

Uncertainty Quantification and Exploration for Reinforcement Learning

Yi Zhu

Department of Industrial Engineering and Management Sciences, Northwestern University, Evanston, IL 60208,
yizhu2020@u.northwestern.edu

Jing Dong

Division, Risk and Operations Division, Columbia Business School, New York, NY 10027, jing.dong@gsb.columbia.edu

Henry Lam

Department of Industrial Engineering and Operations Research, Columbia University, New York, NY 10027,
henry.lam@columbia.edu

We investigate statistical uncertainty quantification for reinforcement learning (RL) and its implications in exploration policy. Despite ever-growing literature on RL applications, fundamental questions about inference and error quantification, such as large-sample behaviors, appear to remain quite open. In this paper, we fill in the literature gap by studying the central limit theorem behaviors of estimated Q-values and value functions under various RL settings. In particular, we explicitly identify closed-form expressions of the asymptotic variances, which allow us to efficiently construct asymptotically valid confidence regions for key RL quantities. Furthermore, we utilize these asymptotic expressions to design an effective exploration strategy, which we call Q-value-based Optimal Computing Budget Allocation (Q-OCBA). The policy relies on maximizing the relative discrepancies among the Q-value estimates. Numerical experiments show superior performances of our exploration strategy than other benchmark policies.

Key words: reinforcement learning, statistical inference, policy exploration, OCBA

1. Introduction

We consider the standard reinforcement learning (RL) setting where the agent interacts with a random environment and aims to maximize the accumulated discounted reward over time. The environment is formulated as a Markov decision process (MDP) and the agent is uncertain about the true dynamics to start with. As the agent interacts with the environment, data about system dynamics are collected and the agent becomes increasingly confident about her decision. With finite data, however, the potential reward from each decision is estimated with errors. The goal of the agent is to make decisions as close to optimal as possible over time, given the collected data at hand.

In this paper, we focus on statistical uncertainty quantification in RL and its implications in exploration policy. More precisely, we investigate the large-sample behaviors in estimating the so-called Q-values (action-value functions) and the associated optimal value functions, in the form

of central limit convergences to Gaussian (or other) distributions. We explicitly characterize the asymptotic variances of the limiting distributions in terms of the system parameters (i.e., transition probabilities, mean rewards, and data-collection policies), from which we can construct statistically valid confidence intervals/regions. These closed-form formulas, moreover, are utilized in a novel procedure that we call Q-value-based Optimal Computing Budget Allocation (Q-OCBA) to efficiently explore the underlying MDP.

The motivation for our investigation is twofold. First, like classical statistical inference, large-sample behavior and confidence interval construction are fundamental in assessing the error and reliability of estimated quantities with respect to the data noises. Despite the ever-growing literature on RL applications, some of these fundamental questions appear to remain quite open. Our first goal in this paper is to fill in the literature gap by studying the large sample behaviors of RL. In the RL context, the key objects that measure performances and determine optimal policies are the Q-values, namely the maximum cumulative rewards that initialize at given state-and-action pairs. Closely related and derivable from the Q-values are the optimal value functions, namely the maximum cumulative rewards initialized at given states. We derive appropriate large-sample asymptotic distributions that allow us to construct confidence regions for Q-value and optimal-value-function estimations when data are collected from the underlying Markov chain. In addition to quantifying the error of each value estimate, our results also allow one to evaluate the assertiveness of performances among different decisions. For example, while one state-action pair may appear better by reading the point estimate of its Q-value, its variability, which is estimated via our approach, can also be larger, and should be properly accounted for when deciding whether the action should be selected given the state.

Our investigation on the limit theorems and confidence region constructions is in line with the so-called input uncertainty problem in stochastic simulation (Barton 2012, Song et al. 2014, Lam 2016). The latter problem aims to quantify the impacts on simulation outputs due to the statistical noises incurred in fitting the input models that generate the fed-in random variates. Quantifying these impacts requires estimating the output variability contributed from the input noises (e.g., Song and Nelson 2015, Cheng and Holland 1997), or constructing output confidence intervals that properly account for input uncertainty (e.g., Cheng and Holland 2004, Barton et al. 2013). In stochastic simulation, the inputs often take the form of parametric models (e.g., Xie et al. 2014, Zhu et al. 2020), a finite collection or mixture of parametric models (e.g., Chick 2001, Zouaoui and Wilson 2004), or in the nonparametric case, empirical distributions where independence of random variates is typically assumed (e.g., Barton and Schruben 2001, Yi and Xie 2017). Our study can be viewed as addressing the input uncertainty problem for MDPs, where the input model now takes the form of Markov transition matrices and reward distributions. In addition to the Markovian

structure, another distinction of our work from conventional input uncertainty is the involvement of optimization in our target quantity (both the Q-values and the optimal value functions). In this regard, our work also relates to the uncertainty quantification of stochastic optimization problems, which aims to understand the sub-optimality of obtained solutions due to the statistical noise in constructing the objective functions (Mak et al. 1999, Bayraksan and Morton 2006). Similar to the sample average approximation and stochastic programming literature (e.g., Shapiro et al. 2014, Higle and Sen 2013), our limit theorems capture the statistical fluctuations of estimated optimal values. However, instead of a finite number of stages or having independent variates across stages, we focus on the error quantification of maximum cumulative rewards, which necessitate the intricate use of the Bellman equation machinery and the stationarity of the underlying Markov chain.

Besides input uncertainty, we describe our methodological advancements relative to the RL literature. Our results complement the existing finite-sample bounds (e.g., Kearns and Singh 1998, Kakade 2003, Munos and Szepesvári 2008) by offering closed-form asymptotic variances that often show up in the first-order terms in these bounds. In constructing confidence intervals (and also in using variance-based exploration strategies like Q-OCBA that we detail further below), finite-sample error bounds are typically conservative. In contrast, in our development, the dependence of asymptotic variance on key parameters, such as the reward and transition variabilities, and the frequency of visits for each state-action pair, is tightly characterized. To this end, Mannor et al. (2004, 2007) are among the first and very few to study asymptotically tight statistical properties of value function estimation. They investigate the bias and variance in value function estimates under a fixed policy. Methodologically, our technique resolves a main technical challenge in Mannor et al. (2004, 2007) that allows us to generalize their variance results to Q-values and optimal value functions. The derivation in Mannor et al. (2004, 2007) hinges on an expansion of the value function in terms of the perturbation of the transition matrix, which (as pointed out by the authors) is not easily extendable from the value function under a fixed-policy to the optimal value function. In contrast, our results utilize an implicit function theorem applied to the Bellman equation which can be verified to be sufficiently smooth. This idea turns out to allow us to obtain gradients for Q-values, translate the asymptotic variance of Q-values to the optimal value function, and furthermore generalize similar asymptotic distributional results for constrained MDPs, approximate value iteration, and kernel estimation. Recently, Devraj and Meyn (2017) and Chen et al. (2020) derive central limit convergence results and characterize the corresponding asymptotic variances for some parametric Q-learning algorithms. Their development builds on an elegant connection between Q-learning and stochastic approximation. On the other hand, our main asymptotic result (Theorem 1) does not impose any parametric assumption on the Q-values.

Our work is also related to the line of studies on dynamic treatment regimes (DTR) applied commonly in medical decision-making (Laber et al. 2014). DTR focuses on the statistical properties of policies on finite horizons (such as two-period). Our infinite-horizon results on the optimal values and Q-values distinguish our developments from the DTR literature. Indeed, Laber et al. (2014) list the statistical properties for the infinite-horizon case as an “open problem”. Moreover, as we will elaborate in the sequel, our result for the non-unique policy case can be demonstrated to correspond to the “non-regularity” concept in DTR, where the true parameters are very close to the decision “boundaries” that switch the optimal policy (motivated by situations of small treatment effects), thus making the obtained policy highly sensitive to estimation noises.

The second motivation of our study is to design good exploration policies by directly using our tight error estimates derived from our limit theorems. We consider the pure exploration setting where an agent is first assigned a period to collect as much experience as possible, and then, with the optimal policy trained offline, starts deployment to gain reward. The goal is to collect as informative data as possible in the first stage via a good exploration strategy such that, at the end of this stage, our recommended policy coincides with the true optimal policy of the MDP with the highest probability. This setting is motivated by recent autonomous-driving applications (Kalashnikov et al. 2018), and differs from classical learning settings that balance the exploration-exploitation trade-off: exploring poorly-understood states and actions in the hope of improving future performance versus exploiting existing knowledge to attain better reward now (Jaksch et al. 2010). Crucially, our setting has a dedicated learning stage in which no rewards are collected, and the focus is to select the best policy using the data collected in the learning stage. The separation of learning and execution stages is particularly attractive when extra safety-protection tools could be incorporated into the learning stage. For example, to derive safe driving policies for autonomous vehicles, “soft crash” (or “near crash”) perceived by high-precision sensors is used to reveal the potential risks of unsafe driving behaviors and avoids real physical damage to the car (Kiran et al. 2020, Amini et al. 2020). The learning stage can also sometimes be carried out through a simulation platform where real maneuvers could be virtually simulated (Schöner 2017, Corso et al. 2020).

Our main contribution in regard to the second motivation is the design of a pure exploration policy by maximizing the worst-case relative discrepancy among the estimated Q-values (ratio of the mean squared difference to the variance). We coin our strategy Q-OCBA, as our policy search criterion resembles the widely known Optimal Computing Budget Allocation (OCBA) procedure in simulation optimization (Chen and Lee 2011). OCBA has been successfully applied to ranking and selection (Glynn and Juneja 2004, Gao et al. 2017) where the search space is a finite number of discrete simulatable alternatives, a variant called targeting and selection (Ryzhov 2018), random search (Chen et al. 2013), stochastically constrained problems and feasibility determination (Lee

et al. 2012, Gao and Chen 2016), and quantile-based selection (Peng et al. 2018a). In sequential OCBA implementations, one divides the computation budget into stages in which one sequentially updates mean and variance estimates, and optimizes next-stage budget allocations according to the worst-case relative discrepancy criterion (Chen et al. 2006). Like OCBA, the worst-case relative discrepancy in Q-OCBA serves as a proxy for the probability of selecting the optimal policy, which is analogous to the probability of correct selection in ranking and selection. Nonetheless, Q-OCBA bears two important distinctions: One is that we utilize the Q-value estimates and asymptotic variance estimates, derived via our limit theorems, in our maximization criterion. The other is that while the allocation in OCBA can be derived and implemented straightforwardly, it is not the case in RL. Because of the underlying Markov chain, the desired allocation has to be implemented through a randomized policy that gives rise to a stationary distribution of state-action pairs that matches the desired allocation. To address this issue, we design a computationally tractable scheme to obtain the randomized policy. More precisely, through our derived asymptotic variance formula that depends explicitly on the frequency of visits to each state-action pair, and by characterizing all the feasible stationary distributions of the MDP as a set of linear constraints, we can derive our Q-OCBA exploration policy by solving a convex optimization problem. We further propose a sequential updating rule to implement the exploration strategy in practice.

To the best of our knowledge, our work is among one of the few to concretely use the variances of Q-value estimates and link budget allocation to develop exploration policies for MDP. The only other work that uses OCBA-based criterion for efficient budget allocation in MDP is Jia (2012), which is for a different setting. Similar to our work, the objective is to maximize the probability of selecting the optimal action at each state. However, they consider a setting where the state-action pair at each “step” can be arbitrarily chosen and the Q-values for different state-action pairs are estimated independently by simulation. This allows them to use standard sample variance for variance estimation. In contrast, we consider the “online” setting where the next state is determined by the current state-action pair through the underlying Markov chain. This poses a substantial challenge to constructing a good variance estimator. Interestingly, one way to further improve the efficiency of the strategy proposed in Jia (2012) is to reuse (at least partially) the sample simulated for the Q-value estimation at one state-action pair to estimate the Q-value at another state-action pair. However, this would require sophisticated derivations of the corresponding variances as in our setting. Another related work is Devraj and Meyn (2017), which uses asymptotic variances to optimize the corresponding Q-learning algorithm with certain parameterization. They consider the classic regret minimization objective, which is different from our pure exploration setting.

More generally, we also position Q-OCBA relative to the online learning literature. In particular, our pure exploration problem resembles the so-called best-arm identification in multi-armed

bandits (Audibert and Bubeck 2010, Kaufmann et al. 2016, Russo 2016), which studies procedures to efficiently identify alternatives with the highest reward. This problem is intricately related to ranking and selection in simulation (Kim and Nelson 2007) in that they both focus on the performance of the selected alternative at the end of the learning period, without concerning the interim rewards, but with differences in the analytical tools and the types of complexity guarantees. The ranking and selection or best-arm identification problems can be viewed as a one-state special case of the pure exploration setting in RL, which is much less studied (the only recent work that we are aware of is Putta and Tulabandhula (2017)). Compared to the former problem that often assumes independence across samples from different alternatives, RL is more complex due to the long-term effect of an action that could determine the future frequency of visits for many state-action pairs. Q-OCBA provides an implementable avenue to obtain a good policy for precisely this purpose. Recently, there are also extensions of ranking and selection to context-dependent settings, where rewards can depend on agent’s covariates (Shen et al. 2021). Although in both RL and contextual ranking and selection, the rewards are state-dependent, in RL, the state variables follow a Markov chain, whereas in contextual ranking and selection, the state variables (covariates) are static.

Most existing RL algorithms consider settings where one solves for the optimal policy while simultaneously learning the dynamics of the underlying MDP. The objective is to minimize the regret, which gives rise to a non-trivial exploration-exploitation trade-off. In this context, policies that focus on more systematic exploration such as posterior sampling reinforcement learning (PSRL) and various versions of the upper confidence bound algorithm (e.g., UCRL2, UCBVI) have been developed and shown to achieve good regret bounds (Osband et al. 2013, Azar et al. 2017, Jin et al. 2018). Recently, Bellemare et al. (2017) advocates learning an approximate value distribution rather than its expectation to obtain more stable learning. In this paper, we consider a very different setting – a pure exploration task. Thus, our policy focuses more on exploration than the aforementioned policies (which were designed for a different objective). To support our approach, we show through numerical experiments that Q-OCBA consistently achieves superior performances compared to an array of benchmark policies in the literature (see Section 6.2.1).

The rest of the paper is organized as follows. Section 2 first describes our MDP setup. Section 3 presents our results on large-sample behaviors. Section 4 demonstrates their use in developing exploration strategies. Section 5 studies two approximations schemes to deal with large state and/or action spaces. Section 6 substantiates our findings with simulation experiments. Section 7 provides some concluding remarks. All proofs are deferred to the appendix.

1.1. Notations

We introduce some notations used throughout the paper. We denote “ \Rightarrow ” as “convergence in distribution”, and $\mathcal{N}(\mu, \Sigma)$ as a multivariate Gaussian distribution with mean vector μ and covariance

matrix Σ . We write I as the identity matrix, and e_i as the i -th unit vector. We also write $\mathbf{0}$ as a zero vector or matrix. The dimension of $\mathcal{N}(\mu, \Sigma)$, I , e_i , and $\mathbf{0}$ should be clear from the context. When not specified, all the vectors are column vectors. For $\epsilon > 0$, we define $a \pm \epsilon := [a - \epsilon, a + \epsilon]$, i.e., it is a closed interval with lower bound $a - \epsilon$ and upper bound $a + \epsilon$. Lastly, $\mathbb{1}(\cdot)$ denotes the indicator function.

2. Problem Setup

Consider an infinite horizon accumulated discounted reward MDP, $\mathcal{M} = (\mathcal{S}, \mathcal{A}, R, P, \gamma, \rho)$, where \mathcal{S} is the state space, \mathcal{A} is the action space, $R(s, a)$ denotes the random reward when the agent is in state $s \in \mathcal{S}$ and selects action $a \in \mathcal{A}$, $P(s'|s, a)$ is the probability of transitioning to state s' in the next epoch given the current state s and taken action a , γ is the discount factor, and ρ is the initial state distribution. The distribution of the reward R and the transition probability P are unknown to the agent.

We assume both \mathcal{S} and \mathcal{A} are finite sets. Without loss of generality, we denote $\mathcal{S} = \{1, 2, \dots, m_s\}$ and $\mathcal{A} = \{1, 2, \dots, m_a\}$. We also make the following stochasticity assumption:

ASSUMPTION 1. $R(s, a)$ has a finite mean $\mu_R(s, a)$ and a finite variance $\sigma_R^2(s, a)$, $\forall (s, a) \in \mathcal{S} \times \mathcal{A}$. For any given $(s, a) \in \mathcal{S} \times \mathcal{A}$, $R(s, a)$ and $s' \sim P(\cdot|s, a)$ are independent random variables.

A policy π is a mapping from each state $s \in \mathcal{S}$ to a probability measure over actions in \mathcal{A} . Specifically, we write $\pi(a|s)$ as the probability of taking action a when the agent is in state s and $\pi(\cdot|s)$ as the m_a -dimensional vector of action probabilities at state s . For convenience, we sometimes write $\pi(s)$ as the realized action given the current state s . The value function associated with a policy π is defined as

$$V^\pi(s) = \mathbb{E}^\pi \left[\sum_{t=0}^{\infty} \gamma^t R(s_t, \pi(s_t)) | s_0 = s \right]$$

with $s_{t+1}|(s_t, \pi(s_t)) \sim P(\cdot|s_t, \pi(s_t))$. The expected value function, under the initial distribution ρ , is denoted by

$$\chi^\pi = \sum_s \rho(s) V^\pi(s).$$

A policy π^* is said to be optimal if $V^{\pi^*}(s) = \max_\pi V^\pi(s)$ for all $s \in \mathcal{S}$. For convenience, we denote $V^* = V^{\pi^*}$ and $\chi^* = \sum_s \rho(s) V^*(s)$. The Q-value at $(s, a) \in \mathcal{S} \times \mathcal{A}$, denoted by $Q(s, a)$, is defined as $Q(s, a) = \mu_R(s, a) + \gamma \mathbb{E}[V^*(s')|s, a]$. Correspondingly, $V^*(s) = \max_a Q(s, a)$ and the Bellman equation for the Q-values takes the form

$$Q(s, a) = \mu_R(s, a) + \gamma \mathbb{E} \left[\max_{a'} Q(s', a') | s, a \right]. \quad (1)$$

Denoting the Bellman operator as $\mathcal{T}_{\mu_R, P}(\cdot)$, Q is a fixed point associated with $\mathcal{T}_{\mu_R, P}$, i.e., $Q = \mathcal{T}_{\mu_R, P}(Q)$. For most part of this paper, we make the following assumption about Q :

ASSUMPTION 2. For any state $s \in \mathcal{S}$, $\arg \max_{a \in \mathcal{A}} Q(s, a)$ is unique.

Under Assumption 2, the optimal policy π^* is unique and deterministic. Let

$$a^*(s) = \arg \max_{a \in \mathcal{A}} Q(s, a).$$

Then $\pi^*(a|s) = \mathbb{1}(a = a^*(s))$. We relax this assumption in Section 3.2.

We next introduce some statistical quantities arising from data. Suppose we have n observations (whose collection mechanism will be made precise later), which we denote as

$$\{(s_t, a_t, r_t(s_t, a_t), s'_t(s_t, a_t)) : 1 \leq t \leq n\},$$

where $r_t(s_t, a_t)$ is the realized reward at time t and $s'_t(s_t, a_t) = s_{t+1}$. We define the sample mean $\hat{\mu}_{R,n}$ and the sample variance $\hat{\sigma}_{R,n}^2$ of the random rewards as

$$\hat{\mu}_{R,n}(s = i, a = j) = \begin{cases} \frac{\sum_{1 \leq t \leq n} r_t(s_t, a_t) \mathbb{1}(s_t = i, a_t = j)}{\sum_{1 \leq t \leq n} \mathbb{1}(s_t = i, a_t = j)}, & \text{if } \sum_{1 \leq t \leq n} \mathbb{1}(s_t = i, a_t = j) > 0 \\ u_R^0(i, j) & \text{if } \sum_{1 \leq t \leq n} \mathbb{1}(s_t = i, a_t = j) = 0; \end{cases} \quad (2)$$

$$\hat{\sigma}_{R,n}^2(s = i, a = j) = \begin{cases} \frac{\sum_{1 \leq t \leq n} r_t(s_t, a_t)^2 \mathbb{1}(s_t = i, a_t = j)}{\sum_{1 \leq t \leq n} \mathbb{1}(s_t = i, a_t = j)} - \hat{\mu}_{R,n}(i, j)^2 & \text{if } \sum_{1 \leq t \leq n} \mathbb{1}(s_t = i, a_t = j) > 0 \\ s_R^0(i, j) & \text{if } \sum_{1 \leq t \leq n} \mathbb{1}(s_t = i, a_t = j) = 0, \end{cases}$$

where $u_R^0(i, j) \in \mathcal{R}$ and $s_R^0(i, j) \geq 0$ are some appropriate initial values for the mean and variance if no data is available.

Similarly, we define the empirical transition matrix \hat{P}_n as

$$\hat{P}_n(s' = k | s = i, a = j) = \begin{cases} \frac{\sum_{1 \leq t \leq n} \mathbb{1}(s_t = i, a_t = j, s'_t(s_t, a_t) = k)}{\sum_{1 \leq t \leq n} \mathbb{1}(s_t = i, a_t = j)} & \text{if } \sum_{1 \leq t \leq n} \mathbb{1}(s_t = i, a_t = j) > 0 \\ p^0(k | i, j) & \text{if } \sum_{1 \leq t \leq n} \mathbb{1}(s_t = i, a_t = j) = 0, \end{cases} \quad (3)$$

where p^0 with $p^0(k | i, j) \geq 0$ and $\sum_k p^0(k | i, j) = 1$ is some appropriate initialization of the transition matrix if no data is available.

Note that $u^0(i, j)$, $s^0(i, j)$, and $p^0(k | i, j)$ above can be any reasonable initialization when the corresponding state-action pairs have not been visited yet. For example, we can initialize $\hat{P}_n(s' = k | s = i, a = j) = 1/m_s$, $\hat{\mu}_{R,n}(s = i, a = j) = 0$, and $\hat{\sigma}_{R,n}^2(s = i, a = j) = 1$ if there is no good prior knowledge.

