

---

# Logarithmic Regret for Adversarial Online Control

---

Dylan J. Foster<sup>1</sup> Max Simchowitz<sup>2</sup>

## Abstract

We introduce a new algorithm for online linear-quadratic control in a known system subject to adversarial disturbances. Existing regret bounds for this setting scale as  $\sqrt{T}$  unless strong stochastic assumptions are imposed on the disturbance process. We give the first algorithm with logarithmic regret for arbitrary adversarial disturbance sequences, provided the state and control costs are given by known quadratic functions. Our algorithm and analysis use a characterization for the optimal offline control law to reduce the online control problem to (delayed) online learning with approximate *advantage functions*. Compared to previous techniques, our approach does not need to control movement costs for the iterates, leading to logarithmic regret.

## 1. Introduction

Reinforcement learning and control consider the behavior of an agent making decisions in a dynamic environment in order to suffer minimal loss. In light of recent practical breakthroughs in data-driven approaches to continuous RL and control (Lillicrap et al., 2016; Mnih et al., 2015; Silver et al., 2017), there is great interest in applying these techniques in real-world decision making applications. However, to reliably deploy data-driven RL and control in physical systems such as self-driving cars, it is critical to develop principled algorithms with provable safety and robustness guarantees. At the same time, algorithms should not be overly pessimistic, and should be able to take advantage of benign environments whenever possible.

In this paper we develop algorithms for online linear-quadratic control which ensure robust worst-case performance while optimally adapting to the environment at hand. Linear control has traditionally been studied in settings where the dynamics of the environment are either governed

by a well-behaved stochastic process or driven by a worst-case process to which the learner must remain robust in the  $\mathcal{H}_\infty$  sense. We consider an intermediate approach introduced by Agarwal et al. (2019a) in which disturbances are non-stochastic but performance is evaluated in terms of *regret*. This benchmark forces the learner’s control policy to achieve near optimal performance on any specific disturbance process encountered.

Concretely, we consider a setting in which the state evolves according to linear dynamics:

$$x_{t+1} = Ax_t + Bu_t + w_t, \quad (1)$$

where  $x_t \in \mathbb{R}^{d_x}$  are states,  $u_t \in \mathbb{R}^{d_u}$  are inputs, and  $A \in \mathbb{R}^{d_x \times d_x}$  and  $B \in \mathbb{R}^{d_x \times d_u}$  are system matrices known to the learner. We refer to  $w_t \in \mathbb{R}^{d_x}$  as the *disturbance* (or, “noise”), which we assume is selected by an *adaptive* adversary and satisfies  $\|w_t\| \leq 1$ ; we let  $w$  refer to the entire sequence  $w_{1:T}$ . We consider fixed quadratic costs of the form  $\ell(x, u) := x^\top R_x x + u^\top R_u u$ , where  $R_x, R_u \geq 0$  are given. This model encompasses noise which is uncorrelated ( $\mathcal{H}_2$ ), worst-case ( $\mathcal{H}_\infty$ ), or governed by some non-stationary stochastic process. The model also approximates control techniques such as feedback linearization and trajectory tracking (Slotine & Li, 1991), where  $A$  and  $B$  are the result of linearizing a known nonlinear system and the disturbances arise due to systematic errors in linearization rather than from a benign noise process.

For any policy  $\pi$  that selects controls based on the current state and disturbances observed so far, we measure its performance over a time horizon  $T$  by

$$J_T(\pi; w) = \sum_{t=1}^T \ell(x_t^\pi, u_t^\pi),$$

the total cost incurred by following  $u_t = \pi_t(x_t, w_{1:t-1})$ . Letting  $\pi^K$  denote a state-feedback control law of the form  $\pi_t^K(x) = -Kx$  for all  $t$ , the learning algorithm’s goal is to minimize

$$\text{Reg}_T = J_T(\pi^{\text{alg}}; w) - \inf_{K \in \mathcal{K}} J_T(\pi^K; w),$$

where  $\pi^{\text{alg}}$  denotes the learner’s policy and  $\mathcal{K}$  is an appropriately defined set of stabilizing controllers. Thus,  $\pi^{\text{alg}}$  has low regret when its performance nearly matches the optimal

---

<sup>1</sup>Massachusetts Institute of Technology <sup>2</sup>UC Berkeley. Correspondence to: Dylan Foster <dylanf@mit.edu>.

controller  $K \in \mathcal{K}$  on the specific, realized noise sequence. While the class  $\mathcal{K}$  contains the optimal  $\mathcal{H}_\infty$  and  $\mathcal{H}_2$  control policies, we also develop algorithms to compete with a more general class of stabilizing linear controllers, which may fare better for certain noise sequences (Appendix B).

**Logarithmic regret in online control.** Agarwal et al. (2019a) introduced the adversarial LQR setting we study and provided an efficient algorithm with  $\sqrt{T}$ -regret. Subsequent works (Agarwal et al., 2019b; Simchowitz et al., 2020) have shown that logarithmic regret is possible when the disturbances follow a semi-adversarial process with persistent excitation. Our main result is to achieve logarithmic regret for fully adversarial disturbances, provided that costs are known and quadratic.

### 1.1. Contributions

We introduce Riccatitron (Algorithm 1), a new algorithm for online linear control with adversarial disturbances which attains polylogarithmic regret.

**Theorem 1 (informal).** Riccatitron attains regret  $\mathcal{O}(\log^3 T)$ , where  $\mathcal{O}$  hides factors polynomial in relevant problem parameters.

Riccatitron has comparable computational efficiency to previous methods. We show in Appendix B that the algorithm also extends to a more general benchmark class of linear controllers with internal state, and to “tracking” loss functions of the form  $\ell_t(x, u) := \ell(x - a_t, u - b_t)$ . Some conceptual contributions are as follows.

**When is logarithmic regret possible in online control?** Simchowitz & Foster (2020) and Cassel et al. (2020) independently show that logarithmic regret is impossible in a minimax sense if the system matrices  $(A, B)$  are unknown, even when disturbances are i.i.d. gaussian. Conversely, our result shows that if  $A$  and  $B$  are known, logarithmic regret is possible even when disturbances are adversarial. Together, these results paint a clear picture of when logarithmic regret is achievable in online linear control. We note, however, that our approach heavily leverages the structure of linear control with *strongly convex, quadratic* costs. We refer the to the related work section for discussion of further structural assumptions that facilitate logarithmic regret.

**Addressing trajectory mismatch.** Riccatitron represents a new approach to a problem we call *trajectory mismatch* that arises when considering policy regret in online learning problems with state. In dynamic environments, different policies inevitably visit different state trajectories. Low-regret algorithms must address the mismatch between the performance of the learner’s policy  $\pi^{\text{alg}}$  on its own realized trajectory and the performance of each benchmark policy  $\pi$

on the alternative trajectories it induces. Most algorithms with policy regret guarantees (Even-Dar et al., 2009; Zimin & Neu, 2013; Abbasi-Yadkori et al., 2013; Arora et al., 2012; Anava et al., 2015; Abbasi-Yadkori et al., 2014; Cohen et al., 2018; Agarwal et al., 2019a;b; Simchowitz et al., 2020) adopt an approach to addressing this trajectory mismatch that we refer to as “online learning with stationary costs”, or OLwS. At each round  $t$ , the learner’s adaptive policy  $\pi^{\text{alg}}$  commits to a policy  $\pi^{(t)}$ , typically from a benchmark class  $\Pi$ . The goal is to ensure that the iterates  $\pi^{(t)}$  attain low regret on a proxy sequence of *stationary* cost functions  $\pi \mapsto \lambda_t(\pi)$  that describe the loss the learner would suffer at stage  $t$  under the *fictional* trajectory that would arise if she had played the policy  $\pi$  at all stages up to time  $t$  (or in some cases, on the corresponding steady-state trajectory as  $t \rightarrow \infty$ ). Since the stationary cost does not depend on the learner’s state, low regret on the sequence  $\{\lambda_t\}$  can be obtained by feeding these losses directly into a standard online learning algorithm. To relate regret on the proxy sequence back to regret on the true sequence, most approaches use that the iterates produced by the online learner are sufficiently slow-moving.

The main technical challenge Riccatitron overcomes is that for the stationary costs that arise in our setting, no known algorithm produces iterates which move sufficiently slowly to yield logarithmic regret via OLwS (Appendix C.4). We adopt a new approach for online control we call *online learning with advantages*, or OLwA, which abandons stationary costs, and instead considers the control-theoretic advantages of actions relative to the unconstrained offline optimal policy  $\pi^*$ . Somewhat miraculously, we find that these advantages remove the explicit dependence on the learner’s state, thereby eliminating the issue of trajectory mismatch described above. In particular, unlike OLwS, we *do not* need to verify that the iterates produced by our algorithm change slowly.

### 1.2. Our approach: Online learning with advantages

In this section we sketch the *online learning with advantages* (OLwA) technique underlying Riccatitron. Let  $\pi^*$  denote the optimal unconstrained policy given knowledge of the entire disturbance sequence  $\mathbf{w}$ , and let  $\mathbf{Q}_t^*(x, u; \mathbf{w})$  be the associated Q-function (this quantity is formally defined in Definition 3). The *advantage*<sup>1</sup> with respect to  $\pi^*$ ,  $\mathbf{A}_t^*(u; x, \mathbf{w}) := \mathbf{Q}_t^*(x, u; \mathbf{w}) - \mathbf{Q}_t^*(w, u, \pi^*(x); \mathbf{w})$ , describes the difference between the total cost accumulated by selecting action  $u$  in state  $x$  at time  $t$  and subsequently playing according to the optimal policy  $\pi^*$ , versus choosing  $u_t = \pi_t^*(x; \mathbf{w})$  as well. By the well-known performance dif-

<sup>1</sup>Since we use losses rather than rewards, “advantage” refers to the advantage of  $\pi^*$  over  $u$  rather than the advantage of  $u$  over  $\pi^*$ ; the latter terminology is more common in reinforcement learning.

ference lemma (Kakade, 2003), the relative cost of a policy is *equal* the sum of the advantages under the states visited by said policy:<sup>2</sup>

$$J_T(\pi; \mathbf{w}) - J_T(\pi^*; \mathbf{w}) = \sum_{t=1}^T \mathbf{A}_t^*(u_t^\pi; x_t^\pi, \mathbf{w}). \quad (2)$$

With this observation, the regret  $\text{Reg}_T(\pi^{\text{alg}}; \mathbf{w}, \Pi)$  of any algorithm  $\pi^{\text{alg}}$  to a policy class  $\Pi$  can be expressed as:

$$\sum_{t=1}^T \mathbf{A}_t^*(u_t^{\text{alg}}; x_t^{\text{alg}}, \mathbf{w}) - \inf_{\pi \in \Pi} \sum_{t=1}^T \mathbf{A}_t^*(u_t^\pi; x_t^\pi, \mathbf{w}). \quad (3)$$

The expression (3) suggests that a reasonable approach might be to run an online learner on the functions  $\pi \mapsto \mathbf{A}_t^*(u_t^\pi; x_t^\pi, \mathbf{w})$ . However, there are two issues. First, the advantages in the first sum are evaluated on the states  $x_t^{\text{alg}}$  under  $\pi^{\text{alg}}$ , and in the second sum under the comparator trajectories  $x_t^\pi$  (trajectory mismatch). Second, like  $\pi^*$  itself, the advantages require knowledge of all future disturbances, which are not yet known to the learner at time  $t$ . We show that if the control policies are parametrized using a particular optimal control law, the advantages do not depend on the state, and can be approximated using only finite lookahead.

**Theorem 2 (informal).** *For control policies  $\pi$  with a suitable parametrization, the mapping  $\pi \mapsto \mathbf{A}_t^*(u_t^\pi; x_t^\pi, \mathbf{w})$  can be arbitrarily well approximated by a function  $\pi \mapsto \widehat{\mathbf{A}}_{t,h}(\pi; w_{1:t+h})$  which (1) does not depend on the state, (2) can be determined by the learner at time  $t+h$ , and (3) has a simple quadratic structure.*

The ‘‘magic’’ behind this theorem is that the functional dependence of the unconstrained optimal policy  $\pi^*(x; \mathbf{w})$  on the state  $x$  is linear, and does not depend  $\mathbf{w}$  (Theorem 3). As a consequence, the state-dependent portion of  $\pi^*$  can be built into the controller parametrization, leaving only the  $\mathbf{w}$ -dependent portion up to the online learner. In light of this result, we use online learning to ensure low regret on the sequence of loss functions  $f_t(\pi) := \widehat{\mathbf{A}}_{t,h}(\pi; w_{1:t+h})$ ; we address the fact that  $f_t$  is only revealed to the learner after a delay of  $h$  steps via a standard reduction (Joulani et al., 2013). We then show that for an appropriate controller parameterization  $f_t(\pi)$  is exp-concave with respect to the learner’s policy and hence second-order online learning algorithms attain logarithmic regret (Hazan et al., 2007).

We refer the reader to Appendix C for an in-depth overview of the OLwS framework, its relationship to OLwA, and challenges associated with using these techniques to achieve logarithmic regret.

<sup>2</sup>See Lemma D.12 in Appendix E for a general statement of the performance difference lemma. The invocation of the performance difference lemma here is slightly different from other results on online learning in MDPs such as Even-Dar et al. (2009), in that the role of  $\pi$  and  $\pi^*$  is swapped.

### 1.3. Preliminaries

We consider the linear control setting in (1). For normalization, we assume  $\|w_t\| \leq 1 \forall t$ . We also assume  $x_1 = 0$ .

**Policies and trajectories.** We consider policies  $\pi$  parameterized as functions of  $x_t$  and  $\mathbf{w}$  via  $u_t = \pi_t(x_t; \mathbf{w})$ . We assume that, when selecting action  $u_t$  at time  $t$ , the learner has access to all states  $x_{1:t}, u_{1:t-1}$ , as well as  $w_{1:t-1}$  (the latter assumption is without loss of generality by the identity  $w_s = x_{s+1} - Ax_s - Bu_s$ ). Thus, a policy is said to be *executable* if  $\pi_t(x; \mathbf{w})$  depends only on  $x$  and  $w_{1:t-1}$ , i.e.  $\pi(x; \mathbf{w}) = \pi(w; w_{1:t-1})$ . For analysis purposes, we also consider *non-executable* whose value at time  $t$  may depend on the entire sequence  $\mathbf{w}$ . For a policy  $\pi$  and sequence  $\mathbf{w}$ , we let  $x_t^\pi(\mathbf{w}), u_t^\pi(\mathbf{w})$  denote the resulting states and input trajectories (which we note depend only on  $w_{1:t-1}$ ). For simplicity, we often write  $x_t^\pi$  and  $u_t^\pi$ , suppressing the  $\mathbf{w}$ -dependence. We shall let  $\pi^{\text{alg}}$  refer to the policy selected by the learner’s algorithm, and use the shorthand  $x_t^{\text{alg}}(\mathbf{w}), u_t^{\text{alg}}(\mathbf{w})$  to denote the corresponding trajectories. Given a class of policies  $\Pi$ , the regret of the policy  $\pi^{\text{alg}}$  is given by

$$\text{Reg}_T(\pi^{\text{alg}}; \Pi, \mathbf{w}) = J_T(\pi^{\text{alg}}; \mathbf{w}) - \inf_{\pi \in \Pi} J_T(\pi; \mathbf{w}).$$

We consider a benchmark class of policies induced by state feedback control laws  $\pi_t^K(x) = -Kx$ , indexed by matrices  $K \in \mathbb{R}^{d_u \times d_x}$ .

**Linear control theory.** We say that a linear controller  $K \in \mathbb{R}^{d_u \times d_x}$  is *stabilizing* if  $A - BK$  is *stable*, that is  $\rho(A - BK) < 1$  where  $\rho(\cdot)$  denotes the spectral radius.<sup>3</sup> We assume the system  $(A, B)$  is *stabilizable* in the sense that there exists a stabilizing controller  $K$ . For any stabilizable system, there is a unique positive semidefinite solution  $P_\infty \geq 0$  to the *discrete algebraic Riccati equation* (henceforth, DARE),

$$P = A^\top P A + R_x - A^\top P B (R_u + B^\top P B)^{-1} B^\top P A. \quad (4)$$

The solution  $P_\infty$  to (4) is an intrinsic property of the system (1) with  $(A, B)$  and characterizes the optimal infinite-horizon cost for control in the absence of noise (Bertsekas, 2005). Our algorithms and analysis make use of this parameter, as well as the corresponding optimal state feedback controller  $K_\infty := (R_u + B^\top P_\infty B)^{-1} B^\top P_\infty A$ . We also use the steady-state covariance matrix  $\Sigma_\infty := R_u + B^\top P_\infty B$  and closed-loop dynamics matrix  $A_{\text{cl}, \infty} := A - BK_\infty$ .

**Competing with state feedback.** While  $K_\infty$  represents the (asymptotically) optimal control law in the presence of uncorrelated, unbiased stochastic noise,  $\pi^{K_\infty}$  may not be

<sup>3</sup>For a possibly asymmetric matrix  $A$ ,  $\rho(A) = \max\{|\lambda| \mid \lambda \text{ is an eigenvalue for } A\}$ .

the optimal state feedback policy in hindsight for a given sequence of adversarial perturbations  $w_t$ . Hence, we compete with linear controllers that satisfy a quantitative version of the stability property.

**Definition 1** (Strong Stability (Cohen et al., 2018)). *We say that  $A - BK \in \mathbb{R}^{d_x \times d_x}$  is  $(\kappa, \gamma)$ -strongly stable if there exists matrices  $H, L \in \mathbb{R}^{d_x \times d_x}$  such that  $A - BK = HLH^{-1}$ ,  $\|H\|_{\text{op}}\|H\|_{\text{op}}^{-1} \leq \kappa$  and  $\|L\|_{\text{op}} \leq \gamma$ .*

Given parameters  $(\kappa_0, \gamma_0)$ , we consider the benchmark class

$$\mathcal{K}_0 = \{ \|K\|_{\text{op}} \leq \kappa_0 : A - BK \text{ is } (\kappa_0, \gamma_0)\text{-strongly stable} \}.$$

**Lemma D.1** (Appendix D.1) shows that the closed-loop dynamics for  $K_\infty$  are always  $(\kappa_\infty, \gamma_\infty)$ -strongly stable for suitable  $\gamma_\infty, \kappa_\infty$ . We assume that  $\mathcal{K}_0$  is chosen such that  $\kappa_\infty \leq \kappa_0$  and  $\gamma_\infty \leq \gamma_0$ .<sup>4</sup> Our algorithms minimize policy regret to the class of induced policies for  $\mathcal{K}_0$ :

$$\mathcal{K}_0\text{-Reg}_T(\pi^{\text{alg}}; \mathbf{w}) := J_T(\pi^{\text{alg}}; \mathbf{w}) - \inf_{K \in \mathcal{K}_0} J_T(\pi^K; \mathbf{w}).$$

**Problem parameters.** Our regret bounds depend on the following basic parameters for the LQR problem:

$$\Psi_* := \max\{1, \|A\|_{\text{op}}, \|B\|_{\text{op}}, \|R_x\|_{\text{op}}, \|R_u\|_{\text{op}}\}, \beta_* := \max\{1, \lambda_{\min}^{-1}(R_u), \lambda_{\min}^{-1}(R_x)\}, \Gamma_* := \max\{1, \|P_\infty\|_{\text{op}}\}.$$

**Additional notation.** We adopt non-asymptotic big-oh notation: For functions  $f, g : \mathcal{X} \rightarrow \mathbb{R}_+$ , we write  $f = \mathcal{O}(g)$  if there exists some constant  $C > 0$  such that  $f(x) \leq Cg(x)$  for all  $x \in \mathcal{X}$ . We use  $\tilde{\mathcal{O}}(\cdot)$  to suppress logarithmic dependence on system parameters, and we use  $\mathcal{O}_*(\cdot)$  to suppress *all* dependence on system parameters. For a vector  $x \in \mathbb{R}^d$ , we let  $\|x\|$  denote the euclidean norm and  $\|x\|_\infty$  denote the element-wise  $\ell_\infty$  norm. For a matrix  $A$ , we let  $\|A\|_{\text{op}}$  denote the operator norm. If  $A$  is symmetric, we let  $\lambda_{\min}(A)$  denote the minimum eigenvalue. When  $P > 0$  is a positive definite matrix, we let  $\|x\|_P = \sqrt{\langle x, Px \rangle}$  denote the induced weighted euclidean norm. We let  $\mathbf{w}_{t-1} = (w_{t-1}, w_{t-2}, \dots, w_1, \mathbf{0}, \mathbf{0}, \dots)$  denote a sequence of past  $w$ s, terminating in an infinite sequence of zeros. To simplify indexing, we let  $w_s \equiv \mathbf{0}$  for  $s \leq 0$ , so that  $\mathbf{w}_{t-1} = (w_{t-1}, w_{t-2}, \dots)$ . We also let  $w_s \equiv \mathbf{0}$  for  $s > T$ .

## 1.4. Organization

Section 2 introduces the Riccatitron algorithm, states its formal regret guarantee, and gives an overview of the algorithm’s building blocks and proof techniques. Section 3 gives a high-level proof of the key “approximate advantage” theorem used by the algorithm. Omitted proofs are deferred to Appendix E and Appendix F, and additional technical tools stated and proven in Appendix D.

<sup>4</sup>This assumption only serves to keep notation compact.

Appendix A gives a detailed survey of related work. Appendix B sketches extensions of Riccatitron to more general settings, and Appendix C gives a detailed survey of challenges associated with applying previous approaches to online reinforcement learning to obtain logarithmic regret in our setting.

## 2. Logarithmic regret for online linear control

Our main algorithm, Riccatitron, is described in Algorithm 1. The algorithm combines several ideas.

1. Following Agarwal et al. (2019a), we move from linear policies of the form  $\pi^K(x; \mathbf{w}) = -Kx$ , to a relaxed set of *disturbance-action* (DAP) policies of the form  $\pi_t^{(M)}(x; \mathbf{w}) = -K_\infty x - q_t^M(\mathbf{w}_{t-1})$ , where

$$q_t^M(\mathbf{w}_{t-1}) = \sum_{i=1}^m M^{[i]} w_{t-i},$$

and where  $K_\infty$  is linear controller from the DARE (4).

2. We show that the optimal unconstrained policy with full knowledge of the sequence  $\mathbf{w}$  takes the form  $\pi_t^*(x; \mathbf{w}) = -K_t x - q_t^*(\mathbf{w}_{t:T})$ , where  $(K_t)$  is a particular sequence of linear controllers that arises from the so-called *Riccati recursion*. We then show that for *any* policy of the form  $\pi_t(x; \mathbf{w}) = -K_\infty x - q_t(\mathbf{w})$ —in particular, for the DAP parameterization above—the advantage functions  $\mathbf{A}_t^*(u_t^\pi; x_t^\pi; \mathbf{w})$  can be well approximated by simple quadratic functions of the form

$$\|q_t(\mathbf{w}) - q_t^*(\mathbf{w}_{t:T})\|_{\Sigma_\infty}^2.$$

This essentially removes the learner’s state from the equation, and reduces the problem of control to that of predicting the optimal controller’s bias vector  $q_t^*(\mathbf{w}_{t:T})$ . The remaining challenge is that the optimal bias vectors depend on the future disturbances, which are not available to the learner at time  $t$ .

3. We show that the advantages can be truncated to require only finite lookahead, thereby reducing the problem to *online learning with delayed feedback*. We then apply a reduction from delayed online learning to classical online learning (Joulani et al., 2013), which proceeds by running multiple copies of a base online learning algorithm over separate subsequences of rounds.
4. Finally—using the structure of the disturbance-action parameterization—we show that the resulting online learning problem is exp-concave. As a result, we can use a second-order online learning algorithm—either online Newton step (ONS, Hazan et al. (2007)) given in Algorithm 2, or Vovk-Azoury-Warmuth (VAW, Vovk (1998); Azoury & Warmuth (2001)) given in Algorithm 3—as our base learner to obtain logarithmic regret.



**Algorithm 1** Riccatitron

- 1: **parameters:**  
 Horizon  $h$ , DAP length  $m$ , radius  $R$ , decay factor  $\gamma$ .  
 Online Newton parameters  $\eta_{\text{ONS}}, \varepsilon_{\text{ONS}}$ ,  
 or Vovk-Azoury-Warmuth parameter  $\varepsilon_{\text{VAW}}$ .
- 2: **initialize:**  
 Let  $\mathcal{M}_0 \leftarrow \mathcal{M}(m, R, \gamma)$  (Eq. 5).  
 Instantiate base learners  $\text{BL}^{(1)}, \dots, \text{BL}^{(h+1)}$  as either  
 $\text{ONS}(\varepsilon_{\text{ONS}}, \eta_{\text{ONS}}, \mathcal{M}_0)$  or  $\text{VAW}(\varepsilon_{\text{VAW}}, \mathcal{M}_0, \Sigma_\infty)$ .
- 3: Let  $\tau_t = (t-1) \bmod (h+1) + 1 \in [h+1]$ .
- 4: **for**  $t = 1, \dots, T$ : **do**  
 // Predict using base learner  $\tau_t$ .  
- 5: Let  $M_t$  denote the  $k_t$ -th iterate produced  
 by  $\text{BL}^{(\tau_t)}$  where  $k_t \leftarrow \lfloor t/(h+1) \rfloor$ .
- 6: Play  $u_t = -K_\infty x_t - q^{M_t}(\mathbf{w}_{t-1})$  (Definition 2).
- 7: Observe  $x_{t+1}$  and  $w_t$ .  
 // Update base learner  $\tau_{t+1}$ .
- 8: **if**  $t \geq h+1$  **then**  
 // Approximate advantage from Eq. (10).  
- 9: Update  $\text{BL}^{(\tau_{t+1})}$  with  $\hat{\mathbf{A}}_{t-h:h}(M; \mathbf{w}_t)$ .

Together, these components give rise to the scheme in Algorithm 1. At time  $t$ , the algorithm plays the action  $u_t = -K_\infty x_t - q^{M_t}(\mathbf{w}_{t-1})$ , where  $M_t$  is provided by the ONS (or VAW) instance responsible for the current round. The algorithm then observes  $w_t$  and uses this to form the approximate advantage function for time  $t-h$ , where  $h$  is the lookahead distance. The advantage is then used to update the ONS/VAW instance responsible for the next round. The main regret guarantee for this approach is as follows.

**Theorem 1.** *For an appropriate choice of parameters, Riccatitron ensures*

$$\mathcal{K}_0\text{-Reg}_T \leq \mathcal{O}_*(d_{\mathbf{x}}d_{\mathbf{u}} \log^3 T),$$

where  $\mathcal{O}_*$  suppresses polynomial dependence on system parameters. Suppressing only logarithmic dependence on system parameters, the regret is at most

$$\tilde{\mathcal{O}}(d_{\mathbf{x}}d_{\mathbf{u}} \log^3 T \cdot \beta_*^{11} \Psi_*^{19} \Gamma_*^{11} \kappa_0^8 (1-\gamma_0)^{-4}).$$

In the remainder of this section we overview the algorithmic building blocks of Riccatitron and the key ideas of the proof.

### 2.1. Disturbance-action policies

Cost functionals parametrized by state feedback controllers (e.g.,  $K \mapsto J_T(\pi^K; \mathbf{w})$ ) are generally non-convex (Fazel et al., 2018). To enable the use of tools from online convex optimization, we use a convex *disturbance-action* controller parameterization introduced by Agarwal et al. (2019a).

**Algorithm 2** Online Newton Step (ONS( $\varepsilon, \eta, \mathcal{C}, \Sigma$ ))

- 1: **parameters:** Learning rate  $\eta > 0$ , regularization parameter  $\varepsilon > 0$ , convex constraint set  $\mathcal{C}$ .  
 // OCO with exp-concave costs  $f_k(z)$ , where  $z \in \mathcal{C} \subset \mathbb{R}^d$ .
- 2: **initialize:**  $d \leftarrow \dim(\mathcal{C})$ ,  $z_1 \in \mathcal{C}$ ,  $E_0 \leftarrow \varepsilon \cdot I_d$ .
- 3: **for**  $k = 1, 2, \dots$ : **do**
- 4: Play  $z_k$  and **receive** gradient  $\nabla_k := \nabla f_k(z_k)$ .
- 5:  $E_k \leftarrow E_{k-1} + \nabla_k \nabla_k^\top$ .
- 6:  $\tilde{z}_{k+1} \leftarrow z_k - \eta E_k^{-1} \nabla_k$ .
- 7:  $z_{k+1} \leftarrow \arg \min_{z \in \mathcal{C}} \|z - \tilde{z}_{k+1}\|_{E_k}^2$ .

