

---

# On the Statistical Benefits of Temporal Difference Learning

---

David Cheikhi<sup>1</sup> Daniel Russo<sup>1</sup>

## Abstract

Given a dataset on actions and resulting long-term rewards, a direct estimation approach fits value functions that minimize prediction error on the training data. Temporal difference learning (TD) methods instead fit value functions by minimizing the degree of temporal inconsistency between estimates made at successive time-steps. Focusing on finite state Markov chains, we provide a crisp asymptotic theory of the statistical advantages of this approach. First, we show that an intuitive *inverse trajectory pooling coefficient* completely characterizes the percent reduction in mean-squared error of value estimates. Depending on problem structure, the reduction could be enormous or nonexistent. Next, we prove that there can be dramatic improvements in estimates of the difference in value-to-go for two states: TD’s errors are bounded in terms of a novel measure — the problem’s *trajectory crossing time* — which can be much smaller than the problem’s time horizon.

## 1. Introduction

Temporal difference learning is a distinctive approach to estimation in long-term optimization problems. Its importance to reinforcement learning is hard to overstate. In their seminal book, Sutton & Barto (2018) write: *If one had to identify one idea as central and novel to reinforcement learning, it would undoubtedly be temporal difference learning.*

Competing with temporal difference (TD) learning is a straightforward direct-estimation approach. There, one proceeds by collecting data on past decisions and the cumulative long-term ‘reward’ that followed them. If actions were chosen with some experimental randomness, then regression of long-term rewards on the draw of actions would – with enough data – correctly identify actions’ causal impacts.

---

<sup>\*</sup>Equal contribution <sup>1</sup>Columbia University. Correspondence to: David Cheikhi <d.cheikhi@columbia.edu>.

The direct approach has two significant drawbacks. The first is delay: actions can only be evaluated after their full long-term effects realize. The second is variance: long-term outcomes can be extremely noisy and individual actions often have a small impact on them.

TD aims to alleviate these challenges by leveraging data on intermediate outcomes – those observed after the decision but before final outcomes are realized. The availability of such data is increasingly common. Robots collect regular sensor measurements, recommendation systems log sequential user interactions, and digital devices can track patient’s health metrics across time. Sutton (1988) observed that successive predictions, updated as information is gathered, should be *temporally consistent*. He proposed to fit maximally consistent value estimates by iteratively minimizing ‘temporal difference errors’. TD has since become an intellectual pillar of the reinforcement learning literature. It is used in most successful applications.

**Our Contributions.** We aim to provide crisp insight into the statistical benefits of TD. This paper focuses on the simplest possible setting, where training data consists of a batch of trajectories sampled independently from a finite Markov reward process. We compare the asymptotic scaling of mean squared error under TD and direct value estimation. Two main findings, different in nature, emerge:

1. The relative benefits of TD are determined by a natural ‘inverse trajectory pooling coefficient.’ TD uses value-to-go at intermediate states as a surrogate (Athey et al., 2019; Prentice, 1989). This is beneficial exactly when trajectories that originate with distinct states/actions tend to reach common intermediate states. We present simple examples illustrating when the benefits of TD are enormous and when they vanish entirely.
2. TD is especially beneficial for advantage estimation; That is, when estimating the difference in value-to-go from one state/action versus another. TD estimates of the value-to-go at different states are not independent and the coupling of errors leads to this improvement. While the mean squared error of direct advantage estimation generally scales with the length of the problem horizon, we show that TD’s errors are bounded by a smaller *trajectory crossing time*. This novel notion of

effective horizon can be small even in some problems with unbounded mixing time.

**On the Markov assumption and state representation.**

Our focus on Markov models is standard in the academic literature on reinforcement learning. This is, in a certain technical sense, an innocuous assumption. One could always use the entire sequence of observations so far as a Markov state (Puterman, 2014). But practical algorithms need to use appropriate compression of the history. The choice of representation has a subtle interplay with the benefit of TD. Indeed, we comment in section 9 that the benefits of TD can vanish when the state representation is too rich and trajectory pooling vanishes.

**2. Related works**

TD has been a central idea in RL since it was first proposed. It is deceptively simple, has intriguing connections to neuroscience (Schultz et al., 1997), and seems to be routed in dynamic programming theory. In the 1990s, researchers gathered both limited convergence guarantees (Dayan & Sejnowski, 1994; Jaakkola et al., 1993) and examples of divergence (Baird, 1995). Tsitsiklis & Van Roy (1997) offered a clarifying theory of when TD converges, and characterized the TD fixed point it reaches. Maei et al. (2009); Sutton et al. (2009) proposed methods to reach the TD fixed point in off-policy settings or when nonlinear function approximation is used.

While illuminating, this theory does not clarify why TD should be preferred over direct value function estimation (dubbed ‘Monte-Carlo’ or ‘MC’). In fact, the main guarantee is convergence to an approximate value function whose mean-squared error is *larger* than the one MC reaches. Folklore, intuition, and experiments suggest TD often converges to its limit at a faster rate.

The literature has emphasized the distinction between online and batch TD algorithms. The convergence speed of batch TD methods, like LSTD (Bradtke & Barto, 1996), is a statistical question. With purely online algorithms, each observation is used to make a single stochastic gradient type update and then immediately discarded, so issues of memory, compute, and data efficiency are conflated. The deep RL literature has adopted experience replay (Andrychowicz et al., 2017; Mnih et al., 2015; Schaul et al., 2015; Wang et al., 2016), which blurs the line between batch and online implementations by recording observations in a dataset and resampling them many times.

We give a complete and intuitive characterization of the efficiency benefits of TD in the simplest possible setting: a batch variant applied without function approximation in a finite state Markov Reward Process. Here it is straightforward to show TD is more efficient than MC, but more subtle to

understand when the efficiency gains are large. Grunewalder et al. (2007) and Grünwälder & Obermayer (2011) make progress in this direction. They prove that LSTD is at least as statistically efficient as MC, without quantifying the improvement. They also display cases where the two procedures have the same performance. Textbooks by Sutton & Barto (2018) and Szepesvári (2010) give illuminating examples, but no theory.

A number of papers bound the data requirements of TD, mainly in settings with function approximation. See for example (Mannor et al., 2004), (Lu, 2005), (Lazaric et al., 2010), (Pires & Szepesvári, 2012), (Tagorti & Scherrer, 2015), (Bhandari et al., 2018), (Pananjady & Wainwright, 2020), (Khamaru et al., 2020), (Chen et al., 2020), or (Farias et al., 2022). These show certain problem instances have low data requirements, but do not clarify when enforcing temporal consistency in value estimates produces large efficiency gains. To the best of our knowledge, our insights about advantage estimation are new (See Sec 8).

**3. Problem Formulation**

We first describe the problem of value function estimation in Markov reward processes (MRPs). We then observe that after appropriate relabeling of the state variables, this can also represent the problem of evaluating the long-term impact of actions. Most mathematical results are stated for MRPs, but the alternative interpretation enriches the results.

**3.1. Value function estimation**

A *trajectory* in a terminating Markov reward process is a Markovian sequence  $\tau = (S_0, R_1, S_1, R_2, S_2, \dots, S_{T-1}, R_T, \emptyset)$ , consisting of a sequence of states  $(S_t)_{t \in [T]} \subset \mathcal{S}$ , rewards  $(R_t)_{t \in [T]} \subset \mathbb{R}$ , and termination time  $T$ . The termination time is the first time at which  $S_T = \emptyset$ , where  $\emptyset$  is thought of as a special terminal state. Assume the distribution of  $R_t$  is independent of past rewards given the current state  $S_t$ .

The law of a Markov reward process (MRP) is specified by the tuple  $\mathcal{M} = (\mathcal{S} \cup \{\emptyset\}, P, R, d)$  consisting of a state space, transition kernel, reward distribution, and initial state distribution. Here  $P$  is a transition matrix over the augmented state space  $\mathcal{S} \cup \emptyset$ , specifying a probability  $P(s' | s)$  of transitioning from  $s$  to  $s'$ . We assume terminal state is absorbing and reachable. That is,  $P(\emptyset | \emptyset) = 1$  and for every state  $s$  there is some  $t$  such that the  $t$  step transition  $P^t(\emptyset | s)$  is strictly positive. The object  $R$  specifies the draw of rewards conditioned on a state transition as  $R(dr | s, s') = \mathbb{P}(R_t = dr | S_t = s, S_{t+1} = s')$ . Throughout we use the notation  $r(s, s')$  for the mean of  $R(\cdot | s, s')$ . We assume  $R(\emptyset, \emptyset) = 0$ . The initial state distribution  $d$  is a probability distribution over  $\mathcal{S}$  from which  $S_0$  is drawn.

The value function

$$\begin{aligned} V(s) &= \mathbb{E} \left[ \sum_{t=1}^{\infty} R_t \mid S_0 = s \right] \\ &= \mathbb{E} \left[ \sum_{t=1}^T r(S_t, S_{t+1}) \mid S_0 = s \right] \end{aligned}$$

specifies the expected future reward earned prior to termination. It is immediate from our formulation that  $V(\emptyset) = 0$ . Our formulation is the Markov reward process analogue of stochastic shortest path problems (Bertsekas & Tsitsiklis, 1991). Discounted problems are a special case where there is a constant probability of termination  $P(\emptyset \mid s) = 1 - \gamma$  for each non-terminal state  $s \in \mathcal{S}$ . In that case, the horizon  $T$  follows a geometric distribution with mean  $1/(1 - \gamma)$ .

A related quantity measures the value-to-go differences between states,

$$\mathbb{A}(s, s') = V(s) - V(s').$$

We call this the *advantage* of  $s$  over  $s'$ , since, as revealed in the next subsection, it is closely related to the advantage function in RL (Baird III, 1993).

We consider the problem of estimating the value-to-go at initial states. We compare methods that produce estimates  $\hat{V}$  based on  $n$  independent trajectories

$$\mathcal{D} = \left( \tau_i = (S_0^{(i)}, R_1^{(i)}, S_1^{(i)}, \dots, S_{T^{(i)}-1}^{(i)}, R_{T^{(i)}}^{(i)}, \emptyset) \right)_{i=1, \dots, n}$$

by their mean squared error  $\mathbb{E} \left[ \left( V(s) - \hat{V}(s) \right)^2 \right]$  or  $\mathbb{E} \left[ \left( \mathbb{A}(s, s') - \hat{\mathbb{A}}(s, s') \right)^2 \right]$ , where the expectation is over the randomness in  $\mathcal{D}$ . We assume that all states have a non-zero probability of being visited in the dataset.

### 3.2. Heterogenous treatment effect estimation

This section demonstrates theory developed in this paper directly applies to settings seemingly more general than the one described in 3.1. Specifically, we explain how evaluating the long-term impact of a chosen decision in a specific context falls in the scope of our setting with a simple appropriate relabeling of variables.

Here we consider a special case of our formulation where the continuing state space  $\mathcal{S} = \mathcal{S}_0 \cup \mathcal{S}_I$  is partitioned into a set of initial states  $\mathcal{S}_0$  and a set of intermediate states  $\mathcal{S}_I$ . With probability 1,  $S_0 \in \mathcal{S}_0$ ,  $S_T = \emptyset$ , and at intermediate times  $t \in \{1, \dots, T-1\}$ ,  $S_t \in \mathcal{S}_I$ .

We give the initial states a special interpretation. We think of them as consisting of an initial context  $X_0$  and a decision  $A_0$  and write  $S_0 = (X_0, A_0)$ . Using the more familiar

notation  $Q(X_0, A_0) \equiv V(S_0)$ , we have

$$Q(x, a) = \mathbb{E} \left[ \sum_{t=1}^T R_t \mid X_0 = x, A_0 = a \right].$$

The initial distribution  $d_0$  is determined by an initial context distribution  $\mathbb{P}(X_0 = x)$  and a *logging policy*  $\mathbb{P}(A_0 = x \mid X_0 = x)$  which determines the frequency with which actions are observed in the data.

Of particular interest in this setting is the advantage

$$\mathbb{A}((x, a), (x, a')) = Q(x, a) - Q(x, a'),$$

which measures the performance difference between actions  $a$  and  $a'$  in context  $x$ . The term ‘advantage’ is common in RL, but in causal inference one might call this ‘heterogenous treatment effect’ estimation.

Since policy gradient methods typically involve computing the expectation of weighted advantages, we expect that insights developed in this paper could apply to these methods.

## 4. Algorithms

### 4.1. Direct approach: First visit monte-Carlo (MC)

For any candidate value function, we can evaluate its accuracy by comparing the future value it predicts from a given state visited in the data and the actual cumulative reward observed in the remainder of that trajectory. This suggests a natural *direct* value estimation approach: over a candidate class of value functions, pick one that minimizes mean squared prediction error on the dataset. This method is called *Monte Carlo* in the RL literature.

