
Markovian Experimental Design under Concept Drift

Ayberk Yarkın Yıldız¹ Lili Su¹ Carlee Joe-Wong² Edmund Yeh¹ Stratis Ioannidis¹

¹Northeastern University, Boston, MA ²Carnegie Mellon University, Pittsburgh, PA

{yildiz.ay, l.su, e.yeh, ioannidis}@northeastern.edu
cjoewong@andrew.cmu.edu

Abstract

We study how to optimally select training samples in the presence of concept drift. We propose a *Markovian experimental design* framework, in which a learner sequentially collects samples and re-estimates a model through Bayesian linear regression. We show that, when model drift is mean-reverting, i.e., governed by an unknown mean fixed-point coupled with Gaussian deviations, model posterior estimates are determined by a Kalman filter and the system is asymptotically stable. We also show that the steady-state model posterior can be computed as the solution of the so-called Discrete Algebraic Riccati Equation (DARE), and that D-optimal steady-state and myopic sample selection policies can be computed by solving tractable convex optimization problems. We do so by characterizing the curvature of the DARE solution, even though the latter generally has no closed form. Finally, we show that our approach generalizes to the non-linear model setting by replacing the Kalman filter with an extended Kalman filter.

1 Introduction

In many practical machine learning (ML) settings, algorithms trained over offline data and deployed in realistic conditions are subject to *concept drift* [45, 5]: the model, i.e., the relationship between output and input variables, changes as time progresses. Applications in which concept drift is observed, as well as causes for such drifts, are numerous: in recommender systems, concept drift arises due to changes in content popularity and temporal trends [11, 44], in wireless communications, due to slowly varying channel conditions [1, 57, 60], in environmental monitoring, due to seasonal and meteorological cycles [14], in financial forecasting, due to shifting market and volatility regimes [51], and so forth. In practice, concept drift is often addressed by continuously retraining models [33, 10]. The sheer volume of data streams available necessitates retraining over a subset of all data available at a given time. Privacy regulations and consequent data retention practices [15, 67] dictate that data is erased after some time, restricting retraining only to recently collected data. This gives rise to the following question: which subset of the data should one use to improve the quality of a predictive model, in light of (a) the presence of concept drift and (b) the data available at retraining time?

We propose a *Markovian experimental design* framework that allows us to study this question. In our proposed framework, a learner keeps track of a model drifting according to linear dynamics coupled with Gaussian noise. The learner sequentially selects samples to collect and subsequently re-estimates the model, e.g., via maximum a posteriori estimation over the new data. Our proposed mathematical framework allows us to address the data selection question in a precise and rigorous fashion; in particular, we make the following contributions:

- We model the drift as a specific, mean-reverting Gaussian process: in our drift dynamics, model evolution is characterized by a mean fixed-point coupled with Gaussian deviations. Concretely, the drifting model $\beta_k \in \mathbb{R}^d$ evolves as

$$\beta_{k+1} = \rho \beta_k + (1 - \rho) \beta_* + \mathbf{w}_k, \quad \text{where } \rho \in (0, 1), \quad \mathbf{w}_k \sim \mathcal{N}(\mathbf{0}, \sigma_{\mathbf{w}}^2 \mathbf{I}) \quad (1)$$

i.e., oscillating around the (unknown) mean fixed-point β_* . These dynamics are simple but also relevant from a practical standpoint: in a real-life scenario, the mean fixed-point captures “on-average” or “status-quo” behavior, as it corresponds to the model learned if one ignored the presence of drift. In turn, Gaussian deviations correspond to temporal concept drifts, that need to be accounted for.

- We consider first a Bayesian linear regression setting. We show that, under these model dynamics, (a) model posterior estimates are governed by a Kalman filter, (b) the mean fixed-point can be estimated arbitrarily well, and (c) the system is asymptotically stable under stationary selection policies. In particular, the system converges to a steady-state model posterior, whose covariance can be computed as the solution of the so-called Discrete Algebraic Riccati Equation (DARE) [4]. We also prove that D-optimal sample selection policies, both myopic/one-step lookahead as well as in steady-state, can be computed efficiently by solving appropriately-defined convex optimization problems.
- Finally, we extend our analysis to the non-linear case through the use of an Extended-Kalman Filter (EKF) [8], which linearizes the model function.

From a technical standpoint, even though the DARE for arbitrary drift dynamics does not generally have a closed-form solution, we construct an explicit characterization of the spectrum of its solution under our dynamics. We accomplish this by coupling the eigen-decomposition of covariance limit under our proposed mean fixed-point dynamics with the eigen-decomposition of the design matrix. This paves the way to optimizing the stationary policy in steady state: armed with this characterization, we prove that the D-Optimal design objective is Schur-concave w.r.t. this spectrum, which in turn implies concavity over the positive semi-definite cone of covariance matrices. Our mean fixed-point dynamics are also crucial to computing a steady-state policy in the EKF setting, as they enable us to compute the model Jacobian around the mean fixed-point limit.

2 Related Work

Experimental design [12, 9, 55, 28], i.e., the process of determining which samples to collect/experiments to conduct prior to model inference, has a long history in both machine learning and statistical inference. In the context of linear regression, the so-called “alphabet” criteria [55, 12] aim at minimizing different scalarizations of the model covariance; Bayesian versions of the latter also exist [12], and the D-optimality criterion we consider here falls under this category. The interest in sequential experimental design procedures, in which experiments are conducted sequentially, can be traced back to seminal works by Wald [68] and Robbins [58]. Bayesian assumptions naturally fit this setting, and many classic sequential experimental design works follow such an approach [12, 18, 48, 3, 17]. Significant computational challenges arise when estimating high-dimensional, non-linear models [23, 29, 30]. Experimental design has also been used in the estimation of model parameters of a dynamical system; see, e.g., [64, 65, 46]. Studied problems and approaches are orthogonal to the ones we consider here, namely, selecting which samples to collect; conversely, they would be useful for estimating parameters governing drift dynamics in our setting (e.g., matrix $\mathbf{A}$ in App. E). Nevertheless, both this and the classic regime assume that model parameters to be estimated are fixed, and not subject to concept drift across stages; as a result, our consideration of drift constitutes a significant departure from prior experimental design literature.

The prevalence of concept drifts in real-life settings has given rise to online learning algorithms designed to adapt a model in the presence of drift [45, 71]. This is also a rich research area that, not surprisingly, again traces its roots to Robbins & Monro [59] and includes classic stochastic approximation algorithms [7, 39], window-based and forgetting-factor approaches [72, 36, 43], and ensemble methods [35, 70]. Beyond model adaptation methods, techniques combining drift detection [5, 16] with online learning have been proposed [69, 42, 38, 6]. These methods primarily focus on adapting model parameters to track evolving data distributions. Though studying a rather simple setting (primarily linear regression) compared to literature on concept drift, our work departs by focusing on the question of how to combine online learning (i.e., model adaptation) with experimental design (i.e., the choice of experiments to conduct/samples to collect).

A more closely related thread considers experimental design for Hidden Markov Models (HMMs), where the model parameters (the hidden state) take one out of finitely many possible values, and transitions are Markovian. Several experimental design techniques for HMMs use objectives akin to

ours (e.g., expected information gain [54] or entropy reduction in latent variable models [31]) and study myopic, one-step-ahead policies [54, 31, 56, 66]. Alali & Imani [2] study both steady state and infinite-horizon policies using deep RL, to approximate a Bayesian lookahead policy for data collection in HMMs. Our approach differs from this line of work by departing from the finite hidden state assumption, considering a dynamical system evolution of drift instead. This leads to an infinite state space, but is also more realistic when model changes are smooth.

Another related thread studies sensor scheduling and selection for Kalman filtering, where one selects which sensors to activate so as to minimize the steady-state estimation error. Dutta et al. [19] give a unified full-horizon formulation solved by mixed-integer programming, without optimality guarantees; Ye et al. [73] study steady-state sensor selection under the trace (A-optimality) criterion, prove it is NP-hard and inapproximable, and resort to greedy heuristics. Our setting differs in two respects. First, our “state” is a drifting learned model with an unknown fixed point β_* estimated online, rather than a known-dynamics physical state. Second, we optimize a continuous, fractional design over the simplex rather than a combinatorial subset of sensors; together with our $\rho\mathbf{I}$ plus isotropic-noise structure, this yields a closed-form characterization of the DARE spectrum and a concave steady-state D-optimal objective. As a result, our design is a poly-time solvable convex program, circumventing the above hardness results. We include the greedy Kalman information-gain heuristic [61] as a competitor in our experiments.

To the best of our knowledge, the work closest to ours is a recent paper by Si Salem [62], which considers which samples to retain in an online learning setting. Value of data is captured by a transductive loss which relates to the expected prediction error (Eq. (33)) in our paper. Si Salem [62] provides a data retention algorithm with bounded regret guarantees in this setting. We depart by modeling drift in a Markovian fashion; this implies that estimation via a Kalman filter is optimal and, therefore, no data needs to be retained past the current stage. This also allows us to identify optimal (as opposed to bounded-regret) myopic and steady-state decisions in our setting.

3 Markovian Experimental Design under Concept Drift

3.1 Problem Setting

We consider a dynamic learning setting in which a learner tries to learn a model subject to *concept drift*. The learning process proceeds in stages and, in each stage, the learner chooses a set of experiments to conduct and observes labels governed by the model active at this stage. The learner updates its model estimates using these collected labels and proceeds to the next stage, whereupon the model drifts.

Dynamic Label Collection. Given a candidate set of feature vectors $\mathcal{X}_k \subset \mathbb{R}^d$ with $|\mathcal{X}_k| = p_k \in \mathbb{N}$, the learner selects experiments to conduct at stage k by creating a vector $\mathbf{n}_k = [n_{ik}]_{i=1}^{p_k} \in \mathbb{N}^{p_k}$ with $\mathbf{n}_k^\top \mathbf{1} = n_k$. Under selection vector $\mathbf{n}_k$, the experiments to be conducted correspond to the design matrix $\mathbf{X}_k \equiv \mathbf{X}(\mathbf{n}_k) \in \mathbb{R}^{n_k \times d}$, where the map $\mathbf{X} : \mathbb{N}^{p_k} \rightarrow \mathbb{R}^{n_k \times d}$ stacks n_{ik} copies of vector $\mathbf{x}_i \in \mathcal{X}_k$, for $i = 1, \dots, p_k$. The learner collects:

$$\mathbf{y}_k = h(\mathbf{X}_k, \beta_k) + \epsilon_k, \quad \text{where } \epsilon_k \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}), \quad (2)$$

and $h : \mathbb{R}^{n_k \times d} \times \mathbb{R}^{d_\beta} \rightarrow \mathbb{R}^{n_k}$ is a function parameterized by $\beta_k \in \mathbb{R}^{d_\beta}$ (a.k.a. the *model*) at stage k . Our goal is to determine how the learner should (a) choose its *design selection policy* $\mathbf{n}_k \in \mathbb{N}^{p_k}$ and (b) *estimate* $\beta_k \in \mathbb{R}^{d_\beta}$ at each stage k , given that the model is subject to drift. We discuss drift next.

Concept Drift. A simple way to model drift is via: $\beta_{k+1} = \beta_k + \mathbf{w}_k$, where $\mathbf{w}_k \sim \mathcal{N}(\mathbf{0}, \sigma_w^2 \mathbf{I})$, with the first model sampled from a known prior, i.e., $\beta_1 \sim \mathcal{N}(\beta_0, \Sigma_0)$. Model parameters are thus subject to Gaussian perturbations of variance σ_w^2 , capturing temporally-localized trends. A problem with this simple model is that model variance drifts to infinity: this is not realistic in standard training tasks, where both labels and features stay bounded. We thus assume the following dynamics:

$$\beta_{k+1} = \rho \beta_k + (1 - \rho) \beta_* + \mathbf{w}_k, \quad (3)$$

where $\rho \in (0, 1)$ is a positive constant, $\mathbf{w}_k \sim \mathcal{N}(\mathbf{0}, \sigma_w^2 \mathbf{I})$ are i.i.d. perturbations, and $\beta_* \in \mathbb{R}^d$ is an (unknown) mean fixed-point.¹ Intuitively, the dynamics in Eq. (3) capture a model that oscillates

¹It is easy to see that $\lim_{k \rightarrow \infty} \mathbb{E}[\beta_k] = \beta_*$ under Eq. (3). Note that, in contrast, $\lim_{k \rightarrow \infty} \beta_k$ *does not* converge in the classic sense: it does, however, converge *in distribution* to a Gaussian r.v. centered at β_* .

around β_* . The latter corresponds to mean behavior, or the “status-quo”: it is the model the learner discovers when ignoring temporal drifts altogether (see also Thm. 2 below).

Linear Setting. Our main focus is the linear setting, where $d_\beta = d$, $h(\mathbf{X}_k, \beta_k) \equiv \mathbf{X}_k \beta_k$, and:

$$\mathbf{y}_k = \mathbf{X}_k \beta_k + \epsilon_k, \quad \text{where } \epsilon_k \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}). \quad (4)$$

We show that, in the linear setting, Maximum a Posteriori (MAP) model estimation corresponds to a Kalman filter, and subsequently select designs using a classic experimental design criterion. For completeness, we review Kalman filters in App. A and experimental design in App. B.

3.2 Model Estimation in the Linear Setting

Before we optimize design selection, we first discuss how to estimate β_k under arbitrary design selection policies. We will subsequently reason about the quality of such estimates to build optimal design selection policies. As a warm-up (but also instructive) step, we momentarily consider estimation when the mean fixed point β_* in Eq. (3) is a-priori known.

Known β_* . We use the common Kalman filter notation $(\cdot)_{k|k'}$ to denote a quantity estimated at stage k , having access to samples up to stage $k' \leq k$. Then, the Bayesian estimate of β_k is as follows:

Theorem 1. *Consider the system evolving via Eqs. (3)–(4). Assume that the mean fixed-point β_* is known, and that design matrices $\mathbf{X}_k \in \mathbb{R}^{n_k \times d}$ are selected by an arbitrary (possibly non-stationary) policy. Then, at stage k , the posterior distribution of β_k after sample collection is Gaussian with mean and covariance given by:*

$$\hat{\beta}_{k|k} = \hat{\beta}_{k|k-1} + \mathbf{K}_k \left(\mathbf{y}_k - \mathbf{X}_k \hat{\beta}_{k|k-1} \right), \quad \text{and} \quad \Sigma_{\xi_{k|k}} = (\mathbf{I} - \mathbf{K}_k \mathbf{X}_k) \Sigma_{\xi_{k|k-1}}, \quad (5)$$

respectively, where $\hat{\beta}_{k|k-1}$, $\Sigma_{\xi_{k|k-1}}$ are the a priori estimates before label collection:

$$\hat{\beta}_{k|k-1} = \rho \hat{\beta}_{k-1|k-1} + (1 - \rho) \beta_*, \quad \text{and} \quad \Sigma_{\xi_{k|k-1}} = \rho^2 \Sigma_{\xi_{k-1|k-1}} + \sigma_w^2 \mathbf{I}, \quad (6)$$

and $\mathbf{K}_k$ is the so-called Kalman gain:

$$\mathbf{K}_k = \Sigma_{\xi_{k|k-1}} \mathbf{X}_k^\top \left(\mathbf{X}_k \Sigma_{\xi_{k|k-1}} \mathbf{X}_k^\top + \sigma^2 \mathbf{I} \right)^{-1}. \quad (7)$$

We prove this in App. C. Note that the presence of the β_* term in Eq. (3) implies that our setting slightly deviates from the classic Kalman filter setting (see App. A). We nevertheless reduce estimation to the classic setting through an appropriate transformation (namely, by subtracting β_*).

Estimating β_* . Next, we turn our attention to the case where β_* is not known. We show that β_* can be estimated arbitrarily well by averaging per-stage linear regression estimates:

Theorem 2. *Consider the system evolving via Eqs. (3)–(4). Assume that the eigenvalues of matrices $\{\mathbf{X}_k^\top \mathbf{X}_k\}_{k=1}^\infty$ are bounded from above as well as bounded away from zero. Then:*

$$\hat{\beta}_* = \frac{1}{k} \sum_{k'=1}^k \hat{\beta}_{k'}^{\text{LR}}, \quad \text{for } \hat{\beta}_{k'}^{\text{LR}} \equiv \left(\mathbf{X}_{k'}^\top \mathbf{X}_{k'} \right)^{-1} \mathbf{X}_{k'}^\top \mathbf{y}_{k'} \quad (8)$$

converges both w.p. 1 and in expectation to the mean fixed-point β_* as $k \rightarrow \infty$.

The proof is in App. D. Combining Thms. 1 and 2, we can estimate the model at stage k using Eqs. (5)–(8), replacing β_* with $\hat{\beta}_*$. Note that, in Thm. 1, β_* is used to obtain the posterior mean/MAP estimate $\hat{\beta}_{k|k}$ but not the posterior covariance $\Sigma_{\xi_{k|k}}$. Hence, the resulting plug-in estimate $\hat{\beta}_{k|k}$ is asymptotically unbiased and the covariance is correctly estimated for an arbitrary design selection policy, so long as design matrices are well-conditioned.² This also implies that we can design experiments that minimize the model covariance without exact knowledge of β_* . We exploit this in identifying optimal myopic and steady-state design selection policies next.

²It also suffices that they are well-conditioned infinitely often, with β_* estimated only when they do. Note that this well-conditioning assumption holds throughout our experiments; nevertheless, it can also be enforced without affecting convexity guarantees by adding a small exploration component to the design (see App. E).

3.3 Optimal Design Selection in the Linear Setting

We next turn our attention to the question of how the learner is to perform the experimental design selection (i.e., how the learner picks $\mathbf{n}_k$ at every stage). To that end, we use the so-called *D-Optimality criterion* from classic experimental design (see App. B). We describe how to design D-optimal design selection policies in both the *myopic* and *steady-state* case.

3.3.1 Optimal Myopic Design Selection

We first focus on a *myopic* policy: the learner optimizes the design selection at stage k in a *stage-wise* (i.e., one stage at a time) fashion, having access to the past history at stage $k - 1$. In D-optimal design [55, 9], the learner would select experiments that minimize the log det of the model covariance $\Sigma_{\xi_{k|k-1}}$. Intuitively, this minimizes the entropy of the (normal) distribution of β_k ; it acts as a proxy of the expected prediction error over an unseen sample (see App. B). As in classic experimental design [55, 9], we assume that n_k is large enough so that we can describe the selection through a fractional policy: our first result is that a D-optimal fractional stage-wise design policy can be obtained by solving a convex optimization problem.

Theorem 3. *Consider the system evolving via Eqs. (3)–(4). The myopic D-optimal fractional policy at stage k can be computed as a solution to the following relaxed optimization problem:*

$$\text{Maximize: } G\left(\pi_k; \sigma^2, \Sigma_{\xi_{k|k-1}}\right) \equiv \log \det \left[\frac{n_k}{\sigma^2} \mathbf{B}(\pi_k) + \Sigma_{\xi_{k|k-1}}^{-1} \right], \quad (9a)$$

$$\text{s.t.: } \pi_k^\top \mathbf{1} \leq 1, \pi_k \geq \mathbf{0}, \quad (9b)$$

where $\pi_k \in [0, 1]^{p_k}$ is the fractional policy at stage $k \in \mathbb{N}$, $\mathbf{B}(\pi) = \sum_{i=1}^{p_k} \pi_i \mathbf{x}_i \mathbf{x}_i^\top$, and $\Sigma_{\xi_{k|k-1}}^{-1} \in \mathbb{R}^{d \times d}$ is the inverse of the a priori covariance given by Eq. (6).

The proof is in App. E; in fact, we show that this even extends to more general model drift dynamics than Eq. (3). The convexity of Prob. (9) implies that it can be solved with standard convex optimization techniques [9]. Note that the monotonicity of G implies that the upper inequality constraint binds. Moreover, the fractional solution can be rounded producing an approximate integral solution s.t. $\mathbf{n}_k \approx n_k \cdot \pi_k$ via standard techniques; we review rounding schemes in App. B.

3.3.2 Optimal Steady-State Selection

Next, we turn our attention to stationary policies. To that end, we assume hereon that $\mathcal{X}_k = \mathcal{X}$ and $p_k = p$ (i.e., experiment options are fixed), and $\mathbf{n}_k = \mathbf{n} \in \mathbb{N}^p$, with $\mathbf{n}^\top \mathbf{1} = n$. We seek stationary policies $\mathbf{n} \in \mathbb{N}^p$ that are optimal in *steady-state*. We first establish that a steady-state indeed exists.

