# Online Control of Unknown Time-Varying Dynamical Systems 

Edgar Minasyan<br>Princeton University, Google AI Princeton<br>minasyan@princeton.edu

Paula Gradu<br>UC Berkeley<br>pgradu@berkeley.edu

Max Simchowitz<br>MIT<br>msimchow@mit.edu

Elad Hazan<br>Princeton University, Google AI Princeton<br>ehazan@princeton.edu


#### Abstract

We study online control of time-varying linear systems with unknown dynamics in the nonstochastic control model. At a high level, we demonstrate that this setting is qualitatively harder than that of either unknown time-invariant or known timevarying dynamics, and complement our negative results with algorithmic upper bounds in regimes where sublinear regret is possible. More specifically, we study regret bounds with respect to common classes of policies: Disturbance Action (SLS), Disturbance Response (Youla), and linear feedback policies. While these three classes are essentially equivalent for LTI systems, we demonstrate that these equivalences break down for time-varying systems. We prove a lower bound that no algorithm can obtain sublinear regret with respect to the first two classes unless a certain measure of system variability also scales sublinearly in the horizon. Furthermore, we show that offline planning over the state linear feedback policies is NP-hard, suggesting hardness of the online learning problem. On the positive side, we give an efficient algorithm that attains a sublinear regret bound against the class of Disturbance Response policies up to the aforementioned system variability term. In fact, our algorithm enjoys sublinear adaptive regret bounds, which is a strictly stronger metric than standard regret and is more appropriate for time-varying systems. We sketch extensions to Disturbance Action policies and partial observation, and propose an inefficient algorithm for regret against linear state feedback policies.


## 1 Introduction

The control of linear time-invariant (LTI) dynamical systems is well-studied and understood. This includes classical methods from optimal control such as LQR and LQG, as well as robust $H_{\infty}$ control. Recent advances study regret minimization and statistical complexity for online linear control, in both stochastic and adversarial perturbation models. Despite this progress, rigorous mathematical guarantees for nonlinear control remain elusive: nonlinear control is both statistically and computationally intractable in general.
In the face of these limitations, recent research has begun to study the rich continuum of settings which lie between LTI systems and generic nonlinear ones. The hope is to provide efficient and robust algorithms to solve the most general control problems that are tractable, and at the same time, to characterize precisely at which degree of nonlinearity no further progress can be made.

This paper studies the control of linear, time-varying (LTV) dynamical systems as one such point along this continuum. This is because the first-order Taylor approximation to the dynamics of any smooth nonlinear system about a given trajectory is an LTV system. These approximations are widely popular because they allow for efficient planning, as demonstrated by the success of iLQR and iLQG methods for nonlinear receding horizon control. We study online control of discrete-time LTV systems, with dynamics and time-varying costs

$$
\begin{equation*}
x_{t+1}=A_{t} x_{t}+B_{t} u_{t}+w_{t}, \quad c_{t}\left(x_{t}, u_{t}\right):(x, u) \rightarrow \mathbb{R} \tag{1.1}
\end{equation*}
$$

Above, $x_{t}$ is the state of the system, $u_{t}$ the control input, $w_{t}$ the disturbances, and $A_{t}, B_{t}$ the system matrices. Our results extend naturally to partial-state observation, where the controller observes linear projections of the state $y_{t}=C_{t} x_{t}$. We focus on the challenges introduced when the system matrices $A_{t}, B_{t}$ and perturbations $w_{t}$ are not known to the learner in advance, and can only be determined by live interaction with the changing systems.

In this setting, we find that the overall change in system dynamics across time characterizes the difficulty of controlling the unknown LTV system. We define a measure, called system variability, which quantifies this. We show both statistical and computational lower bounds as well as algorithmic upper bounds in terms of the system variabilility. Surprisingly, system variability does not impede the complexity of control when the dynamics are known [16].

### 1.1 Contributions

We consider the recently popularized nonstochastic model of online control, and study regret bounds with respect to common classes of policies: Disturbance Action (DAC/SLS [44]), Disturbance Response (Drc/Youla [46]), and linear feedback policies. Planning over the third class of feedback policies in LTI systems admits efficient convex relaxations via the the first two parametrizations, DAC and DRC. This insight has been the cornerstone of both robust [49, 44] and online [3, 39] control.

Separation of parametrizations. For linear time-varying systems, however, we find that equivalences between linear feedback, DAC and DRC fail to hold: we show that there are cases where any one of the three parametrizations exhibits strictly better control performance than the other two.

Regret against convex parametrizations. Our first set of results pertain to DAC and DRC parametrizations, which are convex and admit efficient optimization. We demonstrate that no algorithm can obtain sublinear regret with respect to these classes when faced with unknown, LTV dynamics unless a certain measure of system variability also scales sublinearly in the horizon. This is true even under full observation, controllable dynamics, and fixed control cost. This finding is in direct contrast to recent work which shows sublinear regret is attainable over LTV system dynamics if they are known [16].

We give an efficient algorithm that attains sublinear regret against these policy classes up to an additive penalty for the aforementioned system variability term found in our lower bound. When the system variability is sufficiently small, our algorithm recovers state-of-the-art results for unknown LTI system dynamics up to logarithmic factors.

In fact, our algorithm enjoys sublinear adaptive regret [21], a strictly stronger metric than standard regret which is more appropriate for time-varying systems. We also show that the stronger notion of adaptivity called strongly adaptive regret [11] is out of reach in the partial information setting.
Regret against state feedback. Finally, we consider the class of state feedback policies, which are linear feedback with memory length one. We show that full-information optimization over state feedback policies is computationally hard. This suggests that obtaining sublinear regret relative to these policies may be computationally prohibitive, though does not entirely rule out the possibility of improper learning. However, improper learning cannot be done via the DRC or DAC relaxations in light of our policy class separation results. Finally, we include an inefficient algorithm which attains sublinear (albeit nonparametric-rate) regret against state feedback control policies.

## Paper Structure

Discussion of relevant literature and relation to our work can be found in Section 1.2. In Section 2, we formally introduce the setting of LTV nonstochastic control, the policy classes we study and our key result regarding their non-equivalence in the LTV setting (Theorem 2.1). Motivated by this
non-equivalence, the remainder of the paper is split into the study of convex policies (Section 3) and of state feedback policies (Section 4). In Section 3, we show that regret against the DAC and DrC classes cannot be sublinear unless the metric system variability (Definition 3.1) itself is sublinear (Theorem 3.1), and also propose Algorithm 2 whose adaptive regret scales at the rate of our lower bound plus a $T^{2 / 3}$ term (Theorem 3.4). On the other hand, in Section 4 we show sublinear regret against state feedback policies is technically possible (Theorem 4.1) with a computationally inefficient algorithm, but also provide a computational lower bound (Theorem 4.2) for planning which reveals significant difficulties imposed by the LTV dynamics in this scenario as well. Finally, in Section 5 we pose several future directions, concerning both questions in LTV control, as well as the extension to nonlinear control.

### 1.2 Related Work

Our study of LTV systems is motivated by the widespread practical popularity of iterative linearization for nonlinear receding horizon control; e.g., the iLQR [40], iLC [29], and iLQG [41] algorithms. Recent research has further demonstrated that near-optimal solutions to LTV approximations of dynamics confer stability guarantees onto the original nonlinear system of interest [45].

Low-Regret Control: We study algorithms which enjoy sublinear regret for online control of LTV systems; that is, whose performance tracks a given benchmark of policies up to a term which is vanishing relative to the problem horizon. [1] initiated the study of online control under the regret benchmark by introducing the online LQR problem: where a learner is faced with an unknown LTI system, fixed costs and i.i.d. Gaussian disturbances, and must attain performance relative to the LQR-optimal policy. Bounds for this setting were later improved and refined in [12, 26, 10, 38], and extended to partial-state observation in [25, 24]. Our work instead adopts the nonstochastic control setting [3], where the adversarially chosen (i.e. non-Gaussian) noise is considered to model the drift terms that arise in linearizations of nonlinear terms, and where costs may vary with time. [3] consider known system dynamics, later extended to unknown systems under both full-state [20] and partial-state observation [39, 37]. The study of nonstochastic control of known LTV dynamics was taken up in [16], with parallel work by [32] considering known LTV dynamics under stochastic noise.

Unknown LTV dynamics: Our work is the first to consider online (low-regret) control of unknown LTV systems in any model. There is, however, a rich body of classical work on adaptive control of LTV systems [28, 42]. These guarantees focus more heavily on error sensitivity and stability; they only permit dynamical recovery up to error that scales linearly in system noise, and thus guarantee only (vacuous) linear-in-horizon regret. More recent work has studied identification (but not online control) of an important LTV class called switching systems [31, 35].
Online Convex Optimization: We make extensive use of techniques from the field of online convex optimization [9, 18]. Most relevant to our work is the literature on adapting to changing environments in online learning, which starts from the works of [22,6]. The notion of adaptive regret was introduced in [21] and significantly studied since as a metric for adaptive learning in OCO [2, 47]. [11] proposed to strengthen adaptive regret and the stronger metric has been shown to imply results over dynamic regret [48].

Recent nonlinear control literature: Recent research has also studied provably guarantees in various complementary (but incomparable) models: planning regret in nonlinear control [4], adaptive nonlinear control under linearly-parameterized uncertainty [5], online model-based control with access to non-convex planning oracles [23], and control with nonlinear observation models [27, 13].

## 2 Problem Setting

We study control of a linear time-varying (LTV) system Eq. (1.1) with state $x_{t} \in \mathbb{R}^{d_{x}}$, control input $u_{t} \in \mathbb{R}^{d_{u}}$ chosen by the learner, and the external disturbance $w_{t} \in \mathbb{R}^{d_{x}}$ chosen by Nature. The system is characterized by time-varying matrices $A_{t} \in \mathbb{R}^{d_{x} \times d_{x}}, B_{t} \in \mathbb{R}^{d_{x} \times d_{u}}$. For simplicity, the initial state is $x_{1}=0$. At each time $t$, oblivious ${ }^{1}$ adversary picks the system matrices $\left(A_{t}, B_{t}\right)$, disturbances $w_{t}$ and cost functions $c_{t}: \mathbb{R}^{d_{x}} \times \mathbb{R}^{d_{u}} \rightarrow \mathbb{R}$. The dynamics $\left(A_{t}, B_{t}\right)$ are unknown to the learner: one observes only the next state $x_{t+1}$ and current cost $c_{t}(\cdot, \cdot)$ after playing control $u_{t}$.

[^0]Adaptive Regret. The goal of the learner is to minimize regret w.r.t. a policy class $\Pi$, i.e. the difference between the cumulative cost of the learner and the best policy $\pi^{\star} \in \Pi$ in hindsight. Formally, the regret of an algorithm $\mathcal{A}$ with control inputs $u_{1: T}$ and corresponding states $x_{1: T}$, over an interval $I=[r, s] \subseteq[T]$, is defined as

$$
\begin{equation*}
\operatorname{Regret}_{I}(\mathcal{A} ; \Pi)=\sum_{t \in I} c_{t}\left(x_{t}, u_{t}\right)-\inf _{\pi \in \Pi} \sum_{t \in I} c_{t}\left(x_{t}^{\pi}, u_{t}^{\pi}\right) \tag{2.1}
\end{equation*}
$$

Here $u_{t}^{\pi}, x_{t}^{\pi}$ indicate the control input and the corresponding state when following policy $\pi$. For a randomized algorithm $\mathcal{A}$, we consider the expected regret. In this work, we focus on designing control algorithms that minimize adaptive regret, i.e. guarantee a low regret relative to the best-in-hindsight policy $\pi_{I}^{\star} \in \Pi$ on any interval $I \subseteq[T]$. This performance metric of adaptive regret is more suitable for control over LTV dynamics given its agility to compete against different local optimal policies $\pi_{I}^{\star} \in \Pi$ at different times [16].
Key objects. A central object in our study is the sequence of Nature's x's $x_{1: T}^{\text {nat }}$ that arises from playing zero control input $u_{t}=0$ at each $t \in[T]$, i.e. $x_{t+1}^{\text {nat }}=A_{t} x_{t}^{\text {nat }}+w_{t}$ [39]. define the following operators for all $t$,

$$
\Phi_{t}^{[0]}=\mathbb{I}, \quad \forall h \in[1, t), \Phi_{t}^{[h]}=\prod_{k=t}^{t-h+1} A_{k}, \quad \forall i \in[0, t), G_{t}^{[i]}=\Phi_{t}^{[i]} B_{t-i}
$$

where the matrix product $\prod_{s}^{r}$ with $s \geq r$ is taken in the indicated order $k=s, \ldots, r$. the following identities give an alternative representation for the Nature's x's $x_{t}^{\text {nat }}$ and state $x_{t}$ with control input $u_{t}$ in terms of the Markov operator at time $t, G_{t}=\left[G_{t}^{[i]}\right]_{i \geq 0}$ :

$$
x_{t+1}^{\mathrm{nat}}=\sum_{i=0}^{t-1} \Phi_{t}^{[i]} w_{t-i}, \quad x_{t+1}=x_{t+1}^{\mathrm{nat}}+\sum_{i=0}^{t-1} G_{t}^{[i]} u_{t-i}
$$

These operators and the alternative representation capture the dynamics by decoupling the disturbance and the control action effects.

Assumptions. We make the three basic assumptions: we require from (i) the disturbances to not blow up the system with no control input, (ii) the system to have decaying effect over time, and (iii) the costs to be well-behaved and admit efficient optimization. Formally, these assumptions are:
Assumption 1. For all $t \in[T]$, assume $\left\|x_{t}^{\text {nat }}\right\| \leq R_{\text {nat }}$.
Assumption 2. Assume there exist $R_{G} \geq 1$ and $\rho \in(0,1)$ s.t. for any $h \geq 0$ and for all $t \in[T]$

$$
\sum_{i \geq h}\left\|G_{t}^{[i]}\right\|_{\mathrm{op}} \leq R_{G} \cdot \rho^{h}:=\psi(h)
$$

Assumption 3. Assume the costs $c_{t}: \mathbb{R}^{d_{x}} \times \mathbb{R}^{d_{u}} \rightarrow \mathbb{R}$ are general convex functions that satisfy the conditions $0 \leq c_{t}(x, u) \leq L \max \left\{1,\|x\|^{2}+\|u\|^{2}\right\}$, and $\left\|\nabla c_{t}(x, u)\right\| \leq L \max \{1,\|x\|+\|u\|\}$ for some constant $L>0$, where $\nabla$ denotes any subgradient [7].

The conditions in Assumption 3 allow for functions whose values and gradient grow as quickly as quadratics (e.g. the costs in LQR ), and the $\max \{1, \cdot\}$ term ensures the inclusion of standard bounded and Lipschitz functions as well. Assumptions 1 and 2 arise from the assumption our LTV system is open-loop stable; Appendix A. 2 extends to the case where a nominal stabilizing controller is known, as in prior work $[3,39]$. While these two assumptions may seem unnatural at first, they can be derived from the basic conditions of disturbance norm bound and sequential stability.
Lemma 2.1. Suppose that there exist $C_{1} \geq 1, \rho_{1} \in(0,1)$ such that $\left\|\Phi_{t}^{[h]}\right\|_{\mathrm{op}} \leq C_{1} \rho_{1}^{h}$ for any $h \geq 0$ and all $t \in[T]$, and suppose that $\max _{t}\left\|w_{t}\right\| \leq R_{w}$. Then, Assumption 1 holds with $R_{\mathrm{nat}}=\frac{C_{1}}{1-\rho_{1}} R_{w}$, and Assumption 2 holds with $\rho=\rho_{1}$ and $R_{G}=\max \left\{1, \max _{t}\left\|B_{t}\right\|_{\mathrm{op}} \cdot \frac{C_{1}}{1-\rho_{1}}\right\}$.

Note that Assumption 2 implies that $\left\|G_{t}\right\|_{\ell_{1}, \mathrm{op}}=\sum_{i \geq 0}\left\|G_{t}^{[i]}\right\|_{\mathrm{op}} \leq R_{G}$. It also suggests that for a sufficiently large $h$ the effect of iterations before $t-h$ are negligible at round $t$. This prompts introducing a truncated Markov operator: denote $\bar{G}_{t}^{h}=\left[G_{t}^{[i]}\right]_{i<h}$ to be the $h$-truncation of the true Markov operator $G_{t}$. It follows that their difference is $\left\|\bar{G}_{t}^{h}-G_{t}\right\|_{\ell_{1}, \mathrm{op}}=\sum_{i>h}\left\|G_{t}^{[i]}\right\|_{\mathrm{op}} \leq \psi(h)$ negligible in operator norm for a sufficiently large $h$. Define the bounded set of $\bar{h}$-truncated Markov operators to be $\mathcal{G}\left(h, R_{G}\right)=\left\{G=\left[G^{[i]}\right]_{0 \leq i<h}:\|G\|_{\ell_{1}, \text { op }} \leq R_{G}\right\}$ with $\bar{G}_{t}^{h} \in \mathcal{G}\left(h, R_{G}\right)$ for all $t$.

### 2.1 Benchmarks and Policy Classes

The performance of an algorithm, measured by Eq. (2.1), directly depends on the policy class $\Pi$ that is chosen as a benchmark to compete against. In this work, we consider the following three policy classes: DRC, DAC, and linear feedback. DRC parameterizes control inputs in terms of Nature's x's $x_{t}^{\text {nat }}$, DAC does so in terms of the disturbances $w_{t}$ and linear feedback in terms of the states $x_{t}$. We express all three in terms of a length- $m$ parameter $M=\left[M^{[i]}\right]_{i<m}$ in a bounded ball $\mathcal{M}\left(m, R_{M}\right)$ :

$$
\mathcal{M}\left(m, R_{M}\right)=\left\{\left(M^{[0]}, \ldots, M^{[m-1]}\right): \sum_{i=0}^{m-1}\left\|M^{[i]}\right\|_{\mathrm{op}} \leq R_{M}\right\}
$$

Definition 2.1 (DRC policy class). A DRC control policy $\pi_{\mathrm{drc}}^{M}$ of length $m$ is given by $u_{t}^{M}=$ $\sum_{i=0}^{m-1} M^{[i]} x_{t-i}^{\text {nat }}$ where $M=\left[M^{[i]}\right]_{i<m}$ is the parameter of the policy. Define the bounded DRC policy class as $\Pi_{\mathrm{drc}}\left(m, R_{M}\right)=\left\{\pi_{\mathrm{drc}}^{M}: M \in \mathcal{M}\left(m, R_{M}\right)\right\}$.
Definition 2.2 (DAC policy class). A DAC control policy $\pi_{\text {dac }}^{M}$ of length $m$ is given by $u_{t}^{M}=$ $\sum_{i=0}^{m-1} M^{[i]} w_{t-i}$ where $M=\left[M^{[i]}\right]_{i<m}$ is the parameter of the policy. Define the bounded DAC policy class as $\Pi_{\mathrm{dac}}\left(m, R_{M}\right)=\left\{\pi_{\mathrm{dac}}^{M}: M \in \mathcal{M}\left(m, R_{M}\right)\right\}$.
Definition 2.3 (Feedback policy class). A feedback control policy $\pi_{\text {feed }}^{M}$ of length $m$ is given by $u_{t}^{M}=\sum_{i=0}^{m-1} M^{[i]} x_{t-i}$ where $M=\left[M^{[i]}\right]_{i<m}$ is the parameter of the policy. Define the bounded feedback policy class as $\Pi_{\text {feed }}\left(m, R_{M}\right)=\left\{\pi_{\text {feed }}^{M}: M \in \mathcal{M}\left(m, R_{M}\right)\right\}$. In the special case of memory $m=1$, denote the state feedback policy class as $\Pi_{\text {state }}=\Pi_{\text {feed }}(m=1)$.

Convexity. Both the DrC and DAC policy classes are convex parametrizations: a policy $\pi \in$ $\Pi_{\mathrm{drc}} \cup \Pi_{\mathrm{dac}}$ outputs controls $u_{t}$ that are linear in the policy-independent sequences $x_{1: T}^{\text {nat }}$ and $w_{1: T}$, and thus the mapping from parameter $M$ to resulting states and inputs (resp. costs) is affine (resp. convex). Hence, we refer to these as the convex classes. In contrast, feedback policies select inputs based on policy-dependent states, and are therefore non-convex [15].
We drop the arguments $m, R_{M}$ when they are clear from the context. The state feedback policies $\Pi_{\text {state }}$ encompass the $\mathcal{H}_{2}$ and $\mathcal{H}_{\infty}$ optimal control laws under full observation. For LTI systems, DRC and DAC are equivalent $[46,44]$ and approximate all linear feedback policies to arbitrarily high precision [3, 39]. However, we show that these relationships between the classes break down for LTV systems: there exist scenarios where any one of the three classes strictly outperforms the other two.
Theorem 2.1 (Informal). For each class $\Pi$ in $\left\{\Pi_{\mathrm{drc}}, \Pi_{\mathrm{dac}}, \Pi_{\mathrm{feed}}\right\}$ there exists a sequence of wellbehaved $\left(A_{t}, B_{t}, w_{t}, c_{t}\right)$ such that a policy $\pi^{\star} \in \Pi$ suffers 0 cumulative cost, but each of the other two classes $\Pi^{\prime} \in\left\{\Pi_{\mathrm{drc}}, \Pi_{\mathrm{dac}}, \Pi_{\text {feed }}\right\} \backslash \Pi$ suffers $\Omega(T)$ cost on all their constituent policies $\pi \in \Pi^{\prime}$.

The formal theorem that includes the definition of a well-behaved instance sequence and the final statement dependence on $m, R_{M}$ along with its proof can be found in Appendix F.1.
Notation. The norm $\|\cdot\|$ refers to Euclidean norm unless otherwise stated, $[n]$ is used as a shorthand for $[1, n], T$ is used as a subscript shorthand for $[T]$. The asymptotic notation $\mathcal{O}(\cdot), \Omega(\cdot)$ suppress all terms independent of $T, \widetilde{\mathcal{O}}(\cdot)$ additionally suppresses terms logarithmic in $T$. We define $\widetilde{\mathcal{O}}^{\star}(\cdot)$ to suppress absolute constants, polynomials in $R_{\text {nat }}, R_{G}, R_{M}$ and logarithms in $T$.

## 3 Online Control over Convex Policies

This section considers online control of unknown LTV systems so as to compete with the convex DrC and DAC policy classes. The fundamental quantity which appears throughout our results is the system variability, which measures the variation of the time-varying Markov operators $G_{t}$ over intervals $I$.
Definition 3.1. Define the system variability of an LTV dynamical system with Markov operators $\mathbf{G}=G_{1: T}$ over a contiguous interval $I \subseteq[T]$ to be

$$
\operatorname{Var}_{I}(\mathbf{G})=\min _{G} \frac{1}{|I|} \sum_{t \in I}\left\|G-G_{t}\right\|_{\ell_{2}, F}^{2}=\frac{1}{|I|} \sum_{t \in I}\left\|G_{I}-G_{t}\right\|_{\ell_{2}, F}^{2},
$$

where $\|\cdot\|_{\ell_{2}, F}$ indicates the $\ell_{2}$ norm of the fully vectorized operator and $G_{I}=|I|^{-1} \sum_{t \in I} G_{t}$ is the


Our results in this section for both upper and lower bounds focus on expected regret: high probability results are possible as well with more technical effort using standard techniques.

### 3.1 A Linear Regret Lower Bound

Our first contribution is a negative one: that the regret against the class of either DAC or DRC policies cannot scale sublinearly in the time horizon. Informally, our result shows that the regret against these classes scales as $T \sigma$, where $\sigma^{2}$ is the system variability.

More precisely, for any $\sigma^{2} \in(0,1 / 8]$, we construct a distribution $\mathcal{D}_{\sigma}$ over sequences $\left(A_{t}, B_{t}, c_{t}, w_{t}\right)$, formally specified in Appendix F.2. Here, we list the essential properties of $\mathcal{D}_{\sigma}:(i) A_{t} \equiv 0$, (ii) $c_{t} \equiv c$ is a fixed cost satisfying Assumption 3 with $L \leq 4$, (iii) the matrices $\left(B_{t}\right)$ are i.i.d., with $\left\|B_{t}\right\|_{\mathrm{op}} \leq 2$ almost surely, and $\mathbb{E}\left[\left\|B_{t}-\mathbb{E}\left[B_{t}\right]\right\|_{\mathrm{F}}^{2}\right]=\sigma^{2}$, and $(i v)\left\|w_{t}\right\| \leq 4$ for all $t$. These conditions imply that Assumptions 1 and 2 hold for $R_{G}=2, \rho=0, R_{\text {nat }}=4$. The condition $A_{t} \equiv 0$ implies that $x_{t}^{\text {nat }}=w_{t}$ for all $t$, so the classes DRC and DAC are equivalent and the lower bound holds over both. Moreover, by Jensen's inequality, this construction ensures that

$$
\begin{aligned}
\mathbb{E}\left[\operatorname{Var}_{I}(\mathbf{G})\right] & =|I|^{-1} \mathbb{E}\left[\min _{G} \sum_{t \in I}\left\|G-G_{t}\right\|_{\ell_{2}, F}^{2}\right] \\
& =|I|^{-1} \mathbb{E}\left[\min _{B} \sum_{t \in I}\left\|B-B_{t}\right\|_{\mathrm{F}}^{2}\right] \leq \mathbb{E}\left[\left\|B_{t}-\mathbb{E}\left[B_{t}\right]\right\|_{\mathrm{F}}^{2}\right]=\sigma^{2} .
\end{aligned}
$$

In particular, $\mathbb{E}\left[\operatorname{Var}_{T}(\mathbf{G})\right] \leq \sigma^{2}$. For the described construction, we show the following lower bound:

Theorem 3.1. Let $C$ be a universal, positive constant. For any $\sigma \in(0,1 / 8]$ and any online control algorithm $\mathcal{A}$, there exists a DRC policy $\pi^{\star} \in \Pi_{\mathrm{drc}}(1,1)$ s.t. expected regret incurred by $\mathcal{A}$ under the distribution $\mathcal{D}_{\sigma}$ and cost $c(x, u)$ is at least

$$
\mathbb{E}_{\mathcal{D}_{\sigma}, \mathcal{A}}\left[\operatorname{Regret}_{T}\left(\mathcal{A} ;\left\{\pi^{\star}\right\}\right)\right] \geq C \cdot T \sigma \geq C \cdot T \cdot \sqrt{\mathbb{E}\left[\operatorname{Var}_{T}(\mathbf{G})\right]}
$$

A full construction and proof of Theorem 3.1 is given in Appendix F.2. In particular, for $\sigma=1 / 8$, we find that no algorithm can attain less than $\Omega(T)$ expected regret; a stark distinction from either unknown LTI [39, 20] or known LTV [16] systems.

### 3.2 Estimation of Time-Varying Vector Sequences

To devise an algorithmic upper bound that complements the result in Theorem 3.1, we first consider the setting of online prediction under a partial information model. This setting captures the system identification phase of LTV system control and is used to derive the final control guarantees. Formally, consider the following repeated game between a learner and an oblivious adversary: at each round $t \in[T]$, the adversary picks a target vector $\mathbf{z}_{t}^{\star} \in \mathcal{K}$ from a convex decision set $\mathcal{K}$ contained in a 0 -centered ball of radius $R_{z}$; simultaneously, the learner selects an estimate $\hat{\mathbf{z}}_{t} \in \mathcal{K}$ and suffers quadratic loss $\ell_{t}\left(\hat{\mathbf{z}}_{t}\right)=\left\|\hat{\mathbf{z}}_{t}-\mathbf{z}_{t}^{\star}\right\|^{2}$. The only feedback the learner has access to is via the following noisy and costly oracle.
Oracle 1 (Noisy Costly Oracle). At each time $t \in[T]$, the learner selects a decision $b_{t} \in\{0,1\}$ indicating whether a query is sent to the oracle. If $b_{t}=1$, the learner receives an unbiased estimate $\tilde{\mathbf{z}}_{t}$ as response such that $\left\|\tilde{\mathbf{z}}_{t}\right\| \leq \tilde{R}_{z}$ and $\mathbb{E}\left[\tilde{\mathbf{z}}_{t} \mid \mathcal{F}_{t}, b_{t}=1\right]=\mathbf{z}_{t}^{\star}$. The filtration $\mathcal{F}_{t}$ is the sigma algebra generated by $\tilde{\mathbf{z}}_{1: t-1}, b_{1: t-1}$ and the choices of the oblivious adversary $\mathbf{z}_{1: T}^{\star}$. A completed query results in a unit cost $\lambda>0$ for the learner.

The performance metric of an online prediction algorithm $\mathcal{A}_{\text {pred }}$ is expected quadratic loss regret along with the extra cumulative oracle query cost. It is defined over each interval $I=[r, s] \subseteq[T]$ as

$$
\begin{equation*}
\operatorname{Regret}_{I}\left(\mathcal{A}_{\text {pred }} ; \lambda\right)=\mathbb{E}_{\mathcal{F}_{1: T}}\left[\sum_{t \in I} \ell_{t}\left(\hat{\mathbf{z}}_{t}\right)-\min _{\mathbf{z} \in \mathcal{K}} \sum_{t \in I} \ell_{t}(\mathbf{z})+\lambda \sum_{t \in I} b_{t}\right] \tag{3.1}
\end{equation*}
$$

To attain adaptive regret, i.e. bound Eq. (3.1) for each interval $I$, we propose Algorithm 1 constructed as follows. First, suppose we wanted non-adaptive (i.e. just $I=[T]$ ) guarantees. In this special case, we propose to sample $b_{t} \sim \operatorname{Bernoulli}(p)$ for an appropriate parameter $p \in(0,1)$, and perform a gradient descent update on the importance-weighted square $\operatorname{loss} \tilde{\ell}_{t}(\mathbf{z})=\frac{1}{2 p} \mathbb{I}\left\{b_{t}=1\right\}\left\|\mathbf{z}-\tilde{\mathbf{z}}_{t}\right\|^{2}$. To

```
Algorithm 1 Adaptive Estimation Algorithm (ADA-PRED)
    Input: parameter \(p\), decision set \(\mathcal{K}\)
    Initialize: \(\hat{\mathbf{z}}_{1}^{(1)} \in \mathcal{K}\), working dictionary \(\mathcal{S}_{1}=\left\{\left(1: \hat{\mathbf{z}}_{1}^{(1)}\right)\right\}, q_{1}^{(1)}=1\), parameter \(\alpha=\frac{p}{\left(R_{z}+\tilde{R}_{z}\right)^{2}}\)
    for \(t=1, \ldots, T\) do
        Play iterate \(\hat{\mathbf{z}}_{t}=\sum_{\left(i, \hat{\mathbf{z}}_{t}^{(i)}\right) \in \mathcal{S}_{t}} q_{t}^{(i)} \hat{\mathbf{z}}_{t}^{(i)}\)
        Draw/Receive \(b_{t} \sim \operatorname{Bernoulli}(p)\)
        if \(b_{t}=1\) then
            Request estimate \(\tilde{\mathbf{z}}_{t}\) from Oracle 1
            Let \(\tilde{\ell}_{t}(\mathbf{z})=\frac{1}{2 p}\left\|\mathbf{z}-\tilde{\mathbf{z}}_{t}\right\|^{2}\) and \(\tilde{\nabla}_{t}=\frac{1}{p}\left(\mathbf{z}_{t}-\tilde{\mathbf{z}}\right)\)
        else
            Let \(\tilde{\mathbf{z}}_{t} \leftarrow \emptyset\) and \(\tilde{\ell}_{t}(\mathbf{z})=0\) and \(\tilde{\nabla}_{t}=0\)
        Update predictions \(\hat{\mathbf{z}}_{t+1}^{(i)} \leftarrow \operatorname{Proj}_{\mathcal{K}}\left(\hat{\mathbf{z}}_{t}^{(i)}-\eta_{t}^{(i)} \tilde{\nabla}_{t}\right)\) for all \(\left(i, \hat{\mathbf{z}}_{t}^{(i)}\right) \in S_{t}\)
        Form new dictionary \(\tilde{\mathcal{S}}_{t+1}=\left(i, \hat{\mathbf{z}}_{t+1}^{(i)}\right)_{i \in \operatorname{keys}\left(\mathcal{S}_{t}\right)}\)
        Construct proxy new weights \(\bar{q}_{t+1}^{(i)}=\frac{t}{t+1} \cdot \frac{q_{t}^{(i)} e^{-\alpha \tilde{\ell}_{t}\left(\hat{\mathbf{z}}_{t}^{(i)}\right)}}{\sum_{j \in \operatorname{keys}\left(\mathcal{S}_{t}\right)} q_{t}^{(j)} e^{-\alpha \tilde{\ell}_{t}\left(\hat{\mathbf{z}}_{t}^{(j)}\right)}}\) for all \(i \in \operatorname{keys}\left(\mathcal{S}_{t}\right)\)
        Add new instance \(\tilde{\mathcal{S}}_{t+1} \leftarrow \tilde{\mathcal{S}}_{t+1} \cup\left(t+1, \hat{\mathbf{z}}_{t+1}^{(t+1)}\right)\) for arbitrary \(\hat{\mathbf{z}}_{t+1}^{(t+1)} \in \mathcal{K}\) with \(\bar{q}_{t+1}^{(t+1)}=\frac{1}{t+1}\)
        Prune \(\tilde{\mathcal{S}}_{t+1}\) to form \(\mathcal{S}_{t+1}\) (see Appendix C.1)
        Normalize \(q_{t+1}^{(i)}=\frac{\bar{q}_{t+1}^{(i)}}{\sum_{j \in \operatorname{keys}\left(\mathcal{S}_{t+1}\right)} \bar{q}_{t+1}^{(j)}}\)
```

extend this method to enjoy adaptive regret guarantees, we adopt the approach of [21]: the core idea in this approach is to initiate an instance of the base method at each round $t$ and use a weighted average of the instance predictions as the final prediction (Line 4). The instance weights are multiplicatively updated according to their performance (Line 13). To ensure computational efficiency, the algorithm only updates instances from a working dictionary $\mathcal{S}_{t}$ (Line 11). These dictionaries are pruned each round (Line 15) such that $\left|\mathcal{S}_{t}\right|=O(\log T)$ (see Appendix C. 1 for details).

