Distributionally Robust Chance Constrained Data-enabled Predictive Control

Journal Article

Author(s):

Coulson, Jeremy; Lygeros, John ; Dörfler, Florian

Publication date:

2022-07

Permanent link:

<https://doi.org/10.3929/ethz-b-000496838>

Rights / license:

[In Copyright - Non-Commercial Use Permitted](#)

Originally published in:

IEEE Transactions on Automatic Control 67(7), <https://doi.org/10.1109/TAC.2021.3097706>

Funding acknowledgement:

787845 - Optimal control at large (EC)

Distributionally Robust Chance Constrained Data-enabled Predictive Control

Jeremy Coulson

John Lygeros

Florian Dörfler

Accepted for publication: IEEE Transactions on Automatic Control. (DOI: 10.1109/TAC.2021.3097706) © 2021 IEEE.

Abstract—We study the problem of finite-time constrained optimal control of unknown stochastic linear time-invariant systems, which is the key ingredient of a predictive control algorithm – albeit typically having access to a model. We propose a novel distributionally robust data-enabled predictive control (DeePC) algorithm which uses noise-corrupted input/output data to predict future trajectories and compute optimal control inputs while satisfying output chance constraints. The algorithm is based on (i) a non-parametric representation of the subspace spanning the system behaviour, where past trajectories are sorted in Page or Hankel matrices; and (ii) a distributionally robust optimization formulation which gives rise to strong probabilistic performance guarantees. We show that for certain objective functions, DeePC exhibits strong out-of-sample performance, and at the same time respects constraints with high probability. The algorithm provides an end-to-end approach to control design for unknown stochastic linear time-invariant systems. We illustrate the closed-loop performance of the DeePC in an aerial robotics case study.

Index Terms—Data-driven control, predictive control, distributionally robust optimization.

I. INTRODUCTION

OPTIMAL control of unknown systems can be approached in two ways: model-based and data-driven. In model-based control, a predictive model for the system of interest is first identified from data and subsequently used for control design. What have come to be known as “data-driven methods”, on the other hand, aim to design controllers directly from data, without explicitly identifying a predictive model. These methods are suitable for applications where first-principle models are not conceivable (e.g., in human-in-the-loop applications), when models are too complex for control design (e.g., in fluid dynamics), and when thorough modelling and parameter identification is too costly (e.g., in robotics). Data-driven control has recently gained a lot of popularity, but most methods cannot be applied (respectively, lack formal certificates) for real-time control of safety-critical systems.

In this work, we focus on a data-driven control technique for unknown, stochastic, and constrained linear systems. In particular, we present a method for finite-horizon optimal predictive control using input/output data from the unknown system, where the system behaviour is characterized by a data matrix time series. This method was first presented for deterministic systems in [1] and was later extended to stochastic systems in [2]. These works were motivated by [3] in which a unique method of direct data-driven open-loop control was conceived based on the seminal work on behavioural systems theory [4].

All authors are with the Department of Information Technology and Electrical Engineering at ETH Zürich, Switzerland {jcoulson, lygeros, dorfler}@control.ee.ethz.ch. The research was supported by the ERC under project OCAL, grant agreement 787845, and ETH Zürich funds.

A key challenge for such data-driven control methods is ensuring the performance and safety of the system in the presence of uncertainties, corrupted data, and noise. In this work, we present an end-to-end optimal data-driven control approach which comes with such guarantees but without explicitly identifying a predictive model and is agnostic to the particular probabilistic uncertainty. The approach combines the so-called data-enabled predictive control (DeePC) method from [1], [2] and distributionally robust optimization techniques from [5].

Data-driven control has been historically approached using, e.g., iterative feedback tuning and virtual reference feedback tuning [6], [7]. More recently the default approach is often reinforcement learning. There are many approaches to reinforcement learning, many of which are episodic, where learning and control alternate; see [8] for a review of methods and challenges. Here, we follow different lines of literature.

Behavioural Framework for Control: Instead of learning a parametric system representation, one can describe the entire subspace of possible trajectories of a linear time-invariant (LTI) system only in terms of raw data sorted into a Hankel matrix. This result became known as the fundamental lemma [4], has been inspired by subspace system identification [9], and was first leveraged in [3] for the computation of open-loop control and for simulating system responses. The result was used in [1] in which a predictive control algorithm was proposed and was extended for stochastic systems in [2]. Robust closed-loop stability guarantees have been provided in [10]. Additionally, numerical case studies have illustrated that the algorithm performs robustly on some stochastic and weakly nonlinear systems and often outperforms system identification followed by conventional model predictive control (MPC) [11]–[13]. The behavioural framework was used in [14] to construct explicit feedback controllers. Since then, the behavioural framework has become popular for control design giving rise to numerous methods [15]–[20].

Learning-based MPC: In learning-based MPC the unknown system dynamics are substituted with a learned model mapping inputs to output predictions; see [21] for a comprehensive survey. A learning MPC approach for iterative tasks is presented in [22]. Other approaches conceptually related to ours exploit previously measured trajectories (i.e., motion primitives or trajectory dictionaries) based on which they synthesize new trajectories [23]–[25]. These methods, however, require learning a dictionary of libraries that best fits a dataset, but do not take into account the control objective.

Sequential System Identification and Control: System identification (ID) produces a nominal model and an uncertainty estimate allowing for robust control design. Most approaches offer asymptotic guarantees, but some provide finite sample

certificates [26]–[29]. In this spirit, an end-to-end ID and control pipeline is given in [30], [31] and arrives at a data-driven control solution with guarantees on the sample efficiency, stability, performance, and robustness. We refer also to the earlier surveys on identification for control [32], [33]. Well-known short-comings of sequential ID and control are as follows [32]–[35]: The system ID step is known to be the most time consuming part of model-based control design. Furthermore, most system ID techniques seek the model that best fits the data, but do not take into account the control objective, possibly leading to unsatisfactory performance. Finally, practitioners often demand end-to-end automated solutions. Lastly, we mention approaches related to subspace identification and (in hindsight) also to the behavioural perspective: [36] derives a linear model from a Hankel matrix to be used for predictive control (see also [37] for an overview of subspace methods for identification and predictive control). Connections between indirect data-driven control (sequential system ID and control) and direct data-driven control are discussed in [38].

Stochastic MPC: To account for the stochastic uncertainty in the system dynamics, stochastic MPC approaches usually consider minimizing the expectation of the objective function, while satisfying chance constraints [39]. While chance constraints may be unsuitable for applications where constraint violation has catastrophic consequences, they are often used when transient violations are allowed or when violations are associated with financial penalties. Closest to our work is [40] in which a stochastic MPC approach is developed for the case when the probability distribution of the stochastic disturbance is unknown. The authors consider a distributionally robust approach to safe-guard against the unknown distribution assuming that an accurate system model is given.

The approach of this paper follows our earlier work [1], [2]. Our contributions are as follows:

Distributionally robust data-driven control formulation: We formulate a novel distributionally robust data-enabled predictive control problem based on behavioural systems theory and distributionally robust optimization.

Probabilistic guarantees on performance: We show that a distributionally robust optimal control problem admits a tractable reformulation, and with high confidence, its solutions exhibit strong out-of-sample guarantees. That is, we prove that the optimal value of the tractable formulation is an upper confidence bound on the out-of-sample performance of an optimal worst-case solution. We provide sample complexity results, and are able to leverage additional data measurements in comparison to [2]. Furthermore, the distributional nature of the robustness means that the method is robust against a set of systems compatible with the data collected, which can include non-Gaussian and non-additive noise as well as (weakly) nonlinear systems.

Safety through chance constraint satisfaction: We enforce chance constraints on the outputs of the system using a distributionally robust conditional value-at-risk formulation. We provide a tractable reformulation of the chance constraints and provide sample complexity results such that they hold with high confidence and are agnostic to the underlying probabilistic uncertainty.

Tight guarantees due to new data structure: We provide a novel formulation of the so-called fundamental lemma from behavioural system theory by proving that the subspace of trajectories of an LTI system can be spanned by trajectories organized in a Page matrix. This is in contrast with all of the literature leveraging the behavioural framework for control in which a Hankel matrix is used. In particular, this provides a novel contribution relative to [1], [2] as well as the body of literature leveraging the behavioural framework for control. Furthermore, we show that this alternative matrix structure gives tighter guarantees and results in better performance. The guarantees provided are for the finite-horizon optimal control problem. Providing guarantees for a closed-loop receding horizon implementation is a subject for future work. We refer to [10] for related results in this direction.

Section II reviews preliminaries on behavioural systems and presents non-parametric representations. In Section III, we present the data-enabled predictive control algorithm for deterministic systems. In Section IV we present a distributionally robust data-enabled predictive control algorithm for stochastic systems. In Section V, we illustrate the main results in simulation on a quadcopter and present a detailed analysis of the hyperparameters involved. We make concluding remarks in Section VI. All proofs have been deferred to the Appendix.

Notation: We denote by $\mathbb{Z}_{\geq 0}$, and $\mathbb{Z}_{> 0}$ the set of non-negative and positive integers respectively. Given $x, y \in \mathbb{R}^n$, $\langle x, y \rangle := x^\top y$ denotes the usual inner product on $\mathbb{R}^n$. We denote the associated dual norm of a norm $\|\cdot\|$ on $\mathbb{R}^n$ by $\|x\|_* := \sup_{\|y\| \leq 1} \langle x, y \rangle$. The convex conjugate of a function $f: X \rightarrow \mathbb{R}$ is denoted by $f^*(\theta) := \sup_{x \in X} \langle \theta, x \rangle - f(x)$. We define the positive part of a real-valued function f as $f_+(x) := \max\{f(x), 0\}$. We denote by δ_x the Dirac distribution at x . Given a signal $u: \mathbb{Z} \rightarrow \mathbb{R}^m$, we denote the restriction of the signal to an interval by $u_{[1, T]} := \text{col}(u_1, \dots, u_T)$, where $\text{col}(u_1, \dots, u_T)$ denotes the stacked column vector $(u_1^\top, \dots, u_T^\top)^\top$. We use the $\hat{\cdot}$ symbol to denote recorded data samples and to indicate that objects depend on data samples.

II. BEHAVIOURAL SYSTEMS

A. Preliminaries and Notation

Behavioural system theory is a natural way of viewing a dynamical system when one is not concerned with a particular system representation. This is in contrast to classical system theory, where a particular parametric system representation (such as a state-space model) is used to describe the behaviour, and system properties are derived by studying the chosen representation. Following [41], we define a dynamical system and its properties in terms of its behaviour.

Definition 2.1: A *dynamical system* is a 3-tuple $(\mathbb{Z}_{\geq 0}, \mathbb{W}, \mathcal{B})$ where $\mathbb{Z}_{\geq 0}$ is the discrete-time axis, $\mathbb{W}$ is a signal space, and $\mathcal{B} \subseteq \mathbb{W}^{\mathbb{Z}_{\geq 0}}$ is the behaviour.

Definition 2.2: Let $(\mathbb{Z}_{\geq 0}, \mathbb{W}, \mathcal{B})$ be a dynamical system.

- (i) $(\mathbb{Z}_{\geq 0}, \mathbb{W}, \mathcal{B})$ is *linear* if $\mathbb{W}$ is a vector space and $\mathcal{B}$ is a linear subspace of $\mathbb{W}^{\mathbb{Z}_{\geq 0}}$.
- (ii) $(\mathbb{Z}_{\geq 0}, \mathbb{W}, \mathcal{B})$ is *time invariant* if $\mathcal{B} \subseteq \sigma \mathcal{B}$ where σ is the backward time shift defined by $(\sigma w)(t) = w(t+1)$ and $\sigma \mathcal{B} = \{\sigma w \mid w \in \mathcal{B}\}$.

- (iii) $(\mathbb{Z}_{\geq 0}, \mathbb{W}, \mathcal{B})$ is *complete* if $\mathcal{B}$ is closed in the topology of pointwise convergence.

Note that if a dynamical system satisfies (i)-(ii) then (iii) is equivalent to finite dimensionality of $\mathbb{W}$ (see [41, Section 7.1]). We denote the class of systems $(\mathbb{Z}_{\geq 0}, \mathbb{R}^{m+p}, \mathcal{B})$ satisfying (i)-(iii) by $\mathcal{L}^{m+p}$, where $m, p \in \mathbb{Z}_{\geq 0}$. With slight abuse of notation and terminology, we denote a dynamical system in $\mathcal{L}^{m+p}$ only by its behaviour $\mathcal{B}$.

Definition 2.3: The *restricted behaviour in the interval* $[1, T]$ is the set $\mathcal{B}_T = \{w \in (\mathbb{R}^{m+p})^T \mid \exists v \in \mathcal{B} \text{ s.t. } w_t = v_t, 1 \leq t \leq T\}$. A vector $w \in \mathcal{B}_T$ is called a *T-length trajectory of the dynamical system* $\mathcal{B}$.

Definition 2.4: A system $\mathcal{B} \in \mathcal{L}^{m+p}$ is *controllable* if for every $T \in \mathbb{Z}_{>0}$, $w^1 \in \mathcal{B}_T$, $w^2 \in \mathcal{B}$ there exists $w \in \mathcal{B}$ and $T' \in \mathbb{Z}_{>0}$ such that $w_t = w_t^1$ for $1 \leq t \leq T$ and $w_t = w_{t-T-T'}^2$ for $t > T + T'$.

In other words, a behavioural system is controllable if any two trajectories can be patched together in finite time.

B. Parametric system representation

Without loss of generality, any trajectory $w \in \mathcal{B}$ can be written as $w = \text{col}(u, y)$, where $\text{col}(u, y) := (u^\top, y^\top)^\top$ (see [42, Theorem 2]). In what follows, we will associate u and y with inputs and outputs. There are several equivalent ways of representing a behavioural system $\mathcal{B} \in \mathcal{L}^{m+p}$, including the classical input/output/state representation denoted by $\mathcal{B}(A, B, C, D) = \{\text{col}(u, y) \in (\mathbb{R}^{m+p})^{\mathbb{Z}_{\geq 0}} \mid \exists x \in (\mathbb{R}^n)^{\mathbb{Z}_{\geq 0}} \text{ s.t. } \sigma x = Ax + Bu, y = Cx + Du\}$. The input/output/state representation of smallest order (i.e., smallest state dimension) is called a *minimal representation*, and we denote its order by $n(\mathcal{B})$. Another important property of a system $\mathcal{B} \in \mathcal{L}^{m+p}$ is the *lag* defined by the smallest integer $\ell \in \mathbb{Z}_{>0}$ such that the observability matrix $\mathcal{O}_\ell(A, C) := \text{col}(C, CA, \dots, CA^{\ell-1})$ has rank $n(\mathcal{B})$. We denote the lag by $\ell(\mathcal{B})$ (see [41, Section 7.2] for equivalent definitions). The lower triangular Toeplitz (impulse response) matrix consisting of $T_f \in \mathbb{Z}_{>0}$ system Markov parameters is denoted by

$$\mathcal{T}_{T_f}(A, B, C, D) := \begin{pmatrix} D & 0 & \cdots & 0 \\ CB & D & \cdots & 0 \\ \vdots & \ddots & \ddots & \vdots \\ CA^{T_f-2}B & \cdots & CB & D \end{pmatrix}.$$

Lemma 2.1: ([3, Lemma 1]): Let $\mathcal{B} \in \mathcal{L}^{m+p}$ and $\mathcal{B}(A, B, C, D)$ a minimal input/output/state representation. Let $T_{\text{ini}}, T_f \in \mathbb{Z}_{>0}$ with $T_{\text{ini}} \geq \ell(\mathcal{B})$ and $\text{col}(u_{\text{ini}}, u, y_{\text{ini}}, y) \in \mathcal{B}_{T_{\text{ini}}+T_f}$. Then there exists a unique $x_{\text{ini}} \in \mathbb{R}^{n(\mathcal{B})}$ such that

$$y = \mathcal{O}_{T_f}(A, C)x_{\text{ini}} + \mathcal{T}_{T_f}(A, B, C, D)u.$$

In other words, given a sufficiently long window of initial system data $\text{col}(u_{\text{ini}}, y_{\text{ini}})$, the initial state x_{ini} is unique and can be computed given knowledge of A, B, C, D and $\text{col}(u, y)$.