Furthermore, we define the $m_s \times m_s$ sampling covariance matrix $\Sigma_{P_{s,a}}$ (with one sample point of $\mathbb{1}(s_t = s, a_t = a)$) as

$$\Sigma_{P_{s,a}}(k_1, k_2) = \begin{cases} P(k_1 | s, a)(1 - P(k_1 | s, a)) & k_1 = k_2 \\ -P(k_1 | s, a)P(k_2 | s, a) & k_1 \neq k_2. \end{cases}, \text{ for } 1 \leq k_1 \leq m_s, 1 \leq k_2 \leq m_s. \quad (4)$$

With the data, we construct our estimate of Q , called \hat{Q}_n , via the empirical fixed point of the Bellman operator, i.e.,

$$\hat{Q}_n = \mathcal{T}_{\hat{\mu}_{R,n}, \hat{P}_n}(\hat{Q}_n).$$

Correspondingly, we write

$$\hat{V}_n^*(s) = \max_{a \in \mathcal{A}} \hat{Q}_n(s, a) \quad \text{and} \quad \hat{\chi}_n^* = \sum_{s \in \mathcal{S}} \rho(s) \hat{V}_n^*(s).$$

Throughout the paper we focus on the estimation errors due to the noise of collected data, and assume the MDP or Q-value evaluation given the estimated average reward and transaction matrix can be obtained exactly.

3. Quantifying Asymptotic Estimation Errors

In this section, we present an array of results regarding the asymptotic behaviors of \hat{Q}_n and \hat{V}_n^* as $n \rightarrow \infty$. We start with the standard case where each state has a unique optimal action, i.e., Assumption 2 holds. We then discuss two important extensions: non-unique optimal action and constrained MDP.

To prepare, we first make an assumption on our exploration policy π to gather data. Define the extended transition probability \tilde{P}^π under policy π as

$$\tilde{P}^\pi(s', a' | s, a) = P(s' | s, a) \pi(a' | s').$$

ASSUMPTION 3. *The Markov chain with transition probability \tilde{P}^π is irreducible.*

Under Assumption 3, as we have a finite state space and action space, \tilde{P}^π has a unique stationary distribution, denoted by w , which is equal to the long run frequency of visiting each state-action pair, i.e., for $i \in \mathcal{S}$ and $j \in \mathcal{A}$

$$w(i, j) = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{1 \leq t \leq n} \mathbb{1}(s_t = i, a_t = j),$$

where all $w(i, j)$'s are positive. Note that Assumption 3 is satisfied if for any two states s, s' , there exists a sequence of actions such that s' is attainable from s under P , and if π is sufficiently mixed, e.g., π satisfies $\pi(a | s) > 0$ for all s, a . We define

$$\hat{w}_n(i, j) := \frac{1}{n} \sum_{1 \leq t \leq n} \mathbb{1}(s_t = i, a_t = j),$$

which is the empirical frequency of visiting each state-action pair.

Let $N = m_s m_a$. In our subsequent algebraic derivations, we need to re-arrange μ_R , Q , and w as N -dimensional vectors. With a little abuse of notation, we define the following indexing rule: $(s = i, a = j)$ is re-indexed as $(i - 1)m_a + j$, e.g.,

$$\mu_R(i, j) = \mu_R((i - 1)m_a + j).$$

We also need to re-arrange \tilde{P}^π as an $N \times N$ matrix following the same indexing rule, e.g.,

$$\tilde{P}^\pi(i', j' | i, j) = \tilde{P}^\pi((i - 1)m_a + j, (i' - 1)m_a + j').$$

3.1. Limit Theorems under Sufficient Exploration

We first establish the asymptotic Normality of \hat{Q}_n under an exploration policy π :

THEOREM 1. *Under Assumptions 1 and 2, if the data are collected according to π satisfying Assumption 3, $\hat{Q}_n \rightarrow Q$ almost surely (a.s.) as $n \rightarrow \infty$. Moreover,*

$$\sqrt{n}(\hat{Q}_n - Q) \Rightarrow \mathcal{N}(\mathbf{0}, \Sigma) \quad \text{as } n \rightarrow \infty,$$

where

$$\Sigma = (I - \gamma \tilde{P}^{\pi^*})^{-1} W^{-1} (D_R + \gamma^2 D_Q) ((I - \gamma \tilde{P}^{\pi^*})^{-1})^T, \quad (5)$$

W , D_R and D_Q are $N \times N$ diagonal matrices with

$$\begin{aligned} W((i-1)m_a + j, (i-1)m_a + j) &= w(i, j), \\ D_R((i-1)m_a + j, (i-1)m_a + j) &= \sigma_R^2(i, j), \\ D_Q((i-1)m_a + j, (i-1)m_a + j) &= (V^*)^T \Sigma_{P_{i,j}} V^*, \end{aligned} \quad (6)$$

where $\Sigma_{P_{i,j}}$ is defined in (4).

In addition to the asymptotic Normality, a key result in Theorem 1 is an explicit characterization of the asymptotic variance Σ , which is derived using the delta method (Serfling 2009). Intuitively, it is the product of the sensitivities (i.e., gradient) of Q with respect to its parameters and the variances of the parameter estimates. Here, the parameters are μ_R and P and the corresponding gradients are $(I - \gamma \tilde{P}^{\pi^*})^{-1}$ and $(I - \gamma \tilde{P}^{\pi^*})^{-1} V^*$ respectively. The variances of these parameter estimates (i.e. (2) and (3)) involve $\sigma_R^2(i, j)$, $\Sigma_{P_{i,j}}$, and $w(i, j)$ (i.e., the proportion of samples to each state-action pair).

Using the relations that $V^*(s) = \max_{a \in \mathcal{A}} Q(s, a)$ and $\hat{V}_n^*(s) = \max_{a \in \mathcal{A}} \hat{Q}_n(s, a)$, we can leverage Theorem 1 to further establish the asymptotic Normality of \hat{V}_n^* and $\hat{\chi}_n^*$:

COROLLARY 1. *Under Assumptions 1, 2 and 3,*

$$\sqrt{n}(\hat{V}_n^* - V^*) \Rightarrow \mathcal{N}(0, \Sigma_V) \quad \text{and} \quad \sqrt{n}(\hat{\chi}_n^* - \chi^*) \Rightarrow \mathcal{N}(0, \sigma_\chi^2) \quad \text{as } n \rightarrow \infty$$

where

$$\Sigma_V = (I - \gamma P^{\pi^*})^{-1} (W^{\pi^*})^{-1} [D_R^{\pi^*} + \gamma^2 D_V^{\pi^*}] ((I - \gamma P^{\pi^*})^{-1})^T,$$

and

$$\sigma_\chi^2 = \rho^T \Sigma_V \rho,$$

where P^{π^*} is an $m_s \times m_s$ transition matrix with $P^{\pi^*}(i, j) = P(j|i, a^*(i))$, W^{π^*} , $D_R^{\pi^*}$ and $D_V^{\pi^*}$ are $m_s \times m_s$ diagonal matrices with

$$W^{\pi^*}(i, i) = w(i, a^*(i)), \quad D_R^{\pi^*}(i, i) = \sigma_R^2(i, a^*(i)), \quad \text{and} \quad D_V^{\pi^*}(i, i) = (V^*)^T \Sigma_{P_{i, a^*(i)}} V^*$$

respectively.

Theorem 1 and Corollary 1 can be used immediately for statistical inference. In particular, we can construct confidence regions for subsets of the Q-values jointly, or confidence intervals for linear combinations of the Q-values. For example, a quantity of interest that we will later utilize to design good exploration policies is $Q(s, a_1) - Q(s, a_2)$, i.e., the difference between action a_1 and a_2 when the agent is in state s . Define $\sigma_{\Delta Q}^2$ as

$$\sigma_{\Delta Q}^2(s, a_1, a_2) = (e_{(s-1)m_a+a_1} - e_{(s-1)m_a+a_2})^T \Sigma (e_{(s-1)m_a+a_1} - e_{(s-1)m_a+a_2}) \quad (7)$$

and its estimator $\hat{\sigma}_{\Delta Q, n}^2$ by replacing Q, V^*, σ_R^2, w, P with $\hat{Q}_n, \hat{V}_n^*, \hat{\sigma}_{R, n}^2, \hat{w}_n, \hat{P}_n$ in Σ . Then, the $100(1 - \alpha)\%$ confidence interval for $Q(s, a_1) - Q(s, a_2)$ takes the form

$$C_n(s, a_1, a_2; \alpha) = \left(\hat{Q}_n(s, a_1) - \hat{Q}_n(s, a_2) \right) \pm z_{\alpha/2} \hat{\sigma}_{\Delta Q, n}(s, a_1, a_2),$$

where $z_{\alpha/2}$ is the $(1 - \alpha/2)$ -quantile of $\mathcal{N}(0, 1)$, i.e., $\mathbb{P}(\mathcal{N}(0, 1) \leq z_{\alpha/2}) = 1 - \alpha/2$. This confidence interval is asymptotically valid in the sense that

$$\lim_{n \rightarrow \infty} \mathbb{P}(Q(s, a_1) - Q(s, a_2) \in C_n(s, a_1, a_2; \alpha)) = 1 - \alpha.$$

Following a similar approach as the proof of Theorem 1, we can also establish the asymptotic Normality for the estimated value function under any given policy $\tilde{\pi}$. In this case, the value function $V^{\tilde{\pi}}$ satisfies the equation

$$V^{\tilde{\pi}}(s) = \sum_a \mu_R(s, a) \tilde{\pi}(a|s) + \gamma \sum_a \tilde{\pi}(a|s) \sum_{s'} P(s'|s, a) V^{\tilde{\pi}}(s').$$

Denote the estimator of $V^{\tilde{\pi}}$ as $\hat{V}_n^{\tilde{\pi}}$, which satisfies

$$\hat{V}_n^{\tilde{\pi}}(s) = \sum_a \hat{\mu}_{R, n}(s, a) \tilde{\pi}(a|s) + \gamma \sum_a \tilde{\pi}(a|s) \sum_{s'} \hat{P}_n(s'|s, a) \hat{V}_n^{\tilde{\pi}}(s').$$

COROLLARY 2. *Under Assumptions 1 and 3,*

$$\sqrt{n}(\hat{V}_n^{\tilde{\pi}} - V^{\tilde{\pi}}) \Rightarrow \mathcal{N}(0, \Sigma_V^{\tilde{\pi}}) \text{ as } n \rightarrow \infty$$

where

$$\Sigma_V^{\tilde{\pi}} = (I - \gamma P^{\tilde{\pi}})^{-1} D^{\tilde{\pi}} ((I - \gamma P^{\tilde{\pi}})^{-1})^T,$$

$P^{\tilde{\pi}}$ is an $m_s \times m_s$ transition matrix with $P^{\tilde{\pi}}(i, j) = \sum_{1 \leq a \leq m_a} P(j|i, a) \tilde{\pi}(a|i)$, $D^{\tilde{\pi}}$ is an $m_s \times m_s$ diagonal matrix with

$$D^{\tilde{\pi}}(i, i) = \sum_{1 \leq j \leq m_a} \frac{\tilde{\pi}(j|i)^2}{w(i, j)} [(\gamma V^{\tilde{\pi}})^T \Sigma_{P_{i, j}} (\gamma V^{\tilde{\pi}}) + \sigma_R^2(i, j)].$$

Corollary 2 essentially recovers Corollary 4.1 in Mannor et al. (2007). Different from Mannor et al. (2007), we derive our results by using an implicit function theorem on the corresponding fixed-point equation to obtain the gradient of Q , viewing the latter as the solution to the equation and as a function of μ_R, P . This approach is able to generalize the results for a fixed policy in Mannor et al. (2007) to the optimal value functions, and provide distributional statements as Theorem 1 and Corollary 1 above. We also note that an alternative route to obtain our results is to conduct perturbation analysis on the linear programming (LP) representation of the MDP, which would also give gradient information of V^* (and hence Q as well). The implicit function approach, nonetheless, is more elementary and intuitive.

3.2. Non-Unique Optimal Policy

Theorem 1 above is developed assuming the optimal solution for the MDP is unique. A natural question to ask is what would happen if Assumption 2 does not hold. When the optimal solution for the MDP is not unique, the estimated \hat{Q}_n and \hat{V}_n^* may “jump” around different optimal actions, leading to a more complicated large-sample behavior. We elaborate on this next. Let

$$U = \{u \in \mathbb{R}^{m_s N + N} : \|u\| = 1\}.$$

THEOREM 2. *Suppose Assumptions 1 and 3 hold, but the optimal policy is not unique. Then there exists $K \geq 1$ distinct $m_s \times (Nm_s + N)$ matrices $\{G_k\}_{1 \leq k \leq K}$ and a deterministic partition of U , $\{U_k\}_{1 \leq k \leq K}$, i.e., $U = \cup_{1 \leq k \leq K} U_k$, such that*

$$\sqrt{n}(\hat{V}_n^* - V^*) \Rightarrow \sum_{k=1}^K G_k \mathbf{1}(Z/\|Z\| \in U_k) Z \quad \text{as } n \rightarrow \infty,$$

where $Z \sim \mathcal{N}(\mathbf{0}, \Sigma_{R,P})$,

$$\Sigma_{R,P} = \begin{pmatrix} W^{-1} D_R & \mathbf{0} \\ \mathbf{0} & D_P \end{pmatrix}, \quad D_P = \begin{pmatrix} \frac{\Sigma_{P_{1,1}}}{w(0m_a+1)} & & & \\ & \ddots & & \\ & & \frac{\Sigma_{P_{i,j}}}{w((i-1)m_a+j)} & \\ & & & \ddots \\ & & & & \frac{\Sigma_{P_{m_s,m_a}}}{w((m_s-1)m_a+m_a)} \end{pmatrix}, \quad (8)$$

and W and D_R are defined in (6).

To see Theorem 2, note that the non-uniqueness of the optimal solution can be formalized as a non-degeneracy in the LP representation of the MDP. In particular, when the optimal policy is not unique, the LP representation of the optimal value function has multiple optimal bases. When perturbing $[P, \mu_R]$ along a specific direction $u \in U$ for a small enough amount, at least one of the

bases remains optimal. Let B_u denote an optimal basis under the perturbation u and let \tilde{G}_{B_u} denote the derivative of the optimal solution to the perturbed LP with respect to $[P, \mu_R]$. Then, the corresponding directional derivative takes the form $u^T \tilde{G}_{B_u}$. Based on the above discussion, K in Theorem 2 is the number of distinct derivatives among the optimal bases, $\{G_k\}_{1 \leq k \leq K}$ denotes these derivatives, and $\{U_k\}_{1 \leq k \leq K}$ denotes the corresponding partition of U . In particular, for $u \in U_k$, the directional derivative of the optimal value takes the form $u^T G_k$. See Appendix A.4 for more details.

We note from Theorem 2 that if $K = 1$, we recover the case of Corollary 1. However, if $K > 1$, the limit distribution becomes non-Gaussian. This arises because the sensitivity to P or μ_R can be very different depending on the perturbation direction, which is a consequence of non-uniqueness of the solution. Lastly, we comment that Theorem 2 can be viewed as an MDP analog to the non-Gaussian limit in sample average approximation where an expected-value optimization may have multiple optimal solutions (Shapiro et al. 2014).

3.3. Constrained Problems

The MDP introduced in Section 2 does not have any constraint. In recent years, motivated by budgeted decision-making (Boutilier and Lu 2016) and safety-critical applications (Achiam et al. 2017, Chow et al. 2017), there are growing interests in constrained MDPs. The goal is to maximize the long-run accumulated discounted reward, $V^\pi(s)$, while ensure that a long-run accumulated discounted cost, denoted as $L^\pi(s) = \mathbb{E}[\sum_{t=0}^{\infty} \gamma^t C(s_t, \pi(s_t)) | s_0 = s]$, is less than some given budget η , i.e.,

$$\max_{\pi} \sum_s \rho(s) V^\pi(s) \quad \text{subject to} \quad \sum_s \rho(s) L^\pi(s) \leq \eta \quad (9)$$

We refer to $L^\pi(s)$ as the loss function. Assume data is generated as before. In addition, we have observations on the incurred cost $\{c_t(s_t, a_t) : 1 \leq t \leq n\}$. Define the sample mean and sample variance of the random costs as

$$\hat{\mu}_{C,n}(s=i, a=j) = \begin{cases} \frac{\sum_{1 \leq t \leq n} c_t(s_t, a_t) \mathbf{1}(s_t=i, a_t=j)}{\sum_{1 \leq t \leq n} \mathbf{1}(s_t=i, a_t=j)}, & \text{if } \sum_{1 \leq t \leq n} \mathbf{1}(s_t=i, a_t=j) > 0 \\ u_C^0(i, j) & \text{if } \sum_{1 \leq t \leq n} \mathbf{1}(s_t=i, a_t=j) = 0; \end{cases}$$

$$\hat{\sigma}_{C,n}^2(s=i, a=j) = \begin{cases} \frac{\sum_{1 \leq t \leq n} c_t(s_t, a_t)^2 \mathbf{1}(s_t=i, a_t=j)}{\sum_{1 \leq t \leq n} \mathbf{1}(s_t=i, a_t=j)} - \hat{\mu}_{C,n}(i, j)^2 & \text{if } \sum_{1 \leq t \leq n} \mathbf{1}(s_t=i, a_t=j) > 0 \\ s_C^0(i, j) & \text{if } \sum_{1 \leq t \leq n} \mathbf{1}(s_t=i, a_t=j) = 0, \end{cases}$$

where $u_C^0(i, j) \in \mathcal{R}$ and $s_C^0(i, j) \geq 0$ are some appropriate initial values for the mean and variance if no data is available.

We follow our paradigm to solve the empirical counterpart of the problem, namely to find a policy $\hat{\pi}_n^*$ that solves (9) by using $\hat{V}_n^\pi(s)$ and $\hat{L}_n^\pi(s)$ instead of $V^\pi(s)$ and $L^\pi(s)$, where $\hat{V}_n^\pi(s)$ and $\hat{L}_n^\pi(s)$ are the value functions and loss functions estimates using the empirical $\hat{\mu}_{R,n}$, $\hat{\mu}_{C,n}$, \hat{P}_n . We

focus on the estimation error of the optimal values (instead of the feasibility, which could also be important but not pursued here).

To understand the error, we first utilize an optimality characterization of constrained MDPs. In general, an optimal policy for (9) is a “split” policy (Feinberg and Rothblum 2012), namely, a policy that is deterministic except at one state where a randomization between two different actions is allowed. This characterization can be deduced from the associated LP using the occupancy measure (Altman 1999). We refer to the randomization probability as the mixing parameter and denote the optimal mixing parameter as α^* . In particular, whenever this state, say s_r , is visited, action $a_1^*(s_r)$ is chosen with probability α^* and action $a_2^*(s_r)$ is chosen with probability $1 - \alpha^*$.

THEOREM 3. *Suppose Assumptions 1 and 3 hold and there is a unique optimal policy. Moreover, assume that there is no deterministic policy π that satisfies $\sum_s \rho(s) L^\pi(s) = \eta$. Then*

$$\sqrt{n}(\hat{V}_n^* - V^*) \Rightarrow N(0, \Sigma_c) \text{ as } n \rightarrow \infty,$$

where one of the following two cases hold:

1. *The optimal policy is deterministic. In this case, $\Sigma_c = \Sigma_V$ where Σ_V is defined in Corollary 1.*
2. *The optimal policy is deterministic, except at one state, s_r , where a randomization between two actions, $a_1^*(s_r)$ and $a_2^*(s_r)$, occurs, with the mixing parameter α^* . In this case,*

$$\begin{aligned} \Sigma_c = & \left((I - \gamma P^{\pi^*})^{-1} [G^{\pi^*}, 0, H_V^{\pi^*}] - \frac{(I - \gamma P^{\pi^*})^{-1} h_V \rho^T (I - \gamma P^{\pi^*})^{-1} [0, G^{\pi^*}, H_L^{\pi^*}]}{\rho^T (I - \gamma P^{\pi^*})^{-1} h_L} \right) \\ & \times \Sigma_{R,C,P} \left((I - \gamma P^{\pi^*})^{-1} [G^{\pi^*}, 0, H_V^{\pi^*}] - \frac{(I - \gamma P^{\pi^*})^{-1} h_V \rho^T (I - \gamma P^{\pi^*})^{-1} [0, G^{\pi^*}, H_L^{\pi^*}]}{\rho^T (I - \gamma P^{\pi^*})^{-1} h_L} \right)^T, \end{aligned}$$

where

$$\Sigma_{R,C,P} = \begin{pmatrix} W^{-1} D_R & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & W^{-1} D_C & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & D_P \end{pmatrix},$$

D_C is an $N \times N$ diagonal matrix with $D_C((i-1)m_a + j, (i-1)m_a + j) = \sigma_C^2(i, j)$, W and D_R are defined in (6), and D_P is defined in (8). h_V and h_L are m_s -dimensional vectors. When $s = s_r$

$$\begin{aligned} h_V(s) &= (\mu_R(s, a_1^*(s)) - \mu_R(s, a_2^*(s))) + \sum_{j=1}^{m_s} \gamma V^{\pi^*}(j) (P(j|s, a_1^*(s)) - P(j|s, a_2^*(s))), \\ h_L(s) &= (\mu_C(s, a_1^*(s)) - \mu_C(s, a_2^*(s))) + \sum_{j=1}^{m_s} \gamma L^{\pi^*}(j) (P(j|s, a_1^*(s)) - P(j|s, a_2^*(s))). \end{aligned}$$

when $s \neq s_r$, $q_V(s) = q_L(s) = 0$.

$$G^{\pi^*} = \begin{pmatrix} \pi^*(\cdot|1)^T & & \\ & \ddots & \\ & & \pi^*(\cdot|m_s)^T \end{pmatrix}, H_V^{\pi^*} = \begin{pmatrix} q_V^{\pi^*}(1)^T & & \\ & \ddots & \\ & & q_V^{\pi^*}(m_s)^T \end{pmatrix}, H_L^{\pi^*} = \begin{pmatrix} q_L^{\pi^*}(1)^T & & \\ & \ddots & \\ & & q_L^{\pi^*}(m_s)^T \end{pmatrix},$$

where $q_V^{\pi^*}(i)$ and $q_L^{\pi^*}(i)$ are N -dimensional row vectors:

$$q_V^{\pi^*}(i) = \gamma \left[\pi^*(1|i)(V^{\pi^*})^T, \dots, \pi^*(j|i)(V^{\pi^*})^T, \dots, \pi^*(m_a|i)(V^{\pi^*})^T \right]^T,$$

$$q_L^{\pi^*}(i) = \gamma \left[\pi^*(1|i)(L^{\pi^*})^T, \dots, \pi^*(j|i)(L^{\pi^*})^T, \dots, \pi^*(m_a|i)(L^{\pi^*})^T \right]^T.$$

Case 1 in Theorem 3 corresponds to the case where the constraint in (9) is not binding. This effectively reduces the constrained MDP to the unconstrained scenario in Corollary 1, since a small perturbation of μ_R, μ_C, P does not affect feasibility. Case 2 is when the constraint is binding. In this case, α^* must be chosen such that the split policy ensures equality in the constraint. When μ_R, μ_C, P are perturbed, the estimated $\hat{\alpha}_n^*$ would adjust accordingly. In this case, \hat{V}_n^* incurs two sources of noises, one from the uncertainty in $\hat{\mu}_{R,n}, \hat{P}_n$ that appears also in unconstrained problems, and the other from the uncertainty in calibrating $\hat{\alpha}_n^*$ that is affected by $\hat{\mu}_{C,n}, \hat{P}_n$. This latter source of noise leads to the extra terms in the asymptotic variance expression.

