
Adaptive Principal Component Regression with Applications to Panel Data

Anish Agarwal*
Department of IEOR
Columbia University
New York, NY 10027
aa5194@columbia.edu

Keegan Harris
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213
keeganh@cs.cmu.edu

Justin Whitehouse
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213
jwhiteho@cs.cmu.edu

Zhiwei Steven Wu
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213
zhiweiw@cs.cmu.edu

Abstract

Principal component regression (PCR) is a popular technique for fixed-design *error-in-variables* regression, a generalization of the linear regression setting in which the observed covariates are corrupted with random noise. We provide the first time-uniform finite sample guarantees for online (regularized) PCR whenever data is collected *adaptively*. Since the proof techniques for analyzing PCR in the fixed design setting do not readily extend to the online setting, our results rely on adapting tools from modern martingale concentration to the error-in-variables setting. As an application of our bounds, we provide a framework for experiment design in panel data settings when interventions are assigned adaptively. Our framework may be thought of as a generalization of the synthetic control and synthetic interventions frameworks, where data is collected via an adaptive intervention assignment policy.

1 Introduction

An omnipresent task in machine learning, statistics, and econometrics is that of making predictions about outcomes of interest given an action and conditioned on observable covariates. An often overlooked aspect of the prediction task is that in many settings the learner only has access to *imperfect observations* of the covariates, due to e.g. measurement error or inherent randomness in the problem domain. Such settings are sometimes formulated as *error-in-variables* regression: a learner is given access to a collection of data $(Z_n, a_n, Y_n)_{n \geq 1}$, where $Z_n \in \mathbb{R}^d$ are the *observed covariates*, $a_n \in \{1, \dots, A\}$ is the *action taken*, and $Y_n \in \mathbb{R}$ is the *outcome* for each observation n . Typically, the outcomes are assumed to be generated by a linear model $Y_n := \langle \theta(a_n), X_n \rangle + \xi_n$ and $Z_n := X_n + \epsilon_n$, where $X_n \in \mathbb{R}^d$ are the *true covariates*, $\epsilon_n \in \mathbb{R}^d$ is the covariate noise, $\theta(a_n) \in \mathbb{R}^d$ is an *unknown slope vector* associated with action a_n , and $\xi_n \in \mathbb{R}$ is the response noise. Note that the learner does not get to see the true covariates X_n . Observe that when $\epsilon_n = 0$ we recover the traditional linear regression setting. Such an error-in-variables model can encompass many forms of data corruption including measurement error, missing values, discretization, and differential privacy—see [10, 5] for details.

Our point of departure from previous work is that we allow the sequence of data $(Z_n, a_n, Y_n)_{n \geq 1}$ to be chosen *adaptively*. In other words, we provide bounds for learning in the error-in-variables regres-

*For part of this work, Anish was a postdoc at Amazon, Core AI.

sion setting when the data seen at the current round n is allowed to depend on the previously-seen data $(Z_m, a_m, Y_m)_{1 \leq m < n}$. Adaptive data collection occurs when the choices of future observations can depend on the inference from previous data, which is common in learning paradigms such as multi-armed bandits [58, 78], active learning [74], and time-series analysis [76, 36]. Similar to prior work on adaptive data collection that shows that valid statistical inference can be done when the true covariates are observed [34, 67, 43, 91, 92], our work provides the first time-uniform finite sample guarantees for error-in-variables regression using adaptively collected data.

Concretely, we focus on analyzing *principal component regression (PCR)* [54], a method that has been shown to be effective for learning from noisy covariates [10, 7] and a central tool for learning from panel data [8, 5, 7]. At a high level, PCR “de-noises” the sequence of observed covariates $(Z_n)_{n \geq 1}$ as $(\hat{Z}_n)_{n \geq 1}$ by performing hard singular value thresholding, after which a linear model is learned using the observed outcomes $(Y_n)_{n \geq 1}$ and the denoised covariates $(\hat{Z}_n)_{n \geq 1}$. See Section 3.2 for further technical background on PCR.

1.1 Contributions

1. We derive novel time-uniform bounds for an online variant of regularized PCR when the sequence of covariates is chosen adaptively. The techniques used to derive bounds for PCR in the fixed-sample regime [7] do not extend to the setting in which data is collected adaptively; thus, we require new tools and ideas to obtain our results. Specifically, our results rely on applying recent advances in martingale concentration [50, 51], as well as more classical results on self-normalized concentration [3, 32, 33] which are commonly applied to online regression problems, to the error-in-variables setting. As an example of the bounds we obtain, consider the task of estimating the underlying relationship $\theta(a)$ between true (i.e. noiseless) covariates and observations, given access to n adaptively-chosen noisy covariates and their corresponding actions and observations. The ℓ_2 estimation error of the adaptive PCR estimator $\hat{\theta}_n(a)$ can be bounded as

$$\|\hat{\theta}_n(a) - \theta(a)\|_2^2 = \tilde{O} \left(\frac{1}{\text{snr}_n(a)^2} \kappa(\mathbf{X}_n(a))^2 \right)$$

with high probability, where $\text{snr}_n(a)$ is the *signal-to-noise ratio* associated with action a at round n (Definition 3.5), a measure of how well the true covariates stand out from the noise. $\kappa(\mathbf{X}_n(a))$ is the *condition number* of the true covariates. Intuitively, if $\text{snr}_n(a)$ is high the true covariates can be well-separated from the noise, and therefore PCR accurately estimates $\theta(a)$ as long as the true covariates are well-conditioned.

Despite the harder setting we consider, our PCR bounds for adaptively-collected data largely match the bounds currently known for the fixed-sample regime, and even improve upon them in two important ways: (1) Our bounds are *computable*, i.e. they depend on *known* constants and quantities available to the algorithm. (2) Unlike Agarwal et al. [7], our bounds do not depend on the ℓ_1 -norm of $\theta(a)$, i.e., we do not require approximate sparsity of $\theta(a)$ for the bounds to imply consistency. This is important because PCR is a rotationally-invariant algorithm, and so its performance guarantees should not depend on the orientation of the basis representation of the space to be learned. The price we pay for adaptivity is that $\text{snr}_n(a)$ is defined with respect to a bound on the *total* amount of noise seen by the algorithm so far, instead of just the noise associated with the rounds that a is taken. As a result, our bound for $\hat{\theta}_n(a)$ may not be tight if a is seldomly selected.

2. We apply our PCR results to the problem of online experiment design with *panel data*. In panel data settings, the learner observes repeated, noisy measurements of *units* (e.g. medical patients, subpopulations, geographic locations) under different *interventions* (e.g. medical treatments, discounts, socioeconomic policies) over *time*. This is an ubiquitous method of data collection, and as a result, learning from panel data has been the subject of significant interest in the econometrics and statistics communities (see Section 2). A popular framework for learning from panel data is *synthetic control (SC)* [1, 2], which uses historical panel data to estimate counterfactual unit measurements under control. Synthetic interventions (SI) [8] is a recent generalization of the SC framework which allows for counterfactual estimation under treatment, in addition to control. By leveraging online PCR, we can perform counterfactual estimation of unit-specific treatment effects under both

treatment and control, as in the SI framework. However, unlike the traditional SI framework, we are the first to establish statistical rates for counterfactual unit outcome estimates under different interventions *while allowing for both units and interventions to be chosen adaptively*. Such adaptivity may naturally occur when treatments are prescribed to new units based on the outcomes of previous units. For example, this is the case when the intervention chosen for each unit is the one which appears to be best based on observations in the past.

2 Related work

Error-in-variables regression There is a rich literature on error-in-variables regression (e.g. [42, 57, 28, 44, 84, 47, 41]), with research focusing on topics such as high-dimensional [62, 56, 31, 72] and Bayesian settings [70, 81, 40]. Principal component regression (PCR) [54, 22, 10, 7] is a popular method for error-in-variables regression. The results of Agarwal et al. [7] are of particular relevance to us, as they provide finite sample guarantees for the fixed design (i.e. non-adaptive) version of the setting we consider.

Self-normalized concentration There has been a recent uptick in the application of self-normalized, martingale concentration to online learning problems. In short, self-normalized concentration aims to control the growth of processes that have been normalized by a random, or empirical, measure of accumulated variance [32, 33, 50, 51, 88]. Self-normalized concentration has led to breakthroughs in wide-ranging areas of machine learning such as differential privacy [85, 86], PAC-Bayesian learning [30], convex divergence estimation [63], and online learning [87, 29, 3]. Of particular importance for our work are the results of Abbasi-Yadkori et al. [3], which leverage self-normalized concentration results for vector-valued processes [32, 33] to construct confidence ellipsoids for online regression tasks. We take inspiration from these results when constructing our estimation error bounds for PCR in the sequel.

Learning in panel data settings Our application to panel data builds off of the SI framework [8, 9], which itself is a generalization of the canonical SC framework for learning from panel data [1, 2, 52, 35, 17, 60, 89, 13, 14, 59, 16, 20, 23, 27, 39]. In both frameworks, a *latent factor model* is often used to encode structure between units and time-steps [25, 61, 15, 18, 19, 68, 64, 65]. Specifically, it is assumed that unit outcomes are the product of unit- and time/intervention-specific latent factors, which capture the heterogeneity across time-steps, units, and interventions, and allows for the estimation of unit-specific counterfactuals under treatment and control. Other extensions of the SI framework include applications in biology [79], network effects [11], combinatorially-many interventions [4], and intervening under incentives [45, 66]. Finally, there is a growing line of work at the intersection of online learning and panel data. Chen [26] views the problem of SC as an instance of online linear regression, which allows them to apply the regret guarantees of the online learning algorithm *follow-the-leader* [55] to show that the predictions of SC are comparable to those of the best-in-hindsight weighted average of control unit outcomes. Farias et al. [38] build on the SC framework to estimate treatment effects in adaptive experimental design, while minimizing the regret associated with experimentation. The results of Farias et al. [38] are part of a growing line of work on counterfactual estimation and experimental design using multi-armed bandits [69, 77, 90, 24].

3 Setting and background

Notation We use boldface symbols to represent matrices. For $N \in \mathbb{N}$, we use the shorthand $[N] := \{1, \dots, N\}$. Unless specified otherwise, $\|v\|$ denotes the ℓ_2 -norm of a vector v , and $\|\mathbf{A}\|_{op}$ the operator norm of matrix \mathbf{A} . We use $\text{diag}(a_1, \dots, a_k)$ to represent a $k \times k$ diagonal matrix with entries a_1, \dots, a_k . For two numbers $a, b \in \mathbb{R}$, we use $a \wedge b$ as shorthand for $\min\{a, b\}$, and $a \vee b$ to mean $\max\{a, b\}$. Finally, \mathbb{S}^{d-1} denotes the d -dimensional unit sphere.

3.1 Problem setup

We now describe our error-in-variables setting. We consider a *learner* who interacts with an *environment* over a sequence of rounds. At the start of each round $n \geq 1$, the environment generates covariates $X_n \in W^* \subset \mathbb{R}^d$, where W^* is a low-dimensional linear subspace of dimension $\dim(W^*) = r < d$.

We assume that r (but not W^*) is known to the learner. Such “low rank” assumptions are reasonable whenever, e.g. data is generated according to a *latent factor model*, a popular assumption in high-dimensional statistical settings [53, 12, 48]. As we will see in Section 5, analogous assumptions are often also made in panel data settings. The learner then observes *noisy* covariates $Z_n := X_n + \epsilon_n$, where $\epsilon_n \in \mathbb{R}^d$ is a random noise vector. Given Z_n , the learner selects an *action* $a_n \in [A]$ and observes $Y_n := \langle \theta(a_n), X_n \rangle + \xi_n$, where ξ_n is random noise and $\theta(a)$ for $a \in [A]$ are unknown slope vectors in W^* that parameterize action choices such that $\|\theta(a)\|_2 \leq L$ for some $L \in \mathbb{R}$. We require that the covariate noise is “well-behaved” according to one of the two the following assumptions:

Assumption 3.1 (SubGaussian Covariate Noise). *For any $n \geq 1$, the noise variable ϵ_n satisfies (a) ϵ_n is σ -subGaussian, (b) $\mathbb{E}\epsilon_n = 0$, and (c) $\|\mathbb{E}\epsilon_n\epsilon_n^\top\|_{op} \leq \gamma$, for some constant $\gamma > 0$.*

Assumption 3.2 (Bounded Covariate Noise). *For any $n \geq 1$, the noise variable ϵ_n satisfies (a) $\|\epsilon_n\| \leq \sqrt{Cd}$, (b) $\mathbb{E}\epsilon_n = 0$, and (c) $\mathbb{E}\epsilon_n\epsilon_n^\top = \Sigma$, for some positive-definite matrix Σ satisfying $\|\Sigma\|_{op} \leq \gamma$, for some constant $\gamma > 0$.*

Note that Assumption 3.2 is a special case of Assumption 3.1, which allows us to get stronger results in some settings. We also impose the following constraint on the noise in the outcomes.

Assumption 3.3 (SubGaussian Outcome Noise). *For any $n \geq 1$, the noise variable ξ_n satisfies (a) $\mathbb{E}\xi_n = 0$, (b) ξ_n is η -subGaussian, and (c) $\mathbb{E}\xi_n^2 \leq \alpha$, for some constant α .*

Under this setting, the goal of the learner is to estimate $\theta(a)$ for $a \in [A]$ given an (possibly adaptively-chosen) observed sequence $(Z_n, a_n, Y_n)_{n \geq 1}$. For $n \geq 1$, we define the matrix $\mathbf{Z}_n \in \mathbb{R}^{n \times d}$ to be the matrix of *observed* (i.e. noisy) covariates, with Z_1, \dots, Z_n as its rows. Similarly, $\mathbf{X}_n = (X_1, \dots, X_n)^\top \in \mathbb{R}^{n \times d}$ is the matrix of *noiseless* covariates (which are unobserved), and $\mathcal{E}_n = (\epsilon_1, \dots, \epsilon_n)^\top \in \mathbb{R}^{n \times d}$, $\mathbf{Y}_n = (Y_1, \dots, Y_n)^\top \in \mathbb{R}^{n \times 1}$, and $\Xi_n = (\xi_1, \dots, \xi_n)^\top \in \mathbb{R}^{n \times 1}$ are defined analogously. For any action $a \in [A]$, let $N_n(a) := \{s \leq n : a_s = a\}$ be the *set of rounds* up to and including round n on which action a was chosen. Likewise, let $c_n(a) := |N_n(a)|$ denote the *number of rounds* by round n on which action a was chosen. For $a \in [A]$, we enumerate $N_n(a)$ in increasing order as $i_1 \leq \dots \leq i_{c_n(a)}$. Finally, we define $\mathbf{Z}_n(a) \in \mathbb{R}^{c_n(a) \times d}$ to be $\mathbf{Z}(a) = (Z_{i_1}, \dots, Z_{i_{c_n(a)}})^\top$, and define $\mathbf{X}_n(a)$, $\mathcal{E}_n(a)$, $\mathbf{Y}_n(a)$, and $\Xi_n(a)$ analogously.

3.2 Principal component regression

Background on singular value decomposition Any matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$ may be written in terms of its singular value decomposition $\mathbf{A} = \mathbf{U}\Sigma\mathbf{V}^\top$, where $\mathbf{U} \in \mathbb{R}^{n \times d \wedge n}$ and $\mathbf{V} \in \mathbb{R}^{d \times d \wedge n}$ are matrices with orthonormal columns, and $\Sigma = \text{diag}(\sigma_1(\mathbf{A}), \dots, \sigma_{d \wedge n}(\mathbf{A})) \in \mathbb{R}^{(d \wedge n) \times (d \wedge n)}$ is a diagonal matrix containing the singular values of \mathbf{A} , where we assume $\sigma_1(\mathbf{A}) \geq \dots \geq \sigma_{d \wedge n}(\mathbf{A}) \geq 0$. Given a *truncation level* k , we define the truncation of \mathbf{A} onto its top k principal components as $\mathbf{A}_k := \mathbf{U}_k \text{diag}(\sigma_1(\mathbf{A}), \dots, \sigma_{k \wedge d \wedge n}(\mathbf{A})) \mathbf{V}_k^\top$, where $\mathbf{U}_k \in \mathbb{R}^{n \times k \wedge d \wedge n}$ is the matrix with the first $k \wedge d \wedge n$ columns of \mathbf{U} , and $\mathbf{V}_k \in \mathbb{R}^{n \times k \wedge d \wedge n}$ is defined analogously. Given such a singular value decomposition, we can define the projection matrix onto the subspace spanned by the top k right singular vectors as $\mathbf{P}_k \in \mathbb{R}^{d \times d}$ given by $\mathbf{P}_k := \mathbf{V}_k \mathbf{V}_k^\top$.

For $n \geq 1$, $a \in [A]$, and $\mathbf{Z}_n(a)$ as defined above, we write the k -truncated singular value decomposition of $\mathbf{Z}_n(a)$ as $\mathbf{Z}_{n,k}(a) = \widehat{\mathbf{U}}_{n,k}(a) \text{diag}(\sigma_1(\mathbf{Z}_n(a)), \dots, \sigma_{k \wedge n \wedge d}(\mathbf{Z}_n(a))) \widehat{\mathbf{V}}_{n,k}^\top(a)$, and the corresponding projection onto the top k right singular vectors of $\mathbf{Z}_n(a)$ as $\widehat{\mathbf{P}}_{n,k}(a)$. When $k = r$, we leverage the simplified notation $\widehat{\mathbf{P}}_n(a) := \widehat{\mathbf{P}}_{n,r}(a)$. (Recall $r = \dim(W^*)$.) By \mathbf{P} , we denote the projection matrix onto the true, underlying subspace W^* . While \mathbf{P} is never known, our results leverage the fact that $\widehat{\mathbf{P}}_n(a)$ converges to \mathbf{P} nicely over time. We define the projected noisy covariate matrix matrix to be $\widehat{\mathbf{Z}}_n(a) := \mathbf{Z}_n(a) \widehat{\mathbf{P}}_n(a)$, and define $\widehat{\mathbf{X}}_n(a)$, $\widehat{\mathcal{E}}_n(a)$ similarly. Any quantity with a “ $\widehat{\cdot}$ ” is defined equivalently to quantities with “ \cdot ”, except with \mathbf{P} in place of $\widehat{\mathbf{P}}_n(a)$. We are now ready to introduce our procedure for estimating $\theta(a)$ for $a \in [A]$, called *adaptive* (or *online*) principal component regression.

