

Learning Linear Dynamical Systems with Semi-Parametric Least Squares

Max Simchowitz

University of California - Berkeley

MSIMCHOW@BERKELEY.EDU

Ross Boczar

University of California - Berkeley

BOCZAR@BERKELEY.EDU

Benjamin Recht

University of California - Berkeley

BRECHT@BERKELEY.EDU

Editors: Alina Beygelzimer and Daniel Hsu

Abstract

We analyze a simple prefiltered variation of the least squares estimator for the problem of estimation with biased, *semi-parametric* noise, an error model studied more broadly in causal statistics and active learning. We prove an oracle inequality which demonstrates that this procedure provably mitigates the variance introduced by long-term dependencies. We then demonstrate that prefiltered least squares yields, to our knowledge, the first algorithm that provably estimates the parameters of partially-observed linear systems that attains rates which do not incur a worst-case dependence on the rate at which these dependencies decay. The algorithm is provably consistent even for systems which satisfy the weaker *marginal stability* condition obeyed by many classical models based on Newtonian mechanics. In this context, our semi-parametric framework yields guarantees for both stochastic and worst-case noise.

Keywords: Linear Dynamical Systems, Statistical Learning Theory, Variance Reduction, Time Series, Semi-Parametric Statistics.

1. Introduction

Serial data are ubiquitous in machine learning, control theory, reinforcement learning, and the physical and social sciences. A major challenge is that such data exhibit long-term correlations, which often obscure the effects of design variables on measured observations and drive up the variance of statistical estimators.

In the study of linear, time-invariant dynamical (LTI) systems, for example, a vast literature of both classical and contemporary work typically assumes that the system exhibits a property called *strict stability*, which ensures that long term correlations decay geometrically (Verhaegen, 1993). While recent works show this condition can be removed in the special case when the full system state is perfectly observed (Simchowitz et al., 2018; Sarkar and Rakhlin, 2018; Faradonbeh et al., 2018), it is not known whether the condition is necessary in general. Moreover, it is not fully understood whether the rate of decay of correlations is the right quantity to parametrize the difficulty of learning, even for strictly stable systems.

Among the many challenges introduced by both non-strictly stable and almost-unstable LTI systems is that the more one probes, the more the long-term correlations compound to yield measurements with large magnitudes, thereby driving up the variance of statistical estimators. This

problem of growing variance arises in many other problem domains as well: for example, the reinforcement learning community has produced a great body of work dedicated to reducing variance in the presence of long time horizons and large reward baselines (Sutton and Barto, 1998; Greensmith et al., 2004).

This work intervenes by making two contributions. First, we analyze a simple prefiltered variation of the least squares estimator (*PF-LS*) for the problem of estimation with biased, *semi-parametric* noise, an error model studied more broadly in causal statistics and active learning (Chernozhukov et al., 2017; Krishnamurthy et al., 2018). We prove an oracle inequality which demonstrates that this procedure provably mitigates the variance introduced by long-term dependencies. Second, we demonstrate that prefiltered least squares yields, to our knowledge, the first algorithm that provably estimates the parameters of partially-observed linear systems that attains rates which do not incur a worst-case dependence on the rate at which these dependencies decay. The algorithm is provably consistent even for systems which satisfy the weaker *marginal stability* condition obeyed by many classical models based on Newtonian mechanics. In this context, our semi-parametric framework yields guarantees for both stochastic and worst-case noise.

1.1. Problem Statement

We consider the problem of regressing a sequence of observations $(\mathbf{y}_t) \subset \mathbb{R}^m$ to a sequence of inputs $(\mathbf{u}_t) \subset \mathbb{R}^p$ for times $t \in [N]$. For a fixed $T \in \mathbb{N}$, we define the concatenated input vectors $\bar{\mathbf{u}}_t = [\mathbf{u}_t | \dots | \mathbf{u}_{t-T+1}] \in \mathbb{R}^{Tp}$, and assume an *serial, semi-parametric* relationship between \mathbf{y}_t and $\bar{\mathbf{u}}_t$; that is, there exists a filtration $\{\mathcal{F}_t\}$ and a $G_\star \in \mathbb{R}^{m \times Tp}$ for which

$$\mathbf{y}_t = G_\star \bar{\mathbf{u}}_t + \boldsymbol{\delta}_t, \quad \mathbf{u}_t \in \mathcal{F}_t, \quad \boldsymbol{\delta}_t \in \mathcal{F}_{t-T}. \quad (1.1)$$

We choose inputs such that $\mathbf{u}_t | \mathcal{F}_{t-1} \sim \mathcal{N}(0, I_p)$ which ensures (i) the *Neyman-orthogonality* condition $\mathbb{E}[\boldsymbol{\delta}_t \bar{\mathbf{u}}_t^\top] = 0$, indispensable for consistent semi-parametric estimation (Chernozhukov et al., 2017) and (ii) that the inputs (\mathbf{u}_t) have well-conditioned covariance. This setting is motivated by the problem of estimating the parameters $(A_\star, B_\star, C_\star, D_\star)$ of a discrete-time linear system, which evolves according to the updates

$$\begin{aligned} \mathbf{x}_{t+1} &= A_\star \mathbf{x}_t + B_\star \mathbf{u}_t + B_w \mathbf{w}_t \\ \mathbf{y}_t &= C_\star \mathbf{x}_t + D_\star \mathbf{u}_t + D_z \mathbf{z}_t, \end{aligned} \quad (1.2)$$

where $(\mathbf{u}_t) \subset \mathbb{R}^p$ is the sequence of inputs, $(\mathbf{y}_t) \subset \mathbb{R}^m$ the sequence of outputs, $(\mathbf{x}_t) \subset \mathbb{R}^n$ is a sequence of states, $(\mathbf{w}_t) \subset \mathbb{R}^{d_w}$ and $(\mathbf{z}_t) \subset \mathbb{R}^{d_z}$ are sequences of process noise and sensor noise, respectively, and all above matrices are of appropriate dimension.¹ Crucially, we *do not* observe the system states \mathbf{x}_t or the noises \mathbf{z}_t and \mathbf{w}_t . We shall assume that the process and sensor noises can be chosen *semi-adversarially* in the sense that $\mathbf{w}_t, \mathbf{z}_t \in \mathcal{F}_{t-T}$ (i.e. an adversary may only act with a T -step delay). This model captures both stochastic and worst-case oblivious noise. A simple recursion shows that this condition implies that (1.1) holds in this setting, with G_\star equal to

$$G_\star := [D_\star, C_\star B_\star, C_\star A_\star B_\star, C_\star A_\star^2 B_\star, \dots, C_\star A_\star^{T-2} B_\star] \in \mathbb{R}^{m \times Tp},$$

the length- T response function from the inputs $\mathbf{u}_{t-T+1}, \dots, \mathbf{u}_t$ to the observation \mathbf{y}_t . Examining the dynamical equations (1.2), we see that the error $\boldsymbol{\delta}_t$ corresponds to the residual part of \mathbf{y}_t which does not depend on $\bar{\mathbf{u}}_t$, and is therefore $\{\mathcal{F}_{t-T}\}$ -adapted.

1. We do not estimate B_w and D_z , which are in general unidentifiable without a specific noise model.

An important recent result due to [Oymak and Ozay \(2018\)](#) demonstrates that for these LTI systems, a consistent estimate of G_\star can produce a consistent estimate $(\widehat{A}, \widehat{B}, \widehat{C}, \widehat{D})$ of an *equivalent realization* of $(A_\star, B_\star, C_\star, D_\star)$.² Furthermore, [Oymak and Ozay \(2018\)](#) show that if the operator norm $\|A_\star\|_{\text{op}}$ is strictly less than one, then ordinary least squares (OLS) can efficiently recover G_\star from the inputs $\mathbf{u}_1, \dots, \mathbf{u}_N$. Formally, if the process and sensor noises are i.i.d normally distributed, the least squares estimator which uses samples $\mathbf{u}_{N_1-T+1}, \dots, \mathbf{u}_N$ for some $N_1 \leq N/10$,

$$\widehat{G}_{\text{LS}} := \arg \min_{G \in \mathbb{R}^{m \times Tp}} \sum_{t=N_1}^N \|G\bar{\mathbf{u}}_t - \mathbf{y}_t\|_2^2, \quad (1.3)$$

converges to G_\star at a rate of $\mathcal{O}(N^{-1/2})$. Unfortunately, the condition $\|A_\star\|_{\text{op}} < 1$ is quite stringent, and the learning rates degrade as $\|A_\star\|_{\text{op}}$ approaches 1. Indeed, many systems of interest do not even satisfy a weaker condition known as *strict stability*: $\rho(A_\star) < 1$, where $\rho(\cdot)$ denotes the spectral radius. For example, simple oscillators, integrators, and elementary systems that arise from Newton's laws yield realizations where $\rho(A_\star) = 1$. For example, the LTI system corresponding to the discretization of the differential equation $F = m\ddot{x}$, with sampling time $\Delta > 0$, includes the matrix $A_\star = \exp\left(\Delta \cdot \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}\right) = \begin{bmatrix} 1 & \Delta \\ 0 & 1 \end{bmatrix}$. This matrix violates the strict stability condition, yet satisfies $\rho(A_\star) = 1$. As mentioned above, non-asymptotic bounds for learning LTI systems typically yield rates which depend on the *inverse stability gap* $1/(1 - \rho(A_\star))$ ([Oymak, 2018](#); [Hardt et al., 2016](#); [Shah et al., 2012](#)); for example, [Oymak \(2018\)](#) requires one to select a horizon length T for which $\|A_\star\|_{\text{op}}^T \leq .99$, which necessitates that $T \gtrsim \frac{1}{1 - \rho(A_\star)}$.

This work, on the other hand, suggests that a dependence on stability gap can be avoided in many cases, and the difficulty of learning can instead be parametrized by quantities that are often less conservative. Specifically, our results will extend to all *marginally* stable systems; that is, systems which satisfy $\rho(A_\star) \leq 1$, including possibly $\rho(A_\star) = 1$.

1.2. Prefiltered Least Squares (PF-LS)

In light of the limitations of ordinary least squares, we analyze PF-LS, a simple prefiltering step to improve the estimation of the matrix G_\star in the general semi-parametric setting (1.1). In Section 2, we specialize our analysis to establish consistent recovery of any linear dynamical system for which $\rho(A_\star) \leq 1$. Prefiltering mitigates the magnitude of the errors δ_t by learning a coarse linear filter of future outputs, denoted $\phi_{\text{rdg}} \in \mathbb{R}^{m \times \bar{L}}$, for $\bar{L} \in \mathbb{N}$. This filter uses a sequence $(\mathbf{k}_t) \in \mathbb{R}^{\bar{L}}$ encoding past observations to estimate (\mathbf{y}_t) via the prediction $\phi_{\text{rdg}} \cdot \mathbf{k}_t$; while we shall first describe our results for general features, we shall ultimately specialize our bounds to the features $\mathbf{k}_t := [\mathbf{y}_{t-T}^\top | \mathbf{y}_{t-2T}^\top | \dots | \mathbf{y}_{t-TL}^\top]^\top \in \mathbb{R}^{Lm}$, with $\bar{L} = Lm$ (see (2.7)).

We then estimate G_\star by regressing the filtered observations $(\mathbf{y}_t - \phi_{\text{rdg}} \cdot \mathbf{k}_t)$ to $(\bar{\mathbf{u}}_t)$. Concretely, our procedure is achieved with the following two steps of least squares.

2. That is, a pair $(\bar{A}, \bar{B}, \bar{C}, \bar{D})$ such that $\bar{D} = D_\star$, and for all $j \in \{0, 1, 2, \dots\}$, $C_\star A_\star^j B_\star = \bar{C} \bar{A}^j \bar{B}$.

$$\phi_{\text{rdg}} \leftarrow \arg \min_{\phi \in \mathbb{R}^{m \times \bar{L}}} \sum_{t=N_1}^N \|\mathbf{y}_t - \phi \cdot \mathbf{k}_t\|_2^2 + \mu^2 \|\phi\|_{\text{F}}^2 \quad (1.4)$$

$$\widehat{G}_{\text{PF}} \leftarrow \arg \min_{G \in \mathbb{R}^{m \times Tp}} \sum_{t=N_1}^N \|(\mathbf{y}_t - \phi_{\text{rdg}} \cdot \mathbf{k}_t) - G \bar{\mathbf{u}}_t\|_2^2 \quad (1.5)$$

Throughout, we let $\tilde{N} = N - N_1 + 1$, and we use the notation $\mathbf{\Delta} \in \mathbb{R}^{\tilde{N} \times m}$ to denote the matrix whose rows are $\boldsymbol{\delta}_{N_1}^\top, \dots, \boldsymbol{\delta}_N^\top$ and $\mathbf{K} \in \mathbb{R}^{\tilde{N} \times \bar{L}}$ the matrix whose rows are $\mathbf{k}_{N_1}^\top, \dots, \mathbf{k}_N^\top$. Our first contribution is the following inequality. Throughout, we use \lesssim to denote inequality up to universal multiplicative constants.

Theorem 1.1 (Prefiltering Oracle Inequality, Informal) *Consider the general semi-parametric setting described in Section 1.1, and suppose that $\mathbf{k}_t \in \mathcal{F}_{t-T}$. Then, with high probability,*

$$\|\widehat{G}_{\text{PF}} - G_\star\|_{\text{op}} \lesssim \frac{\text{Opt}_\mu + \text{Ovfit}_\mu}{N} \cdot \tilde{\mathcal{O}} \left(\sqrt{T(p + m + \bar{L})} \right).$$

Here, $\tilde{\mathcal{O}}(\cdot)$ hides logarithmic terms in N , $\|\mathbf{K}\|_{\text{op}}$, and $1/\mu$. The term

$\text{Ovfit}_\mu \sim \tilde{\mathcal{O}}(\sqrt{T(p + \bar{L})} \|G_\star\|_{\text{op}})$ captures the extent to which prefiltering overfits to $(\bar{\mathbf{u}}_t)$, and

$$\text{Opt}_\mu := \min_{\phi \in \mathbb{R}^{m \times \bar{L}}} \|\mathbf{\Delta} - \mathbf{K}\phi^\top\|_{\text{op}} + \mu \|\phi\|_{\text{op}} \quad (1.6)$$

describes the data-dependent prediction error of the best filter.

We defer a precise statement of Theorem 1.1 to Theorem 3.3. Note that this result makes no assumptions on the structure of the noise $\boldsymbol{\delta}_t$ or the features \mathbf{k}_t , other than the measurability assumptions that $\mathbf{k}_t, \mathbf{w}_t, \mathbf{z}_t \in \mathcal{F}_{t-T}$. Moreover, the term Opt_μ captures to the actual sequence of errors $\boldsymbol{\delta}_t$, rather than an a priori upper bound. For a sense of scaling, the overfitting term Ovfit_μ is typically dominated by Opt_μ , and by setting $\phi = 0$, $\text{Opt}_\mu \leq \|\mathbf{\Delta}\|_{\text{op}}$. When $\boldsymbol{\delta}_t = \mathcal{O}(1)$ on average, this terms behaves as $\sim \sqrt{\tilde{N}}$, and thus $\|\widehat{G}_{\text{PF}} - G_\star\|_{\text{op}}$ decays at a rate of $\tilde{\mathcal{O}}(N^{-1/2})$. In general, we only need to ensure $\text{Opt}_\mu \sim \sqrt{\tilde{N}}$.

1.3. Organization

Section 2 presents a precise statement of the PF-LS oracle inequality for LTI systems, Proposition 2.1, as well as bounds for the associated term Opt_μ in terms of the *phase rank* of A_\star (Definition 2.3). Consistency of estimation for marginally stable LTI systems is presented as a consequence (Corollary B.1). Section 3 walks the reader through the analysis of the PF-LS estimator, culminating in a formal statement of the oracle inequality, Theorem 3.3. Section 4 provides a proof sketch for results described in Section 2, and Section 5 addresses related work. Complete proofs are deferred to the Appendix, which is divided into three parts: Part I contains graphical illustrations of phase rank (Appendix A), the proof of Corollary B.1, and the lower bound for OLS, Theorem C.1. Part II pertains to the the oracle inequality and related material from Section 3. Part III addresses results specific to LTI systems. The appendix begins with a [preface](#) which consolidates notation and outlines the organization of the subsequent appendices in greater detail.

2. Rates for Learning LTI Systems

In the setting of marginally stable systems, we can not guarantee that $\|\Delta\|_{\text{op}}$ grows as \sqrt{N} due to the possible accumulation of system inputs. Indeed, Theorem C.1 in the appendix shows that the OLS estimator is inconsistent whenever $\rho(A_\star) \geq 1$ and the system satisfies a weak identifiability criterion. Even for strictly systems, it is not clear whether the inverse stability gap $(1 - \rho(A_\star))^{-1}$, correctly describes the difficulty of estimation. What we show is that by choosing a large enough filter length $L \in \mathbb{N}$ and features

$$\mathbf{k}_t := [\mathbf{y}_{t-T}^\top | \mathbf{y}_{t-2T}^\top | \cdots | \mathbf{y}_{t-TL}^\top]^\top \in \mathbb{R}^{Lm}, \quad (2.7)$$

we can ensure both that $\text{Opt}_\mu \sim N^{1/2}$ for marginally stable systems and that Opt_μ need not depend on the stability gap $1 - \rho(A_\star)$. Our choice of \mathbf{k}_t in (2.7) corresponds to filtering a subsampled history of the outputs in order to predict \mathbf{y}_t . This linear prefiltering step is in the spirit of many schemes detailed in the system identification and time-series literature (see Section 5), many of which ensure performance in both stochastic and adversarial settings for strictly stable systems. From the perspective of prefiltered semi-parametric learning, we observe that \bar{L} corresponds to $L \cdot m$, and that the dynamical equations (1.2) imply that features (2.7) are $\{\mathcal{F}_{t-T}\}$ -adapted.

We begin the task of deriving explicit estimation rates for LTI systems by first establishing an oracle inequality for the PF-LS scheme given by (1.4) and (1.5), for two particular noise models.

Definition 2.1 (Noise Models) *In the stochastic noise model, $\mathbf{w}_t | \mathcal{F}_{t-T}$ and $\mathbf{z}_{t-T} | \mathcal{F}_{t-T}$ are conditionally 1-subgaussian.³ In the adversarial noise model, the noise processes satisfy the bounds $\|\mathbf{w}_t\|_2^2 \leq d_w$ and $\|\mathbf{z}_t\|_2^2 \leq d_z$ with probability 1.*

Observe that shaped and/or scaled noise can be addressed by altering B_w and D_z appropriately. The conditions $\|\mathbf{w}_t\|_2^2 \leq d_w$ and $\|\mathbf{z}_t\|_2^2 \leq d_z$ make the adversarial and stochastic noise models comparable, as in the stochastic noise model we have $\mathbb{E}[\|\mathbf{w}_t\|_2^2] \lesssim d_w$ and $\mathbb{E}[\|\mathbf{z}_t\|_2^2] \lesssim d_z$ by subgaussianity. We shall assume $\mathbf{x}_1 = 0$ for the rest of the section, and we address general \mathbf{x}_1 in the Appendix. Now, we introduce the following parameter, which illustrates the dependence of our bounds on the eigenstructure of A_\star , the conditioning of the eigenvalues, and the magnitude of the noises encoded in B_w and D_z .

Definition 2.2 (Magnitude Bound) *Let $A_\star = S J_\star S^{-1}$ denote the Jordan-normal decomposition of A_\star . We define $\bar{M} := \|C_\star S\|_{\text{op}} (\|S^{-1} B_\star\|_{\text{op}} + \|S^{-1} B_w\|_{\text{op}}) + \|D_\star\|_{\text{op}} + \|D_z\|_{\text{op}}$.*

Lastly, our results apply once N satisfies a moderate lower bound. Specifically, we define $N_{\min} = \max\{10TL, cTp \log^4(2Tp)\}$, where c is a sufficiently large constant. Our result bounds $\|\hat{G}_{\text{PF}} - G_\star\|_{\text{op}}$ for LTI systems in terms of Opt_μ , \bar{M} , and dimension quantities. Furthermore, we let $\log_+(x) := \max\{1, \log(x)\}$. With these quantities defined, we state a specialized version of our general oracle inequality, Theorem 3.3.

Proposition 2.1 *Fix T and L , and suppose that $N \geq N_{\min}$, $N_1 = TL$, $\rho(A_\star) \leq 1$, and that the largest Jordan block of A_\star is of size k . Choosing some $\mu \geq 1$ and defining*

$$\bar{\mathbf{d}} := \tilde{p} + Lm (\log_+ \bar{M} + k \log_+ N) = \tilde{\mathcal{O}}(p + Lmk), \quad (2.8)$$

3. That is, $\mathbb{E}[\exp(\lambda(v, \mathbf{w}_t)) | \mathcal{F}_{t-T}] \leq \exp(\lambda^2 \|v\|_2^2)$ for all $v \in \mathbb{R}^{d_w}$, and analogously for \mathbf{z}_t

it holds with probability at least $1 - \delta - (2Np)^{-\log^2(2Tp) \log^2(2Np)}$ in the stochastic noise model that

$$\|\widehat{G}_{\text{PF}} - G_\star\|_{\text{op}} \lesssim N^{-1/2} \text{Opt}_\mu \cdot \sqrt{\frac{T(\bar{\mathbf{d}} + \log \frac{1}{\delta})}{N}} + \frac{T(\|G_\star\|_{\text{op}} + \mu)(\bar{\mathbf{d}} + \log \frac{1}{\delta})}{N}.$$

In the adversarial noise model, we instead take $\bar{\mathbf{d}} := \tilde{p} + Lm (\log_+(\bar{M} + d_z + d_w) + k \log_+ N)$. Moreover, we have $\|G_\star\|_{\text{op}} \lesssim \bar{M}T^{k-1/2}$.

We remark that the logarithmic terms in $\bar{\mathbf{d}}$ can be refined further, but we state the above bound for its relative simplicity. In the following subsection, we show that for any marginally stable system, one can ensure that $\text{Opt}_\mu \sim N^{1/2}$ as long as L is chosen to be sufficiently large. The parameter L can be chosen according to a structural minimization procedure (see Appendix H.4); more detailed discussion of the choices of T and L is deferred to Appendix B.1.

2.1. Learning without the stability gap

We shall now aim to answer the more question: how accurately does the inverse stability gap $(1 - \rho(A_\star))^{-1}$ describe the intrinsic difficulty of learning an LTI system? To this end, we describe an alternative criterion we call *phase rank*, which does not depend on the stability radius in the worst case. We use phase rank to bound the term Opt_μ , and thus, through Proposition 2.1, upper bound the learning rate of the estimator \widehat{G}_{PF} . As a second alternative to the inverse stability gap, in Appendix K we define a condition called *strong observability*, related to the classical notion of observability in control theory (Hautus, 1983). This condition also allows us to bound estimation rates in terms of quantities that do not degrade (in the worst case) as $\rho(A_\star) \rightarrow 1$.

Phase Rank: Many approaches in the recent learning theory literature have developed bounds which depend directly on the spectrum of A_\star and magnitudes of the state-space realization matrices $(A_\star, B_\star, C_\star, D_\star)$; however, many existing works have incurred dependencies on the minimal polynomial of A_\star , which is exponentially large in the worst case. Our first approach adopts a new measure of complexity we call *phase rank*, inspired by Hazan et al. (2018), derives bounds from the spectrum of A_\star without paying for the size of the minimal polynomial. Rather than capturing the effects of all eigenvalues, the phase rank groups together eigenvalues with approximately the same phase, and only considers the eigenvalues of A_\star which lie near the boundary of the unit disk. Formally, let $\mathbb{D} : \{z \in \mathbb{C} : |z| \leq 1\}$ denote the complex unit disk, and for a marginally stable A_\star , let $\text{blkspec}(A_\star) \subset \mathbb{C} \times \mathbb{N}$ denote the set⁴ of all pairs (λ, k) , where λ is an eigenvalue of A_\star and k is a size of an associated Jordan block. We define phase rank as follows:

Definition 2.3 (Phase Rank) *Let $\alpha \geq 1$. We say that A_\star has (α, T) -phase rank d if there exists $\mu_1, \dots, \mu_d \in \mathbb{D}$ such that, for any $(\lambda, k) \in \text{blkspec}(A_\star)$ with $|\lambda| \geq 1 - ((1 + \alpha)T)^{-1}$, there exists at least k elements $\mu_{i_1}, \dots, \mu_{i_k} \subset \{\mu_1, \dots, \mu_d\}$ satisfying*

$$\max_{j \in [k]} \min_{\tilde{\mu}: \tilde{\mu}^T = \mu_{i_j}^T} |\lambda - \tilde{\mu}| \leq \alpha(1 - |\lambda|).$$

4. As we will be taking the maximum over $\text{blkspec}(A_\star)$, set is equivalent to multiset for our purposes.

In the above definition, the parameter T allows us to group together eigenvalues having approximately the same phase mod $2\pi/T$, and the α parameter controls the ‘width’ of the approximation. The phase rank is typically small for many systems of interest; for example, for real diagonalizable systems, it is at most 2, *regardless of the stability radius*, thereby obviating a dependence on $1 - \rho(A_\star)$ in the worst case. Phase rank can also take advantage of benign systems which do exhibit stability; for example, phase rank is equal to 0 for strictly stable systems with $\rho(A_\star) < 1 - ((1 + \alpha)T)^{-1}$. In Appendix A, we give some visual diagrams to aid the intuition for this condition. Moreover, the $(1, T)$ phase-rank is at most the degree of the ‘‘minimal-phase polynomial’’ of A_\star , the measure of complexity studied by Hazan et al. (2018) which inspired this condition. With this definition in hand, we provide the following bound on Opt_μ .

Proposition 2.2 (Bounds for Phase Rank) *Suppose that A_\star has (α, T) phase rank d and maximum Jordan block size k . Then, for any $\delta \in (0, 1)$; with $L \geq d$, $N \geq T(d+1+\alpha) \max\{m, \log(1/\delta)\}$, and $N_1 \geq TL$; it holds with probability $1 - \delta$ under the stochastic noise setting of Definition 2.1 that, with $C_{\alpha,d,k} := 2^d(k^2(1 + \alpha)^{k-\frac{1}{2}} + d^{k-\frac{1}{2}})$,*

$$N^{-1/2}\text{Opt}_\mu \lesssim (\bar{M} + \mu N^{-1/2}) \cdot T^{k-1/2} C_{\alpha,d,k}.$$

We note that for strictly stable systems, one could instead establish bounds directly in terms of the stability radius, rather than the phase rank. These bounds would roughly scale as $N^{-1/2}\text{Opt}_\mu \lesssim k^2(1 - \rho(A_\star))^{k-\frac{1}{2}}$, and not depend directly on T . For systems where the eigendirections of A_\star can be ‘‘disentangled’’ when observed by C_\star , we show in Appendix I.1.4 that one may decouple the eigenmodes of A_\star to only have to consider the phase rank restricted to smaller portions of the spectrum of A_\star . For example, if there exists a well-conditioned matrix V for which $C_\star V$ and $V^{-1}A_\star$ are diagonal with blocks (C_1, C_2) , (A_1, A_2) , respectively, then we only incur a penalty for the maximum of the phase ranks of A_1 and A_2 . Appendix I also gives refinements of Proposition 2.2 that take more granular aspects of $\text{blkspec}(A_\star)$ into account. For constant phase rank d , we have the bound $N^{-1/2}\text{Opt}_\mu \sim \mathcal{O}(T^{k-1/2})$, which roughly matches the dependence on $\|F_\star\|_{\text{op}}$ in the bounds in Oymak and Ozay (2018). Lastly, as mentioned above, Theorem C.1 shows that OLS inconsistently estimates G_\star when $\rho(A_\star) = 1$, even with no process or sensor noise.

2.2. Recovering the System Parameters

Given our consistent estimates of G_\star , we can recover the system parameters $(A_\star, B_\star, C_\star, D_\star)$ at a rate of $\tilde{\mathcal{O}}(1/N^{1/4})$, even for marginally stable systems. We prove Corollary B.1, a precise corollary with all parameter dependencies made explicit, and discuss choices of parameters L and T in Appendix B.

3. Oracle Inequality for Prefiltered Least Squares

The goal of this section is to present Theorem 3.3, a technical version of the general oracle inequality Theorem 1.1 for the estimator \hat{G}_{PF} . To guide the proofs and intuition, we first consider the performance of the estimator which uses an arbitrary, fixed filter ϕ , rather than the data-dependent filter chosen by (1.4):

$$\hat{G}_{\text{fil}}(\phi) := \arg \min_{G \in \mathbb{R}^{m \times Tp}} \sum_{t=N_1}^N \|(\mathbf{y}_t - \phi \cdot \mathbf{k}_t) - G\bar{\mathbf{u}}_t\|_2. \quad (3.9)$$

Note that for $\phi = 0$, $\widehat{G}_{\text{fil}}(\phi) = \widehat{G}_{\text{LS}}$ is the OLS estimator. To analyze $\widehat{G}_{\text{fil}}(\phi)$, we must define the associated error $\delta_{\phi,t} := \delta_t - \phi \cdot \mathbf{k}_t$. Since δ_t, \mathbf{k}_t are \mathcal{F}_{t-T} -adapted, $\delta_{\phi,t}$ is as well. Thus, $\mathbf{y}_t - \phi \cdot \mathbf{k}_t = G_\star \bar{\mathbf{u}}_t + \delta_{\phi,t}$ is a semi-parametric model describing the relationship between $\mathbf{y}_t - \phi \cdot \mathbf{k}_t$ and $\bar{\mathbf{u}}_t$. We shall now establish two crucial properties which ensure estimation. The first shows that the matrix $\bar{\mathbf{U}} := [\bar{\mathbf{u}}_{N_1} | \dots | \bar{\mathbf{u}}_N]^\top$ is well-conditioned with high probability.

Lemma 3.1 (Lemma C.2 in Oymak and Ozay (2018)) *Define the event $\mathcal{E}_{\bar{\mathbf{U}}}$ and number $\delta_{\bar{\mathbf{U}}}$ as*

$$\mathcal{E}_{\bar{\mathbf{U}}} := \left\{ \frac{N}{2} I \preceq \bar{\mathbf{U}}^\top \bar{\mathbf{U}} \preceq 2I \right\} \quad \delta_{\bar{\mathbf{U}}} := (2Np)^{-\log^2(2Tp) \log^2(2Np)}.$$

Then, if $N_1 \leq \frac{1}{10}N$ and the sample size N satisfies $N \geq c'Tp \log^2(2Tp) \log^2(2Np)$ for a sufficiently large $c' > 0$, it holds that $\mathbb{P}[\mathcal{E}_{\bar{\mathbf{U}}}^c] \leq \delta_{\bar{\mathbf{U}}}$.

Inverting, we see that it suffices to take $N \geq N_{\min}$ for a possibly larger constant c to ensure the condition $N \geq c'Tp \log^2(2Tp) \log^2(2Np)$ holds. The second property we shall use is *Neyman orthogonality*, which states that

$$\mathbb{E}[\bar{\mathbf{u}}_t \delta_{\phi,t}^\top] = \mathbb{E}[\mathbb{E}[\bar{\mathbf{u}}_t | \mathcal{F}_{t-T}] \delta_{\phi,t}^\top] = 0.$$

This property is satisfied in our setting, since $\delta_{\phi,t}$ is \mathcal{F}_{t-T} adapted and $\mathbb{E}[\mathbf{u}_t | \mathcal{F}_{t-1}] = 0$. Chernozhukov et al. (2017) show that under general conditions, Neyman orthogonality generally ensures consistency of least squares, and Krishnamurthy et al. (2018) recently demonstrated that this idea implies consistency in a time-series setting. Our first result adapts this argument to handle the fact that our regression variables, $\bar{\mathbf{u}}_t$ have structure; namely, they are concatenated subsequences of (\mathbf{u}_t) . Denoting $\Delta_\phi \in \mathbb{R}^{\tilde{N} \times m}$ to be the matrix $[\delta_{\phi,N_1}^\top | \dots | \delta_{\phi,N}^\top]^\top$, we show the following bound.

Proposition 3.2 (Error Bound for Fixed Filter) *For any fixed filter $\phi \in \mathbb{R}^{m \times Lm}$, $\delta \in (0, 1)$, and $\kappa > 0$, it holds with probability at least $1 - \delta$ that on $\mathcal{E}_{\bar{\mathbf{U}}}$,*

$$\|\widehat{G}_{\text{fil}}(\phi) - G_\star\|_{\text{op}} \lesssim \frac{(\|N^{-1/2} \Delta_\phi\|_{\text{op}} + \kappa) T^{1/2}}{\sqrt{N}} \sqrt{\tilde{p} + m + \log \frac{1}{\delta} + \text{lil}\left(\frac{\|\Delta_\phi\|_{\text{op}}}{\kappa N^{1/2}}\right)}, \quad (3.10)$$

where $\text{lil}(x) := \log_+(\log_+(x))$ and $\tilde{p} := p \min\{T, \log^2(eTp) \log^2(Tp)\}$. In particular, the complement of (3.10) occurs with probability at most $\delta + \delta_{\bar{\mathbf{U}}}$.

For a sense of scaling, observe that whenever the sequence $(\delta_{t,\phi})$ is $\mathcal{O}(1)$ in magnitude on average, then $\frac{\|\Delta_\phi\|_{\text{op}}}{\sqrt{N}}$ is $\mathcal{O}(1)$ with high probability, yielding estimation rates of $\tilde{\mathcal{O}}\left(\sqrt{\frac{T(p+m)}{N}}\right)$.

Proof Sketch: Proposition 3.2 is derived as a special case of a more general result, Theorem E.1, which relies on self-normalized tail bounds for martingale sequences due to Abbasi-Adkori (2011). This theorem is similar in spirit to the tail bounds obtained by Krishnamurthy et al. (2018) for semi-parametric contextual bandits. The parameter $\kappa > 0$ arises from the use of these tools, but it can be chosen quite small due to the doubly-logarithmic dependence in $1/\kappa$. The novelty of our bound comes from a careful chaining argument in Appendix F.1 specific to semi-parametric regression with the concatenated sequence $(\bar{\mathbf{u}}_t)$, based on the techniques in Krahmer et al. (2014). This bound yields a dependence on \tilde{p} instead of the larger quantity Tp . In service of this argument, we give a recipe for applying Talagrand's chaining (Talagrand, 2014) to self-normalized martingale tail bounds in Appendix F, which may be of general interest.

3.1. Statement of the Oracle Inequality

In Proposition 3.2, we bounded the error for the least squares estimate associated with a fixed predictor, $\widehat{G}_{\text{fl}}(\phi)$, in terms of its associated error $\|\Delta_\phi\|_{\text{op}}$. Specifically, Proposition 3.2 implied that when $\|\Delta_\phi\|_{\text{op}}$ grows as $\widetilde{\mathcal{O}}(N^{1/2})$, $\|\widehat{G}_{\text{fl}}(\phi) - G_\star\|_{\text{op}}$ decays as $\widetilde{\mathcal{O}}(N^{-1/2})$.

In many cases, such as our setting of marginally stable systems, it is not possible to select a filter ϕ a priori in such a way that $\|\Delta_\phi\|_{\text{op}} \leq \widetilde{\mathcal{O}}(N^{1/2})$. Instead, in light of Proposition 3.2, one would like to choose the filter $\widehat{\phi}$ which minimizes $\|\Delta_\phi\|_{\text{op}}$, and pay for the magnitude of its associated error $\|\Delta_{\widehat{\phi}}\|_{\text{op}}$. Our main result of this section is an oracle inequality, proved in Appendix D, which shows that prefiltering the output sequence (\mathbf{y}_t) essentially accomplishes this goal.

Theorem 3.3 (PF-LS Oracle Inequality) *Let Opt_μ be as in (1.6), and define*

$$\begin{aligned} \text{Ovfit}_\mu(\delta) &:= \|G_\star\|_{\text{op}} \cdot \min \left\{ N^{1/2}, T^{1/2} \sqrt{\log \frac{1}{\delta} + \widetilde{p} + \log \det(I + \mu^{-2} \mathbf{K} \mathbf{K}^\top)^{1/2}} \right\} \\ \mathbf{d}_{\text{eff}}(\text{Opt}, \bar{L}, \mu) &:= \widetilde{p} + m + \text{lil} \frac{\text{Opt}}{\mu} + \bar{L} \log_+(\text{Opt} + \frac{\sqrt{N} \|\mathbf{K}\|_{\text{op}}}{\mu^2}). \end{aligned}$$

Then for any $\delta \in (0, 1)$, then following inequality holds probability with $1 - \delta - \delta_{\bar{U}}$, provided N satisfies the conditions of Lemma 3.1:

$$\|\widehat{G}_{\text{PF}} - G_\star\|_{\text{op}} \lesssim \frac{N^{-1/2}(\text{Opt}_\mu + \text{Ovfit}_\mu(\delta) + \mu)}{\sqrt{N}} \cdot \sqrt{T (\log \frac{1}{\delta} + \mathbf{d}_{\text{eff}}(\text{Opt}_\mu + \text{Ovfit}_\mu(\delta), \bar{L}, \mu))}.$$

Here, the term Opt_μ corresponds to the error achieved by the best filter ϕ , and the term \mathbf{d}_{eff} captures the ‘‘effective dimension’’ of the estimation problem, totaling the dimensions of the filter class, inputs (\mathbf{u}_t) , and observations (\mathbf{y}_t) . For the case of linear systems where $\mathbf{k}_t = [\mathbf{y}_{t-T}^\top \dots \mathbf{y}_{t-TL}^\top]^\top$ and $\bar{L} = Lm$, in Appendix H.4 we give an algorithm for selecting the parameter L which admits an oracle inequality, Proposition H.4. Moreover, the bound of Theorem 3.3 depends only logarithmically on $1/\mu$, so μ may be taken to be very small.

Proof Sketch: In proving Theorem 3.3, we first obtain an intermediate but analogous result in terms of the intermediate quantity $\|\Delta_{\phi_{\text{rdg}}}\|_{\text{op}} + \mu \|\phi_{\text{rdg}}\|_{\text{op}}$, which we bound in Appendix D.2 by $\text{Opt}_\mu + \text{Ovfit}_\mu$ using KKT arguments and a variant of Proposition 3.2. To prove the the intermediate result, Appendix D.1 considers ‘‘slices’’ $v^\top \phi$ along directions $v \in \mathcal{S}^{m-1}$, and establish uniform bounds with respect to a hierarchy of coverings of \mathbb{R}^{Lm} , each with a different scale and granularity, such that the bounds hold for each covering in the hierarchy simultaneously. This lets us tailor the granularity of the covering for each specific filter ϕ , which (a) yields bounds depending on the data-dependent errors $\|\Delta_\phi\|_{\text{op}}$, (b) ensures tighter control on filters ϕ with smaller norm, and (c) tolerates logarithmically more error as $\|\phi\|_{\text{op}}$ grows. Specializing the uniform bound to ϕ_{rdg} requires loose control over $\|\phi_{\text{rdg}}\|_{\text{op}}$ (hence the regularization in (1.4)), which adding the regularization term allows us to control. The choice of μ trades off between the magnitude of Opt_μ and the quantity $\frac{\|\mathbf{K}\|_{\text{op}}}{\kappa\mu}$ inside the logarithmic term. This is somewhat of an artifact of the proof and is unnecessary if $\sigma_{\min}(\mathbf{K})$ is bounded from below, for example. Again, since $1/\mu$ appears in a logarithmic term, we shall typically take μ to be quite small (e.g. $1/N$).

4. Proof Sketch for Bounding Opt_μ

Since $\text{Opt}_\mu = \min_\phi \|\Delta_\phi\|_{\text{op}} + \mu\|\phi\|_{\text{op}}$, it suffices to exhibit some ϕ with reasonable operator norm for which $\|\Delta_\phi\|_{\text{op}}$ grows as \sqrt{N} . To this end, we define the auxiliary signal $\tilde{\mathbf{x}}_{n;t}$ and associated observation $\tilde{\mathbf{y}}_{n;t}$ via

$$\tilde{\mathbf{y}}_{n;t} = C_\star \tilde{\mathbf{x}}_{n;t}, \quad \tilde{\mathbf{x}}_{n;t} := \begin{cases} A_\star^{n-(t-TL)} \mathbf{x}_{t-TL} & n \geq t - TL \\ \mathbf{x}_n & n \leq t - TL \end{cases}.$$

Here, $\tilde{\mathbf{x}}_t$ is the state as if the noise and inputs had been ‘‘shut off’’ at time $t - TL$. We further define the features $\tilde{\mathbf{k}}_t := [\tilde{\mathbf{y}}_{t-T;t}^\top \mid \tilde{\mathbf{y}}_{t-2T;t}^\top \mid \cdots \mid \tilde{\mathbf{y}}_{t-TL;t}^\top]^\top$ and decompose the error term as $\delta_{\phi,t} = \delta_t - \phi \cdot \mathbf{k}_t = \text{Err}_{\phi,t}^{(1)} + \text{Err}_{\phi,t}^{(2)}$, where

$$\text{Err}_{\phi,t}^{(1)} := \tilde{\mathbf{y}}_{t;t} - \phi \cdot \tilde{\mathbf{k}}_t \quad \text{and} \quad \text{Err}_{\phi,t}^{(2)} := (\delta_t - \tilde{\mathbf{y}}_{t;t}) - \phi \cdot (\mathbf{k}_t - \tilde{\mathbf{k}}_t).$$

Here, $\text{Err}_{\phi,t}^{(1)}$ describes the approximation error $\tilde{\mathbf{y}}_{t;t} - \phi \cdot \tilde{\mathbf{k}}_t$ incurred in predicting $\tilde{\mathbf{y}}_{t;t}$ from the shut-off sequence $\tilde{\mathbf{y}}_{t-T;t}, \tilde{\mathbf{y}}_{t-2T;t}, \dots, \tilde{\mathbf{y}}_{t-TL;t}$, and $\text{Err}_{\phi,t}^{(2)}$ accounts for the additional noise induced by the shut-off sequence. In Propositions 4.1 (resp. G.2), we prove bounds on the total contributions of these two errors under stochastic (resp. adversarial) noise models outlined in Assumption 2.1. For any fixed ϕ , the error terms $\text{Err}_{\phi,t}^{(2)}$ do not grow with time, since they only account for the contribution of noise over TL time steps; thus, the contribution of $\text{Err}_{\phi,t}^{(2)}$ to $\|\Delta_\phi\|_{\text{op}}$ grows as \sqrt{N} .

The terms $\text{Err}_{\phi,t}^{(1)}$, on the other hand, may grow with time because they depend on the state \mathbf{x}_{t-TL} , which can grow in magnitude for marginally stable systems under consistent excitation. Fortunately, by the Cayley-Hamilton theorem, we can observe that for large enough L , there always exists a ϕ for which $\text{Err}_{\phi,t}^{(1)} = 0$ for all t . Indeed, let $f(z) = z^d + f_1 z^{d-1} + \cdots + f_d$ denote the minimal polynomial of A_\star^T , and let $\phi_f = -[f_1 I_m \mid f_2 I_m \mid \cdots \mid f_d I_m \mid \mathbf{0}]$. Then, if $L \geq d$, a short computation shows that $\text{Err}_{\phi_f,t}^{(1)} = 0$. Unfortunately, $\phi = \phi_f$ requires L to be at least the degree of the minimal polynomial of A_\star^T , which can be as large as n in general. Moreover, the minimal polynomial f may have exponentially large coefficients, which can amplify the effect of noise in $\text{Err}_{\phi_f,t}^{(2)}$ and also affect the contribution of the regularization term $\mu\|\phi\|_{\text{op}}$. As introduced in Section 2.1, bounded phase rank ensures that there exists a smaller (both in length and in norm) filter ϕ than one would obtain by applying the minimal polynomial. Throughout, we shall consider the stochastic case; the adversarial case is similar and deferred to the Appendix.

Applying Phase Rank: Proposition 2.2 in the introduction gave an explicit bound on Opt_μ in terms of the largest Jordan block k of A_\star and the (α, T) phase rank of the system. The formal proof of this bound is deferred to Appendix I. Here, we shall instead provide an informal intuition about why phase rank is also a natural quantity. Consider the state transition matrix $A_\star = \begin{bmatrix} 1 & 0 \\ 0 & 1-\epsilon \end{bmatrix}$, and suppose $B_\star = B_w = I_2$. The first coordinate corresponds to a marginally unstable eigenvalue 1, and the second corresponds to an eigenvalue which is strictly stable by a small margin ϵ . Taking $L = 1$, we see the filter $\phi = I_m$ corresponding to the polynomial $g(z) = z - 1$ not only exactly cancels the first mode, it also *downweights* the second mode by a factor of $(1 - (1 - \epsilon)^T)$, as $\text{Err}_{\phi,t}^{(1)} = C_\star[2](1 - (1 - \epsilon)^T) \cdot \tilde{\mathbf{x}}_{t-T}[2]$. Moreover, we can express the second coordinate as $\tilde{\mathbf{x}}_t[2] = \sum_{s=1}^t (1 - \epsilon)^{t-s} (\mathbf{u}_s[2] + \mathbf{w}_s[2])$. Due to the geometric decay, this sum roughly depends on only the last $\mathcal{O}(1/\epsilon)$ terms in the sum. Therefore, even for adversarial noise, $|\tilde{\mathbf{x}}_t[2]|$ should be at most

$\mathcal{O}(1/\epsilon)$ on average. With this observation in hand, $\|\text{Err}_{\phi,t}^{(1)}\| \lesssim (1 - (1 - \epsilon)^T) \cdot \frac{1}{\epsilon} \lesssim T$ on average: this bound depends neither on the time step t nor the parameter ϵ .

Now, how does this connect to phase rank? We show in Appendix I that the salient feature of our choice of ϕ was that the corresponding polynomial $g(z)$ had a root with the same *phase* as the eigenvalues of A_* , which exactly offset the magnitude of the state along the corresponding eigendirections. Generalizing to systems with Jordan blocks and multiple phases, we prove that small phase rank lets us construct small-norm filters ϕ which yield small $\|\text{Err}_{\phi,t}^{(1)}\|$.

