

A Sharp Characterization of Linear Estimators for Offline Policy Evaluation

Juan C. Perdomo
jcperdomo@berkeley.edu
University of California, Berkeley

Peter Bartlett
peter@berkeley.edu
University of California, Berkeley

Akshay Krishnamurthy
akshaykr@microsoft.com
Microsoft Research

Sham Kakade
sham@seas.harvard.edu
Harvard University

March 8, 2022

Abstract

Offline policy evaluation is a fundamental statistical problem in reinforcement learning that involves estimating the value function of some decision-making policy given data collected by a potentially different policy. In order to tackle problems with complex, high-dimensional observations, there has been significant interest from theoreticians and practitioners alike in understanding the possibility of *function approximation* in reinforcement learning. Despite significant study, a sharp characterization of when we might expect offline policy evaluation to be tractable, even in the simplest setting of *linear* function approximation, has so far remained elusive, with a surprising number of strong negative results recently appearing in the literature.

In this work, we identify simple control-theoretic and linear-algebraic conditions that are necessary and sufficient for classical methods, in particular Fitted Q-iteration (FQI) and least squares temporal difference learning (LSTD), to succeed at offline policy evaluation. Using this characterization, we establish a precise hierarchy of regimes under which these estimators succeed. We prove that LSTD works under strictly weaker conditions than FQI. Furthermore, we establish that if a problem is not solvable via LSTD, then it cannot be solved by a broad class of linear estimators, even in the limit of infinite data. Taken together, our results provide a complete picture of the behavior of linear estimators for offline policy evaluation (OPE), unify previously disparate analyses of canonical algorithms, and provide significantly sharper notions of the underlying statistical complexity of OPE.

1 Introduction

A central component of a practical sequential decision making system is its ability to cope with high-dimensional and complex data sources. While feature engineering or discretization techniques can in principle be used to address the challenges associated with complex data, these approaches require significant domain expertise and suffer from a curse-of-dimensionality phenomenon that limit their practical relevance. Instead, the use of more general *function approximation methods* for reinforcement learning (RL) promises to avoid these drawbacks. Consequently, understanding these methods has long been a topic of interest to theoreticians and practitioners alike.

While the use of nonlinear methods is by now common in the empirical reinforcement learning literature, the much simpler linear function approximation setting remains somewhat poorly understood theoretically, despite decades of study. Indeed, recently there has been a surge of research effort focusing on necessary and sufficient conditions for reinforcement learning with linear function approximation, including the first provably efficient algorithms for online exploration [Yang and Wang, 2020, Jin et al., 2020] and a number of

surprising statistical lower bounds that hold even under strong assumptions [Wang et al., 2021c, Weisz et al., 2021a,b]. This line of work represents substantial progress, yet we still lack a clear picture as to precisely when and why RL with linear function approximation is tractable.

As a step towards providing this clarity, in this paper we focus on the simpler *offline policy evaluation* problem (OPE), under the assumption that the action-value function is *linearly realizable* in some known features. Here, rather than interacting with an environment to maximize reward as in the standard RL formulation, the goal is to estimate the performance of a given decision-making policy by leveraging an observational dataset collected by a potentially different policy. OPE is perhaps the simplest, non-trivial setting in which to study function approximation in RL. It is also practically relevant in its own right: both OPE and the closely-related offline policy optimization problem represent a promising avenue toward applying RL in safety-critical domains where active exploration is infeasible. Moreover, the principles developed for OPE are routinely used in online RL algorithms.

Fitted Q-iteration (FQI) [Ernst et al., 2005, Riedmiller, 2005] and least squares temporal difference learning (LSTD) [Bradtke and Barto, 1996, Boyan, 1999, Nedić and Bertsekas, 2003] are canonical algorithms for offline policy evaluation with function approximation. These simple, moment-based methods are some of the most popular approaches in practice and have served as inspiration for recent empirical breakthroughs in RL [Mnih et al., 2015]. They have also been the subject of intense theoretical investigation, with early results on convergence and instability described by Bertsekas and Tsitsiklis [1995], Tsitsiklis and Van Roy [1996] as well as several more recent results [Antos et al., 2008, Chen and Jiang, 2019, Lazaric et al., 2012]. Nevertheless, a sharp finite-sample characterization of the behavior of FQI and LSTD, even in the linear realizability setting, remains undeveloped.

In this paper, we identify necessary and sufficient conditions for FQI and LSTD to succeed at offline policy evaluation under linear realizability. In doing so, we establish a precise hierarchy of conditions under which these methods work; in particular, we prove that LSTD succeeds under strictly weaker assumptions than FQI. Moreover, if an offline policy evaluation problem is not solvable via LSTD, then it cannot be solved by any linear, moment-based method (see Definition 4.1) even in the limit of infinite data. Our characterization draws upon ideas from the theory of Lyapunov stability and provides a new, unifying perspective on the statistical complexity of offline policy evaluation. In particular, we show how traditional quantities, such as the “effective horizon”, fail to capture the true complexity of the problem (Sections 3.1 and 4.1) and propose instance-dependent measures which are significantly sharper. Furthermore, our results unify previously disparate analyses for FQI and LSTD as our conditions are implied by prior assumptions (Sections 3.2 and 4.2). Taken together, our results provide a complete picture of the possibilities and limitations of linear estimators for offline policy evaluation under linear realizability.

1.1 Linear estimators & the offline policy evaluation problem

Let $\mathcal{M} := (\mathcal{S}, \mathcal{A}, P, R, \gamma)$ denote an infinite horizon, γ -discounted MDP where \mathcal{S} is the set of states, \mathcal{A} is the set of actions, $R : \mathcal{S} \times \mathcal{A} \rightarrow \Delta([-1, 1])$ is the random reward function, and $P : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$ is the transition operator, which defines a distribution over states for every pair (s, a) . The action-value function Q^π captures the expected total reward achieved by a randomized policy $\pi : \mathcal{S} \rightarrow \Delta(\mathcal{A})$ from an initial state-action pair (s, a) when the trajectory is generated such that for each time step h , $a_h \sim \pi(s_h)$ and $s_{h+1} \sim P(\cdot \mid s_h, a_h)$.

$$Q^\pi(s, a) := \mathbb{E} \left[\sum_{h=0}^{\infty} \gamma^h r(s_h, a_h) \mid (s_0, a_0) = (s, a), \pi \right]. \quad (1.1)$$

In the offline policy evaluation problem, we are given a policy π and a dataset $\{(s_i, a_i, r_i(s_i, a_i), s'_i, a'_i)\}_{i=1}^n$ of observed transitions and rewards, where the initial pair (s_i, a_i) is sampled from some *arbitrary* distribution \mathcal{D} , $r_i(s_i, a_i) \sim R(s_i, a_i)$, the next state is sampled from the transition operator $s'_i \sim P(\cdot \mid s_i, a_i)$, and the next action $a'_i \sim \pi(s'_i)$ is sampled according to π .¹ Our goal is to return an estimate \hat{Q}^π of Q^π . For concreteness, we measure performance via $\mathbb{E}_{(s,a) \sim \mathcal{D}} |\hat{Q}^\pi(s, a) - Q^\pi(s, a)|$ and we ask that this quantity is vanishingly small

¹We “augment” the dataset to include the next state action $a' \sim \pi(s')$ purely for notational convenience.

with high probability over the draw of the dataset. For simplicity, we assume that samples are drawn i.i.d. via the procedure described above.²

As we would like to develop methods that scale to settings where the cardinalities of the sets \mathcal{S} and \mathcal{A} are large or infinite, our focus is on understanding policy evaluation using linear function approximation, as per the following definition:

Assumption 1 (Linear Realizability). Q^π is linearly realizable³ in a known feature map $\phi : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^d$ if there exists a vector $\theta_\gamma^* \in \mathbb{R}^d$ such that for all $(s, a) \in \mathcal{S} \times \mathcal{A}$, $Q^\pi(s, a) = \phi(s, a)^\top \theta_\gamma^*$.

Fitted Q-iteration. As mentioned previously, fitted Q-iteration is one of the most popular algorithms for policy evaluation in practice and can in principle work with any function approximation method. In the linear case, given a dataset $\{(s_i, a_i, r_i(s_i, a_i), s'_i, a'_i)\}_{i=1}^n$ and an initial vector θ_0 , FQI iteratively solves least squares regression problems of the form

$$\hat{\theta}_{t+1} \in \arg \min_{\theta} \sum_{i=1}^n \left(\phi(s_i, a_i)^\top \theta - r(s_i, a_i) - \gamma \phi(s'_i, a'_i)^\top \hat{\theta}_t \right)^2, \quad (1.2)$$

for some number of rounds T and returns the estimator $\hat{Q}^\pi(s, a) := \phi(s, a)^\top \hat{\theta}_T$.

Least squares temporal difference learning. In the linear function approximation setting, the vector θ_γ^* which realizes Q^π in the feature mapping ϕ satisfies the fixed point equation,

$$\Sigma_{\text{cov}} \theta_\gamma^* = \gamma \Sigma_{\text{cr}} \theta_\gamma^* + \theta_{\phi, r}. \quad (1.3)$$

Here, Σ_{cov} is the offline feature covariance matrix, Σ_{cr} is the cross-covariance matrix between time-adjacent features, and $\theta_{\phi, r}$ is the mean feature-reward vector. (see Eqs. (1.5) and (2.3) for a formal definition). LSTD tries to approximate θ_γ^* by computing the plug-in estimate to the closed form solution to the equation above,

$$\hat{\theta}_{\text{LS}} := (I - \gamma \hat{\Sigma}_{\text{cov}}^{-1} \hat{\Sigma}_{\text{cr}})^\dagger \hat{\Sigma}_{\text{cov}}^{-1} \hat{\theta}_{\phi, r} = (\hat{\Sigma}_{\text{cov}} - \gamma \hat{\Sigma}_{\text{cr}})^\dagger \hat{\theta}_{\phi, r}. \quad (1.4)$$

and returns $\hat{Q}^\pi(s, a) := \phi(s, a)^\top \hat{\theta}_{\text{LS}}$ [Bradtke and Barto, 1996]. We focus on the unregularized variant of both of these algorithms. However, similar insights apply to the regularized cases (see Appendix A.7).

1.2 Our contributions

The main result of our work is that we identify simple linear algebraic conditions which exactly characterize when linear estimators will succeed at offline policy evaluation under linear realizability of Q^π . Under these conditions, which we introduce below, we establish upper bounds on the sample complexity necessary to perform policy evaluation which scale with: (1) for FQI, the operator norm of the solution to a particular discrete-time Lyapunov equation, and (2) for LSTD, the minimum singular value of an instance-dependent matrix. In both cases, we illustrate how our results unify previously disparate analyses of these algorithms, and demonstrate how our new instance-dependent quantities provide sharper notions of the statistical complexity of OPE when compared to bounds that explicitly depend on traditional parameters such as the “effective horizon”, i.e., $1/(1 - \gamma)$.

Our conditions can be introduced rather succinctly. For FQI, the key definitions and assumption are:

$$\Sigma_{\text{cov}} := \mathbb{E}_{(s, a) \sim \mathcal{D}} [\phi(s, a) \phi(s, a)^\top], \quad \Sigma_{\text{cr}} := \mathbb{E}_{(s, a) \sim \mathcal{D}} [\phi(s, a) \phi(s', a')^\top]. \quad (1.5)$$

$s' \sim P(\cdot | s, a), a' \sim \pi(s')$

²Extensions to Markovian data are fairly well-understood, see e.g., Mou et al. [2021], Nagaraj et al. [2020].

³Note that realizability of Q^π does not imply that the rewards are linearly realizable.

Assumption 2 (Stability). The matrix Σ_{cov} is full rank and $\rho(\gamma\Sigma_{\text{cov}}^{-1}\Sigma_{\text{cr}}) < 1$.

Here, Σ_{cov} is the offline state-action covariance, Σ_{cr} is the cross-covariance, $\gamma\Sigma_{\text{cov}}^{-1/2}\Sigma_{\text{cr}}\Sigma_{\text{cov}}^{-1/2}$ is the *whitened cross-covariance*,⁴ $\rho(A) = \max_i |\lambda_i(A)|$ is the *spectral radius* of the matrix A . The assumption that Σ_{cov} is full rank is not fundamental and is included primarily to simplify the presentation.⁵ If [Assumption 2](#) holds, we let P_γ be the unique solution (over X) to the Lyapunov equation,

$$X = (\gamma\Sigma_{\text{cov}}^{-1/2}\Sigma_{\text{cr}}\Sigma_{\text{cov}}^{-1/2})^\top X (\gamma\Sigma_{\text{cov}}^{-1/2}\Sigma_{\text{cr}}\Sigma_{\text{cov}}^{-1/2}) + I.$$

Our first main result is that, under stability, FQI satisfies the following error guarantee:

Theorem 1 (Informal). *Let $\hat{Q}^\pi(s, a) = \phi(s, a)^\top \hat{\theta}_T$, where $\hat{\theta}_T$ is the T -step FQI solution. Under [Assumptions 1 and 2](#), as well as standard regularity assumptions for linear regression, for n large enough,*

$$\mathbb{E}_{\mathcal{D}} |Q^\pi(s, a) - \hat{Q}^\pi(s, a)| \lesssim \text{cond}(P_\gamma) \|P_\gamma\|_{\text{op}}^2 \sqrt{\frac{d \log(1/\delta)}{n}} + \mathcal{O}(\exp(-T)),$$

with probability $1 - \delta$. Here, $\text{cond}(\cdot)$ and $\|\cdot\|_{\text{op}}$ denote the condition number and operator norm.

For the sake of clarity, we have suppressed dependence on universal constants and other quantities which arise from standard analysis of linear regression in the informal statement of the upper bound. Since $P_\gamma \succeq I$, $\text{cond}(P_\gamma)$ can always be crudely upper bounded by the operator norm, so that the primary factor, beyond the standard $\sqrt{d/n}$ term for linear regression, is the dependence on $\|P_\gamma\|_{\text{op}}$. We show in [Section 3.2](#) that, for settings where FQI was previously shown to succeed (e.g., under low distribution shift or Bellman completeness [[Wang et al., 2021a](#)]), stability always holds and $\|P_\gamma\|_{\text{op}}$ is never much larger than $1/(1 - \gamma)$, demonstrating how our bound recovers and unifies prior results. However, we also find that, in general, this quantity provides a much *sharper* notion of complexity for OPE. Indeed, there are simple instances where $\|P_\gamma\|_{\text{op}}$ is $\mathcal{O}(1)$ for all $\gamma \in (0, 1)$, but of course, $1/(1 - \gamma)$ can be arbitrarily large.

The key insight behind this result is that, in the linear setting, FQI can be written as a power series in the empirical versions of the second moment matrices described in [Eq. \(1.5\)](#). More precisely, $\hat{\theta}_T = \sum_{k=0}^T (\gamma \hat{\Sigma}_{\text{cov}}^{-1} \hat{\Sigma}_{\text{cr}})^k \hat{\Sigma}_{\text{cov}}^{-1} \hat{\theta}_{\phi, r}$ where $\hat{\theta}_{\phi, r}$ is obtained by solving a regression for the rewards. The behavior of the algorithm is governed by the growth of these matrix powers. Using ideas from Lyapunov theory, we show that if stability holds, then these decay at a geometric rate governed by $\|P_\gamma\|_{\text{op}}$ and FQI succeeds. On the other hand, if the spectral radius is greater than one, then these matrix powers grow exponentially, and FQI will drastically amplify any estimation errors. This leads to the *necessity* of stability for FQI:

Proposition 3.4 (Informal). *If $\rho(\gamma\Sigma_{\text{cov}}^{-1}\Sigma_{\text{cr}}) > 1$, the variance of the FQI solution grows exponentially with the number of regression rounds T .*

Turning to LSTD, while the solution is defined in terms of similar moment quantities to those relevant for FQI, it solves for θ_γ^* in a more direct manner and hence its behavior is somewhat different. We prove that LSTD succeeds if the following condition holds:

Assumption 3 (Invertibility). The matrices Σ_{cov} and $I - \gamma\Sigma_{\text{cov}}^{-1}\Sigma_{\text{cr}}$ are both full rank.

Our main result for LSTD is that under invertibility, θ_γ^* is identifiable via LSTD as per the following informal theorem statement:

Theorem 2 (Informal). *Let $\hat{Q}^\pi(s, a) = \phi(s, a)^\top \hat{\theta}_{\text{LS}}$, where $\hat{\theta}_{\text{LS}}$ is the LSTD solution. Under [Assumptions 1 and 3](#), as well as standard regularity assumptions for linear regression, if n is large enough,*

$$\mathbb{E}_{\mathcal{D}} |Q^\pi(s, a) - \hat{Q}^\pi(s, a)| \lesssim \frac{1}{\sigma_{\min}(I - \gamma\Sigma_{\text{cov}}^{-1/2}\Sigma_{\text{cr}}\Sigma_{\text{cov}}^{-1/2})^2} \sqrt{\frac{d \log(1/\delta)}{n}}$$

with probability $1 - \delta$. Here, σ_{\min} denotes the minimum singular value of a matrix.

⁴For any matrix A and invertible matrix L , the eigenvalues of A and $L^{-1}AL$ are identical. Therefore, one could equivalently state [Assumptions 2](#) and [3](#) in terms of $\gamma\Sigma_{\text{cov}}^{-1/2}\Sigma_{\text{cr}}\Sigma_{\text{cov}}^{-1/2}$.

⁵For example, the results carry over if all features $\phi(s, a)$ lie in a low dimensional subspace.

This result follows somewhat directly from an analysis of linear regression, since by inspecting [Eq. \(1.4\)](#), we can see that $\hat{\theta}_{LS}$ is the solution to an ordinary least squares problem. The dependence on the minimum singular value arises from translating from the norm induced by the design matrix of the regression problem to the norm under which we measure the errors in \hat{Q}^π .

Perhaps surprisingly, we will see that invertibility is strictly weaker than stability ([Assumption 2](#)), which highlights a fundamental distinction between these two methods. This comparison also reveals that stability cannot be a necessary condition in any algorithm-independent sense, since LSTD can succeed without stability. However, complementing [Theorem 2](#), we prove that invertibility is necessary for a large class of natural estimators, specifically those that rely on low-order moments of the features and the regression function between the features and the rewards (this includes FQI and LSTD). The following lower bound shows that the value function is *unidentifiable* by these linear estimators if invertibility does not hold.