Theorem 3.2. Given access to queries from Oracle 1 and with stepsizes $\eta_{t}^{(i)}=\frac{1}{t-i+1}$, Algorithm 1 enjoys the following adaptive regret guarantee: for all $I=[r, s] \subseteq[T]$,

$$
\begin{equation*}
\operatorname{Regret}_{I}(\operatorname{ADA}-\operatorname{PRED} ; \lambda) \leq \frac{2\left(R_{z}+\tilde{R}_{z}\right)^{2}(1+\log s \cdot \log |I|)}{p}+\lambda p|I| \tag{3.2}
\end{equation*}
$$

When $I=[T]$, the optimal choice of parameter $p=\log T / \sqrt{\lambda T}$ yields regret scaling roughly as $\sqrt{\lambda T \log ^{2} T}$. Unfortunatelly, this gives regret scaling as $\sqrt{T}$ for all interval sizes: to attain $\sqrt{|I|}$ regret on interval $I$, the optimal choice of $p$ would yield $\sim T / \sqrt{|I|}$ regret on $[T]$, which is considerably worse for small $|I|$. One may ask if there exists a strongly adaptive algorithm which adapts $p$ as well, so as to enjoy regret polynomial in $|I|$ for all intervals $I$ simultaneously [11]. The following result shows this is not possible:
Theorem 3.3 (Informal). For all $\gamma>0$ and $\lambda>0$, there exists no online algorithm $\mathcal{A}$ with feedback access to Oracle 1 that enjoys strongly adaptive regret of $\operatorname{Regret}_{I}(\mathcal{A} ; \lambda)=\tilde{O}\left(|I|^{1-\gamma}\right)$.

Hence, in a sense, Algorithm 1 is as adaptive as one could hope for: it ensures a regret bound for all intervals $I$, but not a strongly adaptive one. The lower bound construction, formal statement, and proof of Theorem 3.3 are given in Appendix C.2.

### 3.3 Adaptive Regret for Control of Unknown Time-Varying Dynamics

We now apply our adaptive estimation algorithm (Algorithm 1) to the online control problem. Our proposed algorithm, Algorithm 2, takes in two sub-routines: a prediction algorithm $\mathcal{A}_{\text {pred }}$ which enjoys low prediction regret in the sense of the previous section, and a control algorithm $\mathcal{A}_{\text {ctrl }}$ which has low regret for control of known systems. Our master algorithm trades off between the two in epochs $\tau=1,2, \ldots$ of length $h$ : each epoch corresponds to one step of $\mathcal{A}_{\text {pred }}$ indexed by $[\tau]$.

At each epoch, the algorithm receives Markov operator estimates from $\mathcal{A}_{\text {pred }}$ (Line 3) and makes a binary decision $b_{[\tau]} \sim \operatorname{Bernoulli}(p)$. If $b_{[\tau]}=1$, then it explores using i.i.d. Rademacher inputs (Line 7), and sends the resulting estimator to $\mathcal{A}_{\text {pred }}$ (Line 14). This corresponds to one query from Oracle 1. Otherwise, it selects inputs in line with $\mathcal{A}_{\text {ctrl }}$ (Line 9), and does not give a query to $\mathcal{A}_{\text {pred }}$ (Line 16). Regardless of exploration decision, the algorithm feeds costs, current estimates of the Markov operator and Nature's x's based on the Markov operator estimates to $\mathcal{A}_{\text {ctrl }}$ (Lines 10-12), which it uses to select inputs and update its parameter.
The prediction algorithm $\mathcal{A}_{\text {pred }}$ is taken to be ADA-PrED with the decision set $\mathcal{K}=\mathcal{G}\left(h, R_{G}\right)$ : the projection operation onto it and the ball $\mathbb{B}_{R_{\text {nat }}}$ is done by clipping when the norm of the argument exceeds the indicated bound. The control algorithm $\mathcal{A}_{\text {ctrl }}$ is taken to be DRC-OgD [39] for known systems.
Theorem 3.4. For $h=\frac{\log T}{\log \rho^{-1}}, p=T^{-1 / 3}$ and $m \leq \sqrt{T}$, on any contiguous interval $I \subseteq[T]$,
Algorithm 2 enjoys the following adaptive regret guarantee:

$$
\begin{equation*}
\mathbb{E}\left[\operatorname{Regret}_{I}(\text { ADA-CTRL }) ; \Pi_{\mathrm{drc}}\left(m, R_{M}\right)\right] \leq \widetilde{\mathcal{O}}^{\star}\left(\operatorname{Lm}\left(|I| \sqrt{\mathbb{E}\left[\operatorname{Var}_{I}(\mathbf{G})\right]}+d_{u} T^{2 / 3}\right)\right) \tag{3.3}
\end{equation*}
$$

Proof Sketch. The analysis proceeds by reducing the regret incurred to that over a known system, accounting for: 1) the additional exploration penalty $(O(p|I|))$, 2) the system misspecification induced error $\left(\sim \sum_{t \in I}\left\|\hat{G}_{t}-\bar{G}_{t}^{h}\right\|_{\ell_{1}, \mathrm{op}}\right)$, and 3) truncation errors $(\sim \psi(h)|I|)$. Via straightforward computations, the system misspecification error can be expressed in terms of the result in Theorem 3.2, ultimately leading to an error contribution $\sim|I| \sqrt{\mathbb{E}\left[\operatorname{Var}_{I}(\mathbf{G})\right]}+p^{-1 / 2}|I|^{1 / 2}$. The analysis is finalized by noting that the chosen $p$ ideally balances $p|I|$ and $p^{-1 / 2}|I|^{1 / 2}$, and that the chosen $h$ ensures that the truncation error is negligible. The full proof can be found in Appendix D.

The adaptive regret bound in Eq. (3.3) has two notable terms. Note that the first term $|I| \sqrt{\mathbb{E}\left[\operatorname{Var}_{I}(\mathbf{G})\right]}$ for $I=[T]$ matches the regret lower bound in Theorem 3.1. Furthermore, our algorithm is adaptive in this term for all intervals $I$. On the other hand, for unknown LTI systems with $\operatorname{Var}_{I}(\mathbf{G})=0$, the algorithm recovers the state-of-the-art bound of $T^{2 / 3}$ [20]. However, the $T^{2 / 3}$ term is not adaptive to the intervals $I$ consistent with the lower bound against strongly adaptive algorithms in Theorem 3.3.

```
Algorithm 2 DRC-OGD with Adaptive Exploration (ADA-CTRL)
    Input: \(p, h, \mathcal{A}_{\text {pred }} \leftarrow \operatorname{ADA}-\operatorname{PrED}\left(p, \hat{G}_{0}, \mathcal{G}\left(h, R_{G}\right)\right), \mathcal{A}_{\text {ctrl }} \leftarrow \operatorname{DRC-OGD}\left(m, R_{\mathcal{M}}\right)\)
    for \(\tau=1, \ldots, T / h\) do \(\quad \triangleright\) let \(t_{\tau}=(\tau-1) h+1\)
        Set \(\hat{G}_{t_{\tau}}, \hat{G}_{t_{\tau}+1}, \ldots, \hat{G}_{t_{\tau}+h-1}\) equal to \(\tau\)-th iterate \(\hat{G}_{[\tau]}\) from \(\mathcal{A}_{\text {pred }}\)
        Draw \(b_{[\tau]} \sim \operatorname{Bernoulli}(p)\)
        for \(t=t_{\tau}, \ldots, t_{\tau}+h-1\) do
                if \(b_{[\tau]}=1\) then
                    Play control \(u_{t} \sim\{ \pm 1\}^{d_{u}}\)
                else
                    Play control \(u_{t}\) according to the \(t\)-th input chosen by \(\mathcal{A}_{\text {ctrl }}\)
                Suffer cost \(c_{t}\left(x_{t}, u_{t}\right)\), observe new state \(x_{t+1}\)
                Extract \(\hat{x}_{t+1}^{\text {nat }}=\operatorname{Proj}_{\mathbb{B}_{R_{\text {nat }}}}\left(x_{t+1}-\sum_{i=0}^{h-1} \hat{G}_{t}^{[i]} u_{t-i}\right)\)
                Feed cost, Markov operator and Nature's x estimates \(\left(c_{t}, \hat{G}_{t}, \hat{x}_{t+1}^{\mathrm{nat}}\right)\) to \(\mathcal{A}_{\text {ctrl }}\).
        if \(b_{[\tau]}=1\) then
            Feed \(\left(b_{[\tau]}, \tilde{G}_{[\tau]}\right)\) to \(\mathcal{A}_{\text {pred }}\), where \(\tilde{G}_{[\tau]}^{[i]}=x_{t_{\tau}+h} u_{t_{\tau}+h-i}^{\top}, i=0,1, \ldots, h-1\).
        else
            Feed \(\left(b_{[\tau]}, \tilde{G}_{[\tau]}\right)\) to \(\mathcal{A}_{\text {pred }}\), where \(\tilde{G}_{[\tau]} \leftarrow \emptyset\).
```


## 4 Online Control over State Feedback

Given the impossibility of sublinear regret against DRC/DAC without further restrictions on system variability, this section studies whether sublinear regret is possible against the class of linear feedback
policies. For simplicity, we focus on the state feedback policies $u_{t}=K x_{t}$, that is, linear feedback policies with memory $m=1$ (Definition 2.3). We note that state feedback policies were the class which motivated the relaxation to DAC policies in the first study of nonstochastic control [3].

We present two results, rather qualitative in nature. First, we show that obtaining sublinear regret is, in the most literal sense, possible. The following result considers regret relative to a class $\mathcal{K}$ of static feedback controllers which satisfy the restrictive assumption that each $K \in \mathcal{K}$ stabilizes the time varying dynamics $\left(A_{t}, B_{t}\right)$; see Appendix E for the formal algorithm, assumptions, and guarantees. We measure the regret against this class $\mathcal{K}$ :

$$
\operatorname{Regret}_{T}(\mathcal{K}):=\sum_{t=1}^{T} c_{t}\left(x_{t}, u_{t}\right)-\inf _{K \in \mathcal{K}} \sum_{t=1}^{T} c_{t}\left(x_{t}^{K}, u_{t}^{K}\right)
$$

where $\left(x_{t}^{K}, u_{t}^{K}\right)$ are the iterates arising under the control law $u_{t}=K x_{t}$.
Theorem 4.1 (Sublinear regret against state-feedback). Under a suitable stabilization assumption, there exists a computationally inefficient control algorithm which attains sublinear expected regret:

$$
\mathbb{E}\left[\operatorname{Regret}_{T}(\mathcal{K})\right] \leq e^{\Omega\left(d_{x} d_{u} / 2\right)} \cdot T^{1-\frac{1}{2\left(d_{x} d_{u}+3\right)}}
$$

Above, $\Omega(\cdot)$ suppresses a universal constant and exponent base, both of which are made explicit in a formal theorem statement in Appendix E. The bound follows by running the EXP3 bandit algorithm on a discretization of the set $\mathcal{K}$ (high probability regret can be obtained by instead using Exp3.P [8]). The guarantee in Theorem 4.1 is neither practical nor sharp; its sole purpose is to confirm the possibility of sublinear regret. Due to the bandit reduction and exponential size of the cover of $\mathcal{K} \subset \mathbb{R}^{d_{u} \times d_{x}}$, the algorithm is computationally inefficient and suffers a nonparametric rate of regret [33]: $\epsilon$-regret requires $T=\epsilon^{-\Omega(\text { dimension })}$.

One may wonder if one can do much better than this naive bandit reduction. For example, is there structure that can be leveraged? For LTV systems, we show that there is strong evidence to suggest that, at least from a computational standpoint, attaining polynomial regret (e.g. $T^{1-\alpha}$ for $\alpha>0$ independent of dimension) is computationally prohibitive.
Theorem 4.2. There exists a reduction from MAX-3SAT on $m$-clauses and $n$-literals to the problem of finding a state-feedback controler $K$ which is within a small constant factor of optimal for the cost $\sum_{t=1}^{T} c_{t}\left(x_{t}^{K}, x_{t}^{K}\right)$ on a sequence of sequentially stable LTV systems and convex costs $\left(A_{t}, B_{t}, c_{t}\right)$ with no disturbance ( $w_{t} \equiv 0$ ), with state dimension $n+1$, input dimension 2 , and horizon $T=\Theta(m n)$. Therefore, unless $\mathrm{P}=\mathrm{NP}$, the latter cannot be solved in time polynomial in $n$ [17].

A more precise statement, construction, and proof are given in Appendix F.4. Theorem 4.2 demonstrates that solving the offline optimization problem over state feedback controllers $K$ to within constant precision is NP-Hard. In particular, this means that any sublinear regret algorithm which is proper and convergent, in the sense that $u_{t}=K_{t} x_{t}$ for some sequence $K_{t}$ converges to a limit as $T \rightarrow \infty$, must be computationally inefficient. This is true even if the costs and dynamics are known in advance. Our result suggests it is computationally hard to obtain sublinear regret, but it does not rigorously imply it. For example, there may be more clever convex relaxations (other than DRC and DAC, which provably cannot work) that yield efficient and sublinear regret. Secondly, this lower bound does not rule out the possibility of an computationally inefficient algorithm which nevertheless attains polynomial regret.

## 5 Discussion and Future Work

This paper provided guarantees for and studied the limitations of sublinear additive regret in online control of an unknown, linear time-varying (LTV) dynamical system.
Our setting was motivated by the fact that the first-order Taylor approximation (Jacobian linearization) of smooth, nonlinear systems about any smooth trajectory is LTV. One would therefore hope that low-regret guarantees against LTV systems may imply convergence to first-order stationary points of general nonlinear control objectives [34], which in turn may enjoy stability properties [45]. Making this connection rigorous poses several challenges. Among them, one would need to extend our low-regret guarantees against oblivious adversaries to hold against adaptive adversaries, the latter
modeling how nonlinear system dynamics evolve in response to the learner's control inputs. This may require parting from our current analysis, which leverages the independence between exploratory inputs and changes in system dynamics.

Because we show that linear-in- $T$ regret is unavoidable for changing systems with large system variability, at least for the main convex policy parametrizations, it would be interesting to study our online setting under other measures of performance. In particular, the competive ratio, or the ratio of total algorithm cost to optimal cost in hindsight (as opposed to the difference between the two measured by regret) may yield a complementary set of tradeoffs, or lead to new and exciting principles for adaptive controller design. Does system variability play the same deciding roles in competive analysis as it does in regret? And, in either competitive or regret analyses, what is the correct measure of system variability (e.g. variability in which norm/geometry, or of which system parameters) which best captures sensitivity of online cost to system changes?

## Acknowledgments

Elad Hazan and Edgar Minasyan have been supported in part by NSF grant \#1704860. This work was done in part when Paula Gradu was at Google AI Princeton and Princeton University. Max Simchowitz is generously supported by an Open Philanthropy AI fellowship.

## References

[1] Yasin Abbasi-Yadkori and Csaba Szepesvári. Regret bounds for the adaptive control of linear quadratic systems. In Proceedings of the 24th Annual Conference on Learning Theory, pages 1-26, 2011.
[2] Dmitry Adamskiy, Wouter M Koolen, Alexey Chernov, and Vladimir Vovk. A closer look at adaptive regret. The Journal of Machine Learning Research, 17(1):706-726, 2016.
[3] Naman Agarwal, Brian Bullins, Elad Hazan, Sham Kakade, and Karan Singh. Online control with adversarial disturbances. In International Conference on Machine Learning, pages 111119, 2019.
[4] Naman Agarwal, Elad Hazan, Anirudha Majumdar, and Karan Singh. A regret minimization approach to iterative learning control. arXiv preprint arXiv:2102.13478, 2021.
[5] Nicholas M. Boffi, Stephen Tu, and Jean-Jacques E. Slotine. Regret bounds for adaptive nonlinear control, 2020.
[6] Olivier Bousquet and Manfred K Warmuth. Tracking a small set of experts by mixing past posteriors. Journal of Machine Learning Research, 3(Nov):363-396, 2002.
[7] Stephen Boyd, Stephen P Boyd, and Lieven Vandenberghe. Convex optimization. Cambridge university press, 2004
[8] Sébastien Bubeck and Nicolo Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. arXiv preprint arXiv:1204.5721, 2012.
[9] Nicolo Cesa-Bianchi and Gábor Lugosi. Prediction, learning, and games. Cambridge university press, 2006.
[10] Alon Cohen, Tomer Koren, and Yishay Mansour. Learning linear-quadratic regulators efficiently with only $\sqrt{T}$ regret. In International Conference on Machine Learning, pages 1300-1309, 2019.
[11] Amit Daniely, Alon Gonen, and Shai Shalev-Shwartz. Strongly adaptive online learning. In Francis Bach and David Blei, editors, Proceedings of the 32nd International Conference on Machine Learning, volume 37 of Proceedings of Machine Learning Research, pages 1405-1411, Lille, France, 07-09 Jul 2015. PMLR.
[12] Sarah Dean, Horia Mania, Nikolai Matni, Benjamin Recht, and Stephen Tu. Regret bounds for robust adaptive control of the linear quadratic regulator. In Advances in Neural Information Processing Systems, pages 4188-4197, 2018.
[13] Sarah Dean and Benjamin Recht. Certainty equivalent perception-based control. arXiv preprint arXiv:2008.12332, 2020.
[14] Eyal Even-Dar, Sham M Kakade, and Yishay Mansour. Experts in a markov decision process. In Advances in neural information processing systems, pages 401-408, 2005.
[15] Maryam Fazel, Rong Ge, Sham Kakade, and Mehran Mesbahi. Global convergence of policy gradient methods for the linear quadratic regulator. In International Conference on Machine Learning, pages 1467-1476, 2018.
[16] Paula Gradu, Elad Hazan, and Edgar Minasyan. Adaptive regret for control of time-varying dynamics. arXiv preprint arXiv:2007.04393, 2020.
[17] Johan Håstad. Some optimal inapproximability results. Journal of the ACM (JACM), 48(4):798859, 2001.
[18] Elad Hazan. Introduction to online convex optimization. Foundations and Trends $®$ ) in Optimization, 2(3-4):157-325, 2016.
[19] Elad Hazan. Introduction to online convex optimization, 2019.
[20] Elad Hazan, Sham Kakade, and Karan Singh. The nonstochastic control problem. In Proceedings of the 31st International Conference on Algorithmic Learning Theory, pages 408-421. PMLR, 2020.
[21] Elad Hazan and Comandur Seshadhri. Efficient learning algorithms for changing environments. In Proceedings of the 26th annual international conference on machine learning, pages 393-400. ACM, 2009.
[22] Mark Herbster and Manfred K Warmuth. Tracking the best expert. Machine learning, 32(2):151178, 1998.
[23] Sham Kakade, Akshay Krishnamurthy, Kendall Lowrey, Motoya Ohnishi, and Wen Sun. Information theoretic regret bounds for online nonlinear control. arXiv preprint arXiv:2006.12466, 2020.
[24] Sahin Lale, Kamyar Azizzadenesheli, Babak Hassibi, and Anima Anandkumar. Logarithmic regret bound in partially observable linear dynamical systems, 2020.
[25] Sahin Lale, Kamyar Azizzadenesheli, Babak Hassibi, and Anima Anandkumar. Regret minimization in partially observable linear quadratic control, 2020.
[26] Horia Mania, Stephen Tu, and Benjamin Recht. Certainty equivalence is efficient for linear quadratic control. In Advances in Neural Information Processing Systems, pages 10154-10164, 2019.
[27] Zakaria Mhammedi, Dylan J Foster, Max Simchowitz, Dipendra Misra, Wen Sun, Akshay Krishnamurthy, Alexander Rakhlin, and John Langford. Learning the linear quadratic regulator from nonlinear observations. arXiv preprint arXiv:2010.03799, 2020.
[28] Richard H Middleton and Graham C Goodwin. Adaptive control of time-varying linear systems. IEEE Transactions on Automatic Control, 33(2):150-155, 1988.
[29] Kevin L Moore. Iterative learning control for deterministic systems. Springer Science \& Business Media, 2012.
[30] Samet Oymak and Necmiye Ozay. Non-asymptotic identification of liti systems from a single trajectory. In 2019 American control conference (ACC), pages 5655-5661. IEEE, 2019.
[31] Necmiye Ozay, Constantino Lagoa, and Mario Sznaier. Set membership identification of switched linear systems with known number of subsystems. Automatica, 51:180-191, 2015.
[32] Guannan Qu, Yuanyuan Shi, Sahin Lale, Anima Anandkumar, and Adam Wierman. Stable online control of ltv systems stable online control of linear time-varying systems. arXiv preprint arXiv:2104.14134, 2021.
[33] Alexander Rakhlin and Karthik Sridharan. Online non-parametric regression. In Conference on Learning Theory, pages 1232-1264. PMLR, 2014.
[34] Vincent Roulet, Siddhartha Srinivasa, Dmitriy Drusvyatskiy, and Zaid Harchaoui. Iterative linearized control: stable algorithms and complexity guarantees. In International Conference on Machine Learning, pages 5518-5527. PMLR, 2019.
[35] Tuhin Sarkar, Alexander Rakhlin, and Munther Dahleh. Nonparametric system identification of stochastic switched linear systems. In 2019 IEEE 58th Conference on Decision and Control (CDC), pages 3623-3628. IEEE, 2019.
[36] Tuhin Sarkar, Alexander Rakhlin, and Munther A Dahleh. Finite-time system identification for partially observed lti systems of unknown order. arXiv preprint arXiv:1902.01848, 2019.
[37] Max Simchowitz. Making non-stochastic control (almost) as easy as stochastic, 2020.
[38] Max Simchowitz and Dylan Foster. Naive exploration is optimal for online lqr. In International Conference on Machine Learning, pages 8937-8948. PMLR, 2020.
[39] Max Simchowitz, Karan Singh, and Elad Hazan. Improper learning for non-stochastic control, 2020.
[40] Y. Tassa, T. Erez, and E. Todorov. Synthesis and stabilization of complex behaviors through online trajectory optimization. In 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems, pages 4906-4913, 2012.
[41] Emanuel Todorov and Weiwei Li. A generalized iterative lqg method for locally-optimal feedback control of constrained nonlinear stochastic systems. In Proceedings of the 2005, American Control Conference, 2005., pages 300-306. IEEE, 2005.
[42] Kostas S Tsakalis and Petros A Ioannou. Linear time-varying systems: control and adaptation. Prentice-Hall, Inc., 1993.
[43] Roman Vershynin. High-dimensional probability: An introduction with applications in data science, volume 47. Cambridge university press, 2018.
[44] Yuh-Shyang Wang, Nikolai Matni, and John C Doyle. A system-level approach to controller synthesis. IEEE Transactions on Automatic Control, 64(10):4079-4093, 2019.
[45] Tyler Westenbroek, Max Simchowitz, Michael I Jordan, and S Shankar Sastry. On the stability of nonlinear receding horizon control: a geometric perspective. arXiv preprint arXiv:2103.15010, 2021.
[46] Dante Youla, Hamid Jabr, and Jr Bongiorno. Modern wiener-hopf design of optimal controllerspart ii: The multivariable case. IEEE Transactions on Automatic Control, 21(3):319-338, 1976.
[47] Lijun Zhang, Tie-Yan Liu, and Zhi-Hua Zhou. Adaptive regret of convex and smooth functions. arXiv preprint arXiv:1904.11681, 2019.
[48] Lijun Zhang, Tianbao Yang, rong jin, and Zhi-Hua Zhou. Dynamic regret of strongly adaptive methods. In Jennifer Dy and Andreas Krause, editors, Proceedings of the 35th International Conference on Machine Learning, volume 80 of Proceedings of Machine Learning Research, pages 5882-5891. PMLR, 10-15 Jul 2018.
[49] Kemin Zhou, John Comstock Doyle, Keith Glover, et al. Robust and optimal control, volume 40. Prentice hall New Jersey, 1996.

## Contents

1 Introduction ..... 1
1.1 Contributions ..... 2
1.2 Related Work ..... 3
2 Problem Setting ..... 3
2.1 Benchmarks and Policy Classes ..... 5
3 Online Control over Convex Policies ..... 5
3.1 A Linear Regret Lower Bound ..... 6
3.2 Estimation of Time-Varying Vector Sequences ..... 6
3.3 Adaptive Regret for Control of Unknown Time-Varying Dynamics ..... 7
4 Online Control over State Feedback ..... 8
5 Discussion and Future Work ..... 9
A Extensions ..... 15
A. 1 Affine Offsets ..... 15
A. 2 Changing Stabilizing Controllers ..... 15
A. 3 Partial Observation ..... 16
A. 4 The DAC parametrization ..... 16
B Adapive Regret for Time-Varying DRC-OGD ..... 17
B. 1 Adaptive Regret of OGD for functions with memory ..... 17
B. 2 Proof of Theorem B. 1 ..... 19
C Estimation of Time-Varying Vector Sequences ..... 22
C. 1 Adaptive Regret Bound ..... 24
C.1.1 Working Set Construction ..... 26
C. 2 No Strong Adaptivity ..... 26
D Adaptive Regret for Control of Changing Unknown Dynamics ..... 27
D. 1 Estimation of the Markov Operator ..... 28
D. 2 Error Sensitivity ..... 31
D. 3 Proof of Theorem D. 1 ..... 34
E Sublinear Regret for State Feedback ..... 35
E. 1 Proof of Theorem E. 1 ..... 36
E. 2 Ommited Proofs ..... 40
F Lower Bounds and Separations ..... 41
F. 1 Separation between policy classes ..... 41
F. 2 Linear Regret Against DrC and DAC ..... 44
F.2.1 Omitted proofs ..... 46
F. 3 Lower Bound without Stability ..... 47
F. 4 Hardness of Computing Best State Feedback Controller ..... 48

## A Extensions

## A. 1 Affine Offsets

For many systems, performance improves dramatically for controllers with constant affine terms, that is

$$
u_{t}^{M}=\bar{u}^{M}+\sum_{t=0}^{h-1} M^{[i]} x_{t-i}^{\mathrm{nat}}
$$

where $\bar{u}^{M}$ is a constant affine term encoded by $M=\left(M^{[0]}, \ldots, M^{[h-1]}, \bar{u}^{M}\right)$. All our arguments apply more generally to control policies of this form. Moreover, we can even allow linear combinations of time varying terms:

$$
\begin{aligned}
u_{t}^{M}= & \sum_{t=0}^{h-1} M^{[i]} x_{t-i}^{\mathrm{nat}}+\sum_{t=0}^{h^{\prime}-1} M^{[i]^{\prime}} \psi_{i}(t) \\
& M=\left(M^{[0]}, \ldots, M^{[h-1]}, M^{[0]^{\prime}}, \ldots, M^{\left[h^{\prime}-1\right]^{\prime}}\right)
\end{aligned}
$$

where now $\psi_{i}(t)$ are fixed, possibly time varying basis functions (which do not depend on $M$ ). The case of constant affine terms corresponds to $h^{\prime}=1$, and $\psi_{i}(t)=1$ for all $t$.

## A. 2 Changing Stabilizing Controllers

Our results extend naturally to the following setting: for each time $t=1,2, \ldots, T$, the algorithm has access to a static feedback control policy $K_{t}^{\text {stb }}$ such that the closed loop matrices $\left(A_{t}+B_{t} K_{t}^{\text {stb }}\right)$ are sequentially stable, that is

$$
\Phi_{s+h, s}^{\mathrm{stb}}:=\prod_{i=s}^{s+h}\left(A_{t}+B_{t} K_{t}^{\mathrm{stb}}\right)
$$

has geometric decay. We let $x_{t}^{\mathrm{nat}}$ denote the iterates produced by the updates

$$
x_{t+1}^{\mathrm{nat}}=\left(A_{t}+B_{t} K_{t}^{\mathrm{stb}}\right) x_{t}^{\mathrm{nat}}, u_{t}^{\mathrm{nat}}=K x_{t}^{\mathrm{nat}}, \quad x_{1}^{\mathrm{nat}}=0 .
$$

We compute the stabilized policies of the form

$$
u_{t}^{M}=K_{t}^{\mathrm{stb}} x_{t}+\sum_{i=0}^{h-1} M^{[i]} x_{t-i}^{\mathrm{nat}}
$$

To facillicate the extension, we define the stabilized Markov operator

$$
G_{\star, t}^{[0]}=\left[\begin{array}{c}
0 \\
I_{d_{u}}
\end{array}\right], \quad G_{\star, t}^{[i]}:=\left[\begin{array}{c}
I_{d_{x}} \\
K_{t}
\end{array}\right] \Phi_{t, t-i+1}^{\mathrm{stb}} B_{t-i},
$$

This Markov operator satisfies

$$
\left[\begin{array}{l}
x_{t} \\
u_{t}
\end{array}\right]=\left[\begin{array}{l}
x_{t}^{\mathrm{nat}} \\
u^{\mathrm{nat}}
\end{array}\right]+\sum_{i=0}^{t-1} G_{\star, t}^{[i]} \tilde{u}_{t}, \quad \text { s.t. } u_{t}=K_{t} x_{t}+\tilde{u}_{t} .
$$

With similar techniques, we obtain estimates $\widehat{G}_{t}$, back out estimates of the Nature's sequence $\left(x^{\text {nat }}, u^{\text {nat }}\right)$ via

$$
\left[\begin{array}{l}
\hat{x}_{t}^{\mathrm{nat}} \\
\hat{u}^{\mathrm{nat}}
\end{array}\right]=\operatorname{clip}_{r}\left(\left[\begin{array}{l}
x_{t} \\
u_{t}
\end{array}\right]-\sum_{i=0}^{h} \widehat{G}_{t}^{[i]}\left(u_{t-i}-K_{t-i}^{\mathrm{stb}} x_{t-i}\right)\right),
$$

for a truncation radius $r>0$ suitably chosen. Recall that we apply this $\operatorname{clip}_{r}(\cdot)$ operator to ensure the estimates of the Nature's sequenc does not grow unbounded and exert feedback. We then select inputs

$$
\hat{u}_{t}(M)=K_{t}^{\mathrm{stb}} x_{t}+\sum_{i=0}^{h-1} M^{[i]} \hat{x}_{t-i}^{\mathrm{nat}}
$$

Markov operator with low adaptive regret, and apply our OcO-with-memory algorithm to losses

$$
\widehat{f}_{t}(M)=c_{t}\left(\left[\begin{array}{c}
\hat{x}_{t}^{\text {nat }} \\
\hat{u}_{t}^{\text {nat }}
\end{array}\right]-\sum_{i=0}^{h} \widehat{G}_{t}^{[i]} \hat{u}_{t-i}(M)\right) .
$$

## A. 3 Partial Observation

Our results further extend to partially observed systems. We explain this extension for sequentially stable systems; extensions to sequentialy stabilized systems by time varying linear dynamic controllers follows from the exposition in [39], Appendix C.

For partially observed systems, we have the same state transition dynamics $x_{t+1}=A_{t} x_{t}+B_{t} u_{t}+w_{t}$, but, for a time-varying observation matrix $C_{t}$ and process noise $e_{t}$, we observe outputs $y_{t} \in \mathbb{R}^{d_{y}}$

$$
y_{t}=C_{t} x_{t}+e_{t} .
$$

Costs $c_{t}\left(y_{t}, u_{t}\right)$ are suffered on input and outputs. As for full observation, the Nature's sequence $x^{\text {nat }}$ and $y^{\text {nat }}$ correspond to states $x_{t}$ and outputs $y_{t}$ which arise under identially zero input $u_{t} \equiv 0$. The DRC parametrization selects linear conmbinations of Nature's y's:

$$
u_{t}^{M}=\sum_{i=0}^{m-1} M^{[i]} y_{t-i}^{\mathrm{nat}}
$$

Recalling

$$
\Phi_{s+h, s}:=\prod_{i=s}^{s+h}\left(A_{t}+B_{t}\right)
$$

the relevant Markov operators $G_{\star t}$ are the ones mapping inputs to outputs:

$$
G_{\star, t}^{[0]}=0, \quad G_{\star, t}^{[i]}=C_{t} \Phi_{t, t-i+1} B_{t-i}
$$

With similar techniques, we obtain estimates $\widehat{G}_{t}$, back out estimates of the Nature's sequence $y^{\text {nat }}$ via

$$
\hat{y}^{\mathrm{nat}}=\operatorname{clip}_{r}\left(y_{t}-\sum_{i=0}^{h} \widehat{G}_{t}^{[i]} u_{t-i}\right)
$$

for trunctation radius $r>0$ suitably chosen, we select inputs

$$
\left.\hat{u}_{t}(M)=\sum_{i=0}^{h-1} M^{[i]} \hat{y}_{t-i}^{\mathrm{nat}}\right)
$$

update parameters with the Oco-with-memory losses

$$
\widehat{f}_{t}(M)=\ell_{t}\left(\hat{y}_{t}^{\mathrm{nat}}-\sum_{i=0}^{h} \widehat{G}_{t}^{[i]} M^{[i]} \hat{u}_{t-i}(M), u_{t}\right)
$$

## A. 4 The DAC parametrization

Here we sketch an algorithm to compete with DAC-parametrized control policies [3]. For simplicity, we focus on sequentially stable system, though the discussion extends to systems sequentially stabilized by sequences of controllers $\left(K_{t}^{\text {stb }}\right)$. Note that DAC does not apply under partial observation.
Recall that, in the DAC parametrization, the inputs are selected as linear combinations of past disturbances:

$$
u_{t}^{M}=\sum_{i=0}^{h-1} M^{[i]} w_{t-i-1}
$$

To implement DAC, we therefore need empiricals estimate $\hat{w}_{t}$ of $w_{t}$. As per [20], it suffices to construct estimates $\left(\hat{A}_{t}, \hat{B}_{t}\right)$ of $\left(A_{t}, B_{t}\right)$, and choose

$$
\hat{w}_{t}=\operatorname{clip}_{r}\left(x_{t}-\hat{A}_{t-1} x_{t-1}-\hat{B}_{t-1} u_{t-1}\right),
$$

again clipped at a suitable radius $r>0$ to block compounding feedback. Given these estimates, our algorithm extends to DAC control in the expected way.