To formally describe the algorithm, define the random time  $T(s) = \min\{t : S_t = s \vee S_t = \emptyset\}$  to be the first hitting time of state  $s$  if it is reached, or  $T$  otherwise. Let  $I(s) = \{i \in [n] : s \in \tau_i\}$  to be set of trajectories that visit state  $s$ . Form a dataset

$$\mathcal{D}^{\text{MC}} = \bigcup_{s \in \mathcal{S}} \bigcup_{i \in I(s)} \left\{ \left( s, \sum_{t=T^{(i)}(s)+1}^{T^{(i)}} R_t^{(i)} \right) \right\}.$$

that records pairs of states and the cumulative rewards earned following the first visit to the state. Given a parameterized class of value functions  $\{V_\theta : \theta \in \Theta\}$ , a direct value estimation approach is to solve the least squares problem

$$\min_{\theta \in \Theta} \sum_{(s,v) \in \mathcal{D}^{\text{MC}}} (V_\theta(s) - v)^2$$

We focus on tabular representations, where the space of parameterized value functions spans all possible functions. In that case

$$\hat{V}_{\text{MC}}(s) = \mathbb{E}_{(S,V) \sim \mathcal{D}^{\text{MC}}} [V \mid S = s],$$

where the notation  $(S, V) \sim \mathcal{D}^{\text{MC}}$  means that state/value pairs are sampled uniformly at random from the dataset. The value estimate at state  $s$  is simply the average reward earned after visits to state  $s$  in the dataset.

What we described is called *first visit* Monte-carlo in the literature. It is an unbiased estimator because it only includes the first time a state was visited during the an episode. We focus on this version for analytical simplicity, but many of our key examples focus on cases where initial states are never revisited (See e.g. Subsection 3.2) and it coincides with an “every visit” Monte Carlo estimate.

## 4.2. Indirect approach: TD learning

Temporal difference learning uses a reformatted dataset consisting of tuples of reward realizations and state transitions:

$$\mathcal{D}^{\text{TD}} = \{(S_t^{(i)}, R_{t+1}^{(i)}, S_{t+1}^{(i)})_{t=1, \dots, T^{(i)}, i=1, \dots, n}\}.$$

Define the temporal difference error between two candidate value functions  $V, V'$  to be the average gap in Bellman’s equation:

$$\ell_{\text{TD}}(V, V') = \mathbb{E}_{(S, R, S') \sim \mathcal{D}^{\text{TD}}} [V(S) - (R + V'(S'))],$$

where the notation  $(S, R, S') \sim \mathcal{D}^{\text{TD}}$  means that tuples are sample uniformly at random from the dataset. Given a parameterized class of value functions,  $V_{\Theta} = \{V_{\theta} : \theta \in \Theta\}$ , batch TD algorithms iteratively generate parameters  $(\theta_1, \theta_2, \dots)$  by solving  $\min_{\theta_{i+1}} \ell(V_{\theta_{i+1}}, V_{\theta_i})$ . (Online TD algorithms, combined with experience replay, make SGD updates instead.) For linear function approximation (Tsitsiklis & Van Roy, 1996), or neural networks in the ‘Neural tangent kernel’ regime (Cai et al., 2019), value functions are known to converge to a fixed point

$$\hat{V}_{\text{TD}} = \arg \min_{V \in V_{\Theta}} \ell(V, \hat{V}_{\text{TD}}),$$

which, in a sense, maximizes feasible temporal consistency.

Again, we focus on tabular representations, where the space of parameterized value functions spans all possible functions. In that case, TD solves the *empirical Bellman equation*

$$\hat{V}_{\text{TD}}(s) = \mathbb{E}_{(S, R, S') \sim \mathcal{D}^{\text{TD}}} [R + \hat{V}_{\text{TD}}(S') \mid S = s]$$

## 5. Intuition: surrogacy and intermediate outcomes

A lot of the intuition regarding TD can be gained through the simple example in Figure 1. Imagine our goal is to select the website design among 100 alternatives that leads to the largest sale rate (of some product). Users arrive and are randomly assigned to one of the 100 pages. They either click

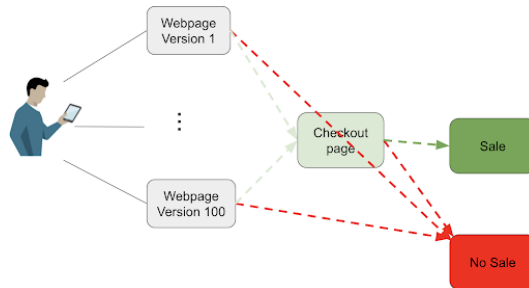


Figure 1. Modeling a user’s behavior

to purchase and proceed to the checkout page or navigate away from the site without purchasing. Among those who click, only a small fraction complete the sale.

Assume, for simplicity, that we have no access to personal information that distinguishes users from one another. (There is only one possible  $x$  in Section 3.2.) There are 100 possible *initial states*, corresponding to the webpage version, and the user is equally likely to start at each of those. A sale and a no-sale immediately precede termination. We call the checkout page an *intermediate state*. The sale state is associated with a reward of 1 and all others are associated with a reward of 0.

What we called the Monte-Carlo estimate of the value function would directly estimate the value of an impression of each webpage to be the fraction who purchased among that cohort of users.

Due to the directed nature of state transitions, TD estimation can be thought of in two steps. We first estimate  $\hat{V}(\text{checkout})$  to be the fraction who purchased among users who visit the checkout page. We then estimate the value of an impression of webpage  $i$  by

$$\hat{V}^{\text{TD}}(\text{webpage } i) = \text{CTR}(i) \times \hat{V}(\text{checkout}),$$

the empirical click-through rate among those shown webpage  $i$  times the estimated sale rate on the checkout page.

Monte-carlo estimation is unbiased, but it may be difficult to reliably estimate the efficacy of each webpage. If only a small fraction click initially, and among those who do only a small fraction convert to a sale, one would need to show each webpage to an enormous number of users. With TD, we *pool data* from across users who were shown any of the 100 webpages when estimating  $\hat{V}(\text{checkout})$ , greatly reducing variance. In this example, there is a lot of data pooling because trajectories that begin at distinct states quickly converge to the intermediate state. In fact, our theory reveals that certain intuitive measures of trajectory pooling exactly determine degree of statistical benefit TD provides.

Another view of TD is that it uses the intermediate click/no-

click outcome as a surrogate or proxy-metric (Prentice, 1989). Recognizing this, TD’s potential downsides become as transparent as its benefits. If the conversion probability among users who visit the checkout page depends strongly on which page design they saw, the Markov property does not hold and TD is biased. We discuss this example again in the conclusion, mentioning how the risks and benefits interplay with the choice of state representation.

## 6. Empirical illustration

### 6.1. The benefits of TD

To illustrate how much TD can improve over MC, we explore an example: we consider a layered MRP as describe in Figure 2. A layered MRP with width  $W$  and horizon  $T$  has  $W \times (T - 1)$  states split in  $T$  layers. States in layer  $t$  can only transition to states in layer  $t + 1$  and states in the last layer  $T - 1$  always transition to the terminal state.


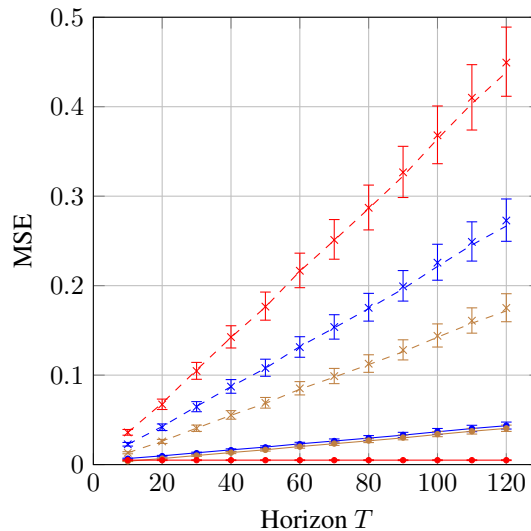


Figure 2. Layered MRP with width  $W$  and horizon  $T$ . Transitions are chosen randomly and rewards are uniform on  $[r(s, s') - 1; r(s, s') + 1]$  where  $r(s, s')$  is chosen uniformly between -1 and 1.


We consider a Layered MRP with width  $W = 5$ . We focus on state  $s_1^{(1)}$  and  $s_1^{(2)}$  and study the accuracy of the estimates of their value as we vary the horizon  $T$  of the MRP. We also study the accuracy of the estimate of the advantage  $\mathbb{A}(s_1^{(1)}, s_1^{(2)}) = V(s_1^{(1)}) - V(s_1^{(2)})$ . Figure 3 displays the Mean Square Error (MSE) of the TD and MC estimates for these quantities when the dataset contains  $n = 2000$  independent trajectories. MSE calculations involve 10000 Monte-Carlo replications. Alongside the observed MSE, we plot projected MSE based on the central limit theorem from Proposition B.4 and Proposition B.5. There is almost perfect alignment between asymptotic and finite sample results.

We highlight three main take-aways from this example:

1. **TD can vastly outperform MC.** For the chosen states  $s_1^{(1)}$  and  $s_1^{(2)}$ , the MSE is about 5 times smaller when using TD instead of MC for a Layered MRP with width  $W = 5$  and horizon  $T = 120$ .



(a) Full Y-Axis



(b) Truncated Y-Axis




Figure 3. MSE of different MC and TD estimates on Layered MRP with  $W = 5$  and varying horizon  $T$

2. **TD benefits are enhanced for advantage estimation.** In this example, TD performs up to 100 times better than MC for the advantage estimation. This example also shows that the MSE of the TD estimate of the ATE is smaller than the MSE of individual estimates of the value of  $s_1^{(1)}$  and  $s_1^{(2)}$  when the horizon  $T$  is larger than 20. On the other hand, the MSE of the MC estimate of the advantage is larger than the individual MSE of the

estimates of the value.

3. **TD effectively truncates the horizon.** While the MSE of the MC estimate of the advantage scales linearly with the horizon  $T$ , the MSE of the TD estimate is constant with respect to the horizon  $T$ . This is all the more striking that the variance of the total reward along a trajectory scales linearly with the horizon.

## 6.2. Dependence on the MRP structure

We have seen in Section 6.1 that TD vastly outperforms MC in the case of Layered MRP. However, different MRP structures lead to different level of improvement of TD over MC. To illustrate this, we introduce a new class of




Figure 4. MRP with meeting horizon  $H$

MRPs described in Figure 4. Each of the  $k$  initial states  $s_1^{(1)}, \dots, s_1^{(k)}$  lead to disjoint trajectories for the first  $H - 1$  steps before reaching a common state on the  $H$ th step. We are interested in seeing how TD and MC perform when the crossing time  $H$  varies. Figure 5 displays the ratio of the MSE of the TD and MC estimates for the values of  $s_1^{(1)}, s_1^{(2)}$  and for the advantage as the crossing time  $H$  varies. The MSE used to compute the ratios have been computed using  $n = 200$  independent trajectories and 1000 Monte-Carlo replications. These ratios are plotted alongside the asymptotic ratio from Theorem 7.2. As the crossing time  $H$  gets closer to the horizon  $T$ , the advantage of TD over MC vanishes until  $H = T$ , when TD and MC produce the exact same estimates. To convey intuition, we first focus on the two extreme cases:

- In the case where  $H = 2$ , depicted in Figure 6a, all initial states directly transition to the same state. This mimics the webpage example in Figure 1. In this example, apart from the first reward, trajectories do not depend of the initial state. TD pools trajectories across actions which allows to highly reduce the variance of the estimate of the value at  $s_2$ . This low variance estimate is then used as a surrogate to estimate the value-to-go at initial states. On the other hand, MC does not leverage the structure of the MRP and produces




Figure 5. Ratio of variance between TD and MC as a function of the meeting horizon  $H$  for  $T = 20$

an independent estimate for each initial state, using only trajectories starting at a given state to evaluate this state. In this case, TD will significantly improve over MC.

- At the other extreme, consider  $H = T$ . Then, no two initial states can lead to a common state before the terminal state, as shown in Figure 6b. There is no opportunity to pool trajectories across actions so TD strictly reduces to MC in this case.




Figure 6. 6a is an instance on which TD leverages pooling to improve considerably over MC. 6b is an instance on which TD and MC output the same estimate.

## 6.3. Organization of the results

In Section 7, we characterize the ratio in variance between TD and MC estimates for value estimation depending on the MRP structure. This characterization enables an intuitive understanding of which MRP structures lead to a large

improvement of TD and, conversely, for which MRP structures TD and MC perform similarly. In Section 8, we show that the TD estimate of advantages scales with an effective horizon that can be much smaller than the horizon.

## 7. Value estimation and the pooling coefficient

Recall that  $T(s)$  is the first hitting time of state  $s$  if it is reached, or  $T$  otherwise. The variables

$$N(s') = \sum_{t=0}^T \mathbb{1}(S_t = s'), \quad N(s \rightarrow s') = \sum_{t=T(s)}^T \mathbb{1}(S_t = s'),$$

respectively measure the total number of visits to state  $s'$  and the number of visits to  $s'$  which occur after a visit to state  $s$ .

Define the coupling coefficient between  $s$  and  $s'$  by

$$C(s, s') = \frac{\mathbb{E}[N(s \rightarrow s')]}{\mathbb{E}[N(s')]},$$

with  $C(s, s') = 0$  if  $\mathbb{E}[N(s')] = 0$ . Implicitly, it is understood that  $S_0$  is drawn from the MRP's initial distribution  $d$ . Among all trajectories which reach state  $s'$ , the coupling coefficient measures the fraction which first pass through state  $s$ . If the coupling coefficient is large, it means  $s$  and  $s'$  are strongly coupled.

