Beyond Black-Box Advice: Learning-Augmented Algorithms for MDPs with Q-Value Predictions

Tongxin Li
School of Data Science
CUHK-SZ, China
litongxin@cuhk.edu.cn

Yiheng Lin
Computing + Mathematical Sciences
Caltech, USA
yihengl@caltech.edu

Shaolei Ren
Electrical & Computer Engineering
UC Riverside, USA
shaolei@ucr.edu

Adam Wierman
Computing + Mathematical Sciences
Caltech, USA
adamw@caltech.edu

Abstract

We study the tradeoff between consistency and robustness in the context of a single-trajectory time-varying Markov Decision Process (MDP) with untrusted machine-learned advice. Our work departs from the typical approach of treating advice as coming from black-box sources by instead considering a setting where additional information about how the advice is generated is available. We prove a first-of-its-kind consistency and robustness tradeoff given Q-value advice under a general MDP model that includes both continuous and discrete state/action spaces. Our results highlight that utilizing Q-value advice enables dynamic pursuit of the better of machine-learned advice and a robust baseline, thus result in near-optimal performance guarantees, which provably improves what can be obtained solely with black-box advice.

1 Introduction

Machine-learned predictions and hand-crafted algorithmic advice are both crucial in online decision-making problems, driving a growing interest in *learning-augmented algorithms* [1, 2] that exploit the benefits of predictions to improve the performance for typical problem instances while bounding the worst-case performance [3, 4]. To this point, the study of learning-augmented algorithms has primarily viewed machine-learned advice as potentially untrusted information generated by black-box models. Yet, in many real-world problems, additional knowledge of the machine learning models used to produce advice/predictions is often available and can potentially improve the performance of learning-augmented algorithms.

A notable example that motivates our work is the problem of minimizing costs (or maximizing rewards) in a single-trajectory Markov Decision Process (MDP). More concretely, a value-based machine-learned policy $\widetilde{\pi}$ can be queried to provide suggested actions as advice to the agent at each step [5–7]. Typically, the suggested actions are chosen to minimize (or maximize, in case of rewards) estimated cost-to-go functions (known as Q-value predictions) based on the current state.

Naturally, in addition to suggested actions, the Q-value function itself can also provide additional information (e.g., the long-term impact of choosing a certain action) potentially useful to the design of a learning-augmented algorithm. Thus, this leads to two different designs for learning-augmented algorithms in MDPs: *black-box* algorithms and *grey-box* algorithms. A learning-augmented algorithm using $\widetilde{\pi}$ is black-box if $\widetilde{\pi}$ provides only the suggested action $\widetilde{u}$ to the learning-augmented algorithm,

37th Conference on Neural Information Processing Systems (NeurIPS 2023).

whereas it is value-based (a.k.a., grey-box) if $\widetilde{\pi}$ provides an estimate of the Q-value function $\widetilde{Q}$ (that also implicitly includes a suggested action $\widetilde{u}$ obtained by minimizing $\widetilde{Q}$) to the learning-augmented algorithm.

Value-based policies $\widetilde{\pi}$ often perform well empirically in stationary environments in practice [5, 6]. However, they may not have performance guarantees in all environments and can perform poorly at times due to a variety of factors, such as non-stationary environments [8–11], policy collapse [12], sample inefficiency [13], and/or when training data is biased [14]. As a consequence, such policies often are referred to as "untrusted advice" in the literature on learning-augmented algorithms, where the notion of "untrusted" highlights the lack of performance guarantees. In contrast, recent studies in competitive online control [15–21] have begun to focus on worst-case analysis and provide control policies $\overline{\pi}$ with strong performance guarantees even in adversarial settings, referred to as *robustness*, i.e., $\overline{\pi}$ provides "trusted advice." Typically, the goal of a learning-augmented online algorithm [1, 3] is to perform nearly as well as the untrusted advice when the machine learned policy performs well, a.k.a., achieve *consistency*, while also ensuring worst-case robustness. Combining the advice of an untrusted machine-learned policy $\widetilde{\pi}$ and a robust policy $\overline{\pi}$ naturally leads to a tradeoff between consistency and robustness. In this paper, we explore this tradeoff in a time-varying MDP setting and seek to answer the following key question for learning-augmented online algorithms:

*Can Q-value advice from an untrusted machine-learned policy, $\widetilde{\pi}$, in a **grey-box** scenario provide more benefits than the **black-box** action advice generated by $\widetilde{\pi}$ in the context of **consistency and robustness tradeoffs** for MDPs?*

1.1 Contributions

We answer the question above in the affirmative by presenting and analyzing a unified projection-based learning-augmented online algorithm (PROjection Pursuit policy, simplified as PROP in Algorithm 1) that combines action feedback from a trusted, robust policy $\overline{\pi}$ with an untrusted ML policy $\widetilde{\pi}$. In addition to offering a consistency and robustness tradeoff for MDPs with black-box advice, our work moves beyond the black-box setting. Importantly, by considering the grey-box setting, the design of PROP demonstrates that the *structural information* of the untrusted machine-learned advice can be leveraged to determine the trust parameters dynamically, which would otherwise be challenging (if not impossible) in a black-box setting. To our best knowledge, PROP is the first-of-its-kind learning-augmented algorithm that applies to general MDP models, which allow continuous or discrete state and action spaces.

Our main results characterize the tradeoff between consistency and robustness for both black-box and grey-box settings in terms of the ratio of expectations, RoE, built upon the traditional consistency and robustness metrics in [3, 22, 23, 4] for the competitive ratio. We show in Theorem 5.2 that for the black-box setting, PROP is $(1 + \mathcal{O}((1-\lambda)\gamma))$-consistent and $(\mathsf{ROB} + \mathcal{O}(\lambda\gamma))$-robust where $0 \leq \lambda \leq 1$ is a hyper-parameter. Moreover, for the black-box setting, PROP cannot be both $(1 + o(\lambda\gamma))$-consistent and $(\mathsf{ROB} + o((1-\lambda)\gamma))$-robust for any $0 \leq \lambda \leq 1$ where γ is the diameter of the action space. In sharp contrast, by using a careful design of a robustness budget parameter in PROP with Q-value advice (grey-box setting), PROP is 1-consistent and $(\mathsf{ROB} + o(1))$-robust.

Our result highlights the benefits of exploiting the additional information informed by the estimated Q-value functions, showing that the ratio of expectations can approach the better of the two policies $\widetilde{\pi}$ and $\overline{\pi}$ for any single-trajectory time-varying, and even possibly adversarial environments — if the value-based policy $\widetilde{\pi}$ is near-optimal, then the worst-case $\mathsf{RoE}(\mathrm{PROP})$ can approach 1 as governed by a consistency parameter; otherwise, $\mathsf{RoE}(\mathrm{PROP})$ can be bounded by the ratio of expectations of $\overline{\pi}$ subject to an additive term $o(1)$ that decreases when the time horizon T increases.

A key technical contribution of our work is to provide the first quantitative characterization of the consistency and robustness tradeoff for a learning-augmented algorithm (PROP) in a general MDP model, under both standard black-box and novel grey-box settings. Importantly, PROP is able to leverage a broad class of robust policies, called *Wasserstein robust* policies, which generalize the well-known contraction principles that are satisfied by various robust policies [24] and have been used to derive regrets for online control [19, 25]. A few concrete examples of Wasserstein robust policies applicable for PROP are provided in Table 1(Section 3.1).

1.2 Related Work

Learning-Augmented Algorithms with Black-Box Advice. The concept of integrating black-box machine-learned guidance into online algorithms was initially introduced by [26]. [3] coined terms "robustness" and "consistency" with formal mathematical definitions based on the competitive ratio. Over the past few years, the consistency and robustness approach has gained widespread popularity and has been utilized to design online algorithms with black-box advice for various applications, including ski rental [3, 22, 23], caching [27–29], bipartite matching [30], online covering [31, 32], convex body chasing [4], nonlinear quadratic control [33]. The prior studies on learning-enhanced algorithms have mainly focused on creating meta-strategies that combine online algorithms with black-box predictions, and typically require manual setting of a trust hyper-parameter to balance consistency and robustness. A more recent learning-augmented algorithm in [33] investigated the balance between competitiveness and stability in nonlinear control in a black-box setting. However, this work limits the robust policy to a linear quadratic regulator and does not provide a theoretical basis for the selection of the trust parameters. [34] generalized the black-box advice setting by considering distributional advice.

Online Control and Optimization with Structural Information. Despite the lack of a systematic analysis, recent studies have explored the usage of structural information in online control and optimization problems. Closest to our work, [7] considered a related setting where the Q-value function is available as advice, and shows that such information can be utilized to reduce regret in a tabular MDP model. In contrast, our analysis applies to more general models that allow continuous state/action spaces. In [17], the dynamical model and the predictions of disturbances in a linear control system are shown to be useful in achieving a near-optimal consistency and robustness tradeoff. The predictive optimization problem solved by MPC [35, 36, 16, 37] can be regarded as a special realization of grey-box advice, where an approximated cost-to-go function is constructed from structural information that includes the (predicted) dynamical model, costs, and disturbances.

MDP with External Feedback. Feedback from external sources such as control baselines [38, 39], visual explanations [40], and human experts [41–43] is often available in MDP. This external feedback can be beneficial for various purposes, such as ensuring safety [44], reducing variance [38], training human-like chatbots [41], and enhancing overall trustworthiness [45], among others. The use of control priors has been proposed by [38] as a way to guarantee the Lyapunov stability of the training process in reinforcement learning. They used the Temporal-Difference method to tune a coefficient that combines a RL policy and a control prior, but without providing a theoretical foundation. Another related area is transfer learning in RL, where external Q-value advice from previous tasks can be adapted and utilized in new tasks. Previous research has shown that this approach can outperform an agnostic initialization of Q, but these results are solely based on empirical observations and lack theoretical support [46–48].

2 Problem Setting

We consider a finite-horizon, single-trajectory, time-varying MDP with T discrete time steps. The state space $\mathcal{X}$ is a subset of a normed vector space embedded with a norm $\|\cdot\|_{\mathcal{X}}$. The actions are chosen from a convex and compact set $\mathcal{U}$ in a normed vector space characterized by some norm $\|\cdot\|_{\mathcal{U}}$. Notably, $\mathcal{U}$ can represent either continuous actions or the probability distributions used when choosing actions from a finite set.[1] The diameter of the action space $\mathcal{U}$ is denoted by $\gamma := \max_{u\in\mathcal{U}} \|u\|_{\mathcal{U}}$. Denote $[T] \coloneqq \{0, \ldots, T-1\}$. For each time step $t \in [T]$, let $P_t : \mathcal{X} \times \mathcal{U} \to \mathcal{P}_{\mathcal{X}}$ be the transition probability, where $\mathcal{P}_{\mathcal{X}}$ is a set of probability measures on $\mathcal{X}$. We consider time-varying costs $c_t : \mathcal{X} \times \mathcal{U} \to \mathbb{R}_+$, while rewards can be treated similarly by adding a negative sign. An initial state $x_0 \in \mathcal{X}$ is fixed. This MDP model is compactly represented by $\mathsf{MDP}(\mathcal{X}, \mathcal{U}, T, P, c)$.

The goal of a policy in this MDP setting is to minimize the total cost over all T steps. The policy agent has no access to the full MDP. At each time step $t \in [T]$, only the incurred cost value $c_t(x_t, u_t)$ and the next state $x_{t+1} \sim P_t(\cdot|x_t, u_t)$ are revealed to the agent after playing an action $u_t \in \mathcal{U}$. We denote a policy by $\pi = (\pi_t : t \in [T])$ where each $\pi_t : \mathcal{X} \to \mathcal{U}$ chooses an action u_t when observing

[1] The action space $\mathcal{U}$ is assumed to be a continuous, convex, and compact set for more generality. When the actions are discrete, $\mathcal{U}$ can be defined as the set of all probability distributions on a finite action space. We relegate the detailed discussions in Appendix A.2 and F.

x_t at step $t \in [T]$. Note that our results can be generalized to the setting when π_t is stochastic and outputs a probability distribution on $\mathcal{U}$. Given $\mathsf{MDP}(\mathcal{X}, \mathcal{U}, T, P, c)$, we consider an optimization with time-varying costs and transition dynamics. Thus, our goal is to find a policy π that minimizes the following expected total cost:

$$J(\pi) \coloneqq \mathbb{E}_{P,\pi}\Big[\sum_{t\in[T]} c_t\left(x_t, \pi_t(x_t)\right)\Big] \tag{1}$$

where the randomness in $\mathbb{E}_{P,\pi}$ is from the transition dynamics $P = (P_t : t \in [T])$ and the policy $\pi = (\pi_t : t \in [T])$. We focus our analysis on the expected dynamic regret and the ratio of expectations, defined below, as the performance metrics for our policy design.

Definition 1 (Expected dynamic regret)**.** *Given* $\mathsf{MDP}(\mathcal{X}, \mathcal{U}, T, P, c)$*, the (expected) dynamic regret of a policy* $\pi = (\pi_t : t \in [T])$ *is defined as the difference between the expected cost induced by the policy* π*,* $J(\pi)$ *in* (1)*, and the optimal expected cost* $J^\star \coloneqq \inf_\pi J(\pi)$*, i.e.,* $\mathsf{DR}(\pi) \coloneqq J(\pi) - J^\star$*.*

Dynamic regret is a more general (and often more challenging to analyze) measure than classical static regret, which has been mostly used for stationary environments [49, 50]. The following definition of the ratio of expectations [51, 52] will be used as an alternative performance metric in our main results.

Definition 2 (Ratio of expectations)**.** *Given* $\mathsf{MDP}(\mathcal{X}, \mathcal{U}, T, P, c)$*, the ratio of expectations of a policy* $\pi = (\pi_t : t \in [T])$ *is defined as* $\mathsf{RoE}(\pi) \coloneqq J(\pi)/J^\star$ *where* $J(\pi)$ *and* $J^\star$ *are the same as in Definition 1.*

Dynamic regret and the ratio of expectations defined above also depend on the error of the untrusted ML advice; we make this more explicit in Section 3.2. Next, we state the following continuity assumption, which is standard in MDPs with continuous action and state spaces [53–55]. Note that our analysis can be readily adapted to general Hölder continuous costs with minimal modifications.

Assumption 1 (Lipschitz costs)**.** *For any time step* $t \in [T]$*, the cost function* $c_t : \mathcal{X} \times \mathcal{U} \to \mathbb{R}_+$ *is Lipschitz continuous with a Lipschitz constant* $L_{\mathrm{C}} < \infty$*, i.e., for any* $t \in [T]$*,* $|c_t(x, u) - c_t(x', u')| \leq L_C\left(\|x - x'\|_{\mathcal{X}} + \|u - u'\|_{\mathcal{U}}\right)$*. Moreover,* $0 < c_t(x, u) < \infty$ *for all* $t \in [T]$*,* $x \in \mathcal{X}$*, and* $u \in \mathcal{U}$*.*

3 Consistency and Robustness in MDPs

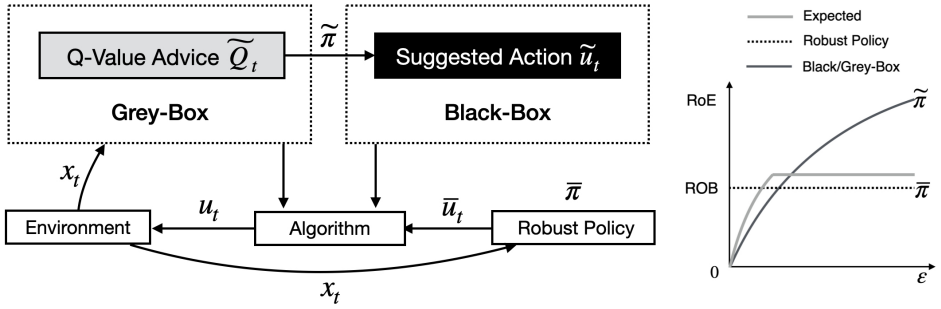

Figure 1: *Left*: Overview of settings in our problem. *Right*: consistency and robustness tradeoff, with RoE and ε defined in Definition 2 and Equation (4).

Our objective is to achieve a balance between the worst-case guarantees on cost minimization in terms of dynamic regret provided by a robust policy, $\overline{\pi}$, and the average-case performance of a valued-based policy, $\widetilde{\pi}$, in the context of $\mathsf{MDP}(\mathcal{X}, \mathcal{U}, T, P, c)$. In particular, we denote by $\mathsf{ROB} \geq 1$ a ratio of expectation bound of the robust policy $\overline{\pi}$ such that the worst case $\mathsf{RoE}(\overline{\pi}) \leq \mathsf{ROB}$. In the learning-augmented algorithms literature, these two goals are referred to as consistency and robustness [3, 1]. Informally, robustness refers to the goal of ensuring worst-case guarantees on cost minimization comparable to those provided by $\overline{\pi}$ and consistency refers to ensuring performance nearly as good as $\widetilde{\pi}$ when $\widetilde{\pi}$ performs well (e.g., when the instance is not adversarial). Learning-augmented algorithms seek to achieve consistency and robustness by combining $\overline{\pi}$ and $\widetilde{\pi}$, as illustrated in Figure 1.

Table 1: Examples of models covered in this paper and the associated control baselines. For the right column, bounds on the ratio of expectations RoE are exemplified, where ROB is defined in Section 3 and $\mathcal{O}$ omits inessential constants.

Model	**Robust Baseline $\overline{\pi}$**	RoE
Time-varying MDP (Our General Model)	Wasserstein Robust Policy (Definition 3)	ROB
Discrete MDP (Appendix A.2)	Any Policy that Induced a Regular Markov Chain	—
Time-Varying LQR (Appendix A.1)	MPC with Robust Predictions (Algorithm 3)	$\mathcal{O}(1)$

Our focus in this work is to design robust and consistent algorithms for two types of advice: black-box advice and grey-box advice. The type of advice that is nearly always the focus in the learning-augmented algorithm literature is black-box advice — only providing a suggested action $\widetilde{u}_t$ without additional information. In contrast, on top of the action $\widetilde{u}_t$, grey-box advice can also reveal the internal state of the learning algorithm, e.g., the Q-value $\widetilde{Q}_t$ in our setting. This contrast is illustrated in Figure 1.

