Distributionally Robust Infinite-horizon Control: from a pool of samples to the design of dependable controllers

Jean-Sébastien Brouillon*, Andrea Martin*, John Lygeros, Florian Dörfler, and Giancarlo Ferrari-Trecate

***Abstract*— We study control of constrained linear systems when faced with only partial statistical information about the disturbance processes affecting the system dynamics and the sensor measurements. Specifically, given a finite collection of disturbance realizations, we consider the problem of designing a stabilizing control policy with provable safety and performance guarantees in face of the inevitable mismatch between the true and the empirical distributions. We capture this discrepancy using Wasserstein ambiguity sets, and we formulate a distributionally robust (DR) optimal control problem, which provides guarantees on the expected cost, safety, and stability of the system. To solve this problem, we first present new results for DR optimization of quadratic objectives using convex programming, showing that strong duality holds under mild conditions. Then, by combining our results with the system level parametrization (SLP) of linear feedback policies, we show that the design problem can be reduced to a semidefinite optimization problem (SDP).**

I. Introduction

As modern engineered systems become increasingly complex and interconnected, classical control methods based on stochastic optimization face the challenge of overcoming the lack of a precise statistical description of the uncertainty. In fact, the probability distribution of the uncertainty is generally unknown and only indirectly observable through a finite number of independent samples. In addition, replacing the true distribution with a nominal estimate in the spirit of certainty equivalence often proves unsatisfactory; the optimization process amplifies any statistical error in the distribution inferred from data, resulting in solutions that are prone to yielding poor out-of-sample performance [1]–[3].

Motivated by these observations, the paradigm of distributionally robust optimization (DRO) considers a minimax stochastic optimization problem over a neighborhood of the nominal distribution defined in terms of a distance in the probability space. In this way, the solution becomes robust to the most averse distribution that is sufficiently close to the nominal distribution, while the degree of conservatism of the underlying optimization can be regulated by adjusting the radius of the ambiguity set.

While several alternatives have been proposed to measure the discrepancy between probability distributions, including the Kullback–Leibler divergence and the total variation distance [4], recent literature has shown that working with ambiguity sets defined using the Wasserstein metric [5] offers a number of advantages in terms of expressivity, computational tractability, and statistical out-of-sample guarantees [1]–[3]. Thanks to these properties, Wasserstein DRO has found application in a wide variety of domains, ranging from finance and machine learning to game theory, see, e.g., [2], [3], [6]–[8].

Similarly, Wasserstein ambiguity sets have recently been interfaced with the dynamic environments and continuous actions spaces typical of control. In [9], the authors consider a generalization of classical linear quadratic Gaussian (LQG) control, where the noise distributions belong to Wasserstein balls centered at nominal Gaussian distributions. Motivated by the idea of leveraging uncertainty samples for data-driven decision-making under general distributions, a parallel line of research instead considers ambiguity sets centered at nominal empirical distributions. Among other contributions exploiting the greater expressivity provided by this data-driven approach, [10]–[12] consider the design of tube-based predictive control schemes, [13] and [14] address infinite-horizon problems using dynamic programming, [15] and [16] focus on filtering and state estimation problems. More fundamentally, [17] and [18] provide exact characterizations of how Wasserstein ambiguity sets propagate through the system dynamics, shedding light on the role of feedback in controlling shape and size of the ambiguity sets resulting from distributional uncertainty.

Despite these advances, it remains unclear how the availability of samples can drive the design of a control policy that guarantees safety and performance in face of distributional uncertainty while simultaneously ensuring stability of the closed-loop system. Motivated by this challenge, we first establish novel strong duality results for DRO of quadratic functions, which are routinely encountered in control, by explicitly accounting for the (possibly bounded) support of the uncertainty. Then, leveraging the system level parametrization (SLP) of linear dynamic controllers [19], we present a convex reformulation of the Distributionally Robust Infinite-horizon Controller (DRInC) synthesis problem, which exploits a finite impulse response (FIR) approximation of the system closed-loop maps. As key advantages, our optimization-based approach guarantees stability of the closed-loop interconnection by design, and only requires one-shot offline computations. As such, our solution bypasses the computational bottleneck that would result by recomputing the optimal control policy online

Jean-Sébastien Brouillon, Andrea Martin, and Giancarlo Ferrari-Trecate are with the Institute of Mechanical Engineering, EPFL, Switzerland. E-mail addresses: {jean-sebastien.brouillon, andrea.martin, giancarlo.ferraritrecate}@epfl.ch.

John Lygeros and Florian Dörfler are with the Department of Information Technology and Electrical Engineering, ETH Zürich, Switzerland. E-mail addresses: {jlygeros, dorfler}@ethz.ch.

Research supported by the Swiss National Science Foundation (SNSF) under the NCCR Automation (grant agreement 51NF40_80545).

*Jean-Sébastien Brouillon and Andrea Martin contributed equally.

according to a receding horizon strategy [10]–[12]. In fact, as the complexity of the synthesis problem increases with the number of considered uncertainty samples, solving the policy optimization problem in real-time becomes prohibitive whenever the nominal empirical distribution is estimated using a sufficiently large number of uncertainty samples. Further, differently from [13] and [14], which consider infinite-horizon DRC in unconstrained scenarios, our approach naturally extends to include satisfaction of probabilistic safety constraints expressed as distributionally robust conditional value-at-risk (CVaR) constraints. Lastly, the proposed optimization perspective allows us to seamlessly study the partially observed setting, extending the recent results [9], [20], [21] on output-feedback DRC to the infinite horizon case. As we comment throughout the paper, our formulation encompasses several control problems considered in the literature, providing a unified perspective on stochastic and robust control objectives.

II. Problem Statement

A. System dynamics and uncertainty description

We consider controllable and observable linear dynamical systems described by the state-space equations:

$$x_{t+1} = Ax_t + Bu_t + w_t\,,\ y_t = Cx_t + v_t\,, \tag{1}$$

where $x_t \in \mathbb{R}^n$, $u_t \in \mathbb{R}^m$, $y_t \in \mathbb{R}^p$, $w_t \in \mathbb{R}^n$ and $v_t \in \mathbb{R}^p$ are the system state, the control input, the observable output, and the stochastic disturbances modeling process and measurement noise, respectively. We study infinite-horizon control when only partial statistical information about the distribution of the joint disturbance process $\xi_t = (w_t, v_t)$ is available. Specifically, we assume availability of $N \in \mathbb{N}$ independent observations $\boldsymbol{\xi}_T^{(1)}, \dots, \boldsymbol{\xi}_T^{(N)}$, where each sample

$$\boldsymbol{\xi}_T^{(i)} = (\boldsymbol{w}_T^{(i)}, \boldsymbol{v}_T^{(i)}) = (w_0^{(i)}, \dots, w_T^{(i)}, v_0^{(i)}, \dots, v_T^{(i)})\,, \tag{2}$$

constitutes a trajectory of length $T \in \mathbb{N}$ of w_t and v_t. As no performance or safety guarantee can be established if the samples in (2) are not representative of the asymptotic statistics of $\boldsymbol{w}$ and $\boldsymbol{v}$, we start by formulating the following stationarity assumption, see, e.g., [22, p. 154].

Assumption 1: For all $t \in \mathbb{N}$, the stochastic process that generates the joint disturbance vector $\xi_t = (w_t, v_t)$ is stationary of order T, i.e., $\mathbb{P}(\xi_0, \dots, \xi_T) = \mathbb{P}(\xi_t, \dots, \xi_{t+T})$. We note that Assumption 1 subsumes the usual setting where each realization of the disturbance processes is independent and identically distributed, and more generally allows modeling temporal correlation between samples that are separated by up to T time steps. Further, as the order T can theoretically be arbitrarily large, this assumption is relatively mild, albeit, in practice, an upper bound on the order T is often dictated by computational complexity concerns.

Throughout the paper, we denote by $\Xi \subseteq \mathbb{R}^d$, with $d = (n+p)(T+1)$, the support of the unknown probability distribution $\mathbb{P}$, and we make the following assumption.

Assumption 2: The support set $\Xi = \{\boldsymbol{\xi} \in \mathbb{R}^d : H\boldsymbol{\xi} \leq h\}$ is full-dimensional, that is, Ξ contains a d-dimensional ball with strictly positive radius.

We mainly focus on the case where Ξ is a compact polyhedron. Nevertheless, as we will highlight in the following, our results naturally extend to the most studied case $\Xi = \mathbb{R}^d$.

Remark 1: Reconstructing $\boldsymbol{w}_T^{(i)}$ and $\boldsymbol{v}_T^{(i)}$ online, that is, given the corresponding input and output signals $(\boldsymbol{u}_T^{(i)}, \boldsymbol{y}_T^{(i)})$ only, is in general not possible. Still, the samples in (2) can be reconstructed from a series of offline experiments conducted in a laboratory environment, where the availability of additional sensors allows measuring the entire state trajectory $\boldsymbol{x}_T^{(i)}$ of the system. Alternatively, if $\boldsymbol{w}$ and $\boldsymbol{v}$ represent the effect of complex physical phenomena, e.g., wind gusts and turbulences, and sensor inaccuracy, respectively, the samples in (2) can also be generated using high-fidelity simulators.

B. Control objectives, policies, and uncertainty propagation

We consider the problem of designing offline a stabilizing feedback policy that retains probabilistic safety and performance guarantees over an infinite horizon. Specifically, given $D \succeq 0$, we measure the control cost that a policy $\boldsymbol{u} = \boldsymbol{\pi}(\boldsymbol{y})$ incurs whenever the joint disturbance sequence $\boldsymbol{\xi}$ realizes as:

$$J(\boldsymbol{\pi}, \boldsymbol{\xi}) = \lim_{T' \to \infty} \frac{1}{T'} \sum_{t=0}^{T'} \begin{bmatrix} x_t^\top & u_t^\top \end{bmatrix} D \begin{bmatrix} x_t \\ u_t \end{bmatrix}\,,$$

and we define polytopic safe sets $\mathcal{X} \subseteq \mathbb{R}^n$ and $\mathcal{U} \subseteq \mathbb{R}^m$ for the system state and input signals, respectively, as:

$$\mathcal{X} = \{x \in \mathbb{R}^n : g_x(x) = \max_{j \in [J_x]}\ G_{xj}^\top x + g_{xj} \leq 0\,,\ J_x \in \mathbb{N}\}\,,$$
$$\mathcal{U} = \{u \in \mathbb{R}^m : g_u(u) = \max_{j \in [J_u]}\ G_{uj}^\top u + g_{uj} \leq 0\,,\ J_u \in \mathbb{N}\}\,,$$

where $[J_x]$ denotes the set $\{1, \dots, J_x\} \subset \mathbb{N}$ and similarly for $[J_u]$. Then, given a safety parameter $\gamma \in (0, 1)$ to control the level of acceptable constraint violations, we formulate the following chance constrained stochastic optimization problem:

$$\boldsymbol{\pi}^\star = \arg\min_{\boldsymbol{\pi}}\ \mathbb{E}_{\mathbb{P}}\left[J(\boldsymbol{\pi}, \boldsymbol{\xi})\right] \tag{3a}$$
$$\text{subject to}\ \ \mathrm{CVaR}_\gamma^{\mathbb{P}}(\max\{g_x(x_t(\boldsymbol{\xi})), g_u(u_t(\boldsymbol{\xi}))\}) \leq 0\,, \tag{3b}$$

where CVaR constraints are defined according to

$$\mathrm{CVaR}_\gamma^{\mathbb{P}}(g(\boldsymbol{\xi})) = \inf_{\tau \in \mathbb{R}}\ \tau + \frac{1}{\gamma}\mathbb{E}_{\mathbb{P}}[\max\{g(\boldsymbol{\xi}) - \tau, 0\}]\,, \tag{4}$$

for any measurable function $g : \mathbb{R}^d \to \mathbb{R}$. We note that, besides implying that $\mathbb{P}[x_t \in \mathcal{X}\,, u_t \in \mathcal{U}] \geq 1 - \gamma$, (3b) also accounts for the expected amount of constraint violation in the γ percent of cases where any such violation occurs. As such, the CVaR formulation reflects the observation that, in most control applications, severe breaches of the safety constraints often have far more detrimental consequences than mild violations. As the probability distribution $\mathbb{P}$ is fundamentally unknown, however, we cannot address the decision problem (3) directly, and we instead rely on the following approximations.