The inverse *trajectory pooling coefficient* measures the average coupling coefficient  $C(s, s')$  over a distribution of possible successor states  $s'$ . The right distribution over which to average turns out to be  $\mu_s(\cdot)$ , identified in the definition below. That distribution weighs highly states that are likely to be visited following a visit to  $s$  (high  $\mathbb{E}[N(s') \mid S_0 = s]$ ) and contribute heavily to estimator variance (measured through the one-step variance  $\text{Var}(R_{t+1} + V(S_{t+1}) \mid S_t = s')$ ).

**Definition 7.1** (Inverse trajectory pooling coefficient). For any state  $s \in \mathcal{S}$  define

$$C(s) = \mathbb{E}_{s' \sim \mu_s} [C(s, s')],$$

where  $\mu_s(\cdot)$  a probability distribution over states defined by  $\mu_s(s') \propto \mathbb{E}[N(s') \mid S_0 = s] \times \text{Var}[R_t + V(S_{t+1}) \mid S_t = s']$ .

Again, the inverse trajectory pooling coefficient is small when there is a lot of trajectory pooling. The next theorem compares the asymptotic mean squared error of the value estimated under TD and a direct approach. The asymptotic ratio of their mean squared errors is equal to the inverse trajectory pooling coefficient.

**Theorem 7.2.** For any  $s \in \mathcal{S}$ ,

$$\lim_{n \rightarrow \infty} \frac{\mathbb{E} \left[ \left( \hat{V}_{TD}(s) - V(s) \right)^2 \right]}{\mathbb{E} \left[ \left( \hat{V}_{MC}(s) - V(s) \right)^2 \right]} = C(s).$$

Let us interpret this result. Recall that TD updates value prediction at state  $s$  using value predictions at successor states. The theorem shows this is helpful precisely when there is a lot of trajectory pooling, resulting in a small inverse trajectory pooling coefficient. When this holds, and the dataset  $\mathcal{D}$  is large, there will be many trajectories which reach an important possible successor  $s'$  of  $s$ , but never cross  $s$  first. TD leverages these trajectories to learn about  $s'$  and then properly incorporates that knowledge to better evaluate  $s$ . Direct estimation approaches only use sub-trajectories originating at  $s$  to evaluate  $s$  and forego the trajectory pooling benefit. We already developed this intuition by discussing the simple example in Figure 1. The theorem confirms that this interpretation of TD's advantages is exactly the right one.

Figure 6a describes an instance with extreme trajectory pooling. Trajectories that start in distinct states tend to immediately reach common successors, so TD understands value-to-go from successors quite well. Figure 6b is a case with no trajectory pooling at all (i.e.  $C(s) = 1$ ).

## 8. Horizon truncation in advantage estimation

Section 6 previewed two of the paper's key insights: TD's benefits are enhanced for advantage estimation and, in that setting, it effectively truncates the problem's time horizon. Theory in this section formalizes these insights.

**The MSE of direct advantage estimates scales with the horizon.** The variance of the total reward along a trajectory typically scales with the horizon. Therefore, it would not be surprising that the mean squared error of the estimate of the advantage also scales with the horizon. We show that it is indeed the case for MC by stating a lower bound on the mean squared error.

**Proposition 8.1.** For  $s, s' \in \mathcal{S}$  such that  $\mathbb{P}[s \in \tau \wedge s' \in \tau] = 0$ ,

$$\begin{aligned} \lim_{n \rightarrow \infty} n \cdot \mathbb{E} \left[ \left( \hat{\mathbb{A}}_{MC}(s, s') - \mathbb{A}(s, s') \right)^2 \right] \\ \geq \sigma_{\min}^2 \left( \frac{\mathbb{E}[T \mid S_0 = s]}{\mathbb{P}[s \in \tau]} + \frac{\mathbb{E}[T \mid S_0 = s']}{\mathbb{P}[s' \in \tau]} \right). \end{aligned}$$

where  $\sigma_{\min}^2 = \min_{s \in \mathcal{S}} \text{Var}[R_t + V(S_{t+1}) \mid S_t = s]$

The condition  $\mathbb{P}[s \in \tau \wedge s' \in \tau] = 0$  guarantees that no trajectory can visit both  $s$  and  $s'$ . It ensures that the MC estimate of the value at  $s$  and  $s'$  are independent. It is verified when considering the heterogeneous treatment effect, described in Section 3.2, where a single action is chosen for every trajectory. The scaling in the inverse probability of visiting  $s$  and  $s'$  appears because  $n\mathbb{P}[s \in \tau]$  and  $n\mathbb{P}[s' \in \tau]$  are asymptotically the number of trajectories available for the Monte-Carlo estimation.

**The MSE of TD’s advantage estimates scales with a smaller trajectory crossing time.** Rather than scale with problem’s time horizon, the mean squared error of TD’s advantage estimate is bounded by a smaller notion of the problem’s effective horizon. To formally capture this phenomenon, we introduce the *trajectory crossing time*  $H(s, s')$  for two states  $s$  and  $s'$  to be the expected time for trajectories starting at  $s$  and  $s'$  to cross under the most optimistic coupling. Two trajectories always cross once both have terminated, as in that case both have visited the terminal state  $\emptyset$ , but they could cross much sooner. Intuition for this definition is provided below.

**Definition 8.2.** The set of distributions  $\Psi(s, s')$  is the set of all joint distributions over trajectories  $(\tau, \tilde{\tau})$  such that the marginal distributions of  $\tau$  and  $\tilde{\tau}$  are those of trajectories starting at  $s$  and  $s'$ , respectively.

**Definition 8.3** (Trajectory crossing time). The trajectory crossing time of two states  $s$  and  $s'$  is the expected time for trajectories originating from  $s$  and  $s'$  to cross under the best coupling that preserves the trajectories’ marginal distributions:

$$H(s, s') = \min_{\psi \in \Psi(s, s')} \mathbb{E}_{(\tau, \tilde{\tau}) \sim \psi} [\inf\{t | \mathcal{C}_t(S, S') \neq \emptyset\}]$$

where  $\mathcal{C}_t(S, S') = \{S_0, \dots, S_t\} \cap \{S'_1, \dots, S'_t\}$  is the set of states visited by both trajectories at time  $t$ .

The following theorem establishes that the mean squared error of the TD estimate of the advantage scales with the trajectory crossing time instead of the full horizon:

**Theorem 8.4.** For  $s, s' \in \mathcal{S}$ ,

$$\begin{aligned} \lim_{n \rightarrow \infty} n \cdot \mathbb{E} \left[ \left( \hat{\mathbb{A}}_{\text{TD}}(s, s') - \mathbb{A}(s, s') \right)^2 \right] \\ \leq 2 \left( \frac{\sigma_{\max}^2}{\min(\mathbb{P}[s \in \tau], \mathbb{P}[s' \in \tau])} \right) \cdot H(s, s') \end{aligned}$$

with  $\sigma_{\max}^2 = \max_{s \in \mathcal{S}} \text{Var}[R_t + V(S_{t+1}) | S_t = s]$ .

Figure 3, in Section 6, provides an empirical illustration of this result.

This result can actually understate the benefits the TD. Any trajectory pooling that happens before the trajectories cross further helps reducing the variance, but is not reflected in the upper bound of Theorem 8.4. In particular, we show in Appendix B.5 that trajectory pooling ensures that TD estimates are at least as accurate as MC estimates, even when the crossing time matches the full horizon. We also discuss the tightness of these bounds in Appendix B.6.

**Comparison with a coupling time.** Trajectories are said to cross if one of the trajectories reaches a state already visited by the other one, potentially at an earlier time. It




Figure 7. An example with no coupling but rapid crossing.

is different from the coupling time where trajectories have to reach a common state simultaneously. In particular, the crossing time is always smaller than the coupling time. The MRP defined in Figure 7 illustrates this. Trajectories starting at states  $s_1^{(1)}$  and  $s_1^{(2)}$  only couple when reaching the terminal state, after  $m + 1$  timesteps. However, they cross in two timesteps, that is  $H(s_1^{(1)}, s_1^{(2)}) = 2$ .

**Intuition for the result.** Let us give intuition for Theorem 8.4. Under the MRP structure, two trajectories reaching a common state have the same future expected reward. Hence, when estimating the difference in expected total rewards along two trajectories, one starting at  $s$ , the other starting at  $s'$ , it is only useful to estimate them up until the state where trajectories cross. By computing estimates at every state, TD leverages this property: two trajectories reaching a common state (the crossing state) use the same estimate (the value at the crossing state) to update predictions along both trajectories. Since the same estimate is used, its value cancels out when computing the difference in values at  $s$  and  $s'$ . Whether the value at the crossing state is accurately estimated doesn’t affect the estimation of the advantage.

## 9. Open questions

There is a subtle interplay between the choice of state representation and the benefits of imposing temporal consistency in value estimates. Consider again the problem in Figure 1. In that case, we chose to represent the ‘checkout page’ as a state, implying that the purchase probability at the checkout page does not depend on the initial webpage shown to the user. This makes a strong surrogacy assumption, which TD leverages to greatly improve data efficiency. An alternative representation of the state in the second period is of the form  $s = (\text{website version } i, \text{checkout})$ , retaining information about how the user navigated to the checkout. In this case, there is no trajectory pooling and our theory indicates that TD behaves as MC would. By using a very rich representation, which recalls much of the past, the benefits of TD disappear. Clearly, we want representations that are accurate, to avoid severe approximation errors. But we have shown that representations that are forgetful of aspects of the past offer enormous benefits; this lets value-to-go from intermediate states serve as surrogate outcomes. How



should one balance this tradeoff while learning representations from data? This question closely relates to the one of understanding temporal consistency with function approximation.

## References

- Andrychowicz, M., Wolski, F., Ray, A., Schneider, J., Fong, R., Welinder, P., McGrew, B., Tobin, J., Pieter Abbeel, O., and Zaremba, W. Hindsight experience replay. *Advances in neural information processing systems*, 30, 2017.
- Athey, S., Chetty, R., Imbens, G. W., and Kang, H. The surrogate index: Combining short-term proxies to estimate long-term treatment effects more rapidly and precisely. Technical report, National Bureau of Economic Research, 2019.
- Baird, L. Residual algorithms: Reinforcement learning with function approximation. In *Machine Learning Proceedings 1995*, pp. 30–37. Elsevier, 1995.
- Baird III, L. C. Advantage updating. Technical report, WRIGHT LAB WRIGHT-PATTERSON AFB OH, 1993.
- Bertsekas, D. P. and Tsitsiklis, J. N. An analysis of stochastic shortest path problems. *Mathematics of Operations Research*, 16(3):580–595, 1991.
- Bhandari, J., Russo, D., and Singal, R. A finite time analysis of temporal difference learning with linear function approximation. In *Conference on learning theory*, pp. 1691–1692. PMLR, 2018.
- Bradtke, S. J. and Barto, A. G. Linear least-squares algorithms for temporal difference learning. *Machine learning*, 22(1):33–57, 1996.
- Cai, Q., Yang, Z., Lee, J. D., and Wang, Z. Neural temporal-difference learning converges to global optima. *Advances in Neural Information Processing Systems*, 32, 2019.
- Chen, S., Devraj, A., Busic, A., and Meyn, S. Explicit mean-square error bounds for monte-carlo and linear stochastic approximation. In *International Conference on Artificial Intelligence and Statistics*, pp. 4173–4183. PMLR, 2020.
- Dayan, P. and Sejnowski, T. J. Td ( $\lambda$ ) converges with probability 1. *Machine Learning*, 14(3):295–301, 1994.
- Farias, V. F., Li, A. A., Peng, T., and Zheng, A. T. Markovian interference in experiments. *arXiv preprint arXiv:2206.02371*, 2022.
- Grünewälder, S. and Obermayer, K. The optimal unbiased value estimator and its relation to lstd, td and mc. *Machine learning*, 83(3):289–330, 2011.
- Grunewalder, S., Hochreiter, S., and Obermayer, K. Optimality of lstd and its relation to mc. In *2007 International Joint Conference on Neural Networks*, pp. 338–343. IEEE, 2007.
- Jaakkola, T., Jordan, M., and Singh, S. Convergence of stochastic iterative dynamic programming algorithms. *Advances in neural information processing systems*, 6, 1993.
- Khamaru, K., Pananjady, A., Ruan, F., Wainwright, M. J., and Jordan, M. I. Is temporal difference learning optimal? an instance-dependent analysis. *arXiv preprint arXiv:2003.07337*, 2020.
- Lazaric, A., Ghavamzadeh, M., and Munos, R. Finite-sample analysis of lstd. In *ICML-27th International Conference on Machine Learning*, pp. 615–622, 2010.
- Lu, F. Error bounds in reinforcement learning policy evaluation. In *Conference of the Canadian Society for Computational Studies of Intelligence*, pp. 438–449. Springer, 2005.
- Maei, H., Szepesvari, C., Bhatnagar, S., Precup, D., Silver, D., and Sutton, R. S. Convergent temporal-difference learning with arbitrary smooth function approximation. *Advances in neural information processing systems*, 22, 2009.
- Mannor, S., Simester, D., Sun, P., and Tsitsiklis, J. N. Bias and variance in value function estimation. In *Proceedings of the twenty-first international conference on Machine learning*, pp. 72, 2004.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al. Human-level control through deep reinforcement learning. *nature*, 518(7540): 529–533, 2015.
- Pananjady, A. and Wainwright, M. J. Instance-dependent  $\ell_\infty$ -bounds for policy evaluation in tabular reinforcement learning. *IEEE Transactions on Information Theory*, 67(1):566–585, 2020.
- Pires, B. A. and Szepesvári, C. Statistical linear estimation with penalized estimators: an application to reinforcement learning. *arXiv preprint arXiv:1206.6444*, 2012.
- Prentice, R. L. Surrogate endpoints in clinical trials: definition and operational criteria. *Statistics in Medicine*, 8(4): 431–440, 1989.
- Puterman, M. L. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014.