We assume in Theorem 3 that there is no deterministic policy π satisfying $\sum_s \rho(s) L^\pi(s) = \eta$. This requirement is imposed to ensure that the corresponding LP formulation of the constrained MDP is non-degenerate. Similar to Assumption 2, when it does not hold, the limiting distribution can be non-Gaussian (see, e.g., Theorem 2).

4. Exploration Policy

In this section, we utilize our results in Section 3 to design good exploration policies. We focus on the setting in which an agent is assigned a period to collect data by running the state transition with an exploration policy. The goal is to obtain the best policy at the end of the period in a probabilistic sense, i.e., minimize the probability of selecting a suboptimal policy. We restrict our development to the standard case where each state has a unique optimal action as discussed in Section 3.1.

We first define, for $i \in \mathcal{S}$, $j \in \mathcal{A}$ and $j \neq a^*(i)$, the relative discrepancy as

$$h_{ij} = (Q(i, a^*(i)) - Q(i, j))^2 / \sigma_{\Delta Q}^2(i, a^*(i), j), \quad (10)$$

where $\sigma_{\Delta Q}^2(i, a^*(i), j)$ is defined in (7). Intuitively, h_{ij} captures the relative “difficulty” in obtaining the optimal policy given the estimation errors of Q-values. In particular, if the Q-values are far apart, or if the estimation variance is small, then h_{ij} is large which signifies an “easy” problem, and vice versa.

Our proposed strategy attempts to maximize the worst of h_{ij} ’s, i.e.,

$$\max_{w \in \mathcal{W}} \min_{i \in \mathcal{S}} \min_{j \in \mathcal{A}, j \neq a^*(i)} h_{ij}, \quad (11)$$

where w denotes the proportion of visits to each state-action pair, within some allocation set \mathcal{W} which will be defined in (12) below. Based on our interpretation of h_{ij} 's, criterion (11) aims to make the problem the “easiest” to differentiate the optimal policy. Alternatively, one can also interpret (11) from a large deviation point of view (Glynn and Juneja 2004, Dong and Zhu 2016). Suppose the Q-values for state i between two different actions $a^*(i)$ and j are very close. Then, one can show that the probability of a suboptimal selection of j has roughly an exponential in n decay rate controlled by h_{ij} . Obviously, there can be many more comparisons to consider, but the exponential form dictates that the smallest decay rate dominates the calculation, thus leading to the inner minimizations in (11). Criterion like (11) is motivated by the OCBA procedure in simulation optimization, which in general considers simple mean-value alternatives (Chen and Lee 2011). Here, we consider the estimation of Q-values. Thus, we refer to our procedure as Q-OCBA.

Implementing criterion (11) requires addressing two additional considerations. First, solving (11) needs the model primitives Q , P , and σ_R^2 , which appear in the expression of h_{ij} . These quantities are unknown a priori, but as we collect more data, they can be sequentially estimated. This leads to a multi-stage parameter update plus optimization scheme. Second, since data are collected through running a Markov chain, not all allocation w is *admissible*, i.e., realizable as the stationary distribution of some Markov chain. To resolve this issue, we next derive a convenient characterization for admissibility.

Accordingly to Assumption 3, we call a policy $\pi(\cdot|s)$ admissible if the Markov Chain with transition probability \tilde{P}^π is irreducible, and we denote w_π as its stationary distribution. Define

$$\mathcal{W} = \left\{ w > 0 : \sum_{1 \leq j \leq m_a} w((i-1)m_a + j) = \sum_{1 \leq k \leq m_s} \sum_{1 \leq l \leq m_a} w((k-1)m_a + l) P(i|k, l) \right. \\ \left. \forall 1 \leq i \leq m_s, \sum_{1 \leq i \leq m_s} \sum_{1 \leq j \leq m_a} w((i-1)m_a + j) = 1 \right\}. \quad (12)$$

The following lemma provides a characterization for the set of admissible policies.

LEMMA 1. *For any admission policy π , $w_\pi \in \mathcal{W}$. For any $w \in \mathcal{W}$, π_w with $\pi_w(a = j|s = i) = w((i-1)m_a + j) / (\sum_{k=1}^{m_a} w((i-1)m_a + k))$ is an admissible policy.*

Lemma 1 implies that optimizing over the set of admissible policies is equivalent to optimizing over the set of stationary distributions. The latter is much more tractable thanks to the linear structure of \mathcal{W} . In practice, we use $\mathcal{W}_\eta = \mathcal{W} \cap \{w(i) \geq \eta, i = 1, \dots, N\}$ for some small η to ensure closedness of the set (In our numerical experiments, we use $\eta = 10^{-6}$).

To elaborate the criterion (11), we note that it can be equivalently written as $\min_{w \in \mathcal{W}} \max_{i \in \mathcal{S}} \max_{j \in \mathcal{A}, j \neq a^*(i)} 1/h_{ij}$. Then, plugging the formula for $\sigma_{\Delta Q}^2(i, a^*(i), j)$ (i.e., (7)) in h_{ij} , we have

$$\min_{w \in \mathcal{W}} \max_{i \in \mathcal{S}} \max_{j \in \mathcal{A}, j \neq a^*(i)} \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} \frac{c_{ij}(s, a)}{w_{s,a}}, \quad (13)$$

where

$$c_{ij}(s, a) = \frac{(H_{ij}((s-1)m_a + a))^2 (\sigma_R^2(s, a) + (V^*)^T \Sigma_{P_{s,a}} V^*)}{(Q(i, a^*(i)) - Q(i, j))^2} \quad (14)$$

and

$$H_{ij} = (e_{(i-1)m_a + a^*(i)} - e_{(i-1)m_a + j})^T (I - \gamma \tilde{P}^{\pi^*})^{-1}.$$

Note that $c_{ij}(s, a)$ can be easily estimated with plug-in estimators. Next, note that the objective function in (13) is convex with N variables and the feasible region \mathcal{W}_η can be characterized by $m_s + 1 + N$ linear constraints. Thus, (13) can be solved by standard convex optimization solvers efficiently (in our numerical experiments in Section 6, solving it takes less than 0.01 seconds).

We are now ready to introduce our Q-OCBA implementation – Algorithm 1.

Input: Number of iterations K , cumulative number of data collected after stage k , B_k , $k = 1, \dots, K$, with $B_K = n$, i.e., n is the total sampling budget (and we define $B_0 = 0$), auxiliary parameters $\eta > 0$ (a small constant), $C_u > 0$ (a large constant), and $C_l > 0$ (a small constant), and initial exploration policy π_1 .

Initialization: $k = 1$;

while $k \leq K$ **do**

Run π_k for $B_k - B_{k-1}$ steps and collect the corresponding data; Calculate \hat{P}_{B_k} , $\hat{\mu}_{R, B_k}$, and $\hat{\sigma}_{R, B_k}^2$ based on the B_k data points collected;

Apply value-iteration using \hat{P}_{B_k} and $\hat{\mu}_{R, B_k}$ to obtain \hat{Q}_{B_k} ;

Solve the following adaption of (13) for the optimal w_k

$$\min_{w \in \mathcal{W}_\eta} \max_{i \in \mathcal{S}} \max_{j \in \mathcal{A}, j \neq a^*(i)} \sum_{(s, a) \in \mathcal{S} \times \mathcal{A}} \frac{\max(C_l, \min(C_u, \hat{c}_{ij, B_k}(s, a)))}{w_{s, a}}, \quad (15)$$

where $\hat{c}_{ij, B_k}(s, a)$ denotes the plug-in estimator of $c_{ij}(s, a)$ defined in (14);

Set $\pi_{k+1}(a = j | s = i) = w_k((i-1)m_a + j) / \sum_{l=1}^{m_a} w_k((i-1)m_a + l)$ and $k = k + 1$;

end

Algorithm 1: Q-OCBA sequential updating rule for exploration

Let π^e denote the optimal exploration policy (here e stands for exploration) and w^* denote the corresponding stationary distribution, i.e., w^* solves (13). Note that π^e is different from π^* . The former is used to collect data, while the latter is used to collect reward in the deployment stage and is to be learned offline using the data collected under π^e . We also denote $\hat{\pi}_n$ as the exploration policy derived from Algorithm 1 after collecting n data points. The following theorem establishes the asymptotic consistency of the Q-OCBA sequential updating rule.

THEOREM 4. Under Assumptions 1, 2, and 3, for Q-OCBA defined in Algorithm 1, suppose one of the following two cases holds as we send the sampling budget $n = B_K$ to infinity:

1. The number of iterations K is fixed. This implies that there exists at least one $\hat{k} \in \{1, \dots, K\}$, such that $B_{\hat{k}} - B_{\hat{k}-1} \rightarrow \infty$ as $n \rightarrow \infty$.
2. The number of iterations $K \rightarrow \infty$ as $n \rightarrow \infty$ and $B_k - B_{k-1} \geq N$ for all $k = 1, 2, \dots$.

Then,

$$\hat{Q}_n \rightarrow Q \text{ and } \hat{\pi}_n \rightarrow \pi^e \text{ a.s. as } n \rightarrow \infty.$$

Moreover, in Case 2, if there exists a constant $\gamma > 0$ such that $B_k - B_{k-1} \geq \gamma k$, $k = 1, 2, \dots$, then

$$\frac{1}{n} \sum_{t=1}^n \mathbb{1}(s_t = s, a_t = a) \rightarrow w^*(s, a) \text{ a.s. for any } (s, a) \in \mathcal{S} \times \mathcal{A} \text{ as } n \rightarrow \infty.$$

The first result in Theorem 4 indicates that as we collect more data in Q-OCBA, our estimated Q -values will get closer to the true Q -values and our estimated exploration policy will get closer to the optimal exploration policy. This is achieved by showing that each state-action pair is visited infinitely often under our policy. The second result shows that if the number of iterations is increasing with the sampling budget and the batch size $B_k - B_{k-1}$ is properly increasing with k , the overall proportion of samples at each state-action pair is also getting closer to the optimal proportion as we collect more data. Recall that w_k denotes the stationary distribution under policy π_k in Algorithm 1. From the first result, in Case 2, we have $w_k \rightarrow w^*$ as $k \rightarrow \infty$. Thus, for the second result, we only need to show that the convergence happens sufficiently fast. To characterize the convergence rate, we employ the large deviations theory for the empirical measure of the underlying Markov chain.

Theorem 4 lays out fairly general conditions for K and B_k 's, under which Algorithm 1 achieves consistency. However, with a finite budget in practice, we need to be more mindful about the choice of these parameters. In Section 6.2.2, we numerically investigate how to tune Q-OCBA. In general, we suggest setting K between 6 to 10. As for B_k 's, when n or K is small, we suggest evenly splitting the sampling budget among the K stages, i.e., $B_k - B_{k-1} = n/K$. When n and K are large, we can use gradually increasing batch sizes, e.g., $B_k - B_{k-1} = \gamma k$ where $\gamma = 2n/(K(K+1))$.

5. Approximations for Large State Space and/or Action Space

In many applications, the state space \mathcal{S} and/or the action space \mathcal{A} can be very large. In those settings, updating an $m_s \times m_a$ look-up table via $\mathcal{T}_{\mu_R, P}$ can be computationally infeasible, and approximation/parameterization is often employed. In this section, we study two approximation schemes: approximate value iteration and kernel representation.

5.1. Approximate Value Iteration

When m_s is large, approximate value iteration can be employed. It operates by applying a mapping M over $\mathcal{T}_{\mu_R, P}$. In many cases, $M = M_g \circ M_I^{S_0}$, where $M_I^{S_0}$ is a dimension-reducing “inherit” mapping $\mathbb{R}^{m_s m_a} \rightarrow \mathbb{R}^{m_{s_0} m_a}$ and M_g is the “generalization” mapping $\mathbb{R}^{m_{s_0} m_a} \rightarrow \mathbb{R}^{m_s m_a}$ that lifts back to the full dimension. By selecting a “representative” subset $S_0 \subset \mathcal{S}$ with cardinality $m_{s_0} \ll m_s$, $M_I^{S_0}$ is defined as $M_I^{S_0}(x) = [x(i, j)]_{i \in S_0, 1 \leq j \leq m_a}$ where $[x_i]_{i \in I}$ denotes the set of x_i ’s whose index $i \in I$. The idea is that we update the Q-values for $s \in S_0$ only, and use the mapping M_g to extrapolate the other Q-values. In this setup, we define Q^M as a fixed point of the operator $M \circ \mathcal{T}_{\mu_R, P}(\cdot)$, and $V^M(s) = \max_a Q^M(s, a)$. We also define $Q_{S_0}^M = M_I^{S_0} \circ \mathcal{T}_{\mu_R, P}(Q^M)$ as the dimension-reduced Q-values.

We next study the large-sample behavior of an approximate value iteration scheme. We first introduce a few assumptions. Assume \mathcal{S} is a complete vector space. Define the max norm as $\|a\|_\infty = \max_i |a_i|$. To guarantee the existence of Q^M , we make the following assumption on the generalization map M_g :

ASSUMPTION 4. M_g is a max-norm non-expansion mapping in \mathcal{S} , i.e.,

$$\|M_g(x) - M_g(y)\|_\infty < \|x - y\|_\infty \text{ for any } x, y \in \mathcal{S}.$$

Under Assumption 4, $M \circ \mathcal{T}_{\mu_R, P}(\cdot)$ is still a contraction mapping (Theorem 3.1 in Gordon 1995). Thus, Q^M is well defined. On the other hand, when this assumption is not satisfied, there exist MDPs for which $M \circ \mathcal{T}_{\mu_R, P}(\cdot)$ does not have a unique fixed point (Gordon 1995). Assumption 4 is generally satisfied by “local” approximation methods such as linear interpolation, k -nearest neighbors, and local weighted averaging. Moreover, Gordon (1995) defines the notion of an averager and shows that any M_g associated with an averager is max-norm non-expansion (Theorem 3.2 in Gordon 1995). In particular, M_g is an averager if the fitted values (i.e., $Q(s, a)$ for $s \in \mathcal{S}$) are the weighted average of some target values (i.e., $Q(s, a)$ for $s \in S_0$) and some predetermined constants.

We also need the following analogs of Assumptions 2 and 3 to Q^M and S_0 :

ASSUMPTION 5. For any state $s \in \mathcal{S}$, $\arg \max_{a \in \mathcal{A}} Q^M(s, a)$ is unique.

ASSUMPTION 6. For the Markov Chain with transition probability \tilde{P}^π , the set of states $\{(s, a) : s \in S_0, a \in \mathcal{A}\}$ is irreducible.

Let $N_0 = m_{s_0} m_a$ and $I_{S_0} = \{(i-1)m_a + j : i \in S_0, j \in \mathcal{A}\}$. With Assumption 6, we denote $\tilde{P}_{S_0}^M$ as a sub-matrix of \tilde{P}^π that only contains rows with indices in I_{S_0} . We also denote $S_0(i)$ as the i -th element (state) in S_0 . Define \hat{Q}_n^M as the empirical estimate of Q^M built on n observations.

THEOREM 5. Under Assumptions 4, 5, and 6, if M_g is continuously differentiable, then

$$\sqrt{n}(\hat{Q}_n^M - Q^M) \Rightarrow \mathcal{N}(0, \Sigma_{S_0}^M) \text{ as } n \rightarrow \infty,$$

where

$$\begin{aligned} \Sigma_{S_0}^M &= (I - \gamma \nabla M_g(Q_{S_0}^M) \tilde{P}_{S_0}^M)^{-1} \nabla M_g(Q_{S_0}^M) (W^{S_0})^{-1} [D_R^{S_0} + \gamma^2 D_Q^{S_0}] \\ &\quad \nabla M_g(Q_{S_0}^M)^T ((I - \gamma \nabla M_g(Q_{S_0}^M) \tilde{P}_{S_0}^M)^{-1})^T, \end{aligned}$$

∇M_g is the Jacobian of the mapping M_g , and W^{S_0} , $D_R^{S_0}$, $D_Q^{S_0}$ are $N_0 \times N_0$ diagonal matrices with

$$\begin{aligned} W^{S_0}((i-1)m_a + j, (i-1)m_a + j) &= w(S_0(i), j), \\ D_R^{S_0}((i-1)m_a + j, (i-1)m_a + j) &= \sigma_R^2(S_0(i), j), \\ D_Q^{S_0}((i-1)m_a + j, (i-1)m_a + j) &= (V^M)^T \Sigma_{P_{S_0(i),j}} V^M. \end{aligned}$$

An important class of approximation scheme that satisfies Assumption 4 is the nearest neighbors (Shah and Xie 2018). Given data $q \in \mathbb{R}^{m_{s_0} m_a}$, the nearest-neighbor averaging can be expressed as

$$M_g(q)(s, a) = \sum_{i=1}^{m_{s_0}} K(s, S_0(i)) q(S_0(i), a), \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A}, \quad (16)$$

where $K(\cdot, \cdot) \geq 0$ is the weighting kernel function satisfying $\sum_{i=1}^{m_{s_0}} K(s, S_0(i)) = 1$ for any $s \in \mathcal{S}$ and $K(x, y) = 0$ if $\|x - y\| > h$, where h is the width parameter. Then, for any $s \in \mathcal{S}$, $Q^M(s, a) = M_g(Q_{S_0}^M)$; and for any $s_0 \in \mathcal{S}_0$,

$$Q_{S_0}^M(s_0, a) = \mu_R(s_0, a) + \gamma \mathbb{E} \left[\max_{a' \in \mathcal{A}} M_g(Q_{S_0}^M)(s', a') | s_0, a \right].$$

We note from (16) that the nearest-neighbor averaging is an averager and is thus max-norm non-expansion.

5.2. Kernel Representation

When $N = m_s m_a$ is large, a common practice is to parameterize the high-dimensional Q-values or policy functions using a set features. In this subsection, we study one particular parameterization scheme – the kernel representation (Yang and Wang 2020), where we assume the transition probability P can be fully embedded in a kernel space:

ASSUMPTION 7. For each $(s, a) \in \mathcal{S} \times \mathcal{A}$ and $s' \in \mathcal{S}$, there exists a transition core matrix $M^* \in \mathbb{R}^{k_1} \times \mathbb{R}^{k_2}$, such that

$$P(s'|s, a) = \phi(s, a)^T M^* \psi(s')$$

where $\phi(s, a) \in \mathbb{R}^{k_1}$ and $\psi(s') \in \mathbb{R}^{k_2}$ are feature functions. We assume that $\phi(s, a)$ and $\psi(s')$ are known, while M^* has to be estimated from data.

k_1 and k_2 in Assumption 7 are typically much smaller than N . Thus, we can approximate P by a lower dimensional manifold. Similar assumptions are made for the distribution of the reward:

ASSUMPTION 8. *For each $(s, a) \in \mathcal{S} \times \mathcal{A}$, there exist parameters $\theta_\mu \in R^{d_1}$ and $\theta_\sigma \in R^{d_2}$, such that $\mathbb{E}[R(s, a)] = \phi_\mu(s, a)^T \theta_\mu$ and $\mathbb{E}[R^2(s, a)] = \phi_\sigma(s, a)^T \theta_\sigma$ where $\phi_\mu(s, a) \in \mathcal{R}^{d_1}$ and $\phi_\sigma(s, a) \in \mathcal{R}^{d_2}$ are feature functions. We assume that $\phi_\mu(s, a)$ and $\phi_\sigma(s, a)$ are known, while θ_μ and θ_σ have to be estimated from data.*

We next define a few more notations to facilitate our subsequent discussion. Define M_v^* as a vector by concatenating columns of M^* , i.e, $M_v^* = [(M^*)_1^T, \dots, (M^*)_{k_2}^T]^T$, where $(M^*)_i$ is the i -th column of M^* . Then, $P(s'|s, a)$ can be written as

$$\begin{aligned} P(s'|s, a) &= \left(\begin{pmatrix} \phi(s, a)^T & & \\ & \ddots & \\ & & \phi(s, a)^T \end{pmatrix} M_v^* \right)^T \psi(s') \\ &= \psi(s')^T \begin{pmatrix} \phi(s, a)^T & & \\ & \ddots & \\ & & \phi(s, a)^T \end{pmatrix} M_v^* =: F_{\phi, \psi}(s, a, s') M_v^*, \end{aligned}$$

where

$$F_{\phi, \psi}(s, a, s') = \psi(s')^T \begin{pmatrix} \phi(s, a)^T & & \\ & \ddots & \\ & & \phi(s, a)^T \end{pmatrix}.$$

When n data points are available, the least-squares estimation of M^* takes the form:

$$\hat{M}_n = A_n^{-1} \sum_{t=1}^n \phi(s_t, a_t) \psi(s'_t)^T K_\psi^{-1},$$

where $K_\psi = \sum_s \psi(s) \psi(s)^T$ and $A_n = I + \sum_{t=1}^n \phi(s_t, a_t) \phi(s_t, a_t)^T$. Define the vector $\hat{M}_{n,v}$ by concatenating columns of \hat{M}_n , i.e, $\hat{M}_{n,v} = [(\hat{M}_n)_1^T, \dots, (\hat{M}_n)_{k_2}^T]^T$, where $(\hat{M}_n)_i$ is the i -th column of \hat{M}_n . Similar to $P(s'|s, a)$, we can define

$$\hat{P}_n^K(s'|s, a) = F_{\phi, \psi}(s, a, s') \hat{M}_{n,v}. \quad (17)$$

Next, consider an extended Markov chain $\bar{X} = (s, a, s')$ with transition kernel

$$\bar{P}^\pi((s_2, a_2, s'_2) | (s_1, a_1, s'_1)) = \begin{cases} \pi(a_2 | s_2) P(s'_2 | s_2, a_2), & \text{if } s'_1 = s_2; \\ 0, & \text{otherwise.} \end{cases}$$

Let $\bar{\mathcal{S}} = \{x_1, x_2, \dots, x_{m_s^2 m_a}\}$ denote the state space of \bar{X} . The states are indexed in the same order as they appear in the rows of \bar{P}^π . For any state x in $\bar{\mathcal{S}}$, denote its original (s, a, s') representation as $(s(x), a(x), s'(x))$. Let $\bar{\omega}$ denote the stationary distribution of \bar{X} . Then, $\bar{\omega}_x = w_{s(x), a(x)} P(s'(x) | s(x), a(x))$. Lastly, define

$$G(x) = \phi(s(x), a(x)) \psi^T(s'(x)) K_\psi^{-1} - \phi(s(x), a(x)) \phi(s(x), a(x))^T M^*$$

and the vector $G_v(x)$ by concatenating columns of $G(x)$, i.e., $G_v(x) = [(G(x))_1^T, \dots, (G(x))_{k_2}^T]^T$, where $(G(x))_i$ is i -th column of $G(x)$.