**Definition 2** (Disturbance-action policy (DAP)). *Let  $M = (M^{[i]})_{i=1}^m$  denote a sequence of matrices  $M^{[i]} \in \mathbb{R}^{d_{\mathbf{u}} \times d_{\mathbf{x}}}$ . We define the corresponding disturbance-action policy  $\pi^{(M)}$  as  $\pi_t^{(M)}(x; \mathbf{w}) = -K_\infty x - q^M(\mathbf{w}_{t-1})$ , where  $q^M(\mathbf{w}_{t-1}) = \sum_{i=1}^m M^{[i]} w_{t-i}$ .*

We work with DAPs for which the sequence  $M$  belongs to the set

$$\mathcal{M}(m, R, \gamma) := \{M = (M^{[i]})_{i=1}^m : \|M^{[i]}\|_{\text{op}} \leq R\gamma^{i-1}\}, \quad (5)$$

where  $m$ ,  $R$ , and  $\gamma$  are algorithm parameters. We note that DAPs can be defined with general stabilizing controllers  $K \neq K_\infty$ , but the choice  $K = K_\infty$  is critical in the design and analysis of our main algorithm.

The first lemma we require is a variant of a result of Agarwal et al. (2019a), which shows that disturbance-action policies are sufficiently rich enough to approximate all state feedback laws.

**Lemma 2.1** (Expressivity of DAP). *Suppose we choose our set of disturbance-action matrices as  $\mathcal{M}_0 := \mathcal{M}(m, R_*, \gamma_0)$ , where  $m = (1-\gamma_0)^{-1} \log((1-\gamma_0)^{-1}T)$  and  $R_* = 2\beta_* \Psi_*^2 \Gamma_* \kappa_0^2$ . Then for all  $\mathbf{w}$ , we have*

$$\inf_{M \in \mathcal{M}_0} J_T(\pi^{(M)}; \mathbf{w}) \leq \inf_{K \in \mathcal{K}_0} J_T(\pi^K; \mathbf{w}) + C_{\text{apx}},$$

where  $C_{\text{apx}} \leq \mathcal{O}(\beta_*^2 \Psi_*^8 \Gamma_*^2 \kappa_0^7 (1-\gamma_0)^{-2})$ .

We refer the reader to Appendix E.2 for a proof. Going forward, we define

$$D_q = \tilde{\mathcal{O}}\left(\beta_*^{5/2} \Psi_*^3 \Gamma_*^{5/2} \kappa_0^2 (1-\gamma_0)^{-1}\right), \quad (6)$$

which serves as an upper bound on  $\|q_t^M\|$  for  $M \in \mathcal{M}_0$ , as well as other certain other bias vector sequences that arise in the subsequent analysis. In light of Lemma 2.1, the remainder of our discussion will directly bound regret with respect to DAPs:

$$\mathcal{M}_0\text{-Reg}_T(\pi; \mathbf{w}) := J_T(\pi; \mathbf{w}) - \inf_{M \in \mathcal{M}_0} J_T(\pi^{(M)}; \mathbf{w}). \quad (7)$$

We note in passing that DAPs are actually rich enough to compete with a broader class of linear control policies with internal state; this extension is addressed in [Appendix B.2](#).

## 2.2. Advantages in linear control

To proceed, we adopt the OLWA paradigm, which minimizes approximations to the *advantages* (or, differences between the Q-functions) relative to the optimal unconstrained policy  $\pi^*$  given access to the entire sequence  $\mathbf{w}$ . Recalling  $\ell(x, u) = \|x\|_{R_x}^2 + \|u\|_{R_u}^2$ , we define the optimal controller  $\pi^*$  and associated Q-functions and advantages by induction.

**Definition 3.** *The optimal Q-function and policy at time  $T$  are given by  $\mathbf{Q}_T^*(x, u; \mathbf{w}) = \ell(x, u)$ ,  $\pi_T^*(x; \mathbf{w}) = \min_u \mathbf{Q}_T^*(x, u; \mathbf{w}) = 0$ , and  $\mathbf{V}_T^*(x; \mathbf{w}) = \ell(x, 0) = \|x\|_{R_x}^2$ . For each timestep  $t < T$ , the optimal Q-function and policy are given by*

$$\begin{aligned} \mathbf{Q}_t^*(x, u; \mathbf{w}) &= \|x\|_Q^2 + \|u\|_R^2 + \mathbf{V}_{t+1}^*(Ax + Bu + w_t; \mathbf{w}), \\ \pi_t^*(x; \mathbf{w}) &= \arg \min_{u \in \mathbb{R}^{d_u}} \mathbf{Q}_t^*(x, u; \mathbf{w}), \\ \mathbf{V}_t^*(x; \mathbf{w}) &= \min_{u \in \mathbb{R}^{d_u}} \mathbf{Q}_t^*(x, u; \mathbf{w}) = \mathbf{Q}_t^*(x, \pi_t^*(x; \mathbf{w}); \mathbf{w}). \end{aligned}$$

The advantage function for the optimal policy is  $\mathbf{A}_t^*(u; x, \mathbf{w}) := \mathbf{Q}_t^*(x, u; \mathbf{w}) - \mathbf{Q}_t^*(x, \pi_t^*(x; \mathbf{w}); \mathbf{w})$ .

The advantage function  $\mathbf{A}_t^*(u; x, \mathbf{w})$  represents the total excess cost incurred by selecting a control  $u \neq \pi_t^*(x; \mathbf{w})$  at state  $x$  and time  $t$ , assuming we follow  $\pi^*$  for the remaining rounds. We have  $\mathbf{A}_t^*(u; x, \mathbf{w}) \geq 0$  since, by Bellman’s optimality condition,  $\pi_t^*(x; \mathbf{w})$  is a minimizer of  $\mathbf{Q}^*(x, u; \mathbf{w})$ .

The advantages arise in our setting through application of the performance difference lemma ([Lemma D.12](#)), which we recall states that for any policy  $\pi$ , the regret to  $\pi^*$  is equal to the sum of advantages under the trajectory induced by  $\pi$ , i.e.  $J_T(\pi; \mathbf{w}) - J_T(\pi^*; \mathbf{w}) = \sum_{t=1}^T \mathbf{A}_t^*(u_t^\pi; x_t^\pi, \mathbf{w})$ . To analyze Riccatitron, we apply this identity to obtain the regret decomposition

$$\begin{aligned} \mathcal{M}_0\text{-Reg}_T(\pi; \mathbf{w}) &= \sum_{t=1}^T \mathbf{A}_t^*(u_t^\pi; x_t^\pi, \mathbf{w}) \\ &\quad - \inf_{M \in \mathcal{M}_0} \sum_{t=1}^T \mathbf{A}_t^*(u_t^{\pi^{(M)}}; x_t^{\pi^{(M)}}, \mathbf{w}) \quad (8) \end{aligned}$$

This decomposition is *exact*, and avoids the pitfalls of the usual stationary cost-based regret decomposition associated with the classical OLWS approach (cf. [Appendix C](#)). Our goal going forward will be to treat these advantages as “losses” that can be fed into an appropriate online learning algorithm to select controls. However, this approach presents three challenges: (a) the advantages for the policy  $\pi$  are evaluated on the trajectory  $x_t^\pi$ , while the advantages

for comparator are evaluated under the trajectory induced by  $\pi^{(M)}$ ; (b) the advantage is a difference in Q-functions that considers *all* future expected reward. In particular,  $\mathbf{A}_t^*(\cdot; \cdot, \mathbf{w})$  depends on all future  $w_t$ s, including those not yet revealed to the learner; (c) the functional form of the advantages is opaque, and it is not clear that any online learning algorithm can achieve logarithmic regret even if they were able to evaluate  $\mathbf{A}_t^*$  at time  $t$ .

## 2.3. Approximate advantages

Our main structural result—and the starting point for Riccatitron—is the following observation. Let  $\pi$  be any policy of the form  $\pi_t(x; \mathbf{w}_{t-1}) = -K_\infty x - q^{M_t}(\mathbf{w}_{t-1})$ , where  $M_t = M_t(\mathbf{w}_{t-1})$  are arbitrary functions of past  $\mathbf{w}$ , and where  $K_\infty$  is the infinite horizon Riccati optimal controller. Then  $\mathbf{A}_t^*(u_t^\pi; x_t^\pi, \mathbf{w})$  is well-approximated by an *approximate advantage function*  $\widehat{\mathbf{A}}_{t;h}(M; \mathbf{w}_{t+h})$  which (a) does not depend on the state, and (b) depends on only a small horizon  $h$  of future disturbances, and (c) is a pure quadratic function of  $M$ , and thereby amenable to fast (logarithmic) rates for online learning. Let  $h$  be a horizon/lookahead parameter. Defining

$$q_{\infty;h}^*(w_{1:h+1}) := \sum_{i=1}^{h+1} \Sigma_\infty^{-1} B^\top (A_{\text{cl},\infty}^\top)^{i-1} P_\infty w_i, \quad (9)$$

the approximate advantage function is

$$\widehat{\mathbf{A}}_{t;h}(M; \mathbf{w}_{t+h}) := \|q^M(\mathbf{w}_{t-1}) - q_{\infty;h}^*(w_{t:t+h})\|_{\Sigma_\infty}^2. \quad (10)$$

The following theorem facilitates the use of the approximate advantages.

**Theorem 2.** *Let  $\pi$  be any policy of the form  $\pi_t(x; \mathbf{w}) = -K_\infty x - q^{M_t}(\mathbf{w}_{t-1})$ , where  $M_t = M_t(\mathbf{w}) \in \mathcal{M}_0$ . Then, by choosing  $h = 2(1 - \gamma_\infty)^{-1} \log(\kappa_\infty^2 \beta_\star^2 \Psi_\star \Gamma_\star^2 T^2)$  as the horizon parameter, we have*

$$\sum_{t=1}^T |\mathbf{A}_t^*(u_t^\pi; x_t^\pi, \mathbf{w}) - \widehat{\mathbf{A}}_{t;h}(M_t; \mathbf{w}_{t+h})| \leq C_{\text{adv}},$$

where  $C_{\text{adv}} = \widetilde{\mathcal{O}}(\beta_\star^{11} \Psi_\star^{19} \Gamma_\star^{11} \kappa_0^8 (1 - \gamma_0)^{-4} \log^2 T)$ .

The proof of this theorem constitutes a primary technical contribution of our paper, and is proven in [Section 3](#). Briefly, the idea behind the result is to use that the optimal policy  $\pi^*$  itself satisfies  $\pi_t^*(x; \mathbf{w}) \approx -K_\infty x - q_{\infty;h}^*(w_{t:t+h})$  whenever  $h$  is sufficiently large and  $t \leq T - \mathcal{O}_*(\log T)$ , and that  $\mathbf{A}_t^*$  has a simple quadratic structure. This characterization for is why it is essential to consider advantages with respect to the *optimal* policy  $\pi^*$ , and why our DAPs use the controller  $K_\infty$  as opposed to an arbitrary stabilizing controller as in [Agarwal et al. \(2019a\)](#).

## 2.4. Online learning with delays

An immediate consequence of [Theorem 2](#) is that for any algorithm (in particular, Riccatitron) which selects  $\pi_t(x; \mathbf{w}) =$

$-K_\infty x - q^{M_t}(\mathbf{w}_{t-1})$ , the regret  $\mathcal{M}_0\text{-Reg}_T(\pi; \mathbf{w})$  is at most

$$\sum_{t=1}^T \widehat{\mathbf{A}}_{t;h}(M_t; \mathbf{w}_{t+h}) - \inf_{M \in \mathcal{M}_0} \sum_{t=1}^T \widehat{\mathbf{A}}_{t;h}(M; \mathbf{w}_{t+h}) + 2C_{\text{adv}}. \quad (11)$$

This is simply an online convex optimization problem with  $M_1, \dots, M_T$  as iterates—the only catch is that the “loss” at time  $t$ ,  $\widehat{\mathbf{A}}_{t;h}(M_t; \mathbf{w}_{t+h})$ , can only be evaluated after observing  $\mathbf{w}_{t+h}$ , which will not be revealed to the learner until after round  $t+h$ . This is therefore an instance of online learning with *delays*, namely, the loss function suffered at time  $t$  is only available at time  $t+h+1$  (since  $\mathbf{w}_t$  is revealed at time  $t+1$ ). To reduce the problem of minimizing regret on the approximate advantages in (11) to classical online learning without delays, we use a simple black-box reduction.

Consider a generic online convex optimization setting where, at each time  $t$ , the learner proposes an iterate  $z_t$ , then suffers cost  $f_t(z_t)$  and observes  $f_t$  (or some function of it). Suppose we have an algorithm for this non-delayed setting that guarantees that for every sequence,  $\sum_{t=1}^T f_t(z_t) - \inf_{z \in \mathcal{C}} \sum_{t=1}^T f_t(z) \leq R(T)$ , where  $R$  is increasing in  $T$ . Now consider the same setting with delay  $h$ , and let  $\tau(t) = (t-1) \bmod (h+1) + 1 \in [h+1]$ . We use the following strategy: Make  $h+1$  copies of the base algorithm. At round  $t$ , observe  $z_t$ , predict  $z_t$  using the output of instance  $\tau(t)$ , then update instance  $\tau(t+1)$  using the loss  $f_{t-h}(z_{t-h})$  (which is now available).

**Lemma 2.2** (cf. Joulani et al. (2013)). The generic delayed online learning reduction has regret at most

$$\sum_{t=1}^T f_t(z_t) - \inf_{z \in \mathcal{C}} \sum_{t=1}^T f_t(z) \leq (h+1)R(T/(h+1)),$$

where  $R(T)$  is the regret of the base instance.

Lemma 2.2 shows that minimizing the regret in (11) is as easy as minimizing regret in the non-delayed setting, up to a factor of  $h = \mathcal{O}_*(\log T)$ . For completeness, we provide a proof Appendix E.4. All that remains is to specify the base algorithm for the reduction.

## 2.5. Exp-concave online learning

We have reduced the problem of obtaining logarithmic regret for online control to obtaining logarithmic regret for online learning with approximate advantages of the form in (11). A sufficient condition to obtain fast rates in online learning is strong convexity of the loss Hazan (2016), but while the advantages  $\widehat{\mathbf{A}}_{t;h}(M; \mathbf{w}_{t+h})$  are strongly convex with respect to  $q^M(\mathbf{w})$ , they are not strongly convex with respect to the parameter  $M$ . Itself. Fortunately, logarithmic regret can also be achieved for loss functions that satisfy a weaker condition called exp-concavity (Hazan et al., 2007; Cesa-Bianchi & Lugosi, 2006).

**Definition 4.** A function  $f : \mathcal{C} \rightarrow \mathbb{R}$  is  $\alpha$ -exp-concave if  $\nabla^2 f(z) \succeq \alpha(\nabla f(z))(\nabla f(z))^\top$  for all  $z \in \mathcal{C}$ .

Intuitively, an exp-concave function  $f$  exhibits strong curvature along the directions of its gradient, which are precisely the directions along which  $f$  is sensitive to change. This property holds for linear regression-type losses, as the following standard lemma (Appendix E.4) shows.

**Lemma 2.3.** Let  $A \in \mathbb{R}^{d_1 \times d_2}$ , and consider the function  $f(z) = \|Az - b\|_\Sigma^2$ , where  $\Sigma \succeq 0$ . If we restrict to  $z \in \mathbb{R}^{d_2}$  for which  $f(z) \leq R$ , then  $f$  is  $(2R)^{-1}$ -exp-concave.

Observe that the approximate advantage functions  $\widehat{\mathbf{A}}_{t;h}(M; \mathbf{w}_{t+h})$  indeed have the form  $f(z) = \|Az - b\|_\Sigma^2$  (viewing the map  $M \mapsto q^M(\mathbf{w}_{t-1})$  as a linear operator), and thus satisfy exp-concavity for appropriate  $\alpha > 0$ . To take advantage of this property we use online Newton step (ONS, Algorithm 2), a second-order online convex optimization algorithm which guarantees logarithmic regret for exp-concave losses.

**Lemma 2.4** (Hazan (2016)). Suppose that  $\sup_{z, z' \in \mathcal{C}} \|z - z'\| \leq D$ ,  $\sup_{z \in \mathcal{C}} \|\nabla f_t(z)\| \leq G$ , and that each loss  $f_k$  is  $\alpha$ -exp-concave. Then by setting  $\eta = 2 \max\{4GD, \alpha^{-1}\}$  and  $\varepsilon = \eta^2/D$ , the online Newton step algorithm guarantees

$$\sum_{k=1}^T f_k(z_k) - \inf_{z \in \mathcal{C}} \sum_{k=1}^T f_k(z) \leq 5(\alpha^{-1} + GD) \cdot d \log T.$$

**Putting everything together.** With the regret decomposition in terms of approximate advantages (Theorem 2) and the blackbox-reduction for online learning with delays (Lemma 2.2), the design and analysis of Riccatitron (Algorithm 1) is rather simple. In view of Lemma 2.1, we initialize the set  $\mathcal{M}_0$  sufficiently large to compete with the appropriate state-feedback controllers (Line 2). Using Theorem 2, our goal is to obtain a regret bound for the approximate advantages in (11). In view of the delayed online learning reduction Lemma 2.2, we initialize  $h+1$  base online learners (Line 2). Since the approximate advantages  $\widehat{\mathbf{A}}_t$  are pure quadratics, we use online Newton step for the base learner, which ensures logarithmic regret via Lemma 2.4.

## 2.6. Sharpening the regret bound

With online Newton step as the base algorithm, Riccatitron has regret  $\mathcal{O}_*(d_x d_u \sqrt{d_x} \wedge d_u \log^3 T)$ . The  $d_x d_u$  factor comes from the hard dependence on  $\dim(\mathcal{C})$  in the ONS regret bound (Lemma 2.4), while the  $\sqrt{d_x} \wedge d_u$  factor is an upper bound on the Frobenius norm for each  $M \in \mathcal{M}_0$ . We can obtain improved dimension dependence by replacing ONS with a vector-valued variant of the classical Vovk-Azoury-Warmuth algorithm (VAW), described in Algorithm 3 (Appendix E.3). The VAW algorithm goes beyond the generic exp-concave online learning setting and

exploits the quadratic structure of the approximate advantages. [Theorem 5](#) in [Appendix E.3](#) shows that its regret depends only logarithmically on the Frobenius norm of the parameter vectors, so it avoids the  $\sqrt{d_x \wedge d_u}$  factor paid by ONS (up to a log term). This leads to a final regret bound of  $\mathcal{O}_*(d_x d_u \log^3 T)$  for Riccatitron. The runtime for both algorithms is identical.

The calculation for the final regret bound is carried out in [Appendix E.1](#).

### 3. Advantages without states

We now prove the key ‘‘approximate advantage’’ theorem ([Theorem 2](#)) used in the analysis of Riccatitron. The roadmap for the proof is as follows:

1. In [Section 3.1](#), we show that the unconstrained optimal policy takes the form  $\pi_t^*(x; \mathbf{w}) = -K_t x_t - q_t^*(\mathbf{w})$ , where  $q_t^*(\mathbf{w})$  depends on all future disturbances, and where  $K_t$  is the finite-horizon solution to the Riccati recursion ([Definition 5](#)).
2. Next, [Section 3.2](#) presents an intermediate version of the approximate advantage theorem for policies of the form  $\hat{\pi}_t(x; \mathbf{w}) = -K_t x_t - q^{M_t}(\mathbf{w}_{t-1})$ . Because any such policy has the same state dependence as the optimal policy  $\pi^*$ , we are able to show that  $\mathbf{A}_t^*(u_t^{\hat{\pi}}; x_t^{\hat{\pi}}, \mathbf{w})$  has *no state dependence*. Moreover, the linear structure of the dynamics and quadratic structure of the losses ensures that  $\mathbf{A}_t^*(u_t^{\hat{\pi}}; x_t^{\hat{\pi}}, \mathbf{w})$  is a quadratic of the form  $\|q^{M_t}(\mathbf{w}_{t-1}) - q_t^*(\mathbf{w}_{t:T})\|_{\Sigma_t}^2$ , where  $\Sigma_t$  is a finite-horizon approximation to  $\Sigma_\infty$ , and  $q_t^*(\mathbf{w}_{t:T})$  is the bias vector of the optimal controller.
3. Finally ([Section 3.3](#)), we use stability of the Riccati recursion to show that  $q_t^*(\mathbf{w})$  can be replaced with a term that depends only on  $\mathbf{w}_{t+h}$ , up to a small error. Similarly, we show that  $\Sigma_t$  can be replaced by  $\Sigma_\infty$  and  $K_t$  by  $K_\infty$ .

This argument implies that a slightly modified analogue of Riccatitron which replaces infinite-horizon quantities ( $K_\infty, \Sigma_\infty, \dots$ ) with finite-horizon analogues from the Riccati recursion attains a similar regret. We state Riccatitron with the infinite horizon analogues to simplify presentation, as well as implementation.

#### 3.1. A closed form for the true optimal policy

Our first result characterizes the optimal unconstrained optimal controller  $\pi^*$  given full knowledge of the disturbance sequence  $\mathbf{w}$ , as well as the corresponding value function. To begin, we introduce a variant of the classical *Riccati recursion*.

**Definition 5** (Riccati recursion). *Define  $P_{T+1} = 0$  and  $c_{T+1} = 0$  and consider the recursion:*

$$\begin{aligned} P_t &= R_x + A^\top P_{t+1} A - A^\top P_{t+1} B \Sigma_t^{-1} B^\top P_{t+1} A, \\ \Sigma_t &= R_u + B^\top P_{t+1} B, \\ K_t &= \Sigma_t^{-1} B^\top P_{t+1} A, \\ c_t(\mathbf{w}_{t:T}) &= (A - BK_t)^\top (P_{t+1} \mathbf{w}_t + c_{t+1}(\mathbf{w}_{t+1:T})). \end{aligned}$$

We also define corresponding closed loop matrices via  $A_{\text{cl},t} = A - BK_t$ .

For i.i.d. disturbances with  $\mathbb{E}[w_t] = 0$  for all times  $t$ , the optimal controller is the state feedback law  $\pi_t(x) = -K_t x_t$ , and  $K_t \rightarrow K_\infty$  as  $t \rightarrow -\infty$ . The following theorem shows that for arbitrary disturbances the optimal controller applies the same state feedback law, but with an extra bias term that depends on the disturbance sequence.

**Theorem 3.** *The optimal controller is given by  $\pi_t^*(x, \mathbf{w}) = -K_t x - q_t^*(\mathbf{w}_{t:T})$ , where*

$$q_t^*(\mathbf{w}_{t:T}) = \sum_{i=t}^{T-1} \Sigma_t^{-1} B^\top \left( \prod_{j=t+1}^i A_{\text{cl},j}^\top \right) P_{i+1} \mathbf{w}_i. \quad (12)$$

Moreover, for each time  $t$  we have

$$\mathbf{V}_t^*(x; \mathbf{w}) = \|x\|_{P_t}^2 + 2\langle x, c_t(\mathbf{w}_{t:T}) \rangle + f_t(\mathbf{w}_{t:T}), \quad (13)$$

where  $f_t$  is a function that does not depend on the state  $x$ .

[Theorem 3](#) is a special case of a more general result, [Theorem 4](#), proven in [Appendix D](#).

#### 3.2. Removing the state

We now use the characterization of  $\pi^*$  to show that the advantages  $\mathbf{A}_t^*(u_t^{\hat{\pi}}; x_t^{\hat{\pi}}, \mathbf{w})$  have a particularly simple structure when we consider policies of the form  $\hat{\pi}_t(x; \mathbf{w}) = -K_t x_t - q_t(\mathbf{w}_{t-1})$ , where  $q_t(\mathbf{w})$  is an arbitrary function of  $\mathbf{w}$ . For such policies,  $\mathbf{A}_t^*$  is a quadratic function which does not depend explicitly on the state.

**Lemma 3.1.** Consider a policy  $\hat{\pi}_t(x)$  of the form  $\hat{\pi}_t(x; \mathbf{w}) = -K_t x_t - q_t(\mathbf{w})$ . Then, for all  $x$ ,

$$\mathbf{A}_t^*(\hat{\pi}_t(x; \mathbf{w}); x, \mathbf{w}) = \|q_t(\mathbf{w}) - q_t^*(\mathbf{w}_{t:T})\|_{\Sigma_t}^2.$$

*Proof.* Since  $\mathbf{Q}_t^*(x, \cdot; \mathbf{w})$  is a strongly convex quadratic, and since  $\pi_t^*(x; \mathbf{w}) = \arg \min_{u \in \mathbb{R}^{d_u}} \mathbf{Q}_t^*(x, u; \mathbf{w})$ , first-order optimality conditions imply that for any  $u$ ,

$$\begin{aligned} \mathbf{A}_t^*(u; x, \mathbf{w}) &= \mathbf{Q}_t^*(x, u; \mathbf{w}) - \mathbf{Q}_t^*(x, \pi_t^*(x; \mathbf{w}); \mathbf{w}) \\ &= \|u - \pi_t^*(x; \mathbf{w})\|_{\nabla_u^2 \mathbf{Q}_t^*(x, u; \mathbf{w})}^2. \end{aligned}$$

A direct computation based on [\(13\)](#) reveals that  $\nabla_u^2 \mathbf{Q}_t^*(x, u; \mathbf{w}) = R + B^\top P_{t+1} B = \Sigma_t$ , so that  $\mathbf{A}_t^*(u; x, \mathbf{w}) = \|u - \pi_t^*(x; \mathbf{w})\|_{\Sigma_t}^2$ . Finally, since  $\pi_t^*(x; \mathbf{w}) = -K_t x - q_t^*(\mathbf{w}_{t:T})$ , we have that if  $u = \hat{\pi}_t(x; \mathbf{w}) = -K_t x_t - q_t(\mathbf{w})$ , then the states in the expression  $u - \pi_t^*(x; \mathbf{w})$  cancel, leaving  $u - \pi_t^*(x; \mathbf{w}) = -(q_t(\mathbf{w}) - q_t^*(\mathbf{w}_{t:T}))$ .  $\square$



### 3.3. Truncating the future and passing to infinite horizon

The next lemma—proven in [Appendix F](#)—shows that we can truncate  $q_t^*(w_{t:T})$  to only depend on disturbances at most  $h$  steps in the future.

**Lemma 3.2.** For any  $h \in [T]$  define a truncated version of  $q_t^*$  as follows:

$$q_{t:t+h}^*(w_{t:t+h}) = \sum_{i=t}^{(t+h) \wedge T-1} \Sigma_t^{-1} B^\top \left( \prod_{j=t+1}^i A_{cl,j}^\top \right) P_{i+1} w_i. \quad (14)$$

Then for any  $t$  such that  $t+h < T - \tilde{O}(\beta_* \Psi_*^2 \Gamma_*)$ , setting  $\bar{\gamma}_\infty = \frac{1}{2}(1 + \gamma_\infty) < 1$ , we have the bound  $\|q_{t:t+h}^*(w_{t:t+h}) - q_t^*(w_{t:T})\| \leq \kappa_\infty^2 \beta_*^2 \Psi_* \Gamma_*^2 (T-h) \bar{\gamma}_\infty^h$ , which is geometrically decreasing in  $h$ .

Going forward we use that both  $q_t^*$  and  $q_{t:t+h}^*$  have norm at most  $\beta_* \Psi_* \Gamma_* \kappa_\infty (1 - \gamma_\infty)^{-1} =: D_{q^*}$  ([Lemma D.6](#)). As an immediate corollary of [Lemma 3.2](#), we approximate the advantages using finite lookahead.

**Lemma 3.3.** Consider a policy  $\hat{\pi}_t(x; \mathbf{w}) = -K_t x_t - q_t(\mathbf{w})$ , and suppose that  $\|q_t\| \leq D_q$ , where  $D_q \geq D_{q^*}$ . If we choose  $h = 2(1 - \gamma_\infty)^{-1} \log(\kappa_\infty^2 \beta_*^2 \Psi_* \Gamma_*^2 T^2)$ , we are guaranteed that

$$\sum_{t=1}^T |\mathbf{A}_t^*(u_t^{\hat{\pi}}; x_t^{\hat{\pi}}, \mathbf{w}) - \|q_t(\mathbf{w}) - q_{t:t+h}^*(w_{t:t+h})\|_{\Sigma_t}^2| \leq C_{\text{trunc}},$$

where  $C_{\text{trunc}} \leq \tilde{O}(D_q^2 \beta_* \Psi_* \Gamma_*^2 (1 - \gamma_\infty)^{-1} \log T)$ .

