

A. Algorithm Pseudocodes

This section contains the pseudocodes for the algorithms used in the experiments in this paper:

- Section 3 - Empirical Results for Control:
Differential Q-learning and RVI Q-learning
- Section 5 - Empirical Results for Prediction:
Differential TD-learning and Average Cost TD learning
- Section 6 - Estimating the Actual Differential Value Function:
Centered Differential Q-learning

Algorithm 1: Differential Q-learning (one-step off-policy control)

Input: The policy b to be used (e.g., ϵ -greedy)

Algorithm parameters: step-size parameters α, η

```

1 Initialize  $Q(s, a) \forall s, a; \bar{R}$  arbitrarily (e.g., to zero)
2 Obtain initial  $S$ 
3 while still time to train do
4    $A \leftarrow$  action given by  $b$  for  $S$ 
5   Take action  $A$ , observe  $R, S'$ 
6    $\delta = R - \bar{R} + \max_a Q(S', a) - Q(S, A)$ 
7    $Q(S, A) = Q(S, A) + \alpha\delta$ 
8    $\bar{R} = \bar{R} + \eta\alpha\delta$ 
9    $S = S'$ 
10 end
11 return  $Q$ 

```

Algorithm 2: RVI Q-learning (one-step off-policy control)

Input: The policy b to be used (e.g., ϵ -greedy)

Algorithm parameters: step-size parameter α

```

1 Initialize  $Q(s, a) \forall s, a$  arbitrarily (e.g., to zero)
2 Choose function  $f(Q)$  (e.g., a single reference state-action pair —  $f(Q) = Q(s_0, a_0)$ )
3 Obtain initial  $S$ 
4 while still time to train do
5    $A \leftarrow$  action given by  $b$  for  $S$ 
6   Take action  $A$ , observe  $R, S'$ 
7    $\delta = R - f(Q) + \max_a Q(S', a) - Q(S, A)$ 
8    $Q(S, A) = Q(S, A) + \alpha\delta$ 
9    $S = S'$ 
10 end
11 return  $Q$ 

```

Algorithm 3: Differential TD-learning (one-step off-policy prediction)

Input: The policy π to be evaluated, and b to be used

Algorithm parameters: step-size parameters α, η

1 Initialize $V(s) \forall s, \bar{R}$ arbitrarily (e.g., to zero)

2 **while** *still time to train* **do**

3 $A \leftarrow$ action given by b for S

4 Take action A , observe R, S'

5 $\delta = R - \bar{R} + V(S') - V(S)$

6 $\rho = \pi(A|S) / b(A|S)$

7 $V(S) = V(S) + \alpha\rho\delta$

8 $\bar{R} = \bar{R} + \eta\alpha\rho\delta$

9 $S = S'$

10 **end**

11 return V

Algorithm 4: Average Cost TD-learning (one-step on-policy prediction)

Input: The policy π to be evaluated

Algorithm parameters: step-size parameters α, η

1 Initialize $V(s) \forall s, \bar{R}$ arbitrarily (e.g., to zero)

2 **while** *still time to train* **do**

3 $A \leftarrow$ action given by π for S

4 Take action A , observe R, S'

5 $\delta = R - \bar{R} + V(S') - V(S)$

6 $V(S) = V(S) + \alpha\delta$

7 $\bar{R} = \bar{R} + \eta\alpha(R - \bar{R})$

8 $S = S'$

9 **end**

10 return V

Algorithm 5: Centered Differential Q-learning

Input: The policy b to be used (e.g., ϵ -greedy)

Algorithm parameters: step-size parameters $\alpha, \eta, \beta, \kappa$

1 Initialize $Q(s, a), F(s, a) \forall s, a; \bar{R}, \bar{Q}$ arbitrarily (e.g., to zero)

2 Obtain initial S

3 **while** *still time to train* **do**

4 $A \leftarrow$ action given by b for S

5 Take action A , observe R, S'

6 $\delta = R - \bar{R} + \max_a Q(S', a) - Q(S, A)$

7 $Q(S, A) = Q(S, A) + \alpha\delta$

8 $\bar{R} = \bar{R} + \eta\alpha\delta$

9 $\Delta = Q(S, A) - \bar{Q} + F(S', \operatorname{argmax}_a Q(S', a)) - F(S, A)$

10 $F(S, A) = F(S, A) + \beta\Delta$

11 $\bar{Q} = \bar{Q} + \kappa\beta\Delta$

12 $S = S'$

13 **end**

14 return $Q - \bar{Q}e$, where e is a $|S \times \mathcal{A}|$ vector of all ones.

B. Convergence Proofs

In this section, we present the convergence proofs of Differential Q-learning and Differential Q-planning in subsection B.1, of Differential TD-learning and Differential TD-planning in subsection B.2, and that of the centered version of these algorithms in subsection B.3.

For convenience, the following notations are used for all the proofs:

- Given any vector x , $\sum x$ denotes sum of all elements in x . Formally, $\sum x \doteq \sum_i x(i)$.
- e denotes an all-ones vector, whose length may be $|\mathcal{S} \times \mathcal{A}|$ or $|\mathcal{S}|$ depending on the context.
- Finally, $\exp(\cdot)$ is used instead of $e^{(\cdot)}$ to denote the exponential function.

B.1. Proof of Differential Q-learning and Differential Q-planning

In this section, we analyze a general algorithm that includes both Differential Q-learning and Differential Q-planning cases. We call it *General Differential Q*. We first formally define it and then explain why both Differential Q-learning and Differential Q-planning are special cases of General Differential Q. We then provide assumptions and the convergence theorem of General Differential Q. The theorem would lead to convergence of Differential Q-learning and Differential Q-planning. Finally, we provide a proof for the theorem.

Given a MDP $\mathcal{M} \doteq (\mathcal{S}, \mathcal{A}, \mathcal{R}, p)$, for each state $s \in \mathcal{S}$ action $a \in \mathcal{A}$ and discrete step $n \geq 0$, let $R_n(s, a), S'_n(s, a) \sim p(\cdot, \cdot | s, a)$ denote a sample of resulting state and reward. We hypothesize a set-valued process $\{Y_n\}$ taking values in the set of nonempty subsets of $\mathcal{S} \times \mathcal{A}$ with the interpretation: $Y_n = \{(s, a) : (s, a) \text{ component of } Q \text{ was updated at time } n\}$. Let $\nu(n, s, a) \doteq \sum_{k=0}^n I\{(s, a) \in Y_k\}$, where I is the indicator function. Thus $\nu(n, s, a)$ = the number of times the (s, a) component was updated up to step n . The update rules of General Differential Q are

$$Q_{n+1}(s, a) \doteq Q_n(s, a) + \alpha_{\nu(n, s, a)} \delta_n(s, a) I\{(s, a) \in Y_n\}, \quad \forall s \in \mathcal{S}, a \in \mathcal{A}, \quad (\text{B.1})$$

$$\bar{R}_{n+1} \doteq \bar{R}_n + \eta \sum_{s, a} \alpha_{\nu(n, s, a)} \delta_n(s, a) I\{(s, a) \in Y_n\}, \quad (\text{B.2})$$

where

$$\delta_n(s, a) \doteq R_n(s, a) - \bar{R}_n + \max_{a'} Q_n(S'_n(s, a), a') - Q_n(s, a). \quad (\text{B.3})$$

Here $\alpha_{\nu(n, s, a)}$ is the stepsize at step n for state–action pair (s, a) . The quantity $\alpha_{\nu(n, s, a)}$ depends on the sequence $\{\alpha_n\}$, which is an algorithmic design choice, and also depends on the visitation of state–action pairs $\nu(n, s, a)$. To obtain the stepsize, the algorithm could maintain a $|\mathcal{S} \times \mathcal{A}|$ -size table counting the number of visitations to each state–action pair, which is exactly $\nu(\cdot, \cdot, \cdot)$. Then the stepsize $\alpha_{\nu(n, s, a)}$ can be obtained as long as the sequence $\{\alpha_n\}$ is specified.

Now we show Differential Q-learning is a special case of General Differential Q. Consider a sequence of real experience $\dots, S_t, A_t, R_{t+1}, S_{t+1}, \dots$. By choosing step $n = \text{time step } t$,

$$Y_t(s, a) = 1, \text{ if } s = S_t, a = A_t, \\ Y_t(s, a) = 0 \text{ otherwise,}$$

and $S'_n(S_t, A_t) = S_{t+1}, R_n(S_t, A_t) = R_{t+1}$, update rules (B.1), (B.2), and (B.3) become

$$Q_{t+1}(S_t, A_t) \doteq Q_t(S_t, A_t) + \alpha_{\nu(t, S_t, A_t)} \delta_t, \text{ and } Q_{t+1}(s, a) \doteq Q_t(s, a), \forall s \neq S_t, a \neq A_t, \quad (\text{B.4})$$

$$\bar{R}_{t+1} \doteq \bar{R}_t + \eta \alpha_{\nu(t, S_t, A_t)} \delta_t, \quad (\text{B.5})$$

$$\delta_t \doteq R_{t+1} - \bar{R}_t + \max_{a'} Q_t(S_{t+1}, a') - Q_t(S_t, A_t), \quad (\text{B.6})$$

which are Differential Q-learning's update rules with stepsize at time t being $\alpha_{\nu(t, S_t, A_t)}$.

Similarly we can show Differential Q-planning is a special case of General Differential Q. Consider a sequence of simulated experience $\dots, \hat{S}_n, \hat{A}_n, \hat{R}_n, \hat{S}'_n, \dots$. By choosing step n to be the planning step,

$$Y_n(s, a) = 1, \text{ if } s = \hat{S}_n, a = \hat{A}_n, \\ Y_n(s, a) = 0, \text{ otherwise,}$$

and $S'_n(\hat{S}_n, \hat{A}_n) = \hat{S}'_n$, $R_n(\hat{S}_n, \hat{A}_n) = \hat{R}_n$, update rules (B.1), (B.2), and (B.3) become

$$Q_{n+1}(\hat{S}_n, \hat{A}_n) \doteq Q_n(\hat{S}_n, \hat{A}_n) + \alpha_{\nu(n, \hat{S}_n, \hat{A}_n)} \delta_n, \text{ and } Q_{n+1}(s, a) \doteq Q_n(s, a), \forall s \neq \hat{S}_n, a \neq \hat{A}_n, \quad (\text{B.7})$$

$$\bar{R}_{n+1} \doteq \bar{R}_n + \eta \alpha_{\nu(n, \hat{S}_n, \hat{A}_n)} \delta_n, \quad (\text{B.8})$$

$$\delta_n \doteq \hat{R}_{n+1} - \bar{R}_n + \max_{a'} Q_n(\hat{S}_{n+1}, a') - Q_n(\hat{S}_n, \hat{A}_n), \quad (\text{B.9})$$

which are Differential Q-planning's update rules with stepsize $\alpha_{\nu(n, \hat{S}_n, \hat{A}_n)}$.

We now specify assumptions on General Differential Q, which are required by our convergence theorem.

Assumption B.1 (Communicating Assumption). *The MDP \mathcal{M} has a single communicating class, that is, each state in \mathcal{M} is accessible from every other state under some deterministic stationary policy.*

Assumption B.2 (Action-Value Function Uniqueness). *There exists a unique solution of q only up to a constant in (5).*

Assumption B.3 (Stepsize Assumption). $\alpha_n > 0$, $\sum_{n=0}^{\infty} \alpha_n = \infty$, $\sum_{n=0}^{\infty} \alpha_n^2 < \infty$.

Assumption B.4 (Asynchronous Stepsize Assumption A). *Let $[\cdot]$ denote the integer part of (\cdot) , for $x \in (0, 1)$,*

$$\sup_i \frac{\alpha_{[xi]}}{\alpha_i} < \infty$$

and

$$\frac{\sum_{j=0}^{[yi]} \alpha_j}{\sum_{j=0}^i \alpha_j} \rightarrow 1$$

uniformly in $y \in [x, 1]$.

Assumption B.5 (Asynchronous Stepsize Assumption B). *There exists $\Delta > 0$ such that*

$$\liminf_{n \rightarrow \infty} \frac{\nu(n, s, a)}{n+1} \geq \Delta,$$

a.s., for all $s \in \mathcal{S}$, $a \in \mathcal{A}$. Furthermore, for all $x > 0$, let

$$N(n, x) = \min \left\{ m > n : \sum_{i=n+1}^m \alpha_i \geq x \right\},$$

the limit

$$\lim_{n \rightarrow \infty} \frac{\sum_{i=\nu(n, x, s, a)}^{\nu(N(n, x), s, a)} \alpha_i}{\sum_{i=\nu(n, s', a')}^{\nu(N(n, x), s', a')} \alpha_i}$$

exists a.s. for all s, s', a, a' .

We now explain the meanings of these assumptions.

Assumption B.1 is the standard communicating assumption for the MDP. If this is not satisfied (i.e., there exist states from which it is impossible to get back to the others), no learning algorithm can be guaranteed to learn differential value function up to an additive constant for any policy in that MDP using a single stream of experience. The reward rate of a learned policy can still converge to the optimal reward rate under a slightly weaker *weakly communicating* assumption, which assumes that the MDP has a single communicating class and some additional *transient* states. Whenever a weakly communicating MDP starts from a transient state, eventually it will never visit that state again under any policy. The differential values of states in the communicating class can be learned well using some algorithms but that of transient states can not. Both the theorem and proof of convergence under the weakly communicating assumption would require distinguishing between these two class of states. For a simpler analysis, we use the communicating assumption here. In case of our control planning problem, transient states can appear in the simulated experience for an infinite number of times, and thus differential values of transient states can be learned accurately. Therefore the communicating assumption for the planning algorithm can be

replaced by the more general weakly communicating assumption. However, to give a simple theorem and proof which cover both our learning and planning algorithms, we choose to present our result using the communicating MDP assumption.

Assumption B.2 is required by average-reward learning and planning algorithms to guarantee convergence of estimates of Q to a unique solution (up to a constant). A necessary and sufficient condition for Assumption B.2 is provided by Schweitzer & Federgruen (1978). The condition is that there exists a randomized stationary optimal policy that induces a single recurrent class of states \mathcal{C} such that recurrent states induced by any randomized stationary optimal policy are members of \mathcal{C} .

Assumptions B.3, B.4, and B.5 originate from the another result showing convergence of stochastic approximation algorithms (Borkar 1998) and were also required by the convergence theorem of RVI Q-learning. Assumptions B.3 and B.4 can be satisfied if the sequence $\{\alpha_n\}$ decreases to 0 appropriately. The sequence $\{\alpha_n\}$ could be, for example, $1/n$, $1/(n \log n)$, or $\log n/n$ (Abounadi, Bertsekas, and Borkar 2001). The first part of Assumption B.5 requires that, for each state–action pair, the limiting ratio of the number of visitations to the pair and the number of visitations to all pairs is greater than or equal to any fixed positive number. The second part of the assumption requires that the relative update frequency between any two elements is finite. For example, Borkar (personal communication; see also page 403 by Bertsekas and Tsitsiklis 1996) showed that with a common $\alpha_n = 1/n$, Assumption B.5 can be satisfied (Assumption B.3 and B.4 can also be satisfied with this stepsize).

It is easy to verify that under the communicating assumption the following system of equations:

$$q(s, a) = \sum_{s', r} p(s', r | s, a) (r - \bar{r} + \max_{a'} q(s', a')), \text{ for all } s \in \mathcal{S}, a \in \mathcal{A}, \quad (\text{B.10})$$

$$r_* - \bar{R}_0 = \eta \left(\sum q - \sum Q_0 \right), \quad (\text{B.11})$$

has a unique solution for q . Denote the solution as q_∞ .