LEMMA 2. Assume $\mathbb{E}_{\bar{\omega}}[G_v^2(\bar{X})] < \infty$. Then, for any $(s, a, s') \in \bar{\mathcal{S}}$,

$$\sqrt{n}(\hat{P}_n^K(s'|s, a) - P(s'|s, a)) \Rightarrow \mathcal{N}(\mathbf{0}, F_{\phi, \psi}(s, a, s') \Upsilon_E^{-1} \Sigma_G \Upsilon_E^{-1} (F_{\phi, \psi}(s, a, s'))^T),$$

where

$$\Upsilon_E = \begin{pmatrix} I + \sum_{s,a} w_{s,a} \phi(s, a) \phi(s, a)^T & & \\ & \ddots & \\ & & I + \sum_{s,a} w_{s,a} \phi(s, a) \phi(s, a)^T \end{pmatrix}$$

and

$$\Sigma_G = \text{Var}_{\bar{\omega}}(G_v(\bar{X}_0)) + 2 \sum_{t=1}^{\infty} \text{Cov}_{\bar{\omega}}(G_v(\bar{X}_0), G_v(\bar{X}_t)).$$

Here, $\text{Var}_{\bar{\omega}}(G_v(\bar{X}_0))$ is the covariance matrix of $G_v(\bar{X}_0)$ where \bar{X}_0 follows the stationary distribution $\bar{\omega}$ and $\text{Cov}_{\bar{\omega}}(G_v(\bar{X}_0), G_v(\bar{X}_t))$ is a matrix with $\{\text{Cov}_{\bar{\omega}}(G_v(\bar{X}_0), G_v(\bar{X}_t))\}_{s,r} = \text{Cov}_{\bar{\omega}}((G_v(\bar{X}_0))_s, (G_v(\bar{X}_t))_r)$. Specifically,

$$\text{Var}_{\bar{\omega}}(G_v(\bar{X}_0)) = \sum_{(s,a,s') \in \bar{\mathcal{S}}} w_{s,a} P(s'|s, a) G_v(s, a, s') G_v(s, a, s')^T$$

and

$$\begin{aligned} & \sum_{t=1}^{\infty} \text{Cov}_{\bar{\omega}}(G_v(\bar{X}_0), G_v(\bar{X}_t)) \\ &= \sum_{x_k, x_l \in \bar{\mathcal{S}}} w_{s(x_k), a(x_k)} P(s'(x_k)|s(x_k), a(x_k)) G_v(x_k) G_v(x_l)^T \{(I - \bar{P}^\pi)^{-1}\}_{k,l}. \end{aligned}$$

Note that the asymptotic variance $F_{\phi, \psi}(s, a, s') \Psi_E^{-1} \Sigma_G \Psi_E^{-1} (F_{\phi, \psi}(s, a, s'))^T$ can be estimated using plug-in estimators. The sampling policy will affect Σ_G through the transition matrix \bar{P}^π and the stationary distribution $\bar{\omega}$.

We next study the estimation of the mean reward μ_R . Let $\Phi_{\mu,n}$ denote the observed feature matrix for the reward, i.e., the t -th row of $\Phi_{\mu,n}$ is $\phi_\mu(s_t, a_t)^T$, $t = 1, \dots, n$. Similarly, let $\Phi_{\sigma,n}$ denote the observed feature matrix for the squared reward, i.e., the t -th row of $\Phi_{\sigma,n}$ is $\phi_\sigma(s_t, a_t)^T$, $t = 1, \dots, n$. We also denote $Y_{\mu,n}$ as the observed reward vector, i.e., the t -th row of $Y_{\mu,n}$ is $r_t(s_t, a_t)$, and $Y_{\sigma,n}$ as the observed squared reward vector whose t -th entry is $r_t(s_t, a_t)^2$, $t = 1, \dots, n$. Then, the least-squares estimators for θ_μ and θ_σ take the form

$$\hat{\theta}_{\mu,n} = (\Phi_{\mu,n}^T \Phi_{\mu,n})^{-1} \Phi_{\mu,n}^T Y_{\mu,n} \quad \text{and} \quad \hat{\theta}_{\sigma,n} = (\Phi_{\sigma,n}^T \Phi_{\sigma,n})^{-1} \Phi_{\sigma,n}^T Y_{\sigma,n} \quad \text{respectively.}$$

In this case, $\mu_R(s, a)$ and $\sigma_R^2(s, a)$ can be estimated via

$$\hat{\mu}_{R,n}^K(s, a) = \phi_\mu(s, a)^T \hat{\theta}_{\mu,n} \tag{18}$$

and

$$(\hat{\sigma}_{R,n}^K(s, a))^2 = \phi_\sigma(s, a)^T \hat{\theta}_{\sigma,n} - (\phi_\mu(s, a)^T \hat{\theta}_{\mu,n})^2$$

respectively.

Consider a Markov chain $\tilde{X} = (s, a)$ with transition kernel

$$\tilde{P}^\pi(s_2, a_2 | s_1, a_1) = P(s_2 | s_1, a_1) \pi(a_2 | s_2).$$

We denote $\tilde{\mathcal{S}} = \{x_1, \dots, x_{m_s m_a}\}$ as the state space of \tilde{X} , where the states are indexed in the same order as they appear in the rows of \tilde{P}^π . For any state $\tilde{x} \in \tilde{\mathcal{S}}$, denote its (s, a) representation as $(s(\tilde{x}), a(\tilde{x}))$. Recall that w is the stationary distribution of \tilde{X} . Let

$$H(\tilde{X}) = \left(r(s(\tilde{X}), a(\tilde{X})) - \mu_R(s(\tilde{X}), a(\tilde{X})) \right) \phi_\mu(s(\tilde{X}), a(\tilde{X})).$$

LEMMA 3. Assume $\mathbb{E}_w[(H(\tilde{X}))^2] < \infty$. Then, for any $(s, a) \in \tilde{\mathcal{S}}$,

$$\sqrt{n}(\hat{\mu}_{R,n}^K(s, a) - \mu_R(s, a)) \Rightarrow \mathcal{N}(\mathbf{0}, \phi_\mu(s, a)^T \Upsilon_\mu^{-1} \Sigma_H \Upsilon_\mu^{-1} \phi_\mu(s, a)),$$

where

$$\Upsilon_\mu = \mathbb{E}_w[\phi_\mu(\tilde{X}_0)(\phi_\mu(\tilde{X}_0))^T]$$

and

$$\Sigma_H = \text{Var}_w(H(\tilde{X}_0)) + 2 \sum_{t=1}^{\infty} \text{Cov}_w(H(\tilde{X}_0), H(\tilde{X}_t)).$$

Specifically,

$$\begin{aligned} \text{Var}_w(H(\tilde{X}_0)) &= \sum_{(s,a) \in \tilde{\mathcal{S}}} w_{s,a} H(s, a) H(s, a)^T \\ &= \sum_{(s,a) \in \tilde{\mathcal{S}}} w_{s,a} \sigma_R^2(s, a) \phi_\mu(s, a) \phi_\mu(s, a)^T \\ &= \sum_{(s,a) \in \tilde{\mathcal{S}}} w_{s,a} (\phi_\sigma(s, a)^T \theta_\sigma) \phi_\mu(s, a) \phi_\mu(s, a)^T \end{aligned}$$

and

$$\sum_{t=1}^{\infty} \text{Cov}_w(H(\tilde{X}_0), H(\tilde{X}_t)) = \sum_{x_k, x_l \in \tilde{\mathcal{S}}} w_{s(x_k), a(x_k)} H(s(x_k), a(x_k)) H(s(x_l), a(x_l)) \left\{ (I - \tilde{P}^\pi)^{-1} \right\}_{k,l}.$$

Note that the asymptotic variance $\phi_\mu(s, a)^T \Phi_\mu^{-1} \Sigma_H \Phi_\mu^{-1} \phi_\mu(s, a)$ can be estimated using plug-in estimators. Here the sampling policy will affect Σ_H through the transition matrix \tilde{P}^π and the stationary distribution w .

Based on the above analysis, the kernel-based Q-value estimates can be expressed as the empirical fixed point of the Bellman operator

$$\hat{Q}_n^K = \mathcal{T}_{\hat{P}_n^K, \hat{\mu}_{R,n}^K}(\hat{Q}_n^K),$$

where \hat{P}_n^K is defined in (17) and $\hat{\mu}_{R,n}^K$ is defined in (18). Combining Lemmas 2 and 3, we have the following central limit theorem for \hat{Q}_n^K .

THEOREM 6. *Under assumptions 7 and 8, assume $\mathbb{E}_{\tilde{\omega}}[G_v^2(X)] < \infty$ and $\mathbb{E}_w[H^2(\tilde{X})] < \infty$. Then,*

$$\sqrt{n}(\hat{Q}_n^K - Q) \Rightarrow \mathcal{N}(\mathbf{0}, \Sigma^K) \quad \text{as } n \rightarrow \infty,$$

where

$$\Sigma^K = (I - \gamma \tilde{P}^{\pi^*})^{-1} (\phi_\mu^T \Upsilon_\mu^{-1} \Sigma_H \Upsilon_\mu^{-1} \phi_\mu + \gamma^2 V_D F_{\phi, \psi} \Upsilon_E^{-1} \Sigma_G \Upsilon_E^{-1} F_{\phi, \psi}^T V_D) ((I - \gamma \tilde{P}^{\pi^*})^{-1})^T, \quad (19)$$

ϕ_μ is an $(m_s m_a) \times d_1$ matrix whose i -th row is $\phi_\mu(s_i, a_i)$, $F_{\phi, \psi}$ is an $(m_s^2 m_a) \times k_1$ matrix whose i -th row is $F_{\phi, \psi}(s_i, a_i, s'_i)$,

$$V_D = \begin{pmatrix} (V^*)^T & & & & \\ & \ddots & & & \\ & & (V^*)^T & & \\ & & & \ddots & \\ & & & & (V^*)^T \end{pmatrix},$$

Υ_E and Σ_G are defined in Lemma 2, Υ_μ and Σ_H are defined in Lemma 3.

5.3. Q-OCBA Variants

In this section, we develop the Q-OCBA variants for the nearest-neighbors approximation and the kernel representation.

We first consider the nearest-neighbors approximate value iteration. Based on Theorem 5, let

$$h_{ij}^M = (Q^M(i, a^*(i)) - Q^M(i, j))^2 / \sigma_{\Delta Q^M}^2(i, a^*(i), j),$$

where

$$\sigma_{\Delta Q^M}^2(i, a^*(i), j) = (e_{(i-1)m_a + a^*(i)} - e_{(i-1)m_a + j})^T \Sigma_{S_0}^M (e_{(i-1)m_a + a^*(i)} - e_{(i-1)m_a + j}).$$

Then, we can modify Q-OCBA to Q^M -OCBA, where our goal is to find a sampling policy that solves

$$\max_{w \in \mathcal{W}_\eta} \min_{i \in \mathcal{S}} \min_{j \in \mathcal{A}, j \neq a^*(i)} h_{ij}^M. \quad (20)$$

Next, we consider the kernel representation. Based on Theorem 6, let

$$h_{ij}^K = (Q^K(i, a^*(i)) - Q^K(i, j))^2 / \sigma_{\Delta Q^K}^2(i, a^*(i), j),$$

where

$$\sigma_{\Delta Q^K}^2(i, a^*(i), j) = (e_{(i-1)m_a + a^*(i)} - e_{(i-1)m_a + j})^T \Sigma^K (e_{(i-1)m_a + a^*(i)} - e_{(i-1)m_a + j})$$

Then, we can modify Q-OCBA to Q^K -OCBA, where our goal is to find an exploration policy that solves

$$\max_{w \in \mathcal{W}_\eta} \min_{i \in \mathcal{S}} \min_{j \in \mathcal{A}, j \neq a^*(i)} h_{ij}^K. \quad (21)$$

In both cases discussed above, we can modify sequential updating rule in Algorithm 1 to develop the corresponding algorithmic implementations. In particular, for Q^M -OCBA, we replace (15) in Algorithm 1 with (20). For Q^K -OCBA, we replace (15) with (21).


6. Numerical Experiments

In this section, we conduct numerical experiments to support our large-sample results in Section 3, demonstrate the performance of Q-OCBA against some benchmark exploration policies, and analyze how to tune the parameters when implementing Q-OCBA. We use the RiverSwim problem with m_s states and two actions at each state: swimming left (0) and swimming right (1). Figure 1 provides a pictorial illustration of the problem. The triplet above each arc in Figure 1 represents i) the action, ii) the transition probability to the next state given the current state and action, and iii) the reward under the current state and action. For example, in $(1, 0.3, 0)$, 1 is the action of swimming right, 0.3 is the probability of moving to the next state as the directed arc indicates, and 0 is the reward for the current state-action pair. Note that, in this problem, rewards are only given at the left and right boundary states, where the left-boundary reward, $r_L < 10$, will be varied in our experiments. We set the discount factor $\gamma = 0.95$.

Note that swimming to the right (against the current of the river) will more often than not leave the agent in the same state, but will sometimes move the agent to the right, and with a much smaller probability move the agent to the left. Swimming to the left (with the current) always succeeds in moving the agent to the left, until the leftmost state is reached at which point swimming to the left yields a small reward (i.e., $r_L < 10$). The agent receives a much larger reward for reaching the rightmost state. This MDP requires a sequence of appropriate actions in order to explore effectively to the right, and is challenging to learn when nothing is known at the beginning. Thus, it is a classic example to study exploration policies (see, e.g., Strehl and Littman 2008, Osband et al. 2013).

6.1. Statistical Quality of Interval Estimations

We first demonstrate the validity of our large-sample asymptotic results. All coverage rates reported in our numerical experiments are estimated using 10^3 independent repetitions of the procedure. The estimation errors (half width of the 95% confidence intervals (CIs)) are around 0.01. Our target confidence level (coverage rate) is 95%.

Figure 1 RiverSwim Problem

We start with the RiverSwim problem with a small state space, i.e., $m_s = 6$ as depicted in Figure 1, and $r_L = 1$. Table 1 reports the coverage rates of the asymptotically valid 95% CI for Q , V^* , and χ^* , constructed based on Theorem 1 and Corollary 1. The data collection policy we employ is a random exploration policy under which we swim right with probability 0.8 at each state, i.e., $\pi(1|s) = 0.8$. We observe that the coverage rate approaches the nominal 95% as the number of observations n increases. Specifically, when n is 10^2 , the coverage rates are all around 30%, when n is 5×10^2 , they are around 90%, and when $n \geq 10^3$, they are around 95%. These suggest a sample size of 10^3 is enough to elicit our asymptotic results in this problem.

Table 1 Coverage for $Q(s, a)$, V^* and χ^* values using exact tabular update

n	10^2	5×10^2	10^3	10^4
$Q(1, 0)$	0.2	0.91	0.95	0.95
$Q(3, 1)$	0.31	0.89	0.95	0.95
$Q(6, 0)$	0.30	0.87	0.94	0.95
$V^*(2)$	0.29	0.89	0.94	0.96
$V^*(4)$	0.30	0.90	0.94	0.95
$V^*(5)$	0.29	0.88	0.94	0.95
χ^*	0.30	0.88	0.94	0.95

We next consider the RiverSwim problem with a large state space, i.e., $m_s = 31$, and $r_L = 1$. We use the linear interpolation approximation with $\mathcal{S}_0 = \{1, 4, \dots, 28, 31\}$. Table 2 reports the coverage rates of the asymptotically valid 95% CI for Q , V^* , and χ^* , constructed based on Theorem 5, with different sample size and under different exploration policies. In particular, we vary the values of $\pi(1|s)$ from 0.8 to 0.9.

Compared to the exact update with a small state space, the coverage-rate convergence for the approximate update with a larger state space appears slower. Specifically, comparing Tables 1 and 2 that use the same random exploration with $\pi(1|s) = 0.8$, we note that while the nominal coverage

is obtained when $n = 10^4$ in the exact update with a smaller state space for all studied quantities, this sample size is not enough for approximate update with $m_{s_0} = 11$, where it appears that we need n to be of order 10^7 to obtain the nominal coverage.

We also note that, when the coverage is very far from the nominal level, discrepancies can show up among the estimates of Q , V^* , and χ^* . For example, when $\pi(1|s) = 0.8$ and $n = 10^4$, the coverage rates of Q and V^* are around 47% – 49% but that of χ^* is as low as 2%, and when $\pi(1|s) = 0.9$ and $n = 10^6$, the coverage rates of Q and V^* are around 33% – 35% but that of χ^* is only 4% . However, when the coverage rate is close to 95%, all these quantities appear to attain this accuracy simultaneously in all the cases considered. Nonetheless, the convergence behaviors predicted by Theorem 5 are observed to hold.

Furthermore, Table 2 shows that the rates of convergence to the nominal coverage are quite different for different values of $\pi(1|s)$. The convergence rate when $\pi(1|s) = 0.85$ seems to be the fastest, with the coverage rate close to 95% already when $n = 10^5$. On the other hand, when $\pi(1|s) = 0.8$, the coverage rate is close to 95% only when $n = 10^6$, and when $\pi(1|s) = 0.9$, even $n = 10^7$ is not large enough to have the target coverage rate. These caution that estimation quality can be quite sensitive to the exploration policy (the quality of data collected). We investigate the efficiency of different exploration policies further in the next subsection.

Table 2 Linear interpolation in approximate value iteration

	n	10^4	10^5	10^6	10^7
$\pi(1 s) = 0.8$	Average Q coverage	0.47	0.73	0.93	0.95
	Average V^* coverage	0.49	0.74	0.94	0.95
	χ^* coverage	0.02	0.37	0.93	0.94
$\pi(1 s) = 0.85$	Average Q coverage	0.45	0.92	0.94	0.95
	Average V^* coverage	0.48	0.93	0.94	0.95
	χ^* coverage	0.12	0.91	0.95	0.94
$\pi(1 s) = 0.9$	Average Q coverage	0.30	0.37	0.33	0.69
	Average V^* coverage	0.34	0.39	0.35	0.70
	χ^* coverage	0.01	0.01	0.04	0.55

6.2. Efficiency of Exploration Policies

In this section, we investigate the efficiency of Q-OCBA defined in Algorithm 1. We first compare Q-OCBA to other benchmark policies for the task of learning the optimal policy. In particular, given a sampling budget, we compare which policy would be able to collect a more “informative” data set to learn the optimal policy. Note that Q-OCBA and other benchmark policies are used to collect data, and the optimal policy is trained offline using the data collected under these policies (i.e., by solving the empirical fixed point of the Bellman operator). We then discuss how to tune Q-OCBA, e.g., choosing the number of iterations K .

6.2.1. Comparison with other benchmark policies We compare Q-OCBA to four benchmark policies: i) ϵ -greedy with $\epsilon = 0.2$, ii) random exploration (RE) with $\pi(1|s) = 0.6$ and $\pi(1|s) = 0.8$, iii) UCRL2 (a variant of UCRL) with $\delta = 0.05$ (Jaksch et al. 2010), and iv) PSRL (Osband et al. 2013). We do sequential updating for Q-OCBA, PSRL, and ϵ -greedy.