Theorem 3 (Informal). *Even in the limit of infinite data, any OPE problem for which invertibility does not hold cannot be solved by a broad class of linear estimators, including FQI and LSTD.*

Together with our previous results, this result completes our analysis of linear estimators for offline policy evaluation under linear realizability.

1.3 Related work

RL with function approximation. Analyses of function approximation in reinforcement learning can be traced to the seminal papers of [Bellman and Dreyfus \[1959\]](#), [Bellman \[1961\]](#), as well as [Reetz \[1977\]](#) and [Whitt \[1978\]](#). [Schweitzer and Seidmann \[1985\]](#) were one of the first to consider approximating value functions using linear combinations of some known set of features. More recently, a number of modeling assumptions—typically involving strong representational conditions on both the MDP and the features—that enable statistically efficient online RL with linear function approximation have been proposed, along with corresponding algorithms [[Zanette et al., 2020](#), [Yang and Wang, 2020](#), [Jin et al., 2020](#)].

FQI. Introduced by [Ernst et al. \[2005\]](#) and extended by [Riedmiller \[2005\]](#), fitted Q-iteration has been analyzed several times in the context of offline policy evaluation. Building off previous studies of approximate methods in dynamic programming [[Antos et al., 2008](#), [Munos, 2007](#), [Gordon, 1999](#)], [Chen and Jiang \[2019\]](#) establish sample complexity upper bounds for FQI assuming that the corresponding distributions and MDP satisfy concentrability [[Munos, 2003](#)] and Bellman completeness [[Szepesvári and Munos, 2005](#)]. Both of these conditions are significantly stronger than mere *realizability* of value functions. More recent work by [Wang et al. \[2021a,b\]](#) adapts these results to the linear setting and additionally shows that a “low distribution shift” condition suffices for linear FQI.

LSTD. Initial analysis of least squares temporal difference learning (LSTD) date back to the work of [Baird \[1995\]](#), [Bradtke and Barto \[1996\]](#), [Boyan \[1999\]](#) and [Nedić and Bertsekas \[2003\]](#). Since then, the finite sample performance of the algorithm has been analyzed by [Lazaric et al. \[2012\]](#), [Bhandari et al. \[2018\]](#), [Duan et al. \[2021\]](#) and its behavior in the offline setting studied by [Yu \[2010\]](#), [Li et al. \[2021\]](#), [Mou et al. \[2020, 2021\]](#). [Tu and Recht \[2018\]](#) analyze on-policy LSTD for the LQR setting. We evaluate our contributions in light of these previous works in [Section 4.2](#).

Other OPE methods. Apart from these methods, researchers have studied “min-max” algorithms for OPE which estimate the value of the underlying policy using ideas from the importance sampling literature [[Liu et al., 2018](#), [Uehara et al., 2020](#)]. Furthermore, significant attention has been devoted to the problem of offline policy evaluation under unobserved confounding [[Namkoong et al., 2020](#), [Xu et al., 2021](#), [Kallus and Zhou, 2020](#)], where the value function cannot be estimated perfectly given the observed variables. [Xie and Jiang \[2021\]](#) establish formal guarantees for the BVFT algorithm which carries out policy evaluation for general *nonlinear* function classes assuming realizability, albeit under stronger notions of data coverage (see [Assumption 7](#)). Recent work by [Zhan et al. \[2022\]](#) extends this line of research. They introduce a

new algorithm which works under weaker data coverage assumptions than those in Xie and Jiang [2021]. However, to do so they require additional assumptions on the expressivity of the underlying class of function approximators. In particular, Zhan et al. [2022], and the class of minimax algorithms more broadly, rely on a function class that can (at a minimum) realize the state-occupancy density ratio between the distribution induced by the policy π and the offline distribution \mathcal{D} , which is a distinct condition from linear realizability of Q^π .

Lower bounds under linear realizability. For the finite-horizon, policy evaluation setting, Wang et al. [2021a] illustrate how exponential dependence on the horizon is unavoidable, even if the offline covariance matrix is robustly full rank. Since then, these bounds have been extended to the discounted, infinite horizon case by Amortila et al. [2020] and Zanette [2021]. Importantly, Amortila et al. [2020] establish that OPE can be information-theoretically intractable, even if: 1) all features are bounded, 2) Σ_{cov} is full rank, and 3) the learner has access to infinitely many samples drawn as in Section 1.1. Analogous negative results for online or generative-model settings have been shown to hold even in the presence of a constant suboptimality gap [Wang et al., 2021c] or polynomially large action sets [Weisz et al., 2021a,b]. Duan et al. [2020] prove lower bounds for OPE which hold for general function classes. Foster et al. [2021] illustrate that polynomially many samples in the size of the state space are necessary for offline policy evaluation, even if concentrability and realizability both hold. In short, a clean characterization of when offline policy evaluation is tractable using linear function approximation has, so far, proven to be quite elusive.

2 Preliminaries

Before delving into our main results, we review some of the relevant definitions and preliminaries.

Notation. We use $s \in \mathcal{S}$ and $a \in \mathcal{A}$ to denote states and actions, \top to denote vector or matrix transposes, and \dagger to denote pseudoinverses. For a matrix X , we let $\text{cond}(X) := \sigma_{\max}(X)/\sigma_{\min}(X)$ denote its condition number, the ratio between the largest and smallest singular values. For symmetric matrices, A and B , we use $A \succeq B$ if $A - B$ is positive semidefinite. We let $\rho(X) := \max_i |\lambda_i(X)|$ be the spectral radius of a matrix X where λ_i are the eigenvalues. We say that a matrix is *stable* if its spectral radius is strictly smaller than 1. For square, stable matrices A , we let $\text{dlyap}(A)$ be the solution, over X , to the discrete-time Lyapunov equation: $X = A^\top X A + I$. This equation has a solution if and only if $\rho(A) < 1$. If the solution exists, it admits the closed-form expression $\text{dlyap}(A) = \sum_{j=0}^{\infty} (A^\top)^j A^j$. Lastly, we say $a \lesssim b$ if $a \leq c \cdot b$ for some universal constant c .

We define the next state-action covariance Σ_{next} and the distribution shift coefficient \mathcal{C}_{ds} as

$$\Sigma_{\text{next}} := \mathbb{E}_{\substack{(s,a) \sim \mathcal{D} \\ s' \sim P(\cdot|s,a), a' \sim \pi(s')}} [\phi(s', a') \phi(s', a')^\top], \quad \mathcal{C}_{\text{ds}} := \inf\{\beta > 0 : \Sigma_{\text{next}} \preceq \beta \Sigma_{\text{cov}}\}. \quad (2.1)$$

Note that \mathcal{C}_{ds} is guaranteed to be finite if Σ_{cov} is full rank. Given a dataset $\{(s_i, a_i, r(s_i, a_i), s'_i, a'_i)\}_{i=1}^n$ of n i.i.d. data points drawn according to the data generating process described in Section 1.1, we define the empirical counterparts of the quantities in Eq. (1.5),

$$\hat{\Sigma}_{\text{cov}} := \frac{1}{n} \sum_{i=1}^n \phi(s_i, a_i) \phi(s_i, a_i)^\top, \quad \hat{\Sigma}_{\text{cr}} := \frac{1}{n} \sum_{i=1}^n \phi(s_i, a_i) \phi(s'_i, a'_i)^\top, \quad (2.2)$$

as well as the true, and empirical, mean feature-reward vectors:

$$\theta_{\phi, r} := \mathbb{E}_{\mathcal{D}} \phi(s, a) r(s, a), \quad \hat{\theta}_{\phi, r} := \frac{1}{n} \sum_{i=1}^n \phi(s_i, a_i) r(s_i, a_i). \quad (2.3)$$

Linear regression. Next, we introduce moment-type quantities that arise in our analysis of linear regression. We adopt the approach from [Hsu et al. \[2012\]](#) and make use of the statistical leverages ρ_s and $\rho_{s'}$, although other approaches for analyzing linear regression will yield the same qualitative results:

$$\rho_s := \sup_{(s,a) \in \text{supp}(\mathcal{D})} \|\Sigma_{\text{cov}}^{-1/2} \phi(s, a)\|_2, \quad \rho_{s'} := \sup_{\substack{(s,a) \in \text{supp}(\mathcal{D}), \\ s' \in \text{supp}(P(\cdot|s,a)), a' \in \text{supp}(\pi(s'))}} \|\Sigma_{\text{cov}}^{-1/2} \phi(s', a')\|_2. \quad (2.4)$$

In addition, we define the variances σ_{cov}^2 and σ_{cr}^2 where,

$$\sigma_{\text{cov}}^2 := \|\mathbb{E}(\Sigma_{\text{cov}}^{-1/2} \phi(s, a) \phi(s, a)^\top \Sigma_{\text{cov}}^{-1/2})^2 - I\|_{\text{op}}, \quad (2.5)$$

and σ_{cr}^2 is the maximum of the following two quantities,

$$\sup_{\|v\|=1} \mathbb{E} \left(v^\top \Sigma_{\text{cov}}^{-1/2} \phi(s', a') \right)^2 \|\Sigma_{\text{cov}}^{-1/2} \phi(s, a)\|_2^2 - \|\Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{cr}}^\top \Sigma_{\text{cov}}^{1/2} v\|^2 \quad (2.6)$$

$$\sup_{\|v\|=1} \mathbb{E} \left(v^\top \Sigma_{\text{cov}}^{-1/2} \phi(s, a) \right)^2 \|\Sigma_{\text{cov}}^{-1/2} \phi(s', a')\|_2^2 - \|\Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{cr}} \Sigma_{\text{cov}}^{1/2} v\|^2. \quad (2.7)$$

In [Appendix C.3](#), we prove that σ_{cr}^2 and σ_{cov}^2 can always be upper bounded in terms of the statistical leverages and the coefficient \mathcal{C}_{ds} . However, they can be much smaller in some settings,⁶ so for the sake of generality, we opt to state our bounds in terms of these quantities. Throughout our analysis of methods for offline policy evaluation, we will repeatedly make use of the following concentration result:

Lemma 2.1. *For all $n \gtrsim \rho_s^2 \log(d/\delta)$, define the estimation errors,*

$$\varepsilon_{\text{op}} := \|\Sigma_{\text{cov}}^{1/2} (\gamma \widehat{\Sigma}_{\text{cov}}^{-1} \widehat{\Sigma}_{\text{cr}}) \Sigma_{\text{cov}}^{-1/2} - \gamma \Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{cr}} \Sigma_{\text{cov}}^{-1/2}\|_{\text{op}}, \quad \varepsilon_r := \|\Sigma_{\text{cov}}^{1/2} (\widehat{\Sigma}_{\text{cov}}^{-1} \widehat{\theta}_{\phi,r} - \Sigma_{\text{cov}}^{-1} \theta_{\phi,r})\|_2. \quad (2.8)$$

With probability $1 - \delta$, ε_r and ε_{op} satisfy the following inequalities:

$$\begin{aligned} \varepsilon_{\text{op}} &\lesssim \sqrt{\frac{\max(\sigma_{\text{cr}}^2, \rho_s^2 \mathcal{C}_{\text{ds}}) \log(d/\delta)}{n}} + \frac{\max(\mathcal{C}_{\text{ds}} \rho_s^2, \rho_s \rho_{s'}) \log(d/\delta)}{n} \\ \varepsilon_r &\lesssim \sqrt{\frac{\max(\|\Sigma_{\text{cov}}^{-1/2} \theta_0^*\|^2 \sigma_{\text{cov}}^2, d) \log(d/\delta)}{n}} + \frac{\|\Sigma_{\text{cov}}^{-1/2} \theta_{\phi,r}\|_2 \rho_s^2 \log(d/\delta)}{n}. \end{aligned}$$

Later on, we state our upper bounds on the policy evaluation error of FQI and LSTD in terms of these regression errors ε_{op} , ε_r , with the understanding that they satisfy the high probability upper bounds above.

3 Fitted Q-Iteration

In this section, we present our first set of results illustrating how stability ([Assumption 2](#)) characterizes the success of fitted Q-iteration for OPE under linear realizability of Q^π . Following some initial remarks regarding the functional form of the FQI solution, in [Section 3.1](#), we present our upper bound on the estimation error of FQI. Later on, in [Section 3.2](#), we illustrate how our Lyapunov stability analysis unifies previous studies of when FQI succeeds and conclude by discussing lower bounds and limitations of the algorithm in [Section 3.3](#).

FQI preliminaries. From examining the definition of FQI in [Eq. \(1.2\)](#), we see that, at the population level, the algorithm develops the recursion:

$$\theta_{t+1} = \gamma \Sigma_{\text{cov}}^{-1} \Sigma_{\text{cr}} \theta_t + \Sigma_{\text{cov}}^{-1} \theta_{\phi,r}.$$

⁶For example, tighter bounds can be achieved if the distributions are hypercontractive, see [Appendix C.3](#).

Unrolling the recursion above, and setting $\theta_0 = 0$,⁷ the T -step regression vector is equal to:

$$\theta_T = \sum_{k=0}^T (\gamma \Sigma_{\text{cov}}^{-1} \Sigma_{\text{cr}})^k \Sigma_{\text{cov}}^{-1} \theta_{\phi, r}. \quad (3.1)$$

Linear realizability of Q^π (Assumption 1) implies that the true weight vector θ_γ^* satisfies the equation,

$$\Sigma_{\text{cov}} \theta_\gamma^* = \theta_{\phi, r} + \gamma \Sigma_{\text{cr}} \theta_\gamma^*. \quad (3.2)$$

Hence, if $I - \gamma \Sigma_{\text{cov}}^{-1} \Sigma_{\text{cr}}$ is invertible, $\theta_\gamma^* = (I - \gamma \Sigma_{\text{cov}}^{-1} \Sigma_{\text{cr}})^{-1} \Sigma_{\text{cov}}^{-1} \theta_{\phi, r}$. We now recall the following fact:

Fact 3.1. *If $\rho(A) < 1$, then the matrix $(I - A)$ is invertible. Moreover, $(I - A)^{-1} = \sum_{k=0}^{\infty} A^k$.*

Using this, along with the observation that the spectrum of a matrix is invariant to the choice of basis, we see that if stability (Assumption 2) holds, the vector θ_γ^* can also be written as a power series:

$$\theta_\gamma^* = \sum_{k=0}^{\infty} (\gamma \Sigma_{\text{cov}}^{-1} \Sigma_{\text{cr}})^k \Sigma_{\text{cov}}^{-1} \theta_{\phi, r}. \quad (3.3)$$

One of the key insights tying stability and FQI is that, regardless of whether $\gamma \Sigma_{\text{cov}}^{-1} \Sigma_{\text{cr}}$ is stable, the FQI solution at the population level is *always* equal to the power series in Eq. (3.1). If stability holds, performing infinitely many exact FQI updates converges to θ_γ^* . However, θ_γ^* is (in general) *only* equal to this power series if stability holds, which hints at the necessity of this condition. With these connections between stability and the functional forms of FQI and θ_γ^* in mind, we now present our upper bounds on the performance of this algorithm.

3.1 Stability is sufficient for fitted Q-iteration

Theorem 1. *Assume that Q^π is linearly realizable (Assumption 1) and that stability holds (Assumption 2). For $\varepsilon_{\text{op}}, \varepsilon_r$ defined as in Eq. (2.8), if $n \gtrsim \rho_s^2 \log(d/\delta)$ and $\varepsilon_{\text{op}} \leq 1/(6\|P_\gamma\|_{\text{op}}^2)$, T -step FQI satisfies,*

$$\begin{aligned} \|\Sigma_{\text{cov}}^{1/2}(\hat{\theta}_T - \theta_\gamma^*)\|_2 &\lesssim \text{cond}(P_\gamma)^{1/2} \|P_\gamma\|_{\text{op}} \cdot \varepsilon_r + \text{cond}(P_\gamma) \|P_\gamma\|_{\text{op}}^2 \cdot \|\Sigma_{\text{cov}}^{-1/2} \theta_{\phi, r}\| \cdot \varepsilon_{\text{op}} \\ &\quad + \text{cond}(P_\gamma) \|P_\gamma\|_{\text{op}} \cdot \|\Sigma_{\text{cov}}^{-1/2} \theta_{\phi, r}\|_2 \cdot \exp\left(-\frac{T+1}{2\|P_\gamma\|_{\text{op}}}\right). \end{aligned}$$

Let $\hat{Q}^\pi(s, a) := \phi(s, a)^\top \hat{\theta}_T$. Much like in standard analyses of linear regression, from Theorem 1 we immediately obtain: (1) a bound on $\mathbb{E}_{\mathcal{D}} |Q^\pi(s, a) - \hat{Q}^\pi(s, a)|$ via Jensen's inequality since $\mathbb{E}_{\mathcal{D}} (Q^\pi(s, a) - \hat{Q}^\pi(s, a))^2 = \|\Sigma_{\text{cov}}^{1/2}(\hat{\theta}_T - \theta_\gamma^*)\|_2^2$ and (2) a bound on $|Q^\pi(s, a) - \hat{Q}^\pi(s, a)|$ for any (s, a) pair since $|Q^\pi(s, a) - \hat{Q}^\pi(s, a)| \leq \|\Sigma_{\text{cov}}^{-1/2} \phi(s, a)\|_2 \|\Sigma_{\text{cov}}^{1/2}(\hat{\theta}_T - \theta_\gamma^*)\|_2$ via Cauchy-Schwarz.

We defer the full proof to Appendix A.1 and instead summarize the key steps here. The theorem is essentially a perturbation bound which distinguishes between two sources of error in policy evaluation for FQI: ε_r which captures errors in learning the rewards, and the dominant error, ε_{op} , which comes from estimating the transitions. Since under stability, we can write the true vector θ_γ^* as a power series in second moment matrices (see Eq. (3.3)), and since $\hat{\theta}_T$ is by definition a truncated power series in the empirical counterparts of these matrices, we can show that the error between θ_γ^* and $\hat{\theta}_T$ is bounded by the operator norm of two power series: one in $(\gamma \Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{cr}} \Sigma_{\text{cov}}^{-1/2})^k$ and the other in $(\gamma \hat{\Sigma}_{\text{cov}}^{-1/2} \hat{\Sigma}_{\text{cr}} \hat{\Sigma}_{\text{cov}}^{-1/2})^k$. Lyapunov arguments directly show that the powers of $(\gamma \Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{cr}} \Sigma_{\text{cov}}^{-1/2})$ decay exponentially in k since the matrix is stable. For the empirical version, we use the fact that any stable matrix A has nontrivial *stability margin*: for small enough perturbations Δ , matrices of the form $A + \Delta$ satisfy similar decay rates to A . Thus, we can bound the two power series by simple geometric series and the perturbation bound follows.

