymbol{Z}(t))\right\|_{1}\right] \tag{27}
\end{equation*}
$$

We define set $B_{k}$ as the set of parameters $\theta$ for which the transition kernel $P_{\theta}(\cdot \mid \boldsymbol{z})$ is close to the empirical transition kernel $P_{\hat{\theta}_{k}}(\cdot \mid \boldsymbol{z})$ at episode $k$ for every state-action pair $\boldsymbol{z}=(\boldsymbol{x}, a) \in \mathcal{X} \times \mathcal{A}$, or

$$
B_{k}=\left\{\theta:\left\|P_{\theta}(\cdot \mid \boldsymbol{z})-P_{\hat{\theta}_{k}}(\cdot \mid \boldsymbol{z})\right\|_{1} \leq \beta_{k}(\boldsymbol{z}), \boldsymbol{z}=(\boldsymbol{x}, a) \in\{0,1, \cdots, h T\}^{d} \times \mathcal{A}\right\}
$$

where $\beta_{k}(\boldsymbol{z})=\sqrt{\frac{14 \prod_{i=1}^{d}\left(x_{i}+h\right)}{\max \left(1, N_{t_{k}}(\boldsymbol{z})\right)} \log \left(\frac{2|\mathcal{A}| T}{\tilde{\delta}}\right)}$ for $\boldsymbol{x}=\left(x_{1}, \ldots, x_{d}\right)$ and some $0<\tilde{\delta}<1$, which will be determined later. We simplify the $\ell_{1}$-difference of the real and empirical transition kernels as follows

$$
\begin{aligned}
& \left\|P_{\boldsymbol{\theta}^{*}}(\cdot \mid \boldsymbol{Z}(t))-P_{\hat{\theta}_{k}}(\cdot \mid \boldsymbol{Z}(t))\right\|_{1} \\
& =\mathbb{I}_{\left\{\boldsymbol{\theta}^{*} \notin B_{k}\right\}}\left\|P_{\theta_{*}}(\cdot \mid \boldsymbol{Z}(t))-P_{\hat{\theta}_{k}}(\cdot \mid \boldsymbol{Z}(t))\right\|_{1}+\mathbb{I}_{\left\{\boldsymbol{\theta}^{*} \in B_{k}\right\}}\left\|P_{\theta_{*}}(\cdot \mid \boldsymbol{Z}(t))-P_{\hat{\theta}_{k}}(\cdot \mid \boldsymbol{Z}(t))\right\|_{1} \\
& \leq 2 \mathbb{I}_{\left\{\boldsymbol{\theta}^{*} \notin B_{k}\right\}}+\beta_{k}(\boldsymbol{Z}(t)) .
\end{aligned}
$$

Similarly, we have

$$
\left\|P_{\theta_{k}}(\cdot \mid \boldsymbol{Z}(t))-P_{\hat{\theta}_{k}}(\cdot \mid \boldsymbol{Z}(t))\right\|_{1} \leq 2 \mathbb{I}_{\left\{\theta_{k} \notin B_{k}\right\}}+\beta_{k}(\boldsymbol{Z}(t))
$$

Substituting in 27, we get

$$
\begin{equation*}
R_{2} \leq \mathbb{E}\left[\sum_{k=1}^{K_{T}} \sum_{t=t_{k}}^{t_{k+1}-1} 2 v_{M}\left[\mathbb{I}_{\left\{\boldsymbol{\theta}^{*} \notin B_{k}\right\}}+\mathbb{I}_{\left\{\theta_{k} \notin B_{k}\right\}}\right]\right]+\mathbb{E}\left[\sum_{k=1}^{K_{T}} \sum_{t=t_{k}}^{t_{k+1}-1} 2 v_{M} \beta_{k}(\boldsymbol{Z}(t))\right] \tag{28}
\end{equation*}
$$

We first find an upper bound for $v_{M}=\max \underset{\substack{1 \leq k \leq K_{T} \\\|\boldsymbol{x}\|_{\infty} \leq M_{\boldsymbol{\theta}^{*}}^{T}}}{ }\left|v\left(\boldsymbol{x} ; \theta_{k}\right)\right|$ using the bounds derived in (13) and 14). From (13),

$$
\begin{align*}
v\left(\boldsymbol{x} ; \theta_{k}\right) & \leq \mathbb{E}_{\boldsymbol{x}}^{\pi_{\theta_{k}}^{*}}\left[K d\left(\|\boldsymbol{x}\|_{\infty}+h \tau_{0^{d}}\right)^{r} \tau_{0^{d}}\right] \\
& \leq \mathbb{E}_{\boldsymbol{x}}^{\pi_{\theta_{k}}^{*}}\left[2^{r} K d\left(\|\boldsymbol{x}\|_{\infty}^{r}+h^{r}\left(\tau_{0^{d}}\right)^{r}\right) \tau_{0^{d}}\right] \\
& =K d\left(2\|\boldsymbol{x}\|_{\infty}\right)^{r} \mathbb{E}_{\boldsymbol{x}}^{\pi_{\theta_{k}}^{*}}\left[\tau_{0^{d}}\right]+K d(2 h)^{r} \mathbb{E}_{\boldsymbol{x}}^{\pi_{\theta_{k}}^{*}}\left[\left(\tau_{0^{d}}\right)^{r+1}\right] \\
& \leq K d 2^{r}\left(\|\boldsymbol{x}\|_{\infty}^{r}+h^{r}\right) \mathbb{E}_{\boldsymbol{x}}^{\pi_{\theta_{k}}^{*}}\left[\left(\tau_{0^{d}}\right)^{r+1}\right] \\
& \leq K d(r+1) 2^{r}\left(\|\boldsymbol{x}\|_{\infty}^{r}+h^{r}\right) \phi_{\theta_{k}}^{p}(r+1)\left(V_{\theta_{k}}^{p}(\boldsymbol{x})+b_{\theta_{k}}^{p} \alpha_{C_{\theta_{k}}^{p}}^{p}\right) \\
& \leq K d(r+1) 2^{r}\left(\|\boldsymbol{x}\|_{\infty}^{r}+h^{r}\right) \phi_{\theta_{k}}^{p}(r+1)\left(s_{*}^{p}\|\boldsymbol{x}\|_{\infty}^{r_{*}^{p}}+b_{*}^{p}\left(K_{*}\right)^{-1}\right), \tag{29}
\end{align*}
$$

where the second line follows from the inequality $(a+b)^{r} \leq 2^{r}\left(a^{r}+b^{r}\right)$, the fifth line from Lemma 10, and the last line from Assumption 4 and 21 . We further have

$$
\begin{aligned}
\phi_{\theta_{1}, \theta_{2}}^{p}(r+1) & =\prod_{j=1}^{r+1} \frac{1}{\beta_{\theta_{1}, \theta_{2}}^{\eta_{j}}}\left(2^{j-1}+(j-1) \alpha_{C_{\theta_{1}, \theta_{2}}^{p}} b_{\theta_{1}, \theta_{2}}^{\eta_{j}}\right) \\
& \leq \prod_{j=1}^{r+1} \frac{r+1}{\min \left(1, \beta_{*}^{p}\right)}\left(2^{j-1}+(j-1)\left(K_{*}\right)^{-1} b_{\theta_{1}, \theta_{2}}^{\eta_{j}}\right)
\end{aligned}
$$

where using the definition of $b_{\theta_{1}, \theta_{2}}^{\eta_{j}}$ in 38,

$$
b_{\theta_{1}, \theta_{2}}^{\eta_{j}}=\left(b_{\theta_{1}, \theta_{2}}^{p}\right)^{\eta_{j}}+\eta_{j} \tilde{\beta}_{\theta_{1}, \theta_{2}}^{p} \max \left(1,\left(\tilde{\beta}_{\theta_{1}, \theta_{2}}^{p}\right)^{\left(\alpha_{\theta_{1}, \theta_{2}}^{p}+\eta_{j}-1\right) /\left(1-\alpha_{\theta_{1}, \theta_{2}}^{p}\right)}\right) \leq 1+b_{*}^{p}+\beta_{*}^{p}
$$

We also define

$$
\phi_{*}^{p}(r+1):=\prod_{j=1}^{r+1} \frac{r+1}{\min \left(1, \beta_{*}^{p}\right)}\left(2^{j-1}+(j-1)\left(K_{*}\right)^{-1}\left(1+b_{*}^{p}+\beta_{*}^{p}\right)\right) .
$$

We next find a lower bound for $v\left(\boldsymbol{x} ; \theta_{k}\right)$ using (14) as follows:

$$
v\left(\boldsymbol{x} ; \theta_{k}\right) \geq-J^{*} \mathbb{E}_{\boldsymbol{x}}^{\pi_{\theta_{k}}^{*}}\left[\tau_{0^{d}}\right] \geq-\frac{J^{*}}{\beta_{*}^{p}}\left(s_{*}^{p}\|\boldsymbol{x}\|_{\infty}^{r_{*}^{p}}+\frac{b_{*}^{p}}{K_{*}}\right)
$$

Combining (29) and the above equation, we get a uniform upper bound for $\left|v\left(\boldsymbol{x} ; \theta_{k}\right)\right|$ over $\Theta$, which we use to upper bound $v_{M}=\max \quad 1 \leq k \leq K_{T}\left|v\left(\boldsymbol{x} ; \theta_{k}\right)\right|$ as below

$$
\|\boldsymbol{x}\|_{\infty} \leq M_{\boldsymbol{\theta}^{*}}^{T}
$$

$$
\begin{align*}
v_{M} & \leq\left(J^{*}+K d(r+1) 2^{r}\right) \phi_{*}^{p}(r+1)\left(\left(M_{\boldsymbol{\theta}^{*}}^{T}\right)^{r}+h^{r}\right)\left(s_{*}^{p}\left(M_{\boldsymbol{\theta}^{*}}^{T}\right)^{r_{*}^{p}}+b_{*}^{p}\left(K_{*}\right)^{-1}\right) \\
& =c_{p_{1}}\left(\left(M_{\boldsymbol{\theta}^{*}}^{T}\right)^{r}+h^{r}\right)\left(s_{*}^{p}\left(M_{\boldsymbol{\theta}^{*}}^{T}\right)^{r_{*}^{p}}+b_{*}^{p}\left(K_{*}\right)^{-1}\right) \\
& \leq c_{p_{2}}\left(M_{\boldsymbol{\theta}^{*}}^{T}\right)^{r+r_{*}^{p}}, \tag{30}
\end{align*}
$$

where the constant terms are defined as

$$
c_{p_{1}}:=\left(J^{*}+K d(r+1) 2^{r}\right) \phi_{*}^{p}(r+1), \quad c_{p_{2}}:=\max \left(1, c_{p_{1}}\left(h^{r}+1\right)\left(s_{*}^{p}+b_{*}^{p}\left(K_{*}\right)^{-1}\right)\right) .
$$

A deterministic upper bound on $v_{M}$ can also be found from the above equation. Noting that from Assumption 2, until time $T$ only states with each component less than or equal to $h T$ are visited, we have

$$
v_{M} \leq c_{p_{2}}\left(M_{\boldsymbol{\theta}^{*}}^{T}\right)^{r+r_{*}^{p}} \leq c_{p_{2}}(T h)^{r+r_{*}^{p}}:=Q(T),
$$

where $Q(T)$ is a polynomial defined as above. Using the bounds derived for $v_{M}$, we bound $R_{2}$ starting with the first term on the right-hand side of (28). We have

$$
\begin{align*}
\mathbb{E}\left[\sum_{k=1}^{K_{T}} \sum_{t=t_{k}}^{t_{k+1}-1} 2 v_{M}\left[\mathbb{I}_{\left\{\boldsymbol{\theta}^{*} \notin B_{k}\right\}}+\mathbb{I}_{\left\{\theta_{k} \notin B_{k}\right\}}\right]\right] & \leq 2 Q(T) \mathbb{E}\left[\sum_{k=1}^{K_{T}} \sum_{t=t_{k}}^{t_{k+1}-1} \mathbb{I}_{\left\{\boldsymbol{\theta}^{*} \notin B_{k}\right\}}+\mathbb{I}_{\left\{\theta_{k} \notin B_{k}\right\}}\right] \\
& \leq 2 T Q(T) \mathbb{E}\left[\sum_{k=1}^{K_{T}} \mathbb{I}_{\left\{\boldsymbol{\theta}^{*} \notin B_{k}\right\}}+\mathbb{I}_{\left\{\theta_{k} \notin B_{k}\right\}}\right] \\
& \leq 2 T Q(T) \sum_{k=1}^{T} \mathbb{E}\left[\mathbb{I}_{\left\{\boldsymbol{\theta}^{*} \notin B_{k}\right\}}+\mathbb{I}_{\left\{\theta_{k} \notin B_{k}\right\}}\right] \\
& \leq 4 T Q(T) \sum_{k=1}^{T} \mathbb{P}\left\{\boldsymbol{\theta}^{*} \notin B_{k}\right\} \tag{31}
\end{align*}
$$

where the last inequality follows from (20) and the fact that set $B_{k}$ is $\mathcal{H}_{t_{k}}$-measurable. To further simplify the first term in (28), we find an upper bound for $\mathbb{P}\left\{\boldsymbol{\theta}^{*} \notin B_{k}\right\}$ using [64]. For a fixed $\boldsymbol{z}=(\boldsymbol{x}, a)$ and $n$ independent samples of the distribution $P_{\boldsymbol{\theta}^{*}}(. \mid \boldsymbol{z})$, the $L^{1}$-deviation of the true distribution $P_{\boldsymbol{\theta}^{*}}(. \mid \boldsymbol{z})$ and empirical distribution at the end of episode $k, P_{\hat{\theta}_{k}}(. \mid \boldsymbol{z})$, is bounded in [10] as

$$
\mathbb{P}\left\{\left\|P_{\boldsymbol{\theta}^{*}}(\cdot \mid \boldsymbol{z})-P_{\hat{\theta}_{k}}(\cdot \mid \boldsymbol{z})\right\|_{1} \geq \sqrt{\frac{14 \prod_{i=1}^{d}\left(x_{i}+h\right)}{n} \log \left(\frac{2|\mathcal{A}| T}{\tilde{\delta}}\right)}\right\} \leq \frac{\tilde{\delta}}{20|\mathcal{A}| T^{7} \prod_{i=1}^{d}\left(x_{i}+h\right)}
$$

Therefore,

$$
\mathbb{P}\left\{\left\|P_{\boldsymbol{\theta}^{*}}(\cdot \mid \boldsymbol{z})-P_{\hat{\theta}_{k}}(\cdot \mid \boldsymbol{z})\right\|_{1} \geq \beta_{k}(\boldsymbol{z}) \mid N_{t_{k}}(\boldsymbol{z})=n\right\} \leq \frac{\tilde{\delta}}{20|\mathcal{A}| T^{7} \prod_{i=1}^{d}\left(x_{i}+h\right)},
$$

and

$$
\begin{aligned}
& \mathbb{P}\left\{\left\|P_{\boldsymbol{\theta}^{*}}(\cdot \mid \boldsymbol{z})-P_{\hat{\theta}_{k}}(\cdot \mid \boldsymbol{z})\right\|_{1} \geq \beta_{k}(\boldsymbol{z})\right\} \\
& =\sum_{n=1}^{T} \mathbb{P}\left\{\left\|P_{\boldsymbol{\theta}^{*}}(\cdot \mid \boldsymbol{z})-P_{\hat{\theta}_{k}}(\cdot \mid \boldsymbol{z})\right\|_{1} \geq \beta_{k}(\boldsymbol{z}) \mid N_{t_{k}}(\boldsymbol{z})=n\right\} \mathbb{P}\left\{N_{t_{k}}(\boldsymbol{z})=n\right\} \\
& \leq \frac{\tilde{\delta}}{20|\mathcal{A}| T^{6} \prod_{i=1}^{d}\left(x_{i}+h\right)}
\end{aligned}
$$

The probability that at episode $k \leq T$, the true parameter $\boldsymbol{\theta}^{*}$ does not belong to the confidence set $B_{k}$ can be bounded using the above and union bound as

$$
\begin{aligned}
\mathbb{P}\left\{\boldsymbol{\theta}^{*} \notin B_{k}\right\} & \leq \sum_{\boldsymbol{z} \in\{0,1, \cdots, h T\}^{d} \times \mathcal{A}} \mathbb{P}\left\{\left\|P_{\boldsymbol{\theta}^{*}}(\cdot \mid \boldsymbol{z})-P_{\hat{\theta}_{k}}(\cdot \mid \boldsymbol{z})\right\|_{1} \geq \beta_{k}(\boldsymbol{z})\right\} \\
& \leq \sum_{\boldsymbol{z} \in\{0,1, \cdots, h T\}^{d} \times \mathcal{A}} \frac{\tilde{\delta}}{20|\mathcal{A}| T^{6} \prod_{i=1}^{d}\left(x_{i}+h\right)} \\
& =\sum_{\boldsymbol{x} \in\{0,1, \cdots, h T\}^{d}} \frac{\tilde{\delta}}{20 T^{6} \prod_{i=1}^{d}\left(x_{i}+h\right)} \\
& \leq \frac{\tilde{\delta}}{20 T^{6}}(\log (h(T+1))+1)^{d} \\
& \leq \frac{\tilde{\delta}}{20 k^{6}}(\log (h(T+1))+1)^{d}
\end{aligned}
$$

In the summation in the above equation, we have simplified the expression by summing over $x_{i} \leq h T$ instead of considering the more detailed summation over $x_{i} \leq M_{\boldsymbol{\theta}^{*}}^{T}$. However, this simplification does not affect the final evaluation of regret, as this term is not dominant and only contributes to a logarithmic term in the regret bound. Substituting in (31,

$$
\begin{align*}
\mathbb{E}\left[\sum_{k=1}^{K_{T}} \sum_{t=t_{k}}^{t_{k+1}-1} 2 v_{M}\left[\mathbb{I}_{\left\{\boldsymbol{\theta}^{*} \notin B_{k}\right\}}+\mathbb{I}_{\left\{\theta_{k} \notin B_{k}\right\}}\right]\right] & \leq 4 T Q(T) \sum_{k=1}^{T} \mathbb{P}\left\{\boldsymbol{\theta}^{*} \notin B_{k}\right\} \\
& \leq \frac{\tilde{\delta}(\log (h(T+1))+1)^{d} T Q(T)}{5} \sum_{k=1}^{\infty} \frac{1}{k^{6}} \\
& <\tilde{\delta}(\log (h(T+1))+1)^{d} T Q(T) . \tag{32}
\end{align*}
$$

We now upper bound the second term in (28). From (30),

$$
\begin{equation*}
\mathbb{E}\left[\sum_{k=1}^{K_{T}} \sum_{t=t_{k}}^{t_{k+1}-1} 2 v_{M} \beta_{k}(\boldsymbol{Z}(t))\right] \leq 2 c_{p_{2}} \mathbb{E}\left[\left(M_{\boldsymbol{\theta}^{*}}^{T}\right)^{r+r_{*}^{p}} \sum_{k=1}^{K_{T}} \sum_{t=t_{k}}^{t_{k+1}-1} \beta_{k}(\boldsymbol{Z}(t))\right] \tag{33}
\end{equation*}
$$

To bound the regret term resulting from the summation of $\beta_{k}(\boldsymbol{Z}(t))$, we note that from the second stopping criterion, $N_{t}(\boldsymbol{Z}(t)) \leq 2 N_{t_{k}}(\boldsymbol{Z}(t))$ for all $t_{k} \leq t<t_{k+1}$ and

$$
\begin{align*}
& \sum_{k=1}^{K_{T}} \sum_{t=t_{k}}^{t_{k+1}-1} \beta_{k}(\boldsymbol{Z}(t)) \\
& =\sum_{k=1}^{K_{T}} \sum_{t=t_{k}}^{t_{k+1}-1} \sqrt{\frac{14 \prod_{i=1}^{d}\left(\boldsymbol{X}_{i}(t)+h\right)}{\max \left(1, N_{t_{k}}(\boldsymbol{Z}(t))\right)} \log \left(\frac{2|\mathcal{A}| T}{\tilde{\delta}}\right)} \\
& \leq \sqrt{14 \log \left(\frac{2|\mathcal{A}| T}{\tilde{\delta}}\right)}\left[\sum_{k=1}^{K_{T}} \sum_{t=t_{k}}^{\tilde{t}_{k+1}-1} \sqrt{\frac{2 \prod_{i=1}^{d}\left(\boldsymbol{X}_{i}(t)+h\right)}{\max \left(1, N_{t}(\boldsymbol{Z}(t))\right)}}+\sum_{k=1}^{K_{T}} \sum_{t=\tilde{t}_{k+1}}^{t_{k+1}-1} \sqrt{\prod_{i=1}^{d}\left(\boldsymbol{X}_{i}(t)+h\right)}\right] . \tag{34}
\end{align*}
$$

The first summation can be simplified as

$$
\begin{aligned}
\sum_{k=1}^{K_{T}} \sum_{t=t_{k}}^{\tilde{t}_{k+1}-1} \sqrt{\frac{2 \prod_{i=1}^{d}\left(\boldsymbol{X}_{i}(t)+h\right)}{\max \left(1, N_{t}(\boldsymbol{Z}(t))\right)}} & \leq \sqrt{2\left(M_{\boldsymbol{\theta}^{*}}^{T}+h\right)^{d}} \sum_{k=1}^{K_{T}} \sum_{t=t_{k}}^{\tilde{t}_{k+1}-1} \frac{1}{\sqrt{\max \left(1, N_{t}(\boldsymbol{Z}(t))\right)}} \\
& \leq 3 \sqrt{2\left(M_{\boldsymbol{\theta}^{*}}^{T}+h\right)^{d}} \sum_{\boldsymbol{z} \in\left\{0,1, \cdots, M_{\boldsymbol{\theta}^{*}}^{T}\right\}^{d} \times \mathcal{A}} \sqrt{N_{T+1}(\boldsymbol{z})} \\
& \leq 3 \sqrt{\left.2|\mathcal{A}|\left(M_{\boldsymbol{\theta}^{*}}^{T}+h\right)^{d} \sqrt{\sum_{\boldsymbol{z} \in\left\{0,1, \cdots, M_{\boldsymbol{\theta}^{*}}^{T}\right\}^{d} \times \mathcal{A}} N_{T+1}(\boldsymbol{z})}\right]} \\
& \leq 3 \sqrt{2|\mathcal{A}| T}\left(M_{\boldsymbol{\theta}^{*}}^{T}+h\right)^{d}
\end{aligned}
$$

where the second inequality is due to the following arguments,

$$
\begin{aligned}
\sum_{k=1}^{K_{T}} \sum_{t=t_{k}}^{\tilde{t}_{k+1}-1} \frac{1}{\sqrt{\max \left(1, N_{t}(\boldsymbol{Z}(t))\right)}} & =\sum_{\boldsymbol{z} \in\left\{0,1, \cdots, M_{\boldsymbol{\theta}^{*}}^{T}\right\}^{d} \times \mathcal{A}}\left(\mathbb{I}_{\left\{N_{T+1}(\boldsymbol{z})>0\right\}}+\sum_{i=1}^{N_{T+1}(\boldsymbol{z})-1} \frac{1}{\sqrt{i}}\right) \\
& \leq 3 \sum_{\boldsymbol{z} \in\left\{0,1, \cdots, M_{\boldsymbol{\theta}^{*}}^{T}\right\}^{d} \times \mathcal{A}} \sqrt{N_{T+1}(\boldsymbol{z})} .
\end{aligned}
$$

For the second term in (34), we get

$$
\begin{aligned}
\sum_{k=1}^{K_{T}} \sum_{t=\tilde{t}_{k+1}}^{t_{k+1}-1} \sqrt{\prod_{i=1}^{d}\left(\boldsymbol{X}_{i}(t)+h\right)} & =\sqrt{\left(M_{\boldsymbol{\theta}^{*}}^{T}+h\right)^{d}} \sum_{k=1}^{K_{T}} E_{k} \\
& \leq K_{T}\left(\max _{1 \leq i \leq T} \tau_{0^{d}}^{(i)}\right) \sqrt{\left(M_{\boldsymbol{\theta}^{*}}^{T}+h\right)^{d}} \\
& \leq 2 \sqrt{|\mathcal{A}| T \log _{2} T}\left(\max _{1 \leq i \leq T} \tau_{0^{d}}^{(i)}\right)\left(M_{\boldsymbol{\theta}^{*}}^{T}+h\right)^{d}
\end{aligned}
$$

where $E_{k}=T_{k}-\tilde{T}_{k}$, and $K_{T}$ is bounded from Lemma 8. Thus $\sum_{k=1}^{K_{T}} \sum_{t=t_{k}}^{t_{k+1}-1} \beta_{k}(\boldsymbol{Z}(t))$ is bounded as

$$
\sum_{k=1}^{K_{T}} \sum_{t=t_{k}}^{t_{k+1}-1} \beta_{k}(\boldsymbol{Z}(t)) \leq 24 \sqrt{|\mathcal{A}| T \log _{2} T \log \left(\frac{2|\mathcal{A}| T}{\tilde{\delta}}\right)}\left(\max _{1 \leq i \leq T} \tau_{0^{d}}^{(i)}\right)\left(M_{\boldsymbol{\theta}^{*}}^{T}+h\right)^{d}
$$

Substituting the above bound in (33),

$$
\begin{aligned}
& \mathbb{E}\left[\sum_{k=1}^{K_{T}} \sum_{t=t_{k}}^{t_{k+1}-1} 2 v_{M} \beta_{k}(\boldsymbol{Z}(t))\right] \\
& \leq 48 c_{p_{2}} \sqrt{|\mathcal{A}| T \log _{2} T \log \left(\frac{2|\mathcal{A}| T}{\tilde{\delta}}\right)} \mathbb{E}\left[\left(M_{\boldsymbol{\theta}^{*}}^{T}\right)^{r+r_{*}^{p}}\left(M_{\boldsymbol{\theta}^{*}}^{T}+h\right)^{d}\left(\max _{1 \leq i \leq T} \tau_{0^{d}}^{(i)}\right)\right] \\
& \leq c_{p_{3}} \sqrt{|\mathcal{A}| T \log _{2} T \log \left(\frac{2|\mathcal{A}| T}{\tilde{\delta}}\right)} \mathbb{E}\left[\left(M_{\boldsymbol{\theta}^{*}}^{T}+h\right)^{d+r+r_{*}^{p}}\left(\max _{1 \leq i \leq T} \tau_{0^{d}}^{(i)}\right)\right]
\end{aligned}
$$

where $c_{p_{3}}:=48 c_{p_{2}}$. Finally, from the above equation, (32), and 28,

$$
\begin{aligned}
R_{2} & \leq \tilde{\delta}(\log (h(T+1))+1)^{d} T Q(T) \\
& +c_{p_{3}} \sqrt{|\mathcal{A}| T \log _{2} T \log \left(\frac{2|\mathcal{A}| T}{\tilde{\delta}}\right)} \mathbb{E}\left[\left(M_{\boldsymbol{\theta}^{*}}^{T}+h\right)^{d+r+r_{*}^{p}}\left(\max _{1 \leq i \leq T} \tau_{0^{d}}^{(i)}\right)\right]
\end{aligned}
$$

By choosing $\tilde{\delta}=\frac{1}{T Q(T)}$, we get
$R_{2}$
$\leq(\log (h(T+1))+1)^{d}+c_{p_{3}} \sqrt{|\mathcal{A}| T \log _{2} T \log \left(2|\mathcal{A}| T^{2} Q(T)\right)} \mathbb{E}\left[\left(M_{\boldsymbol{\theta}^{*}}^{T}+h\right)^{d+r+r_{*}^{p}}\left(\max _{1 \leq i \leq T} \tau_{0^{d}}^{(i)}\right)\right]$,
$\leq(\log (h(T+1))+1)^{d}+c_{p_{3}} \sqrt{|\mathcal{A}| T} \log _{2}\left(2|\mathcal{A}| T^{2} Q(T)\right) \mathbb{E}\left[\left(M_{\boldsymbol{\theta}^{*}}^{T}+h\right)^{d+r+r_{*}^{p}}\left(\max _{1 \leq i \leq T} \tau_{0^{d}}^{(i)}\right)\right]$,
where $Q(T)=c_{p_{2}}(T h)^{r+r_{*}^{p}}$.

