
Lookahead-Bounded Q-Learning

Ibrahim El Shar¹ Daniel R. Jiang^{1,2}

Abstract

We introduce the lookahead-bounded Q-learning (LBQL) algorithm, a new, provably convergent variant of Q-learning that seeks to improve the performance of standard Q-learning in stochastic environments through the use of “lookahead” upper and lower bounds. To do this, LBQL employs previously collected experience and each iteration’s state-action values as dual feasible penalties to construct a sequence of sampled information relaxation problems. The solutions to these problems provide estimated upper and lower bounds on the optimal value, which we track via stochastic approximation. These quantities are then used to constrain the iterates to stay within the bounds at every iteration. Numerical experiments confirm the fast convergence of LBQL as compared to the standard Q-learning algorithm and several related techniques. Our approach is particularly appealing in problems that require expensive simulations or real-world interactions.

1 Introduction

Since its introduction by Watkins (1989), Q-learning has become one of the most widely-used reinforcement learning (RL) algorithms (Sutton & Barto, 2018), due to its conceptual simplicity, ease of implementation, and convergence guarantees (Jaakkola et al., 1994; Tsitsiklis, 1994; Bertsekas & Tsitsiklis, 1996). However, practical implementations of Q-learning on real-world problems is difficult due to the (often) high cost of obtaining data and experience from real environments, along with other issues such as overestimation bias (Szepesvári, 1998; Even-Dar & Mansour, 2003; Lee & Powell, 2019).

In this paper, we address these challenges by focusing on a specific class of problems with *partially known* models in the following sense: writing the transition dynamics as $s_{t+1} = f(s_t, a_t, w_{t+1})$, where s_t and a_t are the current

state and action, s_{t+1} is the next state, and w_{t+1} is random noise, we consider problems where f is known but w_{t+1} can only be observed through interactions with the environment. This type of model is the norm in the control (Bertsekas & Tsitsiklis, 1996) and operations research (Powell, 2007) communities. We propose and analyze a new RL algorithm called *lookahead-bounded Q-learning* (LBQL), which exploits knowledge of f to improve the efficiency of Q-learning and address overestimation bias. It does so by making better use of the observed data through estimating upper and lower bounds using a technique called *information relaxation* (IR) (Brown et al., 2010).

Indeed, there are abundant real-world examples that fall into this subclass of problems, as we now illustrate with a few examples. In *inventory control*, the transition from one inventory state to the next is a well-specified function f given knowledge of a stochastic demand w_{t+1} (Kunnumkal & Topaloglu, 2008). For *vehicle routing*, f is the often just the routing decision itself, while w_{t+1} are exogenous demands that require observation (Secomandi, 2001). In *energy systems operations*, a typical setting is to optimize storage that behaves through linear transitions f together with unpredictable renewable supply, w_{t+1} (Kim & Powell, 2011). In *portfolio optimization*, f is the next portfolio, and w_{t+1} represents random prices (Rogers, 2002). In Section 5 of this paper, we discuss another application that follows this paradigm: *spatial dynamic pricing for car sharing*.

Although we specialize to problems with partially known transition dynamics, this should not be considered restrictive: in fact, our proposed algorithm can be integrated seamlessly with the model-based RL framework to handle the standard model-free RL setting, where f is constantly being learned. We leave this extension to future work.

Main Contributions. We make the following methodological and empirical contributions in this paper.

1. We propose a novel algorithm that takes advantage of IR theory and Q-learning to generate upper and lower bounds on the optimal value. This allows our algorithm to mitigate the effects of maximization bias, while making better use of the collected experience and the transition function f . A variant of the algorithm based on experience replay is also given.
2. We prove that our method converges almost surely to

¹University of Pittsburgh ²Facebook Core Data Science. Correspondence to: Ibrahim El Shar <ije8@pitt.edu>.

the optimal action-value function. The proof requires a careful analysis of several interrelated stochastic processes (upper bounds, lower bounds, and the Q-factors themselves).

3. Numerical experiments on five test problems show superior empirical performance of LBQL compared to Q-learning and other widely-used variants. Interestingly, sensitivity analysis shows that LBQL is more robust to learning rate and exploration parameters.

The rest of the paper is organized as follows. In the next section, we review the related literature. In Section 3, we introduce the notation and review the basic theory of IR. In Section 4, we present our algorithm along with its theoretical results. In Section 5, we show the numerical results where LBQL is compared to other Q-learning variants. Finally, we state our conclusion and future work in Section 6.

2 Related Literature

Upper and lower bounds on the optimal value have recently been used by optimism-based algorithms, e.g., [Dann et al. \(2018\)](#) and [Zanette & Brunskill \(2019\)](#). These papers focus on finite horizon problems, while we consider infinite horizon case. Their primary use of the lower and upper bounds is to achieve better exploration, while our work is focused on improving the action-value estimates by mitigating overestimation and enabling data re-use.

In the context of real-time dynamic programming (RTDP) ([Barto et al., 1995](#)), Bounded RTDP ([McMahan et al., 2005](#)), Bayesian RTDP ([Sanner et al., 2009](#)) and Focused RTDP ([Smith & Simmons, 2006](#)) propose extensions of RTDP where a lower bound heuristic and an upper bound are maintained on the value function. These papers largely use heuristic approaches to obtain bounds, while we use the principled idea of IR duality.

More closely related to our paper is the work of [He et al. \(2016\)](#), which exploits multistep returns to construct bounds on the optimal action-value function, before utilizing constrained optimization to enforce those bounds. However, unlike our work, no theoretical guarantees are provided. To the best of our knowledge, we provide the first asymptotic proof of convergence to the general approach of enforcing dynamically computed (noisy) bounds. Thus, our analysis can serve as a base for deriving theoretical guarantees of similar algorithms.

There are also two papers that utilize IR bounds in the related setting of *finite horizon dynamic programming*. [Jiang et al. \(2020\)](#) use IR dual bounds in a tree search algorithm in order to ignore parts of the tree. Recent work by [Chen et al. \(2019\)](#) uses IR duality in a duality-based dynamic programming algorithm that converges monotonically to the optimal value function through a series of “subsolu-

tions” under more restrictive assumptions (e.g., knowledge of probability distributions).

3 Background

In this section, we first introduce some definitions and concepts from Markov decision process (MDP) theory. Then, we describe the basic theory of information relaxations and duality, which is the main tool used in our LBQL approach.

3.1 MDP Model

Consider a discounted, infinite horizon MDP with a finite state space \mathcal{S} , and a finite action space \mathcal{A} , and a disturbance space \mathcal{W} . Let $\{w_t\}$ be a sequence of independent and identically distributed (i.i.d.) random variables defined on a probability space $(\Omega, \mathcal{F}, \mathbf{P})$, where each w_t is supported on the set \mathcal{W} . Let $s_t \in \mathcal{S}$ be the state of the system at time t . We also define a state transition function $f : \mathcal{S} \times \mathcal{A} \times \mathcal{W} \rightarrow \mathcal{S}$, such that if action a_t is taken at time t , then the next state is governed by $s_{t+1} = f(s_t, a_t, w_{t+1})$. This “transition function” model of the MDP is more convenient for our purposes, but we note that it can easily be converted to the standard model used in RL, where the transition probabilities, $p(s_{t+1} | s_t, a_t)$, are modeled directly. For simplicity and ease of notation, we assume that w is independent¹ from (s, a) . Let $r(s_t, a_t)$ be the expected reward when taking action $a_t \in \mathcal{A}$ in state $s_t \in \mathcal{S}$. We assume that the rewards $r(s_t, a_t)$ are uniformly bounded by R_{\max} and for simplicity in notation, that the feasible action set \mathcal{A} does not depend on the current state. As usual, a deterministic Markov policy $\pi \in \Pi$ is a mapping from states to actions, such that $a_t = \pi(s_t)$ whenever we are following policy π . We let Π be the set of all possible policies (or the set of all “admissible” policies).

Given a discount factor $\gamma \in (0, 1)$ and a policy $\pi \in \Pi$, the *value* and the *action-value* functions are denoted respectively by

$$V^\pi(s) = \mathbf{E} \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) \mid \pi, s_0 = s \right] \quad \text{and}$$

$$Q^\pi(s, a) = \mathbf{E} \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) \mid s_0 = s, a_0 = a, \pi \right],$$

where the notation of “conditioning on π ” refers to actions a_t selected by $\pi(s_t)$. The expectation \mathbf{E} , here and throughout the paper, is taken with respect to \mathbf{P} . Our objective is to find a policy $\pi \in \Pi$ such that starting from any state s , it achieves the optimal expected discounted cumulative reward. The value of this optimal policy π^* for a state s is called the optimal value function and is denoted by $V^*(s) = \max_{\pi} V^\pi(s)$. Specifi-

¹However, we can also allow for (s, a) -dependent w with essentially no fundamental changes to our approach.

cally, it is well-known that an optimal policy selects actions according to $\pi^*(s) = \arg \max_{a \in \mathcal{A}} Q^*(s, a)$, where $Q^*(s, a) = \max_{\pi} Q^{\pi}(s, a)$ is the optimal action-value function (Puterman, 2014). The Bellman optimality equation gives the following recursion:

$$Q^*(s_t, a_t) = r(s_t, a_t) + \gamma \mathbf{E} \left[\max_{a_{t+1}} Q^*(s_{t+1}, a_{t+1}) \right].$$

The goal in many RL algorithms, including Q -learning (Watkins, 1989), is to approximate Q^* .

3.2 Information Relaxation Duality

Now let us give a brief review of the theory behind information relaxation duality from Brown et al. (2010), which is a way of computing an *upper bound* on the value and action-value functions. This generalizes work by Rogers (2002), Haugh & Kogan (2004), and Andersen & Broadie (2004) on pricing American options. Note that any feasible policy provides a lower bound on the optimal value, but computing an upper bound is less straightforward. The *information relaxation* approach proposes to relax the “non-anticipativity” constraints on policies, i.e., it allows them to depend on realizations of future uncertainties when making decisions. Naturally, optimizing in the class of policies that can “see the future” provides an upper bound on the best admissible policy. We focus on the special case of *perfect information relaxation*, where full knowledge of the future uncertainties, i.e., the sample path (w_1, w_2, \dots) , is used to create upper bounds. The naive version of the perfect information bound is simply given by

$$V^*(s_0) \leq \mathbf{E} \left[\max_{\mathbf{a}} \left\{ \sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) \right\} \right],$$

which, in essence, is an interchange of the expectation and max operators; the interpretation here is that an agent who is allowed to adjust her actions *after* the uncertainties are realized achieves higher reward than an agent who acts sequentially. As one might expect, perfect information can provide upper bounds that are quite loose. The central idea of the information relaxation approach to strengthen these upper bounds is to simultaneously (1) allow the use of future information but (2) also penalize the agent for doing so.

A penalty function is said to be *dual feasible* if it does not penalize any admissible policy $\pi \in \Pi$ in expectation. Let $s_{t+1} = f(s_t, a_t, w_{t+1})$ be the next state, $\varphi : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ be a bounded function, and w have the same distribution as w_{t+1} . Then, penalties involving terms of the form

$$\begin{aligned} z_t^{\pi}(s_t, a_t, w_{t+1} | \varphi) &:= \gamma^{t+1} \left(\varphi(s_{t+1}, \pi(s_{t+1})) \right. \\ &\quad \left. - \mathbf{E} \left[\varphi(f(s_t, a_t, w), \pi(f(s_t, a_t, w))) \right] \right) \end{aligned} \quad (1)$$

are dual feasible because

$$\mathbf{E} \left[\sum_{t=0}^{\infty} z_t^{\pi}(s_t, a_t, w_{t+1} | \varphi) \right] = 0.$$

This is a variant of the types of penalties introduced in Brown & Haugh (2017), extended to the case of action-value functions. Intuitively, if φ is understood to be an estimate of the optimal action-value function Q^* and π an estimate of the optimal policy π^* , then z_t^{π} can be thought of as the one-step value of future information (i.e., knowing w_{t+1} versus taking an expectation over its distribution).

These terms, however, may have negative expected value for policies that violate non-anticipativity constraints. Let π_{φ} be the policy that is greedy with respect to the bounded function φ (taken as an action-value function). Consider the problem constructed by subtracting the penalty term from the reward in each period and relaxing non-anticipativity constraints by interchanging maximization and expectation:

$$\begin{aligned} Q^U(s_0, a_0) &= \mathbf{E} \left[\max_{\mathbf{a}} \left\{ \sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) \right. \right. \\ &\quad \left. \left. - z_t^{\pi_{\varphi}}(s_t, a_t, w_{t+1} | \varphi) \right\} \right]. \end{aligned} \quad (2)$$

Brown & Haugh (2017) shows that the objective value of this problem, $Q^U(s_0, a_0)$, is an upper bound on $Q^*(s_0, a_0)$, where $\mathbf{a} := (a_0, a_1, \dots)$ is an infinite sequence of actions. Notice that in principle, it is possible to estimate the problem in (2) using Monte Carlo simulation. To do this, we generate infinitely long sample paths of the form $\mathbf{w} = (w_1, w_2, \dots)$, and for each fixed \mathbf{w} , we solve the inner deterministic dynamic programming (DP) problem. Averaging over the results produces an estimate of the upper bound of $Q^U(s_0, a_0)$. Our new approach to Q -learning will take advantage of this idea, with φ and π being continuously updated.

3.3 Absorption Time Formulation

In practice, however, we cannot simulate infinitely long sample paths \mathbf{w} . One solution is to use an *equivalent* (random) finite horizon formulation (see for e.g. Proposition 5.3.1 of Puterman 2014) where instead of discounting, a new absorbing state \tilde{s} with zero reward is added to the state space \mathcal{S} . This new state \tilde{s} can be reached from every state and for any feasible action with probability $1 - \gamma$. We define a new state transition function h , which transitions to \tilde{s} with probability $1 - \gamma$ from every (s, a) , but conditioned on not absorbing (i.e., with probability γ), h is identical to f . We refer to this as the *absorption time formulation*, where the horizon length $\tau := \min\{t : s_t = \tilde{s}\}$ has a geometric distribution with parameter $1 - \gamma$ and the state transitions are governed by the state transition function h instead of f . Let \mathcal{Q} be the set of bounded functions φ such that $\varphi(\tilde{s}, a) = 0$ for all $a \in \mathcal{A}$. The penalty terms for the absorption time formulation are defined in a similar way as (1), except we now consider $\varphi \in \mathcal{Q}$:

$$\begin{aligned} \zeta_t^{\pi}(s_t, a_t, w_{t+1} | \varphi) &:= \varphi(s_{t+1}, \pi(s_{t+1})) \\ &\quad - \mathbf{E} \left[\varphi(h(s_t, a_t, w), \pi(h(s_t, a_t, w))) \right], \end{aligned} \quad (3)$$

where $s_{t+1} = h(s_t, a_t, w_{t+1})$. We now state a proposition that summarizes the information relaxation duality results, which is a slight variant of results in Proposition 2.2 of Brown & Haugh (2017).

Proposition 1 (Duality Results, Proposition 2.2 in Brown & Haugh (2017)). *The following duality results are stated for the absorption time formulation of the problem.*

(i) *Weak Duality: For any $\pi \in \Pi$ and $\varphi \in \mathcal{Q}$,*

$$Q^\pi(s_0, a_0) \leq \mathbf{E} \left[\max_{\mathbf{a}} \sum_{t=0}^{\tau-1} \left(r(s_t, a_t) - \zeta_t^{\pi, \varphi}(s_t, a_t, w_{t+1} | \varphi) \right) \right] \quad (4)$$

(ii) *Strong Duality: It holds that*

$$Q^*(s_0, a_0) = \inf_{\varphi \in \mathcal{Q}} \mathbf{E} \left[\max_{\mathbf{a}} \sum_{t=0}^{\tau-1} \left(r(s_t, a_t) - \zeta_t^{\pi, \varphi}(s_t, a_t, w_{t+1} | \varphi) \right) \right], \quad (5)$$

with the infimum attained at $\varphi = Q^*$.

The DP inside the expectation of the right hand side of (4) is called the *inner problem*. Weak duality tells us that by using a dual feasible penalty, we can get an estimated upper bound on the optimal action-value function $Q^*(s_0, a_0)$ by simulating multiple sample paths and averaging the optimal value of the resulting inner problems. Strong duality suggests that the information gained from accessing the future is perfectly cancelled out by the optimal dual feasible penalty.

For a given sample path $\mathbf{w} = (w_1, \dots, w_\tau)$, each of the inner DP problems can be solved via the backward recursion

$$Q_t^U(s_t, a_t) = r(s_t, a_t) - \zeta_t^{\pi, \varphi}(s_t, a_t, w_{t+1} | \varphi) + \max_a Q_{t+1}^U(s_{t+1}, a), \quad (6)$$

for $t = \tau - 1, \tau - 2, \dots, 0$ with $s_{t+1} = h(s_t, a_t, w_{t+1})$ and $Q_\tau^U \equiv 0$ (as there is no additional reward after entering the absorbing state \bar{s}). The optimal value of the inner problem is given by $Q_0^U(s_0, a_0)$.

