TACKLING PERFORMATIVITY IN DISCRETE-TIME DYNAMICAL SYSTEMS: AN ITERATIVE REFINEMENT APPROACH

BY

HELING ZHANG

THESIS

Submitted in partial fulfillment of the requirements
for the degree of Master of Science in Electrical and Computer Engineering
in the Graduate College of the
University of Illinois Urbana-Champaign, 2023

Urbana, Illinois

Advisor:

Assistant Professor Roy Dong

Abstract

In many real-world dynamical systems, obtaining precise prior knowledge about system noise remains a challenge. This uncertainty complicates traditional control strategies, such as stochastic and robust control, especially when the noise exhibits "performativity"— an explicit dependence on control inputs. Addressing this challenge, this paper presents a novel iterative method tailored for such systems. Our approach finds the open-loop control law that minimizes the worst-case loss, given that the noise induced by this control lies in its $(1-p)$-confidence set for a predetermined p. At each iteration, we harness conformal prediction techniques to empirically estimate the confidence set shaped by the preceding control law. These derived confidence sets offer empirical constraints on the system's noise, guiding a robust control design that targets worst-case loss minimization. Under specific regularity conditions, our method is shown to converge to a near-optimal open-loop control. While our focus is on open-loop controls, the adaptive, data-driven nature of our approach suggests its potential applicability across diverse scenarios and extensions.

Table of contents

List of Symbols

$x(t) \in \mathbb{R}^n$	State vector at time step t.
$w(t) \in \mathbb{R}^n$	Noise vector at time step t.
$u(t) \in \mathbb{R}^m$	Control input at time step t.
$\mathbf{x}$	Concatenation of state trajectory, $(x(0), \ldots, x(T))$.
$\mathbf{w}$	Concatenation of noise trajectory, $(w(0), \ldots, x(T-1))$.
$\mathbf{u}$	Concatenation of control trajectory, $(u(0), \ldots, u(T-1))$.
$\nabla_x f(x, y)$	Gradient of f in x.
$\partial_{\mathbf{u}} L(\mathbf{w}, \mathbf{u})$	Subgradient of L in $\mathbf{u}$.
$\lvert \cdot \rvert$	Euclidean distance, short-hand for $\lVert \cdot \rVert_2$.
$\mathbb{P}(\cdot)$	Probability.
$\lceil \cdot \rceil$	Ceiling operator.

Chapter 1

Introduction

The concept of performativity traces back to British philosopher John L. Austin's work in 1962, where he posited that language and other forms of expression not only describe the world but can also influence and change it [3]. This idea was later adapted to the realm of supervised learning by Perdomo et al., highlighting scenarios where predictive models inadvertently alter the very data distributions they aim to learn [13].

This performative phenomenon is not exclusive to language or machine learning; it also manifests in control systems. Consider a system governed by the following stochastic differential equation (SDE)

$$dX_t = f(t, X_t, u_t)\, dt + g(t, X_t, u_t)\, dW_t, \tag{1.1}$$

where $(X_t)_{0\leq t\leq 1}$ is the n-dimensional state process, $(u_t)_{0\leq t\leq 1}$ is the m-dimensional control process, and $(W_t)_{0\leq t\leq 1}$ is a k-dimensional Brownian motion. When $g(X_t, u_t)$ has an explicit dependence on u_t, we say that the control trajectory $(u_t)_{0\leq t\leq 1}$ is performative. Similarly, we can also define performativity for the discrete-time counterpart of the system in Equation 1.1

$$x_{t+1} = f(t, x_t, u_t, w_t), \tag{1.2}$$

where w_t is a member of some probability space $(W, \mathcal{F})$ and has distribution $p(dw_t|x_t, u_t)$. We say that above system is performative if $p(dw_t|x_t, u_t)$ has an explicit dependence on u_t.

Historically, such systems have been the focus of stochastic control, where the objective is to minimize the expected loss, exemplified by:

$$J := \mathbb{E}\left[\int_0^1 q(X_s, u_s)\, ds + r(X_1) \,\middle|\, X_0 = x_0 \right]. \tag{1.3}$$

Robust control offers another perspective, especially for discrete-time systems like Equation 1.2. If $w(t)$ always bounded, i.e. $w(t) \in C(u(t))$ for some closed and bounded $C(u(t))$, then the focus shifts to minimizing the worst-case loss:

$$J = \max \left\{ \sum_{i=0}^{T} q(x(t), u(t)) + r(X(T)) \,\middle|\, x(0) = 0, w(t) \in C(u(t)) \right\}.$$

Both paradigms—stochastic and robust control— rely on explicit knowledge of the noise distribution. Yet, real-world systems often lack this clarity. Addressing this gap, our study focus on scenarios in which no prior knowledge about noise distribution is available. Instead, we harness the power of empirical observations, drawing from sample trajectories to inform and refine control laws.

Designing controllers with safety guarantees is a common theme when dealing with noisy systems. When the noise is completely known, forward reachability and related methods can provide provable guarantees of safety [14]. In cases where the noise is unknown, it is common to seek online learning methods to learn the noise empirically, as in [1], [8]. This paper falls into the same category as the above works, but provides a different type of guarantee.

A key component of our method is *conformal prediction* [2], [15], a technique that allows us to construct confidence sets with rigorous probability from finite samples. There exist many prior works that apply conformal prediction in control settings. For instance, conformal prediction is used in [4], [6], [11] to estimate reachable sets, which is then used for safety planning of autonomous systems. In [7], the authors use conformal prediction to derive probabilistic guarantees on the satisfication of certain system specifications. In [5], [10], [16], conformal prediction is used to construct predictors with probabilistic guarantees, which is then used to design model predictive controllers. Our work is another instance of applying conformal prediction on control problems, although in a way different from all prior works.

Methodologically, our work is inspired by the works on *performative prediction* [13], which is the study of performativity in supervised learning. Conceptually, our problem shares several similarities with performative prediction: 1. the solutions to our problems alter the noise we encounter, and 2. we both seek the solution that gives the best performance in the noise it induces. As a result, our method adopts a similar iterative structure, in which each iteration builds upon the noise induced by the result from the previous iteration, but eventually converges to a fixed-point under certain regularity conditions.

Chapter 2

Tackling Performative Control

2.1 Problem Statement

This paper delves into the robust optimal control of a discrete-time, finite time horizon dynamical system subject to noise. Consider the system given by

$$x(t+1) = f(t, x(t), u(t)) + w(t), \quad t = 0, 1, \ldots, T-1, \tag{2.1}$$

where $x(t) \in X_t \subseteq \mathbb{R}^n$ is the system state, $u(t) \in U_t \subseteq \mathbb{R}^m$ is the control input, and $w(t) \in \mathbb{R}^n$ is a zero-mean process noise whose distribution depends on $x(t)$ and $u(t)$, i.e. $w(t) \sim \mathcal{D}(x(t), u(t))$. We also assume without loss of generality that the system starts at a fixed state $x(0) = x_0$. With a slight abuse of notation, we use x, u, w to denote associated trajectories

$$\begin{aligned}
\mathbf{x} &= (x(0), \ldots, x(T)) \in \mathbb{R}^{n(T+1)}, \\
\mathbf{u} &= (u(0), \ldots, u(T-1)) \in \mathbb{R}^{mT}, \\
\mathbf{w} &= (w(0), \ldots, w(T-1)) \in \mathbb{R}^{nT}.
\end{aligned}$$

The objective is to determine an open-loop control law $u = (u(0), ...u(T-1))$ that optimally governs the system's behavior. To quantify the performance of a control law, we introduce a loss function J dependent on the entire trajectory of states and control inputs. This loss function can be represented as

$$J(\mathbf{x}, \mathbf{u}) := \phi(x(T)) + \sum_{t=0}^{T-1} \varphi(x(t), u(t)).$$

However, given the stochastic nature of the noise $w(t)$, the loss function J becomes a random object. This makes it impossible to directly minimize the value of the loss function. Traditional approaches either aim to minimize the expected loss, requiring knowledge of the noise distribution, or minimize the worst-case loss for bounded noise, requiring knowledge of the extreme values of the noise. Both approaches require some prior knowledge of the behavior of the noise. This work addresses scenarios where no such prior knowledge is available, except its zero-mean property.

When sampling is feasible, we can construct confidence sets for the noise with probabilistic guarantees using techniques like conformal prediction. The primary focus of this work is to minimize the worst-case loss given these confidence sets. Formally, for a designated probability level

$$p \in (0,1), \tag{2.2}$$

let $C_t(\mathbf{u})$ be the confidence set for $w(t)$ such that $w(t) \in C_t(\mathbf{u})$ with probability at least $1 - p/T$. Applying the union bound, this would give us $\mathbf{w} \in C_0(\mathbf{u}) \times \cdots \times C_{T-1}(\mathbf{u})$ with probability at least $1 - p$. We are interested in the following optimization problem

$$\begin{aligned} \min_{\mathbf{u}} \quad & \max_{\mathbf{w}} J(\mathbf{x}, \mathbf{u}) \\ \text{s.t} \quad & \mathbf{w} \in C_0(\mathbf{u}) \times \cdots \times C_{T-1}(\mathbf{u}) \\ & \mathbf{u} \in U_0 \times \cdots \times U_{T-1}. \end{aligned} \tag{2.3}$$

For clarity, we can also equivalently express the loss as $L(\mathbf{w}, \mathbf{u}) = J(\mathbf{x}, \mathbf{u})$, emphasizing its dependence on the noise and control. Further, let

$$\begin{aligned} C(\mathbf{u}) &:= C_0(\mathbf{u}) \times \cdots \times C_{T-1}(\mathbf{u}) \\ U &:= U_0 \times \cdots \times U_{T-1}, \end{aligned}$$

then we can write the optimization problem in a more concise form

$$\begin{aligned} \min_{\mathbf{u}} \quad & \max_{\mathbf{w} \in C(\mathbf{u})} L(\mathbf{w}, \mathbf{u}) \\ \text{s.t} \quad & \mathbf{u} \in U. \end{aligned} \tag{2.4}$$

This problem can be interpreted as a middle-ground between stochastic control and robust control. By concentrating on the worst-case scenario within these confidence sets, we ensure robustness against the most adverse noise scenarios while avoiding being excessively conservative.

