

VALUE-AWARE IMPORTANCE WEIGHTING FOR OFF-POLICY REINFORCEMENT LEARNING

Kristopher De Asis

Department of Computing Science
University of Alberta
Edmonton, Canada
kldeasis@ualberta.ca

Eric Graves

Department of Computing Science
University of Alberta
Edmonton, Canada
graves@ualberta.ca

Richard S. Sutton

Department of Computing Science
University of Alberta
Edmonton, Canada
rsutton@ualberta.ca

ABSTRACT

Importance sampling is a central idea underlying off-policy prediction in reinforcement learning. It provides a strategy for re-weighting samples from a distribution to obtain unbiased estimates under another distribution. However, importance sampling weights tend to exhibit extreme variance, often leading to stability issues in practice. In this work, we consider a broader class of importance weights to correct samples in off-policy learning. We propose the use of *value-aware importance weights* which take into account the sample space to provide lower variance, but still unbiased, estimates under a target distribution. We derive how such weights can be computed, and detail key properties of the resulting importance weights. We then extend several reinforcement learning prediction algorithms to the off-policy setting with these weights, and evaluate them empirically.

1 OFF-POLICY REINFORCEMENT LEARNING

Value functions are central to reinforcement learning. However, most algorithms for learning a value function are limited to *on-policy* learning, where a specific policy of interest must be executed in order to estimate its value function. This paradigm can be inefficient, especially in the case of long-lived agents seeking to learn about many possible ways of behaving (Sutton et al., 2011; White, 2015). On-policy learning can also be expensive or dangerous in real-world applications such as robotics, recommendation systems (Maystre et al., 2023), and healthcare (Liao et al., 2021).

In contrast, algorithms capable of *off-policy* learning can learn the value function for a policy from data collected by a different policy. This flexibility allows a lifelong learning agent to learn about many possible policies at the same time, in parallel (Sutton et al., 1999; Klissarov & Precup, 2021). It also allows an agent to learn about the optimal policy while following an exploratory policy, as in Q-learning (Watkins & Dayan, 1992). Off-policy learning can also improve sample efficiency and stability by allowing an agent to learn from data generated by older policies, e.g., as in experience replay (Lin, 1992; Mnih et al., 2013; Schaul et al., 2016). Taken to the extreme, off-policy learning can completely remove the need for environmental interaction during training, instead allowing the agent to learn offline from datasets generated by human experts, existing control algorithms, or even random behaviour (Levine et al., 2020).

A common strategy for extending algorithms to use off-policy data is called *importance sampling* (Kahn, 1950a;b; Kloek & Van Dijk, 1978). When an action is chosen according to a different *behaviour policy* than the *target policy* the agent seeks to learn about, importance sampling scales the observed outcome by the ratio of the action’s probability under the target and behaviour policies, known as the *importance sampling ratio*. Scaling an observed outcome (e.g., returns, policy gradients, etc.) in this way leads to an unbiased estimate of its expectation under the target policy.

However, off-policy estimators based on importance sampling tend to exhibit extreme variance in practice. In the worst case the variance of importance sampling-based estimators can scale exponentially with the horizon of the problem (Liu et al., 2020). Several approaches have been introduced to reduce variance, such as per-decision importance sampling (Precup et al., 2000), clipped importance sampling ratios (Munos et al., 2016; Espeholt et al., 2018), self-normalized (or weighted) importance sampling (Rubinstein, 1981; Mahmood et al., 2014), and the DICE family of algorithms that learn parametric models of importance weights (Nachum et al., 2019; Zhang et al., 2020). However, none of these approaches is strictly better than the others, with weighted importance sampling, clipping, and the DICE algorithms introducing bias and/or dependence on function approximation, and the variance of per-decision importance sampling can still be unacceptably large in practice (Precup et al., 2000).

In this work, we propose the use of *value-aware importance weighting*. Like importance sampling, we correct samples by a multiplicative weight, but further take into account knowledge of the sample space (the possible outcomes of a random variable) to produce weights of lower variance. Our key contributions are as follows:

- Introduction of value-aware importance weighting as a general framework for off-policy corrections
- Derivation and discussion of an instance of value-aware importance weights with desirable properties
- Extension of existing algorithms with the derived importance weights, and empirical evaluation of the resulting algorithms.

We emphasize our focus on correcting the update target, but not the state-visitation distribution, as it is a simpler setting where variance concerns already manifest. Nevertheless, we present preliminary results with Emphatic TD methods in Appendix D showcasing the generality of our estimator and its compatibility with methods which re-weight updates.

2 BACKGROUND

Reinforcement learning is typically modeled as a Markov decision process (MDP). At each discrete time step t , an agent observes the current state $S_t \in \mathcal{S}$ and selects an action $A_t \in \mathcal{A}(S_t)$, where \mathcal{S} is the set of all states in the MDP, and $\mathcal{A}(s)$ is the set of available actions in state s . After selecting an action, the agent receives information about the next state $S_{t+1} \in \mathcal{S}$, along with a reward $R_{t+1} \in \mathbb{R}$, sampled according to the MDP’s transition model $p(s', r|s, a) = P(S_{t+1} = s', R_{t+1} = r|S_t = s, A_t = a)$. Actions are selected according to a policy $\pi(a|s) = P(A_t = a|S_t = s)$, and we are interested in *returns*:

$$G_t = \sum_{k=0}^{T-t-1} \gamma^k R_{t+k+1} \quad (1)$$

where $\gamma \in [0, 1]$ and T being the final time step of an episode, or $\gamma \in [0, 1]$ and $T = \infty$ in a continuing problem. Reinforcement learning control aims to compute an optimal policy π^* which maximizes the *expected* return.

Value-based methods for reinforcement learning aim to compute value functions, which are defined to be expected returns conditioned on various quantities. Here we emphasize the *action-value function*, which conditions on state s , immediate action a , and policy π to follow thereafter:

$$q_\pi(s, a) = \mathbb{E}_\pi[G_t|S_t = s, A_t = a] \quad (2)$$

Computing value functions for a policy is denoted *policy evaluation*, after which a value function may inform decision making through *policy improvement*: A policy π' derived to be greedy with respect to q_π is an improved policy, where $q_{\pi'}(s, a) \geq q_\pi(s, a), \forall s, a$. *Policy iteration* is the process of alternating policy evaluation and policy improvement to approach an optimal policy.

Exact value function computation may be impractical due to unreasonably large state spaces and imperfect or unknown transition models. As such, solution methods typically compute approximate value functions, $Q \approx q_\pi$, informed by sampled transitions generated by interaction with the environment, e.g., with an incremental update:

$$Q(S_t, A_t) \leftarrow Q(S_t, A_t) + \alpha(\hat{G}_t - Q(S_t, A_t))$$

where \hat{G}_t is a sample-based estimate of the return from (S_t, A_t) and following π , and $\alpha \in (0, 1]$ is the step-size.

Off-policy learning focuses on policy evaluation, but considers samples from the environment generated by a *behavior policy* μ , while aiming to evaluate a *target policy* π . One technique to account for this discrepancy is the use of *importance sampling* (Rubinstein, 1981; Hammersley & Handscomb, 1964), where a sample drawn from one distribution is re-weighted to be as if under another distribution:

$$\begin{aligned} \rho_t &= \frac{\pi(A_t|S_t)}{\mu(A_t|S_t)} \\ \mathbb{E}_\mu[\rho_t X(S_t, A_t)] &= \mathbb{E}_\pi[X(S_t, A_t)] \end{aligned} \quad (3)$$

where ρ_t is the importance weight, and we assume every action π may select must also be possible under μ . To correct the sequence of actions involved in a return, one can re-weight it with a product of these importance weights:

$$\prod_{k=t}^{T-1} \frac{\pi(A_k|S_k)p(S_{k+1}|S_k, A_k)}{\mu(A_k|S_k)p(S_{k+1}|S_k, A_k)} = \prod_{k=t}^{T-1} \frac{\pi(A_k|S_k)}{\mu(A_k|S_k)} = \prod_{k=t}^{T-1} \rho_k$$

$$\mathbb{E}_\mu \left[\left(\prod_{k=t}^{T-1} \rho_k \right) G_t \right] = \mathbb{E}_\pi [G_t]$$

While such re-weighted samples provide unbiased estimates under the target policy, importance weights are well known to suffer large variance (Sutton & Barto, 2018). Nevertheless, many off-policy methods for reinforcement learning are built on the idea of importance sampling (Precup et al., 2000), many of which propose ways to manage its large variance, often at the cost of increased bias (Mahmood et al., 2014; Munos et al., 2016).

3 VALUE-AWARE IMPORTANCE WEIGHTING

The importance weights given by Equation 3 can be seen to only depend on the behavior and target policies. This has appeal in its intuitive simplicity, as well as being generally applicable regardless of a random variable’s sample space. In this section, we consider importance weights which additionally depend on a state’s immediate action-values (denoted $\check{\rho}$). We formulate it as a constrained optimization problem subject to the constraints:

1. $\sum_a \mu_a \check{\rho}_a Q_a = \sum_a \pi_a Q_a$
2. $\sum_a \mu_a \check{\rho}_a = 1$

where for notational convenience, subscripts condition on a specific action, and everything is implicitly conditioned on the same (arbitrary) state. The first constraint ensures that the weighted action-values sampled under the behavior policy averages into the expected value under the target policy. The second constraint constrains the expected importance weight to be 1. This was specified because it’s a property of importance sampling, and avoids arbitrary attenuation or amplification in anticipation of needing to correct a sequence. Given the variance concerns with importance sampling, a reasonable objective is to minimize $\text{Var}_\mu(\check{\rho})$. This choice leads to the following Lagrangian:

$$\mathcal{L}(\check{\rho}, \lambda_0, \lambda_1) = \sum_a \mu_a (\check{\rho}_a - 1)^2 + \lambda_0 \left(\sum_a \pi_a Q_a - \sum_a \mu_a \check{\rho}_a Q_a \right) + \lambda_1 \left(1 - \sum_a \mu_a \check{\rho}_a \right) \quad (4)$$

Setting $\nabla \mathcal{L}(\check{\rho}, \lambda_0, \lambda_1) = 0$ and solving for $\check{\rho}$, we get:

$$\check{\rho}_a = 1 + \frac{Q_a - \mathbb{E}_\mu[Q]}{\mathbb{E}_\mu[Q^2] - \mathbb{E}_\mu[Q]^2} (\mathbb{E}_\pi[Q] - \mathbb{E}_\mu[Q]) \quad (5)$$

See Appendix A for a complete derivation. We denote $\check{\rho}$ the *Sparho* estimator. Interestingly, the above expression has a form where an action-value is mean-centered and normalized by its variance (under μ), and then scaled by a measure of off-policyness. We emphasize that these weights are not learned or approximated, and are computed in closed-form given a state’s behavior and target policy probabilities, and its immediate action-values. Interestingly, the estimator does not directly work with the policy probabilities, and only uses expectations computed under them. This may be related to regression importance sampling (Hanna et al., 2018) in that a *most likely* behavior policy (given the observed data thus far) might be implicitly used. It’s unclear however if a direct connection can be made, as the above weights typically can’t be expressed as a ratio of two policies.