C. Hankel and Page Matrices

We introduce two important matrix structures: the Hankel matrix and the Page matrix.

Definition 2.5: Let $L, T \in \mathbb{Z}_{>0}$. Let $u_{[1,T]} = \{u_j\}_{j=1}^T \subset \mathbb{R}^m$ be a sequence of vectors.

- We define the *Hankel matrix*¹ of depth L as

$$\mathcal{H}_L(u_{[1,T]}) := \begin{pmatrix} u_1 & u_2 & \cdots & u_{T-L+1} \\ u_2 & u_3 & \cdots & u_{T-L+2} \\ \vdots & \vdots & \ddots & \vdots \\ u_L & u_{L+1} & \cdots & u_T \end{pmatrix}.$$

- We define the *Page matrix* of depth L as

$$\mathcal{P}_L(u_{[1,T]}) := \begin{pmatrix} u_1 & u_{L+1} & \cdots & u_{(\lfloor \frac{T}{L} \rfloor - 1)L + 1} \\ u_2 & u_{L+2} & \cdots & u_{(\lfloor \frac{T}{L} \rfloor - 1)L + 2} \\ \vdots & \vdots & \ddots & \vdots \\ u_L & u_{2L} & \cdots & u_{\lfloor \frac{T}{L} \rfloor L} \end{pmatrix},$$

where $\lfloor \cdot \rfloor$ is the floor function which rounds its argument down to the nearest integer.

Note that if T is a multiple of L , then $\lfloor \frac{T}{L} \rfloor L = T$ and the above expressions simplify accordingly. Hankel matrices have a long history in subspace identification [9] and more recently in data-driven control [1], [3], [10], [14]. Page matrices have also been used as an alternative to the classical Hankel matrix [43], [44].

Definition 2.6: Let $L, T, M \in \mathbb{Z}_{>0}$ and $u_{[1,T]} = \{u_j\}_{j=1}^T \subset \mathbb{R}^m$ be a sequence of vectors. We call $u_{[1,T]}$:

- *Hankel exciting of order L* if the matrix $\mathcal{H}_L(u_{[1,T]})$ has full row rank.
- *L-Page exciting of order M* if the matrix

$$\begin{pmatrix} \mathcal{P}_L(u_{[1, T-(M-1)L]}) \\ \mathcal{P}_L(u_{[L+1, T-(M-2)L]}) \\ \vdots \\ \mathcal{P}_L(u_{[L(M-1)+1, T]}) \end{pmatrix}$$

has full row rank.

Roughly speaking, the terms Hankel and Page exciting refer to a collection of inputs sufficiently rich and long to yield an output sequence representative for the system's behaviour. Note that for a sequence of vectors to be Hankel exciting of order L , one requires that $T \geq L(m+1) - 1$. For a sequence of vectors to be L -Page exciting of order M , one requires that $T \geq L((mL+1)M - 1)$. We also observe that the definition of an L -Page exciting sequence of order M depends on two indices, L and M (the number of block rows of the Page matrices and the number of vertically concatenated Page matrices, respectively), whereas the definition of a Hankel exciting sequence of order L depends only on a single index L (the number of block rows of the Hankel matrix).

The definition for Hankel exciting appeared in [4] under the name *persistently exciting*. A more general notion of persistency of excitation appeared in [18] in the case when the data matrix was a mosaic Hankel matrix and was termed collective persistency of excitation. Our notion of Page exciting is a modification of these notions; the main difference is that there are no repeated entries in each of the data matrices. This

¹The classical definition of a Hankel matrix requires it to be square. We slightly abuse this classical terminology, allowing for general dimensions.

has important implications when the entries of the matrix are corrupted by noise; see Section V-B, the case-study [12], and references [43], [44]. In short, the independence of the Page matrix entries leads to statistically and algorithmically favourable properties, e.g., singular-value thresholding can be used for de-noising. We will observe later that certain robustness and optimality guarantees become tight for Page matrices, and they lead to superior control performance. The price-to-pay is an increasing number of data samples. Almost all results within this paper hold for both Hankel and Page matrix structures, and we will explicitly comment on when the results differ depending on the matrix structure chosen.

D. Non-parametric System Representation

We now present a result known in behavioural systems theory as the *Fundamental Lemma* [3]. The result first appeared in [4, Theorem 1] in the case when the Hankel matrix was used, and was recently extended in [18, Theorem 2] for more general mosaic Hankel matrices. We present the result using the Page matrix data structure introduced in Definition 2.5.

Theorem 2.1: Consider a controllable system $\mathcal{B} \in \mathcal{L}^{m+p}$. Let $L, T \in \mathbb{Z}_{>0}$ with $L \geq n(\mathcal{B})$. Let $(\hat{u}_{[1,T]}, \hat{y}_{[1,T]}) = \{(\hat{u}_j, \hat{y}_j)\}_{j=1}^T \subset \mathbb{R}^{(m+p)}$ be a T -length trajectory of $\mathcal{B}$. Assume $\hat{u}_{[1,T]}$ to be L -Page exciting of order $n(\mathcal{B}) + 1$. Then $(u_{[1,L]}, y_{[1,L]}) = \text{col}(u_1, \dots, u_L, y_1, \dots, y_L) \in \mathcal{B}_L$ if and only if there exists a vector $g \in \mathbb{R}^{\lfloor \frac{T}{L} \rfloor}$ such that

$$\begin{pmatrix} \mathcal{P}_L(\hat{u}_{[1,T]}) \\ \mathcal{P}_L(\hat{y}_{[1,T]}) \end{pmatrix} g = \begin{pmatrix} u_{[1,L]} \\ y_{[1,L]} \end{pmatrix} \quad (1)$$

The original result [4, Theorem 1] required Hankel excitation of order $L + n(\mathcal{B})$, and it coincides with Theorem 2.1 for $L = 1$. Theorem 2.1 replaces the need for a model or system identification process and allows for any trajectory of a controllable LTI system to be constructed using a finite number of data samples generated by a sufficiently rich (i.e., L -page exciting) input sequence. Each column of the Page matrix in (1) is a trajectory of the system, and can be thought of as a motion primitive. By linearly combining elements of this trajectory library, we recover the whole space of trajectories.

In a sense, the Page matrix in (1) is a non-parametric predictive model based on raw data. To be precise in behavioural language, the Page matrix in (1) is an image representation of the restricted behaviour $\mathcal{B}_L$. This is in contrast to kernel representations (with latent variables) such as state-space models which parameterize $\mathcal{B}_L$ by means of its orthogonal complement. In what follows, we will express a linear system $\mathcal{B}$ by its *data matrix representation* in (1). We refer to the Page matrix on the left-hand side of (1) as the *data matrix*.

Remark 2.1: When $L \geq \ell(\mathcal{B})$, it can be shown that the rank of the data matrix in the data matrix representation given by (1) is $mL + n(\mathcal{B})$ where m is the dimension of the inputs. Hence, the column span of the data matrix carves out a low-dimensional subspace of $\mathbb{R}^{L(m+p)}$ which coincides with the restricted behaviour $\mathcal{B}_L$, the space of L -length trajectories. Since in general $n(\mathcal{B})$ is unknown, it suffices to replace $n(\mathcal{B})$ in Theorem 2.1 with an upper bound on the system order. •

For the rest of the paper, we think of an LTI system $\mathcal{B}$ only in terms of its data matrix representation. We see next that Theorem 2.1 allows to implicitly estimate the state, predict the future behaviour, and design optimal control inputs [3].

III. DETERMINISTIC DEEPC

Consider the controllable LTI system $\mathcal{B} \in \mathcal{L}^{m+p}$ whose model is unknown. We address the problem of finite-horizon optimal control, where the goal is to design a finite sequence of control inputs that result in desirable outputs. In particular, let $t \in \mathbb{Z}_{\geq 0}$ be the current time, let $T_f \in \mathbb{Z}_{>0}$ be the prediction horizon, and let $f_1: \mathbb{R}^{mT_f} \rightarrow \mathbb{R}_{\geq 0}$ (respectively, $f_2: \mathbb{R}^{pT_f} \rightarrow \mathbb{R}_{\geq 0}$) be a cost function on the future inputs (respectively, outputs). We wish to design an input sequence $\text{col}(u_t, \dots, u_{t+T_f-1}) \in \mathbb{R}^{mT_f}$ such that the corresponding output sequence $\text{col}(y_t, \dots, y_{t+T_f-1}) \in \mathbb{R}^{pT_f}$ minimizes the cost $f_1 + f_2$. Furthermore, we wish for the inputs and outputs to lie in the constraint sets $\mathcal{U} \subseteq \mathbb{R}^{mT_f}$ and $\mathcal{Y} \subseteq \mathbb{R}^{pT_f}$, respectively.

The conventional formulation of this finite-time optimal control problem uses a parametric input/output/state representation $\mathcal{B}(A, B, C, D)$ for prediction and estimation:

$$\begin{aligned} \min_{u, y, x} \quad & f_1(u) + f_2(y) \\ \text{s.t.} \quad & x_{k+1} = Ax_k + Bu_k \quad \forall k \in \{0, \dots, T_f - 1\} \\ & y_k = Cx_k + Du_k \quad \forall k \in \{0, \dots, T_f - 1\} \\ & x_0 = \hat{x}_t \\ & u \in \mathcal{U}, y \in \mathcal{Y}. \end{aligned} \quad (2)$$

Here, $\hat{x}_t$ is the current state at time t , which in the deterministic case can be obtained, e.g., by propagating the model equations backward in time by $T_{\text{ini}} \geq \ell(\mathcal{B})$ steps; see Lemma 2.1. When optimization problem (2) is implemented in a receding horizon fashion it is known as model-predictive control (MPC) and is a widely celebrated control technique. Typically, in receding horizon MPC problems the constraints take the form of separable stage constraints with an additional terminal constraint (similarly for the cost functions) [45].

Using Theorem 2.1, we can see that raw data collected from system $\mathcal{B}$ can be used to perform implicit state estimation at the same time as predict forward trajectories of $\mathcal{B}$. Indeed, let $T_{\text{ini}}, T_f \in \mathbb{Z}_{>0}$ and let $(\hat{u}_{[1,T]}, \hat{y}_{[1,T]}) = \{(\hat{u}_j, \hat{y}_j)\}_{j=1}^T \subset \mathbb{R}^{(m+p)}$ be a T -length trajectory of $\mathcal{B}$ collected offline. Assume that $\hat{u}_{[1,T]}$ is $(T_{\text{ini}} + T_f)$ -Page exciting of order $n(\mathcal{B}) + 1$. We partition the data matrices into two parts; one will be used to perform the implicit state estimation and the other to perform forward predictions. More formally, we use the language of subspace system identification [9] and define

$$\begin{pmatrix} \hat{U}_p \\ \hat{U}_f \end{pmatrix} := \mathcal{P}_{T_{\text{ini}}+T_f}(\hat{u}_{[1,T]}), \quad \begin{pmatrix} \hat{Y}_p \\ \hat{Y}_f \end{pmatrix} := \mathcal{P}_{T_{\text{ini}}+T_f}(\hat{y}_{[1,T]}), \quad (3)$$

where $\hat{U}_p \in \mathbb{R}^{mT_{\text{ini}} \times \lfloor \frac{T}{T_{\text{ini}}+T_f} \rfloor}$ consists of the first T_{ini} block rows of $\mathcal{P}_{T_{\text{ini}}+T_f}(\hat{u}_{[1,T]})$ and $\hat{U}_f \in \mathbb{R}^{mT_f \times \lfloor \frac{T}{T_{\text{ini}}+T_f} \rfloor}$ the last T_f block rows (similarly for $\hat{Y}_p$ and $\hat{Y}_f$). Let $t \in \mathbb{Z}_{\geq 0}$ be the current time and denote by $\text{col}(\hat{u}_{\text{ini}}, \hat{y}_{\text{ini}})$ the most recent T_{ini} -length trajectory measured from $\mathcal{B}$. Then given a sequence of

inputs $u = \text{col}(u_t, \dots, u_{T_f-1})$ of length $T_f \in \mathbb{Z}_{>0}$, one may solve for a vector $g \in \mathbb{R}^{\lfloor \frac{T}{T_{\text{ini}}+T_f} \rfloor}$ satisfying

$$\begin{pmatrix} \hat{U}_p \\ \hat{Y}_p \\ \hat{U}_f \end{pmatrix} g = \begin{pmatrix} \hat{u}_{\text{ini}} \\ \hat{y}_{\text{ini}} \\ u \end{pmatrix}. \quad (4)$$

By Theorem 2.1, the predicted T_f -length output of $\mathcal{B}$ is then given by $y = \hat{Y}_f g$. Note that in general the g satisfying (4) is non-unique. However, by Lemma 2.1, if $T_{\text{ini}} \geq \ell(\mathcal{B})$ then the predicted output y is unique. The top two block equations of (4) can be thought of as implicitly fixing the initial state from which the future trajectory will depart.

We now recall the deterministic DeePC method presented in [1] (using Hankel matrices) for controlling the system $\mathcal{B}$.

Data Collection (offline): Apply to $\mathcal{B}$ a T -length input sequence $\hat{u}_{[1,T]}$ that is $(T_{\text{ini}}+T_f)$ -Page exciting of order $n(\mathcal{B})+1$. Measure the corresponding output trajectory $\hat{y}_{[1,T]}$. Form data matrices $\hat{U}_p, \hat{U}_f, \hat{Y}_p, \hat{Y}_f$ according to (3).

