

OPERATOR SHIFTING FOR MODEL-BASED POLICY EVALUATION*

XUN TANG[†], LEXING YING[‡], AND YUHUA ZHU[§]

Abstract. In model-based reinforcement learning, the transition matrix and reward vector are often estimated from random samples subject to noise. Even if the estimated model is an unbiased estimate of the true underlying model, the value function computed from the estimated model is biased. We introduce an operator shifting method for reducing the error introduced by the estimated model. When the error is in the residual norm, we prove that the shifting factor is always positive and upper bounded by $1 + O(1/n)$, where n is the number of samples used in learning each row of the transition matrix. We also propose a practical numerical algorithm for implementing the operator shifting.

Keywords. Operator shifting; Model-based Reinforcement Learning; policy evaluation; noisy matrices.

AMS subject classifications. 90C40; 15B51.

1. Introduction

Reinforcement learning (RL) has received much attention following recent successes, such as AlphaGo and AlphaZero [25, 26]. One of the fundamental problems of RL is policy evaluation [29]. When the transition dynamics are unknown, one learns the dynamics model from observed data in model-based RL. However, even if the learned model is an unbiased estimate of the true dynamics, the policy evaluation under the learned model is biased. The question of interest in this paper is whether one can increase the accuracy of the policy evaluation given an estimated dynamics model.

We consider a discounted Markov decision process (MDP) $\mathcal{M} = (\mathbb{S}, \mathbb{A}, P, r, \gamma)$ with discrete state space \mathbb{S} and discrete action space \mathbb{A} . $|\mathbb{S}|$ and $|\mathbb{A}|$ are used to denote the size of \mathbb{S} and \mathbb{A} , respectively. P is a third-order tensor, where for each action $a \in \mathbb{A}$, $P^a \in \mathbb{R}^{|\mathbb{S}| \times |\mathbb{S}|}$ is the transition matrix between the states. r is a second-order tensor such that $r_{s,a}$ is the reward at state $s \in \mathbb{S}$ if action $a \in \mathbb{A}$ is taken. Finally, $\gamma \in (0, 1)$ is the discount factor. A policy π is a second-order tensor, where for each state $s \in \mathbb{S}$, π_s represents the probability distribution over \mathbb{A} . At each time step t , one observes a state $s_t \in \mathbb{S}$ and takes an action $a_t \in \mathbb{A}$ according to the policy π_{s_t} . The environment returns the next state s_{t+1} according to the distribution $P_{s_t}^{a_t}$, and an associated reward r_{s_t, a_t} . The state value function $v^\pi \in \mathbb{R}^{|\mathbb{S}|}$ is the expected discounted cumulative reward if one starts from an initial state s and follows a policy π , i.e., the s -th component is

$$v_s^\pi = \mathbb{E}_{\substack{a_t \sim \pi_{s_t} \\ s_{t+1} \sim P_{s_t}^{a_t}}} \left[\sum_{t \geq 0} \gamma^t r_{s_t, a_t} \mid s_0 = s \right].$$

Given a policy π , the goal of policy evaluation in MDP is to solve for v^π . Let $b^\pi \in \mathbb{R}^{|\mathbb{S}|}$, $P^\pi \in \mathbb{R}^{|\mathbb{S}| \times |\mathbb{S}|}$ be the reward vector and the transition matrix under policy π ,

*Received: October 04, 2022; Accepted (in revised form): March 02, 2023. Communicated by Jianfeng Lu.

[†]Institute for Computational and Mathematical Engineering, Stanford University, Stanford, CA 94305, USA (xuntang@stanford.edu).

[‡]Department of Mathematics and Institute for Computational and Mathematical Engineering, Stanford University, Stanford, CA 94305, USA (lexing@stanford.edu).

[§]Department of Mathematics and Halicioğlu Data Science Institute, University of California - San Diego, La Jolla, CA 92093, USA (yuz244@ucsd.edu).

i.e.,

$$b^\pi = \sum_a r_{sa} \pi_s^a, \quad P^\pi = \sum_a P_{ss'}^a \pi_s^a. \quad (1.1)$$

The value function v^π satisfies the Bellman equation [29] $(I - \gamma P^\pi)v^\pi = b^\pi$. For notational simplicity, we drop the dependency on π and write this system as

$$(I - \gamma P)v = b. \quad (1.2)$$

In practice, the true transition matrix P and the reward vector b are often inaccessible. In the model-based RL, one approximates the transition matrix P and the reward vector b by the empirical data \hat{P} and \hat{b} estimated from samples, respectively [6, 16, 22, 28, 31]. A naive approach is to solve

$$(I - \gamma \hat{P})\hat{v} = \hat{b}. \quad (1.3)$$

Even if \hat{P} and \hat{b} are unbiased estimates for P and b , $\hat{v} = (I - \gamma \hat{P})^{-1} \hat{b}$ is a biased estimate for v , i.e., $\mathbb{E}_{\hat{P}, \hat{b}} \hat{v} \neq v$.

The operator shifting idea was introduced in [9, 10] to address this issue. The paper [10] considers the noisy symmetric elliptic systems, while the follow-up paper [9] addresses the asymmetric setting under the assumption that \hat{b} is isotropic, i.e., $\mathbb{E}[\hat{b}\hat{b}^\top] = I$. However, this isotropic condition often fails to hold in RL. In this paper, we extend the operator shifting framework to general MDPs of form (1.2). When applying this framework to the MDP setting, we add an appropriately chosen matrix \hat{K} to the operator $(I - \gamma \hat{P})^{-1}$ so that the shifted estimate $\tilde{v} = \left[(I - \gamma \hat{P})^{-1} - \beta \hat{K} \right] \hat{b}$ is a better estimate than \hat{v} in the sense that,

$$\mathbb{E}_{\hat{P}, \hat{b}} \|\tilde{v} - v\|^2 < \mathbb{E}_{\hat{P}, \hat{b}} \|\hat{v} - v\|^2 \quad (1.4)$$

for a certain norm $\|\cdot\|$.

Contributions. We derive a stable shifted operator for model-based policy evaluation without assumptions on the underlying transition dynamics or reward vectors. When the approximated transition matrix \hat{P} follows the multinomial distribution and n samples are used to learn each row of the transition matrix P ,

- we prove that the optimal shifting factor is always positive and upper bounded by $1 + O(\frac{1}{n})$ for any P and b , which guarantees the stability of the shifted operator, and
- we propose a numerical algorithm to find the optimal shifting factor, which is more efficient and accurate than the bootstrapping method proposed in [10].

Related work. Our problem is a special instance of the larger field of uncertainty quantification (UQ). In most UQ problems, one assumes that the operator (linear or non-linear) and the source term are generated from known distributions, and the task is to estimate certain quantities (such as moments, tail bounds) of the distribution of the solution. A large variety of numerical methods have been developed in UQ for this purpose in the last two decades [8, 11, 13, 15, 17, 23, 27, 33], including Monte-Carlo and quasi Monte-Carlo methods [4, 7, 12, 18, 20], stochastic collocation methods [1, 3, 21, 34], stochastic Galerkin methods [2, 5, 19, 35], and etc. The problem that we face is somewhat different: since the true P and b are unknown, one does not know the distributions of the empirical data \hat{P} or \hat{b} . As a result, the solution relies more on statistical techniques such as shrinkage [14] rather than the traditional UQ techniques.

Contents. Section 2 derives the oracle estimators and proves their lower and upper bounds. Section 3 proposes a practical estimator and demonstrates its performance with a few numerical examples. Section 4 contains all the proofs of the main results.

2. Operator shifting for policy evaluation

2.1. Problem setup. As mentioned above, P and b are the *unknown* underlying transition matrix and reward vector, while \hat{P} and \hat{b} are unbiased estimates for P and b , respectively. For notational convenience, we introduce A and \hat{A}

$$A = (I - \gamma P), \quad \hat{A} = (I - \gamma \hat{P}). \tag{2.1}$$

Since the transition dynamics is not symmetric in general, both P and A are non-symmetric. The norm of interest is a slightly generalized version of the residual norm

$$\|x\|_M^2 = x^\top A^\top M A x, \tag{2.2}$$

where M is a symmetric positive definite matrix. This paper mainly discusses two cases: (1) $M = I$, which means $\|\cdot\|_M$ is the usual residual norm, and (2) $M = A^{-\top} A^{-1}$, which means $\|\cdot\|_M$ is the l_2 norm.

In this paper, we choose the shifting matrix $\hat{K} = \hat{A}^{-1}$, which implies that the shifted estimate is $(1 - \beta)\hat{A}^{-1}\hat{b}$. By using $\varepsilon = (1 - \beta)$ instead as the shifting parameter, one can write the above estimate as $\tilde{v}_\varepsilon \equiv \varepsilon \hat{A}^{-1}\hat{b}$, and the objective is to minimize the following mean square error over ε ,

$$\text{MSE}(\varepsilon) \equiv \mathbb{E}_{\hat{P}, \hat{b}} \|v - \tilde{v}_\varepsilon\|_M^2. \tag{2.3}$$

The minimizer ε^* to (2.3) is referred as the *optimal shifting factor*.

Since (2.3) is a quadratic minimization, one can explicitly write out the optimal shifting factor ε^* :

$$\varepsilon^* = \frac{\mathbb{E}_{\hat{P}, \hat{b}} [b^\top M A \hat{v}]}{\mathbb{E}_{\hat{P}, \hat{b}} [\hat{v}^\top A^\top M A \hat{v}]}, \quad \text{where } \hat{v} = \hat{A}^{-1}\hat{b}. \tag{2.4}$$

The oracle estimate (2.4) is not easy to work with as it depends on the unknown matrix A . Our immediate goal is to derive a closed-form approximation of ε^* , which is accurate and allows for efficient implementation. To achieve this, we introduce a second-order approximation ε° to ε^* . We show that ε° takes a simple closed-form without approximating any expectations under the following mild assumption:

ASSUMPTION 2.1. *The i -th row \hat{p}_i^\top of \hat{P} is a random vector $\frac{1}{n_i} X_i$, where n_i is the number of samples for state i and X_i follows the multinomial distribution with $\mathbb{E}[X_i] = n_i p_i^\top$. Moreover, X_i is independent from X_j whenever $i \neq j$. The i -th entry of \hat{b} is an average of observed reward at state i .*

The part of Assumption 2.1 on the estimation of P is equivalent to that \hat{p}_i^\top follows the normalized multinomial distribution, which holds when a tabular maximum likelihood model [28] is used to estimate the transition dynamics P . That is, one generates sufficiently many transitions according to P and lets $\hat{P}_{ii'} = n_{ii'}/n_i$, where $n_{ii'}$ denotes the number of transitions observed from i to i' , and $n_i = \sum_{i'} n_{ii'}$.

Throughout this paper, we assume for simplicity that the number of samples n_i of each state is the same, i.e., for any $i \in \mathbb{S}$, $n_i \equiv n$. The sample size n plays an important

role in determining the magnitude of the operator shifting factor ε^* and the performance of the operator shifting algorithm. If the value of n_i depends on $i \in \mathbb{S}$, all the theoretical results still hold with slight modification (see Remark 2.1 for details).

2.2. Second-order approximation for ε^* . To simplify the discussion, we introduce \hat{Z} and \hat{Y}

$$\hat{Z} = A - \hat{A} = \gamma(P - \hat{P}), \quad \hat{Y} = \hat{Z}A^{-1} = (A - \hat{A})A^{-1}, \tag{2.5}$$

where A and \hat{A} are defined in (2.1). Some basic algebraic manipulations lead to the following lemma.