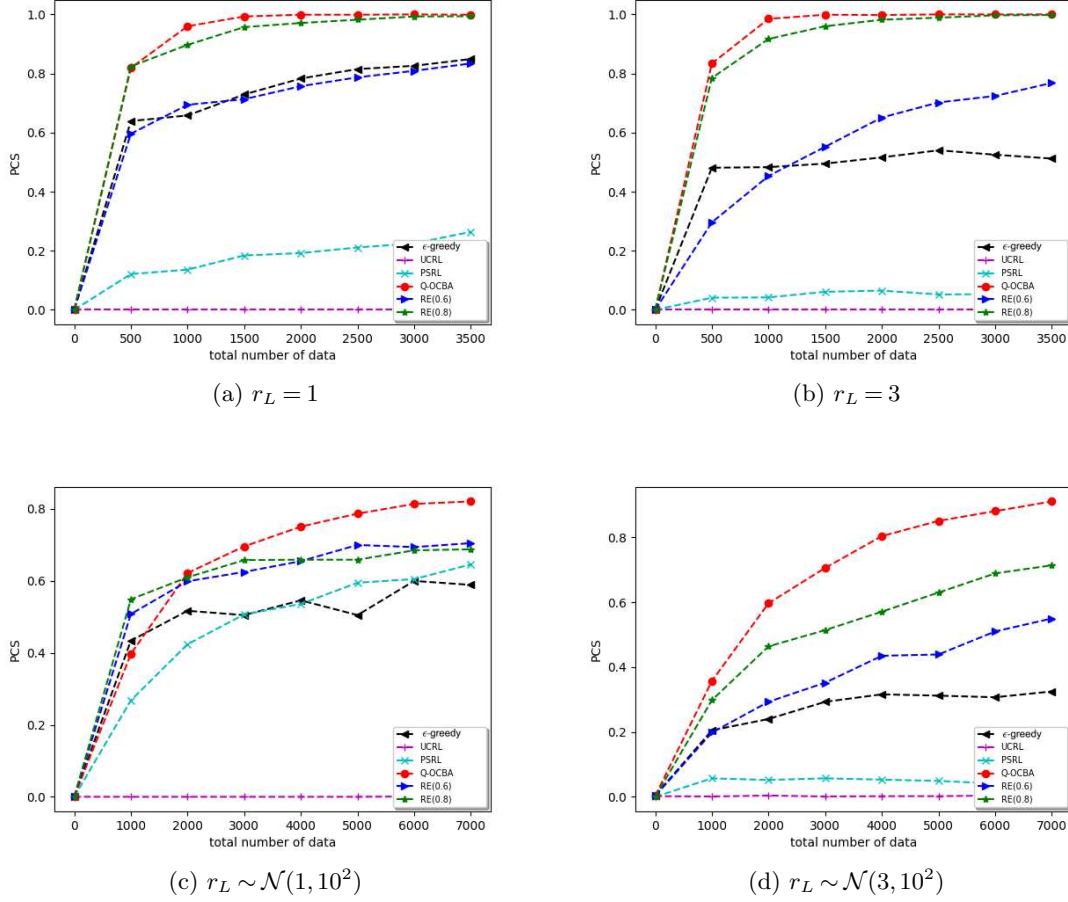
We first set $m_s = 6$, and vary the values of r_L . r_L is either set as a constant or a Gaussian random variable. $\hat{P}_{B_0}(s'|s, a)$ is initialized to be $1/m_s$, and $\hat{\mu}_{R, B_0}(s, a)$ and $\hat{\sigma}_{R, B_0}^2(s, a)$ are initialized to be 1 for all (s, a) pairs. Figure 2 compares the probability of obtaining the optimal policy (probability of correct selection) using data collected under different policies. The probability of correct selection is estimated using 10^3 independent repetitions of the procedures and the estimation errors (half width of the 95% confidence intervals) are around 0.01. We observe that Q-OCBA outperforms the other policies in almost all cases. For example, when $r_L = 1$ and $n = 2000$, the probability of correct selection using Q-OCBA is 4 times higher than that using PSRL, and 25% higher than that using ϵ -greedy. By carefully tuning the RE parameter $\pi(1|s)$, we find that $\pi(1|s) = 0.8$ achieves comparable performance as Q-OCBA. However, when $\pi(1|s)$ is not optimally chosen, e.g., $\pi(1|s) = 0.6$, the performance of RE is significantly worse than Q-OCBA. For example, when $r_L = 1$ and $n = 2 \times 10^3$, RE(0.8) leads to a 95% probability of correct selection, while RE(0.6) only gives a 72% probability. Note that there is no systematic way to tune the parameter of RE other than trial and error. In this sense Q-OCBA is more robust than RE.

We also observe that ϵ -greedy and PSRL perform much worse when $r_L = 3$ compared to when $r_L = 1$. One possible explanation is that for $r_L = 1$, the (s, a) pairs that need to be explored more also tend to have larger Q-values. However, as r_L increases, the corresponding changes in the Q-values would change the exploration “preference” of ϵ -greedy and PSRL. On the other hand, as the underlying stochasticity of the system does not change with r_L , the states that need more exploration remain unchanged. Thus, there is a misalignment between the Q-values and the (s, a) pairs that need more exploration. In contrast, the performance of Q-OCBA is very stable against different values of r_L . The superiority of Q-OCBA in these experiments comes as no surprise to us. The benchmark policies like UCRL and PSRL are designed to minimize regret which involves balancing the exploration-exploitation trade-off. On the other hand, Q-OCBA focuses on efficient exploration only, i.e., our goal is to minimize the probability of incorrect policy selection. This is achieved by carefully utilizing the variance information gathered from the previous stages, and is made possible by our derived asymptotic variance formulas.

Lastly, we observe that randomness in reward contaminates the estimation. For example, when $r_L = 1$ (Figure 2(a)), Q-OCBA is able to achieve a higher than 90% probability of correct selection when the sample size is 10^3 . However, when $r_L \sim \mathcal{N}(1, 10^2)$ (Figure 2(c)), even with a sample size of 7×10^3 , the probability of correct selection under Q-OCBA is only around 80%. Nevertheless,

even with a high variability in reward, Q-OCBA still outperforms other benchmark policies. For example, when $r_L \sim \mathcal{N}(1, 10^2)$ and $n = 7 \times 10^3$, Q-OCBA achieves a 14% higher probability of correct selection than the second best, which in this case is RE(0.6).

Figure 2 Probability of correct selection (PCS) with different sampling budgets and different values of r_L .



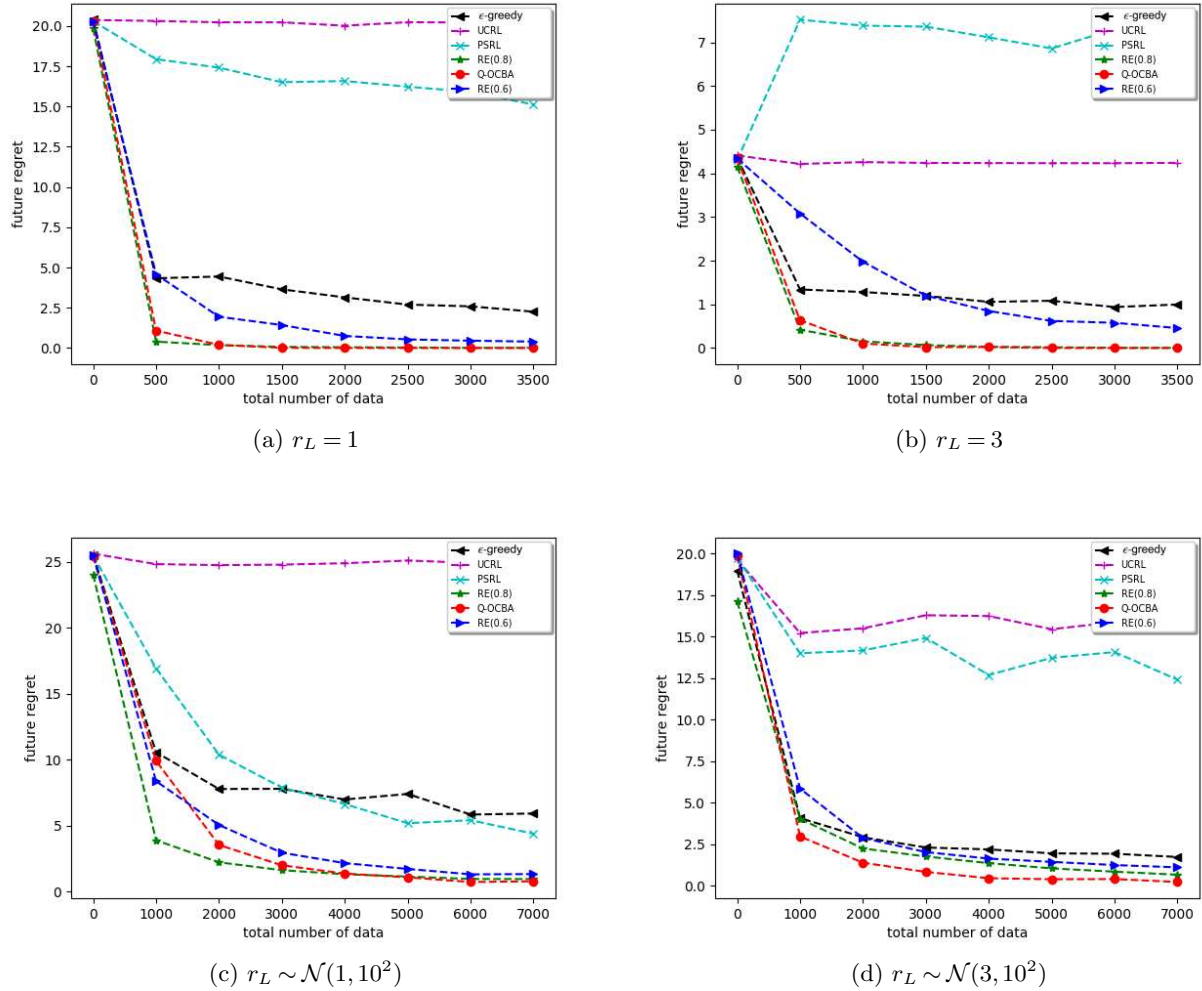
In practice, when we have a two-stage implementation where we collect experience in the first stage and deploy the trained policy to collect reward in the second stage, we would care more about the regret from the trained policy than the probability of obtaining the optimal policy itself. Let $\hat{\pi}_n^*$ denote the estimated optimal policy trained offline using a sample of size n collected via a proper exploration policy, e.g., Q-OCBA. In this case, we define the future regret as

$$R = \rho^T (V^* - \mathbb{E}[V^{\hat{\pi}_n^*}]).$$

In particular, R measures the average loss from deploying the estimated optimal policy in the second stage.

In Figure 3, we compare the future regret for $\hat{\pi}_n^*$'s trained based on data collected under different first-stage exploration policies using the examples as those in Figure 2. We observe that Q-OCBA still outperforms most benchmark policies. For example, when $r_L = 1$ and $n = 2 \times 10^3$, the regret of the policy trained using data collected under Q-OCBA is only 1/160 of that using PSRL, and 1/30 of that using ϵ -greedy. Even though Q-OCBA is not designed to minimize R , it is intuitive that the probability of selecting the optimal policy is positively corrected with $V^{\hat{\pi}_n^*}$, which matches our empirical observations.

Figure 3 Future regret of Q-OCBA vs benchmarks as n increases

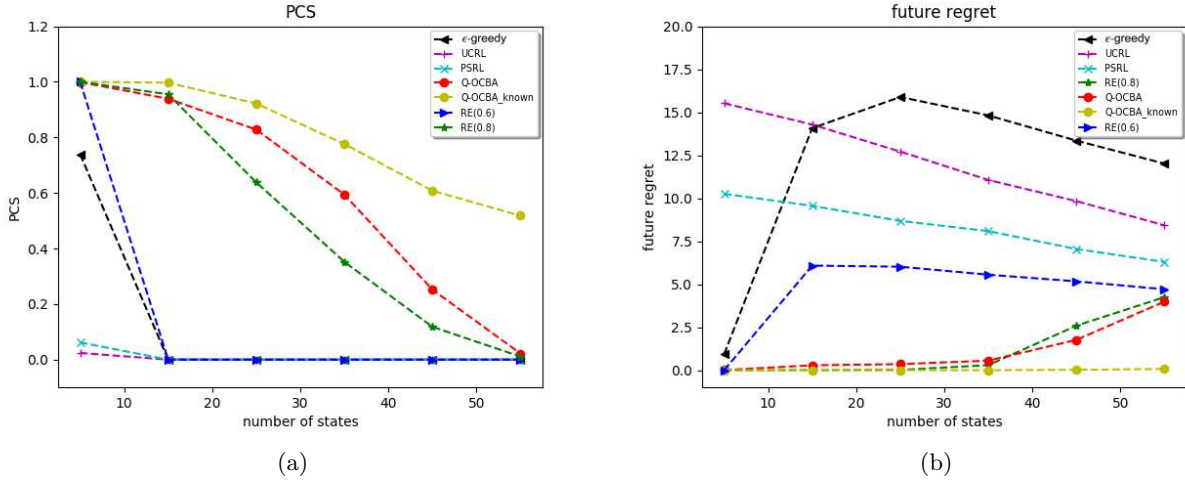


Next, we vary the value of m_s to test the performance of the algorithm for large state spaces. In particular, we fix $n = 10^5$ and assume $r_L \sim \mathcal{N}(2, 1)$. For the sequential update (for Q-OCBA, PSRL, and ϵ -greedy), we set $K = 10$. For the initialization, we sample \hat{P}_{B_0} , $\hat{\mu}_{R, B_0}$ and $\hat{\sigma}_{R, B_0}^2$ from Uniform[0, 1] (\hat{P}_{B_0} needs further normalization to be a valid transition matrix).

Figure 4 compares the performance (probability of correct selection and future regret) under different data collecting policies for different values of m_s . We also calculate the performance of Q-OCBA with known system parameters (i.e., h_{ij} 's defined in (10) can be evaluated exactly). This oracle policy is referred to as Q-OCBA_{known}.

We observe from Figure 4 that with a fixed sampling budget, the performance of all policies deteriorates as m_s increases. This is because the problem is more difficult to learn as the state space grows larger. Q-OCBA_{known} performs the best among all policies tested. Even when $m_s = 55$, it still achieves a 52% probability of correct selection. Q-OCBA consistently outperforms the other benchmark policies. For example, when $m_s = 45$, Q-OCBA (with estimated parameters) achieves a 25% probability of correct selection, while the best benchmark policy, which in this case is RE(0.8), only has an 11% probability of correct selection. We also observe that the performance of UCRL, PSRL, and ϵ -greedy deteriorate very quickly as m_s increases. When $m_s \geq 15$, their probability of correct selection is almost zero. This suggests that with a limited sampling budget, when the state space is large, it is very important to optimize the exploration policy to collect informative data for policy training. In this regard, Q-OCBA shows the potential to outperform in this example.

Figure 4 Probability of correct selection (PCS) and future regret with different m_s



6.2.2. Tuning Q-OCBA Q-OCBA requires very minimum tuning. Given a sampling budget, we need to specify the number of iterations K and the budget for each iteration. In addition, we need to provide some reasonable initializations for $\hat{P}(\cdot|s, a)$ and $\hat{\mu}_R(s, a)$.

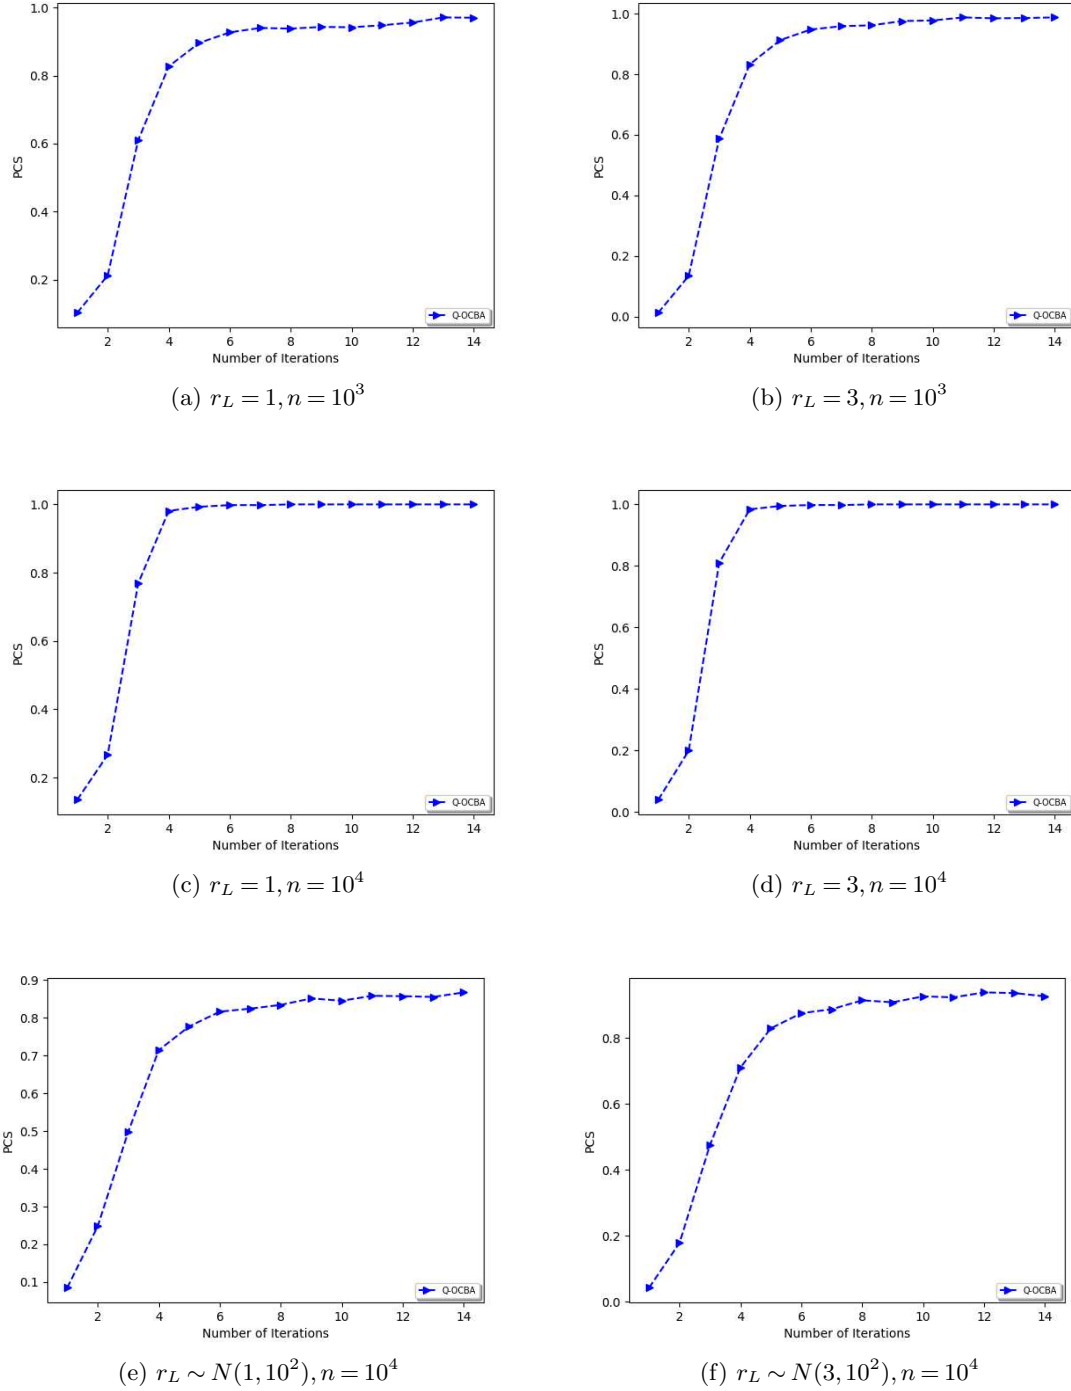
We note from our experiments that the performance of Q-OCBA is quite robust to different choices of initializations. This robustness stems from the updating criterion (15). In particular, (15) leads to stationary distributions that are reasonably far from zero for all the state-action

pairs. This stationary measure induces efficient exploration even when the initializations deviate substantially from the true values. In Table 3, we report the probability of correct selection under different initialization for RiverSwim with $m_s = 6$ for different values of r_L . We set the initialization for $\hat{P}(\cdot|s, a)$ as uniform, i.e., for $1 \leq s \leq 4$, $p^0(s'|s, a) = 1/3$ for $s' \in \{s, s+1, s-1\}$, for $s = 0$, $p^0(0|0, a) = 1/2$ and $\hat{P}(1|0, a) = 1/2$, and for $s = 5$, $p^0(5|5, a) = 1/2$ and $p^0(4|5, a) = 1/2$. We vary the initializations for $\hat{\mu}_R(s, a)$, for all $(s, a) \in \mathcal{S} \times \mathcal{A}$, from 0 to 100. The sampling budget is $n = 10^3$, the number of iterations is $K = 10$, and the sampling budget is equally distributed among the 10 iterations. We observe from Table 3 that across all problem instances and initializations, Q-OCBA achieves a very high probability of correct selection. Specifically, even in the extreme case where we set $u_R^0(s, a) = 100$, when the true $r_L = 1$, the probability of correct selection is 0.97. We also note that the probability of correct selection is not monotone in the value of r_L , with $r_L = 2$ leading to the lowest probability of correct selection.

Table 3 Probability of correct selection using Q-OCBA under different initializations, $n = 10^3$ and $K = 10$

r_L	$u_R^0(s, a) = 0$	$u_R^0(s, a) = 1$	$u_R^0 = 10$	$u_R^0(s, a) = 100$
1	0.98	0.95	0.96	0.97
2	0.91	0.91	0.87	0.84
3	0.98	0.97	0.99	0.98

We next study the choice of K – the number of iterations, which determines how often the exploration policy is updated during data collection. Figure 5 compares the probability of correct selection for different values of K based on RiverSwim with $m_s = 6$. To focus on the effect of K , the sampling budget is equally distributed among the K iterations. We observe that as K increases, the probability of correct selection under Q-OCBA increases. Specifically, with 10^4 samples, when $K = 2$, the probability of correct selection is only around 18% – 25%; when $K = 10$, the probability of correct selection increases to above 85% in all cases, with the probability almost equal to 1 for $r_L = 1$ or 3. However, there is a diminishing benefit of increasing K . For example, when increasing K from 2 to 4, the probability of correct selection increases by more than 0.5. In contrast, when increasing K from 6 to 8, the probability only increases by only 0.03. We also note that when the sampling budget n is small or when there is a lot of uncertainty in r_L , we benefit more from having a larger value of K . In particular, when comparing Figure 5(a) to 5(c), $K = 4$ leads to a probability of correct selection of 0.82 when $n = 10^3$ versus 0.98 when $n = 10^4$. When $n = 10^3$, to achieve a higher than 0.9 probability of correct selection, we need $K \geq 6$. As a general rule of thumb, we suggest setting K between 6 and 10, as updating the policy at each iteration, i.e., solving (15), incurs some computational cost.

Figure 5 Number of stages against probability of correct selection (PCS)

7. Concluding Remarks

In this paper, we investigate the large-sample behaviors in the estimation of Q-values and related quantities in RL. We establish the corresponding central limit theorem with an explicit character-

ization of the asymptotic variance. Utilizing the asymptotic variance obtained, we further develop a pure exploration policy, Q-OCBA, that maximizes the worst-case relative discrepancy among the estimated Q-values relative to the optimal action. In our numerical experiments, we compare Q-OCBA with various benchmark policies in terms of the probability of correctly selecting the optimal policy and the future regret. We observe that Q-OCBA outperforms benchmark methods according to these pure exploration criteria. Our work is among the few that investigate statistical uncertainty quantification for RL and develop policies for pure exploration purposes.