First, we construct the empirical probability distribution

$$\widehat{\mathbb{P}} = \frac{1}{N} \sum_{i=1}^{N} \delta_{\boldsymbol{\xi}_T^{(i)}}\,, \tag{5}$$

where $\delta_{\boldsymbol{\xi}_T^{(i)}}$ denotes the Dirac delta distribution at $\boldsymbol{\xi}_T^{(i)}$. In order to immunize against any error in $\widehat{\mathbb{P}}$, we replace the nominal objective (3a) with the minimization of the worst-case expected loss over the set of distributions $\mathbb{B}_\epsilon(\widehat{\mathbb{P}}) \subseteq \mathcal{P}(\boldsymbol{\Xi})$ that are supported on $\boldsymbol{\Xi}$ and are sufficiently close to the empirical estimate $\widehat{\mathbb{P}}$.[1] More formally, we define

$$\mathbb{B}_\epsilon(\widehat{\mathbb{P}}) = \{\mathbb{Q} \in \mathcal{P}(\boldsymbol{\Xi}) : W(\widehat{\mathbb{P}}, \mathbb{Q}) \leq \epsilon\}\,, \tag{6}$$

where $\epsilon \geq 0$ is the radius of the ambiguity set $\mathbb{B}_\epsilon(\widehat{\mathbb{P}})$, and $W(\widehat{\mathbb{P}}, \mathbb{Q})$ is the Wasserstein distance between $\widehat{\mathbb{P}}$ and $\mathbb{Q}$, i.e.,

$$W(\widehat{\mathbb{P}}, \mathbb{Q}) = \inf_{\pi \in \Pi} \int_{\boldsymbol{\Xi}^2} \|\boldsymbol{\xi} - \boldsymbol{\xi}'\|_2^2 \, \pi(\mathrm{d}\boldsymbol{\xi}, \mathrm{d}\boldsymbol{\xi}')\,, \tag{7}$$

where Π denotes the set of joint probability distributions of $\boldsymbol{\xi}$ and $\boldsymbol{\xi}'$ with marginal distributions $\widehat{\mathbb{P}}$ and $\mathbb{Q}$, respectively [1], [2]. In (7), the decision variable π encodes a transportation plan for moving a mass distribution described by $\widehat{\mathbb{P}}$ to a distribution described by $\mathbb{Q}$. Thus, $\mathbb{B}_\epsilon(\widehat{\mathbb{P}})$ can be interpreted as the set of distributions onto which $\widehat{\mathbb{P}}$ can be reshaped at a cost of at most ϵ, where the cost of moving a unit probability from $\boldsymbol{\xi}$ to $\boldsymbol{\xi}'$ is given by $\|\boldsymbol{\xi} - \boldsymbol{\xi}'\|_2^2$.

Second, since dynamic programming solutions are generally computationally intractable, we restrict our attention to policies $\boldsymbol{\pi} \in \boldsymbol{\Pi}_{\mathrm{L}}$ that are linear in the past observations $\boldsymbol{y}$, that is, $\boldsymbol{u} = \boldsymbol{\pi}(\boldsymbol{y}) = \boldsymbol{K}(z)\boldsymbol{y}$ for some real-rational proper transfer function $\boldsymbol{K}(z)$. Besides computational advantages, our choice is supported by recent advances in DRC, which show that linear policies are globally optimal for a generalization of the classical unconstrained LQG problem, where the noise distributions belong to a Wasserstein ambiguity set (6), centered at a nominal Gaussian distribution $\widehat{\mathbb{P}}$ [9].

We are now in a position to state our problem of interest as:

$$\inf_{\boldsymbol{\pi} \in \boldsymbol{\Pi}_{\mathrm{L}}} \quad \sup_{\mathbb{Q} \in \mathbb{B}_\epsilon(\widehat{\mathbb{P}})} \quad \mathbb{E}_{\mathbb{Q}}\left[J(\boldsymbol{\pi}, \boldsymbol{\xi})\right] \tag{8a}$$

$$\text{subject to} \quad \sup_{\mathbb{Q} \in \mathbb{B}_\epsilon(\widehat{\mathbb{P}})} \quad \mathrm{CVaR}_\gamma^{\mathbb{Q}}(g_t(\boldsymbol{\xi})) \leq 0\,,\ \forall t \in \mathbb{N}\,, \tag{8b}$$

where $g_t(\boldsymbol{\xi}) = \max\{g_x(x_t(\boldsymbol{\xi})), g_u(u_t(\boldsymbol{\xi}))\}$ for compactness. Note that the worst-case distributions in (8a) and (8b) may not coincide. Despite the fact that in practice the uncertainty distribution is unique, the formulation in (8) proves necessary to ensure safety for all distributions in $\mathbb{B}_\epsilon(\widehat{\mathbb{P}})$ and not simply for the one maximizing the expected control cost.

C. Expressivity of the problem formulation and related work

The solution to the DRO problem (8) depends on the radius ϵ defining (6). In particular, we argue that (8) generalizes classical $\mathcal{H}_2$ and $\mathcal{H}_\infty$ control problems, which correspond to the limit cases of ϵ approaching 0 and ∞, respectively.

If $\epsilon = 0$, the Wasserstein ball $\mathbb{B}_\epsilon(\widehat{\mathbb{P}})$ reduces to the singleton $\{\widehat{\mathbb{P}}\}$ and the supremum disappears. This gives a simple Monte-Carlo-based control design problem [24], [25]. Moreover, because $J(\boldsymbol{\pi}, \boldsymbol{\xi})$ is quadratic, the resulting optimal controller is the LQG designed for $\mathbb{P}_{\mathcal{N}} = \mathcal{N}(\mathbb{E}_{\xi\sim\widehat{\mathbb{P}}}[\xi], \mathrm{var}_{\xi\sim\widehat{\mathbb{P}}}[\xi])$ in the absence of constraints [26]. Indeed, because both the dynamics and the controller are linear, one has[2]

$$\mathbb{E}_{\mathbb{P}_{\mathcal{N}}}\left[J(\boldsymbol{\pi}, \boldsymbol{\xi})\right] = \mathbb{E}_{\widehat{\mathbb{P}}}\left[J(\boldsymbol{\pi}, \boldsymbol{\xi})\right],$$

which means that the $\arg\min_\pi$ of both expectations is also the same.

If ϵ is very large and $\boldsymbol{\Xi}$ is compact, (8) can also be seen as a generalization of $\mathcal{H}_\infty$ synthesis methods [26], [27]. In fact, in the limit case of $\epsilon \to \infty$ and no matter how $\widehat{\mathbb{P}}$ is constructed, (6) contains all distributions $\mathcal{P}(\boldsymbol{\Xi})$ supported on $\boldsymbol{\Xi}$, including the degenerate distribution taking value at the most-averse $\boldsymbol{\xi}$ almost surely.

Intermediate values of ϵ instead yield solutions that leverage the observations (2) to trade-off robustness to adversarial perturbations or distribution shifts against performance under distributions in a neighborhood of $\widehat{\mathbb{P}}$.

We conclude this section by remarking that, differently from [9], we do not assume that the nominal distribution $\widehat{\mathbb{P}}$ is Gaussian, and instead use the empirical estimate (5) to provide greater design flexibility. In fact, if $\mathbb{P}$ is, e.g., bimodal, then the Wasserstein distance between $\mathbb{P}$ and its closest Gaussian distribution $\mathbb{G}$ will generally be larger than the Wasserstein distance between $\mathbb{P}$ and its empirical estimate $\widehat{\mathbb{P}}$. In turn, this implies that a larger radius ϵ needs to be used to ensure that $\mathbb{P} \in \mathbb{B}_\epsilon(\mathbb{G})$ with high probability, leading to a more conservative design.

III. Background

In this section, we recall useful technical preliminaries, and we discuss the design assumptions that will allow us to compute an approximate solution to (8) through convex programming. In particular, we start by reviewing the system level approach to controller synthesis [19], and then present recent duality results from the DRO literature [3].

A. System level synthesis

The system level synthesis framework provides a convex parameterization of the non-convex set of internally stabilizing controllers $\boldsymbol{K}(z)$, allowing one to reformulate many control problems as optimization over the closed-loop responses $\boldsymbol{\Phi}_{xw}(z), \boldsymbol{\Phi}_{xv}(z), \boldsymbol{\Phi}_{uw}(z)$ and $\boldsymbol{\Phi}_{uv}(z)$ that map $\boldsymbol{w}$ and $\boldsymbol{v}$ to $\boldsymbol{x}$ and $\boldsymbol{u}$. To define these maps, we first combine the linear output feedback policy $\boldsymbol{u} = \boldsymbol{K}(z)\boldsymbol{y}$ with the z transform of the state dynamics in (1) to obtain:

$$(zI - (A + B\boldsymbol{K}(z)C))\boldsymbol{x} = \boldsymbol{w} + B\boldsymbol{K}(z)\boldsymbol{v}\,.$$

[1] It is well-known that solving (3) upon naively replacing $\mathbb{P}$ with $\widehat{\mathbb{P}}$, that is, setting ϵ to zero in (6), may lead to decisions that are unsafe or exhibit poor out-of-sample performance, as the optimization process often amplifies any estimation error in $\widehat{\mathbb{P}}$. Instead, for any $\beta > 0$, if $\mathbb{P}$ is light-tailed and the radius ϵ is chosen as a sublinearly growing function of $\frac{\log(1/\beta)}{N}$, then results from measure concentration theory ensure that $\mathbb{P}$ lie inside the ambiguity set (6) with confidence $1 - \beta$, see, [23, Theorem 2] and [2, Theorem 18]. Therefore, in this case, any solution to (8) retains finite-samples probabilistic guarantees in terms of out-of-samples control cost and constraint satisfaction.

[2] Both expectations are equal to the same linear transformation of the first and second moments of $\widehat{\mathbb{P}}$ and $\mathbb{P}_{\mathcal{N}}$, which are equal.

Then, since the transfer matrix $(zI - (A + B\boldsymbol{K}(z)C))$ is invertible for any proper controller $\boldsymbol{K}(z)$, we have

$$\begin{aligned}\begin{bmatrix}\boldsymbol{x}\\ \boldsymbol{u}\end{bmatrix} &= \begin{bmatrix}\boldsymbol{\Phi}_{xw}(z) & \boldsymbol{\Phi}_{xv}(z)\\ \boldsymbol{\Phi}_{uw}(z) & \boldsymbol{\Phi}_{uv}(z)\end{bmatrix}\begin{bmatrix}\boldsymbol{w}\\ \boldsymbol{v}\end{bmatrix} = \boldsymbol{\Phi}_{\xi}(z)\boldsymbol{\xi}\,,\\ &= \begin{bmatrix}(zI-(A+B\boldsymbol{K}(z)C))^{-1} & \boldsymbol{\Phi}_{xw}(z)B\boldsymbol{K}(z)\\ \boldsymbol{K}(z)C\boldsymbol{\Phi}_{xw}(z) & \boldsymbol{\Phi}_{uw}(z)+z\boldsymbol{K}(z)\end{bmatrix}\boldsymbol{\xi}\,.\end{aligned}$$

In particular, we note that causality of $\boldsymbol{K}(z)$ implies causality of $\boldsymbol{\Phi}_{uv}$ and strict causality of $\boldsymbol{\Phi}_{xw}, \boldsymbol{\Phi}_{xv}$ and $\boldsymbol{\Phi}_{uw}$. Further, one can show that the affine subspace defined by

$$\begin{bmatrix}zI - A & -B\end{bmatrix}\boldsymbol{\Phi}_{\xi}(z) = \begin{bmatrix}I & 0\end{bmatrix}\,, \tag{9a}$$

$$\boldsymbol{\Phi}_{\xi}(z)\begin{bmatrix}zI-A\\ -C\end{bmatrix} = \begin{bmatrix}I\\ 0\end{bmatrix}\,, \tag{9b}$$

characterizes all and only the system responses $\boldsymbol{\Phi}_{\xi}(z)$ that are achievable by an internally stabilizing controller $\boldsymbol{K}(z)$ [19]. Despite the fact that (9) defines a convex feasible set, minimizing a given convex objective with respect to the closed-loop transfer matrix $\boldsymbol{\Phi}_{\xi}(z) = \sum_{k=0}^{\infty}\Phi(k)z^{-k}$ proves challenging, as the resulting optimization problem remains infinite dimensional. Therefore, to recover tractability and following [19], [28], we rely on a FIR approximation of $\boldsymbol{\Phi}_{\xi}(z)$, i.e., we restrict our attention to the truncated system response $\boldsymbol{\Phi}_{\xi}^{T}(z) = \sum_{k=0}^{T}\Phi(k)z^{-k}$. We remark that controllability and observability of (1) ensure that (9) admits a FIR solution [19, Theorem 4]. At the same time, since $\boldsymbol{\Phi}_{\xi}(z)$ represents a stable map, the effect of this FIR approximation becomes negligible if T is sufficiently large; for the case of LQR regulators, for instance, it was shown that the performance degradation relative to the solution to the infinite-horizon problem decays exponentially with T, see [29, Section 5].