## B. 7 Proof of Theorem 1

Proof. Lemmas 1, 2, and 3 along with Cauchy-Schwarz inequality showed that the regret terms $R_{0}$ and $R_{2}$ are of the order $\tilde{O}\left(K r d J^{*} h^{d+2 r+r_{*}^{p}} \sqrt{|\mathcal{A}| T}\right)$ and the term $R_{1}$ is $\tilde{O}\left(J^{*}(h)^{r_{*}^{p}}\right)$. Therefore, from $R\left(T, \pi_{T S D E}\right)=R_{0}+R_{1}+R_{2}$, the regret of Algorithm 1, $R\left(T, \pi_{T S D E}\right)$, is $\tilde{O}\left(K r d J^{*} h^{d+2 r+r_{*}^{p}} \sqrt{|\mathcal{A}| T}\right)$.

## B. 8 Requirement of an optimal policy oracle.

To implement our algorithm, we need to find the optimal policy for each model sampled by the algorithm-optimal policy for Theorem 1 and optimal policy within policy class $\Pi$ for Corollary 1 . this has also been used in past work [23, 24, 36]. In the finite state-space setting, [49] provides a schedule of $\epsilon$ values and selects $\epsilon$-optimal policies to obtain $\tilde{O}(\sqrt{T})$ regret guarantees. The issue with extending the analysis of [49] to the countable state-space setting is that we need to ensure (uniform) ergodicity for the chosen $\epsilon$-optimal policies; the lim sup or liminf of the time-average expected reward (used to define the average cost problem) being finite doesn't imply ergodicity. In other words, we must formulate (and verify) ergodicity assumptions for a potentially large set of close-to-optimal algorithms whose structure is undetermined. Another issue is that, to the best of our knowledge, there isn't a general structural characterization of all $\epsilon$-optimal stationary policies for countable state-space MDPs or even a characterization of the policy within this set that is selected by any computational procedure in the literature; current results only discuss existence and characterization of the stationary optimal policy. In the absence of such results, stability assumptions with the same uniformity across models as in our submission will be needed, which are likely too strong to be useful.
If we could verify the stability requirements of Assumptions 3 and 4 for a subset of policies, the optimal oracle is not needed, and instead, by choosing approximately optimal policies within this subset, we can follow the same proof steps as [49] to guarantee regret performance similar to Corollary 1 (without knowledge of model parameters). To theoretically analyze the performance of the algorithm that follows an approximately optimal policy rather than the optimal one, we assume that for a specific sequence of $\left\{\epsilon_{k}\right\}_{k=1}^{\infty}$, an $\epsilon_{k}$-optimal policy is given, which is defined below.

Definition 1. Policy $\pi \in \Pi$ is called an $\epsilon$-optimal policy iffor every $\theta \in \Theta$,

$$
c(\boldsymbol{x}, \pi(\boldsymbol{x}))+\sum_{\boldsymbol{y} \in \mathcal{X}} P_{\theta}(\boldsymbol{y} \mid \boldsymbol{x}, \pi(\boldsymbol{x})) v(\boldsymbol{y} ; \theta) \leq c\left(\boldsymbol{x}, \pi_{\theta}^{*}(\boldsymbol{x})\right)+\sum_{\boldsymbol{y} \in \mathcal{X}} P_{\theta}\left(\boldsymbol{y} \mid \boldsymbol{x}, \pi_{\theta}^{*}(\boldsymbol{x})\right) v(\boldsymbol{y} ; \theta)+\epsilon,
$$

where $\pi_{\theta}^{*}$ is the optimal policy in the policy class $\Pi$ corresponding to parameter $\theta$ and $v(. ; \theta)$ is the solution to Poisson equation (5).

Given $\epsilon$-optimal policies that satisfy Assumptions 3 and 4, in Theorem 2 we extend the regret guarantees of Corollary 1 to the algorithm employing $\epsilon$-optimal policy, instead of the best-in-class policy, and show that the same regret upper bounds continue to apply.

Theorem 3. Consider a non-negative sequence $\left\{\epsilon_{k}\right\}_{k=1}^{\infty}$ such that for every $k \in \mathbb{N}, \epsilon_{k}$ is bounded above by $\frac{1}{k+1}$ and an $\epsilon_{k}$-optimal policy satisfying Assumptions 3 and 4 is given. The regret incurred by Algorithm $\left\lceil\right.$ while using the $\epsilon_{k}$-optimal policy during any episode $k$ is $\tilde{O}\left(d h^{d} \sqrt{|\mathcal{A}| T}\right)$.

Proof. For the $\epsilon_{k}$-optimal policy used in episode $k$, shown by $\pi^{\epsilon_{k}}$, we have

$$
\begin{aligned}
c\left(\boldsymbol{x}, \pi^{\epsilon_{k}}(\boldsymbol{x})\right)+\sum_{\boldsymbol{y} \in \mathcal{X}} P_{\theta_{k}}\left(\boldsymbol{y} \mid \boldsymbol{x}, \pi^{\epsilon_{k}}(\boldsymbol{x})\right) v\left(\boldsymbol{y} ; \theta_{k}\right) & \leq c\left(\boldsymbol{x}, \pi_{\theta_{k}}^{*}(\boldsymbol{x})\right)+\sum_{\boldsymbol{y} \in \mathcal{X}} P_{\theta_{k}}\left(\boldsymbol{y} \mid \boldsymbol{x}, \pi_{\theta_{k}}^{*}(\boldsymbol{x})\right) v\left(\boldsymbol{y} ; \theta_{k}\right)+\epsilon_{k} \\
& =J\left(\theta_{k}\right)+v\left(\boldsymbol{x} ; \theta_{k}\right)+\epsilon_{k}
\end{aligned}
$$

Thus,

$$
\begin{aligned}
& R\left(T, \pi_{T S D E}\right)=\mathbb{E}\left[\sum_{k=1}^{K_{T}} \sum_{t=t_{k}}^{t_{k+1}-1} c\left(\boldsymbol{X}(t), \pi^{\epsilon_{k}}(\boldsymbol{X}(t))\right)\right]-T \mathbb{E}\left[J\left(\boldsymbol{\theta}^{*}\right)\right]=R_{0}+R_{1}+R_{2}+\mathbb{E}\left[\sum_{k=1}^{K_{T}} T_{k} \epsilon_{k}\right] \\
& \text { with } R_{0}=\mathbb{E}\left[\sum_{k=1}^{K_{T}} T_{k} J\left(\theta_{k}\right)\right]-T \mathbb{E}\left[J\left(\boldsymbol{\theta}^{*}\right)\right], \\
& R_{1}=\mathbb{E}\left[\sum_{k=1}^{K_{T}} \sum_{t=t_{k}}^{t_{k+1}-1}\left[v\left(\boldsymbol{X}(t) ; \theta_{k}\right)-v\left(\boldsymbol{X}(t+1) ; \theta_{k}\right)\right]\right], \\
& R_{2}=\mathbb{E}\left[\sum_{k=1}^{K_{T}} \sum_{t=t_{k}}^{t_{k+1}-1}\left[v\left(\boldsymbol{X}(t+1) ; \theta_{k}\right)-\sum_{\boldsymbol{y} \in \mathcal{X}} P_{\theta_{k}}\left(\boldsymbol{y} \mid \boldsymbol{X}(t), \pi^{\epsilon_{k}}(\boldsymbol{X}(t))\right) v\left(\boldsymbol{y} ; \theta_{k}\right)\right]\right] .
\end{aligned}
$$

We assumed that given $\epsilon$-optimal policies satisfy Assumptions 3 and 4 As a result, we can utilize the proof of Theorem 1 to deduce that the term $R_{0}+R_{1}+R_{2}$ is of the order $\tilde{O}\left(d h^{d} \sqrt{|\mathcal{A}| T}\right)$. Moreover, we can simplify the term $\mathbb{E}\left[\sum_{k=1}^{K_{T}} T_{k} \epsilon_{k}\right]$ as below:

$$
\begin{equation*}
\mathbb{E}\left[\sum_{k=1}^{K_{T}} T_{k} \epsilon_{k}\right]=\mathbb{E}\left[\sum_{k=1}^{K_{T}} \tilde{T}_{k} \epsilon_{k}\right]+\mathbb{E}\left[\sum_{k=1}^{K_{T}} E_{k} \epsilon_{k}\right] \tag{35}
\end{equation*}
$$

From the second stopping condition of Algorithm 1 , we have $\tilde{T}_{k} \leq \tilde{T}_{k-1}+1 \leq \ldots \leq k+1$ and

$$
\mathbb{E}\left[\sum_{k=1}^{K_{T}} T_{k} \epsilon_{k}\right] \leq \mathbb{E}\left[K_{T}\right]
$$

where we have used the assumption that $\epsilon_{k} \leq \frac{1}{k+1}$. For the second term of 35), from 25)

$$
\begin{equation*}
\mathbb{E}\left[\sum_{k=1}^{K_{T}} E_{k} \epsilon_{k}\right] \leq \mathbb{E}\left[\sum_{k=1}^{K_{T}} \frac{E_{k}}{k+1}\right] \leq \mathbb{E}\left[\max _{1 \leq i \leq T} \tau_{0^{d}}^{(i)} \sum_{k=1}^{K_{T}} \frac{1}{k+1}\right] \leq \mathbb{E}\left[\max _{1 \leq i \leq T} \tau_{0^{d}}^{(i)} \log \left(K_{T}+1\right)\right] \tag{36}
\end{equation*}
$$

where in the last inequality we have used $\sum_{i=1}^{n} \frac{1}{n} \leq 1+\log (n)$. Finally, as a result of Lemma 6 and Lemma 8 , the result follows.

## C Bounds on hitting times under polynomial and geometric ergodicity

## C. 1 Polynomial upper bounds for the moments of hitting time of state $0^{d}$

For any $\theta_{1}, \theta_{2} \in \Theta$, consider the Markov process with transition kernel $P_{\theta_{1}}^{\pi_{\theta_{2}}^{*}}$ obtained from the $\operatorname{MDP}\left(\mathcal{X}, \mathcal{A}, c, P_{\theta_{1}}\right)$ by following policy $\pi_{\theta_{2}}^{*}$. [29], Lemma 3.5] establishes that if the process is polynomially ergodic, equivalently satisfies $\sqrt[4]{ }$, then for every $0<\eta \leq 1$, there exists constants $\beta_{\theta_{1}, \theta_{2}}^{\eta}, b_{\theta_{1}, \theta_{2}}^{\eta}>0$ such that the following holds:

$$
\begin{equation*}
\Delta\left(V_{\theta_{1}, \theta_{2}}^{p}\right)^{\eta}(\boldsymbol{x}) \leq-\beta_{\theta_{1}, \theta_{2}}^{\eta}\left(V_{\theta_{1}, \theta_{2}}^{p}(\boldsymbol{x})\right)^{\alpha_{\theta_{1}, \theta_{2}}^{p}+\eta-1}+b_{\theta_{1}, \theta_{2}}^{\eta} \mathbb{I}_{C_{\theta_{1}, \theta_{2}}^{p}}(\boldsymbol{x}), \quad \boldsymbol{x} \in \mathcal{X} \tag{37}
\end{equation*}
$$

where for $\eta \in(0,1), \tilde{\beta}_{\theta_{1}, \theta_{2}}^{p}:=\min \left(\beta_{\theta_{1}, \theta_{2}}^{p}, 1\right)$ and

$$
\begin{equation*}
\beta_{\theta_{1}, \theta_{2}}^{\eta}=\eta \tilde{\beta}_{\theta_{1}, \theta_{2}}^{p}, b_{\theta_{1}, \theta_{2}}^{\eta}=\left(b_{\theta_{1}, \theta_{2}}^{p}\right)^{\eta}+\eta \tilde{\beta}_{\theta_{1}, \theta_{2}}^{p} \max \left(1,\left(\tilde{\beta}_{\theta_{1}, \theta_{2}}^{p}\right)^{\left(\alpha_{\theta_{1}, \theta_{2}}^{p}+\eta-1\right) /\left(1-\alpha_{\theta_{1}, \theta_{2}}^{p}\right)}\right) \tag{38}
\end{equation*}
$$

and for $\eta=1, \beta_{\theta_{1}, \theta_{2}}^{\eta}=\beta_{\theta_{1}, \theta_{2}}^{p}$ and $b_{\theta_{1}, \theta_{2}}^{\eta}=b_{\theta_{1}, \theta_{2}}^{p}$. Consequently, the following result is immediate from the proof of [29, Theorem 3.6]; for completeness, we provide the proof in Appendix D. 1

Lemma 9. Suppose a finite set $C_{\theta_{1}, \theta_{2}}^{p}$, constants $\beta_{\theta_{1}, \theta_{2}}^{p}, b_{\theta_{1}, \theta_{2}}^{p}>0, r /(r+1) \leq \alpha_{\theta_{1}, \theta_{2}}^{p}<1$, and a function $V_{\theta_{1}, \theta_{2}}^{p}: \mathcal{X} \rightarrow[1,+\infty)$ exist such that (4) holds. Then, there exist a sequence of non-negative functions $V_{\theta_{1}, \theta_{2}}^{i}: \mathcal{X} \rightarrow[1,+\infty)$ for $i=0, \ldots, r+1$ that satisfy the following system of drift equations for finite sets $C_{\theta_{1}, \theta_{2}}^{i}$, constants $b_{\theta_{1}, \theta_{2}}^{i} \geq 0$ and $\beta_{\theta_{1}, \theta_{2}}^{i}>0$ :

$$
\begin{equation*}
\Delta V_{\theta_{1}, \theta_{2}}^{i-1}(\boldsymbol{x}) \leq-\beta_{\theta_{1}, \theta_{2}}^{i} V_{\theta_{1}, \theta_{2}}^{i}(\boldsymbol{x})+b_{\theta_{1}, \theta_{2}}^{i} \mathbb{I}_{C_{\theta_{1}, \theta_{2}}^{i}}(\boldsymbol{x}), \quad \boldsymbol{x} \in \mathcal{X}, i=1, \ldots, r+1 \tag{39}
\end{equation*}
$$

Notice that $r$ is the maximum degree of the cost function $c$ defined in Assumption 1 . Following the proof and approach of [29] and using the set of equations 39], we can find an upper-bound for $\mathbb{E}_{\boldsymbol{x}}\left[\tau_{0^{d}}^{i}\right]$ for $i=1, \ldots, r+1$ in Lemma 10 In order to establish upper bounds for the first $r+1$ moments of $\tau_{0^{d}}$, it is crucial to choose the value of $\alpha_{\theta_{1}, \theta_{2}}^{p}$ greater than or equal to $\frac{r}{r+1}$, as demonstrated in the proof of Lemma 10 in Appendix D. 2
Lemma 10. For $i=1, \ldots, r+1$, and for all $\boldsymbol{x} \in \mathcal{X}$

$$
\mathbb{E}_{\boldsymbol{x}}^{\pi_{\theta_{2}}^{*}}\left[\left(\tau_{0^{d}}\right)^{i}\right] \leq i \phi_{\theta_{1}, \theta_{2}}^{p}(i)\left(V_{\theta_{1}, \theta_{2}}^{p}(\boldsymbol{x})+b_{\theta_{1}, \theta_{2}}^{p} \alpha_{C_{\theta_{1}, \theta_{2}}^{p}}\right),
$$

where $\phi_{\theta_{1}, \theta_{2}}^{p}(i):=\prod_{j=1}^{i} \frac{1}{\beta_{\theta_{1}, \theta_{2}}^{\eta_{j}}}\left(2^{j-1}+(j-1) \alpha_{C_{\theta_{1}, \theta_{2}}^{p}} b_{\theta_{1}, \theta_{2}}^{\eta_{j}}\right), \eta_{i}=1-(i-1)\left(1-\alpha_{\theta_{1}, \theta_{2}}^{p}\right)$, $b_{\theta_{1}, \theta_{2}}^{\eta_{i}}$ and $\beta_{\theta_{1}, \theta_{2}}^{\eta_{i}}$ defined in (38), and $\alpha_{C_{\theta_{1}, \theta_{2}}^{p}}=\left(\min _{\boldsymbol{y} \in C_{\theta_{1}, \theta_{2}}^{p}} K_{\theta_{1}, \theta_{2}}(\boldsymbol{y})\right)^{-1}$.
Based on Lemma 10, we impose the conditions of Assumption 4 to obtain uniform (over model class) and polynomial (in norm of the state) upper-bounds on the moments of hitting times to $0^{d}$. Moreover, these conditions lead to a uniform characterization of parameters of Lemma 10 over all models in our class.

## C. 2 Distribution of return times to state $0^{d}$

For any $\theta_{1}, \theta_{2} \in \Theta$, consider the Markov process with transition kernel $P_{\theta_{1}}^{\pi_{\theta_{2}}^{*}}$ obtained from the MDP $\left(\mathcal{X}, \mathcal{A}, c, P_{\theta_{1}}\right)$ by following policy $\pi_{\theta_{2}}^{*}$. In the following lemma, we show that the tail probabilities of the return times to the common state $0^{d}$, again $\tau_{0^{d}}$, converge geometrically fast to 0 , and characterize the convergence parameters in terms of the constants given in Assumption 3 . Explicitly, we show

$$
\mathbb{P}_{0^{d}}\left(\tau_{0^{d}}>n\right) \leq c_{\theta_{1}, \theta_{2}}^{g}\left(\tilde{\gamma}_{\theta_{1}, \theta_{2}}^{g}\right)^{n}
$$

for problem and policy dependent constants $c_{\theta_{1}, \theta_{2}}^{g}$ and $\tilde{\gamma}_{\theta_{1}, \theta_{2}}^{g}$. We will follow the method outlined in [27] with the goal to identify problem dependent parameters that will be relevant to our results. Proof of the following lemma is given in Appendix D.3 and follows the methodology of [27].
Lemma 11. For every $\theta_{1}, \theta_{2} \in \Theta$ in the Markov process obtained from the Markov decision process $\left(\mathcal{X}, \mathcal{A}, c, P_{\theta_{1}}\right)$ following policy $\pi_{\theta_{2}}^{*}$, the return time to state 0 starting from state 0 satisfies the following:

$$
\mathbb{P}_{0^{d}}\left(\tau_{0^{d}}>n\right) \leq c_{\theta_{1}, \theta_{2}}^{g}\left(\tilde{\gamma}_{\theta_{1}, \theta_{2}}^{g}\right)^{n}
$$

where

$$
c_{\theta_{1}, \theta_{2}}^{g}=\frac{b_{\theta_{1}, \theta_{2}}^{g}\left(\tilde{b}_{\theta_{1}, \theta_{2}}^{g}\right)^{2}}{\tilde{b}_{\theta_{1}, \theta_{2}}^{g}-1} \quad \text { and } \quad \tilde{\gamma}_{\theta_{1}, \theta_{2}}^{g}=1-\frac{1}{\tilde{b}_{\theta_{1}, \theta_{2}}^{g}},
$$

with

$$
\tilde{b}_{\theta_{1}, \theta_{2}}^{g}=\frac{3 b_{\theta_{1}, \theta_{2}}^{g}+1}{1-\gamma_{\theta_{1}, \theta_{2}}^{g}}\left(\left|C_{\theta_{1}, \theta_{2}}^{g}\right|^{2} \max \left(1, \max _{\boldsymbol{u} \in C_{\theta_{1}, \theta_{2}}^{g} \backslash\left\{0^{d}\right\}} \mathbb{E}_{\boldsymbol{u}}^{\pi_{\theta_{2}}^{*}}\left[\tau_{0^{d}}\right]\right)\right)
$$

Based on Lemma 11, it is necessary to impose the conditions in Assumption 3 to obtain uniform tail probability bounds on $\tau_{0^{d}}$ for all model parameters and policy choices in $\Theta$. Moreover, these conditions lead to a uniform characterization of $c_{\theta_{1}, \theta_{2}}^{g}$ and $\tilde{\gamma}_{\theta_{1}, \theta_{2}}^{g}$ over $\Theta$. Furthermore, as a result of Lemma 10 and uniformity conditions of Assumption $4, \mathbb{E}_{\boldsymbol{u}}^{\pi_{\theta_{2}}^{*}}\left[\tau_{0^{d}}\right]$ has a uniform bound over $\Theta$ and $C_{\theta_{1}, \theta_{2}}^{g} \backslash\left\{0^{d}\right\}$, which can be characterized in terms of the polynomial Lyapunov function.

## D Proofs of hitting time bounds

## D. 1 Proof of Lemma 9

Proof. In the proof, to avoid cumbersome notation we will drop the indices $\theta_{1}, \theta_{2}$. Following the proof of Theorem 3.6 in [29], we choose $\eta_{i}=1-(i-1)\left(1-\alpha^{p}\right)$ for $i=1, \ldots, r+1$ and note that as $\alpha^{p} \in\left[\frac{r}{r+1}, 1\right)$, we have $\eta_{i} \in\left[\frac{1}{r+1}, 1\right]$. As a result, we can apply 37) to each $\eta_{i}$ to get

$$
\Delta\left(V^{p}\right)^{\eta_{i}}(\boldsymbol{x}) \leq-\beta^{\eta_{i}}\left(V^{p}(\boldsymbol{x})\right)^{i \alpha^{p}-i+1}+b^{\eta_{i}} \mathbb{I}_{C^{p}}(\boldsymbol{x}), \quad i=1, \ldots, r+1
$$

Thus, the system of drift equations (39) hold for

$$
\begin{array}{ll}
V_{i}=\left(V^{p}\right)^{1-i\left(1-\alpha^{p}\right)}, & i=0, \ldots, r+1 \\
\beta_{i}=\beta^{\eta_{i}}, & i=1, \ldots, r+1 \\
b_{i}=b^{\eta_{i}}, & i=1, \ldots, r+1 \\
C_{i}=C^{p}, & i=1, \ldots, r+1
\end{array}
$$

where $\beta^{\eta_{i}}$ and $b^{\eta_{i}}$ are defined in (38).