Theorem B.1 (Convergence of General Differential Q). *If Assumption B.1–Assumption B.5 hold, then the General Differential Q algorithm (Equations B.1–B.3) converges a.s. \bar{R}_n to r_* , Q_n to q_∞ , and $r(\pi_t)$ to r_* where π_t is any greedy policy w.r.t. Q_t .*

We now prove this theorem.

B.1.1. PROOF OF THEOREM B.1

At each step, the increment to \bar{R}_n is η times the increment to Q_n and $\sum Q_n$. Therefore, the cumulative increment can be written

$$\begin{aligned} \bar{R}_n - \bar{R}_0 &= \eta \sum_{i=0}^{n-1} \sum_{s, a} \alpha_{\nu(i, s, a)} \delta_i(s, a) I\{(s, a) \in Y_i\} \\ &= \eta \left(\sum Q_n - \sum Q_0 \right) \\ \implies \bar{R}_n &= \eta \sum Q_n - \eta \sum Q_0 + \bar{R}_0 = \eta \sum Q_n - c, \end{aligned} \quad (\text{B.12})$$

$$\text{where } c \doteq \eta \sum Q_0 - \bar{R}_0. \quad (\text{B.13})$$

Now substituting \bar{R}_n in (B.1) with (B.12), we have $\forall s \in \mathcal{S}, a \in \mathcal{A}$:

$$\begin{aligned} Q_{n+1}(s, a) &= Q_n(s, a) + \alpha_{\nu(n, s, a)} \left(R_n(s, a) + \max_{a'} Q_n(S'_n(s, a), a') - Q_n(s, a) - \eta \sum Q_n + c \right) I\{(s, a) \in Y_n\} \\ &= Q_n(s, a) + \alpha_{\nu(n, s, a)} \left(\tilde{R}_n(s, a) + \max_{a'} Q_n(S'_n(s, a), a') - Q_n(s, a) - \eta \sum Q_n \right) I\{(s, a) \in Y_n\}, \end{aligned} \quad (\text{B.14})$$

where $\tilde{R}_n(s, a) \doteq R_n(s, a) + c$.

(B.14) is in the same form as the RVI Q-learning's update (Equation 2.7 by Abounadi, Bertsekas, and Borkar (2001), see also (8)) with $f(Q_n) = \eta \sum Q_n$, for a MDP $\tilde{\mathcal{M}}$ whose rewards are all shifted by c from the original MDP \mathcal{M} .

This transformed MDP has the same state and action space as the original MDP and has the transition probability defined as

$$\tilde{p}(s', r + c | s, a) \doteq p(s', r | s, a). \quad (\text{B.15})$$

In other words, $\tilde{\mathcal{M}} \doteq (\mathcal{S}, \mathcal{A}, \mathcal{R}, \tilde{p})$.

Note that the communicating assumption we made for the original MDP is still valid for the transformed MDP. For this transformed MDP, denote the best possible average reward rate as \tilde{r}_* . Then

$$\tilde{r}_* = r_* + c \quad (\text{B.16})$$

because the reward in the transformed MDP is shifted by c compared with the original MDP. Combining (B.16), (B.11), and (B.13), we have

$$\tilde{r}_* = \eta \sum q_\infty. \quad (\text{B.17})$$

Furthermore, because

$$\begin{aligned} q_\infty(s, a) &= \sum_{s', r} p(s', r | s, a) (r + \max_{a'} q_\infty(s', a') - r_*) \quad (\text{from (B.10)}) \\ &= \sum_{s', r} p(s', r | s, a) (r + c + \max_{a'} q_\infty(s', a') - \tilde{r}_*) \quad (\text{from (B.16)}) \\ &= \sum_{s', r} \tilde{p}(s', r | s, a) (r + \max_{a'} q_\infty(s', a') - \tilde{r}_*) \quad (\text{from (B.15)}), \end{aligned} \quad (\text{B.18})$$

q_∞ is a solution of q in the action-value Bellman equations for not only the original MDP \mathcal{M} but also the transformed MDP $\tilde{\mathcal{M}}$.

If the convergence theorem of the RVI Q-learning applies, then $Q_n \rightarrow q_\infty$ and $\eta \sum Q_n \rightarrow \tilde{r}_*$. However, in general, $f(x) \doteq \eta \sum x$ does not satisfy some requirements on f by Abounadi, Bertsekas, and Borkar (2001). In particular,

$$f(e) = 1, \text{ and } f(x + ce) = f(x) + c, \forall x \in \mathbb{R}^{|\mathcal{S} \times \mathcal{A}|} \quad (\text{B.19})$$

in Assumption 2.2 (Abounadi, Bertsekas, and Borkar 2001) are violated. In the next section, Theorem B.2, we extend the RVI Q-learning family of algorithms by replacing (B.19) with the following weaker assumptions:

$$\exists u > 0 \text{ s.t. } f(e) = u, \text{ and } f(x + ce) = f(x) + cu, \forall x \in \mathbb{R}^{|\mathcal{S} \times \mathcal{A}|}. \quad (\text{B.20})$$

It can be seen that (B.19) is a special case of (B.20) when $u = 1$. Therefore RVI Q-learning family is a subset of the extended RVI Q-learning family.

Because $f(x) = \eta \sum x$ satisfies assumptions on f required by Theorem B.2 and Assumption B.1-Assumption B.5 also hold for the transformed MDP $\tilde{\mathcal{M}}$, (B.14) converges a.s. Q_n to q_∞ , which is the solution of

$$\begin{aligned} q(s, a) &= \sum_{s', r} \tilde{p}(s', r | s, a) (r - \bar{r} + \max_{a'} q(s', a')), \text{ for all } s \in \mathcal{S}, a \in \mathcal{A} \\ \eta \sum q &= \tilde{r}_* \end{aligned}$$

by (B.18) and (B.17).

Now consider \bar{R}_n . Combining (B.12) and $Q_n \rightarrow q_\infty$, we have $\bar{R}_n \rightarrow \eta \sum q_\infty - c$. In addition, because $\eta \sum q_\infty = \tilde{r}_*$ (Equation B.17), we have $\bar{R}_n \rightarrow \tilde{r}_* - c$. Because $\tilde{r}_* = r_* + c$ (Equation B.16), we have

$$\bar{R}_n \rightarrow r_* \text{ a.s. as } n \rightarrow \infty. \quad (\text{B.21})$$

Finally consider $r(\pi_t)$ where π_t is a greedy policy w.r.t. Q_t . From Theorem 8.5.5 by Puterman (1994), we have,

$$\min_{s, a} (TQ_t(s, a) - Q_t(s, a)) \leq r(\pi_t) \leq r_* \leq \max_{s, a} (TQ_t(s, a) - Q_t(s, a)) \quad (\text{B.22})$$

$$\implies |r_* - r(\pi_t)| \leq sp(TQ_t - Q_t) \quad (\text{B.23})$$

where $TQ(s, a) \doteq \sum_{s', r} \tilde{p}(s', r | s, a)(r + \max_{a'} Q(s', a'))$. Because $Q_t \rightarrow q_\infty$ a.s., and $sp(TQ_t - Q_t)$ is a continuous function of Q_t , by continuous mapping theorem, $sp(TQ_t - Q_t) \rightarrow sp(Tq_\infty - q_\infty) = 0$ a.s. Therefore we conclude that $r(\pi_t) \rightarrow r_*$.

Theorem B.1 is proved.

Theorem B.2 (Convergence of the Extended RVI Q-learning). *For any $Q_0 \in \mathbb{R}^{|\mathcal{S} \times \mathcal{A}|}$, let $R_n, Y_n, \alpha_{\nu(n, s, a)}$ be defined as aforementioned, consider an update rule*

$$Q_{n+1}(s, a) = Q_n(s, a) + \alpha_{\nu(n, s, a)} \left(R_n(s, a) + \max_{a'} Q_n(S'_n(s, a), \cdot) - Q_n(s, a) - f(Q_n) \right) I\{(s, a) \in Y_n\}, \quad (\text{B.24})$$

if

1. Assumption B.1-Assumption B.5 hold,
2. $f : \mathbb{R}^{|\mathcal{S} \times \mathcal{A}|} \rightarrow \mathbb{R}$ is Lipschitz and there exists some $u > 0$ such that $\forall c \in \mathbb{R}$ and $x \in \mathbb{R}^{|\mathcal{S} \times \mathcal{A}|}$, $f(e) = u$, $f(x + ce) = f(x) + cu$ and $f(cx) = cf(x)$,

then Q_n converges a.s. to q_* , where q_* is the solution to action-value optimality equation (Equation (B.10)) satisfying $f(q_*) = r_*$.

If we set $u = 1$ in the above theorem, then we recover the convergence result of RVI Q-learning.

The rest part of this section proves the above theorem. We use arguments similar to those of RVI Q-learning.

First, note that (B.24) is in the same form as the asynchronous update (Equation 7.1.2) by Borkar (2009). We apply the result in Section 7.4 of the same text (Borkar 2009) (see also Theorem 3.2 by Borkar (1998)), which shows convergence for Equation 7.1.2, to show convergence of (B.14). This result, given Assumption B.4, B.5, only requires showing the convergence of the following *synchronous* version of (B.24):

$$Q_{n+1}(s, a) = Q_n(s, a) + \alpha_n (R_n(s, a) + g(Q_n(S'_n(s, a), \cdot)) - Q_n(s, a) - f(Q_n)), \text{ for all } s \in \mathcal{S}, a \in \mathcal{A}, \quad (\text{B.25})$$

Like the proof of RVI Q-learning, first define operators T, T_1, T_2 :

$$\begin{aligned} T(Q)(s, a) &\doteq \sum_{s', r} p(s', r | s, a)(r + g(Q(s', \cdot))), \\ T_1(Q) &\doteq T(Q) - r_*e, \\ T_2(Q) &\doteq T(Q) - f(Q)e = T_1(Q) + (r_* - f(Q))e. \end{aligned}$$

Consider two ordinary differential equations (ODEs):

$$\dot{y}_t = T_1(y_t) - y_t, \quad (\text{B.26})$$

$$\dot{x}_t = T_2(x_t) - x_t. \quad (\text{B.27})$$

Note that by the properties of T_1 and T_2 , both (B.26) and (B.27) have Lipschitz r.h.s.'s and thus are well-posed.

The next two lemmas are the same as Lemma 3.1 and Lemma 3.2 by Abounadi, Bertsekas, and Borkar (2001). Their proofs do not rely on properties of f and therefore they hold with our more general f function.

Lemma B.1. *Let \bar{y} be an equilibrium point of the ODE defined in (B.26). Then $\|y_t - \bar{y}\|_\infty$ is nonincreasing, and $y_t \rightarrow y_*$ for some equilibrium point y_* of (B.26) that may depend on y_0 .*

Lemma B.2. (B.27) *has a unique equilibrium at q_* .*

We then show the relation between x_t and y_t using the following lemma. It shows that the difference between x_t and y_t is a vector with identical elements and this vector satisfies a new ODE.

Lemma B.3. *Let $x_0 = y_0$, then $x_t = y_t + z_t e$, where z_t satisfies the ODE $\dot{z}_t = -uz_t + (r_* - f(y_t))$.*

Proof. The proof of $x_t = y_t + z_t e$ is the same with the Lemma 3.3 by Abounadi, Bertsekas, and Borkar (2001).

Now we show $\dot{z}_t = -uz_t + (r_* - f(y_t))e$. Note that $f(x_t) = f(y_t + z_t e) = f(y_t) + uz_t$. In addition, $T_1(x_t) - T_1(y_t) = T_1(y_t + z_t e) - T_1(y_t) = T_1(y_t) + z_t e - T_1(y_t) = z_t e$, therefore we have, for $z_t \in \mathbb{R}$:

$$\begin{aligned}
 \dot{z}_t e &= \dot{x}_t - \dot{y}_t \\
 &= (T_1(x_t) - x_t + (r_* - f(x_t))e) - (T_1(y_t) - y_t) \quad (\text{from (B.26) and (B.27)}) \\
 &= -(x_t - y_t) + (T_1(x_t) - T_1(y_t)) + (r_* - f(x_t))e \\
 &= -z_t e + z_t e + (r_* - f(x_t))e \\
 &= -uz_t e + uz_t e + (r_* - f(x_t))e \\
 &= -uz_t e + (r_* - f(y_t))e \\
 \implies \dot{z}_t &= -uz_t + (r_* - f(y_t)).
 \end{aligned}$$

□

With the above lemmas, we have:

Lemma B.4. q_* is the globally asymptotically stable equilibrium for (B.27).

Proof. We have shown that q_* is the unique equilibrium in Lemma B.2.

With that result, we first prove Lyapunov stability. That is, we need to show that given any $\epsilon > 0$, we can find a $\delta > 0$ such that $\|q_* - x_0\|_\infty \leq \delta$ implies $\|q_* - x_t\|_\infty \leq \epsilon$ for $t \geq 0$.

First, from Lemma B.3 we have $\dot{z}_t = -uz_t + (r_* - f(y_t))e$. By variation of parameters and $z_0 = 0$, we have

$$z_t = \int_0^t \exp(u(\tau - t)) (r_* - f(y_\tau)) d\tau.$$

Then

$$\begin{aligned}
 \|q_* - x_t\|_\infty &= \|q_* - y_t - z_t u e\|_\infty \\
 &\leq \|q_* - y_t\|_\infty + u |z_t| \\
 &\leq \|q_* - y_0\|_\infty + u \int_0^t \exp(u(\tau - t)) |r_* - f(y_\tau)| d\tau \\
 &= \|q_* - x_0\|_\infty + u \int_0^t \exp(u(\tau - t)) |f(q_*) - f(y_\tau)| d\tau \quad (\text{from (B.17)}). \tag{B.28}
 \end{aligned}$$

Because f is L -lipschitz, we have

$$\begin{aligned}
 |f(q_*) - f(y_\tau)| &\leq L \|q_* - y_\tau\|_\infty \\
 &\leq L \|q_* - y_0\|_\infty \quad (\text{from Lemma B.1}) \\
 &= L \|q_* - x_0\|_\infty,
 \end{aligned}$$

$$\begin{aligned}
 \int_0^t \exp(u(\tau - t)) |f(q_*) - f(y_\tau)| d\tau &\leq \int_0^t \exp(u(\tau - t)) L \|q_* - x_0\|_\infty d\tau \\
 &= L \|q_* - x_0\|_\infty \int_0^t \exp(u(\tau - t)) d\tau \\
 &= L \|q_* - x_0\|_\infty \frac{1}{u} (1 - \exp(-ut)) \\
 &= \frac{L}{u} \|q_* - x_0\|_\infty (1 - \exp(-ut))
 \end{aligned}$$

Substituting the above equation in (B.28), we have

$$\|q_* - x_t\|_\infty \leq (1 + L) \|q_* - x_0\|_\infty.$$

Lyapunov stability follows.