By construction, we have:

1. $\check{\rho} = \rho$ is the unique solution to the constraints when there are 2 or fewer actions.
2. $\check{\rho}Q \sim \mu$ is an unbiased estimator of $\mathbb{E}_\pi[Q]$. Despite the dependence on Q , by satisfying the constraints it provides an immediately unbiased estimate for the current value estimates.
3. $\text{Var}_\mu(\check{\rho}) \leq \text{Var}_\mu(\rho)$. Ordinary importance sampling weights are in the space of solutions which satisfy the constraints, and variance is explicitly minimized.
4. $\check{\rho}_a = \rho_a = 1$ when $\mu = \pi$. This trivially minimizes the objective when on-policy, maintaining strict generalization of the on-policy setting akin to importance sampling.

Despite being unbiased in estimating the expected action-value, inaccuracies in the value-function will still result in a biased estimate of the return. This is akin to the bias introduced by temporal difference methods (Sutton, 1988), and does not incur additional bias when correcting sampled action-values (e.g., off-policy Sarsa(0)), but will introduce bias if importance weights computed with inaccurate values are used to correct sampled rewards. Also of note, it may be possible for the importance weights to be negative. This may seem counter-intuitive in that it may lead to “undoing” updates, but should still give the correct expectation with lower variance. Nevertheless, in our empirical evaluation we found that the weights remained all positive the majority of the time, likely due to the second constraint tending the solutions toward the positive direction. Lastly, computing $\hat{\rho}$ requires first computing $\mathbb{E}_\pi[Q]$. It may seem redundant to compute the final solution only to form a sample-based estimator for the same quantity, but we argue that the benefit lies in the need to correct a sequence, as opposed to solely an immediate action (Detailed in Section 4). We also emphasize that the importance weights presented are for one choice of objective. We explore alternative objectives in Appendix E, but focus on the above estimator in this work.

Here we empirically evaluate the above estimator in an off-policy bandit prediction setting to gain intuition on the potential variance reduction, and observe how it might scale with the size of the action space. Bandit instances are randomly generated by sampling μ and π according to $\text{softmax}(\mathcal{N}_{|\mathcal{A}| \times 1}(0, \beta))$, and Q according to $\beta + \mathcal{N}_{|\mathcal{A}| \times 1}(0, \beta)$, with β tuning the determinism of the policies and the spread of the action-values. Using $\beta = 2$, we swept over action space sizes from 2 to 32,768. For each action space, we generated 10,000 bandit instances and measured $\text{Var}(\rho Q)$ with ρ as the importance sampling and Sparho estimators. We additionally compare clipped versions of the above estimator, akin to Retrace and Vtrace estimators (Munos et al., 2016; Espeholt et al., 2018) For the clipped versions of the previous estimators, we additionally measure $\text{Bias}^2(\rho Q)$ and $\mathbb{E}_\mu[\rho]$. Results are shown in Figure 1.

It can be seen that the Sparho estimator has substantially lower variance than the importance sampling estimator. Perhaps surprisingly, we observe a *decrease* in variance with the size of the action space, despite the exponential growth in that of the importance sampling estimator. With the clipped estimators, we see a substantial decrease in variance. However, the clipped Sparho estimator had lower variance, substantially lower bias, and maintained a larger expected importance weight than the clipped importance sampling estimator. Of note, the clipped Sparho estimator similarly exhibited very favorable scaling with the size of the action space.

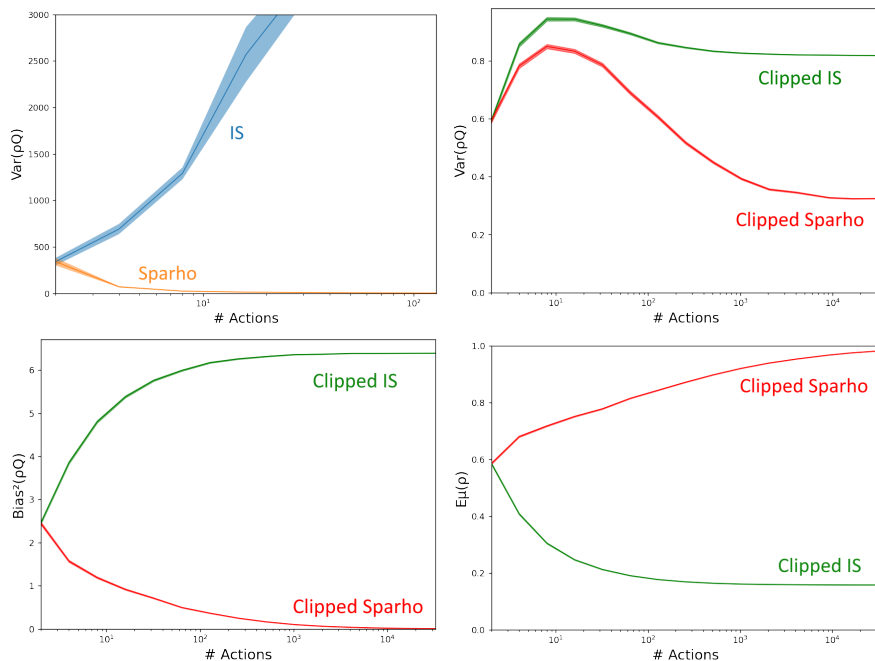


Figure 1: Various statistics against size of the action space, measured across randomly generated bandit instances. Results are averaged over 10,000 instances, and shaded regions represent one standard error.

The above results assume access to the true action-values, computing relevant statistics in closed form for a random bandit instance. We note that results are similar with learned values, and show a representative example in Appendix D.

4 MULTI-STEP CORRECTIONS

The action-value definition can be expressed in terms of successor values, denoted the *Bellman equation for q_π* :

$$q_\pi(s, a) = \sum_{s', r} p(s', r | s, a) \left(r + \gamma \sum_{a'} \pi(a' | s') q_\pi(s', a') \right) \quad (6)$$

We have that $\sum_a \pi(a | s) q_\pi(s, a) = \sum_a \mu(a | s) \check{\rho}(s, a) q_\pi(s, a)$, giving:

$$q_\pi(s, a) = \sum_{s', r} p(s', r | s, a) \left(r + \gamma \sum_{a'} \mu(a' | s') \check{\rho}(s', a') q_\pi(s', a') \right) \quad (7)$$

Should we expand the recursion multiple times, and then draw samples according to behavior policy μ , we get a multi-step estimator analogous to per-decision importance sampling:

$$\begin{aligned} G_t &= R_{t+1} + \gamma \check{\rho}_{t+1} G_{t+1} \\ G_t &= R_{t+1} + \gamma \check{\rho}_{t+1} R_{t+2} + \gamma^2 \check{\rho}_{t+1} \check{\rho}_{t+2} R_{t+3} + \gamma^3 \check{\rho}_{t+1} \check{\rho}_{t+2} \check{\rho}_{t+3} R_{t+4} + \dots \end{aligned} \quad (8)$$

Such a form, like importance sampling, conveniently lends itself to online, incremental algorithms (e.g., TD(λ)). An important note, however, is that $\check{\rho}_k$ is computed based on the current value estimates for state S_k , and if values are inaccurate, the expansion of the Bellman operator may result in a different quantity than what was used to compute the importance weights. This suggests that additional bias may be introduced in this multi-step extension, but also that the true value function remains the fixed-point. Of note, it's unclear whether the resulting operator is still a contraction, but we have not come across any example of multiple fixed-points in our empirical evaluation.

In the one-step case, a sample of Equation 7 gives Expected Sarsa, which is convergent under standard assumptions. The multi-step case is a bit more involved due to the form of the Sparho estimator, but in a loop-free episodic MDP, we have that a state-action pair preceding a terminal state will be convergent as it only predicts the immediate reward. By induction, we get a cascade of one-step Expected Sarsa estimators with eventual stable bootstrap targets. It is uncertain whether we have such guarantee for general MDPs due to the relatively complex expression for the importance weights, the existence of negative weights, and the non-stationarity of the weights as Q changes.

We compute the *expected every-visit* Monte Carlo update for the per-decision Sparho estimator given by Equation 8 from various starting values in a two-state MDP. This lets us visualize update directions and gain intuition for the bias introduced in the Monte Carlo extreme of the multi-step case. We use a step size $\alpha = 0.1$, and show the importance sampling estimator for reference. Figure 2 precisely specifies the MDP, and Figure 3 visualizes the expected updates.




Figure 2: Two-state MDP for visualizing expected multi-step updates.

As one might expect, the unbiased importance sampling estimator points directly at the fixed-point, and the arrows would directly jump there if $\alpha = 1$. The multi-step Sparho estimator generally has similar magnitude updates, but points in slightly different directions relating to its bias. Visualizing trajectories, the values seem to rush to an area near the fixed-point, before turning to approach it. Such a path may be an artifact of projecting higher-dimensional action-values into two state-values for visualization. Of particular note is the top right corner, where large jumps occur despite the small step size. We suspect that such large magnitude updates are due to interference between the transition matrix P_μ and the importance weights in the matrix $(I - \gamma P_\mu \text{diag}(\check{\rho}))^{-1}$. Under importance sampling, we have that $P_\pi = P_\mu \text{diag}(\check{\rho})$, but not with $\check{\rho}$. It's difficult to comment on the generality and extent of the phenomenon, as $\check{\rho}$ keeps changing, but in stochastic empirical evaluation (See Section 5) such jumps appeared indistinguishable from typical return variance, and did not appear to cause divergence.




Figure 3: Expected every-visit Monte Carlo updates for the per-decision importance sampling and Sparho estimators. The red circle denotes the fixed-point v_π , and sample trajectories of the Sparho estimator are visualized on the right.

