Multiple-Step Greedy Policies in Online and Approximate Reinforcement Learning

Yonathan Efroni*

jonathan.efroni@gmail.com

Gal Dalal* gald@campus.technion.ac.il

Bruno Scherrer† bruno.scherrer@inria.fr Shie Mannor* shie@ee.technion.ac.il

Abstract

Multiple-step lookahead policies have demonstrated high empirical competence in Reinforcement Learning, via the use of Monte Carlo Tree Search or Model Predictive Control. In a recent work [5], multiple-step greedy policies and their use in vanilla Policy Iteration algorithms were proposed and analyzed. In this work, we study multiple-step greedy algorithms in more practical setups. We begin by highlighting a counter-intuitive difficulty, arising with soft-policy updates: even in the absence of approximations, and contrary to the 1-step-greedy case, monotonic policy improvement is not guaranteed unless the update stepsize is sufficiently large. Taking particular care about this difficulty, we formulate and analyze online and approximate algorithms that use such a multi-step greedy operator.

1 Introduction

The use of the 1-step policy improvement in Reinforcement Learning (RL) was theoretically investigated under several frameworks, e.g., Policy Iteration (PI) [18], approximate PI [2, 9, 13], and Actor-Critic [10]; its practical uses are abundant [22, 12, 25]. However, single-step based improvement is not necessarily the optimal choice. It was, in fact, empirically demonstrated that multiple-step greedy policies can perform conspicuously better. Notable examples arise from the integration of RL and Monte Carlo Tree Search [4, 28, 23, 3, 25, 24] or Model Predictive Control [15, 6, 27].

Recent work [5] provided guarantees on the performance of the multiple-step greedy policy and generalizations of it in PI. Here, we establish it in the two practical contexts of online and approximate PI. With this objective in mind, we begin by highlighting a specific difficulty: *softly updating* a policy with respect to (w.r.t.) a multiple-step greedy policy does not necessarily result in improvement of the policy (Section 4). We find this property intriguing since monotonic improvement is guaranteed in the case of soft updates w.r.t. the 1-step greedy policy, and is central to the analysis of many RL algorithms [10, 9, 22]. We thus engineer several algorithms to circumvent this difficulty and provide some non-trivial performance guarantees, that support the interest of using multi-step greedy operators. These algorithms assume access to a generative model (Section 5) or to an approximate multiple-step greedy policy (Section 6).

2 Preliminaries

Our framework is the infinite-horizon discounted Markov Decision Process (MDP). An MDP is defined as the 5-tuple (S, A, P, R, γ) [18], where S is a finite state space, A is a finite action space,

^{*}Department of Electrical Engineering, Technion, Israel Institute of Technology

[†]INRIA, Villers les Nancy, France

 $P \equiv P(s'|s,a)$ is a transition kernel, $R \equiv r(s,a)$ is a reward function, and $\gamma \in (0,1)$ is a discount factor. Let $\pi: \mathcal{S} \to \mathcal{P}(\mathcal{A})$ be a stationary policy, where $\mathcal{P}(\mathcal{A})$ is a probability distribution on \mathcal{A} . Let $v^{\pi} \in \mathbb{R}^{|\mathcal{S}|}$ be the value of a policy π , defined in state s as $v^{\pi}(s) \equiv \mathbb{E}^{\pi}[\sum_{t=0}^{\infty} \gamma^{t} r(s_{t}, \pi(s_{t})) | s_{0} = s]$. For brevity, we respectively denote the reward and value at time t by $r_{t} \equiv r(s_{t}, \pi_{t}(s_{t}))$ and $v_{t} \equiv v(s_{t})$. It is known that $v^{\pi} = \sum_{t=0}^{\infty} \gamma^{t} (P^{\pi})^{t} r^{\pi} = (I - \gamma P^{\pi})^{-1} r^{\pi}$, with the component-wise values $[P^{\pi}]_{s,s'} \triangleq P(s' \mid s,\pi(s))$ and $[r^{\pi}]_{s} \triangleq r(s,\pi(s))$. Lastly, let

$$q^{\pi}(s,a) = \mathbb{E}^{\pi}\left[\sum_{t=0}^{\infty} \gamma^{t} r(s_{t}, \pi(s_{t})) \mid s_{0} = s, a_{0} = a\right]. \tag{1}$$

Our goal is to find a policy π^* yielding the optimal value v^* such that

$$v^* = \max(I - \gamma P^{\pi})^{-1} r^{\pi} = (I - \gamma P^{\pi^*})^{-1} r^{\pi^*}.$$
 (2)

This goal can be achieved using the three classical operators (equalities hold component-wise):

$$\forall v, \pi, \ T^{\pi}v = r^{\pi} + \gamma P^{\pi}v,$$

$$\forall v, \ Tv = \max_{\pi} T^{\pi}v,$$

$$\forall v, \ \mathcal{G}(v) = \{\pi : T^{\pi}v = Tv\},$$

where T^{π} is a linear operator, T is the optimal Bellman operator and both T^{π} and T are γ -contraction mappings w.r.t. the max norm. It is known that the unique fixed points of T^{π} and T are v^{π} and v^* , respectively. The set $\mathcal{G}(v)$ is the standard set of 1-step greedy policies w.r.t. v.

3 The h- and κ -Greedy Policies

In this section, we bring forward necessary definitions and results on two classes of multiple-step greedy policies: h- and κ -greedy [5]. Let $h \in \mathbb{N} \setminus \{0\}$. The h-greedy policy π_h outputs the first optimal action out of the sequence of actions solving a non-stationary, h-horizon control problem as follows:

$$\forall s \in \mathcal{S}, \ \pi_h(s) \in \arg\max_{\pi_0} \max_{\pi_1, \dots, \pi_{h-1}} \mathbb{E}^{\pi_0 \dots \pi_{h-1}} \left[\sum_{t=0}^{h-1} \gamma^t r(s_t, \pi_t(s_t)) + \gamma^h v(s_h) \mid s_0 = s \right].$$

Since the h-greedy policy can be represented as the 1-step greedy policy w.r.t. $T^{h-1}v$, the set of h-greedy policies w.r.t. v, $\mathcal{G}_h(v)$, can be formally defined as follows:

$$\forall v, \pi, \ T_h^{\pi} v = T^{\pi} T^{h-1} v,$$

$$\forall v, \ \mathcal{G}_h(v) = \{\pi : T_h^{\pi} v = T^h v\}.$$

Let $\kappa \in [0,1]$. The set of κ -greedy policies w.r.t. a value function v, $\mathcal{G}_{\kappa}(v)$, is defined using the following operators:

$$\forall v, \pi, \ T_{\kappa}^{\pi}v = (I - \kappa \gamma P^{\pi})^{-1}(r^{\pi} + (1 - \kappa)\gamma P^{\pi}v)$$

$$\forall v, \ T_{\kappa}v = \max_{\pi} T_{\kappa}^{\pi}v = \max_{\pi} (I - \kappa \gamma P^{\pi})^{-1}(r^{\pi} + (1 - \kappa)\gamma P^{\pi}v)$$

$$\forall v, \ \mathcal{G}_{\kappa}(v) = \{\pi : T_{\kappa}^{\pi}v = T_{\kappa}v\}.$$
(3)

Remark 1. A comparison of (2) and (3) reveals that finding the κ -greedy policy is equivalent to solving a $\kappa\gamma$ -discounted MDP with shaped reward $r_{v,\kappa}^{\pi} \stackrel{def}{=} r^{\pi} + (1-\kappa)\gamma P^{\pi}v$.

In [5, Proposition 11], the κ -greedy policy was explained to be interpolating over all geometrically κ -weighted h-greedy policies. It was also shown that for $\kappa=0$, the 1-step greedy policy is restored, while for $\kappa=1$, the κ -greedy policy is the optimal policy.

Both T_{κ}^{π} and T_{κ} are ξ_{κ} contraction mappings, where $\xi_{\kappa} = \frac{\gamma(1-\kappa)}{1-\gamma\kappa} \in [0,\gamma]$. Their respective fixed points are v^{π} and v^{*} . For brevity, where there is no risk of confusion, we shall denote ξ_{κ} by ξ . Moreover, in [5] it was shown that both the h- and κ -greedy policies w.r.t. v^{π} are strictly better then π , unless $\pi = \pi^{*}$.




Figure 1: The Tightrope Walking MDP used in the counter example of Theorem 1.

Next, let

$$q_{\kappa}^{\pi}(s,a) = \max_{\pi'} \mathbb{E}^{\pi'} \left[\sum_{t=0}^{\infty} (\kappa \gamma)^{t} (r(s_{t}, \pi'(s_{t})) + \gamma(1-\kappa)v^{\pi}(s_{t+1}) \mid s_{0} = s, a_{0} = a \right].$$
 (4)

The latter is the optimal q-function of the surrogate, $\gamma \kappa$ -discounted MDP with v^{π} -shaped reward (see Remark 1). Thus, we can obtain a κ -greedy policy, $\pi_{\kappa} \in \mathcal{G}_{\kappa}(v^{\pi})$, directly from q_{κ}^{π} :

$$\pi_{\kappa}(s) \in \arg\max_{a} q_{\kappa}^{\pi}(s, a), \ \forall s \in \mathcal{S}.$$

See that the greedy policy w.r.t. $q_{\kappa=0}^{\pi}(s,a)$ is the 1-step greedy policy since $q_{\kappa=0}^{\pi}(s,a)=q^{\pi}(s,a)$.

4 Multi-step Policy Improvement and Soft Updates

In this section, we focus on policy improvement of multiple-step greedy policies, performed with soft updates. Soft updates of the 1-step greedy policy have proved necessary and beneficial in prominent algorithms [10, 9, 22]. Here, we begin by describing an intrinsic difficulty in selecting the step-size parameter $\alpha \in (0,1]$ when updating with multiple-step greedy policies. Specifically, denote by π' such multiple-step greedy policy w.r.t. v^{π} . Then, $\pi_{\text{new}} = (1 - \alpha)\pi + \alpha\pi'$ is not necessarily better than π

Theorem 1. For any MDP, let π be a policy and v^{π} its value. Let $\pi_{\kappa} \in \mathcal{G}_{\kappa}(v^{\pi})$ and $\pi_h \in \mathcal{G}_h(v^{\pi})$ with $\kappa \in [0, 1]$ and h > 1. Consider the mixture policies with $\alpha \in (0, 1]$,

$$\pi(\alpha, \kappa) \stackrel{def}{=} (1 - \alpha)\pi + \alpha \pi_{\kappa},$$
$$\pi(\alpha, h) \stackrel{def}{=} (1 - \alpha)\pi + \alpha \pi_{h}.$$

Then we have the following equivalences:

- 1. The inequality $v^{\pi(\alpha,\kappa)} \geq v^{\pi}$ holds for all MDPs if and only if $\alpha \in [\kappa, 1]$.
- 2. The inequality $v^{\pi(\alpha,h)} \geq v^{\pi}$ holds for all MDPs if and only if $\alpha = 1$.

The above inequalities hold entry-wise, with strict inequality in at least one entry unless $v^{\pi} = v^*$.

Proof sketch. See Appendix A for the full proof. Here, we only provide a counterexample demonstrating the potential non-monotonicity of $\pi(\alpha, \kappa)$ when the stepsize α is not big enough. One can show the same for $\pi(\alpha, h)$ with the same example.

Consider the Tightrope Walking MDP in Fig. 1. It describes the act of walking on a rope: in the initial state s_0 the agent approaches the rope, in s_1 the walking attempt occurs, s_2 is the goal state and s_3 is repeatedly met if the agent falls from the rope, resulting in negative reward.

First, notice that by definition, $\forall v, \ \pi^* \in \mathcal{G}_{\kappa=1}(v)$. We call this policy the "confident" policy. Obviously, for any discount factor $\gamma \in (0,1)$, $\pi^*(s_0) = a_1$ and $\pi^*(s_1) = a_1$. Instead, consider the "hesitant" policy $\pi_0(s) \equiv a_0 \ \forall s$. We now claim that for any $\alpha \in (0,1)$ and

$$c > \frac{\alpha}{1 - \alpha} \tag{5}$$

the mixture policy, $\pi(\alpha, \kappa=1)=(1-\alpha)\pi_0+\alpha\pi^*$, is not strictly better than π_0 . To see this, notice that $v^{\pi_0}(s_1)<0$ and $v^{\pi_0}(s_0)=0$; i.e., the agent accumulates zero reward if she does not climb the rope. Thus, while $v^{\pi_0}(s_0)=0$, taking any mixture of the confident and hesitant policies can result in $v^{\pi(\alpha,\kappa=1)}(s_0)<0$, due to the portion of the transition to s_1 and its negative contribution. Based on this construction, let $\kappa\in[0,1]$. To ensure $\pi^*\in\mathcal{G}_\kappa(v^\pi)$, we find it is necessary that

$$c \le \frac{\kappa}{1 - \kappa}.\tag{6}$$

To conclude, if both (5) and (6) are satisfied, the mixture policy does not improve over π_0 . Due to the monotonicity of $\frac{x}{1-x}$, such a choice of c is indeed possible for $\alpha < \kappa$. \square

Theorem 1 guarantees monotonic improvement for the 1-step greedy policy as a special case when $\kappa=0$. Hence, we get that for any $\alpha\in(0,1]$, the mixture of any policy π and the 1-step greedy policy w.r.t. v^{π} is monotonically better then π . To the best of our knowledge, this result was not explicitly stated anywhere. Instead, it appeared within proofs of several famous results, e.g, [10, Lemma 5.4], [9, Corollary 4.2], and [21, Theorem 1].

In the rest of the paper, we shall focus on the κ -greedy policy and extend it to the online and the approximate cases. The discovery that the κ -greedy policy w.r.t. v^{π} is not necessarily strictly better than π will guide us in appropriately devising algorithms.

5 Online κ -Policy Iteration with Cautious Soft Updates

In [5], it was shown that using the κ -greedy policy in the improvement stage leads to a convergent PI procedure – the κ -PI algorithm. This algorithm repeats i) finding the optimal policy of small-horizon surrogate MDP with shaped reward, and ii) calculating the value of the optimal policy and use it to shape the reward of next iteration. Here, we devise a practical version of κ -PI, which is model-free, online and runs in two timescales; i.e, it performs i) and ii) simultaneously.