DeePC (online): Let $f_1: \mathbb{R}^{mT_f} \rightarrow \mathbb{R}_{\geq 0}$, $f_2: \mathbb{R}^{pT_f} \rightarrow \mathbb{R}_{\geq 0}$ be cost functions describing the costs on future inputs and outputs of $\mathcal{B}$, respectively. Collect the most recent T_{ini} -length trajectory $\text{col}(\hat{u}_{\text{ini}}, \hat{y}_{\text{ini}})$ measured from $\mathcal{B}$. Solve the finite-horizon optimal control problem:

$$\begin{aligned} \min_g \quad & f_1(\hat{U}_f g) + f_2(\hat{Y}_f g) \\ \text{s.t.} \quad & \begin{pmatrix} \hat{U}_p \\ \hat{Y}_p \end{pmatrix} g = \begin{pmatrix} \hat{u}_{\text{ini}} \\ \hat{y}_{\text{ini}} \end{pmatrix} \\ & \hat{U}_f g \in \mathcal{U} \\ & \hat{Y}_f g \in \mathcal{Y}. \end{aligned} \quad (5)$$

Similar to MPC, the online DeePC optimization can be implemented in a receding horizon fashion. Note that the optimization problem (5) only requires input/output measurements from the system and does not require the identification of a model. The equivalence of DeePC to classical model predictive control when a system model for $\mathcal{B}$ is given, is established in [1] in the case when the data matrix is a Hankel matrix.

Proposition 3.1: Consider an LTI system $\mathcal{B}$ whose input/output/state representation is given by $\mathcal{B}(A, B, C, D)$. Consider the MPC optimization problem (2) and the optimization problem (5) with $T_{\text{ini}} \geq \ell(\mathcal{B})$. Then g satisfying the constraints of (5) implies that $(u, y) = (\hat{U}_f g, \hat{Y}_f g)$ satisfies the constraints of (2). Conversely, (u, y) satisfying the constraints of (2) implies the existence of g satisfying the constraints of (5) where $(u, y) = (\hat{U}_f g, \hat{Y}_f g)$.

As a direct corollary of Proposition 3.1, DeePC problem (5) and MPC problem (2) result in equivalent and unique (assuming convexity) open-loop (resp., closed-loop) behaviour when applied in an open-loop (resp., receding horizon) fashion to a deterministic LTI system $\mathcal{B}$. As a major difference, most MPC approaches assume direct state measurements, which allows certain methods to be applied (e.g., inclusion of terminal state constraints). We remark that $(u_{\text{ini}}, y_{\text{ini}}, u, y)$ are coordinates of a (generally non-minimal) realization allowing for similar methodologies. See [10] in which stability and recursive feasibility for a variant of (5) have been shown.

IV. DISTRIBUTIONALLY ROBUST DEEPC

We now consider the case where system $\mathcal{B}$ is subject to stochastic disturbances, which result in noisy output measurements in the data matrices. We begin with modifying the objective function and constraints of the deterministic DeePC problem (5) to robustify against the stochastic disturbances. This leads to a semi-infinite optimization problem which we reformulate under reasonable assumptions into a finite, convex, and tractable program termed the distributionally robust DeePC problem. Finally, we show that the distributionally robust DeePC problem enjoys robust performance guarantees.

A. Problem Setup

We view the output data matrices $\hat{Y}_p$ and $\hat{Y}_f$ given by (3) as particular realizations of random variables which we denote by Y_p and Y_f , respectively. Recall that we use the $\hat{\cdot}$ notation to denote recorded data samples and to indicate that objects depend on recorded data samples. Define the random variable $\xi = (\xi_1^\top, \dots, \xi_{p(T_{\text{ini}}+T_f)}^\top)^\top$ where $\xi_i^\top \in \mathbb{R}^{\lfloor \frac{T}{T_{\text{ini}}+T_f} \rfloor}$ denotes the i -th row of the matrix $(Y_p^\top \ Y_f^\top)^\top$. We denote the support set of the random variable ξ by Ξ . We denote the probability distribution of the random variable ξ by $\mathbb{P}$. Note that this distribution is induced by the system $\mathcal{B}$ itself, and any stochastic disturbance acting on the system.

Due to the stochastic nature of the system, the future trajectory predictions obtained from the deterministic DeePC optimization problem (5) are uncertain. Furthermore, the realized data matrix (1) may not describe a low-dimensional subspace as described in Remark 2.1. In fact, the data matrix in (1) will likely have full rank and can thus predict arbitrary future trajectories. Finally, the distribution $\mathbb{P}$ itself is unknown since we do not have models of the system and disturbances. Hence, we make some changes to the objective function and the constraints in (5) to robustify against the noisy data.

Objective Function: Since the constraint used to obtain the initial condition of the future trajectory given by $\hat{Y}_p g = \hat{y}_{\text{ini}}$ is affected by the stochastic disturbance and may not be feasible, we lift it into the objective function using an estimation function $f_3: \mathbb{R}^{pT_{\text{ini}}} \rightarrow \mathbb{R}_{\geq 0}$ mapping $Y_p g - \hat{y}_{\text{ini}}$ to a penalty. Including f_3 in the objective function can be thought of as an implicit estimation of the initial condition from which the predicted future trajectory must evolve. Choosing $f_3(\cdot) = \|\cdot\|_2$ would result in a least-squares type initial condition estimate reminiscent of moving horizon estimation [45].

Since the system is stochastic, we focus on minimizing the expectation of the objective function in (5), where the expectation is taken with respect to the true distribution $\mathbb{P}$, i.e., we wish to minimize $\mathbb{E}_{\mathbb{P}}[f_1(\hat{U}_f g) + f_2(\hat{Y}_f g) + f_3(Y_p g - \hat{y}_{\text{ini}})]$.

Recall that the distribution $\mathbb{P}$ pertains to the offline samples $(\hat{u}_{[1,T]}, \hat{y}_{[1,T]})$ and not to the real-time measurement $\hat{y}_{\text{ini}}$ addressed through the above least-square estimation. As discussed before, the distribution $\mathbb{P}$ itself is unknown, and must be estimated from data. In order to be robust to errors in this estimate, we use distributionally robust optimization techniques; namely, we construct a so-called ambiguity set $\hat{\mathcal{P}}$ depending on the collected data (specified precisely later) such that the true distribution lies in the ambiguity set with

high confidence. We then optimize the expectation of the cost function, where the expectation is taken with respect to the worst-case distribution in $\hat{\mathcal{P}}$. This leads us to the distributionally robust objective

$$\sup_{\mathbb{Q} \in \hat{\mathcal{P}}} \mathbb{E}_{\mathbb{Q}} \left[f_1(\hat{U}_f g) + f_2(Y_f g) + f_3(Y_p g - \hat{y}_{ini}) \right].$$

We note that this formulation leads to robustness against a “set of systems” compatible with the (possibly noisy) data samples, and that this “set of systems” is broader than mere LTI systems with additive process and measurement noise.

Constraints: Requiring the future outputs to lie in a particular subset $\mathcal{Y}$ almost surely as in (5) may not be possible. Thus, we relax the hard output constraint to a so-called conditional value at risk (CVaR) constraint. Let the output constraint be written as $\mathcal{Y} = \{y \mid h(y) \leq 0\}$, where $h: \mathbb{R}^{pT_t} \rightarrow \mathbb{R}$ describes desired constraints on future trajectories y . We define

$$\text{CVaR}_{1-\alpha}^{\mathbb{P}}(h(y)) := \inf_{\tau \in \mathbb{R}} \left\{ \tau + \frac{1}{\alpha} \mathbb{E}_{\mathbb{P}}[(h(y) - \tau)_+] \right\},$$

where $\alpha \in (0, 1)$ is a user-chosen confidence parameter. It is well known that the constraint $\text{CVaR}_{1-\alpha}^{\mathbb{P}}(h(y)) \leq 0$ is a convex relaxation of the classical chance constraint (or VaR) given by $\mathbb{P}(h(y) \leq 0) \geq 1 - \alpha$ (see, e.g., [46]). In other words, the CVaR constraint ensures that the constraint $y \in \mathcal{Y}$ will hold with high probability. In the case when the stochastic disturbance is a continuous random variable, the CVaR constraint can be interpreted as penalizing the expected violation in the α percent of cases where violations do occur (see Figure 1). The latter point is particularly important in control applications where large violations of constraints could lead to catastrophic system behaviour. Hence, the CVaR constraint is not merely a relaxation of the classical VaR constraint, but often the more reasonable constraint formulation for control problems.

Fig. 1: Depiction of the conditional value at risk at level $\alpha = 0.15$ for a Gaussian random variable X . The shaded region accounts for $\alpha \times 100\%$ of the mass of the Gaussian distribution.

Similar to minimizing the worst case expected objective function, we ask that the CVaR constraint be satisfied for the worst-case distribution in $\hat{\mathcal{P}}$. This leads us to the robust constraint

$$\sup_{\mathbb{Q} \in \hat{\mathcal{P}}} \text{CVaR}_{1-\alpha}^{\mathbb{Q}}(h(y)) \leq 0.$$

One particular ambiguity set $\hat{\mathcal{P}}$ that offers strong performance guarantees as well as mathematical tractability while allowing the user to adjust its conservatism is the *Wasserstein ambiguity set* centred around the so-called *empirical distribution*, i.e., the sample distribution [5]. In particular, let $\mathcal{M}(\Xi)$ be the set of all distributions $\mathbb{Q}$ supported on Ξ such that $\mathbb{E}_{\mathbb{Q}}[\|\xi\|_r] < \infty$, where $\|\cdot\|_r$ is the r -norm for some $r \in [1, \infty]$.

Definition 4.1: Let $r \in [1, \infty]$. The *Wasserstein metric* $d_W: \mathcal{M}(\Xi) \times \mathcal{M}(\Xi) \rightarrow \mathbb{R}_{\geq 0}$ is defined as

$$d_W(\mathbb{Q}_1, \mathbb{Q}_2) := \inf_{\Pi} \left\{ \int_{\Xi^2} \|\xi_1 - \xi_2\|_r \Pi(d\xi_1, d\xi_2) \right\},$$

where Π is a joint distribution of ξ_1 and ξ_2 with marginal distributions $\mathbb{Q}_1 \in \mathcal{M}(\Xi)$ and $\mathbb{Q}_2 \in \mathcal{M}(\Xi)$, respectively.

The Wasserstein metric can be viewed as a distance between probability distributions, where the distance is calculated via an optimal mass transport plan Π . For $\epsilon \geq 0$, we denote the *Wasserstein ball of radius ϵ centred around distribution $\mathbb{Q}$* by

$$B_{\epsilon}(\mathbb{Q}) := \{\mathbb{Q}' \in \mathcal{M}(\Xi) \mid d_W(\mathbb{Q}, \mathbb{Q}') \leq \epsilon\}.$$

We chose the Wasserstein ball as the ambiguity set as opposed to other popular ambiguity sets such as f -divergence or moment-based ambiguity sets since the latter require prior knowledge and assumptions on the true data generating distribution (absolute continuity with respect to a nominal distribution or known moments).

We now present the distributionally robust DeePC method for controlling stochastic systems.

Data Collection (offline): We collect $N \in \mathbb{Z}_{>0}$ trajectories via repeated identical experiments. We start by fixing an input sequence $\hat{u}_{[1,T]}$ that is $(T_{ini} + T_f)$ -Page exciting of order $n(\mathcal{B}) + 1$. For each experiment $i \in \{1, \dots, N\}$, we assume that the system is initialized to the same (unknown) state, apply inputs $\hat{u}_{[1,T]}$, and measure the corresponding output trajectory $\hat{y}_{[1,T]}^{(i)}$. The superscript (i) denotes the trajectory obtained from the i -th experiment. For each data batch $i \in \{1, \dots, N\}$, build data matrices $\hat{U}_p, \hat{U}_f, \hat{Y}_p^{(i)}$, and $\hat{Y}_f^{(i)}$ using (3). The output data matrices implicitly define N data samples which we denote by $\hat{\xi}^{(i)}$ for $i \in \{1, \dots, N\}$. Lastly, we use the data collected to build the *empirical distribution* $\hat{\mathbb{P}}_N = \frac{1}{N} \sum_{i=1}^N \delta_{\hat{\xi}^{(i)}}$. Note that collecting only one T -length trajectory (i.e., $N = 1$) is possible. As we will see in Section V-B though, choosing N larger can lead to less uncertainty and better performance.

Distributionally Robust DeePC (online): We consider the Wasserstein ball of radius ϵ (which will be chosen as a hyperparameter quantifying a desired robustness level) around the empirical distribution $\hat{\mathbb{P}}_N$ as the ambiguity set. Hence, the distributionally robust DeePC problem is given by

$$\begin{aligned} \min_g \quad & \sup_{\mathbb{Q} \in B_{\epsilon}(\hat{\mathbb{P}}_N)} \mathbb{E}_{\mathbb{Q}} \left[f_1(\hat{U}_f g) + f_2(Y_f g) + f_3(Y_p g - \hat{y}_{ini}) \right] \\ \text{s.t.} \quad & \hat{U}_p g = \hat{u}_{ini} \\ & \hat{U}_f g \in \mathcal{U} \\ & \sup_{\mathbb{Q} \in B_{\epsilon}(\hat{\mathbb{P}}_N)} \text{CVaR}_{1-\alpha}^{\mathbb{Q}}(h(Y_f g)) \leq 0 \end{aligned} \tag{6}$$

By assuming that the system is initialized to the same (unknown) state at the beginning of each experiment and applying an identical input sequence for each experiment, we ensure that the collection of data matrices $\hat{Y}_p^{(i)}$ and $\hat{Y}_f^{(i)}$ describe N independent and identically distributed (i.i.d.) samples of the random variable ξ . Before their realization, the collection $\{\hat{\xi}^{(i)}\}_{i=1}^N$ can be viewed as a random object governed by the N -fold product distribution $\mathbb{P}^N$. The empirical distribution $\hat{\mathbb{P}}_N$

is our sample average estimate for the true distribution $\mathbb{P}$ and is set as the centre of the ambiguity set in (6).

Compared to MPC or deterministic DeePC discussed in Section III, the distributional robust program (6) features the additional robustness parameter ϵ , whose selection will be discussed later, and α whose value is usually chosen to be 0.1, 0.05, or 0.01. One notable feature of (6) is that it is robust to all probability distributions within the ambiguity set that could describe the data matrices Y_p and Y_f . Hence, the above formulation is robust to a “set of systems” which captures more than LTI systems with additive noise. In fact, for a sufficiently large robustness parameter ϵ , the ambiguity set considers all time series originating from linear or nonlinear systems as long as they have “finite expectation”, i.e., the integrals defining the Wasserstein metric exist. This follows from the fact that the Wasserstein ball contains all Dirac distributions generated by time series ξ satisfying $d_W(\delta_\xi, \hat{\mathbb{P}}_N) \leq \epsilon$. Hence, for ϵ large enough, the time series ξ can be arbitrary.