Compared to black-box advice, grey-box advice has received much less attention in the literature, despite its potential to improve tradeoffs between consistency and robustness as recently shown in [34, 17]. Nonetheless, the extra information on top of the suggested action in a grey-box setting potentially allows the learning-augmented algorithm to make a better-informed decision based on the advice, thus achieving a better tradeoff between consistency and robustness than otherwise possible.

In the remainder of this section, we discuss the robustness properties for the algorithms we consider in our learning-augmented framework (Section 3.1), and introduce the notions of consistency in our grey-box and black-box models in Section 3.2.

3.1 Locally Wasserstein-Robust Policies

We begin with constructing a novel notion of robustness for our learning-augmented framework based on the Wasserstein distance as follows. Denote the robust policy by $\overline{\pi} \coloneqq (\overline{\pi}_t : t \in [T])$, where each $\overline{\pi}_t$ maps a system state to a deterministic action (or a probability of actions in the stochastic setting). Denote by $\rho_{t_1:t_2}(\rho)$ the joint distribution of the state-action pair $(x_t, u_t) \in \mathcal{X} \times \mathcal{U}$ at time $t_2 \in [T]$ when implementing the baselines $\overline{\pi}_{t_1}, \ldots, \overline{\pi}_{t_2}$ consecutively with an initial state-action distribution ρ. We use $\|\cdot\|_{\mathcal{X}\times\mathcal{U}} \coloneqq \|\cdot\|_{\mathcal{X}} + \|\cdot\|_{\mathcal{U}}$ as the included norm for the product space $\mathcal{X} \times \mathcal{U}$. Let $W_p(\mu, \nu)$ denote the Wasserstein p-distance between distributions μ and ν whose support set is $\mathcal{X} \times \mathcal{U}$:

$$W_p(\mu,\nu) \coloneqq \left(\inf_{J\in\mathcal{J}(\mu,\nu)} \int \|(x,u)-(x',u')\|^p_{\mathcal{X}\times\mathcal{U}} \mathrm{d}J\left((x,u),(x',u')\right)\right)^{1/p}$$

where $p \in [1, \infty)$ and $\mathcal{J}(\mu, \nu)$ denotes a set of all joint distributions J with a support set $\mathcal{X} \times \mathcal{U}$ that have marginals μ and ν. Next, we define a robustness condition for our learning-augmented framework.

Definition 3 (r-locally p-Wasserstein robustness)**.** *A policy $\overline{\pi} = (\pi_t : t \in [T])$ is r-**locally p-Wasserstein-robust** if for any $0 \le t_1 \le t_2 < T$ and any pair of state-action distributions ρ, ρ' where the the p-Wasserstein distance between them is bounded by $W_p(\rho, \rho') \le r$, for some radius $r > 0$, the following inequality holds:*

$$W_p\left(\rho_{t_1:t_2}(\rho), \rho_{t_1:t_2}(\rho')\right) \le s(t_2 - t_1) W_p\left(\rho, \rho'\right) \tag{2}$$

for some function $s : [T] \to \mathbb{R}_+$ satisfying $\sum_{t\in[T]} s(t) \le C_s$ where $C_s > 0$ is a constant.

Our robustness definition is naturally more relaxed than the usual contraction property in the control/optimization literature [25, 35] — if any two different state-action distributions converge exponentially with respect to the Wasserstein p-distance, then a policy $\overline{\pi}$ is *r-locally p-Wasserstein-robust*. This is illustrated in Figure 2. Note that, although the Wasserstein robustness in Definition 3 well captures a variety of distributional robustness metrics such as the total variation robustness defined on finite state/action spaces, it can also be further generalized to other metrics for probability distributions.

As shown in Appendix A (provided in the supplementary material), by establishing a connection between the Wasserstein distance and the total variation metric, any policy that induces a regular

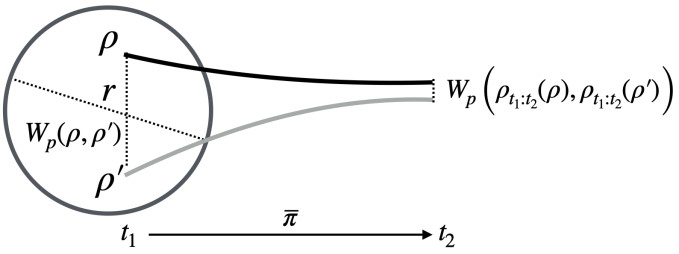

Figure 2: An illustration of an *r-locally p-Wasserstein-robust* policy.

Markov chain satisfies the fast mixing property and the state-action distribution will converge with respect to the total variation distance to a stationary distribution [56]. A more detailed discussion can be found in Appendix A.2. Moreover, the Wasserstein-robustmess in Definition 3 includes a set of contraction properties in control theory as special cases. For example, for a locally Wasserstein-robust policy, if the transition kernel P and the baseline policy $\overline{\pi}$ are deterministic, then the state-action distributions become point masses, reducing Definition 3 to a state-action perturbation bound in terms of the ℓ_2-norm when implementing the policy $\overline{\pi}$ from different starting states [35, 19].

The connections discussed above highlight the existence of several well-known robust policies that satisfy Definition 3. Besides the case of discrete MDPs discussed in Appendix A.2, another prominent example is model predictive control (MPC), for which robustness follows from the results in [19] (see Appendix A.1 for details). The model assumption below will be useful in our main results.

Assumption 2. *There exists a γ-locally p-Wasserstein-robust baseline control policy (Definition 3) $\overline{\pi}$ for some $p \geq 1$, where γ is the diameter of the action space $\mathcal{U}$.*

3.2 Consistency and Robustness for RoE

In parallel with the notation of "consistency and robustness" in the existing literature on learning-augmented algorithms [3, 1], we define a new metric of consistency and robustness in terms of RoE. To do so, we first introduce an optimal policy $\pi^\star$. Based on $\mathsf{MDP}(\mathcal{X}, \mathcal{U}, T, P, c)$, let $\pi_t^\star = (\pi_t^\star : t \in [T])$ denote the optimal policy at each time step $t \in [T]$, whose optimal Q-value function is

$$Q_t^\star(x, u) \coloneqq \inf_\pi \mathbb{E}_{P,\pi}\left[\sum_{\tau=t}^{T-1} c_\tau\left(x_\tau, u_\tau\right) \middle| x_t = x, u_t = u\right],$$

where $\mathbb{E}_{P,\pi}$ denotes an expectation with respect to the randomness of the trajectory $\{(x_t, u_t) : t \in [T]\}$ obtained by following a policy π and the transition probability P at each step $t \in [T]$. The Bellman optimality equations can then be expressed as

$$Q_t^\star(x, u) = \left(c_t + \mathbb{P}_t V_{t+1}^\star\right)(x, u), \qquad V_t^\star(x) = \inf_{v \in \mathcal{U}} Q_t^\star(x, v), \qquad V_t^\star(x) = 0 \tag{3}$$

for all $(x, u) \in \mathcal{X} \times \mathcal{U}$, $t \in [T]$ and $t \in [T]$, where we write $\left(\mathbb{P}_t V^\star\right)(x, u) \coloneqq \mathbb{E}_{x' \sim P_t(\cdot|x,u)}\left[V^\star(x')\right]$. This indicates that for each time step $t \in [T]$, $\pi_t^\star$ is the greedy policy with respect to its optimal Q-value functions $(Q_t^\star : t \in [T])$. Note that for any $t \in [T]$, $Q_t^\star(x, u) = 0$. Given this setup, the value-based policies $\widetilde{\pi} \coloneqq (\widetilde{\pi}_t : t \in [T])$ take the following form. For any $t \in [T]$, a value-based policy $\widetilde{\pi}_t : \mathcal{X} \to \mathcal{U}$ produces an action $\widetilde{u}_t \in \arg\min_{v \in \mathcal{U}} \widetilde{Q}_t\left(x_t, v\right)$ by minimizing an estimate of the optimal Q-value function $\widetilde{Q}_t$.

We make the following assumption on the machine-learned untrusted policy $\widetilde{\pi}$ and the Q-value advice.

Assumption 3. *The machine-learned untrusted policy $\widetilde{\pi}$ is value-based. The Q-value advice $\widetilde{Q}_t : \mathcal{X} \times \mathcal{U} \to \mathbb{R}$ is Lipschitz continuous with respect to $u \in \mathcal{U}$ for any $x \in \mathcal{X}$, with a Lipschitz constant L_Q for all $t \in [T]$. Moreover, $\widetilde{Q}_t\left(x, u\right) - Q_t^\star\left(x, u\right) = o(T)$ for all $(x, u) \in \mathcal{X} \times \mathcal{U}$ and $t \in [T]$.*

We can now define a consistency measure for Q-value advice $\widetilde{Q}_t$, which measures the error of the estimates of the Q-value functions due to approximation error and time-varying environments, etc. Let $p \in (0, \infty]$. Fix a sequence of distributions $\rho = (\rho_t : t \in [T])$ whose support set is $\mathcal{X} \times \mathcal{U}$ and let ϕ_t be the marginal distribution of ρ_t on $\mathcal{X}$. We define a quantity representing the error of the Q-value

advice

$$\varepsilon(p,\rho) \coloneqq \sum_{t\in[T]} \left(\left\| \widetilde{Q}_t - Q_t^\star \right\|_{p,\rho_t} + \left\| \inf_{v\in\mathcal{U}} \widetilde{Q}_t - \inf_{v\in\mathcal{U}} Q_t^\star \right\|_{p,\phi_t} \right) \tag{4}$$

where $\|\cdot\|_{p,\rho} \coloneqq \left(\int |\cdot|^p \,\mathrm{d}\rho\right)^{1/p}$ denotes the $L_{p,\rho}$-norm. A policy with Q-value functions $\{Q_t : t \in [T]\}$ is said to be (ε, p, ρ)-*consistent* if there exists an ε satisfying (4). In addition, a policy is $(0,\infty)$-consistent if $\widetilde{Q}_t$ is a Lebesgue-measurable function for all $t \in [T]$ and (∞,ε)-consistent if the L_∞-norm satisfies $\sum_{t\in[T]} \|\widetilde{Q}_t - Q_t^\star\|_\infty \le \varepsilon$. The consistency error of a policy in (4) quantifies how the Q-value advice is close to optimal Q-value functions. It depends on various factors such the function approximation error or training error due to the distribution shift, and has a close connection to a rich literature on value function approximation [57–61]. The results in [59] generalized the worst-case L_∞ guarantees to arbitrary $L_{p,\rho}$-norms under some mixing assumptions via policy iteration for a stationary Markov decision process (MDP) with a continuous state space and a discrete action space. Recently, approximation guarantees for the average case for parametric policy classes (such as a neural network) of value functions have started to appear [57, 58, 60]. These bounds are useful in lots of supervised machine learning methods such as classification and regression, whose bounds are typically given on the expected error under some distribution. These results exemplify richer instances of the consistency definition (see (4)) and a summary of these bounds can be found in [61].

Now, we are ready to introduce our definition of consistency and robustness with respect to the ratio of expectations, similar to the growing literature on learning-augmented algorithms [3, 22, 23, 4]. We write the ratio of expectations $\mathsf{RoE}(\varepsilon)$ of a policy π as a function of the Q-value advice error ε in terms of the L_∞ norm, defined in (4).

Definition 4 (Consistency and Robustness)**.** *An algorithm π is said to be k-**consistent** if its worst-case (with respect to the MDP model $\mathsf{MDP}(\mathcal{X},\mathcal{U},T,P,c)$) ratio of expectations satisfies $\mathsf{RoE}(\varepsilon) \le k$ for $\varepsilon = 0$. On the other hand, it is l-**robust** if $\mathsf{RoE}(\varepsilon) \le l$ for any $\varepsilon > 0$.*

4 The Projection Pursuit Policy (PROP)

In this section we introduce our proposed algorithm (Algorithm 1), which achieves near-optimal consistency while bounding the robustness by leveraging a robust baseline (Section 3.1) in combination with value-based advice (Section 3.2). A key challenge in the design is how to exploit the benefits of good value-based advice while avoiding following it too closely when it performs poorly. To address this challenge, we propose to judiciously project the value-based advice into a neighborhood of the robust baseline. By doing so, the actions we choose can follow the value-based advice for consistency while staying close to the robust baseline for robustness. More specifically, at each step $t \in [T]$, we choose $u_t = \mathrm{Proj}_{\overline{\mathcal{U}}_t}(\widetilde{u}_t)$ where a projection operator $\mathrm{Proj}_{\overline{\mathcal{U}}_t}(\cdot) : \mathcal{U} \to \mathcal{U}$ is defined as

$$\mathrm{Proj}_{\overline{\mathcal{U}}_t}(u) \coloneqq \operatorname*{arg\,min}_{v\in\mathcal{U}} \|u - v\|_{\mathcal{U}} \text{ subject to } \|v - \overline{\pi}_t(x_t)\|_{\mathcal{U}} \le R_t, \tag{5}$$

corresponding to the projection of u onto a ball $\overline{\mathcal{U}}_t \coloneqq \{u \in \mathcal{U} : \|u - \overline{\pi}_t(x_t)\|_{\mathcal{U}} \le R_t\}$. Note that when the optimal solution of (5) is not unique, we choose the one on the same line with $\overline{\pi}_t(x_t) - u$.

The PROjection Pursuit policy, abbreviated as PROP, can be described as follows. For a time step $t \in [T]$, let $\widetilde{\pi}_t : \mathcal{X} \to \mathcal{U}$ denote a policy that chooses an action $\widetilde{u}_t$ (arbitrarily choose one if there are multiple minimizers of $\widetilde{Q}_t$), given the current system state x_t at time $t \in [T]$ and step $t \in [T]$. An action $u_t = \mathrm{Proj}_{\overline{\mathcal{U}}_t}(\widetilde{u}_t(x_t))$ is selected by projecting the machine-learned action $\widetilde{u}_t(x_t)$ onto a norm ball $\overline{\mathcal{U}}_t$ defined by the robust policy $\overline{\pi}$ given a radius $R_t \ge 0$. Finally, PROP applies to both black-box and grey-box settings (which differ from each other in terms of how the radius R_t is decided). The results under both settings are provided in Section 5, revealing a tradeoff between consistency and robustness.

The radii $(R_t : t \in [T])$ can be interpreted as *robustness budgets* and are key design parameters that determine the consistency and robustness tradeoff. Intuitively, the robustness budgets reflect the trustworthiness on the value-based policy $\widetilde{\pi}$ — the larger budgets, the more trustworthiness and hence the more freedom for PROP to follow $\widetilde{\pi}$. How the robustness budget is chosen differentiates the grey-box setting from the black-box one.

Algorithm 1 PROjection **P**ursuit Policy (**PROP**)

Initialize : Untrusted policy $\widetilde{\pi} = (\widetilde{\pi}_t : t \in [T])$ and baseline policy $\overline{\pi} = (\overline{\pi}_t : t \in [T])$

for $t = 0, \ldots, T-1$ **do**
  //Implement black-box (Section 4.1) or grey-box (Section 4.2) procedures
  $(\widetilde{u}_t, R_t) \leftarrow$BLACK-BOX$(x_t)$ or $(\widetilde{u}_t, R_t) \leftarrow$GREY-BOX$(x_t)$
  Set action $u_t = \mathrm{Proj}_{\overline{\mathcal{U}}_t}(\widetilde{u}_t)$ where $\overline{\mathcal{U}}_t \coloneqq \{u \in \mathcal{U} : \|u - \overline{\pi}_t(x_t)\|_{\mathcal{U}} \leq R_t\}$
  Sample next state $x_{t+1} \sim P_t(\cdot|x_t, u_t)$
end

4.1 Black-Box Setting

In the black-box setting, the only information provided by $\widetilde{\pi}$ is a suggested action $\widetilde{u}$ for the learning-augmented algorithm. Meanwhile, the robust policy $\overline{\pi}$ can also be queried to provide advice $\overline{u}$. Thus, without additional information, a natural way to utilize both $\widetilde{\pi}$ and $\overline{\pi}$ is to decide a projection radius at each time based on the how the obtained $\widetilde{u}$ and $\overline{u}$. More concretely, at each time $t \in [T]$, the robustness budget R_t is chosen by the following BLACK-BOX Procedure, where we set $R_t = \lambda\eta_t$ with $\eta_t \coloneqq \|\widetilde{u}_t - \overline{u}_t\|_{\mathcal{U}}$ representing the difference between the two advice measured in terms of the norm $\|\cdot\|_{\mathcal{U}}$ and $0 \leq \lambda \leq 1$ being a tradeoff hyper-parameter that measures the trustworthiness on the machine-learned advice. The choice of $R_t = \lambda\eta_t$ can be explained as follows. The value of η_t indicates the intrinsic discrepancy between the robust advice and the machine-learned untrusted advice — the larger discrepancy, the more difficult to achieve good consistency and robustness simultaneously. Given a robust policy and an untrusted policy, by setting a larger λ, we allow the actual action to deviate more from the robust advice and to follow the untrusted advice more closely, and vice versa. λ is a crucial hyper-parameter that can be pre-determined to yield a desired consistency and robustness tradeoff. The computation of R_t is summarized in Procedure 1 below.