Of note, n -step temporal difference (TD) estimators can attenuate the momentary large magnitude updates:

$$\begin{aligned} G_{t:h} &= R_{t+1} + \gamma \check{\rho}_{t+1} G_{t+1:h} \\ G_{t:t} &= Q(S_t, A_t) \end{aligned} \tag{9}$$

where $h = t + n$. Such attenuation is expected as it interpolates Monte Carlo returns with the one-step extreme, which was previously acknowledged to be Expected Sarsa. Figure 4 explores this for $n \in \{4, 8, 16\}$ for the above MDP. Perhaps surprisingly, the 16-step TD update has no noticeable large jumps, and appears closer to the 4-step updates than Monte Carlo despite the MDP having an expected trajectory length of around 2 steps. The top-right sample trajectory required n close to 100 to start overshooting the fixed-point akin to the Monte Carlo visualization.




Figure 4: Expected n -step TD updates for the per-decision Sparho estimator. The red circle denotes the fixed-point v_π , with sample trajectories from a variety of starting points are visualized.

5 COMPARISON WITH ONLINE OFF-POLICY METHODS

In this section, we empirically compare the multi-step Sparho estimator with importance sampling in a variety of environments. We specifically focus on the online, incremental setting with eligibility traces (Sutton, 1988; Precup et al., 2000), starting with the following λ -return:

$$G_t^\lambda = R_{t+1} + \gamma \left(\lambda (\rho_{t+1} G_{t+1}^\lambda + \mathbb{E}_\pi [Q(S_{t+1}, \cdot)] - \rho_{t+1} Q(S_{t+1}, A_{t+1})) + (1 - \lambda) \mathbb{E}_\pi [Q(S_{t+1}, \cdot)] \right) \tag{10}$$

which can be re-written as the following sum of one-step Expected Sarsa TD errors, weighted by an eligibility trace:

$$G_t^\lambda = Q(S_t, A_t) + \sum_{k=t}^{\infty} (R_{k+1} + \gamma \mathbb{E}_\pi [Q(S_{k+1}, \cdot)] - Q(S_k, A_k)) \prod_{i=t+1}^k \gamma \lambda \rho_i \tag{11}$$

We denote the incremental algorithm using the above λ -return as $Q(\lambda)$. It corresponds with the n -step per-decision estimator given by Equation 8, but with $\mathbb{E}_\pi[Q(S_{t+1}, \cdot)] - \rho_{t+1}Q(S_{t+1}, A_{t+1})$ appended to every reward beyond the immediate one. The appended terms have zero expected value and can be viewed as control variates, leading to a doubly-robust estimator (Jiang & Li, 2016; Thomas & Brunskill, 2016; De Asis & Sutton, 2018). They’re central to the presence of one-step Expected Sarsa TD errors in Equation 11. Without such terms, we would get Sarsa(λ) with per-decision importance sampling (Precup et al., 2000), where $\mathbb{E}_\pi[Q(S_{k+1}, \cdot)] \rightarrow \rho_{k+1}Q(S_{k+1}, A_{k+1})$, leading to additional variance on the order of what is presented in Figure 1. This choice of λ -return is in anticipation of an analogous Sparho estimator, which requires computing the expected action-value under the target policy anyways, leading to a fairer comparison. We denote the Sparho variant of the above estimator as Sparho(λ):

$$G_t^\lambda = Q(S_t, A_t) + \sum_{k=t}^{\infty} (R_{k+1} + \gamma \mathbb{E}_\pi[Q(S_{k+1}, \cdot)] - Q(S_k, A_k)) \prod_{i=t+1}^k \gamma \lambda \tilde{\rho}_i \quad (12)$$

Additionally, we compare against clipped variants of each estimator, where the importance weights in the above λ -returns are clipped to the range $[0, 1]$. Clipping them in Equation 11 gives Retrace(λ) (Munos et al., 2016), and we denote the clipped variant of Equation 12 as ReSparho(λ).

First, we evaluate the aforementioned algorithms in *Path World*. The environment can be thought of as a fully-connected feed-forward neural network with one input and one output. Each neuron is a state and each connection is an action, with the input being the initial state and the output the terminal state. The width of the network represents the size of the action space, $|\mathcal{A}|$, and the depth gives the number of decisions per episode. Each transition is assigned a fixed reward, and every run, behavior and target policy probabilities are generated per-state with the procedure of Section 3’s bandit experiments with $\beta = 1$. We fix the depth to 5, consider $|\mathcal{A}| \in \{8, 32\}$, and performed a systematic sweep over step size α and trace-decay rate λ . The root-mean-square (RMS) error between the learned values and the true values was measured after 10,000 steps for $|\mathcal{A}| = 8$, and after 500,000 steps for $|\mathcal{A}| = 32$. 30 independent runs were performed and results can be seen in Figure 5.

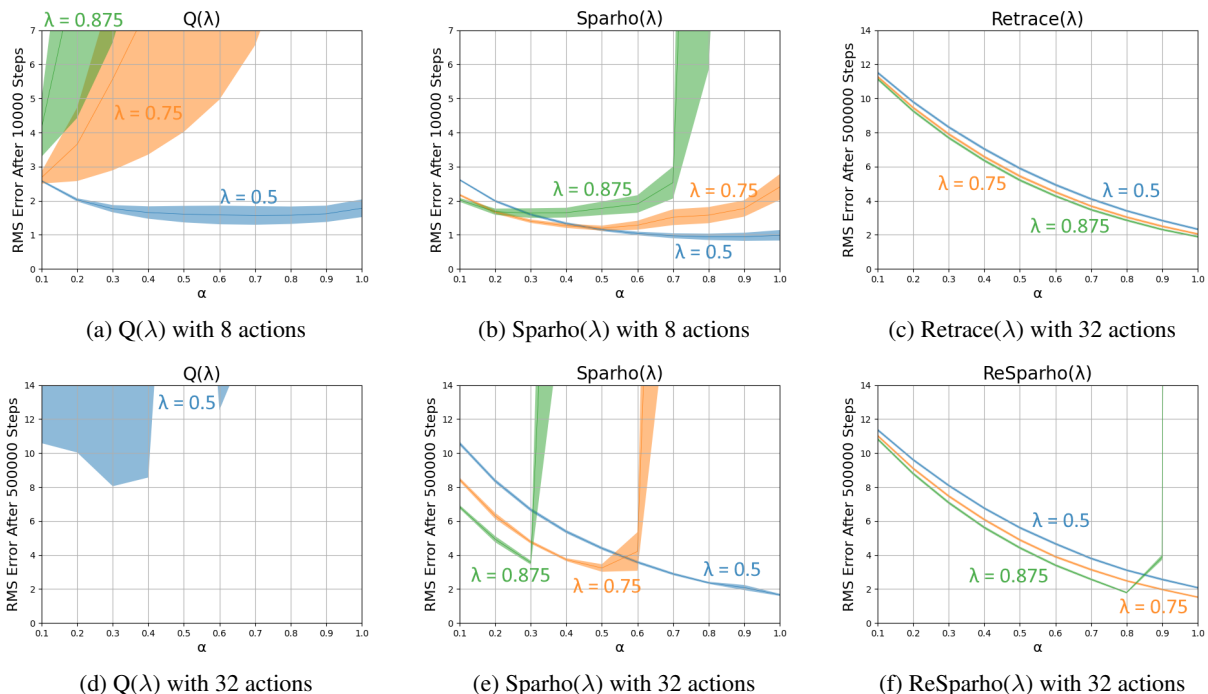


Figure 5: Results of an ablation study in the Path World environment for $|\mathcal{A}| \in \{8, 32\}$. Results are averaged over 30 runs, and shaded regions represent one standard error.

It can be seen that $Q(\lambda)$ suffers variance issues for $\lambda > 0.5$, and struggles further as the action space grows. In contrast, Sparho(λ)’s reduced variance enables the use of larger λ , but may be unstable for large (α, λ) pairs. For the clipped variants, we see substantially reduced variance, but a slower rate of learning due to the additional bias introduced by an expected importance weight < 1 . Of note, ReSparho(λ) still outperforms Retrace(λ) across most parameter settings,

consistent with the bandit results demonstrating a larger expected importance weight, but exhibits similar instability should the (α, λ) pair get too large. Such instability with large (α, λ) pairs are likely due to the potential momentary large-magnitude updates observed in Figure 3, from the interference between P_μ and $\text{diag}(\bar{\rho})$.

Next, we consider a 5×5 grid world with the center as the start state and two opposite corners as terminal states. Moving off of the grid keeps the agent in place, and a reward of -1 is received at each step. We first use a tabular representation with 4-directional movement, and we consider behavior and target policies which favor a specified action with probability $1 - \epsilon$, and behave uniformly random otherwise. In contrast with the Path World environment, this MDP sees considerably many state revisitations within an episode. Performing a systematic sweep over step size and trace-decay rate, we present learning curves for the best parameter setting in terms of RMS error after 20,000 steps. Results are averaged across 100 runs, and can be seen in Figure 6a.

To see how results scale to more difficult problems, we use a variant of the environment with 8-directional movement, and use feature vectors for use with linear function approximation. For each run, a random 16-bit binary feature vector is generated per state consisting of exactly 8 ones. Such a choice prevents perfectly representing the true values, and allows us to see the extent of potential bias the multi-step Sparho operator may introduce in the linear fixed-point. A similar parameter sweep was performed, and learning curves are presented for the best parameter setting in terms of RMS error after 50,000 steps. Results are averaged across 100 runs, and can be seen in Figure 6b.




Figure 6: Results in the 5×5 Grid World environment with tabular and linear representations. Results are averaged over 100 runs, and shaded regions represent one standard error.

With the tabular representation, we see that $Q(\lambda)$ was slowest to converge. $\text{Retrace}(\lambda)$ substantially improves over it, with both $\text{Sparho}(\lambda)$ and $\text{ReSparho}(\lambda)$ performing even better. Of note, improved performance here largely comes from the ability to tolerate larger trace-decay rates without becoming unstable, and explains why the curves for the Sparho estimators may appear to have higher variance. In the case of the clipped estimators, assuming $\lambda \leq 1$, the rate of convergence is bottle-necked by the expected importance weight magnitude. With the linear representation, we see a similar trend, but with $\text{Retrace}(\lambda)$ learning quicker than $\text{Sparho}(\lambda)$. Of note, despite the aforementioned bias introduced by the multi-step Sparho operator, there is no noticeable impact on the final RMS error achieved, with every algorithm but $Q(\lambda)$ converging to a similar distance to the true value function. This is likely due to the expectation-correction behavior (Harutyunyan et al., 2016) of the control variates in Equation 10.

