Data-driven distributionally robust MPC for constrained stochastic systems

Peter Coppens and Panagiotis Patrinos†

Abstract

In this paper we introduce a novel approach to distributionally robust optimal control that supports online learning of the ambiguity set, while guaranteeing recursive feasibility. We introduce conic representable risk, which is useful to derive tractable reformulations of distributionally robust optimization problems. Specifically, to illustrate the techniques introduced, we utilize risk measures constructed based on data-driven ambiguity sets, constraining the second moment of the random disturbance. In the optimal control setting, such moment-based risk measures lead to tractable optimal controllers when combined with affine disturbance feedback. Assumptions on the constraints are given that guarantee recursive feasibility. The resulting control scheme acts as a robust controller when little data is available and converges to the certainty equivalent controller when a large sample count implies high confidence in the estimated second moment. This is illustrated in a numerical experiment.

I. Introduction

DISTRIBUTIONALLY *robust optimization (DRO)* has gained traction recently as a technique that balances robustness with performance in an intuitive fashion. From a theoretical point of view such techniques act as regularizers [1] and in a data-driven setting, DRO acts at the interface between stochastic and robust optimization [2].

In the control community the potential of such techniques has not gone unnoticed [3], [4]. Here one would ideally solve stochastic optimal control problems like

$$\begin{aligned} \underset{\boldsymbol{\pi}\in\boldsymbol{\Pi}}{\text{minimize}} \quad & \mathbb{E}\left[\textstyle\sum_{t=0}^{N-1} \ell_t(x_t, \pi_t(w_0,\ldots,w_{t-1})) + \ell_N(x_N)\right] \\ \text{subj. to} \quad & x_{t+1} = f(x_t, \pi_t(w_0,\ldots,w_{t-1}), w_t),\ t \in \mathbb{N}_{0:N-1} \\ & \mathbb{P}[\phi(x_t) \leq 0] \geq 1-\varepsilon,\ t \in \mathbb{N}_{1:N-1} \\ & \psi(x_N) \leq 0 \text{ a.s., } x_0 \text{ given,} \end{aligned}$$

where $x_t \in \mathbb{R}^{n_x}$ denotes the state. Parametrized, causal policies $\boldsymbol{\pi} \in \boldsymbol{\Pi}$ map disturbances to inputs. That is, an element π_t of the sequence $\boldsymbol{\pi} = \{\pi_t\}_{t=0}^{N-1}$, maps $\{w_i\}_{i=0}^{t-1}$ to inputs in $\mathbb{R}^{n_u}$ for $t \geq 1$ and $\pi_0 \in \mathbb{R}^{n_u}$. Here, the disturbances $w_t \in \mathbb{R}^{n_w}$ are i.i.d. random vectors, the distribution of which is unknown, usually introducing the need for robust

†P. Coppens and P. Patrinos are with the Department of Electrical Engineering (ESAT-STADIUS), KU Leuven, Kasteelpark Arenberg 10, 3001 Leuven, Belgium. Email: peter.coppens@kuleuven.be, panos.patrinos@kuleuven.be

This work was supported by: the Research Foundation Flanders (FWO) PhD grant 11E5520N and research projects G0A0920N, G086518N and G086318N; Research Council KU Leuven C1 project No. C14/18/068; Fonds de la Recherche Scientifique – FNRS and the FWO – Vlaanderen under EOS project no 30468160 (SeLMA); EU's Horizon 2020 research and innovation programme: Marie Skłodowska-Curie grant No. 953348.

approaches. DRO then improves upon classical robust control by using available data to infer properties of the distribution, while retaining guarantees.

The core construct in DRO is the *ambiguity set*, a set of distributions against which one should robustify. Several ambiguity sets have been examined with varying success. The most common are moment-based, ϕ-divergence and Wasserstein ambiguity sets [5]. Such ambiguity sets are connected to so-called *risk measures* by duality. Hence this approach is directly related to risk-averse optimization [6].

Throughout the paper we rely on *conic representable ambiguity and risk* to derive tractable problems, similar to the methodology presented in [7], [8]. The main contributions are then as follows: *(i)* we derive tight, data-driven, moment-based ambiguity sets that are conic representable and shrink when more data becomes available; *(ii)* we extend conic risks to the multi-stage setting and use them to model *average value-at-risk* constraints; *(iii)* we synthesize the controller such that it is recursive feasible when it is applied in a receding horizon fashion; *(iv)* we illustrate how our framework leads to tractable controllers based on affine disturbance feedback policies [9], which are evaluated in numerical experiments.

Similar results were achieved in [10] for a tube-based approach with Wasserstein ambiguity, the radius of which is not data-driven; in [11] for relaxed, robust constraints; in [4] for moment-based ambiguity which is not data-driven and does not guarantee recursive feasibility; and [12], [13] for discrete distributions. DR control of Markov decision processes with finite state-spaces was also considered in [14]. Our framework supports online learning of truly data-driven ambiguity sets and risk constraints within a continuous state-space, while guaranteeing recursive feasibility.

This section continues with notation and preliminaries. Next §II introduces conic and data-driven ambiguity in a single-stage setting and §III introduces multi-stage extensions as well as the optimal control problem that we want to solve. Then §IV shows how to construct a controller such that recursive feasibility is guaranteed. Finally §V illustrates how our techniques lead to tractable controllers and contains numerical experiments.

A. Notation and preliminaries

Let $\mathbb{N}$ denote the integers and $\mathbb{R}$ ($\mathbb{R}_+$) the (nonnegative) reals. We denote by $\mathbb{S}^d$ the symmetric d by d matrices and by $\mathbb{S}^d_{++}$ ($\mathbb{S}^d_+$) the positive (semi)-definite matrices. For two matrices of compatible dimensions X, Y we use $[X; Y]$ ($[X, Y]$) for vertical (horizontal) concatenation. We use $\|\cdot\|_2$ to denote the spectral norm (Euclidean norm) for matrices (vectors) and $[\cdot]_+ := \max(0, \cdot)$. For matrices (or vectors) X, Y and cone $\mathcal{K}$, let $X \preccurlyeq_{\mathcal{K}} Y$ ($X \succcurlyeq_{\mathcal{K}} Y$) be $Y - X \in \mathcal{K}$ ($X - Y \in \mathcal{K}$). When $\mathcal{K} = \mathbb{S}^d_+$ we use $\preceq$ ($\succeq$).

Meanwhile, $(X, Y) := [\mathrm{vec}(X); \mathrm{vec}(Y)]$ interprets X, Y as column vectors in vertical concatenation. Let $\mathrm{diag}(X, Y)$ be a (block) diagonal matrix and let $I_d \in \mathbb{S}^d$ denote the identity. For a vector $x \in \mathbb{R}^d$, $[x]_i$ denotes the i'th element.

Slice notation: We introduce $\mathbb{N}_{a:b} = \{a, \ldots, b\}$. Similarly we use $\boldsymbol{w}_{a:b}$ to denote the sequence $\{w_i\}_{i \in \mathbb{N}_{a:b}}$. For a sequence of length N, index a (b) is omitted when 0 ($N-1$) is implied (when both are omitted we write $\boldsymbol{w}$).

Interpreting $\boldsymbol{w}_{0:N-1}$ with $w_i \in \mathbb{R}^{n_w}$ as an element of $\mathbb{R}^{Nn_w}$, consider affine maps $\boldsymbol{x}_{0:M-1} = \boldsymbol{A}\boldsymbol{w}_{0:N-1} + \boldsymbol{a}_{0:M-1}$ ($\boldsymbol{A} \in \mathbb{R}^{Mn_x \times Nn_w}$). Introducing homogeneous coordinates $\overline{\boldsymbol{w}} = (\boldsymbol{w}, 1)$ gives $\boldsymbol{x} = \overline{\boldsymbol{A}}\overline{\boldsymbol{w}}$, with $\overline{\boldsymbol{A}} = [\boldsymbol{A}, \boldsymbol{a}]$.

For a matrix $\boldsymbol{A}$ acting on sequences, the slice $\boldsymbol{A}_{i:j,k:\ell}$ describes the part mapping $\boldsymbol{w}_{k:\ell}$ to $\boldsymbol{x}_{i:j}$. So we take block rows and block columns, with blocks of size $\mathbb{R}^{n_x \times n_w}$.

Risk measures and ambiguity: Given some measurable space $(\mathcal{W}, \mathcal{B})$ with $\mathcal{W}$ a compact subset of $\mathbb{R}^{n_w}$ and $\mathcal{B}$ the associated Borel sigma-algebra, we use $\mathfrak{M}_+(\mathcal{W})$ ($\mathfrak{M}(\mathcal{W})$) to denote the space of finite (signed) measures on $(\mathcal{W}, \mathcal{B})$, making the dependency on $\mathcal{W}$ explicit. Similarly, let $\mathfrak{P}(\mathcal{W})$ denote the space of probability measures.

We also consider the space $\mathcal{Z} := \mathcal{C}(\mathcal{W})$ of continuous (bounded) $\mathcal{B}$-measurable functions $z\colon \mathcal{W} \to \mathbb{R}$. Elements of $\mathcal{Z}$ act as random loss functions. Notation $z \sim \mu$ means z has distribution $\mu \in \mathfrak{P}(\mathcal{W})$. The space $\mathfrak{M}(\mathcal{W})$ and $\mathcal{Z}$ are paired by the bilinear form [6, §2.2], for $z \in \mathcal{Z}$, $\mu \in \mathfrak{M}(\mathcal{W})$,

$$\langle z, \mu \rangle := \int_{\mathcal{W}} z(w) \mathrm{d}\mu(w).$$

We endow $\mathfrak{M}(\mathcal{W})$ with the weak* topology.

We write $z \succcurlyeq 0$ ($\mu \succcurlyeq 0$) to imply $z(w) \geq 0$ ($\mu(w) \geq 0$), $\forall w \in \mathcal{W}$. Note that, since $\mathcal{Z}$ and $\mathfrak{M}(\mathcal{W})$ are linear spaces, we can use the usual notation of linear operators (e.g., let $E\colon \mathfrak{M}(\mathcal{W}) \to \mathbb{R}^n$, then $E\mu = (\langle \epsilon_0, \mu \rangle, \dots \langle \epsilon_{n-1}, \mu \rangle)$ for some random variables $\epsilon_i \in \mathcal{Z}$, $i \in \mathbb{N}_{0:n-1}$). For each linear operator E we have the adjoint $E^*\colon \mathbb{R}^n \to \mathcal{Z}$, with $E^*\lambda := (\epsilon_0, \dots, \epsilon_{n-1}) \cdot \lambda$, where $\cdot$ is the usual inner product between vectors. After all,

$$\begin{aligned} E\mu \cdot \lambda &= (\langle \epsilon_0, \mu \rangle, \dots \langle \epsilon_{n-1}, \mu \rangle) \cdot \lambda \\ &= \left(\int_{\mathcal{W}} (\epsilon_0(w), \dots, \epsilon_{n-1}(w)) \mathrm{d}\mu(w) \right) \cdot \lambda \\ &= \int_{\mathcal{W}} ((\epsilon_0(w), \dots, \epsilon_{n-1}(w)) \cdot \lambda) \, \mathrm{d}\mu(w) \\ &= \langle (\epsilon_0, \dots, \epsilon_{n-1}) \cdot \lambda, \mu \rangle = \langle E^*\lambda, \mu \rangle. \end{aligned} \tag{1}$$

We define risk based on its ambiguity as in most DRO literature [6]. Specifically, we say that $\mathcal{A} \subseteq \mathfrak{M}_+(\mathcal{W})$ is an *ambiguity set* if it is a non-empty, closed and convex subset of $\mathfrak{P}(\mathcal{W})$. The associated risk measure $\rho_{\mathcal{A}}\colon \mathcal{Z} \to \mathbb{R}$ is then [6, §2]

$$\rho_{\mathcal{A}}[z] = \max_{\mu \in \mathcal{A}} \langle z, \mu \rangle = \max_{\mu \in \mathcal{A}} \mathbb{E}_\mu[z], \tag{2}$$

where $\mathbb{E}_\mu[\cdot]$ denotes the expected value w.r.t. $\mu \in \mathfrak{P}(\mathcal{W})$. and constitutes a mapping from random loss functions to the real line, which (similarly to expectation) can be used to deterministically compare random variables.

Our definition of an ambiguity set is directly related to that of coherent risk [15].

