Safe learning LQR of linear dynamics with multiplicative noise

Peter Coppens and Panagiotis Patrinos†

***Abstract*— Control of linear dynamics with multiplicative noise naturally introduces robustness against dynamical uncertainty. Moreover, many physical systems are subject to multiplicative disturbances. In this work we show how these dynamics can be identified from state trajectories. The least-squares scheme enables exploitation of prior information and comes with practical data-driven confidence bounds and sample complexity guarantees. We complement this scheme with an associated control synthesis procedure for LQR which robustifies against distributional uncertainty, guaranteeing stability with high probability and converging to the true optimum at a rate inversely proportional with the sample count. Throughout we exploit the underlying multi-linear problem structure through tensor algebra and completely positive operators. The scheme is validated through numerical experiments.**

I. Introduction

Recently the control community has gained renewed interest in learning control [1]. This can be viewed as a reaction to the practical success of machine learning — specifically *Reinforcement Learning (RL)* — when applied to control dynamical systems. An opportunity arises there in the lack of theoretical guarantees, particularly related to safety. On that regard, control theory has much to offer [1], [2].

One approach towards learning with guarantees has been to restrict the problem of learning control to specific classes of dynamics. The *Linear Quadratic Regulator (LQR)* [3, §4.1] has proven a valuable subject for such investigations. We categorize two approaches: *top-down* takes a RL algorithm and applies it to LQR, proving convergence, stability, etc. [4]–[6]; and *bottom-up* develops learning control schemes specialized to LQR with guarantees [7]–[9].

The latter approach is usually the classic *model-based* [10, p. 360], [11, p. 537], [3, §6.1], implying that the dynamics are identified first, before synthesizing a controller. As suggested in [1] under the term *coarse-id* learning, such model-based approaches — when combined with tools from statistical learning theory — produce provably safe controllers [7]. We also present such a coarse-id method here.

Specifically we apply this framework to *linear dynamics with multiplicative noise* and develop a *distributionally robust (DR)* generalization [12] of the coarse-id framework for deterministic LQR, first providing a least-squares estimator of the dynamics, paired with tight error bounds, before presenting a control synthesis scheme. We thereby extend the developments in our previous work [13] and interpolate between robust and stochastic control, becoming less conservative as data is gathered, while guaranteeing stability with high probability.

LQR under multiplicative noise was first considered by Wonham [14], who developed a generalized *Riccati equation*. Similarly many recent developments in learning LQR are generalizable to multiplicative noise [13], [15]–[17]. This is interesting for two reasons.

First, several real-world systems are influenced by multiplicative noise. These models occur frequently when disturbances enter a model through the parameters (e.g., through vibrations, nonlinear effects in spring systems, thermal and aerodynamic influences, etc.). A list of examples with references is provided in [18, §1.9] with concrete applications in aerospace and vehicle control. Multiplicative noise has also been observed in biological applications like sensorimotor control [19], [20], cell populations [21], [22], population migration [22] and immune systems [23]. Other applications include nuclear fission [22], power grids [24], communication channels [25] and climate models [26].

Second, the use of multiplicative noise induces robustness against parametric uncertainty. In [27] this is motivated by the lack of stability margins in output-feedback LQG [28]. This same fact was argued in a historical review by J. Doyle [29] to have motivated the development of $\mathcal{H}_\infty$ methods. In fact, as we formalize later, there is a strong relation between *mean square stability (mss)* for multiplicative noise systems and robust stability [27], [30], [31].

Next, we provide a detailed overview of related work.

A. Related work

Historical Overview: The history of multiplicative noise is extensive, and it has been investigated under several names. It was first considered by [14] under the name *linear dynamics with state-dependent noise*, for which an optimal control scheme was developed. This resulted in several other early works, an overview of which is given in [18, p.6]. Note that the name *stochastic bilinear systems* has also been used to describe such systems [22], [32]. Most recent developments related to multiplicative noise have been related to control synthesis with indefinite stage costs [33], numerical solution of generalized Riccati equations [18] and to applying $\mathcal{H}_\infty$ methods to stochastic control [18], [34]. We similarly consider robustification against perturbations to the distribution of the disturbance (similar to [35]), but do so in a data-driven fashion.

†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

The work of the first author is supported by: the Research Foundation Flanders (FWO) PhD grant 11E5520N and research projects G086518N and G086318N; Research Council KU Leuven C1 project No. C14/18/068; Fonds de la Recherche Scientifique – FNRS and the Fonds Wetenschappelijk Onderzoek – Vlaanderen under EOS project no 30468160 (SeLMA).

Learning control for multiplicative noise: In the top-down approach several results from LQR have been generalized to multiplicative noise already. Q-learning has been investigated in [15] and policy gradient in [16], [17].

Developments in the bottom-up approach remain limited, mostly since multiplicative noise dynamics are especially challenging to identify. With some exceptions considering partial identification [36], system identification has only been investigated recently. In [37], [38] a scheme based on deterministic linear system identification was developed by averaging over trajectories. The disadvantage over our method is that prior model structure can not be used and the confidence bounds are loose. The same authors also considered an identification scheme, where the mean dynamics are identified first and the noise is characterized using bootstrapping [39]. Similarly, this scheme comes with few guarantees. Finally [40] considers a least-squares estimator similar to ours. They consider scalar measurements, thwarting complete identification. Their confidence bounds are valid for single-trajectories (while we use multiple). We show experimentally that our method is also applicable in the single-trajectory case, but leave a formal proof to further work.

Mathematical Tools: Throughout this work we illustrate what mathematical tools are applicable in the analysis of linear systems with multiplicative noise.

We show how the underlying bilinear structure is expressed using tensor algebra [41], [42], making the formulation of least squares estimators and studying their accuracy convenient.

We also exploit the connection with *completely positive (CP)* operators used to model quantum channels [31]. They were linked to multiplicative noise dynamics before in [18] and to stability of linear systems and Lyapunov operators in [43]–[45]. The equivalence between mss and robust stability too was developed originally for CP operators [31].

Our estimator is similar to those considered in compressed covariance sensing [46] and those in bilinear estimation [47].

To develop our concentration inequalities we use modern tools from matrix concentration [48]. Our controller synthesis uses results from robust optimization [49]. Specifically we use *distributionally robust* optimization [12] to interpolate between robust and stochastic control synthesis.

B. Contributions

The main contributions of this work are as follows.

(i) We show how *linear quadratic (LQ)* control of multiplicative noise is equivalent to LQ control of CP dynamics;

(ii) We consider an identification procedure that can both include prior knowledge about the dynamics and also performs well in the model-free setting. The optimal estimate is accompanied with a tight confidence set;

(iii) We introduce a synthesis procedure for DR LQR;

(iv) We analyse the sample complexity of the scheme.

As mentioned above, CP operators have been applied in multiplicative noise before [44]. A formal equivalence of LQ control has not been shown however. The other results all strongly rely on this equivalence.

A similar problem to (ii) was previously considered in [37], [38], however a practical ambiguity set is not derived and there is no possibility to include prior information about the dynamics. In [40] a partial ambiguity set is derived, but the full system dynamics are never identified. We consider (ii) as the most important contribution.

We also generalize our previous work to the model-free setting: (iii) extends upon [13] and (iv) extends upon [50], and applies the result to the DR control synthesis procedure.

We illustrate practical applicability of our work through numerical experiments.

The remainder of this paper describes the construction, solution and analysis of DR control for LQR under multiplicative noise. We provide a problem statement in §II and preliminary results in §III. We then consider the nominal case for LQR under multiplicative noise, when the dynamics are known. Next, in §V, we describe an ambiguity set based on state measurements; in §VI-A we solve the resulting DR problem; and in §VI-B we study convergence of the solution to the nominal controller as more data is gathered. Then we show how to add more prior information to the model — related to a deterministic term — in §VII. We end the paper with numerical results in §VIII and a conclusion in §IX.

C. Notation overview

Below we provide an overview of the notation.

Let $\mathbb{N}$ denote the integers and $\mathbb{R}$ ($\overline{\mathbb{R}}$) the (extended) reals. We introduce the shorthand $\mathbb{N}_{a:b} = \{a, \dots, b\}$ for $a, b \in \mathbb{N}$, with $a \leq b$. For $x \in \mathbb{R}^m$ let $\|x\|_2 = \sqrt{x^\top x}$ denote the Euclidean norm and $\mathcal{B}_n \subset \mathbb{R}^n$ the unit-norm ball. Assuming column vectors, let $(x, y) = [x; y]$ denote the vertical concatenation between vectors.

For a random vector w let $\mathbb{E}_\nu[w]$ denote the expectation with respect to the measure ν. The subscript ν is omitted when the true distribution is implied.

Matrices: For $Z \in \mathbb{R}^{m\times n}$ let $\lambda(Z) = (\lambda_1, \dots, \lambda_m)$ (when $m = n$) denote the vector of eigenvalues in descending order and $\rho(Z) = \lambda_1(X)$ the spectral radius. Similarly $\sigma(Z) = (\sigma_1, \dots, \sigma_r)$ denotes the vector of non-zero singular values in descending order with $r = \mathrm{rk}(Z)$ the rank. Denote by $Z^\dagger$ and $Z^\top$ the pseudo-inverse and the transpose respectively. Also let $\|Z\|_2$ ($\|Z\|_F$) be the spectral (Frobenius) norm. Let $\otimes$ denote the Kronecker product and $\mathrm{vec}(X) \in \mathbb{R}^{mn}$ the vectorization of X resulting from stacking the columns, and $\mathrm{unvec}(\cdot)$ its reverse. When $m = n$ let $\mathrm{Tr}(Z)$ denote the trace.

For conformable matrices X, Y we use $[X; Y]$ ($[X, Y]$) for vertical (horizontal) concatenation, let $\mathrm{diag}(X, Y)$ be a (block) diagonal and let $I_d \in \mathbb{R}^{d\times d}$ be the identity. Elements of matrices $X \in \mathbb{R}^{m\times n}$ (and vectors $x \in \mathbb{R}^d$) are indexed using $[X]_{ij}$ ($[x]_i$) for $i \in \mathbb{N}_{1:m}$, $j \in \mathbb{N}_{1:n}$ ($i \in \mathbb{N}_{1:d}$). A slice $a : b$ (:) substitutes index i when it takes values in $\mathbb{N}_{a:b}$ ($\mathbb{N}_{1:m}$).

We denote by $\mathbb{S}^d$ the set of symmetric d by d matrices and by $\mathbb{S}^d_{++}$ ($\mathbb{S}^d_+$) the positive (semi)definite matrices. Then, for

$P, Q \in \mathbb{S}^d$, we write $P \succ Q$ ($P \succeq Q$) To signify $P - Q \in \mathbb{S}^d_{++}$ ($P - Q \in \mathbb{S}^d_+$). Similarly $P \succneqq 0$ implies $0 \neq P \succeq 0$. Let $\mathfrak{d}_d = (d+1)d/2$ denote the degrees of freedom of elements of $\mathbb{S}^d$ and $\mathrm{vec_s}(X) \in \mathbb{R}^{\mathfrak{d}_d}$ denotes the symmetrized vectorization (cf. Def. III.1) such that $\mathrm{vec_s}(X)^\top \mathrm{vec_s}(X) = \|X\|_F^2$, with $\mathrm{unvec_s}(\cdot)$ its inverse. Let $\otimes_{\mathrm{s}}$ the symmetrized Kronecker product (cf. Def. III.2).

Tensors: We use bold calligraphic letters to denote (third-order) tensors $\boldsymbol{\mathcal{A}} \in \mathbb{R}^{p\times q\times r}$. We follow the notation of [41] and provide standalone definitions in §III-B. A tensor is indexed as $[\boldsymbol{\mathcal{A}}]_{ijk}$ with $i \in \mathbb{N}_{1:p}$, $j \in \mathbb{N}_{1:q}$, $k \in \mathbb{N}_{1:r}$, where an index i is replaced by $a : b$ (:) when i takes values in $\mathbb{N}_{a:b}$ ($\mathbb{N}_{1:p}$). We use $(\overline{\times}_n)$ $\times_n$ to denote the *n-mode (vector) product* and $\boldsymbol{\mathcal{A}}_{(n)}$ for the *n-mode matricization*. Let $[\![\boldsymbol{\mathcal{A}}; X_1, X_2, X_3]\!] := \boldsymbol{\mathcal{A}} \times_1 X_1 \times_2 X_2 \times_3 X_3$ denote the Tucker operator (cf. Rem. III.10).

II. Problem Statement

In this work we consider linear systems with input- and state-multiplicative noise given for all $t \in \mathbb{N}$ as

$$x_{t+1} = A(v_t)x_t + B(v_t)u_t, \tag{1}$$

where $x_t \in \mathbb{R}^{n_x}$ is the state, $u_t \in \mathbb{R}^{n_u}$ the input and $v_t \in \mathbb{R}^{n_v}$ is the disturbance, which follows an i.i.d. random process. The matrices are given by $A(v) := \sum_{i=1}^{n_v} [v]_i A_i$ and $B(v) := \sum_{i=1}^{n_v} [v]_i B_i$.

The primary goal is to study solutions of the following stochastic LQR problem:

$$\underset{u_0,u_1,\dots}{\text{minimize}} \quad \mathbb{E}\left[\sum_{t=0}^{\infty} x_t^\top Q x_t + u_t^\top R u_t\right], \tag{$\mathcal{LQR}_\star$}$$

subject to (1), for given x_0 and $Q \succ 0$ and $R \succ 0$[1].

In practical applications the dynamics (1) are often unknown. Hence we introduce a system identification scheme to estimate (1), supporting both the setting where the modes A_i, B_i are unknown and the setting where they are (partially) known. This information will be encoded through the use of a mode tensor $\boldsymbol{\mathcal{M}}$ described in §IV-A.

Besides the modes, characterizing the distribution of v_t is a more considerable challenge. We however are aided by the following fact. By linearity, the cost of ($\mathcal{LQR}_\star$) only involves second moments of the states $X_t = \mathbb{E}[x_t x_t^\top]$ and the inputs $U_t = \mathbb{E}[u_t u_t^\top]$. In this work we thus tackle LQR in terms of the moment dynamics:

$$X_{t+1} = \mathcal{E}(Z_t), \quad \text{with } Z_t = \mathbb{E}[z_t z_t^\top], \tag{2}$$

for augmented state $z_t = (x_t, u_t) \in \mathbb{R}^{n_z}$. These dynamics are easier to estimate, compared to (1), yet contain all information required to solve the stochastic LQR problem. We view (2) as an implicit definition for the operator $\mathcal{E}$ and argue in §IV-A how the fundamental operator $\mathcal{E}$ *(i)* is linear in Z_t, characterizing its matrix $\boldsymbol{E}$ in terms of the quantities (1), *(ii)* is *completely positive (CP)*, *(iii)* has a bilinear structure since $\boldsymbol{E}$ depends linearly on $V = \mathbb{E}[v_t v_t^\top]$. These properties serve as a central theme of this work.

Since we will never get a completely accurate estimate of the distribution it makes sense in safety critical scenarios to use data to construct a set of likely distributions $\mathscr{A}$ containing the true distribution with high probability (see the formal setup in [51]). Given $\mathscr{A}$ we then solve the DR LQR problem:

$$\underset{u_0,u_1,\dots}{\text{minimize}} \ \underset{\nu\in\mathscr{A}}{\text{maximize}} \ \mathbb{E}_\nu\left[\sum_{t=0}^{\infty} x_t^\top Q x_t + u_t^\top R u_t\right], \tag{$\overline{\mathcal{LQR}}$}$$

which upper bounds the cost of ($\mathcal{LQR}_\star$) with high probability. In this work, only the second moment $V = \mathbb{E}[v_t v_t^\top]$ affects the problem. So we formulate the ambiguity set over moments.

III. Preliminary results

We introduce the main tools used throughout our work here. Specifically this section presents operators over the space of symmetric matrices in §III-A and tensor algebra in §III-B.

A. Operators over symmetric matrices

As argued later, the second moment of the disturbance v is identified by solving a linear system of equations, where the variable of interest is a symmetric matrix.

Useful tools when dealing with such problems are the Kronecker product $\otimes$ [52, Def. 4.2.1], the vector operator $\mathrm{vec}(\cdot)$ [52, Def. 4.2.9] and its inverse $\mathrm{unvec}(\cdot)$. These are connected through the fundamental property [52, Lem. 4.3.1]

$$\mathrm{vec}(UXV^\top) = (V \otimes U)\mathrm{vec}(X), \tag{3}$$

for any $U \in \mathbb{R}^{m\times n}, V \in \mathbb{R}^{\ell\times k}$ and $X \in \mathbb{R}^{n\times k}$.

This property is often used to rewrite linear equations over matrices in terms of vectors. When $X \in \mathbb{S}^d$ however, we often run into non-invertible systems since $\mathrm{vec}(X) \in \mathbb{R}^{d^2}$ has unnecessary degrees of freedom. After all, a symmetric matrix is uniquely determined by $\mathfrak{d}_d := d(d+1)/2$ values.

To exploit this fact, we will use [53, Def. E.1]:

Definition III.1. *For any $X \in \mathbb{S}^d$, define* $\mathrm{vec_s} \colon \mathbb{S}^d \to \mathbb{R}^{\mathfrak{d}_d}$,

$$\begin{aligned}\mathrm{vec_s}(X) := ([X]_{11}, \sqrt{2}[X]_{21}, \dots, \sqrt{2}[X]_{d1},\\ [X]_{22}, \sqrt{2}[X]_{32}, \dots, \sqrt{2}[X]_{d2}, \dots, [X]_{dd})\end{aligned}$$

and let $Q_d \in \mathbb{R}^{\mathfrak{d}_d \times d}$ be such that $\mathrm{vec_s}(X) = Q_d \mathrm{vec}(X)$.

It is easy to verify that Q_d is a unique matrix for which $Q_d Q_d^\top = I_{\mathfrak{d}_d}$ [53, §E.2]. We can then introduce the symmetrized Kronecker product [53, Def. E.3],

Definition III.2. *For any $U, V \in \mathbb{R}^{m\times n}$, let*

$$V \otimes_{\mathrm{s}} U = \frac{1}{2} Q_m (U \otimes V + V \otimes U) Q_n^\top.$$

By definition, we have a generalization of (3) [53, Def. E.3]

$$\mathrm{vec_s}(ZXZ^\top) = (Z \otimes_{\mathrm{s}} Z)\mathrm{vec_s}(X). \tag{4}$$

[1] These assumptions are not necessary for solvability [33], yet are added here for convenience.

We will also consider more general linear operators on symmetric matrices $\mathcal{S}\colon \mathbb{S}^n \to \mathbb{S}^m$, which can be expressed for any $X \in \mathbb{S}^n$ as

$$\mathcal{S}(X) = \mathrm{unvec}_{\mathrm{s}}(\boldsymbol{S}\mathrm{vec}_{\mathrm{s}}(X)),$$

for some $\boldsymbol{S} \in \mathbb{R}^{\eth_m \times \eth_n}$. The set of such operators is denoted SL_n^m. The adjoint S^*, defined to satisfy $\mathrm{Tr}[\mathcal{S}(X)Y] = \mathrm{Tr}[X\mathcal{S}^*(Y)]$, has matrix $\boldsymbol{S}^\top$. Similar to [50] we will often need to bound the spectral norm of the image of such matrix operators. As such we introduce the following norm on SL_n^m:

$$\|\mathcal{S}\|_2 := \max_{X \in \mathbb{S}^d} \{\|\mathcal{S}(X)\|_2 \colon \|X\|_2 \leq 1\},$$

for which we have $\|\mathcal{S}(X)\|_2 \leq \|\mathcal{S}\|_2 \|X\|_2$ for any $X \in \mathbb{S}^n$. This norm is bounded in terms of the spectral norm of $\boldsymbol{S}$:

Lemma III.3. *Let $\mathcal{S} \in \mathrm{SL}_n^m$ with matrix $\boldsymbol{S}$. Then,*

$$\|\mathcal{S}\|_2 \leq \sqrt{m}\|\boldsymbol{S}\|_2.$$

Proof. Proof deferred to Appendix A. □

We are especially interested in linear operators that preserve positive semidefiniteness of the argument. That is, a map $\mathcal{S}\colon \mathbb{S}^n \to \mathbb{S}^m$ is *positive* iff $\mathcal{S}(X) \succeq 0$ for all $X \succeq 0$. More specifically we consider a specific type of positive maps:

Definition III.4. *A $\mathcal{S} \in \mathrm{SL}_n^m$ is* completely positive (CP) *[54] if, for some $A_i \in \mathbb{R}^{m \times n}$ with $i \in \mathbb{N}_{1:r}$,*

$$\mathcal{S}(X) = \sum_{i=1}^{r} A_i X A_i^\top. \tag{5}$$

The matrix of $\mathcal{S}$ is $\boldsymbol{S} = \sum_{i=1}^r (A_i \otimes_{\mathrm{s}} A_i)$. We write $\mathcal{S} \in \mathrm{CP}_n^m$.

We specialize Lem. III.3 for CP operators [43, Eq. 2.1]:

Lemma III.5. *Let $\mathcal{S} \in \mathrm{CP}_n^m$. Then, $\|\mathcal{S}\|_2 = \|\mathcal{S}(I)\|_2$.*

Like for linear dynamics, stability of CP operators can be defined in terms of the spectral radius.

Lemma III.6. *For any $\mathcal{S} \in \mathrm{CP}_n^m$ the following assertions are equivalent:*

(a) $\rho(\mathcal{S}) := \rho(\boldsymbol{S}) < 1$;

(b) $\|\mathcal{S}^t(X)\| \to 0$ *as* $t \to \infty$, $\forall X \succeq 0$.[2]

Proof. The proof is analogous to [32, Prop. 1]. □

The spectral radius of a CP map (5) is related to the *outer spectral radius* $\hat{\rho}(A_1, \cdots, A_r) := \sqrt{\rho(\sum_{i=1}^r A_i \otimes A_i)}$ for square matrices $\{A_i\}_{i=1}^r$. This radius bounds the *joint spectral radius*

$$\rho(A_1, \cdots, A_r) := \lim_{t \to \infty} \left\{ \sup_{1 \leq i_1, \ldots, i_t \leq r} \|A_{i_1} \ldots A_{i_t}\|_2^{1/t} \right\}.$$

Specifically [31, Cor. 1.5]:

Theorem III.7. *Consider the matrices $\{A_i\}_{i=1}^r$. Then*

$$\hat{\rho}(A_1, \ldots, A_r)/\sqrt{r} \leq \rho(A_1, \ldots, A_r) \leq \hat{\rho}(A_1, \ldots, A_r).$$

[2]The specific norm is left unspecified since the convergence would remain the same by equivalence of matrix norms [55, Ex. 5.6.P23].

Note that the joint spectral radius describes stability the dynamics of switched linear system. That is, dynamics $x_{t+1} = A_t x_t$ are stable for any sequence of matrices A_t from $\{A_i\}_{i=1}^r$ iff $\rho < 1$ [56, Cor. 1.1]. Therefore Thm. III.7 gives a strong relation between stability of such systems and stability of $\mathcal{S}$ (and as we will see later mean square stability of linear dynamics with multiplicative noise).

We finally consider a Lyapunov criterion for stability:

Proposition III.8. *Let $\mathcal{S} \in \mathrm{CP}_d^d$. Then*

(i) If $\rho(\mathcal{S}) < 1$, then $\forall H \succ 0$

$$\exists P \succeq 0 \colon P - \mathcal{S}^*(P) = H. \tag{6}$$

Moreover, P is unique, positive definite and $\mathrm{Tr}[PX] = \mathrm{Tr}[H \sum_{t=0}^\infty \mathcal{S}^t(X)]$, $\forall X \succeq 0$.

(ii) If (6) for some $H \succ 0$, then $\rho(\mathcal{S}) < 1$.

Proof. For linear systems with multiplicative noise this statement was shown in [32, Lem. 1-2]. The proof for CP operators is similar.

(i): Note that (6) iff

$$\exists P \succeq 0 \colon (I - \boldsymbol{S}^\top)\mathrm{vec}_{\mathrm{s}}(P) = \mathrm{vec}_{\mathrm{s}}(H),$$

since the matrix of $\mathcal{S}^*$ is $\boldsymbol{S}^\top$. When $\rho(\mathcal{S}) = \rho(\boldsymbol{S}) < 1$, then $(I - \boldsymbol{S}^\top)^{-1} = \sum_{t=0}^\infty (\boldsymbol{S}^\top)^t = \sum_{t=0}^\infty (\boldsymbol{S}^t)^\top$. So, we take $P = \sum_{t=0}^\infty (\mathcal{S}^t)^*(H)$. Then, $\forall X \succneqq 0$,

$$\mathrm{Tr}[PX] = \mathrm{Tr}[\textstyle\sum_{t=0}^\infty (\mathcal{S}^t)^*(H)X] = \mathrm{Tr}[H \sum_{t=0}^\infty \mathcal{S}^t(X)] > 0,$$

where the last inequality follows from $\sum_{t=0}^\infty \mathcal{S}^t(X) = X + \sum_{t=1}^\infty \mathcal{S}^t(X) \succneqq 0$ and $H \succ 0$, which implies $\mathrm{Tr}[HX] \succ 0$ by Lem. F.1. The same lemma then implies $P \succ 0$. Uniqueness then follows by invertibility of $(I - \boldsymbol{S})$.

