

Global Optimality Guarantees For Policy Gradient Methods

Jalaj Bhandari* and Daniel Russo

Columbia University

Abstract

Policy gradient methods are perhaps the most widely used class of reinforcement learning algorithms. These methods apply to complex, poorly understood, control problems by performing stochastic gradient descent over a parameterized class of policies. Unfortunately, even for simple control problems solvable by classical techniques, policy gradient algorithms face non-convex optimization problems and are widely understood to converge only to local minima. This work identifies structural properties – shared by finite MDPs and several classic control problems – which guarantee that policy gradient objective function has no suboptimal local minima despite being non-convex. When these assumptions are relaxed, our work gives conditions under which any local minimum is near-optimal, where the error bound depends on a notion of the expressive capacity of the policy class.

1 Introduction

Many recent successes in reinforcement learning are driven by a class of algorithms called policy gradient methods. These methods search over a parameterized class of policies by performing stochastic gradient descent on a cost function capturing the cumulative expected cost incurred. Specifically, they aim to optimize over a smooth, and often stochastic, class of parametrized policies $\{\pi_\theta\}_{\theta \in \mathbb{R}^d}$. For discounted or episodic problems, they treat the scalar cost function $\ell(\theta) = \int J_{\pi_\theta}(s) d\rho(s)$, which averages the total cost-to-go function J_{π_θ} over a random initial state distribution ρ . Policy gradient methods perform stochastic gradient descent on $\ell(\cdot)$, following the iteration

$$\theta_{k+1} = \theta_k - \alpha_k (\nabla \ell(\theta_k) + \text{noise}).$$

Unfortunately, even for simple control problems solvable by classical methods, the total cost ℓ is a non-convex function of θ . Typical of results concerning the black-box optimization of non-convex functions, policy gradient methods are widely understood to converge asymptotically to a stationary point or a local minimum. Important theory guarantees this under technical conditions [4, 36, 57] and it is widely repeated in textbooks and surveys [21, 43, 56].

The reinforcement learning literature seems to provide almost no guarantees into the *quality* of the points to which policy gradient methods converge. Although these methods can be applied to a very broad class of problems, it is not clear whether they adequately address even simple and classical dynamic programming problems. Inspired by this disconnect, important recent work of Fazel et al. [16], showed that policy gradient on the space of linear policies for deterministic linear quadratic control problem converges to the global optimum, despite the non-convexity of the objective. The authors provided an intricate analysis in this case, leveraging a variety of closed form expressions available for linear-quadratic problems. Separate from the RL literature, Kunnumkal and Topaloglu [32] propose a stochastic approximation method for setting base-stock levels in inventory control. Surprisingly, despite non-convexity of the objective, an intricate analysis quite different that from Fazel et al. [16] establishes convergence to the global optimum.

*<http://www.columbia.edu/~jb3618/>

Our work aims to construct a simple and more general understanding of the global convergence properties of policy gradient methods. As a consequence of our general framework, we can show that for several classic dynamic programming problems, policy gradient methods performed with respect to natural structured policy classes faces no suboptimal local minima. More precisely, despite its non-convexity, any stationary point¹ of the policy gradient cost function is a global optimum. The examples we treat include:

Example 1. Softmax policies applied in finite state and action MDPs: *Here, with n states and k actions, $\theta \in \mathbb{R}^{kn}$. The policy π_θ associates each state s with a probability distribution $(\pi_\theta(s, 1), \dots, \pi_\theta(s, k))$ over actions, with $\pi_\theta(s, i) = e^{\theta_{si}} / \sum_{j=1}^k e^{\theta_{sj}}$. This set of policies contains all possible stochastic policies and its closure contains all possible policies.*

Example 2. Linear policies applied in linear quadratic control: *Here, actions $a_t \in \mathbb{R}^k$ and states $s_t \in \mathbb{R}^n$ are vectors, states evolve according to $s_{t+1} = As_t + Ba_t + w_t$ where w_t is i.i.d Gaussian noise², and the goal is to minimize the cumulative discounted cost $\mathbb{E} \sum_{t=0}^{\infty} \gamma^t (a_t^T R a_t + s_t^T K s_t)$ for positive definite matrices R and K . It is known that a linear policy of the form $\pi_\theta(s) = \theta s$ for $\theta \in \mathbb{R}^{k \times n}$, is optimal for this problem. We assume (A, B) are controllable, in which case the set of stable linear policies $\Theta_{\text{stable}} := \{\theta \in \mathbb{R}^{k \times n} : \max_{x: \|x\|_2=1} \|(A + B\theta)x\|_2 < 1\}$ is nonempty.*

Example 3. Threshold policies applied in an optimal stopping problem: *One classic optimal stopping problem is an asset selling problem, where at every time t , an agent observes i.i.d offers $y_t \in \mathbb{R}$ and chooses a stopping time τ with the goal of maximizing $\mathbb{E}[\gamma^\tau y_\tau]$. We consider a somewhat richer contextual variant of this problem. In each round, the agent passively observes contextual information, x_t which evolves according to an uncontrolled Markov chain with finite state space. The context reflects variables like the weather or economic indicators, which are not influenced by the offers but inform the likelihood of receiving high offers. Conditioned on the context x_t , y_t is drawn i.i.d from some bounded distribution $q_{x_t}(\cdot)$. The agent's objective is to solve $\sup_\tau \mathbb{E}[\gamma^\tau y_\tau]$ where the supremum is taken over stopping times adapted to the observations $\{(x_t, y_t)\}$. There are standard ways to cast such a stopping problem as an MDP with a particular state-space. (See [5] or Appendix E.2.) The optimal policy in this setting has a threshold for each context x , and accepts an offer in that context if and only if it exceeds the threshold. To accommodate cases where the set of possible offers is discrete, while still using smooth policies, we consider randomized policies that map a state to a probability of accepting the offer, $\pi_\theta(x, y) \in [0, 1]$. For a vector $\theta = (\theta_0^x, \theta_1^x)_{x \in \mathcal{X}}$ we set $\pi_\theta(x, y) = f(\theta_0^x + \theta_1^x y)$ where $f(z) \equiv 1/(1 + e^{-z})$ is the logistic function. While this policy is similar to the one in Example 1, it leverages the structure of the problem and hence has only $d := 2|\mathcal{X}|$ parameters even if the set of possible offers is infinite.*

Example 4. Base-stock policies applied in finite horizon inventory control: *The example we treat is known as a multi-period newsvendor problem with backlogged demands. The state of a seller's inventory evolves according to $s_{t+1} = s_t + a_t - w_t$ where a_t is the quantity of inventory ordered from a supplier and $w_t \in [0, w_{\max}]$ is the random demand at time t . Negative values of s_t indicate backlogged demand that must be filled in later periods. We allow for continuous inventory and order levels. Here we consider a finite horizon objective of minimizing $\mathbb{E} \sum_{t=1}^{H-1} (ca_t + b \max\{s_t + a_t - w_t, 0\} + p \max\{-s_t + a_t - w_t, 0\})$, where c is per-unit ordering cost, b is a per-unit holding cost, and p is a per-unit cost for backlogged demand. Only non-negative orders $a_t \geq 0$ are feasible. For a finite horizon problem, we consider the class of time-inhomogenous base-stock policies, which are known to contain the optimal policy. Here $\theta = (\theta_1, \dots, \theta_{H-1})$ is a vector, and at time t such a policy orders inventory $a_t = \max\{0, \theta_t - s_t\}$. That is, it orders enough inventory to reach a target level θ_t , whenever feasible.*

¹Any point with $\nabla f(x) = 0$ is a stationary point of the function f

²The work of Fazel et al. [16] considers LQ control with a random initial state but does not consider noisy dynamics. Their objective is the total undiscounted cost-to-go over an infinite horizon. With noisy dynamics, this objective is infinite under all policies. We introduce discounting to keep the total cost-to-go finite.

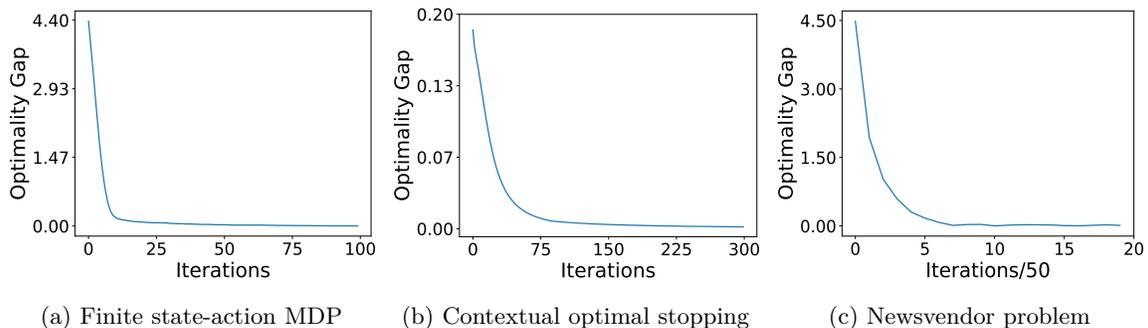


Figure 1: Global convergence of gradient descent on $\ell(\cdot)$ for Examples 1, 3 and 4.

For each of these examples, simple experiments show that gradient descent with backtracking line search performed on $\ell(\theta)$ converges rapidly to the global minimum. Sample plots for three of the problems are shown in Figure 1. For linear quadratic control we refer readers to Figure 1 in [16]. We have shared code [here](#) for reproducibility and full experiment details are also given in Appendix F.

Our work aims to understand this phenomenon. Why does gradient descent on a non-convex function reach the global minimum? These examples share important structural properties. Consider a linear quadratic control problem. Starting with a linear policy and performing a policy iteration step yields another linear policy. That is, the policy class is *closed* under policy improvement. In addition, although the cost-to-go function is a nasty non-convex function of the policy, the policy iteration update involves just solving a quadratic minimization problem. In fact, for each of the first three examples, the policy class is closed under policy improvement and the policy-iteration objective (i.e. the Q -function) is smooth and convex in the chosen action. Similar ideas, apply to the fourth example, but as shown in Theorem 2, weaker conditions are needed to ensure convergence for some finite-horizon problems. Given this insight, strikingly simple proofs show that any stationary point of the cost function $\ell(\theta)$ is a global minimum.

In our view, these canonical control problems provide an important benchmark and sanity check for policy gradient methods. At the same time, one hopes that the insights developed from considering these problems extend to more complex scenarios. To spur progress in this direction, we take a first step in Section 5 where we relax the assumption that the policy class is closed under policy improvement. Our theory gives conditions under which any stationary point of $\ell(\cdot)$ is nearly optimal, where the error bound depends on a notion of the expressive capacity of the policy class.

Beyond RL, this work connects to a large body of work on first-order methods in non-convex optimization. Under broad conditions, these methods are guaranteed to converge asymptotically to stationary points of the objective function under a variety of noise models [9, 10]. The ubiquity of non-convex optimization problems in machine learning and especially deep learning has sparked a slew of recent work [1, 14, 24, 34] giving rates of convergence and ensuring convergence to approximate local minima rather than saddle points. A complementary line of research studies the optimization landscape of specific problems to essentially ensure that local minima are global, [11, 18, 19, 28, 55]. Taken together, these results show interesting non-convex optimization problems can be efficiently solved using gradient descent. Our work contributes to the second line of research, offering insight into the optimization landscape of $\ell(\cdot)$ for classic dynamic programming problems.

Challenges with policy gradient methods and the scope of this work. There are many reasons why practitioners may find simple policy gradient methods, like the classic REINFORCE algorithm reviewed in Appendix A, offer poor performance. In an effort to clarify the scope of our contribution, and its place in the literature, let us briefly review some of these challenges.

1. *Non-convexity of the loss function:* Policy gradient methods apply (stochastic) gradient descent

on a non-convex loss function. Such methods are usually expected to converge toward a stationary point of the objective function. Unfortunately, a general non-convex function could have many stationary points that are far from optimal.

2. *Unnatural policy parameterization:* It is possible for parameters that are far apart in Euclidean distance to describe nearly identical policies. Precisely, this happens when the Jacobian matrix of the policy $\pi_\theta(\cdot | s)$ vanishes or becomes ill conditioned. Researchers have addressed this challenge through natural gradient algorithms [2, 26], which perform steepest descent in a different metric. The issue can also be alleviated with regularized policy gradient algorithms [50, 52].
3. *Insufficient exploration:* Although policy gradients are often applied with stochastic policies, convergence with this kind of naive random exploration can require a number of iterations that scales exponentially with the number of states in the MDP. Kakade and Langford [25] provide a striking example. Combining efficient exploration methods with policy gradient algorithms is challenging, but is an active area of research [see e.g. 40, 44].
4. *Large variance of stochastic gradients:* The variance of estimated policy gradients generally increases with the problem’s effective time horizon, usually expressed in terms of a discount factor or the average length of an episode. Considerable research is aimed at alleviating this problem through the use of actor-critic methods [31, 36, 57] and appropriate baselines [37, 51].