## D. 2 Proof of Lemma 10

The proof of Lemma 10 uses the following lemma.
Lemma 12 (Proposition 11.3.2, [43]). Suppose for nonnegative functions $f, g$, and $V$ on the state space $\mathcal{X}$ and every $k \in \mathbb{Z}_{+}$, the following holds:

$$
\mathbb{E}\left[V\left(X_{k+1}\right) \mid \mathcal{F}_{k}\right] \leq V\left(X_{k}\right)-f\left(X_{k}\right)+g\left(X_{k}\right) .
$$

Then, for any initial condition $x$ and stopping time $\tau$

$$
\mathbb{E}_{x}\left[\sum_{k=0}^{\tau-1} f\left(X_{k}\right)\right] \leq V(x)+\mathbb{E}_{x}\left[\sum_{k=0}^{\tau-1} g\left(X_{k}\right)\right]
$$

Proof of Lemma 10 Following [29], the proof uses an induction argument. We will use the notation of Lemma 9 for simplicity. Similarly, in this proof we will also denote $\phi_{\theta_{1}, \theta_{2}}^{p}(i)$ as $\phi(i), K_{\theta_{1}, \theta_{2}}(\cdot)$ as $K(\cdot)$, and $\vec{V}_{\theta_{1}, \theta_{2}}^{i}, b_{\theta_{1}, \theta_{2}}^{i}, \beta_{\theta_{1}, \theta_{2}}^{i}, C_{\theta_{1}, \theta_{2}}^{i}$ as $V_{i}, b_{i}, \beta_{i}, C_{i}$.
From irreducibility, for all $\boldsymbol{x} \in \mathcal{X}, K(\boldsymbol{x})$ is positive and finite. Considering the system of drift equations found in Lemma $9, C_{i}=C^{p}$ is a finite set for all $i=1, \ldots, r+1$. Thus, $\min _{\boldsymbol{y} \in C_{i}} K(\boldsymbol{y})$ is strictly positive. For all $\boldsymbol{x} \in \mathcal{X}$ and $i=1, \ldots, r+1$, we have

$$
\begin{equation*}
\mathbb{I}_{C_{i}}(\boldsymbol{x}) \leq\left(\min _{\boldsymbol{y} \in C_{i}} K(\boldsymbol{y})\right)^{-1} K(\boldsymbol{x}) \tag{40}
\end{equation*}
$$

We set $\alpha_{C^{p}}:=\left(\min _{\boldsymbol{y} \in C_{i}} K(\boldsymbol{y})\right)^{-1}=\left(\min _{\boldsymbol{y} \in C^{p}} K(\boldsymbol{y})\right)^{-1}$. From Lemma 9 , for $j=1$ and $\boldsymbol{x} \in \mathcal{X}$

$$
\Delta V_{0}(\boldsymbol{x}) \leq-\beta_{1} V_{1}(\boldsymbol{x})+b_{1} \mathbb{I}_{C_{1}}(\boldsymbol{x})
$$

By applying Lemma 12 , for all $\boldsymbol{x} \in \mathcal{X}$ we get

$$
\begin{equation*}
\beta_{1} \mathbb{E}_{\boldsymbol{x}}\left[\sum_{k=0}^{\tau_{0}{ }^{d}-1} V_{1}\left(\boldsymbol{X}_{k}\right)\right] \leq V_{0}(\boldsymbol{x})+b_{1} \mathbb{E}_{\boldsymbol{x}}\left[\sum_{k=0}^{\tau_{0}-1} \mathbb{I}_{C_{1}}\left(\boldsymbol{X}_{k}\right)\right] \tag{41}
\end{equation*}
$$

Using (40) and (41), followed by noting that

$$
K(\boldsymbol{x})=\sum_{n=0}^{\infty} 2^{-n-2} P^{n}\left(\boldsymbol{x}, 0^{d}\right)=\sum_{n=0}^{\infty} 2^{-n-2} \mathbb{E}_{\boldsymbol{x}}\left[\mathbb{I}_{0^{d}}\left(\boldsymbol{X}_{n}\right)\right]
$$

we get

$$
\begin{aligned}
\mathbb{E}_{\boldsymbol{x}}\left[\sum_{k=0}^{\tau_{0^{d}}-1} V_{1}\left(\boldsymbol{X}_{k}\right)\right] & \leq \frac{1}{\beta_{1}} V_{0}(\boldsymbol{x})+\frac{b_{1} \alpha_{C^{p}}}{\beta_{1}} \mathbb{E}_{\boldsymbol{x}}\left[\sum_{n=0}^{\infty} 2^{-n-2} \sum_{k=0}^{\tau_{0^{d}}-1} \mathbb{I}_{0^{d}}\left(\boldsymbol{X}_{k+n}\right)\right] \\
& =\frac{1}{\beta_{1}} V_{0}(\boldsymbol{x})+\frac{b_{1} \alpha_{C^{p}}}{\beta_{1}} \mathbb{E}_{\boldsymbol{x}}\left[\sum_{n=0}^{\infty} 2^{-n-2} \sum_{k=n}^{\tau_{0^{d}}-1+n} \mathbb{I}_{0^{d}}\left(\boldsymbol{X}_{k}\right)\right] \\
& \leq \frac{1}{\beta_{1}} V_{0}(\boldsymbol{x})+\frac{b_{1} \alpha_{C^{p}}}{\beta_{1}} \mathbb{E}_{\boldsymbol{x}}\left[\sum_{n=0}^{\infty} 2^{-n-2} \sum_{k=n \vee \tau_{0^{d}}}^{\tau_{0^{d}}-1+n} \mathbb{I}_{0^{d}}\left(\boldsymbol{X}_{k}\right)\right] \\
& \leq \frac{1}{\beta_{1}} V_{0}(\boldsymbol{x})+\frac{b_{1} \alpha_{C^{p}}}{\beta_{1}} \sum_{n=0}^{\infty} 2^{-n-2}(n+1) \\
& =\frac{1}{\beta_{1}} V_{0}(\boldsymbol{x})+\frac{b_{1} \alpha_{C^{p}}}{\beta_{1}} .
\end{aligned}
$$

As $V_{1}(\boldsymbol{x}) \geq 1$, this gives us a bound on $\mathbb{E}_{\boldsymbol{x}}\left[\tau_{0^{d}}\right]$ as follows:

$$
\mathbb{E}_{\boldsymbol{x}}\left[\tau_{0^{d}}\right] \leq \frac{1}{\beta_{1}} V_{0}(\boldsymbol{x})+\frac{b_{1} \alpha_{C^{p}}}{\beta_{1}}
$$

Assume for $i \geq 1$, by the induction assumption we have

$$
\begin{equation*}
\mathbb{E}_{\boldsymbol{x}}\left[\sum_{k=0}^{\tau_{0^{d}}-1}(k+1)^{i-1} V_{i}\left(\boldsymbol{X}_{k}\right)\right] \leq \phi(i)\left(V_{0}(\boldsymbol{x})+b_{1} \alpha_{C^{p}}\right) \tag{42}
\end{equation*}
$$

Set $j=i+1$ in (39), which yields

$$
\Delta V_{i}(\boldsymbol{x}) \leq-\beta_{i+1} V_{i+1}(\boldsymbol{x})+b_{i+1} \mathbb{I}_{C^{p}}(\boldsymbol{x})
$$

Define $Z_{k}=k^{i} V_{i}\left(\boldsymbol{X}_{k}\right)$. From the above equation, we have

$$
\begin{aligned}
\mathbb{E}\left[Z_{k+1} \mid \boldsymbol{X}_{k}\right] & \leq(k+1)^{i}\left(V_{i}\left(\boldsymbol{X}_{k}\right)-\beta_{i+1} V_{i+1}\left(\boldsymbol{X}_{k}\right)+b_{i+1} \mathbb{I}_{C^{p}}\left(\boldsymbol{X}_{k}\right)\right) \\
& \leq Z_{k}+2^{i}(k+1)^{i-1} V_{i}\left(\boldsymbol{X}_{k}\right)+(k+1)^{i} b_{i+1} \mathbb{I}_{C^{p}}\left(\boldsymbol{X}_{k}\right)-(k+1)^{i} \beta_{i+1} V_{i+1}\left(\boldsymbol{X}_{k}\right)
\end{aligned}
$$

By applying Lemma 12 to the above equation, we get

$$
\begin{align*}
& \beta_{i+1} \mathbb{E}_{\boldsymbol{x}}\left[\sum_{k=0}^{\tau_{0^{d}}-1}(k+1)^{i} V_{i+1}\left(\boldsymbol{X}_{k}\right)\right] \\
& \leq 2^{i} \mathbb{E}_{\boldsymbol{x}}\left[\sum_{k=0}^{\tau_{0}-1}(k+1)^{i-1} V_{i}\left(\boldsymbol{X}_{k}\right)\right]+b_{i+1} \mathbb{E}_{\boldsymbol{x}}\left[\sum_{k=0}^{\tau_{0}-1}(k+1)^{i} \mathbb{I}_{C^{p}}\left(\boldsymbol{X}_{k}\right)\right] \\
& \leq 2^{i} \phi(i)\left(V_{0}(\boldsymbol{x})+b_{1} \alpha_{C^{p}}\right)+\alpha_{C^{p}} b_{i+1} \mathbb{E}_{\boldsymbol{x}}\left[\left(\tau_{0^{d}}\right)^{i}\right] \tag{43}
\end{align*}
$$

where the second inequality follows from (40) and the induction hypothesis 42). Thereafter, from (42) (by using integral lower bound after using $V_{i} \geq 1$ ), we have

$$
\frac{1}{i} \mathbb{E}_{\boldsymbol{x}}\left[\left(\tau_{0^{d}}\right)^{i}\right] \leq \mathbb{E}_{\boldsymbol{x}}\left[\sum_{k=0}^{\tau_{0^{d}}-1}(k+1)^{i-1} V_{i}\left(\boldsymbol{X}_{k}\right)\right] \leq \phi(i)\left(V_{0}(\boldsymbol{x})+b_{1} \alpha_{C^{p}}\right)
$$

Substituting in 433, we get

$$
\begin{aligned}
\beta_{i+1} \mathbb{E}_{\boldsymbol{x}}\left[\sum_{k=0}^{\tau_{0^{d}}-1}(k+1)^{i} V_{i+1}\left(\boldsymbol{X}_{k}\right)\right] & \leq 2^{i} \phi(i)\left(V_{0}(\boldsymbol{x})+b_{1} \alpha_{C^{p}}\right)+i b_{i+1} \alpha_{C^{p}} \phi(i)\left(V_{0}(\boldsymbol{x})+b_{1} \alpha_{C^{p}}\right) \\
& =\left(2^{i}+i b_{i+1} \alpha_{C^{p}}\right) \phi(i)\left(V_{0}(\boldsymbol{x})+b_{1} \alpha_{C^{p}}\right) \\
& =\beta_{i+1} \phi(i+1)\left(V_{0}(\boldsymbol{x})+b_{1} \alpha_{C^{p}}\right) .
\end{aligned}
$$

This completes the proof.

## D. 3 Proof of Lemma 11

Proof. In the proof, to avoid cumbersome notation we will drop the indices $\theta_{1}, \theta_{2}$. Based on Assumption 33, there exists a finite set $C^{g}$, constants $b^{g}, \gamma^{g} \in(0,1)$, and a function $V^{g}: \mathcal{X} \rightarrow$ $[1,+\infty)$ satisfying

$$
\begin{equation*}
\Delta V^{g}(\boldsymbol{x}) \leq-\left(1-\gamma^{g}\right) V^{g}(\boldsymbol{x})+b^{g} \mathbb{I}_{C^{g}}(\boldsymbol{x}), \quad \boldsymbol{x} \in \mathcal{X} \tag{44}
\end{equation*}
$$

For $n \geq 1$, define the $n$-step taboo probabilities [43] as

$$
{ }_{A} P_{\boldsymbol{x} B}^{n}=\mathbb{P}_{\boldsymbol{x}}\left(\boldsymbol{X}_{n} \in B, \tau_{A}>n\right),
$$

$\underset{\tilde{V}}{\text { where }} A, B \subseteq \mathcal{X}$, and $\tau_{A}$ is the first hitting time of set $A$. We also let ${ }_{A} P_{\boldsymbol{x} B}^{0}=\mathbb{I}_{B}(\boldsymbol{x})$ and $\tilde{V}^{g}=\sum_{n=0}^{\infty}{ }_{0} P^{n} V^{g}$. Applying the last exit decomposition on $C^{g} \backslash\left\{0^{d}\right\}$ for all $x \in \mathcal{X}$, we obtain

$$
\begin{align*}
& \tilde{V}^{g}(\boldsymbol{x}) \\
& =\sum_{n=0}^{\infty} \sum_{\boldsymbol{y} \in \mathcal{X}}{ }_{0^{d}} P_{\boldsymbol{x} \boldsymbol{y}}^{n} V^{g}(\boldsymbol{y}) \\
& =V^{g}(\boldsymbol{x})+\sum_{n=1}^{\infty} \sum_{\boldsymbol{y} \in \mathcal{X}} C^{g} P_{\boldsymbol{x} \boldsymbol{y}}^{n} V^{g}(\boldsymbol{y}) \\
& +\sum_{n=1}^{\infty} \sum_{\boldsymbol{y} \in \mathcal{X}} \sum_{m=1}^{n-1} \sum_{\boldsymbol{z} \in C^{g} \backslash\left\{0^{d}\right\}}{ }_{0^{d}} P_{\boldsymbol{x} \boldsymbol{z}}^{m} C^{g} P_{\boldsymbol{z} \boldsymbol{y}}^{n-m} V^{g}(\boldsymbol{y})+\sum_{n=1}^{\infty} \sum_{\boldsymbol{y} \in \mathcal{X}} \sum_{\boldsymbol{z} \in C^{g} \backslash\left\{0^{d}\right\}} 0^{d} P_{\boldsymbol{x} \boldsymbol{z}}^{n} C^{g} P_{\boldsymbol{z} \boldsymbol{y}}^{0} V^{g}(\boldsymbol{y}) \\
& =V^{g}(\boldsymbol{x})+\sum_{n=1}^{\infty} \sum_{\boldsymbol{y} \in \mathcal{X}} C^{g} P_{\boldsymbol{x} \boldsymbol{y}}^{n} V^{g}(\boldsymbol{y})  \tag{45}\\
& +\sum_{\boldsymbol{y} \in \mathcal{X}} \sum_{\boldsymbol{z} \in C^{g} \backslash\left\{0^{d}\right\}}\left(\sum_{m=1}^{\infty}{ }_{0^{d}} P_{\boldsymbol{x} \boldsymbol{z}}^{m}\right)\left(\sum_{n=1}^{\infty} C^{g} P_{\boldsymbol{z} \boldsymbol{y}}^{n} V^{g}(\boldsymbol{y})\right) \tag{46}
\end{align*} \underbrace{}_{\sum_{n=1}^{\sum_{\boldsymbol{z} \in C^{g} \backslash\left\{0^{d}\right\}}^{\infty} \sum_{\text {Term } 2}} 0^{d} P_{\boldsymbol{x} \boldsymbol{z}}^{n} V^{g}(\boldsymbol{z})}
$$

where we break up the trajectories starting at state $\boldsymbol{x}$ and reaching state $\boldsymbol{y}$ while avoiding state $0^{d}$ into two: ones that never visit the set $C^{g}$, and the others that visit $C^{g} \backslash\left\{0^{d}\right\}$ up until time $m$ but not afterwards and exit $C^{g} \backslash\left\{0^{d}\right\}$ at time $m$.
We first bound Term 1 in 46) by finding an upper bound for the probability term $\sum_{m=10^{d}}^{\infty} P_{x z}^{m}$ using the first entrance decomposition on $C^{g} \backslash\left\{0^{d}\right\}$ while noting that $\boldsymbol{z} \in C^{g} \backslash\left\{0^{d}\right\}$ :

$$
\begin{align*}
\sum_{m=1}^{\infty}{ }_{0^{d}} P_{\boldsymbol{x} \boldsymbol{z}}^{m} & =\sum_{m=1}^{\infty} \sum_{l=1}^{m} \sum_{\substack{\left.\boldsymbol{u} \in C^{g} \backslash 0^{d}\right\} \\
\boldsymbol{v} \notin C^{g}}} C^{g} P_{\boldsymbol{x} \boldsymbol{v}}^{l-1} P_{\boldsymbol{v} \boldsymbol{u} 0^{d}} P_{\boldsymbol{u} \boldsymbol{z}}^{m-l} \\
& =\sum_{\boldsymbol{u} \in C^{g} \backslash\left\{0^{d}\right\}}\left(\sum_{l=0}^{\infty} \sum_{\boldsymbol{v} \notin C^{g}} C^{g} P_{\boldsymbol{x} \boldsymbol{v}}^{l} P_{\boldsymbol{v} \boldsymbol{u}}\right)\left(\sum_{m=0}^{\infty}{0^{d}} P_{\boldsymbol{u} \boldsymbol{z}}^{m}\right) \\
& \leq \sum_{\boldsymbol{u} \in C^{g} \backslash\left\{0^{d}\right\}} \sum_{m=0}^{\infty}{ }_{0^{d}} P_{\boldsymbol{u} \boldsymbol{z}}^{m} \\
& \leq \sum_{\boldsymbol{u} \in C^{g} \backslash\left\{0^{d}\right\}} \sum_{m=0}^{\infty} \mathbb{P}_{\boldsymbol{u}}\left(\tau_{0^{d}}>m\right)  \tag{47}\\
& \leq\left|C^{g}\right| \max _{\boldsymbol{u} \in C^{g} \backslash\left\{0^{d}\right\}} \mathbb{E}_{\boldsymbol{u}}\left[\tau_{0^{d}}\right],
\end{align*}
$$

where the third line follows from the fact that $\sum_{l=0}^{\infty} \sum_{\boldsymbol{v} \notin C^{g} C^{g}} P_{\boldsymbol{x} \boldsymbol{v}}^{l} P_{\boldsymbol{v} \boldsymbol{u}}$ is the probability of entrance to $C^{g}$ through $\boldsymbol{u} \in C^{g} \backslash\{0\}$, so it is less than 1. Irreducibility and positive recurrence combined with $\left|C^{g}\right|<\infty$ imply that $\max _{\boldsymbol{u} \in C^{g} \backslash\left\{0^{d}\right\}} \mathbb{E}_{\boldsymbol{u}}\left[\tau_{0^{d}}\right]<\infty$, which shows $\sum_{m=00^{d}}^{\infty} P_{\boldsymbol{x} \boldsymbol{z}}^{m}$ is finite. Next,
by induction we prove that for $n \geq 1$ and $\boldsymbol{z} \in C^{g} \backslash\left\{0^{d}\right\}$ we have

$$
\begin{equation*}
\sum_{\boldsymbol{y} \in \mathcal{X}} C^{g} P_{\boldsymbol{z} \boldsymbol{y}}^{n} V^{g}(\boldsymbol{y}) \leq\left(\gamma^{g}\right)^{n-1} b^{g} \tag{48}
\end{equation*}
$$

For $n=1$, we have using Assumption 3 that

$$
\sum_{\boldsymbol{y} \in \mathcal{X}} C^{g} P_{\boldsymbol{z} \boldsymbol{y}} V^{g}(\boldsymbol{y}) \leq \sum_{\boldsymbol{y} \in \mathcal{X}} P_{\boldsymbol{z} \boldsymbol{y}} V^{g}(\boldsymbol{y}) \leq b^{g} .
$$

Assuming that (48) holds for $n$, for $n+1$ we have

$$
\begin{array}{rlr}
\sum_{\boldsymbol{y} \in \mathcal{X}} C^{g} P_{\boldsymbol{z} \boldsymbol{y}}^{n+1} V^{g}(\boldsymbol{y}) & \leq \sum_{\substack{\boldsymbol{y} \in \mathcal{X} \\
\boldsymbol{v} \notin C^{g}}} C^{g} P_{\boldsymbol{z} \boldsymbol{v}}^{n} P_{\boldsymbol{v} \boldsymbol{y}} V^{g}(\boldsymbol{y}) \leq \gamma^{g} \sum_{\boldsymbol{v} \notin C^{g}} C^{g} P_{\boldsymbol{z} \boldsymbol{v}}^{n} V^{g}(\boldsymbol{v}) \quad \text { (Using (44)) } \\
& \leq \gamma \sum_{\boldsymbol{v} \in \mathcal{X}} C^{g} P_{\boldsymbol{z} \boldsymbol{v}}^{n} V^{g}(\boldsymbol{v}) \leq\left(\gamma^{g}\right)^{n} b^{g}, \quad \quad \text { (By induction step) }
\end{array}
$$

so (48) is shown. We collect these bounds later on for our result on Term 2.
We now simplify the summation in (45). Similar to previous arguments, we will use induction for $n \geq 1$ and show for all $\boldsymbol{x} \in \mathcal{X}$

$$
\begin{equation*}
\sum_{\boldsymbol{y} \in \mathcal{X}} C^{g} P_{\boldsymbol{x} \boldsymbol{y}}^{n} V^{g}(\boldsymbol{y}) \leq\left(\gamma^{g}\right)^{n-1}\left(\gamma^{g} V^{g}(\boldsymbol{x})+b^{g}\right) \tag{49}
\end{equation*}
$$

For $n=1$, we have

$$
\sum_{\boldsymbol{y} \in \mathcal{X}} C^{g} P_{\boldsymbol{x} \boldsymbol{y}} V^{g}(\boldsymbol{y}) \leq \sum_{\boldsymbol{y} \in \mathcal{X}} P_{\boldsymbol{x} \boldsymbol{y}} V^{g}(\boldsymbol{y}) \leq \gamma^{g} V^{g}(\boldsymbol{x})+b^{g} .
$$

Assuming that 49 holds for $n$, for $n+1$ we have

$$
\begin{aligned}
\sum_{\boldsymbol{y} \in \mathcal{X}} C^{g} P_{\boldsymbol{x} \boldsymbol{y}}^{n+1} V^{g}(\boldsymbol{y}) & \leq \sum_{\boldsymbol{z} \notin C^{g}} C^{g} P_{\boldsymbol{x} \boldsymbol{z}}^{n} \sum_{\boldsymbol{y} \in \mathcal{X}} P_{\boldsymbol{z} \boldsymbol{y}} V^{g}(\boldsymbol{y}) \leq \gamma^{g} \sum_{\boldsymbol{z} \notin C^{g}} C^{g} P_{\boldsymbol{x} \boldsymbol{z}}^{n} V^{g}(\boldsymbol{z}) \\
& \leq \gamma^{g} \sum_{\boldsymbol{z} \in \mathcal{X}} C^{g} P_{\boldsymbol{x} \boldsymbol{z}}^{n} V^{g}(\boldsymbol{z}) \leq\left(\gamma^{g}\right)^{n}\left(\gamma^{g} V^{g}(\boldsymbol{x})+b^{g}\right),
\end{aligned}
$$

where the first and second inequalities follow from the definition of taboo probabilities and 44 . Thus, $\sqrt{49}$ is proved. Lastly, for Term 2 in (46), we note

$$
\begin{aligned}
\sum_{n=1}^{\infty} \sum_{\boldsymbol{z} \in C^{g} \backslash\left\{0^{d}\right\}} 0^{d} P_{\boldsymbol{x} \boldsymbol{z}}^{n} V^{g}(\boldsymbol{z}) & \leq \max _{\boldsymbol{y} \in C^{g} \backslash\left\{0^{d}\right\}} V^{g}(\boldsymbol{y}) \sum_{\boldsymbol{z} \in C^{g} \backslash\left\{0^{d}\right\}} \sum_{n=1}^{\infty} 0^{d} P_{\boldsymbol{x} \boldsymbol{z}}^{n} \\
& \leq b^{g}\left|C^{g}\right|^{2} \max _{\boldsymbol{u} \in C^{g} \backslash\left\{0^{d}\right\}} \mathbb{E}_{\boldsymbol{u}}\left[\tau_{0^{d}}\right]
\end{aligned}
$$