How does one obtain the estimates $\left(\hat{A}_{t}, \hat{B}_{t}\right)$ ? First, we observe that since

$$
G_{\star, t}^{[0]}=0, \quad G_{\star, t}^{[i]}=\Phi_{t, t-i+1} B_{t-i}, \quad \Phi_{s+h, s}=\prod_{i=s}^{s+h} A_{i},
$$

we have $B_{t}=G_{\star, t}^{[1]}$. Hence, we can select $\hat{B}_{t}$ as $\widehat{G}_{t}^{[1]}$.
The estimate of $A_{t}$ is more involved. For linear, time-invariant systems, $A$ can be recovered from $G_{\star}$ via the Ho-Kalman procedure as is does in [20] (see also [30, 36]). For time-varying systems, this become more challenging. Ommitting details in the interest of brevity, one can use the robustness properties of Ho-Kalman to argue that if the system matrices are slow moving (an assumption required for low regret), $G_{\star, t}$ is close to a stationarized analogue $\bar{G}_{\star, t}$ given by

$$
\bar{G}_{\star, t}^{[i]}:=A_{t}^{i-1} B_{t} .
$$

Hence, we can view any estimate $\widehat{G}_{t}$ of $G_{\star, t}$ as an estimate of $\bar{G}_{\star, t}$, and apply Ho-Kalman to the latter.

## B Adapive Regret for Time-Varying DRC-OGD

We first extend the DRC-OGD algorithm from [39] to the setting of known linear time-varying dynamics. We spell out Algorithm 3 and prove it attains $\tilde{\mathcal{O}}(\sqrt{T})$ adaptive regret over general convex costs under fully adversarial noise (Theorem B.1) with respect to the DRC policy class. The main technique, using OGD over the DRC parametrization, remains unchanged from the original paper and we show it generalizes naturally to LTV systems.

```
Algorithm 3 Disturbance Response Control via Online Gradient Descent (DRC-OGD)
    Input: stepsize \(\eta\), memory \(m\), radius \(R_{\mathcal{M}}\)
    Initialize \(M_{1} \in \mathcal{M}\left(m, R_{\mathcal{M}}\right)\) arbitrarily
    Receive initial state \(x_{1}\), set \(x_{1}^{\text {nat }}=x_{1}\) and \(x_{\leq 0}^{\text {nat }}=0\)
    for \(t=1 \ldots T\) do
        Play control \(u_{t}=\sum_{i=0}^{m-1} M_{t}^{[i]} x_{t-i}^{\mathrm{nat}}\)
        Suffer \(c_{t}\left(x_{t}, u_{t}\right)\) and observe cost function \(c_{t}(\cdot, \cdot)\)
        Construct \(f_{t}\left(M_{0}, \ldots, M_{h}\right) \doteq c_{t}\left(\hat{x}_{t}\left(M_{0: h-1}\right), u_{t}\left(M_{h}\right)\right)\) and let \(\tilde{f}_{t}(M) \doteq f_{t}(M, \ldots, M)\)
        Update \(M_{t+1} \leftarrow \Pi_{\mathcal{M}}\left(M_{t}-\eta \nabla \tilde{f}_{t}\left(M_{t}\right)\right)\)
        Receive new system \(G_{t}\)
        Receive new state \(x_{t+1}\) and extract \(x_{t+1}^{\mathrm{nat}}=x_{t+1}-\sum_{i=0}^{t-1} G_{t}^{[i]} u_{t-i}\) or receive \(x_{t+1}^{\mathrm{nat}}\)
```

Theorem B.1. Running Algorithm 3 with $\eta=\frac{\sqrt{d_{\min }} R_{\mathcal{M}}^{2}}{2 L R_{\mathrm{sys}}^{2}(h+1)^{5 / 4} \sqrt{T}}$ guarantees the following regret bound on every interval $I=[r, s]$ :

$$
\sum_{t=r}^{s} c_{t}\left(x_{t}, u_{t}\right)-\min _{\pi \in \Pi_{\mathrm{drc}}} \sum_{t=r}^{s} c_{t}\left(x_{t}^{\pi}, u_{t}^{\pi}\right) \leq 6 L R_{\mathrm{sys}}^{2}\left(3 \sqrt{d_{\min }} m(h+1)^{5 / 4} \sqrt{T}+\psi(h)|I|\right)
$$

where $R_{\mathrm{sys}}=R_{G} R_{\mathcal{M}} R_{\mathrm{nat}}$ and $d_{\mathrm{min}}=\min \left\{d_{x}, d_{u}\right\}$.

## B. 1 Adaptive Regret of OGD for functions with memory

We first prove that OGD with a fixed stepsize attains $O(\sqrt{T})$ adaptive regret for functions with memory.

```
Algorithm 4 Online Gradient Descent for OCOwMem (Mem-OGD)
    Input: stepsize \(\eta\), memory \(m\), set \(\mathcal{K}\)
    Initialize \(x_{1} \in \mathcal{K}\) arbitrarily, set \(x_{\leq 0}=x_{1}\) by convention
    for \(t=1 \ldots T\) do
        Play \(x_{t}\)
        Suffer \(f_{t}\left(x_{t-h}, \ldots, x_{t}\right)\) and observe loss function \(f_{t}(\cdot, \ldots, \cdot)\)
        Construct proxy loss function \(\tilde{f}_{t}(x) \doteq f_{t}(x, \ldots, x)\)
        Update \(x_{t+1}=\Pi_{\mathcal{K}}\left(x_{t}-\eta \nabla \tilde{f}_{t}\left(x_{t}\right)\right)\)
```

Theorem B.2. Let $\left\{f_{t}: \mathcal{K}^{h+1} \rightarrow[0,1]\right\}_{t=1}^{T}$ be a sequence of $L$ coordinate-wise Lipschitz loss functions with memory such that $\tilde{f}_{t}$ (Line 6) is convex. Then, on any interval $I=[r, s] \subseteq[T]$, Algorithm 4 enjoys the following adaptive policy regret guarantee:

$$
\sum_{t=r}^{s} f_{t}\left(x_{t-h}, \ldots, x_{t}\right)-\min _{x \in \mathcal{K}} \sum_{t=r}^{s} \tilde{f}_{t}(x) \leq \frac{D^{2}}{\eta}+2 \eta L^{2}(h+1)^{5 / 2}|I|
$$

where $D=\operatorname{diam}(\mathcal{K})$.
First we state and prove the following well-known fact about vanilla projected OGD over (memoryless) loss functions:
Fact B.1. Let $\left\{\tilde{f}_{t}: \mathcal{K} \rightarrow[0,1]\right\}_{t=1}^{T}$ be a sequence of convex loss functions with $\|\nabla \tilde{f}(x)\| \leq G$. Then, on any interval $I=[r, s] \subseteq[T]$, projected OGD enjoys the following guarantee:

$$
\sum_{t=r}^{s} \tilde{f}_{t}\left(x_{t}\right)-\min _{x \in \mathcal{K}} \sum_{t=r}^{s} \tilde{f}(x) \leq \frac{D^{2}}{\eta}+\eta|I| G^{2}
$$

where $D=\operatorname{diam}(\mathcal{K})$.
Proof. Consider an arbitary interval $I=[r, s] \subseteq[T]$ Let $\mathbf{x}^{\star}=\arg \min _{x \in \mathcal{K}} \sum_{t=r}^{s} \tilde{f}_{t}(x)$ and denote $\nabla_{t} \doteq \nabla \tilde{f}_{t}\left(x_{t}\right)$ for simplicity. By convexity we have

$$
\begin{equation*}
\tilde{f}_{t}\left(x_{t}\right)-\tilde{f}_{t}\left(x^{\star}\right) \leq \nabla_{t}^{\top}\left(x_{t}-x^{\star}\right) \tag{B.1}
\end{equation*}
$$

By the Pythagorean theorem

$$
\begin{align*}
\left\|x_{t+1}-x^{\star}\right\|^{2} & =\left\|\Pi_{\mathcal{K}}\left(x_{t}-\eta \nabla_{t}\right)-x^{\star}\right\|^{2} \\
& \leq\left\|x_{t}-\eta \nabla_{t}-x^{\star}\right\|^{2} \\
& \leq\left\|x_{t}-x^{\star}\right\|^{2}+\eta^{2}\left\|\nabla_{t}\right\|^{2}-2 \eta \nabla_{t}^{\top}\left(x_{t}-x^{\star}\right) \tag{B.2}
\end{align*}
$$

Hence we can bound the interval regret as:

$$
\begin{align*}
\Rightarrow 2 \sum_{t=r}^{s}\left(f_{t}\left(x_{t}\right)-f_{t}\left(x^{\star}\right)\right) & \leq 2 \sum_{t=r}^{s} \nabla_{t}^{\top}\left(x_{t}-x^{\star}\right)  \tag{B.1}\\
& \leq \sum_{t=r}^{s}\left(\frac{\left\|x_{t}-x^{\star}\right\|^{2}-\left\|x_{t+1}-x^{\star}\right\|^{2}}{\eta}+\eta G^{2}\right)  \tag{B.2}\\
& =\frac{\left\|x_{r}-x^{\star}\right\|^{2}}{\eta}+|I| \eta G^{2} \\
& \leq \frac{D^{2}}{\eta}+\eta|I| G^{2}
\end{align*}
$$

which yields the desired adaptive regret bound.
Using this simple fact and Lipschitzness we are able to easily prove the desired guarantee for Algorithm 4.

Proof of Theorem B.2. First note that Algorithm 4 is just doing gradient descent on the proxy convex loss functions $\tilde{f}_{t}: \mathcal{K} \rightarrow[\underset{\sim}{0}, 1]$. Hence, as long as we identify the gradient bound we can apply Fact B. 1 to get a bound on $\tilde{f}$-regret. Observe that

$$
\begin{aligned}
\left|\tilde{f}_{t}(x)-\tilde{f}_{t}(y)\right| & =\left|f_{t}(x, \ldots, x)-f_{t}(y, \ldots, y)\right| \\
& \leq\left|f_{t}(x, \ldots, x)-f_{t}(y, \ldots, x)\right|+\left|f_{t}(y, x, \ldots, x)-f_{t}(y, y, \ldots, x)\right|+ \\
& \ldots+\left|f_{t}(y, \ldots, y, x)-f_{t}(y, \ldots, y)\right| \\
& \leq L(h+1)\|x-y\|
\end{aligned}
$$

so $\tilde{f}_{t}$ is $L(h+1)$-Lipschitz and hence has a gradient bound of $L(h+1)$. So we can apply Fact B. 1 to get

$$
\begin{equation*}
\sum_{t=r}^{s} \tilde{f}_{t}\left(x_{t}\right)-\min _{x \in \mathcal{K}} \sum_{t=r}^{s} \tilde{f}(x) \leq \frac{D^{2}}{\eta}+\eta|I| L^{2}(h+1)^{2} \tag{B.3}
\end{equation*}
$$

We can use this to bound the adaptive policy regret. First note that

$$
\left\|x_{t}-x_{t-1}\right\|=\left\|\eta \nabla \tilde{f}_{t}\left(x_{t-1}\right)\right\| \leq \eta L(h+1)
$$

and by the triangle inequality

$$
\begin{equation*}
\left\|x_{t}-x_{t-i}\right\| \leq \sum_{j=1}^{i}\left\|x_{t-j+1}-x_{t-j}\right\| \leq \eta L(h+1) \cdot i \tag{B.4}
\end{equation*}
$$

Using Eq. (B.4) and Lipschitzness we have:

$$
\begin{align*}
f_{t}\left(x_{t-h}, \ldots, x_{t}\right)-\tilde{f}_{t}\left(x_{t}\right) & \leq L\left\|\left(x_{t-h}, \ldots, x_{t}\right)-\left(x_{t}, \ldots, x_{t}\right)\right\| \\
& \leq L \sqrt{\sum_{i=1}^{h}\left\|x_{t}-x_{t-i}\right\|^{2}} \\
& \leq \eta L^{2}(h+1) \cdot \sqrt{\sum_{i=1}^{h} i^{2}} \\
& \leq \eta L^{2}(h+1)^{5 / 2} \tag{B.5}
\end{align*}
$$

Combining everything we get

$$
\begin{aligned}
\sum_{t=r}^{s} f_{t}\left(x_{t-h}, \ldots, x_{t}\right)-\min _{x \in \mathcal{K}} \sum_{t=r}^{s} \tilde{f}_{t}(x) & =\underbrace{\sum_{t=r}^{s} f_{t}\left(x_{t-h}, \ldots, x_{t}\right)-\tilde{f}_{t}\left(x_{t}\right)}_{\text {(dist. to proxy loss [Eq. (B.5)]) }}+\underbrace{\sum_{t=r}^{s} \tilde{f}_{t}\left(x_{t}\right)-\min _{x \in \mathcal{K}} \sum_{t=r}^{s} \tilde{f}_{t}\left(x_{t}\right)}_{(\tilde{f} \text {-regret [Eq. (B.3)]) }} \\
& \leq 2 \eta L^{2}(h+1)^{5 / 2}|I|+\frac{D^{2}}{\eta}
\end{aligned}
$$

## B. 2 Proof of Theorem B. 1

We first prove that the constructed loss function satisfies key properties for efficient optimization.
Lemma B. 2 (Convexity). The loss functions $\tilde{f}_{t}$ constructed in Appendix B of Algorithm 3 are convex in $M$.

Proof. By definition we have that:

$$
\begin{align*}
\hat{x}_{t}(M) & =x_{t}^{\mathrm{nat}}+\sum_{i=1}^{h} G_{t}^{[i]} u_{t-i} \\
& =x_{t}^{\mathrm{nat}}+\sum_{i=1}^{h} G_{t}^{[i]}\left(\sum_{j=0}^{m-1} M^{[j]} x_{t-i-j}^{\mathrm{nat}}\right) \tag{B.6}
\end{align*}
$$

which is affine in $M$. Even more simply, we have $u_{t}(M)=\sum_{i=0}^{m-1} M^{[i]} x_{t-i}^{\mathrm{nat}}$.
Since $\hat{x}_{t}(M)$ and $u_{t}(M)$ are affine, and, respectively, linear functions of $M$ and composition with the convex cost $c_{t}$ preserves convexity we get the desired property.

Lemma B. 3 (Lipschitzness). The loss functions $f_{t}$ constructed in Appendix B of Algorithm 3 are $L_{f}$ coordinate-wise Lipschitz for $L_{f}=3 L R_{n a t}^{2} R_{G}^{2} R_{\mathcal{M}} \sqrt{m}$.

Proof. Observe that by Eq. (B.6) we have $\| \hat{x}_{t}\left(M_{0: h)} \| \leq R_{\text {nat }}\left(1+R_{G} R_{\mathcal{M}}\right)\right.$. Straightforwardly, $\left\|u_{t}\left(M_{h}\right)\right\| \leq R_{\text {nat }} R_{\mathcal{M}}$ as well.
For an arbitrary $i \in \overline{0, h}$, denoting $M_{0: h}=\left(M_{0}, \ldots, M_{i}, \ldots M_{h}\right), \tilde{M}_{0: h}=\left(M_{0}, \ldots, \tilde{M}_{i}, \ldots, M_{h}\right)$ and using the sub-quadratic lipschitzness of the costs we have:

$$
\begin{aligned}
\left|f_{t}\left(M_{0: h}\right)-f_{t}\left(\tilde{M}_{0: h}\right)\right| & =\left|c_{t}\left(\hat{x}_{t}\left(M_{0: h}\right), u_{t}\left(M_{0: h}\right)\right)-c_{t}\left(\hat{x}_{t}\left(\tilde{M}_{0: h}\right), u_{t}\left(\tilde{M}_{0: h}\right)\right)\right| \\
& \leq 3 L R_{\mathrm{nat}} R_{G} R_{\mathcal{M}}\left\|G_{t}^{[h-i]}\left(\sum_{j=0}^{m-1}\left(M_{i}^{[j]}-\tilde{M}_{i}^{[j]}\right) x_{t-i-j}^{\mathrm{nat}}\right)\right\|(i<h) \\
& \text { or } \leq 3 L R_{\mathrm{nat}} R_{G} R_{\mathcal{M}}\left(\sum_{j=0}^{m-1}\left(M_{i}^{[j]}-\tilde{M}_{i}^{[j]}\right) x_{t-i-j}^{\mathrm{nat}}\right) \quad(i=h) \\
& \leq 3 L R_{\mathrm{nat}}^{2} R_{G}^{2} R_{\mathcal{M}} \sqrt{m}\left\|M_{i}-\tilde{M}_{i}\right\|_{F}
\end{aligned}
$$

so the function is coordinate-wise Lipschitz with constant $L_{f}=3 L R_{\text {nat }}^{2} R_{G}^{2} R_{\mathcal{M}} \sqrt{m}$.
Lemma B. 4 (Euclidean Diameter). The euclidean diameter of $\mathcal{M}\left(m, R_{\mathcal{M}}\right)$ is at most $D=$ $2 \sqrt{m \cdot \min \left\{d_{x}, d_{u}\right\}} R_{\mathcal{M}}$.

Proof. That for an arbitrary $M \in \mathcal{M}\left(m, R_{\mathcal{M}}\right)$, we have

$$
\begin{aligned}
\|M\|_{F} & =\sqrt{\sum_{i=0}^{m-1}\left\|M^{[i]}\right\|_{F}^{2}} \\
& \leq \sqrt{m \cdot \max _{i \in[m-1]}\left\|M^{[i]}\right\|_{F}^{2}} \\
& =\sqrt{m} \max _{i \in[m-1]} \min \left\{d_{x}, d_{u}\right\}\left\|M^{[i]}\right\|_{o p} \\
& \leq \sqrt{m \cdot \min \left\{d_{x}, d_{u}\right\}}\|M\|_{\ell_{1}, o p} \\
& \leq \sqrt{m \cdot \min \left\{d_{x}, d_{u}\right\}} R_{\mathcal{M}}
\end{aligned}
$$

and the euclidean diameter is at most twice the maximal euclidean norm, concluding our statement.

The three lemmas above will allow us to use the results in Appendix B. 1 to obtain adaptive regret guarantees in terms of $f_{t}$ which truncated the effect on the state of actions further than $h$ in the past. To convert guarantees in terms of $f_{t}$ to ones in terms of $c_{t}$, we prove that the effect of the past is minimal:

Lemma B. 5 (Truncation Error). For a changing DRC policy that acts according to $M_{1}, \ldots, M_{t}$ up to time $t$ we have that:

$$
c_{t}\left(x_{t}, u_{t}\right)-c_{t}\left(\hat{x}_{t}\left(M_{t-h: t-1}\right), u_{t}\left(M_{t}\right)\right) \leq 3 L R_{n a t}^{2} R_{\mathcal{M}}^{2} R_{G} \psi(h)
$$

Proof. By the sub-quadratic Lipschitzness (and noting $\left\|u_{t}\right\|_{2} \leq R_{\mathcal{M}} R_{\text {nat }},\left\|x_{t}\right\|_{2} \leq R_{n} a t(1+$ $R_{G} R_{\mathcal{M}}$ ), and $u_{t}=u_{t}\left(M_{t}\right)$ ) we have:

$$
\begin{aligned}
c_{t}\left(x_{t}, u_{t}\right)-c_{t}\left(\hat{x}_{t}\left(M_{t-h: t-1}\right), u_{t}\left(M_{t}\right)\right) & \leq 3 L R_{\mathrm{nat}} R_{G} R_{\mathcal{M}} \| x_{t}\left(M_{1: t-1}-\hat{x}_{t}\left(M_{t-h: t-1}\right) \|\right. \\
& =3 L R_{\mathrm{nat}} R_{G} R_{\mathcal{M}}\left\|\sum_{i=h}^{t-1} G_{t-1}^{[i]} u_{t-1-i}\right\| \\
& \leq 3 L R_{\mathrm{nat}}^{2} R_{\mathcal{M}}^{2} R_{G} \psi(h)
\end{aligned}
$$

Having proven all these preliminary results the proof of the main theorem is immediate:
Proof of Theorem B.1. By the definition of the proxy loss in Line 7 of Algorithm 3, we can expand the regret of Algorithm 3 over interval $I=[r, s]$ as:

$$
\begin{aligned}
\text { Regret } & =\sum_{t=r}^{s} c_{t}\left(x_{t}, u_{t}\right)-\min _{M \in \mathcal{M}} \sum_{t=r}^{s} c_{t}\left(x_{t}^{M}, u_{t}^{M}\right) \\
& =\underbrace{\sum_{t=r}^{s} c_{t}\left(x_{t}, u_{t}\right)-\sum_{t=r}^{s} c_{t}\left(\hat{x}_{t}\left(M_{t-h: t-1}\right), u_{t}\left(M_{t}\right)\right)}_{\text {(truncation error I) }}+\underbrace{\sum_{t=r}^{s} f_{t}\left(M_{t-h: t}\right)-\min _{M \in \mathcal{M}} \tilde{f}_{t}(M)}_{\text {(truncation error II) }} \\
& +\underbrace{\min _{M \in \text { ( }} \sum_{t=r}^{s} c_{t}\left(\hat{x}_{t}(M), u_{t}(M)\right)-\min _{M \in \mathcal{M}} \sum_{t=r}^{s} c_{t}\left(x_{t}^{M}, u_{t}^{M}\right)}_{M \in \mathcal{M}}
\end{aligned}
$$

The first truncation error is bounded directly by Lemma B.5. For the second truncation error, let $M^{\star}=\arg \min _{M \in \mathcal{M}} \sum_{t=r}^{s} c_{t}\left(x_{t}^{M}, u_{t}^{M}\right)$. Clearly we have

$$
\min _{M \in \mathcal{M}} \sum_{t=r}^{s} c_{t}\left(\hat{x}_{t}(M), u_{t}(M)\right) \leq \sum_{t=r}^{s} c_{t}\left(\hat{x}_{t}\left(M^{\star}\right), u_{t}\left(M^{\star}\right)\right)
$$

and hence we can apply Lemma B. 5 to bound:

$$
\begin{aligned}
\text { truncation error II } & =\min _{M \in \mathcal{M}} \sum_{t=r}^{s} c_{t}\left(\hat{x}_{t}(M), u_{t}(M)\right)-\min _{M \in \mathcal{M}} \sum_{t=r}^{s} c_{t}\left(x_{t}^{M}, u_{t}^{M}\right) \\
& \leq \sum_{t=r}^{s} c_{t}\left(\hat{x}_{t}\left(M^{\star}\right), u_{t}\left(M^{\star}\right)\right)-\sum_{t=r}^{s} c_{t}\left(x_{t}^{M^{\star}}, u_{t}^{M^{\star}}\right) \\
& \leq 3 L R_{\mathrm{nat}}^{2} R_{\mathcal{M}}^{2} R_{G} \psi(h)|I|
\end{aligned}
$$

Finally, due to Lemma B.2, Lemma B. 3 and Lemma B. 4 we can apply Theorem B. 2 to get

$$
\text { f-regret } \leq \frac{D^{2}}{\eta}+2 \eta L_{f}^{2}(h+1)^{5 / 2}|I|
$$

Summing everything up and plugging in the Lipschitz and diameter constants, we have:

Regret $\leq 6 L R_{\text {nat }}^{2} R_{\mathcal{M}}^{2} R_{G} \psi(h)|I|+\frac{4 m \min \left\{d_{x}, d_{u}\right\} R_{\mathcal{M}}^{2}}{\eta}+18 \eta\left(L^{2} R_{\text {nat }}^{4} R_{G}^{4} R_{\mathcal{M}}^{2}\right) m(h+1)^{5 / 2}|I|$
Setting $\eta \doteq \frac{\sqrt{\min \left\{d_{x}, d_{u}\right\}}}{2 L R_{G}^{2} R_{\text {nat }}^{2}(h+1)^{5 / 4} \sqrt{T}}$, we get

$$
\begin{aligned}
\text { Regret } & \leq 6 L R_{\text {nat }}^{2} R_{\mathcal{M}}^{2} R_{G} \psi(h)|I|+17 \sqrt{\min \left\{d_{x}, d_{u}\right\}}\left(L R_{G}^{2} R_{\mathrm{nat}}^{2} R_{\mathcal{M}}^{2}\right) \cdot(h+1)^{5 / 4} m \cdot \sqrt{T} \\
& \leq 6 L R_{\mathrm{sys}}^{2}\left(3 \sqrt{d_{\min }} m(h+1)^{5 / 4} \sqrt{T}+\psi(h)|I|\right)
\end{aligned}
$$

where we denote $R_{\text {sys }} \doteq R_{G} R_{\mathcal{M}} R_{\text {nat }}$ and $d_{\text {min }} \doteq \min \left\{d_{x}, d_{u}\right\}$.

## C Estimation of Time-Varying Vector Sequences

In this section we segway into the setting of online prediction under a partial information model. The goal is to estimate a sequence of vectors under limited noisy feedback where the feedback access is softly restricted via additional cost. As shown in the following section, this setting captures the system identification phase of controlling an unknown time-varying dynamical system. We first extensively study the simplified setting as below, and afterwards transfer our findings into meaningful results in control.

Formally, consider the following repeated game between a learner and an oblivious adversary: at each round $t \in[T]$, the adversary picks a target vector $\mathbf{z}_{t}^{\star} \in \mathcal{K}$ from a convex decision set $\mathcal{K}$ contained in a 0 -centered ball of radius $R_{z}$; simultaneously, the learner selects an estimate $\hat{\mathbf{z}}_{t} \in \mathcal{K}$ and suffers quadratic loss $f_{t}\left(\hat{\mathbf{z}}_{t}\right)=\left\|\hat{\mathbf{z}}_{t}-\mathbf{z}_{t}^{\star}\right\|^{2}$. The only feedback the learner has access to is via the following noisy and costly oracle.
Oracle 2 (Noisy Costly Oracle). At each time $t \in[T]$, the learner selects a decision $b_{t} \in\{0,1\}$ indicating whether a query is sent to the oracle. If $b_{t}=1$, the learner receives an unbiased estimate $\tilde{\mathbf{z}}_{t}$ as response such that $\left\|\tilde{\mathbf{z}}_{t}\right\| \leq \tilde{R}_{z}$ and $\mathbb{E}\left[\tilde{\mathbf{z}}_{t} \mid \mathcal{F}_{t}, b_{t}=1\right]=\mathbf{z}_{t}^{\star}$ where $\mathcal{F}_{t}$ is the filtration sigma algebra generated by the entire sequence $\mathbf{z}_{1: T}^{\star}$ and the past $\tilde{\mathbf{z}}_{1: t-1}, b_{1: t-1}$. A completed query results in a unit cost for the learner denoted $b_{t}$ as well by abuse of notation.

The idea behind this setting is to model a general estimation framework for a time-varying system which focuses only on exploration. Committing to exploration, however, cannot realistically be free hence the additional cost for the number of calls to Oracle 2. Our goal is to design an algorithm $\mathcal{A}$ that minimizes the quadratic loss regret along with the extra oracle cost, defined over each interval $I=[r, s] \subseteq[T]$ as

$$
\begin{equation*}
\operatorname{Regret}_{I}(\mathcal{A} ; \lambda)=\mathbb{E}\left[\sum_{t \in I} f_{t}\left(\hat{\mathbf{z}}_{t}\right)\right]-\min _{\mathbf{z} \in \mathcal{K}} \sum_{t \in I} f_{t}(\mathbf{z})+\lambda \mathbb{E}\left[\sum_{t \in I} b_{t}\right] \tag{C.1}
\end{equation*}
$$

where $\lambda \geq 0$ is a scaling constant independent of the horizon $T$. For the $I=[T]$ entire interval, we use $T$ as a subscript instead of $[T]$. The expectation above is taken over both the (potential) randomness of the algorithm and the stochasticity of the oracle responses; it is taken in the round order $t=1, \ldots, T$ at each round conditioning on the past iterations.
In terms of estimation itself, the metric to consider over interval $I$ is given by $\operatorname{Regret}_{I}(\mathcal{A} ; 0)$ that ignores the oracle call costs. Furthermore, we observe that the best-in-hindsight term in (C.1) is in fact a fundamental quantity of the vector sequence as defined below. This formulation will be used, and is more appropriate, when transferring our findings to the setting of control.
Definition C.1. Define the variability of a time-varying vector sequence $\mathbf{z}_{1: T}$ over an interval $I \subseteq[T]$ to be

$$
\operatorname{Var}_{I}\left(\mathbf{z}_{1: T}\right)=\frac{1}{|I|} \min _{\mathbf{z} \in \mathcal{K}} \sum_{t \in I}\left\|\mathbf{z}-\mathbf{z}_{t}\right\|^{2}=\frac{1}{|I|} \sum_{t \in I}\left\|\overline{\mathbf{z}}_{I}-\mathbf{z}_{t}\right\|^{2}
$$

where $\overline{\mathbf{z}}_{I}=|I|^{-1} \sum_{t \in I} \mathbf{z}_{t} \in \mathcal{K}$ is the empirical average of the members of the sequence that correspond to $I$.