Now in order to prove the asymptotic stability, in addition to Lyapunov stability, we need to show that there exists $\delta > 0$ such that if $\|x_0 - q_*\|_\infty < \delta$, then $\lim_{t \rightarrow \infty} \|x_t - q_*\|_\infty = 0$. Note that

$$\begin{aligned} \lim_{t \rightarrow \infty} z_t &= \lim_{t \rightarrow \infty} \int_0^t \exp(u(\tau - t)) (r_* - f(y_\tau)) d\tau \\ &= \lim_{t \rightarrow \infty} \frac{\int_0^t \exp(u\tau) (r_* - f(y_\tau)) d\tau}{\exp(ut)} \\ &= \lim_{t \rightarrow \infty} \frac{\exp(ut) (r_* - f(y_t))}{u \exp(ut)} \quad (\text{by L'Hospital's rule}) \\ &= \frac{r_* - f(y_*)}{u} \quad (\text{by Lemma B.1}). \end{aligned}$$

Because $x_t = y_t + z_t e$ (Lemma B.3) and $y_t \rightarrow y_*$ (Lemma B.1), we have $x_t \rightarrow y_* + (r_* - f(y_*))e/u$, which must coincide with q_* because that is the only equilibrium point for (B.27) (Lemma B.2). Therefore $\lim_{t \rightarrow \infty} \|x_t - q_*\|_\infty = 0$ for any x_0 . Asymptotic stability is shown and the proof is complete. \square

Lemma B.5. Equation B.25 converges a.s. Q_n to q_* as $n \rightarrow \infty$.

Proof. The proof uses Theorem 2 in Section 2 of Borkar (2009) and is essentially the same as Lemma 3.8 by Abounadi, Bertsekas and Borkar (2001). For completeness, we repeat the proof (with more details) here.

First write the synchronous update (B.25) as

$$Q_{n+1} = Q_n + \alpha_n (h(Q_n) + M_{n+1})$$

where

$$\begin{aligned} h(Q_n)(s, a) &\doteq \sum_{s', r} p(s', r | s, a) (r + \max_{a'} Q_n(s', a')) - Q_n(s, a) - f(Q_n) \\ &= T(Q_n)(s, a) - Q_n(s, a) - f(Q_n) \\ &= T_2(Q_n)(s, a) - Q_n(s, a), \\ M_{n+1}(s, a) &\doteq R_n(s, a) + \max_{a'} Q_n(S'_n(s, a), a') - T(Q_n)(s, a). \end{aligned}$$

Theorem 2 requires verifying the following conditions and concludes that Q_n converges to a (possibly sample path dependent) compact connected internally chain transitive invariant set of ODE $\dot{x}_t = h(x_t)$. This is exactly the ODE defined in (B.27). Lemma B.2 and B.4 conclude that this ODE has q_∞ as the unique globally asymptotically stable equilibrium. Therefore the (possibly sample path dependent) compact connected internally chain transitive invariant set is a singleton set containing only the unique globally asymptotically stable equilibrium. Thus Theorem 2 concludes that $Q_n \rightarrow q_\infty$ a.s. as $n \rightarrow \infty$. We now list conditions required by Theorem 2:

- **(A1)** The function h is Lipschitz: $\|h(x) - h(y)\| \leq L \|x - y\|$ for some $0 < L < \infty$.
- **(A2)** The sequence $\{\alpha_n\}$ satisfies $\alpha_n > 0$, and $\sum \alpha_n = \infty$, $\sum \alpha_n^2 < \infty$.
- **(A3)** $\{M_n\}$ is a martingale difference sequence with respect to the increasing family of σ -fields

$$\mathcal{F}_n \doteq \sigma(Q_i, M_i, i \leq n), n \geq 0$$

That is

$$\mathbb{E}[M_{n+1} \mid \mathcal{F}_n] = 0 \quad \text{a.s., } n \geq 0.$$

Furthermore, $\{M_n\}$ are square-integrable

$$\mathbb{E}[\|M_{n+1}\|^2 \mid \mathcal{F}_n] \leq K(1 + \|Q_n\|^2) \quad \text{a.s., } n \geq 0,$$

for some constant $K > 0$.

- **(A4)** $\sup_n \|Q_n\| \leq \infty$ a.s..

Let us verify these conditions now.

(A1) is satisfied as both T and \sum operators are Lipschitz.

(A2) is satisfied by Assumption B.3.

(A3) is also satisfied because for any $s \in \mathcal{S}, a \in \mathcal{A}$

$$\begin{aligned} \mathbb{E}[M_{n+1}(s, a) \mid \mathcal{F}_n] &= \mathbb{E} \left[R_n(s, a) + \max_{a'} Q_n(S'_n(s, a), a') - T(Q_n)(s, a) \mid \mathcal{F}_n \right] \\ &= \mathbb{E} \left[R_n(s, a) + \max_{a'} Q_n(S'_n(s, a), a') \mid \mathcal{F}_n \right] - T(Q_n)(s, a) \\ &= 0 \end{aligned}$$

and $\mathbb{E}[\|M_{n+1}\|^2 \mid \mathcal{F}_n] \leq K(1 + \|Q_n\|^2)$ for a suitable constant $K > 0$ can be verified by a simple application of triangle inequality.

To verify (A4), we apply Theorem 7 in Section 3 by Borkar (2009), which shows $\sup_n \|Q_n\| \leq \infty$ a.s., if (A1), (A2), and (A3) are all satisfied and in addition we have the following condition satisfied:

(A5) The functions $h_d(x) \doteq h(dx)/d, d \geq 1, x \in \mathbb{R}^k$, satisfy $h_d(x) \rightarrow h_\infty(x)$ as $d \rightarrow \infty$, uniformly on compacts for some $h_\infty \in C(\mathbb{R}^k)$. Furthermore, the ODE $\dot{x}_t = h_\infty(x_t)$ has the origin as its unique globally asymptotically stable equilibrium.

Note that

$$\begin{aligned} h_\infty(x) &= \lim_{d \rightarrow \infty} h_d(x) = \lim_{d \rightarrow \infty} (T(dx) - dx - f(dx)e) / d \\ &= T_0(x) - x - f(x)e \end{aligned}$$

where

$$T_0(x) \doteq \sum_{s', r} p(s', r \mid s, a) \max_{a'} x(s', a').$$

The function h_∞ is clearly continuous in every $x \in \mathbb{R}^k$ and therefore $h_\infty \in C(\mathbb{R}^k)$.

Now consider the ODE $\dot{x}_t = h_\infty(x_t) = T_0(x_t) - x_t - f(x_t)e$. Clearly the origin is an equilibrium. This ODE is a special case of (B.27), corresponding to the reward being always zero, therefore Lemma B.2 and B.4 also apply to this ODE and the origin is the unique globally asymptotically stable equilibrium.

(A1), (A2), (A3), (A4) are all verified and therefore

$$Q_n \rightarrow q_* \quad \text{a.s. as } n \rightarrow \infty. \tag{B.29}$$

□

B.2. Proof of Differential TD-learning and Differential TD-planning

The proof is similar to that of Differential Q-learning and Differential Q-planning. We consider a General Differential TD algorithm which includes both Differential TD-learning and Differential TD-planning.

Given a MDP $\mathcal{M} \doteq (\mathcal{S}, \mathcal{A}, \mathcal{R}, p)$, a behavior policy b , and a target policy π , for any state $s \in \mathcal{S}$ and discrete step $n \geq 0$, let $A_n(s) \sim b(\cdot | s)$, $R_n(s, A_n(s))$, $S'_n(s, A_n(s)) \sim p(\cdot, \cdot | s, A_n(s))$. We hypothesize a set-valued process $\{Y_n\}$ taking values in the set of nonempty subsets of \mathcal{S} with the interpretation: $Y_n = \{s : s \text{ component of } V \text{ was updated at time } n\}$. Define $\nu(n, s) = \sum_{i=0}^n I\{s \in Y_i\}$ where I is the indicator function. Thus $\nu(n, s)$ = the number of times $V(s)$ was updated up to time n . Then the update rules of General Differential TD are, for $n \geq 0$:

$$V_{n+1}(s) \doteq V_n(s) + \alpha_{\nu(n,s)} \rho_n(s) \delta_n(s) I\{s \in Y_n\} \quad \forall s \in \mathcal{S} \quad (\text{B.30})$$

$$\bar{R}_{n+1} \doteq \bar{R}_n + \sum_s \alpha_{\nu(n,s)} \rho_n(s) \delta_n(s) I\{s \in Y_n\}, \quad (\text{B.31})$$

where

$$\delta_n(s) \doteq R_n(s, A_n(s)) + V_n(S'_n(s, A_n(s))) - V_n(s) - \bar{R}_n, \quad (\text{B.32})$$

and $\rho_n(s) \doteq \pi(A_n(s) | s) / b(A_n(s) | s)$ is the importance sampling ratio (this is always well-defined given Assumption B.7).

The quantity $\alpha_{\nu(n,s)}$ is the stepsize at step n for state s and can be obtained the same way as introduced in B.1. It can be shown, using similar arguments as those in B.1, that Differential TD-learning and Differential TD-planning are special cases of General Differential TD. And therefore we only need to prove the convergence of General Differential TD. We now specify required assumptions for the convergence proof.

Assumption B.6. *The Markov chain induced by the target policy is unichain.*

Assumption B.7 (Coverage Assumption). *$b(a | s) > 0$ if $\pi(a | s) > 0$ for all $s \in \mathcal{S}$, $a \in \mathcal{A}$.*

The above assumption requires that the behavior policy covers all possible state–action pairs the target policy may incur. To guarantee the full coverage, we will need that the behavior policy visit all states for an infinite number of times.

Assumption B.8 (Asynchronous Stepsize Assumption B). *There exists $\Delta > 0$ such that*

$$\liminf_{n \rightarrow \infty} \frac{\nu(n, s)}{n+1} \geq \Delta,$$

a.s., for all $s \in \mathcal{S}$. Furthermore, for all $x > 0$, and

$$N(n, x) = \min \left\{ m \geq n : \sum_{i=n+1}^m \alpha_i \geq x \right\},$$

the limit

$$\lim_{n \rightarrow \infty} \frac{\sum_{i=\nu(n,s)}^{\nu(N(n,x),s)} \alpha_i}{\sum_{i=\nu(n,s')}^{\nu(N(n,x),s')} \alpha_i}$$

exists a.s. for all s, s' .

It can be easily verified that

$$v(s) = \sum_a \pi(a | s) \sum_{s', r} p(s', r | s, a) (r - \bar{r} + v(s')), \quad \text{for all } s \in \mathcal{S}, \quad (\text{B.33})$$

$$r(\pi) - \bar{R}_0 = \eta \left(\sum v - \sum V_0 \right) \quad (\text{B.34})$$

has a unique solution of v . Denote the solution as v_∞ .

Theorem B.3 (Convergence of General Differential TD). *If Assumptions B.6, B.3, B.4, B.8, and B.7 hold, then General Differential TD (Equations B.30-B.32) converges a.s., \bar{R}_n to $r(\pi)$ and V_n to v_∞ .*

We now prove this theorem.

B.2.1. PROOF OF THEOREM B.3

Similar as what we did in the proof of General Differential Q, we can combine update rules (B.30)-(B.32) to obtain a single update rule.

$$\begin{aligned}
 \bar{R}_n - \bar{R}_0 &= \eta \sum_{i=0}^{n-1} \sum_s \alpha_{\nu(i,s)} \rho_i(s) \delta_i(s) I\{s \in Y_k\} \\
 &= \eta \left(\sum V_n - \sum V_0 \right) \\
 &\implies \\
 \bar{R}_n &= \eta \sum V_n - \eta \sum V_0 + \bar{R}_0 = \eta \sum V_n - c, \tag{B.35}
 \end{aligned}$$

$$\text{where } c \doteq \eta \sum V_0 - \bar{R}_0. \tag{B.36}$$

Substituting \bar{R}_n in (B.30) with (B.35) we have, $\forall s \in \mathcal{S}$:

$$\begin{aligned}
 V_{n+1}(s) &= V_n(s) + \alpha_{\nu(n,s)} \rho_n(s) \left(R_n(s, A_n(s)) + V_n(S'_n(s, A_n(s))) - V_n(s) - \eta \sum V_n + c \right) I\{s \in Y_n\} \\
 &= V_n(s) + \alpha_{\nu(n,s)} \rho_n(s) \left(\tilde{R}_n(s, A_n(s)) + V_n(S'_n(s, A_n(s))) - V_n(s) - \eta \sum V_n \right) I\{s \in Y_n\}, \tag{B.37}
 \end{aligned}$$

where $\tilde{R}_n(s, A_n(s)) \doteq R_n(s, A_n(s)) + c$. Now (B.37) is in the same form with the asynchronous update (Equation 7.1.2) studied by Borkar (2009). Again we can apply the result in Section 7.4 by Borkar (2009) to show convergence of (B.37). This result, given Assumption B.4 and B.8, only requires showing the convergence of the following *synchronous* version of General Differential TD:

$$V_{n+1}(s) = V_n(s) + \alpha_n \rho_n(s) \left(\tilde{R}_n(s, A_n(s)) + V_n(S'_n(s, A_n(s))) - V_n(s) - \eta \sum V_n \right), \quad \forall s \in \mathcal{S}. \tag{B.38}$$

This transformed MDP has the same state and action space as the original MDP and has the transition probability defined as

$$\tilde{p}(s', r + c | s, a) \doteq p(s', r | s, a). \tag{B.39}$$

Note that the unichain assumption (Assumption B.1) and the coverage assumption (Assumption B.7) we made for the original MDP is still valid for the transformed MDP. For this transformed MDP, denote the average reward rate following policy π as $\tilde{r}(\pi)$. Then

$$\tilde{r}(\pi) = r(\pi) + c \tag{B.40}$$

because the reward in the transformed MDP is shifted by c compared with the original MDP.

Combining (B.40), (B.34) and (B.36), we have

$$\tilde{r}(\pi) = \eta \sum v_\infty. \tag{B.41}$$

Furthermore,

$$\begin{aligned}
 v_\infty(s) &= \sum_a \pi(a | s) \sum_{s', r} p(s', r | s, a) (r + v_\infty(s') - r(\pi)) \quad (\text{from (B.33)}) \\
 &= \sum_a \pi(a | s) \sum_{s', r} p(s', r | s, a) (r + c + v_\infty(s') - \tilde{r}(\pi)) \quad (\text{from (B.40)}) \\
 &= \sum_a \pi(a | s) \sum_{s', r} \tilde{p}(s', r | s, a) (r + v_\infty(s') - \tilde{r}(\pi)),
 \end{aligned}$$

therefore v_∞ is a solution of v in the state-value Bellman equations for not only the original MDP \mathcal{M} but also the transformed MDP $\tilde{\mathcal{M}}$.

We now show $V_n \rightarrow v_\infty$ and $\eta \sum V_n \rightarrow \tilde{r}(\pi)$. First, define operators T, T_1, T_2 :

$$\begin{aligned} T(V)(s) &\doteq \sum_a \pi(a | s) \sum_{s', r} \tilde{p}(s', r | s, a) (r + V(s')), \\ T_1(V) &\doteq T(V) - \tilde{r}(\pi)e, \\ T_2(V) &\doteq T(V) - \left(\eta \sum V\right) e = T_1(V) + \left(\tilde{r}(\pi) - \eta \sum V\right) e. \end{aligned}$$

Consider two ODEs:

$$\dot{y}_t = T_1(y_t) - y_t, \quad (\text{B.42})$$

$$\dot{x}_t = T_2(x_t) - x_t. \quad (\text{B.43})$$

Note that by the properties of T, T_1, T_2 , both (B.42) and (B.43) have Lipschitz R.H.S.'s and thus are well-posed.

The next lemma is similar to Lemma 3.1 by Abounadi, Bertsekas, and Borkar (2001) and is a special case of Theorem 3.1 and Lemma 3.2 by Borkar and Soumyanath (1997).