3.4 Lower Bounds using IR

The penalty function approach also allows for using a feasible policy to estimate a *lower bound* on the optimal value, such that when using a common sample path, this lower bound is guaranteed to be less than the corresponding estimated upper bound, a crucial aspect of our theoretical analysis. Specifically, given a sample path $(w_1, w_2, \dots, w_\tau)$, the inner problem used to evaluate a feasible policy $\pi \in \Pi$ is given by

$$Q_t^L(s_t, a_t) = r(s_t, a_t) - \zeta_t^\pi(s_t, a_t, w_{t+1} | \varphi) + Q_{t+1}^L(s_{t+1}, \pi(s_{t+1})), \quad (7)$$

for $t = 0, \dots, \tau - 1$, with $s_{t+1} = h(s_t, a_t, w_{t+1})$ and $Q_\tau^L \equiv 0$. It follows that $\mathbf{E}[Q_0^L(s_0, a_0)] = Q^\pi(s_0, a_0)$, because penalty terms $\zeta_t^\pi(s_t, a_t, w_{t+1} | \varphi)$ have zero mean.

4 QL with Lookahead Bounds

We now introduce our proposed approach, which integrates the machinery of IR duality with Q -learning in a unique way. An outline of the essential steps is given below.

1. On a given iteration, we first experience a realization of the exogenous information w_{t+1} and make a standard Q -learning update.
2. We then set φ to be the newly updated Q -iterate and compute upper and lower bounds on the true Q^* , which are then tracked and averaged using a stochastic approximation step.
3. Finally, we project the Q -iterate so that it satisfies the averaged upper and lower bounds and return to Step 1.

Figure 1 shows an illustration of each of these steps at a given iteration of the algorithm. Since we are setting φ to be the current Q -iterate at every iteration, the information relaxation bounds are computed using a dynamic sequence of penalty functions and averaged together using stochastic approximation. The idea is that as our approximation of Q^* improves, our upper and lower bounds also improve. As the upper and lower bounds improve, the projection step ensures that the Q -iterates also improve. It is this back-and-forth feedback between the two processes that has the potential to yield rapid convergence toward the optimal Q^* .

The primary drawback of our approach is that in the computation of the information relaxation dual bounds, expectations need to be computed. We first show an *idealized* version of the algorithm where these expectations are estimated using unbiased samples of w_{t+1} from a black-box simulator. Later, we relax the need for a black-box simulator and show how our algorithm can be implemented with a replay-buffer. Both versions are analyzed theoretically and convergence results are provided.

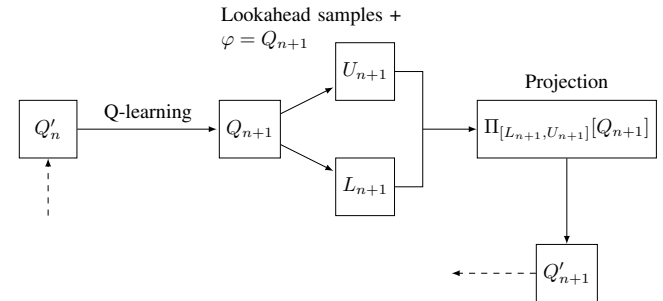


Figure 1: Illustration of LBQL Algorithm at iteration n .

4.1 An Idealized Algorithm

Let $\{w_{t+1}^1, w_{t+1}^2, \dots, w_{t+1}^K\}$ be a *batch* (as opposed to a sample path) of K samples from the distribution of the exogenous information w_{t+1} (i.e., from a black-box simulator).

An empirical version of (3) is simply given by:

$$\hat{\zeta}_t^\pi(s_t, a_t, w_{t+1} | \varphi) := \varphi(s_{t+1}, \pi(s_{t+1})) - \frac{1}{K} \sum_{k=1}^K \varphi(h(s_t, a_t, w_{t+1}^k), \pi(h(s_t, a_t, w_{t+1}^k))), \quad (8)$$

where $s_{t+1} = h(s_t, a_t, w_{t+1})$. Given a sample path $\mathbf{w} = (w_1, w_2, \dots, w_\tau)$ of the absorption time formulation of the problem, analogues to (6) and (7) using $\hat{\zeta}_t^\pi$, where in (7) we set $\pi = \pi_\varphi$ (i.e., the lower bound on the optimal value is constructed by approximately evaluating the feasible policy π_φ) are given by

$$\hat{Q}_t^U(s_t, a_t) = r(s_t, a_t) - \hat{\zeta}_t^{\pi_\varphi}(s_t, a_t, w_{t+1} | \varphi) + \max_a \hat{Q}_{t+1}^U(s_{t+1}, a) \quad (9)$$

$$\hat{Q}_t^L(s_t, a_t) = r(s_t, a_t) - \hat{\zeta}_t^{\pi_\varphi}(s_t, a_t, w_{t+1} | \varphi) + \hat{Q}_{t+1}^L(s_{t+1}, \pi_\varphi(s_{t+1})) \quad (10)$$

for $t = 0, 1, \dots, \tau - 1$, where $s_{t+1} = h(s_t, a_t, w_{t+1})$, $\hat{Q}_\tau^U \equiv \hat{Q}_\tau^L \equiv 0$, and we assume that each call to $\hat{\zeta}_t^\pi$ uses a fresh batch of K samples.

Proposition 2. *The valid upper and lower bound properties continue to hold in the empirical case:*

$$\mathbf{E}[\hat{Q}_0^L(s, a)] \leq Q^*(s, a) \leq \mathbf{E}[\hat{Q}_0^U(s, a)],$$

for any state-action pair (s, a) .

We include the proof in Appendix A.1. The proof is similar to that of Proposition 2.3(iv) of Brown et al. (2010), except extended to the infinite horizon setting with the absorption time formulation. A detailed description of the LBQL algorithm is given in Algorithm 1, where we use ‘ n ’ for the iteration index in order to avoid confusion with the ‘ t ’ used in the inner DP problems. We use $\Pi_{[a,b]}[x]$ to denote x projected onto $[a, b]$, i.e., $\Pi_{[a,b]}[x] = \max\{\min\{x, b\}, a\}$, where either a or b could be infinity. Let $\rho = R_{\max}/(1 - \gamma)$, the initial lower and upper bounds estimates are set such that $L_0(s, a) = -\rho$ and $U_0(s, a) = \rho$ for all $(s, a) \in \mathcal{S} \times \mathcal{A}$. The initial action-value Q_0 is set arbitrarily such that $L_0(s, a) \leq Q_0(s, a) \leq U_0(s, a)$ for all $(s, a) \in \mathcal{S} \times \mathcal{A}$.




Figure 2: A simple stochastic MDP.

Example 1. We demonstrate the idealized LBQL algorithm using the simple MDP shown in Figure 2. The MDP has two non-terminal states A and B. Each episode starts in state A, with a choice of two actions: right and left

Algorithm 1 Lookahead-Bounded Q-Learning

Input: Initial estimates $L_0 \leq Q_0 \leq U_0$, batch size K , and stepsize rules $\alpha_n(s, a)$, $\beta_n(s, a)$.

Output: Approximations $\{L_n\}$, $\{Q'_n\}$, and $\{U_n\}$.

Set $Q'_0 = Q_0$ and choose an initial state s_0 .

for $n = 0, 1, 2, \dots$ **do**

Choose an action a_n via some behavior policy (e.g., ϵ -greedy). Observe w_{n+1} . Let

$$Q_{n+1}(s_n, a_n) = Q'_n(s_n, a_n) + \alpha_n(s_n, a_n) \left[r_n(s_n, a_n) + \gamma \max_a Q'_n(s_{n+1}, a) - Q'_n(s_n, a_n) \right]. \quad (11)$$

Set $\varphi = Q_{n+1}$. Using one sample path \mathbf{w} , compute $\hat{Q}_0^U(s_n, a_n)$ and $\hat{Q}_0^L(s_n, a_n)$ using (9) and (10).

Update and enforce upper and lower bounds:

$$U_{n+1}(s_n, a_n) = \Pi_{[-\rho, \infty]} \left[U_n(s_n, a_n) + \beta_n(s_n, a_n) \left[\hat{Q}_0^U(s_n, a_n) - U_n(s_n, a_n) \right] \right], \quad (12)$$

$$L_{n+1}(s_n, a_n) = \Pi_{[\infty, \rho]} \left[L_n(s_n, a_n) + \beta_n(s_n, a_n) \left[\hat{Q}_0^L(s_n, a_n) - L_n(s_n, a_n) \right] \right], \quad (13)$$

$$Q'_{n+1}(s_n, a_n) = \Pi_{[L_{n+1}(s_n, a_n), U_{n+1}(s_n, a_n)]} [Q_{n+1}(s_n, a_n)]. \quad (14)$$

end for

denoted by R and L respectively. The rewards and transition probabilities of taking an action in each state are shown on the edges in the figure. Assume that the transitions are governed by the outcome of a fair coin. If the outcome is Head then we transition in the direction of our chosen action and in the opposite direction for a Tail outcome. For a discount factor $\gamma = 0.95$, the optimal policy is to go right at both A and B. The optimal action-values are given by $Q^*(A, R) = Q^*(B, R) = 0$, $Q^*(A, L) = Q^*(B, L) = -1$. Consider applying the idealized version of LBQL described in Algorithm 1.




Figure 3: An illustration of LBQL iterates for Example 1.

We let $\alpha_n = 0.1$, $\beta_n = 0.05$ for all n . Figure 3 illustrates two iterations from the first and the 58th episodes. Initially $Q_0(s, a) = 0$ and $\rho = 20$. After one episode the bounds are still loose so we have $Q_1(A, R) = Q'_1(A, R) = -0.1$. At episode 58 (281 iterations): learning has occurred for the

lower and upper bounds values for the right action at \mathbb{A} and \mathbb{B} . We see that the bounds are effective already in keeping the Q -iterate close to Q^* . Interestingly, the upper bound is enforced at \mathbb{A} and the lower bound is enforced at \mathbb{B} . Note that these are the results of a real simulation.

4.2 Analysis of Convergence

In this section, we analyze the convergence of the idealized version of the LBQL algorithm to the optimal action-value function Q^* . We start by summarizing and developing some important technical results that will be used in our analysis. All proofs are presented in Appendix A.

The following proposition establishes the boundedness of the action-value iterates and asymptotic bounds on the L_n and U_n iterates of Algorithm 1, which are needed in our proof of convergence. The proof of this proposition is presented in Section A.2 in the Appendix.

Proposition 3 (Boundedness). *For all $(s, a) \in \mathcal{S} \times \mathcal{A}$, we have the following:*

- (i) *The iterates $Q_n(s, a)$ and $Q'_n(s, a)$, remains bounded for all $(s, a) \in \mathcal{S} \times \mathcal{A}$ and for all n .*
- (ii) *For every $\eta > 0$, and with probability one, there exists some finite iteration index n_0 such that $L_n(s, a) \leq Q^*(s, a) + \eta$ and $Q^*(s, a) - \eta \leq U_n(s, a)$ for all iterations $n \geq n_0$.*

Proposition 3(i) ensures that at each iteration n the action-value iterates Q_n and Q'_n are bounded. This allows us to set $\varphi = Q_{n+1}$ at each iteration of Algorithm 1 and is required to establish convergence in general. The proof is based on establishing an inductive relationship that connects Q_n and Q'_n to the previous lower and upper bounds iterates. Specifically, we show that both action-value iterates are bounded below by the preceding upper bounded iterates and above by the preceding lower bound iterates. Proposition 3(ii) ensures that there exists a finite iteration after which the lower and upper bound iterates L_n and U_n are lower and upper bounds on the optimal action-value function Q^* with an error margin of at most an arbitrary amount $\eta > 0$. In the proof of Proposition 3(ii), we bound the lower and upper bound iterates by a noise process and another sequence that converges to Q^* . We show that the noise process possesses some properties that help to eliminate the effect of the noise asymptotically. With the effects of the noise terms vanishing, the boundedness of the lower and upper bound iterates by Q^* is achieved. Examining the update equations (12) and (13) for U_{n+1} and L_{n+1} in Algorithm 1, we remark that they are not “standard” stochastic approximation or stochastic gradient updates because \hat{Q}_0^U and \hat{Q}_0^L are computed with iteration-dependent penalty functions generated by $\varphi = Q_{n+1}$. In other words, the noiseless function itself is changing over time. The proof of Proposition 3(ii) essentially uses the fact that even though these updates are being

performed with respect to different underlying functions, as long as we can apply Proposition 2 in every case, then after the noise is accounted for, the averaged values U_{n+1} and L_{n+1} are eventually bounded below and above by Q^* , respectively. The following lemma derives some guarantees on the lower and upper bound iterates of Algorithm 1, whose proof appears in Section A.3 of the Appendix.

Lemma 1 (Consistency of Bounds). *If $L_0(s, a) \leq U_0(s, a)$, then $L_n(s, a) \leq U_n(s, a)$ for all iterations n and for all $(s, a) \in \mathcal{S} \times \mathcal{A}$.*

In particular, Lemma 1 shows that the upper and lower bound iterates do not interchange roles and become inconsistent. This is an important property; otherwise, the projection step of Algorithm 1 loses its meaning and would require additional logic to handle inconsistent bounds. The results of Lemma 1 follows mainly by the fact that we are using the same sample path to solve the upper and lower bound inner problems, (9) and (10), respectively. Before stating our convergence results, we first state a typical assumption on the stepsizes and the state visits.

Assumption 1. We assume that:

- (i) $\sum_{n=0}^{\infty} \alpha_n(s, a) = \infty$, $\sum_{n=0}^{\infty} \alpha_n^2(s, a) < \infty$,
 $\sum_{n=0}^{\infty} \beta_n(s, a) = \infty$, $\sum_{n=0}^{\infty} \beta_n^2(s, a) < \infty$,
- (ii) Each state $s \in \mathcal{S}$ is visited infinitely often with probability one.

We now state one of our main theoretical results.

Theorem 1 (Convergence of LBQL). *Under Assumption 1, the following hold with probability 1:*

- (i) *$Q'_n(s, a)$ in Algorithm 1 converges to the optimal action-value function $Q^*(s, a)$ for all state-action pairs (s, a) .*
- (ii) *If the penalty terms are computed exactly, i.e. as per (3), then the iterates $L_n(s, a)$, $Q'_n(s, a)$, $U_n(s, a)$ in Algorithm 1 converge to the optimal action-value function $Q^*(s, a)$ for all state-action pairs (s, a) .*

The primary challenge in the analysis is to handle the interdependent feedback between Q , U , and L , which is *unique to our algorithm* (it is not immediately obvious that the proposed scheme does not diverge).

4.3 LBQL with Experience Replay

We now introduce a more practical version of LBQL that uses experience replay in lieu of a black-box simulator. Here, we use a noise buffer \mathcal{B} to record the unique noise values w that are observed at every iteration. We further assume that the noise space \mathcal{W} is finite, a reasonable assumption for a finite MDP. The buffer \mathcal{B} is used to in two ways: (1) to generate the sample path \mathbf{w} and (2) to estimate the expectation in the penalty function. Here, we track and update the

distribution of the noise w after every iteration as opposed to using a batch of size K . Note that this is equivalent to using a sample of batch size equal to the current size of the buffer whenever the bounds are updated if we were recording the noise observations at each iteration. To illustrate how this can be done, suppose $\mathcal{W} = \{w_a, w_b, w_c, w_d\}$ and that at iteration n we observe $w_{n+1} = w_a$. Let p_a denote the probability of observing w_a , and $N_n(w_a)$ the number of times w_a is observed in the first n iterations, then the empirical estimate of p_a is given by $\hat{p}_n(w_a) = N_n(w_a)/n$.² We denote by $\hat{\mathbf{E}}_n[\cdot]$ the expectation computed using the empirical distribution \hat{p}_n . To differentiate the penalty and the action-values (solutions to the inner problems used in updating both lower and upper bounds) that use noise sampled from the buffer from those defined in the Idealized-version of the algorithm, we define the following reciprocal versions of each, respectively:

$$\begin{aligned} \tilde{\zeta}_t^\pi(s_t, a_t, w | \varphi) &:= \varphi(s_{t+1}, \pi(s_{t+1})) \\ &- \hat{\mathbf{E}}_n[\varphi(h(s_t, a_t, w), \pi(h(s_t, a_t, w)))], \end{aligned} \quad (15)$$

and given a sample path $\mathbf{w} = (w_1, w_2, \dots, w_\tau)$ the inner problems analogous to (9) and (10) are given by

$$\begin{aligned} \tilde{Q}_t^U(s_t, a_t) &= r(s_t, a_t) - \tilde{\zeta}_t^{\pi\varphi}(s_t, a_t, w_{t+1} | \varphi) \\ &+ \max_a \tilde{Q}_{t+1}^U(s_{t+1}, a) \end{aligned} \quad (16)$$

$$\begin{aligned} \tilde{Q}_t^L(s_t, a_t) &= r(s_t, a_t) - \tilde{\zeta}_t^{\pi\varphi}(s_t, a_t, w_{t+1} | \varphi) \\ &+ \tilde{Q}_{t+1}^L(s_{t+1}, \pi_\varphi(s_{t+1})) \end{aligned} \quad (17)$$

for $t = 0, 1, \dots, \tau - 1$, where $s_{t+1} = h(s_t, a_t, w_{t+1})$ and $\tilde{Q}_\tau^U \equiv \tilde{Q}_\tau^L \equiv 0$. The pseudo-code of LBQL with experience replay is shown in Algorithm 2 in Appendix B.