The method is depicted in Algorithm 1. It is similar to the asynchronous PI analyzed in [16], except for two major differences. First, the fast timescale tracks both q^π, q^π_κ and not just q^π . Thus, it enables access to both the 1-step-greedy and κ -greedy policies. The 1-step greedy policy is attained via the q^π estimate, which is plugged into a q-learning [29] update rule for obtaining the κ -greedy policy. The latter essentially solves the surrogate $\kappa\gamma$ -discounted MDP (see Remark 1). The second difference is in the slow timescale, in which the policy is updated using a new operator, b_s , as defined below. To better understand this operator, first notice that in Stochastic Approximation methods such as Algorithm 1, the policy is improved using soft updates with decaying stepsizes. However, as Theorem 1 states, monotonic improvement is not guaranteed below a certain stepsize value. Hence, for $q, q_\kappa \in \mathbb{R}^{|\mathcal{S} \times \mathcal{A}|}$ and policy π , we set $b_s(q, q_\kappa, \pi)$ to be the κ -greedy policy only when assured to have improvement:

$$b_s(q, q_{\kappa}, \pi) = \begin{cases} a_{\kappa}(s) & \text{if } q(s, a_{\kappa}) \ge v^{\pi}(s), \\ a_{1\text{-step}}(s) & \text{else}, \end{cases}$$

where $a_{\kappa}(s) \stackrel{\text{def}}{=} \arg\max_{a} q_{\kappa}(s,a), \ a_{\text{1-step}}(s) \stackrel{\text{def}}{=} \arg\max_{a} q(s,a), \ \text{and} \ v^{\pi}(s) = \sum_{a} \pi(a \mid s) q(s,a).$

We respectively denote the state and state-action-pair visitation counters after the n-th time-step by $\nu_n(s) \stackrel{\text{def}}{=} \sum_{k=1}^n \mathbbm{1}_{s=s_k}$ and $\phi_n(s,a) \stackrel{\text{def}}{=} \sum_{k=1}^n \mathbbm{1}_{(s,a)=(s_k,a_k)}$. The stepsize sequences $\mu_f(\cdot), \mu_s(\cdot)$ satisfy the common assumption (B2) in [16], among which $\lim_{n\to\infty} \mu_s(n)/\mu_f(n) \to 0$. The second moments of $\{r_n\}$ are assumed to be bounded. Furthermore, let ν be some measure over the state space, s.t. $\forall s \in \mathcal{S}, \ \nu(s) > 0$. Then, we assume to have a generative model $\mathbb{G}(\nu, \pi)$, using which we sample state $s \sim \nu$, sample action $s \sim \pi(s)$, apply action $s \sim \pi(s)$ and receive reward $s \sim \pi(s)$ and next state $s \sim \pi(s)$.

The fast-timescale update rules in lines 6 and 8 can be jointly written as the sum of $H_{\kappa}^{\pi}(q, q_{\kappa})$ (defined below) and a martingale difference noise.

Algorithm 1 Two-Timescale Online κ -Policy-Iteration

```
1: initialize: \pi_0, q_0, q_{\kappa,0}.

2: for n = 0, ... do

3: s_n, a_n, r_n, s'_n \sim \mathbb{G}(\nu, \pi_n)

4: # Fast-timescale updates

5: \delta_n = r_n + \gamma v_n^{\pi}(s'_n) - q_n(s_n, a_n)

6: q_{n+1}(s_n, a_n) \leftarrow q_n(s_n, a_n) + \mu_f(\phi_{n+1}(s_n, a_n))\delta_n

7: \delta_{\kappa,n} = r_n + \gamma(1 - \kappa)v_n^{\pi}(s'_n) + \kappa\gamma \max_{a'} q_{\kappa,n}(s'_n, a') - q_{\kappa,n}(s_n, a_n)

8: q_{\kappa,n+1}(s_n, a_n) \leftarrow q_{\kappa,n}(s_n, a_n) + \mu_f(\phi_{n+1}(s_n, a_n))\delta_{\kappa,n}

9: # Slow-timescale updates

10: \pi_{n+1}(s_n) \leftarrow \pi_n(s_n) + \mu_s(\nu_{n+1}(s_n))(b_{s_n}(q_{n+1}, q_{\kappa,n+1}, \pi_n) - \pi_n(s_n))

11: end for

12: return: \pi
```

Definition 1. Let $q, q_{\kappa} \in \mathbb{R}^{|\mathcal{S}||\mathcal{A}|}$. The mapping $H_{\kappa}^{\pi} : \mathbb{R}^{2|\mathcal{S}||\mathcal{A}|} \to \mathbb{R}^{2|\mathcal{S}||\mathcal{A}|}$ is defined as follows $\forall (s, a) \in \mathcal{S} \times \mathcal{A}$.

$$H_{\kappa}^{\pi}(q,q_{\kappa})(s,a) \stackrel{\textit{def}}{=} \begin{bmatrix} r(s,a) + \gamma \mathbb{E}_{s',a^{\pi}} q(s',a^{\pi}) \\ r(s,a) + \gamma (1-\kappa) \mathbb{E}_{s',a^{\pi}} q(s',a^{\pi}) + \kappa \gamma \mathbb{E}_{s'} \max_{a'} q_{\kappa}(s',a') \end{bmatrix},$$

where $s' \sim P(\cdot \mid s, a), a^{\pi} \sim \pi(s')$.

The following lemma shows that, given a fixed π , H_{κ}^{π} is a contraction, equivalently to [16, Lemma 5.3] (see Appendix B for the proof).

Lemma 2. H_{κ}^{π} is a γ -contraction in the max-norm. Its fixed point is $[q^{\pi}, q_{\kappa}^{\pi}]^{\top}$, as defined in (1), (4).

Finally, based on several intermediate results given in Appendix C and relaying on Lemma 2, we establish the convergence of Algorithm 1.

Theorem 3. The coupled process $(q_n, q_{\kappa,n}, \pi_n)$ in Algorithm 1 converges to the limit (q^*, q^*, π^*) , where q^* is the optimal q-function and π^* is the optimal policy.

For $\kappa=1$, the fast-timescale update rule in line 8 corresponds to that of q-learning [29]. For that κ , Algorithm 1 uses an estimated optimal q-function to update the current policy when improvement is assured. For $\kappa<1$, the estimated κ -dependent optimal q-function (see (4)) is used, again with the 'cautious' policy update. Moreover, Algorithm 1 combines an off-policy algorithm, i.e., q-learning, with an on-policy Actor-Critic algorithm. To the best of our knowledge, this is the first appearance of these two approaches combined in a single algorithm.

6 Approximate κ -Policy Iteration with Hard Updates

Theorem 1 establishes the conditions required for guaranteed monotonic improvement of softly-updated multiple-step greedy policies. The algorithm in Section 5 then accounts for these conditions to ensure convergence. Contrarily, in this section, we derive and study algorithms that perform hard policy-updates. Specifically, we generalize the prominent Approximate Policy Iteration (API) [13, 7, 11] and Policy Search by Dynamic Programming (PSDP) [1, 19]. For both, we obtain performance guarantees that exhibit a tradeoff in the choice of κ , with optimal performance bound achieved with $\kappa > 0$. That is, our approximate κ -generalized PI methods outperform the 1-step greedy approximate PI methods in terms of best known guarantees.

For the algorithms here we assume an oracle that returns a κ -greedy policy with some error. Formally, we denote by $\mathcal{G}_{\kappa,\delta,\nu}(v)$ the set of approximate κ -greedy policies w.r.t. v, with δ approximation error under some measure ν .

Definition 2 (Approximate κ -greedy policy). Let $v: \mathcal{S} \to \mathbb{R}$ be a value function, $\delta \geq 0$ a real number and ν a distribution over \mathcal{S} . A policy $\pi \in \mathcal{G}_{\kappa,\delta,\nu}(v)$ if $\nu T_{\kappa}^{\pi} v \geq \nu T_{\kappa} v - \delta$.

Such a device can be implemented using existing approximate methods, e.g., Conservative Policy Iteration (CPI) [9], approximate PI or VI [7], Policy Search [21], or by having an access to an approximate model of the environment. The approximate κ -greedy oracle assumed here is less

restrictive than the one assumed in [5]. There, a uniform error over states was assumed, whereas here, the error is defined w.r.t. a specific measure, ν . For practical purposes, ν can be thought of as the initial sampling distribution to which the MDP is initialized. Lastly, notice that the larger κ is, the harder it is to solve the surrogate $\kappa\gamma$ -discounted MDP since the discount factor is bigger [17, 26, 8]; i.e., the computational cost of each call to the oracle increases.

Using the concept of *concentrability coefficients* introduced in [13] (there, they were originally termed "diffusion coefficients"), we follow the line of work in [13, 14, 7, 19, 11] to prove our performance bounds. This allows a direct comparison of the algorithms proposed here with previously studied approximate 1-step greedy algorithms. Namely, our bounds consist of concentrability coefficients $C^{(1)}$, $C^{(2)}$, $C^{(2,k)}$ and $C^{\pi^*(1)}$ from [19, 11], as well as two new coefficients $C^{\pi^*}_{\kappa}$ and $C^{\pi^*(1)}_{\kappa}$.

Definition 3 (Concentrability coefficients [19, 11])). Let μ, ν be some measures over S. Let $\{c(i)\}_{i=0}^{\infty}$ be the sequence of the smallest values in $[1, \infty) \cup \{\infty\}$ such that for every i, for all sequences of deterministic stationary policies $\pi_1, \pi_2, ..., \pi_i, \mu \prod_{j=1}^i P^{\pi_j} \leq c(i)\nu$. Let $C^{(1)}(\mu, \nu) = (1-\gamma)\sum_{i=0}^{\infty} \gamma^i c(i)$ and $C^{(2,k)}(\mu, \nu) = (1-\gamma)^2 \sum_{i,j=0}^{\infty} \gamma^{i+j} c(i+j+k)$. For brevity, we denote $C^{(2,0)}(\mu, \nu)$ as $C^{(2)}(\mu, \nu)$. Similarly, let $\{c^{\pi^*}(i)\}_{i=0}^{\infty}$ be the sequence of the smallest values in $[1,\infty) \cup \{\infty\}$ such that for every $i,\mu(P^{\pi^*})^i \leq c^{\pi^*}(i)\nu$. Let $C^{\pi^*(1)}(\mu, \nu) = (1-\gamma)\sum_{i=0}^{\infty} \gamma^i c^{\pi^*}(i)$.

We now introduce two new concentrability coefficients suitable for bounding the worst-case performance of PI algorithms with approximate κ -greedy policies.

Definition 4 (κ -Concentrability coefficients). Let $C_{\kappa}^{\pi^*(1)}(\mu,\nu) = \frac{\xi}{\gamma}C^{\pi^*(1)}(\mu,\nu) + (1-\xi)\kappa c(0)$. Also, let $C_{\kappa}^{\pi^*}(\mu,\nu) \in [1,\infty) \cup \{\infty\}$ be the smallest value s.t. $d_{\kappa,\mu}^{\pi^*} \leq C_{\kappa}^{\pi^*}(\mu,\nu)\nu$, where $d_{\kappa,\mu}^{\pi^*} = (1-\xi)\mu(I-\xi D_{\kappa}^{\pi^*}P^{\pi^*})^{-1}$ is a probability measure and $D_{\kappa}^{\pi} = (1-\kappa\gamma)(I-\kappa\gamma P^{\pi})^{-1}$ is a stochastic matrix.

In the definitions above, ν is the measure according to which the approximate improvement is guaranteed, while μ specifies the distribution on which one measures the loss $\mathbb{E}_{s\sim\mu}[v^*(s)-v^{\pi_k}(s)]=\mu(v^*-v^{\pi_k})$ that we wish to bound. From Definition 4 it holds that $C^{\pi^*}_{\kappa=0}(\mu,\nu)=C^{\pi^*}(\mu,\nu)$; the latter was previously defined in, e.g, [19, Definition 1].

Before giving our performance bounds, we first study the behavior of the coefficients appearing in them. The following lemma sheds light on the behavior of $C_{\kappa}^{\pi^*}(\mu,\nu)$. Specifically, it shows that under certain constructions, $C_{\kappa}^{\pi^*}(\mu,\nu)$ decreases³ as κ increases (see proof in Appendix D).

Lemma 4. Let $\nu(\alpha) = (1 - \alpha)\nu + \alpha\mu$. Then, for all $\kappa' > \kappa$, there exists $\alpha^* \in (0,1)$ such that $C_{\kappa'}^{\pi^*}(\mu,\nu(\alpha^*)) \leq C_{\kappa}^{\pi^*}(\mu,\nu)$. The inequality is strict for $C_{\kappa}^{\pi^*}(\mu,\nu) > 1$. For $\mu = \nu$ this implies that $C_{\kappa}^{\pi^*}(\nu,\nu)$ is a decreasing function of κ .

Definition 4 introduces two coefficients with which we shall derive our bounds. Though traditional arithmetic relations between them do not exist, they do comply to some notion of ordering.

Remark 2 (Order of concentrability coefficients). In [19], an order between the concentrability coefficients was introduced: a coefficient A is said to be strictly better than B-a relation we denote with $A \prec B-if$ and only if i) $B < \infty$ implies $A < \infty$ and ii) there exists an MDP for which $A < \infty$ and $B = \infty$. Particularly, it was argued that

$$C^{\pi^*}(\mu,\nu) \prec C^{\pi^*(1)}(\mu,\nu) \prec C^{(1)}(\mu,\nu) \prec C^{(2)}(\mu,\nu),$$
 and $C^{(2,k_1)}(\mu,\nu) \prec C^{(2,k_2)}(\mu,\nu)$ if $k_2 < k_1$.

In this sense, $C_{\kappa}^{\pi^*(1)}(\mu,\nu)$ is analogous to $C^{\pi^*(1)}(\mu,\nu)$, while its definition might suggest improvement as κ increases. Moreover, combined with the fact that $C_{\kappa}^{\pi^*}(\mu,\nu)$ improves as κ increases, as Lemma 4 suggests, $C_{\kappa}^{\pi^*}(\mu,\nu)$ is better than all previously defined concentrability coefficients.

6.1 κ -Approximate Policy Iteration

A natural generalization of API [13, 19, 11] to the multiple-step greedy policy is κ -API, as given in Algorithm 2. In each of its iterations, the policy is updated to the approximate κ -greedy policy w.r.t. $v^{\pi_{k-1}}$; i.e, a policy from the set $\mathcal{G}_{\kappa,\delta,\nu}(v^{\pi_{k-1}})$.

