2025 IEEE 64th Conference on Decision and Control (CDC)
December 10-12, 2025. Rio de Janeiro, Brazil

Learning with Imperfect Models: When Multi-step Prediction Mitigates Compounding Error

Anne Somalwar[1], Bruce D. Lee, George J. Pappas, Nikolai Matni

2025 IEEE 64th Conference on Decision and Control (CDC) | 979-8-3315-2627-6/25/$31.00 ©2025 IEEE | DOI: 10.1109/CDC57313.2025.11312354

***Abstract*— Compounding error, where small prediction mistakes accumulate over time, presents a major challenge in learning-based control. For example, this issue often limits the performance of model-based reinforcement learning and imitation learning. One common approach to mitigate compounding error is to train multi-step predictors directly, rather than relying on autoregressive rollout of a single-step model. However, it is not well understood when the benefits of multi-step prediction outweigh the added complexity of learning a more complicated model. In this work, we provide a rigorous analysis of this trade-off in the context of linear dynamical systems. We show that when the model class is well-specified and accurately captures the system dynamics, single-step models achieve lower asymptotic prediction error. On the other hand, when the model class is misspecified due to partial observability, direct multi-step predictors can significantly reduce bias and thus outperform single-step approaches. These theoretical results are supported by numerical experiments, wherein we also (a) empirically evaluate an intermediate strategy which trains a single-step model using a multi-step loss and (b) evaluate performance of single step and multi-step predictors in a closed loop control setting.**

I. Introduction

A typical approach to time series forecasting is to fit a one-step ahead prediction model and apply it recursively to obtain predictions over multiple time steps. In doing so, small errors may compound over time, leading to poor long-horizon prediction. This issue hinders the application of such single-step models in e.g., controller design.

By directly training multi-step models to predict longer horizons, the issue of compounding error can be mitigated. The main drawback of doing so is that the number of parameters for a direct multi-step predictor scales with the prediction horizon, thus potentially requiring more data to achieve a desired prediction performance. While this tradeoff between prediction horizon accuracy and data requirements is broadly known to exist, it is primarily studied from an empirical perspective [1, 2]. We therefore lack principled guidance for exactly when direct multi-step prediction should be preferred over autoregressive rollout of single-step models. Motivated by this challenge, we provide a rigorous comparison of the sample efficiency of learning multi-step predictors with that of learning single-step predictors in the setting of a linear dynamical system.

A. Related Work

a) Multi-step Identification: The goal of system identification is to use data to learn a model that can be used for forecasting or control [3]. To this end, one typically wants to select a model from the hypothesis class that minimizes the simulation error, i.e., the cumulative prediction error over all future time steps. Due to the computational challenge of doing so, it is much more common to instead learn a model which minimizes the single-step prediction error and apply it autoregressively [4]. However, such approaches tend to generalize poorly if the underlying data generating process does not belong to the hypothesis class. This has motivated the application of algorithms such as Data as Demonstrator (DaD) which use approaches from imitation learning to train a predictor to self-correct its prediction error at each step, leading to improved empirical performance [5, 6].

Direct learning of prediction models for each time-step in the prediction horizon partially bypasses the issue of compounding error, and has proven successful for both model predictive control [7–9] and value approximation in MDPs [10]. Empirical studies of single-step and multi-step dynamics models learned with various neural network architectures have also been conducted. Namely, Lambert et al. [2] investigate the performance of recursive application of single-step models parameterized by neural networks in a handful of examples and characterize circumstances which may worsen the effects of compounding error. Chandra et al. [11] give a comparison of various deep learning architectures for predicting multiple steps of a time series and show the efficacy of bidirectional and encoder-decoder LSTM networks.

These empirical studies underscore the importance of careful consideration when designing a model for multi-step prediction. Our work studies this question in stylized settings, allowing us to clearly demonstrate the potential drawbacks and benefits of direct multi-step prediction as compared to more traditional single-step approaches.

b) Learning-Enabled Control: While the issue of compounding error has long been studied in system identification and control, it has resurfaced as a prominent issue in the learning community, exacerbated by the use of neural networks as function approximators. In model-based reinforcement learning, it has been observed that synthesized controllers may exploit compounding errors of the learned model [12, 13], motivating numerous heuristics for accounting for the model error during policy synthesis [14–17]. Lambert et al. [1] instead propose learning a direct multi-step predictor parametrized by a low dimensional decision variable which may be optimized online. The issue of a mismatch between the hypothesis class and the underlying data generating process also poses a challenge for behavior cloning.

[1] All authors are with the University of Pennsylvania. Emails : {somalwar, brucele, nmatni, pappasg}@seas.upenn.edu

Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 13:03:03 UTC from IEEE Xplore. Restrictions apply.

For example, incorrectly assuming that the demonstrator is Markovian can lead to a policy that deviates substantially from the demonstrator [18, 19]. This can be remedied in part by replacing the standard behavior cloning objective with an objective that predicts the expert actions for multiple timesteps to maintain temporal consistency, resulting in so-called *action-chunking* based approaches [19, 20]. We draw inspiration from these studies, and consider instances of linear systems with either Markovian or non-Markovian observations to compare direct multi-step prediction and autoregressive evaluation of single-step predictors.

B. Contributions

We provide the first quantitative comparison of the multi-step prediction error incurred by directly learned multi-step predictors against that incurred by autoregressive evaluation of a single-step predictor. In particular:

- We provide an asymptotic characterization of the multi-step prediction error for the two methods in the setting of a fully observed dynamical system. Our results show that for stable systems with a small spectral radius, the prediction error of autoregressive evaluation decays significantly faster with increasing data than that of direct multi-step prediction. This benefit diminishes as the spectral radius increases to one.
- We characterize the prediction error for the two methods applied to a partially observed dynamical system which is incorrectly assumed to be fully observed by the hypothesis class, thus addressing the issue of trajectory prediction under misspecification due to an unjustified Markovian assumption. These results demonstrate that multi-step predictors may enjoy significantly lower bias in the face of model misspecification.
- We conduct numerical experiments that compare the aforementioned approaches with an additional standard baseline: fitting a single-step model that minimizes the multi-step prediction error.
- We empirically evaluate the performance of autoregressive rollouts of single step predictors and multi-step predictors in a closed loop control setting.