From the above equation, (47), (48), and (49), we bound $\tilde{V}^{g}(\boldsymbol{x})$ as follows:

$$
\begin{aligned}
& \tilde{V}^{g}(\boldsymbol{x}) \\
& \leq V^{g}(\boldsymbol{x})+\left(\gamma^{g} V^{g}(\boldsymbol{x})+b^{g}\right) \sum_{n=1}^{\infty}\left(\gamma^{g}\right)^{n-1}+\left|C^{g}\right|^{2} b^{g} \max _{\boldsymbol{u} \in C^{g} \backslash\left\{0^{d}\right\}} \mathbb{E}_{\boldsymbol{u}}\left[\tau_{0^{d}}\right]\left(1+\sum_{n=1}^{\infty}\left(\gamma^{g}\right)^{n-1}\right) \\
& \leq \frac{V^{g}(\boldsymbol{x})}{1-\gamma^{g}}+\frac{3\left|C^{g}\right|^{2} b^{g}}{1-\gamma^{g}} \max \left(1, \max _{\boldsymbol{u} \in C^{g} \backslash\left\{0^{d}\right\}} \mathbb{E}_{\boldsymbol{u}}\left[\tau_{0^{d}}\right]\right) \\
& \leq V^{g}(\boldsymbol{x})\left(\frac{3 b^{g}+1}{1-\gamma^{g}}\left(\left|C^{g}\right|^{2} \max \left(1, \max _{\boldsymbol{u} \in C^{g} \backslash\left\{0^{d}\right\}} \mathbb{E}_{\boldsymbol{u}}\left[\tau_{0^{d}}\right]\right)\right)\right),
\end{aligned}
$$

where the last line is due to $V^{g}(\boldsymbol{x}) \geq 1$. Taking

$$
\tilde{b}^{g}:=\frac{3 b^{g}+1}{1-\gamma^{g}}\left(\left|C^{g}\right|^{2} \max \left(1, \max _{\boldsymbol{u} \in C^{g} \backslash\left\{0^{d}\right\}} \mathbb{E}_{\boldsymbol{u}}\left[\tau_{0^{d}}\right]\right)\right)>1,
$$

we have shown that

$$
\begin{equation*}
\tilde{V}^{g}(\boldsymbol{x}) \leq \tilde{b}^{g} V^{g}(\boldsymbol{x}), \quad \boldsymbol{x} \in \mathcal{X} \tag{50}
\end{equation*}
$$

We now upper-bound $\mathbb{P}_{0^{d}}\left(\tau_{0^{d}}>n\right)$ for all $n \geq 1$ in an inductive manner, starting with $\mathbb{P}_{0^{d}}\left(\tau_{0^{d}}>1\right)$. As a part of showing this, for every $\boldsymbol{x} \neq 0^{d}$ we argue that for all $n \geq 1$

$$
\begin{equation*}
\mathbb{P}_{\boldsymbol{x}}\left(\tau_{0^{d}}>n\right) \leq \tilde{V}^{g}(\boldsymbol{x})\left(1-\frac{1}{\tilde{b}^{g}}\right)^{n} \tag{51}
\end{equation*}
$$

First note that

$$
\begin{equation*}
\tilde{V}^{g}(\boldsymbol{x}) \geq V^{g}(\boldsymbol{x}) \geq 1 \tag{52}
\end{equation*}
$$

Thus,

$$
\begin{align*}
\mathbb{P}_{\boldsymbol{x}}\left(\tau_{0^{d}}>1\right) & =\sum_{\boldsymbol{y} \in \mathcal{X}}{0^{d}}^{d} P_{\boldsymbol{x} \boldsymbol{y}} \leq \sum_{\boldsymbol{y} \in \mathcal{X}}{ }_{0^{d}} P_{\boldsymbol{x} \boldsymbol{y}} \tilde{V}^{g}(\boldsymbol{y}) \\
& =\sum_{\boldsymbol{y} \in \mathcal{X}}{ }_{0^{d}} P_{\boldsymbol{x} \boldsymbol{y}} \sum_{n=0}^{\infty} \sum_{\boldsymbol{z} \in \mathcal{X}}{ }_{0^{d}} P_{\boldsymbol{y} \boldsymbol{z}}^{n} V^{g}(\boldsymbol{z})=\sum_{\boldsymbol{z} \in \mathcal{X}} \sum_{n=1}^{\infty}{ }_{0^{d}} P_{\boldsymbol{x} \boldsymbol{z}}^{n} V^{g}(\boldsymbol{z}) . \tag{53}
\end{align*}
$$

We now apply the bound in (50) to get

$$
\begin{equation*}
\mathbb{P}_{\boldsymbol{x}}\left(\tau_{0^{d}}>1\right) \leq \sum_{\boldsymbol{z} \in \mathcal{X}} \sum_{n=1}^{\infty}{ }_{0^{d}} P_{\boldsymbol{x} \boldsymbol{z}}^{n} V^{g}(\boldsymbol{z})=\tilde{V}^{g}(\boldsymbol{x})-V^{g}(\boldsymbol{x}) \leq \tilde{V}^{g}(\boldsymbol{x})\left(1-\frac{1}{\tilde{b}^{g}}\right) \tag{54}
\end{equation*}
$$

With the base of induction established, we assume the statement in 51 is true for $n$, and show that it continues to hold for $n+1$ as follows:

$$
\begin{aligned}
\mathbb{P}_{\boldsymbol{x}}\left(\tau_{0^{d}}>n+1\right) & =\sum_{\boldsymbol{y} \neq 0^{d}} P_{\boldsymbol{x} \boldsymbol{y}} \mathbb{P}_{\boldsymbol{y}}\left(\tau_{0^{d}}>n\right) \\
& \leq\left(1-\frac{1}{\tilde{b}^{g}}\right)^{n} \sum_{\boldsymbol{y} \neq 0^{d}} P_{\boldsymbol{x} \boldsymbol{y}} \tilde{V}^{g}(\boldsymbol{y}) \\
& \leq \tilde{V}^{g}(\boldsymbol{x})\left(1-\frac{1}{\tilde{b}^{g}}\right)^{n+1}
\end{aligned}
$$

where the final inequality uses the same arguments as in (53) and (54).
Finally, using the tail probabilities of hitting time of state $0^{d}$ from any state $\boldsymbol{x} \neq 0^{d}$, we bound the tail probability of the return time to state $0^{d}$ (starting from $0^{d}$ ) as follows

$$
\begin{aligned}
\mathbb{P}_{0^{d}}\left(\tau_{0^{d}}>n+1\right) & =\sum_{\boldsymbol{x} \neq 0^{d}} P_{0 \boldsymbol{x}} \mathbb{P}_{\boldsymbol{x}}\left(\tau_{0^{d}}>n\right) \leq\left(1-\frac{1}{\tilde{b}^{g}}\right)^{n} \sum_{\boldsymbol{x} \neq 0^{d}} P_{0 \boldsymbol{x}} \tilde{V}^{g}(\boldsymbol{x}) \\
& \leq \tilde{b}^{g}\left(1-\frac{1}{\tilde{b}^{g}}\right)^{n} \sum_{\boldsymbol{x} \neq 0^{d}} P_{0 \boldsymbol{x}} V^{g}(\boldsymbol{x}) \leq b^{g} \tilde{b}^{g}\left(1-\frac{1}{\tilde{b}^{g}}\right)^{n}
\end{aligned}
$$

where the final inequality follows from the definition of $b^{g}$, and we have

$$
\tilde{\gamma}^{g}=1-\frac{1}{\tilde{b}^{g}}, \text { and } c^{g}=\frac{b^{g}\left(\tilde{b}^{g}\right)^{2}}{\tilde{b}^{g}-1}
$$

and the proof is complete.

## E Queueing model examples

## E. 1 Model 1: Two-server queueing system with a common buffer

We consider a continuous-time queueing system with two heterogeneous servers with unknown service rate vector $\boldsymbol{\theta}^{*}=\left(\theta_{1}^{*}, \theta_{2}^{*}\right)$ and a common infinite buffer, shown in Figure 2 a . Arrivals to the system are according to a Poisson process with rate $\lambda$ and service times are exponentially distributed
with parameter $\theta_{i}^{*}$, depending on the assigned server. The service rate vector $\boldsymbol{\theta}^{*}$ is sampled from the prior distribution $\nu_{0}$ defined on the space $\Theta$ given as

$$
\begin{equation*}
\Theta=\left\{\left(\theta_{1}, \theta_{2}\right) \in \mathbb{R}_{+}^{2}: \frac{\lambda}{\theta_{1}+\theta_{2}} \leq \frac{1-\delta}{1+\delta}, 1 \leq \frac{\theta_{1}}{\theta_{2}} \leq R\right\} \tag{55}
\end{equation*}
$$

for fixed $\delta \in(0,0.5)$ and $R \geq 1$. Note that for any $\left(\theta_{1}, \theta_{2}\right) \in \Theta$, we have $\theta_{1} \geq \theta_{2}$ and the stability requirement $\lambda<\theta_{1}+\theta_{2}$ holds. The countable state space $\mathcal{X}$ is defined as $\mathcal{X}=\left\{\boldsymbol{x}=\left(x_{0}, x_{1}, x_{2}\right): x_{0} \in \mathbb{N} \cup\{0\}, x_{1}, x_{2} \in\{0,1\}\right\}$, in which $x_{0}$ is the length of the queue, and $x_{i}, i=1,2$ is equal to 1 if server $i$ is busy serving a job. At each time instance $r \in \mathbb{R}_{+}$, the dispatcher can assign jobs from the (non-empty) buffer to an available server. Thus, the action space $\mathcal{A}$ is equal to

$$
\mathcal{A}=\{h, b, 1,2\}
$$

where $h$ indicates no action, $b$ sends a job to both of the servers, and $i=1,2$ assigns a job to server $i$. The goal of the dispatcher is to minimize the expected sojourn time of customers, which by Little's law [52] is equivalent to minimizing the average number of customers in the system, or

$$
\begin{equation*}
\inf _{\pi \in \Pi} \limsup _{T \rightarrow \infty} \frac{1}{T} \int_{0}^{T}\|\boldsymbol{X}(r)\|_{1} d r \tag{56}
\end{equation*}
$$

where $\boldsymbol{X}(r)$ is the state of the system at time $r \in \mathbb{R}_{+}$, immediately after the arrival/departure and just before the action is taken. In [38], it is argued that from uniformization [39] and sampling the continuous-time Markov process at a rate of $\lambda+\theta_{1}^{*}+\theta_{2}^{*}$, a discrete-time Markov chain is obtained, which converts the original continuous-time problem shown in (56) to an equivalent discrete-time problem as below:

$$
\begin{equation*}
\inf _{\pi \in \Pi} \limsup _{T \rightarrow \infty} \frac{1}{T} \int_{0}^{T}\|\boldsymbol{X}(r)\|_{1} d r=\inf _{\pi \in \Pi} \limsup _{T \rightarrow \infty} \frac{1}{T} \sum_{i=0}^{T-1}\|\boldsymbol{X}(i)\|_{1} \tag{57}
\end{equation*}
$$

To obtain a uniform sampling rate of $\lambda+\theta_{1}^{*}+\theta_{2}^{*}$, the continuous-time system is sampled at arrivals, real and dummy customer departures. In [38], it is further shown that the optimal policy that achieves the infimum in (57) is a threshold policy $\pi_{t}$ with the optimal finite threshold $t(\theta) \in \mathbb{N}$, with the policy defined as below:

$$
\pi_{t}(\boldsymbol{x})=\left\{\begin{array}{l}
\text { h if }\left\{x_{0}=0\right\} \text { or }\left\{\|\boldsymbol{x}\|_{1} \leq t, x_{1}=1\right\} \text { or }\left\{x_{1}=x_{2}=1\right\} \\
1 \text { if }\left\{x_{0} \geq 1, x_{1}=0\right\} \\
2 \text { if }\left\{x_{0} \geq 1,\|\boldsymbol{x}\|_{1} \geq t+1, x_{1}=1, x_{2}=0\right\}
\end{array}\right.
$$

note that action $b$ is not used. Policy $\pi_{t}$ assigns a job to the faster (first) server whenever there is a job waiting in the queue and the first server is available. In contrast, $\pi_{t}$ dispatches a job to the second server only if the number of jobs in the system are greater than threshold $t$ and the second server is available. If neither of these conditions hold, no action or $h$ is taken. Consequently, we can restrict the set of all policies $\Pi$ in (57) to the set $\Pi_{t}$, which is the set of all possible threshold policies corresponding to some $t \in \mathbb{N}$.
In the rest of this subsection, our aim is to show that Assumptions 1-5 are satisfied for the discrete-time Markov process obtained by uniformization of the described queueing system and hence, conclude that Algorithm 11 can be used to learn the unknown service rate vector $\boldsymbol{\theta}^{*}$ with the expected regret of order $\tilde{O}(\sqrt{T})$.
Assumption 1. Cost function is given as $c(\boldsymbol{x}, a)=\|\boldsymbol{x}\|_{1}$, which satisfies Assumption 1 with $f_{c}(\boldsymbol{x})=x_{0}+x_{1}+x_{2}$ and $K=r=1$.

Assumption 2. For any state-action pair $(\boldsymbol{x}, a)$ and $\theta \in \Theta$, we have $P_{\theta}(A(\boldsymbol{x}) ; \boldsymbol{x}, a)=0$ where $A(\boldsymbol{x})=\left\{\boldsymbol{y} \in \mathcal{X}:\left|\|\boldsymbol{y}\|_{1}-\|\boldsymbol{x}\|_{1}\right|>1\right\}$; thus, Assumption 2 holds with $h=1$.
Assumption 3. Consider a queueing system with parameter $\theta$ following threshold policy $\pi_{t}$ for some $t \in \mathbb{N}$. The uniformized discrete-time Markov chain is irreducible and aperiodic on a subset of state space given as $\mathcal{X}_{t}=\mathcal{X} \backslash(\{(i, 0,0): i \geq \min (t, 2)\} \cup\{(0,1,1)\})$. In [38], it is proved that for every $t$, the chain consists of a single positive recurrent class and the corresponding average number of customers, depicted by $J^{t}(\theta)$, is calculated. Moreover, it is shown that for every $\theta \in \Theta$ the optimal threshold $t(\theta)$ can be numerically found as the smallest $i \in \mathbb{N}$ for which $J^{i}(\theta)<J^{i+1}(\theta)$. Define the set $T^{*}$ as the set of all optimal thresholds corresponding to at least one $\theta \in \Theta$, or

$$
T^{*}=\{t: t=t(\theta) \text { for } \theta \in \Theta\} .
$$

Remark 7. There is a discrepancy between the class of MDPs defined in this section and in Section 2 as in the former the MDPs are not irreducible in the whole state space $\mathcal{X}$. Specifically, for every Markov process generated by a queueing system with parameter $\theta$ following threshold policy $\pi_{t}$, irreducibility holds on $\mathcal{X}_{t} \subset \mathcal{X}$. Nevertheless, the results of Section 4 are valid as starting from state $(0,0)$, the visited states are positive recurrent; see Remark 4

In the following proposition, we verify the geometric ergodicity of the discrete-time chain governed by any parameter $\theta \in \Theta$ and obtained by following any threshold policy $\pi_{t}$ for $t \in T^{*}$; proof is given in Appendix F. 1

Proposition 1. The discrete-time Markov process obtained from the queueing system governed by parameter $\theta=\left(\theta_{1}, \theta_{2}\right) \in \Theta$ and following threshold policy $\pi_{t}$ for some $t \in T^{*}$ is geometrically ergodic. Equivalently, the following holds

$$
\Delta V_{\theta, t}^{g}(\boldsymbol{x}) \leq-\left(1-\gamma_{\theta, t}^{g}\right) V_{\theta, t}^{g}(\boldsymbol{x})+b_{\theta, t}^{g} \mathbb{I}_{C_{\theta, t}^{g}}(\boldsymbol{x}), \quad \boldsymbol{x} \in \mathcal{X}_{t}
$$

for

$$
\begin{align*}
& V_{\theta, t}^{g}(\boldsymbol{x})=\exp \left(-\log (1-\delta)\|\boldsymbol{x}\|_{1}\right) \\
& C_{\theta, t}^{g}=\left\{\left(x_{0}, x_{1}, 0\right): x_{0}<t\right\} \cup\{(0,0,1)\}  \tag{58}\\
& b_{\theta, t}^{g}=\max _{\boldsymbol{x} \in C_{\theta, t}^{g}} \exp \left(-\log (1-\delta)\left(\|\boldsymbol{x}\|_{1}+1\right)\right)  \tag{59}\\
& \gamma_{\theta, t}^{g}=\frac{1}{2}-\frac{1}{2\left(\theta_{1}+\theta_{2}+\lambda\right)}\left(\left(\theta_{1}+\theta_{2}\right)(1-\delta)+\lambda(1-\delta)^{-1}\right) . \tag{60}
\end{align*}
$$

Having described all the terms explicitly, we verify the rest of the conditions of Assumption 3, which lead to uniform (over model class) upper-bounds on the moments of hitting time to $0^{d}$ as follows:

1. From (60), $\sup _{\theta \in \Theta, t \in T^{*}} \gamma_{\theta, t}^{g} \leq 1 / 2<1$.
2. From (58), we can see that state $(0,0)$ belongs to $C_{\theta, t}^{g}$ for all $\theta \in \Theta$ and $t \in T^{*}$. In order for $C_{*}^{g}=\cup_{\theta \in \Theta, t \in T^{*}} C_{\theta, t}^{g}$ to be a finite set, the supremum of the optimal threshold $t(\theta)$ over $\Theta$ should be finite. In [37] with service rate vector $\left(\theta_{1}, \theta_{2}\right)$, it is shown that the optimal threshold is bounded above by $\sqrt{2} \theta_{1} / \theta_{2}$, which further gives

$$
\begin{equation*}
t(\theta) \leq \sqrt{2} \frac{\theta_{1}}{\theta_{2}} \leq \sqrt{2} R \tag{61}
\end{equation*}
$$

Thus, $\sup _{\theta \in \Theta} t(\theta) \leq \sqrt{2} R$, which is finite. To confirm a uniform upper bound for $b_{\theta, t}^{g}$, we note that from (59),

$$
\sup _{\theta \in \Theta, t \in T^{*}} b_{\theta, t}^{g}=\frac{2-\delta}{1-\delta} \max _{x \in C_{*}^{g}} \exp \left(-\log (1-\delta)\|x\|_{1}\right)
$$

which is finite as $\left|C_{*}^{g}\right|<\infty$.
Assumption 4. To find an upper bound on the second moment of hitting times, we verify Assumption 4 and show that there exists a finite set $C_{\theta, t}^{p}$, constants $\beta_{\theta, t}^{p}, b_{\theta, t}^{p}>0, r /(r+1) \leq \alpha_{\theta, t}^{p}<1$, and a function $V_{\theta, t}^{p}: \mathcal{X}_{t} \rightarrow[1,+\infty)$ satisfying

$$
\begin{equation*}
\Delta V_{\theta, t}^{p}(\boldsymbol{x}) \leq-\beta_{\theta, t}^{p}\left(V_{\theta, t}^{p}(\boldsymbol{x})\right)^{\alpha_{\theta, t}^{p}}+b_{\theta, t}^{p} \mathbb{I}_{C_{\theta, t}^{p}}(\boldsymbol{x}), \quad \boldsymbol{x} \in \mathcal{X}_{t} . \tag{62}
\end{equation*}
$$

Proposition 2. The discrete-time Markov process obtained from the queueing system governed by parameter $\theta=\left(\theta_{1}, \theta_{2}\right) \in \Theta$ and following threshold policy $\pi_{t}$ for some $t \in T^{*}$ is polynomially
ergodic. This is true because 62) holds for

$$
\begin{align*}
& V_{\theta, t}^{p}(\boldsymbol{x})=\|\boldsymbol{x}\|_{1}^{2}  \tag{63}\\
& C_{\theta, t}^{p}=\left\{\left(x_{0}, x_{1}, 0\right): x_{0}<t\right\} \cup\left\{\left(x_{0}, x_{1}, x_{2}\right): x_{0}<\frac{2 \lambda}{\theta_{1}+\theta_{2}-\lambda}, x_{1}+x_{2} \geq 1\right\}  \tag{64}\\
& \left.b_{\theta, t}^{p}=\max _{\boldsymbol{x} \in C_{\theta, t}^{p}}\left(\|\boldsymbol{x}\|_{1}+1\right)^{2}\right)  \tag{65}\\
& \beta_{\theta, t}^{p}=1-\frac{2 \lambda}{\theta_{1}+\theta_{2}+\lambda}  \tag{66}\\
& \alpha_{\theta, t}^{p}=\frac{1}{2} \tag{67}
\end{align*}
$$

Proof of Proposition 2 is given in Appendix F.2. We define the normalized rates as $\tilde{\lambda}=\frac{\lambda}{\lambda+\theta_{1}+\theta_{2}}$ and $\tilde{\theta}_{i}=\frac{\theta_{i}}{\lambda+\theta_{1}+\theta_{2}}$, for $i=1,2$. From the choice of parameter space $\Theta$, we have $\tilde{\lambda} \leq 0.5-0.5 \delta$, $\tilde{\theta}_{1}+\tilde{\theta}_{2} \geq 0.5+0.5 \delta$, and $\tilde{\theta}_{1} \geq 0.25+0.25 \delta$. We verify the remaining conditions of Assumption 4 as follows:

1. From (63), the first condition holds with $r_{*}^{p}=2$ and $s_{*}^{p}=2$.
2. From (64), we can see that state $(0,0)$ belongs to $C_{\theta, t}^{p}$ for all $\theta \in \Theta$ and $t \in T^{*}$. Furthermore,

$$
\sup _{\theta \in \Theta, t \in T^{*}} \frac{2 \lambda}{\theta_{1}+\theta_{2}-\lambda} \leq \frac{1-\delta}{\delta}
$$

which follows from the stability condition $\tilde{\lambda} \leq 0.5-0.5 \delta$. Thus, from the definition of $C_{\theta, t}^{p}$ in (64), and the fact that $\sup _{\theta \in \Theta} t(\theta) \leq \sqrt{2} R$ as argued in in 61, $C_{*}^{p}=\cup_{\theta \in \Theta, t \in T^{*}} C_{\theta, t}^{p}$ is a finite set. We also note that $\sup _{\theta \in \Theta, t \in T^{*}} b_{\theta, t}^{p}$ is finite as $\left|C_{*}^{p}\right|<\infty$. It remains to show that $\inf _{\theta \in \Theta, t \in T^{*}} \beta_{\theta, t}^{p}$ is positive, which is equivalent to verifying that $\sup _{\theta \in \Theta, t \in T^{*}} \tilde{\lambda}<1 / 2$, which follows from the stability condition $\tilde{\lambda} \leq 0.5-0.5 \delta$.
3. We need to show that $K_{\theta, t}(\boldsymbol{x}):=\sum_{n=0}^{\infty} 2^{-n-2}\left(P_{\theta}^{t}\right)^{n}\left(\boldsymbol{x}, 0^{d}\right)$ is strictly bounded away from zero. We notice that from any non-zero state $\boldsymbol{x}$, the queueing system hits $0^{d}$ in $\|\boldsymbol{x}\|_{1}$ transitions only if all transitions are real departures. Hence,

$$
\begin{aligned}
K_{\theta, t}(\boldsymbol{x}) & \geq 2^{-\|\boldsymbol{x}\|_{1}-2}\left(P_{\theta}^{t}\right)^{\|\boldsymbol{x}\|_{1}}\left(\boldsymbol{x}, 0^{d}\right) \\
& \geq 2^{-\|\boldsymbol{x}\|_{1}-2}\left(\tilde{\theta}_{1}\right)^{\|\boldsymbol{x}\|_{1}}\left(\tilde{\theta}_{2}\right)^{\|\boldsymbol{x}\|_{1}} \\
& \geq 2^{-\|\boldsymbol{x}\|_{1}-2} R^{-\|\boldsymbol{x}\|_{1}}\left(\tilde{\theta}_{1}\right)^{2\|\boldsymbol{x}\|_{1}} \\
& \geq 2^{-\|\boldsymbol{x}\|_{1}-2} R^{-\|\boldsymbol{x}\|_{1}}\left(\frac{1}{4}+\frac{\delta}{4}\right)^{2\|\boldsymbol{x}\|_{1}}
\end{aligned}
$$

where the third and fourth inequalities follow from the definition of $\Theta$ in (55). Thus, the infimum of $K_{\theta, t}(\boldsymbol{x})$ over the finite set $C_{*}^{p}$ and sets $\Theta$ and $T^{*}$ is strictly greater than zero.