We now highlight some of the salient aspects of the bound.

⁷We initialize at 0 for simplicity, but this is not fundamental for the overall analysis of FQI.

Throughout this section, we will present corollaries of this result, which can be viewed as specializations of [Theorem 1](#) to particular settings. In each case, we will focus on discussing variants of the perturbation bound, [Eq. \(3.7\)](#), which hold under the specific assumptions.

3.2.1 Low distribution shift implies stability

Recent work by [Wang et al. \[2021b\]](#) shows that FQI succeeds at OPE for infinite horizon, discounted problems if there is low distribution shift. More formally, they prove offline evaluation is tractable if the offline covariance Σ_{cov} has good coverage over the next state covariance Σ_{next} as per the following assumption.

Assumption 4 (Low Distribution Shift). There is low distribution shift if $\mathcal{C}_{\text{ds}} < 1/\gamma^2$.

Note that if \mathcal{D} is the stationary measure for π , then $\Sigma_{\text{cov}} = \Sigma_{\text{next}}$ and [Assumption 4](#) holds with $\mathcal{C}_{\text{ds}} = 1$ (recall the definition of \mathcal{C}_{ds} in [Eq. \(2.1\)](#)). Under this low distribution shift condition, we prove:

Corollary 3.1. *If there is low distribution shift ([Assumption 4](#)) and if Σ_{cov} is full rank, then for all $j \geq 0$,*

$$\|(\gamma \Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{cr}} \Sigma_{\text{cov}}^{-1/2})^j\|_{\text{op}} \leq (\sqrt{\mathcal{C}_{\text{ds}} \gamma^2})^j. \quad (3.8)$$

Hence, $\|P_\gamma\|_{\text{op}} \leq 1/(1 - \gamma\sqrt{\mathcal{C}_{\text{ds}}})$, [Assumption 2](#) holds,⁸. Furthermore, for $\gamma_{\text{ds}} := \gamma\sqrt{\mathcal{C}_{\text{ds}}}$, if Q^π is linearly realizable ([Assumption 1](#)), $n \gtrsim \rho_s^2 \log(d/\delta)$, and $\varepsilon_{\text{op}} \leq 1/2(1 - \gamma_{\text{ds}})$, T -step FQI satisfies:

$$\|\Sigma_{\text{cov}}^{1/2}(\hat{\theta}_T - \theta_\gamma^*)\|_2 \lesssim \frac{1}{1 - \gamma_{\text{ds}}} \varepsilon_r + \frac{1}{(1 - \gamma_{\text{ds}})^2} \|\Sigma_{\text{cov}}^{-1/2} \theta_{\phi, r}\| \cdot \varepsilon_{\text{op}} + \frac{1}{1 - \gamma_{\text{ds}}} \gamma_{\text{ds}}^{T+1}.$$

While low distribution shift implies stability, the converse is not true. It is not hard to come up with examples where $\gamma \Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{cr}} \Sigma_{\text{cov}}^{-1/2}$ is stable, yet the distribution shift coefficient is larger than $1/\gamma^2$.

3.2.2 Bellman completeness implies stability

In addition to the low-distribution shift setting, FQI is known to succeed in both finite horizon and discounted, infinite horizon settings under a representational condition known as Bellman completeness [[Szepesvári and Munos, 2005](#), [Wang et al., 2021a,b](#)]:

Assumption 5 (Bellman completeness). A feature map ϕ is Bellman complete for an MDP \mathcal{M} , if for all $\theta \in \mathbb{R}^d$, there exists a vector θ' such that for all $(s, a) \in \mathcal{S} \times \mathcal{A}$,

$$\phi(s, a)^\top \theta' = \mathbb{E}[r(s, a)] + \gamma \mathbb{E}_{s' \sim P(\cdot | s, a)} \phi(s', \pi'(s))^\top \theta.$$

Intuitively, completeness asserts that Bellman backups of linear functions of the features again lie in the span of the features. It has previously been observed [[Wang et al., 2021a,b](#)] that completeness implies a certain “non-expansiveness” of Bellman backups. This non-expansiveness is the key step towards establishing the connection to stability and is formalized in the following result:

Corollary 3.2. *If ϕ is Bellman complete ([Assumption 5](#)) and Σ_{cov} is full rank, then for all $j \geq 0$,*

$$\|(\gamma \Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{cr}} \Sigma_{\text{cov}}^{-1/2})^j\|_{\text{op}} \leq \rho_s \gamma^j. \quad (3.9)$$

Hence, $\|P_\gamma\|_{\text{op}} \leq \rho_s/(1 - \gamma)$, and [Assumption 2](#) holds. Furthermore, if Q^π is linearly realizable ([Assumption 1](#)), $n \gtrsim \rho_s^2 \log(d/\delta)$, and $\varepsilon_{\text{op}} \leq (1 - \gamma)/(2\rho_s)$, T -step FQI satisfies:⁹

$$\|\Sigma_{\text{cov}}^{1/2}(\hat{\theta}_T - \theta_\gamma^*)\|_2 \lesssim \frac{\rho_s}{1 - \gamma} \cdot \varepsilon_r + \frac{\rho_s^2}{(1 - \gamma)^2} \varepsilon_{\text{op}} + \frac{\rho_s^2}{1 - \gamma} \gamma^{T+1}.$$

As with low distribution shift, the converse is not true. It is not hard to find examples where stability holds yet Bellman completeness does not.

⁸A matrix A is stable if and only if $\lim_{k \rightarrow \infty} A^k = 0$.

⁹Completeness implies realizability of rewards which in turn implies $\|\Sigma_{\text{cov}}^{-1/2} \theta_{\phi, r}\|_2^2 = \mathbb{E}r(s, a)^2 \leq 1$, see [Lemma A.6](#).

3.3 Stability is necessary for fitted Q-iteration

We conclude our analysis of FQI by showing that our characterization of when the algorithm succeeds is exactly sharp, in an instance-dependent sense. If stability fails that is, $\rho(\gamma \Sigma_{\text{cov}}^{-1} \Sigma_{\text{cr}}) > 1$, then estimation procedures of this sort are guaranteed to have exponentially large variance.

Proposition 3.4. *Let \mathcal{M} be any infinite-horizon, discounted MDP with corresponding offline distribution \mathcal{D} which satisfies the following properties: Σ_{cov} is full rank and $\gamma \Sigma_{\text{cov}}^{-1} \Sigma_{\text{cr}}$ has an eigenvalue λ with $|\lambda| > 1$. Then, approximations of the T -step FQI solution, $\hat{Q}^\pi(s, a) = \phi(s, a)^\top \hat{\theta}_T$ where,*

$$\hat{\theta}_T := \sum_{k=0}^T (\gamma \Sigma_{\text{cov}}^{-1} \Sigma_{\text{cr}})^k \Sigma_{\text{cov}}^{-1} \hat{\theta}_{\phi, r}, \quad \hat{\theta}_{\phi, r} := \theta_{\phi, r}^* + z,$$

and z is a zero-mean, random vector satisfying $\Lambda := \mathbb{E} z z^\top \succ 0$, have exponentially large variance,

$$\mathbb{E} \|\hat{\theta}_T - \mathbb{E} \hat{\theta}_T\|_2^2 \geq \sigma_{\min}(\Lambda) \cdot \left(\frac{\lambda^{T+1} - 1}{\lambda - 1} \right)^2.$$

This proposition corroborates empirical findings on the instability of FQI by Wang et al. [2021b] and shows that an idealized variant of FQI incurs exponentially large variance (in the number of rounds T) for an instance that results in an unstable “backup operator” $\gamma \Sigma_{\text{cov}}^{-1} \Sigma_{\text{cr}}$. By standard bias-variance decomposition, this directly implies exponentially large error for estimating θ_γ^* . Although, note that since stability does not hold, there is no guarantee that θ_γ^* can be written as a power series, so it may not even be the limiting solution of population FQI as discussed at the beginning of this section.

The algorithm is idealized in two senses, both of which are relatively minor. First, it has perfect knowledge of Σ_{cov} and Σ_{cr} which does not happen in practice, but is favorable to the algorithm, resulting in a stronger lower bound. Second, the error in estimating the reward is assumed to have a full-rank covariance; this arises naturally whenever rewards are perturbed with centered Gaussian noise since Σ_{cov} is full rank. Thus, the result shows that even when the dynamics are known, errors in estimating the rewards will be exponentially magnified, resulting in overall divergence of the algorithm.

While the theorem does not consider the marginally stable case where $\rho(\gamma \Sigma_{\text{cov}}^{-1} \Sigma_{\text{cr}}) = 1$, we note in the proof that if the spectral radius is exactly one, the variance can grow at least linearly with T . However, marginal stability introduces other issues as we illustrate later on.

At this point, it is natural to wonder whether stability is necessary not just for the success of this algorithm, but rather for the success of *any* algorithm at offline policy evaluation. It turns out that this is not the case. As we will show in the following section, least squares temporal differencing works under strictly weaker conditions than fitted Q-iteration.

4 Least Squares Temporal Difference Learning

Building on our analysis of FQI, we now analyze how a closely related algorithm, least squares temporal difference learning, overcomes some of its shortcomings in the context of offline policy evaluation. Similarly to the previous section, we start by illustrating how invertibility is sufficient for LSTD in Section 4.1, and discuss connections to previous sufficient conditions in Section 4.2. Lastly, we conclude in Section 4.3 by presenting lower bounds which show that if invertibility does not hold, then the offline policy evaluation problem cannot be solved using linear estimators (FQI and LSTD being special cases), even asymptotically.

4.1 Invertibility is sufficient for LSTD

Theorem 2. *Assume that realizability and invertibility (Assumptions 1 and 3) both hold and let $\varepsilon_r, \varepsilon_{\text{op}}$ be defined as in Eq. (2.8). If $n \gtrsim \rho_s^2 \log(d/\delta)$ and $\varepsilon_{\text{op}} \leq \sigma_{\min}(I - \gamma \Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{cr}} \Sigma_{\text{cov}}^{-1/2})/2$, then the LSTD solution,*

$$\hat{\theta}_{\text{LS}} := (I - \gamma \hat{\Sigma}_{\text{cov}}^{-1} \hat{\Sigma}_{\text{cr}})^\dagger \hat{\Sigma}_{\text{cov}}^{-1} \hat{\theta}_{\phi, r},$$

satisfies the following error guarantee:

$$\|\Sigma_{\text{cov}}^{1/2}(\theta_\gamma^* - \hat{\theta}_{\text{LS}})\|_2 \lesssim \frac{1}{\sigma_{\min}(I - \gamma\Sigma_{\text{cov}}^{-1/2}\Sigma_{\text{cr}}\Sigma_{\text{cov}}^{-1/2})} \cdot \varepsilon_r + \frac{1}{\sigma_{\min}(I - \gamma\Sigma_{\text{cov}}^{-1/2}\Sigma_{\text{cr}}\Sigma_{\text{cov}}^{-1/2})^2} \|\Sigma_{\text{cov}}^{-1/2}\theta_{\phi,r}\|_2 \cdot \varepsilon_{\text{op}}.$$

As per our discussion immediately following [Theorem 1](#), the upper bound on $\|\Sigma_{\text{cov}}^{1/2}(\theta_\gamma^* - \hat{\theta}_\gamma)\|_2$ again directly implies guarantees on $|Q^\pi(s, a) - \hat{Q}^\pi(s, a)|$, both pointwise and in expectation, where now $\hat{Q}^\pi(s, a) = \phi(s, a)^\top \hat{\theta}_{\text{LS}}$. On a technical level, the proof follows from standard perturbation bounds on matrix inverses.

Our upper bound for LSTD has qualitatively similar properties to that presented for FQI in [Theorem 1](#).

A sharper notion of problem complexity. Much like $\|P_\gamma\|_{\text{op}}$ for FQI, the magnitude of our upper bound for the policy evaluation error of LSTD is determined an instance-dependent quantity: $1/\sigma_{\min}(I - \gamma\Sigma_{\text{cov}}^{-1/2}\Sigma_{\text{cr}}\Sigma_{\text{cov}}^{-1/2})$. As before this term is: (1) never much larger than $1/(1 - \gamma)$ for settings where OPE was previously shown to be tractable (see the next subsection for further discussion of this point), and (2) is often significantly smaller. For example, for the OPE instance detailed in [\(3.5\)](#), if $p \leq .7$, then $1/\sigma_{\min}(I - \gamma\Sigma_{\text{cov}}^{-1/2}\Sigma_{\text{cr}}\Sigma_{\text{cov}}^{-1/2}) \leq 4$ for all $\gamma \in (0, 1)$.

Coordinate invariance. From [Proposition 3.2](#), we know that for any choice of full rank matrix L and features $\tilde{\phi}(\cdot) = L\phi(\cdot)$, the whitened cross-covariance in these new features, $\gamma\tilde{\Sigma}_{\text{cov}}^{-1/2}\tilde{\Sigma}_{\text{cr}}\tilde{\Sigma}_{\text{cov}}^{-1/2}$ (see definition in [Eq. \(3.4\)](#)) is equal to $\gamma U\Sigma_{\text{cov}}^{-1/2}\Sigma_{\text{cr}}\Sigma_{\text{cov}}^{-1/2}U^\top$ for some orthogonal matrix U . Since conjugating by an orthogonal matrix preserves singular values, $1/\sigma_{\min}(I - \gamma\Sigma_{\text{cov}}^{-1/2}\Sigma_{\text{cr}}\Sigma_{\text{cov}}^{-1/2})$ is invariant to the choice of coordinates.

4.2 Contextualizing Invertibility

Paralleling our discussion of stability for FQI, we now discuss how our notion of invertibility relates to previous conditions analyzed in the literature. Furthermore, we will present how stability implies invertibility, establishing a precise “nesting” between the classes of OPE problems which satisfy each condition.


4.2.1 Stability \subsetneq Invertibility

Proposition 4.1. *If Σ_{cov} is full rank and $\gamma\Sigma_{\text{cov}}^{-1/2}\Sigma_{\text{cr}}\Sigma_{\text{cov}}^{-1/2}$ is stable ([Assumption 2](#)), then $I - \gamma\Sigma_{\text{cov}}^{-1/2}\Sigma_{\text{cr}}\Sigma_{\text{cov}}^{-1/2}$ is invertible ([Assumption 3](#)). Furthermore,*

$$\frac{1}{\sigma_{\min}(I - \gamma\Sigma_{\text{cov}}^{-1/2}\Sigma_{\text{cr}}\Sigma_{\text{cov}}^{-1/2})} \lesssim \text{cond}(P_\gamma)^{1/2}\|P_\gamma\|_{\text{op}}. \quad (4.1)$$

The main message of this proposition is twofold. First, for the case of linear function approximation, any OPE problem that is solvable via FQI, must also be solvable via LSTD. Second, from [Eq. \(4.1\)](#) we see that main complexity measure for [Theorem 2](#), $1/\sigma_{\min}(I - \gamma\Sigma_{\text{cov}}^{-1/2}\Sigma_{\text{cr}}\Sigma_{\text{cov}}^{-1/2})$ is never larger than the corresponding upper bound for FQI in [Theorem 1](#), $\text{cond}(P_\gamma)^{1/2}\|P_\gamma\|_{\text{op}}$.

Interestingly enough, while stability implies invertibility. The converse is not true. There exist problems for which $I - \gamma\Sigma_{\text{cov}}^{-1/2}\Sigma_{\text{cr}}\Sigma_{\text{cov}}^{-1/2}$ is invertible, but $\gamma\Sigma_{\text{cov}}^{-1/2}\Sigma_{\text{cr}}\Sigma_{\text{cov}}^{-1/2}$ is not stable. For example, consider the following 2 state MDP, with no actions:



If we set $R(s_0) = R(s_1) = \text{Unif}(\{\pm 1\})$, and $\phi(s_0) = 1$, $\phi(s_1) = 2$, then this OPE instance is trivially linearly realizable with $\theta_\gamma^* = 0$. If the offline distribution \mathcal{D} places mass p on s_0 and $1 - p$ on s_1 , it is easy to see that $I - \gamma\Sigma_{\text{cov}}^{-1/2}\Sigma_{\text{cr}}\Sigma_{\text{cov}}^{-1/2}$ is invertible for all $p, \gamma \in (0, 1)$. However, for $p = \gamma = .9$, $\gamma\Sigma_{\text{cov}}^{-1/2}\Sigma_{\text{cr}}\Sigma_{\text{cov}}^{-1/2}$ is at

least $3/2$, hence stability does not hold and FQI will necessarily diverge. Together, these results establish a separation between the set of problems solvable via FQI and those solvable via LSTD.¹⁰

Moreover, for the set of previously analyzed settings where stability holds, we can establish quantitative upper bounds on $1/\sigma_{\min}(I - \gamma \Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{cr}} \Sigma_{\text{cov}}^{-1/2})$ illustrating how this quantity is comparable to $1/(1 - \gamma)$.

Corollary 4.1. *If there is low distribution shift (Assumption 4), then for $\gamma_{\text{ds}} := \gamma \sqrt{\mathcal{C}_{\text{ds}}}$,*

$$\frac{1}{\sigma_{\min}(I - \gamma \Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{cr}} \Sigma_{\text{cov}}^{-1/2})} \leq \frac{1}{1 - \gamma_{\text{ds}}}.$$

Moreover, if Bellman completeness holds (Assumption 5), then

$$\frac{1}{\sigma_{\min}(I - \gamma \Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{cr}} \Sigma_{\text{cov}}^{-1/2})} \leq \frac{\rho_s}{1 - \gamma}.$$

This result follows from observing that $1/\sigma_{\min}(I - \gamma \Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{cr}} \Sigma_{\text{cov}}^{-1/2}) = \|(I - \gamma \Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{cr}} \Sigma_{\text{cov}}^{-1/2})^{-1}\|_{\text{op}}$. Since stability holds for both of these settings, we can use Fact 3.1 to write $(I - \gamma \Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{cr}} \Sigma_{\text{cov}}^{-1/2})^{-1}$ as an infinite power series in $\gamma \Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{cr}} \Sigma_{\text{cov}}^{-1/2}$. Applying the triangle inequality and the bounds from Eqs. (3.8) and (3.9) on the powers of $\gamma \Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{cr}} \Sigma_{\text{cov}}^{-1/2}$ finishes the proof of this corollary.