Our results, while limited to stylized settings, capture key properties of real systems—such as partial observability and misspecification—and thus provide useful guidance to practitioners navigating the complex design space of data-driven multi-step predictors.

Notation: $\mathcal{N}(\mu, \sigma^2)$ denotes a normal distribution with mean μ and variance σ^2, $\rho(\cdot)$ denotes the spectral radius of a matrix, $\|\cdot\|$ denotes the vector Euclidean norm, $\|\cdot\|_F$ denotes the matrix Frobenius norm, $\mathsf{vec}(\cdot)$ denotes the vectorization of a matrix, and $\otimes$ denotes the Kronecker product.

Extended manuscript with full proofs can be found at: https://arxiv.org/abs/2504.01766

II. Problem Formulation

Consider the linear time invariant dynamical system

$$\begin{aligned} x_{t+1} &= Ax_t + Bu_t + B_w w_t \qquad t \in \mathbb{Z}^+ \\ y_t &= Cx_t + D_v v_t, \qquad\qquad t \in \mathbb{Z}^+ \end{aligned} \tag{1}$$

with state $x_t \in \mathbb{R}^{d_x}$, input $u_t \in \mathbb{R}^{d_u}$, observation $y_t \in R^{d_y}$, process noise $w_t \overset{iid}{\sim} \mathcal{N}(0, I_{d_x})$, sensor noise $v_t \overset{iid}{\sim} \mathcal{N}(0, I_{d_y})$, and initial condition $x_0 = 0$. We assume that (A, C) is observable, that $(A, \begin{bmatrix} B, B_w \end{bmatrix})$ is controllable, $\rho(A) < 1$, and that the control inputs are selected randomly as $u_t \overset{iid}{\sim} \mathcal{N}(0, I_{d_u})$.

We assume that the dynamics (1) are unknown, and our goal is to learn a predictor that forecasts a horizon H of future observations using past observations. To this end, we suppose that we are given a dataset $\mathcal{D}_N = \{(y_t, u_t)\}_{t=1}^N$ collected from a training rollout of (1) which will be used to determine a function $\hat{f}_H$ belonging to a hypothesis class $\mathcal{F}_H$. This function will be used to predict $y_{t+1:t+H}$ given $y_{1:t}$ and $u_{1:t+H-1}$.

The quality of the learned function will be measured by the loss

$$L(\hat{f}_H) \triangleq \bar{\mathbf{E}} \left\| y_{t+1:t+H} - \hat{f}_H(y_{1:t}, u_{1:t+H-1}) \right\|^2, \tag{2}$$

where the operator $\bar{\mathbf{E}}$ is defined as

$$\bar{\mathbf{E}}[f(t)] \triangleq \lim_{T\to\infty} \mathbf{E} \frac{1}{T} \sum_{t=0}^{T} f(t),$$

and the expectation is taken over an evaluation rollout of system (1) that is independent of the dataset $\mathcal{D}_N$. By ergodicity of the process (1), this is equivalent to taking an expectation under the steady state distribution for the system.

To provide rigorous understanding of situations where multi-step prediction does or does not help, we consider a simplified setting in which the hypothesis class consists of static linear predictors, i.e., a function $f_H \in \mathcal{F}_H$, is given by

$$f_H(y_{1:t}, u_{1:t+H-1}) = G \begin{bmatrix} y_t \\ u_{t:t+H-1} \end{bmatrix}, \tag{3}$$

for a matrix $G \in S \subseteq \mathbb{R}^{Hd_y \times (d_y + Hd_u)}$. Here the subspace S encodes whether we are fitting a multi-step or single-step model: we provide explicit parameterizations for these model-classes in the next subsections. Our restriction to static linear predictors of the form (3) assumes that the observation sequence is Markovian, i.e., that a history of observations is unnecessary to predict the future trajectory. In the sequel, we slightly abuse notation and denote the loss (2) incurred by a predictor (3) defined by matrix $\hat{G}$ by $L(\hat{G})$.

We consider two settings: one where the Markovian assumption is justified ($C = I$ and $D_v = 0$), resulting in a well-specified problem, and one where it is not justified ($C \neq I$ and $D_v \succ 0$), resulting in a misspecified problem. In these two settings, we compare the H step prediction error (2) incurred by a learned single-step model rolled out for H timesteps to that incurred by a directly learned H-step model.

Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 13:03:03 UTC from IEEE Xplore. Restrictions apply.

A. Single-step Predictors

The single-step approach first solves

$$\begin{bmatrix}\hat{G}_y & \hat{G}_u\end{bmatrix} = \underset{\substack{G_y\in\mathbb{R}^{d_y\times d_y}\\ G_u\in\mathbb{R}^{d_y\times d_u}}}{\operatorname{argmin}} \sum_{t=1}^{N-1} \left\| y_{t+1} - \begin{bmatrix}G_y & G_u\end{bmatrix}\begin{bmatrix}y_t\\ u_t\end{bmatrix}\right\|^2. \tag{4}$$

Using this model, one can predict $y_{t+1:t+H}$ by rolling out $\begin{bmatrix}\hat{G}_y & \hat{G}_u\end{bmatrix}$ autoregressively:

$$\begin{aligned}\hat{y}_{t+1} &= \begin{bmatrix}\hat{G}_y & \hat{G}_u\end{bmatrix}\begin{bmatrix}y_t\\ u_t\end{bmatrix}\\ \hat{y}_{t+2} &= \begin{bmatrix}\hat{G}_y & \hat{G}_u\end{bmatrix}\begin{bmatrix}\hat{y}_{t+1}\\ u_{t+1}\end{bmatrix}\\ &\vdots\\ \hat{y}_{t+H} &= \begin{bmatrix}\hat{G}_y & \hat{G}_u\end{bmatrix}\begin{bmatrix}\hat{y}_{t+H-1}\\ u_{t+H-1}\end{bmatrix}.\end{aligned} \tag{5}$$

The resulting H-step predictor can be composed to form a direct mapping from the data to the predicted trajectory as

$$\hat{G}_N^{SS} = \begin{bmatrix}\hat{G}_y & \hat{G}_u & 0 & \dots & 0\\ \hat{G}_y^2 & \hat{G}_y\hat{G}_u & \hat{G}_u & \dots & 0\\ \vdots & & & \ddots & \\ \hat{G}_y^H & \hat{G}_y^{H-1}\hat{G}_u & \hat{G}_y^{H-2}\hat{G}_u & \dots & \hat{G}_u\end{bmatrix}. \tag{6}$$

As past predictions become part of the regressor for future predictions, this approach often suffers from compounding error.