According to the discussed FIR approximation, we let:

$$\begin{aligned}\boldsymbol{\Phi}_x &= [\Phi_{xw}(T),\ldots,\Phi_{xw}(0),\Phi_{xv}(T),\ldots,\Phi_{xv}(0)]\,,\\ \boldsymbol{\Phi}_u &= [\Phi_{uw}(T),\ldots,\Phi_{uw}(0),\Phi_{uv}(T),\ldots,\Phi_{uv}(0)]\,,\end{aligned}$$

and we define $\boldsymbol{\Phi} = [\boldsymbol{\Phi}_x^\top, \boldsymbol{\Phi}_u^\top]^\top$ for compactness. With this notation in place, for any $t \geq T$, we have that:

$$x_t = \boldsymbol{\Phi}_x\boldsymbol{\xi}_{t-T:t}\,,\ u_t = \boldsymbol{\Phi}_u\boldsymbol{\xi}_{t-T:t}\,, \tag{10}$$

where $\boldsymbol{\xi}_{t-T:t} = [w_{t-T},\ldots,w_t,v_{t-T},\ldots,v_t]^\top$ collects the last $T+1$ realizations of the process and measurement noises.

The following proposition, for which we provide a proof in Appendix A for the sake of comprehensiveness, shows how to implement a controller that achieves a given pair of system responses $\boldsymbol{\Phi}_x$ and $\boldsymbol{\Phi}_u$.

Proposition 1: If the closed loop map $\boldsymbol{\Phi}$ is achievable, the corresponding control policy $\boldsymbol{\pi}(\boldsymbol{\Phi})$ can be implemented as a linear system with dynamics

$$\delta_t = -\boldsymbol{\Phi}_x\boldsymbol{\phi}_{t-T:t},\ u_t = \boldsymbol{\Phi}_u\boldsymbol{\phi}_{t-T:t} + \Phi_{uv}(0)C\delta_t, \tag{11}$$

where $\boldsymbol{\phi}_{t-T:t} = [\delta_{t-T+1}^\top,\ldots,\delta_{t-1}^\top,0_{2n},y_{t-T}^\top,\ldots,y_t^\top]^\top$.

B. A stationarity control problem

As we consider an infinite horizon control problem, we focus on the steady state behavior of the system, and we are instead less interested in the transient behavior [30]. Motivated by this and to take full advantage of the stationarity properties of ξ_t in Assumption 1, we focus on designing an optimal safe controller to operate the system for $t \geq T$ only. In this setting, we proceed to show that the distributionally robust worst-case control cost and CVaR constraints admit finite-dimensional representations

Assumption 3: The system is initialized by an external controller with $x_0,\ldots,x_{T-1} \in \mathcal{X}$ and $u_0,\ldots,u_{T-1} \in \mathcal{U}$.

We therefore redefine the optimization cost J in (8a) as

$$J_T(\boldsymbol{\pi}(\boldsymbol{\Phi}),\boldsymbol{\xi}) = \lim_{T'\to\infty}\frac{1}{T'-T}\sum_{t=T}^{T'}\boldsymbol{\xi}_{t-T:t}^\top\boldsymbol{\Phi}^\top D\boldsymbol{\Phi}\boldsymbol{\xi}_{t-T:t}\,.$$

Note that due to the stationarity of $\mathbb{Q}$ (see Assumption 1), J_T satisfies

$$\begin{aligned}&\mathbb{E}_{\mathbb{Q}}J(\boldsymbol{\pi}(\boldsymbol{\Phi}),\boldsymbol{\xi})\\ &\quad= \lim_{T'\to\infty}\mathbb{E}_{\substack{\xi_{0:R}\sim\mathbb{Q}\\ \vdots\\ \xi_{T'-T:T'}\sim\mathbb{Q}}}\frac{1}{T'-T}\sum_{t=T}^{T'}\boldsymbol{\xi}_{t-T:t}^\top\boldsymbol{\Phi}^\top D\boldsymbol{\Phi}\boldsymbol{\xi}_{t-T:t},\\ &\quad= \mathbb{E}_{\xi_T\sim\mathbb{Q}}\boldsymbol{\xi}_T^\top\boldsymbol{\Phi}^\top D\boldsymbol{\Phi}\boldsymbol{\xi}_T.\end{aligned} \tag{12}$$

The problem statement (8) for DRInC synthesis can be reformulated as finding the optimal FIR map $\boldsymbol{\Phi}^\star$ of length $T+1$ given by

$$\boldsymbol{\Phi}^\star = \underset{\boldsymbol{\Phi}\text{ achievable}}{\arg\min}\ \sup_{\mathbb{Q}\in\mathbb{B}_\epsilon(\widehat{\mathbb{P}})}\mathbb{E}_{\boldsymbol{\xi}_T\sim\mathbb{Q}}\boldsymbol{\xi}_T^\top\boldsymbol{\Phi}^\top D\boldsymbol{\Phi}\boldsymbol{\xi}_T\,, \tag{13}$$

while satisfying the achievability constraints (9) as well as conditional value-at-risk constraints

$$\sup_{\mathbb{Q}\in\mathbb{B}_\epsilon(\widehat{\mathbb{P}})}\mathrm{CVaR}_{1-\gamma}^{\boldsymbol{\xi}_T\sim\mathbb{Q}}(G_j^\top\boldsymbol{\Phi}\boldsymbol{\xi}_T + g_j) \leq 0, \forall j\in[J], \tag{14}$$

where $J = J_x + J_u$ and $[J] = \{1,\ldots,J\}$ enumerates all the constraints on $[x^\top, u^\top]$, which are defined by

$$G = \begin{bmatrix}G_x & 0\\ 0 & G_u\end{bmatrix},\ g = \begin{bmatrix}g_x\\ g_u\end{bmatrix}.$$

We highlight that while (8a) is an infimum problem, the minimum in (13) is attained. Indeed, as $\boldsymbol{\Xi}$ is full-dimensional per Assumption 2, there is always a distribution $\widehat{\mathbb{Q}}$ such that $\mathbb{E}_{\widehat{\mathbb{Q}}}J(\boldsymbol{\pi}(\boldsymbol{\Phi}),\boldsymbol{\xi})$ is strongly convex in $\boldsymbol{\Phi}$ (e.g., an empirical distribution containing samples that form a basis for $\mathbb{R}^d$). Moreover, since $\mathbb{E}_{\widehat{\mathbb{Q}}}J(\boldsymbol{\pi}(\boldsymbol{\Phi}),\boldsymbol{\xi}) \leq \sup_{\mathbb{Q}\in\mathbb{B}_\epsilon(\widehat{\mathbb{P}})}\mathbb{E}_{\mathbb{Q}}J(\boldsymbol{\pi}(\boldsymbol{\Phi}),\boldsymbol{\xi})$ by definition, the supremum in (13) is strongly convex and the minimizer $\boldsymbol{\Phi}^\star$ is attainable. However, both the control cost grow quadratically, which can render $\sup_{\mathbb{Q}\in\mathbb{B}_\epsilon(\widehat{\mathbb{P}})}\ \mathbb{E}_{\mathbb{Q}}\left[J(\boldsymbol{\pi}(\boldsymbol{\Phi}),\boldsymbol{\xi})\right]$ unattainable [3][3]. In

[3]The ratio between the growth rates of the loss function and the transport cost is crucial in DRO problems. If the control cost grows faster than the transport cost, the adversary can make the control cost diverge by moving an infinitesimal amount of mass very far away from the empirical distribution. Reversely, if the control cost grows slower, there is always be a point at which it is not worth for the adversary to keep moving and the supremum is attained. This is the case for the constraints, as their cost grows linearly.

what follows, we use the recent advances in DRO theory presented in [3] to reformulate the control design problem as a finite-dimensional and tractable problem.

C. Strong duality for DRO of piecewise linear objectives

The minimization (13) subject to (14) is infinite-dimensional and therefore cannot be directly solved. The next proposition, which serves as a starting point for our derivations in Section IV, shows how DRO of piecewise linear objectives can be recast as a finite-dimensional convex program.

Proposition 2: Let $a_j \in \mathbb{R}^d$ and $b_j \in \mathbb{R}$ constitute a piece-wise linear cost with J pieces. If Assumption 2 holds and $\epsilon > 0$, then the risk:

$$\sup_{\mathbb{Q}\in\mathbb{B}_\epsilon} \mathbb{E}_{\xi_T\sim\mathbb{Q}} \max_{j\in[J]} a_j^\top \boldsymbol{\xi}_T + b_j \,, \tag{15}$$

can be equivalently computed as:

$$\inf_{\lambda\geq 0,\kappa_{ij}\geq 0} \lambda\epsilon + \frac{1}{N}\sum_{i\in[N]} s^{(i)} \,, \text{ subject to} \tag{16a}$$

$$s^{(i)} \geq b_j + \frac{\|a_j\|_2^2}{4\lambda} - a_j^\top \boldsymbol{\xi}_T^{(i)} \tag{16b}$$

$$+\frac{1}{4\lambda}\kappa_{ij}^\top HH^\top\kappa_{ij} - \frac{1}{2\lambda}a_j^\top H^\top\kappa_{ij} + \left(H\boldsymbol{\xi}_T^{(i)} + h\right)^\top \kappa_{ij},$$

for all $i = 1, \ldots, N$ and $j = 1, \ldots, J$.

Proof: This proposition is a direct consequence of [3, Proposition 2.12]. For the sake of clarity, we report detailed derivations in Appendix C. ■

Proposition 2 uses strong duality to establish an equivalence between (16) and (15). In particular, the decision variables λ and κ_{ij} in (16) correspond to the Lagrange multipliers associated with the constraints $\mathbb{Q} \in \mathbb{B}_\epsilon$ and $\boldsymbol{\xi}_T \in \Xi$, respectively. The optimal value of λ can thus be interpreted as the shadow cost of robustification, i.e., the amount by which the risk $\mathbb{E}_{\xi_T\sim\mathbb{Q}} \max_{j\in[J]} a_j^\top \boldsymbol{\xi}_T + b_j$ increases for each unit increase of ϵ. The variables $s^{(i)}$ instead represent the empirical Lagrangian for each sample.

IV. Main Results

In this section, we present our main results. Motivated by the observation that the operational costs of engineering applications usually relate to energy consumption and are thus often modeled using quadratic functions, we first extend the results of Proposition 2 beyond piecewise linear objectives.

A. Non-convexity challenges

While [3, Proposition 2.12] holds for general transport costs and no matter if Ξ is bounded or not, this strong duality result does not directly apply to (13), as the objective $J(\boldsymbol{\pi}(\boldsymbol{\Phi}), \boldsymbol{\xi})$ is not piece-wise concave. An extension of current state-of-the-art results in DRO is therefore required to minimize a risk of the form

$$\mathcal{R}(Q) := \sup_{\mathbb{Q}\in\mathbb{B}_\epsilon} \mathbb{E}_{\xi_T\sim\mathbb{Q}} \; \boldsymbol{\xi}_T^\top Q \boldsymbol{\xi}_T \,. \tag{17}$$

where Ξ does not necessarily equal $\mathbb{R}^d$ and $Q \succeq 0$.

We start by observing that if the loss is not concave with respect to $\boldsymbol{\xi}_T$, then the optimization problem in (17) may not be convex. In fact, while [2] shows that there is a hidden convexity when $\Xi = \mathbb{R}^d$, this result does not hold in general. To illustrate this point, consider for example the situation drawn in Fig. 1. One can observe that if the constraint $\mathbb{Q} \in \mathbb{B}_\epsilon(\delta)$ is active, then the problem (17) amounts to a Quadratically Constrained Quadratic Program (QCQP), which admits a tight convex relaxation as a Semi-Definite Program (SDP) [31]. Conversely, however, when the constraint $\mathbb{Q} \in \mathbb{B}_\epsilon(\delta)$ is not active, the adversary must maximize a convex Quadratic Program (QP), which is not convex.