(ii): We assume (6) holds. So, $\forall X_t \succeq 0$,

$$\mathrm{Tr}[X_t(P - \mathcal{S}^*(P))] = \mathrm{Tr}[(X_t - \mathcal{S}(X_t))P]$$
$$\mathrm{Tr}[X_t P] - \mathrm{Tr}[X_{t+1}P] = \mathrm{Tr}[X_t H],$$

with $X_{t+1} = \mathcal{S}(X_t)$. Telescoping over $t \in \mathbb{N}$ gives

$$\mathrm{Tr}[X_0 P] - \lim_{t \to \infty} \mathrm{Tr}[X_t P] = \mathrm{Tr}[H \textstyle\sum_{t=0}^\infty X_t], \ \forall X_0 \succeq 0.$$

Since $X_t \succeq 0$ and $P \succeq 0$ we have $\lim_{t\to\infty} \mathrm{Tr}[X_t P] \geq 0$. Hence $\infty > \mathrm{Tr}[X_0 P] \geq \mathrm{Tr}[H \sum_{t=0}^\infty X_t]$. Therefore the terms $\mathrm{Tr}[HX_t] \to 0$. By $H \succ 0$ and Lem. F.1 this implies $X_t \to 0$. Thus $\rho(\mathcal{S}) < 1$ by Lem. III.6. □

B. Tensor algebra

The dynamics in (1) inherently have bilinear structure. This warrants the use of multi-linear structures and specifically third order tensors. In this section we introduce tensors and show how they are used to model CP operators.

Let $\boldsymbol{\mathcal{T}} \in \mathbb{R}^{q_1 \times q_2 \times q_3}$ denote a generic third order tensor. A good introduction is given in [41], [57, §12.4]. We restate some parts here for completeness.

We index tensors using $[\boldsymbol{\mathcal{T}}]_{ijk} \in \mathbb{R}$. Subtensors are selected with *colon notation* [57, §12.4] (e.g. $[\boldsymbol{\mathcal{T}}]_{:jk} \in \mathbb{R}^{q_1}$, $[\boldsymbol{\mathcal{T}}]_{i:k} \in \mathbb{R}^{q_2}$ and $[\boldsymbol{\mathcal{T}}]_{::k} \in \mathbb{R}^{q_1 \times q_2}$). We take a *slice* along mode-n for $n \in \{1, 2, 3\}$ using $[\boldsymbol{\mathcal{T}}]_{n|i} \in \mathbb{R}^{\times_{i \neq n} q_i}$ (e.g., $[\boldsymbol{\mathcal{T}}]_{1|i} = [\boldsymbol{\mathcal{T}}]_{i::}$).

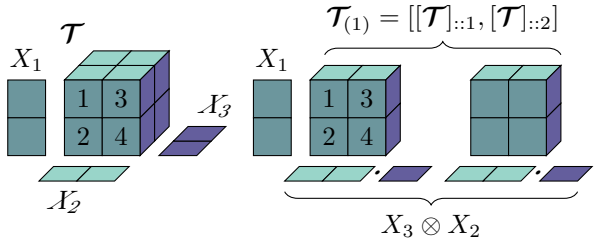


Fig. 1. Visualization of 1-mode matricization.

Next consider *n-mode matricizations* $\mathcal{T}_{(n)} \in \mathbb{R}^{q_n \times (\Pi_{i \neq n} q_i)}$:

$$\mathcal{T}_{(n)} := [\mathrm{vec}([\mathcal{T}]_{n|1}), \dots, \mathrm{vec}([\mathcal{T}]_{n|q_n})]^\top. \tag{7}$$

These matricizations extract the linearity of the multilinear tensor operator in one of its arguments. This is illuminated by introducing the *mode-n product*. This operation multiplies $\mathcal{T} \in \mathbb{R}^{q_1 \times q_2 \times q_3}$ along axis n with matrix $X \in \mathbb{R}^{p_n \times q_n}$ resulting in a tensor where q_n is replaced by p_n. The definition can be given in terms of matricizations:

$$(\mathcal{T} \times_n X)_{(n)} := X \mathcal{T}_{(n)}. \tag{8}$$

Since matricization is one-to-one, this uniquely defines the tensor $\mathcal{T} \times_n X$. Alternatively, e.g. for $n = 2$, we have $[\mathcal{T} \times_2 X]_{ijk} = \sum_{\ell=1}^{q_2} [X]_{j\ell} [\mathcal{T}]_{i\ell k}$ with $i \in \mathbb{N}_{1:q_1}$, $j \in \mathbb{N}_{1:p_2}$ and $k \in \mathbb{N}_{1:q_3}$. The extraction of linearity in one of the arguments is most explicit in (8). Similarly the *mode-n vector product* is $\mathcal{T} \mathbin{\overline{\times}}_n x = \sum_{j=1}^{q_n} [x]_j [\mathcal{T}]_{n|j}$. So the order of the tensor is reduced to $(\mathcal{T} \mathbin{\overline{\times}}_n x) \in \mathbb{R}^{\times_{i \neq n} q_i}$.

We can extend the single mode-n product to the multilinear one, similarly to (8).

Definition III.9. *For $X_i \in \mathbb{R}^{p_i \times q_i}$ with $i \in \{1,2,3\}$ let*

$$\begin{aligned} \mathcal{Y} &:= ((\mathcal{T} \times_3 X_3) \times_2 X_2) \times_1 X_1 \\ &=: [\![\mathcal{T}; X_1, X_2, X_3]\!] \in \mathbb{R}^{p_1 \times p_2 \times p_3}. \end{aligned} \tag{9}$$

When $p_n = 1$ which is the case when $X_n = x_n^\top$ for (column) vector $x_n \in \mathbb{R}^{q_n}$ then that axis is removed, reducing the order of $\mathcal{Y}$. That is, $\times_n$ is replaced by $\mathbin{\overline{\times}}_n$. As common in literature we omit the transposes in this case writing $[\![\mathcal{T}; x_1, x_2, x_3]\!]$.

Remark III.10. The order of operations in (9) is important when acting on vectors. After all since $\mathcal{T} \mathbin{\overline{\times}}_1 x_1 \in \mathbb{R}^{q_2 \times q_3}$ writing $(\mathcal{T} \mathbin{\overline{\times}}_1 x_1) \mathbin{\overline{\times}}_3 x_3$ is not defined. The Tucker operator helps to avoid such complexities, since we assume the order is reduced after *all* mode-n products have been completed.

We frequently use a generalization of (8) [42, Prop. 3.7]:

Proposition III.11. *For $\mathcal{Y} \in \mathbb{R}^{p_1 \times p_2 \times p_3}$ as in (III.9):*

$$\mathcal{Y}_{(n)} = X_n \mathcal{T}_{(n)} (X_j \otimes X_i)^\top, \tag{10}$$

for $n, i, j \in \{1,2,3\}$ with $j > i$ and $n \neq i \neq j$.

A graphical illustration of (10) is given in Fig. 1. Note how the values along the third axis are split up both by the matricization and the Kronecker product.

We can describe CP maps using tensors:

Lemma III.12. *Any $\mathcal{S} \in \mathrm{CP}_n^m$ can be expressed as:*

$$\mathcal{S}(X) = \boldsymbol{\mathcal{A}}_{(1)} (W \otimes X) \boldsymbol{\mathcal{A}}_{(1)}^\top \tag{11}$$

for some $\boldsymbol{\mathcal{A}} \in \mathbb{R}^{m \times n \times r}$ and $W \in \mathbb{S}_+^r$. We will refer to $\boldsymbol{\mathcal{A}}$ as the mode tensor *and W as the* parameter *matrix.*

Proof. Proof deferred to Appendix A. □

Remark III.13. Dependency of $\mathcal{S}$ on W in (11) can be made explicit by writing $\mathcal{S}(W; X)$. That is $\mathcal{S}(W; \cdot) \in \mathrm{SL}_n^m$. We distinguish the parameters of the map from its arguments using a semicolon. So $\mathcal{S}^*(W; P)$ denotes the adjoint w.r.t. the argument, not the parameter W.

We show that $\mathcal{S}(W; X)$ is also CP in its parameter W.

Corollary III.14. *Let $\mathcal{S}(W; X)$ as in (11). Then*

(i) $\forall X \succeq 0$, $\mathcal{S}(\cdot; X) \in \mathrm{CP}_n^m$;

(ii) $\mathcal{S}(\overline{W}; X) \succeq \mathcal{S}(W; X)$ *if* $\overline{W} \succeq W$.

Proof. Proof deferred to Appendix A. □

We will see later in the setting of multiplicative noise dynamics that Cor. III.14*(ii)* is useful for finding sufficient conditions for stability. Specifically, it allows constructing dynamics whose second moments dominate those of the true dynamics in the psd. sense.

There is a bilinear structure present in (11). This becomes clear when considering $\mathrm{Tr}[P\mathcal{S}(W; X)]$, which is linear in P, X and W. Hence, there is a version of Prop. III.11 for CP maps. To prove it we need some more preliminaries.

We define Kronecker products on $\mathcal{T} \in \mathbb{R}^{q_1 \times q_2 \times q_3}$ [58]:

$$[\mathcal{T} \otimes \mathcal{T}]_{(i+(j-1)q_1)::} = [\mathcal{V}]_{j::} \otimes [\mathcal{V}]_{i::} \text{ for } i, j \in \mathbb{N}_{1:q_1}. \tag{12}$$

Also let $\mathcal{V} \otimes_{\mathrm{s}} \mathcal{V} = [\![\mathcal{V} \otimes \mathcal{V}; Q_{q_1}, Q_{q_2}, Q_{q_3}]\!]$.

Next we state a generalization of a property for the Kronecker product between matrices, which is used to rewrite CP operators.

Lemma III.15. *Suppose $\mathcal{T} \in \mathbb{R}^{q_1 \times q_2 \times q_3}$ and $X_i \in \mathbb{R}^{p_i \times q_i}$ for $i \in \{1,2,3\}$ and let $\mathcal{Y} = [\![\mathcal{T}; X_1, X_2, X_3]\!]$. Then*

$$\mathcal{Y} \otimes_{\mathrm{s}} \mathcal{Y} = [\![\mathcal{T} \otimes_{\mathrm{s}} \mathcal{T}; X_1 \otimes_{\mathrm{s}} X_1, X_2 \otimes_{\mathrm{s}} X_2, X_3 \otimes_{\mathrm{s}} X_3]\!]$$

Proof. Proof deferred to Appendix A. □

Corollary III.16. *Let $\mathcal{S}(W; X)$ be CP and parameterized as in (11) with some $\mathcal{V}$. Then $\mathcal{S}^*(W; X) = \mathcal{V}_{(2)} (W \otimes X) \mathcal{V}_{(2)}^\top$. Moreover its matrix $\mathcal{S} = (\mathcal{V} \otimes_{\mathrm{s}} \mathcal{V}) \mathbin{\overline{\times}}_3 \mathrm{vec}_{\mathrm{s}}(W)$.*

Proof. Proof deferred to Appendix A. □

IV. Nominal Setting

We consider the nominal case for multiplicative noise, where the distribution is known. We describe properties of the (CP) operator $\mathcal{E}$ mentioned in (2), which fully describes the second moment dynamics of (1) and argue below how identifying $\mathcal{E}$ is sufficient to describe both stability and quadratically optimal control.

A. The model tensor

We first illustrate how the previous tools interface with multiplicative noise. Consider the true mode tensor $\boldsymbol{\mathcal{V}} \in \mathbb{R}^{n_x \times n_z \times n_v}$, which is populated as $[\boldsymbol{\mathcal{V}}]_{::i} = [A_i, B_i]$ for $i \in \mathbb{N}_{1:n_v}$ using the modes in (1). This splits up the dynamics into two unknowns. The tensor $\boldsymbol{\mathcal{V}}$ and the distribution of v_t.

To encode known information about $\boldsymbol{\mathcal{V}}$ we introduce the mode tensor $\boldsymbol{\mathcal{M}} \in \mathbb{R}^{n_x \times n_z \times n_w}$ and the auxiliary random vector $w_t \in \mathbb{R}^{n_w}$. When the modes are known, we can simply pick $\boldsymbol{\mathcal{M}} = \boldsymbol{\mathcal{V}}$ and $w_t = v_t$. However, when $\boldsymbol{\mathcal{V}}$ is only partially known we can still select a $\boldsymbol{\mathcal{M}}$ that has sufficient modeling capacity. Specifically we would like that for any realization of the disturbance v_t and hence of the random matrices $A(v_t)$ and $B(v_t)$, there exists a realization of w_t producing the same matrices by the modes encoded in $\boldsymbol{\mathcal{M}}$.

A sufficient assumption for this property is:

Assumption I (Model Equivalence)**.** *Consider $\boldsymbol{\mathcal{V}} \in \mathbb{R}^{n_x \times n_z \times n_v}$ with $n_z := n_x + n_u$ s.t. $[\boldsymbol{\mathcal{V}}]_{::i} = [A_i, B_i]$ for $i \in \mathbb{N}_{1:n_v}$ and $\{A_i, B_i\}_{i=1}^{n_v}$ as in (1).*

Assume $\boldsymbol{\mathcal{M}} \in \mathbb{R}^{n_x \times n_z \times n_w}$ with $n_w \in \mathbb{N}$ is such that

$$\operatorname{rk}(\boldsymbol{\mathcal{M}}_{(3)}) = \operatorname{rk}([\boldsymbol{\mathcal{M}}_{(3)}; \boldsymbol{\mathcal{V}}_{(3)}]) = n_w.$$

Here $\boldsymbol{\mathcal{M}}_{(3)} \in \mathbb{R}^{n_w \times n_x n_z}$ and $\boldsymbol{\mathcal{V}}_{(3)} \in \mathbb{R}^{n_v \times n_x n_z}$ denote the 3-mode matricizations as in (7).

Note that $\boldsymbol{\mathcal{M}}$ can be partitioned into $\boldsymbol{\mathcal{A}} := [\boldsymbol{\mathcal{M}}]_{:,1:n_x,:}$ and $\boldsymbol{\mathcal{B}} := [\boldsymbol{\mathcal{M}}]_{:,n_x+1:n_z,:}$ to get a generalization of A and B for discrete linear time invariant systems. When $\boldsymbol{\mathcal{V}} = \boldsymbol{\mathcal{M}}$ we have $[\boldsymbol{\mathcal{A}}]_{::i} = A_i$ and $[\boldsymbol{\mathcal{B}}]_{::i} = B_i$ for $i \in \mathbb{N}_{1:n_v}$.

Given Asm. I and this partitioning, we can show that

$$\begin{aligned} x_{t+1} &= [\![\boldsymbol{\mathcal{A}}; I_{n_x}, x_t, w_t]\!] + [\![\boldsymbol{\mathcal{B}}; I_{n_x}, u_t, w_t]\!] \\ &= [\![\boldsymbol{\mathcal{M}}; I_{n_x}, z_t, w_t]\!], \end{aligned} \tag{13}$$

describes the same dynamics as (1), given a suitable choice of w_t. We introduced the augmented state $z_t = (x_t, u_t) \in \mathbb{R}^{n_z}$. We formalize equivalence of (13) and (1) below.

Lemma IV.1. *When Asm. I holds and $w_t = (\boldsymbol{\mathcal{M}}_{(3)}^\top)^\dagger \boldsymbol{\mathcal{V}}_{(3)}^\top v_t$, $\forall t \in \mathbb{N}$, then (13) produces the same trajectories as (1).*

Proof. By Asm. I, $\boldsymbol{\mathcal{M}}_{(3)}^\top w_t = \boldsymbol{\mathcal{V}}_{(3)}^\top v_t$. By using (8) on both sides and reshaping we have $(\boldsymbol{\mathcal{M}} \mathbin{\overline{\times}}_3 w_t) = (\boldsymbol{\mathcal{V}} \mathbin{\overline{\times}}_3 v_t)$. By the definitions of $\mathbin{\overline{\times}}_3$, $\boldsymbol{\mathcal{V}}$ and $A(v)$, $B(v)$ we have respectively

$$\begin{aligned} \boldsymbol{\mathcal{V}} \mathbin{\overline{\times}}_3 v &= \textstyle\sum_{i=1}^{n_v} [v]_i [\boldsymbol{\mathcal{V}}]_{::i} \\ &= \textstyle\sum_{i=1}^{n_v} [v]_i [A_i, B_i] = [A(v), B(v)]. \end{aligned}$$

By definition of the Tucker product $[\![\boldsymbol{\mathcal{M}}; I_{n_x}, z_t, w_t]\!] = (\boldsymbol{\mathcal{M}} \mathbin{\overline{\times}}_3 w_t) \mathbin{\overline{\times}}_2 z_t = [A(v_t), B(v_t)] z_t = A(v_t) x_t + B(v_t) u_t$. So the produced trajectories will be the same. □

Example IV.2. We illustrate the need for Asm. I. Assume $n_x = 2$, $n_u = 1$ and $n_v = 3$ with true modes:

$$A_1 = \begin{bmatrix} 2 & 0 \\ 1 & 0 \end{bmatrix}, A_2 = \begin{bmatrix} 0 & 3 \\ 0 & 0 \end{bmatrix}, A_3 = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}, B_3 = \begin{bmatrix} 0 \\ 2 \end{bmatrix},$$

and $B_1 = B_2 = 0$. In this case $\boldsymbol{\mathcal{V}}_{(3)}$ is given as

$$\left[\begin{array}{cccc|cc} 2 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 3 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 2 \end{array}\right] = \left[\begin{array}{c|c} \operatorname{vec}^\top(A_1) & \operatorname{vec}^\top(B_1) \\ \operatorname{vec}^\top(A_2) & \operatorname{vec}^\top(B_2) \\ \operatorname{vec}^\top(A_3) & \operatorname{vec}^\top(B_3) \end{array}\right]$$

Say a designer, through some erroneous prior analysis, assumed that the top-left component in A_1 was equal to zero instead of two. This corresponds to the following mode tensor:

$$\boldsymbol{\mathcal{M}}_{(3)} = \left[\begin{array}{cccc|cc} \mathbf{0} & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 3 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 2 \end{array}\right].$$

for this choice Asm. I fails and so will Lem. IV.1. To see this, take $v_t = (1, 0, 0)$. Then $\operatorname{vec}([A(v), B(v)]) = (2, 1, 0, 0, 0, 0)$. Observe that such a vector cannot be constructed by any linear combination of the rows of $\boldsymbol{\mathcal{M}}_{(3)}$. So there is no distribution for w_t that reproduces the true dynamics.

Informally Asm. I requires that the span of $\boldsymbol{\mathcal{M}}_{(3)}^\top$ should be sufficiently large to model the outputs of $A(v)$ and $B(v)$ given by realizations of $v \in \mathbb{R}^{n_v}$.

Note that this is always possible, by taking $n_w = n_x(n_x + n_u)$ and $\boldsymbol{\mathcal{M}}_{(3)} = I_{n_w}$. We refer to this as the *model-free basis*, since no prior structural assumptions are imposed on the dynamics. If we know the *true modes* A_i and B_i, we can simply select $\boldsymbol{\mathcal{M}} = \boldsymbol{\mathcal{V}}$. As such, $\boldsymbol{\mathcal{M}}$ is used to introduce prior information into the dynamics.

Since the sequence v_t is iid, the sequence w_t as defined in Lem. IV.1 is iid. too. We thus omit t as the distribution remains constant. We can then use the transformation in Lem. IV.1 to find the true $W = \mathbb{E}[ww^\top]$ in terms of $V = \mathbb{E}[vv^\top]$:

$$W = (\boldsymbol{\mathcal{M}}_{(3)}^\top)^\dagger \boldsymbol{\mathcal{V}}_{(3)}^\top V \boldsymbol{\mathcal{V}}_{(3)} \boldsymbol{\mathcal{M}}_{(3)}^\dagger. \tag{14}$$

This allows translating the *ground truth* dynamics to another mode tensor and will aid in validation of our method during the experiments in §VIII.

We next illustrate the connection with CP operators by considering the dynamics of $X_t = \mathbb{E}[x_t x_t^\top]$ as in (2). By Prop. III.11 $[\![\boldsymbol{\mathcal{M}}; I, z_t, w_t]\!] = \boldsymbol{\mathcal{M}}_{(1)}(w_t \otimes z_t)$. Taking the outer product and the expectation allows us to derive:

$$X_{t+1} = \mathcal{E}(W; Z_t) := \boldsymbol{\mathcal{M}}_{(1)}(W \otimes Z_t)\boldsymbol{\mathcal{M}}_{(1)}^\top, \tag{15}$$

where $X_t = \mathbb{E}[x_t x_t^\top]$, $Z_t = \mathbb{E}[z_t z_t^\top]$ and $W = \mathbb{E}[w_t w_t^\top]$. We write $\mathcal{E}(\cdot)$ when the specific value of W is unimportant. The operator $\mathcal{E} \colon \mathbb{S}^{n_z} \to \mathbb{S}^{n_x}$ is clearly CP by Lem. III.12. It is linear with its matrix characterized by Cor. III.16. Moreover it has a bilinear structure, depending linearly on W (and by (14) on V). So we confirmed the claims from the problem statement. We provide a brief illustration of how $\mathcal{E}$ encodes the dynamics and then discuss nominal stability and LQR, before introducing the system identification scheme.

Example IV.3. Continuing on from Ex. IV.2, we now assume $\boldsymbol{\mathcal{M}} = \boldsymbol{\mathcal{V}}$ and $W = I$. We will examine $\mathcal{E}$. First note that

$$\boldsymbol{\mathcal{M}}_{(1)} = \left[\begin{array}{ccc|ccc|ccc} 2 & 0 & 0 & 0 & 3 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 2 \end{array}\right] = \left[\begin{array}{cc|cc|cc} A_1 & B_1 & A_2 & B_2 & A_3 & B_3 \end{array}\right].$$

If we partition $Z_t = [X_t, V_t; V_t^\top, U_t]$ with $X_t = \mathbb{E}[x_t x_t^\top]$, $V_t = \mathbb{E}[x_t u_t^\top]$ and $U_t = \mathbb{E}[u_t u_t^\top]$ then

$$\begin{aligned}\mathcal{E}(Z_t) &= \textstyle\sum_{i=1}^3 \left[A_i,\, B_i\right] Z_t \left[A_i^\top;\, B_i^\top\right] \\ &= \textstyle\sum_{i=1}^3 A_i X_t A_i^\top + A_i V_t B_i^\top + B_i V_t^\top A_i^\top + B_i U_t B_i^\top.\end{aligned}$$

For $u_t = 0$ this reduces to $\sum_{i=1}^3 \mathbb{E}[(A_i x_t)(A_i x_t)^\top]$.

The matrix associated with $\mathcal{E}$ as described in Cor. III.16 is given as follows:

$$\boldsymbol{E} = Q_2 \begin{bmatrix} 4 & 0 & 0 & 0 & 9 & 0 & 0 & 0 & 0 \\ 2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 1 & 2 & 0 & 2 & 4 \end{bmatrix} Q_2^\top.$$

with

$$Q_2 = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1/\sqrt{2} & 1/\sqrt{2} & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix},$$

which is added as prescribed in the definition of $\otimes_{\mathrm{s}}$.

B. Stability

We consider stability of the autonomous case of (13):

$$x_{t+1} = [\![\boldsymbol{\mathcal{A}}; x_t, w_t]\!]. \tag{16}$$

Stability of stochastic systems is intimately related to convergence of random variables, which can be defined in several ways [59, §7.2]. We specifically consider convergence to zero of the second moment [59, Def. 1.a],

Definition IV.4. *Consider the autonomous system* (16)*. It is considered* mean-square stable (mss) *if*

$$\mathbb{E}[x_t x_t^\top] \to 0 \quad \textit{as} \quad t \to \infty.$$

From the definition, mss is equivalent to the stability of

$$X_{t+1} = \mathcal{F}(X_t) := \boldsymbol{\mathcal{A}}_{(1)}(W \otimes X_t)\boldsymbol{\mathcal{A}}_{(1)}^\top,$$

where $X_t = \mathbb{E}[x_t x_t^\top]$ as before and $\mathcal{F}$ is the autonomous version of $\mathcal{E}$. Clearly this system is stable for all $X_0 \succeq 0$ iff $\rho(\mathcal{F}) < 1$ by Lem. III.6. Therefore:

Corollary IV.5. *The system* (16) *is mss iff* $\rho(\mathcal{F}) < 1$.

We can verify stability with the Lyapunov condition Prop. III.8 and, as mentioned earlier, stability of (16) is equivalent with robust stability of the switching system with modes $\{A_i\}_{i=1}^{n_w}$ in the sense of Thm. III.7. This equivalence was observed before in [30] and [27], yet not fully characterized.

C. Linear Quadratic Regulation

The goal of this section is twofold *(i)* we argue how LQ control of CP dynamics is equivalent to ($\mathcal{LQR}_\star$); and *(ii)* we show how the optimal policy is determined by solving a SDP, generalizing the result of [60].

We start by stating the CP equivalent of ($\mathcal{LQR}_\star$)

$$\begin{aligned} \underset{Z_t}{\text{minimize}} \quad & \sum_{t=0}^{\infty} \mathrm{Tr}[Z_t H] \\ \text{subj. to} \quad & Z_t = \begin{bmatrix} X_t & V_t \\ V_t^\top & U_t \end{bmatrix} \succeq 0, \\ & X_{t+1} = \mathcal{E}(Z_t), \quad \forall t \in \mathbb{N}, \end{aligned} \tag{$\mathcal{LQR}_{\mathrm{cp}}$}$$

with $H := \mathrm{diag}(Q, R) \succ 0$. The psd. constraint on Z_t ensures that it acts like the second moment of a random vector.