Procedure 1 BLACK-BOX Procedure at $t \in [T]$ (Input: state x_t and hyper-parameter $0 \leq \lambda \leq 1$)

Implement $\widetilde{\pi}_t$ and $\overline{\pi}_t$ to obtain $\widetilde{u}_t$ and $\overline{u}_t$, respectively.
Set robustness budget $R_t = \lambda\eta_t$ where $\eta_t \coloneqq \|\widetilde{u}_t - \overline{u}_t\|_{\mathcal{U}}$; Return $(\widetilde{u}_t, R_t)$

4.2 Grey-Box Setting

In the grey-box setting, along with the suggested action $\widetilde{u}$, the value-based untrusted policy $\widetilde{\pi}$ also provides an estimate of the Q-value function $\widetilde{Q}$ that indicates the long-term cost impact of an action. To utilize such additional information informed by $\widetilde{Q}_t$ at each time $t \in [T]$, we propose a novel algorithm that dynamically adjusts the budget R_t to further improve the consistency and robustness tradeoff. More concretely, let us consider the Temporal-Difference (TD) error $\mathsf{TD}_t = c_{t-1} + \mathbb{P}_{t-1}\widetilde{V}_t - \widetilde{Q}_{t-1}$. Intuitively, if a non-zero TD-error is observed, the budget R_t needs to be decreased so as to minimize the impact of the learning error. However, the exact TD-error is difficult to compute in practice, since it requires complete knowledge of the transition kernels $(P_t : t \in [T])$. To address this challenge, we use the following estimated TD-error based on previous trajectories:

$$\delta_t(x_t, x_{t-1}, u_{t-1}) \coloneqq c_{t-1}(x_{t-1}, u_{t-1}) + \inf_{v \in \mathcal{U}} \widetilde{Q}_t(x_t, v) - \widetilde{Q}_{t-1}(x_{t-1}, u_{t-1}). \tag{6}$$

Denote by $\beta > 0$ a hyper-parameter. Based on the estimated TD-error in (6), the *robustness budget* in Algorithm 1 is set as

$$R_t \coloneqq \left[\underbrace{\|\widetilde{\pi}_t(x_t) - \overline{\pi}_t(x_t)\|_{\mathcal{U}}}_{\text{Decision Discrepancy } \eta_t} - \frac{\beta}{L_Q} \sum_{s=1}^{t} \underbrace{\delta_s(x_s, x_{s-1}, u_{s-1})}_{\text{Approximate TD-Error}} \right]^+, \tag{7}$$

which constitutes two terms. The first term $\eta_t \coloneqq \|\widetilde{\pi}_t(x_t) - \overline{\pi}_t(x_t)\|$ measures the *decision discrepancy* between the untrusted policy $\widetilde{\pi}$ and the baseline policy $\overline{\pi}$, which normalizes the total budget, similar to the one used in the black-box setting in Procedure 1. The second term is the approximate TD-error, which is normalized by the Lipschitz constant L_Q of Q-value functions. With these terms defined, the GREY-BOX Procedure below first chooses a suggested action $\widetilde{u}_t$ by minimizing $\widetilde{Q}_t$ and then decides a robustness budget R_t using (7).

Procedure 2 GREY-BOX Procedure at $t \in [T]$ (Input: state x_t and hyper-parameter $0 \leq \beta \leq 1$)

Obtain advice $\widetilde{Q}_t$ and $\widetilde{u}_t$ where $\widetilde{u}_t \in \arg\inf_{v \in \mathcal{U}} \widetilde{Q}_t(x_t, v)$
Implement $\overline{\pi}_t$ and obtain $\overline{u}_t$
Set robustness budget R_t as (7); Return $(\widetilde{u}_t, R_t)$

5 Main Results

We now formally present the main results for both the black-box and grey-box settings. Our results not only quantify the tradeoffs between consistency and robustness formally stated in Definition 4 with respect to the ratio of expectations, but also emphasize a crucial role that additional information about the estimated Q-values plays toward improving the consistency and robustness tradeoff.

5.1 Black-Box Setting

In the existing learning-augmented algorithms, the untrusted machine-learned policy $\widetilde{\pi}$ is often treated as a black-box that generates action advice $\widetilde{u}_t$ at each time $t \in [T]$. Our first result is the following general dynamic regret bound for the black-box setting (Section 4.1). We utilize the Big-O notation, denoted as $\mathcal{O}(\cdot)$ and $o(\cdot)$ to disregard inessential constants.

Theorem 5.1. *Suppose the machine-learned policy $\widetilde{\pi}$ is (∞, ε)-consistent. For any MDP model satisfying Assumption 1,2, and 3, the expected dynamic regret of* PROP *with the* BLACK-BOX *Procedure is bounded by* $\mathsf{DR}(\mathsf{PROP}) \leq \min\{\mathcal{O}(\varepsilon) + \mathcal{O}((1-\lambda)\gamma T), \mathcal{O}\left((\mathsf{ROB} + \lambda\gamma - 1)\, T\right)\}$ *where ε is defined in* (4)*, γ is the diameter of the action space $\mathcal{U}$, T is the length of the time horizon,* ROB *is the ratio of expectations of the robust baseline $\overline{\pi}$, and $0 \leq \lambda \leq 1$ is a hyper-parameter.*

When λ increases, the actual action can deviate more from the robust policy, making the dynamic regret potentially closer to that of the value-based policy. While the regret bound in Theorem 5.1 clearly shows the role of λ in terms of controlling how closely we follow the robust policy, the dynamic regret given a fixed $\lambda \in [0, 1]$ grows linearly in $\mathcal{O}(T)$. In fact, the linear growth of dynamic regret holds even if the black-box policy $\widetilde{\pi}$ is consistent, i.e., ε is small. This can be explained by noting the lack of dynamically tuning λ to follow the better of the two policies — even when one policy is nearly perfect, the actual action still always deviates from it due to the fixed choice of λ.

Consider any MDP model satisfying Assumptions 1,2, and 3. Following the classic definitions of consistency and robustness (see Definition 4), we summarize the following characterization of PROP, together with a negative result in Theorem 5.3. Proofs of Theorem 5.1, 5.2, and 5.3 are detailed in Appendix C.

Theorem 5.2 (BLACK-BOX Consistency and Robustness)**.** PROP *with the* BLACK-BOX *Procedure is $(1 + \mathcal{O}((1-\lambda)\gamma))$-consistent and $(\mathsf{ROB} + \mathcal{O}(\lambda\gamma))$-robust where $0 \leq \lambda \leq 1$ is a hyper-parameter.*

Theorem 5.3 (BLACK-BOX Impossibility)**.** PROP *with the* BLACK-BOX *Procedure cannot be both $(1 + o((1-\lambda)\gamma))$-consistent and $(\mathsf{ROB} + o(\lambda\gamma))$-robust for any $0 \leq \lambda \leq 1$.*

5.2 Grey-Box Setting

To overcome the impossibility result in the black-box setting, we dynamically tune the robustness budgets by tapping into additional information informed by the estimated Q-value functions using the GREY-BOX Procedure (Section 4.2). By setting the robustness budgets in (7), an analogous result of Theorem 5.1 is given in Appendix D, which leads to a dynamic regret bound of PROP in the grey-box setting (Theorem D.1 in Appendix D). Consider any MDP model satisfying Assumptions 1,2, and 3. Our main result below indicates that knowing more structural information about a black-box policy can indeed bring additional benefits in terms of the consistency and robustness tradeoff, even if the black-box policy is untrusted.

Theorem 5.4 (GREY-BOX Consistency and Robustness)**.** PROP *with the* GREY-BOX *Procedure is 1-consistent and $(\mathsf{ROB} + o(1))$-robust for some $\beta > 0$.*

Theorem 5.3 implies that using the BLACK-BOX Procedure, PROP cannot be 1-consistent and $(\mathsf{ROB} + o(1))$-robust, while this can be achieved using the GREY-BOX Procedure. On one hand, this theorem validates the effectiveness of the PROP policy with value-based machine-learned advice that

may not be fully trusted. On the other hand, this sharp contrast between the black-box and grey-box settings reveals that having access to information of value function can improve the tradeoff between consistency and robustness (see Definition 4) non-trivially. A proof of Theorem 5.4 can be found in Appendix D. Applications of our main results are discussed in Appendix A.

6 Concluding Remarks

Our results contribute to the growing body of literature on learning-augmented algorithms for MDPs and highlight the importance of considering consistency and robustness in this context. In particular, we have shown that by utilizing the *structural information* of machine learning methods, it is possible to achieve improved performance over a black-box approach. The results demonstrate the potential benefits of utilizing value-based policies as advice; however, there remains room for future work in exploring other forms of structural information.

Limitations and Future Work. One limitation of our current work is the lack of analysis of more general forms of black-box procedures. Understanding and quantifying the available structural information in a more systematic way is another future direction that could lead to advances in the design of learning-augmented online algorithms and their applications in various domains.

Acknowledgement

We would like to thank the anonymous reviewers for their helpful comments. This work was supported in part by the National Natural Science Foundation of China (NSFC) under grant No. 72301234, the Guangdong Key Lab of Mathematical Foundations for Artificial Intelligence, and the start-up funding UDF01002773 of CUHK-Shenzhen. Yiheng Lin was supported by the Caltech Kortschak Scholars program. Shaolei Ren was supported in part by the U.S. U.S. National Science Foundation (NSF) under grant CNS–1910208. Adam Wierman was supported in part by the U.S. NSF under grants CNS–2146814, CPS–2136197, CNS–2106403, NGSDI–2105648.

References

[1] Michael Mitzenmacher and Sergei Vassilvitskii. Algorithms with predictions. *Communications of the ACM*, 65(7):33–35, 2022.

[2] Tongxin Li. *Learning-Augmented Control and Decision-Making: Theory and Applications in Smart Grids*. PhD thesis, California Institute of Technology, 2023.

[3] Manish Purohit, Zoya Svitkina, and Ravi Kumar. Improving online algorithms via ml predictions. *Advances in Neural Information Processing Systems*, 31, 2018.

[4] Nicolas Christianson, Tinashe Handina, and Adam Wierman. Chasing convex bodies and functions with black-box advice. In *Conference on Learning Theory*, pages 867–908. PMLR, 2022.

[5] Ofir Nachum, Mohammad Norouzi, Kelvin Xu, and Dale Schuurmans. Bridging the gap between value and policy based reinforcement learning. *Advances in neural information processing systems*, 30, 2017.

[6] Zhuoran Yang, Chi Jin, Zhaoran Wang, Mengdi Wang, and Michael I Jordan. On function approximation in reinforcement learning: optimism in the face of large state spaces. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, pages 13903–13916, 2020.

[7] Noah Golowich and Ankur Moitra. Can q-learning be improved with advice? In *Conference on Learning Theory*, pages 4548–4619. PMLR, 2022.

[8] Chen-Yu Wei and Haipeng Luo. Non-stationary reinforcement learning without prior knowledge: An optimal black-box approach. In *Conference on Learning Theory*, pages 4300–4354. PMLR, 2021.

[9] Weichao Mao, Kaiqing Zhang, Ruihao Zhu, David Simchi-Levi, and Tamer Basar. Near-optimal model-free reinforcement learning in non-stationary episodic mdps. In *International Conference on Machine Learning*, pages 7447–7458. PMLR, 2021.

[10] Yuwei Luo, Varun Gupta, and Mladen Kolar. Dynamic regret minimization for control of non-stationary linear dynamical systems. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 6(1):1–72, 2022.

[11] Peng Zhao, Long-Fei Li, and Zhi-Hua Zhou. Dynamic regret of online markov decision processes. In *International Conference on Machine Learning*, pages 26865–26894. PMLR, 2022.

[12] Christian Scheller, Yanick Schraner, and Manfred Vogel. Sample efficient reinforcement learning through learning from demonstrations in minecraft. In *NeurIPS 2019 Competition and Demonstration Track*, pages 67–76. PMLR, 2020.

[13] Matthew Botvinick, Sam Ritter, Jane X Wang, Zeb Kurth-Nelson, Charles Blundell, and Demis Hassabis. Reinforcement learning, fast and slow. *Trends in cognitive sciences*, 23(5):408–422, 2019.

[14] Xueying Bai, Jian Guan, and Hongning Wang. A model-based reinforcement learning with adversarial training for online recommendation. *Advances in Neural Information Processing Systems*, 32, 2019.

[15] Guanya Shi, Yiheng Lin, Soon-Jo Chung, Yisong Yue, and Adam Wierman. Online optimization with memory and competitive control. *Advances in Neural Information Processing Systems*, 33:20636–20647, 2020.

[16] Gautam Goel and Babak Hassibi. Competitive control. *IEEE Transactions on Automatic Control*, 2022.

[17] Tongxin Li, Ruixiao Yang, Guannan Qu, Guanya Shi, Chenkai Yu, Adam Wierman, and Steven Low. Robustness and consistency in linear quadratic control with untrusted predictions. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 6(1):1–35, 2022.

[18] Oron Sabag, Sahin Lale, and Babak Hassibi. Optimal competitive-ratio control. *arXiv preprint arXiv:2206.01782*, 2022.

[19] Yiheng Lin, Yang Hu, Guannan Qu, Tongxin Li, and Adam Wierman. Bounded-regret mpc via perturbation analysis: Prediction error, constraints, and nonlinearity. *Advances in Neural Information Processing Systems*, 35:36174–36187, 2022.

[20] Tongxin Li, Yue Chen, Bo Sun, Adam Wierman, and Steven H Low. Information aggregation for constrained online control. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 5(2):1–35, 2021.

[21] Tongxin Li, Bo Sun, Yue Chen, Zixin Ye, Steven H Low, and Adam Wierman. Learning-based predictive control via real-time aggregate flexibility. *IEEE Transactions on Smart Grid*, 12(6):4897–4913, 2021.

[22] Alexander Wei and Fred Zhang. Optimal robustness-consistency trade-offs for learning-augmented online algorithms. *Advances in Neural Information Processing Systems*, 33:8042–8053, 2020.

[23] Shom Banerjee. Improving online rent-or-buy algorithms with sequential decision making and ml predictions. *Advances in Neural Information Processing Systems*, 33:21072–21080, 2020.

[24] Stephen Tu, Alexander Robey, Tingnan Zhang, and Nikolai Matni. On the sample complexity of stability constrained imitation learning. In *Learning for Dynamics and Control Conference*, pages 180–191. PMLR, 2022.

[25] Hiroyasu Tsukamoto, Soon-Jo Chung, and Jean-Jaques E Slotine. Contraction theory for nonlinear stability analysis and learning-based control: A tutorial overview. *Annual Reviews in Control*, 52:135–169, 2021.

[26] Mohammad Mahdian, Hamid Nazerzadeh, and Amin Saberi. Online optimization with uncertain information. *ACM Transactions on Algorithms (TALG)*, 8(1):1–29, 2012.

[27] Dhruv Rohatgi. Near-optimal bounds for online caching with machine learned advice. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1834–1845. SIAM, 2020.

[28] Thodoris Lykouris and Sergei Vassilvitskii. Competitive caching with machine learned advice. *Journal of the ACM (JACM)*, 68(4):1–25, 2021.

[29] Sungjin Im, Ravi Kumar, Aditya Petety, and Manish Purohit. Parsimonious learning-augmented caching. In *International Conference on Machine Learning*, pages 9588–9601. PMLR, 2022.

[30] Antonios Antoniadis, Christian Coester, Marek Elias, Adam Polak, and Bertrand Simon. Online metric algorithms with untrusted predictions. In *International Conference on Machine Learning*, pages 345–355. PMLR, 2020.

[31] Etienne Bamas, Andreas Maggiori, and Ola Svensson. The primal-dual method for learning augmented algorithms. *Advances in Neural Information Processing Systems*, 33:20083–20094, 2020.

[32] Keerti Anand, Rong Ge, Amit Kumar, and Debmalya Panigrahi. Online algorithms with multiple predictions. In *International Conference on Machine Learning*, pages 582–598. PMLR, 2022.

[33] Tongxin Li, Ruixiao Yang, Guannan Qu, Yiheng Lin, Adam Wierman, and Steven H Low. Certifying black-box policies with stability for nonlinear control. *IEEE Open Journal of Control Systems*, 2023.

[34] Ilias Diakonikolas, Vasilis Kontonis, Christos Tzamos, Ali Vakilian, and Nikos Zarifis. Learning online algorithms with distributional advice. In *International Conference on Machine Learning*, pages 2687–2696. PMLR, 2021.

[35] Yiheng Lin, Yang Hu, Guanya Shi, Haoyuan Sun, Guannan Qu, and Adam Wierman. Perturbation-based regret analysis of predictive control in linear time varying systems. *Advances in Neural Information Processing Systems*, 34:5174–5185, 2021.

[36] Yiheng Lin, Yang Hu, Guannan Qu, Tongxin Li, and Adam Wierman. Bounded-regret mpc via perturbation analysis: Prediction error, constraints, and nonlinearity. In *Advances in Neural Information Processing Systems*, 2022.

[37] David Hoeller, Farbod Farshidian, and Marco Hutter. Deep value model predictive control. In *Conference on Robot Learning*, pages 990–1004. PMLR, 2020.

[38] Richard Cheng, Abhinav Verma, Gabor Orosz, Swarat Chaudhuri, Yisong Yue, and Joel Burdick. Control regularization for reduced variance reinforcement learning. In *International Conference on Machine Learning*, pages 1141–1150. PMLR, 2019.

[39] Lukas Brunke, Melissa Greeff, Adam W Hall, Zhaocong Yuan, Siqi Zhou, Jacopo Panerati, and Angela P Schoellig. Safe learning in robotics: From learning-based control to safe reinforcement learning. *Annual Review of Control, Robotics, and Autonomous Systems*, 5:411–444, 2022.

[40] Lin Guan, Mudit Verma, Suna Sihang Guo, Ruohan Zhang, and Subbarao Kambhampati. Widening the pipeline in human-guided reinforcement learning with explanation and context-aware data augmentation. *Advances in Neural Information Processing Systems*, 34:21885–21897, 2021.

[41] Paul F Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep reinforcement learning from human preferences. *Advances in neural information processing systems*, 30, 2017.

[42] James MacGlashan, Mark K Ho, Robert Loftin, Bei Peng, Guan Wang, David L Roberts, Matthew E Taylor, and Michael L Littman. Interactive learning from policy-dependent human feedback. In *International Conference on Machine Learning*, pages 2285–2294. PMLR, 2017.

[43] Leo Gao, John Schulman, and Jacob Hilton. Scaling laws for reward model overoptimization. In *International Conference on Machine Learning*, pages 10835–10866. PMLR, 2023.

[44] Felix Berkenkamp, Matteo Turchetta, Angela Schoellig, and Andreas Krause. Safe model-based reinforcement learning with stability guarantees. *Advances in neural information processing systems*, 30, 2017.