This definition concludes the setup of our abstraction to general estimation of vector sequences. Regarding algorithmic results, we first present a base method that achieves logarithmic regret over the entire trajectory $[1, T]$. The idea is for the learner to uniformly query Oracle 2 with probability $p$ : once an estimate $\tilde{\mathbf{z}}_{t}$ is received, construct a stochastic gradient with expectation equal to the true gradient, and perform a gradient update. The algorithm is described in detail in Algorithm 5, and its guarantee given in the theorem below.

```
Algorithm 5 Base Estimation Algorithm
    Input: \(p, \hat{\mathbf{z}}_{1} \in \mathcal{K}\)
    for \(t=1, \ldots, T\) do
        Play iterate \(\hat{\mathbf{z}}_{t}\)
        Draw/Receive \(b_{t} \sim \operatorname{Bernoulli}(p)\)
        if \(b_{t}=1\) then
            Receive estimate \(\tilde{\mathbf{z}}_{t}\) from Oracle 2
            Construct importance weighted gradient \(\tilde{\nabla}_{t}:=\frac{1}{p}\left(\hat{\mathbf{z}}_{t}-\tilde{\mathbf{z}}_{t}\right)\)
        else
            Set \(\tilde{\nabla}_{t}=0\).
        Update \(\hat{\mathbf{z}}_{t+1}=\operatorname{Proj}_{\mathcal{K}}\left(\hat{\mathbf{z}}_{t}-\eta_{t} \tilde{\nabla}_{t}\right), \eta_{t}=\frac{1}{t}\).
```

Theorem C.1. Given access to queries from Oracle 2, with stepsizes $\eta_{t}=\frac{1}{t}$, Algorithm 5 enjoys the following regret guarantee:

$$
\begin{equation*}
\operatorname{Regret}_{T}(\text { Algorithm } 5 ; \lambda) \leq \frac{\left(R_{z}+\tilde{R}_{z}\right)^{2}(1+\log T)}{p}+\lambda p T \tag{C.2}
\end{equation*}
$$

Proof of Theorem C.1. To prove the bound in the theorem, we construct the following proxy loss functions: if $b_{t}=1$ denote $\tilde{f}_{t}(\mathbf{z})=\frac{1}{2 p}\left\|\mathbf{z}-\tilde{\mathbf{z}}_{t}\right\|^{2}$, otherwise for $b_{t}=0$ denote $\tilde{f}_{t}(\mathbf{z})=0$. The stochastic gradients of these functions at the current iterate can be written as $\nabla_{\mathbf{z}} \tilde{f}_{t}\left(\hat{\mathbf{z}}_{t}\right)=\tilde{\nabla}_{t}=$ $\frac{\mathbb{I}\left\{b_{t}=1\right\}}{p}\left(\hat{\mathbf{z}}_{t}-\tilde{\mathbf{z}}_{t}\right)$ and are used by the algorithm in the update rule. The idealized gradients are $\nabla_{t}=\left(\hat{\mathbf{z}}_{t}-\mathbf{z}_{t}^{\star}\right)$ which we would use given access to the true targets $\mathbf{z}_{t}^{\star}$. Recall that $\mathcal{F}_{t}$ denotes the sigma-algebra generated by the true target sequence $\mathbf{z}_{1: T}^{\star}$, as well as randomness of the past rounds $b_{1: t-1}$ and $\tilde{\mathbf{z}}_{1: t-1}$. Note then that $\hat{\mathbf{z}}_{t}$ is $\mathcal{F}_{t}$ measurable. We characterize two essential properties of the stochastic gradients:
Lemma C.1. Let $\overline{\mathbf{z}}^{\star} \in \mathcal{K}$ be the minimizer of $\sum_{t=1}^{T}\left\|\mathbf{z}-\mathbf{z}_{t}^{\star}\right\|^{2}$, i.e. empirical average of $\mathbf{z}_{1: T}^{\star}$. Then,

$$
\mathbb{E}\left[\left\langle\tilde{\nabla}_{t}, \hat{\mathbf{z}}_{t}-\overline{\mathbf{z}}^{\star}\right\rangle\right]=\mathbb{E}\left[\left\langle\nabla_{t}, \hat{\mathbf{z}}_{t}-\overline{\mathbf{z}}^{\star}\right\rangle\right] .
$$

Moreover, $\mathbb{E}\left[\left\|\tilde{\nabla}_{t}\right\|^{2}\right] \leq\left(R_{z}+\tilde{R}_{z}\right)^{2} / p$.
Proof. Using the Oracle 2 assumption on $\tilde{\mathbf{z}}_{t}$, we get

$$
\begin{aligned}
\mathbb{E}\left[\tilde{\nabla}_{t} \mid \mathcal{F}_{t}\right] & =\frac{1}{p} \mathbb{E}\left[\mathbb{I}\left\{b_{t}=1\right\}\left(\hat{\mathbf{z}}_{t}-\tilde{\mathbf{z}}_{t}\right) \mid \mathcal{F}_{t}\right] \\
& =\frac{1}{p} \mathbb{E}\left[\mathbb{I}\left\{b_{t}=1\right\} \cdot \hat{\mathbf{z}}_{t} \mid \mathcal{F}_{t}\right]-\frac{1}{p} \mathbb{E}\left[\mathbb{I}\left\{b_{t}=1\right\} \cdot \tilde{\mathbf{z}}_{t} \mid \mathcal{F}_{t}\right] \\
& =\hat{\mathbf{z}}_{t}-\mathbb{E}\left[\tilde{\mathbf{z}}_{t} \mid \mathcal{F}_{t}, b_{t}=1\right] \stackrel{(i)}{=} \hat{\mathbf{z}}_{t}-\mathbf{z}_{t}^{\star}=\nabla_{t}
\end{aligned}
$$

where $(i)$ uses the unbiasedness property of Oracle 2 . Next, since $\overline{\mathbf{z}}^{\star}$ is determined by $\mathbf{z}_{1: T}^{\star}$ it is therefore $\mathcal{F}_{t}$ measurable for all $t$. Thus, $\hat{\mathbf{z}}_{t}-\overline{\mathbf{z}}^{\star}$ is $\mathcal{F}_{t}$ measurable, so

$$
\mathbb{E}\left[\left\langle\tilde{\nabla}_{t}, \hat{\mathbf{z}}_{t}-\overline{\mathbf{z}}^{\star}\right\rangle\right]=\mathbb{E}\left[\left\langle\mathbb{E}\left[\tilde{\nabla}_{t} \mid \mathcal{F}_{t}\right], \hat{\mathbf{z}}_{t}-\overline{\mathbf{z}}^{\star}\right\rangle\right]=\mathbb{E}\left[\left\langle\nabla_{t}, \hat{\mathbf{z}}_{t}-\overline{\mathbf{z}}^{\star}\right\rangle\right] .
$$

Finally, using the norm bound $\left\|\hat{\mathbf{z}}_{t}\right\| \leq R_{z}$ since $\hat{\mathbf{z}}_{t} \in \mathcal{K}$ and the assumption that $\left\|\tilde{\mathbf{z}}_{t}\right\| \leq \tilde{R}_{z}$ from Oracle 2, we conclude

$$
\mathbb{E}\left[\left\|\tilde{\nabla}_{t}\right\|^{2}\right]=\frac{1}{p^{2}} \mathbb{E}\left[\mathbb{I}\left\{b_{t}=p\right\}\left\|\tilde{\mathbf{z}}_{t}-\hat{\mathbf{z}}_{t}\right\|^{2}\right] \leq \frac{1}{p^{2}} \mathbb{E}\left[\mathbb{I}\left\{b_{t}=p\right\}\left(R_{z}+\tilde{R}_{z}\right)^{2}\right]=\frac{\left(R_{z}+\tilde{R}_{z}\right)^{2}}{p}
$$

The rest of the theorem proof mirrors that of Theorem 3.3 in [19] but accounting for the stochastic gradient. We can view Algorithm 5 as running online stochastic gradient descent over strongly convex functions on losses $\frac{1}{2} f_{t}(\mathbf{z})=\frac{1}{2}\left\|\mathbf{z}-\mathbf{z}_{t}^{\star}\right\|^{2}$ with true gradient $\nabla_{t}$ and stochastic gradient $\tilde{\nabla}_{t}$ at the iterate $\hat{\mathbf{z}}_{t}$. Since the losses $\frac{1}{2} f_{t}$ are 1 -strongly convex, $\mathbb{E}\left[b_{t}\right]=p$ and using the claim from Lemma C. 1 we get,

$$
\begin{aligned}
\frac{1}{2} \operatorname{Regret}_{T}=\frac{1}{2} \mathbb{E}\left[\sum_{t=1}^{T}\left(f_{t}\left(\hat{\mathbf{z}}_{t}\right)-f_{t}\left(\overline{\mathbf{z}}^{\star}\right)+\lambda \cdot b_{t}\right)\right] & \leq \mathbb{E}\left[\sum_{t=1}^{T}\left(\left\langle\nabla_{t}, \hat{\mathbf{z}}_{t}-\overline{\mathbf{z}}^{\star}\right\rangle-\frac{1}{2}\left\|\hat{\mathbf{z}}_{t}-\overline{\mathbf{z}}^{\star}\right\|^{2}\right)\right]+\frac{1}{2} \lambda p T \\
& =\frac{1}{2} \mathbb{E}\left[\sum_{t=1}^{T}\left(2\left\langle\tilde{\nabla}_{t}, \hat{\mathbf{z}}_{t}-\overline{\mathbf{z}}^{\star}\right\rangle-\left\|\hat{\mathbf{z}}_{t}-\overline{\mathbf{z}}^{\star}\right\|^{2}\right)\right]+\frac{1}{2} \lambda p T
\end{aligned}
$$

The update rule is given as $\hat{\mathbf{z}}_{t+1}=\operatorname{Proj}_{\mathcal{K}}\left(\hat{\mathbf{z}}_{t}-\eta_{t} \tilde{\nabla}_{t}\right)$, so from the Pythagorean theorem for the projection

$$
\begin{aligned}
\left\|\hat{\mathbf{z}}_{t+1}-\overline{\mathbf{z}}^{\star}\right\|^{2} & \leq\left\|\hat{\mathbf{z}}_{t}-\eta_{t} \tilde{\nabla}_{t}-\overline{\mathbf{z}}^{\star}\right\|^{2}=\left\|\hat{\mathbf{z}}_{t}-\overline{\mathbf{z}}^{\star}\right\|^{2}+\eta_{t}^{2}\left\|\tilde{\nabla}_{t}\right\|^{2}-2 \eta_{t}\left\langle\tilde{\nabla}_{t}, \hat{\mathbf{z}}_{t}-\overline{\mathbf{z}}^{\star}\right\rangle . \\
2\left\langle\tilde{\nabla}_{t}, \hat{\mathbf{z}}_{t}-\overline{\mathbf{z}}^{\star}\right\rangle & \leq \frac{\left\|\hat{\mathbf{z}}_{t}-\overline{\mathbf{z}}^{\star}\right\|^{2}-\left\|\hat{\mathbf{z}}_{t+1}-\overline{\mathbf{z}}^{\star}\right\|^{2}}{\eta_{t}}+\eta_{t}\left\|\tilde{\nabla}_{t}\right\|^{2} .
\end{aligned}
$$

Combining the above bounds results in

$$
\frac{1}{2} \text { Regret }_{T} \leq \frac{1}{2} \mathbb{E}\left[\sum_{t=1}^{T}\left(\frac{\left\|\hat{\mathbf{z}}_{t}-\overline{\mathbf{z}}^{\star}\right\|^{2}-\left\|\hat{\mathbf{z}}_{t+1}-\overline{\mathbf{z}}^{\star}\right\|^{2}}{\eta_{t}}-\left\|\hat{\mathbf{z}}_{t}-\overline{\mathbf{z}}^{\star}\right\|^{2}\right)+\sum_{t=1}^{T} \eta_{t}\left\|\tilde{\nabla}_{t}\right\|^{2}\right]+\frac{1}{2} \lambda p T .
$$

The telescoping sum inside the parentheses is equal to 0 , the gradient term is bounded $\mathbb{E}\left[\left\|\tilde{\nabla}_{t}\right\|^{2}\right] \leq$ $\frac{\left(R_{z}+\tilde{R}_{z}\right)^{2}}{p}$ according to Lemma C. 1 and the stepsize sum is bounded by $\sum_{t=1}^{T} \eta_{t} \leq 1+\log T$, yielding the final result

$$
\operatorname{Regret}_{T}(\text { Algorithm } 5 ; \lambda) \leq \frac{\left(R_{z}+\tilde{R}_{z}\right)^{2}(1+\log T)}{p}+\lambda p T
$$

## C. 1 Adaptive Regret Bound

The guarantee in Theorem C. 1 ensures that the predicted sequence $\hat{\mathbf{z}}_{1: T}$ performs comparably to the empirical mean $\overline{\mathbf{z}}^{\star}$ of the entire target sequence $\mathbf{z}_{1: T}^{\star}$. However, that doesn't imply much about the performance of Algorithm 5 on a given local interval $I \subseteq[T]$ since $\overline{\mathbf{z}}_{I}^{\star}$ can be very different from $\overline{\mathbf{z}}^{\star}$. Hence, we would like to extend our results to hold for any interval $I$, i.e. derive adaptive regret results as introduced in [21]. To do so we will use the approach of [21] using Algorithm 5 as a subroutine. The resulting algorithm, presented in Algorithm 6, suffers only a logarithmic computational overhead over Algorithm 5 with its performance guarantee stated in the theorem below.
Theorem C.2. Taking the base estimation algorithm $\mathcal{A}$ to be Algorithm 5 and given access to queries from Oracle 2, Algorithm 6 enjoys the following guarantee:

$$
\begin{equation*}
\forall I=[r, s] \subseteq[T], \quad \operatorname{Regret}_{I}(\text { Algorithm } 6 ; \lambda) \leq \frac{2\left(R_{z}+\tilde{R}_{z}\right)^{2}(1+\log s \cdot \log |I|)}{p}+\lambda p|I| \tag{C.3}
\end{equation*}
$$

Corollary C.1. The estimation error over each interval $I=[r, s] \subseteq[T]$ is bounded as follows,

$$
\mathbb{E}\left[\sum_{t \in I}\left\|\hat{\mathbf{z}}_{t}-\mathbf{z}_{t}^{\star}\right\|^{2}\right] \leq \operatorname{Var}_{I}\left(\mathbf{z}_{1: T}^{\star}\right)+\frac{2\left(R_{z}+\tilde{R}_{z}\right)^{2}(1+\log s \cdot \log |I|)}{p}
$$

Proof of Theorem C.2. First observe that $\tilde{\ell}_{t}$ is $\alpha$-exp concave with $\alpha=\frac{p}{\left(R_{z}+\tilde{R}_{z}^{2}\right)}$. This is evident given its construction: $\tilde{\ell}_{t}(\mathbf{z})=\frac{1}{2 p}\left\|\mathbf{z}-\tilde{\mathbf{z}}_{t}\right\|^{2}$ with $\|\mathbf{z}\| \leq R_{z}$ since $\mathbf{z} \in \mathcal{K}$ and $\left\|\tilde{\mathbf{z}}_{t}\right\| \leq \tilde{R}_{z}$ according

```
Algorithm 6 Adaptive Estimation Algorithm
    Input: parameter \(p\), decision set \(\mathcal{K}\), base estimation algorithm \(\mathcal{A}, \hat{\mathbf{z}}_{1} \in \mathcal{K}\)
    Initialize: \(\mathcal{A}_{1} \leftarrow \mathcal{A}\left(p, \hat{\mathbf{z}}_{1}\right)\), working set \(\mathcal{S}_{1}=\{1\}, q_{1}^{(1)}=1\), parameter \(\alpha=\frac{p}{\left(R_{z}+\tilde{R}_{z}\right)^{2}}\)
    for \(t=1, \ldots, T\) do
            Compute predictions \(\hat{\mathbf{z}}_{t}^{(i)} \leftarrow \mathcal{A}_{i}\) for \(i \in \mathcal{S}_{t}\)
            Play iterate \(\hat{\mathbf{z}}_{t}=\sum_{i \in \mathcal{S}_{t}} q_{t}^{(i)} \hat{\mathbf{z}}_{t}^{(i)}\)
            Draw/Receive \(b_{t} \sim \operatorname{Bernoulli}(p)\)
            if \(b_{t}=1\) then
                Request estimate \(\tilde{\mathbf{z}}_{t}\) from Oracle 2
                Let \(\tilde{\ell}_{t}(\mathbf{z})=\frac{1}{2 p}\left\|\mathbf{z}-\tilde{\mathbf{z}}_{t}\right\|^{2}\)
            else
                    Let \(\tilde{\mathbf{z}}_{t} \leftarrow \emptyset\) and \(\tilde{\ell}_{t}(\mathbf{z})=0\)
            Update expert algorithms \(\mathcal{A}_{i}\left(b_{t}, \tilde{\mathbf{z}}_{t}\right)\) for all \(i \in \mathcal{S}_{t}\)
            Form new set \(\tilde{\mathcal{S}}_{t+1}=(i)_{i \in \mathcal{S}_{t}}\)
            Construct proxy new weights \(\bar{q}_{t+1}^{(i)}=\frac{t}{t+1} \cdot \frac{q_{t}^{(i)} e^{-\alpha \tilde{\ell}_{t}\left(\mathbf{z}_{t}^{(i)}\right)}}{\sum_{j \in \mathcal{S}_{t}} q_{t}^{(j)} e^{-\alpha \tilde{t}_{t}\left(\hat{\mathbf{z}}_{t}^{(j)}\right)}}\) for all \(i \in \mathcal{S}_{t}\)
            Add new instance \(\tilde{\mathcal{S}}_{t+1} \leftarrow \tilde{\mathcal{S}}_{t+1} \cup t+1\) for arbitrary \(\mathcal{A}_{t+1} \leftarrow \mathcal{A}\left(p, \hat{\mathbf{z}}_{1}^{(t+1)}=\hat{\mathbf{z}}_{1}\right)\) with
    \(\bar{q}_{t+1}^{(t+1)}=\frac{1}{t+1}\)
        Prune \(\tilde{\mathcal{S}}_{t+1}\) to form \(\mathcal{S}_{t+1}\)
        Normalize \(q_{t+1}^{(i)}=\frac{\bar{q}_{t+1}^{(i)}}{\sum_{j \in \mathcal{S}_{t+1}} \bar{q}_{t+1}^{(j)}}\)
```

to Oracle 2. The rest of the algorithm uses the approach of [21], in particular Algorithm 1, over exp concave functions to derive the guarantee in the theorem statement.
We note that Claim 3.1 in [21] holds identically in our case, i.e. for any $I=[r, s]$ the regret of Algorithm 6 with respect to $\mathcal{A}_{r}$ is bounded by $\frac{2}{\alpha}(\ln r+\ln |I|)$ if $\mathcal{A}_{r}$ stays in the working set. We combine this fact with the bound given in Theorem C. 1 to get that Algorithm 6 enjoys regret $\frac{3}{\alpha}(\log r+\log |I|)$ over $I=[r, s]$ if $\mathcal{A}_{r}$ stays in the working set $\mathcal{S}_{t}$ throughout $I$. Finally, an induction argument along with the working set properties detailed in Appendix C.1.1 identical to that of Lemma 3.2 in [21] yields the desired result for $\tilde{\ell}_{t}$. Notice that this is our desired result in expectation,

Observation C.2. We have the following identity for any $t$ and $r \leq t$ :

$$
\mathbb{E}\left[\tilde{\ell}_{t}\left(\hat{\mathbf{z}}_{t}\right)-\tilde{\ell}_{t}\left(\hat{\mathbf{z}}_{t}^{(r)}\right)\right]=\mathbb{E}\left[\ell_{t}\left(\hat{\mathbf{z}}_{t}\right)-\ell_{t}\left(\hat{\mathbf{z}}_{t}^{(r)}\right)\right]
$$

Proof of Observation C.2. We can expand:

$$
\begin{aligned}
2 \mathbb{E}\left[\tilde{\ell}_{t}\left(\hat{\mathbf{z}}_{t}\right)-\tilde{\ell}_{t}\left(\hat{\mathbf{z}}_{t}^{(r)}\right)\right] & =2 p \cdot \mathbb{E}\left[\tilde{\ell}_{t}\left(\hat{\mathbf{z}}_{t}\right)-\tilde{\ell}_{t}\left(\hat{\mathbf{z}}_{t}^{(r)}\right) \mid b_{t}=1\right]+2(1-p) \cdot \mathbb{E}\left[\tilde{\ell}_{t}\left(\hat{\mathbf{z}}_{t}\right)-\tilde{\ell}_{t}\left(\hat{\mathbf{z}}_{t}^{(r)}\right) \mid b_{t}=0\right] \\
& =p \cdot \mathbb{E}\left[\frac{1}{p} \cdot\left(\left\|\hat{\mathbf{z}}_{t}\right\|^{2}+\left\langle\hat{\mathbf{z}}_{t}, \tilde{\mathbf{z}}_{t}\right\rangle+\left\|\tilde{\mathbf{z}}_{t}\right\|^{2}-\left\|\hat{\mathbf{z}}_{t}^{(r)}\right\|^{2}-\left\langle\hat{\mathbf{z}}_{t}^{(r)}, \tilde{\mathbf{z}}_{t}\right\rangle-\left\|\tilde{\mathbf{z}}_{t}\right\|^{2}\right)\right]+0 \\
& =\mathbb{E}\left[\left\|\hat{\mathbf{z}}_{t}\right\|^{2}-\left\|\hat{\mathbf{z}}_{t}^{(r)}\right\|^{2}\right]+\mathbb{E}\left[\left\langle\hat{\mathbf{z}}_{t}-\hat{\mathbf{z}}_{t}^{(r)}, \tilde{\mathbf{z}}_{t}\right\rangle\right]
\end{aligned}
$$

By the linearity of expectation, the fact that $\hat{\mathbf{z}}_{t}, \hat{\mathbf{z}}_{t}^{(r)}$ are completely determined given $\mathcal{F}_{t-1}$, and the law of total expectation we have that

$$
\begin{aligned}
\mathbb{E}\left[\left\langle\hat{\mathbf{z}}_{t}-\hat{\mathbf{z}}_{t}^{(r)}, \tilde{\mathbf{z}}_{t}\right\rangle\right] & =\mathbb{E}\left[\mathbb{E}\left[\left\langle\hat{\mathbf{z}}_{t}-\hat{\mathbf{z}}_{t}^{(r)}, \tilde{\mathbf{z}}_{t}\right\rangle \mid \mathcal{F}_{t-1}\right]\right] \\
& =\mathbb{E}\left[\left\langle\hat{\mathbf{z}}_{t}-\hat{\mathbf{z}}_{t}^{(r)}, \mathbf{z}_{t}^{\star}\right\rangle\right]
\end{aligned}
$$

Plugging this in above we, adding and subtracting $\left\|\mathbf{z}_{t}^{\star}\right\|^{2}$, and rearranging we have:

$$
\begin{aligned}
2 \mathbb{E}\left[\tilde{\ell}_{t}\left(\hat{\mathbf{z}}_{t}\right)-\tilde{\ell}_{t}\left(\hat{\mathbf{z}}_{t}^{(r)}\right)\right] & =\mathbb{E}\left[\left\|\hat{\mathbf{z}}_{t}\right\|^{2}-\left\|\hat{\mathbf{z}}_{t}^{(r)}\right\|^{2}+\left\langle\hat{\mathbf{z}}_{t}-\hat{\mathbf{z}}_{t}^{(r)}, \mathbf{z}_{t}^{\star}\right\rangle+\|\mathbf{z}\|^{2}-\|\mathbf{z}\|^{2}\right] \\
& =\mathbb{E}\left[\left\|\hat{\mathbf{z}}_{t}\right\|^{2}+\left\langle\hat{\mathbf{z}}_{t}, \mathbf{z}_{t}^{\star}\right\rangle+\|\mathbf{z}\|^{2}-\left\|\hat{\mathbf{z}}_{t}^{(r)}\right\|^{2}-\left\langle\hat{\mathbf{z}}_{t}^{(r)}, \mathbf{z}_{t}^{\star}\right\rangle-\|\mathbf{z}\|^{2}\right] \\
& =2 \mathbb{E}\left[\ell_{t}\left(\hat{\mathbf{z}}_{t}\right)-\ell_{t}\left(\hat{\mathbf{z}}_{t}^{(r)}\right)\right]
\end{aligned}
$$

as desired.
Combining Observation C. 2 with the fact that $\tilde{\ell}_{t}$ are $\alpha$-exp concave for $\alpha=\frac{p}{\left(R_{z}+\tilde{R}_{z}\right)^{2}}$ we conclude the final statement of Theorem C.2.

## C.1.1 Working Set Construction

Our Algorithm 6 makes use of the working sets $\left\{\mathcal{S}_{t}\right\}_{t \in[T]}$ along with its properties in Claim C.3. In this section, we show the explicit construction of these working sets as in [21] and prove the claim.
Claim C.3. The following properties hold for the working sets $S_{t}$ for all $t \in[T]:(i)\left|\mathcal{S}_{t}\right|=O(\log T)$; (ii) $[s,(s+t) / 2] \cap \mathcal{S}_{t} \neq \emptyset$ for any $s \in[t]$; (iii) $\mathcal{S}_{t+1} \backslash \mathcal{S}_{t}=\{t+1\}$; (iv) $\left|\mathcal{S}_{t} \backslash \mathcal{S}_{t+1}\right| \leq 1$.

For any $i \in[T]$, let it be given as $i=r 2^{k}$ with $r$ odd and $k$ nonnegative. Denote $m=2^{k+2}+1$, then $i \in S_{t}$ if and only if $t \in[i, i+m]$. This fully describes the construction of the working sets $\left\{S_{t}\right\}_{t \in[T]}$, and we proceed to prove its properties.

Proof of Claim C.3. For all $t \in[T]$ we show the following properties of the working sets $S_{t}$.
(i) $\left|S_{t}\right|=O(\log T)$ : if $i \in S_{t}$ then $1 \leq i=r 2^{k} \leq t$ which implies that $0 \leq k \leq \log _{2} t$. For each fixed $k$ in this range, if $r 2^{k}=i \in S_{t}$ then $i \in\left[t-2^{k+2}-1, t\right]$ by construction. Since $\left[t-2^{k+2}-1, t\right]$ is an interval of length $2^{k+2}+2=4 \cdot 2^{k}+2$, it can include at most 3 numbers of the form $r 2^{k}$ with $r$ odd. Thus, there is at most 3 numbers $i=r 2^{k} \in S_{t}$ for each $0 \leq k \leq \log _{2} t$ which means that $\left|S_{t}\right|=O(\log t)=O(\log T)$.
(ii) $[s,(s+t) / 2] \cap S_{t} \neq \emptyset$ for all $s \in[t]$ : this trivially holds for $s=t-1, t$. Let $2^{l} \leq(t-s) / 2$ be the largest such exponent of 2 . Since the size of the interval $[s,(s+t) / 2]$ is $\lfloor(t-s) / 2\rfloor$, then there exists $u \in[s,(s+t) / 2]$ that divides $2^{l}$. This means that the corresponding $m \geq 2^{l+2}+1>t-s$ for $u \geq s$ is large enough so that $t \in[u, u+m]$, and consequently, $u \in S_{t}$.
(iii) $S_{t+1} \backslash S_{t}=\{t+1\}$ : let $i \in S_{t+1}$ and $i \notin S_{t}$, which is equivalent to $t+1 \in[i, i+m]$ and $t \notin[i, i+m]$. Clearly, $i=t+1$ satisfies these conditions and is the only such number.
(iv) $\left|S_{t} \backslash S_{t+1}\right| \leq 1$ : suppose there exist two $i_{1}, i_{2} \in S_{t} \backslash S_{t+1}$. This implies that $i_{1}+m_{1}=t=$ $i_{2}+m_{2}$ which in turn means $2^{k_{1}}\left(r_{1}+4\right)=2^{k_{2}}\left(r_{2}+4\right)$. Since both $r_{1}+4, r_{2}+4$ are odd, then $k_{1}=k_{2}$, and consequently, $r_{1}=r_{2}$ resulting in $i_{1}=i_{2}$. Thus, there can not exist two different members of $S_{t} \backslash S_{t+1}$ which concludes that $\left|S_{t} \backslash S_{t+1}\right| \leq 1$.

## C. 2 No Strong Adaptivity

Notice that even though the guarantee of Theorem C. 2 applies to all intervals $I$, it does not entail meaningful guarantees for all. The reason is the choice of parameter $p$ : if one wishes to optimize Regret $_{T}$ then $p=\mathcal{O}\left(T^{-1 / 2}\right)$ implies $\mathcal{O}(\sqrt{T})$ regret, but this choice is meaningless for intervals with length $|I| \ll \sqrt{T}$; on the other hand, optimizing the bound for small intervals leads to large bounds for the entire horizon. One might then ask whether there exist methods with strongly adaptive guarantees, and we answer this question with a negative.
Theorem C.3. For any $\gamma>0$ and oracle cost $\lambda>0$, there exists no online algorithm $\mathcal{A}$ with feedback access to Oracle 2 that enjoys the following strongly adaptive regret guarantee: $\operatorname{Regret}_{I}(\mathcal{A} ; \lambda)=$ $\tilde{\mathcal{O}}\left(|I|^{1-\gamma}\right)$.

Proof. The proof of this impossibility results follows a simple construction: the idea behind it is that strongly adaptive guarantees imply both large and small amount of exploration. Let us suppose
there exists such an algorithm $\mathcal{A}$ and arrive at a contradiction: $\forall I=[r, s] \subseteq[T]$ algorithm $\mathcal{A}$ has a regret bound $\operatorname{Regret}_{I}(\mathcal{A} ; \lambda) \leq C \cdot|I|^{1-\gamma}=\tilde{\mathcal{O}}\left(|I|^{1-\gamma}\right)$ over any oblivious sequence $\mathbf{z}_{1: T}^{\star}$ where $C$ depends on problem parameters, $\lambda$ and $\log T$.
Construct the following oblivious sequence: let $k=T^{1-\gamma / 2}$ and $I_{1}, \ldots, I_{k}$ be consecutive disjoint intervals such that $\cup_{j \in[k]} I_{j}=[T], I_{j} \cap I_{l}=\emptyset$ for all $j \neq l$, and $\left|I_{j}\right|=T / k=T^{\gamma / 2}$ for all $j \in[k]$ (w.l.o.g. we assume $T$ divides $k$ ). Now for each interval $I_{j}, j \in[k]$, sample a fresh $q_{j} \in\{ \pm 1\} \sim \operatorname{Rad}(1 / 2)$ and let $\mathbf{z}_{t}^{\star}=q_{j}$ for all $t \in I_{j}$.
According to the assumed guarantee, the overall regret is bounded as $\operatorname{Regret}_{T}(\mathcal{A} ; \lambda) l e q C \cdot T^{1-\gamma}$ which by definition implies that $\sum_{t=1}^{T} b_{t} \leq \frac{C}{\lambda} T^{1-\gamma}<k$ where the last inequality is true for sufficiently large horizon $T$. Since there are $k$ consecutive disjoint intervals $I_{1}, \ldots, I_{k}$ and less than $k$ overall calls to Oracle 2, there exists an interval $I \in\left\{I_{1}, \ldots, I_{k}\right\}$ such that $\sum_{t \in I} b_{t}=0$.
On the other hand, the assumed guarantee for $\mathcal{A}$ implies that the interval $I$ of size $|I|=T^{\gamma / 2}$ enjoys sublinear regret, i.e. $\operatorname{Regret}_{I}(\mathcal{A} ; \lambda)=o(|I|)$. We show that this is a contradiction given that $\sum_{t \in I} b_{t}=0$. As there were no oracle calls for the interval $I$, the predictions of $\mathcal{A}, \hat{\mathbf{z}}_{t}$ over $t \in I$, are independent from the Rademacher sample of the interval $q_{I}$ : this is true since the samples for each interval in $I_{1}, \ldots, I_{k}$ are independent. Therefore, $\hat{\mathbf{z}}_{t} \perp q_{I}$ for all $t \in I$ which means that since the best loss in hindsight over $I$ is equal to 0 as $\mathbf{z}_{t}^{\star}=q_{I}$ for all $t \in I$,

$$
\operatorname{Regret}_{I}(\mathcal{A} ; \lambda) \geq \mathbb{E}_{q_{I}}\left[\sum_{t \in I} \ell_{t}\left(\hat{\mathbf{z}}_{t}\right)\right]=\sum_{t \in I} \mathbb{E}_{q_{I}}\left[\left\|\hat{\mathbf{z}}_{t}-q_{I}\right\|^{2}\right]=\Omega(|I|)
$$

Hence, for the interval $I$, the regret of $\mathcal{A}$ cannot be sublinear, which contradict the assumption that $\mathcal{A}$ exhibits strongly adaptive guarantees. This concludes that no strongly adaptive online algorithm exists in the described partial information model.

## D Adaptive Regret for Control of Changing Unknown Dynamics

In this section we give our full control algorithm which attains sublinear regret with respect to $\Pi_{d r c}$ up to an additive system variability term. A key component is the system estimation for which we will use Algorithm 6 and its guarantees from Appendix C. More specifically, our algorithm is based on the canonical explore-exploit approach: it explores with some probability $p$ by inputting random controls into the system, and otherwise outputs a control according to DRC-OGD (Algorithm 3). Note that due to the long-term consequences which appear in control, we need to explore for $h$ consecutive steps in order to get an estimate for the $h$-truncation of the Markov operator. Hence, our algorithm will determine whether it explores or exploits in blocks of length $h$. Furthermore, we will define the set of Markov operator of length $h$ and $\ell_{1}$,op-norm bounded by $R_{G}$ as $\mathcal{G}\left(h, R_{G}\right)$ :

$$
\mathcal{G}\left(h, R_{G}\right) \doteq\left\{G=\left(G^{[0]}, \ldots, G^{[h-1]}\right) \in \mathbb{R}^{h \times d_{x} \times d_{u}} \text { s.t. } \sum_{i=0}^{h-1}\|G\|_{o p} \leq R_{G}\right\}
$$

Remark D.1. Note that the radius of $\mathcal{G}\left(h, R_{G}\right)$ is bounded by $\bar{R}_{G}=\sqrt{h \cdot d_{\min }} R_{G}$ where $d_{\text {min }}=$ $\min \left\{d_{x}, d_{u}\right\}$.

Proof of Remark D.1. $\forall G \in \mathcal{G}\left(h, R_{G}\right)$, we have:

$$
\begin{aligned}
\|G\|_{F} & =\sqrt{\sum_{i=0}^{h-1}\left\|G^{[i]}\right\|_{F}^{2}} \\
& \leq \sqrt{h} \sqrt{\max _{i}\left\|G^{[i]}\right\|_{F}^{2}} \\
& \leq \sqrt{h \min \left\{d_{x}, d_{u}\right\}} \max _{i}\left\|G^{[i]}\right\|_{o p} \\
& \leq \sqrt{h \min \left\{d_{x}, d_{u}\right\}} R_{G}
\end{aligned}
$$

and denoting $d_{\text {min }}=\min \left\{d_{x}, d_{u}\right\}$ yields the result.

Remark D.2. By abuse of notation, we will consider $G^{[i]} \doteq 0$ for $G \in \mathcal{G}\left(h, R_{G}\right)$.
Remark D.3. For simplicity, we assume $T$ divisible by $h$ (this is w.l.o.g. up to an extra $O(h)$ cost which for us is negligble).