LEMMA 2.1. *When $\mathbb{E}[\hat{b}] = b$, the optimal shifting factor ε^* defined in (2.4) has the form*

$$\varepsilon^* = \frac{\mathbb{E}_{\hat{P}} \left[b^\top M (I - \hat{Y})^{-1} b \right]}{\mathbb{E}_{\hat{P}} \left[\text{tr} \left((\text{cov}[\hat{b}] + b^\top b) (I - \hat{Y})^{-\top} M (I - \hat{Y})^{-1} \right) \right]}, \tag{2.6}$$

where \hat{Y} is defined in (2.5). Moreover, if the values of the reward at state s and s' (i.e. \hat{b}_s and $\hat{b}_{s'}$) are uncorrelated, the matrix $\text{cov}[\hat{b}]$ is diagonal.

Next, we approximate the value of ε^* using a Neumann expansion of the matrix $(I - \hat{Y})^{-1}$

$$(I - \hat{Y})^{-1} = I + \hat{Y} + \hat{Y}^2 + O\left(\frac{\rho(\hat{Y})^3}{1 - \rho(\hat{Y})}\right), \tag{2.7}$$

when the spectral radius $\rho(\hat{Y}) < 1$. In fact, a modest requirement on n guarantees $\rho(\hat{Y}) < 1$ with high probability, as shown in Appendix A.1. The denominator term in (2.6) admits the approximation

$$(I - \hat{Y})^{-\top} M (I - \hat{Y})^{-1} \approx M + M(\hat{Y} + \hat{Y}^2) + (\hat{Y}^\top + (\hat{Y}^\top)^2)M + \hat{Y}^\top M \hat{Y}. \tag{2.8}$$

Assumption 2.1 implies $\mathbb{E}[\hat{Y}] = \mathbb{E}[\hat{Y}^\top] = 0$ as a simple consequence of \hat{P} being an unbiased estimator. Therefore, after taking an expectation, the first order terms of \hat{Y} in (2.7) and (2.8) disappear.

We can further approximate the shifting factor ε^* by expanding $(I - \hat{Y})^{-1}$ in the numerator and denominator of (2.6) up to the second order in \hat{Y} . When Assumption 2.1 holds and $\rho(\hat{Y}) < 1$, the approximated optimal shifting factor ε^* defined in (2.4) has a second-order approximation

$$\varepsilon^* \approx \varepsilon^\circ \equiv \frac{\mathbb{E}_{\hat{P}} \left[b^\top \left(M + \frac{M\hat{Y}^2 + (\hat{Y}^\top)^2 M}{2} \right) b \right]}{\mathbb{E}_{\hat{P}} \left[\text{tr} \left((\text{cov}[\hat{b}] + b^\top b) \left(M + \hat{Y}^\top M \hat{Y} + M\hat{Y}^2 + (\hat{Y}^\top)^2 M \right) \right) \right]}. \tag{2.9}$$

The derivation of (2.9) is deferred to Section 4.2.

Under Assumption 2.1, this second-order approximation can be written in a form without explicit expectation. This expectation-free form depends on the transition matrix P , the expected reward b , and the reward covariance $\text{cov}[\hat{b}]$. Let \hat{p}_i be random vectors corresponding to the i -th row of \hat{P} and $p_i = \mathbb{E}[\hat{p}_i]$, i.e., $\{\hat{p}_i\}_{i=1}^{|\mathcal{S}|}$ and $\{p_i\}_{i=1}^{|\mathcal{S}|}$ are the row vectors of \hat{P} and P , respectively:

$$\hat{P} = \begin{bmatrix} \hat{p}_1^\top \\ \vdots \\ \hat{p}_{|\mathcal{S}|}^\top \end{bmatrix}, \quad P = \begin{bmatrix} p_1^\top \\ \vdots \\ p_{|\mathcal{S}|}^\top \end{bmatrix}. \tag{2.10}$$

THEOREM 2.1. *The second-order approximation ε° in (2.9) admits the expectation-free form*

$$\varepsilon^\circ = \theta(b, P) \equiv \frac{b^\top (M + H/2)b}{b^\top (M + G + H)b + \text{tr}(\text{cov}[\hat{b}](M + G + H))}, \tag{2.11}$$

where

$$\begin{cases} B_i = \frac{1}{n} (\text{diag } p_i - p_i p_i^\top); \\ G = \mathbb{E}_{\hat{P}} [\hat{Y}^\top M \hat{Y}] = \gamma^2 A^{-\top} \left(\sum_{i=1}^{|\mathcal{S}|} [M]_{ii} B_i \right) A^{-1}; \\ H = \mathbb{E}_{\hat{P}} [(\hat{Y}^\top)^2 M + M \hat{Y}^2] \\ = \gamma^2 \left[\sum_{i=1}^{|\mathcal{S}|} A^{-\top} B_i A^{-1} \text{diag}(e_i) \right] M + \gamma^2 M \left[\sum_{i=1}^{|\mathcal{S}|} \text{diag}(e_i) A^{-\top} B_i A^{-1} \right]. \end{cases} \tag{2.12}$$

Here $\text{diag}(e_i) \in \mathbb{R}^{|\mathcal{S}| \times |\mathcal{S}|}$ is a matrix with elements 0 except for 1 on the (i, i) -th entry, p_i are row vectors of P as defined in (2.10), and the matrices A and \hat{Y} depend on P .

The proof of the above theorem is given in Section 4.2.

REMARK 2.1. Theorem 2.1 still holds under conditions weaker than Assumption 2.1. Assuming the rows of \hat{P} and entries for \hat{b} are independent unbiased estimators, then the second-order approximation in (2.9) is still valid. Moreover, the expectation-free form in (2.11) holds when one replaces the definition of B_i in (2.12) by $B_i = \text{cov}[\hat{p}_i]$. This slightly more general statement is presented in Lemma 4.1, from which Theorem 2.1 is derived as a special case. In particular, if the state i receives n_i samples, then (2.11) will still hold with B_i in (2.12) replaced by $\frac{1}{n_i} (\text{diag } p_i - p_i p_i^\top)$.

Theorem 2.1 also proves that the choice of ε° is asymptotically as powerful as ε^* with $n \rightarrow \infty$. For $\text{MSE}(\varepsilon)$ defined in (2.3), the following estimation holds, with the proof deferred to Section 4.3.

LEMMA 2.2. *The MSE in (2.3) can be approximated by*

$$\text{MSE}(\varepsilon) = (1 - \varepsilon)^2 \|b\|_M^2 + (g + h + t)\varepsilon^2 - h\varepsilon + O\left(n^{-\frac{3}{2}}\right), \tag{2.13}$$

where $g = b^\top G b, h = b^\top H b, t = \text{tr}(\text{cov}[\hat{b}](M + G + H))$ with all symbols defined in Theorem 2.1. In addition,

$$\varepsilon^* - \varepsilon^\circ = O\left(n^{-\frac{3}{2}}\right). \tag{2.14}$$

The relative error reduction factor $\eta \equiv \frac{\text{MSE}(1) - \text{MSE}(\varepsilon^\circ)}{\text{MSE}(1)}$, with MSE representing the mean square error defined in (2.3), is a useful measure for improvements. Below is a corollary of Lemma 2.2 regarding η .

COROLLARY 2.1. *Define relative error reduction factor as $\eta \equiv \frac{\text{MSE}(1) - \text{MSE}(\varepsilon^\circ)}{\text{MSE}(1)}$. For n sufficiently large, η is positive, and decays as follows*

$$\eta = O(1/n), \quad (2.15)$$

where MSE is defined in (2.3).

The proof is given in Section 4.3. The numerical results also verify the relationship in the above corollary.

2.3. Lower and upper bounds for ε° . In this section, we aim to provide bounds to show that ε° will approximately fall in the $(0, 1)$ range. Throughout this subsection, we conduct the analysis in the residual norm, which is the case with $M = I$. We present an upper bound and a lower bound for ε° in Theorem 2.2. The relevant parameters are $|\mathbb{S}|$, n and γ , which are the number of states, the number of samples per state and the discount factor, respectively.

THEOREM 2.2. *Let $p_M = \max P_{i,j}$ and $b_M = \max_i |b_i|$. If $\frac{p_M}{n} \frac{\gamma^2}{(1-\gamma)^2} \left(\frac{(1-\gamma)}{\gamma} + \frac{\sqrt{|\mathbb{S}|} b_M}{\|b\|_2} \right)^2 \leq \frac{1}{2}$, then ε° is bounded by*

$$0 < \varepsilon^\circ \leq 1 + \frac{p_M}{n} \frac{\gamma^2}{(1-\gamma)^2} \left((1-\gamma) + \gamma \frac{\sqrt{|\mathbb{S}|} b_M}{\|b\|_2} \right)^2.$$

The bound comes from technique using the spectral structure of the covariance matrix of a multinomial distribution and a tight bound for $(I - \gamma P)^{-1}b$. We defer the proof to Section 4.

We now discuss the implication of Theorem 2.2. The reward vector b is *spread* if $\frac{\max_i |b_i|}{\|b\|_2} \sim O(|\mathbb{S}|^{-1/2})$. Similarly, the transition matrix P is *spread* if $\max P_{ij} \sim o(|\mathbb{S}|^{-1/2})$. If both b and P are spread, it follows that $\varepsilon^\circ - 1$ is upper bounded by a term in $O(n^{-1} |\mathbb{S}|^{-1/2})$.

3. Practical algorithm

3.1. Algorithm. In practice, we do not have direct access to P or b . Therefore, the second-order estimate ε° derived in (2.11) is an oracle estimator. One can address this issue by bootstrapping the distribution of \hat{P} . More specifically, let P be a transition matrix and denote $\mathcal{M}_n(P)$ as the normalized multinomial distribution that the estimated transition matrix \hat{P} follows according to Assumption 2.1. Since one only has access to a single observation \hat{P} , $\mathcal{M}_n(P)$ is approximated by $\mathcal{M}_n(\hat{P})$ in the numerical implementation. In the usual bootstrapping procedure, one needs to simulate i.i.d. samples $\{\hat{P}_{(j)}\}_{j=1}^l \sim \mathcal{M}_n(\hat{P})$. By setting $\tilde{Y}_{(j)} = \gamma(\hat{P}_{(j)} - \hat{P})\hat{A}^{-1}$ and following the form in Theorem 2.1, one can approximate ε° in (2.9) by replacing the expectation with an empirical mean:

$$\varepsilon^\circ \approx \frac{\hat{b}^\top (M + \tilde{H}/2)\hat{b}}{\hat{b}^\top (M + \tilde{G} + \tilde{H})\hat{b} + \text{tr}(\tilde{\Sigma}(M + \tilde{G} + \tilde{H}))}, \quad (3.1)$$

with

$$\begin{cases} \tilde{G} = \frac{1}{l} \sum_{j=1}^l \tilde{Y}_{(j)}^\top M \tilde{Y}_{(j)} \approx \mathbb{E}_{\hat{P}} \left[\hat{Y}^\top M \hat{Y} \right] = G; \\ \tilde{H} = \frac{1}{l} \sum_{j=1}^l M \tilde{Y}_{(j)}^2 + \left(\tilde{Y}_{(j)}^\top \right)^2 M \approx \mathbb{E}_{\hat{P}} \left[M \hat{Y}^2 + (\hat{Y}^\top)^2 M \right] = H; \\ \tilde{\Sigma} \approx \text{cov} \left[\hat{b} \right]. \end{cases} \tag{3.2}$$

However, there is a major drawback to this scheme. In addition to the error caused by the difference between $\mathcal{M}_n(\hat{P})$ and $\mathcal{M}_n(P)$, the scheme introduces additional errors due to the empirical mean in place of the expectation. The empirical mean errors $\tilde{G} - G$ and $\tilde{H} - H$ are of order $O(l^{-1/2})$. In addition, the procedure in (3.1) has a computational cost of order $O(l|\mathbb{S}|^3)$.