- Schaul, T., Quan, J., Antonoglou, I., and Silver, D. Prioritized experience replay. *arXiv preprint arXiv:1511.05952*, 2015.
- Schultz, W., Dayan, P., and Montague, P. R. A neural substrate of prediction and reward. *Science*, 275(5306): 1593–1599, 1997.
- Sutton, R. S. Learning to predict by the methods of temporal differences. *Machine learning*, 3(1):9–44, 1988.
- Sutton, R. S. and Barto, A. G. *Reinforcement learning: An introduction*. MIT press, 2018.
- Sutton, R. S., Maei, H. R., Precup, D., Bhatnagar, S., Silver, D., Szepesvári, C., and Wiewiora, E. Fast gradient-descent methods for temporal-difference learning with linear function approximation. In *Proceedings of the 26th annual international conference on machine learning*, pp. 993–1000, 2009.
- Szepesvári, C. Algorithms for reinforcement learning. *Synthesis lectures on artificial intelligence and machine learning*, 4(1):1–103, 2010.
- Tagorti, M. and Scherrer, B. On the rate of convergence and error bounds for lstd ( $\lambda$ ). In *International Conference on Machine Learning*, pp. 1521–1529. PMLR, 2015.
- Tsitsiklis, J. and Van Roy, B. Analysis of temporal-difference learning with function approximation. *Advances in neural information processing systems*, 9, 1996.
- Tsitsiklis, J. N. and Van Roy, B. An analysis of temporal-difference learning with function approximation. *IEEE TRANSACTIONS ON AUTOMATIC CONTROL*, 42(5), 1997.
- Wang, Z., Bapst, V., Heess, N., Mnih, V., Munos, R., Kavukcuoglu, K., and de Freitas, N. Sample efficient actor-critic with experience replay. *arXiv preprint arXiv:1611.01224*, 2016.

### A. Additional experiments

Figure 3 and 5 aims at displaying the difference in precision of MC estimates and TD estimates on minimal working examples, to underline what is driving the difference. In this section, we complement the experiments in the main body to illustrate that the theory holds in more complex settings. Specifically, we show empirical results on an augmented version of the Layered MRP introduced in Figure 2, where we introduce cycles. Specifically, a cyclic Layered MRP with width  $W$ , horizon  $T$  and backward arc probability  $p$  has  $T - 1$  layers of  $W$  states, each of which has a non-zero probability to each state in the next layer and with probability  $p$ , a transition to a state in a previous or the current layer. The structure is illustrated in Figure 8.




Figure 8. Layered MRP with width  $W$  and horizon  $T$  and backwards transitions. There is at most one backward transition per node. The probability of adding a backward transition is fixed  $p$  and the destination of a backward transition is chosen uniformly at random. Transitions are chosen randomly and rewards are uniform on  $[r(s, s') - 1; r(s, s') + 1]$  where  $r(s, s')$  is chosen uniformly between -1 and 1.

We first show that the Central Limit Theorems governing the MC and TD estimates still hold empirically when cycles are introduced. We used 10000 Monte-Carlo replications, each of which used 2000 samples to compute the value-to-go.

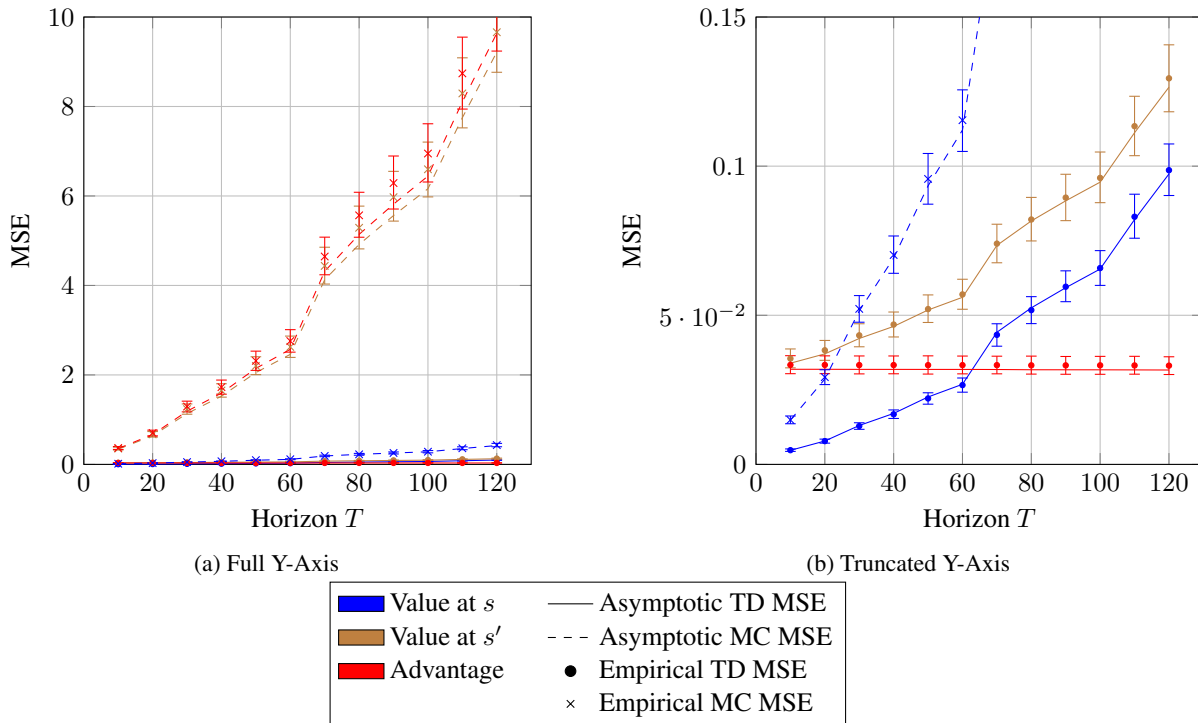


Figure 9. MSE of different MC and TD estimates on Cyclic Layered MRP with  $W = 5$ ,  $p = 0.1$  and varying horizon  $T$

Next, we show how accurate the limiting approximation is in the finite sample regime. To do so, we compute the empirical

MSE for the MC and TD estimates for a range of sample sizes and plot alongside the normal approximation. The MRP used for these experiment is a Layered MRP with cycles with parameter  $W = 5, T = 120, p = 0.1$ . We used 10000 Monte-Carlo replications.

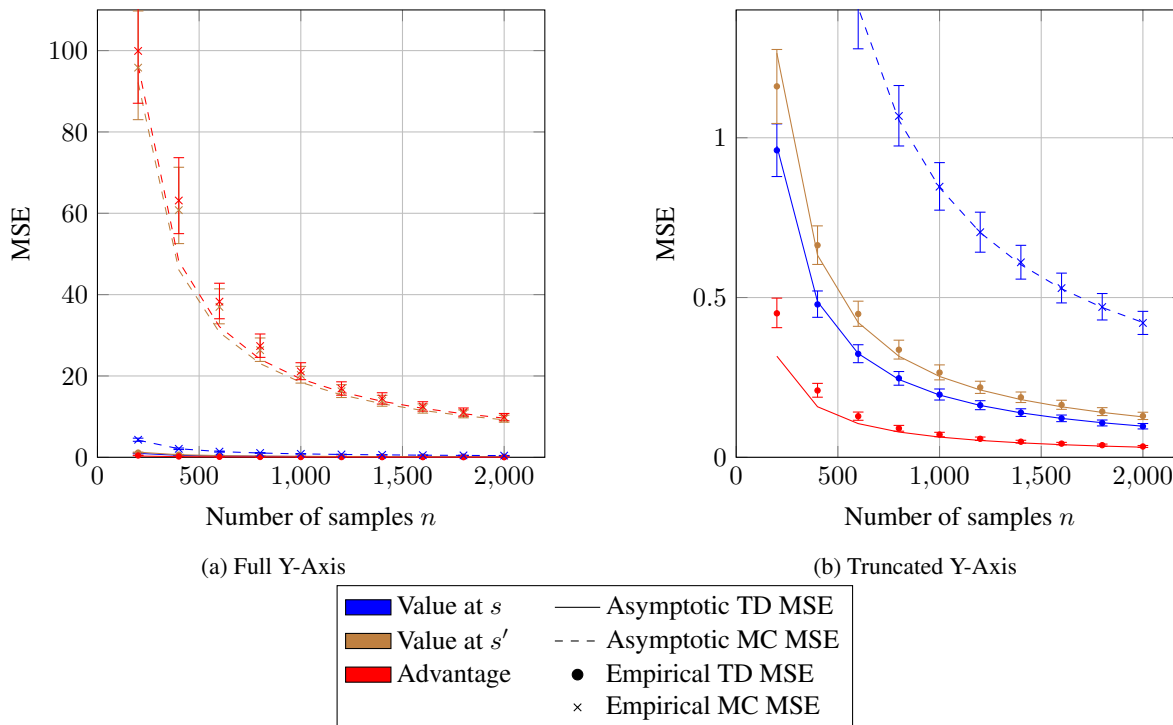


Figure 10. MSE of MC and TD estimators when varying the number of samples, compared with the normal approximation

An alternate view on the impact of the number of samples is the one of regret: when trying to choose what state has the highest value function between two states  $s$  and  $s'$ , how often would an estimator indicate the wrong decision. That is, the regret of an estimator  $\hat{V}$  is  $\mathbb{E} \left[ \mathbb{1} \left( \hat{V}(s) > \hat{V}(s') \right) \right]$ , assuming  $V(s') > V(s)$ . We plot the empirical regret when using MC estimates versus TD estimates, as a function of the number of samples, along with the normal approximation derived from the Central Limit Theorems.

## B. Proofs.

We state and prove results in the context of weighted value function which is a linear combination of the value function evaluated at individual states.

**Definition B.1** (Weighted value function). For a weighting over states  $\pi$ , the extended value function is defined as

$$J(\pi) = \sum_{s \in \mathcal{S}} \pi(s) V(s)$$

By setting  $\pi$  to be a mass point at a single state, we recover the value function at this state. When interpreting initial states as actions (as in Section 3.2), we recover randomized policy when using a distribution over actions as the weighting. In this case, the weighted value function is the expected value when playing according to the randomized policy. Note that our definition of weighted value function allows for any weighting of states, including negative weights. This will be useful for analyzing advantages  $V(s) - V(s')$  by setting  $\pi(s) = 1$  and  $\pi(s') = -1$ .

We also extend our definition of expected number of visits to weightings over initial states.

**Definition B.2** (Weighted expected number of visits). For a weighting over states  $\pi$ , we write  $\eta_\pi(s)$  the weighted number




Figure 11. Regret for TD and MC estimates on Cyclic Layered MRP with  $W = 5, p = 0.1, T = 120$  as the number of samples used vary.

of visits to  $s$ :

$$\eta_\pi(s) = \sum_{s' \in \mathcal{S}} \pi(s') \mathbb{E}[N(s) | S_0 = s']$$

Similarly to the weighted value function,  $\pi$  is not enforced to be a distribution over state, allowing even for negative values. In the case where  $\pi$  is a distribution over state, we recover the probabilistic interpretation:  $\eta_\pi(s)$  is the expected number of visits to state  $s$  when the initial distribution is  $\pi$ .

**Definition B.3** (One-step variance).

$$\sigma_V^2(s) = \text{Var}[R_t + V(S_{t+1}) | S_t = s].$$

We extend trajectories into infinite horizon trajectories that stay in the terminating state and stop collecting rewards once the terminating state is reached:  $S_t = \emptyset$  and  $R_{t+1} = 0$  for all  $t \geq T$ . Equivalently, we define the transition  $P(\emptyset | \emptyset) = 1$  and the reward  $R(\emptyset, \emptyset) = 0$  a.s.

We start by stating and proving Central Limit Theorems (CLT) for the convergence of both TD and MC estimates. We then use these two results as building blocks to prove the main theorems.