Lemma I.1. *Suppose that $\mathcal{A} \subseteq \mathfrak{P}(\mathcal{W})$ is non-empty, closed and convex. Then $\rho_{\mathcal{A}}$ in (2) is coherent. Specifically, $\forall z, z' \in \mathcal{Z}$ and $\alpha \in \mathbb{R}$, $\rho_{\mathcal{A}}$ is*

(i) convex, proper, and lower semi–continuous;

(ii) monotonous: $\rho_{\mathcal{A}}(z) \geq \rho_{\mathcal{A}}(z')$ if $z \succcurlyeq z'$;

(iii) translation equivariant: $\rho_{\mathcal{A}}(z + \alpha) = \rho_{\mathcal{A}}(z) + \alpha$;

(iv) positive homogeneous $\rho_{\mathcal{A}}(\alpha z) = \alpha \rho_{\mathcal{A}}(z)$ if $\alpha > 0$.

Moreover, $\mathcal{A}$ is compact and equal to the domain of $\rho^*_{\mathcal{A}}$ and $\rho_{\mathcal{A}}[z]$ is finite, where $\rho^*_{\mathcal{A}}$ denotes the convex conjugate.

Proof. Let $\chi_{\mathcal{A}}$ denote the indicator function of $\mathcal{A}$ (i.e. $\chi_{\mathcal{A}}[\mu] = +\infty$ if $\mu \notin \mathcal{A}$ and 0 otherwise). Then, by (2),

$$\rho[z] = \chi^* = \sup_{\mu \in \mathfrak{M}(\mathcal{W})} \{\langle z, \mu \rangle - \chi[\mu]\}, \tag{3}$$

where we omit the subscript of $\rho_{\mathcal{A}}$ and $\chi_{\mathcal{A}}$ for convenience. Since χ is an indicator function, it is convex ($\mathcal{A}$ is convex); lower semi–continuous ($\mathcal{A}$ is closed); and proper ($\mathcal{A}$ is nonempty). Therefore, by [16, Prop.2.112] and (3), ρ is proper, convex and lower semi–continuous (its epigraph is an intersection of closed halfspaces). Therefore (i) holds.

Since χ is convex and lower semi–continuous, we apply [16, Thm. 2.133] to show $\chi = \chi^{**} = \rho^*$, where the second equality follows by (3). Hence the domain of ρ^* is $\mathcal{A}$.

Compactness of $\mathfrak{P}(\mathcal{W})$ follows by Prohorov's theorem [17, p.13]. Since $\mathcal{A}$ is a closed subset of $\mathfrak{P}(\mathcal{W})$ it is also compact.

The results (ii)–(iv) follow directly from [15, Thm. 2.2]. Specifically (ii) from $\mathcal{A} \subset \mathfrak{M}_+(\mathcal{W})$, (iii) from $\mu(\mathcal{W}) = 1$ for all $\mu \in \mathcal{A}$ and (iv) from (2).

Next, we show that for any $z \in \mathcal{Z}$,

$$\langle z, \mu \rangle \leq \alpha, \, \forall \mu \in \mathfrak{P}(\mathcal{W}) \quad \Leftrightarrow \quad z \preccurlyeq \alpha, \tag{4}$$

where the inequality on the right holds pointwise over $\mathcal{W}$ (i.e. $z(w) \leq \alpha, \forall w \in \mathcal{W}$). The argument for (4) is as follows [17, Eq. 3.7]. Since $\mu(\mathcal{W}) = 1$, $\langle z, \mu \rangle \leq \alpha$ iff $\langle \alpha - z, \mu \rangle \geq 0$, which holds if $\alpha - z \succcurlyeq 0$ (since $\mu \succcurlyeq 0$). For the converse note $\delta_w \in \mathfrak{P}(\mathcal{W})$ for any $w \in \mathcal{W}$ with δ_w a dirac measure. So $\langle \alpha - z, \delta_w \rangle = \alpha - z(w) \geq 0$ for any $w \in \mathcal{W}$ is a necessary condition. So we have shown (4)

From (4) we can conclude $\langle z, \mu \rangle \leq \sup_{w \in \mathcal{W}} z(w)$, $\forall z \in \mathcal{Z}, \mu \in \mathcal{A}$. Hence, by (2), $\rho[z] \leq \sup_{w \in \mathcal{W}} z(w)$. Since $z(w)$ is finite for any $w \in \mathcal{W}$, $\rho[z]$ is finite (cf. [6, §2.2]). □

II. Single-stage problems

Given the dual formulation of a risk measure in (2), it is clear that the choice of $\mathcal{A}$ is a critical design decision. In this section we introduce how ambiguity sets, using moment information, are derived from data. We also introduce *conic representable risk*, used to derive tractable problems.

A. Data-driven risk

In DRO the reasoning is usually as follows. Consider a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and the optimization problem

$$\underset{u \in \mathcal{U}}{\text{minimize}} \qquad \mathbb{E}_{\mu_\star}[f(u, w)], \tag{5}$$

with $u \in \mathbb{R}^{n_u}$ some decision variable, f some loss function and $w \colon \Omega \to \mathcal{W}$ a random variable with $\mathcal{W} \subset \mathbb{R}^{n_w}$ the (compact) support of w. The main difficulty in solving the stochastic optimization problem (5) is that the *distribution* (or *push-forward measure*), $\mu_\star \in \mathfrak{P}(\mathcal{W})$ defined on the sample space $(\mathcal{W}, \mathcal{B})$ as $\mu_\star(O) = \mathbb{P}[w^{-1}(O)]$ for all $O \in \mathcal{B}$ and with $w^{-1}(O)$ the pre–image of O, is unknown.

Hence, instead one introduces an ambiguity set $\mathcal{A} \subseteq \mathfrak{P}(\mathcal{W})$, which contains $\mu_\star$ with some confidence. To do so one can estimate some statistic θ based on data. In the case of ϕ-divergence [12] and Wasserstein ambiguity [18], this θ is the empirical distribution,

while for moment-based ambiguity, θ encapsulates moment information. We will consider this final case in §II-B. To summarize:

Definition II.1. *Consider random variable $w\colon \Omega \to \mathcal{W}$ with distribution $\mu_\star$ and i.i.d. samples $\boldsymbol{w}_{0:M-1}\colon \Omega \to \mathcal{W}^M$. Let $\theta\colon \mathcal{W}^M \to \Theta$ denote a statistic for a set Θ and let $\beta \in \mathbb{R}$ be some radius*[1]*. Then a* data-driven ambiguity *$\mathcal{A}\colon \Theta \times \mathbb{R} \rightrightarrows \mathfrak{P}(\mathcal{W})$ with confidence $\delta \in (0,1)$ maps $(\theta(\boldsymbol{w}), \beta)$ to an ambiguity set $\mathcal{A}_\beta(\theta(\boldsymbol{w})) \subseteq \mathfrak{P}(\mathcal{W})$ such that*

$$\mathbb{P}[\mu_\star \in \mathcal{A}_\beta(\theta(\boldsymbol{w}_{0:M-1}))] \geq 1 - \delta. \tag{6}$$

In [12] this is referred to as a *learning system*.

Based on (6) we minimize $\rho_{\mathcal{A}_\beta(\hat{\theta})}[f(u,w)]$ instead of (5). The result upper bounds (5) with probability at least $1-\delta$.

B. Moment-based ambiguity

As mentioned before, we focus on the case where θ encapsulates moment information. Such ambiguity sets have the advantage that [18] *(i)* they can contain measures with support not limited to the observed samples (unlike most ϕ-divergence based sets); *(ii)* the radius is estimated with reasonable accuracy based on known information of the distribution (unlike for Wasserstein-based sets); and *(iii)* problem complexity does not grow with the sample count.

To ensure that an ambiguity set satisfying (6) can be derived, we assume that $\mathcal{W}$ is bounded, which is often the case in control applications and is therefore the usual assumption in robust control. Other common choices are that w is multivariate Gaussian or that it satisfies some concentration properties (e.g., sub-Gaussian) [19]. We have:

Lemma II.2. *Let $\mathcal{W} = \{w \in \mathbb{R}^{n_w} \colon \|w\|_2 \leq r\}$ and $R_w = \operatorname{diag}(I_{n_w}, cr)$ with $c \in \mathbb{R}$. Assume we have a set of i.i.d. samples $\boldsymbol{w}_{0:M-1}$ of $w \sim \mu_\star$ and let $\hat{C} := \sum_{i=0}^{M-1} \overline{w}_i \overline{w}_i^\top / M$. Then,*

$$\mathcal{A}_\beta(\hat{C}) = \left\{ \mu \in \mathfrak{P}(\mathcal{W}) \colon \left\| R_w(\hat{C} - \mathbb{E}_\mu[\overline{w}\,\overline{w}^\top]) R_w^\top \right\|_2 \leq \beta \right\},$$

satisfies (6) when $\beta = 0.5r^2(1+\sqrt{1+16c^2})\sqrt{2\log(2(n_w+1)/\delta)/M}$.

Proof. We use a matrix Hoeffding bound [20, Thm. 1.3] with improved constants. See App. A for the full proof. □

Remark II.3. In the numerical experiments towards the end of the paper we select $c = 1/4$. This choice results in a relatively simple expression for the radius

$$\beta = 0.5(1+\sqrt{2})r^2\sqrt{2\log(2(n_w+1)/\delta)/M}$$

and performed well in experiments.

[1] Some moment-based ambiguity set can have multiple radii (cf. [2]).

C. Conic-representable ambiguity

We introduce *conic representable ambiguity* (similar to the framework in [7], [8]) below and show how such risk is related to robust optimization through conic duality.

Definition II.4. *Consider a compact sample space $\mathcal{W} \subset \mathbb{R}^{n_w}$ and $\mathcal{Z} = C(\mathcal{W})$. An ambiguity set $\mathcal{A}$ is* conic representable *if, for some $E, F\colon \mathfrak{M}(\mathcal{W}) \to \mathbb{R}^{n_b}$ and $b \in \mathbb{R}^{n_b}$,*

$$\mathcal{A} = \{\mu \in \mathfrak{P}(\mathcal{W})\colon \exists \nu \in \mathfrak{M}_+(\mathcal{W}), E\mu + F\nu \preccurlyeq_{\mathcal{K}} b\}\,,$$

with ν some auxiliary measure and $\mathcal{K}$ a closed, convex cone. Usually we assume $F = 0$. When $F \neq 0$ we refer to the ambiguity as ν-conic representable*. Similarly we refer to $\rho_{\mathcal{A}}$, as in (2), as (ν-)conic representable risk (conic for short).*

The parameters of $\mathcal{A}$ should be selected such that it is an ambiguity set (i.e., a nonempty, closed and convex subset of $\mathfrak{P}(\mathcal{W})$). Since we usually want an $\mathcal{A}$ satisfying (6), it will be non-empty as it should at least contain the true distribution. The random variables used to construct E and F, are all continuous. Therefore [21, Thm. 15.5] E and F are continuous mappings. Thus $\mathcal{A}$ is the intersection between the closed set $\mathfrak{P}(\mathcal{W})$ and the pre–image of a closed set under a continuous mapping, which is also closed. Hence $\mathcal{A}$ is a closed subset of $\mathfrak{P}(\mathcal{W})$. Convexity of $\mathcal{A}$ then follows, since E and F are linear and $\mathcal{K}$ is convex.

In [8] it was shown that both the *average* and *entropic* value-at-risk are conic whenever $\mathcal{W}$ is finite. Many more risks fall under this framework [7], [22].

Direct application of conic linear duality [17] gives:

Lemma II.5. *A risk $\rho_{\mathcal{A}}[z]$ as in Def. II.4 is equal to the optimal value of*

$$\begin{array}{ll} \underset{\lambda \succcurlyeq_{\mathcal{K}^*} 0, \tau}{\text{minimize}} & \tau + b \cdot \lambda \\ \text{subj. to} & E^*\lambda + \tau \succcurlyeq z,\ F^*\lambda \succcurlyeq 0, \end{array} \tag{D}$$

where the functional inequalities should hold pointwise for all $w \in \mathcal{W}$, E^* and F^* denote the adjoint operators (cf. §I-A), and $\mathcal{K}^*$ the dual cone.