We emphasize that this paper is focused on the first challenge and on understanding the risks posed by spurious local minima. Such an investigation is relevant to many strategies for searching locally over the policy space, including policy gradient methods, natural gradient methods [26], finite difference methods [47], random search [35], and evolutionary strategies [49]. For concreteness, one can mostly have in mind the idealized policy gradient iteration $\theta_{k+1} = \theta_k - \alpha_k \nabla \ell(\theta_k)$. As in the REINFORCE algorithm in Appendix A, we imagine applying policy gradient algorithms in simulation, where an appropriate restart distribution ρ provides sufficient exploration.

A natural direction for future work would be to analyze the rate of convergence of specific algorithms that follow noisy gradient steps on $\ell(\cdot)$. Fazel et al. [16] give an impressive analysis of several exact gradient-based algorithms for deterministic linear quadratic control, along with an extension to zeroth order optimization for approximate gradients. We leave this for future work.

2 Problem formulation

Consider a Markov decision process (MDP), which is a six-tuple $\mathcal{M} := (\mathcal{S}, \mathcal{A}, g, P, \gamma, \rho)$. In some of our examples, the state space \mathcal{S} is a convex subset of \mathbb{R}^n , but to ease notation we present some of the notations below assuming \mathcal{S} is countable, trusting readers can substitute sums for integrals when needed. The initial state distribution ρ is a probability distribution supported on the entire state space. For each $s \in \mathcal{S}$, the set of feasible actions is $\mathcal{A}(s) \subset \mathbb{R}^k$. When action a is executed in state s , the agent incurs some immediate cost and transitions to a new state. The instantaneous cost function $g(s, a)$ specifies the expected cost incurred and the transition kernel specifies probability $P(s' | s, a)$ of transitioning to state s' in the next period. We assume expected costs are non-negative for all feasible state-action pairs. A policy π is a mapping from states to feasible actions. For each π , the cost-to-go function $J_\pi(s) = \mathbb{E}_\pi [\sum_{t=0}^{\infty} \gamma^t g(s_t, \pi(s_t)) | s_0 = s]$ encodes the expected total discounted cost incurred when applying policy π from initial state s . Here the expectation is taken over the sequence of states $(s_0, s_1, s_2 \dots)$ visited under π , since the function g already integrates over any randomness in instantaneous costs. The state-action cost-to-go function

$$Q_\pi(s, a) = g(s, a) + \gamma \sum_{s' \in \mathcal{S}} P(s' | s, a) J_\pi(s')$$

measures the cumulative expected cost of taking action a in state s and applying π thereafter. We let $J^*(s) = \inf_{\pi} J_{\pi}(s)$ denote the optimal cost-to-go function. For every problem we consider, J^* is the unique solution to the Bellman equation $J = TJ$, where the Bellman operator T associates each function $J : \mathcal{S} \rightarrow \mathbb{R}$ with another function $TJ : \mathcal{S} \rightarrow \mathbb{R}$ defined as

$$TJ(s) = \min_{a \in \mathcal{A}(s)} \left[g(s, a) + \gamma \sum_{s' \in \mathcal{S}} P(s' | s, a) J(s') \right].$$

Similarly, define $T_{\pi}J(s) = g(s, \pi(s)) + \gamma \sum_{s' \in \mathcal{S}} P(s' | s, \pi(s)) J(s')$. We assume throughout that $\mathcal{A}(s)$ is convex. In some settings, like linear quadratic control problems, this is natural in all problem formulations. In others, like MDPs with a finite set of actions, the action set is convexified by randomization. In particular, when there are k deterministic actions $\{1, \dots, k\}$ feasible in each state, we will take $\mathcal{A}(s) = \Delta^{k-1}$ to be the $k - 1$ dimensional probability simplex. Cost and transition functions are naturally extended to functions on the simplex defined by $g(s, a) = \sum_{i=1}^k g(s, i) a_i$ and $P(s' | s, a) = \sum_{i=1}^k P(s' | s, i) a_i$. Policy gradient methods search over a parameterized class of policies $\Pi = \{\pi_{\theta}(\cdot) : \theta \in \mathbb{R}^d\}$. When considering softmax policies for finite state and action MDPs as in Example 1, we take Π to consist of the closure of this set, in which case it contains all stationary policies. We assume throughout that $\pi_{\theta}(s)$ is differentiable as a function of θ . We overload notation, writing $J_{\theta}(s) := J_{\pi_{\theta}}(s)$ and $Q_{\theta}(s, a) := Q_{\pi_{\theta}}(s, a)$ for each $\theta \in \mathbb{R}^d$. Although classical dynamic programming methods seek a policy that minimizes the expected cost incurred *from every* initial state, for policy gradient methods it is more natural to study a scalar loss function $\ell(\theta) = \sum_{s \in \mathcal{S}} J_{\theta}(s) \rho(s)$ under which states are weighted by their initial probabilities under ρ . The discounted state-occupancy measure under ρ and π_{θ} is defined as $\eta_{\theta}(s) = (1 - \gamma) \sum_{t=0}^{\infty} \gamma^t \mathbb{P}_{\pi_{\theta}}(s_t = s | s_0 \sim \rho)$ where the subscript indicates that transition probabilities are evaluated under the Markov chain that results from applying π_{θ} . We often consider the weighted 1-norm, $\|J\|_{1, \eta_{\theta}} = \sum_s |J(s)| \eta_{\theta}(s)$.

3 General results

The introduction described in words some special structural properties shared by our motivating examples. This section states formal assumptions capturing that intuition and culminates in a strikingly simple proof that such conditions ensure that $\ell(\cdot)$ has no suboptimal stationary points.

Assumption 1 (Closure under policy improvement). *For any $\pi \in \Pi$, there is $\pi_+ \in \Pi$ such that for every $s \in \mathcal{S}$, $\pi_+(s) = \arg \min_{a \in \mathcal{A}(s)} Q_{\pi}(s, a)$.*

Assumption 2 (Convexity of policy improvement steps). *For every $\pi \in \Pi$ and $s \in \mathcal{S}$, $Q_{\pi}(s, a)$ is a convex function of $a \in \mathcal{A}(s)$.*

Next, we assume the policy class Π is convex, ensuring that a soft policy-iteration update from the policy π to the $(1 - \alpha)\pi + \alpha\pi_+$ is feasible. In addition to this somewhat stringent assumption on the policy class, we need a mild regularity property of the parametrization. To make this assumption more transparent, let's look at Examples 1 and 2 with softmax and linear policies respectively. Consider any two policies π_{θ} and $\pi_{\theta'}$. The goal is to find a direction u in the parameter space such that the directional derivative of π_{θ} along u points in the direction of $\pi_{\theta'}$. Since the map $\theta \mapsto \pi_{\theta}$ is one-to-one³ and convexity ensures $(1 - \alpha)\pi_{\theta} + \alpha\pi_{\theta'} \in \Pi$, we can pick θ^{α} such that $\pi_{\theta^{\alpha}} = (1 - \alpha)\pi_{\theta} + \alpha\pi_{\theta'}$. Varying $\alpha \in [0, 1]$ traces out a line segment in policy space and a smooth curve in parameter space. Then the desired direction u satisfies $u = \lim_{\alpha \rightarrow 0} \frac{\theta^{\alpha} - \theta}{\alpha}$. In the case of linear quadratic control, the direction can be expressed simply as $u = \theta' - \theta$. For softmax policies, the existence of u follows from an inverse function theorem, which ensures the differentiable map $\theta \mapsto \pi_{\theta}(s)$ has differentiable inverse. There, the direction $u = (u^s)_{s \in \mathcal{S}}$ is a concatenation of $|\mathcal{S}|$ vectors of length $k = |\mathcal{A}|$. Each

³Softmax policies are one-to-one with a common parameterization that fixes a single component of θ per state. Otherwise, we can follow the argument above, with an appropriate rule for selecting θ^{α} when multiple exist.

u^s solves the linear system $[\frac{\partial \pi_\theta(s)}{\partial \theta}]u^s = \pi_{\theta'}(s) - \pi_\theta(s)$ where $\frac{\partial \pi_\theta(s)}{\partial \theta}$ is the Jacobian matrix. This parallels the construction of natural gradient directions [26].

Assumption 3 (Convexity of the policy class). *Assume Π is convex. Moreover, for any policy $\pi \in \Pi$ and any $\theta \in \mathbb{R}^d$, there exists $u \in \mathbb{R}^d$ such that for every $s \in \mathcal{S}$, $\frac{d}{d\alpha} \pi_{\theta+\alpha u}(s)|_{\alpha=0} = \pi(s) - \pi_\theta(s)$.*

Finally, the policy gradient theorem [36, 57] requires that certain limits and integrals can be interchanged. In specific applications, this is often easy to justify. Here we state some general, though potentially stringent, regularity conditions that allows us to simply apply a general policy gradient theorem due to [53], stated for directional derivatives in Lemma 1. For specific applications like linear quadratic control, the interchange of limits and expectations can be easily verified and we don't need this assumption.

Assumption 4 (Regularity conditions by [53]). *$\mathcal{S} \subset \mathbb{R}^m$ is compact, and $\nabla_a P(s' | s, a)$, $\nabla_\theta \pi_\theta(s)$, $g(s, a)$, $\nabla_a g(s, a)$, and $\rho(s)$ exists and are jointly continuous in θ, s, a and s' .*

Lemma 1 (Policy gradients for directional derivatives). *Under assumption 4, for any $\theta, u \in \mathbb{R}^d$,*

$$\frac{d}{d\alpha} \ell(\theta + \alpha u) \Big|_{\alpha=0} = (1 - \gamma)^{-1} \sum_{s \in \mathcal{S}} \eta_\theta(s) \frac{d}{d\alpha} Q_\theta(s, \pi_{\theta+\alpha u}(s)) \Big|_{\alpha=0}. \quad (1)$$

The following theorem shows that $\ell(\cdot)$ has no suboptimal stationary points by constructing a specific descent direction. The direction is chosen to point toward the policy gradient update, and we show the corresponding directional derivative scales with the average magnitude of the Bellman error $J_\theta(s) - TJ_\theta(s)$ weighted under the state occupancy distribution η_θ .

Theorem 1 (No spurious local minima). *Under Assumptions 1-4, for some policy π_θ , let $\pi_+ \in \Pi$ be a policy iteration update defined by $\pi_+(s) \in \arg \min_{a \in \mathcal{A}(s)} Q_\theta(s, a)$. Take u to satisfy*

$$\frac{d}{d\alpha} \pi_{\theta+\alpha u}(s) = \pi_+(s) - \pi_\theta(s) \quad \forall s \in \mathcal{S}. \quad (2)$$

Then,

$$\frac{d}{d\alpha} \ell(\theta + \alpha u) \Big|_{\alpha=0} \leq -(1 - \gamma)^{-1} \|J_\theta - TJ_\theta\|_{1, \eta_\theta}.$$

Proof. After applying the policy gradient theorem stated in Lemma 1, our goal is to bound $\frac{d}{d\alpha} Q_\theta(s, \pi_{\theta+\alpha u}(s)) \Big|_{\alpha=0}$. Let $\langle x, y \rangle = x^\top y$ denote the standard inner product. We have

$$\begin{aligned} \frac{d}{d\alpha} Q_\theta(s, \pi_{\theta+\alpha u}(s)) \Big|_{\alpha=0} &= \left\langle \frac{\partial}{\partial a} Q_\theta(s, a) \Big|_{a=\pi_\theta(s)}, \frac{d}{d\alpha} \pi_{\theta+\alpha u}(s) \Big|_{\alpha=0} \right\rangle && \text{Chain rule} \\ &= \left\langle \frac{\partial}{\partial a} Q_\theta(s, a) \Big|_{a=\pi_\theta(s)}, \pi_+(s) - \pi_\theta(s) \right\rangle && \text{By (2)} \\ &\leq Q_\theta(s, \pi_+(s)) - Q_\theta(s, \pi_\theta(s)) && \text{Convexity of } Q_\theta(s, \cdot) \\ &= TJ_\theta(s) - J_\theta(s). \end{aligned}$$

The final inequality follows from the first order condition for convex differentiable function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ which implies $f(y) \geq f(x) + \nabla f(x)^\top (y - x)$. We use the fact that $J_\theta(s) = T_{\pi_\theta} J_\theta(s) \geq TJ_\theta(s) \forall s \in \mathcal{S}$ in conjunction with Lemma 1 to get,

$$\frac{d}{d\alpha} \ell(\theta + \alpha u) \Big|_{\alpha=0} \leq -(1 - \gamma)^{-1} \sum_{s \in \mathcal{S}} \eta_\theta(s) |TJ_\theta(s) - J_\theta(s)| = -(1 - \gamma)^{-1} \|TJ_\theta - J_\theta\|_{1, \eta_\theta}.$$

□

An immediate corollary is that, under Assumptions 1-4, if $\nabla\ell(\theta) = 0$ then $J_\theta(s) = TJ_\theta(s)$ almost surely over s drawn from ρ . Textbooks on dynamic programming provide different technical conditions under which any policy whose cost-to-go function solves Bellman’s equation must be optimal. This holds when T is a contraction, but also in many settings where T is not a contraction [7]. This applies immediately to Examples 1 and 3. For linear quadratic control as formulated in Example 2, any *stable* linear policy π_θ satisfies $J_\theta = TJ_\theta$ if and only if it is the optimal policy.