Problem (6) is a semi-infinite optimization problem due to the supremums over the space of distributions in the cost and in the constraints. Moreover, the Wasserstein metric (resp. CVaR) are themselves defined via infinite (resp. finite) programs. Thus, (6) does not immediately present itself as a tractable program. Below, we show that under reasonable assumptions on the objective and constraint function, the semi-infinite optimization problem above admits a tractable reformulation.

B. Main results

We provide a tractable reformulation of the distributionally robust DeePC problem (6) which depends on two intermediate lemmas (Lemma A.3 and A.4 in the Appendix) providing separate reformulations for the objective function and constraints. In the statement of the result below, we make use of the dual norm $\|\cdot\|_q = \|\cdot\|_{r,*}$, where, q and r satisfy $\frac{1}{r} + \frac{1}{q} = 1$.

Theorem 4.1: (Tractable Reformulation): Assume that f_1 is convex and f_2, f_3 , and h are convex and Lipschitz continuous. Specifically, let $L_{\text{obj}} > 0$ (resp., $L_{\text{con}} > 0$) be the Lipschitz constant with respect to the r -norm of the mapping $(x, y) \mapsto f_2(x) + f_3(y)$ (resp., h). Then the optimal value of problem (6) is upper bounded by the optimal value of

$$\begin{aligned} \min_{g, \tau, s_i} \quad & f_1(\hat{U}_f g) + \frac{1}{N} \sum_{i=1}^N \left(f_2(\hat{Y}_f^{(i)} g) + f_3(\hat{Y}_p^{(i)} g - \hat{y}_{\text{ini}}) \right) \\ & + L_{\text{obj}} \epsilon \|g\|_{r,*} \\ \text{s.t.} \quad & \hat{U}_p g = \hat{u}_{\text{ini}} \\ & \hat{U}_f g \in \mathcal{U} \\ & -\tau \alpha + L_{\text{con}} \epsilon \|g\|_{r,*} + \frac{1}{N} \sum_{i=1}^N s_i \leq 0 \\ & \tau + h(\hat{Y}_f^{(i)} g) \leq s_i \quad \forall i \leq N \\ & s_i \geq 0 \quad \forall i \leq N. \end{aligned} \quad (7)$$

Moreover, the solution $\hat{g}^*$ to the above will satisfy $\text{CVaR}_{1-\alpha}^{\mathbb{Q}}(h(Y_f \hat{g}^*)) \leq 0$ for all $\mathbb{Q} \in B_\epsilon(\hat{\mathbb{P}}_N)$. The objective of (7) coincides with the objective of (6) when $\Xi = \mathbb{R}^{p(T_{\text{ini}}+T_f) \lfloor \frac{T}{T_{\text{ini}}+T_f} \rfloor}$.

Observe that when the data matrices in (3) are Hankel matrices, then Ξ is always a strict subspace of $\mathbb{R}^{p(T_{\text{ini}}+T_f) \lfloor \frac{T}{T_{\text{ini}}+T_f} \rfloor}$ encoding the Hankel structure. When the matrices are Page matrices, the support set Ξ may coincide with $\mathbb{R}^{p(T_{\text{ini}}+T_f) \lfloor \frac{T}{T_{\text{ini}}+T_f} \rfloor}$ (e.g., if the measurements are affected by Gaussian noise). A similar result holds for piecewise affine constraint functions by combining Lemma A.3 and A.5. In this case, the reformulated constraint set can be made tight at the cost of making the objective an upperbound (see Remark A.1). Since, we treat the objective and constraints separately (via Lemma A.3 and A.4), it is possible to choose different values of ϵ in the constraint and objective of (6), though we refrain from doing so for clarity of exposition. Furthermore, by treating the objective and the constraints separately, their worst case distributions may differ leading to a more conservative solution.

The above result shows that the distributionally robust objectives can be reformulated as the sample average of the objective function plus an additional regularization term. To the best of our knowledge this observation was first made in [47] in the context of machine learning problems. Hence, being robust in the trajectory space (data space) in the sense of the r -norm requires regularizing the sample average objective with the *dual* q -norm, where the weight of regularization depends on the Lipschitz constant of the objective function and the radius of the Wasserstein ball. For example, the 1-norm regularization adopted in [1] corresponds to ∞ -norm robustness in the trajectory space.

We now show that the robust DeePC problem (7) exhibits strong probabilistic guarantees. We need the following auxiliary measure concentration result under a light-tailedness assumption.

Theorem 4.2: (Measure Concentration [48, Theorem 2]): Assume that $\mathbb{E}_{\mathbb{P}}[\exp(\|\xi\|_r^a)] < \infty$ for some $a > 1$. Then

$$\begin{aligned} & \mathbb{P}^N \left\{ d_W(\mathbb{P}, \hat{\mathbb{P}}_N) \geq \epsilon \right\} \\ & \leq \begin{cases} c_1 \exp(-c_2 N \epsilon^{\lfloor \frac{T}{T_{\text{ini}}+T_f} \rfloor}) & \text{if } \epsilon \leq 1 \\ c_1 \exp(-c_2 N \epsilon^a) & \text{if } \epsilon > 1 \end{cases} \end{aligned}$$

for all $N \geq 1$, and $\epsilon > 0$, where $\mathbb{P}^N$ is the N -fold product distribution, and c_1, c_2 are positive constants depending on a , $\lfloor \frac{T}{T_{\text{ini}}+T_f} \rfloor$ and the value of $\mathbb{E}_{\mathbb{P}}[\exp(\|\xi\|_r^a)]$.

We can use Theorem 4.2 in order to compute the minimum Wasserstein radius ϵ such that the true data-generating distribution $\mathbb{P}$ lives inside the Wasserstein ball $B_\epsilon(\hat{\mathbb{P}}_N)$ with confidence at least $1 - \beta$. Indeed, by inverting the above inequalities the minimum radius is given by

$$\epsilon(\beta, N) = \begin{cases} \left(\frac{\log(c_1 \beta^{-1})}{c_2 N} \right)^{\frac{1}{\lfloor \frac{T}{T_{\text{ini}}+T_f} \rfloor}} & \text{if } N \geq \frac{\log(c_1 \beta^{-1})}{c_2} \\ \left(\frac{\log(c_1 \beta^{-1})}{c_2 N} \right)^{\frac{1}{a}} & \text{if } N < \frac{\log(c_1 \beta^{-1})}{c_2} \end{cases} \quad (8)$$

Theorem 4.3: (Robust Performance Guarantee): Assume that $\mathbb{E}_{\mathbb{P}}[e^{\|\xi\|_r^a}] < \infty$ for some $a > 1$. Let $\beta \in (0, 1)$ be the desired level of confidence. Let $J(g) := \mathbb{E}_{\mathbb{P}}[f_1(\hat{U}_f g) + f_2(Y_f g) + f_3(Y_p g - \hat{y}_{\text{ini}})]$. Let $\hat{J}(g)$ denote the value of the

objective function in (7) evaluated at g and $\hat{g}^*$ denote an optimizer of problem (7) with $\epsilon(\beta, N)$ chosen as in (8). Then

$$\mathbb{P}^N \left\{ J(\hat{g}^*) \leq \hat{J}(\hat{g}^*) \right\} \geq 1 - \beta, \text{ and}$$

$$\mathbb{P}^N \left\{ \text{CVaR}_{1-\alpha}^{\mathbb{P}}(h(Y_{\hat{f}}\hat{g}^*)) \leq 0 \right\} \geq 1 - \beta.$$

Theorem 4.3 shows that with probability $1 - \beta$ (with respect to the N -fold product distribution $\mathbb{P}^N$ of $\mathbb{P}$), the optimal value of the robust DeePC problem (7) is an upper bound for the expected cost with respect to the true distribution $\mathbb{P}$. Furthermore, the CVaR constraint with respect to the true distribution $\mathbb{P}$ holds with probability $1 - \beta$.

Remark 4.1: The Wasserstein radius given by (8) decays with a rate of $N^{\frac{1}{T/(T_{\text{ini}}+T_f)}}$. To decrease the radius by 50%, the number of samples N must increase by $2^{\lfloor \frac{T}{T_{\text{ini}}+T_f} \rfloor}$. This rate (albeit unfortunate) is tight [49]. In practice, the Wasserstein ball radius given by (8) is often larger than necessary, i.e., $\mathbb{P} \notin B_\epsilon(\mathbb{P}_N)$ with probability much less than β . Furthermore, even when $\mathbb{P} \in B_\epsilon(\mathbb{P}_N)$, the robust quantity $\hat{J}(\hat{g}^*)$ may still serve as an upper bound for the out-of-sample performance $J(\hat{g}^*)$ [5]. In addition to this conservativeness, we note that the Wasserstein radius provided by (8) is hardly quantifiable as it depends on constants which are unknown to us in our data-driven setting. For practical purposes, one should choose the radius of $B_\epsilon(\mathbb{P}_N)$ in a data-driven fashion (e.g., as we have done via tuning in experiments in Section V-B). Alternatively, one could use standard cross-validation techniques (some of which are stated in [5]) or more advanced techniques (albeit only providing asymptotic guarantees) provided in [50]. •

The robust DeePC Algorithm 1 implements (7) in a receding horizon fashion. We note that all results in this section hold when the data matrices are Hankel matrices in (3). However, the upper bound and set inclusions presented in Lemmas A.3, A.4 and Theorem 4.1 are not tight because the Hankel structure restricts Ξ to a strict subspace of $\mathbb{R}^{p(T_{\text{ini}}+T_f)\lfloor \frac{T}{T_{\text{ini}}+T_f} \rfloor}$. If the constraints are piecewise affine (see Lemma A.5), then the constraint reformulation can be tight for Hankel matrices. In the following case study we will compare performance of Hankel and Page matrix structures.

Algorithm 1 Robust DeePC

Input: data trajectories $\text{col}(\hat{u}, \hat{y}^{(i)}) \in \mathbb{R}^{(m+p)T}$ for $i \in \{1, \dots, N\}$ with $\hat{u}(T_{\text{ini}}+T_f)$ -Page exciting of order $n(\mathcal{B})+1$, most recent input/output measurements $\text{col}(\hat{u}_{\text{ini}}, \hat{y}_{\text{ini}}) \in \mathbb{R}^{(m+p)T_{\text{ini}}}$

- 1: Solve (7) for $\hat{g}^*$.
 - 2: Compute optimal input sequence $u^* = \hat{U}_{\hat{f}}\hat{g}^*$.
 - 3: Apply optimal input sequence $(u_t, \dots, u_{t+\nu-1}) = (u_1^*, \dots, u_\nu^*)$ for some $\nu \leq T_f$.
 - 4: Set t to $t + \nu$ and update $\hat{u}_{\text{ini}}$ and $\hat{y}_{\text{ini}}$ to the T_{ini} most recent input/output measurements.
 - 5: Return to 1.
-

V. NUMERICAL RESULTS

A. Nonlinear and Stochastic Aerial Robotics Case Study

We illustrate the performance of distributionally robust DeePC with a high-fidelity nonlinear quadcopter simulation [13]. The output measurements are the 3 spatial coordinates (p_x, p_y, p_z) of the quadcopter. We denote the output at time t by $y_t = (p_{x,t}, p_{y,t}, p_{z,t}) \in \mathbb{R}^3$ and the i -th component of y_t by $y_{t,i}$. The inputs are the total thrust produced by all 4 rotors f_{tot} , and the angular body rates around the x and y body axes of the quadcopter, ω_x, ω_y respectively. We denote the input at time t by $u_t = (f_{\text{tot},t}, \omega_{x,t}, \omega_{y,t}) \in \mathbb{R}^3$ and the i -th component of u_t by $u_{t,i}$. The output measurements are affected by additive zero-mean Gaussian noise during the data collection phase (offline) and the control phase (online) in which the distributionally robust DeePC Algorithm 1 is implemented. Statistics of the noise were chosen to closely match the experimental setup [13].

The output measurements were taken at a rate of 25Hz (i.e., every 40ms). We performed $N = 10$ offline experiments each of which yielded 15500 input/output measurements to populate the Page matrices $\hat{U}_{\text{p}}, \hat{U}_{\text{f}}, \hat{Y}_{\text{p}}^{(i)}$, and $\hat{Y}_{\text{f}}^{(i)}$ for $i \in \{1, \dots, 10\}$. The inputs used to excite the system were drawn from a Gaussian distribution. The same randomly generated inputs were used for every repeated experiment. Due to the fact that the quadcopter system is open-loop unstable, these randomly generated inputs were added to an existing control that maintains the quadcopter around a hover state. This existing controller was only used during the offline data collection phase, and was not used during the online control phase. Since the randomly generated inputs were much larger in magnitude than the adjustments made by the existing stabilizing controller during data collection, the matrices $\hat{Y}_{\text{p}}^{(i)}, \hat{Y}_{\text{f}}^{(i)}$ are approximately i.i.d.

The future prediction horizon was chosen as $T_f = 25$ (1 second in real time), and the time horizon used to implicitly estimate the initial condition was set to $T_{\text{ini}} = 6$. The inputs were constrained to the set $[0.1597, 0.4791] \times [-\frac{\pi}{2}, \frac{\pi}{2}]^2$. The output constraint function h was chosen such that (p_x, p_y, p_z) was constrained in the set $[-0.6, 1.6]^3$ with $\alpha = 0.1$ the confidence parameter in the CVaR constraint. The composite objective function consisted of weighted 2-norms

$$f_1(u_{[1,T_f]}) = 16\|u_{[1,T_f],1} - f_{\text{ref}}\|_2 + 4\|(u_{[1,T_f],2}, u_{[1,T_f],3})\|_2$$

$$f_2(y_{[1,T_f]})^2 = 1600^2 \sum_{t=1}^{T_f} \|y_t - y_{\text{ref}}\|_2^2$$

$$f_3(\sigma) = 750000\|\sigma\|_2,$$

where $f_{\text{ref}} = (0.27, \dots, 0.27)$ is the thrust needed to hover the quadcopter and $y_{\text{ref}} = (p_{x,\text{ref}}, p_{y,\text{ref}}, p_{z,\text{ref}}) = (0.5, 0.5, 1.5)$. The Wasserstein norm was chosen as the 2-norm and $\epsilon = 0.003$. The optimization problem (7) was implemented in a receding horizon fashion with control horizon $\nu = 1$ (see Algorithm 1). A representative trajectory of the nonlinear and stochastic closed-loop system is illustrated in Figure 2(a), where the distributionally robust DeePC succeeds in steering the quadcopter to the reference trajectory while satisfying output constraints.

B. Choosing the Hyperparameters

In this section we study the effect of the hyperparameters on the performance of distributionally robust DeePC. To compare the performance of Algorithm 1 for different hyperparameters, we simulated the nonlinear stochastic quadcopter in receding horizon with control horizon $\nu = 1$. We computed the tracking error $\sum_{t=0}^{T_{\text{sim}}} \|y_t - y_{\text{ref}}\|_2^2$ where $T_{\text{sim}} = 250$ (corresponding to 10 seconds in real time).