Proof. By (2) the primal problem is

$$\begin{array}{ll} \underset{\mu, \nu \in \mathfrak{M}_+(\mathcal{W})}{\text{maximize}} & \langle z, \mu \rangle \\ \text{subj. to} & E\mu + F\nu \preccurlyeq_{\mathcal{K}} b \\ & \langle 1, \mu \rangle = 1, \end{array} \tag{P}$$

with $\operatorname{val}(P) = \rho[z]$ (where we omit the subscript for convenience). We refer to the minimization problem (D) as the dual problem. Let $\tau \in \mathbb{R}$ and $\lambda \in \mathbb{R}^{n_b}$. Then the Lagrangian is

$$\begin{aligned} \varphi[\mu, \nu, \lambda] &:= \langle z, \mu \rangle + (1 - \langle 1, \mu \rangle) \cdot \tau + (b - E\mu - F\nu) \cdot \lambda \\ &= \tau + b \cdot \lambda - \langle \tau + E^*\lambda - z, \mu \rangle - \langle F^*\lambda, \nu \rangle, \end{aligned}$$

where we can use (1) to construct the adjoints. We have

$$\mathcal{K}^* := \{\lambda \in \mathbb{R}^{n_b}\colon \lambda^* \cdot \lambda \geq 0, \forall \lambda^* \in \mathcal{K}\}.$$

Hence $\max_{\nu\in\mathfrak{M}_+(\mathcal{W})}\min_{\lambda\in\mathcal{K}^*,\tau}\{(b-E\mu-F\nu)\cdot\lambda\}=-\chi[\mu]$, where χ is the indicator of $\mathcal{A}$. Therefore

$$\max_{\mu,\nu\in\mathfrak{M}_+(\mathcal{W})}\min_{\lambda\in\mathcal{K}^*,\tau}\{\varphi[\mu,\nu,\lambda]\}=\max_{\mu\in\mathfrak{M}_+(\mathcal{W})}\{\langle z,\mu\rangle-\chi[\mu]\}=\rho[z].$$

Similarly note that [17, Eq. 3.7]

$$\begin{aligned}\mathfrak{M}_+^*(\mathcal{W}) &:= \{z\in\mathcal{Z}\colon \langle z,\mu\rangle\geq 0,\forall\mu\in\mathfrak{M}_+(\mathcal{W})\}\\ &= \{z\in\mathcal{Z}\colon z(w)\geq 0,\forall w\in\mathcal{W}\}=\{z\in\mathcal{Z}\colon z\succcurlyeq 0\},\end{aligned}$$

which follows from a similar argument as (4). As such

$$\min_{\lambda\in\mathcal{K}^*,\tau}\max_{\mu,\nu\in\mathfrak{M}_+(\mathcal{W})}\{\varphi[\mu,\nu,\lambda]\}=\operatorname{val}(D),$$

since λ gives a finite cost iff $\tau+E^*\lambda-z\in\mathfrak{M}_+^*(\mathcal{W})$ and $F^*\lambda\in\mathfrak{M}_+^*(\mathcal{W})$ (i.e. $\tau+E^*\lambda-z\succcurlyeq 0$ and $F^*\lambda\succcurlyeq 0$).

All that is left is to show strong duality (i.e. $\operatorname{val}(D)=\operatorname{val}(P)=\rho[z]$). This follows directly from coherence of ρ (specifically ρ being proper, implying consistency of (P)), compactness of $\mathcal{W}$ and [17, Cor. 3.1]. □

Note that constraints in the dual are robust constraints, since they hold for all $w\in\mathcal{W}$. Hence, techniques from robust optimization enable finding tractable reformulations.

Example II.6. The ambiguity set $\mathcal{A}_\beta(\hat{C})$ of Lem. II.2 is conic with $n_b=3n_w^2$, $E\mu=(\pm\langle R_w\overline{w}\,\overline{w}^\top R_w^\top,\mu\rangle)$, $b=(R_w\hat{C}R_w^\top\pm\beta I)$ and $\mathcal{K}=\mathbb{S}_+^{n_w+1}\times\mathbb{S}_+^{n_w+1}$.

Moreover, letting $\lambda=(\Lambda,\mathrm{V})$ with $\Lambda,\mathrm{V}\in\mathbb{S}^{n_w+1}$ and $\tau\in\mathbb{R}$ while using Lem. II.5, means

$$\begin{aligned}\rho_{\mathcal{A}_\beta(\hat{C})}[z]=\min_{\Lambda,\mathrm{V}\succeq 0,\tau}\quad & \tau+\operatorname{Tr}[\Lambda(R_w\hat{C}R_w^\top+\beta I)]+\operatorname{Tr}[\mathrm{V}(R_w\hat{C}R_w^\top-\beta I)]\\ \text{s.t.}\quad & \tau+E^*\lambda\succcurlyeq z,\end{aligned}$$

where the adjoint $E^*\colon\mathbb{R}^{n_b}\to\mathcal{Z}$, is (cf. (1))

$$\begin{aligned}(\tau+E^*\lambda)(w)&=\tau+\operatorname{Tr}[R_w\overline{w}\,\overline{w}^\top R_w^\top(\Lambda-\mathrm{V})]\\ &=\overline{w}^\top R_w^\top(\Lambda-\mathrm{V})R_w\overline{w}+\tau.\end{aligned}$$

If the constraint $\tau+E^*\lambda\succcurlyeq z$ is LMI representable, then $\rho_{\mathcal{A}_\beta(\hat{C})}[z]$ can be evaluated by solving a SDP. For example if $z=\overline{w}^\top P\overline{w}$. Then, since $w^\top w\leq r^2$, we can apply the S-Lemma [23, Thm. B.2.1.] to show that $\tau+E^*\lambda\succcurlyeq z$ iff.,

$$\exists s\geq 0,\ R_w^\top(\Lambda-\mathrm{V})R_w+\operatorname{diag}(sI,\tau-sr^2)-P\succeq 0. \tag{7}$$

We also consider ambiguity with only support constraints.

Example II.7. Ambiguity $\mathfrak{P}(\mathcal{W})$ is conic representable with $n_b=0$. Hence, $\rho_{\mathfrak{P}(\mathcal{W})}[z]=\min_\tau\{\tau\colon\tau\succcurlyeq z\}$, corresponds to $\rho_{\mathfrak{P}(\mathcal{W})}[z]=\max_{w\in\mathcal{W}}z(w)$ and only considers the support as is common in robust optimization.

III. Multi-stage problems

In this section we show how conic single-stage risk can be extended to a multi-stage setting, which is required to develop distributionally robust MPC controllers. Specifically, we will consider risk measures operating on the dynamics

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

with $x_t \in \mathbb{R}^{n_x}$ ($u_t \in \mathbb{R}^{n_u}$) the state (input) and $w_t \in \mathbb{R}^{n_w}$ the disturbance, which follows a random process. For $t \in \mathbb{N}_{0:N-1}$ we consider $\ell_t \colon \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \to \mathbb{R}_+$ a stage cost function, and $\ell_N \colon \mathbb{R}^{n_x} \to \mathbb{R}_+$ the terminal cost.

For each stage t, the trajectory up to that time $\boldsymbol{w}_{0:t-1}$ is an element of $\mathcal{W}^t$. For each $\mathcal{W}^t$, $\mathcal{B}^t$ is the accompanying Borel sigma-algebra, ($\mathfrak{M}^t$) $\mathfrak{M}^t_+$ the set of (signed) measures and $\mathfrak{P}^t$ the set of probability measures on $(\mathcal{W}^t, \mathcal{B}^t)$. For brevity we henceforth omit the explicit dependency on $\mathcal{W}^t$. Also consider the paired spaces of continuous functions $\mathcal{Z}_t = \mathcal{C}(\mathcal{W}^t)$.

We can then consider multistage ambiguity sets $\mathcal{A}^t$, which are nonempty, closed and convex subsets of $\mathfrak{P}^t$. These in turn define a multistage analog to risk measures[2], *multistage risk measures* [6, §4.2], $\rho_{\mathcal{A}^t} \colon \mathcal{Z}_t \to \mathbb{R}$. Since this is simply a usual risk measure, but defined on $\mathcal{Z}_t$, the properties of Lem. I.1 generalize. We specifically consider coherent multistage risk

$$\rho_{\mathcal{A}^t}[z_t] = \max_{\mu^t \in \mathcal{A}^t} \langle z_t, \mu^t \rangle = \max_{\mu^t \in \mathcal{A}^t} \mathbb{E}_{\mu^t}[z_t]. \tag{8}$$

Given such risks, the goal is to solve, for a given x_0,

$$\begin{aligned}
\underset{\boldsymbol{\pi} \in \boldsymbol{\Pi}}{\text{minimize}} \quad & \rho_{\mathcal{A}^N}\left[\sum_{t=0}^{N-1} \ell_t(x_t, \pi_t(\boldsymbol{w}_{:t-1})) + \ell_N(x_N)\right] && \text{(9a)}\\
\text{subj. to} \quad & x_{t+1} = f(x_t, \pi_t(\boldsymbol{w}_{:t-1}), w_t),\ t \in \mathbb{N}_{0:N-1} && \text{(9b)}\\
& \mathfrak{r}^t_{\mathcal{A}}[\phi(x_t)] \leq 0,\ t \in \mathbb{N}_{1:N-1} && \text{(9c)}\\
& \psi(x_N) \leq 0 \text{ a.s.}, && \text{(9d)}
\end{aligned}$$

where $\boldsymbol{\Pi}$ denotes a set of parametrized, continuous, causal policies. The risk constraints (9c) involve the multistage risk measures $\mathfrak{r}^t_{\mathcal{A}}$ and are discussed in detail in §IV. We illustrate how (9) interpolates between the robust setting and (1) in §V.

Remark III.1. Problem (9) is not exact as we optimize over parametrized policies (cf. §V), for tractability. As such, time-consistency [24] cannot be guaranteed (i.e. a policy computed at $t = 0$ may not be optimal at $t = 1$ after realization of w_0). Hence a receding horizon scheme is used.

A. Product ambiguity

To enforce independence of the disturbances w_t we introduce *product ambiguity* [6, §4.2]. For a sequence of single-stage ambiguity factors $\mathcal{A}_i$ for $i \in \mathbb{N}_{0:t-1}$, consider

$$\bigtimes_{i=0}^{t-1} \mathcal{A}_i = \mathcal{A}_0 \times \cdots \times \mathcal{A}_{t-1}, \tag{10}$$

[2]Multistage risk is often constructed using nested conditional risk measures. We avoid such a construction for conciseness and tractability. The consequences of this are discussed in [6, §4].

where some $\mu^t \in \mathcal{A}_0 \times \cdots \times \mathcal{A}_{t-1}$ if it is constructed as a product measure of some $\mu_i \in \mathcal{A}_i$, for $i \in \mathbb{N}_{0:t-1}$ (denoted by $\mu^t = \mu_0 \times \cdots \times \mu_{t-1}$). We show that in certain cases such ambiguities are conic representable.

Before doing so we need to extend linear operators $E_i \colon \mathfrak{M} \to \mathbb{R}^{n_b}$ to take arguments in $\mathfrak{M}^t$ in a natural way. To do so, note that for any $E_i \colon \mathfrak{M} \to \mathbb{R}^{n_b}$ and $\mu \in \mathfrak{M}$ we have $E_i\mu = \int_{w \in W} e_i(w) \mathrm{d}\mu(w)$ for some $e_i \colon \mathcal{W} \to \mathbb{R}^{n_b}$, by definition. Measures $\mu^t \in \mathfrak{M}^t$ take arguments $\boldsymbol{w}_{:t-1} = (w_0, \dots, w_{t-1})$, so we introduce $E|_i \colon \mathfrak{M}^t \to \mathbb{R}^{n_b}$ such that

$$E|_i\mu^t = \int_{\mathcal{W}^t} e_i(w_i) \mathrm{d}\mu^t(\boldsymbol{w}_{:t-1}), \ \forall \mu^t \in \mathfrak{M}^t. \tag{11}$$

With these new operators we have

Lemma III.2. *Let $\mathcal{A}_i$ be conic representable with parameters $E_i, b_i, \mathcal{K}_i$ for $i \in \mathbb{N}_{0:t-1}$. Then $\times_{i=0}^{t-1}\mathcal{A}_i$ is also conic representable with parameters*

$$E\mu^t = (E|_0\mu^t, \dots, E|_{t-1}\mu^t), \qquad b = (b_0, \dots, b_{t-1}),$$

and $\mathcal{K} = \mathcal{K}_0 \times \dots \mathcal{K}_{t-1}$. Moreover,

$$\rho_{\times_{i=0}^{t-1}\mathcal{A}_i}[z_t] = \min_{\lambda_i \succcurlyeq_{\mathcal{K}_i^*} 0, \tau} \left\{ \tau + \textstyle\sum_{i=0}^{t-1} b_i \cdot \lambda_i \colon \tau + \sum_{i=0}^{t-1} E|_i^* \lambda_i \succcurlyeq z_t \right\}.$$

Proof. Let $\mu^t = \mu_0 \times \cdots \times \mu_{t-1}$, with $\mu_i \in \mathfrak{P}_i$. Then, following the notation in (11),

$$\begin{aligned} E|_i\mu^t &:= \int_{\mathcal{W}^t} e_i(w_i) \mathrm{d}\mu^t(\mathrm{d}\boldsymbol{w}_{0:t-1}) \\ &\overset{(i)}{=} \int_{\mathcal{W}} \cdots \int_{\mathcal{W}} e_i(w_i) \mathrm{d}\mu_0(w_0) \dots \mathrm{d}\mu_{t-1}(w_{t-1}) = \int_{\mathcal{W}} e_i(w_i) \mathrm{d}\mu_i(w_i), \end{aligned}$$

where *(i)* follows from $\mu^t = \mu_0 \times \cdots \times \mu_{t-1}$ and $\mu^t \in \mathfrak{P}^t$. Hence $E|_i\mu \preccurlyeq_{\mathcal{K}_i} b_i$ iff $E_i\mu_i \preccurlyeq_{\mathcal{K}_i} b_i$. Repeating the same argument for each i proves that $\mathcal{A}^t := \times_{i=0}^{t-1} \mathcal{A}_i$ is conic representable. Since $\mathcal{A}_i$ are all non-empty, $\mathcal{A}^t$ is also nonempty. Convexity and closedness follow from the arguments below Def. II.4. The dual then follows from applying Lem. II.5. □

B. Risk constraints

Ideally constraints like (9c) would require the state to lie within some set almost surely. Since such a constraint in a stochastic setting can be very conservative, we will instead implement average value-at-risk constraints, for $\alpha \in (0, 1)$,

$$\mathrm{AV@R}_\alpha^\mu[z] := \inf_{\tau \in \mathbb{R}} \left\{ \tau + \alpha^{-1} \mathbb{E}_\mu[z - \tau]_+ \right\} \leq 0. \tag{12}$$

Such constraints *(i)* act as a convex relaxation of chance constraints [25]; *(ii)* penalize the expected violation in the α quantile where violations do occur. In control applications (12) is natural, since it penalizes large violations more.