2.2 Methods

In this section, we introduce empirical iterative refinement of performative control (E-IRPC), an algorithm that targets the problem described in the previous section.

2.2.1 Empirical Iterative Refinement of Performative Control

The core of our methodology is an iterative process designed to progressively refine the control law to achieve optimal performance in the presence of noise. Our method consists of the following steps:

1. **Initialization:** We start by finding the optimal open-loop control for the nominal system (i.e. the system without noise) with loss function J as defined in Equation 2.1

$$\begin{aligned} \min_{\mathbf{u}} \quad & J(\mathbf{x}, \mathbf{u}) \\ \text{s.t} \quad & x(t+1) = f(t, x(t), u(t)) \\ & \mathbf{u} \in U_0 \times \cdots \times U_{T-1}. \end{aligned} \tag{2.5}$$

   We denote the resulting control law as $\mathbf{u}_0 = (u_0(0), \ldots, u_0(T-1))$.

2. **Sampling and Confidence Set Construction:** With the control law from the previous step fixed, sample N_1 independent trajectories of system states, represented as $\mathbf{x}_1^{(j)}$ for $j = 1, \dots, N_1$. This also gives us N_1 trajectories of the noise, $\mathbf{w}_1^{(j)}$ for $j = 1, \dots, N_1$. Notice that since we are using the same control input, at each time step t, $(w_1^{(j)}(t))_{j=1}^{N_1}$ are i.i.d. We can then use conformal prediction to construct
$$\mathbb{P}(w(t) \in \hat{C}_t(\mathbf{u})) \geq 1 - \frac{p}{T}. \tag{2.6}$$
Again, for simplicity, we denote $\hat{C}(\mathbf{u}) := \hat{C}_0(\mathbf{u}) \times \cdots \times \hat{C}_{T-1}(\mathbf{u})$. The details of such construction will be outlined in the next subsection.

3. **Robust Control Problem Formulation:** With the confidence sets in hand, we can formulate a robust control problem similar to Equation 2.3
$$\begin{aligned} \min_{\mathbf{u}} \quad & \max_{\mathbf{w}} J(\mathbf{x}, \mathbf{u}) \\ \text{s.t} \quad & \mathbf{w} \in \hat{C}_0(u_0) \times \cdots \times \hat{C}_{T-1}(u_0) \\ & \mathbf{u} \in U_0 \times \cdots \times U_{T-1}. \end{aligned}$$
Solve this problem to obtain a new control law $\mathbf{u}_1$.

4. **Iterative Refinement:** The new control law is then used as the initial control law in step 2, and the process repeats.

In summary, for $i = 0, 1, ...,$ we fix the control law $\mathbf{u}_i$, sample N_i trajectories $(\mathbf{x}_{i+1}^{(j)})_{j=1}^{N_i}$ and $(\mathbf{w}_{i+1}^{(j)})_{j=1}^{N_i}$ to construct a confidence set $\hat{C}(\mathbf{u}_i)$, then formulate and solve a robust control problem to obtain u_{i+1}. To simplify our notation, define $\hat{\mathcal{A}} : U \to U$ as
$$\hat{\mathcal{A}}(\mathbf{u}) \in \arg\min \left\{ \max\{L(\mathbf{w}, \mathbf{u}) : \mathbf{w} \in \hat{C}(\mathbf{u})\} : \mathbf{u} \in U \right\}, \tag{2.7}$$
then our method is summarized in Algorithm 1.

The iterative nature of this method is essential due to the performative aspect of the control: the control law influences the distribution of the noise, necessitating repeated refinement.

Algorithm 1 E-IRPC

```
u_0 ← Solution to Problem 2.5
while i > 0 do
    Sample (x_{i+1}^{(j)})_{j=1}^{N_i} and (w_{i+1}^{(j)})_{j=1}^{N_i}
    Construct Ĉ(u_{i−1}) from samples
    u_i ← Â(u_{i−1})                    ▷ where Â is as defined in Equation 2.7.
end while
```

2.2.2 Confidence Set Construction using conformal prediction

In this section, we will outline the method we use to construct the confidence sets in Step 2. As described in the previous section, for each timestep t, we have N_i samples of $w(t)$. Let's sort their Euclidean norm in the increasing order as follows
$$\left| w^{(\eta_1)}(t) \right| \leq \cdots \leq \left| w^{(\eta_{N_i})}(t) \right|.$$

Since these samples are i.i.d, we can easily show that the Euclidean norm of another i.i.d sample $w(t)$ is equally likely to fall between that of any existing samples. Formally, this gives us

$$\mathbb{P}\left(|w(t)| \leq |w^{(\eta_k)}(t)|\right) = \frac{k}{N_i + 1} \tag{2.8}$$

for any $k = 1, \ldots, N_i$. This allows us to construct the confidence sets with desired probabilistic guarantee. Specifically, for the probability level p defined in 2.2, pick $k = \left\lceil (N_i + 1)\left(1 - \frac{p}{T}\right)\right\rceil$ and define

$$\hat{C}_t(\mathbf{u}) := \left\{w(t) : |w(t)| \leq |w^{(\eta_k)}|\right\}, \tag{2.9}$$

by Equation 2.8, we have

$$\mathbb{P}\left(w(t) \in \hat{C}_t(\mathbf{u})\right) = \frac{\left\lceil (N_i + 1)\left(1 - \frac{p}{T}\right)\right\rceil}{N_i + 1} \geq 1 - \frac{p}{T},$$

which means that Equation 2.9 gives us the desired confidence set. It is worth noting that when constructing these confidence sets, we can replace the Euclidean norm with any other score function. This will allow us to construct confidence sets with the same probabilistic guarantee but different shapes. For example, if we can chose $s(w(t)) := \sqrt{w^\intercal(t) H w(t)}$ for some positive definite definite H, and rank the sample noises according the score

$$s\left(w^{(\eta_1)}(t)\right) \leq \cdots \leq s\left(w^{(\eta_{N_i})}(t)\right)$$

the resulting confidence set $\tilde{C}_t(\mathbf{u}) = \left\{w(t) : s(w(t)) \leq s\left(w^{(\eta_k)}(t)\right)\right\}$ will still have the same probabilistic guarantee, but takes the shape of an ellipsoid instead of a ball.

2.2.3 Ideal Iterative Refinement of Performative Control

So far, we have detailed our primary method, which inherently relies on estimating specific properties of the noise distribution using finite samples. Naturally, the algorithm's performance is influenced by the number of samples taken. In this section, we introduce an idealized variant of our method, termed the *Ideal Iterative Refinement of Performative Control* (I-IRPC). This version sidesteps the variability introduced by finite sample trajectories. As we will discuss, the I-IRPC can be conceptualized as the infinite sample counterpart of our E-IRPC.

Formally, let $\mathbf{u}$ be the control input for the system defined by 2.1, and $w(t)$ be the noise at time t whose distribution is given by $\mathcal{D}(x(t), u(t))$. We define the *ideal conformal set* for $w(t)$ as

$$C_t(\mathbf{u}) = \{w(t) : |w(t)| \leq Q_t(\mathbf{u})\}, \tag{2.10}$$

where $Q_t(\mathbf{u})$ is the quantile function of $|w(t)|$ defined as

$$Q_t(\mathbf{u}) := \inf\left\{r : \mathbb{P}(|w(t)| \leq r) \geq 1 - \frac{p}{T}\right\} \tag{2.11}$$

with p as the overall probability level predefined by 2.2. The set $C_t(\mathbf{u})$ is the idealized counterpart of $\hat{C}_t(\mathbf{u})$ defined in 2.9, which is impossible to obtain without knowing the exact distribution of $w(t)$. In the I-IRPC, we simply assume that there is an oracle that gives $C_t(\mathbf{u})$ for any $\mathbf{u}$ we choose.

Notice that Q_t depends solely on $\mathbf{u}$ because $\mathbf{u}$ uniquely defines the distribution of $w(t)$. Now, let

$C(\mathbf{u}) := C_0(\mathbf{u}) \times \cdots \times C_{T-1}(\mathbf{u})$ and define $\mathcal{A} : U \to U$ as

$$\mathcal{A}(\mathbf{u}) \in \arg\min \left\{ \max\{L(\mathbf{w}, \mathbf{u}) : \mathbf{w} \in C(\mathbf{u})\} : \mathbf{u} \in U \right\}. \tag{2.12}$$

Then the I-IRPC is given by Algorithm 2. The I-IRPC gives us a starting point for analysis. Specifically, due

Algorithm 2 I-IRPC

```
u_0 ← Solution to Problem 2.5
while i > 0 do
    Get C(u) from the oracle
    u_i ← A(u_{i-1})                                    ▷ where A is as defined in Equation 2.12.
end while
```

to its deterministic nature, it allows us to make the following key definitions.

Definition 1 (Performatively Stable Control)**.** A *performatively stable control* is a control input $\mathbf{u}_{PS} \in U$ such that

$$\mathcal{A}(\mathbf{u}_{PS}) = \mathbf{u}_{PS}.$$

Definition 2 (Performatively Optimal Control)**.** A *performatively optimal control* is a control input $\mathbf{u}_{PO} \in U$ such that

$$\mathbf{u}_{PO} \in \arg\min_{\mathbf{u} \in U} \max_{\mathbf{w} \in C(\mathbf{u})} L(\mathbf{w}, \mathbf{u}).$$

where $C(\mathbf{u})$ is defined in Equation 2.10.

Conceptually, the performatively stable control is the control law that performs optimally in the noise it induces. The performatively optimal control, on the other hand, gives the best performance among all feasible control laws in the presence of self-induced noise. Notice that the performative optimal control, in contrast to performative stable control, may not be the optimal control law in the noise that it induces. As we will see shortly, these concepts play a central role in our theoretical analysis of our proposed method.