4.2.2 Other connections

Recent work by Mou et al. [2020] analyzes oracle inequalities for solving projected fixed point equations, of which the Bellman equation, Eq. (3.2), is a special case. For the offline policy evaluation setting, they prove that a stochastic approximation variant of LSTD succeeds if the following condition holds:

Assumption 6. $\gamma \Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{cr}} \Sigma_{\text{cov}}^{-1/2}$ satisfies $\kappa := \frac{1}{2} \lambda_{\max}(\gamma \Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{cr}} \Sigma_{\text{cov}}^{-1/2} + (\gamma \Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{cr}} \Sigma_{\text{cov}}^{-1/2})^\top) < 1$.

Here, λ_{\max} denotes the maximal eigenvalue of a matrix.¹¹ In their paper, the authors remark how Assumption 6 directly implies that $I - \gamma \Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{cr}} \Sigma_{\text{cov}}^{-1/2}$ is invertible. Amongst other quantities, their bounds scale with $1/(1 - \kappa)$. This quantity is always at least as large as $1/\sigma_{\min}(I - \gamma \Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{cr}} \Sigma_{\text{cov}}^{-1/2})$.

Proposition 4.2. *If Assumption 6 holds, then $I - \gamma \Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{cr}} \Sigma_{\text{cov}}^{-1/2}$ is invertible and*

$$\frac{1}{\sigma_{\min}(I - \gamma \Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{cr}} \Sigma_{\text{cov}}^{-1/2})} \leq \frac{1}{1 - \kappa}.$$

Recent work by Li et al. [2021] extends the stochastic approximation analysis from Mou et al. [2020] to incorporate variance reduction techniques. Their upper bounds directly assume invertibility, but also have explicit dependence $1/(1 - \gamma)$ which can be quite loose in certain settings as detailed earlier.

4.3 Invertibility is necessary for all linear estimators

We finish our presentation of LSTD by proving that invertibility is not just sufficient, it is also strictly necessary for LSTD, as well as for all other linear estimators. To do so, we first formally define what we mean by linear estimators:

Definition 4.1 (Population Linear Estimator). Let Alg be a deterministic algorithm which given an MDP \mathcal{M} , a distribution \mathcal{D} over $\mathcal{S} \times \mathcal{A}$, and a policy π returns a function $\hat{Q}^\pi : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$. Furthermore, let $(\mathcal{M}, \mathcal{D}, \pi)$ and $(\bar{\mathcal{M}}, \bar{\mathcal{D}}, \bar{\pi})$ be two OPE instances such that:

¹⁰The careful reader might observe that the main reason why FQI fails in this example is that the algorithm is sensitive to the scale of the next state features. For instance, stability (and realizability) would hold if $|\phi(s_1)| < 1$.

¹¹The matrix in Assumption 6 is symmetric so all eigenvalues are real and the maximum is well defined.

- The corresponding action value functions Q^π, \bar{Q}^π are both linearly realizable in a feature map ϕ .
- The covariance, cross-covariance, mean feature-reward vectors, and expected rewards (as defined in Eqs. (1.5) and (2.3)) are identical in $(\mathcal{M}, \mathcal{D}, \pi)$ and $(\bar{\mathcal{M}}, \bar{\mathcal{D}}, \bar{\pi})$:

$$\bar{\Sigma}_{\text{cov}} = \Sigma_{\text{cov}}, \quad \bar{\Sigma}_{\text{cross}} = \Sigma_{\text{cr}}, \quad \bar{\theta}_{\phi, r} = \theta_{\phi, r}, \quad \mathbb{E}_{\bar{\mathcal{D}}} \bar{r}(s, a) = \mathbb{E}_{\mathcal{D}} r(s, a).$$

We say that Alg is a *population linear estimator* if $\text{Alg}(\mathcal{M}, \mathcal{D}, \pi) = \text{Alg}(\bar{\mathcal{M}}, \bar{\mathcal{D}}, \bar{\pi})$.

While our focus has been on studying the finite sample performance of estimators for OPE, in this definition we choose to catalogue algorithms based on their asymptotic behavior so as to neatly abstract technical modifications like variance reduction. These techniques introduce differences in finite-sample performance, but are not essential to the overall *identifiability* concerns that are the focus of this subsection. Intuitively, linear estimators are those whose population-level solution depends on the low-order moments which are relevant for common estimation procedures like ordinary least-squares. From Eqs. (1.4) and (3.1), we see that LSTD and FQI both satisfy this definition. We now present our lower bound characterizing the limits of linear estimators for OPE.

Theorem 3. *Let $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, R, \gamma)$ be any MDP with associated offline distribution \mathcal{D} such that:*

- $Q^\pi(s, a)$ is linearly realizable in ϕ .¹²
- Σ_{cov} is full rank.
- $I - \gamma \Sigma_{\text{cov}}^{-1} \Sigma_{\text{cr}}$ is rank deficient.

Then, there exists an MDP $\bar{\mathcal{M}} = (\mathcal{S}, \mathcal{A}, P, \bar{R}, \gamma)$, differing only in the reward function, such that for the same offline distribution \mathcal{D} :

- *The Q -function for π in $\bar{\mathcal{M}}, \bar{Q}^\pi$, is linearly realizable in the same feature mapping ϕ .*
- *The covariance, cross-covariance, mean-feature reward vectors, and expected rewards in $\bar{\mathcal{M}}$ are identical to their counterparts in \mathcal{M} :*

$$\bar{\Sigma}_{\text{cov}} = \Sigma_{\text{cov}}, \quad \bar{\Sigma}_{\text{cross}} = \Sigma_{\text{cr}}, \quad \bar{\theta}_{\phi, r} = \theta_{\phi, r}, \quad \mathbb{E}_{\bar{\mathcal{D}}} \bar{r}(s, a) = \mathbb{E}_{\mathcal{D}} r(s, a).$$

- *The Q functions are different: $\mathbb{E}_{\mathcal{D}}(Q^\pi(s, a) - \bar{Q}^\pi(s, a))^2 \geq \sigma_{\min}(\Sigma_{\text{cov}})$.*

Consequently, if we define LE as the set of population linear estimators,

$$\inf_{\mathcal{A} \in \text{LE}} \sup_{(\mathcal{M}', \mathcal{D}', \pi') \in \mathcal{N}} \mathbb{E}_{\mathcal{D}}(Q'^\pi(s, a) - \hat{Q}^\pi(s, a))^2 \gtrsim \sigma_{\min}(\Sigma_{\text{cov}}),$$

where $\hat{Q}^\pi = \text{Alg}(\mathcal{M}', \mathcal{D}', \pi')$ and $\mathcal{N} = \{(\mathcal{M}, \mathcal{D}, \pi), (\bar{\mathcal{M}}, \bar{\mathcal{D}}, \bar{\pi})\}$

In other words, this theorem states that for *any* OPE instance where $I - \gamma \Sigma_{\text{cov}}^{-1} \Sigma_{\text{cr}}$, or equivalently, $I - \gamma \Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{cr}} \Sigma_{\text{cov}}^{-1/2}$, is rank deficient, we can perturb the rewards to construct an alternative instance with matching low order moments. Consequently, any population linear estimator (such as LSTD or FQI) will return the same estimate \hat{Q}^π in both cases. Yet, since the Q -functions are distinct, they will necessarily converge to the wrong answer in one case. Together with Theorem 2, this result illustrates how our characterization of the settings where LSTD succeeds is exactly sharp in an instance-dependent (local) sense.

¹²Without loss of generality, we can assume ϕ has a constant feature.

4.3.1 Going beyond linear estimators

Of course, linear estimators are by no means the only methods for offline policy evaluation. Given the negative result from [Theorem 3](#), it is natural to ask: what are the algorithm-independent limits for OPE under linear realizability? We close with a brief discussion of how our work provides insight into this question.


Introduced by [Xie and Jiang \[2021\]](#), the BVFT algorithm is a statistically (but not computationally) efficient algorithm for OPE using a *general* function class \mathcal{F} under two assumptions: (1) Q^π is realizable in \mathcal{F} and (2) the offline distribution \mathcal{D} and the MDP dynamics satisfy a strong data coverage condition referred to as *pushforward concentrability*.

Assumption 7 (Pushforward Concentrability, [Xie and Jiang \[2021\]](#)). An MDP \mathcal{M} and offline distribution \mathcal{D} satisfy pushforward concentrability if:

- The offline distribution \mathcal{D} has strictly positive mass on all $(s, a) \in \mathcal{S} \times \mathcal{A}$: $P_{\mathcal{D}}(s, a) > 0$.
- There exists a constant $1 \leq C_A < \infty$ such that for any $(s, a) \in \mathcal{S} \times \mathcal{A}$, $P_{\mathcal{D}}(a | s) \geq 1/C_A$.
- There exists a constant C_S such that for all $s, s' \in \mathcal{S}$ and $a \in \mathcal{A}$.¹³

$$\frac{P(s' | s, a)}{P_{\mathcal{D}}(s')} \leq C_S.$$

In the linear function approximation setting, realizability of Q^π in \mathcal{F} reduces to our realizability condition ([Assumption 1](#)). However, pushforward concentrability is in general distinct from stability or invertibility. That is, for problems that are linearly realizable, pushforward concentrability does not imply, nor is implied by, the assumption that $\sigma_{\min}(I - \gamma \Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{cr}} \Sigma_{\text{cov}}^{-1/2}) > 0$. Therefore, there exist settings where linear estimators may fail, yet BVFT can succeed and vice versa. In particular, consider the following MDP:



There are no actions, and the feature map is defined as $\phi(s_0) = 1$ and $\phi(s_1) = 2/\gamma$. If we set the rewards to have nonzero variance and satisfy $\mathbb{E}r(s_1) = r^*$, $\mathbb{E}r(s_0) = \frac{-1}{2(1-\gamma)}r^*$, then this MDP is linearly realizable with $\theta_\gamma^* = \frac{\gamma}{2(1-\gamma)}r^*$. For any $\gamma \in (0, 1)$, a simple continuity argument proves that there always exists a $p \in (0, 1)$ such that if the offline distribution places mass p on s_0 and $1 - p$ on s_1 , Σ_{cov} is full rank and $\gamma \Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{cr}} \Sigma_{\text{cov}}^{-1/2} = 1$. Therefore, realizability and pushforward concentrability both hold, but invertibility does not. For the converse direction, it is not hard to see how one might construct examples where linear realizability and invertibility both hold, but [Assumption 7](#) does not. The first condition asserting that \mathcal{D} be supported on all states and actions is particularly stringent.¹⁴

Recall from the construction in [Theorem 3](#), that for any OPE instance where invertibility fails, the alternative $\bar{\mathcal{M}}$ has exactly the same states and transitions. Therefore, any estimator that outperforms linear methods must necessarily consider nonlinear or higher-order interactions between features and rewards. Interestingly enough, a simple tabular method which ignores the feature mapping ϕ and directly estimates the rewards successfully approximates the value function in this example.

5 Discussion

In this work, we characterize the exact limits of linear estimators for offline policy evaluation, under the assumption that the value function is linearly realizable in some known set of features. Our stability and

¹³We omit the last assumption on the initial state distribution from [Xie and Jiang \[2021\]](#) as it is not essential for the purposes of our discussion.

¹⁴In this construction, we have departed from our assumption that $\sup_{s,a} |r(s, a)| < 1$ since $\mathbb{E}r(s, a)$ is on the order of $\Omega((1 - \gamma)^{-1})$. However, the magnitude of the rewards should not affect the *identifiability* of Q^π , only the estimation rate for quantities like ε_{op} and ε_r .

invertibility based analyses introduce new, sharper notions of complexity for this classical setting and provide a simple, unifying perspective which brings together previously disparate analysis of popular algorithms.

Two extensions to our results pertain to the finite horizon setting and to policy optimization. As a starting point, we have focused on the infinite-horizon, discounted setting as the conditions there are cleaner than in the finite horizon case. Nevertheless, we conjecture that Lyapunov stability and invertibility can be used to analyze finite horizon problems as well. Regarding policy optimization, understanding when this task is possible under linear realizability is an important direction for future work. We hope that our characterization of linear estimators for policy evaluation provides a useful perspective on this closely related problem. Apart from these extensions, it would be value to study quantitative, instance-dependent lower bounds on the sample complexity necessary for offline policy evaluation under linear realizability.

6 Acknowledgments

We gratefully acknowledge the support of Microsoft through the BAIR Open Research Commons. JCP is supported by an NSF Graduate Research Fellowship. Sham Kakade acknowledges funding from the National Science Foundation Grant under award #CCF-1703574 and the Office of Naval Research under award N00014-21-1-2822.

References

- Philip Amortila, Nan Jiang, and Tengyang Xie. A variant of the Wang-Foster-Kakade lower bound for the discounted setting. *arXiv:2011.01075*, 2020.
- András Antos, Csaba Szepesvári, and Rémi Munos. Learning near-optimal policies with Bellman-residual minimization based fitted policy iteration and a single sample path. *Machine Learning*, 2008.
- Leemon Baird. Residual algorithms: Reinforcement learning with function approximation. In *Machine Learning*, 1995.
- Richard Bellman. On the approximation of curves by line segments using dynamic programming. *Communications of the ACM*, 1961.
- Richard Bellman and Stuart Dreyfus. Functional approximations and dynamic programming. *Mathematical Tables and Other Aids to Computation*, 1959.
- Dimitri P Bertsekas and John N Tsitsiklis. Neuro-dynamic programming: An overview. In *IEEE Conference on Decision and Control*, 1995.
- Jalaj Bhandari, Daniel Russo, and Raghav Singal. A finite time analysis of temporal difference learning with linear function approximation. In *Conference on Learning Theory*, 2018.
- Justin A Boyan. Least-squares temporal difference learning. In *International Conference on Machine Learning*, 1999.
- Steven J Bradtke and Andrew G Barto. Linear least-squares algorithms for temporal difference learning. *Machine learning*, 1996.
- Jinglin Chen and Nan Jiang. Information-theoretic considerations in batch reinforcement learning. In *International Conference on Machine Learning*, 2019.
- Yeshwanth Cherapanamjeri, Samuel B Hopkins, Tarun Kathuria, Prasad Raghavendra, and Nilesch Tripuraneni. Algorithms for heavy-tailed statistics: Regression, covariance estimation, and beyond. In *Symposium on Theory of Computing*, 2020.

- Yaqi Duan, Zeyu Jia, and Mengdi Wang. Minimax-optimal off-policy evaluation with linear function approximation. In *International Conference on Machine Learning*, 2020.
- Yaqi Duan, Mengdi Wang, and Martin J Wainwright. Optimal policy evaluation using kernel-based temporal difference methods. *arXiv:2109.12002*, 2021.
- Damien Ernst, Pierre Geurts, and Louis Wehenkel. Tree-based batch mode reinforcement learning. *Journal of Machine Learning Research*, 2005.
- Dylan J Foster, Akshay Krishnamurthy, David Simchi-Levi, and Yunzong Xu. Offline reinforcement learning: Fundamental barriers for value function approximation. *arXiv:2111.10919*, 2021.
- Geoffrey J Gordon. *Approximate solutions to Markov decision processes*. PhD thesis, Carnegie Mellon University, 1999.
- Daniel Hsu, Sham M Kakade, and Tong Zhang. Random design analysis of ridge regression. In *Conference on Learning Theory*, 2012.
- Chi Jin, Zhuoran Yang, Zhaoran Wang, and Michael I Jordan. Provably efficient reinforcement learning with linear function approximation. In *Conference on Learning Theory*, 2020.
- Nathan Kallus and Angela Zhou. Confounding-robust policy evaluation in infinite-horizon reinforcement learning. *Advances in Neural Information Processing Systems*, 2020.
- Alessandro Lazaric, Mohammad Ghavamzadeh, and Rémi Munos. Finite-sample analysis of least-squares policy iteration. *Journal of Machine Learning Research*, 2012.
- Tianjiao Li, Guanghui Lan, and Ashwin Pananjady. Accelerated and instance-optimal policy evaluation with linear function approximation. *arxiv:2112.13109*, 2021.
- Qiang Liu, Lihong Li, Ziyang Tang, and Dengyong Zhou. Breaking the curse of horizon: Infinite-horizon off-policy estimation. *Advances in Neural Information Processing Systems*, 2018.
- Stanislav Minsker. On some extensions of bernstein’s inequality for self-adjoint operators. *Statistics & Probability Letters*, 2017.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, Stig Petersin, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dhharshan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassibis. Human-level control through deep reinforcement learning. *Nature*, 2015.
- Wenlong Mou, Ashwin Pananjady, and Martin J. Wainwright. Optimal oracle inequalities for solving projected fixed-point equations. *arXiv:2012.05299*, 2020.
- Wenlong Mou, Ashwin Pananjady, Martin J Wainwright, and Peter L Bartlett. Optimal and instance-dependent guarantees for markovian linear stochastic approximation. *arXiv:2112.12770*, 2021.
- Rémi Munos. Error bounds for approximate policy iteration. In *International Conference on Machine Learning*, 2003.
- Rémi Munos. Performance bounds in l_p -norm for approximate value iteration. *SIAM journal on control and optimization*, 2007.
- Dheeraj Nagaraj, Xian Wu, Guy Bresler, Prateek Jain, and Praneeth Netrapalli. Least squares regression with markovian data: Fundamental limits and algorithms. *Advances in Neural Information Processing Systems*, 2020.