Plug-in estimate. Luckily in our case, Assumption 2.1 (i.e., \hat{P} follows the normalized multinomial distribution) allows for a direct formula for ε° , which automatically removes the error in the empirical mean. We can simply set

$$\tilde{\varepsilon}^\circ := \theta(\hat{b}, \hat{P}), \tag{3.3}$$

where $\theta(b, P)$ is defined in Theorem 2.1. The complete numerical algorithm is presented in Algorithm 1. The right-hand side of (3.1) converges to $\theta(\hat{b}, \hat{P})$ as $l \rightarrow \infty$. In addition, the computational cost is reduced from $O(l|\mathbb{S}|^3)$ to $O(|\mathbb{S}|^3)$. This complete removal of empirical mean error is what sets the multinomial MDP case apart from general operator shifting. Moreover, since both ε° in (2.11) and $\tilde{\varepsilon}^\circ$ in (3.3) share the same functional form, the lower and upper bounds in Section 2.3 automatically apply to both ε° and $\tilde{\varepsilon}^\circ$. In all the following numerical examples, we use the approximated factor $\tilde{\varepsilon}^\circ$, which does not rely on oracle access.

3.2. Numerical examples.

Policy evaluation of an MDP over a circle. We first consider an MDP over a discrete state space $\mathbb{S} = \{k\}_{k=0}^{N-1}$ with $N = 64$ and $\gamma = 0.9$. The transition dynamics are given as below,

$$\begin{aligned} s_{t+1} &\leftarrow s_t + (1 + Z_t) a_t \pmod{N}, \\ r_{s_t, a_t} &= \sin\left(\frac{2\pi s_t}{N}\right) + a_t \cos\left(\frac{2\pi s_t}{N}\right) / 10 + X_t, \end{aligned}$$

where $a_t \in \{\pm 1\}$ is drawn from a policy $\pi(a_t | s_t) = \frac{1}{2} + \frac{1}{5} a_t \sin\left(\frac{2\pi s_t}{N}\right)$, and $X_t \sim N(0, \delta)$ with $\delta \in \{0, 0.1, 0.2\}$. When $\delta = 0$, the reward is deterministic. Here Z_t is a random integer taking values in the set $\{-\sigma, \dots, \sigma\}$ with equal probability, where $\sigma \in \{1, 2, 4\}$. A larger σ means each state could transit to more neighboring states under one step. Figure 3.1 illustrates the distribution of the estimated $v(x_s)$ for $x_s := 0$, see caption for implementation detail. One can see that the shifted value is more concentrated around the ground truth value than the naive solution. One can also bootstrap to reduce the l_2 error. Specifically, the l_2 error refers to the case when $M = A^{-\top} A^{-1}$ in the norm $\|\cdot\|_M$ defined in (2.2). Despite a lack of access to A , one can use $M = \hat{A}^{-\top} \hat{A}^{-1}$ for l_2 error minimization, which works well empirically. The error reduction trend remains the same (see Figure 3.2 for details). Overall, the error reduction for l_2 norm is less significant than the residual norm, though it is still significant for small n .

In the following examples, we focus on the results for the residual norm.

Algorithm 1 Operator Shifting for estimating MDP (Multinomial)

Outputs the bootstrapped $\tilde{\varepsilon}^\circ = \theta(\hat{b}, \hat{P})$. The function can also output the true value by $\varepsilon^\circ = \theta(b, P)$ if one has oracle access to P, b .

Require: \hat{P} : Estimated transition matrix

Require: \hat{b} : Estimated expected reward

Require: n : Sample data size per state

Require: $\tilde{\Sigma}$: Estimated covariance matrix of \hat{b}

Require: γ : Discount factor

Require: M : The chosen norm matrix

```

1: function  $\theta(\hat{P}, \hat{b})$ 
2:    $\hat{C}, \hat{D} \leftarrow \text{zeros}(|\mathbb{S}|)$  ▷  $\text{zeros}(n)$ : A zero matrix of size  $n$ 
3:   for  $i \leftarrow 1$  to  $|\mathbb{S}|$  do
4:      $\hat{p}_i \leftarrow \hat{P}^\top e_i$  ▷ Get the  $i$ -th row of  $\hat{P}$ 
5:      $\hat{B}_i \leftarrow \frac{1}{n} (\hat{p}_i \hat{p}_i^\top - \text{diag}(\hat{p}_i))$  ▷ Estimate covariance
6:      $\hat{C} \leftarrow \hat{C} + \hat{B}_i \hat{A}^{-1} \text{diag}(e_i)$ 
7:      $\hat{D} \leftarrow \hat{D} + [M]_{ii} \hat{B}_i$ 
8:   end for
9:    $\hat{C} \leftarrow \gamma^2 \hat{A}^{-\top} \hat{C}$  ▷ Approximate  $\mathbb{E} \left[ (\hat{Y}^\top)^2 \right]$ 
10:   $\hat{G} \leftarrow \gamma^2 \hat{A}^{-\top} \hat{D} \hat{A}^{-1}$  ▷ Approximate  $\mathbb{E} \left[ \hat{Y}^\top M \hat{Y} \right]$ 
11:   $\hat{H} \leftarrow \hat{C} M + M \hat{C}^\top$  ▷ Approximate  $\mathbb{E} \left[ (\hat{Y}^\top)^2 M + M \hat{Y}^2 \right]$ 
12:   $\tilde{\varepsilon}^\circ \leftarrow \frac{\hat{b}^\top (M + \hat{H}/2) \hat{b}}{\hat{b}^\top (M + \hat{G} + \hat{H}) \hat{b} + \text{tr}(\tilde{\Sigma} (M + \hat{G} + \hat{H}))}$ 
13:  return  $\tilde{\varepsilon}^\circ$ 
14: end function

```

MDPs generated by random graphs. To test the robustness of Algorithm 1, here we apply the operator shifting method to different underlying transition matrices. For consistency, we set $|\mathbb{S}| = 64$.

As discussed in the 1D circle case, the randomness in \hat{b} usually boosts the performance of the operator shifting method. Here we take out the randomness in the reward and instead let \hat{b} be deterministic, i.e., $\hat{b} = b$ and $\text{cov}[\hat{b}] = 0$. To test different b , we assume that b is randomly generated according to $\mathcal{N}(0, I)$. The transition matrix P corresponds to the random walk on a directed random graph $G = (V, E, w)$, where $V = \mathbb{S}$ is the vertex set, E is the edge set, and the edge weight is $w: E \rightarrow \mathbb{R}_{\geq 0}$.

Two types of random graphs are considered. In the first *dense* case, the graph G is considered to be fully connected, and the weight $w(e)$ on each edge e is an i.i.d. random variable following $w(e) \sim \mathcal{U}(0, 1)$. In the second *sparse* case, a sparse graph is considered. In order to generate a random sparse graph, one initializes with a graph containing an empty edge set,

$$G \leftarrow G_0 := (V = \mathbb{S}, E = \emptyset).$$

For each vertex $v \in V$, two vertices v_1, v_2 are randomly selected from the set $\mathbb{S} \setminus \{v\}$ that excludes v itself with equal probability, and then

$$E \leftarrow E \cup \{(v_1, v), (v, v_2)\}.$$




Fig. 3.1: 1D circle example. Histogram of $v_{sol}(x_s) - v(x_s)$, where $v(x_s)$ denotes the true return at x_s , and $v_{sol}(x_s)$ denotes the estimated return from either the naive solution or the operator shifting solution. The shifted solution is more concentrated around the ground truth. The histogram is obtained via 20,000 simulations, and we set $n=8$, $\delta=0.2$, and $\sigma=4$.

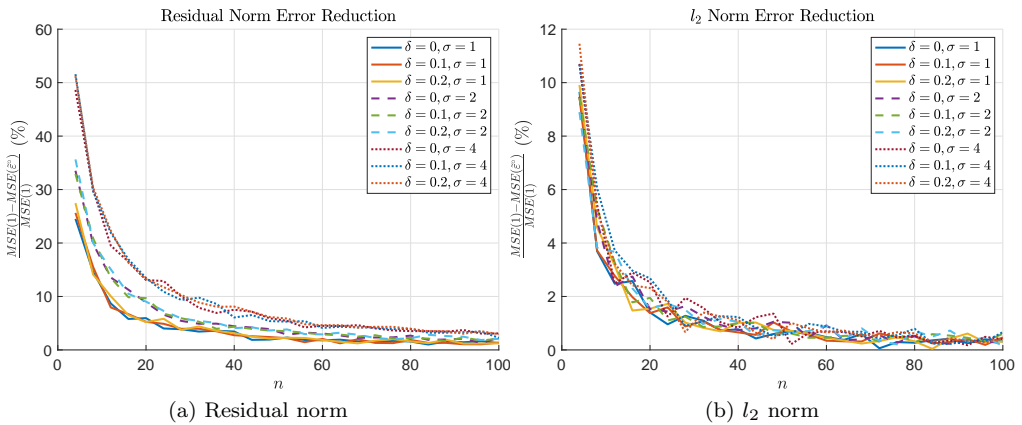


Fig. 3.2: 1D circle example. Error reduction as a function of n for the residual norm (left) and the l_2 norm (right). The y-axis is the error reduction rate in MSE, relative to the error of the naive solution. The error reduction rate is inversely proportional to n across different choices of σ and δ . For the residual norm case, the error reduction is heavily influenced by the choice of parameters, where a larger σ or δ implies a larger reduction in error.

After enumerating over all vertices, one then assigns a weight of one to all existing edges in G . This construction ensures that none of the vertices is a well or sink node, that is, each vertex has at least one indegree and one outdegree, but the transition matrix is still quite sparse. Figure 3.3 shows that the same MSE reduction pattern holds in the random directed graph cases. The operator shifting solution still consistently outperforms the naive solution.

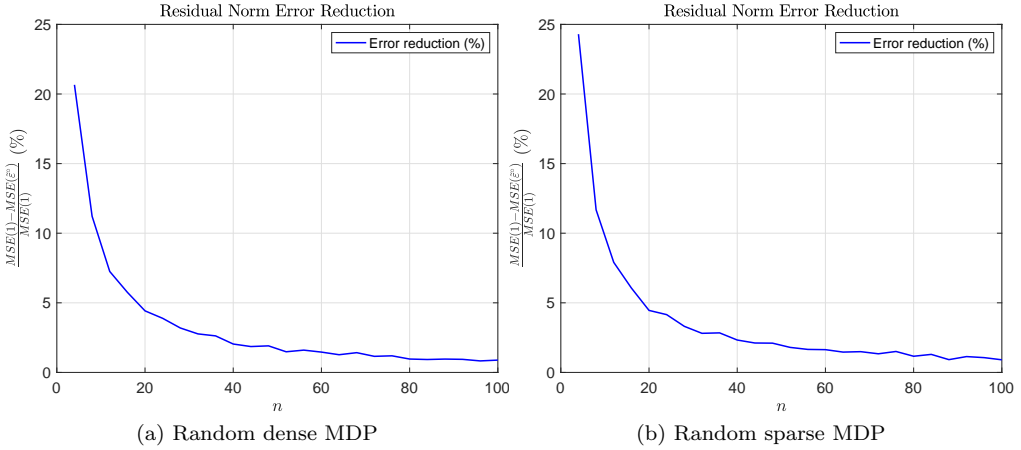


Fig. 3.3: Random directed graphs. Error reduction as a function of sample size n . Left: the random dense graph. Right: the random sparse graph.




Fig. 3.4: 2D torus example. Error reduction as a function of sample size n .

Policy evaluation of an MDP over a torus. We now consider an MDP with a discrete state space $\mathbb{S} = \{s_{ij} = (i, j)\}_{i,j=0}^{N-1}$ with $N = 8$ and $\gamma = 0.9$. Note that the size of the state space $|\mathbb{S}|$ is still 64. Let $(s)_k$ stand for the first or second entry of the vector s with $k = 1$ or 2 . The transition dynamics and reward are given by

$$s_{t+1} \leftarrow s_t + (1 + Z_t)a_t \pmod N,$$

$$r_{s_t, a_t} \leftarrow -2 + \sin\left(\frac{2\pi(s_t)_1}{N}\right) + \cos\left(\frac{2\pi(s_t)_2}{N}\right) + X_t,$$

where $a_t \in \mathbb{A} = \{(\pm 1, 0), (0, \pm 1)\}$, $X_t \sim N(0, \delta)$ with $\delta \in \{0, 0.1, 0.2\}$. Here Z_t is a random integer taking values in the set $\{-\sigma, \dots, \sigma\}$ with equal probability, where $\sigma \in \{1, 2, 4\}$.