Now, let's revisit E-IRPC from the lens of I-IRPC. Define the empirical quantile function by

$$\hat{Q}_t(\mathbf{u}) := \inf \left\{ r : \frac{1}{N_i} \sum_{j=1}^{N_i} \mathbf{1} \left\{ \left| w^{(j)}(t) \right| \le r \right\} \ge 1 - \frac{p}{T} \right\} = \left| w^{(\eta_k)}(t) \right|. \tag{2.13}$$

As the name suggests, the empirical quantile function is a finite sample approximation of the actual quantile function defined in Equation 2.11. The conformal sets constructed in Section 2.2.2 is then given by

$$\hat{C}_t(\mathbf{u}) = \left\{ w(t) : |w(t)| \le \hat{Q}_t(\mathbf{u}) \right\},$$
$$\hat{C}_u = \hat{C}_0(\mathbf{u}) \times \cdots \times \hat{C}_{T-1}(\mathbf{u}).$$

In this sense, the E-IRPC algorithm is but a finite-sample approximation of I-IRPC. This connection is central in the upcoming analysis.

2.3 Preliminaries

We start by presenting some standard assumptions on the loss function $L(\mathbf{w}, \mathbf{u})$.

Assumption 1 (λ-strong convexity in $\mathbf{u}$)**.** *We assume that the loss function $L(\mathbf{w}, \mathbf{u})$ is λ-strongly convex in $\mathbf{u}$, i.e. for any $\mathbf{u}, \mathbf{u}' \in U$, and any $\mathbf{w} \in \mathbb{R}^{nT}$,*

$$L(\mathbf{w}, \mathbf{u}) \geq L(\mathbf{w}, \mathbf{u}') + (\mathbf{u} - \mathbf{u}')^\intercal s_{\mathbf{u}'} + \frac{\lambda}{2}|\mathbf{u} - \mathbf{u}'|^2,$$

where $s_{\mathbf{u}'} \in \partial_{\mathbf{u}'} L(\mathbf{w}, \mathbf{u}')$.

Assumption 2 (β-smoothness in $\mathbf{w}$)**.** *We assume that the loss function $L(\mathbf{w}, \mathbf{u})$ is β-smooth in $\mathbf{w}$, i.e. $L(\mathbf{w}, \mathbf{u})$ is differentiable in $\mathbf{w}$ and the gradient $\nabla_w L(\mathbf{w}, \mathbf{u})$ is β-Lipschitz in $\mathbf{w}$, that is, for any $\mathbf{w}, \mathbf{w}' \in \mathbb{R}^{nT}$ and any $\mathbf{u} \in U$,*

$$|\nabla_w L(\mathbf{w}, \mathbf{u}) - \nabla_{w'} L(\mathbf{w}', \mathbf{u})| \leq \beta |\mathbf{w} - \mathbf{w}'|.$$

We now introduce an important assumption regarding the noise distribution.

Assumption 3 (Lipschitz continuity of quantile functions)**.** *Let $Q_t(\mathbf{u})$ be the $(1 - p/T)$-quantile function for $|w(t)|$ when $\mathbf{u}$ is the control input and p is the predefined probability level given by Equation 2.2 (the definition of $Q_t(\mathbf{u})$ is given in Equation 2.11). We assume that $Q_t(\mathbf{u})$ is ϵ_t-Lipschitz, i.e. for all $\mathbf{u}, \mathbf{u}' \in U$,*

$$|Q_t(\mathbf{u}) - Q_t(\mathbf{u}')| \leq \epsilon_t |\mathbf{u} - \mathbf{u}'|.$$

Next, we introduce some existing results needed to establish our results. We will not prove them here, but interested readers can find detailed proof from the reference we provide.

Proposition 1 (Maximum Theorem [12, Chapter E.3])**.** *Let X and Θ be topological spaces, $f : X \times \Theta \to \mathbb{R}$ be a continuous function and $C : \Theta \to 2^X$ is a set-valued function such that $C(\theta)$ is compact and non-empty for all $\theta \in \Theta$. Then the function $f^* : \Theta \to \mathbb{R}$ defined by*

$$f^*(\theta) = \sup\{f(x, \theta) : x \in C(\theta)\},$$

is continuous, and the maximizers defined by

$$C^*(\theta) = \arg\max\{f(x, \theta) : x \in C(\theta)\} = \{x \in C(\theta) : f(x, \theta) = f^*(\theta)\}$$

is upper hemicontinuous with nonempty and compact values.

2.4 Main Results

In this section, we delve into a rigorous theoretical analysis of our proposed method. Our exploration commences with an in-depth examination of the deterministic counterpart of our method, termed the I-IRPC. Initially, we delineate the sufficient conditions that guarantee the existence and uniqueness of its fixed points, which correspond to the performatively stable control. Subsequently, we outline the conditions under which this method converges. Furthermore, under mild supplementary assumptions, we demonstrate that the performatively stable control closely approximates the performative optimal control. Transitioning from the ideal to the practical, we extend these insights to our primary method, the E-IRPC. This method can be conceptualized as an empirical adaptation of the I-IRPC. To set the stage for our detailed analysis, we first elucidate some foundational assumptions and preliminary findings that underpin our subsequent discussions.

2.4.1 Existence of fixed-points and optima

We are now ready to establish our main results. We starting by investigating the existence of performatively stable control and performatively optimal control. To declutter notation, define

$$g(\mathbf{u},\mathbf{v}) = \max\{L(\mathbf{w},\mathbf{v}) : \mathbf{w} \in C(\mathbf{u})\}. \tag{2.14}$$

Theorem 1 (Existence of performatively stable control)**.** *Suppose the following assumptions hold:*

- $L(\mathbf{w},\mathbf{u})$ *is strictly convex in* $\mathbf{u}$*,*
- $g(\mathbf{u},\mathbf{v})$ *jointly continuous in* $\mathbf{u}$ *and* $\mathbf{v}$*,*
- *the quantile functions* $Q_t(\mathbf{u})$ *are* ϵ_t *in* $\mathbf{u}$ *for all* $t = 0, \dots, T-1$ *(Assumption 3), and*
- U *is compact and convex.*

Then there exists at least one performatively stable control.

Proof. Since $g(\mathbf{u},\mathbf{v})$ jointly continuous in $\mathbf{u}$ and $\mathbf{v}$ and the constant set-valued mapping $\mathbf{v} \mapsto U$ is continuous, by Proposition 1, we know that the set valued mapping given by

$$\mathbf{u} \mapsto \arg\min\{g(\mathbf{u},\mathbf{v}) : \mathbf{v} \in U\}$$

is upper hemicontinuous. Further, by strict convexity of $L(\mathbf{w},\mathbf{u})$ in $\mathbf{u}$, $g(\mathbf{u},\mathbf{v})$ is also strictly convex in $\mathbf{u}$. Therefore, the set $\arg\min\{g(\mathbf{u},\mathbf{v}) : \mathbf{v} \in U\}$ contains only one element, that is

$$\arg\min\{g(\mathbf{u},\mathbf{v}) : \mathbf{v} \in U\} = \{\mathcal{A}(\mathbf{u})\}.$$

Then, the continuity of $\mathcal{A}$ follows directly from the upper hemicontinuity of $\mathbf{u} \mapsto \{\mathcal{A}(\mathbf{u})\}$. □

Uniqueness of the performatively stable control requires a bit more conditions, as we will see shortly. The existence and uniqueness of the performatively optimal control, on the other hand, is much easier to establish.

Theorem 2 (Existence and Uniqueness of performatively optimal control)**.** *Suppose the following assumptions hold:*

- $L(\mathbf{w},\mathbf{u})$ *is jointly continuous in* $\mathbf{u}$ *and*
- U *is compact and convex.*

Then there exists a unique performatively optimal control.

Before proving this, we first establish the following lemma that is necessary for this proof.

Lemma 1. *If the quantile functions* $Q_t(\mathbf{u})$ *are* ϵ_t*-Lipschitz (Assumption 3), then the set-valued mapping* $C(\mathbf{u})$ *defined by*

$$C(\mathbf{u}) = C_0(\mathbf{u}) \times \cdots \times C_{T-1}$$

with

$$C_t(\mathbf{u}) = \{w(t) : |w(t)| \leq Q_t(\mathbf{u})\}, \quad t = 0, \dots, T-1$$

is continuous.

Proof. We have to show that $C(\mathbf{u})$ is both upper hemicontinuous and lower hemicontinuous.

- To show upper hemicontinuity, we need to show that for every $\mathbf{u} \in U$ and any sequence $(\mathbf{u}^k)_{k=1}^{\infty}$ with $\lim_{k\to\infty} \mathbf{u}^k = \mathbf{u}$, the limit of any convergent sequence $(\mathbf{w}^k)_{k=1}^{\infty}$ with $\mathbf{w}^k \in C(\mathbf{u}^k)$ lies in $C(\mathbf{u})$.

  Suppose for contradiction that there exists a sequence $(\mathbf{w}^k)_{k=1}^{\infty}$ with $\mathbf{w}^k \in C(\mathbf{u}^k)$ and $\lim_{k\to\infty} \mathbf{w}^k = \mathbf{w} \notin C(\mathbf{u})$. By the definition of $C(\mathbf{u})$, this means that there exists some integer $t \in [0, T-1]$ and $\delta > 0$ such that $|w(t)| \geq Q_t(\mathbf{u}) + \delta$. Since $\lim_{k\to\infty} \mathbf{w}^k = \mathbf{w}$, we can pick a constant $K_1(\delta/2)$ such that

  $$|w^k(t)| - |w(t)| \leq |w^k(t) - w(t)||\mathbf{w}^k - \mathbf{w}| \leq \frac{\delta}{2}$$

  for all $k > K_1(\delta/2)$. Similarly, pick $K_2(\delta/2\epsilon_t)$ with

  $$|Q_t(\mathbf{u}^k) - Q_t(\mathbf{u})|\epsilon_t|\mathbf{u}^k - \mathbf{u}| < \frac{\delta}{2}$$

  for all $k > K_2(\delta/w\epsilon_t)$, we have

  $$\begin{aligned} |w^k(t)| - Q_t(\mathbf{u}^k) &= \big(|w^k(t)| - |w(t)|\big) + \big(|w(t)| - Q_t(\mathbf{u}^k)\big) + \big(Q_t(\mathbf{u}^k) - Q_t(\mathbf{u})\big) \\ &\geq \big(|w(t)| - Q_t(\mathbf{u}^k)\big) - \big||w^k(t)| - |w(t)|\big| - \big|Q_t(\mathbf{u}^k) - Q_t(\mathbf{u})\big| \\ &> \delta - \delta/2 - \delta/2 \\ &= 0, \end{aligned}$$

  thus $\mathbf{w}^k \notin C(\mathbf{u}^k)$ for all $k > \max\{K_1(\delta/2), K_2(\delta/2\epsilon_t)\}$. We have a contradiction.