To evaluate the expectation in (12), true knowledge about the distribution is needed. Hence, we will operate on the distributionally robust AV@R constraint instead:

$$\mathbf{r}\text{-}\mathrm{AV@R}_\alpha^{\mathcal{A}}[z] := \max_{\mu \in \mathcal{A}} \mathrm{AV@R}_\alpha^\mu[z] \leq 0, \tag{13}$$

with $\mathcal{A}$ the *core ambiguity*. If $\mathcal{A}$ satisfies (6), then (13) implies the chance constraint $\mathbb{P}[z_t \leq 0] \geq 1-\varepsilon$ holds with $1-\varepsilon \leq (1-\delta)(1-\alpha)$. Moreover, whenever $\mathcal{A}$ is conic, then robust AV@R is ν-conic.

Lemma III.3. *Let $\mathcal{A}$ be conic with parameters $E_{\mathbf{c}}, b_{\mathbf{c}}, \mathcal{K}_{\mathbf{c}}$. Then $\mathbf{r}$-AV@R$_\alpha^{\mathcal{A}}$ in (13) is ν-conic with*

$$E\mu = (E_{\mathbf{c}}\mu, \langle 1, \mu\rangle), \qquad F\nu = (E_{\mathbf{c}}\nu, \langle 1, \nu\rangle), \qquad b = (b_{\mathbf{c}}, 1)/\alpha,$$

and $\mathcal{K} = \mathcal{K}_{\mathbf{c}} \times \{0\}$. Moreover, $\mathbf{r}$-AV@R$_\alpha^\rho[z]$ equals

$$\min_{\lambda \succcurlyeq_{\mathcal{K}_{\mathbf{c}}^*} 0, \tau, \tau_{\mathbf{c}}} \left\{ \tau + \alpha^{-1}(\tau_{\mathbf{c}} + b_{\mathbf{c}} \cdot \lambda) \colon E_{\mathbf{c}}^*\lambda + \tau_{\mathbf{c}} \succcurlyeq 0,\ E_{\mathbf{c}}^*\lambda + \tau_{\mathbf{c}} + \tau \succcurlyeq z \right\}. \tag{14}$$

Proof. This proof generalizes the methodology of [26] to arbitrary conic representable risk. First note that

$$\max_{\mu \in \mathcal{A}} \inf_{\tau \in \mathbb{R}} \left\{ \tau + \alpha^{-1}\mathbb{E}_\mu[z-\tau]_+ \right\} = \inf_{\tau \in \mathbb{R}} \left\{ \tau + \alpha^{-1}\rho_{\mathcal{A}}[z-\tau]_+ \right\}, \tag{15}$$

by [27, Thm. 2.1]. Specifically let $\phi(\tau, w) = \tau + \alpha^{-1}[z(w) - \tau]_+$. Then we have *(i)* $\phi(\tau, \cdot) \in \mathcal{Z}$, implying that it is μ integrable and measurable; *(ii)* $\phi(\cdot, w)$ is convex for any $w \in \mathcal{W}$; *(iii)* $\rho_{\mathcal{A}}[z-\tau]_+$ is finite (Lem. I.1); *(iv)* the set $\mathcal{A} \subseteq \mathfrak{P}(\mathcal{W})$ is compact (Lem. I.1); *(v)* $\phi(\tau, \cdot)$ is continuous and hence bounded for any $\tau \in \mathbb{R}$ on $\mathcal{W}$. Under these properties as well as $\mathcal{A}$ being convex, [27, Thm. 2.1] states that strong duality holds, allowing us to exchange the inf and the max.

Applying Lem. II.5 to $\rho_{\mathcal{A}}$ on the r.h.s., results in (14) (where $[\cdot]_+$ produces two separate constraints and $\tau_{\mathbf{c}}$, $\lambda_{\mathbf{c}}$ act as the Lagrangian multipliers for the constraint $\mu_{\mathbf{c}} \in \mathcal{A}_{\mathbf{c}}$). Again applying Lem. II.5 gives the original ν-conic representation.

The second application of Lem. II.5 requires the resulting set of measures (denoted $\mathcal{A}_\alpha$ below) to be a nonempty, closed and convex subset of $\mathfrak{P}(\mathcal{W})$. By construction we already have $\mathcal{A}_\alpha \subseteq \mathfrak{P}(\mathcal{W})$. Next, we show that $\mathcal{A}_\alpha$ is larger than $\mathcal{A}$. After all, for any $\mu_{\mathbf{c}} \in \mathcal{A}$, take $\mu = \mu_{\mathbf{c}}$ and $\nu = \alpha^{-1}(1-\alpha)\mu_{\mathbf{c}} \succcurlyeq 0$, since $\alpha \in (0,1)$. Moreover, since $\mathcal{K}_{\mathbf{c}}$ is a cone,

$$E_{\mathbf{c}}\mu + E_{\mathbf{c}}\nu \preccurlyeq_{\mathcal{K}_{\mathbf{c}}} b_{\mathbf{c}}/\alpha \quad \Leftrightarrow \quad E_{\mathbf{c}}(\alpha\mu + \alpha\nu) \preccurlyeq_{\mathcal{K}_{\mathbf{c}}} b_{\mathbf{c}},$$

with $\alpha\mu + \alpha\nu = \alpha\mu_{\mathbf{c}} + (1-\alpha)\mu_{\mathbf{c}} = \mu_{\mathbf{c}}$. For the same reason we have $\langle 1, \mu\rangle + \langle 1, \nu\rangle = \alpha^{-1}\langle 1, \mu_{\mathbf{c}}\rangle = \alpha^{-1}$. Therefore $\mu_{\mathbf{c}} \in \mathcal{A}_\alpha$ for each $\mu_{\mathbf{c}} \in \mathcal{A}$. Hence, since $\mathcal{A}$ is nonempty, so is $\mathcal{A}_\alpha$. Closedness and convexity then follow by the arguments below Def. II.4. So using Lem. II.5 is justified. □

IV. Recursive feasibility

We show how one can configure the constraints of (9) such that recursive feasibility is ensured. To do so we assume

(A1) $\imath_{\mathcal{A}}^t[z_t] := \mathbf{r}\text{-AV@R}_\alpha^{\mathfrak{P}^{t-1} \times \mathcal{A}}[z_t]$, $\forall t \in \mathbb{N}_{0:N-1}, z_t \in \mathcal{Z}^t$;

(A2) $\mathcal{A}$ is updated based on measurements as $g(\mathcal{A}, w)$ (e.g., following Lem. II.2) and $\mathcal{A}^+ := g(\mathcal{A}, w) \subseteq \mathcal{A}$ a.s.

We introduce the terminal set $\mathcal{X}_N := \{x \colon \psi(x) \leq 0\}$. Let $\mathcal{V}_N^{\mathcal{A}}(x_0)$ denote the minimum of (9) for some x_0 and let $\mathcal{D}_N(\mathcal{A})$ denote its domain. Then consider the set of feasible policies $\mathbf{\Pi}_N(x_0, \mathcal{A}) := \{\boldsymbol{\pi} \in \mathbf{\Pi} \colon (9b), (9c), (9d)\}$.

We begin with the following definition.

Definition IV.1 (Recursive Feasibility)**.** *Let* $x_0 \in \mathcal{D}_N(\mathcal{A})$ *and* $\boldsymbol{\pi} \in \boldsymbol{\Pi}_N(x_0, \mathcal{A})$*. If,* $f(x_0, \pi_0, w_0) \in \mathcal{D}_N(g(\mathcal{A}, w_0))$ *a.s., then* (9) *is* recursive feasible *(RF).*

We can then prove the following theorem:

Theorem IV.2. *Assume (A1), (A2) and that we are given some terminal policy* $\pi_f(x_N)$ *such that*

(A3) $\mathcal{X}_N \subseteq \{x \in \mathbb{R}^{n_x} : \phi(f(x, \pi_f(x), w)) \leq 0,\ \forall w \in \mathcal{W}\}$*;*

(A4) $\mathcal{X}_N$ *is robust positive invariant (RPI) for* π_f *(i.e.,* $f(x, \pi_f(x), w) \in \mathcal{X}_N$ *for each* $(x, w) \in \mathcal{X}_N \times \mathcal{W}$*);*

(A5) $\forall \boldsymbol{\pi}_{0:N-1} \in \boldsymbol{\Pi}$*, let* $\pi_N = \pi_f(x_N)$*, depending on* $\boldsymbol{w}_{:N}$ *through* x_N*. Then the shifted policy* $\boldsymbol{\pi}^+_{0:N-1} = \boldsymbol{\pi}_{1:N}(w_0, \cdot)$ *for any fixed* $w_0 \in \mathcal{W}$*, lies in* $\boldsymbol{\Pi}$*.*

Then, (9) *is recursive feasible.*

Proof. We will consider any fixed $w_0 \in \mathcal{W}$ and show that, given that (9) is feasible for some x_0, it will also be feasible for the next time step starting from $x_0^+ = f(x_0, \pi_0, w_0)$ (cf. Def. IV.1). Here, we consider the feasible policy $\boldsymbol{\pi}_{0:N-1} \in \boldsymbol{\Pi}_N(x_0, \mathcal{A})$, to which we append $\pi_N = \pi_f(x_N)$. Propagating the dynamics with this policy gives the sequence of states $\boldsymbol{x}_{0:N+1}$, depending on $\boldsymbol{w}_{:N}$ through (9b).

We then define the shifted sequence of states as

$$\boldsymbol{x}^+_{0:N}(\boldsymbol{w}_{1:N-1}) := (\boldsymbol{x}_1(w_0, \boldsymbol{w}_{1:N}), \dots, \boldsymbol{x}_{N+1}(w_0, \boldsymbol{w}_{1:N})),$$

where w_0 is considered fixed and $\boldsymbol{w}_{1:N}$ is left variable. We can analogously define the shifted policy $\boldsymbol{\pi}^+_{0:N-1}$ as

$$\boldsymbol{\pi}^+_{0:N-1}(\boldsymbol{w}_{1:N-1}) := (\boldsymbol{\pi}_1(w_0, \boldsymbol{w}_{1:N-1}), \dots, \boldsymbol{\pi}_f(x_N(w_0, \boldsymbol{w}_{1:N-1}))).$$

By construction, these shifted sequences satisfy (9b) and we can consider risk measures over (continuous) functions of these, where integration is performed over $\boldsymbol{w}_{1:N}$.

Using this coupling between the feasible problem and the shifted problem we show $\boldsymbol{\pi}^+ \in \boldsymbol{\Pi}_N(x_0, \mathcal{A}_+)$. That is, the candidate policy $\boldsymbol{\pi}^+$ is feasible for the shifted problem.