Convergence to Steady State. Under a stationary policy, the system determined by dynamics (3)–(4) is asymptotically stable:

Theorem 4. *Consider the system evolving via Eqs. (3)–(4) under a stationary policy $\mathbf{n} \in \mathbb{N}^p$, and assume that $\mathbf{X}(\mathbf{n}) \in \mathbb{R}^{n \times d}$ is rank d . Then, as $k \rightarrow \infty$, the error covariance $\Sigma_{\xi_{k|k}}$ converges to $\bar{\Sigma}_\xi \in \mathbb{R}^{d \times d}$ that is a solution to the Discrete Algebraic Riccati Equation (DARE):*

$$\bar{\Sigma}_\xi^{-1} = \frac{1}{\sigma^2} \mathbf{X}(\mathbf{n})^\top \mathbf{X}(\mathbf{n}) + (\rho^2 \bar{\Sigma}_\xi + \sigma_w^2 \mathbf{I})^{-1}, \quad (10)$$

while the Kalman gain $\mathbf{K}_k$ converges to:

$$\bar{\mathbf{K}} = \bar{\Sigma}_\xi \mathbf{X}(\mathbf{n})^\top (\mathbf{X}(\mathbf{n}) \bar{\Sigma}_\xi \mathbf{X}(\mathbf{n})^\top + \sigma^2 \mathbf{I})^{-1}. \quad (11)$$

The proof is in App. F. Together, Thms. 2 and 4 imply the following. First, given a stationary policy $\mathbf{n} \in \mathbb{N}^p$, the posterior covariance $\Sigma_{\xi_{k|k}}$ converges to the solution of the DARE (10) *irrespective of whether β_* is known or not*. This is because neither $\Sigma_{\xi_{k|k}}$ nor the DARE (10) depend on β_* . Second, this convergence $\Sigma_{\xi_{k|k}} \rightarrow \bar{\Sigma}_\xi$ is *decoupled from the convergence $\hat{\beta}_* \rightarrow \beta_*$ in Thm. 2*: the covariance converges irrespective of whether $\hat{\beta}_*$ converges (or is even used in the estimation of β_k), while $\hat{\beta}_*$ converges w.p. 1, for any stationary selection policy, as long as $\text{rank}(\mathbf{X}(\mathbf{n})) = d$.

Optimal Stationary Policy. Armed with Thms. 1–4, we seek a stationary policy that minimizes the steady-state covariance, as given by DARE (10). We stress again that covariance dynamics are identical regardless of whether the mean fixed-point β_* is known or estimated via Thm. 2. Formally, we wish to select $\pi \in [0, 1]^p$ that maximizes $\log \det [\bar{\Sigma}_\xi^{-1}(\pi)]$, where, by Thm. 4, $\bar{\Sigma}_\xi(\pi)$ is the solution to the (relaxed) DARE:

$$\bar{\Sigma}_\xi^{-1} = \frac{n}{\sigma^2} B(\pi) + (\rho^2 \bar{\Sigma}_\xi + \sigma_w^2 \mathbf{I})^{-1}, \quad (12)$$

with $B(\pi) \equiv \sum_{i=1}^p \pi_i x_i x_i^\top$. Hence, finding the optimal fractional policy in steady state amounts to solving the following maximization problem:

$$\text{Max.}: G(\pi; \sigma^2, \sigma_w^2) \equiv \log \det [\bar{\Sigma}_\xi^{-1}(\pi)], \quad (13a)$$

$$\text{s.t.}: \pi^\top \mathbf{1} \leq 1, \pi \geq \mathbf{0}. \quad (13b)$$

Note that the map $\pi \mapsto \bar{\Sigma}_\xi(\pi)$ is defined *implicitly*, through the non-linear system of equations in Eq. (12), rather than through an explicit formula. To design a stationary policy that is optimal in steady state, we first try to relate how the steady-state covariance $\bar{\Sigma}_\xi$ depends on the stationary fractional policy $\pi \in [0, 1]^p$. To that end, we prove an auxiliary lemma, establishing that matrices $\bar{\Sigma}_\xi$ and $B(\pi)$ share the same eigenvectors.

Lemma 1. Consider $\Sigma \in \mathbb{R}^{d \times d}$ that is a solution to

$$\Sigma^{-1} = \frac{n}{\sigma^2} B + (\rho^2 \Sigma + \sigma_w^2 \mathbf{I})^{-1}, \quad (14)$$

where $B \in \mathbb{S}_+^d$ is a PSD matrix. For $i = 1, \dots, d$, let λ_i, e_i , be the eigenvalues and eigenvectors of B , respectively. Then, $\Sigma = \sum_{i=1}^d \lambda'_i e_i e_i^\top$, where, for $i = 1, \dots, d$,

$$\lambda'_i = 2\sigma_w^2 \frac{\sigma^2}{n} \cdot \left[\sqrt{\left(\frac{\sigma^2}{n} (1 - \rho^2) + \sigma_w^2 \lambda_i \right)^2 + 4\lambda_i \rho^2 \sigma_w^2 \frac{\sigma^2}{n}} + \left(\frac{\sigma^2}{n} (1 - \rho^2) + \sigma_w^2 \lambda_i \right) \right]^{-1}. \quad (15)$$

We stress that Lemma 1, which we prove in App. G, *does not generally hold for DARE solutions of arbitrary drift dynamics (see Eq. (29))*. One of our main technical contributions is thus to observe that, under dynamics (3), we can solve the DARE *explicitly*.

Lemma 1 also paves the way for determining the curvature of the D-Optimal design objective in steady state. In particular, leveraging this explicit construction, we show that *Prob. (13) is a convex optimization problem*:

Theorem 5. The objective function $G : \mathbb{R}^p \rightarrow \mathbb{R}$ defined in Eq. (13a) is non-decreasing and concave.

The proof is in App. H. In light of the intricate relationship between π and $\bar{\Sigma}_\xi$ through Lemma 1, the proof is technically challenging; we provide a proof outline in the beginning of App. H. In short, we prove that the D-Optimal design objective is non-decreasing and Schur-concave w.r.t. the eigenvalues of $B(\pi)$, which in turn implies concavity over the positive semi-definite cone of covariance matrices. The same machinery extends to the A-optimality (trace) criterion, which is likewise convex (App. I). An immediate consequence is that the optimal steady-state policy π can be computed efficiently via convex optimization techniques. A key challenge lies in computing the gradient ∇G , as this requires differentiating the steady-state covariance defined implicitly through the DARE systems of equations (Eq. (12)). In App. J, we show how the gradient can be computed via the implicit function theorem [37], and solve Prob. (13) efficiently using the Frank–Wolfe algorithm [24].

3.4 Non-linear Setting

Our approach generalizes to the case when h is non-linear using the so-called *Extended Kalman filter (EKF)* [8], which linearizes h around the current model estimate at each stage. Formally, the EKF linearizes h through its first-order Taylor approximation, so that

$$h(\mathbf{X}_k, \beta_k) \approx h(\mathbf{X}_k, \hat{\beta}_{k|k-1}) + \mathbf{H}_k(\beta_k - \hat{\beta}_{k|k-1}), \quad (16)$$

Table 1: Dataset features and model parameter dimensions. The dataset-feature block (p : candidate experiments per stage; n : training-set size per stage; d : feature dimension; T : number of stages) is uniform across methods. The non-linear d_β block reports the model parameter dimension per architecture (ResNet/CNN/MLP); in the linear setting $d_\beta = d$.

Dataset	Dataset Features				Non-linear d_β		
	p	n	d	T	ResNet	CNN	MLP
Synthetic	40	50	30	2000	97	98	97
Beijing PM2.5 [14]	19	30	7	1826	39	33	31
S&P 500 [51]	16	25	6	62	39	33	29
Bike Sharing [21]	19	30	8	724	39	33	33
California Housing [52]	40	55	6	412	39	33	29
AT&T Downlink [13]	16	25	12	335	39	33	41
AT&T Uplink [13]	16	25	8	328	39	33	33

where $\mathbf{H}_k = \nabla_{\beta} h(\mathbf{X}_k, \hat{\beta}_{k|k-1}) \in \mathbb{R}^{n_k \times d}$ is the Jacobian of h . In practice, this leads to the EKF adaptation equations that, compared to a standard Kalman filter, replace the design matrix $\mathbf{X}_k$ with the Jacobian $\mathbf{H}_k$. Our approach naturally extends to this setting, as discussed in App. K. A caveat is that, in the non-linear case, least-squares estimation is non-convex, so the mean fixed-point β_* may not be identifiable. Nevertheless, we can still produce a “status-quo” estimate, by minimizing the square-error loss ignoring dynamics. Just as in the linear case, this estimate does not play a role in the design of myopic policies. In contrast to the linear case, however, it does play a role in the design of a steady-state policy as, to compute the steady-state Jacobian, we exploit the fact that the model is centered around β_* under the dynamics in Eq. (3). As we discuss in the next section, both of these approximations (the linearization via Eq. (16) and our β_* estimate) work very well in practice.

4 Numerical Evaluation

4.1 Experimental Setup

Our experimental setup, including methods, datasets, and metrics, is described in detail in App. L. We briefly review it here. Our source code is publicly available³.

Datasets. We use a synthetic dataset and six publicly available datasets, summarized in Table 1. We generate four versions of the synthetic dataset, all consistent with the dynamics in Eqs. (2)–(3), each using different model functions h : a linear model, ResNet [27], a convolutional neural network (CNN), and a Multi-Layer Perceptron (MLP). The Beijing PM2.5 dataset [14] contains 5 years of hourly air quality measurements, where pollution dynamics are driven by seasonal weather patterns and policy interventions. The S&P dataset [51] contains daily stock prices for S&P 500 companies, which drift due to changing market regimes and volatility cycles. The Bike Sharing dataset [21] contains hourly rental counts from Washington, D.C., over 2011–2012; demand patterns drift with seasonal weather and growing ridership adoption. The California Housing dataset [52] contains aggregated 1990 US Census statistics in California; we induce drift by assigning different geographic regions to different stages. The AT&T Downlink and AT&T Uplink datasets contain a multi-carrier cellular study [13], collected via a cross-country drive (from Boston, MA to Los Angeles, CA) measurement campaign, traversing approximately 4,500 km across 13 US states.

Methods. We couple different *design selection* methods with different *model estimation* methods. For design selection, we consider five approaches: **Stage-wise Optimal Selection (STG)**, our optimal myopic design, described in Prob. (9); **Stationary Optimal Selection (STS)**, our optimal steady-state design policies for linear and non-linear settings, described in Prob. (13); **Greedy Kalman Information-Gain Selection (GKS)** [61], an external greedy sensor-selection baseline that sequentially selects the candidates of highest one-step information gain; **Adam Selection (ADM)**, which performs projected gradient ascent to solve the Prob. (9) at each stage using the Adam optimizer [34]; and **Random Selection (RND)**, which samples feature vectors u.a.r. For

³<https://github.com/neu-spiral/MarkovExp>

Table 2: Test RMSE on synthetic and real datasets for Linear, ResNet, CNN, and MLP models. **Bold** indicates the best result per model and underline the second-best. Mean $\pm$ std reported over 100 trials with different random seeds.

Model	Method	Beijing PM2.5	S&P 500	Bike Sharing	Cal. Housing	AT&T Downlink	AT&T Uplink	Synthetic
Linear	RND-SGD	148.42 ± 1.20	2.95 ± 0.75	280.78 ± 2.18	4.11 ± 0.04	158.66 ± 1.85	25.41 ± 1.25	7.75 ± 0.09
	RND-SE	144.80 ± 0.87	2.77 ± 0.72	273.01 ± 0.57	3.81 ± 0.03	152.38 ± 1.27	23.07 ± 1.15	4.22 ± 0.18
	RND-AE	142.53 ± 0.50	2.71 ± 0.73	267.70 ± 1.00	3.72 ± 0.01	146.30 ± 1.69	20.77 ± 0.95	3.20 ± 0.05
	RND-RLS	142.48 ± 0.55	2.68 ± 0.71	265.57 ± 1.04	3.76 ± 0.03	147.43 ± 1.71	21.55 ± 0.92	5.48 ± 0.12
	ADM-SGD	147.91 ± 1.15	2.58 ± 0.68	280.63 ± 2.16	3.94 ± 0.05	153.44 ± 1.75	18.48 ± 1.35	7.57 ± 0.12
	ADM-SE	143.52 ± 0.76	2.43 ± 0.59	272.93 ± 0.42	3.76 ± 0.01	149.35 ± 1.25	13.12 ± 1.15	2.78 ± 0.10
	ADM-AE	141.46 ± 0.70	1.79 ± 0.57	265.40 ± 0.31	3.70 ± 0.02	142.34 ± 1.09	12.17 ± 0.58	2.74 ± 0.08
	ADM-RLS	142.15 ± 0.61	1.84 ± 0.62	264.42 ± 0.38	3.71 ± 0.02	145.31 ± 1.42	12.58 ± 0.74	4.92 ± 0.09
	GKS-SGD	146.05 ± 1.01	2.74 ± 0.71	278.97 ± 2.00	3.88 ± 0.06	156.18 ± 1.81	19.52 ± 1.31	6.98 ± 0.08
	GKS-SE	144.55 ± 0.69	2.61 ± 0.65	271.34 ± 0.48	3.78 ± 0.03	151.83 ± 1.32	17.63 ± 1.17	2.92 ± 0.10
	GKS-AE	141.65 ± 0.52	1.88 ± 0.55	263.01 ± 0.41	3.71 ± 0.02	143.51 ± 1.14	15.67 ± 0.93	2.79 ± 0.09
	GKS-RLS	141.13 ± 0.48	1.81 ± 0.51	264.44 ± 0.35	3.72 ± 0.02	145.57 ± 1.27	16.10 ± 1.08	7.15 ± 0.14
	STG-SGD	143.94 ± 1.01	2.51 ± 0.59	278.93 ± 2.00	3.89 ± 0.03	150.84 ± 1.03	17.85 ± 0.93	7.53 ± 0.11
	STG-SE	143.57 ± 0.69	2.55 ± 0.63	270.57 ± 0.47	3.88 ± 0.02	140.44 ± 1.30	13.18 ± 0.97	2.76 ± 0.07
	STG-RLS	140.53 ± 0.57	1.79 ± 0.55	260.73 ± 0.33	3.70 ± 0.02	139.14 ± 0.98	13.05 ± 0.84	5.59 ± 0.17
	STG-AE (ours)	138.61 ± 0.60	1.68 ± 0.51	252.97 ± 0.39	3.66 ± 0.01	138.97 ± 0.85	11.10 ± 0.81	2.66 ± 0.08
STS-AE (ours)	133.82 ± 0.63	1.75 ± 0.53	244.18 ± 0.31	3.68 ± 0.01	132.78 ± 0.90	10.51 ± 0.78	2.56 ± 0.06	
ResNet	RND-SGD	104.97 ± 1.23	1.78 ± 0.66	187.55 ± 1.11	2.51 ± 0.94	131.90 ± 0.93	13.46 ± 1.23	1.41 ± 0.09
	RND-SE	100.99 ± 3.00	0.80 ± 0.23	180.10 ± 2.45	1.85 ± 0.70	130.51 ± 0.98	8.30 ± 1.93	2.59 ± 0.50
	RND-AE	99.24 ± 0.39	0.76 ± 0.03	180.01 ± 0.35	1.32 ± 0.15	131.51 ± 0.33	7.66 ± 0.27	1.29 ± 0.10
	RND-RLS	99.89 ± 0.78	0.77 ± 0.05	182.88 ± 0.51	1.48 ± 0.23	129.76 ± 1.12	7.90 ± 0.39	2.18 ± 0.31
	ADM-SGD	107.95 ± 1.21	1.58 ± 0.66	187.55 ± 1.12	2.51 ± 0.94	131.89 ± 0.93	13.17 ± 1.23	1.37 ± 0.08
	ADM-SE	97.69 ± 3.02	0.77 ± 0.18	180.12 ± 2.46	1.85 ± 0.70	130.51 ± 0.99	8.21 ± 1.98	2.46 ± 0.47
	ADM-AE	96.31 ± 0.32	0.75 ± 0.02	177.88 ± 0.34	1.23 ± 0.12	127.44 ± 0.41	6.86 ± 0.23	1.29 ± 0.09
	ADM-RLS	97.50 ± 0.55	0.74 ± 0.05	175.46 ± 0.38	1.30 ± 0.21	129.98 ± 0.52	8.07 ± 0.37	2.14 ± 0.34
	GKS-SGD	107.56 ± 1.31	1.38 ± 0.66	187.34 ± 1.19	2.50 ± 0.94	131.72 ± 0.97	13.42 ± 1.22	1.34 ± 0.08
	GKS-SE	96.75 ± 1.19	1.03 ± 0.14	180.53 ± 2.37	1.69 ± 0.60	130.23 ± 0.96	12.14 ± 1.53	2.47 ± 0.29
	GKS-AE	94.89 ± 1.38	0.81 ± 0.05	179.93 ± 0.62	1.29 ± 0.19	129.45 ± 0.55	10.95 ± 0.86	1.11 ± 0.12
	GKS-RLS	93.42 ± 1.29	0.79 ± 0.11	178.13 ± 1.14	1.32 ± 0.18	130.63 ± 0.91	12.74 ± 0.95	2.57 ± 0.15
	STG-SGD	103.55 ± 1.29	1.48 ± 0.66	186.34 ± 1.18	2.50 ± 0.94	131.72 ± 0.97	13.45 ± 1.22	1.19 ± 0.08
	STG-SE	98.91 ± 2.99	0.75 ± 0.19	179.44 ± 2.47	1.96 ± 0.62	129.83 ± 0.97	8.40 ± 1.93	2.54 ± 0.45
	STG-RLS	95.26 ± 0.98	0.71 ± 0.03	175.95 ± 0.31	1.26 ± 0.19	128.35 ± 0.35	7.91 ± 0.43	2.24 ± 0.12
	STG-AE (ours)	92.22 ± 0.38	0.67 ± 0.03	174.81 ± 0.27	1.24 ± 0.17	123.32 ± 0.29	6.74 ± 0.32	0.98 ± 0.07
STS-AE (ours)	87.28 ± 0.32	0.60 ± 0.04	177.88 ± 0.33	1.19 ± 0.16	122.43 ± 0.35	6.60 ± 0.34	0.89 ± 0.08	
CNN	RND-SGD	106.66 ± 2.03	1.99 ± 1.17	187.14 ± 1.64	2.59 ± 1.05	132.06 ± 0.98	13.11 ± 1.16	1.46 ± 0.09
	RND-SE	97.66 ± 4.08	0.80 ± 0.11	181.95 ± 3.17	1.91 ± 0.59	130.27 ± 1.17	8.27 ± 1.98	2.78 ± 0.52
	RND-AE	96.32 ± 0.21	0.76 ± 0.02	180.92 ± 0.19	1.32 ± 0.10	131.46 ± 0.30	7.79 ± 0.27	1.39 ± 0.07
	RND-RLS	95.86 ± 0.58	0.79 ± 0.07	181.21 ± 0.28	1.28 ± 0.15	129.35 ± 0.47	7.95 ± 0.35	2.17 ± 0.26
	ADM-SGD	106.59 ± 2.03	1.99 ± 1.17	187.20 ± 1.58	2.59 ± 1.05	132.04 ± 0.99	13.11 ± 1.16	1.45 ± 0.09
	ADM-SE	97.64 ± 4.08	0.79 ± 0.09	180.99 ± 3.13	1.94 ± 0.56	130.27 ± 1.17	8.23 ± 0.21	3.01 ± 0.54
	ADM-AE	95.26 ± 0.24	0.74 ± 0.01	179.96 ± 0.16	1.25 ± 0.09	128.43 ± 0.25	6.76 ± 0.25	1.37 ± 0.07
	ADM-RLS	96.12 ± 0.39	0.77 ± 0.05	180.11 ± 0.19	1.27 ± 0.08	129.14 ± 0.53	7.51 ± 0.29	2.03 ± 0.23
	GKS-SGD	104.51 ± 2.11	1.98 ± 1.17	185.67 ± 2.17	2.58 ± 1.01	131.71 ± 0.99	13.05 ± 1.11	1.39 ± 0.08
	GKS-SE	97.15 ± 4.02	0.81 ± 0.08	181.70 ± 3.13	2.09 ± 0.41	130.43 ± 1.13	8.30 ± 1.70	2.61 ± 0.35
	GKS-AE	95.31 ± 0.22	0.73 ± 0.03	178.92 ± 0.18	1.26 ± 0.15	128.45 ± 0.24	6.89 ± 0.28	1.38 ± 0.08
	GKS-RLS	96.05 ± 0.61	0.76 ± 0.07	178.41 ± 0.18	1.25 ± 0.13	128.34 ± 0.35	7.38 ± 0.32	2.15 ± 0.28
	STG-SGD	104.52 ± 3.13	1.98 ± 1.17	184.87 ± 2.67	2.58 ± 1.05	131.71 ± 0.99	13.05 ± 1.10	1.42 ± 0.08
	STG-SE	98.12 ± 4.07	0.80 ± 0.14	181.63 ± 3.15	2.12 ± 0.48	130.42 ± 1.13	8.17 ± 1.93	2.97 ± 0.56
	STG-RLS	94.98 ± 0.25	0.71 ± 0.03	178.98 ± 0.20	1.24 ± 0.13	125.86 ± 0.21	6.42 ± 0.24	2.03 ± 0.21
	STG-AE (ours)	93.23 ± 0.17	0.66 ± 0.04	177.93 ± 0.15	1.20 ± 0.12	123.47 ± 0.26	5.86 ± 0.22	1.39 ± 0.12
STS-AE (ours)	88.23 ± 0.13	0.62 ± 0.04	173.93 ± 0.17	1.24 ± 0.10	121.39 ± 0.27	5.78 ± 0.26	1.35 ± 0.10	
MLP	RND-SGD	104.59 ± 1.43	1.53 ± 0.66	187.34 ± 1.32	2.46 ± 1.03	134.37 ± 1.49	13.06 ± 1.19	3.93 ± 0.57
	RND-SE	96.46 ± 5.37	0.82 ± 0.13	181.75 ± 2.22	1.38 ± 0.09	132.21 ± 1.96	9.95 ± 2.31	4.61 ± 0.74
	RND-AE	95.70 ± 0.65	0.78 ± 0.03	180.80 ± 0.26	1.35 ± 0.12	128.36 ± 0.35	9.74 ± 0.34	3.86 ± 0.62
	RND-RLS	96.10 ± 0.78	0.80 ± 0.04	180.79 ± 0.29	1.35 ± 0.18	128.89 ± 0.53	9.88 ± 0.36	4.27 ± 0.50
	ADM-SGD	104.59 ± 1.43	1.53 ± 0.66	186.35 ± 1.32	2.43 ± 1.03	132.19 ± 1.50	13.02 ± 1.09	3.70 ± 0.53
	ADM-SE	94.46 ± 5.37	0.79 ± 0.13	180.80 ± 2.14	1.37 ± 0.09	130.44 ± 1.88	8.89 ± 2.39	4.34 ± 0.67
	ADM-AE	93.65 ± 0.76	0.75 ± 0.02	179.99 ± 0.32	1.34 ± 0.11	127.46 ± 0.31	7.95 ± 0.41	3.61 ± 0.55
	ADM-RLS	93.18 ± 0.80	0.77 ± 0.03	179.46 ± 0.38	1.33 ± 0.14	128.12 ± 0.38	8.40 ± 0.39	4.06 ± 0.50
	GKS-SGD	104.51 ± 1.48	1.51 ± 0.62	187.15 ± 1.28	2.42 ± 1.01	132.07 ± 1.30	13.05 ± 1.16	3.91 ± 0.54
	GKS-SE	94.69 ± 5.30	0.92 ± 0.10	181.47 ± 2.14	1.38 ± 0.12	130.58 ± 1.81	8.91 ± 1.71	4.48 ± 0.71
	GKS-AE	93.88 ± 0.89	0.81 ± 0.08	179.85 ± 0.30	1.34 ± 0.11	128.10 ± 0.35	8.46 ± 0.35	3.59 ± 0.54
	GKS-RLS	94.12 ± 0.77	0.79 ± 0.06	178.92 ± 0.28	1.33 ± 0.17	127.85 ± 0.35	8.53 ± 0.31	4.18 ± 0.47
	STG-SGD	104.59 ± 1.43	1.53 ± 0.66	187.07 ± 1.47	2.41 ± 1.03	132.08 ± 1.29	12.89 ± 1.11	3.98 ± 0.57
	STG-SE	94.33 ± 5.34	0.76 ± 0.13	181.14 ± 2.14	1.37 ± 0.10	130.18 ± 1.87	8.98 ± 2.32	4.90 ± 0.82
	STG-RLS	92.98 ± 0.80	0.74 ± 0.03	178.12 ± 0.29	1.33 ± 0.15	126.65 ± 0.31	8.25 ± 0.30	4.06 ± 0.58
	STG-AE (ours)	91.65 ± 0.86	0.69 ± 0.02	177.97 ± 0.32	1.28 ± 0.12	124.29 ± 0.34	8.21 ± 0.33	2.98 ± 0.65
STS-AE (ours)	89.70 ± 0.68	0.72 ± 0.03	174.99 ± 0.33	1.31 ± 0.16	121.45 ± 0.29	7.90 ± 0.29	2.63 ± 0.55	