- To prove lower hemicontinuity, we need to show that for any $\mathbf{u} \in U$ and any sequence $(\mathbf{u}^k)_{k=1}^{\infty}$ with $\lim_{k\to\infty} \mathbf{u}^k = \mathbf{u}$, for any $\mathbf{w} \in C(\mathbf{u})$, there exists $(\mathbf{w}^k)_{k=1}^{\infty}$ with $\mathbf{w}^k \in C(\mathbf{u}^k)$ such that $\lim_{k\to\infty} \mathbf{w}^k = \mathbf{w}$. Pick

  $$w^k(t) = \min\{Q_t(\mathbf{u}^k), |w(t)|\}\frac{w(t)}{|w(t)|},$$

  clearly $\mathbf{w}^k \in C(\mathbf{u}^k)$. Now since

  $$\begin{aligned} |\mathbf{w}^k - w| &= \sqrt{\sum_{t=0}^{T-1} |w^k(t) - w(t)|^2} \\ &\leq \sqrt{\sum_{t=0}^{T-1} |Q_t(\mathbf{u}^k) - Q_t(\mathbf{u})|^2} \\ &\leq \sqrt{\epsilon_t^2 |\mathbf{u}^k - \mathbf{u}|^2} \\ &= \left(\sum_{t=0}^{T-1} \epsilon_t^2\right) |\mathbf{u}^k - \mathbf{u}| \\ &\xrightarrow{k\to\infty} 0, \end{aligned}$$

  we have $\lim_{k\to\infty} \mathbf{w}^k = \mathbf{w}$. This completes the proof.

$\square$

Now we are ready to prove Theorem 2.

Proof of Theorem 2. Define

$$h(\mathbf{u}) := \max_{\mathbf{w}\in C(\mathbf{u})} L(\mathbf{w},\mathbf{u}) = g(\mathbf{u},\mathbf{u})$$

Notice the performative optimal control is the minimizer of $h(\mathbf{u})$ over U. Since $\mathbf{u} \mapsto C(\mathbf{u})$ is continuous, by Proposition 1, $h(\mathbf{u})$ is continuous in $\mathbf{u}$. The existence of the performative optimal control then follows from the Weierstrass Extreme Value Theorem. □

2.4.2 Convergence Guarantees with Aligned Worst-Case Noise

We are now ready to investigate the convergence properties of the I-IRPC algorithm. Our analysis relies on the following additional assumption.

Definition 3. We say that the worst case noise align for the loss function L and the confidence set mapping C if for any $\mathbf{v}, \mathbf{u}, \mathbf{u}' \in U$, align if there exists $\mathbf{w}_1^* \in \arg\max_{\mathbf{w}\in C(\mathbf{u})} L(\mathbf{w},\mathbf{v})$ and $\mathbf{w}_2^* \in \arg\max_{\mathbf{w}\in C(\mathbf{u}')} L(\mathbf{w},\mathbf{v})$ such that

$$\frac{\mathbf{w}_1^*}{|\mathbf{w}_1^*|} = \frac{\mathbf{w}_2^*}{|\mathbf{w}_2^*|}.$$

Assumption 4 (Alignment of worst-case noises)**.** *We assume that the worst-case noises align.*

Theorem 3 (Convergence of I-IRPC to $\mathbf{u}_{PS}$)**.** *Suppose the following assumptions hold:*

- $L(\mathbf{w},\mathbf{u})$ *is λ-strongly convex in* $\mathbf{u}$ *(Assumption 1),*
- $L(\mathbf{w},\mathbf{u})$ *is β-smooth in* $\mathbf{w}$ *(Assumption 2),*
- *the quantile function $Q_t(\mathbf{u})$ is ϵ_t-Lipschitz in $\mathbf{u}$ for all $t = 0,\ldots,T-1$ with $\sum_{t=1}^{T-1} \epsilon_t^2 < \frac{\lambda^2}{\beta^2}$ (Assumption 3) and*
- *the worst case noises align (Assumption 4).*

Then the I-IRPC converges to a unique performative stable control $\mathbf{u}_{PS}$ at a linear rate:

$$|\mathbf{u}_i - \mathbf{u}_{PS}| \leq \delta \text{ for all } i \geq \left(1 - \frac{\beta\sqrt{\sum_{t=0}^{T-1}\epsilon_t^2}}{\lambda}\right)^{-1} \log\left(\frac{|\mathbf{u}_0 - \mathbf{u}_{PS}|}{\delta}\right). \tag{2.15}$$

Proof. Let $\mathbf{u}, \mathbf{u}' \in U$ be two different control inputs. Since $L(\mathbf{w},\mathbf{u})$ is λ-strongly convex in $\mathbf{u}$, for any $\mathbf{u} \in U$, $g(\mathbf{u},\mathbf{v})$ defined by Equation 2.14 is λ-strongly convex in $\mathbf{u}$. Let $\partial_{\mathbf{u}} g(\mathbf{u},\mathbf{v})$ denote the subgradient of $g(\mathbf{u},\mathbf{v})$ with respect to $\mathbf{u}$. Then by strong convexity, for any $s_{\mathbf{u}} \in \partial_{\mathcal{A}(\mathbf{u})} g(\mathbf{u},\mathcal{A}(\mathbf{u}))$ and any $s_{\mathbf{u}'} \in \partial_{\mathcal{A}(\mathbf{u}')} g(\mathbf{u},\mathcal{A}(\mathbf{u}'))$, we have

$$g(\mathbf{u},\mathcal{A}(\mathbf{u})) - g(\mathbf{u},\mathcal{A}(\mathbf{u}')) \geq (\mathcal{A}(\mathbf{u}) - \mathcal{A}(\mathbf{u}'))^\intercal s_u + \frac{\lambda}{2}|\mathcal{A}(\mathbf{u}) - \mathcal{A}(\mathbf{u}')|^2$$
$$g(\mathbf{u},\mathcal{A}(\mathbf{u}')) - g(\mathbf{u},\mathcal{A}(\mathbf{u})) \geq (\mathcal{A}(\mathbf{u}') - \mathcal{A}(\mathbf{u}))^\intercal s_{\mathbf{u}'} + \frac{\lambda}{2}|\mathcal{A}(\mathbf{u}) - \mathcal{A}(\mathbf{u}')|^2,$$

Adding both sides we get

$$\lambda|\mathcal{A}(\mathbf{u}) - \mathcal{A}(\mathbf{u}')|^2 \leq (\mathcal{A}(\mathbf{u}) - \mathcal{A}(\mathbf{u}'))^\intercal s_u + (\mathcal{A}(\mathbf{u}') - \mathcal{A}(\mathbf{u}))^\intercal s_{\mathbf{u}'} \tag{2.16}$$

Further, by definition of subgradients, we also have

$$0 \geq g(\mathbf{u},\mathcal{A}(\mathbf{u})) - g(\mathbf{u},\mathcal{A}(\mathbf{u}')) \geq (\mathcal{A}(\mathbf{u}) - \mathcal{A}(\mathbf{u}'))^\intercal s_u \tag{2.17}$$

Similarly, for any $s'_{\mathbf{u}'} \in \partial_{\mathcal{A}(\mathbf{u}')} g(\mathbf{u}', \mathcal{A}(\mathbf{u}')$ (notice that the first argument in g is u' instead of $\mathbf{u}$)

$$0 \geq g(\mathbf{u}', \mathcal{A}(\mathbf{u})) - g(\mathbf{u}', \mathcal{A}(\mathbf{u}')) \geq (\mathcal{A}(\mathbf{u}) - \mathcal{A}(\mathbf{u}'))^\intercal s_u \tag{2.18}$$

Combining these with Equation 2.16, we get

$$\begin{aligned}
\lambda|\mathcal{A}(\mathbf{u}) - \mathcal{A}(\mathbf{u}')|^2 &\leq (\mathcal{A}(\mathbf{u}) - \mathcal{A}(\mathbf{u}'))^\intercal s_u + (\mathcal{A}(\mathbf{u}') - \mathcal{A}(\mathbf{u}))^\intercal s_{\mathbf{u}'} \\
&\leq (\mathcal{A}(\mathbf{u}') - \mathcal{A}(\mathbf{u}))^\intercal s_{\mathbf{u}'} && \text{[by Equation 2.17]} \\
&\leq (\mathcal{A}(\mathbf{u}') - \mathcal{A}(\mathbf{u}))^\intercal s_{\mathbf{u}'} - (\mathcal{A}(\mathbf{u}') - \mathcal{A}(\mathbf{u}))^\intercal s'_{\mathbf{u}'} && \text{[by Equation 2.18]} \\
&\leq |\mathcal{A}(\mathbf{u}') - \mathcal{A}(\mathbf{u})|\, |s_{\mathbf{u}'} - s'_{\mathbf{u}'}|, && \text{[by Cauchy-Schwartz inequality]}
\end{aligned}$$

or equivalently,

$$|\mathcal{A}(\mathbf{u}) - \mathcal{A}(\mathbf{u}')| \leq \frac{1}{\lambda}|s_{\mathbf{u}'} - s'_{\mathbf{u}'}|. \tag{2.19}$$