We spell out the full procedure in Algorithm 7 and give its guarantee below in Theorem D.1.

```
Algorithm 7 DRC-OGD with Exploration
    Input: \(p, h, \hat{G}_{0}, \mathcal{A} \leftarrow\) Algorithm \(6\left(p, \hat{G}_{0}, \mathcal{G}\left(h, R_{G}\right)\right), \mathcal{C} \leftarrow\) Algorithm \(3\left(\eta, m, R_{\mathcal{M}}\right)\)
    for \(\tau_{1}=0, \ldots, T / h-1\) do
        Request \(\hat{G}_{\tau_{1} \cdot h+1} \leftarrow \mathcal{A}\) and set \(\hat{G}_{\tau_{1} \cdot h+2}, \ldots, \hat{G}_{\left(\tau_{1}+1\right) \cdot h} \leftarrow \hat{G}_{\tau_{1} \cdot h+1}\)
        Draw \(b_{\tau_{1}+1} \sim \operatorname{Bernoulli}(p)\)
        for \(\tau_{2}=1, \ldots, h\) do \(\quad \triangleright\) let \(t \doteq \tau_{1} \cdot h+\tau_{2}\)
            if \(b_{\tau_{1}}=1\) then
                    Play control \(u_{t} \sim\{ \pm 1\}^{d_{u}}\)
            else
                    Play control \(u_{t} \leftarrow \mathcal{C}\)
            Suffer cost \(c_{t}\left(x_{t}, u_{t}\right)\), observe new state \(x_{t+1}\)
            Extract \(\hat{x}_{t+1}^{\text {nat }}=\operatorname{Proj}_{\mathbb{B}_{R_{\text {nat }}}}\left[x_{t+1}-\sum_{i=0}^{h-1} \hat{G}_{t}^{[i]} u_{t-i}\right]\)
            Update \(\mathcal{C} \leftarrow\left(c_{t}, \hat{G}_{t}, \hat{x}_{t+1}^{\text {nat }}\right)\)
        if \(b_{\tau_{1}}=1\) then
            Let \(\tilde{G}_{\left(\tau_{1}+1\right) \cdot h}^{[i]}=x_{\left(\tau_{1}+1\right) \cdot h+1} u_{\left(\tau_{1}+1\right) \cdot h-i}^{\top}\) for \(i=\overline{0, h-1}\)
        else
            Let \(\tilde{G}_{\left(\tau_{1}+1\right) \cdot h} \leftarrow \emptyset\)
        Update \(\mathcal{A} \leftarrow\left(b_{\tau_{1}+1}, \tilde{G}_{\left(\tau_{1}+1\right) \cdot h}\right)\)
```

Theorem D.1. For $h=\frac{\log T}{\log \rho^{-1}}, p=T^{-1 / 3}$ and $m \leq \sqrt{T}$, on any contiguous interval $I \subseteq[T]$, Algorithm 2 enjoys the following adaptive regret guarantee ${ }^{2}$ :

$$
\mathbb{E}\left[\operatorname{Regret}_{I}(\text { ADA-CTRL }) ; \Pi_{\mathrm{drc}}\left(m, R_{M}\right)\right] \leq \widetilde{\mathcal{O}}^{\star}\left(\operatorname{Lm}\left(|I| \sqrt{\mathbb{E}\left[\operatorname{Var}_{I}(\mathbf{G})\right]}+d_{u} T^{2 / 3}\right)\right)
$$

The proof of this theorem will proceed in terms of a quantity which we call total system variability which captures the total (rather than average) deviation from the mean operator for each interval. More precisely,
Definition D.1. Define the total system variability of an LTV dynamical system with Markov operators $\mathbf{G}=G_{1: T}$ over a contiguous interval $I \subseteq[T]$ to be

$$
\operatorname{Var}_{I}^{\mathrm{tot}}(\mathbf{G})=\min _{G} \sum_{t \in I}\left\|G-G_{t}\right\|_{\ell_{2}, F}^{2}=\sum_{t \in I}\left\|G_{I}-G_{t}\right\|_{\ell_{2}, F}^{2}
$$

where $\|\cdot\|_{\ell_{2}, F}$ indicates the $\ell_{2}$ norm of the fully vectorized operator and $G_{I}=|I|^{-1} \sum_{t \in I} G_{t}$ is the empirical average of the operators that correspond to $I$.

## D. 1 Estimation of the Markov Operator

Note that the estimation component of Algorithm 7 directly operates in the setting of Appendix C and effectively solves the problem of adaptively estimating the sequence $\bar{G}_{1 \cdot h}, \ldots, \bar{G}_{(T / h) \cdot h}$ of the $h$-truncations of the true Markov operators $G_{1 \cdot h}, \ldots, G_{(T / h) \cdot h}$. To formally be able to apply Theorem C.2, we first show that the estimates sent to Algorithm 6 satisfy the properties of Oracle 2.
Claim D.1. The estimators $\tilde{G}_{\tau_{1} \cdot h}, \tau_{1}=\overline{1, T / h}$ satisfy the properties of Oracle 2 with $\tilde{R}_{G}=$ $\sqrt{h \cdot d_{u}}\left(R_{\text {nat }}+R_{G} \max \left\{\sqrt{d_{u}}, R_{\text {nat }} R_{\mathcal{M}}\right\}\right)$.

Proof. There are only two things to prove:

[^1]1. Boundedness. Because we clip the nature's x estimates $\hat{x}_{t}^{\text {nat }}$ to $R_{\text {nat }}$ and when the DAC policy lies in $\mathcal{M}\left(m, R_{\mathcal{M}}\right)$, we have that if $b_{t}=0,\left\|u_{t}\right\| \leq R_{\text {nat }} R_{\mathcal{M}}$. If $u_{t}$ is an exploratory action then $\left\|u_{t}\right\|=\sqrt{d_{u}} \leq \max \left\{\sqrt{d_{u}}, R_{\text {nat }} R_{\mathcal{M}}\right\}$ by design. By the equation of the progression of the state, we have that

$$
\begin{equation*}
\left\|x_{t}\right\| \leq R_{\text {nat }}+R_{G} \max \left\{\sqrt{d_{u}}, R_{\text {nat }} R_{\mathcal{M}}\right\} \tag{D.1}
\end{equation*}
$$

and by Cauchy Scwarz we get

$$
\left\|\tilde{G}_{\tau_{1} \cdot h}^{[i]}\right\|_{F}=\left\|u_{\tau_{1} \cdot h-i}^{\top} x_{\tau_{1} \cdot h+1}\right\| \leq \sqrt{d_{u}}\left(R_{\mathrm{nat}}+R_{G} \max \left\{\sqrt{d_{u}}, R_{\mathrm{nat}} R_{\mathcal{M}}\right\}\right)
$$

and hence

$$
\left\|\tilde{G}_{\tau_{1} \cdot h}\right\|_{F} \leq \sqrt{h \cdot d_{u}}\left(R_{\mathrm{nat}}+R_{G} \max \left\{\sqrt{d_{u}}, R_{\mathrm{nat}} R_{\mathcal{M}}\right\}\right)
$$

2. Unbiasedness. Plugging in $x_{\tau_{1} \cdot h}=x_{\tau_{1} \cdot h}^{\text {nat }}+\sum_{i=0}^{t} G_{\tau_{1} \cdot h}^{[i]} u_{\tau_{1} \cdot h-i}$, we get exactly that $\mathbb{E}\left[\tilde{G}_{\tau_{1} \cdot h}^{[i]}\right]=G_{\tau_{1} \cdot h}^{[i]}$. Since this holds for the selected truncation $h$, we have $\mathbb{E}\left[\tilde{G}_{\tau_{1} \cdot h}\right]=\bar{G}_{\tau_{1} \cdot h}$ for $G$ as defined.

Hence we can simply apply the guarantees of Appendix C to obtain the following guarantee:
Corollary D.1. On any interval $J=[k, l] \subseteq[T / h]$, we have that:

$$
\mathbb{E}\left[\sum_{\tau=k}^{l}\left\|\tilde{G}_{\tau \cdot h}-\bar{G}_{\tau \cdot h}\right\|_{F}^{2}\right] \leq \operatorname{Var}_{J \cdot h}^{\mathrm{tot}}\left(\bar{G}_{1: T}\right)+\operatorname{Regret}_{J}\left(\mathcal{A}\left(p, \bar{R}_{G}, \tilde{R}_{G} ; 0\right)\right)
$$

where we use $J \cdot h$ to denote the set $[k \cdot h, \ldots, l \cdot h]$.
However, to properly analyze the additional regret introduced in the control framework by our estimation error, we need to convert Corollary D. 1 into a guarantee in terms of $\ell_{1}$, op norm which holds for any contiguous interval $I=[r, s] \subseteq[T]$. This step is rather straightforward and only relies on a few basic properties/observations which we collect in Observation D. 2 below.
Observation D.2. We will make the following observations:

1. $\|A\|_{o p} \leq\|A\|_{F}$ for any matrix $A$,
2. $\operatorname{Var}_{J}^{\text {tot }}\left(\mathbf{z}_{1: T}\right) \leq \operatorname{Var}_{I}^{\text {tot }}\left(\mathbf{z}_{1: T}\right)$ for any set of indices $J \subseteq I$, and any sequence $\mathbf{z}_{1: T}$,
3. $\left\|\hat{G}_{t}-\bar{G}_{t}\right\|_{F}^{2}=\sum_{i=1}^{h}\left\|\hat{G}_{t}^{[i]}-\bar{G}_{t}^{[i]}\right\|_{F}^{2}$,
4. $\left\|G_{t}-\bar{G}_{t}\right\|_{\ell_{1}, o p} \leq \psi(h)$,
5. For any $I=[r, s], \sum_{t=r+1}^{s}\left\|\bar{G}_{t}-\bar{G}_{t-1}\right\|_{F}^{2} \leq 4 \operatorname{Var}_{I}^{\mathrm{tot}}\left(\bar{G}_{1: T}\right) \leq 4 \operatorname{Var}_{I}^{\mathrm{tot}}(\mathbf{G})$.

Proof. Properties 1-4 follow from the definitions of the relevant quantities, or are general well-known facts. For 5 , by the triangle inequality, we have that, for any $\bar{G}$,

$$
\begin{gathered}
\left\|\bar{G}_{t}-\bar{G}_{t-1}\right\|_{F} \leq\left\|\bar{G}_{t}-\bar{G}\right\|_{F}+\left\|\bar{G}_{t-1}-\bar{G}\right\|_{F} \\
\Rightarrow\left\|\bar{G}_{t}-\bar{G}_{t-1}\right\|_{F}^{2} \leq 2\left(\left\|\bar{G}_{t}-\bar{G}\right\|_{F}^{2}+\left\|\bar{G}_{t-1}-\bar{G}\right\|_{F}^{2}\right)
\end{gathered}
$$

summing the above from $r+1$ to $s$ and taking $\bar{G}$ to be the sample mean over $I$ yields the desired result. The second inequality simply follows by the fact that truncation can only decrease variance.

As a first step, we first bound the Frobenius norm error of the truncated operators over an arbitrary contiguous interval $I=[r, s] \subseteq[T]$.

Lemma D.3. On any interval $I=[r, s] \subseteq[T]$ with $r-1, s$ divisible by $h^{3}$, we can bound the Frobenius estimation error of the truncated operators as:

$$
\mathbb{E}\left[\sum_{t=r}^{s}\left\|\hat{G}_{t}-\bar{G}_{t}\right\|_{F}^{2}\right] \leq 10 h^{2} \operatorname{Var}_{I}^{\mathrm{tot}}(\mathbf{G})+2 h \operatorname{Regret}_{I}\left(\mathcal{A}\left(p, \bar{R}_{G}, \tilde{R}_{G} ; 0\right)\right)
$$

Proof. By Algorithm 7, we can write

$$
\begin{aligned}
\sum_{t=r}^{s}\left\|\hat{G}_{t}-\bar{G}_{t}\right\|_{F}^{2} & =\sum_{\tau_{1}=(r-1) / h}^{s / h-1} \sum_{\tau_{2}=1}^{h}\left\|\hat{G}_{\tau_{1} \cdot h+\tau_{2}}-\bar{G}_{\tau_{1} \cdot h+\tau_{2}}\right\|_{F}^{2} \\
& =\sum_{\tau_{1}=(r-1) / h}^{s / h-1} \sum_{\tau_{2}=1}^{h}\left\|\hat{G}_{\left(\tau_{1}+1\right) \cdot h}-\bar{G}_{\tau_{1} \cdot h+\tau_{2}}\right\|_{F}^{2} \\
& \leq 2 \sum_{\tau_{1}=(r-1) / h}^{s / h-1} \sum_{\tau_{2}=1}^{h}\left(\left\|\hat{G}_{\left(\tau_{1}+1\right) \cdot h}-\bar{G}_{\left(\tau_{1}+1\right) \cdot h}\right\|_{F}^{2}+\left\|\bar{G}_{\left(\tau_{1}+1\right) \cdot h}-\bar{G}_{\tau_{1} \cdot h+\tau_{2}}\right\|_{F}^{2}\right) \\
& =2 h \underbrace{s / h}_{\text {(Algorithm } 6 \text { estimation error) }}\left\|\hat{G}_{\tau \cdot h}^{s / h}-\bar{G}_{\tau \cdot h}\right\|_{F}^{2}
\end{aligned} \underbrace{\sum_{\tau_{1}=(r-1) / h}^{s / h-1} \sum_{\tau_{2}=1}^{h}\left\|\bar{G}_{\left(\tau_{1}+1\right) \cdot h}-\bar{G}_{\tau_{1} \cdot h+\tau_{2}}\right\|_{F}^{2}}_{(\bar{G} \text { movement) }}
$$

For the first term, we will simply apply the above corollary after taking expectation. We will therefore focus on bounding the $\bar{G}$ movement term.

$$
\begin{align*}
\forall \tau_{2}: \quad\left\|\bar{G}_{\left(\tau_{1}+1\right) \cdot h}-\bar{G}_{\tau_{1} \cdot h+\tau_{2}}\right\|_{F}^{2} & \leq\left(\sum_{i=1}^{h-1}\left\|\bar{G}_{\tau_{1} \cdot h+i+1}-\bar{G}_{\tau_{1} \cdot h+i}\right\|_{F}\right)^{2} \\
& \leq h \sum_{i=1}^{h-1}\left\|\bar{G}_{\tau_{1} \cdot h+i+1}-\bar{G}_{\tau_{1} \cdot h+i}\right\|_{F}^{2} \tag{C.S.}
\end{align*}
$$

This implies that

$$
\sum_{\tau_{2}=1}^{h}\left\|\bar{G}_{\left(\tau_{1}+1\right) \cdot h}-\bar{G}_{\tau_{1} \cdot h+\tau_{2}}\right\|_{F}^{2} \leq h^{2} \sum_{i=1}^{h-1}\left\|\bar{G}_{\tau_{1} \cdot h+i+1}-\bar{G}_{\tau_{1} \cdot h+i}\right\|_{F}^{2}
$$

So finally we have that

$$
\begin{aligned}
\bar{G} \text {-movement } & \leq h^{2} \sum_{t=r+1}^{s}\left\|\bar{G}_{t}-\bar{G}_{t}\right\|_{F}^{2} \\
& \leq 4 h^{2} \operatorname{Var}_{I}^{\text {tot }}(\mathbf{G})
\end{aligned}
$$

Observation D. 2 (5)
Finally, taking expectation, plugging in Corollary D. 1 and noting that for $J={ }_{\tilde{R}}[r / h, s / h]$, $\operatorname{Var}_{J \cdot h}^{\text {tot }}\left(\bar{G}_{1: T}\right) \leq \operatorname{Var}_{I}^{\text {tot }}(\mathbf{G})$ (Observation D. 2 (2)) and $\operatorname{Regret}_{J}\left(\mathcal{A}\left(p, \bar{R}_{G}, \tilde{R}_{G} ; 0\right)\right) \leq$ $\operatorname{Regret}_{I}\left(\mathcal{A}\left(p, \bar{R}_{G}, \tilde{R}_{G} ; 0\right)\right)$, we get:

$$
\mathbb{E}\left[\sum_{t=r}^{s}\left\|\hat{G}_{t}-\bar{G}_{t}\right\|_{F}^{2}\right] \leq 10 h^{2} \mathbb{E}\left[\operatorname{Var}_{I}^{\mathrm{tot}}(\mathbf{G})\right]+2 h \operatorname{Regret}_{I}\left(\mathcal{A}\left(p, \bar{R}_{G}, \tilde{R}_{G} ; 0\right)\right)
$$

[^2]Lemma D.4. On any interval $I=[r, s] \subseteq[T]$ with $r-1, s$ divisible by $h$, we can bound the squared $\ell_{1}$, op estimation error of the truncated operators as:

$$
\mathbb{E}\left[\sum_{t=r}^{s}\left\|\hat{G}_{t}-\bar{G}_{t}\right\|_{\ell_{1}, o p}^{2}\right] \leq 20 h^{3} \mathbb{E}\left[\operatorname{Var}_{I}^{\mathrm{tot}}(\mathbf{G})\right]+4 h^{2} \operatorname{Regret}_{I}\left(\mathcal{A}\left(p, \bar{R}_{G}, \tilde{R}_{G} ; 0\right)\right)
$$

Proof. We have that

$$
\begin{array}{rlr}
\sum_{t=r}^{s}\left\|\hat{G}_{t}-\bar{G}_{t}\right\|_{\ell_{1}, o p}^{2} & \leq \sum_{t=r}^{s}\left(\sum_{i=0}^{h-1}\left\|\hat{G}_{t}^{[i]}-\bar{G}_{t}^{[i]}\right\|_{o p}\right)^{2} & \Delta \text {-ineq. } \\
& \leq 2 \sum_{t=r}^{s}\left(\sum_{i=0}^{h-1}\left\|\hat{G}_{t}^{[i]}-\bar{G}_{t}^{[i]}\right\|_{o p}\right)^{2} & \\
& \leq 2 h \sum_{t=r}^{s} \sum_{i=0}^{h-1}\left\|\hat{G}_{t}^{[i]}-\bar{G}_{t}^{[i]}\right\|_{F}^{2} & \text { Observation D.2 (1) \& C.S. } \\
& =2 h \sum_{t=r}^{s}\left\|\hat{G}_{t}-\bar{G}_{t}\right\|_{F}^{2} & \text { Observation D.2 (3) } \tag{3}
\end{array}
$$

Taking expectation and plugging in the bound in Lemma D. 3 yields the promised result.

Finally we can use Lemma D. 4 and Cauchy-Schwarz to get a result in terms of the linear (rather than squared) $\ell_{1}$, op error accumulated over an interval:

Proposition D.5. On any interval $I=[r, s] \subseteq[T]$ with $r-1, s$ divisible by $h$, we can bound the squared $\ell_{1}$, op estimation error of the truncated operators as:

$$
\mathbb{E}\left[\sum_{t=r}^{s}\left\|\hat{G}_{t}-\bar{G}_{t}\right\|_{\ell_{1}, o p}\right] \leq 5 h^{3 / 2}|I|^{1 / 2} \sqrt{\operatorname{Var}_{I}^{\mathrm{tot}}(\mathbf{G})}+2 h|I|^{1 / 2} \operatorname{Regret}_{I}^{1 / 2}\left(\mathcal{A}\left(p, \bar{R}_{G}, \tilde{R}_{G} ; 0\right)\right)
$$

Proof. By Cauchy-Schwarz and Jensen (since $\sqrt{ }$. is concave) we have:

$$
\begin{aligned}
\mathbb{E}\left[\sum_{t=r}^{s}\left\|\hat{G}_{t}-\bar{G}_{t}\right\|_{\ell_{1}, o p}\right] & \leq|I|^{1 / 2} \mathbb{E}\left[\left(\sum_{t=r}^{s}\left\|\hat{G}_{t}-\bar{G}_{t}\right\|_{\ell_{1}, o p}^{2}\right)^{1 / 2}\right] \\
& \leq|I|^{1 / 2}\left(\mathbb{E}\left[\sum_{t=r}^{s}\left\|\hat{G}_{t}-\bar{G}_{t}\right\|_{\ell_{1}, o p}^{2}\right]\right)^{1 / 2} \\
& \leq|I|^{1 / 2} \sqrt{20 h^{3} \mathbb{E}\left[\operatorname{Var}_{I}^{\mathrm{tot}}(\mathbf{G})\right]+4 h^{2} \operatorname{Regret}_{I}\left(\mathcal{A}\left(p, \bar{R}_{G}, \tilde{R}_{G} ; 0\right)\right)} \\
& \leq 5 h^{3 / 2}|I|^{1 / 2} \sqrt{\mathbb{E}\left[\operatorname{Var}_{I}^{\mathrm{tot}}(\mathbf{G})\right]}+2 h|I|^{1 / 2} \operatorname{Regret}_{I}^{1 / 2}\left(\mathcal{A}\left(p, \bar{R}_{G}, \tilde{R}_{G} ; 0\right)\right)
\end{aligned}
$$

## D. 2 Error Sensitivity

We now analyze concretely how the $G_{t}$ estimation errors induce additional regret over the case of known systems. We can decompose the expected regret over an interval $I=[r, s]$ as:

$$
\begin{aligned}
\operatorname{Regret}_{I}= & \sum_{t=r}^{s} c_{t}\left(x_{t}, u_{t}\right)-\min _{\pi \in \Pi_{\mathrm{drc}}} \sum_{t=r}^{s} c_{t}\left(x_{t}^{\pi}, u_{t}^{\pi}\right) \\
= & \underbrace{\sum_{t=r}^{s} c_{t}\left(x_{t}, u_{t}\right)-c_{t}\left(\hat{x}_{t}\left(M_{t-h: t-1}\right), \hat{u}_{t}\left(M_{t}\right)\right)}_{\text {(realized iterate error) }} \\
& +\underbrace{\sum_{t=r}^{s} c_{t}\left(\hat{x}_{t}\left(M_{t-h: t-1}\right), \hat{u}_{t}\left(M_{t}\right)-\inf _{M \in \mathcal{M}} \sum_{t=r}^{s} c_{t}\left(\hat{x}_{t}(M), \hat{u}_{t}(M)\right)\right.}_{\text {(regret := } \overbrace{\operatorname{Regret}}^{I} \text { ) }} \\
& +\underbrace{\inf _{M \in \mathcal{M}} \sum_{t=r}^{s} c_{t}\left(\hat{x}_{t}(M), \hat{u}_{t}(M)\right)-\inf _{M \in \mathcal{M}} \sum_{t=r}^{s} c_{t}\left(x_{t}(M), u_{t}(M)\right)}_{\text {(comparator error) }}
\end{aligned}
$$

First let us bound the realized iterate error which bounds the difference between what actually happened and what would have happened in the fictive $\left(\hat{G}_{t}, \hat{x}_{t}^{\text {nat }}\right)$ system (without exploration).
Lemma D.6. We can bound the difference between the true $x_{t}^{\mathrm{nat}}$ and the extracted $\hat{x}_{t}^{\mathrm{nat}}$ as:

$$
\left\|x_{t}^{\mathrm{nat}}-\hat{x}_{t}^{\mathrm{nat}}\right\| \leq R_{\mathrm{nat}} R_{\mathcal{M}}\left(\left\|\bar{G}_{t-1}-\hat{G}_{t-1}\right\|_{\ell_{1}, o p}+\psi(h)\right)
$$

Proof. Since $x^{\text {nat }} \in \mathbb{B}_{R_{\text {nat }}}$ and because (as argued earlier) $\left\|u_{t}\right\| \leq R_{\text {nat }} R_{\mathcal{M}}$, we have:

$$
\begin{aligned}
\left\|x_{t}^{\mathrm{nat}}-\hat{x}_{t}^{\mathrm{nat}}\right\| & =\left\|x_{t}^{\mathrm{nat}}-\operatorname{Proj}_{\mathbb{B}_{R_{\mathrm{nat}}}}\left[x_{t}-\sum_{i=0}^{h-1} \hat{G}_{t-1}^{[i]} u_{t-1-i}\right]\right\| \\
& \leq\left\|x_{t}^{\mathrm{nat}}-x_{t}+\sum_{i=0}^{h-1} \hat{G}_{t-1}^{[i]} u_{t-1-i}\right\| \\
& =\left\|\sum_{i=0}^{h-1}\left(G_{t-1}^{[i]}-\hat{G}_{t-1}^{[i]}\right) u_{t-1-i}+\sum_{i=h}^{t-2} G_{t-1}^{[i]} u_{t-1-i}\right\| \\
& \leq R_{\mathrm{nat}} R_{\mathcal{M}}\left(\left\|\bar{G}_{t-1}-\hat{G}_{t-1}\right\|_{\ell_{1}, o p}+\psi(h)\right)
\end{aligned}
$$

Lemma D. 7 (Realized Iterate Error). For $I=[r, s]$ with $r-1$, $s$ divisible by $h$, we can bound the realized iterate error as:
$\sum_{t=r}^{s} c_{t}\left(x_{t}, u_{t}\right)-c_{t}\left(\hat{x}_{t}\left(M_{t-h: t-1}\right), \hat{u}_{t}\left(M_{t}\right)\right) \leq 2 h \sum_{\tau_{1}=(r-1) / h}^{s / h} b_{\tau_{1}-1}+4 L \sqrt{d_{u}} \frac{R_{\mathrm{sys}}^{2}}{R_{G}}\left(\sum_{t=r}^{s}\left\|\bar{G}_{t-1}-\hat{G}_{t-1}\right\|_{\ell_{1}, o p}+\psi(h)|I|\right)$
where we denote $R_{\mathrm{sys}} \doteq R_{G} R_{\mathrm{nat}} R_{\mathcal{M}}$.

Proof. Consider the cost difference on a $h$-block indexed by $t=\overline{\tau_{1} \cdot h+1, \tau_{1} \cdot h+h}$. Consider the following three cases:

1. $b_{\tau_{1}}=1$ : in this case we are exploring and cannot give a better guarantee than

$$
c_{t}\left(x_{t}, u_{t}\right)-c_{t}\left(\hat{x}_{t}\left(M_{t-h: t-1)}, \hat{u}_{t}\left(M_{t}\right)\right) \leq 1\right.
$$

2. $b_{\tau_{1}}=0, b_{\tau_{1}-1}=1$ : while we are not exploring during the current round, and hence $u_{t}=\hat{u}_{t}\left(M_{t}\right)$ for $t=\overline{\tau_{1} \cdot h+1, \tau_{1} \cdot h+h}$, we have explored in the previous round and
therefore $u_{t}$ and $\hat{u}_{t}\left(M_{t}\right)$ may be arbitrarily far for $t \leq \tau_{1} \cdot h$. This can induce $x_{t}$ and $\hat{x}_{t}\left(M_{t-h: t-1}\right)$ to be quite far, especially the closer we get to $\tau_{1} \cdot h+1$. As such, in this event, we will also simply bound

$$
c_{t}\left(x_{t}, u_{t}\right)-c_{t}\left(\hat{x}_{t}\left(M_{t-h: t-1)}, \hat{u}_{t}\left(M_{t}\right)\right) \leq 1\right.
$$

3. $b_{\tau_{1}}=0, b_{\tau_{1}-1}=0$ : finally, in this case we have that $u_{t}=\hat{u}_{t}\left(M_{t}\right)$ for $t=$ $\overline{\tau_{1} \cdot h-h+1, \tau_{1} \cdot h+h}$. We can expand:

$$
x_{t}=x_{t}^{\mathrm{nat}}+\sum_{i=0}^{h-1} G_{t-1}^{[i]} u_{t-1-i}+\sum_{i=h}^{t-2} G_{t-1}^{[i]} u_{t-1-i}
$$

and

$$
\hat{x}_{t}\left(M_{t-h: t-1}\right)=\hat{x}_{t}^{\mathrm{nat}}+\sum_{i=0}^{h-1} \hat{G}_{t-1}^{[i]} \hat{u}_{t-1-i}\left(M_{t-1-i}\right)
$$

By the observation above, for $t=\overline{\tau_{1} \cdot h+1, \tau_{1} \cdot h+h}$, we have

$$
\begin{aligned}
\left\|x_{t}-\hat{x}_{t}\left(M_{t-h: t-1}\right)\right\| & \leq\left\|x_{t}^{\mathrm{nat}}-\hat{x}_{t}^{\mathrm{nat}}\right\|+R_{\mathrm{nat}} R_{\mathcal{M}}\left(\left\|\bar{G}_{t-1}-\hat{G}_{t-1}\right\|_{\ell_{1}, o p}+\psi(h)\right) \\
& \leq 2 R_{\mathrm{nat}} R_{\mathcal{M}}\left(\left\|\bar{G}_{t-1}-\hat{G}_{t-1}\right\|_{\ell_{1}, o p}+\psi(h)\right) \quad \text { Lemma D.6 }
\end{aligned}
$$

Hence, by the sub-quadratic Lipschitzness of the cost and Eq. (D.1) we have

$$
c_{t}\left(x_{t}, u_{t}\right)-c_{t}\left(\hat{x}_{t}\left(M_{t-h: t-1)}, \hat{u}_{t}\left(M_{t}\right)\right) \leq 4 L \sqrt{d_{u}} R_{G} R_{\mathrm{nat}}^{2} R_{\mathcal{M}}^{2}\left(\left\|\bar{G}_{t-1}-\hat{G}_{t-1}\right\|_{\ell_{1}, o p}+\psi(h)\right)\right.
$$

So for any $\tau_{1}=\overline{0, T / h}$, we have

$$
\begin{aligned}
\sum_{t=\tau_{1} \cdot h+1}^{\tau_{1} \cdot h+h} c_{t}\left(x_{t}, u_{t}\right)-c_{t}\left(\hat{x}_{t}\left(M_{t-h: t-1}\right), \hat{u}_{t}\left(M_{t}\right)\right) \leq & \mathbf{1}_{b_{\tau_{1}}=1} \cdot h+\mathbf{1}_{b_{\tau_{1}}=0, b_{\tau_{1}-1}=1} \cdot h \\
& +\mathbf{1}_{b_{\tau_{1}}=0, b_{\tau_{1}-1}=0} 4 L \sqrt{d_{u}} R_{G} R_{\mathrm{nat}}^{2} R_{\mathcal{M}}^{2} \sum_{t=\tau_{1} \cdot h+1}^{\tau_{1} \cdot h+h}\left\|\bar{G}_{t-1}-\hat{G}_{t-1}\right\|_{\ell_{1}, o p} \\
& +\mathbf{1}_{b_{\tau_{1}}=0, b_{\tau_{1}-1}=0} 4 L \sqrt{d_{u}} R_{G} R_{\mathrm{nat}}^{2} R_{\mathcal{M}}^{2} h \cdot \psi(h) \\
& \leq\left(b_{\tau_{1}}+b_{\tau_{1}-1}\right) \cdot h \\
& +4 L \sqrt{d_{u}} R_{G} R_{\mathrm{nat}}^{2} R_{\mathcal{M}}^{2}\left(\sum_{t=\tau_{1} \cdot h}^{\left(\tau_{1}+1\right) \cdot h-1}\left\|\bar{G}_{t}-\hat{G}_{t}\right\|_{\ell_{1}, o p}+h \psi(h)\right)
\end{aligned}
$$

summing over $\tau_{1}$ yields the desired result.