[45] Mengdi Xu, Zuxin Liu, Peide Huang, Wenhao Ding, Zhepeng Cen, Bo Li, and Ding Zhao. Trustworthy reinforcement learning against intrinsic vulnerabilities: Robustness, safety, and generalizability. *arXiv preprint arXiv:2209.08025*, 2022.

[46] Matthew E Taylor and Peter Stone. Transfer learning for reinforcement learning domains: A survey. *Journal of Machine Learning Research*, 10(7), 2009.

[47] Irina Higgins, Arka Pal, Andrei Rusu, Loic Matthey, Christopher Burgess, Alexander Pritzel, Matthew Botvinick, Charles Blundell, and Alexander Lerchner. Darla: Improving zero-shot transfer in reinforcement learning. In *International Conference on Machine Learning*, pages 1480–1490. PMLR, 2017.

[48] Elynn Y Chen, Michael I Jordan, and Sai Li. Transferred q-learning. *arXiv preprint arXiv:2202.04709*, 2022.

[49] Peter Auer, Thomas Jaksch, and Ronald Ortner. Near-optimal regret bounds for reinforcement learning. *Advances in neural information processing systems*, 21, 2008.

[50] Mohammad Gheshlaghi Azar, Ian Osband, and Rémi Munos. Minimax regret bounds for reinforcement learning. In *International Conference on Machine Learning*, pages 263–272. PMLR, 2017.

[51] Allan Borodin and Ran El-Yaniv. *Online computation and competitive analysis*. cambridge university press, 2005.

[52] Nikhil R Devanur and Thomas P Hayes. The adwords problem: online keyword matching with budgeted bidders under random permutations. In *Proceedings of the 10th ACM conference on Electronic commerce*, pages 71–78, 2009.

[53] Naman Agarwal, Brian Bullins, Elad Hazan, Sham Kakade, and Karan Singh. Online control with adversarial disturbances. In *International Conference on Machine Learning*, pages 111–119. PMLR, 2019.

[54] Elad Hazan, Sham Kakade, Karan Singh, and Abby Van Soest. Provably efficient maximum entropy exploration. In *International Conference on Machine Learning*, pages 2681–2691. PMLR, 2019.

[55] Paula Gradu, John Hallman, and Elad Hazan. Non-stochastic control with bandit feedback. *Advances in Neural Information Processing Systems*, 33:10764–10774, 2020.

[56] Sheldon M Ross. *Stochastic processes*. John Wiley & Sons, 1995.

[57] Dimitri P Bertsekas and John N Tsitsiklis. Neuro-dynamic programming: an overview. In *Proceedings of 1995 34th IEEE conference on decision and control*, volume 1, pages 560–564. IEEE, 1995.

[58] Rémi Munos. Error bounds for approximate policy iteration. In *ICML*, volume 3, pages 560–567. Citeseer, 2003.

[59] András Antos, Csaba Szepesvári, and Rémi Munos. Learning near-optimal policies with bellman-residual minimization based fitted policy iteration and a single sample path. *Machine Learning*, 71(1):89–129, 2008.

[60] Matthieu Geist, Bruno Scherrer, and Olivier Pietquin. A theory of regularized markov decision processes. In *International Conference on Machine Learning*, pages 2160–2169. PMLR, 2019.

[61] Alekh Agarwal, Sham M Kakade, Jason D Lee, and Gaurav Mahajan. On the theory of policy gradient methods: Optimality, approximation, and distribution shift. *J. Mach. Learn. Res.*, 22(98):1–76, 2021.

[62] Yiheng Lin, James A Preiss, Emile Timothy Anand, Yingying Li, Yisong Yue, and Adam Wierman. Online adaptive policy selection in time-varying systems: No-regret via contractive perturbations. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.

[63] Runyu Zhang, Yingying Li, and Na Li. On the regret analysis of online LQR control with predictions. In *2021 American Control Conference (ACC)*, pages 697–703. IEEE, 2021.

[64] James R Norris. *Markov chains*. Cambridge university press, 1998.

[65] Steve Lalley. Markov chains: Basic theory. `https://galton.uchicago.edu/~lalley/Courses/312/MarkovChains.pdf`, 2016.

[66] Chi Jin, Zeyuan Allen-Zhu, Sebastien Bubeck, and Michael I Jordan. Is q-learning provably efficient? *Advances in neural information processing systems*, 31, 2018.

[67] Yingying Li and Na Li. Online learning for markov decision processes in nonstationary environments: A dynamic regret analysis. In *2019 American Control Conference (ACC)*, pages 1232–1237. IEEE, 2019.

[68] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.

[69] David Silver, Guy Lever, Nicolas Heess, Thomas Degris, Daan Wierstra, and Martin Riedmiller. Deterministic policy gradient algorithms. In *International conference on machine learning*, pages 387–395. Pmlr, 2014.

[70] Leonid Kantorovitch. On the translocation of masses. *Management science*, 5(1):1–4, 1958.

[71] Niranjan Srinivas, Andreas Krause, Sham Kakade, and Matthias Seeger. Gaussian process optimization in the bandit setting: no regret and experimental design. In *Proceedings of the 27th International Conference on International Conference on Machine Learning*, pages 1015–1022, 2010.

A Application Examples

In this section, we delve deeper into the practical applications of our main results, which provide a general consistency and robustness tradeoff. By presenting concrete examples, we aim to demonstrate the versatility and relevance of our findings to various real-world problems and scenarios. We consider the settings summarized in Table 1. These examples illustrate how our results can be applied to optimize tradeoffs between consistency and robustness for more specific models that can be represented as a nonstationary MDP. Additionally, these examples highlight the significance of considering the tradeoff between consistency and robustness in the design and implementation of decision-making algorithms in the learning-augmented framework, and the impact of the structural information in the grey-box setting.

A.1 MPC baseline in time-varying dynamical systems

The first application is an online optimal control problem, which is a special case of the general MDP in Section 2. Suppose that the dynamics and cost function in time $t \in [T]$ are given by

$$x_{t+1} = A_t x_t + B_t u_t + w_t \tag{8}$$

and

$$c_t(x_t, u_t) = \frac{1}{2}\left((x_t)^\top Q_t x_t + (u_t)^\top R_t u_t\right). \tag{9}$$

Here, $(w_t : t \in [T])$ is a sequence of bounded and oblivious disturbances that is unknown to the online controller. [2] At each time step, the controller observes (A_t, B_t, Q_t, R_t) for future k time steps but all future disturbances are unknown. Since we assume $c_T \equiv 0$, the optimal u_{T-1} is always 0 and the online control problem in episode t actually terminates after the state x_{T-1} is revealed.

We show how to apply Model Predictive Control (MPC) with robust predictions as the robust baseline in our framework [17, 62]. To define MPC with robust predictions, we first need to define the finite-time optimal control problem (FTOCP) solved by MPC at every step: For $t, t' \in [T]$, we define

$$\begin{aligned}\psi_{t,t'}\left(x_t, \left(w_{\tau|t} : \tau \in [t : t'-1]\right); P_{t'}\right) = \underset{u_{t:(t'-1)|t}}{\arg\min} & \sum_{\tau=t}^{t'-1} c_\tau\left(x_{\tau|t}, u_{\tau|t}\right) + \frac{1}{2} x_{t'|t}^\top P_{t'} x_{t'|t} \\ \text{s.t. } & x_{\tau+1|t} = A_t x_{\tau|t} + B_t u_{\tau|t} + w_{\tau|t}, \forall \tau \in [t : t'-1]; \\ & x_{t|t} = x_t,\end{aligned}$$

Here, $\left(w_{\tau|t} : \tau \in [t : t'-1]\right)$ can be viewed as the predicted future disturbances, and MPC with robust predictions sets them to zero vectors, therefore becomes robust against (potentially) adversarial environments with large disturbances $(w_t : t \in [nt])$. The term $x_{t'|t}^\top P_{t'} x_{t'|t}/2$ is a terminal cost that regularizes the last predictive state. To simplify the notation, we use the shorthand $\psi_{t,t'}(x_t; P) \coloneqq \psi_{t,t'}(x_t, 0_{\times(t'-t)}; P)$, where $0_{\times(t'-t)}$ denotes a sequence of $(t'-t)$ zeros, and

$$\psi_{t,t'}(x_t) \coloneqq \begin{cases} \psi_{t,t'}(x_t; P_{t'}) & \text{if } t' < T-1, \\ \psi_{t,t'}(x_t; Q_{t'}) & \text{if } t' = T-1. \end{cases}$$

With this notation, we define MPC with robust predictions formally in Algorithm 3. At each time step, it solves a k-step predictive FTOCP and commits the first control action in the optimal solution. Since the future disturbances are unknown, MPC predicts them to zero vectors. The terminal cost matrices P_t are pre-determined.

We make some standard assumptions, following those in the literature of online control [35, 63, 19]. The first assumption is that the cost functions are well-conditioned and the dynamical matrices are uniformly bounded.

Assumption 4. *For any $t \in [T]$, we have $\|A_t\| \le a, \|B_t\| \le b, \|w_t\| \le d$, and*

$$\mu I_n \preceq Q_t \preceq \ell I_n, \mu I_m \preceq R_t \preceq \ell I_m, \mu I_n \preceq P \preceq \ell I_n.$$

[2] By oblivious, we mean the sequence $(w_t : t \in [T])$ is determined by the environment before the game starts and are not random.

Algorithm 3 MPC with Robust Predictions (MPC_k)

Initialize : Prediction horizon k and terminal cost matrix P.

for $t = 0, \ldots, T-1$ **do**
  Set $t' \leftarrow \min\{t+k, T-1\}$
  Observe x_t and $(A_\tau, B_\tau, Q_\tau, R_\tau : \tau \in [t : t'-1])$
  Set action $u_t = \psi_{t,t'}(x_t)[u_{t|t}]$
end

The second assumption guarantees that for arbitrary bounded disturbance sequences $(w_t : t \in [T])$, there exists a controller that can stabilize the system.

Assumption 5. *For any $t, t' \in [T], t \le t'$, define a block matrix*

$$\Xi^t_{t,t'} := \begin{bmatrix} I & & & & \\ -A_t & -B_t & I & & \\ & & \ddots & & \\ & & -A_{t'-1} & -B_{t'-1} & I \end{bmatrix}.$$

We assume $\sigma_{\min}(\Xi_{t,t'}) \ge \sigma$ for some positive constant σ, where $\sigma_{\min}(\cdot)$ denotes the smallest singular value of a matrix.

The interpretation of Assumption 5 is as follows. It holds with σ if and only if for any sequence $x_t, w_{t:t'-1}$ that satisfies $\left\|\left((x_t)^\top, (w_t)^\top, \ldots, (w_{t'-1})^\top\right)\right\| \le 1$, there exists a feasible trajectory $x_t, u_t, x_{t+1}, \ldots, u_{t'-1}, x_{t'}$ subject to

$$\left\|\left((x_t)^\top, (u_t)^\top, (x_{t+1})^\top, \ldots, (u_{t'-1})^\top, (x_{t'})^\top\right)\right\| \le \frac{1}{\sigma}.$$

Thus, Assumption 5 holds provided that there is an exponentially stabilizing controller. With these assumptions, we are ready to present our main result about MPC_k in Theorem A.1.

Theorem A.1. *Suppose Assumptions 4 and 5 hold. Consider the case when the robust baseline policy $\overline{\pi}$ is MPC_k (Algorithm 3), and for some $\overline{\lambda} \in (\lambda, 1)$, the prediction horizon k satisfies that*

$$k \ge \min\left\{T, \frac{1}{2}\log\left(C^3 ba\lambda/(\overline{\lambda}-\lambda)\right)/\log(1/\lambda)\right\}.$$

We also assume that $x_0 = 0$, and $R_t \le \overline{R}$. Then, the following holds for the robust baseline $\overline{\pi}$:

(i) The Wasserstein robustness (Definition 3) holds globally with $s(t) = C(1+C)(a+b)\overline{\lambda}^{t-1}$.

(ii) The PROP *controller (Algorithm 1) is always stable in the sense that*

$$\|x_t\| \le \overline{R}_x := \frac{C(d+b\overline{R})}{1-\overline{\lambda}} \text{ and } \|u_t\| \le \overline{R}_u := C\overline{R}_x + \overline{R}.$$

(iii) The competitive ratio of the robust baseline MPC_k satisfies that

$$\mathsf{ROB} \le \frac{2\ell C^2(1+C^2)(1+a^2+b^2)}{\mu(1-\overline{\lambda})^2}.$$

Here, the coefficients C and λ are given by $\lambda = \left(\frac{\overline{\sigma}-\underline{\sigma}}{\overline{\sigma}+\underline{\sigma}}\right)^{\frac{1}{2}}, C = \frac{4(\ell+1+a+b)}{\underline{\sigma}^2\cdot\lambda}$, where

$$\underline{\sigma} := \min(\mu, 1)\cdot(a+b+1)\cdot\sqrt{\frac{\ell}{2\mu\ell+\mu\sigma^2}}, \text{ and } \overline{\sigma} := \sqrt{2}(\ell+a+b+1).$$

The first result of Theorem A.1 shows that MPC_k satisfies the Wasserstein robustness (see Definition 3), which is the critical assumption we require for any robust baseline policy. The second result guarantees that PROP (with MPC_k as the robust baseline) will always stay in a bounded ball in the

Euclidean space as long as the radius R_t is uniformly upper bounded. Thus, we can assume $\mathcal{X}$ and $\mathcal{U}$ are compact without loss of generality. The third result gives an upper bound of the robust competitive ratio ROB, which in this application is a special deterministic case of the considered ratio of expectation (ROE) in the general results. With the settings above, we conclude that in the grey-box setting, PROP with GREY-BOX Procedure can be 1-consistent and $\left(\frac{2\ell C^2(1+C^2)(1+a^2+b^2)}{\mu(1-\overline{\lambda})^2}+o(1)\right)$-robust. We defer the detailed proof of Theorem A.1 to Appendix E.

A.2 Baseline Policies for MDPs with Finite State/Action Spaces

Our second example focuses on an MDP environment $(\mathcal{S}, \mathcal{A}, (\mathbb{P}_t : t \in [T]), (c_t : t \in [T]), T)$ with a finite state space $\mathcal{S}$ and a finite action space $\mathcal{A}$. Given a policy $\overline{\pi}_t : \mathcal{S} \to \Delta(\mathcal{A})$ for $t \in [T]$, let $\left(\overline{\mathbb{P}}_t\right)$ denote the state transition probability that maps $\mathcal{S}$ to $\Delta(\mathcal{S})$, which is defined as

$$\overline{\mathbb{P}}_t(s; s') = \sum_{a \in \mathcal{A}} \overline{\pi}_t(s; a)\mathbb{P}_t(s, a; s').$$

We consider the setting when every entry of $\overline{\mathbb{P}}_t$ is strictly positive. Under this assumption, one can show that the one-step transition probability $\overline{\mathbb{P}}_t$ is a contractive mapping in total variance distance [64]. We state this result formally in Lemma 1. To simplify the notation, for any $0 \le t \le t' < T$, we define the multi-step transition matrix as $\overline{\mathbb{P}}_{t:t'} \coloneqq \overline{\mathbb{P}}_t \overline{\mathbb{P}}_{t+1} \cdots \overline{\mathbb{P}}_{t'}$.

Lemma 1. *Under the assumption that* $\min_{t\in[T]} \min_{s,s'\in\mathcal{S}} \overline{\mathbb{P}}_t(s; s') \ge \epsilon$*, for any* $0 \le t \le t' < T$ *and distributions* $\mu, \nu \in \Delta(\mathcal{S})$*, we have that*

$$\mathsf{TV}\left(\mu^\top \overline{\mathbb{P}}_{t:t'}, \nu^\top \overline{\mathbb{P}}_{t:t'}\right) \le \lambda^{t'-t} \mathsf{TV}(\mu, \nu), \tag{10}$$

where $\lambda = 1 - |\mathcal{S}|\,\epsilon$.

Lemma 1 follows from Proposition 5 in [65]. In the case that not every entry of $\overline{\mathbb{P}}_t$ is strictly positive, but the entries of $\overline{\mathbb{P}}_{t:t+d}$ are strictly positive for some constant $d \in \mathbb{Z}_+$, we can still obtain a similar contraction property as Lemma 1 with a weaker decay rate λ. Note that the exponential contractive property in Lemma 1 is different with the one in Wasserstein robustness (Definition 3) because the distance between distributions are measured by total variance instead of Wasserstein distance. To convert it into the form required by Wasserstein robustness, we need to define an underlying metric for the discrete state/action space.

Without loss of generality, we assume $\mathcal{X} \coloneqq \{e_i : i = 1, \ldots, |S|\} \subseteq \mathbb{R}^{|\mathcal{S}|}$, where each element of $\mathcal{X}$ corresponds to a unique state in $\mathcal{S}$. Here, each e_i is an indicator vector of $\mathbb{R}^{|\mathcal{S}|}$ defined as

$$e_i(j) = \begin{cases} 1 & \text{if } i = j, \\ 0 & \text{otherwise.} \end{cases}$$

Since a policy in the discrete MDP maps $\mathcal{S}$ to $\Delta(\mathcal{A})$, we set $\mathcal{U} = \Delta(\mathcal{A}) \subseteq \mathbb{R}^{|\mathcal{A}|}$, which denotes the distribution of actions and is compact and convex. To define the Wasserstain distance, we adopt ℓ_1 distance as the metric on the state space $\mathcal{X}$, action space $\mathcal{U}$, and state-action space $\mathcal{X} \times \mathcal{U}$, i.e.,

$$\|(x, u) - (x', u')\|_1 = \|x - x'\|_1 + \|u - u'\|_1, \text{ for all } x, x' \in \mathcal{X}, u, u' \in \mathcal{U}.$$

Using these definitions, we can use the contraction property in the TV distance (Lemma 1) to establish the Wasserstein robustness of the baseline policy $\overline{\pi}$.