4.4 Convergence of LBQL with Experience Replay

In this section, we prove that the version of LBQL with experience replay also converges to the optimal action-value function. We start by stating a lemma that confirms Proposition 3 and Lemma 1 still hold when the penalty terms are computed using (15).

Lemma 2. *If at any iteration n , the penalty terms are computed using the estimate distribution \hat{p}_n , i.e., as per (15), then Proposition 3 and Lemma 1 still hold.*

Theorem 2 (Convergence of LBQL with experience replay). *Under Assumption 1, the following hold with probability 1:*

- (i) $Q'_n(s, a)$ in Algorithm 2 converges to the optimal action-value function $Q^*(s, a)$ for all state-action pairs (s, a) .

²Note that LBQL could, in principle, be adapted to the case of continuous noise (i.e., where w is continuous random variable) using methods like kernel density estimation (KDE).

- (ii) The iterates $L_n(s, a), Q'_n(s, a), U_n(s, a)$ in Algorithm 2 converge to the optimal action-value function $Q^*(s, a)$ for all state-action pairs (s, a) .

The proof is similar to that of Theorem 1, but using the observations collected in the buffer naturally results in an additional bias term in our analysis. The proof of Lemma 2 shows that as we learn the distribution of the noise, this bias term goes to zero and our original analysis in the unbiased case continues to hold.

Notice that the results in part (ii) of the theorem are, in a sense, stronger than that of Theorem 1(ii). While both achieve asymptotic convergence of the lower and upper bounds to the optimal action-value function, Theorem 2(ii) does not require computing the penalty with the true distribution, i.e., using (3). This is because in the experience replay version, the distribution of the noise random variables is also learned.

5 Numerical Experiments

In our numerical experiments we make slight modifications to Algorithm 2 which help reducing its computational requirements. A detailed description of all changes is included in Appendix C. We also open-source a Python package³ for LBQL that reproduces all experiments and figures presented in this paper. We compare LBQL with experience replay with several algorithms: Q-learning (QL), double Q-learning (Double-QL), speedy Q-learning (SQL), and bias-corrected Q-learning (BCQL) (van Hasselt, 2010; Azar et al., 2011; Lee & Powell, 2019). The environments we consider are summarized below. Detailed description of the environments, the parameters used for the five algorithms, and sensitivity analysis are deferred to Appendix D.

Windy Gridworld (WG). This is a well-known variant of the standard gridworld problem discussed in Sutton & Barto (2018). There is an *upward wind* with a random intensity. The agent moves extra steps in the wind direction whenever it reaches an affected square. The reward is -1 until the goal state is reached, and the reward is 0 thereafter.

Stormy Gridworld (SG). We then consider a new domain that adds the additional complexity of rain and *multi-directional* wind to windy gridworld. The location of the rain is random and when it occurs, puddles that provide negative rewards are created. The reward is similar to that of WG, except that puddle states provide a reward of -10 .

Repositioning in Two-Location Car-sharing (2-CS-R). Our next benchmark is a synthetic problem of balancing inventory of cars by repositioning in a car-sharing platform with two stations, (He et al., 2019). The actions are to decide on the number of cars to be repositioned from one station

³https://github.com/ibrahim-elshar/LBQL_ICML2020.

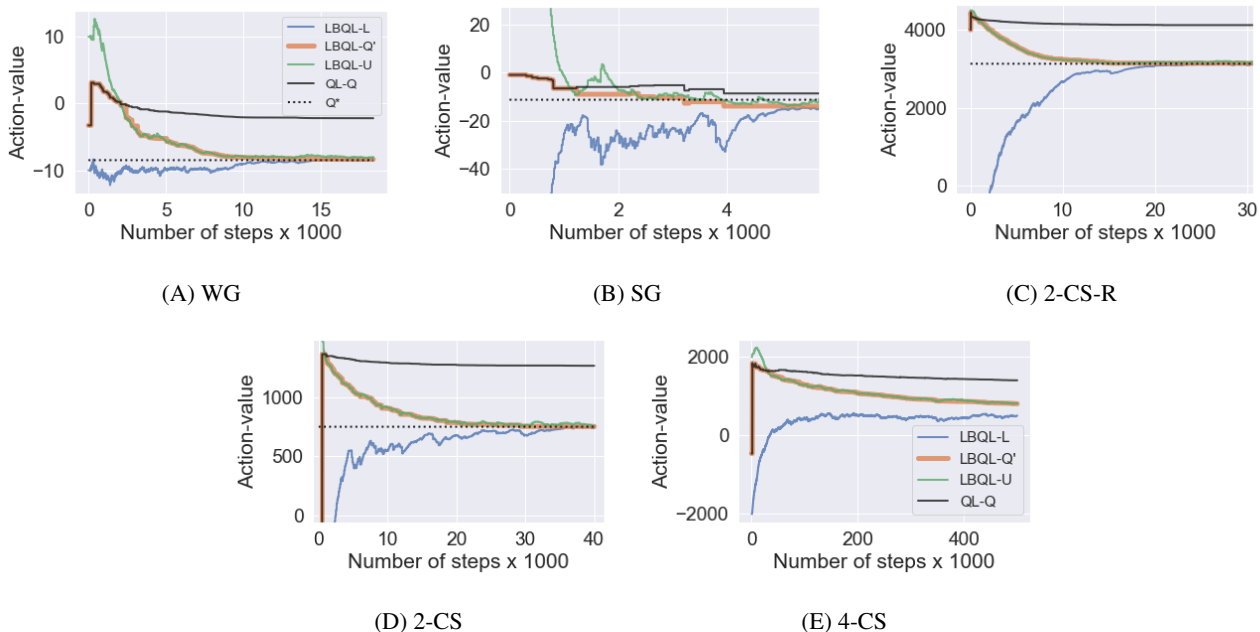


Figure 4: Illustration of LBQL Upper and Lower Bounds.

to the other before random demand is realized. All rentals are one-way (i.e., rentals from station A end up at B, and vice-versa). The goal is to maximize revenue for a fixed rental price subject to lost sales and repositioning costs.

Pricing in Two-Location Car-sharing (2-CS). Here, we consider the benchmark problem of spatial dynamic pricing on a car-sharing platform with two stations, motivated partially by (Bimpikis et al., 2019). The actions are to set a price at each station, which influence the station’s (stochastic) demand for rentals. Rentals are one-way and the goal is to maximize revenue under lost sales cost.

Pricing in Four-Location Car-sharing (4-CS). The final benchmark that we consider is a variant of the above pricing problem with four stations. Now, however, we consider both one way and return trips at each station. In this case, we have two sources of randomness: the noise due to stochastic demand and the noise due to the random distribution of fulfilled rentals between the stations.

First, we illustrate conceptually in Figure 4 how the upper and lower bounds of LBQL can “squeeze” the Q-learning results toward the optimal value (the plots show a particular state-action pair (s, a) for illustrative reasons). For example, in Figure 4A, we observe that the LBQL iterates (orange) match the Q-learning iterates (solid black) initially, but as the upper bound (green) becomes better estimated, the LBQL iterates are pushed toward the optimal value (dotted black). We see that even though the same hyperparameters are used between LBQL and QL, the new approach is able to quickly converge. In the 4-CS example, Figure 4E, Q^*

is not shown since it is computationally difficult to obtain, but the gap between the upper and lower bounds, along with Theorem 2(ii), suggest that LBQL is converging faster than standard Q-learning.

The full results (with 95% confidence intervals) of the numerical experiments are shown in Figure 5. LBQL drastically outperforms the other algorithms in terms of the performance curve on the gridworld domains, but for the car-sharing problems, double Q-learning is superior in the first 20,000 steps. Afterwards, LBQL catches up and remains the best performing algorithm. From the relative error plots (which measure the percent error, in l_2 -norm, of the approximate value function with the true optimal value function, i.e., $\|V_n - V^*\|_2 / \|V^*\|_2$), we see that LBQL has the steepest initial decline. In windy gridworld and car-sharing, LBQL outperforms the rest in terms of relative error, but BCQL and SQL achieve slightly lower relative error than LBQL for stormy gridworld.

We also conducted a set of sensitivity analysis experiments, where we varied the learning rate and exploration hyperparameters across all algorithms (results in Appendix D.3). We examine the number of iterations and CPU time needed to reach 50%, 20%, 5%, and 1% relative error. The results show that LBQL outperforms BCQL, SQL, and QL in terms of both iterations and CPU time in reaching 20%, 5%, and 1% relative error across *all 15 hyperparameter settings* that we tested. For the case of 50% relative error, BCQL outperforms LBQL in five out of 15 cases. This indicates that LBQL is significantly more robust to hyperparameter

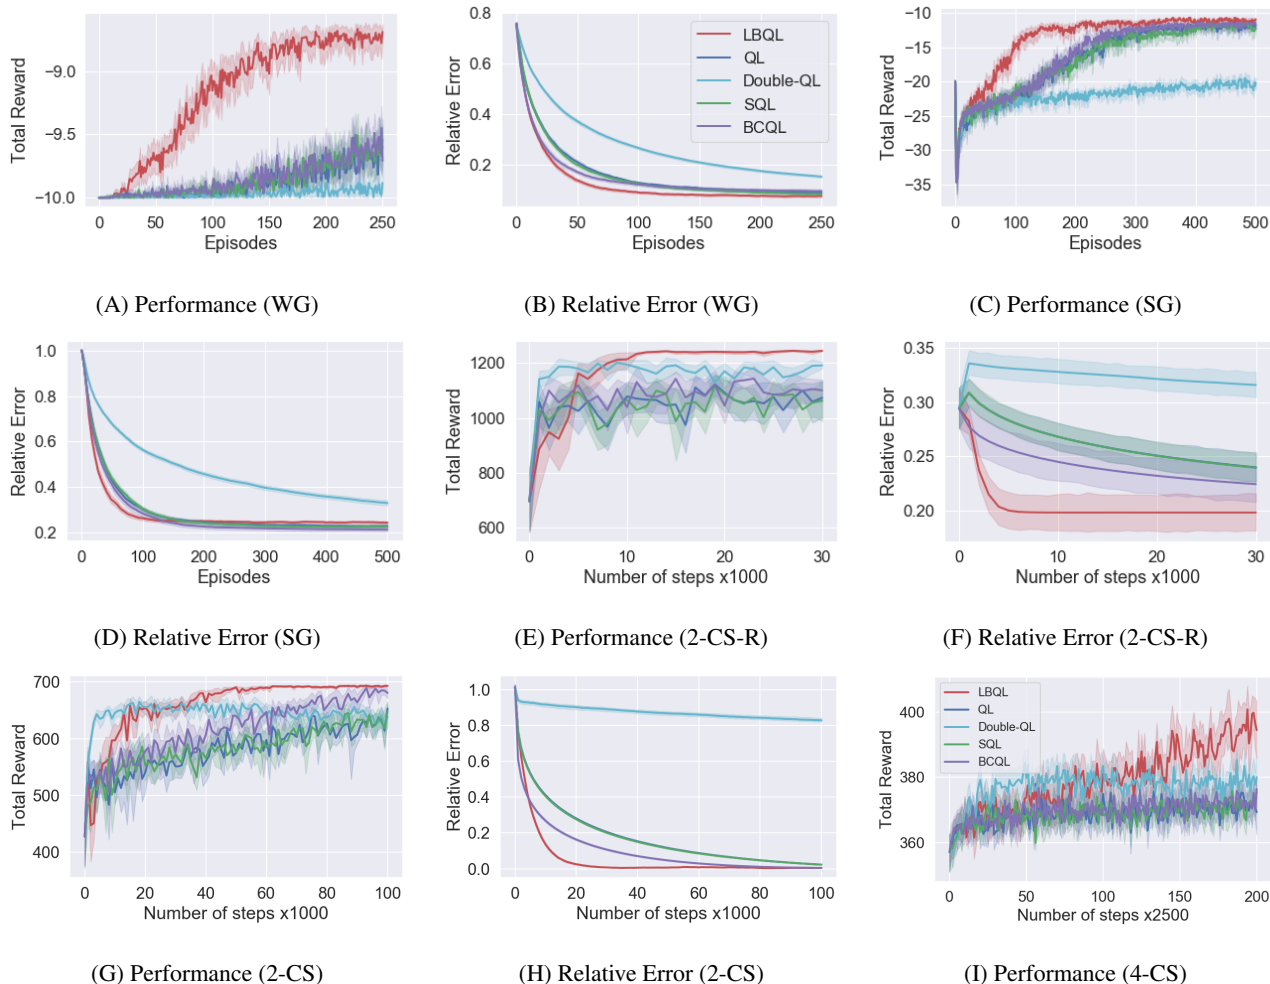


Figure 5: Results from Numerical Experiments.

settings than the other algorithms. Roughly speaking, this robustness might be attributed to the planning aspect of the algorithm where lower and upper bounds are computed.

6 Conclusion

In this paper, we present LBQL algorithm and prove its almost sure convergence to the optimal action-value function. We also propose a practical extension of the algorithm that uses an experience replay buffer. Numerical results illustrate the rapid convergence of our algorithm empirically when compared to a number of well-known variants of Q-learning on five test problems. LBQL is shown to have superior performance, robustness against learning rate and exploration strategies, and an ability to mitigate maximization bias.

Interesting future work is the extension of our new framework to the model-based RL setting, where the transition function f is learned while the policy is optimized. Other interesting future work includes looking beyond the tabu-

lar case and adapting our algorithm to the setting of value function approximations, such as DQN (Mnih et al., 2013).

References

- Andersen, L. and Broadie, M. Primal-dual simulation algorithm for pricing multidimensional American options. *Management Science*, 50(9):1222–1234, 2004.
- Azar, M. G., Munos, R., Ghavamzadeh, M., and Kappen, H. J. Speedy Q-learning. In *Advances in Neural Information Processing Systems 24*, 2011.
- Barto, A. G., Bradtke, S. J., and Singh, S. P. Learning to act using real-time dynamic programming. *Artificial intelligence*, 72(1-2):81–138, 1995.
- Bertsekas, D. P. and Tsitsiklis, J. N. *Neuro-dynamic programming*, volume 5. Athena Scientific Belmont, MA, 1996.

- Bimpikis, K., Candogan, O., and Saban, D. Spatial pricing in ride-sharing networks. *Operations Research*, 2019.
- Brown, D. B. and Haugh, M. B. Information relaxation bounds for infinite horizon markov decision processes. *Operations Research*, 65(5):1355–1379, 2017.
- Brown, D. B., Smith, J. E., and Sun, P. Information relaxations and duality in stochastic dynamic programs. *Operations Research*, 58(4-part-1):785–801, 2010.
- Chen, N., Ma, X., and Liu, Y. Information relaxation and the duality-based dynamic programming (DDP). *Available at SSRN 3416994*, 2019.
- Dann, C., Li, L., Wei, W., and Brunskill, E. Policy certificates: Towards accountable reinforcement learning. *arXiv preprint arXiv:1811.03056*, 2018.
- Even-Dar, E. and Mansour, Y. Learning rates for Q-learning. *Journal of Machine Learning Research*, 5(Dec):1–25, 2003.
- Haugh, M. B. and Kogan, L. Pricing American options: a duality approach. *Operations Research*, 52(2):258–270, 2004.
- He, F. S., Liu, Y., Schwing, A. G., and Peng, J. Learning to play in a day: Faster deep reinforcement learning by optimality tightening. *arXiv preprint arXiv:1611.01606*, 2016.
- He, L., Hu, Z., and Zhang, M. Robust repositioning for vehicle sharing. *Manufacturing & Service Operations Management*, 2019.
- Jaakkola, T., Jordan, M. I., and Singh, S. P. Convergence of stochastic iterative dynamic programming algorithms. In *Advances in Neural Information Processing Systems*, pp. 703–710, 1994.
- Jiang, D. R., Al-Kanj, L., and Powell, W. B. Optimistic Monte Carlo tree search with sampled information relaxation dual bounds. *Operations Research (forthcoming)*, 2020.
- Kim, J. H. and Powell, W. B. Optimal energy commitments with storage and intermittent supply. *Operations Research*, 59(6):1347–1360, 2011.
- Kunnumkal, S. and Topaloglu, H. Using stochastic approximation methods to compute optimal base-stock levels in inventory control problems. *Operations Research*, 56(3):646–664, 2008.
- Kushner, H. and Yin, G. G. *Stochastic approximation and recursive algorithms and applications*, volume 35. Springer Science & Business Media, 2003.
- Lee, D. and Powell, W. B. Bias-corrected Q-learning with multistate extension. *IEEE Transactions on Automatic Control*, 2019.
- McMahan, H. B., Likhachev, M., and Gordon, G. J. Bounded real-time dynamic programming: RTDP with monotone upper bounds and performance guarantees. In *Proceedings of the 22nd international conference on Machine learning*, pp. 569–576. ACM, 2005.
- Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., and Riedmiller, M. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.
- Powell, W. B. *Approximate Dynamic Programming: Solving the curses of dimensionality*, volume 703. John Wiley & Sons, 2007.
- Puterman, M. L. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, 2014.
- Rogers, L. C. G. Monte Carlo valuation of American options. *Mathematical Finance*, 12(3):271–286, 2002.
- Sanner, S., Goetschalckx, R., Driessens, K., and Shani, G. Bayesian real-time dynamic programming. In *Twenty-First International Joint Conference on Artificial Intelligence*, 2009.
- Secomandi, N. A rollout policy for the vehicle routing problem with stochastic demands. *Operations Research*, 49(5):796–802, 2001.
- Smith, T. and Simmons, R. Focused real-time dynamic programming for mdps: Squeezing more out of a heuristic. In *AAAI*, pp. 1227–1232, 2006.
- Sutton, R. S. and Barto, A. G. *Reinforcement Learning: An Introduction*. MIT press, 2018.
- Szepesvári, C. The asymptotic convergence-rate of Q-learning. In *Advances in Neural Information Processing Systems*, pp. 1064–1070, 1998.
- Tsitsiklis, J. N. Asynchronous stochastic approximation and Q-learning. *Machine Learning*, 16(3):185–202, 1994.
- van Hasselt, H. Double Q-learning. In *Advances in Neural Information Processing Systems*, pp. 2613–2621, 2010.
- Watkins, C. J. C. H. *Learning from Delayed Rewards*. PhD thesis, King’s College, Cambridge, UK, 1989.
- Zanette, A. and Brunskill, E. Tighter problem-dependent regret bounds in reinforcement learning without domain knowledge using value function bounds. *arXiv preprint arXiv:1901.00210*, 2019.