Lastly, we evaluate the algorithms in the *Acrobot-v1* environment (Brockman et al., 2016), to demonstrate whether previous results extend toward control problems with non-linear function approximation. We used an ϵ -greedy behavior policy where ϵ was annealed from 1 to 0.1 over 25,000 steps, with a target policy deterministically greedy with respect to the current value estimates. Of note, all of the above algorithms are identical in the on-policy case. As a result, we evaluate each algorithm based on the first 25,000 steps, where there is a larger degree of off-policyness. Specifically, for each step we record the mean of the last 10 episodes, and measure the area under the curve (AUC). A 2-hidden-layer fully-connected neural network with *tanh* activations was used to represent the action-values, and we performed a sweep over step size, trace-decay rate, and hidden layer width. Averaged across 30 runs, the results for the best hidden layer width can be seen in Figure 7.

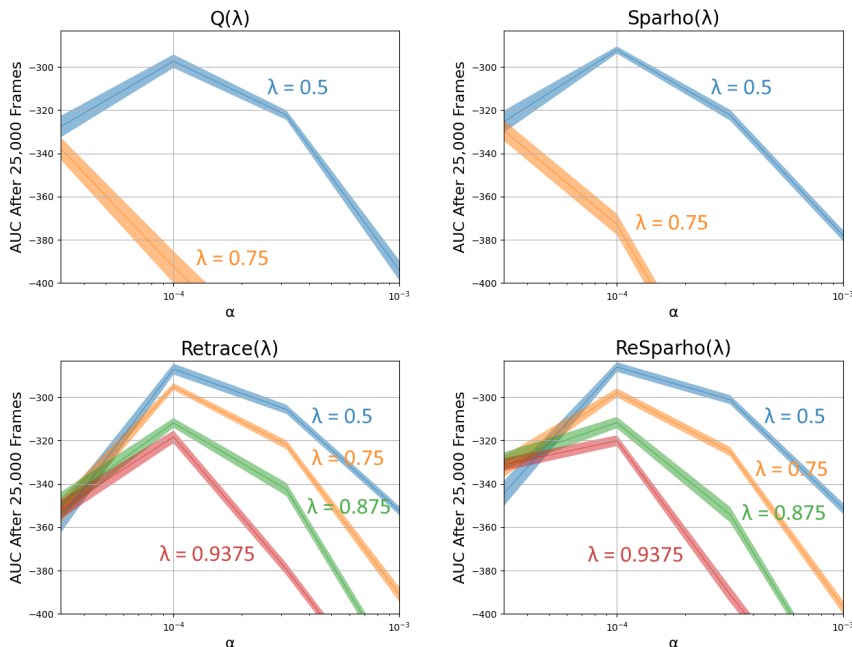


Figure 7: Results in the Acrobot-v1 environment with a hidden layer width of 512. Results are averaged over 30 runs, and shaded regions represent one standard error.

It can be seen that Sparho(λ) outperforms Q(λ), with lower variance in the episodic returns. ReSparho(λ) performs similar to Retrace(λ) when properly tuned, but exhibits flatter sensitivity curves in the direction of smaller α . Similarities in performance are likely due to relatively small degrees of off-policy (greedy vs. ϵ -greedy), the environment only having three actions, and how control performance might not require accurate value estimates. Nevertheless, all of the empirical evaluation taken together suggests that the improvements here are likely due to better value estimation.

6 CONCLUSION AND FUTURE WORK

In this work, we investigated the use of value-aware importance weights as a lower variance alternative to importance sampling. We derived a minimum-variance instance of such weights, empirically characterized the potential variance reduction, and analyzed the behavior of its per-decision multi-step extension. We then evaluated the online learning capabilities of our novel estimator and demonstrated improvements in both off-policy prediction and control, where the difference between pairs of algorithms is the substitution of value-aware importance weights.

This opens several avenues for future work. Our work focused on one instantiation of value-aware importance weights, but there are many alternate objectives one can consider in deriving new estimators, e.g., a non-negative constraint, adapted mixtures of importance sampling and Sparho for higher (α , λ) pairs, etc. Much of our analysis was through empirical evaluation, leaving room for theoretical analysis of the Sparho estimator and other comparable importance weights. We further acknowledge complimentary similarities with the more robust doubly-robust estimator (Farajtabar et al., 2018), where one instead searches the space of value estimates to minimize the estimator’s variance. This suggests the possibility of searching through the combined spaces of importance weights and value-estimates to further reduce variance. Lastly, it’s worth exploring other extensions for the multi-step case. Instead of incurring bias from only considering immediate action-values, a trajectory-aware algorithm could use future sampled information in place of the action-values of selected actions. However, it’s unclear if such a procedure can produce an online algorithm.

ACKNOWLEDGEMENTS

The authors were generously supported by Amii, NSERC, and CIFAR, and would like to thank Arsalan Sharifnassab and Huizhen Yu for insights and discussions contributing to the results presented in this paper. The authors further thank the reviewers for valuable feedback during the review process.

REFERENCES

- G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba. OpenAI gym. *CoRR*, abs/1606.01540, 2016.
- Kristopher De Asis and Richard S. Sutton. Per-decision multi-step temporal difference learning with control variates. In *Proceedings of the Thirty-Fourth Conference on Uncertainty in Artificial Intelligence, UAI 2018*, 07 2018.
- Lasse Espeholt, Hubert Soyer, Rémi Munos, Karen Simonyan, Volodymyr Mnih, Tom Ward, Yotam Doron, Vlad Firoiu, Tim Harley, Iain Dunning, Shane Legg, and Koray Kavukcuoglu. IMPALA: scalable distributed deep-rl with importance weighted actor-learner architectures. In *Proceedings of the Thirty-Fifth International Conference on Machine Learning*, 2018.
- Mehrdad Farajtabar, Mohammad Ghavamzadeh, and Yinlam Chow. More robust doubly robust off-policy evaluation. In *International Conference on Machine Learning*, 2018.
- John M. Hammersley and David C. Handscomb. *Monte Carlo methods*. Methuen London, 1964.
- Josiah P. Hanna, Scott Niekum, and Peter Stone. Importance sampling policy evaluation with an estimated behavior policy. In *International Conference on Machine Learning*, 2018.
- Anna Harutyunyan, Marc G. Bellemare, Tom Stepleton, and Rémi Munos. Q(\$\lambda\$) with off-policy corrections. *CoRR*, abs/1602.04951, 2016. URL <http://arxiv.org/abs/1602.04951>.
- Nan Jiang and Lihong Li. Doubly robust off-policy value evaluation for reinforcement learning. In *Proceedings of the Thirty-Third International Conference on Machine Learning*, 2016.
- Herman Kahn. Random sampling (monte carlo) techniques in neutron attenuation problems. i. *Nucleonics*, 6(See also NSA 3-990), 1950a.
- Herman Kahn. Random sampling (monte carlo) techniques in neutron attenuation problems. ii. *Nucleonics*, 6(See also NSA 4-3795), 1950b.
- Martin Klissarov and Doina Precup. Flexible option learning. *Advances in Neural Information Processing Systems*, 34, 2021.
- Teun Kloek and Herman K. Van Dijk. Bayesian estimates of equation system parameters: an application of integration by monte carlo. *Econometrica: Journal of the Econometric Society*, pp. 1–19, 1978.
- Sergey Levine, Aviral Kumar, George Tucker, and Justin Fu. Offline reinforcement learning: Tutorial, review, and perspectives on open problems. *arXiv preprint arXiv:2005.01643*, 2020.
- Peng Liao, Predrag Klasnja, and Susan Murphy. Off-policy estimation of long-term average outcomes with applications to mobile health. *Journal of the American Statistical Association*, 116(533):382–391, 2021.
- Long-Ji Lin. Self-improving reactive agents based on reinforcement learning, planning and teaching. *Machine Learning*, 8:293–321, 1992.
- Yao Liu, Pierre-Luc Bacon, and Emma Brunskill. Understanding the curse of horizon in off-policy evaluation via conditional importance sampling. In *International Conference on Machine Learning*, pp. 6184–6193. PMLR, 2020.
- A. Rupam Mahmood, Hado P. Van Hasselt, and Richard S. Sutton. Weighted importance sampling for off-policy learning with linear function approximation. *Advances in Neural Information Processing Systems*, 27, 2014.
- Lucas Maystre, Daniel Russo, and Yu Zhao. Optimizing audio recommendations for the long-term: A reinforcement learning perspective. *arXiv preprint arXiv:2302.03561*, 2023.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.
- Rémi Munos, Tom Stepleton, Anna Harutyunyan, and Marc G. Bellemare. Safe and efficient off-policy reinforcement learning. In *Proceedings of the Thirtieth Conference on Neural Information Processing Systems*, 2016.
- Ofir Nachum, Yinlam Chow, Bo Dai, and Lihong Li. Dualdice: Behavior-agnostic estimation of discounted stationary distribution corrections. *Advances in Neural Information Processing Systems*, 32, 2019.

- Doina Precup, Richard S. Sutton, and Satinder Singh. Eligibility traces for off-policy policy evaluation. *Computer Science Department Faculty Publication Series*, 06 2000.
- Reuven Y. Rubinstein. *Simulation and the Monte Carlo Method*. John Wiley & Sons, Inc., 1981. ISBN 0471089176.
- Tom Schaul, John Quan, Ioannis Antonoglou, and David Silver. Prioritized experience replay. In *Proceedings of the 4th International Conference on Learning Representations*, 2016.
- Richard S. Sutton. Learning to predict by the methods of temporal differences. *Machine learning*, 3(1):9–44, 1988.
- Richard S. Sutton and Andy G. Barto. *Reinforcement Learning: An Introduction*. The MIT Press, 2nd edition, 2018.
- Richard S. Sutton, Doina Precup, and Satinder Singh. Between mdps and semi-mdps: A framework for temporal abstraction in reinforcement learning. *Artificial intelligence*, 112(1-2):181–211, 1999.
- Richard S. Sutton, Joseph Modayil, Michael Delp, Thomas Degris, Patrick M. Pilarski, Adam White, and Doina Precup. Horde: A scalable real-time architecture for learning knowledge from unsupervised sensorimotor interaction. In *Proceedings of the 10th International Conference on Autonomous Agents and MultiAgent Systems*, 2011.
- Philip Thomas and Emma Brunskill. Data-efficient off-policy policy evaluation for reinforcement learning. In *Proceedings of the Thirty-Third International Conference on Machine Learning*, 2016.
- Christopher J. C. H. Watkins and Peter Dayan. Q-learning. *Machine learning*, 8(3):279–292, 1992.
- Adam White. *Developing a predictive approach to knowledge*. PhD thesis, University of Alberta, 2015.
- Ruiyi Zhang, Bo Dai, Lihong Li, and Dale Schuurmans. Gendice: Generalized offline estimation of stationary values. In *International Conference on Learning Representations*, 2020.