I: By (A5), $\boldsymbol{\pi}^+ \in \boldsymbol{\Pi}$;

II: We show that (9c) at t implies (9c) in the shifted problem at $t-1$. That is, $\mathfrak{r}^t_{\mathcal{A}}[\phi(x_t)] \geq \mathfrak{r}^{t-1}_{\mathcal{A}}[\phi(x^+_{t-1})]$ for any $w_0 \in \mathcal{W}$. So, letting $z = \phi(x_t)$, by (15),

$$\mathfrak{r}^t_{\mathcal{A}}[z] = \max_{\mu^t \in \mathfrak{P}^{t-1} \times \mathcal{A}} \inf_{\tau \in \mathbb{R}} \left\{\tau + \alpha^{-1} \mathbb{E}_{\mu^t}[z_\tau]\right\} = \inf_{\tau \in \mathbb{R}} \left\{\tau + \alpha^{-1} \rho_{\mathfrak{P}^{t-1} \times \mathcal{A}}[z_\tau]\right\},$$

with $z_\tau = [z - \tau]_+$. For $z^+ := \phi(x^+_{t-1}) = z(w_0, \boldsymbol{w}_{1:t-1})$ ($z^+_\tau = [z^+ - \tau]_+$), we replace $\rho_{\mathfrak{P}^{t-1} \times \mathcal{A}}[z_\tau]$ with $\rho_{\mathfrak{P}^{t-2} \times \mathcal{A}}[z^+_\tau]$ for $\mathfrak{r}^{t-1}_{\mathcal{A}}[z^+]$. Writing out $\rho_{\mathfrak{P}^{t-1} \times \mathcal{A}}$ gives

$$\begin{aligned}
&\max_{\mu^t \in \mathfrak{P}^{t-1} \times \mathcal{A}} \left\{ \int_{\mathcal{W}^t} z_\tau(\boldsymbol{w}_{:t-1}) \mathrm{d}\mu^t(\boldsymbol{w}_{:t-1}) \right\} \\
&\overset{(i)}{=} \max_{\mu_{t-1} \in \mathcal{A}} \max_{\mu^{t-1} \in \mathfrak{P}^{t-1}} \left\{ \int_{\mathcal{W}^{t-1}} \overbrace{\left(\int_{\mathcal{W}} z_\tau(\boldsymbol{w}_{:t-1}) \mathrm{d}\mu_{t-1}(w_{t-1}) \right)}^{h(\boldsymbol{w}_{:t-2})} \mathrm{d}\mu^{t-1}(\boldsymbol{w}_{:t-2}) \right\} \\
&\overset{(ii)}{=} \max_{\mu_{t-1} \in \mathcal{A}} \max_{\boldsymbol{w}_{:t-2} \in \mathcal{W}^{t-1}} h(\boldsymbol{w}_{:t-2}) \overset{(iii)}{\geq} \max_{\mu_{t-1} \in \mathcal{A}} \max_{\boldsymbol{w}_{1:t-2} \in \mathcal{W}^{t-2}} h(w_0, \boldsymbol{w}_{1:t-2}).
\end{aligned}$$

Noting that $\mu^t = \mu^{t-1} \times \mu_{t-1}$ with $\mu^{t-1} \in \mathfrak{P}^{t-1}$ and $\mu_{t-1} \in \mathcal{A}$, before splitting up the max and the integrals, gives *(i)*. The inner integral (i.e. $h(\boldsymbol{w}_{:t-2})$) then acts as a continuous random variable $\mathcal{W}^{t-1} \to \mathbb{R}$ for any fixed μ_{t-1} (cf. App. B). Hence we can apply the reasoning within Ex. II.7 to maximize over $\boldsymbol{w}_{:t-2} \in \mathcal{W}^{t-1}$ instead of over measures resulting in *(ii)*. It is clear that, fixing the value of w_0 results in the inequality *(iii)*. Reverting the steps *(i)* and *(ii)* to get a maximization over $\mu^{t-1} \in \mathfrak{P}^{t-2} \times \mathcal{A}$ shows that the final expression after *(iii)* equals $\rho^+_{\mathcal{A}^{t-1}}[z^+ - \tau]_+$. Hence $\rho_{\mathfrak{P}^{t-1} \times \mathcal{A}}[z - \tau]_+ \geq \rho_{\mathfrak{P}^{t-2} \times \mathcal{A}}[z^+ - \tau]_+$ for all $\tau \in \mathbb{R}$, $z \in \mathcal{Z}^t$ and $w_0 \in \mathcal{W}$. Therefore $\imath^t_{\mathcal{A}}[\phi(x_t)] \geq \imath^{t-1}_{\mathcal{A}}[\phi(x^+_{t-1})]$. Since $\mathcal{A}^+ \subseteq \mathcal{A}$ by (A2), $\imath^t_{\mathcal{A}}[\phi(x_t)] \geq \imath^{t-1}_{\mathcal{A}^+}[\phi(x^+_{t-1})]$. Hence (9c) holds for all $t \in \mathbb{N}_{1:N-2}$ in the shifted problem. For $t = N-1$ we rely on (A3) and (A4).

III: The terminal constraint (9d) follows directly from (A4).

We have thus shown that $\boldsymbol{\pi}^+$ is a feasible policy. □

Remark IV.3. Note that (A1) is essential since RF is a robust property, holding a.s. It acts as a convex relaxation of chance constraints conditioned on previous time steps (i.e., $\mathbb{P}[\phi(x_t) \leq 0 \mid x_{t-1}] \leq \varepsilon$ which holds a.s., hence $\forall \mu^{t-1} \in \mathfrak{P}^{t-1}$ by Ex. II.7). Due to the reduction of the policy space $\boldsymbol{\Pi}$ (cf. Rem. III.1) it is harder to satisfy such constraints for larger t. Other (less conservative) reformulations exist in the stochastic MPC literature, which impose all constraints at the first time step using a (maximal) RPI set (cf. [28]).

V. Affine disturbance feedback

To make the reformulations above more concrete, we show how (9) is solved. In general this is intractable, since we need to optimize over infinite dimensional policies $\boldsymbol{\pi}$, under robust constraints associated with the risks (cf. Rem. III.1). Hence, we use affine disturbance feedback. The resulting optimization problem is a SDP. Different ambiguity sets and policies would give other reformulations (e.g., [11], [8]).

Consider linear dynamics, quadratic losses and constraints:

$$\begin{aligned} f(x,u,w) &= Ax + Bu + Ew, & \pi_f(x) &= K_f x, \\ \ell_t(x,u) &= x^\top Q x + u^\top R u, & \ell_N(x) &= x^\top Q_f x, \\ \phi(x) &= x^\top G x + 2g^\top x + \gamma, & \psi(x) &= x^\top G_f x + 2g_f^\top x + \gamma_f. \end{aligned}$$

with $Q \succeq 0$, $R \succeq 0$, $Q_f \succeq 0$, $G \succeq 0$, $G_f \succeq 0$. We could include (hard) input constraints as well or multiple state constraints (cf. discussion in [4, §1] on modeling joint chance constraints), but abstain from doing so for conciseness.

In this setting affine disturbance feedback [9] has been applied to solve many robust optimal control problems (and even some DRO problems [4]). The idea is to let

$$\boldsymbol{\pi}(\boldsymbol{w}) = \boldsymbol{F}\boldsymbol{w} + \boldsymbol{f},$$

where $\boldsymbol{F} \colon \mathbb{R}^{(N+1)n_w} \to \mathbb{R}^{(N+1)n_x}$ (defined in App. C), has a structure that enforces causality of $\boldsymbol{\pi}$. Note $\boldsymbol{x}_{:N} \in \mathbb{R}^{(N+1)n_x}$, $\boldsymbol{w}_{:N-1} \in \mathbb{R}^{Nn_w}$ and $\boldsymbol{u}_{:N-1} \in \mathbb{R}^{Nn_u}$.

The state trajectory then depends on the disturbance as

$$\boldsymbol{x} = (\boldsymbol{B}\boldsymbol{F} + \boldsymbol{E})\boldsymbol{w} + (\boldsymbol{A}x_0 + \boldsymbol{B}\boldsymbol{f}) = \overline{\boldsymbol{H}}\,\overline{\boldsymbol{w}},$$

with $\boldsymbol{A}, \boldsymbol{B}, \boldsymbol{E}$ defined in App. C, $\overline{\boldsymbol{F}} = [\boldsymbol{F}, \boldsymbol{f}]$ and $\overline{\boldsymbol{H}} = [\boldsymbol{H}, \boldsymbol{h}] = [\boldsymbol{B}\boldsymbol{F} + \boldsymbol{E}, \boldsymbol{A}x_0 + \boldsymbol{B}\boldsymbol{f}]$. Here the linear part, $\boldsymbol{H}$, can be interpreted as the sensitivity of the state to the disturbance,

while $\boldsymbol{h} = (h_0, \dots, h_N)$ is the deterministic part of the state (i.e., the state trajectory when $\boldsymbol{w} = 0$).

We first show how the cost (9a) is implemented, before doing the same for the risk constraints (9c) and the terminal constraint (9d).

A. Reformulation of the cost

We will assume the risk in (9a) has a product ambiguity, $\mathcal{A}^N = \mathcal{A}_\beta(\hat{C}) \times \cdots \times \mathcal{A}_\beta(\hat{C})$ with $\mathcal{A}_\beta(\hat{C})$ as in Lem. II.2. Letting $z := \boldsymbol{x}^\top \boldsymbol{Q}\boldsymbol{x} + \boldsymbol{u}^\top \boldsymbol{R}\boldsymbol{u}$, (9a) is

$$\rho_{\mathcal{A}^N}[z] = \rho_{\mathcal{A}^N}\left[\overline{\boldsymbol{w}}^\top \left(\overline{\boldsymbol{H}}^\top \boldsymbol{Q}\overline{\boldsymbol{H}} + \overline{\boldsymbol{F}}^\top \boldsymbol{R}\overline{\boldsymbol{F}}\right)\overline{\boldsymbol{w}}\right] \tag{17}$$

where $\boldsymbol{Q}$ and $\boldsymbol{R}$ are defined in App. C.

Note, by Ex. II.6, that the factors $\mathcal{A}_i = \mathcal{A}_\beta(\hat{C})$ for $i \in \mathbb{N}_{0:N-1}$ are conic with

$$E_i^\rho \mu = (\pm\langle R_w \overline{w}\, \overline{w}^\top R_w^\top, \mu\rangle), \quad b_i^\rho = (R_w \hat{C} R_w^\top \pm \beta I), \quad \mathcal{K}_i^\rho = \mathbb{S}_+^{n_w+1} \times \mathbb{S}_+^{n_w+1}.$$

Therefore, by Lem. III.2, $\mathcal{A}^N = \times_{i=0}^{N-1} \mathcal{A}_i$ is conic with

$$E^\rho \mu^N = (E^\rho|_0 \mu^N, \dots, E^\rho|_{N-1} \mu^N), \qquad b^\rho = (b_0^\rho, \dots, b_{N-1}^\rho)$$

and $\mathcal{K}^\rho = \mathcal{K}_0^\rho \times \cdots \times \mathcal{K}_{N-1}^\rho$. Let $\Lambda_i^\rho, \mathrm{V}_i^\rho \succeq 0$, $\tau \in \mathbb{R}$, $\lambda_i^\rho = (\Lambda_i^\rho, \mathrm{V}_i^\rho)$ and $b = (R_w \hat{C} R_w^\top \pm \beta I)$ for all $i \in \mathbb{N}_{0:N-1}$. Then, by Lem. III.2,

$$\rho_{\mathcal{A}^N}[z] = \min_{\Lambda_i^\rho \succeq 0, \mathrm{V}_i^\rho \succeq 0, \tau} \quad \tau + \textstyle\sum_{i=0}^{N-1} b \cdot \lambda_i^\rho \tag{18a}$$

$$\text{s.t.} \quad \tau + \textstyle\sum_{i=0}^{N-1} E^\rho|_i^* \lambda_i^\rho \succcurlyeq z, \tag{18b}$$

with $E^\rho|_i^* \lambda_i^\rho = \overline{w}_i^\top R_w^\top (\Lambda_i^\rho - \mathrm{V}_i^\rho) R_w \overline{w}_i$ for $i \in \mathbb{N}_{0:N-1}$.

