
Online learning with dynamics: A minimax perspective

Kush Bhatia
EECS, UC Berkeley
kush@cs.berkeley.edu

Karthik Sridharan
CS, Cornell University
sridharan@cs.cornell.edu

Abstract

We study the problem of *online learning with dynamics*, where a learner interacts with a stateful environment over multiple rounds. In each round of the interaction, the learner selects a policy to deploy and incurs a cost that depends on both the chosen policy and current state of the world. The state-evolution dynamics and the costs are allowed to be time-varying, in a possibly adversarial way. In this setting, we study the problem of minimizing policy regret and provide non-constructive upper bounds on the minimax rate for the problem.

Our main results provide sufficient conditions for online learnability for this setup with corresponding rates. The rates are characterized by: 1) a complexity term capturing the expressiveness of the underlying policy class under the dynamics of state change, and 2) a dynamic stability term measuring the deviation of the instantaneous loss from a certain counterfactual loss. Further, we provide matching lower bounds which show that both the complexity terms are indeed necessary.

Our approach provides a unifying analysis that recovers regret bounds for several well studied problems including online learning with memory, online control of linear quadratic regulators, online Markov decision processes, and tracking adversarial targets. In addition, we show how our tools help obtain tight regret bounds for a new problems (with non-linear dynamics and non-convex losses) for which such bounds were not known prior to our work.

1 Introduction

Machine learning systems deployed in the real-world interact with people through their decision making. Such systems form a feedback loop with their environment: they learn to make decisions from real-world data and decisions made by these systems in turn affect the data that is collected. In addition, people often learn to adapt to such automated decision makers in an attempt to maximize their own utility rendering any assumption on the data generation process futile. Motivated by these aspects of decision making, we propose the problem of *online learning with dynamics* which involves repeated interaction between a learner and an environment with an underlying state. The decisions made by the learner affect this state of the environment which evolves as a dynamical system. Further, we place no distributional assumptions on the learning data and allow this to be adversarial.

Given such a setup, a natural question to ask is how does one measure the performance of the learner? Classical online learning studies one such notion of performance known as regret. This measure compares the performance of the learner to that of a fixed best policy in hindsight, when evaluated on the *same* states which were observed by the learner. Such a measure of performance clearly does not work for the above setup: if we would have deployed a different policy, we would have observed different states of the environment. To overcome this, we study a counterfactual notion of regret, called *Policy Regret*, where the comparator term is the performance of a policy on the states one would have observed if this policy was deployed from the beginning of time.

Such a notion of regret has been studied in the online learning literature for understanding memory based adversaries [20, 4, 5] and more recently, for the study of specific reinforcement learning models [11, 1, 10]. However, a vast majority of these works have focused on known and fixed models of state evolution, often restricting the scope to linear dynamical systems. Further, these works have focused on simplistic policy classes as the comparators in their notion of policy regret. Contrast this with the vast literature on statistical learning [28, 6] and classical online learning [22] which study the question of learnability in full generality; for arbitrary losses and general function classes.

Our work is a step towards addressing this gap. We study the problem of learnability for a class of online learning problems with underlying states evolving as a dynamical system in its full generality. Our main results provide sufficient conditions (along with non-asymptotic upper bounds) on when such problems are learnable, that is, can have vanishing policy regret. Our approach is non-constructive and provides a complexity term that provides upper bounds on the minimax rates for these problems. Further, we provide lower bounds showing that for a large class of problems, our upper bounds are tight up to constant factors. By studying the problem in full generality, we show how several well-studied problems in the literature comprising online Markov decision processes [11], online adversarial tracking [1], online linear quadratic regulator [10], online control with adversarial noise [2], and online learning with memory [5, 4] can be seen as specific examples of our general framework. We recover the best known rates for a majority of these problems, often times even generalizing these setups. We also provide examples where, to the best of our knowledge, previous techniques are not able to obtain useful bounds on regret; however using our minimax tools, we are able to provide tight bounds on the policy regret for these examples.

Formally, we consider the setup where \mathcal{X} denotes an arbitrary set of states, Π an arbitrary class of policies and \mathcal{Z} an arbitrary instance space. Given this, the interaction between the learner and nature can be expressed as a T round protocol where on each round $t \in [T]$, the learner picks a policy $\pi_t \in \Pi$, the adversary simultaneously picks instance $(z_t, \zeta_t) \in \mathcal{Z}$. The learner suffers loss $\ell(\pi_t, x_t, z_t)$ and the state of the system evolves¹ as $x_{t+1} \leftarrow \Phi(x_t, \pi_t, \zeta_t)$, where Φ is known to the learner. The goal of the learner is to minimize policy regret

$$\text{Reg}_T^{\text{pol}} = \sum_{t=1}^T \ell(\pi_t, x_t, z_t) - \inf_{\pi \in \Pi} \sum_{t=1}^T \ell(\pi, x_t[\pi^{(t-1)}, \zeta_{1:t-1}], z_t) ,$$

where x_t are the states of the system based on learners choices of policies and $x_t[\pi^{(t-1)}, \zeta_{1:t-1}]$ represents the state of the system at time t if the policy π was used the previous $t-1$ rounds. We refer to the loss $\ell(\pi, x_t[\pi^{(t-1)}, \zeta_{1:t-1}], z_t)$ as the *counterfactual loss* of policy π . Notice that dynamics Φ being fixed or known in advance to the learner is not really restrictive since an adversary can encode arbitrary state dynamics mapping in ζ_t 's and Φ can just be seen as an applicator of these mapping.

Our contributions. We are interested in the following question: for a given problem instance $(\Pi, \mathcal{Z}, \Phi, \ell)$, is the problem learnable, that is, does there exists a learning algorithm such that policy regret is such that $\text{Reg}_T^{\text{pol}} = o(T)$. Below we highlight some of the key contributions of this paper.

1. We show that the minimax policy regret for any problem specified by $(\Pi, \mathcal{Z}, \Phi, \ell)$ can be upper bounded by sum of two terms: i) a sequential Rademacher complexity like term for the class of counterfactual losses of the policy class, and ii) a term we refer to as dynamic stability term for the Empirical Risk Minimizer (ERM) (or regularized ERM) algorithm.
2. We analyze the problem in the dual game. While in most cases ERM does not even have low classical regret let alone policy regret, we show that ERM like strategy in the dual game can lead to the two term decomposition of minimax policy regret we mention above.
3. Ours is the first work that studies arbitrary online dynamical systems, and provides an analysis for general policy classes and loss functions (possibly non-convex).
4. We provide lower bounds that show that our sufficient conditions are tight for a large class of problem instances showing that both the terms in our upper bounds are indeed necessary.
5. We delineate a number of previously studied problems including online linear quadratic regulator, and online learning with memory for which we recover rates. More importantly, we provide examples of new non-convex and general online learning with dynamics problems and obtain tight regret bounds. For these examples, none of the previous methods are able to obtain any non-degenerate regret bounds.

¹while we consider deterministic dynamics here, Section 2 considers general dynamics with stochastic noise

Related work. The classical online learning setup [8] considers a repeated interactive game between a learner and an environment without any notion of underlying dynamics. Sequential complexity measures were introduced in [22] to get tight characterization of minimax regret rates for the classical online learning setting. They showed that for the class of online supervised learning problems, one can upper and lower bound minimax rate in terms of a sequential Rademacher complexity of the predictor class. The works [18, 7] provided an analog of VC theory for online classification and the sequential complexity measures in work [22] provided such a theory for general supervised online learning. This paper can be seen as deriving such characterization of learnability and tight rates for the problem of online learning with dynamics. In the more general setting we consider, while the main mathematical tools introduced in [22] are useful, they are not by themselves sufficient because of the complexities of policy regret and the state dynamics. This is evident from our upper bound which consists of two terms (both of which we show are necessary) and only one of them is a sequential Rademacher complexity type term.

Another line of work closes related ours is that on the theory of optimal control (see [17] for a review). Linear dynamical systems with simple zero mean noise models like Gaussian noise for state dynamics have been extensively studied (see the surveys [19] and [14] for an extensive review). While majority of the work in control have focused on linear dynamics with fixed noise models, H_∞ control (and more generally robust control) literature has aimed at extending the setting to worst case perturbations (see [26]). However these works focus on cumulative costs and are often not practical for machine learning scenarios where such algorithms tend to be overly conservative.

There has been recent work dealing with adversarial costs and linear dynamics with either stochastic or adversarial noise. Online Markov decision processes [11], Online Adversarial Tracking [1] and Online Linear Quadratic Regulator [10, 25] are all examples of such work that deal with specific form of possibly adversarially chosen cost functions, albeit the loss functions in these problems are very specific and the dynamics are basically linear with either fixed stochastic noise or no noise. Perhaps the closest comparison to our work is the work by Agarwal et. al [2] (and also [3, 12]) where adversarial but convex costs, linear policies and linear dynamics with an adversarial component are considered. In contrast, we consider arbitrary class of policies, both adversarially chosen costs (possibly non-convex) and dynamics that are presented on the fly and arbitrary state space.

2 Online learning with dynamics

We now formally define the online learning with dynamics problem. We let \mathcal{X} represent the state space, Π denote the set of learner policies, and $\mathcal{Z} = \mathcal{Z}_\ell \times \mathcal{Z}_\Phi$ denote the space of adversary's moves.

2.1 Problem setup

The problem of online learning with dynamics proceeds as a repeated game between a learner and an adversary played over T rounds. The state of the system at time t , denoted by $x_t \in \mathcal{X}$, evolves according to a stochastic dynamical system as $x_{t+1} = \Phi(x_t, \pi_t, \zeta_t) + w_t$, where $\Phi : \mathcal{X} \times \Pi \times \mathcal{Z}_\Phi \rightarrow \mathcal{X}$ is the transition function and $w_t \sim \mathcal{D}_w$ is a zero-mean additive noise. The transition function Φ is allowed to depend on adversary's action ζ_t allowing the dynamics to change across time steps. We assume that the dynamics function Φ and distribution \mathcal{D}_w are fixed apriori and are known to the learner before the game begins.

Given these dynamics, the repeated online game between the learner and the adversary starts at an initial state x_1 proceeds via the following interactive protocol:

On round $t = 1, \dots, T$,

- the learner picks policy $\pi_t \in \Pi$, adversary simultaneously selects instance $(z_t, \zeta_t) \in \mathcal{Z}$
- the learner receives payoff (loss) signal $\ell(\pi_t, x_t, z_t)$
- the state of the system transitions to $x_{t+1} = \Phi(x_t, \pi_t, \zeta_t) + w_t$

We consider the *full information* version of the above game: the learner gets to observe the instances (z_t, ζ_t) at time t . The objective of the learner is to minimize, in expectation, the *policy regret*

$$\text{Reg}_T^{\text{pol}} = \sum_{t=1}^T \mathbb{E}[\ell(\pi_t, x_t[\pi_{1:t-1}, w_{1:t-1}, \zeta_{1:t-1}], z_t)] - \inf_{\pi \in \Pi} \mathbb{E} \left[\sum_{t=1}^T \ell(\pi, x_t[\pi^{(t-1)}, w_{1:t-1}, \zeta_{1:t-1}], z_t) \right] \quad (1)$$

with respect to a policy class Π and dynamics model Φ . In the above definition, the notation $x_t[\pi_{1:t-1}, w_{1:t-1}, \zeta_{1:t-1}]$ makes the dependence of the state x_t explicit on the previous policies, noise and adversarial actions. For notational convenience, we will often drop dependencies on the noise variables w and adversarial actions ζ when it is clear from context.

Observe that in the definition of policy regret, the loss depends on the state of the system at that instance which can be potentially very different for the learner and a comparator policy π . This lends an additional source of complexity to the interactive game, and can make the problem *much* harder than its counterpart without a dynamics.

The problem of online learning with dynamics generalizes the online learning problem where the loss functions $\ell(\pi, x, z) = \tilde{\ell}(\pi, z)$ are independent of the underlying state variables. Indeed, our notion of policy regret in equation (1) reduces to the notion of external regret studied in the online learning literature. Also, the problem of online learning with memory involves adversaries which have bounded memory of length m and thus the loss incurred by the learner at any time is a function of its past m moves. By setting the state variable $x_t = [\pi_{t-m}, \dots, \pi_{t-1}]$, the dynamics function $\Phi(x_t, \pi_t) = [\pi_{t-m+1}, \dots, \pi_t]$, and the noise disturbances $w_t = 0$, we can see that the bounded memory adversaries can be seen as a special case of our problem with dynamics.

2.2 Minimax Policy Regret

Given the setup of the previous section, we study the online learning with dynamics game between the learner and the adversary through a minimax perspective. Studying this minimax value allows one to understand the limits of learnability for a tuple $(\Pi, \mathcal{Z}, \Phi, \ell)$: upper bounds on this value imply existence of algorithms with corresponding rates while lower bounds on this values represent the information-theoretic limits of learnability. In the following lemma, we formally define the value $\mathcal{V}_T(\Pi, \mathcal{Z}, \Phi, \ell)$ of the minimax policy regret for a given problem which informally is the policy regret of best learning algorithm against the worst case adversary.

Proposition 1 (Value \rightarrow Dual Game). *Let \mathcal{Q} and \mathcal{P} denote the sets of probability distributions over the policy class Π and the adversarial actions \mathcal{Z} respectively, satisfying the necessary conditions for the minimax theorem to hold. Then, we have that²*

$$\mathcal{V}_T(\Pi, \mathcal{Z}, \Phi, \ell) := \left\langle \left\langle \inf_{q_t \in \mathcal{Q}} \sup_{(z_t, \zeta_t) \in \mathcal{Z}} \mathbb{E} \right\rangle_{t=1}^T \right\rangle \left[\text{Reg}_T^{\text{pol}} \right] = \left\langle \left\langle \sup_{p_t \in \mathcal{P}} \inf_{\pi_t} \mathbb{E} \right\rangle_{t=1}^T \right\rangle \left[\text{Reg}_T^{\text{pol}} \right]. \quad (2)$$

The proof of the proposition is deferred to Appendix A. The proof proceeds via a repeated application of von Neumann's minimax theorem (for instance see [23, Appendix A]). Notice that the minimax theorem changes the order of the online sequential game defined in the setup above: at every time step t , the adversary proceeds first and outputs a distribution p_t over instances and the learner responds back with π_t *after* having observed the distribution. The actual loss instance (z_t, ζ_t) is then sampled from the revealed distribution p_t . On the other hand, the comparator remains the same as before: the best policy $\pi \in \Pi$ in hindsight. This reversed game, termed the *Dual Game*, forms the basis of our analysis and allows us to study the complexity of the online learning with dynamics problem.

3 Upper bounds on value of the game

Our main result in this section concerns an upper bound on the value of the sequential game $\mathcal{V}_T(\Pi, \mathcal{Z}, \Phi, \ell)$ relating it to the study of certain stability properties of empirical minimizers and stochastic processes associated with them. Before we proceed to describe the main result, we revisit some preliminaries and setup notation which would be helpful in describing the main result.

Sequential Rademacher complexity. The notion of Sequential Rademacher Complexity, introduced in [22], is a natural generalization of the Rademacher complexity for online learning. However, observe that the loss of the comparator term in the definition of policy regret in equation (1) depends on the adversarial actions $\zeta_{1:t-1}$ through the dynamics and z_t through the loss function ℓ . We define the following version of sequential Rademacher complexity for such dynamics based losses.

² $\langle \dots \rangle_{t=1}^T$ denotes interleaved application of the sequence of operator inside. For example, for $T = 2$, $\langle \langle \sup_{p_t} \inf_{q_t} \rangle_{t=1}^2 \rangle [\cdot] = \sup_{p_1} \inf_{q_1} \sup_{p_2} \inf_{q_2} [\cdot]$

Definition 1. *The Sequential Rademacher Complexity of a policy class Π with respect to loss function $\ell : \Pi \times \mathcal{X} \times \mathcal{Z}_\ell \mapsto \mathbb{R}$ and dynamics $\Phi : \mathcal{X} \times \Pi \times \mathcal{Z}_\Phi \rightarrow \mathcal{X}$ is defined as*

$$\mathfrak{R}_T^{\text{seq}}(\ell \circ \Pi) := \sup_{(\mathbf{z}, \zeta)} \mathbb{E}_\epsilon \left[\sup_{\pi \in \Pi} \sum_{t=1}^T \epsilon_t \ell(\pi, x_t[\zeta_1(\epsilon), \dots, \zeta_{t-1}(\epsilon)], \mathbf{z}_t(\epsilon)) \right],$$

where the outer supremum is taken over $\mathcal{Z} = \mathcal{Z}_\ell \times \mathcal{Z}_\Phi$ -valued trees³ of depth T and $\epsilon = (\epsilon_1, \dots, \epsilon_T)$ is a sequence of i.i.d. Rademacher random variables.

A similar definition was also used by Han et al. [13] in the context of online learning with strategies where the notion of regret was defined w.r.t. a set of strategies rather than a fixed action. As compared with the classical online learning problem, the above comprises problems where the loss at time t depends on the complete history $(\zeta_1, \dots, \zeta_{t-1})$ of the adversarial choices along with z_t . As noted by [13], such dependence on the the adversary's history can often make the online learning problem harder to learn compared with the online learning problem.

Empirical Risk Minimization (ERM). Given a sequence of loss functions $\ell_t : \mathcal{F} \mapsto \mathbb{R}$ for $t \in [T]$, the ERM with respect to a function class \mathcal{F} is defined to be the minimizer of the cumulative loss with $f_{\text{ERM}, T} \in \operatorname{argmin}_{f \in \mathcal{F}} \sum_{t=1}^T \ell_t(f)$. In the statistical learning setup, the problems of supervised classification and regression are known to be learnable with respect to a function class \mathcal{F} if and only if the empirical risks uniformly converge over this class \mathcal{F} to the population risks. In contrast, our results provide *sufficient* conditions for learnability in terms of certain stability properties of such empirical risk minimizers.

Dynamic stability. We introduce the notion of dynamic stability which captures the stability of an algorithm's interaction with the underlying dynamics Φ . In order to do so, we define a notion of counterfactual loss ℓ_t^Φ of a policy π as the loss incurred by a learner which selects π for time $1 : t$.

Definition 2 (Counterfactual Losses). *Given a sequence of adversarial actions $\zeta_{1:t-1}, z_t$, dynamics function Φ , and noise distribution \mathcal{D}_w , the counterfactual loss of a policy π at time t is*

$$\ell_t^\Phi(\pi, \zeta_{1:t-1}, z_t) := \mathbb{E}_{w_s \sim \mathcal{D}_w} \left[\ell(\pi, x_t[\pi^{(t-1)}, w_{1:t-1}, \zeta_{1:t-1}], z_t) \right].$$

With this definition, observe that the comparator term in the value \mathcal{V}_T in equation (2) is in fact a cumulative sum of counterfactual losses for a policy π . Any algorithm \mathcal{A} that plays a sequence of policies $\{\pi_t\}$ in the online game incurs an instantaneous loss $\ell(\pi_t, x_t[\pi_{1:t-1}, \zeta_{1:t-1}], z_t)$ at time t . In comparison, the counterfactual loss $\ell^\Phi(\pi_t, \zeta_{1:t-1}, z_t)$ represents a scenario where the algorithm commits to the policy π_t from the beginning of the game. Our notion of dynamic stability of an algorithm is precisely the deviation between these two types of losses: instantaneous and counterfactual.

Definition 3 (Dynamic Stability). *An algorithm \mathcal{A} is said to be $\{\beta_t\}$ -dynamically stable if for all sequences of adversarial actions $[(z_1, \zeta_1), \dots, (z_T, \zeta_T)]$ and time instances $t \in [T]$*

$$\left| \mathbb{E}_{w_{1:t-1}} [\ell(\pi_t, x_t[\pi_{1:t-1}, w_{1:t-1}, \zeta_{1:t-1}], z_t)] - \ell^\Phi(\pi_t, \zeta_{1:t-1}, z_t) \right| \leq \beta_t \quad \text{where } \pi_t = \mathcal{A}((z_{1:t-1}, \zeta_{1:t-1})).$$

It is interesting to note that if that loss functions are independent of the underlying states, that is $\ell(\pi, x, z) = \tilde{\ell}(\pi, z)$, then any algorithm is dynamically stable in a trivial manner with the stability parameters $\beta_t = 0$ for all time instances t .

With these definitions, we now proceed to describe our main result. Recall that Proposition 1 translates the problem of studying the value of the game $\mathcal{V}_T(\Pi, \mathcal{Z}, \Phi, \ell)$ to that of studying the policy regret in a dual game. In this dual game, the learner has access to the set of adversaries distribution $\{p_s\}_{s=1}^t$ at time t and the policy π_t can be a function of these. For a regularization function $\Omega : \Pi \mapsto \mathbb{R}_+$, we denote the regularized ERMs with respect to function class Π and counterfactual losses ℓ^Φ by

$$\pi_{\text{RERM}, t} \in \operatorname{argmin}_{\pi \in \Pi} \sum_{s=1}^t \mathbb{E}_{z_s} [\ell^\Phi(\pi, \zeta_{1:s-1}, z_s)] + \lambda \cdot \Omega(\pi), \quad (3)$$

where $\lambda \geq 0$ is the regularization parameter. The following theorem provides an upper bound on the value \mathcal{V}_T in terms of the dynamic stability parameters of the regularized ERMs above as well a sequential Rademacher complexity of the *effective* loss class $\ell^\Phi \circ \Pi := \{\ell^\Phi(\pi, \cdot) : \pi \in \Pi\}$.

³A \mathcal{Z} -valued tree \mathbf{z} of depth d is defined as a sequence $(\mathbf{z}_1, \dots, \mathbf{z}_d)$ of mappings $\mathbf{z}_t : \{\pm 1\}^{t-1} \mapsto \mathcal{Z}$ (see [21])

Theorem 1 (Upper bound on value). *For any online learning with dynamics instance $(\Pi, \mathcal{Z}, \Phi, \ell)$, consider the set of regularized ERMs given by eq. (3) with regularization function Ω and parameter $\lambda \geq 0$ having dynamic stability parameters $\{\beta_{\text{RERM},t}\}_{t=1}^T$. Then, we have that the value of the game*

$$\mathcal{V}_T(\Pi, \mathcal{Z}, \Phi, \ell) \leq \sum_{t=1}^T \beta_{\text{RERM},t} + 2\mathfrak{R}_{\mathcal{T}}^{\text{seq}}(\ell^\Phi \circ \Pi) + 2\lambda \cdot \sup_{\pi \in \Pi} \Omega(\pi). \quad (4)$$

The complete proof of the above theorem can be found in Appendix B. A few comments on Theorem 1 are in order. The theorem provides sufficient conditions to ensure learnability of the online learning with dynamics problem. In particular, the two terms Term (I) = $\sum_{t=1}^T \beta_{\text{RERM},t}$ and Term (II) = $\mathfrak{R}_{\mathcal{T}}^{\text{seq}}(\ell^\Phi \circ \Pi)$ contain the main essence of the upper bound. Term (I) concerns the dynamic mixability property of the regularized ERM in the dual game. If there exist *approximate* minimizers (regularized) of the sequence of counterfactual losses within the policy class Π such that $\pi_{\text{RERM},t}$ is uniformly close to $\pi_{\text{RERM},t+1}$ the dynamic stability parameters can be made to be small. Term (II) comprises of the sequential Rademacher complexity of the loss class $\ell^\Phi \circ \Pi$ which involves the underlying policy class Π as well as the counterfactual loss ℓ^Φ . This measure of complexity can be seen as one which corresponds to an effective online game where the the loss at time t depends on the adversarial actions up to time t . Compare this to the instantaneous loss $\ell(\pi_t, x_t[\pi_{1:t-1}, \zeta_{1:t-1}], z_t)$ which depends on both the policies as well as the adversarial actions up to time t . Observe that for the classical online learning setup without dynamics, the dynamic stability parameters $\beta_{\text{RERM},t} \equiv 0$. On setting the value of regularization parameter $\lambda = 0$, we recover back the learnability result of Rakhlin et al. [22].

We would like to highlight that the complexity-based learnability guarantees of Theorem 1 are non-constructive in nature. In particular, the theorem says that any non-trivial upper bounds on the stability and sequential complexity terms would guarantee the *existence* of an online learning algorithm with the corresponding policy regret. Our minimax perspective on the problem allows us to study the problem in full generality without making assumptions with respect to the policy class Π , adversarial actions \mathcal{Z} and the underlying (possibly adversarial) dynamics Φ , and provide sufficient conditions for learnability.

Given the upper bound on the value $\mathcal{V}_T(\Pi, \mathcal{Z}, \Phi, \ell)$, one can observe that there is a possible tension between the two complexity terms: while dynamic stability term promotes using policies which are ‘‘similar’’ across time steps, the regularized complexity term seeks policies which are minimizers of cumulative losses and might vary across time steps. In order to balance similar trade-offs, a natural *Mini-Batching Algorithm* has been proposed in various works on online learning with memory [5] and online learning with switching costs [9]. The key idea is that the learner divides the time T into intervals of length $\tau > 0$ and commits to playing the same strategy over this time period.