Assumption 5. We finally verify Assumption 5, which asserts that $\sup _{\theta \in \Theta} J(\theta)$ is finite. We have

$$
J(\theta)=\mathbb{E}_{\boldsymbol{X} \sim \mu_{\theta, t(\theta)}}[c(\boldsymbol{X})]=\mathbb{E}_{\boldsymbol{X} \sim \mu_{\theta, t(\theta)}}\left[\|\boldsymbol{X}\|_{1}\right]=\mathbb{E}_{\boldsymbol{X} \sim \mu_{\theta, t(\theta)}}\left[\sqrt{V_{\theta, t(\theta)}^{p}(\boldsymbol{X})}\right]
$$

where $\mu_{\theta, t(\theta)}$ is the stationary distribution of the discrete-time process governed by parameter $\theta$ and following the optimal policy according to $\theta$. From (62) and [43, Theorem 14.3.7],

$$
\mu_{\theta, t(\theta)}\left(\sqrt{V_{\theta, t(\theta)}^{p}(\boldsymbol{X})}\right) \leq \frac{b_{*}^{p}}{\beta_{*}^{p}}
$$

which is finite from the the previously verified assumption. Consequently,

$$
\sup _{\theta \in \Theta} J(\theta) \leq \frac{b_{*}^{p}}{\beta_{*}^{p}}<\infty
$$

## E. 2 Model 2: Two heterogeneous parallel queues

We consider two parallel queues with infinite buffers, each with its own single server, and unknown service rate vector $\boldsymbol{\theta}^{*}=\left(\theta_{1}^{*}, \theta_{2}^{*}\right)$, shown in Figure 2 b The service rate vector $\boldsymbol{\theta}^{*}$ is sampled from the prior distribution $\nu_{0}$ defined on the space $\Theta$ given as

$$
\begin{equation*}
\Theta=\left\{\left(\theta_{1}, \theta_{2}\right) \in \mathbb{R}_{+}^{2}: \frac{\lambda}{\theta_{1}+\theta_{2}} \leq \frac{1-\delta}{1+\delta}, 1 \leq \frac{\theta_{1}}{\theta_{2}} \leq R\right\} \tag{68}
\end{equation*}
$$

for fixed $\delta \in(0,0.5)$ and $R \geq 1$, which ensures the stability of the queueing system. Consider the discrete-time $\operatorname{MDP}\left(\mathcal{X}, \mathcal{A}, P_{\boldsymbol{\theta}^{*}}, c\right)$ obtained by sampling the queueing system at the Poisson arrival sequence. The countably infinite state space $\mathcal{X}$ is defined as below

$$
\mathcal{X}=\left\{\boldsymbol{x}=\left(x_{1}, x_{2}\right): x_{i} \in \mathbb{N} \cup\{0\}\right\},
$$

where the state of the system is the number of jobs in the server-queue pair $i$ just before an arrival. Furthermore, the action space $\mathcal{A}$ is equal to

$$
\mathcal{A}=\{1,2\},
$$

where action $i \in \mathcal{A}$ indicates the arrival dispatched to queue $i$. The unbounded cost function $c$ : $\mathcal{X} \times \mathcal{A} \rightarrow \mathbb{N} \cup\{0\}$ is defined as the total number of jobs in the queueing system, i.e., $c(\boldsymbol{x}, a)=\|\boldsymbol{x}\|_{1}$. For every $\omega \in \mathbb{R}_{+}$, we define policy $\pi_{\omega}: \mathcal{X} \rightarrow \mathcal{A}$, which routes the arrival according to the weighted queue lengths, as

$$
\pi_{\omega}(\boldsymbol{x})=\arg \min \left(1+x_{1}, \omega\left(1+x_{2}\right)\right)
$$

where the tie is broken in favor of the first server. We also define policy class $\tilde{\Pi}$ as the set of policies $\pi_{\omega}$ such that $\omega$ belongs to a compact interval; in other words,

$$
\tilde{\Pi}=\left\{\pi_{\omega} ; \omega \in\left[\frac{1}{c_{R} R}, c_{R} R\right]\right\}
$$

where $R$ is defined in 68 and $c_{R} \geq 1$. We aim to minimize the infinite-horizon average cost in the policy class $\tilde{\Pi}$, that is,

$$
\begin{equation*}
J(\theta)=\inf _{\pi \in \tilde{\Pi}} \limsup _{T \rightarrow \infty} \frac{1}{T} \mathbb{E}\left[\sum_{t=1}^{T} c(\boldsymbol{X}(t), A(t))\right] \tag{69}
\end{equation*}
$$

where $\boldsymbol{X}(t)=\left(X_{1}(t), X_{2}(t)\right)$ is the occupancy vector of the queueing system just before arrival $t$. Even with the controlled Markov process transition kernel fully-specified (by the values of the arrival rate and the two service rates), the optimal policy ${ }^{1}$ that satisfies (69) in policy class $\tilde{\Pi}$ is not known except when $\theta_{1}=\theta_{2}$ where the optimal value is $\omega=1$, and so, to learn it, we will use Proximal Policy Optimization for countable state-space controlled Markov processes as developed in [18]. Note that [18] requires full knowledge of the controlled Markov process, which holds in our learning scheme since we use the parameters sampled from the posterior for determining the policy at the beginning of each episode. Furthermore, for each policy in the set of applicable policies $\tilde{\Pi}$, [18] also requires that the resulting Markov process be geometrically ergodic, which we will establish below.
Proposition 3. The discrete-time Markov process obtained from the queueing system governed by parameter $\theta=\left(\theta_{1}, \theta_{2}\right) \in \Theta$ and following policy $\pi_{\omega} \in \tilde{\Pi}$ is geometrically ergodic. Equivalently, the following holds

$$
\begin{equation*}
\Delta V_{\theta, \omega}^{g}(\boldsymbol{x}) \leq-\left(1-\gamma_{\theta, \omega}^{g}\right) V_{\theta, \omega}^{g}(\boldsymbol{x})+b_{\theta, \omega}^{g} \mathbb{I}_{C_{\theta, \omega}^{g}}(\boldsymbol{x}), \quad \boldsymbol{x} \in \mathcal{X} \tag{70}
\end{equation*}
$$

[^0]\[

$$
\begin{align*}
& V_{\theta, \omega}^{g}(\boldsymbol{x})=\frac{\omega}{\omega+1} \exp \left(a_{\theta, \omega}^{g} \frac{x_{1}+1}{\omega}\right)+\frac{1}{\omega+1} \exp \left(a_{\theta, \omega}^{g}\left(x_{2}+1\right)\right) \\
& a_{\theta, \omega}^{g}=\min \left(\omega \log (1+\delta), \log (1+\delta), \omega \log \frac{1-0.5 \delta}{1-\delta}, \log \frac{1-0.5 \delta}{1-\delta}, \frac{\delta\left(1-\delta^{2}\right)}{4 c_{R} R(1-0.5 \delta)}\right)  \tag{71}\\
& C_{\theta, \omega}^{g}=\left\{\left(x_{1}, x_{2}\right) \in \mathcal{X}: x_{i} \leq \max \left(x_{i, \theta, \omega}^{g_{j}}, 0\right), i, j=1,2\right\}  \tag{72}\\
& b_{\theta, \omega}^{g}=\max _{x \in C_{\theta, \omega}^{g}}\left(\frac{2 \omega}{\omega+1} \exp \left(a_{\theta, \omega}^{g} \frac{x_{1}+2}{\omega}\right)+\frac{2}{\omega+1} \exp \left(a_{\theta, \omega}^{g}\left(x_{2}+2\right)\right)\right)  \tag{73}\\
& \gamma_{\theta, \omega}^{g}=\frac{1}{2}+\frac{1}{2} \max \left(\zeta_{1, \theta, \omega}, \zeta_{2, \theta, \omega}, \frac{\zeta_{1, \theta, \omega} \omega}{1+\omega} \exp \left(\frac{a_{\theta, \omega}^{g}}{\omega}\right)+\frac{\zeta_{2, \theta, \omega}}{1+\omega}, \frac{\zeta_{1, \theta, \omega} \omega}{1+\omega}+\frac{\zeta_{2, \theta, \omega}}{1+\omega} \exp \left(a_{\theta, \omega}^{g}\right)\right) \tag{74}
\end{align*}
$$
\]

and problem-dependent constants $x_{i, \theta, \omega}^{g_{j}}$ and $\zeta_{i, \theta, \omega}$ for $i, j=1,2$.

Proof of Proposition 3 is given in Appendix F. 3 In the rest of this subsection, our aim is to show that Assumptions 1-5 are satisfied for the discrete-time MDP and conclude that Algorithm 1 can be used to learn the unknown service rate vector $\boldsymbol{\theta}^{*}$ with expected regret of order $\tilde{O}(\sqrt{T})$.
Assumption 1. Cost function is given as $c(\boldsymbol{x}, a)=\|\boldsymbol{x}\|_{1}$, which satisfies Assumption 1 with $f_{c}(\boldsymbol{x})=x_{0}+x_{1}+x_{2}$ and $K=r=1$.
Assumption 2. For any state-action pair $(\boldsymbol{x}, a)$ and $\theta \in \Theta$, we have $P_{\theta}(A(\boldsymbol{x}) ; \boldsymbol{x}, a)=0$ where $A(\boldsymbol{x})=\left\{\boldsymbol{y} \in \mathcal{X}:\|\boldsymbol{y}\|_{1}-\|\boldsymbol{x}\|_{1}>1\right\}$; thus, the MDP is skip-free to the right with $h=1$. Moreover, from any $(\boldsymbol{x}, a)$, the finite set $\left\{\boldsymbol{y} \in \mathcal{X}:\|\boldsymbol{y}\|_{1} \leq\|\boldsymbol{x}\|_{1}+1\right\}$ is only accessible in one step; thus, Assumption 2 holds.
Assumption 3. In Proposition 3, we verified the geometric ergodicity of the discrete-time chain governed by parameter $\theta=\left(\theta_{1}, \theta_{2}\right) \in \Theta$ and following policy $\pi_{\omega} \in \tilde{\Pi}$ and thus, it only remains to verify the uniform model conditions. We define the normalized rates as $\tilde{\lambda}=\frac{\lambda}{\lambda+\theta_{1}+\theta_{2}}$ and $\tilde{\theta}_{i}=\frac{\theta_{i}}{\lambda+\theta_{1}+\theta_{2}}$, for $i=1,2$. From the choice of parameter space $\Theta$, we have $\tilde{\lambda} \leq 0.5-0.5 \delta$, $\tilde{\theta}_{1}+\tilde{\theta}_{2} \geq 0.5+0.5 \delta$, and $\tilde{\theta}_{1} \geq 0.25+0.25 \delta$.

1. We first argue that $\zeta_{1, \theta, \omega}$ is bounded away from 1 as follows

$$
\begin{aligned}
1-\zeta_{1, \theta, \omega} & =1-\frac{\frac{\lambda}{\theta_{1}+\lambda}}{1-\exp \left(-\frac{a_{\theta, \omega}^{g}}{\omega}\right) \frac{\theta_{1}}{\theta_{1}+\lambda}}=\frac{\frac{\theta_{1}}{\theta_{1}+\lambda}\left(1-\exp \left(-\frac{a_{\theta, \omega}^{g}}{\omega}\right)\right)}{1-\exp \left(-\frac{a_{\theta, \omega}^{g}}{\omega}\right) \frac{\theta_{1}}{\theta_{1}+\lambda}} \\
& \geq \frac{\theta_{1}}{\theta_{1}+\lambda}\left(1-\exp \left(-\frac{a_{\theta,\left(c_{R} R\right)^{-1}}^{g}}{c_{R} R}\right)\right)>\tilde{\theta}_{1}\left(1-\exp \left(-\frac{\left.\left.a_{\theta,\left(c_{R} R\right)^{-1}}^{c_{R} R}\right)\right)}{}\right.\right. \\
& >(0.25+0.25 \delta)\left(1-\exp \left(-\frac{a_{\theta,\left(c_{R} R\right)^{-1}}^{g}}{c_{R} R}\right)\right)
\end{aligned}
$$

where the first line follows from the definition of $\zeta_{1, \theta, \omega}$ in Appendix F.3, the second line from 71) and the definition of policy class $\tilde{\Pi}$. As $a_{\theta, \omega}^{g}$ does not depend on $\theta, \sup _{\theta \in \Theta, \omega \in\left[\frac{1}{c_{R} R}, c_{R} R\right]} \zeta_{1, \theta, \omega}<1$. Furthermore, by similar arguments it can be shown that $\zeta_{2, \theta, \omega}$ is bounded away from 1. We next argue that $\frac{\zeta_{1, \theta, \omega} \omega}{1+\omega} \exp \left(\frac{a_{\theta, \omega}^{g}}{\omega}\right)+\frac{\zeta_{2, \theta, \omega}}{1+\omega}$ is bounded away from 1 using an upper bound found in

Appendix F. 3 as below,

$$
\begin{align*}
& 1-\frac{\zeta_{1, \theta, \omega} \omega}{1+\omega} \exp \left(\frac{a_{\theta, \omega}^{g}}{\omega}\right)-\frac{\zeta_{2, \theta, \omega}}{1+\omega} \\
& \geq 1-\frac{\frac{\lambda}{1+\omega}\left(\omega+a_{\theta, \omega}^{g} \zeta_{4}\right)}{\lambda+\frac{\theta_{1} a_{\theta, \omega}^{g} \zeta_{3}}{\omega}}-\frac{\frac{\lambda}{1+\omega}}{\lambda+\theta_{2} a_{\theta, \omega}^{g} \zeta_{3}} \\
& =\frac{a_{\theta, \omega}^{g}\left(-a_{\theta, \omega}^{g} \zeta_{3} \theta_{2}\left(\lambda \zeta_{4}-\frac{\zeta_{3} \theta_{1}(1+\omega)}{\omega}\right)+\lambda \zeta_{3}\left(\theta_{1}+\theta_{2}\right)-\lambda^{2} \zeta_{4}\right)}{(1+\omega)\left(\lambda+\theta_{1} a_{\theta, \omega}^{g} \zeta_{3} \omega^{-1}\right)\left(\lambda+\theta_{2} a_{\theta, \omega}^{g} \zeta_{3}\right)} \\
& >\frac{\left(\zeta_{3} a_{\theta, \omega}^{g}\right)^{2} \tilde{\theta}_{1} \tilde{\theta}_{2}}{\omega\left(\tilde{\lambda}+\tilde{\theta}_{1} a_{\theta, \omega}^{g} \zeta_{3} \omega^{-1}\right)\left(\tilde{\lambda}+\tilde{\theta}_{2} a_{\theta, \omega}^{g} \zeta_{3}\right)} \\
& >\frac{\left(\zeta_{3} a_{\left.\theta,\left(c_{R} R\right)^{-1}\right)^{2}(0.25+0.25 \delta)^{2}}^{c_{R} R^{2}\left(1+c_{R} R \zeta_{3} a_{\theta, c_{R} R}^{g}\right)^{2}}\right.}{} \tag{75}
\end{align*}
$$

where $\zeta_{3}=(1+\delta)^{-1}, \zeta_{4}=\frac{1-0.5 \delta}{1-\delta}$, and we have used the arguments of Appendix F. 3 and the definition of $\Theta$. Using a similar argument, we can show that $\frac{\zeta_{1, \theta, \omega} \omega}{1+\omega}+\frac{\zeta_{2, \theta, \omega}}{1+\omega} \exp \left(a_{\theta, \omega}^{g}\right)$ is bounded away from one, and finally, we conclude that $\sup _{\theta \in \Theta, \omega \in\left[\frac{1}{c_{R} R}, c_{R} R\right]} \gamma_{\theta, \omega}^{g}<1$.
2. From (72), we can see that state $(0,0)$ belongs to $C_{\theta, \omega}^{g}$ for all $\theta \in \Theta$ and $\omega \in\left[\frac{1}{c_{R} R}, c_{R} R\right]$. In order for $C_{*}^{g}$ to be a finite set, the supremum of $x_{i, \theta, \omega}^{g_{j}}$ over $\Theta$ and $\tilde{\Pi}$ should be finite. From the definition of $x_{1, \theta, \omega}^{g_{1}}$ in Appendix F. 3 .

$$
\begin{aligned}
x_{1, \theta, \omega}^{g_{1}} & =\frac{\omega}{a_{\theta, \omega}^{g}} \log \frac{\left(c_{R} R+1\right) \exp \left(c_{R} R a_{\theta, \omega}^{g}\right)}{(\omega+1) \gamma_{\theta, \omega}^{g}-\omega \zeta_{1, \theta, \omega} \exp \left(\frac{a_{\theta, \omega}^{g}}{\omega}\right)-\zeta_{2, \theta, \omega}} \\
& \leq \frac{c_{R} R}{a_{\theta,\left(c_{R} R\right)^{-1}}^{g}} \log \frac{\left(c_{R} R+1\right) \exp \left(c_{R} R a_{\theta, c_{R} R}^{g}\right)}{(\omega+1) \gamma_{\theta, \omega}^{g}-\omega \zeta_{1, \theta, \omega} \exp \left(\frac{a_{\theta, \omega}^{g}}{\omega}\right)-\zeta_{2, \theta, \omega}}
\end{aligned}
$$

and we can derive a lower bound for the denominator from (75). Similarly, we can show that $\sup _{\theta \in \Theta, \omega \in\left[\frac{1}{c_{R^{R}}}, c_{R} R\right]} x_{2, \theta, \omega}^{g_{2}}$ is finite. We next find a uniform upper bound for $x_{2, \theta, \omega}^{g_{1}}$ from Appendix F. 3 .

$$
\begin{aligned}
& x_{2, \theta, \omega}^{g_{1}} \\
& =\frac{1}{a_{\theta, \omega}^{g}} \log \frac{\left(c_{R} R+1\right) \exp \left(c_{R} R a_{\theta, \omega}^{g}\right)+\omega \exp \left(a_{\theta, \omega}^{g} \frac{x_{1, \theta, \omega}^{g_{1}}+1}{\omega}\right)\left(\zeta_{1, \theta, \omega} \exp \left(\frac{a_{\theta, \omega}^{g}}{\omega}\right)-\gamma_{\theta, \omega}^{g}\right)}{\gamma_{\theta, \omega}^{g}-\zeta_{2, \theta, \omega}} \\
& \leq \frac{1}{a_{\theta,\left(c_{R} R\right)^{-1}}^{g}} \log \frac{\left(2 c_{R} R+1\right) \exp \left(c_{R} R a_{\theta, c_{R} R}^{g}\left(x_{1, \theta, \omega}^{g_{1}}+2\right)\right)}{1-\gamma_{\theta, \omega}^{g}}
\end{aligned}
$$

which is uniformly bounded as $\gamma_{\theta, \omega}^{g}$ is unformly bounded away from 1 and the second line follows from (74) and the fact that $\gamma_{\theta, \omega}^{g}-\zeta_{2, \theta, \omega} \geq 1-\gamma_{\theta, \omega}^{g}$. Arguments verifying the finiteness of the supremum of $x_{1, \theta, \omega}^{g_{2}}$ follow similarly, and we conclude that $\left|C_{*}^{g}\right|<\infty$. To confirm a uniform upper bound for $b_{\theta, \omega}^{g}$, we note that from (73),

$$
\sup _{\theta \in \Theta, \omega \in\left[\frac{1}{c_{R} R}, c_{R} R\right]} b_{\theta, \omega}^{g} \leq \max _{x \in C_{*}^{g}}\left(2 \exp \left(c_{R} R a_{\theta, c_{R} R}^{g}\left(x_{1}+2\right)\right)+2 \exp \left(a_{\theta, c_{R} R}^{g}\left(x_{2}+2\right)\right)\right)
$$

which is finite as $a_{\theta, c_{R} R}^{g}$ is independent of the choice of $\theta$ and $\left|C_{*}^{g}\right|<\infty$.
Assumption 4. We next verify Assumption 4 and show that there exists a finite set $C_{\theta, \omega}^{p}$, constants $\beta_{\theta, \omega}^{p}, b_{\theta, \omega}^{p}>0, r /(r+1) \leq \alpha_{\theta, \omega}^{p}<1$, and a function $V_{\theta, \omega}^{p}: \mathcal{X} \rightarrow[1,+\infty)$ satisfying

$$
\begin{equation*}
\Delta V_{\theta, \omega}^{p}(\boldsymbol{x}) \leq-\beta_{\theta, \omega}^{p}\left(V_{\theta, \omega}^{p}(\boldsymbol{x})\right)^{\alpha_{\theta, \omega}^{p}}+b_{\theta, t}^{p} \mathbb{I}_{C_{\theta, \omega}^{p}}(\boldsymbol{x}), \quad \boldsymbol{x} \in \mathcal{X} \tag{76}
\end{equation*}
$$

Proposition 4. The discrete-time Markov process obtained from the queueing system governed by parameter $\theta=\left(\theta_{1}, \theta_{2}\right) \in \Theta$ and following policy $\pi_{\omega} \in \tilde{\Pi}$ is polynomially ergodic. This follow because (76) holds for

$$
\begin{align*}
& V_{\theta, \omega}^{p}(\boldsymbol{x})=\frac{x_{1}^{2}}{\omega}+x_{2}^{2},  \tag{77}\\
& C_{\theta, \omega}^{p}=\left\{\left(x_{1}, x_{2}\right) \in \mathcal{X}: x_{i} \leq\left(16 c_{R}^{2} R^{3-i}+101 c_{R} R\right) \frac{\lambda+\theta_{i}}{\theta_{i}}, i=1,2\right\},  \tag{78}\\
& \beta_{\theta, \omega}^{p}=\min \left(\frac{\theta_{2}}{2\left(\theta_{2}+\lambda\right) \sqrt{\omega+1}}, \frac{\theta_{1}+\theta_{2}-\lambda}{\left(\theta_{1}+\theta_{2}+\lambda\right) \sqrt{\omega+1}}, \frac{\theta_{2}}{2\left(\theta_{2}+\lambda\right)}, \frac{\theta_{1}}{2\left(\theta_{1}+\lambda\right) \sqrt{\omega}}\right),  \tag{79}\\
& b_{\theta, \omega}^{p}=\left(\beta_{\theta, \omega}^{p}+1\right) \max _{x \in C_{\theta, \omega}^{p}}\left(\frac{\left(x_{1}+1\right)^{2}}{\omega}+\left(x_{2}+1\right)^{2}\right),  \tag{80}\\
& \alpha_{\theta, \omega}^{p}=\frac{1}{2} . \tag{81}
\end{align*}
$$

Proof of Proposition 4 is given in Appendix F.4. Next, we verify the remaining conditions of Assumption 4

1. From (77) and the fact that $\omega \in\left[\frac{1}{c_{R} R}, c_{R} R\right]$, the first condition holds with $r_{*}^{p}=2$ and $s_{*}^{p}=$ $\sup _{\theta \in \Theta, \omega \in\left[\frac{1}{c_{R} R}, c_{R} R\right]} s_{\theta, \omega}=c_{R} R+1$.
2. From (78), state $(0,0)$ belongs to $C_{\theta, \omega}^{p}$ for all $\theta \in \Theta$ and $\omega \in\left[\frac{1}{c_{R} R}, c_{R} R\right]$. Furthermore, for $i=1,2$,

$$
\begin{equation*}
\sup _{\theta \in \Theta, \omega \in\left[\frac{1}{c_{R}^{R}}, c_{R} R\right]} \frac{\lambda+\theta_{i}}{\theta_{i}} \leq \sup _{\theta \in \Theta} \frac{1}{\tilde{\theta}_{2}} \leq \sup _{\theta \in \Theta} \frac{R}{\frac{\theta_{1}}{\tilde{\theta}_{1}} \leq \frac{4 R}{1+\delta}, ., ~} \tag{82}
\end{equation*}
$$

which follows from the fact that $\theta_{1} \leq R \theta_{2}$ and $\tilde{\theta}_{1} \geq 0.25+0.25 \delta$. Thus, from the definition of $C_{\theta, \omega}^{p}$ in (78), $C_{*}^{p}=\cup_{\theta \in \Theta, \omega \in\left[\frac{1}{c_{R} R}, c_{R} R\right]} C_{\theta, \omega}^{p}$ is a finite set. We next verify that the infimum of $\beta_{\theta, \omega}^{p}$, found in (79), is positive. In (82), we showed that infimum of $\frac{\lambda+\theta_{i}}{\theta_{i}}$ over $\Theta$ is lower bounded by $\frac{1+\delta}{4}$. From this, the fact that $\omega$ belongs to a compact set, and $\theta_{1}+\theta_{2}+\lambda \geq \delta$, it follows that $\inf _{\theta \in \Theta, \omega \in\left[\frac{1}{c_{R^{R}}}, c_{R} R\right]} \beta_{\theta, \omega}^{p}>0$. Furthermore, it is easy to see that $\beta_{\theta, \omega}^{p} \leq \sqrt{c_{R} R}$. Hence, from 80),

$$
\begin{aligned}
\sup _{\theta \in \Theta, \omega \in\left[\frac{1}{c_{R} R}, c_{R} R\right]} b_{\theta, \omega}^{p} & =\sup _{\theta \in \Theta, \omega \in\left[\frac{1}{c_{R} R}, c_{R} R\right]}\left(\beta_{\theta, \omega}^{p}+1\right) \max _{x \in C_{\theta, \omega}^{p}}\left(\frac{\left(x_{1}+1\right)^{2}}{\omega}+\left(x_{2}+1\right)^{2}\right) \\
& \leq\left(\sqrt{c_{R} R}+1\right) \max _{x \in C_{*}^{p}}\left(C_{R} R\left(x_{1}+1\right)^{2}+\left(x_{2}+1\right)^{2}\right),
\end{aligned}
$$

which is finite as $\left|C_{*}^{p}\right|<\infty$.
3. We need to show that $K_{\theta, \omega}(\boldsymbol{x}):=\sum_{n=0}^{\infty} 2^{-n-2}\left(P_{\theta}^{\pi_{\omega}}\right)^{n}\left(\boldsymbol{x}, 0^{d}\right)$ is strictly bounded away from zero. We show this using the fact that from any state $\boldsymbol{x}$, the queueing system hits $(0,0)$ in one step with positive probability. Take $x_{i, \theta, \omega}=\max _{\boldsymbol{x} \in C_{\theta, \omega}} x_{i}$ for $i=1,2$. We have

$$
\begin{aligned}
\inf _{\theta \in \Theta, \omega \in\left[\frac{1}{c_{R} R}, c_{R} R\right]} \min _{\boldsymbol{x} \in C_{\theta, \omega}} K(\boldsymbol{x}) & \geq \inf _{\theta \in \Theta, \omega \in\left[\frac{1}{c_{R} R}, c_{R} R\right]} \min _{\boldsymbol{x} \in C_{\theta, \omega}} P\left(\boldsymbol{x}, 0^{d}\right) \\
& \geq \inf _{\theta \in \Theta, \omega \in\left[\frac{1}{c_{R} R}, c_{R} R\right]} P\left(\left(x_{1, \theta, \omega}, x_{2, \theta, \omega}\right), 0^{d}\right) .
\end{aligned}
$$

The infimum in the right-hand side of the above equation is attained for the minimum normalized service rates possible for each server, or $\tilde{\theta}_{1}=\frac{1+\delta}{4}$ and $\tilde{\theta}_{2}=\frac{1+\delta}{4 R}$. Therefore, the infimum of $K_{\theta, \omega}(\boldsymbol{x})$ over the finite set $C_{*}^{p}, \Theta$, and interval $\left[\frac{1}{c_{R} R}, c_{R} R\right]$ is strictly greater than zero.