Lemma B.6. *Let \bar{y} be an equilibrium point of (B.42). Then $\|y_t - \bar{y}\|_\infty$ is nonincreasing, and $y_t \rightarrow y_*$ for some equilibrium point y_* of (B.42) that may depend on y_0 .*

The next lemma is similar to Lemma B.2 and the proof of it is almost the same as the proof of Lemma B.2. The only changes are to replace \tilde{r}_*, q_∞ and q with $\tilde{r}(\pi), v_\infty$ and v respectively.

Lemma B.7. (B.43) has a unique equilibrium at v_∞ .

The next two lemmas are almost the same as Lemma B.3 and B.4. Their proofs can be easily obtained from the proofs of Lemma B.3 and B.4 by replacing \tilde{r}_* with $\tilde{r}(\pi)$.

Lemma B.8. *Let $x_0 = y_0$, then $x_t = y_t + z_t e$, where z_t satisfies the ODE $\dot{z}_t = -kz_t + (\tilde{r}(\pi) - k \sum y_t)$, and $k \doteq |\mathcal{S}|$.*

Lemma B.9. v_∞ is the unique globally asymptotically stable equilibrium for (B.43).

Lemma B.10. *Synchronous General Differential TD (Equation B.38) converges a.s., V_n to v_∞ as $n \rightarrow \infty$.*

Proof. Similar as what we did in the proof of Lemma B.5, we use Theorem 2 in Section 2 by Borkar (2009) to show the convergence of this lemma.

We first write the synchronous update rule (B.38) as

$$V_{n+1} = V_n + \alpha_n (h(V_n) + M_{n+1}), \quad (\text{B.44})$$

where

$$\begin{aligned} h(V_n)(s) &\doteq \sum_a \pi(a | s) \sum_{s', r} \tilde{p}(s', r | s, a) (r + V_n(s')) - V_n(s) - \eta \sum V_n \\ &= T(V_n)(s) - V_n(s) - \eta \sum V_n \\ &= T_2(V_n)(s) - V_n(s), \end{aligned} \quad (\text{B.45})$$

$$M_{n+1}(s) \doteq \rho_n(s) \left(\tilde{R}_n(s, A_n(s)) + V_n(S'_n(s, A_n(s))) - V_n(s) - \eta \sum V_n \right) - h(V_n)(s). \quad (\text{B.46})$$

Similar as the proof of Lemma B.5, we only need to verify conditions (A1) - (A4) in order to conclude that V_n converges v_∞ a.s. as $n \rightarrow \infty$.

(A1) is satisfied as both T and \sum operators are Lipschitz.

(A2) is satisfied by Assumption B.3.

(A3) is also satisfied because for any $s \in \mathcal{S}$

$$\begin{aligned}
 \mathbb{E}[M_{n+1}(s) \mid \mathcal{F}_n] &= \mathbb{E} \left[\rho_n(s) \left(\tilde{R}_n(s, A_n(s)) + V_n(S'_n(s, A_n(s))) - V_n(s) - \eta \sum V_n \right) - h(V_n)(s) \mid \mathcal{F}_n \right] \\
 &= \mathbb{E} \left[\rho_n(s) \left(\tilde{R}_n(s, A_n(s)) + V_n(S'_n(s, A_n(s))) - V_n(s) - \eta \sum V_n \right) \mid \mathcal{F}_n \right] - h(V_n)(s) \\
 &= \mathbb{E} \left[\rho_n(s) \left(\tilde{R}_n(s, A_n(s)) + V_n(S'_n(s, A_n(s))) \right) \mid \mathcal{F}_n \right] - V_n(s) - \eta \sum V_n - h(V_n)(s) \\
 &= \mathbb{E}[\rho_n(s)(\tilde{R}_n(s, A_n(s)) + V_n(S'_n(s, A_n(s)))) \mid \mathcal{F}_n] - T(V_n)(s) \\
 &= 0
 \end{aligned}$$

and $\mathbb{E}[\|M_{n+1}\|^2 \mid \mathcal{F}_n] \leq K(1 + \|V_n\|^2)$ for a suitable constant $K > 0$ can be verified by applying triangle inequality given the boundedness of the second moment of the importance sampling ratio, reward and V_n .

To verify (A4), again we only need to verify (A5). Note that

$$h_\infty(x) = \lim_{a \rightarrow \infty} h_a(x) = \lim_{a \rightarrow \infty} \frac{T(ax) - ax - \eta(\sum ax)e}{a} = T_0(x) - x - \eta\left(\sum x\right)e,$$

where

$$T_0(x) \doteq \sum_a \pi(a \mid s) \sum_{s', r} \tilde{p}(s', r \mid s, a) x(s').$$

The function h_∞ is clearly continuous in every $x \in \mathbb{R}^k$ and therefore $h_\infty \in C(\mathbb{R}^k)$.

Now consider the ODE $\dot{x}_t = h_\infty(x_t) = T_0(x_t) - x_t - \eta(\sum x_t)e$, clearly the origin is an equilibrium. This ODE is a special case of (B.43), corresponding to the reward being always zero, therefore Lemma B.7 and Lemma B.9 also apply to this ODE and the origin is the unique globally asymptotically stable equilibrium.

(A1), (A2), (A3), (A4) are all verified and therefore $V_n \rightarrow v_\infty$ a.s. as $n \rightarrow \infty$.

□

Given the convergence of V_n in the synchronous update rule (B.38), the convergence of V_n in the original update rule (B.37) follows immediately using results introduced in Chapter 7 of Borkar (2009) under Assumption B.4, B.8.

Finally consider \bar{R}_n . Because $\bar{R}_n = \eta \sum V_n - c$ (Equation B.35) and $V_n \rightarrow v_\infty$, we have $\bar{R}_n \rightarrow \eta \sum v_\infty - c$. In addition, because $\tilde{r}(\pi) = \eta \sum v_\infty$, we have $\bar{R}_n \rightarrow \tilde{r}(\pi) - c$. Finally, because $\tilde{r}(\pi) = r(\pi) + c$, we have

$$\bar{R}_n \rightarrow r(\pi). \tag{B.47}$$

a.s. as $n \rightarrow \infty$.

Theorem B.3 is proved.

B.3. Centered Algorithms

This section serves as a supplement of Section 6 of the main text. We introduce 4 algorithms: Centered Differential TD-learning, Centered Differential TD-planning, Centered Differential Q-learning, and Centered Differential Q-planning. All these algorithms are shown to converge to the centered (actual) differential value function rather than the differential value function plus some offset.

The next lemma is useful in the convergence proofs for the centered algorithms.

Lemma B.11. *Let π be a stationary Markov policy. Assume that the induced Markov chain under π is unichain. Let d_π be the stationary distribution following policy π . Then*

1) $(v, \bar{r}) = (v_\pi, r(\pi))$ is the unique solution of (B.33) and

$$\sum_s d_\pi(s) v(s) = 0, \tag{B.48}$$

and

2) if $v = v_\pi + ce$ then $c = \sum_s d_\pi(s) v(s)$.

Proof. Let P_π denote the $|\mathcal{S}| \times |\mathcal{S}|$ transition probability matrix under policy π , i.e., $P_\pi(s, s') \doteq \sum_{a,r} \pi(a | s) p(s', r | s, a)$ and let $P_\pi^* \doteq \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{t=1}^N P_\pi^{t-1}$. Because \mathcal{S} is finite, the limit exists and P_π^* is a stochastic matrix (has row sums equal to 1). Because the Markov chain induced by π is unichain, all rows of P_π^* are identical and are all equal to d_π^\top . Let $r_\pi(s) \doteq \sum_{a,r,s'} \pi(a | s) p(s', r | s, a) r$ denote the expected one-step reward under π . Then the average reward rate following π can be written as

$$r(\pi) = d_\pi^\top r_\pi, \quad (\text{B.49})$$

and the differential value function following policy π can be written as

$$v_\pi(s) = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{k=0}^{N-1} \sum_{t=0}^k P_\pi^t (r_\pi - r(\pi))(s),$$

or $v_\pi = H_{P_\pi} r_\pi$ in vector form, where $H_{P_\pi} \doteq \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{k=0}^{N-1} \sum_{t=0}^k (P_\pi^t - P_\pi^*)$.

The differential value function v_π satisfies (B.33) due to Theorem 8.2.6 (a) by Puterman (1994).

To see that v_π satisfies the equation (B.48), we apply Equation A.18 in Appendix A by Puterman (1994), which is $P_\pi^* H_{P_\pi} = 0$. Therefore we have $d_\pi^\top H_{P_\pi} = 0$ because all rows of P_π^* are d_π^\top . Because $v_\pi = H_{P_\pi} r_\pi$, we have $d_\pi^\top v_\pi = d_\pi^\top H_{P_\pi} r_\pi = 0$.

To verify that v_π is the unique solution of (B.33) and (B.48), suppose there exists another vector $v' \neq v_\pi$ satisfying (B.33) and (B.48), then $v' = v_\pi + ce$ for some $c \neq 0$ (any two solutions of (B.33) differ by a constant). Substituting this into (B.48), we have

$$d_\pi^\top v' = d_\pi^\top (v_\pi + ce) = d_\pi^\top v_\pi + cd_\pi^\top e = c$$

To satisfy (B.48), we must have $c = 0$. Therefore, v_π is the unique solution of (B.33) and (B.48).

To prove the second part, consider $v = v_\pi + ce$, then we have $\sum_s d_\pi(s) v(s) = \sum_s d_\pi(s) (v_\pi + ce)(s) = c$. \square

B.3.1. CENTERED DIFFERENTIAL TD-LEARNING AND DIFFERENTIAL TD-PLANNING

Centered Differential TD-learning is already presented in Section 6 of the main text. The planning version of Centered Differential TD-learning is called *Centered Differential TD-planning*. It uses simulated experience just as in Differential TD-planning. In addition, just like Differential TD-planning, Centered Differential TD-planning maintains V_n and \bar{R}_n . Centered Differential TD-planning also maintains an auxiliary table of estimates $F_n(s, a)$, $\forall s \in \mathcal{S}, a \in \mathcal{A}$ and an offset estimate \bar{V}_n , and updates them just as in Centered Differential TD-learning, using S_n, A_n, R_n, S'_n instead of $S_t, A_t, R_{t+1}, S_{t+1}$.

Just as we did in Section B.1 and B.2, we now present a general algorithm that includes both Centered Differential TD-learning and Centered Differential TD-planning. We call it *General Centered Differential TD*. Using arguments that are similar as those in Section B.1, it can be shown that both Centered Differential TD-learning and Centered Differential TD-planning are special cases of General Centered Differential TD.

The data is generated the same way as it in B.2. Also, we use same notations introduced in B.2. In addition to update rules of General Differential TD (Equation B.30-B.32), General Centered Differential TD has two more update rules:

$$F_{n+1}(s) \doteq F_n(s) + \beta_{\nu(n,s)} \rho_n(s) \Delta_n(s) I\{s \in Y_n\} \quad \forall s \in \mathcal{S}, \quad (\text{B.50})$$

$$\bar{V}_{n+1} \doteq \bar{V}_n + \kappa \beta_{\nu(n,s)} \sum_s \rho_n(s) \Delta_n(s) I\{s \in Y_n\}, \quad (\text{B.51})$$

where

$$\Delta_n(s) \doteq V_n(s) + F_n(S'_n(s), A_n(s)) - F_n(s) - \bar{V}_n. \quad (\text{B.52})$$

Here $\beta_{\nu(n,s)}$ is the stepsize and κ is a positive number. $\beta_{\nu(n,s)}$ and κ doesn't need to be equal to $\alpha_{\nu(s,a)}$ and η .

Theorem B.4 (Convergence of Centered Differential TD). *If Assumption B.1 holds, Assumption B.3, B.4, B.8, and B.7 hold for both α_n and β_n , then General Centered Differential TD (Equations B.30-B.32, B.50-B.52) converges a.s., \bar{R}_n to $r(\pi)$ and $V_n - \bar{V}_n e$ to v_π .*

Proof. To show this theorem, we use the last extension of Section 2.2 by Borkar (2009), which states that a deterministic or random bounded $o(1)$ noise will not influence the convergence.

Because (B.30)-(B.32) will not be influenced by (B.50)-(B.52), we have $V_n \rightarrow v_\infty$ and $\bar{R}_n \rightarrow r(\pi)$ according to Theorem B.3.

Now consider (B.50)-(B.52). Similar as the proof of Theorem B.3, we can combine (B.50)-(B.52) and obtain a single update rule:

$$F_{n+1}(s) = F_n(s) + \beta_{\nu(n,a)} \rho_n(s) \left(\tilde{V}_n(s) + F_n(S'_n(s, A_n(s))) - F_n(s) - \kappa \sum F_n \right) I\{s \in Y_n\}, \quad \forall s \in \mathcal{S}, \quad (\text{B.53})$$

where $\tilde{V}_n(s) \doteq V_n(s) + c$ and $c \doteq \kappa \sum F_0 - \bar{V}_0$. As we discussed above, given Assumption B.4 and B.8, to obtain $V_n - \bar{V}_n e \rightarrow v_\pi$, it only remains to show the convergence of the following synchronous update rule:

$$F_{n+1}(s) = F_n(s) + \beta_n \rho_n(s) \left(\tilde{V}_n(s) + F_n(S'_n(s, A_n(s))) - F_n(s) - \kappa \sum F_n \right), \quad \forall s \in \mathcal{S}. \quad (\text{B.54})$$

Now we rewrite the above equation:

$$F_{n+1}(s) = F_n(s) + \beta_n \rho_n(s) \left(v_\infty(s) + c + (V_n(s) - v_\infty(s)) + F_n(S'_n(s, A_n(s))) - F_n(s) - \kappa \sum F_n \right), \quad \forall s \in \mathcal{S}.$$

Let

$$F_{n+1} = F_n + \beta_n (h(F_n) + M_{n+1} + \epsilon_n), \quad (\text{B.55})$$

where

$$\begin{aligned} h(F_n)(s) &= \sum_a \pi(a | s) \sum_{s'} p(s' | s, a) (v_\infty(s) + c + F_n(S'_n(s, A_n(s)))) - F_n(s) - \kappa \sum F_n, \\ M_{n+1}(s) &= \rho_n(s) \left(v_\infty(s) + c + F_n(S'_n(s, A_n(s))) - F_n(s) - \kappa \sum F_n \right) - h(F_n)(s), \\ \epsilon_n(s) &= \rho_n(s) (V_n(s) - v_\infty(s)). \end{aligned}$$

We first show that without ϵ_n , (B.55) converges. Then we need to show that ϵ_n is bounded and is $o(1)$ so that the last extension of the Section 2.2 of Borkar (2009) can be applied to conclude the convergence of (B.55) with ϵ_n .

Given a new table of estimates $F'_n(s, a), \forall s \in \mathcal{S}, a \in \mathcal{A}$, and the following update rule

$$F'_{n+1} \doteq F'_n + \alpha_n (h(F'_n) + M'_{n+1}), \quad (\text{B.56})$$

where $M'_{n+1}(s) \doteq \rho_n(s) (v_\infty(s) + c + F'_n(S'_n(s, A_n(s))) - F'_n(s) - \kappa \sum F'_n) - h(F'_n)(s)$, and $F'_0 \doteq F_0$.