Figure 1: Representative results for the linear measurement model on the synthetic dataset. (a) Estimation error between $\hat{\beta}_*$ and the true mean fixed-point β_* ; lower is better. (b) Test RMSE for the STS-AE method when β_* is unknown (estimated online) vs. known; lower is better. Shaded regions represent ± 1 standard deviation from the mean.

(SGD) employs a standard SGD update for the model estimation ignoring the impact of drift, run to convergence at each stage. Our methods are **STG-AE** and **STS-AE**.

Performance Metrics. We evaluate two metrics. In both synthetic and real datasets, we compute the **test RMSE** between predicted and test labels, evaluated over per-stage hold-out test sets. This metric is directly tied to our design objective: for a fresh input $\mathbf{x}$, the expected prediction error decomposes as $\text{EPE}(\mathbf{x}) = \sigma^2 + \mathbf{x}^\top \Sigma_{\xi} \mathbf{x}$ (App. B, Eq. (33)), so a design that shrinks the posterior covariance Σ_{ξ} , equivalently that attains a higher log-determinant objective, lowers the expected test error. This explains why, throughout our experiments, methods with a higher design objective attain a lower test RMSE (see App. M.2). In synthetic experiments, we also measure **Model Error** as the Euclidean distance between the true model β_k and its estimate $\hat{\beta}_{k|k}$. All metrics are averaged over 100 runs with different random seeds, and we report mean and standard deviation.

4.2 Results

Table 2 reports the test RMSE for all methods across the four model architectures (Linear, ResNet, CNN, MLP) and seven datasets; to reduce clutter, we omit -K variants and report them in App. M. Across all settings, our methods STG-AE and STS-AE consistently achieve the lowest test RMSE. Adaptive estimation (-AE) is superior to static (-SE), SGD-based, and forgetting-factor RLS (-RLS) estimation, confirming that the (extended) Kalman filter effectively tracks the drifting model. Among design selection strategies, our optimal policies (STG, STS) outperform greedy information-gain (GKS), Adam-based (ADM) and random (RND) selection, with the advantage most pronounced for non-linear models. Comparing architectures, the non-linear models achieve substantially lower error than the linear model on all real-world datasets, e.g., on Beijing PM 2.5, the best non-linear method (STS-AE with ResNet, 87.28) improves over the best linear method (133.82) by over 34%. CNN and ResNet perform similarly, while MLP shows slightly higher variance on the synthetic scenario. Model error performance for synthetic scenarios is in App. M.

Fig. 1 shows the convergence of model estimate $\hat{\beta}_*$ to the true model β_* for the linear model in the synthetic scenario, when experiments are selected by STS-AE, STG-AE, ADM-AE, RND-AE, and GKS-AE. We also compare the test RMSE as a function of stage k for STS-AE and STS-AE-K, where β_* is known. We report additional comparisons between methods that utilize adaptive estimation with their -K variants, as well as the convergence of $\hat{\beta}_*$ to β_* , in App. M. We observe that the methods closely track -K variants, and that the estimator approximates β_* .

In Fig. 2, over the ResNet model, we evaluate the impact of varying the system parameters ρ , σ , and σ_w for the synthetic dataset, each varied independently while keeping the others fixed at their default values (see Table 4). Across all parameter settings, our proposed methods achieve lower test RMSE and model error compared to the baselines, demonstrating robustness under varying system dynamics. In particular, as ρ decreases (faster drift), the gap between adaptive (-AE) and the (-SE), (-SGD), and (-RLS) baselines widens, highlighting the importance of Kalman-based tracking when drift is strong.




Figure 2: Test RMSE (Eq. (145), top row) and model error (Eq. (148), bottom row) comparisons with respect to different values of ρ , σ , and σ_w on the synthetic dataset for the ResNet model; other parameters as in Table 4. Lower is better. Shaded regions represent ± 1 standard deviation from the mean.

Similarly, as σ_w increases (larger drift noise), our optimal selection policies (STG, STS) maintain a clear advantage over GKS, ADM and RND, indicating that informed sample selection becomes increasingly valuable when the model changes rapidly. We report the impacts of these parameters on the other models in App. M.

5 Conclusion

We propose *Markovian Experimental Design*, to determine an optimal data selection strategy over a model that is drifting according to linear dynamics coupled with Gaussian noise. The underlying stage-wise and steady-state optimization problems can be solved using convex optimization techniques. Several interesting research questions remain open. More complex model drift dynamics, beyond Eq. (3), require reasoning about the concavity of the D-Optimality objective, or another scalarization of the covariance, without explicitly solving the discrete algebraic Riccati equation. Another question of interest is whether the end-to-end optimization of the design selection admits a tractable, non-myopic, dynamic policy. Dynamic policies can also be studied in the adversarial setting, akin to Si Salem [62], which would remove stochasticity behind model dynamics. We discuss limitations in App. N.

Acknowledgments and Disclosure of Funding

We gratefully acknowledge the support from the National Science Foundation under grants 2112471, 2106891, 2107062, and the National Science Foundation CAREER award under grant 2340482.

References

- [1] Amani Al-Shawabka, Francesco Restuccia, Salvatore D’Oro, Tong Jian, Bruno Costa Rendon, Nasim Soltani, Jennifer Dy, Stratis Ioannidis, Kaushik Chowdhury, and Tommaso Melodia. Exposing the fingerprint: Dissecting the impact of the wireless channel on radio fingerprinting. In *IEEE INFOCOM 2020-IEEE Conference on Computer Communications*, pages 646–655. IEEE, 2020.
- [2] Mohammad Alali and Mahdi Imani. Deep reinforcement learning data collection for Bayesian inference of hidden Markov models. *IEEE transactions on artificial intelligence*, 2024.

- [3] Billy Amzal, Frédéric Y Bois, Eric Parent, and Christian P Robert. Bayesian-optimal design via interacting particle systems. *Journal of the American Statistical association*, 101(474):773–785, 2006.
- [4] Brian DO Anderson and John B Moore. *Optimal filtering*. Courier Corporation, 2005.
- [5] Stephen Bach and Mark Maloof. A Bayesian approach to concept drift. *Advances in neural information processing systems*, 23, 2010.
- [6] Manuel Baena-Garcia, José del Campo-Ávila, Raul Fidalgo, Albert Bifet, Ricard Gavaldà, and Rafael Morales-Bueno. Early drift detection method. In *Fourth international workshop on knowledge discovery from data streams*, volume 6, pages 77–86, 2006.
- [7] Albert Benveniste, Michel Métivier, and Pierre Priouret. *Adaptive algorithms and stochastic approximations*, volume 22. Springer Science & Business Media, 2012.
- [8] Gary Bishop, Greg Welch, et al. An introduction to the Kalman filter. *Proc of SIGGRAPH, Course*, 8(27599-23175):41, 2001.
- [9] Stephen Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.
- [10] Julian Bright, Alessandro Cerè, Georgios Schinas, and Theiss Heilker. Automate model retraining with Amazon SageMaker Pipelines when drift is detected. <https://aws.amazon.com/blogs/machine-learning/automate-model-retraining-with-amazon-sagemaker-pipelines-when-drift-is-detected/>, 2021.
- [11] Huanhuan Cao, Enhong Chen, Jie Yang, and Hui Xiong. Enhancing recommender systems under volatile user interest drifts. In *Proceedings of the 18th ACM conference on Information and knowledge management*, pages 1257–1266, 2009.
- [12] Kathryn Chaloner and Isabella Verdinelli. Bayesian experimental design: A review. *Statistical science*, pages 273–304, 1995.
- [13] Fukun Chen, Moinak Ghoshal, Enfu Nan, Phuc Dinh, Imran Khan, Z Jonny Kong, Y Charlie Hu, and Dimitrios Koutsounikolas. A large-scale study of the potential of multi-carrier access in the 5g era. In *International Conference on Passive and Active Network Measurement*, pages 469–484. Springer, 2025.
- [14] Song Chen. Beijing PM2.5. UCI Machine Learning Repository, 2015. DOI: <https://doi.org/10.24432/C5JS49>.
- [15] Susan Corbett. The retention of personal information online: A call for international regulation of privacy law. *Computer Law & Security Review*, 29(3):246–254, 2013.
- [16] Anton Dries and Ulrich Rückert. Adaptive concept drift detection. *Statistical Analysis and Data Mining: The ASA Data Science Journal*, 2(5-6):311–327, 2009.
- [17] Hovav A Dror and David M Steinberg. Sequential experimental designs for generalized linear models. *Journal of the American Statistical Association*, 103(481):288–298, 2008.
- [18] Christopher C Drovandi, James M McGree, and Anthony N Pettitt. Sequential monte carlo for Bayesian sequentially designed experiments for discrete data. *Computational Statistics & Data Analysis*, 57(1):320–335, 2013.
- [19] Shamak Dutta, Nils Wilde, and Stephen L Smith. A unified approach to optimally solving sensor scheduling and sensor selection problems in kalman filtering. In *2023 62nd IEEE Conference on Decision and Control (CDC)*, pages 1170–1176. IEEE, 2023.
- [20] Ky Fan. On a theorem of Weyl concerning eigenvalues of linear transformations i. *Proceedings of the National Academy of Sciences*, 35(11):652–655, 1949.
- [21] Hadi Fanaee-T and Joao Gama. Event labeling combining ensemble detectors and background knowledge. *Progress in Artificial Intelligence*, 2(2):113–127, 2014.

- [22] William Feller. *An introduction to probability theory and its applications*, volume 2. John Wiley & Sons, 1991.
- [23] Adam Foster, Desi R Ivanova, Ilyas Malik, and Tom Rainforth. Deep adaptive design: Amortizing sequential Bayesian experimental design. In *International conference on machine learning*, pages 3384–3395. PMLR, 2021.
- [24] Marguerite Frank and Philip Wolfe. An algorithm for quadratic programming. *Naval research logistics quarterly*, 3(1-2):95–110, 1956.
- [25] Robert G Gallager. *Stochastic processes: Theory for Applications*. Cambridge University Press, 2013.
- [26] William W Hager. Updating the inverse of a matrix. *SIAM review*, 31(2):221–239, 1989.
- [27] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [28] Xun Huan, Jayanth Jagalur, and Youssef Marzouk. Optimal experimental design: Formulations and computations. *Acta Numerica*, 33:715–840, 2024.
- [29] Xun Huan and Youssef M Marzouk. Simulation-based optimal Bayesian experimental design for nonlinear systems. *Journal of Computational Physics*, 232(1):288–317, 2013.
- [30] Xun Huan and Youssef M Marzouk. Sequential Bayesian optimal experimental design via approximate dynamic programming. *arXiv preprint arXiv:1604.08320*, 2016.
- [31] Aditi Jha, Zoe C Ashwood, and Jonathan W Pillow. Active learning for discrete latent variable models. *Neural computation*, 36(3):437–474, 2024.
- [32] Rudolph Emil Kalman. A new approach to linear filtering and prediction problems. *Journal of basic Engineering*, 82(1):35–45, 1960.
- [33] Jarek Kazmierczak, Khalid Salama, and Valentin Huerta. MLOps: Continuous delivery and automation pipelines in machine learning. <https://cloud.google.com/architecture/mlops-continuous-delivery-and-automation-pipelines-in-machine-learning>, 2024. Google Cloud Architecture center.
- [34] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [35] J Zico Kolter and Marcus A Maloof. Dynamic weighted majority: An ensemble method for drifting concepts. *The Journal of Machine Learning Research*, 8:2755–2790, 2007.
- [36] Ivan Koychev. Gradual forgetting for adaptation to concept drift. Proceedings of ECAI 2000 Workshop on Current Issues in Spatio-Temporal Reasoning, 2000.
- [37] Steven George Krantz and Harold R Parks. *The implicit function theorem: history, theory, and applications*. Springer Science & Business Media, 2002.
- [38] Ludmila I Kuncheva. Classifier ensembles for changing environments. In *International workshop on multiple classifier systems*, pages 1–15. Springer, 2004.
- [39] Harold J Kushner and G George Yin. *Stochastic approximation and recursive algorithms and applications*. Springer, 2003.
- [40] Lap Chi Lau, Robert Wang, and Hong Zhou. Experimental design using interlacing polynomials. In *2025 Symposium on Simplicity in Algorithms (SOSA)*, pages 452–464. SIAM, 2025.
- [41] Lap Chi Lau and Hong Zhou. A local search framework for experimental design. *SIAM Journal on Computing*, 51(4):900–951, 2022.
- [42] Jinpeng Li, Hang Yu, Zhenyu Zhang, Xiangfeng Luo, and Shaorong Xie. Concept drift adaptation by exploiting drift type. *ACM Transactions on Knowledge Discovery from Data*, 18(4):1–22, 2024.

- [43] Lennart Ljung and Torsten Söderström. *Theory and practice of recursive identification*. MIT press, 1983.
- [44] Andreas Lommatzsch, Benjamin Kille, and Sahin Albayrak. Incorporating context and trends in news recommender systems. In *Proceedings of the international conference on web intelligence*, pages 1062–1068, 2017.
- [45] Jie Lu, Anjin Liu, Fan Dong, Feng Gu, Joao Gama, and Guangquan Zhang. Learning under concept drift: A review. *IEEE transactions on knowledge and data engineering*, 31(12):2346–2363, 2018.
- [46] Horia Mania, Ali Jadbabaie, Devavrat Shah, and Suvrit Sra. Time varying regression with hidden linear dynamics. In *Learning for Dynamics and Control Conference*, pages 858–869. PMLR, 2022.
- [47] Albert W Marshall, Ingram Olkin, and Barry C Arnold. *Inequalities: theory of majorization and its applications*. Springer, 1979.
- [48] Peter Müller, Don A Berry, Andy P Grieve, Michael Smith, and Michael Krams. Simulation-based sequential Bayesian design. *Journal of statistical planning and inference*, 137(10):3140–3150, 2007.
- [49] Kevin P Murphy. *Machine learning: a probabilistic perspective*. MIT press, 2012.
- [50] Aleksandar Nikolov, Mohit Singh, and Uthaipon Tantipongpipat. Proportional volume sampling and approximation algorithms for a-optimal design. *Mathematics of Operations Research*, 47(2):847–877, 2022.
- [51] Cam Nugent. S&P 500 stock data. Kaggle, 2018. <https://www.kaggle.com/datasets/camnugent/sandp500>.
- [52] R Kelley Pace and Ronald Barry. Sparse spatial autoregressions. *Statistics & Probability Letters*, 33(3):291–297, 1997.
- [53] Thrasyvoulos Pappas, Alan Laub, and Nils Sandell. On the numerical solution of the discrete-time algebraic Riccati equation. *IEEE Transactions on Automatic Control*, 25(4):631–641, 1980.
- [54] Thomas Parr, Karl Friston, and Peter Zeidman. Active data selection and information seeking. *Algorithms*, 17(3):118, 2024.
- [55] Friedrich Pukelsheim. *Optimal design of experiments*. SIAM, 2006.
- [56] Randy J Read, Robert D Oeffner, and Airlie J McCoy. Measuring and using information gained by observing diffraction data. *Biological Crystallography*, 76(3):238–247, 2020.
- [57] Francesco Restuccia, Salvatore D’Oro, Amani Al-Shawabka, Mauro Belgiovine, Luca Angioloni, Stratis Ioannidis, Kaushik Chowdhury, and Tommaso Melodia. Deepradioid: Real-time channel-resilient optimization of deep learning-based radio fingerprinting algorithms. In *Proceedings of the twentieth ACM international symposium on mobile ad hoc networking and computing*, pages 51–60, 2019.
- [58] Herbert Robbins. Some aspects of the sequential design of experiments. 1952.
- [59] Herbert Robbins and Sutton Monro. A stochastic approximation method. *The annals of mathematical statistics*, pages 400–407, 1951.
- [60] Kunal Sankhe, Mauro Belgiovine, Fan Zhou, Shamnaz Riyaz, Stratis Ioannidis, and Kaushik Chowdhury. Oracle: Optimized radio classification through convolutional neural networks. In *IEEE INFOCOM 2019-IEEE conference on computer communications*, pages 370–378. IEEE, 2019.
- [61] Manohar Shamaiah, Siddhartha Banerjee, and Haris Vikalo. Greedy sensor selection: Leveraging submodularity. In *49th IEEE conference on decision and control (CDC)*, pages 2572–2577. IEEE, 2010.

- [62] Tareq Si Salem. Adaptive transductive inference via sequential experimental design with contextual retention. In *NeurIPS 2024 Workshop on Bayesian Decision-making and Uncertainty*, 2024.
- [63] Mohit Singh and Weijun Xie. Approximate positive correlated distributions and approximation algorithms for d-optimal design. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2240–2255. SIAM, 2018.
- [64] Arno Strouwen, Bart M Nicolaï, and Peter Goos. Adaptive and robust experimental design for linear dynamical models using Kalman filter. *Statistical Papers*, 64(4):1209–1231, 2023.
- [65] Dries Telen, Dominique Vercammen, Filip Logist, and Jan Van Impe. Robustifying optimal experiment design for nonlinear, dynamic (bio) chemical systems. *Computers & chemical engineering*, 71:415–425, 2014.
- [66] Christopher Tosh, Mauricio Tec, and Wesley Tansey. Targeted active learning for probabilistic models. *arXiv preprint arXiv:2210.12122*, 2022.
- [67] Alexander Tsesis. Data subjects’ privacy rights: regulation of personal data retention and erasure. *U. Colo. L. Rev.*, 90:593, 2019.
- [68] Abraham Wald. *Sequential analysis*. Courier Corporation, 1947.
- [69] Ke Wan, Yi Liang, and Susik Yoon. Online drift detection with maximum concept discrepancy. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 2924–2935, 2024.
- [70] Haixun Wang, Wei Fan, Philip S Yu, and Jiawei Han. Mining concept-drifting data streams using ensemble classifiers. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 226–235, 2003.
- [71] Shuo Wang, Leandro L Minku, and Xin Yao. A systematic study of online class imbalance learning with concept drift. *IEEE transactions on neural networks and learning systems*, 29(10):4802–4821, 2018.
- [72] Gerhard Widmer and Miroslav Kubat. Learning in the presence of concept drift and hidden contexts. *Machine learning*, 23(1):69–101, 1996.
- [73] Lintao Ye, Sandip Roy, and Shreyas Sundaram. On the complexity and approximability of optimal sensor selection for kalman filtering. In *2018 Annual American Control Conference (ACC)*, pages 5049–5054. IEEE, 2018.

