
An Alternate Policy Gradient Estimator for Softmax Policies

Shivam Garg¹ Samuele Tosatto¹ Yangchen Pan² Martha White^{1*} A. Rupam Mahmood^{1*}

¹University of Alberta. ²Noah’s Ark Lab, Huawei (work done while at UofA).

*CIFAR AI Chair, Alberta Machine Intelligence Institute (Amii).

Abstract

Policy gradient (PG) estimators are ineffective in dealing with softmax policies that are sub-optimally saturated, which refers to the situation when the policy concentrates its probability mass on sub-optimal actions. Sub-optimal policy saturation may arise from bad policy initialization or sudden changes in the environment that occur after the policy has already converged. Current softmax PG estimators require a large number of updates to overcome policy saturation, which causes low sample efficiency and poor adaptability to new situations. To mitigate this problem, we propose a novel PG estimator for softmax policies that utilizes the bias in the critic estimate and the noise present in the reward signal to escape the saturated regions of the policy parameter space. Our theoretical analysis and experiments, conducted on bandits and various reinforcement learning environments, show that this new estimator is significantly more robust to policy saturation.

1 INTRODUCTION

Policy gradient (PG) algorithms aim to optimize a sequential decision making problem defined on a set of parameterized policies: the policy parameters are optimized using standard stochastic gradient-based update to maximize the net reward received by the agent. A salient feature of the PG methods is that they can be used with both discrete and continuous action spaces, and have been successfully deployed in a variety of real world tasks such as robotics (Haarnoja et al., 2018; Mahmood, et al., 2018) and large scale simulated problems (Berner et al., 2019).

Proceedings of the 25th International Conference on Artificial Intelligence and Statistics (AISTATS) 2022, Valencia, Spain. PMLR: Volume 151. Copyright 2022 by the author(s).

When the action-space is discrete, policies are typically represented using a categorical distribution parameterized by a softmax function. However, softmax policies have some inherent issues. Even with access to the true gradients, they can be slow in responding to non-stationarity (Hennes et al., 2019) and their performance heavily depends on the initialization of the policy parameters (Li et al., 2021; Mei et al. 2020a; Mei et al. 2020b). Both these problems arise from an issue, which we call *sub-optimal policy saturation*.

Sub-optimal policy saturation refers to the situation when the policy places a high probability mass on sub-optimal actions. An agent with a saturated policy will not be able to explore other actions and may continue to remain in the sub-optimal region if an appropriate measure is not taken. Sub-optimal policy saturation is a significant issue that arises in multiple scenarios: (1) Non-stationarity: when the agent is acting in a constantly changing environment, and what was once an optimal strategy may no longer work well. (2) Pre-training / transfer learning: deep reinforcement learning systems are often pre-trained on different tasks before being used on the main task. The subtle differences in the structure of these tasks might lead to a bad policy initialization which is sub-optimally saturated. And (3) Stochastic updates: policy gradient updates suffer from high variance (Williams, 1992) and therefore, it is possible for a policy to become saturated on sub-optimal actions during the course of learning.

PG methods with softmax policies are particularly susceptible to the issue of saturation. Prior work has suggested different approaches to mitigate this issue: entropic regularization (Ding et al., 2021; Ahmed et al., 2019; Haarnoja et al., 2018; Peters et al., 2010), natural PG algorithms (Kakade, 2001; Peters and Schaal, 2008), and log barrier regularization (Agarwal et al., 2021) can all make the policy more explorative. Explicit reward bonuses can also encourage the exploration of low-density state-action pairs (Auer et al. 2002). However, entropic regularization and explorative bonuses introduce additional terms in the objective, and therefore the resulting optimal policy can

be different from the original one. Another solution is to replace the softmax with a function such as the escort transform (Mei et al., 2020b). In this article, we take a different approach to address this issue. Instead of augmenting the optimization objective with bonuses or regularizers, we introduce a PG estimator that inherently helps to escape the sub-optimally saturated regions. Furthermore, the above augmentations can still be applied on top of our estimator.

Our proposed estimator is a simple yet effective approach for dealing with sub-optimally saturated policies thereby making PG algorithms more robust. The classic likelihood ratio estimator for softmax policies, which we call the *regular estimator*, takes a frustratingly large amount of experience to escape sub-optimally saturated policy regions. The regular estimator produces near zero gradients at saturation as both its expectation and variance, even with reward noise, are vanishingly small at those regions. In contrast, our proposed estimator, which we call the *alternate estimator*, has a non-zero variance in the same scenario that can be utilized to escape the sub-optimal regions. Further, the alternate estimator naturally utilizes the bias in the critic estimate to increase the policy’s entropy, thereby encouraging exploration. As the critic estimate improves, this effect reduces, allowing the policy to saturate towards the optimal actions.

Contributions: We present the alternate PG estimator for softmax policies and show that it works by utilizing reward noise as well as the bias in the critic to escape saturated regions in the policy space. We theoretically justify this behavior in the bandit setting complemented by simple and intuitive examples. We also provide empirical verification¹ of these properties on (1) problems with artificially saturated policies and (2) on non-stationary tasks where sub-optimal policy saturation naturally arises. Our results show that this new estimator works well in both the bandit and the MDP settings with various PG algorithms such as REINFORCE and Actor-Critic (Sutton et al., 2000), and is superior (or competitive) to the regular estimator under tabular, linear, and neural representations.

2 PRELIMINARIES

In reinforcement learning, the decision making task is described using Markov decision processes (MDPs) $\mathcal{M} := (\mathcal{S}, \mathcal{A}, \mathcal{R}, \mu, p, \gamma)$, where \mathcal{S} is the set of states, \mathcal{A} is the discrete set of actions, $\mathcal{R} \subset \mathbb{R}$ is the set of rewards, $\mu \in \Delta(\mathcal{S})^2$ is the start state distribution, $p :$

$\mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S} \times \mathcal{R})$ is the transition dynamics, and $\gamma \in [0, 1]$ is the discount factor. The agent maintains a policy $\pi : \mathcal{S} \rightarrow \Delta(\mathcal{A})$ that describes its interaction with the environment. This interaction is represented using the history $H_T = \{S_0, A_0, R_1, S_1, A_1, \dots, R_T, S_T\}$ where $S_0 \sim \mu$, $A_t \sim \pi(\cdot|S_t)$, and $S_{t+1}, R_{t+1} \sim p(\cdot, \cdot|S_t, A_t)$ for $t \in \{0, \dots, T-1\}$ ³. The agent uses such interactions to learn a policy that maximizes the expected return $\mathcal{J} = \mathbb{E}_\pi[\sum_0^{T-1} \gamma^t R_{t+1}]$. PG methods provide one way to accomplish this task. We now describe these methods in two typical settings.

Gradient Bandits

Discrete bandits are a simplification of the full MDP framework. They have no states and at each timestep the agent picks an action A_t from a discrete action set \mathcal{A} and obtains a reward $R_t \sim p(\cdot|A_t)$. In this setting, the agent’s goal reduces to maximizing the expected immediate reward $\mathcal{J} := \mathbb{E}_\pi[R_t] =: r_\pi$. Gradient bandit algorithms maximize \mathcal{J} using gradient ascent. Given a softmax policy $\pi_\theta(a) = e^{\theta_a} / \sum_{b \in \mathcal{A}} e^{\theta_b}$, where θ_a is the action preference for action a , the gradient $\nabla \mathcal{J}$ (§2.8, Sutton and Barto, 2018) is given by

$$[\nabla_{\theta} \mathcal{J}]_i = r(a_i) \cdot \pi(a_i) \cdot (1 - \pi(a_i)), \quad (1)$$

where $r(a) := \mathbb{E}[R|a]$ is the reward function. However, the agent usually does not have access to r for all actions at the same time and must resort to using sample based gradient estimators $\hat{\mathbf{g}}$ such that $\nabla_{\theta} \mathcal{J} = \mathbb{E}_{A \sim \pi; R \sim p(\cdot|A)}[\hat{\mathbf{g}}(A, R)]$. A typical choice is to use what we call the *regular estimator*

$$[\hat{\mathbf{g}}^{\text{REG}}(A, R)]_a = (R - b)(\mathbb{I}(A = a) - \pi(a)), \quad (2)$$

where \mathbb{I} is the indicator function and b is an estimate of the average reward r_π .

Policy Gradient Methods

In an episodic MDP, the PG objective is $\mathcal{J} = \mathbb{E}_\pi[\sum_0^{T-1} \gamma^t R_{t+1}]$. The policy gradient theorem (Sutton et al., 2000) states that

$$\nabla \mathcal{J} = \sum_{s \in \mathcal{S}} \nu_\pi(s) \sum_{a \in \mathcal{A}} \nabla \pi(a|s) q_\pi(s, a), \quad (3)$$

where $\nu_\pi(s) := \sum_{k=0}^{\infty} \gamma^k \mathbb{P}(S_k = s)$ is the γ -discounted state occupancy measure under policy π . Eq. 3 can also be written as an expectation,

$$\nabla \mathcal{J} = \mathbb{E}_{\nu_\pi, \pi} \left[\underbrace{\nabla \log \pi(A|S) (q_\pi(S, A) - v_\pi(S))}_{=: \mathbf{g}^{\text{REG}}(S, A)} \right]. \quad (4)$$

¹Code for the experiments is available at https://github.com/svmgrg/alternate_pg.

²The object $\Delta(\mathcal{X})$ denotes the set of all possible probability distributions over the set \mathcal{X} .

³In this article, we only consider episodic tasks with the random variable T denoting the episode termination length. However, all the arguments presented here can be extended for continuing tasks as well.

This expression gives us the sample based gradient estimator \mathbf{g}^{REG} . Again, the agent will use $\hat{\mathbf{g}}^{\text{REG}}(S, A)$ with value function estimates \hat{q}_π and \hat{v}_π .

3 ALTERNATE GRADIENT BANDITS

In this section, we introduce and analyze the alternate PG estimator. We begin by describing the vector notation which makes our calculations succinct and easier to follow. For a k -armed bandit problem, let $\boldsymbol{\pi}, \mathbf{r} \in \mathbb{R}^k$ denote the policy and the reward vectors with $[\boldsymbol{\pi}]_k = \pi(a_k)$ and $[\mathbf{r}]_k = r(a_k)$. Using these vectors, the expectations of scalar quantities become vector inner products: $\mathcal{J}_\pi \equiv r_\pi = \mathbb{E}_\pi[R] = \sum_a \pi(a)r(a) = \boldsymbol{\pi}^\top \mathbf{r}$. Further, the softmax policy can be written as $\boldsymbol{\pi} = e^\theta / (\mathbf{1}^\top e^\theta)$, where the vector $e^\theta \in \mathbb{R}^k$ is defined as $[e^\theta]_a := e^{\theta_a}$. And $\nabla_\theta \boldsymbol{\pi} = (\mathbf{I} - \boldsymbol{\pi} \mathbf{1}^\top) \text{diag}(\boldsymbol{\pi})$. Hence, the regular policy gradient can be derived as follows:

$$\begin{aligned} \nabla_\theta \mathcal{J} &= \nabla_\theta (\boldsymbol{\pi}^\top \mathbf{r}) = (\mathbf{I} - \boldsymbol{\pi} \mathbf{1}^\top) \text{diag}(\boldsymbol{\pi}) \mathbf{r} \\ &= \mathbb{E}_{A \sim \pi} [r(A) (\mathbf{e}_A - \boldsymbol{\pi})] \\ &= \mathbb{E}_{A \sim \pi, R \sim p(\cdot|A)} \underbrace{[(R - r_\pi) (\mathbf{e}_A - \boldsymbol{\pi})]}_{=: \mathbf{g}^{\text{REG}}(A, R)}, \end{aligned} \quad (5)$$

where $\mathbf{e}_A \in \mathbb{R}^k$ represents the basis vector with the A th element equal to one. The approximate version of it, as given before in Eq. 2, is:

$$\hat{\mathbf{g}}^{\text{REG}}(A, R) = (R - b) (\mathbf{e}_A - \boldsymbol{\pi}), \quad (6)$$

where b is an estimate of r_π . Also note that this estimator remains unbiased: $\nabla \mathcal{J} = \mathbb{E}[\hat{\mathbf{g}}^{\text{REG}}(A, R)]$.

The Alternate Gradient Bandit Estimator is derived by using an alternate form of the true gradient:

$$\begin{aligned} \nabla_\theta \mathcal{J} &= (\mathbf{I} - \boldsymbol{\pi} \mathbf{1}^\top) \text{diag}(\boldsymbol{\pi}) \mathbf{r} = (\text{diag}(\mathbf{r}) - \boldsymbol{\pi} \mathbf{r}^\top) \boldsymbol{\pi} \\ &= \text{diag}(\mathbf{r}) \boldsymbol{\pi} - \boldsymbol{\pi}^\top \mathbf{r} \boldsymbol{\pi} = (\text{diag}(\mathbf{r}) - \boldsymbol{\pi}^\top \mathbf{r} \mathbf{I}) \boldsymbol{\pi} \\ &= \mathbb{E}_{A \sim \pi, R \sim p(\cdot|A)} [(R - r_\pi) \mathbf{e}_A], \end{aligned} \quad (7)$$

where $\mathbf{g}^{\text{ALT}}(A, R) := (R - r_\pi) \mathbf{e}_A$ is the *alternate* gradient bandit estimator. Using a baseline b to estimate the average reward r_π , gives the following estimator:

$$\hat{\mathbf{g}}^{\text{ALT}}(A, R) = (R - b) \mathbf{e}_A. \quad (8)$$

The alternate estimator is faster per-timestep than the regular estimator: At each time-step, the regular estimator updates all the components of the action preference vector $\boldsymbol{\theta}$, whereas the alternate estimator only updates a single component. Due to this, a gradient bandit algorithm employing the regular estimator will have a per-timestep computational complexity of $\mathcal{O}(|\mathcal{A}|)$. However, as we show in §B.1, gradient bandits utilizing the alternate estimator can use

a tree based sampling procedure to have a per-time-step complexity of $\mathcal{O}(\log |\mathcal{A}|)$. When the action space is large, this difference becomes substantial. Now we discuss how the alternate estimator utilizes the reward noise and the bias in the critic estimate to escape saturated regions in the policy space. Proofs of the results are given in §B.2 and §B.3.

Alternate Estimator Utilizes Reward Noise

Whenever the policy is saturated, i.e. it places a high probability mass on some actions, the expected gradient becomes close to zero, and consequently the policy weights are not updated. Further, for the regular estimator, the variance at these ‘‘corners’’ of the probability simplex is also zero. Therefore, there is neither any gradient signal nor any noise, and the agent is unable to escape the sub-optimal region. In contrast, the alternate estimator, despite being zero in expectation, has a non-zero variance and therefore the agent can utilize the reward noise to escape the sub-optimal region. We illustrate this point with Figure 1 (left and middle) and an accompanying example:

Example 1. Consider a 3-armed bandit with $\mathcal{A} = \{a_0, a_1, a_2\}$, the expected reward $\mathbf{r} = [0 \ 0 \ 1]^\top$, and the policy $\boldsymbol{\pi} = [1 \ 0 \ 0]^\top$. The rewards are perturbed with a normally distributed noise $\epsilon \sim \mathcal{N}(0, \sigma^2)$, so that upon picking action A , the agent receives the reward $r(A) + \epsilon$. Since $r_\pi = 0$, $R - r_\pi = 0 + \epsilon - 0 = \epsilon$. And because the agent samples $A = a_0$ at each timestep, we obtain

$$\begin{aligned} \mathbf{g}^{\text{REG}}(A, R) &= (R - r_\pi) (\mathbf{e}_A - \boldsymbol{\pi}) \\ &= \epsilon \cdot ([1 \ 0 \ 0]^\top - [1 \ 0 \ 0]^\top) = [0 \ 0 \ 0]^\top, \text{ and} \\ \mathbf{g}^{\text{ALT}}(A, R) &= (R - r_\pi) \mathbf{e}_A = \epsilon \cdot [1 \ 0 \ 0]^\top = [\epsilon \ 0 \ 0]^\top. \end{aligned}$$

This example shows that at a sub-optimal corner, the agent with the alternate estimator effectively does a random walk until it escapes that region and can choose other actions to get a non-zero reward signal. Further, doing a random walk at the optimal corner is not problematic, since there is an attractive gradient signal as soon as the agent moves away from this corner. The following propositions formalize these points.

Proposition 1. Define $\mathcal{I}_c := \{a \mid r(a) = c\}$ for some constant $c \in \mathbb{R}$. Assume that $\exists c$ such that $\mathcal{I}_c \neq \emptyset$ and that the policy is saturated on the actions in the set \mathcal{I}_c : $\sum_{a \in \mathcal{I}_c} \pi(a) = 1$. Then the expected policy gradient for softmax policies is zero: $\nabla_\theta \mathcal{J} = \mathbb{E}[\mathbf{g}^{\text{REG}}(A, R)] = \mathbb{E}[\hat{\mathbf{g}}^{\text{REG}}(A, R)] = \mathbb{E}[\mathbf{g}^{\text{ALT}}(A, R)] = \mathbf{0}$.

Proposition 2. Let $\sigma(a)^2 := \mathbb{V}[R|A = a]$ denote the variance⁴ of the reward corresponding to action a . As

⁴We compute the variance of vector random variables elementwise; for details, see §A.




Figure 1: Policy updates and the variance for PG estimators on 3-armed bandits plotted on the probability simplex. The policy-update plots were drawn by updating the policy using the corresponding stochastic PG estimator, assuming that action a_0 was taken and a noise of constant magnitude was added to the reward. The blue and grey arrows correspond to the policy change direction for $+ve$ and $-ve$ reward noises. The variance plots show the heatmap for the zeroth element of the gradient estimator. Last three plots show policy updates using expectation of the estimators. We include more of these plots for other settings in §B.5.

sume that the policy is saturated on the action c , i.e., $\pi(c) = 1$. Then, the variance of the regular PG estimator (with or without the baseline subtraction) is zero: $\mathbb{V}[\mathbf{g}^{REG}(A)] = \mathbb{V}[\hat{\mathbf{g}}^{REG}(A, R)] = \mathbf{0}$. Whereas, the variance of the alternate PG estimator is non-zero: $\mathbb{V}[\mathbf{g}^{ALT}(A, R)] = \sigma(c)^2 \mathbf{e}_c$.

Note that the example and the propositions presented above require the policy to lie at the boundary of the probability simplex which is unsatisfiable for softmax policies, as this would require some of the action preferences to be infinitely large. Despite that, using continuity arguments, we can still reason that for the regular estimator, both the expected gradient and its variance vanish in the proximity of the simplex boundary, whereas the alternate estimator will have non-zero variance⁵. In addition to this, the stochastic update for the alternate estimator is much higher than that for the regular estimator. This can be seen from Figure 1 (left): look at the length of the update arrows near the bottom-right corner on the simplex.

Alternate Estimator Utilizes its Biasedness

Apart from utilizing the reward noise, the alternate estimator can also utilize the bias in the critic estimate b . To see this, consider a baseline $b \neq r_\pi$. Then the alternate estimator becomes biased:

$$\nabla_{\theta} \mathcal{J} \neq \mathbb{E}[\hat{\mathbf{g}}^{ALT}(A, R)] = \mathbb{E}[(R - b) \mathbf{e}_A] = \boldsymbol{\pi} \odot (\mathbf{r} - b\mathbf{1}).$$

However, this biased update $\mathbb{E}[\hat{\mathbf{g}}^{ALT}(A, R)]$ is not the gradient of any function (refer §B.3). Therefore, in order to understand its behavior, we look at its fixed point and under what conditions it acts as an attractor

⁵In practice, PG methods are often used with stochastic gradients averaged over mini-batches of data. An increase in batch-size would reduce the variance of alternate, thereby decreasing the random walk effect. However, the biasedness of the alternate estimator is independent of the batch-size and can still help in escaping policy saturation.

or a repeller. Figure 1 (right) illustrates this fixed point and its behavior based on how the baseline is initialized for one specific instance of a bandit problem.

We now state a theorem that formalizes this property. Let $r_1 \leq r_2 \leq \dots \leq r_k$ represent an ordering of the true rewards for the bandit problem. Let $\pi_t = e^{\boldsymbol{\theta}^{(t)}} / \mathbf{1}^\top e^{\boldsymbol{\theta}^{(t)}}$ represent the softmax policy at timestep t . Here, $\boldsymbol{\theta}^{(t)}$ represents action preference vector at timestep t , which is updated using the rule $\boldsymbol{\theta}^{(t+1)} = \boldsymbol{\theta}^{(t)} + \alpha \mathbb{E}[\hat{\mathbf{g}}^{ALT}(A, R)]$ for $\alpha > 0$.

Lemma 2.1. Fixed points of the biased gradient bandit update. (1) If there exists an $n \in \{1, 2, \dots, k-1\}$ such that $r_1 \leq \dots \leq r_n < b < r_{n+1} \leq \dots \leq r_k$, then $\mathbb{E}_\pi[\hat{\mathbf{g}}^{ALT}(A, R)]$ is never equal to zero. (2) If $b = r(a)$ for at least one action, then $\mathbb{E}_\pi[\hat{\mathbf{g}}^{ALT}(A, R)] = 0$ at any point on the face of the probability simplex given by $\sum_{a \in \mathcal{I}_b} \pi(a) = 1$ with $\mathcal{I}_b := \{a | r(a) = b\}$. (3) If $b < r_1$ or $b > r_k$, then $\mathbb{E}_\pi[\hat{\mathbf{g}}^{ALT}(A, R)] = 0$ at a point π^* within the simplex boundary given by $\pi^*(a) = \frac{1}{r(a)-b} \left(\sum_{c \in A} \frac{1}{r(c)-b} \right)^{-1}$, $\forall a \in A$.

Theorem 1. Nature of the fixed point π^* . Assume that $\pi_t \neq \pi^*$. If the baseline is pessimistic, i.e. $b < r_1$, then for any $\alpha > 0$, the fixed point π^* acts as a repeller: $D_{KL}(\pi^* || \pi_{t+1}) > D_{KL}(\pi^* || \pi_t)$, where D_{KL} denotes the KL-divergence. And if the baseline is optimistic, i.e. $b > r_k$, then given a sufficiently small positive stepsize α the fixed point π^* acts as an attractor: $D_{KL}(\pi^* || \pi_{t+1}) < D_{KL}(\pi^* || \pi_t)$.

The above theorem illustrates an important property of the alternate estimator: with an optimistically initialized baseline, the agent is updated towards a more uniform distribution π^* . And for an optimistic baseline, this distribution has a higher probability of picking the action with the maximum reward as compared to other actions; this can be seen from Figure 1 (right) and the expression for π^* in Lemma 2.1. Therefore,

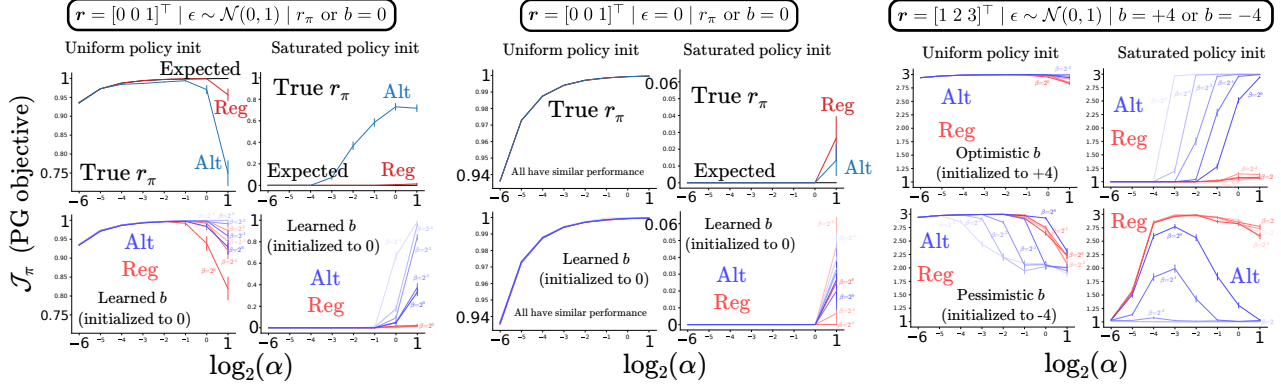


Figure 2: Parameter sensitivity plots for different PG estimators on 3-armed bandit problems. The headers give the specific experimental settings. The algorithms were run for 1000 timesteps and the plots show the mean performance during the last 50 timesteps, averaged over 150 runs.

if the agent were stuck in a sub-optimal corner of the probability simplex, an optimistic baseline would make its policy more uniform and encourage exploration⁶. And even though π^* is different from the optimal policy, the agent with an alternate estimator can still reach the optimal policy, because as the agent learns and improves its baseline estimate, the alternate PG estimator becomes asymptotically unbiased. And hopefully by this time, the agent has already escaped the saturated policy region.

4 BANDIT EXPERIMENTS

In this section, we present experiments with the gradient bandit algorithm on 3-armed bandit tasks, with a fixed reward structure \mathbf{r} and a normally distributed reward noise ϵ . The experimental results serve as a demonstration and the verification of the properties of the alternate and the regular estimators discussed above. We trained five different agents on the bandit task. The policy weights were updated using either the expected gradient $\nabla \mathcal{J} = \pi \odot (\mathbf{r} - r_\pi)$, the regular estimator with true r_π (Eq. 5) or a learned baseline (Eq. 6), or the alternate estimator with true r_π (Eq. 7) or a learned baseline (Eq. 8). The baseline was learned using a running average: $b_{t+1} = (1 - \beta)b_t + \beta R_t$. Additional details are given in §D.

Experiment 1 demonstrates that the alternate estimator performs competitively with the regular estimator for uniform policy initialization and clearly outperforms it in the case of a saturated policy initialization. We set $\mathbf{r} = [0\ 0\ 1]^T$ with noise $\epsilon \sim \mathcal{N}(0, 1)$. The action preferences were initialized to $\theta_a = 0$, $\forall a$ for uniform policy, and to $\theta_{a_0} = 10$ and $\theta_{a_1} = \theta_{a_2} = 0$ for the sat-

⁶On the flip side, with a pessimistically initialized baseline, the alternate estimator can pre-maturely saturate towards a sub-optimal corner; see Figure 1 (right).




Figure 3: Learning curves for the gradient bandit algorithm. The results were averaged over 150 independent runs; shaded region shows standard error.

urated policy. The baseline b was initialized to zero. Figure 3 shows the learning curves for the best parameter configuration. Figure 2 (left) shows the sensitivity plots; in particular it shows the final performance of the methods for different parameter settings.

For saturated policy initialization, we observe that the alternate estimators (both r_π and a learned b) learned good behaviors, whereas the expected gradient and the regular estimators failed to learn anything. Further, alternate with a learned baseline converged faster than alternate with r_π . We attribute this result to a better exploration afforded by the baseline.

Experiment 2 retained the previous experimental setting, but set the reward noise $\epsilon = 0$. Figure 2 (middle) shows the sensitivity plots for both uniform and saturated policy initializations. These results show that for uniform initialization, all the methods had a similar performance. For the saturated initializations, none of the methods were able to learn a good policy. This experiment verifies our claim that reward noise helps the alternate estimator to escape the sub-optimal regions in the policy space; without reward noise, the agent is unable to perform a random walk.

Experiment 3 studies the effect of initializing the reward baseline optimistically ($b = +4$) or pessimistically ($b = -4$). The reward was set to $\mathbf{r} = [1 \ 2 \ 3]^\top$ with noise $\epsilon \sim \mathcal{N}(0, 1)$. Figure 2 (right) shows the sensitivity plots for both the estimators. From the plots, we see that optimistic initialization greatly helped the alternate estimator in escaping the saturated policy. Whereas, a pessimistic initialization significantly hampered the performance of the alternate estimator, making it worse than the regular estimator even for the uniform policy case. In contrast, the baseline initialization affected the regular estimator in a milder way.

5 ALTERNATE PG ESTIMATOR FOR MDPs

The alternate estimator readily extends to MDPs, where it enjoys properties analogous to the bandit case. We start by deriving the regular estimator for MDPs using the vector notation. Given a state s , let $\boldsymbol{\pi}(\cdot|s) \in \mathbb{R}^{|\mathcal{A}|}$ be the policy vector with $[\boldsymbol{\pi}(\cdot|s)]_a = \pi(a|s)$ and $[\mathbf{q}_\pi(s, \cdot)]_a = q_\pi(s, a)$ be the action value vector. The softmax policy for MDPs is $\pi(a|s) = e^{[\boldsymbol{\theta}_\mathbf{w}(s)]_a} / \sum_{b \in \mathcal{A}} e^{[\boldsymbol{\theta}_\mathbf{w}(s)]_b}$, where $\boldsymbol{\theta}_\mathbf{w} : \mathcal{S} \rightarrow \mathbb{R}^{\mathcal{A}}$ denotes the action preference vector, parameterized by \mathbf{w} , and $[\boldsymbol{\theta}_\mathbf{w}(s)]_a$ is its a th element. In practice, $\boldsymbol{\theta}_\mathbf{w}$ could be implemented using a neural network. For brevity, we will drop \mathbf{w} , i.e. $\boldsymbol{\theta} \equiv \boldsymbol{\theta}_\mathbf{w}$. Now, using the chain rule of differentiation (see §A) $\nabla_{\mathbf{w}} \boldsymbol{\pi}(\cdot|s) = [\nabla_{\mathbf{w}} \boldsymbol{\theta}(s)] [\nabla_{\boldsymbol{\theta}} \boldsymbol{\pi}(\cdot|s)]$, we can re-write the PG theorem (Eq. 3) to explicitly separate the gradient of the action likelihood $\nabla_{\boldsymbol{\theta}} \boldsymbol{\pi}(\cdot|s)$ and the preference $\nabla_{\mathbf{w}} \boldsymbol{\theta}(s)$:

$$\begin{aligned} \nabla \mathcal{J}_\pi &= \sum_s \nu_\pi(s) \sum_a \nabla \pi(a|s) q_\pi(s, a) \\ &= \sum_s \nu_\pi(s) \nabla \boldsymbol{\pi}(\cdot|s) \mathbf{q}_\pi(s, \cdot) \\ &= \sum_s \nu_\pi(s) [\nabla_{\mathbf{w}} \boldsymbol{\theta}(s)] \nabla_{\boldsymbol{\theta}} \boldsymbol{\pi}(\cdot|s) \mathbf{q}_\pi(s, \cdot). \end{aligned} \quad (9)$$

In the analysis that follows, we will analytically work out the gradient of the action likelihood for the softmax policy. We will leave the gradient of the preference which depends on the specific function approximator used and could be computed later using, say, an automatic differentiation package. We can write

$$\begin{aligned} \nabla_{\boldsymbol{\theta}} \boldsymbol{\pi}(\cdot|s) \mathbf{q}_\pi(s, \cdot) &= (\mathbf{I} - \boldsymbol{\pi}(\cdot|s) \mathbf{1}^\top) \text{diag}(\boldsymbol{\pi}(\cdot|s)) \mathbf{q}_\pi(s, \cdot) \\ &= \text{diag}(\boldsymbol{\pi}(\cdot|s)) \mathbf{q}_\pi(s, \cdot) - \boldsymbol{\pi}(\cdot|s) \mathbf{1}^\top \text{diag}(\boldsymbol{\pi}(\cdot|s)) \mathbf{q}_\pi(s, \cdot) \\ &= \boldsymbol{\pi}(\cdot|s) \odot \mathbf{q}_\pi(s, \cdot) - \boldsymbol{\pi}(\cdot|s) \boldsymbol{\pi}(\cdot|s)^\top \mathbf{q}_\pi(s, \cdot) \end{aligned} \quad (10)$$

$$\begin{aligned} &= \text{diag}(\mathbf{q}_\pi(s, \cdot)) \boldsymbol{\pi}(\cdot|s) - \boldsymbol{\pi}(\cdot|s) \mathbf{q}_\pi(s, \cdot)^\top \boldsymbol{\pi}(\cdot|s) \\ &= [\text{diag}(\mathbf{q}_\pi(s, \cdot)) - \boldsymbol{\pi}(\cdot|s) \mathbf{q}_\pi(s, \cdot)^\top] \boldsymbol{\pi}(\cdot|s) \\ &= \sum_a \pi(a|s) q_\pi(s, a) (\mathbf{e}_A - \boldsymbol{\pi}(\cdot|s)). \end{aligned} \quad (11)$$

Putting Eq. 11 into Eq. 9 gives us that $\nabla_{\mathbf{w}} \mathcal{J}_\pi$

$$\begin{aligned} &= \sum_s \nu_\pi(s) \sum_a \pi(a|s) \nabla_{\mathbf{w}} \boldsymbol{\theta}(s) (\mathbf{e}_A - \boldsymbol{\pi}(\cdot|s)) q_\pi(s, a) \\ &= \mathbb{E} \left[\underbrace{\left(\nabla_{\mathbf{w}} [\boldsymbol{\theta}(S)]_A - \sum_a \pi(a|S) \nabla_{\mathbf{w}} [\boldsymbol{\theta}(S)]_a \right)}_{=: \mathbf{g}^{\text{REG}}(S, A)} h_\pi(S, A) \right], \end{aligned}$$

where we introduced the advantage function $h_\pi(S, A) := q_\pi(S, A) - v_\pi(S)$ by subtracting an action independent v_π from q_π . Note that despite the unfamiliar expression, this is the regular estimator: $\mathbf{g}^{\text{REG}}(S, A) = \nabla_{\mathbf{w}} \log \pi(A|S) h_\pi(S, A)$ (refer §C).