Lemma B.10 shows that F'_n converges to some point a.s. and (B.47) shows that $\bar{V}'_n \doteq \kappa \sum F'_n - c$ converges to the reward rate following policy π in a new MDP whose transition dynamics is the same as it of the original MDP but the reward from state s is $v_\infty(s)$ instead of $R_n(s, A_n(s))$. From (B.49), the reward rate in the new MDP is $d_\pi^\top v_\infty$. Therefore $\bar{V}'_n = \kappa \sum F'_n - c$ converges to $d_\pi^\top v_\infty$ a.s..

We now show that ϵ_n is bounded and is $o(1)$. ϵ_n is bounded because ρ_n is bounded due to the finite state and action space and Assumption B.7, and V_n is bounded as shown in the proof of Theorem B.3. In addition, because $V_n \rightarrow v_\infty$ and ρ_n is bounded, ϵ_n converges to 0 and thus ϵ_n is $o(1)$.

Given the above results, the last extension of the Section 2.2 of Borkar (2009) applies. In other words, the noise ϵ_n does not change the convergence of F'_n (i.e., $\lim_{n \rightarrow \infty} F'_n = \lim_{n \rightarrow \infty} F_n$). Therefore we conclude that almost surely, F_n converges to some point and $\bar{V}_n = \kappa \sum F_n - c$ converges to $d_\pi^\top v_\infty$.

Because $V_n \rightarrow v_\infty$ and $\bar{V}_n \rightarrow d_\pi^\top v_\infty$, $V_n - \bar{V}_n e \rightarrow v_\infty - d_\pi^\top v_\infty e$. Because $d_\pi^\top v_\infty e$ is a vector with all equal elements, $v_\infty - d_\pi^\top v_\infty e$ satisfies the state-value Bellman equation (B.33). In addition, because $\sum_{s'} (d_\pi(s')(v_\infty(s') - d_\pi^\top v_\infty)) = 0$, from Lemma B.11 we have $v_\infty - d_\pi^\top v_\infty e = v_\pi$. Therefore $V_n - \bar{V}_n e \rightarrow v_\pi$ a.s., as $n \rightarrow \infty$.

The original update rule (B.53) and the synchronous update rule (B.54) converge to the same point. Therefore using the original update rule, $V_n - \bar{V}_n e \rightarrow v_\pi$ a.s. \square

B.3.2. CENTERED DIFFERENTIAL Q-LEARNING AND DIFFERENTIAL Q-PLANNING

Our *Centered Differential Q-learning* maintains, in addition to the first estimator (Equations 5-7), a second estimator in which the reward is the value estimate of the first estimator. The second estimator maintains a scalar offset estimate \bar{Q}_t , an auxiliary table of estimates $F_t(s, a), \forall s \in \mathcal{S}, a \in \mathcal{A}$, and uses the following update rules:

$$F_{t+1}(S_t, A_t) \doteq F_t(S_t, A_t) + \beta_t \Delta_t, \quad \text{and} \quad F_{t+1}(s, a) \doteq F_t(s, a), \forall s \neq S_t, a \neq A_t, \quad (\text{B.57})$$

$$\bar{V}_{t+1} \doteq \bar{V}_t + \kappa \beta_t \Delta_t, \quad (\text{B.58})$$

where

$$\Delta_t \doteq Q_t(S_t, A_t) - \bar{Q}_t + F_t(S_{t+1}, \underset{a'}{\operatorname{argmax}} Q_t(S_{t+1}, a')) - F_t(S_t, A_t), \quad (\text{B.59})$$

is the TD error of the second estimator, $\{\beta_t\}$ is a step size sequence, and κ is a positive constant. β_t and κ can be different from α_t and η . We call (5)-(7) plus (B.57)-(B.59) Centered Differential Q-learning.

The planning version of Centered Differential Q-learning is called *Centered Differential Q-planning*. It uses simulated experience just as in Differential Q-planning. Just like Differential Q-planning, Centered Differential Q-planning maintains Q_n and \bar{R}_n . In addition, Centered Differential Q-planning maintains an auxiliary table of estimates $F_n(s, a)$, for all $s \in \mathcal{S}, a \in \mathcal{A}$, and an offset estimate \bar{Q}_n , and updates them just as in Centered Differential Q-learning, using S_n, A_n, R_n, S'_n instead of $S_t, A_t, R_{t+1}, S_{t+1}$.

Just as we did in Section B.1 and B.2, we now present a general algorithm that includes both Centered Differential Q-learning and Centered Differential Q-planning cases. We call it *General Centered Differential Q*. Using arguments that are similar as those in Section B.1, it can be shown that both Centered Differential Q-learning and Centered Differential Q-planning are special cases of General Centered Differential Q.

For General Centered Differential Q, let the data be generated the same way as it in B.1. Also, we use same notations introduced in B.1. In addition to update rules of General Differential Q (Equation B.1-B.3), General Centered Differential Q has two more update rules:

$$F_{n+1}(s, a) = F_n(s, a) + \alpha_{\nu(n, s, a)} \Delta_n(s, a) I\{(s, a) \in Y_n\} \quad \forall s \in \mathcal{S}, a \in \mathcal{A}, \quad (\text{B.60})$$

$$\bar{Q}_{n+1} = \bar{Q}_n + \sum_{s, a} \alpha_{\nu(n, s, a)} \Delta_n(s, a) I\{(s, a) \in Y_n\}, \quad (\text{B.61})$$

where

$$\Delta_n(s, a) = Q_n(s, a) + F_n(S'_n(s, a), \underset{a'}{\operatorname{argmax}} Q_n(S'_n(s, a), a')) - F_n(s, a) - \bar{Q}_n. \quad (\text{B.62})$$

We now present a convergence theorem for General Centered Differential Q. Unlike the previous theorems, this theorem requires that the optimal policy is unique. The reason is, if there are multiple optimal policies all achieving the optimal average reward, the greedy policy w.r.t. Q_n will jump between these optimal policies even in the limit so the second estimator can not evaluate any particular optimal policy. In addition, unlike the discounted case, where different optimal policies all correspond to the same unique optimal value function, in the average reward case, optimal policies correspond to different differential value functions. Therefore, in order to use the second estimator to evaluate some policy derived from Q_n , that policy must converge as $n \rightarrow \infty$.

In practice, our algorithms can still deal with problems with multiple optimal policies. This can be achieved by choosing a small threshold $\epsilon > 0$, and then replace the $\operatorname{argmax}_a Q(s, a)$ in our algorithms with the first action \tilde{a} satisfying $Q(s, \tilde{a}) \geq \max_a Q(s, a) - \epsilon$. The resulting policy will converge to an optimal policy if ϵ is sufficiently small.

Theorem B.5 (Convergence of General Centered Differential Q). *If Assumption B.1 holds, Assumption B.3, B.4, and B.5 hold for both α_n and β_n , and the optimal policy is unique, denote the differential value function for the optimal policy as q_* , then General Centered Differential Q (Equations B.1-B.3, B.60-B.62) converges, almost surely, \bar{R}_n to r_* and $Q_n - \bar{Q}_n e$ to q_* .*

Proof. Similar as what we did to show Theorem B.4 we use the last extension of section 2.2 of (Borkar 2009) to show this theorem.

Because (B.1)-(B.3) will not be influenced by (B.60)-(B.62), we have $Q_n \rightarrow q_\infty$ and $\bar{R}_n \rightarrow r_*$ a.s., according to Theorem B.1.

Now consider (B.60)-(B.62). Similar as the proof of Theorem B.1, we can combine (B.60)-(B.62) and obtain a single update rule:

$$\begin{aligned} F_{n+1}(s, a) &= F_n(s, a) + \beta_{\nu(n, s, a)} \\ &\left(\tilde{Q}_n(s, a) + F_n(S'_n(s, a), \operatorname{argmax}_{a'} Q_n(S'_n(s, a), a')) - F_n(s, a) - \kappa \sum F_n \right) I\{(s, a) \in Y_n\}, \\ &\forall s \in \mathcal{S}, a \in \mathcal{A}. \end{aligned} \quad (\text{B.63})$$

where $\tilde{Q}_n(s, a) \doteq Q_n(s, a) + c$ and $c \doteq \sum F_0 - \bar{Q}_0$. As we discussed above, given Assumption B.4 and B.5, to obtain $Q_n - \bar{Q}_n \rightarrow q_*$ a.s., it only remains to show the convergence of the following synchronous update rule:

$$\begin{aligned} F_{n+1}(s, a) &= F_n(s, a) + \beta_n \\ &\left(\tilde{Q}_n(s, a) + F_n(S'_n(s, a), \operatorname{argmax}_{a'} Q_n(S'_n(s, a), a')) - F_n(s, a) - \kappa \sum F_n \right), \\ &\forall s \in \mathcal{S}, a \in \mathcal{A}. \end{aligned} \quad (\text{B.64})$$

Rewriting the above equation, we have

$$\begin{aligned} F_{n+1}(s, a) &= F_n(s, a) + \beta_n \\ &\left(q_\infty(s, a) + c + \sum_{s'} p(s' | s, a) F_n(s', \pi_*(s')) - F_n(s, a) - \kappa \sum F_n \right. \\ &+ F_n(S'_n(s, a), \pi_*(S'_n(s, a))) - \sum_{s'} p(s' | s, a) F_n(s', \pi_*(s')) \\ &+ Q_n(s, a) - q_\infty(s, a) + F_n(S'_n(s, a), \operatorname{argmax}_{a'} Q_n(S'_n(s, a), a')) \\ &\left. - F_n(S'_n(s, a), \pi_*(S'_n(s, a))) \right), \end{aligned}$$

where π_* is the unique greedy policy w.r.t. q_∞ , as there is only one optimal policy by our assumption.

Now, let

$$F_{n+1} = F_n + \beta_n (h(F_n) + M_{n+1} + \epsilon_n), \quad (\text{B.65})$$

where

$$\begin{aligned} h(F_n)(s, a) &\doteq q_\infty(s, a) + c + \sum_{s'} p(s' | s, a) F_n(s', \pi_*(s')) - F_n(s, a) - \kappa \sum F_n, \\ M_{n+1}(s, a) &\doteq F_n(S'_n(s, a), \pi_*(S'_n(s, a))) - \sum_{s'} p(s' | s, a) F_n(s', \pi_*(s')), \\ \epsilon_n(s, a) &\doteq Q_n(s, a) - q_\infty(s, a) + F_n(S'_n(s, a), \operatorname{argmax}_{a'} Q_n(S'_n(s, a), a')) \\ &\quad - F_n(S'_n(s, a), \pi_*(S'_n(s, a))). \end{aligned}$$

We will first show that without ϵ_n , (B.65) converges a.s.. Then we will propose a variant of the last extension of the Section 2.2 of (Borkar 2009) and use that show the convergence of (B.65) with ϵ_n .

Given a new table of estimates $F'_n(s, a)$, for all $s \in \mathcal{S}, a \in \mathcal{A}$, consider the following update rule

$$F'_{n+1} = F'_n + \beta_n (h(F'_n) + M'_{n+1}),$$

where $M'_{n+1}(s, a) \doteq F'_n(S'_n(s, a), \pi_*(S'_n(s, a))) - \sum_{s'} p(s' | s, a) F'_n(s', \pi_*(s'))$, and $F'_0 \doteq F_0$.

The above update can be viewed as a special case of (B.44) with $\rho_n = 1$ for a new Markov Reward Process. The state space for this MRP is $\mathcal{S} \times \mathcal{A}$. The transition dynamics of the MRP is defined as $\tilde{p}((s', a') | (s, a)) \doteq \sum \pi_*(a | s) \sum_{s'} p(s' | s, a) \mathbb{I}(a' = \pi_*(s'))$ while the reward starting from (s, a) is $\tilde{r}((s, a)) \doteq q_\infty(s, a)$.

Therefore Lemma B.5 applies and we have that the update F'_n converges to some point satisfying the state-value Bellman equation for this new MRP. By (B.21), $\tilde{Q}'_n = \kappa \sum F'_n - c$ converges to the reward rate for this new MRP, which is $\sum_{s,a} d_{\pi_*}(s, a) q_\infty(s, a)$, the offset in q_∞ w.r.t. q_* by Lemma B.11.

Now, we propose a variant of the last extension of the Section 2.2 of Borkar (2009) and apply it to show that the additional noise ϵ_n does not affect the convergence and therefore $\lim_{n \rightarrow \infty} F_n = \lim_{n \rightarrow \infty} F'_n$ as $n \rightarrow \infty$ and \tilde{Q}_n also converges to $\sum_{s,a} d_{\pi_*}(s, a) q_\infty(s, a)$. The extension of the Section 2.2 of Borkar (2009) requires that ϵ_n is bounded and is $o(1)$. The variant we propose also requires that ϵ_n is $o(1)$, however instead of requiring ϵ_n being bounded, it requires a weaker condition

$$\|\epsilon_n\|_\infty \leq K(1 + \|F_n\|_\infty), \quad (\text{B.66})$$

where K is a positive constant.

This can be shown with the following arguments. 1) If the boundedness of F_n holds, then the conclusion of Lemma 1 of Section 2 of Borkar (2009) will not be affected and therefore the convergence of F_n remains unchanged. 2) The boundedness of F_n can be shown with the following three modifications of the proofs in Section 3 of Borkar (2009):

1. It can be seen that the claim of Lemma 4 in Section 3.2 of Borkar (2009) remains unchanged with this additional noise ϵ_n .
2. A result similar to Lemma 5 in Section 3.2 of Borkar (2009) can be shown for this additional noise. That is, the sequence $\tilde{\zeta}'_n \doteq \sum_{k=0}^{n-1} a_k \tilde{\epsilon}_k$, $n \geq 1$ is a.s. convergent, where a_k are the stepsizes, $\tilde{\epsilon}_k = \epsilon_k / r(n)$ for $m(n) \leq k < m(n+1)$ and $r(\cdot)$ and $m(\cdot)$ are defined in Section 3.2 of Borkar (2009). This is due to Assumption B.3 and also ϵ_n being $o(1)$.
3. Lemma 6 of Section 3.2 of Borkar (2009) holds with the additional ϵ_n .

Now let us verify if (B.66) holds and if ϵ_n is $o(1)$. It can be seen that (B.66) is satisfied because Q_n is bounded as we showed in the proof of Lemma B.5. In addition, because $Q_n \rightarrow q_\infty$ a.s., $\epsilon_n \rightarrow 0$ a.s. as $n \rightarrow \infty$ and thus ϵ_n is $o(1)$. Therefore a.s., F_n converges and \tilde{Q}_n converges to $\sum_{s,a} d_{\pi_*}(s, a) q_\infty(s, a)$, the offset of q_∞ . Therefore $Q_n - \tilde{Q}_n e \rightarrow q_*$ a.s..

Q_n in the original update rule (B.63) and Q_n in the synchronous update rule (B.64) converge to the same point. Therefore using the original update rule, $Q_n - \tilde{Q}_n e \rightarrow q_*$ a.s.. \square

C. Additional Experiments and Experimental Details

Here we provide the remaining details of all the experiments reported in this paper. We also present additional experiments that are pertinent to this paper. In particular, the following sections contain:

1. Details of the control experiments in Section 3
2. An experiment demonstrating the max reference function is not always the best choice of the reference function for RVI Q-learning
3. An experiment demonstrating RVI Q-learning diverges when the reference state is transient
4. Details of the prediction experiments in Section 5 (both on- and off-policy)
5. An empirical demonstration of the centering technique introduced in Section 6

All the experiment code is available at <https://github.com/abhisheknaik96/average-reward-methods>.

C.1. Details of the Control Experiments on the Access-Control Queuing Task

In this section, we provide the rest of the experimental details for the control experiments on the Access-Control Queuing Task in Section 3 of the main text.