We show that ($\mathcal{LQR}_\star$) is equivalent to ($\mathcal{LQR}_{\mathrm{cp}}$)

Proposition IV.6. *Let* $\mathrm{Val}(\cdot)$ *denote the optimal value. Then*

$$\mathrm{Val}(\mathcal{LQR}_\star) = \mathrm{Val}(\mathcal{LQR}_{\mathrm{cp}}),$$

given Asm. I. Moreover sequence $(Z_t)_{t\in\mathbb{N}}$ *is feasible for* ($\mathcal{LQR}_{\mathrm{cp}}$) *iff there is a feasible sequence in* ($\mathcal{LQR}_\star$) $u_t = K_t x_t + \delta_t$ *s.t.* $Z_t = \mathbb{E}[z_t z_t^\top]$ *where* $z_t = (x_t, u_t)$*. Here* $K_t \in \mathbb{R}^{n_x \times n_u}$ *and* $\delta_t \in \mathbb{R}^{n_u}$ *is a random vector* $\forall t \in \mathbb{N}$.

Proof. The proof is deferred to Appendix B. □

To find the exact solution of ($\mathcal{LQR}_{\mathrm{cp}}$), similarly to usual LQR, we need to consider a Riccati equation:

$$\mathcal{R}(P) := Q + \mathcal{F}^*(P) - (\mathcal{H}^*)^\top(P)(R + \mathcal{G}^*(P))^{-1}\mathcal{H}^*(P). \tag{17}$$

where $\mathcal{F}^*(P) = \boldsymbol{\mathcal{A}}_{(2)}(W \otimes P)\boldsymbol{\mathcal{A}}_{(2)}^\top$, $\mathcal{G}^*(P) = \boldsymbol{\mathcal{B}}_{(2)}(W \otimes P)\boldsymbol{\mathcal{B}}_{(2)}^\top$ and $\mathcal{H}^*(P) = \boldsymbol{\mathcal{B}}_{(2)}(W \otimes P)\boldsymbol{\mathcal{A}}_{(2)}^\top$ with tensors $\boldsymbol{\mathcal{A}}$ and $\boldsymbol{\mathcal{B}}$ as in (13). We use adjoint CP operators to be consistent with Cor. III.16.

Remark IV.7. The operators $\mathcal{F}^*$, $\mathcal{H}^*$ and $\mathcal{G}^*$ form the blocks of $\mathcal{E}^*$ when it is partitioned similarly to Z_t (cf. Lem. B.1).

In the deterministic setting with $n_w = 1$ and $w_t = 1$ (so $\mathbb{E}[W] = 1$), we have $\mathcal{F}^*(P) = A_1^\top P A_1$, $\mathcal{H}^*(P) = B_1^\top P A_1$ and $G^*(P) = B_1^\top P B_1$. Hence we recover the classical Riccati operator as a special case.

We introduce the *closed-loop operators* as

$$\Pi_K(X) := [I; K] X [I, K^\top],\ \mathcal{E}_K(X) := \mathcal{E}(\Pi_K(X)), \tag{18}$$

The policy $Z_t = \Pi_K(X_t)$ is recovered when taking $u_t = K x_t$ in the equivalent ($\mathcal{LQR}_\star$) (cf. Prop. IV.6).

We can then state the following theorem.

Theorem IV.8. *For* ($\mathcal{LQR}_{\mathrm{cp}}$)*, assume* $\exists K\colon \rho(\mathcal{E}_K(\cdot)) < 1$ *and* $H := \mathrm{diag}(Q, R) \succ 0$*. Then* $Z_t = \Pi_{K_\star}(X_t)$ *is the optimal policy, with* $K_\star = -(R + \mathcal{G}^*(P_\star))^{-1}\mathcal{H}^*(P_\star)$ *and* $P_\star \succ 0$ *solves the Riccati equation* $P_\star = \mathcal{R}(P_\star)$*. Moreover*

(i) $\rho(\mathcal{E}_{K_\star})) < 1$*;*

(ii) $\mathrm{Val}(\mathcal{LQR}_{\mathrm{cp}}) = \mathrm{Tr}[P_\star X_0]$*;*

Proof. See Appendix B for the full proof. □

Example IV.9. Let us consider a deterministic system with $n_w = 1$ and

$$A_1 = \begin{bmatrix} 1 & 0.1 \\ 0 & 1 \end{bmatrix}, B_1 = \begin{bmatrix} 0 \\ 0.1 \end{bmatrix}.$$

This will help us understand how ($\mathcal{LQR}_{\text{cp}}$) generalizes the usual notions of LQR. We pick $H = \text{diag}(1, 1, 0.1)$. Then the solution described by Thm. IV.8 is approximately:

$$K = [-1.75, -2.64].$$

One can easily verify that this is the usual LQR solution in the deterministic setting. Let's spend some time noting what happens when we set $X_0 = x_0 x_0^\top$. In that case, for some u_0, let $Z_0 = (x_0, u_0)(x_0, u_0)^\top$. Then (similarly to Ex. IV.3):

$$\begin{aligned}\mathcal{E}(Z_0) = A_1 x_0 x_0^\top A_1^\top + A_1 x_0 u_0^\top B_1^\top \\ + B_1 u_0 x_0^\top A_1^\top + B_1 u_0 u_0^\top B_1^\top.\end{aligned}$$

It is not difficult to verify that this equals $x_1 x_1^\top$ for $x_1 = A_1 x_0 + B_1 u_0$. We could repeat this for each time step noting that $X_t = x_t x_t^\top$ and $Z_t = (x_t, u_t)(x_t, u_t)^\top$ in ($\mathcal{LQR}_{\text{cp}}$). Trivially we have $Z_t \succeq 0$, so the trajectory is feasible. Moreover the cost becomes $\text{Tr}[Z_t H] = x_t^\top I x_t + 0.1 u_t^\top I u_t$. So we can see that the classical LQR controller is at least a feasible solution to ($\mathcal{LQR}_{\text{cp}}$) where we then have

$$\begin{aligned}Z_t = \Pi_K(X_t) = \begin{bmatrix} x_t x_t^\top & x_t x_t^\top K^\top \\ K x_t x_t^\top & K x_t x_t^\top K^\top \end{bmatrix} \\ = (x_t, Kx_t)(x_t, Kx_t)^\top.\end{aligned}$$

Note that we have shown in Thm. IV.8 that such policies are in fact optimal. In terms of classical LQR this implies that, a random policy $u_t = Kx_t + \delta_t$ for some random δ_t will not improve performance.

Even though we directly tackled second moment dynamics in Thm. IV.8, the optimal policy $Z_t = \Pi_{K_\star}(X_t)$ is realizable for multiplicative noise dynamics using $u_t = K_\star x_t$. So this controller is also optimal for ($\mathcal{LQR}_\star$) as is well known (cf. [13, Prop. 3]) and achieves the equality in Prop. IV.6.

Thus we can conclude that LQR of CP dynamics is (i) identical to LQR of multiplicative noise dynamics; (ii) a natural generalization of classical LQR; and (iii) a relaxation for LQR of switching systems in the sense of Thm. III.7.

The final question is how to find a $P_\star$ satisfying (17). Numerical solution of such equations is considered in detail in [44]. We consider a reformulation as a SDP, akin to [60].

Theorem IV.10. *Let $P_\star$ solve the Riccati equation (17). Then $P_\star$ is a maximizer of*

$$\begin{aligned}&\underset{P \succeq 0}{\text{maximize}} \quad \text{Tr}[P X_0] \\ &\text{subj. to} \quad \begin{bmatrix} P - \mathcal{F}^*(P) & -(\mathcal{H}^*)^\top(P) \\ -\mathcal{H}^*(P) & -\mathcal{G}^*(P) \end{bmatrix} \preceq \begin{bmatrix} Q & \\ & R \end{bmatrix},\end{aligned}$$

with $\mathcal{F}^*$, $\mathcal{H}^*$ and $\mathcal{G}^*$ as defined below (17). Moreover, when $X_0 \succ 0$, then $P_\star$ is the unique solution.

Proof. The proof is analogous to the one in [60] and [13, Prop. 3]. The full proof is deferred to Appendix B. □

V. Identifying the second moment

Our goal in this work is to generalize Thm. IV.8 to the data-driven DR setting ($\mathcal{LQR}$), while retaining stability for a linear controller. This problem was considered in [13], for sub-Gaussian normalized disturbance ξ (i.e. $\xi := \text{Var}[w]^{-1/2}(w - \mathbb{E}[w])$). Therefore [13] supports among others Gaussian w or bounded ξ. This setup had two main limitations: *(i)* bounded w does not imply bounded ξ, so the link with classical robust control (as exploited for additive noise in [51]) is lost; *(ii)* the setup is inapplicable when only state measurements are available. This section resolves these issues.

A. Measurement model

Considering the dynamics (13), the goal is to estimate the *true*[3] $W_\star := \mathbb{E}[ww^\top]$. Unlike in Lem. V.3 we do not assume access to samples $(w_i)_{i=1}^N$. Instead we measure a sequence of N independent samples $(x_{i+1}, z_i)_{i=1}^N$, with $z_i = (x_i, u_i)$ the augmented state, satisfying the measurement model:

$$x_{i+1} = [\![\boldsymbol{\mathcal{M}}; I_{n_x}, z_i, w_i]\!], \tag{19}$$

where $(w_i)_{i=1}^N$ denotes the noise sequence.

There are two important properties that (19) should satisfy: *(i)* the 'span' of the *(mixed) mode tensor* $\boldsymbol{\mathcal{M}}$ should describe the true dynamics (i.e. Asm. I); and *(ii)* the data should be sufficiently rich to render W observable.

Model tensor: To know that the measurement model is valid, we need Asm. I since in that case Lem. IV.1 holds. Importantly, this also means that it is sufficient to identify $\mathbb{E}[ww^\top]$ to describe the moment dynamics (15).

Informativity: To guarantee a property akin to *persistency of excitation*, we assume that the support of z_i is not degenerate as is formalized in the following assumption.

Assumption II (Nondegenerate support)**.** *We assume that $\{z_i\}_{i \in \mathbb{N}_{1:N}}$ is a sequence of i.i.d. copies of a random vector z, with a distribution that is dominated by the Lebesgue measure.*

The result is that, by Lem. D.1, if $N \geq n_z$,

$$\mathbb{P}[\text{rk}([z_1, \ldots, z_N]) = n_z] = 1. \tag{20}$$

Remark V.1. Independence is guaranteed for each i by either *(i)* taking a known random $z_i = (x_0, u_0)$, updating (1) once and taking $x_{i+1} = x_1$; or *(ii)* taking a known x_0, updating (1) T times for random u_t and taking $x_{i+1} = x_T$ and $z_i = (x_{T-1}, u_{T-1})$. We refer to *(ii)* as *rollout* as in [7], [37] and to *(i)* as *random initialization*.

For rollout it is difficult to verify (20). Usually a notion of controllability and persistency of excitation is employed. Full generalizations of these concepts for multiplicative noise do not exist to our knowledge. Note however that the event in (20) is related to the zeros of a polynomial. So an argument similar to the one in Prop. V.2 below is applicable. In practice one can sample u_t from a distribution dominated by the Lebesgue measure and test whether (20) holds. The result should generalize for more data with probability one.

[3] We will henceforth denote the *true* values using a star subscript.

B. Noise observability

In some cases, the disturbances can be inferred exactly from state measurements. Specifically (19) is rewritten as:

$$x_{i+1} = (\boldsymbol{\mathcal{M}} \,\overline{\times}_2\, z_i) w_i = (\textstyle\sum_{j=1}^{n_z} [z_i]_j [\boldsymbol{\mathcal{M}}]_{:j:}) w.$$

Thus if $(\boldsymbol{\mathcal{M}} \,\overline{\times}_2\, z_i)$ is left invertible, we can uniquely identify $(w_i)_{i=1}^N$ from $(x_{i+1}, z_i)_{i=1}^N$. Invertibility of $(\boldsymbol{\mathcal{M}} \,\overline{\times}_2\, z_i)$ however is analogous to invertibility of a linear subspace of matrices (or linear forms), which is an unsolved problem in general (cf. [61] for 3×3 matrices). We have the following:

Proposition V.2. *Let z denote a random vector with non-degenerate support (cf. Asm. II) and $n_x \geq n_w$. Then, $(\boldsymbol{\mathcal{M}} \,\overline{\times}_2\, z)$ is left invertible either with probability one or zero.*

Proof. A high-level proof is given in [62, p.2] in the setting of invertibility of linear forms. □

This suggests that it is often sufficient to sample $(\boldsymbol{\mathcal{M}} \,\overline{\times}_2\, z)$ for a random z and check its invertibility. The conclusion will then generalize with probability one. When successful, we can determine w_i exactly and we can use:

Lemma V.3. *Let $\mathcal{W} = \{w \in \mathbb{R}^{n_w} : \|w\|_2 \leq r_w\}$. Assume we have a set of i.i.d. samples $\{w_i\}_{i \in \mathbb{N}_{1:M}}$ of a random vector w and let $\hat{W} := \sum_{i=1}^N w_i w_i^\top / N$. Then*

$$\mathbb{P}[\|\hat{W} - W_\star\|_2 \leq \beta_W] \geq 1 - \delta$$

when $\beta_W = r_w^2 \sqrt{2\ln(2n_w/\delta)/N}$.

Proof. The result follows from a direct application of Lem. C.1 by noting that the terms in the error satisfy $-r_w^2 I \preceq -w_i w_i^\top \preceq W_\star - w_i w_i^\top \preceq W_\star \preceq r_w^2 I$. □

We also want to estimate $W_\star$ in the other case (i.e., when $n_x < n_w$). For example, for a model-free basis we have $n_w = n_x n_z > n_x$. To do so we design least squares estimators.

C. Least squares estimator

To get a *least square (LS)* estimator, we need to convert (19) into an expression linear in $w_i w_i^\top$. We can do so by using Lem. III.15 on (19), which gives:

$$x_{i+1} \otimes_{\mathrm{s}} x_{i+1} = [\![\boldsymbol{\mathcal{M}} \otimes_{\mathrm{s}} \boldsymbol{\mathcal{M}}; z_i \otimes_{\mathrm{s}} z_i, w_i \otimes_{\mathrm{s}} w_i]\!]. \tag{21}$$

Note $\mathbb{E}[w_i \otimes_{\mathrm{s}} w_i] = \mathbb{E}[\mathrm{vec}_{\mathrm{s}}(w_i w_i^\top)] =: \mathrm{vec}_{\mathrm{s}}(W_\star)$. We can write this in terms of linear equations by expanding the Tucker operator (cf. Def. III.9):

$$\begin{aligned} x_{i+1} \otimes_{\mathrm{s}} x_{i+1} = (\boldsymbol{\mathcal{W}} \,\overline{\times}_2\, (z_i \otimes_{\mathrm{s}} z_i))\mathrm{vec}_{\mathrm{s}}(W_\star) \\ + (\boldsymbol{\mathcal{W}} \,\overline{\times}_2\, (z_i \otimes_{\mathrm{s}} z_i))\eta_i, \end{aligned}$$

with $\boldsymbol{\mathcal{W}} = \boldsymbol{\mathcal{M}} \otimes_{\mathrm{s}} \boldsymbol{\mathcal{M}}$ and noise $\eta_i = w_i \otimes_{\mathrm{s}} w_i - \mathrm{vec}_{\mathrm{s}}(W_\star)$.

As in usual least squares, stacking the equations gives:

$$Y_N = \boldsymbol{Z}_N \mathrm{vec}_{\mathrm{s}}(W_\star) + E_N, \tag{22}$$

where we stack that data as follows

$$Y_N = (x_2 \otimes_{\mathrm{s}} x_2, \ldots, x_{N+1} \otimes_{\mathrm{s}} x_{N+1})$$

$$\boldsymbol{Z}_N = [\boldsymbol{\mathcal{W}} \,\overline{\times}_2\, (z_1 \otimes_{\mathrm{s}} z_1); \ldots; \boldsymbol{\mathcal{W}} \,\overline{\times}_2\, (z_N \otimes_{\mathrm{s}} z_N)] \tag{23a}$$

$$E_N = ([\![\boldsymbol{\mathcal{W}}; z_1 \otimes_{\mathrm{s}} z_1, \eta_1]\!], \ldots, [\![\boldsymbol{\mathcal{W}}; z_N \otimes_{\mathrm{s}} z_N, \eta_N]\!]). \tag{23b}$$

Note that $\mathbb{E}[\eta_i] = 0$ so $\mathbb{E}[E_N] = 0$, which suggests the LS estimate:

$$\mathrm{vec}_{\mathrm{s}}(\hat{W}) = \boldsymbol{Z}_N^\dagger Y_N. \tag{24}$$

A similar estimator to $\hat{W}$ is used in compressed covariance sensing [46], where it is observed that it acts as a good heuristic for the *maximum likelihood estimator* when w is Gaussian. Noting that each element of the vector equation (21) constitutes a bilinear form of z_i, reveals a connection with bilinear estimation [47]. There, adjusted LS estimators exist that are consistent even when z_i is perturbed by random noise. For simplicity we consider exact state measurements here, which enables the use of ordinary LS, for which it is convenient to derive concentration inequalities.

D. Error analysis

We analyse the error of the LS estimator (24) under Asm. I–II. The discussion somewhat differs from the classical case since the estimate (24) is biased. This is inconsequential, as we are interested in estimating the second moment dynamics. Specifically, we show that $\mathcal{E}(\hat{W}; \cdot)$ as an estimate of the second moment dynamics $\mathcal{E}_\star := \mathcal{E}(W_\star; \cdot)$ is unbiased. This phenomenon is common in system identification of multiplicative noise and was observed previously in [38], [40]. As this complicates the error analysis, instead of giving all details, we only provide the intuition here. Formal statements, to which we refer in the text, are deferred to Appendix E.

The overall estimation error is given as

$$\mathrm{vec}_{\mathrm{s}}(\hat{W} - W_\star) = (I - \boldsymbol{Z}_N^\dagger \boldsymbol{Z}_N)\mathrm{vec}_{\mathrm{s}}(W_\star) + \boldsymbol{Z}_N^\dagger E_N. \tag{25}$$

We first discuss the nuances associated with the first term on a high level, before formally stating the error bound.

Estimator bias: Noting that the second term is zero mean, the first term in (25) constitutes the bias in the estimate. This bias is characterized exactly in Lem. E.1 and is often non-zero. The reason for this can be given in terms of the matrix associated with $\mathcal{E}(W; \cdot)$ denoted as $\boldsymbol{E} \in \mathbb{R}^{\eth_{n_x} \times \eth_{n_z}}$, which uniquely determines the second moment dynamics. Our procedure parametrized $\boldsymbol{E}$ in terms of $W \in \mathbb{S}^{n_w}$. In other words, we use $\eth_{n_w}$ parameters to describe $\eth_{n_x}\eth_{n_z}$ degrees of freedom. In the model-free case (i.e. when $n_w = n_z n_x$) we would always over parametrize $\boldsymbol{E}$ for $n_x > 1$.

One might therefore argue that $\boldsymbol{E}$ should be parametrized directly. We instead use W for three reasons: *(i)* often a description like (1) with a bound on $\|w\|_2$ is more natural in practical applications and such bounds do not translate well into bounds on $\boldsymbol{E}$; *(ii)* control synthesis using a confidence set over W is more tractable; and *(iii)* the bias is inconsequential for the LQR cost and stability analysis.

Here *(iii)* is caused by the fact that the bias always lies in the kernel of the second moment dynamics (cf. (15)) $\mathcal{E}(W;Z)$ w.r.t. W. This kernel in fact is also the cause of the bias in the first place as formalized in Lem. E.2.

Data-driven error: We are now ready to state a data-driven bound, suitable for control synthesis. To do so we first consider some auxiliary operators. Letting $\boldsymbol{\mathcal{M}}$ denote the mode tensor (cf. Asm. I) and $\boldsymbol{\mathcal{W}} := \boldsymbol{\mathcal{M}} \otimes_{\mathrm{s}} \boldsymbol{\mathcal{M}}$ as before, we introduce $\mathcal{W}(Z) := \boldsymbol{\mathcal{W}}_{(3)}((Z \otimes_{\mathrm{s}} Z) \otimes I_{\delta_{n_x}})\boldsymbol{\mathcal{W}}_{(3)}^{\top}$ and

$$\boldsymbol{H}_i := \left[\textstyle\sum_{j=1}^{N} \mathcal{W}\left(z_j z_j^{\top}\right)\right]^{\dagger} \mathcal{W}\left(z_i z_i^{\top}\right), \tag{26}$$

where $\mathcal{H}_i \in \mathrm{SL}_{n_z}^{n_w}$, the operator associated with $\boldsymbol{H}_i$ has norm $\|\mathcal{H}_i\|_2 \leq \sqrt{n_w}\|\boldsymbol{H}_i\|_2$ (cf. Lem. III.3). One can show that (cf. Cor. E.4) the second term in (25) equals

$$\mathrm{vec}_{\mathrm{s}}\left(\textstyle\sum_{i=1}^{N} \mathcal{H}_i(w_i w_i^{\top} - W_\star)\right),$$

which is an i.i.d. sum of bounded random matrices. Hence a matrix Hoeffding bound is applicable:

Theorem V.4. *Let $\mathcal{E}_\star$ describe the true moment dynamics as in (15). Assume access to samples $(x_{i+1}, z_i)_{i=1}^N$ satisfying (19) and Asm. II and a mode tensor $\boldsymbol{\mathcal{M}}$ satisfying Asm. I.*

If $\hat{W}$ is the LS estimate (24) and $N \succeq \delta_{n_z}$. Then

$$\mathbb{P}[\mathcal{E}(\underline{W};Z) \preceq \mathcal{E}_\star(Z) \preceq \mathcal{E}(\overline{W};Z)] \geq 1-\delta$$

for all $Z \in \mathbb{S}^{n_z}$ with $\underline{W} := \hat{W} - \beta_W I_{n_w}$ and $\overline{W} := \hat{W} + \beta_W I_{n_w}$ where $\beta_W = r_w^2 \zeta_W \sqrt{2\log(2n_w/\delta)}$ and $\zeta_W^2 := \sum_{i=1}^N \|\mathcal{H}_i\|_2^2$, where $\mathcal{H}_i$ denote the operators associated with $\boldsymbol{H}_i$ in (26).

Proof. Proof deferred to Appendix E. □

Sample complexity: Due to independence one may expect the matrices $\mathcal{W}(z_i z_i^{\top})$ to behave similarly. Hence, one can expect $\|\mathcal{H}_i\|_2$ to decrease with $1/N$. So $\zeta_W^2 = \sum_{i=1}^N \|\mathcal{H}_i\|_2^2$ would decrease with $1/N$ too. We formalize this intuition in a simplified setting below.

Lemma V.5. *Following the setting of Thm. V.4 and assuming additionally that $\|z_i\|_2 \leq r_z$ a.s. for all $i \in \mathbb{N}_{1:N}$. Then, if $N \geq (\|\boldsymbol{\mathcal{W}}\|_2^2 \tau_W/\gamma_W)^2$,*

$$\mathbb{P}\left[\zeta_W \leq \frac{\|\boldsymbol{\mathcal{W}}\|_2^2}{\sqrt{N}\gamma_W - \|\boldsymbol{\mathcal{W}}\|_2^2 \tau_W}\right] \leq \delta,$$

with $d_W = \mathrm{rk}(\boldsymbol{\mathcal{W}}_{(3)})$, $\tau_W := \sqrt{2\ln(2d_W/\delta)}$ and

$$\gamma_W := \lambda_{d_W}\left[\boldsymbol{\mathcal{W}}_{(3)}\left(\mathbb{E}[zz^{\top} \otimes zz^{\top}]/r_z^4 \otimes I\right)\boldsymbol{\mathcal{W}}_{(3)}^{\top}\right],$$

with $\boldsymbol{\mathcal{W}} = \boldsymbol{\mathcal{M}} \otimes_{\mathrm{s}} \boldsymbol{\mathcal{M}}$.

Proof. Proof deferred to Appendix E. □

Remark V.6. Notes on bounding the *tensor spectral norm* $\|\boldsymbol{\mathcal{W}}\|_2$ are given in [63].

Remark V.7. Note that convergence of ζ_W depends on the kurtosis of z (i.e. the fourth order moment) through γ_W, which can be difficult to quantify in practice. We observe that the value of γ_W is invariant to scaling of z and is larger when $z \otimes_{\mathrm{s}} z$ is spread out close to the boundary of its support, which will imply faster convergence.

VI. Distributionally robust control synthesis

In this section we design approximate reformulations of ($\overline{\mathcal{LQR}}$). The goal is to preserve the properties of Thm. IV.8. Specifically we want to synthesize a distributionally robustly stabilizing controller, by solving a SDP as in Thm. IV.10. The synthesis procedure is described in §VI-A. We also briefly comment on the sample complexity in §VI-B.

A. Synthesis

In this section, we use our confidence bound in Thm. V.4 to synthesize a controller that stabilizes the true system with high probability. Similarly, to §IV-C we do so by investigating ($\mathcal{LQR}_{\mathrm{cp}}$). Specifically consider the CP version of ($\overline{\mathcal{LQR}}$):

$$\underset{Z_t \succeq 0}{\text{minimize}} \ \underset{\underline{W} \preceq W \preceq \overline{W}}{\text{maximize}} \ \textstyle\sum_{t=0}^{\infty} \mathrm{Tr}[Z_t H]. \tag{$\overline{\mathcal{LQR}}_{\mathrm{cp}}$}$$

subject to $Z_t = [X_t, V_t; V_t^{\top}, U_t] \succeq 0$, $X_{t+1} = \mathcal{E}(W;Z_t)$ and with $\underline{W}$ and $\overline{W}$ as in Thm. V.4. The constraints on W define the equivalent of the ambiguity set in ($\overline{\mathcal{LQR}}$).