Let us denote any such mini-batching algorithm by \mathcal{A}_τ and the corresponding minimax value restricted to this class of algorithms by $\mathcal{V}_{T,\tau}(\Pi, \mathcal{Z}, \Phi, \ell)$ where the infimum in equation 2 is taken over all mini-batching algorithms \mathcal{A}_τ . Similar to the regularized ERM of equation (3), we define the following mini-batched ERMs:

$$\pi_{\text{ERM},t}^\tau = \begin{cases} \pi_{\text{ERM}}(t) & \text{for } t \equiv 0 \pmod{\tau} \\ \pi_{\text{ERM}}(\tau \lfloor \frac{t}{\tau} \rfloor) & \text{otherwise} \end{cases}, \quad (5)$$

where we have used the notation $\pi_{\text{ERM}}(t) := \pi_{\text{ERM},t}$. In the following proposition, we prove an upper bound analogous to that of Theorem 1 for this class of mini-batching algorithms⁴.

Proposition 2 (Mini-batching algorithms.). *For any online learning with dynamics game $(\Pi, \mathcal{Z}, \Phi, \ell)$, consider the set of mini-batch ERMs given by equation (5) having dynamic stability parameters $\{\beta_{\text{ERM},t}^\tau\}_{t=1}^T$. Then, we have that the value of the game*

$$\mathcal{V}_T(\Pi, \mathcal{Z}, \Phi, \ell) \leq \inf_{\tau > 0} \mathcal{V}_{T,\tau}(\Pi, \mathcal{Z}, \Phi, \ell) \leq \inf_{\tau > 0} \left(\sum_{t=1}^T \beta_{\text{ERM},t}^\tau + 2\tau \cdot \sup_{s \in [\tau]} \mathfrak{R}_{\mathcal{T}/\tau}^{\text{seq}}(\ell_s^\Phi \circ \Pi) \right), \quad (6)$$

where ℓ_s^Φ is the counterfactual loss for the s^{th} batch.

We defer the proof of the above proposition to Appendix B. In comparison with the upper bound of Theorem 1, this bound concerns the dynamic stability of the mini-batched ERMS as compared to

⁴For this class of mini-batching algorithms, we consider an oblivious adversary which cannot adapt to the randomness of the learner.

their regularized counterparts. Often times, obtaining bounds on the stability parameters $\{\beta_{\text{ERM},t}^\tau\}_{t=1}^T$ can be much easier than the ones for regularized ERMS. For instance, it is easy to see that for the problem of online learning with memory with adversaries having memory m , one can bound $\sum_{t=1}^T \beta_{\text{ERM},t}^\tau = O(\frac{mT}{\tau})$ whenever the losses are bounded, providing a natural trade-off between the two complexity terms.

4 Lower bounds on value of the game

Having established sufficient conditions for the learnability of the online learning with dynamics problem in the previous section, we now turn to address the optimality of these conditions. Recall that Theorem 1 and Proposition 2 established upper bounds on the value $\mathcal{V}_T(\Pi, \mathcal{Z}, \Phi, \ell)$ for instances of our problem. The following theorem shows that both the upper bounds of equations (4) and (6) are indeed tight upto constant factors.

Theorem 2 (Lower Bound). *For the online learning with dynamics problem, there exist problem instances $\{(\Pi, \mathcal{Z}, \Phi, \ell_i)\}_{i=1}^3$, a regularization function Ω and a universal constant $c > 0$ such that*

$$\mathcal{V}_T(\Pi, \mathcal{Z}, \Phi, \ell_1) \geq c \cdot \mathfrak{R}_T^{\text{seq}}(\ell_1^\Phi \circ \Pi) \quad (7a)$$

$$\mathcal{V}_T(\Pi, \mathcal{Z}, \Phi, \ell_2) \geq c \cdot \inf_{\lambda > 0} \left(\sum_{t=1}^T \beta_{\text{RERM},t} + \lambda \cdot \sup_{\pi \in \Pi} \Omega(\pi) \right) \quad (7b)$$

$$\mathcal{V}_T(\Pi, \mathcal{Z}, \Phi, \ell_3) \geq c \cdot \inf_{\tau > 0} \left(\sum_{t=1}^T \beta_{\text{ERM},t}^\tau + 2\tau \mathfrak{R}_{T/\tau}^{\text{seq}}(\ell_3^\Phi \circ \Pi) \right), \quad (7c)$$

where $\beta_{\text{RERM},t}$ and $\beta_{\text{ERM},t}^\tau$ are the dynamic mixability parameters of the regularized ERM w.r.t. ℓ_2 (eq. (3)) and mini-batching ERM w.r.t. ℓ_3 (eq. (5)) respectively.

A few comments on Theorem 2 are in order. The theorem exhibits that the sufficiency conditions from Theorem 1 and Proposition 2 are indeed necessary by exhibiting instances whose value is lower bounded by these terms. In particular, equation (7a) shows that the sequential Rademacher term is necessary, (7b) establishes necessity for the dynamic stability of the regularized ERM, while (7c) shows that the mini-batching upper bound is also tight. It is worth noting that these lower bounds are not instance dependent but rather construct specific examples to demonstrate the tightness of our upper bound from the previous section. We next present the key idea for the proof of the theorem and defer the complete details to Appendix C.

Proof sketch. We now describe the example instances which form the crux of the proof for Theorem 2. Consider the online learning with dynamics game between a learner and an adversary with the state space $\mathcal{X} = \{x \in \mathbb{R}^d \mid \|x\|_2 \leq 1\}$ and the set of adversarial actions $\mathcal{Z}_\ell^{\text{Lin}} = \{z \in \mathbb{R}^d \mid \|z\|_2 \leq 1\}$. Further, we consider the constant policy class $\Pi_{\text{Lin}} = \{\pi_f \mid \pi(x) = f \text{ for all states } x \text{ with } f \in \mathbb{B}_d(1)\}$, consisting of policies π_f which select the same action f at each state x . With a slight abuse of notation, we represent the policy π_t played by the learner at time by the corresponding d -dimensional vector f_t . Further, we let the dynamics function $\Phi_{\text{Lin}}(x_t, f_t, \zeta_t) = f_t$. We now define the loss function which consists of two parts, a linear loss and a L -Lipschitz loss involving the dynamics:

$$\ell_L(f_t, x_t, z_t) = \langle f_t, z_t \rangle + \sigma(f_t, x_t) \quad \text{where} \quad \sigma(f_t, x_t) = \begin{cases} L \|f_t - x_t\|_2 & \text{for } \|f_t - x_t\|_2 \leq \frac{1}{L} \\ 1 & \text{otherwise} \end{cases}. \quad (8)$$

Observe that this example constructs a family of instances one for each value of the Lipschitz constant of L of the function σ . For this family of instances, we establish that the value

$$\mathcal{V}_T(\Pi_{\text{Lin}}, \mathcal{Z}, \Phi_{\text{Lin}}, \ell_L) \geq \begin{cases} \sqrt{T} & \text{for } 0 < L < 1 \\ \sqrt{LT} & \text{for } 1 \leq L \leq (4T)^{\frac{1}{3}} \\ 2^{\frac{1}{3}} T^{\frac{2}{3}} & \text{for } L > (4T)^{\frac{1}{3}} \end{cases}.$$

The proof finally connects these lower bounds to the upper bounds of Theorem 1 and Proposition 2.♣

With the lower bounds given in Theorem 2, it is natural to ask whether the sufficient conditions in Theorem 1 and Proposition 2 are indeed necessary for every instance of the online learning with dynamics problem. The answer to this question is unsurprisingly *No* given the generality in which we study this problem. Consider the following simple instance of the problem:

$$\Pi = \mathcal{X}, \quad \ell(\pi, x, z) = \tilde{\ell}(\pi, z) + \mathbb{I}[\pi = x], \quad \text{and} \quad x_{t+1} = \pi_t,$$

for any non-negative bounded loss $\tilde{\ell}(\pi, z) \in [0, 1]$ for all $\pi \in \Pi, z \in \mathcal{Z}_\ell$. Consider any policy class for which $\mathfrak{R}_T^{\text{seq}}(\ell^\Phi \circ \Pi) > 0$. Both bounds (4) and (6) suggest that the problem is learnable with rate at least $\mathfrak{R}_T^{\text{seq}}(\ell^\Phi \circ \Pi)$. However, observe that the indicator term in the loss is quite severe on the comparator; it ensures that the comparator term is at least T . Thus, *any* algorithm which selects a policy from Π at every instance can ensure that the policy regret is at most 0! While the above example establishes that the sufficient conditions are not necessary in an instance dependent manner, our next proposition establishes that they are indeed tight for large class of problems instances.

Proposition 3 (Instance-dependent lower bound). *a) Given any online learning problem $(\mathcal{F}, \mathcal{Z}_\ell, \ell)$ with a bounded loss function $\ell : \mathcal{F} \times \mathcal{Z}_\ell \mapsto [-1, 1]$, there exists an online learning with dynamics problem $(\Pi_{\mathcal{F}}, \mathcal{Z}_\ell \times \{-1, 1\}, \Phi, \tilde{\ell})$ and a universal constant $c > 0$ such that*

$$\mathcal{V}_T(\Pi_{\mathcal{F}}, \mathcal{Z}_\ell \times \{-1, 1\}, \Phi, \tilde{\ell}) \geq c \cdot \inf_{\tau > 0} \left(\sum_{t=1}^T \beta_{\text{ERM}, t}^\tau + 2\tau \mathfrak{R}_{T/\tau}^{\text{seq}}(\ell^\Phi \circ \Pi) \right),$$

where $\beta_{\text{ERM}, t}^\tau$ are the dynamic mixability parameters of the mini-batching ERM w.r.t. ℓ (eq. (5)).

b) Given a policy class Π and dynamics function Φ , there exists an online learning with dynamics problem $(\Pi, \mathcal{Z}, \Phi, \ell)$ and a universal constant $c > 0$ such that

$$\mathcal{V}_T(\Pi, \mathcal{Z}, \Phi, \ell) \geq c \cdot \mathfrak{R}_T^{\text{seq}}(\ell^\Phi \circ \Pi).$$

We defer the proof of the proposition to Appendix C. This proposition can be seen as a strengthening of the lower bounds (7a) and (7c) showing that for a very large class of problems, the upper bound given by the mini-batching algorithm and the sequential complexity terms are in fact necessary.

5 Examples

In this section, we look at specific examples of the online learning with dynamics problem and obtain learnability guarantees for these instances using our upper bounds from Theorem 1. For clarity of exposition, our focus in this section on the scaling of the value $\mathcal{V}_T(\Pi, \mathcal{Z}, \Phi, \ell)$ with the time horizon T . The proofs in Appendix D explicitly detail out all the problem dependent parameters.

Example 1: Online Isotron with dynamics. Single Index Models (SIM) are class of semi-parametric models widely studied in the econometric and operations research community. Kalai and Sastry [16] introduced the Isotron algorithm for learning SIMs and Rakhlin et al [23] established that the online version of this problem is learnable. Here, we introduce a version of this problem with a state variable that requires a component of the model to vary slowly across time:

$$\mathcal{X} = \mathbb{R}, \mathcal{F} = \{f = (\sigma, \mathbf{w} = (w_1, w)) \mid \sigma : [-1, 1] \mapsto [-1, 1] \text{ 1-Lipschitz}, \mathbf{w} \in \mathbb{R}^{d+1} \mid |w_1| \leq 1, \|w\|_2 \leq 1\},$$

$$\Pi_{\mathcal{F}} = \{\pi_f \mid \pi \in \mathcal{F}, \pi_f(x) = f \text{ for all } x \in \mathcal{X}\}, \mathcal{Z}_\ell = [-1, 1]^{d+1} \times [-1, 1], \mathcal{Z}_\Phi = \phi,$$

$$\ell(\pi_f, x, z = (z_1, \mathbf{x}, y)) = (y - \sigma(\langle \mathbf{x}, w \rangle))^2 + (z_1 - w_1)^2 + (x - w_1)^2, \Phi(x, \pi_{(\sigma, w)}, \zeta) = w_1, \mathcal{D}_w = 0. \quad (9)$$

The following corollary establishes the learnability of the online Isotron problem with dynamics.

Corollary 1. *For the online Isotron problem with dynamics given by $(\Pi_{\mathcal{F}}, \mathcal{Z}, \Phi, \ell)$ in equation (5), we have that the minimax value*

$$\mathcal{V}_{\text{Iso}, T}(\Pi_{\mathcal{F}}, \mathcal{Z}, \Phi, \ell) = \tilde{\mathcal{O}}(\sqrt{T}).$$

It is important to note that as with the online learning problem [23], it is not clear whether a computationally efficient method attaining the above guarantee exists. ♣

Example 2: Online Markov decision processes [11]. This example considers the problem of Online Markov Decision Processes (MDPs) studied in Even-Dar et al. [11]. The setup consists of a finite state space $|\mathcal{X}| = S$, a finite action space $|\mathcal{U}| = A$, and

$$\mathcal{Z}_\ell = \{z \mid z \in [0, 1]^{S \times A}\}, \mathcal{Z}_\Phi = \phi, \Pi_{\text{MDP}} = \{\pi \mid \pi : \mathcal{X} \mapsto \Delta(\mathcal{U})\}, \ell(\pi, x, z) = z(x, \pi(x)),$$

$$\Phi \text{ given by } P : \mathcal{X} \times \mathcal{U} \mapsto \Delta(\mathcal{U}) \text{ with } x' \sim P(x, \pi(x)). \quad (10)$$

With this setup, we now provide a bound on the minimax value $\mathcal{V}_T^{\text{MDP}}$, assuming, as in [11], that the underlying MDP is unichain and satisfies a mixability assumption (see Appendix D for details).

Corollary 2. *For the online MDP problem given by $(\Pi_{\text{MDP}}, \mathcal{Z}, \Phi, \ell)$ in equation (5), we have that the minimax value*

$$\mathcal{V}_{\text{MDP}, T}(\Pi_{\text{MDP}}, \mathcal{Z}, \Phi, \ell) = \mathcal{O}(\sqrt{T}).$$

The above corollary helps one recover the same $\mathcal{O}(\sqrt{T})$ regret bound that was obtained by [11]. ♣

Example 3: Online Linear Quadratic Regulator [10]. The online Linear Quadratic Regulator (LQR) setup studied in this section was first studied in Cohen et al. [10]. The setup consists of a LQ system - with linear dynamics and quadratic costs - where the cost functions are chosen adversarially. The comparator class consists of a set of strongly stable linear policies (see Appendix D).

$$\mathcal{X} = \mathbb{R}^d, \mathcal{Z}_\ell = \{(Q, R) \mid Q, R > 0, \text{tr}(Q), \text{tr}(R) \leq C\}, \Pi_{\text{LQR}} = \{K \in \mathbb{R}^{k \times d} \mid K \text{ is } (\kappa, \gamma) \text{-strongly stable}\}, \\ \mathcal{Z}_\Phi = \phi, \ell(\pi, x, z) = x^\top Qx + (Kx)^\top R(Kx), \Phi(x, K) = (A + BK)x, \mathcal{D}_w = \mathcal{N}(0, I). \quad (11)$$

With this setup, we now establish the learnability of this problem in the following corollary.

Corollary 3. *For the online LQR problem given by $(\Pi_{\text{LQR}}, \mathcal{Z}, \Phi, \ell)$ in equation (5), we have that the minimax value*

$$\mathcal{V}_{\text{LQR}, T}(\Pi_{\text{LQR}}, \mathcal{Z}, \Phi, \ell) = \tilde{\mathcal{O}}(\sqrt{T}).$$

Note that [10] obtained a similar policy regret bound of $\mathcal{O}(\sqrt{T})$ but their analysis only worked for an oblivious adversary whereas the guarantee of Corollary 3 holds for an adaptive adversary. ♣

Example 4: Online Control with Adversarial Disturbances [2]. In this example, we consider a simplified version of the setup from Agarwal et al. [2] where the adversary is allowed to perturb the dynamics at each time instance along with the loss functions. We consider the LQ version of the problem where the dynamics are linear and the cost functions quadratic. Similar to the example above, the comparator class consists of a set of strongly stable linear policies (see Appendix D).

$$\mathcal{X} = \mathbb{R}^d, \mathcal{Z}_\ell = \{(Q, R) \mid Q, R > 0, \text{tr}(Q), \text{tr}(R) \leq C\}, \Pi_{\text{LQR}} = \{K \in \mathbb{R}^{k \times d} \mid K \text{ is } (\kappa, \gamma) \text{-strongly stable}\}, \\ \mathcal{Z}_\Phi = \{\zeta \mid \|\zeta\|_2 \leq W\}, \ell(\pi, x, z) = x^\top Qx + (Kx)^\top R(Kx), \Phi(x, K, \zeta) = (A + BK)x + \zeta, \mathcal{D}_w = \phi. \quad (12)$$

With this setup, we now establish the learnability of this problem in the following corollary.

Corollary 4. *For the Online Control with Adversarial Disturbances problem instance $(\Pi_{\text{LQR}}, \mathcal{Z}, \Phi, \ell)$ detailed in equation (5), we have that the minimax value*

$$\mathcal{V}_{\text{adv}, T}(\Pi_{\text{LQR}}, \mathcal{Z}, \Phi, \ell) = \tilde{\mathcal{O}}(\sqrt{T}).$$

In contrast to Example 3 above, the disturbances in the dynamics are actually controlled by the adversary instead of being random. Our result above shows that this harder version of the problem is also learnable at \sqrt{T} rate. ♣

In addition to these four examples, in Appendix D we consider additional examples of the framework including online adversarial tracking [1] and a non-linear generalization of the LQ problem [10, 2] which we call online non-linear control.

Broader Impact

This work is mainly theoretical in nature and hopes to provide theoretical foundations for learning under dynamical systems. The work is expected to have a broader impact in the future by opening up research on learning and non-linear dynamical systems with complex policy classes. In the future, we hope that our work will enable ML systems to be deployed reliably in more reactive environments.

Acknowledgments and Disclosure of Funding

We would like to thank Dylan Foster, Mehryar Mohri and Ayush Sekhari for helpful discussions. KB is supported by a JP Morgan AI Fellowship. KS would like to acknowledge NSF CAREER Award 1750575 and Sloan Research Fellowship

References

- [1] Y. Abbasi-Yadkori, P. Bartlett, and V. Kanade. Tracking adversarial targets. In *International Conference on Machine Learning*, pages 369–377, 2014.
- [2] N. Agarwal, B. Bullins, E. Hazan, S. M. Kakade, and K. Singh. Online control with adversarial disturbances. *arXiv preprint arXiv:1902.08721*, 2019.
- [3] N. Agarwal, E. Hazan, and K. Singh. Logarithmic regret for online control. In *Advances in Neural Information Processing Systems*, pages 10175–10184, 2019.
- [4] O. Anava, E. Hazan, and S. Mannor. Online learning for adversaries with memory: price of past mistakes. In *Advances in Neural Information Processing Systems*, pages 784–792, 2015.
- [5] R. Arora, O. Dekel, and A. Tewari. Online bandit learning against an adaptive adversary: From regret to policy regret. In *Proceedings of the 29th International Conference on Machine Learning*, 2012.
- [6] P. L. Bartlett and S. Mendelson. Rademacher and gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3, 2002.
- [7] S. Ben-David, D. Pál, and S. Shalev-Shwartz. Agnostic online learning.
- [8] N. Cesa-Bianchi and G. Lugosi. *Prediction, learning, and games*. Cambridge university press, 2006.
- [9] L. Chen, Q. Yu, H. Lawrence, and A. Karbasi. Minimax regret of switching-constrained online convex optimization: No phase transition. *arXiv preprint arXiv:1910.10873*, 2019.
- [10] A. Cohen, A. Hassidim, T. Koren, N. Lazic, Y. Mansour, and K. Talwar. Online linear quadratic control. *arXiv preprint arXiv:1806.07104*, 2018.
- [11] E. Even-Dar, S. M. Kakade, and Y. Mansour. Online markov decision processes. *Mathematics of Operations Research*, 34(3):726–736, 2009.
- [12] D. J. Foster and M. Simchowitz. Logarithmic regret for adversarial online control. *arXiv preprint arXiv:2003.00189*, 2020.
- [13] W. Han, A. Rakhlin, and K. Sridharan. Competing with strategies. In *Conference on Learning Theory*, pages 966–992, 2013.
- [14] M. Hardt, T. Ma, and B. Recht. Gradient descent learns linear dynamical systems. *The Journal of Machine Learning Research*, 19(1), 2018.
- [15] E. Hazan et al. Introduction to online convex optimization. *Foundations and Trends® in Optimization*, 2(3-4):157–325, 2016.
- [16] A. T. Kalai and R. Sastry. The isotron algorithm: High-dimensional isotonic regression. In *COLT*. Citeseer, 2009.
- [17] D. E. Kirk. *Optimal control theory: an introduction*. Courier Corporation, 2004.
- [18] N. Littlestone. Learning quickly when irrelevant attributes abound: A new linear-threshold algorithm. *Machine learning*, 2(4), 1988.
- [19] L. Ljung. System identification. *Wiley encyclopedia of electrical and electronics engineering*, 1999.
- [20] N. Merhav, E. Ordentlich, G. Seroussi, and M. J. Weinberger. On sequential strategies for loss functions with memory. *IEEE Transactions on Information Theory*, 48(7):1947–1958, 2002.
- [21] A. Rakhlin and K. Sridharan. *Statistical Learning and Sequential Prediction*. Lecture Notes, 2014.
- [22] A. Rakhlin, K. Sridharan, and A. Tewari. Online learning: Random averages, combinatorial parameters, and learnability. In *Advances in Neural Information Processing Systems*, 2010.
- [23] A. Rakhlin, K. Sridharan, and A. Tewari. Online learning via sequential complexities. *Journal of Machine Learning Research*, 16(2):155–186, 2015.
- [24] S. Shalev-Shwartz et al. Online learning and online convex optimization. *Foundations and Trends® in Machine Learning*, 4(2):107–194, 2012.
- [25] M. Simchowitz and D. J. Foster. Naive exploration is optimal for online lqr. *arXiv preprint arXiv:2001.09576*, 2020.

- [26] R. F. Stengel. *Optimal control and estimation*. Courier Corporation, 1994.
- [27] A. S. Suggala and P. Netrapalli. Online non-convex learning: Following the perturbed leader is optimal. *arXiv preprint arXiv:1903.08110*, 2019.
- [28] V. Vapnik and A. Y. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability & Its Applications*, 16, 1971.

A Proof of Proposition 1

The minimax value of the policy regret for the online learning with dynamics protocol is achieved when at every time t , the learner picks the best distribution q_t , the adversary picks the worst-case z_t and a sample of policy π_t is then drawn from q_t . This can be succinctly represented as a sequence of infimum, supremum and expectations as

$$\begin{aligned} \mathcal{V}_T(\Pi, \mathcal{Z}, \Phi, \ell) &= \left\langle \left\langle \inf_{q_t \in \mathcal{Q}} \sup_{(z_t, \zeta_t) \in \mathcal{Z}} \mathbb{E} \right\rangle_{t=1}^{T-1} \right\rangle \inf_{q_T \in \mathcal{Q}} \sup_{(z_T, \zeta_T) \in \mathcal{Z}} \mathbb{E} \left[\text{Reg}_T^{\text{pol}} \right] \\ &\stackrel{(i)}{=} \left\langle \left\langle \inf_{q_t \in \mathcal{Q}} \sup_{(z_t, \zeta_t) \in \mathcal{Z}} \mathbb{E} \right\rangle_{t=1}^{T-1} \right\rangle \sup_{p_T \in \mathcal{P}} \inf_{\pi_T} \mathbb{E}_{(z_T, \zeta_T) \sim p_T} \left[\text{Reg}_T^{\text{pol}} \right] \\ &\stackrel{(ii)}{=} \left\langle \left\langle \sup_{p_t \in \mathcal{P}} \inf_{\pi_t} \mathbb{E}_{(z_t, \zeta_t) \sim p_t} \right\rangle_{t=1}^T \right\rangle \left[\text{Reg}_T^{\text{pol}} \right], \end{aligned}$$

where (i) follows from an application of the von Neumann's minimax theorem for the distributions \mathcal{Q} and \mathcal{P} (see [23, Appendix A]) and (ii) follows from repeatedly performing the same step for $t = \{1, \dots, T-1\}$. This establishes the desired claim. \square

B Proofs of upper bounds

B.1 Proof of Theorem 1

Recall from equation (3) that the dual regularized ERM for a regularization function Ω and parameter $\lambda \geq 0$ is

$$\pi_{\text{RERM}, t} \in \underset{\pi \in \Pi}{\text{argmin}} \sum_{s=1}^t \mathbb{E} \left[\ell^\Phi(\pi, \zeta_{1:s-1}, z_s) \right] + \lambda \cdot \Omega(\pi).$$

The proof of our main result relies on the following intermediate result which relates the performance of the above RERM with that of any policy $\pi \in \Pi$ when compared on the counterfactual losses ℓ^Φ .