Definition 3.4 (Adaptive Principal Component Regression). *Given regularization parameter $\rho \geq 0$ and truncation level $k \in \mathbb{N}$, for $a \in [A]$ and $n \geq 1$ let $\widehat{\mathbf{Z}}_n(a) := \mathbf{Z}_n(a) \widehat{\mathbf{P}}_{n,k}(a)$ and $\widehat{\mathbf{V}}_n(a) := \widehat{\mathbf{Z}}_n(a)^\top \widehat{\mathbf{Z}}_n(a) + \rho \widehat{\mathbf{P}}_{n,k}(a)$. Regularized principal component regression estimates $\theta(a)$ as*

$$\widehat{\theta}_n(a) := \widehat{\mathbf{V}}_n(a)^{-1} \widehat{\mathbf{Z}}_n(a) \mathbf{Y}_n(a).$$

Setting $\rho = 0$ recovers the version of PCR used in Agarwal et al. [7]. In words, (unregularized) PCR “denoises” the observed covariates by projecting them onto the subspace given by their k -truncation, before estimating $\theta(a)$ via linear regression using the projected covariates. We choose to regularize since it is known that regularization increases the stability of regression-style algorithms. This added stability from regularization allows us to exploit the online regression bounds of Abbasi-Yadkori et al. [3] to measure the performance of our estimates. Throughout the sequel, we only consider adaptive PCR with truncation level $k = r$.

3.3 Signal to noise ratio

We now introduce the concept of *signal to noise ratio*, which will be integral in stating and proving our results. The signal to noise ratio provides a measure of how strongly the true covariates (this is the “signal” of the problem, measured through $\sigma_r(\mathbf{X}_n(a))$) stand out sequentially with respect to the “noise” induced by \mathcal{E}_n (which we will measure through the relevant high probability bounds on $\|\mathcal{E}_n\|_{op}$).

Definition 3.5 (Signal to Noise Ratio). We define the signal to noise ratio associated with an action $a \in [A]$ at round n as

$$\text{snr}_n(a) := \frac{\sigma_r(\mathbf{X}_n(a))}{U_n},$$

where $(U_n)_{n \geq 1}$ is a noise-dependent sequence growing as $U_n = O\left(\sqrt{n} + \sqrt{d} + \sqrt{\log\left(\frac{1}{\delta} \log(n)\right)}\right)$.

The price we pay for adaptivity is encoded directly into the definition of the signal to noise ratio, $\text{snr}_n(a)$. While one may imagine defining $\text{snr}_n(a)$ as the ratio between $\sigma_r(\mathbf{X}_n(a))$ and $\|\mathcal{E}_n(a)\|_{op}$, bounding $\|\mathcal{E}_n(a)\|_{op}$ is a nontrivial task as the rows of $\mathcal{E}_n(a)$ may be strongly correlated. To circumvent this, we apply the trivial bound $\|\mathcal{E}_n(a)\|_{op} \leq \|\mathcal{E}_n\|_{op}$. Thus, *the price of adaptivity in our setting is that the signal from covariates associated with an action a must stand out with respect to the total covariate noise by time n* . The growth condition on U_n presented in Definition 3.5 is motivated as follows: w.h.p $\mathbb{E} \|\mathcal{E}_n\|_{op} \approx \sqrt{d} + \sqrt{n}$, and the extra additive $\sqrt{\log\left(\frac{1}{\delta} \log(n)\right)}$ factor is the price we pay for having high probability control of $\|\mathcal{E}_n\|_{op}$ uniformly over rounds. Below we provide an exact definition for U_n , as this choice leads to bounds with known constants and simple conditions on $\text{snr}_n(a)$ for validity.

We consider the following two sequences $(U_n)_{n \geq 1}$ in defining signal to noise ratio, which both satisfy $\|\mathcal{E}_n\|_{op} \leq U_n, \forall n \geq 1$ with probability at least $1 - \delta$, per Lemma B.3.

$$U_n^2 := \begin{cases} \beta \left(3\sqrt{n\ell_\delta/2\mathcal{N}(n)} + 5\ell_\delta/2\mathcal{N}(n) \right) + n\gamma & \text{when Assumption 3.1 holds} \\ \frac{3}{2}\sqrt{nCd\gamma\ell_\delta(n)} + \frac{7}{3}Cd\ell_\delta(n) + n\gamma & \text{when Assumption 3.2 holds.} \end{cases}$$

In the above, $\delta \in (0, 1)$, $\ell_\delta(n) := 2 \log \log(2n) + \log\left(\frac{d\pi^2}{12\delta}\right)$, $\beta = 32\sigma^2 e^2$, and $\mathcal{N} = 17^d$ is an upper bound on the $1/8$ -covering number of \mathbb{S}^{d-1} . While the exact forms of the above sequences $(U_n)_{n \geq 1}$ may appear complicated, it is helpful to realize that, under either Assumption 3.1 or 3.2, we have $U_n = O\left(\sqrt{n} + \sqrt{d} + \sqrt{\log\left(\frac{1}{\delta} \log(n)\right)}\right)$, i.e., the growth condition on U_n presented in Definition 3.5 is satisfied.

We can likewise define the *empirical signal to noise ratio associated with action a* as $\widehat{\text{snr}}_n(a) := \frac{\sigma_r(\mathbf{Z}_n(a))}{U_n}$. Note that unlike the (true) signal to noise ratio $\text{snr}_n(a)$, the empirical version $\widehat{\text{snr}}_n(a)$ is *computable* by the learner. Thus, it will be integral in stating our empirical-style bounds in the section that follows.

We conclude this section by comparing our notion of signal to noise ratio to that of Agarwal et al. [7], who define $\text{snr}_n(a)$ instead as $\frac{\sigma_r(\mathbf{X}_n(a))}{\sqrt{n} + \sqrt{d}}$, i.e. the ratio of the “signal” in the covariates to the *expected* operator norm of covariate noise $\mathbb{E}\|\mathcal{E}_n\|_{op}$. Since the goal of our work is high-probability (not in-expectation) estimation guarantees for PCR, we believe using high probability bounds on $\|\mathcal{E}_n\|_{op}$ is more natural when defining $\text{snr}_n(a)$.

4 Adaptive bounds for online (regularized) PCR

We now present the main results of this work—high-probability, time- and action-uniform bounds measuring the convergence of the PCR estimates $\hat{\theta}_n(a)$ to the true slope vectors $\theta(a)$. Unlike existing results [6–8], our bounds are valid when the covariates $(X_n)_{n \geq 1}$ and actions $(a_n)_{n \geq 1}$ are determined in an online (potentially adversarial) manner.

We first point out why the analysis of Agarwal et al. [7, 8] breaks down in the setting of adaptive (or online) PCR. First, many of the concentration inequalities leveraged in Agarwal et al. [7] do not hold in the adaptive design setting. As a particular example, the authors leverage the Hanson-Wright inequality [82, 73] for quadratic forms to study how the noisy covariate matrix \mathbf{Z}_n concentrates around the true matrix \mathbf{X}_n . This inequality fails to hold when the design points $(X_n)_{n \geq 1}$ depend on the previous observations. Second, the techniques leveraged by Agarwal et al. [8] to extend the convergence guarantees of PCR to the multiple action setting fail to hold when the n -th action a_n is selected based on previous observations. Lastly, the bounds presented in [7] are inherently fixed-time in nature—a simple way to convert existing fixed-time bounds to time-uniform ones would be to perform a union bound over time steps, but that introduces looseness in the bounds.

We are able to construct our bounds by exploiting connections between online PCR and self-normalized martingale concentration [50, 51, 32, 33]. In particular, we combine martingale-based results for constructing confidence ellipsoids for online regression [3, 32, 33] with methods for high-dimensional covariance estimation [83, 80] to prove our results. Exploiting this connection is what allows us to extend the results of Agarwal et al. [7] to the adaptive design, time-uniform setting. We begin with a bound which, up to constants and polylogarithmic factors, captures the rate of convergence of online PCR in terms of (a) the underlying signal to noise ratio and (b) the conditioning of the observed data.

Theorem 4.1 (Rate of Convergence for Online PCR). *Let $\delta \in (0, 1)$ be an arbitrary confidence parameter. Let $\rho > 0$ be chosen to be sufficiently small, as detailed in Appendix F. Further, assume that there is some $n_0 \geq 1$ such that $\text{rank}(\mathbf{X}_{n_0}(a)) = r$ and $\text{snr}_n(a) \geq 2$ for all $n \geq n_0$. Then, with probability at least $1 - O(A\delta)$, simultaneously for all actions $a \in [A]$ and time steps $n \geq n_0$, we have*

$$\|\hat{\theta}_n(a) - \theta(a)\|_2^2 = \tilde{O}\left(\frac{1}{\text{snr}_n(a)^2} \kappa(\mathbf{X}_n(a))^2\right),$$

where $\kappa(\mathbf{X}_n(a)) := \frac{\sigma_1(\mathbf{X}_n(a))}{\sigma_r(\mathbf{X}_n(a))}$ is the condition number (ignoring zero singular values) of $\mathbf{X}_n(a)$.

Theorem 4.1 is proved in Appendix E. We begin by comparing our bounds to those of Agarwal et al. [7, 8]. At any fixed time, our bounds take on roughly the same form as those of the aforementioned authors, having an inverse quadratic dependence on the signal to noise ratio. To make their bounds non-vacuous, the authors need to make the “soft sparsity” assumption of $\|\theta(a)\|_1 = O(\sqrt{d})$. Our bound, on the other hand, suffers no dependence on the ℓ_1 -norm of the $\theta(a)$ ’s. This makes intuitive sense, as the specific choice of a basis should not impact the rate of convergence of PCR. However, our bounds pay a price for adaptivity—in particular, the signal to noise ratio associated with an action is defined with respect to a bound on the *total* operator norm of the matrix \mathcal{E}_n . If an action is selected very infrequently, the above bound may become loose.

While the above bound is stated in terms of signal to noise ratio, if we make additional assumptions, we can obtain bounds directly in terms of d , n , and r . In particular, the following “well-balancing” assumptions suffice.

Assumption 4.2 (Well-balancing assumptions). *For all $n \geq n_0$, the following hold: (a) $\sigma_i(\mathbf{X}_n(a)) = \Theta\left(\sqrt{\frac{c_n(a)d}{r}}\right)$ for all $i \in [r]$, (b) $c_n(a) = \Theta(c_n(a'))$ for all $a, a' \in [A]$, and (c) $A = O(r)$.*

Corollary 4.3. *Assume the same setup as Theorem 4.1, and further assume Assumption 4.2 holds. Then with probability at least $1 - O(A\delta)$, simultaneously for all actions $a \in [A]$ and time steps $n \geq n_0$, we have*

$$\|\hat{\theta}_n(a) - \theta(a)\|_2^2 = \tilde{O}\left(\frac{r^2}{d \wedge n}\right).$$

Corollary 4.3 shows that Theorem 4.1 obtains the same estimation rate as Theorem 4.1 of Agarwal et al. [7] if assumption Assumption 4.2 holds. This “well-balancing” assumption says roughly that all non-zero singular values of \mathbf{X}_n are of the same order, each action is selected with the same frequency, and that the number of actions is, at most, proportional to dimension of the true, unknown subspace. As noted by Agarwal et al. [7], the assumption of a “well-balanced spectrum” (for \mathbf{X}_n) is common in many works in econometrics and robust statistics, and additionally holds with high probability if the entries of \mathbf{X}_n are i.i.d. [62, 21, 37]. Further, it is often the case that there only few available actions (for instance, in the synthetic control literature there are only two actions [2, 1, 38]), justifying the assumption that $A = O(r)$. Lastly, ensuring that each action is played (very roughly) the same number of times can be viewed as a price for adaptivity.

The proof of Theorem 4.1 is immediate as a corollary from the following, more complicated bound. Theorem 4.4 below measures the convergence of $\hat{\theta}_n(a)$ to $\theta(a)$ in terms of empirical (i.e. observed) quantities. We imagine this bound to be the most practically relevant of our results, as, unlike the results of Agarwal et al. [7], it is directly computable by the learner, involves known constants, and places minimal conditions on the signal to noise ratio.

Theorem 4.4 (Empirical Guarantees for Online PCR). *Let $\delta \in (0, 1)$ be an arbitrary confidence parameter. Let $\rho > 0$ be chosen to be sufficiently small, as detailed in Appendix F. Further, assume that there is some $n_0 \geq 1$ such that $\text{rank}(\mathbf{X}_{n_0}(a)) = r$ and $\text{snr}_n(a) \geq 2$ for all $n \geq n_0$. Then, with probability at least $1 - O(A\delta)$, simultaneously for all actions $a \in [A]$ and time steps $n \geq n_0$, we have*

$$\left\| \hat{\theta}_n(a) - \theta(a) \right\|_2^2 \leq \frac{L^2}{\widehat{\text{snr}}_n(a)^2} [74 + 216\kappa(\mathbf{Z}_n(a))^2] + \frac{2\text{err}_n(a)}{\sigma_r(\mathbf{Z}_n(a))^2},$$

where $\kappa(\mathbf{Z}_n(a)) := \frac{\sigma_1(\mathbf{Z}_n(a))}{\sigma_r(\mathbf{Z}_n(a))}$, $\|\theta(a)\|_2 \leq L$, and in the above we define the “error” term $\text{err}_n(a)$ to be

$$\begin{aligned} \text{err}_n(a) &:= 32\rho L^2 + 64\eta^2 \left(\log\left(\frac{A}{\delta}\right) + r \log\left(1 + \frac{\sigma_1(\mathbf{Z}_n(a))^2}{\rho}\right) \right) \\ &\quad + 6\eta^2 \sqrt{2c_n(a)\ell_\delta(c_n(a))} + 10\eta^2 \ell_\delta(c_n(a)) + 6c_n(a)\alpha. \end{aligned}$$

We see that the above bound, with the exception of the third term, more or less resembles the bound presented in Theorem 4.1, just written in terms of the observed covariates $\mathbf{Z}_n(a)$ instead of the true covariates $\mathbf{X}_n(a)$. We view the third term as a slowly growing “error” term. In particular, all terms in the quantity $\text{err}_n(a)$ are either constant, logarithmic in the singular values of $\mathbf{Z}_n(a)$, or linear in $c_n(a)$, the number of times by round n action a has been selected. This implies that $\text{err}_n(a) = \tilde{O}(n + d)$, ensuring $\text{err}_n(a)$ is dominated by other terms in the asymptotic analysis. We now provide the proof of Theorem 4.4. The key application of self-normalized, martingale concentration comes into play in bounding the quantities that appear in the upper bounds of terms T_1 and T_2 (to be defined below).

Proof. Observe the decomposition, for any $n \geq 1$ and $a \in [A]$

$$\hat{\theta}_n(a) - \theta(a) = \hat{\mathbf{P}}_n(a)(\hat{\theta}_n(a) - \theta(a)) + (\mathbf{P}^\perp - \hat{\mathbf{P}}_n^\perp(a))\theta(a),$$

where \mathbf{P}^\perp is the projection onto the subspace orthogonal to W^* and $\hat{\mathbf{P}}_n^\perp(a)$ is the projection onto the subspace orthogonal to the learned subspace (i.e. that spanned by $\mathbf{Z}_{n,r}(a)$). Since $\hat{\mathbf{P}}_n(a)(\hat{\theta}_n(a) - \theta(a))$ and $(\mathbf{P}^\perp - \hat{\mathbf{P}}_n^\perp(a))\theta(a)$ are orthogonal vectors, we have

$$\left\| \hat{\theta}_n(a) - \theta(a) \right\|_2^2 = \left\| \hat{\mathbf{P}}_n(a)(\hat{\theta}_n(a) - \theta(a)) \right\|_2^2 + \left\| (\hat{\mathbf{P}}_n^\perp(a) - \mathbf{P}^\perp)\theta(a) \right\|_2^2.$$

We bound these two terms separately, beginning with the second term. Going forward, fix an action $a \in [A]$. Observe that with probability at least $1 - \delta$, simultaneously for all $n \geq n_0(a)$,

$$\begin{aligned} \left\| (\hat{\mathbf{P}}_n^\perp(a) - \mathbf{P}^\perp)\theta(a) \right\|_2^2 &\leq \left\| \hat{\mathbf{P}}_n^\perp(a) - \mathbf{P}^\perp \right\|_{op}^2 \|\theta(a)\|_2^2 \\ &\leq L^2 \left\| \hat{\mathbf{P}}_n^\perp(a) - \mathbf{P}^\perp \right\|_{op}^2 = L^2 \left\| \hat{\mathbf{P}}_n(a) - \mathbf{P} \right\|_{op}^2 \\ &\leq \frac{4L^2 U_n^2}{\sigma_r(\mathbf{X}_n(a))^2} \leq \frac{6L^2 U_n^2}{\sigma_r(\mathbf{Z}_n(a))^2}, \end{aligned}$$

where the equality in the above comes from observing $\|\widehat{\mathbf{P}}_n^\perp(a) - \mathbf{P}^\perp\|_{op} = \|\widehat{\mathbf{P}}_n(a) - \mathbf{P}\|_{op}$, the second-to-last inequality comes from applying Lemma B.4, and the last inequality follows from the second part of Lemma B.6.