Next we solve ($\overline{\mathcal{LQR}}$) by generalizing (IV.8). We will use the closed-loop policy Π_K as defined in (18) and:

$$\mathcal{E}_K(W;\cdot) = \mathcal{E}(W;\Pi_K(\cdot)).$$

We have the following:

Theorem VI.1. *Assume $H = \mathrm{diag}(Q,R) \succ 0$ and*

$$\exists K : \rho(\mathcal{E}_K(W;\cdot)) < 1, \ \forall W : \underline{W} \preceq W \preceq \overline{W}.$$

Then the policy $Z_t = \Pi_{\overline{K}}(X_t)$ is optimal for ($\overline{\mathcal{LQR}}_{\mathrm{cp}}$), with $(\overline{P}, \overline{K})$ the optimal solution as described in Thm. IV.8 where we assume $\mathcal{E} = \mathcal{E}(\overline{W},\cdot)$. Moreover

(i) $\rho(\mathcal{E}_K(W;\cdot)) < 1$, $\forall W : \underline{W} \preceq W \preceq \overline{W}$;

(ii) $\mathrm{Val}(\overline{\mathcal{LQR}}_{\mathrm{cp}}) = \mathrm{Tr}[\overline{P} X_0]$.

Proof. We do not repeat the full proof of Thm. IV.8. Instead we consider only the Bellman operator and show that it has a similar structure to the one of ($\mathcal{LQR}_{\mathrm{cp}}$).

The analogous Bellman operator to (44) is given as

$$\underset{Z \succeq 0}{\text{minimize}} \ \underset{\underline{W} \preceq W \preceq \overline{W}}{\text{maximize}} \{\mathrm{Tr}[ZH] + \mathrm{Tr}[P_k \mathcal{E}(W;Z)]\}. \tag{27}$$

Observing that (cf. Cor. III.14) $\mathcal{E}(Z) = \mathcal{E}(W;Z) \preceq \mathcal{E}(\overline{W};Z)$ if $W \preceq \overline{W}$ shows that $\mathcal{E}(\overline{W};Z)$ achieves the maximum in (27) $\forall Z$. Therefore, value iteration becomes:

$$\mathrm{Tr}[P_{k+1}X] = \min_{Z \succeq 0}\{\mathrm{Tr}[ZH] + \mathrm{Tr}[P_k \mathcal{E}(\overline{W};Z)]\}.$$

The proof is then completed by following the remainder of the proof of Thm. IV.8, for $\mathcal{E}(\overline{W};\cdot)$ instead of $\mathcal{E}(\cdot)$. □

Thm. V.4 implies the true $\mathcal{E}_\star$ satisfies $\mathcal{E}(\underline{W};\cdot) \preceq \mathcal{E}_\star \preceq \mathcal{E}(\overline{W};\cdot)$ with probability at least $1-\delta$. So we have:

Corollary VI.2. *Let $\overline{K}$ denote the DR policy as in Thm. VI.1 and assume the setup of Thm. V.4 holds. Moreover, let $\mathcal{E}_\star$ denote the true moment dynamics.*

Then with probability at least $1-\delta$,

(i) $\rho(\mathcal{E}_\star(\Pi_K(\cdot)))<1$;

(ii) $\mathrm{Val}(\overline{\mathcal{LQR}}_{\mathrm{cp}}) \geq \mathrm{Val}(\mathcal{LQR}_{\mathrm{cp}}) = \mathrm{Val}(\mathcal{LQR}_\star)$.

Proof. By Thm. V.4, $\mathbb{P}[\mathcal{E}(\underline{W};\cdot) \preceq \mathcal{E}_\star \preceq \mathcal{E}(\overline{W};\cdot)] \geq 1-\delta$. Thus, $\exists W\colon \underline{W} \preceq W \preceq \overline{W}$ such that $\mathcal{E}(W;\cdot) = \mathcal{E}_\star(\cdot)$. So stability *(i)* follows from Thm. VI.1*(i)* and *(ii)* follows by maximization over W in ($\overline{\mathcal{LQR}}_{\mathrm{cp}}$) for the inequality and Prop. IV.6 for the equality. □

Remark VI.3. As in Prop. IV.6, we have $\mathrm{Val}(\overline{\mathcal{LQR}}_{\mathrm{cp}}) = \mathrm{Val}(\overline{\mathcal{LQR}})$. To save space we omit a proof and instead claim only $\mathrm{Val}(\overline{\mathcal{LQR}}_{\mathrm{cp}}) \geq \mathrm{Val}(\mathcal{LQR}_\star)$. This is sufficient for safety critical applications where an upper bound is sufficient for stability. We show consistency of ($\mathcal{LQR}_{\mathrm{cp}}$) later to further strengthen the result.

B. Sample complexity

As shown in the previous section, the solution of the DR problem §V is described by a perturbed Riccati equation. As such, we use perturbation analysis to bound the sub-optimality of the DR controller when applied to the true dynamics. We leverage a previous result by the authors in [50] to do so.

The final complexity bound is as follows:

Theorem VI.4. *Let* $\overline{K}$ *denote the optimal controller as in Thm. VI.1. Moreover, assume Asm. I–II. Then, for sufficient samples* N *and with probability at least* $1-\delta$,

$$\mathrm{Tr}\left[\sum_{t=0}^{\infty} Z_t H\right] - \mathrm{Val}(\mathcal{LQR}_{\mathrm{cp}}) = \mathcal{O}\left(\frac{1}{N}\right),$$

where the first term is the cost of ($\mathcal{LQR}_{\mathrm{cp}}$) *achieved for* $Z_t = \Pi_{\overline{K}}(X_t)$ *with true moment dynamics* $X_{t+1} = \mathcal{E}_\star(Z_t)$.

Moreover by Prop. IV.6, for sufficient samples N *and with probability at least* $1-\delta$,

$$\mathbb{E}\left[\sum_{t=0}^{\infty} x_t^\top Q x_t + u_t R u_t\right] - \mathrm{Val}(\mathcal{LQR}_\star) = \mathcal{O}\left(\frac{1}{N}\right),$$

where $x_{t+1} = A(v_t)x_t + B(v_t)u_t$ *as in (1),* $u_t = \overline{K}x_t$. *So the first term is the cost achieved when applying* $\overline{K}$ *to the true multiplicative noise dynamics (1).*

Proof. We suggest some modifications of the proof of [50, Cor. III.4] to show the second result. The first result then follows directly from Prop. IV.6.

Modifications are required as [50, Asm. II.2] assumes a relative bound like $\underline{\alpha}W_\star \preceq \hat{W} - W_\star \preceq \bar{\alpha}W_\star$, with $\underline{\alpha} \leq \bar{\alpha} = \mathcal{O}(1/\sqrt{N})$ with high probability, $W_\star$ the true value of $\mathbb{E}[ww^\top]$ and $\hat{W}$ some estimate. Instead we have $0 \preceq \overline{W} - W_\star \preceq 2\beta_W I$ with $\beta_W = \mathcal{O}(1/\sqrt{N})$ by Thm. V.4 and Lem. V.5 (where we ignore the effect of the bias, without loss of generality due to Lem. E.2). So the proofs need to be modified. Specifically we replace [50, Lem. IV.3] with a bound like

$$\|\boldsymbol{\mathcal{A}}_{(2)}(\overline{W} - W_\star \otimes P)\boldsymbol{\mathcal{A}}_{(2)}^\top\|_2 \leq 2\beta_W \|\boldsymbol{\mathcal{A}}_{(2)}(I \otimes P)\boldsymbol{\mathcal{A}}_{(2)}^\top\|_2.$$

The proof is analogous and is based on Cor. III.14. Following the proof of [50, Lem IV.3] and its dependencies, clearly only the constants in [50, Table. 1] are affected. The rate shown in [50, Cor. III.4] remains the same. □

Remark VI.5. Since the estimation error bounds are of the same order, we get the same sample complexity in the certainty equivalent setting [10, §3.1], where we use $\hat{W}$ instead of $\overline{W}$. The proof is analogous.

VII. Imposing prior structure

The tools we introduce for identification are also applicable when prior structure is imposed on the dynamics (1). Let

$$x_{t+1} = \left(\sum_{i=2}^{n_v+1} A_i[v_t]_i + A_1\right) x_t + \left(\sum_{i=2}^{n_v+1} B_i[v_t]_i + B_1\right) u_t, \tag{28}$$

where A_i and B_i are assumed known for $i \in \mathbb{N}_{1:n_v+1}$. We were able to handle the terms linear in v_t in the previous sections by adequate selection of $\boldsymbol{\mathcal{M}}$ in Asm. I. However, the introduction of the known terms A_1 and B_1 are not supported[4]. In this section we illustrate how this prior information can be exploited in a DR control scheme.

In this setting we slightly alter Asm. I:

Assumption III. *Consider* $\boldsymbol{\mathcal{V}} \in \mathbb{R}^{n_x \times n_z \times n_v+1}$ *s.t.* $[\boldsymbol{\mathcal{V}}]_{::i} = [A_i, B_i]$ *for* $i \in \mathbb{N}_{1:n_v+1}$ *and* $\{A_i, B_i\}_{i=1}^{n_v+1}$ *as in (28).*

Assume $\boldsymbol{\mathcal{M}} \in \mathbb{R}^{n_x \times n_z \times n_w}$ *is s.t.* $[\boldsymbol{\mathcal{M}}]_{::1} = [\boldsymbol{\mathcal{V}}]_{::1}$ *and*

$$\mathrm{rk}(\tilde{\boldsymbol{\mathcal{M}}}_{(3)}) = \mathrm{rk}([\tilde{\boldsymbol{\mathcal{M}}}_{(3)}; \tilde{\boldsymbol{\mathcal{V}}}_{(3)}]) = n_w,$$

with $\tilde{\boldsymbol{\mathcal{M}}} := [\boldsymbol{\mathcal{M}}]_{:,:,2:n_w}$ *(similarly for* $\tilde{\boldsymbol{\mathcal{V}}}$*) the non-structured part of the mode tensor.*

We can also partition $\boldsymbol{\mathcal{M}}$ similarly as before in $\boldsymbol{\mathcal{A}}$ and $\boldsymbol{\mathcal{B}}$ and write the dynamics

$$\begin{aligned} x_{t+1} &= [\![\boldsymbol{\mathcal{A}}; x_t, w_t]\!] + [\![\boldsymbol{\mathcal{B}}; u_t, w_t]\!] \\ &= [\![\boldsymbol{\mathcal{M}}; z_t, w_t]\!]. \end{aligned} \tag{29}$$

Then, Lem. IV.1 generalizes to

Lemma VII.1. *When Asm. I holds and* $w_t = (1, \tilde{w}_t)$ *with* $\tilde{w}_t = (\tilde{\boldsymbol{\mathcal{M}}}_{(3)}^\top)^\dagger \tilde{\boldsymbol{\mathcal{V}}}_{(3)}^\top v_t$, $\forall t \in \mathbb{N}$. *Then* (29) *produce the same trajectories as* (28).

Proof. The proof is analogous to that of Lem. IV.1. □

The second moment dynamics for (28) are then:

$$X_{t+1} = \boldsymbol{\mathcal{M}}_{(1)}(W \otimes_{\mathrm{s}} Z_t)\boldsymbol{\mathcal{M}}_{(1)}^\top =: \mathcal{E}(W; Z_t). \tag{30}$$

where $X_t = \mathbb{E}[x_t x_t^\top]$ and $Z_t = \mathbb{E}[z_t z_t^\top]$ with

$$W := \mathrm{diag}(0, \Sigma) + (1, \mu)(1, \mu)^\top,$$

with $\tilde{W} := \mathbb{E}[\tilde{w}_t \tilde{w}_t^\top]$, $\Sigma := \tilde{W} - \mu\mu^\top$ and $\tilde{\mu} := \mathbb{E}[\tilde{w}_t]$. Similar to before we just write $\mathcal{E}(\cdot) = \mathcal{E}(W;\cdot)$ when the specific value of W is unimportant.

[4] They can be supported by assuming A_0 and B_0 are multiplied by an unknown random value. However, by doing so, structural information is lost.

When estimating the second moment dynamics, estimating $\tilde{W}$ is similar to before. However we now also estimate the mean of μ. We provide details on estimation in §VII-A. Then we design a DR controller in §VII-B.

A. Mean identification

The measurement model (19), under Asm. III is:

$$x_{i+1} = [\![\boldsymbol{\mathcal{M}}; z_i, (1, \tilde{w}_i)]\!], \quad \forall i \in \mathbb{N}_{1:N}. \tag{31}$$

By expanding the Tucker product, (31) is equivalent to $y_{i+1} := x_{i+1} - [A_1, B_1]z_i = [\![\tilde{\boldsymbol{\mathcal{M}}}; z_i, w_i]\!]$. This is identical to the original measurement model (19), but with a perturbed value of x_{i+1} and truncated tensor $\tilde{\boldsymbol{\mathcal{M}}}$. Nonetheless, we can still identify $\tilde{W}$ using (24).

Moreover, identification of the mean is also possible:

$$y_{i+1} = (\tilde{\boldsymbol{\mathcal{M}}} \,\overline{\times}_2\, z_i)(\mu_\star + \epsilon_i),$$

with $\epsilon_i = \tilde{w}_i - \mu_\star$ a zero-mean error term and with $\mu_\star$ the true value of the mean μ.

As such, introducing the stacked operators:

$$Y_N^\mu = (y_2, \dots, y_{N+1}),\ \boldsymbol{Z}_N^\mu = [\tilde{\boldsymbol{\mathcal{M}}} \,\overline{\times}_2\, z_1; \dots; \tilde{\boldsymbol{\mathcal{M}}} \,\overline{\times}_2\, z_N]$$

The LS estimator then becomes:

$$\hat{\mu} = (\boldsymbol{Z}_N^\mu)^\dagger Y_N^\mu. \tag{32}$$

Data-driven bound: Compared to identification of $\tilde{W}$, developing a data-driven bound for the is considerably easier as the estimator is unbiased. We have the following error model:

$$\mu_\star - \hat{\mu} = \textstyle\sum_{i=1}^N G_i \epsilon_i, \tag{33}$$

where we introduce $G_i := (\sum_{j=1}^N \mathcal{M}(z_j z_j^\top))^\dagger \mathcal{M}(z_i z_i^\top)$ and $\mathcal{M}(Z) := \tilde{\boldsymbol{\mathcal{M}}}_{(3)}(Z \otimes I_{n_x})\tilde{\boldsymbol{\mathcal{M}}}_{(3)}^\top$.

So, under Asm. II, the error is a sum of norm-bounded zero-mean i.i.d. random vectors. Hence Lem. C.2 is applicable:

Lemma VII.2. *Following Asm. II–III with $N \geq n_z$. Let $\hat{\mu} = (\boldsymbol{Z}_N^\mu)^\dagger Y_N^\mu$ and $\zeta_\mu^2 := \sum_{i=1}^N \|G_i\|_2^2$. Then,*

$$\mathbb{P}[\|\mu_\star - \hat{\mu}\|_2 \leq \beta_\mu] \geq 1 - \delta,$$

with $\beta_\mu = r_w \zeta_\mu (2 + \sqrt{2\ln(1/\delta)})$.

Proof. A more formal proof is provided in Appendix E. □

Moment ambiguity: We will synthesize a controller that is robust for all $\mathcal{E}$ in the following ambiguity set

Definition VII.3. *Consider the structured ambiguity set:*

$$\mathscr{W} := \left\{ W \in \mathbb{S}^{n_w} : \begin{array}{l} W = \operatorname{diag}(0, \Sigma) + (1, \mu)(1, \mu)^\top, \\ \|\hat{\mu} - \mu\|_2 \leq \beta_\mu,\ \|\hat{\Sigma} - \Sigma\|_2 \leq \beta_\Sigma \end{array} \right\}, \tag{34}$$

with $\beta_\mu = r_w \zeta_\mu (2 + \sqrt{2\ln(1/\delta_\mu)})$, $\beta_\Sigma = \beta_W + \beta_\mu(\beta_\mu + 2\|\hat{\mu}\|_2)$ and $\beta_W = r_w^2 \zeta_W \sqrt{2\log(2n_w/\delta_W)}$ as in Thm. V.4 for probability δ_W. Moreover $\hat{\mu}$ is as in (32) and $\hat{\Sigma} := \hat{\tilde{W}} - \hat{\mu}\hat{\mu}^\top$ for $\hat{\tilde{W}}$ as in (24), but with the truncated mode tensor and y_{i+1} instead of x_{i+1}.

The ambiguity set is valid in the following sense:

Theorem VII.4. *Let $\mathcal{E}_\star$ denote the true moment dynamics and let Asm. II–III hold. Then with probability at least $1 - \delta_\mu - \delta_W$:*

$$\exists W \in \mathscr{W} : \mathcal{E}_\star(Z) = \mathcal{E}(W; Z), \forall Z \in \mathbb{S}^{n_z}$$

with $\mathscr{W}$ as in Def. VII.3.

Proof. The full proof deferred to Appendix E. It combines Thm. V.4 and Lem. VII.2 with some triangle inequalities. □

B. Control synthesis

Using the learned ambiguity set we can then proceed with control synthesis. Unlike in Thm. VI.1 we cannot find the exact solution. Instead we solve a relaxed SDP by leveraging a result from robust optimization [49].

In the non-structured setting, we used Thm. IV.10 to synthesize a controller by solving a SDP. There is however no guarantee that a feasible solution to that SDP produces a stabilizing controller. This is only guaranteed for the optimum.

So instead we will consider:

$$\begin{array}{ll} \underset{(P,K)}{\text{minimize}} & \operatorname{Tr}[PX_0] \\ \text{subj. to} & P \succeq 0,\ P - \mathcal{E}_K^*(P) \succeq Q + K^\top R K. \end{array} \tag{35}$$

The second constraint is a Lyapunov inequality for $\mathcal{E}_K(P)$, which is defined as $\mathcal{E}(\Pi_K(\cdot))$ as in (18). Hence any feasible K is stabilizing by Lem. III.6.

Theorem VII.5. *The problem (35) satisfies:*

(i) $\rho(\mathcal{E}_K) < 1$ for any feasible K;

(ii) $(P_\star, K_\star)$ is optimal and uniquely so if $X_0 \succ 0$.

Here $P_\star$ and $K_\star$ those of Thm. IV.8.

Proof. Full proof deferred to Appendix B. □

We can now robustify (35). Specifically, we introduce the robust constraint

$$P - \mathcal{E}_K^*(W; P) \succeq Q + K^\top R K, \quad \forall W \in \mathscr{W} \tag{36}$$

with $\mathscr{W}$ as in Def. VII.3 and $\mathcal{E}_K^*(W; P) = \Pi_K^*(\mathcal{E}^*(W; P))$ (similar to Lem. B.1). The non-linear SDP (35) with this robust constraint is relaxed as stated below.

Theorem VII.6. *Following the notation of Def. VII.3 and letting $\{(A_i, B_i)\}_{i=1}^{n_w}$ denote the modes in (28), we introduce $\Theta \in \mathbb{S}^{n_x}$, $\Gamma \in \mathbb{R}^{n_u \times n_x}$, $\overline{\Sigma} := \hat{\Sigma} + \beta_\Sigma I$ and $F_i = A_i\Theta + B_i\Gamma$ for $i \in \mathbb{N}_{1:n_w}$. Moreover, let $\boldsymbol{F} := [F_2; \dots; F_{n_w}]$ and $\hat{F} = F_1 + \sum_{i=2}^{n_w} [\hat{\mu}]_i F_i$. Let $X_0 = C_0 C_0^\top$ denote the Cholesky factorization of $X_0 \succeq 0$.*

Consider the SDP:

$$\begin{array}{ll} \underset{\Theta,\Gamma,\Phi,\Psi,\Lambda}{\text{minimize}} & \operatorname{Tr}[\Phi] \\ \text{subj. to} & \Theta \succ 0, \Phi \succ 0, \Psi \succ 0, \Lambda \succ 0, \\ & \begin{bmatrix} \Phi & U_0^\top \\ U_0 & \Theta \end{bmatrix} \succeq 0,\ \begin{bmatrix} \Lambda & \beta_\mu \boldsymbol{F}^\top \\ \beta_\mu \boldsymbol{F} & I \otimes \Psi \end{bmatrix} \succeq 0, \\ & T \succeq \operatorname{diag}(\Lambda, 0, \Psi, 0, 0), \end{array}$$

where we use the linear map:

$$T := \begin{bmatrix} \Theta & \boldsymbol{F}^\top & \hat{F}^\top & \Theta & \Gamma^\top \\ \boldsymbol{F} & \bar{\Sigma}^{-1} \otimes \Theta & & & \\ \hat{F} & & \Theta & & \\ \Theta & & & Q^{-1} & \\ \Gamma & & & & R^{-1} \end{bmatrix}.$$

If we assume Asm. II–III. Let $\mathcal{E}_\star$ denote the true second moment dynamics as in (30) and Π_K the closed-loop policy as in (18). Then given a feasible (Θ, Γ), the pair (P, K) with $P = \Theta^{-1}$ and $K = \Gamma\Theta^{-1}$ satisfies

(i) $\rho(\mathcal{E}_\star(\Pi_K(\cdot))) < 1$*; and*

(ii) $\mathrm{Tr}[\Phi] \geq \mathrm{Tr}[PX_0] \geq \mathrm{val}(\mathcal{LQR}_{\mathrm{cp}}) = \mathrm{val}(\mathcal{LQR}_\star)$*,*

with probability at least $1 - \delta_W - \delta_\mu$.

Proof. The proof is similar to that of [13, Thm. 1]. Note that $\mathcal{E}_K^*(P) = \Pi_K^*(\mathcal{E}^*(P))$ by Lem. B.1. So (36) becomes

$$P - \Pi_K^*(\mathcal{E}^*(W; P)) \succeq \Pi_K^*(H), \tag{37}$$

with $H = \mathrm{diag}(Q, R)$. We then plug in the specific form of W as described in Def. VII.3. That is $W = \mathrm{diag}(0, \Sigma) + (1, \mu)(1, \mu)^\top$. Thus the left-hand side of (37) becomes

$$P - \Pi_K^*(\mathcal{E}^*(\mathrm{diag}(0, \Sigma; P))) - \Pi_K^*(\mathcal{E}^*((1, \mu)(1, \mu)^\top; P)). \tag{38}$$

Note that, for all Z,

$$\mathcal{E}\left(\mathrm{diag}(0, \Sigma); Z\right) = \boldsymbol{\mathcal{M}}_{(3)}(\mathrm{diag}(0, \Sigma), Z)\boldsymbol{\mathcal{M}}_{(3)}^\top.$$

We have $\boldsymbol{\mathcal{M}}_{(3)} = [[\boldsymbol{\mathcal{M}}]_{::1}, \tilde{\boldsymbol{\mathcal{M}}}_{(3)}]$. Thus

$$\mathcal{E}(\mathrm{diag}(0, \Sigma); Z) = \tilde{\mathcal{E}}(\Sigma; Z) := \tilde{\boldsymbol{\mathcal{M}}}_{(3)}(\Sigma \otimes Z)\tilde{\boldsymbol{\mathcal{M}}}_{(3)}^\top. \tag{39}$$

This enables us to further simplify (38) to show that

$$P - \Pi_K^*(\tilde{\mathcal{E}}^*(\Sigma; P)) - \Pi_K^*(\mathcal{E}^*((1, \mu)(1, \mu)^\top; P)) \succeq \Pi_K^*(H)$$

holding for all Σ s.t. $\|\Sigma - \hat{\Sigma}\|_2 \leq \beta_\Sigma$ and for all $\mu \in (\{\hat{\mu}\} + r_\mu \mathcal{B}_{n_w})$, where $\mathcal{B}_{n_w}$ denotes the unit ball implies (36).

Clearly, $\tilde{\mathcal{E}}^*(\bar{\Sigma}; P) \succeq \tilde{\mathcal{E}}^*(\Sigma; P)$ (cf. Cor. III.14) for all Σ s.t. $\|\Sigma - \hat{\Sigma}\|_2 \leq \beta_\Sigma$. Hence (36) holds if:

$$P - \Pi_K^*(\mathcal{E}^*(\mu_0\mu_0^\top; P)) - \Pi_K^*(\tilde{\mathcal{E}}^*(\bar{\Sigma}; P)) \succeq \Pi^*(H), \tag{40}$$

for all $\mu \in (\{\hat{\mu}\} + r_\mu \mathcal{B}_{n_w})$. By Lem. B.1,

$$\Pi_K^*(\mathcal{E}^*(W; P)) = [I, K^\top](\boldsymbol{\mathcal{M}})_{(2)}(W \otimes P)(\boldsymbol{\mathcal{M}})_{(2)}^\top[I; K].$$

We can expand the definition of the 2-mode matricization (7):

$$(\boldsymbol{\mathcal{M}})_{(2)} = \begin{bmatrix} A_1^\top & \dots & A_{n_w}^\top \\ B_1^\top & \dots & B_{n_w}^\top \end{bmatrix}.$$

Hence, letting $\Theta = P^{-1}$ and $\Gamma = K\Theta$ and left-right multiplying (40), plugging in the definition of $\boldsymbol{F}$ and F_1 gives:

$$\begin{aligned} W - [F_1; \boldsymbol{F}]^\top(\mu_0 \otimes I)P(\mu_0^\top \otimes I)[F_1; \boldsymbol{F}] \\ - \boldsymbol{F}^\top(\bar{\Sigma} \otimes I)\boldsymbol{F} \succeq \Theta Q \Theta + \Gamma^\top R \Gamma, \end{aligned}$$

for all $\mu \in r_\mu \mathcal{B}_{n_w}$. Applying Schur complements and introducing $\theta = (\hat{\mu} - \mu)/r_\mu$, allows us to rewrite the final display as an inequality of the type $\mathcal{U}[\beta_\mu] \succeq 0$ in Lem. F.2, with T as in the theorem above. Applying the Lemma gives the second and third LMI.