The alternate estimator for MDPs. From Eq. 10,

$$\begin{aligned} &\nabla_{\boldsymbol{\theta}} \boldsymbol{\pi}(\cdot|s) \mathbf{q}_\pi(s, \cdot) \\ &= \text{diag}(\mathbf{q}_\pi(s, \cdot)) \boldsymbol{\pi}(\cdot|s) - \boldsymbol{\pi}(\cdot|s)^\top \mathbf{q}_\pi(s, \cdot) \boldsymbol{\pi}(\cdot|s) \\ &= [\text{diag}(\mathbf{q}_\pi(s, \cdot)) - v_\pi(s) \mathbf{I}] \boldsymbol{\pi}(\cdot|s) \\ &= \sum_a \pi(a|s) (q_\pi(s, a) - v_\pi(s)) \mathbf{e}_A, \end{aligned} \quad (12)$$

where we used that $v_\pi(s) = \sum_a \pi(a|s) q_\pi(s, a) = \boldsymbol{\pi}(\cdot|s)^\top \mathbf{q}_\pi(s, \cdot)$. Now put Eq. 12 in Eq. 9 to obtain

$$\begin{aligned} \nabla_{\mathbf{w}} \mathcal{J}_\pi &= \sum_s \nu_\pi(s) \sum_a \pi(a|s) \nabla_{\mathbf{w}} \boldsymbol{\theta}(s) \mathbf{e}_A h_\pi(S, A) \\ &= \mathbb{E} \left[\underbrace{\nabla_{\mathbf{w}} [\boldsymbol{\theta}(S)]_A}_{=: \mathbf{g}^{\text{ALT}}(S, A)} h_\pi(S, A) \right]. \end{aligned} \quad (13)$$

The alternate gradient estimator is equivalent to the gradient of the softmax policy, if we drop the normalization constant; i.e. $\nabla_{\mathbf{w}} \log \frac{e^{[\boldsymbol{\theta}(s)]_A}}{1} = \nabla_{\mathbf{w}} [\boldsymbol{\theta}(s)]_A$. Also a note that even though the regular and the alternate PG estimators are equal in expectation, in general they are not equal for an arbitrary state-action pair: $\mathbf{g}^{\text{REG}}(S, A) \neq \mathbf{g}^{\text{ALT}}(S, A)$. Finally, it is straightforward to adapt the alternate estimator given in Eq. 13 to work with different PG methods such as REINFORCE, Actor-Critic, TRPO, or PPO (see §C.2).

6 TABULAR-MDP EXPERIMENTS

In this section, we demonstrate that the alternate estimator for MDPs enjoys similar benefits as the alternate estimator for bandits. We use the chain environment, shown in Figure 4 (bottom left), which is an episodic MDP where the expected rewards are zero everywhere expect at the rightmost transition. We train five different agents using REINFORCE (Williams, 1992). All the agents maintain a tabular policy (and in some cases additionally a tabular value function estimate). The policy is learned either using the expected PG update, or using the REINFORCE algorithm using the regular estimator (true v_π or learned baseline) or




Figure 4: Learning curves and parameter sensitivity plots for REINFORCE on the five state chain MDP. Each agent was run for 100 episodes. The parameter sensitivity plots show the mean performance during the last 10 episodes, averaged over 150 runs. Policy was initialized, for all states, to $\theta_{\text{right}} = 0$; and $\theta_{\text{left}} = 0$ (for uniform), $\theta_{\text{left}} = 3$ (for saturated), and $\theta_{\text{left}} = 1$ (for saturated + pessimistic baseline).

the alternate estimator (again with true v_π or learned baseline). The critic is estimated using Monte-Carlo sampling. Further details (including algorithm pseudocode) and additional experiments are given in §E.

Experiment 4 showcases that the alternate PG estimator is competitive with the regular estimator in case of uniform policy initialization and superior to it in case of sub-optimally saturated initialization. The results are essentially the same as that from the bandit experiments. Figure 4 (bottom-left) shows the learning curves for the best performing parameter configuration, and Figure 4 (middle) shows the stepsize sensitivity corresponding to the final performance for each parameter setting. Observe that for the saturated initialization case, alternate with baseline is vastly superior to all the methods, and alternate with true v_π is a little better than the regular estimators. We attribute the superior performance of alternate estimator with baseline to the bias of the estimator combined with utilizing the noise in the returns, which probably allows it to have better exploration; this point is also reinforced from the stepsize sensitivity plots: smaller β values have higher final performance. We show the performance of the methods with no reward noise in §E: the results are similar to the bandit setting.

Experiment 5 studies the performance of the alternate estimator with optimistically and pessimistically initialized baselines. Figure 4 (right) shows that having an optimistic baseline significantly helps the alternate estimator; in particular, compare the performance of the alternate estimator on saturated policies with ($b = +4$) and without ($b = 0$) optimism. Whereas, a pessimistic baseline hurts its performance. Also note that the alternate estimator with the optimistic baseline prefers smaller critic stepsizes (allowing it to enjoy the optimism for longer), and vice-versa

for the pessimistic baseline. Even though we found that, in general, an optimistic baseline helps the performance of the alternate estimator, for certain MDPs (where a uniform policy is bad for exploration; see §E.5) it can also hurt its performance.

7 FUNCTION APPROXIMATION EXPERIMENTS

In this section, we demonstrate that the desirable properties of the alternate estimator continue to hold with function approximation. We run experiments with the online Actor-Critic algorithm (Degrís et al., 2012) on classical MDPs using linear function approximation (with tile-coding) and neural networks. Unlike the previous experiments, we did not saturate the policy artificially. Instead, we induced non-stationarity in the environment which naturally lead to sub-optimally saturated policies. Additional details and an extensive empirical analysis (including artificially induced policy saturation) are given in §F and §G.

Experiment 6 studies the online AC algorithm on MountainCar (Moore, 1990) and Acrobot (Sutton, 1996) control tasks. To induce non-stationarity in either task, we switched the **left** and **right** actions after half-time. Figure 39 (left) shows the learning curves for the best parameter configuration. For each estimator, we selected two sets of parameter configurations: one that had the best performance right before the non-stationarity hit, and another one that had the best performance at the end of the experiment (denoted by a **final** in the subscript). For MountainCar, the alternate estimator is superior to the regular estimator for both sets of stepsizes. Remarkably, the best performing parameter set for regular (**red curve**) at timestep 100k was unable to recover from




Figure 5: Learning curves, entropy plots, and sensitivity plots for Online Actor-Critic with linear function approximation on MountainCar (200k timesteps) and Acrobot (400k timesteps). Learning curves are for the best performing stepsize configurations at the time non-stationarity was introduced and at the final timestep. Entropy plots show entropy of the policy on the exact states encountered during each episode. The sensitivity plots show the mean performance during the last 5000 timesteps. All the results were averaged over 50 runs.




Figure 6: **(Left)** Parameter sensitivity of online AC with linear function approximation on RealisticMountainCar (100k timesteps, 50 runs); plot shows average performance in the last 5000 steps. **(Right)** Learning curves for online Actor-Critic with neural networks on DotReacher (100k timesteps, 50 runs).

the non-stationarity; whereas alternate (blue curve), despite having similar performance as regular at 100k timestep, was able to recover. On Acrobot, the difference in performance is still there but relatively smaller.

We attribute the superior performance of the alternate estimator to the bias in the critic estimate. To see this, consider MountainCar: at 100k timesteps, the critic would have converged to predict a return of about -200 . But when the non-stationarity hit, the agent would have started receiving returns much lower than -200 . But this means that at that time, the critic estimate became optimistic and encouraged exploration by pushing the policy towards a more uniform distribution. Figure 39 (middle) corroborates this point: the policy entropy for alternate (but not for regular) jumps right around the timestep when the non-stationarity hit. Figure 39 (right) shows the parameter sensitivity plots for these tasks.

Experiment 7 considers a difficult control environ-

ment called RealisticMountainCar (see §F). As Figure 6 (left) shows even with a uniformly initialized policy, the alternate estimator performed superior to the regular estimator; clearly, the optimistically initialized baseline ($b = 0$) helped in exploration.

Experiment 8 indicates that the alternate estimator works with neural policies as well. Figure 6 (right) shows the learning curves of online AC with neural networks on the DotReacher task (see §G). The results follow the same pattern as Experiment 6: the alternate estimator is still competitive or superior, although the difference is less pronounced.

Experiment 9 studies two approaches designed to handle policy saturation: entropy regularization (Ahmed et al., 2019) and escort transform (Mei et al., 2020b). Entropy regularization aims to overcome policy saturation by encouraging exploration. It augments the rewards received from the environment with the action-entropy of the policy scaled by a constant parameter τ . A larger τ encourages a more uniform policy. Escort transform (ET) is an alternative to softmax policies, given by $\pi(a|s) = \frac{[\theta(s)]_a^p}{\sum_{a'} [\theta(s)]_{a'}^p}$, where $p \geq 1$ is the escort parameter. ET has inherently better convergence properties for escaping policy saturation as compared to softmax. For further details, see §C.4 and §C.5.

We trained online AC with entropy regularization and softmax policy (both regular and alternate estimators) and separately with ET on non-stationary MountainCar and Acrobot tasks. Figure 7 shows the learning curves for the best performing parameter settings. These learning curves show that adding an entropy bonus didn't seem to improve the performance of the regular estimator (red curve). Regular with entropy performed worse even when compared to alter-

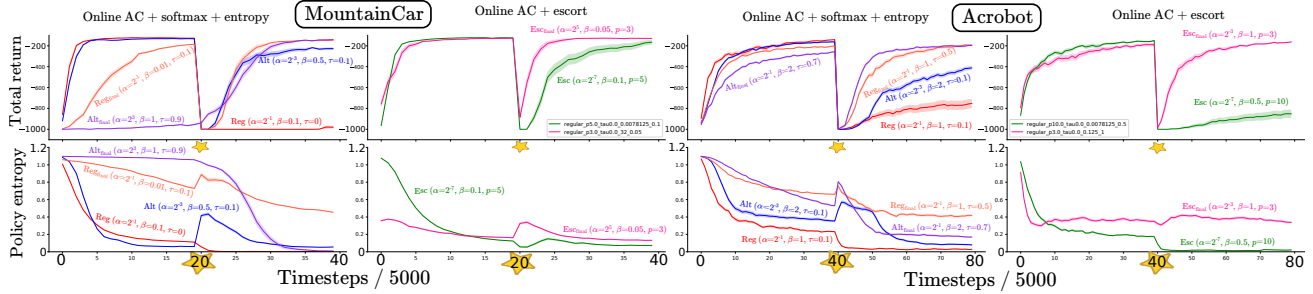


Figure 7: Learning curves and entropy plots for Online Actor-Critic with the softmax policy parameterization (regular and alternate estimators) and the escort policy parameterization. We chose the best performing stepsize configurations at the time non-stationarity was introduced and an additional set of stepsizes that performed best at the final timestep. Entropy plots show mean entropy of the policy on the exact states encountered during each episode. All the results were averaged over 50 runs.

nate without an entropy bonus (from Experiment 6). On the other hand, ET seemed to handle the non-stationarity somewhat well. Particularly, on MountainCar, escort (green curve) slightly outperformed alternate (blue curve). However, on Acrobot, alternate remained superior. From the parameter sensitivity plots for these experiments (given in §F.3), we further found that alternate outperforms regular with entropy for almost all the parameters considered. We also found that escort transform’s performance is somewhat sensitive to the value of the parameter p , which highlights another advantage of the alternate estimator: it has one less parameter than escort transform.

8 RELATED WORKS

We now discuss three prior works that bear a strong resemblance to the alternate estimator. Young (2019) studied the unbiased alternate estimator with a focus on saving compute by only calculating the gradient of a single action preference per timestep (weakly related to §B.1). In contrast, we focused on using the biased alternate estimator to overcome policy saturation. The alternate estimator looks very similar to the NeuRD (Hennes et al., 2019) update: $\nabla_{\mathbf{w}} \mathcal{J} = \sum_s \nu_{\pi}(s) \sum_a \nabla_{\mathbf{w}} [\theta(S)]_A h_{\pi}(S, A)$. However, NeuRD is missing the policy weighting from Eq. 13 (which would make it more aggressive than the unbiased alternate estimator). Also, NeuRD doesn’t consider learned (and hence biased) baselines and was motivated from a game theoretic perspective. The attractive/repulsive effect of baselines on the alternate estimator (Theorem 1) was previously reported (termed as the *non-committal* property) for the regular estimator by Chung et al. (2021). (Our experiments in Figure 14 also demonstrate this for the regular estimator.) However, because of the unbiasedness of the regular esti-

imator, this effect is much weaker compared to what is experienced by the (biased) alternate estimator.

9 CONCLUSIONS

We proposed an alternate policy gradient estimator for softmax policies that, as we demonstrated theoretically and empirically, effectively utilizes the reward noise and the bias in the critic to escape sub-optimally saturated regions in the policy space. Our analysis, conducted on multiple bandit and MDP tasks, suggests that this estimator works well with different PG algorithms and different function approximation schemes. The alternate estimator makes existing PG methods more viable for non-stationary problems, and by extension for many practical real-life control tasks.

Limitations and Future Work: Inherently, the scope of this work is limited to discrete action spaces and the log-likelihood family of PG estimators. In this article, we provided a sufficient proof-of-concept that motivates a further study of the alternate estimator. Our theoretical analysis only considered the tabular bandit setting where the results were asymptotic in nature. It would be useful to extend them to MDPs with finite time analysis. A particularly interesting direction is to see, possibly using a two-timescale analysis, how the bias in the alternate estimator affects the policy optimization dynamics as the baseline estimate becomes more accurate. Empirically, our results did not explore the batch PG algorithms such as PPO (Schulman et al., 2017) or A2C (Mnih et al., 2016). We had good reasons for avoiding them: batch algorithms have a large number of interacting parts which introduces tangential factors making a careful empirical analysis much more difficult. Having said that, an important future work would be to study the alternate estimator with these algorithms on large MDPs.

Acknowledgments

The authors gratefully acknowledge funding from the Canada CIFAR AI Chairs program, the Reinforcement Learning and Artificial Intelligence (RLAI) laboratory, the Alberta Machine Intelligence Institute (Amii), and the Natural Sciences and Engineering Research Council (NSERC) of Canada. Shivam Garg also gratefully acknowledges support from Csaba Szepesvári during part of the duration of this project, and thanks Compute Canada for providing computational resources.

References

- Agarwal, A., Kakade, S. M., Lee, J. D., Mahajan, G. (2021). On the theory of policy gradient methods: Optimality, approximation, and distribution shift. *Journal of Machine Learning Research*, 22(98), 1-76.
- Agarwal, A., Jiang, N., Kakade, S. M., Sun, W. (2019). Reinforcement Learning: Theory and Algorithms. *CS Dept., UW Seattle, Seattle, WA, USA, Technical Report*.
- Ahmed, Z., Le Roux, N., Norouzi, M., Schuurmans, D. (2019). Understanding the impact of entropy on policy optimization. In *International conference on machine learning* (pp. 151-160).
- Auer, P., Cesa-Bianchi, N., Fischer, P. (2002). Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2), 235-256.
- Barnard, E. (1993). Temporal-difference methods and Markov models. *IEEE Transactions on Systems, Man, and Cybernetics*, 23(2), 357-365.
- Berner, C., Brockman, G., Chan, B., Cheung, V., Debiak, P., Dennison, C., Farhi, D., Fischer, Q., Hashme, S., Hesse, C., Józefowicz, R. (2019). Dota 2 with Large Scale Deep Reinforcement Learning. *arXiv preprint arXiv:1912.06680*.
- Chung, W., Thomas, V., Machado, M. C., Le Roux, N. (2021). Beyond variance reduction: Understanding the true impact of baselines on policy optimization. In *International Conference on Machine Learning* (pp. 1999-2009).
- Ciosek, K., Whiteson, S. (2020). Expected Policy Gradients for Reinforcement Learning. *Journal of Machine Learning Research*, 21(52).
- Cormen, T. H., Leiserson, C. E., Rivest, R. L., Stein, C. (2009). *Introduction to algorithms*. MIT press.
- Degrís, T., Pilarski, P. M., Sutton, R. S. (2012). Model-free reinforcement learning with continuous action in practice. In *American Control Conference (ACC)* (pp. 2177-2182). IEEE.
- Ding, Y., Zhang, J., Lavaei, J. (2021). Beyond Exact Gradients: Convergence of Stochastic Soft-Max Policy Gradient Methods with Entropy Regularization. *arXiv preprint arXiv:2110.10117*.
- Hennes, D., Morrill, D., Omidshafiei, S., Munos, R., Perolat, J., Lanctot, M., Gruslys, A., Lespiau, J.B., Parmas, P., Duenez-Guzman, E., Tuyls, K., (2019). Neural replicator dynamics. *arXiv preprint arXiv:1906.00190*.
- Kakade, S. M. (2001). A natural policy gradient. *Advances in neural information processing systems*, 14.
- Kingma, D. P., Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Lafferty, J., Liu, H., Wasserman, L. (2008). Concentration of measure. *Available online: <http://www.stat.cmu.edu/~larry/=sml/Concentration.pdf>*.
- Li, G., Wei, Y., Chi, Y., Gu, Y., Chen, Y. (2021). Softmax policy gradient methods can take exponential time to converge. *arXiv preprint arXiv:2102.11270*.
- Mahmood, A. R., Korenkevych, D., Vasan, G., Ma, W., Bergstra, J. (2018). Benchmarking reinforcement learning algorithms on real-world robots. *Conference on robot learning*.
- Mei, J., Xiao, C., Szepesvári, C., Schuurmans, D. (2020a). On the global convergence rates of softmax policy gradient methods. In *International Conference on Machine Learning* (pp. 6820-6829).
- Mei, J., Xiao, C., Dai, B., Li, L., Szepesvári, C., Schuurmans, D. (2020b). Escaping the Gravitational Pull of Softmax. *Advances in Neural Information Processing Systems*, 33.
- Mei, J., Dai, B., Xiao, C., Szepesvári, C., Schuurmans, D. (2021). Understanding the Effect of Stochasticity in Policy Optimization. *arXiv preprint arXiv:2110.15572*.
- Mnih, V., Badia, A. P., Mirza, M., Graves, A., Lillicrap, T., Harley, T., Silver, D., Kavukcuoglu, K.

- (2016). Asynchronous methods for deep reinforcement learning. In *International conference on machine learning* (pp. 1928-1937).
- Moore, A. W. (1990). Efficient memory-based learning for robot control.
- Pardo, F., Tavakoli, A., Levdik, V., Kormushev, P. (2018). Time limits in reinforcement learning. In *International Conference on Machine Learning* (pp. 4045-4054).
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S. (2019). Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 8026-8037.
- Peters, J., Schaal, S. (2008). Natural actor-critic. *Neurocomputing*, 71(7-9), 1180-1190.
- Peters, J., Mulling, K., Altun, Y. (2010). Relative entropy policy search. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence*.
- Schulman, J., Levine, S., Abbeel, P., Jordan, M., Moritz, P. (2015). Trust region policy optimization. In *International conference on machine learning* (pp. 1889-1897).
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., Klimov, O. (2017). Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Singh, S. P., Sutton, R. S. (1996). Reinforcement learning with replacing eligibility traces. *Machine learning*, 22(1), 123-158.
- Sutton, R. S. (1996). Generalization in reinforcement learning: Successful examples using sparse coarse coding. *Advances in neural information processing systems*, 1038-1044.
- Sutton, R. S., McAllester, D. A., Singh, S. P., Mansour, Y. (2000). Policy gradient methods for reinforcement learning with function approximation. In *Advances in neural information processing systems* (pp. 1057-1063).
- Sutton, R. S., Barto, A. G. (2018). *Reinforcement Learning: An Introduction*, Second Edition. MIT Press.
- Williams, R. J. (1992). Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4), 229-256.
- Young, K. (2019). Learning What to Remember: Strategies for Selective External Memory in Online Reinforcement Learning Agents. *University of Alberta, M.Sc. Thesis*.

Supplementary Material: An Alternate Policy Gradient Estimator for Softmax Policies

Contents

1	INTRODUCTION	1
2	PRELIMINARIES	2
3	ALTERNATE GRADIENT BANDITS	3
4	BANDIT EXPERIMENTS	5
5	ALTERNATE PG ESTIMATOR FOR MDPs	6
6	TABULAR-MDP EXPERIMENTS	6
7	FUNCTION APPROXIMATION EXPERIMENTS	7
8	RELATED WORKS	9
9	CONCLUSIONS	9
	References	10
A	Minor Background	14
A.1	Differentiation of Vector Variables	14
A.2	Expectation and Variance of Vector Variables	15
B	Analysis of the Alternate Estimator in the Bandit Setting	15
B.1	Per-timestep Computational Complexity	15
B.2	PROOFS: The Alternate Estimator Utilizes Reward Noise	18
B.3	PROOFS: The Alternate Estimator Utilizes its Biasedness	20
B.4	Example: Optimistic Initialization helps the Alternate Estimator	24
B.5	Visualizing the Stochastic PG Estimators for Bandits	25
C	Softmax PG Estimators for MDPs	29
C.1	Remark about the Regular Softmax PG Estimator for MDPs	29
C.2	Adapting the Alternate Estimator to various PG Algorithms	29
C.3	Additional Details for Implementing Online Linear AC	31

C.4	Entropy Regularized Policy Gradient	32
C.5	Escort Transform	33
D	Gradient Bandits: Full Experiments	34
D.1	Experimental Details	34
D.2	Alternate Estimator can Escape Saturated Policy Regions	34
D.3	Gradient Bandit Estimators with Increasingly Saturated Policy Initializations	36
D.4	Effect of Bias in the Reward Baselines on the Estimators	38
D.5	Can the Regular Estimator be Modified to Deal with Saturated Policies?	40
E	REINFORCE with Tabular Representation: Full Experiments	41
E.1	Experimental Details	42
E.2	Performance of PG Estimators against Increasingly Saturated Policy Initializations	43
E.3	Optimistic and Pessimistic Value Function Estimates	44
E.4	Effect of adding Gradient Noise to the Estimators	46
E.5	A Uniform Policy Initialization can be Sub-optimal too!	47
F	Online Actor-Critic with Linear Function Approximation	47
F.1	Experimental Details for Linear Function Approximation	49
F.2	Non-stationary MountainCar and Acrobot	49
F.3	Online AC with Entropy Regularization and Escort Transform on Non-stationary Tasks	51
F.4	Experiments with Saturated Linear Softmax Policies	53
F.5	Experiments on RealisticMountainCar	55
G	Online AC with Neural Networks	56
G.1	Details about Experiments with Neural Networks	56
G.2	Non-stationary DotReacher	56
G.3	Experiments with Saturated Neural Policies on DotReacher	57
H	A Digressive Remark: Softmax Policies Learn even with Extremely Large Stepsizes	57

A Minor Background

A.1 Differentiation of Vector Variables

In this section, we first explain the vector notation we use and then state some useful vector differentiation rules. We assume that $\boldsymbol{\theta} \in \mathbb{R}^m$ is a vector and θ_i denotes the i th element of $\boldsymbol{\theta}$, $\mathbf{c} \in \mathbb{R}^m$ is a vector independent of $\boldsymbol{\theta}$, $\text{diag}(\boldsymbol{\theta})$ denotes the $m \times m$ diagonal matrix with θ_j as the j th element on the main diagonal, the function $\alpha : \mathbb{R}^m \rightarrow \mathbb{R}$ defined as $\alpha(\boldsymbol{\theta}) := \alpha(\theta_1, \theta_2, \dots, \theta_m) \in \mathbb{R}$ is a scalar function dependent on the vector $\boldsymbol{\theta}$, the function $\mathbf{x} : \mathbb{R}^m \rightarrow \mathbb{R}^n$ defined as $\mathbf{x}(\boldsymbol{\theta}) := [x_1(\boldsymbol{\theta}) \ x_2(\boldsymbol{\theta}) \ \dots \ x_n(\boldsymbol{\theta})]^\top \in \mathbb{R}^n$ is a vector function of $\boldsymbol{\theta}$, and $\mathbf{y}(\mathbf{x}(\boldsymbol{\theta})) \in \mathbb{R}^p$ is a vector function which depends on \mathbf{x} and consequently on $\boldsymbol{\theta}$ as well. Finally, for a function $f : \mathbb{R} \rightarrow \mathbb{R}$, we define $f(\boldsymbol{\theta}) := [f(\theta_1) \ f(\theta_2) \ \dots \ f(\theta_m)]^\top$ as the vector obtained by applying the function f to the vector $\boldsymbol{\theta}$ elementwise. For example, $e^\boldsymbol{\theta} = [e^{\theta_1} \ e^{\theta_2} \ \dots \ e^{\theta_m}]^\top$.

The gradient of a scalar $\alpha(\boldsymbol{\theta})$ with respect to $\boldsymbol{\theta}$ is the vector defined as

$$\nabla_{\boldsymbol{\theta}} \alpha(\boldsymbol{\theta}) := \left[\frac{\partial \alpha(\boldsymbol{\theta})}{\partial \theta_1} \quad \frac{\partial \alpha(\boldsymbol{\theta})}{\partial \theta_2} \quad \dots \quad \frac{\partial \alpha(\boldsymbol{\theta})}{\partial \theta_m} \right]_{m \times 1}^\top.$$

We overload the gradient operator to take gradient of a vector $\mathbf{x}(\boldsymbol{\theta})$ with respect to $\boldsymbol{\theta}$ as follows:

$$\begin{aligned} \nabla_{\boldsymbol{\theta}} \mathbf{x}(\boldsymbol{\theta}) &:= \left[\nabla_{\boldsymbol{\theta}} x_1(\boldsymbol{\theta}) \quad \nabla_{\boldsymbol{\theta}} x_2(\boldsymbol{\theta}) \quad \dots \quad \nabla_{\boldsymbol{\theta}} x_n(\boldsymbol{\theta}) \right]_{m \times n} \\ &\equiv \begin{bmatrix} \frac{\partial x_1(\boldsymbol{\theta})}{\partial \theta_1} & \frac{\partial x_2(\boldsymbol{\theta})}{\partial \theta_1} & \dots & \frac{\partial x_n(\boldsymbol{\theta})}{\partial \theta_1} \\ \frac{\partial x_1(\boldsymbol{\theta})}{\partial \theta_2} & \frac{\partial x_2(\boldsymbol{\theta})}{\partial \theta_2} & \dots & \frac{\partial x_n(\boldsymbol{\theta})}{\partial \theta_2} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial x_1(\boldsymbol{\theta})}{\partial \theta_m} & \frac{\partial x_2(\boldsymbol{\theta})}{\partial \theta_m} & \dots & \frac{\partial x_n(\boldsymbol{\theta})}{\partial \theta_m} \end{bmatrix}_{m \times n}. \end{aligned}$$

According to the above definition, $\nabla_{\boldsymbol{\theta}} \mathbf{x}(\boldsymbol{\theta})$ is the transpose of the usual Jacobian matrix, and readily reduces to the gradient vector if $\mathbf{x}(\boldsymbol{\theta})$ is a scalar (i.e. $n = 1$). Further, using this convention, if $\boldsymbol{\theta}$ is a scalar (i.e. $m = 1$), then

$$\nabla_{\boldsymbol{\theta}} \mathbf{x}(\boldsymbol{\theta}) = \left[\frac{\partial x_1(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \quad \frac{\partial x_2(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \quad \dots \quad \frac{\partial x_n(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \right]_{1 \times n}.$$

Finally, if both $\mathbf{x}(\boldsymbol{\theta})$ and $\boldsymbol{\theta}$ are scalars (i.e. $m = n = 1$), then

$$\nabla_{\boldsymbol{\theta}} x(\boldsymbol{\theta}) = \frac{\partial x(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}}.$$

We are now ready to state some useful propositions, along with their proofs, regarding the differentiation of vector variables.

Proposition 3. $\nabla_{\boldsymbol{\theta}}(\boldsymbol{\theta}^\top \mathbf{c}) = \nabla_{\boldsymbol{\theta}}(\mathbf{c}^\top \boldsymbol{\theta}) = \mathbf{c}$.

Proof.

$$\frac{\partial}{\partial \theta_k} \boldsymbol{\theta}^\top \mathbf{c} = \frac{\partial}{\partial \theta_k} \sum_{i=1}^m \theta_i c_i = c_k \quad \Rightarrow \quad \nabla_{\boldsymbol{\theta}} \boldsymbol{\theta}^\top \mathbf{c} = \mathbf{c}.$$

□

Proposition 4. $\nabla_{\boldsymbol{\theta}} [\alpha(\boldsymbol{\theta}) \mathbf{x}(\boldsymbol{\theta})] = \mathbf{x}(\boldsymbol{\theta}) \nabla_{\boldsymbol{\theta}} \alpha(\boldsymbol{\theta})^\top + \alpha(\boldsymbol{\theta}) \nabla_{\boldsymbol{\theta}} \mathbf{x}(\boldsymbol{\theta})$.

Proof.

$$\begin{aligned} \left(\nabla_{\boldsymbol{\theta}} [\alpha(\boldsymbol{\theta}) \mathbf{x}(\boldsymbol{\theta})] \right)_{ij} &= \frac{\partial}{\partial \theta_i} [\alpha(\boldsymbol{\theta}) x_j(\boldsymbol{\theta})] = \frac{\partial \alpha(\boldsymbol{\theta})}{\partial \theta_i} x_j(\boldsymbol{\theta}) + \alpha(\boldsymbol{\theta}) \frac{\partial x_j(\boldsymbol{\theta})}{\partial \theta_i} \\ \Rightarrow \quad \nabla_{\boldsymbol{\theta}} [\alpha(\boldsymbol{\theta}) \mathbf{x}(\boldsymbol{\theta})] &= \nabla_{\boldsymbol{\theta}} \alpha(\boldsymbol{\theta}) \mathbf{x}(\boldsymbol{\theta})^\top + \alpha(\boldsymbol{\theta}) \nabla_{\boldsymbol{\theta}} \mathbf{x}(\boldsymbol{\theta}). \end{aligned}$$

□

Proposition 5. $\nabla_{\boldsymbol{\theta}} \mathbf{y}(\mathbf{x}(\boldsymbol{\theta})) = \left[\nabla_{\boldsymbol{\theta}} \mathbf{x}(\boldsymbol{\theta}) \right] \left[\nabla_{\mathbf{x}(\boldsymbol{\theta})} \mathbf{y}(\mathbf{x}(\boldsymbol{\theta})) \right]$.