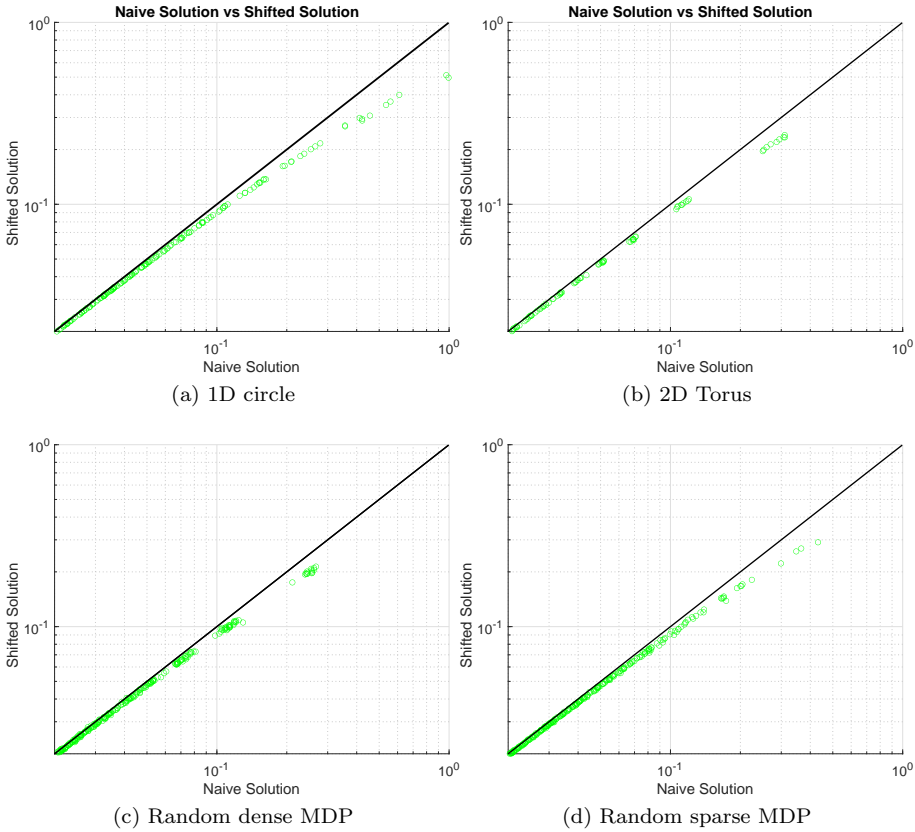


Fig. 3.5: The normalized MSE of the operator shifting solution is plotted against that of the naive solution. All data points are below or close to the diagonal, showing that the operator shifting solution outperforms the naive solution in all data points collected.

We use the policy

$$\pi(a_t = (a_1, a_2) | s_t) = \frac{1}{4} + \frac{1}{20} \left(a_1 \cos\left(\frac{2\pi(s_t)_1}{N}\right) + a_2 \sin\left(\frac{2\pi(s_t)_2}{N}\right) \right). \quad (3.4)$$

Figure 3.4 summarizes the performance and exhibits a similar error reduction trend. Contrary to the role of the parameters in the 1D circle case, different choices of σ and δ do not change the performance of the operator shifting method.

Summary of numerical experiments. Figure 3.5 plots the normalized MSE of the naive solution against the operator shifting solution. In the torus and circle cases, the data points are obtained by varying the sample size n , the reward variance δ , and the transition parameter σ . In the randomly generated MDP case, the data points are obtained by sampling random MDPs and varying the value of the sample size n . The vast majority of the data points are below the diagonal line, suggesting that operator shifting consistently reduces the MSE.

As a further remark, the numerical result shows that the bounds in Theorem 2.2 for ε° are quite pessimistic. In practice, ε° almost always falls in the $(0,1)$ range, even for small n .

4. Proofs

4.1. Proof of Lemma 2.1.

Proof. From (2.5), $(I - \hat{Y})A\hat{v} = (\hat{A}A^{-1})A\hat{v} = \hat{A}\hat{v} = \hat{b}$ and

$$A\hat{v} = (I - \hat{Y})^{-1}\hat{b}.$$

Hence (2.4) can be written as

$$\varepsilon^* = \frac{\mathbb{E}_{\hat{P}, \hat{b}} \left[b^\top M (I - \hat{Y})^{-1} \hat{b} \right]}{\mathbb{E}_{\hat{P}, \hat{b}} \left[\hat{b}^\top (I - \hat{Y})^{-\top} M (I - \hat{Y})^{-1} \hat{b} \right]}.$$

From Assumption 2.1, $\mathbb{E}[\hat{b}] = b$. Moreover, it follows from Assumption 2.1 that \hat{P} is independent to \hat{b} . Hence one can write the numerator as

$$\mathbb{E}_{\hat{P}, \hat{b}} \left[b^\top M (I - \hat{Y})^{-1} \hat{b} \right] = \mathbb{E}_{\hat{P}} \left[b^\top M (I - \hat{Y})^{-1} b \right],$$

and the denominator as

$$\begin{aligned} & \mathbb{E}_{\hat{P}, \hat{b}} \left[\hat{b}^\top (I - \hat{Y})^{-\top} M (I - \hat{Y})^{-1} \hat{b} \right] \\ = & \mathbb{E}_{\hat{P}} \left[b^\top (I - \hat{Y})^{-\top} M (I - \hat{Y})^{-1} b \right] + \mathbb{E}_{\hat{P}, \hat{b}} \left[(b - \hat{b})^\top (I - \hat{Y})^{-\top} M (I - \hat{Y})^{-1} (b - \hat{b}) \right]. \end{aligned}$$

One can rewrite the second term in terms of the variance of \hat{b} by the trace property

$$\mathbb{E}_{\hat{P}, \hat{b}} \left[(b - \hat{b})^\top (I - \hat{Y})^{-\top} M (I - \hat{Y})^{-1} (b - \hat{b}) \right] = \mathbb{E}_{\hat{P}} \left[\text{tr} \left(\text{cov} [\hat{b}] (I - \hat{Y})^{-\top} M (I - \hat{Y})^{-1} \right) \right],$$

and likewise one has

$$\mathbb{E}_{\hat{P}} \left[b^\top (I - \hat{Y})^{-\top} M (I - \hat{Y})^{-1} b \right] = \mathbb{E}_{\hat{P}} \left[\text{tr} \left(b^\top b (I - \hat{Y})^{-\top} M (I - \hat{Y})^{-1} \right) \right].$$

The entries of \hat{b} are uncorrelated because as defined in (1.1),

$$\mathbb{P}(b_{s_1} = r_{s_1}^{a_1}, b_{s_2} = r_{s_2}^{a_2}) = \pi_{s_1, a_2} \pi_{s_1, a_2} = \mathbb{P}(b_{s_1} = r_{s_1}^{a_1}) \mathbb{P}(b_{s_2} = r_{s_2}^{a_2}).$$

As a result, $\text{cov} [\hat{b}]$ is a diagonal matrix as claimed. □

4.2. Proof of Theorem 2.1 and derivation of (2.9).

Derivation of (2.9). We first show the derivation of (2.9). First one inserts the truncated Neumann series into the definition of ε^* in (2.6). According to (2.7),

$$(I - \hat{Y})^{-1} \approx I + \hat{Y} + \hat{Y}^2.$$

One has the following series of approximations by truncating out terms beyond order two

$$M (I - \hat{Y})^{-1} \approx M + M\hat{Y} + M\hat{Y}^2,$$

$$(I - \hat{Y})^{-\top} M (I - \hat{Y})^{-1} \approx M + M(\hat{Y} + \hat{Y}^2) + (\hat{Y}^\top + (\hat{Y}^\top)^2)M + \hat{Y}^\top M \hat{Y}.$$

Note that $\mathbb{E}[\hat{P}] = P$ due to Assumption 2.1. Thus $\mathbb{E}[\hat{Y}] = 0$. Therefore, taking expectation of the above two terms gives

$$\mathbb{E}_{\hat{P}} \left[M (I - \hat{Y})^{-1} \right] \approx M + M \mathbb{E}_{\hat{P}} [\hat{Y}^2], \tag{4.1}$$

$$\mathbb{E}_{\hat{P}} \left[(I - \hat{Y})^{-\top} M (I - \hat{Y})^{-1} \right] \approx M + M \mathbb{E}_{\hat{P}} [\hat{Y}^2] + \mathbb{E}_{\hat{P}} [(\hat{Y}^\top)^2] M + \mathbb{E}_{\hat{P}} [\hat{Y}^\top M \hat{Y}]. \tag{4.2}$$

Plugging (4.1) and (4.2) into (2.6) leads to

$$\begin{aligned} \varepsilon^* &= \frac{\mathbb{E}_{\hat{P}} \left[b^\top M (I - \hat{Y})^{-1} b \right]}{\mathbb{E}_{\hat{P}} \left[\text{tr} \left((\text{cov} [\hat{b}] + b^\top b) (I - \hat{Y})^{-\top} M (I - \hat{Y})^{-1} \right) \right]} \\ &\approx \frac{\mathbb{E}_{\hat{P}} \left[b^\top (M + M \hat{Y}^2) b \right]}{\mathbb{E}_{\hat{P}} \left[\text{tr} \left((\text{cov} [\hat{b}] + b^\top b) (M + \hat{Y}^\top M \hat{Y} + M \hat{Y}^2 + (\hat{Y}^\top)^2 M) \right) \right]}, \end{aligned}$$

where the numerator term can be symmetrized so as to get (2.9).

Proof of Theorem 2.1. Let $N = |\mathcal{S}|$. Denote by $\{\hat{p}_i\}_{i=1}^N$ and $\{p_i\}$ the row vectors of \hat{P} and P , respectively:

$$\hat{P} = \begin{bmatrix} \hat{p}_1^\top \\ \vdots \\ \hat{p}_N^\top \end{bmatrix}, \quad P = \begin{bmatrix} p_1^\top \\ \vdots \\ p_N^\top \end{bmatrix}.$$

To show that ε° follows the formula in Theorem 2.1, it suffices to prove the following auxiliary lemma:

LEMMA 4.1. *Assume the following two conditions hold.*

- (a): \hat{P}, \hat{b} are unbiased estimators of P, b .
- (b): X_i is independent to X_j whenever $i \neq j$.

Then one has

$$\mathbb{E} [\hat{Y}^\top M \hat{Y}] = \gamma^2 A^{-\top} \left(\sum_{i=1}^N [M_{ii}] \text{cov} [\hat{p}_i] \right) A^{-1}, \tag{4.3}$$

$$\mathbb{E} [\hat{Y}^2] = \gamma^2 \sum_{i=1}^N \text{diag} (e_i) A^{-\top} \text{cov} [\hat{p}_i] A^{-1}, \tag{4.4}$$

$$\mathbb{E} [(\hat{Y}^\top)^2] = \gamma^2 \sum_{i=1}^N A^{-\top} \text{cov} [\hat{p}_i] A^{-1} \text{diag} (e_i), \tag{4.5}$$

where \hat{p}_i is the random vector corresponding to the i -th row of \hat{P} .

Both conditions in Lemma 4.1 are satisfied under Assumption 2.1. In Theorem 2.1, one has $X_i \sim \text{multinomial}(n, p_i)$ with the following covariance structure

$$\text{cov}[\hat{p}_i] = \frac{1}{n} [\text{diag} p_i - p_i p_i^\top] = B_i. \tag{4.6}$$

Plugging in (4.6) in Lemma 4.1 immediately gives the expectation-free form in Lemma 2.1, which proves Theorem 2.1.