### B.1. Central Limit Theorems

**Proposition B.4** (Central Limit Theorem for MC). For  $s \in \mathcal{S}$ ,

$$\sqrt{n}(\hat{V}_{\text{MC}}(s) - V(s)) \Rightarrow \mathcal{N}\left(0, \frac{1}{\mathbb{P}[s \in \mathcal{T}]} \sum_{s' \in \mathcal{S}} \mathbb{E}[N(s') | S_0 = s] \sigma_V^2(s')\right)$$

*Proof.* We recall that, for tabular representation, the MC estimator takes the form

$$\begin{aligned} \hat{V}_{\text{MC}}(s) &= \mathbb{E}_{(S, V) \sim \mathcal{D}^{\text{MC}}}[V | S = s] \\ &= \frac{1}{|I(s)|} \sum_{i \in I(s)} \sum_{t=T^{(i)}(s)}^{T^{(i)}} R_{t+1}^{(i)}, \end{aligned}$$

where  $I(s)$  is the set of trajectories that visit state  $s$  and  $T^{(i)}(s)$  is the first visit to state  $s$  in trajectory  $i$ . Since we consider first-visit MC, each trajectory appears at most once in the summation. We start by rewriting the error  $\hat{V}_{\text{MC}}(s) - V(s)$  as the product of a scaling factor and the average of i.i.d. random variables:

$$\begin{aligned}\hat{V}_{\text{MC}}(s) - V(s) &= \frac{1}{|I(s)|} \sum_{i \in I(s)} \left( \sum_{t=T^{(i)}(s)}^{T^{(i)}} R_{t+1}^{(i)} - V(s) \right) \\ &= \frac{n}{|I(s)|} \cdot \frac{1}{n} \sum_{i=1}^n \left( \mathbb{1}(s \in \tau^{(i)}) \left( \sum_{t=T^{(i)}(s)}^{T^{(i)}} R_{t+1}^{(i)} - V(s) \right) \right)\end{aligned}$$

Recall that if  $s$  is not visited in trajectory  $i$ ,  $T^{(i)}(s)$  is defined to be  $T^{(i)}(s) = T^{(i)}$ .

- We start by proving a Central Limit Theorem on

$$\frac{1}{n} \sum_{i=1}^n \left( \mathbb{1}(s \in \tau^{(i)}) \sum_{t=T^{(i)}(s)}^{T^{(i)}} R_{t+1}^{(i)} \right).$$

The variables  $\left( \mathbb{1}(s \in \tau^{(i)}) \left( \sum_{t=T^{(i)}(s)}^{T^{(i)}} R_{t+1}^{(i)} - V(s) \right) \right)_{i=1, \dots, n}$  are  $n$  i.i.d., zero mean random variables. We now compute their variance.

$$\begin{aligned}\text{Var} \left[ \mathbb{1}(s \in \tau) \left( \sum_{t=T(s)}^T R_{t+1} - V(s) \right) \right] &= \mathbb{P}[s \in \tau] \text{Var} \left[ \mathbb{1}(s \in \tau) \left( \sum_{t=T(s)}^T R_{t+1} - V(s) \right) \mid s \in \tau \right] \\ &= \mathbb{P}[s \in \tau] \text{Var} \left[ \sum_{t=T(s)}^T R_{t+1} - V(s) \mid s \in \tau \right]\end{aligned}$$

Since the summation starts at the stopping time defined by the first visit to state  $s$ , the Strong Markov Property enables to re-index the summation in the following way:

$$\begin{aligned}\text{Var} \left[ \mathbb{1}(s \in \tau) \left( \sum_{t=T(s)}^T R_{t+1} - V(s) \right) \right] &= \mathbb{P}[s \in \tau] \text{Var} \left[ \sum_{t=0}^T R_{t+1} - V(s) \mid S_0 = s \right] \\ &= \mathbb{P}[s \in \tau] \text{Var} \left[ \sum_{t=0}^{\infty} R_{t+1} - V(s) \mid S_0 = s \right]\end{aligned}$$

where we allowed the sum to run to infinity since  $(R_{t+1})_t$  is a.s. stationary at 0 for  $t \geq T$ . Similarly, we use the fact that  $(V(S_t) - V(S_{t+1}))$  is a.s. stationary at 0 to write  $V(S_0) = \sum_{t=0}^{\infty} (V(S_t) - V(S_{t+1}))$ . Plugging in the previous expression gives:

$$\text{Var} \left[ \mathbb{1}(s \in \tau) \left( \sum_{t=T(s)}^T R_{t+1} - V(s) \right) \right] = \mathbb{P}[s \in \tau] \text{Var} \left[ \sum_{t=0}^{\infty} (R_{t+1} + V(S_{t+1}) - V(S_t)) \mid S_0 = s \right]$$

Notice that  $(R_{t+1} + V(S_{t+1}) - V(S_t))_t$  are martingale differences with respect to the filtration  $\mathcal{F}_t = \{S_0, \dots, S_t\}$ . Using that martingale differences are uncorrelated:

$$\text{Var} \left[ \mathbb{1}(s \in \tau) \left( \sum_{t=T(s)}^T R_{t+1} - V(s) \right) \right] = \mathbb{P}[s \in \tau] \sum_{t=0}^{\infty} \mathbb{E} [(V(S_{t+1}) + R_{t+1} - V(S_t))^2 | S_0 = s]$$

We then group the terms in the sum by the value of  $S_t$ :

$$\begin{aligned} \text{Var} \left[ \mathbb{1}(s \in \tau) \left( \sum_{t=T(s)}^T R_{t+1} - V(s) \right) \right] &= \\ &= \mathbb{P}[s \in \tau] \sum_{t=0}^{\infty} \sum_{s' \in \mathcal{S}} \mathbb{P}[S_t = s' | S_0 = s] \mathbb{E} [(V(S_{t+1}) + R_{t+1} - V(S_t))^2 | S_t = s'] \\ &= \mathbb{P}[s \in \tau] \sum_{t=0}^{\infty} \sum_{s' \in \mathcal{S}} \mathbb{P}[S_t = s' | S_0 = s] \sigma_V^2(s') \\ &= \mathbb{P}[s \in \tau] \sum_{s' \in \mathcal{S}} \sigma_V^2(s') \sum_{t=0}^{\infty} \mathbb{P}[S_t = s' | S_0 = s] \\ &= \mathbb{P}[s \in \tau] \sum_{s' \in \mathcal{S}} \mathbb{E}[N(s') | S_0 = s] \sigma_V^2(s') \end{aligned}$$

Using the Central Limit Theorem, we obtain the following convergence:

$$\sqrt{n} \frac{1}{n} \sum_{i=1}^n \left( \mathbb{1}(s \in \tau^{(i)}) \sum_{t=T^{(i)}(s)}^{T^{(i)}} R_{t+1}^{(i)} \right) \Rightarrow \mathcal{N} \left( 0, \mathbb{P}[s \in \tau] \sum_{s' \in \mathcal{S}} \mathbb{E}[N(s') | S_0 = s] \sigma_V^2(s') \right).$$

- The Strong Law of Large Number ensures:

$$\frac{n}{|I(s)|} \xrightarrow[n \rightarrow \infty]{} \frac{1}{\mathbb{P}[s \in \tau]} \quad \text{a.s..}$$

Finally, using Slutsky's Theorem, the product converges:

$$\frac{n}{|I(s)|} \cdot \frac{1}{n} \sum_{i=1}^n \left( \mathbb{1}(s \in \tau^{(i)}) \left( \sum_{t=T^{(i)}(s)}^{T^{(i)}} R_{t+1}^{(i)} - V(s) \right) \right) \Rightarrow \mathcal{N} \left( 0, \frac{1}{\mathbb{P}[s \in \tau]} \sum_{s' \in \mathcal{S}} \mathbb{E}[N(s') | S_0 = s] \sigma_V^2(s') \right).$$

□

**Proposition B.5** (Central Limit Theorem for TD). *For any weighting  $\pi$ ,*

$$\sqrt{n}(\hat{J}_{\text{TD}}(\pi) - J(\pi)) \Rightarrow \mathcal{N} \left( 0, \sum_{s' \in \mathcal{S}} \frac{\eta_{\pi}^2(s') \sigma_V^2(s')}{\mathbb{E}[N(s')]} \right).$$

**Corollary B.6.** *For  $s \in \mathcal{S}$*

$$\sqrt{n}(\hat{V}_{\text{TD}}(s) - V(s)) \Rightarrow \mathcal{N} \left( 0, \sum_{s' \in \mathcal{S}} \frac{\mathbb{E}[N(s') | S_0 = s]^2 \sigma_V^2(s')}{\mathbb{E}[N(s')]} \right).$$

## B.2. Proof of Proposition B.5

**Notations** The idea of the proof is to use the fixed point identities verified by  $V$  and  $\hat{V}_{\text{TD}}$  to obtain a recursive formula for the error  $\hat{V}_{\text{TD}} - V$ . We then analyze the two quantities that appear when iterating this recursive formula: one is a sum of empirical one-step error and we show the other is of second order.

We adopt a vectorial representation of the value function  $V = (V(s))_{s \in \mathcal{S}}$  such that  $J(\pi) = \langle V, \pi \rangle$ . Let  $T$  be the Bellman operator:

$$\mathcal{T}(f)(s) = \mathbb{E}_{S' \sim P(\cdot | s)} [R(s, S') + f(S')]$$

The value function is the unique fixed point of the Bellman operator:  $V = T(V)$  to verify  $V(\emptyset) = 0$ . We also write the transition operator as follow:

$$P(f)(s) = \mathbb{E}_{S' \sim P(\cdot | s)} [f(S')] = \langle P(\cdot | s), f \rangle.$$

As discussed in Section 4.2, by trying to minimize temporal differences, TD solves the *empirical Bellman equation*:

$$\hat{V}_{\text{TD}}(s) = \mathbb{E}_{(S, R, S') \sim \mathcal{D}^{\text{TD}}} \left[ R + \hat{V}^{\text{TD}}(S') \mid S = s \right].$$

Solving the *empirical Bellman equation* can be viewed as being a fixed point (with  $f(\emptyset) = 0$ ) of the empirical Bellman operator that we define as follow:

$$\hat{\mathcal{T}}(f)(s) = \mathbb{E}_{(S, R, S') \sim \mathcal{D}^{\text{TD}}} [R + f(S') \mid S = s].$$

Similarly, we note the empirical transition operator

$$\hat{P}(f)(s) = \mathbb{E}_{(S, R, S') \sim \mathcal{D}^{\text{TD}}} [f(S') \mid S = s].$$

Finally, we introduce an explicit notation of the empirical Bellman operator that will ease proofs. To do so, we first need to introduce  $B(s)$ , the set of all visits to state  $s$ :

$$B(s) = \{(i, t) \mid S_t^{(i)} = s\}$$

Using this notation, the empirical Bellman operator can be written as follow:

$$\hat{\mathcal{T}}(f)(s_0) = \frac{1}{|B(s_0)|} \sum_{(i, t) \in B(s_0)} \left( f(S_{t+1}^{(i)}) + R_{t+1}^{(i)} \right).$$

**Analysis** We start by expanding the error vector  $\hat{V}_{\text{TD}} - V$  into a sum of empirical one-step errors and a second term that we later show to be of second order.

### Lemma B.7.

$$\sqrt{n} \left( \hat{V}_{\text{TD}} - V \right) = \sqrt{n} \sum_{t=0}^{\infty} P^t (\hat{\mathcal{T}}V - V) + \sum_{t=0}^{\infty} P^t \left( \sqrt{n} (\hat{P} - P) (\hat{V}_{\text{TD}} - V) \right)$$

Note that, since  $P^t(S_0) = S_t$  is stationary at  $\emptyset$  almost surely, the quantities summed in Lemma B.7 are ultimately zero almost surely.

*Proof.* By expanding the difference  $\hat{V}_{\text{TD}} - V$  using the fixed point identities, we obtain a recursive identity:



$$\begin{aligned}
 \hat{V}_{\text{TD}} - V &= \hat{\mathcal{T}}\hat{V}_{\text{TD}} - \mathcal{T}V \\
 &= (\hat{\mathcal{T}} - \mathcal{T})\hat{V}_{\text{TD}} + \mathcal{T}\hat{V}_{\text{TD}} - \mathcal{T}V \\
 &= (\hat{P} - P)(\hat{V}_{\text{TD}} - V) + (\hat{\mathcal{T}} - \mathcal{T})V + P(\hat{V}_{\text{TD}} - V) \\
 &= (\hat{P} - P)(\hat{V}_{\text{TD}} - V) + (\hat{\mathcal{T}}V - V) + P(\hat{V}_{\text{TD}} - V).
 \end{aligned}$$

Iterating this identity and multiplying by  $\sqrt{n}$  on both sides gives

$$\sqrt{n}(\hat{V}_{\text{TD}} - V) = \sum_{t=0}^{\infty} P^t \left( \sqrt{n}(\hat{P} - P)(\hat{V}_{\text{TD}} - V) \right) + \sqrt{n} \sum_{t=0}^{\infty} P^t (\hat{\mathcal{T}}V - V). \quad (1)$$

□

We now state two lemmas that analyze the two terms that appear in equation (1). They are proved later.