Proof. We show this result taking the total derivative⁷ of \mathbf{y} with respect to $\boldsymbol{\theta}$.

$$\begin{aligned} \left(\nabla_{\boldsymbol{\theta}} \mathbf{y}(\mathbf{x}(\boldsymbol{\theta})) \right)_{ij} &= \frac{\partial y_j(\mathbf{x}(\boldsymbol{\theta}))}{\partial \theta_i} = \sum_{k=1}^m \frac{\partial y_j(\mathbf{x}(\boldsymbol{\theta}))}{\partial x_k(\boldsymbol{\theta})} \frac{\partial x_k(\boldsymbol{\theta})}{\partial \theta_i} \\ \Rightarrow \left(\nabla_{\boldsymbol{\theta}} \mathbf{y}(\mathbf{x}(\boldsymbol{\theta})) \right)_{.j} &= \sum_{k=1}^m \frac{\partial y_j(\mathbf{x}(\boldsymbol{\theta}))}{\partial x_k(\boldsymbol{\theta})} \nabla_{\boldsymbol{\theta}} x_k(\boldsymbol{\theta}) = \nabla_{\boldsymbol{\theta}} \mathbf{x}(\boldsymbol{\theta}) \nabla_{\mathbf{x}(\boldsymbol{\theta})} y_j(\mathbf{x}(\boldsymbol{\theta})) \\ \Rightarrow \nabla_{\boldsymbol{\theta}} \mathbf{y}(\mathbf{x}(\boldsymbol{\theta})) &= \left[\nabla_{\boldsymbol{\theta}} \mathbf{x}(\boldsymbol{\theta}) \right] \left[\nabla_{\mathbf{x}(\boldsymbol{\theta})} \mathbf{y}(\mathbf{x}(\boldsymbol{\theta})) \right], \end{aligned}$$

where $\left(\nabla_{\boldsymbol{\theta}} \mathbf{y}(\mathbf{x}(\boldsymbol{\theta})) \right)_{.j}$ denotes the j th column of the matrix $\nabla_{\boldsymbol{\theta}} \mathbf{y}(\mathbf{x}(\boldsymbol{\theta}))$. □

Proposition 6. Let f be a function, applied elementwise to a vector $\boldsymbol{\theta}$, i.e.

$$f(\boldsymbol{\theta}) := [f(\theta_1) \ f(\theta_2) \ \dots \ f(\theta_m)]^\top.$$

Then

$$\nabla_{\boldsymbol{\theta}} f(\boldsymbol{\theta}) = \text{diag} \left(\left[\frac{\partial f(\theta_1)}{\partial \theta_1} \quad \frac{\partial f(\theta_2)}{\partial \theta_2} \quad \dots \quad \frac{\partial f(\theta_m)}{\partial \theta_m} \right]^\top \right).$$

A.2 Expectation and Variance of Vector Variables

We define the expectation of a vector random variable (or equivalently that of a vector function of random variables, as discussed in the previous section) in the standard way: for a d -dimensional random variable \mathbf{X} , $\mathbb{E}[\mathbf{X}] := \sum_{\mathbf{X}} p(\mathbf{X}) \mathbf{X} = \sum_{x_1, x_2, \dots, x_d} p(X_1 = x_1, X_2 = x_2, \dots, X_d = x_d) \mathbf{X}$. Note that this definition is equivalent to taking elementwise expectation, i.e. $\mathbb{E}[\mathbf{X}] \equiv [\mathbb{E}[X_1] \ \mathbb{E}[X_2] \ \dots \ \mathbb{E}[X_d]]^\top$. Using this definition, the expectation of a vector function, also see §A.1, $\mathbf{f}(\cdot)$ operating on a (scalar or vector) random variable \mathbf{Y} , is given by $\mathbb{E}[\mathbf{f}(\mathbf{Y})] = \sum_{\mathbf{Y}} p(\mathbf{Y}) \mathbf{f}(\mathbf{Y})$.

In a similar way, the variance and covariance of vector random are computed elementwise as well: $\mathbb{V}[\mathbf{X}] = \mathbb{E}[\mathbf{X} \odot \mathbf{X}] - \mathbb{E}[\mathbf{X}] \odot \mathbb{E}[\mathbf{X}]$ and $\text{Cov}(\mathbf{X}, \mathbf{Y}) = \mathbb{E}[\mathbf{X} \odot \mathbf{Y}] - \mathbb{E}[\mathbf{X}] \odot \mathbb{E}[\mathbf{Y}]$, where \mathbf{X} and \mathbf{Y} are vector random variables (they are not matrices!). Also, to make the notation less cumbersome, we will often use $\mathbf{x}^2 := \mathbf{x} \odot \mathbf{x}$ for any vector \mathbf{x} ; therefore, we can equivalently write $\mathbb{V}[\mathbf{X}] = \mathbb{E}[\mathbf{X}^2] - \mathbb{E}[\mathbf{X}]^2$.

B Analysis of the Alternate Estimator in the Bandit Setting

In this section, we begin by discussing the properties of the alternate bandit estimator in the bandit setting. In particular, we show that the alternate gradient bandit estimator has a per-timestep computational complexity of $\log(|\mathcal{A}|)$. We also give proofs of the statements regarding how the alternate estimator utilizes the reward noise and the bias in the critic estimate. After that, we first present another example (similar to Example 1) that shows how an optimistic baseline makes the policy more uniform. Finally, we give the complete policy update and variance plots (analogous to those in Figure 1) for three armed bandit problems with different reward structures.

B.1 Per-timestep Computational Complexity

We now show that the per-timestep cost of a discrete gradient bandit algorithm using the alternate estimator has a weaker dependence on the number of actions than the regular estimator. The gradient bandit algorithms involves multiple operations during each timestep. At the start of learning, the agent has to initialize the action preferences. Then at each timestep, it needs to compute the gradient estimate, update the action preferences and the policy, and sample an action from the policy. The sampling process itself requires the agent to maintain a sampling data structure which also needs to be updated each time the policy changes. The total per-timestep computational cost of the algorithm depends on all of these operations, which we summarize in Table 1.

⁷See <https://mathworld.wolfram.com/TotalDerivative.html>.

Computational Process	Regular Estimator	Alternate Estimator
	+ Vose’s Alias Sampling	+ Tree Sampling
Initialize θ	$\mathcal{O}(\mathcal{A})$	$\mathcal{O}(\mathcal{A})$
Compute \hat{g} and update θ	$\mathcal{O}(\mathcal{A})$	$\mathcal{O}(1)$
Compute π	$\mathcal{O}(\mathcal{A})$	–
Update the sampling data structure	$\mathcal{O}(\mathcal{A})$	$\mathcal{O}(\log \mathcal{A})$
Sample $A \sim \pi$	$\mathcal{O}(1)$	$\mathcal{O}(\log \mathcal{A})$
Total cost	$\mathcal{O}(\mathcal{A} + n_{\text{iter}} \mathcal{A})$	$\mathcal{O}(\mathcal{A} + n_{\text{iter}} \log \mathcal{A})$

Table 1: The key steps involved, along with their respective computational cost, while running the gradient bandit algorithm with the regular or the alternate estimator. We denote the total number of iterations by n_{iter} . Also, it is reasonable to assume that $n_{\text{iter}} \gg |\mathcal{A}|$.

The regular estimator updates each component of the action preference vector θ at each timestep, which has a complexity that is linear in the number of actions. Let us assume, for illustration purposes, that the agent with the regular estimator uses Vose’s alias sampling method⁸ for obtaining actions from the policy. As a result, the agent needs to build an *alias table* which takes a linear time. Once this table is built, the agent can then sample an action in constant time from this table. However, updating the alias table each time has a linear cost as well. Therefore, the total per-timestep cost for this agent becomes $\mathcal{O}(|\mathcal{A}|)$. In contrast, the alternate estimator updates only a single component of the action preference vector at each timestep. Because of this, a gradient bandit instance using the alternate estimator can be sped up to have a total per-time-step complexity of $\mathcal{O}(\log |\mathcal{A}|)$. To achieve this, we employ a sampling procedure based on a balanced tree data-structure, which we call the *sampling tree*. Although, the sampling tree method needs logarithmic time for sampling a single action, it can be updated in logarithmic time as well if only a single action preference changes at each timestep (as opposed to the alias method which would require a linear time even if only a single action preference were to change).

Using a Balanced Binary Tree for Sampling Actions

The sampling tree is a balanced binary tree (for a reference on our notation, see Chapter 12, Cormen et al., 2009). It has one node corresponding to each action that stores a reference and the exponentiated action preference corresponding to that action. It also stores the sum of all the exponentiated action preferences corresponding to the actions lying in the left subtree of this node. Figure 8 shows a schematic diagram of such a node, and one possible sampling tree⁹. The cost for building the sampling tree for the first time is linear in the number of actions. The agent would need to initialize the action preferences, create a balanced tree (using Algorithm 2), and then calculate the aggregated exponentiated preferences of actions in the left-subtree (using Algorithm 3). And all of these operations are clearly linear in the size of the action space.

We summarize the complete operation of building the sampling tree in Algorithm 1, and further explain it with an example. Assume that we are given a randomly sorted list of actions $a_{\text{list}} = \{a_3, a_0, a_2, a_1\}$. Let us arbitrarily initialize our action preference list to $\theta_{\text{list}} = \{0.4, 0.1, 0.3, 0.2\}$. Now according to Algorithm 1, we will randomly assign either a_0 or a_2 as the head node; we choose a_2 as the head node. Next, we create a left sub-tree from $\{a_3, a_0\}$, this time randomly choosing a_3 to be the head node and consequently a_0 as its right child. And finally, we create the right sub-tree from a_1 . This gives us the exact tree shown in Figure 8. Note that the nodes in the resulting sampling tree are not ordered in any fashion, and nor is the tree a search tree based on the action preferences. The only required property is that it has to be balanced.

We now discuss how the agent uses this data structure for sampling actions. The sampling process itself takes logarithmic time. To sample an action from the policy, the agent first samples a key from a uniform distribution¹⁰

⁸For a reference on different sampling methods, see <https://www.keithschwarz.com/darts-dice-coins/>.

⁹The sampling tree is not necessarily a search tree. However, one can apply procedures similar to those discussed here without changing the asymptotic computational complexity, to balanced search trees such as the Red-black trees (Chapter 13, Cormen et al., 2009) to obtain a sampling tree that is also a search tree. This might be required if one needs to maintain a list over the actions ordered by their probability; say, to pick the greedy action at each timestep.

¹⁰It is standard to assume that sampling a single value from a uniform distribution takes constant time.

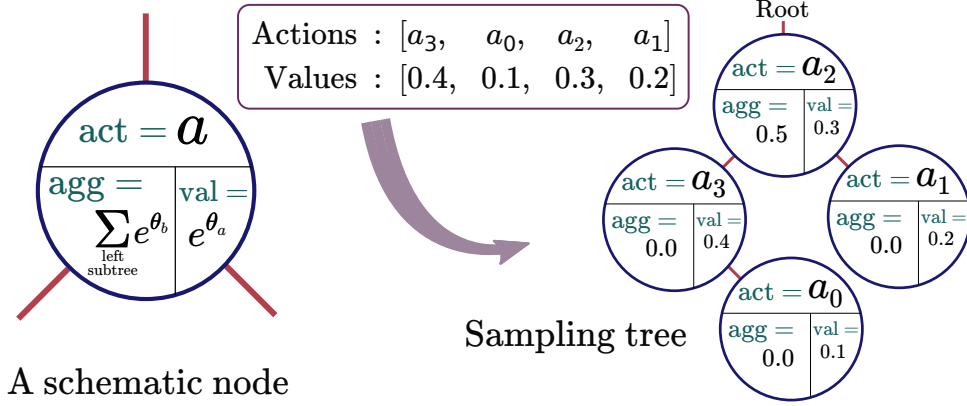


Figure 8: A typical sampling tree node and one possible tree for a softmax policy with four actions. The values refer to exponentiated action preference e^{θ_a} for that particular action.

Algorithm 1 Full Procedure for Initializing the Sampling Tree

Input $a_{\text{list}} = \{a_{i_0}, a_{i_1}, \dots, a_{i_{n-1}}\}$, where i_0, i_1, \dots, i_{n-1} is some permutation of $0, 1, \dots, n-1$
 Initialize $\theta_{\text{list}} = \{\theta_{i_0}, \theta_{i_1}, \dots, \theta_{i_{n-1}}\}$
 head node of the tree \leftarrow BUILD-SAMPLING-TREE($a_{\text{list}}, \theta_{\text{list}}, \text{nil}$)
 Call COMPUTE-AGGREGATE(head node of the tree)

Algorithm 2 Building the Sampling Tree

```

function BUILD-SAMPLING-TREE( $a_{\text{list}}, \theta_{\text{list}}, \text{parent}$ )
     $n \leftarrow \text{size}(a_{\text{list}})$ 
    if  $n = 0$  then
        return  $\text{nil}$ 
    if  $n$  is odd then # odd number of actions
         $m \leftarrow (n - 1)/2$ 
    else # even number of actions
        Randomly set  $m$  equal to one of  $(n - 1)/2$  or  $(n + 1)/2$ 
    Create a new node  $z$ 
     $z.\text{act} \leftarrow a_{\text{list}}[m]$ 
     $z.\text{val} \leftarrow \exp(\theta_{\text{list}}[m])$ 
     $z.p \leftarrow \text{parent}$ 
    # assume  $\text{list}[i \rightarrow j] := \{\text{list}[i], \text{list}[i + 1], \dots, \text{list}[j - 1]\}$ 
     $z.\text{left} \leftarrow \text{BUILD-SAMPLING-TREE}(a_{\text{list}}[0 \rightarrow m], \theta_{\text{list}}[0 \rightarrow m], z)$ 
     $z.\text{right} \leftarrow \text{BUILD-SAMPLING-TREE}(a_{\text{list}}[m + 1 \rightarrow n], \theta_{\text{list}}[m + 1 \rightarrow n], z)$ 
    return  $z$ 
    
```

Algorithm 3 Calculate Left Sub-tree Aggregate

```

function COMPUTE-AGGREGATE( $z$ )
    if  $z = \text{nil}$  then # if  $z$  is a leaf node
        return 0
     $L \leftarrow \text{COMPUTE-AGGREGATE}(z.\text{left})$ 
     $R \leftarrow \text{COMPUTE-AGGREGATE}(z.\text{right})$ 
     $z.\text{agg} \leftarrow L$ 
    return  $L + z.\text{val} + R$ 
    
```

over the interval $\left[0, \sum_{i=1}^{|\mathcal{A}|} e^{\theta_i}\right)$, and then calls the ACTION-SELECT function (see Algorithm 4) on the sampling tree with this key. Note that the ACTION-SELECT function runs in $\log(|\mathcal{A}|)$ time for a balanced sampling tree. The next proposition explains the functioning of the ACTION-SELECT function.

Algorithm 4 Selecting an Action from the Sampling Tree

```

Generate  $x \sim \mathcal{U}\left[0, \sum_{i=1}^{|\mathcal{A}|} e^{\theta_i}\right)$  #  $\mathcal{U}$  refers to the uniform distribution
Call ACTION-SELECT(head node of the tree,  $x$ )
# function definition
function ACTION-SELECT( $z, x$ )
  if  $x < z.agg$  then
    return ACTION-SELECT( $z.left, x$ )
  else if  $z.agg \leq x < z.agg + z.val$  then
    return  $z.act$ 
  else # this refers to the case when  $x \geq z.agg + z.val$ 
     $x_{new} \leftarrow x - z.agg - z.val$ 
    return ACTION-SELECT( $z.right, x_{new}$ )

```

Proposition 7. Let $T(z_1, z_2, \dots, z_n)$ denote a sampling tree of size n with nodes z_1, z_2, \dots, z_n arranged in an inorder traversal (see Chapter 12, Cormen et al., 2009). Also, without loss of generality assume that the node z_i corresponds to the action a_i and stores the value e^{θ_i} . Then given the sampling tree T and a search key $x \in [0, \sum_{i=1}^n e^{\theta_i})$ as the input, the ACTION-SELECT function returns the action $z_k.act = a_k$ such that the index k satisfies the relation $\sum_{i=1}^{k-1} e^{\theta_i} \leq x < \sum_{i=1}^k e^{\theta_i}$.

Proof. We will use induction to prove this statement. First, consider a tree of size one, $T(z_j)$ with $z_j.val = e^{\theta_j}$ and $z_j.agg = 0$. By the assumption on the search key, $x \in [0, e^{\theta_j})$, i.e. the second condition in Algorithm 4 is true, and therefore the function returns $z_j.act$. Next, assume that the function does return the intended action for any tree of size $l \leq m$ where m is some arbitrary positive number for an appropriate key x . We now prove the proposition statement for a key $x \in [0, \sum_{i=1}^{l+1} e^{\theta_i})$, and an arbitrary tree of size $l + 1$, $T(z_1, z_2, \dots, z_{r-1}, z_r, z_{r+1}, \dots, z_{l+1})$ where z_r represents the root node and $1 \leq r \leq l + 1$. Then, depending on the value of x , we get three cases:

- *Case 1* ($x < z_r.agg = \sum_{i=1}^{r-1} e^{\theta_i}$): the function gets called again on the sub-tree $T(z_1, \dots, z_{r-1})$, of size $r - 1 \leq l$, with the key $x \in [0, \sum_{i=1}^{r-1} e^{\theta_i})$.
- *Case 2* ($\sum_{i=1}^{r-1} e^{\theta_i} = z_r.agg \leq x < z_r.agg + z_r.val = \sum_{i=1}^r e^{\theta_i}$): the function returns $z_r.act$.
- *Case 3* ($x \geq z_r.agg + z_r.val = \sum_{i=1}^r e^{\theta_i}$): the function gets called on $T(z_{r+1}, \dots, z_{l+1})$, a sub-tree of size $l - r + 1 \leq l$, with the modified key $x_{new} \in [0, \sum_{i=r+1}^{l+1} e^{\theta_i})$.

This closes the induction, and thus the proposition is true for sampling trees of all sizes. □

The above proposition shows that the action a_k is selected whenever the search key x lies in the range $[\sum_{i=1}^{k-1} e^{\theta_i}, \sum_{i=1}^k e^{\theta_i})$. However, since $x \sim \mathcal{U}\left[0, \sum_{i=1}^{|\mathcal{A}|} e^{\theta_i}\right)$, the probability of action a_k being chosen is equal to $\frac{e^{\theta_k}}{\sum_{i=1}^{|\mathcal{A}|} e^{\theta_i}}$. Therefore, using Algorithm 4 is indeed equivalent to sampling from a softmax policy with action preferences $\theta_1, \theta_2, \dots, \theta_{|\mathcal{A}|}$.

When an agent, using the alternate estimator, updates the action preference for the sampled action, it only needs to update the exponentiated preference for that action in the tree and the aggregated values for all nodes lying on the path from this nodes upto the root of the tree. The procedure shown in Algorithm 5 accomplishes this in a logarithmic time. Also note that the structure of the tree remains unchanged during all these operations.

B.2 PROOFS: The Alternate Estimator Utilizes Reward Noise

In this section, we first show that the expected policy gradient at saturated regions, or equivalently the *corners* of the probability simplex boundary (see Figure 10), is zero. We then show that at these corners, the variance of the regular estimator is zero as well, but the variance of the alternate estimator is equal to the variance of the reward noise which in general is non-zero. The consequence of these two facts is that whenever the policy

Algorithm 5 Fix Aggregates Upstream after an Update

Call FIX-UPSTREAM(sampling tree, node corresponding to the sampled action a , $\theta_a^{\text{new}} - \theta_a^{\text{old}}$)
function definition
function FIX-UPSTREAM(T, z, δ)
 repeat # repeat until node z is the root node
 if $z.p.\text{left} = z$ then # update the parent's cumulant if z is the left child
 $z.p.\text{agg} \leftarrow z.p.\text{agg} + \delta$
 $z \leftarrow z.p$ **# move up the tree**
 until $z.p \neq T.\text{nil}$

is saturated, i.e. it places a high probability mass on some actions, the expected gradient becomes close to zero, and therefore the policy weights are not updated. Further, since the variance of the regular estimator at these corners is also zero, there is neither any gradient signal nor any noise, and the agent is unable to escape the sub-optimal region. In contrast, the alternate estimator, despite being zero in expectation, has a non-zero variance and therefore the agent can utilize the reward noise to escape the sub-optimal region.

Proposition 8. Define $\mathcal{I}_c := \{a \mid r(a) = c\}$ for some $c \in \mathbb{R}$. Consider a constant c such that¹¹ $\mathcal{I}_c \neq \emptyset$ and that the policy is saturated on the actions in the set \mathcal{I}_c , i.e. $\sum_{a \in \mathcal{I}_c} \pi(a) = 1$. Then the expected policy gradient for softmax policies is zero:

$$\nabla_{\theta} \mathcal{J} = \mathbb{E}[\mathbf{g}^{\text{REG}}(A, R)] = \mathbb{E}[\hat{\mathbf{g}}^{\text{REG}}(A, R)] = \mathbb{E}[\mathbf{g}^{\text{ALT}}(A, R)] = \mathbf{0}.$$

Proof. We set $\nabla_{\theta} \mathcal{J} = \boldsymbol{\pi} \odot (\mathbf{r} - r_{\boldsymbol{\pi}} \mathbf{1}) = \mathbf{0}$ to obtain a system of $|\mathcal{A}| + 1$ linear equations in $|\mathcal{A}|$ variables:

$$\pi(a) \cdot (r(a) - r_{\boldsymbol{\pi}}) = 0, \forall a \in \mathcal{A} \quad \text{and} \quad \sum_{a \in \mathcal{A}} \pi(a) = 1. \quad (14)$$

Consider $\mathcal{I}_c = \{a \mid r(a) = c\}$ where c is equal to the expected reward for some action. Then for a policy satisfying $\sum_{a \in \mathcal{I}_c} \pi(a) = 1$, we immediately see that $\pi(a) = 0, \forall a \notin \mathcal{I}_c$ and $r_{\boldsymbol{\pi}} = c = r(a), \forall a \in \mathcal{I}_c$. Consequently, this policy would also satisfy Eqs. 14 and thereby $\nabla_{\theta} \mathcal{J} = \mathbf{0}$. \square

The above proposition shows that the expected gradient bandit update $\nabla_{\theta} \mathcal{J}_{\boldsymbol{\pi}}$ is zero only at the simplex boundary which can happen either at any one of the corners, or at a point on one of the face (i.e. a sub-space) of the simplex such that all the actions on that face have the same expected reward. The next proposition considers the variance of the estimators at the simplex corners.

Proposition 9. Let $\sigma(a)^2 := \mathbb{V}[R|A = a]$ denote the conditional variance of the reward corresponding to action a . Assume that the policy is saturated on the action c , i.e., $\pi(c) = 1$. Also assume that $\sigma(c)^2 \neq 0$. Then, the variance of the regular PG estimator (with or without a baseline) is zero: $\mathbb{V}[\mathbf{g}^{\text{REG}}(A)] = \mathbb{V}[\hat{\mathbf{g}}^{\text{REG}}(A, R)] = \mathbf{0}$. Whereas, the variance of the alternate PG estimator (with or without a baseline) is non-zero: $\mathbb{V}[\mathbf{g}^{\text{ALT}}(A, R)] = \mathbb{V}[\hat{\mathbf{g}}^{\text{ALT}}(A, R)] = \sigma(c)^2 \mathbf{e}_c$.

Proof. As defined in the proof statement, the conditional variance of the rewards given an action a is denoted using $\sigma(a)^2$, i.e. $\sigma(a)^2 := \mathbb{V}[R|A = a]$. Further, define $\boldsymbol{\sigma}^2$ to be the conditional variance vector with $[\boldsymbol{\sigma}^2]_k = \sigma(a_k)^2$. Also for brevity, let $\mathbf{x}^2 := \mathbf{x} \odot \mathbf{x}$ for any vector \mathbf{x} . Then, it is straightforward to show the following

¹¹We can always find such a c by choosing it to be equal to the expected reward of some action.

statements:

$$\mathbb{E}[R] = \sum_a \pi(a) \sum_x p(x|a)x = \sum_a \pi(a)r(a) = r_\pi, \quad (15)$$

$$\begin{aligned} \mathbb{E}[R^2] &= \mathbb{E}_A \left[\mathbb{E}_R[R^2|A] \right] = \mathbb{E}_A \left[\mathbb{V}[R|A] + \mathbb{E}[R|A]^2 \right] = \sum_a \pi(a) \left(\sigma(a)^2 + r(a)^2 \right) \\ &= \boldsymbol{\pi}^\top (\boldsymbol{\sigma}^2 + \mathbf{r}^2), \end{aligned} \quad (16)$$

$$\mathbb{E}[\mathbf{e}_A] = \sum_a \pi(a) \mathbf{e}_a = \boldsymbol{\pi}, \quad (17)$$

$$\mathbb{E}[\mathbf{e}_A R] = \sum_a \pi(a) \mathbf{e}_a \sum_x p(x|a)x = \sum_a \pi(a) \mathbf{e}_a r(a) = \boldsymbol{\pi} \odot \mathbf{r}, \quad (18)$$

$$\mathbb{E}[\mathbf{e}_A R^2] = \mathbb{E}_A \left[\mathbf{e}_A \mathbb{E}_R[R^2|A] \right] = \sum_a \pi(a) \mathbf{e}_a \left(\sigma(a)^2 + r(a)^2 \right) = \boldsymbol{\pi} \odot (\boldsymbol{\sigma}^2 + \mathbf{r}^2). \quad (19)$$

Next, let $\boldsymbol{\eta}$ be an action independent vector and b be an action independent scalar. And assume that the variance for vector variables is calculated elementwise, i.e. the variance and covariance of arbitrary vector random variables \mathbf{X} and \mathbf{Y} are defined as $\mathbb{V}[\mathbf{X}] := \mathbb{E}[\mathbf{X} \odot \mathbf{X}] - \mathbb{E}[\mathbf{X}] \odot \mathbb{E}[\mathbf{X}]$ and $\text{Cov}(\mathbf{X}, \mathbf{Y}) := \mathbb{E}[\mathbf{X} \odot \mathbf{Y}] - \mathbb{E}[\mathbf{X}] \odot \mathbb{E}[\mathbf{Y}]$ respectively. Then,

$$\begin{aligned} &\mathbb{V}[(\mathbf{e}_A - \boldsymbol{\eta})(R - b)] \\ &= \mathbb{V}[\mathbf{e}_A R - \mathbf{e}_A b - \boldsymbol{\eta} R + \boldsymbol{\eta} b] = \mathbb{V}[\mathbf{e}_A R - \mathbf{e}_A b - \boldsymbol{\eta} R] \\ &= \mathbb{V}[\mathbf{e}_A R] + \mathbb{V}[\mathbf{e}_A b] + \mathbb{V}[\boldsymbol{\eta} R] - 2 \text{Cov}(\mathbf{e}_A R, \mathbf{e}_A b) - 2 \text{Cov}(\mathbf{e}_A b, \boldsymbol{\eta} R) + 2 \text{Cov}(\mathbf{e}_A R, \boldsymbol{\eta} R) \\ &= \mathbb{V}[\mathbf{e}_A R] + \mathbb{V}[\mathbf{e}_A] b^2 + \mathbb{V}[R] \boldsymbol{\eta}^2 - 2 \text{Cov}(\mathbf{e}_A R, \mathbf{e}_A) b - 2 \text{Cov}(\mathbf{e}_A R, R) \odot \boldsymbol{\eta} + 2 \text{Cov}(\mathbf{e}_A, R) \odot \boldsymbol{\eta} b. \end{aligned} \quad (20)$$

Simplifying the above expression for a policy saturated at action c , i.e. $\boldsymbol{\pi} = \mathbf{e}_c$, gives us

$$\begin{aligned} &\mathbb{V}[\mathbf{e}_A R] + \mathbb{V}[\mathbf{e}_A] b^2 + \mathbb{V}[R] \boldsymbol{\eta}^2 - 2 \text{Cov}(\mathbf{e}_A R, \mathbf{e}_A) b - 2 \text{Cov}(\mathbf{e}_A R, R) \odot \boldsymbol{\eta} + 2 \text{Cov}(\mathbf{e}_A, R) \odot \boldsymbol{\eta} b \\ &= \mathbf{e}_c \odot \boldsymbol{\sigma}^2 + \mathbf{0} + \sigma(c)^2 \boldsymbol{\eta}^2 - \mathbf{0} - 2 \boldsymbol{\sigma} \odot \mathbf{e}_c \odot \boldsymbol{\eta} + \mathbf{0} \\ &= \sigma(c)^2 \mathbf{e}_c + \sigma(c)^2 \boldsymbol{\eta}^2 - 2\sigma(c)^2 \boldsymbol{\eta}(c)^2 \mathbf{e}_c = \begin{cases} \sigma(c)^2 \mathbf{e}_c & \text{if } \boldsymbol{\eta} = \mathbf{0} \text{ (i.e. for alternate estimator), or} \\ \mathbf{0} & \text{if } \boldsymbol{\eta} = \boldsymbol{\pi} = \mathbf{e}_c \text{ (i.e. for the regular estimator).} \end{cases} \end{aligned}$$

This concludes the proof. \square

Note that these two propositions again require the policy to have support only at a subset of all the actions. This condition cannot be fulfilled for softmax policies since the action preferences for these actions will need to be infinitely large. Despite that, using continuity arguments, we can still reason that for the regular estimator, both the expected gradient and its variance vanish in the proximity of the simplex boundary, whereas the alternate estimator will have non-zero variance.

B.3 PROOFS: The Alternate Estimator Utilizes its Biasedness

The alternate estimator can also utilize the bias in the critic estimate b to escape the saturated regions of the policy space. If the critic estimate b is not equal to r_π , then the alternate estimator becomes biased, whereas the regular estimator does not. And for the alternate estimator, this bias helps make the policy more uniform thereby moving the agent out of the saturated region. To see this biasedness, consider a baseline $b \neq r_\pi$ and a vector $\boldsymbol{\eta}$ which we will set equal to either $\boldsymbol{\pi}$ for the regular estimator, or $\mathbf{0}$ for the alternate estimator. Then it is straightforward to see that

$$\begin{aligned} \mathbb{E}[(\mathbf{e}_A - \boldsymbol{\eta})(R - b)] &= \mathbb{E}[\mathbf{e}_A R] - \mathbb{E}[\mathbf{e}_A b] - \mathbb{E}[R] \boldsymbol{\eta} + \boldsymbol{\eta} b = \boldsymbol{\pi} \odot \mathbf{r} - \boldsymbol{\pi} b - r_\pi \boldsymbol{\eta} + \boldsymbol{\eta} b \\ &= \begin{cases} \boldsymbol{\pi} \odot (\mathbf{r} - r_\pi \mathbf{1}) & \text{if } \boldsymbol{\eta} = \boldsymbol{\pi}, \quad \text{and} \\ \boldsymbol{\pi} \odot (\mathbf{r} - b \mathbf{1}) & \text{if } \boldsymbol{\eta} = \mathbf{0}. \end{cases} \end{aligned} \quad (21)$$

Therefore,

$$\nabla_{\boldsymbol{\theta}} \mathcal{J} = \mathbb{E}[\hat{\mathbf{g}}^{\text{REG}}(A, R)] = \boldsymbol{\pi} \odot (\mathbf{r} - r_\pi \mathbf{1}) \neq \boldsymbol{\pi} \odot (\mathbf{r} - b \mathbf{1}) = \mathbb{E}[\hat{\mathbf{g}}^{\text{ALT}}(A, R)].$$

Interestingly, this biased update $\mathbb{E}[\hat{\mathbf{g}}^{\text{ALT}}(A, R)]$ is not the gradient of any function. To show this, we adopt the same process as Barnard (1993) did for the TD update.