Proof. (Proof of Lemma 4.1.) We first calculate $\mathbb{E}[\hat{Y}^\top M \hat{Y}]$. To do this, we rely on the assumption that the i -th row \hat{p}_i is independent to \hat{p}_j whenever $i \neq j$. As a consequence, the rows of \hat{Z} are independent. Then, for any matrix M , one has

$$\mathbb{E}[\hat{Y}^\top M \hat{Y}] = \mathbb{E}[A^{-\top} \hat{Z}^\top M \hat{Z} A^{-1}] = A^{-\top} \mathbb{E}[\hat{Z}^\top M \hat{Z}] A^{-1}.$$

By denoting the rows of \hat{Z} by $\hat{z}_1^\top, \dots, \hat{z}_N^\top$,

$$\hat{Z}^\top M \hat{Z} = [\hat{z}_1 \dots \hat{z}_N] M \begin{bmatrix} \hat{z}_1^\top \\ \dots \\ \hat{z}_N^\top \end{bmatrix} = \sum_{i,j=1}^N \hat{z}_i M_{ij} \hat{z}_j^\top.$$

By taking the expectation, the only non-zero terms are the ones with $i = j$. Hence,

$$\mathbb{E}[\hat{Z}^\top M \hat{Z}] = \sum_{i=1}^N M_{ii} \mathbb{E}[\hat{z}_i \hat{z}_i^\top].$$

Then by definition of \hat{Z} one has

$$\mathbb{E}[\hat{z}_i \hat{z}_i^\top] = \gamma^2 \mathbb{E}[(\hat{p}_i - \mathbb{E}[\hat{p}_i])(\hat{p}_i - \mathbb{E}[\hat{p}_i])^\top] = \gamma^2 \text{cov}[\hat{p}_i]. \tag{4.7}$$

Hence one can get the first part of Lemma 4.1, which is

$$\mathbb{E}[\hat{Y}^\top M \hat{Y}] = A^{-\top} \mathbb{E}[\hat{Z}^\top M \hat{Z}] A^{-1} = \gamma^2 A^{-\top} \left(\sum_{i=1}^N M_{ii} \text{cov}[\hat{p}_i] \right) A^{-1}.$$

Now we move on to proving the form of $\mathbb{E}[\hat{Y}^2]$. Writing out \hat{Y}^2 explicitly

$$\hat{Y}^2 = \hat{Z} A^{-1} \hat{Z} A^{-1} = \begin{bmatrix} \hat{z}_1^\top A^{-1} \\ \dots \\ \hat{z}_N^\top A^{-1} \end{bmatrix} \begin{bmatrix} \hat{z}_1^\top \\ \dots \\ \hat{z}_N^\top \end{bmatrix} A^{-1} = \sum_{i=1}^N \sum_{j=1}^N \begin{bmatrix} 0 \\ \dots \\ \hat{z}_i^\top A^{-1} \\ \dots \\ 0 \end{bmatrix} \begin{bmatrix} 0 \\ \dots \\ \hat{z}_j^\top \\ \dots \\ 0 \end{bmatrix} A^{-1}.$$

After the expectation, the only non-zero terms are $i = j$. Thus one has

$$\mathbb{E}[\hat{Y}^2] = \sum_{i=1}^N \mathbb{E} \left[\begin{bmatrix} 0 \\ \dots \\ \hat{z}_i^\top A^{-1} \\ \dots \\ 0 \end{bmatrix} \begin{bmatrix} 0 \\ \dots \\ \hat{z}_i^\top \\ \dots \\ 0 \end{bmatrix} \right] A^{-1},$$

with

$$\left(\begin{bmatrix} 0 \\ \vdots \\ \hat{z}_i^\top A^{-1} \\ \vdots \\ 0 \end{bmatrix} \begin{bmatrix} 0 \\ \vdots \\ \hat{z}_i^\top \\ \vdots \\ 0 \end{bmatrix} \right)_{jk} = \begin{cases} \sum_{l=1}^N \hat{Z}_{il} A_{li}^{-1} \hat{Z}_{ik} & j=i \\ 0 & j \neq i. \end{cases}$$

For the matrix $A^{-\top} \hat{z}_i \hat{z}_i^\top$, note that

$$[A^{-\top} \hat{z}_i \hat{z}_i^\top]_{jk} = (A^{-\top} \hat{z}_i)_j (\hat{z}_i^\top)_k = \sum_{l=1}^N A_{jl}^{-\top} \hat{Z}_{il} \hat{Z}_{ik} = \sum_{l=1}^N A_{lj}^{-1} \hat{Z}_{il} \hat{Z}_{ik}.$$

Applying $j=i$ leads to

$$[\text{diag}(e_i) A^{-\top} \hat{z}_i \hat{z}_i^\top]_{jk} = \begin{cases} \sum_{l=1}^N \hat{Z}_{il} A_{li}^{-1} \hat{Z}_{ik} & j=i \\ 0 & j \neq i. \end{cases}$$

Hence we have

$$\begin{aligned} \mathbb{E} [\hat{Y}^2] &= \sum_{i=1}^N \mathbb{E} \left[\begin{bmatrix} 0 \\ \vdots \\ \hat{z}_i^\top A^{-1} \\ \vdots \\ 0 \end{bmatrix} \begin{bmatrix} 0 \\ \vdots \\ \hat{z}_i^\top \\ \vdots \\ 0 \end{bmatrix} \right] A^{-1} = \sum_{i=1}^N \mathbb{E} \left[\text{diag}(e_i) \begin{bmatrix} 0 \\ \vdots \\ \hat{z}_i^\top A^{-1} \\ \vdots \\ 0 \end{bmatrix} \begin{bmatrix} 0 \\ \vdots \\ \hat{z}_i^\top \\ \vdots \\ 0 \end{bmatrix} \right] A^{-1} \\ &= \sum_{i=1}^N \mathbb{E} [\text{diag}(e_i) A^{-\top} \hat{z}_i \hat{z}_i^\top] A^{-1} = \sum_{i=1}^N \text{diag}(e_i) A^{-\top} \mathbb{E} [\hat{z}_i \hat{z}_i^\top] A^{-1}. \end{aligned}$$

Taking transpose results in

$$\mathbb{E} [(\hat{Y}^\top)^2] = \sum_{i=1}^N A^{-\top} \mathbb{E} [\hat{z}_i \hat{z}_i^\top] A^{-1} \text{diag}(e_i).$$

□

4.3. Proof of Lemma 2.2 and Corollary 2.1. We first prove Lemma 2.2.

Proof. (Proof of Lemma 2.2.) Going back to the original quadratic optimization problem, one has

$$\mathbb{E}_{\hat{P}, \hat{b}} \|b - \varepsilon A \hat{v}\|_M^2 = \varepsilon^2 \mathbb{E}_{\hat{P}, \hat{b}} \|A \hat{v}\|_M^2 - 2\varepsilon \mathbb{E}_{\hat{P}, \hat{b}} [b^\top M A \hat{v}] + \|b\|_M^2. \tag{4.8}$$

Using Lemma 2.1 and the second-order approximation in equation (2.9), one has

$$\mathbb{E}_{\hat{P}, \hat{b}} [b^\top M A \hat{v}] = \mathbb{E}_{\hat{P}} \left[b^\top \left(M + \frac{M \hat{Y}^2 + (\hat{Y}^\top)^2 M}{2} \right) b \right] + h.o.t.,$$

and

$$\begin{aligned} \mathbb{E}_{\hat{P}, \hat{b}} \|A \hat{v}\|_M^2 &= \mathbb{E}_{\hat{P}} \left[b^\top (M + \hat{Y}^\top M \hat{Y} + M \hat{Y}^2 + (\hat{Y}^\top)^2 M) b \right] \\ &\quad + \mathbb{E}_{\hat{P}} \left[\text{tr} \left(\text{cov} [\hat{b}] (M + \hat{Y}^\top M \hat{Y} + M \hat{Y}^2 + (\hat{Y}^\top)^2 M) \right) \right] + h.o.t., \end{aligned}$$

where *h.o.t.* stands for high order terms.

We now show that $h.o.t. = O\left(n^{-\frac{3}{2}}\right)$. The expectation of the third order or higher terms in \hat{Y} is computed by moments of third order or higher in \hat{Z} . Under Assumption 2.1, rows of \hat{Z} are independent, and hence moments of \hat{Z} are linear combinations of moments in multinomial distribution. Each row of matrix \hat{Z} is an average of n random variables with mean zero, which is why its moments of third order or higher decay at the rate of at least $O\left(n^{-\frac{3}{2}}\right)$ by use of the Marcinkiewicz-Zygmund inequality.

One can then plug in the explicit formula from Lemma 4.1 and (4.6), leading to

$$\mathbb{E}_{\hat{P}, \hat{b}} [b^\top M A \hat{v}] = \|b\|_M^2 + b^\top H b / 2 + O\left(n^{-\frac{3}{2}}\right) = \|b\|_M^2 + h / 2 + O\left(n^{-\frac{3}{2}}\right). \tag{4.9}$$

and

$$\begin{aligned} \mathbb{E}_{\hat{P}, \hat{b}} \|A \hat{v}\|_M^2 &= \|b\|_M^2 + b^\top G b + b^\top H b + \text{tr}\left(\text{cov}\left[\hat{b}\right] (M + G + H)\right) + O\left(n^{-\frac{3}{2}}\right) \\ &= \|b\|_M^2 + h + g + t + O\left(n^{-\frac{3}{2}}\right). \end{aligned} \tag{4.10}$$

Plugging it into (4.8) results in (2.13).

We now move to prove (2.14). From (4.8) and the results before, one has

$$\varepsilon^* = \frac{\|b\|_M^2 + h / 2 + O\left(n^{-\frac{3}{2}}\right)}{\|b\|_M^2 + h + g + t + O\left(n^{-\frac{3}{2}}\right)}.$$

On the other hand, Theorem 2.1 proves that ε° follows the following form:

$$\varepsilon^\circ = \frac{\|b\|_M^2 + h / 2}{\|b\|_M^2 + h + g + t}. \tag{4.11}$$

Moreover, from Theorem 2.1, it follows that $G, H, \text{cov}\left[\hat{b}\right] \propto \frac{1}{n}$, and therefore $g, h, t \propto \frac{1}{n}$. Hence (2.14) holds and ε° is asymptotically optimal. \square

Proof. (Proof of Corollary 2.1.) Throughout this proof, we use the fact that $g, h, t \propto \frac{1}{n}$ as in the proof of Lemma 2.2. Without loss of generality, assume that $\|b\|_M^2 = 1$. The proof is organized as follows. First, we prove that $\text{MSE}(1) - \text{MSE}(\varepsilon^*) = O\left(n^{-2}\right)$. Second, we prove that $\text{MSE}(\varepsilon^*) - \text{MSE}(\varepsilon^\circ) = O\left(n^{-3}\right)$. As a consequence, one obtains $\text{MSE}(1) - \text{MSE}(\varepsilon^\circ) = O\left(n^{-2}\right)$. Then, note that

$$\text{MSE}(1) = (g + t) + O\left(n^{-\frac{3}{2}}\right) = O(1/n), \tag{4.12}$$

and so the relative error is of order $O(1/n)$ as claimed, and η is positive for n sufficiently large.