We now bound the first term. Observe that we can write

$$\begin{aligned} \left\| \widehat{\mathbf{P}}_n(a) \left(\widehat{\theta}_n(a) - \theta(a) \right) \right\|_2^2 &\leq \frac{1}{\sigma_r(\mathbf{Z}_n(a))^2} \left\| \widehat{\mathbf{Z}}_n(a) \left(\widehat{\theta}_n(a) - \theta(a) \right) \right\|_2^2 \\ &\leq \frac{2}{\sigma_r(\mathbf{Z}_n(a))^2} \left[\underbrace{\left\| \widehat{\mathbf{Z}}_n(a) \widehat{\theta}_n(a) - \mathbf{X}_n(a) \theta(a) \right\|_2^2}_{T_1} + \underbrace{\left\| \mathbf{X}_n(a) \theta(a) - \widehat{\mathbf{Z}}_n(a) \theta(a) \right\|_2^2}_{T_2} \right], \end{aligned} \quad (1)$$

where the first inequality follows from the fact that $\widehat{\mathbf{P}}_n(a) \preceq \frac{1}{\sigma_r(\widehat{\mathbf{Z}}_n(a))^2} \widehat{\mathbf{Z}}_n(a)^\top \widehat{\mathbf{Z}}_n(a)$ and $\sigma_r(\mathbf{Z}_n(a)) = \sigma_r(\widehat{\mathbf{Z}}_n(a))$, and the second inequality comes from applying the parallelogram inequality. First we bound T_1 . We have, with probability at least $1 - O(\delta)$, simultaneously for all $n \geq n_0(a)$

$$\begin{aligned} T_1 &\leq 8 \left\| \check{\mathbf{Y}}_n(a)^{1/2} \left(\check{\theta}_n(a) - \theta(a) \right) \right\|_2^2 + 6 \|\Xi_n(a)\|_2^2 + 8 \left\| \widehat{\mathbf{Z}}_n(a) \theta(a) - \mathbf{X}_n(a) \theta(a) \right\|_2^2 \\ &\leq 32\rho L^2 + 64\eta^2 \left(\log \left(\frac{A}{\delta} \right) + r \log \left(1 + \frac{\sigma_1(\mathbf{Z}_n(a))^2}{\rho} \right) \right) + 16L^2 U_n^2 \\ &\quad + 6\eta^2 \sqrt{2c_n(a)\ell_\delta(c_n(a))} + 10\eta^2 \ell_\delta(c_n(a)) + 6c_n(a)\alpha + 8T_2, \end{aligned} \quad (2)$$

where the first inequality follows from Lemma D.1, and the second inequality follows from applying Lemmas C.1 and D.3. $\ell_\delta(n) = 2 \log \log(2n) + \log \left(\frac{d\pi^2}{12\delta} \right)$, as defined in Lemma A.2. We now bound T_2 . With probability at least $1 - O(\delta)$ simultaneously for all $n \geq n_0$, we have

$$\begin{aligned} T_2 &\leq 2L^2 \sigma_1(\mathbf{Z}_n(a))^2 \left\| \mathbf{P} - \widehat{\mathbf{P}}_n(a) \right\|_{op}^2 + 2L^2 \|\mathcal{E}_n\|_{op}^2 \\ &\leq \frac{8L^2 \sigma_1(\mathbf{Z}_n(a))^2 U_n^2}{\sigma_r(\mathbf{X}_n(a))^2} + 2L^2 U_n^2 \\ &\leq \frac{12L^2 \sigma_1(\mathbf{Z}_n(a))^2 U_n^2}{\sigma_r(\mathbf{Z}_n(a))^2} + 2L^2 U_n^2. \end{aligned} \quad (3)$$

The first inequality follows from Lemma D.2, the second inequality follows from applying Lemmas B.4 and B.3, and the final inequality follows from Lemma B.6.

Piecing the above inequalities together yields the desired result, which can be checked via the argument at the end of Appendix D. A union bound over actions then yields that the desired inequality holds over all actions $a \in [A]$ with probability at least $1 - O(A\delta)$. \square

5 Application to causal inference with panel data

We now apply our bounds for adaptive PCR to online experiment design in the context of panel data. In this setting, the learner is interested in estimating *unit-specific counterfactuals* under different *interventions*, given a sequence of unit *outcomes* (or *measurements*) over *time*. Units can range from medical patients, to subpopulations or geographic locations. Examples of interventions include medical treatments, discounts, and socioeconomic policies. *Synthetic control (SC)* is a popular framework used to estimate counterfactual unit outcomes in panel data settings, had they not been treated (i.e. remained under *control*) [1, 2]. In SC, there is a notion of a *pre-intervention* time period in which all units are under control, followed by a *post-intervention* time period, in which every unit undergoes one of several interventions (including control). At a high level, SC fits a model of a unit's pre-treatment outcomes using pre-treatment data from units who remained under control in the post-intervention time period. It then constructs a "synthetic control" by using the learned model to predict the unit's post-intervention outcomes, had they remained under control. *Synthetic interventions (SI)* is a recent generalization of the SC framework, which allows for counterfactual estimation of unit outcomes under different interventions, in addition to control [8]. Using our

bounds from Section 4, we show how to generalize the SI framework of Agarwal et al. [8] to settings where interventions are assigned via an *adaptive intervention assignment policy*.

As a motivating example, consider an online e-commerce platform (learner) which assigns discounts (interventions) to users (units) with the goal of maximizing total user engagement on the platform. For concreteness, suppose that the e-commerce platform assigns discounts *greedily* with respect to the discount level which appears to be best at the current round (i.e. maximizes total engagement for the current user), given the sequence of previously observed (user, discount level, engagement level) tuples. Under such a setting, the intervention assigned at the current round n will be correlated with the observed engagement levels at previous rounds, thus breaking the requirement of the SI framework [8] that the intervention assignment is not adaptive to previously observed outcomes.

Formally, we consider a panel data setting in which the principal observes units over a sequence of rounds. In each round n , the learner observes a unit n under *control* for $T_0 \in \mathbb{N}$ time steps, followed by one of A *interventions* (including control, which we denote by 0) for the remaining $T - T_0$ time steps, where $T \in \mathbb{N}$. Overloading notation to be consistent with the literature on panel data, we denote the potential outcome of unit n at time t under intervention a by $Y_{n,t}^{(a)} \in \mathbb{R}$, the set of unit n 's pre-treatment outcomes (under control) by $Y_{n,pre} := [Y_{n,1}^{(0)}, \dots, Y_{n,T_0}^{(0)}]^T \in \mathbb{R}^{T_0}$, and their post-intervention potential outcomes under intervention a by $Y_{n,post}^{(a)} := [Y_{n,T_0+1}^{(a)}, \dots, Y_{n,T}^{(a)}]^T \in \mathbb{R}^{T-T_0}$. We use a to refer to an arbitrary intervention in $\{0, \dots, A-1\}$ and a_n to denote the *realized* intervention unit n actually receives in the post-intervention time period. We posit that potential outcomes are generated by the following *latent factor model* over units, time steps, and interventions.

Assumption 5.1 (Latent Factor Model). *Suppose the outcome for unit n at time step t under treatment $a \in \{0, \dots, A-1\}$ takes the form*

$$Y_{n,t}^{(a)} = \langle U_t^{(a)}, V_n \rangle + \epsilon_{n,t}^{(a)},$$

where $U_t^{(a)} \in \mathbb{R}^r$ is a latent factor which depends only on the time step t and intervention a , $V_n \in \mathbb{R}^r$ is a latent factor which only depends on unit n , and $\epsilon_{n,t}^{(a)}$ is zero-mean SubGaussian random noise with variance at most σ^2 . We assume, without loss of generality, that $|\langle U_t^{(a)}, V_n \rangle| \leq 1$ for all $n \geq 1$, $t \in [T]$, $a \in \{0, \dots, A-1\}$.

Note that the learner observes $Y_{n,t}^{(a)}$ for only the intervention a_n that unit n is under at time step t , and never observes $U_t^{(a)}$, V_n , or $\epsilon_{n,t}^{(a)}$. Such ‘‘low rank’’ assumptions are ubiquitous within the panel data literature (see references in Section 2). We assume that r is known to the learner, although principled heuristics exist for estimating r in practice from data (see, e.g. Section 3.2 of Agarwal et al. [8]). Additionally, we make the following ‘‘causal transportability’’ assumption on the latent factors.

Assumption 5.2 (Linear span inclusion). *For any post-intervention time step $t \in [T_0 + 1, T]$ and intervention $a \in \{0, \dots, A-1\}$, we assume that $U_t^{(a)} \in \text{span}(\{U_t^{(0)} : t \in [T_0]\})$.*

Intuitively, Assumption 5.2 allows for information to be inferred about the potential outcomes in the post-intervention time period using pre-treatment observations. The goal of the learner is to estimate unit-specific counterfactual outcomes under different interventions *when the sequence of units and interventions is chosen adaptively*. In line with previous work in SI and SC, our target causal parameter is the (counterfactual) *average expected post-intervention outcome*.

Definition 5.3. *The average expected post-intervention outcome of unit n under intervention a is*

$$\mathbb{E}\bar{Y}_{n,post}^{(a)} := \frac{1}{T - T_0} \sum_{t=T_0+1}^T \mathbb{E}Y_{n,t}^{(a)},$$

where the expectation is taken with respect to $(\epsilon_{n,t}^{(a)})_{T_0 < t \leq T}$.

While we consider the *average* post-intervention outcome, our results may be readily extended to settings in which the target causal parameter is any *linear* combination of post-intervention outcomes. Next we show that under Assumption 5.1 and Assumption 5.2, $\mathbb{E}\bar{Y}_{n,post}^{(a)}$ may be written as a linear combination of unit n 's *pre-intervention* outcomes. We note that similar observations have previously been made in the panel data literature (e.g. [46]), but we include the following lemma for completeness' sake.

Lemma 5.4 (Reformulation of average expected post-intervention outcome). *Under Assumption 5.1 and Assumption 5.2, there exists slope vector $\theta(a) \in \mathbb{R}^{T_0}$, such that the average expected post-intervention outcome of unit n under intervention a is expressible as*

$$\mathbb{E}\bar{Y}_{n,\text{post}}^{(a)} = \frac{1}{T - T_0} \langle \theta(a), \mathbb{E}Y_{n,\text{pre}} \rangle.$$

$\theta(a)$ may be interpreted as a unit-independent measure of the causal relationship between pre- and post-intervention outcomes. Using this reformulation, adaptive guarantees for the estimation of causal effects over time may now be obtained by applying our online PCR results of Section 4. Overloading the notation of Section 3, we let $\mathbf{X}_n = (\mathbb{E}Y_{1,\text{pre}}, \dots, \mathbb{E}Y_{n,\text{pre}})^T$, $\mathbf{Z}_n = (Y_{1,\text{pre}}, \dots, Y_{n,\text{pre}})^T$, $\epsilon_{n,\text{pre}} = (\epsilon_{n,1}^{(0)}, \dots, \epsilon_{n,T_0}^{(0)})^T$, $\mathcal{E}_n = (\epsilon_{1,\text{pre}}, \dots, \epsilon_{n,\text{pre}})^T$, $\xi_n = \sum_{t=T_0+1}^T \epsilon_{n,t}^{(a_n)}$, $\Xi_n = (\xi_1, \dots, \xi_n)^T$, and

$$\mathbf{Y}_n = \left(\frac{1}{T - T_0} \sum_{t=1}^{T_0} Y_{1,t}^{(a_1)}, \dots, \frac{1}{T - T_0} \sum_{t=1}^{T_0} Y_{n,t}^{(a_n)} \right)^T.$$

Finally, we define quantities such as $\mathbf{Z}_n(a)$, $\mathbf{X}_n(a)$, $\mathbf{Y}_n(a)$ analogously to Section 3. We now turn to bounding our primary quantity of interest in the panel data setting: prediction error for the average expected post-intervention outcome.

Theorem 5.5 (Prediction error of average expected post-intervention outcome). *Let $\delta \in (0, 1)$ be an arbitrary confidence parameter and $\rho > 0$ be chosen to be sufficiently small, as detailed in Appendix F. Further, assume that Assumptions 5.1 and 5.2 are satisfied, there is some $n_0 \geq 1$ such that $\text{rank}(\mathbf{X}_{n_0}(a)) = r$, and $\text{sur}_n(a) \geq 2$ for all $n \geq n_0$. If $T_0 \leq \frac{1}{2}T$ and $r \leq \sqrt{T_0 \wedge n}$, then under Assumption 4.2 with probability at least $1 - O(A\delta)$, simultaneously for all interventions $a \in \{0, \dots, A - 1\}$*

$$|\widehat{\mathbb{E}}\bar{Y}_{n,\text{post}}^{(a)} - \mathbb{E}\bar{Y}_{n,\text{post}}^{(a)}| = \tilde{O} \left(\frac{L}{\sqrt{T - T_0}} + \frac{r}{\sqrt{T_0 \wedge n}} + \frac{r(L \vee 1)}{\sqrt{(T - T_0)(T_0 \wedge n)}} \right)$$

where $\widehat{\mathbb{E}}\bar{Y}_{n,\text{post}}^{(a)} := \frac{1}{T - T_0} \cdot \langle \widehat{\theta}_n(a), Y_{n,\text{pre}} \rangle$ is the estimated average post-intervention outcome for unit n under intervention a .

A more complicated expression which does not require Assumption 4.2 or $T_0 \leq \frac{1}{2}T$ may be found in Appendix G. Observe that $|\widehat{\mathbb{E}}\bar{Y}_{n,\text{post}}^{(a)} - \mathbb{E}\bar{Y}_{n,\text{post}}^{(a)}| \rightarrow 0$ with high probability as $T, T_0, n \rightarrow \infty$.

We conclude this section by comparing our results with those of the (non-adaptive) synthetic interventions framework of Agarwal et al. [8]. Since we are regressing over *time*, our method for estimating $\widehat{\mathbb{E}}\bar{Y}_{n,\text{post}}^{(a)}$ is known as a *horizontal* regression method in the panel data literature. This is in contrast to *vertical* regression methods, which regress over *units*. See Shen et al. [75] for more details on the similarities and differences between horizontal and vertical regression methods in panel data settings. While we do not exactly match the bound of Agarwal et al. [8] since their synthetic interventions framework of uses a vertical regression method, the two bounds are similar, with the notable differences being as follows: (1) The bound of [8] contains a “slow” $\tilde{O}(r^{1/2}T_0^{-1/4})$ term which does not appear in our analysis. (2) The non-adaptive SI bound also contains a term which scales as $\tilde{O}(\frac{\sqrt{n}}{\sqrt{T - T_0}})$ in the worst case, while our bound has no such dependence. However, this comes at the price of slow rate whenever L is large compared to $\sqrt{T - T_0}$.

6 Conclusion

We obtain the first adaptive bounds for principal component regression and apply them to the problem of online experiment design in the context of panel data, where we allow for interventions to be assigned according to an adaptive policy. Exciting directions for future work include applications of our results to domains such as differential privacy, and using our bounds to obtain contextual bandit algorithms (e.g. based on LinUCB [3]) capable of regret minimization when given access to noisy contexts.

Acknowledgements

KH is supported in part by an NDSEG Fellowship. KH, JW, and ZSW are supported in part by NSF FAI Award #1939606. The authors would like to thank the anonymous NeurIPS reviewers for valuable feedback.

References

- [1] Alberto Abadie and Javier Gardeazabal. The economic costs of conflict: A case study of the basque country. *American economic review*, 93(1):113–132, 2003.
- [2] Alberto Abadie, Alexis Diamond, and Jens Hainmueller. Synthetic control methods for comparative case studies: Estimating the effect of california’s tobacco control program. *Journal of the American statistical Association*, 105(490):493–505, 2010.
- [3] Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. *Advances in neural information processing systems*, 24, 2011.
- [4] Abhineet Agarwal, Anish Agarwal, and Suhas Vijaykumar. Synthetic combinations: A causal inference framework for combinatorial interventions. *arXiv preprint arXiv:2303.14226*, 2023.
- [5] Anish Agarwal and Rahul Singh. Causal inference with corrupted data: Measurement error, missing values, discretization, and differential privacy. *arXiv preprint arXiv:2107.02780*, 2021.
- [6] Anish Agarwal, Devavrat Shah, Dennis Shen, and Dogyoon Song. On robustness of principal component regression. *Advances in Neural Information Processing Systems*, 32, 2019.
- [7] Anish Agarwal, Devavrat Shah, and Dennis Shen. On principal component regression in a high-dimensional error-in-variables setting. *arXiv preprint arXiv:2010.14449*, 2020.
- [8] Anish Agarwal, Devavrat Shah, and Dennis Shen. Synthetic interventions. *arXiv preprint arXiv:2006.07691*, 2020.
- [9] Anish Agarwal, Munther Dahleh, Devavrat Shah, and Dennis Shen. Causal matrix completion. *arXiv preprint arXiv:2109.15154*, 2021.
- [10] Anish Agarwal, Devavrat Shah, Dennis Shen, and Dogyoon Song. On robustness of principal component regression. *Journal of the American Statistical Association*, 116(536):1731–1745, 2021. doi: 10.1080/01621459.2021.1928513.
- [11] Anish Agarwal, Sarah Cen, Devavrat Shah, and Christina Lee Yu. Network synthetic interventions: A framework for panel data with network interference. *arXiv preprint arXiv:2210.11355*, 2022.
- [12] Deepak Agarwal and Bee-Chung Chen. Regression-based latent factor models. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 19–28, 2009.
- [13] Muhammad Jehangir Amjad, Devavrat Shah, and Dennis Shen. Robust synthetic control. *Journal of Machine Learning Research*, 19:1–51, 2018.
- [14] Muhummad Amjad, Vishal Mishra, Devavrat Shah, and Dennis Shen. mrsc: Multi-dimensional robust synthetic control. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 3(2), 2019.
- [15] Manuel Arellano and Bo Honore. Panel data models: Some recent developments. *Handbook of Econometrics*, 02 2000.
- [16] Dmitry Arkhangelsky, Susan Athey, David A. Hirshberg, Guido W. Imbens, and Stefan Wager. Synthetic difference in differences, 2020.
- [17] Susan Athey, Mohsen Bayati, Nikolay Doudchenko, Guido Imbens, and Khashayar Khosravi. Matrix completion methods for causal panel data models. *Journal of the American Statistical Association*, 116(536):1716–1730, 2021.
- [18] Jushan Bai. Inferential theory for factor models of large dimensions. *Econometrica*, 71(1):135–171, 2003. ISSN 00129682, 14680262. URL <http://www.jstor.org/stable/3082043>.
- [19] Jushan Bai. Panel data models with interactive fixed effects. *Econometrica*, 77(4):1229–1279, 2009. ISSN 00129682, 14680262. URL <http://www.jstor.org/stable/40263859>.