Lemma 1. *For any policy $\pi \in \Pi$ and any sequence of distributions $\{p_t\}_{t=1}^T$ over instance space \mathcal{Z} , we have*

$$\sum_{t=1}^T \mathbb{E}_{z_t} \left[\ell^\Phi(\pi_{\text{RERM}, t}, \zeta_{1:t-1}, z_t) \right] \leq \sum_{t=1}^T \mathbb{E}_{z_t} \left[\ell^\Phi(\pi, \zeta_{1:t-1}, z_t) \right] + \lambda \cdot (\Omega(\pi) - \Omega(\pi_{\text{RERM}, 1})). \quad (13)$$

Taking this lemma as given, let us proceed to the proof of the theorem statement. For the purpose of this proof, we will use the notation $\hat{\pi}_t := \pi_{\text{RERM}, t}$. Let us begin by considering the value of the game and its equivalence to the dual game established by Proposition 1 as⁵

$$\begin{aligned} \mathcal{V}_T(\Pi, \mathcal{Z}, \Phi, \ell) &= \left\langle \left\langle \sup_{p_t \in \mathcal{P}} \inf_{\pi_t} \mathbb{E}_{(z_t, \zeta_t) \sim p_t} \right\rangle_{t=1}^T \right\rangle \left(\sum_{t=1}^T \mathbb{E} \left[\ell(\pi_t, x_t[\pi_{1:t-1}, \zeta_{1:t-1}], z_t) \right] - \inf_{\pi \in \Pi} \mathbb{E} \left[\sum_{t=1}^T \ell(\pi, x_t[\pi^{(t-1)}, \zeta_{1:t-1}], z_t) \right] \right) \\ &\stackrel{(i)}{\leq} \left\langle \left\langle \sup_{p_t \in \mathcal{P}} \mathbb{E}_{(z_t, \zeta_t) \sim p_t} \right\rangle_{t=1}^T \right\rangle \left(\sum_{t=1}^T \mathbb{E} \left[\ell(\hat{\pi}_t, x_t[\hat{\pi}_{1:t-1}, \zeta_{1:t-1}], z_t) \right] - \mathbb{E}_{\zeta_{1:t-1}, z_t} \left[\ell^\Phi(\hat{\pi}_t, \zeta_{1:t-1}, z_t) \right] \right. \\ &\quad \left. \sum_{t=1}^T \mathbb{E}_{\zeta_{1:t-1}, z_t} \left[\ell^\Phi(\hat{\pi}_t, \zeta_{1:t-1}, z_t) \right] - \inf_{\pi \in \Pi} \mathbb{E} \left[\sum_{t=1}^T \ell(\pi, x_t[\pi^{(t-1)}, \zeta_{1:t-1}], z_t) \right] \right) \\ &\stackrel{(ii)}{\leq} \left\langle \left\langle \sup_{p_t \in \mathcal{P}} \mathbb{E}_{(z_t, \zeta_t) \sim p_t} \right\rangle_{t=1}^T \right\rangle \left(\sum_{t=1}^T \mathbb{E} \left[\ell(\hat{\pi}_t, x_t[\hat{\pi}_{1:t-1}, \zeta_{1:t-1}], z_t) \right] - \mathbb{E}_{\zeta_{1:t-1}, z_t} \left[\ell^\Phi(\hat{\pi}_t, \zeta_{1:t-1}, z_t) \right] \right) \quad [\text{Term (I)}] \\ &\quad + \left\langle \left\langle \sup_{p_t \in \mathcal{P}} \mathbb{E}_{(z_t, \zeta_t) \sim p_t} \right\rangle_{t=1}^T \right\rangle \left(\sum_{t=1}^T \mathbb{E}_{\zeta_{1:t-1}, z_t} \left[\ell^\Phi(\hat{\pi}_t, \zeta_{1:t-1}, z_t) \right] - \inf_{\pi \in \Pi} \mathbb{E} \left[\sum_{t=1}^T \ell(\pi, x_t[\pi^{(t-1)}, \zeta_{1:t-1}], z_t) \right] \right) \quad [\text{Term (II)}], \end{aligned} \quad (14)$$

where (i) follows from upper bounding the infimum over the policies π_t by the choice of $\pi_t = \hat{\pi}_t$ and (ii) follows from the linearity of the expectation and sub-additivity of the supremum function.

⁵we suppress the dependence of the state $x_t[\pi_{1:t-1}, \zeta_{1:t-1}]$ on the random noise $w_{1:t-1}$.

Focussing on the first term in the above decomposition,

$$\begin{aligned}
\text{Term (I)} &= \left\| \left\| \sup_{p_t \in \mathcal{P}} \mathbb{E} \right\|_{(z_t, \zeta_t) \sim p_t} \right\|_{t=1}^T \left(\sum_{t=1}^T \mathbb{E} [\ell(\hat{\pi}_t, x_t[\hat{\pi}_{1:t-1}, \zeta_{1:t-1}], z_t)] - \mathbb{E} [\ell^\Phi(\hat{\pi}_t, \zeta_{1:t-1}, z_t)] \right) \\
&= \left\| \left\| \sup_{p_t \in \mathcal{P}} \mathbb{E} \right\|_{\zeta_t} \right\|_{t=1}^T \left(\sum_{t=1}^T \mathbb{E} [\ell(\hat{\pi}_t, x_t[\hat{\pi}_{1:t-1}, \zeta_{1:t-1}], z_t)] - \mathbb{E} [\ell^\Phi(\hat{\pi}_t, \zeta_{1:t-1}, z_t)] \right) \\
&\stackrel{(i)}{=} \left\| \left\| \sup_{p_t \in \mathcal{P}} \mathbb{E} \right\|_{\zeta_t} \right\|_{t=1}^T \left(\sum_{t=1}^T \sup_{z_t} |\mathbb{E} [\ell(\hat{\pi}_t, x_t[\hat{\pi}_{1:t-1}, \zeta_{1:t-1}], z_t)] - [\ell^\Phi(\hat{\pi}_t, \zeta_{1:t-1}, z_t)]| \right) \\
&\stackrel{(ii)}{\leq} \sum_{t=1}^T \beta_{\text{RERM}, t}, \tag{15}
\end{aligned}$$

where (i) follows from an application of Hölder's inequality and (ii) follows from the definition of the dynamic stability of the regularized ERM algorithm.

Having established the upper bound on the first term, we now proceed to the second term of equation (14).

$$\begin{aligned}
\text{Term (II)} &\stackrel{(i)}{=} \left\| \left\| \sup_{p_t \in \mathcal{P}} \mathbb{E} \right\|_{(z_t, \zeta_t) \sim p_t} \right\|_{t=1}^T \sup_{\pi \in \Pi} \left(\sum_{t=1}^T \mathbb{E} [\ell^\Phi(\hat{\pi}_t, \zeta_{1:t-1}, z_t)] - \sum_{t=1}^T \ell^\Phi(\pi, \zeta_{1:t-1}, z_t) \right) \\
&\stackrel{(ii)}{\leq} \left\| \left\| \sup_{p_t \in \mathcal{P}} \mathbb{E} \right\|_{(z_t, \zeta_t), z'_t} \right\|_{t=1}^T \sup_{\pi \in \Pi} \left(\sum_{t=1}^T \ell^\Phi(\pi, \zeta_{1:t-1}, z'_t) - \sum_{t=1}^T \ell^\Phi(\pi, \zeta_{1:t-1}, z_t) \right) + \lambda \sup_{\pi \in \Pi} \Omega(\pi) \\
&\stackrel{(iii)}{\leq} \left\| \left\| \sup_{p_t \in \mathcal{P}} \mathbb{E} \right\|_{(z_t, \zeta_t), z'_t} \right\|_{t=1}^T \mathbb{E} \sup_{\pi \in \Pi} \left(\sum_{t=1}^T \epsilon_t (\ell^\Phi(\pi, \zeta_{1:t-1}, z'_t) - \ell^\Phi(\pi, \zeta_{1:t-1}, z_t)) \right) + \lambda \sup_{\pi \in \Pi} \Omega(\pi) \\
&\leq 2 \left\| \left\| \sup_{p_t \in \mathcal{P}} \mathbb{E} \right\|_{(z_t, \zeta_t)} \right\|_{t=1}^T \mathbb{E} \sup_{\pi \in \Pi} \left(\sum_{t=1}^T \epsilon_t \ell^\Phi(\pi, \zeta_{1:t-1}, z_t) \right) + \lambda \sup_{\pi \in \Pi} \Omega(\pi),
\end{aligned}$$

where (i) follows from rewriting the comparator in terms of the counterfactual loss ℓ^Φ , (ii) follows from Lemma 1, and in (iii) we introduce the Rademacher variables ϵ_t . Using Jensen's inequality, we can obtain a further upper bound on Term (II) as

$$\begin{aligned}
\text{Term (II)} &\leq 2 \left\| \left\| \sup_{p_t \in \mathcal{P}} \mathbb{E} \right\|_{(z_t, \zeta_t), \epsilon_t} \right\|_{t=1}^T \sup_{\pi \in \Pi} \left(\sum_{t=1}^T \epsilon_t \ell^\Phi(\pi, \zeta_{1:t-1}, z_t) \right) + \lambda \sup_{\pi \in \Pi} \Omega(\pi) \\
&\leq 2 \sup_{\mathbf{z}, \zeta} \mathbb{E} \sup_{\pi \in \Pi} \left(\sum_{t=1}^T \epsilon_t \ell^\Phi(\pi, [\zeta_1(\epsilon), \dots, \zeta_{t-1}(\epsilon)], \mathbf{z}_t(\epsilon)) \right) + \lambda \sup_{\pi \in \Pi} \Omega(\pi) \tag{16}
\end{aligned}$$

where in the last line, we have replaced the worst case joint distributions over the \mathcal{Z} space by the corresponding worst case \mathcal{Z} -valued trees (see [13, 23] for more details). The upper bound on the value $\mathcal{V}_T(\Pi, \mathcal{Z}, \Phi, \ell)$ now follows from combining the bounds obtained in equations (15) and (16). \square

Proof of Lemma 1. For the purpose of this proof, we will use the short hand $\hat{\pi}_t := \pi_{\text{RERM}, t}$. We will prove the statement of the lemma via an inductive argument on the number of time steps t .

Base Case: For time step $t = 1$, we have that $\pi_{\text{RERM}, 1}$ is the minimizer of the regularized loss implying

$$\mathbb{E} [\ell^\Phi(\hat{\pi}_1, \zeta_\phi, z_1)] \leq \mathbb{E} [\ell^\Phi(\pi, \zeta_\phi, z_1)] + \lambda(\Omega(\pi) - \Omega(\hat{\pi}_1))$$

for any $\pi \in \Pi$.

Inductive Step: Assume that the equation (13) holds for some time step s and consider the cumulative loss at time step $s + 1$

$$\begin{aligned}
\sum_{t=1}^{s+1} \mathbb{E}_{z_t} [\ell^\Phi(\hat{\pi}_t, \zeta_{1:t-1}, z_t)] &= \sum_{t=1}^s \mathbb{E}_{z_t} [\ell^\Phi(\hat{\pi}_t, \zeta_{1:t-1}, z_t)] + \mathbb{E}_{z_{s+1}} [\ell^\Phi(\hat{\pi}_{s+1}, \zeta_{1:s}, z_{s+1})] \\
&\stackrel{(i)}{\leq} \sum_{t=1}^s \mathbb{E}_{z_t} [\ell^\Phi(\hat{\pi}_{s+1}, \zeta_{1:t-1}, z_t)] + \lambda(\Omega(\hat{\pi}_{s+1}) - \Omega(\hat{\pi}_1)) \\
&\quad + \mathbb{E}_{z_{s+1}} [\ell^\Phi(\hat{\pi}_{s+1}, \zeta_{1:s}, z_{s+1})] \\
&\stackrel{(ii)}{\leq} \sum_{t=1}^{s+1} \mathbb{E}_{z_t} [\ell^\Phi(\pi, \zeta_{1:t-1}, z_t)] + \lambda(\Omega(\pi) - \Omega(\hat{\pi}_1)),
\end{aligned}$$

where (i) follows from the induction hypothesis for time s and applying it for $\pi = \hat{\pi}_{s+1}$, and (ii) follows from the fact that $\hat{\pi}_{s+1}$ is the minimizer of the regularized objective at time $s + 1$. This concludes the proof of the lemma. \square

B.2 Proof of Proposition 2

For the purpose of this proof, we restrict our attention to an oblivious adversary wherein the adversary selects instances $\{z_t\}_{t=1}^T$ before the game begins. Several recent works [5, 9] have studied specific versions of a mini-batching algorithms under such an oblivious adversary.

For any mini-batching algorithm with parameter τ , we consider denote by $\hat{T} = T/\tau$ as the effective time horizon⁶ of the game. We now look at the mini-batched value of the game

$$\mathcal{V}_{T,\tau}(\Pi, \mathcal{Z}, \Phi, \ell) \leq \left\langle \left\langle \inf_{q_t} \left\langle \sup_{z_t, s, \zeta_t, s} \right\rangle \right\rangle_{s=1}^{\tau} \mathbb{E}_{\pi_t, s \sim q_t} \right\rangle_{t=1}^{\hat{T}} \mathbb{E}_w[\text{Reg}_T^{\text{pol}}]$$

represents the minimax policy regret for any such mini-batching algorithm \mathcal{A}_τ in the presence of an oblivious adversary. Let us denote the comparator term by

$$\psi(\zeta_{1:T}, z_{1:T}) := \inf_{\pi \in \Pi} \left[\mathbb{E}_w \sum_{t=1}^T \ell(\pi, x_t[\pi^{(t-1)}, \zeta_{1:t-1}], z_t) \right].$$

Following a repeated application of von Neumann's minimax theorem similar to the proof of Proposition 1, we upper bound the value

$$\mathcal{V}_{T,\tau}(\Pi, \mathcal{Z}, \Phi, \ell) \leq \left\langle \left\langle \sup_{\bar{p}_t} \inf_{\pi_t} \mathbb{E}_{(\bar{z}_t, \bar{\zeta}_t) \sim \bar{p}_t} \right\rangle \right\rangle_{t=1}^{\hat{T}} \left[\mathbb{E}_w \sum_{t=1}^{\hat{T}} \sum_{s=1}^{\tau} \ell(\pi_t, x_t[\pi_{1:t-1}^{(\tau)}, \pi_t^{(s-1)}, \bar{\zeta}_{1:t-1}, \bar{\zeta}_t^{1:s-1}], \bar{z}_t^s) - \psi(\bar{\zeta}_{1:\hat{T}}, \bar{z}_{1:\hat{T}}) \right], \tag{17}$$

where the distribution $\bar{p} \in \mathcal{P}^\tau$ is a joint distribution over instances $(\bar{z}, \bar{\zeta}) \in \mathcal{Z}^\tau$ and we have explicitly indicated the dependence of the state variable on the past sequence of policies and adversarial instances. Define the mini-batched loss at time t

$$L_\tau(\pi_t, \bar{\zeta}_{1:t-1}, \bar{z}_t, \bar{\zeta}_t; \pi_{1:t-1}) := \mathbb{E}_w \sum_{s=1}^{\tau} \ell(\pi_t, x_t[\pi_{1:t-1}^{(\tau)}, \pi_t^{(s-1)}, \bar{\zeta}_{1:t-1}, \bar{\zeta}_t^{1:s-1}], \bar{z}_t^s),$$

and the corresponding mini-batched counterfactual loss

$$L_\tau^\Phi(\pi_t, \bar{\zeta}_{1:t-1}, \bar{z}_t, \bar{\zeta}_t) := \mathbb{E}_w \sum_{s=1}^{\tau} \ell(\pi_t, x_t[\pi_t^{(\tau)}, \pi_t^{(s-1)}, \bar{\zeta}_{1:t-1}, \bar{\zeta}_t^{1:s-1}], \bar{z}_t^s).$$

Given these definitions, we can rewrite equation (17) as

$$\mathcal{V}_{T,\tau}(\Pi, \mathcal{Z}, \Phi, \ell) \leq \left\langle \left\langle \sup_{\bar{p}_t} \inf_{\pi_t} \mathbb{E}_{(\bar{z}_t, \bar{\zeta}_t) \sim \bar{p}_t} \right\rangle \right\rangle_{t=1}^{\hat{T}} \left[\sum_{t=1}^{\hat{T}} L_\tau(\pi_t, \bar{\zeta}_{1:t-1}, \bar{z}_t, \bar{\zeta}_t; \pi_{1:t-1}) - \psi(\bar{\zeta}_{1:\hat{T}}, \bar{z}_{1:\hat{T}}) \right]. \tag{18}$$

⁶We assume T/τ to be an integer; if not, it affects the bound by an additive factor of τ .

The expression on the right can be seen as a dual game between a learner and an adversary of \hat{T} rounds. At each round, the adversary reveals a joint distribution \bar{p}_t over the instances and the learner selects a policy q_t . The learner then receives the loss L_τ for that round. We further bound the value by selecting the mini-batched dual ERM strategies for the learner, given by

$$\hat{\pi}_t^\tau := \pi_{\text{ERM},t}^\tau = \operatorname{argmin}_\pi \left(\sum_{s=1}^t \mathbb{E}_{\bar{z}_s, \bar{\zeta}_s} [L_\tau^\Phi(\pi, \bar{\zeta}_{1:s-1}, \bar{z}_s, \bar{\zeta}_s)] \right).$$

Substituting the above mini-batched policies in equation (18) and following a similar set of steps as in proof of Theorem 1, we get,

$$\begin{aligned} \mathcal{V}_{T,\tau}(\Pi, \mathcal{Z}, \Phi, \ell) &\leq \left\langle \left\langle \sup_{\bar{p}_t} \mathbb{E}_{(\bar{z}_t, \bar{\zeta}_t) \sim \bar{p}_t} \right\rangle \right\rangle_{t=1}^{\hat{T}} \left[\sum_{t=1}^{\hat{T}} L_\tau(\hat{\pi}_t^\tau, \bar{\zeta}_{1:t-1}, \bar{z}_t, \bar{\zeta}_t; \hat{\pi}_{1:t-1}^\tau) - \psi(\bar{\zeta}_{1:\hat{T}}, \bar{z}_{1:\hat{T}}) \right] \\ &\leq \left\langle \left\langle \sup_{\bar{p}_t} \mathbb{E}_{(\bar{z}_t, \bar{\zeta}_t) \sim \bar{p}_t} \right\rangle \right\rangle_{t=1}^{\hat{T}} \left[\sum_{t=1}^{\hat{T}} L_\tau(\pi_t, \bar{\zeta}_{1:t-1}, \bar{z}_t, \bar{\zeta}_t; \pi_{1:t-1}) - L_\tau^\Phi(\pi_t, \bar{\zeta}_{1:t-1}, \bar{z}_t, \bar{\zeta}_t) \right] \\ &\quad + \left\langle \left\langle \sup_{\bar{p}_t} \mathbb{E}_{(\bar{z}_t, \bar{\zeta}_t) \sim \bar{p}_t} \right\rangle \right\rangle_{t=1}^{\hat{T}} \left[\sum_{t=1}^{\hat{T}} \mathbb{E}_{\bar{z}_t, \bar{\zeta}_t} [L_\tau^\Phi(\pi_t, \bar{\zeta}_{1:t-1}, \bar{z}_t, \bar{\zeta}_t)] - \psi(\bar{\zeta}_{1:\hat{T}}, \bar{z}_{1:\hat{T}}) \right] \\ &\stackrel{(i)}{\leq} \sum_{t=1}^T \beta_{\text{ERM},t}^\tau + \left\langle \left\langle \sup_{\bar{p}_t} \mathbb{E}_{(\bar{z}_t, \bar{\zeta}_t) \sim \bar{p}_t} \right\rangle \right\rangle_{t=1}^{\hat{T}} \left[\sum_{t=1}^{\hat{T}} \mathbb{E}_{\bar{z}_t, \bar{\zeta}_t} [L_\tau^\Phi(\pi_t, \bar{\zeta}_{1:t-1}, \bar{z}_t, \bar{\zeta}_t)] - \psi(\bar{\zeta}_{1:\hat{T}}, \bar{z}_{1:\hat{T}}) \right], \end{aligned}$$

where step (i) follows from upper bounding the sequence of joint distributions by the worst-case sequence of the adversary instances $\bar{z}_t, \bar{\zeta}_t$. The second term in the above expression can be upper bounded by using an induction argument, similar to that used in Lemma 1. The resulting bound is given by

$$\mathcal{V}_{T,\tau}(\Pi, \mathcal{Z}, \Phi, \ell) \leq \sum_{t=1}^T \beta_{\text{ERM},t}^\tau + \left\langle \left\langle \sup_{\bar{p}_t} \mathbb{E}_{(\bar{z}_t, \bar{\zeta}_t)} \right\rangle \right\rangle_{t=1}^{\hat{T}} \sup_{\pi \in \Pi} \left[\sum_{t=1}^{\hat{T}} \mathbb{E}_{\bar{z}_t, \bar{\zeta}_t} [L_\tau^\Phi(\pi, \bar{\zeta}_{1:t-1}, \bar{z}_t, \bar{\zeta}_t)] - L_\tau^\Phi(\pi, \bar{\zeta}_{1:t-1}, \bar{z}_t, \bar{\zeta}_t) \right].$$

Symmetrizing the above expression and introducing Rademacher variables, we get,

$$\begin{aligned} \mathcal{V}_{T,\tau}(\Pi, \mathcal{Z}, \Phi, \ell) &\leq \sum_{t=1}^T \beta_{\text{ERM},t}^\tau + 2 \left\langle \left\langle \sup_{\bar{p}_t} \mathbb{E}_{(\bar{z}_t, \bar{\zeta}_t) \in \epsilon_t} \right\rangle \right\rangle_{t=1}^{\hat{T}} \sup_{\pi \in \Pi} \left[\sum_{t=1}^{\hat{T}} \epsilon_t L_\tau^\Phi(\pi, \bar{\zeta}_{1:t-1}, \bar{z}_t, \bar{\zeta}_t) \right] \\ &\stackrel{(i)}{\leq} \sum_{t=1}^T \beta_{\text{ERM},t}^\tau + 2 \left\langle \left\langle \sup_{\bar{p}_t} \mathbb{E}_{(\bar{z}_t, \bar{\zeta}_t) \in \epsilon_t} \right\rangle \right\rangle_{t=1}^{\hat{T}} \sup_{\pi \in \Pi} \sum_{s=1}^{\tau} \left[\left| \sum_{t=1}^{\hat{T}} \epsilon_t \ell_{s,t}^\Phi(\pi, \bar{\zeta}_{1:t-1}, \bar{z}_t, \bar{\zeta}_t) \right| \right] \\ &\leq \sum_{t=1}^T \beta_{\text{ERM},t}^\tau + 2 \sum_{s=1}^{\tau} \left\langle \left\langle \sup_{\bar{p}_t} \mathbb{E}_{(\bar{z}_t, \bar{\zeta}_t) \in \epsilon_t} \right\rangle \right\rangle_{t=1}^{\hat{T}} \sup_{\pi \in \Pi} \left[\left| \sum_{t=1}^{\hat{T}} \epsilon_t \ell_{s,t}^\Phi(\pi, \bar{\zeta}_{1:t-1}, \bar{z}_t, \bar{\zeta}_t) \right| \right] \\ &\leq \sum_{t=1}^T \beta_{\text{ERM},t}^\tau + 2\tau \cdot \sup_{s \in [\tau]} \mathfrak{N}_{T/\tau}^{\text{seq}}(\ell_s^\Phi \circ \Pi), \end{aligned}$$

where step (i) follows from swapping the supremum with the summation and in the last step we have used the definition of sequential Rademacher complexity with an absolute value. This establishes the desired claim. \square

C Proofs of lower bounds

C.1 Proof of Theorem 2

We begin by recalling the example instance described in the proof sketch of Theorem 2. The online learning game between learner and adversary is given comprises of the state space $\mathcal{X} = \{x \in \mathbb{R}^d \mid \|x\|_2 \leq 1\}$ and the set of adversarial actions $\mathcal{Z}_\ell^{\text{Lin}} = \{z \in \mathbb{R}^d \mid \|z\|_2 \leq 1\}$ for some dimension $d \geq 3$. In our setup, the adversarial instance space for the dynamics is empty. Given this state space, our policy class Π_{Lin} is a constant class of policies

$$\Pi_{\text{Lin}} = \{\pi_f \mid \pi(x) = f \text{ for all states } x \text{ with } f \in \mathbb{B}_d(1)\},$$

consisting of policies π_f which select the same action f at each state x . With a slight abuse of notation, we represent the policy π_t played by the learner at time t by the corresponding d -dimensional vector f_t . Further, we let the dynamics function $\Phi_{\text{Lin}}(x_t, f_t, \zeta_t) = f_t$ with the noise distribution $\mathcal{D}_w = 0$. Observe that the dynamics simply remembers the last action played by the learner and sets the next state as $x_{t+1} = f_t$ in a deterministic way with the starting state $x_1 = 0$. We now define the loss function which consists of two parts, a linear loss and a L -Lipschitz loss involving the dynamics:

$$\ell_L(f_t, x_t, z_t) = \langle f_t, z_t \rangle + \sigma(f_t, x_t) \quad \text{where} \quad \sigma(f_t, x_t) = \begin{cases} L\|f_t - x_t\|_2 & \text{for } \|f_t - x_t\|_2 \leq \frac{1}{L} \\ 1 & \text{otherwise} \end{cases}. \quad (19)$$

Observe that this example constructs a family of instances one for each value of the Lipschitz constant of L of the function σ . For this setup, the loss function ℓ^Φ for any time $t > 1$ is just the linear part of the loss

$$\ell^\Phi(f, x[f^{t-1}], z) = \langle f, z \rangle + \underbrace{\sigma(f, x)}_{=0} = \langle f, z \rangle.$$

Let us now break down the lower bound analysis into two cases: that of the Lipschitz constant $L \leq 1$ and $L > 1$.