Appendix to *Lookahead-Bounded Q-Learning*
A Proofs
A.1 Proof of Proposition 2

Proof. We provide a proof that is similar to that of Proposition 2.3 (iv) of (Brown et al., 2010) but for the case of the absorption time formulation of an infinite horizon problem. Here, we define a policy $\pi := \{\pi_t\}_{t \geq 0}$ as a sequence of functions, that maps from $\{w_t\}_{t \geq 1}$ to feasible actions. We may also use stationary policies where π_t is the same for all t and only depends on the current state s_t . Let $\mathbb{G} = \{\mathcal{G}_t\}_{t \geq 0}$ be the perfect information relaxation of the natural filtration $\mathbb{F} = \{\mathcal{F}_t\}_{t \geq 0}$. Under \mathbb{G} , we have $\mathcal{G}_t = \mathcal{F}$, i.e., we have access to the entire future uncertainties at each t . Define by $\Pi_{\mathbb{G}}$ the set of policies that includes the policies that have access to future uncertainties in addition to nonanticipative policies. Let $\hat{\mathbb{G}}$ be a relaxation of \mathbb{G} such that in addition to what is known under \mathbb{G} the estimate penalty terms $\hat{\zeta}_t^{\pi, \varphi}(s_t, a_t, w_{t+1} | \varphi)$ are revealed at time t .

We first prove $\mathbf{E}[\hat{Q}_0^L(s, a)] \leq Q^*(s, a)$. For an admissible policy π , we have

$$\begin{aligned}
 \mathbf{E}[\hat{Q}_0^L(s, a)] &\stackrel{(a)}{=} \mathbf{E} \left[\sum_{t=0}^{\tau-1} r(s_t, \pi(s_t)) - \hat{\zeta}_t^{\pi}(s_t, a_t, w_{t+1} | \varphi) \mid s_0 = s, a_0 = a \right] \\
 &\stackrel{(b)}{=} \mathbf{E} \left[\sum_{t=0}^{\tau-1} (r(s_t, \pi(s_t)) - \hat{\zeta}_t^{\pi}(s_t, a_t, w_{t+1} | \varphi)) \mathbb{1}_{\{\tau < \infty\}} \mid s_0 = s, a_0 = a \right] \\
 &\stackrel{(c)}{=} \sum_{\tau'=1}^{\infty} \mathbf{E} \left[\sum_{t=0}^{\tau'-1} (r(s_t, \pi(s_t)) - \hat{\zeta}_t^{\pi}(s_t, a_t, w_{t+1} | \varphi)) \mathbb{1}_{\{\tau = \tau'\}} \mid s_0 = s, a_0 = a \right] \\
 &\stackrel{(d)}{=} \sum_{\tau'=1}^{\infty} \mathbf{E} \left[\sum_{t=0}^{\tau'-1} (r(s_t, \pi(s_t)) - \mathbf{E}[\hat{\zeta}_t^{\pi}(s_t, a_t, w_{t+1} | \varphi) \mid \mathcal{G}_t]) \mathbb{1}_{\{\tau = \tau'\}} \mid s_0 = s, a_0 = a \right] \\
 &\stackrel{(e)}{=} \sum_{\tau'=1}^{\infty} \mathbf{E} \left[\sum_{t=0}^{\tau'-1} (r(s_t, \pi(s_t)) - \zeta_t^{\pi}(s_t, a_t, w_{t+1} | \varphi)) \mathbb{1}_{\{\tau = \tau'\}} \mid s_0 = s, a_0 = a \right] \\
 &\stackrel{(f)}{=} \mathbf{E} \left[\sum_{t=0}^{\tau-1} (r(s_t, \pi(s_t)) - \zeta_t^{\pi}(s_t, a_t, w_{t+1} | \varphi)) \mathbb{1}_{\{\tau < \infty\}} \mid s_0 = s, a_0 = a \right] \\
 &\stackrel{(g)}{=} \mathbf{E} \left[\sum_{t=0}^{\tau-1} (r(s_t, \pi(s_t)) - \zeta_t^{\pi}(s_t, a_t, w_{t+1} | \varphi)) \mid s_0 = s, a_0 = a \right] \\
 &\leq Q^*(s, a)
 \end{aligned}$$

Equality (a) follows from the definition of $\hat{Q}_0^L(s, a)$. Equality (b) follows since τ has finite mean and r and φ are uniformly bounded. Equalities (c) and (d) follow from the law of total expectations. Equality (e) follows from Lemma A.1 in (Brown et al., 2010) and from the estimated penalty terms being unbiased, i.e., $\mathbf{E}[\hat{\zeta}_t^{\pi, \varphi}(s_t, a_t, w_{t+1} | \varphi) \mid \mathcal{G}_t] = \zeta_t^{\pi, \varphi}(s_t, a_t, w_{t+1} | \varphi)$. Equalities (f) and (g) follow by the law of total expectation and τ being almost surely finite stopping time, respectively. The inequality follows since the expected value of the penalty terms for a feasible policy is zero and the action-value function of a feasible policy, $Q^{\pi}(s, a)$, is less than $Q^*(s, a)$.

Now, we prove $Q^*(s, a) \leq \mathbf{E}[\hat{Q}_0^U(s, a)]$. Let π_G^* be the optimal solution for the dual problem,

$$\max_{\pi_G \in \Pi_{\mathbb{G}}} \mathbf{E} \left[\sum_{t=0}^{\tau-1} (r(s_t, \pi_G) - \zeta_t^{\pi_G}(s_t, a_t, w_{t+1} | \varphi)) \mid s_0 = s, a_0 = a \right]. \quad (18)$$

We have,

$$\begin{aligned}
 \mathbf{E}[\hat{Q}_0^U(s, a)] &\stackrel{(a)}{=} \mathbf{E} \left[\max_{\mathbf{a}} \left\{ \sum_{t=0}^{\tau-1} r(s_t, a_t) - \hat{\zeta}_t^{\pi^\varphi}(s_t, a_t, w_{t+1} | \varphi) \right\} \mid s_0 = s, a_0 = a \right] \\
 &\stackrel{(b)}{=} \max_{\pi_G \in \Pi_{\hat{\mathbb{G}}}} \mathbf{E} \left[\sum_{t=0}^{\tau-1} (r(s_t, \pi_G) - \hat{\zeta}_t^{\pi_G}(s_t, a_t, w_{t+1} | \varphi)) \mid s_0 = s, a_0 = a \right] \\
 &\geq \mathbf{E} \left[\sum_{t=0}^{\tau-1} (r(s_t, \pi_G^*) - \hat{\zeta}_t^{\pi_G^*}(s_t, a_t, w_{t+1} | \varphi)) \mid s_0 = s, a_0 = a \right] \\
 &\stackrel{(c)}{=} \mathbf{E} \left[\sum_{t=0}^{\tau-1} (r(s_t, \pi_G^*) - \hat{\zeta}_t^{\pi_G^*}(s_t, a_t, w_{t+1} | \varphi)) \mathbb{1}_{\{\tau < \infty\}} \mid s_0 = s, a_0 = a \right] \\
 &\stackrel{(d)}{=} \sum_{\tau'=1}^{\infty} \mathbf{E} \left[\sum_{t=0}^{\tau'-1} (r(s_t, \pi_G^*) - \hat{\zeta}_t^{\pi_G^*}(s_t, a_t, w_{t+1} | \varphi)) \mathbb{1}_{\{\tau=\tau'\}} \mid s_0 = s, a_0 = a \right] \\
 &\stackrel{(e)}{=} \sum_{\tau'=1}^{\infty} \mathbf{E} \left[\sum_{t=0}^{\tau'-1} (r(s_t, \pi_G^*) - \mathbf{E}[\hat{\zeta}_t^{\pi_G^*}(s_t, a_t, w_{t+1} | \varphi) | \mathcal{G}_t]) \mathbb{1}_{\{\tau=\tau'\}} \mid s_0 = s, a_0 = a \right] \\
 &\stackrel{(f)}{=} \sum_{\tau'=1}^{\infty} \mathbf{E} \left[\sum_{t=0}^{\tau'-1} (r(s_t, \pi_G^*) - \zeta_t^{\pi_G^*}(s_t, a_t, w_{t+1} | \varphi)) \mathbb{1}_{\{\tau=\tau'\}} \mid s_0 = s, a_0 = a \right] \\
 &\stackrel{(g)}{=} \mathbf{E} \left[\sum_{t=0}^{\tau-1} (r(s_t, \pi_G^*) - \zeta_t^{\pi_G^*}(s_t, a_t, w_{t+1} | \varphi)) \mathbb{1}_{\{\tau < \infty\}} \mid s_0 = s, a_0 = a \right] \\
 &\stackrel{(h)}{=} \mathbf{E} \left[\sum_{t=0}^{\tau-1} (r(s_t, \pi_G^*) - \zeta_t^{\pi_G^*}(s_t, a_t, w_{t+1} | \varphi)) \mid s_0 = s, a_0 = a \right] \\
 &\stackrel{(i)}{=} \max_{\pi_G \in \Pi_{\mathbb{G}}} \mathbf{E} \left[\sum_{t=0}^{\tau-1} (r(s_t, \pi_G) - \zeta_t^{\pi_G}(s_t, a_t, w_{t+1} | \varphi)) \mid s_0 = s, a_0 = a \right] \\
 &\geq Q^*(s, a).
 \end{aligned}$$

Equality (a) and (b) follow from the definition of $\hat{Q}_0^U(s, a)$ and since $\hat{\mathbb{G}}$ is a relaxation of the perfect information relaxation \mathbb{G} which allow us to interchange the maximum and the expectation. The first inequality follows because $\pi_G^* \in \Pi_{\hat{\mathbb{G}}}$ since $\Pi_{\mathbb{G}} \subseteq \Pi_{\hat{\mathbb{G}}}$. Equality (c) follows since r and φ are uniformly bounded and τ has finite mean. Equalities (d) and (e) follow from the law of total expectations. Equality (f) follows from Lemma A.1 in (Brown et al., 2010) and from the estimated penalty terms being unbiased, i.e., $\mathbf{E}[\hat{\zeta}_t^{\pi_G^*}(s_t, a_t, w_{t+1} | \varphi) | \mathcal{G}_t] = \zeta_t^{\pi_G^*}(s_t, a_t, w_{t+1} | \varphi)$. Equalities (g) and (h) follow by the law of total expectation and τ being almost surely finite stopping time, respectively. Equality (i) follows since by definition π_G^* is the optimal solution of (18). The last inequality follows by weak duality (Proposition 1(i)). \square

First, we state a technical lemma that is used in the proof of Proposition 3 and Lemma 1.

Lemma A.1. *For all $n = 1, 2, \dots$, if $L_{n-1}(s, a) \leq U_{n-1}(s, a)$ and $Q'_{n-1} \in \mathcal{Q}$ then $L_n(s, a) \leq U_n(s, a)$ for all $(s, a) \in \mathcal{S} \times \mathcal{A}$.*

Proof. Fix an $(s, a) \in \mathcal{S} \times \mathcal{A}$. Note that the optimal values of the inner problems in (9) and (10), $\hat{Q}_0^U(s, a)$ and $\hat{Q}_0^L(s, a)$ respectively, are computed using the same sample path \mathbf{w} and for each period within the inner DP, the same batch of samples is used for estimating the expectation in both the upper and lower bound problems. For clarity, let us denote the values of $\hat{Q}_0^L(s, a)$ and $\hat{Q}_0^U(s, a)$ at iteration $n = 1, 2, 3, \dots$ by $\hat{Q}_{n,0}^L(s, a)$ and $\hat{Q}_{n,0}^U(s, a)$, respectively. Assume $\alpha_n(s, a) \leq 1$ for all n . We provide a proof by induction. For $n = 1$, we have:

$$Q_1(s_0, a_0) = Q'_0(s_0, a_0) + \alpha_0(s_0, a_0) \left[r(s_0, a_0) + \gamma \max_a Q'_0(s_1, a) - Q'_0(s_0, a_0) \right].$$

Since the rewards $r(s, a)$ are uniformly bounded, $0 < \gamma < 1$ and $|Q'_0(s, a)| \leq \rho$ then Q_1 is bounded. Set $\varphi = Q_1$, since the actions selected by the policy π_{Q_1} are feasible in (9), we have

$$\hat{Q}_{1,0}^U(s, a) - \hat{Q}_{1,0}^L(s, a) \geq 0,$$

and with $L_0(s, a) \leq U_0(s, a)$, it follows that $L_1(s, a) \leq U_1(s, a)$. A similar proof can be used to show the inductive case also holds at iteration n ,

$$\begin{aligned} Q_n(s_{n-1}, a_{n-1}) &= Q'_{n-1}(s_{n-1}, a_{n-1}) \\ &+ \alpha_{n-1}(s_{n-1}, a_{n-1}) \left[r(s_{n-1}, a_{n-1}) + \gamma \max_a Q'_{n-1}(s_n, a) - Q'_{n-1}(s_{n-1}, a_{n-1}) \right]. \end{aligned}$$

By the inductive hypothesis, we have $Q'_{n-1} \in \mathcal{Q}$ and $L_{n-1}(s, a) \leq U_{n-1}(s, a)$. Then, similar to the base case, we have $Q_n \in \mathcal{Q}$ and $\hat{Q}_{n,0}^U(s, a) - \hat{Q}_{n,0}^L(s, a) \geq 0$. Therefore, $L_n(s, a) \leq U_n(s, a)$. Since our choice of (s, a) was arbitrary then the result follows for all $(s, a) \in \mathcal{S} \times \mathcal{A}$. \square

A.2 Proof of Proposition 3

Proof. Part (i): First, note that by (12) and (13) the upper and lower bound estimates $U_n(s, a)$ and $L_n(s, a)$ are bounded below and above by ρ and $-\rho$ respectively for all $(s, a) \in \mathcal{S} \times \mathcal{A}$ and for all n , where $\rho = R_{\max}/(1 - \gamma)$. We assume in this proof that $\alpha_n(s, a) \leq 1$ for all n . Let $\tilde{L}_n = \max_{(s,a)} L_n(s, a)$ and $\tilde{U}_n = \max_{(s,a)} U_n(s, a)$. We claim that for every iteration n , we have that for all (s, a) ,

$$\bar{L}_n \leq Q_n(s, a) \leq \bar{U}_n \quad \text{and} \quad \bar{L}'_n \leq Q'_n(s, a) \leq \bar{U}'_n \quad (19)$$

where

$$\bar{L}_n = \min \left\{ \tilde{U}_{n-1}(1 + \gamma), \dots, \tilde{U}_1 \sum_{i=0}^{n-1} \gamma^i, -M \sum_{i=0}^n \gamma^i \right\} \quad (20)$$

$$\bar{U}_n = \max \left\{ \tilde{L}_{n-1}(1 + \gamma), \dots, \tilde{L}_1 \sum_{i=0}^{n-1} \gamma^i, M \sum_{i=0}^n \gamma^i \right\}, \quad (21)$$

$$\bar{L}'_n = \min \left\{ \tilde{U}_n, \tilde{U}_{n-1}(1 + \gamma), \dots, \tilde{U}_1 \sum_{i=0}^{n-1} \gamma^i, -M \sum_{i=0}^n \gamma^i \right\}, \quad (22)$$

$$\bar{U}'_n = \max \left\{ \tilde{L}_n, \tilde{L}_{n-1}(1 + \gamma), \dots, \tilde{L}_1 \sum_{i=0}^{n-1} \gamma^i, M \sum_{i=0}^n \gamma^i \right\}, \quad (23)$$

and M is a finite positive scalar defined as $M = \max \{ R_{\max}, \max_{(s,a)} Q_0(s, a) \}$.