A SPARHO DERIVATION

We begin with the Lagrangian from Equation 4.

$$\mathcal{L}(\check{\rho}, \lambda_0, \lambda_1) = \sum_a \mu_a (\check{\rho}_a - 1)^2 + \lambda_0 \left(\sum_a \pi_a Q_a - \sum_a \mu_a \check{\rho}_a Q_a \right) + \lambda_1 \left(1 - \sum_a \mu_a \check{\rho}_a \right)$$

We can rewrite this in matrix form as follows:

$$\mathcal{L}(\check{\rho}, \lambda_0, \lambda_1) = (\check{\rho} - \mathbf{1})^T D_\mu (\check{\rho} - \mathbf{1}) + \lambda_0 (\boldsymbol{\pi}^T \mathbf{Q} - (D_\mu \mathbf{Q})^T \check{\rho}) + \lambda_1 (1 - (D_\mu \mathbf{1})^T \check{\rho}) \quad (13)$$

where D_μ is a diagonal matrix with entries of $\boldsymbol{\mu}$ along its diagonal. We assume $\boldsymbol{\mu}$ is soft, i.e., has a non-zero probability for each action, ensuring D_μ 's inverse exists. Setting $\nabla \mathcal{L}(\check{\rho}, \lambda_0, \lambda_1) = 0$ and solving for $\check{\rho}$, we get:

$$\begin{aligned} \nabla \mathcal{L}(\check{\rho}, \lambda_0, \lambda_1) = 0 &= 2D_\mu \check{\rho} - 2\boldsymbol{\mu} - D_\mu \mathbf{Q} \lambda_0 - D_\mu \mathbf{1} \lambda_1 \\ 0 &= D_\mu \check{\rho} - \boldsymbol{\mu} - D_\mu \mathbf{Q} \frac{\lambda_0}{2} - D_\mu \mathbf{1} \frac{\lambda_1}{2} \\ 0 &= D_\mu^{-1} D_\mu \check{\rho} - D_\mu^{-1} \boldsymbol{\mu} - D_\mu^{-1} D_\mu \mathbf{Q} \frac{\lambda_0}{2} - D_\mu^{-1} D_\mu \mathbf{1} \frac{\lambda_1}{2} \\ 0 &= \check{\rho} - \mathbf{1} - \mathbf{Q} \frac{\lambda_0}{2} - \mathbf{1} \frac{\lambda_1}{2} \\ \check{\rho} &= \mathbf{1} + \mathbf{Q} \frac{\lambda_0}{2} + \mathbf{1} \frac{\lambda_1}{2} \end{aligned} \quad (14)$$

Multiply both sides of Equation 14 by $(D_\mu \mathbf{Q})^T$ and apply one of the constraints:

$$\begin{aligned} (D_\mu \mathbf{Q})^T \check{\rho} &= (D_\mu \mathbf{Q})^T \mathbf{1} + (D_\mu \mathbf{Q})^T \mathbf{Q} \frac{\lambda_0}{2} + (D_\mu \mathbf{Q})^T \mathbf{1} \frac{\lambda_1}{2} \\ (D_\mu \mathbf{Q})^T \check{\rho} - (D_\mu \mathbf{Q})^T \mathbf{1} &= (D_\mu \mathbf{Q})^T \mathbf{Q} \frac{\lambda_0}{2} + (D_\mu \mathbf{Q})^T \mathbf{1} \frac{\lambda_1}{2} \\ \boldsymbol{\pi}^T \mathbf{Q} - \boldsymbol{\mu}^T \mathbf{Q} &= \boldsymbol{\mu}^T (\mathbf{Q} \circ \mathbf{Q}) \frac{\lambda_0}{2} + \boldsymbol{\mu}^T \mathbf{Q} \frac{\lambda_1}{2} \end{aligned} \quad (15)$$

Multiply both sides of Equation 14 by $(D_\mu \mathbf{1})^T$ and apply the other constraint:

$$\begin{aligned} (D_\mu \mathbf{1})^T \check{\rho} &= (D_\mu \mathbf{1})^T \mathbf{1} + (D_\mu \mathbf{1})^T \mathbf{Q} \frac{\lambda_0}{2} + (D_\mu \mathbf{1})^T \mathbf{1} \frac{\lambda_1}{2} \\ (D_\mu \mathbf{1})^T \check{\rho} - (D_\mu \mathbf{1})^T \mathbf{1} &= (D_\mu \mathbf{1})^T \mathbf{Q} \frac{\lambda_0}{2} + (D_\mu \mathbf{1})^T \mathbf{1} \frac{\lambda_1}{2} \\ 0 &= \boldsymbol{\mu}^T \mathbf{Q} \frac{\lambda_0}{2} + \frac{\lambda_1}{2} \end{aligned} \quad (16)$$

We can now use Equations 15 and 16 to solve for the Lagrange multipliers:

$$\begin{aligned} \begin{bmatrix} \boldsymbol{\mu}^T (\mathbf{Q} \circ \mathbf{Q}) & \boldsymbol{\mu}^T \mathbf{Q} \\ \boldsymbol{\mu}^T \mathbf{Q} & 1 \end{bmatrix} \begin{bmatrix} \frac{\lambda_0}{2} \\ \frac{\lambda_1}{2} \end{bmatrix} &= \begin{bmatrix} \boldsymbol{\pi}^T \mathbf{Q} - \boldsymbol{\mu}^T \mathbf{Q} \\ 0 \end{bmatrix} \\ \begin{bmatrix} \frac{\lambda_0}{2} \\ \frac{\lambda_1}{2} \end{bmatrix} &= \begin{bmatrix} \boldsymbol{\mu}^T (\mathbf{Q} \circ \mathbf{Q}) & \boldsymbol{\mu}^T \mathbf{Q} \\ \boldsymbol{\mu}^T \mathbf{Q} & 1 \end{bmatrix}^{-1} \begin{bmatrix} \boldsymbol{\pi}^T \mathbf{Q} - \boldsymbol{\mu}^T \mathbf{Q} \\ 0 \end{bmatrix} \\ \frac{\lambda_0}{2} &= \frac{\boldsymbol{\pi}^T \mathbf{Q} - \boldsymbol{\mu}^T \mathbf{Q}}{\boldsymbol{\mu}^T (\mathbf{Q} \circ \mathbf{Q}) - (\boldsymbol{\mu}^T \mathbf{Q})^2} \end{aligned} \quad (17)$$

$$\frac{\lambda_1}{2} = -\boldsymbol{\mu}^T \mathbf{Q} \frac{\boldsymbol{\pi}^T \mathbf{Q} - \boldsymbol{\mu}^T \mathbf{Q}}{\boldsymbol{\mu}^T (\mathbf{Q} \circ \mathbf{Q}) - (\boldsymbol{\mu}^T \mathbf{Q})^2} \quad (18)$$

Substituting Equations 17 and 18 back into Equation 14, we get:

$$\begin{aligned} \check{\rho} &= \mathbf{1} + \mathbf{Q} \frac{\boldsymbol{\pi}^T \mathbf{Q} - \boldsymbol{\mu}^T \mathbf{Q}}{\boldsymbol{\mu}^T (\mathbf{Q} \circ \mathbf{Q}) - (\boldsymbol{\mu}^T \mathbf{Q})^2} - \boldsymbol{\mu}^T \mathbf{Q} \frac{\boldsymbol{\pi}^T \mathbf{Q} - \boldsymbol{\mu}^T \mathbf{Q}}{\boldsymbol{\mu}^T (\mathbf{Q} \circ \mathbf{Q}) - (\boldsymbol{\mu}^T \mathbf{Q})^2} \\ \check{\rho} &= \mathbf{1} + \frac{\mathbf{Q} - \boldsymbol{\mu}^T \mathbf{Q}}{\boldsymbol{\mu}^T (\mathbf{Q} \circ \mathbf{Q}) - (\boldsymbol{\mu}^T \mathbf{Q})^2} (\boldsymbol{\pi}^T \mathbf{Q} - \boldsymbol{\mu}^T \mathbf{Q}) \end{aligned}$$

Replacing the inner products with the expectations they correspond with gives us:

$$\check{\rho} = \mathbf{1} + \frac{\mathbf{Q} - \mathbb{E}_\mu[Q]}{\mathbb{E}_\mu[Q^2] - \mathbb{E}_\mu[Q]^2} (\mathbb{E}_\pi[Q] - \mathbb{E}_\mu[Q]) \quad (19)$$

B λ -RETURN BACKWARD-VIEW DERIVATION

We start with the recursive definition of the λ -return for Sarsa(λ) (Precup et al., 2000), one of the simplest TD(λ) algorithms for off-policy action-value estimation:

$$G_t^\lambda = R_{t+1} + \gamma \left(\lambda \rho_{t+1} G_{t+1}^\lambda + (1 - \lambda) \rho_{t+1} Q(S_{t+1}, A_{t+1}) \right)$$

Adding the control variate terms to the discounted future rewards, we get:

$$\begin{aligned} G_t^\lambda &= R_{t+1} + \gamma \left(\lambda \rho_{t+1} G_{t+1}^\lambda + (1 - \lambda) \rho_{t+1} Q(S_{t+1}, A_{t+1}) + \mathbb{E}_\pi[Q(S_{t+1}, \cdot)] - \rho_{t+1} Q(S_{t+1}, A_{t+1}) \right) \\ G_t^\lambda &= R_{t+1} + \gamma \left(\lambda (\rho_{t+1} G_{t+1}^\lambda + \mathbb{E}_\pi[Q(S_{t+1}, \cdot)] - \rho_{t+1} Q(S_{t+1}, A_{t+1})) + \right. \\ &\quad \left. (1 - \lambda) (\rho_{t+1} Q(S_{t+1}, A_{t+1}) + \mathbb{E}_\pi[Q(S_{t+1}, \cdot)] - \rho_{t+1} Q(S_{t+1}, A_{t+1})) \right) \\ G_t^\lambda &= R_{t+1} + \gamma \left(\lambda (\rho_{t+1} G_{t+1}^\lambda + \mathbb{E}_\pi[Q(S_{t+1}, \cdot)] - \rho_{t+1} Q(S_{t+1}, A_{t+1})) + (1 - \lambda) \mathbb{E}_\pi[Q(S_{t+1}, \cdot)] \right) \end{aligned}$$