- [20] Jushan Bai and Serena Ng. Matrix completion, counterfactuals, and factor analysis of missing data, 2020.
- [21] Jushan Bai and Serena Ng. Matrix completion, counterfactuals, and factor analysis of missing data. *Journal of the American Statistical Association*, 116(536):1746–1763, 2021.
- [22] Eric Bair, Trevor Hastie, Debashis Paul, and Robert Tibshirani. Prediction by supervised principal components. *Journal of the American Statistical Association*, 101(473):119–137, 2006.
- [23] Eli Ben-Michael, Avi Feller, and Jesse Rothstein. The augmented synthetic control method, 2020.
- [24] Aldo Gael Carranza, Sanath Kumar Krishnamurthy, and Susan Athey. Flexible and efficient contextual bandits with heterogeneous treatment effect oracles. In *International Conference on Artificial Intelligence and Statistics*, pages 7190–7212. PMLR, 2023.
- [25] Gary Chamberlain. Panel data. In Z. Griliches† and M. D. Intriligator, editors, *Handbook of Econometrics*, volume 2, chapter 22, pages 1247–1318. Elsevier, 1 edition, 1984. URL <https://EconPapers.repec.org/RePEc:eee:ecochp:2-22>.
- [26] Jiafeng Chen. Synthetic control as online linear regression. *Econometrica*, 91(2):465–491, 2023.
- [27] Victor Chernozhukov, Kaspar Wuthrich, and Yinchu Zhu. Practical and robust t -test based inference for synthetic control and related methods, 2020.
- [28] Andrew Chesher. The effect of measurement error. *Biometrika*, 78(3):451–462, 1991.
- [29] Sayak Ray Chowdhury and Aditya Gopalan. On kernelized multi-armed bandits. In *International Conference on Machine Learning*, pages 844–853. PMLR, 2017.
- [30] Ben Chugg, Hongjian Wang, and Aaditya Ramdas. A unified recipe for deriving (time-uniform) pac-bayes bounds. *arXiv preprint arXiv:2302.03421*, 2023.
- [31] Abhirup Datta and Hui Zou. Cocolasso for high-dimensional error-in-variables regression. 2017.
- [32] Victor H de la Peña, Michael J Klass, and Tze Leung Lai. Self-normalized processes: exponential inequalities, moment bounds and iterated logarithm laws. 2004.
- [33] Victor H de la Peña, Michael J Klass, and Tze Leung Lai. Pseudo-maximization and self-normalized processes. 2007.
- [34] Yash Deshpande, Lester Mackey, Vasilis Syrgkanis, and Matt Taddy. Accurate inference for adaptive linear models. In *International Conference on Machine Learning*, pages 1194–1203. PMLR, 2018.
- [35] N. Doudchenko and G. Imbens. Balancing, regression, difference-in-differences and synthetic control methods: A synthesis. *NBER Working Paper No. 22791*, 2016.
- [36] Walter Enders. *Applied econometric time series*. John Wiley & Sons, 2008.
- [37] Jianqing Fan, Weichen Wang, and Yiqiao Zhong. An ℓ_∞ eigenvector perturbation bound and its application to robust covariance estimation. *Journal of Machine Learning Research*, 18(207): 1–42, 2018.
- [38] Vivek Farias, Ciamac Moallemi, Tianyi Peng, and Andrew Zheng. Synthetically controlled bandits. *arXiv preprint arXiv:2202.07079*, 2022.
- [39] Iván Fernández-Val, Hugo Freeman, and Martin Weidner. Low-rank approximations of nonseparable panel models, 2020.
- [40] Jorge I Figueroa-Zúñiga, Cristian L Bayes, Víctor Leiva, and Shuangzhe Liu. Robust beta regression modeling with errors-in-variables: a bayesian approach and numerical applications. *Statistical Papers*, pages 1–24, 2022.

- [41] Wayne A Fuller. *Measurement error models*. John Wiley & Sons, 2009.
- [42] Zvi Griliches and Vidar Ringstad. Error-in-the-variables bias in nonlinear contexts. *Econometrica: Journal of the Econometric Society*, pages 368–370, 1970.
- [43] Vitor Hadad, David A Hirshberg, Ruohan Zhan, Stefan Wager, and Susan Athey. Confidence intervals for policy evaluation in adaptive experiments. *Proceedings of the national academy of sciences*, 118(15):e2014602118, 2021.
- [44] Daniel B Hall. Measurement error in nonlinear models: a modern perspective, 2008.
- [45] Keegan Harris, Anish Agarwal, Chara Podimata, and Zhiwei Steven Wu. Strategyproof decision-making in panel data settings and beyond. *arXiv preprint arXiv:2211.14236*, 2022.
- [46] Keegan Harris, Dung Daniel T Ngo, Logan Stapleton, Hoda Heidari, and Steven Wu. Strategic instrumental variable regression: Recovering causal relationships from strategic responses. In *International Conference on Machine Learning*, pages 8502–8522. PMLR, 2022.
- [47] Jerry Hausman. Mismeasured variables in econometric analysis: problems from the right and problems from the left. *Journal of Economic perspectives*, 15(4):57–67, 2001.
- [48] Peter D Hoff. Multiplicative latent factor models for description and prediction of social networks. *Computational and mathematical organization theory*, 15(4):261, 2009.
- [49] Jean Honorio and Tommi Jaakkola. Tight bounds for the expected risk of linear classifiers and pac-bayes finite-sample guarantees. In *Artificial Intelligence and Statistics*, pages 384–392. PMLR, 2014.
- [50] Steven R Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. Time-uniform chernoff bounds via nonnegative supermartingales. 2020.
- [51] Steven R Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. Time-uniform, nonparametric, nonasymptotic confidence sequences. 2021.
- [52] Cheng Hsiao, H. Steve Ching, and Shui Ki Wan. A panel data approach for program evaluation: Measuring the benefits of political and economic integration of hong kong with mainland china. *Journal of Applied Econometrics*, 27(5):705–740, 2012. doi: <https://doi.org/10.1002/jae.1230>.
- [53] Rodolphe Jenatton, Nicolas Roux, Antoine Bordes, and Guillaume R Obozinski. A latent factor model for highly multi-relational data. *Advances in neural information processing systems*, 25, 2012.
- [54] Ian T Jolliffe. A note on the use of principal components in regression. *Journal of the Royal Statistical Society Series C: Applied Statistics*, 31(3):300–303, 1982.
- [55] Adam Kalai and Santosh Vempala. Efficient algorithms for online decision problems. *Journal of Computer and System Sciences*, 71(3):291–307, 2005.
- [56] Abhishek Kaul and Hira L Koul. Weighted l1-penalized corrected quantile regression for high dimensional measurement error models. *Journal of Multivariate Analysis*, 140:72–91, 2015.
- [57] In-Won Kim, Michael J Liebman, and Thomas F Edgar. Robust error-in-variables estimation using nonlinear programming techniques. *AIChE Journal*, 36(7):985–993, 1990.
- [58] Tor Lattimore and Csaba Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.
- [59] Kathleen T. Li. Inference for factor model based average treatment effects. *Available at SSRN 3112775*, 2018.
- [60] Kathleen T. Li and David R. Bell. Estimation of average treatment effects with panel data: Asymptotic theory and implementation. *Journal of Econometrics*, 197(1):65 – 75, 2017. ISSN 0304-4076. doi: <https://doi.org/10.1016/j.jeconom.2016.01.011>.
- [61] Kung-Yee Liang and Scott L. Zeger. Longitudinal data analysis using generalized linear models. *Biometrika*, 73(1):13–22, 04 1986. ISSN 0006-3444. doi: 10.1093/biomet/73.1.13. URL <https://doi.org/10.1093/biomet/73.1.13>.

- [62] Po-Ling Loh and Martin J Wainwright. High-dimensional regression with noisy and missing data: Provable guarantees with non-convexity. *Advances in neural information processing systems*, 24, 2011.
- [63] Tudor Manole and Aaditya Ramdas. Martingale methods for sequential estimation of convex functionals and divergences. *IEEE Transactions on Information Theory*, 2023.
- [64] Hyungsik Roger Moon and Martin Weidner. Linear regression for panel with unknown number of factors as interactive fixed effects. *Econometrica*, 83(4):1543–1579, 2015. ISSN 00129682, 14680262. URL <http://www.jstor.org/stable/43616977>.
- [65] Hyungsik Roger Moon and Martin Weidner. Dynamic linear panel regression models with interactive fixed effects. *Econometric Theory*, 33(1):158–195, 2017. doi: 10.1017/S0266466615000328.
- [66] Daniel Ngo, Keegan Harris, Anish Agarwal, Vasilis Syrgkanis, and Zhiwei Steven Wu. Incentive-aware synthetic control: Accurate counterfactual estimation via incentivized exploration. 2023.
- [67] Xinkun Nie, Xiaoying Tian, Jonathan Taylor, and James Zou. Why adaptively collected data have negative bias and how to correct for it. In *International Conference on Artificial Intelligence and Statistics*, pages 1261–1269. PMLR, 2018.
- [68] M. Hashem Pesaran. Estimation and inference in large heterogeneous panels with a multifactor error structure. *Econometrica*, 74(4):967–1012, 2006. ISSN 00129682, 14680262. URL <http://www.jstor.org/stable/3805914>.
- [69] Chao Qin and Daniel Russo. Adaptivity and confounding in multi-armed bandit experiments. *arXiv preprint arXiv:2202.09036*, 2022.
- [70] Park M Reilly and Hugo Patino-Lea1. A bayesian study of the error-in-variables model. *Technometrics*, 23(3):221–231, 1981.
- [71] Steven Roman, S Axler, and FW Gehring. *Advanced linear algebra*, volume 3. Springer, 2005.
- [72] Mathieu Rosenbaum and Alexandre B Tsybakov. Sparse recovery under matrix uncertainty. *The Annals of Statistics*, pages 2620–2651, 2010.
- [73] Mark Rudelson and Roman Vershynin. Hanson-wright inequality and sub-gaussian concentration. 2013.
- [74] Burr Settles. Active learning literature survey. 2009.
- [75] Dennis Shen, Peng Ding, Jasjeet Sekhon, and Bin Yu. A tale of two panel data regressions. *arXiv preprint arXiv:2207.14481*, 2022.
- [76] Robert H Shumway, David S Stoffer, and David S Stoffer. *Time series analysis and its applications*, volume 3. Springer, 2000.
- [77] David Simchi-Levi and Chonghuan Wang. Multi-armed bandit experimental design: Online decision-making and adaptive inference. In *International Conference on Artificial Intelligence and Statistics*, pages 3086–3097. PMLR, 2023.
- [78] Aleksandrs Slivkins et al. Introduction to multi-armed bandits. *Foundations and Trends® in Machine Learning*, 12(1-2):1–286, 2019.
- [79] Chandler Squires, Dennis Shen, Anish Agarwal, Devavrat Shah, and Caroline Uhler. Causal imputation via synthetic interventions. In *Conference on Causal Learning and Reasoning*, pages 688–711. PMLR, 2022.
- [80] Joel A Tropp et al. An introduction to matrix concentration inequalities. *Foundations and Trends® in Machine Learning*, 8(1-2):1–230, 2015.
- [81] Sridhar Ungarala and Bhavik R Bakshi. A multiscale, bayesian and error-in-variables approach for linear dynamic data rectification. *Computers & Chemical Engineering*, 24(2-7):445–451, 2000.

- [82] Roman Vershynin. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press, 2018.
- [83] Martin J Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge university press, 2019.
- [84] Tom Wansbeek and Erik Meijer. Measurement error and latent variables. *A companion to theoretical econometrics*, pages 162–179, 2001.
- [85] Justin Whitehouse, Aaditya Ramdas, Steven Z Wu, and Ryan M Rogers. Brownian noise reduction: Maximizing privacy subject to accuracy constraints. *Advances in Neural Information Processing Systems*, 35:11217–11228, 2022.
- [86] Justin Whitehouse, Aaditya Ramdas, Ryan Rogers, and Steven Wu. Fully-adaptive composition in differential privacy. In *International Conference on Machine Learning*, pages 36990–37007. PMLR, 2023.
- [87] Justin Whitehouse, Zhiwei Steven Wu, and Aaditya Ramdas. On the sublinear regret of GP-UCB. *Advances in Neural Information Processing Systems*, 2023.
- [88] Justin Whitehouse, Zhiwei Steven Wu, and Aaditya Ramdas. Time-uniform self-normalized concentration for vector-valued processes. *arXiv preprint arXiv:2310.09100*, 2023.
- [89] Yiqing Xu. Generalized synthetic control method: Causal inference with interactive fixed effects models. *Political Analysis*, 25(1):57–76, 2017. doi: 10.1017/pan.2016.2.
- [90] Jingwen Zhang, Yifang Chen, and Amandeep Singh. Causal bandits: Online decision-making in endogenous settings. *arXiv preprint arXiv:2211.08649*, 2022.
- [91] Kelly Zhang, Lucas Janson, and Susan Murphy. Statistical inference with m-estimators on adaptively collected data. *Advances in neural information processing systems*, 34:7460–7471, 2021.
- [92] Kelly W Zhang, Lucas Janson, and Susan A Murphy. Statistical inference after adaptive sampling in non-markovian environments. *arXiv preprint arXiv:2202.07098*, 2022.

A Results on martingale concentration

In this section, we discuss the basics of martingale concentration that will be used ubiquitously throughout this work. Recall that a process $(M_n)_{n \geq 0}$ is a martingale with respect to some filtration $(\mathcal{F}_n)_{n \geq 0}$ if (a) $(M_n)_{n \geq 0}$ is adapted to $(\mathcal{F}_n)_{n \geq 0}$, (b) $\mathbb{E}|M_n| < \infty$ for all $n \geq 0$, and (c) $\mathbb{E}(M_{n+1} | \mathcal{F}_n) = M_n$ for all $n \geq 0$. We denote the “increments” of a martingale as $\Delta M_n := M_n - M_{n-1}$. Results on martingale concentration yield a means of providing tight, time-uniform concentration for statistical tasks in which data is adaptively collected. Recent advances in martingale concentration have allowed for advances in disparate fields of statistical theory such as PAC-Bayesian learning [30], composition in differential privacy [86, 85], and the estimation of convex statistical divergences [63].

Martingale concentration emerges naturally in our work in two ways. First, martingale concentration plays an integral role in bounding the deviations between $\tilde{\theta}_t(a)$, the ridge estimate in the true low-dimensional subspace, and $\theta(a)$. In bounding these deviations, we leverage results on the concentration of self-normalized martingale processes, particularly those of de la Peña et al. [32, 33]. These results have long proven useful for calibrating confidence in online linear regression tasks [3, 29], but to the best of our knowledge, we are the first to couple these results with the low-dimensional estimation innate to PCR.

Second, we leverage martingale methods to control the rate at which PCR’s estimate of the projection operator onto the unknown subspace converges. To accomplish this, we couple recent breakthroughs on time-uniform, self-normalized concentration for scalar-valued processes [50, 51] with the covering and matrix-CGF approaches for bounding the error in estimates of covariance matrices [83, 71, 80]. While this aspect of our analysis is, more or less, a straightforward merging of two techniques for concentration of measure, it nonetheless requires care due to the technical nature of the machinery being used.

We start by recounting the time-uniform concentration inequality we leverage for controlling the error in the ridge estimate in the true, low-dimensional subspace. The following result is from Abbasi-Yadkori et al. [3], but is a special case of more general, self-normalized concentration results from de la Peña et al. [32, 33].

Lemma A.1 (Method of Mixtures). *Let $(\mathcal{F}_t)_{t \geq 0}$ be a filtration. Let $S_t = \sum_{s=1}^t \epsilon_s X_s$ where $(\epsilon_t)_{t \geq 1}$ is an $(\mathcal{F}_t)_{t \geq 0}$ -adapted \mathbb{R} -valued process of σ -subGaussian random variables and $(X_t)_{t \geq 1}$ is an $(\mathcal{F}_t)_{t \geq 1}$ -predictable \mathbb{R}^d -valued process. Let $\rho > 0$ be arbitrary, and let*

$$\mathcal{V}_t := \sum_{s=1}^t X_s X_s^T + \rho I_d.$$

Let $\delta \in (0, 1)$ be any confidence parameter. Then we have with probability at least $1 - \delta$, simultaneously for all $t \geq 1$,

$$\|\mathcal{V}_t^{-1/2} S_t\|_2 \leq \sigma \sqrt{2 \log \left(\frac{1}{\delta} \sqrt{\det(\rho^{-1} \mathcal{V}_t)} \right)}.$$

We make several brief comments about the above lemma. First, note that while the word “martingale” doesn’t directly appear, the process $(S_n)_{n \geq 0}$ is, in fact, a martingale with respect to the filtration $(\mathcal{F}_n)_{n \geq 0}$. The concentration inequality follows from “mixing” a family of martingales based on $(S_n)_{n \geq 0}$ with respect to some suitable probability measure. Second, how we leverage Lemma A.1 in conjunction with the low-dimensional structure of the problem comes through the presence of $\det(\rho^{-1} \mathcal{V}_n)$ in the bound. In particular, if the sequence $(X_n)_{n \geq 0}$ lies in some low-dimensional subspace W , with $\dim(W) = r$, then $\rho^{-1} \mathcal{V}_n$ will have at most r non-unit eigenvalues, and hence $\log \det(\rho^{-1} \mathcal{V}_n) \approx r \log(\|\rho^{-1} \mathcal{V}_n\|_{op})$. We exploit this idea further in the sequel.