The result follows from the claim in (19). To see this note that at any iteration n , \bar{L}_n and \bar{L}'_n are bounded below by $-\rho \sum_{i=0}^n \gamma^i$ since each term inside the minimum of (20) and (22) is bounded below by $-\rho \sum_{i=0}^n \gamma^i$. As $n \rightarrow \infty$, we have

$$\frac{-\rho}{1 - \gamma} \leq \liminf_{n \rightarrow \infty} \bar{L}_n \quad \text{and} \quad \frac{-\rho}{1 - \gamma} \leq \liminf_{n \rightarrow \infty} \bar{L}'_n. \quad (24)$$

An analogous argument yields

$$\limsup_{n \rightarrow \infty} \bar{U}_n \leq \frac{\rho}{1 - \gamma} \quad \text{and} \quad \limsup_{n \rightarrow \infty} \bar{U}'_n \leq \frac{\rho}{1 - \gamma}. \quad (25)$$

Boundedness of $Q_n(s, a)$ and $Q'_n(s, a)$ for all $(s, a) \in \mathcal{S} \times \mathcal{A}$ follows from (19), (24) and (25).

Now, we prove our claim in (19) by induction. Since Algorithm 1 is asynchronous then at the n th iteration, the updates for the action-value iterates for (s, a) , $Q_{n+1}(s, a)$ and $Q'_{n+1}(s, a)$, are either according to (11) and (14) (case 1) or set equal to $Q_n(s, a)$ and $Q'_n(s, a)$ respectively (case 2).

We first focus on $Q'_n(s, a) \leq \bar{U}'_n$ part of (19), since $\bar{L}'_n \leq Q'_n(s, a)$ and $\bar{L}_n \leq Q_n(s, a) \leq \bar{U}_n$ proceed in an analogous manner. For $n = 1$, we have $Q'_0(s, a) = Q_0(s, a)$, so if the update is carried out as in case 1,

$$\begin{aligned} Q_1(s, a) &= (1 - \alpha_0(s, a)) Q'_0(s, a) + \alpha_0(s, a) [r(s, a) + \gamma \max_a Q'_0(s', a)] \\ &\leq (1 - \alpha_0(s, a)) M + \alpha_0(s, a) M + \alpha_0(s, a) \gamma M \\ &\leq M(1 + \gamma) \end{aligned}$$

so $Q'_1(s, a) \leq \max\{L_1(s, a), \min\{U_1(s, a), M(1 + \gamma)\}\}$. Now consider the case where $U_1(s, a) \leq M(1 + \gamma)$. Since Q_0 is bounded by ρ and $L_0(s, a) \leq U_0(s, a)$ then by Lemma A.1, we have $L_1(s, a) \leq U_1(s, a)$, so

$$Q'_1(s, a) \leq U_1(s, a) \leq M(1 + \gamma). \quad (26)$$

Otherwise, if $U_1(s, a) \geq M(1 + \gamma)$, we then have

$$Q'_1(s, a) \leq \max\{L_1(s, a), M(1 + \gamma)\}. \quad (27)$$

From (26) and (27), we have

$$\begin{aligned} Q'_1(s, a) &\leq \max\{L_1(s, a), M(1 + \gamma)\} \\ &\leq \max\{\tilde{L}_1, M(1 + \gamma)\}. \end{aligned} \quad (28)$$

If the update is carried out as in case 2, we have,

$$\begin{aligned} Q'_1(s, a) &= Q'_0(s, a) \\ &\leq M \\ &< M(1 + \gamma) \\ &\leq \max\{\tilde{L}_1, M(1 + \gamma)\}. \end{aligned}$$

Thus $\bar{U}'_n(s, a)$ part of (19) is true for $n = 1$. Suppose that it is true for $n = 1, 2, \dots, k$. We will show it for $n = k + 1$. Consider first the instance where the update is carried out according to case 1. We do casework on the inequality

$$Q'_k(s, a) \leq \max\left\{\tilde{L}_k, \tilde{L}_{k-1}(1 + \gamma), \dots, \tilde{L}_1 \sum_{i=0}^{k-1} \gamma^i, M \sum_{i=0}^k \gamma^i\right\}, \quad (29)$$

which holds for all (s, a) . First, let us consider the case where the right-hand-side of (29) is equal to $\tilde{L}_{k'} \sum_{i=0}^{k-k'} \gamma^i$ for some k' such that $1 \leq k' \leq k$. Then, we have

$$\begin{aligned} Q_{k+1}(s, a) &= (1 - \alpha_k(s, a))Q'_k(s, a) + \alpha_k(s, a)[r(s, a) + \gamma \max_a Q'_k(s', a)] \\ &\leq (1 - \alpha_k(s, a))\tilde{L}_{k'} \sum_{i=0}^{k-k'} \gamma^i + \alpha_k(s, a)M + \alpha_k(s, a)\gamma \tilde{L}_{k'} \sum_{i=0}^{k-k'} \gamma^i \\ &\leq (1 - \alpha_k(s, a))\tilde{L}_{k'} \sum_{i=0}^{k-k'} \gamma^i + \alpha_k(s, a)\tilde{L}_{k'} + \alpha_k(s, a)\tilde{L}_{k'} \sum_{i=1}^{k-k'+1} \gamma^i \\ &= (1 - \alpha_k(s, a))\tilde{L}_{k'} \sum_{i=0}^{k-k'} \gamma^i + \alpha_k(s, a)\tilde{L}_{k'} \sum_{i=0}^{k-k'} \gamma^i + \alpha_k(s, a)\tilde{L}_{k'} \gamma^{k-k'+1} \\ &\leq \tilde{L}_{k'} \sum_{i=0}^{k-k'} \gamma^i + \tilde{L}_{k'} \gamma^{k-k'+1} \\ &= \tilde{L}_{k'} \sum_{i=0}^{k-k'+1} \gamma^i \end{aligned} \quad (30)$$

The first inequality holds by the induction assumption, (29). The second inequality holds since in this case we have the right-hand-side of (29) is equal to $\tilde{L}_{k'}(1 + \gamma + \dots + \gamma^{k-k'})$ so it follows that

$$\tilde{L}_{k'}(1 + \gamma + \dots + \gamma^{k-k'}) \geq M(1 + \gamma + \dots + \gamma^k)$$

which in turn implies that $\tilde{L}_{k'} \geq M$. Finally, the third inequality holds by the assumption that $\alpha_n(s, a) \leq 1$ for all n . We have

$$\begin{aligned} Q'_{k+1}(s, a) &= \max\{L_{k+1}(s, a), \min\{U_{k+1}(s, a), Q_{k+1}(s, a)\}\} \\ &\leq \max\{L_{k+1}(s, a), \min\{U_{k+1}(s, a), \tilde{L}_{k'}(1 + \gamma + \dots + \gamma^{k-k'+1})\}\}. \end{aligned}$$

Now, consider the case where $U_{k+1}(s, a) \leq \tilde{L}_{k'}(1 + \gamma + \dots + \gamma^{k-k'+1})$. By the induction assumption $Q'_n(s, a)$ is bounded below by $-\rho \sum_{i=0}^n \gamma^i$ and above by $\rho \sum_{i=0}^n \gamma^i$ for all $(s, a) \in \mathcal{S} \times \mathcal{A}$ and all $n = 1, 2, \dots, k$. Since $L_0(s, a) \leq U_0(s, a)$ then Lemma A.1 can be applied iteratively on $n = 1, \dots, k + 1$, to obtain that $L_{K+1}(s, a) \leq U_{k+1}(s, a)$ for all $(s, a) \in \mathcal{S} \times \mathcal{A}$. Thus, we have

$$Q'_{k+1}(s, a) \leq U_{k+1}(s, a) \leq \tilde{L}_{k'}(1 + \gamma + \dots + \gamma^{k-k'+1}). \quad (31)$$

Otherwise, if $U_{k+1}(s, a) \geq \tilde{L}_{k'}(1 + \gamma + \dots + \gamma^{k-k'+1})$, we have

$$\begin{aligned} Q'_{k+1}(s, a) &\leq \max\{L_{k+1}(s, a), \tilde{L}_{k'}(1 + \gamma + \dots + \gamma^{k-k'+1})\} \\ &\leq \max\{\tilde{L}_{k+1}, \tilde{L}_{k'}(1 + \gamma + \dots + \gamma^{k-k'+1})\}. \end{aligned} \quad (32)$$

Moving on to the case where the right-hand-side of (29) is equal to $M(1 + \gamma + \dots + \gamma^k)$:

$$\begin{aligned} Q_{k+1}(s, a) &= (1 - \alpha_k(s, a)) Q'_k(s, a) + \alpha_k(s, a) [r(s, a) + \gamma \max_a Q'_k(s', a)] \\ &\leq (1 - \alpha_k(s, a)) M \sum_{i=0}^k \gamma^i + \alpha_k(s, a) M + \alpha_k(s, a) \gamma M \sum_{i=0}^k \gamma^i \\ &= (1 - \alpha_k(s, a)) M \sum_{i=0}^k \gamma^i + \alpha_k(s, a) M + \alpha_k(s, a) M \sum_{i=1}^{k+1} \gamma^i \\ &\leq M \sum_{i=0}^k \gamma^i - \alpha_k(s, a) M \sum_{i=0}^k \gamma^i + \alpha_k(s, a) M \sum_{i=0}^k \gamma^i + M \gamma^{k+1} \\ &= M(1 + \gamma + \dots + \gamma^{k+1}). \end{aligned} \quad (33)$$

We have

$$\begin{aligned} Q'_{k+1}(s, a) &= \max\{L_{k+1}(s, a), \min(U_{k+1}(s, a), Q_{k+1}(s, a))\} \\ &\leq \max\{L_{k+1}(s, a), \min(U_{k+1}(s, a), M(1 + \gamma + \dots + \gamma^{k+1}))\} \end{aligned}$$

Now if $U_{k+1}(s, a) \leq M(1 + \gamma + \dots + \gamma^{k+1})$ then by applying Lemma A.1 as before,

$$Q'_{k+1}(s, a) \leq U_{k+1}(s, a) \leq M(1 + \gamma + \dots + \gamma^{k+1}). \quad (34)$$

Otherwise, if $U_{k+1}(s, a) \geq M(1 + \gamma + \dots + \gamma^{k+1})$, we have

$$\begin{aligned} Q'_{k+1}(s, a) &\leq \max\{L_{k+1}(s, a), M(1 + \gamma + \dots + \gamma^{k+1})\} \\ &\leq \max\{\tilde{L}_{k+1}, M(1 + \gamma + \dots + \gamma^{k+1})\}. \end{aligned} \quad (35)$$

Now, if the update is carried out according to case 2,

$$\begin{aligned} Q'_{k+1}(s, a) &= Q'_K(s, a) \\ &\leq \max\{\tilde{L}_k, \tilde{L}_{k-1}(1 + \gamma), \dots, \tilde{L}_1 \sum_{i=0}^{k-1} \gamma^i, M \sum_{i=0}^k \gamma^i\} \\ &\leq \max\{\tilde{L}_{k+1}, (1 + \gamma) \tilde{L}_k, \dots, \tilde{L}_1 \sum_{i=0}^k \gamma^i, M \sum_{i=0}^{k+1} \gamma^i\}. \end{aligned} \quad (36)$$

By (31), (32), (34), (35) and (36), we have $Q'_{k+1}(s, a) \leq \bar{U}'_{k+1}$.

A similar argument can be made to show $\bar{L}_n \leq Q'_n(s, a)$ and $\bar{L}_n \leq Q_n(s, a) \leq \bar{U}_n$, which completes the inductive proof. \square

Proof. Part (ii): Fix an $(s, a) \in \mathcal{S} \times \mathcal{A}$. By part (i) we have the action-value iterates Q_n and Q'_n are bounded for all n . We denote the ‘‘sampling noise’’ term using

$$\xi_n^L(s, a) = \hat{Q}_{n,0}^L(s, a) - \mathbf{E}[\hat{Q}_{n,0}^L(s, a)].$$

We also define an accumulated noise process started at iteration ν by $W_{\nu,\nu}^L(s, a) = 0$, and

$$W_{n+1,\nu}^L(s, a) = (1 - \alpha_n(s, a)) W_{n,\nu}^L(s, a) + \alpha_n(s, a) \xi_{n+1}^L(s, a) \quad \forall n \geq \nu,$$

which averages noise terms together across iterations. Note that since τ is an almost surely finite stopping time, the rewards $r(s, a)$ are uniformly bounded and Q_{n+1} is also bounded (by part (i)). Then, $\hat{Q}_{n,0}^L$ is bounded by some random variable and so is the conditional variance of $\xi_n^L(s, a)$. Hence, Corollary 4.1 in (Bertsekas & Tsitsiklis, 1996) applies and it follows that

$$\lim_{n \rightarrow \infty} W_{n,\nu}^L(s, a) = 0 \quad \forall \nu \geq 0.$$

Let $\tilde{\nu}$ be large enough so that $\alpha_n(s, a) \leq 1$ for all $n \geq \tilde{\nu}$. We also define

$$\begin{aligned} Y_{\tilde{\nu}}^L(s, a) &= \rho, \\ Y_{n+1}^L(s, a) &= (1 - \alpha_n(s, a)) Y_n^L(s, a) + \alpha_n(s, a) Q^*(s, a), \quad \forall n \geq \tilde{\nu}. \end{aligned}$$

It is easy to see that the sequence $Y_n^L(s, a) \rightarrow Q^*(s, a)$. We claim that for all iterations $n \geq \tilde{\nu}$, it holds that

$$L_n(s, a) \leq \min\{\rho, Y_n^L(s, a) + W_{n, \tilde{\nu}}^L(s, a)\}.$$

To prove this claim, we proceed by induction on n . For $n = \tilde{\nu}$, we have

$$Y_{\tilde{\nu}}^L(s, a) = \rho \quad \text{and} \quad W_{\tilde{\nu}, \tilde{\nu}}^L(s, a) = 0,$$

so it is clear that the statement is true for the base case. We now show that it is true for $n + 1$ given that it holds at n :

$$\begin{aligned} L_{n+1}(s, a) &= \min\{\rho, (1 - \alpha_n(s, a)) L_n(s, a) + \alpha_n(s, a) (\hat{Q}_{n,0}^L(s, a) - \mathbf{E}[\hat{Q}_{n,0}^L(s, a)] + \mathbf{E}[\hat{Q}_{n,0}^L(s, a)])\} \\ &= \min\{\rho, (1 - \alpha_n(s, a)) L_n(s, a) + \alpha_n(s, a) \xi_n^L(s, a) + \alpha_n(s, a) \mathbf{E}[\hat{Q}_{n,0}^L(s, a)]\} \\ &\leq \min\{\rho, (1 - \alpha_n(s, a)) (Y_n^L(s, a) + W_{n, \nu_k}^L(s, a)) + \alpha_n(s, a) \xi_n^L(s, a) + \alpha_n(s, a) Q^*(s, a)\} \\ &\leq \min\{\rho, Y_{n+1}^L(s, a) + W_{n+1, \tilde{\nu}}^L(s, a)\}, \end{aligned}$$

where the first inequality follows by the induction hypothesis and $\mathbf{E}[\hat{Q}_{n,0}^L(s, a)] \leq Q^*(s, a)$ follows by Proposition 2. Next, since $Y_n^L(s, a) \rightarrow Q^*(s, a)$, $W_{n, \nu_k}^L(s, a) \rightarrow 0$ and $Q^*(s, a) \leq \rho$, we have

$$\limsup_{n \rightarrow \infty} L_n(s, a) \leq Q^*(s, a).$$

Therefore, since our choice of (s, a) was arbitrary, it follows that for every $\eta > 0$, there exists some time n' such that $L_n(s, a) \leq Q^*(s, a) + \eta$ for all $(s, a) \in \mathcal{S} \times \mathcal{A}$ and $n \geq n'$.