So by Lem. F.2, a feasible pair (Θ, Γ) produces a pair (P, K) that is feasible for the robust constraint (36). Since (36) holds for all $W \in \mathscr{W}$ it holds for $\mathcal{E}_\star$ provided that $\exists W \in \mathscr{W}$ such that $\mathcal{E}_\star(Z) = \mathcal{E}(W; Z)$ for all Z, which is true with probability at least $1 - \delta_\mu - \delta_W$. Hence, by Thm. VII.5*(i)* (i) follows.

The first constraint in the final SDP is equivalent to

$$\Phi \succeq U_0^\top \Theta^{-1} U_0^\top,$$

by a Schur complement. So $\mathrm{Tr}[\Phi] \geq \mathrm{Tr}[U_0^\top \Theta^{-1} U_0^\top] = \mathrm{Tr}[X_0 P]$. Combined with the previous argument about feasibility and Thm. VII.5*(ii)*, we prove (ii). □

VIII. Numerical Experiments

In this section we numerically investigate the methods developed in this paper. We begin with the simple case of *repeated initilialization* (cf. Rem. V.1), comparing both the case where the modes are known and where they are not, i.e. the model-free setting. We show how, in the first setting, our bounds are sufficiently tight to enable control synthesis. Next we use data generated using *rollout* to produce similar estimates and we highlight the differences and challenges associated with doing so. We also show that experimentally our method also works when using only a single trajectory, although theoretical guarantees are not yet available in this setting. We also compared to the averaging rollout approach in [37], [38]. Finally we show how structural information can be exploited to get tighter estimates of the uncertainty.

A. Accuracy of the estimate and bounds

We consider the case $n_x = 2$ and $n_u = 1$ with modes:

$$A_1 = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}, A_2 = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}, A_3 = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}, B_3 = \begin{bmatrix} 0 \\ 1 \end{bmatrix},$$

and $B_1 = B_2 = 0$. The disturbance is sampled uniformly from the ball $\{w \in \mathbb{R}^3 \colon \|w - \mu\|_2 \leq 0.25\}$ with $\mu = (0, 0, 0.1)$.

Repeated initialization: To generate measurements we sample a $z_0 = (x_0, u_0)$ uniformly from an Euclidean ball of radius one and then propagate the dynamics by one step to get a measurement x_1, i.e. using *repeated initilialization* as in Rem. V.1. From this data, an estimate $\hat{W}$ and associated radius β_W are determined as prescribed in Thm. V.4.

To empirically quantify the accuracy of the scheme, the procedure above is repeated $M = 100$ times for each sample count N. For each estimate we compute $\|\hat{W} - W_\star\|_2$ and β_W. A confidence plot is provided in Fig. 2. The middle plot depicts the result when $\boldsymbol{\mathcal{M}}$ is selected, based on true mode info (i.e. $[\boldsymbol{\mathcal{M}}]_{::i} = [A_i, B_i]$ for $i = 1, 2, 3$); and the right-most plot uses the model-free case (i.e. $(\boldsymbol{\mathcal{M}})_{(3)} = I$).

We determined $W_\star$ in the first setting by noting that $\mathbb{E}[ww^\top] = r_w^2 I_{n_w}/(n_w + 2)$ when w is sampled uniformly from $\{w \in \mathbb{R}^{n_w} \colon \|w\|_2 \leq r_w\}$ — which can be verified by symmetry and solving a simple integral — and applying

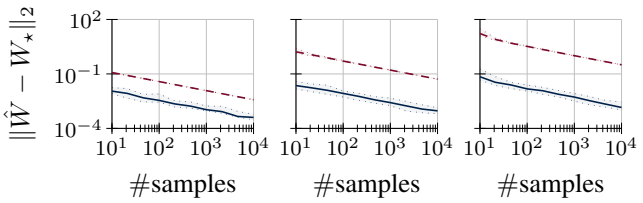


Fig. 2. Error in moment estimation using repeated initialization: (left) when directly observing the disturbance; (mid) when using the true modes in $\boldsymbol{\mathcal{M}}$; (right) when using no mode information, i.e. model-free. The dashed line depicts the radius predicted by Thm. V.4. Each colored area is a 0.1 confidence interval surrounding the median.

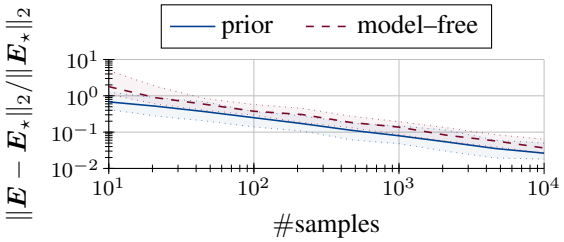


Fig. 3. Estimation accuracy of matrix of second order dynamics and comparison between model-free and true mode based estimation. Each colored area is a 0.1 confidence interval surrounding the median.

the appropriate transformation[5]. Similarly, $W_\star$ for the model-free case is then recovered by applying the transformation in (14).

The left-most plot depicts the error when $\hat{W}$ is estimated directly from measurements of w and the bound is as in Lem. V.3. It is clear from the figure that the empirical error does not increase much when we do not directly observe the disturbance — at least for this simple model. Instead the main loss is in the bound, which is about an order of magnitude looser than the direct sample case. This is to be expected however and, as we confirm later using DR synthesis, the bounds are still practical in low-dimensional settings.

The middle and right-most plot depict the use of prior information and the model-free case respectively. The fact that $W_\star$ is not the same for both cases implies that the absolute errors are dissimilar. For a fair comparison, we additionally use the estimated second moment dynamics $\mathcal{E}$ as in (15). Specifically we plot $\|\hat{\boldsymbol{E}} - \boldsymbol{E}_\star\|_2 / \|\boldsymbol{E}_\star\|_2$ where $\hat{\boldsymbol{E}}$ is the matrix associated with the estimated dynamics and $\boldsymbol{E}_\star$ its true value. These are computed using Cor. III.16. Note that $\boldsymbol{E}_\star$ is the same independent of the selected $\boldsymbol{\mathcal{M}}$. Fig. 3 shows the result.

It is clear that, for the low-dimensional example considered here, exploiting prior mode information has no significant advantage when it comes to estimation accuracy. Instead the main advantage is a smaller ambiguity set. When using the model-free approach, the control synthesis problem remains infeasible for any tested sample count. That is, the ambiguity set exceeds the minimum size that can be

[5]To shift the ball we shift $\tilde{w} = w + \mu$. Hence $\mathbb{E}[\tilde{w}\tilde{w}^\top] = W_\star + \mu\mu^\top$.

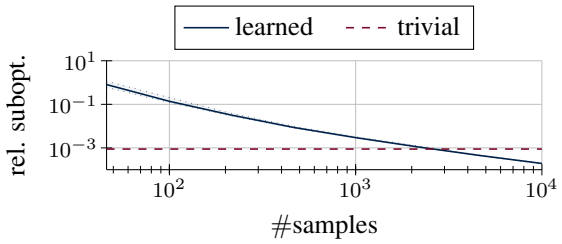


Fig. 4. Evaluation of suboptimality of DR control synthesis. The colored area is a 0.1 confidence interval surrounding the median.

stabilized by a single controller. When exploiting the prior information encapsulated in the modes, this is not the case.

Control synthesis: We can quantify the size of the ambiguity set through controller synthesis as in Thm. VI.1. The performance of a gain $\hat{K}$ is measured through the infinite horizon cost achieved on the true system. This cost equals the optimum of the following SDP, by Prop. III.8 and Lem. B.1,

$$\min_P \{\mathrm{Tr}[PX_0] \colon P - \Pi^*_{\hat{K}}(\mathcal{E}^*_\star(P)) \succeq \mathrm{diag}(Q, R),\ P \succeq 0\}.$$

We select $X_0 = I$ here and in Thm. IV.10 for control synthesis. Also let $Q = I$ and $R = 10$. The relative error with the true optimum of ($\mathcal{LQR}_\star$) is then a metric for the accuracy of the ambiguity set. The result is depicted in Fig. 4. Note that the rate is as predicted in Thm. VI.4.

The dashed line in Fig. 4 depicts the cost for the *trivial ambiguity*, which only uses $\|w\|_2 \leq r_w$ implying $\|W_\star\|_2 \leq r_w^2$. This constraint is also used to synthesize a controller with Thm. VI.1. We say that the learned ambiguity set is *informative* when the performance of the associated controller improves upon the trivial one.

Rollout: Repeated initialization is not realistic in practice, since it assumes we can directly select all states of the system. Rollout instead initializes the system at some — easily realizable — initial state and then applies a sequence of control actions to excite the system. We use the state $x_0 = (1, 1)$ and use the control law $u_t = [-0.5, -0.2]x_t + \delta_t$ with δ_t sampled uniformly from a Euclidian ball of radius 35 at each time step to integrate the dynamics for $T = 25$ time steps. Then $z_i = (x_{T-1}, u_{T-1})$ and $x_{i+1} = x_T$ are used as data points (cf. Rem. V.1) for $i = 1, \ldots, N$. Here N is the amount of rollouts, which we also refer to as the sample count.

Similarly to before we can evaluate the accuracy of the estimate empirically, by resampling data sets $M = 100$ times and computing the errors. The result is depicted in Fig. 5. Comparing with Fig. 2 we see, as expected, that estimation is more challenging when using rollout data. This is likely explained by the distribution of z_i being less suitable for identification as discussed in Rem. V.7

Single trajectory: The right plot in Fig. 5 also depicts the empirical error when we simply use one rollout of length $T = N$. So the data is then (x_i, u_i) and x_{i+1} for $i \in \mathbb{N}_{0:N-1}$. It is clear that the estimate becomes more accurate when the trajectory length is increased, showing — at least empirically

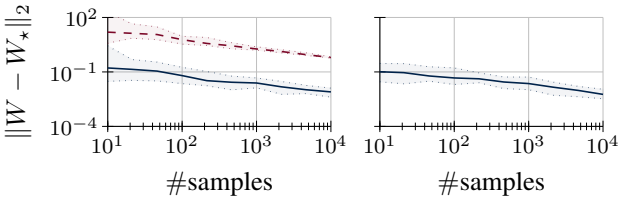


Fig. 5. Empirical result for toy problem using rollout data: (left) using the tail of #nsamples rollouts of length 25, with predicted radius as a dashed line; (right) using one rollout of length #nsamples. Each colored area is a 0.1 confidence interval surrounding the median.

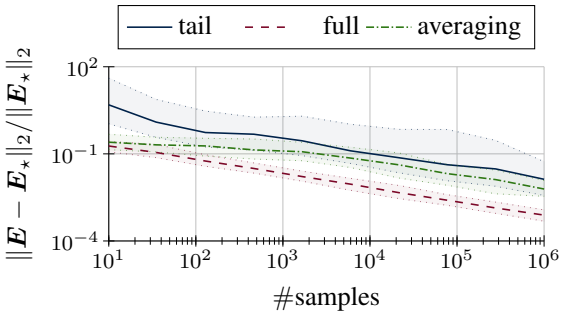


Fig. 6. Comparison of three system identification approaches: (tail) using only the tail of each rollout; (full) using the full rollout; (averaging) when trajectories are averaged over rollouts. Each colored area is a 0.1 confidence interval surrounding the median.

— that our approach can also work for single trajectory identification. This was also observed for a similar setup in [40], yet a convergence proof is still unavailable. It is also unclear which conditions the system and exciting inputs should satisfy (cf. Rem. V.1).

B. Averaged rollouts approach

In this part, we compare our approach with that of [37], [38]. We will refer to this approach as the *averaging rollout approach* (contrasting our *least-squares approach*). In essence, this approach identifies the second moment dynamics (15) directly by using a sample average to construct a single moment trajectory (Z_t, X_t). The second moment dynamics are then estimated using standard identification techniques from LTI identification.

We consider a dynamical system similar[6] to the one considered in [38, §IV.A]. Specifically

$$A_1 = \begin{bmatrix} -0.2 & 0.3 \\ -0.4 & 0.8 \end{bmatrix}, A_2 = \begin{bmatrix} 0.2 & -0.15 \\ -1 & 0 \end{bmatrix}, A_3 = \begin{bmatrix} 1 & 0.1 \\ 0.2 & 0 \end{bmatrix},$$
$$A_4 = \begin{bmatrix} -0.05 & 1 \\ -0.15 & 0 \end{bmatrix}, A_5 = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix},$$
$$B_1^\top = \begin{bmatrix} -1.8 & -0.8 \end{bmatrix}, B_6^\top = \begin{bmatrix} -1 & -0.15 \end{bmatrix}, B_7^\top = \begin{bmatrix} 0.15 & -1 \end{bmatrix}.$$

The matrices A_i and B_i left unspecified are assumed zero. We then sample the disturbance uniformly from a degenerate ellipse such that $W_\star = \mathrm{diag}(1, 0.15, 0.08, 0.02, 0.08, 0.05, 0.2)$ and $[w_t]_1 = 1$ for all $t \in \mathbb{N}$.

We then generate rollouts as in [38] and again evaluate the estimation accuracy in terms of $\boldsymbol{E}$. Fig. 6 shows the result. The rollout length was selected as $T = 12$ and the horizontal axis tracks the amount of rollouts. In the averaging approach, the inputs should not only be random, but their distributions as well. This explains the high variance in the estimates.

The worst error is achieved when we use only the tail of each rollout in our least-squares scheme, which makes sure our convergence guarantees still hold. The averaging approach performs marginally better, and the asymptotic rates are the same (as confirmed by the theory in [38]). One can expect that accuracy improves further when the rollout length is increased[7]. Note however that the approach of [38] does not come with a tight error analysis. If theoretical guarantees are not critical, we can similarly to the single trajectory case considered above, also add other states in the rollout and use the full data set. The resulting estimate clearly outperforms the averaging result. Intuitively this is explained by the latter losing information by averaging.

It is also important to note that the scheme, as introduced in [38] does not support correlation between $A(w)$ and $B(w)$ (as in our first example). Extending the averaging scheme to support this is however trivial.

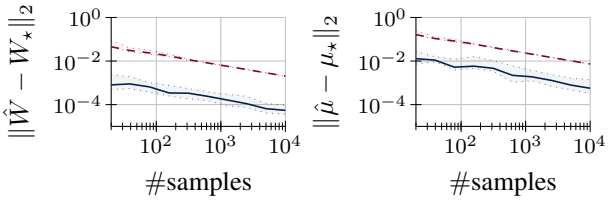


Fig. 7. Error in moment estimation using repeated initialization and prior structure: (left) the second moment estimation error; (right) the mean estimation error. The dashed lines depict the predicted radius. Each colored area is a 0.1 confidence interval surrounding the median.

C. Exploiting prior structure

We consider the structured dynamics (28) with

$$A_1 = \begin{bmatrix} 1 & 0.02 \\ 0 & 0.992 \end{bmatrix}, B_1 = \begin{bmatrix} 0 \\ 0.02 \end{bmatrix}$$
$$A_2 = \begin{bmatrix} 0 & 0 \\ 0 & -0.03 \end{bmatrix}, A_3 = \begin{bmatrix} 0 & -0.03 \\ 0 & 0 \end{bmatrix}, B_4 = \begin{bmatrix} 0 \\ 0.01 \end{bmatrix}.$$

The matrices A_i and B_i left unspecified are assumed zero. The disturbance is distributed uniformly on $\{w \in \mathbb{R}^3 \colon \|w\|_2 \leq 0.05\}$ and data is generated using *repeated initialization*, analogously to the procedure used in §VIII-A. We then estimate $\hat{W}$ and $\hat{\mu}$ using the procedure in §VII. The empirical error and the predicted radii are then predicted in Fig. 7.

Finally, we also evaluate the size of the ambiguity set by using the control synthesis procedure of §VII-B. The LQR

[6]The dynamics in [38] are defined differently. Constructing an exact equivalent that also satisfies Asm. I is challenging. So instead the simpler system we consider is only approximately equivalent.

[7]The analysis performed in [38] however confirms an opposite result. This is likely due to conservativeness in their analysis.

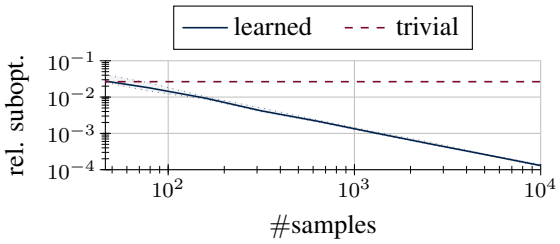


Fig. 8. Evaluation of suboptimality of DR control synthesis with structured prior information. The colored area is a 0.1 confidence interval surrounding the median.

problem is tuned identically to before with $X_0 = I$, $Q = I$ and $R = 10$. We again compare with the true optimum and the trivial controller recovered when using only $\|\mu\|_2 \leq r_w$ and $\|W_\star\|_2 \leq r_w^2 I$ (the transformation to an ambiguity set like Def. VII.3 is analogous to the data-driven case). The resulting suboptimality plot is depicted in Fig. 8. Note that, the horizontal axis starts at $N = 20$. For fewer samples, the synthesis problem is infeasible. It is immediately clear how the structured prior information is exploited by our scheme by noting how little samples are required before the scheme improves upon the trivial controller. Also note that, even though we only proved the $1/N$ rate for the non-structured case, the same decrease is observed here. The same observation was made in [13].

IX. Conclusion

We developed a novel system identification scheme for linear dynamics with state- and input-multiplicative noise. The resulting estimators are shown to converge at a $1/\sqrt{N}$ rate, with N the number of samples, when the data is generated either using rollout or through repeated initialization of the dynamics. We illustrated empirically that — for simple dynamics — the constants in our bounds are practical, i.e. they can be used for DR control synthesis. Moreover, the tightness of the bound is similar to that of a usual matrix Hoeffding bound. Also the DR control synthesis problem was shown to converge at a $1/N$ rate to the true optimum.

For more complex dynamics it is likely that the bounds are not sufficiently tight to enable DR control synthesis. In that case either bootstrapping can be exploited as in [7], [64] or the DR scheme can be used to identify the robustness of the certainty equivalent controller as in [30].

We additionally illustrated how knowledge about deterministic terms can be exploited to accelerate the identification process and tighten the associated bounds (referred to as *structured information*). An additional DR control synthesis scheme was also presented, which exploits the tighter bounds. The scheme was empirically shown to also converge to the true optimum at a $1/N$ rate.

Moreover experiments illustrated that the identification scheme functions well when data is gathered from a single trajectory and compared the results with the averaging rollout identification scheme of [37], [38].

From these tests it is clear that correlation between errors does not accumulate when using single trajectory data. A formal proof of this fact and a sample complexity analysis however proved challenging and is considered an open problem for future work.

References

[1] B. Recht, "A Tour of Reinforcement Learning: The View from Continuous Control," *Annu. rev. control robot. auton. syst.*, vol. 2, no. 1, 2019, arXiv: 1806.09460.

[2] L. Hewing, K. P. Wabersich, M. Menner, and M. N. Zeilinger, "Learning-Based Model Predictive Control: Toward Safe Learning in Control," *Annu. rev. control robot. auton. syst.*, vol. 3, no. 1, 2020.

[3] D. P. Bertsekas, *Dynamic Programming and Optimal Control – Vol I*. Athena scientific, 3rd ed., 2005.

[4] M. Fazel, R. Ge, S. M. Kakade, and M. Mesbahi, "Global Convergence of Policy Gradient Methods for the Linear Quadratic Regulator," *ICML*, vol. 4, 2018.

[5] S. Bradtke, B. Ydstie, and A. Barto, "Adaptive linear quadratic control using policy iteration," in *Proceedings of 1994 American Control Conference - ACC '94*, vol. 3, IEEE, 1994.

[6] F. L. Lewis and D. Vrabie, "Reinforcement learning and adaptive dynamic programming for feedback control," *IEEE Circuits and Systems Magazine*, vol. 9, no. 3, 2009.

[7] S. Dean, H. Mania, N. Matni, B. Recht, and S. Tu, "On the Sample Complexity of the Linear Quadratic Regulator," *FoCM*, 2019.

[8] H. Mania, S. Tu, and B. Recht, "Certainty equivalence is efficient for linear quadratic control," in *NeurIPS*, vol. 32, Curran Associates, Inc., 2019.

[9] M. Abeille and A. Lazaric, "Efficient optimistic exploration in linear-quadratic regulators via lagrangian relaxation," *ICML*, 2020.

[10] K. J. Astrom and B. Wittenmark, *Adaptive Control*. Dover Publications, Inc., 2nd ed., 2008.

[11] L. Ljung, *System Identification: Theory for the User*. Prentice Hall PTR, 1999.

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

[13] P. Coppens, M. Schuurmans, and P. Patrinos, "Data-driven distributionally robust LQR with multiplicative noise," in *L4DC*, vol. 120, PMLR, 2020, arXiv: 1912.09990.

[14] W. M. Wonham, "Optimal Stationary Control of a Linear System with State-Dependent Noise," *SIAM Journal on Control*, vol. 5, no. 3, 1967.

[15] T. Wang, H. Zhang, and Y. Luo, "Stochastic linear quadratic optimal control for model-free discrete-time systems based on Q-learning algorithm," *Neurocomputing*, vol. 312, 2018.

[16] B. Gravell, P. M. Esfahani, and T. Summers, "Learning robust control for LQR systems with multiplicative noise via policy gradient," 2019, arXiv: 1905.13547.

[17] B. Pang and Z.-P. Jiang, "Reinforcement Learning for Adaptive Optimal Stationary Control of Linear Stochastic Systems," 2021, arXiv: 2107.07788.

[18] T. Damm, *Rational Matrix Equations in Stochastic Control*. Lecture Notes in Control and Information Sciences, Springer Berlin Heidelberg, 1 ed., 2004.

[19] E. Todorov and M. Jordan, "Optimal feedback control as a theory of motor coordination," *Nat. Neurosci*, vol. 5, no. 11, 2002.

[20] E. Todorov, "Stochastic Optimal Control and Estimation Methods Adapted to the Noise Characteristics of the Sensorimotor System," *Neural Computation*, vol. 17, no. 5, 2005.

[21] G. Russo and R. Shorten, "On common noise-induced synchronization in complex networks with state-dependent noise diffusion processes," *Physica D: Nonlinear Phenomena*, vol. 369, 2018.

[22] R. R. Mohler and W. J. Kolodziej, "An Overview of Stochastic Bilinear Control Processes," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 10, no. 12, 1980.

[23] R. Mohler, C. Bruni, and A. Gandolfi, "A systems approach to immunology," *Proceedings of the IEEE*, vol. 68, no. 8, 1980.

[24] A. Afshari, M. Karrari, H. R. Baghaee, and G. B. Gharehpetian, "Resilient cooperative control of AC microgrids considering relative state-dependent noises and communication time-delays," *IET Renew. Power Gener.*, vol. 14, no. 8, 2020.

[25] Fan Wang and V. Balakrishnan, "Robust Kalman filters for linear time-varying systems with stochastic parametric uncertainties," *IEEE Trans. Signal Process.*, vol. 50, no. 4, 2002.

[26] P. Sura, M. Newman, C. Penland, and P. Sardeshmukh, "Multiplicative noise and non-Gaussianity: A paradigm for atmospheric regimes?," *J Atmos Sci*, vol. 62, no. 5, 2005.

[27] D. Bernstein, "Robust static and dynamic output-feedback stabilization: Deterministic and stochastic perspectives," *IEEE Trans. Autom. Control*, vol. 32, no. 12, 1987.

[28] J. Doyle, "Guaranteed margins for LQG regulators," *IEEE Trans. Autom. Control*, vol. 23, no. 4, 1978.

[29] J. Doyle, "Robust and optimal control," in *CDC*, vol. 2, IEEE, 1996.

[30] B. Gravell, P. M. Esfahani, and T. Summers, "Robust Control Design for Linear Systems via Multiplicative Noise," 2020, arXiv: 2004.08019.

[31] J. E. Pascoe, "The outer spectral radius and dynamics of completely positive maps," 2019, arXiv: 1905.09895.

[32] C. Kubrusly and O. Costa, "Mean square stability conditions for discrete stochastic bilinear systems," *IEEE Trans. Autom. Control*, vol. 30, no. 11, 1985.

[33] M. A. Rami, X. Chen, and X. Zhou, "Discrete-time Indefinite LQ Control with State and Control Dependent Noises," *J Glob Optim*, vol. 23, no. 3/4, 2002.

[34] D. Hinrichsen and A. J. Pritchard, "Stochastic H^∞," *SIAM J Control Optim*, vol. 36, no. 5, 1998.

[35] L. El Ghaoui, "State-feedback control of systems with multiplicative noise via linear matrix inequalities," *Control Syst. Lett.*, vol. 24, no. 3, 1995.

[36] C. S. Kubrusly, "Identification of discrete-time stochastic bilinear systems," *Int. J. Control*, vol. 33, no. 2, 1981.

[37] Y. Xing, B. Gravell, X. He, K. H. Johansson, and T. Summers, "Linear system identification under multiplicative noise from multiple trajectory data," in *ACC*, 2020.