The task starts with all 10 servers free. With four types of customers, 11 possible number of free servers (0 to 10), and two actions, there are a total of 88 state–action pairs. The value function for both algorithms and the reward rate estimate for Differential Q-learning were initialized to zero. Both algorithms were run with step size α in the range $\{0.0015625, 0.00625, 0.025, 0.1, 0.4\}$. For Differential Q-learning, η was chosen from $\{0.125, 0.25, 0.5, 1, 2\}$. The reference functions were chosen as mentioned in Section 3. Both algorithms used an ϵ -greedy behavior policy with $\epsilon = 0.1$ and no annealing.

The learning curve in Figure 1 corresponds to the parameters for Differential Q-learning that resulted in the largest reward rate averaged over the training period of 80,000 steps: $\alpha = 0.025$ and $\eta = 0.125$. A point on the solid curves denotes the reward rate during training computed over a sliding window of previous 2000 rewards, and the shaded region denotes one standard error.

C.2. Max is Not Always the Best Choice of Reference Function for RVI Q-learning

In Section 5 we pointed out that RVI Q-learning performs well if the reference state (or state–action pair) occurs frequently under an optimal policy. This led to the speculation that perhaps the state–action pair with the highest action-value estimate might be the best choice of a reference state–action pair because under the RL paradigm, an agent seeks to visit the highly-rewarding states. We gave an example in the main text that it is not true in general that the state–action pair with the highest action-value estimate also occurs frequently under an optimal policy. In this section, we show this empirically; max is not always the best choice of the reference function for RVI Q-learning.

The two domains used are variants of the Two Loop MDP (from Section 5). The first variant has the same transitions as in the Two Loop MDP except there is a +10 reward when going from state 8 to state 0 instead of +2. The optimal policy is to take the action `right` in state 0 and obtain a reward rate of +2 per step. The second variant builds on the first variant in that there is an additional state 9. Starting from any state other than state 9, no matter what action is taken, there is a 0.02 probability of moving to state 9 with 0 reward and a 0.98 probability of moving to a state with a reward just as in the first variant. From state 9, the action deterministically leads to state 0 with a reward of +100; this makes state 9 a high-value state. But it rarely occurs under the optimal policy, which is again to take the action `right` in state 0. The optimal reward rate in this second domain is 3.84.

In addition to RVI Q-learning with the max reference function, we also ran Differential Q-learning as a baseline on these two domains. The value function for both algorithms and the reward rate estimate for Differential Q-learning were initialized to zero. Both algorithms were run with step size α in the range $\{0.003125, 0.00625, 0.0125, 0.025, 0.05, 0.1, 0.2, 0.4\}$. For Differential Q-learning, η was chosen from $\{0.125, 0.25, 0.5, 1, 2\}$. Both algorithms used an ϵ -greedy behavior policy with $\epsilon = 0.1$ and no annealing. The experiments were run for 100000 steps and repeated 30 times.

The sensitivity plots for both domains are shown in Figure C.2. The performance of RVI Q-learning was quite different qualitatively in both domains. In the first variant, the max reference function resulted in the best performance. In this case,

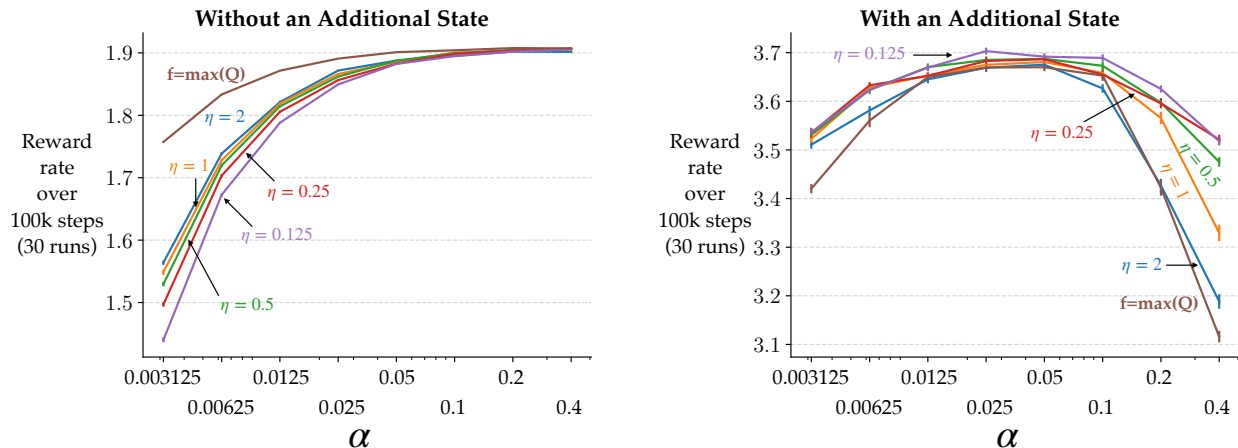


Figure C.2: Parameter studies of RVI Q-learning using max as the reference function and Differential Q-learning with various values of η . *Left*: In the first domain in which the state–action pair with the highest action value occurs frequently under the optimal policy, RVI Q-learning using a max reference function performed well across a wide range of step sizes. *Right*: On the other hand, in the second domain, the state–action pair with the highest action value occurs rarely under the optimal policy, and RVI Q-learning using a max reference function performs well only for a relatively narrow range of step sizes. The domains are described in the text.

the state–action pair with the highest action value (state 8) occurs frequently under the optimal policy of taking the right loop. But in the second variant, the state–action pair corresponding to the highest action value (state 9) occurs rarely under the optimal policy. As expected, the max reference function did not result in good performance in this case. In fact, the rate of learning for Differential Q-learning was better than or equal to that of RVI Q-learning with the max reference function for almost the whole range of parameters tested.

These experiments show that the value of the state–action pair with the highest action value is not in general the best choice of the reference function for RVI Q-learning.

C.3. RVI Q-learning Diverges when the Reference State is Transient

In this experiment, we show that RVI Q-learning diverges if the reference state is a *transient* state of the MDP, that is, the state does not occur more than a finite number of times under any policy. Note that transient states are not allowed in communicating MDPs but are allowed in the unichain MDPs and the more general weakly-communicating MDPs. While our theory was developed for communicating MDPs, it can be extended with some modification to the more general weakly-communicating MDP case (see the discussion right after Assumption B.5 for more details about this extension). The convergence results for RVI Q-learning (Abounadi et al. 2001) were developed for the unichain case, but we show via an experiment that RVI Q-learning can diverge under certain conditions. In particular, when the reference state is transient.

The domain is a simple two-state MDP with the transition and reward dynamics shown in Figure C.3 (left). State 0 is transient under all stationary policies (including the optimal policy), meaning it only occurs a finite number of times before it is never seen again.

The behavior policy is random. The value function for both algorithms and the reward rate estimate for Differential Q-learning was initialized to zero. The reference state–action pair was set to be action a in state 0. The step sizes were all set to a value of 0.01 (an arbitrary choice; this effect can be observed for any positive step size). The starting state was state 0, and the experiments were run for 1000 steps and repeated 50 times.

Figure C.3 (right) shows the evolution of the learned value estimates over time (the standard error is smaller than the width of the line representing the mean). The value of the reference state–action pair $Q(0, a)$ cannot reach the optimal reward rate of 2 and hence the under-estimation leads to divergence in the estimate of the recurrent state which is updated as $Q_{t+1}(2, a) = Q_t(2, a) + \alpha(2 - Q(0, a) + Q_t(2, a) - Q_t(2, a))$ (refer to Algorithm 2).

This simple experiment demonstrates that RVI Q-learning diverges when the reference state–action pair is transient.




Figure C.3: Demonstration of divergence in RVI Q-learning when the reference state is transient. *Left*: The two-state MDP in which state 0 is transient under all policies. *Right*: Comparison of estimated values with RVI Q-learning and Differential Q-learning algorithms on the two-state MDP. The value of the recurrent state diverges in case of RVI Q-learning, whereas all the estimates converge in case of Differential Q-learning. The solid lines denote the mean, and one standard error is less than the width of the lines.

C.4. Details of the Prediction Experiments

This section presents the supplementary material for the experiments in Section 5: the remaining experimental details, how the evaluation metric is computed, the sensitivity plots of the reward-rate error (RRE) for on-policy Differential TD-learning and Average Cost TD-learning, and the sensitivity plots for RMSVE (TVR) and RRE for off-policy Differential TD-learning.

The step size α and the parameter η for all three algorithms (Average Cost TD-learning, on- and off-policy Differential TD-learning) were chosen from $\{0.025, 0.05, 0.1, 0.2, 0.4\}$ and $\{0.125, 0.25, 0.5, 1, 2\}$ respectively. The step sizes were decayed by a factor of 0.9995 at each step. The value estimates and the reward-rate estimate for all algorithms were initialized to zero. The learning curves for on-policy Differential TD-learning and Average Cost TD-learning (blue and orange) on the top-right of Figure 3 correspond to the parameters that minimized the average RMSVE (TVR) over the training period, which reflects their rate of learning: $\alpha = 0.2$ and $\eta = 0.25$ for Differential TD-learning, and $\alpha = 0.1$ and $\eta = 0.125$ for Average Cost TD-learning. The learning curve for off-policy Differential TD (green) in the same plot is plotted for the parameters that resulted in the minimum asymptotic RMSVE (TVR) computed over the last 5000 steps of training: $\alpha = 0.2$ and $\eta = 0.5$. In all the plots, the solid line represents the mean, and the error bars indicate one standard error (which in many cases was less than the width of the solid lines).

We now describe the evaluation metric in more detail — the variant of RMSVE originally proposed by Tsitsiklis and Van Roy (1999) (hence the abbreviation ‘RMSVE (TVR)’). As noted in Section 4, there are multiple solutions to the Bellman equations for the differential value function of the form $v_\pi(s) + c$, where $c \in \mathbb{R}$. All algorithms converge to one of these solutions depending on design choices such as initializations and reference states. Therefore, computing the value error w.r.t. the actual value function v_π does not say much about convergence. Tsitsiklis and Van Roy proposed computing the error w.r.t. the nearest valid solution to the Bellman equations — this error would be zero for any valid solution to the Bellman equations that an algorithm converges to. Mathematically, this error is given by:

$$\inf_c \|v - (v_\pi + ce)\|_{d_\pi} = \|\mathcal{P}v - v_\pi\|_{d_\pi},$$

where \mathcal{P} is a projection operator and d_π is the stationary state distribution corresponding to the policy π . Algorithmically, this translates to computing the offset of the learned value function, subtracting it, and then computing the RMSVE w.r.t. the actual value function v_π . The offset can be computed by simply taking a dot product of the learned value function and d_π : $d_\pi^T(v_\pi + ce) = d_\pi^T v_\pi + c d_\pi^T e = 0 + c = c$, where $d_\pi^T v_\pi = 0$ (from Lemma B.11).

For the target policy π that uniformly randomly picks one of the two actions in state 0, $d_\pi = [0.2, 0.1, 0.1, 0.1, 0.1, 0.1, 0.1, 0.1, 0.1]^T$, and $v_\pi = [-0.2, -1.4, -1.1, -0.8, -0.5, 0.6, 0.9, 1.2, 1.5]^T$, which can be obtained by solving the Bellman equations along with the constraint $d_\pi^T v_\pi = 0$.

The top row in Figure C.4 shows the sensitivity of RRE of on-policy Differential TD-learning and Average Cost TD-learning averaged over the training period w.r.t. its parameters for the experiment in Section 5. On-policy Differential TD-learning was less sensitive to both α and η , and converged to a lower RRE across a large range of its parameters as compared to

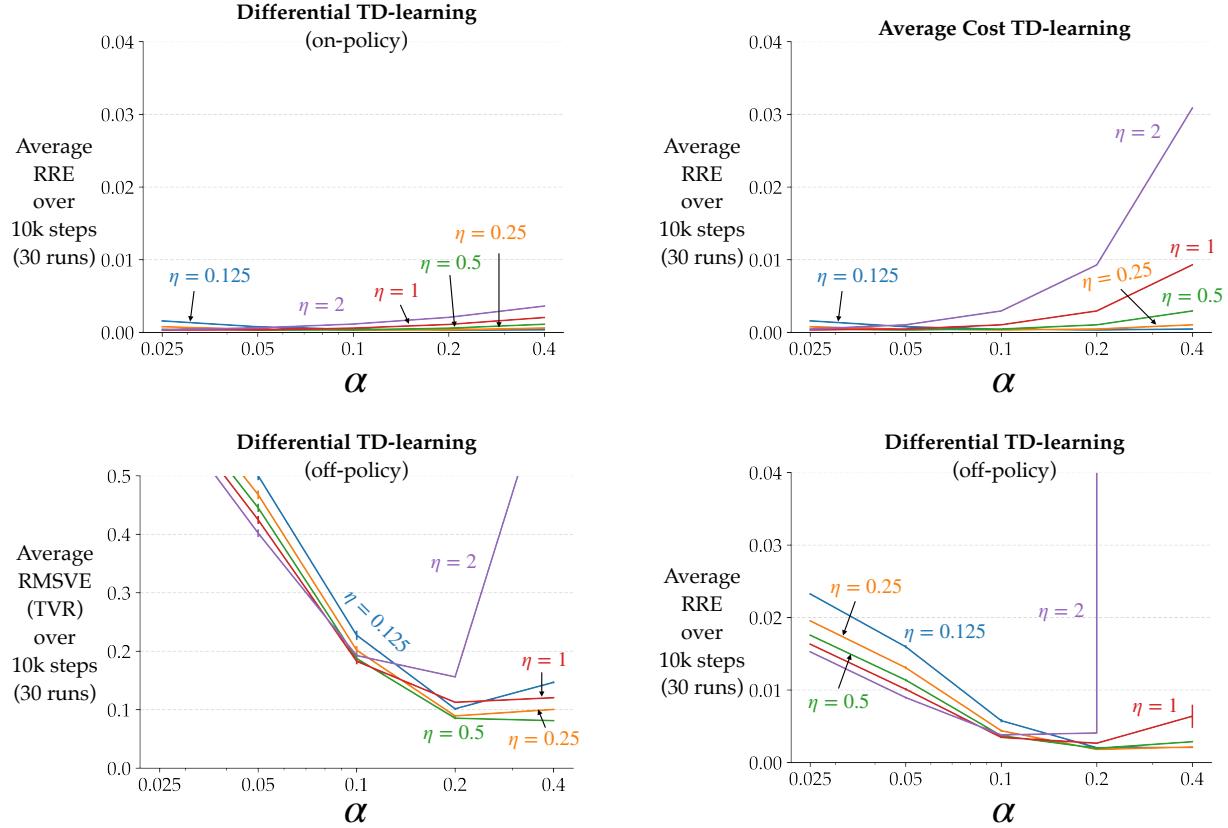


Figure C.4: Parameter studies for the prediction experiments on the Two Loop task. The solid lines denote the mean, and one standard error is less than the width of the lines. *Top*: On-policy Differential TD-learning achieves equal or lower average RRE than Average Cost TD-learning for a broad range of parameters. *Bottom*: Sensitivity of off-policy Differential TD-learning’s performance in terms of average RMSVE (TVR) and average RRE w.r.t. parameters α and η .

Average Cost TD-learning. A similar trend was observed for RMSVE (TVR) and discussed in Section 5.

The bottom row of Figure C.4 shows the sensitivity of both RMSVE (TVR) and RRE of off-policy Differential TD-learning w.r.t. its parameters α and η . The rate of convergence is affected if the parameters are too high or too low, but otherwise it is relatively insensitive to different choices of η .