At this point, we have established an analogue of [Theorem 2](#), except that we are still using state-action controllers  $K_t$  rather than  $K_\infty$ , and the approximate advantages in [Lemma 3.3](#) are using the finite-horizon counterparts of  $\Sigma_\infty$  and  $q_{\infty;h}$ . The following lemmas show that we can pass to these infinite-horizon quantities by paying a small approximation cost.

**Lemma 3.4.** Let policies  $\pi_t(x; \mathbf{w}) = -K_\infty x - q_t(\mathbf{w})$  and  $\hat{\pi}_t(x; \mathbf{w}) = -K_t x - q_t(\mathbf{w})$  be given, where  $q_t$  is arbitrary but satisfies  $\|q_t\| \leq D_q$  for some  $D_q \geq 1$ . Then

$$|J_T(\hat{\pi}, \mathbf{w}) - J_T(\pi, \mathbf{w})| \leq C_{K_\infty},$$

where  $C_{K_\infty} \leq \tilde{O}(\kappa_\infty^4 \beta_*^6 \Psi_*^{13} \Gamma_*^6 (1 - \gamma_\infty)^{-2} D_q^2 \cdot \log(D_q T))$ .

**Lemma 3.5.** Let  $(q_t)_{t=1}^T$  be an arbitrary sequence with  $\|q_t\| \leq D_q$  for some  $D_q \geq D_{q^*}$ . Then it holds that

$$\left| \sum_{t=1}^T \|q_t - q_{t:t+h}^*(w_{t:t+h})\|_{\Sigma_t}^2 - \|q_t - q_{\infty;h}^*(w_{t:t+h})\|_{\Sigma_\infty}^2 \right| \leq C_\infty,$$

where  $C_\infty \leq \tilde{O}(D_q^2 \cdot \beta_*^4 \Psi_*^7 \Gamma_*^4 \kappa_\infty^2 (1 - \gamma_\infty)^{-1} h \log(D_q T))$ .

Combining these results immediately yields the proof of [Theorem 2](#); details are given in [Appendix F](#).

## 4. Conclusion

We have presented the first efficient algorithm with logarithmic regret for online linear control with arbitrary adversarial disturbance sequences. Our result highlights the power of online learning with advantages, and we are hopeful that this framework will find broader use. Numerous questions naturally arise for future work: Does our framework extend to more general loss functions, or to more general classes of dynamical systems in control and reinforcement learning? Can our results be extended to handle partial observed dynamical systems? Can we obtain  $\sqrt{T}$ -regret for adversarial disturbances in unknown systems, as is possible in the stochastic regime?

### Acknowledgements

DF acknowledges the support of NSF TRIPODS award #1740751. MS is generously supported by an Open Philanthropy AI Fellowship. We thank Ali Jadbabaie for helpful discussions.

### References

- Abbasi-Yadkori, Y. and Szepesvári, C. Regret bounds for the adaptive control of linear quadratic systems. In *Proceedings of the 24th Annual Conference on Learning Theory*, pp. 1–26, 2011.
- Abbasi-Yadkori, Y., Bartlett, P. L., Kanade, V., Seldin, Y., and Szepesvári, C. Online learning in markov decision processes with adversarially chosen transition probability distributions. In *Advances in neural information processing systems*, pp. 2508–2516, 2013.
- Abbasi-Yadkori, Y., Bartlett, P., and Kanade, V. Tracking adversarial targets. In *International Conference on Machine Learning*, pp. 369–377, 2014.
- Agarwal, N., Bullins, B., Hazan, E., Kakade, S., and Singh, K. Online control with adversarial disturbances. In *International Conference on Machine Learning*, pp. 111–119, 2019a.
- Agarwal, N., Hazan, E., and Singh, K. Logarithmic regret for online control. In *Advances in Neural Information Processing Systems*, pp. 10175–10184, 2019b.
- Anava, O., Hazan, E., and Mannor, S. Online learning for adversaries with memory: price of past mistakes. In *Advances in Neural Information Processing Systems*, pp. 784–792, 2015.
- Arora, R., Dekel, O., and Tewari, A. Online bandit learning against an adaptive adversary: from regret to policy regret. In *International Conference on Machine Learning (ICML)*, pp. 1747–1754, 2012.

- Azoury, K. S. and Warmuth, M. K. Relative loss bounds for on-line density estimation with the exponential family of distributions. *Machine Learning*, 43(3):211–246, June 2001.
- Bertsekas, D. P. *Dynamic Programming and Optimal Control, Vol. I*. Athena Scientific, 2005.
- Cassel, A., Cohen, A., and Koren, T. Logarithmic regret for learning linear quadratic regulators efficiently. *International Conference on Machine Learning (ICML)*, 2020.
- Cesa-Bianchi, N. and Lugosi, G. *Prediction, Learning, and Games*. Cambridge University Press, 2006.
- Cohen, A., Hasidim, A., Koren, T., Lazic, N., Mansour, Y., and Talwar, K. Online linear quadratic control. In *International Conference on Machine Learning*, pp. 1028–1037, 2018.
- Cohen, A., Koren, T., and Mansour, Y. Learning linear-quadratic regulators efficiently with only  $\sqrt{T}$  regret. In *International Conference on Machine Learning*, pp. 1300–1309, 2019.
- Dean, S., Mania, H., Matni, N., Recht, B., and Tu, S. Regret bounds for robust adaptive control of the linear quadratic regulator. In *Advances in Neural Information Processing Systems*, pp. 4188–4197, 2018.
- Even-Dar, E., Kakade, S. M., and Mansour, Y. Online markov decision processes. *Mathematics of Operations Research*, 34(3):726–736, 2009.
- Faradonbeh, M. K. S., Tewari, A., and Michailidis, G. On optimality of adaptive linear-quadratic regulators. *arXiv preprint arXiv:1806.10749*, 2018.
- Fazel, M., Ge, R., Kakade, S., and Mesbahi, M. Global convergence of policy gradient methods for the linear quadratic regulator. In *International Conference on Machine Learning*, pp. 1467–1476, 2018.
- Hazan, E. Introduction to online convex optimization. *Foundations and Trends® in Optimization*, 2(3-4):157–325, 2016.
- Hazan, E. and Kale, S. Newtron: an efficient bandit algorithm for online multiclass prediction. In *Advances in Neural Information Processing Systems*, pp. 891–899, 2011.
- Hazan, E., Agarwal, A., and Kale, S. Logarithmic regret algorithms for online convex optimization. *Machine Learning*, 69(2-3):169–192, 2007.
- Hazan, E., Kakade, S., and Singh, K. The nonstochastic control problem. In *Proceedings of the 31st International Conference on Algorithmic Learning Theory*, pp. 408–421. PMLR, 2020.
- Joulani, P., Gyorgy, A., and Szepesvári, C. Online learning under delayed feedback. In *International Conference on Machine Learning*, pp. 1453–1461, 2013.
- Kakade, S. M. *On the sample complexity of reinforcement learning*. PhD thesis, University College London (University of London), 2003.
- Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., and Wierstra, D. Continuous control with deep reinforcement learning. *International Conference on Learning Representations (ICLR)*, 2016.
- Lincoln, B. and Rantzer, A. Relaxing dynamic programming. *IEEE Transactions on Automatic Control*, 51(8): 1249–1260, 2006.
- Mania, H., Tu, S., and Recht, B. Certainty equivalence is efficient for linear quadratic control. In *Advances in Neural Information Processing Systems*, pp. 10154–10164, 2019.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540): 529, 2015.
- Orabona, F., Crammer, K., and Cesa-Bianchi, N. A generalized online mirror descent with applications to classification and regression. *Machine Learning*, 99(3):411–435, 2015.
- Rakhlin, A. and Sridharan, K. Online nonparametric regression. In *Conference on Learning Theory*, 2014.
- Rosenberg, A. and Mansour, Y. Online convex optimization in adversarial markov decision processes. In *International Conference on Machine Learning*, pp. 5478–5486, 2019.
- Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., et al. Mastering the game of go without human knowledge. *Nature*, 550(7676):354, 2017.
- Simchowitz, M. and Foster, D. J. Naive exploration is optimal for online LQR. *International Conference on Machine Learning (ICML)*, 2020.
- Simchowitz, M., Singh, K., and Hazan, E. Improper learning for non-stochastic control. *Conference on Learning Theory (COLT)*, 2020.

- Slotine, J.-J. E. and Li, W. *Applied nonlinear control*. Prentice-Hall, 1991.
- Vovk, V. Aggregating strategies. *Proceedings of the conference on computational Learning Theory*, 1990.
- Vovk, V. A game of prediction with expert advice. In *Proceedings of the eighth annual conference on computational learning theory*, pp. 51–60. ACM, 1995.
- Vovk, V. Competitive on-line linear regression. In *NIPS '97: Proceedings of the 1997 conference on advances in neural information processing systems 10*, pp. 364–370, Cambridge, MA, USA, 1998. MIT Press.
- Yu, J. Y., Mannor, S., and Shimkin, N. Markov decision processes with arbitrary reward processes. *Mathematics of Operations Research*, 34(3):737–757, 2009.
- Zimin, A. and Neu, G. Online learning in episodic markovian decision processes by relative entropy policy search. In *Advances in neural information processing systems*, pp. 1583–1591, 2013.

## Organization and notation

Notation	Definition
$A, B$	system matrices (dynamics, Eq. (1))
$w_t, w_{1:t}, \mathbf{w}$	disturbance at time $t$ ; from $1, \dots, t$ ; from $1, 2, \dots$
$\pi$	control policy
$K, \pi^K$	static feedback controller, induced policy
$M, \pi^{(M)}$	DAP controller (Definition 2), induced policy
$\pi^{\text{alg}}$	policy selected by the learner
$d_x, d_u$	state/input dimension
$\mathbf{x}_t^\pi(\mathbf{w}), \mathbf{u}_t^\pi(\mathbf{w})$	state/input at time $t$ under policy $\pi$ and disturbances $\mathbf{w}$
$\mathbf{x}_t^K(\mathbf{w}), \mathbf{u}_t^K(\mathbf{w})$	state/input at time $t$ under policy $\pi^K$
$\ell, R_x, R_u$	cost function $\ell(x, u) = x^\top R_x x + u^\top R_u$
$J_T(\pi; \mathbf{w})$	cost of policy $\pi$ , $\sum_{t=1}^T \ell(\mathbf{x}_t^\pi(\mathbf{w}), \mathbf{u}_t^\pi(\mathbf{w}))$
$\text{Reg}_T(\pi^{\text{alg}}; \Pi, \mathbf{w})$	regret with benchmark $\Pi$ : $J_T(\pi^{\text{alg}}; \mathbf{w}) - \inf_{\pi \in \Pi} J_T(\pi; \mathbf{w})$
$\mathcal{K}_0$	benchmark class of strongly stable feedback controllers
$\mathcal{K}_0\text{-Reg}_T(\pi^{\text{alg}}; \mathbf{w})$	regret benchmark with comparator $\mathcal{K}_0$ : $J_T(\pi^{\text{alg}}; \mathbf{w}) - \inf_{K \in \mathcal{K}_0} J_T(\pi^K; \mathbf{w})$ .
$\mathcal{M}_0$	benchmark class of DAP controllers (parameterized as $\mathcal{M}(m, R, \gamma)$ in Definition 2)
$\mathcal{M}_0\text{-Reg}_T(\pi^{\text{alg}}; \mathbf{w})$	regret relative to $\mathcal{M}_0$ (Eq. (7))
$P_\infty$	solution to the DARE (Eq. (4))
$K_\infty$	optimal infinite horizon LQR controller
$\Sigma_\infty$	optimal infinite horizon LQR covariance
$A_{\text{cl}, \infty}$	closed loop system $A - BK_\infty$ under optimal infinite-horizon controller
$\kappa_\infty, \gamma_\infty$	strong stability parameters for $A_{\text{cl}, \infty}$ (see Definition 1)
$\kappa_0, \gamma_0$	strong stability parameters for $\mathcal{K}_0$ (see Definition 1)
$\Psi_\star$	$\max\{1, \ A\ _{\text{op}}, \ B\ _{\text{op}}, \ R_x\ _{\text{op}}, \ R_u\ _{\text{op}}\}$
$\beta_\star$	$\max\{1, \lambda_{\min}^{-1}(R_u), \lambda_{\min}^{-1}(R_x)\}$
$\Gamma_\star$	$\max\{1, \ P_\infty\ _{\text{op}}\}$
$\pi^\star$	unconstrained optimal policy (Definition 3)
$\mathbf{Q}^\star, \mathbf{V}^\star$	Q-function and value function under $\pi^\star$ (Definition 3)
$\mathbf{A}^\star$	advantage, defined as $\mathbf{Q}_t^\star(x, u; \mathbf{w}) - \mathbf{Q}_t^\star(x, \pi_t^\star(x; \mathbf{w}); \mathbf{w})$
$q_t$	generic bias-predicting term (e.g., $\pi(x, \mathbf{w}) = -K_\infty x - q_t(\mathbf{w})$ )
$q_t^M$	bias-predicting term in DAP (Definition 2)
$q_{\infty; h}^\star(w_{1:h+1})$	truncated approximation for $q_t^\star$ -function in $\pi^\star$ (defined above Eq. (10)).
$\hat{\mathbf{A}}_{t; h}(M; \mathbf{w}_{t+h})$	approximate advantage (Eq. (10))
$\mathcal{C}$	generic constraint set for online optimization
$\text{ONS}, \text{ONS}(\varepsilon, \eta, \mathcal{C})$	Online Newton Step (Algorithm 2) with learning rate $\eta$ , regularization parameter $\varepsilon$
$\text{VAW}, \text{VAW}(\varepsilon, \mathcal{C}, \Sigma)$	Vovk-Azoury-Warmuth (Algorithm 3) with regularization parameter $\varepsilon$ , cost $\Sigma$
$P_t, \Sigma_t, K_t, A_{\text{cl}, t}$	finite-horizon optimal analogues of $P_\infty, \Sigma_\infty, K_\infty, A_{\text{cl}, \infty}$ (Definition 6)
$q_t^\star(w_{t:T})$	bias function for $\pi^\star$ ( $\pi_t^\star(x; w_{t:T}) = -K_t x - q_t^\star(w_{t:T})$ ; Eq.(12))
$H_{\text{cl}, \infty}, L_{\text{cl}, \infty}$	matrices that witness strong stability of $A_{\text{cl}, \infty}$ (Definition 1)
$L_{\text{cl}, t}$	$H_{\text{cl}, \infty}^{-1} A_{\text{cl}, t} H_{\text{cl}, \infty}$ (used to show strong stability of $A_{\text{cl}, t}$ )
$A_{\text{cl}, i \rightarrow t}$	$A_{\text{cl}, t} A_{\text{cl}, t-1} \dots A_{\text{cl}, i+1}$ , with convention $A_{\text{cl}, t \rightarrow t} = I$
$\bar{\gamma}_\infty$	$\frac{1}{2}(1 + \gamma_\infty) < 1$
$\Delta_{\text{stab}}$	$4 \cdot \beta_\star \Psi_\star^2 \Gamma_\star \log(2\Psi_\star \Gamma_\star \kappa_\infty (1 - \gamma_\infty)^{-1})$ (“decay time”)
$T_{\text{stab}}$	$T - \Delta_{\text{stab}}$

Table 1: Summary of notation

This appendix is organized as follows. [Appendix A](#) discusses additional related work. [Appendix B](#) describes extensions of RiccatiTron. [Appendix B.1](#) extends the algorithm to compete with general benchmark policy classes. [Appendix B.2](#)



demonstrates that Riccatitron competes with richer benchmark class that includes arbitrary linear controllers with internal state; [Appendix B.3](#) extends the algorithm to consider “tracking costs” studied by [Abbasi-Yadkori et al. \(2014\)](#); [Appendix B.4](#) explains how the algorithm can accommodate time-varying quadratic costs, provided that they are known to the learner in advance.

[Appendix C](#) explains challenges associated with using online learning with stationary costs (OLWS) to attain logarithmic regret in our setting. This appendix also provides a unifying (albeit informal) treatment of existing OLWS approaches. In addition, [Appendix C.5](#) highlights the differences between Riccatitron and MDP-E ([Even-Dar et al., 2009](#)), a variant of OLWS which is superficially similar to our approach.

The remaining three appendices are dedicated to proving our main results. [Appendix D](#) collects some basic structural results for linear quadratic control which we use throughout the appendix, and [Appendix D.2](#) describes a variant of the performance difference lemma ([Kakade, 2003](#)) which is used in our analysis. [Appendix E](#) provides the missing proofs from [Section 2](#). Importantly, [Appendix E.1](#) proves [Theorem 1](#), and [Appendix E.3](#) establishes a regret guarantee for the vector-valued VAW algorithm ([Algorithm 3](#)). Finally, [Appendix F](#) supplies the missing proofs from [Section 3](#), culminating in the proof of [Theorem 2](#).

Notation used throughout the main paper and appendix is collected in [Table 1](#).

## A. Additional related work

**Linear control for known systems.** [Cohen et al. \(2018\)](#) establish  $\sqrt{T}$  regret for online control of known linear systems under stochastic noise and time varying quadratic cost. [Agarwal et al. \(2019a\)](#) achieve  $\sqrt{T}$ -regret with both *adversarial* disturbances and time varying, adversarially chosen loss functions  $\ell_t$  via a reduction to online convex optimization with memory ([Anava et al., 2015](#)). Their approach adopts a “disturbance-action” policy parameterization (or, DAP), which we utilize as well ([Definition 2](#)). Certain previous results achieve logarithmic regret by making assumptions that ensure stationary costs are strongly convex, allowing for logarithmic regret and movement cost via [Anava et al. \(2015\)](#) or similar arguments. [Abbasi-Yadkori et al. \(2014\)](#) consider an online tracking problem with known system parameters *zero* exogenous noise. The absence of noise enables an approach based on MDP-E (see [Appendix C.5](#)), for which the relevant Q-functions in this setting are strongly convex, leading to logarithmic regret. More recently [Agarwal et al. \(2019b\)](#) showed that in the noisy setting the stationary costs  $\lambda_t$  themselves are strongly convex in a disturbance-action parametrization, provided that the loss functions  $\ell_t$  are strongly convex and the noise covariance is well-conditioned, which also leads to logarithmic regret. [Simchowitz et al. \(2020\)](#) show that this approach extends to “semi-adversarial” disturbances with a well-conditioned stochastic component and a possibly adversarial component. Our results (with the restriction that costs are quadratic) give the first logarithmic regret bounds for the fully adversarial setting and, to the best of our knowledge, give the first instance in online control where an exp-concave but not strongly convex parametrization attains logarithmic regret.

**Linear control for unknown systems.** For *unknown* systems, various works ([Abbasi-Yadkori & Szepesvári, 2011](#); [Faradonbeh et al., 2018](#); [Cohen et al., 2019](#); [Mania et al., 2019](#)) establish  $\sqrt{T}$ -regret for fixed quadratic losses and stationary stochastic noise, which is optimal for this setting ([Simchowitz & Foster, 2020](#); [Cassel et al., 2020](#)). Because of the stochastic nature of these problems, purely statistical techniques suffice. By combining these techniques with OCO with memory ([Anava et al., 2015](#)), other recent works have addressed both unknown dynamics and adversarial noise ([Hazan et al., 2020](#); [Simchowitz et al., 2020](#)). ([Cassel et al., 2020](#)) show that logarithmic regret *is* achievable under stochastic noise for systems  $(A, B)$  where only  $A$  is unknown, or where only  $B$  is unknown and the optimal controller satisfies a non-degeneracy assumption.

**Online reinforcement learning.** Online linear control belongs to a broader line of work on online reinforcement learning in (known or unknown) Markov decision processes with adversarial costs or transitions. Given the staggering breadth of work in this direction from the online learning, control, and RL communities, we focus on past contributions which are most closely related to our setting. As discussed earlier, essentially all prior approaches to online RL abide by the OLWS paradigm. Perhaps the first result in this direction is the MDP-E algorithm of [Even-Dar et al. \(2009\)](#), which attains  $\sqrt{T}$  policy regret in a tabular MDP with known stationary dynamics and adversarially chosen rewards. Subsequent works ([Abbasi-Yadkori et al., 2013](#)) achieves  $\sqrt{T}$ -regret in a tabular setting where the both the rewards and transition kernels are selected by an adversary. A parallel line of work on adversarial tabular MDPs considers the episodic setting ([Zimin & Neu, 2013](#); [Rosenberg & Mansour, 2019](#)), which alleviates the need to bound the movement costs between iterates.

**Policy regret.** All of the approaches described so far can be viewed as special cases of the general problem of minimizing *policy regret* in online learning. A finite-memory formulation of the policy regret benchmark was popularized by Arora et al. (2012). Anava et al. (2015) generalize this result to the *online convex optimization with memory* setting and demonstrate that many popular online learning algorithms naturally produce slow-moving iterates, yielding near-optimal policy regret bounds (see Appendix C.3.1 for detailed discussion). These results have found immediate application in online linear control (Agarwal et al., 2019a; Hazan et al., 2020; Simchowitz et al., 2020). However, the analysis of Anava et al. (2015) does not extend to give fast rates for the exp-concave loss functions which arise in our setting.

## B. Extensions

### B.1. General policy classes

In view of Section 2 and Section 3, it should be clear the disturbance-action parameterization used in Riccatitron serves only to facilitate the use of tools from online convex optimization. By appealing to tools from the more general online learning framework, we can derive rates for generic, potentially nonlinear benchmark policy classes.

Suppose we wish to compete with a benchmark class  $\Pi$  where each  $\pi \in \Pi$  takes the form  $\pi(x; \mathbf{w}) = -K_\infty x - q_t^\pi(\mathbf{w}_{t-1})$ , and suppose that the learner’s policy takes the form  $\pi^{\text{alg}}(x; \mathbf{w}) = -K_\infty x - q_t^{\text{alg}}(\mathbf{w}_{t-1})$ . The development so far implies that as long as  $\|q_t^\pi\|$  is uniformly bounded for all  $\pi \in \Pi$ , we have

$$\text{Reg}_T(\pi^{\text{alg}}; \Pi, \mathbf{w}) = \sum_{t=1}^T \left\| q_t^{\text{alg}}(\mathbf{w}_{t-1}) - q_{\infty;h}^*(\mathbf{w}_{t+h}) \right\|_{\Sigma_\infty}^2 - \inf_{\pi \in \Pi} \sum_{t=1}^T \left\| q_t^\pi(\mathbf{w}_{t-1}) - q_{\infty;h}^*(\mathbf{w}_{t+h}) \right\|_{\Sigma_\infty}^2 + C_{\text{err}}, \quad (15)$$

where  $C_{\text{err}}$  is a logarithmic approximation error term. We can appeal to the generic delayed online learning reduction once more to reduce this problem to online supervised learning. Consider the following protocol for online learning: At time  $t$ : Receive  $w_{t-1}$ , predict  $\hat{q}_t \in \mathbb{R}^{d_u}$ , then receive  $q_t^* \in \mathbb{R}^{d_u}$ . If we have an algorithm for this protocol that ensures

$$\sum_{t=1}^T \|\hat{q}_t - q_t^*\|_{\Sigma_\infty}^2 - \inf_{\pi \in \Pi} \sum_{t=1}^T \|q_t^\pi(\mathbf{w}_{t-1}) - q_t^*\|_{\Sigma_\infty}^2 \leq R_{\text{OSL}}(T), \quad (16)$$

for every sequence, then the delayed online learning reduction enjoys regret  $(h+1)R_{\text{OSL}}(T)/(h+1)$  for the delayed problem (15). For example, since the loss  $\hat{q} \mapsto \|\hat{q} - q^*\|_{\Sigma_\infty}^2$  is exp-concave, we can apply Vovk’s aggregating algorithm (Vovk, 1990; 1995) to guarantee

$$\text{Reg}_T(\pi^{\text{alg}}; \Pi, \mathbf{w}) \leq \mathcal{O}_*(\log|\Pi| \cdot \log T)$$

for any finite class of policies. More generally, one can derive fast rates for arbitrary nonparametric classes of benchmark policies via the offset Rademacher complexity-based minimax bounds given in Rakhlin & Sridharan (2014).

### B.2. Alternative regret benchmarks

Throughout the main paper we only considered benchmarks based on linear feedback controllers of the form  $u_t = -Kx_t$ , where  $K$  is strongly stabilizing. We now show that DAP controllers (and consequently Riccatitron) can be used to compete with a more general class of linear controllers with *internal state*. We use an argument from Simchowitz et al. (2020). Consider  $m_Q \in \mathbb{N}$ , and controller of the form

$$\pi_t^{[Q]}(x; \mathbf{w}) = -K_\infty x + \sum_{i=0}^{m_Q-1} Q^{[i]} x_{t-i}^{K_\infty}(\mathbf{w}), \quad (17)$$

where  $x_t^{K_\infty}(\mathbf{w})$  denotes the state that would arise at time  $t$  if the linear selected the optimal linear control law  $u_s^{K_\infty}(\mathbf{w}) = -K_\infty x_s^{K_\infty}(\mathbf{w})$  for all  $s < t$ . We note that this counterfactual can be computed from  $w_{1:t-1}$ . By Simchowitz et al. (2020), to show that the DAP parameterization competes with controllers with internal state, it suffices to show that the parameterization competes with controllers of the form (17). To see this is indeed the case, observe that since  $K_\infty$  stabilizes the system  $(A, B)$ , we have

$$x_s^{K_\infty}(\mathbf{w}) = \sum_{i=0}^h (A - BK_\infty)^i w_{s-i-1} \pm e^{-\Omega(h(1-\gamma_\infty))},$$

where we use  $\pm$  in an informal, vector-valued sense. Hence, we can render

$$\pi_t^{[Q]}(x; \mathbf{w}) = -K_\infty x + \sum_{i=0}^{m_Q-1} \sum_{j=0}^h Q^{[i]}(A - BK_\infty)^j w_{t-(i+j+1)} \pm e^{-\Omega(h(1-\gamma_\infty))}. \quad (18)$$

It follows that setting  $m = m_Q + h$ , we can approximate the above behavior with an  $m$ -length controller of the form  $M^{[k]} = \sum_{i=0}^{m-1} \sum_{j=0}^h Q^{[i]}(A - BK_\infty)^j \mathbb{1}_{i+j+1=k}$  captures the policy (17).

To formalize the extension, one must also verify that for some reasonable  $R, m$ , the sequence  $M$  above lies in the set

$$\mathcal{M}(m, R, \gamma) := \{M = (M^{[i]})_{i=1}^m : \|M^{[i]}\|_{\text{op}} \leq R\gamma^{i-1}\},$$

that is, the sequence enjoys geometric decay with parameter  $\gamma$ . This decay can be achieved in numerous ways, e.g. taking  $\gamma = 1/m$  and inflating  $R$  by a factor of  $e$ . At the extreme, one can show that the constraint set  $\mathcal{M}(m, R, \gamma)$  can be replaced with a set which *does not* enforce geometric decay,

$$\widetilde{\mathcal{M}}(m, R) := \{M = (M^{[i]})_{i=1}^m : \|M^{[i]}\|_{\text{op}} \leq R, \forall i\},$$

at the expense of suffering a larger polynomial in  $\log T$  in the final regret bound. We omit the details in the interest of brevity.

### B.3. Tracking moving targets

We next show that Riccatitron generalizes to a setting with moving targets (or, ‘‘adversarial targets’’) previously studied without adversarial noise by Abbasi-Yadkori et al. (2014). In this setting, for a sequence of *targets*  $a_{1:T}, b_{1:T}$ , the learner’s loss at time  $t$  is given by

$$\ell_t(x, u) = \ell(x - a_t, u - b_t) = \|x - a_t\|_{R_x}^2 + \|u - b_t\|_{R_u}^2.$$

Let us adopt the shorthand  $\bar{w}_t = (w_t, a_t, b_t)$ , and  $\bar{\mathbf{w}} = (w_{1:T}, a_{1:T}, b_{1:T})$ . Theorem 4—proven in Appendix D—shows that if we define

$$q_t^*(\bar{w}_{t:T}) = \Sigma_t^{-1} \left( -R_u b_t + B^\top \sum_{i=t}^{T-1} \left( \prod_{j=t+1}^i A_{\text{cl},j}^\top \right) P_{i+1} w_i + B^\top \sum_{i=t+1}^{T-1} \left( \prod_{j=t+1}^{i-1} A_{\text{cl},j}^\top \right) (K_i^\top R_u b_i - R_x a_i) \right), \quad (19)$$

where  $K_t, \Sigma_t$ , and so on are given by the Riccati Recursion (Definition 5), then the optimal unconstrained controller is given by  $\pi_t^*(x; \bar{\mathbf{w}}) = -K_t x - q_t^*(\bar{w}_{t:T})$ . Retracing our steps from the special case without moving targets, we have the following generalization of Lemma 3.1.