Now, we discuss the scalar-valued martingale concentration results from Howard et al. [50, 51] that will be used in our work. Before this, recall that a random variable X is said to be σ -subGaussian if $\log \mathbb{E} e^{\lambda X} \leq \frac{\lambda^2 \sigma^2}{2}$ for all $\lambda \in \mathbb{R}$. Likewise, we say X is (σ, c) -subExponential if $\log \mathbb{E} e^{\lambda X} \leq \frac{\lambda^2 \sigma^2}{2}$ for all $|\lambda| < \frac{1}{c}$, and we say X is (σ, c) -subGamma if $\log \mathbb{E} e^{\lambda X} \leq \sigma^2 \psi_{G,c}(\lambda)$ for all $|\lambda| < \frac{1}{c}$, where $\psi_{G,c}(\lambda) := \frac{\lambda^2}{2(1-c\lambda)}$. A particularly useful fact in the sequel is that if X is (σ, c) -subExponential, it is also (σ, c) -subGamma [50].

Lemma A.2. Suppose $(X_n)_{n \geq 1}$ is a sequence of independent, (σ, c) -subGamma random variables. Let $(S_n)_{n \geq 0}$ be given by $S_n := \sum_{m=1}^n X_m$, and let $\delta \in (0, 1)$ be arbitrary. Then, with probability at least $1 - \delta$, simultaneously for all $n \geq 0$, we have

$$S_n \leq \frac{3}{2} \sigma \sqrt{n \ell_\delta(n)} + \frac{5}{2} c \ell_\delta(n),$$

where $\ell_\delta(n) := 2 \log \log(2n) + \log\left(\frac{d\pi^2}{12\delta}\right)$.

Note that we have simplified above bound from Howard et al. [51] in several ways. First, the bound as presented in Howard et al. [51] applies to general classes of processes with potentially correlated increments. We have simplified the bound to the setting of independent, subGamma increments to suit our setting. Further, the original bound has many parameters, each of which can be fine-tuned to fit an application. We have pre-selected parameters so that (a) the bound is legible and (b) constants remain relatively small.

We can couple the above bound with a standard argument for bounding error in covariance estimation [83] to obtain high-probability, time-uniform results. The following result will prove useful in measuring the rapidity at which PCA can learn the true, low-dimensional subspace in which the noiseless contexts and slope vectors lie.

Lemma A.3. Let $(\epsilon_n)_{n \geq 1}$ be a sequence of independent, σ -subGaussian random vectors in \mathbb{R}^d . Then, for any $\delta \in (0, 1)$, with probability at least $1 - \delta$, simultaneously for all $n \geq 1$, we have

$$\left\| \sum_{m=1}^n \epsilon_m \epsilon_m^\top - \mathbb{E} \epsilon_m \epsilon_m^\top \right\|_{op} \leq \beta \left(3 \sqrt{n \ell_{\delta/2\mathcal{N}}(n)} + 5 \ell_{\delta/2\mathcal{N}}(n) \right),$$

where ℓ_δ is as defined in Lemma A.2, $\beta = 32\sigma^2 e^2$, and $\mathcal{N} = 17^d$ is an upper bound on the $1/8$ -covering number of \mathbb{S}^{d-1} .

Proof. Define the process $(M_n)_{n \geq 0}$ by $M_n := \sum_{m=1}^n \epsilon_m \epsilon_m^\top - \mathbb{E} \epsilon_m \epsilon_m^\top$. Using a standard covering argument (more or less verbatim from Wainwright [83]), we have that, if $K \subset \mathbb{S}^{d-1}$ is a minimal $1/8$ -covering of \mathbb{S}^{d-1} , then

$$\|M_n\|_{op} \leq 2 \max_{\nu \in K} \langle \nu, M_n \nu \rangle \quad (4)$$

It is clear that $(M_n)_{n \geq 0}$ is a Hermitian matrix-valued martingale with respect to the natural filtration $(\mathcal{F}_n)_{n \geq 0}$ given by $\mathcal{F}_n := \sigma(\epsilon_m : m \leq n)$. It is thus straightforward to see that, for any $\nu \in \mathbb{S}^{d-1}$, the process $(M_n^\nu)_{n \geq 0}$ given by $M_n^\nu := \nu^\top M_n \nu$ is a real-valued martingale with respect to this same filtration.

Moreover, a standard argument (see the proof of Theorem 6.5 in Wainwright [83]) yields that

$$\log \left[\mathbb{E} e^{\lambda \langle \nu, \Delta M_n \nu \rangle} \right] \leq \frac{\lambda^2 \beta^2}{2} \quad \text{for all } |\lambda| < \frac{1}{\beta}, \quad (5)$$

where $\beta := 32e^2\sigma^2$. In other words, for any $\nu \in \mathbb{S}^{d-1}$, the random variable $\langle \nu, \Delta M_n \nu \rangle$ is (β, β) -sub-Exponential, as outlined above. In particular, per the results of Howard et al. [50, 51], this implies that the random variable $\langle \nu, \Delta M_n \nu \rangle$ is (β, β) -subGamma. Applying Lemma A.2, we have, with probability at least $1 - \delta$, simultaneously for all $n \geq 1$ and $\nu \in K$, that

$$|\langle \nu, M_n \nu \rangle| \leq \frac{3}{2} \beta \sqrt{n \ell_{\delta/2\mathcal{N}}(n)} + \beta \frac{5}{2} \ell_{\delta/2\mathcal{N}}(n).$$

Plugging this into Inequality 4, we have that, with probability at least $1 - \delta$, simultaneously for all $n \geq 1$,

$$\|M_n\|_{op} \leq \beta \left(3 \sqrt{n \ell_{\delta/2\mathcal{N}}(n)} + 5 \ell_{\delta/2\mathcal{N}}(n) \right).$$

□

In the case that the sequence of independent noise variables $(\epsilon_n)_{n \geq 1}$ satisfies $\mathbb{E} \epsilon_n \epsilon_n^\top = \Sigma$ and $\|\epsilon_n\|_2 \leq \sqrt{B}$ uniformly in n , we can obtain significantly tighter bounds (in terms of constants). Of particular import is the following bound from Howard et al. [51], which combines their ‘‘stitching approach’’ to time-uniform concentration with matrix chernoff techniques [50, 80, 83] to obtain time-uniform bounds on covariance estimation.

Lemma A.4. Let $(\epsilon_n)_{n \geq 1}$ be a sequence of mean zero, independent random vectors in \mathbb{R}^d such that, for all $n \geq 1$, we have $\mathbb{E}\epsilon_n\epsilon_n^\top = \Sigma$ and $\|\epsilon_n\|_2 \leq \sqrt{Cd}$ almost surely. Further assume $\|\Sigma\|_{op} \leq \gamma$. Then, for any $\delta \in (0, 1)$, we have, with probability at least $1 - \delta$, simultaneously for all $n \geq 1$

$$\left\| \sum_{m=1}^n \{\epsilon_m\epsilon_m^\top - \Sigma\} \right\|_{op} \leq \frac{3}{2} \sqrt{nCd\gamma\ell_\delta(n)} + \frac{7}{3} Cd\ell_\delta(n)$$

where $\ell_\delta(n) := 2 \log \log(2n) + \log\left(\frac{d\pi^2}{12\delta}\right)$.

B Convergence results for PCA and singular values

In this appendix, we discuss time-uniform convergence results for adaptive principal component analysis — one half of the principal component regression (PCR) algorithm. In addition, we analyze how singular values of \mathbf{Z}_n , the noisy covariate matrix, cluster around those of \mathbf{X}_n , the true covariate matrix. We, in the style of Agarwal et al. [7], reduce our study of these quantities to the study of the operator norm of \mathcal{E}_n , which we control using the martingale concentration results outlined in Appendix A.

Before continuing, we enumerate several linear algebraic facts that are useful in bounding deviations in singular values.

Lemma B.1 (Weyl's Inequality). Let $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{t \times n}$. Then, for any $i \in [t \wedge n]$, we have

$$|\sigma_i(\mathbf{A}) - \sigma_i(\mathbf{B})| \leq \|\mathbf{A} - \mathbf{B}\|_{op}.$$

Lemma B.2 (Wedin's Lemma). Let $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{t \times n}$, and suppose \mathbf{A} and \mathbf{B} have spectral decompositions

$$\mathbf{A} = \mathbf{U}\Sigma\mathbf{V}^T \quad \text{and} \quad \mathbf{B} = \widehat{\mathbf{U}}\widehat{\Sigma}\widehat{\mathbf{V}}^T.$$

Then, for any $r \leq n \wedge d$, we have

$$\max \left\{ \left\| \mathbf{U}_r\mathbf{U}_r^T - \widehat{\mathbf{U}}_r\widehat{\mathbf{U}}_r^T \right\|_{op}, \left\| \mathbf{V}_r\mathbf{V}_r^T - \widehat{\mathbf{V}}_r\widehat{\mathbf{V}}_r^T \right\|_{op} \right\} \leq \frac{2\|\mathbf{A} - \mathbf{B}\|_{op}}{\sigma_r - \sigma_{r+1}},$$

where σ_r (resp. σ_{r+1}) is the r -th (resp. $r + 1$ -st) largest singular value of \mathbf{A} .

We now prove a time-uniform, high probability bound on the operator norm of \mathcal{E}_n . (Recall $\mathcal{E}_n = (\epsilon_1, \dots, \epsilon_n)^T$.) In particular, we prove a bound for two settings — a looser bound (in terms of constants) which holds when the noise variables ϵ_n are assumed to be subGaussian, and a tighter, more practically relevant bound that holds when ϵ_n are assumed to be bounded.

Lemma B.3 (Covariance Noise Bound). Let $(\epsilon_n)_{n \geq 1}$ be a sequence of independent, mean zero random vectors in \mathbb{R}^d such that $\|\mathbb{E}\epsilon_n\epsilon_n^\top\|_{op} \leq \gamma$ for all $n \geq 1$. Let $\delta \in (0, 1)$ be an arbitrary confidence parameter. Then, with probability at least $1 - \delta$, simultaneously for all $n \geq 1$ we have

$$\|\mathcal{E}_n\|_{op}^2 \leq \begin{cases} \beta \left(3\sqrt{n\ell_\delta/2\mathcal{N}(n)} + 5\ell_\delta/2\mathcal{N}(n) \right) + n\gamma & \text{when Assumption 3.1 holds} \\ \frac{3}{2} \sqrt{nCd\gamma\ell_\delta(n)} + \frac{7}{3} Cd\ell_\delta(n) + n\gamma & \text{when Assumption 3.2 holds.} \end{cases}$$

where ℓ_δ is as defined in Lemma A.2, $\beta = 32\sigma^2e^2$, and $\mathcal{N} = 17^d$ is an upper bound on the $1/8$ -covering number of \mathbb{S}^{d-1} .

Proof. Observe the following basic chain of reasoning.

$$\|\mathcal{E}_n\|_{op}^2 = \|\mathcal{E}_n^\top \mathcal{E}_n\|_{op} \leq \|\mathcal{E}_n^\top \mathcal{E}_n - \mathbb{E}\mathcal{E}_n^\top \mathcal{E}_n\|_{op} + \|\mathbb{E}\mathcal{E}_n^\top \mathcal{E}_n\|_{op}.$$

Now, both Assumption 3.1 and 3.2 imply that $\|\mathbb{E}[\epsilon_n\epsilon_n^\top]\|_{op} \leq \gamma$ for all $n \geq 1$. Consequently, noting that $\mathcal{E}_n^\top \mathcal{E}_n = \sum_{m=1}^n \epsilon_m\epsilon_m^\top$, we have that

$$\|\mathbb{E}\mathcal{E}_n^\top \mathcal{E}_n\|_{op} \leq \sum_{m=1}^n \|\mathbb{E}\epsilon_m\epsilon_m^\top\|_{op} \leq n\gamma.$$

Likewise, we have, applying Lemma A.3 in the case the ϵ_n satisfy the subGaussian assumption and Lemma A.4 in the case the noise satisfies the bounded assumption, that, with probability at least $1 - \delta$, simultaneously for all $n \geq 1$,

$$\|\mathcal{E}_n^\top \mathcal{E}_n - \mathbb{E} \mathcal{E}_n^\top \mathcal{E}_n\|_{op} \leq \begin{cases} \beta \left(3\sqrt{n\ell_{\delta/2\mathcal{N}}(n)} + 5\ell_{\delta/2\mathcal{N}}(n) \right) & \text{when Assumption 3.1 holds} \\ \frac{3}{2}\sqrt{nCd\gamma\ell_{\delta}(n)} + \frac{7}{3}Cd\ell_{\delta}(n) & \text{when Assumption 3.2 holds.} \end{cases}$$

Adding these two lines of reasoning together yields the desired conclusion. \square

We can now apply Lemma B.3 to bound the rate at which $\widehat{\mathbf{P}}_n$, the projection operator onto the learned subspace at time n , converges to the projection operator \mathbf{P} onto the true, low-dimensional subspace.

Lemma B.4 (Projection Convergence). *Let $\widehat{\mathbf{P}}_n$ denote the projection operator onto the learned subspace. Further, let \mathbf{P} denote the projection operator onto the true unknown subspace W^* . Assume $\text{rank}(\mathbf{X}_{n_0}) = r$ for some $n_0 \geq 1$. Then, for any $\delta \in (0, 1)$, we have with probability at least $1 - \delta$, simultaneously for all $n \geq n_0$,*

$$\|\widehat{\mathbf{P}}_n - \mathbf{P}\|_{op}^2 \leq \begin{cases} \frac{4\beta(3\sqrt{n\ell_{\delta/2\mathcal{N}}(n)} + 5\ell_{\delta/2\mathcal{N}}(n)) + 4n\gamma}{\sigma_r(\mathbf{X}_n)^2} & \text{when Assumption 3.1 holds} \\ \frac{6\sqrt{nCd\gamma\ell_{\delta}(n)} + \frac{14}{3}Cd\ell_{\delta}(n) + 4n\gamma}{\sigma_r(\mathbf{X}_n)^2} & \text{when Assumption 3.2 holds.} \end{cases}$$

Proof. We write the singular value decompositions of \mathbf{Z}_n and \mathbf{X}_n respectively as

$$\mathbf{Z}_n = \widehat{\mathbf{U}}_n \widehat{\Sigma}_n \widehat{\mathbf{V}}_n^\top \quad \text{and} \quad \mathbf{X}_n = \check{\mathbf{U}}_n \check{\Sigma}_n \check{\mathbf{V}}_n^\top.$$

Since we have assumed $\text{rank}(\mathbf{X}_{n_0}) = r$ for some $n_0 \geq 1$, we have, for all $n \geq n_0$,

$$\widehat{\mathbf{P}}_n = \widehat{\mathbf{V}}_{n,r} \widehat{\mathbf{V}}_{n,r}^\top \quad \text{and} \quad \mathbf{P} = \check{\mathbf{V}}_{n,r} \check{\mathbf{V}}_{n,r}^\top.$$

Now, applying Lemma B.2 and Lemma B.3 we have that, with probability at least $1 - \delta$, simultaneously for all $n \geq n_0$,

$$\begin{aligned} \|\widehat{\mathbf{P}}_n - \mathbf{P}\|_{op}^2 &= \left\| \widehat{\mathbf{V}}_{n,r} \widehat{\mathbf{V}}_{n,r}^\top - \check{\mathbf{V}}_{n,r} \check{\mathbf{V}}_{n,r}^\top \right\|_{op}^2 \\ &\leq \frac{4 \|\mathbf{X}_n - \mathbf{Z}_n\|_{op}^2}{\sigma_r(\mathbf{X}_n)^2} \\ &= \frac{4 \|\mathcal{E}_n\|_{op}^2}{\sigma_r(\mathbf{X}_n)^2} \\ &\leq \begin{cases} \frac{4\beta(3\sqrt{n\ell_{\delta/2\mathcal{N}}(n)} + 5\ell_{\delta/2\mathcal{N}}(n)) + 4n\gamma}{\sigma_r(\mathbf{X}_n)^2} & \text{when Assumption 3.1 holds} \\ \frac{6\sqrt{nCd\gamma\ell_{\delta}(n)} + \frac{14}{3}Cd\ell_{\delta}(n) + 4n\gamma}{\sigma_r(\mathbf{X}_n)^2} & \text{when Assumption 3.2 holds.} \end{cases} \end{aligned}$$

This proves the desired result. \square

We can obtain the following empirical version of Lemma B.4. The proof of the following is identical—the only difference is that, in the application of Wedin’s theorem (Lemma B.2), we put the singular values of \mathbf{Z}_n in the denominator instead of \mathbf{X}_n . This inequality is, in practice, more useful in computing confidence bounds, as the singular values of \mathbf{Z}_n are computable, whereas the singular values of \mathbf{X}_n are not.