Table 3: Summary of Notation

Markovian Experimental Design		Bayesian Linear Regression & Experimental Design	
$n_k \in \mathbb{N}$	# of samples at stage k	$n \in \mathbb{N}$	# of samples
$\mathcal{X}_k \subset \mathbb{R}^d$	candidate feature vectors at stage k	$\mathcal{X} \subset \mathbb{R}^d$	candidate feature vectors
$p_k \in \mathbb{N}$	# of candidate feature vectors $ \mathcal{X}_k $	$p \in \mathbb{N}$	# of candidate feature vectors $ \mathcal{X} $
$\beta_k \in \mathbb{R}^d$	model at stage k	$\beta \in \mathbb{R}^d$	model
$\mathbf{X}_k \in \mathbb{R}^{n_k \times d}$	design matrix at stage k	$\mathbf{X} \in \mathbb{R}^{n \times d}$	design matrix
$\mathbf{y}_k \in \mathbb{R}^{n_k}$	labels at stage k	$\mathbf{y} \in \mathbb{R}^n$	labels
$\mathbf{n}_k \in \mathbb{N}^{p_k}$	integral selection policy at stage k	$\mathbf{n} \in \mathbb{N}^p$	integral selection policy
$\boldsymbol{\pi}_k \in [0, 1]^{p_k}$	fractional selection policy at stage k	$\boldsymbol{\pi} \in [0, 1]^p$	fractional selection policy
$\hat{\beta}_k \in \mathbb{R}^d$	MAP estimate at stage k ($\equiv \hat{\beta}_{k k}$)	$\hat{\beta} \in \mathbb{R}^d$	MAP estimate
$\boldsymbol{\xi}_k \in \mathbb{R}^d$	error $\hat{\beta}_k - \beta_k$	$\boldsymbol{\xi} \in \mathbb{R}^d$	error $\hat{\beta} - \beta$
$\boldsymbol{\epsilon}_k \in \mathbb{R}^{n_k}$	additive i.i.d. noise at stage k	$\boldsymbol{\epsilon} \in \mathbb{R}^n$	additive i.i.d. noise
$\mathbf{w}_k \in \mathbb{R}^d$	Gaussian model drift at stage k	Common Notation	
$\mathbf{A} \in \mathbb{R}^{d \times d}$	linear drift coefficient	$d \in \mathbb{N}$	# of features
$\sigma_{\mathbf{w}}^2 \in \mathbb{R}_+$	Gaussian model drift variance	$\sigma^2 \in \mathbb{R}_+$	additive noise variance
$\hat{\beta}_{k k} \in \mathbb{R}^d$	MAP estimate at stage k after observing $(\mathbf{X}_k, \mathbf{y}_k)$	$\beta_0 \in \mathbb{R}^d$	prior mean
$\boldsymbol{\Sigma}_{\boldsymbol{\xi}_{k k-1}} \in \mathbb{R}^{d \times d}$	covariance estimate at stage k before observing $(\mathbf{X}_k, \mathbf{y}_k)$	$\boldsymbol{\Sigma}_0 \in \mathbb{R}^{d \times d}$	prior covariance
$\boldsymbol{\Sigma}_{\boldsymbol{\xi}_{k k}} \in \mathbb{R}^{d \times d}$	covariance estimate at stage k after observing $(\mathbf{X}_k, \mathbf{y}_k)$	$\mathbf{X}(\cdot)$	design matrix map (Eq. (38))
$\hat{\beta}_{k k-1} \in \mathbb{R}^d$	MAP estimate at stage k before observing $(\mathbf{X}_k, \mathbf{y}_k)$	$\mathbf{B}(\cdot) \in \mathbb{R}^{d \times d}$	Scaled sum of outer products $\mathbf{x}_i \mathbf{x}_i^\top$
$\bar{\boldsymbol{\Sigma}}_{\boldsymbol{\xi}} \in \mathbb{R}^{d \times d}$	steady-state covariance	$G(\cdot)$	D-optimal design objective
$\beta_* \in \mathbb{R}^d$	mean fixed-point		
$\hat{\beta}_* \in \mathbb{R}^d$	mean fixed-point estimate		
$\rho \in [0, 1]$	interpolation parameter		
$\bar{\mathbf{K}} \in \mathbb{R}^{d \times n}$	steady-state Kalman gain		

A Dynamical System Estimation using Kalman Filters

For completeness, we review dynamical system estimation via Kalman filters. The topic is classic; we refer the interested reader to Gallager [25] and Anderson & Moore [4], who we follow in the exposition below.

A.1 Dynamical System Model

Consider a dynamical system, that consists of a sequence of multivariate Gaussian random variables $\{\beta_k\}_{k=1}^\infty \in \mathbb{R}^d$ where β_1 is sampled from a prior normal distribution with mean $\beta_0 \in \mathbb{R}^d$ and covariance $\boldsymbol{\Sigma}_0 \in \mathbb{R}^{d \times d}$, i.e., $\beta_1 \sim \mathcal{N}(\beta_0, \boldsymbol{\Sigma}_0)$, and β_{k+1} evolves from β_k for $k \geq 1$ as follows:

$$\beta_{k+1} = \mathbf{A}_k \beta_k + \mathbf{w}_k, \quad \mathbf{w}_k \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_{\mathbf{w}_k}), \quad (17)$$

where $\mathbf{A}_k \in \mathbb{R}^{d \times d}$ is a known state transition matrix and $\mathbf{w}_k \in \mathbb{R}^d$ is the process noise for $k \geq 1$ with covariance matrix $\boldsymbol{\Sigma}_{\mathbf{w}_k} \in \mathbb{R}^{d \times d}$. In addition, observations $\mathbf{y}_k \in \mathbb{R}^n$ are given by:

$$\mathbf{y}_k = \mathbf{X}_k \beta_k + \mathbf{e}_k, \quad \mathbf{e}_k \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_{\mathbf{e}_k}), \quad (18)$$

where each measurement matrix yields $\mathbf{X}_k = [\mathbf{x}_{ik}^\top]_{i=1}^n \in \mathbb{R}^{n \times d}$ and $\mathbf{e}_k \in \mathbb{R}^n$ is the measurement noise for $k \geq 1$ with covariance matrix $\boldsymbol{\Sigma}_{\mathbf{e}_k} \in \mathbb{R}^{n \times n}$. Assuming the initial state β_1 and the noise vectors at each stage $\mathbf{w}_k$ and $\mathbf{e}_k$ are all mutually independent, Kalman filter evaluates the MAP estimate of parameter β_k at each stage respectively [32, 25]. The Kalman filter is a recursive algorithm for state estimation in systems with uncertain and noisy measurements, optimally combining past observations and a mathematical model to estimate the current state by minimizing the mean squared error between the estimated and true states. Being modeled on a Markov chain, it requires only the estimated state from the previous time step and the current measurement to compute the current estimate, unlike batch estimation techniques that need a history of observations and estimates [32, 25].

A.2 Estimation

Formally, the notation $\hat{\beta}_{k|k'}$ defines each stage estimation of the model vector at stage k given observations up to and including at stage $k' \leq k$. The respective errors are defined as $\xi_{k|k'} = \beta_k - \hat{\beta}_{k|k'}$. Each stage is represented by $\hat{\beta}_{k|k}$ and $\Sigma_{\xi_{k|k}}$, that are the *a posteriori* stage estimate mean and estimate covariance matrix at stage k given the observations up to and including at stage k respectively. Kalman filter equations can be divided into two phases, which are *a priori* (predict) and *a posteriori* (update) phases. The *a priori* phase of the Kalman filter equations is defined as [25, 8]:

$$\hat{\beta}_{k|k-1} = A_k \hat{\beta}_{k-1|k-1}, \quad (19)$$

$$\Sigma_{\xi_{k|k-1}} = A_k \Sigma_{\xi_{k-1|k-1}} A_k^\top + \Sigma_{w_k}. \quad (20)$$

After we see the evidence at stage k , we compute the Kalman gain as [8]:

$$K_k = \Sigma_{\xi_{k|k-1}} X_k^\top \left(X_k \Sigma_{\xi_{k|k-1}} X_k^\top + \Sigma_{e_k} \right)^{-1}, \quad (21)$$

and we update the state and covariance estimate as below, getting the *a posteriori* estimates at stage k [8]:

$$\hat{\beta}_{k|k} = \hat{\beta}_{k|k-1} + K_k \left(y_k - X_k \hat{\beta}_{k|k-1} \right), \quad (22)$$

$$\Sigma_{\xi_{k|k}} = (I - K_k X_k) \Sigma_{\xi_{k|k-1}}. \quad (23)$$

Here, Kalman gain K_k balances the contribution of the measurement X_k to minimize the posterior covariance by adjusting the trade-off between model predictions and new observations. By using the Woodbury matrix identity [26] on Eq. (23), the *a posteriori* covariance is alternatively defined without utilizing the Kalman gain as:

$$\Sigma_{\xi_{k|k}} = \left(X_k^\top \Sigma_{e_k}^{-1} X_k + \Sigma_{\xi_{k|k-1}}^{-1} \right)^{-1}. \quad (24)$$

The *a priori* estimate at $k = 1$ and the respective covariance are given by [25, 8]:

$$\hat{\beta}_{1|0} = \beta_0, \quad (25)$$

$$\Sigma_{\xi_{1|0}} = \Sigma_0. \quad (26)$$

The Kalman filter provides an optimal estimation by assuming that the errors w_k and e_k have a normal distribution. Moreover, regardless of the normal distribution, if the noise covariances Σ_{w_k} and Σ_{e_k} are known, then the Kalman filter serves as the optimal linear estimator [8].

A.3 Steady-State Kalman Filter

We review in this section asymptotic stability results pertaining to Kalman filters. We briefly introduce several definitions first.

Definition 1. A Kalman filter described in Eqs. (17)–(26) is time invariant if the system parameters are independent from time, i.e., the process noise covariance Σ_{w_k} and the measurement noise covariance Σ_{e_k} are constant over time, the measurement matrix X_k and the state transition matrix A_k are constant or converge to constant matrices for all k .

Definition 2 ([4]). A pair of matrices (F, G) is stabilizable if there exists a matrix H such that:

$$|\lambda_i(F + GH)| < 1, \quad (27)$$

for all i , i.e., such reconstructed matrix $(F + GH)$ is stable.

Definition 3 ([4]). A pair of matrices (F, G) is detectable if the dual system pair $(F^\top, G)$ is stabilizable (See Def. 2).

Armed with these definitions, we state here a classic convergence theorem regarding the steady-state of Kalman filters:

Theorem 6 ([4]). *If the Kalman filter described in Eqs. (17)–(26) is time invariant (see Def. 1), and asymptotically stable, i.e., $|\lambda_i(\mathbf{A})| < 1$, then the covariance and the Kalman gain converge to stationary matrices $\bar{\Sigma}_\xi$ and $\bar{\mathbf{K}}$ respectively. If the Kalman filter is time invariant (see Def. 1) but not necessarily asymptotically stable, i.e., $|\lambda_i(\mathbf{A})| \not< 1$, but the pair $[\mathbf{A}, \mathbf{X}^\top]$ is detectable (Def. 3), and the pair $[\mathbf{A}, \mathbf{G}]$ is stabilizable for any $\mathbf{G}$ with $\mathbf{G}\mathbf{G}^\top = \Sigma_w$ (Def. 2), then the covariance and the Kalman gain still converge to stationary matrices $\bar{\Sigma}_\xi$ and $\bar{\mathbf{K}}$ respectively.*

For both of the conditions, the Kalman filter gain and covariance matrices converge to stationary matrices as $k \rightarrow \infty$ [4]. The steady-state covariance $\Sigma_{\xi_{k|k}}$ converges to a constant matrix $\bar{\Sigma}_\xi$, that satisfies the discrete algebraic Riccati equation (DARE) [4]. The equation is derived by equating ($\Sigma_{\xi_{k|k}} = \Sigma_{\xi_{k-1|k-1}}$) via Eqs. (20) and (23):

$$\bar{\Sigma}_\xi = \mathbf{A}\bar{\Sigma}_\xi\mathbf{A}^\top + \Sigma_w - \mathbf{A}\bar{\Sigma}_\xi\mathbf{X}^\top \left(\mathbf{X}\bar{\Sigma}_\xi\mathbf{X}^\top + \Sigma_e \right)^{-1} \mathbf{X}\bar{\Sigma}_\xi\mathbf{A}^\top, \quad (28)$$

where $\mathbf{A}$, $\mathbf{X}$, and Σ_e are constant at steady state. Using the Woodbury-applied alternative of the posterior covariance in Eq. (24), the DARE can be re-written as:

$$\bar{\Sigma}_\xi = \left[\left(\mathbf{A}\bar{\Sigma}_\xi\mathbf{A}^\top + \Sigma_w \right)^{-1} + \mathbf{X}^\top \Sigma_e^{-1} \mathbf{X} \right]^{-1}. \quad (29)$$

This equation provides the unique solution for the steady-state covariance [4], [53]. Additionally, the Kalman gain converges to:

$$\bar{\mathbf{K}} = \bar{\Sigma}_\xi\mathbf{X}^\top \left(\mathbf{X}\bar{\Sigma}_\xi\mathbf{X}^\top + \Sigma_e \right)^{-1}. \quad (30)$$

B Bayesian Linear Regression and Experimental Design

In this section, we review the fundamentals of Bayesian Linear Regression and Experimental Design. We also formalize the connection between the expected prediction error (EPE) in Eq. (33) and the D-optimality criterion in Eq. (39). This topic is classic (see, e.g., Boyd & Vandenberghe [9]).

B.1 Bayesian Linear Regression

In Bayesian linear regression (see, e.g., Murphy [49] and Gallager [25]), a learner is given n samples $(\mathbf{x}_i, y_i)$, where $\mathbf{x}_i$ represents the feature vector and y_i is the corresponding label. There exists a model $\beta \in \mathbb{R}^d$ sampled from prior $\mathcal{N}(\beta_0, \Sigma_0)$ such that:

$$\mathbf{y} = \mathbf{X}\beta + \epsilon, \quad \epsilon \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}), \quad (31)$$

where $\mathbf{y} = [y_i]_{i=1}^n \in \mathbb{R}^n$ is the label vector, $\mathbf{X} = [\mathbf{x}_i^\top]_{i=1}^n \in \mathbb{R}^{n \times d}$ is the design matrix, and $\epsilon = [\epsilon_i]_{i=1}^n \in \mathbb{R}^n$ is the noise vector. The maximum a posteriori (MAP) estimate $\hat{\beta}^{\text{MAP}} = \arg \max_{\beta \in \mathbb{R}^d} \mathbf{P}(\beta | \mathbf{X}, \mathbf{y})$ is:

$$\hat{\beta}^{\text{MAP}} = \beta_0 + \Sigma_0 \mathbf{X}^\top (\mathbf{X} \Sigma_0 \mathbf{X}^\top + \sigma^2 \mathbf{I})^{-1} (\mathbf{y} - \mathbf{X} \beta_0). \quad (32)$$

Expected Prediction Error (EPE). Consider the parameter estimate $\hat{\beta}$ generated via Bayesian linear regression (Eq. (32)). Consider also a prediction at a test point $\mathbf{x} \in \mathbb{R}^d$. Then, the expected prediction error (EPE) at this unseen sample point $(\mathbf{x}, y) \in \mathbb{R}^d \times \mathbb{R}$ under estimate $\hat{\beta}$ is:

$$\mathbb{E} \left[\left(y - \mathbf{x}^\top \hat{\beta} \right)^2 \right] = \sigma^2 + \mathbf{x}^\top \Sigma_\xi \mathbf{x}, \quad (33)$$

where $\Sigma_\xi \in \mathbb{S}_+^d$ is the covariance of the error $\xi = \hat{\beta} - \beta$, i.e.,

$$\Sigma_\xi \equiv \text{cov}(\xi) = \left(\frac{1}{\sigma^2} \mathbf{X}^\top \mathbf{X} + \Sigma_0^{-1} \right)^{-1}. \quad (34)$$

The first term σ^2 of the EPE corresponds to irreducible observation noise, while the second term captures the contribution of estimation uncertainty and depends on both the test direction $\mathbf{x}$ and the geometry of Σ_ξ . Hence, the eigenvalues of the error covariance Σ_ξ capture the general variability of the EPE across all directions; we elaborate on this below.

EPE and Eigenvalues of the Covariance. Let the eigen-decomposition of the error covariance be

$$\Sigma_{\xi} = \mathbf{U} \Lambda \mathbf{U}^{\top}, \quad \Lambda = \text{diag}(\lambda_1, \dots, \lambda_d), \quad (35)$$

where $\{\lambda_i\}_{i=1}^d$ are the eigenvalues of Σ_{ξ} and $\mathbf{U} = [\mathbf{u}_1, \dots, \mathbf{u}_d]$ contains the corresponding orthonormal eigenvectors.

The test point $\mathbf{x}$ can then be written as a linear combination of the eigenvectors, i.e.,

$$\mathbf{x} = \sum_{i=1}^d c_i \mathbf{u}_i, \quad (36)$$

with the coefficients $c_i = \mathbf{u}_i^{\top} \mathbf{x}$. Substituting this representation into (33) yields

$$\mathbf{x}^{\top} \Sigma_{\xi} \mathbf{x} = \sum_{i=1}^d \lambda_i c_i^2. \quad (37)$$

Equation (37) shows that the EPE depends on how strongly the prediction direction $\mathbf{x}$ aligns with the eigenvectors of the error covariance and on the magnitude of the corresponding eigenvalues. Large eigenvalues correspond to directions of high predictive uncertainty, whereas small eigenvalues indicate well-estimated directions. As a result, the eigenvalues of the estimation error covariance capture the variability of the EPE across orthogonal prediction directions.

B.2 Experimental Design

In experimental design [12, 9, 55], a learner determines which $n \in \mathbb{N}$ specific experiments to conduct in order to accurately estimate linear model β with prior $\mathcal{N}(\beta_0, \Sigma_0)$. Formally, there exists a set of candidate feature vectors $\mathcal{X} = \{\mathbf{x}_i\}_{i=1}^p \subset \mathbb{R}^d$ corresponding to p possible experiment settings. The overall process consists of three steps, illustrated in Fig. 3. First, in the *design selection* step, the learner determines a vector $\mathbf{n} = [n_i]_{i=1}^p \in \mathbb{N}^p$ with $\mathbf{n}^{\top} \mathbf{1} \leq n$, comprising the number of times each experiment is performed. The design matrix is thus:

$$\mathbf{X}(\mathbf{n}) = \left[\begin{array}{c} \mathbf{x}_1^{\top} \\ \vdots \\ \mathbf{x}_1^{\top} \\ \vdots \\ \mathbf{x}_p^{\top} \\ \vdots \\ \mathbf{x}_p^{\top} \end{array} \right] \left\{ \begin{array}{l} n_1 \\ \vdots \\ n_p \end{array} \right\} \in \mathbb{R}^{n \times d}. \quad (38)$$

In the *label collection* step, the learner collects the label vector $\mathbf{y} = [y_i]_{i=1}^n \in \mathbb{R}^n$ corresponding to this design matrix, governed by Eq. (31). Finally, in the *model estimation* step, the learner estimates the model using Bayesian linear regression on the collected labels, i.e., via Eq. (32).

In *D-optimal design* [55], the selection $\mathbf{n} \in \mathbb{N}^p$ is determined by minimizing the log det of the posterior covariance, i.e.,

$$\log \det [\Sigma_{\xi}(\mathbf{n})] = -\log \det \left(\frac{1}{\sigma^2} \mathbf{B}(\mathbf{n}) + \Sigma_0^{-1} \right), \quad (39)$$

where $\mathbf{B}(\mathbf{n}) \equiv \sum_{i=1}^p n_i \mathbf{x}_i \mathbf{x}_i^{\top} \in \mathbb{R}^{d \times d}$. Minimizing this criterion is equivalent to minimizing the entropy of the posterior, as well as the EPE in Eq. (33) in all directions. For large n , a standard approach for computing the optimal design is by solving the following convex relaxation [9]:

$$\text{Max.: } G(\boldsymbol{\pi}; \sigma^2, \Sigma_0) \equiv \log \det \left[\frac{n}{\sigma^2} \mathbf{B}(\boldsymbol{\pi}) + \Sigma_0^{-1} \right] \quad (40a)$$

$$\text{s.t.: } \boldsymbol{\pi}^{\top} \mathbf{1} \leq 1, \boldsymbol{\pi} \geq \mathbf{0}, \quad (40b)$$

where $\boldsymbol{\pi} = [\pi_i]_{i=1}^p \in [0, 1]^p$ is a *fractional policy* and $\mathbf{B}(\boldsymbol{\pi}) \equiv \sum_{i=1}^p \pi_i \mathbf{x}_i \mathbf{x}_i^{\top} \in \mathbb{R}^{d \times d}$. The fractional solution can subsequently be rounded (see below), producing an approximate integral solution s.t. $\mathbf{n} \approx n\boldsymbol{\pi}$. As the objective G is concave in $\boldsymbol{\pi}$, Prob. (40) is a convex optimization problem that can be solved using standard techniques.