Theorem A.2. *Suppose the Markov chain on state space* $\mathcal{S}$ *induced by the baseline policy* $\overline{\pi} = (\overline{\pi}_t : t \in [T])$ *satisfies that* $\overline{\mathbb{P}}_t(s; s') \ge \epsilon$ *for all* $t \in [T]$ *and* $s, s' \in \mathcal{S}$*, then Wasserstein robustness (Definition 3) holds globally with* $s(t) = 2\lambda^{t-1}$*, where* $\lambda = 1 - |\mathcal{S}|\,\epsilon$.

We defer the proof of Theorem A.2 to Appendix F. Theorem A.2 shows that the Wasserstein robustness in Definition 3 is general enough to capture a wide class of baseline policies in finite state/action settings. It also enables comparison between our results and previous studies that assume discrete state/action spaces [66, 9, 7].

A.3 Numerical Results

In light of the applications detailed in Appendix A.1, we present two case studies. We consider linear dynamics as a specific instance of an MDP and use the MPC described in Algorithm 3 as our robust baseline.

A.3.1 Basic Settings

Dynamics. We investigate the impact of the hyper-parameter β in the robustness budget R_t in (7) by considering the following update rule:

$$\begin{bmatrix} d_{t+1} \\ v_{t+1} \end{bmatrix} = A \begin{bmatrix} d_t \\ v_t \end{bmatrix} + Bu_t + w_t, \tag{11}$$

which is cast in the canonical form (8). The system matrices A and B are defined as

$$A \coloneqq \begin{bmatrix} 1 & 0 & 0.2 & 0 \\ 0 & 1 & 0 & 0.2 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad B \coloneqq \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 0.2 & 0 \\ 0 & 0.2 \end{bmatrix}, \tag{12}$$

and $w_t \coloneqq Ay_t - y_{t+1}$, where $(y_t : t \in [T])$ specifies an unknown trajectory to be tracked. The choice of A, B and $(w_t : t \in [T])$ specifies a two-dimensional robot tracking problem as detailed in [67, 17]. In this application, the robot controller maneuvers along a fixed but unknown trajectory, given by $(y_t : t \in [T])$. At each time $t \in [T]$, the robot controller needs to decide an acceleration action u_t, without knowing the desired location y_t. It can only access the past trajectories $(y_\tau : \tau \in [t])$. The location of the robot controller at time $t+1$, denoted $l_{t+1} \in \mathbb{R}^2$, is determined by its prior location and its velocity $v_t \in \mathbb{R}^2$ according to $l_{t+1} = l_t + 0.2v_t$. Furthermore, at each subsequent time t, the controller has the ability to apply an adjustment u_t to alter its velocity, resulting in $v_{t+1} = v_t + 0.2u_t$ at the next time step. This system can be reformulated as (11) by letting $x_t = l_t - y_t$, the tracking error between the current location at time t and the desired location y_t.

To efficiently track the trajectory, we use quadratic costs as in (9) with

$$Q \coloneqq \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}, \quad R \coloneqq \begin{bmatrix} 10^{-2} & 0 \\ 0 & 10^{-2} \end{bmatrix}. \tag{13}$$

With the settings above, we encapsulate a Gym environment [68] with action space and state space defined as hyper-cubes in $\mathbb{R}^2$ and $\mathbb{R}^4$ such that each action/state coordinate is within $[-100, 100]$.

MPC Baseline $\overline{\pi}$. With predictions $(\widetilde{w}_t : t \in [T])$ of the perturbations satisfy $\widetilde{w}_t = 0$ for all $t \in [T]$, with a terminal cost matrix P as the solution of the discrete algebraic Riccati equation (DARE), the MPC baseline in Algorithm 3 can be stated as the following linear quadratic regulator $\overline{\pi}_{\mathsf{MPC}}(x_t) = -Kx_t$ where $K \coloneqq (R + B^\top PB)^{-1}B^\top PA$.

Machine-Learned Policy $\widetilde{\pi}$. We use the deep deterministic policy gradient (DDPG) algorithm [69] to generate machine learned advice and Q-value functions, with hyper-parameters set as in Table 2.

Table 2: Hyper-parameters used in DDPG.

Parameter	Value
Maximal number of episodes	10^3
Episode length	10^2
Discount factor	1.0
Actor network learning rate	10^{-3}
Critic network learning rate	10^{-3}
Soft target update parameter	10^{-3}
Replay buffer size	10^6
Minibatch size	128

PROP Implementation In our empirical implementation of PROP, we set $L_Q = 1$ in (6) and use $|\delta_t|$ at each t-th time step, instead of $\sum_{s=1}^{t} \delta_s$ to generate more stable results.

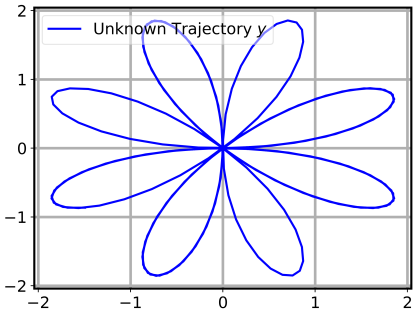

Figure 3: Unknown trajectory y in the case study that illustrates the impact of β.

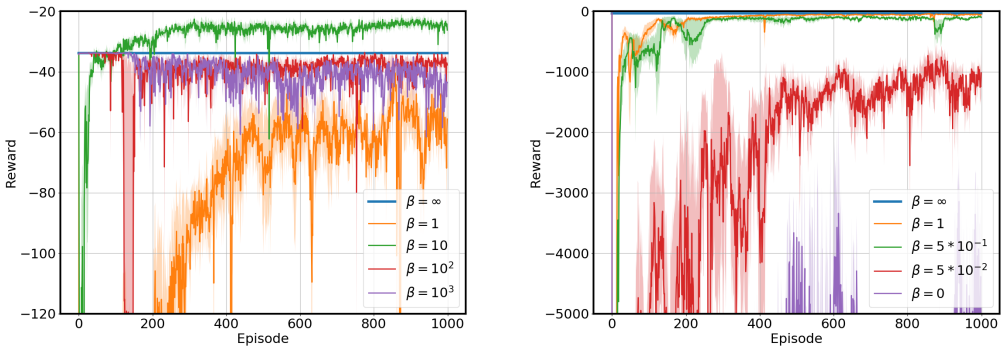

Figure 4: Average awards with varying choices of the hyper-parameter β in the robustness budget of PROP. Shadow area depicts the range of standard deviations for 5 random tests. Left: $\beta = 1, 10, 10^2, 10^3$, and ∞ (directly applying the MPC baseline); Right: $\beta = 0, 0.05, 0.5, 1$, and ∞.

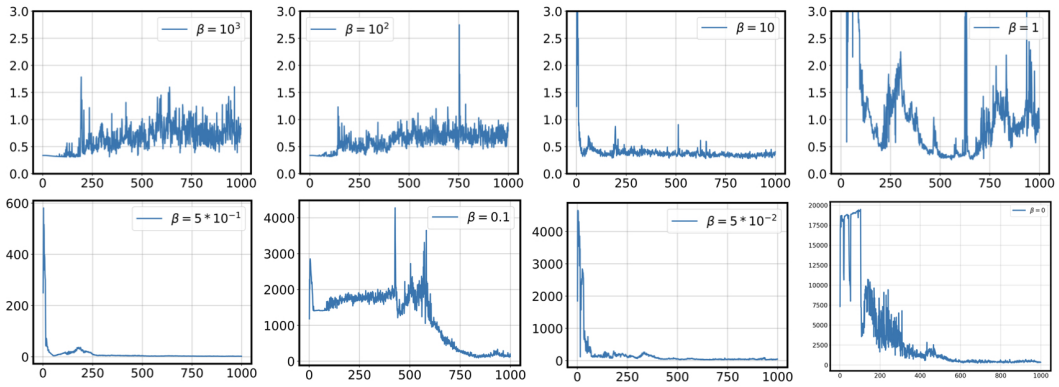

Figure 6: Average approximate TD-error (see (6)) with different choices of the hyper-parameter β in the robustness budget of PROP. Shadow area depicts the range of standard deviations for 5 random tests.

A.3.2 Case Studies

Impact of Hyper-Parameter β. With the basic settings described above, we delve into the effects of the hyper-parameter β on the selection of the robustness budget as in (7). This, in turn, influences the average rewards, projection radii, and the approximate TD-error for PROP.

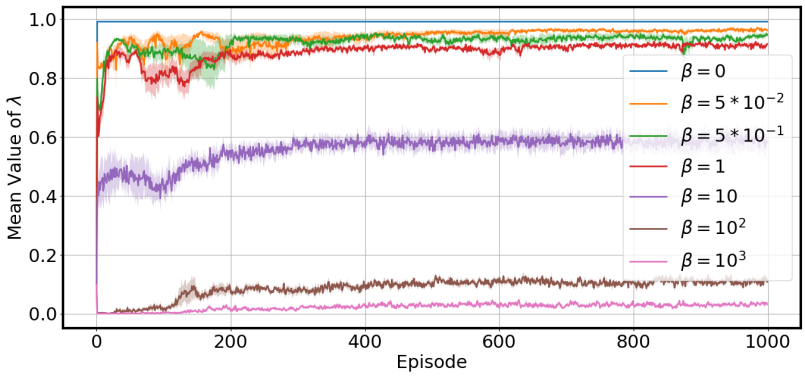

Figure 5: The influence of hyper-parameter β on the projection radii $(R_t : t \in [T])$. The shaded region represents the standard deviation range from 5 random tests. As β increases, the average trust coefficient $\lambda(R_t)$ decreases.

We set the unknown $(y_t : t \in [T])$ to be tracked as a rose-shaped trajectory shown in Figure 3:

$$y_t \coloneqq \begin{bmatrix} 2\cos\left(\frac{t}{20}\right)\sin\left(\frac{t}{5}\right) \\ 2\sin\left(\frac{t}{20}\right)\sin\left(\frac{t}{5}\right) \end{bmatrix}, \quad t \in [T]. \tag{14}$$

We vary the value of β from 0 up to ∞. It is important to highlight that when $\beta = 0$, PROP operates the same as the pure DDPG in our experiments. In contrast, when $\beta = \infty$, PROP is equivalent to the MPC baseline discussed earlier. Arbitrary exploration in the action space will lead to unstable states, causing the pure DDPG to remain non-convergent throughout its training process. From our experiments, we observe that setting β between 5 and 25 yields the largest average reward. The results are summarized in Figure 4.As noted in the proof of Lemma 3 in Appendix B.2, the action given by PROP at each time $t \in [T]$ can be written as

$$u_t = \lambda(R_t)\,\widetilde{\pi}_t(x_t) + (1 - \lambda(R_t))\,\overline{\pi}_t(x_t)$$

where $\lambda(R_t) \coloneqq \min\{1, R_t/\|\widetilde{\pi}_t(x_t) - \overline{\pi}_t(x_t)\|_{\mathcal{U}}\}$ serves as a *trust coefficient* between 0 and 1. Here, R_t is the robustness budget defined in (7). In Figure 5 we illustrate the behavior of $\lambda(R_t)$ averaged over all time steps and tests. Likewise, Figure 6 displays the evolution of the approximate TD-error with various selections of the hyper-parameter β, averaged over all time steps and tests. Notably, a distinct convergence of the approximate TD-error is evident when $\beta = 10$, which also yields high average rewards, shown in Figure 4. It's worth noting, however, that we did not actively optimize for β.

Non-Stationary Environment In a subsequent experiment, we address scenarios where there is a distribution shift in the underlying MDP. We use the same matrices A, B, Q, and R in (12) and (13).

For each w_t in (11), we treat it as an independent Gaussian vector. Specifically, every entry $w_t(i)$ of w_t is considered as an independent Gaussian random variable. For the first 700 episodes, each $w_t(i)$ is sampled from a normal distribution $\mathcal{N}(\mu, \sigma)$ where $\mu = 0.5$ and $\sigma = 0.05$. However, in the last 300 episodes, we adjust μ to -0.5.

In the context of this nonstationary MDP, Figure 7 illustrates the reward recovery after the occurrence of a distribution shift for varying choices of β. The two top figures highlight the average rewards: the top-left figure corresponds to an action space of $\mathcal{U} = [-100, 100]$, while the top-right is set to $\mathcal{U} = [-5, 5]$. On the bottom, the left figure presents the average behavior of $\lambda(R_t)$, and the right one illustrates the average approximate TD-error for the case $\mathcal{U} = [-100, 100]$ With $\beta = 10$ for $\mathcal{U} = [-100, 100]$ and $\beta = 1$ for $\mathcal{U} = [-5, 5]$ respectively, there is an evident near-optimal tradeoff between consistency and robustness. PROP consistently achieves notable average rewards before the distribution shift and showcases a swift recovery afterwards, validating the algorithm's efficacy. Similar to the fist set of experiments, it is worth mentioning that we did not explicitly fine-tune the value of β.

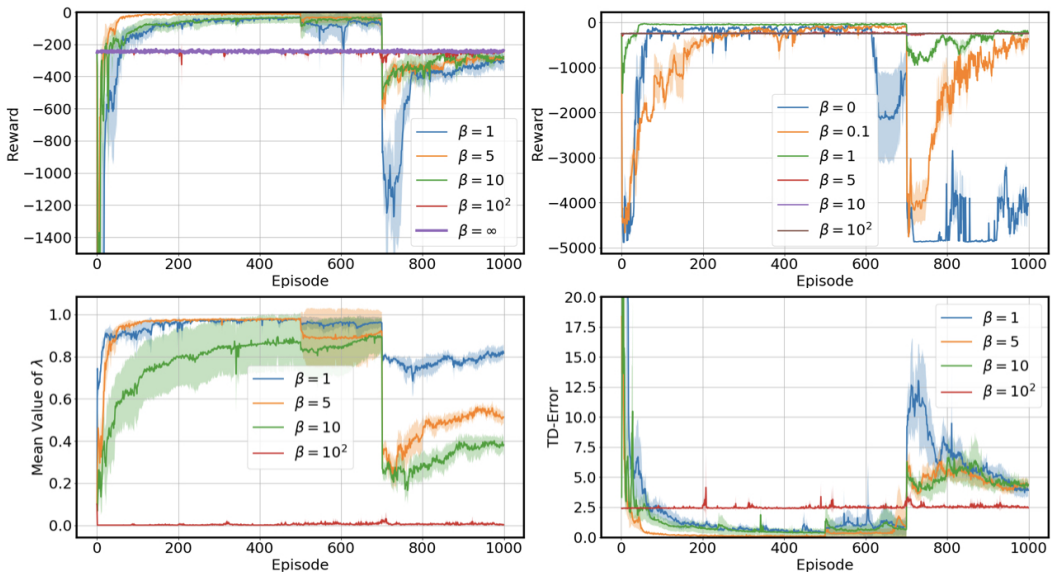

Figure 7: Stability against distribution shifts for different settings of the hyper-parameter β.

B Useful Lemmas

In this appendix, we present results that will be used when proving our main theorems.

B.1 Perturbation Lemma

We first prove the following perturbation lemma as a robustness guarantee, which holds for both the black-box (Section 4.1) and grey-box (Section 4.2) settings.

Lemma 2 (Perturbation Lemma)**.** *Under Assumption 1 and 2, the dynamic regret of Algorithm* (1) *(denoted by* PROP*) can be bounded by* $\mathsf{DR}(\mathsf{PROP}) \leq \mathcal{O}\left((\mathsf{ROB}-1)T\right) + L_{\mathrm{C}} C_s \sum_{t\in[T]} \mathbb{E}\left[(R_t)^p\right]^{1/p}$ *for constants* $L_{\mathrm{C}}, C_s > 0$.

Proof. Our proof consists of two parts. We first bound the Wasserstein distance between joint action-state distributions for the robust baseline and PROP. Next, we bound the dynamic regret.