Case 1: $L \leq 1$. For the case when $L \leq 1$, we lower bound the value of the game by ignoring the dynamics loss σ .

$$\begin{aligned} \mathcal{V}_T(\Pi_{\text{Lin}}, \mathcal{Z}^{\text{Lin}}, \Phi_{\text{Lin}}, \ell) &= \left\langle \left\langle \inf_{q_t} \sup_{z_t} \mathbb{E} \right\rangle_{t=1}^T \right\rangle \sum_{t=1}^T \left(\langle f_t, z_t \rangle + \sigma(f_t, f_{t-1}) - \inf_{f \in \Pi_{\text{Lin}}} \sum_{t=1}^T \langle f, z_t \rangle + \sigma(f, x_1) \right) \\ &\stackrel{(i)}{\geq} \left\langle \left\langle \inf_{q_t} \sup_{z_t} \mathbb{E} \right\rangle_{t=1}^T \right\rangle \left(\sum_{t=1}^T \langle f_t, z_t \rangle - \inf_{f \in \Pi_{\text{Lin}}} \left(\sum_{t=1}^T \langle f, z_t \rangle \right) \right) + 1, \end{aligned}$$

where (i) follows by noting that $\sigma(f_t, f_{t-1}) \in [0, 1]$. The above lower bound reduces the value to that of an online linear game between a learner and an adversary. A lower bound on the value of this game can be shown to be $\sqrt{T}/2$ (see [21]) and thus for the case when $L < 1$, we have that the value $\mathcal{V}_T(\Pi_{\text{Lin}}, \mathcal{Z}^{\text{Lin}}, \Phi_{\text{Lin}}, \ell) \geq c\sqrt{T}$ for some $c = 0.5$.

Case 2: $L > 1$. We now proceed to the case when the Lipschitz constant⁷ $L > 1$. In order to prove the requisite lower bound, we will describe the adversary's choice of action z_t . The adversary's strategy is to stick to some action z and only *switch* to a new action when one of events E1 or E2 happen.

E1 The time $t = \lambda L$ for $\lambda = \{1, \dots, T/L\}$.

E2 Let t_0 denote the last time the adversary had switched and denote the expected deviation from the previous move by $\delta_t := \mathbb{E}_{f_t, f_{t-1}} \|f_t - f_{t-1}\|$. Further, let $\Delta_{t_0, t} = \sum_{s=t_0}^t \delta_s$ denote the cumulative deviation of the moves from time t_0 upto time t . The adversary switches at time t whenever $\Delta_{t_0, t} > \frac{1}{L}$.

Given the above events, we now define the adversarial action when it switches. Let t be a time instance when one of E1 or E2 happens. Then the adversary selects z_t such that

$$\|z_t\|_2 = 1, \quad \langle Z_{t-1}, z_t \rangle = 0, \quad \mathbb{E}_{f_t \sim q_t} \langle f_t, z_t \rangle = 0,$$

where $Z_{t-1} = \sum_{s=1}^{t-1} z_s$ is the cumulative sum of the adversary's past actions. Note that our choice of dimensions $d \geq 3$ ensures that such a z_t will always exist.

In order to understand the performance of any algorithm, let us partition the time interval into T/L blocks each of length L and denote each such block $I_i := [L(i-1) + 1, Li]$. Let k_i denote the number of times the learner causes event E2 to occur in the interval I_i . Observe that the cumulative loss within an interval I_i

$$\left\langle \left\langle \mathbb{E} \right\rangle_{\pi_t} \right\rangle_{t \in I_i} \sum_{t \in I_i} \langle f_t, z_t \rangle + \sigma(f_t, f_{t-1}) \quad \begin{cases} = 0 & \text{if } k_i = 0 \\ \geq k_i - 1 & \text{if } k_i \geq 1 \end{cases} \quad (20)$$

⁷Assume L to be an integer; if not, redefine $L = \lfloor L \rfloor$.

where the lower bound for the case $k_i \geq 1$ follows since at each round the learner can only obtain inner product $\mathbb{E}_{f_t} \langle f_t, z_t \rangle \geq -1/L$ at each round of the interval. As soon as $\mathbb{E}_{f_t} \langle f_t, z_t \rangle < -1/L$, the adversary switches and ensures that $\mathbb{E}_{f_t} \langle f_t, z_t \rangle = 0$ for that time. The lower bound of $k_i - 1$ follows since the total length of the interval is L and each time event E2 occurs, the learner pays a cumulative cost of 1. Note that the case for $k_i = 0$ is equivalent to the case $k_i = 1$ and hence going forward, we assume each $k_i \geq 1$.

Let $K = \sum_{i=1}^{T/L} k_i$ denote the total number of times an algorithm causes event E2 to happen and let $K = \beta T/L$ for some $\beta \in [1, L]$. Then, for any sequence of learner distributions $[q_1, \dots, q_T]$, we have that the policy regret is lower bounded as

$$\text{Reg}_T^{\text{pol}}(\mathcal{A}) = \left\langle \left\langle \mathbb{E} \right\rangle_{f_t \sim q_t} \right\rangle_{t=1}^T \left(\langle f_t, z_t \rangle + \sigma(f_t, f_{t-1}) \inf_{f \in \Pi_{\text{Lin}}} \sum_{t=1}^T \langle f, z_t \rangle \right) \stackrel{(i)}{\geq} K - \frac{T}{L} + \|Z_T\|_2, \quad (21)$$

where (i) follows from an application of the Cauchy-Schwarz inequality and from the bound in (20). In order to lower bound the term $\|Z_T\|_2$, we define the set of times $T_{\text{sw}} = [\hat{t}_1, \dots, \hat{t}_{K_{\text{ad}}}]$ where $K_{\text{ad}} \leq T/L + K$ is the total number of switches that the adversary makes. Further, let us denote by \hat{z}_i the choice of the adversary at time \hat{t}_i and by $\gamma_i := \hat{t}_{i+1} - \hat{t}_i$ as the length of the interval for which the adversary played \hat{z}_i . Then, the squared norm

$$\|Z_T\|_2^2 = \left\| \sum_{i=1}^{K_{\text{ad}}} \gamma_i \hat{z}_i \right\|_2^2 = \sum_{i=1}^{K_{\text{ad}}} \gamma_i^2 \|\hat{z}_i\|_2^2 + 2 \sum_{j=2}^{K_{\text{ad}}} \gamma_j \left\langle \sum_{i=1}^{j-1} \gamma_i \hat{z}_i, \hat{z}_j \right\rangle = \sum_{i=1}^{K_{\text{ad}}} \gamma_i^2,$$

where the last inequality follows from the choice of adversary ensuring that $\langle Z_{i-1}, z_i \rangle = 0$ and noting that $\|z_t\| = 1$ for all time t . We can now obtain a lower bound on $\|Z_T\|_2$ by an application of the cauchy-Schwarz inequality as

$$\|Z_T\|_2 = \sqrt{\sum_{i=1}^{K_{\text{ad}}} \gamma_i^2} \geq \frac{\sum_{i=1}^{K_{\text{ad}}} \gamma_i}{\sqrt{K_{\text{ad}}}} = \frac{T}{\sqrt{K_{\text{ad}}}}.$$

Substituting the above value in equation (21) and taking an infimum over all algorithms, we have that the minimax value

$$\mathcal{V}_T(\Pi_{\text{Lin}}, \mathcal{Z}^{\text{Lin}}, \Phi_{\text{Lin}}, \ell) \geq \inf_{\beta \in [1, L]} \left((\beta - 1) \frac{T}{L} + \frac{\sqrt{LT}}{\sqrt{\beta + 1}} \right),$$

where the inequality above follows from setting $K = \beta T/L$ and the fact that $K_{\text{ad}} \leq T/L + K$. Optimizing for the value of β , we get that the minimax value

$$\mathcal{V}_T(\Pi_{\text{Lin}}, \mathcal{Z}, \Phi_{\text{Lin}}, \ell_L) \geq \begin{cases} \frac{\sqrt{T}}{2} & \text{for } 0 < L < 1 \\ \frac{\sqrt{LT}}{\sqrt{2}} & \text{for } 1 \leq L \leq (32T)^{\frac{1}{3}} \\ 2^{\frac{1}{3}} T^{\frac{2}{3}} & \text{for } L > (32T)^{\frac{1}{3}} \end{cases}. \quad (22)$$

Thus, we have that the value is lower bounded by these three different terms each corresponding to different ranges of the Lipschitz constant L . In order to obtain the requisite lower bounds, we now evaluate each term on the right hand side of equations (7a)- (7c).

Bound (7a). This corresponds to the sequential Rademacher complexity of the class Π which corresponds to the unit Euclidean ball with respect to the linear loss. Following the calculations in Rakhlin and Sridharan (see[21, Chapter 10]), we have that

$$\mathfrak{R}_T^{\text{seq}}(\ell^\Phi \circ \mathcal{F}) \leq \sqrt{T}. \quad (23)$$

Bound (7b). In order to establish an upper bound on the dynamic stability parameters, we consider the regularization given by the squared loss as $\Omega(f) = \frac{\|f\|_2^2}{2}$ with some regularization parameter $\lambda \geq 0$. Given that the form of the counterfactual loss ℓ^Φ , the regularized ERM

$$f_{\text{ERM}, t} = \text{Proj}_{\mathbb{B}_d(1)} \left(\frac{1}{\lambda} \sum_{s=1}^t \mathbb{E} [z_s] \right)$$

for the dual game and the adversarial distributions given by $\{p_t\}$. Consequently, the stability parameters

$$\beta_{\text{RERM},t} = \sigma(f_{\text{RERM},t}, f_{\text{RERM},t-1}) \leq L \|f_{\text{RERM},t} - f_{\text{RERM},t-1}\|_2 \leq \frac{L}{\lambda}.$$

Finally, the bound of equation (7b) can now be evaluated as

$$\inf_{\lambda>0} \left(\sum_{t=1}^T \beta_{\text{RERM},t} + \lambda \cdot \sup_{\pi \in \Pi} \Omega(\pi) \right) \leq \inf_{\lambda \geq 0} \left(\frac{LT}{\lambda} + \frac{\lambda}{2} \right) = \sqrt{\frac{LT}{2}}. \quad (24)$$

Bound (7c). We now proceed to the bound given by the mini-batching ERMs with parameter $\tau > 0$. The stability parameters for the mini-batching ERM can be upper bounded as

$$\beta_{\text{ERM},t}^\tau : \begin{cases} \leq 2 & \text{for } t \equiv 0 \pmod{\tau} \\ = 0 & \text{otherwise} \end{cases},$$

where the first case follows trivially from the fact that two unit norm vectors can have a distance at most 2 and the second case is a consequence of the fact that $\ell = \ell^\Phi$ anytime an algorithm repeats the past two policies. Combining this with the sequential Rademacher bound of equation (23) we have

$$\inf_{\tau>0} \left(\sum_{t=1}^T \beta_{\text{ERM},t}^\tau + 2\tau \mathfrak{R}_{T/\tau}^{\text{seq}}(\ell^\Phi \circ \Pi) \right) \leq \inf_{\tau>0} \frac{2T}{\tau} + 2\tau \sqrt{\frac{T}{\tau}} = 2T^{\frac{2}{3}}. \quad (25)$$

Comparing equations (23), (24) and (25) with the lower bounds on the value \mathcal{V}_T in equation (22), we see that the sequential Rademacher bound is tight up to constant factors in the regime $L \leq 1$, the dynamic stability bounds are tight for the regime $1 < L < (32T)^{\frac{1}{3}}$ and the mini-batching bounds are tight for the range $(32T)^{\frac{1}{3}} \leq L \leq T$. This establishes the desired claim. \square

C.2 Proof of Proposition 3

We establish both parts of the proposition separately. For both the subparts, we lower bound the value \mathcal{V}_T by first describing a problem instance $(\Pi, \mathcal{Z}, \Phi, \ell)$ and compute the value for a specific choice of adversarial actions. We assume that the loss function $|\ell(f, z)| \leq 1$ for all $f \in \mathcal{F}$ and $z \in \mathcal{Z}_\ell$. The bounds for larger loss values can be obtained by a corresponding scaling.

C.2.1 Proof of part (a)

We denote by $K = T/\tau$ the number of times a mini-batching algorithm changes its policy.

Constructing online learning with dynamics instance. Given an instance of the online learning problem $(\mathcal{F}, \mathcal{Z}_\ell, \ell)$, we construct the online learning with dynamics instance with state space $\mathcal{X} = \mathcal{F}$ and policy class

$$\Pi_{\mathcal{F}} = \{\pi_f \mid f \in \mathcal{F}, \pi_f(x) = f \text{ for all } x \in \mathcal{X}\},$$

which plays the same action f for all states $x \in \mathcal{X}$. Going forward, with a slight abuse of notation we use the action f and the constant policy π_f interchangeably.

The adversary's loss instance space is given by $\tilde{\mathcal{Z}}_\ell = \mathcal{Z}_\ell \times \{-1, +1\}$ with the actions $z_t \in \mathcal{Z}_\ell$ and $\epsilon_t \in \{-1, +1\}$. The dynamics function $\Phi(x, \pi_f, \zeta) = f$ represent the deterministic dynamics which remembers the last action played by the learner and is not affect by the adversary. The instantaneous loss $\tilde{\ell}(f_t, x_t, (z_t, \epsilon_t))$ is given as

$$\tilde{\ell}(f_t, x_t, (z_t, \epsilon_t)) = \epsilon_t \ell(f_t, z_t) + \mathbb{I}[f_t \neq x_t].$$

With the above loss function, notice that the counterfactual loss $\ell^\Phi(f_t, (z_t, \epsilon_t)) = \epsilon_t \ell(f_t, z_t)$ for all time $t > 1$ and the dynamic stability parameters for any algorithm $\beta_t = \mathbb{E}_{\mathcal{A}}[\mathbb{I}[f_t \neq f_{t-1}]]$.

Specifying the adversary. Given the online learning with dynamics problem above, we now specify an adversary for this setup. Let $K^* = T/\tau^*$ denote the optimal number of switches given by

$$K^* = \operatorname{argmin}_K \left(K + 2 \frac{T}{K} \mathfrak{R}_K^{\text{seq}}(\ell \circ \Pi_{\mathcal{F}}) \right).$$

Note that such a value of K^* is an equalizer of the two terms and ensures that K^* and $\frac{2T}{K^*} \mathfrak{R}_{K^*}^{\text{seq}}$ are equal. Now, consider the worst case \mathcal{Z}_ℓ -valued tree \mathbf{z}_T of depth T corresponding to the online learning problem $(\mathcal{F}, \mathcal{Z}_\ell, \ell)$

$$\mathbf{z}_T = \operatorname{argsup}_{\mathbf{z}} \mathbb{E}_\epsilon \left[\sup_{f \in \mathcal{F}} \sum_{t=1}^T \epsilon_t \ell(f, \mathbf{z}(\epsilon)) \right].$$

The adversary computes the tree \mathbf{z}_{2K^*} produces instances (z_t, ϵ_t) as

Case 1. Whenever $t = \lambda\tau^*/2$ for $\lambda = \{1, \dots, 2T/K^*\}$, the adversary samples ϵ_t as a Rademacher random variable and sets $z_t = \mathbf{z}_{2K^*}(\epsilon_{1:2(t-1)/\tau^*})$.

Case 2. For any time $t \neq \lambda\tau^*$, the adversary computes the probability of switch $p_t^{\text{sw}} = \mathbb{E}_{\mathcal{A}} \mathbb{I}[f_t \neq f_{t-1}]$ and selects instance (z_t, ϵ_t) as

$$(z_t, \epsilon_t) = \begin{cases} (z_{t-1}, \epsilon_t \sim \text{Rad}) & \text{if } p_t^{\text{sw}} > \frac{1}{2} \\ (z_{t-1}, \epsilon_{t-1}) & \text{otherwise} \end{cases}.$$

Lower bound on the value. For any algorithm \mathcal{A} producing distributions q_1, \dots, q_T , the expected policy regret is

$$\begin{aligned} \mathbb{E}_{\mathcal{A}, \epsilon} [\text{Reg}_T^{\text{pol}}] &\stackrel{(i)}{\geq} \sum_{t=1}^T \mathbb{I}[p_t^{\text{sw}} > 0.5] - \mathbb{E}_\epsilon \inf_{f \in \mathcal{F}} \sum_{t=1}^T \ell^\Phi(f, z_t, \epsilon_t) \\ &= \sum_{t=1}^T \mathbb{I}[p_t^{\text{sw}} > 0.5] + \mathbb{E}_\epsilon \sup_{f \in \mathcal{F}} \sum_{t=1}^T \epsilon_t \ell(f, z_t) \end{aligned}$$

where inequality (i) follows from fact that whenever $p_t^{\text{sw}} > 0.5$, the adversary samples a new rademacher variable ϵ_t . For any algorithm, let $K^{\text{sw}} = \sum_t \mathbb{I}[p_t^{\text{sw}} > 0.5]$ denote the number of time periods for which the switching probability is greater than half. We break the lower bound in two separate cases depending on the value of K^{sw} .

Case 1: $K^{\text{sw}} \geq K^*$. For this case, the policy regret for any algorithm can be lower bounded as

$$\mathbb{E}_{\mathcal{A}, \epsilon} [\text{Reg}_T^{\text{pol}}] \geq K^* \stackrel{(i)}{=} \frac{1}{2} \left(K^* + 2 \frac{T}{K^*} \mathfrak{R}_{K^*}^{\text{seq}}(\ell \circ \Pi_{\mathcal{F}}) \right), \quad (26)$$

where (i) follows from our previous observation that $K^* = \frac{2T}{K^*} \mathfrak{R}_{K^*}^{\text{seq}}$.

Case 2: $K^{\text{sw}} < K^*$. For this case, note that the complete time horizon can be divided into at most $3K^*$ intervals wherein the adversary selects the same instances (z, ϵ) , each of length at most $T/2K^*$. By the pigeonhole principle, we must have at least K^* intervals having length $T/2K^*$ beginning at time $t = \lambda\tau^*/2$ for some integral λ . Denote the collection of times in these intervals by \mathcal{I} . We can now lower bound the policy regret as

$$\begin{aligned} \mathbb{E}_{\mathcal{A}, \epsilon} [\text{Reg}_T^{\text{pol}}] &\geq \mathbb{E} \sup_{\epsilon \in \mathcal{F}} \sum_{t=1}^T \epsilon_t \ell(\pi, \mathbf{z}_{2K^*}(\epsilon)) \\ &\stackrel{(i)}{\geq} \mathbb{E}_{\epsilon_t: t \in \mathcal{I}} \left[\sup_{f \in \mathcal{F}} \sum_{t \in \mathcal{I}} \epsilon_t \ell(\pi, \mathbf{z}_{2K^*}(\epsilon)) \right] \\ &\stackrel{(ii)}{\geq} \frac{T}{2K^*} \cdot \mathfrak{R}_{K^*}^{\text{seq}}(\ell \circ \mathcal{F}) \\ &= \frac{1}{4} \left(K^* + 2 \frac{T}{K^*} \mathfrak{R}_{K^*}^{\text{seq}}(\ell \circ \Pi_{\mathcal{F}}) \right) \end{aligned} \quad (27)$$

where (i) follows from the an application of Jensen's inequality and the fact that the resampled ϵ_t when adversary switched because of the learner are not used to parse the tree \mathbf{z}_{2K^*} and (ii) follows from noting that each pair (z, ϵ) was used exactly $T/2K^*$ times.

Combining equations (26) and (27) along with the observation that the minimax value of the online learning with dynamics $\mathcal{V}_T(\Pi, \mathcal{Z}_\ell \times \{+1, -1\}, \Phi, \tilde{\ell})$ is the minimum policy regret for any algorithm establishes the desired claim. \square

C.2.2 Proof of part (b)

We will prove a slightly stronger version of the lower bound from which the desired statement will follow. We follow a strategy similar to the one used in the proof of part (a) above.

Constructing online learning with dynamics instance. Let the dynamics function be defined over states space \mathcal{X} and adversary instance space \mathcal{Z}_Φ . Consider any loss function $\tilde{\ell} : \Pi \times \mathcal{X} \times \tilde{\mathcal{Z}}_\ell \mapsto \mathbb{R}$ for some instance space $\tilde{\mathcal{Z}}_\ell$. We define the space of adversarial loss actions $\mathcal{Z}_\ell = \tilde{\mathcal{Z}}_\ell \times \{-1, +1\}$ and the corresponding loss $\ell(\pi, x, (z, \epsilon)) = \epsilon \cdot \tilde{\ell}(\pi, x, z)$. This defines an instance of the online learning with dynamics problem $(\Pi, \mathcal{Z} = \mathcal{Z}_\ell \times \mathcal{Z}_\Phi, \Phi, \ell)$.

Specifying the adversary. Consider the $\tilde{\mathcal{Z}}_\ell$ and \mathcal{Z}_Φ valued trees \mathbf{z}_T and ζ_T defined as

$$(\mathbf{z}_T, \zeta_T) = \operatorname{argsup}_{\mathbf{z}, \zeta} \mathbb{E}_\epsilon \left[\sup_{\pi \in \Pi} \sum_{t=1}^T \epsilon_t \ell^\Phi(\pi, \zeta_{1:t-1}(\epsilon), \mathbf{z}(\epsilon)) \right],$$

which correspond to the worst-case trees of the sequential Rademacher complexity of the class $\ell^\Phi \circ \Pi$. At every time t , the adversary selects $(z_t, \epsilon_t, \zeta_t)$ by sampling a uniform Rademacher variable and traversing the two trees as

$$\epsilon_t \sim \text{Rad}, \quad z_t = \mathbf{z}_T(\epsilon_{1:t-1}) \quad \text{and} \quad \zeta_t = \zeta_T(\epsilon_{1:t-1}).$$

Lower bound on the value. For any algorithm \mathcal{A} , the expected policy regret is given by

$$\mathbb{E}_{\mathcal{A}, \epsilon} [\text{Reg}_{\mathcal{G}_T^{\text{pol}}}^{\text{pol}}] \stackrel{(i)}{=} \mathbb{E}_\epsilon \left[\sup_{\pi \in \Pi} \sum_{t=1}^T \epsilon_t \ell^\Phi(\pi, \zeta_{1:t-1}, z_t) \right] \stackrel{(ii)}{=} \mathfrak{R}_T^{\text{seq}}(\ell^\Phi \circ \Pi),$$

where (i) follows from noting that the loss at time t is a zero-mean random variable and (ii) is implied by the definition of the trees \mathbf{z}_T and ζ_T .

Finally, observing that the minimax value is equal to the policy regret of the best algorithm completes the proof. \square

D Details of examples

In this section, we work out the examples mentioned in Section 5 in detail and prove the rates for their respective value functions.

Before proceeding to the examples, we introduce some notation. Most of the examples that we consider have dynamics which are not affected by the adversary, that is, the instance space \mathcal{Z}_Φ is empty. We focus on this special case and derive a few results which will be helpful in deriving bounds for the examples.

Borrowing from the theory of stochastic processes, we next define ergodicity of the dynamics which relates a sequence of instantaneous losses to a notion of stationary loss $\ell_*^\Phi : \Pi \times \mathcal{Z} \mapsto \mathbb{R}$.