Lemma B.5 (Projection Convergence). *Let $\widehat{\mathbf{P}}_n$ denote the projection operator onto the learned subspace. Further, let \mathbf{P} denote the projection operator onto the true unknown subspace W^* . Assume $\text{rank}(\mathbf{X}_{n_0}) = r$ for some $n_0 \geq 1$. Then, for any $\delta \in (0, 1)$, we have with probability at least $1 - \delta$, simultaneously for all $n \geq n_0$,*

$$\|\widehat{\mathbf{P}}_n - \mathbf{P}\|_{op} \leq \begin{cases} \frac{4\beta(3\sqrt{n\ell_{\delta/2\mathcal{N}}(n)} + 5\ell_{\delta/2\mathcal{N}}(n)) + 4n\gamma}{(\sigma_r(\mathbf{Z}_n) - \sigma_{r+1}(\mathbf{Z}_n))^2} & \text{when Assumption 3.1 holds} \\ \frac{6\sqrt{nCd\gamma\ell_{\delta}(n)} + \frac{14}{3}Cd\ell_{\delta}(n) + 2n\gamma}{(\sigma_r(\mathbf{Z}_n) - \sigma_{r+1}(\mathbf{Z}_n))^2} & \text{when Assumption 3.2 holds.} \end{cases}$$

What remains is to provide a concentration inequality bounding the deviations between the singular values of the noisy covariate matrix \mathbf{Z}_n and the true, low rank covariate matrix \mathbf{X}_n . We have the following inequality.

Lemma B.6. *Let $\delta \in (0, 1)$ be arbitrary. Then with probability at least $1 - \delta$ we have, simultaneously for all $i \in [r]$ and $n \geq 1$,*

$$|\sigma_i(\mathbf{Z}_n) - \sigma_i(\mathbf{X}_n)|^2 \leq \begin{cases} \beta (3\sqrt{n\ell_{\delta/2\mathcal{N}}(n)} + 5\ell_{\delta/2\mathcal{N}}(n)) + n\gamma & \text{when Assumption 3.1 holds} \\ \frac{3}{2}\sqrt{nCd\gamma\ell_{\delta}(n)} + \frac{7}{3}Cd\ell_{\delta}(n) + n\gamma & \text{when Assumption 3.2 holds.} \end{cases}$$

In addition, on the same probability at least $1 - \delta$ event, for all $n \geq 1$ such that $\text{snr}_n \geq 2$, we have

$$\frac{\sigma_r(\mathbf{X}_n)}{2} \leq \sigma_r(\mathbf{Z}_n) \leq \frac{3\sigma_r(\mathbf{X}_n)}{2}.$$

Note that the same holds if snr_n is replaced with the action-specific signal to noise ratio $\text{snr}_n(a)$.

Proof. By Lemma B.1, we know that for any $i \in [r]$,

$$|\sigma_i(\mathbf{Z}_n) - \sigma_i(\mathbf{X}_n)| \leq \|\mathbf{Z}_n - \mathbf{X}_n\|_{op} = \|\mathcal{E}_n\|_{op}.$$

Thus, applying Lemma B.3 yields the first part of the theorem.

Now, suppose $\text{snr}_n \geq 2$. Let U_n^2 denote either the first or second line of the already-shown inequality (depending on whether Assumption 3.1 or Assumption 3.2 holds). Then with probability at least $1 - \delta$,

$$\begin{aligned} \sigma_i(\mathbf{Z}_n) &\geq \sigma_i(\mathbf{X}_n) - U_n \\ &= \sigma_i(\mathbf{X}_n) - \sigma_i(\mathbf{X}_n) \frac{U_n}{\sigma_i(\mathbf{X}_n)} \\ &\geq \sigma_i(\mathbf{X}_n) - \sigma_i(\mathbf{X}_n) \frac{U_n}{\sigma_r(\mathbf{X}_n)} \\ &= \sigma_i(\mathbf{X}_n) - \sigma_i(\mathbf{X}_n) \frac{1}{\text{snr}_n} \\ &\geq \frac{\sigma_i(\mathbf{X}_n)}{2}. \end{aligned}$$

□

C Convergence results for regression in the true subspace

In this section, we construct confidence ellipsoids bounding the error between $\check{\theta}_n(a)$ and $\theta(a)$, where

$$\check{\theta}_n(a) := \check{\mathcal{V}}_n(a)^{-1} \check{\mathbf{Z}}_n(a)^\top \mathbf{Y}_n(a),$$

i.e. $\check{\theta}_n(a)$ is the estimate of $\theta(a)$ given access to the true, underlying subspace W^* . While the learner never has direct access to W^* , the quantity $\check{\theta}_n(a)$ proves useful in bounding $\|\widehat{\theta}_n(a) - \theta(a)\|_2^2$, the ℓ_2 error between the PCR estimate of $\theta(a)$ and $\theta(a)$ itself. In other words, the bounds proved in this section provide a theoretical tool for understanding the convergence of PCR in adaptive settings. The main lemma in this appendix is the following.

Lemma C.1 (Error bound in true subspace). *Let $\delta \in (0, 1)$ be an arbitrary confidence parameter. Then, with probability at least $1 - 2\delta$, simultaneously for all $a \in [A]$ and $n \geq 1$, we have*

$$\begin{aligned} \left\| \check{\mathcal{V}}_n(a)^{1/2} \left(\check{\theta}_n(a) - \theta(a) \right) \right\|_2^2 &\leq 4\rho L^2 + 8\eta^2 \left[\log \left(\frac{A}{\delta} \right) + r \log \left(1 + \frac{\sigma_1(\mathbf{Z}_n(a))^2}{\rho} \right) \right] \\ &+ 2L^2 \begin{cases} \beta (3\sqrt{n\ell_{\delta/2\mathcal{N}}(n)} + 5\ell_{\delta/2\mathcal{N}}(n)) + n\gamma & \text{when Assumption 3.1 holds} \\ \frac{3}{2}\sqrt{nCd\gamma\ell_{\delta}(n)} + \frac{7}{3}Cd\ell_{\delta}(n) + n\gamma & \text{when Assumption 3.2 holds.} \end{cases} \end{aligned}$$

where β , ℓ_{δ} , and \mathcal{N} are as defined in Lemma A.3 and A.4.

To make the proof of the above lemma more modular, we introduce several lemmas. The first of these lemmas simply provides an alternative, easier-to-bound representation of the difference (or error) vector $\check{\theta}_n(a) - \theta(a)$

Lemma C.2. *For any $n \in [N]$ and $a \in [A]$, we have*

$$\check{\theta}_n(a) - \theta(a) = \check{\mathcal{V}}_n(a)^{-1} \check{\mathbf{Z}}_n(a)^\top \Xi_n(a) + \check{\mathcal{V}}_n(a)^{-1} \check{\mathbf{Z}}_n(a)^\top \check{\mathcal{E}}_n(a) \theta(a) - \rho \check{\mathcal{V}}_n(a)^{-1} \theta(a).$$

Proof. A straightforward computation yields

$$\begin{aligned} \check{\theta}_n(a) - \theta(a) &= \check{\mathcal{V}}_n(a)^{-1} \check{\mathbf{Z}}_n(a)^\top \mathbf{Y}_n(a) - \theta(a) \\ &= \check{\mathcal{V}}_n(a)^{-1} \check{\mathbf{Z}}_n(a) (\mathbf{X}_n(a) \theta(a) + \Xi_n(a)) - \theta(a) \\ &= \check{\mathcal{V}}_n(a)^{-1} \check{\mathbf{Z}}_n(a) \mathbf{X}_n(a) \theta(a) + \check{\mathcal{V}}_n(a)^{-1} \check{\mathbf{Z}}_n(a)^\top \Xi_n(a) - \theta(a) \\ &\quad \pm \check{\mathcal{V}}_n(a)^{-1} \check{\mathbf{Z}}_n(a) \check{\mathcal{E}}_n(a) \theta(a) \pm \rho \check{\mathcal{V}}_n(a)^{-1} \theta(a) \\ &= \check{\mathcal{V}}_n(a)^{-1} \left[\check{\mathbf{Z}}_n(a)^\top (\mathbf{X}_n(a) + \check{\mathcal{E}}_n(a)) + \rho I_d \right] \theta(a) + \check{\mathcal{V}}_n(a)^{-1} \check{\mathbf{Z}}_n(a)^\top \Xi_n(a) - \theta(a) \\ &\quad - \check{\mathcal{V}}_n(a)^{-1} \check{\mathbf{Z}}_n(a) \check{\mathcal{E}}_n(a) \theta(a) - \rho \check{\mathcal{V}}_n(a)^{-1} \theta(a) \\ &= \check{\mathcal{V}}_n(a)^{-1} \left[\check{\mathbf{Z}}_n(a)^\top \check{\mathbf{Z}}_n(a) + \rho I_d \right] \theta(a) + \check{\mathcal{V}}_n(a)^{-1} \check{\mathbf{Z}}_n(a)^\top \Xi_n(a) - \theta(a) \\ &\quad - \check{\mathcal{V}}_n(a)^{-1} \check{\mathbf{Z}}_n(a) \check{\mathcal{E}}_n(a) \theta(a) - \rho \check{\mathcal{V}}_n(a)^{-1} \theta(a) \\ &= \check{\mathcal{V}}_n(a)^{-1} \check{\mathbf{Z}}_n(a)^\top \Xi_n(a) - \check{\mathcal{V}}_n(a)^{-1} \check{\mathbf{Z}}_n(a)^\top \check{\mathcal{E}}_n(a) \theta(a) - \rho \check{\mathcal{V}}_n(a)^{-1} \theta(a), \end{aligned}$$

which proves the desired result. \square

We leverage the following technical lemma in bounding the determinant of $\check{\mathcal{V}}_n(a)$, the projection of the covariance matrix $\mathcal{V}_n(a)$ onto the true, unknown subspace W^* .

Lemma C.3 (Determinant-Trace Inequality). *Let $\mathbf{A} \in \mathbb{R}^{d \times d}$ be a positive-semidefinite matrix of rank r . Then, for any $\rho > 0$, we have*

$$\log \det(I_d + \rho^{-1} \mathbf{A}) \leq r \log \left(1 + \frac{\sigma_1(\mathbf{A})}{\rho} \right).$$

Proof. We have

$$\begin{aligned} \log \det(I_d + \rho^{-1} \mathbf{A}) &= r \log \left[\det(I_d + \rho^{-1} \mathbf{A})^{1/r} \right] \\ &= r \log \prod_{i=1}^r (1 + \rho^{-1} \sigma_i(\mathbf{A}))^{1/r} \\ &\leq r \log \left(\sum_{i=1}^r \frac{1 + \rho^{-1} \sigma_i(\mathbf{A})}{r} \right) \\ &\leq r \log \left(1 + \frac{\sigma_1(\mathbf{A})}{\rho} \right). \end{aligned}$$

\square

Lemma C.4. *Let $\mathbf{A} \in \mathbb{R}^{t \times d}$ be a matrix, and let $\rho > 0$ be arbitrary. Suppose \mathbf{A} has $1 \leq k \leq t \wedge d$ non-zero singular values. Then,*

$$\left\| (\mathbf{A}^T \mathbf{A} + \rho I_d)^{-1/2} \mathbf{A}^T \right\|_{op} \leq \frac{\sigma_k(\mathbf{A})}{\sqrt{\sigma_k(\mathbf{A})^2 + \rho}} \leq 1$$

Proof. Let us write the singular value decomposition of \mathbf{A} as

$$\mathbf{A} = \mathbf{U} \Sigma \mathbf{V}^T = \mathbf{U}_k \Sigma_k \mathbf{V}_k^T,$$

where the second equality follows from the fact that \mathbf{A} has exactly k non-zero singular values. We have the equality

$$\mathbf{A}^T \mathbf{A} + \rho I_d = \mathbf{V}(\Sigma_k^2 + \rho I_k) \mathbf{V}^T.$$

Thus, we see

$$\begin{aligned} \left\| (\mathbf{A}^T \mathbf{A} + \rho I_d)^{-1/2} \mathbf{A}^T \right\|_{op} &= \left\| \mathbf{V}_k (\Sigma_k^2 + \rho I_k)^{-1/2} \mathbf{V}_k^T \mathbf{V}_k \Sigma_k \mathbf{U}_k^T \right\|_{op} \\ &= \left\| (\Sigma_k^2 + \rho I_k)^{-1/2} \Sigma_k \right\|_{op} \\ &= \max_{i=1}^k \frac{\sigma_i(\mathbf{A})}{\sqrt{\rho + \sigma_i(\mathbf{A})^2}} \leq \frac{\sigma_k(\mathbf{A})}{\sqrt{\rho + \sigma_k(\mathbf{A})^2}}. \end{aligned}$$

□

With the aforementioned technical lemmas, alongside the time-uniform martingale bounds presented in Appendix A and the singular value concentration results proved in Appendix B, we can now prove Lemma C.1.

Proof of Lemma C.1. First, by Lemma C.2, we see that we have, for any $a \in [A]$ and $n \geq 1$,

$$\check{\theta}_n(a) - \theta(a) = \check{\mathcal{V}}_n(a)^{-1} \check{\mathbf{Z}}_n(a)^\top \check{\Xi}_n(a) + \check{\mathcal{V}}_n(a)^{-1} \check{\mathbf{Z}}_n(a)^\top \check{\mathcal{E}}_n(a) \theta(a) - \rho \check{\mathcal{V}}_n(a)^{-1} \theta(a).$$

With this decomposition in hand, we can apply the parallelogram inequality ($\|x - y\|_2^2 + \|x + y\|_2^2 \leq 2\|x\|_2^2 + 2\|y\|_2^2$) twice to see that

$$\begin{aligned} &\left\| \check{\mathcal{V}}_n(a)^{1/2} (\check{\theta}_n(a) - \theta(a)) \right\|_2^2 \\ &\leq 4 \left\| \check{\mathcal{V}}_n(a)^{-1/2} \check{\mathbf{Z}}_n(a)^\top \check{\Xi}_n(a) \right\|_2^2 + 4\rho^2 \left\| \check{\mathcal{V}}_n(a)^{-1/2} \theta(a) \right\|_2^2 + 2 \left\| \check{\mathcal{V}}_n(a)^{-1/2} \check{\mathbf{Z}}_n(a)^\top \check{\mathcal{E}}_n(a) \theta(a) \right\|_2^2 \\ &\leq 4 \left\| \check{\mathcal{V}}_n(a)^{-1/2} \check{\mathbf{Z}}_n(a)^\top \check{\Xi}_n(a) \right\|_2^2 + 4\rho L^2 + 2 \left\| \check{\mathcal{V}}_n(a)^{-1/2} \check{\mathbf{Z}}_n(a)^\top \check{\mathcal{E}}_n(a) \theta(a) \right\|_2^2. \end{aligned}$$

We bound the first and last terms separately. Applying Lemma A.1 for each $a \in [A]$ and taking a union bound over actions yields that, with probability at least $1 - \delta$, simultaneously for all $n \geq 1$ and $a \in [A]$,

$$\begin{aligned} \left\| \check{\mathcal{V}}_n(a)^{-1/2} \check{\mathbf{Z}}_n(a)^\top \check{\Xi}_n(a) \right\|_2 &\leq \eta \sqrt{2 \log \left(\frac{A}{\delta} \sqrt{\det(\rho^{-1} \check{\mathcal{V}}_n(a))} \right)} \\ &\leq \eta \sqrt{2 \left[\log \left(\frac{A}{\delta} \right) + r \log \left(1 + \frac{\sigma_1(\check{\mathbf{Z}}_n(a))^2}{\rho} \right) \right]} \\ &\leq \eta \sqrt{2 \left[\log \left(\frac{A}{\delta} \right) + r \log \left(1 + \frac{\sigma_1(\hat{\mathbf{Z}}_n(a))^2}{\rho} \right) \right]}, \end{aligned}$$

where the second inequality comes from applying Lemma C.3 and the third inequality comes from the fact that $\sigma_i(\hat{\mathbf{Z}}_n(a)) \geq \sigma_i(\check{\mathbf{Z}}_n(a))$ for all $i \in [r]$.

Next, we bound the final term in the above expansion. Observe that

$$\left\| \check{\mathcal{V}}_n(a)^{-1/2} \check{\mathbf{Z}}_n(a)^\top \check{\mathcal{E}}_n(a) \theta(a) \right\|_2 \leq \left\| \check{\mathcal{V}}_n(a)^{-1/2} \check{\mathbf{Z}}_n(a)^\top \right\|_{op} \left\| \check{\mathcal{E}}_n(a) \right\|_{op} \|\theta(a)\|_2.$$

First, note that $\|\theta(a)\|_2 \leq L$ by assumption. Next, Lemma C.4 yields

$$\left\| \check{\mathcal{V}}_n(a)^{-1/2} \check{\mathbf{Z}}_n(a)^\top \right\|_{op} \leq 1.$$

Lastly, note that by Lemma B.3, we have that, with probability at least $1 - \delta$, simultaneously for all $n \geq 1$

$$\left\| \check{\mathcal{E}}_n(a) \right\|_{op}^2 \leq \|\mathcal{E}_n\|^2 \leq \begin{cases} \beta (3\sqrt{n\ell_{\delta/2\mathcal{N}}(n)} + 5\ell_{\delta/2\mathcal{N}}(n)) + n\gamma & \text{when Assumption 3.1 holds} \\ \frac{3}{2}\sqrt{nCd\gamma\ell_{\delta}(n)} + \frac{7}{3}Cd\ell_{\delta}(n) + n\gamma & \text{when Assumption 3.2 holds.} \end{cases}$$

This proves the desired inequality. □

D Technical lemmas for Theorem 4.4

In this appendix, we present several additional technical lemmas that are needed for bounding terms in the main theorem.