Notice again that this hold for any $s_{\mathbf{u}'} \in \partial_{\mathcal{A}(\mathbf{u}')} g(\mathbf{u}, \mathcal{A}(\mathbf{u}'))$ and any $s'_{\mathbf{u}'} \in \partial_{\mathcal{A}(\mathbf{u}')} g(\mathbf{u}', \mathcal{A}(\mathbf{u}'))$. Now, take

$$\begin{aligned}
\mathbf{w}_1^* &\in \arg\max_{\mathbf{w} \in C(\mathbf{u})} L(\mathbf{w}, \mathcal{A}(\mathbf{u}')) \\
\mathbf{w}_2^* &\in \arg\max_{\mathbf{w} \in C(\mathbf{u}')} L(\mathbf{w}, \mathcal{A}(\mathbf{u}')),
\end{aligned}$$

it's easy to show that

$$\begin{aligned}
\nabla_{\mathcal{A}(\mathbf{u}')} L(\mathbf{w}_1^*, \mathcal{A}(\mathbf{u}')) &\in \partial_{\mathcal{A}(\mathbf{u}')} g(\mathbf{u}, \mathcal{A}(\mathbf{u}')) \\
\nabla_{\mathcal{A}(\mathbf{u}')} L(\mathbf{w}_2^*, \mathcal{A}(\mathbf{u}')) &\in \partial_{\mathcal{A}(\mathbf{u}')} g(\mathbf{u}', \mathcal{A}(\mathbf{u}')).
\end{aligned}$$

By β-smoothness of $L(\mathbf{w}, \mathbf{u})$ in $\mathbf{w}$, we have Combining this with Equation 2.19, we get

$$\begin{aligned}
|\mathcal{A}(\mathbf{u}) - \mathcal{A}(\mathbf{u}')| &\leq \frac{1}{\lambda} \left|\nabla_{\mathcal{A}(\mathbf{u}')} L(\mathbf{w}_1^*, \mathcal{A}(\mathbf{u}')) - \nabla_{\mathcal{A}(\mathbf{u}')} L(\mathbf{w}_2^*, \mathcal{A}(\mathbf{u}'))\right| \\
&\leq \frac{\beta}{\lambda}|\mathbf{w}_1^* - \mathbf{w}_2^*| && \text{[by } \beta\text{-smoothness (Assumption 2)]} \\
&= \frac{\beta}{\lambda}\sqrt{\sum_{t=0}^{T-1} |Q_t(\mathbf{u}) - Q_t(\mathbf{u}')|^2} && \text{[by Assumption 4]} \\
&\leq \frac{\beta\sqrt{\sum_{t=0}^{T-1} \epsilon_t^2}}{\lambda}|\mathbf{u} - \mathbf{u}'|. && \text{[by } \epsilon_t\text{-Lipschitz continuity (Assumption 3)]}
\end{aligned}$$

Thus, when

$$\frac{\beta\sqrt{\sum_{t=0}^{T-1} \epsilon_t^2}}{\lambda} < 1,$$

by the Contraction Mapping Theorem (see, for example, [9, Appendix B]), the I-IRPC $u_{i+1} = \mathcal{A}(\mathbf{u}_i)$ converges to a unique performatively stable control $\mathbf{u}_{PS}$. Further, we also have

$$|\mathbf{u}_i - \mathbf{u}_{PS}| \leq \left(\frac{\beta\sqrt{\sum_{t=0}^{T-1} \epsilon_t^2}}{\lambda}\right)^i |\mathbf{u}_0 - \mathbf{u}_{PS}|,$$

which gives us the linear convergence rate. ☐

2.4.3 Relating Performatively Stable Control to Performatively Optimal Control

So far, we have identified the conditions that allow the I-IRPC to converges to a unique performatively stable control, $\mathbf{u}_{PS}$. However, what we really care about is the performatively optimal control, $\mathbf{u}_{PO}$, which is the solution to Problem 2.3. In this section, we will establish the sufficient condition that makes $\mathbf{u}_{PO}$ close to $\mathbf{u}_{PS}$. To do so, we first need to make the following additional assumption.

Assumption 5 (L_w-Lipschitz continuity in $\mathbf{w}$)**.** *We assume that $L(\mathbf{w}, \mathbf{u})$ is L_w-Lipschitz in $\mathbf{w}$, i.e. for any $\mathbf{w}, \mathbf{w}' \in \mathbb{R}^{nT}$ and $\mathbf{u} \in U$,*

$$|L(\mathbf{w}, \mathbf{u}) - L(\mathbf{w}', \mathbf{u})| \leq L_w |\mathbf{w} - \mathbf{w}'|.$$

Theorem 4 (Relating $\mathbf{u}_{PS}$ and $\mathbf{u}_{PO}$)**.** *In addition to assumptions made in Theorem 3, if Assumption 5 is also satisfied, then*

$$|\mathbf{u}_{PS} - \mathbf{u}_{PO}| \leq \frac{2L_w \sqrt{\sum_{t=0}^{T-1} \epsilon_t^2}}{\lambda}. \tag{2.20}$$

Proof. By the definition of $\mathbf{u}_{PO}$ and $\mathbf{u}_{PS}$, we have

$$g(\mathbf{u}_{PO}, \mathbf{u}_{PO}) \leq g(\mathbf{u}_{PS}, \mathbf{u}_{PS}) \leq g(\mathbf{u}_{PS}, \mathbf{u}_{PO}).$$

Since $\mathbf{u}_{PS}$ is a maximizer of $g(\mathbf{u}_{PS}, \mathbf{v})$, we know that $0 \in \partial_{\mathbf{v}} g(\mathbf{u}_{PS}, \mathbf{v})$. Therefore, by λ-strong convexity (Assumption 1), we have

$$g(\mathbf{u}_{PS}, \mathbf{u}_{PO}) - g(\mathbf{u}_{PS}, \mathbf{u}_{PS}) \geq \frac{\lambda}{2} |\mathbf{u}_{PS} - \mathbf{u}_{PO}|^2.$$

Further, let by Assumption 5, we also have

$$g(\mathbf{u}_{PS}, \mathbf{u}_{PO}) - g(\mathbf{u}_{PO}, \mathbf{u}_{PO}) \leq L_w |\mathbf{w}_1^* - \mathbf{w}_2^*|.$$

Combining above equations and applying Assumption 4, we get

$$|\mathbf{u}_{PS} - \mathbf{u}_{PO}|^2 \leq \left(\frac{2L_w}{\lambda} \sqrt{\sum_{t=0}^{T-1} \epsilon_t^2} \right) |\mathbf{u}_{PS} - \mathbf{u}_{PO}|.$$

Simplifying this completes the proof. □

2.4.4 Finite Sample Results

So far, we have focused our analysis on the I-IRPC, which is the ideal version of our proposed method. Now we are finally ready to bridge the gap between them, and establish the convergence properties of our E-IRPC algorithm. The key here is to view the radius of conformal sets obtained in Section 2.2.2 as the empirical approximation of the quantile function defined by Equation 2.11.

Definition 4 (Hazard rate)**.** Let X be a 1-dimensional random variable. The *hazard rate* of X evaluated at point x is defined by

$$h(x) = \lim_{\theta \to 0} \frac{\mathbb{P}(x \leq X \leq x + \theta | X \geq x)}{\theta} = \frac{f(x)}{1 - F(x)},$$

where f is the p.d.f (assuming that it exists) and F is the c.d.f.

Assumption 6 (Additional assumptions on the distribution of noise)**.** *We assume that for any* $\mathbf{u} \in U$*, we have*

- *the hazard rate of* $|w|$ *is positive and non-decreasing*
- *the p.d.f of* $|w|$ *exists and is continuously differentiable.*

This assumption is the combination of [17, Assumption 1 and 2].

Theorem 5 (Convergence of E-IRPC to a neighborhood of $\mathbf{u}_{PS}$)**.** *Suppose the following assumptions hold:*

- $L(\mathbf{w}, \mathbf{u})$ *is* λ*-strongly convex in* $\mathbf{u}$ *(Assumption 1),*
- $L(\mathbf{w}, \mathbf{u})$ *is* β*-smooth in* $\mathbf{w}$ *(Assumption 2),*
- *the quantile function* $Q_t(\mathbf{u})$ *is* ϵ_t*-Lipschitz in* $\mathbf{u}$ *for all* $t = 0, \ldots, T-1$ *with* $\sum_{t=1}^{T-1} \epsilon_t^2 < \frac{\lambda^2}{4\beta^2}$*, (Assumption 3),*
- *the worst case noises align (Assumption 4) and*
- *Assumption 6.*

Then the E-IRPC converges to the neighborhood of a unique performative stable control $\mathbf{u}_{PS}$ *at a linear rate. Specifically, for any* $\delta \in (0,1)$*, if we take* $N_i = \mathcal{O}\left(\frac{4\lambda^2 T^3}{\beta^2 \delta^2 \sum_{t=0}^{T-1} \epsilon_t^2} \log\left(\frac{6p}{\pi^2 i^2 T^2}\right)\right)$ *samples in each iteration, then with probability* $1-p$ *(*p *is the same as the overall probability level defined in Equation 2.2),*

$$|\mathbf{u}_i - \mathbf{u}_{PS}| \leq \delta \text{ for all } i \geq \left(1 - \frac{2\beta\sqrt{\sum_{t=0}^{T-1} \epsilon_t^2}}{\lambda}\right)^{-1} \log\left(\frac{|\mathbf{u}_0 - \mathbf{u}_{PS}|}{\delta}\right). \tag{2.21}$$

Proof. For each timestep t, take the N_i samples of $w(t)$ and sort their Euclidean norm in the increasing order as follows

$$\left|w^{(\eta_1)}(t)\right| \geq \cdots \geq \left|w^{(\eta_{N_i})}(t)\right|.$$