Using Proposition 2, $Q^*(s, a) \leq \mathbf{E}[\hat{Q}_{n,0}^U(s, a)]$, a similar argument as the above can be used to establish that

$$Q^*(s, a) \leq \liminf_{n \rightarrow \infty} U_n(s, a).$$

Hence, there exists some time n'' such that $Q^*(s, a) - \eta \leq U_n(s, a)$ for all (s, a) and $n \geq n''$. Take n_0 to be some time greater than n' and n'' and the result follows. \square

A.3 Proof of Lemma 1

Proof. We use induction on n . Since for all $(s, a) \in \mathcal{S} \times \mathcal{A}$, $L_0(s, a) \leq U_0(s, a)$ and $-\rho \leq Q'_0(s, a) \leq \rho$ then $L_1(s, a) \leq U_1(s, a)$ by Lemma A.1. Suppose that, $L_n(s, a) \leq U_n(s, a)$ holds for all (s, a) for all $n = 1, \dots, k$. For all $(s, a) \in \mathcal{S} \times \mathcal{A}$, we have $Q'_k(s, a)$ is bounded since by Proposition 3(i) $Q'_n(s, a)$ is bounded for all n . We also have $L_k(s, a) \leq U_k(s, a)$ for all (s, a) by the induction assumption. Applying Lemma A.1 again at $n = k + 1$ yields $L_{k+1}(s, a) \leq U_{k+1}(s, a)$ and the inductive proof is complete. \square

A.4 Proof of Theorem 1

Proof. We first prove part (i). We start by writing Algorithm 1 using DP operator notation. Define a mapping H such that

$$(HQ')(s, a) = r(s, a) + \gamma \mathbf{E}[\max_{a'} Q'(s', a')],$$

where $s' = f(s, a, w)$. It is well-known that the mapping H is a γ -contraction in the maximum norm. We also define a random noise term

$$\xi_n(s, a) = \gamma \max_{a'} Q'_n(s', a') - \gamma \mathbf{E}[\max_{a'} Q'_n(s', a')]. \quad (37)$$

The main update rules of Algorithm 1 can then be written as

$$Q_{n+1}(s, a) = (1 - \alpha_n(s, a)) Q'_n(s, a) + \alpha_n(s, a) [(HQ'_n)(s, a) + \xi_{n+1}(s, a)],$$

$$\begin{aligned}
 U_{n+1}(s, a) &= \Pi_{[-\rho, \infty]} \left[(1 - \beta_n(s, a)) U_n(s, a) + \beta_n(s, a) \hat{Q}_0^U(s, a) \right], \\
 L_{n+1}(s, a) &= \Pi_{[\infty, \rho]} \left[(1 - \beta_n(s, a)) L_n(s, a) + \beta_n(s, a) \hat{Q}_0^L(s, a) \right], \\
 Q'_{n+1}(s, a) &= \Pi_{[L_{n+1}(s, a), U_{n+1}(s, a)]} [Q_{n+1}(s, a)].
 \end{aligned} \tag{38}$$

Assume without loss of generality that $Q^*(s, a) = 0$ for all state-action pairs (s, a) . This can be established by shifting the origin of the coordinate system. Note that by (38) at any iteration n and for all (s, a) , we have $L_n(s, a) \leq Q'_n(s, a) \leq U_n(s, a)$.

We proceed via induction. First, note that by Propostion 3(i) the iterates of Algorithm 1 $Q'_n(s, a)$ are bounded in the sense that there exists a D_0 such that $|Q'_n(s, a)| \leq D_0$ for all (s, a) and iterations n . Define the sequence $D_{k+1} = (\gamma + \epsilon) D_k$, such that $\gamma + \epsilon < 1$ and $\epsilon > 0$. Clearly, $D_k \rightarrow 0$. Suppose that there exists some time n_k such that for all (s, a) ,

$$\max\{-D_k, L_n(s, a)\} \leq Q'_n(s, a) \leq \min\{D_k, U_n(s, a)\}, \quad \forall n \geq n_k.$$

We will show that this implies the existence of some time n_{k+1} such that

$$\max\{-D_{k+1}, L_n(s, a)\} \leq Q'_n(s, a) \leq \min\{D_{k+1}, U_n(s, a)\} \quad \forall (s, a), n \geq n_{k+1}.$$

This implies that $Q'_n(s, a)$ converges to $Q^*(s, a) = 0$ for all (s, a) . We also assume that $\alpha_n(s, a) \leq 1$ for all (s, a) and n . Define an accumulated noise process started at n_k by $W_{n_k, n_k}(s, a) = 0$, and

$$W_{n+1, n_k}(s, a) = (1 - \alpha_n(s, a)) W_{n, n_k}(s, a) + \alpha_n(s, a) \xi_{n+1}(s, a), \quad \forall n \geq n_k, \tag{39}$$

where $\xi_n(s, a)$ is as defined in (37). We now use Corollary 4.1 in (Bertsekas & Tsitsiklis, 1996) which states that under the assumptions of Theorem 1 on the step size $\alpha_n(s, a)$, and if $\mathbf{E}[\xi_n(s, a) | \mathcal{F}_n] = 0$ and $\mathbf{E}[\xi_n^2(s, a) | \mathcal{F}_n] \leq A_n$, where the random variable A_n is bounded with probability 1, the sequence $W_{n+1, n_k}(s, a)$ defined in (39) converges to zero, with probability 1. From our definition of the stochastic approximation noise $\xi_n(s, a)$ in (37), we have

$$\mathbf{E}[\xi_n(s, a) | \mathcal{F}_n] = 0 \quad \text{and} \quad \mathbf{E}[\xi_n^2(s, a) | \mathcal{F}_n] \leq C(1 + \max_{s', a'} Q_n'^2(s', a')),$$

where C is a constant. Then, it follows that

$$\lim_{n \rightarrow \infty} W_{n, n_k}(s, a) = 0 \quad \forall (s, a), n_k.$$

Now, for the sake of completeness, we restate a lemma from (Bertsekas & Tsitsiklis, 1996) below, which we will use to bound the accumulated noise.

Lemma A.2 (Lemma 4.2 in (Bertsekas & Tsitsiklis, 1996)). *For every $\delta > 0$, and with probability one, there exists some n' such that $|W_{n, n'}(s, a)| \leq \delta$, for all $n \geq n'$.*

Using the above lemma, let $n_{k'} \geq n_k$ such that, for all $n \geq n_{k'}$ we have

$$|W_{n, n_{k'}}(s, a)| \leq \gamma \epsilon D_k < \gamma D_k.$$

Furthermore, by Proposition 3(ii) let $\nu_k \geq n_{k'}$ such that, for all $n \geq \nu_k$ we have

$$L_n(s, a) \leq \gamma D_k - \gamma \epsilon D_k \quad \text{and} \quad \gamma \epsilon D_k - \gamma D_k \leq U_n(s, a).$$

Define another sequence Y_n that starts at iteration ν_k .

$$Y_{\nu_k}(s, a) = D_k \quad \text{and} \quad Y_{n+1}(s, a) = (1 - \alpha_n(s, a)) Y_n(s, a) + \alpha_n(s, a) \gamma D_k \tag{40}$$

Note that it is easy to show that the sequence $Y_n(s, a)$ in (40) is decreasing, bounded below by γD_k , and converges to γD_k as $n \rightarrow \infty$. Now we state the following lemma.

Lemma A.3. *For all state-action pairs (s, a) and iterations $n \geq \nu_k$, it holds that:*

$$(1) \quad Q'_n(s, a) \leq \min\{U_n(s, a), Y_n(s, a) + W_{n, \nu_k}(s, a)\},$$

$$(2) \max\{L_n(s, a), -Y_n(s, a) + W_{n, \nu_k}(s, a)\} \leq Q'_n(s, a).$$

Proof. We focus on part (1). For the base case $n = \nu_k$, the statement holds because $Y_{\nu_k}(s, a) = D_k$ and $W_{\nu_k, \nu_k}(s, a) = 0$. We assume it is true for n and show that it continues to hold for $n + 1$:

$$\begin{aligned} Q_{n+1}(s, a) &= (1 - \alpha_n(s, a))Q'_n(s, a) + \alpha_n(s, a) [(HQ'_n)(s, a) + \xi_{n+1}(s, a)] \\ &\leq (1 - \alpha_n(s, a)) \min\{U_n(s, a), Y_n(s, a) + W_{n, \nu_k}(s, a)\} \\ &\quad + \alpha_n(s, a) (HQ'_n)(s, a) + \alpha_n(s, a) \xi_{n+1}(s, a) \\ &\leq (1 - \alpha_n(s, a)) (Y_n(s, a) + W_{n, \nu_k}(s, a)) + \alpha_n(s, a) \gamma D_k + \alpha_n(s, a) \xi_{n+1}(s, a) \\ &\leq Y_{n+1}(s, a) + W_{n+1, \nu_k}(s, a), \end{aligned}$$

where we used $(HQ'_n) \leq \gamma \|Q'_n\| \leq \gamma D_k$. Now, we have

$$\begin{aligned} Q'_{n+1}(s, a) &= \Pi_{[L_{n+1}(s, a), U_{n+1}(s, a)]} [Q_{n+1}(s, a)] \\ &\leq \Pi_{[L_{n+1}(s, a), U_{n+1}(s, a)]} [Y_{n+1}(s, a) + W_{n+1, \nu_k}(s, a)] \\ &\leq \min\{U_{n+1}(s, a), Y_{n+1}(s, a) + W_{n+1, \nu_k}(s, a)\}. \end{aligned}$$

The first inequality holds because

$$Q_{n+1}(s, a) \leq Y_{n+1}(s, a) + W_{n+1, \nu_k}(s, a).$$

The second inequality holds because $Y_{n+1}(s, a) + W_{n+1, \nu_k}(s, a) \geq \gamma D_k - \gamma \epsilon D_k$, $L_n(s, a) \leq \gamma D_k - \gamma \epsilon D_k$, and $L_n(s, a) \leq U_n(s, a)$ by Lemma 1. Symmetrically, it can be shown that

$$\max\{L_{n+1}(s, a), -Y_{n+1}(s, a) + W_{n+1, \nu_k}(s, a)\} \leq Q'_{n+1}(s, a),$$

which completes the proof. \square

Since $Y_n(s, a) \rightarrow \gamma D_k$ and $W_{n, \nu_k}(s, a) \rightarrow 0$, we have

$$\limsup_{n \rightarrow \infty} \|Q'_n\| \leq \gamma D_k < D_{k+1}.$$

Therefore, there exists some time n_{k+1} such that

$$\max\{-D_{k+1}, L_n(s, a)\} \leq Q'_n(s, a) \leq \min\{D_{k+1}, U_n(s, a)\} \quad \forall (s, a), n \geq n_{k+1},$$

completing thus the induction.

For part (ii) of the theorem: we fix (s, a) and focus on the convergence analysis of $U_n(s, a)$ to $Q^*(s, a)$. A similar analysis can be done to show $L_n(s, a) \rightarrow Q^*(s, a)$ almost surely. First note that we can write (12) the update equation of $U_n(s, a)$ as:

$$U_{n+1}(s_n, a_n) = \Pi_{[-\rho, \infty]} [U_n(s_n, a_n) + \beta_n(s_n, a_n) [\psi_n(U_n(s_n, a_n), Q_{n+1}(s_n, a_n))]]$$

where $\psi_n(U_n(s_n, a_n), Q_{n+1}(s_n, a_n))$ is the stochastic gradient and in this case is equal to $\hat{Q}_0^U(s_n, a_n) - U_n(s_n, a_n)$. We define the noise terms

$$\bar{\epsilon}_{n+1}(s_n, a_n) = \psi_n(U_n(s_n, a_n), Q^*(s_n, a_n)) - \mathbf{E}[\psi_n(U_n(s_n, a_n), Q^*(s_n, a_n))] \quad (41)$$

$$\bar{\epsilon}_{n+1}(s_n, a_n) = \psi_n(U_n(s_n, a_n), Q_{n+1}(s_n, a_n)) - \psi_n(U_n(s_n, a_n), Q^*(s_n, a_n)). \quad (42)$$

Note here that $\bar{\epsilon}_{n+1}(s_n, a_n)$ represents the error that the sample gradient deviates from its mean when computed using the optimal action-value Q^* and $\bar{\epsilon}_{n+1}(s_n, a_n)$ is the error between the two evaluations of ψ_n due only to the difference between $Q_{n+1}(s_n, a_n)$ and $Q^*(s_n, a_n)$. Thus, we have

$$\psi_n(U_n(s_n, a_n), Q_{n+1}(s_n, a_n)) = \mathbf{E}[\psi_n(U_n(s_n, a_n), Q^*(s_n, a_n))] + \bar{\epsilon}_{n+1}(s_n, a_n) + \bar{\epsilon}_{n+1}(s_n, a_n).$$

Since $Q_n \rightarrow Q^*$ by part (i) of the Theorem, then $\bar{\epsilon}_n(s_n, a_n) \rightarrow 0$ almost surely. It is now convenient to view $U_n(s, a)$ as a stochastic process in n , adapted to the filtration $\{\mathcal{F}_n\}_{n \geq 0}$. By definition of $\bar{\epsilon}_{n+1}(s, a)$, we have that

$$\mathbf{E}[\bar{\epsilon}_{n+1}(s, a) | \mathcal{F}_n] = 0 \quad a.s.$$

Since $\bar{\epsilon}_{n+1}(s, a)$ is unbiased and $\bar{\epsilon}_{n+1}(s, a)$ converges to zero, we can apply Theorem 2.4 of (Kushner & Yin, 2003), a standard stochastic approximation convergence result, to conclude that $U_n(s, a) \rightarrow Q^*(s, a)$ almost surely. Since our choice of (s, a) was arbitrary, this convergence holds for all $(s, a) \in \mathcal{S} \times \mathcal{A}$. \square

A.5 Proof of Lemma 2

Proof. First note that Lemma A.1 still holds in this case. To see this, note that if the requirements of the lemma are satisfied, i.e., if we are at iteration $n = 1, 2, \dots$, and in the previous iteration, we had $L_{n-1}(s, a) \leq U_{n-1}(s, a)$ for all (s, a) and $Q'_{n-1} \in \mathcal{Q}$ then Q_n is bounded using the same argument as before. Since the rewards $r(s, a)$ are uniformly bounded and τ is an almost surely finite stopping time, then $\tilde{Q}_{n,0}^L$ and $\tilde{Q}_{n,0}^U$ are finite. Moreover, since $\tilde{Q}_{n,0}^L$ and $\tilde{Q}_{n,0}^U$ are computed using the same sample path w , it follows that

$$\tilde{Q}_{n,0}^U(s, a) - \tilde{Q}_{n,0}^L(s, a) \geq 0, \quad \text{for all } (s, a) \in \mathcal{S} \times \mathcal{A}.$$

This can be easily seen if we subtract (17) from (16). Notice that the reward and the penalty will both cancel out and we'll have $\tilde{Q}_{n,t}^U - \tilde{Q}_{n,t}^L \geq 0$ for all $t = 0, 1, \dots, \tau - 1$. With $L_{n-1}(s, a) \leq U_{n-1}(s, a)$ and $\tilde{Q}_{n,0}^U(s, a) \geq \tilde{Q}_{n,0}^L(s, a)$ it follows that $L_n(s, a) \leq U_n(s, a)$ for all (s, a) .

Now, we prove Proposition 3 again for the experience buffer case.

For part (i): our original proof still holds since Lemma A.1 still holds.