- Hongseok Namkoong, Ramtin Keramati, Steve Yadlowsky, and Emma Brunskill. Off-policy policy evaluation for sequential decisions under unobserved confounding. *Advances in Neural Information Processing Systems*, 2020.
- A Nedić and Dimitri P Bertsekas. Least squares policy evaluation algorithms with linear function approximation. *Discrete Event Dynamic Systems*, 2003.
- Juan Perdomo, Jack Umenberger, and Max Simchowitz. Stabilizing dynamical systems via policy gradient methods. *Advances in Neural Information Processing Systems*, 2021.
- Dieter Reetz. Approximate solutions of a discounted markovian decision process. *Bonner Mathematische Schriften*, 1977.
- Martin Riedmiller. Neural fitted Q iteration—first experiences with a data efficient neural reinforcement learning method. In *European Conference on Machine Learning*, 2005.
- Paul J Schweitzer and Abraham Seidmann. Generalized polynomial approximations in markovian decision processes. *Journal of Mathematical Analysis and Applications*, 1985.
- Gilbert W Stewart. *Matrix perturbation theory*. Citeseer, 1990.
- Csaba Szepesvári and Rémi Munos. Finite time bounds for sampling based fitted value iteration. In *International Conference on Machine Learning*, 2005.
- Joel A Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of Computational Mathematics*, 2012.
- John N Tsitsiklis and Benjamin Van Roy. Feature-based methods for large scale dynamic programming. *Machine Learning*, 1996.
- Stephen Tu and Benjamin Recht. Least-squares temporal difference learning for the linear quadratic regulator. In *International Conference on Machine Learning*, 2018.
- Masatoshi Uehara, Jiawei Huang, and Nan Jiang. Minimax weight and Q-function learning for off-policy evaluation. In *International Conference on Machine Learning*, 2020.
- Ruosong Wang, Dean Foster, and Sham M. Kakade. What are the statistical limits of offline RL with linear function approximation? In *International Conference on Learning Representations*, 2021a.
- Ruosong Wang, Yifan Wu, Ruslan Salakhutdinov, and Sham M. Kakade. Instabilities of offline RL with pre-trained neural representation. In *International Conference on Machine Learning*, 2021b.
- Yuanhao Wang, Ruosong Wang, and Sham Kakade. An exponential lower bound for linearly realizable MDPs with constant suboptimality gap. *Advances in Neural Information Processing Systems*, 2021c.
- Gellért Weisz, Philip Amortila, and Csaba Szepesvári. Exponential lower bounds for planning in MDPs with linearly-realizable optimal action-value functions. In *Algorithmic Learning Theory*, 2021a.
- Gellért Weisz, Csaba Szepesvári, and András György. Tensorplan and the few actions lower bound for planning in MDPs under linear realizability of optimal value functions. *arXiv:2110.02195*, 2021b.
- Ward Whitt. Approximations of dynamic programs, I. *Mathematics of Operations Research*, 1978.
- Tengyang Xie and Nan Jiang. Batch value-function approximation with only realizability. In *International Conference on Machine Learning*, 2021.
- Liyuan Xu, Heishiro Kanagawa, and Arthur Gretton. Deep proxy causal learning and its application to confounded bandit policy evaluation. *Advances in Neural Information Processing Systems*, 2021.

- Lin Yang and Mengdi Wang. Reinforcement learning in feature space: Matrix bandit, kernels, and regret bound. In *International Conference on Machine Learning*, 2020.
- Huizhen Yu. Convergence of least squares temporal difference methods under general conditions. In *International Conference on Machine Learning*, 2010.
- Andrea Zanette. Exponential lower bounds for batch reinforcement learning: Batch RL can be exponentially harder than online RL. In *International Conference on Machine Learning*, 2021.
- Andrea Zanette, Alessandro Lazaric, Mykel Kochenderfer, and Emma Brunskill. Learning near optimal policies with low inherent bellman error. In *International Conference on Machine Learning*, 2020.
- Wenhao Zhan, Baihe Huang, Audrey Huang, Nan Jiang, and Jason D Lee. Offline reinforcement learning with realizability and single-policy concentrability. *arXiv:2202.04634*, 2022.

A Supporting Arguments for Section 3: FQI

A.1 Proof of Theorem 1: stability is sufficient for FQI

The existence of $\widehat{\Sigma}_{\text{cov}}$ and the upper bounds on the regression errors ε_r and ε_{op} are guaranteed by Lemmas C.3 and C.4. To analyze the error of FQI, we introduce the shorthand,

$$A := \gamma \Sigma_{\text{cov}}^{-1} \Sigma_{\text{cr}}, \quad \widehat{A} = \gamma \widehat{\Sigma}_{\text{cov}}^{-1} \widehat{\Sigma}_{\text{cr}}, \quad \theta_t^* := \sum_{k=0}^t A^k \theta_0^*, \quad \widehat{\theta}_t := \sum_{k=0}^t \widehat{A}^k \widehat{\theta}_0,$$

$$w_t := \widehat{\theta}_t - \theta_t^*, \quad \Delta := \widehat{A} - A,$$

where $\widehat{\theta}_0 = \widehat{\Sigma}_{\text{cov}}^{-1} \widehat{\theta}_{\phi, r}$ and $\theta_0^* = \Sigma_{\text{cov}}^{-1} \theta_{\phi, r}$. Using this notation, by stability, we observe that $\theta_\gamma^* = \theta_\infty^*$, and we can write the errors vectors of the t -step FQI solution as,

$$\Sigma_{\text{cov}}^{1/2} (\theta_\gamma^* - \widehat{\theta}_t) = \Sigma_{\text{cov}}^{1/2} \sum_{k=t+1}^{\infty} A^k \theta_0^* + \Sigma_{\text{cov}}^{1/2} w_t. \quad (\text{A.1})$$

Next, we develop the recursion in w_t ,

$$\begin{aligned} w_{t+1} &= \sum_{j=0}^{t+1} \widehat{A}^j \widehat{\theta}_0 - \sum_{j=0}^{t+1} A^j \theta_0^* \\ &= \widehat{A} \widehat{\theta}_t + \widehat{\theta}_0 - A \theta_t^* - \theta_0^* \\ &= \widehat{A} w_t + \Delta \theta_t^* + w_0. \end{aligned}$$

Unrolling the recursion and multiplying on the left by $\Sigma_{\text{cov}}^{1/2}$, we get that

$$\Sigma_{\text{cov}}^{1/2} w_{t+1} = \sum_{j=0}^{t+1} \left(\Sigma_{\text{cov}}^{1/2} \widehat{A} \Sigma_{\text{cov}}^{-1/2} \right)^j \Sigma_{\text{cov}}^{1/2} w_0 + \sum_{j=0}^t \left(\Sigma_{\text{cov}}^{1/2} \widehat{A} \Sigma_{\text{cov}}^{-1/2} \right)^j \left(\Sigma_{\text{cov}}^{1/2} \Delta \Sigma_{\text{cov}}^{-1/2} \right) \Sigma_{\text{cov}}^{1/2} \theta_{t-j}^*.$$

Note that $\varepsilon_r = \|\Sigma_{\text{cov}}^{1/2} w_0\|$ and $\varepsilon_{\text{op}} = \|\Sigma_{\text{cov}}^{1/2} \Delta \Sigma_{\text{cov}}^{-1/2}\|_{\text{op}}$. Therefore, taking the norm of both sides and applying the triangle inequality,

$$\|\Sigma_{\text{cov}}^{1/2} w_{t+1}\| \leq \sum_{k=0}^{t+1} \left\| \left(\Sigma_{\text{cov}}^{1/2} \widehat{A} \Sigma_{\text{cov}}^{-1/2} \right)^k \right\|_{\text{op}} \|\Sigma_{\text{cov}}^{1/2} w_0\| \quad (\text{A.2})$$

$$\begin{aligned} &+ \sum_{k=0}^t \left\| \left(\Sigma_{\text{cov}}^{1/2} \widehat{A} \Sigma_{\text{cov}}^{-1/2} \right)^k \right\|_{\text{op}} \|\Sigma_{\text{cov}}^{1/2} \Delta \Sigma_{\text{cov}}^{-1/2}\|_{\text{op}} \sup_{0 \leq h \leq t} \|\Sigma_{\text{cov}}^{1/2} \theta_h^*\| \\ &= (\varepsilon_r + \varepsilon_{\text{op}} \sup_{0 \leq h \leq t} \|\Sigma_{\text{cov}}^{1/2} \theta_h^*\|) \cdot \sum_{k=0}^{t+1} \left\| \left(\Sigma_{\text{cov}}^{1/2} \widehat{A} \Sigma_{\text{cov}}^{-1/2} \right)^k \right\|_{\text{op}}. \end{aligned} \quad (\text{A.3})$$

Now, recalling the definition of θ_h^* , we bound:

$$\sup_{0 \leq h \leq t} \|\Sigma_{\text{cov}}^{1/2} \theta_h^*\| \leq \sum_{j=0}^t \left\| \left(\gamma \Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{cr}} \Sigma_{\text{cov}}^{-1/2} \right)^j \right\|_{\text{op}} \|\Sigma_{\text{cov}}^{1/2} \theta_0^*\|. \quad (\text{A.4})$$

Therefore, combining these last two inequalities (A.4), (A.3), and the identity from Eq. (A.1),

$$\begin{aligned} \|\Sigma_{\text{cov}}^{1/2}(\theta_\gamma^* - \hat{\theta}_t)\|_2 &\leq \sum_{k=t+1}^{\infty} \left\| \left(\Sigma_{\text{cov}}^{1/2} A \Sigma_{\text{cov}}^{-1/2} \right)^k \right\|_{\text{op}} \|\Sigma_{\text{cov}}^{1/2} \theta_0^*\|_2 + \|\Sigma_{\text{cov}}^{1/2} w_t\|_2 \\ &\leq \|\Sigma_{\text{cov}}^{1/2} \theta_0^*\|_2 \sum_{k=t+1}^{\infty} \alpha_k + \left(\varepsilon_r + \varepsilon_{\text{op}} \|\Sigma_{\text{cov}}^{1/2} \theta_0^*\|_2 \sum_{k=0}^{t-1} \alpha_k \right) \sum_{k=0}^t \hat{\alpha}_k, \end{aligned} \quad (\text{A.5})$$

where $\hat{\alpha}_k := \left\| \left(\Sigma_{\text{cov}}^{1/2} \hat{A} \Sigma_{\text{cov}}^{-1/2} \right)^j \right\|_{\text{op}}$ and $\alpha_k := \left\| \left(\Sigma_{\text{cov}}^{1/2} A \Sigma_{\text{cov}}^{-1/2} \right)^j \right\|_{\text{op}}$. Since $\varepsilon_{\text{op}} \leq 1/(6\|P_\gamma\|_{\text{op}})$ and $\gamma \Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{cr}} \Sigma_{\text{cov}}^{-1/2}$ is stable, Lemma A.1 tells us that

$$\begin{aligned} \hat{\alpha}_j &= \|P_\gamma^{-1/2} P_\gamma^{1/2} \left(\Sigma_{\text{cov}}^{1/2} \hat{A} \Sigma_{\text{cov}}^{-1/2} \right)^j\|_{\text{op}} \\ &\leq \|P_\gamma^{-1/2}\|_{\text{op}} \|P_\gamma^{1/2} \left(\Sigma_{\text{cov}}^{1/2} \hat{A} \Sigma_{\text{cov}}^{-1/2} \right)^j\|_{\text{op}} \\ &\leq \|P_\gamma^{1/2}\|_{\text{op}} \|P_\gamma^{1/2}\|_{\text{op}} \left(1 - \frac{1}{2\|P_\gamma\|_{\text{op}}} \right)^{j/2} = \text{cond}(P_\gamma)^{1/2} \left(1 - \frac{1}{2\|P_\gamma\|_{\text{op}}} \right)^{j/2}. \end{aligned}$$

Using similar reasoning, we get that

$$\alpha_j \leq \text{cond}(P_\gamma)^{1/2} \left(1 - \frac{1}{\|P_\gamma\|_{\text{op}}} \right)^{j/2}.$$

In conclusion, $\|\Sigma_{\text{cov}}^{1/2}(\theta_\gamma^* - \hat{\theta}_t)\|_2$ is bounded by,

$$\|\Sigma_{\text{cov}}^{1/2} \theta_0^*\|_2 \text{cond}(P_\gamma)^{1/2} \left(1 - \frac{1}{\|P_\gamma\|_{\text{op}}} \right)^{(t+1)/2} \sum_{k=0}^{\infty} \alpha_k + \left(\varepsilon_r + \varepsilon_{\text{op}} \|\Sigma_{\text{cov}}^{1/2} \theta_0^*\|_2 \sum_{k=0}^{\infty} \alpha_k \right) \sum_{k=0}^{\infty} \hat{\alpha}_k.$$

The final bound comes from summing the geometric series, $\sum_{j=0}^{\infty} (1-c)^{j/2} = (1 - \sqrt{1-c})^{-1}$, for $c \in (0, 1)$ and applying the numerical inequality,

$$\left(1 - \sqrt{1 - \frac{1}{2z}} \right)^{-1} \leq 10z,$$

which holds for all $z \geq 1$. To prove the numerical inequality, observe that all critical points of the function $10z - (1 - \sqrt{1 - \frac{1}{2z}})^{-1}$ over the domain $x \geq 0$ lie at values of x less than 1, and the limit as $x \rightarrow \infty$ is ∞ .

Lemma A.1. *Let A be a square, stable matrix and let $P = \text{dlyap}(A)$. Then, for all $k \geq 0$,*

$$\|A^k\|_{\text{op}}^2 \leq \text{cond}(P) \left(1 - \frac{1}{\|P\|_{\text{op}}} \right)^k.$$

Furthermore, for any matrix Δ such that $\|\Delta\|_{\text{op}} \leq 1/(6\|P\|_{\text{op}}^2)$,

$$\|(A + \Delta)^k\|_{\text{op}}^2 \leq \text{cond}(P) \left(1 - \frac{1}{2\|P\|_{\text{op}}} \right)^k.$$

Proof. This particular lemma is almost identical to the one from Perdomo et al. [2021]. However, we include the proof for the sake of providing a self-contained presentation. For the first result, by definition of the

solution to the Lyapunov equation, for any unit vector x ,

$$\begin{aligned} x^\top A^\top P A x &= x^\top P x - x^\top I x \\ &= x^\top P x \left(1 - \frac{\|x\|_2^2}{x^\top P x}\right) \\ &\leq x^\top P x \left(1 - \frac{1}{\|P\|_{\text{op}}}\right). \end{aligned}$$

Hence, $A^\top P A \preceq P(1 - \|P\|_{\text{op}}^{-1})$. By iterating $(A^k)^\top P A^k \preceq P(1 - \|P\|_{\text{op}}^{-1})^k$ and

$$\|P^{1/2} A^k\|_{\text{op}}^2 \leq \|P\|_{\text{op}} (1 - \|P\|_{\text{op}}^{-1})^k.$$

Therefore,

$$\|A^k\|_{\text{op}} = \|P^{-1/2} P^{1/2} A^k\|_{\text{op}} \leq \|P^{-1/2}\|_{\text{op}} \|P^{1/2} A^k\|_{\text{op}} \leq \text{cond}(P)^{1/2} \left(1 - \frac{1}{\|P\|_{\text{op}}}\right)^{k/2}.$$

For the second result, using the insights from above,

$$(A + \Delta)^\top P (A + \Delta) = A^\top P A + A^\top P \Delta + \Delta^\top P A + \Delta^\top P \Delta.$$

Now, $A^\top P A \preceq P(1 - \|P\|_{\text{op}}^{-1})$ and

$$\|A^\top P \Delta\|_{\text{op}} = \|\Delta^\top P A\|_{\text{op}} \leq \|\Delta P^{1/2}\|_{\text{op}} \|P^{1/2} A\|_{\text{op}} \leq \|\Delta P^{1/2}\|_{\text{op}} \|P^{1/2}\|_{\text{op}} \leq \|\Delta\|_{\text{op}} \|P\|_{\text{op}}.$$

Bounding, $\|\Delta^\top P \Delta\|_{\text{op}} \leq \|P\|_{\text{op}} \|\Delta\|_{\text{op}}^2$, and using the fact that $P \succeq I$ we get that for,

$$\|\Delta\|_{\text{op}} \leq 1/(6\|P\|_{\text{op}}^2),$$

the following relationship holds:

$$A^\top P \Delta + \Delta^\top P A + \Delta^\top P \Delta \preceq P \frac{1}{2\|P\|_{\text{op}}}.$$

Therefore,

$$(A + \Delta)^\top P (A + \Delta) \preceq P \left(1 - \frac{1}{2\|P\|_{\text{op}}}\right),$$

and the second result follows by using the same steps as the first. \square

A.2 Proof of **Proposition 3.2**: coordinate invariance of P_γ

If we define the whitened features, $\phi_w(\cdot) = \Sigma_{\text{cov}}^{-1/2} \phi(\cdot)$, then $\tilde{\phi}(\cdot) = L' \phi_w(\cdot)$ where $L' = L \Sigma_{\text{cov}}^{1/2}$. Now, let $U S V^\top$ be the singular value decomposition of L' . Then,

$$\tilde{\Sigma}_{\text{cov}} = \mathbb{E}_{x \sim \mathcal{D}} \tilde{\phi}(x) \tilde{\phi}(x)^\top = L' \mathbb{E}_{x \sim \mathcal{D}} \phi_w(x) \phi_w(x)^\top L'^\top = L' L'^\top = U S^2 U^\top,$$

where we have used the fact that the whitened features have identity covariance. By this calculation, we have that $\tilde{\Sigma}_{\text{cov}}^{1/2} = U S U^\top$. Using similar substitutions, we can also deduce that $\tilde{\Sigma}_{\text{cr}} = L' \Sigma_{\text{cr}}^{(w)} L'^\top$ where $\Sigma_{\text{cr}}^{(w)} = \Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{cr}} \Sigma_{\text{cov}}^{-1/2}$. Therefore,

$$\tilde{\Sigma}_{\text{cov}}^{-1/2} \tilde{\Sigma}_{\text{cr}} \tilde{\Sigma}_{\text{cov}}^{-1/2} = (U S^{-1} U^\top) (U S V^\top) \Sigma_{\text{cr}}^{(w)} (V S U^\top) (U S^{-1} U^\top) = (U V^\top) \Sigma_{\text{cr}}^{(w)} (U V^\top)^\top.$$

Since $(U V^\top)$ is an orthogonal matrix, the equality of condition numbers follows by the fact that for any matrix A and orthogonal matrix M , $M A M^\top = A$ have the same singular values. On the other hand, the invariance of the operator norm of P_γ follows from the following lemma:

Lemma A.2. Let A be a stable matrix and M be any orthogonal matrix, then

$$\|\text{dlyap}(A^\top)\|_{\text{op}} = \|\text{dlyap}(MA^\top M^\top)\|_{\text{op}}.$$

Proof. Let $P = \text{dlyap}(A)$ be the unique solution over X to the matrix equation:

$$X = A^\top X A + I.$$

Likewise, let $P' = \text{dlyap}(MAM^\top)$ be the unique solution (over X') to the equation:

$$X' = MA^\top M^\top X' MAM^\top + I.$$

From this, we can deduce that $M^\top X' M = A^\top M^\top X' M A + I$. Therefore, $P = M^\top X' M = M^\top P' M$. The conclusion follows from the fact that singular values are invariant to conjugation by an orthogonal matrix. \square