Definition 4 (Ergodicity). *We say that the dynamics Φ are ergodic with respect to the loss ℓ if for any policy $\pi \in \Pi$ and adversarial action $z \in \mathcal{Z}_\ell$, the expected loss converges to a stationary loss starting from any state x_1 as*

$$\lim_{t \rightarrow \infty} \mathbb{E}_{\{w_t\}} \ell(\pi, x_t[\pi^{(t-1)}], z) = \ell_*^\Phi(\pi, z).$$

The loss function ℓ_*^Φ can be seen as the limit of the counterfactual losses ℓ^Φ and as we shortly show, the losses and dynamics in most of our examples satisfy this ergodicity assumption. For setups where such a stationary loss exists, we define the ergodic stability parameters β_t^* analogous to the dynamic stability parameters.

Definition 5 (Ergodic Stability). *An algorithm \mathcal{A} is said to be $\{\beta_t^*\}$ -ergodic stable if for all sequences of adversarial actions $[z_1, \dots, z_T]$ and time instances $t \in [T]$*

$$|\mathbb{E}_{w_{1:t-1}} [\ell(\pi_t, x_t[\pi_{1:t-1}, w_{1:t-1}], z_t)] - \ell_*^\Phi(\pi_t, z_t)| \leq \beta_t^* \quad \text{where} \quad \pi_t = \mathcal{A}(z_{1:t-1}).$$

Observe that the ergodic stability parameters are defined with respect to the stationary loss as compared to their dynamic stability counterparts which were defined with respect to the counterfactual losses. Next, we define the set of regularized ERMs π_{RERM}^* with respect to these stationary loss as

$$\pi_{\text{RERM},t}^* = \operatorname{argmin}_{\pi \in \Pi} \sum_{t=1}^T \mathbb{E} \left[\ell_*^\Phi(\pi, z_t) \right] + \lambda \cdot \Omega(\pi), \quad (28)$$

for some regularization function Ω and parameter $\lambda \geq 0$. Given this notation, the following corollary upper bounds the value of the game $\mathcal{V}_T(\Pi, \mathcal{Z}, \Phi, \ell)$ in terms of the sequential Rademacher complexity of the loss class $\ell_*^\Phi \circ \Pi$ and the ergodic stability of the RERMs π_{RERM}^* .

Corollary 5. *For any online learning with dynamics instance $(\Pi, \mathcal{Z}, \Phi, \ell)$ with ergodic dynamics Φ , consider the set of regularized ERMs given by eq. (28) with regularization function Ω and parameter $\lambda \geq 0$ having ergodic stability parameters $\{\beta_{\text{RERM},t}^*\}_{t=1}^T$. Then, we have that the value of the game*

$$\mathcal{V}_T(\Pi, \mathcal{Z}, \Phi, \ell) \leq \sum_{t=1}^T \beta_{\text{RERM},t}^* + 2\mathfrak{N}_T^{\text{seq}}(\ell_*^\Phi \circ \Pi) + 2\lambda \sup_{\pi \in \Pi} \Omega(\pi) + \underbrace{\sup_{\pi \in \Pi} \sum_{t=1}^T \left| \ell^\Phi(\pi, z_t, t) - \ell_*^\Phi(\pi, z) \right|}_{\text{Mixing Gap}}. \quad (29)$$

Compared with the corresponding upper bound in Theorem 1, the above bound has an additional term: the worst case deviation of the counterfactual losses⁸ from the stationary losses. This term, which we call the *Mixing Gap*, captures how quickly the dynamics mix to these stationary stationary losses when the same policy is repeatedly played over a period of time. The proof of the corollary is very similar to that of Theorem 1 and we provide it below for completeness.

Proof of Corollary 5. We begin by considering the value of the game and its equivalence to the dual game established by Proposition 1 as

$$\begin{aligned} \mathcal{V}_T(\Pi, \mathcal{Z}, \Phi) &= \left\langle \sup_{p_t \in \mathcal{P}} \inf_{\pi_t} \mathbb{E} \right\rangle_{z_t \sim p_t} \left[\mathbb{E} \left[\sum_{t=1}^T \ell(\pi_t, x_t[\pi_{1:t-1}], z_t) \right] - \inf_{\pi \in \Pi} \mathbb{E} \left[\sum_{t=1}^T \ell(\pi, x_t[\pi^{(t-1)}], z_t) \right] \right] \\ &\stackrel{(i)}{\leq} \left\langle \sup_{p_t \in \mathcal{P}} \mathbb{E} \right\rangle_{z_t \sim p_t} \left[\mathbb{E} \left[\sum_{t=1}^T \ell(\pi_{\text{RERM},t}, x_t[\pi_{\text{RERM},1:t-1}], z_t) \right] - \inf_{\pi \in \Pi} \mathbb{E} \left[\sum_{t=1}^T \ell(\pi, x_t[\pi^{(t-1)}], z_t) \right] \right] \\ &\stackrel{(ii)}{\leq} \left\langle \sup_{p_t \in \mathcal{P}} \mathbb{E} \right\rangle_{z_t \sim p_t} \left[\mathbb{E} \left[\sum_{t=1}^T \mathbb{E} \left[\ell(\pi_{\text{RERM},t}^*, x_t[\pi_{\text{RERM},1:t-1}^*], z_t) \right] \right] - \inf_{\pi \in \Pi} \left(\sum_{t=1}^T \ell_*^\Phi(\pi, z_t) \right) \right] \quad [\text{Term (I)}] \\ &\quad + \left\langle \sup_{p_t \in \mathcal{P}} \mathbb{E} \right\rangle_{z_t \sim p_t} \left[\sup_{\pi \in \Pi} \left(\sum_{t=1}^T \ell_*^\Phi(\pi, z_t) - \mathbb{E} \left[\sum_{t=1}^T \ell(\pi, x_t[\pi^{(t-1)}], z_t) \right] \right) \right], \end{aligned}$$

where (i) follows from replacing the \inf_{π_t} at every time step with $\pi_{\text{RERM},t}^*$ and (ii) follows from the subadditivity of the sup function and the fact that $\inf_y (g(y) + h(y)) \geq \inf_y g(y) + \inf_y h(y)$. The second term in the expression now corresponds to the worst-case deviation of the stationary loss from the counterfactual losses.

Further, observe that Term (I) above is similar to the term obtained in equation 14 and the desired upper bound can be obtained by following the same sequence of steps as in the proof of Theorem 1. \square

Having established the above corollary, we proceed to studying the examples from Section 5 in detail.

D.1 Online Isotron with dynamics

In this section, we look at the online Isotron with dynamics problem introduced in Section 5. The setup consists of a real valued state space $\mathcal{X} = \mathbb{R}$. The policy class Π is based on a function class \mathcal{F} consisting of a 1-Lipschitz function along with a $d + 1$ unit dimensional vector and is given as

$$\begin{aligned} \mathcal{F} &= \{f = (\sigma, \mathbf{w} = (w_1, w)) \mid \sigma : [-1, 1] \mapsto [-1, 1] \text{ 1-Lipschitz, } \mathbf{w} \in \mathbb{R}^{d+1} \mid |w_1| \leq 1, \|w\|_2 \leq 1\}, \\ \Pi_{\mathcal{F}} &= \{\pi_f \mid \pi \in \mathcal{F}, \pi_f(x) = f \text{ for all } x \in \mathcal{X}\}. \end{aligned}$$

⁸since the dynamics are independent of the adversary, we have added an additional time index t to make explicit the number of times policy π is run in the environment.

The adversary selects instances in the space $\mathcal{Z} = [-1, 1]^{d+1} \times [-1, 1]$ and we represent each instance $z = (z_1, \mathbf{x}, y)$. Given this setup, we now formalize the online learning protocol, starting from initial state $x_1 = 0$.

On round $t = 1, \dots, T$,

- the learner selects a policy $\pi_t \in \Pi_{\mathcal{F}}$ and the adversary selects $z_t \in \mathcal{Z}$
- the learner receives loss $\ell(\pi_t, x_t, z_t) = (y_t - \sigma(\langle \mathbf{x}_t, w_t \rangle))^2 + (z_{t,1} - w_{t,1})^2 + (x_t - w_{t,1})^2$
- the state of the system transitions to $x_{t+1} = w_{t,1}$

Given this setup, the next corollary provides a bound on the value of this game $\mathcal{V}_{\text{Iso}, T}(\Pi_{\mathcal{F}}, \mathcal{Z}, \Phi, \ell)$.

Corollary 6 (Online Isotron with dynamics). *For the online Isotron with dynamics problem, there exists a universal constant $c > 0$ such that*

$$\mathcal{V}_{\text{Iso}, T}(\Pi_{\mathcal{F}}, \mathcal{Z}, \Phi, \ell) \leq c\sqrt{T} \log^{3/2}(T).$$

It is worth recalling that the above game is an dynamical extension of the online Isotron problem instance studied by [23]. We are not aware of any primal algorithm which can get a rate of \sqrt{T} for both the online learning version as well the dynamical version of this game. Our non-constructive analysis on the other hand proved a way to guarantee learnability at this rate for the Isotron problem.

Proof of Corollary 6. We prove the above statement by bounding the mixing gap and the ergodic stability parameters for the appropriate regularized ERMs.

Bound on mixing gap. Note that for any time $t > 1$, the losses ℓ_*^Φ and ℓ^Φ are identical since the state variable only depends on the policy at time $t - 1$. Therefore, one can upper bound the loss by constant $c = 12$.

ERMs. For the dual game, we consider the ERM at time t given by

$$f_{\text{ERM}, t} = (\sigma_t, \mathbf{w}_t) = \underset{\sigma, \mathbf{w}}{\operatorname{argmin}} \left\{ \sum_{s=1}^t \left(\mathbb{E}_{z_s \sim p_s} [(y_s - \sigma(\langle \mathbf{x}_s, w \rangle))^2 + (z_{s,1} - w_1)^2] \right) \right\},$$

and set $\pi_t = \pi_{f_{\text{ERM}, t}}$.

Ergodic stability parameters. Note that objective function in the above equation is strongly-convex with respect to the parameter w_1 and a simple calculation shows that $|w_{t,1} - w_{t-1,1}| \leq \frac{2}{t}$. We can now bound the ergodic stability parameter as

$$\beta_{\text{ERM}, t}^* = |\ell(\pi_t, x_t[\pi_{1:t-1}], z_t) - \ell_*^\Phi(\pi_t, z_t)| = |w_{t-1,1} - w_{t,1}|^2 \leq \frac{4}{t^2}. \quad (30)$$

Bound on the value. Having established bounds on the mixing gap and the ergodic stability parameters of the ERM, we now use Corollary 5 to upper bound the value of the game as

$$\begin{aligned} \mathcal{V}_{\text{Iso}, T}(\Pi_{\mathcal{F}}, \mathcal{Z}, \Phi, \ell) &\stackrel{(i)}{\leq} \sum_{t=1}^T \beta_{\text{ERM}, t}^* + 2\mathfrak{R}_T^{\text{seq}}(\ell_*^\Phi \circ \Pi_{\mathcal{F}}) + 16 \\ &\stackrel{\text{Eq. (30)}}{\leq} 8 + 2\mathfrak{R}_T^{\text{seq}}(\ell_*^\Phi \circ \Pi_{\mathcal{F}}) + 16 \\ &\stackrel{(ii)}{\leq} c\sqrt{T} \log^{3/2}(T), \end{aligned}$$

where (i) follows by the upper bound of 16 on the mixing gap and (ii) follows by the corresponding bound on the sequential Rademacher complexity $2\mathfrak{R}_T^{\text{seq}}(\ell_*^\Phi \circ \Pi_{\mathcal{F}})$ from [23, Proposition 18]. \square

D.2 Online Markov decision processes

In this section, we revisit the problem of Online Markov Decision Processes (MDPs) studied in [11]. The setup consists of a finite state space such that $|\mathcal{X}| = S$ and a finite action space with $|\mathcal{U}| = A$. The policy class Π consists of all stationary policies, that is,

$$\Pi_{\text{MDP}} = \{\pi \mid \pi : \mathcal{X} \mapsto \Delta(\mathcal{U})\},$$

where $\Delta(\mathcal{U})$ represents the set of all probability distributions over the action space. In addition, the transitions are drawn according to a known function $P : \mathcal{X} \times \mathcal{U} \mapsto \Delta(\mathcal{U})$. The sequential game then proceeds as follows, starting from some state $x_1 \sim d$:

On round $t = 1, \dots, T$,

- the learner selects a policy $\pi_t \in \Pi_{\text{MDP}}$ and the adversary selects $z_t \in \mathcal{Z} = [0, 1]^{S \times A}$
- the learner receives loss $\ell(\pi_t, x_t, z_t) = z_t(x_t, \pi_t(x_t))$
- the state of the system transitions to $x_{t+1} \sim P(x_t, u_t)$

For every stationary policy π , we let P^f denote the transition function induced by π , that is,

$$P^f(x, x') := \sum_{u \in \mathcal{U}} \pi^u(x) P^{x'}(x, u),$$

where we have used superscript to denote the relevant coordinate of the vector. As in [11], we make the following mixability assumptions about the underlying MDP.

Assumption 1 (MDP Unichain). *We assume that the underlying MDP given by the transition function P is unichain. Further, there exists $\tau \geq 1$ such that for all policies π and distributions $d, d' \in \Delta(\mathcal{U})$ we have*

$$\|dP^\pi - d'P^f\|_1 \leq e^{-1/\tau} \|d - d'\|_1.$$

The parameter τ is often referred to as the mixing time of the MDP. Since the MDP is assumed to be unichain, every policy π has a well defined unique stationary distribution d_π with the stationary loss given by $\ell_\star^\Phi(\pi, z) = \mathbb{E}_{x \sim d_\pi} \mathbb{E}_{u \sim \pi(x)} z(x, u)$. Given this setup, we can obtain an upper bound on the value $\mathcal{V}_{\text{MDP}, T}$ as follows:

Corollary 7 (Online MDP). *For the online Markov Decision Process sequential game satisfying Assumption 1, the value $\mathcal{V}_{\text{MDP}, T}(\Pi_{\text{MDP}}, \mathcal{Z}, \Phi)$ is bounded by*

$$\mathcal{V}_{\text{MDP}, T}(\Pi_{\text{MDP}}, \mathcal{Z}, \Phi, \ell) \leq 4\tau \sqrt{TS \log A} + 2\tau(1 + e^{1/\tau}).$$

The above corollary helps one recover the same $\mathcal{O}(\sqrt{T})$ regret bound that was obtained by [11]. In terms of the dependence of problem specific parameters, while our bound above shows a \sqrt{S} dependence, their bound was independent of S . However note that while the setting studied by [11] consisted of the weaker oblivious adversary, we consider the stronger adaptive adversary which can adapt to the learners strategy.

Proof of Corollary 7. In order to establish the bound, we begin by bounding the ergodic stability parameters as well as the mixing gap for loss ℓ^Φ and ℓ_\star^Φ .

Bound on mixing gap. Consider any policy $\pi \in \Pi_{\text{MDP}}$ and the associated steady state distribution d_π . The stationary loss for this problem is then

$$\ell_\star^\Phi(\pi, z) = \mathbb{E}_{x \sim d_\pi} \mathbb{E}_{u \sim \pi(x)} [z(x, u)].$$

Consider now the difference between the stationary loss and the counterfactual loss at any time t

$$\begin{aligned} |\ell^\Phi(\pi, z, t) - \ell_\star^\Phi(\pi, z)| &= \left| \mathbb{E}_{x_t^\pi \sim d_\pi^t} \mathbb{E}_{u \sim \pi(x_t^\pi)} [z(x_t^\pi, u)] - \mathbb{E}_{x \sim d_\pi} \mathbb{E}_{u \sim \pi(x)} [z(x, u)] \right| \\ &\stackrel{(i)}{=} \left| \mathbb{E}_{x \sim d_\pi^t} [\tilde{z}_\pi(x)] - \mathbb{E}_{x \sim d_\pi} [\tilde{z}_\pi(x)] \right| \\ &\stackrel{(ii)}{\leq} \|\tilde{z}_\pi\|_\infty \cdot \|d_\pi^t - d_\pi\|_1 \\ &\stackrel{(iii)}{\leq} 2e^{-(t-1)/\tau}, \end{aligned} \tag{31}$$

where in (i), we use the redefined loss function $\tilde{z}_\pi(x) := \mathbb{E}_{u \sim \pi(x)} z(x, u)$, (ii) follows from an application of Hölder's inequality, and (iii) follows from Assumption 1 and the fact the $\|d_1 - d_\pi\| \leq 2$.

Ergodic stability parameters. For this setup, we will be using a regularized ERM and parameterize the policy $\pi_{\text{RERM},t}$ as a distribution over the deterministic policies present in Π_{MDP} . Let us denote this subset of policies by $\Pi_{\text{MDP}}^{\text{det}}$. Note that a distribution q in $\mathcal{Q}_{\text{MDP}}^{\text{det}}$ is randomized policy in the class Π_{MDP} . We will work with the negative entropy function as the regularizer.

$$q_{\text{RERM},t} \in \operatorname{argmin}_{q \in \mathcal{Q}_{\text{MDP}}^{\text{det}}} \left(\mathbb{E}_{\pi \sim q} \left[\sum_{s=1}^t \mathbb{E}_{x \sim d_{\pi}} \mathbb{E}_{u \sim \pi(x)} [\bar{z}_2(x, u)] \right] + \lambda \cdot \sum_{i=1}^{|\Pi_{\text{MDP}}^{\text{det}}|} q_i \ln q_i \right),$$

where we denote by $\bar{z}_s = \mathbb{E}_{z_s \sim p_t} z_s$ the expected loss at time s . Now, we can encode the loss at time s for every policy $\pi \in \Pi_{\text{MDP}}^{\text{det}}$ in a vector $\ell_s^{\text{det}} \in [0, 1]^{|\Pi_{\text{MDP}}^{\text{det}}|}$ where the π^{th} coordinate $\ell_{s,\pi}^{\text{det}}$ is the loss for policy π . Given this, we can show that the distribution $q_{\text{RERM},t}$ is given by:

$$(q_{\text{RERM},t})_{\pi} = \frac{\exp\left(-\frac{1}{\lambda} \sum_{s=1}^t \ell_{s,\pi}^{\text{det}}\right)}{\sum_j \exp\left(-\frac{1}{\lambda} \sum_{s=1}^t \ell_{s,j}^{\text{det}}\right)}.$$

Going forward, we drop the RERM term from the distribution $q_{\text{RERM},t}$ for ease of readability. In addition, the boundedness of the loss function $|\ell_{s,\pi}^{\text{det}}| \leq 1$ ensures that the RERM solutions satisfy the following stability property:

$$\|q_t - q_{t+1}\|_1 \leq \frac{1}{\lambda}. \quad (32)$$

Given the above stability, one can also obtain a bound on the action distribution between the randomized policy $\pi_t = \mathbb{E}_{\pi \sim q_t} [\pi]$ and the corresponding π_{t+1} :

$$\|\pi_t(x) - \pi_{t+1}(x)\|_1 = \left\| \mathbb{E}_{\pi \sim q_t} [f(x)] - \mathbb{E}_{\pi \sim q_{t+1}} [f(x)] \right\|_1 = \|q_t - q_{t+1}\|_1 \leq \frac{1}{\lambda},$$

where the second equality follows from the fact that $\|\pi(x)\|_1 = 1$ since they are distributions over the action space \mathcal{U} . Now, following a similar calculation as Lemma 5.2 in [11], we can obtain a bound on the variation in state distributions while playing policies $q_{1:t-1}$ as compared to the steady state distribution d_{q_t} .

$$\|d[q_{1:t-1}] - d_{q_t}\|_1 \leq \frac{2\tau^2}{\lambda} + 2e^{-t/\tau}.$$

With this bound in place, we can now bound the ergodic stability parameters $\beta_{\text{RERM},t}^*$ for the ERM procedure as

$$\begin{aligned} \beta_{\text{RERM},t}^* &= |\mathbb{E}[\ell(\pi_t, x_t[\pi_{1:t-1}], z)] - \ell_{\star}^{\Phi}(\pi_t, z)| \\ &= \left| \mathbb{E}_{x \sim d[q_{1:t-1}]} \bar{z}_{\pi_t}(x) - \mathbb{E}_{x \sim d_{q_t}} \bar{z}_{\pi_t}(x) \right| \\ &\leq \frac{2\tau^2}{\lambda} + 2e^{-t/\tau}. \end{aligned} \quad (33)$$

Bound on the value. Having established bounds on the mixing gap and the RERM ergodic stability parameters, we now proceed to obtain the requisite bound on the value $\mathcal{V}_{\text{MDP},T}(\Pi_{\text{MDP}}, \mathcal{Z}, \Phi, \ell)$.

$$\begin{aligned} \mathcal{V}_{\text{MDP},T}(\Pi_{\text{MDP}}, \mathcal{Z}, \Phi, \ell) &\stackrel{(i)}{\leq} \sum_{t=1}^T \beta_{\text{RERM},t}^* + 2\mathfrak{R}_T^{\text{seq}}(\ell_{\star}^{\Phi} \circ \Pi_{\text{MDP}}) + \sup_{\pi \in \Pi} \sum_{t=1}^T |\ell^{\Phi}(\pi, z_t, t) - \ell_{\star}^{\Phi}(\pi, z)| + \lambda S \log A \\ &\stackrel{\text{Eq. (31)}}{\leq} \sum_{t=1}^T \beta_{\text{RERM},t}^* + 2\mathfrak{R}_T^{\text{seq}}(\ell_{\star}^{\Phi} \circ \Pi_{\text{MDP}}) + 2\tau e^{1/\tau} + \lambda S \log A \\ &\stackrel{\text{Eq. (33)}}{\leq} \frac{2\tau^2}{\lambda} T + 2\mathfrak{R}_T^{\text{seq}}(\ell_{\star}^{\Phi} \circ \Pi_{\text{MDP}}) + 2\tau(1 + e^{1/\tau}) + \lambda S \log A \\ &\stackrel{(ii)}{\leq} 2\tau \sqrt{TS \log A} + 2\mathfrak{R}_T^{\text{seq}}(\ell_{\star}^{\Phi} \circ \Pi_{\text{MDP}}) + 2\tau(1 + e^{1/\tau}) \end{aligned}$$

where (i) follows since the entropy over the class $\Pi_{\text{MDP}}^{\text{det}}$ is upper bounded by $\log |\Pi_{\text{MDP}}^{\text{det}}|$, and (ii) follows by setting $\lambda = \tau \sqrt{\frac{T}{S \log A}}$. Finally, bounding the sequential Rademacher complexity of the finite loss class $\ell_{\star}^{\Phi} \circ \Pi_{\text{MDP}}$ by $2\sqrt{ST \log(A)}$ completes the proof of the corollary. \square

D.3 Online linear quadratic regulator

The online Linear Quadratic Regulator (LQR) setup studied in this section was first studied in [10]. The setup consists of a LQ system - with linear dynamics and quadratic costs - where the cost functions can be adversarial in nature. The comparator class Π_{LQR} comprises a subset of linear policies K which satisfy the following strong stability property.

Definition 6 (Strongly Stable Policy). *A policy K is (κ, γ) -strongly stable (for $\kappa > 0$ and $0 < \gamma < 1$) if $\|K\|_2 \leq \kappa$, and there exists matrices L and H such that $A + BK = HLH^{-1}$, with $\|L\|_2 \leq 1 - \gamma$ and $\|H\|_2 \|H^{-1}\|_2 \leq \kappa$.*

The policy class Π_{LQR} is then defined as $\Pi_{\text{LQR}} = \{K \mid K \text{ is } (\kappa, \gamma)\text{-strongly stable}\}$. Given this policy class, the sequential protocol for this game proceeds as follows, starting from state $x_0 = 0$

On round $t = 1, \dots, T$,

- the learner selects a policy $K_t \in \Pi_{\text{LQR}}$ and the adversary selects instance $z_t \in \mathcal{Z} = (Q_t, R_t)$ such that $Q_t \geq 0$, $R_t \geq 0$ and $\text{tr}(Q_t), \text{tr}(R_t) \leq C$
- the learner receives loss $\ell(\pi_t, x_t, z_t) = x_t^\top Q_t x_t + u_t^\top R_t u_t$
- the state of the system transitions to $x_{t+1} = Ax_t + Bu_t + w_t$

where we assume that the stochastic noise $w_t \sim \mathcal{N}(0, W)$ with $\|W\|_2 \leq \sigma_w$, $\text{tr}(W) \leq \Psi_w$ and $W \geq \tau_w I$. The transition matrices A and B , as well as the noise covariance matrix W are assumed to be known to both the learner and the adversary in advance. Given this setup, the stationary loss is given by

$$\ell_*^\Phi(K, z) = \langle Q + K^\top RK, X_K \rangle = \text{tr}[(Q + K^\top RK)X_K], \quad (34)$$

where $X_K = (A + BK)X_K(A + BK)^\top + W$.