Relaxing Assumption 1 for finite horizon problems. For finite horizon problems, we can guarantee that there are no spurious local minima for policy gradient under a much weaker condition. Rather than require the policy class is closed under improvement – which would imply the policy class contains the optimal policy – it is sufficient that the policy class contain the optimal policy. For this reason, our theory will cover as special cases a broad variety of finite horizon dynamic programming problems for which structured policy classes are known to be optimal.

Unfortunately, we do not have space in this short paper to develop specialized notation for finite horizon problems. We do so and give a more detailed treatment in Appendix C. We can state our formal result without rewriting our problem formulation, by a well known trick that treats finite-horizon time-inhomogenous MDPs as a special case of infinite horizon MDPs (see e.g. [42]). Under the following assumption, the state space factorizes into $H+1$ components, thought of as stages or time periods of the decision problem. Under any policy, any state will transition to some state in the next stage until stage $H+1$ is reached and the interaction effectively ends. This assumption on the policy class allows us to change the policy in stage h without influencing the policy at other stages, essentially encoding time-inhomogenous policies.

Assumption 5 (Finite horizon). *Suppose the state space factors as $\mathcal{S} = \mathcal{S}_1 \cup \dots \cup \mathcal{S}_H \cup \mathcal{S}_{H+1}$, where for a state $s \in \mathcal{S}_h$ with $h \leq H$, $\sum_{s' \in \mathcal{S}_{h+1}} p(s'|s, a) = 1$ for all $a \in \mathcal{A}(s)$. The final subset $\mathcal{S}_{H+1} = \{T\}$ contains a single costless absorbing state, with $P(T|T, a) = 1$ and $g(T, a) = 0$ for any action a . The policy parameter $\theta = (\theta^1, \dots, \theta^H)$ is the concatenation of H sub-vectors, where for any fixed $s \in \mathcal{S}_h$, $\pi_\theta(s)$ depends only on θ^h .*

Theorem 2. *Under assumptions 3, 4, and 5, if $Q^*(s, a)$ is a convex function of $a \in \mathcal{A}(s)$ for all $s \in \mathcal{S}$ and Π contains an optimal policy, then $\nabla\ell(\theta) = 0$ if and only if $\ell(\theta) = \inf_{\theta' \in \mathbb{R}^d} \ell(\theta')$.*

4 Revisiting Examples 1-4

Softmax policies in finite state and action MDPs. Consider again Example 1. Abusing notation, we could write a stochastic policy for a finite state and action MDP as long vector $\pi = [\pi(s, i)]_{s \in \mathcal{S}, 1 \leq i \leq k}$ where $\pi(s, i)$ is the probability of choosing action i in state s . Softmax policies are $\pi_\theta(s, i) = e^{\theta_{si}} / \sum_{j=1}^k e^{\theta_{sj}}$. We have $\Pi = \text{Closure} \{ \pi_\theta : \theta \in \mathbb{R}^d \}$. This contains all possible policies, so policy class is automatically closed under policy improvement (Assumption 1). In this case, Assumption 2 holds since the Q function is linear: $Q_\theta(s, a) = \sum_{i=1}^k Q_\theta(s, i) a_i$ for a probability vector $a \in \Delta^{k-1}$. The policy class is clearly convex, since the probability simplex is convex. We gave a constructive definition of u above Assumption 3.

Linear policies in linear quadratic control. Since the work of Kleinman [30] and Hewer [22], it has been known that, starting from any stable linear policy, policy iteration solves a sequence of quadratic minimization problems with solutions converging to the optimal linear policy. The conditions needed to apply each step in the proof of Theorem 1 essentially follow immediately from this classic theory. However, like this work, we need to add an appropriate qualifier to rule out unstable linear policies, under which the cost-to-go is infinite from every state and many expressions are not even defined. We provide more details in Appendix E.1, and also discuss when gradient descent on $\ell(\cdot)$ will not leave the class of stable policies.

Lemma 2. Consider Example 2. Choose ρ to be $N(0, I)$. Then, for any $\theta \in \Theta_{\text{stable}}$, if $\nabla \ell(\theta) = 0$ then $\ell(\theta) = \min_{\theta'} \ell(\theta')$.

Threshold policies in optimal stopping. In Example 3, we considered a parameterized class of soft or randomized threshold policies. We take Π to be the closure of the set of such policies, which also contains any deterministic threshold policies. This policy class is closed under policy improvement (Assumption 3): for any $a \in [0, 1]$ denoting a probability of accepting the offer and any $\pi_\theta \in \Pi$, we have

$$Q_\theta((x, y), a) = ay + (1 - a)\gamma \sum_{(x', y') \in \mathcal{S}} p(x'|x)q_x(y')J_\theta((x', y')).$$

For any state $s = (x, y)$, $\pi_+(s) := \arg \max_{a \in [0, 1]} Q_\theta(s, a) = 1$ if and only if the offer y exceeds the continuation value $\gamma \sum_{(x', y') \in \mathcal{S}} p(x'|x)q_x(y')J_\pi((x', y'))$. This means that, starting from a threshold policy, each step of policy iteration yields a new threshold policy, so the convergence of policy iteration implies threshold policies are optimal for this problem. Unfortunately, while we can essentially copy *the proof* of Theorem 1 line by line to establish Lemma 3 as shown below, it does not apply directly to this problem. The challenge is that the policy class Π is not convex, so moving on a line segment toward the policy iteration update π_+ is not a feasible descent direction. However, it is still simple to move the policy π_θ closer to π_+ , in the sense that for small α , $|\pi_{\theta+\alpha u}(s) - \pi_+(s)| < |\pi_\theta(s) - \pi_+(s)|$ at every s . The proof, given in Appendix E.2, essentially writes the formula for such a descent direction u , and then repeats each line of Theorem 1 with this choice of u .

Lemma 3. For the optimal stopping problem formulated in Example 3, $\nabla \ell(\theta) \neq 0$ for any $\theta \in \mathbb{R}^d$.

Base-stock policies in finite horizon inventory control. We consider again the multi-period newsvendor problem with backlogged demand described in Example 4. In this problem, it is known that base-stock-policies are optimal [5, 45], but the policy gradient cost function is a non convex function of the vector of base-stock levels [32]. The following lemma, proved in Appendix E.3, shows that nevertheless any stationary point of the objective function is a global minimum. The result is stated formally here in terms of the notation in Example 4. We establish this claim essentially by modifying one line in the proof of Theorem 2. The modification addresses the fact that, because a policy π_θ only orders in some states, local changes in the base-stock levels θ only changes the actions in those states. This property technically breaks the convexity of the policy class, but does not affect our construction of a descent direction on $\ell(\cdot)$. Global convergence of some online gradient methods for this problem were also established through more direct approaches in [23, 32].

Lemma 4. Consider Example 4 and define for $s \in \mathbb{R}$, $h \in \{1, \dots, H\}$, and $\theta \in \mathbb{R}^H$,

$$J_\theta((h, s)) = \mathbb{E}_{\pi_\theta} \left[\mathbb{E} \sum_{t=h}^{H-1} (ca_t + b \max\{s_t + a_t - w_t, 0\} + p \max\{-s_t + a_t - w_t, 0\}) \mid s_h = s \right].$$

Let ρ be an initial distribution supported on $\{1, \dots, H\} \times \mathbb{R}$ such that all moments of the marginal distribution $\rho(h, \cdot)$ exist for each h . Set $\ell(\theta) = \mathbb{E}_{(h, s) \sim \rho} [J_\theta((h, s))]$. Then for $\theta \in \mathbb{R}^d$, $\nabla \ell(\theta) = 0$ if and only if $\ell(\theta) = \min_{\theta' \in \mathbb{R}^H} \ell(\theta')$.

5 Approximation and expressive policy classes

So far, we have studied some classical dynamic programming problems that are ideally suited to policy iteration. The key property we used is that certain structured policy classes were closed under policy improvement, so that exact policy iteration can be performed when only considering

that policy class. Although simple structured policy classes are common in some applications of stochastic approximation based policy search [e.g. 27, 33, 60], they are not widely used in the RL literature. Instead, flexible policy classes like those parameterized by a deep neural network, a Kernel method [46], or using state aggregation [8, 17, 54] are preferred. In place of a concluding section, we leave the readers with some preliminary but interesting progress toward understanding why, for highly expressive policy classes, any local minimum of the policy gradient cost function might be near-optimal. We conjecture this theory can at least be clearly instantiated in the special case of state aggregation given in Appendix B.

Let $T_\pi(\cdot)$ denote the Bellman operator corresponding to a policy π , defined by $T_\pi J(s) = g(s, \pi(s)) + \gamma \sum_{s' \in \mathcal{S}} P(s'|s, \pi(s))J(s')$. Recall the optimal Bellman operator $T(\cdot)$ is defined by $TJ(s) = \min_{a \in \mathcal{A}(s)} g(s, a) + \gamma \sum_{s' \in \mathcal{S}} P(s'|s, a)J(s')$. Given an expressive policy class Π ,

$$\inf_{\pi \in \Pi} \|T_\pi J_{\pi_\theta} - TJ_{\pi_\theta}\|_{1, \eta_\theta} \quad (3)$$

measures the approximation error of the best approximate policy iteration update to a policy $\pi_\theta \in \Pi$. If Π satisfied Assumption 1, the approximation error would be zero since $T_\pi J_{\pi_\theta}(s) - TJ_{\pi_\theta}(s)$ for every $s \in \mathcal{S}$. Equation (3) measures the deviation from this ideal case, in a norm that weights states by the discounted-state-occupancy distribution η_θ under the policy π_θ . The first part of Theorem 3 shows that if θ is a stationary point of $\ell(\cdot)$, then the Bellman error $TJ_{\pi_\theta} - J_{\pi_\theta}$ measured in this same norm is upper bounded by the approximation error.

But when does a small Bellman error $\|TJ_{\pi_\theta} - J_{\pi_\theta}\|_{1, \eta_\theta}$ imply the policy is near optimal? This is the second part of Theorem 3. We relate the Bellman error in the supremum norm to the average Bellman error over states sampled from the initial distribution ρ . Define $\mathcal{J} = \{J_{\pi_\theta} : \theta \in \mathbb{R}^d\}$ and

$$C_\rho = \inf_{J \in \mathcal{J}} \frac{\|TJ - J\|_{1, \rho}}{\|TJ - J\|_\infty}. \quad (4)$$

The constant C_ρ measures the extent to which errors $TJ(s) - J(s)$ at some state must be detectable by random sampling, which depends both on the initial state distribution ρ and on properties of the set of cost-to-go functions \mathcal{J} . If the state space is finite, $C_\rho \leq 1/(\min_{s \in \mathcal{S}} \rho(s))$ and naturally captures the ability of the distribution ρ to uniformly explore the state space. This is similar to constants that depend on the worst-case likelihood ratio between state occupancy measures [25]. However, those constants can equal zero for continuous state problems. It seems (4) could still be meaningful for such problems⁴ since it also captures regularity properties of \mathcal{J} . The second part of Theorem 2 is reminiscent of results in the study of approximate policy iteration methods, pioneered by [3, 6, 9, 38, 39], among others. The primary differences are that (1) we directly consider an approximate policy class whereas that line of work considers the error in parametric approximations to the Q -function and (2) we make a specific link with the stationary points of a policy gradient method. The abstract framework of Kakade and Langford [25] is also closely related, though they do not study the stationary points of $\ell(\cdot)$. We refer the readers to Appendix D for the proof.

Theorem 3. *Suppose Assumptions 2-4 hold. Then,*

$$\nabla \ell(\theta) = 0 \implies \|TJ_{\pi_\theta} - J_{\pi_\theta}\|_{1, \eta_\theta} \leq \inf_{\pi \in \Pi} \|T_\pi J_{\pi_\theta} - TJ_{\pi_\theta}\|_{1, \eta_\theta} \quad (5)$$

If T is a contraction with respect to $\|\cdot\|_\infty$ with modulus γ and $\ell^ \equiv \int J^*(s)d\rho(s)$, then*

$$\nabla \ell(\theta) = 0 \implies \ell(\theta) \leq \ell^* + \frac{C_\rho}{(1-\gamma)^2} \inf_{\pi \in \Pi} \|T_\pi J_{\pi_\theta} - TJ_{\pi_\theta}\|_{1, \eta_\theta} \quad (6)$$

⁴To give some intuition, consider a quadratic function $f(s) = s^T K s$ on the unit sphere $\mathcal{S} = \{s \in \mathbb{R}^n : \|s\|_2 \leq 1\}$. For ρ denoting the uniform density of \mathcal{S} , we have $\int_{[0,1]} |f(s)|d\rho(s) \geq (1/n)\|f\|_\infty$.