1) ϵ : We varied the Wasserstein radius for fixed $N = 1$ and $N = 10$, $T_{\text{ini}} = 6$, and 300 columns in the data matrix. Figures 2(b) and 2(c) show the relationship between tracking error and ϵ for Page and Hankel matrix structures, respectively. Increasing the number of data batches N increases the range of radii that lead to satisfactory tracking performance. This is because the empirical distribution $\hat{\mathbb{P}}_N$ at the centre of the Wasserstein ball in problem (6) gets closer (in the Wasserstein sense) to the true distribution $\mathbb{P}$ from which the data is drawn. This allows the user to decrease ϵ . Since the optimal value of the distributionally robust DeePC problem (7) is monotonically increasing with ϵ , choosing the minimum epsilon such that $\mathbb{P} \in B_\epsilon(\hat{\mathbb{P}}_N)$, gives the minimum upper bound on the out-of-sample performance that holds with high confidence. Hence, increasing N decreases this minimum ϵ leading to better out-of-sample performance with high confidence. We note that the effect of increasing N is more pronounced for the Hankel matrices. This may be due to the fact that Page matrices already require more data points to construct when compared to a Hankel matrix with the same number of columns.

Performing experiments to obtain the range of ϵ resulting in stable flight (Figures 2(b) and 2(c)) may not always be possible. We refer the reader to [5, Section 7.2.2] for methods

approximating an optimal ϵ using the data already collected.

2) T : We varied the number of data points used in the data matrices constructed in (3). In particular, we studied the effect of adding additional columns into the data matrices with $N = 1$, $T_{\text{ini}} = 6$, and choosing $\epsilon \in \{a \times 10^{-b} \mid a \in \{1, \dots, 9\}, b \in \{2, 3, 4\}\}$ which minimizes the tracking error. In these simulations the number of data points used to construct the Page matrices did not meet the theoretical minimum required by Theorem 2.1, but still exhibits good tracking performance. Figure 2(d) indicates that increasing the amount of data in the data matrices can significantly improve the tracking performance. Past a certain threshold (200 columns for Page matrix structure, 300 columns for Hankel matrix structure), the tracking performance remains approximately constant. This is likely due to the fact that with this amount of data, the data matrices characterize a rich enough subspace of trajectories well approximating the nonlinear dynamics. This intuition is drawn from the fact that nonlinear dynamics (under certain assumptions) can be lifted to large (often infinite) dimensional linear dynamics, and a large enough data set well approximates the dominant modes.

Aside from comparing the number of columns (and thus the size of the optimization variables), we also compare the Hankel and Page matrices given the same fixed amount of data T . A representative comparison is presented in Table I. The mean solve time represents the average time in seconds it took to solve (7) with $N = 1$, $T_{\text{ini}} = 6$ and choosing $\epsilon \in \{a \times 10^{-b} \mid a \in \{1, \dots, 9\}, b \in \{2, 3, 4\}\}$ which minimizes the tracking error. The optimization was performed using MOSEK [51] on a 3.4 GHz Intel Core i5 with 16GB of RAM.

When $T = 527$, the number of columns in the Hankel and

Fig. 2: (a) Step trajectory of quadcopter using distributionally robust DeePC Algorithm 1; (b)–(c) Dependence of tracking error on the Wasserstein radius ϵ ; (d) Dependence of tracking error on the number of columns in the data matrix; (e) Comparison of tracking error for Page and Hankel matrices; (f) Dependence of tracking error on the initial condition horizon T_{ini} .

TABLE I: Comparison of mean solve times and tracking error for Hankel and Page matrices for varying data lengths.

Data length	Mean solve time (s)		Tracking error	
	$T = 527$	$T = 15407$	$T = 527$	$T = 15407$
Hankel	0.182	3.70	33.94	34.35
Page	0.007	0.166	7×10^6	33.97

Page matrices are 497 and 17 respectively. When $T = 15407$, the number of columns in the Hankel and Page matrices are 15377 and 497 respectively. Each additional $T_{\text{ini}} + T_f$ data points allow the addition of one column to the Page matrix and $T_{\text{ini}} + T_f$ columns to the Hankel matrix. Hence, the size of the optimization variable g in (7) when using a Hankel matrix increases $T_{\text{ini}} + T_f$ times faster than that of a Page matrix.

In terms of tracking error, we see that for $T = 527$ the performance of the Page matrix is poor (i.e., the quadcopter crashes) since there are not enough columns in the Page matrix to construct a rich enough subspace of trajectories. Increasing to $T = 15407$ results in good performance for both the Hankel and Page matrices at the cost of a much larger solve time for the Hankel matrix. Furthermore, the tracking errors confirm the observation from Figure 2(d) that there is a threshold on the number of data points above which the tracking errors remain approximately constant.

3) *Data matrix structure*: We performed simulations with hyperparameters $N \in \{1, 2, 5, 10\}$, $T_{\text{ini}} = 6$, number of columns in data matrix ranging from 150 to 500, and $\epsilon \in \{a \times 10^{-b} \mid a \in \{1, \dots, 9\}, b \in \{2, 3, 4\}\}$. Over all simulations, the same data set was used to populate the Page/Hankel matrices constructed in (3). The 273 simulations with tracking error no larger than 200 (i.e., those exhibiting stable flight of the quadcopter) are plotted in the histogram in Figure 2(e). The Page matrix significantly outperforms the Hankel matrix in terms of tracking performance. This observation may be due to the fact that the tractable reformulation of the distributionally robust DeePC objective in (7) is tight for the Page matrix. Additionally, the entries of the Hankel matrix are repeated which may result in a higher sensitivity to noise compared to the Page matrix whose entries are not repeated. Lastly, we note that using a Page matrix allows one to perform singular value decomposition (SVD) to optimally preprocess the data. This is not the case for the Hankel matrix since SVD does not preserve Hankel structure. See [52] for a relevant case study.

4) T_{ini} : We varied the horizon over which the initial condition is implicitly estimated while fixing $N = 1$, and 300 columns in the data matrix. For each value of $T_{\text{ini}} \in \{1, \dots, 10\}$, a simulation was performed for every $\epsilon \in \{a \times 10^{-b} \mid a \in \{1, \dots, 9\}, b \in \{2, 3, 4\}\}$. For the best performing ϵ , the value of the tracking error for each T_{ini} is reported in Figure 2(f). Once T_{ini} is past a certain threshold ($T_{\text{ini}} = 5$), the distributionally robust DeePC algorithm exhibits satisfactory tracking performance. This is because T_{ini} is the main parameter to fix the “system complexity” inside the DeePC algorithm (see Lemma 2.1).

VI. CONCLUSION

We presented a data-driven method for controlling stochastic constrained LTI systems only using raw data collected from

the system and without the need to explicitly identify a model. The method comes with strong out-of-sample performance guarantees and ensures constraint satisfaction with high confidence. Furthermore, we discussed how the performance of the algorithm is affected by the hyperparameters. Future work includes exploring the extension of this method to strongly nonlinear systems, extending to closed-loop guarantees, and comparison to identification-based control approaches. Furthermore, our Page excitation condition is only sufficient, and we seek relaxed conditions. Finally, we are interested in an online adaptation of the data matrix.

ACKNOWLEDGMENT

The authors would like to thank Linbin Huang, Peyman Mohajerin Esfahani, Claudio de Persis, Ivan Markovsky, and Manfred Morari for useful discussions, as well as Ezzat Elokda and Paul Beuchat for providing the simulation model.

APPENDIX

Proof of Theorem 2.1: The “if” part of the statement is obvious by linearity and time-invariance of the system. We now prove the “only if” part. Our proof strategy is partially inspired by arguments in [18]. Let $\bar{u}_{[1,L]} = (\bar{u}_1, \dots, \bar{u}_L)$, $\bar{y}_{[1,L]} = (\bar{y}_1, \dots, \bar{y}_L)$ be a trajectory of $\mathcal{B}$. Hence, there exist matrices A, B, C, D such that $(\bar{u}_{[1,L]}, \bar{y}_{[1,L]})$ is a trajectory of $\mathcal{B}(A, B, C, D)$ with initial state $\bar{x}_1 \in \mathbb{R}^{n(\mathcal{B})}$. For ease of notation, throughout the proof we will denote $n(\mathcal{B})$ by n . Define the matrix of state sequences from x_i to x_{i+jL} as

$$X_{i,j} = (x_i \quad x_{i+L} \quad \dots \quad x_{i+jL}).$$

Then by the system dynamics we have that

$$\begin{pmatrix} \mathcal{P}_L(u_{[1,T]}) \\ \mathcal{P}_L(y_{[1,T]}) \end{pmatrix} = \begin{pmatrix} 0 & I \\ \mathcal{O}_L(A, C) & \mathcal{T}_L(A, B, C, D) \end{pmatrix} \begin{pmatrix} X_{1, \lfloor \frac{T}{L} \rfloor - 1} \\ \mathcal{P}_L(u_{[1,T]}) \end{pmatrix},$$

where I is the identity matrix. Furthermore,

$$\begin{pmatrix} \bar{u}_{[1,L]} \\ \bar{y}_{[1,L]} \end{pmatrix} = \begin{pmatrix} 0 & I \\ \mathcal{O}_L(A, C) & \mathcal{T}_L(A, B, C, D) \end{pmatrix} \begin{pmatrix} \bar{x}_1 \\ \bar{u}_{[1,L]} \end{pmatrix}.$$

If there exists a vector g such that

$$\begin{pmatrix} \bar{x}_1 \\ \bar{u}_{[1,L]} \end{pmatrix} = \begin{pmatrix} X_{1, \lfloor \frac{T}{L} \rfloor - 1} \\ \mathcal{P}_L(u_{[1,T]}) \end{pmatrix} g, \quad (9)$$

then

$$\begin{pmatrix} \bar{u}_{[1,L]} \\ \bar{y}_{[1,L]} \end{pmatrix} = \begin{pmatrix} 0 & I \\ \mathcal{O}_L(A, C) & \mathcal{T}_L(A, B, C, D) \end{pmatrix} \begin{pmatrix} X_{1, \lfloor \frac{T}{L} \rfloor - 1} \\ \mathcal{P}_L(u_{[1,T]}) \end{pmatrix} g = \begin{pmatrix} \mathcal{P}_L(u_{[1,T]}) \\ \mathcal{P}_L(y_{[1,T]}) \end{pmatrix} g,$$

which is what we wanted to show. Hence it remains to prove that there always exists a vector g satisfying (9) for any $\bar{x}_1$ and $\bar{u}_{[1,L]}$. Equivalently, we will show that the matrix $\begin{pmatrix} X_{1, \lfloor \frac{T}{L} \rfloor - 1} \\ \mathcal{P}_L(u_{[1,T]}) \end{pmatrix}$ has full row rank. Let $(\xi \quad \eta)$ be an arbitrary vector in the left kernel of this matrix where $\xi^\top \in \mathbb{R}^n$ and

$\eta^\top \in \mathbb{R}^{mL}$. We will show that $(\xi \ \eta)$ must be the zero vector. To do this, we switch our attention to the matrix

$$\begin{pmatrix} X \\ U \end{pmatrix} := \begin{pmatrix} X_{1, \lfloor \frac{T}{L} \rfloor - n - 1} \\ \mathcal{P}_L(u_{[1, T - nL]}) \\ \mathcal{P}_L(u_{[L+1, T - (n-1)L]}) \\ \vdots \\ \mathcal{P}_L(u_{[nL+1, T]}) \end{pmatrix} \quad (10)$$

and follow arguments from [18] to construct $n+1$ vectors in the left kernel of (10). By definition, $(\xi \ \eta \ \mathbb{0}_{nmL})$ is in the left kernel of (10), where $\mathbb{0}_{nmL}^\top \in \mathbb{R}^{nmL}$ denotes a vector containing nmL zeros. Furthermore, we have that

$$\begin{pmatrix} X_{L+1, \lfloor \frac{T}{L} \rfloor - n - 1} \\ \mathcal{P}_L(u_{[L+1, T - (n-1)L]}) \end{pmatrix} = M \begin{pmatrix} X \\ U \end{pmatrix}, \quad (12)$$

where M is defined in (11) and $\mathbb{0}_{n \times n}$ is the zero matrix. Since $\begin{pmatrix} X_{L+1, \lfloor \frac{T}{L} \rfloor - n - 1} \\ \mathcal{P}_L(u_{[L+1, T - (n-1)L]}) \end{pmatrix}$ is a submatrix of $\begin{pmatrix} X_{1, \lfloor \frac{T}{L} \rfloor - 1} \\ \mathcal{P}_L(u_{[1, T]}) \end{pmatrix}$ then by definition of $(\xi \ \eta)$, we have

$$(\xi \ \eta) \begin{pmatrix} X_{L+1, \lfloor \frac{T}{L} \rfloor - n - 1} \\ \mathcal{P}_L(u_{[L+1, T - (n-1)L]}) \end{pmatrix} = \mathbb{0}_{(\lfloor \frac{T}{L} \rfloor - n - 1)L + 1}.$$

Hence, by (12) we have that

$$(\xi A^L \ \xi A^{L-1}B \ \dots \ \xi AB \ \xi B \ \eta \ \mathbb{0}_{(n-1)mL})$$

is in the left kernel of (10). Note also that

$$\begin{aligned} & X_{2L+1, \lfloor \frac{T}{L} \rfloor - n - 1} \\ &= (A^{2L} \ A^{2L-1}B \ \dots \ AB \ B \ \mathbb{0}_{n \times (n-1)mL}) \begin{pmatrix} X \\ U \end{pmatrix}. \end{aligned}$$

Proceeding as above we find that the vector

$$(\xi A^{2L} \ \xi A^{2L-1}B \ \dots \ \xi AB \ \xi B \ \eta \ \mathbb{0}_{(n-2)mL})$$

is in the left kernel of (10). Continuing in this way, we obtain $n+1$ vectors in the left kernel of (10) given by

$$\begin{aligned} w_0 &= (\xi \ \eta \ \mathbb{0}_{nmL}) \\ w_1 &= (\xi A^L \ \xi A^{L-1}B \ \dots \ \xi AB \ \xi B \ \eta \ \mathbb{0}_{(n-1)mL}) \\ w_2 &= (\xi A^{2L} \ \xi A^{2L-1}B \ \dots \ \xi AB \ \xi B \ \eta \ \mathbb{0}_{(n-2)mL}) \\ &\vdots \\ w_n &= (\xi A^{nL} \ \xi A^{nL-1}B \ \dots \ \xi AB \ \xi B \ \eta) \end{aligned}$$

However, since the input vectors are L -Page exciting of order $n+1$ the left kernel of (10) has dimension at most n . This implies that the vectors $w_0, \dots, w_n$ are linearly dependent. From the structure of the vectors we see that $\eta = \mathbb{0}_{mL}$. By