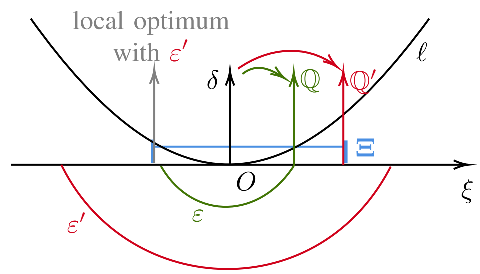

Fig. 1. Illustration of two worst-case distributions $\mathbb{Q} \in \mathbb{B}_\epsilon(\delta)$ and $\mathbb{Q}' \in \mathbb{B}_{\epsilon'}(\delta)$ in different Wasserstein balls around the Dirac delta distribution. The support ξ is represented by the horizontal blue line above the ξ axis, and the left-most Dirac distribution represents a local minima in $\mathbb{B}_{\epsilon'}(\delta)$ for the risk $\mathcal{R}(Q)$ in (17).

Whether the constraint $\mathbb{Q} \in \mathbb{B}_\epsilon(\delta)$ is active or not depends on the value taken at the optimum by its Lagrange multiplier λ, which represents the shadow cost of robustification. The following proposition provides a sufficient condition for the constraint to be active by generalizing the example shown in Fig. 1 to $\mathbb{R}^d$.

Proposition 3: Let $\partial\Xi = \{\boldsymbol{\xi} : \max_{k\in[n_H]} H_k\boldsymbol{\xi} - h_k = 0\}$, where n_H is the number of rows in H, denote the boundary of Ξ. If

$$\frac{1}{N}\sum_{i\in[N]} \min_{\tilde{\boldsymbol{\xi}}\in\partial\Xi} \left\|\boldsymbol{\xi}_T^{(i)} - \tilde{\boldsymbol{\xi}}\right\|_2^2 > \epsilon \,, \tag{18}$$

that is, if the average squared distance between the samples and the border $\partial\Xi$ of the support Ξ is strictly greater than epsilon, then the optimal shadow cost of robustification $\lambda^\star$ is greater than $\lambda_{\max}(Q)$ for any $Q \in \mathbb{R}^{d\times d}$.

Proof: The proof is given in Appendix D. ■

Proposition 3 shows that λ is contingent on the radius ϵ, the support Ξ, and the realizations $\boldsymbol{\xi}^{(i)}$. The radius ϵ is usually small, as the samples should be approximating the real distribution well enough, which means that the condition (18) is often satisfied. In the next section, we utilize the inequality $\lambda^\star \geq \lambda_{\max}(Q)$ to propose a strong dual formulation for (17).

B. Tight convex relaxation for DRO of quadratic objectives

In this section, we present a convex upper bound for (17), and prove that it becomes tight if λ is greater than $\lambda_{\max}(Q)$, the largest eigenvalue of Q.

Lemma 4: Let $Q \in \mathbb{R}^{d\times d}$ be a symmetric and positive definite matrix. Under Assumption 2, if $\epsilon > 0$ and if Ξ is bounded, the risk (17) satisfies

$$\mathcal{R}(Q) \leq \inf_{\substack{\lambda\geq 0,\mu_i\geq 0,\\ \psi_i\geq -\mu_i\\ \alpha\geq 0}} \lambda\epsilon + \frac{1}{N}\sum_{i\in[N]} s^{(i)}\,, \tag{19a}$$

$$\text{subject to}\,, \forall i\in[N]:$$

$$\begin{bmatrix} s^{(i)}-h^\top\psi_i+\lambda\|\boldsymbol{\xi}_T^{(i)}\|_2^2 & \star & \star\\ 2\lambda\boldsymbol{\xi}_T^{(i)}+H^\top\psi_i & 4(\lambda I-Q) & \star\\ H^\top\mu_i & 0 & 4Q\end{bmatrix}\succeq 0, \tag{19b}$$

$$\begin{bmatrix}\alpha & \star\\ H^\top\mu_i & \lambda I-Q\end{bmatrix}\succeq 0. \tag{19c}$$

Moreover, (19a) holds with equality and (19c) is inactive if the optimum $\lambda^\star$ of λ satisfies $\lambda^\star I \succ Q$.

Proof: This result is obtained by taking the limit of (16) when the number J of pieces tends to infinity. The detailed derivations are presented in Appendix E. ■

We stress that our results continue to hold even if $H = 0$ and $h = 0$, that is, if $\Xi = \mathbb{R}^d$. In this case, (19) simplifies substantially.

Corollary 5: Lemma 4 also holds if $\Xi = \mathbb{R}^d$ and (19) simplifies into

$$\mathcal{R}(Q) = \inf_{\lambda\geq 0} \lambda\epsilon + \frac{1}{N}\sum_{i\in[N]} s^{(i)}\,, \tag{20a}$$

$$\text{subject to}\begin{bmatrix} s^{(i)}+\lambda\|\boldsymbol{\xi}_T^{(i)}\|_2^2 & \star\\ \lambda\boldsymbol{\xi}_T^{(i)} & \lambda I-Q\end{bmatrix}\succeq 0. \tag{20b}$$

Proof: If $\Xi = \mathbb{R}^d$, the problem (13) falls into the assumptions of [2, Theorem 11]. Additionally, we observe that, when $H = 0$ and $h = 0$, (20b) has the same Schur complement as (19b) and (19c) is always satisfied. ■

To understand the effect of having restricted our attention to distributions with bounded support, it is of interest to compare (19) with (20). In both problems, the presence of the term $\lambda I - Q$ in (19b) and (20b) implies that any feasible solution has a shadow cost λ greater or equal than $\lambda_{\max}(Q)$. On the other hand, for (20b) to be feasible, λ should be large enough to guarantee $s^{(i)} + \lambda\|\boldsymbol{\xi}_T^{(i)}\|_2^2 \geq 0$, whereas the presence of the additional term $-h^\top\psi_i$ in the top-left entry of (19b) softens this requirement, demonstrating the helpful contribution of the bounded support.

C. Convex formulation of DRInC design

Our results of Section IV-B does not directly allow us to solve (13), as (12) shows that Q depends quadratically on $\boldsymbol{\Phi}$ and may also be rank deficient. In this subsection, we mitigate the issues associated with quadratic matrix inequalities by employing a Schur complement, and we address singularity concerns by examining the behavior of the system as Q approaches singularity, showing that this limit remains well-behaved.

Lemma 6: Under Assumption 2, if $\epsilon > 0$ and Ξ is bounded, the optimal closed loop map $\boldsymbol{\Phi}^\star$ in (13) is given by

$$\boldsymbol{\Phi}^\star = \arg_{\boldsymbol{\Phi}} \min_{\boldsymbol{\Phi}\text{ achievable},Q}\, \lim_{\eta\to 0} \mathcal{R}(Q+|\eta| I)\,, \tag{21a}$$

$$\text{subject to}\begin{bmatrix} Q & \star\\ D^{\frac{1}{2}}\Phi & I\end{bmatrix}\succeq 0\,. \tag{21b}$$

Proof: The proof can be found in Appendix F ■

We continue our derivations by presenting an equivalent convex reformulation of the safety constraints in (14). In particular, in the next proposition, we embed the function $\max\{\cdot - \tau, 0\}$ in (4) as a $(J+1)^{th}$ constraint.

Lemma 7: Under Assumption 2 and if $\epsilon > 0$, the constraints (14) can be reformulated as the following convex LMIs

$$\rho\epsilon + \frac{\gamma-1}{\gamma}\tau + \frac{1}{N}\sum_{i\in[N]}\zeta^{(i)} \leq 0\,,\ \rho\geq 0\,, \tag{22a}$$

$$\forall i\in[N]\,, \forall j\in[J+1]:\ \kappa_{ij}\geq 0\,, \tag{22b}$$

$$\begin{bmatrix}\zeta^{(i)}-\frac{1}{\gamma}(G_j^\top\boldsymbol{\Phi}\boldsymbol{\xi}_T^{(i)}+g_j)-(H\boldsymbol{\xi}_T^{(i)}+h)^\top\kappa_{ij} & \star\\ \frac{1}{\gamma}\boldsymbol{\Phi}^\top G_j - H^\top\kappa_{ij} & 4\rho\gamma^2 I\end{bmatrix}\succeq 0, \tag{22c}$$

where $G_{J+1} = 0$ and $g_{J+1} = \tau$.

Proof: The proof can be found in Appendix G. ■

Leveraging Lemmas 4, (6), and (7), we are now ready to reformulate (13) subject to (14) as SDP.

Theorem 8: Under Assumption 2 and if $\epsilon > 0$, the closed loop map given by

$$\boldsymbol{\Phi}^\star = \arg\min_{\boldsymbol{\Phi}\text{ achievable}} \inf_{\substack{Q,s^{(i)},\zeta^{(i)},\tau,\\ \lambda\geq 0,\rho\geq 0,\alpha\geq 0,\\ \mu_i\geq 0,\kappa_{ij}\geq 0,\\ \psi_i\geq -\mu_i}} \lambda\epsilon + \frac{1}{N}\sum_{i\in[N]} s^{(i)}\,,$$

$$\begin{aligned}&\text{subject to}\\ &(21\text{b}), (22\text{a}),\\ &(19\text{b}), (19\text{c}),\quad \forall i\in[N],\\ &(22\text{c}),\qquad\quad \forall i\in[N], j\in[J+1],\end{aligned}$$

is stable and satisfies the safety constraints (14). Moreover, it optimizes (13) if Ξ is bounded and the optimizer $\lambda^\star$ is greater than $\lambda_{\max}(\boldsymbol{\Phi}^{\star\top} D\boldsymbol{\Phi}^\star)$.

Proof: We first highlight that $\boldsymbol{\Phi}$ is FIR and therefore stable by definition. Second, the safety constraints (14) are equivalent to (22), as shown in Lemma 7. Third, consider a closed loop map $\widehat{\boldsymbol{\Phi}}$, which optimizes the expectation of $\boldsymbol{\xi}_T^\top\widehat{\boldsymbol{\Phi}}^\top D\widehat{\boldsymbol{\Phi}}\boldsymbol{\xi}_T + |\eta|\|\boldsymbol{\xi}_T\|_2^2$ for $\eta\neq 0$. With $Q = \boldsymbol{\Phi}^\top D\boldsymbol{\Phi} + |\eta| I \succ 0$, Lemma 4 shows that $\mathcal{R}(\boldsymbol{\Phi}^\top D\boldsymbol{\Phi})$ is tightly upper-bounded by (19). Fourth and finally, as shown in Lemma 6, taking the limit $\eta\to 0$ yields $\widehat{\boldsymbol{\Phi}}\to\boldsymbol{\Phi}^\star$ from (13), which concludes the proof. ■

We remark that the reformulation proposed in Theorem 8 is exact whenever the true shadow cost of robustification λ is greater or equal than $\lambda_{\max}(Q)$, a condition which is always satisfied for sufficiently small ϵ as per Proposition 3. When λ is lower than $\lambda_{\max}(Q)$, the solution computed using Theorem 8 may instead be suboptimal. Nevertheless, our solution retains safety and stability guarantees in face of the uncertain distribution, since neither (22) nor the achievability constraints depend on λ.

V. Conclusion

We have presented an end-to-end synthesis method from a collection of a finite number of disturbance realizations to the design of a stabilizing linear policy with DR safety and performance guarantees. Our approach consists in estimating an empirical distribution using samples of the uncertainty, and then computing a feedback policy that safely minimizes the worst-case expected cost over all distributions within a Wasserstein ball around the nominal estimate through the solution of an SDP. We have shown that, as the radius of this ambiguity set varies, our problem statement recovers classical control formulations. To address the resulting optimal control problem, we have established a novel tight convex relaxation for DRO of quadratic objectives, and we have combined our results with the system level synthesis framework, presenting conditions under which our design method is non-conservative.

Future work will validate the effectiveness of our approach by means of numerical simulations and real-world experiments.

References

[1] 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.

[2] D. Kuhn, P. M. 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.

[3] S. Shafieezadeh-Abadeh, L. Aolaritei, F. Dörfler, and D. Kuhn, "New perspectives on regularization and computation in optimal transport-based distributionally robust optimization," *arXiv preprint arXiv:2303.03900*, 2023.

[4] A. L. Gibbs and F. E. Su, "On choosing and bounding probability metrics," *International statistical review*, vol. 70, no. 3, pp. 419–435, 2002.

[5] C. Villani *et al.*, *Optimal transport: old and new*. Springer, 2009, vol. 338.

[6] Z. Chen, D. Kuhn, and W. Wiesemann, "Data-driven chance constrained programs over wasserstein balls," *Operations Research*, 2022.

[7] C. Frogner, C. Zhang, H. Mobahi, M. Araya, and T. A. Poggio, "Learning with a wasserstein loss," *Advances in neural information processing systems*, vol. 28, 2015.