Lemma D.1. *Assuming the setup of Theorem 4.4, we have, for all $n \geq 1$,*

$$\left\| \widehat{\mathbf{Z}}_n(a) \widehat{\theta}_n(a) - \mathbf{X}_n(a) \theta(a) \right\|_2^2 \leq 8 \left\| \check{\mathbf{Y}}_n(a)^{1/2} \left(\check{\theta}_n(a) - \theta(a) \right) \right\|_2^2 + 6 \|\Xi_n(a)\|_2^2 + 8 \left\| \widehat{\mathbf{Z}}_n(a) \theta(a) - \mathbf{X}_n(a) \theta(a) \right\|_2^2$$

Proof. A relatively straightforward computation yields

$$\begin{aligned} \left\| \widehat{\mathbf{Z}}_n(a) \widehat{\theta}_n(a) - \mathbf{X}_n(a) \theta(a) \right\|_2^2 &= \left\| \widehat{\mathbf{Z}}_n(a) \widehat{\theta}_n(a) - \mathbf{X}_n(a) \theta(a) + \Xi_n(a) - \Xi_n(a) \right\|_2^2 \\ &\leq 2 \left\| \widehat{\mathbf{Z}}_n(a) \widehat{\theta}_n(a) - \mathbf{Y}_n(a) \right\|_2^2 + 2 \|\Xi_n(a)\|_2^2 \\ &\leq 2 \left\| \widehat{\mathbf{Z}}_n(a) \check{\theta}_n(a) - \mathbf{Y}_n(a) \right\|_2^2 + 2 \|\Xi_n(a)\|_2^2 \\ &\leq 4 \left\| \widehat{\mathbf{Z}}_n(a) \check{\theta}_n(a) - \mathbf{X}_n(a) \theta(a) \right\|_2^2 + 6 \|\Xi_n(a)\|_2^2 \\ &\leq 8 \left\| \widehat{\mathbf{Z}}_n(a) \left(\check{\theta}_n(a) - \theta(a) \right) \right\|_2^2 + 6 \|\Xi_n(a)\|_2^2 + 8 \left\| \widehat{\mathbf{Z}}_n(a) \theta(a) - \mathbf{X}_n(a) \theta(a) \right\|_2^2 \\ &= 8 \left\| \widehat{\mathbf{Z}}_n(a) \mathbf{P} \left(\check{\theta}_n(a) - \theta(a) \right) \right\|_2^2 + 6 \|\Xi_n(a)\|_2^2 + 8 \left\| \widehat{\mathbf{Z}}_n(a) \theta(a) - \mathbf{X}_n(a) \theta(a) \right\|_2^2 \\ &\leq 8 \left\| \check{\mathbf{Z}}_n(a) \left(\check{\theta}_n(a) - \theta(a) \right) \right\|_2^2 + 6 \|\Xi_n(a)\|_2^2 + 8 \left\| \widehat{\mathbf{Z}}_n(a) \theta(a) - \mathbf{X}_n(a) \theta(a) \right\|_2^2 \\ &\leq 8 \left\| \check{\mathbf{Y}}_n(a)^{1/2} \left(\check{\theta}_n(a) - \theta(a) \right) \right\|_2^2 + 6 \|\Xi_n(a)\|_2^2 + 8 \left\| \widehat{\mathbf{Z}}_n(a) \theta(a) - \mathbf{X}_n(a) \theta(a) \right\|_2^2 \end{aligned}$$

where we apply the the parallelogram inequality to obtain the first and third inequalities. The second inequality follows from the characterization of ridge regression in Fact F.2 and the fact ρ is chosen sufficiently small. \square

Lemma D.2. *Assuming the setup of Theorem 4.4, we have, for all $n \geq 1$,*

$$\left\| \mathbf{X}_n(a) \theta(a) - \widehat{\mathbf{Z}}_n(a) \theta(a) \right\|_2^2 \leq 2L^2 \sigma_1(\mathbf{Z}_n(a))^2 \left\| \mathbf{P} - \widehat{\mathbf{P}}_n(a) \right\|_{op}^2 + 2L^2 \|\mathcal{E}_n\|_{op}^2$$

Proof. A relatively simple computation yields

$$\begin{aligned} \left\| \mathbf{X}_n(a) \theta(a) - \widehat{\mathbf{Z}}_n(a) \theta(a) \right\|_2^2 &= \left\| \mathbf{X}_n(a) \theta(a) - \check{\mathbf{Z}}_n(a) \theta(a) + \check{\mathbf{Z}}_n(a) \theta(a) - \widehat{\mathbf{Z}}_n(a) \theta(a) \right\|_2^2 \\ &\leq 2 \left\| \check{\mathbf{Z}}_n(a) \theta(a) - \widehat{\mathbf{Z}}_n(a) \theta(a) \right\|_2^2 + 2 \left\| \check{\mathbf{Z}}_n(a) \theta(a) - \mathbf{X}_n(a) \theta(a) \right\|_2^2 \\ &= 2 \left\| \check{\mathbf{Z}}_n(a) \theta(a) - \widehat{\mathbf{Z}}_n(a) \theta(a) \right\|_2^2 + 2 \left\| \check{\mathcal{E}}_n(a) \theta(a) \right\|_2^2 \\ &= 2 \left\| \mathbf{Z}_n(a) (\mathbf{P} - \widehat{\mathbf{P}}_n(a)) \theta(a) \right\|_2^2 + 2 \left\| \check{\mathcal{E}}_n(a) \theta(a) \right\|_2^2 \\ &\leq 2L^2 \sigma_1(\mathbf{Z}_n(a))^2 \left\| \mathbf{P} - \widehat{\mathbf{P}}_n(a) \right\|_{op}^2 + 2L^2 \left\| \check{\mathcal{E}}_n \right\|_{op}^2, \end{aligned}$$

where we have applied the parallelogram inequality to obtain the first inequality. \square

Lemma D.3. *Suppose $(\xi_n)_{n \geq 1}$ is a sequence of independent, η -subGaussian random variables. Then, for any $\delta \in (0, 1)$, we have simultaneously for all $n \geq 1$*

$$\|\Xi_n\|_2^2 \leq 6\eta^2 \sqrt{2n\ell_\delta(n)} + 10\eta^2 \ell_\delta(n) + n\alpha,$$

where we recall $\Xi_n := (\xi_1, \dots, \xi_n)^\top$.

Proof. First, observe that by the appendix of Honorio and Jaakkola [49], since ξ_n is η -subGaussian, $\xi_n^2 - \mathbb{E}\xi_n^2$ is $(4\sqrt{2}\eta^2, 4\eta^2)$ -subExponential, and thus by our discussion in Appendix A $\xi_n^2 - \mathbb{E}\xi_n^2$ is also $(4\sqrt{2}\eta^2, 4\eta^2)$ -subGamma. Further, by Assumption 3.3, we have $\sum_{m=1}^n \mathbb{E}\xi_m^2 \leq n\alpha$. Recall that Lemma A.2 yields that, with probability at least $1 - \delta$,

$$\sum_{m=1}^n \xi_m^2 - \mathbb{E}\xi_m^2 \leq 6\eta^2 \sqrt{2n\ell_\delta(n)} + 10\eta^2 \ell_\delta(n).$$

Piecing everything together, we have that

$$\begin{aligned} \|\Xi_n\|_2^2 &= \sum_{m=1}^n \xi_m^2 \\ &= \sum_{m=1}^n \xi_m^2 - \mathbb{E}\xi_m^2 + \sum_{m=1}^n \mathbb{E}\xi_m^2 \\ &\leq 6\eta^2 \sqrt{2n\ell_\delta(n)} + 10\eta^2 \ell_\delta(n) + n\alpha. \end{aligned}$$

□

We now mention the finishing steps in proving Theorem 4.4, carrying over from where the proof in the paper ended. We see, with the previously addressed inequalities that, with probability at least $1 - O(A\delta)$, simultaneously for all $n \geq n_0$ and $a \in [A]$,

$$\begin{aligned} \left\| \widehat{\theta}_n(a) - \theta(a) \right\|_2^2 &\leq \frac{6L^2 U_n^2}{\sigma_r(\mathbf{Z}_n(a))^2} + \frac{2}{\sigma_r(\mathbf{Z}_n(a))^2} (T_1 + T_2) \\ &= \frac{6L^2 U_n^2}{\sigma_r(\mathbf{Z}_n(a))^2} + \frac{2}{\sigma_r(\mathbf{Z}_n(a))^2} \left[32\rho L^2 + 64\eta^2 \left(\log\left(\frac{A}{\delta}\right) + r \log\left(1 + \frac{\sigma_1(\mathbf{Z}_n)^2}{\rho}\right) \right) \right. \\ &\quad \left. + 16L^2 U_n^2 + 6\eta^2 \sqrt{2c_n(a)\ell_\delta(c_n(a))} + 10\eta^2 \ell_\delta(c_n(a)) + 6c_n(a)\alpha + \frac{108L^2 \sigma_1(\mathbf{Z}_n(a))^2 U_n^2}{\sigma_r(\mathbf{Z}_n(a))^2} + 18L^2 U_n^2 \right] \\ &= \frac{6L^2}{\widehat{\text{snr}}_n(a)^2} + \frac{L^2}{\widehat{\text{snr}}_n(a)^2} \left[68 + \frac{216\sigma_1(\mathbf{Z}_n(a))^2}{\sigma_r(\mathbf{Z}_n(a))^2} \right] + \frac{2}{\sigma_r(\mathbf{Z}_n(a))^2} \left[32\rho L^2 \right. \\ &\quad \left. + 64\eta^2 \left(\log\left(\frac{A}{\delta}\right) + r \log\left(1 + \frac{\sigma_1(\mathbf{Z}_n(a))^2}{\rho}\right) \right) + 6\eta^2 \sqrt{2c_n(a)\ell_\delta(c_n(a))} + 10\eta^2 \ell_\delta(c_n(a)) + 6c_n(a)\alpha \right] \end{aligned}$$

E Proof of Theorem 4.1

In this section, we prove Theorem 4.1. All that is required in proving this bound is simplifying the results of Theorem 4.4, given we allow ourselves slack to control the bound up to universal constants and poly-logarithmic factors.

Proof of Theorem 4.1. Recall that by Theorem 4.4, we have with probability at least $1 - O(A\delta)$, simultaneously for all $n \geq n_0$

$$\left\| \widehat{\theta}_n(a) - \theta(a) \right\|_2^2 \leq \frac{L^2}{\widehat{\text{snr}}_n(a)^2} [74 + 216\kappa(\mathbf{Z}_n(a))^2] + \frac{2\text{err}_n(a)}{\sigma_r(\mathbf{Z}_n(a))^2},$$

where in the above we define the “error” term $\text{err}_n(a)$ to be

$$\begin{aligned} \text{err}_n(a) &:= \underbrace{24\rho L^2 + 64\eta^2 \left(\log\left(\frac{A}{\delta}\right) + r \log\left(1 + \frac{\sigma_1(\mathbf{Z}_n(a))^2}{\rho}\right) \right)}_{T_1} \\ &\quad + \underbrace{6\eta^2 \sqrt{2c_n(a)\ell_\delta(c_n(a))} + 10\eta^2 \ell_\delta(c_n(a)) + 6c_n(a)\alpha}_{T_2}. \end{aligned}$$

First, by the second part of Lemma B.6 coupled with the assumption that $\text{snr}_n(a) \geq 2$ for all $n \geq n_0$ and $a \in [A]$, we have that $\text{snr}_n(a) = \Theta(\widehat{\text{snr}}_n(a))$. Further, by this same result, we have that $\kappa(\mathbf{Z}_n(a)) = \Theta(\kappa(\mathbf{X}_n(a)))$. From this, it is clear that $\frac{L^2}{\widehat{\text{snr}}_n(a)^2} [74 + 216\kappa(\mathbf{Z}_n(a))^2] = O\left(\frac{1}{\widehat{\text{snr}}_n(a)^2} \kappa(\mathbf{X}_n(a))^2\right)$. What remains is to show that $\frac{2\text{err}_n(a)}{\sigma_r(\widehat{\mathbf{Z}}_n(a))^2} = \tilde{O}\left(\frac{1}{\widehat{\text{snr}}_n(a)^2}\right)$. To this end, it suffices to show that $\text{err}_n(a) = \tilde{O}(n + d)$. But this is trivial, as it is clear that $T_1 = \tilde{O}(1)$ and $T_2 = \tilde{O}(c_n(a)) = \tilde{O}(n + d)$. Thus, we have proved the desired result. \square

F Equivalent formulations of ridge regression

We begin by discussing properties and equivalent formulations of ridge regression, as the estimate produced by (regularized) PCR, $\hat{\theta}_n(a)$, is precisely the ridge estimate of the unknown parameter $\theta(a)$ when restricted the subspace associated with the projection matrix $\hat{\mathbf{P}}_n$.

Fact F.1 (Ridge regression formulation). *Let \widehat{W}_n be the subspace associated with the projection matrix $\hat{\mathbf{P}}_n$. Then, $\hat{\theta}_n(a)$ satisfies*

$$\hat{\theta}_n(a) = \arg \min_{\theta \in \widehat{W}_n} \left\{ \|\mathbf{Z}_n(a)\theta - \mathbf{Y}_n(a)\|_2^2 + \frac{\rho}{2} \|\theta\|_2^2 \right\}.$$

That is, $\hat{\theta}_n(a)$ is the solution to ρ -regularized ridge regression when estimates are restricted to \widehat{W}_n .

Ridge regression may also be represented in the following, constrained optimization format.

Fact F.2 (Constrained formulation of ridge regression). *Let \widehat{W}_n be the subspace associated with the projection matrix $\hat{\mathbf{P}}_n$. Then, $\hat{\theta}_n(a)$ satisfies*

$$\hat{\theta}_n(a) = \arg \min_{\theta \in \widehat{W}_n: \|\theta\|_2 \leq R_\rho} \|\mathbf{Z}_n(a)\theta - \mathbf{Y}_n(a)\|_2^2,$$

where R_ρ is some constant only depending on ρ .

The larger ρ is, the smaller R_ρ must become. Since we know $\|\theta(a)\|_2 \leq L$, for all $a \in [A]$, throughout the main body and appendix of this paper, we assume that ρ is chosen to be sufficiently small such that our estimates $\hat{\theta}_n(a)$ satisfy $\|\hat{\theta}_n(a)\|_2 \leq L$, i.e. we select $\rho > 0$ satisfying $R_\rho \leq L$.

G Proofs for application to panel data

Lemma G.1 (Reformulation of average expected post-intervention outcome). *Under Assumption 5.1 and Assumption 5.2, the average expected post-intervention outcome of unit n under intervention a may be written as*

$$\mathbb{E}[\bar{Y}_{n,\text{post}}^{(a)}] = \frac{1}{T - T_0} \cdot \langle \theta(a), \mathbb{E}[Y_{n,\text{pre}}] \rangle,$$

for some slope vector $\theta(a) \in \mathbb{R}^{T_0}$.

Proof. Similar observations have been made in [45, 75]. For completeness, we include the proof here as well. From Assumption 5.1 and Definition 5.3,

$$\mathbb{E}[\bar{Y}_{n,\text{post}}^{(a)}] = \frac{1}{T - T_0} \cdot \left\langle \sum_{t=T_0+1}^T U_t^{(a)}, V_n \right\rangle.$$

Applying Assumption 5.2, we see that

$$\mathbb{E}[\bar{Y}_{n,\text{post}}^{(a)}] = \frac{1}{T - T_0} \cdot \left\langle \sum_{t=1}^{T_0} \theta(a)_t \cdot U_t^{(0)}, V_n \right\rangle$$

for some $\theta(a) = [\theta_1(a), \dots, \theta_{T_0}(a)]^T \in \mathbb{R}^{T_0}$. \square

The following lemma follows straightforwardly from translating the notation of Theorem 4.4 to the panel data setting. Note that in the panel data setting, $d = T_0$, $\eta = \sigma\sqrt{T - T_0}$, and $\alpha = \sigma^2(T - T_0)$. Additionally, Assumption 3.1 is satisfied with $\gamma = \sigma^2$.