³A smaller coefficient is obviously better. The best value for any concentrability coefficient is 1.

Algorithm 2 κ -API

```
\begin{array}{l} \textbf{initialize} \quad \kappa \in [0,1], \nu, \delta, v^{\pi_0} \\ v \leftarrow v^{\pi_0} \\ \textbf{for} \quad k = 1, ... \textbf{do} \\ \quad \pi_k \leftarrow \mathcal{G}_{\kappa,\delta,\nu}(v) \\ \quad v \leftarrow v^{\pi_k} \\ \textbf{end for} \\ \textbf{return} \quad \pi \end{array}
```

Algorithm 3 κ -PSDP

```
 \begin{array}{l} \textbf{initialize} \quad \kappa \in [0,1], \nu, \delta, v^{\pi_0}, \Pi = [ \ ] \\ v \leftarrow v^{\pi_0} \\ \textbf{for} \quad k = 1, ... \ \textbf{do} \\ \pi_k \leftarrow \mathcal{G}_{\kappa, \delta, \nu}(v) \\ v \leftarrow T_\kappa^{\pi_k} v \\ \Pi \leftarrow & \texttt{Append}(\Pi, \pi_k) \\ \textbf{end for} \\ \textbf{return} \quad \Pi \end{array}
```

The following theorem gives a performance bound for κ -API (see proof in Appendix E), with

$$\begin{split} C_{\kappa-\mathrm{API}}(\mu,\nu) &= (1-\kappa)^2 C^{(2)}(\mu,\nu) + (1-\gamma)\kappa \left((1-\kappa)C^{(1)}(\mu,\nu) + (1-\gamma\kappa)C_{\kappa}^{\pi^*(1)}(\mu,\nu) \right), \\ C_{\kappa-\mathrm{API}}^{(k,1)}(\mu,\nu) &= (1-\kappa\gamma) \left(\kappa(1-\kappa\gamma)C_{\kappa}^{\pi^*}(\mu,\nu) + (1-\kappa)^2 C^{(1)}(\mu,\nu) \right), \\ C_{\kappa-\mathrm{API}}^{(k,2)}(\mu,\nu) &= (1-\kappa)\kappa \left((1-\gamma)C^{(1)}(\mu,\nu) + g(\kappa)(1-\kappa)\gamma^k C^{(2,k)}(\mu,\nu) \right), \end{split}$$

where $g(\kappa)$ is a bounded function for $\kappa \in [0, 1]$.

Theorem 5. Let π_k be the policy at the k-th iteration of κ -API and δ be the error as defined in Definition 2. Then

$$\mu(v^*-v^{\pi_k}) \leq \frac{C_{\kappa-\mathrm{API}}(\mu,\nu)}{(1-\gamma)^2}\delta + \xi^k \frac{R_{\mathrm{max}}}{1-\gamma}.$$
 Also, let $k = \left\lceil \frac{\log \frac{R_{\mathrm{max}}}{\delta(1-\gamma)}}{1-\xi} \right\rceil$. Then $\mu(v^*-v^{\pi_k}) \leq \frac{C_{\kappa-\mathrm{API}}^{(k,1)}(\mu,\nu)}{(1-\gamma)^2} \log \left(\frac{R_{\mathrm{max}}}{(1-\gamma)\delta}\right) \delta + \frac{C_{\kappa-\mathrm{API}}^{(k,2)}(\mu,\nu)}{(1-\gamma)^2} \delta + \delta.$

For brevity, we now discuss the first part of the statement; the same insights are true for the second as well. The bound for the original API is restored for the 1-step greedy case of $\kappa=0$, i.e, $\mu(v^*-v^{\pi_k}) \leq \frac{C^{(2)}(\mu,\nu)}{(1-\gamma)^2}\delta + \frac{\gamma^k R_{\max}}{1-\gamma}$ [19, 11]. As in the case of API, our bound consists of a fixed approximation error term and a geometrically decaying term. As for the other extreme, $\kappa=1$, we first remind that in the non-approximate case, applying $T_{\kappa=1}$ amounts to solving the original γ -discounted MDP in a single step [5, Remark 4]. In the approximate setup we investigate here, this results in the vanishing of the second, geometrically decaying term, since $\xi=0$ for $\kappa=1$. We are then left with a single constant approximation error: $\mu(v^*-v^{\pi_k}) \leq c(0)\delta$. Notice that c(0) is independent of π^* (see Definition 3). It represents the mismatch between μ and ν [9].

Next, notice that, by definition (see Definition 3), $C^{(2)}(\mu,\nu) > (1-\gamma)^2 c(0)$; i.e., $\frac{C^{(2)}(\mu,\nu)}{(1-\gamma)^2}\delta > c(0)\delta$. Given the discussion above, we have that the κ -API performance bound is *strictly* smaller with $\kappa=1$ than with $\kappa=0$. Hence, the bound suggests that κ -API is strictly better than the original API for $\kappa=1$. Since all expressions there are continuous, this behavior does not solely hold point-wise.

Remark 3 (Performance tradeoff). Naively, the above observation would lead to the choice of $\kappa=1$. However, it is reasonable to assume that δ , the error of the κ -greedy step, itself depends on κ , i.e, $\delta \equiv \delta(\kappa)$. The general form of such dependence is expected to be monotonically increasing: as the effective horizon of the surrogate $\kappa\gamma$ -discounted MDP becomes larger, its solution is harder to obtain (see Remark 1). Thus, Theorem 5 reveals a performance tradeoff as a function of κ .

6.2 κ -Policy Search by Dynamic Programming

We continue with generalizing another approximate PI method – PSDP [1, 19]. We name it κ -PSDP and introduce it in Algorithm 3. This algorithm updates the policy differently from κ -API. However, similarly to κ -API, it uses hard updates. We will show this algorithm exhibits better performance than any other previously analyzed approximate PI method [19].

The κ -PSDP algorithm, unlike κ -API, returns a sequence of deterministic policies, Π . Given this sequence, we build a single, non-stationary policy by successively running N_k steps of $\Pi[k]$, followed

by N_{k-1} steps of $\Pi[k-1]$, etc, where $\{N_i\}_{i=1}^k$ are i.i.d. geometric random variables with parameter $1-\kappa$. Once this process reaches π_0 , it runs π_0 indefinitely. We shall refer to this non-stationary policy as $\sigma_{\kappa,k}$. Its value $v^{\sigma_{\kappa,k}}$ can be seen to satisfy

$$v^{\sigma_{\kappa,k}} = T_{\kappa}^{\Pi[k]} T_{\kappa}^{\Pi[k-1]} \dots T_{\kappa}^{\Pi[1]} v^{\pi_0}.$$

This algorithm follows PSDP from [19]. Differently from it, the 1-step improvement is generalized to the κ -greedy improvement and the non-stationary policy behaves randomly. The following theorem gives a performance bound for it (see proof in Appendix F).

Theorem 6. Let $\sigma_{\kappa,k}$ be the policy at the k-th iteration of κ -PSDP and δ be the error as defined in Definition 2. Then

$$\mu(v^*-v^{\sigma_{\kappa,k}}) \leq \frac{C_{\kappa}^{\pi^*(1)}(\mu,\nu)}{1-\xi}\delta + \xi^k \frac{R_{\max}}{1-\gamma}.$$
 Also, let $k = \left\lceil \frac{\log \frac{R_{\max}}{\delta(1-\gamma)}}{1-\xi} \right\rceil$. Then $\mu(v^*-v^{\sigma_{\kappa,k}}) \leq \frac{C_{\kappa}^{\pi^*}(\mu,\nu)}{(1-\xi)^2} \log \left(\frac{R_{\max}}{(1-\gamma)\delta}\right)\delta + \delta.$

Compared to κ -API from the previous section, the κ -PSDP bound consists of a different fixed approximation error and a shared geometrically decaying term. Regarding the former, notice that $C_{\kappa}^{\pi^*(1)}(\mu,\nu) \prec C_{\kappa-\mathrm{API}}(\mu,\nu)$, using the notation from Remark 2. This suggests that κ -PSDP is strictly better than κ -API in the metrics we consider, and is in line with the comparison of the original API to the original PSDP given in [19].

Similarly to the previous section, we again see that substituting $\kappa=1$ gives a tighter bound than $\kappa=0$. The reason is that $\frac{C^{\pi^*(1)}(\mu,\nu)}{1-\gamma}\delta>c(0)\delta$, by definition (see Definition 3); i.e., we have that κ -PSDP is generally better than PSDP. Also, contrarily to κ -API, here we directly see the performance improvement as κ increases due to the decrease of $C_\kappa^{\pi^*}$ prescribed in Lemma 4, for the construction given there. Moreover, the κ tradeoff discussion in Remark 3 applies here as well.

An additional advantage of this new algorithm over PSDP is reduced space complexity. This can be seen from the $1-\xi$ in the denominator in the choice of k in the second part of Theorem 6. It shows that, since ξ is a strictly decreasing function of κ , better performance is guaranteed with significantly fewer iterations by increasing κ . Since the size of stored policy Π is linearly dependent on the number of iterations, larger κ improves space efficiency.

7 Discussion and Future Work

In this work, we introduced and analyzed online and approximate PI methods, generalized to the κ -greedy policy, an instance of a multiple-step greedy policy. Doing so, we discovered two intriguing properties compared to the well-studied 1-step greedy policy, which we believe can be impactful in designing state-of-the-art algorithms. First, successive application of multiple-step greedy policies with a soft, stepsize-based update does not guarantee improvement; see Theorem 1. To mitigate this caveat, we designed an online PI algorithm with a 'cautious' improvement operator; see Section 5.

The second property we find intriguing stemmed from analyzing κ generalizations of known approximate hard-update PI methods. In Section 6, we revealed a performance tradeoff in κ , which can be interpreted as a tradeoff between short-horizon bootstrap bias and long-rollout variance. This corresponds to the known λ tradeoff in the famous $TD(\lambda)$.

The two characteristics above lead to new compelling questions. The first regards improvement operators: would a non-monotonically improving PI scheme necessarily not converge to the optimal policy? Our attempts to generalize existing proof techniques to show convergence in such cases have fallen behind. Specifically, in the online case, Lemma 5.4 in [10] does not hold with multiple-step greedy policies. Similar issues arise when trying to form a κ -CPI algorithm via, e.g., an attempt to generalize Corollary 4.2 in [9]. Another research question regards the choice of the parameter κ given the tradeoff it poses. One possible direction for answering it could be investigating the concentrability coefficients further and attempting to characterize them for specific MDPs, either theoretically or via estimation. Lastly, a next indisputable step would be to empirically evaluate implementations of the algorithms presented here.

Acknowledgments

This work was partially funded by the Israel Science Foundation under contract 1380/16.

References

- [1] J Andrew Bagnell, Sham M Kakade, Jeff G Schneider, and Andrew Y Ng. Policy search by dynamic programming. In *Advances in neural information processing systems*, pages 831–838, 2004.
- [2] Dimitri P Bertsekas and John N Tsitsiklis. Neuro-dynamic programming: an overview. In *Decision and Control, 1995., Proceedings of the 34th IEEE Conference on*, volume 1, pages 560–564. IEEE, 1995.
- [3] Bruno Bouzy and Bernard Helmstetter. Monte-carlo go developments. In *Advances in computer games*, pages 159–174. Springer, 2004.
- [4] Cameron B Browne, Edward Powley, Daniel Whitehouse, Simon M Lucas, Peter I Cowling, Philipp Rohlfshagen, Stephen Tavener, Diego Perez, Spyridon Samothrakis, and Simon Colton. A survey of monte carlo tree search methods. *IEEE Transactions on Computational Intelligence and AI in games*, 4(1):1–43, 2012.
- [5] Yonathan Efroni, Gal Dalal, Bruno Scherrer, and Shie Mannor. Beyond the one step greedy approach in reinforcement learning. *arXiv preprint arXiv:1802.03654*, 2018.
- [6] Damien Ernst, Mevludin Glavic, Florin Capitanescu, and Louis Wehenkel. Reinforcement learning versus model predictive control: a comparison on a power system problem. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 39(2):517–529, 2009.
- [7] Amir-massoud Farahmand, Csaba Szepesvári, and Rémi Munos. Error propagation for approximate policy and value iteration. In *Advances in Neural Information Processing Systems*, pages 568–576, 2010.
- [8] Nan Jiang, Alex Kulesza, Satinder Singh, and Richard Lewis. The dependence of effective planning horizon on model accuracy. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, pages 1181–1189. International Foundation for Autonomous Agents and Multiagent Systems, 2015.
- [9] S.M. Kakade and J. Langford. Approximately Optimal Approximate Reinforcement Learning. In *International Conference on Machine Learning*, pages 267–274, 2002.
- [10] Vijaymohan R Konda and Vivek S Borkar. Actor-critic-type learning algorithms for markov decision processes. *SIAM Journal on control and Optimization*, 38(1):94–123, 1999.
- [11] Alessandro Lazaric, Mohammad Ghavamzadeh, and Rémi Munos. Analysis of classification-based policy iteration algorithms. *The Journal of Machine Learning Research*, 17(1):583–612, 2016.
- [12] 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.
- [13] Rémi Munos. Error bounds for approximate policy iteration. In *Proceedings of the Twentieth International Conference on International Conference on Machine Learning*, pages 560–567. AAAI Press, 2003.
- [14] Rémi Munos. Performance bounds in l_p-norm for approximate value iteration. *SIAM journal on control and optimization*, 46(2):541–561, 2007.
- [15] Rudy R Negenborn, Bart De Schutter, Marco A Wiering, and Hans Hellendoorn. Learning-based model predictive control for markov decision processes. *IFAC Proceedings Volumes*, 38(1):354–359, 2005.
- [16] Steven Perkins and David S Leslie. Asynchronous stochastic approximation with differential inclusions. *Stochastic Systems*, 2(2):409–446, 2013.
- [17] Marek Petrik and Bruno Scherrer. Biasing approximate dynamic programming with a lower discount factor. In Advances in neural information processing systems, pages 1265–1272, 2009.