**Lemma B.8.** For all  $s \in \mathcal{S}$ , as  $n \rightarrow \infty$ ,

$$\sum_{t=0}^{\infty} P^t \left( \sqrt{n}(\hat{P} - P)(\hat{V}_{\text{TD}} - V) \right) (s) \rightarrow 0 \quad \text{a.s.}$$

**Lemma B.9.** As  $n \rightarrow \infty$

$$\sqrt{n} \sum_{s \in \mathcal{S}} \pi(s) \sum_{t=0}^{\infty} P^t (\hat{\mathcal{T}}V - V)(s) \Rightarrow \mathcal{N} \left( 0, \sum_{s \in \mathcal{S}} \frac{\eta_{\pi}(s)^2}{\mathbb{E}[N(s)]} \sigma_V^2(s) \right)$$

Finally, combining Lemma B.8 and Lemma B.9 with Slutsky Lemma concludes the proof.

We now proceed to prove Lemma B.8 and Lemma B.9, starting with Lemma B.8

### B.2.1. PROOF OF LEMMA B.8

*Proof.* We start by showing that, after appropriate rescaling the error in transition estimates converge to a Gaussian distribution. We do not calculate the variance here since it is not needed.

**Lemma B.10.** For each  $s$ ,  $\sqrt{n}(\hat{P} - P)(\cdot|s)$  weakly converges to a normal distribution with mean zero.

*Proof.* We first express  $\hat{P}(s'|s) - P(s'|s)$  as the product of a scaling factor and the average of  $n$  i.i.d random vectors:

$$\begin{aligned}
 \left( \hat{P}(s'|s) - P(s'|s) \right)_{s' \in \mathcal{S}} &= \frac{1}{|B(s)|} \sum_{(i,t) \in B(s)} \left( \mathbb{1}(S_{t+1}^{(i)} = s') - P(s'|s) \right)_{s' \in \mathcal{S}} \\
 &= \frac{n}{|B(s)|} \cdot \frac{1}{n} \sum_{i=1}^n \left( \sum_{t=0}^{\infty} \mathbb{1}(S_t^{(i)} = s) \left( \mathbb{1}(S_{t+1}^{(i)} = s') - P(s'|s) \right) \right)_{s' \in \mathcal{S}}.
 \end{aligned}$$

We decompose this identity in two terms:

- Strong law of large numbers ensures that

$$\frac{n}{|B(s)|} \xrightarrow{n \rightarrow \infty} \frac{1}{\mathbb{E}[N(s)]} \quad \text{a.s..}$$

- Since trajectories are independent,

$$\sum_{t=0}^{\infty} \mathbb{1}(S_t^{(i)} = s) \left( \mathbb{1}(S_{t+1}^{(i)} = s') - P(s'|s) \right)_{s' \in \mathcal{S}}$$

are independent, identically distributed random variables with finite variance and mean 0. The Central Limit Theorem ensures that the rescaled average of these random variables

$$\sqrt{n} \cdot \frac{1}{n} \sum_{i=1}^n \sum_{t=0}^{\infty} \mathbb{1}(S_t^{(i)} = s) \left( \mathbb{1}(S_{t+1}^{(i)} = s') - P(s'|s) \right)_{s' \in \mathcal{S}}$$

weakly converges to a mean 0 normal distribution.

Finally, the product of these two quantities also converges to a mean 0 normal distribution. □

We terminate the proof of Lemma B.8 by combining Lemma B.10 with the fact that  $\hat{V}_{\text{TD}} - V$  converges almost surely to 0 by Slutsky's Theorem. □

### B.2.2. PROOF OF LEMMA B.9

*Proof.* We start by expanding the transition operator  $P^t$ ,

$$\sqrt{n} \sum_{t=0}^{\infty} P^t(\hat{\mathcal{T}}V - V)(s) = \sqrt{n} \sum_{t=0}^{\infty} \mathbb{E} \left[ \hat{\mathcal{T}}V(S_t) - V(S_t) \mid S_0 = s, \hat{\mathcal{T}} \right]$$

This term consists of a sum one-step differences of the form  $\sqrt{n} \left( \hat{\mathcal{T}}V(S_t^{(i)}) - V(S_t^{(i)}) \right)$ . Grouping  $S_t^{(i)}$  that are equal:

$$\begin{aligned} \sqrt{n} \sum_{t=0}^{\infty} P^t(\hat{\mathcal{T}}V - V)(s) &= \sqrt{n} \cdot \sum_{t=0}^{\infty} \mathbb{E} \left[ \sum_{s' \in \mathcal{S}} \mathbb{1}(S_t = s') \left( \hat{\mathcal{T}}V(s') - V(s') \right) \mid S_0 = s, \hat{\mathcal{T}} \right] \\ &= \sqrt{n} \cdot \mathbb{E} \left[ \sum_{s' \in \mathcal{S}} \sum_{t=0}^{\infty} \mathbb{1}(S_t = s') \left( \hat{\mathcal{T}}V(s') - V(s') \right) \mid S_0 = s, \hat{\mathcal{T}} \right] \\ &= \sqrt{n} \sum_{s' \in \mathcal{S}} \mathbb{E} \left[ \sum_{t=0}^{\infty} \mathbb{1}(S_t = s') \mid S_0 = s' \right] \left( \hat{\mathcal{T}}V(s') - V(s') \right) \\ &= \sqrt{n} \sum_{s' \in \mathcal{S}} \mathbb{E} [N(s') \mid S_0 = s] \left( \hat{\mathcal{T}}V(s') - V(s') \right). \end{aligned}$$

To get the result for a general weighting of states  $\pi$ , we take the dot product of the previous identity with  $\pi$ :

$$\sqrt{n} \sum_{s \in \mathcal{S}} \pi(s) \sum_{t=0}^{\infty} P^t(\hat{\mathcal{T}}V - V)(s) = \sqrt{n} \sum_{s \in \mathcal{S}} \eta_{\pi}(s) \left( \hat{\mathcal{T}}V(s) - V(s) \right). \quad (2)$$

We are left with analyzing a linear combination of terms of the form  $\hat{\mathcal{T}}V(s) - V(s)$ . For an individual  $s$ ,  $\hat{\mathcal{T}}V(s) - V(s)$  is an average of i.i.d. variables and its asymptotical behavior can be controlled using a Central Limit Theorem. However,

to find the limiting distribution of the linear combination, we need to prove a vectorial Central Limit Theorem for the random vector  $\left(\hat{\mathcal{T}}V(s) - V(s)\right)_{s \in \mathcal{S}}$ . In particular, we show that  $\left(\hat{\mathcal{T}}V(s) - V(s)\right)_{s \in \mathcal{S}}$  converge to independent normal distributions.

The following lemma is key to prove the independence of the limiting distribution. It states that one-step differences are uncorrelated:

**Lemma B.11.** For  $(s, t) \neq (s', t')$ :

$$\text{Cov} [\mathbb{1}(S_t = s) (V(S_t) - V(S_{t+1}) - R_{t+1}), \mathbb{1}(S_{t'} = s') (V(S_{t'}) - V(S_{t'+1}) - R_{t'+1})] = 0.$$

*Proof.* The result follows from the Markovian property:

$$\begin{aligned} & \text{Cov} [\mathbb{1}(S_t = s) (V(S_t) - V(S_{t+1}) - R_{t+1}), \mathbb{1}(S_{t'} = \tilde{s}) (V(S_{t'}) - V(S_{t'+1}) - R_{t'+1})] \\ &= \mathbb{E} [\text{Cov} [\mathbb{1}(S_t = s) (V(S_t) - V(S_{t+1}) - R_{t+1}), \mathbb{1}(S_{t'} = s') (V(S_{t'}) - V(S_{t'+1}) - R_{t'+1}) \mid S_t, S_{t+1}, S_{t'}]] \\ &= \mathbb{E} [\mathbb{1}(S_t = s) (V(S_t) - V(S_{t+1}) - R_{t+1}) \mathbb{1}(S_{t'} = s') (V(S_{t'}) - \mathbb{E}[V(S_{t'+1}) + R_{t'+1} \mid S_{t'}])] \\ &= 0. \end{aligned}$$

□

Given the previous result, we can now state and prove the joint Central Limit Theorem of empirical one-step differences:

**Lemma B.12.** As  $n \rightarrow \infty$ ,

$$\sqrt{n} \left( \left( \hat{\mathcal{T}}V(s) - V(s) \right) \right)_{s \in \mathcal{S}} \Rightarrow \mathcal{N}(0, \Sigma)$$

where  $\Sigma$  is a diagonal matrix with

$$\Sigma_{s,s} = \frac{1}{\mathbb{E}[N(s)]} \sigma_V^2(s).$$

*Proof.* We use a similar approach as in the proof of Lemma B.10: we express  $\hat{\mathcal{T}}V(s) - V(s)$  as the product of a scaling factor that converges almost surely and an average of i.i.d. random vectors that is controlled by the Central Limit Theorem:

$$\begin{aligned} \sqrt{n} \cdot \left( \hat{\mathcal{T}}V(s) - V(s) \right) &= \sqrt{n} \cdot \frac{1}{|B(s)|} \sum_{(i,t) \in B(s)} (V(S_{t+1}^{(i)}) + R_{t+1}^{(i)} - V(s)) \\ &= \sqrt{n} \cdot \frac{1}{|B(s)|} \sum_{i=1}^n \sum_{t=0}^{\infty} \mathbb{1}(S_t^{(i)} = s) (V(S_{t+1}^{(i)}) + R_{t+1}^{(i)} - V(S_t^{(i)})) \\ &= \frac{n}{|B(s)|} \left( \sqrt{n} \cdot \frac{1}{n} \sum_{i=1}^n \sum_{t=0}^{\infty} \mathbb{1}(S_t^{(i)} = s) (V(S_{t+1}^{(i)}) + R_{t+1}^{(i)} - V(S_t^{(i)})) \right) \end{aligned}$$

- $B(s)$  can be rewritten as the sum of i.i.d. random variables:

$$B(s) = \sum_{i=1}^n \left( \sum_{t=0}^{\infty} \mathbb{1}(S_t^{(i)} = s) \right).$$

The average  $\frac{|B(s)|}{n}$  converges almost surely to  $\mathbb{E}[N(s)]$  by the Strong Law of Large Numbers. Taking the inverse gives:

$$\frac{n}{|B(s)|} \rightarrow \frac{1}{\mathbb{E}[N(s)]} \quad \text{a.s.} \quad (3)$$

- The random vector

$$\sqrt{n} \cdot \frac{1}{n} \sum_{i=1}^n \left( \sum_{t=0}^{\infty} \mathbb{1}(S_t^{(i)} = s) (V(S_{t+1}^{(i)}) + R_{t+1}^{(i)} - V(S_t^{(i)})) \right)_{s \in \mathcal{S}}$$

is the re-scaled average of  $n$  i.i.d., mean 0, vectors of the form  $(\sum_{t=0}^{\infty} \mathbb{1}(S_t = s) (V(S_{t+1}) + R_{t+1} - V(S_t)))_{s \in \mathcal{S}}$ . The Central Limit Theorem ensures that this quantity converges to a normal distribution with mean 0. We now proceed to compute the variance of this distribution.

- Lemma B.11 ensures that entries of this vector are uncorrelated:
- The variance of a single entry is given by

$$\begin{aligned} \text{Var} \left[ \sum_{t=0}^{\infty} \mathbb{1}(S_t = s) (V(S_{t+1}) + R_{t+1} - V(S_t)) \right] &\stackrel{(a)}{=} \sum_{t=0}^{\infty} \text{Var} [\mathbb{1}(S_t = s) (V(S_{t+1}) + R_{t+1} - V(S_t))] \\ &= \sum_{t=0}^{\infty} \mathbb{E} [\mathbb{1}(S_t = s) (V(S_{t+1}) + R_{t+1} - V(S_t))^2] \\ &= \sum_{t=0}^{\infty} \mathbb{E} [\mathbb{1}(S_t = s) \mathbb{E} [(V(S_{t+1}) + R_{t+1} - V(S_t))^2 | S_t]] \\ &= \sum_{t=0}^{\infty} \mathbb{E} [\mathbb{1}(S_t = s) \sigma_V^2(S_t)] \\ &= \mathbb{E} [N(s)] \sigma_V^2(s) \end{aligned}$$

where (a) also follows from Lemma B.11.

From the Central Limit Theorem, we obtain:

$$\sqrt{n} \cdot \frac{1}{n} \sum_{i=1}^n \left( \sum_{t=0}^{\infty} \mathbb{1}(S_t^{(i)} = s) (V(S_{t+1}^{(i)}) + R_{t+1}^{(i)} - V(S_t^{(i)})) \right)_{s \in \mathcal{S}} \Rightarrow \mathcal{N} \left( 0, \text{Diag} \left( (\mathbb{E} [N(s)] \sigma_V^2(s))_{s \in \mathcal{S}} \right) \right) \quad (4)$$

Combining (3) and (4) gives the result. □

Finally, we just need to take the dot product of  $\sqrt{n} \cdot (\hat{\mathcal{T}}V(s) - V(s))_{s \in \mathcal{S}}$  with the weighted occupancy measure  $\eta_\pi$  to get the result:

$$\sqrt{n} \sum_{s \in \mathcal{S}} \pi(s) \sum_{t=0}^{\infty} P^t (\hat{\mathcal{T}}V - V)(s) \Rightarrow \mathcal{N} \left( 0, \sum_{s \in \mathcal{S}} \frac{\eta_\pi(s)^2}{\mathbb{E} [N(s)]} \sigma_V^2(s) \right)$$

□

### B.3. Proof of Theorem 7.2

The proof follows directly from Proposition B.4 and Proposition B.5.