We can re-write this as a recursion of $G_t^\lambda - Q(S_t, A_t)$:

$$\begin{aligned} G_t^\lambda &= R_{t+1} + \gamma \left(\lambda (\rho_{t+1} G_{t+1}^\lambda + \mathbb{E}_\pi[Q(S_{t+1}, \cdot)] - \rho_{t+1} Q(S_{t+1}, A_{t+1})) + (1 - \lambda) \mathbb{E}_\pi[Q(S_{t+1}, \cdot)] \right) \\ G_t^\lambda &= R_{t+1} + \gamma \lambda \rho_{t+1} G_{t+1}^\lambda + \gamma \lambda \mathbb{E}_\pi[Q(S_{t+1}, \cdot)] - \gamma \lambda \rho_{t+1} Q(S_{t+1}, A_{t+1}) + \gamma \mathbb{E}_\pi[Q(S_{t+1}, \cdot)] - \gamma \lambda \mathbb{E}_\pi[Q(S_{t+1}, \cdot)] \\ G_t^\lambda &= R_{t+1} + \gamma \mathbb{E}_\pi[Q(S_{t+1}, \cdot)] + \gamma \lambda \rho_{t+1} G_{t+1}^\lambda - \gamma \lambda \rho_{t+1} Q(S_{t+1}, A_{t+1}) \\ G_t^\lambda &= R_{t+1} + \gamma \mathbb{E}_\pi[Q(S_{t+1}, \cdot)] + \gamma \lambda \rho_{t+1} (G_{t+1}^\lambda - Q(S_{t+1}, A_{t+1})) \\ G_t^\lambda - Q(S_t, A_t) &= R_{t+1} + \gamma \mathbb{E}_\pi[Q(S_{t+1}, \cdot)] - Q(S_t, A_t) + \gamma \lambda \rho_{t+1} (G_{t+1}^\lambda - Q(S_{t+1}, A_{t+1})) \end{aligned}$$

From here, we get a sum of one-step Expected Sarsa TD errors, discounted by $\gamma \lambda \rho_t$:

$$G_t^\lambda = Q(S_t, A_t) + \sum_{k=t}^{\infty} (R_{k+1} + \gamma \mathbb{E}_\pi[Q(S_{k+1}, \cdot)] - Q(S_k, A_k)) \prod_{i=t+1}^k \gamma \lambda \rho_i$$

To provide intuition on the effect of the control variate, following similar steps with the Sarsa(λ)'s λ -return gives:

$$\begin{aligned} G_t^\lambda &= R_{t+1} + \gamma \left(\lambda \rho_{t+1} G_{t+1}^\lambda + (1 - \lambda) \rho_{t+1} Q(S_{t+1}, A_{t+1}) \right) \\ G_t^\lambda &= R_{t+1} + \gamma \lambda \rho_{t+1} G_{t+1}^\lambda + \gamma \rho_{t+1} Q(S_{t+1}, A_{t+1}) - \gamma \lambda \rho_{t+1} Q(S_{t+1}, A_{t+1}) \\ G_t^\lambda &= R_{t+1} + \gamma \rho_{t+1} Q(S_{t+1}, A_{t+1}) + \gamma \lambda \rho_{t+1} G_{t+1}^\lambda - \gamma \lambda \rho_{t+1} Q(S_{t+1}, A_{t+1}) \\ G_t^\lambda &= R_{t+1} + \gamma \rho_{t+1} Q(S_{t+1}, A_{t+1}) + \gamma \lambda \rho_{t+1} (G_{t+1}^\lambda - Q(S_{t+1}, A_{t+1})) \\ G_t^\lambda - Q(S_t, A_t) &= R_{t+1} + \gamma \rho_{t+1} Q(S_{t+1}, A_{t+1}) - Q(S_t, A_t) + \gamma \lambda \rho_{t+1} (G_{t+1}^\lambda - Q(S_{t+1}, A_{t+1})) \\ G_t^\lambda &= Q(S_t, A_t) + \sum_{k=t}^{\infty} (R_{k+1} + \gamma \rho_{k+1} Q(S_{k+1}, A_{k+1}) - Q(S_k, A_k)) \prod_{i=t+1}^k \gamma \lambda \rho_i \end{aligned}$$

We see that the control variate is responsible for the use of one-step Expected Sarsa TD errors being weighted by the eligibility trace, as opposed to those of one-step Sarsa with importance sampling. It's clear that computing the complete expectation is typically favorable over relying on averaging out importance weighted-samples in the bootstrap target, and can be expected to result in increased variance on the order of what is presented in Figure 1.

C ADDITIONAL EXPERIMENTAL DETAILS

C.1 PATH WORLD DETAILS

Given actions $a \in \{0, 1, 2, \dots, |\mathcal{A}|\}$, rewards were specified by $r(s, a) = \frac{1+a}{|\mathcal{A}|}$.

The scope of the parameter sweep is detailed below:

Parameter	Value(s)
Discount Factor γ	1.0
Action Space Size $ \mathcal{A} $	8, 32
Step Size α	0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1.0
Trace Decay λ	0.5, 0.75, 0.875

C.2 5 × 5 GRID WORLD DETAILS

The scope of the parameter sweep for the tabular results are detailed below:

Parameter	Value(s)
Discount Factor γ	1.0
Action Space Size $ \mathcal{A} $	4
Action Commitment ϵ_π	0.5
Action Commitment ϵ_μ	1.0
Step Size α	0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1.0
Trace Decay λ	0.5, 0.75, 0.875, 0.9375, 0.96875, 0.984375, 0.9921875

The scope of the parameter sweep for the linear function approximation results are detailed below:

Parameter	Value(s)
Discount Factor γ	1.0
Action Space Size $ \mathcal{A} $	8
Action Commitment ϵ_π	0.5
Action Commitment ϵ_μ	0.5
Step Size α	0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1.0
Trace Decay λ	0.5, 0.75, 0.875, 0.9375

C.3 ACROBOT-V1 DETAILS

The scope of the parameter sweep is detailed below:

Parameter	Value(s)
Discount Factor γ	0.99
Time Limit	500
Hidden Layers	2
Hidden Layer Width	32, 64, 128, 256, 512
Step Size α	$10^{-4.5}$, 10^{-4} , $10^{-3.5}$, 10^{-3}
Trace Decay λ	0.5, 0.75, 0.875, 0.9375
Initial ϵ	1.0
Final ϵ	0.1
Linear ϵ Decay Steps	25000

D ADDITIONAL EXPERIMENTAL RESULTS

D.1 ONLINE BANDIT RESULTS

Below we consider the random bandit instances in the case where values are being learned, as the results presented in Figure 1 compute the relevant metrics in closed form given the true action-values. In addition to importance sampling, Sparho, and their clipped versions, we further evaluate regression importance sampling (Hanna et al., 2018) applied to both the importance sampling and Sparho estimators. We fix $|\mathcal{A}| = 8$, $\alpha = 0.01$ and note that the trends are reasonably consistent across the parameter space due to the simplicity of the problem. Averaged over 100 runs, we measure the absolute value error across 1,000 steps.




Figure 8: Absolute value error for various estimators across randomly generated bandit instances. Results are averaged over 100 instances, and shaded regions represent one standard error.

We see that even when action-values are being estimated, the Sparho estimator has substantially lower variance than the importance sampling counterpart for all variants tested (original, clipped, RIS). Given the use of a fixed step-size, this further leads to lower final value error as there is less variability to average out. We see general decrease in variance with the use of an estimated behavior policy for both RIS estimators. We further observe substantial bias in the solution of the clipped estimators, though the clipped Sparho estimator exhibited less of it.

D.2 EMPHATIC TD(λ) IN THE 5×5 GRID WORLD

Here we evaluate the substitution of the Sparho estimator in Emphatic TD(λ). Specifically, we evaluate EQ(λ) with importance sampling, Sparho, as well as their clipped versions, in the 5×5 Grid World environment with a tabular representation. We assume uniform interest $i(s) = 1, \forall s \in \mathcal{S}$, perform a sweep over step size α and trace-decay rate λ (detailed below), and run each instance for 100,000 steps. Results are averaged over 100 runs, and the best parameter settings among the sweep in terms of final RMS error are shown in Figure 9.

Parameter	Value(s)
Discount Factor γ	1.0
Action Space Size $ \mathcal{A} $	4
Action Commitment ϵ_π	0.5
Action Commitment ϵ_μ	1.0
Step Size α	0.0001, 0.0005, 0.001, 0.005, 0.01, 0.05, 0.1
Trace Decay λ	0.5, 0.75, 0.875, 0.9375, 0.96875

We observe that despite the relative simplicity of the problem (a small, structured grid world), Emphatic Q(λ) with importance sampling doesn't reliably learn. Substituting $\rho_t \rightarrow \tilde{\rho}_t$ improves performance, but slowly progresses due to the necessity of a small step size to handle the base algorithm's increased variance. Clipping both estimators enables substantially better learning progress, but notably much slower than their non-emphatic counterparts in Figure 6a.




Figure 9: Results in the 5×5 Grid World environment with a tabular representation for Emphatic $Q(\lambda)$ with various choices of importance weights. Results are averaged over 100 runs, and shaded regions represent one standard error.

E ALTERNATE OBJECTIVES

Below we provide importance weight derivations for alternate optimization objectives.