- [18] Martin L Puterman. Markov decision processes: discrete stochastic dynamic programming. John Wiley & Sons, 1994.
- [19] Bruno Scherrer. Approximate policy iteration schemes: a comparison. In *International Conference on Machine Learning*, pages 1314–1322, 2014.
- [20] Bruno Scherrer. Improved and Generalized Upper Bounds on the Complexity of Policy Iteration. INFORMS, February 2016. Markov decision processes; Dynamic Programming; Analysis of Algorithms.
- [21] Bruno Scherrer and Matthieu Geist. Local policy search in a convex space and conservative policy iteration as boosted policy search. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 35–50. Springer, 2014.
- [22] 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.
- [23] Brian Sheppard. World-championship-caliber scrabble. Artificial Intelligence, 134(1-2):241–275, 2002.
- [24] David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dharshan Kumaran, Thore Graepel, et al. Mastering chess and shogi by self-play with a general reinforcement learning algorithm. *arXiv* preprint *arXiv*:1712.01815, 2017.
- [25] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of go without human knowledge. *Nature*, 550(7676):354, 2017.
- [26] Alexander L Strehl, Lihong Li, and Michael L Littman. Reinforcement learning in finite mdps: Pac analysis. *Journal of Machine Learning Research*, 10(Nov):2413–2444, 2009.
- [27] Aviv Tamar, Garrett Thomas, Tianhao Zhang, Sergey Levine, and Pieter Abbeel. Learning from the hindsight plan—episodic mpc improvement. In *Robotics and Automation (ICRA)*, 2017 *IEEE International Conference on*, pages 336–343. IEEE, 2017.
- [28] Gerald Tesauro and Gregory R Galperin. On-line policy improvement using monte-carlo search. In *Advances in Neural Information Processing Systems*, pages 1068–1074, 1997.
- [29] Christopher JCH Watkins and Peter Dayan. Q-learning. Machine learning, 8(3-4):279–292, 1992.

A Proof of Theorem 1

We start with a generalization of a useful lemma; its original version appeared in, e.g., [20, Lemma 10]. **Lemma 7.** Let v be a value function, π a policy, and $\kappa \in [0, 1]$. Then

$$T_{\kappa}^{\pi}v - v = (I - \kappa \gamma P^{\pi})^{-1}(T^{\pi}v - v).$$

Proof. The proof is a straightforward generalization of the proof in [20, Lemma 10], and [9, Remark 6.1].

$$\begin{split} T_{\kappa}^{\pi}v - v &= (I - \kappa\gamma P^{\pi})^{-1}(r^{\pi} + (1 - \kappa)\gamma P^{\pi}v) - v \\ &= (I - \kappa\gamma P^{\pi})^{-1}(r^{\pi} + (1 - \kappa)\gamma P^{\pi}v - (I - \kappa\gamma P^{\pi})v) \\ &= (I - \kappa\gamma P^{\pi})^{-1}(r^{\pi} + \gamma P^{\pi}v - v) \\ &= (I - \kappa\gamma P^{\pi})^{-1}(T^{\pi}v - v). \end{split}$$

This elementary lemma relates the ' κ -advantage' to the 1-step advantage and is useful to prove Theorem 1 and some following results.

First, since $\pi(\alpha, \kappa) = (1 - \alpha)\pi + \alpha\pi_{\kappa}$, we have that

$$P^{\pi(\alpha,\kappa)} = (1-\alpha)P^{\pi} + \alpha P^{\pi_{\kappa}},$$

$$r^{\pi(\alpha,\kappa)} = (1-\alpha)r^{\pi} + \alpha r^{\pi_{\kappa}};$$

thus, since v^{π} is the fixed-point of T^{π} ,

$$T^{\pi(\alpha,\kappa)}v^{\pi} = (1-\alpha)T^{\pi}v^{\pi} + \alpha T^{\pi_{\kappa}}v^{\pi} = (1-\alpha)v^{\pi} + \alpha T^{\pi_{\kappa}}v^{\pi}.$$
 (7)

Using this, we now prove the first statement of Theorem 1.

$$v^{\pi(\alpha,\kappa)} - v^{\pi} = (I - \gamma P^{\pi(\alpha,\kappa)})^{-1} (T^{\pi(\alpha,\kappa)} v^{\pi} - v^{\pi})$$

$$= \alpha (I - \gamma P^{\pi(\alpha,\kappa)})^{-1} (T^{\pi_{\kappa}} v^{\pi} - v^{\pi})$$

$$= \alpha (I - \gamma P^{\pi(\alpha,\kappa)})^{-1} (I - \kappa \gamma P^{\pi_{\kappa}}) (I - \kappa \gamma P^{\pi_{\kappa}})^{-1} (T^{\pi_{\kappa}} v^{\pi} - v^{\pi})$$

$$= \alpha (I - \gamma P^{\pi(\alpha,\kappa)})^{-1} (I - \kappa \gamma P^{\pi_{\kappa}}) (T^{\pi_{\kappa}}_{\kappa} v^{\pi} - v^{\pi})$$

$$= \alpha (I - \gamma P^{\pi(\alpha,\kappa)})^{-1} (I - \gamma P^{\pi(\alpha,\kappa)} + \gamma (P^{\pi(\alpha,\kappa)} - \kappa P^{\pi_{\kappa}})) (T^{\pi_{\kappa}}_{\kappa} v^{\pi} - v^{\pi})$$

$$= \alpha \left(I + \gamma (I - \gamma P^{\pi(\alpha,\kappa)})^{-1}\right) ((1 - \alpha) P^{\pi} + (\alpha - \kappa) P^{\pi_{\kappa}}) (T^{\pi_{\kappa}}_{\kappa} v^{\pi} - v^{\pi}). \tag{8}$$

For the first relation we use Lemma 7 with $\kappa=1$ and the fact that, by definition, $T_{\kappa=1}^{\pi(\alpha,\kappa)}v^{\pi(\alpha,\kappa)}=v^{\pi(\alpha,\kappa)}$. For the second relation we use (7), for the fourth we again use Lemma 7, and for the last relation we use that $P^{\pi(\alpha,\kappa)}-\kappa P^{\pi_{\kappa}}=(1-\alpha)P^{\pi}+(\alpha-\kappa)P^{\pi_{\kappa}}$.

Next, we show that for $\alpha \geq \kappa$, all terms in (8) are component-wise bigger than or equal to zero. First, using a Taylor expansion, $(I-\gamma P^{\pi(\alpha,\kappa)})^{-1} = \sum_t \gamma^t (P^{\pi(\alpha,\kappa)})^t \geq 0$ component-wise, since it is a weighted sum of transition matrices with positive weights. The same applies for $(1-\alpha)P^\pi + (\alpha-\kappa)P^{\pi_\kappa}$, when $\alpha \geq \kappa$. Thus, for $\alpha \geq \kappa$, $(I+\gamma(I-\gamma P^{\pi(\alpha,\kappa)})^{-1}((1-\alpha)P^\pi + (\alpha-\kappa)P^{\pi_\kappa}) \geq 0$ component-wise. Lastly, since $\pi_\kappa \in \mathcal{G}_\kappa(v^\pi)$, $v^\pi = T_\kappa^\pi v^\pi \leq T_\kappa v^\pi = T_\kappa^{\pi_\kappa} v^\pi$, with equality holding if and only if $v^\pi = v^*$ [5, Lemma 3]. Thus, $T_\kappa^{\pi_\kappa} v^\pi - v^\pi \geq 0$. This concludes the proof for the first statement, for the κ -greedy policy.

For the κ -greedy policy part of the proof for the second statement, we now provide more details on the counterexample presented in Section 4. For convenience, we bring the MDP example here again in Fig. 2. Consider the mixture of the "hesitant" and "confident" policies: $\pi(\alpha, \kappa = 1) = (1 - \alpha)\pi_0 + \alpha\pi(\alpha, \kappa = 1)$. It can be shown that its value is

$$v^{\pi(\alpha,\kappa=1)}(s_0) = \frac{\gamma\alpha}{1 - \gamma(1 - \alpha)} v^{\pi(\alpha,\kappa=1)}(s_1),$$

$$v^{\pi(\alpha,\kappa=1)}(s_1) = \gamma \frac{-c(1 - \alpha) + \alpha}{1 - \gamma}.$$

11




Figure 2: The Tightrope Walking MDP used in the proof of Theorem 1. This class of MDPs is parametrized by c>0.

Thus, we deduce that for any $\alpha \in (0,1)$ and

$$c > \frac{\alpha}{1 - \alpha},\tag{9}$$

 $v^{\pi(\alpha,\kappa=1)}(s_0) < v^{\pi}(s_0) = 0$, i.e, the mixture policy, $\pi(\alpha,\kappa=1)$, is not strictly better then π_0 .

We now find the conditions to ensure that the κ -greedy policy w.r.t. v^{π_0} is the optimal policy; this will generalize the above construction, made for $\kappa=1$, to any $\kappa\in[0,1]$. Observe that for any c>0 and κ it holds that $\pi_{\kappa}(s_1)=a_1=\pi^*(s_1)$, where $\pi_{\kappa}\in\mathcal{G}_{\kappa}(v^{\pi_0})$. Thus, we solely need to consider the policy which is different than π^* at state s_0 , $\tilde{\pi}(s_0)=a_0\neq\pi^*(s_0)$ and $\tilde{\pi}(s_1)=\pi^*(s_1)$. To find which condition ensures the κ -greedy policy w.r.t. v^{π_0} is π^* (and not $\tilde{\pi}$), we require

$$T_{\kappa}^{\pi^*} v^{\pi_0}(s_0) \ge T_{\kappa}^{\tilde{\pi}} v^{\pi_0}(s_0).$$
 (10)

Satisfying this condition insures that $\pi^* \in \mathcal{G}_{\kappa}(v^{\pi_0})$. By definition,

$$T_{\kappa}^{\pi^*} v^{\pi_0}(s_0) = \mathbb{E}^{\pi^*} \left[\sum_{t} (\kappa \gamma)^t (r(s_t, \pi^*(s_t)) + \gamma(1 - \kappa) v^{\pi_0}(s_{t+1}) \mid s_{t=0} = s_0 \right]$$

$$= (\kappa \gamma)^0 \left(\gamma(1 - \kappa) v^{\pi_0}(s_1) \right) + (\kappa \gamma)^1 \left(\gamma(1 - \kappa) v^{\pi_0}(s_2) \right) + \sum_{t=2}^{\infty} (\kappa \gamma)^t (1 + v^{\pi_0}(s_2))$$

$$= (\kappa \gamma)^0 \left(\gamma(1 - \kappa) (-\frac{\gamma c}{1 - \gamma}) \right) + (\kappa \gamma)^1 \left(\gamma(1 - \kappa) \frac{1}{1 - \gamma} \right) + \sum_{t=2}^{\infty} (\kappa \gamma)^t (1 + \gamma(1 - \kappa) \frac{1}{1 - \gamma})$$

$$= \gamma(1 - \kappa) (-\frac{\gamma c}{1 - \gamma}) + \kappa \gamma \frac{\gamma}{1 - \gamma}. \tag{11}$$

Similarly, and since $\tilde{\pi}(s_0) = a_0$, we have that

$$T_{\kappa}^{\tilde{\pi}}v^{\pi_0}(s_0) = 0 \tag{12}$$

Plugging (11) and (12) into (10), we get the condition

$$c \le \frac{\kappa}{1 - \kappa}.\tag{13}$$

To finalize the counterexample and show that strict policy improvement is not guaranteed, we choose c such that both (9) and (13) are satisfied. Such feasible choice exists when $\alpha < \kappa$, due to the monotonicity of $\frac{x}{1-x}$.

The monotonic improvement of $\pi(\alpha,h)$ for $\alpha=1$ was proved in [5, Lemma 1]. To build the counter example, again consider the Tightrope MDP. Let π_0 be the 'hesitant' policy. For any $\gamma\in(0,1),\,h>1$, it holds that $\pi^*\in\mathcal{G}_h(v^{\pi_0})$. Thus, it suffices to satisfy (9) alone to show that $\pi(\alpha,h)=(1-\alpha)\pi_0+\alpha\pi^*$ is not monotonically better then π . Large enough c value ensures that.

B Proof of Lemma 2

We start by showing the contraction property of H_{κ}^{π} . Let (s,a) be a fixed state-action pair, $Q_1,Q_2 \in \mathbb{R}^{2|\mathcal{S}\times\mathcal{A}|}$. For any state-action pair $(s,a),Q_i(s,a)$ is a two-component vector. We denote its first component by $q_i(s,a)$ and its second component by $q_{i,\kappa}(s,a)$. See that

$$||q_1 - q_2||_{\infty} \le ||Q_1 - Q_2||_{\infty},\tag{14}$$

$$||q_{1,\kappa} - q_{2,\kappa}||_{\infty} \le ||Q_1 - Q_2||_{\infty}.$$
 (15)

Taking a component-wise absolute value, we have that

$$\begin{aligned} &|H_{\kappa}^{\pi}Q_{1}-H_{\kappa}^{\pi}Q_{2}|(s,a)\\ =&|H_{\kappa}^{\pi}(q_{1},q_{1,\kappa})-H_{\kappa}^{\pi}(q_{2},q_{2,\kappa})|(s,a)\\ =&\gamma\left[\frac{|\mathbb{E}_{s',a^{\pi}}\left[q_{1}(s',a^{\pi}))-q_{2}(s',\pi(s'))\right]|}{|(1-\kappa)\mathbb{E}_{s',a^{\pi}}\left[q_{1}(s',a^{\pi})-q_{2}(s',a^{\pi})\right)\right]+\kappa\mathbb{E}_{s'}\left[\max_{a'}q_{1,\kappa}(s',a')-\max_{a'}q_{2,\kappa}(s',a')\right]|}\right],\\ \text{where } s'\sim P(\cdot\mid s,a), a^{\pi}\sim\pi(s'). \end{aligned}$$

Let us focus on the first component of the above vector. We have that

$$\gamma |\mathbb{E}_{s',a^{\pi}} \left[q_1(s',a^{\pi}) - q_2(s',a^{\pi}) \right] | \leq \gamma ||q_1 - q_2||_{\infty} \leq \gamma ||Q_1 - Q_2||_{\infty},$$

where we used the standard bound, $|\mathbb{E}[X]| \leq ||X||_{\infty}$ and (14). Similarly, for the second component, we have that

$$\gamma \left| \left((1 - \kappa) \mathbb{E}_{s',a^{\pi}} \left[q_{1}(s', a^{\pi}) - q_{2}(s', a^{\pi}) \right] + \kappa \mathbb{E}_{s',a} \left[\max_{a'} q_{1,\kappa}(s', a') - \max_{a'} q_{2,\kappa}(s', a') \right] \right) \right|$$

$$\leq \gamma \left((1 - \kappa) |\mathbb{E}_{s',a^{\pi}} \left[q_{1}(s', a^{\pi}) - q_{2}(s', a^{\pi}) \right] | + \kappa \mathbb{E}_{s',a} \left[|\max_{a'} q_{1,\kappa}(s', a') - \max_{a'} q_{2,\kappa}(s', a') \right] \right]$$

$$\leq \gamma \left((1 - \kappa) |\mathbb{E}_{s',a^{\pi}} \left[q_{1}(s', a^{\pi}) - q_{2}(s', a^{\pi}) \right] | + \kappa \mathbb{E}_{s',a'} \left[\max_{a'} |q_{1,\kappa}(s', a') - q_{2,\kappa}(s', a') \right] \right)$$

$$\leq \gamma \left((1 - \kappa) ||q_{1} - q_{2}||_{\infty} + \kappa ||q_{1,\kappa} - q_{2,\kappa}||_{\infty} \right)$$

$$\leq \gamma \left((1 - \kappa) ||Q_{1} - Q_{2}||_{\infty} + \kappa ||Q_{1} - Q_{2}||_{\infty} \right)$$

where for the first relation we used the triangle inequality, for the second we used the standard bound $|\max_{x\in\mathcal{X}}f(x)-\max_{x\in\mathcal{X}}g(x)|\leq \max_{x\in\mathcal{X}}|f(x)-g(x)|$, for the third we used the bound $|\mathbb{E}[X]|\leq ||X||_{\infty}$, and for the last (14)-(15).