C.5. Estimating the Actual Differential Value Function

In this section, we demonstrate that the technique introduced in Section 6 of the main text results in the estimation of the actual (centered) differential value function. As mentioned earlier, the technique is general and can be used with any average-reward algorithm proposed in this paper. For illustration, we use this technique with Differential Q-learning on the Two Loop task described in Section 5 of the main text.

We first applied RVI Q-learning and (uncentered Differential Q-learning on this problem). RVI Q-learning was run with the reference function which is the value of a single reference state–action pair, $f = Q(s_0, a_0)$, for which we considered all possible state–action pairs as s_0 and a_0 : state 0 with actions `left` and `right`, and states 1–8 with the action a . Both algorithms were run with a range of step sizes $\alpha \in \{0.025, 0.5, 0.1, 0.2, 0.4\}$, each for 30 runs and 10,000 steps. Differential Q-learning was run with a range of η values: $\{0.125, 0.25, 0.5, 1, 2\}$. Both algorithms had an ϵ -greedy behavior policy with $\epsilon = 0.1$ and no annealing. The value function for both the algorithms and the reward rate estimate for Differential Q-learning was initialized to zero. We then applied Centered Differential Q-learning (refer to Algorithm 5 for the pseudocode) on the Two Loop MDP, for 30 runs and 10,000 steps with the parameters β and κ chosen from $\{0.025, 0.05, 0.1, 0.2, 0.4\}$ and $\{0.125, 0.25, 0.5, 1, 2\}$ respectively. α and η were chosen as ones that achieved performed well for (uncentered) Differential Q-learning ($\alpha = 0.4, \eta = 0.5$) because the centering technique in its current form learns the offset separately and does

Table C.1: The stationary state–action distribution and the action-value function for the optimal policy of choosing action right in state 0 in the Two Loop task.

	state–action pair									
	(0,left)	(0,right)	(1,a)	(2,a)	(3,a)	(4,a)	(5,a)	(6,a)	(7,a)	(8,a)
d_π	0	0.2	0	0	0	0	0.2	0.2	0.2	0.2
q_π	-1.8	-0.8	-2.4	-2.0	-1.6	-1.2	-0.4	0.0	0.4	0.8

not affect the value estimates during learning. The learned offset was subtracted from the value estimates every time the RMSVE² was computed. The stationary state–action distribution and the action-value function for the optimal policy is shown in Table C.1, which were again obtained by solving the Bellman equations with the constraint that $d_\pi^T v_\pi = 0$. The estimates of the secondary estimator were also initialized to zero in every run.

A learning curve for each of the algorithms is shown in the left side of Figure C.5. For RVI Q-learning and (uncentered) Differential Q-learning, these are corresponding to the parameter settings that resulted in the largest reward rate averaged over the training period (reference state 8, action a with $\alpha = 0.4$ for RVI Q-learning and $\alpha = 0.4$ with $\eta = 0.5$ for Differential Q-learning). For Centered Differential Q-learning, the learning curve is plotted for the parameters that resulted in the lowest RMSVE averaged over the course of training ($\beta = 0.4, \kappa = 0.125$).

We saw that Centered Differential Q-learning converged to a differential value function with zero RMSVE, in other words, the centering technique proposed in Section 6 of the main text succeeds in estimating the offset correctly, which is subtracted from the value estimates to result in the centered differential value function. RVI Q-learning and (uncentered) Differential Q-learning also converged to some particular value functions with some offset from the centered differential value function. Note that there was a lot of variance in the values estimated by RVI Q-learning till it converged after about 6000 steps. The source of this variance needs further investigation. Also shown on the right of Figure C.5 is the sensitivity of the performance of the centering technique to its two parameters β and κ . We saw that in this task where the transitions were mostly deterministic, larger step sizes β could be used, and the value of κ only had a small effect.

This experiment shows that the technique introduced in Section 6 of the main text can learn the centered differential value function. We demonstrated this with Differential Q-learning, an off-policy control learning algorithm, and we expect this technique to work with other combinations of settings as well: on-policy and off-policy, prediction and control, learning and planning.

²Note that we no longer need Tsitsiklis and Van Roy’s (1999) variant of the RSMVE (which we earlier used and denoted as ‘RMSVE (TVR)’) because now we want to compute the error w.r.t. the actual differential value function.




Figure C.5: *Left*: On the Two Loop task, Centered Differential Q-learning learned the centered differential value function corresponding to the optimal policy, while RVI Q-learning and (uncentered) Differential Q-learning converged to some offset versions of the centered differential value function. *Right*: Parameter study showing the centering technique was relatively robust to the parameter κ and resulted in a low average RMSVE for a broad range of its parameters. All solid curves denote the mean, and the shaded region or error bars denote one standard error.

D. Additional Discussion

D.1. Yang et al.'s (2016) convergence results are incorrect

We show that the proofs of two lemmas leading up to the convergence theorem of CSV-learning are not valid. The proofs would be valid if both the transition and reward dynamics of the MDP are deterministic. However the assumption of the MDP being deterministic is not made in the paper. We begin by presenting the relevant assumptions, definitions, and lemmas from the paper (Yang et al. 2016).

Assumption 1: *The MDP is irreducible, aperiodic, and ergodic; that is, under any stationary policy, the generated Markov chain is communicating and has a recurrent state.*

Definition 1: $d_t(s', s) \doteq \max_{a'} Q_t(s', a') - \max_a Q_t(s, a)$

Lemma 1: *In the Markov chain M_t under π_t , $d_t(s', s) = \rho_{\pi_t} - r(s, \pi_t(s))$.*

Let us see why this lemma is incorrect:

$$\begin{aligned}
 d_t(s', s) &\doteq \max_{a'} Q_t(s', a') - \max_a Q_t(s, a) \\
 &= \max_{a'} Q_t(s', a') - Q_t(s, \pi_t(a)) \\
 &= \max_{a'} Q_t(s', a') - \sum_{s^n, r} p(s^n, r \mid s, \pi_t(s)) [r - \rho_{\pi_t} + \max_{a^n} Q(s^n, a^n)] \\
 &= \max_{a'} Q_t(s', a') - [r(s, \pi_t(s)) - \rho_{\pi_t} + \sum_{s^n, r} p(s^n, r \mid s, \pi_t(s)) \max_{a^n} Q(s^n, a^n)] \\
 &= \rho_{\pi_t} - r(s, \pi_t(s)) + \max_{a'} Q_t(s', a') - \sum_{s^n, r} p(s^n, r \mid s, \pi_t(s)) \max_{a^n} Q(s^n, a^n) \\
 &\neq \rho_{\pi_t} - r(s, \pi_t(s))
 \end{aligned}$$

The equality does not hold in general, only if there is a deterministic transition from state s to s' .

Lemma 2: *If π_t is stable, then $\rho_{\pi_t} \geq \hat{\rho}$. ($\hat{\rho}$ is the constant or fixed estimate of the reward rate used by the CSV-learning algorithm)*

This is also incorrect because:

$$\begin{aligned}
 \Delta Q_t(s, \pi_t(s)) &= \alpha(r - \hat{\rho} + \max_{a'} Q_t(s', a') - Q_t(s, \pi_t(s))) \\
 &= \alpha(r - \hat{\rho} + \max_{a'} Q_t(s', a') - \max_a Q_t(s, a)) \\
 &= \alpha(r - \hat{\rho} + d_t(s', s)) \quad (\text{Definition 1}) \\
 &= \alpha(r - \hat{\rho} + \rho_{\pi_t} - r(s, \pi_t(s))) \quad (\text{Assuming a deterministic transition from } s \text{ to } s') \\
 &\neq \alpha(\rho_{\pi_t} - \hat{\rho})
 \end{aligned}$$

Again, the equality does not hold in general, but only if the rewards are deterministic as well. In other words, the expected reward from a given state after taking an action according to the policy ($r(s, \pi_t(s))$) is equal to the immediate reward (r). Note we also had to assume there is a deterministic transition from s to s' .

Thus, Theorem 1 only holds if the both the transition and reward dynamics of the MDP are deterministic. The paper does not state this assumption³, which invalidates the proof.

D.2. Average Cost TD-learning cannot be extended to the off-policy case by adding an importance-sampling ratio

Average Cost TD-learning cannot be extended to the off-policy setting by simply adding an importance-sampling (IS) ratio as it only corrects the mismatch in targets due to misalignment between actions taken by the target and behavior policies. The IS ratio does not correct the mismatch in distribution of updates, which is also required. However, using the TD error instead of the conventional error as in Differential TD-learning to update the reward-rate estimate only requires correction to

³In any case, assuming MDPs are completely deterministic is a significantly restrictive and unrealistic assumption.

the mismatch in targets and not to the mismatch in distribution of updates, and hence the addition of the IS ratio suffices. Both of these claims are substantiated below.

Consider the update made by the Average Cost TD-learning algorithms to the reward-rate estimate in the on-policy setting:

$$\bar{R}_{t+1} = \bar{R}_t + \eta\alpha_t(R_{t+1} - \bar{R}_t) \quad (\text{D.1})$$

At convergence the expected update is zero:

$$\begin{aligned} 0 &= \mathbb{E}[R_{t+1} - \bar{R}_t] \\ 0 &= \sum_s d_\pi(s) \sum_a \pi(a|s) \sum_{s',r} p(s',r | s,a)(r - \bar{R}_\infty) \\ 0 &= \sum_s d_\pi(s) \sum_a \pi(a|s) \sum_{s',r} p(s',r | s,a) r - \bar{R}_\infty \\ 0 &= r(\pi) - \bar{R}_\infty \quad (\text{by definition}) \\ \implies \bar{R}_\infty &= r(\pi) \end{aligned}$$

where $d_\pi(s)$ is the steady-state distribution over states when following policy π , and \bar{R}_∞ is the point where \bar{R}_t converges, and $r(\pi)$ is the true reward rate of the policy π . The above equations show that in the on-policy setting, the reward-rate estimate of Average Cost TD-learning converges to the true reward rate of the target policy.

Adding an importance-sampling ratio to Average Cost TD-learning to extend it to the off-policy setting does not work because the point of convergence is no longer the reward rate of the target policy π when following behavior policy b .

Proposed off-policy reward-rate update: $\bar{R}_{t+1} = \bar{R}_t + \eta\alpha_t\rho_t(R_{t+1} - \bar{R}_t)$.

Following a similar analysis at the point of convergence:

$$\begin{aligned} 0 &= \mathbb{E}[\rho_t(R_{t+1} - \bar{R}_t)] \\ 0 &= \sum_s d_b(s) \sum_a b(a|s) \sum_{s',r} p(s',r | s,a) \frac{\pi(a|s)}{b(a|s)} (r - \bar{R}_\infty) \\ 0 &= \sum_s d_b(s) \sum_a \pi(a|s) \sum_{s',r} p(s',r | s,a) (r - \bar{R}_\infty) \\ 0 &= \sum_s d_b(s) \sum_a \pi(a|s) \sum_{s',r} p(s',r | s,a) r - \bar{R}_\infty \\ \implies \bar{R}_\infty &= \sum_s d_b(s) \sum_a \pi(a|s) \sum_{s',r} p(s',r | s,a) r \\ &\neq r(\pi) \end{aligned}$$

With the proposed off-policy reward-rate update for Average Cost TD-learning, the point of convergence for the reward-rate estimate is no longer the reward rate of the target policy π because the IS ratio only corrects for the mismatch in targets and not the mismatch in the distribution of updates.

Now consider the update to the reward-rate estimate in case of off-policy Differential TD-learning:

$$\bar{R}_{t+1} = \bar{R}_t + \eta\alpha_t\rho_t(R_{t+1} - \bar{R}_t + V(S_{t+1}) - V(S_t)) \quad (\text{D.2})$$

At convergence,

$$\begin{aligned}
 0 &= \mathbb{E}\left[\rho_t(R_{t+1} - \bar{R}_t + V_t(S_{t+1}) - V_t(S_t))\right] \\
 0 &= \sum_s d_b(s) \sum_a b(a|s) \sum_{s',r} p(s', r | s, a) \frac{\pi(a|s)}{b(a|s)} (r - \bar{R}_\infty + v_\infty(s') - v_\infty(s)) \\
 0 &= \sum_s d_b(s) \sum_a \pi(a|s) \sum_{s',r} p(s', r | s, a) (r - \bar{R}_\infty + v_\infty(s') - v_\infty(s)) \\
 0 &= \sum_s d_b(s) \left[\sum_a \pi(a|s) \sum_{s',r} p(s', r | s, a) (r - \bar{R}_\infty) + \sum_a \pi(a|s) \sum_{s',r} p(s', r | s, a) (v_\infty(s') - v_\infty(s)) \right] \\
 0 &= \sum_s d_b(s) \left[\sum_a \pi(a|s) \sum_{s',r} p(s', r | s, a) (r - \bar{R}_\infty) + \sum_a \pi(a|s) \sum_{s',r} p(s', r | s, a) (r - r(\pi)) \right] \\
 0 &= \sum_s d_b(s) \sum_a \pi(a|s) \sum_{s',r} p(s', r | s, a) (r - \bar{R}_\infty - r + r(\pi)) \\
 0 &= \sum_s d_b(s) \sum_a \pi(a|s) \sum_{s',r} p(s', r | s, a) (r(\pi) - \bar{R}_\infty) \\
 0 &= r(\pi) - \bar{R}_\infty \\
 \implies \bar{R}_\infty &= r(\pi)
 \end{aligned}$$

The above holds due to the v_∞ being a solution to the Bellman equation by Theorem 3: $v_\infty(s) = \sum_a \pi(a|s) \sum_{s',r} p(s', r | s, a) (r - r(\pi) + v_\infty(s'))$.

Hence, Average Cost TD-learning is inherently an on-policy algorithm. It cannot be extended to the off-policy case simply by using an importance-sampling ratio in its updates. On the other hand, the usage of the TD-error to update the reward-rate estimate instead of the conventional error enables Differential TD-learning to converge to the right solution with the usage of the importance-sampling ratio in the off-policy case.

D.3. Discussion about other off-policy prediction methods in the literature

In this section, we shed some light on why existing off-policy prediction methods for average-reward MDPs (Liu et al. 2018, Tang et al. 2019, Mousavi et al. 2020, Zhang et al. 2020a,b) are not guaranteed to converge to the reward rate of the target policy. The primary characteristic of these methods is that they estimate only the reward rate and not the differential value function. To estimate the reward rate, they first estimate the ratio of the stationary distribution under the target policy and under the behavior policy, and then use that to estimate the reward rate using. The first step is difficult, after which the second step is straightforward. These methods use different ways to estimate the ratio using a batch of data.

While these methods are developed for the function approximation setting, none of them can guarantee convergence to the reward rate even in the tabular setting with an infinite-sized batch of data. Liu et al. (2018), Tang et al. (2019), and Mousavi et al. (2020) used a self-normalization trick in their methods, which typically leads to a biased solution (Zhang et al. 2020a). Zhang et al. (2020a) used a primal-dual approach and their algorithm optimizes a min-max objective; however, this objective may not be convex-concave even in the tabular setting (Zhang et al. 2020b). Therefore there could be multiple solutions to that objective and their algorithm will not in general obtain the one corresponding to the true reward rate. Finally, Zhang et al. (2020b) proposed optimizing a convex-concave saddle-point problem that has a unique solution. However, because of a regularization term in the objective, the unique solution of the objective in general does not yield the true reward rate—only a biased one—even in the tabular case.