The following lemma establishes certain structural properties of the stationary loss, namely, boundedness over the policy class Π_{LQR} and Lipschitzness with respect to the operator norm.

Lemma 2. *The loss function $\ell_*^\Phi : \Pi_{\text{LQR}} \times \mathcal{Z} \mapsto \mathbb{R}_+$ described in equation (34) satisfies*

$$\ell_*^\Phi(K, z) \leq B_{\max} \quad \text{for all } K \in \Pi_{\text{LQR}}, z \in \mathcal{Z}$$

$$|\ell_*^\Phi(K_1, z) - \ell_*^\Phi(K_2, z)| \leq L_{\text{Lip}} \|K_1 - K_2\|_2 \quad \text{for all } K_1, K_2 \in \Pi_{\text{LQR}}, z \in \mathcal{Z},$$

where $B_{\max} := C(1 + \kappa^2) \frac{\sigma_w \kappa^2}{\gamma}$ and $L_{\text{Lip}} := 4C(1 + \kappa^2) \frac{\sigma_b \kappa^5 \sigma_w}{\gamma^2}$.

We defer the proof of the lemma to the end of section and now proceed to obtain an upper bound on the value $\mathcal{V}_{\text{LQR}, T}$ for the above problem.

Corollary 8 (Online Linear Quadratic Regulator). *For the online LQR sequential game, the value $\mathcal{V}_{\text{LQR}, T}$ is bounded as*

$$\mathcal{V}_{\text{LQR}, T}(\Pi_{\text{LQR}}, \mathcal{Z}, \Phi, \ell) \leq \mathcal{O}\left(\sqrt{T \log(T)}\right),$$

where the \mathcal{O} notation hides the dependence of the bound on problem-specific parameters (see equation (40) for the exact dependencies).

Proof. As before, our strategy is to establish upper bounds on the mixing gap and the RERM ergodic stability parameter for the LQR problem, and using these with Corollary 5 to establish an upper bound on the value $\mathcal{V}_{\text{LQR}, T}$.

Existence of stationary loss. Consider any stable policy $K \in \Pi_{\text{LQR}}$. It is well known that a repeated application of the policy K in the linear dynamics ensures that the state x_t converges to a steady-state distribution, that is, the distribution of x_t and $(A + BK)x_t + w_t$ is the same. Since the noise w_t is assumed to be $\mathcal{N}(0, W)$, the steady-state distribution will also be a normal distribution with mean 0 and steady-state covariance X_K satisfying the following recurrence equation:

$$X_K = (A + BK)X_K(A + BK)^\top + W \quad \text{or equivalently} \quad X_K = \sum_{s=0}^{\infty} (A + BK)^s W (A + BK)^{s\top},$$

and the corresponding steady-state loss is given by:

$$\ell_*^\Phi(K, z) = \langle Q + K^\top RK, X_K \rangle = \text{tr}[(Q + K^\top RK)X_K].$$

Bound on mixing gap. We now proceed to obtain upper bounds on the mixing gap for this problem instance. Going forward, we define $X_{K,t}$ to be the state-covariance matrix at time t when policy K has been used for all preceding timesteps. For the purpose of readability, we will drop the dependence of the covariance matrix on the underlying policy K when it is clear from the context. We begin by looking at the convergence of X_t to the stationary matrix X :

$$\begin{aligned} \|X_t - X\|_2 &= \left\| \sum_{s=0}^{t-1} (A + BK)^s W (A + BK)^s{}^\top - \sum_{s=0}^{\infty} (A + BK)^s W (A + BK)^s{}^\top \right\|_2 \\ &= \left\| \sum_{s=t}^{\infty} (A + BK)^s W (A + BK)^s{}^\top \right\|_2 \\ &\stackrel{(i)}{\leq} \sigma_w \sum_{s=t}^{\infty} \kappa^2 (1 - \gamma)^{2s} \\ &\leq \frac{\sigma_w \kappa^2 (1 - \gamma)^{2t}}{\gamma} \end{aligned}$$

where (i) follows from the fact that $\|A + BK\|^s \leq \kappa(1 - \gamma)^s$ from the strong-stability of K . The above analysis shows that the covariance matrix X_t converges to its stationary distribution exponentially fast. One can also obtain a bound similar to above on $\text{tr}(X - X_t)$ with σ_w replaced by Ψ_w . Having established this convergence, we establish a bound on the mixing gap as

$$\begin{aligned} |\mathbb{E}[\ell^\Phi(\pi, z, t)] - \ell_*^\Phi(\pi, z)| &= |\langle Q + K^\top RK, X_t - X \rangle| \\ &\leq (\sigma_q + \kappa^2 \sigma_r) \cdot \text{tr}(X - X_t) \\ &\leq (\sigma_q + \kappa^2 \sigma_r) \cdot \frac{\Psi_w \kappa^2 (1 - \gamma)^{2t}}{\gamma}. \end{aligned} \quad (35)$$

Since the above bound is independent of the underlying policy K , we have thus established a bound on the mixing gap for the policy class Π_{LQR} .

Regularized ERMs. We now define the class of RERM's we use for the function class Π_{LQR} . Instead of working with a fixed regularization function, we shall look at random perturbations as regularizations. Such an idea is popular in the study of online learning algorithms and is often termed as Follow the Perturbed Leader (FTPL); for a detailed study, see [24, 15]. Thus, the regularized ERM solutions at time t are given by:

$$K_{t,\sigma} = \underset{K \in \Pi_{\text{LQR}}}{\text{argmin}} \left(\sum_{s=1}^t \mathbb{E} \left[\langle Q_s + K^\top RK, X_K \rangle \right] - \langle \sigma, K \rangle \right),$$

where $\sigma \in \mathbb{R}^{k \times d}$ such that each coordinate of $\sigma \sim \text{Exp}(\lambda)$, the exponential distribution with parameter $\lambda > 0$. It was established by [27] that if each of the loss function above is L_{Lip} -Lipschitz, the iterates produced by the FTPL strategy above satisfy:

$$\mathbb{E} [\|K_{t,\sigma} - K_{t+1,\sigma}\|_1] \leq c\lambda \cdot L_{\text{Lip}}(kd)^2 \kappa := \lambda_K,$$

where the norm above is defined element-wise. In Lemma 2, we establish that the losses given by $\ell_*^\Phi(\pi, z)$ are indeed Lipschitz over the space of policies Π_{LQR} . With these set of regularized empirical minimizers, we proceed to now bound the ergodic stability parameters of these regularized ERM's, each one of which is strongly-stable.

Sequential strong-stability of solutions. We first establish that the set of RERM solutions produced by the algorithm satisfy the sequential strong-stability property (see [10] for details) with the appropriate parameters. Note that since each of the K_t (we drop the dependence on the random noise σ) belongs to the class Π_{LQR} , we have that $\|K_t\|_2 \leq \kappa$.

Let $X_t := X_{K_t}$ be the steady-state covariance of the t^{th} solution and \hat{X}_t denote the covariance of the state reached when policies $\{K_1, \dots, K_{t-1}\}$ are applied at the first t timesteps. Consider the following decomposition for $A + BK_t$:

$$A + BK_t = H_t L_t H_t^{-1} \quad \text{where} \quad L_t = X_t^{-1/2} (A + BK_t) X_t^{-1/2}, \quad H_t = X_t^{1/2}.$$

Bound on $\|H_t\|_2$ and $\|H_t^{-1}\|_2$. Using the recursive definition of X_t , we have:

$$\|X_t\|_2 = \left\| \sum_{s=0}^{\infty} (A + BK_t)^s W ((A + BK)^s)^\top \right\|_2 \leq \frac{\sigma_w \kappa^2}{\gamma} \quad (36)$$

The above equation allows us to bound $\|H_t\|_2 \leq \kappa \sqrt{\sigma_w / \gamma} = \beta_h$. Also, by the definition of the matrix X_t , we have that $X \geq W$ and hence $\|H_t^{-1}\| \leq 1/\sqrt{\tau_w} = 1/\alpha_h$. Define $\tilde{\kappa} = \beta_h/\alpha_h$ and note that $\tilde{\kappa} \geq \kappa$.

Bound on $\|L_t\|_2$. Starting from the recursive definition of X_t , we have,

$$\begin{aligned} I &= X_t^{-1/2} (A + BK) X_t (A + BK)^\top X_t^{-1/2} + X_t^{-1/2} W X_t^{-1/2} \\ &\geq L_t L_t^\top + \tau_w X_t^{-1} \\ &\geq L_t L_t^\top + \frac{\tau_w \gamma}{\sigma_w \max(\kappa^2, 1)} I, \end{aligned}$$

which implies that $\|L_t\| \leq 1 - \tilde{\gamma}$ where $\tilde{\gamma} = \frac{\tau_w \gamma}{2\sigma_w \max(\kappa^2, 1)}$.

Bound on $\|X_t - X_{t+1}\|$. As before, we begin with the recursive definitions of X_t and X_{t+1} to get:

$$\begin{aligned} X_{t+1} - X_t &= (A + BK_{t+1}) X_{t+1} (A + BK_{t+1})^\top - (A + BK_t) X_t (A + BK_t)^\top \\ &= (A + BK_{t+1}) (X_{t+1} - X_t) (A + BK_{t+1})^\top + \underbrace{B \Delta_t X_t (A + BK_{t+1})^\top}_{T_1} + \underbrace{(A + BK_t) (B \Delta_t)^\top}_{T_2} \\ &= \sum_{s=0}^{\infty} (A + BK_{t+1})^s (T_1 + T_2) ((A + BK_{t+1})^s)^\top, \end{aligned}$$

where $\Delta_t = K_{t+1} - K_t$. Taking norms on both sides, we get:

$$\|X_{t+1} - X_t\|_2 \leq \frac{2\sigma_b \kappa^5 \sigma_w}{\gamma^2} \|\Delta_t\|_2. \quad (37)$$

Bound on $\|H_{t+1}^{-1} H_t\|_2$. Recall that $H_t = X_t^{1/2}$. In order to bound the required term, we proceed as follows:

$$\begin{aligned} \mathbb{E} \|X_{t+1}^{-1/2} X_t^{1/2}\|_2^2 &= \mathbb{E} \|X_{t+1}^{-1/2} X_t X_{t+1}^{-1/2}\| \\ &\leq \mathbb{E} \|X_{t+1}^{-1/2} X_{t+1} X_{t+1}^{-1/2}\|_2 + \mathbb{E} \|X_{t+1}^{-1/2} (X_{t+1} - X_t) X_{t+1}^{-1/2}\| \\ &\leq 1 + \frac{\mathbb{E} \|X_{t+1} - X_t\|_2}{\tau_w} \\ &\leq 1 + \frac{2\sigma_b \kappa^5 \sigma_w}{\tau_w \gamma^2} \lambda_K \\ &\stackrel{(i)}{\leq} 1 + \tilde{\gamma} \end{aligned}$$

where we bound the term $\|X_{t+1} - X_t\|_2$ using Eq. (37) and (i) follows by setting $\lambda \leq \frac{\tilde{\gamma} \gamma^2 \tau_w}{c \sigma_b \sigma_w L_{\text{Lip}} \kappa^6 (kd)^2}$. Finally, using the fact that $\sqrt{1+x} \leq 1 + x/2$ for $x \in [0, 1]$, we have that $\mathbb{E} \|H_{t+1}^{-1} H_t\|_2 \leq 1 + \tilde{\gamma}/2$.

Ergodic stability parameters. We now proceed to obtain an upper bound on the ergodic stability parameters. Before doing so, we obtain some auxiliary results which will be useful in establishing the final bound.

Bound on $\|\hat{X}_t - X_t\|_2$. We will now obtain a bound on the difference between the observed covariance \hat{X}_t when a sequence of ERM are played and the steady-state covariance matrix X_t . Let us set some notation before we begin with bounding this.

$$\Delta_{x,t} := H_t^{-1} (\hat{X}_t - X_t) (H_t^{-1})^\top \quad \text{and} \quad \mathbb{E}_\sigma \|X_t - X_{t+1}\|_2 \leq \tilde{\lambda}.$$

We then have the following recursion for the term $\Delta_{x,t}$ with the expectation with respect to the sampling of the noise variable σ :

$$\begin{aligned}
\mathbb{E}\|\Delta_{x,t+1}\|_2 &\leq \mathbb{E}\|(H_{t+1}^{-1}H_tL_t)\Delta_{x,t}(H_{t+1}^{-1}H_tL_t)^\top\|_2 + \mathbb{E}\|(H_{t+1}^{-1})(X_t - X_{t+1})((H_{t+1}^{-1})^\top)\|_2 \\
&\leq \mathbb{E}\|L_t\|_2^2\|H_{t+1}^{-1}H_t\|_2^2\|\Delta_{x,t}\|_2 + \frac{\tilde{\lambda}}{\alpha^2} \\
&\stackrel{(i)}{\leq} \left(1 - \frac{\tilde{\gamma}}{2}\right)^2 \mathbb{E}\|\Delta_{x,t}\|_2 + \frac{\tilde{\lambda}}{\alpha^2} \\
&\leq e^{-\tilde{\gamma}t}\|\Delta_{x,1}\|_2 + \frac{\tilde{\lambda}}{\alpha^2\tilde{\gamma}},
\end{aligned}$$

where (i) follows from the bound on $\|L_t\| \leq 1 - \tilde{\gamma}$ and $\|H_{t+1}^{-1}H_t\|_2 \leq (1 + \tilde{\gamma}/2)$. Substituting the value for $\Delta_{x,t}$ in the above bound, we get that:

$$\mathbb{E}\|X_{t+1} - \hat{X}_{t+1}\|_2 \leq \frac{\beta_h^2}{\alpha_h^2} \left(e^{-\tilde{\gamma}t} \mathbb{E}\|\hat{X}_1 - X_1\| + \frac{\tilde{\lambda}}{\tilde{\gamma}} \right). \quad (38)$$

Let us now bound the ergodic stability parameters $\beta_{\text{RERM},t}^*$ as

$$\begin{aligned}
\beta_{\text{RERM},t}^* &= \left| \mathbb{E}_{\sigma w} [\ell(\pi_t, x_t[\pi_{1:t-1}], z)] - \mathbb{E}_{\sigma} [\ell_*^\Phi(\pi_t, z)] \right| \\
&= \left| \mathbb{E}_{\sigma} [\text{tr}((Q + K_{t,\sigma}^\top R K_{t,\sigma}) (\hat{X}_{t,\sigma} - X_{t,\sigma}))] \right| \\
&\leq d(\sigma_q + \kappa^2 \sigma_r) \mathbb{E}_{\sigma} \|\hat{X}_{t,\sigma} - X_{t,\sigma}\|_2 \\
&\stackrel{\text{Eq. (38)}}{\leq} d(\sigma_q + \kappa^2 \sigma_r) \cdot \frac{\beta_h^2}{\alpha_h^2} \left(e^{-\tilde{\gamma}t} \mathbb{E}\|\hat{X}_1 - X_1\| + \frac{\tilde{\lambda}}{\tilde{\gamma}} \right) \\
&\leq d(\sigma_q + \kappa^2 \sigma_r) \cdot \frac{\beta_h^2}{\alpha_h^2} \left(e^{-\tilde{\gamma}t} \cdot \frac{2\sigma_w \kappa^2}{\gamma} + \lambda \cdot \frac{c\sigma_b \kappa^6 \sigma_w d^2 k^2 L_{\text{Lip}}}{\tilde{\gamma} \gamma^2} \right), \quad (39)
\end{aligned}$$

where $\lambda > 0$ is a free parameter corresponding to the noise in the perturbation σ .

Bound on the value. Having established upper bounds on the mixing gap and the ergodic stability parameters, we now bound the value $\mathcal{V}_{\text{LQR},T}$ as

$$\begin{aligned}
\mathcal{V}_{\text{LQR},T}(\Pi_{\text{LQR}}, \mathcal{Z}, \Phi, \ell) &\stackrel{(i)}{\leq} \sum_{t=1}^T \beta_{\text{RERM},t}^* + 2\mathfrak{R}_T^{\text{seq}}(\ell_*^\Phi \circ \Pi_{\text{LQR}}) + \sup_{\pi \in \Pi_{\text{LQR}}} \sum_{t=1}^T |\ell^\Phi(\pi, z_t, t) - \ell_*^\Phi(\pi, z)| + \frac{\kappa k d}{\lambda} \\
&\stackrel{\text{Eq. (35)}}{\leq} \sum_{t=1}^T \beta_{\text{RERM},t}^{\text{LQR}} + 2\mathfrak{R}_T^{\text{seq}}(\ell_*^\Phi \circ \Pi_{\text{LQR}}) + (\sigma_q + \kappa^2 \sigma_r) \cdot \frac{\Psi_w \kappa^2}{\gamma^2} + \frac{\kappa k d}{\lambda} \\
&\stackrel{\text{Eq. (39)}}{\leq} d(\sigma_q + \kappa^2 \sigma_r) \cdot \frac{\beta_h^2}{\alpha_h^2} \left(\frac{2\sigma_w \kappa^2}{\gamma^2} + \lambda T \cdot \frac{c\sigma_b \kappa^6 \sigma_w d^2 k^2 L_{\text{Lip}}}{\tilde{\gamma} \gamma^2} \right) \\
&\quad + 2\mathfrak{R}_T^{\text{seq}}(\ell_*^\Phi \circ \Pi_{\text{LQR}}) + (\sigma_q + \kappa^2 \sigma_r) \cdot \frac{\Psi_w \kappa^2}{\gamma^2} + \frac{\kappa k d}{\lambda} \quad (40)
\end{aligned}$$

where (i) follows from the fact that $\mathbb{E}[\sigma_i] = 1/\lambda$.

To obtain a bound on the sequential Rademacher complexity of the class, observe the the matrices $K \in \mathbb{R}^{k \times d}$. Also, by Lemma 2, we have that the loss ℓ_*^Φ is bounded by B_{max} and Lipschitz with respect to policies K with constant L_{Lip} . Using a standard covering number argument, one can get an ϵ -net of the class Π_{LQR} in the frobenius norm with atmost $O(dk(\frac{1}{\epsilon})^{dk})$ elements. Given this cover, one can upper bound the complexity as

$$\mathfrak{R}_T^{\text{seq}}(\ell_*^\Phi \circ \Pi_{\text{LQR}}) \leq c B_{\text{max}} \sqrt{k d \cdot T \log(k d T L_{\text{Lip}})}$$

for some universal constant $c > 0$. Setting $\lambda = O(1/\sqrt{T})$ concludes the proof of the corollary. \square

D.3.1 Proof of Lemma 2

We establish both parts of the claim separately.

Boundedness of stationary loss. Consider the loss ℓ_*^Φ given by

$$\begin{aligned}\ell_*^\Phi(K, z) &= \text{tr}[(Q + K^\top RK)X_K] \\ &\stackrel{(i)}{\leq} C(1 + \kappa^2)\|X_K\|_2 \\ &\stackrel{(ii)}{\leq} C(1 + \kappa^2)\frac{\sigma_w \kappa^2}{\gamma},\end{aligned}$$

where inequality (i) follows from an applicaion of von Neumann's trace inequality and the trace bounds on the matrices Q and R , and step (ii) follows from equation (36).

Lipschitzness of stationary loss. For any two matrices $K_1, K_2 \in \Pi_{\text{LQR}}$ and instance $z \in \mathcal{Z}$, consider the difference between the stationary losses

$$\begin{aligned}|\ell_*^\Phi(K_1, z) - \ell_*^\Phi(K_2, z)| &\leq |\text{tr}[Q(X_{K_1} - X_{K_2})]| + |\text{tr}[R(K_1 X_{K_1} K_1^\top - K_2 X_{K_2} K_2^\top)]| \\ &\leq C \left((1 + \kappa^2)\|X_{K_1} - X_{K_2}\|_2 + \frac{2\kappa^3 \sigma_w}{\gamma}\|K_1 - K_2\|_2 \right) \\ &\stackrel{(i)}{\leq} 4C(1 + \kappa^2)\frac{\sigma_b \kappa^5 \sigma_w}{\gamma^2}\|K_1 - K_2\|_2,\end{aligned}$$

where step (i) follows from equation (37). This concludes the proof. \square

D.4 Online adversarial tracking

The problem of online tracking of adversarial targets in Linear Quadratic Regulators was first posed in Abbasi et al. [1]. The problem setup involves a state space given by \mathbb{R}^d and a action space \mathbb{R}^k . The sequential game proceeds as follows starting from state $x_1 = 0$

On round $t = 1, \dots, T$,

- the learner selects a policy $\pi_t \in \Pi_{\text{track}}$ and adversary selects $z_t \in \mathcal{Z} = \mathbb{R}^d$ such that $\|z_t\|_2 \leq c_z$
- the learner receives loss $\ell(\pi_t, x_t, z_t) = (x_t - z_t)^\top Q(x_t - z_t) + \pi_t(x_t)^\top \pi_t(x_t)$
- the state of the system transitions to $x_{t+1} = Ax_t + Bu_t$

where the matrices A, B, Q are known in advance to the learner and the adversary. In addition, the matrix Q is positive definite, the pair (A, B) is assumed to be controllable while the pair $(A, Q^{1/2})$ is assumed to be observable. The comparator policy class Π_{track} is assumed to be the following restricted class of linear policies:

$$\Pi_{\text{track}} = \{ \pi = (K, \eta) \mid \|A + BK\|_2 \leq \rho; \|K\|_2 \leq c_K; \|\eta\|_2 \leq c_\eta \},$$

such that the action is given by $u_t = K_t x_t + \eta_t$. For this setup, as we establish later, the stationary loss for any policy $\pi = (K, \eta)$ is given by:

$$\ell_*^\Phi(\pi, z) = (x_*^\pi - z)^\top Q(x_*^\pi - z) + \|Kx_*^\pi + \eta\|_2^2, \quad \text{where } x_*^\pi = (I - (A + BK))^{-1}B\eta$$

Given these preliminaries, we obtain a bound on the value $\mathcal{V}_{\text{tar}, T}$ through the following corollary.

Corollary 9 (Online Tracking). *For the online adversarial tracking sequential game, the value $\mathcal{V}_{\text{tar}, T}$ is bounded by:*

$$\mathcal{V}_{\text{tar}, T}(\Pi_{\text{track}}, \mathcal{Z}, \Phi) \leq \mathcal{O}\left(\sqrt{T \log(T)}\right),$$

where the \mathcal{O} notation hides the dependence of the bound on problem-specific parameters (see equation (45) for the exact dependencies).

In contrast to the result obtained above, [1] provide an algorithm for which the regret for the above problem is bounded by $\mathcal{O}(\log^2 T)$. Obtaining such fast rates in our general framework is an interesting open problem.

Proof of Corollary 9. Our general strategy is to obtain bounds on the the mixing gap and the ergodic stability paramters for certain regularized ERMs. We then use these upper bounds together with Corollary 5 to establish the required upper bound.

Bound on mixing gap. Consider any policy $\pi = (K, \eta)$. We are interested in obtaining a bound on the mixability for this function as:

$$|\ell^\Phi(\pi, z, t) - \ell_*^\Phi(\pi, z)| \leq \beta_{\pi, t}.$$

Let us abbreviate the state $x_t[\pi^{(t-1)}]$ by x_t^π . If we run any policy with the linear dynamics, a steady state x_*^π is reached with

$$x_*^\pi = (A + BK)x_*^\pi + B\eta \quad \text{and therefore} \quad x_*^\pi = \underbrace{(I - (A + BK))^{-1}B\eta}_{:=M_K} = M_K\eta.$$

Then, the corresponding loss at this stationary point is given as

$$\ell_*^\Phi(\pi, z) = (x_*^\pi - z)^\top Q(x_*^\pi - z) + \|Kx_*^\pi + \eta\|_2^2.$$

In order to obtain a bound on the mixing gap, we analyse the convergence of the state x_{t+1}^π to the stationary state x_*^π .

$$\begin{aligned} \|x_{t+1}^\pi - x_*^\pi\|_2 &= \left\| \sum_{s=0}^{\infty} (A - BK)^s B\eta - \sum_{s=0}^{t-1} (A - BK)^s B\eta \right\|_2 \\ &\stackrel{(i)}{\leq} \rho^t \|M_K \eta\|_2, \end{aligned}$$

where (i) follows from the assumption that $\|A + BK\|_2 \leq \rho$.