B. Multi-step Predictors

The issue of compounding error from autoregressive roll-out of a single-step model motivates direct multi-step approaches which directly minimize the H step prediction error:

$$\hat{G}_N^{MS} = \underset{G\in S}{\operatorname{argmin}} \sum_{t=1}^{N-H} \left\| y_{t+1:t+H} - G\begin{bmatrix}y_t\\ u_{t:t+H-1}\end{bmatrix}\right\|^2, \tag{7}$$

for $S \subseteq \mathbb{R}^{Hd_y\times(d_y+Hd_u)}$. We consider the function class which fits H distinct predictors, one for each step in the prediction horizon. This amounts to setting $S = \mathbb{R}^{Hd_y\times(d_y+Hd_u)}$.[1]

There is a tradeoff induced by fitting multi-step predictors rather than single-step predictors. In particular, the single-step predictor is subject to compounding error, while the complexity of the above identification problem increases for longer horizons. We study this tradeoff in the two aforementioned settings: a system with Markovian observations, and a system with non-Markovian observations. Due to the Markovian assumption for the identification problem, these cases serve as instances where the identification problem is well-specified and misspecified, respectively.

[1]One could impose the causality structure, i.e. that S has a triangular structure. We refrain from doing so for simplicity, and due to the fact that future inputs are independent of the past.

C. Intermediate Formulations

Rather than fitting independent predictors for every timestep, one can instead formulate a hypothesis class for multi-step prediction with lower complexity. In particular, one could impose additional structure on S, e.g.

$$S = \left\{ \begin{bmatrix}G_y & G_u & 0 & \dots & 0\\ G_y^2 & G_yG_u & G_u & \dots & 0\\ \vdots & & & \ddots & \\ G_y^H & G_y^{H-1}G_u & G_y^{H-2}G_u & \dots & G_u\end{bmatrix} \middle| \begin{matrix}G_y\in\mathbb{R}^{d_y\times d_y}\\ G_u\in\mathbb{R}^{d_y\times d_u}\end{matrix} \right\}. \tag{8}$$

This consists of functions which take the form of a single-step predictor that is applied auto-regressively. In contrast to the single-step approach of Section II-A, solving (7) with this choice of S consists of a multi-step loss function for a class of single-step predictors, a common approach to mitigate the compounding error issue [4] without increasing the number of parameters that must be learned. We study the loss (2) of the predictors (7) fit with classes (8) in numerical experiments and leave analytically characterizing the asymptotic prediction error for this predictor to future work.

III. Well-Specified Setting

In this section, we study the well-specified setting in which the Markovian assumption is valid. In particular, we restrict system (1) to be a fully observed system by assuming that $C = I$ and $D_v = 0$ so $y_t = x_t$ for all t.

To compare the single-step and multi-step approaches in this setting, we first observe that either predictor $\hat{f}_H$ is defined in terms of a linear map $\hat{G}$ applied to the vector $\begin{bmatrix}x_t^\top & u_{t:t+H-1}^\top\end{bmatrix}^\top$. Therefore the loss (2) may be written

$$L(\hat{G}) = \bar{\mathbf{E}}\left\| x_{t+1:t+H} - \hat{G}\begin{bmatrix}x_t\\ u_{t:t+H-1}\end{bmatrix}\right\|^2. \tag{9}$$

Rolling out the dynamics, we find that

$$x_{t+1:t+H} = G^\star\begin{bmatrix}x_t\\ u_{t:t+H-1}\end{bmatrix} + \Gamma_w w_{t:t+H-1},$$

where

$$G^\star = \begin{bmatrix}A & B & & \\ A^2 & AB & B & \\ & \vdots & & \ddots \\ A^H & A^{H-1}B & \dots & B\end{bmatrix},$$

and

$$\Gamma_w = \begin{bmatrix}B_w & & & \\ AB_w & B_w & & \\ \vdots & & \ddots & \\ A^{H-1}B_w & \dots & & B_w\end{bmatrix}.$$

Then expanding $x_{t+1:t+H}$ in equation (9), and using the independence of $w_{t:t+H-1}$ from x_t and $u_{t:t+H-1}$, we conclude that

$$L(\hat{G}) = \bar{\mathbf{E}}\left\|(\hat{G} - G^\star)\begin{bmatrix}x_t\\ u_{t:t+H-1}\end{bmatrix}\right\|^2 + \bar{\mathbf{E}}\left\|\Gamma_w w_{t:t+H-1}\right\|^2$$

Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 13:03:03 UTC from IEEE Xplore. Restrictions apply.

$$= \left\|(\hat{G} - G^\star)\Sigma_z^{1/2}\right\|_F^2 + \|\Gamma_w\|_F^2,$$

where $\Sigma_z = \bar{\mathbf{E}}\begin{bmatrix} x_t \\ u_{t:t+H-1} \end{bmatrix}\begin{bmatrix} x_t \\ u_{t:t+H-1} \end{bmatrix}^\top$ is the stationary covariance for the regressor. Consequently, the discrepancy between the single-step and multi-step predictors is contained in the term $\left\|(\hat{G} - G^\star)\Sigma_z^{1/2}\right\|_F^2$. We study the behavior of this term asymptotically, where $\hat{G}$, or equivalently $\hat{G}_N$, is an operator learned on the dataset of size N.[2] In particular, we examine

$$\lim_{N\to\infty} N\,\mathbf{E}\left[\left\|(\hat{G}_N - G^\star)\Sigma_z^{1/2}\right\|_F^2\right],$$

for the single-step and multi-step predictors $\hat{G}_N^{SS}$ and $\hat{G}_N^{MS}$, respectively, where the expectation is taken over the dataset used to fit $\hat{G}_N$.

The reducible error of the multi-step predictor is characterized by the following proposition.