Step 1. Wasserstein Distance between Joint Action-State Distributions. Denote by π and $\overline{\pi}$ the PROP policy (Algorithm 1) and an r-locally p-Wasserstein-robust policy respectively. For any step $t, t' \in [T]$ with $t' \leq t$, denote by $\overline{\rho}_{t|t'}$ the state-action distribution generated by applying the actions given by Algorithm 1 until step t' and applying the actions generated by the r-locally p-Wasserstein-robust policy afterwards until step t. Let ρ_t and $\overline{\rho}_t$ be the state-action distributions corresponding to Algorithm 1 and the r-locally p-Wasserstein robust baseline at each step $t \in [T]$. Using the triangle inequality, the state-action distribution difference between ρ_t and $\overline{\rho}_t$ in terms of the Wasserstein p-distance satisfies:

$$W\left(\rho_t, \overline{\rho}_t\right) \leq \sum_{\tau=0}^{t-1} W\left(\overline{\rho}_{t|t-\tau}, \overline{\rho}_{t|t-\tau-1}\right) \tag{15}$$

with $\overline{\rho}_{t|0} \coloneqq \overline{\rho}_t$ and $\overline{\rho}_{t|t} \coloneqq \rho_t$. Abuse the notation π and denote by $\overline{\pi}_{t-\tau:t}$ an operator on state-action distributions for applying the control baselines $\overline{\pi}_{t-\tau}, \overline{\pi}_{t-\tau+1}, \ldots, \overline{\pi}_t$ consecutively. Continuing from (15), it follows that for any $t \in [T]$ and $0 \leq \tau < t$,

$$\begin{aligned} W\left(\overline{\rho}_{t|t-\tau}, \overline{\rho}_{t|t-\tau-1}\right) &= W\left(\overline{\pi}_{t-\tau:t}\left(\overline{\rho}_{t-\tau|t-\tau}\right), \overline{\pi}_{t-\tau:t}\left(\overline{\rho}_{t-\tau|t-\tau-1}\right)\right) \\ &\leq s(\tau) W\left(\overline{\rho}_{t-\tau|t-\tau}, \overline{\rho}_{t-\tau|t-\tau-1}\right) \end{aligned} \tag{16}$$

where $\overline{\rho}_{t-\tau|t-\tau} \coloneqq \rho_{t-\tau}$ and in (16) we have used the assumption of the r-locally p-Wasserstein-robust policy (Definition 3). The assumption can be applied because by definition, for all $t \in [T]$, the Wasserstein p-distance between $\overline{\rho}_{t-\tau|t-\tau}$ and $\overline{\rho}_{t-\tau|t-\tau-1}$ can be bounded by

$$
\begin{aligned}
& W\left(\overline{\rho}_{t-\tau|t-\tau}, \overline{\rho}_{t-\tau|t-\tau-1}\right) \\
\leq & \mathbb{E}\left[\left\|\left(x_{t-\tau-1}, \pi_{t-\tau}\left(x_{t-\tau-1}\right)\right)-\left(x_{t-\tau-1}, \overline{\pi}_{t-\tau}\left(x_{t-\tau-1}\right)\right)\right\|_{\mathcal{X}\times\mathcal{U}}^{p}\right]^{1/p} && (17)\\
= & \mathbb{E}\left[\left\|\pi_{t-\tau}\left(x_{t-\tau-1}\right)-\overline{\pi}_{t-\tau}\left(x_{t-\tau-1}\right)\right\|_{\mathcal{U}}^{p}\right]^{1/p} && (18)\\
\leq & \mathbb{E}\left[\left(R_{t-\tau}\right)^{p}\right]^{1/p} \leq r, && (19)
\end{aligned}
$$

where in (17) we have used the definition

$$
W_p(\mu, \nu) \coloneqq \left(\inf_{J \in \mathcal{J}(\mu,\nu)} \int \|(x,u)-(x',u')\|_{\mathcal{X}\times\mathcal{U}}^{p} \mathrm{d}J\left((x,u),(x',u')\right)\right)^{1/p}.
$$

Since $\|(x,u)-(x',u')\|_{\mathcal{X}\times\mathcal{U}} \coloneqq \|x-x'\|_{\mathcal{X}} + \|u-u'\|_{\mathcal{U}}$, (18) follows. Finally, we obtain (19) considering the projection constraint in Algorithm 1. Combining (19) with (15),

$$
W_p\left(\rho_t, \overline{\rho}_t\right) \leq \sum_{\tau=0}^{t-1} s(\tau) \mathbb{E}_{P,\pi}\left[\left(R_{t-\tau}\right)^{p}\right]^{1/p}. \tag{20}
$$

Step 2. Dynamic Regret Analysis. Since the cost functions $(c_t : t \in [T])$ are Lipschitz continuous with a Lipschitz constant L_{C}, using the Kantorovich-Rubinstein duality theorem [70], since $W_p(\mu,\nu) \leq W_q(\mu,\nu)$ for all $1 \leq p \leq q < \infty$, for all $t \in [T]$,

$$
\begin{aligned}
& \mathbb{E}_{(x,u)\sim\rho_t}\left[c_t\left(x,u\right)\right] - \mathbb{E}_{(x,u)\sim\overline{\rho}_t}\left[c_t\left(x,u\right)\right] \\
\leq & \sup_{\|f\|_L \leq L_{\mathrm{C}}} \mathbb{E}_{(x,u)\sim\rho_t}\left[f\left(x,u\right)\right] - \mathbb{E}_{(x,u)\sim\overline{\rho}_t}\left[f\left(x,u\right)\right] \leq L_{\mathrm{C}} W_p\left(\rho_t, \overline{\rho}_t\right), && (21)
\end{aligned}
$$

where $\|\cdot\|_L$ denotes the Lipschitz semi-norm and the supremum is over all Lipschitz continuous functions f with a Lipschitz constant L_{C}. Therefore, the difference between the expected cost of Algorithm 1, denoted by π, and the baseline policy $\overline{\pi}$ satisfies

$$
\begin{aligned}
J(\pi) - J(\overline{\pi}) &= \sum_{t\in[T]} \mathbb{E}_{(x,u)\sim\rho_t}\left[c_t\left(x,u\right)\right] - \mathbb{E}_{(x,x)\sim\overline{\rho}_t}\left[c_t\left(x,u\right)\right] \\
&\leq L_{\mathrm{C}} \sum_{t\in[T]} \sum_{\tau=0}^{t-1} s(\tau) \mathbb{E}_{P,\pi}\left[\left(R_{t-\tau}\right)^{p}\right]^{1/p} \\
&\leq L_{\mathrm{C}} C_s \sum_{t\in[T]} \mathbb{E}_{P,\pi}\left[\left(R_t\right)^{p}\right]^{1/p} && (22)
\end{aligned}
$$

where we have used the assumption of the r-locally robustness policy so that $\sum_{t\in[T]} s(t) \leq C_s$ for some constant $C_s > 0$. Moreover, since the robust baseline $\overline{\pi}$ has a ratio of expectations bound such that $\frac{J(\overline{\pi})}{J^\star} \leq \mathsf{ROB}$. Using Assumption 1, from (22), we obtain

$$
\mathsf{DR}(\mathrm{PROP}) \coloneqq J(\pi) - J^\star \leq \mathcal{O}\left(\left(\mathsf{ROB}-1\right)T\right) + L_{\mathrm{C}} C_s \sum_{t\in[T]} \mathbb{E}_{P,\pi}\left[\left(R_t\right)^{p}\right]^{1/p}.
$$

□

B.2 Projection Lemma

The following lemma implies a useful consistency bound. It is worth noting that the lemma also holds if PROP adopts an alternative approach instead of projecting the actions as shown in (5):

$$
u_t \in \arg\min_{v\in\mathsf{U}} \widetilde{Q}_t\left(x_t, v\right) \text{ subject to } \|\overline{u}_t - v\| \leq R_t
$$

Implementing the projection rule in PROP can significantly reduce computational complexity, particularly when dealing with non-convex Q-advice.

Lemma 3 (Projection Lemma)**.** *Under Assumption 3, the actions and states* (x_t, u_t) *at* $t \in [T]$ *and* $t \in [T]$ *generated by* PROP *(Algorithm 1) satisfy*

$$Q_t^\star(x_t, u_t) - \inf_{v \in \mathcal{U}} Q_t^\star(x_t, v) \leq L_Q \left([\eta_t(x_t) - R_t]^+ \right) + \mu_t(x_t, u_t) \tag{23}$$

where $\eta_t(x) \coloneqq \|\widetilde{\pi}_t(x) - \overline{\pi}_t(x)\|_{\mathcal{U}}$*, and* $Q^\star$ *denotes the optimal Q-value functions satisfying the Bellman optimality equations in* (3)*.*

Proof. Let $\left(\widetilde{Q}_t : t \in [T], t \in [T]\right)$ be the Q-value advice used in Algorithm 1, denoted by π. Since $\zeta_t^Q \coloneqq \widetilde{Q}_t(x_t, u_t) - Q_t^\star(x_t, u_t)$, we have, for any $t \in [T]$ and (x_t, u_t) generated by Algorithm 1,

$$Q_t^\star(x_t, u_t) - \inf_{v \in \mathcal{U}} Q_t^\star(x_t, v) = \widetilde{Q}_t(x_t, u_t) - Q_t^\star(x_t, u_t^\star) - \zeta_t^Q. \tag{24}$$

Let $\widetilde{\pi}_t(x_t) \coloneqq \inf_{u \in \mathcal{U}} \widetilde{Q}_t(x_t, u)$. Note that since $\mathcal{U}$ is convex and compact, the projection step in the PROP policy is equivalent to (we choose the u_t on the line formed by $\widetilde{\pi}_t(x_t)$ and $\overline{\pi}_t(x_t)$ if there are ties in the projection solution)

$$u_t = \lambda(R_t)\, \widetilde{\pi}_t(x_t) + (1 - \lambda(R_t))\, \overline{\pi}_t(x_t),$$

where $\lambda(R_t) \coloneqq \min\{1, R_t / \|\widetilde{\pi}_t(x_t) - \overline{\pi}_t(x_t)\|_{\mathcal{U}}\}$. Write $\eta_t(\cdot) \coloneqq \|\widetilde{\pi}_t(\cdot) - \overline{\pi}_t(\cdot)\|_{\mathcal{U}}$. Therefore, the Q-value advice satisfies

$$\begin{aligned}
\widetilde{Q}_t(x_t, u_t) \leq & \widetilde{Q}_t(x_t, \widetilde{\pi}_t(x_t)) + L_Q \|\widetilde{\pi}_t(x_t) - u_t\| \\
\leq & \widetilde{Q}_t(x_t, \widetilde{\pi}_t(x_t)) + L_Q (1 - \lambda(R_t))\, \eta_t(x_t) \\
= & \widetilde{Q}_t(x_t, \widetilde{\pi}_t(x_t)) + L_Q \left(1 - \min\left\{1, \frac{R_t}{\eta_t(x_t)}\right\}\right) \eta_t(x_t) \\
\leq & \widetilde{Q}_t(x_t, \widetilde{\pi}_t(x_t)) + L_Q (\eta_t(x_t) - R_t),
\end{aligned} \tag{25}$$

where (25) follows since by construction $R_t \leq \eta_t$ for all $t \in [T]$. Let $\widetilde{u}_t$ denote $\widetilde{\pi}_t(x_t)$. Continuing from (24),

$$\begin{aligned}
& Q_t^\star(x_t, u_t) - \inf_{v \in \mathcal{U}} Q_t^\star(x_t, v) \\
= & \widetilde{Q}_t(x_t, u_t) - Q_t^\star(x_t, u_t^\star) - \zeta_t^Q \\
\leq & \widetilde{Q}_t(x_t, \widetilde{u}_t) - Q_t^\star(x_t, u_t^\star) + L_Q [\eta_t(x_t) - R_t]^+ - \zeta_t^Q && (26) \\
\leq & L_Q [\eta_t(x_t) - R_t]^+ + \mu_t && (27)
\end{aligned}$$

where (26) follows from (25). Since $\zeta_t^V \coloneqq \widetilde{Q}_t(x_t, \widetilde{u}_t) - Q_t^\star(x_t, u_t^\star)$, $\widetilde{u}_t$ minimizes $\widetilde{Q}_t$, and $\mu_t \coloneqq \zeta_t^V - \zeta_t^Q$, we obtain (27). □

B.3 Analysis of Approximate TD-Error

The following result that rewrites the approximate TD-error (c.f. (6)) is useful.

Lemma 4. *Consider the approximate TD-error in* (6) *such that*

$$\delta_t(u_{t-1}, x_{t-1}, x_t) \coloneqq c_{t-1}(x_{t-1}, u_{t-1}) + \inf_{v \in \mathcal{U}} \widetilde{Q}_t(x_t, v) - \widetilde{Q}_{t-1}(x_{t-1}, u_{t-1}).$$

It follows that for any $t \in [T]$*,*

$$\mathbb{E}_{P,\pi}[\delta_t(u_{t-1}, x_{t-1}, x_t)] = \mathbb{E}_{P,\pi}\left[\zeta_t^V(x_t) - \zeta_{t-1}^Q(x_{t-1}, u_{t-1})\right],$$

where $\zeta_{-1}^Q = 0$*,* ζ_t^Q *and* ζ_t^V *are defined as*

$$\begin{aligned}
\zeta_t^Q(x_t, u_t) &\coloneqq \widetilde{Q}_t(x_t, u_t) - Q_t^\star(x_t, u_t), \\
\zeta_t^V(x_t) &\coloneqq \inf_{v \in \mathcal{U}} \widetilde{Q}_t(x_t, v) - \inf_{v \in \mathcal{U}} Q_t^\star(x_t, v).
\end{aligned}$$

Proof. Taken expectation with randomness over the action and state trajectories, for any $t \in [T]$,

$$\begin{aligned}
&\mathbb{E}_{P,\pi}\left[\delta_t\left(u_{t-1}, x_{t-1}, x_t\right)\right] \\
\coloneqq &\mathbb{E}_{P,\pi}\left[c_{t-1}\left(x_{t-1}, u_{t-1}\right)\right] + \mathbb{E}_{P,\pi}\left[\inf_{v\in\mathcal{U}} \widetilde{Q}_t\left(x_t, v\right)\right] - \mathbb{E}_{P,\pi}\left[\widetilde{Q}_{t-1}\left(x_{t-1}, u_{t-1}\right)\right] \\
= &\mathbb{E}_{P,\pi}\left[\inf_{v\in\mathcal{U}} \widetilde{Q}_t(x_t, v) - \mathbb{P}_{t-1} V_t^\star(x_{t-1}, u_{t-1})\right] + \mathbb{E}_{P,\pi}\left[Q_{t-1}^\star(x_{t-1}, u_{t-1}) - \widetilde{Q}_{t-1}(x_{t-1}, u_{t-1})\right] \qquad (28) \\
= &\mathbb{E}_{P,\pi}\left[\inf_{v\in\mathcal{U}} \widetilde{Q}_t(x_t, v) - \mathbb{E}_{x'\sim P_t(\cdot|x_{t-1},u_{t-1})} \inf_{v\in\mathcal{U}} Q_t^\star(x', v)\right] \\
&\quad + \mathbb{E}_{P,\pi}\left[Q_{t-1}^\star(x_{t-1}, u_{t-1}) - \widetilde{Q}_{t-1}(x_{t-1}, u_{t-1})\right] \\
= &\mathbb{E}_{P,\pi}\left[\zeta_t^V(x_t) - \zeta_{t-1}^Q(x_{t-1}, u_{t-1})\right]
\end{aligned}$$

where we have used the Bellman optimality equations (3) to derive (28).

□

Next, we present our analysis of the black-box setting by proving Theorem 5.1 and Theorem 5.3.

C Black-Box Consistency and Robustness Analysis

C.1 Proof of Theorem 5.1

Consider an MDP model with Assumption 1,2, and 3. We prove the theorem below.

Theorem C.1. *Suppose the machine-learned policy $\widetilde{\pi}$ is (∞, ε)-consistent. The expected dynamic regret of* PROP *with the* BLACK-BOX *Procedure is bounded by*

$$\mathsf{DR}(\mathsf{PROP}) \leq \min\left\{\mathcal{O}(\varepsilon) + \mathcal{O}((1-\lambda)\gamma T), \mathcal{O}\left((\mathsf{ROB} + \lambda\gamma - 1)\, T\right)\right\}$$

where ε is defined in (4)*, γ is the diameter of the action space $\mathcal{U}$, T is the length of the time horizon,* ROB *is the ratio of expectations of the robust baseline $\overline{\pi}$, and $0 \leq \lambda \leq 1$ is a hyper-parameter.*

Consistency Analysis. To show the first bound in Theorem 5.1 regarding the consistency result, we consider the following steps. For any $t \in [T]$, denote by (x_t, u_t) the corresponding state and action generated by the projection pursuit policy PROP, denoted by π. The Bellman optimality equations (3) imply:

$$Q_t^\star\left(x_t, u_t\right) = c_t\left(x_t, u_t\right) + \mathbb{E}_P\left[\inf_{v\in\mathcal{U}} Q_{t+1}^\star\left(x_{h+1}, v\right) \middle| x_t, u_t\right]. \qquad (29)$$

Therefore the dynamic regret of the projection pursuit policy π can be rewritten as

$$\mathsf{DR}(\mathsf{PROP}) = J(\pi) - J^\star = \left(\mathbb{E}_{P,\pi}\left[\sum_{t=0}^{T-1} c_t\left(x_t, u_t\right)\right] - \inf_{v\in\mathcal{U}} Q_{t,0}^\star\left(x_0, v\right)\right). \qquad (30)$$

Combining (30) with (29), we obtain the following cost-difference bound:

$$\begin{aligned}
&\mathsf{DR}(\mathsf{PROP}) \\
= &\sum_{t=0}^{T-1}\left(\mathbb{E}_{P,\pi}\left[Q_t^\star\left(x_t, u_t\right)\right] - \mathbb{E}_P\left[\inf_{v\in\mathcal{U}} Q_{t+1}^\star\left(x_{t+1}, v\right)\right]\right) - \inf_{v\in\mathcal{U}} Q_{t,0}^\star\left(x_0, v\right) \\
= &Q_{t,0}^\star\left(x_0, u_0\right) - \inf_{v\in\mathcal{U}} Q_{t,0}^\star\left(x_0, v\right) + \sum_{t=1}^{T-1}\left(\mathbb{E}_{P,\pi}\left[Q_t^\star\left(x_t, u_t\right)\right] - \mathbb{E}_P\left[\inf_{v\in\mathcal{U}} Q_t^\star\left(x_t, v\right)\right]\right) \\
= &\sum_{t=0}^{T-1} \mathbb{E}_{P,\pi}\left[Q_t^\star\left(x_t, u_t\right) - \inf_{v\in\mathcal{U}} Q_t^\star\left(x_t, v\right)\right].
\end{aligned}$$

Recall that for the BLACK-BOX Procedure in Section 4.1, the robustness budget is set as $R_t = \lambda\eta_t$ for all $t \in [T]$. Applying the bound in Lemma 3 gives the following consistency bound:

$$\mathsf{DR}(\text{PROP}) = \mathcal{O}\left(\sum_{t=0}^{T-1}\left([\eta_t(x_t) - R_t]^+\right) + \mathbb{E}_{P,\pi}\left[\mu_t\right]\right) \le \mathcal{O}\left(\sum_{t=0}^{T-1}\mathbb{E}_{P,\pi}\left[\mu_t\right]\right) + \mathcal{O}\left((1-\lambda)\gamma T\right)$$

since $\eta_t(x_t) \le \gamma$ for all $t \in [T]$, and

$$\begin{aligned}\mu_t = \zeta_t^V - \zeta_t^Q =& \widetilde{Q}_t(x_t, \widetilde{u}_t) - Q_t^\star(x_t, u_t^\star) - \left(\widetilde{Q}_t(x_t, u_t) - Q_t^\star(x_t, u_t)\right)\\ \le& \left\| \inf_{v\in\mathcal{U}} \widetilde{Q}_t(x_t, v) - \inf_{v\in\mathcal{U}} Q_t^\star(x_t, v)\right\|_\infty + \left\|\widetilde{Q}_t(x_t, u_t) - Q_t^\star(x_t, u_t)\right\|_\infty.\end{aligned}$$

Noting that the machine-learned policy $\widetilde{\pi}$ is (∞, ε)-consistent, we obtain $\sum_{t=0}^{T-1}\mathbb{E}_{P,\pi}\left[\mu_t\right] \le \varepsilon$. Hence,

$$\mathsf{DR}(\text{PROP}) = \mathcal{O}\left(\varepsilon\right) + \mathcal{O}\left((1-\lambda)\gamma T\right). \tag{31}$$

Robustness Analysis. Note that for any $t \in [T]$, $\eta_t(x_t) \le \gamma$, where γ is the diameter of the compact action space $\mathcal{U}$. Hence, noting the black-box setting of the robustness budget $R_t = \lambda\eta_t$ for all $t \in [T]$ and applying Lemma 2, the sum of expected discrepancies, over all t can be bounded by

$$\begin{aligned}\mathsf{DR}(\text{PROP}) \le& \mathcal{O}\left((\mathsf{ROB}-1)T\right) + L_{\mathrm{C}} C_s \sum_{t\in[T]} \mathbb{E}_{P,\pi}\left[(\lambda\gamma)^p\right]^{1/p}\\ \le& \mathcal{O}\left((\mathsf{ROB} + \lambda\gamma - 1)T\right).\end{aligned} \tag{32}$$

Combining (31) and (32), we complete the proof.