Figure 3: The classic experimental design setting. A learner having access to a prior distribution $\mathcal{N}(\beta_0, \Sigma_0)$ wishes to determine which experiments to conduct to estimate (hidden) model β . The entire process comprises three steps. In the *design selection* step, the learner selects n experiments to conduct, resulting in design matrix $\mathbf{X}(\mathbf{n}) \in \mathbb{R}^{n \times d}$. In the *label collection* step, corresponding experiments are conducted and labels $\mathbf{y} \in \mathbb{R}^n$ are revealed. In the *model estimation* step, the model is estimated via maximum a posteriori estimation. Only the covariance of the model prior is used in design selection, while the full prior is used in model estimation. In D-Optimal design, the design selection is determined by selecting samples that minimize the entropy of the posterior.

Rounding Schemes. Classic experimental design (e.g., Boyd & Vandenberghe [9]) assumes rounding π^* to integer $\mathbf{n}$ is deterministic. For large budget n , the relaxation is tight: $n_i = \lfloor n\pi_i^* \rfloor$ satisfies the budget up to a unit deviation per coordinate, and the information matrix $\sum_i n_i \mathbf{x}_i \mathbf{x}_i^\top$ is within $O(1/n)$ of the relaxed optimum. For moderate-to-large n , simple rounding is standard and sufficient. There are also more elaborate rounding schemes that preserve feasibility and objective quality. Singh [63] round via positively correlated distributions for D-design, and Nikolov [50] use proportional volume sampling for A-design. Lau [41] unify and improve both for D/A/E-design, including under additional knapsack constraints, and Lau [40] provide deterministic rounding for D/A/E-design that satisfies cardinality constraints with provable approximation guarantees. Even randomized rounding procedures above produce an allocation with high-probability quality guarantees via martingale concentration bounds.

D-Optimality Criterion vs. EPE. Rather than optimizing the EPE for a single test direction, the D-optimality criterion aims to control prediction uncertainty uniformly across all directions. This is achieved by considering the determinant of the information matrix Σ_ξ^{-1} . Since $\det(\Sigma_\xi^{-1}) = \prod_{i=1}^d \lambda_i^{-1}$, where λ_i are the eigenvalues of the error covariance in Eq. (35), it follows that

$$\log \det(\Sigma_\xi^{-1}) = - \sum_{i=1}^d \log \lambda_i. \quad (41)$$

Maximizing Eq. (41) thus penalizes large eigenvalues of the error covariance and promotes a balanced reduction of uncertainty across all eigen-directions. In this sense, the D-optimality criterion summarizes the variability of the expected prediction error over the entire input space, rather than along a single direction. By jointly controlling these eigenvalues through the log-determinant objective, D-optimality thus provides a principled scalar surrogate for minimizing predictive uncertainty when the test distribution is unknown or varies over time.

C Proof of Theorem 1

For $k \geq 1$, consider the following transformed model vector:

$$\tilde{\beta}_k \equiv \beta_k - \beta_*. \quad (42)$$

Eq. (3) implies that:

$$\tilde{\beta}_{k+1} = \rho \tilde{\beta}_k + \mathbf{w}_k, \quad \mathbf{w}_k \sim \mathcal{N}(\mathbf{0}, \sigma_w^2 \mathbf{I}). \quad (43)$$

Consider also the following “differential” labels, defined as:

$$\tilde{\mathbf{y}}_k \equiv \mathbf{y}_k - \mathbf{X}_k \beta_* \stackrel{\text{Eq. (60)}}{=} \mathbf{X}_k \beta_k - \mathbf{X}_k \beta_* + \epsilon_k \stackrel{\text{Eq. (42)}}{=} \mathbf{X}_k \tilde{\beta}_k + \epsilon_k, \quad (44)$$

where $\mathbf{y}_k$ is the original vector of collected labels. Eqs. (43)–(44) thus define identical dynamics and observations, respectively, as Eqs. (59)–(60), with $\mathbf{A} = \rho \mathbf{I}$. Hence, the transformed model $\tilde{\beta}_k$ can be estimated using the standard Kalman filter equations. Specifically, the a priori stage estimate of $\tilde{\beta}_k$ is:

$$\hat{\tilde{\beta}}_{k|k-1} = \rho \hat{\tilde{\beta}}_{k-1|k-1}, \quad (45)$$

and the a posteriori estimate is:

$$\hat{\tilde{\beta}}_{k|k} = \hat{\tilde{\beta}}_{k|k-1} + \mathbf{K}_k \left(\tilde{\mathbf{y}}_k - \mathbf{X}_k \hat{\tilde{\beta}}_{k|k-1} \right), \quad (46)$$

where the error covariance estimates and the Kalman gain are the same as in Eqs. (61b), (62b), and (63). Using Eq. (42), we re-write the a priori phase as:

$$\begin{aligned} \hat{\beta}_{k|k-1} - \beta_* &= \rho \left(\hat{\beta}_{k-1|k-1} - \beta_* \right) \\ &= \rho \hat{\beta}_{k-1|k-1} - \rho \beta_*. \end{aligned} \quad (47)$$

Thus, we have

$$\hat{\beta}_{k|k-1} = \rho \hat{\beta}_{k-1|k-1} + (1 - \rho) \beta_*. \quad (48)$$

Furthermore, we can re-write the a posteriori phase as:

$$\begin{aligned} \hat{\beta}_{k|k} - \beta_* &= \hat{\beta}_{k|k-1} - \beta_* + \mathbf{K}_k \left(\mathbf{y}_k - \mathbf{X}_k \beta_* - \mathbf{X}_k \left(\hat{\beta}_{k|k-1} - \beta_* \right) \right) \\ &= \hat{\beta}_{k|k-1} - \beta_* + \mathbf{K}_k \left(\mathbf{y}_k - \mathbf{X}_k \beta_* - \mathbf{X}_k \hat{\beta}_{k|k-1} + \mathbf{X}_k \beta_* \right) \\ &= \hat{\beta}_{k|k-1} - \beta_* + \mathbf{K}_k \left(\mathbf{y}_k - \mathbf{X}_k \hat{\beta}_{k|k-1} \right). \end{aligned} \quad (49)$$

Therefore, we get

$$\hat{\beta}_{k|k} = \hat{\beta}_{k|k-1} + \mathbf{K}_k \left(\mathbf{y}_k - \mathbf{X}_k \hat{\beta}_{k|k-1} \right). \quad (50)$$

Thus, the parameter β_k in the system defined in Eq. (3) can be estimated using the regular Kalman filter equations in (5) – (6). $\square$

D Proof of Theorem 2

From Eq. (3), we can recursively write the model vector as:

$$\begin{aligned} \beta_k &= \rho^{k-1} \beta_1 + (1 - \rho^{k-1}) \beta_* + \sum_{j=1}^{k-1} \rho^{k-1-j} \mathbf{w}_j \\ &= \beta_* + \rho^{k-1} (\beta_1 - \beta_*) + \sum_{j=1}^{k-1} \rho^{k-1-j} \mathbf{w}_j. \end{aligned} \quad (51)$$

Using the label collection in Eq. (4), the standard linear regression estimate is:

$$\begin{aligned}
\hat{\beta}_k^{\text{LR}} &= \left(\mathbf{X}_k^\top \mathbf{X}_k \right)^{-1} \mathbf{X}_k^\top \mathbf{y}_k \\
&= \left(\mathbf{X}_k^\top \mathbf{X}_k \right)^{-1} \mathbf{X}_k^\top (\mathbf{X}_k \beta_k + \epsilon_k) \\
&= \beta_k + \left(\mathbf{X}_k^\top \mathbf{X}_k \right)^{-1} \mathbf{X}_k^\top \epsilon_k \\
&= \beta_* + \rho^{k-1} (\beta_1 - \beta_*) + \sum_{j=1}^{k-1} \rho^{k-1-j} \mathbf{w}_j + \left(\mathbf{X}_k^\top \mathbf{X}_k \right)^{-1} \mathbf{X}_k^\top \epsilon_k
\end{aligned} \tag{52}$$

Taking the sample mean of the estimator, we have:

$$\frac{1}{k} \sum_{k'=1}^k \hat{\beta}_{k'} = \beta_* + (\beta_1 - \beta_*) \frac{1}{k} \sum_{k'=1}^k \rho^{k'-1} + \frac{1}{k} \sum_{k'=1}^k \sum_{j=1}^{k'-1} \rho^{k'-1-j} \mathbf{w}_j + \frac{1}{k} \sum_{k'=1}^k \left(\mathbf{X}_{k'}^\top \mathbf{X}_{k'} \right)^{-1} \mathbf{X}_{k'}^\top \epsilon_{k'} \tag{53}$$

Consider the first summation term as

$$\frac{1}{k} \sum_{k'=1}^k \rho^{k'-1} = \frac{1 - \rho^k}{k(1 - \rho)}. \tag{54}$$

Thus the (deterministic) term

$$(\beta_1 - \beta_*) \frac{1}{k} \sum_{k'=1}^k \rho^{k'-1} = (\beta_1 - \beta_*) \frac{1 - \rho^k}{k(1 - \rho)}, \tag{55}$$

converges to 0 as $k \rightarrow \infty$. Next, consider the drift noise term:

$$\frac{1}{k} \sum_{k'=1}^k \sum_{j=1}^{k'-1} \rho^{k'-1-j} \mathbf{w}_j = \frac{1}{k} \sum_{j=1}^{k-1} \mathbf{w}_j \sum_{k'=j+1}^k \rho^{k'-1-j} \tag{56}$$

$$= \frac{1}{k} \sum_{j=1}^{k-1} \mathbf{w}_j \frac{1 - \rho^{k-j}}{1 - \rho}. \tag{57}$$

Quantities $\mathbf{w}_j \frac{1 - \rho^{k-j}}{1 - \rho}$ are zero-mean random variables with bounded variance, as $\frac{1 - \rho^{k-j}}{1 - \rho} < \frac{1}{1 - \rho}$. Hence, the drift noise term in Eq. (56) is the sum of independent, zero-mean, non-identically distributed random variables with bounded variance. It thus satisfies the Kolmogorov criterion for the strong law of large numbers (see Feller [22], p. 269) and thus converges w.p. 1 to zero. Convergence in expectation is trivial (as variables $\mathbf{w}_j$ are zero-mean).

Lastly, our assumptions on the boundedness of eigenvalues of $(\mathbf{X}_{k'}^\top \mathbf{X}_{k'})^{-1}$ implies that random variables $(\mathbf{X}_{k'}^\top \mathbf{X}_{k'})^{-1} \mathbf{X}_{k'}^\top \epsilon_{k'}$ are also independent, zero-mean, and have bounded variance, so the term

$$\frac{1}{k} \sum_{k'=1}^k \left(\mathbf{X}_{k'}^\top \mathbf{X}_{k'} \right)^{-1} \mathbf{X}_{k'}^\top \epsilon_{k'}, \tag{58}$$

converges as $k \rightarrow \infty$ to 0, both w.p. 1 and in expectation, for the same reasons as above. $\square$

E Optimal Myopic/One-Step Lookahead Policies under General Linear Dynamics

In this appendix, we provide an optimal myopic/one-step lookahead policy for general linear dynamics. These dynamics, as described by Eq. (59), generalize the dynamics in Eq. (3), assuming a general linear evolution of drift through matrix $\mathbf{A}$. We show that we can design an optimal myopic/one-step lookahead policy for such general dynamics; the optimal myopic policy when the model evolves via Eq. (3) follows as a consequence (see App. E.1.2).

We note that, if matrix $\mathbf{A}$ is unknown, it can be estimated via traditional dynamical system identification techniques (see, e.g., [46]).




Figure 4: Our proposed Markovian experimental design setting. The experimental design scenario is repeated over multiple stages, during which the underlying model is subject to drift, as determined by Eq. (59). Compared to Fig. 3, the posterior from the previous stage (whose parameters are indicated by $(\cdot)_{k-1|k-1}$) is used to compute the prior at the next stage (indicated by $(\cdot)_{k|k-1}$), via Eq. (61). Note that these are not identical, as the model has drifted from stage $k-1$ to stage k . Moreover, model estimation is full-Bayesian, rather than a MAP estimate of β_k : through Eq. (62), the learner estimates the full posterior (i.e., both mean and variance, indicated by $(\cdot)_{k|k}$). This happens through a Kalman filter estimation process. As a result, the posterior estimated is in fact the posterior after seeing the entirety of observations $\{(\mathbf{X}_{k'}, \mathbf{y}_{k'})\}_{k'=1}^k$. Our analysis answers the question of how to optimally perform the design selection step in this setting (i.e., how to choose $\mathbf{n}_k$), either per-stage k or asymptotically, i.e., in steady-state.

E.1 System Model

We consider a dynamic setting in which time progresses in stages $k = 1, 2, \dots$, and the model vector β is subject to *drift* across stages, as illustrated in Fig. 4. Let $\{\beta_k\}_{k=1}^\infty \subseteq \mathbb{R}^d$ denote the models at different stages. The first model β_1 is sampled from a normal prior with mean $\beta_0 \in \mathbb{R}^d$ and covariance $\Sigma_0 \in \mathbb{R}^{d \times d}$, i.e., $\beta_1 \sim \mathcal{N}(\beta_0, \Sigma_0)$, while, for $k \geq 1$, β_{k+1} evolves from β_k as follows:

$$\beta_{k+1} = \mathbf{A}\beta_k + \mathbf{w}_k, \quad \mathbf{w}_k \sim \mathcal{N}(\mathbf{0}, \sigma_w^2 \mathbf{I}), \quad (59)$$

where $\mathbf{A}$ is a known time-invariant state transition matrix, and $\mathbf{w}_k \in \mathbb{R}^d$ is the drift noise that is i.i.d. across stages. At each stage, the learner interacts with the environment via the three steps of standard experimental design setting as presented in Sec. B.2. However, the process is repeated on each stage k , and both the design selection and model estimation steps are modified; we describe each of these steps in detail below.

Design Selection and Label Collection. Given a candidate set of feature vectors $\mathcal{X}_k \subset \mathbb{R}^d$ with $|\mathcal{X}_k| = p_k \in \mathbb{N}$, a learner selects which experiments to conduct at each stage k by creating a vector $\mathbf{n}_k = [n_{ik}]_{i=1}^{p_k} \in \mathbb{N}^{p_k}$ with $\mathbf{n}_k^\top \mathbf{1} = n_k$ at each respective stage. Under selection vector $\mathbf{n}_k$, the experiments to be conducted correspond to the design matrix $\mathbf{X}_k \equiv \mathbf{X}(\mathbf{n}_k)$, where $\mathbf{X}(\cdot)$ is as in Eq. (38).

After determining the design matrix, the learner collects labels, obtaining:

$$\mathbf{y}_k = \mathbf{X}_k \beta_k + \epsilon_k, \quad \epsilon_k \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}), \quad (60)$$

where $\mathbf{y}_k = [y_{ik}]_{i=1}^{n_k} \in \mathbb{R}^{n_k}$, $\mathbf{X}_k \in \mathbb{R}^{n_k \times d}$ are the vector of collected labels and the corresponding design matrix, respectively, and $\epsilon_k \in \mathbb{R}^{n_k}$ is the (i.i.d.) measurement noise. Note that initial state β_1 , the random drift vectors at each stage $\mathbf{w}_k$, and ϵ_k , are all assumed to be mutually independent.

Model Estimation. Finally, the learner estimates the model β_k at stage k given the entire sequence of samples $\{(\mathbf{X}_{k'}, \mathbf{y}_{k'})\}_{k'=1}^k$ up to stage k . Thus, departing from the standard experimental design setting, for each $k \geq 1$, the MAP estimate of β_k now corresponds to a Kalman filter estimation (see App. A for a review of Kalman filters). In particular, the model estimate at stage k can be adapted from the estimate at each stage $k-1$, through a series of recursive equations. We use the common Kalman filter notation $(\cdot)_{k|k'}$ to denote a quantity estimated at stage k , having access to samples up to stage $k' \leq k$. Then, at stage k , before seeing the samples $(\mathbf{X}_k, \mathbf{y}_k)$, the *a priori* model estimate and its covariance of the model β_k at stage k are [25, 8]:

$$\hat{\beta}_{k|k-1} = \mathbf{A}\hat{\beta}_{k-1|k-1}, \quad (61a)$$

$$\Sigma_{\xi_{k|k-1}} = \mathbf{A}\Sigma_{\xi_{k-1|k-1}}\mathbf{A}^\top + \sigma_w^2 \mathbf{I}. \quad (61b)$$

After the learner obtains the samples (i.e., the evidence) at stage k , the *a posteriori* estimate at stage k and its covariance are given by [8]:

$$\hat{\beta}_{k|k} = \hat{\beta}_{k|k-1} + \mathbf{K}_k (\mathbf{y}_k - \mathbf{X}_k \hat{\beta}_{k|k-1}), \quad (62a)$$

$$\Sigma_{\xi_{k|k}} = (\mathbf{I} - \mathbf{K}_k \mathbf{X}_k) \Sigma_{\xi_{k|k-1}}, \quad (62b)$$

where $\mathbf{K}_k \in \mathbb{R}^{d \times n_k}$ is the Kalman gain [8]:

$$\mathbf{K}_k = \Sigma_{\xi_{k|k-1}} \mathbf{X}_k^\top (\mathbf{X}_k \Sigma_{\xi_{k|k-1}} \mathbf{X}_k^\top + \sigma^2 \mathbf{I})^{-1}, \quad (63)$$

and initial estimates $(\cdot)_{1|0}$ are given by the prior.

Eqs. (61a)–(63) establish a recursion through which the learner can compute the model posterior mean (Eq. (62a)) and covariance (Eq. (62b)) at the conclusion of each stage; the former ($\hat{\beta}_{k|k}$) is also the MAP estimate at stage k . Interestingly, no data-retention is required in this setting: the full information obtained from past observations is encapsulated in the model posterior, which becomes a prior for the next stage.

E.1.1 Optimal Stage-wise Experimental Design

Having described our system model, we turn our attention to the question of how the learner is to perform the experimental design selection (i.e., how should n_k be picked at every stage). We first focus on a *greedy/myopic* setting: in this setting, the learner optimizes the design selection at stage k in a stage-wise fashion (i.e., one stage at a time), having access to the history of past dynamics through the model posterior at stage $k-1$.

Our first (straightforward) result is that a D-optimal fractional stage-wise design policy can again be obtained by solving an appropriate convex optimization problem.

Theorem 7. *The greedy D-optimal fractional policy at stage k can be computed as a solution to the following relaxed optimization problem:*

$$\text{Maximize:} \quad G(\pi_k; \sigma^2, \Sigma_{\xi_{k|k-1}}) \equiv \log \det \left[\frac{n_k}{\sigma^2} \mathbf{B}(\pi_k) + \Sigma_{\xi_{k|k-1}}^{-1} \right], \quad (64a)$$

$$\text{s.t.:} \quad \pi_k^\top \mathbf{1} \leq 1, \pi_k \geq \mathbf{0}, \quad (64b)$$

where $\pi_k \in [0, 1]^{p_k}$ is the fractional policy at stage $k \in \mathbb{N}$, $\mathbf{B}(\pi) = \sum_{i=1}^{p_k} \pi_i \mathbf{x}_i \mathbf{x}_i^\top$, and $\Sigma_{\xi_{k|k-1}}^{-1} \in \mathbb{R}^{d \times d}$ is the inverse of the *a priori* covariance given by Eq. (61b).

Proof. Using the Woodbury matrix identity [26], the *a posteriori* covariance in Eq. (62b), can be re-written as:

$$\Sigma_{\xi_{k|k}} = \left(\frac{1}{\sigma^2} \mathbf{X}_k^\top \mathbf{X}_k + \Sigma_{\xi_{k|k-1}}^{-1} \right)^{-1}, \quad (65)$$

where $\mathbf{X}_k \equiv \mathbf{X}(\mathbf{n}_k)$ is the design matrix at stage k and $\Sigma_{\xi_{k|k-1}}^{-1}$ is the a priori covariance. Thus, given all decisions made by the learner so far, the optimal stage-wise design decision $\mathbf{n}_k$ can be found by maximizing the D-optimality criterion, given by:

$$\log \det [\Sigma_{\xi_{k|k}}^{-1}(\mathbf{n}_k)] = \log \det \left[\frac{1}{\sigma^2} \mathbf{X}(\mathbf{n}_k)^\top \mathbf{X}(\mathbf{n}_k) + \Sigma_{\xi_{k|k-1}}^{-1} \right]. \quad (66)$$

The theorem thus follows by considering the relaxed version of this objective. $\square$

Note that, given past decisions prior to stage k , $\Sigma_{\xi_{k|k-1}}^{-1}$ can be computed recursively, and is thus part of the input to Prob. (64a), while the design is fully determined by $\mathbf{n}_k \approx n_k \boldsymbol{\pi}_k$. Finally, Prob. (64a) is the same form as the classic D-optimal experimental design (c.f. Prob. (40)), and can again be solved via standard convex optimization techniques [9].