Explicit Bounds on $\|\Delta_\phi\|_{\text{op}}$: A central technical step in bounding Opt_μ is obtaining explicit upper bounds on $\|\Delta_\phi\|_{\text{op}}$. In what follows, we use bold sans-serif notation $\mathbf{G} = (A, B, C, D)$ to denote a dynamical system, where $\mathbf{G}_* := (A_*, B_*, C_*, D_*)$. With $\phi = [\Psi_1 \mid \dots \mid \Psi_L]$, we define $\|\phi\|_{\text{bop}} := \sum_{\ell=1}^L \|\Psi_\ell\|_{\text{op}}$. Lastly, we define the associated observation matrix $C_\phi := C_* A_*^{LT} - \sum_{\ell=1}^L \Psi_\ell C_* A_*^{(L-\ell)T} \in \mathbb{R}^{m \times n}$, which controls the size of the filtered output sequence $(\mathbf{y}_t - \phi \cdot \mathbf{k}_t)$, as well as the associated LTI systems $\mathbf{G}_\phi := (A_*, B_*, C_\phi, 0)$ and $\mathbf{F}_\phi := (A_*, B_w, C_\phi, 0)$. We remark that the parameter L in C_ϕ depends on the length of the filter ϕ ; e.g. for $\phi \in \mathbb{R}^{m \times dm}$, we replace L by d . To state our bound on $\|\Delta_\phi\|_{\text{op}}$, we shall also need to define, for arbitrary dynamical systems $\mathbf{G} = (A, B, C, D)$, the Markov parameter matrix $\mathcal{M}_k(\mathbf{G}) := [D \mid CB \mid \dots \mid CA^{k-1}B \mid CA^{k-2}B]$. For example, we see that $\mathcal{M}_T(\mathbf{G}) = G_*$. We shall also identify dynamical systems by their discrete-time transfer functions; that is, we associate $\mathbf{G} = (A, B, C, D)$ with the real rational transfer function $\mathbf{G}(z) = C(zI - A)^{-1}B + D$, mapping $\mathbb{C} \rightarrow \mathbb{R}^{m \times p}$. The notation $\mathcal{M}_k(\mathbf{G})$ and $\mathbf{G}(z)$ allows us to define the following control-theoretic norms:

Definition 4.1 (Control Norms) Consider a dynamical system $\mathbf{G} = (A, B, C, D)$ with $\rho(A) \leq 1$. We define the norms $\|\mathbf{G}\|_{\mathcal{H}_\infty} := \sup_{z \in \mathbb{C}: |z|=1} \|\mathbf{G}(z)\|_{\text{op}}$, and $\|\mathcal{M}_\infty(\mathbf{G})\|_{\text{op}} := \lim_{k \rightarrow \infty} \|\mathcal{M}_k(\mathbf{G})\|_{\text{op}}$. We allow these norms to take on the value ∞ .

The \mathcal{H}_∞ -norm admits a variational interpretation. It corresponds the induced $\ell_2^p \rightarrow \ell_2^m$ norm for LTI systems. We also remark that $\|\mathcal{M}_\infty(\mathbf{G})\|_{\text{op}}$ is equal to the square root of the largest eigenvalue of the so-called ‘‘infinite-horizon Gramian’’ and is an operator-norm representation of the \mathcal{H}_2 -norm in control theory; see Zhou et al. (1996, Chapter 4) for a discussion on both the \mathcal{H}_∞ and \mathcal{H}_2 system norms. Note that $\rho(A) < 1$ guarantees that both norms are finite. We are now ready to state our bound on the norm of the error $\|\Delta_\phi\|_{\text{op}}$. The adversarial case is similar and is given in Proposition G.2; both are proven in Appendix G.

Proposition 4.1 (Stochastic Noise Bound) Consider a filter of the form $\phi = [\Psi_1 \mid \dots \mid \Psi_d] \in \mathbb{R}^{m \times dm}$ for some $1 \leq d \leq L$, and suppose that $N \geq Td \max\{m, \log(1/\delta)\}$. Then, in the stochastic noise model of Assumption 2.1, the extended $\tilde{\phi} := [\phi \mid \mathbf{0}_{m \times (L-m)d}] \in \mathbb{R}^{L \times m}$ satisfies the following with probability $1 - \delta$:

$$\begin{aligned} \|\Delta_{\tilde{\phi}}\|_{\text{op}} &\lesssim \sqrt{N}(\|\mathcal{M}_\infty(\mathbf{G}_\phi)\|_{\text{op}} + \|\mathcal{M}_\infty(\mathbf{F}_\phi)\|_{\text{op}}) + \sqrt{m + \log(1/\delta)}(\Gamma_N(\mathbf{G}_\phi) + \Gamma_N(\mathbf{F}_\phi)) \\ &\quad + \sqrt{N}(1 + \|\phi\|_{\text{bop}})(\|\mathcal{M}_{Td}(\mathbf{G}_*)\|_{\text{op}} + \|\mathcal{M}_{Td}(\mathbf{F}_*)\|_{\text{op}} + \|D_z\|_{\text{op}}), \end{aligned}$$

where we define $\Gamma_N(\mathbf{G}) := \min\{\sqrt{N}\|\mathcal{M}_\infty(\mathbf{G})\|_{\text{op}}, \|\mathbf{G}\|_{\mathcal{H}_\infty}\}$.

The parameter d optimizes for filters that arise from the phase rank of A_* and allows for sharper bounds (replacing d by L) in that setting. Note that for this proposition to not be vacuous, we must choose ϕ such that the \mathcal{H}_∞ - and $\|\mathcal{M}_\infty(\cdot)\|_{\text{op}}$ norms of the systems $\mathbf{G}_\phi, \mathbf{F}_\phi$, described above are finite. This requires canceling out the effects of the modulus-1 eigenvalues of A_* . This is why the phase rank definition is at least as large as the number of such eigenvalues.

5. Related Work

Identifying LTI systems from data has a decades-old history in both the time-series and system identification communities (see [Ljung \(1999\)](#); [Verhaegen \(1993\)](#); [Galrinho \(2016\)](#) and references therein) with least squares estimation being a central tool for dozens of algorithms, many of them similar in spirit to PF-LS, (1.4),(1.5). One can regard PF-LS as a specific instance of a prefiltered autoregressive model (such as ARX or ARMAX); much work has been done on explicit filtering and debiasing schemes for these types of models [Spinelli et al. \(2005\)](#); [Ding \(2013\)](#); [Zheng \(2004\)](#); [Guo and Huang \(1989\)](#); [Zhang \(2011\)](#); [Wang \(2011\)](#); [Galrinho et al. \(2014\)](#). However, analyses of these schemes are often (i) asymptotic, (ii) for strictly stable systems only, or (iii) use a limited noise model. A complementary viewpoint comes from a family of techniques known broadly as *subspace identification* (e.g. [Qin \(2006\)](#)), which take a singular value decomposition (SVD) of the raw data; following [Oymak and Ozay \(2018\)](#) and the classical algorithm of [Kung \(1978\)](#), we instead use SVD as a post-processing step via the Ho-Kalman algorithm. It is an interesting direction for future work to explore of SVD-based algorithms can modified to enjoy guarantees for marginally stable systems as well.

There has also been considerable recent work from the machine learning community on non-asymptotic rates for prediction and estimation in LTI systems. While many have shown that strict stability is not necessary when the full system state can be observed ([Simchowitz et al., 2018](#); [Sarkar and Rakhlin, 2018](#); [Faradonbeh et al., 2018](#)), stability been central to other works providing guarantees for when only (y_t) are observed ([Shah et al., 2012](#); [Hardt et al., 2016](#); [Oymak, 2018](#)). Strict stability can be removed at the expense of requiring a number of independent trajectories which grows with desired accuracy [Oymak \(2018\)](#), or for online prediction problems in which A_* is diagonalizable and persistent process noise is minimal ([Hazan et al., 2017, 2018](#)). The regret bounds in [Hazan et al. \(2018\)](#) depended on the ℓ_1 -norm of the minimal phase polynomial, the inspiration for the phase rank condition in this work.

Beyond linear systems, our prefiltering step bears similarities to the *instrumental variables* technique in used in controls ([Viberg et al., 1997](#)), econometrics ([Hansen and Singleton, 1982](#)) and causal statistics ([Angrist et al., 1996](#)), which is used more for debiasing than for denoising. More broadly, variance reduction has become an indispensable component of reinforcement learning ([Weaver and Tao, 2001](#); [Greensmith et al., 2004](#); [Tucker et al., 2017](#); [Sutton and Barto, 1998](#)), including the theoretical study of tabular Markov Decision Processes ([Kakade et al., 2018](#); [Sidford et al., 2018](#)).

Acknowledgments

We thank Vaishaal Shankar for timely PyWren support. We thank Cyril Zhang and Holden Lee for their generous and thorough exposition of [Hazan et al. \(2018\)](#) through personal correspondence. We also thank Samet Oymak for providing the code accompanying [Oymak and Ozay \(2018\)](#). This work was generously supported in part by ONR awards N00014-17-1-2191, N00014-17-1-2401, and N00014-18-1-2833, the DARPA Assured Autonomy (FA8750-18-C-0101) and Lagrange (W911NF-16-1-0552) programs, and an Amazon AWS AI Research Award. MS is also generously supported by a Berkeley Fellowship, sponsored by the Rose Hill Foundation.

References

- Yasin Abbasi-Adkori. Online Least Squares Estimation with Self-Normalized Processes: An Application to Bandit Problems. 2011.
- Joshua D. Angrist, Guido W. Imbens, and Donald B. Rubin. Identification of causal effects using instrumental variables. *Journal of the American Statistical Association*, 91(434):444–455, 1996.
- Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, and Whitney Newey. Double/debiased/neyman machine learning of treatment effects. *American Economic Review*, 107(5):261–65, 2017.
- Feng Ding. Two-stage least squares based iterative estimation algorithm for CARARMA system modeling. *Applied Mathematical Modelling*, 37(7):4798–4808, April 2013. ISSN 0307904X. doi: 10.1016/j.apm.2012.10.014. URL <https://linkinghub.elsevier.com/retrieve/pii/S0307904X12006191>.
- Mohamad Kazem Shirani Faradonbeh, Ambuj Tewari, and George Michailidis. Finite time identification in unstable linear systems. *Automatica*, 96:342–353, 2018.
- Miguel Galrinho. Least squares methods for system identification of structured models. 2016. URL <http://www.diva-portal.org/smash/get/diva2:953835/FULLTEXT01.pdf>.
- Miguel Galrinho, Cristian Rojas, and Håkan Hjalmarsson. A weighted least-squares method for parameter estimation in structured models. In *53rd IEEE Conference on Decision and Control*, December 2014. doi: 10.1109/CDC.2014.7039903.
- Evan Greensmith, Peter L. Bartlett, and Jonathan Baxter. Variance reduction techniques for gradient estimates in reinforcement learning. *Journal of Machine Learning Research*, 5(Nov):1471–1530, 2004.
- Lei Guo and Dawei Huang. Least-squares identification for ARMAX models without the positive real condition. *IEEE Transactions on Automatic Control*, 34(10):1094–1098, October 1989. ISSN 0018-9286. doi: 10.1109/9.35285.
- Lars Peter Hansen and Kenneth J. Singleton. Generalized instrumental variables estimation of nonlinear rational expectations models. *Econometrica: Journal of the Econometric Society*, pages 1269–1286, 1982.
- Moritz Hardt, Tengyu Ma, and Benjamin Recht. Gradient Descent Learns Linear Dynamical Systems. *arXiv:1609.05191*, 2016.
- M.L.J. Hautus. Strong detectability and observers. *Linear Algebra and its applications*, 50:353–368, 1983.
- Elad Hazan, Karan Singh, and Cyril Zhang. Learning Linear Dynamical Systems via Spectral Filtering. In *Neural Information Processing Systems*, 2017.

- Elad Hazan, Holden Lee, Karan Singh, Cyril Zhang, and Yi Zhang. Spectral Filtering for General Linear Dynamical Systems. 2018.
- B. L. Ho and R. E. Kalman. Effective construction of linear state-variable models from input/output functions. *Automatisierungs-Technik*, 14(1-12):545–548, 1966.
- Sham Kakade, Mengdi Wang, and Lin F Yang. Variance reduction methods for sublinear reinforcement learning. *arXiv preprint arXiv:1802.09184*, 2018.
- Felix Kraemer, Shahar Mendelson, and Holger Rauhut. Suprema of chaos processes and the restricted isometry property. *Communications on Pure and Applied Mathematics*, 67(11):1877–1904, 2014.
- Akshay Krishnamurthy, Zhiwei Steven Wu, and Vasilis Syrgkanis. Semiparametric contextual bandits. *arXiv preprint arXiv:1803.04204*, 2018.
- Sun-Yuan Kung. A new identification and model reduction algorithm via singular value decomposition. In *12th Asilomar Conference on Circuits, Systems and Computers, Pacific Grove, CA, November, 1978*, 1978.
- Beatrice Laurent and Pascal Massart. Adaptive estimation of a quadratic functional by model selection. *Annals of Statistics*, pages 1302–1338, 2000.
- Lennart Ljung. *System Identification: Theory for the User*. 1999.
- Samet Oymak. Stochastic gradient descent learns state equations with nonlinear activations. *arXiv preprint arXiv:1809.03019*, 2018.
- Samet Oymak and Necimye Ozay. Non-asymptotic Identification of LTI Systems from a Single Trajectory. 2018.
- S. Joe Qin. An overview of subspace identification. *Computers & Chemical Engineering*, 30(10-12): 1502–1513, September 2006. ISSN 00981354. doi: 10.1016/j.compchemeng.2006.05.045. URL <https://linkinghub.elsevier.com/retrieve/pii/S009813540600158X>.
- Tuhin Sarkar and Alexander Rakhlin. How fast can linear dynamical systems be learned? *arXiv preprint arXiv:1812.01251*, 2018.
- Tuhin Sarkar, Alexander Rakhlin, and Munther A Dahleh. Finite-time system identification for partially observed lti systems of unknown order. *arXiv preprint arXiv:1902.01848*, 2019.
- Parikshit Shah, Badri Narayan Bhaskar, Gongguo Tang, and Benjamin Recht. Linear System Identification via Atomic Norm Regularization. In *Conference on Decision and Control*, 2012.
- John Shawe-Taylor, Peter L. Bartlett, Robert C. Williamson, and Martin Anthony. Structural risk minimization over data-dependent hierarchies. *IEEE transactions on Information Theory*, 44(5): 1926–1940, 1998.

- Aaron Sidford, Mengdi Wang, Xian Wu, and Yinyu Ye. Variance reduced value iteration and faster algorithms for solving markov decision processes. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 770–787. Society for Industrial and Applied Mathematics, 2018.
- Max Simchowitz, Horia Mania, Stephen Tu, Michael I. Jordan, and Benjamin Recht. Learning without mixing: Towards a sharp analysis of linear system identification. In Sébastien Bubeck, Vianney Perchet, and Philippe Rigollet, editors, *Proceedings of the 31st Conference On Learning Theory*, volume 75 of *Proceedings of Machine Learning Research*, pages 439–473. PMLR, 06–09 Jul 2018. URL <http://proceedings.mlr.press/v75/simchowitz18a.html>.
- William Spinelli, Luigi Piroddi, and Marco Lovera. On the role of prefiltering in nonlinear system identification. *IEEE Transactions on Automatic Control*, 50(10):1597–1602, October 2005. ISSN 0018-9286. doi: 10.1109/TAC.2005.856655. URL <http://ieeexplore.ieee.org/document/1516260/>.
- Elias M. Stein and Rami Shakarchi. Princeton lecture in analysis ii. complex analysis, 2003.
- Richard S. Sutton and Andrew G. Barto. Reinforcement Learning. 1998.
- Michel Talagrand. *Upper and Lower Bounds for Stochastic Processes: Modern Methods and Classical Problems*, volume 60. Springer Science & Business Media, 2014.
- Paolo Tilli. Singular values and eigenvalues of non-hermitian block toeplitz matrices. *Linear Algebra and its Applications*, 272(1-3):59–89, 1998.
- George Tucker, Andriy Mnih, Chris J Maddison, John Lawson, and Jascha Sohl-Dickstein. REBAR: Low-variance, unbiased gradient estimates for discrete latent variable models. In *Advances in Neural Information Processing Systems*, pages 2627–2636, 2017.
- Michel Verhaegen. Subspace model identification part 3. analysis of the ordinary output-error state-space model identification algorithm. *International Journal of Control*, 58(3):555–586, 1993.
- Roman Vershynin. *High-dimensional Probability: An introduction with Applications in Data Science*, volume 47. Cambridge University Press, 2018.
- Mats Viberg, Bo Wahlberg, and Björn Ottersten. Analysis of state space system identification methods based on instrumental variables and subspace fitting. *Automatica*, 33(9):1603–1616, 1997.
- D. Q. Wang. Least squares-based recursive and iterative estimation for output error moving average systems using data filtering. *IET Control Theory Applications*, 5(14):1648–1657, September 2011. ISSN 1751-8644. doi: 10.1049/iet-cta.2010.0416.
- Lex Weaver and Nigel Tao. The optimal reward baseline for gradient-based reinforcement learning. In *Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence*, pages 538–545. Morgan Kaufmann Publishers Inc., 2001.

Yong Zhang. Unbiased identification of a class of multi-input single-output systems with correlated disturbances using bias compensation methods. *Mathematical and Computer Modelling*, 53(9-10):1810–1819, May 2011. ISSN 08957177. doi: 10.1016/j.mcm.2010.12.059. URL <https://linkinghub.elsevier.com/retrieve/pii/S0895717711000045>.

Wei Xing Zheng. A revisit to least-squares parameter estimation of ARMAX systems. In *2004 43rd IEEE Conference on Decision and Control (CDC) (IEEE Cat. No.04CH37601)*, volume 4, pages 3587–3592 Vol.4, December 2004. doi: 10.1109/CDC.2004.1429269.

Kemin Zhou, John C. Doyle, and Keith Glover. *Robust and Optimal Control*, volume 40. Prentice Hall, 1996.

Contents

1	Introduction	1
1.1	Problem Statement	2
1.2	Prefiltered Least Squares (PF-LS)	3
1.3	Organization	4
2	Rates for Learning LTI Systems	5
2.1	Learning without the stability gap	6
2.2	Recovering the System Parameters	7
3	Oracle Inequality for Prefiltered Least Squares	7
3.1	Statement of the Oracle Inequality	9
4	Proof Sketch for Bounding Opt_μ	10
5	Related Work	12
	Preface	20
	Notation	21
I	Proof of Secondary Results	22
A	Examples of Phase Rank	22
B	Proof of Corollary B.1	25
B.1	Dependence on and Selection of T and L	26
B.2	Proof of Corollary B.1 (a)	27
B.3	Proof of Corollary B.1 (b)	27
B.4	Proof of Corollary B.1 (c) and (d)	27
C	Lower Bound for OLS for Marginally Stable Systems	28
C.1	Proof of Proposition C.2	29
C.1.1	Proof of Lemma C.4	30
C.2	Proof of Lemma C.3	31
II	General Bound for Prefiltered Least Squares	33
D	General Statement and Analysis PF-LS	33
D.1	Proof of Uniform Bound: Theorem D.1, Part (a)	34
D.2	Proof of $\widehat{\text{Opt}}_\mu$ bound, Theorem D.1, Part (b)	36
D.2.1	Proof of Lemma D.3	37
D.2.2	Proof of Lemma D.4	38

D.2.3	Proof of Lemma D.5	38
E	Semi-Parametric Regression	39
E.1	Proof of Theorem E.1 , Part (a)	40
E.2	Proof of Theorem E.1 , Part (b)	43
E.3	Proof of Lemma E.3	44
E.4	Proof of Lemma E.2	45
F	Chaining for Self-Normalized Tail Inequalities	45
F.1	Proof of Proposition E.5	47
F.1.1	Proof of Lemma F.4	48
F.2	Proof of Theorem F.2	49
III	Prefiltered Least Squares for Linear Dynamical Systems	51
G	Bounds on $\ \Delta_\phi\ _{\text{op}}$	51
G.1	Outline of the Proofs	51
G.2	Bounding $\mathbf{G}^{(1)}$	53
G.3	Bounding $\mathbf{G}^{(2)}$	54
G.4	Additional terms	56
G.4.1	Process noise	56
G.4.2	Output noise	56
G.4.3	Contribution of the initial state	56
H	Definition of \overline{M}, $\overline{M}_{\text{adv}}$, and Proof of Proposition 2.1	57
H.1	Notation	57
H.2	Proof of Proposition 2.1	57
H.3	Proof of Lemma H.2	60
H.4	Selecting the parameter L	61
H.4.1	Proof Sketch of Proposition H.4	63
I	Polynomial Approximations and Phase Rank	64
I.1	Main Results	65
I.1.1	K_1 and K_2 : Controlling Poles and Markov Operator Norms	65
I.1.2	Main Results: Stochastic Noise with Block-Scalar Filters	67
I.1.3	Results for Adversarial Noise	67
I.1.4	Bounds for Disentangling Filters	68
I.2	Polynomial Approximations for Linear Dynamical Systems	69
I.2.1	Approximations Using Block-Scalar Filters: Proof of Proposition I.1	71
I.2.2	Approximations using Disentangling Filters	72
I.3	Proof of Theorem I.4	72

J	Supporting Proofs	73
J.1	Proofs for Section I.2	73
J.1.1	Proof of Theorem I.12	73
J.1.2	Proof of Proposition I.13	74
J.1.3	Proof of Theorem I.14	76
J.1.4	Proof of Proposition I.11	78
J.2	Bounds on Finite System Norms: Proof of Proposition I.2	82
K	Bounds under Strong Observability	83
K.1	Granular Bounds for Strong Observability	85
K.2	Proof of Theorem K.2	85
K.3	Proof of Proposition K.3	87
K.3.1	Proof of Lemma K.7	88
K.3.2	Proof of Lemma K.6	88

Preface

The appendix is divided into three main parts. Part I begins with Appendix A, which provides illustrated examples of the phase rank condition for various systems. This is then followed by proofs of our secondary results: the proof of Corollary B.1 from Proposition 2.1 and the proof of the lower bound for ordinary least squares, Theorem C.1. Part II contains the supporting material for the results in Section 3, as well as generalizations beyond the setting of linear dynamical systems. Our bounds make use of a general recipe for applying chaining to self-normalized martingale inequalities, described in Appendix F. Lastly, Part III provides the analysis underlying the results in Section 4. Appendix G gives the corresponding results bounding $\|\Delta_\phi\|_{\text{op}}$ in terms of various control theoretic quantities. Appendix H gives a detailed proof of Proposition 2.1, a specific version of our generalized oracle inequality for linear dynamical systems. The constant \bar{M} , its analogue \bar{M}_{adv} for adversarial noise, as well as the intermediate constants M_B, M_C, M_D and M_0 , are defined in Appendix H.1. This section also includes Appendix H.4, which presents and analyzes a procedure for selecting the parameter L in a data-dependent fashion, as well as defining refinements of the constant \bar{M} . Appendices I and J give more granular interpretations of our estimation bounds in terms of the phase rank, as well as supporting technical proofs. Finally, Appendix K defines *strong observability* and gives alternative interpretations of our estimation bounds in terms of this quantity.

Notation

General Mathematical Notation
\lesssim denotes inequality up to a universal constant.
$\mathcal{S}^{d-1} := \{v \in \mathbb{R}^d : \ v\ _2 = 1\}$
$\log_+(x) := \max\{1, \log(x)\}$
$\text{lil}(x) := \log_+(\log_+(x))$
$[n] := \{1, \dots, n\}$
$\ \cdot\ _{\text{op}}$ denotes matrix operator norm
$\ \cdot\ _{\text{F}}$ denotes matrix Frobenius norm
$\ \cdot\ _2$ denotes vector two-norm
$\sigma_k(\cdot)$ denotes the k -th largest singular vector
$\sigma_{\min}(A)$ denotes $\sigma_{n \wedge m}$ for $A \in \mathbb{C}^{n \times m}$
$\text{cond}(A) := \frac{\ A\ _{\text{op}}}{\sigma_{\min}(A)}$ denotes the condition number

Semi-Parametric Notation
$N \in \mathbb{N}$ denotes the sample size,
$\mathbf{u}_1, \dots, \mathbf{u}_N \in \mathbb{R}^p$ denote inputs
$\mathbf{y}_1, \dots, \mathbf{y}_N \in \mathbb{R}^m$ denote observations
$\mathbf{k}_{N_1}, \dots, \mathbf{k}_N \in \mathbb{R}^{\bar{L}}$ denote prefiltering features
$\bar{\mathbf{u}}_t := [\mathbf{u}_t^\top \mathbf{u}_{t-1}^\top \dots \mathbf{u}_{t-T}^\top]^\top \in \mathbb{R}^{Tp}$, for length $T \in \mathbb{N}$
$\bar{\mathbf{U}}$ denotes the matrix whose rows are $\bar{\mathbf{u}}_{N_1}, \dots, \bar{\mathbf{u}}_N$
\mathbf{Y} denotes the matrix whose rows are $\mathbf{y}_{N_1}, \dots, \mathbf{y}_N$
\mathbf{K} denotes the matrix whose rows are $\mathbf{k}_{N_1}, \dots, \mathbf{k}_N$
$\boldsymbol{\delta}_t = \mathbf{y}_t - G_* \mathbf{u}_t$ denotes semiparametric error
$\boldsymbol{\Delta}$ denotes the matrix whose rows are $\boldsymbol{\delta}_{N_1}, \dots, \boldsymbol{\delta}_N$
$\{\mathcal{F}_t\}$ denotes our filtration, $(\boldsymbol{\delta}_t)$ is $\{\mathcal{F}_{t-T}\}$ adapted
(\mathbf{u}_t) is $\{\mathcal{F}_t\}$ -adapted, $\mathbf{u}_t \mathcal{F}_{t-1} \sim \mathcal{N}(0, I_p)$
$\mathbf{k}_t \in \mathbb{R}^{\bar{L}}$ are $\{\mathcal{F}_{t-T}\}$ -adapted prefiltering features
N_1 denotes first recorded observation
$\tilde{N} = N - N_1$ is effective sample size
$N_{\min} = cTp \log^4(Tp)$ for a sufficiently large Tp
ϕ_{rdg} is the filter from (1.4)
\hat{G}_{PF} is the estimator from (1.5)
\hat{G}_{LS} is the least squares estimator from (1.3)
$\hat{G}_{\text{fil}}(\phi)$ is the fixed-filter estimator from (3.9)

LTI System Notation
<p> $(\mathbf{u}_t) \subset \mathbb{R}^p$ denote inputs, $(\mathbf{x}_t) \subset \mathbb{R}^n$ denote states, $(\mathbf{y}_t) \subset \mathbb{R}^m$ denote observations $(\mathbf{w}_t) \subset \mathbb{R}^{d_w}$ denote process noise, $(\mathbf{z}_t) \subset \mathbb{R}^{d_z}$ denote sensor noise The parameters $(A_\star, B_\star, C_\star, D_\star)$ and (B_w, D_z) are clarified in (1.2) $G_\star := [D_\star \mid C_\star B_\star \mid C_\star A_\star B_\star \mid \cdots \mid C_\star A_\star^{T-2} B_\star] \in \mathbb{R}^{m \times Tp}$ </p>
<p> $\mathbf{k}_t := [\mathbf{y}_{t-T}^\top \mid \mathbf{y}_{t-2T}^\top \mid \cdots \mid \mathbf{y}_{t-LT}^\top]^\top$. $\mathbf{G} = (A, B, C, D)$ is a place holder variable for dynamical systems $\mathbf{G}(z) = D + C(zI - A)^{-1}B$ for $\mathbf{G} = (A, B, C, D)$ $\mathbf{G}_\star = D_\star + C_\star(zI - A)^{-1}B_\star$ $\mathbf{F}_\star = D_\star + C_\star(zI - A)^{-1}B_w$. $\mathbf{H}_\star = D_\star + C_\star(zI - A)^{-1}\mathbf{x}_1$. $\phi = [\Psi_1 \mid \cdots \mid \Psi_L]$ for $\phi \in \mathbb{R}^{m \times Lm}$. $C_\phi = C_\star - \sum_{\ell=1}^L \Psi_\ell C_\star$. $\mathbf{G}_\phi = (A_\star, B_\star, C_\phi, D_\star)$ $\mathbf{F}_\phi = (A_\star, B_w, C_\phi, D_\star)$ $\mathbf{H}_\phi = (A_\star, B_w, C_\phi, \mathbf{x}_1)$ </p>
<p> $\mathcal{M}_n(\mathbf{G}) = [D \mid CB \mid CAB \mid \cdots \mid CA^{n-2}B]$ for $\mathbf{G} = (A, B, C, D)$ $\ \mathcal{M}_n(\mathbf{G})\ _{\text{op}} = \ \mathcal{M}_n(\mathbf{G})\ _{\text{op}}$ for $\mathbf{G} = (A, B, C, D)$ $\ \mathcal{M}_\infty(\mathbf{G})\ _{\text{op}} = \lim_{n \rightarrow \infty} \ \mathcal{M}_n(\mathbf{G})\ _{\text{op}}$ $\ \mathbf{G}\ _{\mathcal{H}_{\text{inf}}} = \sup_{z \in \mathbb{C}: z =1} \ \mathbf{G}(z)\ _\infty$ for $\mathbf{G} = (A, B, C, D)$ $\Gamma_N(\mathbf{G}) = \min\{\sqrt{N}\ \mathcal{M}_\infty(\mathbf{G})\ _{\text{op}}, \ \mathbf{G}\ _{\mathcal{H}_{\text{inf}}}\}$ $\ \mathbf{G}\ _{\mathcal{H}_2^{\text{op}}} = \max_{v \in S^{m-1}} \sqrt{\frac{1}{2\pi} \int_0^{2\pi} \ v^\top \mathbf{G}(e^{i\pi\theta})\ _2^2} d\theta$ for $\mathbf{G} = (A, B, C, D)$ $\ \mathbf{G}\ _{\mathcal{H}_2^{\text{op}}} = \ \mathcal{M}_\infty(\mathbf{G})\ _{\text{op}}$ (Lemma 1.10) </p>
<p> $\text{blkspec}(A_\star)$ denotes the set of pairs (λ, k) corresponding to eigenvalues of A_\star and corresponding Jordan block sizes k $A_\star = S J_\star S^{-1}$ denotes the Jordan decomposition of A_\star (note $\text{blkspec}(A_\star) = \text{blkspec}(J_\star)$) $M_B, M_C, M_D, M_0, \bar{M}, \bar{M}_{\text{adv}}$ are constants clarified in Appendix H.1. </p>

Part I

Proof of Secondary Results

Appendix A. Examples of Phase Rank

We first recall the definition of phase rank.

Definition 2.3 (Phase Rank) *Let $\alpha \geq 1$. We say that A_\star has (α, T) -phase rank d if there exists $\mu_1, \dots, \mu_d \in \mathbb{D}$ such that, for any $(\lambda, k) \in \text{blkspec}(A_\star)$ with $|\lambda| \geq 1 - ((1 + \alpha)T)^{-1}$, there exists at least k elements $\mu_{i_1}, \dots, \mu_{i_k} \subset \{\mu_1, \dots, \mu_d\}$ satisfying*

$$\max_{j \in [k]} \min_{\tilde{\mu}: \tilde{\mu}^T = \mu_{i_j}^T} |\lambda - \tilde{\mu}| \leq \alpha(1 - |\lambda|).$$

Phase rank represents how many phases are required to cover ‘‘large eigenvalues’’: eigenvalues with magnitude at least $1 - ((1 + \alpha)T)^{-1}$, where $\alpha \geq 1$ and $T \geq 1$ is an integer. Moreover, the

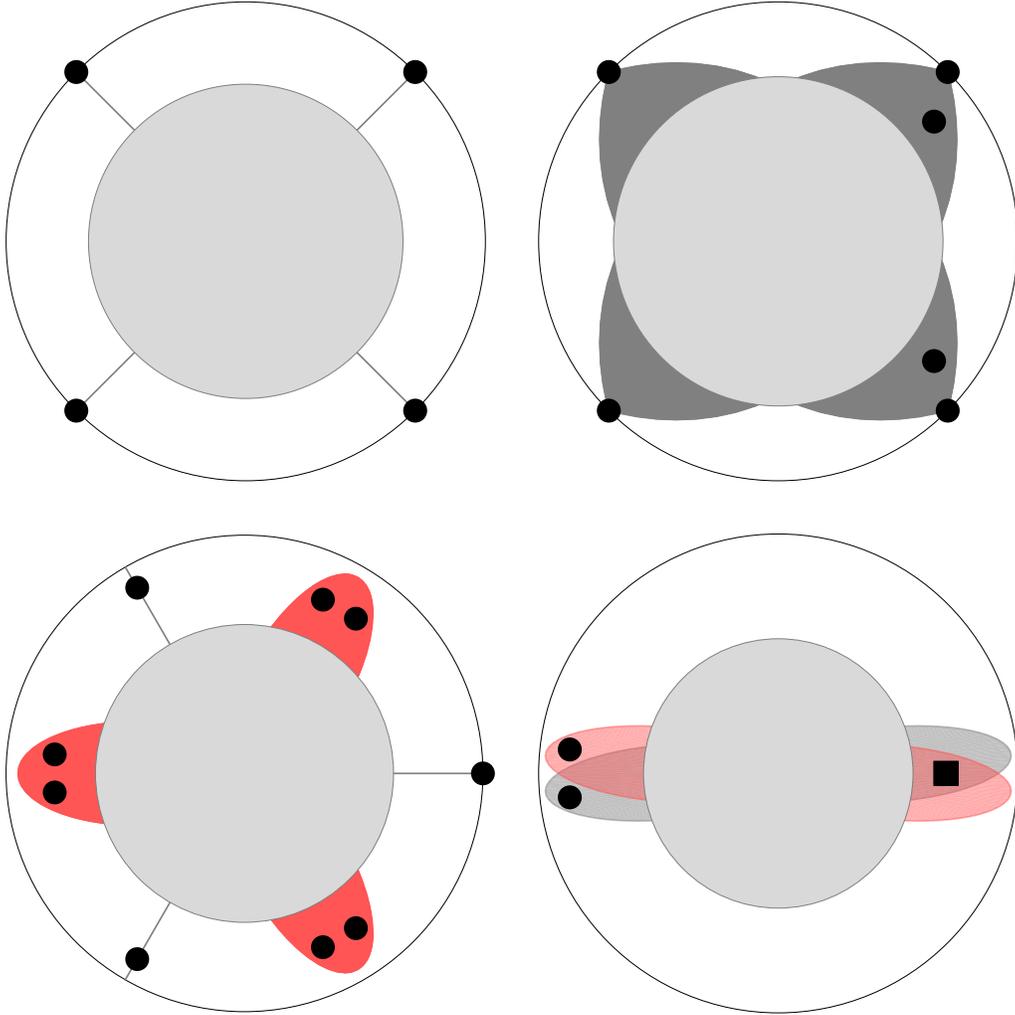


Figure 1: Examples (a), (b), (c), and (d) of phase rank, ordered clockwise from the top-left.

condition $\tilde{\mu}^T = \mu_{i_j}^T$ means we actually only care about phases mod $2\pi/T$. Graphically, this means every μ_i contributes T “bumps” or “spokes” toward covering the eigenvalues of A_* . Figure 1 gives four example spectra for which we will calculate the phase rank. Each example shows the regions in the complex disk we are covering by the choices of μ_i which witness the phase rank conditions. First, we must be clear with graphical notation. The circular gray region⁵ represents the region of “small” eigenvalues, and we will assume spectra lie outside of this region without loss of generality. Single eigenvalues are represented by a dot; doubly repeated eigenvalues by a square. Examples (a), (c), (d) have $\alpha = 1$, whereas (b) depicts $\alpha = 2$ to demonstrate the effect of increasing the parameter.

- (a) In example (a), all eigenvalues lying on the unit circle means we *must* choose each μ_i to have modulus 1. As a result, T equally-spaced spokes are added to the covering region for each μ_i . We then see the $(1, 2)$ phase rank is 2 and the $(1, 4)$ phase rank is 1.

5. These regions have been artificially shrunk in Figure 1 for the sake of legibility, but the examples are morally correct.

- (b) The α parameter controls the width of the covering regions. We see in example (b) that increasing α to 2 allows us to cover the two additional eigenvalues, such that the $(2, 4)$ phase rank is again 1.
- (c) When there are spectra with modulus strictly less than 1, the μ_i witnessing the phase rank condition may also have modulus strictly less than one; this also results in the spoke regions transforming into “bumps”. Example (c) illustrates the $(1, 3)$ phase rank being 2, as the μ_i associated to the red bumps allows covering of the eigenvalues which are not quite at phases $\{\pi/3, -\pi/3, \pi\}$.
- (d) Recall that we need to cover a repeated eigenvalue multiple times as to its multiplicity. However, in example (d), we are able to do so while also covering other eigenvalues; the $(1, 2)$ phase rank is indeed 2.

Appendix B. Proof of Corollary B.1

To state the corollary for system parameter recovery, we first need to define the Hankel matrices used by the Ho-Kalman algorithm.

Definition B.1 (Hankel Matrices) *Given $T \geq 4$, let $T_1 := \lfloor (T - 1)/2 \rfloor$. We define the Hankel matrix $\mathcal{H} \in \mathbb{R}^{T_1 m \times (T_1 + 1)p}$ to be the block matrix⁶ with*

$$\mathcal{H}[i, j] := G_\star[i + j] = C_\star A_\star^{i+j-2} B_\star,$$

and define $\mathcal{H}^- \in \mathbb{R}^{T_1 m \times T_1 p}$ to be the size Hankel matrix created by dropping the last block column of \mathcal{H} .

In order to learn $(A_\star, B_\star, C_\star, D_\star)$ that has state dimension (i.e. system order) n , we need to ensure that $T_1 \geq n$. This assumption, along with an assumed bound on the discrepancy between \mathcal{H} and its estimated version $\widehat{\mathcal{H}}$, are used in [Oymak and Ozay \(2018\)](#) to prove consistency of the Ho-Kalman algorithm (see Appendix B for details). With these assumptions, we can now fully state a corollary for the recovery of system parameters.

Corollary B.1 (Recovery of System Parameters) *Consider a linear dynamical system given by $(A_\star, B_\star, C_\star, D_\star)$, and parameters T and L such that*

- $A_\star \in \mathbb{R}^{n \times n}$, $C_\star \in \mathbb{R}^{m \times n}$, $B_\star \in \mathbb{R}^{n \times p}$, $D_\star \in \mathbb{R}^{n \times p}$.
- Let $(\overline{A}, \overline{B}, \overline{C}, \overline{D})$ denote the realization of $(A_\star, B_\star, C_\star, D_\star)$ corresponding to the applying the Ho-Kalman algorithm to G_\star .
- $(A_\star, B_\star, C_\star, D_\star)$ is minimal ([Definition B.2](#)) with Hankel matrix \mathcal{H} . We also assume that $\lfloor T - 1/2 \rfloor \geq n$, which ensures $\sigma_{\min}(\mathcal{H}^-) > 0$.
- $(\widehat{A}, \widehat{B}, \widehat{C}, \widehat{D})$ is the output of the Ho-Kalman algorithm on \widehat{G}_{PF} , with system order n .
- We have $\rho(A_\star) \leq 1$, the largest Jordan block of A_\star is size at most $k \geq 1$, and A_\star has (α, T) -phase rank at most $k \leq d \leq L$.
- Let \overline{M} be as in [Definition 2.2](#), and define the parameter $\mathbf{d}_L := (p + Lmk)$.

Then, under the stochastic noise model with $\mathbf{x}_1 = 0$, as long as $N_1 = TL$ and

$$N \geq T(d + 1 + \alpha) \max\{m, \log(1/\delta)\} + N_{\min},$$

it holds with probability $1 - \delta - \delta_{\overline{U}}$, there is a unitary matrix $S \in \mathbb{R}^{n \times n}$ such that

$$\begin{aligned} (a) \quad N^{-1/2} \text{Opt}_\mu &\lesssim (\overline{M} + 1) \cdot T^{k-1/2} C_{\alpha, d, k} \\ (b) \quad \|\widehat{G}_{\text{PF}} - G_\star\|_{\text{op}} &\lesssim \underbrace{\frac{T^k C_{\alpha, d, k} (1 + \overline{M}) \widetilde{\mathcal{O}}(\mathbf{d}_L)^{1/2}}{N^{1/2}}}_{\text{'lead term'}} + \underbrace{\frac{T^{k+1/2} (1 + \overline{M}) \widetilde{\mathcal{O}}(\mathbf{d}_L)}{N}}_{\text{'lower order term'}}. \end{aligned}$$

6. Note that this does not include D_\star .

If, in addition, N is sufficiently large such that $\|G_\star - \widehat{G}\|_{\text{op}} \leq \sigma_{\min}(\mathcal{H}^-)/(4\sqrt{T})$, then

$$(c) \quad \max \left\{ \|\widehat{B} - S\overline{B}\|_{\text{F}}, \|\widehat{C} - \overline{C}S^*\|_{\text{F}}, \|\widehat{D} - \overline{D}\|_{\text{F}} \right\} \\ \lesssim n^{\frac{1}{2}}T^{1/4} \left(\frac{T^k C_{\alpha,d,k}(1 + \overline{M}) \widetilde{\mathcal{O}}(\mathbf{d}_L)^{\frac{1}{2}}}{N^{\frac{1}{2}}} + \frac{T^{k+\frac{1}{2}}(1 + \overline{M}) \widetilde{\mathcal{O}}(\mathbf{d}_L)}{N} \right) \\ (d) \quad \|\widehat{A} - \overline{A}\|_{\text{F}} \lesssim \\ \frac{n^{\frac{1}{2}}T^{1/4}\|\mathcal{H}\|_{\text{op}}}{\sigma_{\min}(\mathcal{H}^-)^{\frac{3}{2}}} \left(\frac{(T^k C_{\alpha,d,k}(1 + \overline{M}))^{1/2} \widetilde{\mathcal{O}}(\mathbf{d}_L)^{\frac{1}{4}}}{N^{\frac{1}{4}}} + \sqrt{\frac{T^{k+1/2}(1 + \overline{M}) \widetilde{\mathcal{O}}(\mathbf{d}_L)}{N}} \right),$$

where the notation $\widetilde{\mathcal{O}}(\cdot)$ hides factors at most linear in $\log N$, $\log(1/\delta)$, $\log T$, and/or $\log_+(\overline{M} + d_z + d_w)$. When $\mathbf{x}_1 \neq 0$, we use the more general definition of \overline{M} from Definition H.1.

Under the adversarial noise model, the above bounds hold by (i) replacing \overline{M} with $\overline{M}_{\text{adv}}$ (Definition H.1) and (ii) replacing the factors T^k and $C_{\alpha,d,k}$ in the ‘lead term’ in bound (b) with $T^{k+1/2}$ and $C_{\alpha,d,k}^{\text{adv}} := 2^d \left(\frac{k^2}{2^k} (1 + \alpha)^k + d^k \right)$, respectively, and analogously for the subsequent bounds (c) and (d).

As a consequence, we verify that combining prefiltered least squares with the Ho-Kalman algorithm analyzed in Oymak and Ozay (2018) provides the consistent estimation of the underlying system parameters themselves.

We remark that the condition that $(A_\star, B_\star, C_\star, D_\star)$ is minimal is necessary to ensure identifiability in the manner described by Corollary B.1. When minimality fails, we explain below that there always exists a reduced system model equivalent (in an input-output sense) to $(A_\star, B_\star, C_\star, D_\star)$, which can be estimated in the sense of Corollary B.1.

Remark Concurrent with the preparation of this manuscript, Sarkar and Rakhlin (2018) introduced an improved analysis of the Ho-Kalman algorithm which, in our setting, enables a rate of estimation of $\|\widehat{A} - A_\star\|_2 \leq C\|G_\star - \widehat{G}\|$, which enables a rate of estimation of $\|\widehat{A} - A_\star\| \leq \widetilde{\mathcal{O}}(C'/\sqrt{N})$ for some problem dependent constant C' . We do not pursue this more refined analysis here.

B.1. Dependence on and Selection of T and L

As mentioned above, the recovery of $(A_\star, B_\star, C_\star, D_\star)$ relies on three parameters: the feature length L , the Markov matrix dimension T , and the system order n . As noted above, the bounds are rather insensitive to the choice of regularizer.

Following Oymak and Ozay (2018), we assume that n is known, and that T should be chosen such that $c_1 n \leq \lfloor (T - 1)/2 \rfloor \leq T \leq c_2 n$, for some constants $c_1, c_2 > 1$. These constants ensure that the Hankel matrix \mathcal{H}^- has $\sigma_{\min}(\mathcal{H}^-) > 0$. A more data-dependent approach for selecting T and n , which is particularly useful for systems of unknown order, is pursued by Sarkar et al. (2019); for simplicity, we do not follow this direction here. The parameter L is more straightforward to select. While our given bounds require that $L \geq d$, where d is the phase rank of A , we show in Proposition H.4 that a simple structural risk minimization procedure can be used to tune the parameter as to satisfy an oracle inequality with respect to the optimal choice of L . The details are deferred to Appendix H.4. In particular, the oracle approach implies that for sufficiently large sample sizes N , the bounds in Corollary B.1 hold with L replaced by $\widetilde{\mathcal{O}}(d)$.

B.2. Proof of Corollary B.1 (a)

This bound follows from Proposition 2.2 and choosing (say) $\mu = 1$.

B.3. Proof of Corollary B.1 (b)

We will choose to describe the constant of Corollary B.1 (b) in terms of phase rank rather than strong observability, though both suffice to prove the corollary. Now, recall Proposition 2.1.

Proposition 2.1 *Fix T and L , and suppose that $N \geq N_{\min}$, $N_1 = TL$, $\rho(A_\star) \leq 1$, and that the largest Jordan block of A_\star is of size k . Choosing some $\mu \geq 1$ and defining*

$$\bar{\mathbf{d}} := \tilde{p} + Lm (\log_+ \bar{M} + k \log_+ N) = \tilde{\mathcal{O}}(p + Lmk), \quad (2.8)$$

it holds with probability at least $1 - \delta - (2Np)^{-\log^2(2Tp) \log^2(2Np)}$ in the stochastic noise model that

$$\|\widehat{G}_{\text{PF}} - G_\star\|_{\text{op}} \lesssim N^{-1/2} \text{Opt}_\mu \cdot \sqrt{\frac{T(\bar{\mathbf{d}} + \log \frac{1}{\delta})}{N} + \frac{T(\|G_\star\|_{\text{op}} + \mu)(\bar{\mathbf{d}} + \log \frac{1}{\delta})}{N}}.$$

In the adversarial noise model, we instead take $\bar{\mathbf{d}} := \tilde{p} + Lm (\log_+(\bar{M} + d_z + d_w) + k \log_+ N)$. Moreover, we have $\|G_\star\|_{\text{op}} \lesssim \bar{M} T^{k-1/2}$.