Assumption 5. We finally verify that $\sup _{\theta \in \Theta} J(\theta)$ is finite. We first note that for $\boldsymbol{x}=\left(x_{1}, x_{2}\right)$,

$$
\left(x_{1}+x_{2}\right)^{2} \leq 2 \max \left(\omega^{*}(\theta), 1\right)\left(\frac{x_{1}^{2}}{\omega^{*}(\theta)}+x_{2}^{2}\right)=2 \max \left(\omega^{*}(\theta), 1\right) V_{\theta, \omega^{*}(\theta)}^{p}(\boldsymbol{x})
$$

From the above equation,

$$
\begin{aligned}
J(\theta) & =\mathbb{E}_{\boldsymbol{X} \sim \mu_{\theta, \omega^{*}(\theta)}}[c(\boldsymbol{X})] \\
& =\mathbb{E}_{\boldsymbol{X} \sim \mu_{\theta, \omega^{*}(\theta)}}\left[\|\boldsymbol{X}\|_{1}\right] \\
& \leq \sqrt{2 \max \left(\omega^{*}(\theta), 1\right)} \mathbb{E}_{\boldsymbol{X} \sim \mu_{\theta, \omega^{*}(\theta)}}\left[\sqrt{V_{\theta, \omega^{*}(\theta)}^{p}(\boldsymbol{X})}\right]
\end{aligned}
$$

where $\mu_{\theta, \omega^{*}(\theta)}$ is the stationary distribution of the discrete-time process governed by parameter $\theta$ and following the best in-class policy according to $\theta$, shown by $\pi_{\omega^{*}(\theta)}$. From [43], Theorem 14.3.7,

$$
\mu_{\theta, \omega^{*}(\theta)}\left(\sqrt{V_{\theta, \omega^{*}(\theta)}^{p}(\boldsymbol{X})}\right) \leq \frac{\sup _{\theta \in \Theta, \omega \in\left[\frac{1}{c_{R} R}, c_{R} R\right]} b_{\theta, \omega}^{p}}{\beta_{*}^{p}},
$$

which is finite from the the previous verified assumption. Thus,

$$
\sup _{\theta \in \Theta} J(\theta) \leq \frac{\sqrt{2 c_{R} R}\left(\sup _{\theta \in \Theta, \omega \in\left[\frac{1}{c_{R_{R}^{R}}}, c_{R} R\right]} b_{\theta, \omega}^{p}\right)}{\beta_{*}^{p}}<\infty .
$$

## F Proofs related to the queueing model examples

## F. 1 Proof of Proposition 1

Proof. We define the normalized rates as

$$
\begin{equation*}
\tilde{\lambda}=\frac{\lambda}{\lambda+\theta_{1}+\theta_{2}}, \quad \tilde{\theta}_{i}=\frac{\theta_{i}}{\lambda+\theta_{1}+\theta_{2}}, \tag{83}
\end{equation*}
$$

for $i=1,2$. From the choice of parameter space $\Theta$, we have $\tilde{\lambda} \leq 0.5-0.5 \delta, \theta_{1}+\theta_{2} \geq 0.5+0.5 \delta$, and $\theta_{1} \geq 0.25+0.25 \delta$. To prove geometric ergodicity, from the discussions of Section 2, it suffices to show that there exists a finite set $C_{\theta, t}^{g}$, constants $b_{\theta, t}^{g}>0, \gamma_{\theta, t}^{g} \in(0,1)$, and a function $V_{\theta, t}^{g}: \mathcal{X}_{t} \rightarrow[1,+\infty)$ satisfying

$$
\begin{equation*}
\Delta V_{\theta, t}^{g}(\boldsymbol{x}) \leq-\left(1-\gamma_{\theta, t}^{g}\right) V_{\theta, t}^{g}(\boldsymbol{x})+b_{\theta, t}^{g} \mathbb{I}_{C_{\theta, t}^{g}}(\boldsymbol{x}), \quad \boldsymbol{x} \in \mathcal{X}_{t} \tag{84}
\end{equation*}
$$

Take $V_{\theta, t}^{g}(\boldsymbol{x})=\exp \left(a_{\theta, t}^{g}\|\boldsymbol{x}\|_{1}\right)$ for some $a_{\theta, t}^{g}>0$. For $i \geq 1$ and $\boldsymbol{x}=(i, 1,1)$,

$$
P_{\theta}^{t} V_{\theta, t}^{g}(i, 1,1)=\tilde{\lambda} V_{\theta, t}^{g}(i+1,1,1)+\tilde{\theta}_{1} V_{\theta, t}^{g}(i, 0,1)+\tilde{\theta}_{2} V_{\theta, t}^{g}(i, 1,0)
$$

where $P_{\theta}^{t}$ is the corresponding transition kernel. Thus,

$$
\begin{aligned}
& P_{\theta}^{t} V_{\theta, t}^{g}(i, 1,1)-\left(1-\gamma_{\theta, t}^{g}\right) V_{\theta, t}^{g}(i, 1,1) \\
& =\tilde{\lambda} \exp \left(a_{\theta, t}^{g}(i+3)\right)+\left(\tilde{\theta}_{1}+\tilde{\theta}_{2}\right) \exp \left(a_{\theta, t}^{g}(i+1)\right)-\left(1-\gamma_{\theta, t}^{g}\right) \exp \left(a_{\theta, t}^{g}(i+2)\right) \\
& =\exp \left(a_{\theta, t}^{g}(i+1)\right)\left(\tilde{\lambda} \exp \left(2 a_{\theta, t}^{g}\right)+\tilde{\theta}_{1}+\tilde{\theta}_{2}-\left(1-\gamma_{\theta, t}^{g}\right) \exp \left(a_{\theta, t}^{g}\right)\right) .
\end{aligned}
$$

Take $\tilde{a}_{\theta, t}=\exp \left(a_{\theta, t}^{g}\right)$. We need to find $\tilde{a}_{\theta, t}>1$ and $0<\gamma_{\theta, t}^{g}<1$ such that

$$
\begin{equation*}
\tilde{\lambda} \tilde{a}_{\theta, t}^{2}-\left(1-\gamma_{\theta, t}^{g}\right) \tilde{a}_{\theta, t}+\tilde{\theta}_{1}+\tilde{\theta}_{2}<0 \tag{85}
\end{equation*}
$$

Take $\tilde{a}_{\theta, t}=(1-\delta)^{-1}>1$ and

$$
\tilde{\gamma}_{\theta, t}:=1-\gamma_{\theta, t}^{g}=\frac{1}{2}\left(1+(1-\tilde{\lambda})(1-\delta)+\tilde{\lambda}(1-\delta)^{-1}\right)
$$

We need to have $\tilde{\gamma}_{\theta, t}<1$ which follows from the stability condition $\tilde{\lambda} \leq 0.5-0.5 \delta$ as below:

$$
\begin{aligned}
\tilde{\gamma}_{\theta, t} & =\frac{1}{2}+\frac{1}{2}\left((1-\tilde{\lambda})(1-\delta)+\frac{\tilde{\lambda}}{1-\delta}\right)=\frac{1}{2}+\frac{1}{2}\left(1-\delta-\tilde{\lambda}(1-\delta)+\frac{\tilde{\lambda}}{1-\delta}\right) \\
& =\frac{1}{2}+\frac{1}{2}\left(1-\delta+\tilde{\lambda} \frac{1-(1-\delta)^{2}}{1-\delta}\right)=\frac{1}{2}+\frac{1}{2}\left(1-\delta+\tilde{\lambda} \frac{\delta(2-\delta)}{1-\delta}\right) \\
& \leq \frac{1}{2}+\frac{1}{2}\left(1-\delta+\frac{\delta(2-\delta)}{2}\right)=1-\frac{\delta^{2}}{4}<1 .
\end{aligned}
$$

We now verify 85):

$$
\begin{aligned}
\tilde{\lambda} \tilde{a}_{\theta, t}^{2}-\left(1-\gamma_{\theta, t}^{g}\right) \tilde{a}_{\theta, t}+\tilde{\theta}_{1}+\tilde{\theta}_{2} & =\frac{\tilde{\lambda}}{(1-\delta)^{2}}-\frac{1}{2(1-\delta)}-\frac{1-\tilde{\lambda}}{2}-\frac{\tilde{\lambda}}{2(1-\delta)^{2}}+1-\tilde{\lambda} \\
& =\frac{\tilde{\lambda}}{2(1-\delta)^{2}}+\frac{1-\tilde{\lambda}}{2}-\frac{1}{2(1-\delta)} \\
& =\frac{1}{2(1-\delta)^{2}}\left(\tilde{\lambda}+(1-\tilde{\lambda})(1-\delta)^{2}-(1-\delta)\right) \\
& =\frac{\delta}{2(1-\delta)^{2}}(\delta-1-\tilde{\lambda} \delta+2 \tilde{\lambda}) \\
& =\frac{\delta}{2(1-\delta)^{2}}(\tilde{\lambda}(2-\delta)+\delta-1) \\
& <0
\end{aligned}
$$

where the last line follows from $\tilde{\lambda} \leq 0.5-0.5 \delta<(1-\delta) /(2-\delta)$.
For $\boldsymbol{x}=(i, 0,1)$ and $i \geq 1$, we have

$$
P_{\theta}^{t} V_{\theta, t}^{g}(i, 0,1)=\tilde{\lambda} V_{\theta, t}^{g}(i, 1,1)+\tilde{\theta}_{1} V_{\theta, t}^{g}(i-1,0,1)+\tilde{\theta}_{2} V_{\theta, t}^{g}(i-1,1,0)
$$

and

$$
\begin{aligned}
& P_{\theta}^{t} V_{\theta, t}^{g}(i, 0,1)-\left(1-\gamma_{\theta, t}^{g}\right) V_{\theta, t}^{g}(i, 0,1) \\
& =\tilde{\lambda} \exp \left(a_{\theta, t}^{g}(i+2)\right)+\left(\tilde{\theta}_{1}+\tilde{\theta}_{2}\right) \exp \left(a_{\theta, t}^{g} i\right)-\left(1-\gamma_{\theta, t}^{g}\right) \exp \left(a_{\theta, t}^{g}(i+1)\right) \\
& =\exp \left(a_{\theta, t}^{g} i\right)\left(\tilde{\lambda} \exp \left(2 a_{\theta, t}^{g}\right)+\tilde{\theta}_{1}+\tilde{\theta}_{2}-\left(1-\gamma_{\theta, t}^{g}\right) \exp \left(a_{\theta, t}^{g}\right)\right),
\end{aligned}
$$

which results in the same conditions as previously discussed. When $\boldsymbol{x}=(i, 1,0)$ and $i \geq t$ also same argument holds.
Finally, 84] holds for

$$
\begin{aligned}
C_{\theta, t}^{g} & =\left\{\left(x_{0}, x_{1}, 0\right): x_{0}<t\right\} \cup\{(0,0,1)\} \\
a_{\theta, t}^{g} & =-\log (1-\delta), \\
\gamma_{\theta, t}^{g} & =\frac{1}{2}-\frac{1}{2}\left((1-\tilde{\lambda})(1-\delta)+\tilde{\lambda}(1-\delta)^{-1}\right), \\
V_{\theta, t}^{g}(\boldsymbol{x}) & =\exp \left(a_{\theta, t}^{g}\|\boldsymbol{x}\|_{1}\right), \\
b_{\theta, t}^{g} & =\max _{\boldsymbol{x} \in C_{\theta, t}^{g}} \exp \left(a_{\theta, t}^{g}\|\boldsymbol{x}\|_{1}\right)\left(\exp \left(a_{\theta, t}^{g}\right)+1\right),
\end{aligned}
$$

where the last line holds because $P V_{\theta, t}^{g}(\boldsymbol{x}) \leq V_{\theta, t}^{g}(\boldsymbol{y})$ for $\boldsymbol{y}$ such that $\|\boldsymbol{y}\|_{1}=\|\boldsymbol{x}\|_{1}+1$.

## F. 2 Proof of Proposition 2

Proof. In order to show polynomially ergodicity, we will verify (62). We define $V_{\theta, t}^{p}(\boldsymbol{x})=\|\boldsymbol{x}\|_{1}^{2}$ and $\alpha_{\theta, t}^{p}=1 / 2$, which is equal to $r /(r+1)$ for $r=1 ; r$ is defined in Assumption 1 . For $\boldsymbol{x}=(i, 0,1)$ and $i \geq 1$,

$$
P_{\theta}^{t} V_{\theta, t}^{p}(i, 0,1)=\tilde{\lambda} V_{\theta, t}^{p}(i, 1,1)+\tilde{\theta}_{1} V_{\theta, t}^{p}(i-1,0,1)+\tilde{\theta}_{2} V_{\theta, t}^{p}(i-1,1,0)
$$

in which $\tilde{\lambda}, \tilde{\theta}_{1}$, and $\tilde{\theta}_{2}$ are the normalized rates defined in 83). Thus,

$$
\begin{aligned}
& P_{\theta}^{t} V_{\theta, t}^{p}(i, 0,1)-V_{\theta, t}^{p}(i, 0,1)+\beta_{\theta, t}^{p} \sqrt{V_{\theta, t}^{p}(i, 0,1)} \\
& =\tilde{\lambda}(i+2)^{2}+\left(\tilde{\theta}_{1}+\tilde{\theta}_{2}\right) i^{2}-(i+1)^{2}+\beta_{\theta, t}^{p}(i+1) \\
& =i\left(4 \tilde{\lambda}-2+\beta_{\theta, t}^{p}\right)+4 \tilde{\lambda}-1+\beta_{\theta, t}^{p} .
\end{aligned}
$$

For $\beta_{\theta, t}^{p}=1-2 \tilde{\lambda}$, the right-hand side of above equation is non-positive for $i \geq \frac{2 \tilde{\lambda}}{1-2 \tilde{\lambda}}$. For $\boldsymbol{x}=(i, 1,0)$ and $i \geq t$,

$$
P_{\theta}^{t} V_{\theta, t}^{p}(i, 1,0)=\tilde{\lambda} V_{\theta, t}^{p}(i, 1,1)+\tilde{\theta}_{1} V_{\theta, t}^{p}(i-1,0,1)+\tilde{\theta}_{2} V_{\theta, t}^{p}(i-1,1,0)
$$

Thus,

$$
\begin{aligned}
& P_{\theta}^{t} V_{\theta, t}^{p}(i, 1,0)-V_{\theta, t}^{p}(i, 1,0)+\beta_{\theta, t}^{p} \sqrt{V_{\theta, t}^{p}(i, 1,0)} \\
& =\tilde{\lambda}(i+2)^{2}+\left(\tilde{\theta}_{1}+\tilde{\theta}_{2}\right) i^{2}-(i+1)^{2}+\beta_{\theta, t}^{p}(i+1) \\
& =i\left(4 \tilde{\lambda}-2+\beta_{\theta, t}^{p}\right)+4 \tilde{\lambda}-1+\beta_{\theta, t}^{p}
\end{aligned}
$$

which is also non-positive under the same conditions as the previous case. For $i \geq 1$ and $\boldsymbol{x}=(i, 1,1)$,

$$
P_{\theta}^{t} V_{\theta, t}^{p}(i, 1,1)=\tilde{\lambda} V_{\theta, t}^{p}(i+1,1,1)+\tilde{\theta}_{1} V_{\theta, t}^{p}(i, 0,1)+\tilde{\theta}_{2} V_{\theta, t}^{p}(i, 1,0)
$$

Thus,

$$
\begin{aligned}
& P_{\theta}^{t} V_{\theta, t}^{p}(i, 1,1)-V_{\theta, t}^{p}(i, 1,1)+\beta_{\theta, t}^{p} \sqrt{V_{\theta, t}^{p}(i, 1,1)} \\
& =\tilde{\lambda}(i+3)^{2}+\left(\tilde{\theta}_{1}+\tilde{\theta}_{2}\right)(i+1)^{2}-(i+2)^{2}+\beta_{\theta, t}^{p}(i+2) \\
& =i\left(4 \tilde{\lambda}-2+\beta_{\theta, t}^{p}\right)+8 \tilde{\lambda}-3+2 \beta_{\theta, t}^{p},
\end{aligned}
$$

which is non-positive under the same conditions as the first case. Finally, 62) holds for

$$
\begin{aligned}
C_{\theta, t}^{p} & =\left\{\left(x_{0}, x_{1}, 0\right): x_{0}<t\right\} \cup\left\{\left(x_{0}, x_{1}, x_{2}\right): x_{0}<\frac{2 \tilde{\lambda}}{1-2 \tilde{\lambda}}, x_{1}+x_{2} \geq 1\right\} \\
\beta_{\theta, t}^{p} & =1-2 \tilde{\lambda} \\
\alpha_{\theta, t}^{p} & =\frac{1}{2} \\
V_{\theta, t}^{p}(\boldsymbol{x}) & =\|\boldsymbol{x}\|_{1}^{2} \\
b_{\theta, t}^{p} & \left.=\max _{\boldsymbol{x} \in C_{\theta, t}^{p}}\left(\|\boldsymbol{x}\|_{1}+1\right)^{2}\right),
\end{aligned}
$$

where the last line holds because $P V_{\theta, t}^{p}(\boldsymbol{x}) \leq V_{\theta, t}^{p}(\boldsymbol{y})$ for $\boldsymbol{y}$ such that $\|\boldsymbol{y}\|_{1}=\|\boldsymbol{x}\|_{1}+1$.

## F. 3 Proof of Proposition 3

Proof. To show geometric ergodicity of the chain that follows $\pi_{\omega}$, we verify (70). Take $a_{\theta, \omega}^{g}>0$ and

$$
\begin{equation*}
V_{\theta, \omega}^{g}(\boldsymbol{x})=\frac{\omega}{\omega+1} \exp \left(a_{\theta, \omega}^{g} \frac{x_{1}+1}{\omega}\right)+\frac{1}{\omega+1} \exp \left(a_{\theta, \omega}^{g}\left(x_{2}+1\right)\right) \tag{86}
\end{equation*}
$$

First, we find $P V_{\theta, \omega}^{g}(\boldsymbol{x})$ for the function defined above. We have

$$
\begin{equation*}
P V_{\theta, \omega}^{g}(\boldsymbol{x})=\mathbb{E}_{\boldsymbol{x}}^{\pi_{\omega}}\left[\frac{\omega}{\omega+1} \exp \left(a_{\theta, \omega}^{g} \frac{X_{1}(2)+1}{\omega}\right)\right]+\mathbb{E}_{\boldsymbol{x}}^{\pi_{\omega}}\left[\frac{1}{\omega+1} \exp \left(a_{\theta, \omega}^{g}\left(X_{2}(2)+1\right)\right)\right] \tag{87}
\end{equation*}
$$

where $\boldsymbol{X}(2)=\left(X_{1}(2), X_{2}(2)\right)$ is the state of the system at the second arrival, starting from state $\boldsymbol{x}$. To find the above expectations, we first find the corresponding transition probabilities. If the number of departures from server $i$ during a fixed interval with length $t$ is less than the total number of jobs in the queue of that server, the number of departures follows a Poisson distribution with parameter $\theta_{i} t$. Let $\mathbb{P}\left(\left(x_{1}, x_{2}\right) \rightarrow\left(x_{1}^{\prime}, \mathcal{X}\right)\right)$ be the probability of transitioning from a system with $x_{i}$ jobs in server-queue pair $i$ (just after the assignment of the arrival) to a queueing system with $x_{1}^{\prime}$ jobs in the first server-queue pair (just before the upcoming arrival). For $1 \leq x_{1}^{\prime} \leq x_{1}$, we have
$\mathbb{P}\left(\left(x_{1}, x_{2}\right) \rightarrow\left(x_{1}^{\prime}, \mathcal{X}\right)\right)=\int_{0}^{\infty} \lambda \exp (-\lambda t) \frac{\left(\theta_{1} t\right)^{x_{1}-x_{1}^{\prime}}}{\left(x_{1}-x_{1}^{\prime}\right)!} \exp \left(-\theta_{1} t\right) d t=\frac{\lambda}{\theta_{1}+\lambda}\left(\frac{\theta_{1}}{\theta_{1}+\lambda}\right)^{x_{1}-x_{1}^{\prime}}$,
and

$$
\begin{equation*}
\mathbb{P}\left(\left(x_{1}, x_{2}\right) \rightarrow(0, \mathcal{X})\right)=1-\sum_{i=1}^{x_{1}} \frac{\lambda}{\theta_{1}+\lambda}\left(\frac{\theta_{1}}{\theta_{1}+\lambda}\right)^{x_{1}-i}=\left(\frac{\theta_{1}}{\theta_{1}+\lambda}\right)^{x_{1}} \tag{89}
\end{equation*}
$$

Assume $1+x_{1} \leq \omega\left(1+x_{2}\right)$, which results in the new arrival being assigned to the first server. For the first term in (87), we have

$$
\begin{align*}
& \mathbb{E}_{\boldsymbol{x}}^{\pi_{\omega}}\left[\exp \left(a_{\theta, \omega}^{g} \frac{X_{1}(2)}{\omega}\right)\right] \\
& =\sum_{i=0}^{x_{1}+1} \mathbb{P}\left(\left(x_{1}+1, x_{2}\right) \rightarrow(i, \mathcal{X})\right) \exp \left(a_{\theta, \omega}^{g} \frac{i}{\omega}\right) \\
& =\left(\frac{\theta_{1}}{\theta_{1}+\lambda}\right)^{x_{1}+1}+\sum_{i=1}^{x_{1}+1} \exp \left(a_{\theta, \omega}^{g} \frac{i}{\omega}\right) \frac{\lambda}{\theta_{1}+\lambda}\left(\frac{\theta_{1}}{\theta_{1}+\lambda}\right)^{x_{1}+1-i} \\
& =\left(\frac{\theta_{1}}{\theta_{1}+\lambda}\right)^{x_{1}+1}+\frac{\lambda}{\theta_{1}+\lambda} \exp \left(a_{\theta, \omega}^{g} \frac{x_{1}+1}{\omega}\right) \frac{1-\exp \left(-a_{\theta, \omega}^{g} \frac{x_{1}+1}{\omega}\right)\left(\frac{\theta_{1}}{\theta_{1}+\lambda}\right)^{x_{1}+1}}{1-\exp \left(-\frac{a_{\theta, \omega}^{g}}{\omega}\right) \frac{\theta_{1}}{\theta_{1}+\lambda}} \\
& <\left(\frac{\theta_{1}}{\theta_{1}+\lambda}\right)^{x_{1}+1}+\frac{\lambda}{\theta_{1}+\lambda} \exp \left(a_{\theta, \omega}^{g} \frac{x_{1}+1}{\omega}\right) \frac{1}{1-\exp \left(-\frac{a_{\theta, \omega}^{g}}{\omega}\right) \frac{\theta_{1}}{\theta_{1}+\lambda}} \tag{90}
\end{align*}
$$