Lemma D. 8 (Comparator Error). We can bound the comparator error as:

$$
\text { (comparator error) } \leq 8 L \sqrt{d_{u}} R_{\mathrm{nat}}^{2} R_{G}^{2} R_{\mathcal{M}}^{3} h\left(m \sum_{t=r-h-m}^{s}\left\|\bar{G}_{t}-\hat{G}_{t}\right\|_{\ell_{1}, o p}+\psi(h)\right)
$$

Proof. Let $M^{\star}=\arg \min _{M \in \mathcal{M}} \sum_{t=r}^{s} \sum_{t=r}^{s} c_{t}\left(x_{t}(M), u_{t}(M)\right)$. We have that:

$$
\begin{aligned}
(\text { comparator error) } & \leq \sum_{t=r}^{s} c_{t}\left(\hat{x}_{t}\left(M^{\star}\right), \hat{u}_{t}\left(M^{\star}\right)\right)-c_{t}\left(x_{t}\left(M^{\star}\right), u_{t}\left(M^{\star}\right)\right) \\
& \left.\leq 2 L \sqrt{d_{u}} R_{\mathrm{nat}} R_{G} R_{\mathcal{M}} \sum_{t=r}^{s}\left(\left\|\hat{x}_{t}\left(M^{\star}\right)-x_{t}\left(M^{\star}\right)\right\|+\| \hat{u}_{t}\left(M^{\star}\right)\right)-u_{t}\left(M^{\star}\right) \|\right)
\end{aligned}
$$

We have that

$$
\begin{aligned}
\left\|\hat{u}_{t}\left(M^{\star}\right)-u_{t}\left(M^{\star}\right)\right\| & =\left\|\sum_{i=0}^{m-1} M^{\star,[i]}\left(\hat{x}_{t-i}^{\mathrm{nat}}-x_{t-i}^{\mathrm{nat}}\right)\right\| \\
& \leq R_{\mathrm{nat}} R_{\mathcal{M}}^{2}\left(\sum_{\tau=t-m}^{t-1}\left\|\bar{G}_{\tau}-\hat{G}_{\tau}\right\|_{\ell_{1}, o p}+\psi(h)\right)
\end{aligned}
$$

With a bit more computation, we can also bound the difference in the states. First, let us expand the expression of the states:

$$
\begin{gathered}
x_{t}\left(M^{\star}\right)=x_{t}^{\mathrm{nat}}+\sum_{i=0}^{h-1} G_{t-1}^{[i]} u_{t-1-i}\left(M^{\star}\right)+\sum_{i=h}^{t-2} G_{t-1}^{[i]} u_{t-1-i}\left(M^{\star}\right) \\
\hat{x}_{t}\left(M^{\star}\right)=\hat{x}_{t}^{\mathrm{nat}}+\sum_{i=0}^{h-1} \hat{G}_{t-1}^{[i]} \hat{u}_{t-1-i}\left(M^{\star}\right)
\end{gathered}
$$

The only new thing we need to bound is:

$$
\begin{aligned}
\sum_{i=0}^{h-1} G_{t-1}^{[i]} u_{t-1-i}\left(M^{\star}\right)-\sum_{i=0}^{h-1} \hat{G}_{t-1}^{[i]} \hat{u}_{t-1-i}\left(M^{\star}\right)= & \sum_{i=0}^{h-1} G_{t-1}^{[i]} u_{t-1-i}\left(M^{\star}\right)-\sum_{i=0}^{h-1} \hat{G}_{t-1}^{[i]} u_{t-1-i}\left(M^{\star}\right) \\
& +\sum_{i=0}^{h-1} \hat{G}_{t-1}^{[i]} u_{t-1-i}\left(M^{\star}\right)-\sum_{i=0}^{h-1} \hat{G}_{t-1}^{[i]} \hat{u}_{t-1-i}\left(M^{\star}\right) \\
\leq & R_{\mathrm{nat}} R_{\mathcal{M}}\left\|\bar{G}_{t-1}-\hat{G}_{t-1}\right\|_{\ell_{1}, o p}+R_{G} \sum_{i=1}^{h}\left\|u_{t-i}\left(M^{\star}\right)-\hat{u}_{t-i}\left(M^{\star}\right)\right\| \\
\leq & 2 R_{G} R_{\mathrm{nat}} R_{\mathcal{M}}^{2} h\left(m \sum_{\tau=t-h-m}^{t-1}\left\|\bar{G}_{\tau}-\hat{G}_{\tau}\right\|_{\ell_{1}, o p}+\psi(h)\right)
\end{aligned}
$$

Plugging this into our expressions for $x_{t}\left(M^{\star}\right)$, and using previous bounds we have

$$
\left\|\hat{x}_{t}\left(M^{\star}\right)-x_{t}\left(M^{\star}\right)\right\| \leq 3 R_{G} R_{\mathrm{nat}} R_{\mathcal{M}}^{2} h\left(m \sum_{\tau=t-h-m}^{t-1}\left\|\bar{G}_{\tau}-\hat{G}_{\tau}\right\|_{\ell_{1}, o p}+\psi(h)\right)
$$

Hence we can finalize that:

$$
\text { (comparator error) } \leq 8 L \sqrt{d_{u}} R_{\mathrm{nat}}^{2} R_{G}^{2} R_{\mathcal{M}}^{3} h\left(m \sum_{t=r-h-m}^{s}\left\|\bar{G}_{t}-\hat{G}_{t}\right\|_{\ell_{1}, o p}+\psi(h)\right)
$$

## D. 3 Proof of Theorem D. 1

Proof of Theorem D.1. We have that

$$
\begin{aligned}
\text { Regret }_{I} \leq & \underbrace{18 L \sqrt{d_{\min }} R_{G}^{2} R_{\mathcal{M}}^{2} R_{\mathrm{nat}}^{2} m(h+1)^{5 / 4} \sqrt{T}}_{\text {(Known System Regret) }}+\underbrace{2 h \sum_{\tau=(r-1) / h}^{\sum_{\tau-1}} b_{\tau}}_{\text {(Exploration Penalty) }} \\
& +\underbrace{12 L R_{\mathrm{nat}}^{2} R_{G}^{2} R_{\mathcal{M}}^{3} h m\left(\sum_{t=r}^{s}\left\|\bar{G}_{t}-\hat{G}_{t}\right\|_{\ell_{1}, o p}+2 R_{G}(h+m)\right)}_{\text {(System Misspecification Induced Error) }} \\
& +\underbrace{18 L \sqrt{d_{u}} R_{G}^{2} R_{\mathcal{M}}^{3} R_{\mathrm{nat}}^{2} \psi(h) h(|I|+h+m)}_{\text {(Truncation Error) }}
\end{aligned}
$$

Taking expectation and plugging in Proposition D. 5 we get:

$$
\begin{align*}
\mathbb{E}\left[\operatorname{Regret}_{I}\right] \leq & \widetilde{\mathcal{O}}\left(L \sqrt{d_{u}} R_{\mathrm{sys}}^{2} m \sqrt{T}+p|I|\right. \\
& +L R_{\mathrm{sys}}^{2} R_{\mathcal{M}} m\left(|I|^{1 / 2} \sqrt{\mathbb{E}\left[\operatorname{Var}_{I}^{\mathrm{tot}}(\mathbf{G})\right]}+d_{u} R_{\mathrm{sys}}|I|^{1 / 2} p^{-1 / 2}+R_{G} m\right) \\
& \left.+L \sqrt{d_{u}} R_{\mathrm{sys}}^{2} R_{\mathcal{M}} R_{G} m\right)  \tag{D.2}\\
= & \widetilde{\mathcal{O}}^{\star}\left(L m\left(|I| \sqrt{\mathbb{E}\left[\operatorname{Var}_{I}(\mathbf{G})\right]}+d_{u} T^{2 / 3}\right)\right)
\end{align*}
$$

where $R_{\text {sys }}=R_{G} R_{\mathcal{M}} R_{\text {nat }}$, using $\operatorname{Var}_{I}(\mathbf{G})=|I| \operatorname{Var}_{I}^{\mathrm{tot}}(\mathbf{G})$, and that for the chosen $h$ we have $\psi(h) \leq R_{G} T^{-1}$.

## E Sublinear Regret for State Feedback

We demonstrate that it is (information-theoretically) possible to achieve sublinear (though large) regret against a benchmark of stabilizing static feedback control policies.
We suppose there is a subset $\mathcal{K} \subset \mathbb{R}^{d_{u} \times d_{x}}$ of feedback policies $K$, and our goal is to obtain regret compared to the best $K \in \mathcal{K}$ :

$$
\operatorname{Reg}_{T}(\mathcal{K}):=\sum_{t=1}^{T} c_{t}\left(x_{t}, u_{t}\right)-\inf _{K \in \mathcal{K}} \sum_{t=1}^{T} c_{t}\left(x_{t}^{K}, u_{t}^{K}\right)
$$

where $\left(x_{t}^{K}, u_{t}^{K}\right)$ are the iterates arising under the control law $u_{t}=K x_{t}$.
For this setting, we propose an algorithm the classic Exp3exponential weights algorithm (see, e.g. Chapter 3 of [8]) on an $\varepsilon$-cover $\mathcal{K}_{\varepsilon}$ of $\mathcal{K}$ in the operator norm. We maintain a constant controller $K$ on intervals of length $H$, and feed the losses on those intervals to the Exp3 algorithm. Pseudocode is given in Algorithm 8.
We state our regret bound under the (quite restrictive) assumption that all policies $K \in \mathcal{K}$ are sequentially stabilizing. Formally, given a sequence of controllers $K \in \mathcal{K}$, we define

$$
\Phi_{s: t}(K):=\prod_{i=s}^{t}\left(A_{i}+K_{i} B_{i}\right)=\left(A_{t}+K B_{t}\right) \cdot\left(A_{t-1}+K B_{t-1}\right) \cdots \cdots\left(A_{s}+K B_{s}\right)
$$

We assume that $\Phi_{s: t}(K)$ exhibits geometric decay uniformly over all times for any fixed $K$ :
Assumption 4. There exists $c_{\star} \geq 1$ and $\rho_{\star} \in(1 / 2,1)$ such that for any indices $s \leq t$ and any fixed $K \in \mathcal{K},\left\|\Phi_{s: t}(K)\right\|_{\mathrm{op}} \leq c_{\star} \rho_{\star}^{t-s}$. We define the constant

$$
R_{\mathcal{K}}:=1+\max \{\|K\|: K \in \mathcal{K}\}
$$

Theorem E.1. Suppose Assumptions 3 and 4 holds, and for some $R_{w} \geq 1$ and $R_{B} \geq 0$, $\max _{t}\left\|w_{t}\right\| \leq R_{w}$, and $\max _{t}\left\|B_{t}\right\| \leq R_{B}$. In addition, suppose $T$ is large enough that

```
Algorithm 8 Exponentially Weighted Control
    Input: window length \(H\), step size \(\eta>0\), finite \(\varepsilon\)-cover \(\mathcal{K}_{\varepsilon} \subset \mathcal{K}\), initial estimate \(K_{1} \in \mathcal{K}_{\varepsilon}\)
    Initialize \(\mathcal{L}_{1}(K)=0\) and \(p_{1}(K)=1 /\left|\mathcal{K}_{\varepsilon}\right|, \forall K \in \mathcal{K}_{\varepsilon}\),
    for \(t=1, \ldots, T\) do
        Play select \(u_{t}=K_{t} x_{t}\).
        Recieve \(\hat{c}_{t}=c_{t}\left(x_{t}, u_{t}\right)\)
        if \(t \bmod H=0\) then
            Set \(n=t / H, \ell_{n}=\sum_{i=t-H+1}^{t} \hat{c}_{i}\)
            Set \(\mathcal{L}_{n+1}\left(K_{t}\right)=\frac{1}{p_{n}\left(K_{t}\right)} \ell_{n}+\mathcal{L}_{n}\left(K_{t}\right)\)
            Set \(\mathcal{L}_{n+1}(K)=\mathcal{L}_{n}(K)\) for all \(K \in \mathcal{K}_{\varepsilon} /\left\{K_{t}\right\}\)
            Set \(p_{n+1}(K)=\frac{\exp \left(\eta \mathcal{L}_{n+1}(K)\right)}{\sum_{K^{\prime} \in \mathcal{K}} \exp \left(\eta \mathcal{L}_{n+1}\left(K^{\prime}\right)\right)}\)
            Sample \(K_{t+1} \sim p_{n+1}(\cdot)\).
        else
            Set \(K_{t+1}=K_{t}\)
```

$\mathcal{C}_{\star} \rho_{\star}^{T / 4} \leq 1 / 2$. Then, Algorithm 8 with horizon $H=\left\lceil T^{1 / 4}\right\rceil$, appropriate an step size $\eta$ and minimal $\varepsilon$-covering $\mathcal{K}_{\varepsilon}$ of $\mathcal{K}$ enjoys the following regret bound:

$$
\mathbb{E}\left[\operatorname{Reg}_{T}(\mathcal{K})\right] \leq L \mathcal{C}_{1}\left(5 R_{\mathcal{K}}\right)^{d_{x} d_{u} / 2} \cdot T^{1-\frac{1}{2\left(d_{x} d_{u}+3\right)}}
$$

where $\mathcal{C}_{1}=\mathcal{O}\left(\frac{\mathcal{C}_{\star}^{5}}{\left(1-\rho_{\star}\right)^{3}} R_{\mathcal{K}}^{3} R_{w}^{2}\left(1+R_{B}\right) \sqrt{d_{u} d_{x}}\right)$.
The theorem is established by a reduction to online multi-arm bandits in Appendix E. 1 below.
Remark E. 1 (Extensions of Theorem E.1). The following analysis extends to policies of the form $u_{t}=\left(K_{t}^{\text {stb }}+K\right) x_{t}+v$, where $\left(K_{t}^{\text {stb }}\right)$ is a fixed sequence of control policies determined a priori, $K \in \mathcal{K} \subset \mathbb{R}^{d_{u} \times d_{x}}$ is a feedback parameter, and $v \in \mathcal{V} \subset \mathbb{R}^{d_{u}}$ is a bounded affine term. Letting $\left(x_{t}^{K, c}, u_{t}^{K, c}\right)$ denote the iterates produced by such a policy, our notion of regret is

$$
\operatorname{Reg}_{T}(\mathcal{K} \times \mathcal{V}):=\sum_{t=1}^{T} c_{t}\left(x_{t}, u_{t}\right)-\inf _{(K, c) \in \mathcal{K} \times \mathcal{V}} \sum_{t=1}^{T} c_{t}\left(x_{t}^{K, c}, u_{t}^{K, c}\right),
$$

The only assumptions we require in general is that $\mathcal{V}$ is bounded, and that $\mathcal{K}$, combined with $\left(K_{i}^{\text {stb }}\right)$, are sequentially stabilizing in the sense that, for any $s \leq t$, the fixed ( $K_{i}^{\text {stb }}$ ) sequence, and any $K_{s: t}^{\text {stb }} \in \mathcal{K}^{t-s+1}$, it holds that the products

$$
\Phi_{s: t}^{\mathrm{stb}}\left(K_{s: t}\right):=\prod_{i=s}^{t}\left(A_{i}+\left(K_{i}^{\mathrm{stb}}+K_{i}\right) B_{i}\right)
$$

exhibit geometric decay.

## E. 1 Proof of Theorem E. 1

In what follows, assume that $H=T^{1 / 4}$ evenly divides $T$. For every index $n \in \mathbb{N}$, define $t_{n}=$ $1+(n-1) H$. To avoid confusion, we $x_{t}^{\text {alg }}, u_{t}^{\text {alg }}$ denote the iterates produced by the algorithm. We define the sequence which begins at state $x_{t_{n}}^{\text {alg }}$ at time $t_{n}$, and rolls forward under controller $K$ for future times:

$$
\begin{aligned}
& \bar{x}_{n ; t_{n}}(K)=x_{t_{n}}^{\mathrm{alg}}, \quad \bar{x}_{n ; t+1}(K)=\left(A_{t}+B_{t} K\right) \bar{x}_{n ; t}(K)+w_{t}, \quad t \geq t_{n} \\
& \bar{u}_{n ; t_{n}}(K)=K \bar{x}_{n ; t_{n}}(K) .
\end{aligned}
$$

Observe that, since we select a new controller $K_{t_{n}}$ just before each time $t_{n}$, we have

$$
\left(\bar{x}_{n ; t_{n}}\left(K_{t_{n}}\right), \bar{u}_{n ; t_{n}}\left(K_{t_{n}}\right)\right)=\left(x_{t}^{\text {alg }}, u_{t}^{\text {alg }}\right), \quad \forall t \in\left[t_{n}, t_{n+1}-1\right] .
$$

Therefore, defining the losses,

$$
\ell_{n}(K)=\sum_{t=t_{n}}^{t_{n+1}-1} c_{t}\left(\bar{x}_{n ; t_{n}}(K), \bar{u}_{n ; t_{n}}(K)\right)
$$

we have

$$
\ell_{n}\left(K_{t_{n}}\right)=\sum_{t=t_{n}}^{t_{n+1}-1} c_{t}\left(x_{t}^{\mathrm{alg}}, u_{t}^{\mathrm{alg}}\right)
$$

Therfore, we may decompose the regret as

$$
\begin{aligned}
\mathbb{E}\left[\operatorname{Reg}_{T}(\mathcal{K})\right] & =\underbrace{\mathbb{E}\left[\sum_{n=1}^{T / H} \ell_{n}\left(K_{t_{n}}\right)\right]-\inf _{K \in \mathcal{K}_{\varepsilon}} \mathbb{E}\left[\sum_{n=1}^{T / H} \ell_{n}(K)\right]}_{R_{1}} \\
& +\underbrace{\inf _{K \in \mathcal{K}_{\varepsilon}} \mathbb{E}\left[\sum_{n=1}^{T / H} \ell_{n}(K)\right]-\inf _{K \in \mathcal{K}_{\varepsilon}} \sum_{t=1}^{T} c_{t}\left(x_{t}^{K}, u_{t}^{K}\right)}_{R_{2}} \\
& +\underbrace{\left.\inf _{K \in \mathcal{K}_{\varepsilon}} \sum_{t=1}^{T} c_{t}\left(x_{t}^{K}, u_{t}^{K}\right)\right]-\inf _{K \in \mathcal{K}} \sum_{t=1}^{T} c_{t}\left(x_{t}^{K}, u_{t}^{K}\right)}_{R_{3}} .
\end{aligned}
$$

Here, $R_{1}$ is the simple regret on the $\ell_{n}$ sequence, $R_{2}$ is the extend to which the $\ell_{n}$ sequence approximates regret against controller $K \in \mathcal{K}_{\varepsilon}$ in the covering, and finally $R_{3}$ bounds the regret of the covering against the full set $\mathcal{K}$. Here, expectations are over the randomness in the algorithm, and due to the obliviousness of the adversary, we may assume that $\left(c_{t}, x_{t}^{K}, u_{t}^{K}\right)$ are deterministic and chosen in advance. We bound each of the three terms in sequence. Before proceeding, we use the following estimates:

Lemma E. 1 (Key Term Bounds). Suppose that $H=T^{1 / 4}$ is sufficiently large that $\mathcal{C}_{\star} \rho_{\star}^{H} \leq 1 / 2$. Moreover, let $R_{\star}=\frac{\mathcal{C}_{\star}}{1-\rho_{\star}}$. Then,
(a) For all $K \in \mathcal{K}$ and $t \in[T],\left\|x_{t}^{K}\right\| \leq R_{\star} R_{w}$
(b) For any $t \geq t_{n}$ and $K \in \mathcal{K},\left\|\bar{x}_{n ; t}(K)\right\| \leq 2 \mathcal{C}_{\star} R_{\star} R_{w}$
(c) $\left\|\left(x_{t}^{K}, u_{t}^{K}\right)\right\| \leq R_{\mathcal{K}} R_{\star} R_{w}$ and $\left\|\left(\bar{x}_{n ; t}(K), \bar{u}_{n ; t}(K)\right)\right\| \leq 2 R_{\mathcal{K}} \mathcal{C}_{\star} R_{\star} R_{w}$.

Proof. Part a: Unfolding the dynamics, and bounding $\left\|w_{t}\right\| \leq R_{w}$ and $\Phi$ via Assumption 4,

$$
\left\|x_{t}^{K}\right\|=\left\|\sum_{s=1}^{t-1} \Phi_{s+1: t}(K) w_{t}\right\| \leq R_{w} \mathcal{C}_{\star} \sum_{s \geq 0} \rho_{\star}^{s}=\frac{\mathcal{C}_{\star} R_{w}}{\rho_{\star}}:=R_{\star} R_{w}
$$

Next, we bound $\left\|x_{t_{n}}^{\text {alg }}\right\|$ for some $n$,

$$
\begin{aligned}
\left\|x_{t_{n}}^{\mathrm{alg}}\right\| & =\left\|\sum_{i=1}^{H} \Phi_{t_{n-1}+i: t_{n}-1}\left(K_{t_{n-1}}\right) w_{t_{n}-i}+\Phi_{t_{n-1}: t_{n}}\left(K_{t_{n-1}}\right) x_{t_{n-1}}^{\mathrm{alg}}\right\| \\
& \leq R_{\star} R_{w}+\mathcal{C}_{\star} \rho_{\star}^{H}\left\|x_{t_{n}-1}^{\mathrm{alg}}\right\| .
\end{aligned}
$$

If $H$ is sufficiently large that $\mathcal{C}_{\star} \rho_{\star}^{H} \leq 1 / 2$, then the above is just

$$
\left\|x_{t_{n}}^{\mathrm{alg}}\right\| \leq R_{\star} R_{w}+\frac{1}{2}\left\|x_{t_{n}-1}^{\mathrm{alg}}\right\|
$$

yielding the bound $\left\|x_{t_{n}}^{\text {alg }}\right\| \leq 2 R_{\star} R_{w}$ for all $n$. Part $b$ : Next, let us bound $\left\|\bar{x}_{n ; t}(K)\right\|$ for some $t \geq t_{n}$. We have

$$
\begin{aligned}
\left\|\bar{x}_{n ; t}(K)\right\| & =\left\|\sum_{i=t_{n}+1}^{t} \Phi_{i: t}(K) w_{t} t+\Phi_{t_{n}: t}(K) x_{t_{n}}^{\mathrm{alg}}\right\| \\
& \leq R_{w}\left(\sum_{0}^{t-t_{n}-1} \mathcal{C}_{\star} \rho_{\star}^{i}\right)+\mathcal{C}_{\star} \rho_{\star}^{t_{n}}\left\|x_{t_{n}}^{\mathrm{alg}}\right\| \\
& \leq R_{w} \sum_{0}^{t-t_{n}-1} \mathcal{C}_{\star} \rho_{\star}^{i}+2 \rho_{\star}^{t_{n}} \mathcal{C}_{\star} R_{\star} R_{w} .
\end{aligned}
$$

$\operatorname{Using} R_{\star}=\frac{\mathcal{C}_{\star}}{\rho_{\star}}, \mathcal{C}_{\star} \geq 1$, and $\sum_{0}^{t-t_{n}-1}+\frac{\rho_{\star}^{t_{n}}}{1-\rho_{\star}}=\frac{1}{1-\rho_{\star}}$, the above simplifies to $2 \frac{\mathcal{C}_{\star}^{2} R_{w}}{1-\rho_{\star}}=2 \mathcal{C}_{\star} R_{w} R_{\star}$. Part $c$ : This follows from the fact that, for any $x \in \mathbb{R}^{d_{x}}$ and $K \in \mathcal{K},\|(x, K x)\| \leq(1+\|K\|)\|x\| \leq$ $R_{\mathcal{K}}\|x\|$.

Bounding $R_{1} \quad$ The term $R_{1}$ corresponds to the simple regret on the sequence of losses $\ell_{n}(K)$ over the discrete enumeration of controllers $K \in \mathcal{K}_{\varepsilon}$. Examining Algorithm 8, we simply run the Exp3 algorithm on these losses. By appealing to a standard regret bound for this algorithm with appropriate step size $\eta$, we ensure that

$$
R_{1} \leq 2 B \sqrt{\frac{T}{H}\left|\mathcal{K}_{\varepsilon}\right| \log \left|\mathcal{K}_{\varepsilon}\right|}
$$

provided that, for all $n$ and $K \in \mathcal{K}_{\varepsilon}, \ell_{n}(K) \in[0, B]$. To find the appropriate bound $B$, we note that from the growth condition on the costs, Assumption 3, we have

$$
\begin{aligned}
0 & \leq \ell_{n}(K)=\sum_{t=t_{n}}^{t_{n+1}-1} c_{t}\left(\bar{x}_{n ; t_{n}}(K), \bar{u}_{n ; t_{n}}(K)\right) \leq H \max _{t \geq t_{n}} c_{t}\left(\bar{x}_{n ; t_{n}}(K), \bar{u}_{n ; t_{n}}(K)\right) \\
& \leq L H \max \left\{1,\left\|\left(\bar{x}_{n ; t_{n}}(K), \bar{u}_{n ; t_{n}}(K)\right)\right\|^{2}\right\} \leq 8 L H\left(R_{\mathcal{K}} \mathcal{C}_{\star} R_{\star} R_{w}\right)^{2}
\end{aligned}
$$

where the last inequality uses Lemma E.1. Hence,

$$
R_{1} \leq 16 L\left(R_{\mathcal{K}} \mathcal{C}_{\star} R_{\star} R_{w}\right)^{2} \sqrt{T H\left|\mathcal{K}_{\varepsilon}\right| \log \left|\mathcal{K}_{\varepsilon}\right|}
$$

Bounding $R_{2}$ : To bound $R_{2}$, it suffices to find a probability-one upper bound on

$$
\begin{aligned}
\sup _{K \in \mathcal{K}_{\varepsilon}}\left|\sum_{n=1}^{T / H} \ell_{n}(K)-\sum_{t=1}^{T} c_{t}\left(x_{t}^{K}, u_{t}^{K}\right)\right| & =\sup _{K \in \mathcal{K}_{\varepsilon}}\left|\sum_{n=1}^{T / H} \sum_{t=t_{n}}^{t_{n+1}-1} c_{t}\left(\bar{x}_{n ; t}(K), \bar{u}_{n ; t}(K)\right)-c_{t}\left(x_{t}^{K}, u_{t}^{K}\right)\right| \\
& \leq \frac{T}{H} \sup _{K \in \mathcal{K}_{\varepsilon}} \max _{n} \sum_{t=t_{n}}^{t_{n+1}-1}\left|c_{t}\left(\bar{x}_{n ; t}(K), \bar{u}_{n ; t}(K)\right)-c_{t}\left(x_{t}^{K}, u_{t}^{K}\right)\right|
\end{aligned}
$$

Using the Lipschitz conditions on $c_{t}$, the bounds from Lemma E.1, and the bound $1+\|K\| \leq R_{\mathcal{K}}$,

$$
\begin{aligned}
\sum_{t=t_{n}}^{t_{n+1}-1}\left|c_{t}\left(\bar{x}_{n ; t}(K), \bar{u}_{n ; t}(K)\right)-c_{t}\left(x_{t}^{K}, u_{t}^{K}\right)\right| & \leq L\left(R_{\mathcal{K}} \mathcal{C}_{\star} R_{\star} R_{w}\right) \sum_{t=t_{n}}^{t_{n+1}-1}\left\|\left(\bar{x}_{n ; t}(K), \bar{u}_{n ; t}(K)\right)-\left(x_{t}^{K}, u_{t}^{K}\right)\right\| \\
& =L\left(R_{\mathcal{K}} \mathcal{C}_{\star} R_{\star} R_{w}\right) \sum_{t=t_{n}}^{t_{n+1}-1}\left\|\left(\bar{x}_{n ; t}(K), K \bar{x}_{n ; t}(K)\right)-\left(x_{t}^{K}, K x_{t}^{K}\right)\right\| \\
& =L\left(R_{\mathcal{K}} \mathcal{C}_{\star} R_{\star} R_{w}\right) R_{\mathcal{K}} \sum_{t=t_{n}}^{t_{n+1}-1}\left\|\bar{x}_{n ; t}(K)-x_{t}^{K}\right\|
\end{aligned}
$$

Finally, we can compute that the difference $\bar{x}_{n ; t}(K)-x_{t}^{K}=\Phi_{t_{n}}^{t}\left(x_{t_{n}}^{\mathrm{alg}}-x_{t_{n}}^{K}\right)$ depends only on the response to the state difference at time $t_{n}$. Hence, using Assumption 4 and Lemma E.1,the above is
at most

$$
\begin{aligned}
L\left(R_{\mathcal{K}} \mathcal{C}_{\star} R_{\star} R_{w}\right) R_{\mathcal{K}} \cdot \sum_{i \geq 0} \mathcal{C}_{\star} \rho_{\star}^{i}\left\|\left(x_{t_{n}}^{\mathrm{alg}}-x_{t_{n}}^{K}\right)\right\| & \leq L\left(R_{\mathcal{K}} \mathcal{C}_{\star} R_{\star} R_{w}\right) R_{\mathcal{K}} \cdot \underbrace{\frac{\mathcal{C}_{\star}}{1-\rho_{\star}}}=R_{\star} \cdot \mathcal{C}_{\star} R_{\star} R_{w} \\
& \leq L\left(R_{\mathcal{K}} \mathcal{C}_{\star} R_{\star} R_{w}\right)^{2} R_{\star}
\end{aligned}
$$

Concluding, we find

$$
R_{2} \leq L\left(R_{\mathcal{K}} \mathcal{C}_{\star} R_{\star} R_{w}\right)^{2} \cdot \frac{T R_{\star}}{H}
$$

Bounding $R_{3}$. We now turn to bounding $R_{3}$, which captures the approximation error of approximating $\mathcal{K}$ with $\mathcal{K}_{\varepsilon}$. We require the following technical lemma:
Lemma E.2. Let $K, K^{\prime} \in \mathcal{K}$ satisfy $\left\|K-K^{\prime}\right\|_{\mathrm{op}} \leq \varepsilon$. Then, for all $t \geq 1$,

$$
\left\|x_{t}^{K}-x_{t}^{K^{\prime}}\right\| \leq 4 \varepsilon R_{w} R_{\star}^{2} R_{B}
$$

Hence,

$$
\left\|c_{t}\left(x_{t}^{K}, u_{t}^{K}\right)-c_{t}\left(x_{t}^{K^{\prime}}, u_{t}^{K^{\prime}}\right)\right\| \leq 4 \varepsilon L R_{w}^{2} R_{\mathcal{K}}^{2} R_{\star}^{3}\left(1+R_{B}\right)
$$

Using the fact that $\mathcal{K}_{\varepsilon}$ is an $\varepsilon$-covering of $\mathcal{K}$ in the operator norm means that for any $K \in \mathcal{K}$, we can find a $K^{\prime} \in \mathcal{K}_{\varepsilon}$ for which $\left\|K-K^{\prime}\right\|_{\text {op }} \leq \varepsilon$. Hence, from the above lemma

$$
\left|\sum_{t=1}^{T} c_{t}\left(x_{t}^{K}, u_{t}^{K}\right)-c_{t}\left(x_{t}^{K^{\prime}}, u_{t}^{K^{\prime}}\right)\right| \leq 4 T \varepsilon L R_{w}^{2} R_{\mathcal{K}}^{2} R_{\star}^{3}\left(1+R_{B}\right)
$$

Since $R_{3} \leq \sup _{K \in \mathcal{K}} \inf _{\mathcal{K}^{\prime} \in \mathcal{K}_{\varepsilon}}\left|\sum_{t=1}^{T} c_{t}\left(x_{t}^{K}, u_{t}^{K}\right)-c_{t}\left(x_{t}^{K^{\prime}}, u_{t}^{K^{\prime}}\right)\right|$, we conclude

$$
R_{3} \leq 4 T \varepsilon L R_{w}^{2} R_{\mathcal{K}}^{2} R_{\star}^{3}\left(1+R_{B}\right)
$$

Concluding the proof In sum, we found

$$
\begin{aligned}
\mathbb{E}\left[\operatorname{Reg}_{T}(\mathcal{K})\right] & =R_{1}+R_{2}+R_{3} \\
& \leq \mathcal{O}\left(L\left(R_{\mathcal{K}} \mathcal{C}_{\star} R_{\star} R_{w}\right)^{2}\right)\left(\sqrt{T H\left|\mathcal{K}_{\varepsilon}\right| \log \left|\mathcal{K}_{\varepsilon}\right|}+\frac{T R_{\star}}{H}+\left(1+R_{B}\right) T \varepsilon\right) \\
& \leq \mathcal{O}\left(L\left(R_{\mathcal{K}} \mathcal{C}_{\star} R_{\star} R_{w}\right)^{2}\left(1+R_{B}\right) R_{\star}\right)\left(\sqrt{T H\left|\mathcal{K}_{\varepsilon}\right| \log \left|\mathcal{K}_{\varepsilon}\right|}+\frac{T}{H}+T \varepsilon\right)
\end{aligned}
$$

where in the last line, we use $R_{\star} \geq 1$ and $1+R_{B} \geq 1$. Setting $H=T^{1 / 4}$,

$$
\mathbb{E}\left[\operatorname{Reg}_{T}(\mathcal{K})\right] \leq \mathcal{O}\left(L\left(R_{\mathcal{K}} \mathcal{C}_{\star} R_{\star} R_{w}\right)^{2}\left(1+R_{B}\right) R_{\star}\right) T^{3 / 4} \cdot\left(\sqrt{\left|\mathcal{K}_{\varepsilon}\right| \log \left|\mathcal{K}_{\varepsilon}\right|}+T^{1 / 4} \varepsilon\right)
$$

We bound the cardinality of $\mathcal{K}_{\varepsilon}$. It suffices to ensure $\mathcal{K}_{\varepsilon}$ is an $\varepsilon$-covering in the larger Frobenius norm, which is just the Euclidean norm on $R^{d_{x} d_{u}}$. Since $\mathcal{K}$ is a bounded subset of this space, with radius at most $R_{\mathcal{K}}$, we can find a covering such that $\left|\mathcal{K}_{\varepsilon}\right| \leq\left(\frac{5 R_{\mathcal{K}}}{\varepsilon}\right)^{d_{x} d_{u}}$ (see, e.g. Chapter 4.2 in [43]). This yields

$$
\left|\mathcal{K}_{\varepsilon}\right| \log \left|\mathcal{K}_{\varepsilon}\right| \leq d_{x} d_{u}\left(\frac{5 R_{\mathcal{K}}}{\varepsilon}\right)^{d_{x} d_{u}} \log \left(\frac{5 R_{\mathcal{K}}}{\varepsilon}\right) \leq d_{x} d_{u}\left(\frac{5 R_{\mathcal{K}}}{\varepsilon}\right)^{d_{x} d_{u}+1}=5 R_{\mathcal{K}} \cdot\left(5 R_{\mathcal{K}}\right)^{d_{x} d_{u}} \varepsilon^{-d_{x} d_{u}+1}
$$

where we use $\log x \leq x$. Hence, we can bound
$\mathbb{E}\left[\operatorname{Reg}_{T}(\mathcal{K})\right] \leq \mathcal{O}\left(L R_{\mathcal{K}} R_{\star}\left(R_{\mathcal{K}} \mathcal{C}_{\star} R_{\star} R_{w}\right)^{2}\left(1+R_{B}\right) \sqrt{d_{x} d_{u}}\right)\left(5 R_{\mathcal{K}}\right)^{d_{x} d_{u} / 2} \cdot T^{3 / 4}\left(\left(\frac{1}{\varepsilon}\right)^{\left(1+d_{x} d_{u}\right) / 2}+\varepsilon T^{1 / 4}\right)$
Setting $\varepsilon=T^{-\frac{1}{2\left(d_{x} d_{u}+3\right)}}$ gives

$$
\begin{aligned}
\mathbb{E}\left[\operatorname{Reg}_{T}(\mathcal{K})\right] & \leq \mathcal{O}\left(L R_{\star}\left(R_{\mathcal{K}} \mathcal{C}_{\star} R_{\star} R_{w}\right)^{2}\left(1+R_{B}\right) \sqrt{d_{x} d_{u}}\right)\left(5 R_{\mathcal{K}}\right)^{d_{x} d_{u}+1} \cdot T^{1-\frac{1}{2\left(d_{x} d_{u}+3\right)}} \\
& =\mathcal{O}\left(L \frac{\mathcal{C}_{\star}^{5}}{\left(1-\rho_{\star}\right)^{3}} R_{\mathcal{K}}^{3} R_{w}^{2}\left(1+R_{B}\right) \sqrt{d_{x} d_{u}}\right)\left(5 R_{\mathcal{K}}\right)^{d_{x} d_{u} / 2} \cdot T^{1-\frac{1}{2\left(d_{x} d_{u}+3\right)}}
\end{aligned}
$$