Proposition 10. *The biased expected gradient bandits update is not a gradient of any objective function, i.e. for no function \mathcal{L} , $\nabla_{\theta}\mathcal{L} = \boldsymbol{\pi} \odot (\mathbf{r} - b)$.*

Proof. We prove this statement by contradiction. Assume that the biased update is the gradient of some function \mathcal{L} , i.e. $\frac{\partial \mathcal{L}}{\partial \theta_a} = \pi(a) \cdot (r(a) - b)$. Then,

$$\frac{\partial^2 \mathcal{L}}{\partial \theta_c \partial \theta_a} = \pi(a) [\mathbb{I}(a = c) - \pi(c)] (r(a) - b) \quad \Rightarrow \quad \frac{\partial^2 \mathcal{L}}{\partial \theta_c \partial \theta_a} \neq \frac{\partial^2 \mathcal{L}}{\partial \theta_a \partial \theta_c}, \quad \text{for } a \neq c,$$

which gives a contradiction¹². As a result, the biased update cannot be the gradient of any objective function. \square

The above result suggests that we cannot use the usual optimization analysis, such as setting the gradient equal to zero or computing second order derivatives, to study the properties of this biased update. Therefore, we instead try to understand its behavior by looking at its fixed point. It is now essential to introduce a few more terms. Let $r_1 \leq r_2 \leq \dots \leq r_k$ represent an ordering of the true rewards for the bandit problem. Let $\boldsymbol{\pi}_t = e^{\boldsymbol{\theta}^{(t)}} / \mathbf{1}^\top e^{\boldsymbol{\theta}^{(t)}}$ represent the softmax policy at timestep t . Here, $\boldsymbol{\theta}^{(t)}$ represents the action preference vector at timestep t ; which is updated using the rule $\boldsymbol{\theta}^{(t+1)} = \boldsymbol{\theta}^{(t)} + \alpha \mathbb{E}[\hat{\mathbf{g}}^{\text{ALT}}(A, R)]$ for $\alpha > 0$.

Lemma 10.1. *The fixed points of the biased gradient bandit update can be categorized as follows: (1) If there exists an $n \in \{1, 2, \dots, k-1\}$ such that $r_1 \leq \dots \leq r_n < b < r_{n+1} \leq \dots \leq r_k$, then $\mathbb{E}_{\boldsymbol{\pi}}[\hat{\mathbf{g}}^{\text{ALT}}(A, R)]$ is never equal to zero and has no fixed point. (2) If $b = r(a)$ for at least one action, then $\mathbb{E}_{\boldsymbol{\pi}}[\hat{\mathbf{g}}^{\text{ALT}}(A, R)] = 0$ at any point on the face of the probability simplex given by $\sum_{a \in \mathcal{I}_b} \pi(a) = 1$ with $\mathcal{I}_b = \{a | r(a) = b\}$. (3) If $b < r_1$ or $b > r_k$, then $\mathbb{E}_{\boldsymbol{\pi}}[\hat{\mathbf{g}}^{\text{ALT}}(A, R)] = 0$ at a point $\boldsymbol{\pi}^*$ within the simplex boundary given by $\pi^*(a) = \frac{1}{r(a)-b} \left(\sum_{c \in \mathcal{A}} \frac{1}{r(c)-b} \right)^{-1}$, $\forall a \in \mathcal{A}$.*

Proof. Let $\mathbf{g} := \mathbb{E}[\hat{\mathbf{g}}^{\text{ALT}}(A, R)]$, and recall that we update the policy parameters using $\boldsymbol{\theta}^{\text{new}} = \boldsymbol{\theta}^{\text{old}} + \alpha \mathbf{g}$, and then find the policy by applying the softmax function. Now, at the fixed point, the policy should stop changing, i.e. $\boldsymbol{\pi}^{\text{old}} = \boldsymbol{\pi}^{\text{new}}$. This means that for each action $a \in \mathcal{A}$,

$$\frac{e^{\theta_a^{\text{old}}}}{\sum_b e^{\theta_b^{\text{old}}}} = \frac{e^{\theta_a^{\text{new}}}}{\sum_b e^{\theta_b^{\text{new}}}} = \frac{e^{\theta_a^{\text{old}} + \alpha g_a}}{\sum_b e^{\theta_b^{\text{old}} + \alpha g_b}} = \frac{e^{\theta_a^{\text{old}}} \cdot e^{\alpha g_a}}{\sum_b e^{\theta_b^{\text{old}} + \alpha g_b}} \quad \Leftrightarrow \quad g_a = \frac{1}{\alpha} \log \frac{\sum_b e^{\theta_b^{\text{old}} + \alpha g_b}}{\sum_b e^{\theta_b^{\text{old}}}} =: \kappa, \quad (22)$$

where we introduced κ , a constant independent of the action a . Therefore, at the fixed point of the gradient bandit update, each element of the expected gradient is equal to the constant κ .

Let us use this knowledge to solve for the policy at the fixed point of the biased PG update $\mathbf{g} = \boldsymbol{\pi} \odot (\mathbf{r} - b\mathbf{1})$. We have the following system of $|\mathcal{A}| + 1$ linear equations in $|\mathcal{A}| + 1$ variables:

$$\pi(a) \cdot (r(a) - b) = \kappa, \quad \forall a \in \mathcal{A} \quad \text{and} \quad \sum_{a \in \mathcal{A}} \pi(a) = 1. \quad (23)$$

We now consider the three cases given in the lemma statement separately:

Case 1 (b lies between the expected rewards): We know that there exist two actions a_n and a_{n+1} such that $r(a_n) = r_n < b < r_{n+1} = r(a_{n+1})$, then

$$\pi(a_n) \cdot (r(a_n) - b) < 0 \quad \text{and} \quad \pi(a_{n+1}) \cdot (r(a_{n+1}) - b) > 0.$$

Consequently, for no value of κ , Eqs. 23 could be satisfied, and those equations do not possess any fixed point.

¹²In contrast, for the true gradient bandit update $\nabla_{\theta} \mathcal{J} = \boldsymbol{\pi} \odot (\mathbf{r} - r_{\boldsymbol{\pi}})$,

$$\frac{\partial^2 \mathcal{J}}{\partial \theta_c \partial \theta_a} = -\pi(a)\pi(c)[r(a) + r(c) - 2r_{\boldsymbol{\pi}}] + \pi(a)\mathbb{I}(a = c)[r(a) - r_{\boldsymbol{\pi}}] \quad \Rightarrow \quad \frac{\partial^2 \mathcal{J}}{\partial \theta_c \partial \theta_a} = \frac{\partial^2 \mathcal{J}}{\partial \theta_a \partial \theta_c}, \quad \text{for } a \neq c.$$

Case 2 ($r(a) = b$ for at least one action): If the expected reward for any action a is equal to the baseline b , then $\kappa = 0$, and the fixed point solution corresponds to the face of the simplex where the expected reward for each of the actions is equal to b :

$$\sum_{a \in \mathcal{I}_b} \pi(a) = 1, \quad \text{with} \quad \mathcal{I}_b = \{a \mid r(a) = b\}. \quad (24)$$

Case 3 (b is either less than or greater than all of the expected rewards): We assume that either $b > r(a)$ or $b < r(a)$ for all actions $a \in \mathcal{A}$. Now we solve for κ : From Eqs. 23, $\pi(a) = \frac{\kappa}{r(a)-b}$, $\forall a \in \mathcal{A}$, and therefore

$$\sum_a \pi(a) = 1 \quad \Rightarrow \quad \sum_a \frac{\kappa}{r(a)-b} = 1 \quad \Rightarrow \quad \kappa = \left(\sum_a \frac{1}{r(a)-b} \right)^{-1}.$$

Therefore, the fixed point of the biased update is given by

$$\pi(a) = \frac{1}{r(a)-b} \left(\sum_{c \in \mathcal{A}} \frac{1}{r(c)-b} \right)^{-1} \quad \text{for all actions } a \in \mathcal{A}. \quad (25)$$

This completes the proof¹³. □

Before moving forward, we state a result (Lemma 7.10, Lafferty et al., 2008) which will be useful in proving the next theorem.

Lemma 10.2. *Let X be a bounded random variable with bounds u and l , i.e. $l \leq X \leq u$. Then*

$$\mathbb{E} [e^X] \leq e^{(u-l)^2/8 + \mathbb{E}[X]}.$$

To keep the analysis simple, we assume that for all actions a , $b > r(a)$ or $b < r(a)$. Then, depending on how b is set, the fixed point π^* can act either as an attractor or a repeller.

Theorem 2. *Assume that $\pi_t \neq \pi^*$. If the baseline is pessimistic, i.e. $b < r_1$, then for any $\alpha > 0$, the fixed point π^* acts as a repeller: $D_{KL}(\pi^* \parallel \pi_{t+1}) > D_{KL}(\pi^* \parallel \pi_t)$, where D_{KL} denotes the KL-divergence. And if the baseline is optimistic, i.e. $b > r_k$, then given a sufficiently small positive stepsize α , the fixed point π^* acts as an attractor: $D_{KL}(\pi^* \parallel \pi_{t+1}) < D_{KL}(\pi^* \parallel \pi_t)$.*

Proof. We first express π_{t+1} in terms of π_t . If we define

$$\zeta(a) := \alpha[r(a) - b], \quad (26)$$

then for each action a , the action preference would be updated using $\theta_a^{(t+1)} = \theta_a^{(t)} + \pi_t(a)\zeta(a)$. Applying the softmax function on these updated action preferences gives us π_{t+1} :

$$\pi_{t+1}(a) = \frac{e^{\theta_a^{(t)} + \pi_t(a)\zeta(a)}}{\sum_{c \in \mathcal{A}} e^{\theta_c^{(t)} + \pi_t(c)\zeta(c)}} = \frac{e^{\theta_a^{(t)}} \cdot e^{\pi_t(a)\zeta(a)}}{\sum_c e^{\theta_c^{(t)}} \cdot e^{\pi_t(c)\zeta(c)}} = \frac{\pi_t(a) \cdot e^{\pi_t(a)\zeta(a)}}{\sum_c \pi_t(c) \cdot e^{\pi_t(c)\zeta(c)}}. \quad (27)$$

Let us now compute the KL-divergence between the fixed point π^* and the policy at timestep $t+1$:

$$\begin{aligned} D_{KL}(\pi^* \parallel \pi_{t+1}) &:= \sum_{a \in \mathcal{A}} \pi^*(a) \log \frac{\pi^*(a)}{\pi_{t+1}(a)} = \sum_a \pi^*(a) \log \left(\frac{\pi^*(a)}{\pi_t(a)} \cdot \frac{\sum_c \pi_t(c) \cdot e^{\pi_t(c)\zeta(c)}}{e^{\pi_t(a)\zeta(a)}} \right) \\ &= \sum_a \pi^*(a) \log \frac{\pi^*(a)}{\pi_t(a)} + \sum_a \pi^*(a) \log \left(\sum_c \pi_t(c) \cdot e^{\pi_t(c)\zeta(c)} \right) - \sum_a \pi^*(a) \pi_t(a) \zeta(a) \\ &= D_{KL}(\pi^* \parallel \pi_t) + \log \left(\sum_c \pi_t(c) \cdot e^{\pi_t(c)\zeta(c)} \right) - \left(\sum_c \frac{1}{\zeta(c)} \right)^{-1}, \end{aligned} \quad (28)$$

¹³Interestingly, using this style of reasoning to analyze the true expected gradient bandit would not give us any new information. Even if we assume that $g_a = \pi(a) \cdot (r(a) - b) = \kappa$, $\forall a \in \mathcal{A}$, we find that κ is always zero for the true gradient bandit update. To see this, sum these $|\mathcal{A}|$ equations to get

$$\sum_a \pi(a) \cdot (r(a) - r_\pi) = |\mathcal{A}|\kappa \quad \Rightarrow \quad \sum_a \pi(a)r(a) - r_\pi \sum_a \pi(a) = |\mathcal{A}|\kappa \quad \Rightarrow \quad \kappa = 0.$$

where in the first line we used the expression for π_{t+1} from Eq. 27 and in the last line we used the expression $\pi^*(a) = \frac{1}{r(a)-b} \left(\sum_{c \in \mathcal{A}} \frac{1}{r(c)-b} \right)$ from Eq. 25, to simplify the last term. We now consider the two cases (corresponding to an optimistic and a pessimistic baseline) separately.

Case 1 (Pessimistic baseline): With the assumption that $\alpha > 0$ and $b < r(a)$, we immediately see that $\zeta(a) > 0$ for all actions a . Noting that $\zeta(a) > 0$, we apply Cauchy-Schwartz inequality:

$$\begin{aligned} & \left(\sum_c \pi_t(c)^2 \zeta(c) \right) \left(\sum_c \frac{1}{\zeta(c)} \right) > \left(\sum_c \pi_t(c) \sqrt{\zeta(c)} \cdot \frac{1}{\sqrt{\zeta(c)}} \right)^2 = 1 \\ \Rightarrow & \sum_c \pi_t(c)^2 \zeta(c) > \left(\sum_c \frac{1}{\zeta(c)} \right)^{-1}, \end{aligned} \quad (29)$$

where we could use a strict inequality in the first line because for no constant ρ does $\frac{\pi_t(a)}{\zeta(a)} = \rho$ for all actions a ; if this were true, it would imply that $\pi_t = \pi^*$, which in turn would violate the assumption made in the theorem statement. Next, recall that $\log(\cdot)$ is a concave function and then apply Jensen's inequality followed by the previous result to get:

$$\log \left(\sum_c \pi_t(c) \cdot e^{\pi_t(c)\zeta(c)} \right) > \sum_c \pi_t(c) \log \left(e^{\pi_t(c)\zeta(c)} \right) = \sum_c \pi_t(c)^2 \zeta(c) > \left(\sum_c \frac{1}{\zeta(c)} \right)^{-1}. \quad (30)$$

Above inequality combined with Eq. 28 gives us the desired result:

$$D_{\text{KL}}(\pi^* \|\pi_{t+1}) > D_{\text{KL}}(\pi^* \|\pi_t).$$

Case 2 (Optimistic baseline): In contrast to the previous case, the assumptions $\alpha > 0$ and $b > r(a)$, imply that $\zeta(a) < 0$ for all actions a . We will first use Lemma 10.2 to upper bound the term $\log \left(\sum_c \pi_t(c) \cdot e^{\pi_t(c)\zeta(c)} \right)$ and then obtain the condition on α (as given in Eq. 33) to show that the fixed point is an attractor. Consider a random variable $X : \mathcal{A} \rightarrow \mathbb{R}$ defined as $X(c) = \pi_t(c)\zeta(c)$ and distributed according to π_t , i.e. $\mathbb{P}(X = c) = \pi_t(c)$. Then $\mathbb{E}[X] = \sum_c \pi_t(c)^2 \zeta(c)$ and X is bounded between $l := \alpha \min_c \pi_t(c)[r(c) - b]$ and $u := \alpha \max_c \pi_t(c)[r(c) - b]$. Using Lemma 10.2 along with the fact that $\log(\cdot)$ is an increasing function gives us:

$$\log \left(\sum_c \pi_t(c) e^{\pi_t(c)\zeta(c)} \right) \leq \frac{(u-l)^2}{8} + \sum_c \pi_t(c)^2 \zeta(c). \quad (31)$$

Now the question is whether there exists a stepsize $\alpha > 0$, such that the following inequality holds:

$$\frac{(u-l)^2}{8} + \sum_c \pi_t(c)^2 \zeta(c) < \left(\sum_c \frac{1}{\zeta(c)} \right)^{-1} \quad (32)$$

And the answer is yes! To see this, again apply Cauchy-Schwartz inequality, only this time for $\zeta(c) < 0$, to obtain $\sum_c \pi_t(c)^2 \zeta(c) < \left(\sum_c \frac{1}{\zeta(c)} \right)^{-1}$. Therefore, there does exist an $\alpha > 0$ that satisfies Ineq. 32. With Eq. 28, Ineq. 31, and an appropriate condition on α , the intended result that π^* is an attractor readily follows:

$$\begin{aligned} & \log \left(\sum_c \pi_t(c) e^{\pi_t(c)\zeta(c)} \right) \leq \frac{(u-l)^2}{8} + \sum_c \pi_t(c)^2 \zeta(c) < \left(\sum_c \frac{1}{\zeta(c)} \right)^{-1} \\ \Rightarrow & D_{\text{KL}}(\pi^* \|\pi_{t+1}) < D_{\text{KL}}(\pi^* \|\pi_t). \end{aligned}$$

All that remains is to find the expression for α that satisfies Ineq. 32:

$$\alpha < \frac{8}{(\max_c \pi_t(c)[r(c) - b] - \min_d \pi_t(d)[r(d) - b])^2} \left[\left(\sum_c \frac{1}{r(c) - b} \right)^{-1} - \sum_c \pi_t(c)^2 [r(c) - b] \right]. \quad (33)$$

Note that the above upper bound on α depends on the policy at timestep t , the choice of the baseline, and also the complete bandit reward structure (which is not known to the agent a priori). \square

Values of the different key functions from the proof (left) against a changing baseline, and (middle, right) against a changing stepsize.




Figure 9: The verification of the inequalities used in Theorem 2. We plot the key terms used in the proof: $(\sum_c \frac{1}{\zeta(c)})^{-1}$, $\log(\sum_c \pi_t(c) \cdot e^{\pi_t(c)\zeta(c)})$, its lower bound $\sum_c \pi(c)^2 \zeta(c)$, and its upper bound $\frac{(u-l)^2}{8} + \sum_c \pi(c)^2 \zeta(c)$, for a three armed bandit problem with $\mathbf{r} = [1 \ 2 \ 3]^\top$ and the fixed policy $\boldsymbol{\pi} = [0.2 \ 0.1 \ 0.7]^\top$. The left figure shows the variation of the key terms with a changing baseline and a fixed stepsize α . In the middle and the right plots, we fix the baseline to be pessimistic or optimistic and then vary the stepsize α . The rightmost subplot also serves as the concrete example where the maximum stepsize predicted by Eq. 33, for the fixed point to remain attractive given an optimistic baseline, is not tight.

We present an example in Figure 9, which shows that the upper bound on the stepsize, given in Eq. 33, is not tight. This figure also serves as a verification of the above proof.

The above theorem illustrates an important property of the alternate estimator: with an optimistically initialized baseline, the agent is updated towards a more uniform distribution $\boldsymbol{\pi}^*$. For an optimistic baseline, this distribution has a higher probability of picking the action with the maximum reward as compared to other actions; this can be seen from the expression for $\boldsymbol{\pi}^*$ in Lemma 10.1 or Figure 14 (bottom row, second column). Therefore, if the agent were stuck in a sub-optimal corner of the probability simplex, an optimistic baseline would make the policy more uniform and encourage exploration¹⁴. And even though $\boldsymbol{\pi}^*$ is different from the optimal policy, the agent with an alternate estimator can still reach the optimal policy, because as the agent learns and improves its baseline estimate, the alternate PG estimator becomes asymptotically unbiased. And hopefully by this time, the agent has already escaped the saturated policy region.

One might wonder: but why the biased update is more helpful than the true gradient; after all the true gradient, unlike the biased update, doesn't just move the policy towards a point close to the optimal corner, it updates it exactly towards the optimal corner. The reason for this is that the true gradient update can be vanishingly small in saturated regions and therefore take an excessive amount of time to escape saturation (Mei et al., 2020b; Li et al., 2021). Whereas, as we hypothesize, the biased update will have a much larger update magnitude and therefore result in a faster escape. Even though, we don't theoretically show this, Example 2, the policy update plots shown in Figure 14 (bottom row, second column), and our experiments on bandits with an optimistic baseline initializations (see §D.4) do suggest this.

B.4 Example: Optimistic Initialization helps the Alternate Estimator

The following example considers an optimistically initialized baseline b and shows how the alternate update can use this biased baseline for saturated policies.

Example 2. Consider the 3-armed bandit and the policy given in Example 1. Now assume that the agent doesn't have access to r_π and instead uses an optimistic baseline $b = \delta \gg 0$. Again, the agent will sample $A = a_0$ at

¹⁴On the flip side, with a pessimistically initialized baseline, the alternate estimator can pre-maturely saturate towards a sub-optimal corner; see Figure 14 (bottom right).




Figure 10: The probability simplex for a three armed bandit problem. Each point on the simplex corresponds to a policy. The color map shown is representative of the policy weight for the corresponding action. For instance, in regions that are more blue, the agent’s policy has large probability mass on the the action a_2 . The reward structure for this problem is $\mathbf{r} = [1 \ 2 \ 3]^\top$ with the individual components corresponding to the expected reward on the right, left, and the top corners respectively. We also highlight one corner of the simplex, and the saturated policy region near that corner. In our experiments, we focus on these regions neighbouring the corners, i.e. we focus on policies that are quite saturated but still have non-zero components for all the actions.

each timestep. Therefore:

$$\begin{aligned} \nabla_{\boldsymbol{\theta}} \mathcal{J}_{\pi} &= \boldsymbol{\pi} \odot (\mathbf{r} - r_{\pi} \mathbf{1}) = [1 \ 0 \ 0]^\top \odot [0 \ 0 \ 1]^\top = [0 \ 0 \ 0]^\top, \\ \hat{\mathbf{g}}^{REG}(A, R) &= (R - b)(\mathbf{e}_A - \boldsymbol{\pi}) = (\epsilon - \delta) \cdot ([1 \ 0 \ 0]^\top - [1 \ 0 \ 0]^\top) = [0 \ 0 \ 0]^\top, \text{ and} \\ \hat{\mathbf{g}}^{ALT}(A, R) &= (R - b) \mathbf{e}_A = (\epsilon - \delta) \cdot [1 \ 0 \ 0]^\top = [(\epsilon - \delta) \ 0 \ 0]^\top. \end{aligned}$$

As this example shows, with a suitable baseline (in this case a positive baseline that can override the reward noise), the alternate estimator can use the signal from the baseline to escape the saturated policy region, without relying solely on the reward noise. We formalize these two points in the sections ahead.

B.5 Visualizing the Stochastic PG Estimators for Bandits

In this section, we study the properties of the **stochastic** regular and the alternate estimators. In the previous sections we primarily looked at the expected policy updates, and now we will focus on the sample based updates. We do so by plotting the stochastic policy gradient updates and the variance corresponding to the two estimators for different 3-armed bandit problems and various policies on the probability simplex (see Figure 10). These visualizations also illustrate the properties discussed in the previous sections.

The policy-update plots were created by drawing a vector from the initial policy to the updated policy after one gradient ascent step using the corresponding PG estimator with a fixed stepsize¹⁵. For each estimator, we show four different plots: three for the stochastic estimator, each corresponding to one of the sampled actions (a_0 , a_1 , or a_2); and the fourth one corresponding to the updates made by expected gradient (which itself can be the true policy gradient or the biased one). While calculating the estimator, we also added a noise of constant (unit) magnitude to the expected reward. The blue and grey arrows correspond to the policy change direction for positive and negative reward noises respectively. For clarity, let us illustrate this process with an example. Consider the bandit problem with $\mathcal{A} = \{a_0, a_1, a_2\}$, the reward structure $\mathbf{r} = [0 \ 0 \ 1]^\top$, and a positive reward noise. Let the old policy be $\boldsymbol{\pi}^{\text{old}} = [1/3 \ 1/3 \ 1/3]^\top$ which will be updated using the stochastic regular estimator without baseline with the sampled action being a_2 and a stepsize of $\alpha = 0.4$. The action preferences corresponding to this policy can be computed by taking $\log(\cdot)$ of the action probabilities. This gives us $\boldsymbol{\theta}^{\text{old}} = [\log(1/3) \ \log(1/3) \ \log(1/3)]^\top$. Next, we calculate the gradient update using the regular estimator without the baseline: $\hat{\mathbf{g}}(R, A) = (R - 0) \cdot (\mathbf{e}_A - \boldsymbol{\pi}) = (r(a_2) + 1 - 0) \cdot ([0 \ 0 \ 1]^\top - [1/3 \ 1/3 \ 1/3]^\top) = 2 \cdot [-1/3 \ -1/3 \ 2/3]^\top$. We then update the action preferences: $\boldsymbol{\theta}^{\text{new}} = \boldsymbol{\theta}^{\text{old}} + \alpha \hat{\mathbf{g}}(R, A)$ and in the end take the softmax of these action preferences to obtain the updated policy $\boldsymbol{\pi}^{\text{new}}$. Finally, we plot a blue colored vector (since a positive noise was added to the reward) from $\boldsymbol{\pi}^{\text{old}}$ to $\boldsymbol{\pi}^{\text{new}}$ on the probability simplex. We repeat this process for different initial policies to obtain the final plots (for instance, see Figure 11).

¹⁵We chose stepsizes that resulted in a visually appropriate arrow length. To make the arrow lengths comparable across the estimators, the stepsize was kept same for the regular and the alternate estimators in the same setting (such as for the same choice of the baseline or the bandit reward structure).

The **variance plots** were created by plotting the elementwise variance for the different stochastic PG estimators on the probability simplex. Each hexagonal point on the simplex corresponds to a particular policy value and the color specifies the variance of a particular component of the PG estimator. The variance was calculated using Eq. 20. The top three rows show the variance for the three components and the last row shows the sum of variance along the three components.

Bandit Problem with $\mathbf{r} = [0 \ 0 \ 1]^\top$ and Non-zero Reward Noise

Figure 11 shows the policy update plots and the variance heatmaps for the two estimators on a 3-armed bandit problem with the reward structure $\mathbf{r} = [0 \ 0 \ 1]^\top$ and reward noise $\epsilon \sim \mathcal{N}(0, 1)$. Each column corresponds to the updates and the variance for either the regular estimator with true r_π , the regular estimator without baseline ($b = 0$), the alternate with true r_π , or the alternate estimator without baseline ($b = 0$). We include the zero baseline case because in our experiments, when we learn the baseline, we often initialize it to zero.




Figure 11: Policy updates and the variance plots for the PG estimators on a 3-armed bandit problem with $\mathbf{r} = [0 \ 0 \ 1]^\top$ and a non-zero reward noise. **(Update Plots)** The top three rows show the policy updates corresponding to the different stochastic estimators upon taking a particular action. Last row shows the expected update. Since, the alternate estimator with true r_π and both the regular estimators result in the same expected update, we show only a single plot for these three settings. **(Variance Plots)** The first three rows depict the respective component of the variance vector (of the different gradient estimators) and the last row shows the total variance.

From the update plots, we see that the stochastic alternate estimator has much higher update magnitudes than the regular version. For instance, consider a policy saturated say at action a_0 , i.e. it is near the vicinity of the **right** corner. With the alternate gradient estimator, there is a chance that the policy will move towards the center of the simplex, i.e. it will be updated towards a more uniform distribution. Whereas, with the regular estimator, the agent will likely remain stuck. This observation is also reconciled by the expressions for the estimators: the term $\mathbf{e}_A - \boldsymbol{\pi}$ in the regular estimator becomes very small as $\boldsymbol{\pi}$ saturates on action A ; whereas the alternate estimator only has the term \mathbf{e}_A and its magnitude remains the same irrespective of the saturation. As discussed earlier, the alternate estimator with no baseline is biased. However, its ascent direction points in the same direction as the true gradient update. From the variance plots, we notice that the variance at the corner for both the regular estimators is zero, and for the alternate estimators is non-zero.

So near a corner of the simplex, the alternate estimator utilizes the reward noise to essentially perform a random walk. At sub-optimal corners, this random movement is good since it might make the policy more uniform, from where it is easier to move towards the optimal corner. And at the optimal corners, this random movement is not particularly bad either because optimal corners represent stable equilibrium points: in expectation the gradient signal towards them.




Figure 12: Policy updates and the variance for PG estimators on a 3-armed bandit problem with $\mathbf{r} = [0 \ 0 \ 1]^\top$ and no reward noise.

Bandit Problem with $\mathbf{r} = [0 \ 0 \ 1]^\top$ and Zero Reward Noise

Now we study how these estimators behave when we take away the noise from the rewards in the bandit problem. Figure 12 shows how the policy is updated and the variance of the gradient estimators for a bandit problem with no reward noise. From these plots, we see that the benefit of the alternate estimator (large update magnitudes at the corners) disappears. With neither any reward signal nor any noise, the alternate estimator cannot escape the corners (see the experiments in §D.2). However, the alternate estimator still has a higher update magnitude as compared to the regular estimator near the corners. From the variance plots, we see that, as shown by Proposition 9, the variance for all the estimators is zero at the corners; no reward noise means no variance.

Bandit Problem with $\mathbf{r} = [1 \ 2 \ 3]^\top$ and Non-zero Reward Noise

In the previous settings, there was no reward signal at the sub-optimal corners. Now we consider a case in which there is a non-zero reward signal at the sub-optimal corners and the noise is comparable to it. From Figure 13 (left), we see that the alternate estimator (with true r_π) works better as compared to the regular estimator: it has high noise at sub-optimal corners which might make the policy more uniform. However, with $b = 0$, this is no longer true. In fact, with the alternate estimator, we are driven into the sub-optimal corner faster than the regular without the baseline. Since $b = 0$ is pessimistic for the reward structure $\mathbf{r} = [1 \ 2 \ 3]^\top$, the policy will move away from the center of the simplex (Theorem 2). So initially, the alternate estimator will drive the policy into the corner. Then as its baseline estimate improves and becomes closer to r_π , there is some chance that the agent will escape saturation (although our experiments in Figure 21 show that the alternate estimator with a pessimistic baseline is often inferior to the regular estimator).

In Figure 13 (right), the top row shows the variance of the policy against its entropy, and the bottom row shows the variance of the particular estimator against the policy’s distance (measured using KL divergence) from that corner, i.e. for each action a , we plot $V[g_a]$ against $D_{\text{KL}}(\mathbf{e}_a \parallel \pi)$. The results are similar as the previous variance plots. These entropy figures, unlike the variance heatmaps, can also be plotted for bandit problems with more than three actions.

Bandit Problem with $\mathbf{r} = [1 \ 2 \ 3]^\top$ and an Optimistic or a Pessimistic Baseline

In Figure 14, we show the policy update plots for optimistic or pessimistic baselines. We fix the reward structure to $\mathbf{r} = [1 \ 2 \ 3]^\top$ and do not add any reward noise (the results with non-zero reward noise were effectively similar to these ones, and therefore we skipped them). For an optimistic baseline, we see that both the regular and the alternate stochastic estimators make the policy more uniform. The alternate estimator again has higher update magnitudes compared to the regular estimator. The last row illustrates the point made by Theorem 2, that the alternate estimator with an optimistic baseline attracts the policy to a point on the simplex. Further,

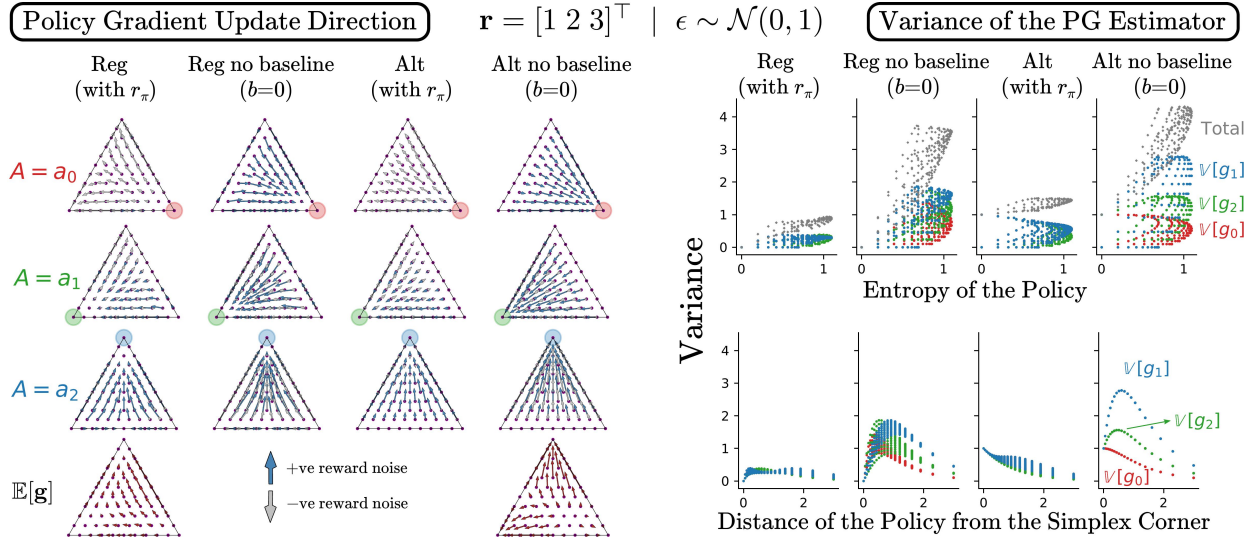


Figure 13: Policy updates and the entropy plots for PG estimators on a 3-armed bandit problem with $\mathbf{r} = [1 \ 2 \ 3]^\top$ and a non-zero reward noise.