Now, taking (say) $\mu = 1$ so that $\mu N^{-1/2} \leq 1$, by Proposition 2.2 we have that

$$N^{-1/2} \text{Opt}_\mu \lesssim (\bar{M} + 1) \cdot T^{k-1/2} C_{\alpha, d, k}.$$

Note that our choice of $\bar{\mathbf{d}}$ ensures that $(\bar{\mathbf{d}} + \log(\frac{1}{\delta}))$ is $\tilde{\mathcal{O}}(\mathbf{d}_L)$. Thus, using the above Proposition (including the bound $\|G_\star\|_{\text{op}} \lesssim \bar{M} T^{k-1/2}$) gives

$$\|\widehat{G}_{\text{PF}} - G_\star\|_{\text{op}} \lesssim \frac{T^k C_{\alpha, d, k} (1 + \bar{M}) \tilde{\mathcal{O}}(\mathbf{d}_L)^{1/2}}{N^{1/2}} + \frac{T^{k+1/2} (1 + \bar{M}) \tilde{\mathcal{O}}(\mathbf{d}_L)}{N}.$$

The required length N then follows from the requirements on N for Propositions 2.2 and 2.1.

B.4. Proof of Corollary B.1 (c) and (d)

We begin by formally introducing standard regularity conditions in control theory, observability, controllability, and minimality:

Definition B.2 (Observability, Controllability, Minimality) *A linear system $(A_\star, B_\star, C_\star, D_\star)$ is said to be controllable if*

$$\text{rank} \left(\begin{bmatrix} B_\star & A_\star B_\star & \cdots & A_\star^{n-1} B_\star \end{bmatrix} \right) = n.$$

A system is said to be observable if

$$\text{rank} \left(\begin{bmatrix} C_\star \\ C_\star A_\star \\ \vdots \\ C_\star A_\star^{n-1} \end{bmatrix} \right) = n.$$

A linear system $(A_\star, B_\star, C_\star, D_\star)$ is said to be minimal if it is both observable and controllable.

Even if $(A_\star, B_\star, C_\star, D_\star)$ is not minimal, there always exists an $n' \leq n$ and an equivalent system $(\bar{A}, \bar{B}, \bar{C}, \bar{D})$ with $\bar{A} \in \mathbb{R}^{n' \times n'}$ such that $(\bar{A}, \bar{B}, \bar{C}, \bar{D})$ is minimal. In this case, the Ho-Kalman algorithm correctly recovers this reduced, minimal system. Section 4 of in [Oymak and Ozay \(2018\)](#) concerns the robustness of the Ho-Kalman algorithm [Ho and Kalman \(1966\)](#), which generates state-space matrices $A_\star, B_\star, C_\star, D_\star$ from the matrix of Markov parameters G_\star . Oymak and Ozay show how these estimates degrade when the matrix of Markov parameters is replaced by a noisy estimate.

Now, recall that $T_1 = \lfloor (T-1)/2 \rfloor$ and note that $2T_1 + 1 \leq T$. Furthermore, recall that we define the Hankel matrix $\mathcal{H} \in \mathbb{R}^{T_1 m \times (T_1+1)p}$ to be the block matrix⁷ with

$$\mathcal{H}[i, j] := G_\star[i + j] = C_\star A_\star^{i+j-2} B_\star,$$

and we now use $\hat{\mathcal{H}}$ to denote its analogous estimated version. We also define $\mathcal{H}^-, \hat{\mathcal{H}}^-$ to be the size (T_1, T_1) Hankel matrices created by dropping the last block column of $\mathcal{H}, \hat{\mathcal{H}}$ respectively.

At this stage, we would like to note that, in contrast to our main results for estimate \hat{G} , the bound requires that $\lfloor (T-1)/2 \rfloor \geq n$, where n is the order of the unknown system. Note that with this choice, under the assumption that $(A_\star, B_\star, C_\star, D_\star)$ is minimal, the Hankel matrix is rank- n and $\sigma_n(\mathcal{H}^-) > 0$ (as noted in Section 4.1 of [Oymak and Ozay \(2018\)](#)). Now, we synthesize their relevant results below, which makes use of our choice $n \leq T_1 \leq T$.

Proposition B.2 (Sections 4.1 and 4.2 of [Oymak and Ozay \(2018\)](#)) *Let $\bar{A}, \bar{B}, \bar{C}, \bar{D}$ be the state-space realization corresponding to the output of the Ho-Kalman algorithm with input G_\star and let $(\hat{A}, \hat{B}, \hat{C}, \hat{D})$ be the state-space realization corresponding to the output of Ho-Kalman with input \hat{G} . Suppose the system $(A_\star, B_\star, C_\star, D_\star)$ is observable and controllable, and suppose*

$$\|\mathcal{H} - \hat{\mathcal{H}}\|_{\text{op}} \leq \sigma_{\min}(\mathcal{H}^-)/4, \tag{B.11}$$

which, as noted by [Oymak and Ozay \(2018, Lemma 4.2\)](#), is satisfied as soon as $\|G_\star - \hat{G}\|_{\text{op}} \leq \sigma_{\min}(\mathcal{H}^-)/(4\sqrt{T})$. Then, there exists a unitary matrix S such that

$$\begin{aligned} \max \left\{ \|\hat{B} - S\bar{B}\|_{\text{F}}, \|\hat{C} - \bar{C}S^*\|_{\text{F}} \right\} &\lesssim n^{1/2}T^{1/4} \sqrt{\|G_\star - \hat{G}\|_{\text{op}}} \\ \|\hat{A} - S\bar{A}S^*\|_{\text{F}} &\lesssim \frac{n^{1/2}T^{1/4} \sqrt{\|G_\star - \hat{G}\|_{\text{op}} \|\mathcal{H}\|_{\text{op}}}}{\sigma_{\min}(\mathcal{H}^-)^{3/2}}. \end{aligned}$$

Since $\|\hat{D} - \bar{D}\|_{\text{F}} \leq \|G_\star - \hat{G}\|_{\text{F}}$, for N sufficiently large we can combine [Proposition B.2](#) and [Section B.3](#) to arrive at [Corollary B.1](#). The bounds (c) and (d) now follow from the above proposition applied to [Corollary B.1](#) (b) and the inequality $\sqrt{x+y} \leq \sqrt{x} + \sqrt{y}$.

Appendix C. Lower Bound for OLS for Marginally Stable Systems

In this section, we show that OLS cannot recover a linear system for a marginally stable system, even with zero process \mathbf{w}_t and sensor noise \mathbf{z}_t . Specifically, we have the following

⁷. Note that this does not include D_\star .

Theorem C.1 *Suppose that A_\star has a Jordan block of magnitude $|\lambda| = 1$ and size k , and that $(A_\star, B_\star, C_\star, D_\star)$ minimal. Suppose also that the process noise \mathbf{w}_t and sensor noise \mathbf{z}_t are identically zero. Then there exists a constant C depending on $A_\star, B_\star, C_\star$, as well as the ambient dimension n and k such that, for all N sufficiently large, $\|\widehat{G}_{\text{LS}} - G_\star\|_{\text{op}} \geq CN^{k-1}$ with constant probability.*

The zero-noise assumption is for simplicity, and the above lower bound can also be demonstrated in the presence of Gaussian noise. The important takeaway is that *even without noise*, the ordinary least squares estimator is inconsistent.

The proof of Theorem C.1 has two components. The first is a lower bound, based on small-ball technicals, which bounds $\|\widehat{G}_{\text{LS}} - G_\star\|_{\text{op}}$ in terms of the Gramian matrices

$$\mathcal{G}_t := \sum_{s=0}^t C_\star A_\star^{t-s} B_\star B_\star^\top A_\star^{t-s} C_\star^\top.$$

The following proposition is proved in Section C.1:

Proposition C.2 *For N sufficiently large, and with no process or sensor noise, there with constant probability*

$$\|\widehat{G}_{\text{LS}} - G_\star\|_{\text{op}} \gtrsim \frac{1}{N} \sqrt{\left\| \sum_{t=N_1}^N \mathcal{G}_{t-T} \right\|_{\text{op}}},$$

where we define the Gramian matrix

The second component, which completes the proof of Theorem C.1, is a lower bound estimate on the operator norm of the sum of the Gramian matrices, proved in Section C.2:

Lemma C.3 *Suppose that A_\star has a Jordan block of magnitude $|\lambda| = 1$ and size k , and that $(A_\star, B_\star, C_\star, D_\star)$ minimal. Then there exists a constant C depending on $A_\star, B_\star, C_\star$, as well as the ambient dimension n and k such that, for all $N \geq \max\{20n^2, 2k\}$,*

$$\sqrt{\left\| \sum_{t=0}^N \mathcal{G}_t \right\|_{\text{op}}} \geq CN^k.$$

C.1. Proof of Proposition C.2

When $N \geq \underline{N}$, Oymak and Ozay (2018, Lemma C.2) yields that $\sigma_{\max}(\overline{\mathbf{U}})^2 \leq 2N$ with probability $1 - \delta_{\overline{\mathbf{U}}}$ for $\delta_{\overline{\mathbf{U}}} = \frac{1}{N^{\omega(1)}}$, and on this event,

$$\|\widehat{G}_{\text{LS}} - G_\star\|_{\text{op}} = \|\overline{\mathbf{U}}^\dagger \Delta\|_{\text{op}} \geq \sigma_{\max}(\overline{\mathbf{U}})^{-2} \|\overline{\mathbf{U}}^\dagger \Delta\|_{\text{op}} \geq \frac{\|\overline{\mathbf{U}}^\top \Delta\|_{\text{op}}}{2N} \geq \frac{\|\mathbf{U}^\top \Delta\|_{\text{op}}}{2N},$$

where we recall that with our notation, \mathbf{U} is the matrix corresponding to the first p columns of $\overline{\mathbf{U}}$. We now lower bound $\|\overline{\mathbf{U}}^\top \Delta\|_{\text{op}}$. Observe that when \mathbf{w}_t and $\mathbf{z}_t = 0$, we can write

$$\delta_t = \sum_{s=1}^{t-T} C_\star A_\star^{t-s} B_\star \mathbf{u}_s.$$

Observe that we can represent this quantity as a quadratic form in a long vector $\mathbf{u}_{1:N} \in \mathbb{R}^{Np}$. We now invoke the follow lemma

Lemma C.4 (Small Ball for Gaussian Quadratic Forms) *There exists a universal constant c such that, for any matrix $M \in \mathbb{R}^{d \times d}$ and vector $\mathbf{z} \sim \mathcal{N}(0, I_d)$,*

$$\mathbb{P}[(\mathbf{z}^\top M \mathbf{z})^2 \geq \frac{1}{2} \mathbb{E}(\mathbf{z}^\top M \mathbf{z})^2] \geq c.$$

We now compute that

$$v^\top \bar{\mathbf{U}}^\top \Delta w = \sum_{t=N_1}^N v^\top \mathbf{u}_t \delta_t^\top w = \sum_{t=N_1}^N \sum_{s=1}^{t-T} \langle v, \mathbf{u}_t \rangle \langle (C_\star A_\star^{t-s} B_\star)^\top w, \mathbf{u}_s \rangle$$

And thus, for $v \in \mathcal{S}^{p-1}$,

$$\begin{aligned} \mathbb{E}(v^\top \bar{\mathbf{U}}^\top \Delta w)^2 &= \mathbb{E} \sum_{t=N_1}^N \sum_{s=1}^{t-T} \langle v, \mathbf{u}_t \rangle^2 \langle (C_\star A_\star^{t-s} B_\star)^\top w, \mathbf{u}_s \rangle^2 + \mathbb{E}[\text{mean zero terms}] \\ &= \sum_{t=N_1}^N \sum_{s=1}^{t-T} \|C_\star A_\star^{t-s} B_\star\|_2^2 w^\top w \\ &= w^\top \left(\sum_{t=N_1}^N \sum_{s=1}^{t-T} C_\star A_\star^{t-s} B_\star B_\star^\top A_\star^{t-s} C_\star^\top \right) w = w^\top \left(\sum_{t=N_1}^N \mathcal{G}_{t-T} \right) w. \end{aligned}$$

Thus, optimizing for w to be a lead eigenvector of $\sum_{t=N_1}^N \mathcal{G}_{t-T}$, we see that

$$\mathbb{P} \left[\|\bar{\mathbf{U}}^\top \Delta\|_{\text{op}} \geq \frac{1}{\sqrt{2}} \sqrt{\left\| \sum_{t=N_1}^N \mathcal{G}_{t-T} \right\|_{\text{op}}} \right] \geq \mathbb{P} \left[v^\top \bar{\mathbf{U}}^\top \Delta w \geq \frac{1}{2} \left\| \sum_{t=N_1}^N \mathcal{G}_{t-T} \right\|_{\text{op}} \right] \geq c.$$

for a universal constant c . Hence, we conclude that, with probability at least $c - \delta_{\bar{\mathbf{U}}}$,

$$\|\hat{G}_{\text{LS}} - G_\star\|_{\text{op}} \geq \frac{\sqrt{\|\bar{\mathbf{U}}^\top \Delta\|_{\text{op}}^2}}{2N} \geq \frac{\sqrt{\left\| \sum_{t=N_1}^N \mathcal{G}_{t-T} \right\|_{\text{op}}}}{2\sqrt{2}N}.$$

C.1.1. PROOF OF LEMMA C.4

We may assume that M is symmetric, since replacing M by $\frac{1}{2}(M + M^\top)$ does not affect the result. Since \mathbf{z} has a unitary invariant distribution, we may also assume that M is diagonal whose vector is a diagonal $\mathbf{a} \in \mathbb{R}^d$. Note that $\|\mathbf{a}\|_2 = \frac{1}{2} \|M + M^\top\|_{\text{F}}$. Then, $\mathbf{z}^\top M \mathbf{z} = \sum_{i=1}^n \mathbf{a}_i z_i^2$. Let $Z = (\mathbf{z}^\top M \mathbf{z})^2$. The Paley-Zygmund inequality states that

$$\mathbb{P}[Z \geq \theta \mathbb{E}[Z]] \geq (1 - \theta^2) \frac{\mathbb{E}[Z]^2}{\mathbb{E}[Z^2]}.$$

First, we have that

$$\mathbb{E}Z = \mathbb{E}\left(\sum_{i=1}^n \mathbf{a}_i^2 \mathbf{z}_i^4\right) + \mathbb{E}\left(\sum_{i \neq j=1}^n \mathbf{a}_i \mathbf{a}_j \mathbf{z}_i^2 \mathbf{z}_j^2\right) = 3 \sum_{i=1}^n \mathbf{a}_i^2 + \sum_{i \neq j=1}^n \mathbf{a}_i \mathbf{a}_j = 2\|\mathbf{a}\|_2^2 + \left(\sum_{i=1}^n \mathbf{a}_i\right)^2.$$

On the other hand, a standard χ^2 -concentration (Laurent and Massart, 2000, Lemma 1) inequality (adapted for nonnegative coefficients) yields:

$$\mathbb{P}\left[\left|\mathbf{z}^\top M \mathbf{z} - \sum_i \mathbf{a}_i\right| \geq 2(\|\mathbf{a}\|_2 \sqrt{t} + \|\mathbf{a}\|_\infty t)\right] \leq 2e^{-t},$$

which crudely implies $\mathbb{P}[Z^2 \gtrsim (\sum_i \mathbf{a}_i)^4 + \|\mathbf{a}\|_2^2 t^4] \leq 2e^{-t}$. Integrating, we find that $\mathbb{E}[Z^2] \lesssim (\sum_i \mathbf{a}_i)^4 + \|\mathbf{a}\|_2^4$. Thus, Paley-Zygmund and the lower bound on $\mathbb{E}Z$ imply for a universal constant c that

$$\mathbb{P}[Z \geq \theta \mathbb{E}[Z]] \geq (1 - \theta^2) \frac{\mathbb{E}[Z]^2}{\mathbb{E}[Z^2]} \geq (1 - \theta^2) \frac{(2\|\mathbf{a}\|_2^2)^2}{\mathbb{E}[Z^2]} \geq c(1 - \theta^2)$$

for a universal constant c . Taking $\theta = 1/2$ concludes.

C.2. Proof of Lemma C.3

If $A_\star \in \mathbb{R}^{n \times n}$, then $\text{rank}(\sum_{t=0}^N \mathcal{G}_t) \leq n$, and thus $\sqrt{\|\sum_{t=0}^N \mathcal{G}_t\|_{\text{op}}} \geq \sqrt{\frac{1}{n} \text{tr}(\sum_{t=0}^N \mathcal{G}_t)}$. We now turn to lower bounding $\text{tr}(\sum_{t=0}^N \mathcal{G}_t)$. We can bound

$$\begin{aligned} \text{tr}\left(\sum_{t=0}^N \mathcal{G}_t\right) &= \text{tr}\left(\sum_{t=1}^N \sum_{s=0}^{t-1} C_\star A_\star^s B_\star B_\star^\top (A_\star^s)^\top C_\star\right) \\ &= \sum_{t=0}^{N-1} \sum_{s=0}^t \|C_\star A_\star^s B_\star\|_{\text{F}}^2 \\ &\geq \sum_{j=0}^{\lfloor (N-1)/n \rfloor - 1} \sum_{p=0}^{n-1} \sum_{s=0}^{nj+p} \|C_\star A_\star^s B_\star\|_{\text{F}}^2 \\ &\geq \sum_{j=0}^{\lfloor (N-1)/n \rfloor - 1} \sum_{p=0}^{n-1} \sum_{s=0}^{nj} \|C_\star A_\star^{s+p} B_\star\|_{\text{F}}^2 \\ &\geq \sum_{j=0}^{\lfloor (N-1)/n \rfloor - 1} \sum_{p=0}^{n-1} \sum_{\ell=0}^{j-1} \sum_{q=0}^{n-1} \|C_\star A_\star^{\ell n + q + p} B_\star\|_{\text{F}}^2 \\ &\geq \sum_{j=0}^{\lfloor (N-1)/n \rfloor - 1} \sum_{\ell=0}^{j-1} \left(\sum_{p=0}^{n-1} \sum_{q=0}^{n-1} \|C_\star A_\star^{\ell n + q + p} B_\star\|_{\text{F}}^2 \right) \end{aligned}$$

Next, we define the controllability matrix

$$\mathcal{C}_n := \text{rank}\left(\begin{bmatrix} B_\star & A_\star B_\star & \cdots & A_\star^{n-1} B_\star \end{bmatrix}\right) = n.$$

and the observability matrix

$$\mathcal{O}_n := \text{rank} \left(\begin{bmatrix} C_\star \\ C_\star A_\star \\ \vdots \\ C_\star A_\star^{n-1} \end{bmatrix} \right),$$

By observability and controllability, $M := \sigma_n(\mathcal{O}_n) \cdot \sigma_n(\mathcal{C}_n) > 0$. Moreover, we observe that

$$\left(\sum_{p=0}^{n-1} \sum_{q=0}^{n-1} \|C_\star A_\star^{\ell n+q+p} B_\star\|_{\mathbb{F}}^2 \right) = \|\mathcal{O}_n A_\star^{kn} \mathcal{C}_n\|_{\mathbb{F}}^2 \geq M^2 \|A_\star^{\ell n}\|_{\mathbb{F}}^2.$$

Therefore,

$$\text{tr} \left(\sum_{t=0}^N \mathcal{G}_t \right) \geq M^2 \sum_{j=0}^{\lfloor (N-1)/n \rfloor - 1} \sum_{\ell=0}^{j-1} \|A_\star^{\ell n}\|_{\mathbb{F}}^2.$$

Next, we lower bound $\|A_\star^{\ell n}\|_{\mathbb{F}}^2$. Let $A_\star = S^{-1} J_\star S$. Then,

$$\|A_\star^{kn}\|_{\mathbb{F}} \geq \|J_\star\|_{\mathbb{F}} \sigma_{\min}(S^{-1}) \sigma_{\min}(S) = \text{cond}(S) \|J_\star^{\ell n}\|_{\mathbb{F}}.$$

Lastly, we note that if J_\star has a Jordan block of eigenvalue $|\lambda| = 1$ and multiplicity, $J_\star^{\ell n}$ then it contains an entry of magnitude $\binom{\ell n}{k-1}$. Thus, we can crudely lower bound

$$\begin{aligned} \text{tr} \left(\sum_{t=0}^N \mathcal{G}_t \right) &\geq M^2 \text{cond}(S)^{-2} \sum_{j=0}^{\lfloor (N-1)/n \rfloor - 1} \sum_{\ell=0}^{j-1} \left(\binom{\ell n}{k} \right)^2 \\ &\gtrsim M^2 \text{cond}(S)^{-2} \frac{N^2}{n^2} \left(\frac{N/2}{k-1} \right)^2, \end{aligned}$$

provided that $N \geq \max\{20n^2, 2k\}$. Hence,

$$\sqrt{\left\| \sum_{t=0}^N \mathcal{G}_t \right\|_{\text{op}}} \gtrsim \frac{M \text{cond}(S) N}{n^{3/2}} \left(\frac{N/2}{k-1} \right) \geq \mathbf{C} N^k,$$

where the last line is by Stirling's approximation (where \mathbf{C} depends on k).

Part II

General Bound for Prefiltered Least Squares

Appendix D. General Statement and Analysis PF-LS

In this section, we a more explicit version of Theorem 3.3, which decouples an oracle bound with a bound on $\|\phi_{\text{rdg}}\|_{\text{op}}$. First, we recall the definitions of $\text{Ovfit}_\mu(\delta)$ and of the optimal term Opt_μ , overfit term $\text{Ovfit}_\mu(\delta)$, and the effective dimension \mathbf{d}_{eff} :

$$\begin{aligned} \text{Opt}_\mu &:= \min_{\phi \in \mathbb{R}^{m \times \bar{L}}} \|\Delta - \mathbf{K}\phi^\top\|_{\text{op}} + \mu\|\phi\|_{\text{op}} \\ \text{Ovfit}_\mu(\delta) &:= \|G_\star\|_{\text{op}} \cdot \min \left\{ N^{1/2}, T^{1/2} \sqrt{\log \frac{1}{\delta} + \tilde{p} + \log \det(I + \mu^{-2}\mathbf{K}\mathbf{K}^\top)^{1/2}} \right\} \\ \mathbf{d}_{\text{eff}}(\text{Opt}, \bar{L}, \mu) &:= \tilde{p} + m + \text{lil} \frac{\text{Opt}}{\mu} + \bar{L} \log_+(\text{Opt} + \frac{\sqrt{N}\|\mathbf{K}\|_{\text{op}}}{\mu^2}). \end{aligned}$$

For our slightly refined bound, we shall use a parameter $\kappa > 0$ as in the bound Proposition 3.2) for a fixed filter, and define a filter-specific effective dimension

$$\widetilde{\mathbf{d}}_{\text{eff}}(\phi; L_{\text{eff}}, \mu, \kappa) := \tilde{p} + m + \text{lil} \frac{\|\Delta_\phi\|_{\text{op}}}{\kappa\sqrt{N}} + L_{\text{eff}} \log_+(\mu\|\phi\|_{\text{op}} + \frac{\mu^{-1}\|\mathbf{K}\|_{\text{op}}}{\kappa + N^{-1/2}\|\Delta_\phi\|_{\text{op}}}).$$

In this setting, our main result is as follows:

Theorem D.1 (General Statement of Theorem 3.3) *Let $\mu, \kappa > 0$ be fixed. Then*

- (a) *Let $\Phi \subset \mathbb{R}^{\bar{L} \times m}$ be a set of filters such that each slice $\Phi_v := \{\phi^\top v : \phi \in \Phi\}$ is contained in a subspace of dimension $L_{\text{eff}} \leq \bar{L}$. Then, with probability $1 - \delta$, the following holds on $\mathcal{E}_{\bar{\mathcal{U}}}$:*

$$\forall \phi \in \Phi, \quad \|G_\star - \widehat{G}_{\text{fil}}(\phi)\|_{\text{op}} \lesssim \frac{N^{-1/2}\|\Delta_\phi\|_{\text{op}} + \kappa}{\sqrt{N}} \cdot T^{1/2} \sqrt{\log \frac{1}{\delta} + \widetilde{\mathbf{d}}_{\text{eff}}(\phi; L_{\text{eff}}, \mu, \kappa)}.$$

- (b) *On $\mathcal{E}_{\bar{\mathcal{U}}}$, the filter $\phi_{\text{rdg}} \in \mathbb{R}^{\bar{L} \times m}$ satisfies $\max\{\mu\|\phi_{\text{rdg}}\|_{\text{op}}, \|\Delta_{\phi_{\text{rdg}}}\|_{\text{op}}\} \leq \text{Opt}_\mu + \text{Ovfit}_\mu(\delta)$. Thus, taking $\Phi = \mathbb{R}^{\bar{L} \times m}$, part (a) implies that on $\mathcal{E}_{\bar{\mathcal{U}}}$, with probability $1 - \delta$,*

$$\begin{aligned} \|G_\star - \widehat{G}_{\text{fil}}(\phi)\|_{\text{op}} &\lesssim \frac{N^{-1/2}(\text{Opt}_\mu + \text{Ovfit}_\mu(\delta)) + \kappa}{\sqrt{N}} \\ &\quad \cdot \sqrt{T \left(\log \frac{1}{\delta} + \widetilde{\mathbf{d}}_{\text{eff}}(\text{Opt}_\mu + \text{Ovfit}_\mu(\delta), \bar{L}, \mu, \kappa) \right)}. \end{aligned}$$

Theorem 3.3 is a direct consequence of combining parts (a) and (b) of the above theorem, setting $\kappa = \mu/\sqrt{N}$, and applying some routine simplifications to bound $\mathbf{d}_{\text{eff}}(\text{Opt}_\mu + \text{Ovfit}_\mu(\delta), \bar{L}, \mu, \kappa) \leq (\text{Opt}_\mu + \text{Ovfit}_\mu(\delta), \bar{L}, \mu)$.

D.1. Proof of Uniform Bound: Theorem D.1, Part (a)

Fix $\mu, \kappa > 0$. Recall that

$$\widehat{G}_{\text{fil}}(\phi) := \arg \min_{G \in \mathbb{R}^{T \times m}} \sum_{t=N_1}^N \|(\mathbf{y}_t - \phi \cdot \mathbf{k}_t) - G \bar{\mathbf{u}}_t\|_2,$$

and thus

$$\widehat{G}_{\text{fil}}(\phi) = G_\star + (\mathbf{Y} - \mathbf{K}\phi^\top) \bar{\mathbf{U}}^\top (\bar{\mathbf{U}}^\top \bar{\mathbf{U}})^{-1} = G_\star + \Delta_\phi \bar{\mathbf{U}}^\dagger.$$

Thus, $\|G_\star - \widehat{G}_{\text{fil}}(\phi)\|_{\text{op}} = \|\Delta_\phi \bar{\mathbf{U}}^\dagger\|_{\text{op}}$.

Step 0: Reduction To $m = 1$ and $\Phi = \mathbb{R}^{L_{\text{eff}}}$: We claim that prove the bound in the setting where $m = 1$ and $\Phi = \mathbb{R}^{L_{\text{eff}}}$. Indeed, given $\phi \in \Phi = \mathbb{R}^{m \times L}$ and a direction $v \in \mathcal{S}^{m-1}$, let $v^\top \phi \in \mathbb{R}^{1 \times L}$ denote the filter along v . Lastly, define

$$\widetilde{\mathbf{d}}_{\text{eff}}(\phi; \mu, \kappa, L_{\text{eff}}, v) := \widetilde{p} + m + \text{lil} \frac{\|v^\top \Delta_\phi\|_2}{\kappa \sqrt{N}} + L_{\text{eff}} \log_+(\mu \|v^\top \phi\|_2 + \frac{\mu^{-1} \|\mathbf{K}\|_{\text{op}}}{\kappa + N^{-1/2} \|v^\top \Delta_\phi\|_2}),$$

Suppose we can prove that for any fixed $v \in \mathcal{S}^{m-1}$, it holds with probability $1 - \delta$ on $\mathcal{E}_{\bar{\mathbf{U}}}$ that

$$\|v^\top \Delta_\phi \bar{\mathbf{U}}^\dagger\|_2 \lesssim \frac{\|v^\top \Delta_\phi\|_2 N^{-1/2} + \kappa}{\sqrt{N}} \cdot T^{1/2} \cdot \sqrt{\log \frac{1}{\delta} + \widetilde{\mathbf{d}}_{\text{eff}}(\phi; \mu, \kappa, v)}. \quad (\text{D.12})$$

Then, noting that $\|v^\top \Delta_\phi\|_2 \leq \|\Delta_\phi\|_{\text{op}}$ and $\widetilde{\mathbf{d}}_{\text{eff}}(\phi; \mu, \kappa, L_{\text{eff}}, v) \leq \widetilde{\mathbf{d}}_{\text{eff}}(\mu, \kappa, L_{\text{eff}})$, the theorem follows from a standard covering argument. Moreover, the bound (D.12) is equivalent to the setting where we observe $\mathbf{y}_t^\top v$ and make predictions $v^\top \phi \cdot \mathbf{k}_t$, where $\phi \in \Phi$, or equivalently, $(v^\top \phi) \in \Phi_v := \{v^\top \phi : \phi \in \Phi\}$. By assumption, there exists an L_{eff} -dimensional subspace of \mathbb{R}^L containing Φ_v . By projecting \mathbf{K} onto this subspace and applying a unitary change of coordinates, we may assume without loss of generality that $\Phi_v \subset \mathbb{R}^{L_{\text{eff}}}$. Finally, since we are proving a uniform bound, it is only stronger to prove the bound for all filter $(v^\top \phi) \in \mathbb{R}^{L_{\text{eff}}}$.

Thus, to simplify notation, we assume $m = 1$ and $\Phi \subset \mathbb{R}^{L_{\text{eff}}}$, and will let $\widetilde{\mathbf{d}}_{\text{eff}}(\dots)$ correspond to these one-dimensional predictions.

Step 1: Pointwise bound Let's start of with a pointwise bound for a fixed filter. By Proposition 3.2, we have with probability at least $1 - \delta$,

$$\|\Delta_\phi \bar{\mathbf{U}}^\dagger\|_2 \leq C \frac{(\|\Delta_\phi\|_2 N^{-1/2} + \kappa) T^{1/2}}{\sqrt{N}} \sqrt{\widetilde{p} + \log \frac{1}{\delta} + \text{lil}(\frac{\|\Delta_\phi\|_{\text{op}}}{\kappa N^{1/2}})}, \quad (\text{D.13})$$

for a sufficiently large constant C .

Step 2: Uniform bound over nets. For $j \geq 1$, let $c_j = e^{e^j}$. Let \mathcal{T}_j denote a $1/c_j$ net of the set $c_j \mathcal{B}_{\mathbb{R}^{L_{\text{eff}}}} / \mu$ in the norm $\mu \|\cdot\|_2$. We shall argue a uniform concentration bound over the nets \mathcal{T}_j , and then uses the nets to approximate bounds for each $\phi \in \mathbb{R}^{L_{\text{eff}}}$. Precisely, define the bound

$$\mathbb{B}(\|\Delta_\phi\|_2, j) := \frac{(N^{-1/2} \|\Delta_\phi\|_2 + \kappa) T^{1/2}}{\sqrt{N}} \sqrt{\widetilde{p} + L_{\text{eff}} \log(1 + 2c_j^2) + \log(\frac{2j^2}{\delta}) + \text{lil} \frac{\|\Delta_\phi\|_2}{\kappa \sqrt{N}}},$$

and, for C as in (D.13), define the event

$$\mathcal{E}_{\text{good}} := \{\forall j, \forall \phi \in \mathcal{T}_j : \|\Delta_\phi \bar{\mathbf{U}}^\dagger\|_2 \leq CB(\|\Delta_\phi\|_2, j)\}, \quad (\text{D.14})$$

By a standard volumetric argument (see e.g. Vershynin (2018, Corollary 4.2.13)), $\log |\mathcal{T}_j| \leq L_{\text{eff}} \log(1 + 2c_j^2)$, and therefore by applying (D.13) with $\delta_j := \frac{\delta}{2|\mathcal{T}_j|j^2}$ for each $\phi \in \mathcal{T}_j$ uniformly, we see that $\mathcal{E}_{\text{good}}$ holds with probability at least $1 - \frac{\pi^2}{12}\delta \geq 1 - \delta$. It remains to show that on $\mathcal{E}_{\text{good}}$, (D.12) holds as well.

Step 3: Uniform bound over $\phi \in \mathbb{R}^{L_{\text{eff}}}$. We shall now need a covering argument to translate (D.14) into a uniform guarantee over $\mathbb{R}^{L_{\text{eff}}}$: To this end, we introduce the shorthand $R_{\mathbf{k}} := \|\mathbf{K}/\mu\|_{\text{op}}$, establish the following covering claim:

Claim D.2 *For any $\phi \in \mathbb{R}^{L_{\text{eff}}}$ with $\mu\|\phi\|_2 \leq c_j$, there exists $\tilde{\phi} \in \mathcal{T}_j$ satisfying $\|\Delta_\phi - \Delta_{\tilde{\phi}}\|_2 \leq R_{\mathbf{k}}/c_j$.*

Proof By definition, $\|\phi\|_2 \in \frac{c_j}{\mu} \mathcal{B}_{\mathbb{R}^{L_{\text{eff}}}}$. Since \mathcal{T}_j is a $1/c_j$ net of $\frac{c_j}{\mu} \mathcal{B}_{\mathbb{R}^{L_{\text{eff}}}}$ in the norm $\frac{1}{\mu} \|\cdot\|_2$, there exists a $\tilde{\phi} \in \mathcal{T}_j$ with $\|\phi - \tilde{\phi}\|_2 \leq \frac{1}{c_j\mu}$. Thus, $\|\Delta_\phi - \Delta_{\tilde{\phi}}\|_2 = \|\mathbf{K}(\phi - \tilde{\phi})\|_2 \leq \|\mathbf{K}\|_{\text{op}} \frac{1}{c_j\mu} = R_{\mathbf{k}}/c_j$. ■

As a consequence, we obtain a uniform bound as follows. Consider any ϕ , and any j for which $\mu\|\phi\|_2 \leq c_j$. For this j , let $\tilde{\phi}$ denote the filter guaranteed by the above claim. Then on $\mathcal{E}_{\text{good}}$,

$$\begin{aligned} \|\Delta_\phi \bar{\mathbf{U}}^\dagger\|_2 &\leq \|\Delta_{\tilde{\phi}} \bar{\mathbf{U}}^\dagger\|_2 + \frac{R_{\mathbf{k}}}{c_j \sigma_{\min}(\bar{\mathbf{U}})} \leq CB(\|\Delta_{\tilde{\phi}}\|_2, j) + \frac{R_{\mathbf{k}}}{c_j \sigma_{\min}(\bar{\mathbf{U}})} \\ &\leq CB(\|\Delta_\phi\|_2 + \frac{R_{\mathbf{k}}}{c_j}, j) + \frac{R_{\mathbf{k}}}{c_j \sigma_{\min}(\bar{\mathbf{U}})}. \end{aligned}$$

Thus we conclude that on $\mathcal{E}_{\text{good}} \cap \mathcal{E}_{\bar{\mathbf{U}}}$, the following holds simultaneously for all $\phi \in \Phi$

$$\begin{aligned} \|\Delta_\phi \bar{\mathbf{U}}^\dagger\|_2 &\lesssim \inf_{j \in \mathbb{N}: \mu\|\phi\|_2 \leq c_j} B\left(\|\Delta_\phi\|_2 + \frac{R_{\mathbf{k}}}{c_j}, j\right) + \frac{R_{\mathbf{k}}}{c_j \sqrt{N}} \\ &\lesssim \inf_{j \in \mathbb{N}: \mu\|\phi\|_2 \leq c_j} B\left(\|\Delta_\phi\|_2 + \frac{\sqrt{N} R_{\mathbf{k}}}{c_j}, j\right), \end{aligned} \quad (\text{D.15})$$

where we note that $B(A + B, j) + B/\sqrt{N} \leq B(A + B + \sqrt{N}B, j) \lesssim B(A + B\sqrt{N}, j)$.

Step 4: Tuning the bound. Lastly, we tune the bound in (D.15) to each specific $\phi \in \Phi$. Note that the infimum over j in (D.15) is taken for each $\phi \in \Phi$, and thus we can tailor j to our filter ϕ . The remainder of the proof is simply choosing j appropriately. Define the *random index*

$$j_0 := \inf \left\{ j \in \mathbb{N} : \max\left\{ \mu\|\phi\|_2, \frac{R_{\mathbf{k}}}{(\kappa + N^{-1/2})\|\Delta_\phi\|_2} \right\} \leq c_j \right\}.$$

Then $\|\Delta_\phi \bar{\mathbf{U}}^\dagger\|_2$ is bounded by (up to constants) by

$$\begin{aligned}
 & \mathsf{B}\left(\|\Delta_\phi\|_2 + \frac{R_{\mathbf{k}}\sqrt{N}}{c_{j_0}}, \mathbf{j}_0\right) \\
 & \stackrel{(i)}{\lesssim} \frac{(N^{-1/2}\|\Delta_\phi\|_2 + \frac{R_{\mathbf{k}}}{c_{j_0}} + \kappa)}{\sqrt{N}} \cdot \sqrt{T} \cdot \sqrt{\tilde{p} + e^{j_0} L_{\text{eff}} + \log \frac{1}{\delta} + \text{lil}\left(\frac{N^{-1/2}\|\Delta_\phi\|_2 + R_{\mathbf{k}}/c_{j_0}}{\kappa}\right)} \\
 & \stackrel{(ii)}{\lesssim} \frac{N^{-1/2}\|\Delta_\phi\|_2 + \kappa}{\sqrt{N}} \cdot \sqrt{T} \cdot \sqrt{\tilde{p} + e^{j_0} L_{\text{eff}} + \log \frac{1}{\delta} + \text{lil}\left(\frac{2N^{-1/2}\|\Delta_\phi\|_2}{\kappa} + 1\right)} \\
 & \stackrel{(iii)}{\lesssim} \frac{N^{-1/2}\|\Delta_\phi\|_2 + \kappa}{\sqrt{N}} \cdot \sqrt{T} \cdot \sqrt{\tilde{p} + e^{j_0} L_{\text{eff}} + \log \frac{1}{\delta} + \text{lil}\left(\frac{\|\Delta_\phi\|_2}{\kappa N^{1/2}}\right)},
 \end{aligned}$$

where (i) follows by combining the definition of B in (D.14) with the chain of inequalities

$$\begin{aligned}
 1 + L_{\text{eff}} \log(1 + 2c_{j_0}^2) + \log\left(\frac{2j_0^2}{\delta}\right) & \lesssim L_{\text{eff}} \log(c_{j_0}) + \log(j_0) + \log(1/\delta) \\
 & \lesssim L_{\text{eff}} e^{j_0} + \log(1/\delta),
 \end{aligned}$$

which uses the facts that $L_{\text{eff}} \geq 1$ and $c_{j_0} = e^{j_0} \geq j_0$. Furthermore, (ii) uses $\frac{R_{\mathbf{k}}}{c_{j_0}} \leq \sqrt{N}\|\Delta_\phi\|_2 + \kappa$ and (iii) uses $\text{lil}(2x + 1) \leq \text{lil}(x) + 1$. Lastly, by the definition of j_0 , we see that

$$e^{j_0} \lesssim \log_+(\mu\|\phi\|_2 + \frac{R_{\mathbf{k}}}{(\kappa + N^{-1/2}\|\Delta_\phi\|_2)}) = \log_+(\mu\|\phi\|_2 + \frac{\|\mathbf{K}\|_{\text{op}}}{\mu(\kappa + N^{-1/2}\|\Delta_\phi\|_2)}),$$

which concludes the proof.

D.2. Proof of $\widetilde{\text{Opt}}_\mu$ bound, Theorem D.1, Part (b)

Let $(\mathbf{k}_t) \subset \mathbb{R}^{\bar{L}}$ be a sequence. For sequences $(\mathbf{a}_t) \subset \mathbb{R}^m$ and regularizer $\mu > 0$, and define the ridge function:

$$\text{ridge}(\mathbf{A}, \mu) := \arg \min_{\phi \in \mathbb{R}^{m \times \bar{L}}} \sum_{t=N_1}^N \|\phi \cdot \mathbf{k}_t - \mathbf{a}_t\|_2^2 + \mu^2 \|\phi\|_{\text{F}}^2,$$

where $\mathbf{A} \in \mathbb{R}^{\tilde{N} \times m}$ is the matrix whose rows are (\mathbf{a}_t) . Note then that

$$\phi_{\text{rdg}} := \text{ridge}(\mathbf{Y}, \mu).$$

To prove the desired bound, we begin by stating some properties of ridge, which are derived from the KKT conditions in Section D.2.1:

Lemma D.3 (Properties of the Ridge Estimator)

(a) $\text{ridge}(\mathbf{A}, \mu)^\top = (\mathbf{K}^\top \mathbf{K} + \mu^2 I)^{-1} \mathbf{K}^\top \mathbf{A}$. In particular, ridge is linear in its first argument.

(b) $\text{ridge}(\mathbf{A}, \mu) \in \arg \min_{\phi \in \mathbb{R}^{m \times \bar{L}}} \|\text{vcat}[\mathbf{K}\phi^\top - \mathbf{A}, \mu\phi^\top]\|_{\text{op}}$, where $\text{vcat}[A, B] := [A^\top | B^\top]^\top$.

The second point of Lemma D.3 implies that $\text{ridge}(\mathbf{A}, \mu)$ is nearly optimal in the operator norm, up to the regularization. In particular, recall

$$\text{Opt}_\mu := \min_{\phi \in \mathbb{R}^{m \times \bar{L}}} \|\Delta_\phi\|_{\text{op}} + \mu \|\phi\|_{\text{op}}.$$

Then it follows from Lemma D.3 that

$$\begin{aligned} \max\{\|\Delta_{\text{ridge}(\Delta, \mu)}\|_{\text{op}}, \mu \|\text{ridge}(\Delta, \mu)\|_{\text{op}}\} &\leq \min_{\phi \in \mathbb{R}^{m \times \bar{L}}} \|\text{vcat} [\mathbf{K}\phi^\top - \Delta, \mu\phi^\top]\|_{\text{op}} \\ &= \min_{\phi \in \mathbb{R}^{m \times \bar{L}}} \|\text{vcat} [\Delta_\phi, \mu\phi^\top]\|_{\text{op}} \leq \text{Opt}_\mu. \end{aligned} \quad (\text{D.16})$$

Unfortunately, our algorithm does not have direct access to Δ , only \mathbf{Y} . The following lemma accounts for this discrepancy.

Lemma D.4 *We have that*

$$\begin{aligned} (a) \quad &\|\Delta_{\phi_{\text{rdg}}}\|_{\text{op}} \leq \text{Opt}_\mu + \|\mathbf{K} \text{ridge}(\bar{\mathbf{U}}G_\star^\top, \mu)^\top\|_{\text{op}} \\ (b) \quad &\mu \|\phi_{\text{rdg}}\|_{\text{op}} \leq \text{Opt}_\mu + \mu \|\text{ridge}(\bar{\mathbf{U}}G_\star^\top, \mu)\|_{\text{op}}. \end{aligned}$$

Lastly, we require a bound on $\|\mathbf{K} \text{ridge}(\bar{\mathbf{U}}G_\star^\top, \mu)^\top\|_{\text{op}}$ and $\mu \|\text{ridge}(\bar{\mathbf{U}}G_\star^\top, \mu)\|_{\text{op}}$. This is accomplished by the following lemma:

Lemma D.5 *With probability at least $1 - \delta$ on $\mathcal{E}_{\bar{\mathbf{U}}}$, $\|\mathbf{K} \text{ridge}(\bar{\mathbf{U}}G_\star^\top, \mu)^\top\|_{\text{op}}$ and $\mu \|\text{ridge}(\bar{\mathbf{U}}G_\star^\top, \mu)\|_{\text{op}}$ are both bounded as $\lesssim \text{Ovfit}_\mu(\delta)$.*

Theorem D.1, Part (b) is now an immediate consequence of the two lemmas above.

D.2.1. PROOF OF LEMMA D.3

The first statement is standard. For the second statement, observe that $\text{ridge}(\mathbf{A}v, \mu) = \text{ridge}(\mathbf{A}, \mu)v$. Hence, we have

$$\begin{aligned} \min_{\phi \in \mathbb{R}^{m \times \bar{L}}} \|\text{vcat} [\mathbf{K}\phi^\top - \mathbf{A}, \mu\phi^\top] v\|_2^2 &= \min_{\phi} \|\mathbf{K}(\phi^\top v) - \mathbf{A}v\|_2^2 + \|\mu\phi^\top v\|_2^2 \\ &= \min_{w \in \mathbb{R}^{\bar{L}}} \|\mathbf{K}w - \mathbf{A}v\|_2^2 + \|\mu w\|_2^2 \\ &= \|\mathbf{K} \text{ridge}(\mathbf{A}v, \mu) - \mathbf{A}v\|_2^2 + \|\mu \text{ridge}(\mathbf{A}v, \mu)\|_2^2 \\ &= \|\mathbf{K} \text{ridge}(\mathbf{A}, \mu)v - \mathbf{A}v\|_2^2 + \|\mu \text{ridge}(\mathbf{A}, \mu)v\|_2^2 \\ &= \|\text{vcat} [\mathbf{K} \text{ridge}(\mathbf{A}, \mu) - \mathbf{A}, \mu \text{ridge}(\mathbf{A}, \mu)] v\|_2^2. \end{aligned} \quad (\text{D.17})$$

To conclude, we see that

$$\begin{aligned}
 \|\text{vcat} [\mathbf{K}\phi_{\text{rdg}}^\top - \mathbf{Y}, \mu\phi_{\text{rdg}}^\top]\|_{\text{op}} &= \|\text{vcat} [\mathbf{K}\text{ridge}(\mathbf{Y}, \mu) - \mathbf{Y}, \mu\text{ridge}(\mathbf{Y}, \mu)]\|_{\text{op}} \\
 &= \max_{v \in \mathcal{S}^{m-1}} \|\text{vcat} [\mathbf{K}\text{ridge}(\mathbf{Y}, \mu) - \mathbf{Y}, \mu\text{ridge}(\mathbf{Y}, \mu)] v\|_2 \\
 &\stackrel{(i)}{=} \max_{v \in \mathcal{S}^{m-1}} \min_{\phi} \|\text{vcat} [\mathbf{K}\phi^\top - \mathbf{Y}, \mu\phi^\top] v\|_2 \\
 &\stackrel{(ii)}{\leq} \min_{\phi} \max_{v \in \mathcal{S}^{m-1}} \|\text{vcat} [\mathbf{K}\phi^\top - \mathbf{Y}, \mu\phi^\top] v\|_2 \\
 &= \min_{\phi} \|\text{vcat} [\mathbf{K}\phi^\top - \mathbf{Y}, \mu\phi^\top]\|_{\text{op}},
 \end{aligned}$$

where (i) uses (D.17) and (ii) uses weak duality. Since, by definition,

$$\min_{\phi} \|\text{vcat} [\mathbf{K}\phi^\top - \mathbf{Y}, \mu\phi^\top]\|_{\text{op}} \leq \|\text{vcat} [\mathbf{K}\phi_{\text{rdg}}^\top - \mathbf{Y}, \mu\phi_{\text{rdg}}^\top]\|_{\text{op}},$$

they must be equal.