Cayley-Hamilton, there exist $\alpha_i \in \mathbb{R}$, $i \in \{0, \dots, n\}$ with $\alpha_n = 1$ such that $\sum_{i=0}^n \alpha_i (A^L)^i = 0$. Define

$$\begin{aligned} v &= \sum_{i=0}^n \alpha_i w_i \\ &= \left(\xi \sum_{i=0}^n \alpha_i (A^L)^i \ \dots \ \xi A^{L-1}B \alpha_n \ \dots \ \xi B \alpha_n \ \mathbb{0}_{mL} \right) \\ &= \left(\mathbb{0}_n \ \xi \sum_{i=1}^n \alpha_i A^{iL-1}B \ \dots \ \xi A^{L-1}B \ \dots \ \xi B \ \mathbb{0}_{mL} \right). \end{aligned}$$

Since v is a linear combination of vectors w_i then it is itself in the left kernel of (10). Hence, $(\xi \sum_{i=1}^n \alpha_i A^{iL-1}B \ \dots \ \xi A^{L-1}B \ \dots \ \xi B \ \mathbb{0}_{mL})$ is in the left kernel of U . However, U has full row rank by the excitation assumption implying the above vector must be the zero vector. Thus, $(\xi A^{L-1}B \ \dots \ B) = \mathbb{0}_{mL}$. By assumption we have that $L \geq n$ and $\mathcal{B}$ is controllable. Thus $\xi = \mathbb{0}_n$. Hence, $(\xi \ \eta) = \mathbb{0}_{n+mL}$ implying that (10) has full row rank which concludes the proof. ■

Proof of Proposition 3.1: We first look at the feasible set of (2). By rewriting the constraints in (2) we obtain

$$y = \mathcal{O}_{T_f}(A, C) \hat{x}_t + \mathcal{T}_{T_f}(A, B, C, D)u, \quad u \in \mathcal{U}, y \in \mathcal{Y},$$

where $\hat{x}_t$ is the state at time t . We now look at the feasible set of (5). Since $\hat{u}_{[1, T]}$ is $(T_{\text{ini}} + T_f)$ -Page exciting of order $n(\mathcal{B})+1$, then by Theorem 2.1 $\text{image}(\text{col}(\hat{U}_p, \hat{Y}_p, \hat{U}_f, \hat{Y}_f)) = \mathcal{B}_{T_{\text{ini}}+T_f}$, where $\hat{U}_p, \hat{Y}_p, \hat{U}_f, \hat{Y}_f$ are defined as in (3). Hence, the feasible set of (5) is equal to $\{(u, y) \in \mathcal{U} \times \mathcal{Y} \mid \text{col}(\hat{u}_{\text{ini}}, u, \hat{y}_{\text{ini}}, y) \in \mathcal{B}_{T_{\text{ini}}+T_f}\}$. Since the system $\mathcal{B}$ yields an equivalent representation given by $\mathcal{B}(A, B, C, D)$, then by Lemma 2.1 the feasible set of (5) can be written as the set of pairs $(u, y) \in \mathcal{U} \times \mathcal{Y}$ satisfying

$$y = \mathcal{O}_{T_f}(A, C) \hat{x}_{\text{ini}} + \mathcal{T}_{T_f}(A, B, C, D)u,$$

where $\hat{x}_{\text{ini}}$ is uniquely determined from $\text{col}(\hat{u}_{\text{ini}}, \hat{y}_{\text{ini}})$ and hence coincides with $\hat{x}_t$. This proves the claim. ■

The following two lemmas will serve as essential cornerstones of our subsequent results.

Lemma A.1: Assume that φ is proper, convex, and lower semicontinuous. Then

$$\sup_{Q \in B_\epsilon(\hat{\mathbb{P}}_N)} \mathbb{E}_Q[\varphi(\xi)] = \inf_{\lambda \geq 0} \lambda \epsilon + \frac{1}{N} \sum_{i=1}^N \sup_{\xi \in \Xi} (\varphi(\xi) - \lambda \|\xi - \hat{\xi}^{(i)}\|_r)$$

Proof: The proof follows from [5, Theorem 4.2], specifically from Equation (12b) in [5] which relies on a marginalization and dualization of the constraint $Q \in B_\epsilon(\hat{\mathbb{P}}_N)$. ■ Next we present a similar result as [2, Lemma 4.1]. The result relaxes the conservative assumption that $\Xi = \prod_{i=1}^{p(T_{\text{ini}}+T_f)} \Xi_i$ for

$$M := \begin{pmatrix} A^L & A^{L-1}B & \dots & AB & B & \mathbb{0}_{n \times mL} & \dots & \mathbb{0}_{n \times mL} \\ \mathbb{0}_{m \times n} & \mathbb{0}_{m \times m} & \dots & \mathbb{0}_{m \times m} & \mathbb{0}_{m \times m} & [I \ \dots \ \mathbb{0}_{m \times m}] & \dots & [\mathbb{0}_{m \times m} \ \dots \ \mathbb{0}_{m \times m}] \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \\ \mathbb{0}_{m \times n} & \mathbb{0}_{m \times m} & \dots & \mathbb{0}_{m \times m} & \mathbb{0}_{m \times m} & [\mathbb{0}_{m \times m} \ \dots \ I] & \dots & [\mathbb{0}_{m \times m} \ \dots \ \mathbb{0}_{m \times m}] \end{pmatrix} \quad (11)$$

appropriately defined Ξ_i imposed in [2]. The proof strategy is partially inspired by arguments in [5].

Lemma A.2: Let $c, d \in \mathbb{Z}_{>0}$ and $\Omega \subseteq \mathbb{R}^{cd}$. Let q be such that $\frac{1}{r} + \frac{1}{q} = 1$. Let $\hat{\zeta} = \text{col}(\hat{\zeta}_1, \dots, \hat{\zeta}_c) \in \Omega$ be given where each $\hat{\zeta}_i \in \mathbb{R}^d$. Let $\lambda \in \mathbb{R}_{>0}$ and $\varphi: \mathbb{R}^c \rightarrow \mathbb{R}$ be a convex and Lipschitz continuous function with Lipschitz constant L_φ with respect to the r -norm. Then for $b \in \mathbb{R}^d$,

$$\begin{aligned} & \sup_{\zeta \in \Omega} \varphi(\zeta_1^\top b, \dots, \zeta_c^\top b) - \lambda \|\zeta - \hat{\zeta}\|_r \\ & \leq \begin{cases} \varphi(\hat{\zeta}_1^\top b, \dots, \hat{\zeta}_c^\top b) & \text{if } L_\varphi \|b\|_q \leq \lambda \\ \infty & \text{otherwise} \end{cases} \end{aligned}$$

The above is an equality when $\Omega = \mathbb{R}^{cd}$.

Proof: We begin by first noting that

$$\begin{aligned} & \sup_{\zeta \in \Omega} \varphi(\zeta_1^\top b, \dots, \zeta_c^\top b) - \lambda \|\zeta - \hat{\zeta}\|_r \\ & \leq \sup_{\zeta \in \mathbb{R}^{cd}} \varphi(\zeta_1^\top b, \dots, \zeta_c^\top b) - \lambda \|\zeta - \hat{\zeta}\|_r \end{aligned}$$

with equality when $\Omega = \mathbb{R}^{cd}$. Let $\Phi(\zeta) = \varphi(\zeta_1^\top b, \dots, \zeta_c^\top b)$. By definition of the conjugate function,

$$\begin{aligned} \Phi^*(z) &= \sup_{\zeta \in \mathbb{R}^{cd}} \langle z, \zeta \rangle - \Phi(\zeta) \\ &= \sup_{\zeta \in \mathbb{R}^{cd}} \langle z, \zeta \rangle - \varphi(\zeta_1^\top b, \dots, \zeta_c^\top b) \\ &= \sup_{\zeta \in \mathbb{R}^{cd}, s \in \mathbb{R}^c} \left\{ \langle z, \zeta \rangle - \varphi(s) \mid s_j = \zeta_j^\top b, \forall j \in \{1, \dots, c\} \right\}, \end{aligned}$$

where $s = \text{col}(s_1, \dots, s_c)$. The Lagrangian of the above is given by

$$\mathcal{L}(s, \zeta, \theta) = \langle z, \zeta \rangle + \langle \theta, (s_1 - \zeta_1^\top b, \dots, s_c - \zeta_c^\top b) \rangle - \varphi(s),$$

By strong duality (see, e.g., [53, Proposition 5.3.1]),

$$\begin{aligned} \Phi^*(z) &= \inf_{\theta} \sup_{\zeta \in \mathbb{R}^{cd}, s} \langle z, \zeta \rangle \\ & \quad + \langle \theta, (s_1 - \zeta_1^\top b, \dots, s_c - \zeta_c^\top b) \rangle - \varphi(s) \\ &= \inf_{\theta} \sup_{\zeta \in \mathbb{R}^{cd}} \langle z, \zeta \rangle - \langle \theta, (\zeta_1^\top b, \dots, \zeta_c^\top b) \rangle + \varphi^*(\theta) \\ &= \inf_{\theta} \sup_{\zeta \in \mathbb{R}^{cd}} \langle (z_1 - \theta_1 b, \dots, z_c - \theta_c b), \zeta \rangle + \varphi^*(\theta) \end{aligned}$$

where $z = \text{col}(z_1, \dots, z_c)$. The above problem may take the value ∞ unless $z_i = \theta_i b$ for all $i \in \{1, \dots, c\}$. Hence, by taking $\Theta := \{\theta \mid \varphi^*(\theta) < \infty\}$ as the effective domain of the conjugate function of φ , we obtain

$$\begin{aligned} \Phi^*(z) &= \inf_{\theta} \{\varphi^*(\theta) \mid z_i = \theta_i b, \forall i \in \{1, \dots, c\}\} \\ &= \inf_{\theta \in \Theta} \{\varphi^*(\theta) \mid z_i = \theta_i b, \forall i \in \{1, \dots, c\}\}. \end{aligned}$$

Since Φ is convex and continuous, the biconjugate Φ^{**} coincides with the function Φ itself. Hence,

$$\begin{aligned} \Phi(\zeta) &= \sup_z \langle z, \zeta \rangle - \Phi^*(z) \\ &= \sup_z \left\{ \langle z, \zeta \rangle - \inf_{\theta \in \Theta} \varphi^*(\theta) \mid z_i = \theta_i b, \forall i \in \{1, \dots, c\} \right\} \\ &= \sup_{\theta \in \Theta} \langle (\theta_1 b, \dots, \theta_c b), \zeta \rangle - \varphi^*(\theta). \end{aligned}$$

Hence,

$$\begin{aligned} & \sup_{\zeta \in \mathbb{R}^{cd}} \Phi(\zeta) - \lambda \|\zeta - \hat{\zeta}\|_r \\ &= \sup_{\zeta \in \mathbb{R}^{cd}} \sup_{\theta \in \Theta} \langle (\theta_1 b, \dots, \theta_c b), \zeta \rangle - \varphi^*(\theta) - \lambda \|\zeta - \hat{\zeta}\|_r \\ &= \sup_{\zeta \in \mathbb{R}^{cd}} \sup_{\theta \in \Theta} \inf_{\|\mu\|_q \leq \lambda} \langle (\theta_1 b, \dots, \theta_c b), \zeta \rangle - \varphi^*(\theta) - \langle \mu, \zeta - \hat{\zeta} \rangle, \end{aligned}$$

where the last equality comes from the definition of the dual norm and homogeneity of the norm. Using the minimax theorem (see, e.g., [53, Proposition 5.5.4]) we switch the supremum and infimum in the above giving

$$\begin{aligned} & \sup_{\zeta \in \mathbb{R}^{cd}} \Phi(\zeta) - \lambda \|\zeta - \hat{\zeta}\|_r \\ &= \sup_{\theta \in \Theta} \inf_{\|\mu\|_q \leq \lambda} \sup_{\zeta \in \mathbb{R}^{cd}} \langle (\theta_1 b, \dots, \theta_c b), \zeta \rangle - \varphi^*(\theta) - \langle \mu, \zeta - \hat{\zeta} \rangle \\ &= \sup_{\theta \in \Theta} \inf_{\|\mu\|_q \leq \lambda} \sup_{\zeta \in \mathbb{R}^{cd}} \langle (\theta_1 b - \mu_1, \dots, \theta_c b - \mu_c), \zeta \rangle \\ & \quad - \varphi^*(\theta) + \langle \mu, \hat{\zeta} \rangle, \end{aligned}$$

where $\mu = \text{col}(\mu_1, \dots, \mu_c)$. Carrying out the supremum over ζ yields

$$\begin{aligned} & \sup_{\zeta \in \mathbb{R}^{cd}} \Phi(\zeta) - \lambda \|\zeta - \hat{\zeta}\|_r \\ &= \sup_{\theta \in \Theta} \inf_{\|\mu\|_q \leq \lambda} \begin{cases} \langle (\theta_1 b, \dots, \theta_c b), \hat{\zeta} \rangle - \varphi^*(\theta) & \text{if } \mu_i = \theta_i b \\ \infty & \text{otherwise} \end{cases} \quad \forall i \in \{1, \dots, c\} \\ &= \sup_{\theta \in \Theta} \begin{cases} \langle (\theta_1 b, \dots, \theta_c b), \hat{\zeta} \rangle - \varphi^*(\theta) & \text{if } \|(\theta_1 b, \dots, \theta_c b)\|_q \leq \lambda \\ \infty & \text{otherwise} \end{cases} \\ &= \begin{cases} \varphi(\hat{\zeta}_1^\top b, \dots, \hat{\zeta}_c^\top b) & \text{if } \sup_{\theta \in \Theta} \|(\theta_1 b, \dots, \theta_c b)\|_q \leq \lambda \\ \infty & \text{otherwise} \end{cases}, \end{aligned}$$

where we used the definition of the biconjugate function and $\varphi^{**} = \varphi$ (due to convexity and continuity). Note that $\|(\theta_1 b, \dots, \theta_c b)\|_q = \|\theta\|_q \|b\|_q$. From [5, Proposition 6.5], we know that $\|\theta\|_q \leq L_\varphi$. In fact, one can show that $\sup_{\theta \in \Theta} \|\theta\|_q = L_\varphi$ [49, Remark 3]. Substituting this into the expression above proves the result. ■

The following lemmas enable our main results and provide separate reformulations for the objective function and constraints. The proofs strategies are partially inspired by [5], [47], [54].