From the above we get that

$$||H_{\kappa}^{\pi}Q_1 - H_{\kappa}^{\pi}Q_2||_{\infty} < \gamma ||Q_1 - Q_2||_{\infty};$$

i.e., the operator H^π_κ is a γ contraction mapping in the max-norm.

It is clear that the fixed point of the first component is q^{π} . The fixed point of the second component is the fixed point of the optimal Bellman operator of the $\kappa\gamma$ -discounted, reward shaped, surrogate MDP (see Remark 1). Its solution is, by construction, q_{κ}^{π} (see (4)).

C Proof of Theorem 3

The proof of Theorem 3 follows the proof in [16, Section 5.1], with several generalizations given below.

C.1 Lipschitzness of the Slow Time Scale Fixed-Point

Before following the main lemmas in [16] and showing they hold for Online κ -PI (Algorithm 1), we shall show that the solution of the fast-time scale ODE (found using a fixed-point argument), $[q^{\pi}, q_{\kappa}^{\pi}]$, is Lipschitz-continuous in the slow time-scale iterate, π .

Lemma 8. Let $\pi: \mathcal{S} \times \mathcal{A} \to [0,1]$ be a stochastic policy. For any π_1, π_2 and $q_1, q_2 \in \mathbb{R}^{|\mathcal{S} \times \mathcal{A}|}$, let

$$||\pi_1 - \pi_2||_{\infty} \stackrel{\text{def}}{=} \max_s \sum_a |\pi_1(a \mid s) - \pi_2(a \mid s)|,$$

$$||q_1 - q_2||_{\infty} \stackrel{def}{=} \max_{s,a} |q_1(s,a) - q_2(s,a)|.$$

Then q^{π} and q_{κ}^{π} are Lipschitz-continuous in π in the max-norm; i.e.,

$$||q^{\pi_1} - q^{\pi_2}||_{\infty} \le L_a ||\pi_1 - \pi_2||_{\infty},$$

$$||q_{\kappa}^{\pi_1} - q_{\kappa}^{\pi_2}||_{\infty} \le L_b ||\pi_1 - \pi_2||_{\infty},$$

where $L_a, L_b > 0$, are functions of $\gamma, \kappa, R_{\text{max}}$.

Proof. We start by proving that $||v^{\pi_1} - v^{\pi_2}||_{\infty} \le L||\pi_1 - \pi_2||_{\infty}$, i.e, v^{π} is Lipschitz in π .

$$||v^{\pi_{1}} - v^{\pi_{2}}||_{\infty} = ||T^{\pi_{1}}v^{\pi_{1}} - T^{\pi_{2}}v^{\pi_{2}}||_{\infty}$$

$$\leq ||T^{\pi_{1}}v^{\pi_{1}} - T^{\pi_{1}}v^{\pi_{2}} + T^{\pi_{1}}v^{\pi_{2}} - T^{\pi_{2}}v^{\pi_{2}}||_{\infty}$$

$$\leq ||T^{\pi_{1}}v^{\pi_{1}} - T^{\pi_{1}}v^{\pi_{2}}||_{\infty} + ||T^{\pi_{1}}v^{\pi_{2}} - T^{\pi_{2}}v^{\pi_{2}}||_{\infty}$$

$$\leq \gamma ||v^{\pi_{1}} - v^{\pi_{2}}||_{\infty} + ||T^{\pi_{1}}v^{\pi_{2}} - T^{\pi_{2}}v^{\pi_{2}}||_{\infty},$$
(16)

where the last relation is due to the fact T^{π_1} is a γ -contraction. We continue by calculating $|T^{\pi_1}v^{\pi_2}-T^{\pi_2}v^{\pi_2}|(s)$.

$$|T^{\pi_1}v^{\pi_2} - T^{\pi_2}v^{\pi_2}|(s) \le |\sum_{a} (\pi_1(a \mid s) - \pi_2(a \mid s))r(s, a)| + \gamma |\sum_{s'} (P_{s', s}^{\pi_1} - P_{s', s}^{\pi_2})v^{\pi_2}(s')|.$$

$$(17)$$

We bound each term in (17). The first term can be bounded by,

$$|\sum_{a} (\pi_{1}(a \mid s) - \pi_{2}(a \mid s))r(s, a)| \leq \sum_{a} |(\pi_{1}(a \mid s) - \pi_{2}(a \mid s))||r(s, a)|$$

$$\leq R_{\max} \max_{s} \sum_{a} |(\pi_{1}(a \mid s) - \pi_{2}(a \mid s))|$$

$$= R_{\max}||\pi_{1} - \pi_{2}||_{\infty}.$$
(18)

In the first relation we used the triangle inequality and in the second inequality the fact that |r(s, a)| is bounded by R_{\max} .

The second term in (17) can be bounded by,

$$|\sum_{s'} (P_{s',s}^{\pi_{1}} - P_{s',s}^{\pi_{2}}) v^{\pi_{2}}(s')| = |\sum_{s',a} P(s' \mid s,a) (\pi_{1}(a \mid s) - \pi_{2}(a \mid s)) v^{\pi_{2}}(s')|$$

$$\leq \sum_{a} \sum_{s'} P(s' \mid s,a) |(\pi_{1}(a \mid s) - \pi_{2}(a \mid s)) v^{\pi_{2}}(s')|$$

$$\leq \sum_{a} \sum_{s'} P(s' \mid s,a) |(\pi_{1}(a \mid s) - \pi_{2}(a \mid s))| |v^{\pi_{2}}(s')|$$

$$\leq \sum_{a} \sum_{s'} P(s' \mid s,a) |(\pi_{1}(a \mid s) - \pi_{2}(a \mid s))| \frac{R_{\max}}{1 - \gamma}$$

$$= \sum_{a} |(\pi_{1}(a \mid s) - \pi_{2}(a \mid s))| \frac{R_{\max}}{1 - \gamma} \sum_{s'} P(s' \mid s,a)$$

$$= \sum_{a} |(\pi_{1}(a \mid s) - \pi_{2}(a \mid s))| \frac{R_{\max}}{1 - \gamma}$$

$$\leq \max_{s} \sum_{a} |(\pi_{1}(a \mid s) - \pi_{2}(a \mid s))| \frac{R_{\max}}{1 - \gamma} = \frac{R_{\max}}{1 - \gamma} ||\pi_{1} - \pi_{2}||_{\infty}$$
(19)

In the first relation we used the triangle inequality, in the forth relation we used the fact that for any π and $s, v^{\pi}(s) \leq \frac{R_{\max}}{1-\gamma}$, and in the fifth relation the fact that for any s and $a, P(s' \mid s, a)$ is a probability function, thus sums to one.

Using (18), (19) to bound (17) yields that for any s,

$$|T^{\pi_1}v^{\pi_2} - T^{\pi_2}v^{\pi_2}|(s) \le \frac{R_{\max}}{1-\gamma}||\pi_1 - \pi_2||_{\infty}.$$

Thus, $||T^{\pi_1}v^{\pi_2} - T^{\pi_2}v^{\pi_2}||_{\infty} \le \frac{R_{\text{max}}}{1-\gamma}||\pi_1 - \pi_2||_{\infty}$. Plugging this bound into (16) and rearranging yields,

$$||v^{\pi_1} - v^{\pi_2}||_{\infty} \le \frac{R_{\text{max}}}{(1 - \gamma)^2} ||\pi_1 - \pi_2||_{\infty},$$
 (20)

giving that $L = \frac{R_{\text{max}}}{(1-\gamma)^2}$.

We continue by analysing $||T_{\kappa}v^{\pi_1}-T_{\kappa}v^{\pi_2}||_{\infty}$. We remind the reader that $T_{\kappa}v^{\pi}$ satisfies the following fixed-point equation:

$$T_{\kappa}v^{\pi}(s) = \max_{a} \left[r(s, a) + \gamma(1 - \kappa) \sum_{s'} P(s' \mid s, a)v^{\pi}(s') + \kappa \gamma \sum_{s'} P(s' \mid s, a)(T_{\kappa}v^{\pi})(s') \right]$$

$$\stackrel{\text{def}}{=} \bar{T}_{\kappa}^{\pi} T_{\kappa}v^{\pi}(s),$$

where we defined the 'optimal' Bellman operator of the surrogate MDP to be \bar{T}_{κ}^{π} (see Remark 1). Furthermore, since this operator is the optimal Bellman operator of a $\kappa\gamma$ -discounted MDP, it is a $\kappa\gamma$ contraction mapping. We now use a similar technique as the above to show $||T_{\kappa}v^{\pi_1}-T_{\kappa}v^{\pi_2}||_{\infty} \leq L_{\kappa}||\pi_1-\pi_2||_{\infty}$, i.e, $T_{\kappa}v^{\pi}$ is Lipschitz in π .

$$\begin{split} ||T_{\kappa}v^{\pi_{1}} - T_{\kappa}v^{\pi_{2}}||_{\infty} &= ||\bar{T}_{\kappa}^{\pi_{1}}T_{\kappa}v^{\pi_{1}} - \bar{T}_{\kappa}^{\pi_{2}}T_{\kappa}v^{\pi_{2}}||_{\infty} \\ &\leq ||\bar{T}_{\kappa}^{\pi_{1}}T_{\kappa}v^{\pi_{1}} - \bar{T}_{\kappa}^{\pi_{1}}T_{\kappa}v^{\pi_{2}}||_{\infty} + ||\bar{T}_{\kappa}^{\pi_{1}}T_{\kappa}v^{\pi_{2}} - \bar{T}_{\kappa}^{\pi_{2}}T_{\kappa}v^{\pi_{2}}||_{\infty} \\ &\leq \kappa\gamma||T_{\kappa}v^{\pi_{1}} - T_{\kappa}v^{\pi_{2}}||_{\infty} + ||\bar{T}_{\kappa}^{\pi_{1}}T_{\kappa}v^{\pi_{2}} - \bar{T}_{\kappa}^{\pi_{2}}T_{\kappa}v^{\pi_{2}}||_{\infty}. \end{split}$$

We now bound the second term.

$$\begin{split} |\bar{T}_{\kappa}^{\pi_{1}} T_{\kappa} v^{\pi_{2}} - \bar{T}_{\kappa}^{\pi_{2}} T_{\kappa} v^{\pi_{2}} |(s) &\leq \max_{a} \gamma(1 - \kappa) |\sum_{s'} P(s' \mid s, a) (v^{\pi_{1}} - v^{\pi_{2}}) (s') | \\ &\leq \max_{a} \gamma(1 - \kappa) \sum_{s'} P(s' \mid s, a) |v^{\pi_{1}} - v^{\pi_{2}}| (s') \\ &\leq \max_{a} \gamma(1 - \kappa) \sum_{s'} P(s' \mid s, a) ||v^{\pi_{1}} - v^{\pi_{2}}||_{\infty} = \gamma(1 - \kappa) ||v^{\pi_{1}} - v^{\pi_{2}}||_{\infty}, \end{split}$$

where we used the definition of \bar{T}_{κ}^{π} and the identity $|\max_{x \in \mathcal{X}} f(x) - \max_{x \in \mathcal{X}} g(x)| \le \max_{x \in \mathcal{X}} |f(x) - g(x)|$ in the first relation and the triangle inequality in the second.

Using (20), we have

$$||T_{\kappa}v^{\pi_{1}} - T_{\kappa}v^{\pi_{2}}||_{\infty} \leq \frac{\gamma(1-\kappa)}{1-\kappa\gamma}||v^{\pi_{1}} - v^{\pi_{2}}||_{\infty}$$
$$\leq \frac{\gamma(1-\kappa)}{1-\kappa\gamma} \frac{R_{\max}}{(1-\gamma)^{2}}||\pi_{1} - \pi_{2}||_{\infty}.$$

These results transform to results on q^{π} and q_{κ}^{π} as follows. Starting with q^{π}

$$|q^{\pi_1} - q^{\pi_2}|(s, a) = |r(s, a) + \gamma \sum_{s'} P(s' \mid s, a) v^{\pi_1} - r(s, a) - \gamma \sum_{s'} P(s' \mid s, a) v^{\pi_2}|$$

$$= \gamma |\sum_{s'} P(s' \mid s, a) (v^{\pi_1} - v^{\pi_2})| \le \gamma ||v^{\pi_1} - v^{\pi_2}||_{\infty}.$$

By taking the max-norm on both sides we get the result since $||v^{\pi_1} - v^{\pi_2}||_{\infty}$ was shown to be Lipschitz in π .