A.3 Proof of Lemma 3.3: FQI under specific growth rates

As discussed in the main body, the proof is identical to that of Theorem 1 except that we specialize to the particular assumptions on the growth of matrix powers. We recall the key inequality from the proof of the main theorem, Eq. (A.5):

$$\|\Sigma_{\text{cov}}^{1/2}(\hat{\theta}_T - \theta_\gamma^*)\|_2 \leq \|\Sigma_{\text{cov}}^{1/2}\theta_0^*\|_2 \sum_{k=t+1}^{\infty} \alpha_k + \left(\varepsilon_r + \varepsilon_{\text{op}} \|\Sigma_{\text{cov}}^{1/2}\theta_0^*\|_2 \sum_{k=0}^{t-1} \alpha_k \right) \sum_{k=0}^t \hat{\alpha}_k.$$

Here, $\alpha_k = \left\| \left(\gamma \Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{cr}} \Sigma_{\text{cov}}^{-1/2} \right)^k \right\|_{\text{op}}$ and $\hat{\alpha}_k := \left\| \left(\Sigma_{\text{cov}}^{1/2} (\gamma \hat{\Sigma}_{\text{cov}}^{-1} \hat{\Sigma}_{\text{cr}}) \Sigma_{\text{cov}}^{-1/2} \right)^k \right\|_{\text{op}}$. By assumption, $\alpha_k \leq \alpha \beta^k$ hence, $\sum_{k=0}^{\infty} \alpha_k \leq \alpha / (1 - \beta)$. Now, by Lemma A.3 since

$$\varepsilon_{\text{op}} = \left\| \Sigma_{\text{cov}}^{1/2} (\gamma \hat{\Sigma}_{\text{cov}}^{-1} \hat{\Sigma}_{\text{cr}}) \Sigma_{\text{cov}}^{-1/2} - \gamma \Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{cr}} \Sigma_{\text{cov}}^{-1/2} \right\|_{\text{op}},$$

we have that:

$$\hat{\alpha}_k \leq \alpha (\beta + \varepsilon_{\text{op}} \alpha)^k.$$

Therefore, as long as $\varepsilon_{\text{op}} < \frac{9}{10} \frac{(1-\beta)}{\alpha}$,

$$\sum_{k=0}^{\infty} \hat{\alpha}_k \leq \alpha \sum_{k=0}^{\infty} (\beta + \varepsilon_{\text{op}} \alpha)^k = \alpha \frac{1}{1 - \beta - \alpha \varepsilon_{\text{op}}} \leq 10 \frac{\alpha}{1 - \beta}.$$

Putting everything together,

$$\|\Sigma_{\text{cov}}^{1/2}(\hat{\theta}_T - \theta_\gamma^*)\|_2 \lesssim \|\Sigma_{\text{cov}}^{1/2}\theta_0^*\|_2 \frac{\alpha^2}{1 - \beta} \cdot \beta^{t+1} + \left(\varepsilon_r + \varepsilon_{\text{op}} \|\Sigma_{\text{cov}}^{1/2}\theta_0^*\|_2 \frac{\alpha}{1 - \beta} \right) \frac{\alpha}{1 - \beta}.$$

Lemma A.3. Let A be a square matrix such that for all nonnegative integers j , $\|A^j\|_{\text{op}} \leq a \cdot b^j$ for scalars $a > 0$ and $b \in (0, 1)$. Then, for any square matrix Δ if we let $\varepsilon := \|\Delta\|_{\text{op}}$ then,

$$\|(A + \Delta)^n\|_{\text{op}} \leq a(b + \varepsilon \cdot a)^n.$$

Proof. We begin by expanding $(A + \Delta)^n$ into monomials $T_{k,j}$,

$$(A + \Delta)^n = \sum_{k=0}^n \sum_{j=1}^{\binom{n}{k}} T_{k,j}, \tag{A.6}$$

where each $T_{k,j}$ has k factors of Δ and $n - k$, A factors. Now, by the submultiplicative property of the operator norm,

$$\|T_{k,j}\|_{\text{op}} \leq \varepsilon^k \prod_{s_i \in S_{k,j}} \|A^{s_i}\|_{\text{op}},$$

where $S_{k,j}$ is a set of positive integers s_i satisfying $\sum_i s_i = n - k$ and $|S| \leq k + 1$. Using our assumption on the growth of $\|A^k\|_{\text{op}}$, we get that,

$$\|T_{k,j}\|_{\text{op}} \leq \varepsilon^k \prod_{s_i \in S_{k,j}} (a \cdot b^{s_i}) \leq a^{k+1} \varepsilon^k b^{n-k}.$$

Going back to the original expansion into monomials, and using the identity,

$$\sum_{k=0}^n \binom{n}{k} x^k = (1+x)^n.$$

We conclude:

$$\|(A + \Delta)^n\|_{\text{op}} \leq a \cdot b^n \sum_{k=0}^n \binom{n}{k} \left(\frac{a\varepsilon}{b}\right)^k = \alpha b^n \left(1 + \frac{a \cdot \varepsilon}{b}\right)^n = \alpha(b + a\varepsilon)^n.$$

□

A.4 Proof of [Corollary 3.1](#): low distribution shift implies stability

Consider the augmented covariance matrix,

$$\mathbb{E} \begin{bmatrix} \phi(s, a) \\ \phi(s', a') \end{bmatrix} \begin{bmatrix} \phi(s, a) \\ \phi(s', a') \end{bmatrix}^\top = \begin{bmatrix} \Sigma_{\text{cov}} & \Sigma_{\text{cr}} \\ \Sigma_{\text{cr}}^\top & \Sigma_{\text{next}} \end{bmatrix} \succeq 0.$$

By a Schur complement argument, $\Sigma_{\text{cr}}^\top \Sigma_{\text{cov}}^{-1} \Sigma_{\text{cr}} \preceq \Sigma_{\text{next}}$. After conjugating by $\Sigma_{\text{cov}}^{-1/2}$ and multiplying by γ^2 , we get that:

$$(\gamma \Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{cr}} \Sigma_{\text{cov}}^{-1/2})^\top (\gamma \Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{cr}} \Sigma_{\text{cov}}^{-1/2}) \preceq \gamma^2 \Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{next}} \Sigma_{\text{cov}}^{-1/2}.$$

Now, by the low distribution shift assumption, $\gamma^2 \Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{next}} \Sigma_{\text{cov}}^{-1/2} \preceq \gamma^2 \Sigma_{\text{cov}}^{-1/2} (\mathcal{C}_{\text{ds}} \Sigma_{\text{cov}}) \Sigma_{\text{cov}}^{-1/2} = \mathcal{C}_{\text{ds}} \gamma^2 I$. Therefore, $(\gamma \Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{cr}} \Sigma_{\text{cov}}^{-1/2})^\top (\gamma \Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{cr}} \Sigma_{\text{cov}}^{-1/2}) \preceq \mathcal{C}_{\text{ds}} \gamma^2 I$. Iterating for $j \geq 0$ gives the first part of the result. The rest follows from [Lemma 3.3](#) by observing that [Eq. \(3.6\)](#) holds with $\alpha = 1, \beta = \sqrt{\mathcal{C}_{\text{ds}} \gamma^2} \in (0, 1)$.

A.5 Proofs of [Corollary 3.2](#): Bellman completeness implies stability

To take advantage of matrix notation, for this result we assume that the state-action space is finite, $|\mathcal{S}||\mathcal{A}| < \infty$. In particular, we introduce the following quantities.

1. Feature matrix $\Phi \in \mathbb{R}^{|\mathcal{S}||\mathcal{A}| \times d}$.
2. Offline distribution vector $\mu \in \mathbb{R}^{|\mathcal{S}||\mathcal{A}|}$.

With this, we have that $\Sigma_{\text{cov}} = \Phi^\top \text{diag}(\mu) \Phi$ and $\Sigma_{\text{cr}} = \Phi^\top \text{diag}(\mu) P^{(\pi)} \Phi$ where $P^{(\pi)}$ is a row stochastic matrix representing the transition operator. [Corollary 3.2](#) follows from the following lemma and [Lemma 3.3](#).

Lemma A.4. *If ϕ is complete ([Assumption 5](#)) and Σ_{cov} is full rank, then for $j \geq 0$,*

$$\|(\Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{cr}} \Sigma_{\text{cov}}^{-1/2})^j\|_{\text{op}} \leq \rho_s.$$

Proof. First, we rewrite the relevant matrix as follows,

$$\begin{aligned} (\Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{cr}} \Sigma_{\text{cov}}^{-1/2})^j &= \Sigma_{\text{cov}}^{1/2} (\Sigma_{\text{cov}}^{-1} \Sigma_{\text{cr}})^j \Sigma_{\text{cov}}^{-1/2} \\ &= \Sigma_{\text{cov}}^{-1/2} \Phi^\top \text{diag}(\mu) \Phi^\top (\Sigma_{\text{cov}}^{-1} \Sigma_{\text{cr}})^j \Sigma_{\text{cov}}^{-1/2}. \end{aligned}$$

Therefore,

$$\|(\Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{cr}} \Sigma_{\text{cov}}^{-1/2})^j\|_{\text{op}} \leq \underbrace{\|\Sigma_{\text{cov}}^{-1/2} \Phi^\top \text{diag}(\mu)^{1/2}\|_{\text{op}}}_{:=T_1} \underbrace{\|\text{diag}(\mu)^{1/2} \Phi (\Sigma_{\text{cov}}^{-1} \Sigma_{\text{cr}})^j \Sigma_{\text{cov}}^{-1/2}\|_{\text{op}}}_{:=T_2}.$$

To bound T_1 , we observe that

$$\|\Sigma_{\text{cov}}^{-1/2} \Phi^\top \text{diag}(\mu)^{1/2}\|_{\text{op}}^2 = \|(\Phi^\top \text{diag}(\mu) \Phi)^{-1/2} \Phi^\top \text{diag}(\mu)^{1/2}\|_{\text{op}}.$$

Letting $A := \text{diag}(\mu)^{1/2} \Phi$, the above expression satisfies,

$$\|(A^\top A)^{-1/2} A^\top\|_{\text{op}}^2 = \sup_{\|v\|=1} v^\top A (A^\top A)^{-1} A^\top v \leq 1,$$

since $A(A^\top A)^{-1} A^\top$ is a projection matrix. Moving onto T_2 , we recall that

$$\|\text{diag}(\mu)^{1/2} \Phi (\Sigma_{\text{cov}}^{-1} \Sigma_{\text{cr}})^j \Sigma_{\text{cov}}^{-1/2}\|_{\text{op}} = \sup_{\|v\|=1} \|\text{diag}(\mu)^{1/2} \Phi (\Sigma_{\text{cov}}^{-1} \Sigma_{\text{cr}})^j \Sigma_{\text{cov}}^{-1/2} v\|.$$

For any fixed vector v , since the entries of μ form a probability measure,

$$\|\text{diag}(\mu)^{1/2} v\| = \sqrt{\sum_{i=1}^d \mu_i v_i^2} \leq \max_i v_i = \|v\|_\infty.$$

Therefore,

$$\|\text{diag}(\mu)^{1/2} \Phi (\Sigma_{\text{cov}}^{-1} \Sigma_{\text{cr}})^j \Sigma_{\text{cov}}^{-1/2}\|_{\text{op}} \leq \sup_{\|v\|=1} \|\Phi (\Sigma_{\text{cov}}^{-1} \Sigma_{\text{cr}})^j \Sigma_{\text{cov}}^{-1/2} v\|_\infty.$$

Then, by repeatedly applying [Lemma A.5](#), we get that

$$\|\Phi (\Sigma_{\text{cov}}^{-1} \Sigma_{\text{cr}})^j \Sigma_{\text{cov}}^{-1/2} v\|_\infty \leq \|\Phi \Sigma_{\text{cov}}^{-1/2} v\|_\infty.$$

Lastly,

$$\|\Phi \Sigma_{\text{cov}}^{-1/2} v\|_\infty = \sup_{(s,a)} |\phi(s,a)^\top \Sigma_{\text{cov}}^{-1/2} v| \leq \sup_{(s,a) \in \mathcal{S} \times \mathcal{A}} \|\Sigma_{\text{cov}}^{-1/2} \phi(s,a)\| = \rho_s.$$

□

Lemma A.5. *If ϕ is complete ([Assumption 5](#)) and Σ_{cov} is full rank, then for all θ ,*

$$\|\Phi \Sigma_{\text{cov}}^{-1} \Sigma_{\text{cr}} \theta\|_\infty \leq \|\Phi \theta\|_\infty.$$

Proof. If we denote the vector of expected rewards by $\vec{r} \in \mathbb{R}^{|\mathcal{S}||\mathcal{A}|}$, then completeness implies that for all θ , there exists a θ' such that

$$\Phi \theta' = \vec{r} + \gamma P^{(\pi)} \Phi \theta.$$

Choosing $\theta = 0$, this means that there exists a vector θ_r such that $\vec{r} = \Phi\theta_r$. Consequently, we deduce that for all θ , there always exists a θ' such that $\Phi\theta' = \gamma P^{(\pi)}\Phi\theta$. Using this realizability condition, for a given distribution μ , θ' must satisfy

$$\begin{aligned}\theta' &= \arg \min_{\bar{\theta}} \mathbb{E}_{(s,a) \sim \mu, s' \sim P(\cdot|s,a)} (\phi(s,a)^\top \bar{\theta} - \gamma \phi(s',a')^\top \theta)^2 \\ &= \gamma \Sigma_{\text{cov}}^{-1} \Sigma_{\text{cr}} \theta.\end{aligned}$$

Together with the previous equation, this implies that for all θ , $\gamma \Phi \Sigma_{\text{cov}}^{-1} \Sigma_{\text{cr}} \theta = \gamma P^{(\pi)} \Phi \theta$. Thus, we conclude that

$$\begin{aligned}\|\Phi \Sigma_{\text{cov}}^{-1} \Sigma_{\text{cr}} \theta\|_\infty &= \|P^{(\pi)} \Phi \theta\|_\infty \\ &\leq \|\Phi \theta\|_\infty,\end{aligned}$$

where we have used the fact that $P^{(\pi)}$ is row stochastic so $\|P^{(\pi)}\|_1 \leq 1$. \square

Lemma A.6. *Assume that the rewards are linearly realizable in the feature mapping ϕ . That is, there exists a vector $\theta_r^* \in \mathbb{R}^d$ such that for all $(s,a) \in \mathcal{S} \times \mathcal{A}$, $\mathbb{E}r(s,a) = \phi(s,a)^\top \theta_r^*$. Then, $\|\Sigma_{\text{cov}}^{-1/2} \theta_{\phi,r}\|_2 \leq 1$.*

Otherwise, if reward realizability does not hold $\|\Sigma_{\text{cov}}^{-1/2} \theta_{\phi,r}\| \leq \sqrt{d}$.

Proof. Expanding out the definition of $\theta_{\phi,r}$,

$$\|\Sigma_{\text{cov}}^{-1/2} \theta_{\phi,r}\|^2 = \text{tr} \left[\Sigma_{\text{cov}}^{-1/2} \mathbb{E}[\phi(s,a)r(s,a)] \mathbb{E}\phi(s,a)^\top r(s,a) \Sigma_{\text{cov}}^{-1/2} \right]$$

Under realizability, $\mathbb{E}[\phi(s,a)r(s,a)] = \mathbb{E}\phi(s,a)\phi(s,a)^\top \theta_r^*$. Hence, the expression above can be rewritten as,

$$\text{tr} \left[\Sigma_{\text{cov}}^{-1} \mathbb{E}[\phi(s,a)\phi(s,a)^\top] \theta_r^* \theta_r^{*\top} \mathbb{E}\phi(s,a)\phi(s,a)^\top \right] = \mathbb{E}(\phi(s,a)^\top \theta_r^*)^2 = \mathbb{E}r(s,a)^2 \leq 1.$$

If the rewards are not linearly realizable in ϕ , then by Jensen's inequality,

$$\begin{aligned}\|\Sigma_{\text{cov}}^{-1/2} \mathbb{E}\phi(s,a)r(s,a)\|^2 &\leq \mathbb{E}\|\Sigma_{\text{cov}}^{-1/2} \phi(s,a)r(s,a)\|^2 \\ &= \mathbb{E} \text{tr} \left[\Sigma_{\text{cov}}^{-1/2} \mathbb{E}\phi(s,a)\phi(s,a)^\top r(s,a)^2 \Sigma_{\text{cov}}^{-1/2} \right] \\ &\leq \sup_{s,a} r(s,a)^2 \text{tr}[I] \\ &\leq d.\end{aligned}$$

\square

A.6 Proof of Proposition 3.4: FQI lower bound

Recall the functional form of the FQI approximation,

$$\hat{\theta}_T = \sum_{k=0}^T (\gamma \Sigma_{\text{cov}}^{-1} \Sigma_{\text{cr}})^{-1} \Sigma_{\text{cov}}^{-1} (\theta_{\phi,r} + z) = \mu + v,$$

where $\mathbb{E}\hat{\theta}_t = \mu := \sum_{k=0}^T (\gamma \Sigma_{\text{cov}}^{-1} \Sigma_{\text{cr}})^{-1} \Sigma_{\text{cov}}^{-1} \theta_{\phi,r}$ and $v := \sum_{k=0}^T (\gamma \Sigma_{\text{cov}}^{-1} \Sigma_{\text{cr}})^{-1} \Sigma_{\text{cov}}^{-1} z$. Expanding out and using $\mathbb{E}v = 0$, we have that

$$\begin{aligned}\mathbb{E}\|\hat{\theta}_T - \mathbb{E}\hat{\theta}_T\|_2^2 &= \mathbb{E}\|\hat{\theta}_T\|_2^2 - \|\mathbb{E}\hat{\theta}_T\|_2^2 \\ &= \mathbb{E}\|\mu\|_2^2 + \mathbb{E}\|v\|_2^2 - \|\mathbb{E}\hat{\theta}_T\|_2^2 \\ &= \mathbb{E}\|v\|_2^2.\end{aligned}$$

Now, letting $A = \gamma \Sigma_{\text{cov}}^{-1} \Sigma_{\text{cr}}$, we have that

$$\mathbb{E} \|v\|_2^2 = \text{tr} \left[\left(\sum_{k=0}^T A^k \right)^\top \Lambda \left(\sum_{k=0}^T A^k \right) \right] \geq \sigma_{\min}(\Lambda) \left\| \sum_{k=0}^T A^k \right\|_{\text{op}}^2 = \sigma_{\min}(\Lambda) \sup_{\|v\|_2=1} v^\top \left(\sum_{k=0}^T A^k \right)^\top \left(\sum_{k=0}^T A^k \right) v,$$

where we have used $\text{tr} [A^\top A] = \|A\|_{\text{F}}^2 \geq \|A\|_{\text{op}}^2$ ($\|\cdot\|_{\text{F}}$ denotes the Frobenius norm of a matrix) and the variational characterization of the operator norm for symmetric matrices. By assumption on the spectral radius, A has an eigenvector u with eigenvalue λ such that $|\lambda| > 1$. Therefore,

$$\sup_{\|v\|_2=1} v^\top \left(\sum_{k=0}^T A^k \right)^\top \left(\sum_{k=0}^T A^k \right) v \geq u^\top \left(\sum_{k=0}^T A^k \right)^\top \left(\sum_{k=0}^T A^k \right) u = \|u\|_2^2 \left(\sum_{k=0}^T \lambda^k \right)^2 = \left(\frac{\lambda^{T+1} - 1}{\lambda - 1} \right)^2.$$

Note that if $|\lambda| = 1$, this series can grow linearly in T (e.g if $\lambda = 1$) or oscillate (if $\lambda = -1$).