We first estimate $\text{MSE}(1) - \text{MSE}(\varepsilon^*)$. Plugging in (2.4), one has

$$\text{MSE}(1) - \text{MSE}(\varepsilon^*) = \frac{\left(\|A \hat{v}\|_M^2 - \mathbb{E}_{\hat{P}, \hat{b}} [b^\top M A \hat{v}]\right)^2}{\|A \hat{v}\|_M^2}. \tag{4.13}$$

Under the assumption that $\|b\|_M^2 = 1$, (4.9) and (4.10) shows

$$\mathbb{E}_{\hat{P}, \hat{b}} \|A \hat{v}\|_M^2 = 1 + h + g + t + O\left(n^{-\frac{3}{2}}\right), \mathbb{E}_{\hat{P}, \hat{b}} [b^\top M A \hat{v}] = 1 + h / 2 + O\left(n^{-\frac{3}{2}}\right).$$

Thus, plugging in the previously derived terms into (4.15), one has

$$\text{MSE}(1) - \text{MSE}(\varepsilon^*) = \frac{\left(g + h/2 + t + O\left(n^{-\frac{3}{2}}\right)\right)^2}{1 + g + h + t + O\left(n^{-\frac{3}{2}}\right)} = O\left(n^{-2}\right). \tag{4.14}$$

We then estimate $\text{MSE}(\varepsilon^*) - \text{MSE}(\varepsilon^\circ)$:

$$\text{MSE}(\varepsilon^*) - \text{MSE}(\varepsilon^\circ) = \left((\varepsilon^*)^2 - (\varepsilon^\circ)^2\right) \mathbb{E}_{\hat{P}, \hat{b}} \|A\hat{v}\|_M^2 - 2(\varepsilon^* - \varepsilon^\circ) \mathbb{E}_{\hat{P}, \hat{b}} [b^\top MA\hat{v}]. \tag{4.15}$$

Then, one uses

$$\varepsilon^* = \frac{\mathbb{E}_{\hat{P}, \hat{b}} [b^\top MA\hat{v}]}{\mathbb{E}_{\hat{P}, \hat{b}} \|A\hat{v}\|_M^2}, \quad \varepsilon^\circ = \frac{1 + h/2}{1 + h + g + t}.$$

By simple algebra, one obtains from (4.15) that

$$\text{MSE}(\varepsilon^*) - \text{MSE}(\varepsilon^\circ) = -\frac{\left(\left(1 + h + g + t\right) \mathbb{E}_{\hat{P}, \hat{b}} [b^\top MA\hat{v}] - \mathbb{E}_{\hat{P}, \hat{b}} \|A\hat{v}\|_M^2 (1 + h/2)\right)^2}{\left(\mathbb{E}_{\hat{P}, \hat{b}} \|A\hat{v}\|_M^2\right) (1 + h + g + t)^2}, \tag{4.16}$$

Importantly, for the numerator term in (4.16), note that

$$\left(1 + h + g + t\right) \mathbb{E}_{\hat{P}, \hat{b}} [b^\top MA\hat{v}] = \left(1 + h + g + t\right) (1 + h/2) + O\left(n^{-\frac{3}{2}}\right),$$

and likewise one has

$$\mathbb{E}_{\hat{P}, \hat{b}} \|A\hat{v}\|_M^2 (1 + h/2) = \left(1 + h + g + t\right) (1 + h/2) + O\left(n^{-\frac{3}{2}}\right).$$

Consequently, the numerator term in (4.16) is of order $O\left(n^{-3}\right)$. Thus

$$\text{MSE}(\varepsilon^*) - \text{MSE}(\varepsilon^\circ) = \frac{O\left(n^{-3}\right)}{\left(1 + g + h + t + O\left(n^{-\frac{3}{2}}\right)\right) (1 + g + h + t)^2} = O\left(n^{-3}\right), \tag{4.17}$$

as is desired. □

4.4. Proof of Theorem 2.2. To prove Theorem 2.2, one first finds a tight bound for $A^{-1}b$ and $A^{-1}\text{diag}(e_i)b$. The tight upper and lower bounds for $A^{-1}b$ are stated in Lemma 4.2. Then, the upper bounds for $\|A^{-1}b\|^2$ and $\sum_i \|A^{-1}\text{diag}(e_i)b\|^2$ are listed in Corollary 4.1. Finally, based on Corollary 4.1, we derive the bound for ε° in Theorem 2.2.

LEMMA 4.2. For any transition matrix $P \in \mathbb{R}^{|\mathcal{S}| \times |\mathcal{S}|}$, vector $b \in \mathbb{R}^{|\mathcal{S}|}$ and $\gamma \in (0, 1)$,

$$b + \frac{\gamma}{1 - \gamma} b_m \mathbf{1} \leq (I - \gamma P)^{-1} b \leq b + \frac{\gamma}{1 - \gamma} b_M \mathbf{1},$$

where $b_m = \min_s b_s, b_M = \max_s b_s$, and inequality between vectors denotes an entry-wise inequality.

Proof. Let $x = (I - \gamma P)^{-1} b$, and $s = \text{argmin}_i x_i$, then the s -th row of $(I - \gamma P)x = b$ is

$$b_s = x_s - \gamma \sum_t P_{st} x_t \leq x_s - \gamma \sum_t P_{st} x_s = (1 - \gamma) x_s,$$

which implies

$$x_s \geq \frac{b_s}{1-\gamma} \geq \frac{b_m}{1-\gamma}, \tag{4.18}$$

where $b_m = \min_s b_s$. For $\forall j \neq s$, one has

$$b_j = x_j - \gamma \sum_t P_{jt} x_t \leq x_j - \gamma \sum_t P_{st} x_s = x_j - \gamma x_s \leq x_j - \gamma \frac{b_m}{1-\gamma},$$

which yields,

$$x_j \geq b_j + \frac{\gamma}{1-\gamma} b_m. \tag{4.19}$$

Combining (4.18) and (4.19) gives

$$x = (I - \gamma P)^{-1} b \geq b + \frac{\gamma}{1-\gamma} b_m \mathbf{1}.$$

On the other hand, let $l = \operatorname{argmax}_i x_i$, then the l -th row of $(I - \gamma P)x = b$ is

$$b_l = x_l - \gamma \sum_t P_{lt} x_t \geq x_l - \gamma \sum_t P_{lt} x_l = (1 - \gamma)x_l,$$

which implies

$$x_l \leq \frac{b_l}{1-\gamma} \leq \frac{b_M}{1-\gamma}, \tag{4.20}$$

where $b_M = \max_s b_s$. For $\forall j \neq l$, one has

$$b_j = x_j - \gamma \sum_t P_{jt} x_t \geq x_j - \gamma \sum_t P_{jt} x_l \geq x_j - \gamma \frac{b_M}{1-\gamma},$$

which yields,

$$x_j \leq b_j + \frac{\gamma}{1-\gamma} b_M. \tag{4.21}$$

Combining (4.20) and (4.21) gives

$$x = (I - \gamma P)^{-1} b \leq b + \frac{\gamma}{1-\gamma} b_M \mathbf{1},$$

which completes the proof. □

LEMMA 4.3. For a vector $v = (v_i)_{i=1}^d$, define $|v| := (|v_i|)_{i=1}^d$, i.e. the entry-wise absolute value of v . Suppose a matrix Q has only non-negative entries. Then, for any vector $v \in \mathbb{R}^{|\mathcal{S}|}$ and $\gamma \in (0, 1)$,

$$|Qv| \leq Q|v|.$$

Proof. Define $x = Qv$. Denote by v_+, v_- the positive and negative parts of v , respectively. That is, $(v_+)_i = v_i \mathbb{1}_{v_i > 0}$ and $(v_-)_i = v_i \mathbb{1}_{v_i \leq 0}$. Since Q has only non-negative entries, it follows $Qv_+ \geq \mathbf{0}$ and $Qv_- \leq \mathbf{0}$. Because $x = Qv_+ + Qv_-$, one has $x \leq Qv_+ - Qv_-$ and $-x \geq -Qv_+ + Qv_-$. Hence

$$|x| \leq Qv_+ - Qv_- = Q|v|.$$

□

COROLLARY 4.1. For any transition matrix P , vector $b \in \mathbb{R}^{|\mathbb{S}|}$ and $\gamma \in (0, 1)$, one has,

$$\begin{aligned} \left\| (I - \gamma P)^{-1} b \right\|_2 &\leq \|k\|_2, \\ \sum_i \left\| (I - \gamma P)^{-1} \text{diag}(e_i) b \right\|_2^2 &\leq \|k\|_2^2, \\ \sum_i \left\| (I - \gamma P)^{-1} \text{diag}(e_i) b \right\|_2 &\leq \sqrt{|\mathbb{S}|} \|k\|_2, \end{aligned}$$

where $k = \frac{1}{1-\gamma} ((1-\gamma)|b| + \gamma b_M \mathbf{1})$ with $b_M = \max_i |b_i|$ and

$$\|k\|_2^2 = \frac{1}{(1-\gamma)^2} \left((1-\gamma)^2 \|b\|_2^2 + |\mathbb{S}| \gamma^2 b_M^2 + 2\gamma(1-\gamma) b_M \|b\|_1 \right).$$

Proof. Note that $(I - \gamma P)^{-1}$ is a matrix with non-negative entries, and therefore one has

$$\left| (I - \gamma P)^{-1} b \right| \leq (I - \gamma P)^{-1} |b|.$$

Then, Lemma 4.2 leads to

$$\left| (I - \gamma P)^{-1} b \right| \leq (I - \gamma P)^{-1} |b| \leq k = \frac{1}{1-\gamma} ((1-\gamma)|b| + \gamma b_M \mathbf{1}),$$

where $b_M = \max_i |b_i|$, which is the first inequality in the corollary.

The second inequality is because

$$\sum_i \left\| (I - \gamma P)^{-1} \text{diag}(e_i) b \right\|_2^2 \leq \left\| \sum_i \left| (I - \gamma P)^{-1} \text{diag}(e_i) b \right| \right\|_2^2 \leq \left\| (I - \gamma P)^{-1} |b| \right\|_2^2 \leq \|k\|_2^2.$$

The third inequality is due to

$$\begin{aligned} &\sum_i \left\| (I - \gamma P)^{-1} \text{diag}(e_i) b \right\|_2 \\ &\leq \sqrt{|\mathbb{S}|} \left\| \sum_i \left| (I - \gamma P)^{-1} \text{diag}(e_i) b \right| \right\|_2 \\ &\leq \sqrt{|\mathbb{S}|} \left\| (I - \gamma P)^{-1} |b| \right\|_2 \\ &\leq \sqrt{|\mathbb{S}|} \|k\|_2. \end{aligned}$$

□

Now we are ready to prove Theorem 2.2.

Proof. (Proof of Theorem 2.2.) Since B_i is the covariance matrix, $B_i \succeq 0$. By the definition of G, H in Theorem 2.1, one has $\text{tr} \left(\text{cov} \left[\hat{b} \right] (I + G + H) \right) \geq 0$. By letting $a_i = (I - \gamma P)^{-1} \text{diag}(e_i) b$ and $d = (I - \gamma P)^{-1} b$,

$$\varepsilon^\circ \leq \frac{\frac{1}{\gamma^2} b^\top b + \sum_i a_i^\top B_i d}{\frac{1}{\gamma^2} b^\top b + \sum_i d^\top B_i d + 2 \sum_i a_i^\top B_i d}$$

$$\begin{aligned}
 &= 1 - \frac{\sum_i d^\top B_i d + \sum_i a_i^\top B_i d}{\frac{1}{\gamma^2} b^\top b + \sum_i d^\top B_i d + 2 \sum_i a_i^\top B_i d} \\
 &= 1 + \frac{\frac{1}{2} \sum_i a_i^\top B_i a_i - \frac{1}{2} \sum_i (a_i + d)^\top B_i (a_i + d) - \frac{1}{2} \sum_i d^\top B_i d}{\frac{1}{\gamma^2} b^\top b + \sum_i (a_i + d)^\top B_i (a_i + d) - \sum_i a_i^\top B_i a_i}. \tag{4.22}
 \end{aligned}$$

The numerator of the second term in (4.22) can be bounded by

$$\frac{1}{2} \sum_i a_i^\top B_i a_i - \frac{1}{2} \sum_i (a_i + d)^\top B_i (a_i + d) - \frac{1}{2} \sum_i d^\top B_i d \leq \frac{1}{2} \lambda_M \sum_i \|a_i\|_2^2,$$

where λ_M is the largest eigenvalue of B_i for $\forall i$.