D.2.2. PROOF OF LEMMA D.4

Proof of (a): By Lemma D.3, ridge is linear in its first argument. Therefore,

$$\phi_{\text{rdg}} = \text{ridge}(\mathbf{Y}, \mu) = \text{ridge}(\mathbf{\Delta} + \overline{\mathbf{U}}G_\star^\top, \mu) = \text{ridge}(\mathbf{\Delta}, \mu) + \text{ridge}(\overline{\mathbf{U}}G_\star^\top, \mu).$$

Therefore,

$$\begin{aligned}
 \|\mathbf{\Delta}_{\phi_{\text{rdg}}}\|_{\text{op}} &= \|\mathbf{\Delta} - \mathbf{K}\phi_{\text{rdg}}^\top\|_{\text{op}} \\
 &= \|\mathbf{\Delta} - \mathbf{K}(\text{ridge}(\mathbf{\Delta}, \mu) + \text{ridge}(\overline{\mathbf{U}}G_\star^\top, \mu))\|_{\text{op}} \\
 &\leq \|\mathbf{\Delta} - \mathbf{K} \cdot \text{ridge}(\mathbf{\Delta}, \mu)\|_{\text{op}} + \|\mathbf{K}\text{ridge}(\overline{\mathbf{U}}G_\star^\top, \mu)\|_{\text{op}} \\
 &= \|\mathbf{\Delta}_{\text{ridge}(\mathbf{\Delta}, \mu)}\|_{\text{op}} + \|\mathbf{K}\text{ridge}(\overline{\mathbf{U}}G_\star^\top, \mu)\|_{\text{op}} \leq \text{Opt}_\mu + \|\mathbf{K}\text{ridge}(\overline{\mathbf{U}}G_\star^\top, \mu)\|_{\text{op}},
 \end{aligned}$$

by (D.16).

Proof of (b): Again, using the linearity of ridge and (D.16), we see that

$$\begin{aligned}
 \mu \|\phi_{\text{rdg}}\|_{\text{op}} &= \mu \|\text{ridge}(\mathbf{\Delta}, \mu) + \text{ridge}(\overline{\mathbf{U}}G_\star^\top, \mu)\|_{\text{op}} \\
 &\leq \mu \|\text{ridge}(\mathbf{\Delta}, \mu)\|_{\text{op}} + \mu \|\text{ridge}(\overline{\mathbf{U}}G_\star^\top, \mu)\|_{\text{op}} \\
 &\leq \text{Opt}_\mu + \mu \|\text{ridge}(\overline{\mathbf{U}}G_\star^\top, \mu)\|_{\text{op}}.
 \end{aligned}$$

D.2.3. PROOF OF LEMMA D.5

For $M \in \{\mathbf{K}, \mu\}$, we have

$$\begin{aligned}
 \|M\text{ridge}(\overline{\mathbf{U}}G_\star^\top, \mu)^\top\|_{\text{op}} &= \|M(\mathbf{K}^\top \mathbf{K} + \mu^2)^{-1} \mathbf{K}^\top \overline{\mathbf{U}}G_\star^\top\|_{\text{op}} \\
 &\leq \|M(\mathbf{K}^\top \mathbf{K} + \mu^2)^{-1/2}\|_{\text{op}} \cdot \|(\mathbf{K}^\top \mathbf{K} + \mu^2)^{-1/2} \mathbf{K}^\top \overline{\mathbf{U}}\|_{\text{op}} \|G_\star\|_{\text{op}} \\
 &\stackrel{(i)}{\leq} \|(\mathbf{K}^\top \mathbf{K} + \mu^2)^{-1/2} \mathbf{K}^\top\|_{\text{op}} \|\overline{\mathbf{U}}\|_{\text{op}} \|G_\star\|_{\text{op}} \\
 &\stackrel{(ii)}{\leq} \|\overline{\mathbf{U}}\|_{\text{op}} \|G_\star\|_{\text{op}}.
 \end{aligned} \tag{D.18}$$

where (i) and (ii) use the fact that $\|\mathbf{K}(\mathbf{K}^\top \mathbf{K} + \mu^2)^{-1/2}\|_{\text{op}}$, $\|\mu(\mathbf{K}^\top \mathbf{K} + \mu^2)^{-1/2}\|_{\text{op}}$, and $\|(\mathbf{K}^\top \mathbf{K} + \mu^2)^{-1/2} \mathbf{K}^\top\|_{\text{op}}$ are all upper bounded by 1. For example,

$$\|\mathbf{K}(\mathbf{K}^\top \mathbf{K} + \mu^2)^{-1/2}\|_{\text{op}} \leq \|\mathbf{K}\|_{\text{op}} \|(\mathbf{K}^\top \mathbf{K} + \mu^2)^{-1/2}\|_{\text{op}} \leq \frac{\|\mathbf{K}\|_{\text{op}}}{\sqrt{\sigma_{\min}(\mathbf{K}^\top \mathbf{K} + \mu^2)}} \leq 1.$$

Thus, $\|M_{\text{ridge}}(\bar{\mathbf{U}}G_\star^\top, \mu)\|_{\text{op}} \leq \|\bar{\mathbf{U}}\|_{\text{op}} \|G_\star\|_{\text{op}}$ which is $\lesssim \sqrt{N} \|G_\star\|_{\text{op}}$ on $\mathcal{E}_{\bar{\mathbf{U}}}$.

For the second bound, starting from (D.18), we use the semi-parametric regression bound, part (a) of Theorem E.1. Since \mathbf{k}_t is \mathcal{F}_{t-T} -adapted, by applying part (a) of Theorem E.1 with $V_0 \leftarrow \mu^2 I$ and $\Delta \leftarrow \mathbf{K}$, we see that with probability at least $1 - \delta$,

$$\|(\mathbf{K}^\top \mathbf{K} + \mu^2)^{-1/2} \mathbf{K}^\top \bar{\mathbf{U}}\|_{\text{op}} \lesssim T^{1/2} \sqrt{\tilde{p} + \log \frac{1}{\delta} + \log \det(I + \mu^{-2} \mathbf{K} \mathbf{K}^\top)^{1/2}}.$$

Concluding we have that, on $\mathcal{E}_{\bar{\mathbf{U}}}$, with probability at least $1 - \delta$,

$$\begin{aligned} \|M_{\text{ridge}}(\bar{\mathbf{U}}G_\star^\top, \mu)\|_{\text{op}} &\lesssim \|G_\star\|_{\text{op}} \min\{T^{1/2} \sqrt{\tilde{p} + \log \frac{1}{\delta} + \log \det(I + \mu^{-2} \mathbf{K} \mathbf{K}^\top)^{1/2}}, \sqrt{N}\} \\ &= \text{Ovfit}_\mu(\delta). \end{aligned}$$

Appendix E. Semi-Parametric Regression

In this section we prove Proposition 3.2 as consequence of a more general setting. Our main theorem is as follows:

Theorem E.1 *Suppose the semi-parametric model (1.1) holds, where $\mathbf{u}_t | \mathcal{F}_{t-1}$ is mean-zero and 1 subgaussian, and let $\tilde{p} := p \min\{T, \log^2(eTp) \log^2(Tp)\}$. Then,*

(a) *For any fixed $\delta \in (0, 1)$ and $V_0 \in \mathbb{S}_{++}^m$, it holds with probability $1 - \delta$ that*

$$\|(\Delta^\top \Delta + V_0)^{-1/2} \Delta^\top \bar{\mathbf{U}}\|_{\text{op}} \lesssim T^{1/2} \sqrt{\tilde{p} + \log \frac{1}{\delta} + \log \det(I + \Delta V_0^{-1} \Delta^\top)^{1/2}}$$

(b) *For any fixed $\delta \in (0, 1)$ and $\kappa > 0$, it holds with probability $1 - \delta$ that*

$$\sigma_{\min}(\bar{\mathbf{U}})^2 \|\bar{\mathbf{U}}^\dagger \Delta\|_{\text{op}} \leq \|\Delta^\top \bar{\mathbf{U}}\|_{\text{op}} \lesssim T^{1/2} (\|\Delta\|_{\text{op}} + \kappa) \sqrt{\tilde{p} + m + \log \frac{1}{\delta} + \text{lil}(\frac{\|\Delta\|_{\text{op}}}{\kappa})},$$

where $\text{lil}(x) := \log(1 + \log(1 + x))$.

Proposition 3.2 corresponds exactly to part (b) of the above theorem, with the substitution $\delta_t \leftarrow \delta_{\phi,t}$. Indeed, we have that (a)

$$\hat{G}_{\text{fil}}(\phi) = (\mathbf{Y} - \mathbf{K}\phi^\top)^\top \bar{\mathbf{U}} (\bar{\mathbf{U}}^\top \bar{\mathbf{U}})^{-1} = G_\star + \Delta_\phi^\top \bar{\mathbf{U}}^\dagger,$$

(b) that the inputs $\mathbf{u}_{t+1} | \mathcal{F}_t \sim \mathcal{N}(0, I_p)$ are 1-subgaussian and $\{\mathcal{F}_t\}$ -adapted, and where we substitute the errors $\delta_{\phi,t}$ are \mathcal{F}_{t-T} -adapted. Hence, whenever $\mathcal{E}_{\bar{\mathbf{U}}}$ occurs, applying Theorem E.1, part (b) with $\kappa \leftarrow \kappa \sqrt{N}$ and $\delta_t \leftarrow \delta_{\phi,t}$ yields

$$N \|\bar{\mathbf{U}}^\dagger \Delta_\phi\|_{\text{op}} \lesssim \sigma_{\min}(\bar{\mathbf{U}})^2 \|\Delta_\phi \bar{\mathbf{U}}^\dagger\|_{\text{op}} \lesssim T^{1/2} (\|\Delta_\phi\|_{\text{op}} + \kappa \sqrt{N}) \sqrt{\tilde{p} + m + \log \frac{1}{\delta} + \text{lil}(\frac{\|\Delta_\phi\|_{\text{op}}}{\sqrt{N}\kappa})},$$

which coincides with the statement of Proposition 3.2 after dividing both sides by N .

For comparison, observe that if one instead regresses (\mathbf{y}_t) to independent white noise $\tilde{u}_t \stackrel{\text{iid}}{\sim} \mathcal{N}(0, I_{Tp}) \in \mathbb{R}^{Tp}$, we would have rates $\|\widehat{G}_{\text{LS}} - G_\star\|_{\text{op}} = \Theta(\sqrt{\frac{1}{N}} \cdot \sqrt{(Tp + m) + \log \frac{1}{\delta}})$. In our setting, martingale tail bounds incurs an addition log-factor (part (a) of Theorem E.1), which we refine to $\text{lil}(\cdot)$ for operator norm error using Lemma E.6, below. In addition, correlation introduced by the concatenated sequence forces us to pay an additional factor of T multiplying the m and $\log \frac{1}{\delta}$ terms as well, which we conjecture is in the worst case. If handled naively, we would also have to pay for $T \cdot (Tp) = pT^2$, and the major effort in the proof of Theorem E.1 is to use a careful chaining argument based on (Krahmer et al., 2014) to ensure that we instead pay for the generally smaller term $T\tilde{p}$. This ensures that our bound is only suboptimal (up to more than log factors) once $m + \log \frac{1}{\delta} \gg p$.

We present the proof of part (a) in E.1, and derive (b) as a consequence in E.2. Part (a) relies on detailed chaining arguments, which we defer to Appendix F.

E.1. Proof of Theorem E.1, Part (a)

We shall also use the notation

$$\tilde{N} = N - N_1 \quad \text{and} \quad N_0 = N_1 - T.$$

Adopting the setting of Theorem E.1, $\{\mathcal{F}_t\}$ be a filtration, and suppose that $(\mathbf{u}_t) \subset \mathbb{R}^p$ is a $\{\mathcal{F}_t\}$ -adapted 1-subgaussian and $(\boldsymbol{\delta}_t) \subset \mathbb{R}^m$ is a sequence of $\{\mathcal{F}_{t-T}\}$ adapted sequence.

We shall begin with the proof of part (a); the proof of part (b) will be derived as a consequence in E.2. Begin by fixing $V_0 \in \mathbb{S}_{++}^m$. For vectors $\bar{v} = [v_1^\top, \dots, v_T^\top]^\top \in \mathbb{R}^{Tp}$, we define the process

$$\mathbf{z}_{\bar{v}} := \|\boldsymbol{\Delta}^\top \bar{\mathbf{U}} \bar{v}\|_{\boldsymbol{\Delta}^\top \boldsymbol{\Delta} + V_0}.$$

We see that

$$\|(\boldsymbol{\Delta}^\top \boldsymbol{\Delta} + V_0)^{-1/2} \boldsymbol{\Delta}^\top \bar{\mathbf{U}}\|_{\text{op}} = \sup_{\bar{v} \in \mathcal{S}^{Tp-1}} \mathbf{z}_{\bar{v}}.$$

Hence, it suffices to show that, with probability $1 - \delta$,

$$\sup_{\bar{v} \in \mathcal{S}^{Tp-1}} \mathbf{z}_{\bar{v}} \lesssim T^{1/2} \sqrt{\log(1/\delta) + \log \det(I + (\boldsymbol{\Delta}^\top \boldsymbol{\Delta} V_0^{-1})^{1/2} + \tilde{p})}$$

The first step is to express $\bar{\mathbf{U}} \bar{v} = \text{Toep}_T(\bar{v}) \vec{\mathbf{u}}$, where $\vec{\mathbf{u}} \in \mathbb{R}^{(N-N_0+1)p}$ is the vector obtained by concatenating $\mathbf{u}_{N_0}, \dots, \mathbf{u}_N$ and $\text{Toep}_T(\bar{v}) \in \mathbb{R}^{\tilde{N} \times (N-N_0+1)p}$ is an $[\tilde{N} \times (N - N_0 + 1)]$ block Toeplitz matrix, where each of the size $(1 \times p)$ blocks on the k -th superdiagonal are equal to v_{T-k}^\top for $0 \leq k \leq T - 1$, i.e.

$$\text{Toep}_T(\bar{v}) := \begin{bmatrix} v_T^\top & \cdots & v_1^\top & & \\ & \ddots & & \ddots & \\ & & v_T^\top & \cdots & v_1^\top \end{bmatrix}.$$

We can then express $\Delta^\top \text{Toep}_T(\bar{v}) \in \mathbb{R}^{m \times ((N-N_0+1)p)}$ with blocks $A_{t-1}(\bar{v})^\top \in \mathbb{R}^{m \times p}$ for $t \in \{N_0, \dots, N\}$, where

$$A_{t-1}(\bar{v})^\top := \sum_{j=0}^{T-1} (\delta_{t+j} \mathbb{I}_{t+j \in [N_1, N]}) v_{T-j}^\top.$$

Since $A_{t-1}(\bar{v})^\top$ depends only on terms δ_s for $s \leq t + T - 1$, we see that $A_{t-1}(\bar{v})$ is \mathcal{F}_{t-1} -adapted. Then we see that

$$\Delta^\top \bar{\mathbf{U}} \bar{v} = \Delta^\top \text{Toep}_T(\bar{v}) \vec{\mathbf{u}} = \sum_{t=N_0}^N A_{t-1}(\bar{v})^\top \mathbf{u}_t.$$

We now introduce the variance process:

$$\mathbb{V}_\Delta(\bar{v}) := \sum_{t=N_0}^N A_{t-1}(\bar{v})^\top A_{t-1}(\bar{v}) = \Delta^\top \text{Toep}_T(\bar{v})^\top \text{Toep}_T(\bar{v}) \Delta,$$

which corresponds to a variance proxy for $\Delta^\top \bar{\mathbf{U}} \bar{v}$. Indeed, consider the simple case where \mathbf{u}_t are deterministic and $\delta_t \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$ are independent Gaussian. Then, then the covariance matrix

$$\mathbb{E}[(\Delta^\top \bar{\mathbf{U}} \bar{v})(\Delta^\top \bar{\mathbf{U}} \bar{v})^\top]$$

would be equal to $\mathbb{V}_\Delta(v)$. In fact, if \mathbf{u}_t is any martingale sequence with $\mathbb{E}[\mathbf{u}_t \mathbf{u}_t^\top] \preceq I$, and \mathbf{u}_t were deterministic, then it still holds $\mathbb{E}[(\Delta^\top \bar{\mathbf{U}} \bar{v})(\Delta^\top \bar{\mathbf{U}} \bar{v})^\top] \preceq \mathbb{V}_\Delta(v)$.

In our general case where both δ_t and δ_t are martingales, and δ_t is subgaussian, we can use $\mathbb{V}_\Delta(v)$ as a *data-dependent* subgaussian variance proxy. Recall that $\delta_t | \mathcal{F}_{t-1}$ is subgaussian and \mathcal{F}_t measurable. Crucially, we will also use that for each \bar{v} , $\{A_t(\bar{v})\}$ is $\{\mathcal{F}_t\}$ -adapted, since $A_t(\bar{v})$ involves the terms only the terms $\delta_{t+1}, \dots, \delta_{t+T}$, and $\{\delta_s\}$ is $\{\mathcal{F}_{s-T}\}$ -adapted. These two points let us invoke the following lemma, which generalizes a bound due to [Abbasi-Adkori \(2011\)](#):

Lemma E.2 (Generalization of Theorem 3 in [Abbasi-Adkori \(2011\)](#)) *Let \mathcal{F}_t denote an arbitrary filtration. Let $\{A_t\} \subset \mathbb{R}^{p \times m}$ and $\{\mathbf{u}_t\} \subset \mathbb{R}^m$ be \mathcal{F}_t adapted, and suppose further that $\mathbf{u}_t | \mathcal{F}_{t-1}$ is mean zero and 1-subgaussian. Define the variance process $\mathbf{V}_k := \sum_{t=1}^k A_{t-1}^\top A_{t-1}$. Then, for any $\{\mathcal{G}_t\}$ -adapted stopping time τ , one has*

$$\left\| \sum_{t=1}^{\tau} A_{t-1}^\top \mathbf{u}_t \right\|_{(\mathbf{V}_\tau + V_0)^{-1}}^2 \leq 2 \log \left(\frac{\det(I + \mathbf{V}_\tau V_0^{-1})^{1/2}}{\delta} \right) \quad \text{w.p. } 1 - \delta.$$

To prove part (a) of the proposition, we shall set $\tau = N$, and $A_{t-1} = A_{t-1}(\bar{v})$ for $t \geq N_0$, and $A_{t-1} = 0$ for $t \leq N_0$. For this choice of A , then $\sum_{t=1}^N A_{t-1}^\top \mathbf{u}_t = \Delta^\top \bar{\mathbf{U}} \bar{v}$ and $\mathbf{V}_\tau = \mathbf{V}_N = \mathbb{V}_\Delta(\bar{v})$. Thus, we have that for any $V_0 \succeq 0$,

$$\left\| \Delta^\top \bar{\mathbf{U}} \bar{v} \right\|_{(\mathbb{V}_\Delta(v) + V_0)^{-1}}^2 \leq 2 \log \left(\frac{\det(I + \mathbb{V}_\Delta(v) V_0^{-1})^{1/2}}{\delta} \right) \quad \text{w.p. } 1 - \delta. \quad (\text{E.19})$$

To use (E.19) to bound $\sup_{\bar{v} \in \mathcal{S}^{Tp-1}} \mathbf{z}_{\bar{v}}$, we shall show that the random variables $\mathbf{z}_{\bar{v}}$ behave like a subgaussian process, with a random offset $\boldsymbol{\mu}$. This will allow us to apply a chaining argument to bound their supremum. To this end, introduce $\|\bar{v}\|_{\ominus} := \|\text{Toep}_T(\bar{v})\|_{\text{op}}$; it is straightforward to check that $\|\cdot\|_{\ominus}$ defines a norm on \mathbb{R}^{Tp} . We shall need the following bound, proved in Section E.3:

Lemma E.3 (Bounds on $\|\cdot\|_{\ominus}$) $\|\bar{v}\|_2 \leq \|\bar{v}\|_{\ominus} \leq \sqrt{T}\|\bar{v}\|_2$.

By definition, $\mathbb{V}_{\Delta}(\bar{v}) = \Delta^{\top} \text{Toep}_T(\bar{v}) \text{Toep}_T(\bar{v})^{\top} \Delta$, and thus the above lemma implies

$$\mathbb{V}_{\Delta}(\bar{v}) \preceq \|\text{Toep}_T(\bar{v})\|_{\text{op}}^2 \Delta^{\top} \Delta = \|\bar{v}\|_{\ominus}^2 \Delta^{\top} \Delta.$$

Thus, we see

$$\begin{aligned} \mathbf{z}_{\bar{v}} &= \|\Delta^{\top} \bar{\mathbf{U}} \bar{v}\|_{(\Delta^{\top} \Delta + V_0)^{-1}} = \|\bar{v}\|_{\ominus} \|\Delta^{\top} \bar{\mathbf{U}} \bar{v}\|_{(\|\bar{v}\|_{\ominus}^2 \Delta^{\top} \Delta + \|\bar{v}\|_{\ominus}^2 V_0)^{-1}} \\ &\leq \|\bar{v}\|_{\ominus} \|\Delta^{\top} \bar{\mathbf{U}} \bar{v}\|_{(\mathbb{V}_{\Delta}(\bar{v}) + \|\bar{v}\|_{\ominus}^2 V_0)^{-1}}. \end{aligned}$$

Thus, by (E.19), we have that with probability $1 - \delta$,

$$\begin{aligned} \|\Delta^{\top} \bar{\mathbf{U}} \bar{v}\|_{(\mathbb{V}_{\Delta}(\bar{v}) + \|\bar{v}\|_{\ominus}^2 V_0)^{-1}} &\leq \sqrt{2 \log \left(\frac{\det(I + \mathbb{V}_{\Delta}(\bar{v})(\|\bar{v}\|_{\ominus}^2 V_0^{-1})^{1/2})}{\delta} \right)} \\ &\stackrel{(i)}{=} \sqrt{2 \log \left(\frac{\det(I + \Delta^{\top} \Delta V_0^{-1})^{1/2}}{\delta} \right)} \\ &\leq \sqrt{2 \log \det(I + \Delta^{\top} \Delta V_0^{-1})^{1/2}} + \sqrt{2 \log(1/\delta)}, \end{aligned}$$

where in (i) we use the fact that if $A_1, B \succ 0$ and $A_2 \succeq A_1$, then $\det(I + A_1 B^{-1}) \leq \det(I + A_2 B^{-1})$, together with the bound $\mathbb{V}_{\Delta}(\bar{v}) \preceq \|\bar{v}\|_{\ominus}^2 \Delta^{\top} \Delta$.

Hence, introducing the random offset $\boldsymbol{\mu} := \sqrt{2 \log \det(I + (\Delta^{\top} \Delta V_0^{-1})^{1/2})}$, which does not depend on \bar{v} , we see that for any $\bar{v} \in \mathbb{R}^{Tp}$,

$$\begin{aligned} &\mathbb{P}[\mathbf{z}_{\bar{v}} \geq \|\bar{v}\|_{\ominus} (\boldsymbol{\mu} + \sqrt{2 \log(1/\delta)})] \\ &= \mathbb{P}[\|\Delta^{\top} \bar{\mathbf{U}} \bar{v}\|_{(\Delta^{\top} \Delta + V_0)^{-1}} \geq \|\bar{v}\|_{\ominus} (\boldsymbol{\mu} + \sqrt{2 \log(1/\delta)})] \leq \delta. \end{aligned} \quad (\text{E.20})$$

Therefore, we find that

$$\begin{aligned} &\mathbb{P} \left[\mathbf{z}_x - \mathbf{z}_y > \|x - y\|_{\ominus} \left(\boldsymbol{\mu} + \sqrt{2 \log(1/\delta)} \right) \right] \\ &= \mathbb{P} \left[\|\Delta^{\top} \bar{\mathbf{U}} x\|_{(\Delta^{\top} \Delta + V_0)^{-1}} - \|\Delta^{\top} \bar{\mathbf{U}} y\|_{(\Delta^{\top} \Delta + V_0)^{-1}} > \|x - y\|_{\ominus} \left(\boldsymbol{\mu} + \sqrt{2 \log(1/\delta)} \right) \right] \\ &\leq \mathbb{P} \left[\|\Delta^{\top} \bar{\mathbf{U}}(x - y)\|_{(\Delta^{\top} \Delta + V_0)^{-1}} > \|x - y\|_{\ominus} \left(\boldsymbol{\mu} + \sqrt{2 \log(1/\delta)} \right) \right] \leq \delta. \end{aligned} \quad (\text{E.21})$$

Intuitively this inequality says that $\mathbf{z}_{\bar{v}}$ has subgaussian tails, modulo the offset $\boldsymbol{\mu}$. This enables us to a chaining argument to the increments $\mathbf{z}_{\bar{v}}$, which is detailed in Appendix F:

Corollary E.4 *There exists universal constants c_1, c_2 such that, with probability $1 - \delta$*

$$\sup_{\bar{v} \in \mathcal{S}^{Tp-1}} \mathbf{z}_{\bar{v}} \leq c_1 \left((\sqrt{\log(c_2/\delta)} + \boldsymbol{\mu}) \max_{x \in \mathcal{S}^{Tp-1}} \|x\|_{\Theta} + \gamma_2(\mathcal{S}^{Tp-1}, \|\cdot\|_{\Theta}) \right),$$

where $\gamma_2(\mathcal{S}^{Tp-1}, \|\cdot\|_{\Theta})$ denotes Talagrand's γ_2 -functional (see, e.g. [Talagrand \(2014\)](#)).

By Lemma E.3, we have $\max_{\bar{v} \in \mathcal{S}^{Tp-1}} \|x\|_{\Theta} \leq \sqrt{T} \max_{\bar{v} \in \mathcal{S}^{Tp-1}} \|\bar{v}\|_2 = \sqrt{T}$. In Section F.1, we sharpen a computation of Dudley's bound due to [Krahmer et al. \(2014\)](#) to control $\gamma_2(\mathcal{S}^{Tp-1}, \|\cdot\|_{\Theta})$:

Proposition E.5 (Control of $\gamma_2(\mathcal{S}^{Tp-1}, \|\cdot\|_{\Theta})$) *We have the bound*

$$\gamma_2(\mathcal{S}^{Tp-1}, \|\cdot\|_{\Theta}) \leq \sqrt{T\tilde{p}}, \quad \text{where } \tilde{p} := p \min\{T, \log^2(eTp) \log^2(eT)\}.$$

Combining the above Proposition with Corollary E.4, we have with probability $1 - \delta$,

$$\begin{aligned} \sup_{\bar{v} \in \mathcal{S}^{Tp-1}} \mathbf{z}_{\bar{v}} &\lesssim \left((\sqrt{\log(c_2/\delta)} + \boldsymbol{\mu})\sqrt{T} + \sqrt{T\tilde{p}} \right) \\ &\lesssim T^{1/2} \sqrt{\log(1/\delta) + \boldsymbol{\mu}^2 + \tilde{p}} \\ &\lesssim T^{1/2} \sqrt{\log(1/\delta) + \log \det(I + \boldsymbol{\Delta}^\top \boldsymbol{\Delta} V_0^{-1})^{1/2} + \tilde{p}} \\ &= T^{1/2} \sqrt{\log(1/\delta) + \log \det(I + \boldsymbol{\Delta}^\top V_0^{-1} \boldsymbol{\Delta})^{1/2} + \tilde{p}}, \quad \text{as needed.} \end{aligned}$$

E.2. Proof of Theorem E.1, Part (b)

We shall bound $\sup_{w \in \mathcal{S}^{m-1}} \|w^\top \boldsymbol{\Delta}^\top \bar{\mathbf{U}}\|_2$. For $w \in \mathcal{S}^{m-1}$, the $m = 1$ case of Theorem E.1 with $\boldsymbol{\Delta} \leftarrow \boldsymbol{\Delta}w$ and $V_0 = \kappa^2 \in \mathbb{R}$ implies that with probability at least $1 - \delta$,

$$\begin{aligned} \frac{\|w^\top \boldsymbol{\Delta}^\top \bar{\mathbf{U}}\|_2}{\sqrt{w^\top \boldsymbol{\Delta}^\top \boldsymbol{\Delta} w + \kappa^2}} &= \left\| \left((\boldsymbol{\Delta}w)^\top (\boldsymbol{\Delta}w) + \kappa^2 \right)^{-1/2} (\boldsymbol{\Delta}w)^\top \bar{\mathbf{U}} \right\|_{\text{op}} \\ &\lesssim T^{1/2} \sqrt{\log(1/\delta) + \log(1 + ((\boldsymbol{\Delta}w)^\top (\boldsymbol{\Delta}w) \kappa^{-2})^{1/2} + \tilde{p}} \\ &\leq T^{1/2} \sqrt{\log(1/\delta) + \log(1 + \frac{\|\boldsymbol{\Delta}\|_{\text{op}}^2}{\kappa^2})^{1/2} + \tilde{p}} \\ &\leq T^{1/2} \sqrt{\log(1/\delta) + \log(1 + \frac{\|\boldsymbol{\Delta}\|_{\text{op}}}{\kappa}) + \tilde{p}}, \end{aligned}$$

where the last line uses $\sqrt{1+x^2} \leq 1+x$ for $x \geq 0$. Since $w^\top \boldsymbol{\Delta}^\top \boldsymbol{\Delta} w \leq \|\boldsymbol{\Delta}\|_{\text{op}}^2$, rearranging shows we have that with probability $1 - \delta$,

$$\begin{aligned} \|w^\top \boldsymbol{\Delta}^\top \bar{\mathbf{U}}\|_2 &\leq \sqrt{\|\boldsymbol{\Delta}\|_{\text{op}}^2 + \kappa^2} T^{1/2} \sqrt{\log(1/\delta) + \log(1 + \frac{\|\boldsymbol{\Delta}\|_{\text{op}}}{\kappa}) + \tilde{p}} \\ &\leq (\|\boldsymbol{\Delta}\|_{\text{op}} + \kappa) T^{1/2} \sqrt{\log(1/\delta) + \log(1 + \frac{\|\boldsymbol{\Delta}\|_{\text{op}}}{\kappa}) + \tilde{p}} \end{aligned}$$

We now proceed to union bound over w . A standard covering argument (see e.g. [Vershynin \(2018, Section 4.2\)](#)) shows that if \mathcal{T} is an ϵ -net of \mathcal{S}^{m-1} , then $\|\Delta^\top \bar{\mathbf{U}}\|_{\text{op}} = \sup_{w \in \mathcal{S}^{m-1}} \|w^\top \Delta^\top \bar{\mathbf{U}}\|_2 \leq \frac{1}{1-\epsilon} \sup_{w \in \mathcal{T}} \|w^\top \Delta^\top \bar{\mathbf{U}}\|_2$. A standard computation (see e.g. [Vershynin \(2018, Corollary 4.2.13\)](#)) lets us choose $|\mathcal{T}| \leq m \log(1 + \frac{2}{\epsilon})$. Setting $\epsilon = 1/2$ and union bounding over $w \in \mathcal{T}$, we have that with probability at least $1 - \delta$,

$$\|\Delta^\top \bar{\mathbf{U}}\|_{\text{op}} \lesssim (\|\Delta\|_{\text{op}} + \kappa) T^{1/2} \sqrt{\log(1/\delta) + \log(1 + \frac{\|\Delta\|_{\text{op}}}{\kappa})} + m + \tilde{p}.$$

Recall the definition $\text{lil}(x) := \log_+(\log_+(x))$. The final bound follows directly from invoking the following lemma with $Z = \|\Delta^\top \bar{\mathbf{U}}\|_{\text{op}}$, $C \lesssim T^{1/2}$, $D = m + \tilde{p}$, and $M = \|\Delta\|_{\text{op}}$:

Lemma E.6 (Iterated Logarithm Conversion) *Let Z be a random variable and suppose that there exists constants C, D and a random variable M such that, for any $\kappa > 0$, it holds that $\mathbb{P}[Z \geq C(\kappa + M) \sqrt{\log(1/\delta) + D + \log(1 + \frac{M}{\kappa})}] \leq \delta$. Then, for any $\beta, \kappa > 0$, one also has*

$$\mathbb{P}\left[Z \geq C((1 + \beta)M + \kappa) \sqrt{\log \frac{2}{\delta} + 2 \text{lil}(\frac{\beta M}{\kappa}) + D + \log\left(1 + \frac{\epsilon}{\beta}\right)}\right] \leq \delta.$$

In particular, if $D \gtrsim 1$, then by setting $\beta = 1$, $\mathbb{P}\left[Z \gtrsim C(M + \kappa) \sqrt{\log \frac{1}{\delta} + \text{lil}(\frac{M}{\kappa}) + D}\right] \leq \delta$.

Proof Define $\kappa_j = e^{j-1} \kappa$ for $j \geq 1$, and $\delta_j = \frac{\delta}{2j^2}$. Then, $\sum_{j \geq 1} \delta_j \leq \delta$, and by a union bound,

$$\mathbb{P}[Z \geq \inf_{j \geq 1} C(\kappa_j + M) \sqrt{\log \frac{1}{\delta_j} + D + \log(1 + \frac{M}{\kappa})}] \leq \delta.$$

In particular, choosing $j = \lceil \log \beta M / \kappa \rceil$, we have that $\frac{\beta M}{e} \leq \kappa_j \leq \max\{\kappa, \beta M\}$, which implies $M + \kappa_j \leq M + \max\{\kappa, \beta M\} \leq (1 + \beta)M + \kappa$ and $\log(1 + \frac{M}{\kappa_j}) \leq \log(1 + \frac{\epsilon}{\beta})$. Moreover, $\delta_j = \frac{\delta}{2j^2} \geq \frac{\delta}{2 \max\{\log \beta M / \kappa, 1\}^2}$, which implies $\log(1/\delta_j) \leq \log(2/\delta) + 2 \log(\max\{\log(\beta M / \kappa), 1\}) \leq \log(2/\delta) + 2 \text{lil}(\beta M / \kappa)$. Hence,

$$\mathbb{P}\left[Z \geq C((1 + \beta)M + \kappa) \sqrt{\log(2/\delta) + 2 \text{lil}(\frac{\beta M}{\kappa}) + D + \log(1 + \frac{\epsilon}{\beta})}\right] \leq \delta.$$

■

E.3. Proof of Lemma E.3

To lower bound $\|\bar{v}\|_{\ominus} = \|\text{Toep}_T(\bar{v})\|_{\text{op}} \geq \|\bar{v}\|_2$. observe that the first row of $\text{Toep}_T(\bar{v})$ consists of the vector $[v_N^\top, v_{N-1}^\top, \dots, v_1^\top] \in \mathbb{R}^T$, followed by zeros. Hence, $\|\text{Toep}_T(\bar{v})\|_{\text{op}} \geq \|[v_N^\top, v_{N-1}^\top, \dots, v_1^\top]\|_2 = \|\bar{v}\|_2$. To upper bound $\|\bar{v}\|_{\ominus} = \|\text{Toep}_T(\bar{v})\|_{\text{op}}$, observe that by the norm-contraction inequality $\|\text{Toep}_T(\bar{v})\|_{\text{op}}$ is bounded by the operator norm of the Toeplitz matrix $M \in \mathcal{R}^{(N-N_1+1) \times (N-N_0+1)}$ (recall $N_0 = N_1 - T$) whose ij -th entry is the norm of the ij -th vector block of $\text{Toep}_T(\bar{v})$. This is an upper triangular Toeplitz matrix with an associated sequence $a_i = \|V_{T-i}\|_2$ for $i \in \{0, \dots, T-1\}$, and $a_i = 0$ for $i \geq N$. By a standard inequality, the operator norm of M is bounded by the ℓ_1 norm of the sequence a_0, a_1, \dots , which is $\sum_{i=0}^{T-1} \|V_{N-i}\|_2$, which in turn is at most $\sqrt{T \sum_{i=0}^{T-1} \|V_{N-i}\|_2^2} = \sqrt{T \|\bar{v}\|_2^2}$.

E.4. Proof of Lemma E.2

The proof is essentially identical to that of Theorem 3 in [Abbasi-Adkori \(2011\)](#), with the exception that the variance-process is matrix-valued, and the noise process is vector-valued. Let $\mathbf{S}_k := \sum_{t=1}^k A_{t-1}^\top \mathbf{u}_t$. We begin by constructing a supermartingale for each direction $w \in \mathbb{R}^m$:

$$M_k(w) := \exp \left(\langle w, \mathbf{S}_k \rangle - \frac{1}{2} \|w\|_{\mathbf{V}_k}^2 \right).$$

Note that $M_0 = 1$. To verify that $M_k(w)$ is a supermartingale with respect to the filtration \mathcal{F}_k , we see use the fact that $\mathbf{S}_k = \mathbf{S}_{k-1} + A_{k-1} \mathbf{u}_k$ and $\mathbf{V}_k = A_{k-1} A_{k-1}^\top + \mathbf{V}_{k-1}$ to write

$$\begin{aligned} \mathbb{E}[M_k(w) | \mathcal{F}_{k-1}] &= \mathbb{E} \left[\exp \left(\langle w, A_{k-1}^\top \mathbf{u}_k + \mathbf{S}_{k-1} \rangle - \frac{1}{2} \|w\|_{\mathbf{V}_{k-1} + A_{k-1}^\top A_{k-1}}^2 \right) \right] \\ &= \mathbb{E} \left[\exp \left(\langle w, A_{k-1}^\top \mathbf{u}_k + \mathbf{S}_{k-1} \rangle - \frac{1}{2} \|w\|_{\mathbf{V}_{k-1}}^2 - \frac{1}{2} \|w\|_{A_{k-1}^\top A_{k-1}}^2 \right) \right] \\ &= \mathbb{E} \left[\exp \left(\langle w, A_{k-1}^\top \mathbf{u}_k \rangle - \frac{1}{2} \|w\|_{A_{k-1}^\top A_{k-1}}^2 \right) M_{k-1}(w) \right] \\ &\stackrel{(i)}{=} M_{k-1}(w) \cdot \mathbb{E} \left[\exp \left(\langle w, A_{k-1}^\top \mathbf{u}_k \rangle - \frac{1}{2} \|w\|_{A_{k-1}^\top A_{k-1}}^2 \right) \right] \\ &= M_{k-1}(w) \cdot \mathbb{E} \left[\exp \left(\langle A_{k-1} w, \mathbf{u}_k \rangle - \frac{1}{2} \|A_{k-1} w\|^2 \right) \right] \stackrel{(ii)}{\leq} M_{k-1}, \end{aligned}$$

where (i) uses the fact that M_{k-1} is \mathcal{F}_k measurable, and (ii) uses that $M_{k-1} \geq 0$ and $\langle A_{k-1} w, \mathbf{u}_k \rangle$ is $\|A_{k-1} w\|^2$ -subgaussian. Since $M_0 = 1$, we conclude by the optional stopping theorem that for any $k \in \mathbb{N}$ and $w \in \mathbb{R}^m$, $\mathbb{E}[M_k(w)] \leq 1$. The remainder of the proof follows that of Theorem 3 in [Abbasi-Adkori \(2011\)](#) verbatim. Specifically, these steps show that

$$\|S_\tau\|_{\bar{\mathbf{V}}_\tau}^2 \leq 2 \log \left(\frac{\det(V_0^{-1/2}) \det(\bar{\mathbf{V}}_\tau^{1/2})}{\delta} \right), \text{ where } \bar{\mathbf{V}}_\tau = \mathbf{V}_\tau + V_0.$$

To conclude, we verify that

$$\begin{aligned} \det(\bar{\mathbf{V}}_\tau^{1/2}) \det(V_0^{-1/2}) &= \sqrt{\det(\bar{\mathbf{V}}_\tau) \det(V_0^{-1})} = \sqrt{\det(\bar{\mathbf{V}}_\tau \cdot V_0^{-1})} \\ &= \sqrt{\det((V_0 + \mathbf{V}_\tau) \cdot V_0^{-1})} = \sqrt{\det(I + \mathbf{V}_\tau V_0^{-1})}. \end{aligned}$$

Appendix F. Chaining for Self-Normalized Tail Inequalities

In this section, we introduce a generic inequality for martingales. Let's consider the general set up. Let (Ω, \mathcal{F}) denote a probability space, \mathcal{X} denote a separable space with metric $d(\cdot, \cdot)$, $\{\mathbf{z}_x\}_{x \in \mathcal{X}}$ denote a real valued random process defined on (Ω, \mathcal{F})

Definition F.1 Let σ, μ denote random variables taking values in $\mathbb{R}_{\geq 0}$. We say that a process $\{\mathbf{z}_x\}_{x \in \mathcal{X}}$ is a (σ, μ) -offset subgaussian process on (\mathcal{X}, d) if, for any $x, y \in \mathcal{X}$ and $u > 0$,

$$\mathbb{P}[\mathbf{z}_x - \mathbf{z}_y \geq d(x, y)(\sigma u + \mu)] \leq \exp(-u^2/2).$$

Note that we do not require $\sigma, \mathbf{u}, \{\mathbf{z}_x\}$ to be independent. We now define Talagrand's γ_2 functional:

Definition F.2 *Let \mathcal{X} be a separable metric space. Talagrand's γ_2 function is defined as*

$$\gamma_2(\mathcal{X}, \mathbf{d}) := \inf_{(\mathcal{A}_n): |\mathcal{A}_n| \leq 2^{2^n}} \sup_{x \in \mathcal{X}} \sum_{n \geq 0} 2^{n/2} \text{diam}(\mathcal{A}_n(x)),$$

where the inf is taken over all sequences of partitions (\mathcal{A}_n) of \mathcal{X} of size $|\mathcal{A}_n| \leq 2^{2^n}$ (with the exception of $|\mathcal{A}_0| = 1$).

Typically, computing the infimum in the definition of γ_2 may be challenging. Fortunately, there exists a easier-to-manage upper bound on γ_2 due to Dudley:

Proposition F.1 (Dudley's Bound) *Let \mathcal{X} be a separable metric space with $\text{diam}(\mathcal{X}) < \infty$. Then,*

$$\gamma_2(\mathcal{X}, \mathbf{d}) \leq \int_0^{\text{diam}(\mathcal{X})} \sqrt{\log \mathcal{N}(\mathcal{X}, \mathbf{d}, u)} du,$$

where $\mathcal{N}(\mathcal{X}, \mathbf{d}, u)$ denotes the cardinality of the minimal u -covering of \mathcal{X} , that is, the cardinality of a minimal subset \mathcal{T} of \mathcal{X} satisfying $\sup_{x \in \mathcal{X}} \inf_{y \in \mathcal{T}} \mathbf{d}(x, y) \leq u$.

Finally, we introduce the main theorem of this section, which extends the generic chaining applied to typical subgaussian processes to (σ, μ) -offset-subgaussian process:

Theorem F.2 *Let $\{\mathbf{z}_x\}_{x \in \mathcal{X}}$ be a (σ, μ) -offset subgaussian process. Then, there exists universal constants c such that*

$$\mathbb{P}\left[\sup_{x, y \in \mathcal{X}} |\mathbf{z}_x - \mathbf{z}_y| \geq c_1 (\sigma \gamma_2(\mathcal{X}, \mathbf{d}) + \text{diam}(\mathcal{X})(\sigma u + \mu))\right] \leq c_2 e^{-u^2}.$$

If in addition \mathcal{X} is normed spaced with $\mathbf{d}(x, y) = \|x - y\|$, and $\mathbb{P}[\mathbf{z}_x \geq \|x\|(\sigma u + \mu)] \leq \exp(-u^2/2)$, then

$$\mathbb{P}\left[\sup_{x, y \in \mathcal{X}} |\mathbf{z}_x| \geq c_1 \left(\sigma \gamma_2(\mathcal{X}, \mathbf{d}) + \left(\min_{x \in \mathcal{X}} \|x\| + \text{diam}(\mathcal{X})\right)(\sigma u + \mu) \right)\right] \leq c_2 e^{-u^2}.$$

With Theorem F.2 in hand, the proof of Corollary E.4 is nearly immediate.