[8] D. O. Adu, T. Başar, and B. Gharesifard, "Optimal transport for a class of linear quadratic differential games," *IEEE Transactions on Automatic Control*, vol. 67, no. 11, pp. 6287–6294, 2022.

[9] B. Taşkesen, D. A. Iancu, Ç. Koçyiğit, and D. Kuhn, "Distributionally robust linear quadratic control," *arXiv preprint arXiv:2305.17037*, 2023.

[10] C. Mark and S. Liu, "Stochastic MPC with distributionally robust chance constraints," *IFAC-PapersOnLine*, vol. 53, no. 2, pp. 7136–7141, 2020.

[11] M. Fochesato and J. Lygeros, "Data-driven distributionally robust bounds for stochastic model predictive control," in *2022 IEEE 61st Conference on Decision and Control (CDC)*. IEEE, 2022, pp. 3611–3616.

[12] L. Aolaritei, M. Fochesato, J. Lygeros, and F. Dörfler, "Wasserstein tube MPC with exact uncertainty propagation," *arXiv preprint arXiv:2304.12093*, 2023.

[13] I. Yang, "Wasserstein distributionally robust stochastic control: A data-driven approach," *IEEE Transactions on Automatic Control*, vol. 66, no. 8, pp. 3863–3870, 2020.

[14] K. Kim and I. Yang, "Distributional robustness in minimax linear quadratic control with Wasserstein distance," *SIAM Journal on Control and Optimization*, vol. 61, no. 2, pp. 458–483, 2023.

[15] V. Krishnan and S. Martínez, "A probabilistic framework for moving-horizon estimation: Stability and privacy guarantees," *IEEE Transactions on Automatic Control*, vol. 66, no. 4, pp. 1817–1824, 2020.

[16] J.-S. Brouillon, F. Dörfler, and G. Ferrari-Trecate, "Regularization for distributionally robust state estimation and prediction," *arXiv preprint arXiv:2304.09921*, 2023.

[17] L. Aolaritei, N. Lanzetti, H. Chen, and F. Dörfler, "Distributional uncertainty propagation via optimal transport," *arXiv preprint arXiv:2205.00343*, 2023.

[18] L. Aolaritei, N. Lanzetti, and F. Dörfler, "Capture, propagate, and control distributional uncertainty," *arXiv preprint arXiv:2304.02235*, 2023.

[19] Y.-S. Wang, N. Matni, and J. C. Doyle, "A system-level approach to controller synthesis," *IEEE Transactions on Automatic Control*, vol. 64, no. 10, pp. 4079–4093, 2019.

[20] J. Coulson, J. Lygeros, and F. Dörfler, "Distributionally robust chance constrained data-enabled predictive control," *IEEE Transactions on Automatic Control*, vol. 67, no. 7, pp. 3289–3304, 2021.

[21] A. Hakobyan and I. Yang, "Wasserstein distributionally robust control of partially observable linear systems: Tractable approximation and performance guarantee," in *2022 IEEE 61st Conference on Decision and Control (CDC)*. IEEE, 2022, pp. 4800–4807.

[22] K. I. Park, M. Park, and James, *Fundamentals of probability and stochastic processes with applications to communications*. Springer, 2018.

[23] 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.

[24] T. S. Badings, A. Abate, N. Jansen, D. Parker, H. A. Poonawala, and M. Stoelinga, "Sampling-based robust control of autonomous systems with non-gaussian noise," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 36, no. 9, 2022, pp. 9669–9678.

[25] L. Blackmore, M. Ono, A. Bektassov, and B. C. Williams, "A probabilistic particle-control approximation of chance-constrained stochastic predictive control," *IEEE transactions on Robotics*, vol. 26, no. 3, pp. 502–517, 2010.

[26] B. Hassibi, A. H. Sayed, and T. Kailath, *Indefinite-quadratic estimation and control: a unified approach to $\mathcal{H}_2$ and $\mathcal{H}_\infty$ theories*. SIAM, 1999.

[27] K. Zhou and J. C. Doyle, *Essentials of robust control*. Prentice hall Upper Saddle River, NJ, 1998, vol. 104.

[28] J. Anderson, J. C. Doyle, S. H. Low, and N. Matni, "System level synthesis," *Annual Reviews in Control*, vol. 47, pp. 364–393, 2019.

[29] S. Dean, H. Mania, N. Matni, B. Recht, and S. Tu, "On the sample complexity of the linear quadratic regulator," *Foundations of Computational Mathematics*, vol. 20, no. 4, pp. 633–679, 2020.

[30] 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.

[31] S. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge university press, 2004.

[32] J. Borwein and A. Lewis, *Convex Analysis and Nonlinear Optimization: Theory and Examples*. Springer New York, 2005.

[33] R. G. Bartle, *The elements of integration and Lebesgue measure*. John Wiley & Sons, 2014.

[34] M. Sion, "On general minimax theorems." *Pacific Journal of Mathematics*, vol. 8, no. 4, pp. 171–176, 1958.

Appendix

A. SLS controller implementation

From [19],

$$\begin{aligned}\boldsymbol{\delta} &= (I - z\boldsymbol{\Phi}_{xw}(z))\boldsymbol{\delta} - \boldsymbol{\Phi}_{xv}(z)\boldsymbol{y},\\ \boldsymbol{u} &= z\boldsymbol{\Phi}_{uw}(z)\boldsymbol{\delta} + \boldsymbol{\Phi}_{uv}(z)\boldsymbol{y},\end{aligned}$$

which means that at each timestep $t \geq T$, one has

$$\delta_t = \delta_t - \sum_{k=1}^{T} \Phi_{xw}(k)\delta_{t-k+1} - \sum_{k=1}^{T} \Phi_{xv}(k) y_{t-k},$$

$$u_t = \sum_{k=1}^{T} \Phi_{uw}(k)\delta_{t-k+1} + \sum_{k=0}^{T} \Phi_{uv}(k) y_{t-k}. \tag{23}$$

The achievability constraints (24) imply $\Phi_{xw}(1) = I$ and $\Phi_{uw}(1) = \Phi_{uv}(0)C$, (see Appendix B). Hence, (23) can be reformulated at as

$$\delta_t = -\sum_{k=1}^{T-1} \Phi_{xw}(k+1)\delta_{t-k} - \sum_{k=1}^{T} \Phi_{xv}(k) y_{t-k},$$

$$u_t = \Phi_{uv}(0)C\delta_t + \sum_{k=1}^{T-1} \Phi_{uw}(k+1)\delta_{t-k} + \sum_{k=0}^{T} \Phi_{uv}(k) y_{t-k}.$$

Writing this controller implementation in matrix form and noting that $\Phi_{xv}(0) = 0$ yields (11).

B. Infinite horizon Achievability

Proposition 9: The achievability constraints (9) are equivalent to

$$[I,0]\mathbf{\Phi}\begin{bmatrix} \mathcal{Z}^{-}\otimes I_n & 0 \\ 0 & \mathcal{Z}^{-}\otimes I_p \end{bmatrix} = [A,B]\mathbf{\Phi}\begin{bmatrix} \mathcal{Z}^{+}\otimes I_n & 0 \\ 0 & \mathcal{Z}^{+}\otimes I_p \end{bmatrix} + [\mathcal{Z}^{+}_{T+1}\otimes I_n, \mathcal{Z}^{+}_{T+1}\otimes 0], \tag{24a}$$

$$\mathbf{\Phi}\begin{bmatrix} \mathcal{Z}^{-}\otimes I_n \\ \mathcal{Z}^{-}\otimes(0C) \end{bmatrix} = \mathbf{\Phi}\begin{bmatrix} \mathcal{Z}^{+}\otimes A \\ \mathcal{Z}^{+}\otimes C \end{bmatrix} + \begin{bmatrix} \mathcal{Z}^{+}_{T+1}\otimes I_n \\ \mathcal{Z}^{+}_{T+1}\otimes(0C) \end{bmatrix}, \tag{24b}$$

where $\mathcal{Z}^{+} = [I_{T+1}, 0]$, $\mathcal{Z}^{-} = [0, I_{T+1}]$ are in $\mathbb{R}^{(T+1)\times(T+2)}$, and $\mathcal{Z}^{+}_{T+1}$ is the last row of $\mathcal{Z}^{+}$.

Proof: By treating $\mathbf{\Phi}_x$ and $\mathbf{\Phi}_u$ as FIR filters:

$$\sum_{k=1}^{T} \Phi_{xw}(k)z^{-k+1} - A\Phi_{xw}(k)z^{-k} - B\Phi_{uw}(k)z^{-k} = I,$$

$$\sum_{k=1}^{T} \Phi_{xv}(k)z^{-k+1} - A\Phi_{xv}(k)z^{-k} - B\Phi_{uv}(k)z^{-k} = B\Phi_{uv}(0),$$

$$\sum_{k=1}^{T} \Phi_{xw}(k)z^{-k+1} - \Phi_{xw}(k)Az^{-k} - \Phi_{xv}(k)Cz^{-k} = I,$$

$$\sum_{k=1}^{T} \Phi_{uw}(k)z^{-k+1} - \Phi_{uw}(k)Az^{-k} - \Phi_{uv}(k)Cz^{-k} = \Phi_{uv}(0)C,$$

which is equivalent to

$$\Phi_{xw}(0) = 0, \Phi_{xv}(0) = 0, \Phi_{uw}(0) = 0, \tag{25a}$$

$$\Phi_{xw}(1) = I, \Phi_{xv}(1) = B\Phi_{uv}(0), \Phi_{uw}(1) = \Phi_{uv}(0)C, \tag{25b}$$

$$\Phi_{xw}(k+1) = A\Phi_{xw}(k) + B\Phi_{uw}(k), \forall k = 1,\dots,T, \tag{25c}$$

$$\Phi_{xv}(k+1) = A\Phi_{xv}(k) + B\Phi_{uv}(k), \forall k = 1,\dots,T, \tag{25d}$$

$$\Phi_{xw}(k+1) = \Phi_{xw}(k)A + \Phi_{xv}(k)C, \forall k = 1,\dots,T, \tag{25e}$$

$$\Phi_{uw}(k+1) = \Phi_{uw}(k)A + \Phi_{uv}(k)C, \forall k = 1,\dots,T, \tag{25f}$$

$$\Phi(T+1) = 0, \tag{25g}$$

$\forall k = 1,\dots,T$ and $\Phi(T+1) = 0$. In matrix form, this yields

$$[I,0]\mathbf{\Phi}\begin{bmatrix} \begin{bmatrix} 0 & I & \dots & 0 & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & \dots & I & 0 & 0 \\ 0 & 0 & \dots & 0 & I & 0 \\ 0 & 0 & \dots & 0 & 0 & I \end{bmatrix} & 0 \\ 0 & \begin{bmatrix} 0 & I & \dots & 0 & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & \dots & I & 0 & 0 \\ 0 & 0 & \dots & 0 & I & 0 \\ 0 & 0 & \dots & 0 & 0 & I \end{bmatrix} \end{bmatrix}$$
$$\underbrace{\quad}_{(25g)}\underbrace{\qquad\quad}_{(25c)}\underbrace{\quad}_{(25b)}\underbrace{\quad}_{(25a)}\;\underbrace{\quad}_{(25g)}\underbrace{\qquad\quad}_{(25d)}\underbrace{\quad}_{(25b)}\underbrace{\quad}_{(25a)}$$

$$= \big[[0,0,\dots,0,I,0],[0,0,\dots,0,0,0]\big] +$$

$$[A,B]\mathbf{\Phi}\begin{bmatrix} \begin{bmatrix} I & 0 & \dots & 0 & 0 & 0 \\ 0 & I & \dots & 0 & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & \dots & I & 0 & 0 \\ 0 & 0 & \dots & 0 & I & 0 \end{bmatrix} & 0 \\ 0 & \begin{bmatrix} I & 0 & \dots & 0 & 0 & 0 \\ 0 & I & \dots & 0 & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & \dots & I & 0 & 0 \\ 0 & 0 & \dots & 0 & I & 0 \end{bmatrix} \end{bmatrix},$$
$$\underbrace{\quad}_{(25g)}\underbrace{\qquad\quad}_{(25c)}\underbrace{\quad}_{(25b)}\underbrace{\quad}_{(25a)}\;\underbrace{\quad}_{(25g)}\underbrace{\qquad\quad}_{(25d)}\underbrace{\quad}_{(25b)}\underbrace{\quad}_{(25a)}$$

and

$$\mathbf{\Phi}\begin{bmatrix} \begin{bmatrix} 0 & I & \dots & 0 & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & \dots & I & 0 & 0 \\ 0 & 0 & \dots & 0 & I & 0 \\ 0 & 0 & \dots & 0 & 0 & I \end{bmatrix} \\ \begin{bmatrix} 0 & 0 & \dots & 0 & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & \dots & 0 & 0 & 0 \\ 0 & 0 & \dots & 0 & 0 & 0 \\ 0 & 0 & \dots & 0 & 0 & 0 \end{bmatrix} \end{bmatrix} = \mathbf{\Phi}\begin{bmatrix} \begin{bmatrix} A & 0 & \dots & 0 & 0 & 0 \\ 0 & A & \dots & 0 & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & \dots & A & 0 & 0 \\ 0 & 0 & \dots & 0 & A & 0 \end{bmatrix} \\ \begin{bmatrix} C & 0 & \dots & 0 & 0 & 0 \\ 0 & C & \dots & 0 & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & \dots & C & 0 & 0 \\ 0 & 0 & \dots & 0 & C & 0 \end{bmatrix} \end{bmatrix}$$
$$\underbrace{\quad}_{(25g)}\underbrace{\qquad\quad}_{(25e),(25f)}\underbrace{\quad}_{(25b)}\underbrace{\quad}_{(25a)}\qquad\underbrace{\quad}_{(25g)}\underbrace{\qquad\quad}_{(25e),(25f)}\underbrace{\quad}_{(25b)}\underbrace{\quad}_{(25a)}$$
$$+ \begin{bmatrix} 0,0,\dots,0,I,0 \\ 0,0,\dots,0,0,0 \end{bmatrix}.$$