References

- [1] Naman Agarwal, Zeyuan Allen-Zhu, Brian Bullins, Elad Hazan, and Tengyu Ma. Finding approximate local minima faster than gradient descent. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 1195–1199. ACM, 2017.
- [2] Shun-Ichi Amari. Natural gradient works efficiently in learning. *Neural computation*, 10(2): 251–276, 1998.
- [3] András Antos, Csaba Szepesvári, and Rémi Munos. Learning near-optimal policies with bellman-residual minimization based fitted policy iteration and a single sample path. *Machine Learning*, 71(1):89–129, 2008.
- [4] Jonathan Baxter and Peter L Bartlett. Infinite-horizon policy-gradient estimation. *Journal of Artificial Intelligence Research*, 15:319–350, 2001.
- [5] Dimitri P Bertsekas. *Dynamic programming and optimal control*, volume 1. Athena scientific Belmont, MA, 1995.
- [6] Dimitri P Bertsekas. Approximate policy iteration: A survey and some new methods. *Journal of Control Theory and Applications*, 9(3):310–335, 2011.
- [7] Dimitri P Bertsekas. *Abstract dynamic programming*. Athena Scientific Belmont, MA, 2013.
- [8] Dimitri P Bertsekas. Feature-based aggregation and deep reinforcement learning: A survey and some new implementations. *IEEE/CAA Journal of Automatica Sinica*, 6(1):1–31, 2019.
- [9] Dimitri P Bertsekas and John N Tsitsiklis. *Neuro-dynamic programming*, volume 5. Athena Scientific Belmont, MA, 1996.
- [10] Dimitri P Bertsekas and John N Tsitsiklis. Gradient convergence in gradient methods with errors. *SIAM Journal on Optimization*, 10(3):627–642, 2000.
- [11] Srinadh Bhojanapalli, Behnam Neyshabur, and Nati Srebro. Global optimality of local search for low rank matrix recovery. In *Advances in Neural Information Processing Systems*, pages 3873–3881, 2016.
- [12] Stephen Boyd and Lieven Vandenbergh. *Convex optimization*. Cambridge university press, 2004.
- [13] Steven J Bradtke, B Erik Ydstie, and Andrew G Barto. Adaptive linear quadratic control using policy iteration. In *Proceedings of 1994 American Control Conference-ACC'94*, volume 3, pages 3475–3479. IEEE, 1994.
- [14] Yair Carmon, John C Duchi, Oliver Hinder, and Aaron Sidford. Accelerated methods for nonconvex optimization. *SIAM Journal on Optimization*, 28(2):1751–1772, 2018.
- [15] Lawrence C Evans. An introduction to mathematical optimal control theory. *Lecture Notes, University of California, Department of Mathematics, Berkeley*, 2005.
- [16] Maryam Fazel, Rong Ge, Sham M Kakade, and Mehran Mesbahi. Global convergence of policy gradient methods for linearized control problems. In *Proceedings of the 35th International Conference on Machine Learning*, pages 1467–1476, 2018.
- [17] Norm Ferns, Prakash Panangaden, and Doina Precup. Metrics for finite markov decision processes. In *Proceedings of the 20th conference on Uncertainty in artificial intelligence*, pages 162–169. AUAI Press, 2004.
- [18] Rong Ge, Furong Huang, Chi Jin, and Yang Yuan. Escaping from saddle points—online stochastic gradient for tensor decomposition. In *Conference on Learning Theory*, pages 797–842, 2015.

- [19] Rong Ge, Jason D Lee, and Tengyu Ma. Matrix completion has no spurious local minimum. In *Advances in Neural Information Processing Systems*, pages 2973–2981, 2016.
- [20] Geoffrey J Gordon. Stable function approximation in dynamic programming. In *Machine Learning Proceedings 1995*, pages 261–268. Elsevier, 1995.
- [21] Ivo Grondman, Lucian Busoniu, Gabriel AD Lopes, and Robert Babuska. A survey of actor-critic reinforcement learning: Standard and natural policy gradients. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 42(6):1291–1307, 2012.
- [22] G Hewer. An iterative technique for the computation of the steady state gains for the discrete optimal regulator. *IEEE Transactions on Automatic Control*, 16(4):382–384, 1971.
- [23] Woonghee Tim Huh and Paat Rusmevichientong. Online sequential optimization with biased gradients: theory and applications to censored demand. *INFORMS Journal on Computing*, 26(1):150–159, 2013.
- [24] Chi Jin, Rong Ge, Praneeth Netrapalli, Sham M Kakade, and Michael I Jordan. How to escape saddle points efficiently. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1724–1732. JMLR. org, 2017.
- [25] Sham Kakade and John Langford. Approximately optimal approximate reinforcement learning. In *ICML*, volume 2, pages 267–274, 2002.
- [26] Sham M Kakade. A natural policy gradient. In *Advances in neural information processing systems*, pages 1531–1538, 2002.
- [27] Itir Karaesmen and Garrett Van Ryzin. Overbooking with substitutable inventory classes. *Operations Research*, 52(1):83–104, 2004.
- [28] Kenji Kawaguchi. Deep learning without poor local minima. In *Advances in neural information processing systems*, pages 586–594, 2016.
- [29] Jack Kiefer. Sequential minimax search for a maximum. *Proceedings of the American mathematical society*, 4(3):502–506, 1953.
- [30] D. Kleinman. On an iterative technique for riccati equation computations. *IEEE Transactions on Automatic Control*, 13:114 – 115, 1968.
- [31] Vijay R Konda and John N Tsitsiklis. Actor-critic algorithms. In *Advances in neural information processing systems*, pages 1008–1014, 2000.
- [32] Sumit Kunnumkal and Huseyin Topaloglu. Using stochastic approximation methods to compute optimal base-stock levels in inventory control problems. *Operations Research*, 56(3):646–664, 2008.
- [33] Pierre L’Ecuyer and Peter W Glynn. Stochastic optimization by simulation: Convergence proofs for the gi/g/1 queue in steady-state. *Management Science*, 40(11):1562–1578, 1994.
- [34] Jason D Lee, Max Simchowitz, Michael I Jordan, and Benjamin Recht. Gradient descent only converges to minimizers. In *Conference on learning theory*, pages 1246–1257, 2016.
- [35] Horia Mania, Aurelia Guy, and Benjamin Recht. Simple random search of static linear policies is competitive for reinforcement learning. In *Advances in Neural Information Processing Systems*, pages 1800–1809, 2018.
- [36] Peter Marbach and John N Tsitsiklis. Simulation-based optimization of markov reward processes. *IEEE Transactions on Automatic Control*, 46(2):191–209, 2001.
- [37] Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, pages 1928–1937, 2016.

- [38] Rémi Munos. Error bounds for approximate policy iteration. In *ICML*, volume 3, pages 560–567, 2003.
- [39] Rémi Munos and Csaba Szepesvári. Finite-time bounds for fitted value iteration. *Journal of Machine Learning Research*, 9(May):815–857, 2008.
- [40] Ofir Nachum, Mohammad Norouzi, and Dale Schuurmans. Improving policy gradient by exploring under-appreciated rewards. *CoRR*, abs/1611.09321, 2017.
- [41] Ronald Ortner and Daniil Ryabko. Online regret bounds for undiscounted continuous reinforcement learning. In *Advances in Neural Information Processing Systems*, pages 1763–1771, 2012.
- [42] Ian Osband, Benjamin Van Roy, Daniel Russo, and Zheng Wen. Deep exploration via randomized value functions. *arXiv preprint arXiv:1703.07608*, 2017.
- [43] Jan Peters and Stefan Schaal. Policy gradient methods for robotics. In *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2219–2225. IEEE, 2006.
- [44] Matthias Plappert, Rein Houthoofd, Prafulla Dhariwal, Szymon Sidor, Richard Y Chen, Xi Chen, Tamim Asfour, Pieter Abbeel, and Marcin Andrychowicz. Parameter space noise for exploration. *arXiv preprint arXiv:1706.01905*, 2017.
- [45] Martin L Puterman. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014.
- [46] Aravind Rajeswaran, Kendall Lowrey, Emanuel V Todorov, and Sham M Kakade. Towards generalization and simplicity in continuous control. In *Advances in Neural Information Processing Systems*, pages 6550–6561, 2017.
- [47] Martin Riedmiller, Jan Peters, and Stefan Schaal. Evaluation of policy gradient methods and variants on the cart-pole benchmark. In *2007 IEEE International Symposium on Approximate Dynamic Programming and Reinforcement Learning*, pages 254–261. IEEE, 2007.
- [48] John Rust. Using randomization to break the curse of dimensionality. *Econometrica: Journal of the Econometric Society*, pages 487–516, 1997.
- [49] Tim Salimans, Jonathan Ho, Xi Chen, Szymon Sidor, and Ilya Sutskever. Evolution strategies as a scalable alternative to reinforcement learning. *arXiv preprint arXiv:1703.03864*, 2017.
- [50] John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy optimization. In *International Conference on Machine Learning*, pages 1889–1897, 2015.
- [51] John Schulman, Philipp Moritz, Sergey Levine, Michael Jordan, and Pieter Abbeel. High-dimensional continuous control using generalized advantage estimation. *arXiv preprint arXiv:1506.02438*, 2015.
- [52] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [53] David Silver, Guy Lever, Nicolas Heess, Thomas Degris, Daan Wierstra, and Martin Riedmiller. Deterministic policy gradient algorithms. In *ICML*, 2014.
- [54] Satinder P Singh, Tommi Jaakkola, and Michael I Jordan. Reinforcement learning with soft state aggregation. In *Advances in neural information processing systems*, pages 361–368, 1995.
- [55] J. Sun, Q. Qu, and J. Wright. Complete dictionary recovery over the sphere i: Overview and the geometric picture. *IEEE Transactions on Information Theory*, 63(2):853–884, Feb 2017. ISSN 0018-9448. doi: 10.1109/TIT.2016.2632162.
- [56] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.

- [57] Richard S Sutton, David A McAllester, Satinder P Singh, and Yishay Mansour. Policy gradient methods for reinforcement learning with function approximation. In *Advances in neural information processing systems*, pages 1057–1063, 2000.
- [58] John N Tsitsiklis and Benjamin Van Roy. Feature-based methods for large scale dynamic programming. *Machine Learning*, 22(1-3):59–94, 1996.
- [59] Benjamin Van Roy. Performance loss bounds for approximate value iteration with state aggregation. *Mathematics of Operations Research*, 31(2):234–244, 2006.
- [60] Garrett Van Ryzin and Gustavo Vulcano. Simulation-based optimization of virtual nesting controls for network revenue management. *Operations Research*, 56(4):865–880, 2008.
- [61] Ward Whitt. Approximations of dynamic programs, i. *Mathematics of Operations Research*, 3(3):231–243, 1978.
- [62] Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256, 1992.

A Background on policy gradient methods for discounted problems

To begin, we provide a brief review of the simplest policy gradient algorithm: the REINFORCE algorithm for episodic tasks first proposed by [62]. What we present below is a special case of this algorithm, tailored to infinite horizon discounted objectives. The algorithm repeatedly interacts with an MDP with uncertain transition probabilities. Playing a policy until period H results in a trajectory of data $(s_0, a_0, c_0, s_1, \dots, s_H, a_H, c_H, s_{H+1})$ consisting of observed states s_t , actions a_t , and rewards $c_t = g(s_t, a_t)$. Policy gradient methods search over a family of policies $\{\pi_\theta : \theta \in \mathbb{R}^d\}$. REINFORCE is restricted to stochastic policies, where $\pi_\theta(a|s)$ is a smooth function of θ that determines the probability of selecting action a in state s . For any $\theta \in \mathbb{R}^d$, we can define the cumulative expected cost-to-go from state s by

$$J_\theta(s) = \mathbb{E}_{\pi_\theta} \left[\sum_{t=0}^{\infty} \gamma^t c_t \mid s_0 = s \right] = \mathbb{E}_{\pi_\theta} \left[\sum_{t=0}^H c_t \mid s_0 = s, H \sim \text{Geom}(1 - \gamma) \right],$$

where H is the number of time steps the policy is executed. The second equality simply notes the well known equivalence between optimizing an infinite horizon discounted objective and optimizing an undiscounted objective over a random geometric time-horizon. REINFORCE with restart distribution ρ can be thought of as performing stochastic gradient descent on the scalar loss $\ell(\theta) := \int J_\theta(s) d\rho(s)$. In particular REINFORCE follows

$$\theta_{k+1} = \theta_k - \alpha_k (\nabla \ell(\theta_k) + \text{noise}).$$

As shown in the algorithm box, generating some noisy but unbiased estimate of $\nabla \ell(\theta)$ is often simple when employing stochastic policies. It is sometimes also feasible for deterministic policies. This is the case when employing actor-critic methods [53] or in some special cases where differential dynamic programming techniques can be employed (See the inventory control example in Subsection F).