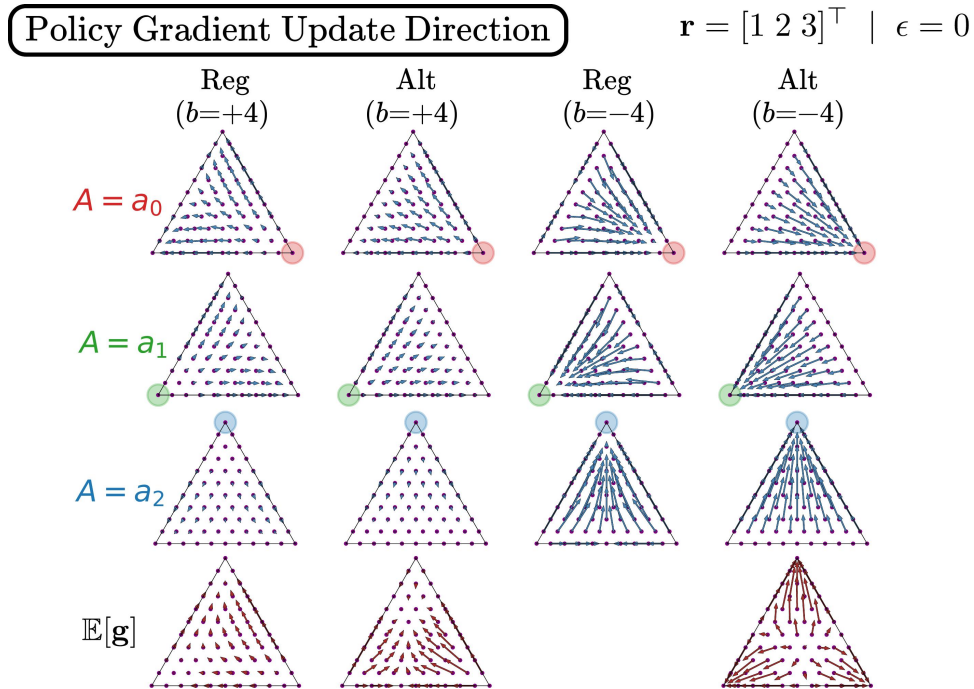


Figure 14: Policy updates for the bandit problem $\mathbf{r} = [1 \ 2 \ 3]^\top$ and no reward noise. The regular and the alternate estimators use an optimistic ($b = +4$) or pessimistic ($b = -4$) baseline. As discussed in §3, both the regular estimators are unbiased, whereas both the alternate estimators are biased.

this attractive signal, as seen from the figure, is much stronger than the expected gradient signal offered by the regular estimator. The expected regular estimator moves the policy towards the optimal corner albeit with very small update magnitudes, especially at the corner. The trends are similar for the pessimistic baseline, except that the policy is updated aggressively towards a corner. The variance plots for this settings looked very similar to those from the previous settings, and therefore we did not include them.

C Softmax PG Estimators for MDPs

In this section, we begin by quickly giving a remark about the regular estimator in the MDP setting. Then we state how to adapt the alternate PG estimator to various PG algorithms (such as PPO), and end by giving implementation details for linear online Actor-Critic algorithm.

C.1 Remark about the Regular Softmax PG Estimator for MDPs

We first show that the regular gradient estimator $\mathbf{g}^{\text{REG}}(S, A)$ (derived in §5) despite the complicated expression is exactly equal to the familiar PG estimator, i.e.

$$\mathbf{g}^{\text{REG}}(S, A) = \nabla_{\mathbf{w}} \log \pi(A|S)(q_{\pi}(S, A) - v_{\pi}(S)). \quad (34)$$

To see this equivalence, note that for a softmax policy π and the state-action pair (S, A) ,

$$\begin{aligned} \nabla_{\mathbf{w}} \log \pi(A|S) &= \nabla_{\mathbf{w}} [\boldsymbol{\theta}(S)]_A - \nabla_{\mathbf{w}} \log \left(\sum_a e^{[\boldsymbol{\theta}(S)]_a} \right) = \nabla_{\mathbf{w}} [\boldsymbol{\theta}(S)]_A - \frac{\sum_a e^{[\boldsymbol{\theta}(S)]_a} \nabla_{\mathbf{w}} [\boldsymbol{\theta}(S)]_a}{\sum_b e^{[\boldsymbol{\theta}(S)]_b}} \\ &= \nabla_{\mathbf{w}} [\boldsymbol{\theta}(S)]_A - \sum_a \pi(a|S) \nabla_{\mathbf{w}} [\boldsymbol{\theta}(S)]_a. \end{aligned}$$

As the reader might observe, therefore the calculation in §5 merely resulted in a different expression for the already familiar PG estimator; it did not yield a new one and therefore was little more than an exercise in manipulating the vector variables $\boldsymbol{\pi}(\cdot|S)$ and $\boldsymbol{\theta}(S)$.

C.2 Adapting the Alternate Estimator to various PG Algorithms

The form for the alternate estimator given by Eq. 13 needs to be adapted to a concrete policy gradient algorithm before it can be used by an RL agent. In this section, we demonstrate this adaption process for the Proximal Policy Optimization algorithm (abbreviated as PPO, Schulman et al., 2017). After that, we directly give the expressions for the alternate estimator adapted to various PG algorithms such as REINFORCE (Williams, 1991), Online Actor-Critic (Sutton and Barto, 2018), and Trust Region Policy Optimization (abbreviated as TRPO, Schulman et al., 2015).

The Alternate PPO Estimator

PPO proposes the following clipped objective that the agent maximizes at each step:

$$\mathcal{J}_{\text{PPO}} := \sum_s \nu_{\pi_{\text{old}}}(s) \sum_a \pi_{\text{old}}(a|s) \cdot \min \left(\frac{\pi_{\mathbf{w}}(a|s)}{\pi_{\text{old}}(a|s)} h_{\pi_{\text{old}}}(s, a), \text{clip} \left[\frac{\pi_{\mathbf{w}}(a|s)}{\pi_{\text{old}}(a|s)}, 1 - \epsilon, 1 + \epsilon \right] h_{\pi_{\text{old}}}(s, a) \right), \quad (35)$$

where $h_{\pi_{\text{old}}}(s, a) := q_{\pi_{\text{old}}}(s, a) - v_{\pi_{\text{old}}}(s)$ is the advantage function under the old policy π_{old} , and ϵ is a fixed parameter. The PPO objective introduces an importance sampling ratio to reuse the data collected by the old policy to make multiple update steps to the new policy¹⁶. Further, the agent only makes updates for those samples which satisfy a particular condition that we define shortly. Using this view, the gradient of the above objective can be equivalently written as

$$\begin{aligned} \nabla_{\mathbf{w}} \mathcal{J}_{\text{PPO}} &= \sum_s \nu_{\pi_{\text{old}}}(s) \sum_a \pi_{\text{old}}(a|s) \cdot \mathbb{I} \left(\text{cond}_{\pi_{\text{old}}}^{\pi_{\mathbf{w}}}(s, a) \right) \frac{\nabla_{\mathbf{w}} \pi_{\mathbf{w}}(a|s)}{\pi_{\text{old}}(a|s)} h_{\pi_{\text{old}}}(s, a) \\ &= \sum_s \nu_{\pi_{\text{old}}}(s) \sum_a \nabla_{\mathbf{w}} \pi_{\mathbf{w}}(a|s) h_{\pi_{\text{old}}}(s, a) \mathbb{I} \left(\text{cond}_{\pi_{\text{old}}}^{\pi_{\mathbf{w}}}(s, a) \right), \end{aligned} \quad (36)$$

where $\text{cond}()$ is the Boolean condition defined as

$$\text{cond}_{\pi_{\text{old}}}^{\pi_{\mathbf{w}}}(s, a) := \left(h_{\pi_{\text{old}}}(s, a) > 0 \wedge \frac{\pi_{\mathbf{w}}(a|s)}{\pi_{\text{old}}(a|s)} < 1 + \epsilon \right) \vee \left(h_{\pi_{\text{old}}}(s, a) < 0 \wedge \frac{\pi_{\mathbf{w}}(a|s)}{\pi_{\text{old}}(a|s)} > 1 - \epsilon \right).$$

¹⁶Despite this reuse of the data collected by the old policy, PPO is still very much an on-policy PG algorithm since it does not correct for the state distribution mismatch between π_{old} and the current policy: $\nu_{\pi_{\text{old}}}(s) \neq \nu_{\pi_{\mathbf{w}}}(s)$.

To reduce clutter, let us define $\tilde{h}_{\pi_{\text{old}}}^{\pi_{\mathbf{w}}}(s, a) := h_{\pi_{\text{old}}}(s, a)\mathbb{I}\left(\text{cond}_{\pi_{\text{old}}}^{\pi_{\mathbf{w}}}(s, a)\right)$. Then the regular PPO estimator can be written as

$$\nabla_{\mathbf{w}}\mathcal{J}_{\text{PPO}} = \mathbb{E}_{S \sim \nu_{\pi_{\text{old}}}; A \sim \pi_{\text{old}}(\cdot|S)} \underbrace{\left[\nabla_{\mathbf{w}} \log \pi_{\mathbf{w}}(A|S) \frac{\pi_{\mathbf{w}}(A|S)}{\pi_{\text{old}}(A|S)} \tilde{h}_{\pi_{\text{old}}}^{\pi_{\mathbf{w}}}(S, A) \right]}_{=: \mathbf{g}^{\text{REG-PPO}}(S, A)}.$$

Now we follow the same procedure that we used to obtain the alternate estimator in Eq. 12, but for the expression that appeared in the PPO gradient (Eq. 36):

$$\sum_a \nabla_{\mathbf{w}} \pi_{\mathbf{w}}(a|s) \tilde{h}_{\pi_{\text{old}}}^{\pi_{\mathbf{w}}}(s, a) = \nabla_{\mathbf{w}} \boldsymbol{\theta}(s) \sum_a \pi_{\mathbf{w}}(a|s) \left[\tilde{h}_{\pi_{\text{old}}}^{\pi_{\mathbf{w}}}(s, a) - \sum_c \pi_{\mathbf{w}}(c|s) \tilde{h}_{\pi_{\text{old}}}^{\pi_{\mathbf{w}}}(s, c) \right] \mathbf{e}_a.$$

Putting the above expression into Eq. 36, we get the alternate estimator for PPO:

$$\begin{aligned} & \nabla_{\mathbf{w}} \mathcal{J}_{\text{PPO}} \\ &= \sum_s \nu_{\pi_{\text{old}}}(s) \sum_a \pi_{\text{old}}(a|s) \nabla_{\mathbf{w}} [\boldsymbol{\theta}(s)]_a \frac{\pi_{\mathbf{w}}(a|s)}{\pi_{\text{old}}(a|s)} \left[\tilde{h}_{\pi_{\text{old}}}^{\pi_{\mathbf{w}}}(s, a) - \sum_c \pi_{\mathbf{w}}(c|s) \tilde{h}_{\pi_{\text{old}}}^{\pi_{\mathbf{w}}}(s, c) \right] \\ &= \mathbb{E}_{S \sim \nu_{\pi_{\text{old}}}; A \sim \pi_{\text{old}}(\cdot|S)} \underbrace{\left[\nabla_{\mathbf{w}} [\boldsymbol{\theta}(S)]_A \frac{\pi_{\mathbf{w}}(A|S)}{\pi_{\text{old}}(A|S)} \left(\tilde{h}_{\pi_{\text{old}}}^{\pi_{\mathbf{w}}}(S, A) - \sum_c \pi_{\mathbf{w}}(c|S) \tilde{h}_{\pi_{\text{old}}}^{\pi_{\mathbf{w}}}(S, c) \right) \right]}_{=: \mathbf{g}^{\text{ALT-PPO}}(S, A)}. \end{aligned} \quad (37)$$

Using a process similar to that demonstrated here, we can extend the alternate estimator to different PG algorithms. We now summarize these estimators.

Alternate REINFORCE Estimator

The regular REINFORCE estimator is

$$\nabla \mathcal{J}_{\pi} = \mathbb{E}_{H_{\infty} \sim \mathbb{P}_{\pi}} \underbrace{\left[\sum_{t=0}^{\infty} \gamma^t \nabla \log \pi(A_t|S_t) (G_t - v_{\pi}(S_t)) \right]}_{=: \mathbf{g}^{\text{REG-REINF}}(H_{\infty})}. \quad (38)$$

And the alternate estimator is given by

$$\nabla_{\mathbf{w}} \mathcal{J}_{\pi} = \mathbb{E}_{H_{\infty} \sim \mathbb{P}_{\pi}} \underbrace{\left[\sum_{t=0}^{\infty} \gamma^t \nabla_{\mathbf{w}} [\boldsymbol{\theta}(S_t)]_{A_t} (G_t - v_{\pi}(S_t)) \right]}_{=: \mathbf{g}^{\text{ALT-REINF}}(H_{\infty})}. \quad (39)$$

Alternate Online AC Estimator

The regular estimator is given by

$$\nabla \mathcal{J}_{\pi} \propto \mathbb{E}_{S, K \sim d_{\pi}; A \sim \pi(\cdot|S); S', R \sim p(\cdot, \cdot|S, A)} \underbrace{\left[\gamma^K \nabla \log \pi(A|S) (R + \gamma v_{\pi}(S') - v_{\pi}(S)) \right]}_{=: \mathbf{g}^{\text{REG-OAC}}(S, A, R, S', K)}, \quad (40)$$

and the alternate estimator is

$$\nabla_{\mathbf{w}} \mathcal{J}_{\pi} \propto \mathbb{E}_{S, K \sim d_{\pi}; A \sim \pi(\cdot|S); S', R \sim p(\cdot, \cdot|S, A)} \underbrace{\left[\gamma^K \nabla_{\mathbf{w}} [\boldsymbol{\theta}(S)]_A (R + \gamma v_{\pi}(S') - v_{\pi}(S)) \right]}_{=: \mathbf{g}^{\text{ALT-OAC}}(S, A, R, S', K)}. \quad (41)$$

TRPO Estimators

TRPO (Eq. 14, Schulman et al., 2015) proposes that the agent solve the following constrained optimization problem at each timestep

$$\max_{\mathbf{w}} \underbrace{\sum_s \nu_{\pi_{\text{old}}}(s) \sum_a \pi_{\mathbf{w}}(a|s) q_{\pi_{\text{old}}}(s, a)}_{=:\mathcal{J}_{\text{TRPO}}} \quad \text{subject to} \quad \sum_s \nu_{\pi_{\text{old}}}(s) \cdot D_{\text{KL}}(\pi_{\text{old}}(\cdot|s) \parallel \pi_{\mathbf{w}}(\cdot|s)) \leq \delta,$$

where δ is a fixed parameter.

We can re-write the gradient of the objective $\mathcal{J}_{\text{TRPO}}$ as follows

$$\nabla_{\mathbf{w}} \mathcal{J}_{\text{TRPO}} = \mathbb{E}_{S \sim \nu_{\pi_{\text{old}}}; A \sim \pi_{\text{old}}(\cdot|S)} \left[\underbrace{\nabla_{\mathbf{w}} \log \pi_{\mathbf{w}}(A|S) \frac{\pi_{\mathbf{w}}(A|S)}{\pi_{\text{old}}(A|S)} q_{\pi_{\text{old}}}(S, A)}_{=:\mathbf{g}^{\text{REG-TRPO}}(S, A)} \right].$$

Then, it is straightforward to see that the alternate estimator would become

$$\nabla_{\mathbf{w}} \mathcal{J}_{\text{TRPO}} = \mathbb{E}_{S \sim \nu_{\pi_{\text{old}}}; A \sim \pi_{\text{old}}(\cdot|S)} \left[\underbrace{\nabla_{\mathbf{w}} [\boldsymbol{\theta}(S)]_A \frac{\pi_{\mathbf{w}}(A|S)}{\pi_{\text{old}}(A|S)} \left(q_{\pi_{\text{old}}}(S, A) - \sum_c \pi_{\mathbf{w}}(c|S) q_{\pi_{\text{old}}}(S, c) \right)}_{=:\mathbf{g}^{\text{ALT-TRPO}}(S, A)} \right]. \quad (42)$$

The alternate estimators proposed here for TRPO and PPO require the computation of the expectation of $q_{\pi_{\text{old}}}$ and $\tilde{h}_{\pi_{\text{old}}}^{\pi_{\mathbf{w}}}$ respectively, under the current policy $\pi_{\mathbf{w}}$ and across the action space. However, if the number of actions is large, this operation becomes prohibitively expensive¹⁷. In such cases, techniques (also see Ciosek and Whiteson, 2020) such as maintaining another function approximator, analogous to a value function, or sampling a subset of actions to approximate this expectation can be used.

C.3 Additional Details for Implementing Online Linear AC

In this section, we provide details that can be used when implementing a PG agent (we used these for the experiments presented in §F) with the estimators presented above. In particular, we calculate the term $\nabla_{\mathbf{w}} \boldsymbol{\theta}(S)$ (see the equations in the previous section) to yield analytical policy gradient update expressions. To keep things concrete, we will focus on the expected policy gradient expression, and the regular and the alternate online AC estimators for softmax policies with linear function approximation. These details can be extended in a straightforward way to work with other PG algorithms such as REINFORCE and different policy parameterizations such as the tabular policies.

Recall that the softmax policy is given by $\pi(a|s) = \frac{e^{[\boldsymbol{\theta}(s)]_a}}{\sum_{c \in \mathcal{A}} e^{[\boldsymbol{\theta}(s)]_c}}$. For linear function approximation and d -dimensional state features, $[\boldsymbol{\theta}(s)]_a := \mathbf{w}_a^\top \mathbf{x}(s)$, with $\mathbf{w}_a \in \mathbb{R}^d$ being the weight vector corresponding to action a , and $\mathbf{x}(s) \in \mathbb{R}^d$ being the feature vector for state s (obtained, say, via tile-coding). Also let $\mathbf{W} := [\mathbf{w}_1 \ \mathbf{w}_2 \ \cdots \ \mathbf{w}_{|\mathcal{A}|}]$ be a $d \times |\mathcal{A}|$ weight matrix. Further, overload the gradient notation, so that for a scalar function $\lambda(\mathbf{W})$, $\nabla_{\mathbf{W}} \lambda$ is a $d \times |\mathcal{A}|$ matrix defined as $[\nabla_{\mathbf{W}} \lambda]_{ij} = \partial \lambda / \partial w_{ij}$. Using this definition, $\nabla_{\mathbf{W}} \lambda = [\nabla_{\mathbf{w}_1} \lambda \ \nabla_{\mathbf{w}_2} \lambda \ \cdots \ \nabla_{\mathbf{w}_{|\mathcal{A}|}} \lambda]$, which in turn means that we can first calculate $\nabla_{\mathbf{w}_c} \mathcal{J}_\pi$ for each action $c \in \mathcal{A}$ and then stack these vectors to obtain $\nabla_{\mathbf{W}} \mathcal{J}_\pi$. The gradient update, for the weights \mathbf{W} , can be written as $\mathbf{W}^{\text{new}} = \mathbf{W}^{\text{old}} + \alpha \nabla_{\mathbf{W}} \mathcal{J}_\pi$ for some $\alpha > 0$. Using this formulation, let us now compute the expressions for the expected policy gradient, the regular estimator, and the alternate estimator.

First note that, using Proposition 3, for an arbitrary state s , and actions a and c we get that

$$\nabla_{\mathbf{w}_c} [\boldsymbol{\theta}(s)]_a = \nabla_{\mathbf{w}_c} (\mathbf{w}_a^\top \mathbf{x}(s)) = \mathbb{I}(c = a) \mathbf{x}(s). \quad (43)$$

¹⁷Moreover, if this expectation is calculated exactly, the alternate (TRPO or PPO) estimator will remain unbiased. Depending on the context this might be a good or a bad property to have. However, in this work, we do not pursue a more involved analysis of the alternate estimators for TRPO and PPO. Also note that the objectives maximized by PPO and TRPO are different from the original PG objective $\mathcal{J}_\pi := \sum_s \nu_{\pi_{\text{old}}}(s) \sum_a \pi_{\mathbf{w}}(a|s) q_{\pi_{\text{old}}}(s, a)$ which would probably make such an analysis more interesting.

Now for computing the expected policy gradient, put Eq. 10 in Eq. 9 to obtain

$$\begin{aligned}
 \nabla_{\mathbf{w}_c} \mathcal{J}_\pi &= \sum_s \nu_\pi(s) [\nabla_{\mathbf{w}_c} \boldsymbol{\theta}(s)] \boldsymbol{\pi}(\cdot|s) \odot (\mathbf{q}_\pi(s, \cdot) - v_\pi(s) \mathbf{1}) \\
 &= \sum_s \nu_\pi(s) \sum_a \nabla_{\mathbf{w}_c} [\boldsymbol{\theta}(s)]_a \pi(a|s) (q_\pi(s, a) - v_\pi(s)) \\
 &= \sum_s \nu_\pi(s) \sum_a \mathbb{I}(c = a) \mathbf{x}(s) \pi(a|s) (q_\pi(s, a) - v_\pi(s)) \\
 &= \sum_s \nu_\pi(s) \mathbf{x}(s) \pi(c|s) (q_\pi(s, c) - v_\pi(s)).
 \end{aligned}$$

Now we stack together the vectors $\nabla_{\mathbf{w}_c} \mathcal{J}_\pi$ for all the actions $c \in \mathcal{A}$ to obtain

$$\nabla_{\mathbf{w}} \mathcal{J}_\pi = \sum_s \nu_\pi(s) \mathbf{x}(s) \left(\boldsymbol{\pi}_{\mathbf{w}}(\cdot|s) \odot (\mathbf{q}_\pi(s, \cdot) - v_\pi(s) \mathbf{1}) \right)^\top. \quad (44)$$

Using chain rule of differentiation for vector variables (Proposition 5), we obtain

$$\begin{aligned}
 \nabla_{\mathbf{w}_c} \pi(a|s) &= [\nabla_{\mathbf{w}_c} \boldsymbol{\theta}(s)] [\nabla_{\boldsymbol{\theta}(s)} \pi(a|s)] = \sum_d \nabla_{\mathbf{w}_c} [\boldsymbol{\theta}(s)]_d \cdot \frac{\partial \pi(a|s)}{\partial [\boldsymbol{\theta}(s)]_d} \\
 &= \sum_d \mathbb{I}(c = d) \mathbf{x}(s) \cdot \pi(a|s) [\mathbb{I}(d = a) - \pi(d|s)] = \pi(a|s) \mathbf{x}(s) [\mathbb{I}(c = a) - \pi(c|s)] \\
 \Rightarrow \nabla_{\mathbf{w}_c} \log \pi(a|s) &= \mathbf{x}(s) [\mathbb{I}(c = a) - \pi(c|s)].
 \end{aligned}$$

Again stacking the vectors $\nabla_{\mathbf{w}_c} \log \pi(a|s)$ for all actions $c \in \mathcal{A}$ gives us

$$\nabla_{\mathbf{w}} \log \pi(a|s) = \mathbf{x}(s) (\mathbf{e}_a - \boldsymbol{\pi}(\cdot|s))^\top. \quad (45)$$

Finally, we can compute the regular estimator for online actor-critic. Using the expression for the regular online AC estimator (given in Eq. 40) for arbitrary $a \in \mathcal{A}$, $r \in \mathcal{R}$, $k \in \{0, 1, \dots\}$, and $s, s' \in \mathcal{S}$ we obtain

$$\begin{aligned}
 \mathbf{g}^{\text{REG-OAC}}(s, a, r, s', k) &:= \gamma^k \nabla_{\mathbf{w}} \log \pi(a|s) (r + \gamma v_\pi(s') - v_\pi(s)) \\
 &= \gamma^k (r + \gamma v_\pi(s') - v_\pi(s)) \mathbf{x}(s) (\mathbf{e}_a - \boldsymbol{\pi}(\cdot|s))^\top.
 \end{aligned} \quad (46)$$

If we stack the vectors $\nabla_{\mathbf{w}_c} [\boldsymbol{\theta}(s)]_a$ from Eq. 43 for different actions c , we would get

$$\nabla_{\mathbf{w}} [\boldsymbol{\theta}(s)]_a = \mathbf{x}(s) \mathbf{e}_a^\top.$$

Then the alternate estimator for online actor-critic (given in Eq. 41) can be computed as

$$\begin{aligned}
 \mathbf{g}^{\text{ALT-OAC}}(s, a, r, s', k) &:= \gamma^k \nabla_{\mathbf{w}} [\boldsymbol{\theta}(s)]_a (r + \gamma v_\pi(s') - v_\pi(s)) \\
 &= \gamma^k (r + \gamma v_\pi(s') - v_\pi(s)) \mathbf{x}(s) \mathbf{e}_a^\top.
 \end{aligned} \quad (47)$$

C.4 Entropy Regularized Policy Gradient

In maximum entropy policy gradient (Ahmed et al., 2019; Haarnoja et al., 2018; Peters et al., 2010), the agent maximizes the entropy regularized objective \mathcal{J}^{ent} defined by:

$$\mathcal{J}^{\text{ent}} := \mathbb{E} \left[\sum_{t=0}^{T-1} \gamma^t \left(R_{t+1} + \tau \mathcal{H}_\pi(S_t) \right) \right], \quad (48)$$

where $\mathcal{H}_\pi(s) = -\sum_{a \in \mathcal{A}} \pi(a|s) \log \pi(a|s)$ is the action-entropy of the policy¹⁸ and τ is the entropy regularization parameter. Note that this objective naturally arises if the agent were to use an entropy augmented reward

¹⁸Further, we explicitly define the entropy function to be zero at the terminal state s_\times , i.e. $\mathcal{H}_\pi(s_\times) := 0$.

$r^{\text{ent}}(R_{t+1}, S_t) = R_{t+1} + \tau \mathcal{H}_\pi(S_t)$. Consequently, we can also define special value functions that track this augmented reward:

$$q_\pi^{\text{ent}}(s, a) := \mathbb{E} \left[\sum_{t=k}^{T-1} \gamma^{t-k} \left(R_{t+1} + \tau \mathcal{H}_\pi(S_t) \right) \middle| S_k = s, A_k = a \right], \text{ and} \quad (49)$$

$$v_\pi^{\text{ent}}(s) := \mathbb{E} \left[\sum_{t=k}^{T-1} \gamma^{t-k} \left(R_{t+1} + \tau \mathcal{H}_\pi(S_t) \right) \middle| S_k = s \right]. \quad (50)$$

Using this objective, the policy gradient can be shown to be equal to

$$\nabla \mathcal{J}^{\text{ent}} = \sum_s \nu_\pi(s) \sum_a \pi(a|s) \left[\nabla \log \pi(a|s) (q_\pi^{\text{ent}}(s, a) - v_\pi^{\text{ent}}(s)) + \tau \nabla \mathcal{H}_\pi(s) \right]. \quad (51)$$

Compare this with the regular policy gradient expressions given in Eqs. 3 and 4. Let us now compute the analytical expression for the gradient of the entropy term:

$$\begin{aligned} \nabla_{\mathbf{w}} \mathcal{H}_\pi(s) &= -\nabla_{\mathbf{w}} \sum_{a \in \mathcal{A}} \pi(a|s) \log \pi(a|s) \\ &= -\sum_a \nabla_{\mathbf{w}} \pi(a|s) \log \pi(a|s) - \sum_a \pi(a|s) \nabla_{\mathbf{w}} \log \pi(a|s) \overset{0}{=} -\sum_a \nabla_{\mathbf{w}} \pi(a|s) \log \pi(a|s) \\ &= -\sum_a \pi(a|s) \nabla_{\mathbf{w}} \log \pi(a|s) \log \pi(a|s) = -\mathbf{x}(s) \sum_a (\mathbf{e}_a - \boldsymbol{\pi}(\cdot|s))^\top \pi(a|s) \log \pi(a|s) \quad (\text{from Eq. 45}) \\ &= -\mathbf{x}(s) \left[\sum_a \mathbf{e}_a \pi(a|s) \log \pi(a|s) - \boldsymbol{\pi}(\cdot|s) \sum_a \pi(a|s) \log \pi(a|s) \right]^\top \\ &= -\mathbf{x}(s) \left[\boldsymbol{\pi}(\cdot|s) \odot \log \boldsymbol{\pi}(\cdot|s) + \mathcal{H}_\pi(s) \boldsymbol{\pi}(\cdot|s) \right]^\top. \end{aligned} \quad (52)$$

With this expression, we can now easily extend the analysis of the previous section to implement, say, entropy regularized online linear actor-critic.

C.5 Escort Transform

The escort transform (Mei et al., 2020b) is a policy parameterization that is theoretically guaranteed to escape policy saturation at a fast rate. An escort policy is given by

$$\pi(a|s) = \frac{|\boldsymbol{\theta}(s)_a|^p}{\sum_{a'} |\boldsymbol{\theta}(s)_{a'}|^p}, \quad (53)$$

where $p \geq 1$ is a tunable parameter. For large p , the escort transform would behave somewhat similarly to softmax, and as $p \rightarrow \infty$, an escort policy would approach the greedy policy. Note that the action preferences for escort transform cannot all simultaneously be equal to zero (also see Theorem 4, Mei et al., 2020b). Let us begin by computing the gradient of this transform with respect to the action preferences (for brevity, we represent $[\boldsymbol{\theta}(s)]_a$ as θ_a in the following calculation):

$$\begin{aligned} \frac{\partial}{\partial \theta_d} \pi(a|s) &= \frac{\partial}{\partial \theta_d} \frac{|\theta_a|^p}{\sum_{a'} |\theta_{a'}|^p} = \frac{\mathbb{I}(a = d) \text{sgn}(\theta_a) \cdot p |\theta_a|^{p-1} \sum_{a'} |\theta_{a'}|^p - \text{sgn}(\theta_d) \cdot p |\theta_d|^{p-1} |\theta_a|^p}{(\sum_{a'} |\theta_{a'}|^p)^2} \\ &= \frac{p}{\sum_{a'} |\theta_{a'}|^p} \left[\mathbb{I}(a = d) \text{sgn}(\theta_a) \cdot |\theta_a|^{p-1} - \text{sgn}(\theta_d) \cdot |\theta_d|^{p-1} \pi(a|s) \right], \end{aligned} \quad (54)$$

where $\text{sgn} : \mathbb{R} \rightarrow \{-1, 0, 1\}$ is the sign function that returns the sign of its input (or zero if the input is zero). Now, we will compute the escort transform analog of Eq. 45:

$$\begin{aligned}
 \nabla_{\mathbf{w}_c} \pi(a|s) &= [\nabla_{\mathbf{w}_c} \boldsymbol{\theta}(s)] [\nabla_{\boldsymbol{\theta}(s)} \pi(a|s)] = \sum_d \nabla_{\mathbf{w}_c} \theta_d \cdot \frac{\partial \pi(a|s)}{\partial \theta_d} \\
 &= \sum_d \mathbb{I}(c = d) \mathbf{x}(s) \cdot \frac{p}{\sum_{a'} |\theta_{a'}|^p} \left[\mathbb{I}(a = d) \text{sgn}(\theta_a) \cdot |\theta_a|^{p-1} - \text{sgn}(\theta_d) \cdot |\theta_d|^{p-1} \pi(a|s) \right] \\
 &= \frac{p \cdot \mathbf{x}(s)}{\sum_{a'} |\theta_{a'}|^p} \left[\mathbb{I}(a = c) \text{sgn}(\theta_a) \cdot |\theta_a|^{p-1} - \text{sgn}(\theta_c) \cdot |\theta_c|^{p-1} \pi(a|s) \right] \\
 \Rightarrow \nabla_{\mathbf{w}_c} \log \pi(a|s) &= p \cdot \mathbf{x}(s) \left[\frac{\mathbb{I}(a = c) \text{sgn}(\theta_a) \cdot |\theta_a|^{p-1}}{|\theta_a|^p} - \text{sgn}(\theta_c) \cdot |\theta_c|^{p-1} \frac{\pi(a|s)}{|\theta_a|^p} \right] \\
 &= p \cdot \mathbf{x}(s) \left[\frac{\mathbb{I}(a = c)}{\theta_a} - \frac{\text{sgn}(\theta_c) \cdot |\theta_c|^{p-1}}{\sum_{a'} |\theta_{a'}|^p} \right].
 \end{aligned}$$

Finally, we stack the vectors $\nabla_{\mathbf{w}_c} \log \pi(a|s)$ for all actions $c \in \mathcal{A}$ to obtain

$$\nabla_{\mathbf{w}} \log \pi(a|s) = p \cdot \mathbf{x}(s) \left[\frac{\mathbf{e}_a}{[\boldsymbol{\theta}(s)]_a} - \frac{\text{sgn}(\boldsymbol{\theta}(s)) \odot |\boldsymbol{\theta}(s)|^{p-1}}{\sum_{a'} [|\boldsymbol{\theta}(s)|_{a'}|^p]} \right]^\top. \quad (55)$$

Using this expression, we can directly implement online linear AC with an escort policy.