Next, for q_{κ}^{π} we have

$$\begin{aligned} &|q_{\kappa}^{\pi_{1}} - q_{\kappa}^{\pi_{2}}|(s, a) \\ &= |\gamma(1 - \kappa) \sum_{s'} P(s' \mid s, a)(v^{\pi_{1}}(s') - v^{\pi_{2}}(s')) + \kappa \gamma \sum_{s'} P(s' \mid s, a)(T_{\kappa}v^{\pi_{1}} - T_{\kappa}v^{\pi_{2}})(s')| \\ &\leq \gamma(1 - \kappa)||v^{\pi_{1}}(s') - v^{\pi_{2}}(s')||_{\infty} + \kappa \gamma||T_{\kappa}v^{\pi_{1}} - T_{\kappa}v^{\pi_{2}}||_{\infty}. \end{aligned}$$

By taking the max-norm on both sides we get the result since, as shown above, both $||v^{\pi_1} - v^{\pi_2}||_{\infty}$ and $||T_{\kappa}v^{\pi_1} - T_{\kappa}v^{\pi_2}||_{\infty}$ are Lipschitz in π . Finally, since the vector space is finite (due to the finite state and action space), all L_p norms are equivalent. Thus, the Lipschitzness result applies in any L_p norm as well.

C.2 Improvement Step

Here, we prove an equivalent lemma to [16, Lemma 5.4] which shows that the mean value of the process improves. Denote $b_s \equiv b_s(q^\pi, q_\kappa^\pi, \pi)$ as the policy defined in the Algrorithm 1. By using Lemma 7 and setting $\kappa = 0$ we have that

$$v^{(1-\alpha)\pi + \alpha b_s} - v^{\pi} = \alpha (I - \gamma P^{(1-\alpha)\pi + \alpha b_s})^{-1} (T^{b_s} v^{\pi} - v^{\pi}).$$

Thus, by taking the limit $\alpha \to 0$ we have

$$\lim_{\alpha \to 0} (v^{(1-\alpha)\pi + \alpha b_s} - v^{\pi}) = \alpha \nabla_{\pi} v^{\pi} (b_s - \pi)$$

$$= \alpha \langle \nabla_{\pi} v^{\pi}, \Delta \pi \rangle$$

$$= \alpha (I - \gamma P^{\pi})^{-1} (T^{b_s} v^{\pi} - v^{\pi}) + \mathcal{O}(\alpha^2) \ge 0,$$

where the last inequality is since $T^{b_s}v^{\pi}-v^{\pi}\geq 0$ by construction and $(I-\gamma P^{\pi})^{-1}\geq 0$ componentwise. We thus get that

$$\frac{1}{\alpha} \lim_{\alpha \to 0} (v^{(1-\alpha)\pi + \alpha b_s} - v^{\pi}) = \langle \nabla_{\pi} v^{\pi}, \Delta \pi \rangle \ge 0.$$

C.3 Convergence of the Algorithm

We define the same Lyapunov function as defined in [16, Lemma 5.5]. Due to previous section it is indeed a Lyapunov function since its derivative is negative and the function is bigger than 0 by construction. The presence of the Lyapunov function leads to the convergence of the policy to the optimal policy, similarly to [16, Corollary 5.6], which leads to the convergence of q^{π} to q^* . Lastly, since $T_{\kappa}v^*=v^*$ [5, Lemma 4] we have that,

$$q_{\kappa}^{\pi^*}(\pi') = r^{\pi'} + \gamma (1 - \kappa) P^{\pi'} v^* + \kappa \gamma P^{\pi'} T_{\kappa} v^*$$
$$= r^{\pi'} + \gamma (1 - \kappa) P^{\pi'} v^* + \kappa \gamma P^{\pi'} v^*$$
$$= r^{\pi'} + \gamma P^{\pi'} v^* = q^*(\pi').$$

which concludes the proof.

D Proof of Lemma 4

We first prove a useful lemma that relates the (unnormalized) future distribution, measured in different κ scales.

Lemma 9. For any policy π and κ , $\kappa' \in [0, 1]$,

$$(I - \xi_{\kappa'} D_{\kappa'}^{\pi} P^{\pi})^{-1} = \frac{\kappa' - \kappa}{1 - \kappa} I + \frac{1 - \kappa'}{1 - \kappa} (I - \xi_{\kappa} D_{\kappa}^{\pi} P^{\pi})^{-1}.$$

Proof. We prove the lemma by using the definition and by some algebraic manipulations.

$$(I - \xi_{\kappa'} D_{\kappa'}^{\pi} P^{\pi})^{-1} = (I - \gamma (1 - \kappa') (I - \kappa \gamma' P^{\pi})^{-1} P^{\pi})^{-1}$$

$$= ((I - \kappa \gamma' P^{\pi})^{-1} (I - \kappa \gamma' P^{\pi} - \gamma (1 - \kappa') P^{\pi}))^{-1}$$

$$= (I - \gamma P^{\pi})^{-1} (I - \gamma \kappa' P^{\pi})$$

$$= (I - \gamma P^{\pi})^{-1} - \kappa' \gamma P^{\pi} (I - \gamma P^{\pi})^{-1}$$

$$= (I - \gamma P^{\pi})^{-1} - \kappa' (I + \gamma P^{\pi} (I - \gamma P^{\pi})^{-1} - I)$$

$$= (I - \gamma P^{\pi})^{-1} - \kappa' ((I - \gamma P^{\pi})^{-1} - I)$$

$$= \kappa' I + (1 - \kappa') (I - \gamma P^{\pi})^{-1}$$

We see that the following relation holds for any $\kappa \in [0, 1]$,

$$(I - \gamma P^{\pi})^{-1} = \frac{1}{1 - \kappa} ((I - \xi_{\kappa} D_{\kappa}^{\pi} P^{\pi})^{-1} - \kappa I).$$

Plugging this relation into the previous one we get,

$$(I - \xi_{\kappa'} D_{\kappa'}^{\pi} P^{\pi})^{-1} = \kappa' I + (1 - \kappa') (I - \gamma P^{\pi})^{-1}$$

$$= \kappa' I + \frac{1 - \kappa'}{1 - \kappa} ((I - \xi_{\kappa} D_{\kappa}^{\pi} P^{\pi})^{-1} - \kappa I)$$

$$= \frac{\kappa' - \kappa}{1 - \kappa} I + \frac{1 - \kappa'}{1 - \kappa} (I - \xi_{\kappa} D_{\kappa}^{\pi} P^{\pi})^{-1}.$$

We are now ready to prove Lemma 4. Assume a constant $C_{\kappa}^{\pi^*}(\mu,\nu) < \infty$ such that,

$$d_{\kappa \mu}^{\pi^*} = (1 - \xi)\mu(I - \xi D_{\kappa}^{\pi^*})^{-1} < C_{\kappa}^{\pi^*}(\mu, \nu)\nu. \tag{21}$$

Given that, we shall calculate $C_{\kappa'}^{\pi^*}(\mu, \nu)$ where $\kappa' > \kappa$.

$$\begin{split} d_{\kappa',\mu}^{\pi^*} &= (1 - \xi_{\kappa'}) \mu (I - \xi D_{\kappa'}^{\pi^*})^{-1} \\ &= (1 - \xi_{\kappa'}) \left(\frac{\kappa' - \kappa}{1 - \kappa} \mu + \frac{1 - \kappa'}{1 - \kappa} \mu ((I - \xi_{\kappa} D_{\kappa}^{\pi} P^{\pi})^{-1}) \right) \\ &\leq \frac{1 - \xi_{\kappa'}}{1 - \kappa} \left((\kappa' - \kappa) \mu + \frac{1 - \kappa'}{1 - \xi_{\kappa}} C_{\kappa}^{\pi^*} (\mu, \nu) \nu \right) \\ &= \frac{1 - \xi_{\kappa'}}{1 - \kappa} (\kappa' - \kappa + \frac{1 - \kappa'}{1 - \xi_{\kappa}} C_{\kappa}^{\pi^*} (\mu, \nu)) (\alpha^* \mu + (1 - \alpha^*) \nu) \stackrel{\text{def}}{=} C_{\kappa'}^{\pi^*} (\mu, \nu(\alpha)) \nu(\alpha) d\alpha \end{split}$$

where we used Lemma 9 in the first line, Equation 21 in the second line, and defined $\alpha^* = (1 + \frac{1-\kappa'}{(1-\xi_\kappa)(\kappa'-\kappa)}C_\kappa^{\pi^*}(\mu,\nu)))^{-1} \in (0,1)$ and $C_{\kappa'}^{\pi^*}(\mu,\nu(\alpha^*)) = \frac{1-\xi_{\kappa'}}{1-\kappa}(\kappa'-\kappa+\frac{1-\kappa'}{1-\xi_\kappa}C_\kappa^{\pi^*}(\mu,\nu))$. By plugging the expressions of $\xi_\kappa, \xi_{\kappa'}$ we see that,

$$C_{\kappa'}^{\pi^*}(\mu,\nu(\alpha^*)) - C_{\kappa}^{\pi^*}(\mu,\nu) = \frac{1-\xi_{\kappa'}}{1-\kappa} (\kappa' - \kappa + (\frac{1-\kappa'}{1-\xi_{\kappa}} - \frac{1-\kappa}{1-\xi_{\kappa'}}) C_{\kappa}^{\pi^*}(\mu,\nu))$$

$$= \frac{1-\xi_{\kappa'}}{1-\kappa} (\kappa' - \kappa) (1 - C_{\kappa}^{\pi^*}(\mu,\nu)). \tag{22}$$

Since $C_{\kappa}^{\pi^*}(\mu,\nu) \geq 1$ and $\frac{1-\xi_{\kappa'}}{1-\kappa}(\kappa'-\kappa) > 0$ we get that $C_{\kappa'}^{\pi^*}(\mu,\nu(\alpha^*)) - C_{\kappa}^{\pi^*}(\mu,\nu) \leq 0$, where the inequality is strict for $C_{\kappa}^{\pi^*}(\mu,\nu) > 1$. Finally, since for $\mu=\nu$ it holds that $\nu(\alpha^*) = (1-\alpha^*)\nu + \alpha^*\nu = \nu$ for, we get that $C_{\kappa}^{\pi^*}(\nu,\nu)$ is a decreasing function of κ .

E Proof of Theorem 5

We first prove two technical lemmas.

Lemma 10. Let π be a policy, $\kappa \in [0,1]$, $\gamma \in (0,1)$ and $i \in \mathbb{N} \setminus \{0\}$. Then

$$(\xi D_{\kappa}^{\pi} P^{\pi})^{i} = \sum_{t=i-1}^{\infty} \frac{t!}{(i-1)!(t-(i-1))!} \gamma^{t+1} (1-\kappa)^{i} \kappa^{t-(i-1)} (P^{\pi})^{t+1},$$

where, as also given in Definition 4, $D_{\kappa}^{\pi} = (1 - \kappa \gamma)(I - \kappa \gamma P^{\pi})^{-1}$.

Proof. First, for any $x \in \mathbb{R}$ s.t |x| < 1 and $i \in \mathbb{N} \setminus \{0\}$ we have that,

$$(1-x)^{-i} = \sum_{t=i-1}^{\infty} \frac{t!}{(i-1)!(t-(i-1))!} x^{t-(i-1)}.$$

Since it holds that $||\gamma \kappa P^{\pi}|| = \gamma \kappa < 1$, where $||\cdot||$ is the spectral norm of the matrix, we can use the same Taylor expansion when replacing x with $\gamma \kappa P^{\pi}$. Thus,

$$(I - \gamma \kappa P^{\pi})^{-i} = \sum_{t=i-1}^{\infty} \frac{t!}{(i-1)!(t-(i-1))!} (\gamma \kappa)^{t-(i-1)} (P^{\pi})^{t-(i-1)}.$$
 (23)

Since $D_{\kappa}^{\pi}=(1-\kappa\gamma)(I-\kappa\gamma P^{\pi})^{-1}$ and any matrix commutes with any function of itself we have that,

$$(\xi D_{\kappa}^{\pi} P^{\pi})^{i} = \gamma^{i} (1 - \kappa)^{i} (D_{\kappa}^{\pi} P^{\pi})^{i} = \gamma^{i} (1 - \kappa)^{i} ((I - \kappa \gamma P^{\pi})^{-1})^{i} (P^{\pi})^{i}.$$

By using (23) and packing the terms we conclude the proof.

$$\begin{split} (\xi D_{\kappa}^{\pi} P^{\pi})^{i} &= \gamma^{i} (1 - \kappa)^{i} (I - \kappa \gamma P^{\pi})^{-i} (P^{\pi})^{i} \\ &= \sum_{t=i-1}^{\infty} \frac{t!}{(i-1)!(t-(i-1))!} \gamma^{t+1} (1 - \kappa)^{i} \kappa^{t-(i-1)} (P^{\pi})^{t+1} \end{split}$$

Lemma 11. Let $\kappa \in [0,1], \ \gamma \in (0,1), n \in \mathbb{N} \cup \{\infty\}$ and $f : \mathbb{N} \to \mathbb{R}$. Then

$$\sum_{l=0}^{\infty} \sum_{i=1}^{n-1} \sum_{t=i-1}^{\infty} \frac{t!}{(i-1)!(t-(i-1))!} \gamma^{t+l+1} \kappa^{t-(i-1)} (1-\kappa)^{i} f(t+1+l)$$

$$\leq (1-\kappa) \sum_{l=0}^{\infty} \sum_{t=0}^{n-2} \gamma^{t+l+1} f(t+1+l) + g(\kappa) (1-\kappa) \kappa \sum_{l=0}^{\infty} \sum_{t=n-1}^{\infty} \gamma^{t+l+1} f(t+1+l),$$

where $g(\kappa)$ is a bounded function of κ . When $n \to \infty$ the second term vanishes.