The matrices can be written in a compact form as (24), which concludes the proof. ∎

C. Proof of Proposition 2

The risk (15) is contingent on three mathematical objects:

1) A loss function $\max_{j\in[J]} \ell_j(\boldsymbol{\xi}_T) = \max_{j\in[J]} a_j^\top \boldsymbol{\xi}_T + b_j$,
2) A transport cost $c(\boldsymbol{\xi}_T, \boldsymbol{\xi}_T^{(i)}) = \|\boldsymbol{\xi}_T - \boldsymbol{\xi}_T^{(i)}\|_2^2$,
3) and a support $\boldsymbol{\Xi} = \left\{\boldsymbol{\xi} : \max_{k\in[n_H]} f_k(\boldsymbol{\xi}) \leq 0\right\}$, where n_H is the number of rows in H and $f_k(\boldsymbol{\xi}) = H_k\boldsymbol{\xi} - h_k$.

Moreover, since the loss is concave and both the transport cost and the support are convex, (15) shows strong duality properties if and only if it is strictly feasible. The strict feasibility is guaranteed by the full-dimensionality of Ξ and the strict positivity of ϵ. The dual problem is given by [3] as

$$\inf_{\lambda\geq 0} \lambda\epsilon + \frac{1}{N}\sum_{i=0}^{N} s^{(i)},$$
$$\text{subject to} \sup_{\boldsymbol{\xi}_T\in\Xi} \ell(\boldsymbol{\xi}_T) - \lambda c(\boldsymbol{\xi}_T, \boldsymbol{\xi}_T^{(i)}) \leq s^{(i)}\,, \forall i \in [N].$$

While the dual problem does not seem much simpler to solve than the primal at first glance, we use [3, Proposition 2.12] to reformulate it using convex conjugates. In our own notation, this gives

$$\inf_{\lambda\geq 0\,,\kappa_{ijk}\geq 0} \lambda\epsilon + \frac{1}{N}\sum_{i\in[N]} s^{(i)}, \text{ subject to}\,, \forall i\in[N], \forall j\in[J]:$$
$$s^{(i)} \geq (-\ell_j)^\star(\zeta_{ij}^\ell) + \lambda c^\star\left(\frac{\zeta_{ij}^c}{\lambda}, \widehat{\boldsymbol{\xi}}_T^{(i)}\right) + \sum_{k\in[n_H]} \kappa_{ijk} f_k^\star\left(\frac{\zeta_{ijk}^f}{\kappa_{ijk}}\right),$$
$$\zeta_{ij}^\ell + \zeta_{ij}^c + \sum_{k\in[n_H]} \zeta_{ijk}^f = 0, \tag{26}$$

where $(-\ell_j)^\star$ is the convex conjugate of the opposite of ℓ_j, $c^\star$ is the convex conjugate of the transport cost c with respect to the first argument, and $f_k^\star$ is the convex conjugate of f_k. Note that the case where $\lambda = 0$ is also well defined in [3] despite the division. All three functions are either linear or quadratic so their conjugates are well-known [32]. Both $-\ell_j$ and f_k are linear so their convex conjugates are b_j and h_k if the conjugates' arguments are equal to $-a_j$ and H_k, respectively, and infinite otherwise. The conjugate of the transport cost is given by $c^\star(\zeta, \boldsymbol{\xi}_T^{(i)}) = \frac{1}{4}\zeta^\top\zeta - \zeta^\top\boldsymbol{\xi}_T^{(i)}$.

In order to minimize (26), one must avoid infinite costs, which adds constraints on λ, ζ_{ij}^ℓ, and ζ_{ijk}^f. This means that (26) is equivalent to

$$\inf_{\lambda\geq 0\,,\kappa_{ijk}\geq 0} \lambda\epsilon + \frac{1}{N}\sum_{i\in[N]} s^{(i)} \leq 0\,, \tag{27a}$$
$$\text{subject to}\,, \forall i \in [N]\,, \forall j \in [J]:$$
$$s^{(i)} \geq b_j + \frac{1}{4\lambda}(\zeta_{ij}^c)^\top\zeta_{ij}^c - \zeta_{ij}^c\boldsymbol{\xi}_T^{(i)} + \sum_{k\in[n_H]} \kappa_{ijk}h_k, \tag{27b}$$
$$\zeta_{ij}^\ell + \zeta_{ij}^c + \sum_{k\in[n_H]} \zeta_{ijk}^f = 0\,, \zeta_{ij}^\ell = -a_j\,, \zeta_{ijk}^f = \kappa_{ijk}H_k. \tag{27c}$$

To conclude the proof, we stack κ_{ijk} for all $k \in [n_H]$ into a vector κ_{ij} and plug the equality constraints (27c) into (27b) to obtain (16).

D. Proof of Proposition 3

We first note that when Ξ is bounded, as mass cannot be moved infinitely far away, the supremum of (17) is attained. This means that $\mathbb{Q}^\star = \arg\max_{\mathbb{Q}\in\mathbb{B}_\epsilon} \mathbb{E}_{\xi_T\sim\mathbb{Q}}\ \boldsymbol{\xi}_T^\top Q\boldsymbol{\xi}_T$ exists. Second, the limited average squared distance between the samples and the border of Ξ implies that no distribution in $\mathbb{B}_\epsilon(\mathbb{P})$ has mass only at the border of Ξ, as the trasport cost would be greater than ϵ. This means that there exists a $\delta > 0$ such that $\mathbb{Q}^\star$ has an amount δ of mass more than $\sqrt{\delta}$ away from the boundary of Ξ. Third and finally, let $w_{\max}(Q)$ be an eigenvector of Q associated with $\lambda_{\max}(Q)$. The distribution $\mathbb{Q}^\star$ satisfies

$$\begin{aligned}\frac{d\mathcal{R}(Q)}{d\epsilon} &= \lim_{d\epsilon\to 0^+}\frac{1}{d\epsilon}\sup_{\mathbb{Q}'\in\mathbb{B}_{d\epsilon}(\mathbb{Q}^\star)} \mathbb{E}_{\substack{\boldsymbol{\xi}'_T\sim\mathbb{Q}'\\ \boldsymbol{\xi}_T\sim\mathbb{Q}^\star}} {\boldsymbol{\xi}'_T}^\top Q\boldsymbol{\xi}'_T - {\boldsymbol{\xi}_T}^\top Q\boldsymbol{\xi}_T \\ &= \lim_{d\epsilon\to 0^+}\frac{1}{d\epsilon}\sup_{\mathbb{Q}'\in\mathbb{B}_{d\epsilon}(\mathbb{Q}^\star)} \mathbb{E}_{\substack{\boldsymbol{\xi}'_T\sim\mathbb{Q}'\\ \boldsymbol{\xi}_T\sim\mathbb{Q}^\star}} (\boldsymbol{\xi}'_T - \boldsymbol{\xi}_T)^\top Q(\boldsymbol{\xi}'_T - \boldsymbol{\xi}_T) \\ &\qquad - 2(\boldsymbol{\xi}'_T - \boldsymbol{\xi}_T)^\top Q\boldsymbol{\xi}_T \\ &\geq \lim_{d\epsilon\to 0^+}\frac{1}{d\epsilon}\max_{\delta\|d\boldsymbol{\xi}\|_2^2\leq d\epsilon} \delta\lambda_{\max}(Q)\|d\boldsymbol{\xi}\|_2^2,\end{aligned} \tag{28}$$

because moving δ of $\mathbb{Q}^\star$'s mass by $\|d\boldsymbol{\xi}\| \leq \sqrt{\delta^{-1}d\epsilon}$ in the direction of $\pm w_{\max}(Q)$ to obtain $\mathbb{Q}'$ remains in $\mathbb{B}_{d\epsilon}(\mathbb{Q}^\star)$ if $d\epsilon \leq \delta^2$, which is true at the limit $d\epsilon \to 0^+$. Hence, the inequality (28) implies that

$$\lambda^\star = \frac{d\mathcal{R}(Q)}{d\epsilon} \geq \lim_{d\epsilon\to 0^+}\frac{1}{d\epsilon}\lambda_{\max}(Q)d\epsilon = \lambda_{\max}(Q),$$

which concludes the proof.

E. Proof of Lemma 4

In order to prove Lemma 4, we first need the following proposition.

Proposition 10: Under the assumptions of Lemma 4, the risk $\mathcal{R}(Q)$ defined in (17) satisfies

$$\mathcal{R}(Q) \leq \inf_{\lambda\geq 0,\kappa_i\geq 0} \lambda\epsilon + \frac{1}{N}\sum_{i\in[N]} s^{(i)}\,, \text{ subject to} \tag{29a}$$
$$s^{(i)} \geq \max_{\bar{\boldsymbol{\xi}}\in\Xi} -\bar{\boldsymbol{\xi}}^\top(Q - \lambda^{-1}Q^2)\bar{\boldsymbol{\xi}} - 2\bar{\boldsymbol{\xi}}^\top Q\boldsymbol{\xi}_T^{(i)} \tag{29b}$$
$$+ \frac{1}{4\lambda}\kappa_i^\top HH^\top\kappa_i - \frac{1}{\lambda}\bar{\boldsymbol{\xi}}^\top QH^\top\kappa_i + \big(H\boldsymbol{\xi}_T^{(i)} + h\big)^\top\kappa_i\,,$$

for all $i = 1, \ldots, N$. Moreover, (19a) holds with equality if the optimum $\lambda^\star$ of λ satisfies $\lambda^\star I \succeq Q$.

Proof: The proof starts by linking the formulation (15) for piece-wise affine costs to $\mathcal{R}(Q)$. To do so, we approximate the quadratic cost using its tangents at each point of a d-dimensional grid $\mathcal{G}_J \subseteq \Xi$, composed of J points. Because the approximation gets closer with more points, this yields

$$\boldsymbol{\xi}_T^\top Q\boldsymbol{\xi}_T = \lim_{J\to\infty}\max_{j\in[J]} 2\boldsymbol{\xi}_T^\top Q\boldsymbol{\xi}_j - \boldsymbol{\xi}_j^\top Q\boldsymbol{\xi}_j\,,$$

where $\boldsymbol{\xi}_j$ is the j^{th} element of $\mathcal{G}_J$. In order to obtain a formulation that fits (15), one must show that the limit operator commutes with the supremum and the expectation.

We show the commutation of the limit with the dominated convergence theorem [33] by finding bounds on the piece-wise affine approximation error

$$\begin{aligned}\Delta_J &= \boldsymbol{\xi}_T^\top Q\boldsymbol{\xi}_T - \max_{j\in[J]} 2\boldsymbol{\xi}_T^\top Q\boldsymbol{\xi}_j - \boldsymbol{\xi}_j^\top Q\boldsymbol{\xi}_j, \\ &= \boldsymbol{\xi}_T^\top Q\boldsymbol{\xi}_T + \min_{j\in[J]} \boldsymbol{\xi}_j^\top Q\boldsymbol{\xi}_j - 2\boldsymbol{\xi}_T^\top Q\boldsymbol{\xi}_j, \\ &= \min_{j\in[J]} (\boldsymbol{\xi}_T - \boldsymbol{\xi}_j)^\top Q(\boldsymbol{\xi}_T - \boldsymbol{\xi}_j)\,.\end{aligned}$$

Note that $\Delta_J \geq 0$ because the tangents of a quadratic function are always below the curve. Moreover, the inequality $\Delta_J \leq \lambda_{\max}(Q) \min_{j\in[J]} \|\boldsymbol{\xi}_T - \boldsymbol{\xi}_j\|_2^2$ is satisfied by definition.