D Gradient Bandits: Full Experiments

We then give the experimental details and the algorithm pseudocode used in §4. Finally, we give additional experiments and discuss them in more detail.

D.1 Experimental Details

This section presents multiple experiments with gradient bandits on different 3-armed bandit tasks. All the bandit problems had the action set $\mathcal{A} = \{a_0, a_1, a_2\}$, a fixed reward structure \mathbf{r} , and a normally distributed reward noise ϵ . The experimental results serve as a demonstration and the verification of the properties of the gradient bandits estimators discussed in the main paper.

By default, in all the experiments, we trained five different gradient bandit agents with a softmax policy: for different agents, the policy weights were updated using either the expected gradient $\nabla \mathcal{J} = \boldsymbol{\pi} \odot (\mathbf{r} - r_\pi)$, the regular estimator with true average reward $\mathbf{g}^{\text{REG}}(R, A) = (R - r_\pi)(\mathbf{e}_A - \boldsymbol{\pi})$ or a learned baseline $\hat{\mathbf{g}}^{\text{REG}}(R, A) = (R - b)(\mathbf{e}_A - \boldsymbol{\pi})$, or the alternate estimator with the true average reward $\mathbf{g}^{\text{ALT}}(R, A) = \mathbf{e}_A(R - r_\pi)$ or a learned baseline $\hat{\mathbf{g}}^{\text{ALT}}(R, A) = \mathbf{e}_A(R - b)$. The baseline in each case was learned using a running average: $b_{t+1} = (1 - \beta)b_t + \beta R_t$, with the stepsize β . The exact pseudocode for gradient bandits (used in the experiments of §4) is given in Algorithm 6. The action preferences for the softmax policy were initialized to $\theta_{a_1} = \theta_{a_2} = 0$; and $\theta_{a_0} = 0$ for uniform policy initialization or $\theta_{a_0} \in \{5, 10, 50\}$ for saturated policy initialization depending on the degree of saturation.

Each experiment was run for 1000 timesteps. For each estimator, we performed a sweep over the policy stepsize $\alpha \in \{2^{-6}, 2^{-5}, \dots, 2^1\}$ and the baseline stepsize $\beta \in \{2^{-4}, 2^{-3}, \dots, 2^0\}$. The sensitivity plots show the mean final performance on the last 50 timesteps averaged over 150 runs for different parameter settings. To reduce clutter, we didn't label all the β s in the sensitivity plots.

For both the estimators, we swept the policy stepsize $\alpha \in \{2^{-6}, 2^{-5}, \dots, 2^1\}$ and the reward baseline stepsize $\beta \in \{2^{-4}, 2^{-3}, \dots, 2^0\}$. Although, we do not always label all the β s in the sensitivity plots to reduce the clutter. The policy was initialized to $\theta_{a_1} = \theta_{a_2} = 0$; and $\theta_{a_0} = 0$ (for uniform), $\theta_{a_0} \in \{0, 5, 10, 50\}$ for saturated policies depending on the degree of saturation.

D.2 Alternate Estimator can Escape Saturated Policy Regions

Experiments in this section demonstrate that the alternate estimator performs competitively with the regular estimator for uniform policy initialization and clearly outperforms it in the case of a saturated policy initialization

Algorithm 6 Gradient Bandit with Regular Estimator

Input a softmax policy $\pi_{\theta}(a) := \frac{e^{\theta_a}}{\sum_{b \in \mathcal{A}} e^{\theta_b}}$ and an average reward function estimate \hat{r}_{π}
 Input the policy stepsize α and the average reward stepsize β
 Initialize θ and \hat{r}_{π}
for total number of timesteps **do**
 Sample an action $A \sim \pi$ and a reward $R \sim p(\cdot|A)$
 for each action $a \in \mathcal{A}$ **do**
 # Calculate the gradient estimator
 $\hat{g}_a^{\text{REG}} = (R - \hat{r}_{\pi})(\mathbb{I}(A = a) - \pi(a))$
 # Update the policy parameter
 $\theta_a \leftarrow \theta_a + \alpha \cdot \hat{g}_a^{\text{REG}}$
 # Update the average reward estimate
 $\hat{r}_{\pi} \leftarrow \beta R + (1 - \beta)\hat{r}_{\pi}$

when there is reward noise present. They also show that in the absence of reward noise, the alternate estimator fails to escape the saturated policy regions. The bandit reward structure is $\mathbf{r} = [0 \ 0 \ 1]^{\top}$ with either normal noise $\epsilon \sim \mathcal{N}(0, 1)$ or zero reward noise.

Learning Curves

Figure 15 shows the learning curves for the best parameter configuration based on final performance (see sensitivity plots) for the five different agents using different PG estimators for two settings: uniform policy initialization and a sub-optimally saturated policy initialization. The regular and alternate estimators, used either the true reward function r_{π} or learned a reward baseline that was initialized to zero. For the uniform policy case, we observe that all the methods achieved a good performance. More interestingly, for saturated policy initialization, the alternate estimators (both r_{π} and with a baseline b) learned good behaviors, whereas the expected PG and the regular estimators failed to learn anything. Further, alternate with a learned baseline converged faster than alternate with r_{π} . We attribute this result to a better exploration afforded by the baseline.




Figure 15: Learning curves for gradient bandits. The reward structure is $\mathbf{r} = [0 \ 0 \ 1]^{\top}$ with noise $\epsilon \sim \mathcal{N}(0, 1)$. The estimators without an accompanying value of β use the true r_{π} . The results are averages over 150 independent runs and the shaded regions represent the standard error.

Parameter Sensitivity Plots and Performance with and without Reward Noise

Figure 16 shows the parameter sensitivity of the gradient estimators for different policy saturations and reward noises. The graphs plot the mean final performance¹⁹ during last 50 steps of the 1000 timestep run of each

¹⁹For small stepsizes, the final performance of all the methods are very similar; it is only at higher stepsizes that we start to see a deviation of the performance of both the stochastic PG estimators from each other and the expected gradient bandits update; for a detailed discussion on stochasticity in PG methods, also see Mei et al. (2021).

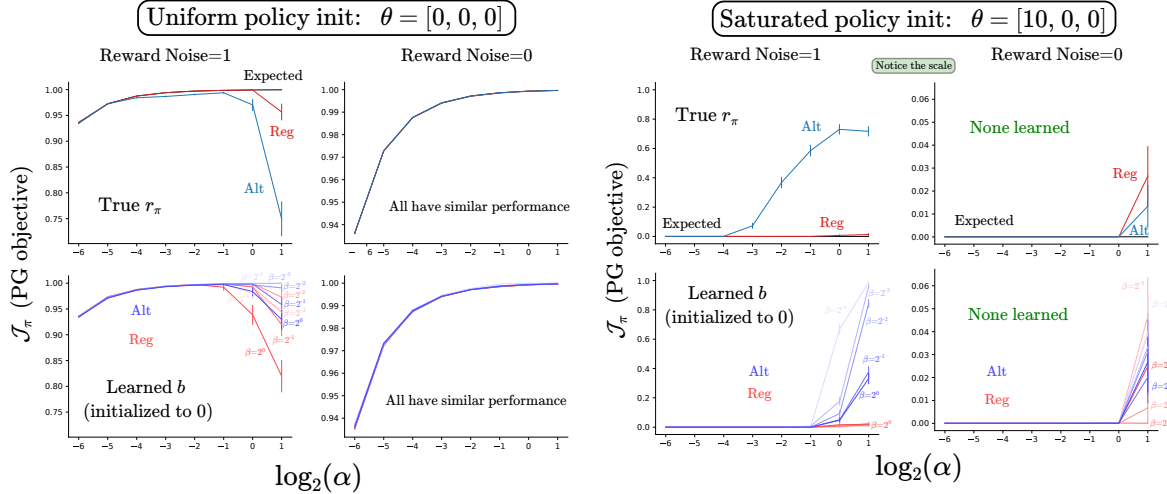


Figure 16: Parameter sensitivity plots for the different estimators on the bandit problem with $\mathbf{r} = [0 \ 0 \ 1]^\top$. We show results for eight different combinations: the estimators using r_π or a learned baseline, uniform or saturated policy initialization, and with or without reward noise. The results were averaged over 150 runs and the bars represent standard error. The x axis shows the policy stepsize α and the curve labels show the reward baseline stepsize β .

method. The performance was measured by analytically calculating $\mathcal{J}_\pi = \mathbb{E}_\pi[R]$. We observe that for uniform initialization, all the methods (irrespective of the magnitude of reward noise) had a similar performance. In contrast, for saturated policies, the alternate estimators with reward noise were able to learn for a range of stepsizes. However, without reward noise, none of the methods escaped the saturation and their policies did not improve. Note that the baseline was initialized to zero which is not optimistic for $\mathbf{r} = [0 \ 0 \ 1]^\top$ and therefore the alternate estimator could not utilize the bias in the baseline (as discussed in §3). This experiment verifies our claim that reward noise helps the alternate estimator to escape the sub-optimal regions in the policy space; without reward noise, the agent is unable to perform a random walk.

Behavior of Action Preferences under Different Estimators

To understand how the action preferences are influenced by different estimators, we plotted their values in Figure 17 (with non-zero reward noise) and Figure 18 (with zero reward noise) using a set of good performing stepsize parameters. We observe that the behavior of the action preferences for the regular estimator with and without reward noise is similar: the preferences increased quite a bit during the first 100 steps and then did not change much, resulting in flat lines. This probably happened because the policy had saturated and the estimator became negligibly small, i.e. with high probability $\mathbf{e}_A \approx \boldsymbol{\pi}$ and $R \approx r_\pi$, and thus $\mathbf{g}^{\text{REG}} = (R - r_\pi)(\mathbf{e}_A - \boldsymbol{\pi}) \approx \mathbf{0}$. In contrast, reward noise affected the behavior of the alternate estimator (with r_π) quite strongly: the graph for θ_2 with reward noise demonstrates the characteristic random walk pattern. Whereas, in the case of no reward noise, this graph becomes smooth and keeps increasing, since even at saturation, because of missing the $-\boldsymbol{\pi}$ term, the alternate estimator $\mathbf{g}^{\text{ALT}} = (R - r_\pi)\mathbf{e}_A$ remains somewhat large. The patterns for alternate with a learned b are much less clear, probably because of additional effects from the bias in the baseline. We also observe that the alternate estimator with reward noise is noisier than the regular variant; this is expected because the random walk style behavior gives it a higher variance than the regular version (refer Figure 11). Alternate without reward noise seems to be less noisy as compared to the regular variant which is again reasonable since in absence of reward noise, alternate has a lower variance than the regular estimator (refer Figure 12).

D.3 Gradient Bandit Estimators with Increasingly Saturated Policy Initializations

In this experiment, we trained the estimators on the bandit task $\mathbf{r} = [1 \ 2 \ 3]^\top$ with noise $\epsilon \sim \mathcal{N}(0, 1)$ and four different policy saturations, ranging from a uniform policy to a policy highly saturated on the worst action. Figure 19 shows the parameter sensitivity plots depicting the final performance of the methods during the last 50 steps of their 1000 timestep run. The results show that as the policy saturation increases, the performance

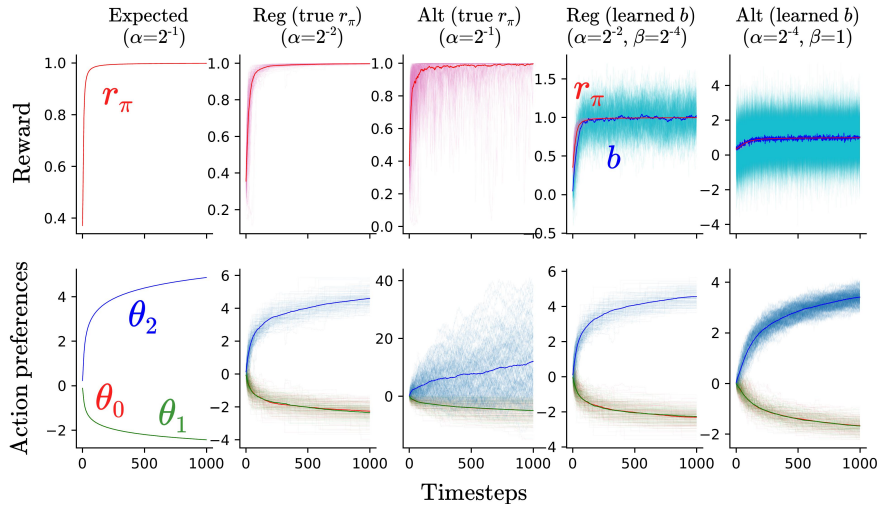


Figure 17: **(Top)** Learning curves for expected reward r_π and value of the baseline b as the policy π is updated using different estimators on the bandit problem $\mathbf{r} = [0 \ 0 \ 1]^\top$ with noise $\epsilon \sim \mathcal{N}(0, 1)$. **(Bottom)** Learning curves for action preferences corresponding to the different estimators. In all cases, the finer lines depict the 150 individual runs and the thick lines show the averages.

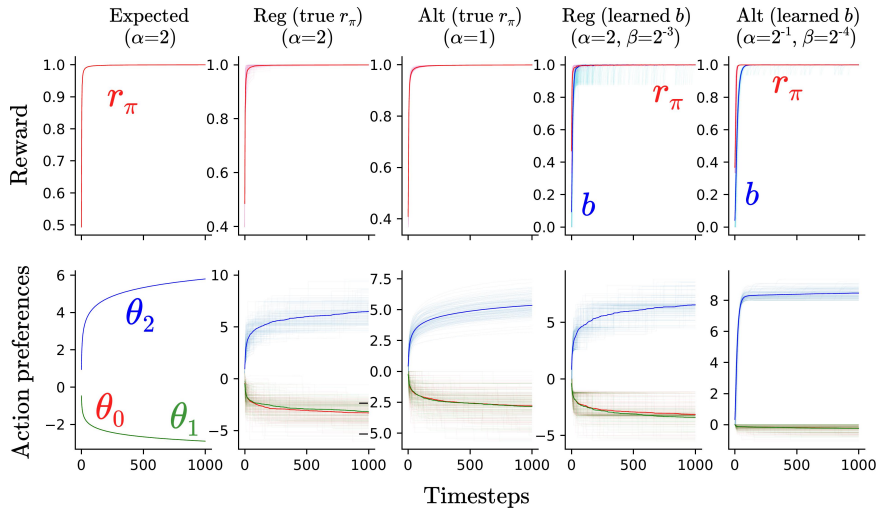


Figure 18: Average reward r_π , value of the learned baseline b , and the action preferences against timesteps for different methods on the bandit problem $\mathbf{r} = [0 \ 0 \ 1]^\top$ and no reward noise.

of all the methods degrade. Further, the alternate estimator (with true r_π) is able to handle high saturations (except for $\theta^{\text{init}} = [5\ 0\ 0]^\top$) better than the regular estimator. Whereas, the alternate estimator with a learned baseline initialized to zero (which is pessimistic for this reward structure) did not perform well. This figure, with a baseline initialized to $b = 0$, also serves as a standard (or a control set) against which the results of the next section can be compared, where we explicitly study the effect of an optimistic ($b = +4$) and a pessimistic ($b = -4$) reward baseline on the two estimators' performance.

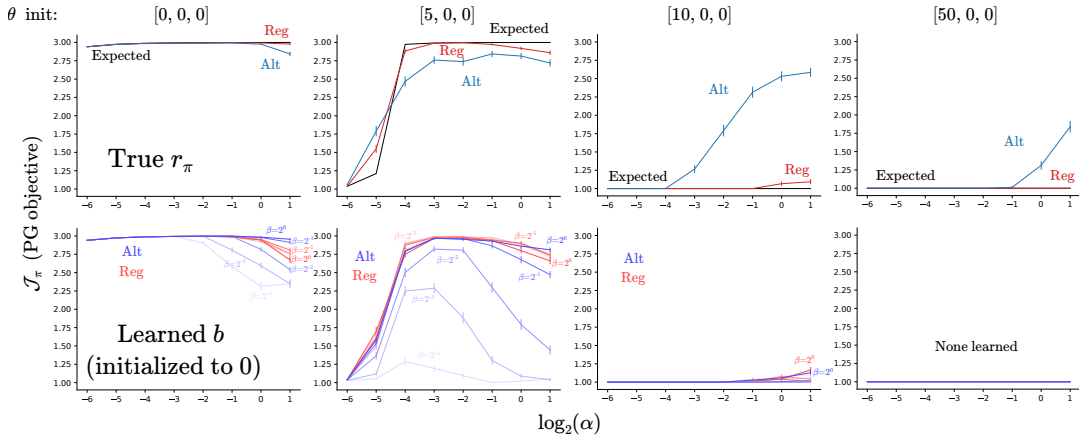


Figure 19: Parameter sensitivity of the PG estimators against increasingly saturated policy initializations on the 3-armed bandit problem $\mathbf{r} = [1\ 2\ 3]^\top$ with $\epsilon \sim \mathcal{N}(0, 1)$. The top row shows the performance of the estimators with the true r_π . Bottom row shows the performance for a learned baseline, initialized to zero. The subplot headers show how the action preferences were initialized.

D.4 Effect of Bias in the Reward Baselines on the Estimators

In this section, we study the effect of initializing the reward baseline optimistically ($b = +4$) or pessimistically ($b = -4$). The reward was set to $\mathbf{r} = [1\ 2\ 3]^\top$ with noise $\epsilon \sim \mathcal{N}(0, 1)$ ²⁰.

Optimistically or Pessimistically Initialized Baseline

Figure 20 and Figure 21 show the sensitivity plots for the regular and alternate estimators for an optimistically and a pessimistically initialized baseline respectively. These results show that the optimistic initialization greatly helped the alternate estimator in escaping saturated policy regions. Whereas, a pessimistic initialization significantly hampered the performance of the alternate estimator, making it worse than the regular estimator even for the uniform policy. In contrast, the baseline initialization affected the regular estimator in a milder way. Also with an optimistic initialization, the alternate estimator preferred smaller stepsizes allowing it to enjoy the optimism effect for longer; and vice-versa for a pessimistic initialization.

These results also verify the fact that the alternate estimator is asymptotically unbiased: the sensitivity plots show that the agent was able to converge to the optimal corner (corresponding to action a_2) and thus obtained a reward of about 3. As the agent's estimate of r_π improved, i.e. the baseline b became closer to the average reward r_π , the bias in the alternate estimator disappeared.

Fixed Baseline

Now we study what happens when we fix the baseline and not learn it, thereby maintaining the biasedness of the alternate estimator. Figure 22 shows the parameter sensitivity plots for this setting for multiple policy initializations. From the results, we see that the regular estimator was not affected by a fixed baseline (since it remains unbiased with any baseline): it was able to learn a good, but not the optimal, policy for a uniform

²⁰The results for the bandit problem with optimistic or pessimistic baselines, and no noise (not shown here) were qualitatively similar to the ones presented here; the signal from the bias in the baseline is quite strong and affects both the estimators much more than the noise in the sampled rewards.

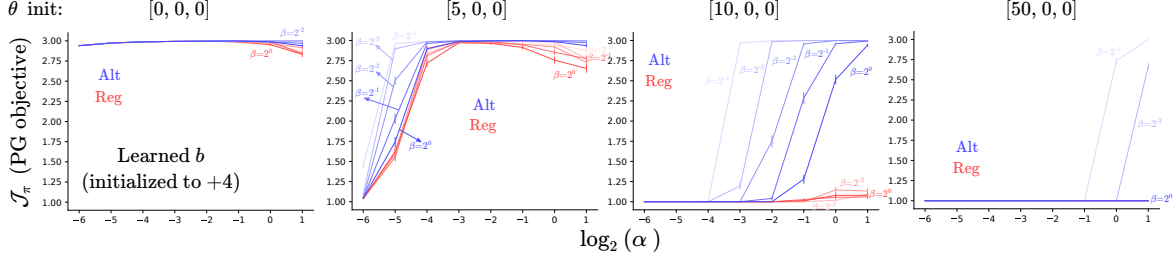


Figure 20: Parameter sensitivity plots for the different estimators for an **optimistically** initialized reward baseline ($b = +4$) on the three armed bandit problem $\mathbf{r} = [1 \ 2 \ 3]^\top$ and $\epsilon \sim \mathcal{N}(0, 1)$.




Figure 21: Parameter sensitivity plots for the different estimators for a **pessimistically** initialized reward baseline ($b = -4$) on the three armed bandit problem $\mathbf{r} = [1 \ 2 \ 3]^\top$ and $\epsilon \sim \mathcal{N}(0, 1)$.

policy initialization. In contrast, the alternate estimator was impacted heavily: for an optimistic baseline, it converged to a value of about 2.4 and for a pessimistic baseline it converged to a value of zero. The values to which the alternate estimator converged remained the same irrespective of the policy initialization. This hints at the fact that the update signal offered by the alternate estimator, due to bias in the critic, is quite strong and can override the high saturation in softmax policies.

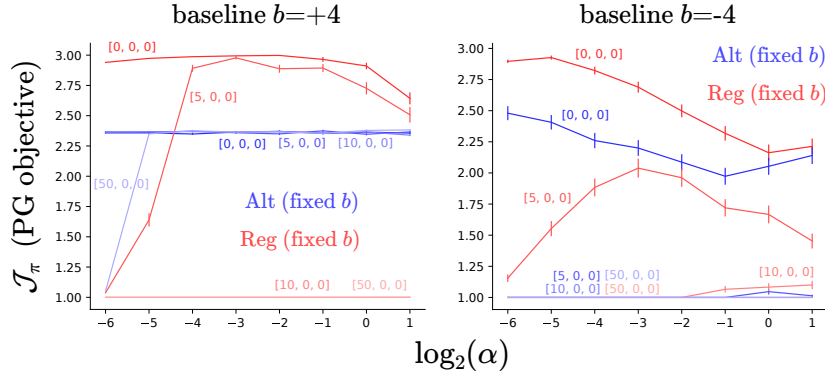


Figure 22: Bias in the alternate PG estimator. Parameter sensitivity of the PG estimators for a fixed baseline on bandit problem $\mathbf{r} = [1 \ 2 \ 3]^\top$ with $\epsilon \sim \mathcal{N}(0, 1)$. The baseline in these experiments was fixed and not updated, and therefore, the alternate estimator remained biased throughout the learning. The individual labels on the curves show the action preference initializations.

This experiment, additionally, serves as a verification of Lemma 10.1. Recall from the lemma that for a fixed baseline, the fixed point of the expected alternate update is given by $\pi^*(a) = \frac{1}{r(a)-b} \left(\sum_c \frac{1}{r(c)-b} \right)^{-1}$. Evaluating this expression for this experiment gives $\boldsymbol{\pi}^* = [0.18 \ 0.27 \ 0.55]^\top$. Using $\mathbf{r} = [1 \ 2 \ 3]^\top$, we can calculate $\mathcal{J}_{\pi^*} = \boldsymbol{\pi}^* \mathbf{r} = 2.36$. And unsurprisingly, this value matches the value shown in Figure 22 (left) to which the alternate estimator converged under various policy initializations. Further, as we had predicted from the update plots shown in Figure 14 (bottom-right), the alternate estimator with a pessimistic baseline converged to the nearest corner. For policies saturated on action a_0 , this corner corresponded to action a_0 as well and thus the agents

obtained a reward of 1 for all the policy saturations considered.

D.5 Can the Regular Estimator be Modified to Deal with Saturated Policies?

In this section we consider whether certain modifications to the regular estimator can make it effective in dealing with saturated softmax policies.

Dropping the Reward Baseline

From our previous experiments, we know that the higher variance of the alternate estimator helps it in escaping the corners of the probability simplex. So we investigate whether a similar increase in variance obtained by removing the reward baseline from the regular estimator will help it in escaping the corners as well. Figure 23 shows the sensitivity of the final mean performance (during the last 50 timesteps of a 1000 timestep run) to the stepsize α for the regular estimator without any baseline on the bandit task $\mathbf{r} = [1 \ 2 \ 3]^\top$ with $\epsilon \sim \mathcal{N}(0, 1)$. The results show that the regular estimator without a baseline cannot be a good competitor to the alternate estimator for escaping the saturated policy regions. Upon careful thought, this should not be surprising: the regular estimator, with or without the baseline, has zero variance at simplex corners (see Proposition 9 and Figure 11). Consequently, there is neither any gradient signal nor any noise near saturated regions and the policy gets stuck. This is unlike the alternate estimator which not just has a higher variance, but a variance structure that is better suited to escaping the saturated policy regions.




Figure 23: Parameter sensitivity plots for the regular estimator without a reward baseline. The individual curve labels show the action preference initializations.

Adding Normally Distributed Noise to the Policy Updates

The reward noise utilized by the alternate estimator ultimately gets translated into noisy updates for the action preferences, whereas this translation doesn’t go through very well for the regular estimator, especially at the corners (refer Example 1). Therefore, in this experiment, we study the effect of adding an artificial noise to the gradient updates for both the estimators. In particular, we trained the different estimators on the bandit problem $\mathbf{r} = [1 \ 2 \ 3]^\top$ with $\epsilon \sim \mathcal{N}(0, 1)$. And to make the updates noisy, we added a normally distributed noise to the gradient updates, i.e. for each action a , the action preferences were updated as follows: $\theta_a^{\text{new}} = \theta_a^{\text{old}} + \alpha(g_a + \xi)$, where $\xi \sim \mathcal{N}(0, \text{grad-noise})$; we also swept over the the parameter $\text{grad-noise} \in \{0, 0.1, 0.5, 1, 2\}$. We experimented with a policy saturated to $\theta = [10 \ 0 \ 0]^\top$ at the beginning of training; and additionally a uniform policy to gauge whether this additional update noise had a detrimental effect on the estimators’ performance, since with a uniform policy there is already a strong gradient signal and the agent does not require any extraneous noise for moving towards the optimal policy.

Figure 24 shows the sensitivity of the final performance for the expected gradient, and the regular and the alternate estimators using the true r_π . Figure 25 shows similar results for the regular and alternate estimators with a learned baseline, initialized to zero. The results from this experiment suggest that adding a normal noise to the gradient bandit update works well for escaping the saturated policy regions in this bandit setting for all the five estimators; in particular, the regular estimator can escape the saturation with a high enough noise. At first, this observation might seem to weaken the case for the alternate estimator. However, note that adding

gradient noise also worsens the performance of all the estimators for the uniform policy case, which is clearly an undesirable property. Now recall that the alternate estimator benefits much more from an optimistic baseline than it does from the reward noise. In contrast to the noise, an optimistic baseline actually offers a clear signal that moves the agent out of saturation and towards a more uniform policy distribution. Then, the obvious follow-up question to ask is whether adding gradient noise to the regular estimator makes it more effective than the alternate estimator with an optimistically initialized baseline. Figure 26, which shows the sensitivity plots for the regular and the alternate estimator with both an optimistic baseline and added gradient noise, answers this question in a negative. Comparing the results from Figure 26 with those from Figure 25 shows that the alternate estimator benefits from an optimistic baseline much more than the regular estimator benefits from added gradient noise. More importantly, even though adding gradient noise helped with this particular bandit instance, as we show in §E.4, adding gradient noise doesn’t help the regular estimator overcome saturated policies in MDPs.




Figure 24: Sensitivity plots for the gradient bandit estimators with true r_π and an additional normally distributed gradient noise. The individual graph labels (0, 0.1, 0.5, 1, 2) show the standard deviation of the normal noise added to the gradient. The subplot headers show the policy saturation at the beginning of the experiment.




Figure 25: Sensitivity plots for the regular and the alternate estimators with added gradient noise and a learned baseline initialized to zero. The figure header shows the standard deviation of the Gaussian gradient noise. Top and bottom rows correspond to uniform and saturated policy initializations respectively. Interestingly, for the uniform policy case and a high standard deviation of the noise, the alternate estimator is more robust to the gradient noise than the regular estimator.

E REINFORCE with Tabular Representation: Full Experiments

In this section, we demonstrate that the alternate estimator for MDPs enjoys similar benefits as the alternate estimator for bandits. We begin by giving the experimental details, and present additional experiments that




Figure 26: Sensitivity plots for the estimators with gradient noise and an optimistic baseline.

complement those presented in §6

E.1 Experimental Details

We use the chain environment²¹, shown in Figure 27, which is an episodic MDP where the expected rewards are zero everywhere except at the rightmost transition. We train five different agents using REINFORCE (exact pseudocode given in Algorithm 7). All the agents maintain a tabular policy and, in some cases additionally, a tabular value function estimate. The policy is learned either using the expected PG, or using the REINFORCE algorithm with either the regular estimator (true v_π or a learned baseline) or the alternate estimator (again with true v_π or a learned baseline). The critic is estimated using Monte-Carlo sampling.




Figure 27: Chain Environment. At the beginning of each episode, the agent is reset to the state s_3 . At each timestep, it can take the action `left` or `right`, which deterministically moves it into the corresponding next state. The episode ends when the agent reaches either of the two terminal states. The expected rewards are zero at all transitions except the one into the right most terminal state. Further, a reward noise $\epsilon \sim \mathcal{N}(0, 1)$ is added to the expected reward at each timestep.

All the methods were run for 100 episodes. The learning curves show the mean performance and standard error, and the sensitivity plots show the mean performance over the last 10 episodes. The results were averaged over 150 independent runs. For getting a uniform policy initialization, we set $\theta_{\text{left}} = \theta_{\text{right}} = 0$ for all the five states. A saturated policy was obtained by initializing $\theta_{\text{right}} = 0$ and $\theta_{\text{left}} \in \{1, 2, 3\}$ depending on the degree of saturation, again for all the five states. The agents either used the true value function v_π (computed analytically) or learned a value function estimate b which was initialized to either zero, $+4$, or -4 for each state of the MDP. We swept over the policy stepsize $\alpha \in \{2^{-6}, 2^{-5}, \dots, 2^1\}$ and the value function baseline stepsize $\beta \in \{2^{-4}, 2^{-3}, \dots, 2^1\}$. We set the discount factor $\gamma = 0.9$. To reduce the runtime of the experiments, if the agent was unable to solve the task after 100 timesteps, we terminated the episode²².

²¹Implementation available here: https://github.com/svmgrg/rl_environments/tree/main/LinearChain.