Our goal is now to find a tight, conservative LMI reformulation of (18b) as an LMI. We use, for any matrix $\Delta \in \mathbb{S}^{n_w+1}$, the partitioning

$$\Delta = \begin{bmatrix} [\Delta]_m & [\Delta]_v \\ [\Delta]_v^\top & [\Delta]_c \end{bmatrix},$$

with $[\Delta]_m \in \mathbb{S}^{n_w}$, $[\Delta]_v \in \mathbb{R}^{n_w}$ and $[\Delta]_c \in \mathbb{R}$. We also introduce $\Delta_i^\rho := R_w^\top (\Lambda_i^\rho - \mathrm{V}_i^\rho) R_w$ for all $i \in \mathbb{N}_{0:N-1}$. Then,

$$\sum_{i=0}^{N-1} E^\rho|_i^* \lambda_i^\rho = \overline{\boldsymbol{w}}^\top \begin{bmatrix} [\Delta_0^\rho]_m & & & [\Delta_0^\rho]_v \\ & \ddots & & \vdots \\ & & [\Delta_{N-1}^\rho]_m & [\Delta_{N-1}^\rho]_v \\ [\Delta_0^\rho]_v^\top & \dots & [\Delta_{N-1}^\rho]_v^\top & \tau + \sum_{i=0}^{N-1} [\Delta_i^\rho]_c \end{bmatrix} \overline{\boldsymbol{w}},$$

which is quadratic in $\boldsymbol{w}$. Note that the argument of $\rho_{\mathcal{A}^N}$ is also quadratic in $\boldsymbol{w}$. Therefore (cf. z in (17)), (18b) is quadratic in $\boldsymbol{w}$. It should hold for all $\boldsymbol{w}$ in the set of quadratic inequalities $\mathcal{W}^N = \{\boldsymbol{w} \colon w_i^\top w_i \leq r^2, \forall i \in \mathbb{N}_{0:N-1}\}$. This is conservatively approximated by a LMI, with the approximate S-Lemma [23, Thm. B.3.1.]. Its effect here is conservatively quantified as increasing the radius r by a factor $9.19\sqrt{\ln(N-1)}$, before

solving the problem exactly. Specifically (18b) holds if there exists a $P^\rho \in \mathbb{S}_+^{N_w+1}$ and $s^\rho \in \mathbb{R}_+^{N-1}$ such that (remembering $\Delta_i^\rho := R_w^\top(\Lambda_i^\rho - \mathrm{V}_i^\rho)R_w$)

$$\begin{bmatrix} [\Delta_0^\rho]_m + [s^\rho]_0 I & & & [\Delta_0^\rho]_v \\ & \ddots & & \vdots \\ & & [\Delta_{N-1}^\rho]_m + [s^\rho]_{N-1} I & [\Delta_{N-1}^\rho]_v \\ [\Delta_0^\rho]_v^\top & \dots & [\Delta_{N-1}^\rho]_v^\top & \tau + \sum_{i=0}^{N-1} [\Delta_i^\rho]_c - [s^\rho]_i r^2 \end{bmatrix} \succeq P^\rho \tag{19}$$

and

$$\begin{bmatrix} P^\rho & \overline{\boldsymbol{H}}^\top \boldsymbol{Q}^{1/2} & \overline{\boldsymbol{F}}^\top \boldsymbol{R}^{1/2} \\ \boldsymbol{Q}^{1/2}\overline{\boldsymbol{H}} & I & \\ \boldsymbol{R}^{1/2}\overline{\boldsymbol{F}} & & I \end{bmatrix} \succeq 0. \tag{20}$$

Here (20) implies $\overline{\boldsymbol{w}}^\top P^\rho \overline{\boldsymbol{w}} \geq z(\boldsymbol{w})$ for all $\boldsymbol{w} \in \mathbb{R}^{N_w}$ by a Schur complement argument. The cost (18a) meanwhile is given as

$$\tau + \textstyle\sum_{i=0}^{N-1} \operatorname{Tr}[\Lambda_i^\rho (R_w \hat{C} R_w^\top + \beta I)] + \operatorname{Tr}[\mathrm{V}_i^\rho (R_w \hat{C} R_w^\top - \beta I)]. \tag{21}$$

B. Reformulation of risk constraints

We consider constraints, $\imath_{\mathcal{A}}^t$ with ambiguity $\mathcal{A} = \mathcal{A}_\beta(\hat{C})$, satisfying (A1). The random variable $\phi(x_t)$ in (9c) equals

$$\overline{x}_t^\top \begin{bmatrix} G & g \\ g^\top & \gamma \end{bmatrix} \overline{x}_t, \quad \text{with } \overline{x}_t = \begin{bmatrix} \boldsymbol{H}_{t,:t-1} & h_t \\ 0 & 1 \end{bmatrix} \overline{\boldsymbol{w}}_{:t-1} \tag{22}$$

and is thus quadratic in $\boldsymbol{w}_{:t-1}$. By construction, (9c) is of type (13) with core ambiguity $\mathfrak{P}^{t-1} \times \mathcal{A}$. Both ambiguity factors are conic by Ex. II.7 and Ex. II.6. We now use $\mu^t = \mu^{t-1} \times \mu_{t-1}$. Then, similarly to before, by Lem. III.2, $\mathfrak{P}^{t-1} \times \mathcal{A}$ is conic for

$$E^\imath \mu^t = E^\imath|_{t-1} \mu^t = (\langle \pm(R_w \overline{w}_{t-1} \overline{w}_{t-1}^\top R_w^\top), \mu^t \rangle),$$

$b^\imath = (R_w \hat{C} R_w^\top \pm \beta I)$ and $\mathcal{K}^\imath = \mathbb{S}_+^{n_w+1} \times \mathbb{S}_+^{n_w+1}$.

Using the duality result of Lem. III.3 shows that (9c) holds if there exists a $\tau_t^\imath, \tau_{\boldsymbol{c},t}^\imath \in \mathbb{R}$, $\Lambda_t^\imath, \mathrm{V}_t^\imath \succeq 0$ and $\lambda_t^\imath = (\Lambda_t^\imath, \mathrm{V}_t^\imath)$ such that

$$\tau_t^\imath + \alpha^{-1} \left(\tau_{\boldsymbol{c},t}^\imath + b^\imath \cdot \lambda_t^\imath \right) \leq 0, \tag{23a}$$

$$\tau_{\boldsymbol{c},t}^\imath + E^\imath|_{t-1}^* \lambda_t^\imath \succcurlyeq 0, \text{ and} \tag{23b}$$

$$\tau_t^\imath + \tau_{\boldsymbol{c},t}^\imath + E^\imath|_{t-1}^* \lambda_t^\imath \succcurlyeq \phi(x_t). \tag{23c}$$

We applied Ex. II.7 for $\mathfrak{P}^{t-1}$ and Ex. II.6 for $\mathcal{A}$. Note that (23a) is a linear constraint. Since $E^\imath|_{t-1}^* \lambda_t^\imath = \overline{w}_{t-1}^\top R_w^\top (\Lambda_t^\imath - \mathrm{V}_t^\imath) R_w \overline{w}_{t-1}$, (23b) and (23c) are of similar structure to (18b). That is, a quadratic inequality that should hold for all $\boldsymbol{w}_{:t-1} \in \mathcal{W}^{t-1}$. Hence we can once again apply the approximate S-Lemma followed by a Schur complement, resulting in a LMI.

Specifically, let $\Delta_t^\imath := R_w^\top (\Lambda_t^\imath - \mathrm{V}_t^\imath) R_w$. Then,

$$E^\imath|_{t-1}^* \lambda_t^\imath = \overline{\boldsymbol{w}}_{:t-1}^\top \begin{bmatrix} [\Delta_t^\imath]_m & [\Delta_t^\imath]_v \\ [\Delta_t^\imath]_v^\top & [\Delta_t^\imath]_c \end{bmatrix} \overline{\boldsymbol{w}}_{:t-1}.$$

Following a similar procedure to before, we apply the approximate S-Lemma for both (23b) and (23c). Specifically, (23b) holds if there exists some $s^{\iota}_{1,t} \in \mathbb{R}^{N-1}_{+}$ such that

$$\begin{bmatrix} [s^{\iota}_{1,t}]_0 I & & & & \\ & \ddots & & & \\ & & [s^{\iota}_{1,t}]_{t-1} I & & \\ & & & [\Delta^{\iota}_t]_m & [\Delta^{\iota}_t]_v \\ & & & [\Delta^{\iota}_t]_v^\top & [\Delta^{\iota}_t]_c + \tau^{\iota}_{\boldsymbol{c},t} - \sum_{i=0}^{t-1} [s^{\iota}_{1,t}]_i r^2 \end{bmatrix} \succeq 0, \tag{24}$$

Meanwhile, (23c) holds if there exists some $P^{\iota}_t \in \mathbb{S}^{tn_w+1}_{+}$ and $s^{\iota}_{2,t} \in \mathbb{R}^{N-1}_{+}$ such that

$$\begin{bmatrix} [s^{\iota}_{2,t}]_0 I & & & & \\ & \ddots & & & \\ & & [s^{\iota}_{2,t}]_{t-1} I & & \\ & & & [\Delta^{\iota}_t]_m & [\Delta^{\iota}_t]_v \\ & & & [\Delta^{\iota}_t]_v^\top & [\Delta^{\iota}_t]_c + \tau^{\iota}_{\boldsymbol{c},t} + \tau^{\iota}_t - \sum_{i=0}^{t-1} [s^{\iota}_{2,t}]_i r^2 \end{bmatrix} \succeq P^{\iota}_t, \tag{25}$$

where, analogously to P^{ρ}, the matrix P^{ι}_t should satisfy

$$\begin{aligned} P^{\iota}_t &\succeq \begin{bmatrix} \boldsymbol{H}_{t,:t-1} & h_t \\ 0 & 1 \end{bmatrix}^\top \begin{bmatrix} G & g \\ g^\top & \gamma \end{bmatrix} \begin{bmatrix} \boldsymbol{H}_{t,:t-1} & h_t \\ 0 & 1 \end{bmatrix} \\ &= \begin{bmatrix} \boldsymbol{H}_{t,:t-1} & h_t \end{bmatrix}^\top G \begin{bmatrix} \boldsymbol{H}_{t,:t-1} & h_t \end{bmatrix} + \begin{bmatrix} 0 & \boldsymbol{H}_{t,:t-1}^\top g \\ g^\top \boldsymbol{H}_{t,:t-1} & 2h_t^\top g + \gamma \end{bmatrix}. \end{aligned} \tag{26}$$

We then apply a Schur complement to find an equivalent LMI for (26):

$$\begin{bmatrix} P^{\iota}_t - \begin{bmatrix} 0 & \boldsymbol{H}_{t,:t-1}^\top g \\ g^\top \boldsymbol{H}_{t,:t-1} & 2h_t^\top g + \gamma \end{bmatrix} & \begin{bmatrix} \boldsymbol{H}_{t,:t-1} & h_t \end{bmatrix}^\top G^{1/2} \\ G^{1/2} \begin{bmatrix} \boldsymbol{H}_{t,:t-1} & h_t \end{bmatrix} & I \end{bmatrix} \succeq 0. \tag{27}$$

Constraint (23a) is equivalent to

$$\tau^{\iota}_t + \alpha^{-1} \left(\tau^{\iota}_{\boldsymbol{c},t} + \mathrm{Tr}[\Lambda^{\iota}_t (R_w \hat{C} R_w^\top + \beta I)] + \mathrm{Tr}[\mathrm{V}^{\iota}_t (R_w \hat{C} R_w^\top - \beta I)] \right) \leq 0. \tag{28}$$

We conclude by reformulating the terminal constraint (9d), which is equivalent to

$$\boldsymbol{w}^\top \begin{bmatrix} \boldsymbol{H}_{N,:} & h_N \\ 0 & 1 \end{bmatrix}^\top \begin{bmatrix} G_f & g_f \\ g_f^\top & \gamma_f \end{bmatrix} \begin{bmatrix} \boldsymbol{H}_{N,:} & h_N \\ 0 & 1 \end{bmatrix} \boldsymbol{w} \leq 0, \tag{29}$$

for all $\boldsymbol{w} \in \mathcal{W}^N$. We apply the approximate S-Lemma, followed by a Schur complement. Specifically (29) holds, if there exists some $s_f \in \mathbb{R}^{N-1}_{+}$ such that

$$\begin{bmatrix} \begin{bmatrix} \mathrm{diag}(s_f) \otimes I_{n_w} & -\boldsymbol{H}_{N,:}^\top g_f \\ -g_f^\top \boldsymbol{H}_{N,:} & -\sum_{i=0}^{N-1} [s_f]_i r^2 - 2h_N^\top g_f - \gamma_f \end{bmatrix} & \begin{bmatrix} \boldsymbol{H}_{N,:} & h_N \end{bmatrix}^\top G_f^{1/2} \\ G_f^{1/2} \begin{bmatrix} \boldsymbol{H}_{N,:} & h_N \end{bmatrix} & I \end{bmatrix} \succeq 0. \tag{30}$$

To summarize, we minimize (21) subject to (19) and (20) (replacing the cumulative stage cost (9a)). Then add constraints (24), (25), (27), (28) for each $t \in \mathbb{N}_{1:N-1}$ (replacing the chance constraints (9c)). The terminal constraint (9d) is then replaced by (30). The resulting minimization problem is a SDP.

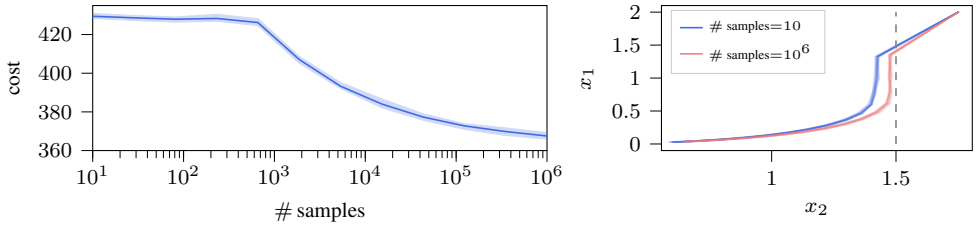

Fig. 1. (left) Range of closed-loop costs when using the matrix-hoeffding bound to construct the ambiguity set for $\boldsymbol{T = 15}$; (right) state trajectories for $\boldsymbol{M = 10, 10^6}$ offline samples.