Proof Let $\mathcal{X} = \mathcal{S}^{Tp-1} \subset \mathbb{R}^{Tp}$, and define the norm $\|\cdot\|_{\ominus} := \|\text{Toep}_T(\bar{v})\|_{\text{op}}$. Recall the inequalities (E.20) and (E.21), restated here for convenience with the $\|\cdot\|_{\ominus}$ notation:

$$\mathbb{P}\left[\mathbf{z}_{\bar{v}} \geq \sqrt{2 \log(1/\delta)} \|\bar{v}\|_{\ominus} \sigma + \mu\right] \leq \delta,$$

where $\mathbf{u} = \mu := \sqrt{2 \log \det(I + (\Delta^\top \Delta V_0^{-1})^{1/2})}$ and $\sigma = 1$. By setting $u = \sqrt{2 \log(1/\delta)}$, we see that \mathbf{z}_x is (σ, μ) -offset normed subgaussian with respect to $\gamma_2(\mathcal{X}, \|\cdot\|_{\ominus})$. The corollary is therefore a direct consequence of Theorem F.2, again substituting in $u = \sqrt{2 \log(1/\delta)}$. \blacksquare

F.1. Proof of Proposition E.5

Proof Our argument follows the proof of Theorem 4.1 in [Krahmer et al. \(2014\)](#), modifying the argument to ensure tighter dependences on T and p . Notably, we remove a dependence on the number of samples N which would arise from invoking [Krahmer et al. \(2014\)](#) without alteration. We shall use Dudley's inequality to bound

$$\begin{aligned} \gamma(\mathcal{S}^{Tp-1}, \|\cdot\|_{\Theta}) &\leq \int_{u=0}^{\text{diam}(\mathcal{S}^{Tp-1}, \|\cdot\|_{\Theta})} \sqrt{\log \mathcal{N}(\mathcal{S}^{Tp-1}, \|\cdot\|_{\Theta}, u)} du \\ &= \int_{u=0}^{\sqrt{2T}} \sqrt{\log \mathcal{N}(\mathcal{S}^{Tp-1}, \|\cdot\|_{\Theta}, u)} du \lesssim \sqrt{Tp(\min\{T, \log^2(eT) \log^2(eTp)\})} \end{aligned}$$

Using the fact that $\|\cdot\|_{\Theta} \leq \sqrt{T} \|\cdot\|_2$, and $\log \mathcal{N}(\mathcal{S}^{Tp-1}, \sqrt{T} \|\cdot\|_2, u) \lesssim Tp \log(1 + \frac{\sqrt{T}}{u})$, one can coarsely bound the above by $\sqrt{Tp} \int_{u=0}^{\sqrt{2T}} \sqrt{\log(1 + \frac{\sqrt{T}}{u})} du \lesssim \sqrt{T^2 p} \int_0^1 \sqrt{\log(1 + \frac{1}{u})} du \lesssim \sqrt{T^2 p}$.

It remains to prove the more refined bound of $\sqrt{Tp \log^2(eT) \log^2(eTp)}$. To do so, we need to control the associated covering numbers $\log \mathcal{N}(\mathcal{S}^{Tp-1}, \|\cdot\|_{\Theta}, u)$. To this end, we shall require two ingredients. The first is known as Maurey's Lemma,

Lemma F.3 [Krahmer et al. \(2014, Lemma 4.2\)](#) *Let \mathcal{U} denote a finite subset of a normed space $(\mathcal{X}, \|\cdot\|)$, and suppose that there exists an $A > 0$ such that, for any $k \in \mathbb{N}$ and sequence $(u_1, \dots, u_k) \in \mathcal{U}^k$, $\mathbb{E}[\|\sum_{i=1}^k \epsilon_i u_i\|] \leq A\sqrt{k}$, where ϵ_i are independent Rademacher random variables. Then,*

$$\log \mathcal{N}(\text{conv}(\mathcal{U}), \|\cdot\|, u) \leq \left(\frac{A}{u}\right)^2 \log |\mathcal{U}|.$$

In order to apply Maurey's lemma, we shall choose $\mathcal{U} = \{\sqrt{Tp} \cdot e_i\}_{i=1}^{Tp}$ and $\mathcal{X} = \mathbb{R}^{Tp}$ with the metric $\|\cdot\|_{\Theta}$. Observe that $\mathcal{S}^{Tp-1} \subset \text{conv}(\mathcal{U})$, and thus $\log \mathcal{N}(\mathcal{S}^{Tp-1}, \|\cdot\|_{\Theta}, u) \leq \log \mathcal{N}(\text{conv}(\mathcal{U}), \|\cdot\|_{\Theta}, u)$. To estimate the latter quantity using Lemma F.3, we shall require the following characterization of $\|\cdot\|_{\Theta}$:

Lemma F.4 *There exists a set $\mathcal{Z} \subset \mathbb{R}^{Tp}$ with $|\mathcal{Z}| \leq 16\pi(Tp)^{3/2}$ such that $\|z\|_{\infty} \leq 1$ for all $z \in \mathcal{Z}$ and $\|\bar{v}\|_{\Theta} \leq 4 \max_{z \in \mathcal{Z}} z^{\top} \bar{v}$ for all $\bar{v} \in \mathbb{R}^{Tp}$.*

Lemma F.4 is proven in the subsection below. We stress that it the lemma is crucial to removing the N -dependence in our final bound; the proof of Theorem 4.1 in [Krahmer et al. \(2014\)](#) effectively renders $\|\bar{v}\|_{\Theta} \leq \max_{z \in \mathcal{Z}'} z^{\top} \bar{v}$ for some $|\mathcal{Z}'| \approx N$.

To obtain a covering estimate, observe that for any sequence $u_1, \dots, u_k \in \mathcal{U}^k$, Lemma F.4 implies

$$\mathbb{E}\left[\left\|\sum_{i=1}^k \epsilon_i u_i\right\|\right] \lesssim \mathbb{E}\left[\max_{z \in \mathcal{Z}} \langle z, \sum_{i=1}^k \epsilon_i u_i \rangle\right] = \mathbb{E}\left[\max_{z \in \mathcal{Z}} \sum_{i=1}^k \epsilon_i \cdot \langle z, u_i \rangle\right]$$

For any fixed $z \in \mathcal{Z}$, we have $|\langle z, u_i \rangle| \leq \|z\|_{\infty} \|u_i\|_1 \leq \sqrt{Tp}$. Thus, by a union bound and Hoeffding's inequality, we have that for any appropriate constant c_1 ,

$$\mathbb{P}\left[\sup_{z \in \mathcal{Z}} \sum_{i=1}^k \epsilon_i \cdot \langle z, u_i \rangle \geq \sqrt{Tp} \cdot \sqrt{ku}\right] \leq \sum_{z \in \mathcal{Z}} \mathbb{P}\left[\sum_{i=1}^k \epsilon_i \cdot \langle z, u_i \rangle \geq \sqrt{Tp} \cdot \sqrt{ku}\right] \leq |\mathcal{Z}| e^{-c_1 u^2}.$$

A standard tail integration argument reveals that $\mathbb{E}[\sup_{z \in \mathcal{Z}} \sum_{i=1}^k \epsilon_i \cdot \langle z, u_i \rangle \geq \sqrt{k}u] \lesssim \sqrt{k} \cdot \sqrt{\log |\mathcal{Z}|} \cdot \sqrt{Tp}$ for an appropriate constant c_2 . Thus, \mathcal{U} and $\|\cdot\|_{\ominus}$ satisfy the conditions of Lemma F.3 with $A \lesssim \sqrt{Tp \log |\mathcal{Z}|}$, and thus

$$\log \mathcal{N}(\mathcal{S}^{Tp-1}, \|\cdot\|_{\ominus}, u) \leq \log \mathcal{N}(\text{conv}(\mathcal{U}), \|\cdot\|_{\ominus}, u) \lesssim \frac{Tp \log |\mathcal{Z}|}{u^2} \cdot \log |\mathcal{U}| \lesssim \frac{Tp \log^2(eTp)}{u^2},$$

where we use the bound $|\mathcal{U}| = Tp$ and $|\mathcal{Z}| \lesssim (Tp)^{3/2}$. For $u \leq 1$, this bound is quite loose, and instead we shall use the bound $\|\cdot\|_{\ominus} \leq T \|\cdot\|_2$ to obtain a standard covering bound:

$$\log \mathcal{N}(\mathcal{S}^{Tp-1}, \|\cdot\|_{\ominus}, u) \leq \log \mathcal{N}(\mathcal{S}^{Tp-1}, Tp \|\cdot\|_2, u) \lesssim Tp \log\left(\frac{eT}{u}\right),$$

Invoking Dudley's bound yields

$$\begin{aligned} \gamma(\mathcal{S}^{Tp-1}, \|\cdot\|_{\ominus}) &\lesssim \int_1^{\sqrt{2T}} \sqrt{\frac{Tp \log^2(eTp)}{u^2}} du + \int_0^1 \sqrt{Tp \log \frac{eT}{u}} \\ &= \sqrt{Tp} \left(\log(eTp) \int_1^{\sqrt{2T}} \frac{du}{u} + \int_0^1 \sqrt{\log \frac{eT}{u}} \right) \\ &\leq \sqrt{Tp} \left(\log(eTp) \int_1^{\sqrt{2T}} \frac{du}{u} + \sqrt{\log eT} + \int_0^1 \sqrt{\log \frac{1}{u}} \right) \\ &= \sqrt{Tp} \left(\log(eTp) \log(\sqrt{2T}) + \sqrt{\log eT} + \int_0^1 \sqrt{\log \frac{1}{u}} \right) \\ &\lesssim \sqrt{Tp} \log(eTp) \log(eT), \end{aligned}$$

where we use the fact that $\int_0^1 \sqrt{\log \frac{1}{u}}$ is at most a universal constant. ■

F.1.1. PROOF OF LEMMA F.4

We can embed $\text{Toep}_T(\bar{v})$ as a submatrix an infinite Toeplitz operator where $a_{-i} = \bar{v}_{i+1}$ for $i \in \{0, \dots, Tp\}$, and 0 elsewhere. It is well known that such a matrix has operator norm bounded by the \mathcal{H}_{∞} norm $\max_{\theta \in [0, 2\pi]} \left| \sum_{k=1}^{Tp} \bar{v}[k] e^{ik\theta} \right|$, where $|\cdot|$ denotes complex magnitude. Define $F(\theta) := \left| \sum_{k=1}^{Tp} \bar{v}[k] e^{ik\theta} \right|$. Our proof will have two ingredients:

1. For all $\theta \in [0, 2\pi]$, we can bound $F(\theta) \leq \sqrt{2} \sup_{\ell \in [4]} \langle z_{\theta, \ell}, \bar{v} \rangle$, where $\|z_{\theta, \ell}\|_{\infty} \leq 1$ for $\ell \in [4]$.
2. We construct a covering \mathcal{T} of $[0, 2\pi]$ such that $\sup_{\theta \in [0, 2\pi]} F(\theta) \leq 2 \max_{\theta \in \mathcal{T}} F(\theta)$, and $|\mathcal{T}| \leq 4\pi(Tp)^{3/2}$

Together, these two imply that, for $\mathcal{Z} = \{z_{\theta, \ell} : \theta \in \mathcal{T}, \ell \in [4]\}$, $\|\bar{v}\|_{\ominus} \leq 2\sqrt{2} \sup_{z \in \mathcal{Z}} \langle z, \bar{v} \rangle$, as well as $\|z\|_{\infty} \leq 1$ for all $z \in \mathcal{Z}$, and $|\mathcal{Z}| = 4|\mathcal{T}| \leq 16\pi(Tp)^{3/2}$.

Proof of Point 1: Given a complex number ω , we have $|\omega| = \sqrt{\Re(\omega)^2 + \Im(\omega)^2} \leq \sqrt{2} \max\{|\Re(\omega)|, |\Im(\omega)|\} = \sqrt{2} \max\{\Re(\omega), -\Re(\omega), \Im(\omega), -\Im(\omega)\}$. Defining the complex vector $z_\theta = (e^{ik\theta})_{k=1}^{Tp}$, and $z_{\theta,1} = \Re(z_\theta)$, $z_{\theta,2} = \Im(z_\theta)$, $z_{\theta,3} = -z_{\theta,1}$, $z_{\theta,4} = -z_{\theta,2}$, we see that $F(\theta) \leq \sqrt{2} \max_{\ell \in [4]} \langle z_{\ell,\theta}, \bar{v} \rangle$. Moreover, $\max_{\ell \in [4]} \|z_{\ell,\theta}\|_\infty \leq 1$, as $\max\{|\Re(e^{ik\theta})|, |\Im(e^{ik\theta})|\} \leq 1$ for all $k \in \mathbb{R}$.

Proof of Point 2: It suffices to show that $F(\theta)$ is $\|\bar{v}\|_2 \cdot (Tp)^{3/2}$ -Lipschitz. Indeed, if this is true then by choosing a $\frac{1}{2(Tp)^{3/2}}$ -net \mathcal{T} of $[0, 2\pi]$, we have

$$\begin{aligned} \|\bar{v}\|_\Theta &= \max_{\theta \in [0, 2\pi]} F(\theta) \leq \max_{\theta \in \mathcal{T}} F(\theta) + \max_{\theta \in [0, 2\pi]} \min_{\theta' \in \mathcal{T}} |F(\theta) - F(\theta')| \\ &\leq \max_{\theta \in \mathcal{T}} F(\theta) + \frac{1}{2(Tp)^{3/2}} \cdot (Tp)^{3/2} \|\bar{v}\|_2 \leq \max_{\theta \in \mathcal{T}} F(\theta) + \frac{1}{2} \|\bar{v}\|_\Theta, \end{aligned}$$

where we used the bound that $\|\bar{v}\|_\Theta \geq \|\bar{v}\|_2$. After rearranging, $\|\bar{v}\|_\Theta \leq \max_{\theta \in \mathcal{T}} 2F(\theta)$. Lastly, we note that we can construct a $\frac{1}{2(Tp)^{3/2}}$ -net of the interval $[0, 2\pi]$ of size at most $\lfloor 4\pi(Tp)^{3/2} \rfloor \leq 4\pi(Tp)^{3/2}$. It remains to show that $F(\theta)$ is Lipschitz. We can bound

$$|F(\theta_1) - F(\theta_2)| \leq \left| \sum_{k=0}^{Tp-1} \bar{v}[k](e^{ik\theta_1} - e^{ik\theta_2}) \right| \leq \|\bar{v}\|_2 \sqrt{\sum_{k=0}^{Tp-1} |e^{ik\theta_1} - e^{ik\theta_2}|^2}$$

by Cauchy Schwartz. Geometrically, $|e^{ik\theta_1} - e^{ik\theta_2}|$ is the distance between the point $(\cos k\theta_1, \sin k\theta_1)$ and the point $(\cos k\theta_2, \sin k\theta_2)$ on the unit sphere, which is at most the arc length $k|\theta_2 - \theta_1|$ between the two points. Hence, $\sqrt{\sum_{k=0}^{Tp-1} |1 - e^{ik(\theta_2 - \theta_1)}|^2} \leq \sqrt{\sum_{k=0}^{Tp-1} k^2(\theta_2 - \theta_1)^2} \leq (Tp)^{3/2} |\theta_2 - \theta_1|$.

F.2. Proof of Theorem F.2

The second part of the theorem is a consequence of the first after noting that

$$\max_{y \in \mathcal{X}} \mathbf{z}_y \leq \min_{x \in \mathcal{X}} \left(|\mathbf{z}_x| + \max_{y \in \mathcal{X}} |\mathbf{z}_y - \mathbf{z}_x| \right) \leq \min_{x \in \mathcal{X}} |\mathbf{z}_x| + \max_{x, y \in \mathcal{X}} |\mathbf{z}_y - \mathbf{z}_x|,$$

and that $\max_{x, y \in \mathcal{X}} |\mathbf{z}_y - \mathbf{z}_x|$ can be bounded by the first part of the theorem, whereas $\min_{x \in \mathcal{X}} \mathbf{z}_x$ can be bounded by the condition $\mathbb{P}[\mathbf{z}_x \geq \|x\|(\boldsymbol{\sigma}u + \boldsymbol{\mu})] \leq \exp(-u^2/2)$.

The proof of the first part of the theorem is analogous to the proof of Theorem 2.2.27 in [Talagrand \(2014\)](#). With a standard separability argument, we may assume without loss of generality that \mathcal{X} is finite. Let $N_n := 2^{2^n}$. By the definition of the γ_2 functional, we may choose a sequence (\mathcal{A}_n) of partitions of \mathcal{X} of size at most $|\mathcal{A}_n| \leq N_n$ satisfying $(\mathcal{A}_n)_{n \geq 0}$ with $\sup_{x \in \mathcal{X}} \sum_{n \geq 0} 2^{n/2} \text{diam}(\mathcal{A}_n(x)) \leq 2\gamma_2(\mathcal{X})$.

The key observation is that this is chosen to (nearly) minimize a bound involving only distances between elements of \mathcal{X} , and is not chosen based on the random variances $\boldsymbol{\sigma}$ or offsets $\boldsymbol{\mu}$.

Let $(\mathcal{X}_n)_{n \geq 0}$ denote any sequence of subsets of \mathcal{X}_n where \mathcal{X}_n comprises of exactly one element of each set in \mathcal{A}_n ; note then that $|\mathcal{X}_n| \leq N_n$. Define the unioned sets $\mathcal{U}_n := \bigcup_{q=0}^n \mathcal{X}_q$. In particular, $\mathcal{U}_0 = \mathcal{X}_0$, and $|\mathcal{U}_n| \leq \sum_{q=0}^n N_q \leq 2N_n$. Lastly, we define the event

$$\mathcal{E}(u) := \{|\mathbf{z}_x - \mathbf{z}_x| \leq 2d(x_1, x_2)(\boldsymbol{\sigma}(u + 2^{n/2}) + \boldsymbol{\mu}), \quad \forall n \geq 1, x, y \in \mathcal{U}_n, \mathbf{z}_x\}.$$

Display (2.60) in Talagrand verifies that $\mathbb{P}[\mathcal{E}(u)^c] \leq 2 \sum_{n \geq 1} |\mathcal{U}_n|^2 \exp(-2(2^n + u^2)) \leq c_2 \exp(-2u^2)$. Therefore, it suffices to show that there exists a constant c_1 such that, if for any fixed $\sigma, \mu \geq 0$, then for any deterministic real-valued process $\{z_x\}_{x \in X}$ satisfying the condition

$$|z_x - z_y| \leq 2d(x, y)(\sigma(u + 2^{n/2}) + \mu), \quad \forall n \geq 1, x, y \in \mathcal{U}_n, \quad (\text{F.22})$$

it holds that $\sup_{x, y \in \mathcal{X}} |z_x - z_y| \leq c_1(\sigma\gamma_2(\mathcal{X}) + \text{diam}(\mathcal{X})(\sigma u + \mu))$.

First consider the case $\sigma = 0$. For any $x, y \in \mathcal{X}$, it suffices to show $|z_x - z_y| \leq 2\mu \text{diam}(\mathcal{X})$. Since $\bigcup_n \mathcal{U}_n = \mathcal{X}$ and \mathcal{X} is finite by assumption, there exists an n large enough for which $x, y \in \mathcal{U}_n$. By (F.22), $|z_x - z_y| \leq 2d(x, y)(\sigma(u + 2^{n/2}) + \mu) = 2d(x, y)\mu \leq 2\mu \text{diam}(\mathcal{X})$.

Next, consider the case $\sigma > 0$. Define $z'_x := z_x/\sigma$, and $u' = u + \mu/\sigma$. Then, the process z'_x satisfies the condition:

$$|z'_x - z'_y| \leq 2d(x, y)(u' + 2^{n/2}), \quad \forall n \geq 1, x, y \in \mathcal{U}_n, \quad (\text{F.23})$$

The proof of Talagrand (2014, Theorem 2.2.27) shows that (F.23) implies that $\sup_{x, y \in \mathcal{X}} |z'_x - z'_y| \leq c_1(\gamma_2(\mathcal{X}) + \text{diam}(\mathcal{X})u')$. Multiplying both sides by σ , we have

$$\sup_{x, y \in \mathcal{X}} |z_x - z_y| \leq c_1(\gamma_2(\mathcal{X}) + \sigma \text{diam}(\mathcal{X})u') = c_1(\gamma_2(\mathcal{X}) + \text{diam}(\mathcal{X})(\sigma u + \mu)).$$

Part III

Prefiltered Least Squares for Linear Dynamical Systems

Appendix G. Bounds on $\|\Delta_\phi\|_{\text{op}}$

In this section, we extend Proposition 4.1 to the case where $\mathbf{x}_1 \neq 0$, and to the adversarial noise regime. Throughout, we assume that \mathbf{x}_1 is deterministic (stochastic bounds can be developed by reasoning over the randomness of \mathbf{x}_1). When $\mathbf{x}_1 \neq 0$, we now have three LTI systems:

$$\mathbf{G}_\phi := (A_\star, B_\star, C_\phi, 0), \quad \mathbf{F}_\phi := (A_\star, B_w, C_\phi, 0), \quad \mathbf{H}_\phi := (A_\star, \mathbf{x}_1, C_\phi, 0). \quad (\text{G.24})$$

Our bound for stochastic noise with general $\mathbf{x}_1 \neq 0$ is as follows:

Proposition G.1 (Stochastic Noise Bound) *Consider a filter of the form $\phi = [\Psi_1 | \dots | \Psi_d] \in \mathbb{R}^{m \times dm}$ for some $1 \leq d \leq L$, and suppose that $N \geq Td \max\{m, \log(1/\delta)\}$. Then, under the stochastic noise model of Assumption 2.1, with probability at least $1 - \delta$ we have that the extended filter $\tilde{\phi} := [\phi | \mathbf{0}_{m \times (L-m)d}] \in \mathbb{R}^{L \times m}$ satisfies*

$$\begin{aligned} \|\Delta_{\tilde{\phi}}\|_{\text{op}} &\lesssim \sqrt{N}(\|\mathcal{M}_\infty(\mathbf{G}_\phi)\|_{\text{op}} + \|\mathcal{M}_\infty(\mathbf{F}_\phi)\|_{\text{op}}) + \|\mathcal{M}_\infty(\mathbf{H}_\phi)\|_{\text{op}} \\ &\quad + \sqrt{m + \log(1/\delta)}(\Gamma_N(\mathbf{G}_\phi) + \Gamma_N(\mathbf{F}_\phi)) \\ &\quad + \sqrt{N}(1 + \|\phi\|_{\text{bop}})(\|\mathcal{M}_{Td}(\mathbf{G})\|_{\text{op}} + \|\mathcal{M}_{Td}(\mathbf{F})\|_{\text{op}} + \|D_z\|_{\text{op}}), \end{aligned}$$

where we define $\Gamma_N(\mathbf{G}) := \min\{\sqrt{N}\|\mathcal{M}_\infty(\mathbf{G})\|_{\text{op}}, \|\mathbf{G}\|_{\mathcal{H}_\infty}\} \dots$

For adversarial noise, we obtain the following analogous bound:

Proposition G.2 (Bound on Opt_λ for adversarial noise) *Consider a filter of the form $\phi = [\Psi_1 | \dots | \Psi_d]$ for some $1 \leq d \leq L$. With probability at least $1 - \delta$, if $N \geq Td \max\{m, \log(1/\delta)\}$, then in the adversarial noise model, the extended filter $\tilde{\phi} := [\phi | \mathbf{0}_{m \times (L-m)d}] \in \mathbb{R}^{L \times m}$ satisfies*

$$\begin{aligned} \|\Delta_{\tilde{\phi}}\|_{\text{op}} &\lesssim \sqrt{N}\|\mathcal{M}_\infty(\mathbf{G}_\phi)\|_{\text{op}} + \sqrt{Nd_w}\Gamma_N(\mathbf{F}_\phi) + \|\mathcal{M}_\infty(\mathbf{H}_\phi)\|_{\text{op}} \\ &\quad + \sqrt{m + \log(1/\delta)}\Gamma_N(\mathbf{G}_\phi) \\ &\quad + \sqrt{N}(1 + \|\phi\|_{\text{bop}}) \left[\|\mathcal{M}_{Td}(\mathbf{G})\|_{\text{op}} + \sqrt{Tdd_w}\|\mathcal{M}_{Td}(\mathbf{F})\|_{\text{op}} + \sqrt{dd_z}\|D_z\|_{\text{op}} \right]. \end{aligned}$$

G.1. Outline of the Proofs

We shall assume without loss of generality that $d = L$, because the extended filter $\tilde{\phi} \in \mathbb{R}^{m \times Lm}$ and the original $\phi \in \mathbb{R}^{m \times dm}$ yields the same errors. We now outline the proofs of Proposition 4.1 and Proposition G.2. Throughout, we fix a filter $\phi \in \mathbb{R}^{m \times Lm}$. We are aiming to control

$$\|\Delta_\phi\|_{\text{op}} = \left\| \Delta - \mathbf{K}\phi^\top \right\|_{\text{op}} = \left\| [\delta_{N_1} - \phi \cdot \mathbf{k}_{N_1} | \delta_{N_1+1} - \phi \cdot \mathbf{k}_{N_1+1} | \dots | \delta_N - \phi \cdot \mathbf{k}_N] \right\|_{\text{op}}.$$

Concretely, we take our features to be L previous output values equally T -spaced, i.e.

$$\mathbf{k}_t := [\mathbf{y}_{t-T}^\top \mid \mathbf{y}_{t-2T}^\top \mid \cdots \mid \mathbf{y}_{t-LT}^\top]^\top.$$

We express our filter as $\phi = [\Psi_1 \mid \dots \mid \Psi_L]$, for $\Psi_j \in \mathbb{R}^{m \times m}$. Thus, our prediction takes the form

$$\phi \cdot \mathbf{k}_t = \sum_{\ell=1}^L \Psi_\ell \mathbf{y}_{t-\ell T}.$$

Recall from Section 4 the auxiliary signal $\tilde{\mathbf{x}}_{n;t}$, associated observation $\tilde{\mathbf{y}}_{n;t}$, and auxiliary features $\tilde{\mathbf{k}}_t$ defined via

$$\begin{aligned} \tilde{\mathbf{x}}_{n;t} &:= \begin{cases} A_\star^{n-(t-LT)} \mathbf{x}_{t-LT} & n \geq t - LT \\ \mathbf{x}_n & n \leq t - LT \end{cases} \\ \tilde{\mathbf{y}}_{n;t} &:= C_\star \tilde{\mathbf{x}}_{n;t} \\ \tilde{\mathbf{k}}_t &:= [\tilde{\mathbf{y}}_{t-T;t}^\top \mid \tilde{\mathbf{y}}_{t-2T;t}^\top \mid \cdots \mid \tilde{\mathbf{y}}_{t-LT;t}^\top]^\top. \end{aligned}$$

We can now decompose the error term as

$$\delta_{\phi,t} = \delta_t - \phi \cdot \mathbf{k}_t = \underbrace{(\tilde{\mathbf{y}}_{t;t} - \phi \cdot \tilde{\mathbf{k}}_t)}_{(\text{Err}_{\phi,t}^{(1)})} + \underbrace{[(\delta_t - \tilde{\mathbf{y}}_{t;t}) - \phi \cdot (\mathbf{k}_t - \tilde{\mathbf{k}}_t)]}_{(\text{Err}_{\phi,t}^{(2)})},$$

where we have suppressed the dependence on ϕ in $\text{Err}_t^{(1)}$ and $\text{Err}_t^{(2)}$, as ϕ is fixed. We further define the stacked errors⁸

$$\text{Err}_{\phi,t_1:t_2}^{(j)} := \left[\text{Err}_{\phi,t_1}^{(j)} \mid \text{Err}_{\phi,t-1}^{(j)} \mid \dots \mid \text{Err}_{\phi,t_2}^{(j)} \right],$$

and our goal will be to control $\|\text{Err}_{\phi,N:N_1}\|_{\text{op}} \leq \|\text{Err}_{\phi,N:N_1}^{(1)}\|_{\text{op}} + \|\text{Err}_{\phi,N:N_1}^{(2)}\|_{\text{op}}$. Toward bounding these two terms, we outline a general strategy to bound their individual components; the full details can be found in the remainder of the section. Let $\mathbf{u}_{N:1} \in \mathbb{R}^{Np}$, $\mathbf{w}_{N:1} \in \mathbb{R}^{Nd_w}$, $\mathbf{z}_{N:1} \in \mathbb{R}^{Nd_z}$ denote the concatenated (from N down to 1) sequences of input, process noise, and sensor noise vectors. By the linearity of the system, we can express

$$\text{Err}_{\phi,N:N_1} = (\mathbf{G}^{(1)} + \mathbf{G}^{(2)})\mathbf{u}_{N:1} + (\mathbf{F}^{(1)} + \mathbf{F}^{(2)})\mathbf{w}_{N:1} + \mathbf{G}_z^{(2)}\mathbf{z}_{N:1} + \mathbf{H}^{(1)}\mathbf{x}_1.$$

where $\mathbf{G}^{(1)}$ encodes the contribution of $\mathbf{u}_{N:1}$ to $\text{Err}_\phi^{(1)}$, $\mathbf{G}^{(2)}$ encodes the contribution of $\mathbf{u}_{N:1}$ to $\text{Err}_\phi^{(2)}$, and so on. For example, $\mathbf{G}^{(1)}$ is a linear operator that maps \mathbb{R}^{Np} to $\mathbb{R}^{\tilde{N} \times m}$. Note that $\mathbf{z}_{N:1}$ only contributes to the second error term and \mathbf{x}_1 only contributes to the first.

Then, the key to bounding the error terms is that we can write, for example,

$$v^\top \mathbf{G}^{(1)} \mathbf{u}_{N:1} = \left(G_v^{(1)} \mathbf{u}_{N:1} \right)^\top$$

8. Here, we reverse the time ordering of the columns (which is norm-preserving) in order to cleanly write many quantities in terms of Toeplitz operators.

for an appropriate block matrix $G_v^{(1)} \in \mathbb{R}^{\tilde{N} \times Np}$. Thus,

$$\|\mathbf{G}^{(1)} \mathbf{u}_{N:1}\|_{\text{op}} = \sup_{v \in \mathcal{S}^{m-1}} \|G_v^{(1)} \mathbf{u}_{N:1}\|_2. \quad (\text{G.25})$$

Our main tool to control this quantity is the following result, due to [Krahmer et al. \(2014\)](#) (somewhat overloading our notation for \mathcal{M}):

Proposition G.3 *Let $\mathcal{M} \subset \mathbb{R}^{n \times m}$ denote a class of matrices, and let $\boldsymbol{\xi} \in \mathbb{R}^m$ denote a subgaussian vector. Then,*

$$\mathbb{P}[\sup_{M \in \mathcal{M}} \|M\boldsymbol{\xi}\|_2 \gtrsim \gamma_2(\mathcal{M}, \|\cdot\|_{\text{op}}) + \text{diam}(\mathcal{M}, \|\cdot\|_{\text{F}}) + \sqrt{\log(1/\delta)} \text{diam}(\mathcal{M}, \|\cdot\|_{\text{op}})] \leq 1 - \delta.$$

We observe that, even with the martingale structure, the vectors $\mathbf{u}_{N:1}$ is a sub-Gaussian vector, as is $\mathbf{w}_{N:1}$ in the stochastic model. Our strategy is outlined as follows:

1. We begin give a detailed proof to bound $\mathbf{G}^{(1)}$ in Section [G.2](#).
2. We begin give a detailed proof to bound $\mathbf{G}^{(2)}$ in Section [G.3](#).
3. We explain how to bound the contributions of the remaining terms in Section [G.4](#). Specifically, we address process noise under both stochastic and adversarial models in [G.4.1](#) and adversarial noise in [G.4.2](#). Lastly, we address the contributions of the initial state in [G.4.3](#).

G.2. Bounding $\mathbf{G}^{(1)}$

Step 1: Recognize a block matrix structure Recalling that $\mathbf{G}_\phi = (A_\star, B_\star, C_\phi, 0)$, for each vector $v \in \mathbb{R}^m$, we define the block Toeplitz matrix

$$\tilde{G}_v^{(1)} := \begin{bmatrix} & v^\top \mathcal{M}_{N-TL}(\mathbf{G}_\phi) & \\ \mathbf{0}_{1 \times p} & v^\top \mathcal{M}_{N-1-TL}(\mathbf{G}_\phi) & \\ \vdots & \vdots & \\ \mathbf{0}_{1 \times (\tilde{N}-2)p} & v^\top \mathcal{M}_{N_1+1-TL}(\mathbf{G}_\phi) & \\ \mathbf{0}_{1 \times (\tilde{N}-1)p} & v^\top \mathcal{M}_{N_1-TL}(\mathbf{G}_\phi) & \end{bmatrix} \in \mathbb{R}^{\tilde{N} \times (N-TL-1)p}.$$

One can verify that $\tilde{G}_v^{(1)}$ is the nonzero portion (i.e. submatrix) of $G_v^{(1)}$, which is sufficient for our goal of bounding matrix norms.

Step 2: Bound the Frobenius and operator norms By considering each row of $\tilde{G}_v^{(1)}$, we immediately see that

$$\|\tilde{G}_v^{(1)}\|_{\text{F}}^2 \leq N \|v^\top \mathcal{M}_N(\mathbf{G}_\phi)\|_{\text{F}}^2 \leq N \|v\|_2^2 \|\mathcal{M}_N(\mathbf{G}_\phi)\|_{\text{op}}^2.$$

Since $\|\mathcal{M}_k(\mathbf{G}_\phi)\|_{\text{op}}$ is increasing in k , we can bound $\|\mathcal{M}_N(\mathbf{G}_\phi)\|_{\text{op}} \leq \|\mathcal{M}_\infty(\mathbf{G}_\phi)\|_{\text{op}}$. However, we would like to note that, in general, $\max_{v \in \mathcal{S}^{m-1}} \|v^\top \mathcal{M}_N(\mathbf{G}_\phi)\|_{\text{F}}$ could be much less than this quantity, and using it instead would sharpen [Proposition 4.1](#).

Now, concerning the operator norm, we have two options. First, we could simply take

$$\|\tilde{G}_v^{(1)}\|_{\text{op}} \leq \|\tilde{G}_v^{(1)}\|_{\text{F}} \leq \|v\| \sqrt{N} \|\mathcal{M}_\infty(\mathbf{G}_\phi)\|_{\text{op}}.$$

On the other hand, we see that

$$\|\tilde{G}_v^{(1)}\|_{\text{op}} \stackrel{(i)}{\leq} \|v^\top \mathbf{G}_\phi\|_{\mathcal{H}_\infty} = \sup_{|z|=1} \|v^\top C_\phi(zI - A_\star)^{-1} B_\star\|_{\text{op}} \leq \|v\|_2 \|\mathbf{G}_\phi\|_{\mathcal{H}_\infty},$$

where (i) comes from the fact that $\tilde{G}_v^{(1)}$ is, in a sense, a ‘‘submatrix’’ of the infinite-dimensional linear operator $v^\top \mathbf{G}_\phi$ (see e.g. [Tilli \(1998\)](#) Corollary 4.2). Thus, recalling that

$$\Gamma_N(\mathbf{G}_\phi) = \min\{\sqrt{N}\|\mathcal{M}_\infty(\mathbf{G}_\phi)\|_{\text{op}}, \|\mathbf{G}_\phi\|_{\mathcal{H}_\infty}\},$$

we have that $\sup_{v \in \mathcal{S}^{m-1}} \|\tilde{G}_v^{(1)}\|_{\text{op}} \leq \Gamma_N(\mathbf{G}_\phi)$.

Step 3: Bound the γ_2 functional To bound $\gamma_2(\mathcal{M}, \|\cdot\|_{\text{op}})$ with $\mathcal{M} = \{G_v^{(1)} : v \in \mathcal{S}^{m-1}\}$, we first note that $\gamma_2(\mathcal{M}, \|\cdot\|_{\text{op}}) \leq \gamma_2(\mathcal{S}^{m-1}, \|G_{(\cdot)}^{(1)}\|_2)$, as the linear map $v \rightarrow G_v^{(1)}$ is only injective, in general. Then, one can use Step 2 above to show (see [Talagrand \(2014\)](#) Exercise 2.2.23) that

$$\gamma_2(\mathcal{S}^{m-1}, \|G_{(\cdot)}^{(1)}\|_2) \leq \Gamma_N(\mathbf{G}_\phi) \gamma_2(\mathcal{S}^{m-1}, \|\cdot\|_2).$$

Finally, upper bounding the right hand side by Dudley’s integral and using the standard covering number bound for the sphere gives

$$\gamma_2(\mathcal{M}, \|\cdot\|_{\text{op}}) \lesssim \Gamma_N(\mathbf{G}_\phi) \sqrt{m}.$$

We remark that one could instead use the generic chaining (see [Talagrand \(2014\)](#) Chapter 2) to bound γ_2 directly and, instead of \sqrt{m} , get the *stable rank* of some matrix, but we will be loose in this aspect.

Putting these three steps together with Proposition [G.3](#) shows that with probability at least $1 - \delta$,

$$\|\mathbf{G}^{(1)} \mathbf{u}_{N:1}\|_{\text{op}} \lesssim \sqrt{N} \|\mathcal{M}_\infty(\mathbf{G}_\phi)\|_{\text{op}} + \sqrt{m + \log(1/\delta)} \Gamma_N(\mathbf{G}_\phi).$$

G.3. Bounding $\mathbf{G}^{(2)}$

Now, we seek to bound $\|\mathbf{G}^{(2)} \mathbf{u}_{N:1}\|_{\text{op}} = \sup_{v \in \mathcal{S}^{m-1}} \|G_v^{(2)} \mathbf{u}_{N:1}\|_2$, and we repeat the previous three steps.

Step 1 The explicit block matrix structure is a bit cumbersome to write out, but the key is to note that the rows of $G_v^{(2)}$ are simply shifted versions of $v^\top G^{(2)}$, where

$$G^{(2)} := \mathcal{M}_{Td}(\mathbf{G}) \cdot \text{blkdiag}(\mathbf{0}_{T \times T}, I_{T(L-1) \times T(L-1)}) - \sum_{k=1}^L \Psi_k[\mathbf{0}_{m \times Tk} \mid \mathcal{M}_{T(L-k)}(\mathbf{G})].$$

Note that the slight complication with blkdiag occurs due to the subtraction of $G_\star \bar{\mathbf{u}}_t$ in δ_t ; this blkdiag term is simply identity in the process noise case of [G.4.1](#).

Step 2 By the same argument as for the first error term, we have that $\|G_v^{(2)}\|_{\text{F}} \leq \sqrt{N}\|v\|_2\|G^{(2)}\|_{\text{op}}$. The operator norm requires slightly more care, and requires the following lemma.

Lemma G.4 *We have that $\|G_v^{(2)}\|_{\text{op}} \lesssim \sqrt{TL}\|v\|_2\|G^{(2)}\|_{\text{op}}$.*

Proof As stated before, $G_v^{(2)}$ is embeddable in a block Toeplitz matrix with shifted versions of $v^\top G^{(2)}$ as its rows, where the nonzero portion of $v^\top G^{(2)}$ has at most TL consecutive blocks. We use the following elementary property of Toeplitz matrices to bound $\|G_v^{(2)}\|_{\text{op}}$.

Lemma G.5 *Let X and Y be Toeplitz matrices generated by $x \in \mathbb{R}^a$ and $y \in \mathbb{R}^b$, where “generated” can include padding x and y with zeros on either end. Then,*

$$\|XY\|_{\text{op}} \leq \|x * y\|_1 \leq (a + b - 1)\|x\|_2\|y\|_2,$$

where $*$ denotes convolution,

Taking $X = Y^\top = G_v^{(2)}$ (and making the appropriate adjustment for the vector case), we see that

$$\|G_v^{(2)}\|_{\text{op}}^2 = \|G_v^{(2)}G_v^{(2)\top}\|_{\text{op}} \lesssim TL\|v^\top G^{(2)}\|_2^2 \leq TL\|v\|_2^2\|G^{(2)}\|_{\text{op}}^2.$$

■

Furthermore, by Hölder’s inequality (with the convention that $\Psi_0 := 1$),

$$\begin{aligned} \|G^{(2)}\|_{\text{op}} &\leq \left\| \sum_{k=0}^L \Psi_k[\mathbf{0}_{m \times Tk} \mid \mathcal{M}_{T(L-k)}(\mathbf{G})] \right\|_{\text{op}} \\ &\leq \left(1 + \sum_{k=1}^L \|\Psi_k\|_{\text{op}} \right) \max_{j \in [L]} \|\mathbf{0}_{m \times Tj} \mid \mathcal{M}_{T(L-j)}(\mathbf{G})\|_{\text{op}} \\ &\leq (1 + \|\phi\|_{\text{bop}})\mathcal{M}_{TL}(\mathbf{G}). \end{aligned}$$

Thus, we have shown that $\|G_v^{(2)}\|_{\text{op}} \lesssim \sqrt{TL}\|v\|_2(1 + \|\phi\|_{\text{bop}})\mathcal{M}_{TL}(\mathbf{G})$.

Step 3 As with the first term, we upper bound Talagrand’s γ_2 functional using the operator norm calculation from Step 2.

Again putting the three steps together and appealing to Proposition G.3, we see that with probability at least $1 - \delta$,

$$\begin{aligned} \|\mathbf{G}^{(2)}\mathbf{u}_{N:1}\|_{\text{op}} &\lesssim \left(\sqrt{N} + \sqrt{TLm} + \sqrt{TL \log(1/\delta)} \right) (1 + \|\phi\|_{\text{bop}})\mathcal{M}_{TL}(\mathbf{G}) \\ &\lesssim \sqrt{N}(1 + \|\phi\|_{\text{bop}})\mathcal{M}_{TL}(\mathbf{G}), \end{aligned}$$

by the assumption on N .

G.4. Additional terms

G.4.1. PROCESS NOISE

In the stochastic noise model, the previous calculations for $\mathbf{G}^{(1)}$ and $\mathbf{G}^{(2)}$ apply directly as the martingale structure allows us to again use Proposition G.3; namely, the vectors $\mathbf{w}_{s:t}$ are 1-subgaussian.

In the adversarial noise model, we can bound (G.25) directly. First, we see that

$$\|\mathbf{F}^{(1)}\mathbf{w}_{N:1}\|_{\text{op}} = \sup_{v \in \mathcal{S}^{m-1}} \|\mathbf{F}_v^{(1)}\mathbf{w}_{N:1}\|_2 \leq \sqrt{Nd_w} \sup_{v \in \mathcal{S}^{m-1}} \|\mathbf{F}_v^{(1)}\|_{\text{op}} \leq \sqrt{Nd_w} \Gamma_N(\mathbf{F}_\phi),$$

by the same arguments as in the preceding section. Second, the same strategy outlined for $\mathbf{G}^{(2)}$ establishes that

$$\|\mathbf{F}^{(2)}\mathbf{w}_{N:1}\|_{\text{op}} \leq \sqrt{NTLd_w}(1 + \|\phi\|_{\text{bop}})\|\mathcal{M}_{TL}(F_\star)\|_{\text{op}}.$$

G.4.2. OUTPUT NOISE

Let us first start with the stochastic noise model. The same calculations as in Section G.3 hold, except now we are concerned with the matrix

$$G_z^{(2)} := \sum_{k=0}^L \Psi_k [\mathbf{0}_{m \times Tk} \mid D_z \mid \mathbf{0}_{m \times T(L-k)-1}].$$

As such, we see that $\|G_{z,v}^{(2)}\|_{\text{F}} \leq \sqrt{N}\|v\|_2\|G_z^{(2)}\|_{\text{op}}$ and $\|G_{z,v}^{(2)}\|_{\text{op}} \lesssim \sqrt{L}\|v\|_2\|G_z^{(2)}\|_{\text{op}}$. Then,

$$\|G_z^{(2)}\|_{\text{op}} = \left\| \sum_{k=0}^L \Psi_k [\mathbf{0}_{m \times Tk} \mid D_z \mid \mathbf{0}_{m \times T(L-k)-1}] \right\|_{\text{op}} \leq (1 + \|\phi\|_{\text{bop}})\|D_z\|_{\text{op}},$$

so by Proposition G.3 we have

$$\begin{aligned} \|\mathbf{G}_z^{(2)}\mathbf{z}_{N:1}\|_{\text{op}} &\lesssim \left(\sqrt{N} + \sqrt{Lm} + \sqrt{L \log(1/\delta)} \right) (1 + \|\phi\|_{\text{bop}})\|D_z\|_{\text{op}} \\ &\lesssim \sqrt{N}(1 + \|\phi\|_{\text{bop}})\|D_z\|_{\text{op}}. \end{aligned}$$

For the adversarial case, we instead have

$$\|\mathbf{G}_z^{(2)}\mathbf{z}_{N:1}\|_{\text{op}} \lesssim \sqrt{NLd_z}(1 + \|\phi\|_{\text{bop}})\|D_z\|_{\text{op}}.$$

G.4.3. CONTRIBUTION OF THE INITIAL STATE

We see that the nonzero contribution of \mathbf{x}_1 to Δ_ϕ is given by

$$\mathbf{H}^{(1)}\mathbf{x}_1 = [C_\phi A^{N-TL-1}\mathbf{x}_1 \mid C_\phi A^{N-TL-2}\mathbf{x}_1 \mid \dots \mid C_\phi A^{N_1-TL-1}\mathbf{x}_1].$$

Thus, $\|\mathbf{H}^{(1)}\mathbf{x}_1\|_{\text{op}} \leq \|\mathcal{M}_N(\mathbf{H}_\phi)\|_{\text{op}}$.

Appendix H. Definition of \overline{M} , $\overline{M}_{\text{adv}}$, and Proof of Proposition 2.1

H.1. Notation

Note that in the body of the text we assumed $\mathbf{x}_1 = 0$. Here we allow $\mathbf{x}_1 \neq 0$, and opt for the more general definitions of \overline{M} . Specifically, we define the constants

$$\begin{aligned} M_0 &:= \|S^{-1}\mathbf{x}_1\| \\ M_B(t) &:= \|S^{-1}B_\star\| + \sqrt{t}\|S^{-1}B_w\|, \quad M_B := M_B(1) \\ M_C &:= \|C_\star S\|_{\text{op}} \\ M_D(t) &:= \|D_\star\|_{\text{op}} + \sqrt{t}\|D_z\|_{\text{op}}, \quad M_D := M_D(1) \end{aligned} \tag{H.26}$$

which inherently reflect the conditioning of the chosen realization of G_\star . We define the general version of \overline{M} and $\overline{M}_{\text{adv}}$ that take \mathbf{x}_1 into account.