Proof. We start by exchanging the summation indices i and t. In order to do so, we decouple the summation to two sums. The range of the indices of the first sum is $t \in \{0,...,n-2\}$ and $i \in \{1,...,t+1\}$ and the range of the indices of the second sum is $t \in \{n-1,...,\infty\}$ and $i \in \{1,...,n-1\}$

$$\sum_{l=0}^{\infty} \sum_{i=1}^{n-1} \sum_{t=i-1}^{\infty} \frac{t!}{(i-1)!(t-(i-1))!} \gamma^{t+l+1} \kappa^{t-(i-1)} (1-\kappa)^{i} f(t+1+l)
= \sum_{l=0}^{\infty} \sum_{t=0}^{n-2} \gamma^{t+l+1} f(t+1+l) \sum_{i=1}^{t+1} \frac{t!}{(i-1)!(t-(i-1))!} \kappa^{t-(i-1)} (1-\kappa)^{i}
+ \sum_{l=0}^{\infty} \sum_{t=n-1}^{\infty} \gamma^{t+l+1} f(t+1+l) \sum_{i=1}^{n-1} \frac{t!}{(i-1)!(t-(i-1))!} \kappa^{t-(i-1)} (1-\kappa)^{i}.$$
(24)

Let us bound the first sum first (24),

$$\sum_{l=0}^{\infty} \sum_{t=0}^{n-2} \gamma^{t+l+1} f(t+1+l) \sum_{i=1}^{t+1} \frac{t!}{(i-1)!(t-(i-1))!} \kappa^{t-(i-1)} (1-\kappa)^{i}$$

$$= \sum_{l=0}^{\infty} \sum_{t=0}^{n-2} \gamma^{t+l+1} f(t+1+l) \sum_{i=0}^{t} \frac{t!}{i!(t-i)!} \kappa^{t-i} (1-\kappa)^{i+1}$$

$$= (1-\kappa) \sum_{l=0}^{\infty} \sum_{t=0}^{n-2} \gamma^{t+l+1} f(t+1+l),$$

where in the first line we changed the index summation $i \leftarrow i-1$ and in the second line we used the binomial identity $\sum_{i=0}^{t} \frac{t!}{i!(t-i)!} \kappa^{t-i} (1-\kappa)^i = (1-\kappa+\kappa)^t = 1$.

In order to bound the second term (25) we define the following function, $\tilde{g}:[n-1,\infty)\to\mathbb{R}$,

$$\tilde{g}(t) \stackrel{\text{def}}{=} \sum_{i=0}^{n-2} \frac{t!}{i!(t-i)!} \kappa^{t-i} (1-\kappa)^i.$$

The function $\tilde{g}(t)$ is a sum of polynomial terms multiplied by a geometric decaying term, κ^t . Thus, this function is bounded from above, i.e, exists $t^* \in [n-1,\infty)$ such that $\tilde{g}(t) \leq \tilde{g}(t^*), \ \forall t \in [n-1,\infty)$.

For such t^* , by construction, we have that

$$\sum_{i=1}^{n-1} \frac{t!}{(i-1)!(t-(i-1))!} \kappa^{t-(i-1)} (1-\kappa)^i = (1-\kappa) \sum_{i=0}^{n-2} \frac{t!}{i!(t-i)!} \kappa^{t-i} (1-\kappa)^i$$

$$\leq (1-\kappa) \sum_{i=0}^{n-2} \frac{t^*!}{i!(t^*-i)!} \kappa^{t^*-i} (1-\kappa)^i$$

$$= (1-\kappa) \kappa^{t^*-(n-2)} \sum_{i=0}^{n-2} \frac{t^*!}{i!(t^*-i)!} \kappa^{(n-2)-i} (1-\kappa)^i$$

$$\leq (1-\kappa) \kappa \sum_{i=0}^{n-2} \frac{t^*!}{i!(t^*-i)!} \kappa^{(n-2)-i} (1-\kappa)^i$$

where the last line holds since for $\kappa \in [0,1]$, $t^* \in [n-1,\infty)$ it holds that $\kappa^{t^*-(n-2)} \leq \kappa$. We now define $g(\kappa) \stackrel{\text{def}}{=} \sum_{i=0}^{n-2} \frac{t^*!}{i!(t^*-i)!} \kappa^{(n-2)-i} (1-\kappa)^i$, and observe that it is a bounded function of $\kappa \in [0,1]$, since it is a sum of positive powers of κ . Thus, (25) is bounded by

$$\sum_{l=0}^{\infty} \sum_{t=n-1}^{\infty} \gamma^{t+l+1} f(t+1+l) \sum_{i=1}^{n-1} \frac{t!}{(i-1)!(t-(i-1))!} \kappa^{t-(i-1)} (1-\kappa)^{i}$$

$$\leq g(\kappa) (1-\kappa) \kappa \sum_{l=0}^{\infty} \sum_{t=n-1}^{\infty} \gamma^{t+l+1} f(t+1+l)$$

Finally, for the case $n=\infty$ observe we can repeat the same analysis we did for the first term (24) without the need to decouple to two sums. Thus, for this case, the bound on the first term, with $n=\infty$, bounds the expression.

We are now ready to prove Theorem 5. The proof strategy is similar to the line of work in [7, 19, 11]: Keeping track of the cumulative error and using the definition of c(i) and $c^{\pi^*}(i)$, we bound the performance loss in the μ -weighted L_1 norm.

Since the policy in each iteration is an approximate κ -greedy policy (see Definition 2), it holds that $\nu T_{\kappa}^{\pi_k} v^{\pi_{k-1}} \geq \nu T_{\kappa} v^{\pi_{k-1}} - \delta$ in each iteration. Let the error vector at the i-th iteration $\bar{\delta}_i$ satisfy $\nu \bar{\delta}_i \leq \delta$. Thus,

$$v^{*} - v^{\pi_{k}} = T_{\kappa}^{\pi^{*}} v^{*} - T_{\kappa}^{\pi^{*}} v^{\pi_{k-1}} + T_{\kappa}^{\pi^{*}} v^{\pi_{k-1}} - v^{\pi_{k}}$$

$$= \xi D_{\kappa}^{\pi^{*}} P^{\pi^{*}} (v^{*} - v^{\pi_{k-1}}) + T_{\kappa}^{\pi^{*}} v^{\pi_{k-1}} - v^{\pi_{k}}$$

$$= \xi D_{\kappa}^{\pi^{*}} P^{\pi^{*}} (v^{*} - v^{\pi_{k-1}}) + T_{\kappa}^{\pi^{*}} v^{\pi_{k-1}} - T_{\kappa}^{\pi_{k}} v^{\pi_{k-1}} + T_{\kappa}^{\pi_{k}} v^{\pi_{k-1}} - v^{\pi_{k}}$$

$$\leq \xi D_{\kappa}^{\pi^{*}} P^{\pi^{*}} (v^{*} - v^{\pi_{k-1}}) + T_{\kappa}^{\pi^{*}} v^{\pi_{k-1}} - \max_{\pi'} T_{\kappa}^{\pi'} v^{\pi_{k-1}} + \bar{\delta}_{i} + T_{\kappa}^{\pi_{k}} v^{\pi_{k-1}} - v^{\pi_{k}}$$

$$\leq \xi D_{\kappa}^{\pi^{*}} P^{\pi^{*}} (v^{*} - v^{\pi_{k-1}}) + \bar{\delta}_{i} + T_{\kappa}^{\pi_{k}} v^{\pi_{k-1}} - v^{\pi_{k}}$$

$$= \xi D_{\kappa}^{\pi^{*}} P^{\pi^{*}} (v^{*} - v^{\pi_{k-1}}) + \bar{\delta}_{i} + \xi D_{\kappa}^{\pi_{\kappa}} P^{\pi_{k}} (v^{\pi_{k-1}} - v^{\pi_{k}}), \tag{26}$$

where we used in the second and last relations that for any policy π , and any value functions $v_1, v_2, T_\kappa^\pi v_1 - T_\kappa^\pi v_2 = \xi D_\kappa^\pi P^\pi (v_1 - v_2)$. This can be seen by using the definition of T_κ^π (see Section 3). Notice that

$$\begin{split} v^{\pi_{k-1}} - v^{\pi_k} &= T_{\kappa}^{\pi_{k-1}} v^{\pi_{k-1}} - v^{\pi_k} \\ &\leq \max_{\pi'} T_{\kappa}^{\pi'} v^{\pi_{k-1}} - v^{\pi_k} \\ &\leq T_{\kappa}^{\pi_k} v^{\pi_{k-1}} - v^{\pi_k} + \bar{\delta}_i \\ &= T_{\kappa}^{\pi_k} v^{\pi_{k-1}} - T_{\kappa}^{\pi_k} v^{\pi_k} + \bar{\delta}_i \\ &= \xi D_{\kappa}^{\pi_k} P^{\pi_k} (v^{\pi_{k-1}} - v^{\pi_k}) + \bar{\delta}_i. \end{split}$$

Hence,

$$(I - \xi D_{\kappa}^{\pi_k} P^{\pi_k})(v^{\pi_{k-1}} - v^{\pi_k}) \le \bar{\delta}_i, \text{ i.e.,}$$

$$v^{\pi_{k-1}} - v^{\pi_k} \le (I - \xi D_{\kappa}^{\pi_k} P^{\pi_k})^{-1} \bar{\delta}_i.$$
(27)

The last equation holds due to [14, Lemma 4.2], combined with the fact that $(I - \xi D_{\kappa}^{\pi_k} P^{\pi_k})^{-1} = \sum_{i=0}^{\infty} \xi D_{\kappa}^{\pi_k} P^{\pi_k} \geq 0$, element-wise.

Plugging (27) into (26), we have that

$$v^* - v^{\pi_k} \le \xi D_{\kappa}^{\pi^*} P^{\pi^*} (v^* - v^{\pi_{k-1}}) + \bar{\delta}_i + \xi D_{\kappa}^{\pi_k} P^{\pi_k} (I - \xi D_{\kappa}^{\pi_k} P^{\pi_k})^{-1} \bar{\delta}_i$$
$$= \xi D_{\kappa}^{\pi^*} P^{\pi^*} (v^* - v^{\pi_{k-1}}) + (I - \xi D_{\kappa}^{\pi_k} P^{\pi_k})^{-1} \bar{\delta}_i,$$

where the second relation holds since for matrix X s.t. ||X|| < 1, $I + X(I - X)^{-1} = (I - X)^{-1}$. We thus get that the errors accumulate as follows.

$$v^* - v^{\pi_k} \le \sum_{i=0}^{k-1} (\xi D_{\kappa}^{\pi^*} P^{\pi^*})^i (I - \xi D_{\kappa}^{\pi_{k-i}} P^{\pi_{k-i}})^{-1} \bar{\delta}_i + (\xi D_{\kappa}^{\pi^*} P^{\pi^*})^k (v^* - v^{\pi_0}).$$

We continue by multiplying both sides with μ and get

$$\mu(v^* - v^{\pi_k}) \le \sum_{i=0}^{k-1} \mu(\xi D_{\kappa}^{\pi^*} P^{\pi^*})^i (I - \xi D_{\kappa}^{\pi_{k-i}} P^{\pi_{k-i}})^{-1} \bar{\delta}_i + \xi^k \frac{R_{\max}}{1 - \gamma}.$$
 (28)

Using Lemma 9 with $\kappa = 0$ and renaming κ' to κ , we have that

$$(I - \xi D_{\kappa}^{\pi_{k-i}} P^{\pi_{k-i}})^{-1} = (1 - \kappa)(I - \gamma P^{\pi_{k-i}})^{-1} + \kappa I.$$

Plugging this relation into (28) gives

$$\mu(v^* - v^{\pi_k}) \leq \sum_{i=0}^{k-1} \mu(\xi D_{\kappa}^{\pi^*} P^{\pi^*})^i ((1 - \kappa)(I - \gamma P^{\pi_{k-i}})^{-1} + \kappa I) \bar{\delta}_i + \xi^k \frac{R_{\max}}{1 - \gamma}$$

$$\leq (1 - \kappa) \sum_{i=0}^{k-1} \mu(\xi D_{\kappa}^{\pi^*} P^{\pi^*})^i (I - \gamma P^{\pi_{k-i}})^{-1} \bar{\delta}_i + \kappa \sum_{i=0}^{k-1} \mu(\xi D_{\kappa}^{\pi^*} P^{\pi^*})^i \bar{\delta}_i + \xi^k \frac{R_{\max}}{1 - \gamma}.$$
(29)

The following two lemmas provide bounds for the first two terms above. The bounds are composed of the concentrability coefficients (see Definition 3 and Definition 4).

Lemma 12. Let $\kappa \in [0,1]$. For any sequence of policies $\{\pi_{k-i}\}_{i=0}^{k-1}$, optimal policy π^* , and error vector which satisfy $\nu \bar{\delta}_i \leq \delta$,

$$\sum_{i=0}^{k-1} \mu(\xi D_{\kappa}^{\pi^*} P^{\pi^*})^i (I - \gamma P^{\pi_{k-i}})^{-1} \bar{\delta}_i \le \left(\frac{(1-\kappa)C^{(2)}(\mu,\nu)}{(1-\gamma)^2} + \frac{\kappa C^{(1)}(\mu,\nu)}{1-\gamma} \right) \delta \tag{30}$$

and

$$\sum_{i=0}^{k-1} \mu(\xi D_{\kappa}^{\pi^*} P^{\pi^*})^{i} (I - \gamma P^{\pi_{k-i}})^{-1} \overline{\delta}_{i} \\
\leq \left(k \frac{(1 - \kappa) C^{(1)}(\mu, \nu)}{1 - \gamma} + \frac{\kappa C^{(1)}(\mu, \nu)}{1 - \gamma} + \frac{g(\kappa)(1 - \kappa) \kappa \gamma^{k} C^{(2, k)}(\mu, \nu)}{(1 - \gamma)^{2}} \right) \delta. \tag{31}$$

Proof. We start with proving (30). Let π' be an arbitrary policy. For i > k-1, we define $\pi_{k-i} = \pi'$ and vectors $\bar{\delta}_i$ s.t. $\nu \bar{\delta}_i \leq \delta$.

$$\sum_{i=0}^{k-1} \mu(\xi D_{\kappa}^{\pi^*} P^{\pi^*})^{i} (I - \gamma P^{\pi_{k-i}})^{-1} \bar{\delta}_{i} \leq \sum_{i=0}^{\infty} \mu(\xi D_{\kappa}^{\pi^*} P^{\pi^*})^{i} (I - \gamma P^{\pi_{k-i}})^{-1} \bar{\delta}_{i}$$

$$= \mu(I - \gamma P^{\pi_{k}})^{-1} \bar{\delta}_{0} + \sum_{i=1}^{\infty} \mu(\xi D_{\kappa}^{\pi^*} P^{\pi^*})^{i} (I - \gamma P^{\pi_{k-i}})^{-1} \bar{\delta}_{i}.$$
(32)

For the first term in (32) we have that

$$\mu(I - \gamma P^{\pi_k})^{-1}\bar{\delta}_0 = \sum_{l=0}^{\infty} \gamma^l \mu(P^{\pi_k})^l \bar{\delta}_0 \le \sum_{l=0}^{\infty} \gamma^l c(l) \nu \bar{\delta}_0 \le \sum_{l=0}^{\infty} \gamma^l c(l) \delta = \frac{C^{(1)}(\mu, \nu)}{1 - \gamma} \delta, \quad (33)$$

where for the second relation we used the definition of the sequence $\{c(i)\}_{i=0}^{\infty}$ (see Definition 3) and in the third relation we used $\nu \bar{\delta}_0 \leq \delta$ (see Definition 2).