C.2 Proof of Theorem 5.2

Let $\mathcal{MDP}$ be the set of all MDP models $\mathsf{MDP}(\mathcal{X}, \mathcal{U}, T, P, c)$ satisfying Assumption 1,2, and 3. To prove Theorem 5.2, noting that by the definitions of consistency and robustness, we apply Theorem 5.1 to derive a bound on the worst-case ratio of expectations:

$$\sup_{\mathcal{MDP}} \mathsf{RoE}(\varepsilon) \le 1 + \sup_{\mathcal{MDP}} \frac{\mathsf{DR}(\text{PROP})}{J^\star} \le \min\left\{1 + \mathcal{O}\left(\frac{\varepsilon}{T}\right) + \mathcal{O}((1-\lambda)\gamma), \mathsf{ROB} + \mathcal{O}(\lambda\gamma)\right\},$$

which implies that PROP with the BLACK-BOX Procedure is $(1 + \mathcal{O}((1-\lambda)\gamma)$-consistent and $(\mathsf{ROB} + \mathcal{O}(\lambda\gamma))$-robust.

C.3 Proof of Theorem 5.3

Proof. According to Lemma 3, the expected dynamic regret of PROP satisfies

$$\mathsf{DR}(\text{PROP}) = \sum_{t\in[T]} \mathbb{E}_{P,\pi}\left[Q_t^\star\left(x_t, u_t\right) - \inf_{v\in\mathcal{U}} Q_t^\star\left(x_t, v\right)\right],$$

where π denotes PROP. For notational simplicity, we introduce the following notation:

$$\begin{aligned}\Delta Q_t^\star(P,\pi) \coloneqq& \mathbb{E}_{P,\pi}\left[Q_t^\star(x_t, u_t) - \inf_{v\in\mathcal{U}} Q_t^\star(x_t, v)\right],\\ \Delta \widetilde{Q}_t(P,\pi) \coloneqq& \mathbb{E}_{P,\pi}\left[\widetilde{Q}_t(x_t, u_t) - \inf_{v\in\mathcal{U}} \widetilde{Q}_t(x_t, v)\right].\end{aligned}$$

With the BLACK-BOX Procedure, we set $R_t = \lambda\eta_t$ with some hyper-parameter $0 \le \lambda \le 1$. Therefore, there exists Lipschitz continuous Q-value predictions $(\widetilde{Q}_t : t \in [T])$ with a Lipschitz constant L_Q such that

$$\mathsf{DR}\left(\text{PROP}(\text{BLACK-BOX})\right) \ge \sum_{t\in[T]}\left(\Delta Q_t^\star(P,\pi) - \Delta\widetilde{Q}_t(P,\pi) + (1-\lambda)L_Q\gamma\right). \tag{33}$$

First, we verify that Wasserstein robust policies exist since we can construct a transition probability P such that the states in different times are independent. Denote by OPT the expected optimal total cost. We can construct cost functions $(c_t : t \in [T])$ that are Lipschitz continuous with a Lipschitz constant L_c and Q-advice $(\widetilde{Q}_t : t \in [T])$ satisfying

$$\frac{\sum_{t\in[T]} \left(\Delta Q_t^\star(P,\pi) - \Delta \widetilde{Q}_t(P,\pi)\right)}{\mathsf{OPT}} \geq \Omega\left(\mathsf{ROB} + \frac{\lambda\gamma L_c}{\mathsf{OPT}} T\right).$$

Note that the corresponding Q-value predictions satisfy Assumption 3. Let $\mathcal{MDP}$ be the set of all MDP models $\mathsf{MDP}(\mathcal{X}, \mathcal{U}, T, P, c)$ satisfying Assumption 1,2, and 3. Combining above with (33), and noting that in Assumption 1, $c_t(x,u) > 0$ for all $t \in [T]$, $x \in \mathcal{X}$, and $u \in \mathcal{U}$, for any $\varepsilon \geq 0$, the ratio of expectations can be bounded by

$$\mathsf{RoE}(\text{PROP}) = 1 + \sup_{\mathcal{MDP}} \frac{\mathsf{DR}(\text{PROP})}{\mathsf{OPT}} = 1 + \Omega\Big((1-\lambda)L_Q\gamma + \min\{\varepsilon, \lambda\gamma L_c + \mathsf{ROB}\}\Big),$$

which implies that PROP cannot be both $(1 + o(\lambda\gamma))$-consistent and $(\mathsf{ROB} + o((1-\lambda)\gamma))$-robust for any $0 \leq \lambda \leq 1$. □

D Grey-Box Consistency and Robustness Analysis

In the following, we present a dynamic regret bound for the grey-box setting (Section 4.2) that is analogous to the one presented in Theorem 5.1 for the black-box scenario.

First, in addition to Definition 4, we further recall the following quantities used in Lemma 4 for notational convenience:

$$\zeta_t^Q(x_t, u_t) \coloneqq \widetilde{Q}_t(x_t, u_t) - Q_t^\star(x_t, u_t), \tag{34}$$

$$\zeta_t^V(x_t) \coloneqq \inf_{v\in\mathcal{U}} \widetilde{Q}_t(x_t, v) - \inf_{v\in\mathcal{U}} Q_t^\star(x_t, v), \tag{35}$$

where by definition, ζ_t^Q and ζ_t^V depend on the random trajectory $((x_t, u_t) : t \in [T])$. Denote $\mu_t \coloneqq \zeta_t^V - \zeta_t^Q$. Note that when the environment is stationary, under some model assumptions and with a Reproducing kernel Hilbert space (RKHS) being the function class, the optimism lemma (Lemma 5.2) in [6] shows that with probability at least $1 - (2T^2H^2)^{-1}$, the generated Q-value functions satisfy $\sum_{(h,t)\in[H]\times[T]} \mathbb{E}_{P,\pi}[\delta_{h,t} + \mu_{h,t}] = \widetilde{O}(H\Gamma_K(T,\lambda)\sqrt{T})$ where H is the number of episodes and $\widetilde{O}(\cdot)$ omits logarithmic terms and $T\Gamma_K(T,\lambda)$ is the maximal information gain [71] that characterizes the intrinsic complexity of the function class.

Denote the by φ_t the per-step cost difference between the robust baseline and the optimal policy at time $t \in [T]$ such that $\sum_{t\in[T]} \varphi_t = \Theta((\mathsf{ROB} - 1)T)$. Suppose the robust baseline $\overline{\pi}$ is γ-locally p-Wasserstein-robust. The following theorem presents a preliminary result that will be used to prove Theorem 5.4.

Theorem D.1 (Grey-Box: Dynamic Regret)**.** *Consider any MDP model satisfying Assumption 1,2, and 3. The expected dynamic regret of* PROP *(Algorithm 1) with the* GREY-BOX *Procedure satisfies the following bound:*

$$\mathsf{DR}(\text{PROP}) \leq \sum_{t\in[T]} \min\left\{ \underbrace{\mathbb{E}_{P,\pi}[\mu_t] + L_Q\mathbb{E}_{P,\pi}(\eta_t(x_t) - R_t)}_{\textit{Consistency Bound (Lemma 3)}}, \underbrace{\varphi_t + L_C C_s \mathbb{E}_{P,\pi}[(R_t)^p]^{1/p}}_{\textit{Robustness Bound (Lemma 2)}} \right\}, \tag{36}$$

where L_Q and L_C are Lipschitz constants, $\mu_t \coloneqq \zeta_t^V - \zeta_t^Q$, and $\eta_t(x) \coloneqq \|\widetilde{\pi}_t(x) - \overline{\pi}_t(x)\|_{\mathcal{U}}$.

D.1 Proof of Theorem D.1

A central step in the proof of Theorem D.1 is to combine the per-step analysis in the consistency and robustness results in Lemma 2 and 3 and apply the selection of budgets in (7) (see Section 4.2). Combining (22) and (23) and summing over all t, (36) follows.

D.2 Proof of Theorem 5.4

Consistency Analysis. We first show that the PROP policy with the GREY-BOX procedure is 1-consistent. Let $\varepsilon(p,\rho)=0$ for some $p\in[0,\infty]$ and let ρ be the trajectory distribution generated by PROP (defined in (4)). Then we must have

$$\mathbb{E}_{P,\pi}\left[\zeta_t^Q\right]=\left[\widetilde{Q}_t\left(x_t,u_t\right)-Q_t^\star\left(x_t,u_t\right)\right]=0,$$
$$\mathbb{E}_{P,\pi}\left[\zeta_t^V\right]=\mathbb{E}_{P,\pi}\left[\inf_{v\in\mathcal{U}}\widetilde{Q}_t\left(x_t,v\right)-\inf_{v\in\mathcal{U}}Q_t^\star\left(x_t,v\right)\right]=0,$$

for any $t\in[T]$. Consider the expectation of the TD-error (with randomness taken over the action and state trajectories). From Lemma 4 we know $\mathbb{E}_{P,\pi}[\delta_t\left(u_{t-1},x_{t-1},x_t\right)]=\mathbb{E}_{P,\pi}[\zeta_t^V(x_t)-\zeta_{t-1}^Q(x_{t-1},u_{t-1})]$. This implies that the TD-error δ_t must satisfy

$$\mathbb{E}_{P,\pi}\left[\delta_t\right]=\mathbb{E}_{P,\pi}[\zeta_t^V-\zeta_{t-1}^Q]=0. \tag{37}$$

Similarly, we have

$$\mathbb{E}_{P,\pi}\left[\mu_t\right]=\mathbb{E}_{P,\pi}\left[\zeta_t^V-\zeta_t^Q\right]=0.$$

Therefore, by the construction of the robustness budget R_t in (7) for the GREY-BOX Procedure,

$$\mathbb{E}_{P,\pi}\left[\eta_t-R_t\right]\le\mathbb{E}_{P,\pi}\left[\frac{\beta}{L_Q}\sum_{s=1}^{t}\delta_s\right]=0.$$

Applying Theorem D.1, we get when $\varepsilon=0$ (i.e., the machine-learned policy $\overline{\pi}$ is optimal), then $\mathsf{RoE}(0)=1$, implying that PROP is 1-consistent.

Robustness Analysis. By Lemma 4, we get for any trajectory $((x_t,u_t):t\in[T])$ and $t\in[T]$, $\mu_t-\delta_t=\zeta_{t-1}^Q-\zeta_t^Q$. Therefore, denoting by $\zeta_{-1}^Q=0$,

$$\sum_{s=0}^{t}\left(\mu_s-\delta_s\right)=\sum_{s=0}^{t}\left(\zeta_{s-1}^Q-\zeta_s^Q\right)=\zeta_t^Q.$$

According to Assumption 3, there exist $\Delta=o(T)$ such that $|\zeta_t^Q|\le\Delta$ for all $t\in[T]$. We consider two cases. First, consider an event $\sum_{t\in[T]}\mu_t\le\Delta$. Let $\mathcal{MDP}$ be the set of all MDP models $\mathsf{MDP}(\mathcal{X},\mathcal{U},T,P,c)$ satisfying Assumption 1,2, and 3. Then, applying Theorem D.1,we derive a bound on the worst-case ratio of expectations:

$$\sup_{\varepsilon\ge0}\sup_{\mathcal{MDP}}\mathsf{RoE}(\varepsilon)\le1+\sup_{\varepsilon\ge0}\sup_{\mathcal{MDP}}\frac{\mathsf{DR(PROP)}}{J^\star}\le1+\mathcal{O}\left(\frac{\Delta}{T}\right)=1+o(1)\le\mathsf{ROB}+o(1).$$

Now, if $\sum_{t\in[T]}\mu_t>\Delta$, then we must have $\sum_{t\in[T]}\delta_t>0$. Since the action space is compact, $\eta_t\le\gamma$, which is bounded. There exists some hyper-parameter $\beta>0$ such that $R_t=0$. Therefore, the PROP with the GREY-BOX Procedure will be switched to the robust baseline $\overline{\pi}$. Without loss of generality, assume $0\le m<T$ is the largest time index such that $\sum_{s=1}^{m}\mu_s>\Delta$ and $\sum_{s=1}^{m-1}\mu_s\le\Delta$. Applying Theorem D.1, we have

$$\mathsf{DR(PROP)}\le\sum_{s=1}^{m-1}\mu_s+\sum_{s=m}^{T-1}\varphi_s+\mathcal{O}(\eta_m)\le\sum_{s=1}^{m-1}\mu_s+\sum_{s=m}^{T-1}\varphi_s+\mathcal{O}(\gamma),$$

implying that

$$\sup_{\varepsilon\ge0}\sup_{\mathcal{MDP}}\mathsf{RoE}(\varepsilon)\le\mathsf{ROB}+\mathcal{O}\left(\frac{\Delta+\gamma}{T}\right)\le\mathsf{ROB}+o(1).$$

E Proof of Theorem A.1

To show Theorem A.1, we first show a technical lemma with respect to MPC_T, which plans until the end of the episode from the first time step.

Lemma 5. *Suppose Assumptions 4 and 5 hold. For each step $t \in [T]$, the control policy of MPC_T can be rewritten as $u_t = \overline{K}_t x_t$, for some matrices $\left(\overline{K}_t : t \in [T]\right)$ satisfy that $\left\|\overline{K}_t\right\| \le C$ for all $t \in [T]$, and*

$$\left\|(A_{t'-1} + B_{t'-1}\overline{K}_{t'-1}) \cdots (A_t + B_t \overline{K}_t)\right\| \le C\lambda^{t'-t}, \forall t, t' \in [T],\ t' \ge t,$$

where λ, C are as defined in Theorem A.1.

Proof. To simplify the notation, we define

$$\Gamma_{t,t'} = \begin{cases} \mathsf{diag}\left(Q_t, R_t, \ldots, R_{t'-1}, Q_{t'}\right) & \text{if } t' = T-1 \\ \mathsf{diag}\left(Q_t, R_t, \ldots, R_{t'-1}, P_{t'}\right) & \text{otherwise} \end{cases}. \tag{38}$$

By the KKT conditions, we see that for any $t \in [T]$, the predictive optimal solution $\psi_{t,T-1}(x_t)$ is given by

$$\begin{pmatrix} x_{t|t} \\ u_{t|t} \\ \vdots \\ \hline x_{T-1|t} \\ \eta_{t|t} \\ \vdots \\ \eta_{T-1|t} \end{pmatrix} = \left(\begin{array}{c|c} \Gamma_{t,T-1} & (\Xi_{t,T-1})^\top \\ \hline \Xi_{t,T-1} & \end{array}\right)^{-1} \begin{pmatrix} 0 \\ \vdots \\ 0 \\ \hline x_t \\ 0 \\ \vdots \\ 0 \end{pmatrix}. \tag{39}$$

Therefore, $u_{t|t}$ is a linear function of x_t, and this relationship defines $\overline{K}_t$. Lemma G.2 in [19] implies that $\left\|\overline{K}_t\right\| \le C$. Note that the block matrix is invertible since $(Q_t : t \in [T])$ and $(R_t : t \in [T])$ are positive definite.

To simplify the notation, we define the state transition matrix

$$\Phi_{t,t'} \coloneqq (A_{t'-1} + B_{t'-1}\overline{K}_{t'-1}) \cdots (A_t + B_t \overline{K}_t).$$

Consider an arbitrary state x_t. Note that $\left(x_{t'|t} : t \le t' < T\right)$ is the optimal trajectory when there is no disturbance after step t. By the principle of optimality, we see that $\left(x_{t'|t} : t \le t' < T\right)$ is identical with the actual trajectory of MPC_T after step t. In other words, for arbitrary x_t, the multi-step transition matrix $\Phi_{t,t'}$ satisfies

$$x_{t'|t} = \Phi_{t,t'} x_t.$$

Lemma G.2 in [19] implies that $\|\Phi_{t,t'}\| \le C\lambda^{t'-t}$. □

Lemma 5 shows that MPC_T has the same effect as a time-varying linear feedback controller that is exponentially stable. We generalize this property to MPC_k with a smaller prediction horizon (Lemma 6) by showing that MPC_k behaves similar to MPC_T when k is sufficiently large.