Definition H.1 Let $A_\star = SJ_\star S^{-1}$ denote the Jordan-normal decomposition of A_\star . We let

$$\begin{aligned} \overline{M} &:= (N^{-1/2}M_0 + M_B)M_C + M_D \\ \overline{M}_{\text{adv}} &:= (N^{-1/2}M_0 + M_B(Tdd_w))M_C + M_D(dd_z). \end{aligned}$$

Note that the above definition reduces to the quantities used in the body of the paper in the case $\mathbf{x}_1 = 0$. We shall require the following bound, which is a corollary of Proposition I.2 and Lemma I.3 in the next section:

Corollary H.1 (Concrete Markov Bounds) Let A_\star have maximum Jordan block size k . Then, for any $n \geq 1$,

$$\|\mathcal{M}_n(\mathbf{G}_\star)\|_{\text{op}} \lesssim \|D_\star\|_{\text{op}} + \|S^{-1}B_\star\|_{\text{op}}\|C_\star S\|_{\text{op}} \cdot k^{1/2}n^{k-\frac{1}{2}},$$

where an analogous bound holds for \mathbf{F}_\star and \mathbf{H}_\star replacing $S^{-1}B_\star$ by $S^{-1}B_w$ and $S^{-1}\mathbf{x}_1$, respectively.

H.2. Proof of Proposition 2.1

We first begin with the following lemma, which gives a generic bound on \mathbf{d}_{eff} and Ovfit_μ in terms of the quantity

$$\widetilde{\text{Opt}} := \|\mathbf{Y}_+\|_{\text{op}} + \|G_\star\|_{\text{op}}, \text{ where } \mathbf{Y}_+ := [\mathbf{y}_N | \dots | \mathbf{y}_1].$$

We have the following bound, in terms of an intermediate error quantity $\overline{\mathbf{d}}_{\text{Opt}}$.

Lemma H.2 On $\mathcal{E}_{\overline{\mathcal{U}}}$, we have

$$\begin{aligned} \mathbf{d}_{\text{eff}}(\text{Opt}_\mu + \text{Ovfit}_\mu, Lm, \mu) &\lesssim \overline{\mathbf{d}}_{\text{Opt}} := \tilde{p} + Lm \log_+(\widetilde{\text{Opt}}) + Lm \log_+(\frac{\sqrt{NL}}{\mu^2}) \\ \text{Ovfit}_\mu &\lesssim \min\{N, \sqrt{T(\overline{\mathbf{d}}_{\text{Opt}} + \log \frac{1}{\delta})}\}\|G_\star\|_{\text{op}}. \end{aligned}$$

We defer the proof of the above lemma until Section H.3. Now, recall that by Theorem 3.3 with $\bar{L} \leftarrow Lm$, on the event $\mathcal{E}_{\bar{U}}$,

$$\|\widehat{G}_{\text{PF}} - G_\star\|_{\text{op}} \lesssim \frac{N^{-1/2}(\text{Opt}_\mu + \text{Ovfit}_\mu(\delta) + \mu)}{\sqrt{N}} \cdot \sqrt{T \left(\log \frac{1}{\delta} + \mathbf{d}_{\text{eff}}(\text{Opt}_\mu + \text{Ovfit}_\mu(\delta), \bar{L}, \mu) \right)}$$

with probability at least $1 - \delta - \delta_{\bar{U}}$. Combining with Lemma H.2, we have

$$\|\widehat{G}_{\text{PF}} - G_\star\|_{\text{op}} \lesssim \frac{(\text{Opt}_\mu + \sqrt{T(\bar{\mathbf{d}}_{\text{Opt}} + \log \frac{1}{\delta})} \|G_\star\|_{\text{op}} + \mu)}{\sqrt{N}} \cdot \sqrt{\frac{T \left(\log \frac{1}{\delta} + \bar{\mathbf{d}}_{\text{Opt}} \right)}{N}}.$$

It now suffices to show that with probability $1 - \delta$ (absorbing union bounds into $\log(1/\delta)$), that when $\mu \geq 1$,

$$\bar{\mathbf{d}}_{\text{Opt}} \lesssim \bar{\mathbf{d}} := \tilde{p} + Lm \left(\log_+ \bar{M} + k \log_+ N \right), \quad (\text{H.27})$$

where k is the largest Jordan block of A_\star . Indeed, with this inequality, we have

$$\begin{aligned} \|\widehat{G}_{\text{PF}} - G_\star\|_{\text{op}} &\lesssim \frac{(\text{Opt}_\mu + \sqrt{T(\bar{\mathbf{d}} + \log \frac{1}{\delta})} \|G_\star\|_{\text{op}} + \mu)}{\sqrt{N}} \cdot \sqrt{\frac{T \left(\log \frac{1}{\delta} + \bar{\mathbf{d}} \right)}{N}} \\ &= N^{-1/2} \text{Opt}_\mu \cdot \sqrt{\frac{T(\bar{\mathbf{d}} + \log \frac{1}{\delta})}{N}} + \frac{T(\|G_\star\|_{\text{op}} + \mu)(\bar{\mathbf{d}} + \log \frac{1}{\delta})}{N}. \end{aligned}$$

Proof of (H.27): For $\mu \geq 1$, using the bound $N \geq L$, we have

$$\begin{aligned} \bar{\mathbf{d}}_{\text{Opt}} &= \tilde{p} + Lm \log_+(\widetilde{\text{Opt}}) + Lm \log_+(\frac{\sqrt{NL}}{\mu^2}) \\ &\lesssim \tilde{p} + Lm \log_+(\widetilde{\text{Opt}}) + Lm \log_+ N. \end{aligned}$$

Thus, it remains to establish the bound

$$\log_+ \widetilde{\text{Opt}} \lesssim \begin{cases} \log_+ \bar{M} + k \log_+ N & (\text{stochastic noise}) \\ \log_+ \bar{M}_{\text{adv}} + k \log_+ N & (\text{adversarial noise}). \end{cases}$$

We begin this task by bounding the random part of $\widetilde{\text{Opt}}$, $\|\mathbf{Y}_+\|_{\text{op}}$. We will do so in terms of the block Toeplitz matrix

$$T_N(\mathbf{G}_\star) = \begin{bmatrix} D_\star & C_\star B_\star & C_\star A_\star B_\star & \dots & C_\star A_\star^{N-2} B_\star \\ \mathbf{0}_{d \times d} & D_\star & C_\star B_\star & \ddots & C_\star A_\star^{N-3} B_\star \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ \vdots & & \ddots & D_\star & C_\star B_\star \\ \mathbf{0}_{d \times d} & \dots & \dots & \mathbf{0}_{d \times d} & D_\star \end{bmatrix} \in \mathbb{R}^{Nm \times Np}$$

and its analog for $T_N(\mathbf{F}_\star)$, as well as $\mathcal{M}_N(\mathbf{H}_\star)$.

Proposition H.3 *Suppose that $N \geq \max\{m, \log(1/\delta)\}$. With probability at least $1 - \delta$, we have that*

(i) *(Stochastic model)*

$$\begin{aligned} \|\mathbf{Y}_+\|_{\text{op}} &\lesssim \sqrt{N} (\|T_N(\mathbf{G}_\star)\|_{\text{op}} + \|T_N(\mathbf{F}_\star)\|_{\text{op}} + \|D_z\|_{\text{op}}) + \|\mathcal{M}_{N+1}(\mathbf{H}_\star)\|_{\text{op}} \\ &\leq \sqrt{N} \left(\sqrt{N} \|\mathcal{M}_N(\mathbf{G}_\star)\|_{\text{op}} + \sqrt{N} \|\mathcal{M}_N(\mathbf{F}_\star)\|_{\text{op}} + \|D_z\|_{\text{op}} \right) + \|\mathcal{M}_{N+1}(\mathbf{H}_\star)\|_{\text{op}} \end{aligned}$$

(ii) *(Adversarial model) $\|\mathbf{Y}_+\|_{\text{op}}$ is bounded as*

$$\begin{aligned} &\lesssim \sqrt{N} \left(\|T_N(\mathbf{G}_\star)\|_{\text{op}} + \sqrt{d_w} \|T_N(\mathbf{F}_\star)\|_{\text{op}} + \sqrt{d_z} \|D_z\|_{\text{op}} \right) + \|\mathcal{M}_{N+1}(\mathbf{H}_\star)\|_{\text{op}} \\ &\leq \sqrt{N} \left(\sqrt{N} \|\mathcal{M}_N(\mathbf{G}_\star)\|_{\text{op}} + \sqrt{d_w} \sqrt{N} \|\mathcal{M}_N(\mathbf{F}_\star)\|_{\text{op}} + \sqrt{d_z} \|D_z\|_{\text{op}} \right) + \|\mathcal{M}_{N+1}(\mathbf{H}_\star)\|_{\text{op}} \end{aligned}$$

Proof The argument mirrors those in Appendix G. The dependence on \mathbf{H}_\star is through the Markov operator $\|\mathcal{M}_N(\mathbf{H}_\star)\|_{\text{op}}$ due to the argument given in Section G.4.3. For the other terms, we bound the \mathbf{G} term as a representative example. Letting $\mathbf{Y}^{(\mathbf{u})}$ being the operator that maps $\mathbf{u}_{N:1}$ to \mathbf{Y}_+ , the key is to again note that $\mathbf{Y}_v^{(\mathbf{u})} = (I \otimes v^\top) T_N(\mathbf{G}_\star)$. Thus, by the now-standard arguments, one can show that

$$\sup_{v \in \mathcal{S}^{m-1}} \|\mathbf{Y}_v^{(\mathbf{u})}\|_{\text{op}} \leq \|T_N(\mathbf{G}_\star)\|_{\text{op}} \quad \text{and} \quad \sup_{v \in \mathcal{S}^{m-1}} \|\mathbf{Y}_v^{(\mathbf{u})}\|_{\text{F}} \leq \sqrt{N} \|T_N(\mathbf{G}_\star)\|_{\text{op}}.$$

Proposition G.3 along with the simplification $N \geq \max\{m, \log(1/\delta)\}$ then gives the desired bounds. Bounding $\|T_N(\mathbf{G}_\star)\|_{\text{op}} \leq \sqrt{N} \|\mathcal{M}_N(\mathbf{G}_\star)\|_{\text{op}}$ follows by considering each row of T_N separately and then noting that each row of T_N is a submatrix of the first. ■

We now bound $\widetilde{\text{Opt}} = \|\mathbf{Y}_+\|_{\text{op}} + \sqrt{N} \|G_\star\|_{\text{op}}$ in the stochastic case using Proposition H.3 and Corollary H.1:

$$\begin{aligned} &\|\mathbf{Y}_+\|_{\text{op}} + \sqrt{N} \|G_\star\|_{\text{op}} \\ &\stackrel{\text{Prop. H.3}}{\lesssim} \sqrt{N} \left(\sqrt{N} \|\mathcal{M}_N(\mathbf{G}_\star)\|_{\text{op}} + \sqrt{N} \|\mathcal{M}_N(\mathbf{F}_\star)\|_{\text{op}} + \|D_z\|_{\text{op}} + \|G_\star\|_{\text{op}} \right) + \|\mathcal{M}_{N+1}(\mathbf{H}_\star)\|_{\text{op}} \\ &\stackrel{(i)}{\leq} \sqrt{N} \left(2\sqrt{N} \|\mathcal{M}_N(\mathbf{G}_\star)\|_{\text{op}} + \sqrt{N} \|\mathcal{M}_N(\mathbf{F}_\star)\|_{\text{op}} + \|D_z\|_{\text{op}} \right) + \|\mathcal{M}_{N+1}(\mathbf{H}_\star)\|_{\text{op}} \\ &\leq 2N \left(\|\mathcal{M}_N(\mathbf{G}_\star)\|_{\text{op}} + \sqrt{N} \|\mathcal{M}_N(\mathbf{F}_\star)\|_{\text{op}} + N^{-1/2} \|\mathcal{M}_{N+1}(\mathbf{H}_\star)\|_{\text{op}} + \|D_z\|_{\text{op}} \right) \\ &\stackrel{\text{Cor. H.1}}{\lesssim} 2N \left(M_C M_B N^{k-1/2} + N^{-1/2} \cdot M_0 M_C N^{k-1/2} + M_D \right) \\ &\leq 2N^{k+1/2} \left(M_C M_B + N^{-1/2} M_0 + M_D \right) \stackrel{\text{Defn. H.1}}{=} 2N^{k+1/2} \overline{M}, \end{aligned}$$

where (i) uses $\|G_\star\|_{\text{op}} = \|\mathcal{M}_T(\mathbf{G}_\star)\|_{\text{op}} \leq \|\mathcal{M}_N(\mathbf{G}_\star)\|_{\text{op}}$ as $N \geq T$, and where in the penultimate, last line M_1 , M_0 , and M_D are the constants in (H.26). Hence, we have that

$$\log_+ \widetilde{\text{Opt}} \lesssim \log_+ 2N^{2k+1} \overline{M} \lesssim k \log N + \log_+ \overline{M}$$

The adversarial case is analogous, where we replace $M_B \leftarrow M_B(d_w)$ and $M_D \leftarrow M_D(d_z)$, which yield the extra factors of $\sqrt{d_w}$ and $\sqrt{d_z}$.

Finally, the bound $\|G_\star\|_{\text{op}} = \|\mathcal{M}_T(\mathbf{G}_\star)\|_{\text{op}} \lesssim \overline{M} T^{k-1/2}$ follows immediately from Proposition I.2 and Lemma I.3.

H.3. Proof of Lemma H.2

Proof Recall the definition

$$\mathbf{d}_{\text{eff}}(\text{Opt}, \bar{L}, \mu) := \tilde{p} + m + \text{lil} \frac{\text{Opt}}{\mu} + \bar{L} \log_+(\text{Opt} + \frac{\sqrt{N} \|\mathbf{K}\|_{\text{op}}}{\mu^2}),$$

which we bound with $\text{Opt} \leftarrow \text{Opt}_\mu + \text{Ovfit}_\mu$ and $\bar{L} \leftarrow Lm$. Since $\text{Ovfit}_\mu \leq \sqrt{N} \|G_\star\|_{\text{op}}$, we can write

$$\begin{aligned} \text{Opt}_\mu + \text{Ovfit}_\mu &\leq \min_{\phi} \|\Delta - \phi \mathbf{K}\|_{\text{op}} + \mu \|\phi\|_{\text{op}} + \sqrt{N} \|G_\star\|_{\text{op}} \\ &\leq \|\Delta\|_{\text{op}} + \sqrt{N} \|G_\star\|_{\text{op}} \quad (\text{taking } \phi = 0) \\ &\leq \|\Delta + G_\star \bar{\mathbf{U}}\|_{\text{op}} + \|G_\star\|_{\text{op}} \|\bar{\mathbf{U}}\|_{\text{op}} + \sqrt{N} \|G_\star\|_{\text{op}} \\ &\lesssim \|\mathbf{Y}\|_{\text{op}} + \sqrt{N} \|G_\star\|_{\text{op}} \quad (\text{on } \mathcal{E}_{\bar{\mathbf{U}}}) \\ &\leq \|\mathbf{Y}_+\|_{\text{op}} + \sqrt{N} \|G_\star\|_{\text{op}} := \widetilde{\text{Opt}}. \end{aligned}$$

Moreover, since \mathbf{K} consists of L submatrices, each of which is a submatrix of \mathbf{Y}_+ , we see that $\|\mathbf{K}\|_{\text{op}} \leq \sqrt{L} \|\mathbf{Y}_+\|_{\text{op}}$. Lastly, recall $\log_+ x \geq 1$. Therefore, we can bound

$$\begin{aligned} \mathbf{d}_{\text{eff}}(\text{Opt}_\mu + \text{Ovfit}_\mu, Lm, \mu) &\lesssim \tilde{p} + \text{lil} \frac{\widetilde{\text{Opt}}}{\mu} + Lm \log_+(\widetilde{\text{Opt}} + \frac{\sqrt{N} \|\mathbf{K}\|_{\text{op}}}{\mu^2}) \quad (\text{absorbing } m) \\ &\leq \tilde{p} + \text{lil} \frac{\widetilde{\text{Opt}}}{\mu} + Lm \log_+(\widetilde{\text{Opt}} + \frac{\sqrt{NL} \|\mathbf{Y}_+\|_{\text{op}}}{\mu^2}) \\ &\leq \tilde{p} + \text{lil} \frac{\widetilde{\text{Opt}}}{\mu} + Lm \log_+(\widetilde{\text{Opt}} (1 + \frac{\sqrt{NL}}{\mu^2})) \\ &\lesssim \tilde{p} + \text{lil} \frac{\widetilde{\text{Opt}}}{\mu} + Lm \log_+(\widetilde{\text{Opt}}) + Lm \log_+(\frac{\sqrt{NL}}{\mu^2}). \end{aligned}$$

Lastly, we observe that since lil is submultiplicative, $\text{lil} \frac{\widetilde{\text{Opt}}}{\mu} \lesssim \text{lil} \widetilde{\text{Opt}} + \text{lil} \frac{1}{\mu} \leq \log_+ \widetilde{\text{Opt}} + \log_+ \frac{1}{\mu} \leq Lm (\log_+ \widetilde{\text{Opt}} + \log_+ \frac{1}{\mu})$. Partially absorbing this lil term, we find that

$$\mathbf{d}_{\text{eff}}(\text{Opt}_\mu + \text{Ovfit}_\mu, Lm, \mu) \lesssim \tilde{p} + Lm \log_+(\widetilde{\text{Opt}}) + Lm \log_+(\frac{\sqrt{NL}}{\mu^2}) + Lm \log_+ \frac{1}{\mu}.$$

Now, note that if $\mu \geq 1$, $Lm \log_+ \frac{1}{\mu} = Lm \leq Lm \log_+(\frac{\sqrt{NL}}{\mu^2})$. On the other hand, if $\mu \leq 1$, then $\frac{\sqrt{NL}}{\mu^2} \geq \frac{1}{\mu}$, and thus $Lm \log_+ \frac{1}{\mu} \leq Lm \log_+(\frac{\sqrt{NL}}{\mu^2})$. In either case, $Lm \log_+ \frac{1}{\mu} \leq Lm \log_+(\frac{\sqrt{NL}}{\mu^2})$, so

$$\mathbf{d}_{\text{eff}}(\text{Opt}_\mu + \text{Ovfit}_\mu, Lm, \mu) \lesssim \tilde{p} + Lm \log_+(\widetilde{\text{Opt}}) + Lm \log_+(\frac{\sqrt{NL}}{\mu^2}).$$

For the second bound of Lemma H.2, we use the fact that if Λ is a matrix of rank Lm ,

$$\log \det(I + \Lambda)^{1/2} = \sum_{i=1}^{Lm} \log \sqrt{1 + \lambda_i(\Lambda)} \leq \sum_{i=1}^{Lm} \log \sqrt{1 + \|\Lambda\|_{\text{op}}} \lesssim Lm \log_+ \|\Lambda\|_{\text{op}}.$$

Applying the above equation with $\Lambda = \mu^{-2}\mathbf{K}\mathbf{K}^\top$, we find that

$$\begin{aligned}
 \text{Ovfit}_\mu(\delta) &:= \|G_\star\|_{\text{op}} \cdot \min \left\{ \sqrt{N}, \sqrt{T} \sqrt{\log \frac{1}{\delta} + \tilde{p} + \log \det(I + \mu^{-2}\mathbf{K}\mathbf{K}^\top)^{1/2}} \right\} \\
 &\lesssim \|G_\star\|_{\text{op}} \cdot \min \left\{ \sqrt{N}, \sqrt{T} \sqrt{\log \frac{1}{\delta} + \tilde{p} + Lm \log_+ \frac{\|\mathbf{K}\|_{\text{op}}^2}{\mu^2}} \right\} \\
 &\leq \|G_\star\|_{\text{op}} \cdot \min \left\{ \sqrt{N}, \sqrt{T} \sqrt{\log \frac{1}{\delta} + \tilde{p} + Lm \log_+ \frac{L\|\mathbf{Y}_+\|_{\text{op}}^2}{\mu^2}} \right\} \\
 &\leq \|G_\star\|_{\text{op}} \cdot \min \left\{ \sqrt{N}, \sqrt{T} \sqrt{\log \frac{1}{\delta} + \tilde{p} + Lm \log_+ \|\mathbf{Y}_+\|_{\text{op}}^2 + Lm \log_+ \frac{L}{\mu^2}} \right\},
 \end{aligned}$$

from which the result follows by taking $Lm \log_+ \|\mathbf{Y}_+\|_{\text{op}}^2 \lesssim Lm \log_+ \|\mathbf{Y}_+\|_{\text{op}} \leq Lm \log_+ \widehat{\text{Opt}}_\mu$ and $Lm \log_+ \frac{L}{\mu^2} \leq Lm \log_+ \frac{\sqrt{NL}}{\mu^2}$, as $N \geq L$. \blacksquare

H.4. Selecting the parameter L

In this section, we give an informal discussion of how to select the parameter L . Observe that the confidence bounds from Theorem 3.3 are *almost* data-dependent, but in fact depends on the quantity $\text{Opt}_\mu + \text{Ovfit}_\mu$, which is not known to the learner. In order to select L , one shall need to replace these quantities with data-dependent ones, and then use a standard procedure (e.g. structural risk minimization) to tune L . First, considering L fixed, define the following empirical proxy for $\text{Opt}_\mu + \text{Ovfit}_\mu$,

$$\widehat{\text{Opt}}_\mu := \|\mathbf{Y} - \mathbf{K}\phi_{\text{rdg}}\|_{\text{op}} + \mu\|\phi_{\text{rdg}}\|_{\text{op}}.$$

Our first main result of this section is that

$$\widehat{\text{Opt}}_\mu \gtrsim \|G_\star\|_{\text{op}}\sqrt{N} + \text{Opt}_\mu$$

with high probability. Formally, our guarantee is

Proposition H.4 *For constants C_1, C_2 , suppose that there exists an $K \in \{e, e^2, \dots\}$ such that L and N satisfy*

- (1) $N \geq C_1 T(\tilde{p} + m + LmK + \log 1/\delta)$
- (2) $K \geq \log \mu\|\phi_{\text{rdg}}\|_{\text{op}}$, and either
- (3) (a): $K \geq \log \frac{C_2\|\mathbf{K}\|_{\text{op}}}{\mu\sqrt{N}\|G_\star\|_{\text{op}}}$ or (b): $K \geq \log \frac{C_2\|\mathbf{K}\|_{\text{op}}}{\mu\widehat{\text{Opt}}_\mu}$.

Then, with probability $1 - 2\delta$, whenever $\mathcal{E}_{\overline{\mathbf{U}}}$ holds,

$$\widehat{\text{Opt}}_\mu \gtrsim \text{Opt}_\mu + \|G_\star\|_{\text{op}}\sqrt{N}.$$

Note that condition (3a) gives a condition which is more amenable to analysis (requiring only a lower bound on $\|G_\star\|_{\text{op}}$), whereas condition (3b) can be evaluated by the learner, as \mathbf{K} and $\widehat{\text{Opt}}$ are empirical quantities. We sketch the proof of the above result in the following subsection. Now, note that by the triangle inequality, we also have

$$\widehat{\text{Opt}}_\mu \lesssim \|G_\star\|_{\text{op}}\sqrt{N} + \text{Opt}_\mu,$$

and moreover, $\|G_\star\|_{\text{op}}\sqrt{N} \leq \text{Ovfit}_\mu$ by definition. Hence we can use structural risk minimization [Shawe-Taylor et al. \(1998\)](#) to select L .

To sketch this approach, let $\widehat{\text{Opt}}_{\mu;L}$, \mathbf{K}_L , Opt_μ^L , $\widehat{G}_{\text{PF}}^L$, and $N_{\min}(L)$ be the corresponding quantities defined for a given L . We define the set

$$\mathcal{S}(\delta) := \left\{ L \in \mathbb{N} : C_1 T (\tilde{p} + m + Lm \left(\log_+ \frac{C_2 \|\mathbf{K}_L\|_{\text{op}}}{\mu \widehat{\text{Opt}}_L} + \log_+ \mu \|\phi_{\text{rdg}}^L\|_{\text{op}} \right) + \log \frac{L}{\delta}) \leq N \right\},$$

which represents the set of admissible lengths L for which [Proposition H.4](#) guarantees that $\widehat{\text{Opt}}_{\mu;L}$ is a good proxy for $\text{Opt}_\mu + \text{Ovfit}_\mu$. We may then select \widehat{L} as

$$\widehat{L} \in \arg \min \{ \text{Conf}(L, \mu) : L \in \mathcal{S}(\delta) \},$$

where we have defined the upper confidence bound

$$\text{Conf}(L, \mu) := \frac{N^{-1/2}(\widehat{\text{Opt}}_{\mu;L} + \mu)}{\sqrt{N}} \cdot \sqrt{T \left(\log \frac{1}{\delta} + \mathbf{d}_{\text{eff}}(\widehat{\text{Opt}}_{\mu;L}, Lm, \mu) \right)}.$$

We can briefly analyze the outcome using a sketch of arguments similar to those in [Section H](#). Denote the L -indexed dimension quantity from [Proposition 2.1](#),

$$\bar{\mathbf{d}}(L) := \tilde{p} + Lm \left(\log_+ \bar{M} + k \log_+ N + \log_+ \right) = \tilde{\mathcal{O}}(p + Lmk).$$

It can be shown that for an appropriate constant C_2 , with probability $1 - \delta$ on $\mathcal{E}_{\overline{\text{T}}}$, it holds that

$$\left\{ L : N \geq C_2 \left(T(\bar{\mathbf{d}}(L) + Lm \log_+ \|G_\star\|_{\text{op}}^{-1} + \log \frac{1}{\delta}) \right) \right\} \subseteq \mathcal{S}(\delta).$$

This can in turn be used to establish the following analogue of [Proposition 2.1](#), whose proof we omit.

Proposition H.5 *Fix a $\delta \in (0, 1)$, and $T, L_{\max} \in \mathbb{N}$. Suppose that $N \geq N_{\min}(L_{\max})$, $N_1 = TL_{\max}$, $\rho(A_\star) \leq 1$, and that the largest Jordan block of A_\star is of size k . Then, once*

$$N \geq C_2 \left(T(\bar{\mathbf{d}}(L_{\max}) + L_{\max} m \log_+ \|G_\star\|_{\text{op}}^{-1} + \log \frac{1}{\delta}) \right),$$

the estimator $\widehat{G}_{\text{PF}}^{\widehat{L}}$, where \widehat{L} is selected in the manner described above, satisfies with probability at least $1 - \delta - (2Np)^{-\log^2(2Tp)} \log^2(2Np)$ in the stochastic noise model that

$$\|\widehat{G}_{\text{PF}}^{\widehat{L}} - G_\star\|_{\text{op}} \lesssim \min_{L \in \{0, 1, \dots, L_{\max}\}} \left(\frac{\text{Opt}_\mu^L + \sqrt{N} \|G_\star\|_{\text{op}} + \mu}{\sqrt{N}} \right) \cdot \sqrt{\frac{T(\bar{\mathbf{d}}(L) + \log \frac{L}{\delta})}{N}}.$$

In the adversarial noise model, we instead take $\bar{\mathbf{d}}(L) := \tilde{p} + Lm \left(\log_+ \bar{M} + \log_+(d_z + d_w) + k \log_+ N \right)$.

We remark that the parameter L_{\max} in the above proposition appears merely in the analysis. Moreover, one can also search L in powers binary powers 2^i for added computational efficiency.

H.4.1. PROOF SKETCH OF PROPOSITION H.4

By the reverse triangle inequality,

$$\widehat{\text{Opt}}_\mu = \|\mathbf{Y} - \mathbf{K}\widehat{\phi}\|_{\text{op}} + \mu\|\widehat{\phi}\|_{\text{op}} \geq \|\Delta - \mathbf{K}\widehat{\phi}\|_{\text{op}} + \mu\|\widehat{\phi}\|_{\text{op}} - \|\mathbf{G}_\star\|_{\text{op}} \geq \text{Opt}_\mu - \|\mathbf{G}_\star\|_{\text{op}}.$$

Thus if $\widehat{\text{Opt}}_\mu \leq \text{Opt}_\mu + \|\mathbf{G}_\star\|_{\text{op}}\sqrt{N}$, then $\text{Opt}_\mu \lesssim \|\mathbf{G}_\star\|_{\text{op}}\sqrt{N}$. Hence, it suffices to show that $\widehat{\text{Opt}}_\mu \gtrsim \|\mathbf{G}_\star\|_{\text{op}}\sqrt{N}$ with probability $1 - \delta$ on $\mathcal{E}_{\overline{\mathcal{U}}}$. Moreover, it suffices to prove the theorem if (3a) holds; indeed, if (3b) holds, then either $\widehat{\text{Opt}} \leq \|\mathbf{G}_\star\|_{\text{op}}\sqrt{N}$, in which case (3a) holds, or $\widehat{\text{Opt}} \geq \|\mathbf{G}_\star\|_{\text{op}}\sqrt{N}$ as desired.

Fix a $v \in \mathcal{S}^{m-1}$ for which $\|\mathbf{G}_\star\|_{\text{op}} = \|v^\top \mathbf{G}_\star\|_2$. In this simplified setting, we show the following lemma.

Lemma H.6 *Fix $v \in \mathcal{S}^{m-1}$. Then with probability $1 - \delta$, for any $j \in \mathbb{N}$ satisfying*

$$e^j \geq \log \left(\frac{8\|\mathbf{K}\|_{\text{op}}}{\mu\|\mathbf{G}_\star v\|_2} \right) \quad \text{and} \quad e^j \geq \log \mu \|\phi_{\text{rdg}} v\|_2,$$

where $N \gtrsim T(\tilde{p} + m + Lme^j + \log 1/\delta)$, then $\|(\mathbf{Y} - \mathbf{K}\phi_{\text{rdg}})v\|_2 \geq \frac{1}{4}\|v^\top \mathbf{G}_\star\|_2$.

This implies Proposition H.4, since $e^j \geq \log \mu \|\phi_{\text{rdg}}\|_{\text{op}}$ implies $e^j \geq \log \mu \|\phi_{\text{rdg}}^\top v\|_2$, and v was chosen so that $\|v^\top \mathbf{G}_\star\|_2 = \|\mathbf{G}_\star\|_{\text{op}}$. We now prove the above lemma:

Proof We shall use the following intermediate lemma, which we prove following the proof of Lemma H.6.

Lemma H.7 *Suppose $m = 1$. Then for any fixed ϕ and any $N \gtrsim T(\tilde{p} + \log \frac{1}{\delta})$, it holds with probability at least $1 - \delta$ on $\mathcal{E}_{\overline{\mathcal{U}}}$ that $\|(\mathbf{Y} - \mathbf{K}\phi)v\|_2 \geq \frac{1}{2\sqrt{2}}\|\mathbf{G}_\star\|\sqrt{N}$.*

Next, we mirror the proof of Theorem D.1, For $j \geq 1$, let $c_j = e^{e^j}$ and let \mathcal{T}_j denote a $1/c_j$ net of the set $c_j \mathcal{B}_{\mathbb{R}^{L_{\text{eff}}}}/\mu$ in the norm $\|\cdot\|_2$. Following the computations in that proof, it holds with probability at least $1 - \delta$ that, for all j which satisfy

$$N \gtrsim T(\tilde{p} + Lme^j + \log 1/\delta),$$

it holds that for all $\tilde{\phi} \in \mathcal{T}_j$ that

$$\|\mathbf{Y} - \mathbf{K}\tilde{\phi}^\top\| \geq \frac{1}{2\sqrt{2}}\|\mathbf{G}_\star\|\sqrt{N}.$$

Hence, if for some $\Delta > 0$ we have

$$e^j \geq \log \left(\frac{\|\mu^{-1}\mathbf{K}\|_{\text{op}}}{\Delta} \right),$$

and if $\phi_{\text{rdg}} \in c_j \mathcal{B}_{\mathbb{R}^{L_{\text{eff}}}}/\mu$, we have with probability $1 - \delta$ on $\mathcal{E}_{\overline{\mathcal{U}}}$,

$$\begin{aligned} \|\mathbf{Y} - \mathbf{K}\tilde{\phi}\| &\geq \|\mathbf{Y} - \mathbf{K}\tilde{\phi}^\top\| - \|\mathbf{K}\|_{\text{op}}\|\phi_{\text{rdg}} - \tilde{\phi}\|_2 \\ &\geq \|\mathbf{Y} - \mathbf{K}\tilde{\phi}^\top\| - \Delta \geq \frac{\|\mathbf{G}_\star\|\sqrt{N}}{2\sqrt{2}} - \Delta. \end{aligned}$$

Finally, setting $\Delta = \frac{\|G_\star\|\sqrt{N}}{8}$ yields the desired bound of Proposition H.4. \blacksquare

Proof [Proof of Lemma H.7] Since $m = 1$, we work with the 2-norm. Here,

$$\begin{aligned} \|\mathbf{Y} - \mathbf{K}_L\phi^\top\|_2^2 &= \|\Delta + \bar{\mathbf{U}}G_\star^\top - \mathbf{K}\phi^\top\|_2^2 \\ &= \|\bar{\mathbf{U}}G_\star^\top\|_2^2 + \|\Delta - \mathbf{K}\phi^\top\|_2^2 + 2\langle \bar{\mathbf{U}}G_\star^\top, \Delta - \mathbf{K}\phi^\top \rangle \\ &= \frac{1}{2}\|\mathbf{U}G_\star^\top\|_2^2 + \frac{1}{2}\|\Delta_\phi\|_2^2 \quad (:= \frac{1}{2}T_1) \\ &\quad + \frac{1}{2}\|\bar{\mathbf{U}}G_\star^\top\|_2^2 + \frac{1}{2}\|\Delta_\phi\|_2^2 + 2\langle \bar{\mathbf{U}}G_\star^\top, \Delta_\phi \rangle^2 \quad (:= T_2). \end{aligned}$$

We shall show that with probability $1 - \delta$ on $\mathcal{E}_{\bar{\mathbf{U}}}$, the term T_2 is nonnegative. This suffices since on $\mathcal{E}_{\bar{\mathbf{U}}}$, $\frac{1}{2}T_1 \geq \frac{1}{4}(\|G_\star\|_2\sqrt{N})^2$. To show that T_2 is nonnegative with high probability, we may assume that $G_\star \neq 0$, for otherwise this holds trivially. On $\mathcal{E}_{\bar{\mathbf{U}}}$, we have that $\|\bar{\mathbf{U}}G_\star^\top\|_2 \geq \frac{1}{\sqrt{2}}\sqrt{N}\|G_\star\| := \alpha$. Now, set $\beta = \|\Delta_\phi\|_2$. We may also assume that $\beta \leq 16\alpha$, since otherwise

$$\frac{1}{2}\|\bar{\mathbf{U}}G_\star^\top\|_2^2 + \frac{1}{2}\|\Delta_\phi\|_2^2 + 2\langle \mathbf{U}G_\star^\top, \Delta_\phi \rangle^2 \geq \frac{\alpha^2}{2} + \frac{\beta^2}{2} - \beta\|G_\star\|\|\mathbf{U}\|_{\text{op}} \geq \frac{\alpha^2}{2} + \frac{\beta^2}{2} - 2\beta\alpha \quad (\text{on } \mathcal{E}_{\bar{\mathbf{U}}}),$$

which is nonnegative for $\beta \geq 16\alpha$. Next, by Theorem E.1 with $\kappa \leftarrow \alpha$, $\Delta \leftarrow \Delta_\phi$, $\|\Delta_\phi\|_{\text{op}} \leftarrow \beta$, and $m = 1$, we have with probability $1 - \delta$ that

$$\|\Delta^\top \mathbf{U}\|_2 \lesssim T^{1/2}(\beta + \alpha)\sqrt{\tilde{p} + \log \frac{1}{\delta} + \text{lil}(\frac{\beta}{\alpha})}.$$

Using the fact that $\beta \leq 16\alpha$, that the above bound is at most $CT^{1/2}\alpha\sqrt{\tilde{p} + \log \frac{1}{\delta}}$ for a universal constant C . Noting that $\|G_\star\| = \sqrt{2/N}\alpha$, we find that with probability $1 - \delta$ on $\mathcal{E}_{\bar{\mathbf{U}}}$,

$$\begin{aligned} T_2(v) &\geq \frac{\alpha^2}{2} - 2\langle \mathbf{U}G_\star^\top, \Delta_\phi \rangle^2 \geq \frac{\alpha^2}{2} - \|G_\star\|_2\|\Delta_\phi^\top \mathbf{U}\|_2 \\ &\geq \frac{\alpha^2}{2} - \sqrt{2/N}\alpha \cdot CT^{1/2}\alpha\sqrt{\tilde{p} + \log \frac{1}{\delta}} = \alpha^2 \left(\frac{1}{2} - C\sqrt{\frac{2T(\tilde{p} + \log \frac{1}{\delta})}{N}} \right), \end{aligned}$$

which is nonnegative as soon as $N \geq C'T(\tilde{p} + \log \frac{1}{\delta})$ for some universal constant C' . \blacksquare

Appendix I. Polynomial Approximations and Phase Rank

In this section, we present demonstrate how to bound Opt_μ using the (α, T) phase rank of A_\star . Our bounds will be in terms of the $M_{(\cdot)}$ -constants defined in Section H.1 above.

1. Section I.1 presents our main findings in terms of two types quantities: $K_1(\dots)$ captures the ‘‘complexity’’ of a polynomial required to cancel out large dynamical modes of A_\star (formalized in Proposition I.1), and $K_2(\cdot)$ describes the growth rate of finite-length Markov parameter matrices (Proposition I.2). We also present findings based on worst-case upper bounds on $K_1(\dots)$ and $K_2(\cdot)$, via Lemma I.3. Specifically,

- (a) Theorems [I.4](#) and [I.5](#) present refined bounds for the stochastic and adversarial noise models respectively, which are simplified into Proposition [2.2](#) for stochastic noise, and Corollary [I.6](#) for adversarial noise.
 - (b) We also derive Theorem [I.8](#), and a simplified consequence Corollary [I.9](#), under the condition that the modes of A_\star can be “disentangled”, even when observed through the matrix C_\star .
2. Section [I.2](#) introduces the main technical tools which . The idea is to show the existence of filters ϕ for control-theoretic norms of the systems $\mathbf{G}_\phi, \mathbf{F}_\phi, \mathbf{H}_\phi$ defined in [\(G.24\)](#) can be bounded by a quantity $H_f^{(q)}(\dots)$, which roughly describes how well a polynomial of bounded degree and coefficient magnitude can “cover” a certain set of poles in the complex plane. We focus on scalar filters (Section [I.2.1](#)), and discuss possibly sharper bounds for richer, non-scalar filters (Section [I.2.2](#)). Section [I.2.1](#) also includes Proposition [I.13](#), which bounds the covering-like quantity $H(\dots)$ with the more transparent $K_1(\alpha, d, T)$.
 3. Finally, Section [I.3](#) gives a proof of Theorem [I.4](#) for stochastic noise; this proof only depends on results stated in the first section of this appendix, [I.1](#).

With the exception of the proof Theorem [I.4](#), all further proofs are deferred to Appendix [J](#). Lastly, again, we recall the definition of phase rank:

Definition 2.3 (Phase Rank) *Let $\alpha \geq 1$. We say that A_\star has (α, T) -phase rank d if there exists $\mu_1, \dots, \mu_d \in \mathbb{D}$ such that, for any $(\lambda, k) \in \text{blkspec}(A_\star)$ with $|\lambda| \geq 1 - ((1 + \alpha)T)^{-1}$, there exists at least k elements $\mu_{i_1}, \dots, \mu_{i_k} \subset \{\mu_1, \dots, \mu_d\}$ satisfying*

$$\max_{j \in [k]} \min_{\tilde{\mu}: \tilde{\mu}^T = \mu_{i_j}^T} |\lambda - \tilde{\mu}| \leq \alpha(1 - |\lambda|).$$

I.1. Main Results

We start by presenting our main results for the stochastic case. Let $A_\star = SJ_\star S^{-1}$ denote the Jordan decomposition of A_\star . Since $\text{blkspec}(A_\star) = \text{blkspec}(J_\star)$, we shall use the two interchangeably.

I.1.1. K_1 AND K_2 : CONTROLLING POLES AND MARKOV OPERATOR NORMS

We begin by introducing two central quantities. First, we introduce a term $K_1(d, T, \alpha, q)$ which reflects how well a d -length linear filter can predict observations of the Jordan-normal linear system J_\star when it has (α, T) phase rank d . Here, prediction is defined by the $\|\mathcal{M}_\infty(\cdot)\|_{\text{op}}$ and \mathcal{H}_∞ norms, indexed by $q \in \{2, \infty\}$, respectively. Formally, we define

$$K_1(d, T, \alpha, q) := \max_{(\lambda, k) \in \text{blkspec}(A_\star)} k^2 c_{\mathcal{H}_q} \begin{cases} 0 & |\lambda| = 1 \\ (T(1 + \alpha))^{k - \frac{\mathbb{I}(q=2)}{2}} 2^{d-k} & |\lambda| \in (1 - \frac{1}{T(1+\alpha)}, 1] \\ \frac{2^d}{(1-|\lambda|)^{k - \frac{\mathbb{I}(q=2)}{2}}} & |\lambda| < 1 - \frac{1}{T(1+\alpha)} \end{cases},$$

$$\text{where } c_{\mathcal{H}_\infty} := 1 \quad \text{and} \quad c_{\mathcal{H}_2} := \sqrt{1 + \frac{2}{\pi}}.$$

Taking $q = \infty$ correspond to the \mathcal{H}_∞ -norm, whereas taking $q = 2$ corresponds to the norm $\|\mathcal{M}_\infty(\cdot)\|_{\text{op}}$. This notation is because $\|\mathcal{M}_\infty(\mathbf{G}_\star)\|_{\text{op}}$ can be rendered as a norm we call $\mathcal{H}_2^{\text{op}}$, defined in (I.32), on the transfer function $\mathbf{G}(z)$. This norm is similar to but slightly sharper than the standard \mathcal{H}_2 -norm in control theory; see Chapter 4 of Zhou et al. (1996) for a discussion on transfer function norms, including \mathcal{H}_2 . In Section I.2.1, we then show the following bound.

Proposition I.1 *Suppose A_\star has (α, T) -phase rank d . Then, there exists a filter $\phi \in \mathbb{R}^{m \times dm}$ with $1 + \|\phi\|_{\text{bop}} \leq 2^d$ such that*

$$\begin{aligned} \|\mathcal{M}_\infty(\mathbf{G}_\phi)\|_{\text{op}} &\leq \|S^{-1}B_\star\|_{\text{op}} \cdot \|C_\star S\|_{\text{op}} K_1(d, T, \alpha, 2) \\ \|\mathbf{G}_\phi\|_{\mathcal{H}_\infty} &\leq \|S^{-1}B_\star\|_{\text{op}} \cdot \|C_\star S\|_{\text{op}} K_1(d, T, \alpha, \infty). \end{aligned}$$

Analogous bounds hold for \mathbf{F}_ϕ and \mathbf{H}_ϕ where $\|S^{-1}B_\star\|_{\text{op}}$ is replaced by $\|S^{-1}B_w\|_{\text{op}}$ and $\|S^{-1}\mathbf{x}_1\|_{\text{op}}$, respectively.

Second, we (somewhat tediously) define a term $K_2(N)$,

$$\widetilde{M}(k, N) := \begin{cases} N^{1/2} & k = 1 \\ N^{k-1/2} \left(\frac{e}{k-1}\right)^{k-1} & 2 \leq k \leq N+1 \\ N^{1/2} 2^N & k \geq N+1 \end{cases} \quad (\text{I.28})$$

$$M(k, \lambda, N) = \begin{cases} \frac{k}{(1-|\lambda|)^{k-\frac{1}{2}}} \wedge \widetilde{M}(k, N) & 0 \leq |\lambda| < 1 \\ \widetilde{M}(k, N) & |\lambda| = 1, \end{cases} \quad (\text{I.29})$$

$$K_2(N) := \max_{(\lambda, k) \in \text{blkspec}(A_\star)} M(k, \lambda, N). \quad (\text{I.30})$$

The term $K_2(N)$ describes the entire magnitude of an length- N trajectory generated by the Jordan-normal linear system J_\star . Indeed, in Section J.2, we prove

Proposition I.2 (Bound on Magnitude of Markov Parameters) *Consider a dynamical system of the form $\mathbf{G} = (A_\star, B, C, D)$, where $A_\star = S J_\star S^{-1}$ is in Jordan normal form. Then, for all $n \geq 1$,*

$$\|\mathcal{M}_n(\mathbf{G})\|_{\text{op}} \leq \|\mathcal{M}_{n+1}(\mathbf{G})\|_{\text{op}} \leq \|D\|_{\text{op}} + \|S^{-1}B\|_{\text{op}} \|CS\|_{\text{op}} K_2(n).$$

It is immediate to then check that K_1 and K_2 admit the following worst case bounds:

Lemma I.3 (Worst-Case Bounds on K_1, K_2) *Suppose that A_\star has $\rho(A_\star) \leq 1$, and has largest Jordan block of size k . Then,*

$$K_2(n) \leq en^{k-\frac{1}{2}}.$$

If further A_\star has (α, T) -phase rank at most $d \geq k$ for some $\alpha \geq 1$, then for $n \geq 1$,

$$K_1(d, T, \alpha, q) \lesssim k^2 (T(1+\alpha))^{k-\frac{1(q-2)}{2}} 2^d.$$

I.1.2. MAIN RESULTS: STOCHASTIC NOISE WITH BLOCK-SCALAR FILTERS

We begin by presenting bounds on Opt_μ for stochastic noise that arise from considering block-scalar filters of the form

$$\tilde{\phi} = [\phi \mid \mathbf{0}] \in \mathbb{R}^{m \times Lm}, \text{ where } \phi = [f_1 I_m \mid f_2 I_m \mid \dots \mid f_d I_m] \in \mathbb{R}^{m \times dm}.$$

We state two bounds: first, a theorem in terms of the more precise bounds $K_1(N)$ and $K_2(N)$, and then a corollary which applies the bounds from Lemma I.3, which is proved in Section I.3:

Theorem I.4 (Bounds for Stochastic Noise) *Suppose that A_\star has (α, T) phase rank at most $1 \leq d \leq L$. Then, for any $\delta \in (0, 1)$ and $N \geq Td \max\{m, \log(1/\delta)\}$, it holds with probability $1 - \delta$ that*

$$\begin{aligned} N^{-1/2} \text{Opt}_\mu &\lesssim M_C(M_B + N^{-1/2}M_0)K_1(d, T, \alpha, 2) \\ &\quad + \min\{K_1(d, T, \alpha, 2), N^{-1/2}K_1(d, T, \alpha, \infty)\}M_C M_B \sqrt{m + \log(1/\delta)} \\ &\quad + 2^d \left(M_C M_B K_2(Td) + M_D + \mu N^{-1/2} \right). \end{aligned}$$

By replacing the above bounds with worst case bounds from Lemma I.3, we obtain the bound

$$\begin{aligned} N^{-1/2} \text{Opt}_\mu &\lesssim M_C(M_B + N^{-1/2}M_0)k^2(T(1 + \alpha))^{k-\frac{1}{2}}2^d \\ &\quad + N^{-1/2}k^2(T(1 + \alpha))^{k-1}2^d M_C M_B \sqrt{m + \log(1/\delta)} \\ &\quad + 2^d \left(M_C M_B (Td)^{k-\frac{1}{2}} + M_D + \mu N^{-1/2} \right). \end{aligned}$$

For $N \geq T(1 + \alpha) \max\{m, \log(1/\delta)\}$, we can absorb the second line into the first term. This yields that $N^{-1/2} \text{Opt}_\mu$ is bounded by \lesssim

$$\begin{aligned} &M_C(M_B + N^{-1/2}M_0)k^2(T(1 + \alpha))^{k-\frac{1}{2}}2^d + 2^d \left(M_C M_B (Td)^{k-\frac{1}{2}} + M_D + \mu N^{-1/2} \right) \\ &\leq 2^d T^{k-\frac{1}{2}} \left(k^2(1 + \alpha)^{k-\frac{1}{2}} + d^{k-\frac{1}{2}} \right) (\bar{M} + \mu N^{-1/2}), \end{aligned}$$

from which we directly obtain Proposition 2.2 as stated in the body of the paper, which we restate here for convenience.