C. Numerical Results

We consider two experiments: one where data is gathered offline and one where data is gathered online. In both settings the controller acts more optimally when more data is available without affecting recursive feasibility.

Offline learning: Consider the linear system

$$x_{t+1} = \begin{bmatrix} 0.9 & 0.2 \\ 0 & 0.8 \end{bmatrix} x + \begin{bmatrix} 0.1 \\ 0.05 \end{bmatrix} u_t + \begin{bmatrix} 0.5 & 0 \\ 0 & 0.1 \end{bmatrix} w_t,$$

with w_t distributed as a truncated Gaussian with covariance $\Sigma = 0.01^2 I$ and $\|w_t\|_2 \leq 0.15$ and $N = 8$. For the cost, we take $Q = \mathrm{diag}(1, 10), R = 10$ and Q_f and K_f and the solution to the discrete algebraic Riccati equation and the corresponding optimal controller respectively. For the state constraints we have $\alpha = 0.2$ and $\phi(x) = [x]_1 - 1.5 \leq 0$, where $[x]_1$ denotes the first element of x. For the terminal constraint we have $\psi(x) = x^\top Q_f x - 9.524 \leq 0$, which is RPI for $\pi_f(x) = K_f x$ and $x_0 = (1.75, 2)$.

We investigate sample counts M between 10 and 10^6. For each sample count we gather an offline data set and compute $\mathcal{A}$ from Lem. II.2 for $\delta = 0.05$. Then we propagate the dynamics, with the receding horizon controller (9) for $T = 15$ time steps, cumulating the stage cost for each step. This experiment is repeated 25 times. The range of resulting costs is shown in Fig. 1. For $M = 10$ and $M = 10^6$, closed-loop state trajectories are depicted. The cost initially remains approximately constant, since the moment constraint is inactive due to insufficient samples. Afterwards, the controller gains confidence, moving closer to the state constraint when more data is available, resulting in a lower cost.

Online learning: We illustrate how our framework can be used to construct controllers that learn online, guaranteeing recursive feasibility. Consider the scalar linear system

$$x_{t+1} = x_t + u_t + w_t,$$

with w_t distributed as a truncated Gaussian with variance $\sigma^2 = 0.05^2$, expectation $\mathbb{E}[w_t] = 0.05$ and $|w_t| \leq 1.0$. Let $Q = 1$, $R = 0.1$, $N = 5$, $\alpha = 0.2$ and $\phi(x) = 1 - x \leq 0$. We use $\pi_f(x) = -0.9x + 2.375$, $\psi(x) = (x - 2.375)^2 - 1.375^2 \leq 0$ and Q_f the solution to the discrete algebraic Riccati equation. To implement learning, the ambiguity update scheme described in App. D is used (which constructs ambiguity from data based on Lem. II.2 for $\delta = 0.05$ and rejects an ambiguity set if it is not a subset of the previously accepted ambiguity).

We consider three cases: *(i)* robust, with $\mathcal{A} = \mathfrak{P}$ and no learning; *(ii)* online, starting from $M = 10$ samples and then updating the ambiguity online while running the

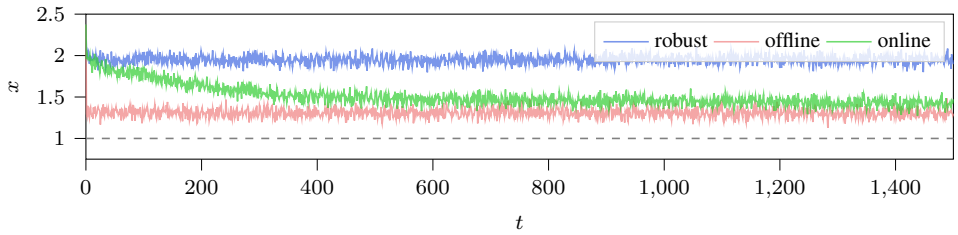

Fig. 2. State trajectories for robust; online-; and offline learning.

controller; and *(iii)* offline, where $M = 10\,000$ initial samples and no updates. Fig. 2 depicts the state trajectory, which approaches the constraint for larger M.

REFERENCES

[1] S. Shafieezadeh-Abadeh, P. M. Esfahani, and D. Kuhn, "Distributionally robust logistic regression," in *NIPS'15*, vol. 1, (Cambridge, MA, USA), p. 1576–1584, MIT Press, 2015.

[2] E. Delage and Y. Ye, "Distributionally Robust Optimization Under Moment Uncertainty with Application to Data-Driven Problems," *Oper. Res.*, vol. 58, no. 3, pp. 595–612, 2010.

[3] K. Kim and I. Yang, "Minimax Control of Ambiguous Linear Stochastic Systems Using the Wasserstein Metric," in *CDC*, pp. 1777–1784, IEEE, 2020.

[4] B. P. G. Van Parys, D. Kuhn, P. Goulart, and M. Morari, "Distributionally Robust Control of Constrained Stochastic Systems," *IEEE TAC*, vol. 61, no. 2, 2015.

[5] H. Rahimian and S. Mehrotra, "Distributionally Robust Optimization: A Review," 2019, arXiv: 1908.05659.

[6] A. Pichler and A. Shapiro, "Mathematical Foundations of Distributionally Robust Multistage Optimization," 2021, arXiv: 2101.02498.

[7] Z. Chen, M. Sim, and H. Xu, "Distributionally Robust Optimization with Infinitely Constrained Ambiguity Sets," *Oper. Res.*, vol. 67, no. 5, pp. 1328–1344, 2019.

[8] P. Sopasakis, M. Schuurmans, and P. Patrinos, "Risk-averse risk-constrained optimal control," in *ECC*, pp. 375–380, IEEE, 2019.

[9] P. Goulart, E. Kerrigan, and J. Maciejowski, "Optimization over state feedback policies for robust control with constraints," *Automatica*, vol. 42, no. 4, pp. 523–533, 2006.

[10] C. Mark and S. Liu, "Stochastic MPC with Distributionally Robust Chance Constraints," 2020, arXiv: 2005.00313.

[11] S. Lu, J. H. Lee, and F. You, "Soft-constrained model predictive control based on data-driven distributionally robust optimization," *AIChE J.*, vol. 66, no. 10, pp. 1–18, 2020.

[12] M. Schuurmans and P. Patrinos, "Learning-based distributionally robust model predictive control of markovian switching systems with guaranteed stability and recursive feasibility," in *CDC*, pp. 4287–4292, IEEE, 2020.

[13] P. Sopasakis, D. Herceg, A. Bemporad, and P. Patrinos, "Risk-averse model predictive control," *Automatica*, vol. 100, pp. 281–289, 2019.

[14] W. Wiesemann, D. Kuhn, and B. Rustem, "Robust Markov Decision Processes," *Math. Oper. Res.*, vol. 38, no. 1, pp. 153–183, 2013.

[15] A. Ruszczyński and A. Shapiro, "Optimization of convex risk functions," *Mathematics of Operations Research*, vol. 31, no. 3, pp. 433–452, 2006.

[16] J. F. Bonnans and A. Shapiro, *Perturbation Analysis of Optimization Problems*. New York, NY: Springer New York, 2000.

[17] A. Shapiro, "On Duality Theory of Conic Linear Problems," in *Semi-Infinite Programming. Nonconvex Optimization and Its Applications, vol 57*, pp. 135–165, Boston, MA: Springer, 2001.

[18] D. Kuhn, P. M. Esfahani, V. A. Nguyen, and S. Shafieezadeh-Abadeh, "Wasserstein Distributionally Robust Optimization," in *Oper. Res. & Mgmt. Sci. in the Age of Analytics*, pp. 130–166, 2019.

[19] P. Coppens, M. Schuurmans, and P. Patrinos, "Data-driven distributionally robust LQR with multiplicative noise," in *L4DC*, vol. 120, (The Cloud), pp. 521–530, PMLR, 2020.

[20] J. A. Tropp, "User-Friendly Tail Bounds for Sums of Random Matrices," *FoCM*, vol. 12, no. 4, pp. 389–434, 2012.

[21] C. D. Aliprantis and K. C. Border, *Infinite Dimensional Analysis*. Berlin/Heidelberg: Springer-Verlag, 2006.
[22] W. Wiesemann, D. Kuhn, and M. Sim, "Distributionally robust convex optimization," *Oper. Res.*, vol. 62, no. 6, pp. 1358–1376, 2014.
[23] A. Ben-Tal, L. El Ghaoui, and A. Nemirovski, *Robust optimization*. Princeton University Press, 2009.
[24] A. Shapiro, "Time consistency of dynamic risk measures," *Oper. Res. Lett.*, vol. 40, no. 6, pp. 436–439, 2012.
[25] A. Shapiro, D. Dentcheva, and A. Ruszczyński, *Lectures on Stochastic Programming*. Soc. Ind. Appl. Math., 2014.
[26] S. Zymler, D. Kuhn, and B. Rustem, "Distributionally robust joint chance constraints with second-order moment information," *Math. Prog.*, vol. 137, no. 1-2, pp. 167–198, 2013.
[27] A. Shapiro and A. Kleywegt, "Minimax analysis of stochastic problems," *Optim Methods Softw.*, vol. 17, no. 3, pp. 523–542, 2002.
[28] M. Lorenzen, F. Dabbene, R. Tempo, and F. Allgower, "Constraint-Tightening and Stability in Stochastic Model Predictive Control," *IEEE TAC*, vol. 62, no. 7, pp. 3165–3177, 2017.
[29] T. Rusten and R. Winther, "A Preconditioned Iterative Method for Saddlepoint Problems," *SIAM Journal on Matrix Analysis and Applications*, vol. 13, pp. 887–904, jul 1992.
[30] J. A. Tropp, "An Introduction to Matrix Concentration Inequalities," *Found. Trends Mach. Learn.*, vol. 8, no. 1-2, pp. 1–230, 2015, arXiv: 1501.01571.
[31] K. B. Athreya and S. N. Lahiri, *Measure Theory and Probability Theory*. Springer Texts in Statistics, Springer New York, 2006.
[32] A. M. Bruckner, J. B. Bruckner, and B. S. Thomson, *Real Analysis*. Prentice-Hall, Inc., 2007.
[33] H. van Waarde, K. Camlibel, and M. Mesbahi, "From noisy data to feedback controllers: non-conservative design via a matrix S-lemma," 2020, arXiv: 2006.00870.

APPENDIX

A. Proof of Lem. II.2

We begin by showing an auxiliary lemma, closely following the proof of [29, Lemma. 2.1].

Lemma A.1. *Consider the following matrix, with $X \in \mathbb{S}^d$ and $x \in \mathbb{R}^d$,*

$$\overline{X} := \begin{bmatrix} X & x \\ x^\top & 0 \end{bmatrix}.$$

Assume $\|X\|_2 \leq r_X^2$ and $\|x\|_2 \leq r_x$. Then $\|X\|_2 \leq (r_X^2 + \sqrt{r_X^4 + 4r_x^2})/2$.

Proof. Assume (v, γ) is an eigenvector of $\overline{X}$ with associated eigenvalue λ. We assume without loss of generality that $\lambda \neq 0$. Then

$$Xv + \gamma x = \lambda v, \tag{31}$$
$$x^\top v = \lambda \gamma. \tag{32}$$

Left multiplying (31) with $v^\top$ gives

$$\begin{aligned} & v^\top X v + \gamma x^\top v = \lambda v^\top v \\ \Leftrightarrow \quad & p(\lambda) := \lambda v^\top X v + v^\top (x x^\top) v - \lambda^2 v^\top v = 0, \end{aligned} \tag{33}$$

where equivalence holds due to (32) and the assumption $\lambda \neq 0$.

(a) Assume $\lambda > 0$: in this case, from (33), $\|X\|_2 \leq r_X^2$ and $\|x\|_2 \leq r_x$ we have

$$(-r_X^2 \lambda - \lambda^2)\|v\|_2^2 \leq p(\lambda) \leq (r_X^2 \lambda + r_x^2 - \lambda^2)\|v\|_2^2.$$

Knowing that $\|v\|_2^2 \geq 0$ and $p(\lambda) = 0$, we should have $(-r_X^2 \lambda - \lambda^2) \leq 0$. This holds trivially since $\lambda \geq 0$. We should also have $(r_X^2 \lambda + r_x^2 - \lambda^2) \geq 0$, which results in a condition on λ by solving the quadratic inequality and taking the positive solution (since we assumed $\lambda \geq 0$). The resulting bound is $\lambda \leq (r_X^2 + \sqrt{r_X^4 + 4r_x^2})/2$.