Proposition III.1. *The reducible asymptotic error of the multi-step predictor $\hat{G}_N^{MS}$ is given by*

$$\lim_{N\to\infty} N\mathbf{E}\left[\left\|(\hat{G}_N^{MS} - G^\star)\Sigma_z^{1/2}\right\|_F^2\right] = \operatorname{tr}\big(\Gamma_w((M_{MS} + Hd_{\mathsf{u}}I_H)\otimes I_{d_{\mathsf{x}}})\Gamma_w^\top\big),$$

where $M_{MS} \in \mathbb{R}^{H\times H}$ is the matrix with entry (i,j) given by $M_{MS}^{ij} = \operatorname{tr}(A^{|i-j|})$.

The above result shows that the error decays asymptotically at a rate of $1/N$. The scaling is characterized by the trace expression, which represents the asymptotic covariance of the estimation error; importantly, it grows with the horizon H (note the $(M_{MS}+Hd_uI_H)$ term). The error of the single-step predictor is characterized below.

Proposition III.2. *The asymptotic error of the single-step predictor $\hat{G}_N^{SS}$ is given by*

$$\lim_{N\to\infty} N\mathbf{E}\left[\left\|(\hat{G}_N^{SS} - G^\star)\Sigma_z^{1/2}\right\|_F^2\right] = \operatorname{tr}(\Gamma_w((M_{SS} + d_{\mathsf{u}}I_H)\otimes I_{d_{\mathsf{x}}})\Gamma_w^T),$$

where $M_{SS} \in \mathbb{R}^{H\times H}$ is the matrix with entry (i,j) given by

$$\operatorname{tr}\left(\left(I - \Sigma_x^{-1}\sum_{\ell=0}^{\min\{i,j\}-2} A^\ell B_w B_w^\top (A^\ell)^\top\right)(A^{|j-i|})^\top\right).$$

Again, the error decays at a rate $1/N$. In contrast to the multi-step predictor, the asymptotic scaling of the single-step prediction error has the quantity $M_{SS} + d_{\mathsf{u}}I_H$ inside the trace. This means that the multi-step predictor suffers an extra factor of H in the input term. Additionally the matrix M_{MS} for the multi-step case has entries which decay as the distance to the diagonal increases, while M_{SS} has entries which decay as the distance to the upper left element increases. Roughly, this indicates that for very stable systems M_{SS} should become smaller than M_{MS}. We make this concrete in the sequel.

[2]We sometimes omit the subscript N on $\hat{G}_N$ to ease notational burden.

A. Comparison

We can express the quadratic form defining the reducible portion of the error in Proposition III.1 as the reducible error in Proposition III.2 plus the additional term $\operatorname{tr}(\Gamma_w((M_{MS} - M_{SS} + (H-1)d_uI_H)\otimes I_{d_X})\Gamma_w^T)$. Note that $(M_{MS} - M_{SS})$ is the matrix with entry (i,j) given by $\operatorname{tr}(\Sigma_x^{-1}\sum_{\ell=0}^{\min\{i,j\}-2} A^\ell B_w B_w^\top (A^\ell)^\top (A^{|j-i|})^\top)$. Let $v_\ell = \mathsf{vec}\left(\Sigma_x^{-1/2}A^\ell B_w\right)$. Then the aforementioned matrix is equal to the gram matrix defined by

$$\begin{bmatrix} 0 \\ v_0^\top \\ v_1^\top \\ \vdots \\ v_{H-1}^\top \end{bmatrix}\begin{bmatrix} 0 \\ v_0^\top \\ v_1^\top \\ \vdots \\ v_{H-1}^\top \end{bmatrix}^\top + \begin{bmatrix} 0 \\ 0 \\ v_0^\top \\ \vdots \\ v_{H-2}^\top \end{bmatrix}\begin{bmatrix} 0 \\ 0 \\ v_0^\top \\ \vdots \\ v_{H-2}^\top \end{bmatrix}^\top + \cdots + \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ v_0^\top \end{bmatrix}\begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ v_0^\top \end{bmatrix}^\top,$$

and is therefore positive semidefinite. As a result, we see that a multi-step predictor is less efficient than a single-step predictor, and that the efficiency gap grows linearly with the prediction horizon H. This scaling quantitatively captures that the direct multi-step predictor has a number of parameters which scales with H.

To better understand the role of system stability in determining this efficiency gap, we consider the special case of a scalar system without inputs. Here the difference between the statistical efficiency of the single-step predictor and the multi-step predictor is characterized by the difference between the matrices

$$M_{SS} = \begin{bmatrix} 1 & a & a^2 & \dots & a^{H-1} \\ a & a^2 & a^3 & \dots & a^H \\ \vdots & & \ddots & & \\ a^{H-1} & & \dots & & a^{2(H-1)} \end{bmatrix}$$

and

$$M_{MS} = \begin{bmatrix} 1 & a & a^2 & \dots & a^{H-1} \\ a & 1 & a & \dots & a^{H-2} \\ \vdots & & \ddots & & \\ a^{H-1} & & \dots & & 1 \end{bmatrix},$$

from which we conclude that the difference between the two diminishes as $|a| \to 1$, i.e. as the system approaches marginal stability.

IV. Misspecified Setting

We again consider a single-step predictor and a multi-step predictor applied to the measurement and sequence of future inputs. However, we now examine the general case in which measurements do not provide full state information by reincorporating partial observations, as specified by $C \in \mathbb{R}^{d_{\mathsf{y}}\times d_{\mathsf{x}}}$ and $D_vD_v^\top \succ 0$, into the dynamics (1). Due to the Markovian assumption made in fitting the predictor, this represents a misspecified setting. To ease notational burden

Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 13:03:03 UTC from IEEE Xplore. Restrictions apply.

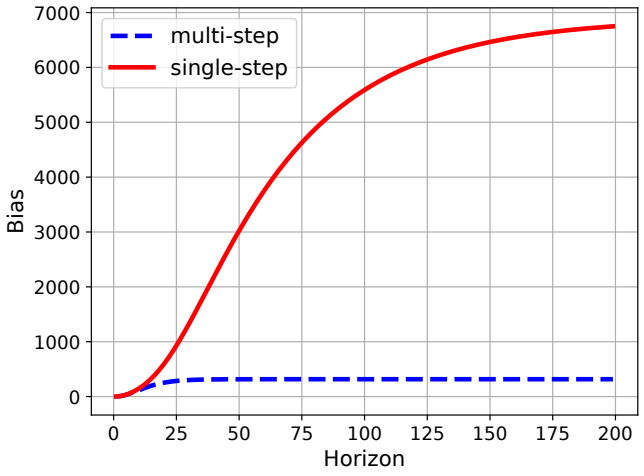


Fig. 1: Comparison of the bias from the single and multi-step estimators across different horizons for the system defined in Example IV.1.

we restrict attention to the setting without inputs and set $B = 0$, although our analysis can be extended naturally to the $B \neq 0$ case.