From, Proposition B.4, we have

$$\lim_{n \rightarrow \infty} \sqrt{n} \cdot \mathbb{E} \left[ \left( \hat{V}_{\text{MC}}(s) - V(s) \right)^2 \right] = \frac{1}{\mathbb{P}[s \in \tau]} \sum_{s' \in \mathcal{S}} \mathbb{E} [N(s') | S_0 = s] \sigma_V^2(s').$$

Similarly, from Proposition B.5, we have

$$\lim_{n \rightarrow \infty} \sqrt{n} \cdot \mathbb{E} \left[ \left( \hat{V}_{\text{TD}}(s) - V(s) \right)^2 \right] = \sum_{s' \in \mathcal{S}} \frac{\mathbb{E} [N(s') | S_0 = s]^2 \sigma_V^2(s')}{\mathbb{E} [N(s')]}.$$

Taking the ratio of these two limits, we obtain:

$$\begin{aligned} \lim_{n \rightarrow \infty} \frac{\mathbb{E} \left[ \left( \hat{V}_{\text{TD}}(s) - V(s) \right)^2 \right]}{\mathbb{E} \left[ \left( \hat{V}_{\text{MC}}(s) - V(s) \right)^2 \right]} &= \frac{\sum_{s' \in \mathcal{S}} \mathbb{E} [N(s') | S_0 = s] \sigma_V^2(s')}{\sum_{s' \in \mathcal{S}} \mathbb{E} [N(s') | S_0 = s] \sigma_V^2(s')} \cdot \frac{\mathbb{E} [N(s') | S_0 = s] \mathbb{P}[s \in \tau]}{\mathbb{E} [N(s')]} \\ &= \mathbb{E}_{s' \sim \mu(s)} \left[ \frac{\mathbb{E} [N(s') | S_0 = s] \mathbb{P}[s \in \tau]}{\mathbb{E} [N(s')]} \right] \end{aligned}$$

where  $\mu(s)$  is defined in Definition 7.1.

Finally,

$$\mathbb{E} [N(s') | S_0 = s] \mathbb{P}[s \in \tau] = \mathbb{E} [N(s \rightarrow s')].$$

#### B.4. Proof of Theorem 8.4

We define  $\pi$  to be such that  $\pi(s) = 1$ ,  $\pi(s') = -1$  and  $\pi(\tilde{s}) = 0$  for all  $\tilde{s}$ . That way  $J(\pi) = V(s) - V(s') = \mathbb{A}(s, s')$ . This allows to use Lemma B.5:

$$\lim_{n \rightarrow \infty} n \cdot \mathbb{E} \left[ \left( \hat{J}_{\text{TD}}(\pi) - J(\pi) \right)^2 \right] = \sum_{\tilde{s} \in \mathcal{S}} \frac{(\mathbb{E} [N(\tilde{s}) | S_0 = s] - \mathbb{E} [N(\tilde{s}) | S_0 = s'])^2 \sigma_V^2(\tilde{s})}{\mathbb{E} [N(\tilde{s})]}$$

We decompose the sum into three terms that we bound separately:

$$\begin{aligned} \lim_{n \rightarrow \infty} n \cdot \mathbb{E} \left[ \left( \hat{J}_{\text{TD}}(\pi) - J(\pi) \right)^2 \right] &\leq \left( \max_{\tilde{s} \in \mathcal{S}} \left| \frac{\mathbb{E} [N(\tilde{s}) | S_0 = s] - \mathbb{E} [N(\tilde{s}) | S_0 = s']}{\mathbb{E} [N(\tilde{s})]} \right| \right) \\ &\quad \times \left( \max_{\tilde{s} \in \mathcal{S}} \sigma_V^2(\tilde{s}) \right) \\ &\quad \times \sum_{\tilde{s} \in \mathcal{S}} |\mathbb{E} [N(\tilde{s}) | S_0 = s] - \mathbb{E} [N(\tilde{s}) | S_0 = s']|. \end{aligned} \tag{5}$$

We start by proving that if trajectories where  $S_0$  is sampled from the initial distribution visit  $s$  frequently often, then the occupancy measure  $\mathbb{E} [N(\tilde{s})]$  and  $\mathbb{E} [N(\tilde{s}) | S_0 = s]$  cannot differ too much (and symmetrically for  $s'$ ).

If  $\mathbb{E} [N(\tilde{s}) | S_0 = s] \geq \mathbb{E} [N(\tilde{s}) | S_0 = s']$  then

$$\left| \frac{\mathbb{E} [N(\tilde{s}) | S_0 = s] - \mathbb{E} [N(\tilde{s}) | S_0 = s']}{\mathbb{E} [N(\tilde{s})]} \right| \leq \frac{\mathbb{E} [N(\tilde{s}) | S_0 = s]}{\mathbb{P}[s \in \tau] \mathbb{E} [N(\tilde{s}) | S_0 = s]} = \frac{1}{\mathbb{P}[s \in \tau]}.$$

When  $\mathbb{E} [N(\tilde{s}) | S_0 = s] < \mathbb{E} [N(\tilde{s}) | S_0 = s']$ , a symmetric argument ensures that

$$\left| \frac{\mathbb{E} [N(\tilde{s}) | S_0 = s] - \mathbb{E} [N(\tilde{s}) | S_0 = s']}{\mathbb{E} [N(\tilde{s})]} \right| \leq \frac{1}{\mathbb{P}[s' \in \tau]}.$$

Combining these two cases gives:

$$\max_{\tilde{s} \in \mathcal{S}} \left| \frac{\mathbb{E} [N(\tilde{s}) | S_0 = s] - \mathbb{E} [N(\tilde{s}) | S_0 = s']}{\mathbb{E} [N(\tilde{s})]} \right| \leq \frac{1}{\min(\mathbb{P}[s \in \tau], \mathbb{P}[s' \in \tau])}.$$

Plugging this result in Equation 5:

$$\lim_{n \rightarrow \infty} n \cdot \mathbb{E} \left[ \left( \hat{J}_{\text{TD}}(\pi) - J(\pi) \right)^2 \right] \leq \frac{\max_{\tilde{s} \in \mathcal{S}} \sigma_V^2(\tilde{s})}{\min(\mathbb{P}[s \in \tau], \mathbb{P}[s' \in \tau])} \cdot \sum_{\tilde{s} \in \mathcal{S}} |\mathbb{E}[N(\tilde{s}) | S_0 = s] - \mathbb{E}[N(\tilde{s}) | S_0 = s']|. \quad (6)$$

Finally, we show that the last sum scales as the crossing time. To give intuition on why this holds, note that we have the following identity:

$$\sum_{\tilde{s} \in \mathcal{S}} \mathbb{E}[N(\tilde{s}) | S_0 = s] = \mathbb{E}[T | S_0 = s].$$

$\mathbb{E}[N(\tilde{s}) | S_0 = s]$  is the expected number of visits to  $\tilde{s}$  when starting at  $s$ : summing over  $\tilde{s}$  is the expected total number of states visited when starting at  $s$  (not counting the terminating state  $\emptyset$ ) which is exactly the horizon. In the case of the comparison of trajectories,  $|\mathbb{E}[N(\tilde{s}) | S_0 = s] - \mathbb{E}[N(\tilde{s}) | S_0 = s']|$  is how many additional times  $\tilde{s}$  has been visited by one of the trajectories compared to the other, in expectation. Summing over  $\tilde{s}$  gives the expected total number of states that have been visited by only one of the trajectories, which is at most twice the number of states visited before both trajectories reach a common state.

**Lemma B.13.**

$$\sum_{\tilde{s} \in \mathcal{S}} |\mathbb{E}[N(\tilde{s}) | S_0 = s] - \mathbb{E}[N(\tilde{s}) | S_0 = s']| \leq 2H(s, s')$$

*Proof.* Let  $\tau, \tau' = (S_0, R_1, \dots, S_{T-1}, R_T, \emptyset), (S'_0, R'_1, \dots, S'_{T-1}, R'_T, \emptyset) \sim \psi$  where  $\psi$  is a joint distribution such that  $\tau$  and  $\tau'$  are marginally distributed like trajectories generated with  $S_0 = s$  and  $S'_0 = s'$ , respectively. We define the crossing time for  $\psi$  is defined as the first time a state has been visited by both trajectories:

$$H_\psi(s, s') = \inf\{t | \{S_0, \dots, S_t\} \cap \{\tilde{S}_0, \dots, \tilde{S}_t\} \neq \emptyset\}.$$

Since we always consider the crossing time of  $s$  and  $s'$ , we omit the  $s$  and  $s'$  dependency and simply write  $H_\psi$  for convenience in the rest of the proof. By definition, either  $S'_{H_\psi} \in \{S_1, \dots, S_{H_\psi}\}$  or  $S_{H_\psi} \in \{S'_1, \dots, S'_{H_\psi}\}$ . Without loss of generality, we assume  $S'_{H_\psi} \in \{s_1, \dots, s_{H_\psi}\}$  holds. Let  $N_\psi = \inf\{t | S_t = S'_{H_\psi}\}$ , that is  $S_{N_\psi} = S'_{H_\psi}$ .

We now construct a new trajectory that follows  $\tau'$  until the crossing state  $S_{N_\psi} = S'_{H_\psi}$  is reached and then follows the trajectory  $\tau$ :  $\hat{\tau} = (S'_0, R'_1, \dots, S'_{H_\psi} = S_{N_\psi}, R_{N+1}, S_{N+1}, \dots, S_{T-1}, R_T, \emptyset)$ . By Markov property,  $\hat{\tau}$  is identically distributed as  $\tau'$ .

We are interested in bounding the difference in occupancy measure:

$$\mathbb{E}[N(\tilde{s}) | S_0 = s] - \mathbb{E}[N(\tilde{s}) | S_0 = s'] = \mathbb{E} \left[ \sum_{t=0}^{\infty} \mathbb{1}(S_t = \tilde{s}) \right] - \mathbb{E} \left[ \sum_{t=0}^{\infty} \mathbb{1}(S'_t = \tilde{s}) \right].$$

Using that  $\hat{\tau}$  and  $\tau'$  are identically distributed, this expression can be rewritten as

$$\begin{aligned} \mathbb{E}[N(\tilde{s}) | S_0 = s] - \mathbb{E}[N(\tilde{s}) | S_0 = s'] &= \mathbb{E} \left[ \sum_{t=0}^{\infty} \mathbb{1}(S_t = \tilde{s}) \right] - \mathbb{E} \left[ \sum_{t=0}^{\infty} \mathbb{1}(\hat{S}_t = \tilde{s}) \right] \\ &= \mathbb{E} \left[ \sum_{t=0}^{N_\psi-1} \mathbb{1}(S_t = \tilde{s}) \right] - \mathbb{E} \left[ \sum_{t=0}^{H_\psi-1} \mathbb{1}(S'_t = \tilde{s}) \right] \end{aligned}$$

Taking absolute values and summing over  $\tilde{s}$  gives

$$\begin{aligned}
 \sum_{\tilde{s} \in \mathcal{S}} |\mathbb{E}[N(\tilde{s}) | S_0 = s] - \mathbb{E}[N(\tilde{s}) | S_0 = s']| &= \sum_{\tilde{s} \in \mathcal{S}} \mathbb{E} \left[ \left| \sum_{t=0}^{N_\psi-1} \mathbb{1}(S_t = \tilde{s}) - \sum_{t=0}^{H_\psi-1} \mathbb{1}(S'_t = \tilde{s}) \right| \right] \\
 &\leq \sum_{s' \in \mathcal{S}} \mathbb{E} \left[ \sum_{t=0}^{N_\psi-1} \mathbb{1}(S_t = \tilde{s}) \right] + \sum_{\tilde{s} \in \mathcal{S}} \mathbb{E} \left[ \sum_{t=0}^{H_\psi-1} \mathbb{1}(S_t = s') \right] \\
 &= \mathbb{E}[N_\psi] + \mathbb{E}[H_\psi] \\
 &\leq 2\mathbb{E}[H_\psi]
 \end{aligned}$$

Taking the infimum over all  $\psi \in \Psi(s, s')$  that conserves marginal distribution gives the result.  $\square$

Finally, plugging Lemma B.13 in Equation 6 leads to the result:

$$\lim_{n \rightarrow \infty} n \cdot \mathbb{E} \left[ \left( \hat{J}_{\text{TD}}(\pi) - J(\pi) \right)^2 \right] \leq 2 \frac{\max_{\tilde{s} \in \mathcal{S}} \sigma_V^2(\tilde{s})}{\min(\mathbb{P}[s \in \tau], \mathbb{P}[s' \in \tau])} \cdot H(s, s')$$