Algorithm 1: REINFORCE Algorithm for infinite horizon discounted MDPs

Input : Policy class π_θ parameterized by θ , discount factor $\gamma \in (0, 1)$, restart distribution $\rho(\cdot)$, step-size sequence $\{\alpha_k\}_{k \in \mathbb{N}}$

Initialize: θ_0 randomly.

for *Episodes*: $k = 0, 1, 2, \dots$ **do**

Sample $H \sim \text{Geometric}(1 - \gamma)$

/* Episode horizon length */

Sample initial state: $s_0 \sim \rho(\cdot)$

/* Using restart distribution */

Play policy π_{θ_k} for $H + 1$ periods.

Observe trajectory $\tau = (s_0, a_0, c_0, \dots, s_H, a_H, c_H, s_{H+1})$

Compute gradient estimate of average cost using a single trajectory:

$$\hat{\nabla}_\theta \ell(\theta_k) = c(\tau) \sum_{t=0}^H \nabla_\theta \log(\pi_\theta(s_t, a_t)); \quad c(\tau) : \text{total cost of trajectory}$$

Take stochastic gradient step: $\theta_{k+1} = \theta_k - \alpha_k \hat{\nabla}_\theta \ell(\theta_k)$

end

B The example of state aggregation

State aggregation is the simplest form of value function approximation employed in reinforcement learning and comes with strong stability properties [20, 58, 59]. It is common across several academic communities [e.g 48, 61]. Numerous theoretical papers carefully construct classes of MDPs with sufficient smooth dynamics, and upper bound the error from planning on a discretized state space [e.g 41]. The following example describes state-aggregation in policy space. It satisfies all of our assumptions other than closure under policy improvement, but we expect it can be shown to be approximately closed under policy improvement.

Example 5 (Softmax policies with state aggregation). *There are a finite number of deterministic actions k , so we take $\mathcal{A} = \Delta^{k-1}$ to be the set of probability distributions over actions. $\mathcal{S} \subset \mathbb{R}^n$ is a bounded convex subset of euclidean space and the dimension n is thought to be small. Reward functions and state transitions probabilities are smooth in s . We therefore expect an effective action in some state s will be effective in another state s' if $\|s - s'\|$ is sufficiently small. We partition the state space $\mathcal{S} = \cup_{i=1}^m \mathcal{S}_i$ into m disjoint subsets. We consider a modified softmax policy π_θ where $\theta = (\theta_{ij} : 1 \leq i \leq m, 1 \leq j \leq k)$. If $s \in \mathcal{S}_i$ lies in the i th subset of the state partition, $\pi_\theta(s)$ plays action $j \in \{1, \dots, k\}$ with probability $e^{\theta_{ij}} / \sum_{\ell=1}^k e^{\theta_{i\ell}}$.*

C Proof of Theorem 2 and formulation of finite horizon problems

In Section 3, we stated our result by treating finite-horizon time-inhomogenous MDPs as a special case of infinite horizon MDPs. For a clearer understanding, we reformulate the finite horizon problem with specialized notation along with restating all the assumptions we need. We then restate Theorem 2 in our new notation and give a proof.

First, let us briefly provide some details to clarify the equivalence. We assumed the state-space factorizes as follows.

Assumption 5 (Finite horizon). *Suppose the state space factors as $\mathcal{S} = \mathcal{S}_1 \cup \dots \cup \mathcal{S}_H \cup \mathcal{S}_{H+1}$, where for a state $s \in \mathcal{S}_h$ with $h \leq H$, $\sum_{s' \in \mathcal{S}_{h+1}} p(s'|s, a) = 1$ for all $a \in \mathcal{A}(s)$. The final subset $\mathcal{S}_{H+1} = \{T\}$ contains a single costless absorbing state, with $P(T|T, a) = 1$ and $g(T, a) = 0$ for any action a . The policy parameter $\theta = (\theta^1, \dots, \theta^H)$ is the concatenation of H sub-vectors, where for any fixed $s \in \mathcal{S}_h$, $\pi_\theta(s)$ depends only on θ^h .*

To simplify the notation, assume for the moment that each set \mathcal{S}_h is finite and $\|\mathcal{S}_1\| = \dots = \|\mathcal{S}_H\| = m$. We could express any state $s \in \cup_{t \leq H} \mathcal{S}_t$ as a unique pair (h, i) such that s is the i 'th element of \mathcal{S}_h . We now rewrite the finite horizon problem in this way.

Consider a finite Markov decision process, represented as $\mathcal{M} := (\mathcal{S}, \mathcal{A}, g, P, H, \rho)$. Over H periods, the state evolves according to a controlled stochastic dynamical system. In each period h , the agent observes the state $s_h \in \mathcal{S}$, chooses the action $a_h \in \mathcal{A}(s)$ which incurs the instantaneous expected costs $g_h(s_h, a_h)$ and transition to a new state s_{h+1} . The transition dynamics are encoded in $P = (P_1, \dots, P_H)$ where $\mathbb{P}(s_{h+1} = s \mid s_h, a_h) = P_h(s_h, a_h, s)$. We continue to assume that $\mathcal{A}(s) \subset \mathbb{R}^k$ is convex, where for finite action spaces this convexity is enforced through randomization. A policy $\pi = (\pi^1, \dots, \pi^H)$ is a sequence of functions, each of which is a mapping from $\mathcal{S} \rightarrow \mathcal{A}$ obeying the constraint $\pi^h(s) \in \mathcal{A}(s)$ for all $s \in \mathcal{S}$. For any π , the associated cost-to-go from period h and state s is defined as

$$J_\pi(h, s) = \mathbb{E} \left[\sum_{t=h}^H g_t(s_t, \pi^t(s_t)) \mid s_h = s \right].$$

We take the Q-function

$$Q_\pi(h, s, a) = g_h(s, a) + \sum_{s' \in \mathcal{S}} P^h(s, a, s') J_\pi(h+1, s'),$$

to denote the expected cumulative cost of taking action a in period h , state s , and continuing to play policy π until the end of the horizon. We set $J_\pi(H+1, s) = 0$ for notational convenience. Let J^* and Q^* denote these cost-to-go functions under the optimal policy. The distribution ρ is a probability distribution supported over $\mathcal{S} \times \{1, \dots, H\}$. (This is exact analogue under Assumption 5 assuming $\rho(s) > 0$ for all s , as we have throughout the paper).

We consider a parameterized family of policies, $\Pi = \{\pi_\theta : \theta \in \mathbb{R}^{dH}\}$ where the parameter $\theta = (\theta_1, \dots, \theta_H)$ is the concatenation of H vectors of length d and $\pi_\theta := (\pi_{\theta_1}^1, \dots, \pi_{\theta_H}^H)$ where $\pi_{\theta_h}^h : \mathcal{S} \rightarrow \mathcal{A}$ is the policy applied in period h . We define the policy class as $\Pi^h = \{\pi_{\theta_h}^h : \theta_h \in \mathbb{R}^d\}$ so $\Pi = \Pi^1 \otimes \dots \otimes \Pi^H$. Policy gradient methods seek to minimize cumulative cost-to-go by minimizing a scalar cost function,

$$\ell(\theta) = \sum_{h \in \{1, \dots, H\}} \sum_{s \in \mathcal{S}} J_\theta(h, s) \rho(h, s).$$

The basic idea for finite horizon problems remains the same - for any suboptimal policy we want to construct a descent direction. We make the following assumptions on the policy class and the Q-function.

Assumption 6 (Correctness of the policy class). *There is some $\pi \in \Pi$ such that for every $s \in \mathcal{S}$ and $h \in \{1, \dots, H\}$,*

$$\pi^h(s) \in \arg \max_{a \in \mathcal{A}(s)} Q^*(h, s, a).$$

In words, we assume that the optimal policy is contained within the policy class. As discussed in the beginning of this section, our results do not require the policy class to be closed under policy improvement, an assumption we needed for infinite horizon problems.

Assumption 7 (Convexity of policy improvement steps). *For every $s \in \mathcal{S}$ and $h \in \{1, \dots, H\}$, $Q^*(h, s, a)$ is convex in $a \in \mathcal{A}(s)$.*

Assumption 8 (Convexity of the policy class). *For each $h \in \{1, \dots, H\}$, Π^h is convex. Moreover, for any $\theta^h \in \mathbb{R}^d$ and $\pi \in \Pi^h$ there exists some $u^h \in \mathbb{R}^d$ such that*

$$\left. \frac{d}{d\alpha} \pi_{\theta^h + \alpha u^h}^h(s) \right|_{\alpha=0} = \pi^h(s) - \pi_{\theta^h}^h(s).$$

We use the policy gradient theorem from [53] as given below.

Lemma 5 (Policy gradient theorem for directional derivatives). *For any $\theta = (\theta^1, \dots, \theta^H) \in \mathbb{R}^{dH}$ and $u = (u^1, \dots, u^H) \in \mathbb{R}^{dH}$*

$$\left. \frac{d}{d\alpha} \ell(\theta + \alpha u) \right|_{\alpha=0} = \sum_{h \in \{1, \dots, H\}} \sum_{s \in \mathcal{S}} \eta_\theta(h, s) \left. \frac{d}{d\alpha} Q^\theta(h, s, \pi_{\theta^h + \alpha u^h}^h(s)) \right|_{\alpha=0}.$$

Using Lemma 5, we can show the following result. Let $\ell^* = \sum_{h=1}^H \sum_{s \in \mathcal{S}} J^*(h, s) \rho(h, s)$ denote the optimal average cost.

Theorem (Restatement of Theorem 2). *For a finite horizon MDP and policy class $\Pi = \{\pi_\theta : \theta \in \mathbb{R}^d\}$ satisfying Assumptions 6-8, $\theta \in \mathbb{R}^d$ satisfies $\nabla \ell(\theta) = 0$ if and only if $\ell(\theta) = \ell^*$.*

Proof. We assume $\ell(\theta) > \ell^*$ and construct a specific descent direction u such that

$$\left. \frac{d}{d\alpha} \ell(\theta + \alpha u) \right|_{\alpha=0} < 0.$$

Since $\ell(\theta) > \ell^*$, there exists some pair (h, s) with $J_\theta(h, s) > J^*(h, s)$. Let t be the last period in which this occurs, i.e. for $h > t$, $J_\theta(h, s) = J^*(h, s)$ for all $s \in \mathcal{S}$. For $t = H$, we have defined $J_\theta(H+1, \cdot) = J^*(H+1, \cdot) = 0$ so this is trivially satisfied. Now, we have that $Q_\theta(t, s, a) = Q^*(t, s, a)$ for every $s \in \mathcal{S}$ and $a \in \mathcal{A}(s)$, since by definition

$$Q_\theta(t, s, a) = g_t(s, a) + \sum_{s' \in \mathcal{S}} P^t(s, a, s') J_\theta(t+1, s') = g_t(s, a) + \sum_{s' \in \mathcal{S}} P^t(s, a, s') J^*(t+1, s').$$

Let π_*^t denote an optimal policy in time period t , defined by $\pi_*^t(s) \in \arg \max_{a \in \mathcal{A}(s)} Q^*(t, s, a)$ for all $s \in \mathcal{S}$. Using Assumption 8, take $u^t \in \mathbb{R}^d$ such that

$$\left. \frac{d}{d\alpha} \pi_{\theta^t + \alpha u^t}^t(s) \right|_{\alpha=0} = \pi_*^t(s) - \pi_{\theta^t}^t(s). \quad (7)$$

Let $u = (u^1, \dots, u^H)$ where $u^i = \mathbf{0} \in \mathbb{R}^d$ for all $i \neq t$. Then, using Lemma 5, we write

$$\left. \frac{d}{d\alpha} \ell(\theta + \alpha u) \right|_{\alpha=0} = \sum_{s \in \mathcal{S}} \eta_\theta(t, s) \left. \frac{d}{d\alpha} Q_\theta(t, s, \pi_{\theta^t + \alpha u^t}^t(s)) \right|_{\alpha=0} = \sum_{s \in \mathcal{S}} \eta_\theta(t, s) \left. \frac{d}{d\alpha} Q^*(t, s, \pi_{\theta^t + \alpha u^t}^t(s)) \right|_{\alpha=0}$$

Repeating the argument from proof of Theorem 1, we get

$$\begin{aligned} \left. \frac{d}{d\alpha} Q^*(t, s, \pi_{\theta^t + \alpha u^t}^t(s)) \right|_{\alpha=0} &= \left\langle \left. \frac{\partial}{\partial a} Q^*(t, s, a) \right|_{a=\pi_{\theta^t}^t(s)}, \left. \frac{d}{d\alpha} \pi_{\theta^t + \alpha u^t}^t(s) \right|_{\alpha=0} \right\rangle && \text{Chain rule} \\ &= \left\langle \left. \frac{\partial}{\partial a} Q^*(t, s, a) \right|_{a=\pi_\theta^t(s)}, \pi_*^t(s) - \pi_{\theta^t}^t(s) \right\rangle && \text{By (7)} \\ &\leq Q^*(t, s, \pi_*^t(s)) - Q(t, s, \pi_{\theta^t}^t(s)) && \text{Convexity of } Q^*(t, s, \cdot) \\ &= J^*(t, s) - J_\theta(t, s) \end{aligned}$$