E. Extensions of our Differential Algorithms to the Setting of Linear Function Approximation

We consider the setting in which at each timestep t , the agent observes a feature vector \mathbf{x}_t representing the state of the environment, takes a discrete action A_t , observes the next feature vector \mathbf{x}'_t and the scalar reward signal R_{t+1} . The agent approximates the action-value function at each timestep t as a linear function of the feature vector: $\hat{q}_t \doteq \mathbf{w}_{A_t}^\top \mathbf{x}_t$ (for feature vectors of dimension d , the action-value function is parameterized with $|\mathcal{A}|$ number of d -dimensional weight vectors).

Algorithm 6: Differential Q-learning with linear function approximation

Algorithm parameters: step-size parameters α, η

```

1 Initialize  $\mathbf{w}_a \in \mathbb{R}^d \forall a$  and  $\bar{R} \in \mathbb{R}$  arbitrarily (e.g., to zero)
2 Obtain initial observation vector  $\mathbf{x}$ 
3 for each timestep do
4     Take action  $A$  (using, say, an  $\epsilon$ -greedy policy w.r.t.  $\hat{q}$ ), obtain  $R, \mathbf{x}'$ 
5      $\delta = R - \bar{R} + \max_a \mathbf{w}_a^\top \mathbf{x}' - \mathbf{w}_A^\top \mathbf{x}$ 
6      $\mathbf{w}_A \leftarrow \mathbf{w}_A + \alpha \delta \mathbf{x}$ 
7      $\bar{R} \leftarrow \bar{R} + \eta \alpha \delta$ 
8      $\mathbf{x} = \mathbf{x}'$ 
9 end
    
```

This algorithm is a straightforward extension of the tabular version of Differential Q-learning (Algorithm 1). Similar extensions exist for Differential Q-planning, Differential TD-learning, Differential TD-planning, and the centered versions of these algorithms.

E.1. Extension of the notion of reference functions to the function approximation setting is not as straightforward

Compared to Differential Q-learning, extensions of the tabular version of RVI Q-learning (Abounadi et al. 2001) to the case of (linear) function approximation are not as straightforward. RVI Q-learning requires the value of a reference function f to be computed at every timestep t , where f is a function over the current estimates of the value estimates \hat{q}_t . Some difficulties that arise with the first attempts of extending the reference functions suggested by Abounadi et al. to the function approximation setting:

- Reference function is the mean of all action-value estimates: $f(\hat{q}_t) \doteq \frac{1}{|\mathcal{S}||\mathcal{A}|} \sum_{s,a} \hat{q}_t(s, a)$
It is easy to see why the computation of this quantity is problematic in the function approximation setting: unlike the tabular setting, the agent does not have access to the underlying states.⁴
- Reference function is the max of all action-value estimates: $f(\hat{q}_t) \doteq \max_{s,a} \hat{q}_t(s, a)$
In the tabular setting, it is straightforward to compute the max of the action-value function over all state–action pairs. In the function approximation setting, computing this quantity would again require access to all the underlying states, which the agent does not have.
- Reference function is the action-value estimate of a single reference state–action pair (s_0, a_0) : $f(\hat{q}_t) \doteq \hat{q}_t(s_0, a_0)$
Again, the agent does not have access to any underlying state in the function approximation setting. Instead, one might consider using a value of a *reference feature vector* with an action as the reference function. The question then becomes what the reference feature vector should be, among the infinite choices in \mathbb{R}^d .
Based on our observations in the tabular setting, we hypothesize that the performance of RVI Q-learning with a reference feature vector would depend on the frequency with feature vectors similar to the reference feature vector occur under the optimal policy for the given problem.

Based on the above discussion, we can attempt to create a couple of reference functions for the function approximation setting, for instance, the action-value estimate corresponding to the *first* feature vector the agent observes along the action of moving left. There is no way to compute the max exactly, but perhaps we can try using the maximum of the set of estimated action values corresponding to the feature vectors when they are observed. These are just some first attempts; further research is required to develop theoretically-grounded reference functions for the function approximation setting.

⁴An alternative problem setting when function approximation is used is when the agent does have access to the underlying states, but they are too many to enumerate (e.g., in a table). In this case $|\mathcal{S}|$ is either unknown, or too large (making $\frac{1}{|\mathcal{S}|}$ too small).

If we have good ways of computing such reference functions at each timestep, the linear function approximation version of RVI Q-learning would look similar to Algorithm 6, except the TD-error term in line 5 would be: $\delta = R - f(\hat{q}) + \max_a \mathbf{w}_a^T \mathbf{x}' - \mathbf{w}_A^T \mathbf{x}$, and there would be no update to a reward-rate estimate like in line 7.

E.2. Preliminary experimental results in the linear function approximation setting

We performed a couple of preliminary experiments in the linear function approximation setting using the PuckWorld and Catcher domains from the PyGame Learning Environment⁵.

In the PuckWorld problem, the agent needs to reach the position marked by the green circle, which moves to a different location after every few timesteps. A still of the environment is shown in the top-left panel of Figure E.1. At every timestep, the agent can take one of four actions — left, right, up, down — which move the agent in that direction by a small amount. Repeated actions in the same direction build some velocity in that direction, which decays at an exponential rate at every timestep. At every timestep, the agent gets a reward proportional to its distance to the goal position. This reward is typically negative and becomes zero when the agent reaches the goal position. At every timestep, the agent observes a six-dimensional feature vector of its horizontal position, vertical position, horizontal velocity, vertical velocity, target’s horizontal position, target’s vertical position. The positions and velocities are scaled to lie in $[0, 1]$ and $-[1, 1]$ respectively. After a regular interval of timesteps, the goal position is uniform-randomly initialized in the two-dimensional space.

⁵<https://github.com/ntasfi/PyGame-Learning-Environment>

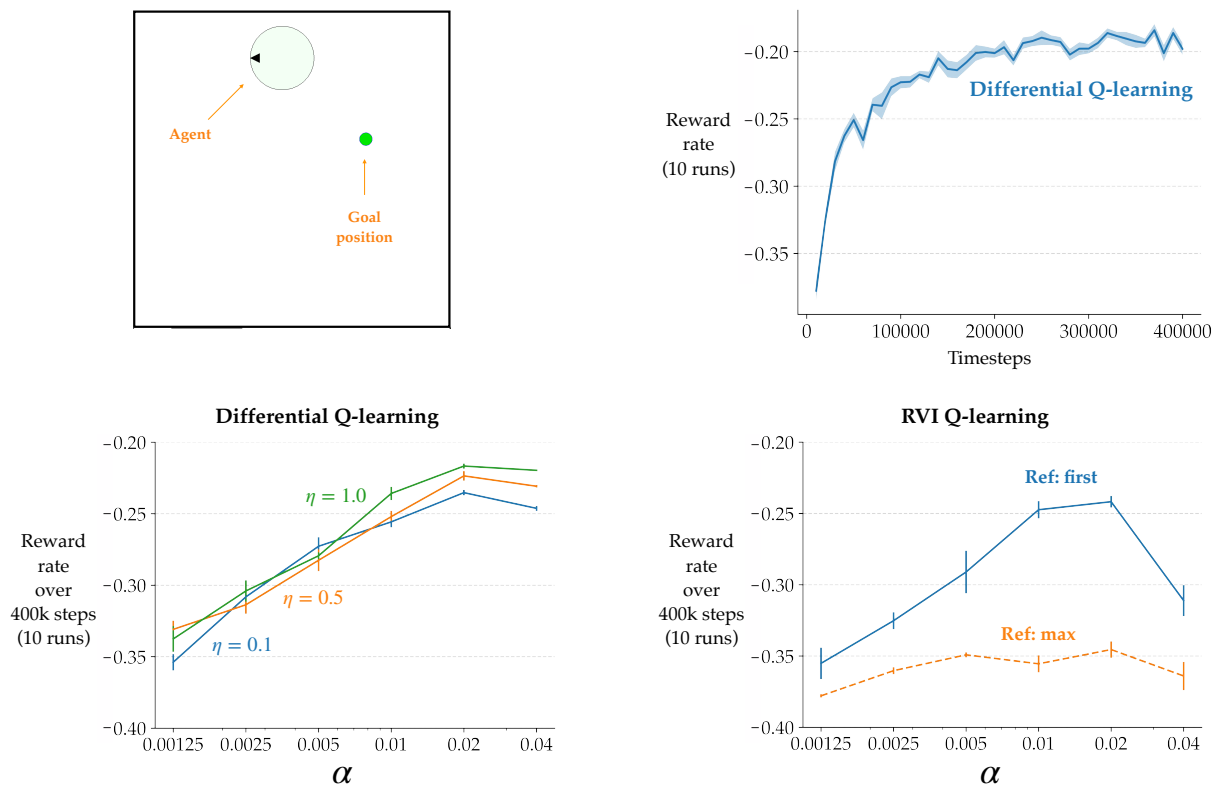


Figure E.1: A learning curve and parameter studies for the linear function approximation versions of Differential Q-learning and RVI Q-learning on the PuckWorld problem. The shaded region and the error bars in the plots represent one standard error. *Top-left:* A still of the PuckWorld domain showing the agent and the goal position. *Top-right:* A typical learning curve started roughly at a reward rate of -0.4 and rose to about -0.19. *Bottom-left:* Parameter studies showing the performance of Differential Q-learning in terms of average reward rate was not very sensitive to the choice of parameters. *Bottom-right:* Parameter studies showing the performance of RVI Q-learning is relatively good when the reference function is the first observed feature vector, and relatively worse for the other reference function for a broad range of step sizes.

We applied the linear function approximation versions of Differential Q-learning and RVI Q-learning on this problem. RVI Q-learning used the two reference functions discussed earlier in this section: (1) the action-value estimate corresponding to the *first* feature vector the agent observed when moving left, and (2) the maximum of the set of estimated action values corresponding to the feature vectors when they are observed.⁶ Both algorithms used tile coding (Sutton & Barto 2018: Section 9.5.4) with 16 symmetric tilings of $2 \times 2 \times 2 \times 2 \times 2$ tiles each. The weight vectors of both algorithms and the reward-rate estimate of Differential Q-learning was initialized to zero. The step-size parameter α was varied for both algorithms in the range $\{0.00125, 0.0025, 0.005, 0.01, 0.02, 0.04\}$. The parameter η for Differential Q-learning was varied in $\{0.1, 0.5, 1.0\}$. Each instance of parameters was applied for 10 runs of 400,000 timesteps each. Both algorithms used an ϵ -greedy policy with $\epsilon = 0.1$ and no annealing.

The top-right panel of Figure E.1 shows a typical learning curve on an instance of this problem where the goal positions is changed after every 100 timesteps. Using an ϵ -greedy policy with $\epsilon = 0.1$, the agent learns a policy that obtains a reward rate (computed over the last 10k steps) of about -0.19. The reward rate of a random policy is around -0.4. This learning curve corresponds to Differential Q-learning with $\alpha = 0.02, \eta = 1.0$. The learned policy was visualized and seen to be good everywhere except at the very edges of the two-dimensional space, which was probably an artifact of tile-coding.

We evaluated the performance of the agents across all the different parameter settings in terms of the average reward rate across the entire 400k timesteps of interaction. This is an indicator of the rate of learning. We observed that Differential Q-learning’s rate of learning was quite robust to the parameter η . Its two parameters did not interact strongly; the best value of α was independent of the choice of η . Moreover, the best performance for different η values was roughly the same. These observations were similar to those in the tabular case (see Section 3 in the main text).

RVI Q-learning also performed well on this problem for one choice of the reference function—the value estimate corresponding to the first feature vector the agent observes (with the ‘left’ action). The performance corresponding to the other reference function tested—tracking the maximum value of the observed feature vectors online—was not as good. This might be because unlike the tabular setting, updating the weights corresponding to one feature vector also modifies the estimate for other feature vectors, making the max hard to track. The best rate of learning corresponding to the better-performing reference function was slightly lower than that with Differential Q-learning.

We now move on to the second experiment in the linear function approximation setting. In the Catcher problem, the agent needs to catch as many falling fruits as possible. A still of the environment is shown in the top-left panel of Figure E.2. ‘Fruits’ fall vertically down from a uniformly-random horizontal position starting at the top of the frame. The agent can control the position of a ‘crate’ at the bottom of the frame using two actions — left and right — which move the crate in that direction by a small amount. If the fruit falls on/in the crate, the agent gets a reward of +40; if the fruit falls anywhere outside at the bottom of the frame, the agent gets -40. The next fruit starts falling only after the previous fruit has reached the bottom of the frame. A fruit takes roughly 40 timesteps to reach the the bottom starting from the top. Hence, the maximal reward rate on this problem is 1. At every timestep, the agent observes a four-dimensional feature vector of the crate’s horizontal position, the crate’s horizontal velocity, the fruit’s horizontal position, and the fruit’s vertical position. The positions and velocity are scaled to lie roughly in $[0, 1]$ and in $[-1, 1]$ respectively.

All the experimental details are the same as for PuckWorld, the only difference being that both algorithms used tile coding with 8 symmetric tilings of $4 \times 4 \times 4 \times 4$ tiles each.

The top-right panel of Figure E.2 shows a typical learning curve on this problem. Using an ϵ -greedy policy with $\epsilon = 0.1$, the agent learns a policy that obtains a reward rate of about 0.85, which is close to the optimal reward rate of 1. The reward rate of a random policy is around -0.3. The learning curve shown corresponds to Differential Q-learning with $\alpha = 0.02, \eta = 1.0$.

Again, we evaluated the rate of the learning of the agents across different parameter settings. We again observed that Differential Q-learning’s rate of learning did not vary much across a broad range of its parameter values. It was also especially robust to η . The linear function approximation version of RVI Q-learning also performed well for one choice of the reference function, not as much with the other. The learned policies corresponding to good parameter values for both algorithms successfully catch almost every fruit.

For RVI Q-learning, using the estimate of the first observed feature vector as a reference value worked better in Catcher than in PuckWorld. This might be because the agent might be observing feature vectors similar to the first one quite frequently, given that the crate has to move across the whole one-dimensional horizontal plane under any optimal policy. On the other hand, the agent moves in a relatively larger two-dimensional space in PuckWorld. In a finite number of agent-environmental

⁶the max was tracked online without storing all the previously observed feature vectors




Figure E.2: A learning curve and parameter studies for the linear function approximation versions of Differential Q-learning and RVI Q-learning on the Catcher problem. The shaded region and the error bars in the plots represent one standard error. *Top-left:* A still of the Catcher domain showing a falling fruit and the crate that the agent controls along the horizontal dimension at the bottom. *Top-right:* A typical learning curve started close to a reward rate of 0 and rose to about 0.9. *Bottom-left:* Parameter studies showing the performance of Differential Q-learning in terms of average reward rate was not very sensitive to the choice of parameters. *Bottom-right:* Parameter studies showing the performance of RVI Q-learning is relatively good when the reference function is the first observed feature vector, and relatively worse for the other reference function for a broad range of step sizes.

interactions, the agent might not visit its starting location that frequently. This suggests that the choice of the reference feature vector can affect the performance of RVI Q-learning differently in different problems. Additionally, in both cases, the other reference function did not result in good performance; this was probably because tracking the maximum action value in the function approximation setting is a poor approximation to the maximum action value across all state–action pairs.

The two experiments showed that the simple extension of the tabular Differential Q-learning to the linear function approximation setting can work rather well in terms of the final performance as well as robustness to different parameter values. The extension of the notion of reference functions to the linear function approximation setting is not as straightforward.