**Lemma B.1** (Advantages for Moving Targets). Consider a policy  $\widehat{\pi}_t(x)$  of the form  $\widehat{\pi}_t(x) = -K_t x - q_t(\bar{\mathbf{w}})$ . For all  $x$ , we have

$$\mathbf{A}_t^*(\widehat{\pi}_t(x); x, \bar{\mathbf{w}}) = \|q_t(\bar{\mathbf{w}}) - q_t^*(\bar{w}_{t:T})\|_{\Sigma_t}^2,$$

where  $q_t^*(\bar{w}_{t:T})$  is given by (19).

To extend Riccatitron to this setting, we define truncated versions of  $q^*$  and  $\mathbf{A}^*$  analogous to to the without-moving-targets case (Eq. (10)). With  $\bar{w}_t := ((w_t, a_t, b_t), (w_{t-1}, a_{t-1}, b_{t-1}), \dots)$ , we define

$$\begin{aligned} q_{\infty;h,\text{move}}^*(\bar{w}_{1:h+1}) &:= \Sigma_\infty^{-1} \left( -R_u b_t + B^\top \sum_{i=1}^{h+1} (A_{\text{cl},\infty}^\top)^{i-1} P_\infty w_i + B^\top \sum_{i=2}^{h+1} (A_{\text{cl},\infty}^\top)^{i-2} (K_i^\top R_u b_i - R_x a_i) \right), \\ \widehat{\mathbf{A}}_{t;h,\text{move}}(M; \bar{w}_{t+h}) &:= \|q_t^M(\mathbf{w}_{t-1}) - q_{\infty;h}^*(\bar{w}_{t+h})\|_{\Sigma_\infty}^2. \end{aligned} \quad (20)$$

We simply run Riccatitron with the new approximate advantage functions  $\widehat{\mathbf{A}}_{t;h,\text{move}}(M; \bar{w}_{t+h})$  from (20) replacing their without-moving-targets variants from (10). Logarithmic regret follows by the same arguments.

#### B.4. Varying quadratic costs

As a final generalization, we show that our analysis generalizes to time varying quadratic losses  $\ell_t(x, u) = x^\top R_{t;x}x + u^\top R_{t;u}u$ , provided the cost matrices  $R_{t;x}$  and  $R_{t;u}$  are known to the learner ahead of time. Of course, this extension generalizes further to “tracking” losses of the form  $\ell_t(x, u) = \|x - a_t\|_{R_{t;x}}^2 + \|u - u_t\|_{R_{t;u}}^2$  as in the previous section.

To perform this generalization, we consider the following variant of the Riccati Recursion.

**Definition 6** (Time-varying Riccati Recursion). *Define  $P_{T+1} = 0$  and  $c_{T+1} = 0$  and consider the recursion:*

$$\begin{aligned} P_t &= R_{t;x} + A^\top P_{t+1}A - A^\top P_{t+1}B\Sigma_t^{-1}B^\top P_{t+1}A, \\ \Sigma_{t;T} &= R_{t;u} + B^\top P_{t+1}B, \\ K_{t;T} &= \Sigma_t^{-1}B^\top P_{t+1}A, \\ c_t(w_{t:T}) &= (A - BK_t)^\top (P_{t+1}w_t + c_{t+1}(w_{t+1:T})). \end{aligned}$$

We similarly define closed-loop matrices  $A_{cl,t} = (A - BK_t)$ . The form of the optimal policy generalizes in the obvious way

$$\pi_t^*(x; \mathbf{w}) = -K_t x - q_t^*(w_{t:T}), \quad \text{and} \quad q_t^*(w_{t:T}) = \sum_{i=t}^{T-1} \Sigma_t^{-1} B^\top \left( \prod_{j=t+1}^i A_{cl,j}^\top \right) P_{i+1} w_i.$$

Advantages take the form

$$\mathbf{A}_{t;T}^*(\widehat{\pi}_t(x); x, \mathbf{w}) = \|q_t(\mathbf{w}) - q_t^*(w_{t:T})\|_{\Sigma_t}^2.$$

Note that compared to the fixed-cost setting, we cannot leverage the existence of the “steady-state” matrix  $P_\infty$  here. Nonetheless, we can still truncate the dependence on the future by using the vectors  $q_{t;T}^*(w_t, \dots, w_{t+h}, 0, \dots, 0)$  to create approximate advantages with finite lookahead, which can then be used within the Riccati scheme.

### C. Limitations of online learning with stationary costs

This section highlights the technical challenges encountered when attempting to apply OLWS to attain logarithmic regret in online control with adversarial disturbances. In addition to highlighting the advantages (no pun intended) of our OLWA approach, this appendix may serve as an informal tutorial of prior approaches for online control problems. The section is organized as follows:

1. [Appendix C.1](#) gives an intuitive overview of the OLWS paradigm, explaining that the regret encountered by the learner incurs a ‘stationarization’ cost reflecting the mismatch between the costs induced by the learner’s actual visited trajectory and the trajectories considered by the stationary costs.
2. [Appendix C.2](#) explains that the standard approach for bounding stationarization cost is in terms of a “movement cost”, which measures the cumulative differences between successive policies  $\pi_t$ : informally,  $\sum_{t=1}^T \|\pi_t - \pi_{t-1}\|$ . Pointing forward to [Appendix C.4](#), we explain that this is the major barrier to obtaining logarithmic regret in our setting. In contrast, stationarization/movement costs *do not* arise in our analysis of OLWA, leading to our main result.
3. [Appendix C.3](#) reviews in greater detail how the OLWS paradigm has been applied to online control with adversarial disturbances. [Appendix C.3.1](#) covers the OCO-with-memory framework due to [Anava et al. \(2015\)](#). [Appendix C.3.2](#) shows how [Agarwal et al. \(2019a\)](#) instantiate this framework for online control with the DAP parametrization, detailing the (approximate) stationary cost functions  $f_{t,h}(M)$  that arise and the corresponding movement cost in the regret analysis. Examining these loss functions, [Appendix C.3.3](#) shows that they are exp-concave but not strongly convex.
4. [Appendix C.4](#) demonstrates that—in the OCO-with-memory framework—the movement cost for sequences of exp-concave but non-strongly convex functions can scale as  $\Omega(\sqrt{T})$  in the worst case. This implies that any analysis which uses a black-box reduction to OCO-with-memory with bounded movement cost cannot guarantee rates faster than  $\mathcal{O}(\sqrt{T})$ .<sup>5</sup>

<sup>5</sup>Note that this argument does not preclude the possibility that OLWS algorithms can attain logarithmic regret; rather, it demonstrates that an analysis which passes through movement costs for arbitrary exp-concave stationary costs is insufficient.



Finally, [Appendix C.5](#) compares our OLwA approach to the MDP-E algorithm proposed by [Even-Dar et al. \(2009\)](#). These two algorithms are superficially similar, because they both consider control-theoretic advantages. Despite these similarities, we note that MDP-E is still an instance of OLwS, and therefore succumbs to the limitations described above. In addition, we highlight that the analysis of MDP-E is ill-suited to settings with adversarial dynamics, such as the one considered in this work.

### C.1. Overview of OLwS

In this section, we give an overview of the online learning with stationary costs (OLwS) framework for online control and discuss some challenges associated with using it to attain logarithmic regret for online linear control. In OLwS, one defines the stationary costs

$$\lambda_t(\pi; \mathbf{w}) := \ell(x_t^\pi(\mathbf{w}), u_t^\pi(\mathbf{w})), \quad (21)$$

which is the cost suffered that would be suffered at time  $t$  had the policy  $\pi$  had used at all previous rounds ([Abbasi-Yadkori et al., 2013](#); [Anava et al., 2015](#); [Agarwal et al., 2019a](#); [Simchowitz et al., 2020](#)).<sup>6</sup> By construction,  $\lambda_t(\pi; \mathbf{w})$  does not depend on the state of the system. Moreover, if  $\pi$  is an executable policy (i.e.,  $\pi_t(x; \mathbf{w})$  depends only on  $x$  and  $w_{1:t-1}$ ), then  $\lambda_t(\pi; \mathbf{w})$  can be determined exactly at time  $t$ . At each round  $t$ , OLwS selects a policy  $\pi^{(t)}$  to minimize regret on the sequence  $\lambda_t(\pi; \mathbf{w})$ , and follows  $u_t = \pi_t^{(t)}(x; \mathbf{w})$ . The total regret is decomposed as

$$\text{Reg}_T(\pi^{\text{alg}}; \mathbf{w}, \Pi) = \underbrace{\left( \sum_{t=1}^T \ell(x_t^{\text{alg}}, u_t^{\text{alg}}) - \lambda_t(\pi^{(t)}; \mathbf{w}) \right)}_{\text{(stationarization cost)}} + \underbrace{\left( \sum_{t=1}^T \lambda_t(\pi^{(t)}; \mathbf{w}) - \inf_{\pi \in \Pi} \sum_{t=1}^T \lambda_t(\pi; \mathbf{w}) \right)}_{\text{(\lambda-regret)}}. \quad (22)$$

### C.2. Avoiding stationarization cost: Our advantage over OLwS

OLwS optimizes stationary costs  $\lambda_t(\pi)$ , which correspond to the loss suffered by the learner at round  $t$  if policy  $\pi$  had been played for every time up to  $t$ . To relate the stationary costs to the learner’s cost, the OLwS proposes the bounding the following movement cost:

$$\text{movement cost} := \sum_t \|\pi^{(t)} - \pi^{(t-1)}\|, \quad (\text{informal}), \quad (23)$$

To our knowledge, all known applications of OLwS bound the stationarization cost via the movement cost (23). When the movement costs are small, the learner’s state at time  $t$ ,  $x_t^{\text{alg}}$ , is similar to the states that would be obtained by selecting  $\pi^{(t)}$  at all time  $s < t$ , namely  $x_s^\pi$ . [Appendix C](#) explains how the cost (23) arises in more detail. While standard online learning algorithms ensure  $\sqrt{T}$ -movement cost, online gradient descent (OGD) has the property that if  $\lambda_t$  is strongly convex (in a suitable parametrization), the movement cost is  $\log T$  ([Anava et al., 2015](#)). Since OGD also ensures logarithmic regret on the  $\lambda_t$ -sequence, the algorithm ensures logarithmic regret overall.

The natural stationary costs that arise in our problem are exp-concave ([Hazan & Kale, 2011](#)), a property that is stronger than convexity but weaker than strong convexity. Exp-concave functions  $\lambda_t$  are strongly convex in the local geometry induced by  $(\nabla \lambda_t)(\nabla \lambda_t)^\top$ , but not necessarily in other directions. This is sufficient for logarithmic regret, but as we explain in [Appendix C.4](#), known methods cannot leverage this property to ensure logarithmic bounds on the relevant movement cost. Herein lies the advantage of OLwA: by considering the *future* costs of an action (by way of the advantage-proxy  $\widehat{\mathbf{A}}$ ) rather than the stationary costs, we avoid the technical challenge of bounding the movement cost in the elusive exp-concave regime.

### C.3. Applying OLwS to online control

#### C.3.1. POLICY REGRET AND ONLINE CONVEX OPTIMIZATION WITH MEMORY

A useful instantiation of the OLwS paradigm is the policy-regret setting introduced by [Arora et al. \(2012\)](#), which considers stationary costs with finite memory. This work considers online learning with loss functions  $f_t(z_t, \dots, z_{t-h})$ , and defines

<sup>6</sup>Many works also consider “steady state” costs obtained by taking  $t \rightarrow \infty$  for a given policy ([Even-Dar et al., 2009](#); [Abbasi-Yadkori et al., 2014](#); [Cohen et al., 2018](#); [Agarwal et al., 2019b](#)), but this formulation is ill-posed in our setting due to the adversarial dynamics.

policy regret for the iterate sequence  $\{z_t\}_{t=1}^T$  as  $\sum_t f(z_t, \dots, z_{t-h}) - \inf_z \sum_t f_t(z, \dots, z)$ . Algorithms for this setting work with a *unary* loss  $\tilde{f}_t(z) = z \mapsto f_t(z, \dots, z)$ , which can be viewed as a special case of the stationary cost  $\lambda_t$  defined above where  $z \in \mathcal{C}$  encodes a policy and  $f_t(z, \dots, z)$  is the loss suffered if  $z$  had been selected throughout the game. Arora et al. (2012) take this approach in an expert setting and Anava et al. (2015) consider a setting where  $\tilde{f}$  is an arbitrary convex loss, which they call *Online Convex Optimization with Memory*. In this setting, the stationarization cost arises via the decomposition

$$(\text{policy regret}) = \underbrace{\left( \sum_{t=1}^T f_t(z_t, \dots, z_{t-h}) - \tilde{f}_t(z_t) \right)}_{(\text{stationarization cost})} + \underbrace{\left( \sum_{t=1}^T \tilde{f}_t(z_t) - \inf_{z \in \mathcal{C}} \sum_{t=1}^T \tilde{f}_t(z) \right)}_{(\lambda\text{-regret})}. \quad (24)$$

When  $f_t$  are Lipschitz in all arguments, Anava et al. (2015) bound the stationarization cost in terms of movement cost for the iterates. They show

$$\sum_{t=1}^T f_t(z_t, \dots, z_{t-h}) - \tilde{f}_t(z_t) \leq \sum_{t=1}^T |f_t(z_t, \dots, z_{t-h}) - \tilde{f}_t(z_t)| \leq \mathcal{O}(h^2 L) \sum_{t=1}^T \|z_t - z_{t-1}\|,$$

where  $L$  is an appropriate Lipschitz constant. Note that this inequality formalizes (23) for this setting.

Anava et al. (2015) demonstrated that many popular online convex optimization algorithms naturally produce slow-moving iterates, leading to policy regret bounds in (24). In particular, they show that applying online gradient descent on the unary losses leads to  $\text{poly}(h) \cdot \sqrt{T}$ -policy regret when  $\tilde{f}_t$  are convex and Lipschitz, and  $\frac{\text{poly}(h)}{\alpha} \cdot \log T$ -policy regret when  $\tilde{f}_t$  are  $\alpha$ -strongly convex. Notably, Anava et al. (2015) do not show that logarithmic regret is attainable for the more general family of exp-concave losses, which are more natural for the setting in this paper.

### C.3.2. OCO WITH MEMORY FOR ONLINE CONTROL

Now, following Agarwal et al. (2019a), we apply OCO with memory to the linear control setting using the DAP parametrization (Definition 2), where  $\pi^{(t)}$  is given by  $\pi^{(M_t)}$ , for a matrix  $M \in \mathcal{M}_0 = \mathcal{M}(m, R, \gamma)$  selected at time  $t$ . We will specialize the OLWS decomposition (24) and explain how to bound each term.

Agarwal et al. (2019a) show stationary costs  $\lambda_t$  in (21) can be approximated up to arbitrarily accuracy by functions  $\tilde{f}_{t,h}(M)$ , which depend only on the most recent  $m+h$  disturbances  $w_{t-(m+h):t}$ , and where  $h = \text{poly}(\log T, \frac{1}{1-\gamma})$ . They also show that the suffered loss  $\ell(x_t^{\text{alg}}, u_t^{\text{alg}})$  can be approximate via  $f_{t,h}(M_{t-h:t})$ , which also depends on recent disturbances, and which specializes to  $\tilde{f}_{t,h}(M)$  when  $M_{t-h:t} = (M, \dots, M)$ . Precisely, for  $M \in \mathcal{M}$ , define the inputs

$$u_s(M; \mathbf{w}) = \sum_{i=1}^m M^{[i]} w_{s-i}$$

Then the functions  $f_{t,h}(M_{t-h:t})$  and  $\tilde{f}_{t,h}(M)$  take the form

$$\begin{aligned} f_{t,h}(M_{t-h:t}) &:= \ell \left( \alpha_t(\mathbf{w}) + \sum_{i=1}^m \Psi_i u_{t-i}(M_{t-i}; \mathbf{w}), u_t(M_t; \mathbf{w}) - K_\infty \sum_{i=1}^m \Psi_i u_{t-i}(M_{t-i}; \mathbf{w}) \right), \\ \tilde{f}_{t,h}(M) &:= \ell \left( \alpha_t(\mathbf{w}) + \sum_{i=1}^m \Psi_i u_{t-i}(M; \mathbf{w}), u_t(M; \mathbf{w}) - K_\infty \sum_{i=1}^m \Psi_i u_{t-i}(M; \mathbf{w}) \right), \end{aligned} \quad (25)$$

where  $\alpha_t(\mathbf{w})$  is a function of  $w_{1:t}$  and  $A, B$  but *not* of the learner's inputs, and  $\Psi_i = (A - BK_\infty)^{i-1} B$ . With this parameterization, the regret decomposition for OCO with memory takes the form

$$\mathcal{M}_0\text{-Reg}_T(\pi^{\text{alg}}; \mathbf{w}) = \underbrace{\left( \sum_{t=1}^T f_{t,h}(M_{t-h:t}) - \tilde{f}_{t,h}(M_t) \right)}_{(\text{stationarization cost})} + \underbrace{\left( \sum_{t=1}^T \tilde{f}_{t,h}(M_t) - \inf_{\pi \in \mathcal{M}_0} \sum_{t=1}^T \tilde{f}_{t,h}(M) \right)}_{(\lambda\text{-regret})} + \tilde{\mathcal{O}}(1). \quad (26)$$

In this setting  $f_{t,h}(\cdot)$  is Lipschitz so—following arguments from Anava et al. (2015)—we have

$$(\text{stationarization cost}) \leq h^2 \cdot \text{poly}(m, R, (1-\gamma)^{-1}) \cdot \sum_{t=1}^T \|M_t - M_{t-1}\|_F,$$

where the right-hand side is a movement cost for the iterates (formalizing (23)), and where we recall  $h$  is the memory horizon,  $m, R, \gamma$  are the parameters defining the set of DAP controllers  $\mathcal{M}_0$ , and  $\|M_t - M_{t-1}\|_F = \sqrt{\sum_{i \geq 0} \|M_t^{[i]} - M_{t-1}^{[i]}\|_F^2}$  (which induces the standard Euclidean geometry for online gradient descent).

In general, the bound on the movement cost will depend on the choice of regret minimization algorithm. Many natural algorithms ensure bounds on the movement cost which are on the same order as their bounds on regret. For example, exponential weights and online gradient descent ensure  $\sqrt{T}$ -bounds (Even-Dar et al., 2009; Yu et al., 2009; Anava et al., 2015), and for strongly convex losses, FTL and online gradient descent ensure logarithmic movement (Abbasi-Yadkori et al., 2014; Anava et al., 2015).

### C.3.3. THE STATIONARY COSTS FOR DAP ARE EXP-CONCAVE BUT NOT STRONGLY CONVEX

For the DAP parametrization, the functions that naturally arise are exp-concave, but not necessarily strongly convex. To see this, consider the loss  $\ell(x, u) = \|x\|^2 + \|u\|^2$ . We observe that the stationary costs considered by Agarwal et al. (2019a), made explicit in (25), are the sum of two quadratic functions of the form considered in Lemma 2.3, and are thus exp-concave.<sup>7</sup> However,  $\tilde{f}_{t;h}(M)$  is *not* strongly convex in general. For example, if the noise sequence is *constant*, say  $w_t = w_{t-1} = \dots = w_1$ , then  $\nabla u_t(M; \mathbf{w})$  is identical for all  $t$  and thus  $\nabla^2 \tilde{f}_{t;h}(M)$  is a rank-one matrix.

### C.4. Movement costs in general exp-concave online learning

In this section, we explain the challenge of achieving low-movement cost in the OCO-with-memory with framework, which elucidates the broader challenge of relating stationary costs to regret in OLwS. We give an informal sketch for an exp-concave OCO-with-memory setting in which the online Newton step algorithm (Algorithm 2) fails to achieve logarithmic regret. Consider a simple class of functions with scalar domain and length-1 memory:

$$f_t(z_1, z_2) = (1 - (w_t z_1 + w_{t-1} z_2))^2, \quad \tilde{f}_t(z) = f_t(z, z),$$

where  $(w_t)$  are parameters chosen by the adversary. We use the constraint set  $z \in \mathcal{C} := [-1/5, 1/5]$ . Policy regret (paralleling (26)) is given by

$$\begin{aligned} \text{(policy regret)} &= \sum_{t=1}^T f_t(z_t, z_{t-1}) - \inf_z \sum_{t=1}^T \tilde{f}_t(z) \\ &= \underbrace{\left( \sum_{t=1}^T f_t(z_t, z_{t-1}) - \tilde{f}_t(z_t) \right)}_{\text{(stationarization cost)}} + \underbrace{\left( \sum_{t=1}^T \tilde{f}_t(z_t) - \inf_z \sum_{t=1}^T \tilde{f}_t(z) \right)}_{\text{(\lambda-regret)}}. \end{aligned}$$

We now construct a sequence of loss functions where the  $\lambda$ -regret for ONS is logarithmic, but where standard upper bounds on stationary cost can grow as  $\Omega(\sqrt{T})$ . Consider the sequence  $w_t = (-1)^t + \frac{\mu}{2}$ , where  $\mu = 1/\sqrt{T}$ . We see that  $\tilde{f}_t(z) = (1 - \mu z)^2$ . We remark that this function is only  $\mu^2 = 1/T$ -strongly convex, so that the guarantees for strongly convex online gradient descent are vacuous Hazan (2016), necessitating the use of ONS.

Let us see what happens if we try to leverage exp-concavity. From Lemma 2.3,  $\tilde{f}_t(z)$  are  $\frac{1}{4}$ -exp-concave on the set  $\mathcal{C}$ . Hence, if we run ONS (Algorithm 2) with an appropriate learning rate,  $\lambda$ -regret scales logarithmically (Hazan, 2016):

$$\sum_{t=1}^T \tilde{f}_t(z_t) - \min_{z \in \mathcal{C}} \sum_{t=1}^T \tilde{f}_t(z) \leq \mathcal{O}(\log T).$$

Let us now turn to the stationarization cost,  $\sum_{t=1}^T f_t(z_t, z_{t-1}) - \tilde{f}_t(z_t)$ . The approach of Anava et al. (2015), is to bound the per-step errors,  $|f_t(z_t, z_{t-1}) - \tilde{f}_t(z_t)|$ . We can directly see that

$$\begin{aligned} f_t(z_t, z_{t-1}) - \tilde{f}_t(z_t) &= (1 - w_t z_t - w_{t-1} z_{t-1})^2 - (1 - w_t z_t - w_{t-1} z_{t-1})^2 \\ &= -2w_{t-1}(1 - w_t z_t)^\top (z_t - z_{t-1}) + w_{t-1}^2 (z_{t-1} + z_t)(z_t - z_{t-1})^2 \\ &= (-2w_{t-1}(1 - w_t z_t)^\top + w_{t-1}^2 (z_{t-1} + z_t))(z_t - z_{t-1}). \end{aligned}$$

<sup>7</sup>The sum of two  $\alpha$ -exp-concave functions is  $\frac{\alpha}{2}$ -exp-concave. For a proof, observe that  $(\nabla(f+g))^\top (\nabla(f+g))^\top \leq 2(\nabla f)(\nabla f)^\top + 2(\nabla g)(\nabla g)^\top$ . Hence, if  $f$  and  $g$  are  $\alpha$ -exp concave, we have  $\nabla^2(f+g) \geq \alpha(\nabla f)(\nabla f)^\top + \alpha(\nabla g)(\nabla g)^\top \geq \frac{\alpha}{2}(\nabla(f+g))^\top (\nabla(f+g))^\top$ .

For  $\mu$  sufficiently small and  $z \in \mathcal{C}$ , we can check that  $|(-2w_{t-1}(1-w_t z_t)^\top + w_{t-1}^2(z_{t-1} + z_t))| \geq \frac{1}{16}$ , so that

$$|f_t(z_t, z_{t-1}) - \tilde{f}_t(z_t)| \geq \frac{|z_t - z_{t-1}|}{16}.$$

Thus, we have

$$\sum_{t=1}^T |f_t(z_t, z_{t-1}) - \tilde{f}_t(z_t)| \geq \frac{1}{16} \cdot (\text{movement cost}), \quad \text{where} \quad (\text{movement cost}) = \frac{1}{6} \sum_{t=1}^T |z_t - z_{t-1}|.$$

We now show that this movement cost is large. For simplicity, we keep our discussion informal to avoid navigating the projection step in ONS. Without projections, we have

$$|z_t - z_{t-1}| = \left| \epsilon + \sum_{s=1}^{t-1} \nabla^2 \tilde{f}_s(z_s) \right|^{-1} (\nabla \tilde{f}_{t-1}(z_{t-1})).$$

Observe that for each  $z \in \mathcal{C}$ ,  $\nabla^2 \tilde{f}_s(z) = \mu^2 = 1/T$ , so that we have  $\epsilon + \sum_{s=1}^{t-1} \nabla^2 \tilde{f}_s(z_s) = (1 + \epsilon)$ . On the other hand, for  $z \in \mathcal{C}$ , we can lower bound  $|\nabla \tilde{f}_{t-1}(z)| \geq \frac{\mu}{2} = \frac{1}{2\sqrt{T}}$ . Hence,

$$(\text{movement cost}) = \frac{1}{16} \sum_{t=1}^T |z_t - z_{t-1}| \geq \frac{\sqrt{T}}{32(1 + \epsilon)}.$$

Here, we note that the standard implementation prescribes  $\epsilon$  to be constant, giving us  $\Omega(\sqrt{T})$  movement. Moreover, increasing  $\epsilon$  will degrade the corresponding regret bound, preventing logarithmic combined regret. Note that increasing  $\epsilon$  to  $1/T^{1/4}$  will partially mitigate the movement cost, but at the expense of increasing the regret on the  $\tilde{f}_t$  sequence.

### C.5. Comparison with MDP-E

MDP-E (Even-Dar et al., 2009) is an instantiation of OLWS for MDPs with known non-adversarial dynamics and time varying adversarial losses  $\ell_t$ . In this setting the stationary costs  $\lambda_t(\pi)$  represent the long-term costs of a policy  $\pi$  under the loss  $\ell_t$  (if one prefers, the loss can be treated as fixed, and  $w_t$  can encode loss information). To achieve low regret on the  $\lambda_t$ -sequence, MDP-E maintains policy iterates  $\{\pi_x^{(t)}\}$  for all states  $x$ , and selects its action according to the policy for the corresponding current state:

$$u_t^{\text{alg}} \leftarrow \pi_{x_t^{\text{alg}}}^{(t)}(x_t^{\text{alg}}).$$

The policy sequence  $\pi_x^{(t)}$  is selected to minimize regret on a certain  $Q$ -function:  $\lambda_{t,x}(\pi) : \pi \mapsto Q^\pi(x, \pi(x))$  (here, policies and  $Q$ -functions are regarded as stationary). Under the assumption that the dynamics under benchmark policies are also stationary, achieving low regret on each  $\{\lambda_{t,x}\}$ -sequence simultaneously for all  $x$  ensures low  $\lambda$ -regret (in the sense of Eq.(22)) over the trajectory  $x_t^{\text{alg}}$ .<sup>8</sup> As a consequence, MDP-E is ill-suited to settings with adversarially changing dynamics. Since OLWA considers  $Q$ -functions and advantages defined with respect to an *fixed* policy  $\pi^*$ , it does not require benchmark policies to have stationary dynamics (which is important, since our adversarial disturbance setting does not have stationary dynamics).

Moreover, like the stationary costs, the functions  $\lambda_{t,x}(\pi)$  describe long-term performance under  $\pi$ , and still need to be related to the learner's realized trajectory, typically via a bound on the movement cost of the policies. As described earlier, the analysis of OLWA does not require bounding the movement cost.

## D. Basic technical results

### D.1. Structural results for LQR

In this section we provide a number of useful structural properties for the optimal controller for linear dynamical systems with quadratic costs and arbitrary bounded disturbances. Even though the results in this section concern the optimal finite-horizon controllers, we prove bounds on various regularity properties for the controllers that depend only on control-theoretic

<sup>8</sup>See the proof of Even-Dar et al. (2009, Theorem 5), which uses that the induced state distribution does not change with  $t$ .



parameters for the optimal infinite-horizon controller in the noiseless setting, which is an intrinsic parameter of the dynamical system. All proofs are deferred to [Appendix D.1.2](#).