Thus,

$$\left. \frac{d}{d\alpha} \ell(\theta + \alpha u) \right|_{\alpha=0} \leq - \sum_{s \in \mathcal{S}} \eta_\theta(t, s) [J_\theta(t, s) - J^*(t, s)] < 0,$$

where we know $J_\theta(t, s) - J^*(t, s) < 0$ for some $s \in \mathcal{S}$ by the construction of t . \square

D Proof of Theorem 3

Theorem 3. *Suppose Assumptions 2-4 hold. Then,*

$$\nabla \ell(\theta) = 0 \implies \|TJ_{\pi_\theta} - J_{\pi_\theta}\|_{1, \eta_\theta} \leq \inf_{\pi \in \Pi} \|T_\pi J_{\pi_\theta} - TJ_{\pi_\theta}\|_{1, \eta_\theta} \quad (5)$$

If T is a contraction with respect to $\|\cdot\|_\infty$ with modulus γ and $\ell^ \equiv \int J^*(s) d\rho(s)$, then*

$$\nabla \ell(\theta) = 0 \implies \ell(\theta) \leq \ell^* + \frac{C_\rho}{(1-\gamma)^2} \inf_{\pi \in \Pi} \|T_\pi J_{\pi_\theta} - TJ_{\pi_\theta}\|_{1, \eta_\theta} \quad (6)$$

Proof. Fix any $\theta \in \mathbb{R}^d$ and the corresponding policy $\pi_\theta \in \Pi$. Consider any other policy $\pi' \in \Pi$. We will eventually consider an approximate policy iteration update by taking an infimum over π' . Using Assumption 2, let u be a direction such that for all $s \in \mathcal{S}$

$$\frac{d}{d\alpha}\pi_{\theta+\alpha u}(s) = \pi'(s) - \pi_\theta(s).$$

By the policy gradient theorem, we have

$$\left. \frac{d}{d\alpha}\ell(\theta + \alpha u) \right|_{\alpha=0} = \frac{1}{1-\gamma} \sum_{s \in \mathcal{S}} \eta_\theta(s) \left. \frac{d}{d\alpha} Q_\theta(s, \pi_{\theta+\alpha u}(s)) \right|_{\alpha=0}.$$

Following the proof of Theorem 1, we get

$$\left. \frac{d}{d\alpha} Q_\theta(s, \pi_{\theta+\alpha u}(s)) \right|_{\alpha=0} \leq T_{\pi'} J_{\pi_\theta}(s) - J_{\pi_\theta}(s).$$

Then, using the fact that $J_{\pi_\theta}(s) = T_{\pi_\theta} J_{\pi_\theta}(s) \geq T J_{\pi_\theta}(s) \forall s \in \mathcal{S}$

$$\begin{aligned} \left. \frac{d}{d\alpha}\ell(\theta + \alpha u) \right|_{\alpha=0} &\leq (1-\gamma)^{-1} \sum_{s \in \mathcal{S}} \eta_\theta(s) [T_{\pi'} J_{\pi_\theta}(s) - J_{\pi_\theta}(s)] \\ &\leq (1-\gamma)^{-1} \sum_{s \in \mathcal{S}} \eta_\theta(s) [T_{\pi'} J_{\pi_\theta}(s) - T J_{\pi_\theta}(s) + T J_{\pi_\theta}(s) - J_{\pi_\theta}(s)] \\ &\leq (1-\gamma)^{-1} \left(\sum_{s \in \mathcal{S}} \eta_\theta(s) |T_{\pi'} J_{\pi_\theta}(s) - T J_{\pi_\theta}(s)| - \sum_{s \in \mathcal{S}} \eta_\theta(s) |T J_{\pi_\theta}(s) - J_{\pi_\theta}(s)| \right) \\ &= (1-\gamma)^{-1} \left(\|T_{\pi'} J_{\pi_\theta} - T J_{\pi_\theta}\|_{1, \eta_\theta} - \|T J_{\pi_\theta} - J_{\pi_\theta}\|_{1, \eta_\theta} \right). \end{aligned}$$

Thus, $\nabla \ell(\theta) = 0$ implies,

$$\|T J_{\pi_\theta} - J_{\pi_\theta}\|_{1, \eta_\theta} \leq \|T_{\pi'} J_{\pi_\theta} - J_{\pi_\theta}\|_{1, \eta_\theta}.$$

Taking the infimum over π' then shows

$$\|T J_{\pi_\theta} - J_{\pi_\theta}\|_{1, \eta_\theta} \leq \inf_{\pi \in \Pi} \|T_\pi J_{\pi_\theta} - T J_{\pi_\theta}\|_{1, \eta_\theta},$$

which establishes the first claim. Now, assuming T is a contraction and using the fact that $J^* = T J^*$, we have $\|J_{\pi_\theta} - J^*\|_\infty \leq \frac{1}{1-\gamma} \|J_{\pi_\theta} - T J_{\pi_\theta}\|_\infty$. This gives,

$$\begin{aligned} \ell(\theta) - \ell^* &\leq \|J_{\pi_\theta} - J^*\|_\infty \leq \frac{1}{1-\gamma} \|J_{\pi_\theta} - T J_{\pi_\theta}\|_\infty \leq \frac{C_\rho}{1-\gamma} \|J_{\pi_\theta} - T J_{\pi_\theta}\|_{1, \rho} \\ &\leq \frac{C_\rho}{(1-\gamma)^2} \|J_{\pi_\theta} - T J_{\pi_\theta}\|_{1, \eta_\theta} \\ &\leq \frac{C_\rho}{(1-\gamma)^2} \left(\inf_{\pi \in \Pi} \|T_\pi J_{\pi_\theta} - T J_{\pi_\theta}\|_{1, \eta_\theta} \right). \end{aligned}$$

□

E Details on examples

E.1 LQ Control

We consider the following LQR problem for states $s_t \in \mathbb{R}^n$ and actions $a_t \in \mathbb{R}^k$ with the goal of minimizing the cumulative discounted cost.

$$\begin{aligned} \min \quad & \mathbb{E} \sum_{t=0}^{\infty} \gamma^t (a_t^T R a_t + s_t^T K s_t) \\ \text{s.t.} \quad & s_{t+1} = A s_t + B a_t + w_t \end{aligned}$$

where w_t is i.i.d Gaussian noise, R, K are positive definite matrices and γ is the discount factor. We assume that the pair (A, B) is controllable which ensures that the optimal cumulative discounted cost is finite. We consider the class of linear policies, $\pi_\theta(s) = \theta s$ for $\theta \in \mathbb{R}^{k \times n}$ which contains the optimal policy for this problem [5, 6, 15]. The set of stable linear policies, defined as

$$\Theta_{\text{stable}} := \{\theta \in \mathbb{R}^{k \times n} : \max_{x: \|x\|_2=1} \|(A + B\theta)x\|_2 < 1,$$

is nonempty since we assumed the system to be controllable. Let's go through all the assumptions. To verify Assumption 1, we refer the readers to [7, 13, 30] which shows that the policy iteration update which solves the quadratic minimization problem above, yields a stable linear policy, $\theta_+ \in \Theta_{\text{stable}}$. In fact, a sequence of policy iteration steps can be shown to converge to the optimal linear policy. Next, to verify Assumption 2, note that the Q-function is a convex quadratic function of a and can be written as

$$Q_\theta(s, a) = a^T R a + \gamma (A s + B a)^T L (A s + B a)$$

for $L = K + \theta^T R \theta + \gamma (A + B\theta)^T L (A + B\theta)$. Finally, the policy class is trivially convex. One might imagine that the policy gradient update may lead to an unstable policy. However, starting from a suboptimal stable policy, θ and choosing the right step-size, say via backtracking line search, one can ensure a strict decrease in the average cost $\ell(\theta)$ ⁵. Thus the gradient update, $\theta' = \theta - \nabla \ell(\theta)$ remains in the sublevel set $\{\tilde{\theta} : \ell(\tilde{\theta}) < \ell(\theta)\}$, which contains only stable policies.

E.2 Threshold policies in optimal stopping

In this section, we give details of the contextual optimal stopping problem along with the proof of Lemma 3. The goal here is to maximize the expected discounted revenue earned by selling a single asset; for this we think of $\ell(\theta)$ or $Q_\theta(s, a)$ as functions we aim to maximize.

In each round, the agent passively observes contextual information, $x_t \in \mathcal{X}$ which evolve according to an uncontrolled Markov chain with transition kernel from x to x' given by $p(x'|x)$. Conditioned on context x_t , the agent receives an offer $y_t \in \mathcal{Y}$ drawn i.i.d from some distribution $q_{x_t}(\cdot)$. Here, \mathcal{X}, \mathcal{Y} denote countable subsets in \mathbb{R} . If the offer is accepted in round t , the process terminates and the decision maker accrues a reward of $\gamma^t y_t$. We assume that the decision maker receives no reward if they continue without accepting any offer. We interpret this setting by taking the context x_t as a model for the state of the economy useful for forecasting future offers, but the specific offer y_t is subject to idiosyncratic randomness beyond what is captured in x_t .

This problem can be formalized as a Markov decision process with state-space $\mathcal{S} = (\mathcal{X} \times \mathcal{Y}) \cup \{T\}$, where T denotes a special absorbing state from which no further value is accrued. The action space is $\{0, 1\}$, where action 1 corresponds to accepting the offer, in which case the decision-maker earns the current value of the item and transitions to the terminal state T . Action 0 corresponds to

⁵Our results show a descent direction exists.

rejecting the offer, the decision-maker earns no reward and transitions to a new state with transition probability:

$$\mathbb{P}(s_{t+1} = (x', y') \mid s_t = (x, y)) = p(x' \mid x)q_{x'}(y').$$

Recall that as shown in Section , the optimal policy in this setting has a threshold for each context x , and accepts an offer in that context if and only if it exceeds the threshold. Mathematically, let $J^*(x, y)$ denote the optimal value function in state (x, y) . Then, it is easy to show that the optimal policy accepts an offer in state (x, y) if and only if

$$y \geq \gamma \sum_{x', y'} p(x' \mid x)q(y \mid x)J^*((x', y')),$$

To accommodate cases where set of possible offers is discrete, while still using smooth policies, we consider a parameterized class of stochastic policies, $\pi_\theta : \mathcal{S} \rightarrow [0, 1]$ that map a state to a probability of accepting the offer. Capturing our knowledge of the structure of the optimal policy⁶, one natural parameterization defines $\theta = (\theta_0^x, \theta_1^x)_{x \in \mathcal{X}}$ to be a vector of length $d := 2|\mathcal{X}|$ and sets

$$\pi_\theta(x, y) = \frac{1}{1 + e^{-(\theta_0^x + \theta_1^x y)}}$$

which can be interpreted as a ‘soft’ threshold on y for each context x . The closure of this policy class contains the optimal policy. Let c_x^* be the optimal threshold for some state (x, \cdot) . In the limit where $\theta_1^x \rightarrow \infty$ and $-(\theta_0^x/\theta_1^x) \rightarrow c_x^*$, we have that $\pi_\theta(x, \cdot)$ converges to the optimal policy with a deterministic threshold c_x^* .

Our next goal is to verify all the assumptions stated in Section for this problem. First note that for $a \in [0, 1]$ denoting the probability of accepting the offer and any policy $\pi_\theta \in \Pi$, we have

$$Q_\theta((x, y), a) = ay + (1 - a)\gamma \sum_{(x', y') \in \mathcal{S}} p(x' \mid x)q_x(y')J_\theta((x', y')) \quad (8)$$

which is linear in a . Next, we show that the class of threshold policies is closed under policy improvement. For any state $s = (x, y)$, consider the policy iteration update

$$\pi_+(s) \in \arg \max_{a \in [0, 1]} Q_\theta(s, a). \quad (9)$$

Clearly, $\pi^+(s) = 1$ if and only if the offer y exceeds the continuation value, $c_\theta(x)$ defined as

$$c_\theta(x) := \gamma \sum_{(x', y') \in \mathcal{S}} p(x' \mid x)q_x(y')J_\pi((x', y')).$$

We now come to Assumption 3 related to convexity of the policy class. Unfortunately, this does not hold for threshold policies as can be seen from a counterexample below. Consider policies $\pi_\theta, \pi_{\theta'} \in \Pi$ such that

$$\pi_\theta(x, y) = \begin{cases} 1 & \text{if } y \geq c_\theta(x), \\ 0 & \text{otherwise} \end{cases} \quad \pi_{\theta'}(x, y) = \begin{cases} 1 & \text{if } y \geq c_{\theta'}(x), \\ 0 & \text{otherwise} \end{cases}$$

Assuming without loss of generality that $c_{\theta'}(x) > c_\theta(x)$, a convex combination of π_θ and $\pi_{\theta'}$

$$\alpha\pi_\theta(x, y) + (1 - \alpha)\pi_{\theta'}(x, y) = \begin{cases} 1 & \text{if } y \geq c_{\theta'}(x) > c_\theta(x), \\ \alpha & \text{if } c_\theta(x) < y < c_{\theta'}(x), \\ 0 & \text{if } y \leq c_\theta(x) < c_{\theta'}(x) \end{cases}$$

⁶That is, the threshold only depends on the context.

does not lie in the class of threshold policies. This is the main challenge in proving a result analogous to Theorem 1 as for any sub-optimal policy π_θ , moving along a line segment towards the policy iteration update is not a feasible direction. But note that Assumption 3 is somewhat stronger than what we need to prove Theorem 1. Essentially, all we need is a feasible direction u in the parameter space to be able to move policy π_θ close to π_+ in the sense that for small α , $|\pi_{\theta+\alpha u}(s) - \pi_+(s)| < |\pi_\theta(s) - \pi_+(s)|$ for all $s \in \mathcal{S}$. For this case, it is simple to find u . We can then show the following result.