Let $k = \left\lfloor \frac{Np}{T} \right\rfloor$, then the empirical quantile function is given by

$$\hat{Q}_t(\mathbf{u}) := \inf\left\{ r : \frac{1}{N_i} \sum_{j=1}^{N_i} \mathbf{1}\left\{\left|w^j(t)\right| \leq r\right\} \geq 1 - \frac{p}{T} \right\} = \left|w^{(\eta_k)}(t)\right|. \tag{2.22}$$

and the conformal sets constructed in Section 2.2.2 is given by

$$\hat{C}_t(\mathbf{u}) = \left\{ w(t) : |w(t)| \leq \hat{Q}_t(\mathbf{u}) \right\},$$
$$\hat{C}_u = \hat{C}_0(\mathbf{u}) \times \cdots \times \hat{C}_{T-1}(\mathbf{u}).$$

Then for any $\mathbf{u}, \mathbf{u}' \in U$, we have

$$\left|\hat{Q}_t(\mathbf{u}) - Q_t(\mathbf{u}')\right| \leq \left|\hat{Q}_t(\mathbf{u}) - Q_t(\mathbf{u})\right| + |Q_t(\mathbf{u}) - Q_t(\mathbf{u}')|.$$

By [17, Theorem 2], for any $\epsilon > 0$, we have

$$\mathbb{P}\left(\left|\hat{Q}_t(\mathbf{u}) - Q_t(\mathbf{u})\right| \geq \epsilon\right) \leq \exp\left(-\frac{\epsilon^2}{2\left(v^r + (c^r + \omega_n)\,\epsilon\right)}\right) + \exp\left(-\frac{\epsilon^2}{2\left(v^l + \omega_n \epsilon\right)}\right)$$

where $v^r = \frac{2}{kL^2}$, $v^l = \frac{2(N_i-k+1)}{(k-1)^2L^2}$, $c^r = \frac{2}{kL}$ and $\omega_n = \frac{b}{N_i}$, with L and b finite for all $\mathbf{u} \in U$. With some algebraic manipulations, we get that with $N_i = \mathcal{O}\left(\frac{4\lambda^2T^3}{\beta^2\delta^2\sum_{t=0}^{T-1}\epsilon_t^2}\log\left(\frac{6p}{\pi^2i^2T^2}\right)\right)$,

$$\mathbb{P}\left(\left|\hat{Q}_t(\mathbf{u}_i) - Q_t(\mathbf{u}_i)\right| \geq \frac{\beta\sqrt{\sum_{t=0}^{T-1}\epsilon_t^2}}{\lambda T}\delta\right) \leq \frac{6p}{\pi^2i^2T}.$$

Following the same arguments as in the proof of Theorem 3, we get

$$|\hat{\mathcal{A}}(\mathbf{u}_i) - \mathcal{A}(\mathbf{u}_i)| \leq \frac{\beta}{\lambda}\sqrt{\sum_{t=0}^{T-1}\left|\hat{Q}_t(\mathbf{u}) - Q_t(\mathbf{u})\right|^2} \leq \frac{\delta\beta\sqrt{\sum_{t=0}^{T-1}\epsilon_t^2}}{\lambda}$$

with probability at least $1 - \frac{6p}{\pi^2i^2}$. Now, applying Theorem 3 we get that when $|\mathbf{u}_i - \mathbf{u}_{PS}| \geq \delta$,

$$\begin{aligned}
|\mathbf{u}_{i+1} - \mathbf{u}_{PS}| &= |\hat{\mathcal{A}}(\mathbf{u}_i) - \mathcal{A}(\mathbf{u}_{PS})| \\
&\leq |\hat{\mathcal{A}}(\mathbf{u}_i) - \mathcal{A}(\mathbf{u}_i)| + |\mathcal{A}(\mathbf{u}_i) - \mathcal{A}(\mathbf{u}_{PS})| \\
&\leq \frac{\beta\sqrt{\sum_{t=0}^{T-1}\epsilon_t^2}}{\lambda}|\mathbf{u}_i - \mathbf{u}_{PS}| + \frac{\beta\sqrt{\sum_{t=0}^{T-1}\epsilon_t^2}}{\lambda}\delta \\
&\leq \frac{2\beta\sqrt{\sum_{t=0}^{T-1}\epsilon_t^2}}{\lambda}|\mathbf{u}_i - \mathbf{u}_{PS}|.
\end{aligned}$$

Therefore, when

$$\frac{2\beta\sqrt{\sum_{t=0}^{T-1}\epsilon_t^2}}{\lambda},$$

applying the union bound, with probability at least $1 - \sum_{i=0}^{\infty}\frac{6p}{\pi^2i^2} = 1 - p$, for all $i = 0, 1, ...$ we have

$$|\mathbf{u}_i - \mathbf{u}_{PS}| \leq \max\left\{\left(\frac{2\beta\sqrt{\sum_{t=0}^{T-1}\epsilon_t^2}}{\lambda}\right)^i|\mathbf{u}_0 - \mathbf{u}_{PS}|, \delta\right\}.$$

This completes the proof. □

2.4.5 Relaxing the Alignment Constraint on the Noise

As one might notice, most of our results so far are based on the assumption that the worst-case noises *align* when we change the control low (Assumption 4). This assumption, however, is not easy to satisfy. In this section, we establish results based on the following weaker assumption.

Definition 5. Given the loss function L and the confidence set mapping C, we say that the worst-case noises *weakly align* if For any $\mathbf{v}, \mathbf{u}, \mathbf{u}' \in U$, there exists $\mathbf{w}_1^* \in \arg\max_{\mathbf{w}\in C(\mathbf{u})} L(\mathbf{w}, \mathbf{v})$ and $\mathbf{w}_2^* \in \arg\max_{\mathbf{w}\in C(\mathbf{u}')} L(\mathbf{w}, \mathbf{v})$ such that for any $t = 0, \ldots, T-1$,

$$(w_1^*(t) - w_2^*(t))^\intercal w_t^*(t) \geq 0$$

Assumption 7 (weak alignment of worst-case noises)**.** *We assume that the worst-case noises weakly align.*

Our first result concerns the convergence of the I-IRPC under the weak alignment condition (Assumption 7), which is the counterpart of Theorem 3.

Theorem 6 (Convergence of I-IRPC)**.** *Suppose the following assumptions hold:*

- $L(\mathbf{w}, \mathbf{u})$ *is* λ*-strongly convex in* $\mathbf{u}$ *(Assumption 1),*
- $L(\mathbf{w}, \mathbf{u})$ *is* β*-smooth in* $\mathbf{w}$ *(Assumption 2),*
- *the quantile function* $Q_t(\mathbf{u})$ *is* ϵ_t*-Lipschitz in* $\mathbf{u}$ *for all* $t = 0, \dots, T-1$ *with* $\sum_{t=1}^{T-1} \epsilon_t^2 < \frac{\lambda^2}{\beta^2}$ *(Assumption 3),*
- *worst case noises weakly align (Assumption 7),*

Then the I-IRPC converges to the neighborhood of a performative stable control $\mathbf{u}_{PS}$ *at a linear rate. Specifically, for any*

$$\delta > \frac{\sum_{t=1}^{T-1} 2\epsilon_t Q_t(\mathbf{u}_{PS})}{\lambda^2/\beta^2 - \sum_{t=1}^{T-1} \epsilon_t^2} + \frac{\sum_{t=0}^{T-1} \epsilon_t Q_t(\mathbf{u}_{PS})}{\sum_{t=0}^{T-1} \epsilon_t^2},$$

we have

$$|\mathbf{u}_i - \mathbf{u}_{PS}| \leq \delta \text{ for all } i \geq \left(1 - \frac{\beta \left(\sqrt{\sum_{t=0}^{T-1} \left(\epsilon_t^2 + 2\epsilon_t \frac{Q_t(\mathbf{u}_{PS})}{\delta}\right)}\right)}{\lambda}\right)^{-1} \log\left(\frac{|\mathbf{u}_0 - \mathbf{u}_{PS}|}{\delta}\right). \tag{2.23}$$

Proof. For any $\mathbf{u}, \mathbf{u}' \in U$, take

$$\mathbf{w}_1^* \in \arg\max_{\mathbf{w} \in C(\mathbf{u})} L(\mathbf{w}, \mathcal{A}(\mathbf{u}'))$$
$$\mathbf{w}_2^* \in \arg\max_{\mathbf{w} \in C(\mathbf{u}')} L(\mathbf{w}, \mathcal{A}(\mathbf{u}')),$$

following the same argument as the proof of Theorem 3 (using λ-strong convexity and β-smoothness), we have

$$\begin{aligned} |\mathcal{A}(\mathbf{u}) - \mathcal{A}(\mathbf{u}')| &\leq \frac{1}{\lambda} \left|\nabla_{\mathcal{A}(\mathbf{u}')} L(\mathbf{w}_1^*, \mathcal{A}(\mathbf{u}')) - \nabla_{\mathcal{A}(\mathbf{u}')} L(\mathbf{w}_2^*, \mathcal{A}(\mathbf{u}'))\right| \\ &\leq \frac{\beta}{\lambda} |\mathbf{w}_1^* - \mathbf{w}_2^*|. \end{aligned}$$

By weak alignment, we have

$$
\begin{aligned}
|\mathbf{w}_1^* - \mathbf{w}_2^*| &= \sqrt{\sum_{t=0}^{T-1} |w_1^*(t) - w_2^*(t)|^2} \\
&\le \sqrt{\sum_{t=0}^{T-1} |Q_t^2(\mathbf{u}) - Q_t^2(u')|} \\
&= \sqrt{\sum_{t=0}^{T-1} |Q_t(\mathbf{u}) - Q_t(\mathbf{u}')||Q_t(\mathbf{u}) + Q_t(\mathbf{u}')|} \\
&\le \sqrt{\sum_{t=0}^{T-1} \left(|Q_t(\mathbf{u}) - Q_t(\mathbf{u}')|^2 + 2|Q_t(\mathbf{u}) - Q_t(\mathbf{u}')|Q_t(\mathbf{u}')\right)} \\
&\le \sqrt{\sum_{t=0}^{T-1} (\epsilon_t^2|\mathbf{u} - \mathbf{u}'|^2 + 2\epsilon_t|\mathbf{u} - \mathbf{u}'|Q_t(\mathbf{u}'))} \\
&= \left(\sqrt{\sum_{t=0}^{T-1} \left(\epsilon_t^2 + 2\epsilon_t \frac{Q_t(\mathbf{u}')}{|\mathbf{u} - \mathbf{u}'|}\right)}\right) |\mathbf{u} - \mathbf{u}'|.
\end{aligned}
\tag{2.24}
$$