Lemma 6. *Suppose Assumptions 4 and 5 hold. Let (C, λ) be the same as Lemma 5. For each step $t \in [T]$, the control policy of MPC_k can be rewritten as $u_t = K_t^k x_t$, for some matrices $\{K_t^k\}_{t\in[T]}$ satisfy that*

$$\left\|K_t^k\right\| \le C, \text{ and } \left\|K_t^k - \overline{K}_t\right\| \le C^2 a \cdot \lambda^{2k}.$$

Further, for any $\widehat{\lambda} > \lambda$, when $k \ge \min\{T, \frac{1}{2}\log\left(C^3 ba\lambda/(\widehat{\lambda} - \lambda)\right)/\log(1/\lambda)\}$, we have

$$\left\|(A_{t'-1} + B_{t'-1}K_{t'-1}^k) \cdots (A_t + B_t K_t^k)\right\| \le C\widehat{\lambda}^{t'-t}, \text{ for any } t, t' \in [T], t' \ge t.$$

Proof. Let $\bar{t} \coloneqq \min\{t+k, T-1\}$. By the KKT conditions, we see that for any $t \in [T]$, the predictive optimal solution $\psi_{t,\bar{t}}(x_t; P_{\bar{t}})$ is given by

$$\begin{pmatrix} x_{t|t} \\ u_{t|t} \\ \vdots \\ \hline x_{\bar{t}|t} \\ \eta_{t|t} \\ \vdots \\ \eta_{\bar{t}|t} \end{pmatrix} = \left(\begin{array}{c|c} \Gamma_{t,\bar{t}} & (\Xi_{t,\bar{t}})^\top \\ \hline \Xi_{t,\bar{t}} & \end{array}\right)^{-1} \begin{pmatrix} 0 \\ \vdots \\ 0 \\ \hline x_t \\ 0 \\ \vdots \\ 0 \end{pmatrix}. \tag{40}$$

Therefore, $u_{t|t}$ is a linear function of x_t, and this relationship defines K_t^k. By Lemma G.2 of [19], we see that $\left\|K_t^k\right\| \le C$.

When $\bar{t} < T-1$, construct an auxiliary disturbance sequence $\widehat{w}_{t:T-2|t}$ with $\widehat{w}_{\bar{t}|t} \coloneqq -A_t \psi_{t,\bar{t}}(x_t)$ and $\widehat{w}_{t'|t} = 0$ for all $t' \neq \bar{t}$. We see that

$$\psi_{t,\bar{t}}(x_t)[u_{t|t}] = \psi_{t,T-1}(x_t, \widehat{w}_{t:T-2|t}; Q_{T-1})[u_{t|t}].$$

Therefore, we see that

$$\begin{aligned} &\left\|\psi_{t,\bar{t}}(x_t)[u_{t|t}] - \psi_{t,T-1}(x_t)[u_{t|t}]\right\| \\ =\ &\left\|\psi_{t,T-1}(x_t, \widehat{w}_{t:T-2|t}; Q_{T-1})[u_{t|t}] - \psi_{t,T-1}(x_t, 0_{\times(T-t-1)}; Q_{T-1})[u_{t|t}]\right\| \\ \le\ & C\lambda^k \left\|\widehat{w}_{\bar{t}|t}\right\| && \text{(41a)} \\ \le\ & C^2 a \cdot \lambda^{2k} \left\|x_t\right\|, && \text{(41b)} \end{aligned}$$

where we have applied the perturbation bounds in Lemma G.2 of [19] in (41a) and (41b). Since this inequality holds for any arbitrary x_t, we see that $\left\|K_t^k - \overline{K}_t\right\| \le C^2 a \cdot \lambda^{2k}$. To simplify the notation, we denote $\epsilon \coloneqq C^2 a \cdot \lambda^{2k}$.

We can derive the following bound in terms of the ℓ_2 norm:

$$\begin{aligned} &\left\|(A_{t'-1} + B_{t'-1}K_{t'-1}^k) \cdots (A_t + B_t K_t^k)\right\| \\ \le\ & \sum_{j=0}^{t'-t} \binom{t'-t}{j} C^{j+1} \lambda^{t-t'} (b\epsilon)^j && \text{(42a)} \\ =\ & C\lambda^{t'-t} \left(1 + Cb\epsilon\right)^{t'-t} \\ \le\ & C\widehat{\lambda}^{t'-t}, && \text{(42b)} \end{aligned}$$

where we use the decomposition that for any $t'' \in \{t, \ldots, t'-1\}$,

$$A_{t''} + B_{t''}K_{t''}^k \le (A_{t''} + B_{t''}\overline{K}_{t''}) + B_{t''}(K_{t''}^k - \overline{K}_{t''})$$

and $\left\|B_{t''}(K_{t''}^k - \overline{K}_{t''})\right\| \le b\epsilon$ in (42a). We also use Lemma 5 in (42a) and the assumption that

$$k \ge \frac{1}{2} \log\left(C^3 ba\lambda / (\widehat{\lambda} - \lambda)\right) / \log(1/\lambda)$$

in (42b). □

To establish a dynamic regret bound that depends on the offline optimal cost, we first need to show a lower bound of J^* that depends on the "power" of the unknown disturbances.

Lemma 7. *The offline optimal cost is lower bounded by*

$$J^* \ge \frac{\mu}{4(1+a^2+b^2)} \sum_{t=0}^{T-2} \left\|w_t\right\|^2.$$

Proof. Note that the dynamics of the LTV system can be rewritten as

$$x_{t+1} - A_t x_t - B_t u_t = w_t.$$

Taking norms on both sides of the equality gives

$$\begin{aligned} \|w_t\| &= \|x_{t+1} - A_t x_t - B_t u_t\| \\ &\le \|x_{t+1}\| + \|A_t x_t\| + \|B_t u_t\| \qquad (43a) \\ &\le \|x_{t+1}\| + a\|x_t\| + b\|u_t\|, \qquad (43b) \end{aligned}$$

where we have used the triangle inequality in (43a) and the definition of the induced matrix ℓ_2-norm in (43b). Taking the squares of both sides and applying the Cauchy-Schwartz inequality together imply

$$\|w_t\|^2 \le (\|x_{t+1}\| + a\|x_t\| + b\|u_t\|)^2 \le \frac{1+a^2+b^2}{\mu}\left(\mu\|x_{t+1}\|^2 + \mu\|x_t\|^2 + \mu\|u_t\|^2\right). \quad (44)$$

By (44) and the assumptions on Q_t and R_t, we obtain that

$$\begin{aligned} \frac{\mu}{2(1+a^2+b^2)} \cdot \sum_{t=0}^{T-1} \|w_t\|^2 &\le \frac{1}{2}\sum_{t=0}^{T-2}\left(\mu\|x_{t+1}\|^2 + \mu\|x_t\|^2 + \mu\|u_t\|^2\right) \\ &\le 2\sum_{t=0}^{T-1} c_t(x_t, u_t). \end{aligned}$$

Since the above inequality holds for any arbitrary trajectory $((x_t, u_t) : t \in [T])$, we conclude that Lemma 7 holds. □

Since the Wasserstein robustness (see Definition 3) is for distributions on the state-action space, we also prove a technical lemma below that helps convert the contraction on deterministic state/action pairs to the contraction on distributions. Let $W_1(\mu, \nu)$ denote the Wasserstein 1-distance between two distributions μ and ν.

Lemma 8. *Suppose $\varphi : \mathcal{Y} \to \mathcal{W}$ is a deterministic function that satisfies $\|\varphi(v) - \varphi(v')\|_{\mathcal{W}} \le \kappa\|v - v'\|_{\mathcal{Y}}$ for any $v, v' \in \mathcal{Y}$. Then, for any pair of distributions ρ and ρ' on $\mathcal{Y}$, we have $W_1(\varphi(\rho), \varphi(\rho')) \le \kappa W_1(\rho, \rho')$.*

Proof. Recall that $W_1(\rho, \rho') \coloneqq \inf_J \int \|v - v'\|_{\mathcal{Y}} \, dJ(v, v')$, where J is a joint distribution on $\mathcal{Y} \times \mathcal{Y}$ with marginals ρ and ρ'. We define a mapping $\Phi : \mathcal{Y} \times \mathcal{Y} \to \mathcal{Z} \times \mathcal{Z}$ as $\Phi(v, v') \coloneqq (\varphi(v), \varphi(v'))$. We see that ΦJ gives a joint distribution on $\mathcal{W} \times \mathcal{W}$ with marginals $\varphi(\rho)$ and $\varphi(\rho')$, and it satisfies

$$\int \|u - u\|_{\mathcal{W}} \, \mathrm{d}(\Phi J)(u, u') = \int \|\varphi(v) - \varphi(v')\|_{\mathcal{Y}} \, \mathrm{d}J(v, v') \le \varepsilon \int \|v - v'\|_{\mathcal{Y}} \, \mathrm{d}J(v, v').$$

Note that the above inequality holds for any J with marginals ρ and ρ'. Thus Lemma 8 holds. □

Now we are ready to show Theorem A.1.

For a state x at time step $t \in [T]$, let $x_{t:t'}(x)$ and $u_{t:t'}(x)$ denote the corresponding state and action of MPC at time step t'. By Lemma 6, we see that for any state-action pairs (x, u) and (x', u') at step t_1, we have

$$\begin{aligned} &\|(x_{t_1+1:t_2}, u_{t_1+1:t_2})(A_{t_1}x + B_{t_1}u + w_{t_1}) - (x_{t_1+1:t_2}, u_{t_1+1:t_2})(A_{t_1}x' + B_{t_1}u' + w_{t_1})\| \\ &\le (1+C)\|x_{t_1+1:t_2}(A_{t_1}x + B_{t_1}u + w_{t_1}) - x_{t_1+1:t_2}(A_{t_1}x' + B_{t_1}u' + w_{t_1})\| \qquad (45a) \\ &\le (1+C)C\widehat{\lambda}^{t_2-t_1-1}\|A_{t_1}(x - x') + B_{t_1}(u - u')\| \qquad (45b) \\ &\le (1+C)C(a+b)\widehat{\lambda}^{t_2-t_1-1}\|(x, u) - (x', u')\|, \qquad (45c) \end{aligned}$$

where we have used Lemma 6 in (45a) and (45b); Moreover, we have applied the assumption that $\|A_{t_1}\| \le a$ and $\|B_{t_1}\| \le b$ and the triangle inequality in (45c). Since (45) establishes a contraction for a deterministic state-action pair and the dynamics is deterministic, applying Lemma 8 finishes the proof of the first conclusion of Theorem A.1.

Using a similar decomposition technique with [19], by Lemma 6, we see that the trajectory $(x_t : t \in [T])$ of PROP satisfies that

$$\|x_t\| \le \sum_{t'=0}^{t-1} \left\|\Phi_{t',t}^k\right\| \cdot (\|w_{t'}\| + b\overline{R}) \le C \sum_{t'=0}^{t-1} \widehat{\lambda}^{t-t'}(d + b\overline{R}) \le \frac{C(d + b\overline{R})}{1-\widehat{\lambda}}, \tag{46}$$

where we denote $\Phi_{t',t}^k \coloneqq (A_{t'-1} + B_{t'-1}K_{t'-1}^k) \cdots (A_t + B_t K_t^k)$ and the assumption that PROP deviates at most $\overline{R}$ from MPC_k's action. We also see that

$$\|u_t\| \le \left\|K_t^k x_t\right\| + \overline{R} \le C\overline{R}_x + \overline{R}. \tag{47}$$

This finishes the proof of the second statement in Theorem A.1.

Let the trajectory of MPC_k when executed without the machine-learned advice be denoted by $(\overline{x}_t : t \in [T])$. We see that

$$\|\overline{x}_t\| = \left\| \sum_{t'=0}^{t-1} \Phi_{t',t}^k w_{t'} \right\| \le C \sum_{t'=0}^{t-1} \widehat{\lambda}^{t-t'} \|w_{t'}\|.$$

Applying the Cauchy-Schwarz inequality, we obtain

$$\begin{aligned}
\|\overline{x}_t\|^2 &\le \left(C \sum_{t'=0}^{t-1} \widehat{\lambda}^{t-t'} \|w_{t'}\| \right)^2 \le C^2 \left(\sum_{t'=0}^{t-1} \widehat{\lambda}^{t-t'} \right) \left(\sum_{t'=0}^{t-1} \widehat{\lambda}^{t-t'} \|w_{t'}\|^2 \right) \\
&\le \frac{C^2}{1-\widehat{\lambda}} \sum_{t'=0}^{t-1} \widehat{\lambda}^{t-t'} \|w_{t'}\|^2 .
\end{aligned} \tag{48}$$

For the control actions of MPC_k, we also see that

$$\|\overline{u}_t\|^2 = \left\|K_t^k \overline{x}_t\right\|^2 \le C^2 \|\overline{x}_t\|^2 \le \frac{C^4}{1-\widehat{\lambda}} \sum_{t'=0}^{t-1} \widehat{\lambda}^{t-t'} \|w_{t'}\|^2 . \tag{49}$$

Therefore, we get the following bound on the total cost:

$$\begin{aligned}
J(\mathsf{MPC}_k) &= \sum_{t=0}^{T-1} \left(\frac{1}{2} (\overline{x}_t)^\top Q_t \overline{x}_t + \frac{1}{2} (\overline{u}_t)^\top R_t \overline{u}_t \right) \\
&\le \frac{\ell}{2} \sum_{t=0}^{T-1} \left(\|\overline{x}_t\|^2 + \|\overline{u}_t\|^2 \right) && \text{(50a)} \\
&\le \frac{C^2(1+C^2)}{2(1-\widehat{\lambda})} \sum_{t=0}^{T-1} \sum_{t'=0}^{t-1} \widehat{\lambda}^{t-t'} \|w_{t'}\|^2 && \text{(50b)} \\
&\le \frac{C^2(1+C^2)}{2(1-\widehat{\lambda})^2} \sum_{t=0}^{T-2} \|w_t\|^2 ,
\end{aligned}$$

where we have used the assumption that $Q_t \preceq \ell I$ and $R_t \preceq \ell I$ in (50a); we have also used the inequalities (48) and (49) in (50b). Combining (50) with the lower bound of J^* in Lemma 7 finishes the proof of the third Statement in Theorem A.1.

F Proof of Theorem A.2

Before showing Theorem A.2, we first state a technical lemma that establishes the relationship between the TV distance and the Wasserstein distance.

Lemma 9. *For any distributions μ, ν on $\mathcal{X}$, we have*

$$W_1(\mu, \nu) = 2\mathsf{TV}(\mu, \nu) = \|\mu - \nu\|_1 .$$

Proof. To see this, note that since $\|x - x'\|_1 = 2$ for any $x \neq x'$, the Wasserstein 1-distance $W_1(\mu, \nu)$ equals 2 times the probability mass we need to transport to convert μ to ν. For every $i \in \{1, \ldots, n\}$ such that $\mu_i > \nu_i$, we need to move out exactly $(\mu_i - \nu_i)$ from the probability mass at e_i to other points $(e_j : j \neq i)$. Therefore, we must have

$$W_1(\mu, \nu) = 2 \sum_{i=1}^{n} \mathbf{1}(\mu_i > \nu_i) \cdot (\mu_i - \nu_i) = 2\mathsf{TV}(\mu, \nu) = \|\mu - \nu\|_1 .$$

□

Note that the MDP's transition kernel acts as a deterministic function. It maps the current state-action pair from $\mathcal{X} \times \mathcal{U}$ to the distribution of the subsequent state in $\mathcal{X}$. Hence, the current state-action distribution on $\mathcal{X} \times \mathcal{U}$ maps to a distribution on $\Delta(\mathcal{X})$. To proceed with this recursion, we require the distribution of the next state, which should be on $\mathcal{X}$. This is in contrast to needing the distribution of the distribution of the next state, which would be on $\Delta(\mathcal{X})$. Therefore, to convert the distributions on $\Delta(\mathcal{X})$ back to distributions on $\mathcal{X}$, we require the following lemma.

Lemma 10. *Let μ, μ' be two distributions on $\Delta(\mathcal{X})$. It follows that $\|\mathbb{E}[\mu] - \mathbb{E}[\mu']\|_1 \leq W_1(\mu, \mu')$.*

Note that $\mathbb{E}[\mu]$ and $\mathbb{E}[\mu']$ are distributions on $\mathcal{X}$.

Proof. By the definition of the Wasserstein distance, we have

$$W_1(\mu, \mu') = \inf_J \int \|x - y\|_1 \, \mathrm{d}J(x, y),$$

where J is a joint distribution on $\Delta(\mathcal{X}) \times \Delta(\mathcal{X})$ with marginals μ and μ'. For any such joint distribution J, we have

$$\int \|x - y\|_1 \, \mathrm{d}J(x, y) \geq \left\| \int (x - y) \mathrm{d}J(x, y) \right\|_1 = \|\mathbb{E}[\mu] - \mathbb{E}[\mu']\|_1 .$$

This finishes the proof of Lemma 10. □

We now resume our discussion with the proof of Theorem A.2.

Given the state-action distribution ρ at step t, let $\mu_{t:t'}(\rho)$ denote the resulting state distribution at step t'. We slightly abuse the notation so that for any pair $(x, u) \in \mathcal{X} \times \mathcal{U}$, $\mu_{t:t'}((x, u))$ still outputs the resulting state distribution at step t'. We see that

$$\|\mu_{t:t+1}((x, u)) - \mu_{t:t+1}((x', u'))\|_1 \leq \|x - x'\|_1 + \|u - u'\|_1 .$$

Therefore, by Lemmas 8 and 10, we see that

$$W(\mu_{t_1:t_1+1}(\rho), \mu_{t_1:t_1+1}(\rho')) \leq W(\rho, \rho'). \tag{51}$$

Note that Lemma 1 and Lemma 9 imply that

$$W(\mu_{t_1:t_2}(\rho), \mu_{t_1:t_2}(\rho')) \leq \lambda^{t_2 - t_1 - 1} W(\mu_{t_1:t_1+1}(\rho), \mu_{t_1:t_1+1}(\rho')).$$

Combining this with (51) gives that

$$W(\mu_{t_1:t_2}(\rho), \mu_{t_1:t_2}(\rho')) \leq \lambda^{t_2 - t_1 - 1} W(\rho, \rho'). \tag{52}$$

For any distributions μ, μ' on $\mathcal{X}$, we also see that $\|\bar{\pi}_{t_2}(\mu) - \bar{\pi}_{t_2}(\mu')\|_1 \leq \|\mu - \mu'\|$, which implies

$$W(\bar{\pi}_{t_2}(\mu), \bar{\pi}_{t_2}(\mu')) \leq W(\mu, \mu').$$

Substituting this into (52) gives that

$$W(\rho_{t_1:t_2}(\rho), \rho_{t_1:t_2}(\rho')) \leq 2\lambda^{t_2 - t_1 - 1} W(\rho, \rho'),$$

validating that the Wasserstein robustness (Definition 3) is satisfied.