Lemma 3. *For the optimal stopping problem formulated in Example 3, $\nabla \ell(\theta) \neq 0$ for any $\theta \in \mathbb{R}^d$.*

Proof. Let π_+ denote the policy in (9). Then $\pi_+((x, y)) = 1$ if and only if $y \geq c_\theta(x)$. We construct a direction $u \in \mathbb{R}^d$ such that moving in the direction of u brings $\pi_\theta(s)$ closer to the action $\pi_+(s)$ in every state s . Take $u = (u_0^x, u_1^x)_{x \in \mathcal{X}} \in \mathbb{R}^d$ where $u_1^x = 1$ and $u_0^x = -c_\theta(x)$. Set $\theta^\alpha = \theta + \alpha u$. Then, for $f(z) = \frac{1}{1+e^{-z}}$, denoting the logistic function, we have

$$\left. \frac{d}{d\alpha} \pi_{\theta+\alpha u}((x, y)) \right|_{\alpha=0} = f'(\theta_0^x + \theta_1^x y) [-c_\theta(x) + y], \quad (10)$$

which is strictly positive for $y > c_\theta(x)$ and strictly negative for $y < c_\theta(x)$ as $f(z)$ is strictly increasing function for $z \in \mathbb{R}$. Using the policy gradient theorem, we get

$$\begin{aligned} \frac{d}{d\alpha} \ell(\theta + \alpha u) &= \sum_{s \in \mathcal{S}} \eta_\theta(s) \left[Q_\theta(s, 1) \left. \frac{d}{d\alpha} \pi_{\theta+\alpha u}(s) \right|_{\alpha=0} + Q_\theta(s, 0) \left. \frac{d}{d\alpha} (1 - \pi_{\theta+\alpha u}(s)) \right|_{\alpha=0} \right] \\ &= \sum_{s \in \mathcal{S}} \eta_\theta(s) [Q_\theta(s, 1) - Q_\theta(s, 0)] \left. \frac{d}{d\alpha} \pi_{\theta+\alpha u}(s) \right|_{\alpha=0} \\ &= \sum_{(x, y)} \eta_\theta((x, y)) [y - c_\theta(x)]^2 f'(\theta_0^x + \theta_1^x y) > 0. \end{aligned}$$

we used the fact that $Q_\theta(s, 1) - Q_\theta(s, 0) = [y - c_\theta(x)]$ from Equation (8). \square

E.3 Base-stock policies in finite horizon inventory control

In this section we give a detailed exposition of the multi-period newsvendor problem with backlogged demands as described in Example 4 along with the proof of Lemma 4. Let $s_t \in \mathbb{R}$ denote the state of the seller's inventory, a_t be the quantity of inventory ordered and $w_t \geq 0$ be the random demand at time t . Note that only positive inventory orders are allowed. So the action space for any state $s \in \mathcal{S}$ is defined as $\mathcal{A}(s) = \{a \in \mathbb{R} ; a \geq 0\}$. The seller's inventory evolves as,

$$s_{t+1} = s_t + a_t - w_t \quad t = 0, \dots, H-1,$$

for a problem with horizon H . For simplicity, we assume the demands w_t to be independent random variables that can take continuous values within a bounded set. Negative values of s_t correspond to backlogged demand that is filled when additional inventory becomes available. Let $c, b, p > 0$ be the per unit costs of ordering, holding and backlogging items respectively. We assume that $p > c$ to avoid degeneracies. If $p < c$, then it would never be optimal to order in the last period and possibly in other periods. Let us define a convex function, $r(\cdot) : \mathbb{R} \rightarrow \mathbb{R}$ as,

$$r(x) = p \max(0, -x) + b \max(0, x), \quad (11)$$

to represent the backlogging/holding cost in any state x . We consider the class of time-inhomogenous base-stock policies parameterized by $\theta = (\theta_0, \dots, \theta_{H-1})$ for each $\theta_t > 0$. At time t such a policy orders inventory $\pi_{\theta_t}^t(s_t) = \max\{0, \theta_t - s_t\}$. For $s \in \mathbb{R}$, $h \in \{1, \dots, H\}$ and $\theta \in \mathbb{R}^H$, define

$$J_\theta((h, s)) = \mathbb{E}_{\pi_\theta} \left[\sum_{t=h}^{H-1} (ca_t + r(s_t + a_t - w_t)) \mid s_h = s \right].$$

The average cost of policy π_θ for any initial distribution, ρ can be defined as,

$$\ell(\theta) = \mathbb{E}_{(h,s) \sim \rho} [J_\theta((h, s))]$$

Our ultimate goal is to prove Lemma 4. For this, we first verify all the assumptions for the finite horizon setting as stated in Section C. First we note that for this finite horizon inventory control problem, it is well known that base-stock-policies are optimal [5, 45] thereby verifying Assumption 6. Next, we verify the convexity of policy improvement steps by showing that $Q^*(t, s, a)$ is convex for all $s \in \mathbb{R}, a \in \mathbb{R}_+$ and $t \in \{0, \dots, H-1\}$. For this, we note that

$$Q^*(t, s, a) = ca + \mathbb{E}_{w_h} [r(s + a - w_t) + J^*(t + 1, s + a)].$$

where we assume that $J^*(h, x)$ is the optimal cost-to-go function for period h and state $x \in \mathcal{S}$ and that $J^*(H, x) = 0$. Note that $r(s + a - w_h)$ is a convex in a for any s and fixed w_t and expectations preserve convexity. A simple inductive argument [5] shows that $J^*(h, x)$ is convex in x . Combining these with simple properties of convex functions, it is easy to see that $Q^*(t, s, a)$ is convex in a .

The main challenge in the proof arises from the fact that because a policy π_θ only orders in some states, local changes in the base-stock levels θ only changes the actions in those states. This property technically breaks the convexity of the policy class as stated in Assumption 8. However, it is easy to construct a descent direction for suboptimal policies. Let $\theta_* = \{\theta_*^1, \dots, \theta_*^H\}$ be the thresholds corresponding to the optimal policy, which are unique by the strict convexity of $Q^*(t, s, \cdot)$. We show in the following lemma that for any policy parameterized by $\theta \in \mathbb{R}^H$ such that $\ell(\theta) > \ell(\theta_*)$, there exists a $u \in \mathbb{R}^H$ such that $\left. \frac{d}{d\alpha} \ell(\theta + \alpha u) \right|_{\alpha=0} < 0$.

Lemma 4. Consider Example 4 and define for $s \in \mathbb{R}, h \in \{1, \dots, H\}$, and $\theta \in \mathbb{R}^H$,

$$J_\theta((h, s)) = \mathbb{E}_{\pi_\theta} \left[\mathbb{E} \sum_{t=h}^{H-1} (ca_t + b \max\{s_t + a_t - w_t, 0\} + p \max\{-s_t + a_t - w_t, 0\}) \mid s_h = s \right].$$

Let ρ be an initial distribution supported on $\{1, \dots, H\} \times \mathbb{R}$ such that all moments of the marginal distribution $\rho(h, \cdot)$ exist for each h . Set $\ell(\theta) = \mathbb{E}_{(h,s) \sim \rho} [J_\theta((h, s))]$. Then for $\theta \in \mathbb{R}^d$, $\nabla \ell(\theta) = 0$ if and only if $\ell(\theta) = \min_{\theta' \in \mathbb{R}^H} \ell(\theta')$.

Proof. Let t be the last period in which the threshold is not optimal: $\theta^t \neq \theta_*^t$. If no such t exists, then we are already at the optimal policy. Else, we have that $Q_\theta(t, s, a) = Q^*(t, s, a)$ for every $s \in \mathcal{S}$ and $a \in \mathcal{A}(s)$. This follows by using the fact that we have the optimal thresholds in the subsequent periods $h \in \{t + 1, \dots, H\}$. Note that exactly the same argument was used in the proof of Theorem 2. Take $u^t = \theta_*^t - \theta^t$ and let $u = (u^1, \dots, u^H)$ where $u^i = 0$ for all $i \neq t$. Then,

$$\left. \frac{d}{d\alpha} \ell(\theta + \alpha u) \right|_{\alpha=0} = \sum_{s \in \mathcal{S}} \eta_\theta(t, s) \left. \frac{d}{d\alpha} Q^*(t, s, \pi_{\theta^t + \alpha u^t}^t(s)) \right|_{\alpha=0}$$

Now,

$$\sum_{s \in \mathcal{S}} \eta(t, s) \left. \frac{d}{d\alpha} Q^*(t, s, \pi_{\theta^t + \alpha u^t}^t(s)) \right|_{\alpha=0} = \sum_{s \in \mathcal{S}} \eta_\theta(t, s) \left(\left. \frac{\partial}{\partial a} Q^*(t, s, a) \right|_{a=\pi_{\theta^t}^t(s)} \right) \left(\left. \frac{d}{d\alpha} \pi_{\theta^t + \alpha u^t}^t(s) \right|_{\alpha=0} \right)$$

Note that,

$$\left. \frac{d}{d\alpha} \pi_{\theta^t + \alpha u^t}^t(s) \right|_{\alpha=0} = \frac{d}{d\alpha} \max \left(0, \theta^t + \alpha(\theta_*^t - \theta^t) - s \right) \Big|_{\alpha=0} = \begin{cases} \theta_*^t - \theta^t & \text{if } s < \theta^t, \\ 0 & \text{if } s > \theta^t \end{cases}$$

This implies,

$$\sum_{s \in \mathcal{S}} \eta(t, s) \frac{d}{d\alpha} Q^*(t, s, \pi_{\theta^t + \alpha u^t}(s)) \Big|_{\alpha=0} = \sum_{s < \theta^t} \eta_{\theta}(t, s) \left(\frac{\partial}{\partial a} Q^*(t, s, a) \Big|_{a=\pi_{\theta^t}^t(s)} \right) (\theta_*^t - \theta^t)$$

Note that $Q^*(t, s, a)$ is a convex function of a in one dimension with a minimum at $a = \pi_*^t(s)$. Thus, if $\theta_*^t < \theta^t$, then we are ordering more than we should, in which case $\frac{\partial}{\partial a} Q^*(t, s, a) \Big|_{a=\pi_{\theta^t}^t(s)} > 0$. On the other hand, if $\theta_*^t > \theta^t$, then we are ordering less than we should, in which case $\frac{\partial}{\partial a} Q^*(t, s, a) \Big|_{a=\pi_{\theta^t}^t(s)} < 0$. In either case, we have

$$\frac{d}{d\alpha} \ell(\theta + \alpha u) \Big|_{\alpha=0} = \sum_{s \in \mathcal{S}} \eta_{\theta}(t, s) \frac{d}{d\alpha} Q^*(t, s, \pi_{\theta^t + \alpha u^t}(s)) \Big|_{\alpha=0} < 0$$

which is a descent direction as desired. \square

F Details regarding numerical experiments

In this section, we provide implementation details about the numerical results shown in Figure 1 for three classic Dynamic Programming problems, as described in Sections 4 and Appendix E.

F.1 Infinite horizon problems:

The tabular MDP and the contextual optimal stopping problem are formulated as infinite horizon discounted problems with finite state and action space. For both, we use a discount factor of $\gamma = 0.9$ and the starting state was sampled uniformly over the finite state space. To compute the gradients exactly, we use the classic Policy Gradient theorem Sutton et al. [57] as stated below.

Theorem 4 (Policy Gradient Sutton et al. [57]). *For infinite horizon discounted MDPs, the policy gradient is given by:*

$$\nabla_{\theta} \ell(\theta) = \sum_{s \in \mathcal{S}} \rho(s) \eta_{\theta}(s) \sum_{a \in \mathcal{A}} Q_{\pi_{\theta}}(s, a) \nabla_{\theta} \pi_{\theta}(s, a)$$

for policy π_{θ} , where $\rho(s)$ denotes the starting state distribution and $\eta_{\theta}(s)$ is the discounted state occupancy measure as defined in Section 2.