[38] Y. Xing, B. Gravell, X. He, K. H. Johansson, and T. Summers, "Identification of Linear Systems with Multiplicative Noise from Multiple Trajectory Data," 2021, arXiv: 2106.16078.

[39] B. Gravell and T. Summers, "Robust learning-based control via bootstrapped multiplicative noise," in *L4DC*, vol. 120, PMLR, 2020.

[40] B. Di and A. Lamperski, "Confidence bounds on identification of linear systems with multiplicative noise," in *ACC*, 2021.

[41] T. G. Kolda, "Multilinear operators for higher-order decompositions.," tech. rep., SNL, 2006.

[42] T. G. Kolda and B. W. Bader, "Tensor Decompositions and Applications," *SIAM Review*, vol. 51, no. 3, 2009.

[43] R. Bhatia and L. Elsner, "Positive linear maps and the lyapunov equation," in *Linear Operators and Matrices*, Birkhäuser Basel, 2002.

[44] T. Damm, "Stability of linear systems and positive semigroups of symmetric matrices," in *Positive Systems*, Springer Berlin Heidelberg, 2003.

[45] G. Lindblad, "On the generators of quantum dynamical semigroups," *Commun. Math. Phys.*, vol. 48, no. 2, 1976.

[46] D. Romero and G. Leus, "Wideband Spectrum Sensing From Compressed Measurements Using Spectral Prior Information," *IEEE Trans. Signal Process.*, vol. 61, no. 24, 2013.

[47] A. Kukush, I. Markovsky, and S. V. Huffel, "Consistent estimation in the bilinear multivariate errors-in-variables model," *Metrika*, vol. 57, no. 3, 2003.

[48] J. A. Tropp, "An Introduction to Matrix Concentration Inequalities," *Found. Trends Mach. Learn.*, vol. 8, no. 1-2, 2015.

[49] A. Ben-Tal, L. El Ghaoui, and A. Nemirovski, "Robustness," in *Handbook of Semidefinite Programming*, 2000.

[50] P. Coppens and P. Patrinos, "Sample Complexity of Data-Driven Stochastic LQR with Multiplicative Uncertainty," in *CDC*, IEEE, 2020.

[51] P. Coppens and P. Patrinos, "Data-driven distributionally robust MPC for constrained stochastic systems," *Control Syst. Lett.*, vol. 6, 2022.

[52] R. A. Horn and C. R. Johnson, *Topics in Matrix Analysis*. Cambridge University Press, 1991.

[53] E. de Klerk, *Aspects of Semidefinite Programming*, vol. 65. Springer US, 2002.

[54] M.-D. Choi, "Positive Linear Maps on Complex Matrices," *Linear Algebra and its Applications*, vol. 10, no. 3, 1975.

[55] R. A. Horn and C. R. Johnson, *Matrix Analysis*. Cambridge University Press, 2 ed., 2012.

[56] R. Jungers, *The Joint Spectral Radius*, vol. 385 of *Lecture Notes in Control and Information Sciences*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009.

[57] G. H. Golub and C. F. V. Loan, *Matrix Computations*. JHU Press, 4th ed., 2013.

[58] N. Lee and A. Cichocki, "Fundamental Tensor Operations for Large-Scale Data Analysis in Tensor Train Formats," 2014, arXiv: 1405.7786.

[59] G. R. Grimmett and D. Stirzaker, *Probability and random processes*. 2001.

[60] V. Balakrishnan and L. Vandenberghe, "Semidefinite programming duality and linear time-invariant systems," *IEEE Trans. Autom. Control*, vol. 48, no. 1, 2003.

[61] D. Testa, "3 × 3 singular matrices of linear forms," 2018, arXiv: 1701.06864v1.

[62] L. Lovász, "Singular spaces of matrices and their application in combinatorics," *Bol. Soc. Bras. Mat.*, vol. 20, no. 1, 1989.

[63] B. Chen and Z. Li, "On the tensor spectral p-norm and its dual norm via partitions," *Comput Optim Appl*, vol. 75, no. 3, 2020.

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

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

[66] J. R. Magnus and H. Neudecker, *Matrix Differential Calculus with Applications in Statistics and Econometrics*. Wiley Series in Probability and Statistics, Wiley, 2019.

[67] D. P. Bertsekas, *Dynamic Programming and Optimal Control – Vol II*. Athena scientific, 3rd ed., 2005.

[68] M. Ledoux and M. Talagrand, *Probability in Banach Spaces*. Springer Berlin Heidelberg, 1991.

[69] S. Boucheron, G. Lugosi, and P. Massart, *Concentration Inequalities*. Oxford University Press, 2013.

[70] M. L. Eaton and M. D. Perlman, "The Non-Singularity of Generalized Sample Covariance Matrices," *The Annals of Statistics*, vol. 1, no. 4, 1973.

[71] T. N. E. Greville, "Note on the Generalized Inverse of a Matrix Product," *SIAM Review*, vol. 8, no. 4, 1966.

APPENDIX A
PROOFS FOR PRELIMINARIES

Proof of Lem. III.3: We begin by noting that

$$\begin{aligned}\|\mathcal{S}\|_2^2 &\le \max_{X\in\mathbb{S}^m}\{\|\mathcal{S}(X)\|_F^2 \colon \|X\|_F \le \sqrt{m}\}\\ &\le m \max_{x\in\mathbb{R}^{3m}}\{x^\top \boldsymbol{S}^\top \boldsymbol{S} x \colon \|x\|_2 \le 1\},\end{aligned}$$

where we used $\|A\|_2 \le \|A\|_F \le \sqrt{\mathrm{rk}(A)}\|A\|_2$ for any $A \in \mathbb{S}^d$ for the first inequality and $\|X\|_F = \|\mathrm{vec_s}(X)\|_2$ for the second. The proof is then concluded by noting that the final expression equals $m\lambda_1(\boldsymbol{S}^\top\boldsymbol{S}) = m\|\boldsymbol{S}\|_2^2$. □

Proof of Lem. III.12: Since $W \succeq 0$, there is some C s.t. $W = CC^\top$. Then,

$$\mathcal{S}(X) = \boldsymbol{\mathcal{A}}_{(1)}(C\otimes I)(I\otimes X)(C^\top\otimes I)\boldsymbol{\mathcal{A}}_{(1)}^\top.$$

Letting $\boldsymbol{\mathcal{V}} = \boldsymbol{\mathcal{A}} \times_1 C^\top$ and applying Prop. III.11, implies $I\boldsymbol{\mathcal{A}}_{(1)}(C\otimes I) = (\llbracket \boldsymbol{\mathcal{A}}; I, I, C^\top\rrbracket)_{(3)} = \boldsymbol{\mathcal{V}}_{(3)}$. Therefore $\mathcal{S}(X) = \boldsymbol{\mathcal{V}}_{(1)}(I\otimes X)\boldsymbol{\mathcal{V}}_{(1)}^\top = \sum_{i=1}^r [\boldsymbol{\mathcal{V}}]_{::i} X [\boldsymbol{\mathcal{V}}]_{::i}^\top$, as in (5). □

Proof of Cor. III.14: Let P denote the *perfect shuffle* [52, Thm. 4.3.8] with $P^\top P = I$. Then

$$\begin{aligned}\mathcal{S}(W;X) &= \boldsymbol{\mathcal{A}}_{(1)}P^\top P(W\otimes X)P^\top P\boldsymbol{\mathcal{A}}_{(1)}^\top\\ &= \boldsymbol{\mathcal{A}}_{(1)}P^\top(X\otimes W)P\boldsymbol{\mathcal{A}}_{(1)}^\top,\end{aligned}$$

by [52, Cor. 4.3.10], which is of the form (11). So the conclusion follows by Lem. III.12. The second result follows from $\mathcal{S}(W;X)$ being CP and linearity. □

Proof of Lem. III.15: From [58, Prop. 2.2] we have

$$\boldsymbol{\mathcal{Y}} \otimes \boldsymbol{\mathcal{Y}} = [\![\boldsymbol{\mathcal{T}} \otimes \boldsymbol{\mathcal{T}}; X_1 \otimes X_1, X_2 \otimes X_2, X_3 \otimes X_3]\!]. \quad (41)$$

For the proof we focus on one argument for conciseness. So let $\boldsymbol{\mathcal{Y}} = [\![\boldsymbol{\mathcal{T}}; X_1; \ldots]\!]$. Then

$$\begin{aligned}\boldsymbol{\mathcal{Y}} \otimes_{\mathrm{s}} \boldsymbol{\mathcal{Y}} &\overset{(a)}{=} [\![\boldsymbol{\mathcal{Y}} \otimes \boldsymbol{\mathcal{Y}}; Q_{p_1}, \ldots]\!] \\ &\overset{(b)}{=} [\![[\![\boldsymbol{\mathcal{T}} \otimes \boldsymbol{\mathcal{T}}; X_1 \otimes X_1, \ldots]\!]; Q_{p_1}, \ldots]\!] \\ &\overset{(c)}{=} [\![\boldsymbol{\mathcal{T}} \otimes \boldsymbol{\mathcal{T}}; Q_{p_1}(X_1 \otimes X_1), \ldots]\!].\end{aligned}$$

Here *(a)* follows by definition of $\otimes_{\mathrm{s}}$, *(b)* from (41) and *(c)* follows by noting $(\boldsymbol{\mathcal{Y}} \times_1 A) \times_1 B = \boldsymbol{\mathcal{Y}} \times_1 BA$ [42, §2.5]. Next for every $H \in \mathbb{S}^{p_1}$ we have $\mathrm{vec_s}^\top(H) Q_{p_1}(X_1 \otimes X_1) = \mathrm{vec}^\top(X_1 H X_1^\top)$, by Def. III.1 and (3). The result is a vectorization of a symmetric matrix in $\mathbb{S}^{q_1}$. Hence, since $Q_{q_1}^\top Q_{q_1}$ acts like the identity on such vectors, we have $h^\top Q_{p_1}(X_1 \otimes X_1) = h^\top Q_{p_1}(X_1 \otimes X_1) Q_{q_1}^\top Q_{q_1}$ for all $h \in \mathbb{R}^{\Delta_{q_1}}$. Hence, reversing the steps of the earlier display, gives

$$\begin{aligned}\boldsymbol{\mathcal{Y}} \otimes_{\mathrm{s}} \boldsymbol{\mathcal{Y}} &= [\![\boldsymbol{\mathcal{T}} \otimes \boldsymbol{\mathcal{T}}; Q_{p_1}(X_1 \otimes X_1) Q_{q_1}^\top Q_{q_1}, \ldots]\!] \\ &= [\![[\![\boldsymbol{\mathcal{T}} \otimes \boldsymbol{\mathcal{T}}; Q_{q_1}, \ldots]\!]; Q_{p_1}(X_1 \otimes X_1) Q_{q_1}^\top, \ldots]\!] \\ &= [\![\boldsymbol{\mathcal{T}} \otimes_{\mathrm{s}} \boldsymbol{\mathcal{T}}; X_1 \otimes_{\mathrm{s}} X_1, \ldots]\!].\end{aligned}$$

Thus we have shown the required result. □

We can use Lem. III.15 to show

Proposition A.1. *Let $\mathcal{S}(W;X)$ be a CP operator parameterized as in (11) with some $\boldsymbol{\mathcal{V}}$. Then $\mathrm{Tr}[P\mathcal{S}(W;X)]$ equals:*

(i) $\mathrm{Tr}[P\boldsymbol{\mathcal{V}}_{(1)}(W \otimes X)\boldsymbol{\mathcal{V}}_{(1)}^\top] = \mathrm{Tr}[X\boldsymbol{\mathcal{V}}_{(2)}(W \otimes P)\boldsymbol{\mathcal{V}}_{(2)}^\top]$;
$= \mathrm{Tr}[W\boldsymbol{\mathcal{V}}_{(3)}(P \otimes X)\boldsymbol{\mathcal{V}}_{(3)}^\top]$

(ii) $[\![\boldsymbol{\mathcal{V}} \otimes \boldsymbol{\mathcal{V}}; \mathrm{vec}(P), \mathrm{vec}(X), \mathrm{vec}(W)]\!]$;

(iii) $[\![\boldsymbol{\mathcal{V}} \otimes_{\mathrm{s}} \boldsymbol{\mathcal{V}}; \mathrm{vec_s}(P), \mathrm{vec_s}(X), \mathrm{vec_s}(W)]\!]$;

for all $P \in \mathbb{S}^{n_p}, X \in \mathbb{S}^{n_x}$ and $W \in \mathbb{S}^{n_w}$.

Proof. By the eigenvalue decomposition, $P = \sum_{i=1}^{n_p} \alpha_i p_i p_i^\top$, $X = \sum_{j=1}^{n_x} \beta_j x_j x_j^\top$ and $W = \sum_{k=1}^{n_w} \gamma_k w_k w_k^\top$.

Writing out $\mathrm{Tr}[P\mathcal{S}(W;X)]$ with Lem. III.12 gives

$$\begin{aligned}&\textstyle\sum_{i,j,k} \alpha_i\beta_j\gamma_k \,\mathrm{Tr}[p_i p_i^\top \boldsymbol{\mathcal{V}}_{(1)}(w_k w_k^\top \otimes x_j x_j^\top)\boldsymbol{\mathcal{V}}_{(1)}^\top] \\ &\quad = \textstyle\sum_{i,j,k} \alpha_i\beta_j\gamma_k [(p_i^\top \boldsymbol{\mathcal{V}}_{(1)}(w_k \otimes x_j))^2].\end{aligned}$$

We show how the expression in square brackets can be rewritten. This, through reversal of the steps above enables proving the results.

By Prop. III.11 $(p^\top \boldsymbol{\mathcal{V}}_{(1)}(w \otimes x))^2 = [\![\boldsymbol{\mathcal{V}}; p, x, w]\!]^2$. Unfolding along $n = 2$ and 3 shows the other equalities in Prop. A.1*(i)*. Moreover since $[\![\boldsymbol{\mathcal{V}}; p, x, w]\!]^2 = [\![\boldsymbol{\mathcal{V}}; p, x, w]\!] \otimes [\![\boldsymbol{\mathcal{V}}; p, x, w]\!]$ we can apply Lem. III.15 to analogously get Prop. A.1*(ii)*. Finally note that $Q_{q_1}^\top Q_{n_p} \mathrm{vec}(P) = Q_{n_p}^\top \mathrm{vec_s}(P)$ since $P \in \mathbb{S}^{n_p}$ and similarly for X and W. From [42, §2.5] $\boldsymbol{\mathcal{T}} \,\overline{\times}_n\, Q^\top x = (\boldsymbol{\mathcal{T}} \times_1 Q) \,\overline{\times}_1\, x$, which combined with the definition of $\boldsymbol{\mathcal{V}} \otimes_{\mathrm{s}} \boldsymbol{\mathcal{V}}$ shows Prop. A.1*(iii)*. □

Proof of Cor. III.16: We can directly apply Prop. A.1*(ii)* to show the expression for $\mathcal{S}^*(W;X)$. For $\boldsymbol{S}$ note that $\mathrm{Tr}[P\mathcal{S}(W;X)] = \mathrm{vec_s}^\top(P)\boldsymbol{S}\mathrm{vec_s}(X)$. We can rewrite the Tucker operator in Prop. A.1*(iii)* as

$$\begin{aligned}&[\![\boldsymbol{\mathcal{V}} \otimes_{\mathrm{s}} \boldsymbol{\mathcal{V}}; \mathrm{vec_s}(P), \mathrm{vec_s}(X), \mathrm{vec_s}(W)]\!] \\ &\quad = (((\boldsymbol{\mathcal{V}} \otimes_{\mathrm{s}} \boldsymbol{\mathcal{V}}) \,\overline{\times}_3\, \mathrm{vec_s}(W)) \,\overline{\times}_1\, \mathrm{vec_s}(P))\mathrm{vec_s}2\mathrm{vec_s}(X) \\ &\quad = \mathrm{vec_s}(P)^\top((\boldsymbol{\mathcal{V}} \otimes_{\mathrm{s}} \boldsymbol{\mathcal{V}}) \,\overline{\times}_3\, \mathrm{vec_s}(W))\mathrm{vec_s}2\mathrm{vec_s}(X).\end{aligned}$$

Since this holds for all $P \in \mathbb{S}^m$, $X \in \mathbb{S}^n$ we have shown $(\boldsymbol{\mathcal{V}} \otimes_{\mathrm{s}} \boldsymbol{\mathcal{V}}) \,\overline{\times}_3\, \mathrm{vec_s}(W) = \boldsymbol{S}$.

Appendix B
Linear Quadratic Regulation

This section gives proves related to LQR. We begin by introducing some auxiliary operators and their adjoints:

Lemma B.1. *Let $\mathcal{E}$ be as in (15), $\Pi_K(X) = [I;K]X[I,K^\top]$, $\mathcal{E}_K(X) = \mathcal{E}(\Pi_K(X))$. Consider the Lyapunov operator $\mathcal{L}(Z) = X - \mathcal{E}(Z)$ and $\mathcal{L}_K(X) = X - \mathcal{E}_K(X)$. Here X is the top-left $n_x \times n_x$ block of Z.*

Then we have the adjoints:

(i) $\Pi_K^*(H) = [I, K^\top]H[I;K]$;

(ii) $\mathcal{E}^*(P) = \begin{bmatrix}\mathcal{F}^*(P) & (\mathcal{H}^*)^\top(P) \\ \mathcal{H}^*(P) & \mathcal{G}^*(P)\end{bmatrix}$;

(iii) $\mathcal{E}_K^*(P) = \Pi_K^*(\mathcal{E}^*(P))$;

(iv) $\mathcal{L}^*(P) = \mathrm{diag}(P,0) - \mathcal{E}^*(P)$;

(v) $\mathcal{L}_K^*(P) = \Pi_K^*(\mathcal{L}^*(P))$.

The elementary adjoints are $\mathcal{F}^*(P) = \boldsymbol{\mathcal{A}}_{(2)}(W \otimes P)\boldsymbol{\mathcal{A}}_{(2)}^\top$, $\mathcal{G}^*(P) = \boldsymbol{\mathcal{B}}_{(2)}(W \otimes P)\boldsymbol{\mathcal{B}}_{(2)}^\top$ and $\mathcal{H}^*(P) = \boldsymbol{\mathcal{B}}_{(2)}(W \otimes P)\boldsymbol{\mathcal{A}}_{(2)}^\top$.

Proof. Only elementary algebra, Lem. III.12 and Cor. III.16 are used. □

Lemma B.2. *Any $Z \succeq 0$ can be partitioned as:*

$$\begin{bmatrix} X & XK^\top \\ KX & KXK^\top + \Delta \end{bmatrix}$$

for $\Delta \succeq 0$ and $X \succeq 0$.

Proof. Let $Z = [X, V; V^\top, U]$. By [65, §A.5.5] $Z \succeq 0$ iff

$$X \succeq 0, \quad \mathrm{im}\, V \subseteq \mathrm{im}\, X, \quad U - V^\top X^\dagger V \succeq 0.$$

So there exists a K s.t. $V = XK^\top$ and $\Delta = U - V^\top X^\dagger V = U - K^\top X K \succeq 0$ □

Proof of Prop. IV.6: We use Asm. I and Lem. IV.1 to consider (13) instead of (1) without loss of generality. We start with a feasible sequence for $(\mathcal{LQR}_\star)$ and construct an (equivalent) feasible sequence for $(\mathcal{LQR}_{\mathrm{cp}})$. Let $H = \mathrm{diag}(Q,R)$, $Z_t = \mathbb{E}[z_t z_t^\top]$ with $z_t = (x_t, u_t)$ and x_t, u_t satisfying (13). Note that, by (15), any sequence Z_t constructed as such satisfies the constraints of $(\mathcal{LQR}_{\mathrm{cp}})$. Then

$$\mathbb{E}\left[\textstyle\sum_{t=0}^\infty x_t^\top Q x_t + u_t^\top R u_t\right] = \textstyle\sum_{t=0}^\infty \mathrm{Tr}[HZ_t]. \quad (42)$$

So the cost of (x_t, u_t) for $(\mathcal{LQR}_\star)$ equals that of the constructed Z_t in $(\mathcal{LQR}_{\mathrm{cp}})$. So $\mathrm{val}(\mathcal{LQR}_{\mathrm{cp}}) \leq \mathrm{val}(\mathcal{LQR}_\star)$.

Next we argue that for any sequence Z_t feasible for ($\mathcal{LQR}_{\text{cp}}$), we can construct a sequence $z_t = (x_t, u_t)$ feasible for ($\mathcal{LQR}_\star$). Note that for any $X \succeq 0$ we can easily construct a random vector x with $\mathbb{E}[xx^\top] = X$. So let x_0 be such a vector for X_0. Then, by Lem. B.2, Z_0 can be partitioned as $[X_0, X_0K_0^\top; K_0X_0, K_0X_0K_0^\top + \Delta_0]$ with $\Delta_0 \succeq 0$ for which we pick a random vector δ_0. So if we take $u_0 = K_0x_0 + \delta_0$ and $z_0 = (x_0, u_0)$ then $\mathbb{E}[z_0z_0^\top] = Z_0$. Repeating the same argument starting at x_1 and continuing for all time steps allows us to construct a feasible trajectory for ($\mathcal{LQR}_\star$). Again by $\mathbb{E}[z_tz_t^\top] = Z_t$ and (42) the cost for both trajectories is the same. Thus $\text{val}(\mathcal{LQR}_{\text{cp}}) \geq \text{val}(\mathcal{LQR}_\star)$. □

We continue with proving auxiliary results for Thm. IV.8. We first consider the finite horizon version of ($\mathcal{LQR}_{\text{cp}}$):

$$\begin{aligned} \underset{Z_t}{\text{minimize}} \quad & \sum_{t=0}^{N} \text{Tr}[Z_tH] + \text{Tr}[PX_N] \\ \text{subj. to} \quad & Z_t = \begin{bmatrix} X_t & V_t \\ V_t^\top & U_t \end{bmatrix} \succeq 0, \\ & X_{t+1} = \mathcal{E}(Z_t), \quad \forall t \in \mathbb{N}_{0:N-1}. \end{aligned} \tag{43}$$

We assume $H = \text{diag}(Q, R) \succ 0$ and denote the optimal value as $\mathcal{J}^N(P; X_0)$ and $\mathcal{J}^\star(X_0) := \lim_{N\to\infty} \mathcal{J}^N(0; X_0)$.

We prove the fundamental property linking $\mathcal{R}$ to (43):

Proposition B.3. *Consider $\mathcal{J}^N$ and $\mathcal{R}$ defined in §IV-C. Let*

$$P_k := \mathcal{R}^k(P),\ K_k = -(R + \mathcal{G}^*(P_k))^{-1}\mathcal{H}^*(P_k),\ k \in \mathbb{N}_{0:N}.$$

Then $Z_t = \Pi_{K_{N-t-1}}(X_t)$ for $t \in \mathbb{N}_{0:N-1}$ achieves the minimum:

$$\mathcal{J}^N(P; X_0) = \text{Tr}[P_NX_0].$$

Proof. The result follows from *dynamic programming* [3, §4.1]. Starting from $\mathcal{J}^0(P; X_0) = \text{Tr}[P_0X_0]$, with $P_0 = P$, dp. constructs $\mathcal{J}^k(P; X_0) = \text{Tr}[P_kX_0]$ using:

$$\text{Tr}[P_{k+1}X] = \min_{Z\succeq 0}\{\text{Tr}[ZH] + \text{Tr}[P_k\mathcal{E}(Z)]\}, \tag{44}$$

Note that $\text{Tr}[P_k\mathcal{E}(Z)] = \text{Tr}[\mathcal{E}^*(P_k)Z]$ by definition of the adjoint. We use Lem. B.1 to write $\mathcal{E}^*$ in terms of $\mathcal{F}^*$, $\mathcal{G}^*$ and $\mathcal{H}^*$. Introducing X, K and Δ to partition Z as in Lem. B.2 allows rewriting the cost of (44) (omitting constant terms) as:

$$\begin{aligned} &\text{Tr}[(\mathcal{H}^*)^\top(P_k)KX] + \text{Tr}[\mathcal{H}^*(P_k)XK^\top] \\ &\quad + \text{Tr}[(R + \mathcal{G}^*(P_k))KXK^\top] + \text{Tr}[(R + \mathcal{G}^*(P_k))\Delta], \end{aligned}$$

which is minimized uniquely at $\Delta = 0$ by $R \succ 0$ and Lem. F.1. We can evaluate the gradient of the remaining terms with respect to K using the identities in [66, Table. 9.4]. This gives the following first-order optimality conditions:

$$2\mathcal{H}^*(P_k)X + 2(\mathcal{G}^*(P_k) + R)KX = 0, \tag{45}$$

for all $X \in \mathbb{S}^{n_x}$. Since $R \succ 0$ this is the case iff

$$K = -(R + \mathcal{G}^*(P_k))^{-1}\mathcal{H}^*(P_k) = K_k.$$

For these values of K and Δ we have $Z = \Pi_{K_k}(X)$, with Π_{K_k} as in Lem. B.1. Plugging into (44):

$$\begin{aligned} \text{Tr}[P_{k+1}X] &= \text{Tr}[\Pi_{K_k}(X)H] + \text{Tr}[P_k\mathcal{E}(\Pi_{K_k}(X))] \\ &= \text{Tr}[X\Pi^*_{K_k}(H + \mathcal{E}^*(P_k))] = \text{Tr}[X\mathcal{R}(P_k)], \end{aligned} \tag{46}$$

where the second equality follows from linearity and adjoints, while the third follows from Lem. B.1 and cancellation of the inverses in K_k. Iterating over k, repeatedly computing the optimal policy K_k by the same procedure as usual in dynamic programming (cf. [3, Prop. 1.3.1]), completes the proof. Note that k counts backwards in time, hence K_{N-t-1} in Z_t. □

The steps in the proof of Prop. B.3 can be used to prove some additional properties of $\mathcal{R}$:

Corollary B.4. *For all $P \succeq 0$, $K' \neq -(R + \mathcal{G}^*(P))^{-1}\mathcal{H}^*(P)$ and $H = \text{diag}(Q, R) \succ 0$:*

(i) $\mathcal{R}(P') \preceq \mathcal{R}(P)$ if $P' \preceq P$;

(ii) $\mathcal{R}(P) \not\succeq \Pi^*(K'; H + \mathcal{E}^*(P))$.