Furthermore, the distance $\min_{j\in[J]} \|\boldsymbol{\xi}_T - \boldsymbol{\xi}_j\|_2^2$ between any $\boldsymbol{\xi}_T \in \boldsymbol{\Xi}$ and the closest point of the grid $\mathcal{G}_J \subseteq \boldsymbol{\Xi}$ can be bounded as

$$\min_{j\in[J]} \|\boldsymbol{\xi}_T - \boldsymbol{\xi}_j\|_2^2 \leq 2r(\boldsymbol{\Xi})\sqrt{d}J^{-\frac{1}{d}}, \forall \boldsymbol{\xi}_T \in \boldsymbol{\Xi},$$

where $r(\boldsymbol{\Xi}) < \infty$ is the radius of a ball containing $\boldsymbol{\Xi}$, which is finite because $\boldsymbol{\Xi}$ is bounded. This gives the following inequality

$$\boldsymbol{\xi}_T^\top Q\boldsymbol{\xi}_T - \Delta_Q J^{-\frac{1}{d}} \leq \max_{j\in[J]} 2\boldsymbol{\xi}_T^\top Q\boldsymbol{\xi}_j - \boldsymbol{\xi}_j^\top Q\boldsymbol{\xi}_j \leq \boldsymbol{\xi}_T^\top Q\boldsymbol{\xi}_T,$$

where $\Delta_Q = 2r(\boldsymbol{\Xi})\sqrt{d}\lambda_{\max}(Q)$. Finally, if all points of a function satisfy an inequality, its supremum must satisfy i as well, hence

$$\mathcal{R}(Q) - \Delta_Q J^{-\frac{1}{d}} \leq \sup_{\mathbb{Q}\in\mathbb{B}_\epsilon} \mathbb{E}_{\xi_T\sim\mathbb{Q}} \max_{j\in[J]} 2\boldsymbol{\xi}_T^\top Q\boldsymbol{\xi}_j - \boldsymbol{\xi}_j^\top Q\boldsymbol{\xi}_j \leq \mathcal{R}(Q).$$

The limit $\lim_{J\to\infty} \mathcal{R}(Q) - \Delta_Q J^{-\frac{1}{d}}$ is equal to $\mathcal{R}(Q)$. Therefore, the supremum of the piece-wise linear approximation is squeezed into the equality

$$\lim_{J\to\infty} \sup_{\mathbb{Q}\in\mathbb{B}_\epsilon} \mathbb{E}_{\xi_T\sim\mathbb{Q}} \max_{j\in[J]} 2\boldsymbol{\xi}_T^\top Q\boldsymbol{\xi}_j - \boldsymbol{\xi}_j^\top Q\boldsymbol{\xi}_j = \mathcal{R}(Q).$$

The second part of the proof aims at bringing the limit back into the problem and evaluating it. Using the previous result and Proposition 2 with $a_j = 2Q\boldsymbol{\xi}_j$ and $b_j = -\boldsymbol{\xi}_j^\top Q\boldsymbol{\xi}_j$, we know that $\mathcal{R}(Q)$ as defined in (17) is equal to

$$\lim_{J\to\infty} \inf_{\lambda\geq 0, \kappa_{ij}\geq 0} \lambda\epsilon + \frac{1}{N}\sum_{i\in[N]} s^{(i)}, \text{ subject to}$$
$$s^{(i)} \geq f(\boldsymbol{\xi}_j, \kappa_{ij}, \lambda), \forall i \in [N], \forall j \in [J],$$

where

$$f(\boldsymbol{\xi}, \kappa, \lambda) = -\boldsymbol{\xi}^\top Q\boldsymbol{\xi} + \frac{1}{\lambda}\boldsymbol{\xi}^\top Q^2\boldsymbol{\xi} - 2\boldsymbol{\xi}^\top Q\boldsymbol{\xi}_T^{(i)} + \frac{1}{4\lambda}\kappa^\top HH^\top\kappa - \frac{1}{\lambda}\boldsymbol{\xi}^\top QH^\top\kappa + \left(H\boldsymbol{\xi}_T^{(i)} + h\right)^\top\kappa,$$

Since there are only existence constraints for κ_{ij}, one can equivalently write

$$\lim_{J\to\infty} \inf_{\lambda\geq 0} \lambda\epsilon + \frac{1}{N}\sum_{i\in[N]} s^{(i)}, \text{ subject to}$$
$$s^{(i)} \geq \min_{\kappa_{ij}\geq 0} f(\boldsymbol{\xi}_j, \kappa_{ij}, \lambda), \forall i \in [N], \forall j \in [J],$$

The constraint holding for all j means that there are infinitely many constraints to satisfy. However, one can collapse all the constraints for a given i into

$$s^{(i)} \geq \max_{j\in[J]} \min_{\kappa_{ij}\geq 0} f(\boldsymbol{\xi}_j, \kappa_{ij}, \lambda), \forall i \in [N],$$

Interestingly, the cost does not depend on J. This means that the limit can be moved into the constraint as

$$\inf_{\lambda\geq 0} \lambda\epsilon + \frac{1}{N}\sum_{i\in[N]} s^{(i)}, \text{ subject to}$$
$$s^{(i)} \geq \lim_{J\to\infty} \max_{j\in[J]} \min_{\kappa_{ij}\geq 0} f(\boldsymbol{\xi}_j, \kappa_{ij}, \lambda), \forall i \in [N],$$

Due to the boundedness of Ξ, the grid $\mathcal{G}_J$ fills the entire set when J tends to ∞. Hence, $\mathcal{R}(Q)$ is equal to

$$\inf_{\lambda\geq 0} \lambda\epsilon + \frac{1}{N}\sum_{i\in[N]} s^{(i)}, \text{ subject to}$$
$$s^{(i)} \geq \max_{\bar{\boldsymbol{\xi}}\in\Xi} \min_{\kappa_i\geq 0} f(\bar{\boldsymbol{\xi}}, \kappa_i, \lambda), \forall i \in [N]. \tag{30a}$$

In general, one has $\max_{\bar{\boldsymbol{\xi}}\in\Xi} \min_{\kappa_i\geq 0} f(\bar{\boldsymbol{\xi}}, \kappa_i, \lambda) \leq \min_{\kappa_i\geq 0} \max_{\bar{\boldsymbol{\xi}}\in\Xi} f(\bar{\boldsymbol{\xi}}, \kappa_i, \lambda)$. This means that (29b) is a stricter constraint than (30a), yielding a larger infimum. Nevertheless, if f is not only convex in κ but also concave in $\boldsymbol{\xi}$, then Sion's minimax theorem proves that the $\max$ and $\min$ operators commute [34, Corollary 3.3]. This means that if $Q - \lambda^{-1}Q^2 \succeq 0$, (29b) and (30a) are equivalent, which concludes the proof ■

Using Proposition 10, we are now ready to prove Lemma 4 by dualizing (29b) to remove the $\max$ operator, and by using Schur's complement to obtain linear inequalities. We start by highlighting that (29b) contains the maximization of the quadratic cost

$$-\underbrace{\bar{\boldsymbol{\xi}}^\top(Q - \lambda^{-1}Q^2)\bar{\boldsymbol{\xi}}}_{\text{quadratic}} - \underbrace{\bar{\boldsymbol{\xi}}^\top(2Q\boldsymbol{\xi}_T^{(i)} + \lambda^{-1}QH^\top\kappa_i)}_{\text{linear}} + \underbrace{\frac{1}{4\lambda}\kappa_i^\top HH^\top\kappa_i + \left(H\boldsymbol{\xi}_T^{(i)} + h\right)^\top\kappa_i}_{\text{constant}},$$

subject to convex polytopic constraints $H\bar{\boldsymbol{\xi}} - h \leq 0$. The dual problem is therefore given by [31] as

$$\min_{\mu_i\geq 0} -\mu_i^\top h + \overbrace{\frac{1}{4\lambda}\kappa_i^\top HH^\top\kappa_i + \left(H\boldsymbol{\xi}_T^{(i)} + h\right)^\top\kappa_i}^{(\clubsuit)} + \frac{1}{4}\left\|H^\top\mu_i - \frac{1}{\lambda}(HQ)^\top\kappa_i - 2Q\boldsymbol{\xi}_T^{(i)}\right\|_{Q_2}^2, \tag{31a}$$
$$\text{subject to } P_\lambda\left(H^\top\mu_i - \frac{1}{\lambda}(HQ)^\top\kappa_i - 2Q\boldsymbol{\xi}_T^{(i)}\right) = 0, \tag{31b}$$

where $\|\cdot\|_{Q_2}^2 = \cdot^\top Q_2 \cdot$, $Q_2 = (Q - \lambda^{-1}Q^2)^\dagger$, and P_λ is the projection on $\text{null}(Q_2^\dagger) = \text{null}(\lambda I - Q)$. Note that $P_\lambda = I - (\lambda I - Q)^\dagger(\lambda I - Q)$ is symmetric, commutes with Q and Q^{-1}, and is equal to both its square and pseudo-inverse. Since we are looking for an upper bound for $\mathcal{R}(Q)$ when $\lambda \leq \lambda_{\max}(Q)$, we can replace (31b) by the stricter constraint

$$P_\lambda H^\top\mu_i = 0, P_\lambda\left(2\lambda\boldsymbol{\xi}_T^{(i)} + H^\top\kappa_i\right) = 0, \tag{32}$$

as it leads to a larger minimum if $P_\lambda \neq 0$ and as it is equivalent for any $\lambda > \lambda_{\max}(Q)$ because $P_\lambda = 0$. Moreover, the last term of (31a) can be split as

$$\frac{1}{4}\left\|2Q\boldsymbol{\xi}_T^{(i)} + \frac{1}{\lambda}(HQ)^\top\kappa_i\right\|_{Q_2}^2 - \frac{1}{2}\left(2Q\boldsymbol{\xi}_T^{(i)} + \frac{1}{\lambda}(HQ)^\top\kappa_i\right)^\top Q_2H^\top\mu_i + \frac{1}{4}\mu_i^\top HQ_2H^\top\mu_i,$$

or equivalently,

$$\frac{1}{4}(2\lambda\boldsymbol{\xi}_T^{(i)}+H^\top\kappa_i)^\top(\lambda^2Q^{-1}-\lambda I)^\dagger(2\lambda\boldsymbol{\xi}_T^{(i)}+H^\top\kappa_i) \tag{33a}$$

$$-\frac{1}{2}\Big(2\lambda\boldsymbol{\xi}_T^{(i)}+H^\top\kappa_i\Big)^\top(\lambda I-Q)^\dagger H^\top\mu_i \tag{33b}$$

$$+\frac{1}{4}\mu_i^\top HQ_2H^\top\mu_i\,, \tag{33c}$$

In order to obtain some simplifications, we use the following Woodbury-like identities:

$$\begin{aligned}(\lambda^2Q^{-1}-\lambda I)^\dagger &=\frac{1}{\lambda^2}\Big(Q^{-1}-\frac{1}{\lambda}I\Big)^\dagger\\
&=\Big(\frac{1}{\lambda}Q^{-1}-\frac{1}{\lambda}Q^{-1}+\frac{1}{\lambda^2}I\Big)\Big(Q^{-1}-\frac{1}{\lambda}I\Big)^\dagger\\
&=(\lambda I-Q)^\dagger+\Big(\frac{1}{\lambda^2}I-\frac{1}{\lambda}Q^{-1}\Big)\Big(Q^{-1}-\frac{1}{\lambda}I\Big)^\dagger\\
&=(\lambda I-Q)^\dagger-\frac{1}{\lambda}(I-P_\lambda) &(34a)\\
Q_2 &= \lambda Q^{-1}(\lambda I-Q)^\dagger, &(34b)\\
&=(I+\lambda Q^{-1}-I)(\lambda I-Q)^\dagger\\
&=(\lambda I-Q)^\dagger+Q^{-1}(I-P_\lambda)\\
&=(\lambda I-Q)^\dagger+Q^{-1}-P_\lambda Q^{-1}P_\lambda. &(34c)\end{aligned}$$