In this setting, the loss is given by

$$L(\hat{G}) = \bar{\mathbf{E}} \left\| y_{t+1:t+H} - \hat{G} y_t \right\|^2. \tag{10}$$

We rewrite the dynamics in innovations form

$$\begin{aligned} \hat{x}_{t+1} &= (A - KC)\hat{x}_t + Ky_t = A\hat{x}_t + KD_e e_t \\ y_t &= C\hat{x}_t + D_e e_t, \end{aligned}$$

where e_t is standard normal noise that is independent across time, K is the Kalman gain defined as $K = ASC^\top (CSC^\top + R)^{-1}$, S is the stabilizing solution to the Riccati equation defined by A, C, $D_w D_w^\top$ and $D_v D_v^\top$, and $D_e = (CSC^\top + D_v D_v^\top)^{1/2}$. Then

$$y_{t+1:t+H} = \Phi \hat{x}_t + G^\star y_t + \Gamma_e e_{t+1:t+H},$$

where

$$\Phi = \begin{bmatrix} C(A-KC) \\ CA(A-KC) \\ \vdots \\ CA^{H-1}(A-KC) \end{bmatrix}, G^\star = \begin{bmatrix} CK \\ CAK \\ \vdots \\ CA^{H-1}K \end{bmatrix},$$

$$\Gamma_e = \begin{bmatrix} D_e & 0 & \dots & 0 \\ CKD_e & D_e & \dots & \\ \vdots & & \ddots & \\ CA^{H-2}KD_e & \dots & CKD_e & D_e \end{bmatrix}.$$

Under these definitions, and exploiting that innovations are independent across time, the error (10) is given by

$$\begin{aligned} L(\hat{G}) &= \bar{\mathbf{E}} \left\| \Phi \hat{x}_t + (G^\star - \hat{G}) y_t + \Gamma_e e_{t+1:t+H} \right\|^2 \\ &= \bar{\mathbf{E}} \left\| \Phi \hat{x}_t + (G^\star - \hat{G}) y_t \right\|^2 + \|\Gamma_e\|_F^2 . \end{aligned}$$

Expanding $y_t = C\hat{x}_t + D_e e_t$,

$$\begin{aligned} &L(\hat{G}) \\ &= \left\| (\Phi + (G^\star - \hat{G})C)\Sigma_{\hat{x}}^{1/2} \right\|_F^2 + \left\| (G^\star - \hat{G}) D_e \right\|_F^2 + \|\Gamma_e\|_F^2 , \end{aligned}$$

where $\Sigma_{\hat{x}}$ is the stationary covariance of $\hat{x}_t$. When $\hat{G}$, or equivalently $\hat{G}_N$, is learned on the dataset of size N, we can decompose this quantity into an irreducible component, and a component which decays to zero as the amount of data $N \to \infty$. Denoting the irreducible component by $B(\hat{G}_N) \triangleq \lim_{N\to\infty} \mathbf{E}\, L(\hat{G}_N)$ and the reducible component by $\varepsilon(\hat{G}_N) \triangleq L(\hat{G}_N) - B(\hat{G}_N)$, we decompose

$$L(\hat{G}_N) = B(\hat{G}_N) + \varepsilon(\hat{G}_N).$$

Unlike the well-specified setting, the irreducible component $B(\hat{G}_N)$ differs depending on whether we fit a single-step model or direct multi-step model. We therefore focus on comparing these bias terms rather than the rate of convergence, since this captures the fundamental difference between the two models. See Section I-C (multi-step) and Section I-D (single-step) of the extended manuscript for characterizations of the rate of decay of the reducible errors $\lim_{N\to\infty} N\, \mathbf{E}[\varepsilon(\hat{G}_N)]$.

The irreducible error for the multi-step predictor is characterized in the following proposition.

Proposition IV.1. *The irreducible error for the multi-step predictor $\hat{G}_N^{MS}$ is given by*

$$B(\hat{G}_N^{MS}) = \operatorname{tr}(\Phi(\Sigma_{\hat{x}} - \Sigma_{\hat{x}} C^\top \Sigma_y^{-1} C \Sigma_{\hat{x}})\Phi^\top) + \|\Gamma_e\|_F^2 .$$

For the single-step predictor, the irreducible error is characterized as follows.

Proposition IV.2. *The irreducible error for the single-step predictor $\hat{G}_N^{SS}$ is given by*

$$\begin{aligned} B(\hat{G}_N^{SS}) &= \operatorname{tr}((\Phi + MC)\Sigma_{\hat{x}}(\Phi + MC)^\top) \\ &\quad + \operatorname{tr}(MD_e D_e^\top M^\top) + \|\Gamma_e\|_F^2 , \end{aligned}$$

where

$$M = G^\star - \begin{bmatrix} CA\Sigma_x C^\top \Sigma_y^{-1} \\ \vdots \\ (CA\Sigma_x C^\top \Sigma_y^{-1})^H \end{bmatrix}, \tag{11}$$

and Σ_x is the stationary covariance of x_t.

A. Comparison

In contrast to the well-specified setting, the dominant discrepancy between the two predictors in the presence of misspecification is the bias term. Note that the bias from the multi-step predictor given in Proposition IV.1 is equal to

$$\min_M \operatorname{tr}((\Phi + MC)\Sigma_{\hat{x}}(\Phi + MC)^\top) + \operatorname{tr}(MD_e D_e^\top M^\top),$$

and therefore, the irreducible error for the direct multi-step never exceeds that of its single-step counterpart.

Due to the dependence of M on powers of $CA\Sigma_x C^\top \Sigma_y^{-1}$ in the single-step model's bias, the spectral radius of this quantity dictates how the bias scales with the horizon.

Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 13:03:03 UTC from IEEE Xplore. Restrictions apply.