For part (ii): first note that since the experience buffer is updated with a new observation of the noise at every iteration then by Borel's law of large numbers, we have our probability estimate $\hat{p}(w)$ for the noise converges to the true noise distribution $p(w)$ as $n \rightarrow \infty$, i.e.,

$$\lim_{n \rightarrow \infty} \hat{p}_n(w) = p(w) \text{ for all } w \in \mathcal{W}. \quad (43)$$

Fix an $(s, a) \in \mathcal{S} \times \mathcal{A}$. By part (i) we have the action-value iterates Q_n and Q'_n are bounded for all n . Now we write the iterate $\tilde{Q}_{n,0}^L(s, a)$ in terms of a noise term and a bias term as follows,

$$\tilde{Q}_{n,0}^L(s, a) = \underbrace{\tilde{Q}_{n,0}^L(s, a) - \mathbf{E}[\tilde{Q}_{n,0}^L(s, a)]}_{\text{noise}} + \underbrace{\mathbf{E}[\tilde{Q}_{n,0}^L(s, a)] - \mathbf{E}[Q_{n,0}^L(s, a)]}_{\text{bias}} + \mathbf{E}[Q_{n,0}^L(s, a)]$$

Now, we define the noise term using

$$\xi_n^L(s, a) = \tilde{Q}_{n,0}^L(s, a) - \mathbf{E}[\tilde{Q}_{n,0}^L(s, a)].$$

Also similar to the original proof we define an accumulated noise process started at iteration ν by $W_{\nu,\nu}^L(s, a) = 0$, and

$$W_{n+1,\nu}^L(s, a) = (1 - \alpha_n(s, a)) W_{n,\nu}^L(s, a) + \alpha_n(s, a) \xi_{n+1}^L(s, a) \quad \forall n \geq \nu,$$

which averages noise terms together across iterations. We have $\mathbf{E}[\tilde{Q}_{n,0}^L(s, a) - \mathbf{E}[\tilde{Q}_{n,0}^L(s, a)] | \mathcal{F}_n] = 0$, so Corollary 4.1 applies and it follows that

$$\lim_{n \rightarrow \infty} W_{n,\nu}^L(s, a) = 0 \quad \forall \nu \geq 0.$$

Let $\tilde{\nu}$ be large enough so that $\alpha_n(s, a) \leq 1$ for all $n \geq \tilde{\nu}$. We denote the bias term by

$$\chi_n(s, a) = \mathbf{E}[\tilde{Q}_{n,0}^L(s, a)] - \mathbf{E}[Q_{n,0}^L(s, a)].$$

Since as $n \rightarrow \infty$ we have $\hat{p}_n(w) \rightarrow p(w)$ then the bias due to sampling from the experience buffer $\chi_n(s, a) \rightarrow 0$. Let $\eta > 0$ and $\bar{\nu} \geq \tilde{\nu}$ be such that $|\chi(s, a)| \leq \frac{\eta}{2}$ for all $n \geq \bar{\nu}$ and all (s, a) . We also define

$$Y_{\bar{\nu}}^L(s, a) = \rho, \\ Y_{n+1}^L(s, a) = (1 - \alpha_n(s, a)) Y_n^L(s, a) + \alpha_n(s, a) Q^*(s, a) + \alpha_n(s, a) \frac{\eta}{2}, \quad \forall n \geq \bar{\nu}.$$

It is easy to see that the sequence $Y_n^L(s, a) \rightarrow Q^*(s, a) + \frac{\eta}{2}$. Now we show that our original claim still holds. Claim: for all iterations $n \geq \bar{\nu}$, it holds that

$$L_n(s, a) \leq \min\{\rho, Y_n^L(s, a) + W_{n,\bar{\nu}}^L(s, a)\}.$$

To prove this claim, we proceed by induction on n . For $n = \bar{\nu}$, we have

$$Y_{\bar{\nu}}^L(s, a) = \rho \quad \text{and} \quad W_{\bar{\nu},\bar{\nu}}^L(s, a) = 0,$$

so the statement is true for the base case. We now show that it is true for $n + 1$ given that it holds at n :

$$\begin{aligned}
 L_{n+1}(s, a) &= \min\{\rho, (1 - \alpha_n(s, a)) L_n(s, a) + \alpha_n(s, a) (\tilde{Q}_{n,0}^L(s, a) - \mathbf{E}[\tilde{Q}_{n,0}^L(s, a)] \\
 &\quad + \mathbf{E}[\tilde{Q}_{n,0}^L(s, a)] - \mathbf{E}[Q_{n,0}^L(s, a)] + \mathbf{E}[Q_{n,0}^L(s, a)])\} \\
 &= \min\{\rho, (1 - \alpha_n(s, a)) L_n(s, a) + \alpha_n(s, a) \xi_n^L(s, a) + \alpha_n(s, a) \chi_n(s, a) \\
 &\quad + \alpha_n(s, a) \mathbf{E}[Q_{n,0}^L(s, a)]\} \\
 &\leq \min\{\rho, (1 - \alpha_n(s, a)) (Y_n^L(s, a) + W_{n,\nu_k}^L(s, a)) + \alpha_n(s, a) \xi_n^L(s, a) \\
 &\quad + \alpha_n(s, a) \frac{\eta}{2} + \alpha_n(s, a) Q^*(s, a)\} \\
 &\leq \min\{\rho, Y_{n+1}^L(s, a) + W_{n+1,\bar{\nu}}^L(s, a)\},
 \end{aligned}$$

where the first inequality follows by the induction hypothesis and $\mathbf{E}[Q_{n,0}^L(s, a)] \leq Q^*(s, a)$. Next, since $Y_n^L(s, a) \rightarrow Q^*(s, a) + \frac{\eta}{2}$, $W_{n,\nu_k}^L(s, a) \rightarrow 0$ then if $Q^*(s, a) + \frac{\eta}{2} \leq \rho$, we have

$$\limsup_{n \rightarrow \infty} L_n(s, a) \leq Q^*(s, a) + \frac{\eta}{2}.$$

Otherwise, if $\rho < Q^*(s, a) + \frac{\eta}{2}$ then

$$\limsup_{n \rightarrow \infty} L_n(s, a) \leq \rho < Q^*(s, a) + \frac{\eta}{2}.$$

Therefore, since our choice of (s, a) was arbitrary, it follows that for every $\eta > 0$, there exists some time n' such that $L_n(s, a) \leq Q^*(s, a) + \eta$ for all $(s, a) \in \mathcal{S} \times \mathcal{A}$ and $n \geq n'$.

Using Proposition 1(i), $Q^*(s, a) \leq \mathbf{E}[Q_{n,0}^U(s, a)]$, a similar argument as the above can be used to establish that

$$Q^*(s, a) - \frac{\eta}{2} \leq \liminf_{n \rightarrow \infty} U_n(s, a).$$

Hence, there exists some time n'' such that $Q^*(s, a) - \eta \leq U_n(s, a)$ for all (s, a) and $n \geq n''$. Take n_0 to be some time greater than n' and n'' and the result follows.

The proof of Lemma 1 when using an experience buffer is similar to that given in appendix A.3 so it is omitted. \square

A.6 Proof of Theorem 2

Proof. The proof of both parts (i) and (ii) are like that of Theorem 1 so we omit them. \square

B LBQL with Experience Replay Algorithm

Algorithm 2 LBQL with Experience Replay

Input: Initial estimates $L_0 \leq Q_0 \leq U_0$, batch size K , stepsize rules $\alpha_n(s, a)$, $\beta_n(s, a)$, and noise buffer \mathcal{B} .

Output: Approximations $\{L_n\}$, $\{Q'_n\}$, and $\{U_n\}$.

Set $Q'_0 = Q_0$ and choose an initial state s_0 .

for $n = 0, 1, 2, \dots$ **do**

 Choose an action a_n via some behavior policy (e.g., ϵ -greedy). Observe w_{n+1} .

 Store w_{n+1} in \mathcal{B} and update $\hat{p}_n(w_{n+1})$.

 Perform a standard Q -learning update:

$$Q_{n+1}(s_n, a_n) = Q'_n(s_n, a_n) + \alpha_n(s_n, a_n) \left[r_n(s_n, a_n) + \gamma \max_a Q'_n(s_{n+1}, a) - Q'_n(s_n, a_n) \right].$$

 Sample randomly a sample path $\mathbf{w} = (w_1, w_2, \dots, w_\tau)$ from \mathcal{B} , where $\tau \sim \text{Geom}(1 - \gamma)$.

 Set $\varphi = Q_{n+1}$. Using \mathbf{w} and the current \hat{p}_n compute $\tilde{Q}_0^U(s_n, a_n)$ and $\tilde{Q}_0^L(s_n, a_n)$, using (16) and (17), respectively.

 Update and enforce upper and lower bounds:

$$U_{n+1}(s_n, a_n) = \Pi_{[-\rho, \infty]} \left[U_n(s_n, a_n) + \beta_n(s_n, a_n) \left[\tilde{Q}_0^U(s_n, a_n) - U_n(s_n, a_n) \right] \right],$$

$$L_{n+1}(s_n, a_n) = \Pi_{[\infty, \rho]} \left[L_n(s_n, a_n) + \beta_n(s_n, a_n) \left[\tilde{Q}_0^L(s_n, a_n) - L_n(s_n, a_n) \right] \right],$$

$$Q'_{n+1}(s_n, a_n) = \Pi_{[L_{n+1}(s_n, a_n), U_{n+1}(s_n, a_n)]} [Q_{n+1}(s_n, a_n)]$$

end for

C Implementation Details of LBQL with Experience Replay

We use a noise buffer \mathcal{B} of size κ to record the noise values w that have been previously observed. The buffer \mathcal{B} is used to generate the sample path \mathbf{w} and the batch sample $\{w_1, \dots, w_K\}$ used to estimate the expectation in the penalty function. Here, it is not necessary that the noise space \mathcal{W} is finite. This is also convenient in problems with a large noise support such as the carsharing problem with four stations where we have two sources of noise. Specifically, the noise due to the distribution of the rentals among the stations has a very large support.

In order to reduce the computational requirements of LBQL, the lower and upper bounds updates are done every m steps and only if the difference between the current values of the bounds is greater than some threshold δ .

Since we can easily obtain inner DP results for all (s, a) each time the DP is solved, we perform the upper and lower bound updates for all (s, a) whenever an update is performed (as opposed to just at the current state-action pair). However, only the action-value of the current (s, a) is projected between the lower and upper bounds, so the algorithm is still asynchronous. The pseudo-code of LBQL with experience replay with these changes, is shown in Algorithm 3.

Algorithm 3 LBQL with Experience Replay (Full Details)

Input: Initial estimates $L_0 \leq Q_0 \leq U_0$, batch size K , stepsize rules $\alpha_n(s, a)$, $\beta_n(s, a)$, noise buffer \mathcal{B} of size κ , number of steps between bound updates m , and threshold δ .

Output: Approximations $\{L_n\}$, $\{Q'_n\}$, and $\{U_n\}$.

Set $Q'_0 = Q_0$ and choose an initial state s_0 .

for $n = 0, 1, 2, \dots$ **do**

 Choose an action a_n via some behavior policy (e.g., ϵ -greedy). Observe w_{n+1} .

 Store w_{n+1} in \mathcal{B} .

 Perform a standard Q-learning update:

$$Q_{n+1}(s_n, a_n) = Q'_n(s_n, a_n) + \alpha_n(s_n, a_n) \left[r_n(s_n, a_n) + \gamma \max_a Q'_n(s_{n+1}, a) - Q'_n(s_n, a_n) \right].$$

if $n \geq \kappa$ and $n \bmod m = 0$ and $|U_n(s_n, a_n) - L(s_n, a_n)| > \delta$ **then**

 Sample randomly a batch $\mathcal{D} = \{w_1, w_2, \dots, w_K\}$ and a sample path $\mathbf{w} = \{w_1, w_2, \dots, w_\tau\}$ from \mathcal{B} , where $\tau \sim \text{Geom}(1 - \gamma)$.

 Set $\varphi = Q_{n+1}$. Using \mathbf{w} and \mathcal{D} , compute $\hat{Q}_0^U(s, a)$ and $\hat{Q}_0^L(s, a)$ for all $(s, a) \in \mathcal{S} \times \mathcal{A}$, using (9) and (10), respectively.

 For all $(s, a) \in \mathcal{S} \times \mathcal{A}$, update upper and lower bounds:

$$U_{n+1}(s_n, a_n) = \Pi_{[-\rho, \infty]} \left[U_n(s_n, a_n) + \beta_n(s_n, a_n) [\hat{Q}_0^U(s_n, a_n) - U_n(s_n, a_n)] \right],$$

$$L_{n+1}(s_n, a_n) = \Pi_{[\infty, \rho]} \left[L_n(s_n, a_n) + \beta_n(s_n, a_n) [\hat{Q}_0^L(s_n, a_n) - L_n(s_n, a_n)] \right],$$

end if

 Enforce upper and lower bounds:

$$Q'_{n+1}(s_n, a_n) = \Pi_{[L_{n+1}(s_n, a_n), U_{n+1}(s_n, a_n)]} [Q_{n+1}(s_n, a_n)]$$

end for

D Numerical Experiments Details

Let $\nu(s, a)$ and $\nu(s)$ be the number of times state-action pair (s, a) and state s , have been visited, respectively. For all algorithms, a polynomial learning rate $\alpha_n(s, a) = 1/\nu_n(s, a)^r$ is used, where $r = 0.5$. Polynomial learning rates have been shown to have a better performance than linear learning rates (van Hasselt, 2010).

We use a discount factor of $\gamma = 0.95$ for the pricing car-sharing/stormy gridworld problems, $\gamma = 0.9$ for the windy gridworld problem and $\gamma = 0.99$ for the repositioning problem. Moreover, we use an ϵ -greedy exploration strategy such that $\epsilon(s) = 1/\nu(s)^e$, where e is 0.4 for the four-stations pricing car-sharing problem and 0.5 for all the other problems. For the car-sharing/windy gridworld problems, the initial state-action values are chosen randomly such that $L_0(s, a) \leq Q_0(s, a) \leq U_0(s, a)$ where

$$L_0(s, a) = -R_{\max}/(1 - \gamma) \quad \text{and} \quad U_0(s, a) = R_{\max}/(1 - \gamma)$$

for all (s, a) . For the stormy gridworld problem, we set the initial state-action values to zero (we find that a random initialization caused all algorithms except LBQL to perform extremely poorly).

We report LBQL parameters used in our numerical experiments in Table 1. Note that for a fair comparison, the parameter K of bias-corrected Q-learning algorithm is taken equal to K of LBQL in all experiments. In addition, the κ steps used to create the buffer for LBQL are included in the total number of steps taken. Results of the gridworld and car-sharing problems are averaged over 50 and 10 runs, respectively. All experiments were run on a 3.5 GHz Intel Xeon processor with 32 GB of RAM workstation.

Table 1: LBQL parameters.

Problem	Parameter				
	β	κ	K	m	δ
2-CS-R	0.01	40	20	10	0.01
2-CS	0.01	40	20	15	0.01
4-CS	0.01	1000	20	200	0.01
WG	0.2	100	10	10	0.01
SG	0.2	500	20	20	0.05

A detailed description of the environments is given in the next two sections.

D.1 Gridworld Examples

First we consider, windy gridworld, a well-known variant of the standard gridworld problem discussed in (Sutton & Barto, 2018). Then we introduce, stormy gridworld, a new environment that is more complicated than windy gridworld. The environments are summarized below.

Windy Gridworld. The environment is a 10×7 gridworld, with a cross wind pointing *upward*, (Sutton & Barto, 2018). The default *wind* values corresponding to each of the 10 columns are $\{0, 0, 0, 1, 1, 1, 2, 2, 1, 0\}$. Allowable actions are {up, right, down, left}. If the agent happens to be in a column whose wind value is different from zero, the resulting next states are shifted upward by the “wind” whose intensity is stochastic, varying from the given values in each column by $\{-1, 0, 1\}$ with equal probability. Actions that corresponds to directions that takes the agent off the grid leave the location of the agent unchanged. The start and goal states are $(3, 1)$ and $(3, 8)$, respectively. The reward is -1 until the goal state is reached, where the reward is 0 thereafter.

Stormy Gridworld. Consider the stochastic windy gridworld environment. Now, however, we allow the wind to blow half the time as before and the other half it can blow from any of the three other directions. The horizontal wind values corresponding to each row from top to bottom are given by $\{0, 0, 1, 1, 1, 1, 0\}$. Also, it can randomly rain with equal probability in any of the central states that are more than two states away from the edges of the grid. The start and goal states are $(3, 1)$ and $(3, 10)$ respectively. Rain creates a puddle which affects the state itself and all of its neighboring states. The reward is as before except when the agent enters a puddle state the reward is -10 .

D.2 Car-sharing Benchmark Examples

In this section, we give the detailed formulations of the two variants of the car-sharing benchmark, repositioning and pricing. The essential difference is that in the pricing version, the decision maker ‘‘repositions’’ by setting prices to induce directional demand (but does not have full control since this demand is random).

D.2.1 REPOSITIONING BENCHMARK FOR CAR-SHARING

We consider the problem of repositioning cars for a two stations car-sharing platform, (He et al., 2019). The action is the number of cars to be repositioned from one station to the other, before random demand is realized. Since repositioning in both directions is never optimal, we use $r > 0$ to denote the repositioned vehicles from station 1 to 2 and $r < 0$ to denote repositioning from station 2 to 1. The stochastic demands at time t are $D_{1,t}$ and $D_{2,t}$ for stations 1 and 2 respectively, are i.i.d., discrete uniform, each supported on $\{3, \dots, 9\}$. The rental prices are $p_1 = 3.5$ and $p_2 = 4$ for stations 1 and 2, respectively. All rentals are one-way (i.e., rentals from station 1 end up at station 2, and vice-versa). The goal is to maximize profit, where unmet demands are charged a lost sales cost $\rho_1 = \rho_2 = 2$ and repositioning cost $c_1 = 1$ for cars reposition from station 1 to 2 and $c_2 = 1.5$ for cars repositioned from station 2 to 1. We assume a total of $\bar{s} = 12$ cars in the system and formulate the problem as an MDP, with state $s_t \in \mathcal{S} = \{0, 1, \dots, 12\}$ representing the number of cars at station 1 at beginning of period t . We denote by $V^*(s_t)$ the optimal value function starting from state s_t . The Bellman recursion is:

$$\begin{aligned}
 V^*(s_t) = \max_{s_t - \bar{s} \leq r_t \leq s_t} \mathbf{E} & \left[\sum_{i \in \{1,2\}} p_i \omega_{it}(D_{i,t+1}) - \sum_{i \in \{1,2\}} \rho_i (D_{i,t+1} - \omega_{it}(D_{i,t+1})) \right. \\
 & \left. - c_1 \max(r_t, 0) + c_2 \min(r_t, 0) + \gamma V^*(s_{t+1}) \right], \quad (44) \\
 \omega_{1t}(D_{1,t+1}) &= \min(D_{1,t+1}, s_t - r_t), \\
 \omega_{2t}(D_{2,t+1}) &= \min(D_{2,t+1}, \bar{s} - s_t + r_t), \\
 s_{t+1} &= s_t - r_t + \omega_{2t}(D_{2,t+1}) - \omega_{1t}(D_{1,t+1}),
 \end{aligned}$$

where $\gamma \in (0, 1)$ is a discount factor. The repositioning problem for two stations is illustrated in Figure 6A. The nodes represent stations, solid arcs represent fulfilled demands, and dashed arcs represent repositioned vehicles.