Similarly, for the second term in (87), we have

$$
\begin{equation*}
\mathbb{E}_{\boldsymbol{x}}^{\pi_{\omega}}\left[\exp \left(a_{\theta, \omega}^{g} X_{2}(2)\right)\right] \leq\left(\frac{\theta_{2}}{\theta_{2}+\lambda}\right)^{x_{2}}+\frac{\lambda}{\theta_{2}+\lambda} \exp \left(a_{\theta, \omega}^{g} x_{2}\right) \frac{1}{1-\exp \left(-a_{\theta, \omega}^{g}\right) \frac{\theta_{2}}{\theta_{2}+\lambda}} \tag{91}
\end{equation*}
$$

To satisfy $\sqrt{70}$, for some $0<\gamma_{\theta, \omega}^{g}<1$ and all but finitely many $\boldsymbol{x}$, the following should hold,

$$
P V_{\theta, \omega}^{g}(\boldsymbol{x}) \leq \gamma_{\theta, \omega}^{g} V_{\theta, \omega}^{g}(\boldsymbol{x})
$$

or from (86) and 87,

$$
\begin{aligned}
& \mathbb{E}_{\boldsymbol{x}}^{\pi_{\omega}}\left[\omega \exp \left(a_{\theta, \omega}^{g} \frac{X_{1}(2)+1}{\omega}\right)\right]+\mathbb{E}_{\boldsymbol{x}}^{\pi_{\omega}}\left[\exp \left(a_{\theta, \omega}^{g}\left(X_{2}(2)+1\right)\right)\right] \\
& \leq \gamma_{\theta, \omega}^{g}\left(\omega \exp \left(a_{\theta, \omega}^{g} \frac{x_{1}+1}{\omega}\right)+\exp \left(a_{\theta, \omega}^{g}\left(x_{2}+1\right)\right)\right)
\end{aligned}
$$

Notice that

$$
\omega\left(\frac{\theta_{1}}{\theta_{1}+\lambda}\right)^{x_{1}+1}+\left(\frac{\theta_{2}}{\theta_{2}+\lambda}\right)^{x_{2}} \leq c_{R} R+1
$$

From (90) and 91, it suffices to have

$$
\begin{align*}
& \left(c_{R} R+1\right) \exp \left(c_{R} R a_{\theta, \omega}^{g}\right)+\frac{\omega \frac{\lambda}{\theta_{1}+\lambda} \exp \left(a_{\theta, \omega}^{g} \frac{x_{1}+2}{\omega}\right)}{1-\exp \left(-\frac{a_{\theta, \omega}^{g}}{\omega}\right) \frac{\theta_{1}}{\theta_{1}+\lambda}}+\frac{\frac{\lambda}{\theta_{2}+\lambda} \exp \left(a_{\theta, \omega}^{g}\left(x_{2}+1\right)\right)}{1-\exp \left(-a_{\theta, \omega}^{g}\right) \frac{\theta_{2}}{\theta_{2}+\lambda}} \\
& \leq \gamma_{\theta, \omega}^{g}\left(\omega \exp \left(a_{\theta, \omega}^{g} \frac{x_{1}+1}{\omega}\right)+\exp \left(a_{\theta, \omega}^{g}\left(x_{2}+1\right)\right)\right) \tag{92}
\end{align*}
$$

Define

$$
\zeta_{1, \theta, \omega}=\frac{\frac{\lambda}{\theta_{1}+\lambda}}{1-\exp \left(-\frac{a_{\theta, \omega}^{g}}{\omega}\right) \frac{\theta_{1}}{\theta_{1}+\lambda}}, \quad \zeta_{2, \theta, \omega}=\frac{\frac{\lambda}{\theta_{2}+\lambda}}{1-\exp \left(-a_{\theta, \omega}^{g}\right) \frac{\theta_{2}}{\theta_{2}+\lambda}}
$$

Simplifying (92), we need the following to hold

$$
\begin{align*}
& \left(c_{R} R+1\right) \exp \left(c_{R} R a_{\theta, \omega}^{g}\right)+\omega \exp \left(a_{\theta, \omega}^{g} \frac{x_{1}+1}{\omega}\right)\left(\zeta_{1, \theta, \omega} \exp \left(\frac{a_{\theta, \omega}^{g}}{\omega}\right)-\gamma_{\theta, \omega}^{g}\right) \\
& +\exp \left(a_{\theta, \omega}^{g}\left(x_{2}+1\right)\right)\left(\zeta_{2, \theta, \omega}-\gamma_{\theta, \omega}^{g}\right) \leq 0 \tag{93}
\end{align*}
$$

As $\zeta_{i, \theta, \omega}<1$, there exists $\gamma_{\theta, \omega}^{g}$ such that

$$
\zeta_{2, \theta, \omega}<\gamma_{\theta, \omega}^{g}<1
$$

From the assumption $1+x_{1} \leq \omega\left(1+x_{2}\right)$ and the above equation, 93) can be further simplified as

$$
\begin{equation*}
\left(c_{R} R+1\right) \exp \left(c_{R} R a_{\theta, \omega}^{g}\right)+\exp \left(a_{\theta, \omega}^{g} \frac{x_{1}+1}{\omega}\right)\left(\omega \zeta_{1, \theta, \omega} \exp \left(\frac{a_{\theta, \omega}^{g}}{\omega}\right)+\zeta_{2, \theta, \omega}-(\omega+1) \gamma_{\theta, \omega}^{g}\right) \leq 0 \tag{94}
\end{equation*}
$$

For the above to hold outside a finite set, we need to have

$$
\begin{equation*}
\frac{\zeta_{1, \theta, \omega} \omega}{1+\omega} \exp \left(\frac{a_{\theta, \omega}^{g}}{\omega}\right)+\frac{\zeta_{2, \theta, \omega}}{1+\omega}<\gamma_{\theta, \omega}^{g} \tag{95}
\end{equation*}
$$

Define

$$
\begin{equation*}
\zeta_{3}=\frac{1}{1+\delta}, \quad \zeta_{4}=\frac{1-0.5 \delta}{1-\delta} \tag{96}
\end{equation*}
$$

Note that $\zeta_{3}<1$ and $\zeta_{4}>1$. Defining function $f(y):=1+\zeta_{4} y-\exp (y)$, we note that for $y \leq \log \zeta_{4}$, $f(y)>0$, where $\log \zeta_{4}$ is the maximizer of $f(y)$. Similarly, taking $g(y):=1-\zeta_{3} y-\exp (-y)$, for $y \leq-\log \zeta_{3}, g(y)>0$, where $-\log \zeta_{3}$ is the maximizer of $g(y)$. Thus, we conclude that for $a_{\theta, \omega}^{g} \leq \min \left(-\omega \log \zeta_{3},-\log \zeta_{3}, \omega \log \zeta_{4}\right)$,

$$
\begin{align*}
\exp (-y) \leq 1-\zeta_{3} y \quad \text { holds for } \quad y & \leq \max \left(\frac{a_{\theta, \omega}^{g}}{\omega}, a_{\theta, \omega}^{g}\right)  \tag{97}\\
\exp (y) \leq 1+\zeta_{4} y \quad \text { holds for } \quad y & \leq \frac{a_{\theta, \omega}^{g}}{\omega} \tag{98}
\end{align*}
$$

To guarantee the existence of $0<\gamma_{\theta, \omega}^{g}<1$ that satisfies (95), we need to ensure the left-hand side of (95) is strictly less than 1 . Using the bounds found in 97) and (98) and the definition of $\zeta_{1, \theta, \omega}$ and $\zeta_{2, \theta, \omega}$, we simplify (95) to get

$$
\frac{\frac{\lambda}{1+\omega}\left(\omega+a_{\theta, \omega}^{g} \zeta_{4}\right)}{\lambda+\frac{\theta_{1} a_{\theta, \omega}^{g} \zeta_{3}}{\omega}}+\frac{\frac{\lambda}{1+\omega}}{\lambda+\theta_{2} a_{\theta, \omega}^{g} \zeta_{3}}<1
$$

which is equivalent to

$$
\begin{equation*}
a_{\theta, \omega}^{g} \zeta_{3} \theta_{2}\left(\lambda \zeta_{4}-\frac{\zeta_{3} \theta_{1}(1+\omega)}{\omega}\right)<\lambda \zeta_{3}\left(\theta_{1}+\theta_{2}\right)-\lambda^{2} \zeta_{4} \tag{99}
\end{equation*}
$$

To make sure there exists $a_{\theta, \omega}^{g}>0$ that satisfies 99 , the right-hand side of 99 needs to be positive, which follows as below:

$$
\begin{align*}
\lambda \zeta_{3}\left(\theta_{1}+\theta_{2}\right)-\lambda^{2} \zeta_{4} & =\lambda\left(\frac{\theta_{1}+\theta_{2}}{1+\delta}-\lambda \frac{1-0.5 \delta}{1-\delta}\right) \\
& =\lambda\left(\theta_{1}+\theta_{2}+\lambda\right)\left(\frac{1-\tilde{\lambda}}{1+\delta}-\tilde{\lambda} \frac{1-0.5 \delta}{1-\delta}\right) \\
& =\lambda\left(\theta_{1}+\theta_{2}+\lambda\right)\left(\frac{1}{1+\delta}-\tilde{\lambda}\left(\frac{1}{1+\delta}+\frac{1-0.5 \delta}{1-\delta}\right)\right) \\
& \geq \lambda\left(\theta_{1}+\theta_{2}+\lambda\right)\left(\frac{1}{1+\delta}-\frac{1-\delta}{2}\left(\frac{1}{1+\delta}+\frac{1-0.5 \delta}{1-\delta}\right)\right) \\
& =\frac{\delta}{4} \lambda\left(\theta_{1}+\theta_{2}+\lambda\right) \tag{100}
\end{align*}
$$

${\underset{\sim}{w}}^{w}$ here $\tilde{\lambda}, \tilde{\theta}_{1}$, and $\tilde{\theta}_{2}$ are the normalized rates defined in (83) and we have used the stability condition $\tilde{\lambda} \leq 0.5-0.5 \delta$. We further simplify the left-hand side of (99) as

$$
\zeta_{3} \theta_{2}\left(\lambda \zeta_{4}-\frac{\zeta_{3} \theta_{1}(1+\omega)}{\omega}\right)<\theta_{2} \lambda \zeta_{3} \zeta_{4}<\frac{1-0.5 \delta}{1-\delta^{2}}\left(\theta_{1}+\theta_{2}+\lambda\right) \lambda
$$

From the above equation and $\sqrt[100]{10}, a_{\theta, \omega}^{g}$ needs to satisfy

$$
a_{\theta, \omega}^{g} \leq \frac{\delta\left(1-\delta^{2}\right)}{8(1-0.5 \delta)}
$$

Finally, we take $a_{\theta, \omega}^{g}$ as

$$
a_{\theta, \omega}^{g}=\min \left(-\omega \log \zeta_{3},-\log \zeta_{3}, \omega \log \zeta_{4}, \frac{\delta\left(1-\delta^{2}\right)}{8(1-0.5 \delta)}\right)
$$

After finding an appropriate $a_{\theta, \omega}^{g}$, we can choose $0<\gamma_{\theta, \omega}^{g}<1$ such that (95) holds or

$$
\gamma_{\theta, \omega}^{g} \geq \frac{1}{2}\left(1+\frac{\zeta_{1, \theta, \omega} \omega}{1+\omega} \exp \left(\frac{a_{\theta, \omega}^{g}}{\omega}\right)+\frac{\zeta_{2, \theta, \omega}}{1+\omega}\right)
$$

Moreover, from (94) a lower bound $x_{1, \theta, \omega}^{g_{1}}$ for $x_{1}$ is derived; In other words, 94) holds for $x_{1}>x_{1, \theta, \omega}^{g_{1}}$. From (93), we can find the corresponding $x_{2, \theta, \omega}^{g_{1}}$ and take $\boldsymbol{x}_{\theta, \omega}^{g_{1}}=\left(x_{1, \theta, \omega}^{g_{1}}, x_{2, \theta, \omega}^{g_{1}}\right)$. By repeating the same arguments when $1+x_{1}<\omega\left(1+x_{2}\right)$, we finally conclude that

$$
\Delta V_{\theta, \omega}^{g}(\boldsymbol{x}) \leq-\left(1-\gamma_{\theta, \omega}^{g}\right) V_{\theta, \omega}^{g}(\boldsymbol{x})+b_{\theta, \omega}^{g} \mathbb{I}_{C_{\theta, \omega}^{g}}(\boldsymbol{x}), \quad \boldsymbol{x} \in \mathcal{X}
$$

for

$$
\begin{aligned}
& V_{\theta, \omega}^{g}(\boldsymbol{x})=\frac{\omega}{\omega+1} \exp \left(a_{\theta, \omega}^{g} \frac{x_{1}+1}{\omega}\right)+\frac{1}{\omega+1} \exp \left(a_{\theta, \omega}^{g}\left(x_{2}+1\right)\right), \\
& a_{\theta, \omega}^{g}=\min \left(\omega \log (1+\delta), \log (1+\delta), \omega \log \frac{1-0.5 \delta}{1-\delta}, \log \frac{1-0.5 \delta}{1-\delta}, \frac{\delta\left(1-\delta^{2}\right)}{4 c_{R} R(1-0.5 \delta)}\right), \\
& C_{\theta, \omega}^{g}=\left\{\left(x_{1}, x_{2}\right) \in \mathcal{X}: x_{i} \leq \max \left(x_{i, \theta, \omega}^{g_{j}}, 0\right), i, j=1,2\right\}, \\
& \gamma_{\theta, \omega}^{g}=\frac{1}{2}+\frac{1}{2} \max \left(\zeta_{1, \theta, \omega}, \zeta_{2, \theta, \omega}, \frac{\zeta_{1, \theta, \omega} \omega}{1+\omega} \exp \left(\frac{a_{\theta, \omega}^{g}}{\omega}\right)+\frac{\zeta_{2, \theta, \omega}}{1+\omega}, \frac{\zeta_{1, \theta, \omega} \omega}{1+\omega}+\frac{\zeta_{2, \theta, \omega}}{1+\omega} \exp \left(a_{\theta, \omega}^{g}\right)\right), \\
& b_{\theta, \omega}^{g}=\max _{x \in C_{\theta, \omega}^{g}}\left(\frac{2 \omega}{\omega+1} \exp \left(a_{\theta, \omega}^{g} \frac{x_{1}+2}{\omega}\right)+\frac{2}{\omega+1} \exp \left(a_{\theta, \omega}^{g}\left(x_{2}+2\right)\right)\right), \\
& \zeta_{1, \theta, \omega}=\frac{\frac{\lambda}{\theta_{1}+\lambda}}{1-\exp \left(-\frac{a_{\theta, \omega}^{g}}{\omega}\right) \frac{\theta_{1}}{\theta_{1}+\lambda}}, \\
& \zeta_{2, \theta, \omega}=\frac{\frac{\lambda}{\theta_{2}+\lambda}}{1-\exp \left(-a_{\theta, \omega}^{g}\right) \frac{\theta_{2}}{\theta_{2}+\lambda}}, \\
& x_{1, \theta, \omega}^{g_{1}}=\frac{\omega}{a_{\theta, \omega}^{g}} \log \frac{\left(c_{R} R+1\right) \exp \left(c_{R} R a_{\theta, \omega}^{g}\right)}{(\omega+1) \gamma_{\theta, \omega}^{g}-\omega \zeta_{1, \theta, \omega} \exp \left(\frac{a_{\theta, \omega}^{g}}{\omega}\right)-\zeta_{2, \theta, \omega}}, \\
& x_{2, \theta, \omega}^{g_{1}}=\frac{1}{a_{\theta, \omega}^{g}} \log \frac{\left(c_{R} R+1\right) \exp \left(c_{R} R a_{\theta, \omega}^{g}\right)+\omega \exp \left(a_{\theta, \omega}^{g} \frac{x_{1, \theta, \omega}^{g_{1}+1}}{\omega}\right)\left(\zeta_{1, \theta, \omega} \exp \left(\frac{a_{\theta, \omega}^{g}}{\omega}\right)-\gamma_{\theta, \omega}^{g}\right)}{\gamma_{\theta, \omega}^{g}-\zeta_{2, \theta, \omega}}, \\
& x_{2, \theta, \omega}^{g_{2}}=\frac{1}{a_{\theta, \omega}^{g}} \log \frac{\left(c_{R} R+1\right) \exp \left(c_{R} R a_{\theta, \omega}^{g}\right)}{(\omega+1) \gamma_{\theta, \omega}^{g}-\omega \zeta_{1, \theta, \omega}-\zeta_{2, \theta, \omega} \exp \left(a_{\theta, \omega}^{g}\right)}, \\
& x_{1, \theta, \omega}^{g_{2}}=\frac{\omega}{a_{\theta, \omega}^{g}} \log \frac{\left(c_{R} R+1\right) \exp \left(c_{R} R a_{\theta, \omega}^{g}\right)+\exp \left(a_{\theta, \omega}^{g}\left(x_{2, \theta, \omega}^{g_{2}}+1\right)\right)\left(\zeta_{2, \theta, \omega} \exp \left(a_{\theta, \omega}^{g}\right)-\gamma_{\theta, \omega}^{g}\right)}{\omega\left(\gamma_{\theta, \omega}^{g}-\zeta_{1, \theta, \omega}\right)} .
\end{aligned}
$$

## F. 4 Proof of Proposition 4

Proof. Define $V_{\theta, \omega}^{p}(\boldsymbol{x})=\frac{x_{1}^{2}}{\omega}+x_{2}^{2}$, and $\alpha_{\theta, \omega}^{p}=1 / 2$. Assume that $x_{1}=0$ and $x_{2}>(1-\omega) / \omega$; which means the new job will be assigned to the first server. The transition probabilities of the discrete-time chain sampled at Poisson arrivals is given in (88) and (89), and we calculate $P V_{\theta, \omega}^{p}(\boldsymbol{x})$ as

$$
\begin{equation*}
P V_{\theta, \omega}^{p}(\boldsymbol{x})=\frac{\lambda}{\omega\left(\lambda+\theta_{1}\right)}+\sum_{i=1}^{x_{2}} i^{2} \frac{\lambda}{\lambda+\theta_{2}}\left(\frac{\theta_{2}}{\theta_{2}+\lambda}\right)^{x_{2}-i}<c_{R} R+\sum_{i=1}^{x_{2}} i^{2} \frac{\lambda}{\lambda+\theta_{2}}\left(\frac{\theta_{2}}{\theta_{2}+\lambda}\right)^{x_{2}-i} \tag{101}
\end{equation*}
$$

We define $d_{i}:=\theta_{i} /\left(\theta_{i}+\lambda\right)$ for $i=1,2$ and

$$
\begin{align*}
& \sum_{i=1}^{x_{2}} i^{2} \frac{\lambda}{\lambda+\theta_{2}}\left(\frac{\theta_{2}}{\theta_{2}+\lambda}\right)^{x_{2}-i} \\
& =\frac{1}{\left(1-d_{2}\right)^{2}}\left(-d_{2}^{x_{2}}\left(d_{2}+d_{2}^{2}\right)+d_{2}^{2}\left(x_{2}^{2}+2 x_{2}+1\right)+d_{2}\left(-2 x_{2}^{2}-2 x_{2}+1\right)+x_{2}^{2}\right) \\
& =\frac{1}{\left(1-d_{2}\right)^{2}}\left(\left(1-d_{2}^{x_{2}}\right)\left(d_{2}+d_{2}^{2}\right)+x_{2}^{2}\left(d_{2}^{2}-2 d_{2}+1\right)+x_{2}\left(2 d_{2}^{2}-2 d_{2}\right)\right) \\
& =x_{2}^{2}-\frac{2 d_{2}}{1-d_{2}} x_{2}+\frac{\left(1-d_{2}^{x_{2}}\right)\left(d_{2}+d_{2}^{2}\right)}{\left(1-d_{2}\right)^{2}} . \tag{102}
\end{align*}
$$

From (101),

$$
P V_{\theta, \omega}^{p}(\boldsymbol{x})-V_{\theta, \omega}^{p}(\boldsymbol{x})+\beta_{\theta, \omega}^{p} x_{2}<\left(-\frac{2 d_{2}}{1-d_{2}}+\beta_{\theta, \omega}^{p}\right) x_{2}+\frac{\left(1-d_{2}^{x_{2}}\right)\left(d_{2}+d_{2}^{2}\right)}{\left(1-d_{2}\right)^{2}}+c_{R} R .
$$

Outside a finite set, we need the above equation to be non-positive; which is equivalent to

$$
\left(-2+\beta_{\theta, \omega}^{p} \frac{1-d_{2}}{d_{2}}\right) x_{2}+\frac{\left(1-d_{2}^{x_{2}}\right)\left(1+d_{2}\right)}{1-d_{2}}+c_{R} R \frac{1-d_{2}}{d_{2}} \leq 0
$$

As $d_{2}<1$,

$$
\begin{equation*}
\frac{1-d_{2}^{y}}{1-d_{2}}=1+d_{2}+\ldots+d_{2}^{y-1} \leq y \quad \text { for } y \geq 1 \tag{103}
\end{equation*}
$$

Thus,

$$
\begin{aligned}
& \left(-2+\beta_{\theta, \omega}^{p} \frac{1-d_{2}}{d_{2}}\right) x_{2}+\frac{\left(1-d_{2}^{x_{2}}\right)\left(1+d_{2}\right)}{1-d_{2}}+c_{R} R \frac{1-d_{2}}{d_{2}} \\
& \leq\left(d_{2}-1+\beta_{\theta, \omega}^{p} \frac{1-d_{2}}{d_{2}}\right) x_{2}+c_{R} R \frac{1-d_{2}}{d_{2}} .
\end{aligned}
$$

By taking $\beta_{\theta, \omega}^{p} \leq d_{2} / 2$, it suffices for the following to be non-positive,

$$
-\frac{1-d_{2}}{2} x_{2}+c_{R} R \frac{1-d_{2}}{d_{2}} \leq 0
$$

which holds for $x_{2} \geq 2 c_{R} R / d_{2}$. Thus, for $x_{1}=0$ and $x_{2} \geq \max \left(2 c_{R} R\left(\lambda+\theta_{2}\right) / \theta_{2},(1-\omega) / \omega\right)=$ $2 c_{R} R\left(\lambda+\theta_{2}\right) / \theta_{2}$, (76) holds. The case of $x_{2}=0$ and non-zero $x_{1}$ follows same arguments and (76) holds for $\beta_{\theta, \omega}^{p} \leq d_{1} / 2 \sqrt{\omega}, x_{2}=0$, and $x_{1} \geq \max \left(2 c_{R} R\left(\lambda+\theta_{1}\right) / \theta_{1}, \omega-1\right)=2 c_{R} R\left(\lambda+\theta_{1}\right) / \theta_{1}$. We now consider the case of $x_{1}, x_{2}>0$ and $x_{1}+1 \leq \omega\left(x_{2}+1\right)$, and note that

$$
\sqrt{V_{\theta, \omega}^{p}(\boldsymbol{x})}=\sqrt{\frac{x_{1}^{2}}{\omega}+x_{2}^{2}} \leq \sqrt{\frac{\left(x_{1}+1\right)^{2}}{\omega}+\left(x_{2}+1\right)^{2}} \leq \sqrt{\omega+1}\left(x_{2}+1\right)
$$

Hence, it suffices to find finite set $C_{\theta, \omega}^{p}$, constants $b_{\theta, \omega}^{p}$ and $\beta_{\theta, \omega}^{p}>0$, such that the following holds for $V_{\theta, \omega}^{p}(\boldsymbol{x})=\frac{x_{1}^{2}}{\omega}+x_{2}^{2}$,

$$
\Delta V_{\theta, \omega}^{p}(\boldsymbol{x}) \leq-\sqrt{\omega+1} \beta_{\theta, \omega}^{p}\left(x_{2}+1\right)+b_{\theta, \omega}^{p} \mathbb{I}_{C_{\theta, \omega}^{p}}^{p}(\boldsymbol{x}) .
$$

As $x_{1}+1 \leq \omega\left(x_{2}+1\right)$, the new arrival is assigned to the first queue and we find $\Delta V_{\theta, \omega}^{p}(\boldsymbol{x})+$ $\sqrt{\omega+1} \beta_{\theta, \omega}^{p}\left(x_{2}+1\right)$ using the same calculations as 102 .