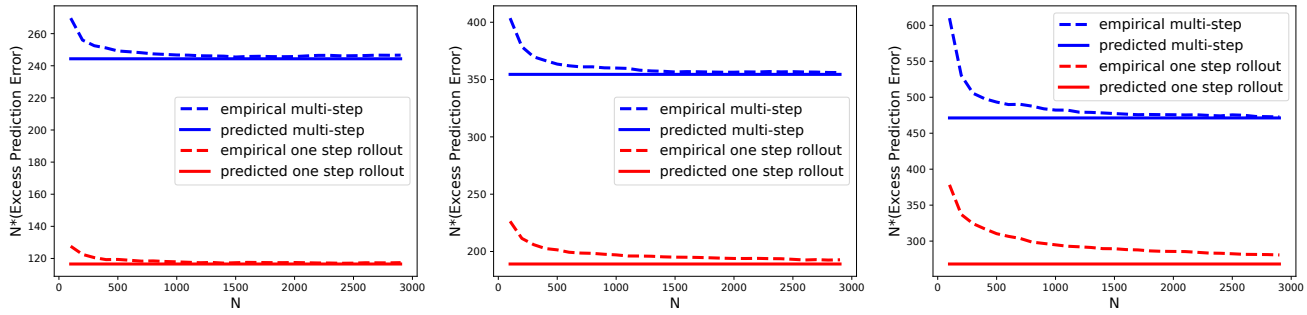


Fig. 2: Convergence of $N\,\mathbf{E}[L(\hat{f}_H)]$ to the reducible prediction errors given in Proposition III.1 (multi-step predictor) and Proposition III.2 (single-step rollout) for the system defined by Equation (12) with $a = 0.5, 0.75, 0.9$ (left to right) and horizon $H = 5$.

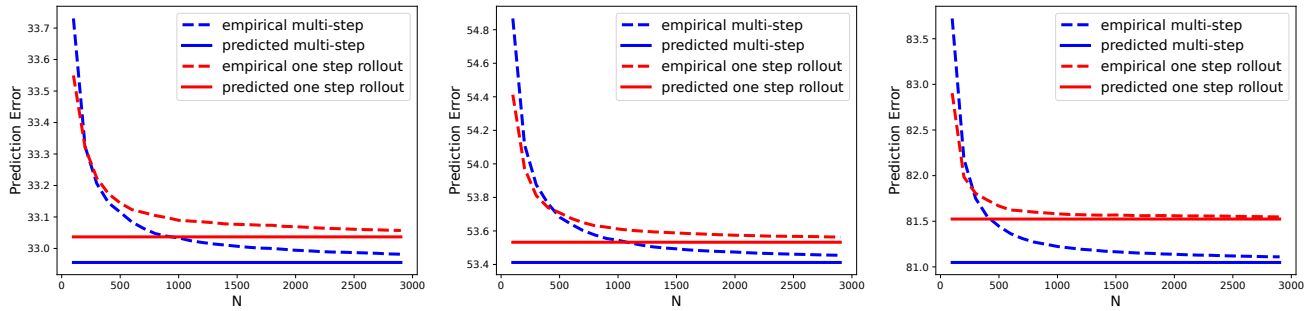


Fig. 3: Convergence of $\mathbf{E}[L(\hat{f}_H)]$ to the irreducible prediction errors given in Proposition IV.1 (multi-step predictor) and Proposition IV.2 (single-step rollout) for the system defined by Equation (12) with $a = 0.5, 0.75, 0.9$ (left to right) and horizon $H = 5$.

Lemma I.1 shows that the quantity $CA\Sigma_x C^\top \Sigma_y^{-1}$ which the estimate $\hat{G}_y$ converges to satisfies $\rho(CA\Sigma_x C^\top \Sigma_y^{-1}) \leq 1$ if A is stable. Despite this, $CA\Sigma_x C^\top \Sigma_y^{-1}$ can feature a spectral radius much larger than that of A, as demonstrated in the following example.

Example IV.1. *Consider the system defined by*

$$A = \begin{bmatrix} 0.9 & 1.0 \\ 0.0 & 0.9 \end{bmatrix}, \Sigma_w = I_2, C = \begin{bmatrix} 1, 0 \end{bmatrix}, \Sigma_v = 1.$$

We find that $\rho(CA\Sigma_x C^\top \Sigma_y^{-1}) = 0.99$*, though* $\rho(A) = 0.9$.

As a consequence of this fact, the gap in bias between the multi-step error and the single-step error can grow with the horizon for moderate H. This is illustrated for the above example in Figure 1.

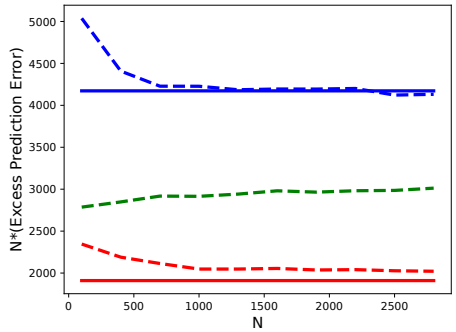
V. Numerical Experiments

To validate the bounds presented in the previous sections, we consider the system defined by

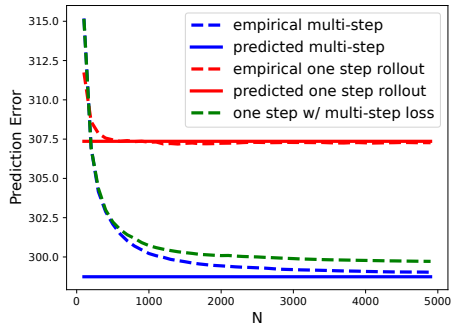
$$A = \begin{bmatrix} a & 1.0 \\ 0.0 & 0.75 \end{bmatrix}, \Sigma_w = I_2 \tag{12}$$

with $B = \begin{bmatrix} 0 & 1 \end{bmatrix}^\top, C = I_2, \Sigma_v = 0$ in the well-specified setting and, alternatively, $B = 0, C = [1, 0], \Sigma_v = 1$ in the misspecified setting.