Thus, at the i-th iteration of I-IRPC, we have

- When $|\mathbf{u}_i - \mathbf{u}_{PS}| > \delta$,

$$
\begin{aligned}
|\mathbf{u}_{i+1} - \mathbf{u}_{PS}| &= |\mathcal{A}(\mathbf{u}_i) - \mathcal{A}(\mathbf{u}_{PS})| \\
&\le \frac{\beta\left(\sqrt{\sum_{t=0}^{T-1}\left(\epsilon_t^2 + 2\epsilon_t\frac{Q_t(\mathbf{u}_{PS})}{\delta}\right)}\right)}{\lambda} |\mathbf{u}_i - \mathbf{u}_{PS}|.
\end{aligned}
$$

- When $|\mathbf{u}_i - \mathbf{u}_{PS}| \le \delta$,

$$
\begin{aligned}
|\mathbf{u}_{i+1} - \mathbf{u}_{PS}| &= |\mathcal{A}(\mathbf{u}_i) - \mathcal{A}(\mathbf{u}_{PS})| \\
&\le \frac{\beta}{\lambda}\sqrt{\sum_{t=0}^{T-1} (\epsilon_t^2\delta^2 + 2\epsilon_t\delta Q_t(\mathbf{u}_{PS}))} \\
&= \frac{\beta}{\lambda}\sqrt{\left(\sqrt{\sum_{t=0}^{T-1}\epsilon_t^2}\delta + \frac{\sum_{t=0}^{T-1}\epsilon_t Q_t(\mathbf{u}_{PS})}{\sqrt{\sum_{t=0}^{T-1}\epsilon_t^2}}\right)^2 - \frac{\left(\sum_{t=0}^{T-1}\epsilon_t Q_t(\mathbf{u}_{PS})\right)^2}{\sum_{t=0}^{T-1}\epsilon_t^2}} \\
&\le \frac{\beta\sqrt{\sum_{t=0}^{T-1}\epsilon_t^2}}{\lambda}\left(\delta + \frac{\sum_{t=0}^{T-1}\epsilon_t Q_t(\mathbf{u}_{PS})}{\sum_{t=0}^{T-1}\epsilon_t^2}\right) \\
&\le \delta + \frac{\sum_{t=0}^{T-1}\epsilon_t Q_t(\mathbf{u}_{PS})}{\sum_{t=0}^{T-1}\epsilon_t^2}.
\end{aligned}
$$

When $\sum_{t=1}^{T-1}\epsilon_t^2 < \frac{\lambda^2}{\beta^2}$, we can pick

$$
\delta > \frac{\sum_{t=1}^{T-1} 2\epsilon_t Q_t(\mathbf{u}_{PS})}{\lambda^2/\beta^2 - \sum_{t=1}^{T-1}\epsilon_t^2} > 0,
$$

to make $\mathcal{A}(\mathbf{u})$ contractive for all $|u - \mathbf{u}_{PS}| > \delta$. This gives us

$$|\mathbf{u}_i - \mathbf{u}_{PS}| \leq \max\left\{\left(\frac{\beta\left(\sqrt{\sum_{t=0}^{T-1}\left(\epsilon_t^2 + 2\epsilon_t\frac{Q_t(\mathbf{u}_{PS})}{\delta}\right)}\right)}{\lambda}\right)^i |\mathbf{u}_0 - \mathbf{u}_{PS}|, \delta + \frac{\sum_{t=0}^{T-1}\epsilon_t Q_t(\mathbf{u}_{PS})}{\sum_{t=0}^{T-1}\epsilon_t^2}\right\},$$

Rearranging this completes the proof. □

Note that this is weaker than Theorem 3 in two ways. First, it does not guarantee the uniqueness of the $\mathbf{u}_{PS}$. In fact, even the existence of $\mathbf{u}_{PS}$ does not come for free—it relies on Theorem 1. Second, we cannot get arbitrarily close to $\mathbf{u}_{PS}$. There is always a gap.

Our second result concerns the relationship between $\mathbf{u}_{PS}$ and $\mathbf{u}_{PO}$ under the weak alignment condition (Assumption 7), which is the counterpart of Theorem 4.

Theorem 7 (Relating $\mathbf{u}_{PS}$ and $\mathbf{u}_{PO}$)**.** *In addition to assumptions made in Theorem 6, if Assumption 5 is also satisfied, then for any $\delta > 0$, we have*

$$|\mathbf{u}_{PO} - \mathbf{u}_{PS}| \leq \max\left\{\frac{2L_w\sqrt{\sum_{t=0}^{T-1}\left(\epsilon_t^2 + 2\epsilon_t\frac{Q_t(\mathbf{u}_{PS})}{\delta}\right)}}{\lambda}, \delta\right\}.$$

Proof. Following the same argument as the proof of Theorem 4, we have

$$\frac{\lambda}{2}|\mathbf{u}_{PO} - \mathbf{u}_{PS}|^2 \leq g(\mathbf{u}_{PS}, \mathbf{u}_{PO}) - g(\mathbf{u}_{PS}, \mathbf{u}_{PS}) \leq g(\mathbf{u}_{PS}, \mathbf{u}_{PO}) - g(\mathbf{u}_{PO}, \mathbf{u}_{PO}) \leq L_w|\mathbf{w}_1^* - \mathbf{w}_2^*|,$$

where

$$\mathbf{w}_1^* \in \underset{\mathbf{w}\in C(\mathbf{u}_{PS})}{\arg\max}\ L(\mathbf{w}, \mathcal{A}(\mathbf{u}_{PO}))$$
$$\mathbf{w}_2^* \in \underset{\mathbf{w}\in C(\mathbf{u}_{PO})}{\arg\max}\ L(\mathbf{w}, \mathcal{A}(\mathbf{u}_{PO})).$$

By weak alignment (as in Equation 2.24), we have

$$|\mathbf{w}_1^* - \mathbf{w}_2^*| \leq \left(\sqrt{\sum_{t=0}^{T-1}\left(\epsilon_t^2 + 2\epsilon_t\frac{Q_t(\mathbf{u}_{PS})}{|\mathbf{u}_{PO} - \mathbf{u}_{PS}|}\right)}\right)|\mathbf{u}_{PO} - \mathbf{u}_{PS}|.$$

combining this with Equation 2.4.5 which gives us

$$|\mathbf{u}_{PO} - \mathbf{u}_{PS}| \leq \frac{2L_w}{\lambda}\sqrt{\sum_{t=0}^{T-1}\left(\epsilon_t^2 + 2\epsilon_t\frac{Q_t(\mathbf{u}_{PS})}{|\mathbf{u}_{PO} - \mathbf{u}_{PS}|}\right)}.$$

Therefore for any $\delta > 0$, we have

$$|\mathbf{u}_{PO} - \mathbf{u}_{PS}| \leq \max\left\{\frac{2L_w\sqrt{\sum_{t=0}^{T-1}\left(\epsilon_t^2 + 2\epsilon_t\frac{Q_t(\mathbf{u}_{PS})}{\delta}\right)}}{\lambda}, \delta\right\}.$$

□

Our third result concerns the convergence of the E-IRPC under the weak alignment condition (Assumption 7), which is the counterpart of Theorem 5.

Theorem 8 (Convergence of E-IRPC)**.** *Suppose the following assumptions hold:*

- $L(\mathbf{w}, \mathbf{u})$ *is* λ*-strongly convex in* $\mathbf{u}$ *(Assumption 1),*
- $L(\mathbf{w}, \mathbf{u})$ *is* β*-smooth in* $\mathbf{w}$ *(Assumption 2),*
- *the quantile function* $Q_t(\mathbf{u})$ *is* ϵ_t*-Lipschitz in* $\mathbf{u}$ *for all* $t = 0, \ldots, T-1$ *with* $\sum_{t=1}^{T-1} \epsilon_t^2 < \frac{\lambda^2}{4\beta^2}$ *(Assumption 3),*
- *the worst case noises weakly align (Assumption 7) and*
- *Assumption 6.*

Then the E-IRPC converges to the neighborhood of a performative stable control $\mathbf{u}_{PS}$ *at a linear rate. Specifically, for any*

$$\delta > \frac{\sum_{t=1}^{T-1} 2\epsilon_t Q_t(\mathbf{u}_{PS})}{\lambda^2/\beta^2 - \sum_{t=1}^{T-1} \epsilon_t^2} + \frac{\sum_{t=0}^{T-1} \epsilon_t Q_t(\mathbf{u}_{PS})}{\sum_{t=0}^{T-1} \epsilon_t^2},$$

if we take $N_i = \mathcal{O}\left(\frac{4\lambda^2 T^3}{\beta^2 \sum_{t=0}^{T-1}(\epsilon_t^2 \delta + 2\epsilon_t Q_t(\mathbf{u}_{PS}))} \log\left(\frac{6p}{\pi^2 i^2 T^2}\right)\right)$ *samples in each iteration, then with probability* $1-p$ *(*p *is the same as the overall probability level defined in Equation 2.2),*

$$|\mathbf{u}_i - \mathbf{u}_{PS}| \le \delta \text{ for all } i \ge \left(1 - \frac{\beta\left(\sqrt{\sum_{t=0}^{T-1}\left(\epsilon_t^2 + 2\epsilon_t \frac{Q_t(\mathbf{u}_{PS})}{\delta}\right)}\right)}{\lambda}\right)^{-1} \log\left(\frac{|\mathbf{u}_0 - \mathbf{u}_{PS}|}{\delta}\right). \tag{2.25}$$