Next, we consider a bound on the norm of the state x_t^π that is reached by any policy.

$$\begin{aligned} \|x_{t+1}^\pi\|_2 &= \|(A + BK)x_t^\pi + B\eta\|_2 \\ &\stackrel{(i)}{=} \left\| \sum_{s=1}^t (A + BK)^{t-s} B\eta \right\|_2 \\ &\stackrel{(ii)}{\leq} \frac{\|B\|c_\eta}{1 - \rho} := c_x, \end{aligned}$$

where (i) follows from recursively applying the definition of the state evolution and the fact that $x_1 = 0$, and (ii) follows from the assumption that $\|A + BK\|_2 \leq \rho$. Having established the above, we now proceed to obtain a bound on the mixing gap as

$$\begin{aligned} |\ell^\Phi(\pi, z, t) - \ell_*^\Phi(\pi, z)| &= |(x_t^\pi - z)^\top Q(x_t^\pi - z) + \|Kx_t^\pi + \eta\|_2^2 - (x_*^\pi - z)^\top Q(x_*^\pi - z) + \|Kx_*^\pi + \eta\|_2^2| \\ &\stackrel{(i)}{\leq} |(x_t^\pi - x_*^\pi)^\top Q(x_t^\pi - z)| + |(x_t^\pi - x_*^\pi)^\top Q(x_*^\pi - z)| + \|K(x_t^\pi - x_*^\pi)\|_2^2 \\ &\quad + 2\langle Kx_*^\pi + \eta, K(x_t^\pi - x_*^\pi) \rangle \\ &\leq 2\|Q\|(c_x + c_z) \cdot \|x_t^\pi - x_*^\pi\|_2 + c_K^2 \cdot \|x_t^\pi - x_*^\pi\|_2^2 + 2c_K(c_K c_x + c_\eta) \cdot \|x_t^\pi - x_*^\pi\|_2 \\ &\leq \rho^{t-1} \cdot \underbrace{c_x(2\|Q\|(c_x + c_z) + 2c_K(c_K c_x + c_\eta))}_{C_{\text{tar},1}} + \rho^{2(t-1)} \cdot \underbrace{c_x^2 c_K^2}_{C_{\text{tar},2}}, \end{aligned} \quad (41)$$

where (i) follows from adding and subtracting x_*^π in both the terms followed by an application of triangle inequality. For ease of presentation, let us represent the above using constants $C_{\text{tar},1}$ and $C_{\text{tar},2}$ with the knowledge that these depend on the underlying problem parameters but independent of the underlying policy π , that is,

$$|\ell^\Phi(\pi, z, t) - \ell_*^\Phi(\pi, z)| \leq \rho^{t-1} C_{\text{tar},1} + \rho^{2(t-1)} C_{\text{tar},2} \quad (42)$$

Ergodic stability parameters. For obtaining a bound on the ergodic stability parameters, we require a few structural results for the loss ℓ_*^Φ defined above. We present these next and defer their proofs to the end of the section.

Lemma 3 (Equivalence of Tracking Cost). *Consider any policy $\pi = (K, \eta)$ and another stable matrix K' . There exists an η' such that we have $\ell_*^\Phi((K, \eta), z) = \ell_*^\Phi((K', \eta'), z)$ such that*

$$\|\eta'\|_2 \leq 2c_\eta \left(\frac{\|B\|c_K}{(1 - \rho)} + 1 \right).$$

Thus going forward, we consider the ERM procedure on the class of functions $\Pi'_{\text{track}}(K)$, parameterized for a fixed stable policy K , where the bias

$$\eta \leq c'_\eta := \max \left(2c_\eta \left(\frac{\|B\|c_K}{1-\rho} + 1 \right), c_\eta \frac{c_z \|Q\| \|B\|}{\sigma_Q \sigma_B^2 (1-\rho)} \right),$$

with σ_X denotes the smallest non-zero singular value of X . Note that the conclusions of Corollary 5 are still valid with the mixability parameters $\beta_{\Pi'_{\text{track}}(K),t}$. The next lemma establishes the stability of the ERM solutions obtained in consecutive rounds.

Lemma 4. *Fix any ρ -stable policy K . The ERM solutions $\pi_{\text{ERM},t} = (K, \eta_t)$ and $\pi_{\text{ERM},t+1} = (K, \eta_{t+1})$ satisfy the following stability bound:*

$$\|\eta_t - \eta_{t+1}\| \leq \frac{2c_z}{t+1} \cdot \frac{c_\eta \|Q\| \|B\|}{\sigma_Q \sigma_B^2 (1-\rho)} := \frac{\psi_\eta}{t+1}$$

Having established the stability bound above, one can proceed in a manner similar to [1, Lemma 8], one can establish that for $t > \rho \log(T)/(1-\rho)$

$$\|x_t[\pi_{\text{ERM},1:t-1}] - x_*^{\pi_{\text{ERM},t}}\|_2 \leq \underbrace{\frac{\|B\|\psi_\eta}{1-\rho} \cdot \frac{2 \log t}{t - \log t} + \rho^{t-1} \frac{\|B\|c'_\eta}{1-\rho}}_{\psi_{x,t}}. \quad (43)$$

Having established the above, we can now obtain a bound on the ergodic stability parameters for the ERM procedure for the online tracking problem. The calculation is similar to the one done for the mixing gap (see Eq. (41)).

$$\begin{aligned} |\ell(\pi_{\text{ERM},t}, x_t[\pi_{\text{ERM},1:t-1}], z_t) - \ell_*^\Phi(\pi_{\text{ERM},t}, z_t)| &\leq 2\|Q\|(c_x + c_z) \cdot \|x_t[\pi_{\text{ERM},1:t-1}] - x_*^{\pi_{\text{ERM},t}}\|_2 \\ &\quad + c_K^2 \cdot \|x_t[\pi_{\text{ERM},1:t-1}] - x_*^{\pi_{\text{ERM},t}}\|_2^2 \\ &\quad + 2c_K(c_K c_x + c'_\eta) \cdot \|x_t[\pi_{\text{ERM},1:t-1}] - x_*^{\pi_{\text{ERM},t}}\|_2 \\ &\leq \psi_{x,t}(2\|Q\|(c_x + c_z) + 2c_K(c_K c_x + c'_\eta)) + \psi_{x,t}^2 c_K^2, \end{aligned}$$

where we have substituted the bound for $\|x_t[\pi_{\text{ERM},1:t-1}] - x_*^{\pi_{\text{ERM},t}}\|$ from Eq. (43). Thus, we that the ERM ergodic stability parameters are

$$\beta_{\text{ERM},t}^* = \psi_{x,t} \cdot (2\|Q\|(c_x + c_z) + 2c_K(c_K c_x + c'_\eta)) + \psi_{x,t}^2 c_K^2. \quad (44)$$

Bound on the value. We now proceed to obtain a bound on the value $\mathcal{V}_{\text{tar},T}$, beginning from the statement of Corollary 5.

$$\begin{aligned} \mathcal{V}_{\text{tar},T}(\Pi_{\text{track}}, \mathcal{Z}, \Phi, \ell) &\leq \sum_{t=1}^T \beta_{\text{ERM},t}^{\text{tar}} + 2\mathfrak{R}_T^{\text{seq}}(\ell_*^\Phi \circ \Pi'_{\text{track}}(K)) + \sup_{\pi \in \Pi_{\text{LQR}}} \sum_{t=1}^T |\ell^\Phi(\pi, z_t, t) - \ell_*^\Phi(\pi, z)| \\ &\stackrel{\text{Eq. (42)}}{\leq} \sum_{t=1}^T \beta_{\text{ERM},t}^{\text{tar}} + 2\mathfrak{R}_T^{\text{seq}}(\ell_*^\Phi \circ \Pi'_{\text{track}}(K)) + \frac{C_{\text{tar},1}}{1-\rho} + \frac{C_{\text{tar},2}}{1-\rho^2} \\ &\stackrel{\text{Eq. (44)}}{\leq} \frac{\rho \log T}{1-\rho} (\|Q\|(c_x + c_z)^2 + (c_K c_x + c'_\eta)^2) + 2\mathfrak{R}_T^{\text{seq}}(\ell_*^\Phi \circ \Pi'_{\text{track}}(K)) + \frac{C_{\text{tar},1}}{1-\rho} + \frac{C_{\text{tar},2}}{1-\rho^2} \\ &\quad + \left(\frac{2 \log^2 T \|B\| \psi_\eta}{1-\rho} + \frac{\|B\|c'_\eta}{(1-\rho)^2} \right) \cdot (2\|Q\|(c_x + c_z) + 4c_K(c_K c_x + c'_\eta)), \end{aligned} \quad (45)$$

where in the last inequality, we have upper bounded the lower order term $\psi_{x,t}^2$ by $\psi_{x,t}$.

Finally, one can obtain a bound on the sequential complexity by noting that the loss ℓ_*^Φ is bounded since the state $\|x_*^\pi\|_2 \leq c_x$ and is Lipschitz in the bias parameter η with respect to the ℓ_2 norm. Using an argument similar to that from the proof of Corollary 8, we have

$$\mathfrak{R}_T^{\text{seq}}(\ell_*^\Phi \circ \Pi'_{\text{track}}(K)) \leq O\left(\sqrt{dT \cdot \log(dT)}\right).$$

Substituting this bound in equation (45) establishes the corollary. \square

D.4.1 Proof of Lemma 3

Let $V_{\pi,z}(x,u)$ represent the value function for state-action pair (x,u) with respect to policy π and loss function z . Following Lemma 12 from [1] we have that:

$$\ell_*^\Phi(\pi',z) - \ell_*^\Phi(\pi,z) = V_{\pi',z}(x_*^\pi, u_{\pi'}) - V_{\pi',z}(x_*^\pi, u_\pi).$$

The action taken by policy π is given by $u_\pi = Kx_*^\pi + \eta$, while that taken by π' is given by $K'x_*^\pi + \eta'$. If we set the value of η' as:

$$\eta' = (K - K')x_*^\pi + \eta \quad \Rightarrow \quad \ell_*^\Phi(\pi',z) = \ell_*^\Phi(\pi,z).$$

Also, note that one can obtain an upper bound on the norm of η' as $\|\eta'\|_2 \leq 2c_x c_K + c_\eta$ using the bounds on the state x_*^π . \square

D.4.2 Proof of Lemma 4

We begin by characterizing the ERM solution η_t as follows:

$$\begin{aligned} \eta_t &= \operatorname{argmin}_\eta \left(\sum_{t=1}^T \mathbb{E}_{z_t \sim p_t} [(x_*^\pi - z_t)^\top Q(x_*^\pi - z_t) + \|Kx_*^\pi + \eta\|_2^2] \right) \\ &= \operatorname{argmin}_\eta \left(\sum_{t=1}^T \mathbb{E}_{z_t \sim p_t} [(x_*^\pi)^\top [Q + K^\top K] x_*^\pi - 2z_t Q x_*^\pi + \eta^\top \eta + \eta^\top K x_*^\pi + (K x_*^\pi)^\top \eta] \right) \\ &= \operatorname{argmin}_\eta \left(\sum_{t=1}^T \mathbb{E}_{z_t \sim p_t} \left[\underbrace{\eta^\top (M_K^\top (Q + K^\top K) M_K + I + K M_K + M_K^\top K^\top) \eta}_{W} - 2z_t^\top Q M_K \eta \right] \right) \\ &= W^{-1} \left(\frac{1}{t} \sum_{s=1}^t \mathbb{E}_{z_s \sim p_s} [z_s] \right) Q M_K, \end{aligned}$$

where the last equality follows by minimizing the quadratic and the existence of the inverse because $W \geq B^\top Q B$. and the fact that η does not lie in the null space of B (it is always better to set it to zero in that case). This ensures that $\|\eta_t\|_2 \leq c_\eta \frac{c_z \|Q\| \|B\|}{\sigma_Q \sigma_B^2 (1-\rho)}$ and hence the policy $\pi_t = (K, \eta_t) \in \Pi'_{\text{track}}$. We can now obtain the stability bounds as:

$$\begin{aligned} \|\eta_t - \eta_{t+1}\|_2 &= \left\| W^{-1} \left(\frac{1}{t} \sum_{s=1}^t \mathbb{E}_{z_s \sim p_s} [z_s] - \frac{1}{t+1} \sum_{s=1}^{t+1} \mathbb{E}_{z_s \sim p_s} [z_s] \right) Q M_K \right\|_2 \\ &\leq \frac{2c_z}{t+1} \cdot \frac{c_\eta \|Q\| \|B\|}{\sigma_Q \sigma_B^2 (1-\rho)}, \end{aligned}$$

where the final inequality follows from using the bound on $\|M_K\|_2$ as well as the fact that $\|z\|_2 \leq c_z$. This concludes the proof of the lemma. \square

D.5 Online non-linear control

In this section, we look at a non-linear control problem: one formed by extending the LQR problem above to have non-linear deterministic dynamics. We parameterize the dynamics using a non-linear function $\sigma_{\text{NL}} : \mathbb{R}^d \mapsto \mathcal{X}$ as follows:

$$x_{t+1} = \sigma_{\text{NL}}[Ax_t + Bu_t],$$

We assume that the function σ_{NL} is 1-Lipschitz and $\|\sigma_{\text{NL}}(x)\| \leq c_x$ for some $c_x > 0$. This is done to ensure that the dynamics satisfy the ergodicity assumption. We now proceed to define the associated policy class Π_{NL} as

$$\Pi_{\text{NL}} = \{ \pi_\theta \mid \theta \in \mathbb{R}^{d_\theta}, \|\theta\|_2 \leq c_\theta, \|[Ax + B\pi_\theta(x)] - [Ax' + B\pi_\theta(x')]\|_2 \leq (1-\gamma)\|x - x'\|_2 \},$$

where the last condition on the function class establishes a stability condition. In addition, we assume that the function class Π_{NL} satisfies a Lipschitz property:

$$\|\pi_\theta(x) - \pi_{\theta'}(x)\|_2 \leq L_\pi \|\theta - \theta'\|_2 \quad \text{for all } x \in \mathcal{X}.$$

The above basically means that if two parameters θ, θ' are close in the parameter space, then the policies parameterized by them are uniformly close for all states. Note that the class of linear policies Π_{track} (without the bias term) defined for the adversarial tracking problem satisfies the above properties. We next outline the learning protocol, with the game starting with $x_1 = 0$.

On round $t = 1, \dots, T$,

- the learner selects policy $\pi_t \in \Pi_{\text{NL}}$ and the adversary selects $z_t \in \mathcal{Z}$.
- the learner receives loss $\ell(\pi_t, x_t, z_t) \in [0, 1]$
- the state of the system transitions to $x_{t+1} = \sigma_{\text{NL}}[Ax_t + Bu_t]$

For the above setup, we shortly establish that the stationary loss for any policy π is given by

$$\ell_*^\Phi(\pi, z) = \ell(\pi, x_*^\pi, z) \quad \text{where} \quad x_*^\pi = \sigma_{\text{NL}}[Ax_*^\pi + B\pi(x_*^\pi)], \quad (46)$$

where the existence of the fixed point is guaranteed by the stability assumption on the function class in conjunction with the Brouwer fixed-point theorem. In the following lemma, we show that the loss function ℓ_*^Φ above is Lipschitz with respect to the parameter θ .

Lemma 5. *The loss function ℓ_*^Φ given in equation (46) satisfies*

$$|\ell_*^\Phi(\pi_{\theta_1}, z) - \ell_*^\Phi(\pi_{\theta_2}, z)| \leq \underbrace{\left(L_{l,\theta} + L_{l,x} \frac{\|B\|_2 L_\pi}{\gamma} \right)}_{L_{\text{Lip}}} \|\theta_1 - \theta_2\|_2 \quad \text{for all } \pi_{\theta_1}, \pi_{\theta_2} \in \Pi_{\text{NL}}, z \in \mathcal{Z}.$$

We prove the lemma at the end of the section. Taking this as given, we now establish the learnability of the function class Π_{NL} in the following corollary.

Corollary 10 (Online Non-Linear Control). *Consider any value of $\lambda > 0$ and loss function ℓ which is $L_{l,x}$ -Lipschitz in the state space and $L_{l,\theta}$ -Lipschitz in the parameter space with respect to the ℓ_2 norm. For the online non-linear control problem described above, we have that the value*

$$\mathcal{V}_{\text{NL},T}(\Pi_{\text{NL}}, \mathcal{Z}, \Phi) \leq \mathcal{O}\left(\sqrt{T \log(T)}\right),$$

where the \mathcal{O} notation hides the dependence of the bound on problem-specific parameters (see equation (52) for the exact dependencies).

Notice that the above corollary establishes an upper bound of $\tilde{\mathcal{O}}(\sqrt{T})$ for the value $\mathcal{V}_{\text{NL},T}$. Thus, despite the fact that the setup does not have the nice structure of the LQR problem, we are able to establish the learnability of the class Π_{NL} in the online learning with dynamics framework.

Proof of Corollary 10. We begin by establishing a bound on the mixing gap for the class Π_{NL} as well as the ERM ergodic stability parameters. Throughout this section, we would often drop the dependence of the function π_θ on the underlying parameter θ when it is clear from the context.

Bound on mixing gap. Consider any policy $\pi \in \Pi_{\text{NL}}$ and the associated stationary loss

$$\ell_*^\Phi(\pi, z) = \ell(\pi, x_*^\pi, z) \quad \text{where} \quad x_*^\pi = \sigma_{\text{NL}}[Ax_*^\pi + B\pi(x_*^\pi)],$$

where the non-linearity σ_{NL} is applied element-wise to its arguments. Consider now the difference between the stationary and the counterfactual loss as

$$\begin{aligned} |\ell^\Phi(\pi, z, t) - \ell(\pi, x_*^\pi, z)| &\stackrel{(i)}{\leq} L_{l,x} \|x_t^\pi - x_*^\pi\|_2 \\ &\leq L_{l,x} \|\sigma_{\text{NL}}[Ax_{t-1}^\pi + B\pi(x_{t-1}^\pi)] - \sigma_{\text{NL}}[Ax_*^\pi + B\pi(x_*^\pi)]\|_2 \\ &\stackrel{(ii)}{\leq} L_{l,x} (1 - \gamma) \|x_{t-1}^\pi - x_*^\pi\|_2 \\ &\stackrel{(iii)}{\leq} 2L_{l,x} c_x (1 - \gamma)^{t-1}, \end{aligned} \quad (47)$$

where (i) follows from the $L_{l,x}$ -Lipschitzness of the loss function, (ii) follows from the fact that $\pi \in \Pi_{\text{NL}}$, and (iii) follows from the boundedness of the states.

Regularized ERM. Similar to the proof for Corollary 8, we consider the following FTPL based ERM

$$\theta_{t,\sigma} = \operatorname{argmin}_{\theta \in \Theta_{\text{NL}}} \left(\sum_{s=1}^t \mathbb{E}_{z_s \sim p_s} [\ell_*^\Phi(\pi_\theta, z)] - \langle \sigma, \theta \rangle \right),$$

where $\sigma \in \mathbb{R}^{d_\theta}$ such that each coordinate of $\sigma \sim \text{Exp}(\lambda)$, the exponential distribution with parameter $\lambda > 0$. We establish in Lemma 5 that the loss functions defined by $\ell_*^\Phi(\cdot, z)$ are L_{NL} -Lipschitz in the parameter θ and hence the iterates satisfy

$$\mathbb{E}_\sigma [\|\theta_{t,\sigma} - \theta_{t+1,\sigma}\|_1] \leq c\lambda \cdot L_{\text{NL}}(d_\theta)^2 \kappa := \lambda_\theta,$$

where the norm above is defined element-wise.

Ergodic stability parameters. With these set of regularized ERMs, we proceed to now bound the stability parameters of these solutions. Consider again the difference between the stationary and the instantaneous loss:

$$\left| \mathbb{E}_\sigma [\ell(\pi_t, x_t[\pi_{1:t-1}]), z] - \mathbb{E}_\sigma [\ell_*^\Phi(\pi_t, z)] \right| \leq L_{l,x} \mathbb{E}_\sigma [\|x_t[\pi_{1:t-1}] - x_t^{\pi_t}\|_2 + \|x_*^{\pi_t} - x_t^{\pi_t}\|_2], \quad (48)$$

where the above inequality follows from the $L_{l,x}$ Lipschitz property of the loss function in the state space. We have dropped the dependence of π on the noise perturbation σ , underlying parameter θ as well as the fact that these are RERM solutions. In order to obtain the stability parameters, we proceed to obtain a bound on the terms on the right.

Bound on $\|x_*^{\pi_t} - x_t^{\pi_t}\|_2$. The upper bound on this difference is similar to the one we obtained while bounding the mixing gap, the only difference being we have to handle the expectation with respect to the random perturbation σ . Consider,

$$\begin{aligned} \mathbb{E}_\sigma [\|x_*^{\pi_t} - x_t^{\pi_t}\|_2] &= \mathbb{E}_\sigma [\|\sigma_{\text{NL}}[Ax_*^{\pi_t} + B\pi_t(x_*^{K_t})] - \sigma_{\text{NL}}[Ax_{t-1}^{\pi_t} + B\pi_t(x_{t-1}^{\pi_t})]\|_2] \\ &\leq (1-\gamma) \mathbb{E}_\sigma [\|x_*^{K_t} - x_{t-1}^{K_t}\|_2] \\ &\leq (1-\gamma)^{t-1} \cdot 2c_x, \end{aligned} \quad (49)$$

where the sequence of inequalities follows since we have $\pi_{\theta,t,\sigma} \in \Pi_{\text{NL}}$ for any sampling of the perturbation variables σ .

Bound on $\|x_t[\pi_{1:t-1}] - x_t^{\pi_t}\|_2$. Consider a parameter $\tau \geq 1$ to be specified later. We can then decompose the desired difference as follows:

$$\begin{aligned} \|x_t[\pi_{1:t-1}] - x_t^{\pi_t}\|_2 &= \sum_{i=1}^{\tau} (\|x_t[\pi_1, \dots, \pi_{t-i}, \pi_t, \dots, \pi_t] - x_t[\pi_1, \dots, \pi_{t-i-1}, K_t, \dots, \pi_t]\|_2) \\ &\quad + \|x_t[\pi_1, \dots, \pi_{t-\tau-1}, \pi_t, \dots, \pi_t] - x_t[\pi_t, \dots, \pi_t]\|_2 \\ &\leq \sum_{i=1}^{\tau} (\|x_t[\pi_1, \dots, \pi_{t-i}, \pi_t, \dots, \pi_t] - x_t[\pi_1, \dots, \pi_{t-i-1}, \pi_t, \dots, \pi_t]\|_2) + 2c_x(1-\gamma)^{\tau-1}, \end{aligned}$$

where the last inequality follows from a similar calculation as in equation (49). We now focus on the terms in the summation above, focussing on a general term i . Let us redefine the state to be $x_0^i = x_{t-i}[\pi_1, \dots, \pi_{t-i-1}]$. Now, denote by $\hat{x}_j = x_{t-i+j}[\pi_{t-i}, \pi_t, \dots, \pi_t]$ to be the state reached when we select π_{t-i} at the $(t-i)^{\text{th}}$ time instance, followed by π_t for $j-1$ steps. Similarly, $\tilde{x}_j = x_{t-i+j}[\pi_t, \pi_t, \dots, \pi_t]$ is the state reached when one begins from x_0^i and selects π_t for the next j time steps. Bounding the sum above is equivalent to bounding the difference $\tilde{x}_i - \hat{x}_i$.

$$\begin{aligned} \|\tilde{x}_i - \hat{x}_i\|_2 &= \|\sigma_{\text{NL}}[A\tilde{x}_{i-1} + B\pi_t(\tilde{x}_{i-1})] - \sigma_{\text{NL}}[A\hat{x}_{i-1} + B\pi_t(\hat{x}_{i-1})]\|_2 \\ &\leq (1-\gamma)^{i-1} \|\tilde{x}_1 - \hat{x}_1\|_2 \\ &= (1-\gamma)^{i-1} \|\sigma_{\text{NL}}[Ax_0^i + B\pi_{t-i}(x_0^i)] - \sigma_{\text{NL}}[Ax_0^i + B\pi_t(x_0^i)]\|_2 \\ &\leq (1-\gamma)^{i-1} \cdot \|B\|_2 \cdot \|\pi_{t-i}(x_0^i) - \pi_t(x_0^i)\| \\ &\stackrel{(i)}{\leq} (1-\gamma)^{i-1} \cdot \|B\|_2 \cdot L_\pi \|\theta_{t-i} - \theta_t\| \\ &\leq i(1-\gamma)^{i-1} \cdot \|B\|_2 \cdot L_\pi \lambda_\theta \end{aligned}$$

where (i) follows from the Lipschitz assumption on the function class in the parameter space. Setting $\tau = t$ and summing up the above inequalities, we get,

$$\|x_t[\pi_{1:t-1}] - x_t^{\pi_t}\|_2 \leq \frac{\|B\|_2 \cdot L_\pi \lambda_\theta}{\gamma^2} + 2c_x(1-\gamma)^{t-1}. \quad (50)$$

Finally, substituting the bounds obtained in eq. (49) and eq. (50) in eq. (48), we get that:

$$\left| \mathbb{E}_\sigma [\ell(\pi_t, x_t[\pi_{1:t-1}]), z] - \mathbb{E}_\sigma [\ell_*^\Phi(\pi_t, z)] \right| \leq L_{l,x} \left(\frac{\|B\|_2 \cdot L_\pi \lambda_\theta}{\gamma^2} + 4c_x(1-\gamma)^{t-1} \right) := \beta_{\text{RERM},t}^* \quad (51)$$

Bound on the value. Having established the mixing gap and the RERM stability paramters, we now upper bound the value $\mathcal{V}_{\text{NL},T}(\Pi_{\text{NL}}, \mathcal{Z}, \Phi, \ell)$

$$\begin{aligned} \mathcal{V}_{\text{NL},T}(\Pi_{\text{NL}}, \mathcal{Z}, \Phi, \ell) &\stackrel{(i)}{\leq} \sum_{t=1}^T \beta_{\text{RERM},t}^* + 2\mathfrak{N}_T^{\text{seq}}(\ell_*^\Phi \circ \Pi_{\text{NL}}) + \sup_{\pi \in \Pi_{\text{NL}}} \sum_{t=1}^T |\ell^\Phi(\pi, z_t, t) - \ell_*^\Phi(\pi, z)| + \frac{c_\theta d_\theta}{\lambda} \\ &\stackrel{\text{Eq. (47)}}{\leq} \sum_{t=1}^T \beta_{\text{RERM},t}^* + 2\mathfrak{N}_T^{\text{seq}}(\ell_*^\Phi \circ \Pi_{\text{NL}}) + \frac{2L_{l,x}c_x}{\gamma} + \frac{c_\theta d_\theta}{\lambda} \\ &\stackrel{\text{Eq. (51)}}{\leq} \frac{L_{l,x}L_\pi\|B\|_2}{\gamma^2} \cdot \lambda_\theta T + 2\mathfrak{N}_T^{\text{seq}}(\ell_*^\Phi \circ \Pi_{\text{NL}}) + \frac{6L_{l,x}c_x}{\gamma} + \frac{c_\theta d_\theta}{\lambda} \end{aligned} \quad (52)$$

where (i) follows from the fact that $\mathbb{E}[\sigma_i] = 1/\lambda$.