## E. 2 Ommited Proofs

Proof of Lemma E.1. Part a: All such iterates can be realized by dynamics of the form $x_{1}=0$, $x_{t+1}=\left(A_{t}+B_{t} K_{t}\right) x_{t}+w_{t}$ and $u_{t}=K_{t} x_{t}$ for any appropriate sequence $\left(K_{1}, K_{2}, \ldots\right)$ of elements of $\mathcal{K}$. For such dynamics, we find

$$
x_{t}=\sum_{s=1}^{t-1}\left(\prod_{i=s+1}^{t}\left(A_{i}+B_{i} K_{i}\right)\right) w_{s}=\sum_{s=1}^{t-1} \Phi_{s+1 ; t}\left(K_{s+1: t}\right) w_{t} .
$$

Using $\left\|w_{s}\right\| \leq R_{w}$ and the assumption $\Phi_{s+1 ; t}\left(K_{s+1: t}\right) w_{t} \leq c_{\star} \rho_{\star}^{t-s-1}$ from Assumption 4, we find

$$
\begin{aligned}
\left\|x_{t}\right\| & \leq R_{w} c_{\star} \sum_{s=1}^{t-1} \rho_{\star}^{t-s-1} \leq \frac{c_{\star} R_{w}}{1-\rho_{\star}}=R_{\star} R_{w} \\
\left\|u_{t}\right\|+\left\|x_{t}\right\| & =\left\|x_{t}\right\|+\left\|K_{t} x_{t}\right\| \leq\left\|x_{t}\right\|\left(1+\left\|K_{t}\right\|\right) \leq R_{\star} R_{\mathcal{K}} R_{w}
\end{aligned}
$$

Part b: Since the closed-loop dynamics for $\overline{\mathbf{x}}_{k ; t}^{K}$ and $x_{t}^{K}$ concincide for $t \geq t_{k}$ and are given by $x_{t+1}=\left(A_{t}+B_{t} K\right) x_{t}+w_{t}$, we can compute

$$
\overline{\mathbf{x}}_{k ; t}^{K}-x_{t}^{K}=\left(\prod_{i=t_{k}}^{t}\left(A_{t}+B_{t} K\right)\right)\left(\overline{\mathbf{x}}_{t_{k} ; k}^{K}-x_{t}^{K}\right)
$$

Bounding $\left.\| \overline{\mathbf{x}}_{k ; t}^{K}-x_{t}^{K}\right) \| \leq 2 R_{w} R_{K}$ from part (a) and $\left\|\left(A_{t}+B_{t} K\right)^{t-t_{k}}\right\| \leq c_{\star} \rho_{\star}^{t-t_{k}}$ from Assumption 4 yields $\left.\| \overline{\mathbf{x}}_{k ; t}^{K}-x_{t}^{K}\right) \| \leq 2 R_{x} c_{\star} \rho_{\star}^{t-t_{k}}$. Summing over $t \geq t_{k}$ yields $\sum_{t \geq t_{k}}\left\|\overline{\mathbf{x}}_{k ; t}^{K}-x_{t}^{K}\right\| \leq$ $2 R_{w} R_{\star}^{2}$. Finally, using $\bar{u}_{k ; t}^{K}=K \overline{\mathbf{x}}_{k ; t}^{K}$ and $u_{t}^{K}=K x_{t}^{K}$ gives

$$
\sum_{t \geq t_{k}}\left\|\overline{\mathbf{x}}_{t ; k}^{K}-x_{t}^{K}\right\|+\left\|\bar{u}_{t ; k}^{K}-u_{t}^{K}\right\| \leq(1+\|K\|) \sum_{t \geq t_{k}}\left\|\overline{\mathbf{x}}_{t ; k}^{K}-x_{t}^{K}\right\| \leq 2 R_{w} R_{\mathcal{K}} R_{\star}^{2}
$$

Proof of Lemma E.2. Introducing the short hand $X_{i}=A_{i}+B_{i} K$ and $Y_{i}=A_{i}+B_{i} K^{\prime}$, and expanding the dynamics, and introducing the short hand

$$
\begin{aligned}
\left\|x_{t}^{K}-x_{t}^{K^{\prime}}\right\| & =\|\sum_{s=1}^{t-1}(\prod_{i=s+1}^{t}(\underbrace{A_{i}+B_{i} K}_{=X_{i}})-\prod_{i=s+1}^{t}(\underbrace{A_{i}+B_{i} K^{\prime}}_{=Y_{i}})) w_{s}\| \\
& \leq R_{w} \sum_{s=1}^{t-1}\left\|\prod_{i=s+1}^{t} X_{i}-\prod_{i=s+1}^{t} Y_{i}\right\|_{\mathrm{op}}
\end{aligned}
$$

Using an elementary matrix telescoping identiy,

$$
\prod_{i=s+1}^{t} X_{i}-\prod_{i=s+1}^{t} Y_{i}=\sum_{j=s+1}^{t}\left(\prod_{i=j+1}^{t} X_{j}\right)\left(X_{j}-Y_{j}\right) \prod_{i=s+1}^{j-1} Y_{j}
$$

Thus, invoking stability assumption, Assumption 4, and setting $R_{B} \geq \max _{t}\left\|B_{t}\right\|_{\mathrm{op}}$,

$$
\begin{aligned}
\left\|\prod_{i=s+1}^{t} X_{i}-\prod_{i=s+1}^{t} Y_{i}\right\|_{\mathrm{op}} & \leq \sum_{j=s+1}^{t}\left\|\prod_{i=j+1}^{t} X_{j}\right\|_{\mathrm{op}}\left\|\prod_{i=s+1}^{j-1} Y_{j}\right\|_{\mathrm{op}}\left\|X_{j}-Y_{j}\right\|_{\mathrm{op}} \\
& =\sum_{j=s+1}^{t} c_{\star} \rho_{\star}^{t-j+1} c_{\star} \rho_{\star}^{j-1-(s+1)}\left\|B_{j}\left(K-K^{\prime}\right)\right\|_{\mathrm{op}} \\
& =\frac{c_{\star}^{2}}{\rho_{\star}^{2}} \sum_{j=s+1}^{t} \rho_{\star}^{t-(s+1)}\left\|B_{j}\left(K-K^{\prime}\right)\right\|_{\mathrm{op}} \\
& \leq \varepsilon R_{B} \frac{c_{\star}^{2}}{\rho_{\star}^{2}}(t-s+1) \rho_{\star}^{t-(s+1)}
\end{aligned}
$$

Thus, we find

$$
\begin{aligned}
\left\|x_{t}^{K}-x_{t}^{K^{\prime}}\right\| & \leq \varepsilon R_{w} R_{B} \frac{c_{\star}^{2}}{\rho_{\star}^{2}} \sum_{s=1}^{t-1}(t-s+1) \rho_{\star}^{t-(s+1)} \\
& \leq \varepsilon R_{w} R_{B} \frac{c_{\star}^{2}}{\rho_{\star}^{2}\left(1-\rho_{\star}\right)^{2}} \leq 4 \varepsilon R_{w} R_{B} \frac{c_{\star}^{2}}{\left(1-\rho_{\star}\right)^{2}}=4 \varepsilon R_{w} R_{\star}^{2} R_{B},
\end{aligned}
$$

where in the last step, we use $\rho_{\star} \geq 1 / 2$. Thus, applying Assumption 3 and ??,

$$
\begin{aligned}
\left\|c_{t}\left(x_{t}^{K}, u_{t}^{K}\right)-c_{t}\left(x_{t}^{K^{\prime}}, u_{t}^{K^{\prime}}\right)\right\| & \leq L \max \left\{1,\left\|\left(x_{t}^{K}, u_{t}^{K}\right)\right\|,\left\|\left(x_{t}^{K^{\prime}}, u_{t}^{K^{\prime}}\right)\right\|\right\} \cdot\left\|\left(x_{t}^{K}, u_{t}^{K}\right)-\left(x_{t}^{K^{\prime}}, u_{t}^{K^{\prime}}\right)\right\| \\
& \leq L R_{w} R_{\mathcal{K}} R_{\star} \cdot\left\|\left(x_{t}^{K}, u_{t}^{K}\right)-\left(x_{t}^{K^{\prime}}, u_{t}^{K^{\prime}}\right)\right\| .
\end{aligned}
$$

Continuing, we bound

$$
\begin{aligned}
\left\|\left(x_{t}^{K}, u_{t}^{K}\right)-\left(x_{t}^{K^{\prime}}, u_{t}^{K^{\prime}}\right)\right\| & =\left\|\left(x_{t}^{K}, K x_{t}^{K}\right)-\left(x_{t}^{K^{\prime}}, K^{\prime} x_{t}^{K^{\prime}}\right)\right\| \\
& \leq\left\|\left(x_{t}^{K}, K x_{t}^{K}\right)-\left(x_{t}^{K^{\prime}}, K x_{t}^{K^{\prime}}\right)\right\|+\left\|\left(K-K^{\prime}\right) x_{t}^{K^{\prime}}\right\| \\
& \leq R_{\mathcal{K}}\left\|x_{t}^{K^{\prime}}-x_{t}^{K^{\prime}}\right\|+\left\|\left(K-K^{\prime}\right) x_{t}^{K^{\prime}}\right\|
\end{aligned}
$$

Finally, using the bound $\left\|x_{t}^{K^{\prime}}-x_{t}^{K^{\prime}}\right\| \leq 4 \varepsilon R_{w} R_{\star}^{2} R_{B}$ derived above, and bounding $\|(K-$ $\left.K^{\prime}\right) x_{t}^{K^{\prime}}\|\leq \varepsilon\| x_{t}^{K^{\prime}} \| \leq \varepsilon R_{w} R_{\star}$ in view of Lemma E.1. Hence,

$$
\begin{aligned}
\left\|c_{t}\left(x_{t}^{K}, u_{t}^{K}\right)-c_{t}\left(x_{t}^{K^{\prime}}, u_{t}^{K^{\prime}}\right)\right\| & \leq L R_{w} R_{\mathcal{K}} R_{\star}\left(4 \varepsilon R_{\mathcal{K}} R_{w} R_{\star}^{2} R_{B}+\varepsilon R_{w} R_{\star}\right) \\
& \leq 4 \varepsilon L R_{w}^{2} R_{\mathcal{K}}^{2} R_{\star}^{3}\left(R_{B}+1\right)
\end{aligned}
$$

where above we use that $R_{\mathcal{K}}, R_{\star} \geq 1$ by assumption.

## F Lower Bounds and Separations

## F. 1 Separation between policy classes

Let $\mathcal{Z}=\left(c_{t}, w_{t}, A_{t}, B_{t}\right)_{t \geq 1}$ denote sequences over costs,disturbances, and dynamics. We let $J_{T}(\pi ; \mathcal{Z})$ denote the cost of policy $\pi$ on the sequence $\mathcal{Z}$. Our lower bounds hold even against sequences which enjoy the following regularity condition.
Definition F.1. We say that $\mathcal{Z}$ is regular if, for all $t, c_{t}(\cdot, \cdot)$ satisfies ?? with $L \leq 1$, and that for all $t$, $\left\|w_{t}\right\| \leq 1,\left\|B_{t}\right\|_{\mathrm{op}} \leq 1$ and $\left\|A_{t}\right\|_{\mathrm{op}} \leq 1 / 2$.

We define the policy classes

$$
\begin{aligned}
\Pi_{\mathrm{drc}}^{+}(h) & :=\left\{\pi: u_{t}^{\pi}=u_{0}+\sum_{i=0}^{h-1} M^{[i]} x_{t-i}^{\mathrm{nat}}, \forall t\right\} \\
\Pi_{\mathrm{dac}}^{+}(h) & :=\left\{\pi: u_{t}^{\pi}=u_{0}+\sum_{i=0}^{h-1} M^{[i]} w_{t-i-1}, \forall t\right\} \\
\Pi_{\text {feed }}^{+}(h) & :=\left\{\pi: u_{t}^{\pi}=u_{0}+\sum_{i=0}^{h-1} M^{[i]} x_{t-i}^{\pi}, \forall t\right\} \\
\Pi_{\mathrm{feed}}(h, R) & :=\left\{\pi: u_{t}^{\pi}=u_{0}+\sum_{i=0}^{h-1} M^{[i]} x_{t-i}^{\pi}, \forall t, \quad \sum_{i=0}^{h-1}\left\|M^{[i]}\right\|_{\mathrm{op}} \leq R .\right\} .
\end{aligned}
$$

That is, $\Pi_{\mathrm{drc}}^{+}(h)$ are all length $h$ DRC policies of unbounded norm and allowing affine offsets, $\Pi_{\mathrm{dac}}^{+}(h)$ are all length $h$ DAC policies of unbounded norm allowing affine offsets, and $\Pi_{\text {feed }}^{+}(h)$ are all static feedback policies of unbounded norm allowing affine offsets, and $\Pi_{\text {feed }}(h, R)$ are feedback policies of bounded norm and horizon.
The following theorem demonstrates that the DAC, DRC, and feedback parametrizations are fundamentally incommensurate.

Theorem F.1. Let $\mathcal{C}_{0}>0$ denote a universal constant. There exists three regular sequences $\mathcal{Z}_{1}, \mathcal{Z}_{2}, \mathcal{Z}_{3}$ in $d_{u}=d_{x}=1$ which separate DAC, DRC, and feedback controllers, in the following sense:
(a) Under $\mathcal{Z}_{1}$, the static feedback policy $\pi$ selecting $u_{t}^{\pi}=\frac{1}{4} x_{t}^{\pi}$ satisfies $J_{T}\left(\pi ; \mathcal{Z}_{1}\right)=0$, but

$$
\inf _{\pi \in \Pi_{\mathrm{drc}}^{+}(h) \cup \Pi_{\mathrm{dac}}^{+}(h)} J_{T}\left(\pi ; \mathcal{Z}_{1}\right) \geq \mathcal{C}_{0}(T-h-2)
$$

(b) Under $\mathcal{Z}_{2}$, the DAC policy $\pi$ selecting $u_{t}^{\pi}=w_{t}$ satisfies $J_{T}\left(\pi ; \mathcal{Z}_{2}\right)=0$, but

$$
\inf _{\pi \in \Pi_{\mathrm{drc}}^{+}(h) \cup \Pi_{\text {feed }}^{+}(h)} J_{T}\left(\pi ; \mathcal{Z}_{2}\right) \geq \mathcal{C}_{0}(T-h-2)
$$

(c) Under $\mathcal{Z}_{3}$, the DRC policy $\pi$ selecting $u_{t}^{\pi}=x_{t}^{\text {nat }}$ satifies $J_{T}\left(\pi ; \mathcal{Z}_{3}\right)=0$, but

$$
\inf _{\pi \in \Pi_{\text {dac }}^{+}(h)} J_{T}\left(\pi ; \mathcal{Z}_{3}\right) \geq \mathcal{C}_{0}(T-h-3)
$$

Moreover, for any $h \in N, R>0$ and $T \geq 10 h$, we have

$$
\inf _{\pi \in \Pi_{\text {feed }}(h, R)} J_{T}\left(\pi ; \mathcal{Z}_{3}\right) \geq \mathcal{C}_{0} \frac{T}{h \max \{R, 1\}}
$$

Proof. We establish the separations for each part with different constant factors. One can choose $\mathcal{C}_{0}$ to be the minimum of all constants which arise.

Proof of part a. We set $\mathcal{Z}_{1}$ to be the sequence with $A_{t}=0$ for all $t, w_{t}=1$ for all $t, c_{t}(x, u)=$ $\frac{1}{8}\left(u-\frac{1}{4} x\right)^{2}$, and

$$
B_{t}= \begin{cases}1 & t \text { odd } \\ -1 & t \text { even }\end{cases}
$$

This sequence is clearly regular, and is clear that the policy $u_{t}^{\pi}=\frac{1}{4} x_{t}^{\pi}$ has $J_{T}\left(\pi ; \mathcal{Z}_{1}\right)=0$. On the other hand, let $\pi \in \Pi_{\text {drc }}^{+}(h) \cup \Pi_{\text {dac }}^{+}(h), J_{T}\left(\pi ; \mathcal{Z}_{1}\right)$. Since $A_{t} \equiv 0, x_{t}^{\text {nat }}=w_{t}$ so $\Pi_{\text {drc }}^{+}(h)=\Pi_{\text {dac }}^{+}(h)$. Moreover, since $w_{t}=1$ for all $t \geq 1$, any $\pi \in \Pi_{\text {dac }}^{+}(h)$ has $u_{t}^{\pi}=\bar{u}$ for some fixed $\bar{u}$ for all $t>h$. Then, for all $t>h+1, x_{t}=1+B_{t-1} \bar{u}$. Thus, $c_{t}\left(x_{t}^{\pi}, u_{t}^{\pi}\right)=\frac{1}{8}\left(\bar{u}-\frac{1+B_{t-1} \bar{u}}{4}\right)^{2}$. Using the definition of $B_{t}$,
$c_{t}\left(x_{t}^{\pi}, u_{t}^{\pi}\right)+c_{t+1}\left(x_{t+1}^{\pi}, u_{t+1}^{\pi}\right)=\frac{1}{8}\left(\bar{u}-\frac{1-\bar{u}}{4}\right)^{2}+\frac{1}{8}\left(\bar{u}-\frac{1+\bar{u}}{4}\right)^{2}=\frac{1}{128}\left((3 \bar{u}+1)^{2}+(5 \bar{u}+1)\right)^{2}=\Omega(1)$.
The bound follows.
Proof of part b. Set $c_{t}(x, u)=\left(u-w_{t-1}\right)^{2}$. Then the DAC policy $u_{t}^{\pi}=w_{t-1}$ has zero cost. Further, set $B_{t} \equiv 0$, thus, $x_{t} \equiv x_{t}^{\text {nat }}$, so $\Pi_{\text {feed }}^{+}(h)$ and $\Pi_{\mathrm{drc}}^{+}(h)$ are equivalent on this system. Finally, let $n=2 m+1$, and set $w_{1}=1$, and for $t \geq 1$, set

$$
\left(A_{t}, w_{t}\right)= \begin{cases}\left(\frac{1}{2}, \frac{1}{2}\right) & t \text { is even } \\ \left(\frac{1}{4}, \frac{3}{4}\right) & t \text { is odd }\end{cases}
$$

Then, one can verify via induction that $x_{t}=x_{t}^{\text {nat }}=1$ for all $t \geq 2$. Hence, for all $t \geq h+1$, any $\pi \in \Pi_{\text {feed }}^{+}(h) \cup \Pi_{\mathrm{drc}}^{+}(h)$ has a constant input $u_{t}^{\pi}=\bar{u}$. However, $c_{t}\left(x_{t}^{\pi}, u_{t}^{\pi}\right)+c_{t+1}\left(x_{t}^{\pi}, u_{t}^{\pi}\right)=\left(\bar{u}-\frac{1}{2}\right)^{2}+\left(\bar{u}-\frac{3}{4}\right)^{2}$, which is greater than a universal constant. Hence, the regret must spaces as $\Omega(T-(h+2))$.

Proof of part c. Fix policity $\pi_{\star}$ to select $u_{t}^{\pi_{\star}}=x_{t-1}^{\text {nat }}$. Denote the sequences that arise from this policy as $\left(x_{t}^{\star}, u_{t}^{\star}\right)$. We set

$$
c_{t}(x, u)=\frac{1}{4}\left(\left(u-\frac{1}{2} u_{t}^{\star}\right)^{2}+\left|x-x_{t}^{\star}\right|\right)
$$

By construction $\pi_{\star}$ has zero cost on $c_{t}$. Now, set $w_{t}=1$ for all $t$, and

$$
\begin{gathered}
A_{t}=\left\{\begin{array}{ccc}
\frac{1}{4} & t & \bmod 3=1 \\
\frac{1}{4} & t & \bmod 3=2 \\
0 & t & \bmod 3=0 .
\end{array}, \quad B_{t}=\left\{\begin{array}{lll}
-\frac{1}{4} & t & \bmod 3=1 \\
-\frac{1}{20} & t & \bmod 3=1 \\
0 & t & \bmod 3=0
\end{array}\right.\right. \\
u_{3 k+1}^{\star}=x_{3 k+1}^{\mathrm{nat}}=1 \\
u_{3 k+2}^{\star}=x_{3 k+2}^{\mathrm{nat}}=\frac{5}{4} \\
u_{3 k+3}^{\star}=x_{3 k+3}^{\mathrm{nat}}=\frac{21}{16} .
\end{gathered}
$$

Hence, a similar argument as in part (b), using the fact that that $w_{t}$ is constant but $u_{t}^{\star}$ is periodic, shows that any $\pi \in \Pi_{\text {dac }}^{+}(h)$ suffers cost $\Omega(T-h-3)$.
Let us now analyze the performance of policies $\pi \in \Pi_{\text {feed }}\left(m, R_{M}\right)$. First, observe that $x_{t}^{\star}=1$ for all $t \geq 2$.

$$
\begin{aligned}
x_{3 k+1}^{\star} & =1 \\
x_{3 k+2}^{\star} & =\frac{5}{4}-\frac{1}{4} u_{3 k+1}^{\star}=1 \\
x_{3 k+2}^{\star} & =\frac{21}{16}-\frac{1}{4} u_{3 k+1}^{\star}-\frac{1}{20} u_{3 k+2}^{\star} \\
& =\frac{21}{16}-\frac{1}{4}-\frac{1}{20} \cdot \frac{5}{4}=1
\end{aligned}
$$

Next, observe that for $\pi \in \Pi_{\text {feed }}(h, R)$, and $t \geq h+1$,

$$
u_{t}^{\pi}=c+\sum_{i=0}^{h-1} M^{[i]} x_{t}^{\pi}=c+\sum_{i=0}^{h-1} M^{[i]} x_{t}^{\star}+\sum_{i=0}^{h-1} M^{[i]}\left(x_{t}^{\pi}-x_{t}^{\star}\right)=\underbrace{\left(c+\sum_{i=0}^{h-1} M^{[i]}\right)}_{:=\bar{u}}+\sum_{i=0}^{h-1} M^{[i]}\left(x_{t}^{\pi}-x_{t}^{\star}\right)
$$

Defining $\epsilon_{t}=u_{t}^{\pi}-\bar{u}$, we have

$$
\left|\epsilon_{t}\right| \leq\left|\sum_{i=0}^{h-1} M^{[i]}\left(x_{t}^{\pi}-x_{t}^{\star}\right)\right| \leq R{\underset{i=0}{h-1}}_{\max _{i=0}}\left|x_{t-i}^{\pi}-x_{t-i}^{\star}\right|
$$

Hence, for integers $k$,

$$
\begin{equation*}
\max _{i \in[3]}\left|\epsilon_{3 k+i}\right| \leq R \max _{t=3 k-h+2}^{3 k+3}\left|x_{t}^{\pi}-x_{t}^{\star}\right| . \tag{F.1}
\end{equation*}
$$

We now argue a dichotomoty on the size of $\max _{i \in[3]}\left|\epsilon_{3 k+i}\right|$. First, we show that if the epsilons are large, the costs incurred on a past window of $h$ must be as well. This is Eq. (F.1) would necessitate that $x_{t}^{\pi}$ differs from $x_{t}^{\star}$ over the previous window.

Claim F.1. Suppose $\max _{i \in[3]}\left|\epsilon_{3 k+i}\right| \geq \frac{1}{32}$. Then, $\sum_{t=3 k-h+2}^{3 k+3} c_{t}\left(x_{t}^{\pi}, u_{t}^{\pi}\right) \geq \frac{1}{2^{7} R}$.

Proof. By Eq. (F.1), we have that if $\max _{i \in[3]}\left|\epsilon_{3 k+i}\right| \geq \frac{1}{32}$, then $\max _{t=3 k-h+2}^{3 k+3}\left|x_{t}^{\pi}-x_{t}^{\star}\right| \geq \frac{1}{32 R}$. Since $c_{t}\left(x_{t}^{\pi}, u_{t}^{\pi}\right) \geq \frac{1}{4}\left|x_{t}^{\pi}-x_{t}^{\star}\right|$, the bound follows by upper bounding the maximum with the sum.

On the other hand, we show that if the $\epsilon$-terms are small, then the costs on $t \in\{3 k+1,3 k+2,3 k+3\}$ are at least a small constant. This is because the inputs selected by $\pi, \bar{u}+\epsilon_{t}$, are close to constant, and therefore can fit the periodic values of $u_{t}^{\pi_{\star}}$.

Claim F.2. Suppose $\max _{i \in[3]}\left|\epsilon_{3 k+i}\right| \leq \frac{1}{32}$. Then, $\sum_{i=1}^{3} c_{3 k+i}\left(x_{t}^{\pi}, u_{t}^{\pi}\right) \geq 2^{-12}$.

Proof. We expand

$$
\begin{aligned}
\sum_{i=1}^{3} c_{3 k+i}\left(x_{t}^{\pi}, u_{t}^{\pi}\right) & \geq \sum_{i=1}^{3} \frac{1}{4}\left(u_{t}^{\star}-u_{t}^{\pi}\right)^{2} \\
& \geq \sum_{i=1}^{3} \frac{1}{4}\left(u_{3 k+i}^{\star}-\bar{u}-\epsilon_{3 k+i}\right)^{2} \\
& =\frac{1}{4}\left(\left(1-\bar{u}-\epsilon_{3 k+1}\right)^{2}+\left(1+\frac{1}{4}-\bar{u}-\epsilon_{3 k+1}\right)^{2}+\left(1+\frac{1}{4}+\frac{1}{16}-\bar{u}-\epsilon_{3 k+i}\right)^{2}\right)
\end{aligned}
$$

In particular, suppose $\max _{i \in[3]}\left|\epsilon_{3 k+i}\right| \leq \frac{1}{32}$. Then unless $|\bar{u}-1| \leq \frac{1}{16}$, the above is at least $\frac{1}{4}\left(\frac{1}{32}\right)^{2}=2^{-12}$. On the other hand, if $|\bar{u}-1| \leq \frac{1}{16}$. Then, $\frac{1}{4}\left(1+\frac{1}{4}-\bar{u}-\epsilon_{3 k+1}\right)^{2} \geq\left(\frac{1}{4}-\frac{1}{16}-\right.$ $\left.\frac{1}{32}\right)^{2} \geq \frac{1}{4}\left(\frac{1}{8}\right)^{2} \geq 2^{-12}$.

Combining both cases, we find that, for all $k$ such that $3 k-h+2 \geq 1$,

$$
\sum_{t=3 k-h+2}^{3 k+3} c_{t}\left(x_{t}^{\pi}, u_{t}^{\pi}\right) \geq \frac{2^{-17}}{\max \{R, 1\}}
$$

In particular, we find that for $J_{T}\left(\pi ; \mathcal{Z}_{3}\right) \geq \Omega\left(\frac{T}{h \max \{1, R\}}\right)$ provided (say) $T \geq 10 h$.

## F. 2 Linear Regret Against DRC and DAC

In this section, we demonstrate linear regret against DRC and DAC. We consider a distribution over instances of the following form

$$
A_{t}=0, \quad B_{t}=\left[\begin{array}{ccc}
1 & 0 & 0  \tag{F.2}\\
0 & \beta_{t} & 0 \\
0 & 0 & 1
\end{array}\right], \quad w_{t}=-\left[\begin{array}{c}
\omega_{t-1} \\
\omega_{t} \\
1
\end{array}\right], t \geq 0
$$

Define the constant $\alpha=1 / 4$. For a given $\sigma \in(0,1 / 8)$, let $\mathcal{D}_{\sigma}$ denote the distribution over ( $A_{t}, B_{t}, w_{t}$ ) induced by drawing

$$
\beta_{t} \stackrel{\text { i.i.d }}{\sim}[1-\sigma, 1+\sigma], \quad \omega_{t} \stackrel{\text { i.i.d }}{\sim}\{1-\alpha \sigma, 1+\alpha \sigma\} .
$$

Note that these instances are (a) controllable, (b) stable, and (c) have variance scaling like $T \sigma^{2}$, and (d) the DRC and DAC parametrizations coincide. Letting $v[i]$ denote the $i$-th coordinate of vectors $v$, we consider cost of the form

$$
\begin{equation*}
c_{f}(x, u)=x[2]^{2}+u[2]^{2}+f(x[1]) \tag{F.3}
\end{equation*}
$$

for either $f(z)=|z|$ or $f(z)=z^{2}$. Note that both choices of $f$ ensure that $c_{f}$ satisfies ??, and the latter choice ensures that $c_{f}(x, u)$ is second order smooth.
Theorem F.2. Let alg be any online learning algorithm. Let $c_{f}(x, u)$ as in (F.3). Then, for any $\sigma>0$, there exist a DRC policy $\pi_{\star} \in \Pi_{\mathrm{drc}}(1,1)^{4}$ such that expected regret incurred by alg under the distribution $\mathcal{D}_{\sigma}$ and cost $c_{f}(x, u)$ is at least

$$
\mathbb{E}_{\mathcal{D}_{\sigma}, \text { alg }}\left[J_{T}(\text { alg })-J_{T}\left(\pi_{\star}\right)\right] \geq \cdot \begin{cases}C_{1} T \sigma^{2} & \text { for } f(z)=z^{2}, T \geq \frac{C_{0}}{\sigma^{2}} \\ C_{1} T \sigma & \text { for } f(z)=z, T \geq \frac{C_{0}}{\sigma}\end{cases}
$$

where above, $C_{0}, C_{1}$ are universal, positive constants.

Proof. In what follows, $\mathbb{E}[\cdot]$ denotes expectation under the nstance from $\mathcal{D}_{\sigma}$, and any algorithmic randomness.

[^3]We expand $x_{t}$ and $u_{t}$ into its coordinates $x_{t}=\left(x_{t ; 1}, x_{t ; 2}, x_{t ; 3}\right)$ and of $u_{t}=\left(u_{t ; 1}, u_{t ; 2}, u_{t ; 3}\right)$. For any policy $\pi$, we can decompose its cost as

$$
\begin{aligned}
J_{T}(\pi)= & \sum_{t=1}^{T} c_{f}\left(x_{t}^{\pi}, u_{t}^{\pi}\right)=\left(u_{T ; 2}^{\pi}\right)^{2}+f\left(x_{T ; 1}^{\pi}\right)+\sum_{t=1}^{T-1} c_{\mathrm{q}}\left(x_{t+1 ; 2}^{\pi}, u_{t ; 2}^{\pi}\right)+f\left(x_{t ; 1}^{\pi}\right) \\
& \text { where } c_{\mathrm{q}}\left(x_{t+1 ; 2}^{\pi}, u_{t ; 2}^{\pi}\right)=\left(x_{t+1 ; 2}^{\pi}\right)^{2}+\left(u_{t ; 2}^{\pi}\right)^{2}
\end{aligned}
$$

Note that $x_{1}=0$ in the above. The following lemma characterizes the conditional expectation of the $c_{\mathrm{q}}$

Lemma F.3. There exist a constant $c_{\mathrm{q} ; \star}$ such that

$$
\mathbb{E}\left[c_{\mathrm{q}}\left(x_{t+1 ; 2}, u_{t ; 2}\right)\right]-c_{\mathrm{q} ; \star}=\left(2+\sigma^{2} / 3\right) \mathbb{E}\left[\left(u_{t ; 2}-\bar{u}\right)^{2}\right],
$$

where $\bar{u}=\frac{1}{2\left(1+\sigma^{2} / 6\right)}$.
The proof of Lemma F. 3 is deferred to the end of the section.
Bounding the cost of $\pi^{\star}$ We select $\pi^{\star}$ to be DAC (or equivalently, DRC since $A_{t} \equiv 0$ ) policy given by

$$
M=\left[\begin{array}{ccc}
0 & 1 & 0 \\
0 & 0 & \bar{u} \\
0 & 0 & 0
\end{array}\right], \quad u_{t}^{M}=M w_{t-1} .
$$

We find that

$$
\begin{aligned}
& \mathbb{E}\left[c_{\mathrm{q}}\left(x_{t+1 ; 2}^{\pi_{\star}}, u_{t ; 2}^{\pi_{\star}}\right)\right]-c_{\mathrm{q} ; \star}=\left(2+\sigma^{2} / 3\right) \mathbb{E}\left[\left(u_{t ; 2}^{M}-\bar{u}\right)^{2}\right]=0, \quad 2 \leq t \leq T-1 \\
& \mathbb{E}\left[f\left(x_{t ; 1}^{M}\right)\right]=\mathbb{E}\left[f\left(u_{t-1 ; 1}^{M}-\omega_{t-2}\right)\right]=0, \quad 3 \leq t \leq T .
\end{aligned}
$$

Since the noise is uniformly bounded independent for all $\sigma, \alpha \leq 1$, we conclude

$$
\begin{align*}
\mathbb{E}\left[J_{T}\left(\pi^{M}\right)\right]-(T-1) c_{\mathrm{q} ; \star} & =\underbrace{\mathbb{E}\left[f\left(x_{t ; 2}^{\pi}\right)\right]}_{\mathbb{E}\left[f\left(\omega_{t}\right)\right]}+\left(u_{T ; 2}^{\pi_{\star}}\right)^{2}+\underbrace{\mathbb{E}\left[f\left(x_{T ; 1}^{\pi}\right)\right]}_{=0}+\underbrace{\mathbb{E}\left[c_{f}\left(x_{1}, u_{1}\right)\right]}_{=0} \\
& =\frac{f(1-\alpha \sigma)+f(1+\alpha \sigma)}{2}+\underbrace{\left(\bar{u} \cdot w_{t-1 ; 3}\right)^{2}}=\bar{u}^{2} \leq 2 ., \tag{F.4}
\end{align*}
$$

where we use $\alpha \leq 1 / 24, \sigma \leq 1 / 8$, and $\bar{u} \leq \frac{1}{2\left(1+\sigma^{2} / 6\right)}$, and $f(z) \leq \max \left\{|z|,|z|^{2}\right\}$ to achieve the bound Eq. (F.4).