Proof. This proof builds on the proof of Theorem 6, in the same way as the proof of Theorem 5 builds on that of Theorem 3. Specifically, by [17, Theorem 2], for any $\epsilon > 0$, we have

$$\mathbb{P}\left(\left|\hat{Q}_t(\mathbf{u}) - Q_t(\mathbf{u})\right| \ge \epsilon\right) \le \exp\left(-\frac{\epsilon^2}{2\left(v^r + \left(c^r + \omega_n\right)\epsilon\right)}\right) + \exp\left(-\frac{\epsilon^2}{2\left(v^l + \omega_n \epsilon\right)}\right)$$

where $v^r = \frac{2}{kL^2}$, $v^l = \frac{2(N_i - k + 1)}{(k-1)^2 L^2}$, $c^r = \frac{2}{kL}$ and $\omega_n = \frac{b}{N_i}$, with L and b finite for all $\mathbf{u} \in U$. Now pick any

$$\delta > \frac{\sum_{t=1}^{T-1} 2\epsilon_t Q_t(\mathbf{u}_{PS})}{\lambda^2/\beta^2 - \sum_{t=1}^{T-1} \epsilon_t^2} + \frac{\sum_{t=0}^{T-1} \epsilon_t Q_t(\mathbf{u}_{PS})}{\sum_{t=0}^{T-1} \epsilon_t^2},$$

as in Theorem 6, with some algebraic manipulations, we get that with $N_i = \mathcal{O}\left(\frac{4\lambda^2 T^3}{\beta^2 \sum_{t=0}^{T-1}(\epsilon_t^2 \delta + 2\epsilon_t Q_t(\mathbf{u}_{PS}))} \log\left(\frac{6p}{\pi^2 i^2 T^2}\right)\right)$,

$$\mathbb{P}\left(\left|\hat{Q}_t(\mathbf{u}_i) - Q_t(\mathbf{u}_i)\right| \ge \frac{\beta\left(\sqrt{\sum_{t=0}^{T-1}\left(\epsilon_t^2 + 2\epsilon_t \frac{Q_t(\mathbf{u}_{PS})}{\delta}\right)}\right)}{\lambda}\delta\right) \le \frac{6p}{\pi^2 i^2 T}.$$

Following the same arguments as in the proof of Theorem 6, we get

$$|\hat{\mathcal{A}}(\mathbf{u}_i) - \mathcal{A}(\mathbf{u}_i)| \leq \frac{\beta}{\lambda}\sqrt{\sum_{t=0}^{T-1}\left|\hat{Q}_t(\mathbf{u}) - Q_t(\mathbf{u})\right|^2} \leq \frac{\beta\left(\sqrt{\sum_{t=0}^{T-1}\left(\epsilon_t^2 + 2\epsilon_t\frac{Q_t(\mathbf{u}_{PS})}{\delta}\right)}\right)}{\lambda}\delta$$

with probability at least $1 - \frac{6p}{\pi^2 i^2}$. Now, applying Theorem 6 we get that when $|\mathbf{u}_i - \mathbf{u}_{PS}| \geq \delta$,

$$\begin{aligned}
|\mathbf{u}_{i+1} - \mathbf{u}_{PS}| &= |\hat{\mathcal{A}}(\mathbf{u}_i) - \mathcal{A}(\mathbf{u}_{PS})| \\
&\leq |\hat{\mathcal{A}}(\mathbf{u}_i) - \mathcal{A}(\mathbf{u}_i)| + |\mathcal{A}(\mathbf{u}_i) - \mathcal{A}(\mathbf{u}_{PS})| \\
&\leq \frac{\beta\left(\sqrt{\sum_{t=0}^{T-1}\left(\epsilon_t^2 + 2\epsilon_t\frac{Q_t(\mathbf{u}_{PS})}{\delta}\right)}\right)}{\lambda}|\mathbf{u}_i - \mathbf{u}_{PS}| + \frac{\beta\left(\sqrt{\sum_{t=0}^{T-1}\left(\epsilon_t^2 + 2\epsilon_t\frac{Q_t(\mathbf{u}_{PS})}{\delta}\right)}\right)}{\lambda}\delta \\
&\leq \frac{2\beta\left(\sqrt{\sum_{t=0}^{T-1}\left(\epsilon_t^2 + 2\epsilon_t\frac{Q_t(\mathbf{u}_{PS})}{\delta}\right)}\right)}{\lambda}|\mathbf{u}_i - \mathbf{u}_{PS}|.
\end{aligned}$$

Therefore, when

$$\frac{2\beta\left(\sqrt{\sum_{t=0}^{T-1}\left(\epsilon_t^2 + 2\epsilon_t\frac{Q_t(\mathbf{u}_{PS})}{\delta}\right)}\right)}{\lambda} < 1,$$

applying the union bound, with probability at least $1 - \sum_{i=0}^{\infty}\frac{6p}{\pi^2 i^2} = 1 - p$, for all $i = 0, 1, ...$ we have

$$|\mathbf{u}_i - \mathbf{u}_{PS}| \leq \max\left\{\left(\frac{2\beta\left(\sqrt{\sum_{t=0}^{T-1}\left(\epsilon_t^2 + 2\epsilon_t\frac{Q_t(\mathbf{u}_{PS})}{\delta}\right)}\right)}{\lambda}\right)^i |\mathbf{u}_0 - \mathbf{u}_{PS}|, \delta\right\}.$$

This completes the proof. □

Chapter 3

Conclusion and Further Discussions

In this study, we present a novel way to approach noisy dynamical systems, particularly those characterized by performative noise. Drawing inspiration from the literature on performative prediction and conformal prediction, we tailored an approach that addresses the challenges posed by real-world noisy discrete-dynamical systems where the characteristics of the noise remain largely unknown.

Our proposed method, E-IRPC, is applicable in settings where we do not have explicit knowledge of the noise distribution. Leveraging empirical observations with conformal prediction techniques, we've showcased an adaptive strategy that converges to near-optimal open-loop controls, while remaining attuned to the system's inherent noise dynamics. By investigating its infinite sample limit, I-IRPC, we are able to derive rigorous convergence bounds under certain regularity conditions, offering a theoretical foundation for future investigation and experiments.

As we reflect on our journey, it is evident that the challenges posed by performative noise in control systems are not insurmountable. With the right blend of theory, empirical observation, and innovative methodologies, we can navigate the complexities of such systems, ensuring stability and optimal performance. While our current focus has been on open-loop controls, our approach holds immense promise for various extensions including closed-loop control.

References

[1] A. K. Akametalu, S. Kaynama, J. F. Fisac, M. N. Zeilinger, J. H. Gillula, and C. J. Tomlin, "Reachability-Based Safe Learning with Gaussian Processes," *53rd IEEE Conference on Decision and Control*, pp. 1424–1431, 2014.

[2] A. N. Angelopoulos and S. Bates, *A Gentle Introduction to Conformal Prediction and Distribution-Free Uncertainty Quantification*, 2022. arXiv: 2107.07511 [cs.LG].

[3] J. Austin, *How to Do Things with Words*. Clarendon Press, 1962.

[4] F. Cairoli, L. Bortolussi, and N. Paoletti, *Neural Predictive Monitoring under Partial Observability*, 2021. arXiv: 2108.07134 [cs.LG].

[5] Y. Chen, U. Rosolia, C. Fan, A. D. Ames, and R. Murray, *Reactive motion planning with probabilistic safety guarantees*, 2020. arXiv: 2011.03590 [cs.RO].

[6] T. G. Dietterich and J. Hostetler, *Conformal Prediction Intervals for Markov Decision Process Trajectories*, 2022. arXiv: 2206.04860 [cs.LG].

[7] C. Fan, X. Qin, Y. Xia, A. Zutshi, and J. Deshmukh, *Statistical Verification of Autonomous Systems using Surrogate Models and Conformal Inference*, 2021. arXiv: 2004.00279 [eess.SY].

[8] J. H. Gillula and C. J. Tomlin, "Reducing Conservativeness in Safety Guarantees by Learning Disturbances Online: Iterated Guaranteed Safe Online Learning," in *Robotics: Science and Systems VIII*, The MIT Press, Jul. 2013.

[9] H. Khalil, *Nonlinear Systems*. Prentice Hall, 2002.

[10] L. Lindemann, M. Cleaveland, G. Shim, and G. J. Pappas, "Safe Planning in Dynamic Environments using Conformal Prediction," *IEEE Robotics and Automation Letters*, 2023.

[11] A. Muthali, H. Shen, S. Deglurkar, *et al.*, *Multi-Agent Reachability Calibration with Conformal Prediction*, 2023. arXiv: 2304.00432 [eess.SY].

[12] E. Ok, *Real Analysis with Economic Applications*. Princeton University Press, 2007.

[13] J. Perdomo, T. Zrnic, C. Mendler-Dünner, and M. Hardt, "Performative Prediction," in *International Conference on Machine Learning*, PMLR, 2020, pp. 7599–7609.

[14] S. Prajna, A. Jadbabaie, and G. J. Pappas, "A Framework for Worst-Case and Stochastic Safety Verification Using Barrier Certificates," *IEEE Transactions on Automatic Control*, vol. 52, no. 8, pp. 1415–1428, 2007.

[15] V. Vovk, A. Gammerman, and G. Shafer, *Algorithmic Learning in a Random World*, 2nd ed. Springer International Publishing, 2022.

[16] S. Yang, G. J. Pappas, R. Mangharam, and L. Lindemann, "Safe Perception-Based Control under Stochastic Sensor Uncertainty using Conformal Prediction," *arXiv preprint arXiv:2304.00194*, 2023.

[17] M. Zhang and C. S. Ong, "Quantile Bandits for Best Arms Identification," in *International Conference on Machine Learning*, PMLR, 2021, pp. 12 513–12 523.