Proof. Result (i) is simply monotonicity of the Bellman operator [67, Lem. 1.1.1]. Specifically, by observing (44) it is clear that the cost is larger for P than for P' since $\text{Tr}[P\mathcal{E}(Z)] \geq \text{Tr}[P'\mathcal{E}(Z)]$ for all Z. Hence the minimum is also larger. Moreover using (46) gives $\text{Tr}[X\mathcal{R}(P)] \geq \text{Tr}[X\mathcal{R}(P')]$. Taking $X = xx^\top$ for any x gives $x^\top\mathcal{R}(P)x \geq x^\top\mathcal{R}(P')x$ thus showing (i). Result (ii) follows from $H = \text{diag}(Q, R) \succ 0$, causing the minimizer we computed for (44) to be unique since it should satisfy (45). □

Proposition B.5. *Assume $\exists K\colon \rho(\mathcal{E}_K) < 1$ and $H = \text{diag}(Q, R) \succ 0$. Then $\forall P \succeq 0$, $\exists P_\star \succeq 0\colon \lim_{t\to\infty} \mathcal{R}^t(P) = P_\star$. Moreover*

(i) $P_\star$ uniquely solves $\mathcal{R}(P_\star) = P_\star$ and $P_\star \succ 0$;

(ii) For $K_\star = -(R + \mathcal{G}^*(P_\star))^{-1}\mathcal{H}^*(P_\star)$, $\rho(\mathcal{E}(K_\star)) < 1$.

Proof. The proof follows that of [3, Prop. 4.4.1]. The sequence $\{\mathcal{R}^k(0)\}_{k\in\mathbb{N}}$ is nondecreasing in the psd. sense (by $H \succ 0$ and Cor. B.4*(i)*) and bounded (by $\exists K\colon \rho(\mathcal{E}_K) < 1$ which implies boundedness of $\mathcal{J}^k(0; X_0)$ and hence, by Prop. B.3 boundedness of $R^k(P)$). The sequence therefore converges to some $P_\star = \mathcal{R}(P_\star)$.

The equality $P_\star = \mathcal{R}(P_\star)$, using (46), can be rewritten as a Lyapunov equality

$$\Pi^*_{K_\star}(H) = P_\star - \mathcal{E}^*_{K_\star}(P_\star). \tag{47}$$

We can therefore use Prop. III.8 to prove stability for $K_\star$ (i.e. (ii)) and $P_\star \succ 0$. To do so we need $H \succ 0$ since then $\Pi^*_{K_\star}(H) \succ 0$ by Lem. B.1.

To show uniqueness (i), note

$$\mathcal{J}^N(P; X_0) = \textstyle\sum_{t=0}^{N-1} \text{Tr}[\Pi_{K_{N-t-1}}(X_t)H] + \text{Tr}[PX_N],$$

with $X_{t+1} = \mathcal{E}_{K_{N-t-1}}(X_t)$ and K_{N-t-1} the optimal gain sequence realizing $\mathcal{J}^N(P; X_0)$ from Prop. B.3. The cost $\mathcal{J}^N(P; X_0)$ is clearly lower bounded by $\mathcal{J}^N(0; X_0)$. Moreover, if we instead apply the gains optimal for $\mathcal{J}^N(0; X_0)$

denoted here as K_k^0 with moment sequence X_t^0, we get an upper bound as well:

$$\mathcal{J}^N(0;X_0) \le \mathcal{J}^N(P;X_0) \le \mathcal{J}^N(0;X_0) + \operatorname{Tr}[PX_N^0].$$

As shown, $\lim_{N\to\infty}\mathcal{J}^N(0;X_0) = \lim_{N\to\infty}\operatorname{Tr}[\mathcal{R}^N(0)X_0]$ is finite $\forall X_0 \succeq 0$. Note that $H \succ 0$ and the cost involves $\operatorname{Tr}[X_t\Pi^*_{K_{N-t-1}}(H)]$ with the second factor being positive definite by earlier arguments. Thus $\lim_{N\to\infty}\mathcal{J}^N(0;X_0)$ must imply that $X_t \to 0$ (otherwise the trace terms remain positive for large N by Lem. F.1). Thus we have stability and $\lim_{k\to\infty}\operatorname{Tr}[PX_k^0] = 0$. Therefore $\mathcal{J}^N(0;X_0) = \mathcal{J}^N(P;X_0)$ when $N\to\infty$. This and Prop. B.3 imply that $\{\mathcal{R}^k(P)\}_{k\in\mathbb{N}}$ converges to the same $P_\star$ independent of P. Say $\exists P' \neq P_\star$: $P' = \mathcal{R}(P')$. Then $\lim_{t\to\infty}\mathcal{R}^t(P') = P' = P_\star$, which is a contradiction. So $P_\star$ is unique, completing the proof. □

We are now ready to prove the main LQR result:

Proof of Thm. IV.8: Stability follows directly from Prop. B.5. Taking $N\to\infty$ in Prop. B.3 shows (ii). Using (47) and Prop. III.8 gives $\operatorname{Tr}[P_\star X_0] = \operatorname{Tr}[H\sum_{t=0}^\infty \mathcal{E}_{K_\star}^t(X_0)]$. That is $Z_t = \Pi_{K_\star}(X_t)$ achieves the optimal cost. □

We conclude this section with the derivation of the SDP reformulations of the Riccati equation (17).

Proof of Thm. IV.10: We follow [60]. Introducing $\mathcal{L}(Z) = X - \mathcal{E}(Z)$ as in Lem. B.1 and its adjoint $\mathcal{L}^*$. It becomes clear from Lem. B.1 that the SDP in Thm. IV.10 is equivalent to

$$\text{maximize}\,\{\operatorname{Tr}[PX_0]\colon \mathcal{L}^*(P) \preceq H, P \succeq 0\}.$$

The Lagrangian dual of this SDP is

$$\text{minimize}\,\{\operatorname{Tr}[HZ]\colon \mathcal{L}(Z) \succeq X_0, Z \succeq 0\}.$$

By [60, Thm. 1] $\exists P \succ 0\colon H - L^*(P) \succ 0$ and $\exists Z \succneqq 0\colon \mathcal{L}(Z) = 0, \operatorname{Tr}[ZH] \preceq 0$ are strong alternatives. The second condition is false since $H \succ 0$. Meanwhile the condition $\exists Z \succ 0\colon \mathcal{L}(Z) \succ X_0$ holds from $\exists K\colon \rho(\mathcal{E}_K) < 1$, $Z = \Pi_K(X)$ and Prop. III.8. Therefore the primal problem is strictly feasible and we thus have strong duality [60, Thm. 4]. By Lem. B.2, we can partition $Z \succeq 0$ as $Z = \Pi_K(X) + \operatorname{diag}(0,\Delta)$ with $X \succeq 0$ and $\Delta \succeq 0$. By complementary slack [60, §II.C]:

$$\begin{bmatrix} I \\ K \end{bmatrix} X \begin{bmatrix} I & K^\top \end{bmatrix} \begin{bmatrix} Q - P + F^*(P) & (\mathcal{H}^*)^\top(P) \\ \mathcal{H}^*(P) & R + \mathcal{G}^*(P) \end{bmatrix} + (R + \mathcal{G}^*(P))\Delta = 0$$

for the optimum. Since $R \succ 0$, the second term is zero iff $\Delta = 0$. A candidate solution $(Z,P) = (\Pi_K(X), P)$ with $X \succ 0$ is then optimal iff

$$\begin{bmatrix} I & K \end{bmatrix} \begin{bmatrix} Q - P + F^*(P) & (\mathcal{H}^*)^\top(P) \\ \mathcal{H}^*(P) & R + \mathcal{G}^*(P) \end{bmatrix} = 0,$$

which holds iff $P = \mathcal{R}(P)$ (i.e. $P = P_\star$) and $K = K_\star$. So $P_\star$ is an optimal solution to the SDP.

If $X_0 \succ 0$, then Z for $\Delta = 0$ is feasible iff $X - \mathcal{E}_K(X) \succeq X_0 \succ 0$. So iff $X \succ 0$, implying $P_\star$ is the only optimum. □

Next we prove the properties of (35).

Proof of Thm. VII.5: Note that, by $H \succ 0$, $\Pi_K^*(H) = Q + K^\top RK \succ 0$. Hence any feasible pair (P,K) satisfies a Lyapunov equation for some $Q' \succeq \Pi_K^*(H)$. From Lem. III.6 stability then follows and

$$\begin{aligned}\forall X\colon\ \operatorname{Tr}[PX] &= \operatorname{Tr}[\textstyle\sum_{t=0}^\infty \mathcal{E}_K^t(X)Q'] \\ &\ge \operatorname{Tr}[\textstyle\sum_{t=0}^\infty \mathcal{E}_K^t(X)\Pi_K^*(H)] \ge \mathcal{J}^\star(X_0).\end{aligned}$$

From (47), $(P_\star, K_\star)$ are feasible for (35) and satisfy $\operatorname{Tr}[P_\star X_0] = \mathcal{J}^\star(X_0)$ by Prop. B.5. Therefore this pair achieves the lower bound and is optimal.

Assume $X_0 \succ 0$ and that there is a feasible $(P',K') \neq (P_\star,K_\star)$: $\operatorname{Tr}[P'X_0] = \operatorname{Tr}[P_\star X_0]$. By feasiblity and Cor. B.4*(ii)*: $P' \succeq \Pi^*_{K'}(H + \mathcal{E}^*(P')) \succeq \mathcal{R}(P')$. So, by repeatedly applying Cor. B.4*(i)* we have $P' \succeq \mathcal{R}(P') \succeq \mathcal{R}^2(P) \succeq \dots$. By induction we then have $P' \succeq \lim_{k\to\infty}\mathcal{R}^k(P') = P_\star$. Hence, by $X_0 \succ 0$, $P' \neq P_\star$ and Lem. F.1 we have $\operatorname{Tr}[P'X_0] > \operatorname{Tr}[P_\star X_0]$ which is a contradiction since we showed $\operatorname{Tr}[P'X_0] = \operatorname{Tr}[P_\star X_0]$ before. Thus $P' = \Pi^*_{K'}(H + \mathcal{E}^*(P')) = \mathcal{R}(P')$, which by Cor. B.4*(ii)* and Prop. B.5 holds iff $(P',K') = (P_\star,K_\star)$.

Appendix C
Vector and Matrix Concentration

Lemma C.1. *For some $\gamma \in \mathbb{R}^N$ and a random, independent sequence $\{X_i\}_{i\in\mathbb{N}_{1:N}} \subset \mathbb{S}^d$ with $\|X_i\|_2 \le [\gamma]_i$ a.s. and $\mathbb{E}[X_i] = 0$. Then, for $\beta \ge 0$, $Y = \sum_{i=1}^N X_i$,*

$$\mathbb{P}\left[\lambda_1(Y) \ge \beta\right] \le \delta; \qquad \mathbb{P}\left[\lambda_d(Y) \le -\beta\right] \le \delta,$$

for $\delta = d\exp(-\beta^2/2\|\gamma\|_2^2)$. Moreover, $\mathbb{P}\left[\|Y\|_2 \ge \beta\right] \le 2\delta$.

Proof. We extend the proof of [51, Lem II.2]. Let $X_i = V_i\Lambda_i V_i^\top$ be the eigenvalue decomposition. Hence

$$\begin{aligned}\mathbb{E}[\exp(\theta X_i)] &= \mathbb{E}[V_i\exp(\theta\Lambda_i)V_i^\top] \\ &\preceq \mathbb{E}\left[\frac{[\gamma]_i I + X_i}{2[\gamma]_i}e^{\theta[\gamma]_i} + \frac{[\gamma]_i I - X_i}{2[\gamma]_i}e^{-\theta[\gamma]_i}\right]\end{aligned}$$

by convexity. Also from $\mathbb{E}[X_i] = 0$ we have, $\mathbb{E}[\exp(\theta X_i)] \preceq \cosh(\theta[\gamma]_i)I \preceq \exp(\theta^2[\gamma]_i^2/2)I$.

Applying [48, Thm. 3.6.1] results in

$$\begin{aligned}\mathbb{P}[\lambda_1(Y) \ge \beta] &\le \inf_{\theta>0} d\exp(\theta\|\gamma\|_2^2/2 - \theta\beta) \\ \mathbb{P}[\lambda_d(Y) \le -\beta] &\le \inf_{\theta<0} d\exp(\theta\|\gamma\|_2^2/2 - \theta\beta).\end{aligned}$$

Minimizing over θ gives the required result for λ_1 and λ_d. Taking a union bound gives the spectral norm bound. □

Lemma C.2. *Let $\gamma \in \mathbb{R}^N$ and consider a random, independent sequence $\{x_i\}_{i\in\mathbb{N}_{1:N}} \subset \mathbb{R}^d$ with $\|x_i\| \le [\gamma]_i$ a.s. Then, for $\beta \ge 0$, $y = \sum_{i=1}^N (x_i - \mathbb{E}[x_i])$,*

$$\mathbb{P}[\|y\| \ge 2\|\gamma\|_2 + \beta] \le \delta,$$

for $\delta = \exp(-\beta^2/2\|\gamma\|_2^2)$.

Proof. We use the techniques of [68, §6]. Introducing a filtration, let $\mathscr{F}_i$ the σ-algebra generated by $\{x_i\}_{i\in\mathbb{N}_{1:N}}$ with $\mathscr{F}_0$ the trivial algebra. Let $\mathbb{E}_i[x] = \mathbb{E}[x \mid \mathscr{F}_i]$, $\Delta\mathbb{E}_i[x] = \mathbb{E}_i[x] - \mathbb{E}_{i-1}[x]$ and $d_i = \Delta E_i[\|y\|_2]$. Then, by telescoping over i, we have $\sum_{i=1}^N d_i = \|y\| - \mathbb{E}[\|y\|]$. Noting that $y - x_i$

does not depend on x_i and letting $\delta_i = \|y\|_2 - \|y - x_i\|_2$ gives

$$\Delta\mathbb{E}_i[\delta_i] = d_i - \mathbb{E}_i[\|y - x_i\|_2] + \mathbb{E}_{i-1}[\|y - x_i\|_2] = d_i.$$

Then, by the triangle inequality, $|\delta_i| = |\|y\|_2 - \|y - x_i\|_2| \leq \|y - y + x_i\|_2 = [\gamma]_i$. Letting $e_i := \mathbb{E}_{i-1}[\delta_i]$ then gives $|d_i - e_i| = |\mathbb{E}_i\delta_i| \leq [\gamma]_i$.

Hence we have constructed a sequence of zero-mean, independent random variables $\{d_i\}_{i\in\mathbb{N}_{1:N}}$ that take their value a.s. in $[-[\gamma]_i + e_i, [\gamma]_i + e_i]$. Applying a classical Hoeffding bound [69, Thm. 2.8] gives

$$\mathbb{P}[\textstyle\sum_{i=1}^N d_i \geq \beta] \leq \exp(-\beta^2/2\sum_{i=1}^N \|\gamma\|_2^2), \tag{48}$$

where, as shown before, $\sum_{i=1}^N d_i = \|y\|_2 - \mathbb{E}[\|y\|_2]$. We introduce a sequence of independent random vectors $\{v_i\}_{i\in\mathbb{N}_{1:N}}$ which is identically distributed to $\{x_i\}_{i\in\mathbb{N}_{1:N}}$. Let $\mathbb{E}_v[x]$ denote the expectation conditioned on $\{x_i\}_{i\in\mathbb{N}_{1:N}}$. Then, by Jensen's inequality, we can bound $\mathbb{E}[\|y\|_2]$ as

$$\begin{aligned}\mathbb{E}[\|\textstyle\sum_{i=1}^N x_i - \mathbb{E}[x_i]\|_2] &= \mathbb{E}[\|\mathbb{E}_v[\sum_{i=1}^N x_i - v_i]\|_2] \\ &\leq \mathbb{E}[\|\textstyle\sum_{i=1}^N (x_i - v_i)\|_2].\end{aligned}$$

Using a classical symmetrization argument, introducing independent Rademacher random variables $\{\epsilon_i\}_{i\in\mathbb{N}_{1:N}}$, gives $\mathbb{E}[\|\sum_{i=1}^N (x_i - v_i)\|_2] \leq 2\mathbb{E}[\|\sum_{i=1}^N x_i\epsilon_i\|_2]$. Applying Jensen once more, followed by a triangle inequality gives

$$\begin{aligned}\mathbb{E}[\|y\|_2] = 2\mathbb{E}[\|\textstyle\sum_{i=1}^N x_i\epsilon_i\|_2] &\leq 2\sqrt{\mathbb{E}[\|\sum_{i=1}^N x_i\epsilon_i\|_2^2]} \\ &\leq 2\sqrt{\mathbb{E}[\textstyle\sum_{i=1}^N \|x_i\|_2^2]} \leq 2\|\gamma\|_2.\end{aligned}$$

Plugging into (48) concludes the proof. □

Appendix D
Rank of Random Matrices

We introduce a fundamental lemma, regarding the span of random i.i.d. vectors [70, Thm. 3.3]:

Lemma D.1. *Given a sequence of i.i.d. copies $\{z_1, \ldots, z_N\}$ of a random vector $z\colon \Omega \to \mathbb{R}^d$ with $N \geq d$, the following statements are equivalent:*

(a) $\forall w \in \mathbb{R}^d,\ w \neq 0,\ \mathbb{P}[w^\top z = 0] = 0$;

(b) $\mathbb{P}[\mathrm{rk}([z_1, \ldots, z_N]) = d] = 1$.

Moreover if either (a) or (b) holds, then $\mathbb{E}[zz^\top] \succ 0$.

Proof. The proof of [70, Thm. 3.3] stands in a more general setting then necessary. We provide a simplified proof for completeness. Without loss of generality we assume $N = d$. Let $Z_k := [z_1, \ldots, z_k]$ and $\mathcal{B}[k]$ denote the event $\mathrm{rk}[Z_k] = k$.

(a) $\Leftarrow$ (b): We proceed by contradiction, so assume $\exists w \neq 0\colon \mathbb{P}[w^\top z = 0] > 0$ and that (b) holds (i.e. $\mathbb{P}[\mathcal{B}[d]] = 1$). Note that $\mathbb{P}[\neg\mathcal{B}[d]] = \mathbb{P}[\exists v \neq 0\colon v^\top Z_d = 0]$ (i.e., the kernel of $Z_d^\top$ is nontrivial). We have,

$$\begin{aligned}&\mathbb{P}[\exists v \neq 0\colon v^\top Z_d = 0] \\ &\quad\geq \mathbb{P}[w^\top Z_d = 0] = \Pi_{i=1}^d \mathbb{P}[w^\top z_i = 0] > 0,\end{aligned}$$

where the first inequality follows by fixing $v = w$, the equality follows by independence and the final inequality follows by the assumption that (a) does not hold. So $\mathbb{P}[\mathcal{B}[d]] \neq 1$ and we have shown (a) $\Leftarrow$ (b) by contradiction.

(a) $\Rightarrow$ (b): We can decompose $\mathbb{P}[\mathcal{B}[d]]$ into

$$\mathbb{P}[z_d \notin \mathrm{im}(Z_{d-1}) \mid \mathcal{B}[d-1]] \cdot \mathbb{P}[\mathcal{B}[d-1]]$$

We know that $\exists w \in \ker(Z_{d-1}^\top)/\{0\}$ since Z_{d-1} is not full rank. Hence the first factor is equal to $\mathbb{P}[w^\top z_d \neq 0 \mid \mathrm{rk}(Z_{d-1}) = d-1]$, where w depends on Z_{d-1}. This factor equals 1 since (a) holds. We can repeatedly apply the same decomposition to the second factor until we arrive at $\mathbb{P}[\mathrm{rk}(z_1) = 1]$, which equals to 1 since (a) also implies $\mathbb{P}[z_1 = 0] = 0$. Hence $\mathbb{P}[\mathrm{rk}(Z_d) = d] = 1$ so (b) holds.

Finally note that $\forall w,\ w^\top\mathbb{E}[zz^\top]w = \mathbb{E}(w^\top z)^2$. Since $\mathbb{P}[w^\top z = 0] = 0$ by (a), $\mathbb{E}[(w^\top z)^2] > 0$. So $\mathbb{E}[zz^\top] \succ 0$, completing the proof. □

We can consider some examples of distributions satisfying (a). A clear sufficient condition is that the density of z is dominated by the Lebesgue measure. Another example are uniform distributions over spheres. Meanwhile, any atomic measure will not satisfy (a).

Corollary D.2. *With the setup of Lem. D.1, assume that z is dominated by the Lebesgue measure. Then, if $N \geq \mathfrak{d}_d$,*

$$\mathbb{P}[\mathrm{rk}([z_1 \otimes_{\mathrm{s}} z_1, \ldots, z_N \otimes_{\mathrm{s}} z_N]) = \mathfrak{d}_d] = 1$$

and $\mathbb{E}[(zz^\top) \otimes_{\mathrm{s}} (zz^\top)] \succ 0$.

Proof. Applying Lem. D.1 to $z \otimes z$ instead of z gives:

$$\mathbb{P}[w^\top(z \otimes_{\mathrm{s}} z) = 0] = 0, \quad \forall w \in \mathbb{R}^{\mathfrak{d}_d}, w \neq 0. \tag{49}$$

The condition $w^\top(z \otimes_{\mathrm{s}} z) = 0 \Leftrightarrow z^\top W z$ describes the roots of a quadratic polynomial, with $W = \mathrm{unvec}_{\mathrm{s}}(w) \in \mathbb{S}^d$. It is well known that such a set is of Lebesgue measure zero, if the polynomial is not identically zero. Since $W \in \mathbb{S}^d$, $z^\top W z = 0$ for all z iff $W = 0$. Hence, by our domination assumption, $\mathbb{P}[w^\top(z \otimes_{\mathrm{s}} z) = 0] = 0$ follows. Applying Lem. D.1 and using $(z\otimes z)(z\otimes z)^\top = (zz^\top)\otimes(zz^\top)$ completes the proof. We can apply similar reasoning to prove the second result. □

Remark D.3. We could have directly used (49) as the assumption on z to get a tight result. We avoided doing so, since (49) is more difficult to verify in practice.

Appendix E
Identification Model

We begin by showing some auxiliary results.

Lemma E.1. *For $\boldsymbol{Z}_N$ as in (23a) and $N \geq \mathfrak{d}_{n_z}$,*

$$\boldsymbol{Z}_N^\dagger \boldsymbol{Z}_N = (\boldsymbol{\mathcal{W}}_{(3)}^\top)^\dagger \boldsymbol{\mathcal{W}}_{(3)}^\top,$$

with probability one over data $(z_i)_{i=1}^N$ satisfying Asm. II.

Proof. We can rewrite (23a) as

$$\boldsymbol{Z}_N = ([z_1 \otimes_{\mathrm{s}} z_1, \ldots, z_N \otimes_{\mathrm{s}} z_N]^\top \otimes I_{\mathfrak{d}_{n_x}})\boldsymbol{\mathcal{W}}_{(3)}^\top, \tag{50}$$

by $[\boldsymbol{\mathcal{T}} \,\overline{\times}_2\, x; \boldsymbol{\mathcal{T}} \,\overline{\times}_2\, y] = \boldsymbol{\mathcal{T}} \times_2 [x, y]^\top$ (cf. (8)) and Prop. III.11.

Let $Z_N := [z_1 \otimes_{\mathrm{s}} z_1, \ldots, z_N \otimes_{\mathrm{s}} z_N]^\top$. Then by $N \geq \mathfrak{d}_{n_z}$, Asm. II and Cor. D.2, $\mathrm{rk}(Z_N) = \mathfrak{d}_{n_z}$ w.p. 1. Hence $Z_N \otimes I$ is also left-invertible [52, Thm. 4.2.15] w.p. 1.

For two conformable matrices U, V with $U^\dagger U = I$, we then draw inspiration from the proof of [71, Thm. 1].

$$\begin{aligned} V^\dagger V &= V^\dagger U^\dagger U V \overset{(a)}{=} V^\dagger U^\dagger (UV)(UV)^\top ((UV)^\top)^\dagger \\ &= V^\dagger V V^\top U^\top ((UV)^\top)^\dagger \\ &\overset{(b)}{=} V^\top U^\top ((UV)^\top)^\dagger = (UV)^\dagger (UV), \end{aligned}$$

where both *(a)* and *(b)* use $A^\dagger A A^\top = A^\top$ (cf. [71, Eq. 2]).