Suppose that $\frac{\lambda_M \|k\|_2^2}{\frac{1}{\gamma^2} b^\top b} \leq \frac{1}{2}$. Then the denominator of the second term of (4.22) can be lower bounded by

$$\frac{1}{\gamma^2} b^\top b + \sum_i (a_i + d)^\top B_i (a_i + d) - \sum_i a_i^\top B_i a_i \geq \frac{1}{\gamma^2} b^\top b - \lambda_M \|k\|_2^2 > 0, \tag{4.23}$$

where $\sum_i \|a_i\|_2^2 \leq \|k\|_2^2$ from Corollary 4.1 is used. Therefore, (4.22) can be bounded by

$$\varepsilon^\circ \leq 1 + \frac{\frac{1}{2} \lambda_M \|k\|_2^2}{\frac{1}{\gamma^2} b^\top b - \lambda_M \|k\|_2^2} \leq 1 + \frac{\lambda_M \|k\|_2^2}{\frac{1}{\gamma^2} b^\top b}.$$

Note that

$$\begin{aligned}
 \frac{\|k\|_2^2}{\frac{1}{\gamma^2} \|b\|_2^2} &\leq \frac{\gamma^2}{(1-\gamma)^2} \frac{\left((1-\gamma)^2 \|b\|_2^2 + |\mathbb{S}| \gamma^2 b_M^2 + 2\gamma(1-\gamma) b_M \|b\|_1 \right)}{\|b\|_2^2} \\
 &\leq \frac{\gamma^2}{(1-\gamma)^2} \left((1-\gamma)^2 + |\mathbb{S}| \gamma^2 \frac{b_M^2}{\|b\|_2^2} + 2\gamma(1-\gamma) \sqrt{|\mathbb{S}|} \frac{b_M}{\|b\|_2} \right) \\
 &= \frac{\gamma^2}{(1-\gamma)^2} \left((1-\gamma) + \gamma \frac{\sqrt{|\mathbb{S}|} b_M}{\|b\|_2} \right)^2, \tag{4.24}
 \end{aligned}$$

where $\|b\|_1 \leq \sqrt{|\mathbb{S}|} \|b\|_2$ is used in the above inequality. The largest eigenvalue λ_M of B_i defined in (2.12) is smaller than $\lambda_M < \frac{p_M}{n}$, where $p_M = \max_{i,j} P_{i,j}$ is the maximum probability of the transition matrix P [30]. This implies that

$$\frac{\lambda_M \|k\|_2^2}{\frac{1}{\gamma^2} \|b\|_2^2} \leq \frac{p_M}{n} \frac{\gamma^2}{(1-\gamma)^2} \left((1-\gamma) + \gamma \frac{\sqrt{|\mathbb{S}|} b_M}{\|b\|_2} \right)^2 \leq \frac{1}{2},$$

where the second inequality follows from the stronger assumption that $\frac{p_M}{n} \frac{\gamma^2}{(1-\gamma)^2} \left(\frac{(1-\gamma)}{\gamma} + \frac{\sqrt{|\mathbb{S}|} b_M}{\|b\|_2} \right)^2 \leq \frac{1}{2}$. Therefore,

$$\varepsilon^\circ \leq 1 + \frac{p_M}{n} \frac{\gamma^2}{(1-\gamma)^2} \left((1-\gamma) \gamma \frac{\sqrt{n} b_M}{\|b\|_2} \right)^2,$$

which completes the proof for the upper bound of ε° .

We now move on to the condition for which $\epsilon^\circ > 0$. From Theorem 2.1, it suffices to show $b^\top b + b^\top Hb/2 > 0$ and $b^\top b + b^\top Hb + b^\top Gb > 0$ to obtain $\epsilon^\circ > 0$. As a consequence of (4.23), one has

$$\begin{aligned} & b^\top b + b^\top Hb + b^\top Gb \\ &= b^\top b + \gamma^2 \sum_i (a_i + d)^\top B_i (a_i + d) - \gamma^2 \sum_i a_i^\top B_i a_i \\ &\geq b^\top b - \gamma^2 \sum_i a_i^\top B_i a_i \\ &\geq b^\top b - \gamma^2 \lambda_M \|k\|_2^2. \end{aligned}$$

Hence the denominator term in ϵ° is positive whenever $b^\top b > \gamma^2 \lambda_M \|k\|_2^2$, which holds when $n > \gamma^2 \lambda_M \frac{\|k\|_2^2}{\|b\|_2^2}$.

For $b^\top b + b^\top Hb/2$, one has

$$b^\top b + b^\top Hb/2 = b^\top b + \gamma^2 \sum_i a_i^\top B_i d \geq b^\top b - \gamma^2 \left\| \sum_i B_i a_i \right\|_1 \|d\|_\infty,$$

where the inequality follows from Holder’s inequality. Moreover, one has $\|d\|_\infty = \|(I - \gamma P)^{-1} b\|_\infty \leq \|(I - \gamma P)^{-1}\|_\infty \|b\|_\infty = \frac{p_M}{1 - \gamma}$. Thus, it suffices to bound the term $\|\sum_i B_i a_i\|_1$. We will show that

$$\left\| \sum_i B_i a_i \right\|_1 \leq 2 \frac{p_M}{n} |\mathbb{S}| \frac{1}{1 - \gamma} b_M. \tag{4.25}$$

Assuming (4.25) holds, one has

$$b^\top b + b^\top Hb/2 \geq b^\top b - \gamma^2 \left\| \sum_i B_i a_i \right\|_1 \|d\|_\infty \geq b^\top b - 2\gamma^2 \frac{1}{n} p_M |\mathbb{S}| \frac{1}{(1 - \gamma)^2} b_M^2,$$

which implies the numerator term $b^\top b + b^\top Hb/2$ is positive whenever

$$n > 2\gamma^2 p_M |\mathbb{S}| \frac{1}{(1 - \gamma)^2} \frac{b_M^2}{\|b\|_2^2}.$$

Thus, for $\epsilon^\circ > 0$, one needs

$$n > \max \left(\gamma^2 \lambda_M \frac{\|k\|_2^2}{\|b\|_2^2}, 2\gamma^2 p_M |\mathbb{S}| \frac{1}{(1 - \gamma)^2} \frac{b_M^2}{\|b\|_2^2} \right).$$

By (4.24), one can see that $n > \gamma^2 \lambda_M \frac{\|k\|_2^2}{\|b\|_2^2}$ holds if $\frac{p_M}{n} \frac{\gamma^2}{(1 - \gamma)^2} \left((1 - \gamma) + \gamma \frac{\sqrt{|\mathbb{S}|} b_M}{\|b\|_2} \right)^2 \leq \frac{1}{2}$.

Moreover, the condition that $n > 2\gamma^2 p_M |\mathbb{S}| \frac{1}{(1 - \gamma)^2} \frac{b_M^2}{\|b\|_2^2}$ is simply a restatement of the condition that $\frac{p_M}{n} \frac{\gamma^2}{(1 - \gamma)^2} \left(\frac{\sqrt{|\mathbb{S}|} b_M}{\|b\|_2} \right)^2 < \frac{1}{2}$.

Then, since $\gamma \in (0, 1)$, the theorem's assumption $\frac{p_M}{n} \frac{\gamma^2}{(1-\gamma)^2} \left(\frac{(1-\gamma)}{\gamma} + \frac{\sqrt{|\mathbb{S}|} b_M}{\|b\|_2} \right)^2 \leq \frac{1}{2}$ implies $\frac{p_M}{n} \frac{\gamma^2}{(1-\gamma)^2} \left((1-\gamma) + \gamma \frac{\sqrt{|\mathbb{S}|} b_M}{\|b\|_2} \right)^2 \leq \frac{1}{2}$ and $\frac{p_M}{n} \frac{\gamma^2}{(1-\gamma)^2} \left(\frac{\sqrt{|\mathbb{S}|} b_M}{\|b\|_2} \right)^2 < \frac{1}{2}$. Therefore, the theorem's assumption implies $\epsilon^\circ > 0$.

For the remainder of the proof, we show that (4.25) holds. Define $B_{i,+} = \frac{1}{n} \text{diag } p_i$ and $B_{i,-} = -\frac{1}{n} p_i p_i^\top$. It follows that $B_i = B_{i,+} + B_{i,-}$, and one has

$$\left\| \sum_i B_i a_i \right\|_1 \leq \left\| \sum_i B_{i,+} a_i \right\|_1 + \left\| \sum_i B_{i,-} a_i \right\|_1. \tag{4.26}$$

We first bound the first term on the right-hand side of (4.26). Let b_+, b_- be the positive and negative parts of b as in Lemma 4.3, and define

$$a_{i,+} := (I - \gamma P)^{-1} \text{diag}(e_i) b_+, \quad a_{i,-} := (I - \gamma P)^{-1} \text{diag}(e_i) b_-.$$

In particular, one has $a_{i,+} \geq \mathbf{0}$ and $a_{i,-} \leq \mathbf{0}$, as $(I - \gamma P)^{-1} \text{diag}(e_i)$ has only non-negative entries. Thus one can further bound by

$$\left\| \sum_i B_{i,+} a_i \right\|_1 \leq \left\| \sum_i B_{i,+} a_{i,+} \right\|_1 + \left\| \sum_i B_{i,+} a_{i,-} \right\|_1.$$

Due to the non-negativity of the entries in $(-a_{i,-})$ and $a_{i,+}$, the right-hand side is a monotonically non-decreasing function in the entries of $B_{i,+}$, and therefore one has

$$\begin{aligned} & \left\| \sum_i B_{i,+} a_{i,+} \right\|_1 + \left\| \sum_i B_{i,+} a_{i,-} \right\|_1 \\ & \leq \left\| \frac{1}{n} \text{diag}(p_M \mathbf{1}) \sum_i a_{i,+} \right\|_1 + \left\| \frac{1}{n} \text{diag}(p_M \mathbf{1}) \sum_i a_{i,-} \right\|_1 \\ & = \frac{p_M}{n} \left(\left\| (I - \gamma P)^{-1} b_+ \right\|_1 + \left\| (I - \gamma P)^{-1} b_- \right\|_1 \right). \end{aligned}$$

Now, note that $\left\| (I - \gamma P)^{-1} b_+ \right\|_1 = \mathbf{1}^\top (I - \gamma P)^{-1} b_+$, and $\left\| (I - \gamma P)^{-1} b_- \right\|_1 = -\mathbf{1}^\top (I - \gamma P)^{-1} b_-$. Therefore,

$$\left\| (I - \gamma P)^{-1} |b| \right\|_1 = \mathbf{1}^\top (I - \gamma P)^{-1} |b| = \mathbf{1}^\top (I - \gamma P)^{-1} (b_+ - b_-).$$

Hence one has

$$\begin{aligned} \left\| \sum_i B_{i,+} a_{i,+} \right\|_1 + \left\| \sum_i B_{i,+} a_{i,-} \right\|_1 & \leq \frac{p_M}{n} \left(\left\| (I - \gamma P)^{-1} b_+ \right\|_1 + \left\| (I - \gamma P)^{-1} b_- \right\|_1 \right) \\ & = \frac{p_M}{n} \left\| (I - \gamma P)^{-1} |b| \right\|_1 \\ & \leq \frac{p_M}{n} |\mathbb{S}| \left\| (I - \gamma P)^{-1} |b| \right\|_\infty \\ & \leq \frac{p_M}{n} |\mathbb{S}| \frac{1}{1-\gamma} b_M. \end{aligned}$$

For the second term in (4.26), one similarly has

$$\left\| \sum_i B_{i,-} a_i \right\|_1 \leq \left\| \sum_i B_{i,-} a_{i,-} \right\|_1 + \left\| \sum_i B_{i,-} a_{i,+} \right\|_1.$$

One can check

$$\begin{aligned} \left\| \sum_i B_{i,-} a_{i,-} \right\|_1 + \left\| \sum_i B_{i,-} a_{i,+} \right\|_1 &= \mathbf{1}^\top \sum_i \frac{1}{n} p_i p_i^\top (I - \gamma P)^{-1} \text{diag}(e_i) |b| \\ &= \sum_i \frac{1}{n} p_i^\top (I - \gamma P)^{-1} \text{diag}(e_i) |b| \\ &\leq \sum_i \frac{p_M}{n} \mathbf{1}^\top (I - \gamma P)^{-1} \text{diag}(e_i) |b| \\ &= \frac{p_M}{n} \mathbf{1}^\top (I - \gamma P)^{-1} |b| \\ &\leq \frac{p_M}{n} \left\| (I - \gamma P)^{-1} |b| \right\|_1 \\ &\leq \frac{p_M}{n} |\mathbb{S}| \frac{1}{1 - \gamma} b_M. \end{aligned}$$

Thus $\left\| \sum_i B_i a_i \right\|_1 \leq 2 \frac{p_M}{n} |\mathbb{S}| \frac{1}{1 - \gamma} b_M$. □

Appendix.