E.1.2 Proof of Thm. 3

We conclude this section by determining the optimal myopic/one-step lookahead policies under the model drift described in Eq. (3). Observe that the *a posteriori* covariance (Eq. (62b)) of the model estimation under the general linear dynamics is exactly the same as the one (Eq. (5)) under the model drift in (3), from Thm. 1. Thus, the optimal stage-wise design for the dynamics in (3) can be found similarly by maximizing the D-optimality criterion, given by (66).

An immediate corollary is that Thm. 7 readily extends to this scenario: we can compute the optimal stage-wise policy by solving an instance of Prob. (64a). Note that determining the optimal selection policy *does not require knowledge of β_** : Prob. (64a) receives candidate samples $\mathcal{X}$ and covariance $\Sigma_{\xi_{k|k-1}}^{-1}$ as input; the latter (by Eqs. (63), (5) and (6)), *does not depend on β_** .

E.1.3 Persistent Excitation

Provided that the set of candidate experiments $\mathcal{X}_k \subset \mathbb{R}^d$ contains d vectors in general position, persistent excitation can be guaranteed by mixing in a small amount of exploration in the policy. This can be done by policies of the form:

$$\boldsymbol{\pi}^\varepsilon = (1 - \varepsilon)\boldsymbol{\pi} + \varepsilon\boldsymbol{\mu} \quad (67)$$

with an exploration distribution $\boldsymbol{\mu}$ satisfying $\mathbf{B}(\boldsymbol{\mu}) \succ \mathbf{0}$. An alternative is to add the convex constraint $\mathbf{B}(\boldsymbol{\pi}) \succeq \delta \mathbf{I}$ to Probs. (9) and (13). As $\mathbf{B}(\boldsymbol{\pi})$ is affine in $\boldsymbol{\pi}$, in both approaches, the convexity of the corresponding optimization problems (Prob. (9) in the stage-wise case, and Prob. (13) in the steady-state case) is preserved.

Note that this is robust to rounding since $\mathbf{B}(\cdot)$ is affine and, hence, Lipschitz. In particular, for R the Euclidian diameter of $\mathcal{X}_k$, a rounded design $\boldsymbol{\pi}$ with $\|\hat{\boldsymbol{\pi}} - \boldsymbol{\pi}^\varepsilon\|_1 \leq \eta_\pi$ satisfies $\mathbf{B}(\hat{\boldsymbol{\pi}}) \succeq (\delta - R^2 \eta_\pi) \mathbf{I} \succ \mathbf{0}$ whenever $\delta > R^2 \eta_\pi$. Hence, adding an exploration term implies that $\text{rank}(\mathbf{X}_k) = d$ even under rounding, when the rounding error η_π is small enough.

F Proof of Theorem 4

The Markovian experimental design system denoted in Eq. (3) reaches a steady-state as the system parameters satisfy the asymptotic conditions that are required to observe convergence in the system. Following Definition 1, the system parameters are defined as $\Sigma_{w_k} = \sigma_w^2 \mathbf{I}$, $\Sigma_{\epsilon_k} = \sigma^2 \mathbf{I}$, where the state transition matrix is defined as $\mathbf{A} = \rho \mathbf{I}$. Additionally, assuming a stationary data collection policy during steady-state, $\mathbf{X}(\mathbf{n})$ does not change over the stages. Therefore, the setting is time invariant. Following Thm. 6, the setting is asymptotically stable as $|\lambda_i(\mathbf{A})| < 1$. Then, these settings yield a stationary covariance $\bar{\Sigma}$ and Kalman gain $\bar{\mathbf{K}}$.

Generalizing beyond the scope of Eq. (3), although the main setting in Eq. (59) is not necessarily asymptotically stable, i.e., not always $|\lambda_i(\mathbf{A})| < 1$, we still observe stability in covariance and Kalman gain as the detectability and stabilizability of $(\mathbf{A}, \mathbf{X}(\mathbf{n})^\top)$ and $(\mathbf{A}, G)$, for any G with $GG^\top = \Sigma_{w_k}$, are satisfied, respectively. Following Def. 3, the detectability of $(\mathbf{A}, \mathbf{X}(\mathbf{n})^\top)$ follows

the stabilizability of $(\mathbf{A}^\top, \mathbf{X}(\mathbf{n})^\top)$ by the duality. Having assumed $\mathbf{X}(\mathbf{n}) \in \mathbb{R}^{n \times d}$ to be rank d , $(\mathbf{A}^\top, \mathbf{X}(\mathbf{n})^\top)$ is stabilizable if the matrix $\mathbf{H}$ in Eq. (27) is denoted as:

$$\mathbf{H} = -(\mathbf{X}(\mathbf{n})^\top)^\dagger \mathbf{A}^\top, \quad (68)$$

where $(\mathbf{X}(\mathbf{n})^\top)^\dagger$ is the Moore-Penrose pseudo-inverse of $\mathbf{X}(\mathbf{n})^\top$. Substituting $\mathbf{H}$, the equation becomes:

$$\mathbf{A}^\top + \mathbf{X}(\mathbf{n})^\top \mathbf{H} = \mathbf{A}^\top - \mathbf{X}(\mathbf{n})^\top (\mathbf{X}(\mathbf{n})^\top)^\dagger \mathbf{A}^\top. \quad (69)$$

Since $\mathbf{X}(\mathbf{n})$ has rank d , the $d \times n$ matrix $\mathbf{X}(\mathbf{n})^\top$ has full row rank, so the projection term satisfies $\mathbf{X}(\mathbf{n})^\top (\mathbf{X}(\mathbf{n})^\top)^\dagger = \mathbf{I}$. Hence $\mathbf{A}^\top + \mathbf{X}(\mathbf{n})^\top \mathbf{H} = \mathbf{A}^\top - \mathbf{A}^\top = \mathbf{0}$, all of whose eigenvalues lie (trivially) strictly within the unit disk. Thus, the equation (27) is always satisfied, i.e., $(\mathbf{A}^\top, \mathbf{X}(\mathbf{n})^\top)$ is stabilizable, which by duality proves that $(\mathbf{A}, \mathbf{X}(\mathbf{n})^\top)$ is detectable. Moreover, following Def. 2, the stabilizability of $(\mathbf{A}, \mathbf{G})$ begins by constructing the matrix $\mathbf{G}$ as:

$$\mathbf{G} = \sigma_w \mathbf{I}, \quad (70)$$

since $\mathbf{G}\mathbf{G}^\top = \sigma_w^2 \mathbf{I} = \Sigma_{w_k}$. As $\mathbf{G} = \sigma_w \mathbf{I}$ is full rank where $\sigma_w > 0$, $(\mathbf{A}, \mathbf{G})$ is stabilizable if the matrix $\mathbf{H}$ in Eq. (27) is denoted as:

$$\mathbf{H} = -\frac{1}{\sigma_w} \mathbf{A}. \quad (71)$$

Substituting $\mathbf{H}$, the equation becomes:

$$\mathbf{A} + \mathbf{G}\mathbf{H} = \mathbf{A} + \sigma_w \mathbf{I} \left(-\frac{1}{\sigma_w} \mathbf{A} \right) = \mathbf{0}, \quad (72)$$

which satisfies the eigenvalues of the equation to be strictly within the unit circle. Thus, the Eq. (27) is always satisfied, i.e., $(\mathbf{A}, \mathbf{G})$ is stabilizable. $\square$

G Proof of Lemma 1

Consider the DARE-like equation:

$$\Sigma^{-1} = \frac{1}{\tilde{\sigma}^2} \mathbf{B} + (\rho^2 \Sigma + \sigma_w^2 \mathbf{I})^{-1}, \quad (73)$$

where $\tilde{\sigma} \equiv \frac{\sigma}{\sqrt{n}}$ and $\mathbf{B} \in \mathbb{S}_+^d$ is a PSD matrix. Let $\lambda_i, \mathbf{e}_i$ be the eigenvalues and eigenvectors of PSD matrix $\mathbf{B}$. We postulate that Eq. (15) is correct, and the eigen-decomposition of the covariance is:

$$\Sigma = \sum_{i=1}^d \lambda'_i \mathbf{e}_i \mathbf{e}_i^\top, \quad (74)$$

We then show that this satisfies the DARE-like equation (73); the lemma thus follows from the uniqueness of the eigen-decomposition of a PSD matrix.

Assuming Eq. (74) holds, we get the following eigen-decompositions:

$$(\rho^2 \Sigma + \sigma_w^2 \mathbf{I})^{-1} = \sum_{i=1}^d (\rho^2 \lambda'_i + \sigma_w^2)^{-1} \mathbf{e}_i \mathbf{e}_i^\top, \quad (75)$$

$$\Sigma^{-1} = \sum_{i=1}^d \lambda_i'^{-1} \mathbf{e}_i \mathbf{e}_i^\top, \quad (76)$$

and the term involving $\mathbf{B}$ satisfies

$$\frac{1}{\tilde{\sigma}^2} \mathbf{B} = \frac{1}{\tilde{\sigma}^2} \sum_i \lambda_i \mathbf{e}_i \mathbf{e}_i^\top. \quad (77)$$

Re-arranging terms, we re-write Eq. (73) as:

$$(\rho^2 \mathbf{\Sigma} + \sigma_w^2 \mathbf{I})^{-1} - \mathbf{\Sigma}^{-1} + \frac{1}{\tilde{\sigma}^2} \mathbf{B} = \mathbf{0}, \quad (78)$$

Re-writing this using Eqs. (75)–(77), we get:

$$\begin{aligned} (\rho^2 \mathbf{\Sigma} + \sigma_w^2 \mathbf{I})^{-1} - \mathbf{\Sigma}^{-1} + \frac{1}{\tilde{\sigma}^2} \mathbf{B} &= \sum_{i=1}^d (\rho^2 \lambda'_i + \sigma_w^2)^{-1} \mathbf{e}_i \mathbf{e}_i^\top - \sum_{i=1}^d \lambda_i^{-1} \mathbf{e}_i \mathbf{e}_i^\top + \frac{1}{\tilde{\sigma}^2} \sum_{i=1}^d \lambda_i \mathbf{e}_i \mathbf{e}_i^\top \\ &= \sum_{i=1}^d \frac{\lambda'_i - \rho^2 \lambda'_i - \sigma_w^2}{\rho^2 \lambda_i'^2 + \sigma_w^2 \lambda'_i} \mathbf{e}_i \mathbf{e}_i^\top + \frac{1}{\tilde{\sigma}^2} \sum_{i=1}^d \lambda_i \mathbf{e}_i \mathbf{e}_i^\top \\ &= \sum_{i=1}^d \frac{\tilde{\sigma}^2 \lambda'_i - \rho^2 \tilde{\sigma}^2 \lambda'_i - \sigma_w^2 \tilde{\sigma}^2 + \rho^2 \lambda_i \lambda_i'^2 + \sigma_w^2 \lambda_i \lambda'_i}{\rho^2 \tilde{\sigma}^2 \lambda_i'^2 + \sigma_w^2 \tilde{\sigma}^2 \lambda'_i} \mathbf{e}_i \mathbf{e}_i^\top \\ &= \mathbf{0}. \end{aligned} \quad (79)$$

The last equation is equivalent to:

$$\tilde{\sigma}^2 \lambda'_i - \rho^2 \tilde{\sigma}^2 \lambda'_i - \sigma_w^2 \tilde{\sigma}^2 + \rho^2 \lambda_i \lambda_i'^2 + \sigma_w^2 \lambda_i \lambda'_i = 0, \quad i = 1, \dots, d. \quad (80)$$

Re-arranging terms, we get the quadratic equations w.r.t. λ'_i :

$$\lambda_i'^2 \rho^2 \lambda_i + \lambda'_i (\tilde{\sigma}^2 (1 - \rho^2) + \sigma_w^2 \lambda_i) - \sigma_w^2 \tilde{\sigma}^2 = 0. \quad (81)$$

In order to solve the quadratic equation, we find Δ as:

$$\Delta = (\tilde{\sigma}^2 (1 - \rho^2) + \sigma_w^2 \lambda_i)^2 + 4\rho^2 \lambda_i \sigma_w^2 \tilde{\sigma}^2. \quad (82)$$

Observe that $\Delta \geq 0$ as the eigenvalues of PSD matrix $\mathbf{B}$ are always non-negative, i.e., $\lambda_i \geq 0$. Moreover, we are interested in positive solutions, as the covariance matrix is also PSD. Therefore, taking the positive root of the two distinct real roots, we write the eigenvalues λ'_i of covariance $\mathbf{\Sigma}$ as:

$$\lambda'_i = \frac{\sqrt{(\tilde{\sigma}^2 (1 - \rho^2) + \sigma_w^2 \lambda_i)^2 + 4\lambda_i \rho^2 \sigma_w^2 \tilde{\sigma}^2} - (\tilde{\sigma}^2 (1 - \rho^2) + \sigma_w^2 \lambda_i)}{2\rho^2 \lambda_i}. \quad (83)$$

Rationalizing this expression by multiplying numerator and denominator by the conjugate $\sqrt{\Delta} + (\tilde{\sigma}^2 (1 - \rho^2) + \sigma_w^2 \lambda_i)$, and using $\Delta - (\tilde{\sigma}^2 (1 - \rho^2) + \sigma_w^2 \lambda_i)^2 = 4\lambda_i \rho^2 \sigma_w^2 \tilde{\sigma}^2$ gives the equivalent form

$$\lambda'_i = \frac{2\sigma_w^2 \tilde{\sigma}^2}{\sqrt{(\tilde{\sigma}^2 (1 - \rho^2) + \sigma_w^2 \lambda_i)^2 + 4\lambda_i \rho^2 \sigma_w^2 \tilde{\sigma}^2} + (\tilde{\sigma}^2 (1 - \rho^2) + \sigma_w^2 \lambda_i)}, \quad (84)$$

which is Eq. (15) (recall $\tilde{\sigma}^2 = \sigma^2/n$). The lemma follows by the uniqueness of the eigen-decomposition of $\mathbf{\Sigma} \in \mathbb{S}_+^d$. Note that Eq. (15) is well-defined for every $\lambda_i \geq 0$: at $\lambda_i = 0$ it equals $\sigma_w^2/(1 - \rho^2) > 0$. $\square$

H Proof of Theorem 5

We begin by giving an outline of the proof. Observe that:

$$G(\boldsymbol{\pi}) = \log \det [\overline{\mathbf{\Sigma}}_{\boldsymbol{\xi}}^{-1}(\boldsymbol{\pi})] = \sum_{i=1}^d \log g(\lambda_i(\boldsymbol{\pi})), \quad (85)$$

where $g : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is the map from $\lambda_i \mapsto \lambda_i'^{-1}$ and λ'_i is given by Eq. (15), and $\lambda_i(\boldsymbol{\pi})$, $i = 1, \dots, d$ are the eigenvalues of $\mathbf{B}(\boldsymbol{\pi})$. The proof proceeds as follows:

1. We first show that the map $g : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is non-decreasing and concave.

2. As a result, the map $\lambda \mapsto \sum_{i=1}^d \log g(\lambda_i)$ is Schur-concave (see Definition 4).
3. This implies that function $\Phi : \mathbb{S}_+^d \rightarrow \mathbb{R}$, defined as

$$\Phi(\mathbf{A}) = \sum_{i=1}^d \log g(\lambda_i(\mathbf{A})), \quad (86)$$

is concave over the positive semi-definite cone $\mathbb{S}_+^d$; moreover, as g is non-decreasing, Φ is non-decreasing in the Loewner order.

4. The theorem then follows from the fact that $G(\pi) = \Phi(\mathbf{B}(\pi))$, and $\mathbf{B}$ is affine and non-decreasing (in the Loewner order) in π .

We prove each of these four constituent steps below.

Concavity of g . Our first step is to show that g is non-decreasing and concave (see Lemma 2). Let $\lambda_i, \mathbf{e}_i$ be the eigenvalues and eigenvectors of the PSD matrix $\mathbf{B}(\pi) = \sum_{i=1}^p \pi_i \mathbf{x}_i \mathbf{x}_i^\top$, and λ'_i (given by Eq. (15)) the eigenvalues of the steady-state covariance $\bar{\Sigma}_\xi$; correspondingly $g(\lambda_i) = 1/\lambda'_i$ are the eigenvalues of $\bar{\Sigma}_\xi^{-1}$. We can thus re-write the objective in Eq. (13a) as:

$$G(\pi) \equiv \log \det [\bar{\Sigma}_\xi^{-1}(\pi)] \quad (87)$$

$$= \log \left(\prod_{i=1}^d g(\lambda_i(\pi)) \right) \quad (88)$$

$$= \sum_{i=1}^d \log g(\lambda_i(\pi)). \quad (89)$$

where the function $g : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is

$$g(\lambda) = 1/\lambda'(\lambda), \quad (90)$$

that defines the eigenvalues of the inverse covariance $\bar{\Sigma}_\xi^{-1}(\pi)$ inside the determinant of the objective.

Lemma 2. *Function $g : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ in Eq. (90) that is defined as:*

$$g(\lambda) = 1/\lambda'(\lambda), \quad (91)$$

is non-decreasing and concave.

Proof. Using the form of λ' in Eq. (15) gives

$$g(\lambda) = \frac{\sqrt{Q(\lambda)} + L(\lambda)}{2\sigma_w^2 \tilde{\sigma}^2}, \quad (92)$$

where $L(\lambda) = \tilde{\sigma}^2(1 - \rho^2) + \sigma_w^2 \lambda$ and $Q(\lambda) = L(\lambda)^2 + 4\lambda \rho^2 \sigma_w^2 \tilde{\sigma}^2$. Here L is affine, and $Q(\lambda) = \sigma_w^4 \lambda^2 + 2\sigma_w^2 \tilde{\sigma}^2(1 + \rho^2)\lambda + \tilde{\sigma}^4(1 - \rho^2)^2$ is a quadratic with positive leading coefficient, strictly positive for $\lambda \geq 0$, with discriminant

$$\Delta_Q = (2\sigma_w^2 \tilde{\sigma}^2(1 + \rho^2))^2 - 4\sigma_w^4 \tilde{\sigma}^4(1 - \rho^2)^2 = 16\rho^2 \sigma_w^4 \tilde{\sigma}^4 > 0. \quad (93)$$

Therefore $(\sqrt{Q})''(\lambda) = -\Delta_Q/(4Q(\lambda)^{3/2}) < 0$, i.e. $\sqrt{Q}$ is concave; being a positive multiple of the sum of the concave $\sqrt{Q}$ and the affine L , g is concave. Finally $L'(\lambda) = \sigma_w^2 > 0$ and $Q'(\lambda) = 2\sigma_w^4 \lambda + 2\sigma_w^2 \tilde{\sigma}^2(1 + \rho^2) > 0$, so $g'(\lambda) = (Q'(\lambda)/(2\sqrt{Q(\lambda)}) + \sigma_w^2)/(2\sigma_w^2 \tilde{\sigma}^2) > 0$; hence g is non-decreasing. Both properties hold for every $\lambda \geq 0$, including the boundary $\lambda = 0$. $\square$

Schur-concavity of map $\lambda \mapsto \sum_{i=1}^d \log g(\lambda_i)$. A consequence of Lemma 2 is that the map $\lambda \mapsto \sum_{i=1}^d \log g(\lambda_i)$ is Schur-concave; to show this, we introduce first the definition of Schur-convex and Schur-concave functions:

Definition 4. For $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$, we say that $\mathbf{x}$ majorizes $\mathbf{y}$, i.e., $\mathbf{x} \succ \mathbf{y}$, if and only if $\mathbf{y}$ is in the convex hull of all vectors obtained by permuting the coordinates of $\mathbf{x}$. By the Birkhoff–von Neumann theorem [47], this is equivalent to the following statement: there exists a doubly stochastic matrix $\mathbf{D}$ such that $\mathbf{y} = \mathbf{D}\mathbf{x}$. A function $\phi : \mathbb{R}^n \rightarrow \mathbb{R}$ is *Schur-convex* if $\mathbf{x} \succ \mathbf{y}$ implies [47]:

$$\phi(\mathbf{x}) \geq \phi(\mathbf{y}), \quad (94)$$

for all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$. Similarly, a function ϕ is *Schur-concave* if $\mathbf{x} \succ \mathbf{y}$ implies $\phi(\mathbf{x}) \leq \phi(\mathbf{y})$.

Armed with this definition, we have the following result:

Lemma 3. *Function $\phi : \mathbb{R}^d \rightarrow \mathbb{R}$ defined as*

$$\phi(\boldsymbol{\lambda}) = \sum_{i=1}^d \log g(\lambda_i) \quad (95)$$

is Schur-concave.