Well-specified: Figure 2 illustrates the well-specified setting. Specifically, we estimate $N\,\mathbf{E}[L(\hat{f}_H)]$ by averaging over $30,000$ datasets D_N for $N \in \{1, ..., 3000\}$ to demonstrate



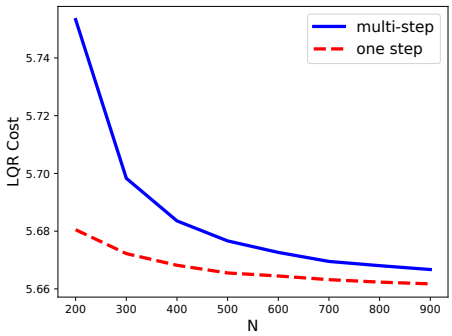
(a) Well-specified case



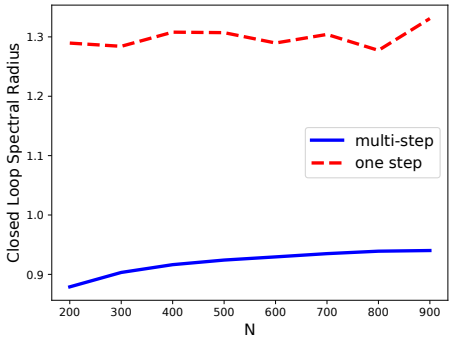
(b) Misspecified case

Fig. 4: Comparison of the rate of decay for the error in the well-specified case (a) and the total error in the misspecified case (b) for the direct multi-step predictor, and the single-step predictor trained with a single-step loss or a multi-step loss.

Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 13:03:03 UTC from IEEE Xplore. Restrictions apply.



(a) Infinite-horizon LQR cost, well-specified case



(b) Spectral radius of the closed loop system, misspecified case

Fig. 5: Infinite-horizon LQR performance in the well-specified case (a) and the misspecified case (b). In (b), closed loop spectral radius greater than 1 for the one step predictor implies infinite LQR cost.

convergence to the reducible errors given in Proposition III.1 and Proposition III.2. In these figures, we fix $H = 5$ and vary a across $0.5, 0.75$, and 0.9.

Misspecified: Figure 3 illustrates the misspecified setting. Specifically, we estimate $\mathbf{E}[L(\hat{f}_H)]$ by averaging over $1,000$ datasets D_N for $N \in \{1, ..., 3000\}$ to demonstrate convergence to the irreducible errors given in Proposition IV.1 and Proposition IV.2. In these figures, we fix $H = 5$ and vary a across $0.5, 0.75$, and 0.9.

Multi-step loss: In Figure 4, we compare the multi-step predictor with the single-step predictor trained using a single-step loss, and a multi-step loss (8) with $a = 0.9$ and $H = 10$. We see that in the well-specified setting, the rate of decay for the prediction error of the single-step model trained with a multi-step loss is lower than that of the multi-step model, but both are outperformed by the single-step model with a single-step loss. However, in the presence of misspecification, the prediction error converges nearly to the level of the direct multi-step predictor. The function class for the predictor (8) strictly less expressive than the direct multi-step predictor, (7), which explains why the direct multi-step loss still incurs less bias. For the single-step model with a multi-step loss, $\hat{G}$ is fit using Adam initialized from the single-step predictor fit with a single-step loss, and using a step size $1e-2$.

Control Performance: In Figure 5, we consider a closed-loop control setting in which the control inputs are selected using predictions from either single-step or multi-step models, each trained on datasets of size N. In particular the control input is selected via model predictive control using a horizon $H = 20$ with stage costs $c(y_t, u_t) = \|y_t\|^2 + \|u_t\|^2$ and y_{t+H} constrained to 0. Panel (a) shows the infinite-horizon LQR cost, $\lim_{T\to\infty} \mathbf{E}\left[\sum_{t=1}^{T} y_t^\top y_t + u_t^\top u_t\right]$ incurred by this controller in the well-specified setting for a system with $a = 0.9$. averaged over 1,000 datasets. The multi-step predictor uses the same horizon, $H = 20$, as the MPC horizon. In the low-data regime, single-step rollouts result in lower cost than multi-step prediction.

Panel (b) shows the spectral radius of the resulting closed-loop system in the misspecified setting for the same system ($a = 0.9$, $H = 20$), averaged over 1,000 datasets. A spectral radius greater than 1 for the single-step rollout indicates that the closed-loop system is unstable, and the associated infinite-horizon LQR cost diverges.

VI. Conclusion

In this work, we present a novel theoretical comparison of the asymptotic prediction error associated with autoregressive rollouts of single-step predictors and direct multi-step predictors. Our analysis offers insight into when each modeling approach is preferable. Specifically, we show that for well-specified model classes, autoregressive rollouts of single-step predictors achieve lower asymptotic prediction error. However, in the presence of model misspecification due to an incorrect Markovian assumption, multi-step predictors can significantly outperform their single-step counterparts.

These findings provide a foundation for more informed model design in learning-based control and forecasting. Promising directions for future work include: (1) developing a rigorous analysis of intermediate approaches, such as the single-step model trained with a multi-step loss, which we investigate only empirically in this work; and (2) extending our analysis beyond the white-noise input assumption to study how each of these prediction approaches performs in a closed-loop control setting.

Acknowledgements

This work is supported in part by NSF Award SLES-2331880, NSF CAREER award ECCS-2045834, and AFOSR Award FA9550-24-1-0102.

References

[1] N. Lambert, A. Wilcox, H. Zhang, K. S. Pister, and R. Calandra, "Learning accurate long-term dynamics for model-based reinforcement learning," in *2021 60th IEEE Conference on decision and control (CDC)*. IEEE, 2021, pp. 2880–2887.

[2] N. Lambert, K. Pister, and R. Calandra, "Investigating compounding prediction errors in learned dynamics models," *arXiv preprint arXiv:2203.09637*, 2022.

[3] L. Ljung, "System identification," in *Signal analysis and prediction*. Springer, 1998, pp. 163–173.

Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 13:03:03 UTC from IEEE Xplore. Restrictions apply.

[4] M. Farina and L. Piroddi, "Some convergence properties of multi-step prediction error identification criteria," in *2008 47th IEEE Conference on Decision and Control*. IEEE, 2008, pp. 756–761.

[5] A. Venkatraman, M. Hebert, and J. Bagnell, "Improving multi-step prediction of learned time series models," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 29, no. 1, 2015.