E.1 MINIMUM ℓ_2 DISTANCE TO 1

The following is the Sparho objective but without the constraint on the expected importance weight being 1. This can have lower variance importance weights which are often close to 1, but may arbitrarily attenuate or amplify a sequence in the multi-step case.

$$\mathcal{L}(\check{\rho}, \lambda) = \sum_a \mu_a (\check{\rho}_a - 1)^2 + \lambda \left(\sum_a \pi_a Q_a - \sum_a \mu_a \check{\rho}_a Q_a \right)$$

$$\mathcal{L}(\check{\rho}, \lambda) = (\check{\rho} - \mathbf{1})^T D_\mu (\check{\rho} - \mathbf{1}) + \lambda_0 (\boldsymbol{\pi}^T \mathbf{Q} - (D_\mu \mathbf{Q})^T \check{\rho})$$

$$\nabla \mathcal{L}(\check{\rho}, \lambda) = 0 = 2D_\mu \check{\rho} - 2\boldsymbol{\mu} - D_\mu \mathbf{Q} \lambda$$

$$0 = D_\mu \check{\rho} - \boldsymbol{\mu} - D_\mu \mathbf{Q} \frac{\lambda}{2}$$

$$0 = D_\mu^{-1} D_\mu \check{\rho} - D_\mu^{-1} \boldsymbol{\mu} - D_\mu^{-1} D_\mu \mathbf{Q} \frac{\lambda}{2}$$

$$0 = \check{\rho} - \mathbf{1} - \mathbf{Q} \frac{\lambda}{2}$$

$$\check{\rho} = \mathbf{1} + \mathbf{Q} \frac{\lambda}{2}$$

$$(D_\mu \mathbf{Q})^T \check{\rho} = (D_\mu \mathbf{Q})^T \mathbf{1} + (D_\mu \mathbf{Q})^T \mathbf{Q} \frac{\lambda}{2}$$

$$\boldsymbol{\pi}^T \mathbf{Q} = \boldsymbol{\mu}^T \mathbf{Q} + \boldsymbol{\mu}^T (\mathbf{Q} \circ \mathbf{Q}) \frac{\lambda}{2}$$

$$\frac{\lambda}{2} = \frac{\boldsymbol{\pi}^T \mathbf{Q} - \boldsymbol{\mu}^T \mathbf{Q}}{\boldsymbol{\mu}^T (\mathbf{Q} \circ \mathbf{Q})}$$

$$\check{\rho} = \mathbf{1} + \frac{\mathbf{Q}}{\boldsymbol{\mu}^T (\mathbf{Q} \circ \mathbf{Q})} (\boldsymbol{\pi}^T \mathbf{Q} - \boldsymbol{\mu}^T \mathbf{Q})$$

$$\check{\rho} = \mathbf{1} + \frac{\mathbf{Q}}{\mathbb{E}_\mu[Q^2]} (\mathbb{E}_\pi[Q] - \mathbb{E}_\mu[Q])$$

E.2 MINIMUM ℓ_2 DISTANCE TO c

This is a generalization of C.1 to have the weights tend toward a specific value c , but again without a hard constraint on the expected importance weight.

$$\mathcal{L}(\check{\rho}, \lambda) = \sum_a \mu_a (\check{\rho}_a - c)^2 + \lambda \left(\sum_a \pi_a Q_a - \sum_a \mu_a \check{\rho}_a Q_a \right)$$

$$\mathcal{L}(\check{\rho}, \lambda) = (\check{\rho} - c)^T D_\mu (\check{\rho} - c) + \lambda_0 (\boldsymbol{\pi}^T \mathbf{Q} - (D_\mu \mathbf{Q})^T \check{\rho})$$

$$\nabla \mathcal{L}(\check{\rho}, \lambda) = 0 = 2D_\mu \check{\rho} - 2c\boldsymbol{\mu} - D_\mu \mathbf{Q} \lambda$$

$$0 = D_\mu \check{\rho} - c\boldsymbol{\mu} - D_\mu \mathbf{Q} \frac{\lambda}{2}$$

$$0 = D_\mu^{-1} D_\mu \check{\rho} - D_\mu^{-1} c\boldsymbol{\mu} - D_\mu^{-1} D_\mu \mathbf{Q} \frac{\lambda}{2}$$

$$0 = \check{\rho} - c - \mathbf{Q} \frac{\lambda}{2}$$

$$\check{\rho} = c + \mathbf{Q} \frac{\lambda}{2}$$

$$(D_\mu \mathbf{Q})^T \check{\rho} = (D_\mu \mathbf{Q})^T c + (D_\mu \mathbf{Q})^T \mathbf{Q} \frac{\lambda}{2}$$

$$\boldsymbol{\pi}^T \mathbf{Q} = c\boldsymbol{\mu}^T \mathbf{Q} + \boldsymbol{\mu}^T (\mathbf{Q} \circ \mathbf{Q}) \frac{\lambda}{2}$$

$$\frac{\lambda}{2} = \frac{\boldsymbol{\pi}^T \mathbf{Q} - c\boldsymbol{\mu}^T \mathbf{Q}}{\boldsymbol{\mu}^T (\mathbf{Q} \circ \mathbf{Q})}$$

$$\check{\rho} = c + \frac{\mathbf{Q}}{\boldsymbol{\mu}^T (\mathbf{Q} \circ \mathbf{Q})} (\boldsymbol{\pi}^T \mathbf{Q} - c\boldsymbol{\mu}^T \mathbf{Q})$$

$$\check{\rho} = c + \frac{\mathbf{Q}}{\mathbb{E}_\mu[Q^2]} (\mathbb{E}_\pi[Q] - c\mathbb{E}_\mu[Q])$$

E.3 MINIMUM VARIANCE AROUND CONSTRAINED LENGTH c

This is a generalization of the minimum-variance Sparho objective, but with a variable expected importance weight magnitude. One possible application of this is to absorb the discount factor into the importance weights to simultaneously discount and correct the distribution of a sample.

$$\mathcal{L}(\check{\rho}, \lambda_0, \lambda_1) = \sum_a \mu_a (\check{\rho}_a - c)^2 + \lambda_0 \left(\sum_a \pi_a Q_a - \sum_a \mu_a \check{\rho}_a Q_a \right) + \lambda_1 \left(c - \sum_a \mu_a \check{\rho}_a \right)$$

$$\mathcal{L}(\check{\rho}, \lambda_0, \lambda_1) = (\check{\rho} - c)^T D_\mu (\check{\rho} - c) + \lambda_0 (\boldsymbol{\pi}^T \mathbf{Q} - (D_\mu \mathbf{Q})^T \check{\rho}) + \lambda_1 (c - (D_\mu \mathbf{1})^T \check{\rho})$$

$$\nabla \mathcal{L}(\check{\rho}, \lambda_0, \lambda_1) = 0 = 2D_\mu \check{\rho} - 2c\boldsymbol{\mu} - D_\mu \mathbf{Q} \lambda_0 - D_\mu \mathbf{1} \lambda_1$$

$$0 = D_\mu \check{\rho} - c\boldsymbol{\mu} - D_\mu \mathbf{Q} \frac{\lambda_0}{2} - D_\mu \mathbf{1} \frac{\lambda_1}{2}$$

$$0 = D_\mu^{-1} D_\mu \check{\rho} - D_\mu^{-1} c\boldsymbol{\mu} - D_\mu^{-1} D_\mu \mathbf{Q} \frac{\lambda_0}{2} - D_\mu^{-1} D_\mu \mathbf{1} \frac{\lambda_1}{2}$$

$$0 = \check{\rho} - c - \mathbf{Q} \frac{\lambda_0}{2} - \mathbf{1} \frac{\lambda_1}{2}$$

$$\check{\rho} = c + \mathbf{Q} \frac{\lambda_0}{2} + \mathbf{1} \frac{\lambda_1}{2}$$

$$\begin{aligned}
(D_\mu \mathbf{Q})^T \check{\rho} &= (D_\mu \mathbf{Q})^T \mathbf{c} + (D_\mu \mathbf{Q})^T \mathbf{Q} \frac{\lambda_0}{2} + (D_\mu \mathbf{Q})^T \mathbf{1} \frac{\lambda_1}{2} \\
(D_\mu \mathbf{Q})^T \check{\rho} - (D_\mu \mathbf{Q})^T \mathbf{c} &= (D_\mu \mathbf{Q})^T \mathbf{Q} \frac{\lambda_0}{2} + (D_\mu \mathbf{Q})^T \mathbf{1} \frac{\lambda_1}{2} \\
\boldsymbol{\pi}^T \mathbf{Q} - c \boldsymbol{\mu}^T \mathbf{Q} &= \boldsymbol{\mu}^T (\mathbf{Q} \circ \mathbf{Q}) \frac{\lambda_0}{2} + \boldsymbol{\mu}^T \mathbf{Q} \frac{\lambda_1}{2} \\
(D_\mu \mathbf{1})^T \check{\rho} &= (D_\mu \mathbf{1})^T \mathbf{c} + (D_\mu \mathbf{1})^T \mathbf{Q} \frac{\lambda_0}{2} + (D_\mu \mathbf{1})^T \mathbf{1} \frac{\lambda_1}{2} \\
(D_\mu \mathbf{1})^T \check{\rho} - (D_\mu \mathbf{1})^T \mathbf{c} &= (D_\mu \mathbf{1})^T \mathbf{Q} \frac{\lambda_0}{2} + (D_\mu \mathbf{1})^T \mathbf{1} \frac{\lambda_1}{2} \\
0 &= \boldsymbol{\mu}^T \mathbf{Q} \frac{\lambda_0}{2} + \frac{\lambda_1}{2} \\
\begin{bmatrix} \boldsymbol{\mu}^T (\mathbf{Q} \circ \mathbf{Q}) & \boldsymbol{\mu}^T \mathbf{Q} \\ \boldsymbol{\mu}^T \mathbf{Q} & 1 \end{bmatrix} \begin{bmatrix} \frac{\lambda_0}{2} \\ \frac{\lambda_1}{2} \end{bmatrix} &= \begin{bmatrix} \boldsymbol{\pi}^T \mathbf{Q} - c \boldsymbol{\mu}^T \mathbf{Q} \\ 0 \end{bmatrix} \\
\begin{bmatrix} \frac{\lambda_0}{2} \\ \frac{\lambda_1}{2} \end{bmatrix} &= \begin{bmatrix} \boldsymbol{\mu}^T (\mathbf{Q} \circ \mathbf{Q}) & \boldsymbol{\mu}^T \mathbf{Q} \\ \boldsymbol{\mu}^T \mathbf{Q} & 1 \end{bmatrix}^{-1} \begin{bmatrix} \boldsymbol{\pi}^T \mathbf{Q} - c \boldsymbol{\mu}^T \mathbf{Q} \\ 0 \end{bmatrix} \\
\frac{\lambda_0}{2} &= \frac{\boldsymbol{\pi}^T \mathbf{Q} - c \boldsymbol{\mu}^T \mathbf{Q}}{\boldsymbol{\mu}^T (\mathbf{Q} \circ \mathbf{Q}) - (\boldsymbol{\mu}^T \mathbf{Q})^2} \\
\frac{\lambda_1}{2} &= -\boldsymbol{\mu}^T \mathbf{Q} \frac{\boldsymbol{\pi}^T \mathbf{Q} - c \boldsymbol{\mu}^T \mathbf{Q}}{\boldsymbol{\mu}^T (\mathbf{Q} \circ \mathbf{Q}) - (\boldsymbol{\mu}^T \mathbf{Q})^2} \\
\check{\rho} &= \mathbf{c} + \mathbf{Q} \frac{\boldsymbol{\pi}^T \mathbf{Q} - c \boldsymbol{\mu}^T \mathbf{Q}}{\boldsymbol{\mu}^T (\mathbf{Q} \circ \mathbf{Q}) - (\boldsymbol{\mu}^T \mathbf{Q})^2} - \boldsymbol{\mu}^T \mathbf{Q} \frac{\boldsymbol{\pi}^T \mathbf{Q} - c \boldsymbol{\mu}^T \mathbf{Q}}{\boldsymbol{\mu}^T (\mathbf{Q} \circ \mathbf{Q}) - (\boldsymbol{\mu}^T \mathbf{Q})^2} \\
\check{\rho} &= \mathbf{c} + \frac{\mathbf{Q} - \boldsymbol{\mu}^T \mathbf{Q}}{\boldsymbol{\mu}^T (\mathbf{Q} \circ \mathbf{Q}) - (\boldsymbol{\mu}^T \mathbf{Q})^2} (\boldsymbol{\pi}^T \mathbf{Q} - c \boldsymbol{\mu}^T \mathbf{Q}) \\
\check{\rho} &= \mathbf{c} + \frac{\mathbf{Q} - \mathbb{E}_\mu[Q]}{\mathbb{E}_\mu[Q^2] - \mathbb{E}_\mu[Q]^2} (\mathbb{E}_\pi[Q] - c \mathbb{E}_\mu[Q])
\end{aligned}$$