Lemma A.3: (Objective Reformulation): Assume that f_2 and f_3 are convex and Lipschitz continuous functions. Specifically, let $L_{\text{obj}} > 0$ be the Lipschitz constant with respect to the r -norm of the mapping $(x, y) \mapsto f_2(x) + f_3(y)$. Then

$$\begin{aligned} & \sup_{\mathbb{Q} \in \mathcal{B}_\epsilon(\mathbb{P}_N)} \mathbb{E}_{\mathbb{Q}} \left[f_1(\widehat{U}_f g) + f_2(Y_f g) + f_3(Y_p g - \hat{y}_{\text{ini}}) \right] \leq \\ & f_1(\widehat{U}_f g) + \frac{1}{N} \sum_{i=1}^N \left(f_2(\widehat{Y}_f^{(i)} g) + f_3(\widehat{Y}_p^{(i)} g - \hat{y}_{\text{ini}}) \right) + L_{\text{obj}} \epsilon \|g\|_{r,*} \end{aligned}$$

The above is an equality when $\Xi = \mathbb{R}^{p(T_{\text{ini}} + T_f) \lfloor \frac{T}{T_{\text{ini}} + T_f} \rfloor}$.

Proof: By definition of ξ , we have $f_2(Y_f g) + f_3(Y_p g - \hat{y}_{\text{ini}}) = f_2(\xi_{p(T_{\text{ini}}+1)}^\top g, \dots, \xi_{p(T_{\text{ini}}+T_f)}^\top g) + f_3((\xi_1^\top g, \dots, \xi_{p(T_{\text{ini}})}^\top g) - \hat{y}_{\text{ini}})$. Define

$$\varphi(\xi_1^\top g, \dots, \xi_{p(T_{\text{ini}}+T_f)}^\top g) := f_2(\xi_{p(T_{\text{ini}}+1)}^\top g, \dots, \xi_{p(T_{\text{ini}}+T_f)}^\top g) + f_3((\xi_1^\top g, \dots, \xi_{p(T_{\text{ini}})}^\top g) - \hat{y}_{\text{ini}})$$

By definition, we have $L_\varphi = L_{\text{obj}}$. Define also $\Phi(\xi) := \varphi(\xi_1^\top g, \dots, \xi_{p(T_{\text{ini}}+T_f)}^\top g)$. Then from Lemma A.1 applied to Φ we obtain

$$\begin{aligned} & \sup_{\mathbb{Q} \in B_\epsilon(\hat{\mathbb{P}}_N)} \mathbb{E}_{\mathbb{Q}} [f_1(\hat{U}_f g) + f_2(Y_f g) + f_3(Y_p g - \hat{y}_{\text{ini}})] \\ &= \inf_{\lambda \geq 0} \lambda \epsilon + f_1(\hat{U}_f g) \\ &+ \frac{1}{N} \sum_{i=1}^N \sup_{\xi \in \Xi} (\varphi(\xi_1^\top g, \dots, \xi_{p(T_{\text{ini}}+T_f)}^\top g) - \lambda \|\xi - \hat{\xi}^{(i)}\|_r) \end{aligned}$$

Applying Lemma A.2 yields

$$\begin{aligned} & \sup_{\mathbb{Q} \in B_\epsilon(\hat{\mathbb{P}}_N)} \mathbb{E}_{\mathbb{Q}} [f_1(\hat{U}_f g) + f_2(Y_f g) + f_3(Y_p g - \hat{y}_{\text{ini}})] \leq \\ & \epsilon L_\varphi \|g\|_{r,*} + f_1(\hat{U}_f g) + \frac{1}{N} \sum_{i=1}^N \left(f_2(\hat{Y}_f^{(i)} g) + f_3(\hat{Y}_p^{(i)} g - \hat{y}_{\text{ini}}) \right). \end{aligned}$$

The above is an equality when $\Xi = \mathbb{R}^{p(T_{\text{ini}}+T_f)} \lfloor \frac{T}{T_{\text{ini}}+T_f} \rfloor$. Substituting $L_\varphi = L_{\text{obj}}$ yields the result. ■

The following two results depend on the constraint set

$$G_{\text{CVaR}} := \left\{ g \mid \sup_{\mathbb{Q} \in B_\epsilon(\hat{\mathbb{P}}_N)} \text{CVaR}_{1-\alpha}^{\mathbb{Q}}(h(Y_f g)) \leq 0 \right\}.$$

Lemma A.4: (Lipschitz Constraint Reformulation): Let h be convex and Lipschitz continuous with Lipschitz constant L_{con} with respect to the r -norm. Then

$$\left\{ g \mid \begin{array}{l} \exists \tau, \lambda, s_i \text{ such that} \\ -\tau\alpha + L_{\text{con}}\epsilon \|g\|_{r,*} + \frac{1}{N} \sum_{i=1}^N s_i \leq 0 \\ \tau + h(\hat{Y}_f^{(i)} g) \leq s_i \\ s_i \geq 0 \\ \forall i \leq N \end{array} \right\} \subseteq G_{\text{CVaR}}.$$

The sets coincide if $\Xi = \mathbb{R}^{p(T_{\text{ini}}+T_f)} \lfloor \frac{T}{T_{\text{ini}}+T_f} \rfloor$ and $h(Y_f g)$ is bounded on Ξ for all g .

Remark A.1: Note that if $\Xi = \mathbb{R}^{p(T_{\text{ini}}+T_f)} \lfloor \frac{T}{T_{\text{ini}}+T_f} \rfloor$, then $h(Y_f g)$ can only be bounded on Ξ for all g if h is constant. Thus, for non-constant constraint functions h , the above will be an inner approximation of G_{CVaR} . •

Proof of Lemma A.4: By definition of CVaR,

$$\text{CVaR}_{1-\alpha}^{\mathbb{P}}(h(y)) \leq 0 \iff \inf_{\tau \in \mathbb{R}} -\tau\alpha + \mathbb{E}_{\mathbb{P}}[(h(y) + \tau)_+] \leq 0$$

This can be shown by multiplying the definition of CVaR by $\alpha > 0$ and changing τ to $-\tau$. Letting $\ell(g, \xi) := h(\xi_{p(T_{\text{ini}}+1)}^\top g, \dots, \xi_{p(T_{\text{ini}}+T_f)}^\top g)$, the constraint $g \in G_{\text{CVaR}}$ is equivalent to g satisfying

$$\sup_{\mathbb{Q} \in B_\epsilon(\hat{\mathbb{P}}_N)} \inf_{\tau \in \mathbb{R}} -\tau\alpha + \mathbb{E}_{\mathbb{Q}}[(\ell(g, \xi) + \tau)_+] \leq 0. \quad (13)$$

We have

$$\begin{aligned} & \sup_{\mathbb{Q} \in B_\epsilon(\hat{\mathbb{P}}_N)} \inf_{\tau \in \mathbb{R}} -\tau\alpha + \mathbb{E}_{\mathbb{Q}}[(\ell(g, \xi) + \tau)_+] \\ & \leq \inf_{\tau \in \mathbb{R}} -\tau\alpha + \sup_{\mathbb{Q} \in B_\epsilon(\hat{\mathbb{P}}_N)} \mathbb{E}_{\mathbb{Q}}[(\ell(g, \xi) + \tau)_+] \end{aligned} \quad (14)$$

by the max-min inequality. From Lemma A.1 we have

$$\begin{aligned} & \sup_{\mathbb{Q} \in B_\epsilon(\hat{\mathbb{P}}_N)} \mathbb{E}_{\mathbb{Q}}[(\ell(g, \xi) + \tau)_+] \\ &= \inf_{\lambda \geq 0} \lambda \epsilon + \frac{1}{N} \sum_{i=1}^N \sup_{\xi \in \Xi} ((\ell(g, \xi) + \tau)_+ - \lambda \|\xi - \hat{\xi}^{(i)}\|_r). \end{aligned}$$

Recalling that $\ell(g, \xi) = h(\xi_{p(T_{\text{ini}}+1)}^\top g, \dots, \xi_{p(T_{\text{ini}}+T_f)}^\top g)$ and using Lemma A.2 gives

$$\begin{aligned} & \sup_{\mathbb{Q} \in B_\epsilon(\hat{\mathbb{P}}_N)} \mathbb{E}_{\mathbb{Q}}[(\ell(g, \xi) + \tau)_+] \\ & \leq \begin{cases} \inf_{\lambda \geq 0} \lambda \epsilon + \frac{1}{N} \sum_{i=1}^N (\ell(g, \hat{\xi}^{(i)}) + \tau)_+ & \text{if } L_{\text{con}} \|g\|_{r,*} \leq \lambda \\ \infty & \text{otherwise} \end{cases} \end{aligned} \quad (15)$$

Hence, by resolving the inf and resorting to an epigraph formulation, we arrive at

$$\begin{aligned} & \inf_{\tau \in \mathbb{R}} -\tau\alpha + \sup_{\mathbb{Q} \in B_\epsilon(\hat{\mathbb{P}}_N)} \mathbb{E}_{\mathbb{Q}}[(\ell(g, \xi) + \tau)_+] \\ & \leq \begin{cases} \inf_{\tau \in \mathbb{R}, s_i \in \mathbb{R}} -\tau\alpha + \epsilon L_{\text{con}} \|g\|_{r,*} + \frac{1}{N} \sum_{i=1}^N s_i \\ \text{s.t. } \tau + \ell(g, \hat{\xi}^{(i)}) \leq s_i \\ s_i \geq 0 \\ \forall i \leq N \end{cases} \end{aligned}$$

Substituting the definition for ℓ and recalling from (13) that we wish for the right hand side of the above inequality to be nonnegative gives

$$\left\{ g \mid \begin{array}{l} \inf_{\tau, s_i} -\tau\alpha + L_{\text{con}}\epsilon \|g\|_{r,*} + \frac{1}{N} \sum_{i=1}^N s_i \leq 0 \\ \tau + h(\hat{Y}_f^{(i)} g) \leq s_i \\ s_i \geq 0 \\ \forall i \leq N \end{array} \right\} \subseteq G_{\text{CVaR}}.$$

We can remove the infimum from the above formulation by noting that there exist variables τ, s_i satisfying the constraints above if and only if the above infimum constraint holds. The “only if” part is obvious. The “if” part can be split into two cases: the infimum is achieved, or the infimum is not achieved. If the infimum is achieved then the optimizer of the infimum satisfies the above. If the infimum is not achieved (i.e., the infimum is $-\infty$), then the first constraint is trivially satisfied, and one can find variables τ, s_i satisfying the remaining constraints. Note that (14) is an equality if $g \mapsto \ell(g, \xi)$ is convex for every ξ and $\xi \mapsto \ell(g, \xi)$ is bounded on Ξ for every g (see, [54, Theorem 2.2]). Furthermore, (15) is an equality if $\Xi = \mathbb{R}^{p(T_{\text{ini}}+T_f)} \lfloor \frac{T}{T_{\text{ini}}+T_f} \rfloor$. Since h is convex, this proves the claim. ■

We present an alternative result for the special case of piecewise affine constraints in which it is possible to obtain a tight set reformulation.

Lemma A.5: (Piecewise Affine Constraint Reformulation): Let $\Xi = \{\xi \mid F\xi \leq d\}$ for some F and d of appropriate dimensions. Let $\ell(g, \xi) := h(\xi_{pT_{\text{ini}}+i}^\top g, \dots, \xi_{p(T_{\text{ini}}+T_i)}^\top g)$. Assume ℓ is piecewise affine in ξ , i.e., $\ell(g, \xi) = \max_{k \leq K} \langle M_k g, \xi \rangle + b_k$ for matrices $M_k, b_k \in \mathbb{R}$ and $K \in \mathbb{Z}_{>0}$. Then

$$\left\{ g \mid \begin{array}{l} \exists \tau, \lambda, s_i, \gamma_{ik} \text{ such that} \\ -\tau\alpha + \lambda\epsilon + \frac{1}{N} \sum_{i=1}^N s_i \leq 0 \\ b_k + \tau + \langle M_k g, \hat{\xi}^{(i)} \rangle \\ \quad + \langle \gamma_{ik}, d - F\hat{\xi}^{(i)} \rangle \leq s_i \\ \|F^\top \gamma_{ik} - M_k g\|_{r,*} \leq \lambda \\ s_i \geq 0 \\ \forall i \leq N, \forall k \leq K \end{array} \right\} \subseteq G_{\text{CVaR}}.$$

The sets coincide when $\ell(g, \xi)$ is bounded on Ξ for all g .

Proof: Recall that the support set of the random variable ξ is defined as $\Xi = \{\xi \mid F\xi \leq d\}$. Substituting $\ell(g, \xi) = \max_{k \leq K} \langle M_k g, \xi \rangle + b_k$ and applying [5, Corollary 5.1] we have that

$$\begin{aligned} & \sup_{\mathbb{Q} \in B_\epsilon(\widehat{\mathbb{P}}_N)} \mathbb{E}_{\mathbb{Q}}[(\ell(g, \xi) + \tau)_+] \\ &= \left\{ \begin{array}{l} \inf_{\lambda, s_i, \gamma_{ik}} \lambda\epsilon + \frac{1}{N} \sum_{i=1}^N s_i \\ \text{s.t. } b_k + \tau + \langle M_k g, \hat{\xi}^{(i)} \rangle + \langle \gamma_{ik}, d - F\hat{\xi}^{(i)} \rangle \leq s_i \\ \|F^\top \gamma_{ik} - M_k g\|_{r,*} \leq \lambda \\ s_i \geq 0 \\ \forall i \leq N, \forall k \leq K \end{array} \right\} \end{aligned}$$

Invoking (13) and (14) and substituting the above yields the following set inclusion

$$\left\{ g \mid \begin{array}{l} \inf_{\tau, \lambda, s_i, \gamma_{ik}} -\tau\alpha + \lambda\epsilon + \frac{1}{N} \sum_{i=1}^N s_i \leq 0 \\ \text{s.t. } b_k + \tau + \langle M_k g, \hat{\xi}^{(i)} \rangle \\ \quad + \langle \gamma_{ik}, d - F\hat{\xi}^{(i)} \rangle \leq s_i \\ \|F^\top \gamma_{ik} - M_k g\|_{r,*} \leq \lambda \\ s_i \geq 0 \\ \forall i \leq N, \forall k \leq K \end{array} \right\} \subseteq G_{\text{CVaR}}.$$

Similarly as in the proof of Lemma A.4, we can replace the infimum from the above formulation with existence of variables $\tau, \lambda, s_i, \gamma_{ik}$ satisfying the constraints above. Hence,

$$\left\{ g \mid \begin{array}{l} \exists \tau, \lambda, s_i, \gamma_{ik} \text{ such that} \\ -\tau\alpha + \lambda\epsilon + \frac{1}{N} \sum_{i=1}^N s_i \leq 0 \\ b_k + \tau + \langle M_k g, \hat{\xi}^{(i)} \rangle \\ \quad + \langle \gamma_{ik}, d - F\hat{\xi}^{(i)} \rangle \leq s_i \\ \|F^\top \gamma_{ik} - M_k g\|_{r,*} \leq \lambda \\ s_i \geq 0 \\ \forall i \leq N, \forall k \leq K \end{array} \right\} \subseteq G_{\text{CVaR}}.$$

If $g \mapsto \ell(g, \xi)$ is convex for every ξ and $\xi \mapsto \ell(g, \xi)$ is bounded on Ξ for every g then (14) reduces to equality (see, [54, Theorem 2.2]) which proves the claim. ■