For the results in this section and the remainder of the appendix we use that  $A_{\text{cl},\infty}$  is  $(\kappa_\infty, \gamma_\infty)$ -strongly stable.

**Lemma D.1.** Let  $\gamma_\infty = \|I - P_\infty^{-1/2} R_x P_\infty^{-1/2}\|_{\text{op}}^{1/2}$ , and  $\kappa_\infty = \|P_\infty^{1/2}\|_{\text{op}} \|P_\infty^{-1/2}\|_{\text{op}}$ . Then the closed loop system  $A_{\text{cl},\infty}$  is  $(\kappa_\infty, \gamma_\infty)$ -strongly stable.

**Proof of Lemma D.1.** Recall ([Bertsekas, 2005](#)) that the infinite-horizon Lyapunov matrix  $P_\infty$  satisfies the equation

$$A_{\text{cl},\infty}^\top P_\infty A_{\text{cl},\infty} - P_\infty + R_x = 0.$$

Since  $P_\infty > 0$ , if we set  $H = P_\infty^{-1/2}$  and  $L = P_\infty^{1/2} A_{\text{cl},\infty} P_\infty^{-1/2}$ , we deduce from this expression that

$$L^\top L - I + P_\infty^{-1/2} R_x P_\infty^{-1/2} = 0,$$

and in particular  $\|L\|_{\text{op}}^2 \leq \|I - P_\infty^{-1/2} R_x P_\infty^{-1/2}\|_{\text{op}} < 1$ . □

**Lemma D.2.** Let  $A$  be  $(\kappa, \gamma)$ -strongly stable. Then for any  $i \geq 0$ ,

$$\|A^i\|_{\text{op}} \leq \kappa \gamma^i.$$

**Proof of Lemma D.2.** Let  $A = H L H^{-1}$ , where  $H$  and  $L$  witness the strong stability property. Then we have

$$\|A^i\|_{\text{op}} \leq \kappa \|L^i\|_{\text{op}} \leq \kappa \gamma^i.$$

□

**Additional notation.** For the remainder of the appendix we adopt the following notation. We let  $H_{\text{cl},\infty}$  and  $L_{\text{cl},\infty}$  denote the matrices that witness strong stability of  $A_{\text{cl},\infty}$ , so that  $A_{\text{cl},\infty} = H_{\text{cl},\infty} L_{\text{cl},\infty} H_{\text{cl},\infty}^{-1}$  and we have  $\|H_{\text{cl},\infty}\|_{\text{op}} \cdot \|H_{\text{cl},\infty}^{-1}\|_{\text{op}} \leq \kappa_\infty$  and  $\|L_{\text{cl},\infty}\|_{\text{op}} \leq \gamma_\infty < 1$ . We also define  $L_{\text{cl},t} = H_{\text{cl},\infty}^{-1} A_{\text{cl},t} H_{\text{cl},\infty}$ , where we recall that  $(A_{\text{cl},t})_{t=1}^T$  denote the closed-loop dynamics arising from the Riccati recursion. We define  $A_{\text{cl},i \rightarrow t} = A_{\text{cl},t} A_{\text{cl},t-1} \cdots A_{\text{cl},i+1}$ , with the convention that  $A_{\text{cl},t \rightarrow t} = I$ . Finally, we define  $\bar{\gamma}_\infty = \frac{1}{2}(1 + \gamma_\infty)$ ,  $\Delta_{\text{stab}} = 4 \cdot \beta_\star \Psi_\star^2 \Gamma_\star \log(2 \Psi_\star \Gamma_\star \kappa_\infty (1 - \gamma_\infty)^{-1})$ , and  $T_{\text{stab}} = T - \Delta_{\text{stab}}$ .

#### D.1.1. PROPERTIES OF THE OPTIMAL POLICY

Recall that [Theorem 3](#) characterizes the optimal unconstrained policy given full knowledge of  $w$ . Rather than directly proving this theorem, we state and prove a more general version, [Theorem 4](#), which generalizes the characterization to the setting of [Appendix B.3](#) in which losses include adversarially chosen targets. The optimal policy for this setting is defined as follows.

**Definition 7** (Optimal policy, Q-function, advantage). Assume  $a_T, b_T = 0$ , and recall that  $\bar{w} = (w_{1:T}, a_{1:T}, b_{1:T})$ . Define  $\mathbf{Q}_T^\star(x, u; \bar{w}) = \ell(x, u)$ ,  $\pi_t^\star(x; \bar{w}) = \min_u \mathbf{Q}_t^\star(x, u) = 0$ , and  $\mathbf{V}_T^\star(x; \bar{w}) = \ell(x, 0)$ . For each  $t < T$  define

$$\begin{aligned} \mathbf{Q}_t^\star(x, u; \bar{w}) &= \|x - a_t\|_Q^2 + \|u - b_t\|_R^2 + \mathbf{V}_{t+1}^\star(Ax + Bu + w_t; \bar{w}), \\ \pi_t^\star(x; \bar{w}) &= \arg \min_{u \in \mathbb{R}^{d_u}} \mathbf{Q}_t^\star(x, u; \bar{w}), \\ \mathbf{V}_t^\star(x; \bar{w}) &= \min_{u \in \mathbb{R}^d} \mathbf{Q}_t^\star(x, u; \bar{w}) = \mathbf{Q}_t^\star(x, \pi_t^\star(x; \bar{w}); \bar{w}). \end{aligned}$$

Finally, define  $\mathbf{A}_t^\star(u; x, \bar{w}) := \mathbf{Q}_t^\star(x, u; \bar{w}) - \mathbf{Q}_t^\star(x, \pi_t^\star(x; \bar{w}); \bar{w})$ .

**Theorem 4** (Generalization of [Theorem 3](#)). Set  $\bar{w}_{t:T} = (w_{t:T}, a_{t:T}, b_{t:T})$ . For each time  $t$ , we have  $\mathbf{V}_t^\star(x; \bar{w}) = \|x\|_{P_t}^2 + 2\langle x, c_t(\bar{w}_{t:T}) \rangle + f_t(\bar{w}_{t:T})$ , where  $f_t$  is a function that does not depend on the state  $x$  and  $c_t$  is defined recursively with  $c_{T+1} = 0$  and

$$c_t(\bar{w}_{t:T}) = (A - BK_t)^\top (P_{t+1} w_t + c_{t+1}(\bar{w}_{t+1:T})) + K_t^\top R_u b_t - R_x a_t.$$

Moreover, if we define

$$q_t^*(\bar{w}_{t:T}) = \Sigma_t^{-1} \left( -R_u b_t + B^\top \sum_{i=t}^{T-1} \left( \prod_{j=t+1}^i A_{\text{cl},j}^\top \right) P_{i+1} w_i + B^\top \sum_{i=t+1}^{T-1} \left( \prod_{j=t+1}^{i-1} A_{\text{cl},j}^\top \right) (K_i^\top R_u b_i - R_x a_i) \right), \quad (27)$$

then the optimal controller is given by  $\pi_t^*(x; \bar{w}) = -K_t x - q_t^*(\bar{w}_{t:T})$ .

**Lemma D.3.** For all  $\tau_1 \leq \tau_2$ , we have

$$\left\| \prod_{t=\tau_1}^{\tau_2} A_{\text{cl},t}^\top \right\|_{\text{op}} \leq \sqrt{\frac{\|P_\infty\|_{\text{op}}}{\lambda_{\min}(R_x)}} \leq \beta_*^{1/2} \Gamma_*^{1/2}.$$

**Lemma D.4.** Let  $\Delta_{\text{stab}} = 4 \cdot \beta_* \Psi_*^2 \Gamma_* \log(2\Psi_* \Gamma_* \kappa_\infty (1 - \gamma_\infty)^{-1}) = \tilde{\mathcal{O}}(\beta_* \Psi_*^2 \Gamma_*)$ , and let  $\bar{\gamma}_\infty = \frac{1}{2}(1 + \gamma_\infty)$ . Then it holds that

$$\|L_{\text{cl},t}\|_{\text{op}} \leq \bar{\gamma}_\infty < 1, \quad \forall t \leq T_{\text{stab}} := T - \Delta_{\text{stab}}.$$

**Lemma D.5.** Let  $\tau_1 \leq \tau_2$  be fixed. Then we have

$$\left\| \prod_{t=\tau_1}^{\tau_2} A_{\text{cl},t}^\top \right\|_{\text{op}} \leq \kappa_\infty \left\| \prod_{t=\tau_1}^{\tau_2} L_{\text{cl},t}^\top \right\|_{\text{op}} \leq \kappa_\infty^2 \beta_*^{1/2} \Gamma_*^{1/2} \cdot \bar{\gamma}_\infty^{\tau_2 \wedge T_{\text{stab}} - \tau_1 \wedge T_{\text{stab}}}.$$

**Lemma D.6.** Let  $w$  be any sequence with  $\|w_t\| \leq 1$ . Let  $t \in [T]$  and  $h \geq 0$  be given. Then we have

$$\|q_t^*(w_{t:T})\| \vee \|q_{t,t+h}^*(w_{t:t+h})\| \leq \tilde{\mathcal{O}}\left(\beta_*^{5/2} \Psi_*^3 \Gamma_*^{5/2} \kappa_\infty^2 (1 - \gamma_\infty)^{-1}\right) =: D_{q^*}, \quad (28)$$

and

$$\|q_{\infty,h}^*(w_{t:t+h})\| \leq \beta_* \Psi_* \Gamma_* \kappa_\infty (1 - \gamma_\infty)^{-1} =: D_{q_\infty^*}. \quad (29)$$

**Lemma D.7.** Let policies  $\pi_t(x; w) = -K_\infty x - q_t(w)$  and  $\hat{\pi}_t(x; w) = -K_t x - q_t(w)$  be given, where  $q_t$  is arbitrary. Then the states for both controllers are given by

$$x_{t+1}^\pi(w) = \sum_{i=1}^t A_{\text{cl},\infty}^{t-i} w_i - \sum_{i=1}^t A_{\text{cl},\infty}^{t-i} B q_i(w) \quad \text{and} \quad x_{t+1}^{\hat{\pi}}(w) = \sum_{i=1}^t A_{\text{cl},i \rightarrow t} w_i - \sum_{i=1}^t A_{\text{cl},i \rightarrow t} B q_i(w).$$

**Lemma D.8.** Let  $\alpha \geq 1$  be given. Define  $\Delta = C \cdot \beta_* \Psi_*^2 \Gamma_* \log(\kappa_\infty^2 \Psi_* \Gamma_* (1 - \gamma_\infty)^{-1} \cdot \alpha T^3)$ , where  $C > 0$  is a numerical constant. If  $C$  is sufficiently large, then for every  $t \leq T - \Delta \leq T_{\text{stab}}$  we are guaranteed that

$$\|K_t - K_\infty\|_{\text{op}} \leq \frac{1}{\kappa_\infty^2 \Psi_* \cdot (\alpha T^3)}, \quad \text{and} \quad \|A_{\text{cl},i \rightarrow t} - A_{\text{cl},\infty}^{t-i}\|_{\text{op}} \leq \frac{1}{\alpha T^2} \quad \forall t \leq T - \Delta. \quad (30)$$

**Lemma D.9.** Let policies  $\pi_t(x; w) = -K_\infty x - q_t(w)$  and  $\hat{\pi}_t(x; w) = -K_t x - q_t(w)$  be given, where  $q_t$  is arbitrary but satisfies  $\|q_t\| \leq D_q$  for some  $D_q \geq 1$ . Then for all  $t \in [T]$ , we have

$$\|x^\pi(w_t)\| \leq 2\kappa_\infty \Psi_* (1 - \gamma_\infty)^{-1} D_q, \quad \text{and} \quad \|u_{t+1}^\pi(w)\| \leq 3\kappa_\infty \beta_* \Psi_*^3 \Gamma_* (1 - \gamma_\infty)^{-1} D_q,$$

as well as

$$\|x_{t+1}^{\hat{\pi}}(w)\| \leq \tilde{\mathcal{O}}\left(\kappa_\infty^2 \beta_*^{3/2} \Psi_*^3 \Gamma_*^{3/2} (1 - \gamma_\infty)^{-1} \cdot D_q\right),$$

and

$$\|u_{t+1}^{\hat{\pi}}(w)\| \leq \tilde{\mathcal{O}}\left(\kappa_\infty^2 \beta_*^{5/2} \Psi_*^5 \Gamma_*^{5/2} (1 - \gamma_\infty)^{-1} \cdot D_q\right).$$

## D.1.2. PROOFS FROM APPENDIX D.1.1

**Proof of Theorem 4.** We first prove that the identity for the value function,

$$\mathbf{V}_t^*(x; \bar{w}_{t:T}) = \|x\|_{P_t}^2 + 2\langle x, c_t(\bar{w}_{t:T}) \rangle + f_t(\bar{w}_{t:T}),$$

holds by induction. Observe that at time  $T$  we indeed have  $\mathbf{V}_T^*(x, w_T) = \|x\|_{R_x}^2 = \|x\|_{P_T}^2$ , where we recall  $a_T, b_T = 0$  by assumption. Now suppose, that at time  $t+1$  we have

$$\mathbf{V}_{t+1}^*(x; \bar{w}_{t+1:T}) = \|x\|_{P_{t+1}}^2 + 2\langle x, c_{t+1}(\bar{w}_{t+1:T}) \rangle + f_{t+1}(\bar{w}_{t+1:T}).$$

We prove that the same holds for time  $t$  using the following lemma.

**Lemma D.10.** Let  $P_1 > 0$ ,  $c_1$ ,  $a_0$ , and  $b_0$  be given and define  $V_1(x) = \|x\|_{P_1}^2 + 2\langle x, c_1 \rangle$  and

$$V_0(x, w, a_0, b_0) = \|x - a_0\|_{R_x}^2 + \min_u \left\{ \|u - b_0\|_{R_u}^2 + V_1(Ax + Bu + w) \right\}. \quad (31)$$

Then we have

$$V_0(x, w, a_0, b_0) = \|x\|_{P_0}^2 + 2\langle x, c_0 \rangle + f(w, a_0, b_0, c_1), \quad (32)$$

where

$$\begin{aligned} P_0 &= R_x + A^\top P_1 A - A^\top P_1 B \Sigma_0^{-1} B^\top P_1 A, \\ \Sigma_0 &= R_u + B^\top P_1 B, \\ K_0 &= \Sigma_0^{-1} B^\top P_1 A, \\ c_0 &= (A - BK_0)^\top (P_1 w + c_1) + K_0^\top R_u b_0 - R_x a_0. \end{aligned}$$

Furthermore, letting  $u^*$  denote the minimizer in (31), we have

$$u^* = -\Sigma_0^{-1} B^\top (P_1 (Ax + w) + c_1 - R_u b_0) = -K_0 x - \Sigma_0^{-1} (B^\top (P_1 w + c_1) - R_u b_0). \quad (33)$$

**Proof of Lemma D.10.** Since the minimization problem in (31) is strongly convex with respect to  $u$ , we conclude from first-order conditions that

$$B^\top P_1 (Ax + Bu^* + w) + R_u (u^* - b_0) + B^\top c_1 = 0,$$

Rearranging,

$$u^* = -(R_u + B^\top P_1 B)^{-1} (B^\top P_1 A + B^\top c_1 + P_1 w - R_u b_0) = -K_0 x - \Sigma_0^{-1} (B^\top (P_1 w + c_1) - R_u b_0),$$

which proves (33). Next, observe that for any  $u$ , we have

$$\begin{aligned} \|u - b_0\|_{R_u}^2 + V_1(Ax + Bu + w) &= u^\top \Sigma_0 u + 2u^\top (B^\top (P_1 Ax + P_1 w + c_1) - R_u b_0) \\ &\quad + x^\top A^\top P_1 Ax + 2x^\top A^\top (P_1 w + c_1) + g(w, c_1, b_0), \end{aligned}$$

where  $g(w, c_1, b_0)$  is a function of  $w, c_1$ , and  $b_0$  but not  $x$  or  $w$ . Next, observe that for any  $\Sigma > 0$  and  $v$ ,  $\min_u u^\top \Sigma u + 2\langle v, u \rangle = -v^\top \Sigma^{-1} v$ . Hence,

$$\begin{aligned} &\min_u \|u - b_0\|_{R_u}^2 + V_1(Ax + Bu + w) \\ &= -\|B^\top (P_1 Ax + P_1 w + c_1) - R_u b_0\|_{\Sigma_0^{-1}}^2 \\ &\quad + x^\top A^\top P_1 Ax + 2x^\top A^\top (P_1 w + c_1) + g(w, c_1, b_0), \\ &= x^\top A^\top (P_1 - P_1 B \Sigma_0^{-1} B^\top P_1) Ax \\ &\quad - 2(B^\top (P_1 w + c_1) - R_u b_0)^\top \Sigma_0^{-1} B^\top P_1 Ax + 2(P_1 w + c_1)^\top Ax + \tilde{g}(w, c_1, b_0), \end{aligned}$$

for an appropriate function  $\tilde{g}$ . We can further simplify the part of this expression that is linear in  $x$  to

$$\begin{aligned} & -2(B^\top(P_1 w + c_1) - R_u b_0)^\top \Sigma_0^{-1} B^\top P_1 A x + 2(P_1 w + c_1)^\top A x \\ & = -2(B^\top(P_1 w + c_1) - R_u b_0)^\top K_0 x + 2(P_1 w + c_1)^\top A x \\ & = 2(P_1 w + c_1)^\top (A - B K_0) x + 2b_0^\top R_u K_0 x, \end{aligned}$$

which yields

$$\begin{aligned} \min_u \left\{ \|u - b_0\|_{R_u}^2 + V_1(Ax + Bu + w) \right\} & = x^\top A^\top (P_1 - P_1 B^\top \Sigma_0^{-1} B P_1) A x \\ & \quad + 2(P_1 w + c_1)^\top (A - B K_0) x + 2b_0^\top R_u K_0 x + \tilde{g}(w, c_1, b_0). \end{aligned}$$

Therefore,

$$V_0(x, w, a_0, b_0) = x^\top P_0 x - 2a_0^\top R_x x + 2b_0^\top R_u K_0 x + 2(P_1 w + c_1)^\top (A - B K_0) x + \tilde{g}(w, c_0, b_0) + \|a_0\|_{R_x}^2.$$

This yields the lemma with  $c_0 = (A - B K_0)^\top (P_1 w + c_1) + K_0^\top R_u b_0 - R_x a_0$ , and  $f(w, a_0, b_0, c_1) = \tilde{g}(w, c_1, b_0) + \|a_0\|_{R_x}^2$ .  $\square$

Applying [Lemma D.10](#) with  $P_1 = P_{t+1}$  and  $c_1 = c_{t+1}(\bar{w}_{t+1:T})$ , and using the definition of  $\mathbf{Q}_t^*$  from [Definition 3](#) we see that we indeed have

$$\mathbf{V}_t^*(x; \bar{w}_{t:T}) = \|x\|_{P_t}^2 + 2\langle x, c_t(\bar{w}_{t:T}) \rangle + f_t(\bar{w}_{t:T}),$$

and that

$$\pi_t^*(x; \mathbf{w}) = -K_t x - \Sigma_t^{-1} (B^\top (P_{t+1} w_t + c_{t+1}(\bar{w}_{t:T})) - R_u b_t).$$

Unfolding the recursion, we also see that for each  $t$ ,

$$c_t(\bar{w}_{t:T}) = \sum_{i=t}^{T-1} \left( \prod_{j=t}^i A_{\text{cl},j}^\top \right) P_{i+1} w_i + \sum_{i=t}^{T-1} \left( \prod_{j=t}^{i-1} A_{\text{cl},j}^\top \right) (K_i^\top R_u b_i - R_x a_i),$$

with the convention that the empty product is equal to 1. Thus, we indeed have

$$q_t^*(\bar{w}_{t:T}) = \Sigma_t^{-1} \left( -R_u b_t + B^\top \sum_{i=t}^{T-1} \left( \prod_{j=t+1}^i A_{\text{cl},j}^\top \right) P_{i+1} w_i + B^\top \sum_{i=t+1}^{T-1} \left( \prod_{j=t+1}^{i-1} A_{\text{cl},j}^\top \right) (K_i^\top R_u b_i - R_x a_i) \right).$$

$\square$

**Proof of [Lemma D.3](#).** Consider the noiseless LQR setup where

$$x_{t+1} = A x_t + B u_t.$$

The optimal policy for this setup is given by  $u_t = -K_t x$ . For each  $t \leq s$ , let  $x_s^*(x_t = x)$  and  $u_s^*(x_t = x)$  respectively denote the value of the state  $x_s$  and control  $u_s$  if we begin with  $x_t = x$  and follow the optimal policy until time  $s$ . Let  $\mathbf{V}_t(x)$  denote the optimal finite-horizon value function for this noiseless setup, which satisfies

$$\mathbf{V}_t(x) \leq \langle P_\infty x, x \rangle,$$

and

$$\mathbf{V}_t(x) = \sum_{s=t}^T \|x_s^*(x_t = x)\|_{R_x}^2 + \|u_s^*(x_t = x)\|_{R_u}^2.$$

Note that  $(x_s^*(x_t = x))^\top = x^\top \prod_{r=t}^{s-1} A_{\text{cl},r}^\top$ , and that we have in particular that

$$\|x_s^*(x_t = x)\|_{R_x}^2 \leq \langle P_\infty x, x \rangle,$$

and so  $\|x_s^*(x_t = x)\|^2 \leq \langle P_\infty x, x \rangle / \lambda_{\min}(R_x)$ . Choosing  $t = \tau_1$  and  $s = \tau_2 + 1$ , we have

$$\left\| \prod_{t=\tau_1}^{\tau_2} A_{\text{cl},t}^\top x \right\|^2 \leq \frac{\langle P_\infty x, x \rangle}{\lambda_{\min}(R_x)}.$$

The result now follows by recalling the definition of the spectral norm.  $\square$

**Proof of Lemma D.4.** First observe that for any  $t$ , we have

$$\begin{aligned} \|L_{\text{cl},t}\|_{\text{op}} &\leq \|L_{\text{cl},\infty}\|_{\text{op}} + \|L_{\text{cl},t} - L_{\text{cl},\infty}\|_{\text{op}} \leq \gamma_\infty + \kappa_\infty \|A_{\text{cl},t} - A_{\text{cl},\infty}\|_{\text{op}} \\ &\leq \gamma_\infty + \kappa_\infty \|B\|_{\text{op}} \|K_t - K_\infty\|_{\text{op}}. \end{aligned}$$

To bound the error between the infinite-horizon optimal controller  $K_\infty$  and the finite-horizon controller  $K_t$ , we appeal to the following lemma.

**Lemma D.11** (Dean et al. (2018), Lemma E.6; Lincoln & Rantzer (2006), Proposition 1). Define

$$\nu = 2\|P_\infty\|_{\text{op}} \cdot \left( \frac{\|A\|_{\text{op}}^2}{\lambda_{\min}(R_x)} \vee \frac{\|B\|_{\text{op}}^2}{\lambda_{\min}(R_u)} \right).$$

Then for all  $0 \leq t \leq T$ , it holds that

$$\|K_t - K_\infty\|_{\text{op}} \leq \|P_t - P_\infty\|_{\Sigma_t}^2 \leq \|P_\infty\|_{\text{op}} \left(1 + \frac{1}{\nu}\right)^{-(T-t+1)}.$$

In particular, for  $\nu_\star := 2\beta_\star \Psi_\star^2 \Gamma_\star$ , we have

$$\|K_t - K_\infty\|_{\text{op}} \leq \|P_t - P_\infty\|_{\Sigma_t}^2 \leq \Gamma_\star \exp\left(-\frac{1}{2\nu_\star}(T-t+1)\right).$$

Lemma D.11 implies that if we set  $\Delta = 2\nu_\star \log(\|P_\infty\|_{\text{op}}/\varepsilon)$ , we have  $\|K_t - K_\infty\|_{\text{op}} \leq \varepsilon$  for all  $t \leq T - \Delta$ . To get the final result, we choose  $\varepsilon = \frac{1}{2}(1 - \gamma_\infty)/(\kappa_\infty(1 \vee \|B\|_{\text{op}}))$ .  $\square$

**Proof of Lemma D.5.** Assume for now that  $\tau_2 \leq T_{\text{stab}}$ ; if not, the result follows trivially from Lemma D.3. We write

$$\left\| \prod_{t=\tau_1}^{\tau_2} L_{\text{cl},t}^\top \right\|_{\text{op}} \leq \left\| \prod_{t=\tau_1}^{\tau_2 \wedge T_{\text{stab}}} L_{\text{cl},t}^\top \right\|_{\text{op}} \cdot \left\| \prod_{t=T_{\text{stab}}+1}^{\tau_2} L_{\text{cl},t}^\top \right\|_{\text{op}}.$$

For the first term, we have

$$\left\| \prod_{t=\tau_1}^{\tau_2 \wedge T_{\text{stab}}} L_{\text{cl},t}^\top \right\|_{\text{op}} \leq \prod_{t=\tau_1}^{\tau_2 \wedge T_{\text{stab}}} \|L_{\text{cl},t}^\top\|_{\text{op}} \leq \bar{\gamma}_\infty^{\tau_2 \wedge T_{\text{stab}} - \tau_1},$$

using Lemma D.4. The second term is bounded using Lemma D.3 as

$$\left\| \prod_{t=T_{\text{stab}}+1}^{\tau_2} L_{\text{cl},t}^\top \right\|_{\text{op}} \leq \kappa_\infty \left\| \prod_{t=T_{\text{stab}}+1}^{\tau_2} A_{\text{cl},t}^\top \right\|_{\text{op}} \leq \kappa_\infty \beta_\star^{1/2} \Gamma_\star^{1/2}.$$

$\square$

**Proof of Lemma D.6.** We first bound  $q_t^\star$  and  $q_{t:t+h}^\star$ . Let  $t \in [T]$  be fixed. Then we have

$$\begin{aligned} \|q_t^\star(w_{t:T})\| &= \left\| \sum_{i=t}^{T-1} \Sigma_t^{-1} B^\top \left( \prod_{j=t+1}^i A_{\text{cl},j}^\top \right) P_{i+1} w_i \right\| \\ &\leq \|\Sigma_t^{-1}\|_{\text{op}} \|B\|_{\text{op}} \max_{i>t} \|P_{i+1}\|_{\text{op}} \sum_{i=t}^{T-1} \left\| \prod_{j=t+1}^i A_{\text{cl},j}^\top \right\|_{\text{op}} \\ &\leq \beta_\star \Psi_\star \Gamma_\star \left( 1 + \sum_{i=t+1}^{T-1} \left\| \prod_{j=t+1}^i A_{\text{cl},j}^\top \right\|_{\text{op}} \right). \end{aligned}$$



Furthermore, the same argument shows that we have

$$\|q_{t;t+h}^*(w_{t:t+h})\| \leq \beta_* \Psi_* \Gamma_* \left( 1 + \sum_{i=t+1}^{T-1} \left\| \prod_{j=t+1}^i A_{\text{cl},j}^\top \right\|_{\text{op}} \right),$$

as well. If  $i > T_{\text{stab}}$ , we trivially bound the summand as  $\beta_*^{1/2} \Gamma_*^{1/2}$  using [Lemma D.3](#). Otherwise, we have  $t+1 \leq i \leq T_{\text{stab}}$ , and we use [Lemma D.5](#), which gives

$$\left\| \prod_{j=t+1}^i A_{\text{cl},j}^\top \right\|_{\text{op}} \leq \kappa_\infty^2 \beta_*^{1/2} \Gamma_*^{1/2} \cdot \bar{\gamma}_\infty^{i-(t+1)}.$$

Summing across the two cases, we have

$$\begin{aligned} \|q_t^*(w_{t:T})\| &\leq \beta_* \Psi_* \Gamma_* \left( 1 + \beta_*^{1/2} \Gamma_*^{1/2} \Delta_{\text{stab}} + \kappa_\infty^2 \beta_*^{1/2} \Gamma_*^{1/2} \sum_{i=t+1}^{T_{\text{stab}}} \bar{\gamma}_\infty^{i-(t+1)} \right) \\ &\leq \beta_* \Psi_* \Gamma_* \left( 1 + \beta_*^{1/2} \Gamma_*^{1/2} \Delta_{\text{stab}} + \kappa_\infty^2 \beta_*^{1/2} \Gamma_*^{1/2} \sum_{i=0}^{\infty} \bar{\gamma}_\infty^i \right) \\ &\leq \beta_* \Psi_* \Gamma_* \left( 1 + \beta_*^{1/2} \Gamma_*^{1/2} \Delta_{\text{stab}} + 2\kappa_\infty^2 \beta_*^{1/2} \Gamma_*^{1/2} (1 - \bar{\gamma}_\infty)^{-1} \right) \\ &\leq 2\beta_*^{3/2} \Psi_* \Gamma_*^{3/2} (\Delta_{\text{stab}} + \kappa_\infty^2 (1 - \bar{\gamma}_\infty)^{-1}). \end{aligned}$$