E.4 MINIMUM VARIANCE OF $\check{\rho}Q$

This objective aims to minimize the variance of $\check{\rho}Q$. Such an objective has intuitive appeal in minimizing the variance of the full update target, as opposed to just the importance weight. However, despite the focus on value-awareness, such an objective may be too reliant on the specific value estimates, and very sensitive to deviations in them. This may lead to substantially more bias, and even an increase in variance should the value estimates not accurately reflect the return.

Let $V = \mathbb{E}_\pi[Q] = \sum_a \pi_a Q_a$:

$$\mathcal{L}(\check{\rho}, \lambda_0, \lambda_1) = \sum_a \mu_a (\check{\rho}_a Q_a - V)^2 + \lambda_0 \left(V - \sum_a \mu_a \check{\rho}_a Q_a \right) + \lambda_1 \left(1 - \sum_a \mu_a \check{\rho}_a \right)$$

$$\mathcal{L}(\check{\rho}, \lambda_0, \lambda_1) = (D_Q \check{\rho} - V)^T D_\mu (D_Q \check{\rho} - V) + \lambda_0 (V - (D_\mu \mathbf{Q})^T \check{\rho}) + \lambda_1 (1 - (D_\mu \mathbf{1})^T \check{\rho})$$

$$\nabla \mathcal{L}(\check{\rho}, \lambda_0, \lambda_1) = 0 = 2D_\mu D_Q D_Q \check{\rho} - 2D_\mu \mathbf{Q} V - D_\mu \mathbf{Q} \lambda_0 - D_\mu \mathbf{1} \lambda_1$$

$$0 = D_\mu D_Q D_Q \check{\rho} - D_\mu \mathbf{Q} V - D_\mu \mathbf{Q} \frac{\lambda_0}{2} - D_\mu \mathbf{1} \frac{\lambda_1}{2}$$

$$0 = D_Q^{-1} D_Q^{-1} D_\mu^{-1} D_\mu D_Q D_Q \check{\rho} - D_Q^{-1} D_Q^{-1} D_\mu^{-1} D_\mu \mathbf{Q} V - D_Q^{-1} D_Q^{-1} D_\mu^{-1} D_\mu \mathbf{Q} \frac{\lambda_0}{2} - D_Q^{-1} D_Q^{-1} D_\mu^{-1} D_\mu \mathbf{1} \frac{\lambda_1}{2}$$

$$0 = \check{\rho} - D_Q^{-1} \mathbf{1} V - D_Q^{-1} \mathbf{1} \frac{\lambda_0}{2} - D_Q^{-1} D_Q^{-1} \mathbf{1} \frac{\lambda_1}{2}$$

$$\check{\rho} = D_Q^{-1} \mathbf{1} V + D_Q^{-1} \mathbf{1} \frac{\lambda_0}{2} + D_Q^{-1} D_Q^{-1} \mathbf{1} \frac{\lambda_1}{2}$$

$$\begin{aligned}
(D_\mu \mathbf{Q})^T \check{\rho} &= (D_\mu \mathbf{Q})^T D_Q^{-1} \mathbf{1} V + (D_\mu \mathbf{Q})^T D_Q^{-1} \mathbf{1} \frac{\lambda_0}{2} + (D_\mu \mathbf{Q})^T D_Q^{-1} D_Q^{-1} \mathbf{1} \frac{\lambda_1}{2} \\
(D_\mu \mathbf{Q})^T \check{\rho} - (D_\mu \mathbf{Q})^T D_Q^{-1} \mathbf{1} V &= (D_\mu \mathbf{Q})^T D_Q^{-1} \mathbf{1} \frac{\lambda_0}{2} + (D_\mu \mathbf{Q})^T D_Q^{-1} D_Q^{-1} \mathbf{1} \frac{\lambda_1}{2} \\
0 &= \frac{\lambda_0}{2} + \boldsymbol{\mu}^T D_\mu^{-1} \mathbf{1} \frac{\lambda_1}{2}
\end{aligned}$$

$$\begin{aligned}
(D_\mu \mathbf{1})^T \check{\rho} &= (D_\mu \mathbf{1})^T D_Q^{-1} \mathbf{1} V + (D_\mu \mathbf{1})^T D_Q^{-1} \mathbf{1} \frac{\lambda_0}{2} + (D_\mu \mathbf{1})^T D_Q^{-1} D_Q^{-1} \mathbf{1} \frac{\lambda_1}{2} \\
(D_\mu \mathbf{1})^T \check{\rho} - (D_\mu \mathbf{1})^T D_Q^{-1} \mathbf{1} V &= (D_\mu \mathbf{1})^T D_Q^{-1} \mathbf{1} \frac{\lambda_0}{2} + (D_\mu \mathbf{1})^T D_Q^{-1} D_Q^{-1} \mathbf{1} \frac{\lambda_1}{2} \\
1 - \boldsymbol{\mu}^T D_Q^{-1} \mathbf{1} V &= \boldsymbol{\mu}^T D_Q^{-1} \mathbf{1} \frac{\lambda_0}{2} + \boldsymbol{\mu}^T D_Q^{-1} D_Q^{-1} \mathbf{1} \frac{\lambda_1}{2}
\end{aligned}$$

$$\begin{aligned}
\begin{bmatrix} 1 & \boldsymbol{\mu}^T D_Q^{-1} \mathbf{1} \\ \boldsymbol{\mu}^T D_Q^{-1} \mathbf{1} & \boldsymbol{\mu}^T D_Q^{-1} D_Q^{-1} \mathbf{1} \end{bmatrix} \begin{bmatrix} \frac{\lambda_0}{2} \\ \frac{\lambda_1}{2} \end{bmatrix} &= \begin{bmatrix} 0 \\ 1 - \boldsymbol{\mu}^T D_Q^{-1} \mathbf{1} V \end{bmatrix} \\
\begin{bmatrix} \frac{\lambda_0}{2} \\ \frac{\lambda_1}{2} \end{bmatrix} &= \begin{bmatrix} 1 & \boldsymbol{\mu}^T D_Q^{-1} \mathbf{1} \\ \boldsymbol{\mu}^T D_Q^{-1} \mathbf{1} & \boldsymbol{\mu}^T D_Q^{-1} D_Q^{-1} \mathbf{1} \end{bmatrix}^{-1} \begin{bmatrix} 0 \\ 1 - \boldsymbol{\mu}^T D_Q^{-1} \mathbf{1} V \end{bmatrix} \\
\frac{\lambda_0}{2} &= \frac{(\boldsymbol{\mu}^T D_Q^{-1} \mathbf{1})^2 V - \boldsymbol{\mu}^T D_Q^{-1} \mathbf{1}}{\boldsymbol{\mu}^T D_Q^{-1} D_Q^{-1} \mathbf{1} - (\boldsymbol{\mu}^T D_Q^{-1} \mathbf{1})^2} \\
\frac{\lambda_1}{2} &= \frac{1 - \boldsymbol{\mu}^T D_Q^{-1} \mathbf{1} V}{\boldsymbol{\mu}^T D_Q^{-1} D_Q^{-1} \mathbf{1} - (\boldsymbol{\mu}^T D_Q^{-1} \mathbf{1})^2}
\end{aligned}$$

$$\begin{aligned}
\check{\rho} &= D_Q^{-1} \mathbf{1} V + D_Q^{-1} \mathbf{1} \frac{(\boldsymbol{\mu}^T D_Q^{-1} \mathbf{1})^2 V - \boldsymbol{\mu}^T D_Q^{-1} \mathbf{1}}{\boldsymbol{\mu}^T D_Q^{-1} D_Q^{-1} \mathbf{1} - (\boldsymbol{\mu}^T D_Q^{-1} \mathbf{1})^2} + D_Q^{-1} D_Q^{-1} \mathbf{1} \frac{1 - \boldsymbol{\mu}^T D_Q^{-1} \mathbf{1} V}{\boldsymbol{\mu}^T D_Q^{-1} D_Q^{-1} \mathbf{1} - (\boldsymbol{\mu}^T D_Q^{-1} \mathbf{1})^2} \\
\check{\rho}_a &= \frac{V}{Q_a} + \frac{\frac{1}{Q_a} (\mathbb{E}_\mu[\frac{1}{Q}]^2 V - \mathbb{E}_\mu[\frac{1}{Q}]) + \frac{1}{Q_a^2} (1 - \mathbb{E}_\mu[\frac{1}{Q}] V)}{\mathbb{E}_\mu[\frac{1}{Q^2}] - \mathbb{E}_\mu[\frac{1}{Q}]^2}
\end{aligned}$$