Proposition 2.2 (Bounds for Phase Rank) *Suppose that A_\star has (α, T) phase rank d and maximum Jordan block size k . Then, for any $\delta \in (0, 1)$; with $L \geq d$, $N \geq T(d + 1 + \alpha) \max\{m, \log(1/\delta)\}$, and $N_1 \geq TL$; it holds with probability $1 - \delta$ under the stochastic noise setting of Definition 2.1 that, with $C_{\alpha, d, k} := 2^d(k^2(1 + \alpha)^{k-\frac{1}{2}} + d^{k-\frac{1}{2}})$,*

$$N^{-1/2} \text{Opt}_\mu \lesssim (\bar{M} + \mu N^{-1/2}) \cdot T^{k-1/2} C_{\alpha, d, k}.$$

I.1.3. RESULTS FOR ADVERSARIAL NOISE

We now present the analogue of Theorem I.4 for adversarial noise; the proof is essentially identical, and omitted in the interest of brevity:

Theorem I.5 (Bounds for Adversarial Noise) *In the setting of Theorem I.4 (with the adversarial noise model), we have that*

$$\begin{aligned} N^{-1/2}\text{Opt}_\mu &\lesssim M_C(\|S^{-1}B_\star\|_{\text{op}} + N^{-1/2}M_0)K_1(d, T, \alpha, 2) \\ &\quad + M_C\|S^{-1}B_\star\|_{\text{op}} \cdot \min\{K_1(d, T, \alpha, 2), N^{-1/2}K_1(d, T, \alpha, \infty)\} \\ &\quad + M_C\|S^{-1}B_w\|_{\text{op}}K_1(d, T, \alpha, \infty)\sqrt{d_w} \\ &\quad + 2^d \left(M_C M_B(Tdd_w)K_2(Td) + M_D(dd_z) + \mu N^{-1/2} \right), \end{aligned}$$

where $M_B(\cdot)$ and $M_D(\cdot)$ are as in (H.26).

Recalling the definition of definition of $\overline{M}_{\text{adv}} := (N^{-1/2}M_0 + M_B(Tdd_w))M_C + M_D(dd_z)$ from Definition H.1, we obtain the following analogue of Proposition 2.2:

Corollary I.6 *In the setting of the previous theorem, where A_\star has (α, T) phase rank d , and maximum Jordan block size k , we have that*

$$\begin{aligned} N^{-1/2}\text{Opt}_\mu &\leq (\overline{M}_{\text{adv}} + \mu N^{-1/2}) \cdot T^k C_{\alpha, d, k}^{\text{adv}}, \quad \text{where} \\ C_{\alpha, d, k}^{\text{adv}} &:= 2^d \left(\frac{k^2}{2^k} (1 + \alpha)^k + d^k \right). \end{aligned}$$

I.1.4. BOUNDS FOR DISENTANGLING FILTERS

For the case when, after a similarity transformation, the invariant subspaces of A_\star can be decomposed onto the rows of C_\star , we can construct individual filters for each element in the decomposition. This is a generalization of the often-studied case (e.g. Sarkar and Rakhlin (2018)) of full-state observation: after transformation, we can observe each mode of A_\star directly. We begin by describing partitions of its associated Jordan matrix J_\star into invariant subspaces.

Definition I.1 (Admissible Spectral Partition) *Let $J \in \mathbb{R}^{n \times n}$ be a matrix in Jordan normal form. We say that a set $\mathcal{S}_{1:r} := \{\mathcal{S}_1, \dots, \mathcal{S}_r\} \subset [n]$ is an admissible spectral partition if, for each $i \in [r]$, the matrix $J(\mathcal{S}_i) := (J_{ab})_{ab \in \mathcal{S}_i \times \mathcal{S}_i} \in \mathbb{C}^{|\mathcal{S}_i| \times |\mathcal{S}_i|}$ is a Jordan matrix.*

In other words, $\mathcal{S}_{1:r} \subset [n]$ is an admissible spectral partition if each \mathcal{S}_i corresponds to coordinates indexing a J -invariant subspace of \mathbb{C}^n .

Next, we introduce a notion under which an admissible spectral partition can be ‘‘disentangled’’ by a transformation V , such that subsets of rows of VC_\star are supported on invariant subspaces of A_\star corresponding to the partition $\{\mathcal{S}_1, \dots, \mathcal{S}_r\}$.

Definition I.2 (Disentangling Matrix) *Let $A_\star = SJ_\star S^{-1}$. We see that an invertible matrix $V \in \mathbb{R}^{m \times m}$ disentangles an admissible spectral partition $\mathcal{S}_{1:r}$ of J_\star if we have the decomposition $VC_\star S = [C_1^\top | C_2^\top | \dots | C_r^\top]^\top$, where each matrix C_i is supported on entries in \mathcal{S}_i . We let $\text{cond}(V)$ denote the condition number of V , and denote the associated quantity*

$$M_C(\mathcal{S}_{1:r}; V) := \|V^{-1}\|_{\text{op}} \max_{v=(v_1, \dots, v_r) \in \mathcal{S}^{m-1}} \left(\sum_{i=1}^r \|v_i^\top C_i\|_2 \right).$$

Finally, we let $K_1(d, T, \alpha, q; \mathcal{S}_i)$ denote the analogue of K_1 restricted to pairs $(\lambda, k) \in \text{blkspec}(J(\mathcal{S}_i))$, and note that $K_1(d, T, \alpha, q; \mathcal{S}_i)$ also satisfies the bound in Lemma I.3. With these definitions in place, we have the following analogue of Proposition I.1, the result motivating the definition of K_1 .

Proposition I.7 *Suppose $\{\mathcal{S}_1, \dots, \mathcal{S}_r\} \subset [n]$ is an admissible partition of J_* , disentangled by a matrix V , and that each $J(\mathcal{S}_i)$ has (α_i, T) phase rank at most d . Then, there exists a filter $\phi \in \mathbb{R}^{m \times dm}$ with $\|\phi\|_{\text{bop}} \leq \kappa(V) \min(r, d)2^d$ such that*

$$\begin{aligned} \|\mathbf{G}_\phi\|_{\mathcal{H}_2^{\text{op}}} &\leq \|S^{-1}B_\star\|_{\text{op}} M_C(\mathcal{S}_{1:r}; V) \cdot \max_i K_1(d, T, \alpha_i, 2; \mathcal{S}_i) \\ \|\mathbf{G}_\phi\|_{\mathcal{H}_\infty} &\leq \|S^{-1}B_\star\|_{\text{op}} M_C(\mathcal{S}_{1:r}; V) \cdot \max_i K_1(d, T, \alpha, \infty; \mathcal{S}_i), \end{aligned}$$

with analogous bounds for \mathbf{F}_ϕ and \mathbf{H}_ϕ .

Following along the lines of the proof of Theorem I.4, we have the following bound for stochastic noise (we omit adversarial noise for brevity).

Theorem I.8 (Bounds for Stochastic Noise with Disentangling Predictors) *Suppose $\{\mathcal{S}_1, \dots, \mathcal{S}_r\} \subset [n]$ is an admissible partition of J_* , disentangled by a matrix V , and that each $J(\mathcal{S}_i)$ has (α_i, T) phase rank at most $1 \leq d \leq L$. Introduce the shorthand*

$$\tilde{K}_1(q) := \max_{i \in [r]} K_1(d, T, \alpha_i, q; \mathcal{S}_i)$$

Then, for any $\delta \in (0, 1)$ and $N \geq Td \max\{m, \log(1/\delta)\}$, it holds with probability $1 - \delta$ that

$$\begin{aligned} N^{-1/2} \text{Opt}_\mu &\lesssim M_C(\mathcal{S}_{1:r}; V) (M_B + N^{-1/2} M_0) \tilde{K}_1(2) \\ &\leq \min\{\tilde{K}_1(2), N^{-1/2} \tilde{K}_1(\infty)\} M_C(\mathcal{S}_{1:r}; V) M_B \sqrt{m + \log(1/\delta)} \\ &\leq \text{cond}(V) \min(r, d) 2^d \left(M_C M_B K_2(N) + M_D + \mu N^{-1/2} \right). \end{aligned} \quad (\text{I.31})$$

Corollary I.9 *Letting $\alpha_{\max} := \max_{i \in [r]} \{\alpha_i\}$, k denote the size of the largest Jordan block of A_\star and supposing $N \geq T(1 + \alpha_{\max})$, (I.31) can be bounded by*

$$\begin{aligned} N^{-1/2} \text{Opt}_\mu &\lesssim T^{k-1/2} \left[\left(k^2 \left(\frac{1 + \alpha_{\max}}{2} \right)^k M_C(\mathcal{S}_{1:r}; V) \left(M_B + \frac{M_0}{\sqrt{N}} \right) \right) \right. \\ &\quad \left. + \left(\text{cond}(V) (r \wedge d) 2^d \left(M_C M_B K_2(N) + M_D + \mu N^{-1/2} \right) \right) \right]. \end{aligned}$$

I.2. Polynomial Approximations for Linear Dynamical Systems

In this section, we present bounds on the terms $\|\mathbf{G}_\phi\|_{\mathcal{H}_\infty}$, $\|\mathbf{F}_\phi\|_{\mathcal{H}_\infty}$, and $\|\mathbf{H}_\phi\|_{\mathcal{H}_\infty}$. Our strategy is to relate these quantities to how well polynomials can approximate a set of complex numbers. To begin, we define $\text{Mon}(L, B)$ as the set of degree- L monic polynomials on \mathbb{C} ,

$$f(z) = z^L + f_1 z^{L-1} + \dots + f_L,$$

with real coefficients and ℓ_1 -norm at most B , i.e. $\|f\|_{L_1} := 1 + \sum_{i=1}^L |f_i| \leq B$. Furthermore, for a finite set $\mathcal{C} \subset \mathbb{D} \times \mathbb{N}$ (usually $\mathcal{C} = \text{blkspec}(A_\star)$), we define the following complexity terms, corresponding to \mathcal{H}_2 and \mathcal{H}_∞ . For $q \in \{2, \infty\}$, and constants $c_{\mathcal{H}_q}$ defined above, we define the complexity terms

$$H_f^{(q)}(\mathcal{C}, T) := \max_{(\lambda, k) \in \mathcal{C}} \begin{cases} 0 & f(z) \text{ has root order } \geq k \text{ at } \lambda \\ \infty & |\lambda| = 1, f(z) \text{ has root order } < k \text{ at } \lambda \\ c_{\mathcal{H}_q} \max_{z: |z-\lambda| \leq 1-|\lambda|} \frac{k^2 |f(z)|}{(1-|\lambda|)^{k-\frac{\mathbb{I}(q=2)}{2}}} & \text{otherwise.} \end{cases}$$

The term $H_f^{(q)}(\text{blkspec}(A_\star), T)$ roughly describes how effectively a polynomial f cancels the poles in A_\star . Due to the T -step subsampling, we shall typically be interested in $H_g^{(q)}$ for polynomials of the form $g(z) = f(z^T)$. We recall the definition of the \mathcal{H}_∞ -norm for a real rational transfer function $\mathbf{G}(z) : \mathbb{C} \rightarrow \mathbb{C}^{m \times p}$:

$$\|\mathbf{G}\|_{\mathcal{H}_\infty} := \sup_{z \in \mathbb{T}} \|\mathbf{G}(z)\|_{\text{op}} .$$

If the poles of A_\star are all strictly inside \mathbb{D} , this quantity is finite. Above, we use the operator norm on $\mathbb{C}^p \rightarrow \mathbb{C}^m$. Now, we define the $\mathcal{H}_2^{\text{op}}$ norm for such a transfer function via

$$\begin{aligned} \|\mathbf{G}\|_{\mathcal{H}_2^{\text{op}}} &:= \max_{v \in \mathcal{S}^{m-1}} \sqrt{\frac{1}{2\pi} \int_{z \in \mathbb{T}} \|v^\top \mathbf{G}(z)\|_2^2} \\ &= \max_{v \in \mathcal{S}^{m-1}} \sqrt{\frac{1}{2\pi} \int_{z \in \mathbb{T}} \text{tr}[(v^\top \mathbf{G}(z))^* (v^\top \mathbf{G}(z))]} \\ &= \max_{v \in \mathcal{S}^{m-1}} \|v^\top \mathbf{G}(z)\|_{\mathcal{H}_2}, \end{aligned} \tag{I.32}$$

where again we use the standard ℓ_2 -norm on \mathbb{C}^p and the definition of the canonical \mathcal{H}_2 -norm (see Section 4.3 of [Zhou et al. \(1996\)](#) for both the frequency-domain and time-domain definitions). Crucially, $\mathcal{H}_2^{\text{op}}$ is equal to the operator norm of the infinite-horizon Markov ‘‘matrix’’.

Lemma I.10 (Equivalence of $\mathcal{H}_2^{\text{op}}$ and $\|\mathcal{M}_\infty(\cdot)\|_{\text{op}}$) *Let $\mathbf{G} = (A, B, C, D)$, and suppose $\rho(A) < 1$. Then,*

$$\|\mathcal{M}_\infty(\mathbf{G})\|_{\text{op}} = \|\mathbf{G}\|_{\mathcal{H}_2^{\text{op}}} .$$

Proof Using (i) to denote block indexing, we see that

$$\begin{aligned} \|v^\top \mathcal{M}_k(\mathbf{G})\|_2^2 &= \text{tr} \left[\sum_{i=0}^{k-2} \mathcal{M}_k^{(i)}(v^\top \mathbf{G}) \mathcal{M}_k^{(i)}(v^\top \mathbf{G})^* \right] \\ \implies \|v^\top \mathcal{M}_\infty(\mathbf{G})\|_2^2 &= \text{tr} \left[\sum_{i=0}^{\infty} \mathcal{M}_\infty^{(i)}(v^\top \mathbf{G}) \mathcal{M}_\infty^{(i)}(v^\top \mathbf{G})^* \right] \\ &\stackrel{(*)}{=} \|v^\top \mathbf{G}\|_{\mathcal{H}_2}^2 \\ \implies \|\mathcal{M}_\infty(\mathbf{G})\|_{\text{op}}^2 &= \|\mathbf{G}\|_{\mathcal{H}_2^{\text{op}}}^2, \end{aligned}$$

where the limiting step holds as $\rho(A) < 1$, and $(*)$ comes from the time-domain characterization of the \mathcal{H}_2 norm. \blacksquare

At the center of our analysis is the following proposition, that demonstrates that $H_f^{(q)}$ does, in fact, describe the ability of polynomials f to cancel poles.

Proposition I.11 (Polynomial Approximation of Jordan Blocks) *Let $f : \mathbb{C} \rightarrow \mathbb{C}$ be an analytic function and $J \in \mathbb{R}^{n \times n}$ be a Jordan block matrix. Then $\|f(J)(zI - J)^{-1}\|_{\mathcal{H}_\infty} \leq H_f^{(\infty)}(\mathcal{C}, 1)$ and $\|f(J)(zI - J)^{-1}\|_{\mathcal{H}_2^{\text{op}}} \leq H_f^{(2)}(\mathcal{C}, 1)$. Moreover, if f is a polynomial of degree at most $d < k$, then the factor of k^2 in $H_f^{(q)}$ can be replaced with $k(d + 1)$.*

The above proposition is proved in Appendix J.1.4. Note that each of the $\mathcal{H}_2^{\text{op}}$ - and \mathcal{H}_∞ -norms is finite as long as its argument has all of its poles strictly inside the unit disk. Therefore, some poles of A_\star with modulus 1 may need to be canceled in order to achieve a finite $\mathcal{H}_2^{\text{op}}$ or \mathcal{H}_∞ -norm; we shall use the (standard) convention that the argument of the norm, as a real rational function of z , should be “evaluated” before computing the norm.

I.2.1. APPROXIMATIONS USING BLOCK-SCALAR FILTERS: PROOF OF PROPOSITION I.1

Our first theorem bounds the \mathcal{H}_∞ - and $\mathcal{H}_2^{\text{op}}$ -norms of \mathbf{F}_ϕ , \mathbf{G}_ϕ and \mathbf{H}_ϕ in terms of the quantity $H_f^{(q)}(\text{blkspec}(A_\star), T)$ by considering simple, block-weighted identity filters of the form

$$\phi = - [f_1 I_{m \times m} | f_2 I_{m \times m} | \dots | f_L I_{m \times m}] \in \mathbb{R}^{m \times Lm}, \quad (\text{I.33})$$

where f_1, \dots, f_L correspond to the coefficients of a polynomial $f \in \text{Mon}(L, B)$.

Theorem I.12 *Let $(A_\star, B_\star, C_\star)$ be a dynamical system, where $A_\star = S J_\star S^{-1}$ denotes the Jordan decomposition of A_\star . Then, for any $f \in \text{Mon}(L, B)$, the filter $\phi \in \mathbb{R}^{m \times Lm}$ from (I.33) satisfies $1 + \|\phi\|_{\text{bop}} \leq B$ and*

$$\begin{aligned} \|\mathbf{G}_\phi\|_{\mathcal{H}_2^{\text{op}}} &\leq \|S^{-1} B_\star\| \cdot \|C_\star S\|_{\text{op}} H_f^{(2)}(\text{blkspec}(J_\star), T) \\ \|\mathbf{G}_\phi\|_{\mathcal{H}_\infty} &\leq \|S^{-1} B_\star\| \cdot \|C_\star S\|_{\text{op}} H_f^{(\infty)}(\text{blkspec}(J_\star), T), \end{aligned}$$

and similarly for \mathbf{F}_ϕ and \mathbf{H}_ϕ , where B_\star is replaced by B_w and \mathbf{x}_1 , respectively.

The theorem above is proven in Section J.1.1. Note that this theorem does not preclude the case where $H_f(\text{blkspec}(J_\star), T) = \infty$, and thus the polynomial f must be chosen appropriately. In particular, by choosing $f(z) = \prod_{i=1}^d (z - \mu_i^T)$, where μ_1, \dots, μ_d are the complex numbers which witness the (α, T) -phase rank condition, we show in Section J.1.2 that for systems of bounded phase rank and Jordan block size, $H_f(\text{blkspec}(J_\star), T)$, is bounded. This is summarized in the following proposition.

Proposition I.13 *Suppose that J_\star has (α, T) -phase rank d . Then there exists a polynomial $f \in \text{Mon}(d, 2^d)$ such that $H_f^{(q)}(\text{blkspec}(J_\star), T) \leq K_1(d, T, \alpha, q)$.*

Proposition I.1 is now a direct consequence of combining Proposition I.13 with Theorem I.12 (with $L = d$).

I.2.2. APPROXIMATIONS USING DISENTANGLING FILTERS

We present an analogue of Theorem I.12 for disentangling filters.

Theorem I.14 *Let $G_\star = (A_\star, B_\star, C_\star)$, and let $A_\star = SJ_\star S^{-1}$ denote the Jordan decomposition of A_\star . Suppose V disentangles a spectral partition $\mathcal{S}_{1:r}$ of J_\star . Then, for any polynomials $f^{(1)}, \dots, f^{(r)} \in \text{Mon}(L, B)$, there exists a filter $\phi \in \mathbb{R}^{m \times Lm}$ with $\|\phi\|_{\text{bop}} \leq \text{cond}(V) \min(r, L)B$ satisfying*

$$\|\mathbf{G}_\phi\|_{\mathcal{H}_\infty} \leq \|V^{-1}\|_{\text{op}} \|S^{-1}B_\star\|_{\text{op}} \max_{v=(v_1, \dots, v_r) \in \mathcal{S}^{m-1}} \sum_{i=1}^r \|v_i^\top C_i\|_2 \left(H_{f^{(i)}}^{(\infty)}(\text{blkspec}(J_\star(\mathcal{S}_i)), T) \right)$$

and similarly for \mathbf{F}_ϕ and \mathbf{H}_ϕ , where B_\star is replaced by B_w and \mathbf{x}_1 , respectively. This also holds for the $\mathcal{H}_2^{\text{op}}$ analogues, replacing $H^{(\infty)}$ by $H^{(2)}$. Here $\text{cond}(V)$ denotes the condition number.

The proof is given in Section J.1.3. Proposition I.7 is now a corollary of this theorem and Proposition I.13.

I.3. Proof of Theorem I.4

We shall prove the stochastic case; the adversarial case follows from essentially the same arguments. Let A_\star have $\rho(A_\star) \leq 1$ and (α, T) phase rank d , and consider the filter ϕ from Proposition I.1, which satisfies $1 + \|\phi\| \leq 2^d$ and

$$\begin{aligned} \|\mathcal{M}_\infty(\mathbf{G}_\phi)\|_{\text{op}} &\leq \|S^{-1}B_\star\|_{\text{op}} \cdot \|C_\star S\|_{\text{op}} K_1(d, T, \alpha, 2) \\ \|\mathbf{G}_\phi\|_{\mathcal{H}_\infty} &\leq \|S^{-1}B_\star\|_{\text{op}} \cdot \|C_\star S\|_{\text{op}} K_1(d, T, \alpha, \infty), \end{aligned} \quad (\text{I.34})$$

and analogously for \mathbf{F}_ϕ and \mathbf{H}_ϕ . Consider the extended filter $\tilde{\phi} = [\phi | \mathbf{0}]$ obtained by embedding ϕ in $\mathbb{R}^{m \times Lm}$. Then, $\|\tilde{\phi}\|_{\text{op}} = \|\phi\|_{\text{op}}$, and thus

$$N^{-1/2} \text{Opt}_\mu \leq N^{-1/2} (\|\Delta_{\tilde{\phi}}\|_{\text{op}} + \mu \|\phi\|_{\text{op}}).$$

Therefore, by the assumption $N \geq Td \max \log(1/\delta)$, m , bounding $N^{-1/2} \|\Delta_{\tilde{\phi}}\|_{\text{op}}$ with Proposition 4.1 implies

$$\begin{aligned} N^{-1/2} \text{Opt}_\mu &\lesssim \underbrace{\|\mathcal{M}_\infty(\mathbf{G}_\phi)\|_{\text{op}} + \|\mathcal{M}_\infty(\mathbf{F}_\phi)\|_{\text{op}} + N^{-1/2} \|\mathcal{M}_\infty(\mathbf{H}_\phi)\|_{\text{op}}}_{(a)} \\ &\quad + \underbrace{\sqrt{\frac{m + \log(1/\delta)}{N}} (\Gamma_N(\mathbf{G}_\phi) + \Gamma_N(\mathbf{F}_\phi))}_{(b)} \\ &\quad + \underbrace{(1 + \|\phi\|_{\text{bop}}) (\|\mathcal{M}_{Td}(\mathbf{G})\|_{\text{op}} + \|\mathcal{M}_{Td}(\mathbf{F})\|_{\text{op}} + \|D_z\|_{\text{op}}) + N^{-1/2} \mu \|\phi\|_{\text{op}}}_{(c)}. \end{aligned}$$

For term (a), we have that

$$\begin{aligned} &(\|\mathcal{M}_\infty(\mathbf{G}_\phi)\|_{\text{op}} + \|\mathcal{M}_\infty(\mathbf{F}_\phi)\|_{\text{op}}) + N^{-1/2} \|\mathcal{M}_\infty(\mathbf{H}_\phi)\|_{\text{op}} \\ &\stackrel{(\text{I.34})}{\lesssim} \|C_\star S\|_{\text{op}} (N^{-1/2} \|S^{-1}\mathbf{x}_1\|_{\text{op}} + \|S^{-1}B_\star\|_{\text{op}} + \|S^{-1}B_w\|_{\text{op}}) K_1(d, T, \alpha, 2) \\ &\leq M_C (M_B + N^{-1/2} M_0) K_1(d, T, \alpha, 2). \end{aligned}$$

Similarly, for the term (b), recalling $\Gamma_N(\mathbf{G}) = \min\{\sqrt{N}\mathcal{M}_N(\mathbf{G}), \|\mathbf{G}\|_{\mathcal{H}_\infty}\}$,

$$\begin{aligned} & \sqrt{\frac{m + \log(1/\delta)}{N}} (\Gamma_N(\mathbf{G}_\phi) + \Gamma_N(\mathbf{F}_\phi)) \\ & \leq \min\{K_1(d, T, \alpha, 2), N^{-1/2}K_1(d, T, \alpha, \infty)\} \cdot \|C_\star S\|_{\text{op}} (\|S^{-1}B_w\| + \|S^{-1}B_\star\|_{\text{op}}) \sqrt{m + \log \frac{1}{\delta}} \\ & = \min\{K_1(d, T, \alpha, 2), N^{-1/2}K_1(d, T, \alpha, \infty)\} M_B M_C \sqrt{m + \log \frac{1}{\delta}}, \end{aligned}$$

where we take $\|\mathcal{M}_N(\mathbf{G}_\phi)\|_{\text{op}} \leq \|\mathcal{M}_\infty(\mathbf{G}_\phi)\|_{\text{op}}$ to use (I.34). For term (c), we have $1 + \|\phi\|_{\text{bop}} \leq 2^d$ and $\|\phi\|_{\text{op}} \leq \|\phi\|_{\text{bop}} \leq 2^d$, so that

$$\begin{aligned} & (1 + \|\phi\|_{\text{bop}}) (\|\mathcal{M}_{Td}(\mathbf{G})\|_{\text{op}} + \|\mathcal{M}_{Td}(\mathbf{F})\|_{\text{op}} + \|D_z\|_{\text{op}}) + N^{-1/2}\mu\|\phi\|_{\text{op}} \\ & \leq 2^d \left(\|\mathcal{M}_{Td}(\mathbf{G})\|_{\text{op}} + \|\mathcal{M}_{Td}(\mathbf{F})\|_{\text{op}} + \|D_z\|_{\text{op}} + N^{-1/2}\mu \right) \\ & \lesssim 2^d \left(M_B M_C K_2(Td) + \|D_\star\|_{\text{op}} + \|D_z\|_{\text{op}} + N^{-1/2}\mu \right) \\ & \leq 2^d \left(M_B M_C K_2(Td) + M_D + N^{-1/2}\mu \right), \end{aligned}$$

where the bound $\|\mathcal{M}_{Td}(\mathbf{G})\|_{\text{op}} + \|\mathcal{M}_{Td}(\mathbf{F})\|_{\text{op}} \lesssim M_B M_C K_2(Td)$ follows from Proposition I.2. Combining parts (a), (b), and (c) then yields Theorem I.4.

Appendix J. Supporting Proofs

J.1. Proofs for Section I.2

J.1.1. PROOF OF THEOREM I.12

We prove the bound for \mathbf{G}_ϕ without loss of generality. Let $f \in \text{Mon}(L, B)$, and define the corresponding filter

$$\phi = - [f_1 I_{m \times m} | f_2 I_{m \times m} | \dots | f_L I_{m \times m}] \in \mathbb{R}^{m \times Lm}.$$

Since $f \in \text{Mon}(L, B)$, we have that

$$1 + \|\phi\|_{\text{bop}} = 1 + \sum_{\ell=1}^L \|f_\ell I_{m \times m}\|_{\text{op}} = 1 + \sum_{\ell=1}^L |f_\ell| = \|f\|_1 \leq B.$$

With this filter, the corresponding observation matrix C_ϕ is given by (with a reasonable abuse of notation)

$$C_\phi = C_\star A_\star^{LT} + C_\star f_1 A_\star^{(L-1)T} + C_\star f_2 A_\star^{(L-2)T} + \dots + C_\star f_L = C_\star f(A_\star^T).$$

Therefore,

$$\mathbf{G}_\phi(z) = C_\star f(A_\star^T) (zI - A_\star)^{-1} B_\star = C_\star S f(J_\star^T) (zI - J_\star)^{-1} S^{-1} B_\star.$$

We can now bound the \mathcal{H}_∞ -norm of \mathbf{G}_ϕ in terms of f :

$$\begin{aligned} \|\mathbf{G}_\phi\|_{\mathcal{H}_\infty} &= \max_{z \in \mathbb{T}} \|C_\star S f(J_\star^T)(zI - J_\star)^{-1} S^{-1} B_\star\|_{\text{op}} \\ &\leq \|C_\star S\|_{\text{op}} \|S^{-1} B_\star\|_{\text{op}} \max_{z \in \mathbb{T}} \|f(J_\star^T)(zI - J_\star)^{-1}\|_{\text{op}} \\ &= \|C_\star S\|_{\text{op}} \|S^{-1} B_\star\|_{\text{op}} \|f(J_\star^T)(zI - J_\star)^{-1}\|_{\mathcal{H}_\infty}. \end{aligned}$$

A similar argument shows that

$$\|\mathbf{G}_\phi\|_{\mathcal{H}_2^{\text{op}}} \leq \|C_\star S\|_{\text{op}} \|S^{-1} B_\star\|_{\text{op}} \|f(J_\star^T)(zI - J_\star)^{-1}\|_{\mathcal{H}_2^{\text{op}}}.$$

To control the \mathcal{H}_∞ and $\mathcal{H}_2^{\text{op}}$ terms at the heart of the above bound, we will recall the following proposition, proved in Section J.1.4 below.

Proposition I.11 (Polynomial Approximation of Jordan Blocks) *Let $f : \mathbb{C} \rightarrow \mathbb{C}$ be an analytic function and $J \in \mathbb{R}^{n \times n}$ be a Jordan block matrix. Then $\|f(J)(zI - J)^{-1}\|_{\mathcal{H}_\infty} \leq H_f^{(\infty)}(\mathcal{C}, 1)$ and $\|f(J)(zI - J)^{-1}\|_{\mathcal{H}_2^{\text{op}}} \leq H_f^{(2)}(\mathcal{C}, 1)$. Moreover, if f is a polynomial of degree at most $d < k$, then the factor of k^2 in $H_f^{(q)}$ can be replaced with $k(d+1)$.*

In particular, we apply Proposition I.11 with the analytic function $\tilde{f}(z) := f(z^T)$. This implies that $\|f(J_\star^T)(zI - J_\star)^{-1}\|_{\mathcal{H}_\infty} \leq H_f^{(\infty)}(\mathcal{C}, T)$. Hence, we find that

$$\|\mathbf{G}_\phi\|_{\mathcal{H}_\infty} \leq \|C_\star S\|_{\text{op}} \|S^{-1} B_\star\|_{\text{op}} H_f^{(\infty)}(\mathcal{C}, T),$$

and similarly for $H_f^{(2)}$.

J.1.2. PROOF OF PROPOSITION I.13

Let μ_1, \dots, μ_d witness the (α, T) -phase rank condition of $\text{blkspec}(A_\star)$. We now consider the corresponding polynomial $f(z) := \prod_{i=1}^d (z - \mu_i^T)$. Note that f is monic and has degree d ; thus, the fact that $f \in \text{Mon}(d, 2^d)$ follow from the following bound on its ℓ_1 -norm.

Lemma J.1 *Let f be a degree- d polynomial whose roots all lie in \mathbb{D} . Then $\max_{z \in \mathbb{D}} |f(z)| \leq \|f\|_{\ell_1} \leq 2^d$.*

Proof The bound $\max_{z \in \mathbb{D}} |f(z)| \leq \|f\|_{\ell_1}$ holds for any polynomial by the triangle inequality. Then, since $|\mu_i| \leq 1$, $\|f\|_{\ell_1} = \sum_{i=0}^d |\sum_{S \in \binom{[d]}{i}} \prod_{i \in S} \mu_i| \leq \sum_{i=0}^d \binom{d}{i} = 2^d$. \blacksquare

Next, for any $(\lambda, k) \in \mathcal{C}$, we shall bound each term in the maximum of $H_f(\mathcal{C}, T)$.

For $|\lambda| = 1$, the approximate phase rank condition implies that there are at least k elements μ_i such that $\lambda^T = \mu_i^T$. Thus, $f(z^T)$ has a root of order $\geq k$ at λ , and the corresponding term in $H_f^{(q)}(\mathcal{C}, T)$ evaluates to zero. We shall therefore show that for $|\lambda| < 1$, one has

$$\max_{z: |z-\lambda| \leq 1-|\lambda|} \frac{k^2 |f(z^T)|}{(1-|\lambda|)^k} \leq \begin{cases} (1+\alpha)^k T^k 2^{d-k} k & |\lambda| \in [1 - \frac{1}{T(1+\alpha)}, 1) \\ \frac{2^d}{(1-|\lambda|)^k} & |\lambda| < 1 - \frac{1}{T(1+\alpha)} \end{cases}. \quad (\text{J.35})$$

This immediately implies the desired bound on $H_f^{(\infty)}(\text{blkspec}(A_\star), T)$ in Proposition I.13. For the correct bound for $q = 2$, we note that since we need only consider $|\lambda| < 1$, by factoring out a $(1 - |\lambda|)^{1/2}$ we have

$$\begin{aligned}
 c_{\mathcal{H}_2}^{-1} H_f^{(2)}(\text{blkspec}(A_\star), T) &= \max_{(\lambda, k) \in \text{blkspec}(A_\star): |\lambda| < 1} \max_{z: |z - \lambda| \leq 1 - |\lambda|} \frac{k^2 |f(z^T)|}{(1 - |\lambda|)^{k-1/2}} \\
 &\stackrel{\text{(J.35)}}{\leq} (1 - |\lambda|)^{1/2} \begin{cases} (1 + \alpha)^k T^k 2^{d-k} k & |\lambda| \in [1 - \frac{1}{T(1+\alpha)}, 1) \\ \frac{2^d}{(1-|\lambda|)^k} & |\lambda| < 1 - \frac{1}{T(1+\alpha)} \end{cases} \\
 &= \begin{cases} (1 - |\lambda|)^{1/2} (1 + \alpha)^k T^k 2^{d-k} k & |\lambda| \in [1 - \frac{1}{T(1+\alpha)}, 1) \\ \frac{2^d}{(1-|\lambda|)^{k-1/2}} & |\lambda| < 1 - \frac{1}{T(1+\alpha)} \end{cases} \\
 &\leq \begin{cases} (1 + \alpha)^{k-1/2} T^{k-1/2} 2^{d-k} k & |\lambda| \in [1 - \frac{1}{T(1+\alpha)}, 1) \\ \frac{2^d}{(1-|\lambda|)^{k-1/2}} & |\lambda| < 1 - \frac{1}{T(1+\alpha)} \end{cases},
 \end{aligned}$$

where the last line uses that $(1 - |\lambda|)^{1/2} \leq (1 + \alpha)^{-1/2} T^{-1/2}$ for $|\lambda| \in [1 - \frac{1}{T(1+\alpha)}, 1)$.

We now turn our attention to the proof of (J.35). For $\lambda < 1 - \frac{1}{(1+\alpha)T}$ and $z \in \mathbb{C}$ such that $|z - \lambda| \leq 1 - |\lambda|$, we note that $z \in \mathbb{D}$. Thus, $z^T \in \mathbb{D}$ implies $f(z^T) \leq 2^d$ by Lemma J.1. Hence,

$$k^2 \max_{z: |z - \lambda| \leq 1 - |\lambda|} \frac{|f(z^T)|}{(1 - |\lambda|)^k} \leq k^2 \frac{2^d}{(1 - |\lambda|)^k}.$$

Lastly, we consider the case $\lambda \in [1 - \frac{1}{(1+\alpha)T}, 1)$. Since $(\lambda, k) \in \mathcal{C}$ and $\{\mu_1, \dots, \mu_d\}$ witnesses the (α, T) -phase rank condition, without loss of generality (by permuting labels) we have $\max_{j \in [k]} \min_{\tilde{\mu}: \tilde{\mu}^T = \mu_j^T} |\lambda - \tilde{\mu}| \leq \alpha(1 - |\lambda|)$. Letting $\tilde{\mu}_i$ denote a complex number satisfying $\tilde{\mu}_i^T = \mu_i^T$ which minimizes $|\lambda - \tilde{\mu}_i|$ (breaking ties arbitrarily), it then follows that

$$\max_{i \in [k]} |\lambda - \tilde{\mu}_i| \leq \alpha(1 - |\lambda|).$$

We shall then factor

$$f(z^T) = \left(\prod_{i=1}^k (z - \tilde{\mu}_i) \right) \cdot F(z)$$

where

$$F(z) := \left(\prod_{i=1}^k \frac{z^T - \mu_i^T}{z - \tilde{\mu}_i} \right) \cdot \left(\prod_{i=k+1}^d (z^T - \mu_i^T) \right).$$

We first estimate the magnitude of $F(z)$.

Lemma J.2 $F(z) \leq T^k 2^{d-k}$ for any $z: |z - \lambda| \leq 1 - |\lambda|$.

Proof First, observe that any $z: |z - \lambda| \leq 1 - |\lambda|$ lies in \mathbb{D} . Hence,

$$|F(z)| \leq \left| \prod_{i=1}^k \frac{z^T - \mu_i^T}{z - \tilde{\mu}_i} \right| \cdot \left| \prod_{i=k+1}^d |z^T| + |\mu_i^T| \right| \leq 2^{d-k} \left| \prod_{i=1}^k \frac{z^T - \mu_i^T}{z - \tilde{\mu}_i} \right|.$$

Next, we observe that for $i \in [k]$, we have $\tilde{\mu}_i^T = \mu_i^T$. Thus,

$$\left| \frac{z^T - \mu_i^T}{z - \tilde{\mu}_i} \right| = \left| \frac{z^T - \tilde{\mu}_i^T}{z - \tilde{\mu}_i} \right| = \left| z^{T-1} + z^{T-2}\tilde{\mu}_i + \dots + \tilde{\mu}_i^{T-1} \right| \leq T.$$

Hence $|F(z)| \leq 2^{d-k}T^k$. ■

Next, we estimate the contribution of $\left| \prod_{i=1}^k (z - \tilde{\mu}_i) \right|$.

Lemma J.3 *Let z satisfy $|z - \lambda| \leq 1 - |\lambda|$. Then, $\left| \frac{\prod_{i=1}^k (z - \tilde{\mu}_i)}{(1 - |\lambda|)^k} \right| \leq (1 + \alpha)^k$.*

Proof For $z : |z - \lambda| \leq 1 - |\lambda|$, we have

$$\begin{aligned} \left| \prod_{i=1}^k (z - \tilde{\mu}_i) \right| &\leq \prod_{i=1}^k (|z - \lambda| + |\lambda - \tilde{\mu}_i|) \\ &\leq \prod_{i=1}^k ((1 - |\lambda|) + \alpha(1 - |\lambda|)) = (1 + \alpha)^k (1 - |\lambda|)^k. \end{aligned}$$
■

Combining these two estimates, we find that

$$k^2 \frac{\max_{z: |z - \lambda| \leq (1 - |\lambda|)} |f(z^T)|}{(1 - |\lambda|)^k} \leq k^2 (1 + \alpha)^k T^k 2^{d-k}.$$

In summary, we have shown that for any $(\lambda, k) \in \mathcal{C}$ with $|\lambda| < 1$, one has

$$\max_{z: |z - \lambda| \leq 1 - |\lambda|} \frac{k^2 |f(z^T)|}{(1 - |\lambda|)^k} \leq k^2 \begin{cases} (1 + \alpha)^k T^k 2^{d-k} & \lambda > 1 - \frac{1}{T(1 + \alpha)} \\ \frac{2^d}{(1 - |\lambda|)^k} & \lambda \leq \frac{1}{T(1 + \alpha)} \end{cases}.$$

J.1.3. PROOF OF THEOREM I.14

We shall prove the bound for $H^{(\infty)}$; the bound for $H^{(2)}$ is similar. As in the statement of the theorem, let $A_\star = S J_\star S^{-1}$, let $\mathcal{S} := \text{blkspec}(A_\star)$, and suppose there exists subsets $\mathcal{S}_1, \dots, \mathcal{S}_r \subset \mathcal{S}$ and an invertible transformation V for which $V C_\star S = [C_1^\top | C_2^\top | \dots | C_r^\top]^\top$, where C_i is supported on entries corresponding to \mathcal{S}_i , and $C_i \in \mathbb{R}^{m_i \times n}$. Finally, let $f^{(1)}, \dots, f^{(r)}$ denote the polynomials in $\text{Mon}(L, B)$, and define the block diagonal matrix

$$X_\ell := \text{blkdiag}(f_\ell^{(1)} I_{m_1 \times m_1}, f_\ell^{(2)} I_{m_2 \times m_2}, \dots, f_\ell^{(r)} I_{m_r \times m_r}),$$

and let

$$\phi = - [V^{-1} X_1 V | V^{-1} X_2 V | \dots | V^{-1} X_L V].$$

We can then compute that

$$\begin{aligned}
 C_\phi &= C_\star A_\star^{TL} + V^{-1} X_1 V C_\star A_\star^{(L-1)T} + \dots + V^{-1} X_L V C_\star \\
 &= V^{-1} \left(V C_\star S S^{-1} A_\star^{TL} S + X_1 V C_\star S S^{-1} A_\star^{(L-1)T} S + \dots + V^{-1} X_L V C_\star S \right) S^{-1} \\
 &= V^{-1} \left(V C_\star S J_\star^{TL} + X_1 V C_\star S J_\star^{T(L-1)} + \dots + V^{-1} X_L V C_\star S \right) S^{-1} \\
 &= V^{-1} \left(\begin{bmatrix} C_1 J_\star^{TL} \\ C_2 J_\star^{TL} \\ \vdots \\ C_r J_\star^{TL} \end{bmatrix} + \begin{bmatrix} f_1^{(1)} C_1 J_\star^{T(L-1)} \\ f_1^{(2)} C_2 J_\star^{T(L-1)} \\ \vdots \\ f_1^{(r)} C_r J_\star^{T(L-1)} \end{bmatrix} + \dots + \begin{bmatrix} f_L^{(1)} C_1 \\ f_L^{(2)} C_2 \\ \vdots \\ f_L^{(r)} C_r \end{bmatrix} \right) S^{-1} \\
 &= V^{-1} \begin{bmatrix} C_1 f^{(1)}(J_\star) \\ C_2 f^{(2)}(J_\star) \\ \vdots \\ C_r f^{(r)}(J_\star) \end{bmatrix} S^{-1}.
 \end{aligned}$$

Hence,

$$\begin{aligned}
 \|\mathbf{G}_\phi\|_{\text{op}} &= \sup_{z \in \mathbb{T}} \left\| V^{-1} \begin{bmatrix} C_1 f^{(1)}(J_\star) \\ C_2 f^{(2)}(J_\star) \\ \vdots \\ C_r f^{(r)}(J_\star) \end{bmatrix} S^{-1} (zI - A_\star)^{-1} B_\star \right\|_{\text{op}} \\
 &= \sup_{z \in \mathbb{T}} \left\| V^{-1} \begin{bmatrix} C_1 f^{(1)}(J_\star) \\ C_2 f^{(2)}(J_\star) \\ \vdots \\ C_r f^{(r)}(J_\star) \end{bmatrix} (zI - J_\star)^{-1} S^{-1} B_\star \right\|_{\text{op}} \\
 &\leq \|V^{-1}\|_{\text{op}} \|S^{-1} B_\star\|_{\text{op}} \sup_{z \in \mathbb{T}} \left\| \begin{bmatrix} C_1 f^{(1)}(J_\star) (zI - J_\star)^{-1} \\ C_2 f^{(2)}(J_\star) (zI - J_\star)^{-1} \\ \vdots \\ C_r f^{(r)}(J_\star) (zI - J_\star)^{-1} \end{bmatrix} \right\|_{\text{op}} \\
 &= \|V^{-1}\|_{\text{op}} \|S^{-1} B_\star\|_{\text{op}} \sup_{z \in \mathbb{T}} \max_{v=(v_1, \dots, v_r) \in \mathcal{S}^{m-1}} \left\| \sum_{i=1}^r v_i^\top C_i f^{(i)}(J_\star) (zI - J_\star)^{-1} \right\|_2. \\
 &\leq \|V^{-1}\|_{\text{op}} \|S^{-1} B_\star\|_{\text{op}} \sup_{z \in \mathbb{T}} \max_{v=(v_1, \dots, v_r) \in \mathcal{S}^{m-1}} \sum_{i=1}^r \|v_i^\top C_i f^{(i)}(J_\star) (zI - J_\star)^{-1}\|_2.
 \end{aligned}$$

Note that the sizes of v_i are given by the admissible spectral partition of A_\star . By assumption, for each i , C_i is supported on coordinates in \mathcal{S}_i . Hence, we see that

$$\|v_i^\top C_i f^{(i)}(J_\star) (zI - J_\star)^{-1}\|_2 = \|v_i^\top C_i f^{(i)}(J(\mathcal{S}_i)) (zI - J(\mathcal{S}_i))^{-1}\|_2,$$

since $J(\mathcal{S}_i)$ is J_\star supported on \mathcal{S}_i . Thus, for any $z \in \mathbb{D}$,

$$\begin{aligned} & \sup_{z \in \mathbb{T}} \max_{v=(v_1, \dots, v_r) \in \mathcal{S}^{m-1}} \sum_{i=1}^r \|v_i^\top C_i f^{(i)}(J_\star)(zI - J_\star)^{-1}\|_2 \\ & \leq \max_{v=(v_1, \dots, v_r) \in \mathcal{S}^{m-1}} \sum_{i=1}^r \|v_i^\top C_i\|_2 \sup_{z \in \mathbb{T}} \|f^{(i)}(J(\mathcal{S}_i))(zI - J(\mathcal{S}_i))^{-1}\|_{\text{op}} \\ & \leq \max_{v=(v_1, \dots, v_r) \in \mathcal{S}^{m-1}} \sum_{i=1}^r \|v_i^\top C_i\|_2 H_{f^{(i)}}^{(\infty)}(\text{blkspec}(J_\star(\mathcal{S}_i)), T), \end{aligned}$$

where for each term we invoke Proposition I.11 and argue as in Theorem I.12. Therefore,

$$\|\mathbf{G}_\phi\|_{\mathcal{H}_\infty} \leq \|V^{-1}\|_{\text{op}} \|S^{-1}B_\star\|_{\text{op}} \cdot \max_{v=(v_1, \dots, v_r) \in \mathcal{S}^{m-1}} \sum_{i=1}^r \|v_i^\top C_i\|_2 H_{f^{(i)}}^{(\infty)}(\text{blkspec}(J_\star(\mathcal{S}_i)), T).$$

Moreover, by the construction in Theorem I.12, we see that

$$\|\phi\|_{\text{bop}} \leq \text{cond}(V) \sum_{\ell=1}^L \|X_\ell\|_{\text{op}} = \text{cond}(V) \sum_{\ell=1}^L \max_{i \in [r]} |f_\ell^{(i)}| \leq \text{cond}(V) \min(r, L) B.$$

J.1.4. PROOF OF PROPOSITION I.11

We wish to show that as long as $f(z)$ has a root of order at least k at each $(\lambda, k) \in \text{blkspec}$ with $|\lambda| = 1$, then

$$\begin{aligned} \|f(J_\star)(zI - J_\star)^{-1}\|_{\mathcal{H}_\infty} & \leq c_{\mathcal{H}_\infty} \max_{w: |w-\lambda| \leq 1-|\lambda|} \frac{k^2 |f(w)|}{(1-|\lambda|)^k}, \quad \text{and} \\ \|f(J_\star)(zI - J_\star)^{-1}\|_{\mathcal{H}_2^{\text{op}}} & \leq c_{\mathcal{H}_2} \max_{w: |w-\lambda| \leq 1-|\lambda|} \frac{k^2 |f(w)|}{(1-|\lambda|)^{k-\frac{1}{2}}}, \end{aligned}$$

where k^2 can be replaced by $k(d+1)$ if f is given by a polynomial of degree d . Note that the right-hand sides of the above displays use argument w instead of z to avoid confusion with the parameter z on the right-hand side. We start with following lemma:

Lemma J.4 *Let $B(z) \in \mathbb{R}^{n \times n}$ be a block-diagonal transfer function with blocks $B_1(z), \dots, B_r(z)$. Then, $\|B(z)\|_{\mathcal{H}_\infty} = \max_{i \in [r]} \|B_i(z)\|_{\mathcal{H}_\infty}$ and $\|B(z)\|_{\mathcal{H}_2^{\text{op}}} = \max_{i \in [r]} \|B_i(z)\|_{\mathcal{H}_2^{\text{op}}}$.*

Proof

For \mathcal{H}_∞ ,

$$\begin{aligned} \|B(z)\|_{\mathcal{H}_\infty} & = \max_{z \in \mathbb{T}} \|B(z)\|_{\text{op}} = \max_{z \in \mathbb{T}} \max_{j \in [r]} \|B_j(z)\|_{\text{op}} \\ & = \max_{j \in [r]} \max_{z \in \mathbb{T}} \|B_j(z)\|_{\text{op}} = \max_{j \in [r]} \|B_j(z)\|_{\mathcal{H}_\infty} \end{aligned}$$

For $\mathcal{H}_2^{\text{op}}$, let $v = (v_1, \dots, v_r)$ be a decomposition of v along the blocks $B_j \in \mathbb{R}^{n_j \times n_j}$. Then:

$$\begin{aligned} \|B(z)\|_{\mathcal{H}_2^{\text{op}}}^2 &= \max_{v \in \mathcal{S}^{n-1}} \frac{1}{2\pi} \int_0^{2\pi} \text{tr}(v^\top B(e^{i\theta})^* B(e^{i\theta}) v) d\theta \\ &= \max_{v \in \mathcal{S}^{n-1}} \sum_{j=1}^r \frac{1}{2\pi} \int_0^{2\pi} \text{tr}(v_j^\top B_j(e^{i\theta})^* B_j(e^{i\theta}) v_j) \\ &\leq \max_{v \in \mathcal{S}^{n-1}} \|v_j\|_2^2 \max_{w_j \in \mathcal{S}^{n_j-1}} \frac{1}{2\pi} \int_0^{2\pi} \sum_{j=1}^r \text{tr}(w_j^\top B_j(e^{i\theta})^* B_j(e^{i\theta}) w_j) \\ &= \max_{v \in \mathcal{S}^{n-1}} \|v_j\|_2^2 \|B_j(z)\|_{\mathcal{H}_2^{\text{op}}}^2 = \max_{j \in [r]} \|B_j(z)\|_{\mathcal{H}_2^{\text{op}}}^2, \end{aligned}$$

where the last holds since $\sum_j \|v_j\|_2^2 = \|v\|_2^2 = 1$. To see the converse holds, one may choose v to be supported on the coordinates of one block B_j . \blacksquare

We now return to the proof of Proposition 1.11. Let f be an analytic function of degree d , where we take $d = \infty$ if f is not a finite-length polynomial. Since $f(J)(zI - J)^{-1}$ has the same Jordan block structure as J , the above lemma lets us assume without loss of generality that J consists of a single Jordan block corresponding to an eigenvalue λ of order k . If f has a zero of order k at λ , then f is divisible by $(z - \lambda)^k$, which by Cayley-Hamilton implies that $f(J) = \mathbf{0}$ and thus $\|f(J)(zI - J)^{-1}\|_{\mathcal{H}_2^{\text{op}}} \leq \|f(J)(zI - J)^{-1}\|_{\mathcal{H}_\infty} = 0$. We shall now show that for any $|\lambda| < 1$ and $z \in \mathbb{T}$,

$$\|f(J)h_z(J)\|_{\text{op}} \leq \frac{1}{|z - \lambda|} \max_{w: |w - \lambda| \leq 1 - |\lambda|} \frac{|f(w)|k(k \wedge d + 1)}{(1 - |\lambda|)^{k-1}}. \quad (\text{J.36})$$

We first show how to conclude the proof assuming the above display (J.36), and then turn to establishing the inequality.