Bounding the cost of adaptive policies Fix any online learning algorithm alg; we lower bound its performance. Because $B_{t}$ and $w_{t}$ are drawn from a fixed probability distribution, and are therefore oblivious to the learner's actions, we may assume without loss of generality that alg is deterministic.

The first step is to argue that any algorithm with small cost must select inputs where are bounded away from zero. Specifically, define the devent $\mathcal{E}_{t}:=\left\{u_{t ; 2}^{\text {alg }} \geq 1 / 6\right\}$. Using Lemma F. 3 together with $\bar{u} \geq 1 / 3$,

$$
\begin{align*}
\mathbb{E}\left[J_{T}(\text { alg })\right]-T c_{\mathrm{q} ; \star} & \geq 2 \sum_{t=1}^{T} \mathbb{E}\left[f\left(x_{t ; 1}+\left(u_{t ; 2}^{\mathrm{alg}}-\bar{u}\right)^{2}\right]\right] \\
& \geq 2 \sum_{t=1}^{T} \mathbb{E}\left[f\left(x_{t ; 1}^{\mathrm{alg}}\right)\right]+\left(\frac{1}{3}-\frac{1}{6}\right)^{2} \mathbb{P}\left[\mathcal{E}_{t}^{c}\right] \\
& =\sum_{t=1}^{T} \mathbb{E}\left[f\left(x_{t ; 1}^{\mathrm{alg}}\right)\right]+\frac{1}{18} \mathbb{P}\left[\mathcal{E}_{t}^{c}\right] \\
& \geq \sum_{t=3}^{T} \mathbb{E}\left[f\left(x_{t ; 1}^{\mathrm{alg}}\right) \mid \mathcal{E}_{t-1}\right] \mathbb{P}\left[\mathcal{E}_{t-2}\right]+\frac{1}{18} \sum_{t=1}^{3} \mathbb{P}\left[\mathcal{E}_{t}^{c}\right] . \tag{F.5}
\end{align*}
$$

We now lower bound $\mathbb{E}\left[f\left(x_{t ; 1}^{\text {alg }}\right) \mid \mathcal{E}_{t-2}\right]$, again deferring the proof to the end of the section.

Lemma F.4. For any executable policy $\pi \mathbb{E}\left[f\left(x_{t ; 1}^{\pi}\right) \mid \mathcal{E}_{t-2}\right] \geq \frac{1}{2} f(\alpha \sigma)$, provided $\alpha \leq 1 / 24$.
Combining Lemma F. 4 with the above bound, we have

$$
\begin{aligned}
\mathbb{E}\left[J_{T}(\mathrm{alg})\right]-(T-1) c_{\mathrm{q} ; \star} & \geq \sum_{t=1}^{T} \frac{f(\alpha \sigma)}{2} \mathbb{P}\left[\mathcal{E}_{t-2}\right]+\frac{1}{18} \sum_{t=1}^{T} \mathbb{P}\left[\mathcal{E}_{t}^{c}\right] \\
& \geq \min \left\{\frac{f(\alpha \sigma)}{2}, \frac{1}{18}\right\} \sum_{t=1}^{T-2} \mathbb{P}\left[\mathcal{E}_{t}\right]+\mathbb{P}\left[\mathcal{E}_{t}^{c}\right] \\
& \geq(T-2) \min \left\{\frac{f(\alpha \sigma)}{2}, \frac{1}{18}\right\}
\end{aligned}
$$

With $\sigma \leq 1$ and $\alpha=\frac{1}{24}, \frac{f(\alpha \sigma)}{2} \leq \frac{1}{18}$. Combining with Eq. (F.4), we have

$$
\mathbb{E}\left[J_{T}(\text { alg })\right]-\mathbb{E}\left[J_{T}\left(\pi^{\star}\right)\right] \geq(T-2) \frac{f(\alpha \sigma)}{2}-2
$$

The bound follows.

## F.2.1 Omitted proofs

Proof of Lemma F.3. Let $\left(\overline{\mathcal{F}}_{t}\right)_{t \geq 1}$ denote the filtration induced by setting $\overline{\mathcal{F}}_{t}$ to be the sigma-algebra generated by $\left(\beta_{1}, \ldots, \beta_{t-1}, \omega_{1}, \ldots, \omega_{t}\right\}$. We have

$$
\begin{aligned}
\mathbb{E}\left[\beta_{t}^{2} \mid \overline{\mathcal{F}}_{t}\right] & =\frac{1}{2 \sigma} \int_{u=1-\sigma}^{1+\sigma} u^{2} \mathrm{~d} u=\frac{(1+\sigma)^{3}-(1-\sigma)^{3}}{6 \sigma} \\
& =\frac{1+3 \sigma+3 \sigma^{2}+\sigma^{3}-\left(1-3 \sigma^{2}+3 \sigma^{2}-\sigma^{3}\right)}{6 \sigma} \\
& =1+\sigma^{3} / 3 .
\end{aligned}
$$

Set $c_{\mathrm{q}}(x, u)=x^{2}+u^{2}$.

$$
\begin{aligned}
\mathbb{E}\left[c_{\mathrm{q}}\left(x_{t+1 ; 2}, u_{t ; 2}\right) \mid \overline{\mathcal{F}}_{t}\right] & =u_{t ; 2}^{2}+\mathbb{E}\left[\left(\beta_{t} u_{t ; 2}-\omega_{t}\right)^{2} \mid \overline{\mathcal{F}}_{t}\right] \\
& =u_{t ; 2}^{2}\left(1+\mathbb{E}\left[\beta_{t}^{2} \mid \overline{\mathcal{F}}_{t}\right]\right)-2 u_{t ; 2} \mathbb{E}\left[\beta_{t} w_{t} \mid \overline{\mathcal{F}}_{t}\right]+\mathbb{E}\left[\omega_{t}^{2} \mid \overline{\mathcal{F}}_{t}\right] \\
& =u_{t ; 2}^{2}\left(2+\sigma^{2} / 3\right)-2 u_{t ; 2}+\left(1+c \sigma^{2}\right)
\end{aligned}
$$

at $u_{\star ; 2}=\frac{1}{2\left(1+\sigma^{2} / 6\right)}$. Define $c_{\mathrm{q} ; \star}:=\min _{u_{t ; 2}} \mathbb{E}\left[\ell_{2}\left(x_{t ; 2}, u_{t ; 2}\right) \mid \overline{\mathcal{F}}_{t}\right]$, we then have

$$
\begin{equation*}
\mathbb{E}\left[c_{\mathbf{q}}\left(x_{t+1 ; 2}, u_{t ; 2}\right) \mid \overline{\mathcal{F}}_{t}\right]-\ell_{2 ; \star}=\left(2+\sigma^{2} / 3\right)\left(u_{t ; 2}-u_{\star ; 2}\right)^{2} \tag{F.6}
\end{equation*}
$$

Proof of Lemma F.4. Let us introduce a second event, $\tilde{\mathcal{E}}$, defined as

$$
\tilde{\mathcal{E}}_{t-1}:=\left\{(1+\alpha \sigma)+u_{t ; 2}^{\mathrm{alg}}(1-\sigma) \leq x_{t-1 ; 2}^{\mathrm{alg}} \leq(1-\alpha \sigma)+u_{t ; 2}^{\mathrm{alg}}(1+\sigma)\right\}
$$

Then, since $f$ is non-negative,

$$
\begin{aligned}
\mathbb{E}\left[f\left(x_{t ; 1}^{\pi}\right) \mid \mathcal{E}_{t-2}\right] & \geq \mathbb{E}\left[\mathbb{I}\left\{\tilde{\mathcal{E}}_{t-1}\right\} f\left(x_{t ; 1}^{\pi}\right) \mid \mathcal{E}_{t-2}\right] \\
& =\mathbb{P}\left[\tilde{\mathcal{E}}_{t-1} \mid \mathcal{E}_{2}\right] \mathbb{E}\left[\mathbb{E}\left[f\left(x_{t ; 1}^{\pi}\right) \mid \tilde{\mathcal{E}}_{t-1} \cap \mathcal{E}_{t-2}\right]\right.
\end{aligned}
$$

Let's first lower bound $\mathbb{P}\left[\tilde{\mathcal{E}}_{t-1} \mid \mathcal{E}_{2}\right]$. Writing $x_{t-1 ; 2}^{\text {alg }}=\beta_{t-2} u_{t ; 2}^{\text {alg }}+\omega_{t-2}, \tilde{\mathcal{E}}_{t-1}$ occurs as soon as

$$
\begin{aligned}
& (1+\alpha \sigma)+u_{t-2 ; 2}^{\mathrm{alg}}(1-\sigma) \leq \beta_{t-2} u_{t-2 ; 2}^{\mathrm{alg}}+\omega_{t-2} \\
& (1-\alpha \sigma)+u_{t-2 ; 2}^{\mathrm{alg}}(1+\sigma) \geq \beta_{t-2} u_{t-2 ; 2}^{\mathrm{alg}}+\omega_{t-2}
\end{aligned}
$$

Using that $1-\alpha \sigma \leq \omega_{t-2} \leq 1+\alpha \sigma$ and rearranging the above, it is enough that

$$
\begin{aligned}
2 \alpha \sigma & \leq\left(\beta_{t}-(1-\sigma)\right) u_{t-2 ; 2}^{\mathrm{alg}} \\
-2 \alpha \sigma & \geq-\left((1+\sigma)-\beta_{t}\right) u_{t-2 ; 2}^{\mathrm{alg}}
\end{aligned}
$$

Now, if $\mathcal{E}_{t-2}$ holds, that $u_{t ; 2}^{\text {alg }} \geq 1 / 6$. Furthermore, by construction $1-\sigma \leq \beta_{t} \leq 1+\sigma$. Therefore, if $\mathcal{E}_{t-2}$ holds, then $\tilde{\mathcal{E}}_{t-1}$ holds as long as

$$
\beta_{t}-(1-\sigma) \geq 12 \alpha \sigma \quad \text { and } \quad(1+\sigma)-\beta_{t} \geq 12 \alpha \sigma
$$

In particular, for $\alpha=\frac{1}{24}$, then

$$
\begin{equation*}
\mathbb{P}\left[\tilde{\mathcal{E}}_{t-1} \mid \mathcal{E}_{2}\right] \geq \mathbb{P}\left[\beta_{t} \in\left[1-\frac{2}{\sigma}, 1+\frac{2}{\sigma}\right]\right]=\frac{1}{2} \tag{F.7}
\end{equation*}
$$

Next, we lower bound $\mathbb{E}\left[f\left(x_{t ; 1}^{\pi}\right) \mid \tilde{\mathcal{E}}_{t-1} \cap \mathcal{E}_{t-2}\right]$. To do so, we observe that $x_{t ; 1}=\omega_{t-2}-u_{t-1 ; 1}^{\text {alg }}$. Moreover, since alg is deterministic (see discussion above), $u_{t-1 ; 1}^{\text {alg }}$ is a deterministic function of $x_{1: t-1}^{\text {alg }}$, and moreover, $\tilde{\mathcal{E}}_{t-1}, \mathcal{E}_{t-2}$ are determinied by $x_{1: t-1}^{\text {alg }}$. Hence,

$$
\begin{aligned}
\mathbb{E}\left[f\left(x_{t ; 1}^{\pi}\right) \mid \tilde{\mathcal{E}}_{t-1}, \mathcal{E}_{t-2}\right] & =\mathbb{E}\left[\mathbb{E}\left[f\left(x_{t ; 1}^{\pi}\right) \mid x_{1: t-1}^{\mathrm{alg}}\right] \mid \tilde{\mathcal{E}}_{t-1} \cap \mathcal{E}_{t-2}\right] \\
& =\mathbb{E}\left[\mathbb{E}\left[f\left(\omega_{t-2}-u_{t-1 ; 1}^{\mathrm{alg}}\right) \mid x_{1: t-1}^{\mathrm{alg}}\right] \mid \tilde{\mathcal{E}}_{t-1} \cap \mathcal{E}_{t-2}\right] \\
& \geq \mathbb{E}\left[\min _{u \in \mathbb{R}} \mathbb{E}\left[f\left(\omega_{t-2}-u\right) \mid x_{1: t-1}^{\mathrm{alg}}\right] \mid \tilde{\mathcal{E}}_{t-1} \cap \mathcal{E}_{t-2}\right]
\end{aligned}
$$

Thus, it suffices to characterize the distribution of $\omega_{t-1} \mid x_{1: t-1}^{\text {alg }}$ whenever the events $\tilde{\mathcal{E}}_{t-1} \cap \mathcal{E}_{t-2}$. Indeed, we claim that $\omega_{t-2} \mid x_{1: t-1}^{\text {alg }}$ is uniform on $\{1-\alpha \sigma, 1+\alpha \sigma\}$ on whenever $\tilde{\mathcal{E}}_{t-1}$ holds.
Claim F.5. Let $\omega^{-}=1-\alpha \sigma$ and $\omega^{+}=1+\alpha \sigma$. then (with probability one) $\mathbb{P}\left[\omega_{t-2}=\omega^{-} \mid\right.$ $\left.x_{1: t-1}^{\mathrm{alg}}\right]=\mathbb{P}\left[\omega_{t-2}=\omega^{+} \mid x_{1: t-1}^{\mathrm{alg}}\right]=\frac{1}{2}$ if $\tilde{\mathcal{E}}_{t-1}$ holds.

Proof. If $\tilde{\mathcal{E}}_{t-1}$ holds, then there exists exactly two values $\beta$ and $\beta^{\prime}$ in $[1-\sigma, 1+\sigma]$ such that

$$
\omega^{+}+\beta u_{t-2 ; 2}^{\mathrm{alg}}=\omega^{-}+\beta^{\prime} u_{t-2 ; 2}^{\mathrm{alg}}=x_{t-1 ; 2}^{\mathrm{alg}}
$$

Even conditioned on $x_{1: t-2}^{\mathrm{alg}}, \omega_{t-2}$ is uniformly distributed on $\left\{\omega^{-}, \omega^{+}\right\}$, and since $\beta$ and $\beta^{\prime}$ have the same probability mass under the uniform distribution of $\beta_{t-2}$, it follows that $\mathbb{P}\left[\omega_{t-2}=\omega^{+} \mid\right.$ $\left.x_{1: t-1}^{\mathrm{alg}}\right]=\left[\omega_{t-2}=\omega^{-} \mid x_{1: t-1}^{\mathrm{alg}}\right]$ when $\tilde{\mathcal{E}}_{t-1}$ holds.

Hence,

$$
\mathbb{E}\left[f\left(x_{t ; 1}^{\pi}\right) \mid \tilde{\mathcal{E}}_{t-1}, \mathcal{E}_{t-2}\right] \geq \min _{u \in \mathbb{R}} \frac{1}{2}(f(1+\alpha \sigma-u)+f(1-\alpha \sigma-u))
$$

For $f(z)=z^{2}$ or $f(z)=|z|$, the minimum is attained at $u=1$, with values $\alpha^{2} \sigma^{2}$ and $\alpha \sigma$, respectively, both equal to $f(\alpha \sigma)$. Thus,combining with Eq. (F.7), we conclude

$$
\mathbb{E}\left[f\left(x_{t ; 1}^{\pi}\right) \mid \mathcal{E}_{t-2}\right] \geq \frac{1}{2} f(\alpha \sigma)
$$

## F. 3 Lower Bound without Stability

Theorem F.3. Consider a scalar LTV system with $A=\rho \in[0,1]$, and $B_{t}$ drawn independently and uniformly at random from $\{-1\}$. Suppose that $w_{1}=1$, and $w_{t}=0$ for all $t>1$. Finally, let $c_{t}(x, u)=x^{2}$ be a fixed costs. Then,
(a) There then DRC policy $u_{t}^{\pi}=-\rho B_{2} x_{t}^{\text {nat }}$, DAC policy $u_{t}^{\pi}=-\rho B_{2} w_{t-1}$, and static-feedback policy $u_{t}^{\pi}=-\rho B_{2}$ A (all chosen with foreknowledge of the $\left(B_{t}\right)_{t \geq 1}$ sequence) all enjoy:

$$
J_{T}(\pi)=1, \quad \text { with probability } 1
$$

(b) Any online learning algorithm without foreknowledge of $\left(B_{t}\right)_{t \geq 1}$ must suffer expected cost

$$
\mathbb{E}\left[J_{T}(\pi)\right]=\Omega\left(\min \left\{T, \frac{1}{1-\rho}\right\}\right) .
$$

Proof. In part (a), all policies choose $u_{2}=-B_{2} \rho w_{1}$, so that $x_{3}=\rho x_{2}-\rho w_{1}=0$. Since 0 is an equilibrium point and $w_{t}=0$ for all $t \geq 2$, the system remains at zero. Hence, the only cost incurred is at times 1 and 2, which are costs of zero and 1 respectively. In part (b), we use the unbiasedness of $B_{t}$ to recurse

$$
\begin{aligned}
\mathbb{E}\left[x_{2}\right] & =A x_{1}+\mathbb{E}\left[B_{1}\right] \mathbb{E}\left[u_{1}\right]+w_{1}=1 \\
\mathbb{E}\left[x_{t+1}\right] & =A \mathbb{E}\left[x_{t}\right]+\mathbb{E}\left[B_{t}\right] \mathbb{E}\left[u_{t}\right]+\underbrace{w_{t}}_{=0}=\rho \mathbb{E}\left[x_{t}\right], \quad \forall t \geq 2,
\end{aligned}
$$

yielding $\mathbb{E}\left[x_{t}\right]=\rho^{t-2}$ for all $t \geq 2$. Hence, the expected cost of the policy is

$$
\mathbb{E}\left[J_{T}(\pi)\right]=\sum_{t=2}^{T} \rho^{(2 t-2)}
$$

Considering the cases where $\rho \geq \frac{1}{T-2}$ and $\rho \leq \frac{1}{T-2}$, we find the above is $\Omega\left(\min \left\{T, \frac{1}{1-\rho}\right\}\right)$.

## F. 4 Hardness of Computing Best State Feedback Controller

Consider a time-varying linear dynamical system with no noise:

$$
x_{t+1}=A_{t} x_{t}+B_{t} u_{t}+w_{t}, \quad \forall t, w_{t} \equiv 0
$$

subject to changing convex costs $c_{t}(x, u)$. We show that even in the no-noise setting properly learning the optimal state feedback policy is computationally hard. This statement holds with the control agent having full prior knowledge of the dynamics $\left(A_{t}, B_{t}, c_{t}\right)$. It relies on a reduction to the MAX-3SAT problem which is NP-hard. Our lower bound is inspired by the analogous one for discrete MDPs by [14].
Theorem F.4. There exists a reduction from MAX-3SAT on m-clauses and $n$-literals to the problem of finding the state feedback policy $K$ optimal for the $\operatorname{cost} \sum_{t=1}^{T} c_{t}\left(x_{t}^{K}, u_{t}^{K}\right)$, over sequentially stable dynamics given by $\left(A_{t}, B_{t}, c_{t}\right)$ : a solution to MAX-3SAT with $k$ value implies optimal cost of at most $-k$, and a solution $K$ to the control problem with $-k-\epsilon$ value implies optimal value of MAX-3SAT at least $k$ for any known $\epsilon>0$.

Let us first describe the construction of the dynamics that reduce the optimal control problem to the MAX-3SAT problem. Consider a 3-CNF formula $\phi$ with $m$ clauses $C_{1}, \ldots, C_{m}$ and $n$ literals $y_{1}, \ldots, y_{n}$. The state space is of dimensionality $d_{x}=n+1$ and the action space is of dimensionality $d_{u}=2$. The control problem is given as a sequence of $m$ episodes corresponding to the clauses of the formula $\phi$.
For a single clause $C_{j}$ with $j \in[m]$, let the dynamics $\left(A_{t}, B_{t}, c_{t}\right)_{t=1}^{n+2}$ be an episode of length $n+2$ constructed as follows. The initial state is $x_{1}=\left[1, \mathbf{0}_{n}\right]^{\top}$. The state transitions are independent of the clause itself given by the following $A_{t} \in \mathbb{R}^{n+1 \times n+1}$ :

- for $1 \leq t<n, A_{t}(t+1)=\left[\mathbf{1}_{n}, 0\right]^{\top}, A_{t}(n+1)=\left[\mathbf{0}_{n}, 1\right]^{\top}, A_{t}(i)=\left[\mathbf{0}_{n+1}\right]^{\top}$ for all other $i \neq t+1, n+1$.
- for $t=n$, it becomes $A_{t}(n+1)=\left[\mathbf{1}_{n+1}\right]^{\top}$ and $A_{t}(i)=\left[\mathbf{0}_{n+1}\right]^{\top}$ for all other $i \neq n+1$.
- for $t=n+1$ take $A_{t}(1)=\left[\mathbf{1}_{n+1}\right]^{\top}$ and $A_{t}(i)=\left[\mathbf{0}_{n+1}\right]^{\top}$ for all other $i \neq 1$; for $t=n+2$ take $A_{t}=\mathbf{0}_{n+1 \times n+1}$ to ensure sequential stability.

The action matrices $B_{t}$ along with the costs $c_{t}$, on the other hand, depend on the content of the clause $C_{j}$ itself. In particular, let $I_{j}$ be the set of indices of the literals that are in clause $C_{j}$. We define the regularity cost to be $c(x, u)=S_{x}(x)+(1-x(n+1))^{2} \cdot S_{u}(u)$, where $S_{x}(\cdot)=\operatorname{dist}\left(\cdot, \Delta_{n+1}\right)$ and $S_{u}(\cdot)=\operatorname{dist}\left(\cdot, \Delta_{2}\right)$ are the distance functions to the simplex sets of corresponding dimensionality. The action matrices and costs are given as follows:

- for $1 \leq t \leq n$ and $t \notin I_{j}, B_{t}=\mathbf{0}_{n+1 \times 2}$ and $c_{t}(x, u)=c(x, u)$.
- for $1 \leq t \leq n$ and $t \in I_{j}$, if $y_{t} \in C_{j}$ : then $B_{t}(t+1)=[-1,0]^{\top}, B_{t}(n+1)=[1,0]^{\top}$ and $B_{t}(i)=[0,0]^{\top}$ for all the other $i \neq t+1, n+1$; the cost is $c_{t}(x, u)=c(x, u)-u(1)$. $(1-x(n+1))$ rewarding the action $[1,0]$ which corresponds to assigning the literal a value $y_{t}=1$.
- for $1 \leq t \leq n$ and $t \in I_{j}$, if $\neg y_{t} \in C_{j}$ : then $B_{t}(t+1)=[0,-1]^{\top}, B_{t}(n+1)=[0,1]^{\top}$ and $B_{t}(i)=[0,0]^{\top}$ for all the other $i \neq t+1, n+1$; the cost is $c_{t}(x, u)=c(x, u)-u(2)$. $(1-x(n+1))$ rewarding the action $[0,1]$ which corresponds to assigning the literal a value $y_{t}=0$.
- for $t=n+1, B_{t}=\mathbf{0}_{n+1 \times 2}$ and for $t=n+2, B_{t}(1)=[1,1]^{\top}$ and $B_{t}(i)=[0,0]^{\top}$ for all other $i \neq 1$; for both $t=n+1, n+2$, the costs are $c_{t}(x, u)=c(x, u)$.

Note that the last two rounds $t=n+1, n+2$ for a clause ensure sequential stability and identical starting state $x_{1}=\left[1, \mathbf{0}_{n}\right]^{\top}$.
Lemma F.6. The described system $\left(A_{t}, B_{t}\right)$ is sequentially stable and the costs $c_{t}$ are convex in $x, u$.
Proof. By the construction of the state matrices $A_{t}$, we know that for any $t \geq n+2$ the operator $\Phi_{t}^{[n+2]}=0$ implying sequential stability of the system. To show convexity, note that the distance function $g(\cdot)=\operatorname{dist}(\cdot, \mathcal{S})$ for any convex and compact set $\mathcal{S}$ is a convex function. More specifically, for $z \in \mathbb{R}^{d_{z}}$ it is given by $g(z)=\min _{w \in \mathcal{S}}\|z-w\|$. This is straightforward to show: take any $z_{1}, z_{2} \in \mathbb{R}^{d_{z}}$, and let $w_{1}, w_{2} \in \mathcal{S}$ be the closest points to $z_{1}, z_{2}$ respectively, i.e. $\left\|w_{1}-z_{1}\right\|=g\left(z_{1}\right)$ and $\left\|w_{2}-z_{2}\right\|=g\left(z_{2}\right)$. For any $\lambda \in[0,1]$, given the convexity of the set $\mathcal{S}$, we know that $\lambda w_{1}+(1-\lambda) w_{2} \in \mathcal{S}$, which concludes the convexity proof for $g$ :

$$
\begin{aligned}
\lambda g\left(z_{1}\right)+(1-\lambda) g\left(z_{2}\right) & =\lambda\left\|z_{1}-w_{1}\right\|+(1-\lambda)\left\|z_{2}-w_{2}\right\| \\
& \geq\left\|\lambda\left(z_{1}-w_{1}\right)+(1-\lambda)\left(z_{2}-w_{2}\right)\right\| \\
& =\left\|\lambda z_{1}+(1-\lambda) z_{2}-\left(\lambda w_{1}+(1-\lambda) w_{2}\right)\right\| \\
& \geq g\left(\lambda z_{1}+(1-\lambda) z_{2}\right) .
\end{aligned}
$$

This means that both $S_{x}(\cdot)$ and $S_{u}(\cdot)$ are convex functions since the simplex is convex in any dimension. The construction of the costs $c_{t}$ is based on these two functions as well as linear components in $x, u$, hence all $c_{t}$ costs are convex in $x, u$.

Lemma F.7. If there exists an assignment of literals $y \leftarrow v$ s.t. the formula $\phi$ has $k \in[1, m]$ satisfied clauses, then there is a corresponding linear policy $K \in \mathbb{R}^{2 \times n+1}$ that suffers the exact cost of $-k$.

Proof. This should be evident from the construction itself. Let $v_{i}=1$ assignment correspond to $\bar{v}_{i}=[1,0]^{\top}$ and $v_{i}=0$ to $\bar{v}_{i}=[0,1]^{\top}$. Then consider the linear policy $K=\left[\bar{v}_{1}, \ldots, \bar{v}_{n}, \mathbf{0}_{2}^{\top}\right]$. Denote $e_{1}, \ldots, e_{n+1} \in \mathbb{R}^{n+1}$ to be the basis vectors of the space. Note that according to the defined $K$, if $x_{t}$ is a basis vector of $\mathbb{R}^{n+1}$, then $u_{t}=K x_{t}$ is a basis vector of $\mathbb{R}^{2}$. It is straightforward to check by our construction that $u_{t}$ being a basis vector of $\mathbb{R}^{2}$ implies $x_{t+1}$ is a basis vector of $\mathbb{R}^{n+1}$. Since $x_{1}=e_{1}$, then the state-action pairs when following policy $K$ are both basis vectors, and satisfy the regularity conditions of $c(x, u)$. This means that the policy $K$ plays $\bar{v}_{t}$ if $x_{t}(t)=1$ and plays $\mathbf{0}_{2}$ if $x_{t}(n+1)=1$. Hence, if the clause $C_{j}$ is satisfied by the assignment $y \leftarrow v$, then the cost of $K$ over the episode is exactly -1 , i.e. once the clause is satisfied, -1 is accrued and the state moves to the sink $e_{n+1}$. If the clause is not satisfied, then the cost is 0 since $c(x, u)$ is 0 throughout. This means that the constructed linear policy $K$ over the whole control sequence suffers cost $-k . \mathrm{ik}$
Lemma F.8. If there exists a state feedback policy $K \in \mathbb{R}^{2 \times n+1}$ s.t. following the actions $u_{t}=K x_{t}$ results in cost at most $-k-\epsilon$ for any $k \in[1, m]$ and any $\epsilon \in(0,1)$, then there is a literal assignment $y \leftarrow v$ s.t. the formula $\phi$ has at least $k$ satisfied clauses.

Proof. Let the linear policy matrix be given as $K=\left[\bar{v}_{1}, \ldots, \bar{v}_{n}, \bar{v}_{n+1}\right]$. The proof consists of two main components: (i) we argue that the policy $K^{*}$ with $\bar{v}_{i}^{*}=\arg \min _{v \in \Delta_{2}}\left\|v-\bar{v}_{i}\right\|$ for $1 \leq i \leq n$ and $\bar{v}_{n+1}^{*}=\mathbf{0}_{2}$ is at least as good as $K$ in terms of cost up to an approximation factor $\epsilon$; (ii) we show that for $K$ that satisfies the constraints, the randomized policy $\hat{K}$ that has $\hat{K}(i)=[1,0]^{\top}$ w.p. $\bar{v}_{i}(1)$ and $\hat{K}(i)=[0,1]^{\top}$ w.p. $\bar{v}_{i}(2)$ (as well as $\hat{K}(n+1)=\mathbf{0}_{2}$ ) suffers expected cost at most that of $K$ itself.

Suppose these two claims are true, then the described randomized linear policy $\hat{K}$ has expected cost at most $-k$, which means that there exists a deterministic linear policy with first $n$ columns as basis vectors, i.e. $[1,0]^{\top}$ or $[0,1]^{\top}$, that suffers cost at most $-k$. It follows that the corresponding assignment of literals given by the first $n$ columns of the linear policy $y \leftarrow v$ satisfies at least $k$ out of the $m$ clauses, so $\phi$ has at least $k$ satisfied clauses.

To prove (i), first notice that the policy $K^{*}$ suffers non-positive cost over the entire horizon since it satisfies the necessary constraints given by the regulatory cost $c(x, u)$. Note also that said $c(x, u)$ can be scaled by any constant $M_{\epsilon}>0$. Now suppose that under the condition $\min _{i}\left\|\bar{v}_{i}-\bar{v}_{i}^{*}\right\| \leq \ell_{\epsilon}$ the cost difference of $K^{*}$ and $K$ is bounded by $\epsilon$ : the choice of $\ell_{\epsilon}$ can depend on any problem parameters, and since the construction is over $T=\Theta(m n)$ overall rounds (finite), such a choice is always possible. Hence, for all such $K^{*}$ we automatically infer that it suffers cost at most $-k$ and satisfies the necessary constraints. On the other hand, if the condition does not hold, then the distance of $\bar{v}_{i}$ from $\Delta_{2}$ is bounded from below by $\ell_{\epsilon}$, meaning that for a sufficiently large choice of $M_{\epsilon}>0$ (given knowledge of $\epsilon$ and all other parameters), the overall cost suffered by $K$ will be positive due to $c(x, u)$, i.e. it will have a higher cost than $K^{*}$. Therefore, any state feedback policy $K$ can be approximately replaced by $K^{*}$ that satisfies the constraints ensured by $c(x, u)$.

To show (ii), for a policy $K$ that does satisfy these constraints, i.e. $\bar{v}_{i} \in \Delta_{2}$ for $1 \leq i \leq n$ and $\bar{v}_{n+1}=\mathbf{0}_{2}$, we show that its randomized version $\hat{K}$ is at least just as good in terms of expected cost. Proving this claim is a matter of unrolling the dynamics for a single clause $C_{j}$. The order, the indices and negation or not of the literals in $C_{j}$ does not affect the cost, so w.l.o.g. assume we have $C_{1}=y_{1} \vee y_{2} \vee y_{3}$. The cost of a general policy $K$ over first 3 iterations is given by

$$
-\bar{v}_{1}(1)-\left(1-\bar{v}_{1}(1)\right)^{2} \cdot \bar{v}_{2}(1)-\left(1-\bar{v}_{1}(1)\right)^{2} \cdot\left(1-\bar{v}_{2}(1)\right)^{2} \cdot \bar{v}_{3}(1)
$$

The alternative randomized linear policy instead suffer expected cost given by

$$
-\bar{v}_{1}(1)-\left(1-\bar{v}_{1}(1)\right) \cdot \bar{v}_{2}(1)-\left(1-\bar{v}_{1}(1)\right) \cdot\left(1-\bar{v}_{2}(1)\right) \cdot \bar{v}_{3}(1)
$$

which is straightforward to show to be not larger than the original cost.
Proof of Theorem F.4. Lemma F. 6 above show that the given LTV system construction along with the costs satisfies the theorem conditions. Lemmas F. 7 and F. 8 indicate that the MAX-3Sat problem can be reduced to the LTV control, in particular proper learning of state feedback policies in this setting. Given that MaX-3Sat is NP-Hard even in its decision form implies the computational hardness of the offline optimization of LTV optimal state feedback policies.


[^0]:    ${ }^{1}$ An oblivious adversary chooses the matrices, costs and perturbations prior to the control trajectory.

[^1]:    ${ }^{2}$ For precise constants please see Eq. (D.2).

[^2]:    ${ }^{3}$ This is w.l.o.g. and only assumed for simplicity of presentation.

[^3]:    ${ }^{4}$ Equivalently, in $\Pi_{\mathrm{dac}}(1,1)$ since $A_{t} \equiv 0$