Recalling the definition of  $\Delta_{\text{stab}}$ , this is at most

$$\tilde{\mathcal{O}}\left(\beta_*^{5/2} \Psi_*^3 \Gamma_*^{5/2} \kappa_\infty^2 (1 - \gamma_\infty)^{-1}\right).$$

To bound  $q_{\infty,h}^*$ , recall that we have

$$q_{\infty,h}^*(w_{h+1}) := \sum_{i=1}^{h+1} \Sigma_\infty^{-1} B^\top (A_{\text{cl},\infty}^\top)^{i-1} P_\infty w_i.$$

It immediately follows that we have

$$\|q_{\infty,h}^*(w_{h+1})\| \leq \left\| \sum_{i=1}^{h+1} \Sigma_\infty^{-1} B^\top (A_{\text{cl},\infty}^\top)^{i-1} P_\infty w_i \right\| \leq \beta_* \Psi_* \Gamma_* \sum_{i=1}^{h+1} \|A_{\text{cl},\infty}^{i-1}\|_{\text{op}}.$$

We may further upper bound this by

$$\kappa_\infty \beta_* \Psi_* \Gamma_* \sum_{i=1}^{h+1} \|L_{\text{cl},\infty}^{i-1}\|_{\text{op}} \leq \kappa_\infty \beta_* \Psi_* \Gamma_* \sum_{i=1}^{h+1} \gamma_\infty^{i-1} \leq \kappa_\infty \beta_* \Psi_* \Gamma_* (1 - \gamma_\infty)^{-1}.$$

□

**Proof of Lemma D.8.** By a change of variables, we have

$$\|A_{\text{cl},i \rightarrow t} - A_{\text{cl},\infty}^{t-i}\|_{\text{op}} \leq \kappa_\infty \|L_{\text{cl},i \rightarrow t} - L_{\text{cl},\infty}^{t-i}\|_{\text{op}}.$$

Let us drop the “cl” subscript to keep notation succinct. Recall that for all  $t \leq T_{\text{stab}}$ ,  $\|L_t\|_{\text{op}} \leq \bar{\gamma}_\infty < 1$ , and that  $\|L_\infty\|_{\text{op}} \leq \gamma_\infty < 1$ . We proceed by a telescoping argument:

$$L_{i \rightarrow t} - L_\infty^{t-i} = L_t (L_{i \rightarrow t-1} - L_\infty^{t-i-1}) + L_\infty^{t-i-1} (L_t - L_\infty),$$

and so

$$\|L_{i \rightarrow t} - L_\infty^{t-i}\|_{\text{op}} = \bar{\gamma}_\infty \|L_{i \rightarrow t-1} - L_\infty^{t-i-1}\|_{\text{op}} + \gamma_\infty^{t-i-1} \|L_t - L_\infty\|_{\text{op}}.$$

Proceedings backwards in the same fashion, we have

$$\begin{aligned}
 \|A_{\text{cl},i \rightarrow t} - A_{\text{cl},\infty}^{t-i}\|_{\text{op}} &\leq \kappa_\infty \bar{\gamma}_\infty^{t-i-1} \sum_{j=i+1}^t \|L_j - L_\infty\|_{\text{op}} \\
 &\leq \kappa_\infty^2 \bar{\gamma}_\infty^{t-i-1} \sum_{j=i+1}^t \|A_{\text{cl},j} - A_{\text{cl},\infty}\|_{\text{op}} \\
 &\leq \kappa_\infty^2 \Psi_\star \bar{\gamma}_\infty^{t-i-1} \sum_{j=i+1}^t \|K_j - K_\infty\|_{\text{op}}.
 \end{aligned}$$

Using [Lemma D.11](#), we are guaranteed that by setting

$$\Delta = C \cdot \beta_\star \Psi_\star^2 \Gamma_\star \log(\kappa_\infty^2 \Psi_\star \Gamma_\star (1 - \gamma_\infty)^{-1} \cdot \alpha T^3) \geq \Delta_{\text{stab}},$$

where  $C$  is a sufficiently large constant, we have

$$\|K_t - K_\infty\|_{\text{op}} \leq \frac{1}{\kappa_\infty^2 \Psi_\star \cdot (\alpha T^3)} \quad \forall t \leq T - \Delta,$$

and in particular,

$$\|A_{\text{cl},i \rightarrow t} - A_{\text{cl},\infty}^{t-i}\|_{\text{op}} \leq \frac{1}{\alpha T^2}.$$

□

**Proof of [Lemma D.9](#).** We first handle the policy  $\pi$ . Observe the state at each step is given by

$$x_{t+1}^\pi(\mathbf{w}) = \sum_{i=1}^t (A - BK_\infty)^{t-i} w_i - \sum_{i=1}^t (A - BK_\infty)^{t-i} Bq_i(\mathbf{w}).$$

Hence, using [Lemma D.2](#), we have

$$\begin{aligned}
 \|x_{t+1}^\pi(\mathbf{w}_t)\| &\leq \kappa_\infty \Psi_\star \sum_{i=1}^t \gamma_\infty^{t-i} (1 + \max_{i \leq t} \|q_i(\mathbf{w})\|) \leq \kappa_\infty \Psi_\star (1 - \gamma_\infty)^{-1} (1 + \max_{i \leq t} \|q_i(\mathbf{w})\|) \\
 &\leq 2\kappa_\infty \Psi_\star (1 - \gamma_\infty)^{-1} D_q.
 \end{aligned}$$

We can now bound the control as

$$\begin{aligned}
 \|u_{t+1}^\pi(\mathbf{w})\| &\leq \|K_\infty x_{t+1}^\pi(\mathbf{w})\| + \|q_{t+1}(\mathbf{w})\| \\
 &\leq 2\kappa_\infty \beta_\star \Psi_\star^3 \Gamma_\star (1 - \gamma_\infty)^{-1} D_q + D_q \\
 &\leq 3\kappa_\infty \beta_\star \Psi_\star^3 \Gamma_\star (1 - \gamma_\infty)^{-1} D_q.
 \end{aligned}$$

where the second inequality uses [\(41\)](#) along with the previous bound on  $x_t^\pi$ .

We now handle the policy  $\hat{\pi}$ . Recall that the state reached after playing any controller of the form  $\hat{\pi}_t(x, \mathbf{w}) = -K_t x - q_t(\mathbf{w})$  for every step is given by

$$x_{t+1}^{\hat{\pi}}(\mathbf{w}) = \sum_{i=1}^t A_{\text{cl},i \rightarrow t} w_i - \sum_{i=1}^t A_{\text{cl},i \rightarrow t} Bq_i(\mathbf{w}),$$

and so

$$\|x_{t+1}^{\hat{\pi}}(\mathbf{w})\| \leq (1 + \Psi_\star \max_{1 \leq i \leq t} \|q_i(\mathbf{w})\|) \cdot \sum_{i=1}^t \|A_{\text{cl},i \rightarrow t}\|_{\text{op}}.$$

By [Lemma D.3](#), we have

$$\begin{aligned}
 \sum_{i=1}^t \|A_{\text{cl},i \rightarrow t}\|_{\text{op}} &\leq \sum_{i=1}^t \kappa_\infty^2 \beta_\star^{1/2} \Gamma_\star^{1/2} \bar{\gamma}_\infty^{T_{\text{stab}} \wedge t - T_{\text{stab}} \wedge (i+1)} \\
 &\leq C \cdot \kappa_\infty^2 \beta_\star^{1/2} \Gamma_\star^{1/2} (\Delta_{\text{stab}} + (1 - \bar{\gamma}_\infty)^{-1}),
 \end{aligned}$$

where  $C$  is a universal constant. Recalling the value for  $\Delta_{\text{stab}}$ , this gives

$$\sum_{i=1}^t \|A_{\text{cl}, i \rightarrow t}\|_{\text{op}} \leq \tilde{\mathcal{O}}(\kappa_{\infty}^2 \beta_{\star}^{3/2} \Psi_{\star}^2 \Gamma_{\star}^{3/2} (1 - \gamma_{\infty})^{-1}).$$

Hence, we can bound the state norm as

$$\|x_{t+1}^{\hat{\pi}}(\mathbf{w})\| \leq \tilde{\mathcal{O}}\left(\kappa_{\infty}^2 \beta_{\star}^{3/2} \Psi_{\star}^3 \Gamma_{\star}^{3/2} (1 - \gamma_{\infty})^{-1} \cdot D_q\right).$$

Finally, we bound the control norm as

$$\|u_{t+1}^{\hat{\pi}}(\mathbf{w})\| \leq \|K_t\|_{\text{op}} \|x_{t+1}^{\hat{\pi}}(\mathbf{w})\| + \|q_{t+1}(\mathbf{w})\|.$$

We use that  $P_t \leq P_{\infty}$  for all  $t$  to bound

$$\|K_t\|_{\text{op}} \leq \beta_{\star} \Psi_{\star}^2 \Gamma_{\star},$$

which gives

$$\|u_{t+1}^{\hat{\pi}}(\mathbf{w})\| \leq \tilde{\mathcal{O}}\left(\kappa_{\infty}^2 \beta_{\star}^{5/2} \Psi_{\star}^5 \Gamma_{\star}^{5/2} (1 - \gamma_{\infty})^{-1} \cdot D_q\right).$$

□

## D.2. Performance difference lemma

Below we state a variant of the performance difference lemma for an abstract MDP setting that generalizes the LQR setting studied in this paper. The setting as follows:

Begin at state  $x_1 \in \mathcal{X}$ . Then, for  $t = 1, \dots, T$ :

- Agent selects control  $u_t \in \mathcal{U}$ .
- Agent observes  $w_t \in \mathcal{W}$  and experiences instantaneous loss  $\ell(x_t, u_t, w_t)$ .
- State evolves as  $x_{t+1} \stackrel{\text{i.i.d.}}{\sim} p(x_t, u_t, w_t)$ , where  $p(x, u, w) \in \Delta(\mathcal{X})$ .

We define the expected loss of a policy  $\pi_t(x; \mathbf{w})$  in this setting as

$$J_T(\pi; \mathbf{w}) = \mathbb{E}_{\pi, \mathbf{w}} \left[ \sum_{t=1}^T \ell(x_t, u_t, w_t) \right], \quad (34)$$

where  $\mathbb{E}_{\pi, \mathbf{w}}$  denotes expectation with respect to the system dynamics with  $\mathbf{w}$  fixed. For each policy  $\pi$ , we define the action-value function for  $\pi$  as follows:

$$\widehat{\mathbf{Q}}_{t;\tau}^{\pi}(x, u; \mathbf{w}_{\tau}) = \mathbb{E}_{\pi, \mathbf{w}_{\tau}} \left[ \sum_{s=t}^{\tau} \ell(x_s, u_s, w_s) \mid x_t = x, u_t = u \right]. \quad (35)$$

The performance difference lemma can now be stated as follows.

**Lemma D.12** (Performance difference lemma). Let  $\hat{\pi}$  and  $\pi$  be any pair of policies of the form  $\pi_t(x; \mathbf{w})$  (i.e., Markovian, but with potentially arbitrary dependence on the sequence  $\mathbf{w}$ ). Then it holds that

$$J_T(\hat{\pi}; \mathbf{w}) - J_T(\pi; \mathbf{w}) = \mathbb{E}_{\pi, \mathbf{w}} \left[ \sum_{t=1}^T \widehat{\mathbf{Q}}_t^{\hat{\pi}}(x_t, \hat{\pi}(x_t; \mathbf{w}); \mathbf{w}) - \widehat{\mathbf{Q}}_t^{\pi}(x_t, \pi(x_t; \mathbf{w}); \mathbf{w}) \right] \quad (36)$$

$$= \mathbb{E}_{\hat{\pi}, \mathbf{w}} \left[ \sum_{t=1}^T \widehat{\mathbf{Q}}_t^{\hat{\pi}}(x_t, \hat{\pi}(x_t; \mathbf{w}); \mathbf{w}) - \widehat{\mathbf{Q}}_t^{\pi}(x_t, \pi(x_t; \mathbf{w}); \mathbf{w}) \right]. \quad (37)$$

**Proof of Lemma D.12.** Let  $t$  be fixed. Observe that for any  $x$ , we have

$$\begin{aligned} & \widehat{\mathbf{Q}}_t^\pi(x, \pi_t(x; \mathbf{w}); \mathbf{w}) \\ &= \ell(x, \pi_t(x; \mathbf{w}), w_t) + \mathbb{E}\left[\widehat{\mathbf{Q}}_{t+1}^\pi(x_{t+1}, \pi_{t+1}(x_{t+1}; \mathbf{w}); \mathbf{w}) \mid x_t = x, u_t = \pi_t(x; \mathbf{w}), \mathbf{w}\right]. \end{aligned}$$

We can alternatively write

$$\begin{aligned} \ell(x, \pi_t(x; \mathbf{w}), w_t) &= \widehat{\mathbf{Q}}_t^{\widehat{\pi}}(x, \pi_t(x; \mathbf{w}); \mathbf{w}) \\ &\quad - \mathbb{E}\left[\widehat{\mathbf{Q}}_{t+1}^{\widehat{\pi}}(x_{t+1}, \widehat{\pi}_{t+1}(x_{t+1}; \mathbf{w}); \mathbf{w}) \mid x_t = x, u_t = \pi_t(x; \mathbf{w}), \mathbf{w}\right]. \end{aligned}$$

Combining these identities, we have

$$\begin{aligned} & \widehat{\mathbf{Q}}_t^{\widehat{\pi}}(x, \widehat{\pi}_t(x; \mathbf{w}); \mathbf{w}) - \widehat{\mathbf{Q}}_t^\pi(x, \pi_t(x; \mathbf{w}); \mathbf{w}) \\ &= \widehat{\mathbf{Q}}_t^{\widehat{\pi}}(x, \widehat{\pi}_t(x; \mathbf{w}); \mathbf{w}) - \widehat{\mathbf{Q}}_t^{\widehat{\pi}}(x, \pi_t(x; \mathbf{w}); \mathbf{w}) \\ &\quad + \mathbb{E}\left[\widehat{\mathbf{Q}}_{t+1}^{\widehat{\pi}}(x_{t+1}, \widehat{\pi}_{t+1}(x_{t+1}; \mathbf{w}); \mathbf{w}) - \widehat{\mathbf{Q}}_{t+1}^\pi(x_{t+1}, \pi_{t+1}(x_{t+1}; \mathbf{w}); \mathbf{w}) \mid x_t = x, u_t = \pi_t(x; \mathbf{w}), \mathbf{w}\right]. \end{aligned} \tag{38}$$

To prove the result, we simply observe that

$$J_T(\widehat{\pi}; \mathbf{w}) - J_T(\pi; \mathbf{w}) = \widehat{\mathbf{Q}}_1^{\widehat{\pi}}(x_1, \widehat{\pi}(x; \mathbf{w}); \mathbf{w}) - \widehat{\mathbf{Q}}_1^\pi(x_1, \pi(x; \mathbf{w}); \mathbf{w}).$$

The equality (36) now follows by applying the identity (38) to the right-hand side above recursively. To prove (37) we use the same argument, except that we replace the one-step identity (38) with

$$\begin{aligned} & \widehat{\mathbf{Q}}_t^{\widehat{\pi}}(x, \widehat{\pi}_t(x; \mathbf{w}); \mathbf{w}) - \widehat{\mathbf{Q}}_t^\pi(x, \pi_t(x; \mathbf{w}); \mathbf{w}) \\ &= \widehat{\mathbf{Q}}_t^{\widehat{\pi}}(x, \widehat{\pi}_t(x; \mathbf{w}); \mathbf{w}) - \widehat{\mathbf{Q}}_t^{\widehat{\pi}}(x, \pi_t(x; \mathbf{w}); \mathbf{w}) \\ &\quad + \mathbb{E}\left[\widehat{\mathbf{Q}}_{t+1}^{\widehat{\pi}}(x_{t+1}, \widehat{\pi}_{t+1}(x_{t+1}; \mathbf{w}); \mathbf{w}) - \widehat{\mathbf{Q}}_{t+1}^\pi(x_{t+1}, \pi_{t+1}(x_{t+1}; \mathbf{w}); \mathbf{w}) \mid x_t = x, u_t = \widehat{\pi}_t(x; \mathbf{w}), \mathbf{w}\right]. \end{aligned}$$

□

## E. Proofs from Section 2

### E.1. Proof of Theorem 1

**Theorem 1.** *For an appropriate choice of parameters, Riccatitron ensures*

$$\mathcal{K}_0\text{-Reg}_T \leq \mathcal{O}_*(d_{\mathbf{x}}d_{\mathbf{u}} \log^3 T),$$

where  $\mathcal{O}_*$  suppresses polynomial dependence on system parameters. Suppressing only logarithmic dependence on system parameters, the regret is at most

$$\widetilde{\mathcal{O}}(d_{\mathbf{x}}d_{\mathbf{u}} \log^3 T \cdot \beta_*^{11} \Psi_*^{19} \Gamma_*^{11} \kappa_0^8 (1 - \gamma_0)^{-4}).$$

**Proof of Theorem 1.** Throughout the proof, we let  $\widehat{\pi}$  denote the policy of Riccatitron, which takes the form  $\widehat{\pi}_t(x, \mathbf{w}_{t-1}) = -K_\infty x - q^{M_t}(\mathbf{w}_{t-1})$ , where  $M_t = M_t(\mathbf{w}_{t-1})$  is selected as in Algorithm 1. The proof is split into multiple subsections.

#### E.1.1. REDUCTION TO ONLINE PREDICTION

As a first step, we appeal to Lemma 2.1 which, by choosing  $\mathcal{M}_0 = \mathcal{M}(m, R_*, \gamma_0)$  for  $m = (1 - \gamma_0)^{-1} \log((1 - \gamma_0)^{-1} T)$ , ensures that

$$J_T(\widehat{\pi}; \mathbf{w}) - \inf_{K \in \mathcal{K}_0} J_T(\pi^K; \mathbf{w}) \leq J_T(\widehat{\pi}; \mathbf{w}) - \inf_{M \in \mathcal{M}_0} J_T(\pi^{(M)}; \mathbf{w}) + C_{\text{app}}.$$

Next, we recall that by the performance difference lemma (2), we have that for any  $M \in \mathcal{M}_0$ ,

$$J_T(\widehat{\pi}; \mathbf{w}) - J_T(\pi^{(M)}; \mathbf{w}) = \sum_{t=1}^T \mathbf{A}_t^*(u_t^{\widehat{\pi}}; x_t^{\widehat{\pi}}, \mathbf{w}) - \mathbf{A}_t^*(u_t^{\pi^{(M)}}; x_t^{\pi^{(M)}}, \mathbf{w}).$$

We apply [Theorem 2](#) to both terms in this summation individually. In particular, by choosing

$$h = 2(1 - \gamma_\infty)^{-1} \log(\kappa_\infty^2 \beta_\star^2 \Psi_\star \Gamma_\star^2 T^2),$$

we are guaranteed that

$$\begin{aligned} & J_T(\widehat{\pi}; \mathbf{w}) - \inf_{M \in \mathcal{M}_0} J_T(\pi^{(M)}; \mathbf{w}) \\ & \leq \sum_{t=1}^T \widehat{\mathbf{A}}_{t,h}(M_t; \mathbf{w}_{t+h}) - \inf_{M \in \mathcal{M}_0} \sum_{t=1}^T \widehat{\mathbf{A}}_{t,h}(M; \mathbf{w}_{t+h}) + C_{\text{adv}}. \end{aligned}$$

Defining a ‘‘loss function’’  $f_t(M) = \widehat{\mathbf{A}}_{t,h}(M_t; \mathbf{w}_{t+h}) = \|q_t^M(\mathbf{w}_{t-1}) - q_{\infty,h}^*(\mathbf{w}_{t+h})\|_{\Sigma_\infty}^2$ , the regret-like quantity above is equivalent to

$$\sum_{t=1}^T f_t(M_t) - \inf_{M \in \mathcal{M}_0} \sum_{t=1}^T f_t(M), \quad (39)$$

where  $\{M_t\}$  are the disturbance-action matrices selected by Riccatitron.

### E.1.2. APPLYING THE ONLINE NEWTON STEP ALGORITHM

As described in the main body, Riccatitron is simply an instance of the generic reduction from online convex optimization with delays to vanilla online convex optimization, with either online Newton step or Vovk-Azoury-Warmuth as the base algorithm in the reduction. For online Newton step, since we have delay  $h$ , [Lemma 2.2](#) ensures that we have

$$\sum_{t=1}^T f_t(M_t) - \inf_{M \in \mathcal{M}_0} \sum_{t=1}^T f_t(M) \leq (h+1)R_{\text{ONS}}(T/(h+1)),$$

where  $R_{\text{ONS}}(T/(h+1))$  is an upper bound on the regret of each ONS instance applied to its respective subsequence of losses. Moreover, by [Lemma 2.4](#) we are guaranteed that if we choose  $\eta_{\text{ons}} = 2 \max\{4G_{\text{oco}}D_{\text{oco}}, \alpha_{\text{oco}}^{-1}\}$  and  $\varepsilon_{\text{ons}} = \eta_{\text{ons}}^2/D_{\text{oco}}$ , then

$$R_{\text{ONS}}(T) \leq 5(\alpha_{\text{oco}}^{-1} + G_{\text{oco}}D_{\text{oco}}) \dim(\mathcal{M}_0) \log T,$$

where  $\alpha_{\text{oco}}$ ,  $G_{\text{oco}}$ , and  $D_{\text{oco}}$  are regularity parameters for the losses  $f_t$  which are specified by the following lemma.

**Lemma E.1.** The weight set  $\mathcal{M}_0$  and loss functions  $f_t(M)$  in (39) satisfy the following properties:

- $\sup_{M, M' \in \mathcal{M}_0} \|M - M'\|_F \leq 4\beta_\star \Psi_\star^2 \Gamma_\star \kappa_0^2 (1 - \gamma_0)^{-1} \cdot \sqrt{d_{\mathbf{x}} \wedge d_{\mathbf{u}}} =: D_{\text{oco}}$ .
- $\sup_{M \in \mathcal{M}_0} \|\nabla f_t(M)\|_F \leq \widetilde{\mathcal{O}}(D_q \Psi_\star^2 \Gamma_\star (1 - \gamma_0)^{-1/2}) =: G_{\text{oco}}$ .
- $f_t$  is  $\alpha_{\text{oco}}$ -exp-concave over  $\mathcal{M}_0$ , where  $\alpha_{\text{oco}} := (4D_q^2 \Psi_\star^2 \Gamma_\star)^{-1}$ .

With this lemma, we can crudely bound the regret of ONS as

$$\begin{aligned} R_{\text{ONS}}(T) &= \widetilde{\mathcal{O}}((G_{\text{oco}}D_{\text{oco}} + \alpha_{\text{oco}}^{-1}) \dim(\mathcal{M}_0) \log T) \\ &= \widetilde{\mathcal{O}}(md_{\mathbf{x}}d_{\mathbf{u}}(G_{\text{oco}}D_{\text{oco}} + \alpha_{\text{oco}}^{-1}) \log T) \\ &= \widetilde{\mathcal{O}}((1 - \gamma_0)^{-1} d_{\mathbf{x}}d_{\mathbf{u}}(G_{\text{oco}}D_{\text{oco}} + \alpha_{\text{oco}}^{-1}) \log^2 T) \\ &= \widetilde{\mathcal{O}}(d_{\mathbf{x}}d_{\mathbf{u}}\sqrt{d_{\mathbf{x}} \wedge d_{\mathbf{u}}} \cdot D_q^2 \kappa_0^2 \beta_\star \Psi_\star^2 \Gamma_\star^2 (1 - \gamma_0)^{-5/2} \log^2 T) \\ &\leq \widetilde{\mathcal{O}}(d_{\mathbf{x}}d_{\mathbf{u}}\sqrt{d_{\mathbf{x}} \wedge d_{\mathbf{u}}} \cdot \kappa_0^6 \beta_\star^6 \Psi_\star^8 \Gamma_\star^7 (1 - \gamma_0)^{-9/2} \log^2 T). \end{aligned}$$

### E.1.3. APPLYING THE VOVK-AZOURY-WARMUTH ALGORITHM

If we use VAW as the base algorithm instead of ONS, then [Lemma 2.2](#) implies that

$$\sum_{t=1}^T f_t(M_t) - \inf_{M \in \mathcal{M}_0} \sum_{t=1}^T f_t(M) \leq (h+1)R_{\text{VAW}}(T/(h+1)),$$



where  $R_{\text{VAW}}(T/(h+1))$  is an upper bound on the regret of each VAW instance. [Theorem 5](#) (detailed in [Appendix E.3](#)) ensures that by setting  $\varepsilon_{\text{vaw}} = \|\Sigma_\infty\|_{\text{op}} D_q^2 D_{\text{oco}}^{-2}$ , we have

$$R_{\text{VAW}}(T) \leq 5\|\Sigma_\infty\|_{\text{op}} D_q \dim(\mathcal{M}_0) \log\left(1 + D_q^{-2} D_{\text{oco}}^2 Q_{\text{oco}} T / \dim(\mathcal{M}_0)\right),$$

where  $D_{\text{oco}}$  is as in [Lemma E.1](#) and

$$Q_{\text{oco}} := \sup_{M \neq 0} \frac{\|q^M(\mathbf{w})\|}{\|M\|_F} \leq \sup_{M \neq 0} \frac{\sum_{i=1}^m \|M^{[i]}\|_{\text{op}}}{\|M\|_F} \leq \sqrt{m}.$$

Recalling that  $\|\Sigma_\infty\|_{\text{op}} \leq 2\Psi_\star^2 \Gamma_\star$ ,  $D_q \leq \tilde{\mathcal{O}}\left(\beta_\star^{5/2} \Psi_\star^3 \Gamma_\star^{5/2} \kappa_0^2 (1-\gamma_0)^{-1}\right)$  and (using the choice of  $m$  from [Lemma 2.1](#))

$$\dim(\mathcal{M}_0) = d_{\mathbf{x}} d_{\mathbf{u}} m = \tilde{\mathcal{O}}(d_{\mathbf{x}} d_{\mathbf{u}} (1-\gamma_0)^{-1} \log T),$$

we can simplify to

$$\begin{aligned} R_{\text{VAW}}(T) &\leq \tilde{\mathcal{O}}\left(\|\Sigma_\infty\|_{\text{op}} D_q d_{\mathbf{x}} d_{\mathbf{u}} m \log T\right) \\ &\leq \tilde{\mathcal{O}}\left(d_{\mathbf{x}} d_{\mathbf{u}} \log^2 T \cdot \beta_\star^{5/2} \Psi_\star^5 \Gamma_\star^{7/2} \kappa_0^2 (1-\gamma_0)^{-2}\right). \end{aligned}$$

#### E.1.4. PUTTING EVERYTHING TOGETHER

We now summarize the development so far. Suppose we choose  $\mathcal{M}_0$  as in [Lemma 2.1](#), using  $m = (1-\gamma_0)^{-1} \log((1-\gamma_0)^{-1} T)$ . [Lemma 2.4](#) implies that if we run VAW as the base algorithm in the reduction using  $\varepsilon_{\text{vaw}} = \|\Sigma_\infty\|_{\text{op}} D_q^2 D_{\text{oco}}^{-2}$  and delay parameter  $h = 2(1-\gamma_\infty)^{-1} \log(\kappa_\infty^2 \beta_\star^2 \Psi_\star \Gamma_\star^2 T^2)$ , we have

$$\begin{aligned} \sum_{t=1}^T f_t(M_t) - \inf_{M \in \mathcal{M}_0} \sum_{t=1}^T f_t(M) &\leq (h+1) R_{\text{VAW}}(T/(h+1)) \\ &\leq \tilde{\mathcal{O}}\left(h \cdot d_{\mathbf{x}} d_{\mathbf{u}} \log^2 T \cdot \beta_\star^{5/2} \Psi_\star^5 \Gamma_\star^{7/2} \kappa_0^2 (1-\gamma_0)^{-2}\right) \\ &\leq \underbrace{\tilde{\mathcal{O}}\left(d_{\mathbf{x}} d_{\mathbf{u}} \log^3 T \cdot \beta_\star^{5/2} \Psi_\star^5 \Gamma_\star^{7/2} \kappa_0^2 (1-\gamma_0)^{-3}\right)}_{=: C_{\text{reg}}}. \end{aligned}$$

In total, we have

$$\mathcal{K}_0\text{-Reg}_T \leq C_{\text{reg}} + C_{\text{apx}} + C_{\text{adv}} \leq \tilde{\mathcal{O}}\left(d_{\mathbf{x}} d_{\mathbf{u}} \log^3 T \cdot \beta_\star^{11} \Psi_\star^{19} \Gamma_\star^{11} \kappa_0^8 (1-\gamma_0)^{-4}\right).$$

□

## E.2. Supporting lemmas

**Lemma 2.1** (Expressivity of DAP). Suppose we choose our set of disturbance-action matrices as  $\mathcal{M}_0 := \mathcal{M}(m, R_\star, \gamma_0)$ , where  $m = (1-\gamma_0)^{-1} \log((1-\gamma_0)^{-1} T)$  and  $R_\star = 2\beta_\star \Psi_\star^2 \Gamma_\star \kappa_0^2$ . Then for all  $\mathbf{w}$ , we have

$$\inf_{M \in \mathcal{M}_0} J_T(\pi^{(M)}; \mathbf{w}) \leq \inf_{K \in \mathcal{K}_0} J_T(\pi^K; \mathbf{w}) + C_{\text{apx}},$$

where  $C_{\text{apx}} \leq \mathcal{O}(\beta_\star^2 \Psi_\star^8 \Gamma_\star^2 \kappa_0^7 (1-\gamma_0)^{-2})$ .