Lemma G.2. *Let $\delta \in (0, 1)$ be an arbitrary confidence parameter. Let $\rho > 0$ be chosen to be sufficiently small, as detailed in Appendix F. Further, assume that there is some $n_0 \geq 1$ such that $\text{rank}(\mathbf{X}_{n_0}) = r$ and $\text{snr}_n \geq 2$ for all $n \geq n_0$. Then, with probability at least $1 - O(A\delta)$, simultaneously for all actions $a \in [A]$ and time steps $n \geq n_0$, we have*

$$\left\| \widehat{\theta}_n(a) - \theta(a) \right\|_2^2 \leq \frac{L^2}{\widehat{\text{snr}}_n(a)^2} [74 + 216\kappa(\mathbf{Z}_n(a))^2] + \frac{2(T - T_0)\text{err}_n(a)}{\sigma_r(\mathbf{Z}_n(a))^2},$$

where $\kappa(\mathbf{Z}_n(a)) := \frac{\sigma_1(\mathbf{Z}_n(a))}{\sigma_r(\mathbf{Z}_n(a))}$, $\|\theta(a)\|_2 \leq L$, and in the above we define the ‘‘error’’ term $\text{err}_n(a)$ to be

$$\begin{aligned} \text{err}_n(a) := & \frac{32\rho L^2}{T - T_0} + 64\sigma^2 \left(\log\left(\frac{A}{\delta}\right) + r \log\left(1 + \frac{\sigma_1(\mathbf{Z}_n(a))^2}{\rho}\right) \right) \\ & + 6\sigma^2 \sqrt{2c_n(a)\ell_\delta(c_n(a))} + 10\sigma^2 \ell_\delta(c_n(a)) + 6\sigma^2 c_n(a). \end{aligned}$$

Theorem G.3 (Prediction error of average expected post-intervention outcome). *Let $\delta \in (0, 1)$ be an arbitrary confidence parameter and $\rho > 0$ be chosen to be sufficiently small, as detailed in Appendix F. Further, assume that Assumptions 5.1 and 5.2 are satisfied, there is some $n_0 \geq 1$ such that $\text{rank}(\mathbf{X}_{n_0}) = r$, and $\text{snr}_n(a) \geq 2$ for all $n \geq n_0$. Then with probability at least $1 - O(A\delta)$, simultaneously for all interventions $a \in \{0, \dots, A - 1\}$*

$$\begin{aligned} |\widehat{\mathbb{E}}\bar{Y}_{n,\text{post}}^{(a)} - \mathbb{E}\bar{Y}_{n,\text{post}}^{(a)}| \leq & \frac{3\sqrt{T_0}}{\widehat{\text{snr}}_n(a)} \left(\frac{L(\sqrt{74} + 12\sqrt{6}\kappa(\mathbf{Z}_n(a)))}{(T - T_0) \cdot \widehat{\text{snr}}_n(a)} + \frac{\sqrt{2\text{err}_n(a)}}{\sqrt{T - T_0} \cdot \sigma_r(\mathbf{Z}_n(a))} \right) \\ & + \frac{2L\sqrt{24T_0}}{(T - T_0) \cdot \widehat{\text{snr}}_n(a)} + \frac{12L\kappa(\mathbf{Z}_n(a))\sqrt{3T_0}}{(T - T_0) \cdot \widehat{\text{snr}}_n(a)} + \frac{2\sqrt{\text{err}_n(a)}}{\sqrt{T - T_0} \cdot \sigma_r(\mathbf{Z}_n(a))} \\ & + \frac{L\sigma\sqrt{\log(A/\delta)}}{\sqrt{T - T_0}} + \frac{L\sigma\sqrt{74\log(A/\delta)}}{\widehat{\text{snr}}_n(a)\sqrt{T - T_0}} + \frac{12\sigma\kappa(\mathbf{Z}_n(a))\sqrt{6\log(A/\delta)}}{\widehat{\text{snr}}_n(a)\sqrt{T - T_0}} \\ & + \frac{\sigma\sqrt{2\text{err}_n(a)\log(A/\delta)}}{\sigma_r(\mathbf{Z}_n(a))}, \end{aligned}$$

where $\widehat{\mathbb{E}}\bar{Y}_{n,\text{post}}^{(a)} := \frac{1}{T - T_0} \cdot \langle \widehat{\theta}_n(a), Y_{n,\text{pre}} \rangle$ is the estimated average post-intervention outcome for unit n under intervention a .

Proof.

$$\begin{aligned} \widehat{\mathbb{E}}\bar{Y}_{n,\text{post}}^{(a)} - \mathbb{E}\bar{Y}_{n,\text{post}}^{(a)} & := \frac{1}{T - T_0} (\langle \widehat{\theta}_n(a), Y_{n,\text{pre}} \rangle - \langle \theta(a), Y_{n,\text{pre}} \rangle) \\ & = \frac{1}{T - T_0} \underbrace{\langle \widehat{\theta}_n(a) - \theta(a), \mathbb{E}Y_{n,\text{pre}} \rangle}_{T_1} + \underbrace{\langle \theta(a), \epsilon_{n,\text{pre}} \rangle}_{T_2} \\ & \quad + \underbrace{\langle \widehat{\theta}_n(a) - \theta(a), \epsilon_{n,\text{pre}} \rangle}_{T_3} \end{aligned}$$

We begin by bounding T_1 . By assumption we have that $\mathbb{E}Y_{n,\text{pre}} \in \text{span}(\mathbb{E}Y_{1,\text{pre}}, \dots, \mathbb{E}Y_{n-1,\text{pre}})$ for all $n \geq n_0$. Therefore,

$$\begin{aligned} \langle \widehat{\theta}_n(a) - \theta(a), \mathbb{E}Y_{n,\text{pre}} \rangle & = \langle \widehat{\theta}_n(a) - \theta(a), \mathbb{E}Y_{n,\text{pre}} \mathbf{P} \rangle \\ & = \langle \mathbf{P}\widehat{\theta}_n(a) - \theta(a), \mathbb{E}Y_{n,\text{pre}} \rangle \\ & \leq \|\mathbb{E}Y_{n,\text{pre}}\|_2 \|\mathbf{P} - \widehat{\mathbf{P}}_n(a) + \widehat{\mathbf{P}}_n(a)\|_2 \|\widehat{\theta}_n(a) - \theta(a)\|_2 \\ & \leq \underbrace{\sqrt{T_0} \cdot \|\mathbf{P} - \widehat{\mathbf{P}}_n(a)\|_2}_{T_{1.1}} \|\widehat{\theta}_n(a) - \theta(a)\|_2 + \underbrace{\sqrt{T_0} \cdot \|\widehat{\mathbf{P}}_n(a)\|_2}_{T_{1.2}} \|\widehat{\theta}_n(a) - \theta(a)\|_2 \end{aligned}$$

By Lemma B.4 and the second part of Lemma B.6, term $T_{1.1}$ may be upper-bounded as

$$\begin{aligned}
T_{1.1} &\leq \sqrt{T_0} \cdot \|\mathbf{P} - \widehat{\mathbf{P}}_n(a)\|_{op} \cdot \|\widehat{\theta}_n(a) - \theta(a)\|_2 \\
&\leq \sqrt{T_0} \cdot \frac{\sqrt{4\beta (3\sqrt{n\ell_{\delta/2\mathcal{N}}(n)} + 5\ell_{\delta/2\mathcal{N}}(n)) + 4n\gamma}}{\sigma_r(\mathbf{X}_n(a))} \cdot \|\widehat{\theta}_n(a) - \theta(a)\|_2 \\
&\leq \sqrt{T_0} \cdot \frac{\sqrt{4\beta (3\sqrt{n\ell_{\delta/2\mathcal{N}}(n)} + 5\ell_{\delta/2\mathcal{N}}(n)) + 4n\gamma}}{\frac{2}{3}\sigma_r(\mathbf{Z}_n(a))} \cdot \|\widehat{\theta}_n(a) - \theta(a)\|_2 \\
&= \frac{3\sqrt{T_0}}{\widehat{\text{snr}}_n(a)} \cdot \|\widehat{\theta}_n(a) - \theta(a)\|_2
\end{aligned}$$

Applying Lemma G.2, we see that

$$\begin{aligned}
\frac{1}{T - T_0} T_{1.1} &\leq \frac{3\sqrt{T_0}}{(T - T_0) \cdot \widehat{\text{snr}}_n(a)} \left(\frac{L^2}{(\widehat{\text{snr}}_n(a))^2} [74 + 216\kappa(\mathbf{Z}_n(a))^2] + \frac{2(T - T_0)\text{err}_n(a)}{\sigma_r(\mathbf{Z}_n(a))^2} \right)^{1/2} \\
&\leq \frac{3\sqrt{T_0}}{(T - T_0) \cdot \widehat{\text{snr}}_n(a)} \left(\frac{L}{\widehat{\text{snr}}_n(a)} \sqrt{74 + 216\kappa(\mathbf{Z}_n(a))^2} + \frac{\sqrt{2(T - T_0)\text{err}_n(a)}}{\sigma_r(\mathbf{Z}_n(a))} \right) \\
&= \frac{3\sqrt{T_0}}{\widehat{\text{snr}}_n(a)} \left(\frac{L(\sqrt{74} + 12\sqrt{6}\kappa(\mathbf{Z}_n(a)))}{(T - T_0) \cdot \widehat{\text{snr}}_n(a)} + \frac{\sqrt{2\text{err}_n(a)}}{\sqrt{T - T_0} \cdot \sigma_r(\mathbf{Z}_n(a))} \right).
\end{aligned}$$

Turning our attention to $T_{1.2}$ and using a line of reasoning nearly identical to equations (1), (2), (3) in the proof of Theorem 4.4, we get that with probability at least $1 - \mathcal{O}(A\delta)$,

$$\begin{aligned}
\frac{1}{T - T_0} T_{1.2} &\leq \frac{2\sqrt{T_0}}{(T - T_0) \cdot \sigma_r(\mathbf{Z}_n(a))} \left(\left\| \widehat{\mathbf{Z}}_n(a)\widehat{\theta}_n(a) - \mathbf{X}_n(a)\theta(a) \right\|_2^2 + \left\| \mathbf{X}_n(a)\theta(a) - \widehat{\mathbf{Z}}_n(a)\theta(a) \right\|_2^2 \right)^{1/2} \\
&\leq \frac{2\sqrt{T_0}}{(T - T_0) \cdot \sigma_r(\mathbf{Z}_n(a))} \left(32\rho L^2 + 64\sigma^2(T - T_0) \left(\log\left(\frac{A}{\delta}\right) + r \log\left(1 + \frac{\sigma_1(\mathbf{Z}_n(a))^2}{\rho}\right) \right) \right) \\
&\quad + 16L^2U_n^2 + 6\sigma^2(T - T_0)\sqrt{2c_n(a)\ell_\delta(c_n(a))} + 10\sigma^2(T - T_0)\ell_\delta(c_n(a)) \\
&\quad + 6\sigma^2(T - T_0)c_n(a) + \frac{108L^2\sigma_1(\mathbf{Z}_n(a))^2U_n^2 + 18L^2U_n^2}{\sigma_r(\mathbf{Z}_n(a))^2} \Big)^{1/2} \\
&= \frac{2\sqrt{T_0}}{(T - T_0) \cdot \sigma_r(\mathbf{Z}_n(a))} \left(34L^2U_n^2 + \frac{108L^2\sigma_1(\mathbf{Z}_n(a))^2U_n^2}{\sigma_r(\mathbf{Z}_n(a))^2} + (T - T_0)\text{err}_n(a) \right)^{1/2} \\
&\leq \frac{2LU_n\sqrt{24T_0}}{(T - T_0) \cdot \sigma_r(\mathbf{Z}_n(a))} + \frac{12L\kappa(\mathbf{Z}_n(a))U_n\sqrt{3T_0}}{(T - T_0) \cdot \sigma_r(\mathbf{Z}_n(a))} + \frac{2\sqrt{\text{err}_n(a)}}{\sqrt{T - T_0} \cdot \sigma_r(\mathbf{Z}_n(a))} \\
&= \frac{2L\sqrt{24T_0}}{(T - T_0) \cdot \widehat{\text{snr}}_n(a)} + \frac{12L\kappa(\mathbf{Z}_n(a))\sqrt{3T_0}}{(T - T_0) \cdot \widehat{\text{snr}}_n(a)} + \frac{2\sqrt{\text{err}_n(a)}}{\sqrt{T - T_0} \cdot \sigma_r(\mathbf{Z}_n(a))}
\end{aligned}$$

Putting our bounds for $T_{1.1}$ and $T_{1.2}$ together, we get that

$$\begin{aligned}
\frac{\langle \widehat{\theta}_n(a) - \theta(a), \mathbb{E}Y_{n,\text{pre}} \rangle}{T - T_0} &\leq \frac{3\sqrt{T_0}}{\widehat{\text{snr}}_n(a)} \left(\frac{L(\sqrt{74} + 12\sqrt{6}\kappa(\mathbf{Z}_n(a)))}{(T - T_0) \cdot \widehat{\text{snr}}_n(a)} + \frac{\sqrt{2\text{err}_n(a)}}{\sqrt{T - T_0} \cdot \sigma_r(\mathbf{Z}_n(a))} \right) \\
&\quad + \frac{2L\sqrt{24T_0}}{(T - T_0) \cdot \widehat{\text{snr}}_n(a)} + \frac{12L\kappa(\mathbf{Z}_n(a))\sqrt{3T_0}}{(T - T_0) \cdot \widehat{\text{snr}}_n(a)} + \frac{2\sqrt{\text{err}_n(a)}}{\sqrt{T - T_0} \cdot \sigma_r(\mathbf{Z}_n(a))}
\end{aligned}$$

Next we bound T_2 . Note that $\langle \theta(a), \epsilon_{n,\text{pre}} \rangle$ is a $\|\theta(a)\|_2\sqrt{T - T_0}\sigma$ -subGaussian random variable. Therefore via a Hoeffding bound, simultaneously for all actions $a \in [A]$, with probability at least $1 - \mathcal{O}(A\delta)$,

$$\frac{\langle \theta(a), \epsilon_{n,\text{pre}} \rangle}{T - T_0} \leq L\sigma\sqrt{\frac{\log(A/\delta)}{T - T_0}}$$

Similarly for T_3 , $\langle \widehat{\theta}_n(a) - \theta(a), \epsilon_{n,pre} \rangle$ is a $\|\widehat{\theta}_n(a) - \theta(a)\|_2 \sqrt{T - T_0} \sigma$ -subGaussian random variable which, after applying a Hoeffding bound and our bound on $\|\widehat{\theta}_n(a) - \theta(a)\|_2$, becomes

$$\begin{aligned} \frac{\langle \widehat{\theta}_n(a) - \theta(a), \epsilon_{n,pre} \rangle}{T - T_0} &\leq \sigma \sqrt{\frac{\log(A/\delta)}{T - T_0}} \left(\frac{L^2}{\widehat{\text{snr}}_n(a)^2} [74 + 216\kappa(\mathbf{Z}_n(a))^2] + \frac{2(T - T_0)\text{err}_n(a)}{\sigma_r(\mathbf{Z}_n(a))^2} \right)^{1/2} \\ &\leq \sigma \sqrt{\frac{\log(A/\delta)}{T - T_0}} \left(\frac{\sqrt{74}L}{\widehat{\text{snr}}_n(a)} + \frac{12\sqrt{6}\kappa(\mathbf{Z}_n(a))}{\widehat{\text{snr}}_n(a)} + \frac{\sqrt{2(T - T_0)\text{err}_n(a)}}{\sigma_r(\mathbf{Z}_n(a))} \right) \\ &= \frac{L\sigma\sqrt{74\log(A/\delta)}}{\widehat{\text{snr}}_n(a)\sqrt{T - T_0}} + \frac{12\sigma\kappa(\mathbf{Z}_n(a))\sqrt{6\log(A/\delta)}}{\widehat{\text{snr}}_n(a)\sqrt{T - T_0}} + \frac{\sigma\sqrt{2\text{err}_n(a)\log(A/\delta)}}{\sigma_r(\mathbf{Z}_n(a))} \end{aligned}$$

Putting everything together, we see that with probability at least $1 - O(A\delta)$,

$$\begin{aligned} |\widehat{\mathbb{E}}\bar{Y}_{n,post}^{(a)} - \mathbb{E}\bar{Y}_{n,post}^{(a)}| &\leq \frac{3\sqrt{T_0}}{\widehat{\text{snr}}_n(a)} \left(\frac{L(\sqrt{74} + 12\sqrt{6}\kappa(\mathbf{Z}_n(a)))}{(T - T_0) \cdot \widehat{\text{snr}}_n(a)} + \frac{\sqrt{2\text{err}_n(a)}}{\sqrt{T - T_0} \cdot \sigma_r(\mathbf{Z}_n(a))} \right) \\ &\quad + \frac{2L\sqrt{24T_0}}{(T - T_0) \cdot \widehat{\text{snr}}_n(a)} + \frac{12L\kappa(\mathbf{Z}_n(a))\sqrt{3T_0}}{(T - T_0) \cdot \widehat{\text{snr}}_n(a)} + \frac{2\sqrt{\text{err}_n(a)}}{\sqrt{T - T_0} \cdot \sigma_r(\mathbf{Z}_n(a))} \\ &\quad + \frac{L\sigma\sqrt{\log(A/\delta)}}{\sqrt{T - T_0}} + \frac{L\sigma\sqrt{74\log(A/\delta)}}{\widehat{\text{snr}}_n(a)\sqrt{T - T_0}} + \frac{12\sigma\kappa(\mathbf{Z}_n(a))\sqrt{6\log(A/\delta)}}{\widehat{\text{snr}}_n(a)\sqrt{T - T_0}} \\ &\quad + \frac{\sigma\sqrt{2\text{err}_n(a)\log(A/\delta)}}{\sigma_r(\mathbf{Z}_n(a))} \end{aligned}$$

Applying Assumption 4.2, the expression simplifies to

$$\begin{aligned} |\widehat{\mathbb{E}}\bar{Y}_{n,post}^{(a)} - \mathbb{E}\bar{Y}_{n,post}^{(a)}| &= \tilde{O} \left(\frac{r\sqrt{T_0}}{\sqrt{T_0} \wedge n} \left(\frac{Lr}{(T - T_0)\sqrt{T_0} \wedge n} + \frac{r}{\sqrt{(T - T_0)(T_0 \wedge n)}} \right) \right. \\ &\quad + \frac{Lr\sqrt{T_0}}{(T - T_0)\sqrt{T_0} \wedge n} + \frac{Lr\sqrt{T_0}}{(T - T_0)\sqrt{T_0} \wedge n} + \frac{r}{\sqrt{(T - T_0)(T_0 \wedge n)}} \\ &\quad \left. + \frac{L}{\sqrt{T - T_0}} + \frac{Lr}{\sqrt{(T - T_0)(T_0 \wedge n)}} + \frac{r}{\sqrt{(T - T_0)(T_0 \wedge n)}} + \frac{r}{\sqrt{T_0 \wedge n}} \right). \end{aligned}$$

Under the assumptions that $T_0 \leq \frac{1}{2}T$ and $r \leq \sqrt{T_0 \wedge n}$, we get that

$$|\widehat{\mathbb{E}}\bar{Y}_{n,post}^{(a)} - \mathbb{E}\bar{Y}_{n,post}^{(a)}| = \tilde{O} \left(\frac{L}{\sqrt{T - T_0}} + \frac{r}{\sqrt{T_0 \wedge n}} + \frac{r(L \vee 1)}{\sqrt{(T - T_0)(T_0 \wedge n)}} \right)$$

□