$$
\begin{align*}
& \Delta V_{\theta, \omega}^{p}(\boldsymbol{x})+\sqrt{\omega+1} \beta_{\theta, \omega}^{p}\left(x_{2}+1\right) \\
& =\frac{1}{\omega}\left(\left(x_{1}+1\right)^{2}-\frac{2 d_{1}}{1-d_{1}}\left(x_{1}+1\right)+\frac{\left(1-d_{1}^{x_{1}+1}\right)\left(d_{1}+d_{1}^{2}\right)}{\left(1-d_{1}\right)^{2}}-x_{1}^{2}\right) \\
& -\frac{2 d_{2}}{1-d_{2}} x_{2}+\frac{\left(1-d_{2}^{x_{2}}\right)\left(d_{2}+d_{2}^{2}\right)}{\left(1-d_{2}\right)^{2}}+\sqrt{\omega+1} \beta_{\theta, \omega}^{p}\left(x_{2}+1\right) \\
& =\frac{x_{1}}{\omega}\left(2-\frac{2 d_{1}}{1-d_{1}}\right)+\frac{1-3 d_{1}}{\omega\left(1-d_{1}\right)}+\frac{\left(1-d_{1}^{x_{1}+1}\right)\left(d_{1}+d_{1}^{2}\right)}{\omega\left(1-d_{1}\right)^{2}}  \tag{104}\\
& +\left(x_{2}+1\right)\left(-\frac{2 d_{2}}{1-d_{2}}+\sqrt{\omega+1} \beta_{\theta, \omega}^{p}\right)+\frac{2 d_{2}}{1-d_{2}}+\frac{\left(1-d_{2}^{x_{2}}\right)\left(d_{2}+d_{2}^{2}\right)}{\left(1-d_{2}\right)^{2}} . \tag{105}
\end{align*}
$$

We next consider two different cases based on the value of $d_{1}$ and analyze them separately.
One. $0.8 \leq d_{1}<1$ : We first notice that the coefficient of $x_{1}$ in (104) is negative, as $d_{1}>1 / 2$. For $x_{1} \geq 1$, 104 is equal to

$$
\begin{aligned}
& \frac{1}{\omega\left(1-d_{1}\right)}\left(\left(2-4 d_{1}\right) x_{1}+1-3 d_{1}+\left(d_{1}+d_{1}^{2}\right) \sum_{i=0}^{x_{1}} d_{1}^{i}\right) \\
& =\frac{1}{\omega\left(1-d_{1}\right)}\left(\left(2-4 d_{1}\right)\left(x_{1}-1\right)+d_{1}^{3}\left(1+d_{1}\right) \sum_{i=0}^{x_{1}-2} d_{1}^{i}+d_{1}\left(1+d_{1}\right)^{2}+3-7 d_{1}\right) \\
& \leq \frac{1}{\omega\left(1-d_{1}\right)}\left(\left(2-4 d_{1}\right)\left(x_{1}-1\right)+d_{1}^{3}\left(1+d_{1}\right)\left(x_{1}-1\right)+d_{1}\left(1+d_{1}\right)^{2}+3-7 d_{1}\right) \\
& =\frac{1}{\omega\left(1-d_{1}\right)}\left(\left(d_{1}^{4}+d_{1}^{3}-4 d_{1}+2\right)\left(x_{1}-1\right)+d_{1}\left(1+d_{1}\right)^{2}+3-7 d_{1}\right) \\
& =\frac{-d_{1}^{3}-2 d_{1}^{2}-2 d_{1}+2}{\omega}\left(x_{1}-1\right)+\frac{-d_{1}^{2}-3 d_{1}+3}{\omega} \\
& <0
\end{aligned}
$$

where the third line follows from (103), and the last line from the fact that when $0.8 \leq d_{1}<1$, both terms $-d_{1}^{3}-2 d_{1}^{2}-2 d_{1}+2$ and $-d_{1}^{2}-3 d_{1}+3$ are negative. Next, we notice that 105 is equal to

$$
\begin{aligned}
& x_{2}\left(-\frac{2 d_{2}}{1-d_{2}}+\sqrt{\omega+1} \beta_{\theta, \omega}^{p}\right)+\sqrt{\omega+1} \beta_{\theta, \omega}^{p}+\frac{\left(1-d_{2}^{x_{2}}\right)\left(d_{2}+d_{2}^{2}\right)}{\left(1-d_{2}\right)^{2}} \\
& \leq x_{2}\left(-\frac{2 d_{2}}{1-d_{2}}+\sqrt{\omega+1} \beta_{\theta, \omega}^{p}\right)+\frac{d_{2}+d_{2}^{2}}{1-d_{2}} x_{2}+\sqrt{\omega+1} \beta_{\theta, \omega}^{p} \\
& =x_{2}\left(-\frac{2 d_{2}}{1-d_{2}}+\frac{d_{2}+d_{2}^{2}}{1-d_{2}}+\sqrt{\omega+1} \beta_{\theta, \omega}^{p}\right)+\sqrt{\omega+1} \beta_{\theta, \omega}^{p} \\
& =x_{2}\left(-d_{2}+\sqrt{\omega+1} \beta_{\theta, \omega}^{p}\right)+\sqrt{\omega+1} \beta_{\theta, \omega}^{p},
\end{aligned}
$$

where the second line follows from 103 . Taking $\beta_{\theta, \omega}^{p} \leq d_{2} / 2 \sqrt{\omega+1}$, we get

$$
x_{2}\left(-\frac{2 d_{2}}{1-d_{2}}+\sqrt{\omega+1} \beta_{\theta, \omega}^{p}\right)+\sqrt{\omega+1} \beta_{\theta, \omega}^{p}+\frac{\left(1-d_{2}^{x_{2}}\right)\left(d_{2}+d_{2}^{2}\right)}{\left(1-d_{2}\right)^{2}} \leq-\frac{d_{2}}{2} x_{2}+\frac{d_{2}}{2}
$$

which is non-positive for $x_{2} \geq 1$. Finally, when $0.8 \leq d_{1}<1, x_{1}, x_{2}>0$, and $x_{1}+1 \leq \omega\left(x_{2}+1\right)$, (76) holds for $\beta_{\theta, \omega}^{p} \leq d_{2} / 2 \sqrt{\omega+1}$.

Two. $d_{1}<0.8$ : Taking $\beta_{\theta, \omega}^{p} \leq \frac{d_{2}}{\sqrt{\omega+1}\left(1-d_{2}\right)}$, we note that the coefficient of $x_{2}$ in (105) is negative. Thus, from $x_{1}+1 \leq \omega\left(x_{2}+1\right)$, 104) and 105,

$$
\begin{align*}
& \Delta V_{\theta, \omega}^{p}(\boldsymbol{x})+\sqrt{\omega+1} \beta_{\theta, \omega}^{p}\left(x_{2}+1\right) \\
& \leq \frac{x_{1}+1}{\omega}\left(2-\frac{2 d_{1}}{1-d_{1}}\right)-\frac{1}{\omega}+\frac{\left(1-d_{1}^{x_{1}+1}\right)\left(d_{1}+d_{1}^{2}\right)}{\omega\left(1-d_{1}\right)^{2}} \\
& +\frac{x_{1}+1}{\omega}\left(-\frac{2 d_{2}}{1-d_{2}}+\sqrt{\omega+1} \beta_{\theta, \omega}^{p}\right)+\frac{2 d_{2}}{1-d_{2}}+\frac{\left(1-d_{2}^{x_{2}}\right)\left(d_{2}+d_{2}^{2}\right)}{\left(1-d_{2}\right)^{2}} \\
& <\frac{x_{1}+1}{\omega}\left(2-\frac{2 d_{1}}{1-d_{1}}-\frac{2 d_{2}}{1-d_{2}}+\sqrt{\omega+1} \beta_{\theta, \omega}^{p}\right)+\frac{2 d_{2}}{1-d_{2}}+\frac{d_{1}+d_{1}^{2}}{\omega\left(1-d_{1}\right)^{2}}+\frac{d_{2}+d_{2}^{2}}{\left(1-d_{2}\right)^{2}} \tag{106}
\end{align*}
$$

As $d_{i}=\tilde{\theta}_{i} /\left(\tilde{\theta}_{i}+\tilde{\lambda}\right)$ in terms of the normalized rates, we get

$$
2-\frac{2 d_{1}}{1-d_{1}}-\frac{2 d_{2}}{1-d_{2}}=2-\frac{2 \tilde{\theta}_{1}}{\tilde{\lambda}}-\frac{2 \tilde{\theta}_{1}}{\tilde{\lambda}}=\frac{-2\left(\tilde{\theta}_{1}+\tilde{\theta}_{2}-\tilde{\lambda}\right)}{\tilde{\lambda}}
$$

which is negative from the stability condition. For $\beta_{\theta, \omega}^{p} \leq \frac{\tilde{\theta}_{1}+\tilde{\theta}_{2}-\tilde{\lambda}}{\tilde{\lambda} \sqrt{\omega+1}}$, from (106) we get

$$
\begin{aligned}
& \Delta V_{\theta, \omega}^{p}(\boldsymbol{x})+\sqrt{\omega+1} \beta_{\theta, \omega}^{p}\left(x_{2}+1\right) \\
& <\frac{-\left(\tilde{\theta}_{1}+\tilde{\theta}_{2}-\tilde{\lambda}\right)}{\omega \tilde{\lambda}}\left(x_{1}+1\right)+\frac{2 d_{2}}{1-d_{2}}+\frac{d_{1}+d_{1}^{2}}{\omega\left(1-d_{1}\right)^{2}}+\frac{d_{2}+d_{2}^{2}}{\left(1-d_{2}\right)^{2}} \\
& =\frac{-\left(\tilde{\theta}_{1}+\tilde{\theta}_{2}-\tilde{\lambda}\right)}{\omega \tilde{\lambda}}\left(x_{1}+1\right)+\frac{2 \tilde{\theta}_{2}}{\tilde{\lambda}}+\frac{\tilde{\theta}_{1}\left(2 \tilde{\theta}_{1}+\tilde{\lambda}\right)}{\omega \tilde{\lambda}^{2}}+\frac{\tilde{\theta}_{2}\left(2 \tilde{\theta}_{2}+\tilde{\lambda}\right)}{\tilde{\lambda}^{2}}
\end{aligned}
$$

which is non-positive for

$$
x_{1}+1 \geq \frac{\tilde{\theta}_{1}\left(2 \tilde{\theta}_{1}+\tilde{\lambda}\right)+\omega \tilde{\theta}_{2}\left(2 \tilde{\theta}_{2}+3 \tilde{\lambda}\right)}{\tilde{\lambda}\left(\tilde{\theta}_{1}+\tilde{\theta}_{2}-\tilde{\lambda}\right)}
$$

As $d_{1}<0.8$, we can see that $\tilde{\lambda}>\tilde{\theta}_{1} / 4$; thus,

$$
\frac{\tilde{\theta}_{1}\left(2 \tilde{\theta}_{1}+\tilde{\lambda}\right)+\omega \tilde{\theta}_{2}\left(2 \tilde{\theta}_{2}+3 \tilde{\lambda}\right)}{\tilde{\lambda}\left(\tilde{\theta}_{1}+\tilde{\theta}_{2}-\tilde{\lambda}\right)}<\frac{4 \tilde{\theta}_{1}\left(2 \tilde{\theta}_{1}+\tilde{\lambda}\right)+4 \omega \tilde{\theta}_{2}\left(2 \tilde{\theta}_{2}+3 \tilde{\lambda}\right)}{\tilde{\theta}_{1}\left(\tilde{\theta}_{1}+\tilde{\theta}_{2}-\tilde{\lambda}\right)}<\frac{4 c_{R} R(1+2 \tilde{\lambda})}{\delta} \leq 4 c_{R} R,
$$

where we have used the fact that $\tilde{\theta}_{1} \geq \tilde{\theta}_{2}, \omega \leq c_{R} R, \tilde{\theta}_{1}+\tilde{\theta}_{2}-\tilde{\lambda} \geq \delta$, and $\tilde{\lambda} \leq 0.5-0.5 \delta$ and it suffices for $x_{1}$ to be greater than or equal to $4 c_{R} R$. For $x_{1}<4 c_{R} R$, 104) can be upper bounded as

$$
\frac{8 c_{R} R}{\omega}+\frac{1-3 d_{1}}{\omega\left(1-d_{1}\right)}+\frac{d_{1}+d_{1}^{2}}{\omega\left(1-d_{1}\right)^{2}} \leq \frac{8 c_{R} R}{\omega}+\frac{2}{\omega\left(1-d_{1}\right)^{2}}<\frac{8 c_{R} R+50}{\omega}
$$

where in the last inequality we have used $d_{1}<0.8$. From (105) and taking $\beta_{\theta, \omega}^{p} \leq d_{2} / 2 \sqrt{\omega+1}$,

$$
\begin{aligned}
& \Delta V_{\theta, \omega}^{p}(\boldsymbol{x})+\sqrt{\omega+1} \beta_{\theta, \omega}^{p}\left(x_{2}+1\right) \\
& \leq \frac{8 c_{R} R+50}{\omega}+\left(-\frac{2 d_{2}}{1-d_{2}}+\frac{d_{2}}{2}\right)\left(x_{2}+1\right)+\frac{2 d_{2}}{1-d_{2}}+\frac{\left(1-d_{2}^{x_{2}}\right)\left(d_{2}+d_{2}^{2}\right)}{\left(1-d_{2}\right)^{2}} \\
& \leq\left(-\frac{2 d_{2}}{1-d_{2}}+\frac{d_{2}}{2}+\frac{d_{2}+d_{2}^{2}}{1-d_{2}}\right) x_{2}+\frac{d_{2}}{2}+\frac{8 c_{R} R+50}{\omega} \\
& =-\frac{d_{2}}{2} x_{2}+\frac{d_{2}}{2}+\frac{8 c_{R} R+50}{\omega}
\end{aligned}
$$

which is negative for

$$
x_{2} \geq 1+\frac{16 c_{R} R+100}{\omega d_{2}}
$$

Finally, when $x_{1}+1 \leq \omega\left(x_{2}+1\right)$ and $x_{1}, x_{2}>0$, 76) holds for $\beta_{\theta, \omega}^{p} \leq$ $\frac{1}{\sqrt{\omega+1}} \min \left(\frac{\tilde{\theta}_{2}}{2\left(\tilde{\theta}_{2}+\tilde{\lambda}\right)}, \tilde{\theta}_{1}+\tilde{\theta}_{2}-\tilde{\lambda}\right), x_{1} \geq 4 c_{R} R$, and $x_{2} \geq 1+\frac{16 c_{R} R+100}{\omega d_{2}}$. Repeating the same arguments when $x_{1}, x_{2}>0$ and $x_{1}+1>\omega\left(x_{2}+1\right)$, 76 holds for $\beta_{\theta, \omega}^{p} \leq$ $\frac{1}{\sqrt{\omega+1}} \min \left(\frac{\tilde{\theta}_{1}}{2\left(\tilde{\theta}_{1}+\tilde{\lambda}\right)}, \tilde{\theta}_{1}+\tilde{\theta}_{2}-\tilde{\lambda}\right), x_{1} \geq 1+\frac{\omega\left(16 c_{R} R^{2}+100\right)}{d_{1}}$, and $x_{2} \geq 4 c_{R} R^{2}$. Finally, (76) holds with

$$
\begin{aligned}
& V_{\theta, \omega}^{p}(\boldsymbol{x})=\frac{x_{1}^{2}}{\omega}+x_{2}^{2} \\
& C_{\theta, \omega}^{p}=\left\{\left(x_{1}, x_{2}\right) \in \mathcal{X}: x_{i} \leq\left(16 c_{R}^{2} R^{3-i}+101 c_{R} R\right) \frac{\lambda+\theta_{i}}{\theta_{i}}, i=1,2\right\} \\
& \beta_{\theta, \omega}^{p}=\min \left(\frac{\tilde{\theta}_{2}}{2\left(\tilde{\theta}_{2}+\tilde{\lambda}\right) \sqrt{\omega+1}}, \frac{\tilde{\theta}_{1}+\tilde{\theta}_{2}-\tilde{\lambda}}{\sqrt{\omega+1}}, \frac{\tilde{\theta}_{2}}{2\left(\tilde{\theta}_{2}+\tilde{\lambda}\right)}, \frac{\tilde{\theta}_{1}}{2\left(\tilde{\theta}_{1}+\tilde{\lambda}\right) \sqrt{\omega}}\right) \\
& b_{\theta, \omega}^{p}=\left(\beta_{\theta, \omega}^{p}+1\right) \max _{x \in C_{\theta, \omega}^{p}}\left(\frac{\left(x_{1}+1\right)^{2}}{\omega}+\left(x_{2}+1\right)^{2}\right) \\
& \alpha_{\theta, \omega}^{p}=\frac{1}{2}
\end{aligned}
$$

where the fourth line holds since $P V_{\theta, \omega}^{p}(\boldsymbol{x}) \leq V_{\theta, \omega}^{p}(\boldsymbol{y})$ for $\boldsymbol{y}=\left(y_{1}, y_{2}\right)$ such that $y_{i}=x_{i}+1$ for $i=1,2$.

## G Numerical results

## G. 1 Comparison of Algorithm 1 with other learning algorithms

We first note that due to the countably infinite state-space setting of our problem, we are unable to directly compare our algorithm to other learning algorithms proposed in the literature. One potential candidate algorithm uses the reward biased maximum likelihood estimation (RBMLE) [33, 34, 11, 42], which estimates the unknown model parameter with the likelihood perturbed a vanishing bias towards parameters with a larger long-term average reward (i.e., optimal value). This scheme also uses the principle of "optimism in the face of uncertainty" in how it perturbs the maximum likelihood estimate. The naive version of the RMBLE algorithm does not apply to our examples due the following key assumption: over all parameters (and the control policies used for them), the transition probabilities are assumed to be mutually absolutely continuous; this is critical for the proofs and also allows the use of log-likelihood functions for computations. Similarly, naive use of the algorithms in [36] and [24] is not possible, again due to a similar absolutely continuity assumption which is critical for the proofs. Our posterior computations avoid such issues as the true parameter always has non-zero mass during the execution of the algorithm: episode $k$ always starts in state $0^{d}$ which is positive recurrent for the Markov chain with true parameter $\boldsymbol{\theta}^{*}$ and policy used $\pi_{\theta_{k}}^{*}$. The RBMLE algorithm has yet another issue in that it requires knowledge of the optimal value function, and hence, for our examples, it may only apply to Model 1 for which the value function is known analytically. Finally, whereas we do get to observe inter-arrival times for both model, we never directly observe completed service times owing to the sampling employed, and this precludes the direct use of Upper-Confidence-Bound based parameter estimation followed by certainty equivalent control algorithms. Owing to these issues, at this point in time, we're unable to perform empirical comparisons of Algorithm 1 to other candidate algorithms with theoretical performance guarantees in a countable state setting.
As discussed in the previous paragraph, learning algorithms with theoretical performance guarantees are established in the finite state setting. One such algorithm is the certainty equivalence control with forcing, which is proposed and discussed in detail in [4]. To assess the finite-time performance of our algorithm, in Figure 4, we compare the performance of our proposed learning algorithm, denoted as TSDE, with the algorithm introduced in [4], referred to as AgrawalTeneketzis. Reference [4] proposes a certainty equivalence control law with forced exploration, which operates in episodes with increasing lengths and a priori fixed sequences of forcing times. Specifically, at the beginning of each episode, all possible stationary control laws are explored for one recurrence interval of state $(0,0)$. Subsequently, based on this exploration, an empirical estimate of the average collected reward


Figure 4: Comparison of the regret performance of Algorithm 1 (referred to as TSDE) with the algorithm proposed by [4] (denoted as AgrawalTeneketzis) for the queueing models of Figure 2 .


Figure 5: Total variation distance between the posterior and real distribution for $\lambda=0.3,0.5,0.7$. The y axis is plotted on a logarithmic scale to display the differences clearly.
is formed, and the control law resulting in the maximum average reward is implemented for the remainder of the episode. The length of the episodes are determined according to sequence $\left\{a_{i}\right\}_{i=0}^{\infty}$ defined as following:

$$
\begin{aligned}
& a_{0}=0 \\
& a_{i}=\sum_{k=1}^{i} b_{k}+i p, \quad \text { for } i \geq 1
\end{aligned}
$$

where $p$ is the number of possible stationary control laws and $b_{i}=\left\lfloor\exp \left(i^{\frac{1}{1+\delta}}\right)\right\rfloor$ for any $\delta>0$. Specifically, episode $i$ terminates after completing additional $a_{i}-a_{i-1}$ recurrence intervals to state $(0,0)$. Both algorithms are implemented in the two queueing systems of Figure 2, where the arrival rate is $\lambda=0.5$ and service rates are distributed according to a Dirichlet prior over $[0.5,1.9]^{2}$. In Figures 4 a and 4 b , we set $\delta=3.5$ and $\delta=3$, respectively. Moreover, in Figure 4b, the goal is to find the optimal weight $w$ in the set $\{1.5,2,2.5,3,3.5\}$. The results in Figure 4 show that both algorithms exhibit a sublinear regret performance. Specifically, Algorithm 1 , TSDE, achieves an $\tilde{O}(\sqrt{T})$ as predicted in our theoretical results of Theorem 1 and Corollary 1 Furthermore, in both queueing models, our proposed algorithm consistently outperforms the algorithm presented in [4] in terms of regret order.


(a) Model 1: Queueing system of Figure 2 a

Figure 6: Optimal policy parameters for different service rate vectors in the two exemplary queuing systems in Model 1 and Model 2 with $\lambda=0.5$.


Figure 7: Estimated average cost of Model 2 for three different service rate vectors.

## G. 2 Model 1: Two-server queueing system with a common buffer.

Figure 3 b illustrates the behavior of the regret of Model 1 for three different arrival rate values and averaged over 2000 simulation runs. In these simulations, the parameter space is selected as

$$
\Theta=\left\{\left(\theta_{1}, \theta_{2}\right) \in[0.5,0.6, \ldots, 1.9]^{2}: \lambda<\theta_{1}+\theta_{2}, \theta_{2}<\theta_{1}\right\}
$$

which results in a prior size of 105. As depicted in Figure 3a the regret has a sub-linear behavior and increases with the arrival rate. The total variation distance between the posterior and real distribution, a point-mass on the random $\boldsymbol{\theta}^{*}$, are plotted in Figure 5 a . As expected, the distance diminishes towards 0 , indicating the learning of the true parameter. As mentioned in Appendix E.1, the optimal policy minimizing the average number of jobs in a system with parameter $\theta$, is a threshold policy $\pi_{t(\theta)}$ with optimal finite threshold $t(\theta) \in \mathbb{N}$, which can be numerically determined as the smallest $i \in \mathbb{N}$ for which $J^{i}(\theta)<J^{i+1}(\theta)$, calculated in [38]. We compute the optimal threshold $t(\theta)$ for every $\theta \in \Theta$ and present the results in Figure 6a. We can see that the threshold increases as the ratio of the service rates grows. Specifically, this is why in Appendix $G$ we imposed conditions on $\Theta$ to ensure that the ratio between the service rates is both upper and lower bounded.

## G. 3 Model 2: Two heterogeneous parallel queues

Figure 3billustrates the behavior of the regret of Model 2 for three different arrival rate values and averaged over 2000 simulation runs. We note that the regret is sub-linear and increases with higher arrival rates. In these simulations, the parameter space is selected as

$$
\Theta=\left\{\left(\theta_{1}, \theta_{2}\right) \in[0.5,0.7, \ldots, 1.9]^{2}: \lambda<\theta_{1}+\theta_{2}, \theta_{2}<\theta_{1}\right\}
$$

which results in a prior size of 28 . As discussed earlier, our goal is to find the average cost minimizing policy within the class of policies $\Pi=\left\{\pi_{\omega} ; \omega \in\left[\left(c_{R} R\right)^{-1}, c_{R} R\right]\right\}, c_{R} \geq 1$, where $\pi_{\omega}(\boldsymbol{x})=$
$\arg \min \left(1+x_{1}, \omega\left(1+x_{2}\right)\right)$ with ties broken for 1 . As discussed before, even with the transition kernel fully specified (by the values of arrival and service rates), the optimal policy in $\Pi$ is not known except when $\theta_{1}=\theta_{2}$ where the optimal value is $\omega=1$, and so, to learn it, we will use Proximal Policy Optimization with approximating martingale-process (AMP) method for countable state-space MDPs [18]. We run the algorithm for 200 policy iterations, using 20 actors for each iteration. We take the state $(0,0)$ as a regeneration state and simulate 1500 independent regenerative cycles per actor in each algorithm iteration. To approximate the value function, we employ a fully connected feed-forward neural network with one hidden layer consisting of $10 \times 10$ units and ReLU activation functions. The AMP method is also employed for variance reduction in value function estimation. The optimal $\omega$ for every $\theta \in \Theta$ is shown in Figure 6b indicating that $\omega$ increases as the ratio of the service rates grows. Therefore, it is necessary to ensure that the ratio between the service rates is bounded from above and below. Furthermore, to evaluate the regret numerically, the value of $J(\theta)$ is required for every $\theta \in \Theta$, which is not known. Thus, after finding the optimal $\omega$ using the PPO algorithm, we perform a separate simulation to approximate the optimal average cost. In Figure 7, we plot the estimated average cost for three different service rate vectors, demonstrating that the optimal average cost decreases as the service rates increase. In Figure 5 b we also depict the total variation distance between the posterior and real distribution, which is a point-mass on the random $\boldsymbol{\theta}^{*}$, and observe that the distance is converging to zero.


[^0]:    ${ }^{1}$ When $\theta_{1}=\theta_{2}$, then the policy with $\omega=1$ (Join-the-Shortest-Queue) is the optimal policy [19] for the underling MDP.