²²It might seem unsatisfying to artificially end an episode after 100 timesteps. However, not doing so resulted in some

Algorithm 7 REINFORCE with Baseline

Input a policy $\pi_{\mathbf{w}}(a|s)$ and a value function estimate $\hat{v}_{\omega}(s)$
 Input the policy stepsize α and the value function stepsize β
 Initialize \mathbf{w} and ω
for total number of episodes **do**
 Sample an episode $(S_0, A_0, R_1, \dots, S_{T-1}, A_{T-1}, R_T, S_T) \sim \mathbb{P}_{\pi}$
 # Calculate the PG estimate
 $\hat{\mathbf{g}}^{\text{REINF}} = \sum_{t=0}^{T-1} \gamma^t (G_t - \hat{v}_{\omega}(S_t)) \nabla_{\mathbf{w}} \log \pi_{\mathbf{w}}(A_t|S_t)$
 # Update the policy
 $\mathbf{w} \leftarrow \mathbf{w} + \alpha \cdot \hat{\mathbf{g}}^{\text{REINF}}$
 for each step of the episode $t = 0, 1, \dots, T - 1$ **do**
 Calculate the return $G_t = R_{t+1} + \gamma R_{t+2} + \dots + \gamma^{T-t-1} R_T$
 # Update the value function estimate
 $\omega \leftarrow \omega + \beta \cdot (G_t - \hat{v}_{\omega}(S_t)) \nabla_{\omega} \hat{v}_{\omega}(S_t)$

E.2 Performance of PG Estimators against Increasingly Saturated Policy Initializations

We begin by showcasing that the alternate PG estimator is competitive with the regular estimator in case of uniform policy initialization and superior to it in case of sub-optimally saturated policy initialization. These results are essentially identical to those from the bandit experiments (§D.2). Figure 28 shows the learning curves for the best performing parameter configuration (chosen from the sensitivity plots), and Figure 29 shows the stepsize sensitivity corresponding to the final performance for each parameter setting. Observe that irrespective of the policy saturation, the alternate estimator with baseline is vastly superior to all the methods (except expected PG), and that the alternate estimator with true v_{π} is a little better than the regular estimators. We attribute the superior performance of alternate estimator with baseline to the bias of the estimator combined with utilizing the noise in the returns, which probably allows it to have better exploration. From the sensitivity plots, we further see that as the initial policy is saturated more towards the left direction (in the chain environment), performance of all the methods worsens.

PG Estimators without Reward Noise

We now study the performance of different estimators when we remove the reward noise from the chain MDP. Figure 30 shows the learning curves and Figure 31 shows the parameter sensitivity plots for different estimators using REINFORCE on the chain MDP with reward noise $\epsilon = 0$. The results are somewhat similar to the bandit setting: in absence of reward noise, the performance of the alternate estimator drops with saturated policy initialization. However, there is one important difference from the bandit setting. For MDPs, the alternate estimator finally cares about the noise in the return estimate which depends on both the reward noise and the sampling noise. Therefore, even without reward noise, in MDPs the alternate estimator can utilize the

of the experiments running for about a million steps, which dramatically increased the runtime of the experiments. To put this number in perspective, an optimal policy can solve the chain environment in three steps. These million timestep long episodes occurred when the policy learned to go right with a high probability in one state and learned to go left in the very next state, effectively creating an unending loop between these two states.

We justify the 100 timestep cutoff by noting that the contribution of states after a hundred steps becomes negligibly small. This can be seen by focusing on the term γ^t in the expression for the REINFORCE estimator \hat{g}^{REINF} (Algorithm 7): for our chosen discount factor, $\gamma^{100} = 0.9^{100} = 2.7 \times 10^{-5}$. We also note that in our experiments, we learned the value function using simple Monte-Carlo estimates and, contrary to the recommendation made by Pardo et al. (2018), did not bootstrap at the state where the episode was artificially timed out. We again justify this by the fact that $\gamma^{100} v_{\pi}(s_{\text{timeout}})$ is negligibly small compared to the returns observed in our experiments (see the scale of the y -axis in Figure 28) and is thus safe to ignore. Yet another argument in favor of keeping a timeout in our experiments is that, unlike timeouts with continuing MDPs where they are guaranteed to occur, in the chain environment the occurrence of the episodes being timed out decreases as the agent’s policy improves.

Also note that we are using a discount factor $\gamma \neq 1$ for an episodic problem despite the fact that episodic problems do not necessarily require discounting. This choice was primarily made to justify timing out the episode after 100 steps. However, even if we ignore the timeout issue, keeping a discount factor around for episodic MDPs can be useful. It caps the maximum return at the fixed value of $\frac{r_{\max}}{1-\gamma}$, and helps us avoid situations such as the infinite loop between two states in the supposedly episodic chain MDP.

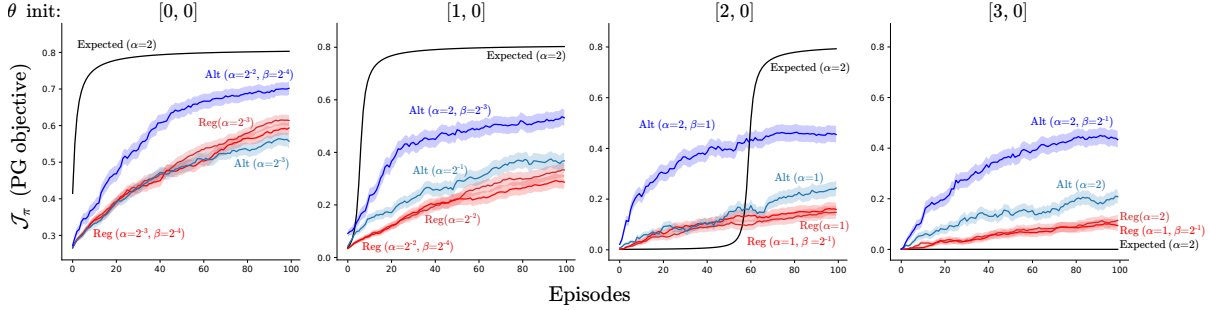


Figure 28: Learning curves for the REINFORCE agent with different estimators against increasingly saturated policy initializations on the five state chain environment. The figure header shows the action preference initialization for θ_{left} and θ_{right} respectively. The estimators which don't mention the baseline stepsize β used the true value function v_π , whereas the others learned a baseline using Monte-Carlo sampling. The baseline was initialized to zero for all the states.

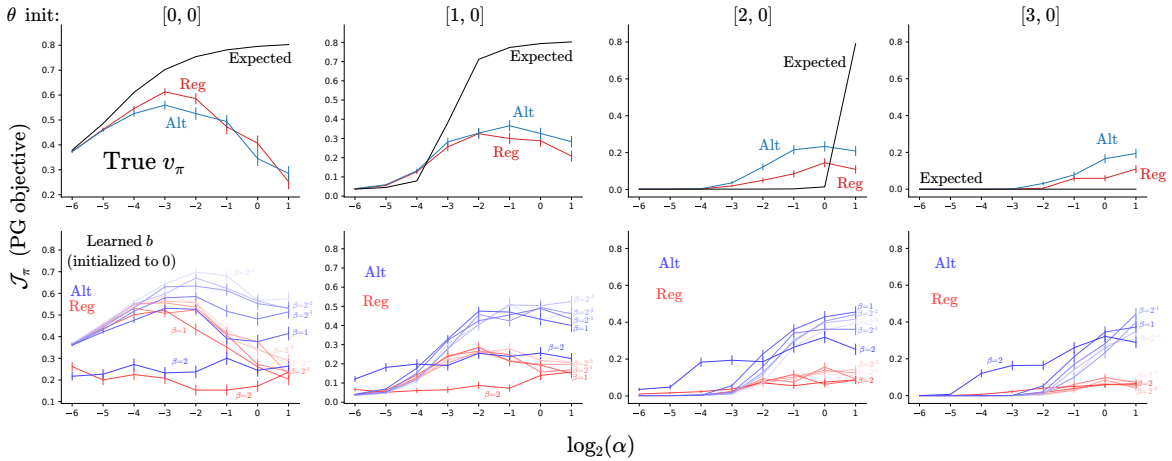


Figure 29: Sensitivity plots for different PG Estimators, showing the mean final performance during the last 10 episodes, against increasingly saturated policy initializations on the chain environment. The different columns correspond to the different degrees of policy saturations and the figure headers show the action preference initializations for the **left** and **right** actions respectively.

noise in estimating the returns from Monte-Carlo sampling. However, for the chain MDP, the transitions were deterministic and therefore the sampling noise was also quite less.

E.3 Optimistic and Pessimistic Value Function Estimates

In this experiment, we study the performance of the alternate estimator when the value function baseline is initialized optimistically or pessimistically and then learned using Monte-Carlo sampling. Figures 32 and 33 show the parameter sensitivity plots for agents with optimistically ($b = +4$) or pessimistically ($b = -4$) initialized baselines respectively, against different degrees of policy saturations. The results, which are yet again similar to the bandit setting, show that having an optimistic baseline significantly helped the alternate estimator; in particular, compare the performance of the alternate estimator on saturated policies with ($b = +4$, as shown in Figure 32) and without ($b = 0$, as shown in Figure 29) optimism. Whereas, a pessimistic baseline hurt its performance. Moreover, the alternate estimator with the optimistic baseline prefers smaller critic stepsizes (allowing it to enjoy the optimism for longer), and vice-versa for the pessimistic baseline.

Fixed Value Function Estimate

In this experiment, we fixed the value function estimate for both the estimators; which made the alternate, but not the regular, estimator biased. Figure 34 shows the sensitivity plots for the two estimators. This experiment




Figure 30: Learning curves for PG estimators without reward noise. The setting is same as Figure 28, except that the reward noise is set to zero. We see that, for saturated policies, the alternate estimator performs worse as compared to its performance in Figure 28 where it had reward noise, and **curiously** alternate with v_π performs better than alternate with a learned baseline.

is the MDP analog of the biased gradient bandit experiment (see Figure 22) and suggests that Theorem 2 might also hold for MDPs: since, we can show that the policy (for the optimistic case) converges to the fixed point $\pi^*(a) = \frac{1}{r(a)-b} \left(\sum_c \frac{1}{r(c)-b} \right)^{-1}$ given by Lemma 10.1. For MDPs, we need to compute this fixed point separately for each state: $\pi(\text{left}|s) = \pi(\text{right}|s) = 0.5$ for $s \in \{s_1, s_2, s_3, s_4\}$, and $\pi(\text{left}|s_5) = 0.43$ and $\pi(\text{right}|s_5) = 0.57$. Then, we analytically calculate $\mathcal{J}_{\pi^*} = v_{\pi^*}(s_3) = 0.28$ which approximately matches the value shown in Figure 34 (left) where the policy for the optimistic fixed baseline actually converged to; therefore, our prediction matches with the actual results. For a pessimistic baseline, the policy converged to the nearest




Figure 31: Sensitivity plots for PG estimators without reward noise. Setting is same as in Figure 29 except that there is no reward noise here.

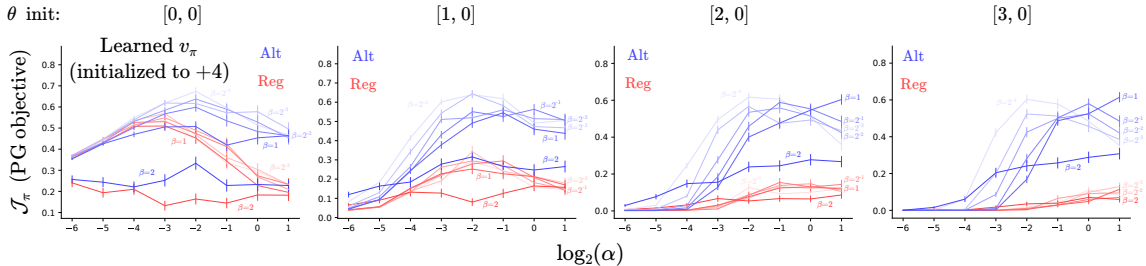


Figure 32: Sensitivity plots for an optimistically initialized value function estimate. The setting is similar to that in Figure 29, except that the baseline is initialized to +4.

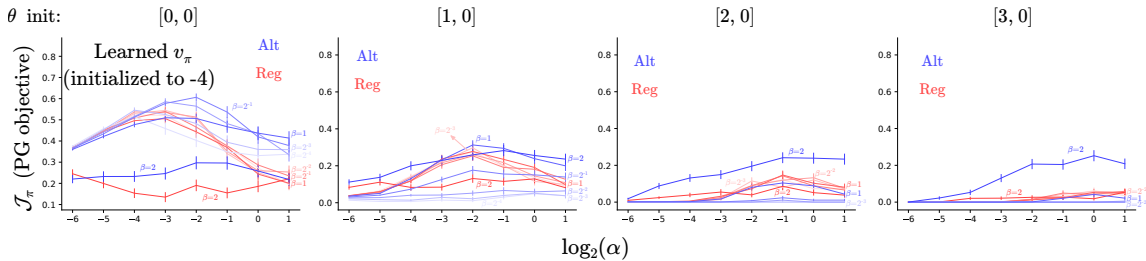


Figure 33: Sensitivity plots for a pessimistically initialized value function estimate. The setting is again similar to that in Figure 29 except that the baseline is initialized to -4 .

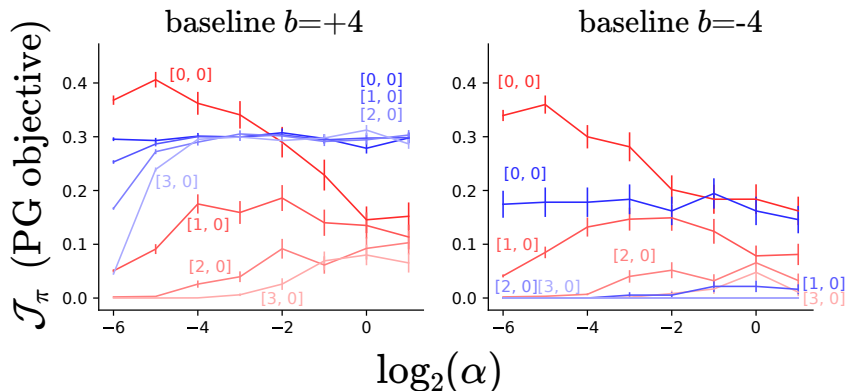


Figure 34: Bias in the alternate PG estimator for MDPs. Sensitivity plots for the REINFORCE agents using regular or alternate estimator with fixed baselines on the chain environment. The value function estimates were fixed to $+4$ (for left subplot) and -4 (for right subplot) and were not subsequently updated. The labels for the individual curves show the action preference initializations $[\theta_{\text{left}}, \theta_{\text{right}}]$ for all the five states.

corner for each state (in this case action left) and obtained zero return.

E.4 Effect of adding Gradient Noise to the Estimators

This experiment studies the efficacy of randomly perturbing the PG estimates, with a normally distributed noise while updating the action preferences, in helping the agent escape saturated policy regions. The preferences, for a given state s and action a , were updated using $\theta_a^{\text{new}}(s) = \theta_a^{\text{old}}(s) + \alpha [g_a(s) + \xi]$ with $\xi \sim \mathcal{N}(0, \text{grad-noise})$. We swept over the noise magnitude $\text{grad-noise} \in \{0, 0.1, 0.5, 1, 2\}$. Both the estimators learned a value function estimate which was initialized to zero²³. Figure 35 shows the sensitivity of the final performance for the two estimators.

These results show that adding gradient noise doesn't significantly help the regular estimator in dealing with saturated policies for MDPs: the performance of the regular estimator with gradient noise was only slightly better than its performance without gradient noise (cf. Figure 29). More importantly, the alternate estimator had a superior final performance compared to the regular estimator for both the uniform and the saturated policy initialization cases. This observation is in contrast to what we saw in the bandit case (§D.5), where the gradient noise seemed to help more. We explain this result by the fact that random noise merely increases the variance of the gradient estimate without providing any effective signal. And while this added variance helped the regular estimator in the simpler bandit problem, it is apparently not sufficient for the MDPs. The alternate estimator on the other hand relies on both the noise from the sampled return and also a much stronger signal from the bias in the critic estimate to escape the saturated policy regions.

²³We do not present the results for expected PG and the regular and alternate REINFORCE estimators using the true v_π , since they were qualitatively similar to the results for the learned baseline case presented here.

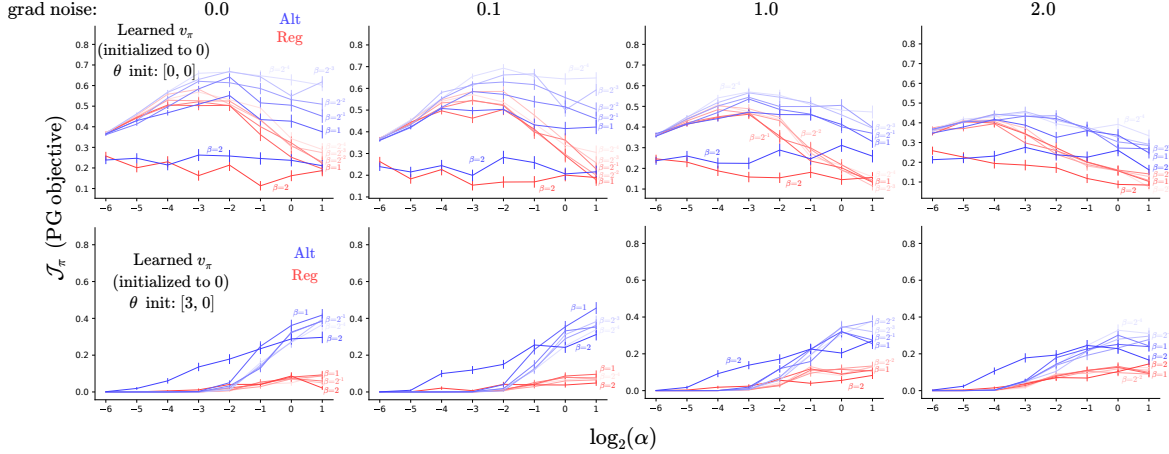


Figure 35: Adding gradient noise to the REINFORCE estimators. The top row shows the sensitivity plots for a uniform policy and the bottom row shows the sensitivity plots for a policy saturated as $\theta_{\text{left}} = 3$ and $\theta_{\text{right}} = 0$. Different columns correspond to different gradient noise magnitudes.

E.5 A Uniform Policy Initialization can be Sub-optimal too!

In all the previous experiments, we found that an optimistically initialized baseline helped the performance of the alternate estimator (by making the policy more uniform). In this section, we consider an MDP, where a uniform policy is bad for exploration, and therefore an optimistic baseline by making the agent’s policy more uniform hurts its performance. Figure 36 shows this MDP, which is a modified form²⁴ of the five state chain environment: instead of having just two actions, this modified chain has four actions for each state, three of which take the agent towards the left direction and the fourth towards the right direction.




Figure 36: Linear chain environment where a uniform policy will explore poorly.

Figure 37 shows the sensitivity plots for the REINFORCE agents with the regular and alternate estimators for four different settings: using the true v_π , and a learned baseline initialized to either of 0, +4, or -4 . The policy for all these settings was initialized uniformly, i.e. $\theta = 0$ for all states and action. The results show that in each case, the alternate estimator had a poorer performance as compared to the regular estimator. In particular, even an optimistically initialized baseline could not improve the alternate estimator’s performance. We reconcile this fact by noting that in this problem the uniform random policy will actually lead to a sub-optimal exploration and an optimistic baseline will hamper the performance of the alternate estimator.

F Online Actor-Critic with Linear Function Approximation

In this section, we demonstrate that the desirable properties of the alternate estimator continue to hold with function approximation. We ran experiments with online actor-critic (see Algorithm 8) on various MDPs using linear function approximation (with tile-coding) and neural networks. The agent was trained using either the regular or the alternate estimator and a critic learned using TD-style updates; in these experiments, we do not have access to the underlying dynamics, and therefore we did not conduct experiments with the true value function

²⁴This MDP was inspired from the “Vanishing gradient example” given in Chapter 12 of Agarwal et al. (2019).

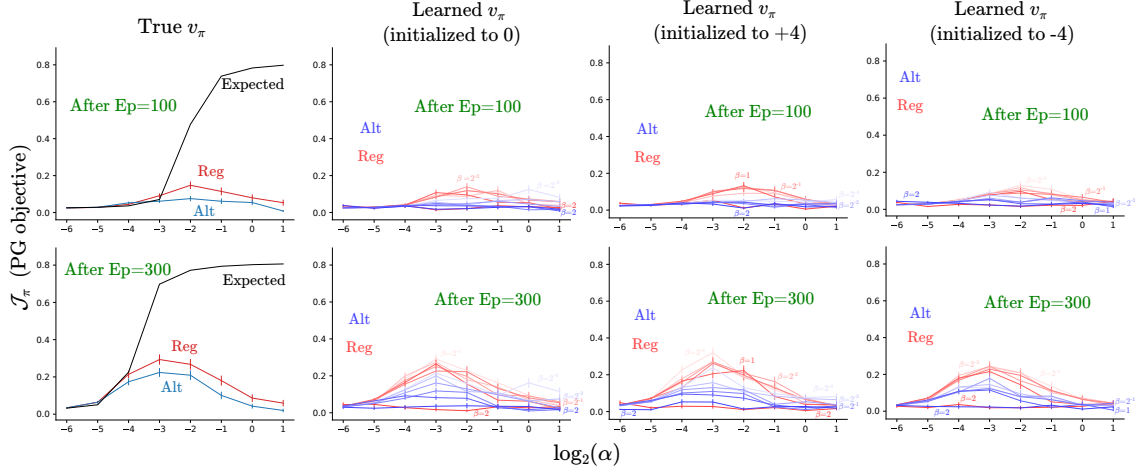


Figure 37: Sensitivity plots for different methods, initialized with a uniform policy, on the difficult chain environment (described in Figure 36). Top row shows the mean final performance (averaged over the last 10 episodes) of the estimators after 100 episodes and the bottom row shows the mean final performance (again averaged over the last 10 episodes) after 300 episodes. Different columns refer to either a true baseline or a learned baseline with different initializations. We see that even after 300 episodes, the algorithms are not able to solve the task (i.e. they are unable to achieve the maximum return of 0.81). Further, even with an optimistic baseline, the alternate estimator performs quite poorly as compared to the regular estimator.

v_π . We considered two different types of experiments: one with artificially introduced policy saturation, and another with non-stationary MDPs which naturally give rise to sub-optimally saturated policies. Experiments where we explicitly saturate the agent’s policy at the beginning of learning are helpful in understanding the behavior of the estimators by providing a controlled environment. The non-stationary MDPs serve as a proof-of-concept that sub-optimally saturated policies are not merely pathological constructs, but can occur in practical problems that a user might care about.

Algorithm 8 Online Actor-Critic

```

Input a policy  $\pi_{\mathbf{w}}(a|s)$  and a value function estimate  $\hat{v}_\omega(s)$ 
Input the policy stepsize  $\alpha$  and the value function stepsize  $\beta$ 
Initialize  $\mathbf{w}$  and  $\omega$ 
for each episode (in total number of episodes) do
    Sample  $S \sim \mu$ 
    Initialize the geometric discounting term  $I_\gamma \leftarrow 1$ 
    while the state  $S$  is not terminal, i.e.  $S \neq s_\times$  do
        Sample  $A \sim \pi_{\mathbf{w}}(\cdot|S)$  and  $S', R \sim p(\cdot, \cdot|S, A)$ 
         $\delta \leftarrow R + \gamma \hat{v}_\omega(S') - \hat{v}_\omega(S)$  (if  $S'$  is terminal, then  $\hat{v}_\omega(S') = \hat{v}_\omega(s_\times) := 0$ )
        # Calculate the gradient estimator
         $\hat{\mathbf{g}}^{\text{OAC}} = I_\gamma \delta \nabla_{\mathbf{w}} \log \pi_{\mathbf{w}}(A|S)$ 
        # Update the policy
         $\mathbf{w} \leftarrow \mathbf{w} + \alpha \cdot \hat{\mathbf{g}}^{\text{OAC}}$ 
         $I_\gamma \leftarrow \gamma I_\gamma$ 
        # Update the value function estimate
         $\omega \leftarrow \omega + \beta \cdot \delta \nabla_{\omega} \hat{v}_\omega(S)$ 
         $S \leftarrow S'$ 
    
```




Figure 38: Schematic diagrams of Acrobot, MountainCar, and the DotReacher environments. For the DotReacher environment, we also plotted multiple trajectories for a particular instance of online actor-critic that learned to solve this task.

F.1 Experimental Details for Linear Function Approximation

For experiments with linear function approximation, we used the environments²⁵ Acrobot (Sutton, 1996) and MountainCar (§4.3, Moore, 1990; Singh and Sutton, 1996), which are depicted in Figure 38. Both of these environments are episodic in nature with a continuous state space and a discrete action space consisting of three actions: **left**, **do-nothing**, and **right**. In Acrobot, the goal is to control an under-actuated double pendulum such that its end effector reaches a certain height, by applying torque on the second joint. In MountainCar, the goal is to move an under-powered car from the bottom to the top of a hill. For both the environments, we fixed the discount factor $\gamma = 1$ and timed out the episode after 1000 timesteps if the agent was unable to solve it by then²⁶.

We performed sweeps over the policy stepsize $\alpha \in \{2^{-13}, 2^{-11}, \dots, 2^5\}$ and the critic stepsize $\beta \in \{0.01, 0.05, 0.1, 0.5, 1, 2\}$. The state for Acrobot is $[\cos(\theta_1) \sin(\theta_1) \cos(\theta_2) \sin(\theta_2) \dot{\theta}_1 \dot{\theta}_2]$, with $\dot{\theta}_1 \in [-4\pi, +4\pi]$ and $\dot{\theta}_2 \in [-9\pi, +9\pi]$. The state for MountainCar is $[\text{position}, \text{velocity}]$ (of the car), with **position** $\in [-1.2, +0.5]$ and **velocity** $\in [-0.07, +0.07]$. Both the policy (for details, see §C.4) and the value function used by the agent were linear in the state features. And the state features were constructed using tile-coding (§9.5.4, Sutton and Barto, 2018). In particular, we used four tiles across each dimension of the state space, and eight different tilings with randomly chosen offsets. We also added an “always on” bias feature and normalized the tile-coding features (by dividing them with nine) to have a unit sum.

F.2 Non-stationary MountainCar and Acrobot

This experiment investigates the performance of the online AC algorithm, with the regular and the alternate estimators, on non-stationary Acrobot and MountainCar. To induce non-stationarity in either task, we switched the **left** and **right** actions after half-time. The policy was initialized uniformly by setting all the action preferences to zero. The critic weights were also initialized to zero. Each agent was trained for 400k timesteps for Acrobot, or 200k timesteps for MountainCar.

Figure 39 (left) shows the learning curves for the two estimators. For each estimator, we selected two different sets of stepsize configurations: one that had the best performance right before the non-stationarity hit, and

²⁵We used the environment implementations available here: <https://github.com/andnp/PyRLEnvs/tree/main/PyRLEnvs/domains>. In this repository, MountainCar is referred to as GymMountainCar.

²⁶For these experiments we bootstrap at the final state of the episode, i.e. if the episode times out, we bootstrap using $\hat{v}_\pi(s_{\text{timeout}})$, and if the episode actually ends, we use $v(s_\times) = 0$ for bootstrapping. This justifies the use of timeouts since (1) this process follows the recommendation made by Pardo et al. (2020), (2) this is an episodic task and as the agent’s performance improves, the episodes will be timed out less, and (3) from our experiments, we observed that 1000 timesteps gave the agent enough time to encounter a true terminal state. Encountering true terminal states is important when learning with TD-style updates, since only at terminal state do we bootstrap from a true value (of zero); for all the other states, we bootstrap from an estimated value. Note that this argument is different from the numerical precision argument given in Footnote 22 of §E. Also note that with $\gamma = 1$ and a timeout of 1000, for the agents that are unable to solve the task within 1000 timesteps, the value function estimates will diverge to negative infinity; for instance, see the value of $\hat{v}_\pi(s_0)$ against time for the regular estimator in Figure 45.

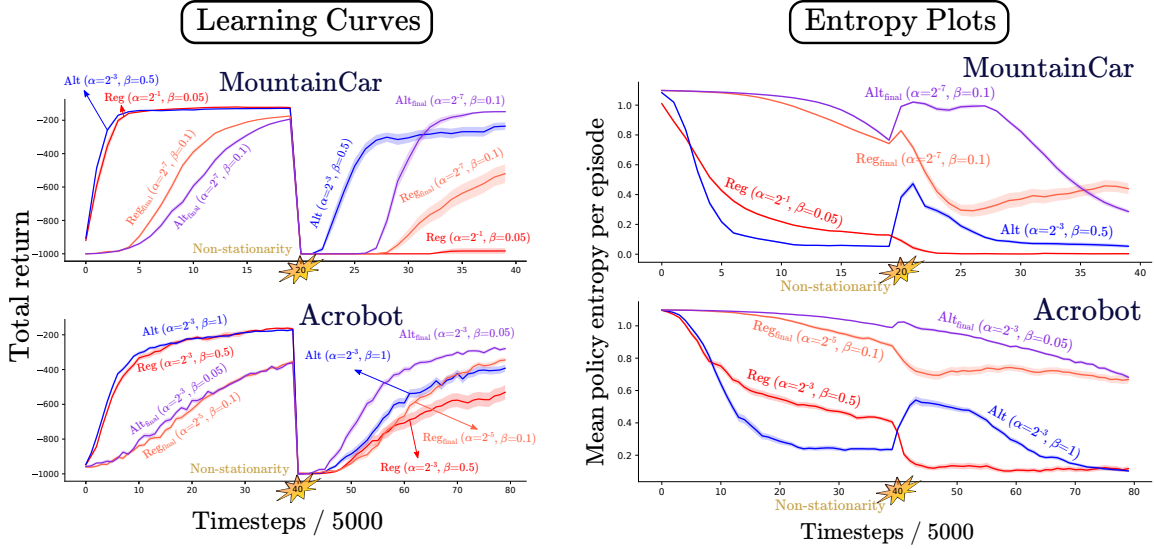


Figure 39: Learning curves and entropy plots for online linear AC with tile-coding on Acrobot and MountainCar. Both the estimators have one learning curve each for the best performing set of stepsizes right before the non-stationarity hit and at the end of learning. The right subplots show the entropy of the policy on the states encountered during each episode. All the curves show the mean values (thick line) and the standard error (shaded region) over 50 independent runs.

another that had the best performance at the end of the experiment (in the figure, we denote this latter set by putting a `final` in the subscript of the corresponding estimators). For MountainCar, the alternate estimator was superior to the regular estimator for both sets of stepsizes. Remarkably, the regular estimator (red curve), with the best performing parameters at timestep $100k$, was unable to recover from the non-stationarity; whereas the alternate estimator (blue curve) was able to recover, despite having a similar performance as the regular estimator at the end of $100k$ timesteps. On Acrobot, the alternate estimator was again superior to the regular estimator, however the difference was less evident.

We attribute the superior performance of the alternate estimator to the bias in the critic estimate. For instance, consider what would happen in MountainCar for an online AC agent using a well performing stepsize configuration. At $100k$ timesteps, the agent would be able to learn a good policy and therefore the critic would learn to predict a return that is greater than -200 . Now, when the non-stationarity hits, the agent’s policy suddenly becomes sub-optimally saturated, and thus the agent starts receiving returns much lower than -200 . This in turn means that at this time, the critic estimate has become optimistic. As a result, for the alternate estimator, the critic would encourage exploration by pushing the policy towards a more uniform distribution. Figure 39 (right) corroborates this point: the policy entropy for the alternate, but not for the regular, estimator has a spike right around the timestep when the non-stationarity was introduced.

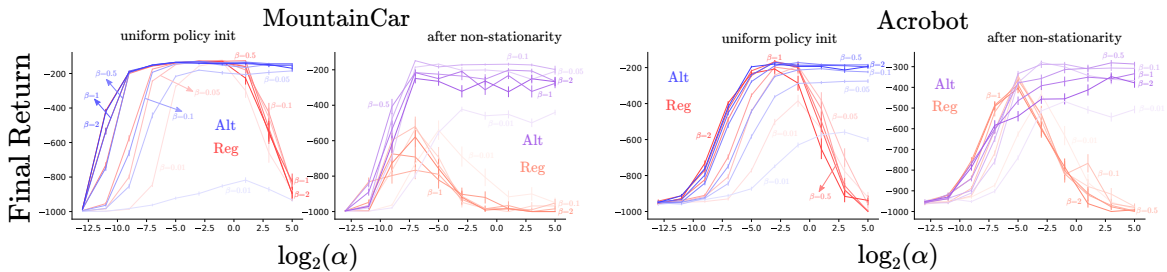


Figure 40: The sensitivity plots for online Actor-Critic with linear function approximation (using tile-coded state features) on MountainCar and Acrobot. The sensitivity plots show the mean performance during the last 5k timesteps, averaged over 50 independent runs.