1. Concluding the proof from (J.36): To bound \mathcal{H}_∞ , we see that

$$\|f(J)h_z(J)\|_{\mathcal{H}_\infty} = \max_{z \in \mathbb{T}} \|f(J)h_z(J)\|_{\text{op}} \leq \max_{w: |w - \lambda| \leq 1 - |\lambda|} \frac{|f(w)|k(k \wedge d + 1)}{(1 - |\lambda|)^k},$$

since we have $|z - \lambda| \geq 1 - |\lambda|$ for all $z \in \mathbb{T}$ and $\lambda \in \mathbb{D}$. Noting that $c_{\mathcal{H}_\infty} = 1$, the above display is precisely the quantity corresponding to $H_f^{(\infty)}$ by taking $k(k \wedge (d + 1)) \leq k^2$, but it allows one to replace k^2 by $k(d + 1)$ for $d < k$.

To bound $\mathcal{H}_2^{\text{op}}$, we use the following lemma, proved at the end of the section.

Lemma J.5 (\mathcal{H}_2 integration) For $c_{\mathcal{H}_2} = \sqrt{1 + \frac{2}{\pi}}$, $\sqrt{\frac{1}{2\pi} \int_0^{2\pi} \frac{1}{|e^{i\theta} - \lambda|^2} d\theta} \leq c_{\mathcal{H}_2} \sqrt{\frac{1}{1 - |\lambda|}}$ for all $\lambda \in \mathbb{D} : |\lambda| < 1$.

We note that in our context, the variable over which we integrate is the subscript z in $f(J)h_z(J)$. Using the integral computation,

$$\begin{aligned}
 \|f(J)h_z(J)\|_{\mathcal{H}_2^{\text{op}}}^2 &= \max_{v \in \mathcal{S}^{n-1}} \frac{1}{2\pi} \int_0^{2\pi} \text{tr}(v^\top (f(J)h_{e^{i\theta}}(J))^* (f(J)h_{e^{i\theta}}(J))v) \\
 &\leq \frac{1}{2\pi} \int_0^{2\pi} \|(f(J)h_{e^{i\theta}}(J))\|_{\text{op}}^2 \\
 &\stackrel{\text{(J.36)}}{\leq} \left(\max_{w:|w-\lambda|\leq 1-|\lambda|} \frac{|f(w)|k(k \wedge (d+1))}{(1-|\lambda|)^{k-1}} \right)^2 \frac{1}{2\pi} \int_0^{2\pi} \frac{1}{|z-\lambda|^2} \\
 &\stackrel{\text{Lem. J.5}}{\leq} \left(\max_{w:|w-\lambda|\leq 1-|\lambda|} \frac{|f(w)|k(k \wedge (d+1))}{(1-|\lambda|)^{k-1}} \right)^2 c_{\mathcal{H}_2}^2 \frac{1}{1-|\lambda|} \\
 &= \left(\max_{w:|w-\lambda|\leq 1-|\lambda|} \frac{|f(w)|k(k \wedge (d+1))}{(1-|\lambda|)^{k-1/2}} \right)^2 c_{\mathcal{H}_2}^2
 \end{aligned}$$

as needed.

2. Proving (J.36): Since $|\lambda| < 1$, the function $h_z(\lambda) := \frac{1}{z-\lambda}$ is analytic on \mathbb{T} , and thus we can write $f(J)(zI - J)^{-1} = f(J)h_z(J)$. By Lemma G.5 and the formula for functions of Jordan block matrices, we see that (dropping the λ argument for brevity) for any $z \in \mathbb{T}$,

$$\begin{aligned}
 \|f(J)h_z(J)\|_{\text{op}} &\leq \sum_{\ell=0}^{k-1} \sum_{j=0}^{\ell} \left| \frac{h_z^{(\ell-j)}}{(\ell-j)!} \right| \left| \frac{f^{(j)}}{j!} \right| \\
 &= \sum_{j=0}^{k-1} \left| \frac{f^{(j)}}{j!} \right| \left| \sum_{\ell=j}^{k-1} \frac{h_z^{(\ell-j)}}{(\ell-j)!} \right| \\
 &= \sum_{j=0}^{k-1} \left| \frac{f^{(j)}}{j!} \right| \left| \sum_{\ell=1}^{k-j} \frac{h_z^{(\ell)}}{\ell!} \right| \\
 &= \sum_{j=0}^{k-1} \left| \frac{f^{(j)}}{j!} \right| \left| \sum_{\ell=1}^{k-j} |z-\lambda|^{-\ell} \right| \\
 &= \sum_{j=0}^{k-1 \wedge d} \left| \frac{f^{(j)}}{j!} \right| \left| \sum_{\ell=1}^{k-j} |z-\lambda|^{-\ell} \right|, \tag{J.37}
 \end{aligned}$$

where the last line uses the fact that if f is a degree d polynomial, all $f^{(j)}$ vanish for $j > d$. Next, we use Cauchy's integral formula to bound the magnitudes of the terms $\left| \frac{f^{(j)}}{j!} \right|$:

Lemma J.6 (Cauchy's Integral Formula, see e.g. Stein and Shakarchi (2003)) *Let $f : \mathbb{C} \rightarrow \mathbb{C}$ be an analytic function. Then $f^{(n)}(\lambda) \leq \frac{n!}{r^n} \max_{a:|\lambda-a|=r} |f(a)|$.*

By setting $r = 1 - |\lambda|$, we have that

$$\left| \frac{f^{(j)}(\lambda)}{j!} \right| \leq \max_{w:|w-\lambda|\leq 1-|\lambda|} \frac{|f(w)|}{(1-|\lambda|)^j}.$$

Furthermore, if f is a polynomial of degree d , moreover, then $f^{(j)}$ vanishes for all $j > d$. Therefore, picking up where we left off from our bound on $\|f(J)h_z(J)\|_{\text{op}}$ in Equation (J.37), we bound

$$\begin{aligned}
 \sum_{j=0}^{d \wedge k-1} \left| \frac{f^{(j)}}{j!} \right| \sum_{\ell=1}^{k-j} |z - \lambda|^{-\ell} &\leq \max_{w: |w-\lambda| \leq 1-|\lambda|} |f(w)| \sum_{j=0}^{d \wedge k-1} \sum_{\ell=1}^{k-j} |z - \lambda|^{-\ell} (1 - |\lambda|)^{-j} \\
 &= \frac{1}{|z - \lambda|} \max_{w: |w-\lambda| \leq 1-|\lambda|} |f(w)| \sum_{j=0}^{k-1 \wedge d} \sum_{\ell=1}^{k-j} |z - \lambda|^{1-\ell} (1 - |\lambda|)^{-j} \\
 &\stackrel{(i)}{\leq} \frac{1}{|z - \lambda|} \max_{w: |w-\lambda| \leq 1-|\lambda|} |f(w)| \sum_{j=0}^{k-1 \wedge d} \sum_{\ell=1}^{k-j} (1 - |\lambda|)^{1-\ell-j} \\
 &\leq \frac{1}{|z - \lambda|} \max_{w: |w-\lambda| \leq 1-|\lambda|} \frac{|f(w)| k(k \wedge d + 1)}{(1 - |\lambda|)^{k-1}},
 \end{aligned}$$

where (i) follows since follows since $\min_{z \in \mathbb{T}} |z - \lambda| = 1 - |\lambda|$ for $\lambda \in \mathbb{D}$, and the last inequality comes from taking the maximum over the at most $\frac{k(k \wedge d + 1)}{2} \leq k(k \wedge d + 1)$ terms in the double sum.

Lastly, we complete the argument by turning to the proof of Lemma J.5.

Proof [Proof of Lemma J.5] By rotation invariance of the integral, we may assume λ is real and non-negative. Fix a $\theta_0 \in (0, \frac{\pi}{2}]$ to be chosen later. Then, we decompose our integral as

$$\begin{aligned}
 \int_0^{2\pi} \frac{1}{|e^{i\theta} - \lambda|^2} d\theta &= \int_0^{2\pi} \frac{1}{\sin^2 \theta + (1 - \lambda \cos \theta)^2} d\theta \\
 &\stackrel{(i)}{=} 2 \int_0^{\pi} \frac{1}{\sin^2 \theta + (1 - \lambda \cos \theta)^2} d\theta \\
 &\stackrel{(ii)}{\leq} 2 \left(\frac{\theta_0}{(1 - \lambda)^2} + \left(\int_{\theta_0}^{\pi/2} \frac{1}{\sin^2 \theta} d\theta \right) + \pi \right) \\
 &\stackrel{(iii)}{=} 2 \left(\frac{\theta_0}{(1 - \lambda)^2} + \cot \theta_0 + \pi \right),
 \end{aligned}$$

where (i) uses the symmetry of the integral, and (ii) breaks the integral into $[0, \theta_0]$, $[\theta_0, \pi]$, and $[\pi, 2\pi]$, bounding the integrand above by $\frac{1}{1-\lambda}$, $\frac{1}{\sin^2 \theta}$, and 1 on each respective portion.

Now, setting $\theta_0 = \arcsin(1 - \lambda) \in (0, \frac{\pi}{2}]$, we have that $\cot \theta_0 \leq 1/\sin \theta_0 = 1/(1 - \lambda)$. Moreover, $\sin \theta_0 = 1 - \lambda$, so

$$\frac{\theta_0}{(1 - \lambda)^2} = \frac{1}{1 - \lambda} \cdot \frac{\theta_0}{\sin \theta_0} \leq \frac{1}{1 - \lambda},$$

since $\sin x \leq x$. Combining the above bounds,

$$\frac{1}{2\pi} \int_0^{2\pi} \frac{1}{|e^{i\theta} - \lambda|^2} d\theta \leq \frac{2}{\pi(1 - \lambda)} + 1 \leq \frac{1}{\lambda} \left(1 + \frac{2}{\pi} \right),$$

as needed. ■

J.2. Bounds on Finite System Norms: Proof of Proposition I.2

Here we prove a bound on the operator norm of a Markov matrix, in terms of the terms M, \widetilde{M}, K_2 in (I.29), (I.28), and (I.30) respectively. We recall the proposition we aim to prove:

Proposition I.2 (Bound on Magnitude of Markov Parameters) *Consider a dynamical system of the form $\mathbf{G} = (A_\star, B, C, D)$, where $A_\star = SJ_\star S^{-1}$ is in Jordan normal form. Then, for all $n \geq 1$,*

$$\|\mathcal{M}_n(\mathbf{G})\|_{\text{op}} \leq \|\mathcal{M}_{n+1}(\mathbf{G})\|_{\text{op}} \leq \|D\|_{\text{op}} + \|S^{-1}B\|_{\text{op}}\|CS\|_{\text{op}}K_2(n).$$

Proof We have that

$$\begin{aligned} \|\mathcal{M}_{n+1}(\mathbf{G})\|_{\text{op}} &= \|[D \mid CB \mid CA_\star B \mid \cdots \mid CA_\star^{n-1} B]\|_{\text{op}} \\ &\leq \|D\|_{\text{op}} + \|[CB \mid CA_\star B \mid \cdots \mid CA_\star^{n-1} B]\|_{\text{op}} \\ &= \|D\|_{\text{op}} + \|[CSS^{-1}B_\star \mid CSJ_\star S^{-1}B \mid \cdots \mid CSJ_\star^{n-1}S^{-1}B_\star]\|_{\text{op}} \\ &= \|D\|_{\text{op}} + \|CS \cdot [I \mid J_\star \mid \cdots \mid J_\star^{n-1}] \cdot (I_{n \times n} \otimes S^{-1}B)\|_{\text{op}} \\ &\leq \|D\|_{\text{op}} + \|CS\|_{\text{op}} \cdot \|[I \mid J_\star \mid \cdots \mid J_\star^{n-1}]\|_{\text{op}} \|I_{n \times n} \otimes S^{-1}B\|_{\text{op}} \\ &= \|D\|_{\text{op}} + \|CS\|_{\text{op}}\|S^{-1}B\|_{\text{op}} \cdot \|[I \mid J_\star \mid \cdots \mid J_\star^{n-1}]\|_{\text{op}}. \end{aligned}$$

Next, we see that since J_\star^t is block diagonal, the operator norm of $[I \mid J_\star \mid \cdots \mid J_\star^{n-1}]$ is equal to the largest operator norm of a block row corresponding to one the blocks of J_\star ⁹. Consequently, it suffices to prove that $J_\star = J$ consists of a single Jordan block $(\lambda, k) \in \text{blkspec}(A_\star)$. We shall first consider a bound that holds for $|\lambda| < 1$, and then a general bound for arbitrary λ . For $|\lambda| < 1$, by Lemma I.10 we have

$$\|[I \mid J \mid \cdots \mid J^{n-1}]\|_{\text{op}} \leq \lim_{n \rightarrow \infty} \|[I \mid J \mid \cdots \mid J^{n-1}]\|_{\text{op}} = \|(zI - J)^{-1}\|_{\mathcal{H}_2^{\text{op}}}.$$

Applying Proposition I.11 with the trivial polynomial $f(z) = 1$, which has degree $d = 0$, gives

$$\|(zI - J)^{-1}\|_{\mathcal{H}_2^{\text{op}}} \leq c_{\mathcal{H}_2} \max_{z: |z-\lambda| \leq 1-|\lambda|} \frac{|f(z)|k(k \wedge d+1)}{(1-|\lambda|)^{k-1/2}} = \frac{k}{(1-|\lambda|)^{k-1/2}}.$$

Now for $\lambda \in [0, 1]$, we need to bound

$$\|[I \mid J \mid \cdots \mid J^{n-1}]\|_{\text{op}}^2 \leq \widetilde{M}(k, n)^2,$$

where we recall the definition

$$\widetilde{M}(k, n) := \begin{cases} n^{1/2} & k = 1 \\ n^{k-1/2} \left(\frac{e}{k-1}\right)^{k-1} & 2 \leq k \leq n+1 \\ n^{1/2} 2^n & k \geq n+1 \end{cases}.$$

9. Indeed, the operator norm is invariant under permutations of rows and columns, and since J_\star^t is block diagonal, one can permute the columns of $[I \mid J_\star \mid \cdots \mid J_\star^{n-1}]$ to render it a block diagonal (rectangular) matrix. It is then well known that the operator norm rectangular block diagonal operators is equal to the operator norm of its largest block.

Recalling the formula for the powers of Jordan blocks, we have that

$$J^t = \begin{bmatrix} \lambda^t & \binom{t}{1}\lambda^{t-1} & \binom{t}{2}\lambda^{t-2} & \dots & \binom{t}{k-1}\lambda^{t-(k-1)} \\ 0 & \lambda^t & \binom{t}{1}\lambda^{t-1} & \dots & \binom{t}{k-2}\lambda^{t-(k-2)} \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ \vdots & & \ddots & \lambda^t & \binom{t}{1}\lambda^{t-1} \\ 0 & \dots & \dots & 0 & \lambda^t \end{bmatrix}, \quad (\text{J.38})$$

where we use the convention $\binom{t}{j} = 0$ for $j > t$. Bounding $\|J^t\|_{\text{op}}$ by the ℓ_1 norm of its first row gives

$$\|J^t\|_{\text{op}} \leq \sum_{j=0}^{k-1} \binom{t}{j} |\lambda|^{t-j} \mathbb{I}(t \geq j) = \sum_{j=0}^{k-1} \alpha_{j,t}, \quad (\text{J.39})$$

where $\alpha_{t,j} = \binom{t}{j} |\lambda|^{t-j} \mathbb{I}(t \geq j) \leq \binom{t}{j} \mathbb{I}(t \geq j)$. Since $\alpha_{t,j}$ is increasing in t , we can use the crude bound

$$\begin{aligned} \|[I | J | \dots | J^{n-2}]\|_{\text{op}}^2 &\leq \sum_{t=0}^{n-1} \|J^t\|_{\text{op}}^2 \leq \sum_{t=0}^{n-1} \left(\sum_{j=0}^{k-1} \alpha_{t,j} \right)^2 \\ &\leq n \left(\sum_{j=0}^{k-1} \alpha_{n,j} \right)^2 = n \left(\sum_{j=0}^{k-1 \wedge n} \binom{n}{j} \right)^2. \end{aligned}$$

For $k \geq n + 1$, $\sum_{j=0}^{k-1} \binom{n}{j} = 2^n$, yielding a bound of $n(2^n)^2$. For $k = 1$, the above sum is n , and for $2 \leq k \leq n + 1$, we have the standard bound $\sum_{j=0}^{k-1} \binom{n}{j} \leq \left(\frac{en}{k-1}\right)^{k-1}$, yielding a bound of $(n^{k-\frac{1}{2}} \left(\frac{e}{k-1}\right)^{k-1})^2$. Taking a square root of each of the three cases concludes the proof. \blacksquare

Appendix K. Bounds under Strong Observability

In this section, we formally define a notion called *strong observability*, inspired by the control theory community, which describes how difficult it is to estimate the hidden state in a linear system with *known* dynamics. We then use this notion to develop a bound on Opt_μ in terms of the quantities introduced. We begin by defining the d -step observability matrix

$$\mathcal{O}_d(A, C) := [C^\top | (CA)^\top | \dots | (CA^{d-1})^\top]^\top \in \mathbb{R}^{md \times n}$$

for conforming A, C . Furthermore, we introduce the definition of an invariant decomposition:

Definition K.1 (Invariant Decomposition) We say that (A_+, A_-, C_+, C_-) is an invariant decomposition of (A_*, C_*) if $A_* = A_+ + A_-$, $C_* = C_+ + C_-$, $A_+A_- = A_-A_+ = 0$, and $\text{rowspace}(C_+) \subset \text{rowspace}(A_+)$ and C_-A_+ .

In other words, A_+ and A_- contain complementary invariant subspaces of A_* , C_+ provides information only about A_+ , and C_- provides no information about A_+ .¹⁰ One should associate A_+ with large dynamical modes we need to filter and A_- with smaller modes we can disregard. Strong observability is then defined as follows.

Definition K.2 (Strong Observability) Given a pair (A_+, C_+) with $\text{rowspace}(C_+) \subseteq \text{rowspace}(A_+)$, we say that (A_+, C_+) is (σ, T, d) -strongly observable if $\sigma_{n_+}(\mathcal{O}_d(A_+^T, C_+)) \geq \sigma > 0$, with $n_+ := \text{rank}(A_+)$.

Here, $\sigma_k(\cdot)$ denotes the k -th largest singular value. Strong observability states that given d observations sampled every T time steps, one can reconstruct the hidden state \mathbf{x}_t in a numerically stable fashion. Restricted to the pair (A_+, C_+) , strong observability is a quantitative version of a fundamental *observability* condition in control theory, and state estimation in particular (see, e.g. Hautus (1983)). As an example, one can show that the transition matrix $A_* = \begin{bmatrix} 1 & \Delta \\ 0 & 1 \end{bmatrix}$ and observation matrix $C_* = [1 \ 0]$, which correspond to Newton's equation $F = m\ddot{x}$ when the position x is observed, satisfies $(\sigma, T, 2)$ -strong observability whenever $T\Delta$ is bounded away from zero.

We begin by stating a simplified bound on Opt_μ under the strong observability condition, in terms of the control-theoretic norm $\|\mathcal{M}_\infty(\cdot)\|_{\text{op}}$.

Proposition K.1 (Bounds for Strong Observability) Let $d \leq L$, $N \geq Ld \log(1/\delta)$ and (A_+, A_-, C_+, C_-) be an invariant decomposition of (A_*, C_*) , with $\rho(A_-) < 1$. Define the systems

$$\mathbf{G}_- = (A_-, B_*, C_-, 0), \quad \text{and} \quad \mathbf{F}_- = (A_-, B_w, C_-, 0).$$

Then if (A_+, A_-) is (σ, T, d) -strongly observable for (A_+, A_-) , then with probability at least $1 - \delta$,

$$N^{-1/2} \text{Opt}_\mu \lesssim d \left(1 + \frac{\|C_+ A_+^T\|_{\text{op}}}{\sigma} \right) \left(\sqrt{(m + \log \frac{1}{\delta}) C_{A_-} + C_{Td} + N^{-1/2} \mu} \right), \quad \text{where}$$

$$C_{A_-} := \|\mathcal{M}_\infty(\mathbf{G}_-)\|_{\text{op}} + \|\mathcal{M}_\infty(\mathbf{F}_-)\|_{\text{op}} \quad \text{and} \quad C_{Td} := \|\mathcal{M}_{Td}(\mathbf{G}_*)\|_{\text{op}} + \|\mathcal{M}_{Td}(\mathbf{F}_*)\|_{\text{op}} + \|D_z\|_{\text{op}}.$$

If L is greater than the degree d_+ of the minimal polynomial of A_+ , then the above bound also holds with $d(1 + \|C_+ A_+^T\|_{\text{op}}/\sigma)$ (resp. d) replaced by 2^{d_+} (resp. d_+), even if strong observability fails.

We note that by choosing the invariant partition (A_+, A_-) to ensure that A_- is stable (i.e., placing all unstable modes into A_+), then the operator norms of the infinite-horizon quantities \mathbf{G}_- and \mathbf{F}_- are finite; moreover, by placing near-unstable modes into A_+ , one can obviate the dependence on instability in these terms as well.

In the following subsection, we shall a state more precise variant of the above bound, including analogues for adversarial noise. Subsequent subsections contain the deferred proofs.

10. We note that if A_+ and A_- satisfies $A_+A_- = A_-A_+ = 0$ and $A_* = A_- + A_+$, then we can obtain an invariant decomposition (A_+, A_-, C_+, C_-) by letting $C_+ = \text{Proj}_{A_*} C$ is the matrix obtained by projecting C onto the rowspace of A_+ , and $C_- = C_* - C_+$ the projection onto its complement.

K.1. Granular Bounds for Strong Observability

In this section, we present bounds under the strong observability criterion, Definition K.2. We define the corresponding *stochastic observability error* term,

$$\begin{aligned} \text{ObsErr}_{\text{stoc}}(\delta, m, N) &:= \sqrt{N}(\|\mathcal{M}_\infty(\mathbf{G}_-)\|_{\text{op}} + \|\mathcal{M}_\infty(\mathbf{F}_-)\|_{\text{op}}) + \|\mathcal{M}_\infty(\mathbf{H}_-)\|_{\text{op}} \\ &\quad + \sqrt{m + \log(1/\delta)}(\Gamma_N(\mathbf{G}_-) + \Gamma_N(\mathbf{F}_-)) \\ &\quad + \sqrt{N}(\|\mathcal{M}_{Td}(\mathbf{G}_\star)\|_{\text{op}} + \|\mathcal{M}_{Td}(\mathbf{F}_\star)\|_{\text{op}} + \|D_z\|_{\text{op}}), \end{aligned}$$

and the *adversarial observability error* term,

$$\begin{aligned} \text{ObsErr}_{\text{adv}}(\delta, m, N) &:= \sqrt{N}\|\mathcal{M}_\infty(\mathbf{G}_-)\|_{\text{op}} + \sqrt{Nd_w}\Gamma_N(\mathbf{F}_-) + \|\mathcal{M}_\infty(\mathbf{H}_-)\|_{\text{op}} \\ &\quad + \sqrt{m + \log(1/\delta)}\Gamma_N(\mathbf{G}_-) \\ &\quad + \sqrt{N}\left[\|\mathcal{M}_{Td}(\mathbf{G}_\star)\|_{\text{op}} + \sqrt{Td_w}\|\mathcal{M}_{Td}(\mathbf{F}_\star)\|_{\text{op}} + \sqrt{dd_z}\|D_z\|_{\text{op}}\right], \end{aligned}$$

where we recall the definition

$$\Gamma_N(\mathbf{G}) := \min\{\sqrt{N}\|\mathcal{M}_\infty(\mathbf{G})\|_{\text{op}}, \|\mathbf{G}\|_{\mathcal{H}_\infty}\}$$

Our main theorem is as follows:

Theorem K.2 *Suppose that $N \geq Td \max\{m, \log(1/\delta)\}$. Then, if (A_+, C_+) is (σ, T, d) -strongly observable for $d \leq L$, then with probability $1 - \delta$, we have that in the stochastic model,*

$$\text{Opt}_\mu \leq (\text{ObsErr}_{\text{stoc}}(\delta, m, N) + \mu) \left(1 + \frac{d\|C_+A_+^d\|_{\text{op}}}{\sigma}\right).$$

where the analogous bound holds with $\text{ObsErr}_{\text{adv}}$ under adversarial noise.

In general, even if (A_+, C_+) may not be (σ, d) -strongly observable, we argue as follows. Let f^+ denote the minimal polynomial of A_+ . If $L \geq \deg(f^+)$, then

$$\begin{aligned} \text{Opt}_\mu &\leq (\text{ObsErr}_{\text{stoc}}(\delta, m, N) + \mu)\|f^+\|_{L_1} \\ &\leq (\text{ObsErr}_{\text{stoc}}(\delta, m, N) + \mu)2^{\deg(f^+)} \leq (\text{ObsErr}_{\text{stoc}}(\delta, m, N) + \mu)2^{\text{rank}(A_+)}, \end{aligned}$$

and analogously for adversarial noise.

Proposition K.1 follows directly by bounding $\Gamma_N(\mathbf{G}) \leq \sqrt{N}\|\mathcal{M}_\infty(\mathbf{G})\|_{\text{op}}$ and dropping the \mathbf{H}_- term under the assumption made in the body that $\mathbf{x}_1 = 0$. We now turn the proof of the theorem. The above theorem is prove in the following subsection.

K.2. Proof of Theorem K.2

Let $\phi = [\Psi_1 | \dots | \Psi_d | \mathbf{0}]$; in view of the discussion in Section G, we can assume $L = d$ for our analysis. Since $A_+A_- = A_-A_+ = 0$, it follows that for any power $k \geq 1$, $A_\star^k = A_+^k + A_-^k$. Moreover, since $\text{rowspan}(C_+) \subset \text{rowspan}(A_+) \in \ker(A_-)^\top$ (and similarly when the signs are swapped), it follows that $C_+A_- = C_-A_+ = 0$. It then follows that

$$C_\star A_\star^k = C_+A_+^k + C_-A_-^k, \quad \text{for all } k \geq 0.$$

This allows us to decompose C_ϕ as $C_\phi := C_\phi^+ + C_\phi^-$, where for $\sigma \in \{+, -\}$

$$C_\phi^\sigma := C_\sigma A_\sigma^{LT} - \sum_{\ell=1}^L \Psi_\ell C_\sigma A_\sigma^{(L-\ell)T} \in \mathbb{R}^{m \times n}.$$

The following proposition gives us a control on $\|\Delta_\phi\|_{\text{op}}$ in terms of the systems C_ϕ^σ . Its proof is deferred to the following subsection:

Proposition K.3 *Suppose that $\phi \in \mathbb{R}^{m \times md}$ satisfies $C_\phi[A_+] = 0$ and that $N \geq Td \max\{m, \log(1/\delta)\}$. Then, in the stochastic noise model,*

$$\|\Delta_\phi\|_{\text{op}} \lesssim (1 + \|\phi\|_{\text{bop}}) \text{ObsErr}_{\text{stoc}}(\delta, m, N),$$

and in the adversarial noise model,

$$\|\Delta_\phi\|_{\text{op}} \lesssim (1 + \|\phi\|_{\text{bop}}) \text{ObsErr}_{\text{adv}}(\delta, m, N).$$

Hence, for stochastic noise, we get that with probability at least $1 - \delta$,

$$\text{Opt}_\mu \leq \mu \|\phi\|_{\text{op}} + \|\Delta_\phi\|_{\text{op}} \leq (1 + \|\phi\|_{\text{bop}})(\mu + \text{ObsErr}_{\text{stoc}}(\delta, m, N)),$$

and similarly for adversarial noise. To conclude, it remains to bound $\|\phi\|_{\text{bop}}$. Considering an invariant partition $A_\star = A_+ + A_-$, we invoke the following lemma, which is a consequence of the Moore-Penrose pseudoinverse.

Lemma K.4 *If (A_+, C_+) is (σ, T, d) -strongly observable, then there exists a matrix $\phi = [\Psi_1 | \dots | \Psi_d] \in \mathbb{R}^{m \times md}$ satisfying $C_+ A_+^{Td} - \sum_{\ell=1}^d \Psi_\ell C_+ A_+^{T\ell} = 0$ and $\|\phi\|_{\text{op}} \leq \frac{\|C_+ A_+^{Td}\|_{\text{op}}}{\sigma}$.*

Using this lemma (proved in Appendix K.3.1) and the error bounds from Proposition 4.1 below, we conclude the proof of Proposition K.1 in Appendix K.3.2. This lemma directly yields the first part of our theorem, since

$$1 + \|\phi\|_{\text{bop}} \leq (1 + d\|\phi\|_{\text{op}}) \leq 1 + \frac{d\|C_+ A_+^d\|_{\text{op}}}{\sigma}.$$

Proof [Proof of Lemma K.4] Recall the observability matrix

$$\mathcal{O}_d(A_+, C_+) = \begin{bmatrix} C_+ \\ C_+ A_+ \\ \vdots \\ C_+ A_+^{d-1} \end{bmatrix}.$$

Observe that since $\text{rowspace}(C_+) \subset \text{rowspace}(A_+)$, $\mathcal{O}_d(A_+, C_+)$ has rank at most $\dim(\text{rowspace}(A_+)) = \text{rank}(A_+) = n_+$. By assumption, $\mathcal{O}_d(A_+, C_+)$ has rank at least n_+ as well. It follows that $\text{rowspace}(\mathcal{O}_d(A_+, C_+)) = \text{rowspace}(A_+)$, and therefore $C_+ A_+^d \in \text{range}(\mathcal{O}_d(A_+, C_+))$. This implies that for the filter

$$\phi = C_+ A_+^d \mathcal{O}_d(A_+, C_+)^{\dagger} \in \mathbb{R}^{m \times dm},$$

we have

$$\phi \mathcal{O}_d(A_+, C_+) = C_+ A_+^d \mathcal{O}_d(A_+, C_+)^\dagger \mathcal{O}_d(A_+, C_+) = C_+ A_+^d.$$

The above then implies that

$$C_\phi^+ = C_+ A_+^d - \phi \cdot \mathcal{O}_d(A_+, C_+) = 0.$$

Moreover, as we have established that $\text{rank}(\mathcal{O}_d(A_+, C_\star)) = n_+$, we have

$$\begin{aligned} \|\phi\|_{\text{op}} &= \|C_+ A_+^d \mathcal{O}_d(A_+, C_+)^\dagger\|_{\text{op}} \\ &\leq \|C_+ A_+^d\|_{\text{op}} \|\mathcal{O}_d(A_+, C_+)^\dagger\|_{\text{op}} \leq \|C_+ A_+^d\|_{\text{op}} \sigma_{n_+}(\mathcal{O}_d(A_+, C_+))^{-1}. \end{aligned}$$

Recalling that our strong observability assumption implies that $\sigma_{n_+}(\mathcal{O}_d(A_+, C_+))^{-1} \geq \sigma$, we find that the above quantity is at most $\|C_+ A_+^d\|_{\text{op}}/\sigma$ by assumption. \blacksquare

The second part of our theorem follows by constructing ϕ to use the minimal polynomial of A_+ .

Lemma K.5 *There exists a ϕ with $1 + \|\phi\|_{\text{bop}} = \|f^+\|_{\ell_1} \leq 2^{\deg(f_+)} \leq 2^{n_+}$ and $C_\phi^+ = 0$, where f^+ is the the minimal polynomial of A_+ .*

Proof Let f^+ denote the minimal polynomial of A_+ , and let

$$\phi = [f_1^+ I_m | f_2^+ I_m | \cdots | f_{\deg(f_+)}^+ I_m | \mathbf{0}].$$

Then, $C_\phi^+ = C_+ f_+(A_+) = 0$. Since f^+ has all of its roots in the complex disk (as the spectrum of A_+ is a subset of the spectrum of A_\star), Lemma J.1 implies that $1 + \|\phi\|_{\text{bop}} = \|f^+\|_{\ell_1} \leq 2^{\deg(f_+)} \leq 2^{n_+}$. \blacksquare

These two possible methods of bounding $\|\phi\|_{\text{bop}}$ conclude the proof.

K.3. Proof of Proposition K.3

Again, we assume $L = d$. Using the properties of an invariant decomposition, we can modify our error calculations as follows; the proof is stated in Section K.3.2.

Lemma K.6 *Suppose that $C_\phi[A_+] = 0$. Then, the conclusions of Propositions G.1 and G.2 hold with \mathbf{G}_ϕ , \mathbf{F}_ϕ and \mathbf{H}_ϕ replaced by the following systems:*

$$\mathbf{G}_\phi^- := (A_-, B_\star, C_\phi^-, 0), \quad \mathbf{F}_\phi^- := (A_-, B_w, C_\phi^-, 0), \quad \mathbf{H}_\phi^- := (A_-, \mathbf{x}_1, C_\phi^-, 0).$$

We can now bound the control norms of these ‘‘minus’’-systems (proof in Section K.3.1):

Lemma K.7 $\|\mathcal{M}_\infty(\mathbf{G}_\phi^-)\|_{\text{op}} \leq (1 + \|\phi\|_{\text{bop}})\|\mathcal{M}_\infty(\mathbf{G}_-)\|_{\text{op}}$ and $\|(\mathbf{G}_\phi^-)^-\|_{\mathcal{H}_\infty} \leq (1 + \|\phi\|_{\text{bop}})\|\mathbf{G}_-\|_{\mathcal{H}_\infty}$, and similarly for the corresponding \mathbf{F} and \mathbf{H} systems.

Now bounding $\Gamma_N(\mathbf{G}) \leq \sqrt{N}\|\mathcal{M}_N(\mathbf{G})\|_{\text{op}}$, we can simplify the bound in the stochastic model to obtain

$$\begin{aligned} \frac{\|\Delta_\phi\|_{\text{op}}}{1 + \|\phi\|_{\text{bop}}} &\lesssim \sqrt{N(m + \log \frac{1}{\delta})} (\|\mathcal{M}_\infty(\mathbf{G}_-)\|_{\text{op}} + \|\mathcal{M}_\infty(\mathbf{F}_-)\|_{\text{op}}) \\ &\quad + \sqrt{N} (\|\mathcal{M}_{Td}(\mathbf{G}_\star)\|_{\text{op}} + \|\mathcal{M}_{Td}(\mathbf{F}_\star)\|_{\text{op}} + \|D_z\|_{\text{op}}). \end{aligned}$$

K.3.1. PROOF OF LEMMA K.7

Let us take \mathbf{G}_ϕ^- as a representative example. Using the formula

$$C_\phi^- := C_- A_-^{LT} - \sum_{\ell=1}^L \Psi_\ell C_- A_-^{(L-\ell)T},$$

we have (with the convention $\Psi_0 = I_m$),

$$\begin{aligned} \|\mathcal{M}_\infty(\mathbf{G}_\phi^-)\|_{\text{op}} &= \lim_{n \rightarrow \infty} \|\mathcal{M}_n(\mathbf{G}_\phi^-)\|_{\text{op}} \\ &= \lim_{n \rightarrow \infty} \sum_{\ell=0}^L \|\mathcal{M}_n((A_-, B_\star, \Psi_{L-\ell} C_- A_-^\ell, 0))\|_{\text{op}} \\ &= \lim_{n \rightarrow \infty} \sum_{\ell=0}^L \|\Psi_{L-\ell}\|_{\text{op}} \|\mathcal{M}_n((A_-, B_\star, C_- A_-^\ell, 0))\|_{\text{op}} \\ &\stackrel{(i)}{\leq} \lim_{n \rightarrow \infty} \sum_{\ell=0}^L \|\Psi_{L-\ell}\|_{\text{op}} \|\mathcal{M}_{n+\ell}((A_-, B_\star, C_-, 0))\|_{\text{op}} \\ &= \lim_{n \rightarrow \infty} \sum_{\ell=0}^L \|\Psi_{L-\ell}\|_{\text{op}} \|\mathcal{M}_{n+\ell}(\mathbf{G}_-)\|_{\text{op}} \\ &= \sum_{\ell=0}^L \|\Psi_{L-\ell}\|_{\text{op}} \lim_{n \rightarrow \infty} \|\mathcal{M}_{n+\ell}(\mathbf{G}_-)\|_{\text{op}} \\ &= \|\mathcal{M}_\infty(\mathbf{G}_-)\|_{\text{op}} \cdot \sum_{\ell=0}^L \|\Psi_{L-\ell}\|_{\text{op}} = \|\mathcal{M}_\infty(\mathbf{G}_-)\|_{\text{op}} (1 + \|\phi\|_{\text{bop}}), \end{aligned}$$

where in (i) we have used the fact that $\mathcal{M}_n((A_-, B_\star, C_- A_-^\ell, 0))$ is a submatrix of the matrix $\mathcal{M}_{n+\ell}((A_-, B_\star, C_-, 0))$. This argument can be applied for the \mathcal{H}_∞ -norm, viewed as the asymptotic limit of the operator norm of the associated Toeplitz operator (see e.g. [Tilli \(1998\)](#) Corollary 4.2).

K.3.2. PROOF OF LEMMA K.6

Examining the arguments in Section G, it suffices to modify the control of the term $\text{Err}^{(1)}$. Recall the shut-off sequence

$$\begin{aligned} \tilde{\mathbf{x}}_{n;t} &:= \begin{cases} A_\star^{n-(t-LT)} \mathbf{x}_{t-LT} & n \geq t - LT \\ \mathbf{x}_n & n \leq t - LT \end{cases} \\ \tilde{\mathbf{y}}_{n;t} &:= C_\star \tilde{\mathbf{x}}_{n;t} \\ \tilde{\mathbf{k}}_t &:= [\tilde{\mathbf{y}}_{t-T;t}^\top \mid \tilde{\mathbf{y}}_{t-2T;t}^\top \mid \cdots \mid \tilde{\mathbf{y}}_{t-LT;t}^\top]^\top, \end{aligned}$$

Defining the terms $\gamma_t := B_w \mathbf{w}_t + B_* \mathbf{u}_t$, since $C_\phi[A_*] = C_\phi^-$ we have that

$$\begin{aligned}
 \text{Err}_t^{(1)} &= \tilde{\mathbf{y}}_{t;t} - \phi \cdot \tilde{\mathbf{k}}_t = C_\phi[A_*] \tilde{\mathbf{x}}_{t-TL} = C_\phi^- \tilde{\mathbf{x}}_{t-TL} \\
 &= C_\phi^- \left(\sum_{i=0}^{t-LT-2} A_*^i \gamma_{t-1-i} + A_*^{t-TL-1} \mathbf{x}_1 \right) \\
 &= C_\phi^- \left(\gamma_{t-1} + \sum_{i=1}^{t-LT-2} A_*^i \gamma_{t-1-i} + A_*^{t-TL-1} \mathbf{x}_1 \right) \\
 &= C_\phi^- \left(\gamma_{t-1} + \sum_{\sigma \in \{+, -\}} \sum_{i=1}^{t-LT-2} (A_\sigma^i) \gamma_{t-1-i} + A_\sigma^{t-TL-1} \mathbf{x}_1 \right).
 \end{aligned}$$

Observe now that the term corresponding to $\sigma = +$ is canceled by C_ϕ^- , because it involves only terms which have the products $A_- A_+$ or $C_- A_+$. Thus,

$$\text{Err}_t^{(1)} = C_\phi'[A_-] \left(\gamma_{t-1} + A_-^i \gamma_{t-1-i} + A_-^{t-TL-1} \mathbf{x}_1 \right).$$

This term can be then be controlled analogously to the term $\text{Err}_t^{(1)}$ in Section G, replacing A_* with A_- and C_ϕ with C_ϕ^- .