Proof. Using convexity composition rules (see Boyd & Vandenberghe [9]), $\log g$ is concave as $\log$ is concave and non-decreasing and g is concave from Lemma 2. We make use of the following lemma:

Lemma 4 (Marshall et al. [47]). *A function $\phi : \mathbb{R}^n \rightarrow \mathbb{R}$ that is defined as*

$$\phi(\mathbf{x}) = \sum_{i=1}^n g(x_i) \quad (96)$$

is Schur-concave if the function $g \in \mathbb{R} \rightarrow \mathbb{R}$ is concave.

Following this lemma, $\phi = \sum_{i=1}^d \log g(\lambda_i)$ is Schur-concave as $\log g$ is concave. $\square$

Concavity and monotonicity of Φ over the PSD cone. Next, we define a function $\Phi : \mathbb{S}_+^d \rightarrow \mathbb{R}$ as

$$\Phi(\mathbf{A}) = \sum_{i=1}^d \log g(\lambda_i(\mathbf{A})), \quad (97)$$

where $\mathbf{A} \in \mathbb{S}_+^d$. Recall that $\mathbf{B}(\boldsymbol{\pi}) = \sum_{i=1}^p \pi_i \mathbf{x}_i \mathbf{x}_i^\top$. Thus, we can write the objective function as

$$G(\boldsymbol{\pi}) = \sum_{i=1}^d \log g(\lambda_i(\boldsymbol{\pi})) = \Phi(\mathbf{B}(\boldsymbol{\pi})). \quad (98)$$

Next, we establish the monotonicity and concavity properties of map Φ .

Lemma 5. *Function $\Phi : \mathbb{S}_+^d \rightarrow \mathbb{R}$ in Eq. (97) that is defined as*

$$\Phi(\mathbf{A}) = \sum_{i=1}^d \log g(\lambda_i(\mathbf{A})), \quad (99)$$

is non-decreasing in the Loewner order (i.e., $\Phi(\mathbf{A}) \geq \Phi(\mathbf{B})$ whenever $\mathbf{A} \succeq \mathbf{B}$, that is, $\mathbf{A} - \mathbf{B} \in \mathbb{S}_+^d$) and concave over the positive-semidefinite cone $\mathbb{S}_+^d$.

Proof. Since function ϕ , defined in Eq. (95), is Schur-concave by Lemma 3, for all $\boldsymbol{\lambda}, \boldsymbol{\lambda}' \in \mathbb{R}^d$ with $\boldsymbol{\lambda} \prec \boldsymbol{\lambda}'$, we have:

$$\boldsymbol{\lambda} \prec \boldsymbol{\lambda}' \Rightarrow \phi(\boldsymbol{\lambda}) \geq \phi(\boldsymbol{\lambda}') \Rightarrow \sum_{i=1}^d \log g(\lambda_i) \geq \sum_{i=1}^d \log g(\lambda'_i), \quad (100)$$

where the majorization ' $\succ$ ' is defined in Definition 4. Moreover, Fan's Theorem [47, 20] states that, for all matrices $\mathbf{A}, \mathbf{B} \in \mathbb{S}_+^d$, we have:

$$\boldsymbol{\lambda}(\mathbf{A} + \mathbf{B}) \prec \boldsymbol{\lambda}(\mathbf{A}) + \boldsymbol{\lambda}(\mathbf{B}). \quad (101)$$

Therefore, the following inequality is satisfied

$$\sum_{i=1}^d \log g(\lambda_i(\mathbf{A} + \mathbf{B})) \geq \sum_{i=1}^d \log g(\lambda_i(\mathbf{A}) + \lambda_i(\mathbf{B})) \quad (102)$$

for all matrices $\mathbf{A}, \mathbf{B} \in \mathbb{S}_+^d$. Thus, the function Φ is concave as:

$$\begin{aligned}
\Phi(\theta \mathbf{A} + (1 - \theta) \mathbf{B}) &= \sum_{i=1}^d \log g(\lambda_i(\theta \mathbf{A} + (1 - \theta) \mathbf{B})) \\
&\stackrel{(102)}{\geq} \sum_{i=1}^d \log g(\theta \lambda_i(\mathbf{A}) + (1 - \theta) \lambda_i(\mathbf{B})) \\
&\stackrel{\text{Lemma 2}}{\geq} \sum_{i=1}^d (\theta \log g(\lambda_i(\mathbf{A})) + (1 - \theta) \log g(\lambda_i(\mathbf{B}))) \\
&= \theta \sum_{i=1}^d \log g(\lambda_i(\mathbf{A})) + (1 - \theta) \sum_{i=1}^d \log g(\lambda_i(\mathbf{B})) \\
&= \theta \Phi(\mathbf{A}) + (1 - \theta) \Phi(\mathbf{B}),
\end{aligned} \tag{103}$$

where the second inequality follows from the concavity of g . Finally, Φ is non-decreasing in the Loewner order: for $\mathbf{A}, \mathbf{B} \in \mathbb{S}_+^d$ with $\mathbf{A} \succeq \mathbf{B}$, write $\mathbf{A} = \mathbf{B} + \mathbf{P}$, where $\mathbf{P} = \mathbf{A} - \mathbf{B} \in \mathbb{S}_+^d$. Then

$$\Phi(\mathbf{A}) = \sum_{i=1}^d \log g(\lambda_i(\mathbf{B} + \mathbf{P})) \stackrel{(102)}{\geq} \sum_{i=1}^d \log g(\lambda_i(\mathbf{B}) + \lambda_i(\mathbf{P})) \geq \sum_{i=1}^d \log g(\lambda_i(\mathbf{B})) = \Phi(\mathbf{B}), \tag{104}$$

where the second inequality holds as $\lambda_i(\mathbf{P}) \geq 0$ and $\log g$ is non-decreasing (g is positive and non-decreasing by Lemma 2). $\square$

Concavity and monotonicity of G . Lemma 5 allows us to prove that the steady-state objective function $G : \mathbb{R}^p \rightarrow \mathbb{R}$ in Eq. (98) is indeed non-decreasing and concave. To see this, recall from Eq. (98) that we can write the objective as $G(\boldsymbol{\pi}) = \Phi(\mathbf{B}(\boldsymbol{\pi}))$, where $\Phi : \mathbb{S}_+^d \rightarrow \mathbb{R}$ is defined in Eq. (97) and $\mathbf{B}(\boldsymbol{\pi}) = \sum_{i=1}^p \pi_i \mathbf{x}_i \mathbf{x}_i^\top$. Observe that $\mathbf{B} : \mathbb{R}^p \rightarrow \mathbb{S}_+^d$ is affine in $\boldsymbol{\pi}$. Since, by Lemma 5, Φ is concave, function G is concave as a composition of a concave and an affine function. Moreover, $\mathbf{B}$ is non-decreasing in the Loewner order: for $\mathbf{0} \leq \boldsymbol{\pi} \leq \boldsymbol{\pi}'$ (componentwise), $\mathbf{B}(\boldsymbol{\pi}') - \mathbf{B}(\boldsymbol{\pi}) = \sum_{i=1}^p (\pi'_i - \pi_i) \mathbf{x}_i \mathbf{x}_i^\top \succeq \mathbf{0}$, as each $\mathbf{x}_i \mathbf{x}_i^\top \succeq \mathbf{0}$. Since, by Lemma 5, Φ is non-decreasing in the Loewner order, $G(\boldsymbol{\pi}') = \Phi(\mathbf{B}(\boldsymbol{\pi}')) \geq \Phi(\mathbf{B}(\boldsymbol{\pi})) = G(\boldsymbol{\pi})$, i.e., G is non-decreasing. $\square$

I Extension to A-Optimality

The D-optimality criterion we adopt already tracks the test RMSE closely (the expected prediction error decomposes as $\sigma^2 + \mathbf{x}^\top \boldsymbol{\Sigma}_\xi \mathbf{x}$, cf. Eq. (33)). Nonetheless, the same machinery shows that the *A-optimality* criterion, minimizing the trace (i.e. the average posterior variance) of the steady-state covariance, also yields a convex program.

Using the closed-form spectrum of Lemma 1, the A-optimal steady-state objective is

$$\text{tr } \bar{\boldsymbol{\Sigma}}_\xi(\boldsymbol{\pi}) = \sum_{i=1}^d \lambda'_i(\boldsymbol{\pi}) = \sum_{i=1}^d \phi(\lambda_i(\mathbf{B}(\boldsymbol{\pi}))), \quad \phi(\lambda) := \lambda'(\lambda) = \frac{1}{g(\lambda)}, \tag{105}$$

where λ'_i are the eigenvalues of $\bar{\boldsymbol{\Sigma}}_\xi$ and $g = 1/\lambda'$ is the positive, non-decreasing, concave map of Lemma 2.

Lemma 6. *The map $\phi(\lambda) = 1/g(\lambda)$ is non-increasing and convex on $\mathbb{R}_+$.*

Proof. Since $g > 0$ is non-decreasing, $\phi = 1/g$ is non-increasing. Since $g > 0$ is concave, $\phi = 1/g$ is convex: $\phi'' = (2(g')^2 - g g'')/g^3 \geq 0$, as $g'' \leq 0$ and $g > 0$. $\square$

Because ϕ is convex, the separable symmetric map $\boldsymbol{\lambda} \mapsto \sum_{i=1}^d \phi(\lambda_i)$ is *Schur-convex* [47], the exact mirror of the Schur-concavity of $\boldsymbol{\lambda} \mapsto \sum_i \log g(\lambda_i)$ used for D-optimality (Definition 4).

Repeating the argument in the proof of Theorem 5 (Fan’s theorem together with the affinity of $B(\pi) = \sum_i \pi_i \mathbf{x}_i \mathbf{x}_i^\top$ in π), the objective $\pi \mapsto \text{tr } \bar{\Sigma}_\xi(\pi) = \sum_{i=1}^d \phi(\lambda_i(B(\pi)))$ is convex over the PSD cone. Hence *minimizing* the A-optimal objective is a convex program, exactly mirroring the D-optimal case of App. H. Moreover, the implicit-gradient machinery of App. J carries over verbatim, with the log-determinant gradient $\partial \log \det$ replaced by the trace gradient ∂tr .

J Solving Prob. (13)

In our numerical evaluations, we find the optimal steady-state policy by solving Prob. (13). As established in Thm. 5, this is a convex optimization problem and thus admits efficient standard techniques; in practice, we implemented a solver using the Frank–Wolfe algorithm [24].

Both in this solver, but also in any optimization approach that uses a gradient evaluation (e.g., projected gradient descent, primal-dual methods, etc.), a key challenge lies in computing the gradient of the objective function in Eq. (13a). This requires differentiating the steady-state covariance defined implicitly through the DARE systems of equations (Eq. (12)). We first describe how this gradient is computed efficiently via the implicit function theorem [37], then explain how it is used within the Frank–Wolfe iterations, and evaluate the computational complexity of each iteration within the algorithm.

J.1 Estimating the Gradient of the Steady-State Objective (12) via the Implicit Function Theorem

Recall that the steady-state objective is written as $G(\pi) = \log \det(\bar{\Sigma}_\xi^{-1}(\pi))$. Define the inverse steady-state covariance

$$\bar{Q}(\pi) \triangleq \bar{\Sigma}_\xi^{-1}(\pi) \in \mathbb{S}_{++}^d, \quad (106)$$

so that $G(\pi) = \log \det(\bar{Q}(\pi))$. The matrix $\bar{Q}(\pi)$ is not available in closed form. Instead, it is defined implicitly as the unique solution to a Riccati-type fixed-point equation. To formalize this dependence, we define the mapping

$$F(Q, \pi) := Q - (\rho^2 Q^{-1} + \sigma_w^2 \mathbf{I})^{-1} - \frac{n}{\sigma^2} \sum_{i=1}^p \pi_i \mathbf{x}_i \mathbf{x}_i^\top. \quad (107)$$

This formulation corresponds to the information-form steady-state Riccati equation and is algebraically equivalent to Eq. (12). We then note that the inverse steady-state covariance $\bar{Q}(\pi)$ satisfies $F(\bar{Q}(\pi), \pi) = \mathbf{0}$. For a fixed allocation π , we compute the matrix $\bar{Q}(\pi)$ numerically via fixed-point iteration applied to the above equation until convergence. Since $F(Q, \pi)$ is continuously differentiable and its Jacobian with respect to Q is invertible at the steady-state, the implicit function theorem [37] guarantees that $\bar{Q}(\pi)$ is differentiable with respect to π . Moreover, for each coordinate π_j ,

$$\frac{\partial \bar{Q}}{\partial \pi_j} = - \left(\frac{\partial F}{\partial Q} \bigg|_{\bar{Q}} \right)^{-1} \frac{\partial F}{\partial \pi_j}, \quad (108)$$

where $\frac{\partial F}{\partial Q} \in \mathbb{S}^d$ denotes a linear operator acting on the space $\mathbb{S}^d$ of symmetric $d \times d$ matrices. The partial derivative with respect to π_j is given explicitly by $\frac{\partial F}{\partial \pi_j} = -\frac{n}{\sigma^2} \mathbf{x}_j \mathbf{x}_j^\top \in \mathbb{S}^d$. On the other hand, the Jacobian $\frac{\partial F}{\partial Q} \in \mathbb{S}^d$ admits the following Fréchet representation: for any perturbation $\Delta \in \mathbb{S}^d$,

$$\Delta \mapsto \Delta - \rho^2 M \bar{Q}^{-1} \Delta \bar{Q}^{-1} M, \quad M = (\rho^2 \bar{Q}^{-1} + \sigma_w^2 \mathbf{I})^{-1} \in \mathbb{S}_{++}^d. \quad (109)$$

Consequently, evaluating $\left(\frac{\partial F}{\partial Q} \right)^{-1} C \in \mathbb{S}^d$ for a given $C \in \mathbb{S}^d$ amounts to solving the two-sided Sylvester (Lyapunov-type) linear matrix equation

$$\Delta - \rho^2 M \bar{Q}^{-1} \Delta \bar{Q}^{-1} M = C, \quad \Delta \in \mathbb{S}^d, \quad (110)$$

which admits a unique solution under standard detectability conditions. Each sensitivity matrix $\frac{\partial \bar{\mathbf{Q}}}{\partial \pi_j} \in \mathbb{S}^d$ is obtained by solving this equation numerically. Finally, applying the chain rule yields the gradient of the steady-state objective:

$$\frac{\partial G(\boldsymbol{\pi})}{\partial \pi_j} = \text{tr}\left(\bar{\mathbf{Q}}^{-1} \frac{\partial \bar{\mathbf{Q}}}{\partial \pi_j}\right) = \text{tr}\left(\bar{\boldsymbol{\Sigma}}_{\boldsymbol{\xi}} \frac{\partial \bar{\mathbf{Q}}}{\partial \pi_j}\right) \in \mathbb{R}^p. \quad (111)$$

This procedure provides an exact first-order characterization of $\nabla G(\boldsymbol{\pi}) \in \mathbb{R}^p$, which is subsequently used by the Frank–Wolfe solver.

Efficient gradient via a single adjoint solve. Evaluating Eq. (108) separately for each coordinate solves the Lyapunov-type equation above p times, costing $O(p d^3)$. This reduces to a *single* solve. The operator $\Delta \mapsto \Delta - \rho^2 \mathbf{M} \bar{\mathbf{Q}}^{-1} \Delta \bar{\mathbf{Q}}^{-1} \mathbf{M}$ is self-adjoint with respect to the trace inner product (as $\mathbf{M}, \bar{\mathbf{Q}}^{-1} \in \mathbb{S}^d$), so by the chain rule and $\frac{\partial F}{\partial \pi_j} = -\frac{n}{\sigma^2} \mathbf{x}_j \mathbf{x}_j^\top$,

$$\frac{\partial G(\boldsymbol{\pi})}{\partial \pi_j} = \text{tr}\left(\bar{\boldsymbol{\Sigma}}_{\boldsymbol{\xi}} \frac{\partial \bar{\mathbf{Q}}}{\partial \pi_j}\right) = -\text{tr}\left(\boldsymbol{\Lambda} \frac{\partial F}{\partial \pi_j}\right) = \frac{n}{\sigma^2} \mathbf{x}_j^\top \boldsymbol{\Lambda} \mathbf{x}_j, \quad (112)$$

where the adjoint matrix $\boldsymbol{\Lambda} \in \mathbb{S}^d$ is the unique solution of the *single* Lyapunov-type equation

$$\boldsymbol{\Lambda} - \rho^2 \mathbf{M} \bar{\mathbf{Q}}^{-1} \boldsymbol{\Lambda} \bar{\mathbf{Q}}^{-1} \mathbf{M} = \bar{\boldsymbol{\Sigma}}_{\boldsymbol{\xi}} = \bar{\mathbf{Q}}^{-1}. \quad (113)$$

The entire gradient $\nabla G(\boldsymbol{\pi}) \in \mathbb{R}^p$ is thus obtained from one $d \times d$ solve for $\boldsymbol{\Lambda}$ ($O(d^3)$) followed by p quadratic forms $\mathbf{x}_j^\top \boldsymbol{\Lambda} \mathbf{x}_j$ ($O(p d^2)$), reducing the per-iteration cost from $O(p d^3)$ to $O(d^3 + p d^2)$.

J.2 Solving Prob. (13) using the Frank–Wolfe algorithm

Using the gradient characterization above, we solve Prob. (13) via the Frank–Wolfe algorithm [24] over the probability simplex

$$\Delta_p := \left\{ \boldsymbol{\pi} \in \mathbb{R}_{\geq 0}^p : \sum_{i=1}^p \pi_i = 1 \right\}. \quad (114)$$

The algorithm is initialized from a feasible point $\boldsymbol{\pi}_0 \in \Delta_p$, which we take to be the uniform allocation in our experiments. At each Frank–Wolfe iteration t , the objective is linearized around the current iterate $\boldsymbol{\pi}_t$, yielding the linear subproblem

$$\mathbf{s}_t \in \arg \max_{\mathbf{s} \in \Delta_p} \nabla G(\boldsymbol{\pi}_t)^\top \mathbf{s}. \quad (115)$$

Since the feasible set is a simplex, this subproblem admits a closed-form solution: $\mathbf{s}_t$ is the unit vector corresponding to the largest component of $\nabla G(\boldsymbol{\pi}_t) \in \mathbb{R}^p$. That is, the algorithm selects the coordinate $j^* = \arg \max_j [\nabla G(\boldsymbol{\pi}_t)]_j$, and assigns all mass to that index. The iterate is then updated via the convex combination

$$\boldsymbol{\pi}_{t+1} = (1 - \gamma_t) \boldsymbol{\pi}_t + \gamma_t \mathbf{s}_t, \quad (116)$$

where the step size $\gamma_t \in [0, 1]$ is selected by performing a one-dimensional line search along the Frank–Wolfe direction $\mathbf{s}_t - \boldsymbol{\pi}_t$. In our implementation, we employ exact line search on the original objective $G(\boldsymbol{\pi})$, which we compute using the *minimize_scalar* method from *scipy.optimize* with box constraints $[0, 1]$.

We emphasize, however, that exact line search is not required for the convergence of the Frank–Wolfe algorithm. In practice, comparable performance can be achieved using inexact line search rules (e.g., backtracking or Armijo line search) or standard diminishing step-size schedules. The use of exact line search in our implementation is primarily for numerical stability and ease of implementation, rather than algorithmic necessity. The algorithm terminates when the Frank–Wolfe duality gap

$$g_t = \nabla G(\boldsymbol{\pi}_t)^\top (\mathbf{s}_t - \boldsymbol{\pi}_t) \quad (117)$$

falls below a tolerance, indicating approximate first-order stationarity.

J.3 Computational Complexity

Each Frank–Wolfe iteration is dominated by the cost of evaluating the gradient $\nabla G(\pi_t) \in \mathbb{R}^p$. Computing the steady-state covariance $\bar{\Sigma}_\xi(\pi_t)$ requires solving a Riccati-type fixed-point equation, which in practice is obtained via fixed-point iteration and involves repeated matrix inversions of dimension d , incurring a cost of $O(d^3)$ per iteration.

Given $\bar{\Sigma}_\xi(\pi_t)$, each coordinate of the gradient requires solving a Lyapunov/Sylvester-type linear matrix equation of dimension $d \times d$, yielding a total cost of $O(p d^3)$ per gradient evaluation. Using the adjoint formulation of Eqs. (112)–(113), this drops to $O(d^3 + p d^2)$: a single $d \times d$ solve for the adjoint Λ followed by p quadratic forms. The remaining $O(d^3)$ term is the $d \times d$ covariance/DARE operation and becomes the bottleneck only for very high-dimensional (e.g., deep) models; there, standard low-rank or diagonal covariance approximations from Bayesian deep learning keep the *effective* d small. By contrast, the Frank–Wolfe linear subproblem admits a closed-form solution and amounts to identifying the largest component of the gradient, which costs $O(p)$.

When exact line search is used, each step-size selection requires multiple evaluations of the steady-state objective along a single direction. Since each such evaluation involves recomputing the steady-state covariance, the total cost of exact line search introduces an additional multiplicative factor. In practice, this overhead can be avoided by using inexact line search rules or predefined step-size schedules, which preserve convergence guarantees while substantially reducing computational cost.