Figure 40 shows the parameter sensitivity plots for these tasks. From the sensitivity plots, we see that the alternate estimator is much less sensitive to the stepsizes and achieves a superior or competitive final performance compared to the regular estimator, both before and after the non-stationarity was introduced.

F.3 Online AC with Entropy Regularization and Escort Transform on Non-stationary Tasks

We now study how the alternate estimator compares to entropy regularization and escort transform, two techniques that are designed to help policy gradient methods deal with policy saturation. In these experiments, we train the agent with the online Actor-Critic algorithm on the non-stationary MountainCar and Acrobot tasks (task description given in the previous section). For entropy regularization, we considered the softmax policy parameterization with either the regular or the alternate PG estimator. For the escort transform, we considered the simple online AC algorithm (without any entropy bonus).




Figure 41: The sensitivity plots for online Actor-Critic with linear function approximation (using tile-coded state features) and entropy regularization on non-stationary MountainCar and Acrobot. The sensitivity plots show the mean performance during the last $5k$ timesteps, averaged over 50 independent runs. The x -axis refer to the policy stepsize α and the differently shaded lines correspond to different critic stepsizes β s. The different rows correspond to different values of the entropy regularization parameter τ .

Since, both entropy regularization and escort transform have additional hyperparameters: the regularization factor τ for the first, and the escort power parameter p for the latter, we swept over $\tau \in \{0.0, 0.1, 0.3, 0.5, 0.7, 0.9\}$

and $p \in \{1, 2, 3, 5, 10\}$. We initialized the value function weights to zero for both the experiments, and the policy uniformly randomly. In particular, we set the actor weights so that the initial action preference for softmax policy was $\theta^{\text{init}} = [0 \ 0 \ 0]^\top$ and for the escort policy was $\theta^{\text{init}} = [1 \ 1 \ 1]^\top$ (since the action preferences for escort transform cannot all be initialized to zero).

Let us now look at the parameter sensitivity plots for this experiment. (We discussed the learning curves for this experiment in Figure 7 of the main paper.) Figure 41 shows the parameter sensitivity plots for entropy regularization. From this figure, we note that adding an entropy bonus to the alternate softmax estimator improved its performance slightly (especially for Acrobot) for small values of τ . However, for larger values of τ , the performance of the alternate estimator dropped. On the other hand, adding an entropy bonus to the regular estimator improved its sensitivity to the policy stepsize α quite a lot. Entropy regularization also significantly improved the performance of the regular estimator after non-stationarity: regular with entropy had higher final returns at the end of the experiment when compared with regular without entropy. Also, note that the alternate estimator (with or without the entropy regularization) remains superior (or competitive) to the regular estimator with entropy regularization. More importantly, recall from the learning curves (Figure 7) that the entropy bonus was not able to help the best performing regular estimator (red curve) to recover from the non-stationarity, whereas the alternate estimator (blue curve) was able to recover from it.

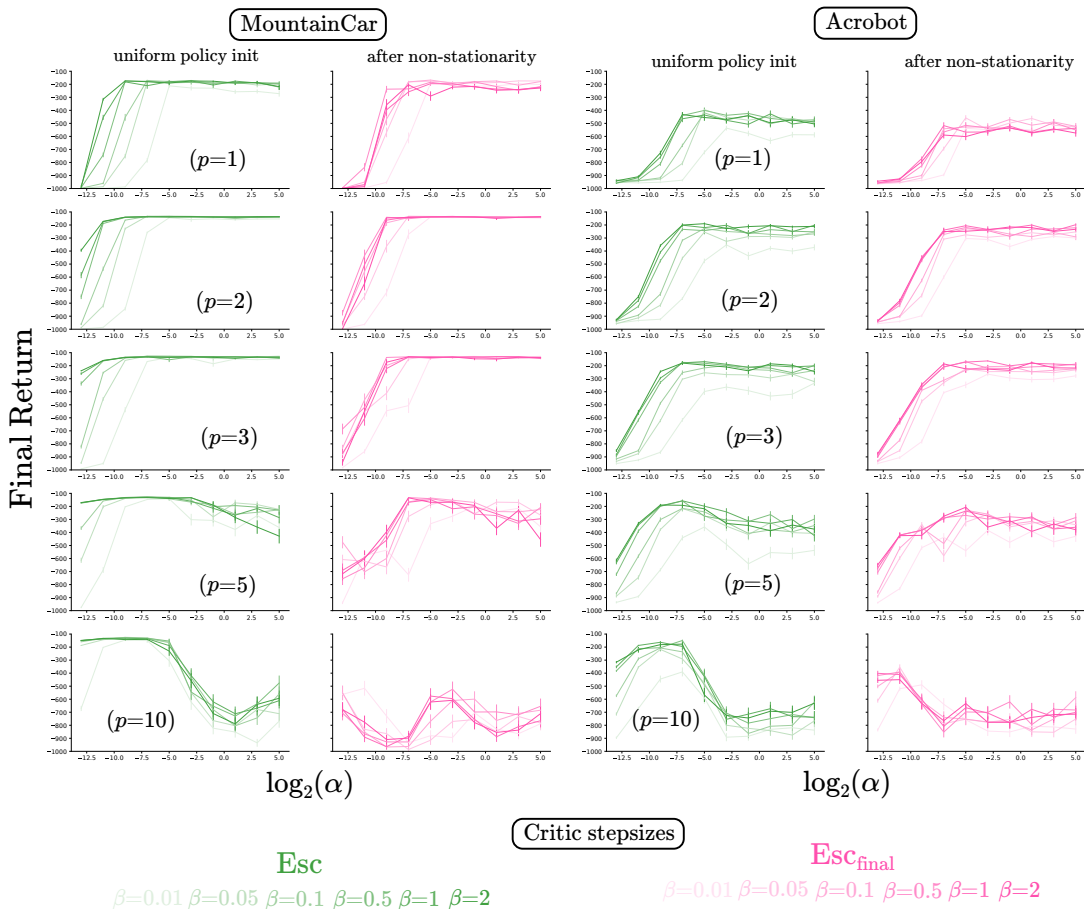


Figure 42: The sensitivity plots for escort transform with online Actor-Critic and linear function approximation (using tile-coded state features) on non-stationary MountainCar and Acrobot. The sensitivity plots show the mean performance during the last $5k$ timesteps, averaged over 50 independent runs. Different rows correspond to different values of the escort power parameter p .

We show the parameter sensitivity plots for escort transform (ET) in Figure 42. These plots show that ET worked really well for $p = 2$ and $p = 3$. For these values of the escort power parameter, ET had a very low sensitivity to both the policy stepsize α and the critic stepsize β , and achieved a good performance for most of the parameter settings. However, for $p = 1$ and $p = 10$, its performance was much worse. These results suggest

that ET is a great choice for handling non-stationarities in the environment. In fact, for the top performing choices of p , ET was able to achieve a better performance on Acrobot as compared to the alternate estimator. However, the results from the learning curves (Figure 7) showed that the best performing parameter setting (before the non-stationarity hit) for ET (green curve) was not able to recover well from the non-stationarity. And, Figure 42 hints that ET might be sensitive to the value of the parameter p , and a more extensive empirical evaluation with various PG algorithms and different environments is needed to study this phenomenon further.

F.4 Experiments with Saturated Linear Softmax Policies

We now consider the performance of the two PG estimators on Acrobot and MountainCar with artificially saturated policies. Our results show the same trends for alternate and regular estimators in the linear function approximation setting as we saw for the bandit and the tabular setting. For obtaining saturated policies, we initialized the action preference corresponding to the `do-nothing` action with a value of either 0, 5, 10, or 100 depending on the degree of policy saturation, whereas the preferences for the actions `left` and `right` were always initialized with zero. The baseline was initialized with different values (0, +500, or -500) imparting an optimistic or a pessimistic effect to the estimators. Each agent was trained for $50k$ timesteps for both the environments

Figure 43 shows the learning curves for for both the environments where the policy was initialized uniformly and the critic was initialized to zero. From this figure, we note that both the estimators were able to learn a good policy (a return of more than -200 is considered good on these environments). Further, the alternate estimator learned faster as compared to the regular estimator. We explain this behavior by noting that a baseline initialization of zero is effectively optimistic for a random policy (which will achieve large negative returns). Therefore, with an optimistic baseline, the alternate baseline was probably able to get to the fixed point of the biased update at a fast rate. Once it was near this fixed point, as its critic estimate improved, it slowly moved towards the optimal policy. Whereas, with the regular estimator, the agent just slowly moved towards the optimal policy, and thus had a lower sample efficiency.




Figure 43: Learning curves for the best stepsize parameter configurations based on final performance on the two environments. The policy was initialized uniformly randomly and the baseline was initialized to zero. The curve show the mean performance and standard error over 50 runs.

We show the sensitivity plots for these estimators, with different policy saturations and baseline initializations, on both Acrobot and MountainCar in Figure 44. The results follow similar trends that we observed for 3-armed bandits and the tabular chain environment. And impressively, these plots show that the alternate estimator can effectively escape saturated policies, even with a high policy saturation of $\theta^{\text{init}} = [0 \ 100 \ 0]^T$.

From the results in this figure, we observe that the alternate estimator had a lower parameter sensitivity and achieved a superior final performance compared to the regular estimator, for almost all the policy saturations and baseline initializations considered in this experiment; except for the (uniform policy and $b = -500$) initialization case, where the regular estimator had a better performance. We also observe that as the baseline was initialized to more negative values (to -500 instead of 0), the performance of the alternate estimator became worse. The baseline initialization affects the regular estimator much less. And as the policy initialization became more saturated, the performance of both the methods worsened. Another thing to note is that the regular estimator had a lower sensitivity to the critic stepsize β ; the alternate estimator performed poorly for lower critic stepsizes than it did for higher critic stepsizes. We believe that the reason for this is that with lower critic stepsizes,

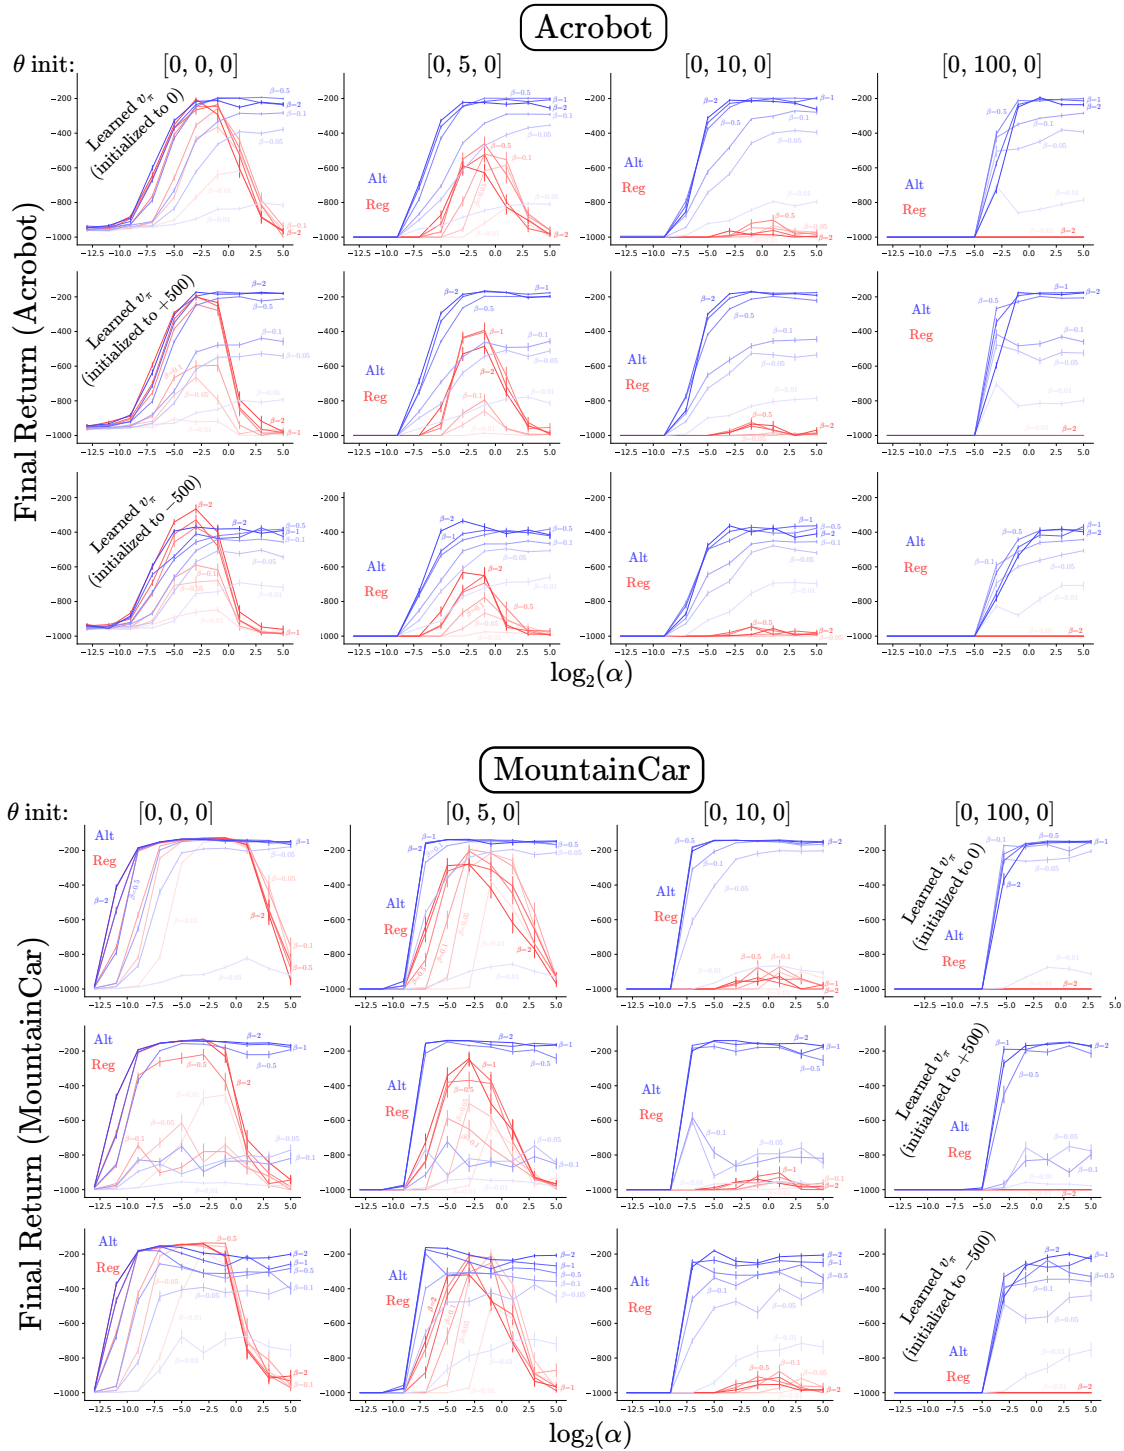


Figure 44: Parameter sensitivity plots for online AC with the two estimators on Acrobot (top) and MountainCar (bottom). The columns correspond to policies with different saturations and the rows correspond to differently initialized critic estimates. Each agent was run for $50k$ timesteps. The plots show the mean final performance during the last $5k$ steps, averaged over 50 runs.

the critic takes more time to track the value function due to which the alternate estimator remains biased for longer; and this probably affects the final performance in a negative way. We make one final comment about this experiment: the initializations (0, +500, and -500) are all effectively optimistic for a saturated policy (which gets a return of less than -1000); this probably explains why the alternate estimator had such a good final performance, for most settings, in both the tasks.

F.5 Experiments on RealisticMountainCar

We now consider the performance of the two estimators on a difficult variant of the usual MountainCar environment which we call the RealisticMountainCar²⁷. Unlike the original environment that had an instantaneous dynamics model (§4.3, Moore, 1990), i.e. an acceleration applied to the car instantaneously changed its velocity, RealisticMountainCar introduces a delay between when the acceleration is applied and when the car’s velocity increases, thereby making the dynamics more realistic. Moreover, as we find from our experiments, this delay makes the environment more difficult to solve for the regular PG estimator. The state-action space for this environment remains the same as the original MountainCar environment. The discount factor used was $\gamma = 1$ and the episode was cutoff after 1000 timesteps if the agent was unable to solve it by then. We performed sweeps over the policy stepsize $\alpha \in \{2^{-8}, 2^{-6}, \dots, 2^{16}\}$ and the critic stepsize $\beta \in \{0.01, 0.05, 0.1, 0.5, 1, 2\}$. All the agents were trained for 100k timesteps and 50 independent runs.




Figure 45: Learning curves for online AC on the RealisticMountainCar task for the best performing parameters. We also plot the value estimate of the initial state of the episode, and the mean policy entropy on the states encountered during different episodes. Even though the alternate estimator was able to learn an effective policy, the regular estimator failed to do so. Note that the regular estimator achieved its best performance at a very high stepsize that almost instantaneously pushed the policy in a highly saturated sub-optimal region; hence the learning curve for the regular estimator is flat (corresponding to an unchanging policy) and the entropy plot for the regular estimator has zero entropy throughout the learning period.

Figure 45 shows the performance curves, the value function estimate of the initial states, and the policy entropy for the two estimators as the agent learned. From these results, we see that the alternate estimator was effectively able to solve this task, whereas the regular estimator was unable to learn a good policy. From the individual runs for the regular estimator, we see that some of them resulted in a good policy (achieving a return of > -200), whereas most of them didn’t learn (achieving the maximum possible return of -1000), giving it a net average return of about -700 . We also included the plots for $\hat{v}_\pi(s_0)$ where \hat{v}_π is the critic estimate and s_0 is the start state observed at the beginning for each episode. From these plots, we see that the value function of the regular estimator diverges to negative infinity because we used a discount factor $\gamma = 1$ and the agent, for most runs, was unable to solve this task within 1000 timesteps.

We show the sensitivity plots for the mean final performance of the two estimators during the last 5k timesteps in Figure 46. These results show that the alternate estimator learns a good policy for a large range of stepsizes. Most curiously, for this experiment, the alternate estimator with a baseline initialization of +500 had a worse performance than the initialization of 0 and -500. We attribute this result to the fact that overcoming an optimistic initialization of +500 takes the longest time (during which the estimator remains biased). Further,

²⁷RealisticMountainCar was proposed by Andy Patterson. We used the implementations available here: <https://github.com/andnp/PyR1Envs/tree/main/PyR1Envs/domains>. In this repository, MountainCar is referred to as GymMountainCar and RealisticMountainCar is referred to as MountainCar.




Figure 46: Parameter sensitivity of different PG estimators with online AC on RealisticMountainCar. The policy was initialized uniformly, and the value function was initialized to different value in each column. We see that the alternate estimator is fairly robust across parameters whereas the regular estimator exhibits poor performance consistently across different parameter settings.

since 0 and -500 are already optimistic for this problem, they had a good performance as well. This observation leads us to conclude that the critic for the alternate estimator should be initialized optimistically, but not too optimistically.

G Online AC with Neural Networks

In this section, we study the online AC algorithm with a neural policy and a neural critic on an environment called DotReacher. We again experimented separately with artificially saturated policy initializations and a non-stationary environment. In either case, we found that the results were similar to the linear function approximation case. We now specify the experimental details and present the empirical results.

G.1 Details about Experiments with Neural Networks

DotReacher²⁸, shown in Figure 38 is a grid world environment where the agent is spawned uniformly randomly in a 2×2 unit squared arena. For the stationary version of DotReacher, the goal is to reach the $(0, 0)$ state within a tolerance of 0.1. The agent has 9 different actions: eight corresponding to eight different directions and one for no movement. Each action deterministically moves the agent 0.03 units in the corresponding direction. The agent gets a reward of -0.01 until it reaches the goal state. We fixed $\gamma = 1$ and the episode cutoff length as 1000. For the non-stationary variant, we initially set the goal state to $(-1, -1)$ and then shifted it to $(+1, +1)$ after $50k$ steps.

For both the experiments involving the stationary the non-stationary DotReacher, we swept the policy stepsize $\alpha \in \{2^{-17}, 2^{-15}, \dots, 2^{-3}\}$, the critic stepsize $\beta \in \{2^{-17}, 2^{-15}, \dots, 2^{-7}\}$, and used the Adam optimizer (Kingma and Ba, 2014) with the default PyTorch (Paszke et al., 2019) configuration²⁹. We maintained two separate neural networks (both with ReLU) for the policy and the critic: the policy network dimensions were $2 \times 10 \times 10 \times 9$ and the critic network dimensions were $2 \times 10 \times 10 \times 1$. For saturating the policy, we increased the action preference corresponding to the no movement action by modifying the bias units of the last layer of the neural network.

G.2 Non-stationary DotReacher

Figure 47 shows the learning curves and the stepsize sensitivity plots for the online AC algorithm with the two estimators on the DotReacher environment. The policy was initialized uniformly and the critic estimate was initialized to zero. The results follow the same pattern as the experiments on the non-stationary Acrobot and MountainCar. The alternate estimator is still competitive or superior than the regular estimator, although the difference is less pronounced.

²⁸DotReacher was proposed by Rupam Mahmood. We used the implementation available here: https://github.com/svmgrg/rl_environments/tree/main/DotReacher.

²⁹Refer here: <https://pytorch.org/docs/stable/generated/torch.optim.Adam.html>.

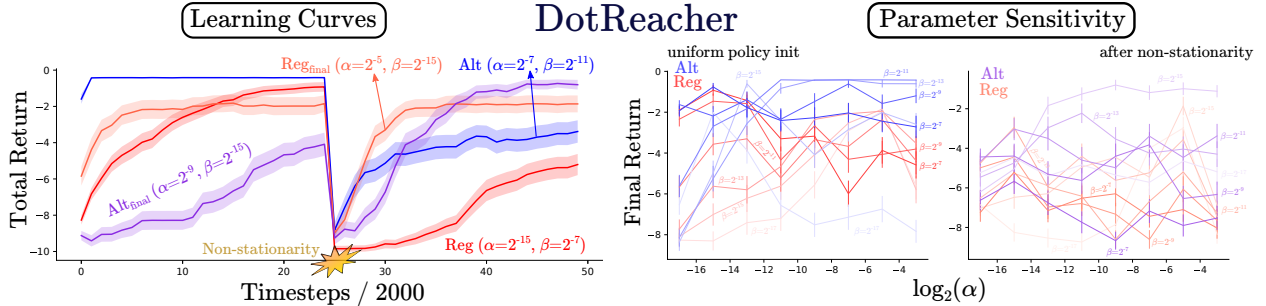


Figure 47: **(Left)** Learning curves for online AC with neural networks on DotReacher (100k timesteps, 50 runs) for two sets of best performing parameter configuration: one that obtained the best performance right before the non-stationarity occurred, and the other that obtained the best final performance at the end of learning. **(Right)** The parameter sensitivity plots showing the mean final performance (during the last 2k timesteps) at the end of 50k timesteps and the mean final performance at the end of 100k timesteps. The mean final performance was computed by averaging over the last 2000 timesteps either at half-time or at the end of training.

G.3 Experiments with Saturated Neural Policies on DotReacher

We now consider the experiments with artificially induced policy saturations on the DotReacher environment. Figure 48 shows the sensitivity plots for both the estimators for different policy saturations and critic initializations. In this problem, we consider a return of more than -1 as being good. The results are same as those discussed in §F.4: the alternate estimator had a superior performance and exhibited less sensitivity to the stepsize α as compared to the regular estimator. Interestingly, for these results, we observe that for lower amounts of policy saturations ($\theta_{\text{no-action}} \in \{5, 10\}$), the regular estimator could learn policies almost as good as the alternate estimator. But for the highest saturation case, only the alternate estimator was able to escape the saturation and obtain high returns. We also see that an optimistic baseline initialization (for this problem, $b = 0$ and $b = +25$) helped the alternate estimator achieve a good performance, whereas the pessimistic baseline initialization ($b = -25$) made the performance of the alternate estimator even worse than the regular estimator (except in the case of super high policy saturation).

H A Digressive Remark: Softmax Policies Learn even with Extremely Large Stepsizes

In this section, we mention a peculiar observation: online AC with the alternate estimator was able to learn a good policy even with extremely large stepsizes. The reader might have noticed from the parameter sensitivity plots presented in this section that the alternate estimator has a fairly good performance for a big range of the policy stepsize α . In particular, the sensitivity plots, given in Figure 44, for Acrobot and MountainCar show that for the stepsize range we tried (with the maximum value of α being $2^5 = 32$), the alternate estimator had a good performance throughout. More interestingly, for RealisticMountainCar (sensitivity plots given in Figure 46), we tried stepsizes as large as $2^{16} = 65536$, and the alternate estimator still received a fairly high final return (only slightly lower than the what the alternate estimator with the best parameter setting achieved and much higher than what the regular estimator received). Even if we take into account the normalization factor from tile-coding (which was 9), the effective stepsize still remains quite high: $2^{16}/9 = 7282$. To put this number in perspective, the default choice of the stepsize for training neural networks with the Adam optimizer (Kingma and Ba, 2014) is 0.0003. One might wonder, whether this only happens for online AC with linear function approximation setting, and the answer is no. We show the stepsize sensitivity plots for REINFORCE on the chain environment (the same experiment as discussed in §E.2 and §E.3) with a wide range of the policy stepsize in Figure 49. And these plots display a similar pattern: the alternate estimator learns for extremely large values of the policy stepsize as well. The first two rows of Figure 49 are exactly the same as Figure 29 except that it includes a larger sweep for the policy stepsize $\alpha \in \{2^{-6}, 2^{-4}, \dots, 2^{14}\}$. Similarly, the third and fourth rows are identical to Figure 32 and Figure 33 respectively. We verified that none of these results were due to bugs in our experiments and now discuss some plausible explanations for this behavior.




Figure 48: Online AC + neural networks with different baseline initializations and policy saturations on the DotReacher task. The plots show the mean final performance during the last $5k$ timesteps of the agent’s $50k$ timestep lifetime. The figure header shows the action preference initialization for the no movement action; the rest of the preferences were initialized to zero. Different rows show differently initialized value function estimates (0, +25, or -25). We observe that for almost all the settings, the alternate estimator performed better than the regular estimator in recovering from saturated policies, especially for higher saturation magnitudes.

We ascribe these observations to the nature of softmax function and the policy optimization problem over the action preferences. Note that (1) the softmax function transforms any set of action preferences, irrespective of how large they are, to a proper probability distribution; and (2) for softmax policies, the optimal point which maximizes the PG objective, lies at infinity on the action preference landscape. The second point is true because for MDPs, the optimal policy is greedy, which requires the action preference for the greedy action to be infinitely large. This makes the optimization problem over the action preferences very different from, say, an unconstrained convex optimization problem. For convex functions, an update step with a high stepsize, even in the correct direction, can make the function value worse than its present value. In contrast, a larger stepsize with softmax functions will only drive the function value closer to the optimum at a faster rate. We illustrate this point in Figure 50. Therefore, as long as the update direction is correct, having large stepsizes should not be a concern for softmax policies.

For concreteness, let us consider the stepsize sensitivity behavior of expected PG in Figure 49: a larger stepsize for expected PG update resulted in a faster convergence to the optimal policy and consequently a higher final return. From gradient based optimization literature, we know that the expected (policy) gradient is the the update direction that results in the fastest local improvement in the (policy) objective, and therefore using the reasons given in the previous paragraph, we would expect a large stepsize to drive the policy towards this local optimum faster. We can also observe this trend for the expected gradient bandit algorithm. Figure 13 (bottom-left) shows that for a bandit problem, a large policy update stepsize should still make the policy converge towards the optimal (top) corner, albeit via a different route. But for that bandit problem, there is only one optimum that is also globally optimal. Also, intuitively, the policy gradient update in the (tabular) MDP setting can be considered as (kind of) making separate gradient bandit updates for each state.

However, due to the sampling noise, the stochastic PG estimators might not always give an update direction that points towards the local optimum. Consequently, a very large stepsize could saturate the policy in a sub-optimal manner. We posit two reasons for why the alternate estimator might be better suited to handle very large policy




Figure 49: Sensitivity plots showing the final performance of **tabular** REINFORCE on the chain environment. The different columns correspond to an increasing degree of policy saturation and the figure headers show the action preference initializations for the **left** and **right** actions respectively. The first row shows the results for the true baseline v_π , and the next three rows show results for a learned baseline, which was initialized to either zero, +4, or -4 at the beginning of learning.




Figure 50: Effect of the stepsize magnitude on optimizing a function value. We consider the optimization problem on a convex function and the logistic function. Note that the logistic function $y = \frac{e^x}{e^x + 1}$ is essentially equivalent (except for a translation along the x -axis) to the softmax function $\pi(a) = \frac{e^{\theta_a}}{e^{\theta_a} + \sum_{c \in \mathcal{A} \setminus a} e^{\theta_c}}$. This figure shows that given the correct update direction (positive direction for both of these problems), an update step with a very large stepsize on the convex function can make the function value worse. Whereas, in the case of a softmax function, the larger the stepsize, the more the function value increases with a single update step.

stepsizes than the regular estimator. The first reason is that the alternate estimator can handle and escape sub-optimally saturated policies much better than the regular estimator. With very large stepsizes, the action preferences for both the estimators would be driven to large values and the policy would become saturated. For such cases, the regular gradient estimator becomes vanishingly small: $\hat{\mathbf{g}}^{\text{REG}}(S, A) = (R - b)(\mathbf{e}_A - \boldsymbol{\pi}) \approx 0$, since near the corners $\mathbf{e}_A \approx \boldsymbol{\pi}$ with a high probability. Therefore, even with a very large stepsize α , the net update to the action preferences $\boldsymbol{\theta}^{\text{new}} = \boldsymbol{\theta}^{\text{old}} + \alpha \hat{\mathbf{g}}^{\text{REG}}(S, A)$ remains small, and the policy doesn't change much. In contrast, for the alternate estimator $\hat{\mathbf{g}}^{\text{ALT}}(S, A) = (R - b) \mathbf{e}_A$, irrespective of the degree of policy saturation, the estimator remains non-zero, which when combined with a large stepsize can be effective in changing even very large action preferences. Although this explanation used the gradient bandit estimators, similar arguments should hold for the MDP policy gradient estimators as well.

The second point is that, since the alternate estimator updates the preference corresponding to only the sampled action, there is much less generalization in the preference values across actions when using the alternate estimator. In contrast, the regular estimator updates all the action preferences and therefore has a high generalization across the action space. This distinction becomes important when the agent updates the policy using an incorrect direction and a large stepsize. With the alternate estimator, only some action preferences for some states get saturated, whereas with the regular estimator, all the preferences for a set of states get saturated. This generalization argument is also supported by the experiments presented in this document. With large stepsizes, the performance of the estimators, especially the regular version, becomes worse as we increase the generalization of the policies: compare the performance of the two estimators in Figure 49 (tabular policy: no generalization); Figures 44 and 46 (linear policy with tile-coding: generalization across 8 tilings); and Figure 48 (neural policy: high generalization). With neural networks, none of the estimators performed well with large stepsizes because updating even a single action preference (in the wrong direction) changes all the weights lying in the initial layers of the network, ultimately affecting all the action preferences.

One experiment to test the above hypothesis could be to run online AC with both the estimators on Acrobot and MountainCar with a tabular policy. This could be achieved with a tile-coding that has super-fine meshes (that is a large number of tiles per dimension) and only a single tiling. If the performance of the regular estimator with high stepsizes improves as compared to Figure 44, then that would serve as a strong verification for this hypothesis. Unfortunately, a detailed investigation of this phenomenon is out of scope of this work, and therefore we do not pursue it any further.