A.7 Extensions to ridge regression

One might wonder whether adding ℓ_2 regularization, that is, an $\lambda \|\theta\|_2^2$, $\lambda > 0$ additive penalty to the FQI or LSTD objective in Eq. (1.2), could help mitigate the divergence phenomenon outlined in Proposition 3.4 or the limits of linear estimators from Theorem 3.

For finite-dimensional problems with full rank covariance, typical analyses of ridge regression set the regularizer λ to shrink with the number of samples n . In this case, the ridge estimator achieves consistent parameter recovery and asymptotically returns the same solution as just performing ordinary least squares. Therefore, we can expect similar blowup if stability fails (in fact, this phenomenon is verified empirically by Wang et al. [2021b]). On the other hand, if the parameter λ is lower bounded by a constant, then ridge regression will have constant bias which will then be amplified by the number of rounds T . Hence, adding regularization does not avoid the need for stability when performing fitted Q-iteration. Similar arguments demonstrate why regularization is unlikely to overcome the limitations of least squares temporal differencing learning (or other linear estimators) in settings where invertibility does not hold.

B Supporting Arguments for Section 4: LSTD

B.1 Proof of Theorem 2: invertibility is sufficient for LSTD

Recall the closed form expression of the empirical LSTD estimator:

$$\hat{\theta}_{\text{LS}} = (I - \gamma \hat{\Sigma}_{\text{cov}}^{-1} \hat{\Sigma}_{\text{cr}})^\dagger \Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{cov}}^{1/2} \hat{\Sigma}_{\text{cov}}^{-1} \hat{\theta}_{\phi, r}.$$

Multiplying on the left by Σ_{cov}^{-1} ,

$$\begin{aligned} \Sigma_{\text{cov}}^{1/2} \hat{\theta}_{\text{LS}} &= \Sigma_{\text{cov}}^{1/2} (I - \gamma \hat{\Sigma}_{\text{cov}}^{-1} \hat{\Sigma}_{\text{cr}})^\dagger \Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{cov}}^{1/2} \hat{\Sigma}_{\text{cov}}^{-1} \hat{\theta}_{\phi, r} \\ &= \left(\Sigma_{\text{cov}}^{1/2} (I - \gamma \hat{\Sigma}_{\text{cov}}^{-1} \hat{\Sigma}_{\text{cr}}) \Sigma_{\text{cov}}^{-1/2} \right)^\dagger (\Sigma_{\text{cov}}^{1/2} \hat{\Sigma}_{\text{cov}}^{-1} \hat{\theta}_{\phi, r}), \end{aligned}$$

where we have used the identity $(ABA^{-1})^\dagger = AB^\dagger A^{-1}$ for any invertible A and B . Similarly,

$$\Sigma_{\text{cov}}^{1/2} \theta_\gamma^* = \left(\Sigma_{\text{cov}}^{1/2} (I - \gamma \Sigma_{\text{cov}}^{-1} \Sigma_{\text{cr}}) \Sigma_{\text{cov}}^{-1/2} \right)^{-1} (\Sigma_{\text{cov}}^{1/2} \Sigma_{\text{cov}}^{-1} \theta_{\phi, r}).$$

Now defining the following quantities,

$$A := \Sigma_{\text{cov}}^{1/2} (I - \gamma \Sigma_{\text{cov}}^{-1} \Sigma_{\text{cr}}) \Sigma_{\text{cov}}^{-1/2}, \quad \hat{A} := \Sigma_{\text{cov}}^{1/2} (I - \gamma \hat{\Sigma}_{\text{cov}}^{-1} \hat{\Sigma}_{\text{cr}}) \Sigma_{\text{cov}}^{-1/2}$$

$$b := \Sigma_{\text{cov}}^{1/2} \Sigma_{\text{cov}}^{-1} \theta_{\phi, r}, \quad \hat{b} := \Sigma_{\text{cov}}^{1/2} \hat{\Sigma}_{\text{cov}}^{-1} \hat{\theta}_{\phi, r}.$$

We can rewrite the above expression as:

$$\Sigma_{\text{cov}}^{1/2}(\theta_\gamma^\star - \hat{\theta}_\gamma) = (A^{-1} - \hat{A}^\dagger)b + \hat{A}^\dagger(b - \hat{b}).$$

Therefore,

$$\|\Sigma_{\text{cov}}^{1/2}(\theta_\gamma^\star - \hat{\theta}_\gamma)\|_2 \leq \|A^{-1} - \hat{A}^\dagger\|_{\text{op}}\|b\|_2 + \|\hat{A}^\dagger\|_{\text{op}}\|b - \hat{b}\|_2.$$

Using [Lemma B.1](#), since $\varepsilon_{\text{op}} \leq \frac{1}{2}\sigma_{\min}(I - \gamma\Sigma_{\text{cov}}^{-1}\Sigma_{\text{cr}})$:

$$\|\Sigma_{\text{cov}}^{1/2}(\theta_\gamma^\star - \hat{\theta}_\gamma)\|_2 \lesssim \frac{\varepsilon_{\text{op}}}{\sigma_{\min}(I - \gamma\Sigma_{\text{cov}}^{-1/2}\Sigma_{\text{cr}}\Sigma_{\text{cov}}^{-1/2})^2} \|\Sigma_{\text{cov}}^{-1/2}\theta_{\phi,r}\|_2 + \frac{\varepsilon_r}{\sigma_{\min}(I - \gamma\Sigma_{\text{cov}}^{-1/2}\Sigma_{\text{cr}}\Sigma_{\text{cov}}^{-1/2})}.$$

Lemma B.1 (Theorem 3.8 in [Stewart \[1990\]](#)). *Let $A \in \mathbb{R}^{m \times n}$, with $m \geq n$ and let $\tilde{A} = A + E$. Then*

$$\|\tilde{A}^\dagger - A^\dagger\|_{\text{op}} \leq \frac{1 + \sqrt{5}}{2} \max\{\|\tilde{A}^\dagger\|_{\text{op}}^2, \|A^\dagger\|_{\text{op}}^2\} \|E\|_{\text{op}}.$$

Furthermore, if $\|E\|_{\text{op}} \leq \frac{1}{2}\sigma_{\min}(A)$, then

$$\|\tilde{A}^\dagger - A^\dagger\|_{\text{op}} \lesssim \|A^\dagger\|_{\text{op}}^2 \|E\|_{\text{op}}.$$

B.2 Proof of [Proposition 4.1](#): Relating stability and invertibility

The first part of the proposition follows directly from [Fact 3.1](#). For the second, again using [Fact 3.1](#):

$$\begin{aligned} 1/\sigma_{\min}(I - \gamma\Sigma_{\text{cov}}^{-1/2}\Sigma_{\text{cr}}\Sigma_{\text{cov}}^{-1/2}) &= \|(I - \gamma\Sigma_{\text{cov}}^{-1/2}\Sigma_{\text{cr}}\Sigma_{\text{cov}}^{-1/2})^{-1}\|_{\text{op}} \\ &= \left\| \sum_{k=0}^{\infty} (\gamma\Sigma_{\text{cov}}^{-1/2}\Sigma_{\text{cr}}\Sigma_{\text{cov}}^{-1/2})^k \right\|_{\text{op}} \\ &\leq \sum_{k=0}^{\infty} \|(\gamma\Sigma_{\text{cov}}^{-1/2}\Sigma_{\text{cr}}\Sigma_{\text{cov}}^{-1/2})^k\|_{\text{op}} \\ &\leq \sum_{k=0}^{\infty} \text{cond}(P_\gamma)^{1/2} \left(1 - \frac{1}{\|P_\gamma\|_{\text{op}}}\right)^{k/2} \end{aligned}$$

Here, we've used [Lemma A.1](#) in the last line. The final bound follows from applying the final argument from the proof of [Theorem 1](#).

B.3 Proof of [Proposition 4.2](#): Relationship to [Mou et al. \[2020\]](#)

The result follows from the proof of Corollary 1 in [Mou et al. \[2020\]](#). We include the calculation for the sake of completeness. For any unit vector u ,

$$(1 - \kappa)\|u\|_2^2 \leq u^\top (I - \gamma\Sigma_{\text{cov}}^{-1/2}\Sigma_{\text{cr}}\Sigma_{\text{cov}}^{-1/2})u \leq \|(I - \gamma\Sigma_{\text{cov}}^{-1/2}\Sigma_{\text{cr}}\Sigma_{\text{cov}}^{-1/2})u\|_{\text{op}}\|u\|_2.$$

Therefore, $\|(I - \gamma\Sigma_{\text{cov}}^{-1/2}\Sigma_{\text{cr}}\Sigma_{\text{cov}}^{-1/2})^{-1}\|_{\text{op}} = 1/\sigma_{\min}(I - \gamma\Sigma_{\text{cov}}^{-1/2}\Sigma_{\text{cr}}\Sigma_{\text{cov}}^{-1/2}) \leq 1/(1 - \kappa)$.

B.4 Proof of [Theorem 3](#): necessity of invertibility for LSTD

We begin by proving two auxiliary claims and then move on to proving each part of the theorem separately.

Claim B.2. *If $I - \gamma\Sigma_{\text{cov}}^{-1}\Sigma_{\text{cr}}$ is rank deficient, then there exists a real vector $v \in \mathbb{R}^d$ such that:*

$$\mathbb{E}_{(s,a) \sim \mathcal{D}, s' \sim P(\cdot|s,a), a' \sim a'} \phi(s,a) \langle \gamma \cdot \phi(s',a') - \phi(s,a), v \rangle = 0.$$

Proof. The matrix being rank deficient implies that there exists a vector v such that $(I - \gamma \Sigma_{\text{cov}}^{-1} \Sigma_{\text{cr}})v = 0$, or equivalently, the matrix $\gamma \Sigma_{\text{cov}}^{-1} \Sigma_{\text{cr}}$ has an eigenvector v with eigenvalue 1. Because the matrix and eigenvalue are both real, we can also take v to be real. From here, $v = \gamma \Sigma_{\text{cov}}^{-1} \Sigma_{\text{cr}} v$. Hence, $\Sigma_{\text{cov}} v = \gamma \Sigma_{\text{cr}} v$. Expanding out the definitions of these matrices,

$$\mathbb{E} \phi(s, a) \langle \phi(s, a), v \rangle = \gamma \mathbb{E} \phi(s, a) \langle \phi(s', a'), v \rangle.$$

□

Claim B.3. For any $(s, a) \in \mathcal{S} \times \mathcal{A}$,

$$\phi(s, a) = -\mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t (\gamma \phi(s_{t+1}, a_{t+1}) - \phi(s_t, a_t)) \mid (s_0, a_0) = (s, a), \pi \right].$$

Proof. The sum telescopes and $\lim_{t \rightarrow \infty} \gamma^t \mathbb{E} \phi(s_t, a_t) = 0$. □

We conclude with the proof of [Theorem 3](#):

Alternate reward. As per the presentation of theorem, the only difference between \mathcal{M} and $\overline{\mathcal{M}}$ is the unknown reward. In particular, we define the new reward function \bar{R} as

$$R'(s, a) = R(s, a) + \langle \gamma \cdot \phi(s', a') - \phi(s, a), v \rangle \quad (\text{B.1})$$

where v is as in [Claim B.2](#), $s' \sim P(\cdot \mid s, a)$, and $a' \sim \pi(s)$.

Proof of identical moments. Since the features, offline distribution, and transitions are all the same, then $\Sigma_{\text{cov}} = \bar{\Sigma}_{\text{cov}}$ and $\Sigma_{\text{cr}} = \bar{\Sigma}_{\text{cross}}$. Next, by expanding out the new reward function:

$$\begin{aligned} \bar{\theta}_{\phi, r} &= \mathbb{E}_{(s, a) \sim \mathcal{D}} \phi(s, a) r(s, a) \\ &= \mathbb{E} \phi(s, a) r(s, a) + \mathbb{E} \phi(s, a) \langle \gamma \cdot \phi(s', a') - \phi(s, a), v \rangle \\ &= \mathbb{E} \phi(s, a) r(s, a) + 0, \end{aligned}$$

where the last line follows from [Claim B.2](#). Furthermore, because ϕ has a constant feature, again using [Claim B.2](#), it must be the case that,

$$\mathbb{E} \langle \gamma \cdot \phi(s', a') - \phi(s, a), v \rangle = 0.$$

Hence, $\mathbb{E} \bar{r}(s, a) = \mathbb{E} r(s, a)$.

Proof of realizability. Expanding out the definition of \bar{Q}^π ,

$$\begin{aligned} \bar{Q}^\pi(s, a) &= \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t \cdot r'(s_t, a_t) \mid (s_0, a_0) = (s, a), \pi \right] \\ &= \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t \cdot r(s_t, a_t) \mid (s_0, a_0) = (s, a), \pi \right] + \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t \cdot \langle \gamma \cdot \phi(s_{t+1}, \pi(s_{t+1})) - \phi(s_t, a_t), v \rangle \mid (s_0, a_0) = (s, a), \pi \right] \\ &= Q^\pi(s, a) - \phi(s, a)^\top v \\ &= \phi(s, a)^\top (\theta_\gamma^* - v), \end{aligned}$$

where in the 3rd line we have used [Claim B.3](#) and in the last one used the assumption that Q^π is linearly realizable. In short, \bar{Q}^π is linearly realizable with weight vector $\theta_\gamma^* - v$.

Proof of different Q functions By the previous part establishing the realizability of \bar{Q}^π ,

$$\mathbb{E}_{\mathcal{D}}(Q^\pi(s, a) - Q'^\pi(s, a))^2 = v^\top \Sigma_{\text{cov}} v \geq \sigma_{\min}(\Sigma_{\text{cov}}) \|v\|_2^2.$$

The precise statement follows from the fact that v has unit length.

C Concentration Analysis: Proof of Lemma 2.1

Lemma C.1 (Matrix Bernstein, [Tropp \[2012\]](#)). *Let $S_1, \dots, S_n \in \mathbb{R}^{d_1 \times d_2}$ be random, independent matrices satisfying $\mathbb{E}[S_k] = 0$, $\max\{\|\mathbb{E}[S_k S_k^\top]\|_{\text{op}}, \|\mathbb{E}[S_k^\top S_k]\|_{\text{op}}\} \leq \sigma^2$, and $\|S_k\|_{\text{op}} \leq L$ almost surely for all k . Then, with probability at least $1 - \delta$ for any $\delta \in (0, 1)$,*

$$\left\| \frac{1}{n} \sum_{k=1}^n S_k \right\|_{\text{op}} \leq \sqrt{\frac{2\sigma^2 \log((d_1 + d_2)/\delta)}{n}} + \frac{2L \log((d_1 + d_2)/\delta)}{3n}.$$

Lemma C.2 (Vector Bernstein, [Minsker \[2017\]](#)). *Let v_1, \dots, v_n be independent vectors in \mathbb{R}^d such that $\mathbb{E}v_k = 0$, $\mathbb{E}\|v_k\|^2 \leq \sigma^2$, and $\|v_k\| \leq L$ almost surely for all k . Then, with probability $1 - \delta$ for any $\delta \in (0, 1)$,*

$$\left\| \frac{1}{n} \sum_{i=1}^n v_i \right\| \leq \sqrt{\frac{2\sigma^2 \log(28/\delta)}{n}} + \frac{2L \log(28/\delta)}{3n}.$$

To shorten the notation in our concentration analysis, we use $x_i = \phi(s_i, a_i)$ and $y_i = \phi(s'_i, a'_i)$, and $r_i = r(s_i, a_i)$. With this shorthand:

$$\Sigma_{\text{cov}} = \mathbb{E}xx^\top, \quad \hat{\Sigma}_{\text{cov}} = \frac{1}{n} \sum_{i=1}^n x_i x_i^\top, \quad \Sigma_{\text{cr}} = \mathbb{E}xy^\top, \quad \hat{\Sigma}_{\text{cr}} = \frac{1}{n} \sum_{i=1}^n x_i y_i^\top, \quad (\text{C.1})$$

$$\theta_{\phi, r} = \mathbb{E}xr, \quad \hat{\theta}_{\phi, r} = \frac{1}{n} \sum_{i=1}^n x_i r_i.$$

C.1 Bounding ε_{op}

Lemma C.3. *If $n \gtrsim \rho_s^2 \log(d/\delta)$ then, with probability $1 - \delta$,*

$$\|\Sigma_{\text{cov}}^{1/2} (\gamma \hat{\Sigma}_{\text{cov}}^{-1} \hat{\Sigma}_{\text{cr}}) \Sigma_{\text{cov}}^{-1/2} - \gamma \Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{cr}} \Sigma_{\text{cov}}^{-1/2}\|_{\text{op}} \lesssim \sqrt{\frac{\max(\sigma_{\text{cr}}^2, \rho_s^2 \mathcal{C}_{\text{ds}}) \log(d/\delta)}{n}} + \frac{\max(\mathcal{C}_{\text{ds}} \rho_s^2, \rho_s \rho_{s'}) \log(d/\delta)}{n}.$$

Proof. Let $\hat{A} := \gamma \hat{\Sigma}_{\text{cov}}^{-1} \hat{\Sigma}_{\text{cr}}$. We start by using the following error decomposition,

$$\begin{aligned} & \|\Sigma_{\text{cov}}^{1/2} \hat{A} \Sigma_{\text{cov}}^{-1/2} - \gamma \Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{cr}} \Sigma_{\text{cov}}^{-1/2}\|_{\text{op}} \\ & \leq \gamma \|\Sigma_{\text{cov}}^{1/2} \hat{\Sigma}_{\text{cov}}^{-1} \Sigma_{\text{cov}}^{1/2} \cdot \Sigma_{\text{cov}}^{-1/2} (\hat{\Sigma}_{\text{cr}} - \Sigma_{\text{cr}}) \Sigma_{\text{cov}}^{-1/2}\|_{\text{op}} + \gamma \|\Sigma_{\text{cov}}^{1/2} (\hat{\Sigma}_{\text{cov}}^{-1} - \Sigma_{\text{cov}}^{-1}) \Sigma_{\text{cov}}^{1/2} \cdot \Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{cr}} \Sigma_{\text{cov}}^{-1/2}\|_{\text{op}} \\ & \leq \underbrace{\gamma \|\Sigma_{\text{cov}}^{1/2} \hat{\Sigma}_{\text{cov}}^{-1} \Sigma_{\text{cov}}^{1/2}\|_{\text{op}}}_{:=T_1} \cdot \underbrace{\|\Sigma_{\text{cov}}^{-1/2} (\hat{\Sigma}_{\text{cr}} - \Sigma_{\text{cr}}) \Sigma_{\text{cov}}^{-1/2}\|_{\text{op}}}_{:=T_2} \\ & \quad + \underbrace{\|\Sigma_{\text{cov}}^{1/2} (\hat{\Sigma}_{\text{cov}}^{-1} - \Sigma_{\text{cov}}^{-1}) \Sigma_{\text{cov}}^{1/2}\|_{\text{op}}}_{:=T_3} \cdot \underbrace{\|\gamma \Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{cr}} \Sigma_{\text{cov}}^{-1/2}\|_{\text{op}}}_{:=T_4}. \end{aligned}$$

We now bound each of these terms separately.