**Proof of Lemma 2.1.** Let  $K \in \mathcal{K}_0$  be fixed. Consider a policy

$$\pi_t^{(M)}(x; \mathbf{w}_{t-1}) = -K_\infty x - q^M(\mathbf{w}_{t-1}),$$

Following ([Agarwal et al., 2019a](#)), we set

$$M^{[i]} = (K - K_\infty)(A - BK)^{i-1}.$$

Suppose for now that  $\pi^{(M)}$  and  $\pi^K$  have  $\|x_t\| \vee \|u_t\| \leq \tilde{D}$  for all  $t$ . Then Lemma 5.2 of (Agarwal et al., 2019a) implies that

$$J_T(\tilde{\pi}^M; \mathbf{w}) \leq J_T(\pi^K; \mathbf{w}) + \mathcal{O}(\tilde{D}\Psi_\star^3\kappa_0^5 \cdot m\gamma_0^{m+1}T). \quad (40)$$

Let us bound the norms for the matrices  $M^{[i]}$  that achieve this bound. First observe that

$$\|K_\infty\|_{\text{op}} \leq \|\Sigma_\infty\|_{\text{op}}^{-1}\|A\|_{\text{op}}\|B\|_{\text{op}}\|P_\infty\|_{\text{op}} \leq \beta_\star\Psi_\star^2\Gamma_\star, \quad \text{and} \quad \|K\|_{\text{op}} \leq \kappa_0. \quad (41)$$

Consequently, Lemma D.2 implies that

$$\|M^{[i]}\|_{\text{op}} \leq (\|K\|_{\text{op}} + \|K_\infty\|_{\text{op}})\kappa_0\gamma_0^{i-1} \leq 2\kappa_0^2\beta_\star\Psi_\star^2\Gamma_\star\gamma_0^{i-1}.$$

Hence, if the use controller  $\pi^{(M)}$ , it would suffice to take

$$\mathcal{M}_0 = \left\{ M = \{M^{[i]}\}_{i \in [m]} \mid \|M^{[i]}\|_{\text{op}} \leq 2\beta_\star\Psi_\star^2\Gamma_\star\kappa_0^2\gamma_0^{i-1} \right\}.$$

To conclude the proof, we provide a bound on  $\tilde{D}$ . To begin, note that each  $M \in \mathcal{M}_0$  has

$$\|q_i^M(\mathbf{w}_{i-1})\| \leq \sum_{i=1}^m \|M^{[i]}\|_{\text{op}} \leq 2\beta_\star\Psi_\star^2\Gamma_\star\kappa_0^2(1-\gamma_0)^{-1} =: D_{\mathcal{M}}. \quad (42)$$

We now provide a bound on  $\tilde{D}$ . First, observe that when  $\pi$  is the static linear controller  $\pi^K$ , we have

$$x_{t+1}(\mathbf{w}_t) = \sum_{i=1}^t (A - BK)^{t-i} w_i,$$

and so, use Lemma D.2, we have

$$\|x_t(\mathbf{w}_{t-1})\|_{\text{op}} \leq \kappa_0 \sum_{i=1}^t \gamma_0^{t-i} \leq \kappa_0(1-\gamma_0)^{-1},$$

and  $\|u_t(\mathbf{w}_{t-1})\| = \|Kx_t(\mathbf{w}_{t-1})\|_{\text{op}} \leq \kappa_0^2(1-\gamma_0)^{-1}$ . To bound the radius for the policies  $\pi^{(M)}$ , we use Lemma D.9, along with the bound (42) to get the following result.

**Corollary 1.** For any  $M \in \mathcal{M}_0$ , the controller  $\pi^{(M)}$  has

$$\|x_{t+1}^{\pi^{(M)}}(\mathbf{w}_t)\| \leq 2\beta_\star\Psi_\star^3\Gamma_\star\kappa_0^3(1-\gamma_0)^{-2}, \quad \text{and} \quad \|u_{t+1}^{\pi^{(M)}}(\mathbf{w}_t)\| \leq 3\beta_\star^2\Psi_\star^5\Gamma_\star^2\kappa_0^3(1-\gamma_0)^{-2}.$$

Hence, we may take

$$\tilde{D} = 2\beta_\star^2\Psi_\star^5\Gamma_\star^2\kappa_0^3(1-\gamma_0)^{-2},$$

and so (40) yields

$$J_T(\pi^{(M)}; \mathbf{w}) \leq J_T(\pi^K; \mathbf{w}) + \mathcal{O}(\beta_\star^2\Psi_\star^8\Gamma_\star^2(1-\gamma_0)^{-2}\kappa_0^7 \cdot m\gamma_0^{m+1}T).$$

By choosing  $m = (1-\gamma_0)^{-1} \log((1-\gamma_0)^{-1}T)$ , we are guaranteed that

$$J_T(\pi^{(M)}; \mathbf{w}) \leq J_T(\pi^K; \mathbf{w}) \leq C_{\text{apx}}.$$

As a closing remark, we observe that (42) implies that we may take  $D_q = \max\{2\kappa_0^2\beta_\star\Psi_\star^2\Gamma_\star(1-\gamma_0)^{-1}, D_{q^\star}\}$ , as the radius for the predictions  $q_t^M$  by the learner, benchmark class, and optimal policy. Hence, recalling the value for  $D_{q^\star}$  from Lemma D.6, we may take

$$D_q \leq \tilde{\mathcal{O}}\left(\beta_\star^{5/2}\Psi_\star^3\Gamma_\star^{5/2}\kappa_0^2(1-\gamma_0)^{-1}\right).$$

□

**Lemma E.1.** The weight set  $\mathcal{M}_0$  and loss functions  $f_t(M)$  in (39) satisfy the following properties:

- $\sup_{M, M' \in \mathcal{M}_0} \|M - M'\|_F \leq 4\beta_\star\Psi_\star^2\Gamma_\star\kappa_0^2(1-\gamma_0)^{-1} \cdot \sqrt{d_{\mathbf{x}} \wedge d_{\mathbf{u}}} =: D_{\text{oco}}.$

- $\sup_{M \in \mathcal{M}_0} \|\nabla f_t(M)\|_F \leq \tilde{\mathcal{O}}(D_q \Psi_*^2 \Gamma_* (1 - \gamma_0)^{-1/2}) =: G_{\text{oco}}$ .
- $f_t$  is  $\alpha_{\text{oco}}$ -exp-concave over  $\mathcal{M}_0$ , where  $\alpha_{\text{oco}} := (4D_q^2 \Psi_*^2 \Gamma_*)^{-1}$ .

**Proof of Lemma E.1.** For the first property, observe that for each  $M \in \mathcal{M}_0$ , we have

$$\begin{aligned} \|M\|_F &= \sqrt{\sum_{i=1}^m \|M^{[i]}\|_F^2} \leq \sqrt{d_x \wedge d_u} \sqrt{\sum_{i=1}^m \|M^{[i]}\|_{\text{op}}^2} \\ &\leq \sqrt{d_x \wedge d_u} \cdot 2\kappa^2 \beta_* \Psi_*^2 \Gamma_* \sqrt{\sum_{i=1}^m \gamma_0^{2(i-1)}} \\ &\leq \sqrt{d_x \wedge d_u} \cdot 2\kappa^2 \beta_* \Psi_*^2 \Gamma_* (1 - \gamma_0)^{-1}. \end{aligned}$$

The bound for  $D_{\text{oco}}$  now follows by triangle inequality.

For the second property, we directly prove that  $f_t$  is Lipschitz as follows: For any  $M, M' \in \mathcal{M}_0$ ,

$$\begin{aligned} &\|q_{t:t+m}^*(w_{t:t+m}) - q^M(\mathbf{w}_{t-1})\|_{\Sigma_t}^2 - \|q_{t:t+m}^*(w_{t:t+m}) - q^{M'}(\mathbf{w}_{t-1})\|_{\Sigma_t}^2 \\ &\leq 2\|\Sigma_t\|_{\text{op}} D_q \|q^M(\mathbf{w}_{t-1}) - q^{M'}(\mathbf{w}_{t-1})\| \\ &= 2\|\Sigma_t\|_{\text{op}} D_q \left\| \sum_{i=1}^m (M^{[i]} - M'^{[i]}) \mathbf{w}_{t-i} \right\|. \end{aligned}$$

We finish the bound as follows:

$$\sum_{i=1}^m \|M^{[i]} - M'^{[i]}\|_{\text{op}} \leq \sum_{i=1}^m \|M^{[i]} - M'^{[i]}\|_F \leq \sqrt{m} \|M - M'\|_F.$$

To simplify the bound, we use that  $\|\Sigma_t\|_{\text{op}} \leq \Psi_*^2 \Gamma_*$  and that  $\sqrt{m} = \tilde{\mathcal{O}}((1 - \gamma_0)^{-1/2})$ .

For the third property, we observe that  $f_t(M)$  obeys the structure in Lemma 2.3, since  $q^M(\mathbf{w}_{t-1})$  is a linear mapping from  $\prod_{i=1}^m \mathbb{R}^{d_u \times d_x}$  to  $\mathbb{R}^{d_u}$ , and since  $\Sigma_t > 0$ . Thus, to prove the exp-concave property, we simply bound the range of the loss as

$$\|q_{t:t+m}^*(w_{t:t+m}) - q^M(\mathbf{w}_{t-1})\|_{\Sigma_t}^2 \leq 2D_q^2 \|\Sigma_t\|_{\text{op}} \leq 2D_q^2 \Psi_*^2 \Gamma_*.$$

□

### E.3. Vector-valued Vovk-Azoury-Warmuth algorithm

In this section we develop a variant of the Vovk-Azoury-Warmuth algorithm (Vovk, 1998; Azoury & Warmuth, 2001) for a vector-valued online regression setting. At each timestep  $t = 1, \dots, T$ , the learner receives a matrix  $A_t \in \mathbb{R}^{d_1 \times d_2}$ , predicts  $z_t \in \mathbb{R}^{d_2}$ , then receives  $b_t \in \mathbb{R}^{d_1}$  and experiences loss  $f_t(z_t)$ , where  $f_t(z) = \|A_t z - b_t\|_{\Sigma}^2$  and  $\Sigma > 0$  is a known matrix. The goal of the learner is to attain low regret

$$\sum_{t=1}^T f_t(z_t) - \inf_{z \in \mathcal{C}} \sum_{t=1}^T f_t(z),$$

where  $\mathcal{C}$  is a convex constraint set. Recall from Algorithm 3 that VAW is the algorithm which, at time  $t$ , predicts with

$$z_t = \arg \min_{z \in \mathcal{C}} \left\{ \langle z, -2 \sum_{i=1}^{t-1} A_i^\top \Sigma b_i \rangle + \|z\|_{E_t}^2 \right\}, \quad (43)$$

where  $E_t = \varepsilon I + \sum_{i=1}^t A_i^\top \Sigma A_i$ .

**Theorem 5.** Let  $\|\Sigma\|_{\text{op}} \leq S$ . Suppose that we run the VAW strategy (Algorithm 3) with parameter  $\varepsilon$ , and that for all  $t$  we have  $\|b_t\| \leq Y$  and  $\|A_t\|_{\text{op}} \leq R$ . Then we are guaranteed that for all  $z \in \mathcal{C}$ ,

$$\sum_{t=1}^T f_t(z_t) - \sum_{t=1}^T f_t(z) \leq \varepsilon \|z\|^2 + 4SY^2 \cdot d_2 \log(1 + SR^2 T / (d_2 \varepsilon)). \quad (44)$$

**Algorithm 3** Vector-valued Vovk-Azoury-Warmuth (VAW( $\varepsilon, \mathcal{C}, \Sigma$ ))

- 1: **parameters:** Regularization parameter  $\varepsilon > 0$ , convex constraint set  $\mathcal{C}$ , cost matrix  $\Sigma > 0$ .  
*// OCO with costs  $f_k(z) := \|A_k z - b_k\|_{\Sigma}^2$ , where  $A_k \in \mathbb{R}^{d_1 \times d_2}$ ,  $b_k \in \mathbb{R}^{d_1}$  and  $z \in \mathcal{C} \subset \mathbb{R}^{d_2}$ .*
- 2: **initialize:**  
 Let  $d_2 = \dim(\mathcal{C})$ .  
 Set  $E_0 = \varepsilon \cdot I_{d_2}$ .
- 3: **for**  $k = 1, 2, \dots$ : **do**
- 4:     **receive** matrix  $A_k \in \mathbb{R}^{d_1 \times d_2}$ .
- 5:      $E_k \leftarrow E_{k-1} + A_k^{\top} \Sigma A_k$ .
- 6:      $z_k \leftarrow \arg \min_{z \in \mathcal{C}} \left\{ \langle z, -2 \sum_{i=1}^{k-1} A_i^{\top} \Sigma b_i \rangle + \|z\|_{E_k}^2 \right\}$
- 7:     **Play**  $z_k$  and **receive** feedback  $b_k \in \mathbb{R}^{d_1}$ .

In particular, if  $\sup_{z \in \mathcal{C}} \|z\| \leq B$ , then by setting  $\varepsilon = SY^2/B^2$  we are guaranteed that

$$\sum_{t=1}^T f_t(z_t) - \inf_{z \in \mathcal{C}} \sum_{t=1}^T f_t(z) \leq 5SY^2 \cdot d_2 \log(1 + B^2 R^2 Y^{-2} T / d_2). \quad (45)$$

**Proof of Theorem 5.** We assume  $\Sigma = I$  without loss of generality by reparameterizing via  $A'_t = \Sigma^{1/2} A_t$  and  $b'_t = \Sigma^{1/2} b_t$ , with  $Y$  and  $R$  scaled up by a factor of  $S^{1/2}$ .

Our proof follows the treatment of VAW in Orabona et al. (2015), which views the algorithm as an instance of online mirror descent with a sequence of time-varying regularizers. Consider the following algorithm parameterized by a sequence of convex regularizers  $\mathcal{R}_t : \mathcal{C} \rightarrow \mathbb{R}$ .

- Initialize  $\theta_1 = 0$ .
- For  $t = 1, \dots, T$ :
  - Let  $z_t = \arg \min_{z \in \mathcal{C}} \{ \langle z, \theta_t \rangle + \mathcal{R}_t(z) \}$ .
  - Receive  $g_t$  and set  $\theta_{t+1} = \theta_t + g_t$ .

The following lemma bounds the regret of this strategy for online linear optimization.

**Lemma E.2** (Orabona et al. (2015), Lemma 1). Suppose that each function  $\mathcal{R}_t$  is  $\beta$ -strongly convex with respect to a norm  $\|\cdot\|_t$ , and let  $\|\cdot\|_{t,*}$  denote the dual norm. Then the online mirror descent algorithm ensures that for every sequence  $g_1, \dots, g_T$ , for all  $z \in \mathcal{C}$ ,

$$\sum_{t=1}^T \langle g_t, z_t - z \rangle \leq \mathcal{R}_T(z) + \sum_{t=1}^T \left( \frac{\|g_t\|_{t,*}^2}{2\beta} + \mathcal{R}_{t-1}(z_t) - \mathcal{R}_t(z_t) \right). \quad (46)$$

Observe that the VAW algorithm (43) is equivalent to running online mirror descent with  $g_t = -2A_t^{\top} b_t$  and  $\mathcal{R}_t(z) = \|z\|_{E_t}^2$ . We use this observation to bound the regret through Lemma E.2. In particular, letting  $\|\cdot\|_t = \|\cdot\|_{E_t}$ , we may take  $\beta = 1$ , which gives

$$\begin{aligned} \sum_{t=1}^T f_t(z_t) - f_t(z) &= \sum_{t=1}^T \|A_t z_t - b_t\|^2 - \|A_t z - b_t\|^2 \\ &= \sum_{t=1}^T 2 \langle -A_t^{\top} b_t, z_t - z \rangle + \sum_{t=1}^T \|A_t z_t\|^2 - \sum_{t=1}^T \|A_t z\|^2 \\ &= \sum_{t=1}^T \langle g_t, z_t - z \rangle + \sum_{t=1}^T \|A_t z_t\|^2 - \mathcal{R}_T(z) + \varepsilon \|z\|^2 \\ &\leq \mathcal{R}_T(z) + \sum_{t=1}^T \left( \frac{1}{2} \|g_t\|_{E_t^{-1}}^2 + \mathcal{R}_{t-1}(z_t) - \mathcal{R}_t(z_t) \right) + \sum_{t=1}^T \|A_t z_t\|^2 - \mathcal{R}_T(z) + \varepsilon \|z\|^2 \\ &= \sum_{t=1}^T \left( \frac{1}{2} \|g_t\|_{E_t^{-1}}^2 + \mathcal{R}_{t-1}(z_t) - \mathcal{R}_t(z_t) \right) + \sum_{t=1}^T \|A_t z_t\|^2 + \varepsilon \|z\|^2, \end{aligned}$$

where the inequality uses Lemma E.2, along with the fact that the dual norm for  $\|\cdot\|_t$  is  $\|\cdot\|_{E_t^{-1}}$ . To simplify further, we observe that  $\mathcal{R}_{t-1}(z_t) - \mathcal{R}_t(z_t) = -\|A_t z_t\|^2$ , so that

$$\begin{aligned} \sum_{t=1}^T f_t(z_t) - f_t(z) &\leq \varepsilon \|z\|^2 + \frac{1}{2} \sum_{t=1}^T \|g_t\|_{E_t^{-1}}^2 \\ &= \varepsilon \|z\|^2 + 2 \sum_{t=1}^T \|A_t^\top b_t\|_{E_t^{-1}}^2 \leq \varepsilon \|z\|^2 + 2Y^2 \sum_{t=1}^T \|E_t^{-1/2} A_t^\top\|_{\text{op}}^2. \end{aligned}$$

To bound the right-hand side we use a generalization of the usual log-determinant potential argument. Throughout the argument we use that since  $E_t > A_t^\top A_t$ ,  $0 \leq \|E_t^{-1/2} A_t^\top\|_{\text{op}} < 1$ . To begin, observe that for each  $t$ , we have

$$\det(E_{t-1}) = \det(E_t - A_t^\top A_t) = \det(E_t) \cdot \det(I - E_t^{-1/2} A_t^\top A_t E_t^{-1/2}).$$

Consequently,

$$\begin{aligned} \frac{\det(E_t)}{\det(E_{t-1})} &= \frac{1}{\det(I - E_t^{-1/2} A_t^\top A_t E_t^{-1/2})} \\ &= \frac{1}{\prod_{i=1}^{d_2} (1 - \lambda_i(E_t^{-1/2} A_t^\top A_t E_t^{-1/2}))} = \prod_{i=1}^{d_2} \frac{1}{1 - \lambda_i(E_t^{-1/2} A_t^\top A_t E_t^{-1/2})}. \end{aligned}$$

Next we observe that since  $0 \leq \|E_t^{-1/2} A_t^\top\|_{\text{op}} < 1$ , we are guaranteed that  $\frac{1}{1 - \lambda_i(E_t^{-1/2} A_t^\top A_t E_t^{-1/2})} \geq 1$  for all  $i$ , and consequently

$$\prod_{i=1}^{d_2} \frac{1}{1 - \lambda_i(E_t^{-1/2} A_t^\top A_t E_t^{-1/2})} \geq \frac{1}{1 - \lambda_{\max}(E_t^{-1/2} A_t^\top A_t E_t^{-1/2})} \geq 1 + \lambda_{\max}(E_t^{-1/2} A_t^\top A_t E_t^{-1/2}),$$

where the second inequality uses that  $\frac{1}{1-x} \geq 1+x$  for  $x \in [0, 1)$ . Since  $\log x$  is increasing, this establishes that

$$\log\left(1 + \|E_t^{-1/2} A_t^\top\|_{\text{op}}^2\right) = \log\left(1 + \lambda_{\max}(E_t^{-1/2} A_t^\top A_t E_t^{-1/2})\right) \leq \log\left(\frac{\det(E_t)}{\det(E_{t-1})}\right).$$

Next we use that since  $\|E_t^{-1/2} A_t^\top\|_{\text{op}} \leq 1$ , we have

$$\|E_t^{-1/2} A_t^\top\|_{\text{op}}^2 \leq 2 \cdot \log\left(1 + \|E_t^{-1/2} A_t^\top\|_{\text{op}}^2\right),$$

using the elementary inequality  $x \leq 2 \log(1+x)$  for all  $x \in [0, 1]$ . Altogether, this gives

$$\sum_{t=1}^T \|E_t^{-1/2} A_t^\top\|_{\text{op}}^2 \leq 2 \sum_{t=1}^T \log\left(\frac{\det(E_t)}{\det(E_{t-1})}\right) = 2 \log\left(\frac{\det(E_T)}{\det(E_0)}\right),$$

where we recall  $E_0 = \varepsilon I$ . Finally, we have

$$\log\left(\frac{\det(E_T)}{\det(E_0)}\right) = \sum_{i=1}^{d_2} \log\left(1 + \lambda_i\left(\sum_{t=1}^T A_t^\top A_t\right)/\varepsilon\right) \leq d_2 \log(1 + R^2 T / (d_2 \varepsilon)).$$

□

#### E.4. Supporting lemmas for online learning

**Lemma 2.2** (cf. Joulani et al. (2013)). The generic delayed online learning reduction has regret at most

$$\sum_{t=1}^T f_t(z_t) - \inf_{z \in \mathcal{C}} \sum_{t=1}^T f_t(z) \leq (h+1)R(T/(h+1)),$$

where  $R(T)$  is the regret of the base instance.



**Proof of Lemma 2.2.** Let  $\mathcal{I}_i$  denote the rounds in which instance  $i$  was used. Then we have

$$\begin{aligned}
 \text{Reg}_T &= \sup_{z \in \mathcal{C}} \left\{ \sum_{t=1}^T f_t(z_t) - \sum_{t=1}^T f_t(z) \right\} \\
 &= \sup_{z \in \mathcal{C}} \left\{ \sum_{i=1}^{h+1} \sum_{t \in \mathcal{I}_i} f_t(z_t) - f_t(z) \right\} \\
 &\leq \sum_{i=1}^{h+1} \sup_{z \in \mathcal{C}} \left\{ \sum_{t \in \mathcal{I}_i} f_t(z_t) - f_t(z) \right\} \\
 &\leq \sum_{i=1}^{h+1} R(T/(h+1)) \\
 &= (h+1)R(T/(h+1)).
 \end{aligned}$$

□

**Lemma 2.3.** Let  $A \in \mathbb{R}^{d_1 \times d_2}$ , and consider the function  $f(z) = \|Az - b\|_{\Sigma}^2$ , where  $\Sigma \geq 0$ . If we restrict to  $z \in \mathbb{R}^{d_2}$  for which  $f(z) \leq R$ , then  $f$  is  $(2R)^{-1}$ -exp-concave.

**Proof of Lemma 2.3.** Recall that the function  $f$  is  $\alpha$ -exp-concave if and only if

$$\nabla^2 f(z) \geq \alpha \nabla f(z) \nabla f(z)^\top.$$

We have

$$\nabla f(z) = 2A^\top \Sigma (Az - b), \quad \text{and} \quad \nabla^2 f(z) = 2A^\top \Sigma A.$$

Hence

$$\nabla f(z) \nabla f(z)^\top \leq 4A^\top \Sigma A \|b - Az\|_{\Sigma}^2 \leq 2R \cdot \nabla^2 f(z).$$

□

## F. Proofs from Section 3

### F.1. Proof of Theorem 2

We restate [Theorem 2](#) here for reference.

**Theorem 2.** Let  $\pi$  be any policy of the form  $\pi_t(x; \mathbf{w}) = -K_\infty x - q^{M_t}(\mathbf{w}_{t-1})$ , where  $M_t = M_t(\mathbf{w}) \in \mathcal{M}_0$ . Then, by choosing  $h = 2(1 - \gamma_\infty)^{-1} \log(\kappa_\infty^2 \beta_\star^2 \Psi_\star \Gamma_\star^2 T^2)$  as the horizon parameter, we have

$$\sum_{t=1}^T \left| \mathbf{A}_t^\star(u_t^\pi; x_t^\pi, \mathbf{w}) - \widehat{\mathbf{A}}_{t,h}(M_t; \mathbf{w}_{t+h}) \right| \leq C_{\text{adv}},$$

where  $C_{\text{adv}} = \tilde{\mathcal{O}}(\beta_\star^{11} \Psi_\star^{19} \Gamma_\star^{11} \kappa_0^8 (1 - \gamma_0)^{-4} \log^2 T)$ .

**Proof of Theorem 2.** To begin, recall that by taking  $D_q$  as in (6), we have  $\|q_t^M\| \leq D_q$  for all  $M \in \mathcal{M}_0$ , and we also have  $D_{q^\star} \leq D_q$ .

For the first step, let  $\pi$  be any policy of the form  $\pi_t(x; \mathbf{w}) = -K_\infty x - q_t^{M_t}(\mathbf{w}_{t-1})$ , and let  $\widehat{\pi}_t(x; \mathbf{w}) = -K_t x - q_t^{M_t}(\mathbf{w}_{t-1})$  be the corresponding controller that uses the finite-horizon state-feedback matrices  $\{K_t\}_{t=1}^T$ . To begin, using the performance difference lemma (2) along with [Lemma 3.4](#),

$$\left| \sum_{t=1}^T \mathbf{A}_t^\star(u_t^\pi; x_t^\pi, \mathbf{w}) - \mathbf{A}_t^\star(u_t^{\widehat{\pi}}; x_t^{\widehat{\pi}}, \mathbf{w}) \right| \leq C_{K_\infty}.$$

Next, using Lemma 3.1, we have

$$\sum_{t=1}^T \mathbf{A}_t^*(u_t^{\widehat{\pi}}; x_t^{\widehat{\pi}}, \mathbf{w}) = \sum_{t=1}^T \|q^{M_t}(\mathbf{w}_{t-1}) - q_t^*(w_{t:T})\|_{\Sigma_t}^2.$$

Using Lemma 3.3, the choice of  $h$  in the theorem statement guarantees that

$$\sum_{t=1}^T \left| \mathbf{A}_t^*(u_t^{\widehat{\pi}}; x_t^{\widehat{\pi}}, \mathbf{w}) - \|q^{M_t}(\mathbf{w}_{t-1}) - q_{t;t+h}^*(w_{t:t+h})\|_{\Sigma_t}^2 \right| \leq C_{\text{trunc}},$$

and finally Lemma 3.5 ensures that

$$\begin{aligned} & \left| \sum_{t=1}^T \|q^{M_t}(\mathbf{w}_{t-1}) - q_{t;t+h}^*(w_{t:t+h})\|_{\Sigma_t}^2 - \|q^{M_t}(\mathbf{w}_{t-1}) - q_{\infty;h}^*(w_{t:t+h})\|_{\Sigma_\infty}^2 \right| \\ &= \left| \sum_{t=1}^T \|q^{M_t}(\mathbf{w}_{t-1}) - q_{t;t+h}^*(w_{t:t+h})\|_{\Sigma_t}^2 - \widehat{\mathbf{A}}_{t;h}(M_t; \mathbf{w}_{t+h}) \right| \\ &\leq C_\infty. \end{aligned}$$

Summing up all the error terms, the total error is proportional to

$$\begin{aligned} & C_{K_\infty} + C_{\text{trunc}} + C_\infty \\ &= \widetilde{\mathcal{O}}(\kappa_\infty^4 \beta_\star^6 \Psi_\star^{13} \Gamma_\star^6 (1 - \gamma_\infty)^{-2} D_q^2 \cdot \log(D_q T)) + \widetilde{\mathcal{O}}(D_q^2 \beta_\star \Psi_\star^4 \Gamma_\star^2 (1 - \gamma_\infty)^{-1} \log T) \\ &\quad + \widetilde{\mathcal{O}}(D_q^2 \cdot \beta_\star^4 \Psi_\star^7 \Gamma_\star^4 \kappa_\infty^2 (1 - \gamma_\infty)^{-1} h \log(D_q T)). \end{aligned}$$

Using the value for  $D_q$  from (6) and that  $h = \widetilde{\mathcal{O}}((1 - \gamma_\infty)^{-1} \log T)$ , we upper bound the total error as

$$\widetilde{\mathcal{O}}(\beta_\star^{11} \Psi_\star^{19} \Gamma_\star^{11} \kappa_0^8 (1 - \gamma_0)^{-4} \log^2 T).$$

□

## F.2. Supporting lemmas

**Lemma 3.2.** For any  $h \in [T]$  define a truncated version of  $q_t^*$  as follows:

$$q_{t;t+h}^*(w_{t:t+h}) = \sum_{i=t}^{(t+h) \wedge T-1} \Sigma_t^{-1} B^\top \left( \prod_{j=t+1}^i A_{\text{cl},j}^\top \right) P_{i+1} w_i. \quad (14)$$

Then for any  $t$  such that  $t + h < T - \widetilde{\mathcal{O}}(\beta_\star \Psi_\star^2 \Gamma_\star)$ , setting  $\bar{\gamma}_\infty = \frac{1}{2}(1 + \gamma_\infty) < 1$ , we have the bound  $\|q_{t;t+h}^*(w_{t:t+h}) - q_t^*(w_{t:T})\| \leq \kappa_\infty^2 \beta_\star^2 \Psi_\star \Gamma_\star^2 (T - h) \bar{\gamma}_\infty^h$ , which is geometrically decreasing in  $h$ .