There are several limitations of our current developments that would lead to interesting future research directions. First, compared to existing finite-sample bounds for RL estimates, our asymptotics result in tighter bounds and explicit characterization of the constants (i.e., the asymptotic variances). However, uncertainty quantification based on asymptotics also has its limitations. With a finite sample size, we may not be able to achieve the desired confidence level and it often requires trial and error to decide how many samples are required to achieve close-to-asymptotic performance.

Second, our exploration policy requires estimating the transition matrix and the reward distribution for each state-action pair. This may need a large amount of data and can be computationally expensive when the state space and/or action space is large. We consider two approximation schemes: nearest neighbors and kernel representation. However, in many large-scale RL problems, it is highly non-trivial to find a good approximation or lower-dimensional representation of the underlying MDP. How to find good approximations for large-scale problems and develop efficient exploration policies accordingly require substantial future developments. We believe our framework provides a promising starting point.

Lastly, the relative discrepancy maximization criteria in Q-OCBA has its own limitations. The mean-variance trade-off in OCBA could lead to misleading results in certain low-confidence scenarios (see, for example, Shin et al. (2016), Peng et al. (2017)). Other criteria for information gaining such as knowledge gradient (Ryzhov et al. 2012) or posterior probability of correct selection (Peng et al. 2018b) may be employed to achieve better performance. However, how to adapt these criteria to the RL setting requires further investigation. In addition, there is in general a lack of finite-sample performance bounds for OCBA algorithms. Developing meaningful performance guarantees when dealing with a relatively small sample size would also be an interesting research direction. This can provide further guidance on how to fine-tune the algorithm parameters.

Appendix A: Proofs of the results in Section 3.1

In all the subsequent proofs, we shall treat P and \hat{P}_n as Nm_s -dimensional vectors following the index rule: $P((i-1)N + (j-1)m_s + k) = P(k|i, j)$, $\hat{P}_n((i-1)N + (j-1)m_s + k) = P_n(k|i, j)$

A.1. Proof of Theorem 1

The proof of Theorem 1 is based on the delta method and the implicit function theorem.

We first note that under Assumptions 1 and 3, we have

$$[\hat{\mu}_{R,n}, \hat{P}_n] \rightarrow [\mu_R, P] \text{ a.s.} \quad (22)$$

and

$$\frac{1}{n} \sum_{1 \leq t \leq n} \mathbb{1}(s_t = i, a_t = j) \rightarrow w((i-1)m_a + j) \text{ a.s. as } n \rightarrow \infty.$$

By Slutsky's theorem and Proposition 3 in (Trevezas and Limnios 2009),

$$\sqrt{n}([\hat{\mu}_{R,n}, \hat{P}_n] - [\mu_R, P]) \Rightarrow \mathcal{N}(0, \Sigma_{R,P}), \quad (23)$$

where $\Sigma_{R,P}$ is defined in (8) and is an $(N + Nm_s) \times (N + Nm_s)$ matrix.

Second, define $F(Q', r', P')$ as a mapping from $\mathbb{R}^N \times \mathbb{R}^N \times \mathbb{R}^{Nm_s}$ to \mathbb{R}^N , representing the fixed point equation for the Bellman operator (1):

$$\begin{aligned} F(Q', r', P')((i-1)m_a + j) &= Q'((i-1)m_a + j) - r'((i-1)m_a + j) \\ &\quad - \gamma \sum_{1 \leq k \leq m_s} P'((i-1)N + (j-1)m_s + k) g_k(Q') \end{aligned}$$

for $1 \leq i \leq m_s$ and $1 \leq j \leq m_a$, where $g_k(Q') = \max_{1 \leq l \leq m_a} Q'((k-1)m_a + l)$, for $1 \leq k \leq m_s$.

Under Assumption 2, there exists an open neighborhood of Q , which we denote as Ω , such that for any $Q' \in \Omega$, $\arg \max_{1 \leq l \leq m_a} Q'((k-1)m_a + l)$ is still unique for each $1 \leq k \leq m_s$. Then, $g_k(Q')$ has all its partial derivatives exist and continuous. This implies that $F(Q', r', P')$ is continuously differentiable in $\Omega \times \mathbb{R}^N \times \mathbb{R}^{Nm_s}$.

Denote the partial derivatives of F as

$$\frac{\partial F}{\partial(Q', r', P')} = \left[\frac{\partial F}{\partial Q'}, \frac{\partial F}{\partial r'}, \frac{\partial F}{\partial P'} \right].$$

Note that $\frac{\partial F}{\partial Q'}$ is an $N \times N$ matrix. Denote its element at the $((i-1)m_a + j)$ -th row, $((k-1)m_a + l)$ -th column by

$$\begin{aligned} \frac{\partial F_{(i-1)m_a+j}}{\partial Q'_{(k-1)m_a+l}} &= \mathbb{1}(i=k, j=l) - \gamma P'((i-1)N + (j-1)m_s + k) \\ &\quad \times \mathbb{1} \left(Q'((k-1)m_a + l) = \max_{1 \leq u \leq m_a} Q'((k-1)m_a + u) \right). \end{aligned}$$

Define \tilde{P}' is an $N \times N$ matrix with

$$\begin{aligned} &\tilde{P}'((i-1)m_a + j, (k-1)m_a + l) \\ &= P'((i-1)N + (j-1)m_s + k) \mathbb{1} \left(Q'((k-1)m_a + l) = \max_{1 \leq u \leq m_a} Q'((k-1)m_a + u) \right) \end{aligned}$$

for $1 \leq i \leq m_s$, $1 \leq j \leq m_a$, $1 \leq k \leq m_s$, and $1 \leq l \leq m_a$. Then we have

$$\frac{\partial F}{\partial Q'} = I - \gamma \tilde{P}'.$$

Since all rows of \tilde{P}' sum up to one, \tilde{P}' can be interpreted as the transition matrix of a Markov chain with state space $\{(i, j) : 1 \leq i \leq m_s, 1 \leq j \leq m_a\}$. Then, $\frac{\partial F}{\partial Q'}$ is invertible for any $Q' \in \Omega$. This allows us to apply the implicit function theorem to the equation $F(Q, \mu_R, P) = 0$. In particular, there exists an open set $U \subset \mathbb{R}^N \times \mathbb{R}^{Nm_s}$, around $\mu_R \times P$, and a unique continuously differentiable function $\phi: U \rightarrow \mathbb{R}^N$, such that $\phi(\mu_R, P) = Q$, and for any $r' \times P' \in U$

$$F(\phi(r', P'), r', P') = 0.$$

In addition, the partial derivatives of ϕ satisfy

$$\nabla \phi(\mu_R, P) := \frac{\partial \phi}{\partial(r', P')} \Big|_{r'=\mu_R, P'=P} = - \left[\frac{\partial F}{\partial Q'} \right]^{-1} \left[\frac{\partial F}{\partial r'}, \frac{\partial F}{\partial P'} \right] \Big|_{Q'=Q, r'=\mu_R, P'=P}.$$

For $\frac{\partial F}{\partial r'}$, we have

$$\frac{\partial F}{\partial r'} \Big|_{Q'=Q, r'=\mu_R, P'=P} = -I_{N \times N}.$$

For $\frac{\partial F}{\partial P'}$, because

$$\begin{aligned} \frac{\partial F_{(i-1)m_a+j}}{\partial P'_{(k-1)N+(l-1)m_s+v}} \Big|_{Q'=Q, r'=\mu_R, P'=P} &= -\gamma \max_{0 \leq u \leq m_a} Q((v-1)m_a + u) \mathbf{1}(k=i, j=l) \\ &= -\gamma V^*(v) \mathbf{1}(k=i, j=l), \end{aligned}$$

for $1 \leq i \leq m_s$, $1 \leq j \leq m_a$, $1 \leq k \leq m_s$, $1 \leq l \leq m_a$, and $1 \leq v \leq m_s$, we can write

$$\frac{\partial F}{\partial P'} \Big|_{Q'=Q, r'=\mu_R, P'=P} = -\gamma \begin{pmatrix} (V^*)^T & & & \\ & \ddots & & \\ & & (V^*)^T & \\ & & & \ddots & \\ & & & & (V^*)^T \end{pmatrix},$$

which is an $N \times Nm_s$ matrix.

Now, from (22), by continuous mapping theorem, we have

$$\phi(\hat{\mu}_{R,n}, \hat{P}_n) - \phi(\mu_R, P) \rightarrow 0 \text{ a.s. as } n \rightarrow \infty,$$

which implies that $\hat{Q}_n \rightarrow Q$ a.s..

From (23), using the delta method, we have

$$\sqrt{n}(\hat{Q}_n - Q) = \sqrt{n}(\phi(\hat{\mu}_{R,n}, \hat{P}_n) - \phi(\mu_R, P)) \Rightarrow \mathcal{N}(0, \nabla \phi(\mu_R, P) \Sigma_{R,P} \nabla(\phi(\mu_R, P))^T) \text{ as } n \rightarrow \infty.$$

Plugging in the formula for $\nabla \phi(\mu_R, P)$ and $\Sigma_{R,P}$ derived above, we have

$$\nabla \phi(\mu_R, P) \Sigma_{R,P} \nabla \phi(\mu_R, P)^T = (I - \gamma \tilde{P}^{\pi*})^{-1} W^{-1} [D_R + \gamma^2 D_Q] ((I - \gamma \tilde{P}^{\pi*})^{-1})^T.$$

□

A.2. Proof of Corollary 1

Define $g_V(Q): \mathbb{R}^N \rightarrow \mathbb{R}^{m_s}$ as

$$g_V(Q) = (g_1(Q), \dots, g_{m_s}(Q)) = (V^{\pi^*}(1), \dots, V^{\pi^*}(m_s)),$$

which is continuously differentiable in an open neighborhood of Q . Then we can apply the delta method to get

$$\sqrt{n}(\hat{V}_n^* - V^{\pi^*}) = \sqrt{n}(g_V(\hat{Q}_n) - g_V(Q)) \Rightarrow \mathcal{N}(0, \nabla g_V(Q) \Sigma (\nabla g_V(Q))^T) \text{ as } n \rightarrow \infty.$$

Note that $\nabla g_V(Q)$ is a $m_s \times N$ matrix with $\nabla g_V(Q)(i, (j-1)m_a + k) = \mathbb{1} (i = j, k = a^*(i))$.

To see where the explicit expression for Σ_V is from, we will first rearrange the indexes such that we can write

$$\tilde{P}^{\pi^*} = \begin{pmatrix} P^{\pi^*} & \mathbf{0} \\ \# & \mathbf{0} \end{pmatrix},$$

where we use $\#$ as a generic placeholder for quantities that we do not need to characterize explicitly. Using this new indexing, we can write $\nabla g_V(Q) = [I, \mathbf{0}]$, and

$$D_R = \begin{pmatrix} D_R^{\pi^*} & \mathbf{0} \\ \mathbf{0} & \# \end{pmatrix}, W = \begin{pmatrix} W^{\pi^*} & \mathbf{0} \\ \mathbf{0} & \# \end{pmatrix}, D_Q = \begin{pmatrix} D_V^{\pi^*} & \mathbf{0} \\ \mathbf{0} & \# \end{pmatrix}.$$

Then

$$\begin{aligned} \nabla g_V(Q)(I - \gamma \tilde{P}^{\pi^*})^{-1} &= [I, \mathbf{0}] \sum_{i=0}^{\infty} \gamma^i (\tilde{P}^{\pi^*})^i \\ &= [I, \mathbf{0}] \sum_{i=0}^{\infty} \gamma^i \begin{pmatrix} (P^{\pi^*})^i & \mathbf{0} \\ \# & \mathbf{0} \end{pmatrix} \\ &= \sum_{i=0}^{\infty} \gamma^i [(P^{\pi^*})^i, \mathbf{0}] = [(I - \gamma P^{\pi^*})^{-1}, \mathbf{0}]. \end{aligned}$$

And thus,

$$\begin{aligned} &\nabla g_V(Q) \Sigma (\nabla g_V(Q))^T \\ &= [(I - \gamma P^{\pi^*})^{-1}, \mathbf{0}] \begin{pmatrix} W^{\pi^*} & \mathbf{0} \\ \mathbf{0} & \# \end{pmatrix}^{-1} \left[\begin{pmatrix} D_R^{\pi^*} & \mathbf{0} \\ \mathbf{0} & \# \end{pmatrix} + \begin{pmatrix} D_V^{\pi^*} & \mathbf{0} \\ \mathbf{0} & \# \end{pmatrix} \right] [(I - \gamma P^{\pi^*})^{-1}, \mathbf{0}]^T \\ &= (I - \gamma P^{\pi^*})^{-1} (W^{\pi^*})^{-1} [D_R^{\pi^*} + D_V^{\pi^*}] ((I - \gamma P^{\pi^*})^{-1})^T. \end{aligned}$$

Lastly, the asymptotic normality of $\hat{\chi}_n^{\pi^*}$ follows from the delta method as well. \square

A.3. Proof of Corollary 2

The convergence results for $\hat{\mu}_{R,n}$ and \hat{P}_n still hold in this case. Similar to the proof of Theorem 1, we define $F^{\tilde{\pi}}$ as a mapping $\mathbb{R}^{m_s} \times \mathbb{R}^N \times \mathbb{R}^{Nm_s} \rightarrow \mathbb{R}^{m_s}$ to represent the corresponding fixed point equation, i.e.,

$$F^{\tilde{\pi}}(V', r', P')(s) = V'(s) - \sum_{1 \leq a \leq m_a} r'(s, a) \tilde{\pi}(a|s) - \gamma \sum_{1 \leq a \leq m_a} \tilde{\pi}(a|s) \sum_{1 \leq s' \leq m_s} P'(s'|s, a) V'(s').$$

Note that $F^{\tilde{\pi}}(V^{\tilde{\pi}}, \mu_R, P) = 0$, $F^{\tilde{\pi}}$ is continuously differentiable and $I - \gamma P^{\tilde{\pi}}$ is invertible. We can thus apply the implicit function theorem. In particular, there exists an open set $U^{\tilde{\pi}}$ around $\mu_R \times P \in \mathbb{R}^N \times \mathbb{R}^{Nm_s}$, and a unique continuously differentiable function $\phi^{\tilde{\pi}}: U^{\tilde{\pi}} \rightarrow \mathbb{R}^N$, such that $\phi^{\tilde{\pi}}(\mu_R, P) = V^{\tilde{\pi}}$ and for any $r' \times P' \in U^{\tilde{\pi}}$,

$$F^{\tilde{\pi}}(\phi^{\tilde{\pi}}(r', P'), r', P') = 0.$$

In addition, the partial derivatives of $\phi^{\tilde{\pi}}$ satisfies

$$\nabla \phi^{\tilde{\pi}}(\mu_R, P) = \left. \frac{\partial \phi^{\tilde{\pi}}}{\partial (r', P')} \right|_{r'=\mu_R, P'=P} = - \left[\frac{\partial F^{\tilde{\pi}}}{\partial V'} \right]^{-1} \left[\frac{\partial F^{\tilde{\pi}}}{\partial r'}, \frac{\partial F^{\tilde{\pi}}}{\partial P'} \right] \Big|_{V'=V^{\tilde{\pi}}, r'=\mu_R, P'=P}$$

where

$$\left. \frac{\partial F^{\tilde{\pi}}}{\partial V'} \right|_{V'=V^{\tilde{\pi}}, r'=\mu_R, P'=P} = I - \gamma P^{\tilde{\pi}},$$

$$\left. \frac{\partial F^{\tilde{\pi}}}{\partial r'} \right|_{V'=V^{\tilde{\pi}}, r'=\mu_R, P'=P} = - \begin{pmatrix} \tilde{\pi}(\cdot|1)^T & & & \\ & \ddots & & \\ & & \tilde{\pi}(\cdot|i)^T & \\ & & & \ddots \\ & & & & \tilde{\pi}(\cdot|m_s)^T \end{pmatrix},$$

where $\tilde{\pi}(\cdot|i)^T = [\tilde{\pi}(1|i), \dots, \tilde{\pi}(j|i), \dots, \tilde{\pi}(m_a|i)]$

and

$$\left. \frac{\partial F^{\tilde{\pi}}}{\partial P'} \right|_{V'=V^{\tilde{\pi}}, r'=\mu_R, P'=P} = - \begin{pmatrix} (q_1^{\tilde{\pi}})^T & & & \\ & \ddots & & \\ & & (q_i^{\tilde{\pi}})^T & \\ & & & \ddots \\ & & & & (q_{m_s}^{\tilde{\pi}})^T \end{pmatrix},$$

where $(q_i^{\tilde{\pi}})^T = \gamma [\tilde{\pi}(1|i)(V^{\tilde{\pi}})^T, \dots, \tilde{\pi}(j|i)(V^{\tilde{\pi}})^T, \dots, \tilde{\pi}(m_a|i)(V^{\tilde{\pi}})^T]$, which is an N -dimensional vector.

Applying the delta method, we have

$$\begin{aligned} \sqrt{n}(\hat{V}_n^{\tilde{\pi}} - V^{\tilde{\pi}}) &= \sqrt{n}(\phi^{\tilde{\pi}}(\hat{\mu}_{R,n}, \hat{P}_n) - \phi^{\tilde{\pi}}(\mu_R, P)) \\ &\Rightarrow \mathcal{N}\left(0, \nabla \phi^{\tilde{\pi}}(\mu_R, P) \Sigma_{R,P} (\nabla \phi^{\tilde{\pi}}(\mu_R, P))^T\right) \text{ as } n \rightarrow \infty, \end{aligned}$$

where

$$\nabla \phi^{\tilde{\pi}}(\mu_R, P) \Sigma_{R,P} \nabla \phi^{\tilde{\pi}}(\mu_R, P)^T = (I - \gamma P^{\tilde{\pi}})^{-1} D^{\tilde{\pi}} ((I - \gamma P^{\tilde{\pi}})^{-1})^T.$$

□

A.4. Proof of Theorem 2

To analyze MDPs with non-unique optimal policies, we utilize the LP representation of the MDP. Consider ρ with $\rho(i) > 0$, for $i = 1, \dots, m_s$, the optimal value function V^* is the optimal solution of the following LP.

$$\begin{aligned} \min \quad & \sum_s \rho(s) V(s) \\ \text{subject to} \quad & V(s) \geq \mu_R(s, a) + \gamma \sum_{s' \in S} P(s'|s, a) V(s'), \quad \forall s, a \end{aligned}$$

The dual of the LP takes the form

$$\begin{aligned} \max \quad & \sum_{s,a} \mu_R(s, a) x_{s,a} \\ \text{subject to} \quad & \sum_a x_{s,a} - \gamma \sum_{s',a} P(s|s', a) x_{s',a} = \rho(s), \quad \forall s \\ & x_{s,a} \geq 0, \quad \forall s, a \end{aligned}$$

where the decision variables, $x_{s,a}$'s, are known as the occupancy measure of the MDP (Puterman 2014). If the MDP has more than one optimal policies, the dual problem has more than one optimal solutions, which further implies that the primal problem is degenerate. The degeneracy of the LP means that some constraints are redundant at the primal optimal solution, i.e., the optimal solution is at the intersection of more than m_s hyperplanes. Since the rows of the primal LP are linearly independent, in this case, there are multiple choices for the set of basic variables at the optimal solution. In addition, for any choice of ρ with positive entries, V^* is always the primal optimal solution (Puterman 2014). We denote v_k as the k -th optimal basis and A_k the corresponding column vectors, $1 \leq k \leq \bar{K}$, for some $\bar{K} > 1$.

For any $u \in U$, we denote the directional Jacobian of V with respect to P and μ_R as $D_u(P, \mu_R)$. We next show that the directional Jacobian is well defined. Vectorize the matrix $(I - \gamma P)$ row by row and denote the corresponding $m_s N$ -dimensional vector as Γ_P . We write $u = (u_P, u_R) \in \mathbb{R}^{m_s N + N}$ where $u_P \in \mathbb{R}^{m_s N}$ and $u_R \in \mathbb{R}^N$. We also define the mapping from Γ_P, μ_R to the k -th optimal basis as ψ_k , $1 \leq k \leq K$. In particular, $v_k = \psi_k(\Gamma_P, \mu_R) = A_k^{-1} \mu_R$. We first note that $\psi_k(\Gamma_P + t u_P, \mu_R + t u_R)$ is continuous in t . Next, fix $u \in U$, and consider the perturbed LP with constraint coefficients $\Gamma_P + t u_P$ and $\mu_R + t u_R$. Due to the continuity of ψ_k 's, there exists $t_u > 0$, such that at least one basis is still optimal for all $t \in (0, t_u)$. In particular, there exists k_u , $1 \leq k_u \leq \bar{K}$, such that $\psi_{k_u}(\Gamma_P + t u_P, \mu_R + t u_R)$ is the optimal solution to the perturbed LP for any $t \in (0, t_u)$. Lastly, we note that the directional derivative of ψ_{k_u} at (Γ_P, μ_R) takes the form $u^T G_{k_u}$, where

$$G_{k_u} = \left. \frac{\partial \psi_{k_u}}{\partial (\Gamma'_P, \mu'_R)} \right|_{\Gamma'_P = \Gamma_P, \mu'_R = \mu_R}.$$

This implies that $D_u(P, \mu_R) = u^T G_{k_u}$ is well defined.

Because there are \bar{K} optimal basis, there are $K \leq \bar{K}$ distinct G_k 's. Note that we allow $K < \bar{K}$ since some G_k 's may be equal. With a little abuse of notation (which might involve rearranging the

index), we write G_1, \dots, G_K as the distinct G_k 's. Then, we can partition the set U into K subsets, U_1, \dots, U_K , such that if $u \in U_k$, $D_u(P, \mu_R) = u^T G_k$.