Next, we bound the second term in (32).

$$\sum_{i=1}^{\infty} \mu(\xi D_{\kappa}^{\pi^{*}} P^{\pi^{*}})^{i} (I - \gamma P^{\pi_{k-i}})^{-1} \bar{\delta}_{i}
= \sum_{l=0}^{\infty} \sum_{i=1}^{\infty} \gamma^{l} \mu(\xi D_{\kappa}^{\pi^{*}} P^{\pi^{*}})^{i} (P^{\pi_{k-i}})^{l} \bar{\delta}_{i}
= \sum_{l=0}^{\infty} \sum_{i=1}^{\infty} \sum_{t=i-1}^{\infty} \frac{t!}{(i-1)!(t-(i-1))!} \gamma^{l+t+1} \kappa^{t-(i-1)} (1-\kappa)^{i} \mu(P^{\pi^{*}})^{t+1} (P^{\pi_{k-i}})^{l} \bar{\delta}_{i}
\leq \sum_{l=0}^{\infty} \sum_{i=1}^{\infty} \sum_{t=i-1}^{\infty} \frac{t!}{(i-1)!(t-(i-1))!} \gamma^{l+t+1} \kappa^{t-(i-1)} (1-\kappa)^{i} c(t+1+l) \delta
\leq (1-\kappa) \sum_{l=0}^{\infty} \sum_{t=0}^{\infty} \gamma^{l+t+1} c(t+1+l) \delta
= (1-\kappa) \left(\sum_{l=0}^{\infty} \sum_{t=1}^{\infty} \gamma^{l+t} c(t+l) + \sum_{l=0}^{\infty} \gamma^{l} c(l) - \sum_{l=0}^{\infty} \gamma^{l} c(l) \right) \delta
= (1-\kappa) \left(\sum_{l=0}^{\infty} \sum_{t=0}^{\infty} \gamma^{l+t} c(t+l) - \sum_{l=0}^{\infty} \gamma^{l} c(l) \right) \delta = (1-\kappa) \left(\frac{C^{(2)}(\mu,\nu)}{(1-\gamma)^{2}} - \frac{C^{(1)}(\mu,\nu)}{1-\gamma} \right) \delta. \quad (36)$$

For the first relation we used the Taylor expansion $(I-\gamma P^{\pi_{k-i}})^{-1}=\sum_{l=0}^{\infty}\gamma^l\left(P^{\pi_{k-i}}\right)^l$, for the second we used Lemma 10, for the third we used the definition of the sequence $\{c(i)\}_{i=0}^{\infty}$ and $\nu\bar{\delta}_i\leq \delta$, for the fourth we applied Lemma 11 with $n=\infty$ and $f(\cdot)=c(\cdot)$, and for the fifth we shifted the summation index $t\leftarrow t+1$.

We bound (32) by summing the bounds in (33) and (36) to obtain the first statement of the lemma, (30).

To prove the second statement, (31), we again split expression of interest, similarly to (32).

$$\sum_{i=0}^{k-1} \mu(\xi D_{\kappa}^{\pi^*} P^{\pi^*})^i (I - \gamma P^{\pi_{k-i}})^{-1} \bar{\delta}_i \le \mu(I - \gamma P^{\pi_k})^{-1} \bar{\delta}_0 + \sum_{i=1}^{k-1} \mu(\xi D_{\kappa}^{\pi^*} P^{\pi^*})^i (I - \gamma P^{\pi_{k-i}})^{-1} \bar{\delta}_i.$$
(37)

As in (33), the first term in (37) is bounded by

$$\mu(I - \gamma P^{\pi_k})^{-1} \bar{\delta}_0 \le \frac{C^{(1)}(\mu, \nu)}{1 - \gamma} \delta.$$
 (38)

Next, we bound the second term in (37).

$$\sum_{i=1}^{k-1} \mu(\xi D_{\kappa}^{\pi^*} P^{\pi^*})^{i} (I - \gamma P^{\pi_{k-i}})^{-1} \bar{\delta}_{i}$$

$$= \sum_{l=0}^{\infty} \sum_{i=1}^{k-1} \gamma^{l} \mu(\xi D_{\kappa}^{\pi^*} P^{\pi^*})^{i} (P^{\pi_{k-i}})^{l} \bar{\delta}_{i}$$

$$\leq (1 - \kappa) \sum_{l=0}^{\infty} \sum_{t=0}^{k-2} \gamma^{t+1+l} c(t+1+l) \delta + g(\kappa) (1-\kappa) \kappa \sum_{l=0}^{\infty} \sum_{t=k-1}^{\infty} \gamma^{t+1+l} c(t+1+l) \delta$$

$$= (1 - \kappa) \sum_{t=0}^{k-2} \sum_{l=0}^{\infty} \gamma^{t+1+l} c(t+1+l) \delta + g(\kappa) (1-\kappa) \kappa \gamma^{k} \sum_{l=0}^{\infty} \sum_{t=0}^{\infty} \gamma^{t+l} c(t+l+k) \delta$$

$$\leq (k-1) \frac{(1-\kappa) C^{(1)}(\mu,\nu)}{1-\gamma} \delta + \frac{g(\kappa) (1-\kappa) \kappa \gamma^{k} C^{(2,k)}(\mu,\nu)}{(1-\gamma)^{2}} \delta. \tag{39}$$

In the first relation we used the Taylor expansion of $(I - \gamma P^{\pi_{k-i}})$. For the second relation we perform the same steps as from (34) to (35), where this time we used Lemma 11 with finite n = k.

Summing the terms in (38) and (39), we obtain the second statement of the lemma, (31).

Lemma 13. Let $\kappa \in [0, 1]$. For any sequence of policies $\{\pi_{k-i}\}_{i=0}^{k-1}$, optimal policy π^* , and error vectors which satisfy $\nu \bar{\delta}_i \leq \delta$,

$$\sum_{i=0}^{k-1} \mu(\xi D_{\kappa}^{\pi^*} P^{\pi^*})^i \bar{\delta}_i \le \frac{1 - \kappa \gamma}{1 - \gamma} C_{\kappa}^{\pi^*(1)}(\mu, \nu) \delta \tag{40}$$

and

$$\sum_{i=0}^{k-1} \mu(\xi D_{\kappa}^{\pi^*} P^{\pi^*})^i \bar{\delta}_i \le k \frac{1 - \kappa \gamma}{1 - \gamma} C_{\kappa}^{\pi^*} (\mu, \nu) \delta. \tag{41}$$

Proof. We begin proving the first statement. For i > k - 1, we define vectors $\bar{\delta}_i$ s.t. $\nu \bar{\delta}_i \leq \delta$. Thus,

$$\sum_{i=0}^{k-1} \mu(\xi D_{\kappa}^{\pi^*} P^{\pi^*})^i \bar{\delta}_i \le \mu \bar{\delta}_0 + \sum_{i=1}^{\infty} \mu(\xi D_{\kappa}^{\pi^*} P^{\pi^*})^i \bar{\delta}_i. \tag{42}$$

For the first term in (42),

$$\mu \bar{\delta}_0 \le c(0)\nu \bar{\delta}_0 \le c(0)\delta,\tag{43}$$

where we used Definition 3 and then Definition 2.

For the second term in (42), we have

$$\sum_{i=1}^{\infty} \mu(\xi D_{\kappa}^{\pi^*} P^{\pi^*})^{i} \bar{\delta}_{i}$$

$$= \sum_{i=1}^{\infty} \sum_{t=i-1}^{\infty} \frac{t!}{(i-1)!(t-(i-1))!} \gamma^{t+1} (1-\kappa)^{i} \kappa^{t-(i-1)} \mu(P^{\pi^*})^{t+1} \bar{\delta}_{i}$$

$$\leq \sum_{i=1}^{\infty} \sum_{t=i-1}^{\infty} \frac{t!}{(i-1)!(t-(i-1))!} \gamma^{t+1} (1-\kappa)^{i} \kappa^{t-(i-1)} c^{\pi^*} (t+1) \delta$$

$$\leq (1-\kappa) \sum_{t=0}^{\infty} \gamma^{t+1} c^{\pi^*} (t+1) \delta$$

$$= (1-\kappa) \sum_{t=0}^{\infty} \gamma^{t} c^{\pi^*} (t) \delta - (1-\kappa) c(0) \delta = \frac{(1-\kappa) C^{\pi^*} (1)}{1-\gamma} (\mu, \nu)}{1-\gamma} \delta - (1-\kappa) c(0) \delta. \tag{44}$$

For the first relation we apply Lemma 10, for the second we use the definition of $\{c^{\pi^*}(i)\}_{i=0}^{\infty}$ and use $\nu \bar{\delta}_i \leq \delta$. For the third relation we apply Lemma 11 with $n=\infty$, $f(\cdot)=c^{\pi^*}(\cdot)$ and drop the l summation.

Summing the terms in (43) and (44), we get

$$\sum_{i=0}^{k-1} \mu(\xi D_{\kappa}^{\pi^*} P^{\pi^*})^i \bar{\delta}_i \le \frac{1}{1-\gamma} \left((1-\kappa) C^{\pi^*(1)}(\mu,\nu) + (1-\gamma)\kappa c(0) \right) \delta = \frac{1-\kappa\gamma}{1-\gamma} C_{\kappa}^{\pi^*(1)}(\mu,\nu) \delta,$$

where we identify $C_{\kappa}^{\pi^*(1)}(\mu,\nu)$ to be the same expression as in Definition 4.

For the second statement of the lemma, (41), we use the identity $(\xi D_{\kappa}^{\pi^*} P^{\pi^*})^i \leq (I - \xi D_{\kappa}^{\pi^*} P^{\pi^*})^{-1}$:

$$\begin{split} \sum_{i=0}^{k-1} \mu(\xi D_{\kappa}^{\pi^*} P^{\pi^*})^i \bar{\delta}_i &\leq \sum_{i=0}^{k-1} \mu(I - \xi D_{\kappa}^{\pi^*} P^{\pi^*})^{-1} \bar{\delta}_i \\ &\leq \sum_{i=0}^{k-1} \frac{C_{\kappa}^{\pi^*}(\mu, \nu)}{1 - \xi} \nu \bar{\delta}_i \leq k \frac{C_{\kappa}^{\pi^*}(\mu, \nu)}{1 - \xi} \delta = k \frac{1 - \kappa \gamma}{1 - \gamma} C_{\kappa}^{\pi^*}(\mu, \nu) \delta, \end{split}$$

where the second relation holds due to the definition of $C_{\kappa}^{\pi^*}(\mu,\nu)$.

So far, the proof went as follows. First, we expressed the cumulative error in (29) as the sum of three terms. Bounding the first and second terms is done with Lemmas 12 and 13, respectively. Each of those two lemmas gives bounds of two forms. These two forms correspond to the two statements in Theorem 5. We now apply the bounds so as to obtain the first statement. Specifically, plugging (30) and (40) into (29) gives the first statement in Theorem 5.

To obtain the second statement of Theorem 5, we apply the second form of the bounds in Lemmas 12 and 13. Specifically, we plug (31) and (41) into (29). This gives

$$\begin{aligned}
&\mu(v^* - v^{n_k}) \\
&\leq \left(k \frac{\kappa C_{\kappa}^{\pi^*}(\mu, \nu)}{1 - \xi} + k \frac{(1 - \kappa)^2 C^{(1)}(\mu, \nu)}{1 - \gamma} + \frac{(1 - \kappa)\kappa C^{(1)}(\mu, \nu)}{1 - \gamma} + \frac{g(\kappa)(1 - \kappa)^2 \kappa \gamma^k C^{(2, k)}(\mu, \nu)}{(1 - \gamma)^2} \right) \delta \\
&+ \xi^k \frac{R_{\max}}{1 - \gamma}.
\end{aligned}$$

We now carefully choose the iteration number k to make the last term smaller than δ :

$$k^* = \left\lceil \frac{\log \frac{R_{max}}{\delta(1-\gamma)}}{1-\xi} \right\rceil = \left\lceil \frac{(1-\kappa\gamma)\log \frac{R_{max}}{\delta(1-\gamma)}}{1-\gamma} \right\rceil. \tag{45}$$

By doing so we see that $\xi^{k^*} \frac{R_{\max}}{1-\gamma} < \delta$ and obtain the second statement of the result.

F Proof of Theorem 6

Here, we merely follow the arguments of [19, Appendix A], while using the operators T_{κ}^{π} instead of T^{π} and the approximate operator defined in Definition 2. As in Section E, we define the componentwise error at the *i*-th iteration, $\bar{\delta}_i$, which satisfies $\nu \bar{\delta}_i \leq \delta$. We have that for all k,

$$v^* - v^{\sigma_{\kappa,k}} = T_{\kappa}^{\pi^*} v^* - T_{\kappa}^{\pi^*} v^{\sigma_{k-1}} + T_{\kappa}^{\pi^*} v^{\sigma_{k-1}} - T_{\kappa}^{\pi_k} v^{\sigma_{k-1}}$$
$$\leq \xi D_{\kappa}^{\pi^*} P^{\pi^*} (v^* - v^{\sigma_{k-1}}) + \bar{\delta}_k.$$

Thus, by induction on k, we obtain:

$$v^* - v^{\sigma_{\kappa,k}} \le \sum_{i=0}^{k-1} (\xi D_{\kappa}^{\pi^*} P^{\pi^*})^i \bar{\delta}_i + (\xi D_{\kappa}^{\pi^*} P^{\pi^*})^k (v^* - v^{\pi_0})$$
$$\le \sum_{i=0}^{k-1} (\xi D_{\kappa}^{\pi^*} P^{\pi^*})^i \bar{\delta}_i + \xi^k \frac{R_{\max}}{1 - \gamma}$$

We can directly bound this term by applying Lemma 13. The two statements in that lemma lead to the two statements in Theorem 6. Again, for the second statement, we set k as in (45).