Bound on T_2 . We apply the Matrix Bernstein inequality on $\Sigma_{\text{cov}}^{-1/2} (\hat{\Sigma}_{\text{cr}} - \Sigma_{\text{cr}}) \Sigma_{\text{cov}}^{-1/2}$. Here we define

$$S_k = \Sigma_{\text{cov}}^{-1/2} (x_k y_k^\top - \Sigma_{\text{cr}}) \Sigma_{\text{cov}}^{-1/2}$$

which is centered and satisfies:

$$\begin{aligned} \|S_k\| &\leq \|\Sigma_{\text{cov}}^{-1/2} x_k y_k^\top \Sigma_{\text{cov}}^{-1/2}\| + \mathbb{E}_{\mathcal{D}} \|\Sigma_{\text{cov}}^{-1/2} x y^\top \Sigma_{\text{cov}}^{-1/2}\| \leq 2 \sup_{(x,y) \in \text{supp}(\mathcal{D})} \|\Sigma_{\text{cov}}^{-1/2} x y^\top \Sigma_{\text{cov}}^{-1/2}\| \\ &\leq 2 \sup_{(x,y) \in \text{supp}(\mathcal{D})} \|\Sigma_{\text{cov}}^{-1/2} x\| \cdot \|\Sigma_{\text{cov}}^{-1/2} y\| \leq 2\rho_s \rho_{s'}. \end{aligned}$$

Therefore for σ_{cr}^2 defined as in Eq. (2.6), , we get that with probability $1 - \delta$,

$$T_2 \leq \sqrt{\frac{2\sigma_{\text{cr}}^2 \log(2d/\delta)}{n}} + \frac{4\rho_s \rho_{s'} \log(2d/\delta)}{3n}.$$

Bound on T_1 and T_3 . Essentially the same argument as for the bound on T_2 reveals that,

$$\|\Sigma_{\text{cov}}^{-1/2} (\hat{\Sigma}_{\text{cov}} - \Sigma_{\text{cov}}) \Sigma_{\text{cov}}^{-1/2}\|_{\text{op}} \leq \sqrt{\frac{2\sigma_x^2 \log(2d/\delta)}{n}} + \frac{2\rho_s^2 \log(2d/\delta)}{3n} =: \tau. \quad (\text{C.2})$$

This inequality directly implies that

$$1 - \tau \leq \lambda_{\min}(\Sigma_{\text{cov}}^{-1/2} \hat{\Sigma}_{\text{cov}} \Sigma_{\text{cov}}^{-1/2}) \leq \lambda_{\max}(\Sigma_{\text{cov}}^{-1/2} \hat{\Sigma}_{\text{cov}} \Sigma_{\text{cov}}^{-1/2}) \leq 1 + \tau,$$

which in particular implies that $\Sigma_{\text{cov}}^{-1/2} \hat{\Sigma}_{\text{cov}} \Sigma_{\text{cov}}^{-1/2}$ is invertible whenever $\tau < 1/2$, a fact that is ensured by our lower bound on n . Therefore:

$$T_1 = \|\Sigma_{\text{cov}}^{1/2} \hat{\Sigma}_{\text{cov}}^{-1} \Sigma_{\text{cov}}^{1/2}\| = \frac{1}{\lambda_{\min}(\Sigma_{\text{cov}}^{-1/2} \hat{\Sigma}_{\text{cov}} \Sigma_{\text{cov}}^{-1/2})} \leq \frac{1}{1 - \tau}. \quad (\text{C.3})$$

More generally, we have that:

$$1 - 2\tau \leq \lambda_{\min}(\Sigma_{\text{cov}}^{1/2} \hat{\Sigma}_{\text{cov}}^{-1} \Sigma_{\text{cov}}^{1/2}) \leq \lambda_{\max}(\Sigma_{\text{cov}}^{1/2} \hat{\Sigma}_{\text{cov}}^{-1} \Sigma_{\text{cov}}^{1/2}) \leq 1 + 2\tau.$$

Using the fact that $1/(1 + \tau) \geq 1 - 2\tau$ and $1/(1 - \tau) \leq 1 + 2\tau$ for $\tau \leq 1/2$, this directly yields

$$T_3 = \|\Sigma_{\text{cov}}^{1/2} (\hat{\Sigma}_{\text{cov}}^{-1} - \Sigma_{\text{cov}}^{-1}) \Sigma_{\text{cov}}^{1/2}\| \leq 2\tau. \quad (\text{C.4})$$

Thus, we have bounded T_1 and T_3 . In particular, for $\tau < 1/2$, $T_1 \leq 2$, and $T_3 \leq 2\tau$.

Bound on T_4 . For T_4 , no concentration argument is required. Instead, a Schur complement argument implies that,

$$\|\Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{cr}} \Sigma_{\text{cov}}^{-1/2}\|_{\text{op}}^2 \leq \|\Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{next}} \Sigma_{\text{cov}}^{-1/2}\|_{\text{op}} \leq \mathcal{C}_{\text{ds}},$$

where we've used $\Sigma_{\text{next}} \preceq \mathcal{C}_{\text{ds}} \Sigma_{\text{cov}}$.

Wrapping up. Taking a union bound, we obtain that

$$\varepsilon_{\text{op}} \lesssim \sqrt{\frac{\max(\sigma_{\text{cr}}^2, \sigma_{\text{cov}}^2 \mathcal{C}_{\text{ds}}) \log(d/\delta)}{n}} + \frac{\max(\mathcal{C}_{\text{ds}} \rho_s^2, \rho_s \rho_{s'}) \log(d/\delta)}{n}.$$

□

C.2 Bounding ε_r

Lemma C.4. *If $n \gtrsim \rho_s^2 \log(d/\delta)$ then, with probability $1 - \delta$,*

$$\|\Sigma_{\text{cov}}^{1/2} \widehat{\Sigma}_{\text{cov}}^{-1} \widehat{\theta}_{\phi,r} - \Sigma_{\text{cov}}^{-1/2} \theta_{\phi,r}\| \lesssim \sqrt{\frac{\max(\|\Sigma_{\text{cov}}^{-1/2} \theta_0^*\|^2 \sigma_{\text{cov}}^2, d) \log(d/\delta)}{n}} + \frac{\|\Sigma_{\text{cov}}^{-1/2} \theta_{\phi,r}\| 2\rho_s^2 \log(d/\delta)}{n}.$$

Proof. The ideas are very similar to Lemma C.3. In this case, the relevant error decomposition is,

$$\begin{aligned} \varepsilon_r &= \|\Sigma_{\text{cov}}^{1/2} \widehat{\Sigma}_{\text{cov}}^{-1} \widehat{\theta}_{\phi,r} - \Sigma_{\text{cov}}^{-1/2} \theta_{\phi,r}\| \\ &\leq \underbrace{\|\Sigma_{\text{cov}}^{1/2} \widehat{\Sigma}_{\text{cov}}^{-1} \Sigma_{\text{cov}}^{1/2}\|_{\text{op}}}_{:=T_1} \underbrace{\|\Sigma_{\text{cov}}^{-1/2} (\theta_{\phi,r} - \widehat{\theta}_{\phi,r})\|}_{:=T_2} + \underbrace{\|(\Sigma_{\text{cov}}^{1/2} \widehat{\Sigma}_{\text{cov}}^{-1} \Sigma_{\text{cov}}^{1/2} - I)\|_{\text{op}}}_{:=T_3} \|\Sigma_{\text{cov}}^{-1/2} \theta_{\phi,r}\|. \end{aligned}$$

Bound on T_1 and T_3 . Whenever τ , defined as in Eq. (C.2), is strictly less than $1/2$, the analysis therein (in particular, Eq. (C.4) and Eq. (C.3)) proves that $T_1 \leq 2$ and $T_3 \leq 2\tau$.

Bound on T_2 . We apply the vector Bernstein inequality, Lemma C.2, on the vectors

$$v_i = \Sigma_{\text{cov}}^{-1/2} x_i r_i - \Sigma_{\text{cov}}^{-1/2} \theta_{\phi,r}.$$

Note that, since the rewards have magnitude bounded by 1,

$$\sup_i \|v_i\| \leq \sup_i \|\Sigma_{\text{cov}}^{-1/2} x_i r_i\| + \|\Sigma_{\text{cov}}^{-1/2} \theta_{\phi,r}\| \leq \|\Sigma_{\text{cov}}^{-1/2} x_i r_i\| + \mathbb{E} \|\Sigma_{\text{cov}}^{-1/2} x r\| \leq 2\rho_s$$

and,

$$\mathbb{E} \|v_i\|^2 = \mathbb{E} \|\Sigma_{\text{cov}}^{-1/2} x_i r_i\|^2 - \|\Sigma_{\text{cov}}^{-1/2} \theta_{\phi,r}\|^2 \leq \text{tr} \left[\Sigma_{\text{cov}}^{-1/2} \mathbb{E} r_i^2 x_i x_i^\top \Sigma_{\text{cov}}^{-1/2} \right] \leq d.$$

Applying vector Bernstein,

$$T_2 \leq \sqrt{\frac{2d \log(28/\delta)}{n}} + \frac{4\rho_s \log(28/\delta)}{3n}.$$

Wrapping up. Combining these, we get that,

$$\varepsilon_r \lesssim \sqrt{\frac{\max(\|\Sigma_{\text{cov}}^{-1/2} \theta_0^*\|^2 \sigma_{\text{cov}}^2, d) \log(d/\delta)}{n}} + \frac{\|\Sigma_{\text{cov}}^{-1/2} \theta_{\phi,r}\| 2\rho_s^2 \log(d/\delta)}{n}.$$

□

C.3 Bounding variances

Bounding σ_{cr}^2 . Again using the notation from Eq. (C.1), and letting

$$S_k = \Sigma_{\text{cov}}^{-1/2} (x_k y_k^\top - \Sigma_{\text{cr}}) \Sigma_{\text{cov}}^{-1/2}$$

bounding σ_{cr}^2 is equivalent to bounding the operator norms of:

$$\begin{aligned} \mathbb{E}[S_k S_k^\top] &= \mathbb{E}[\|\Sigma_{\text{cov}}^{-1/2} y\|_2^2 (\Sigma_{\text{cov}}^{-1/2} x)(\Sigma_{\text{cov}}^{-1/2} x)^\top] - \Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{cr}} \Sigma_{\text{cr}}^\top \Sigma_{\text{cov}}^{-1/2} \\ \mathbb{E}[S_k^\top S_k] &= \mathbb{E}[\|\Sigma_{\text{cov}}^{-1/2} x\|_2^2 (\Sigma_{\text{cov}}^{-1/2} y)(\Sigma_{\text{cov}}^{-1/2} y)^\top] - \Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{cr}}^\top \Sigma_{\text{cr}} \Sigma_{\text{cov}}^{-1/2}. \end{aligned}$$

We will subsequently show that, for any vector $v \in \mathbb{R}^d$, we have

$$v^\top (\mathbb{E}[S_k S_k^\top]) v \geq 0, \quad v^\top (\mathbb{E}[S_k^\top S_k]) v \geq 0. \quad (\text{C.5})$$

Additionally, for any random variables $(a, b) \in \mathbb{R} \times \mathbb{R}^d$ from some joint distribution, Holder's inequality implies that

$$\begin{aligned}\|\mathbb{E}[a^2 b b^\top]\|_{\text{op}} &= \sup_{v, \|v\|_2=1} \mathbb{E}[a^2 (v^\top b)^2] \leq \min\{\sup\{a\} \sup_v \mathbb{E}[(v^\top b)^2], \sup_{b,v}\{(v^\top b)^2\} \mathbb{E}[a^2]\} \\ &= \min\{\sup\{a\} \|\mathbb{E}[b b^\top]\|_{\text{op}}, \sup\{\|b\|_2^2\} \mathbb{E}[a^2]\}.\end{aligned}$$

Using these two facts and positive semi-definiteness, we have that

$$\|\mathbb{E}[S_k S_k^\top]\|_{\text{op}} \leq \|\mathbb{E}[\|\Sigma_{\text{cov}}^{-1/2} y\|_2^2 (\Sigma_{\text{cov}}^{-1/2} x) (\Sigma_{\text{cov}}^{-1/2} x)^\top]\|_{\text{op}} \leq \sup_y \|\Sigma_{\text{cov}}^{-1/2} y\|_2^2 \|\mathbb{E}[(\Sigma_{\text{cov}}^{-1/2} x) (\Sigma_{\text{cov}}^{-1/2} x)^\top]\| \leq \rho_{s'}^2.$$

Essentially the same proof yields a similar bound on $\|\mathbb{E}[S_k^\top S_k]\|$:

$$\|\mathbb{E}[S_k^\top S_k]\|_{\text{op}} \leq \|\mathbb{E}[\|\Sigma_{\text{cov}}^{-1/2} x\|_2^2 (\Sigma_{\text{cov}}^{-1/2} y) (\Sigma_{\text{cov}}^{-1/2} y)^\top]\|_{\text{op}} \leq \rho_0^2 \|\Sigma_{\text{cov}}^{-1/2} \Sigma_{\text{next}} \Sigma_{\text{cov}}^{-1/2}\|_{\text{op}} \leq \rho_s^2 \mathcal{C}_{\text{ds}}.$$

Alternatively, we can get

$$\begin{aligned}\|\mathbb{E}[S_k^\top S_k]\|_{\text{op}} &\leq \|\mathbb{E}[\|\Sigma_{\text{cov}}^{-1/2} x\|_2^2 (\Sigma_{\text{cov}}^{-1/2} y) (\Sigma_{\text{cov}}^{-1/2} y)^\top]\|_{\text{op}} \leq \mathbb{E}[\|\Sigma_{\text{cov}}^{-1/2} x\|_2^2 \|(\Sigma_{\text{cov}}^{-1/2} y) (\Sigma_{\text{cov}}^{-1/2} y)^\top\|] \\ &= \mathbb{E}[\|\Sigma_{\text{cov}}^{-1/2} x\|_2^2 \|\Sigma_{\text{cov}}^{-1/2} y\|_2^2] \leq \rho_s^2 d.\end{aligned}$$

Let us now verify (C.5). Rebinding $\tilde{x} = \Sigma_{\text{cov}}^{-1/2} x$, $\tilde{y} = \Sigma_{\text{cov}}^{-1/2} y$, we have

$$v^\top (\mathbb{E}[S_k S_k^\top]) v = \mathbb{E}[(v^\top \tilde{x})^2 \|\tilde{y}\|_2^2] - (\mathbb{E}(v^\top \tilde{x}) \tilde{y})^\top (\mathbb{E}(v^\top \tilde{x}) \tilde{y}) = \mathbb{E}[(v^\top \tilde{x}) \tilde{y}]_2^2 - \|\mathbb{E}[(v^\top \tilde{x}) \tilde{y}]\|_2^2 \geq 0,$$

where the last inequality is by convexity. In conclusion,

$$\sigma_{\text{cr}}^2 \leq \max(\rho_{s'}^2, \min(\rho_s^2 \mathcal{C}_{\text{ds}}, \rho_s^2 d)).$$

Bounding σ_{cov}^2 . For $\tilde{x} = \Sigma_{\text{cov}}^{-1/2} \phi(s, a)$, the variance σ_{cov}^2 is equal to

$$\sigma_{\text{cov}}^2 = \|\mathbb{E}[\tilde{x} \tilde{x}^\top] \tilde{x} \tilde{x}^\top - I\|_{\text{op}} = \|\mathbb{E}[\|\tilde{x}\|_2^2 \tilde{x} \tilde{x}^\top - I]\|_{\text{op}}.$$

While this quantity is always less than ρ_s^2 , one can achieve tighter bounds if the offline distribution is *hypercontractive* as per the following definition:

Definition C.1. A distribution \mathcal{D} over random vectors x is *L8-L2 hypercontractive* if there exists a positive constant L such that for all unit vectors u ,

$$\mathbb{E}_{x \sim \mathcal{D}}((x - \mathbb{E}x)^\top u)^8 \leq L^2 (\mathbb{E}_{x \sim \mathcal{D}}((x - \mathbb{E}x)^\top u)^2)^4.$$

Gaussians or strongly log-concave distributions are some examples of probability measures that satisfy this condition. If $\Sigma_{\text{cov}}^{-1/2} \phi(s, a)$ is *L8-L2 hypercontractive*, then one can show that

$$\sigma_{\text{cov}}^2 \lesssim L \text{tr}[I + \mu \mu^\top] \|I + \mu \mu^\top\|_{\text{op}},$$

where $\mu := \Sigma_{\text{cov}}^{-1/2} \mathbb{E}_{(s,a) \sim \mathcal{D}} \phi(s, a)$. We point the interested reader to Lemma A.3 in Cherapanamjeri et al. [2020] for a more formal derivation.