Lastly, define $\hat{u}_n = (\hat{P}_n - P, \hat{\mu}_{R,n} - \mu_R) / \sqrt{\|\hat{P}_n - P\|^2 + \|\hat{\mu}_{R,n} - \mu_R\|^2}$. We have

$$\hat{V}_n^* - V^* = \sum_{k=1}^K G_k \mathbb{1}(\hat{u}_n \in U_k) (\hat{P}_n - P, \hat{\mu}_{R,n} - \mu_R) + o_P(\|(\hat{P}_n - P, \hat{\mu}_{R,n} - \mu_R)\|).$$

The convergence, $\sqrt{n}(\hat{V}_n^* - V^*) \Rightarrow \sum_{k=1}^K G_k \mathbb{1}(Z/\|Z\| \in U_k) Z$, then follows from the continuous mapping theorem. \square

A.5. Proof of Theorem 3

We use the LP representation of the constrained MDP. Define $x_{s,a}$ as the occupancy measure

$$x_{s,a} = \sum_{t=0}^{\infty} \gamma^t \mathbb{P}(S_t = s | S_0 \sim \rho),$$

where $S_0 \sim \rho$ denotes that the initial condition of the underlying Markov chain has probability mass function ρ . Then, $x_{s,a}$ satisfies the LP

$$\begin{aligned} & \max \quad \sum_{s,a} \mu_R(s,a) x_{s,a} \\ & \text{subject to} \quad \sum_{s,a} \mu_C(s,a) x_{s,a} \leq \eta \\ & \quad \sum_a x_{s,a} - \gamma \sum_{s',a} P(s|s',a) x_{s',a} = \rho(s), \quad \forall s \\ & \quad x_{s,a} \geq 0, \quad \forall s, a \end{aligned} \tag{24}$$

((24) is the dual formulation in the proof of Theorem 2 with an extra constraint.) The objective and the first constraint correspond to the objective and the constraint in the constrained MDP formulation. The second constraint can be deduced by a one-step analysis on the definition of occupancy measure. Let $(x_{s,a}^*)_{s,a}$ denote the optimal solution of the LP (24). Then, the optimal policy can be expressed as

$$\pi^*(a|s) = \frac{x_{s,a}^*}{\sum_a x_{s,a}^*}.$$

Note that the LP has $m_s + 1$ constraints (excluding the non-negativity constraints). Thus, a basis has $m_s + 1$ basic variables. Moreover, by our assumptions, the optimal solution is unique, which implies that perturbing the parameters μ_R, μ_C , and P does not immediately imply an overshoot to negative values for the reduced costs of the non-basic variables. In particular, when the perturbation is small enough, we still retain the same optimal basis. Next, we consider two cases depending on whether the first constraint is binding or not.

In the first case, the first constraint is non-binding. In this case, the optimal policy is deterministic, i.e., for any s , $x_{s,a} > 0$ for only one a . A small perturbation of the parameters still retains the same basic and non-basic variables, and the derived perturbed policy still retains the first constraint non-binding. The analysis then reduces to that of Corollary 2.

In the second case, the first constraint is binding. In this case, $x_{s,a} > 0$ for only one a , for all s except one state, s_r , where we can have $x_{s_r, a_1^*(s_r)} > 0$ and $x_{s_r, a_2^*(s_r)} > 0$ for two distinct actions $a_1^*(s_r), a_2^*(s_r)$. We denote the mixing parameter by $\alpha^* := \pi^*(a_1^*(s_r)|s_r)$, and so $\pi^*(a_2^*(s_r)|s_r) = 1 - \alpha^*$. Again, perturbing the parameters retains the same basic and non-basic variables. In particular, the first constraint remains binding in the perturbation, and the perturbed optimal policy π^* is still split at the same state and between the same actions.

We next make a few observations for case 2. First, we define

$$\begin{aligned} F_L^{\pi^*}(L, c, P', \alpha') = & L(s) - \left[c(s, a^*(s)) + \gamma \sum_{s'} L(s') P(s'|s, a^*(s)) \right] \mathbb{1}(s \neq s_r) \\ & - \left[\left(c(s_r, a_1^*(s_r)) + \gamma \sum_{s'} L(s') P(s'|s, a_1^*(s_r)) \right) \alpha' \right. \\ & \left. + \left(c(s_r, a_2^*(s_r)) + \gamma \sum_{s'} L(s') P(s'|s, a_2^*(s_r)) \right) (1 - \alpha') \right] \mathbb{1}(s = s_r) \end{aligned}$$

Then we have $F_L^{\pi^*}(L^{\pi^*}, \mu_C, P, \alpha^*) = 0$ and $I - \gamma P^{\pi^*}$ is invertible. Thus, by the implicit function theorem, there exists a continuously differentiable function ϕ_L such that $L^{\pi^*} = \phi_L(\mu_C, P, \alpha^*)$. Next, define

$$F_\alpha^{\pi^*}(c, P', \alpha') = \eta - \rho^T \phi_L(c, P', \alpha').$$

By applying the implicit function theorem again, we have there exists a continuously differentiable function ϕ_α such that $\alpha^* = \phi_\alpha(\mu_C, P)$. Following the same line of arguments as above, we define

$$\begin{aligned} F_V^{\pi^*}(V', r', P', \alpha') = & V'(s) - \left[r'(s, a^*(s)) + \gamma \sum_{s'} V'(s') P(s'|s, a^*(s)) \right] \mathbb{1}(s \neq s_r) \\ & - \left[\left(r'(s_r, a_1^*(s_r)) + \gamma \sum_{s'} V'(s') P(s'|s, a_1^*(s_r)) \right) \alpha' \right. \\ & \left. + \left(r'(s_r, a_2^*(s_r)) + \gamma \sum_{s'} V'(s') P(s'|s, a_2^*(s_r)) \right) (1 - \alpha') \right] \mathbb{1}(s = s_r) \end{aligned}$$

Then, by the implicit function theorem, V^* can be viewed as a continuously differentiable function of μ_R , P and α^* , i.e., $V^* = \phi_I(\mu_R, P, \alpha^*)$. Since α^* can be viewed as a continuously differentiable function of μ_C and P , V^* can also be viewed as a function of μ_R , μ_C , and P , i.e., $V^* = \phi_V(\mu_R, \mu_C, P)$.

Lastly, using the delta method, we have

$$\sqrt{n}(\hat{V}_n^* - V^*) \Rightarrow N(0, \nabla \phi_V(\mu_R, \mu_C, P) \Sigma_{R,C,P} (\nabla \phi_V(\mu_R, \mu_C, P))^T) \text{ as } n \rightarrow \infty,$$

where $\nabla \phi_V(\mu_R, \mu_C, P)$ is the Jacobian of ϕ_V evaluated at μ_R, μ_C, P , and $\Sigma_{R,C,P}$ is the estimation covariance matrix of μ_R, μ_C, P .

To see where the explicit expression for Σ_c is from, we note that

$$\frac{\partial \phi_V}{\partial(r', c', P')} = \frac{\partial \phi_I}{\partial(r', c', P')} + \frac{\partial \phi_I}{\partial \alpha'} \frac{\partial \phi_\alpha}{\partial(r', c', P')}.$$

For ϕ_I , we have

$$\left. \frac{\partial \phi_I}{\partial(r', P', \alpha')} \right|_{r'=\mu_R, P'=P, \alpha'=\alpha^*} = - \left[\frac{\partial F_V^{\pi^*}}{\partial V'} \right]^{-1} \left[\frac{\partial F_V^{\pi^*}}{\partial r'}, \frac{\partial F_V^{\pi^*}}{\partial P'}, \frac{\partial F_V^{\pi^*}}{\partial \alpha'} \right] \Big|_{V'=V^*, r'=\mu_R, P'=P, \alpha'=\alpha^*}$$

and

$$\frac{\partial \phi_I}{\partial c'} = \mathbf{0},$$

where

$$\begin{aligned} \left. \frac{\partial F_V^{\pi^*}}{\partial V'} \right|_{V'=V^*, r'=\mu_R, P'=P, \alpha'=\alpha^*} &= I - \gamma P^{\pi^*}, & \left. \frac{\partial F_V^{\pi^*}}{\partial r'} \right|_{V'=V^*, r'=\mu_R, P'=P, \alpha'=\alpha^*} &= -G^{\pi^*}, \\ \left. \frac{\partial F_V^{\pi^*}}{\partial P'} \right|_{V'=V^*, r'=\mu_R, P'=P, \alpha'=\alpha^*} &= -H_V^{\pi^*}, \text{ and } & \left. \frac{\partial F_V^{\pi^*}}{\partial \alpha'} \right|_{V'=V^*, r'=\mu_R, P'=P, \alpha'=\alpha^*} &= -h_V. \end{aligned}$$

Similarly, for ϕ_α , we have

$$\left. \frac{\partial \phi_\alpha}{\partial(c', P')} \right|_{c'=\mu_C, P'=P} = - \left[\frac{\partial F_\alpha^{\pi^*}}{\partial \alpha'} \right]^{-1} \left[\frac{\partial F_\alpha^{\pi^*}}{\partial c'}, \frac{\partial F_\alpha^{\pi^*}}{\partial P'} \right] \Big|_{\alpha'=\alpha^*, c'=\mu_C, P'=P} \text{ and } \frac{\partial \phi_\alpha}{\partial r'} = \mathbf{0},$$

where

$$\begin{aligned} \left. \frac{\partial F_\alpha^{\pi^*}}{\partial \alpha'} \right|_{\alpha'=\alpha^*, c'=\mu_C, P'=P} &= -\rho^T (I - \gamma P^{\pi^*})^{-1} h_L, & \left. \frac{\partial F_\alpha^{\pi^*}}{\partial c'} \right|_{V'=V^*, r'=\mu_R, P'=P, \alpha'=\alpha^*} &= -\rho^T (I - \gamma P^{\pi^*})^{-1} G^{\pi^*}, \\ \text{and } \left. \frac{\partial F_\alpha^{\pi^*}}{\partial P'} \right|_{V'=V^*, r'=\mu_R, P'=P, \alpha'=\alpha^*} &= -\rho^T (I - \gamma P^{\pi^*})^{-1} H_L^{\pi^*}. \end{aligned}$$

□

Appendix B: Proof of the results in Section 4

B.1. Proof of lemma 1

For any given policy π , by the balance equation for Markov Chains, its induced stationary distribution w_π satisfies

$$\sum_{k,l} w_\pi((k-1)m_a + l) P(i|s=k, a=l) \pi(a=j|s=i) = w_\pi((i-1)m_a + j)$$

for any $i \in \mathcal{S}, j \in \mathcal{A}$. Summing up across j 's for each i , we have

$$\begin{aligned} & \sum_j w_\pi((i-1)m_a + j) \\ &= \sum_j \sum_{k,l} w_\pi((k-1)m_a + l) P(i|s=k, a=l) \pi(a=j|s=i) \\ &= \sum_{k,l} w_\pi((k-1)m_a + l) P(i|s=k, a=l) \end{aligned}$$

On the other hand, for any w in \mathcal{W} , π_w satisfies

$$\begin{aligned} & \sum_{k,l} w((k-1)m_a + l) P(i|s=k, a=l) \pi_w(a=j|s=i) \\ &= \sum_{k,l} w((k-1)m_a + l) P(i|s=k, a=l) w((i-1)m_a + j) / \sum_u w((i-1)m_a + u) \\ &= \sum_u w((i-1)m_a + u) w((i-1)m_a + j) / \sum_u w((i-1)m_a + u) = w((i-1)m_a + j) \end{aligned}$$

for all $i \in S$. Thus, w is the stationary distribution of the Markov chain with transition matrix \tilde{P}^{π_w} . \square

B.2. Proof of Theorem 4

We use the superscript n to mark the dependence of B_k 's and K on n , e.g., B_k^n and K^n when the sampling budget is n .

In the first case in the first result, after collecting B_{k-1}^n data points, as $C_l > 0$, $w_{\hat{k}}(s, a) > 0$ for any $(s, a) \in \mathcal{S} \times \mathcal{A}$. Then, as $B_k^n - B_{k-1}^n \rightarrow \infty$, each (s, a) is visited infinitely often. By Theorem 1, $\hat{Q}_{B_k^n} \rightarrow Q$ a.s. as $n \rightarrow \infty$. Next, because $\hat{c}_{ij, B_k^n} \rightarrow c_{ij}$ a.s. as $n \rightarrow \infty$, by continuous mapping theorem, we have $\hat{\pi}_{B_k^n} \rightarrow \pi^e$ a.s. as $n \rightarrow \infty$.

In the second case in the first result, because $C_l > 0$, there exists $\epsilon > 0$ such that $\min_{(s,a) \in \mathcal{S} \times \mathcal{A}} w_k(s, a) > \epsilon$ for all $k = 1, 2, \dots$. Because $B_k - B_{k-1} \geq N$, for any $(s, a) \in \mathcal{S} \times \mathcal{A}$, there is a positive probability of visiting (s, a) at each stage. Thus, as $K \rightarrow \infty$, each (s, a) is visited infinitely often. The rest follows the same line of argument as in the first case.

In the second result, we first introduce a few more notations. Let $b_k = B_k - B_{k-1}$, i.e., the number of steps in the k -th iteration. Let $s_{k,i}$, $1 \leq i \leq b_k$, denote the state visited at the i -th step in the k -th iteration and $a_{k,i}$ denote the action taken then.

Note that w_k is the stationary distribution of \tilde{P}^{π_k} , i.e., the Markov chain under the exploration policy for iteration k . Then, by the large deviations theory for the empirical measure of a Markov chain (Theorem 1.2 in Ellis et al. (1988)), we have for any given $\epsilon > 0$,

$$\sup_{(s_{k,0}, a_{k,0}) \in \mathcal{S} \times \mathcal{A}} \mathbb{P} \left(\sup_{(s,a) \in \mathcal{S} \times \mathcal{A}} \left| \frac{1}{b_k} \sum_{i=1}^{b_k} \mathbb{1}(s_{k,i} = s, a_{k,i} = a) - w_k(s, a) \right| \geq \epsilon \mid s_{k,0}, a_{k,0} \right) \leq \exp(-b_k I(\epsilon, w_k))$$

for some large deviations rate function I , where $I(\epsilon, w_k) > 0$ when $\epsilon > 0$ and I is continuous in w_k (I is continuous in \tilde{P}^{π_k} , which is in turn continuous in w_k). Because $w_k \rightarrow w^*$ as $k \rightarrow \infty$ and $I(\epsilon, w^*) > 0$, there exists $\bar{I}(\epsilon) > 0$ such that $\min_{k \geq 1} I(\epsilon, w_k) \geq \bar{I}(\epsilon)$. Then, as $b_k \geq \gamma k$,

$$\begin{aligned} & \sum_{k=1}^{\infty} \sup_{(s_{k,0}, a_{k,0}) \in \mathcal{S} \times \mathcal{A}} \mathbb{P} \left(\sup_{(s,a) \in \mathcal{S} \times \mathcal{A}} \left| \frac{1}{b_k} \sum_{i=1}^{b_k} \mathbb{1}(s_{k,i} = s, a_{k,i} = a) - w_k(s, a) \right| \geq \epsilon \mid s_{k,0}, a_{k,0} \right) \\ & \leq \sum_{k=1}^{\infty} \exp(-\gamma k \bar{I}(\epsilon)) < \infty \end{aligned}$$

By Borel-Cantelli Lemma,

$$\frac{1}{b_k} \sum_{i=1}^{b_k} \mathbb{1}(s_{k,i} = s, a_{k,i} = a) - w_k(s, a) \rightarrow 0 \text{ a.s. as } k \rightarrow \infty.$$

Next, note that for any $(s, a) \in \mathcal{S} \times \mathcal{A}$,

$$\begin{aligned} & \frac{1}{B_K} \sum_{k=1}^K \sum_{i=1}^{b_k} \mathbb{1}(s_{k,i} = s, a_{k,i} = a) - w^*(s, a) \\ &= \sum_{k=1}^K \frac{b_k}{B_K} \left(\frac{1}{b_k} \sum_{i=1}^{b_k} \mathbb{1}(s_{k,i} = s, a_{k,i} = a) - w_k(s, a) + w_k(s, a) - w^*(s, a) \right) \\ &= \frac{1}{B_K} \sum_{k=1}^K \left(\frac{1}{b_k} \sum_{i=1}^{b_k} \mathbb{1}(s_{k,i} = s, a_{k,i} = a) - w_k(s, a) \right) b_k + \sum_{k=1}^K \frac{b_k}{B_K} (w_k(s, a) - w^*(s, a)) \end{aligned}$$

As

$$\frac{1}{b_k} \sum_{i=1}^{b_k} \mathbb{1}(s_{k,i} = s, a_{k,i} = a) - w_k(s, a) \rightarrow 0 \text{ and } w_k \rightarrow w^* \text{ a.s. as } k \rightarrow \infty,$$

by Silverman–Toeplitz Theorem,

$$\frac{1}{B_K} \sum_{k=1}^K \sum_{i=1}^{b_k} \mathbb{1}(s_{k,i} = s, a_{k,i} = a) - w^*(s, a) \rightarrow 0 \text{ a.s. as } n \rightarrow \infty.$$

□

Appendix C: Proof of the Results in Section 6

C.1. Proof of Theorem 5

Denote

$$[\hat{\mu}_{R,n}, \hat{P}_n]_{S_0} = [\hat{\mu}_{R,n}((i-1)m_a + j), \hat{P}_n((i-1)N + (j-1)m_s + k)]_{i \in S_0, 1 \leq j \leq m_a, 1 \leq k \leq m_s},$$

and

$$[\mu_R, P]_{S_0} = [\mu_R((i-1)m_a + j), P((i-1)N + (j-1)m_s + k)]_{i \in S_0, 1 \leq j \leq m_a, 1 \leq k \leq m_s}.$$

By Assumption 6, we have

$$\sqrt{n}([\hat{\mu}_{R,n}, \hat{P}_n]_{S_0} - [\mu_R, P]_{S_0}) \Rightarrow \mathcal{N}(0, \Sigma_{R,P,S_0}) \text{ where } \Sigma_{R,P,S_0} = \begin{pmatrix} W_{S_0}^{-1} D_R^{S_0} & \mathbf{0} \\ \mathbf{0} & D_P^{S_0} \end{pmatrix}.$$

Notice $M_g \circ M_I^{S_0} \circ \mathcal{T}_{\hat{\mu}_{R,n}, \hat{P}_n}$ only involves random variables $[\hat{\mu}_{R,n}, \hat{P}_n]_{S_0}$. Changing the distribution of $[\hat{\mu}_{R,n}, \hat{P}_n]_{S \setminus S_0}$ will not change the distribution of \hat{Q}_n^M . We can thus assign auxiliary random variables to $\hat{\mu}_{R,n}$ and \hat{P}_n for all $i \notin S_0$, $1 \leq j \leq m_a$, $1 \leq k \leq m_s$. In particular, we use independent random variables for each $i \notin S_0$ by letting

$$\hat{\mu}_{R,n}((i-1)m_a + j) \stackrel{D}{=} \frac{1}{\sqrt{n}} \mathcal{N}(\mu_R((i-1)m_a + j), 1)$$

$$\hat{P}_n((i-1)N + (j-1)m_s + k) \stackrel{D}{=} \frac{1}{\sqrt{n}} \mathcal{N}(P((i-1)N + (j-1)m_s + k), 1).$$

Doing so, we extend the $m_{s_0}m_a$ -dimensional random variable $[\hat{\mu}_{R,n}, \hat{P}_n]_{S_0}$ to an $m_s m_a$ -dimensional random variable $[\hat{\mu}_{R,n}, \hat{P}_n]_S$ and

$$\sqrt{n}([\hat{\mu}_{R,n}, \hat{P}_n]_S - [\mu_R, P]_S) \Rightarrow \mathcal{N}(0, \Sigma_{R,P,S}) \text{ where } \Sigma_{R,P,S} = \begin{pmatrix} \Sigma_{R,P,S_0} & \mathbf{0} \\ \mathbf{0} & I \end{pmatrix}.$$

Similar to the proof of Theorem 1, define

$$F_M(Q', r', P') = Q' - M_g \circ M_I^{S_0} \circ \mathcal{T}_{r', P'}(Q').$$

By Assumption 4, M_g is max-norm non-expansion. Then, $M_g \circ M_I^{S_0}$ is also max-norm non-expansion, which implies that $\nabla(M_g \circ M_I^{S_0})$ has all its eigenvalues less than or equal to 1. Thus,

$$\frac{\partial F_M}{\partial Q'} = \nabla M(\mathcal{T}_{r', P'}(Q'))(I - \gamma \tilde{P}')$$

is invertible. By Assumption 4, we have $F_M(Q^M, \mu_R, P) = 0$. By Assumption 5, there exists a neighborhood Ω_M around (Q^M, μ_R, P) , such that F_M is continuously differentiable on Ω_M . Then, applying the implicit function theorem, we have that there exists an open set $E_M \subset \Omega_M$ and a continuously differentiable function ϕ_M on E_M , such that $\phi_M(\mu_R, P) = Q^M$ and

$$\nabla \phi_M(\mu_R, P) = - \left[\frac{\partial F_M}{\partial Q'} \right]^{-1} \left[\frac{\partial F_M}{\partial r'}, \frac{\partial F_M}{\partial P'} \right] \Bigg|_{Q'=Q^M, r'=\mu_R, P'=P}.$$

Using the delta method, we have

$$\sqrt{n}(\hat{Q}_n^M - Q^M) \Rightarrow \mathcal{N}(0, \nabla \phi_M(\mu_R, P) \Sigma_{R,P,S} (\nabla \phi_M(\mu_R, P))^T) = \mathcal{N}(0, \Sigma_{S_0}^M) \text{ as } n \rightarrow \infty.$$

□

C.2. Proof of Lemma 2

We have

$$\begin{aligned} \sqrt{n}(\hat{M}_n - M^*) &= A_n^{-1} \sqrt{n} \left(\sum_{i=1}^n \phi(s_i, a_i) \psi(s'_i) K_\psi^{-1} - \phi(s_i, a_i) \phi^T(s_i, a_i) M^* \right) - \sqrt{n} A_n^{-1} M^* \\ &= \left(\frac{1}{n} A_n \right)^{-1} \sqrt{n} \left(\frac{1}{n} \sum_{i=1}^n G(x_i) \right) - \left(\frac{1}{n} A_n \right)^{-1} \frac{1}{\sqrt{n}} M^*. \end{aligned}$$