For each case, the exact Q-values for all state-action pairs for a given policy, π , can be computed by solving the linear system of Bellman equations:

$$Q_{\pi}(s, a) = g(s, a) + \gamma \sum_{s'} P(s'|s, a) \sum_{a' \in \mathcal{A}} \pi(s', a') Q_{\pi}(s', a')$$

In matrix format, the system becomes

$$Q_{\pi} = g + \gamma \mathbf{P} \mathbf{\Pi}_{\pi} Q_{\pi}, \tag{12}$$

which can be solved either analytically or iteratively to obtain the exact Q-values. Here g is the vector of rewards, $\mathbf{P} \in \mathbb{R}^{|\mathcal{S}| \times |\mathcal{A}| \times |\mathcal{A}|}$ is the transition matrix of the process with $\mathbf{P}((s, a), s') = P(s'|s, a)$ and $\mathbf{\Pi}_{\pi} \in \mathbb{R}^{|\mathcal{S}| \times |\mathcal{S}| \times |\mathcal{A}|}$ describes the policy π as $\mathbf{\Pi}_{\pi}(s, (s, a)) = \pi(s, a)$. We use backtracking line search [12], as shown in the pseudo code below, to adaptively choose step-sizes for gradient steps at any iteration: $\theta_{k+1} = \theta_k + \alpha_k \nabla_{\theta} \ell(\theta_k) \quad \forall k \in \mathbb{N}$.

Although, typical implementations set α to be some constant (say 1), we found that scaling α by the inverse of the gradient norm in our implementation was immensely helpful in rapid convergence.

Algorithm 2: Backtracking line search

Input : $\alpha = 1/\|\nabla_{\theta}\ell(\theta_k)\|_2$, $\beta : 0 < \beta < 1$, current parameter value θ_k , gradient $\nabla_{\theta}\ell(\theta_k)$, cost function $\ell(\cdot)$.

Initialize: step-size t as $t = \alpha$.

while $\ell(\theta_k - t\nabla_{\theta}\ell(\theta_k)) > \ell(\theta_k) - \frac{t}{2}\|\nabla_{\theta}\ell(\theta_k)\|_2^2$ **do**

 | Update $t = t\beta$

end

return Step-size t .

This was because in our implementations, we found the gradient norms to be small and searching over large step-sizes allowed greater progress per iteration. As mentioned in Section 1.2, we believe that this numerical issue of vanishing gradients is an artifact of the local geometry of the parameter space and principled approaches, like natural gradient approaches, might alleviate this phenomenon. As this is not the primary focus of our work, we sidestep this issue by scaling α .

The optimal policy, π^* was computed using Policy Iteration algorithm for which the Q-values were computed using Equation (12). We plot the optimality gap: $\ell(\pi_{\theta_k}) - \ell(\pi^*) \quad \forall k \in \mathbb{N}$ in Figure 1. Below, we give specific details of numerical experiments run for the tabular MDP and the contextual optimal stopping problem.

- **Tabular MDP:** For simulation, we take the number of states and actions per state to be 100 and 20 respectively. The probability transition matrix, $P(s'|s, a) \in \mathbb{R}^{|\mathcal{S}||\mathcal{A}| \times |\mathcal{S}|}$ and reward vector, $g(s, a) \in \mathbb{R}^{|\mathcal{S}||\mathcal{A}|}$ were drawn uniformly at random. Recall that for a softmax policy, π^{θ} , parameterized by $\theta \in \mathbb{R}^{|\mathcal{S}||\mathcal{A}|}$, we have the expression for probability of taking action $i \in \mathcal{A}$ in state $s \in \mathcal{S}$ along with the expression for partial derivative as:

$$\pi_{\theta}(s, i) = \frac{e^{\theta_{si}}}{\sum_{k \in \mathcal{A}} e^{\theta_{sk}}}; \quad \frac{d}{d\theta_{s', j}} \pi_{\theta}(s, i) = \begin{cases} 0 & \text{if } s' \neq s \\ \pi_{\theta}(s, i)[1 - \pi_{\theta}(s, i)] & \text{if } s' = s, i = j \\ -\pi_{\theta}(s, i)\pi_{\theta}(s, j) & \text{if } s' = s, i \neq j \end{cases} .$$

This enables a straightforward computation of $\nabla_{\theta}\pi_{\theta}(s, a)$ which is used in the policy gradient theorem as shown above.

- **Contextual Optimal Stopping problem:** For our numerical experiments, we take the number of contexts/latent variables, $|\mathcal{X}| = 10$ and the number of offers, $|\mathcal{Y}| = 50$, with the associated offer value sampled from a uniform distribution. The transition probability matrix for context vectors, $p(x'|x) \in \mathbb{R}^{|\mathcal{X}| \times |\mathcal{X}|}$ as well as the emission probability matrix, $q(y|x) \in \mathbb{R}^{|\mathcal{X}| \times |\mathcal{Y}|}$ were randomly generated. Recall that the policy in this case is parameterized by $\theta = (\theta_0^x, \theta_1^x)_{x \in \mathcal{X}}$ with the probability of accepting offer y in state x along with the expression for partial derivative as

$$\pi_{\theta}(x, y) = \frac{1}{1 + e^{-(\theta_0^x + \theta_1^x y)}}; \quad \frac{d}{d\theta_0^x} \pi_{\theta}(x, y) = \pi_{\theta}(x, y)[1 - \pi_{\theta}(x, y)]; \quad \frac{d}{d\theta_1^x} \pi_{\theta}(x, y) = y(\pi_{\theta}(x, y)[1 - \pi_{\theta}(x, y)])$$

which is used in conjunction with the policy gradient theorem above.

F.2 Finite horizon undiscounted MDP: multi-period newsvendor problem

As that following the problem formulation in Section E.3, the average cost function can be written as:

$$\ell(\theta) = \sum_{h \in \{1, \dots, H\}} \sum_{s \in \mathcal{S}} J_{\theta}(h, s) \rho(h, s)$$

where ρ is the initial distribution for the starting state and the horizon and is supported on $\{1, \dots, H\} \times \mathbb{R}$. One way of approximating the gradients is to use the policy gradient theorem [53, 57] which requires approximating the state-action Q-values, $Q_{\pi_\theta}(s, a)$. Instead, we opt for using differential dynamic programming to compute the gradients as follows. Let w_1, \dots, w_H be i.i.d continuous random variables denoting the demand. We use the notation $(x)^+$ to denote $\max(0, x)$. Then, we have $s_{h+1} = f(s_h, a_h, w_h) = s_h + a_h - w_h$, where the inventory order, $a_h = \pi_{\theta_h}^h(s_h) = (\theta_h - s_h)^+$ and θ^h are the thresholds in period h . We use a Monte Carlo approach by writing the cost function as the expectation: $\ell(\theta) = \mathbb{E}_{w, i, s}[\ell(\theta, w, i, s)]$. Starting from an initial state s , time period i , and a given realization of random demands

$$\ell(\theta, w, i, s) = \sum_{h=i}^{H+1} g(s_h, a_h) - r(s_i)$$

where $g(s_h, a_h) = ca_h + r(s_h)$, $a_{H+1} = 0$ and $r(\cdot)$ represents the holding and backlog cost as defined in Equation (11). We can easily compute $\nabla \ell(\theta) = \mathbb{E}_{w, i, s}[\nabla \ell(\theta, w, i, s)]$. The key then is to use that, almost surely, future states and actions are differentiable functions of θ_h that can be computed through the chain rule of calculus. For our implementation, we take the start distribution to only put weight on the first period ($h=1$) and not randomize over the horizon. In this special case, the cost function becomes $\ell(\theta) = \mathbb{E}[\ell(\theta, w, s)]$ where

$$\ell(\theta, w, s) = \left[\sum_{h=1}^H g(s_h, a_h) \middle| s_1 = s \right]. \quad (13)$$

and we compute the gradients as shown below⁷.

$$\frac{\partial}{\partial \theta_i} \ell(\theta, w, s) = \sum_{h=i}^H \frac{\partial}{\partial \theta_i} g(s_h, a_h) = \frac{\partial a_i}{\partial \theta_i} \left[\frac{\partial}{\partial a_i} g(s_i, a_i) + \sum_{h=i+1}^H \left(\frac{\partial}{\partial s_h} g(s_h, a_h) \cdot \frac{\partial s_h}{\partial a_i} + \frac{\partial}{\partial a_h} g(s_h, a_h) \cdot \frac{\partial a_h}{\partial s_h} \cdot \frac{\partial s_h}{\partial a_i} \right) \right].$$

For inventory control with the class of threshold policies, $a_i = (\theta_i - s_i)^+$ we have

$$\frac{\partial a_i}{\partial \theta_i} = \begin{cases} 0 & \text{if } s_i > \theta_i \\ 1 & \text{if } s_i < \theta_i \end{cases}$$

Now, let τ_i be the first time after i when inventory is ordered, and set $\tau_i = \infty$ if no inventory is ordered till the end of the horizon. For every period h after i but before inventory is ordered, increasing the order a_i by da_i increases the inventory on hand s_h by da_i . A change in a_i does not influence the inventory on-hand after the next restock. This implies,

$$\frac{\partial s_h}{\partial a_i} = \begin{cases} 0 & \text{if } h > \tau_i \\ 1 & \text{if } h \leq \tau_i. \end{cases}$$

Also note that given for threshold policies, $\frac{\partial a_h}{\partial s_h} = -1$ if inventory is ordered in period h and 0 otherwise. Putting this all together, we have that

$$\frac{\partial}{\partial \theta_i} \ell(\theta, w) = \underbrace{\frac{\partial}{\partial a_i} g(s_i, a_i)}_{\text{change in ordering cost}} + \underbrace{\sum_{h=i+1}^{\tau_i \wedge H} \frac{\partial}{\partial s_h} g(s_h, a_h)}_{\text{change in holding cost}} - \underbrace{\mathbf{1}(\tau_i < \infty) \frac{\partial}{\partial a_{\tau_i}} g(s_{\tau_i}, a_{\tau_i})}_{\text{change in future ordering cost}},$$

⁷Note that we did not subtract the term $r(s_1)$ from the sum in the above definition. As s_1 is the initial state and does not depend on the thresholds, this term does not contribute to the gradient computation anyways.

when $s_i < \theta_i$ and $\frac{\partial}{\partial \theta_i} \ell(\theta, w) = 0$ if $s_i > \theta_i$. Recall that we have $g(a, s) = ca + r(s)$, where $r(s) = p \max(0, -s) + b \max(0, s)$. This implies that $\frac{\partial}{\partial a} g(ua, s) = c$ and $r'(s) = \frac{\partial}{\partial s} [r(s) + ca] = b \cdot \mathbf{1}(s > 0) - p \cdot \mathbf{1}(s < 0)$. We find

$$\frac{\partial}{\partial \theta_i} \ell(\theta, w) = \begin{cases} 0, & \theta_i < s_i \\ \sum_{h=i+1}^{\tau_i} r'(s_h), & \theta_i > s_i, \tau_i < \infty \\ c + \sum_{h=i+1}^H r'(s_h), & \theta_i > s_i, \tau_i = \infty \end{cases}$$

An efficient way to compute $\nabla \ell(\theta, w, s)$ is to forward simulate the states and orders throughout the episode and approximate $\frac{\partial}{\partial \theta_i} \ell(\theta) \approx \frac{1}{n} \frac{1}{m} \sum_{l=1}^n \sum_{k=1}^m \frac{\partial}{\partial \theta_i} \ell(\theta, w^k, s^l)$ where the w^k denote i.i.d vectors drawn from the demand distribution and s^l are drawn i.i.d from the start state distribution $\rho(s)$.

To track the optimality gap, we need to compute the optimal thresholds, θ_* . Given θ_* , it is easy to estimate $\ell(\theta_*)$ using Monte Carlo simulations as shown in Equation (13). To get the optimal thresholds, we essentially do approximate dynamic programming as explained briefly below. For a detailed exposition, we refer the readers to Section 3.2 of [5].

Let $J^*(h, s)$ and $J(h, s)$ denote the (optimal) cost-to-go functions from period h and state s . Suppose we have the optimal thresholds from period h to H as $\{\theta_*^h, \dots, \theta_*^H\}$, then we can find θ_*^h by essentially minimizing a convex function in one dimension. Define

$$G_h(\theta) = \mathbb{E}_{s_h} \left[g(s_h, \pi_\theta^h(s_h)) + \mathbb{E}_{w_h} [J^*(h+1)(s_h + \pi_\theta^h(s_h) - w_h)] \right]$$

which can be easily estimated using Monte-Carlo approximation for the expectations. It is easy to show that $G_h(\theta)$ is convex in θ . We use golden-section search [29], a zeroth order technique to approximately optimize $G_h(\theta)$ to get θ_*^h .