D.2.2 PRICING BENCHMARK FOR CAR-SHARING

Suppose that a vehicle sharing manager is responsible for setting the rental price for the vehicles at the beginning of each period in an infinite planning horizon. We model a car sharing system with N stations. The goal is to optimize the prices to set for renting a car at each of the N stations; let the price at station i and time t be p_{it} for $i \in [N] := \{1, \dots, N\}$. Demands are nonnegative, independent and depends on the vehicle renting price according to a stochastic demand function

$$D_{it}(p_{it}, \epsilon_{i,t+1}) := \kappa_i(p_{it}) + \epsilon_{i,t+1},$$

where $D_{it}(p_{it}, \epsilon_{i,t+1})$ is the demand in period t , $\epsilon_{i,t+1}$ are random perturbations that are revealed at time $t + 1$ and $\kappa_i(p_{it})$ is a deterministic demand function of the price p_{it} that is set at the beginning of period t at station $i \in [N]$. The random variables $\epsilon_{i,t+1}$ are independent with $\mathbf{E}[\epsilon_{i,t+1}] = 0$ without loss of generality. Furthermore, we assume that the expected demand $\mathbf{E}[D_{it}(p_{it}, \epsilon_{i,t+1})] = \kappa_i(p_{it}) < \infty$ is strictly decreasing in the rental price p_{it} which is restricted to a set of feasible price levels $[\underline{p}_i, \bar{p}_i]$ for all $i \in [N]$, where $\underline{p}_i, \bar{p}_i$ are the minimum and the maximum prices that can be set at station i , respectively. This assumption implies a one-to-one correspondence between the rental price p_{it} and the expected demand $d_{it} \in \mathcal{D} := [\underline{d}_i, \bar{d}_i]$ for all $p_{it} \in [\underline{p}_i, \bar{p}_i]$ where $\underline{d}_i = \kappa_i(\bar{p}_i)$ and $\bar{d}_i = \kappa_i(\underline{p}_i)$.

The problem can be formulated as an MDP with state \mathbf{s}_t , which is a vector whose components represent the number of available cars at each of the N stations at beginning of period t . The state space is \mathcal{S}^{N-1} with $\mathcal{S} = \{0, 1, \dots, \bar{s}\}$ and \bar{s} is the maximum number of cars in the vehicle sharing system. We assume that a customer at station i goes to station j with probability ϕ_{ij} for all $i, j \in [N]$. Let $Y_{ik,t+1}$ be a random variable taking values in $[N]$ that represents the random destination of customer k at station i , which is only observed at the beginning of period $t + 1$. We have $Y_{ik,t+1} = j$ with probability ϕ_{ij} , so $Y_{ik,t+1}$ are i.i.d. for each customer k . Denote by l_{ij} the distance from station i to j , for all $i, j \in [N]$. We penalize unmet demands by a lost sales unit cost ρ_i , $i \in [N]$. The decision vector is $\mathbf{p}_t = \{p_{it} \in [\underline{p}_i, \bar{p}_i], \forall i \in [N]\}$.

Let $V^*(\mathbf{s}_t)$ be the revenue-to-go function with number of available vehicles \mathbf{s}_t . Thus, we have the Bellman recursion

$$\begin{aligned}
 V^*(\mathbf{s}_t) &= \max_{\mathbf{P}^t} \mathbf{E} \left[\sum_{i \in [N]} p_{it} \sum_{j \in [N]} l_{ij} \omega_{ijt}(\epsilon_{i,t+1}) - \sum_{i \in [N]} \rho_i \left(\kappa_i(p_{it}) + \epsilon_{i,t+1} - \omega_{it}(\epsilon_{i,t+1}) \right) + \gamma V^*(\mathbf{s}_{t+1}) \right] \\
 \omega_{it}(\epsilon_{i,t+1}) &= \min(\kappa_i(p_{it}) + \epsilon_{i,t+1}, s_{it}) \quad \forall i \in [N], \\
 \omega_{ijt}(\epsilon_{i,t+1}) &= \sum_{k=1}^{\omega_{it}(\epsilon_{i,t+1})} \mathbf{1}_{\{Y_{ik,t+1}=j\}} \quad \forall i, j \in [N], \\
 s_{i,t+1} &= s_{it} + \sum_{j \in [N]} \omega_{jit}(\epsilon_{i,t+1}) - \omega_{it}(\epsilon_{i,t+1}), \quad \forall i \in [N],
 \end{aligned} \tag{45}$$

where $\gamma \in (0, 1)$ is a discount factor. Note that the MDP in (45) can be reformulated using the action-value function $Q(s_t, p_t)$ instead of $V(s_t)$. The quantity $\omega_{it}(\epsilon_{i,t+1})$ represents the total fulfilled customer trips from station i at time t for a given realization of the noise $\epsilon_{i,t+1}$. Notice that in (45) there are two sources of randomness: the noise due to stochastic demand represented by ϵ_i , for all $i \in [N]$ and the noise due to the random distribution of fulfilled rentals between the stations, i.e., due to the random variables $Y_{i1}, \dots, Y_{i\omega_{it}(\epsilon_{i,t+1})}$ for all $i \in [N]$. Due to the high dimensionality involved in the state, action, and noise spaces, solving (45) is computationally challenging.

Spatial Pricing in Two-Location Car-sharing. We first consider the pricing problem on two stations and 12 cars in total. The state space is $\mathcal{S} = \{0, 1, \dots, 12\}$ representing the number of cars at station 1. All rentals are one-way. The prices, at each period t , are restricted to $p_{1t} \in [1, 6]$ and $p_{2t} \in [1, 7]$. The stochastic demand functions at period t are given by: $D_{1t}(p_{1t}, \epsilon_{1,t+1}) := 9 - p_{1t} + \epsilon_{1,t+1}$ and $D_{2t}(p_{2t}, \epsilon_{2,t+1}) := 10 - p_{2t} + \epsilon_{2,t+1}$ for stations 1 and 2 respectively. The random variables $\epsilon_{1,t+1}$ and $\epsilon_{2,t+1}$ are independent, discrete uniform, each supported on $\{-3, -2, \dots, 3\}$. We use the discretized expected demands, as our actions: $d_{1t} \in \{3, \dots, 8\}$ and $d_{2t} \in \{3, \dots, 9\}$. The lost sales cost is 2 at both stations.

Spatial Pricing in Four-Location Car-sharing. Consider the car-sharing problem for four stations with $\bar{s} = 20$ cars and $d_{it} \in \{3, 4\}$ for each station. In total there are 1771 states and 16 actions. The random variables $\epsilon_{i,t+1}$ are independent, discrete uniform, each supported on $\{-3, -2, \dots, 3\}$. We consider both one way and return trips at each station. Figure 6B shows an illustration of the stations (nodes) and the rentals between the stations (arcs). The probabilities $\phi_{ij} = 0.25$ for all $i, j \in \{1, 2, 3, 4\}$ and the lost sales costs (ρ_i) are 1.7, 1.2, 1.5, 2 at stations 1, 2, 3, 4, respectively. The distance between the stations are taken such that $l_{ij} = 1$ if $i = j$, and the other distances being symmetrical, meaning $l_{ij} = l_{ji}$ with $l_{12} = 1.8$, $l_{13} = 1.5$, $l_{14} = 1.4$, $l_{23} = 1.6$, $l_{24} = 1.1$, and $l_{34} = 1.2$.




Figure 6: Illustrations of the repositioning and pricing car-sharing problems.

D.3 Sensitivity Analysis

We also do sensitivity analysis on the five algorithms with respect to the learning rate and exploration parameters r and e for the car-sharing problem with two stations. Here, r controls the polynomial learning rate defined by, $\alpha_n(s, a) = 1/\nu_n(s, a)^r$ and e controls the ϵ -greedy exploration strategy, where ϵ is annealed according to $\epsilon(s) = 1/\nu(s)^e$. We use $\nu(s, a)$ and $\nu(s)$ to denote the number of times a state-action pair (s, a) and state s , have been visited, respectively. We report our results in Table 2. These results show the average number of iterations and CPU time until each algorithm first reach 50%, 20%, 5%, 1% relative error for each case of the parameters e and r while keeping all other parameters as before. The “-” indicates that the corresponding % relative error for the corresponding case was not achieved during the course of training. The values in the table are obtained by averaging five independent runs for each case. Except for the few cases where BCQL performs slightly better, LBQL once again drastically outperforms the other algorithms and exhibits *robustness* against the learning rate and exploration parameters, an important practical property that the other algorithms seem to lack.

The effect of varying parameters m and K of LBQL is presented in Figure 7. These plots are obtained by tuning parameters $m \in \{1, 10, 50, 150, 200\}$ and $K \in \{1, 5, 10, 100, 1000\}$ of LBQL algorithm in the car-sharing problem with two stations. All other parameters are kept the same as before. Figures 7A and 7C show the mean total reward with a 95% CI. Figures 7B and 7D show the mean and 95% CI of the relative error given by: $\|V_n - V^*\|_2 / \|V^*\|_2$. The results are obtained from 10 independent runs. Using larger values of m reduces the strength of LBQL in both performance and relative error measures as shown in Figures 7A and 7B. This is expected since the effect of the bounds fades as we update the bounds less frequently. Interestingly, we can see from the performance plot that $m = 10$ strikes a good balance between how often to do the bounds and Q-learning updates and achieves a performance that is slightly better and more stable than that of $m = 1$ after about half of the training process (50,000 steps). In terms of the sample size K , Figures 7C and 7D clearly show that larger values of K improve the performance of LBQL in terms of performance and relative error measures. This is not unexpected because a larger sample yields a better approximation of the penalty.

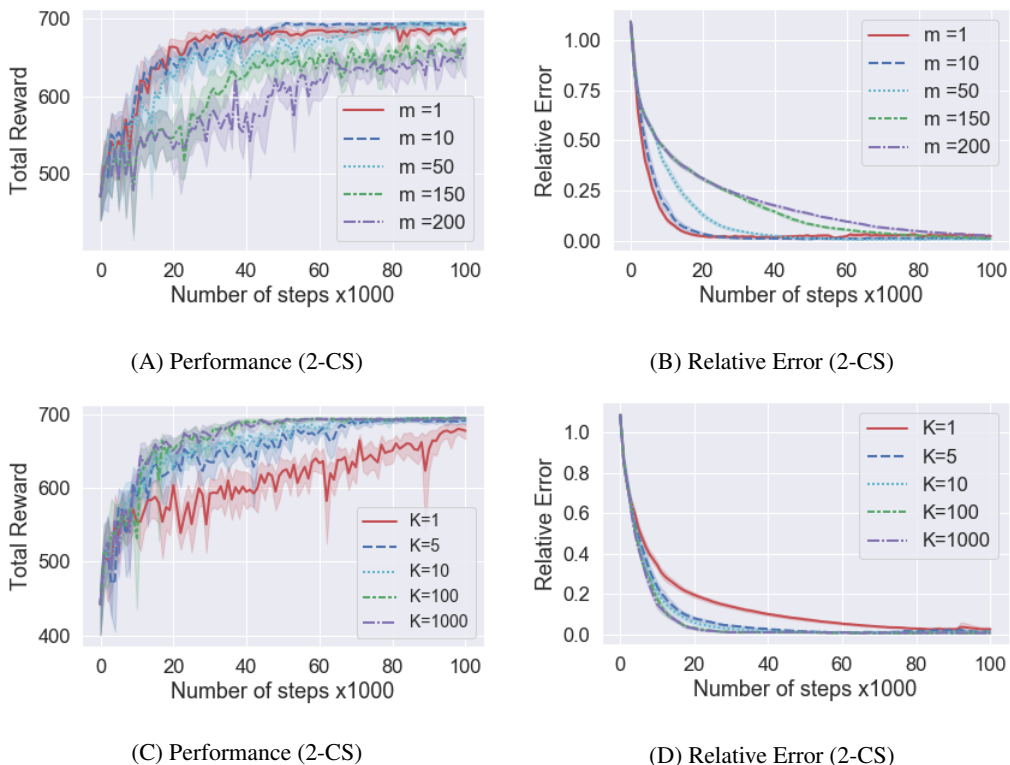


Figure 7: Plots showing the effect of tuning the parameter M and K of LBQL algorithm.

Lookahead-Bounded Q-Learning

Table 2: Computational results for different exploration & learning rate parameters. Bold numbers indicate the best performing algorithm.

		% Relative error								
e	r	50%		20%		5%		1%		
		n	t (s)	n	t (s)	n	t (s)	n	t (s)	
LBQL	0.4	0.5	3,672.6	1.9	9,323.6	4.6	18,456.6	8.9	33,054.0	15.7
		0.6	3,632.0	1.7	9,147.4	4.4	18,270.0	8.6	39,624.8	18.6
		0.7	3,725.2	1.8	9,087.4	4.3	18,217.8	8.6	41,941.2	19.7
		0.8	3,698.2	1.8	9,321.8	4.4	20,860.0	9.9	53,752.6	25.6
		0.9	3,992.2	1.9	10,119.2	4.9	23,070.2	11.0	80,252.8	38.4
	0.5	0.5	3,316.0	1.6	8,040.2	3.8	15,050.2	7.2	27,912.8	13.2
		0.6	3,514.4	1.7	8,529.4	4.0	16,595.6	7.9	36,100.2	17.0
		0.7	3,531.8	1.7	8,712.8	4.1	17,835.6	8.6	46,010.0	21.9
		0.8	3,449.2	1.6	8,571.8	4.0	18,152.6	8.5	65,007.6	30.4
		0.9	3,398.4	1.6	8,346.4	3.9	18,844.8	8.8	99,820.2	46.6
	0.6	0.5	2,877.4	1.3	7,129.0	3.3	13,046.0	6.2	23,822.0	11.2
		0.6	3,182.8	1.5	8,066.0	3.8	15,421.4	7.3	33,286.0	15.6
		0.7	2,979.4	1.4	7,625.6	3.6	15,414.6	7.2	34,238.0	15.9
		0.8	3,272.6	1.6	8,431.0	4.1	17,809.2	8.5	114,032.8	54.2
		0.9	3,185.6	1.5	8,480.0	4.1	19,242.4	9.2	123,331.2	58.8
BCQL	0.4	0.5	3,200.8	1.0	22,455.0	7.0	65,329.0	20.3	107,785.6	33.7
		0.6	4,618.2	1.5	43,724.6	13.6	159,662.6	49.6	292,421.0	34.9
		0.7	8,059.4	2.5	123,484.2	38.4	-	-	-	-
		0.8	17,287.0	5.3	-	-	-	-	-	-
		0.9	67,162.2	20.9	-	-	-	-	-	-
	0.5	0.5	2,209.6	0.7	15,604.0	4.9	48,715.6	15.3	80,317.4	25.2
		0.6	3,274.2	1.0	31,422.8	9.8	124,319.6	38.7	243,101.4	75.6
		0.7	5,619.6	1.8	89,857.0	27.8	-	-	-	-
		0.8	11,417.0	3.6	-	-	-	-	-	-
		0.9	42,605.4	13.1	-	-	-	-	-	-
	0.6	0.5	1,830.4	0.6	11,639.6	3.6	35,763.0	11.1	61,249.0	19.0
		0.6	2,612.4	0.8	23,571.6	7.4	92,101.4	28.8	177,127.6	55.5
		0.7	4,371.2	1.3	66,526.0	20.5	-	-	-	-
		0.8	9,028.2	2.8	297,368.6	17.9	-	-	-	-
		0.9	31,673.6	9.8	-	-	-	-	-	-
SQL	0.4	0.5	7,750.6	1.9	37,889.8	9.1	93,820.0	22.5	141,171.0	33.8
		0.6	11,329.4	2.8	75,364.0	18.3	233,422.0	56.4	-	-
		0.7	20,131.4	4.8	212,767.0	51.0	-	-	-	-
		0.8	46,986.8	11.3	-	-	-	-	-	-
		0.9	182,890.0	43.8	-	-	-	-	-	-
	0.5	0.5	6,122.2	1.5	30,944.8	7.4	79,167.4	19.0	120,527.8	29.0
		0.6	9,166.6	2.2	62,540.6	14.9	201,822.4	48.2	-	-
		0.7	15,835.6	3.8	174,233.6	42.0	-	-	-	-
		0.8	36,548.8	8.7	-	-	-	-	-	-
		0.9	157,029.0	37.7	-	-	-	-	-	-
	0.6	0.5	4,984.0	1.2	24,989.0	6.0	64,605.4	15.4	98,554.6	23.6
		0.6	7,396.2	1.8	50,282.4	12.0	165,574.2	39.8	-	-
		0.7	13,018.8	3.1	143,142.6	34.1	-	-	-	-
		0.8	29,201.0	6.9	-	-	-	-	-	-
		0.9	122,335.6	29.2	-	-	-	-	-	-

(continued on next page)