We first note that as $n \rightarrow \infty$,

$$\frac{1}{\sqrt{n}} M^* \Rightarrow 0 \quad \text{and} \quad \frac{1}{n} A_n \Rightarrow \mathbb{E}_{\bar{w}}[\phi(s(\bar{X}), a(\bar{X})) \phi^T(s(\bar{X}), a(\bar{X}))] = \sum_{s,a} w_{s,a} \phi(s, a) \phi(s, a)^T.$$

We next note that because

$$\mathbb{E}_{\bar{w}}[G_v(\bar{X})] = \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} w_{s,a} \left(\phi(s,a) \phi(s,a)^T M^* \left(\sum_{s' \in \mathcal{S}} \psi(s') \psi(s')^T \right) K_\psi^{-1} - \phi(s,a) \phi(s,a)^T M^* \right) = 0.$$

By Markov Chain Central Limit Theorem, we have

$$\sqrt{n} \left(\frac{1}{n} \sum_{t=1}^n G(\bar{X}_t) \right) \Rightarrow \mathcal{N}(0, \Sigma_G) \text{ as } n \rightarrow \infty,$$

where

$$\Sigma_G = \text{Var}(G_v(X_0)) + 2 \sum_{i=1}^{\infty} \text{Cov}(G_v(X_0), G_v(X_i)).$$

Then, by Slutsky's theorem,

$$\sqrt{n}(M_{v,n} - M_v^*) \Rightarrow N(0, \Upsilon_E^{-1} \Sigma_G \Upsilon_E^{-1}) \text{ as } n \rightarrow \infty.$$

Because $\hat{P}_n^K(s'|s,a) = F_{\phi,\psi}(s,a,s') \hat{M}_{n,v}$,

$$\sqrt{n}(\hat{P}_n^K(s'|s,a) - P(s'|s,a)) \Rightarrow \mathcal{N}(0, F_{\phi,\psi}(s,a,s') \Upsilon_E^{-1} \Sigma_G \Upsilon_E^{-1} (F_{\phi,\psi}(s,a,s'))^T) \text{ as } n \rightarrow \infty.$$

□

C.3. Proof of Lemma 3

We have

$$\begin{aligned} \sqrt{n}(\hat{\theta}_{\mu,n} - \theta_\mu) &= (\Phi_{\mu,n}^T \Phi_{\mu,n})^{-1} \sqrt{n} \sum_{i=1}^n \phi_\mu(s_i, a_i) (r_i - \phi_\mu(s_i, a_i)^T \theta_\mu) \\ &= \left(\frac{1}{n} \Phi_{\mu,n}^T \Phi_{\mu,n} \right)^{-1} \sqrt{n} \left(\frac{1}{n} \sum_{i=1}^n H(\tilde{x}_i) \right). \end{aligned}$$

We first note that

$$\frac{1}{n} \Phi_{\mu,n}^T \Phi_{\mu,n} \Rightarrow \Upsilon_\mu \text{ as } n \rightarrow \infty.$$

Next, because $\mathbb{E}_w[H(\tilde{X})] = 0$, by Markov Chain Central Limit Theorem, we have

$$\sqrt{n} \left(\frac{1}{n} \sum_{i=1}^n H(\tilde{x}_i) \right) \Rightarrow \mathcal{N}(0, \Sigma_H) \text{ as } n \rightarrow \infty,$$

where

$$\Sigma_H = \text{Var}_w(H(\tilde{X}_0)) + 2 \sum_{t=1}^{\infty} \text{Cov}_w(H(\tilde{X}_0), H(\tilde{X}_t)).$$

Then, by Slutsky's theorem,

$$\sqrt{n}(\hat{\theta}_{\mu,n} - \theta_\mu) \Rightarrow \mathcal{N}(0, \Upsilon_\mu^{-1} \Sigma_H \Upsilon_\mu^{-1})$$

Lastly, because $\hat{\mu}_{R,n}(s,a) = \phi_\mu(s,a)^T \hat{\theta}_{\mu,n}$,

$$\sqrt{n}(\hat{\mu}_{R,n}(s,a) - \mu_R(s,a)) \Rightarrow N(0, \phi_\mu(s,a)^T \Upsilon_\mu^{-1} \Sigma_H \Upsilon_\mu^{-1} \phi_\mu(s,a)) \text{ as } n \rightarrow \infty.$$

□

C.4. Proof of Theorem 6

The proof of Theorem 6 is similar to that of Theorem 1. We first note that under Assumption 1, by Lemmas 2 and 3,

$$\sqrt{n}([\hat{\mu}_{R,n}, P_n] - [\mu_R, P]) \Rightarrow \mathcal{N}(0, \Sigma_{R,P}^K), \text{ as } n \rightarrow \infty,$$

where

$$\Sigma_{R,P}^K = \begin{pmatrix} D_R^K & \mathbf{0} \\ \mathbf{0} & D_P^K \end{pmatrix}, \quad D_P^K = F_{\phi,\psi} \Upsilon_E^{-1} \Sigma_G \Upsilon_E^{-1} F_{\phi,\psi}^T, \quad D_R^K = \phi_\mu^T \Upsilon_\mu^{-1} \Sigma_H \Upsilon_\mu^{-1} \phi_\mu.$$

The covariance block $\mathbf{0}$ in $\Sigma_{R,P}^K$ is due to the fact that under Assumption 1,

$$\mathbb{E}_{\bar{w}} \left[(r(\bar{X}) - \phi_\mu(\bar{X})^T \theta_\mu) (\phi(s(\bar{X}), a(\bar{X})) \psi(s'(\bar{X})) K_\psi^{-1} - \phi(s(\bar{X}), a(\bar{X})) \phi(s(\bar{X}), a(\bar{X}))^T M^*) \right] = 0.$$

Next, define $F(Q', r', P')$ as a mapping from $\mathbb{R}^N \times \mathbb{R}^N \times \mathbb{R}^{Nm_s}$ to \mathbb{R}^N , representing the fixed point equation for the Bellman operator (1). Specifically,

$$\begin{aligned} F(Q', r', P')((i-1)m_a + j) = & Q'((i-1)m_a + j) - r'((i-1)m_a + j) \\ & - \gamma \sum_{1 \leq k \leq m_s} P'((i-1)N + (j-1)m_s + k) g_k(Q'), \end{aligned}$$

for $1 \leq i \leq m_s$ and $1 \leq j \leq m_a$, where $g_k(Q') = \max_l Q'((k-1)m_a + l)$, for $1 \leq k \leq m_s$. We can then apply the implicit function theorem to the equation $F(Q, \mu_R, P) = 0$. In particular, there exists an open set $U \subset \mathbb{R}^N \times \mathbb{R}^{Nm_s}$, around (μ_R, P) , and a unique continuously differentiable function $\nu: U \rightarrow \mathbb{R}^N$, such that $\nu(\mu_R, P) = Q$, and for any $(r', P') \in U$

$$F(\nu(r', P'), r', P') = 0.$$

Using the delta method, we have

$$\sqrt{n}(\hat{Q}_n^K - Q) = \sqrt{n}(\nu(\hat{\mu}_{R,n}, \hat{P}_n^K) - \nu(\mu_R, P)) \Rightarrow \mathcal{N}(0, \nabla \nu(\mu_R, P) \Sigma_{R,P}^K (\nabla \nu(\mu_R, P))^T) \text{ as } n \rightarrow \infty.$$

Plugging in the formulas for $\nabla \nu(\mu_R, P)$ and $\Sigma_{R,P}^K$, we have

$$\nabla \nu(\mu_R, P) \Sigma_{R,P}^K \nabla \nu(\mu_R, P)^T = (I - \gamma \tilde{P}^{\pi^*})^{-1} [D_R^K + \gamma^2 V_D D_P^K V_D] ((I - \gamma \tilde{P}^{\pi^*})^{-1})^T.$$

□

Appendix D: Optimal Exploration under a Different Objective

In Q-OCBA, the sequential exploration policy is derived by maximizing the worst-case relative discrepancy among all Q-value estimates. If one is interested in obtaining the best estimate of χ^* (i.e., the optimal value function initialized according to ρ), then it would be more natural to solve

$$\min_{w \in \mathcal{W}_\eta} \sigma_\chi^2. \tag{25}$$

From (25), the optimal exploration policy π_w can be derived based on Lemma 1. The motivation is that by doing so, the half-width of the CI for χ^* is minimized. We refer to this policy as χ^* -OCBA

Table 4 compares the half-width of 95% CI for χ^* constructed using data collected under different exploration policies. The half-width is estimated based on 10^3 independent replications of the procedures and the estimation errors (half width of the 95% CIs) are around 0.01. We vary r_L from 1 to 3. The total sampling budget is $n = 10^4$. 30% of the budget is used to collect some initial data under RE with $\pi(1|s) = 0.8$. We then use the initial data to train ϵ -greedy with $\epsilon = 0.2$, UCRL, PSRL, Q-OCBA, and χ^* -OCBA. This two-stage implementation creates a warm start to help UCRL and PSRL achieve the correct coverage rate, i.e., 95%. In particular, using this implementation, all policies in Table 4 are able to achieve the nominal coverage.

We observe that χ^* -OCBA leads to shorter CIs in χ^* estimation compared to benchmark policies. For example, when $r_L = 3$, the half-width of the CI constructed from χ^* -OCBA is at least 40% less than those constructed using PSRL or UCRL. The performances of χ^* -OCBA and Q-OCBA appear quite similar, with χ^* -OCBA moderately better. For example, when $r_L = 3$, the half-width of the CI constructed from χ^* -OCBA is 24% less than that using Q-OCBA. We attribute the improvement to the targeted criterion on χ^* estimation used in our policy design of χ^* -OCBA. We also note that the performance of χ^* -OCBA is robust against different values of r_L . This is not the case for some benchmark policies such as PSRL, which is designed to strike a balance between exploration and exploitation. For example, when $r_L = 1$, PSRL can perform as well as Q-OCBA, with the half-width of the CI being 2.16 for PSRL, and 2.2 for χ^* -OCBA. However, when $r_L = 2$ or 3, χ^* -OCBA demonstrates significant advantages over PSRL. Specifically, the half width of the CI using χ^* -OCBA is less than 58% of that using PSRL. This observation is consistent with our observation from Figure 2 where we consider the probability of correct selection.

Table 4 Comparison of CI half-widths for different exploration policies

r_L	0.2-greedy	RE(0.6)	PSRL	UCRL	Q-OCBA	χ^* -OCBA
1	3.50	3.47	2.16	6.25	2.58	2.20
2	3.12	2.07	3.26	3.31	1.38	1.35
3	1.47	1.48	1.40	1.40	1.06	0.81

References

- Achiam J, Held D, Tamar A, Abbeel P (2017) Constrained policy optimization. *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, 22–31 (JMLR. org).
- Altman E (1999) *Constrained Markov Decision Processes*, volume 7 (CRC Press).
- Amini A, Gilitschenski I, Phillips J, Moseyko J, Banerjee R, Karaman S, Rus D (2020) Learning robust control policies for end-to-end autonomous driving from data-driven simulation. *IEEE Robotics and Automation Letters* 5(2):1143–1150.

-
- Audibert JY, Bubeck S (2010) Best arm identification in multi-armed bandits. *COLT-23th Conference on learning theory-2010*, 13–p.
- Azar MG, Osband I, Munos R (2017) Minimax regret bounds for reinforcement learning. *International Conference on Machine Learning*, 263–272 (PMLR).
- Barton RR (2012) Tutorial: Input uncertainty in output analysis. Laroque C, Himmelspace J, Pasupathy R, Rose O, Uhrmacher A, eds., *Proceedings of the 2012 Winter Simulation Conference*, 1–12 (Piscataway, New Jersey: IEEE).
- Barton RR, Nelson BL, Xie W (2013) Quantifying input uncertainty via simulation confidence intervals. *INFORMS Journal on Computing* 26(1):74–87.
- Barton RR, Schruben LW (2001) Resampling methods for input modeling. Peters BA, Smith JS, Medeiros DJ, Rohrer MW, eds., *Proceedings of the 2001 Winter Simulation Conference*, volume 1, 372–378 (Piscataway, New Jersey: IEEE).
- Bayraksan G, Morton DP (2006) Assessing solution quality in stochastic programs. *Mathematical Programming* 108(2-3):495–514.
- Bellemare MG, Dabney W, Munos R (2017) A distributional perspective on reinforcement learning. *International Conference on Machine Learning*, 449–458 (PMLR).
- Boutillier C, Lu T (2016) Budget allocation using weakly coupled, constrained markov decision processes. *Proceedings of the Thirty-Second Conference on Uncertainty in Artificial Intelligence*, 52–61 (AUAI Press).
- Chen CH, He D, Fu M (2006) Efficient dynamic simulation allocation in ordinal optimization. *IEEE Transactions on Automatic Control* 51(12):2005–2009.
- Chen Ch, Lee LH (2011) *Stochastic Simulation Optimization: An Optimal Computing Budget Allocation*, volume 1 (World scientific).
- Chen S, Devraj A, Busic A, Meyn S (2020) Explicit mean-square error bounds for monte-carlo and linear stochastic approximation. *International Conference on Artificial Intelligence and Statistics*, 4173–4183 (PMLR).
- Chen W, Gao S, Chen CH, Shi L (2013) An optimal sample allocation strategy for partition-based random search. *IEEE Transactions on Automation Science and Engineering* 11(1):177–186.
- Cheng RC, Holland W (1997) Sensitivity of computer simulation experiments to errors in input data. *Journal of Statistical Computation and Simulation* 57(1-4):219–241.
- Cheng RC, Holland W (2004) Calculation of confidence intervals for simulation output. *ACM Transactions on Modeling and Computer Simulation* 14(4):344–362.
- Chick SE (2001) Input distribution selection for simulation experiments: Accounting for input uncertainty. *Operations Research* 49(5):744–758.
- Chow Y, Ghavamzadeh M, Janson L, Pavone M (2017) Risk-constrained reinforcement learning with percentile risk criteria. *The Journal of Machine Learning Research* 18(1):6070–6120.
- Corso A, Moss RJ, Koren M, Lee R, Kochenderfer MJ (2020) A survey of algorithms for black-box safety validation. *arXiv preprint arXiv:2005.02979* .

- Devraj AM, Meyn SP (2017) Zap q-learning. *Proceedings of the 31st International Conference on Neural Information Processing Systems*, 2232–2241.
- Dong J, Zhu Y (2016) Three asymptotic regimes for ranking and selection with general sample distributions. *Proceedings of the 2016 Winter Simulation Conference*, 277–288 (IEEE Press).
- Ellis RS, et al. (1988) Large deviations for the empirical measure of a markov chain with an application to the multivariate empirical measure. *The Annals of Probability* 16(4):1496–1508.
- Feinberg EA, Rothblum UG (2012) Splitting randomized stationary policies in total-reward markov decision processes. *Mathematics of Operations Research* 37(1):129–153.
- Gao S, Chen W (2016) Efficient feasibility determination with multiple performance measure constraints. *IEEE Transactions on Automatic Control* 62(1):113–122.
- Gao S, Xiao H, Zhou E, Chen W (2017) Robust ranking and selection with optimal computing budget allocation. *Automatica* 81:30–36.
- Glynn P, Juneja S (2004) A large deviations perspective on ordinal optimization. *Proceedings of the 36th conference on Winter Simulation Conference*, 577–585.
- Gordon GJ (1995) Stable function approximation in dynamic programming. *Machine Learning Proceedings 1995*, 261–268 (Elsevier).
- Higle JL, Sen S (2013) *Stochastic decomposition: a statistical method for large scale stochastic linear programming*, volume 8 (Springer Science & Business Media).
- Jaksch T, Ortner R, Auer P (2010) Near-optimal regret bounds for reinforcement learning. *Journal of Machine Learning Research* 11(Apr):1563–1600.
- Jia QS (2012) Efficient computing budget allocation for simulation-based policy improvement. *IEEE Transactions on Automation Science and Engineering* 9(2):342–352.
- Jin C, Allen-Zhu Z, Bubeck S, Jordan MI (2018) Is q-learning provably efficient? *Advances in Neural Information Processing Systems*, 4863–4873.
- Kakade SM (2003) On the sample complexity of reinforcement learning, PhD Thesis, University College London.
- Kalashnikov D, Irpan A, Pastor P, Ibarz J, Herzog A, Jang E, Quillen D, Holly E, Kalakrishnan M, Vanhoucke V, et al. (2018) Qt-opt: Scalable deep reinforcement learning for vision-based robotic manipulation. *arXiv preprint arXiv:1806.10293*.
- Kaufmann E, Cappé O, Garivier A (2016) On the complexity of best-arm identification in multi-armed bandit models. *The Journal of Machine Learning Research* 17(1):1–42.
- Kearns M, Singh S (1998) Finite-sample convergence rates for Q-learning and indirect algorithms. *Proceedings of the conference on Advances in neural information processing systems II*, 996–1002.
- Kim SH, Nelson BL (2007) Recent advances in ranking and selection. *2007 Winter Simulation Conference*, 162–172 (IEEE).
- Kiran BR, Sobh I, Talpaert V, Mannion P, Sallab AAA, Yogamani S, Pérez P (2020) Deep reinforcement learning for autonomous driving: A survey. *arXiv preprint arXiv:2002.00444*.
- Laber EB, Lizotte DJ, Qian M, Pelham WE, Murphy SA (2014) Dynamic treatment regimes: Technical challenges and applications. *Electronic journal of statistics* 8(1):1225.

-
- Lam H (2016) Advanced tutorial: Input uncertainty and robust analysis in stochastic simulation. *2016 Winter Simulation Conference (WSC)*, 178–192 (IEEE).
- Lee LH, Pujowidianto NA, Li LW, Chen CH, Yap CM (2012) Approximate simulation budget allocation for selecting the best design in the presence of stochastic constraints. *IEEE Transactions on Automatic Control* 57(11):2940–2945.
- Mak WK, Morton DP, Wood RK (1999) Monte carlo bounding techniques for determining solution quality in stochastic programs. *Operations Research Letters* 24(1-2):47–56.
- Mannor S, Simester D, Sun P, Tsitsiklis JN (2004) Bias and variance in value function estimation. *Proceedings of the 21st International Conference on Machine Learning*, 72 (ACM).
- Mannor S, Simester D, Sun P, Tsitsiklis JN (2007) Bias and variance approximation in value function estimates. *Management Science* 53(2):308–322.
- Munos R, Szepesvári C (2008) Finite-time bounds for fitted value iteration. *Journal of Machine Learning Research* 9(May):815–857.
- Osband I, Russo D, Van Roy B (2013) (More) efficient reinforcement learning via posterior sampling. *Advances in Neural Information Processing Systems*, 3003–3011.
- Peng Y, Chen CH, Fu MC, Hu JQ (2017) Gradient-based myopic allocation policy: An efficient sampling procedure in a low-confidence scenario. *IEEE Transactions on Automatic Control* 63(9):3091–3097.
- Peng Y, Chen CH, Fu MC, Hu JQ, Ryzhov IO (2018a) Efficient sampling allocation procedures for selecting alternative with optimal quantile. *INFORMS Journal on Computing* .
- Peng Y, Chong EK, Chen CH, Fu MC (2018b) Ranking and selection as stochastic control. *IEEE Transactions on Automatic Control* 63(8):2359–2373.
- Puterman ML (2014) *Markov decision processes: discrete stochastic dynamic programming* (John Wiley & Sons).
- Putta SR, Tulabandhula T (2017) Pure exploration in episodic fixed-horizon markov decision processes. *AAMAS*, 1703–1704.
- Russo D (2016) Simple bayesian algorithms for best arm identification. *Conference on Learning Theory*, 1417–1418.
- Ryzhov IO (2018) The local time method for targeting and selection. *Operations Research* 66(5):1406–1422.
- Ryzhov IO, Powell WB, Frazier PI (2012) The knowledge gradient algorithm for a general class of online learning problems. *Operations Research* 60(1):180–195.
- Schöner H (2017) The role of simulation in development and testing of autonomous vehicles. *Driving Simulation Conference, Stuttgart*.
- Serfling RJ (2009) *Approximation Theorems of Mathematical Statistics*, volume 162 (John Wiley & Sons).
- Shah D, Xie Q (2018) Q-learning with nearest neighbors. *Advances in Neural Information Processing Systems*, 3111–3121.
- Shapiro A, Dentcheva D, Ruszczyński A (2014) *Lectures on stochastic programming: modeling and theory* (SIAM).
- Shen H, Hong LJ, Zhang X (2021) Ranking and selection with covariates for personalized decision making. *INFORMS Journal on Computing* .

- Shin D, Broadie M, Zeevi A (2016) Tractable sampling strategies for quantile-based ordinal optimization. *2016 Winter Simulation Conference (WSC)*, 847–858 (IEEE).
- Song E, Nelson BL (2015) Quickly assessing contributions to input uncertainty. *IIE Transactions* 47(9):893–909.
- Song E, Nelson BL, Pegden CD (2014) Advanced tutorial: Input uncertainty quantification. Tolk A, Diallo S, Ryzhov I, Yilmaz L, Buckley S, Miller J, eds., *Proceedings of the 2014 Winter Simulation Conference*, 162–176 (Piscataway, New Jersey: IEEE).
- Strehl AL, Littman ML (2008) An analysis of model-based interval estimation for markov decision processes. *Journal of Computer and System Sciences* 74(8):1309–1331.
- Trevezas S, Limnios N (2009) Variance estimation in the central limit theorem for markov chains. *Journal of Statistical Planning and Inference* 139(7):2242–2253.
- Xie W, Nelson BL, Barton RR (2014) A bayesian framework for quantifying uncertainty in stochastic simulation. *Operations Research* 62(6):1439–1452.
- Yang L, Wang M (2020) Reinforcement learning in feature space: Matrix bandit, kernels, and regret bound. *International Conference on Machine Learning*, 10746–10756 (PMLR).
- Yi Y, Xie W (2017) An efficient budget allocation approach for quantifying the impact of input uncertainty in stochastic simulation. *ACM Transactions on Modeling and Computer Simulation (TOMACS)* 27(4):25.
- Zhu H, Liu T, Zhou E (2020) Risk quantification in stochastic simulation under input uncertainty. *ACM Transactions on Modeling and Computer Simulation (TOMACS)* 30(1):1–24.
- Zouaoui F, Wilson JR (2004) Accounting for input-model and input-parameter uncertainties in simulation. *IIE Transactions* 36(11):1135–1151.