Applying this to $\boldsymbol{Z}_N$ then gives the required result. □

Lemma E.2. *Take* $\mathrm{vec_s}(\tilde{W}) = (I - (\boldsymbol{\mathcal{W}}_{(3)}^\top)^\dagger \boldsymbol{\mathcal{W}}_{(3)}^\top)\mathrm{vec_s}(W)$ *for any* $W \in \mathbb{S}^{n_w}$. *Consider* $\mathcal{E}$ *as in (15). Then*

$$\mathcal{E}(W + \tilde{W}; Z) = \mathcal{E}(W; Z),$$

for all $W \in \mathbb{S}^{n_w}$ *and* $Z \in \mathbb{S}^{n_z}$.

Proof. From Prop. A.1*(iii)* and Prop. III.11 we have

$$\begin{aligned} \mathrm{Tr}[P\mathcal{E}(\tilde{W}; Z)] &= \left[\!\!\left[\boldsymbol{\mathcal{M}} \otimes_\mathrm{s} \boldsymbol{\mathcal{M}}; \mathrm{vec_s}(P), \mathrm{vec_s}(Z), \mathrm{vec_s}(\tilde{W}) \right]\!\!\right] \\ &= (\mathrm{vec_s}(P) \otimes \mathrm{vec_s}(Z)) \boldsymbol{\mathcal{W}}_{(3)}^\top \mathrm{vec_s}(\tilde{W}), \end{aligned}$$

for all $P \in \mathbb{S}^{n_x}$ and $Z \in \mathbb{S}^{n_z}$. By assumption, and $AA^\dagger A = A$,

$$\boldsymbol{\mathcal{W}}_{(3)}^\top \mathrm{vec_s}(\tilde{W}) = \boldsymbol{\mathcal{W}}_{(3)}^\top (I - (\boldsymbol{\mathcal{W}}_{(3)}^\top)^\dagger \boldsymbol{\mathcal{W}}_{(3)}^\top) \mathrm{vec_s}(W) = 0.$$

Therefore $\mathrm{Tr}[P\mathcal{E}(\tilde{W}; Z)] = 0$ for all P and Z, which is only possible if $\mathcal{E}(\tilde{W}; Z) = 0$ for all $Z \in \mathbb{S}^{n_z}$ (otherwise we can take $P = \mathrm{unvec}(e_i)$, with e_i the canonical basis vector and i the index of the nonzero element in $\mathrm{vec}(\mathcal{E}(\tilde{W}; Z))$). □

Next, we derive properties of $\mathcal{W}$ and $\mathcal{H}_i$ defined in (26).

Lemma E.3. *Let* $\mathcal{W}(Z) := \boldsymbol{\mathcal{W}}_{(3)}((Z \otimes_\mathrm{s} Z) \otimes I_{\mathfrak{d}_{n_x}}) \boldsymbol{\mathcal{W}}_{(3)}^\top$. *Then*

(i) $\mathcal{W}(zz^\top) = (\boldsymbol{\mathcal{W}} \,\overline{\times}_2\, (z \otimes_\mathrm{s} z))^\top (\boldsymbol{\mathcal{W}} \,\overline{\times}_2\, (z \otimes_\mathrm{s} z))$;

(ii) $\sum_{i=1}^N \mathcal{W}(z_i z_i^\top) = \boldsymbol{Z}_N^\top \boldsymbol{Z}_N$;

(iii) $\|\mathcal{W}\|_2 = \sup_z \{\|\mathcal{W}(zz^\top)\|_2 \colon \|z\|_2 \leq 1\} \leq \|\boldsymbol{\mathcal{W}}\|_2$,

with $\boldsymbol{Z}_N$ *as in (23a).*

Proof. The first result (i) follows by using Prop. III.11 to write

$$\begin{aligned} &(\boldsymbol{\mathcal{W}} \,\overline{\times}_2\, (z \otimes_\mathrm{s} z))^\top (\boldsymbol{\mathcal{W}} \,\overline{\times}_2\, (z \otimes_\mathrm{s} z)) \\ &\quad = \boldsymbol{\mathcal{W}}_{(3)} [(z \otimes_\mathrm{s} z) \otimes I][(z \otimes_\mathrm{s} z)^\top \otimes I] \boldsymbol{\mathcal{W}}_{(3)}^\top. \end{aligned}$$

Using $(A \otimes Y)(B \otimes Y) = AB \otimes Y$ [53, Thm. E.1.3] and [53, Lem. E.1.2] to argue $(z \otimes_\mathrm{s} z)(z \otimes_\mathrm{s} z)^\top = (zz^\top \otimes_\mathrm{s} zz^\top)$, proves (i). We can show (ii) by using (50) and applying similar tricks. Finally (iii) is shown by using (i) to argue

$$\begin{aligned} \|\mathcal{W}\|_2 &= \sup_{z \in \mathcal{B}} \{\|(\boldsymbol{\mathcal{W}} \,\overline{\times}_2\, (z \otimes_\mathrm{s} z))\|_2^2\} \\ &= \sup_{x,w,z \in \mathcal{B}} \{(x^\top (\boldsymbol{\mathcal{W}} \,\overline{\times}_2\, (z \otimes_\mathrm{s} z)) w)^2\}, \end{aligned} \tag{51}$$

where we used the variational representation of the spectral norm for the second equality and with $x \in \mathbb{R}^{\mathfrak{d}_{n_x}}$, $w \in \mathbb{R}^{\mathfrak{d}_{n_w}}$ and $\mathcal{B}$ the unit Euclidean ball of generic dimension. The squared quantity can be rewritten by noting

$$x^\top (\boldsymbol{\mathcal{W}} \,\overline{\times}_2\, (z \otimes_\mathrm{s} z)) w = [\![\boldsymbol{\mathcal{W}}; x, z \otimes_\mathrm{s} z, w]\!] .$$

Noting that $(z \otimes_\mathrm{s} z)^\top (z \otimes_\mathrm{s} z) = \|z\|_2^4$ implies that we can relax (51) to take the supremum over $z' \in \mathcal{B}_{\mathfrak{d}_{n_z}}$ instead of over $z \otimes_\mathrm{s} z$. So

$$\|\mathcal{W}\|_2 \leq \left(\sup_{x,w,z' \in \mathcal{B}} \{[\![\boldsymbol{\mathcal{W}}; x, z', w]\!]\} \right)^2,$$

which, plugging in the definition of the tensor spectral norm [63], implies (iii). □

A direct consequence of Lem. E.3 is that

Corollary E.4. *Given* $\boldsymbol{Z}_N$, E_N *as in (23a) and (23b) respectively and* $\mathcal{W}$ *as in Lem. E.3. Then*

$$\begin{aligned} \boldsymbol{Z}_N^\dagger E_N &= [\textstyle\sum_{j=1}^N \mathcal{W}(z_j z_j^\top)]^\dagger (\sum_{i=1}^N \mathcal{W}(z_i z_i^\top)\eta_i) \\ &= \mathrm{vec_s}(\textstyle\sum_{i=1}^N \mathcal{H}_i(w_i w_i^\top - W)), \end{aligned}$$

with $\mathcal{H}_i$ *the linear operator with matrix* $\boldsymbol{H}_i$ *as in (26).*

Proof. Note that

$$\boldsymbol{Z}_N^\dagger E_N = (\boldsymbol{Z}_N^\top \boldsymbol{Z}_N)^\dagger \boldsymbol{Z}_N^\top E_N.$$

Then, the first equality follows by application of Lem. E.3*(ii)* to replace $(\boldsymbol{Z}_N^\top \boldsymbol{Z}_N)$ and Lem. E.3*(i)* to show that $\boldsymbol{Z}_N^\top E_N = \sum_{i=1}^N \mathcal{W}(z_i z_i^\top)\eta_i$. The second equality follows by definition of $\boldsymbol{H}_i$ and $\eta_i = \mathrm{vec_s}(w_i w_i^\top - W)$. □

We are now ready to prove the data-driven bound.

Proof of Thm. V.4: By the classical LS error equation (25):

$$\mathrm{vec_s}(\hat{W} - W) = (I - \boldsymbol{Z}_N^\dagger \boldsymbol{Z}_N)\mathrm{vec_s}(W) + \boldsymbol{Z}_N^\dagger E_N,$$

where we need Asm. I to imply (19) holds using Lem. IV.1. Note that, by Cor. E.4 with $\mathcal{H}_i$ as in (26):

$$\mathrm{unvec_s}(\boldsymbol{Z}_N^\dagger E_N) = \textstyle\sum_{i=1}^N \mathcal{H}_i(w_i w_i^\top - W).$$

Note that $-r_w^2 I \preceq -W \preceq w_i w_i^\top - W \preceq w_i w_i^\top \preceq r_w^2 I$. Hence, by definition of $\|\mathcal{H}_i\|_2$, the spectral norm of $X_i := \mathcal{H}_i(w_i w_i^\top - W)$ is bounded as $\|X_i\|_2 \leq \|\mathcal{H}_i\|_2 r_w^2 := [\gamma]_i$. The matrices X_i are therefore bounded and i.i.d. (by Asm. II). Applying Lem. C.1 and solving for β shows

$$\mathbb{P}[\|\textstyle\sum_{i=1}^N \mathcal{H}_i(w_i w_i^\top - W)\|_2 \leq \beta] \geq 1 - \delta. \tag{52}$$

We have by Asm. II and Lem. E.1 that

$$(I - \boldsymbol{Z}_N^\dagger \boldsymbol{Z}_N)\mathrm{vec_s}(W) = (I - (\boldsymbol{\mathcal{W}}_{(3)}^\top)^\dagger \boldsymbol{\mathcal{W}}_{(3)}^\top)\mathrm{vec_s}(W).$$

Applying Lem. E.2 thus implies:

$$\mathcal{E}(\hat{W} - W; Z) = \mathcal{E}(\mathrm{unvec_s}(\boldsymbol{Z}_N^\dagger E_N); Z).$$

The bound $\|\mathrm{unvec_s}(\boldsymbol{Z}_N^\dagger E_N)\|_2 = \|\sum_{i=1}^N \mathcal{H}_i(w_i w_i^\top - W)\|_2$ in (52) therefore proves the result. □

We can also derive the sample complexity for the simplified setting where $\|z\|_2 \leq r_z$.

Proof of Lem. V.5: To deal with non-invertibility issues, we introduce $U \in \mathbb{R}^{n_w \times d_W}$, a unitary matrix whose columns span $\operatorname{im}(\boldsymbol{\mathcal{W}}_{(3)})$. For all $i \in \mathbb{N}_{1:N}$, we consider the i.i.d. random matrices $Y_i = U^\top \mathcal{W}(z_i z_i^\top) U$ and $Y = \sum_{i=1}^N Y_i$.

By Asm. II and Lem. E.1: $\operatorname{im}(\boldsymbol{Z}_N) = \operatorname{im}(\boldsymbol{\mathcal{W}}_{(3)})$. So $\lambda_{d_W}(Y) = \lambda_{d_W}(U^\top \boldsymbol{Z}_N^\top \boldsymbol{Z}_N U) = \lambda_{d_W}(\boldsymbol{Z}_N^\top \boldsymbol{Z}_N)$. Moreover $Y \succ 0$. Similarly $\mathbb{E}[\mathcal{W}(zz^\top)] = \boldsymbol{\mathcal{W}}_{(3)}(\mathbb{E}[(z \otimes_{\mathrm{s}} z)(z \otimes_{\mathrm{s}} z)^\top] \otimes I)\boldsymbol{\mathcal{W}}_{(3)}^\top$, with $\mathbb{E}[(z \otimes_{\mathrm{s}} z)(z \otimes_{\mathrm{s}} z)^\top] \succ 0$ by Asm. II and Cor. D.2. So $\lambda_{d_W}(\mathbb{E}[Y_i]) = \lambda_{d_W}(\mathbb{E}[\mathcal{W}(zz^\top)])$.

Moreover, since U is unitary, $\|Y_i\|_2 \leq r_z^4 \|\mathcal{W}\|_2 \leq r_z^4 \|\boldsymbol{\mathcal{W}}\|_2^2$ where we use Lem. E.3*(iii)* for the second inequality.

Therefore the matrices Y_i are i.i.d. and bounded as $\|Y_i\|_2 \leq r_z^4 \|\boldsymbol{\mathcal{W}}\|_2^2$ with $\lambda_{d_W}(\mathbb{E}[Y_i]) = \lambda_{d_W}(\mathbb{E}[\mathcal{W}(zz^\top)]) = \gamma_W r_z^4$. Therefore, (C.1) can be applied to show

$$\mathbb{P}[\lambda_{d_W}(Y - \mathbb{E}[Y]) \geq -\beta] \geq 1 - \delta,$$

with $\beta = r_z^2 \|\boldsymbol{\mathcal{W}}\|_2^2 \sqrt{N} \tau_W$ and $\tau_W = \sqrt{2\log(2d_W/\delta)}$. Hence,

$$\lambda_{d_W}(Y) \geq \lambda_{d_W}(\mathbb{E}[Y]) - \beta = N\gamma_W r_z^4 - \beta.$$

So, since $\|Y^\dagger\|_2 = \lambda_1(Y^\dagger) = (\lambda_{d_W}(Y))^{-1}$,

$$\|Y^\dagger\|_2 \leq \frac{r_z^{-4}}{N\gamma_W - \sqrt{N}\|\boldsymbol{\mathcal{W}}\|_2^2 \tau_W}.$$

Using Lem. III.3 we have

$$\zeta_W^2 = \textstyle\sum_{i=1}^N \|\mathcal{H}_i\|_2^2 \leq N \max_i \|\mathcal{H}_i\|_2^2,$$

where

$$\begin{aligned}\|\mathcal{H}_i\|_2^2 &\leq n_w \|Y^\dagger Y_i\|_2^2 \\ &\leq n_w \|Y^\dagger\|_2^2 \|Y_i\|_2^2 \leq n_w (\|Y^\dagger\|_2 r_z^4 \|\boldsymbol{\mathcal{W}}\|_2^2)^2.\end{aligned}$$

Therefore $\|\mathcal{H}_i\|_2^2 \leq (\|\boldsymbol{\mathcal{W}}\|_2^4/(N\gamma_W - \sqrt{N}\|\boldsymbol{\mathcal{W}}\|_2^2\tau_W))^2$. Multiplying by N then gives us the claimed bound for ζ_W^2. □

We continue with the proofs of §VII-A.

Proof of Lem. VII.2: By the classical LS error equation:

$$\mu_\star - \hat{\mu} = (I - (\boldsymbol{Z}_N^\mu)^\dagger \boldsymbol{Z}_N^\mu)\mu_\star + (\boldsymbol{Z}_N^\mu)^\dagger E_N^\mu, \tag{53}$$

with $E_N^\mu = [[\tilde{\boldsymbol{\mathcal{M}}}; z_1, \epsilon_1]], \ldots, [[\tilde{\boldsymbol{\mathcal{M}}}; z_N, \epsilon_N]]]$ and ϵ_i i.i.d. zero-mean random vectors. Equation (53) holds by Asm. III and Lem. VII.1. We can generalize Lem. E.1, which requires Asm. II, to $\boldsymbol{Z}_N^\mu$ to show $(\boldsymbol{Z}_N^\mu)^\dagger \boldsymbol{Z}_N^\mu = (\tilde{\boldsymbol{\mathcal{M}}}_{(3)}^\top)^\dagger \tilde{\boldsymbol{\mathcal{M}}}_{(3)}^\top$, which equals the identity by Asm. I as $\tilde{\boldsymbol{\mathcal{M}}}_{(3)}$ is full rank. Expanding the second term of (53), inserting the definition of G_i and $\mathcal{M}$ gives (33):

$$\mu_\star - \hat{\mu} = \textstyle\sum_{i=1}^N G_i \epsilon_i = \sum_{i=1}^N \mathbb{E}[G_i \tilde{w}_i] - G_i \tilde{w}_i.$$

Note that $G_i \tilde{w}_i$ is a sequence of i.i.d. random vectors with $\|G_i \tilde{w}_i\|_2 \leq \|G_i\|_2 r_w$, by the definition of the spectral norm. Hence Lem. C.2 is applicable. The final result is then recovered by solving for β. □

Proof of Thm. VII.4: Again, let $\mu_\star$ and $\tilde{W}_\star$ denote the true value of $\mathbb{E}[\tilde{w}]$ and $\mathbb{E}[\tilde{w}\tilde{w}^\top]$ respectively and $\hat{\mu}$ and $\hat{\tilde{W}}$ their estimates. Defining $\hat{\Sigma} := \hat{\tilde{W}} - \hat{\mu}\hat{\mu}^\top$ and $\Sigma_\star := \tilde{W}_\star - \mu_\star \mu_\star^\top$, allows us to expand $\mathcal{E}_\star$ as:

$$\begin{aligned}\mathcal{E}_\star(Z) &= \mathcal{E}\left(\begin{bmatrix} 1 & \mu_\star^\top \\ \mu_\star & \Sigma_\star + \mu_\star \mu_\star^\top \end{bmatrix}, Z\right) \\ &= \tilde{\mathcal{E}}(\Sigma_\star + \mu_\star \mu_\star^\top, Z) + \mathcal{E}\left(\begin{bmatrix} 1 & \mu_\star \\ \mu_\star^\top & 0 \end{bmatrix}, Z\right),\end{aligned}$$

by linearity and where $\tilde{\mathcal{E}}(\tilde{W}; Z) := \tilde{\boldsymbol{\mathcal{M}}}_{(3)}(\tilde{W} \otimes Z)\tilde{\boldsymbol{\mathcal{M}}}_{(3)}^\top$. Clearly $\tilde{\mathcal{E}}(\Sigma_\star + \mu_\star\mu_\star^\top; Z) = \tilde{\mathcal{E}}(\tilde{W}_\star; Z)$.

We reconsider (25):

$$\operatorname{vec}_{\mathrm{s}}(\tilde{W}_\star - \hat{\tilde{W}}) = (I - \boldsymbol{Z}_N^\dagger \boldsymbol{Z}_N)\operatorname{vec}_{\mathrm{s}}(\tilde{W}_\star) + \boldsymbol{Z}_N^\dagger E_N,$$

and label the first term ΔW_b and the second ΔW_η. Then

$$\tilde{\mathcal{E}}(\tilde{W}_\star; Z) = \tilde{\mathcal{E}}(\hat{\tilde{W}}; Z) + \tilde{\mathcal{E}}(\Delta W_\eta; Z) + \tilde{\mathcal{E}}(\Delta W_b; Z).$$

By Lem. E.2 the final term is zero. So without loss of generality we assume $\Delta W_b = 0$.

So it is sufficient to show that

$$\|\Sigma_\star - \hat{\Sigma}\|_2 = \|\Delta W_\eta - (\mu_\star\mu_\star^\top - \hat{\mu}\hat{\mu}^\top)\|_2 \leq \beta_\Sigma \tag{54}$$

and $\|\mu_\star - \hat{\mu}\|_2 \leq \beta_\mu$ with probability at least $1 - \delta_\mu - \delta_W$. The second bound follows directly by Lem. VII.2 (which claims (i) $\mathbb{P}[\|\mu_\star - \hat{\mu}\|_2 \leq \beta_\mu] \geq 1 - \delta_\mu$), while we have (ii) $\mathbb{P}[\|\Delta W_\eta\|_2 \leq \beta_W] \geq 1 - \delta_W$ by (52) in the proof of Thm. V.4. A union bound shows that both (i) and (ii) hold w.p. at least $1 - \delta_\mu - \delta_W$.

We have, by a triangle inequality and (54) that

$$\begin{aligned}\|\Sigma_\star - \hat{\Sigma}\|_2 &\leq \|\Delta W_\eta\|_2 + \|\mu_\star\mu_\star^\top - \hat{\mu}\hat{\mu}^\top\|_2 \\ &\leq \delta_W + \|\mu_\star\mu_\star^\top - \hat{\mu}\hat{\mu}^\top\|_2.\end{aligned}$$

We can rewrite $\mu_\star\mu_\star^\top - \hat{\mu}\hat{\mu}^\top$ as

$$(\mu_\star - \hat{\mu})(\mu_\star - \hat{\mu})^\top + (\mu_\star - \hat{\mu})\hat{\mu}^\top + \hat{\mu}(\mu_\star - \hat{\mu})^\top.$$

So, by another triangle inequality

$$\|\Sigma_\star - \hat{\Sigma}\|_2 \leq \beta_W + \beta_\mu(\beta_\mu + 2\|\hat{\mu}\|_2) = \beta_\Sigma.$$

We have therefore shown the claimed result. □

Appendix F
Semidefinite Programming

We will frequently use the following technical lemma:

Lemma F.1. *For some $P \in \mathbb{S}^d$ then*

$$\operatorname{Tr}[PX] > 0,\ \forall X \gneq 0 \quad \Leftrightarrow \quad P \succ 0.$$

Proof. ($\Rightarrow$) Assume P is not positive definite, which implies $\exists x\colon x^\top P x \leq 0$. So $X = xx^\top \gneq 0$ gives a contradiction.

($\Leftarrow$) Any $X \gneq 0$ can be written as $\sum_{i=1}^r x_i x_i^\top$ for some $r > 0$ and nonzero x_i. Then by the cyclic property of the trace $\operatorname{Tr}[PX] = \sum_{i=1}^r x_i^\top P x_i > 0$ by $P \succ 0$. □

We extend the result of [49, Thm. 6.2.1, 6.2.2]. This result was previously integrated in [13, Thm. 10].

Lemma F.2. *Consider the parametric family of matrices,*

$$\mathcal{U}[\rho] = \left\{ T + \rho \textstyle\sum_{i=1}^{d} [\theta]_i (L^\top F_i R + R^\top F_i^\top L) \colon \theta \in \mathcal{B}_d, \right\}$$

with $T \in \mathbb{S}^\ell$, $F_i \in \mathbb{R}^{m \times n}$. *Consider the predicate*

$$\exists \Lambda \in \mathbb{S}_+^n, \Gamma \in \mathbb{S}_+^m \colon T \succeq L^\top \Gamma L + R^\top \Lambda R, \tag{55a}$$

$$\begin{bmatrix} \Lambda & \rho F_1^\top & \dots & \rho F_d^\top \\ \rho F_1 & \Gamma & & \\ \vdots & & \ddots & \\ \rho F_d & & & \Gamma \end{bmatrix} \succeq 0. \tag{55b}$$

Then we have the following implication:

$$\textit{if } (55) \textit{ holds then } \mathcal{U}[\rho] \succeq 0;$$

Proof. We follow the proof of [49, Thm. 6.2.1, Thm. 6.2.2]. Assume (55) holds for some $\Lambda, \Gamma \succeq 0$. We assume $\Lambda = \Lambda^{1/2}\Lambda^{1/2}$ and $\Gamma = \Gamma^{1/2}\Gamma^{1/2}$ are invertible without loss of generality. Otherwise, use a Cholesky decomposition to construct a right invertible substitute for $\Lambda^{1/2}$ and $\Gamma^{1/2}$ and use these instead.

For $x \in \mathbb{R}^\ell$, let $u = \sqrt{2}\Lambda^{1/2} R x$ and $v = \sqrt{2}\Gamma^{1/2} L x$. Then,

$$2x^\top (\textstyle\sum_{i=1}^{d} [\theta]_i L^\top F_i R) x = v^\top (\sum_{i=1}^{d} [\theta]_i \Gamma^{-1/2} F_i \Lambda^{-1/2}) u.$$

Therefore $\mathcal{U}[\rho] \succeq 0$ if

$$x^\top A x + \rho v^\top (\textstyle\sum_{i=1}^{d} [\theta]_i \Gamma^{-1/2} F_i \Lambda^{-1/2}) u \geq 0, \tag{56}$$

for all $x \in \mathbb{R}^\ell$, $u \in \mathbb{R}^n$, $v \in \mathbb{R}^m$ and $\theta \in \mathcal{B}_d$.

We lower bound the second term of (56) as

$$\begin{aligned} &\rho v^\top (\textstyle\sum_{i=1}^{d} [\theta]_i \Gamma^{-1/2} F_i \Lambda^{-1/2}) u \\ &\quad \overset{(a)}{\geq} -\rho \|v\|_2 (\textstyle\sum_{i=1}^{d} |[\theta]_i| \|\Gamma^{-1/2} F_i \Lambda^{-1/2} u\|_2) \\ &\quad \overset{(b)}{\geq} -\rho \|v\|_2 \sqrt{\textstyle\sum_{i=1}^{d} u^\top \Lambda^{-1/2} F_i^\top \Gamma^{-1} F_i \Lambda^{-1/2} u}. \end{aligned}$$

Here (a) used the triangle inequality and (b) used $|\theta|^\top \beta \geq -\||\theta|\|_2 \|\beta\|_2 = -\|\beta\|_2$ for $[\beta]_i = \|\Gamma^{-1/2} F_i \Lambda^{-1/2} u\|_2$.

Note that (55b) implies $\rho^2 \sum_{i=1}^{d} \Lambda^{-1/2} H_i^\top \Gamma^{-1} F_i \Lambda^{-1/2} \preceq I$, by invertibility of Λ, Γ and a Schur complement. Combining this with (56) gives that $\mathcal{U}[\rho] \succeq 0$ holds if

$$x^\top T x \geq \|v\|_2 \|u\|_2,$$

for all $x \in \mathbb{R}^\ell$, $u \in \mathbb{R}^n$, $v \in \mathbb{R}^m$. By the inequality of arithmetic and geometric means (i.e., $(\alpha+\beta)/2 \geq \sqrt{\alpha\beta}$) we have the final sufficient condition for $\mathcal{U}[\rho] \succeq 0$:

$$x^\top T x \geq (\|v\|_2^2 + \|u\|_2^2)/2 = x^\top (L^\top \Gamma L + R^\top \Lambda R) x,$$

which holds $\forall x \in \mathbb{R}^\ell$ by (55a).

□