**Proof of Lemma 3.2.** Let  $\tau = t + h$ . Then we have

$$q_{t;\tau}^*(w_{t;\tau}) - q_{t:T}^*(w_{t:T}) = \sum_{i=\tau+1}^{T-1} \Sigma_t^{-1} B^\top \left( \prod_{j=t+1}^i A_{\text{cl},j}^\top \right) P_{i+1} w_i,$$

Hence we can bound the error as

$$\|q_{t;\tau}^*(w_{t;\tau}) - q_{t:T}^*(w_{t:T})\| = \beta_\star \Psi_\star \Gamma_\star \sum_{i=\tau+1}^{T-1} \left\| \prod_{j=t+1}^i A_{\text{cl},j}^\top \right\|_{\text{op}}.$$

We bound each term in the sum as

$$\left\| \prod_{j=t+1}^i A_{\text{cl},j}^\top \right\|_{\text{op}} \leq \left\| \prod_{j=t+1}^{\tau+1} A_{\text{cl},j}^\top \right\|_{\text{op}} \left\| \prod_{j=\tau+1}^i A_{\text{cl},j}^\top \right\|_{\text{op}}.$$

Applying Lemma D.5 to the first term and Lemma D.3 to the second, this is at most

$$\leq \kappa_\infty^2 \beta_\star \Gamma_\star \bar{\gamma}_\infty^{\tau-t}.$$

The result follows by summing.  $\square$

**Lemma 3.3.** Consider a policy  $\hat{\pi}_t(x; \mathbf{w}) = -K_t x_t - q_t(\mathbf{w})$ , and suppose that  $\|q_t\| \leq D_q$ , where  $D_q \geq D_{q^\star}$ . If we choose  $h = 2(1 - \gamma_\infty)^{-1} \log(\kappa_\infty^2 \beta_\star^2 \Psi_\star \Gamma_\star^2 T^2)$ , we are guaranteed that

$$\sum_{t=1}^T \left| \mathbf{A}_t^\star(u_t^{\hat{\pi}}; x_t^{\hat{\pi}}, \mathbf{w}) - \|q_t(\mathbf{w}) - q_{t;t+h}^\star(w_{t:t+h})\|_{\Sigma_t}^2 \right| \leq C_{\text{trunc}},$$

where  $C_{\text{trunc}} \leq \tilde{\mathcal{O}}(D_q^2 \beta_\star \Psi_\star^4 \Gamma_\star^2 (1 - \gamma_\infty)^{-1} \log T)$ .

**Proof of Lemma 3.3.** First recall that we have  $\|\Sigma_t\|_{\text{op}} \leq \|R\|_{\text{op}} + \|B\|_{\text{op}}^2 \|P_\infty\|_{\text{op}} \leq 2\Psi_\star^2 \Gamma_\star =: D_\Sigma$ . Let  $h$  be fixed, and let  $T_{\text{trunc}} := T_{\text{stab}} - h$ , so that  $t + h \leq T_{\text{stab}}$  for all  $t \leq T_{\text{trunc}}$ . We begin by writing off all of the timesteps after  $T_{\text{trunc}}$ :

$$\begin{aligned} & \sum_{t=1}^T \left| \mathbf{A}_t^\star(\hat{\pi}_t(x_t^{\hat{\pi}}); x_t^{\hat{\pi}}, \mathbf{w}) - \|q_t(\mathbf{w}) - q_{t;t+h}^\star(w_{t:t+h})\|_{\Sigma_t}^2 \right| \\ &= \sum_{t=1}^T \left| \|q_t(\mathbf{w}) - q_t^\star(w_{t:T})\|_{\Sigma_t}^2 - \|q_t(\mathbf{w}) - q_{t;t+h}^\star(w_{t:t+h})\|_{\Sigma_t}^2 \right| \\ &\leq \sum_{t=1}^{T_{\text{trunc}}} \left| \|q_t(\mathbf{w}) - q_t^\star(w_{t:T})\|_{\Sigma_t}^2 - \|q_t(\mathbf{w}) - q_{t;t+h}^\star(w_{t:t+h})\|_{\Sigma_t}^2 \right| + 4D_q^2 D_\Sigma (\Delta_{\text{stab}} + h) \\ &\leq 4D_q D_\Sigma \sum_{t=1}^{T_{\text{trunc}}} \left\| q_t^\star(w_{t:T}) - q_{t;t+h}^\star(w_{t:t+h}) \right\|_{\Sigma_t} + 4D_q^2 D_\Sigma (\Delta_{\text{stab}} + h). \end{aligned}$$

Since  $t + h \leq T_{\text{stab}}$  for all  $t$  in the last summation, Lemma 3.2 implies that

$$\sum_{t=1}^{T_{\text{trunc}}} \left\| q_t^\star(w_{t:T}) - q_{t;t+h}^\star(w_{t:t+h}) \right\|_{\Sigma_t} \leq \kappa_\infty^2 \beta_\star^2 \Psi_\star \Gamma_\star^2 T^2 \bar{\gamma}_\infty^h \leq \kappa_\infty^2 \beta_\star^2 \Psi_\star \Gamma_\star^2 T^2 \exp(-(1 - \bar{\gamma}_\infty)h).$$

By choosing  $h = (1 - \bar{\gamma}_\infty)^{-1} \log(\kappa_\infty^2 \beta_\star^2 \Psi_\star \Gamma_\star^2 T^2)$ , the total error from this term is  $\mathcal{O}(1)$ . Combining this with the previous bound, we see that the total error is at most

$$\mathcal{O}(D_q D_\Sigma + D_q^2 D_\Sigma (\Delta_{\text{stab}} + h)).$$

Lastly, we simplify by using that  $\Delta_{\text{stab}} = \tilde{\mathcal{O}}(\beta_\star \Psi_\star^2 \Gamma_\star)$  and expanding  $h$  and  $D_\Sigma$ , so that the final error term is at most

$$\tilde{\mathcal{O}}(D_q^2 \beta_\star \Psi_\star^4 \Gamma_\star^2 (1 - \gamma_\infty)^{-1} \log T).$$

$\square$

**Lemma 3.4.** Let policies  $\pi_t(x; \mathbf{w}) = -K_\infty x - q_t(\mathbf{w})$  and  $\hat{\pi}_t(x; \mathbf{w}) = -K_t x - q_t(\mathbf{w})$  be given, where  $q_t$  is arbitrary but satisfies  $\|q_t\| \leq D_q$  for some  $D_q \geq 1$ . Then

$$|J_T(\hat{\pi}, \mathbf{w}) - J_T(\pi, \mathbf{w})| \leq C_{K_\infty},$$

where  $C_{K_\infty} \leq \tilde{\mathcal{O}}(\kappa_\infty^4 \beta_\star^6 \Psi_\star^{13} \Gamma_\star^6 (1 - \gamma_\infty)^{-2} D_q^2 \cdot \log(D_q T))$ .

**Proof of Lemma 3.4.** To begin, suppose that the states under both controllers satisfy  $\|x\| \leq D_x$  and the actions satisfy  $\|u\| \leq D_u$ , where  $D_x, D_u \geq 1$ . Then, we immediately have

$$|J_T(\hat{\pi}, \mathbf{w}) - J_T(\pi, \mathbf{w})| \leq 2\Psi_\star \sum_{t=1}^T D_x \|x_t^{\hat{\pi}}(\mathbf{w}) - x_t^\pi(\mathbf{w})\| + D_u \|u_t^{\hat{\pi}}(\mathbf{w}) - u_t^\pi(\mathbf{w})\|,$$

which follows because the function  $x \mapsto \|x\|^2$  is  $2C$ -Lipschitz whenever  $\|x\| \leq C$ . We will first bound the state and action errors on the right hand side, then give appropriate bounds on  $D_x$  and  $D_u$  at the end of the proof.

Let  $\Delta_0$  be fixed, and let  $T_0 = T - \Delta_0$ . Then we can bound the error further as

$$\begin{aligned} & |J_T(\widehat{\pi}; \mathbf{w}) - J_T(\pi; \mathbf{w})| \\ & \leq 2\Psi_* \sum_{t=1}^T D_x \|x_t^{\widehat{\pi}}(\mathbf{w}) - x_t^\pi(\mathbf{w})\| + D_u \|u_t^{\widehat{\pi}}(\mathbf{w}) - u_t^\pi(\mathbf{w})\| \\ & \leq 4\Psi_*(D_x^2 + D_u^2)\Delta_0 + 2\Psi_* \sum_{t=1}^{T_0} D_x \|x_t^{\widehat{\pi}}(\mathbf{w}) - x_t^\pi(\mathbf{w})\| + D_u \|u_t^{\widehat{\pi}}(\mathbf{w}) - u_t^\pi(\mathbf{w})\|, \end{aligned}$$

For the control error term, we further have

$$\begin{aligned} \sum_{t=1}^{T_0} \|u_t^{\widehat{\pi}}(\mathbf{w}) - u_t^\pi(\mathbf{w})\| & = \sum_{t=1}^{T_0} \|K_t x_t^{\widehat{\pi}}(\mathbf{w}) - K_\infty x_t^\pi(\mathbf{w})\| \\ & \leq \sum_{t=1}^{T_0} \|K_\infty\| \|x_t^{\widehat{\pi}}(\mathbf{w}) - x_t^\pi(\mathbf{w})\| + D_x \sum_{t=1}^{T_0} \|K_t - K_\infty\|_{\text{op}} \\ & \leq \sum_{t=1}^{T_0} \beta_* \Psi_*^2 \Gamma_* \|x_t^{\widehat{\pi}}(\mathbf{w}) - x_t^\pi(\mathbf{w})\| + D_x \sum_{t=1}^{T_0} \|K_t - K_\infty\|_{\text{op}}. \end{aligned}$$

In total, this gives us

$$\begin{aligned} & |J_T(\widehat{\pi}; \mathbf{w}) - J_T(\pi; \mathbf{w})| \\ & \leq 4\Psi_*(D_x^2 + D_u^2)\Delta_0 + 2\Psi_* \sum_{t=1}^{T_0} (D_x + D_u \beta_* \Psi_*^2 \Gamma_*) \|x_t^{\widehat{\pi}}(\mathbf{w}) - x_t^\pi(\mathbf{w})\| + D_x D_u \|K_t - K_\infty\|_{\text{op}}. \end{aligned}$$

To bound the state error, we recall that from [Lemma D.7](#), we have

$$x_{t+1}^{\widehat{\pi}}(\mathbf{w}_t) - x_{t+1}^\pi(\mathbf{w}_t) = \sum_{i=1}^t (A_{\text{cl},i \rightarrow t} - A_{\text{cl},\infty}^{t-i})(w_i - Bq_i(\mathbf{w}_{i-1})),$$

and so

$$\|x_{t+1}^{\widehat{\pi}}(\mathbf{w}_t) - x_{t+1}^\pi(\mathbf{w}_t)\| \leq 2\Psi_* D_q \sum_{i=1}^t \|A_{\text{cl},i \rightarrow t} - A_{\text{cl},\infty}^{t-i}\|_{\text{op}}.$$

To bound the error, we recall [Lemma D.8](#), restated here.

**Lemma D.8.** Let  $\alpha \geq 1$  be given. Define  $\Delta = C \cdot \beta_* \Psi_*^2 \Gamma_* \log(\kappa_\infty^2 \Psi_* \Gamma_* (1 - \gamma_\infty)^{-1} \cdot \alpha T^3)$ , where  $C > 0$  is a numerical constant. If  $C$  is sufficiently large, then for every  $t \leq T - \Delta \leq T_{\text{stab}}$  we are guaranteed that

$$\|K_t - K_\infty\|_{\text{op}} \leq \frac{1}{\kappa_\infty^2 \Psi_* \cdot (\alpha T^3)}, \quad \text{and} \quad \|A_{\text{cl},i \rightarrow t} - A_{\text{cl},\infty}^{t-i}\|_{\text{op}} \leq \frac{1}{\alpha T^2} \quad \forall t \leq T - \Delta. \quad (30)$$

We set  $\alpha = 4D_x D_u \beta_* \Psi_*^4 \Gamma_* D_q$ , which ensures that

$$|J_T(\widehat{\pi}; \mathbf{w}) - J_T(\pi; \mathbf{w})| \leq 4\Psi_*(D_x^2 + D_u^2)\Delta_0 + C',$$

where  $C'$  is an absolute numerical constant. To conclude, we recall from [Lemma D.9](#) that it suffices to take  $D_x \leq \widetilde{\mathcal{O}}(\kappa_\infty^2 \beta_*^{3/2} \Psi_*^3 \Gamma_*^{3/2} (1 - \gamma_\infty)^{-1} \cdot D_q)$  and  $D_u \leq \widetilde{\mathcal{O}}(\kappa_\infty^2 \beta_*^{5/2} \Psi_*^5 \Gamma_*^{5/2} (1 - \gamma_\infty)^{-1} \cdot D_q)$ . □

**Lemma 3.5.** Let  $(q_t)_{t=1}^T$  be an arbitrary sequence with  $\|q_t\| \leq D_q$  for some  $D_q \geq D_{q^*}$ . Then it holds that

$$\left| \sum_{t=1}^T \|q_t - q_{t;t+h}^*(w_{t:t+h})\|_{\Sigma_t}^2 - \|q_t - q_{\infty;t+h}^*(w_{t:t+h})\|_{\Sigma_\infty}^2 \right| \leq C_\infty,$$

where  $C_\infty \leq \widetilde{\mathcal{O}}(D_q^2 \cdot \beta_*^4 \Psi_*^7 \Gamma_*^4 \kappa_\infty^2 (1 - \gamma_\infty)^{-1} h \log(D_q T))$ .

**Proof of Lemma 3.5.** Before diving into the proof, we recall that, since  $P_t \leq P_\infty$ , we have

$$\|\Sigma_t\|_{\text{op}} \leq \|\Sigma_\infty\|_{\text{op}} = \|R_x + B^\top P_\infty B\|_{\text{op}} \leq 2\Psi_\star^2 \Gamma_\star =: D_\Sigma.$$

We also recall that  $D_q \geq D_{q^\star}$ . Now let  $\Delta_0 \in \mathbb{N}$  be a fixed constant to be chosen later, and let  $T_0 = T - \Delta_0$ . We immediately upper bound the error as

$$\begin{aligned} & \left| \sum_{t=1}^T \|q_t - q_{t;t+h}^\star(w_{t:t+h})\|_{\Sigma_t}^2 - \|q_t - q_{\infty;h}^\star(w_{t:t+h})\|_{\Sigma_\infty}^2 \right| \\ & \leq \left| \sum_{t=1}^{T_0} \|q_t - q_{t;t+h}^\star(w_{t:t+h})\|_{\Sigma_t}^2 - \|q_t - q_{\infty;h}^\star(w_{t:t+h})\|_{\Sigma_\infty}^2 \right| + 4D_\Sigma D_q^2 \Delta_0. \end{aligned}$$

Now, let  $t \leq T_0$  be fixed. We upper bound the error for each time as

$$\begin{aligned} & \left| \|q_t - q_{t;t+h}^\star(w_{t:t+h})\|_{\Sigma_t}^2 - \|q_t - q_{\infty;h}^\star(w_{t:t+h})\|_{\Sigma_\infty}^2 \right| \\ & \leq \left| \|q_t - q_{t;t+h}^\star(w_{t:t+h})\|_{\Sigma_t}^2 - \|q_t - q_{t;t+h}^\star(w_{t:t+h})\|_{\Sigma_\infty}^2 \right| \\ & \quad + \left| \|q_t - q_{t;t+h}^\star(w_{t:t+h})\|_{\Sigma_\infty}^2 - \|q_t - q_{\infty;h}^\star(w_{t:t+h})\|_{\Sigma_\infty}^2 \right| \\ & \leq D_q^2 \underbrace{\|\Sigma_t - \Sigma_\infty\|_{\text{op}}}_{\mathcal{E}_1} + 4D_q D_\Sigma \underbrace{\|q_{t;t+h}^\star(w_{t:t+h}) - q_{\infty;h}^\star(w_{t:t+h})\|}_{\mathcal{E}_2}. \end{aligned}$$

**Bounding  $\mathcal{E}_1$ .** Expanding the definition of  $\Sigma_t$  and  $\Sigma_\infty$ , we immediately see that  $\|\Sigma_t - \Sigma_\infty\|_{\text{op}} \leq \Psi_\star^2 \|P_{t+1} - P_\infty\|_{\text{op}}$ . Using Lemma D.11, we have

$$\|P_{t+1} - P_\infty\|_{\text{op}}^2 \leq \beta_\star \Gamma_\star (1 + \nu_\star^{-1})^{-(T-t)},$$

where  $\nu_\star = 2\beta_\star \Psi_\star^2 \Gamma_\star$ . Hence, summing across all rounds, we have

$$\sum_{t=1}^T \|P_t - P_\infty\|_{\text{op}} \leq \sum_{t=1}^T \beta_\star^{1/2} \Gamma_\star^{1/2} (1 + \nu_\star^{-1})^{-(T-t)/2}.$$

Since  $\nu_\star^{-1} \leq 1$  and  $1 + x \geq e^{x/2}$  for  $x \in [0, 1]$ , we can upper bound by

$$\begin{aligned} & \leq \beta_\star^{1/2} \Gamma_\star^{1/2} \sum_{t=1}^T e^{-\nu_\star^{-1}(T-t)/4} \\ & \leq O(\beta_\star^{1/2} \Gamma_\star^{1/2} \nu_\star) \\ & = O(\beta_\star^{3/2} \Psi_\star^2 \Gamma_\star^{3/2}), \end{aligned} \tag{47}$$

and  $\sum_{t=1}^T \|\Sigma_t - \Sigma_\infty\|_{\text{op}} \leq O(\beta_\star^{3/2} \Psi_\star^4 \Gamma_\star^{3/2})$ .

**Bounding  $\mathcal{E}_2$ .** Let  $t \leq T_0 \leq T - h$  be fixed, then we have

$$\begin{aligned} & \|q_{t;t+h}^\star(w_{t:t+h}) - q_{\infty;h}^\star(w_{t:t+h})\| \\ & = \left\| \sum_{i=t}^{t+h} \Sigma_t^{-1} B^\top \left( \prod_{j=t+1}^i A_{\text{cl},j}^\top \right) P_{i+1} w_i - \sum_{i=t}^{t+h} \Sigma_\infty^{-1} B^\top (A_{\text{cl},\infty}^\top)^{i-t} P_\infty w_i \right\| \\ & = \left\| \sum_{i=t}^{t+h} \Sigma_t^{-1} B^\top A_{\text{cl},t \rightarrow i}^\top P_{i+1} w_i - \sum_{i=t}^{t+h} \Sigma_\infty^{-1} B^\top (A_{\text{cl},\infty}^\top)^{i-t} P_\infty w_i \right\| \\ & \leq \sum_{i=t}^{t+h} \left\| \Sigma_t^{-1} B^\top A_{\text{cl},t \rightarrow i}^\top P_{i+1} - \Sigma_\infty^{-1} B^\top (A_{\text{cl},\infty}^\top)^{i-t} P_\infty \right\|_{\text{op}}. \end{aligned}$$



Note that for each timestep we have

$$\begin{aligned} & \left\| \Sigma_t^{-1} B^\top A_{\text{cl},t \rightarrow i}^\top P_{i+1} - \Sigma_\infty^{-1} B^\top (A_{\text{cl},\infty}^\top)^{i-t} P_\infty \right\|_{\text{op}} \\ & \leq \left\| (\Sigma_t^{-1} - \Sigma_\infty^{-1}) B^\top A_{\text{cl},t \rightarrow i}^\top P_{i+1} \right\|_{\text{op}} + \left\| \Sigma_\infty^{-1} B^\top (A_{\text{cl},t \rightarrow i}^\top - (A_{\text{cl},\infty}^\top)^{i-t}) P_{i+1} \right\|_{\text{op}} \\ & \quad + \left\| \Sigma_\infty^{-1} B^\top (A_{\text{cl},\infty}^\top)^{i-t} (P_{i+1} - P_\infty) \right\|_{\text{op}}. \end{aligned}$$

If we select  $T_0 \leq T_{\text{stab}} - h$ , then we are guaranteed by [Lemma D.5](#) that  $\|A_{\text{cl},t \rightarrow i}\|_{\text{op}} \leq \beta_\star^{1/2} \Gamma_\star^{1/2} \kappa_\infty^2 \bar{\gamma}_\infty^{i-t}$ , and we also know that  $\|A_{\text{cl},\infty}^{i-t}\|_{\text{op}} \leq \kappa_\infty \bar{\gamma}_\infty^{i-t}$ . Hence, we can upper bound the errors above by

$$\begin{aligned} & \beta_\star^{1/2} \Psi_\star \Gamma_\star^{3/2} \kappa_\infty^2 \bar{\gamma}_\infty^{i-t} \|\Sigma_t^{-1} - \Sigma_\infty^{-1}\|_{\text{op}} + \beta_\star \Psi_\star \Gamma_\star \|A_{\text{cl},t \rightarrow i}^\top - (A_{\text{cl},\infty}^\top)^{i-t}\|_{\text{op}} \\ & \quad + \beta_\star \Psi_\star \kappa_\infty \bar{\gamma}_\infty^{i-t} \|P_{i+1} - P_\infty\|_{\text{op}}. \end{aligned}$$

Furthermore, recall that  $\Sigma_t = R_x + B^\top P_{t+1} B \geq R_x$  and  $\Sigma_\infty = R_x + B^\top P_\infty B \geq R_x$ , and so we have

$$\|\Sigma_t^{-1} - \Sigma_\infty^{-1}\|_{\text{op}} \leq \beta_\star^2 \Psi_\star^2 \|P_{t+1} - P_\infty\|_{\text{op}}.$$

Putting everything together this gives

$$\begin{aligned} & \sum_{t=1}^{T_0} \left\| q_{t;t+h}^\star(w_{t:t+h}) - q_{\infty;h}^\star(w_{t:t+h}) \right\| \\ & \leq 2\beta_\star^{5/2} \Psi_\star^3 \Gamma_\star^{3/2} \kappa_\infty^2 (1 - \gamma_\infty)^{-1} (h+1) \sum_{t=1}^{T_0} \|P_{t+1} - P_\infty\|_{\text{op}} + \beta_\star \Psi_\star \Gamma_\star \sum_{t=1}^{T_0} \sum_{i=t}^{t+h} \|A_{\text{cl},t \rightarrow i}^\top - (A_{\text{cl},\infty}^\top)^{i-t}\|_{\text{op}} \\ & \quad + \beta_\star \Psi_\star \kappa_\infty \sum_{t=1}^{T_0} \sum_{i=t}^{t+h} \gamma_\infty^{i-t} \|P_{i+1} - P_\infty\|_{\text{op}} \\ & \leq 4\beta_\star^{5/2} \Psi_\star^3 \Gamma_\star^{3/2} \kappa_\infty^2 (1 - \gamma_\infty)^{-1} (h+1) \sum_{t=1}^{T_0} \|P_{t+1} - P_\infty\|_{\text{op}} + \beta_\star \Psi_\star \Gamma_\star \sum_{t=1}^{T_0} \sum_{i=t}^{t+h} \|A_{\text{cl},t \rightarrow i}^\top - (A_{\text{cl},\infty}^\top)^{i-t}\|_{\text{op}}. \end{aligned}$$

Recalling [\(47\)](#), we can further upper bound the first term:

$$\leq O\left(\beta_\star^4 \Psi_\star^5 \Gamma_\star^3 \kappa_\infty^2 (1 - \gamma_\infty)^{-1} h\right) + \beta_\star \Psi_\star \Gamma_\star \sum_{t=1}^{T_0} \sum_{i=t}^{t+h} \|A_{\text{cl},t \rightarrow i}^\top - (A_{\text{cl},\infty}^\top)^{i-t}\|_{\text{op}}.$$

To bound the last term, we recall [Lemma D.8](#).

**Lemma D.8.** Let  $\alpha \geq 1$  be given. Define  $\Delta = C \cdot \beta_\star \Psi_\star^2 \Gamma_\star \log(\kappa_\infty^2 \Psi_\star \Gamma_\star (1 - \gamma_\infty)^{-1} \cdot \alpha T^3)$ , where  $C > 0$  is a numerical constant. If  $C$  is sufficiently large, then for every  $t \leq T - \Delta \leq T_{\text{stab}}$  we are guaranteed that

$$\|K_t - K_\infty\|_{\text{op}} \leq \frac{1}{\kappa_\infty^2 \Psi_\star \cdot (\alpha T^3)}, \quad \text{and} \quad \|A_{\text{cl},i \rightarrow t} - A_{\text{cl},\infty}^{t-i}\|_{\text{op}} \leq \frac{1}{\alpha T^2} \quad \forall t \leq T - \Delta. \quad (30)$$

We choose  $\alpha = \beta_\star \Psi_\star \Gamma_\star D_q D_\Sigma$ , and set  $\Delta_0 = C \cdot \beta_\star \Psi_\star^2 \Gamma_\star \log(\kappa_\infty^2 \Psi_\star \Gamma_\star (1 - \gamma_\infty)^{-1} \cdot \alpha T^3) \vee \Delta_{\text{stab}} + h$ , so we are ensured that

$$\beta_\star \Psi_\star \Gamma_\star \sum_{t=1}^{T_0} \sum_{i=t}^{t+h} \|A_{\text{cl},t \rightarrow i}^\top - (A_{\text{cl},\infty}^\top)^{i-t}\|_{\text{op}} \leq C \cdot \frac{1}{D_q D_\Sigma}.$$

Putting everything together leads to a final bound of

$$\begin{aligned} & O\left(D_q^2 \beta_\star^{3/2} \Psi_\star^4 \Gamma_\star^{3/2}\right) + O\left(D_q D_\Sigma \beta_\star^4 \Psi_\star^5 \Gamma_\star^3 \kappa_\infty^2 (1 - \gamma_\infty)^{-1} h\right) + O\left(D_\Sigma D_q^2 \Delta_0\right) \\ & = \tilde{O}\left(D_q^2 \beta_\star^{3/2} \Psi_\star^4 \Gamma_\star^{3/2} + D_q \beta_\star^4 \Psi_\star^7 \Gamma_\star^4 \kappa_\infty^2 (1 - \gamma_\infty)^{-1} h + D_q^2 \cdot \beta_\star \Psi_\star^4 \Gamma_\star^2 \log(D_q T)\right). \end{aligned}$$

□