[6] A. Venkatraman, R. Capobianco, L. Pinto, M. Hebert, D. Nardi, and J. A. Bagnell, "Improved learning of dynamics models for control," in *2016 International Symposium on Experimental Robotics*. Springer, 2017, pp. 703–713.

[7] E. Terzi, L. Fagiano, M. Farina, and R. Scattolini, "Learning multi-step prediction models for receding horizon control," in *2018 European Control Conference (ECC)*. IEEE, 2018, pp. 1335–1340.

[8] H. Balim, A. Carron, M. N. Zeilinger, and J. Köhler, "Stochastic data-driven predictive control: Chance-constraint satisfaction with identified multi-step predictors," *IEEE Control Systems Letters*, 2024.

[9] J. Köhler, K. P. Wabersich, J. Berberich, and M. N. Zeilinger, "State space models vs. multi-step predictors in predictive control: Are state space models complicating safe data-driven designs?" in *2022 IEEE 61st conference on decision and control (CDC)*. IEEE, 2022, pp. 491–498.

[10] K. Asadi, D. Misra, S. Kim, and M. L. Littman, "Combating the compounding-error problem with a multi-step model," *arXiv preprint arXiv:1905.13320*, 2019.

[11] R. Chandra, S. Goyal, and R. Gupta, "Evaluation of deep learning models for multi-step ahead time series prediction," *Ieee Access*, vol. 9, pp. 83 105–83 123, 2021.

[12] K. Chua, R. Calandra, R. McAllister, and S. Levine, "Deep reinforcement learning in a handful of trials using probabilistic dynamics models," *Advances in neural information processing systems*, vol. 31, 2018.

[13] S. Levine, A. Kumar, G. Tucker, and J. Fu, "Offline reinforcement learning: Tutorial, review, and perspectives on open problems," *arXiv preprint arXiv:2005.01643*, 2020.

[14] M. Janner, J. Fu, M. Zhang, and S. Levine, "When to trust your model: Model-based policy optimization," *Advances in neural information processing systems*, vol. 32, 2019.

[15] T. Yu, G. Thomas, L. Yu, S. Ermon, J. Y. Zou, S. Levine, C. Finn, and T. Ma, "Mopo: Model-based offline policy optimization," *Advances in Neural Information Processing Systems*, vol. 33, pp. 14 129–14 142, 2020.

[16] T. Yu, A. Kumar, R. Rafailov, A. Rajeswaran, S. Levine, and C. Finn, "Combo: Conservative offline model-based policy optimization," *Advances in neural information processing systems*, vol. 34, pp. 28 954–28 967, 2021.

[17] R. Kidambi, A. Rajeswaran, P. Netrapalli, and T. Joachims, "Morel: Model-based offline reinforcement learning," *Advances in neural information processing systems*, vol. 33, pp. 21 810–21 823, 2020.

[18] C. Chi, Z. Xu, S. Feng, E. Cousineau, Y. Du, B. Burchfiel, R. Tedrake, and S. Song, "Diffusion policy: Visuomotor policy learning via action diffusion," *The International Journal of Robotics Research*, p. 02783649241273668.

[19] A. Block, A. Jadbabaie, D. Pfrommer, M. Simchowitz, and R. Tedrake, "Provable guarantees for generative behavior cloning: Bridging low-level stability and high-level behavior," *Advances in Neural Information Processing Systems*, vol. 36, pp. 48 534–48 547, 2023.

[20] T. Z. Zhao, V. Kumar, S. Levine, and C. Finn, "Learning fine-grained bimanual manipulation with low-cost hardware," *arXiv preprint arXiv:2304.13705*, 2023.

APPENDIX I

Lemma I.1. *Assume* $\rho(A) < 1$. *Then,*

$$\rho(CA\Sigma_x C^\top \Sigma_y^{-1}) \leq 1.$$

Proof of Lemma I.1. Note that $CA\Sigma_X C^\top \Sigma_y^{-1} = \bar{\mathbf{E}}[y_{t+1}y_t^\top]\bar{\mathbf{E}}[y_t y_t^\top]^{-1}$. From the stationarity of the process,

$$\bar{\mathbf{E}}\begin{bmatrix} y_t \\ y_{t+1} \end{bmatrix}\begin{bmatrix} y_t \\ y_{t+1} \end{bmatrix}^\top = \begin{bmatrix} \Sigma_y & \Sigma_{y+} \\ \Sigma_{y+} & \Sigma_y \end{bmatrix}.$$

By a Schur complement, $\Sigma_y - \Sigma_{y+}\Sigma_y^{-1}\Sigma_{y+} \succeq 0$, or $(\Sigma_y^{-1/2}\Sigma_{y+}\Sigma_Y^{-1/2})^2 \preceq I$. Then $\left\|\Sigma_y^{-1/2}\Sigma_{y+}\Sigma_Y^{-1/2}\right\| \leq 1$. For $i = 1, \ldots, d_{\mathrm{y}}$, it holds that $\left|\lambda_i(\Sigma_{y+}\Sigma_Y^{-1})\right| = \left|\lambda_i(\Sigma_y^{-1/2}\Sigma_{y+}\Sigma_Y^{-1/2})\right| \leq 1$, and thus $\rho(\Sigma_{y+}\Sigma_y^{-1}) \leq 1$. □

For the proofs of the four main propositions in the main paper, we use the following facts. It holds by the Birkoff-Khinchin theorem that $\lim_{N\to\infty} \frac{1}{N}\sum_{t=1}^N z_t z_t^\top \to \Sigma_z$. By Slutsky's theorem, if $\lim_N X_N \overset{d}{=} X$ and $\lim_N Y_N = c$ for X a random variable and c a constant, then $\lim_N X_n Y_n \overset{d}{=} Xc$. It holds by the dominated convergence theorem (DCT) that $\lim_N \mathbf{E}\left\|(\hat{G}_N - G^\star)\Sigma_z^{1/2}\right\|_F^2 = \mathbf{E}\left[\lim_N \left\|(\hat{G}_N - G^\star)\Sigma_z^{1/2}\right\|_F^2\right]$.

The calculations involved to determine the expressions in the main propositions are deferred to the extended manuscript.

Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 13:03:03 UTC from IEEE Xplore. Restrictions apply.