We plug (34a), (34b), and (34c) into (33a), (33b), and (33c), respectively, which gives

$$(33)=\frac{1}{4}\Big\|2\lambda\boldsymbol{\xi}_T^{(i)}+H^\top\kappa_i\Big\|^2_{(\lambda I-Q)^\dagger}\overbrace{-\frac{1}{4\lambda}\Big\|2\lambda\boldsymbol{\xi}_T^{(i)}+H^\top\kappa_i\Big\|^2_2}^{(\spadesuit)}$$

$$\left.+\frac{1}{4\lambda}\Big\|2\lambda\boldsymbol{\xi}_T^{(i)}+H^\top\kappa_i\Big\|^2_{P_\lambda}\right\}(\bigstar) \tag{35a}$$

$$-\frac{1}{2}\Big(H^\top\kappa_i+2\lambda\boldsymbol{\xi}_T^{(i)}\Big)^\top(\lambda I-Q)^\dagger H^\top\mu_i \tag{35b}$$

$$+\frac{1}{4}\|H^\top\mu_i\|^2_{(\lambda I-Q)^\dagger}+\underbrace{\frac{1}{4}\|H^\top\mu_i\|^2_{Q^{-1}}}_{(\blacklozenge)}\underbrace{-\frac{1}{4}\|P_\lambda H^\top\mu_i\|^2_{Q^{-1}}}_{(\bigstar)}. \tag{35c}$$

The terms ($\clubsuit$) in (31a) can be grouped by completing the squares as

$$\underbrace{\frac{1}{4\lambda}\Big\|2\lambda\boldsymbol{\xi}_T^{(i)}+H^\top\kappa_i\Big\|^2_2}_{(\spadesuit)}-\lambda\|\boldsymbol{\xi}_T^{(i)}\|_2^2+\kappa_i^\top h\,, \tag{36}$$

We remark that the terms marked by ($\spadesuit$) in (36) and (35c) cancel out, and that the terms marked by ($\bigstar$) in (35) can be factorized as

$$(2\lambda\boldsymbol{\xi}_T^{(i)}+H^\top\kappa_i+H^\top\mu_i)^\top P_\lambda\Big(2\boldsymbol{\xi}_T^{(i)}+\frac{1}{\lambda}H^\top\kappa_i-Q^{-1}H^\top\mu_i\Big), \tag{37}$$

because the cross terms are in the null space of $(\lambda I-Q)$. The constraint (31b) implies that (37) is zero, so the terms marked by ($\bigstar$) in (35) cancel out. Finally, all remaining terms besides ($\blacklozenge$) in (35c) can be factorized. Hence, the dual problem (31a) is equal to

$$\begin{aligned}\min_{\mu_i\geq 0} h^\top(\kappa_i-\mu_i)&+\frac{1}{4}\|H^\top\mu_i\|^2_{Q^{-1}}-\lambda\|\boldsymbol{\xi}_T^{(i)}\|_2^2 &(38)\\
&+\frac{1}{4}\Big\|2\lambda\boldsymbol{\xi}_T^{(i)}+H^\top(\kappa_i-\mu_i)\Big\|^2_{(\lambda I-Q)^\dagger},\end{aligned}$$

In general, the right-hand side of (29b) is smaller than (38), which means that $s^{(i)}\geq$ (38) implies (29b). Moreover, if $\lambda I-Q\succeq 0$, the problem (38) is a convex and strictly feasible QP. Strong duality therefore shows that the right-hand side of (29b) is equal to (38) in this case. Finally, we replace the upper bound on a minimum by an existence constraint and perform the change of variable $\psi_i=\kappa_i-\mu_i$ to rewrite (29b) as

$$\begin{aligned}s^{(i)}\geq h^\top\psi_i-\lambda\|\boldsymbol{\xi}_T^{(i)}\|_2^2+&\overbrace{\frac{1}{4}\mu_i^\top HQ^{-1}H^\top\mu_i}\\
&\underbrace{+\frac{1}{4}\big(2\lambda\boldsymbol{\xi}_T^{(i)}+H^\top\psi_i\big)^\top(\lambda I-Q)^\dagger\big(2\lambda\boldsymbol{\xi}_T^{(i)}+H^\top\psi_i\big)}.\end{aligned}$$

Applying Schur's lemma to the two terms highlighted with brackets and with (32), we obtain

$$\begin{aligned}\mathcal{R}(Q)\leq &\inf_{\substack{\lambda\geq 0,\mu_i\geq 0,\\ \psi_i\geq-\mu_i}}\lambda\epsilon+\frac{1}{N}\sum_{i\in[N]}s^{(i)}\,,\\
&\text{subject to}\,,\forall i\in[N]:\\
&P_\lambda H^\top\mu_i=0,\end{aligned}$$

$$\begin{bmatrix}s^{(i)}-h^\top\psi_i+\frac{1}{4\lambda}\Big\|2\lambda\boldsymbol{\xi}_T^{(i)}+H^\top\kappa_i\Big\|^2_{P_\lambda}+\lambda\|\boldsymbol{\xi}_T^{(i)}\|_2^2 & \star & \star\\ 2\lambda\boldsymbol{\xi}_T^{(i)}+H^\top\psi_i & 4(\lambda I-Q) & \star\\ H^\top\mu_i & 0 & 4Q\end{bmatrix}\succeq 0,$$

where the equality holds when $\lambda I-Q\succeq 0$. We highlight that $P_\lambda=0$ if $\lambda I-Q\succ 0$. Moreover, $P_\lambda(2\lambda\boldsymbol{\xi}_T^{(i)}+H^\top\kappa_i)=0$ because both $P_\lambda H^\top\mu_i$ and $P_\lambda(2\lambda\boldsymbol{\xi}_T^{(i)}+H^\top\psi_i)$ are zero. Finally, the constraint $P_\lambda H^\top\mu_i=0$ can be enforced as LMI using Schur's complement of $\alpha-\mu_i^\top H(\lambda I-Q)^\dagger H^\top\mu_i$ with an arbitrarily large α, which concludes the proof.

F. Proof of Lemma 6

The proof is conducted in three parts. First, we rewrite the quadratic form $\boldsymbol{\Phi}^\top D\boldsymbol{\Phi}$ as a matrix Q to obtain linear constraints. Second, we analyze the suboptimality when $Q\succ 0$ and show that it vanishes when $Q\to\boldsymbol{\Phi}^\top D\boldsymbol{\Phi}$. Third and finally, we rewrite all the constraints as LMIs.

We start by showing that

$$\mathcal{R}(\boldsymbol{\Phi}^\top D\boldsymbol{\Phi})=\min_{Q\succeq\boldsymbol{\Phi}^\top D\boldsymbol{\Phi}}\mathcal{R}(Q). \tag{39}$$

Recall the definition

$$\mathcal{R}(Q):=\sup_{\mathbb{Q}\in\mathbb{B}_\epsilon}\mathbb{E}_{\xi_T\sim\mathbb{Q}}\;\boldsymbol{\xi}_T^\top Q\boldsymbol{\xi}_T,$$

and note that for any $\boldsymbol{\xi}_T\in\boldsymbol{\Xi}$, if $Q\succeq\boldsymbol{\Phi}^\top D\boldsymbol{\Phi}$ the following inequality holds

$$\boldsymbol{\xi}_T^\top Q\boldsymbol{\xi}_T\geq\boldsymbol{\xi}_T^\top\boldsymbol{\Phi}^\top D\boldsymbol{\Phi}\boldsymbol{\xi}_T.$$

Hence, because probability distributions are non-negative and integrals preserve the order, one has

$$\mathbb{E}_{\boldsymbol{\xi}_T\sim\mathbb{Q}}\left[\boldsymbol{\xi}_T^\top Q\boldsymbol{\xi}_T\right]\geq\mathbb{E}_{\boldsymbol{\xi}_T\sim\mathbb{Q}}[\boldsymbol{\xi}_T^\top\boldsymbol{\Phi}^\top D\boldsymbol{\Phi}\boldsymbol{\xi}_T],$$

for any probability distribution $\mathbb{Q}$ and therefore also for the worst one. Hence, $Q\succeq\boldsymbol{\Phi}^\top D\boldsymbol{\Phi}$ implies that $\mathcal{R}(Q)\geq$

$\mathcal{R}(\mathbf{\Phi}^\top D\mathbf{\Phi})$. Moreover, the equality is attained because $\mathbf{\Phi}^\top D\mathbf{\Phi} \in \arg\min_{Q\succeq \mathbf{\Phi}^\top D\mathbf{\Phi}} \mathcal{R}(Q)$.

The proof continues by showing

$$\mathcal{R}(\mathbf{\Phi}^\top D\mathbf{\Phi}) = \min_{Q\succeq \mathbf{\Phi}^\top D\mathbf{\Phi}} \lim_{\eta\to 0} \mathcal{R}(Q+|\eta| I). \tag{40}$$

Note that $R(Q) = \mathcal{R}(\lim_{\eta\to 0} Q + |\eta| I)$, where one can take the limit out of the risk using the inequality

$$\mathcal{R}(Q) + |\eta| \max_{\boldsymbol{\xi}_T\in\mathbf{\Xi}} \|\boldsymbol{\xi}_T\|_2^2 \geq \mathcal{R}(Q+|\eta| I) \geq \mathcal{R}(Q), \tag{41}$$

which holds if $\mathbf{\Xi}$ is bounded. This means that the limit for $\eta \to 0$ is squeezed between two values that tend towards $\mathcal{R}(Q)$.

We finish the proof by expressing $Q \succeq \mathbf{\Phi}^\top D\mathbf{\Phi}$ as a Schur complement. This yields

$$\begin{bmatrix} Q-\eta I & \Phi^\top D^{\frac{1}{2}} \\ D^{\frac{1}{2}}\Phi & \alpha I \end{bmatrix} \succeq 0\,. \tag{42}$$

Combining (40) and (42) yields (21), which concludes the proof.

G. Proof of Lemma 7

Lemma 7 is a direct consequence of applying Proposition 2 to the definition (4). Indeed, with $G_{J+1} = 0$ and $g_{J+1} = \tau$, one can rewrite (14) as (15) by setting $a_j = \gamma^{-1} G_j \mathbf{\Phi}, b_j = \gamma^{-1}(g_j - \tau + \gamma\tau)$. This means that (14) is equivalent to

$$\begin{aligned}
&\inf_{\rho\geq 0,\kappa_{ij}\geq 0} \rho\epsilon + \frac{1}{N}\sum_{i\in[N]} s^{(i)} \leq 0\,,\\
&\text{subject to}\,, \forall i\in[N]\,, \forall j\in[J+1]:\\
&s^{(i)} \geq \frac{1}{\gamma}(g_j - \tau + \gamma\tau) - \frac{1}{\gamma} G_j^\top \mathbf{\Phi}\boldsymbol{\xi}_T^{(i)} + \big(H\boldsymbol{\xi}_T^{(i)} + h\big)^\top \kappa_{ij}\\
&\qquad + \frac{\|H^\top\kappa_{ij}\|_2^2}{4\rho} - \frac{1}{2\rho\gamma} G_j^\top \mathbf{\Phi} H^\top \kappa_{ij} + \frac{\|\mathbf{\Phi}^\top G_j\|_2^2}{4\rho\gamma^2}.
\end{aligned}$$

One can factorize the last three terms of the constraint and do the change of variable $\zeta^{(i)} = s^{(i)} + \gamma^{-1}\tau - \tau$, which gives

$$\inf_{\rho\geq 0,\kappa_{ij}\geq 0} \rho\epsilon - \frac{1}{\gamma}\tau + \tau + \frac{1}{N}\sum_{i\in[N]} \zeta^{(i)} \leq 0\,, \tag{43a}$$

$$\text{subject to}\,, \forall i\in[N]\,, \forall j\in[J+1]:$$

$$\begin{aligned}
\zeta^{(i)} \geq\ & \frac{1}{\gamma} g_j - \frac{1}{\gamma} G_j^\top \mathbf{\Phi}\boldsymbol{\xi}_T^{(i)} + \big(H\boldsymbol{\xi}_T^{(i)} + h\big)^\top \kappa_{ij} \qquad (43\text{b})\\
&+ \frac{1}{4\rho\gamma^2}(\mathbf{\Phi}^\top G_j - \gamma H^\top\kappa_{ij})^\top(\mathbf{\Phi}^\top G_j - \gamma H^\top\kappa_{ij}).
\end{aligned}$$

Finally, a zero upper-bound constraint on an infimum is equivalent to an existence constaint. Moreover, because $\rho \geq 0$, (43b) can be written as an LMI using Schur's complement, which concludes the proof.