Proof of Theorem 4.1: Applying Lemma A.3 to the objective function and Lemma A.4 to the constraint of optimization problem (6) gives the result. ■

Proof of Theorem 4.3: Choosing $\epsilon(\beta, N)$ as in (8) ensures that $\mathbb{P}^N(\mathbb{P} \in B_\epsilon(\widehat{\mathbb{P}}_N)) \geq 1 - \beta$. Hence, $J(g)$ is upper bounded

by $\sup_{\mathbb{Q} \in B_\epsilon(\widehat{\mathbb{P}}_N)} \mathbb{E}_{\mathbb{Q}} [f_1(\widehat{U}_f g) + f_2(Y_f g) + f_3(Y_p g - \hat{y}_{\text{ini}})]$ with confidence $1 - \beta$. Furthermore, by Theorem 4.1, $J(g)$ is upper bounded by $\widehat{J}(g)$ with confidence $1 - \beta$ for all g . Choosing $g = \hat{g}^*$ proves the first statement. From Lemma A.4, the constraint set of (7) is a subset of the set of vectors g satisfying $\sup_{\mathbb{Q} \in B_\epsilon(\widehat{\mathbb{P}}_N)} \text{CVaR}_{1-\alpha}^{\mathbb{Q}}(h(Y_f g)) \leq 0$. Invoking the fact that $\mathbb{P}^N(\mathbb{P} \in B_\epsilon(\widehat{\mathbb{P}}_N)) \geq 1 - \beta$ proves the second statement. ■

REFERENCES

- [1] J. Coulson, J. Lygeros, and F. Dörfler, “Data-enabled predictive control: In the shallows of the DeePC,” in *2019 18th European Control Conference (ECC)*, 2019, pp. 307–312.
- [2] —, “Regularized and distributionally robust data-enabled predictive control,” *2019 IEEE Conference on Decision and Control (CDC)*, 2019.
- [3] I. Markovsky and P. Rapisarda, “Data-driven simulation and control,” *International Journal of Control*, vol. 81, no. 12, pp. 1946–1959, 2008.
- [4] J. C. Willems, P. Rapisarda, I. Markovsky, and B. L. De Moor, “A note on persistency of excitation,” *Systems & Control Letters*, vol. 54, no. 4, pp. 325–329, 2005.
- [5] P. Mohajerin Esfahani and D. Kuhn, “Data-driven distributionally robust optimization using the Wasserstein metric: Performance guarantees and tractable reformulations,” *Mathematical Programming*, vol. 171, no. 1-2, pp. 115–166, 2018.
- [6] H. Hjalmarsson, M. Gevers, S. Gunnarsson, and O. Lequin, “Iterative feedback tuning: Theory and applications,” *IEEE control systems magazine*, vol. 18, no. 4, pp. 26–41, 1998.
- [7] M. C. Campi, A. Lecchini, and S. M. Savarese, “Virtual reference feedback tuning: A direct method for the design of feedback controllers,” *Automatica*, vol. 38, no. 8, pp. 1337–1346, 2002.
- [8] B. Recht, “A tour of reinforcement learning: The view from continuous control,” *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 2, pp. 253–279, 2019.
- [9] P. Van Overschee and B. De Moor, *Subspace identification for linear systems: Theory–Implementation–Applications*. Springer Science & Business Media, 2012.
- [10] J. Berberich, J. Köhler, M. A. Muller, and F. Allgower, “Data-driven model predictive control with stability and robustness guarantees,” *IEEE Transactions on Automatic Control*, 2020.
- [11] L. Huang, J. Coulson, J. Lygeros, and F. Dörfler, “Data-enabled predictive control for grid-connected power converters,” in *2019 IEEE Conference on Decision and Control (CDC)*, 2019.
- [12] L. Huang, J. Coulson, J. Lygeros, and F. Dörfler, “Data-driven wide-area control,” *arXiv preprint arXiv:1911.12151*, 2019.
- [13] E. Elokda, J. Coulson, P. Beuchat, J. Lygeros, and F. Dörfler, “Data-enabled predictive control for quadcopters,” *ETH Zurich, Research Collection:10.3929/ethz-b-000415427*, 2019.
- [14] C. De Persis and P. Tesi, “Formulas for data-driven control: Stabilization, optimality and robustness,” *IEEE Transactions on Automatic Control*, 2019.
- [15] J. Berberich, A. Romer, C. W. Scherer, and F. Allgöwer, “Robust data-driven state-feedback design,” *arXiv preprint arXiv:1909.04314*, 2019.
- [16] J. Berberich, J. Köhler, M. A. Müller, and F. Allgöwer, “Robust constraint satisfaction in data-driven MPC,” *arXiv preprint arXiv:2003.06808*, 2020.
- [17] A. Koch, J. Berberich, J. Köhler, and F. Allgöwer, “Determining optimal input-output properties: A data-driven approach,” *arXiv preprint arXiv:2002.03882*, 2020.
- [18] H. J. van Waarde, C. De Persis, M. K. Camlibel, and P. Tesi, “Willems’ fundamental lemma for state-space systems and its extension to multiple datasets,” *IEEE Control Systems Letters*, vol. 4, no. 3, pp. 602–607, 2020.
- [19] H. J. Van Waarde, J. Eising, H. L. Trentelman, and M. K. Camlibel, “Data informativity: A new perspective on data-driven analysis and control,” *IEEE Transactions on Automatic Control*, 2020.
- [20] C. D. Persis and P. Tesi, “Low-complexity learning of linear quadratic regulators from noisy data,” *arXiv preprint arXiv:2005.01082*, 2020.
- [21] L. Hewing, K. P. Wabersich, M. Menner, and M. N. Zeilinger, “Learning-based model predictive control: Toward safe learning in control,” *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 3, 2019.

- [22] U. Rosolia and F. Borrelli, "Learning model predictive control for iterative tasks. A data-driven control framework," *IEEE Transactions on Automatic Control*, vol. 63, no. 7, pp. 1883–1896, 2018.
- [23] E. Kaiser, J. N. Kutz, and S. L. Brunton, "Sparse identification of nonlinear dynamics for model predictive control in the low-data limit," *Proceedings of the Royal Society A*, vol. 474, no. 2219, p. 20180335, 2018.
- [24] J. R. Salvador, D. Ramirez, T. Alamo, D. M. de la Peña, and G. García-Marín, "Data driven control: An offset free approach," in *2019 18th European Control Conference (ECC)*. IEEE, 2019, pp. 23–28.
- [25] R. Krug and D. Dimitrov, "Model predictive motion control based on generalized dynamical movement primitives," *Journal of Intelligent & Robotic Systems*, vol. 77, no. 1, pp. 17–35, 2015.
- [26] M. C. Campi and E. Weyer, "Finite sample properties of system identification methods," *IEEE Transactions on Automatic Control*, vol. 47, no. 8, pp. 1329–1334, 2002.
- [27] M. Vidyasagar and R. L. Karandikar, "A learning theory approach to system identification and stochastic adaptive control," in *Probabilistic and randomized methods for design under uncertainty*. Springer, 2006, pp. 265–302.
- [28] S. Tu, R. Boczar, A. Packard, and B. Recht, "Non-asymptotic analysis of robust control from coarse-grained identification," *arXiv preprint arXiv:1707.04791*, 2017.
- [29] T. Sarkar, A. Rakhlin, and M. A. Dahleh, "Finite-time system identification for partially observed LTI systems of unknown order," *arXiv preprint arXiv:1902.01848*, 2019.
- [30] R. Boczar, N. Matni, and B. Recht, "Finite-data performance guarantees for the output-feedback control of an unknown system," *arXiv preprint arXiv:1803.09186*, 2018.
- [31] S. Dean, H. Mania, N. Matni, B. Recht, and S. Tu, "On the sample complexity of the linear quadratic regulator," *Foundations of Computational Mathematics*, pp. 1–47, 2019.
- [32] H. Hjalmarsson, "From experiment design to closed-loop control," *Automatica*, vol. 41, no. 3, pp. 393–438, 2005.
- [33] M. Gevers, "Identification for control: From the early achievements to the revival of experiment design," *European Journal of Control*, vol. 11, pp. 1–18, 2005.
- [34] B. A. Ogunnaike, "A contemporary industrial perspective on process control theory and practice," *Annual Reviews in Control*, vol. 20, pp. 1–8, 1996.
- [35] J. Richalet, "Industrial applications of model based predictive control," *Automatica*, vol. 29, no. 5, pp. 1251–1274, 1993.
- [36] W. Favoreel, B. De Moor, and M. Gevers, "SPC: Subspace predictive control," *IFAC Proceedings Volumes*, vol. 32, no. 2, pp. 4004–4009, 1999.
- [37] B. Huang and R. Kadali, *Dynamic modeling, predictive control and performance monitoring: A data-driven subspace approach*. Springer, 2008.
- [38] F. Dörfler, J. Coulson, and I. Markovsky, "Bridging direct & indirect data-driven control formulations via regularizations and relaxations," *arXiv preprint arXiv:2101.01273*, 2021.
- [39] A. Mesbah, "Stochastic model predictive control: An overview and perspectives for future research," *IEEE Control Systems Magazine*, vol. 36, no. 6, pp. 30–44, 2016.
- [40] B. P. Van Parys, D. Kuhn, P. J. Goulart, and M. Morari, "Distributionally robust control of constrained stochastic systems," *IEEE Transactions on Automatic Control*, vol. 61, no. 2, pp. 430–442, 2015.
- [41] I. Markovsky, J. C. Willems, S. Van Huffel, and B. De Moor, *Exact and Approximate Modeling of Linear Systems: A Behavioral Approach*. SIAM, 2006.
- [42] J. C. Willems, "From time series to linear system – Part I. Finite dimensional linear time invariant systems," *Automatica*, vol. 22, no. 5, pp. 561–580, 1986.
- [43] A. Damen, P. Van den Hof, and A. Hajdasinski, "Approximate realization based upon an alternative to the Hankel matrix: the Page matrix," *Systems & Control Letters*, vol. 2, no. 4, pp. 202–208, 1982.
- [44] A. Agarwal, M. J. Amjad, D. Shah, and D. Shen, "Model agnostic time series analysis via matrix estimation," *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, vol. 2, no. 3, pp. 1–39, 2018.
- [45] F. Borrelli, A. Bemporad, and M. Morari, *Predictive Control for Linear and Hybrid Systems*. Cambridge University Press, 2017.
- [46] A. Nemirovski and A. Shapiro, "Convex approximations of chance constrained programs," *SIAM Journal on Optimization*, vol. 17, no. 4, pp. 969–996, 2006.
- [47] S. Shafieezadeh-Abadeh, D. Kuhn, and P. Mohajerin Esfahani, "Regularization via mass transportation," *Journal of Machine Learning Research*, vol. 20, no. 103, pp. 1–68, 2019.
- [48] N. Fournier and A. Guillin, "On the rate of convergence in Wasserstein distance of the empirical measure," *Probability Theory and Related Fields*, vol. 162, no. 3–4, pp. 707–738, 2015.
- [49] D. Kuhn, P. Mohajerin Esfahani, V. A. Nguyen, and S. Shafieezadeh-Abadeh, "Wasserstein distributionally robust optimization: Theory and applications in machine learning," in *Operations Research & Management Science in the Age of Analytics*. INFORMS, 2019, pp. 130–166.
- [50] J. Blanchet, Y. Kang, and K. Murthy, "Robust Wasserstein profile inference and applications to machine learning," *Journal of Applied Probability*, vol. 56, no. 3, pp. 830–857, 2019.
- [51] M. ApS, *The MOSEK optimization toolbox for MATLAB manual. Version 9.0.*, 2019. [Online]. Available: <http://docs.mosek.com/9.0/toolbox/index.html>
- [52] L. Huang, J. Coulson, J. Lygeros, and F. Dörfler, "Decentralized data-enabled predictive control for power system oscillation damping," *arXiv preprint arXiv:1911.12151*, 2019.
- [53] D. P. Bertsekas, *Convex optimization theory*. Athena Scientific Belmont, 2009.
- [54] A. R. Hota, A. Cherukuri, and J. Lygeros, "Data-driven chance constrained optimization under Wasserstein ambiguity sets," *arXiv preprint arXiv:1805.06729*, 2018.

Jeremy Coulson is a PhD student with the Automatic Control Laboratory at ETH Zürich. He received his Master of Applied Science in Mathematics & Engineering from Queen's University, Canada in August 2017. He received his B.Sc.Eng degree in Mechanical Engineering & Applied Mathematics from Queen's University in 2015. His research interests include data-driven control methods, stochastic optimization, and control of partial differential equations.

John Lygeros completed a B.Eng. degree in electrical engineering in 1990 and an M.Sc. degree in Systems Control in 1991, both at Imperial College of Science Technology and Medicine, London, U.K.. In 1996 he obtained a Ph.D. degree from the Electrical Engineering and Computer Sciences Department, University of California, Berkeley. During the period 1996–2000 he held a series of post-doctoral researcher appointments at the Laboratory for Computer Science, M.I.T., and the Electrical Engineering and Computer Sciences Department at U.C. Berkeley. Between 2000 and 2003 he was a University Lecturer at the Department of Engineering, University of Cambridge, U.K., and a Fellow of Churchill College. Between 2003 and 2006 he was an Assistant Professor at the Department of Electrical and Computer Engineering, University of Patras, Greece. In July 2006 he joined the Automatic Control Laboratory at ETH Zurich, where he is currently serving as the Head of the Automatic Control Laboratory. His research interests include modelling, analysis, and control of hierarchical, hybrid, and stochastic systems, with applications to biochemical networks, automated highway systems, air traffic management, power grids and camera networks. John Lygeros is a Fellow of the IEEE, and a member of the IET and the Technical Chamber of Greece; since 2013 he serves as the Treasurer of the International Federation of Automatic Control.

Florian Dörfler is an Associate Professor at the Automatic Control Laboratory at ETH Zürich. He received his Ph.D. degree in Mechanical Engineering from the University of California at Santa Barbara in 2013, and a Diplom degree in Engineering Cybernetics from the University of Stuttgart in 2008. From 2013 to 2014 he was an Assistant Professor at the University of California Los Angeles. His primary research interests are centered around control, optimization, and system theory with applications in network systems such as electric power grids, robotic

coordination, and social networks. He is a recipient of the distinguished young research awards by IFAC (Manfred Thoma Medal 2020) and EUCA (European Control Award 2020). His students were winners or finalists for Best Student Paper awards at the European Control Conference (2013, 2019), the American Control Conference (2016), and the PES PowerTech Conference (2017). He is furthermore a recipient of the 2010 ACC Student Best Paper Award, the 2011 O. Hugo Schuck Best Paper Award, the 2012-2014 Automatica Best Paper Award, the 2016 IEEE Circuits and Systems Guillemin-Cauer Best Paper Award, and the 2015 UCSB ME Best PhD award.