Following the proof technique of Corollary 8, it suffices to establish the stationary loss is bounded (by definition) and is Lipschitz with respect to the underlying parameter (Lemma 5). Combining this with the fact that the parameter $\theta \in \mathbb{R}^d$, we have that the sequential Rademacher complexity is bounded as

$$\mathfrak{N}_T^{\text{seq}}(\ell_*^\Phi \circ \Pi_{\text{LQR}}) \leq c\sqrt{d \cdot T \log(dTL_{\text{Lip}})},$$

for some universal constant $c > 0$. Setting a value of $\lambda = 1/\sqrt{T}$ then concludes the proof of the corollary. \square

D.5.1 Proof of Lemma 5

For any policies $\pi_1 := \pi_{\theta_1} \in \Pi_{\text{LQR}}$ and $\pi_2 := \pi_{\theta_2} \in \Pi_{\text{LQR}}$, and instance $z \in \mathcal{Z}$, consider the difference in the stationary loss

$$\begin{aligned} |\ell_*^\Phi(\pi_1, z) - \ell_*^\Phi(\pi_2, z)| &\leq |\ell(\pi_1, x_*^{\pi_1}, z) - \ell(\pi_1, x_*^{\pi_2}, z)| + |\ell(\pi_1, x_*^{\pi_2}, z) - \ell(\pi_2, x_*^{\pi_2}, z)| \\ &\stackrel{(i)}{\leq} L_{l,\theta} \|\theta_1 - \theta_2\|_2 + L_{l,x} \|x_*^{\pi_1} - x_*^{\pi_2}\|_2 \\ &\stackrel{(ii)}{\leq} \left(L_{l,\theta} + L_{l,x} \frac{\|B\|_2 L_\pi}{\gamma} \right) \|\theta_1 - \theta_2\|_2, \end{aligned}$$

where inequality (i) follows from the Lipschitz property of the loss function with respect to the policy and state space while inequality (ii) follows from the Lipschitz property of the policy. This establishes the desired claim. \square

D.6 Online LQR with adversarial disturbances

In this section, we consider the example of an online learning with dynamics problem where the adversary is allowed to perturb the dynamics in addition to the adversarial losses at each time step. We will focus on the Linear-Quadratic setup where the dynamics function is linear and the costs quadratic in the state x_t and action u_t . Agarwal et al. [2] studied a general version of this problem where they considered the convex cost functions with linear dynamics.

As in the Online LQR example in Section D.3, we consider the class of linear policies Π_{LQR} which are (κ, γ) -strongly stable. Given this policy class, the online learning with dynamics game proceeds as follows, starting from state $x_0 = 0$

On round $t = 1, \dots, T$,

- the learner selects a policy $K_t \in \Pi_{\text{LQR}}$ and the adversary selects instance $z_t = (Q_t, R_t)$ such that $Q_t \geq 0, R_t \geq 0$ and $\text{tr}(Q_t), \text{tr}(R_t) \leq C$ and ζ_t such that $\|\zeta_t\|_2 \leq W$
- the learner receives loss $\ell(\pi_t, x_t, z_t) = x_t^\top Q_t x_t + u_t^\top R_t u_t$ where action $u_t = K_t x_t$
- the state of the system transitions to $x_{t+1} = Ax_t + Bu_t + \zeta_t$

where we assume that the transition matrices A and B are known to both the learner and adversary in advance. Observe that in this case, a stationary loss ℓ_*^Φ does not exist because of the adversarial perturbations ζ_t in the dynamic; indeed, if a learner repeatedly plays the same policy $K \in \Pi_{\text{LQR}}$, the state of the system is not guaranteed to converge to a unique stationary state. We now proceed to obtain an upper bound on the value $\mathcal{V}_{\text{adv}, T}$ in the following corollary, by directly controlling the dynamic stability parameters $\{\beta_{\text{RERM}, t}\}$ for this policy class Π_{LQR} with a similar FTPL based regularized ERM as used in the proof of Corollary 8.

Corollary 11 (LQR with adversarial disturbances). *For the online LQR with adversarial disturbances problem, the value $\mathcal{V}_{\text{adv}, T}$ is bounded as*

$$\mathcal{V}_{\text{adv}, T} \leq \mathcal{O}(\sqrt{T \log(T)}),$$

where the \mathcal{O} notation hides the dependence on problem-specific parameters.

Proof. As we discussed above, the stationary losses ℓ_*^Φ do not exist for this setup. Instead, we will work directly with the counterfactual losses ℓ^Φ for this setup. Recall from Definition 2, the counterfactual loss at time t for some linear policy $K \in \Pi_{\text{LQR}}$ is defined as

$$\ell_t^\Phi(K_t, \zeta_{1:t}, z_t) = \ell(K_t, x_t[K_t^{(t-1)}, \zeta_{1:t-1}], z_t).$$

To instantiate the above counterfactual for the LQR problem, we will define some notation. Let us denote by $x_t = x_t[K_{1:t-1}, \zeta_{1:t-1}]$ the state at time reached by playing the sequence of policies $K_{1:t-1}$ and by $\tilde{x}_t = x_t[K_t^{(t-1)}, \zeta_{1:t-1}]$ the state when the learner plays policies K_t for the first $t-1$ time steps. Further, let us denote by $X_t = x_t x_t^\top$ the rank 1 covariance matrix at time t for state x_t and similarly $\tilde{X}_t = \tilde{x}_t \tilde{x}_t^\top$ for state \tilde{x}_t . With this notation, we have the losses

$$\ell_t(K_t, x[K_{1:t-1}, \zeta_{1:t-1}], z_t) = \text{tr}((Q_t + K_t^\top R_t K_t) X_t) \quad \text{and} \quad \ell_t^\Phi(K_t, \zeta_{1:t-1}, z_t) = \text{tr}((Q_t + K_t^\top R_t K_t) \tilde{X}_t). \quad (53)$$

We now proceed to define the regularized ERM that we shall use and derive an upper bound on the dynamic stability parameters.

Regularized ERM. As in the proof of Corollary 8, we will consider the class of dual regularized ERM derived by the FTPL strategy

$$K_{t, \sigma} = \underset{K \in \Pi_{\text{LQR}}}{\text{argmin}} \left(\sum_{s=1}^t \mathbb{E} [\langle Q_s + K^\top R K, \tilde{X}_s \rangle] - \langle \sigma, K \rangle \right) \quad (54)$$

where $\sigma \in \mathbb{R}^{d \times k}$ such that each coordinate of σ is sampled i.i.d. from the exponential distribution with parameter $\lambda > 0$. Following a similar argument as the one in the proof of Corollary 8, we have

$$\mathbb{E}_\sigma [\|K_{t, \sigma} - K_{t+1, \sigma}\|_1] \leq c\lambda \cdot L_{\text{Lip}}(kd)^2 \kappa := \lambda_K, \quad (55)$$

where the constant L_{Lip} represents the Lipschitz constant of the function ℓ^Φ (see Lemma 6).

Dynamic stability parameters. For any time $t > 0$ and the policies $\{K_t\}$ given by equation (54) (we drop the dependence on the noise σ) and any sequence of adversarial instances $(\zeta_{1:t}, z_{1:t})$, we have

$$|\ell_t^\Phi(K, \zeta_{1:t-1}, z_t) - \ell_t(K_t, x_t, z_t)| = |\langle Q_t + K_t^\top R_t K_t, X_t - \tilde{X}_t \rangle| \leq \text{tr}(Q_t + K_t^\top R_t K_t) \cdot \|X_t - \tilde{X}_t\|_2. \quad (56)$$

Thus, in order to obtain a bound on the dynamic stability parameters, we need to obtain a bound on the spectral norm of the difference $X_t - \tilde{X}_t$. To do so, we begin by bounding the distance between

the states x_t and \tilde{x}_t as

$$\begin{aligned}
x_t - \tilde{x}_t &= (A + BK_{t-1})x_{t-1} + \zeta_{t-1} - (A + BK_t)\tilde{x}_{t-1} - \zeta_{t-1} \\
&= (A + BK_t)(x_{t-1} - \tilde{x}_{t-1}) + B(K_{t-1} - K_t)x_{t-1} \\
&= (A + BK_t)^{t-1}(x_1 - \tilde{x}_1) + \sum_{s=2}^{t-1} (A + BK_t)^{t-s} B(K_s - K_t)x_s, \tag{57}
\end{aligned}$$

where the final inequality follows by unrolling the recursion. Observe that the first term in the above equality is 0 since both $x_1 = \tilde{x}_1 = \zeta_1$. Taking the ℓ_2 norm on both sides, we get,

$$\begin{aligned}
\|x_t - \tilde{x}_t\|_2 &\stackrel{(i)}{\leq} C_x \sigma_B \kappa \lambda_K \sum_{s=2}^{t-1} (1 - \gamma)^{t-s} (t - s) \\
&\stackrel{(ii)}{\leq} \frac{C_x \sigma_B \kappa \lambda_K}{\gamma} =: C_{x,2} \tag{58}
\end{aligned}$$

where inequality (i) follows by using the fact that K_t is (κ, γ) -strongly stable and the bound on the norm of the state $\|x_s\|_2 \leq \frac{\kappa W}{\gamma} := C_x$ and (ii) follows by summing up the series. With this bound, we obtain an expression for the difference between the covariances at time $t + 1$ as

$$\begin{aligned}
\tilde{X}_{t+1} - X_{t+1} &= \left(\zeta_t ((A + BK_{t+1})\tilde{x}_t)^\top + (A + BK_{t+1})\tilde{x}_t \zeta_t^\top + (A + BK_{t+1})\tilde{X}_t (A + BK_{t+1})^\top \right) \\
&\quad - \left(\zeta_t ((A + BK_t)x_t)^\top + (A + BK_t)x_t \zeta_t^\top + (A + BK_t)X_t (A + BK_t)^\top \right) \\
&= \underbrace{\zeta_t (\tilde{x}_t - x_t)^\top (A + BK_{t+1})^\top + (A + BK_{t+1})(\tilde{x}_t - x_t) \zeta_t^\top}_{\Delta_1^t} \\
&\quad + \underbrace{\zeta_t x_t^\top (K_{t+1} - K_t)^\top B^\top + B(K_{t+1} - K_t)x_t \zeta_t^\top}_{\Delta_2^t} \\
&\quad + \underbrace{B(K_{t+1} - K_t)X_t (A + BK_t)^\top + (A + BK_{t+1})X_t (B(K_{t+1} - K_t))^\top}_{\Delta_3^t} \\
&\quad + (A + BK_{t+1})(\tilde{X}_t - X_t)(A + BK_{t+1})^\top.
\end{aligned}$$

Let us denote by $\Delta_X^{t+1} := \tilde{X}_{t+1} - X_{t+1}$ the difference between the covariance at time $t + 1$ and by $\tilde{K}_{t+1} := A + BK_{t+1}$. With this notation, we can rewrite the above as

$$\begin{aligned}
\Delta_X^{t+1} &= (\tilde{K}_{t+1})\Delta_X^t(\tilde{K}_{t+1})^\top + \sum_{i=1}^3 \Delta_i^t \\
&= \tilde{K}_{t+1}^2 \Delta_X^{t-1} (\tilde{K}_{t+1}^2)^\top + \tilde{K}_{t+1} \sum_{i=1}^3 \Delta_i^{t-1} \tilde{K}_{t+1}^\top + \sum_{i=1}^3 \Delta_i^t \\
&= \tilde{K}_{t+1}^t \Delta_X^1 (\tilde{K}_{t+1}^t)^\top + \sum_{i=1}^3 \sum_{s=2}^t \tilde{K}_{t+1}^{t-s+1} \Delta_i^s (\tilde{K}_{t+1}^{t-s+1})^\top, \tag{59}
\end{aligned}$$

where in the last equality observe that $\Delta_X^1 = 0$. In order to bound the deviation $\|\Delta_X^{t+1}\|_2$, we will now bound each of three terms in the above equation separately.

Bound for Δ_1 . To obtain a bound on the term with the error Δ_1 , recall from equation (58) that we have $\|x_t - \tilde{x}_t\|_2 \leq C_{x,2}$. With this, we have,

$$\begin{aligned}
\sum_{s=2}^t \tilde{K}_{t+1}^{t-s+1} \Delta_1^s (\tilde{K}_{t+1}^{t-s+1})^\top &\leq 2W C_{x,2} \kappa^3 \cdot \sum_{s=2}^t (1 - \gamma)^{2(t-s+1)} \\
&\leq \frac{2W \kappa^3}{\gamma} \cdot C_{x,2}. \tag{60}
\end{aligned}$$

Bound for Δ_2 . For the error term corresponding to Δ_2 , recall that we have $\|K_{t+1} - K_s\|_2 \leq (t+1-s) \cdot \lambda_K$. Substituting this in the error term, we have,

$$\begin{aligned} \sum_{s=2}^t \tilde{K}_{t+1}^{t-s+1} \Delta_2^s (\tilde{K}_{t+1}^{t-s+1})^\top &\leq 2\sigma_B W C_x \lambda_K \kappa^2 \sum_{s=2}^t (1-\gamma)^{2(t-s+1)} (t-s+1) \\ &\leq \frac{2\sigma_B W C_x \kappa^2}{\gamma^2} \cdot \lambda_K. \end{aligned} \quad (61)$$

Bound for Δ_3 . Finally, for the error term Δ_3 , observe that the spectral norm of the covariance $\|X_t\|_2 \leq C_x^2$, and substituting this in the sum, we have

$$\begin{aligned} \sum_{s=2}^t \tilde{K}_{t+1}^{t-s+1} \Delta_3^s (\tilde{K}_{t+1}^{t-s+1})^\top &\leq 2C_x^2 \sigma_B \kappa^3 \lambda_K \sum_{s=2}^t (1-\gamma)^{2(t-s+1)} (t-s+1) \\ &\leq \frac{C_x^2 \sigma_B \kappa^3}{\gamma^2} \cdot \lambda_K. \end{aligned} \quad (62)$$

Substituting the bounds obtained in equations (60), (61) and (62) in the upper bound on the stability parameters in equation (56), we have that the dynamic stability parameters $\beta_{\text{RERM},t} = c_\beta \lambda_K$, where the constant c_β depends on problem dependent parameters and can be obtained from the above equations. Having established a bound on the dynamic stability, we now upper bound the value for this problem.

Bound on the value. The value of the online LQR with adversarial disturbance problem is

$$\begin{aligned} \mathcal{V}_{\text{adv},T}(\Pi_{\text{LQR}}, \mathcal{Z}, \Phi, \ell) &\stackrel{(i)}{\leq} \sum_{t=1}^T \beta_{\text{RERM},t} + 2\mathfrak{R}_T^{\text{seq}}(\ell^\Phi \circ \Pi_{\text{LQR}}) + \frac{\kappa k d}{\lambda} \\ &\stackrel{(ii)}{\leq} c_\beta \lambda_K \cdot T + 2\mathfrak{R}_T^{\text{seq}}(\ell^\Phi \circ \Pi_{\text{LQR}}) + \frac{\kappa k d}{\lambda}, \end{aligned} \quad (63)$$

where (i) follows from the fact that $\mathbb{E}[\sigma_i] = 1/\lambda$ and (ii) follows from noting that each of the dynamic stability parameters is upper bounded by $c_\beta \lambda_K$.

To obtain a bound on the sequential Rademacher complexity of the class, observe that by Lemma 6, we have that the loss ℓ^Φ is bounded by B_{max} and Lipschitz with respect to policies K with constant L_{Lip} . Using a standard covering number argument, one can get an ϵ -net of the class Π_{LQR} in the frobenius norm with atmost $O(dk(\frac{1}{\epsilon})^{dk})$ elements. Given this cover, one can upper bound the complexity as

$$\mathfrak{R}_T^{\text{seq}}(\ell_*^\Phi \circ \Pi_{\text{LQR}}) \leq c B_{\text{max}} \sqrt{kd \cdot T \log(kdT L_{\text{Lip}})}$$

for some universal constant $c > 0$. Setting $\lambda = O(1/\sqrt{T})$ concludes the proof of the corollary. \square

Lemma 6. For the counterfactual loss $\ell_t^\Phi(K, \zeta_{1:t-1}, z_t)$ defined in equation (53) and policy class Π_{LQR} , we have

a ℓ_t^Φ is bounded by $|\ell^\Phi(K, \zeta_{1:t-1}, z_t)| \leq C(1 + \kappa^2) \cdot \left(\frac{\kappa W}{\gamma}\right)^2 := B_{\text{max}}$.

b ℓ_t^Φ is Lipschitz with respect to K with

$$|\ell_t^\Phi(K_1) - \ell_t^\Phi(K_2)| \leq C \left(2C_x^2 \kappa + (\kappa^2 + 1) \left(\frac{2\kappa^4 C_x \sigma_B W}{\gamma^2} + \frac{\kappa^3 C_x^2 \sigma_B}{\gamma} \right) \right) \cdot \|K_1 - K_2\|_2.$$

Proof. We will establish both the parts separately.

Proof for part (a). Consider the counterfactual loss ℓ_t^Φ at time t

$$\begin{aligned} |\ell^\Phi(K, \zeta_{1:t-1}, z_t)| &= \langle Q_t + K^\top R_t K, \tilde{x}_t \tilde{x}_t^\top \rangle \\ &\leq \|Q_t + K^\top R_t K\|_2 \cdot \|\tilde{x}_t\|_2^2 \\ &\leq C(1 + \kappa^2) \cdot \left(\frac{\kappa W}{\gamma}\right)^2, \end{aligned} \quad (64)$$

where the last inequality follows by using the fact that the policy K is (κ, γ) -strongly stable and that $\|x\|_2 \leq C_x := \frac{\kappa W}{\gamma}$.

Proof for part (b). Consider any two linear policies $K_1, K_2 \in \Pi_{\text{LQR}}$. The difference in the counterfactual losses is given by

$$\begin{aligned} |\ell_t^\Phi(K_1, \zeta_{1:t-1}, z_t) - \ell_t^\Phi(K_2, \zeta_{1:t-1}, z_t)| &= |\langle Q + K_1^\top R K_1, \tilde{X}_1 \rangle - \langle Q + K_2^\top R K_2, \tilde{X}_2 \rangle| \\ &\leq C(\kappa^2 + 1) \|\tilde{X}_{1,t} - \tilde{X}_{2,t}\|_2 + 2C_x^2 C \kappa \|K_1 - K_2\|_2, \end{aligned} \quad (65)$$

where we have used the notation $\tilde{X}_t = \tilde{x}_t \tilde{x}_t^\top$ to denote the covariance at time t and the final inequality follows by noting that $\text{tr}(Q)$, $\text{tr} R \leq C$ and $\|K_i\|_2 \leq \kappa$ for $i = \{1, 2\}$. Let us denote by $\tilde{K} = A + BK$ the effective state transition matrix. We now focus on the term corresponding to the difference of the covariances $\tilde{X}_{1,t} - \tilde{X}_{2,t}$.

$$\begin{aligned} \tilde{X}_{1,t} - \tilde{X}_{2,t} &= (\tilde{K}_1 \tilde{x}_{1,t-1} \zeta_{t-1}^\top + \zeta_{t-1} \tilde{x}_{1,t-1}^\top \tilde{K}_1^\top + \tilde{K}_1 \tilde{X}_{1,t-1} \tilde{K}_1^\top) - (\tilde{K}_2 \tilde{x}_{2,t-1} \zeta_{t-1}^\top + \zeta_{t-1} \tilde{x}_{2,t-1}^\top \tilde{K}_2^\top + \tilde{K}_2 \tilde{X}_{2,t-1} \tilde{K}_2^\top) \\ &= \underbrace{(\tilde{K}_1 - \tilde{K}_2) \tilde{x}_{1,t-1} \zeta_{t-1}^\top + \zeta_{t-1} \tilde{x}_{1,t-1}^\top (\tilde{K}_1 - \tilde{K}_2)^\top}_{\Delta_{1,t-1}} + \underbrace{\tilde{K}_2 (\tilde{x}_{1,t-1} - \tilde{x}_{2,t-1}) \zeta_{t-1}^\top + \zeta_{t-1} (\tilde{x}_{1,t-1} - \tilde{x}_{2,t-1})^\top \tilde{K}_2^\top}_{\Delta_{2,t-1}} \\ &\quad + \underbrace{(\tilde{K}_1 - \tilde{K}_2) \tilde{X}_{1,t-1} \tilde{K}_1^\top + \tilde{K}_2 \tilde{X}_{1,t-1} (\tilde{K}_1 - \tilde{K}_2)^\top}_{\Delta_{3,t-1}} + \tilde{K}_2 (\tilde{X}_{1,t-1} - \tilde{X}_{2,t-1}) \tilde{K}_2^\top \\ &= \sum_{i=1}^3 \sum_{s=2}^{t-1} \tilde{K}_2^{t-1+s} \Delta_{i,s} (\tilde{K}_2^{t-1+s})^\top. \end{aligned} \quad (66)$$

In order to show establish the Lipschitzness of the loss function ℓ^Φ , we will obtain a bound on each of the three error terms comprising Δ_i separately now.

Bound on Δ_1 . For the term corresponding to Δ_1 , observe that both the state \tilde{x} and the disturbance ζ are bounded vectors. Using the (κ, γ) -strong stability of the policy K_1 , we have

$$\left\| \sum_{s=2}^{t-1} \tilde{K}_2^{t-1+s} \Delta_{1,s} (\tilde{K}_2^{t-1+s})^\top \right\|_2 \leq \frac{\kappa^2 C_x \sigma_B W}{\gamma} \cdot \|K_1 - K_2\|_2. \quad (67)$$

Bound on Δ_2 . For the second term, observe that

$$\|\tilde{x}_{1,t} - \tilde{x}_{2,t}\|_2 \leq \frac{\kappa \sigma_B C_x}{\gamma} \cdot \|K_1 - K_2\|_2.$$

With this, we can bound the second term in equation (66) as

$$\left\| \sum_{s=2}^{t-1} \tilde{K}_2^{t-1+s} \Delta_{2,s} (\tilde{K}_2^{t-1+s})^\top \right\|_2 \leq \frac{\kappa^4 C_x \sigma_B W}{\gamma^2} \cdot \|K_1 - K_2\|_2. \quad (68)$$

Bound on Δ_3 . For the final term, note that $\|\tilde{X}_t\|_2 \leq C_x^2$. With this, we can bound the term corresponding to Δ_3 as

$$\left\| \sum_{s=2}^{t-1} \tilde{K}_2^{t-1+s} \Delta_{3,s} (\tilde{K}_2^{t-1+s})^\top \right\|_2 \leq \frac{\kappa^3 C_x^2 \sigma_B}{\gamma} \cdot \|K_1 - K_2\|_2. \quad (69)$$

Combining the bounds obtained in equations (67), (68) and (69), with the upper bound in equation (66) establishes the desired claim. \square