#### B.4.1. PROOF OF PROPOSITION 8.1

We now prove that the MC estimate of the advantage can scale as the full horizon, no matter how small the crossing time is. We focus on the advantage of  $s$  over  $s'$  when  $\mathbb{P}[s \in \tau \wedge s' \in \tau] = 0$ , that is no trajectory can visit both  $s$  and  $s'$ . No trajectories visiting both  $s$  and  $s'$  means that the MC value estimates at  $s$  and  $s'$  are independent since they rely on disjoint sets of trajectories. In terms of variance, this implies:

$$\text{Var} \left[ \hat{\mathbb{A}}_{\text{MC}}(s, s') \right] = \text{Var} \left[ \hat{V}_{\text{MC}}(s) - \hat{V}_{\text{MC}}(s') \right] = \text{Var} \left[ \hat{V}_{\text{MC}}(s) \right] + \text{Var} \left[ \hat{V}_{\text{MC}}(s') \right].$$

It now suffices to prove that for any  $s$ ,

$$\lim_{n \rightarrow \infty} \text{Var} \left[ \hat{V}_{\text{MC}}(s) \right] \geq \frac{\sigma_{\min}^2}{\mathbb{P}[s \in \tau]} \mathbb{E}[T | S_0 = s].$$

From Proposition B.4, we know the variance of the MC estimate is

$$\lim_{n \rightarrow \infty} n \cdot \text{Var} \left[ \hat{V}_{\text{MC}}(s) \right] = \frac{1}{\mathbb{P}[s \in \tau]} \sum_{\tilde{s} \in \mathcal{S}} \mathbb{E}[N(\tilde{s}) | S_0 = s] \sigma_V^2(\tilde{s}).$$

Using that  $\sigma_V^2(\tilde{s}) \geq \sigma_{\min}^2$  for all  $\tilde{s} \in \mathcal{S}$ , this expression simplifies to:

$$\lim_{n \rightarrow \infty} n \cdot \text{Var} \left[ \hat{V}_{\text{MC}}(s) \right] \geq \frac{\sigma_{\min}^2}{\mathbb{P}[s \in \tau]} \sum_{\tilde{s} \in \mathcal{S}} \mathbb{E}[N(\tilde{s}) | S_0 = s].$$

Finally,

$$\begin{aligned}
 \sum_{s' \in \mathcal{S}} \mathbb{E}[N(s') | S_0 = s] &= \sum_{\tilde{s} \in \mathcal{S}} \mathbb{E} \left[ \sum_{t=0}^T \mathbb{1}(S_t = \tilde{s}) | S_0 = s \right] \\
 &= \mathbb{E} \left[ \sum_{t=0}^T \sum_{\tilde{s} \in \mathcal{S}} \mathbb{1}(S_t = \tilde{s}) | S_0 = s \right] \\
 &= \mathbb{E}[T | S_0 = s]
 \end{aligned}$$

### B.5. TD always improves over MC

Theorem 8.4 and Proposition 8.1 show that the error of the TD estimate of the advantage scales with the crossing time while the error of the MC estimate of the advantage scales with the full horizon, hinting that TD can provide significantly more precise advantage estimates. However, in some cases, as described in Figure 6b, the crossing time is as large as the full horizon. One can then wonder if there are any setting in which MC can outperform TD for advantage estimation. Theorem 7.2 shows that this is never the case for value-to-go estimation. For completeness, we show that, under a technical condition, this is also never the case for advantage estimation.

**Proposition B.14.** For  $s, s' \in \mathcal{S}$  such that  $\mathbb{P}[s \in \tau \wedge s' \in \tau] = 0$ ,

$$\lim_{n \rightarrow \infty} \frac{\mathbb{E} \left[ \left( \hat{\mathbb{A}}_{\text{TD}}(s, s') - \mathbb{A}(s, s') \right)^2 \right]}{\mathbb{E} \left[ \left( \hat{\mathbb{A}}_{\text{MC}}(s, s') - \mathbb{A}(s, s') \right)^2 \right]} \leq 1$$

*Proof.* From Lemma B.5 and the proof of Proposition 8.1, we have:

$$\begin{aligned} \lim_{n \rightarrow \infty} n \cdot \mathbb{E} \left[ \left( \hat{J}_{\text{TD}}(\pi) - J(\pi) \right)^2 \right] &= \sum_{\tilde{s} \in \mathcal{S}} \frac{(\mathbb{E}[N(\tilde{s}) | S_0 = s] - \mathbb{E}[N(\tilde{s}) | S_0 = s'])^2}{\mathbb{E}[N(\tilde{s})]} \sigma_V^2(\tilde{s}) \\ \lim_{n \rightarrow \infty} n \cdot \mathbb{E} \left[ \left( \hat{J}_{\text{MC}}(\pi) - J(\pi) \right)^2 \right] &= \sum_{\tilde{s} \in \mathcal{S}} \left( \frac{\mathbb{E}[N(\tilde{s}) | S_0 = s]}{\mathbb{P}[s \in \tau]} + \frac{\mathbb{E}[N(\tilde{s}) | S_0 = s']}{\mathbb{P}[s' \in \tau]} \right) \sigma_V^2(\tilde{s}) \end{aligned}$$

Since  $\mathbb{P}[s \in \tau \wedge s' \in \tau] = 0$ , we have:

$$\begin{aligned} \mathbb{E}[N(\tilde{s})] &\geq \mathbb{E}[N(\tilde{s}) | s \in \tau] \mathbb{P}[s \in \tau] + \mathbb{E}[N(\tilde{s}) | s' \in \tau] \mathbb{P}[s' \in \tau] \\ &\geq \mathbb{E}[N(\tilde{s}) | S_0 = s] \mathbb{P}[s \in \tau] + \mathbb{E}[N(\tilde{s}) | S_0 = s'] \mathbb{P}[s' \in \tau] \end{aligned}$$

Plugging this into the limiting error of the TD estimates, we have:

$$\begin{aligned} \lim_{n \rightarrow \infty} n \cdot \mathbb{E} \left[ \left( \hat{J}_{\text{TD}}(\pi) - J(\pi) \right)^2 \right] &\leq \sum_{\tilde{s} \in \mathcal{S}} \frac{(\mathbb{E}[N(\tilde{s}) | S_0 = s] - \mathbb{E}[N(\tilde{s}) | S_0 = s'])^2}{\mathbb{E}[N(\tilde{s}) | S_0 = s] \mathbb{P}[s \in \tau] + \mathbb{E}[N(\tilde{s}) | S_0 = s'] \mathbb{P}[s' \in \tau]} \\ &\leq \sum_{\tilde{s} \in \mathcal{S}} \max \left( \frac{\mathbb{E}[N(\tilde{s}) | S_0 = s]}{\mathbb{P}[s \in \tau]} + \frac{\mathbb{E}[N(\tilde{s}) | S_0 = s']}{\mathbb{P}[s' \in \tau]} \right) \\ &\leq \lim_{n \rightarrow \infty} n \cdot \mathbb{E} \left[ \left( \hat{J}_{\text{MC}}(\pi) - J(\pi) \right)^2 \right] \end{aligned}$$

□

### B.6. On tightness of bounds of advantage estimation

We start by showing that the lower bound on the MSE of the MC estimate of the advantage is tight.

**Proposition B.15.** Let  $\mathcal{M}$  be a MRP such that  $\mathbb{E}[R(s, s')] = 0$ ,  $\text{Var}[R(s, s')] = V$  for all  $s, s'$ . For all  $s, s' \in \mathcal{S}$  such that  $\mathbb{P}[s \in \tau \wedge s' \in \tau] = 0$ ,

$$\lim_{n \rightarrow \infty} n \cdot \mathbb{E} \left[ \left( \hat{\mathbb{A}}_{\text{MC}}(s, s') - \mathbb{A}(s, s') \right)^2 \right] = \sigma_{\min}^2 \left( \frac{\mathbb{E}[T | S_0 = s]}{\mathbb{P}[s \in \tau]} + \frac{\mathbb{E}[T | S_0 = s']}{\mathbb{P}[s' \in \tau]} \right).$$

*Proof.* The only inequality in the proof of Proposition 8.1 is

$$\sigma_V^2(\tilde{s}) \geq \sigma_{\min}^2$$



We show that this inequality is an equality when  $\mathbb{E}[R(s, s')] = 0$  and  $\text{Var}[R(s, s')] = V$  for all  $s, s'$ . In particular, that  $V(s) = 0$  for all  $s$ .

$$\begin{aligned}\sigma_V^2(\tilde{s}) &= \text{Var}[R_0 + V(S_1)|S_0 = s] \\ &= \mathbb{E}[\text{Var}[R_0 + V(S_1)|S_0 = s, S_1]|S_0 = s] + \text{Var}[\mathbb{E}[R_0 + V(S_1)|S_0 = s]|S_0 = s] \\ &= \mathbb{E}[\text{Var}[R(S_0, S_1)]|S_0 = s] \\ &= V\end{aligned}$$

Hence,  $\sigma_V^2(\tilde{s})$  is constant equal to  $V$ . We therefore have the equality

$$\sigma_V^2(\tilde{s}) = \sigma_{\min}^2.$$

□

In the case of the advantage estimation using TD learning, we can prove a lower bound on the MSE. The gap between the lower bound and the upper bound is a factor 2.

**Proposition B.16.** *There exists a MRP  $\mathcal{M} = \{\mathcal{S} \cup \{\emptyset\}, P, R, d\}$  and  $s, s' \in \mathcal{S}$  such that  $\mathbb{P}[s \in \tau \wedge s' \in \tau] = 0$  and*

$$\lim_{n \rightarrow \infty} n \cdot \mathbb{E} \left[ \left( \hat{\mathbb{A}}_{\text{TD}}(s, s') - \mathbb{A}(s, s') \right)^2 \right] \geq \sigma_{\max}^2 \left( \frac{1}{\min(\mathbb{P}[s \in \tau], \mathbb{P}[s' \in \tau])} \right) \cdot H(s, s')$$

*Proof.* We define formally a set of MRP on which the bound is achieved. The structure is pictured in Figure 7. Let  $\mathcal{M}_{W,T,H}$  be the MRP with

- States  $\mathcal{S}_{W,T,H} = \{s_{(w,t)} | w \in \{1, \dots, W\}, t \in \{1, \dots, H-1\}\} \cap \{s_t | t \in \{H, \dots, T-1\}\}$ .
- Transitions  $P(s_{(w,t+1)} | s_{(w,t)}) = 1$  for all  $1 \leq w \leq W, 1 \leq t \leq H-2$ ,  $P(s_H | s_{(w,H-1)}) = 1$  for all  $1 \leq w \leq W$ ,  $P(s_{t+1} | s_t) = 1$  for all  $H \leq t \leq T-2$  and  $P(s_{T-1} | \emptyset) = 1$
- $\text{Var}[R(s, s')] = V$  and  $\mathbb{E}[R(s, s')] = 0$  for all  $s, s'$

First, we note that the TD estimate of the advantage between  $s_{(i,1)}$  and  $s_{(j,1)}$  in  $\mathcal{M}_{W,T,H}$  has the same distribution as the MC estimate of the advantage in  $\mathcal{M}_{W,H,H}$ . However, from Proposition B.15, we know that

$$\lim_{n \rightarrow \infty} n \cdot \mathbb{E} \left[ \left( \hat{\mathbb{A}}_{\text{MC}, \mathcal{M}_{W,H,H}}(s_{(i,1)}, s_{(j,1)}) - \mathbb{A}_{\mathcal{M}_{W,H,H}}(s_{(i,1)}, s_{(j,1)}) \right)^2 \right] = HV \left( \frac{1}{\mathbb{P}[s_{(i,1)} \in \tau]} + \frac{1}{\mathbb{P}[s_{(j,1)} \in \tau]} \right).$$

Noting that  $\mathbb{A}_{\mathcal{M}_{W,H,H}}(s_{(i,1)}, s_{(j,1)}) = \mathbb{A}_{\mathcal{M}_{W,T,H}}(s_{(i,1)}, s_{(j,1)})$  (since all value-to-go are zero), we obtain:

$$\begin{aligned}\lim_{n \rightarrow \infty} n \cdot \mathbb{E} \left[ \left( \hat{\mathbb{A}}_{\text{TD}, \mathcal{M}_{W,T,H}}(s, s') - \mathbb{A}(s, s') \right)^2 \right] &= HV \left( \frac{1}{\mathbb{P}[s_{(i,1)} \in \tau]} + \frac{1}{\mathbb{P}[s_{(j,1)} \in \tau]} \right) \\ &\geq HV \frac{1}{\min(\mathbb{P}[s_{(i,1)} \in \tau], \mathbb{P}[s_{(j,1)} \in \tau])} \\ &\geq \sigma_{\max}^2 \left( \frac{1}{\min(\mathbb{P}[s_{(i,1)} \in \tau], \mathbb{P}[s_{(j,1)} \in \tau])} \right) \cdot H(s_{(i,1)}, s_{(j,1)})\end{aligned}$$

□

As mentioned in Section 8, the upper bound on the MSE of TD estimates does not take into account any pooling that can happen before the crossing time. Hence, it is generally loose. The tightness result is achieved on MRPs with no pooling before the crossing time. On the other hand, the tightness of the lower bound on MC estimates only requires to control the distribution of the rewards and the transition structure of the MRP is not important.