(b) Assume $\lambda < 0$: in this case we have

$$(r_X^2\lambda - \lambda^2)\|v\|_2^2 \leq p(\lambda) \leq (-r_X^2\lambda + r_x^2 - \lambda^2)\|v\|_2^2.$$

Following similar arguments as above, we see that the first inequality holds trivially, while the second holds iff $\lambda \geq -(r_X^2 + \sqrt{r_X^4 + 4r_x^2})/2$.

We have thus shown $-(r_X^2 + \sqrt{r_X^4 + 4r_x^2})/2 \leq \lambda \leq (r_X^2 + \sqrt{r_X^4 + 4r_x^2})/2$, thereby proving the claimed result. □

Using this lemma we can now proceed with the main proof.

Let $\overline{X_i} := R_w(\overline{w}_i\overline{w}_i^\top - \mathbb{E}_{\mu_\star}[\overline{w}\,\overline{w}^\top])R_w^\top$ and its eigenvalue decomposition $\overline{X_i} = V_iD_iV_i^\top$. Clearly this matrix is structured as $\overline{X}$ in Lem. A.1, where $X = w_iw_i^\top - \mathbb{E}_{\mu_\star}[ww^\top]$ and $x = cr(w_i - \mathbb{E}_{\mu_\star}[w])$. We know that a.s., $-r^2I \preceq -\mathbb{E}_{\mu_\star}[ww^\top] \preceq X \preceq w_iw_i^\top \preceq r^2I$, hence $\|X\|_2 \leq r^2$. From a triangle inequality we also know that $\|x\|_2 \leq 2cr^2$ a.s. Applying Lem. A.1 then shows $\|\overline{X}_i\|_2 \leq r^2(1 + \sqrt{1 + 16c})/2 =: \overline{\lambda}$ a.s. Therefore the eigenvalues are bounded as $|d_i| \leq \overline{\lambda}$.

As such, we can produce a matrix sub-Gaussian bound:

$$\begin{aligned}\mathbb{E}[\exp(\theta\overline{X}_i)] &= \mathbb{E}[V_i\exp(\theta D_i)V_i^\top] \\ &\overset{(i)}{\preceq} \mathbb{E}\left[\tfrac{(\overline{\lambda}I+\overline{X_i})}{2\overline{\lambda}}\exp(\theta\overline{\lambda}) + \tfrac{(\overline{\lambda}I-\overline{X_i})}{2\overline{\lambda}}\exp(-\theta\overline{\lambda})\right] \\ &\overset{(ii)}{=} \tfrac{\exp(\theta\overline{\lambda})+\exp(-\theta\overline{\lambda})}{2}I = \cosh(\theta\overline{\lambda})I \preceq \exp(\theta^2\overline{\lambda}^2/2)I,\end{aligned}$$

where *(i)* follows from convexity and $V_iV_i^\top = I$, and *(ii)* from $\mathbb{E}[\overline{X_i}] = 0$.

From independence and [30, Lem. 3.5.1] we have, $\mathbb{E}[\mathrm{Tr}[\exp(\sum_{i=0}^{M-1}\overline{X_i})]] \leq d\exp(M\theta^2\overline{\lambda}^2/2)$, with $d := n_w + 1$. This allows us to apply [30, Prop. 3.2.1]:

$$\mathbb{P}[\lambda_{\max}(\textstyle\sum_{i=0}^{M-1}\overline{X_i}) \geq \beta] \leq \inf_{\theta>0} d\exp(-\theta\beta + M\theta^2\overline{\lambda}^2/2) = d\exp(-\beta^2/(2M\overline{\lambda}^2)).$$

Similarly bounding $\lambda_{\min}$ and applying a union bound gives

$$\mathbb{P}[\|\textstyle\sum_{i=0}^{M-1}\overline{X_i}\| \geq \beta] \leq 2d\exp(-\beta^2/(2M\overline{\lambda}^2)).$$

This is equivalent to

$$\mathbb{P}[\|\textstyle\sum_{i=0}^{M-1}\overline{X_i}/M\| \geq \beta] \leq 2d\exp(-M\beta^2/(2\overline{\lambda}^2)),$$

by a change of variable for β. Setting the right-hand-side equal to δ and solving for β gives the value in Lem. II.2, showing $\mathbb{P}[\|\hat{C} - C\|_2 \leq \beta] \geq 1 - \delta$, implying (6). □

B. Continuity of iterated integration

We argue why $h(\boldsymbol{w}_{:t-2})$ in the proof of Thm. IV.2 is a measurable, continuous random variable. It is given by

$$h(\boldsymbol{w}_{:t-2}) = \int_{\mathcal{W}} z_\tau(\boldsymbol{w}_{:t-2}, w_{t-1})\mathrm{d}\mu(w_{t-1}),$$

for continuous $z_\tau \in \mathcal{Z}^t$. Measurability follows by Tonelli's theorem [31, Thm. 5.2.1], compactness of $\mathcal{W}$ and continuity of z_τ (implying σ-finiteness of the measure spaces and

integrability). Meanwhile, continuity follows by an epsilon-delta argument [32, §9.3]. Specifically, Jensen's inequality gives

$$\begin{aligned}|h(\boldsymbol{w}_{:t-2}) - h(\boldsymbol{w}'_{:t-2})| &\leq \int_{\mathcal{W}} |z_\tau(\boldsymbol{w}_{:t-2}, w_{t-1}) - z_\tau(\boldsymbol{w}'_{:t-2}, w_{t-1})| \mathrm{d}\mu_{t-1}(w_{t-1}), \\ &\leq \max_{w_{t-1}\in\mathcal{W}} \left\{|z_\tau(\boldsymbol{w}_{:t-2}, w_{t-1}) - z_\tau(\boldsymbol{w}'_{:t-2}, w_{t-1})|\right\} \end{aligned} \tag{34}$$

for any $\boldsymbol{w}_{:t-2}, \boldsymbol{w}'_{:t-2} \in \mathcal{W}^{t-1}$. The second inequality holds since $|z_\tau(\boldsymbol{w}_{:t-2}, w_{t-1}) - z_\tau(\boldsymbol{w}'_{:t-2}, w_{t-1})|$ is continuous in w_{t-1} by continuity of z_τ, which enables the application of Ex. II.7 to replace integration over $\mu_{t-1} \in \mathfrak{P}(\mathcal{W})$ by a maximization over $w_{t-1} \in \mathcal{W}$.

From continuity of z_τ we have

$$\forall \varepsilon \geq 0, \exists \delta \geq 0 \text{ s.t. } \|\boldsymbol{w} - \boldsymbol{w}'\| \leq \delta \Rightarrow |z(\boldsymbol{w}) - z(\boldsymbol{w}')| \leq \epsilon.$$

So picking $\boldsymbol{w} = (\boldsymbol{w}_{:t-2}, w_{t-1})$ and $\boldsymbol{w}' = (\boldsymbol{w}'_{:t-2}, w_{t-1})$ and plugging into (34) gives

$$\|\boldsymbol{w}_{:t-2} - \boldsymbol{w}'_{:t-2}\| \leq \delta \Rightarrow |h(\boldsymbol{w}_{:t-2}) - h(\boldsymbol{w}'_{:t-2})| \leq \varepsilon. \tag{35}$$

Hence, for any $\varepsilon \geq 0$ we have a $\delta \geq 0$ s.t. (35) holds, implying continuity of h. □

C. Affine disturbance feedback notation

Let $N_x := (N+1)n_x$, $N_w := Nn_w$ and $N_u := Nn_u$. Following [9], we consider the matrices, $\boldsymbol{A} \in \mathbb{R}^{N_x \times n_x}$, $\boldsymbol{B} \in \mathbb{R}^{N_x \times N_u}$, $\boldsymbol{E} \in \mathbb{R}^{N_x \times N_w}$ and $\boldsymbol{F} \in \mathbb{R}^{N_u \times N_w}$ such that

$$\boldsymbol{A} := \begin{bmatrix} I_{n_x} \\ A \\ A^2 \\ \vdots \\ A^N \end{bmatrix}, \qquad \boldsymbol{B} := \begin{bmatrix} 0 & 0 & \dots & 0 \\ B & 0 & \dots & 0 \\ AB & B & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ A^{N-1}B & A^{N-2}B & \dots & B \end{bmatrix}, \tag{36}$$

$$\boldsymbol{E} := \begin{bmatrix} 0 & 0 & \dots & 0 \\ E & 0 & \dots & 0 \\ AE & E & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ A^{N-1}E & A^{N-2}E & \dots & E \end{bmatrix}, \qquad \boldsymbol{F} := \begin{bmatrix} 0 & 0 & \dots & 0 \\ F_{1,0} & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ F_{N-1,0} & \dots & F_{N-1,N-2} & 0 \end{bmatrix}, \tag{37}$$

$\boldsymbol{Q} = \mathrm{diag}(Q, \dots, Q, Q_f)$ and $\boldsymbol{R} = \mathrm{diag}(R, \dots, R)$.

D. Updating the ambiguity

In this section we examine how (A2) can be guaranteed in practice. More specifically, assume we have M offline samples used to construct $\mathcal{A}_0$ according to Lem. II.2. At time step t we then will have gathered $M + t$ samples, which can be used to construct $\hat{\mathcal{A}}_t$. We then use the update rule

$$\mathcal{A}_{t+1} = \begin{cases} \hat{\mathcal{A}}_{t+1}, & \text{if } \hat{\mathcal{A}}_{t+1} \subseteq \mathcal{A}_t \\ \mathcal{A}_t, & \text{else.} \end{cases} \tag{38}$$

Hence we only need to check $\hat{\mathcal{A}}_{t+1} \subseteq \mathcal{A}_t$. Letting $\hat{C}_+$ (β_+), $\hat{C}$ (β) denote the center (radius) of $\hat{\mathcal{A}}_{t+1}$ and $\mathcal{A}_t$ respectively. Then a sufficient condition for $\hat{\mathcal{A}}_{t+1} \subseteq \mathcal{A}_t$ is,

$$\|R_w(\hat{C} - C)R_w^\top\|_2 \leq \beta, \quad \forall C \succeq 0, \text{ s.t., } \|R_w(\hat{C}_+ - C)R_w^\top\|_2 \leq \beta_+.$$

This condition is only sufficient since the bounded support would imply $\mathrm{Tr}([C]_{:n_w,:n_w}) \leq r^2$, where $[C]_{:n_w,:n_w}$ is a matrix containing the first n_w rows and columns of C (i.e.,

$\mathbb{E}[ww^\top]$). The condition that μ is a probability measure implies $[C]_{n_w+1,n_w+1} = 1$, where $[C]_{n_w+1,n_w+1}$ denotes the bottom-right element of C. Based on numerical experiments however, dropping these constraints does not result in significant conservativeness.

Condition (38) is not trivial to check however. It requires maximizing the convex function $\|R_w(\hat{C}-C)R_w^\top\|_2$ over a convex set, which is NP hard in general. We however have some structure we can exploit. Note how

$$\|R_w(\hat{C}-C)R_w^\top\|_2 \leq \beta \qquad \Leftrightarrow \qquad \begin{bmatrix} C \\ I \end{bmatrix}^\top \begin{bmatrix} -I & \hat{C} \\ \hat{C} & \beta^2(R_w^\top R_w)^{-2} - \hat{C}^2 \end{bmatrix} \begin{bmatrix} C \\ I \end{bmatrix} \succeq 0.$$

This, should hold for all $C \succeq 0$ such that

$$\begin{bmatrix} C \\ I \end{bmatrix}^\top \begin{bmatrix} -I & \hat{C}_+ \\ \hat{C}_+ & \beta_+^2(R_w^\top R_w)^{-2} - \hat{C}_+^2 \end{bmatrix} \begin{bmatrix} C \\ I \end{bmatrix} \succeq 0.$$

A sufficient condition for which is that,

$$v^\top \begin{bmatrix} -I & \hat{C} \\ \hat{C} & \beta^2(R_w^\top R_w)^{-2} - \hat{C}^2 \end{bmatrix} v \geq 0,$$

for all $v \in \mathbb{R}^{2(n_w+1)}$, such that

$$v^\top \begin{bmatrix} -I & \hat{C}_+ \\ \hat{C}_+ & \beta_+^2(R_w^\top R_w)^{-2} - \hat{C}_+^2 \end{bmatrix} v \geq 0 \text{ and } v^\top \begin{bmatrix} & I \\ I & \end{bmatrix} v \geq 0.$$

In fact this condition is sufficient and necessary for (38) (cf. [33]). Applying the Approximate S-Lemma [23, Thm. B.3.1.] (which is exact in this case) allows us to formulate this final condition as a LMI.