A.1. Condition for Convergence of Neumann Series. The spectral radius $\rho(\hat{Y})$ of \hat{Y} can be bounded by the size of state space $|\mathbb{S}|$, the number of samples n used to learn the model P and the number of possible transitions. We define κ as the largest number of transitions among all states,

$$\kappa = \max_{s \in \mathbb{S}} \{k : k = \|P_s\|_0, P_s \text{ is the } s\text{-th row of } P\}. \tag{A.1}$$

The following lemma gives the condition for $\rho(\hat{Y}) < 1$ with high probability. The proof relies on the concentration inequality of l_1 -norm of the multinomial distribution.

LEMMA A.1. *Under Assumption 2.1, for any $C > 0$ and any positive integer $q > 1$, if $n \geq \frac{2C^2 \gamma^2 \kappa}{(1 - \gamma)^2} \log(2 |\mathbb{S}|^q)$,*

$$\mathbb{P} \left[\rho(\hat{Y}) < \frac{1}{C} \right] \geq 1 - \frac{1}{|\mathbb{S}|^{q-1}}.$$

Proof. We have

$$\rho(\hat{Y}) \leq \left\| \hat{Y} \right\|_\infty \leq \left\| \hat{Z} \right\|_\infty \|A^{-1}\|_\infty = \frac{\gamma}{1 - \gamma} \left\| P - \hat{P} \right\|_\infty.$$

By the concentration inequality in [24, 32], for arbitrary $r \in [0, 1]$

$$\mathbb{P} \left[\left\| e_i^\top (\hat{P} - P) \right\|_1 \geq \frac{\sqrt{2\kappa \log 2/r}}{\sqrt{n}} \right] \leq r. \tag{A.2}$$

Taking union bound and setting $r = \frac{1}{|\mathbb{S}|^q}$ leads to

$$\mathbb{P} \left[\left\| P - \hat{P} \right\|_{\infty} \geq \frac{\sqrt{2\kappa \log 2 |\mathbb{S}|^q}}{\sqrt{n}} \right] \leq 1 - (1 - 1/|\mathbb{S}|^q)^{|\mathbb{S}|} \leq 1/|\mathbb{S}|^{q-1},$$

where the second inequality is by Bernoulli's inequality: for $r \geq 1$ and $x \leq 1$,

$$(1-x)^r \geq 1-rx.$$

The proof is completed by noticing

$$\frac{\sqrt{2\kappa \log 2 |\mathbb{S}|^q}}{\sqrt{n}} \leq \frac{1}{C} \iff n \geq \frac{2C^2 \gamma^2 \kappa}{(1-\gamma)^2} \log(2|\mathbb{S}|^q).$$

□

REMARK A.1. In particular, our goal is to show a bound of n to ensure that $\rho(\hat{Y}) < 1$ with high probability. In this case, the sample size requirement is

$$n \geq \frac{2\gamma^2 \kappa}{(1-\gamma)^2} \log(2|\mathbb{S}|^q).$$

The requirement of sample size n only grows at the rate of $O(\kappa \log(|\mathbb{S}|))$. Even though κ may grow proportionally to $|\mathbb{S}|$, one can generally assume that κ grows sublinearly with respect to $|\mathbb{S}|$. In practice, the numerical examples are more well-behaved if κ is large, and usually the convergence of the Taylor series needs only $n=1$. The bound on the spectral radius is intended for ill-behaved MDP with small κ .

REFERENCES

- [1] I. Babuska, F. Nobile, and R. Tempone, *A stochastic collocation method for elliptic partial differential equations with random input data*, SIAM J. Numer. Anal., **45(3):1005–1034**, 2007. [1](#)
- [2] I. Babuska, R. Tempone, and G.E. Zouraris, *Galerkin finite element approximations of stochastic elliptic partial differential equations*, SIAM J. Numer. Anal., **42(2):800–825**, 2004. [1](#)
- [3] J. Bäck, F. Nobile, L. Tamellini, and R. Tempone, *Stochastic spectral Galerkin and collocation methods for PDEs with random coefficients: a numerical comparison*, in J. Hesthaven and E. Ronquist (eds.), Spectral and High Order Methods for Partial Differential Equations, Springer, **43–62**, 2011. [1](#)
- [4] A. Barth, C. Schwab, and N. Zollinger, *Multi-level Monte Carlo finite element method for elliptic PDEs with stochastic coefficients*, Numer. Math., **119(1):123–161**, 2011. [1](#)
- [5] A. Cohen, R. DeVore, and C. Schwab, *Convergence rates of best N -term Galerkin approximations for a class of elliptic sPDEs*, Found. Comput. Math., **10(6):615–646**, 2010. [1](#)
- [6] M. Deisenroth and C.E. Rasmussen, *PILCO: A model-based and data-efficient approach to policy search*, Proceedings of the 28th International Conference on Machine Learning (ICML-11), Bellevue, WA, USA, **465–472**, 2011. [1](#)
- [7] J. Dick, F.Y. Kuo, and I.H. Sloan, *High-dimensional integration: the quasi-Monte Carlo way*, Acta Numer., **22:133–288**, 2013. [1](#)
- [8] D. Estep and D. Neckels, *Fast and reliable methods for determining the evolution of uncertain parameters in differential equations*, J. Comput. Phys., **213(2):530–556**, 2006. [1](#)
- [9] P. Etter and L. Ying, *Operator shifting for general noisy matrix systems*, SIAM J. Math. Data Sci., **4(4):1320–1346**, 2022. [1](#)
- [10] P.A. Etter and L. Ying, *Operator augmentation for noisy elliptic systems*, arXiv preprint, [arXiv:2010.09656](#), 2020. [1](#), [1](#)
- [11] R.G. Ghanem and P.D. Spanos, *Stochastic Finite Elements: A Spectral Approach*, Courier Corporation, 2003. [1](#)

- [12] I.G. Graham, F.Y. Kuo, D. Nuyens, R. Scheichl, and I.H. Sloan, *Quasi-Monte Carlo methods for elliptic PDEs with random coefficients and applications*, J. Comput. Phys., **230(10):3668–3694**, 2011. 1
- [13] M.D. Gunzburger, C.G. Webster, and G. Zhang, *Stochastic finite element methods for partial differential equations with random input data*, Acta Numer., **23:521–650**, 2014. 1
- [14] W. James and C. Stein, *Estimation with quadratic loss*, in S. Kotz and N.L. Johnson (eds.), Breakthroughs in Statistics, Springer, **443–460**, 1992. 1
- [15] O. Le Maître and O.M. Knio, *Spectral Methods for Uncertainty Quantification: With Applications to Computational Fluid Dynamics*, Springer Science & Business Media, **2010**. 1
- [16] S. Levine and V. Koltun, *Guided policy search*, Int. Conf. Mach. Learn., Atlanta, Georgia, USA, **1–9**, 2013. 1
- [17] Y. Marzouk, T. Moselhy, M. Parno, and A. Spantini, *Sampling via measure transport: an introduction*, in R. Ghanem, D. Higdon, and H. Owhadi (eds.), Handbook of Uncertainty Quantification, Springer, Cham, **1–41**, 2016. 1
- [18] S. Mishra and Ch. Schwab, *Sparse tensor multi-level Monte Carlo finite volume methods for hyperbolic conservation laws with random initial data*, Math. Comput., **81(280):1979–2018**, 2012. 1
- [19] H.N. Najm, *Uncertainty quantification and polynomial chaos techniques in computational fluid dynamics*, Annu. Rev. Fluid Mech., **41:35–52**, 2009. 1
- [20] H. Niederreiter, P. Hellekalek, G. Larcher, and P. Zinterhof, *Monte Carlo and quasi-Monte Carlo methods 1996*, in H. Niederreiter, P. Hellekalek, G. Larcher, and P. Zinterhof (eds.), Proceedings of a conference at the University of Salzburg, Austria, Springer Science & Business Media, **July 9–12, 1996**. 1
- [21] F. Nobile, R. Tempone, and C.G. Webster, *A sparse grid stochastic collocation method for partial differential equations with random input data*, SIAM J. Numer. Anal., **46(5):2309–2345**, 2008. 1
- [22] J. Oh, X. Guo, H. Lee, R. Lewis, and S. Singh, *Action-conditional video prediction using deep networks in Atari games*, Neur. Inf. Process. Syst., **2863–2871**, 2015. 1
- [23] B. Peherstorfer, K. Willcox, and M. Gunzburger, *Survey of multifidelity methods in uncertainty propagation, inference, and optimization*, SIAM Rev., **60(3):550–591**, 2018. 1
- [24] J. Qian, R. Fruit, M. Pirotta, and A. Lazaric, *Concentration inequalities for multinoulli random variables*, arXiv preprint, [arXiv:2001.11595](https://arxiv.org/abs/2001.11595), 2020. A.1
- [25] D. Silver, A. Huang, C.J. Maddison, A. Guez, L. Sifre, G. van den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, et al., *Mastering the game of Go with deep neural networks and tree search*, Nature, **529(7587):484–489**, 2016. 1
- [26] D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton, et al., *Mastering the game of Go without human knowledge*, Nature, **550(7676):354–359**, 2017. 1
- [27] A.M. Stuart, *Inverse problems: a Bayesian perspective*, Acta Numer., **19:451–559**, 2010. 1
- [28] R.S. Sutton, *Dyna, an integrated architecture for learning, planning, and reacting*, ACM Sigart Bulletin, **2(4):160–163**, 1991. 1, 2.1
- [29] R.S. Sutton and A.G. Barto, *Reinforcement Learning: An Introduction*, MIT Press, **2018**. 1, 1
- [30] G.S. Watson, *Spectral decomposition of the covariance matrix of a multinomial*, J. R. Stat. Soc., B: Stat. Methodol., **58(1):289–291**, 1996. 4.4
- [31] M. Watter, J.T. Springenberg, J. Boedecker, and M. Riedmiller, *Embed to control: A locally linear latent dynamics model for control from raw images*, Neur. Inf. Process. Syst., **2746–2754**, 2015. 1
- [32] T. Weissman, E. Ordentlich, G. Seroussi, S. Verdu, and M.J. Weinberger, *Inequalities for the L_1 deviation of the empirical distribution*, Hewlett-Packard Labs, Tech. Rep., **2003**. A.1
- [33] D. Xiu, *Numerical Methods for Stochastic Computations*, Princeton University Press, **2010**. 1
- [34] D. Xiu and J.S. Hesthaven, *High-order collocation methods for differential equations with random inputs*, SIAM J. Sci. Comput., **27(3):1118–1139**, 2005. 1
- [35] D. Xiu and G. Em Karniadakis, *The Wiener–Askey polynomial chaos for stochastic differential equations*, SIAM J. Sci. Comput., **24(2):619–644**, 2002. 1