Overall, the computational complexity of each Frank–Wolfe iteration is dominated by gradient evaluation, while the linear subproblem and update steps incur negligible additional overhead.

K Non-linear Setting

In this appendix, we provide the full details of the non-linear generalization summarized in Sec. 3.4. Throughout, we consider the system defined by the drift dynamics in Eq. (3) and the non-linear observation model in Eq. (2):

$$\beta_{k+1} = \rho \beta_k + (1 - \rho) \beta_* + \mathbf{w}_k, \quad \mathbf{w}_k \sim \mathcal{N}(\mathbf{0}, \sigma_w^2 \mathbf{I}), \quad (118)$$

$$\mathbf{y}_k = h(\mathbf{X}_k, \beta_k) + \epsilon_k, \quad \epsilon_k \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}), \quad (119)$$

where $h : \mathbb{R}^{n \times d} \times \mathbb{R}^{d_\beta} \rightarrow \mathbb{R}^n$ is a function parameterized by $\beta_k \in \mathbb{R}^{d_\beta}$.

K.1 Model Estimation in the Non-linear Setting

As the measurement function h is non-linear, the standard Kalman filter no longer applies. We instead employ the *Extended Kalman filter (EKF)* [8], which linearizes h around the current model estimate at each stage. Formally, the EKF linearizes h through its first-order Taylor approximation, so that

$$h(\mathbf{X}_k, \beta_k) \approx h(\mathbf{X}_k, \hat{\beta}_{k|k-1}) + \mathbf{H}_k(\beta_k - \hat{\beta}_{k|k-1}). \quad (120)$$

where

$$\mathbf{H}_k = \left. \frac{\partial h(\mathbf{X}_k, \beta)}{\partial \beta} \right|_{\beta = \hat{\beta}_{k|k-1}} \in \mathbb{R}^{n_k \times d}. \quad (121)$$

is the Jacobian matrix. In practice, this leads to the EKF adaptation equations that, compared to a standard Kalman filter, involve replacing the design matrix $\mathbf{X}_k$ with the Jacobian $\mathbf{H}_k$. Our approach naturally extends to this setting, as discussed below.

Known β_* . We can again show that, under our dynamics as well as approximation in (120), estimation reduces to an EKF when the mean fixed-point model β_* is known.

Theorem 8. *Consider the system evolving via Eqs. (2)–(3). Assume that the mean fixed-point β_* is known. Then, the Extended Kalman Filter maintains a Gaussian approximation $\mathcal{N}(\hat{\beta}_{k|k}, \Sigma_{k|k})$ to β_k , with the posterior mean and covariance:*

$$\hat{\beta}_{k|k} = \hat{\beta}_{k|k-1} + \mathbf{K}_k \left(\mathbf{y}_k - h(\mathbf{X}_k, \hat{\beta}_{k|k-1}) \right), \quad (122a)$$

$$\Sigma_{\xi_{k|k}} = (\mathbf{I} - \mathbf{K}_k \mathbf{H}_k) \Sigma_{\xi_{k|k-1}}, \quad (122b)$$

respectively, where $\hat{\beta}_{k|k-1}$, $\Sigma_{\xi_{k|k-1}}$ are the a priori estimates at stage k given in Eq. (6), $\mathbf{K}_k$ is the Kalman gain:

$$\mathbf{K}_k = \Sigma_{\xi_{k|k-1}} \mathbf{H}_k^\top \left(\mathbf{H}_k \Sigma_{\xi_{k|k-1}} \mathbf{H}_k^\top + \sigma^2 \mathbf{I} \right)^{-1}, \quad (123)$$

and $\mathbf{H}_k$ is the Jacobian matrix:

$$\mathbf{H}_k = \left. \frac{\partial h(\mathbf{X}_k, \beta)}{\partial \beta} \right|_{\beta = \hat{\beta}_{k|k-1}} \in \mathbb{R}^{n_k \times d}. \quad (124)$$

Proof. The prediction step follows directly from the dynamics in Eq. (3), which is identical to the linear case (Eq. (6)). For the posterior, we linearize the observation model around $\hat{\beta}_{k|k-1}$:

$$h(\mathbf{X}_k, \beta_k) \approx h(\mathbf{X}_k, \hat{\beta}_{k|k-1}) + \mathbf{H}_k (\beta_k - \hat{\beta}_{k|k-1}). \quad (125)$$

Under this first-order approximation, the observation model becomes locally linear: $\mathbf{y}_k \approx h(\mathbf{X}_k, \hat{\beta}_{k|k-1}) + \mathbf{H}_k (\beta_k - \hat{\beta}_{k|k-1}) + \epsilon_k$. Defining the *non-linear innovation* $\nu_k \equiv \mathbf{y}_k - h(\mathbf{X}_k, \hat{\beta}_{k|k-1})$ and applying the standard Kalman filter update with $\mathbf{H}_k$ in place of $\mathbf{X}_k$ yields Eqs. (122a)–(122b). The proof of Thm. 1 (App. C) carries over directly, since the transformation $\tilde{\beta}_k = \beta_k - \beta_*$ yields the same structure with $A = \rho \mathbf{I}$ in the state equation and $\mathbf{H}_k$ replacing $\mathbf{X}_k$ in the linearized observation equation. $\square$

As in the linear setting, the covariance recursion does not depend on β_* ; we can compute the optimal myopic policy without knowing β_* . We describe the optimal myopic policies for the non-linear setting in App. K.2.1.

Estimating β_* . As in the linear setting, we need an estimate of the mean fixed-point β_* . We obtain this via cumulative non-linear least squares:

$$\hat{\beta}_* \in \arg \min_{\beta \in \mathbb{R}^d} \sum_{t=1}^k \|\mathbf{y}_t - h(\mathbf{X}_t, \beta)\|_2^2. \quad (126)$$

Unlike the linear case (Thm. 2), where the least-squares problem is convex and the estimator provably converges to β_* , this non-linear estimation is generally non-convex, so the minimizer may not be unique. In practice, we observe that $\hat{\beta}_*$ stabilizes reliably across all model architectures considered in our experiments (see also Fig. 6 in App. M). We use this estimate as the linearization point for the steady-state design selection below.

K.2 Optimal Design Selection in the Non-linear Setting

We now describe how to perform D-optimal design selection in the non-linear setting. The key idea is that, under the linearization (120), the Jacobian vectors play the role of the feature vectors in the linear setting: the optimization problems retain the same convex structure, with the Jacobian evaluated at an appropriate linearization point.

K.2.1 Optimal Myopic Design Selection

At each stage k , the EKF has computed the a priori estimate $\hat{\beta}_{k|k-1}$ and covariance $\Sigma_{k|k-1}$ in Eq. (6). We linearize the observation model around $\hat{\beta}_{k|k-1}$, which is the best available estimate before observing the data at stage k . For each candidate $\mathbf{x}_{ik} \in \mathcal{X}_k$, define the Jacobian vector:

$$\mathbf{g}(\mathbf{x}_i, \hat{\beta}_{k|k-1}) = \left. \frac{\partial h(\mathbf{x}_i, \beta)}{\partial \beta} \right|_{\beta = \hat{\beta}_{k|k-1}} \in \mathbb{R}^{d_\beta}. \quad (127)$$

Applying the Woodbury matrix identity to the a posteriori covariance (122b) yields:

$$\Sigma_{\xi_{k|k}} = \left(\frac{1}{\sigma^2} \mathbf{H}_k^\top \mathbf{H}_k + \Sigma_{\xi_{k|k-1}}^{-1} \right)^{-1}. \quad (128)$$

where, under design vector $\mathbf{n}_k$, we have $\mathbf{H}_k^\top \mathbf{H}_k = \sum_{i=1}^p n_{ik} \mathbf{g}_i \mathbf{g}_i^\top$. The D-optimal myopic design thus amounts to solving:

$$\text{Maximize: } G\left(\boldsymbol{\pi}_k; \sigma^2, \boldsymbol{\Sigma}_{\boldsymbol{\xi}_{k|k-1}}, \hat{\boldsymbol{\beta}}_{k|k-1}\right) \equiv \log \det \left[\frac{n}{\sigma^2} \mathbf{B}(\boldsymbol{\pi}_k; \hat{\boldsymbol{\beta}}_{k|k-1}) + \boldsymbol{\Sigma}_{\boldsymbol{\xi}_{k|k-1}}^{-1} \right], \quad (129a)$$

$$\text{s.t.: } \boldsymbol{\pi}_k^\top \mathbf{1} \leq 1, \boldsymbol{\pi}_k \geq \mathbf{0}, \quad (129b)$$

where $\mathbf{B}(\boldsymbol{\pi}_k; \hat{\boldsymbol{\beta}}_{k|k-1}) = \sum_{i=1}^p \pi_{ik} g(\mathbf{x}_i, \hat{\boldsymbol{\beta}}_{k|k-1}) g(\mathbf{x}_i, \hat{\boldsymbol{\beta}}_{k|k-1})^\top$ is the linearized Fisher information. Since $\mathbf{g}_i$ is fixed at the current estimate $\hat{\boldsymbol{\beta}}_{k|k-1}$, the objective is log det of an affine function of $\boldsymbol{\pi}_k$ and is therefore concave. Prob. (129) is thus a convex optimization problem of the same form as the linear case (Prob. (9)), and can be solved using the same techniques. We stress again that, as in the linear case, the myopic policy *does not require knowledge of* $\boldsymbol{\beta}_*$.

K.2.2 Optimal Steady State Selection

For the steady-state analysis, recall that under the dynamics (3), the model $\boldsymbol{\beta}_k$ oscillates around the mean fixed-point $\boldsymbol{\beta}_*$ in steady state. This motivates us to analyze the steady state using the following approximations: first, we linearize the observation model around $\hat{\boldsymbol{\beta}}_*$, the estimate of the mean fixed-point obtained from (126); second, we treat this linearization as fixed for the purpose of the steady-state analysis. Concretely, we assume that the Jacobian $\mathbf{H}_k \approx \mathbf{H} \equiv \partial h(\mathbf{X}(\mathbf{n}), \boldsymbol{\beta}) / \partial \boldsymbol{\beta}|_{\hat{\boldsymbol{\beta}}_*}$ is approximately constant across stages. Under this assumption, the EKF covariance recursion becomes a time-invariant linear system, to which the standard convergence theory applies.

Convergence to Steady State. Under a stationary selection policy $\mathbf{n} \in \mathbb{N}^p$ and the linearizability assumption above, the EKF covariance recursion becomes time-invariant with state transition matrix $\mathbf{A} = \rho \mathbf{I}$, observation matrix $\mathbf{H}$, process noise $\sigma_w^2 \mathbf{I}$, and measurement noise $\sigma^2 \mathbf{I}$. Assuming $\mathbf{H}$ has rank d_β , Thm. 4 applies directly (with $\mathbf{H}$ replacing $\mathbf{X}(\mathbf{n})$): the pair $(\rho \mathbf{I}, \mathbf{H})$ is detectable and $(\rho \mathbf{I}, \sigma_w^2 \mathbf{I})$ is stabilizable, following from the same arguments as in App. F. Hence, the EKF covariance converges to $\bar{\boldsymbol{\Sigma}}_\xi \in \mathbb{R}^{d_\beta \times d_\beta}$ satisfying the DARE:

$$\bar{\boldsymbol{\Sigma}}_\xi^{-1} = \frac{1}{\sigma^2} \mathbf{H}^\top \mathbf{H} + (\rho^2 \bar{\boldsymbol{\Sigma}}_\xi + \sigma_w^2 \mathbf{I})^{-1}, \quad (130)$$

and the Kalman gain converges to $\mathbf{K} = \bar{\boldsymbol{\Sigma}}_\xi \mathbf{H}^\top (\mathbf{H} \bar{\boldsymbol{\Sigma}}_\xi \mathbf{H}^\top + \sigma^2 \mathbf{I})^{-1}$. Unlike the linear setting, where the covariance dynamics are fully decoupled from $\boldsymbol{\beta}_*$, the non-linear DARE (130) depends on $\boldsymbol{\beta}_*$ through the Jacobian $\mathbf{H}$. We use $\hat{\boldsymbol{\beta}}_*$ for two reasons. First, $\boldsymbol{\beta}_*$ is the point around which the model oscillates in steady state (cf. Eq. (3)), making it the most representative linearization point for the long-run behavior. Second, as $\hat{\boldsymbol{\beta}}_*$ stabilizes, the Jacobian $\mathbf{H}_k \rightarrow \mathbf{H} \equiv \partial h / \partial \boldsymbol{\beta}|_{\hat{\boldsymbol{\beta}}_*}$, and the DARE solution converges accordingly.

Optimal Stationary Policy. Under the linearizability assumption, the steady-state design optimization takes the same form as in the linear case. For each candidate $\mathbf{x}_i$, define the Jacobian vector:

$$g(\mathbf{x}_i, \hat{\boldsymbol{\beta}}_*) = \left. \frac{\partial h(\mathbf{x}_i, \boldsymbol{\beta})}{\partial \boldsymbol{\beta}} \right|_{\hat{\boldsymbol{\beta}}_*} \in \mathbb{R}^{d_\beta}, \quad (131)$$

and the linearized Fisher information $\mathbf{B}(\boldsymbol{\pi}; \hat{\boldsymbol{\beta}}_*) = \sum_{i=1}^p \pi_i g(\mathbf{x}_i, \hat{\boldsymbol{\beta}}_*) g(\mathbf{x}_i, \hat{\boldsymbol{\beta}}_*)^\top$. The optimal stationary policy solves:

$$\text{Maximize: } G\left(\boldsymbol{\pi}; \sigma^2, \sigma_w^2, \hat{\boldsymbol{\beta}}_*\right) \equiv \log \det \left[\bar{\boldsymbol{\Sigma}}_\xi^{-1}(\boldsymbol{\pi}) \right], \quad (132a)$$

$$\text{s.t.: } \boldsymbol{\pi}^\top \mathbf{1} \leq 1, \boldsymbol{\pi} \geq \mathbf{0}. \quad (132b)$$

where $\bar{\boldsymbol{\Sigma}}_\xi(\boldsymbol{\pi})$ solves the relaxed DARE:

$$\bar{\boldsymbol{\Sigma}}_\xi^{-1}(\boldsymbol{\pi}) = \frac{n}{\sigma^2} \mathbf{B}(\boldsymbol{\pi}; \hat{\boldsymbol{\beta}}_*) + (\rho^2 \bar{\boldsymbol{\Sigma}}_\xi(\boldsymbol{\pi}) + \sigma_w^2 \mathbf{I})^{-1}, \quad (133)$$

The map $\boldsymbol{\pi} \mapsto \bar{\boldsymbol{\Sigma}}_\xi(\boldsymbol{\pi})$ is defined implicitly through Eq. (133), precisely as in the linear case. The crucial observation is that, under the linearizability assumption, $\mathbf{B}(\boldsymbol{\pi}; \hat{\boldsymbol{\beta}}_*) = \sum_{i=1}^p \pi_i \mathbf{g}_i \mathbf{g}_i^\top$ has the

same algebraic structure as $B(\pi) = \sum_{i=1}^p \pi_i x_i x_i^\top$ in the linear setting. Thus Lemma 1 applies: the eigenvectors of $\bar{\Sigma}_\xi$ and $B(\pi; \hat{\beta}_*)$ are shared, and the eigenvalues of $\bar{\Sigma}_\xi$ are given by Eq. (15) as a function of the eigenvalues of $B(\pi; \hat{\beta}_*)$. This paves the way for establishing concavity:

Corollary 1 (Concavity of the Non-linear Steady-State Objective). *Under the linearizability assumption, the objective $G : \mathbb{R}^p \rightarrow \mathbb{R}$ defined in (132a) is concave. Consequently, Prob. (132) is a convex optimization problem.*

Proof. Under the linearizability assumption, $B(\pi; \hat{\beta}_*)$ is positive semi-definite and affine in π , with the same structure as in the linear case. Since Lemma 1 holds for any positive semi-definite matrix B , the eigenvalue characterization in Eq. (15) applies. The proof of Thm. 5 (App. H) then carries over verbatim: the map $g : \lambda_i \mapsto \lambda_i'^{-1}$ remains non-decreasing and concave, the function $\Phi(A) = \sum_{i=1}^{d_\beta} \log g(\lambda_i(A))$ is Schur-concave, and $G(\pi) = \Phi(B(\pi; \hat{\beta}_*))$ inherits concavity from the affinity of B in π . $\square$

As a result, under the linearizability assumption, the non-linear steady-state design problem follows the same steps as Thm. 5 in the linear setting, with the Jacobian vectors $g(x_i, \hat{\beta}_*)$ replacing the feature vectors $\{x_i\}$. Consequently, Prob. (132) is a convex optimization problem that can be solved using the optimization approach described in App. J, namely, the Frank-Wolfe algorithm, coupled with implicit differentiation to compute the gradient of the DARE solution w.r.t. the policy π .

L Experimental Setup

L.1 Datasets

We perform numerical evaluations over a synthetic dataset as well as over six publicly available datasets spanning environmental monitoring, finance, transportation, housing, and wireless networks, summarized in Table 1. For the real datasets, to enable re-sampling during the label collection stage, we reveal the ground truth label with additive zero-mean Gaussian noise with variance 10^{-1} . For the real datasets, in both the linear and non-linear settings, the system parameters (ρ, σ, σ_w) are estimated directly from data following the procedure described in App. L.3; the resulting values are summarized in Table 4. In all experiments across all datasets, we run every method 100 times, with a different seed each time, and report the mean and standard deviation of the metrics described in App. L.5.

L.1.1 Synthetic Dataset

We generate four synthetic datasets that are consistent with the dynamics in Eqs. (2)–(3), using different model functions h . We simulated four functions, namely a linear model, a ResNet model, a convolutional neural network (CNN), and a Multi-Layer Perceptron (MLP); the three non-linear architectures are described in App. L.2. Note that, in the linear case, $d_\beta = d$. In all cases, the set of candidate experiments $\mathcal{X}_k \subset \mathbb{R}^d$, i.e., the pool of feature vectors that the learner selects, varies with time. We generate these in an i.i.d. fashion per stage by setting $d = 30$, $p = 40$ and $n = 50$ and sample $x_{ik} \in \mathbb{R}^d$ in $\mathcal{X}_k$ from a normal distribution with zero mean and unit variance, i.e., $x_{ik} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. We generate the ground truth labels via (2) under a model β_k drifting according to Eq. (3). We set the parameters in the model dynamics (Eq. (3)) as $\rho = 0.9$, and the noise parameters to $\sigma = 0.9$ and $\sigma_w = 0.9$ (see Table 4). The prior mean is randomly sampled as $\beta_0 \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and the prior covariance is constructed as $\Sigma_0 = c\mathbf{I}$, where $c = 50$. The first stage model is sampled as $\beta_1 \sim \mathcal{N}(\beta_0, \Sigma_0)$. Moreover, the mean fixed-point is randomly sampled as $\beta_* \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. Beyond the aforementioned set of candidate feature vectors, we create another synthetic set of feature vectors $\mathbf{X}_k^{\text{test}} = \{x_{ik}\}_{i=1}^{n/4} \subset \mathbb{R}^d$, coupled with labels $\mathbf{y}^{\text{test}} \in \mathbb{R}^{n/4}$ generated via Eq. (2). We use this test set to measure the predictive performance of our estimates (see also Sec. L.5). We run these synthetic experiments for a total of $T = 2000$ stages.

L.1.2 Beijing PM2.5 [14].

This dataset consists of hourly air quality measurements collected at a single monitoring station (the US Embassy in Beijing, China) from January 1, 2010 through December 31, 2014, spanning



Figure 5: Illustration of the AT&T Downlink/Uplink dataset’s measurement route (source: [13]). The trace corresponds to a cross-country drive from Boston, MA to Los Angeles, CA in August 2022, spanning approximately 4,500 km and traversing 13 US states. As the vehicle traverses different regions, cellular network conditions change due to varying base-station density, available technology (4G LTE vs. 5G low/mid/mmWave bands), and terrain, inducing geographic concept drift in the feature–throughput relationship.

a total of 5 years (60 months). The regression task is to predict the PM2.5 concentration given the remaining air quality and meteorological measurements. Each sample corresponds to a single hourly measurement with $d = 12$ features that include date-time information (year, month, day, hour) and meteorological variables. The dataset contains 43824 samples in total, corresponding to roughly 8760 hourly observations per year over the 5-year span.

We pre-process the data by applying ordinal encoding to the categorical feature “combined wind direction”. We handled the missing PM2.5 values (4.72% of the whole data) by imputation using linear interpolation. We group together daily measurements as feature vectors, resulting in 24-hour long stages with 24 samples each. Having a total 43824 samples in the data, we end up with $T = 1826$ stages in this scenario. After utilizing the daily measurements, we exclude the date-time information from the data. For each stage, we randomly split 20% of the total measurements as test set $(\mathbf{X}_k^{\text{test}}, \mathbf{y}^{\text{test}})$, and the rest are used as training set: their respective features are utilized as the candidate feature vectors $\mathcal{X}_k$. Thus, the number of experiments becomes $p = 19$ at each stage.

L.1.3 S&P 500 [51].

This dataset contains daily historical stock-price records for all companies listed on the S&P 500 index of the US public equity market, covering a 5-year window from approximately February 2013 through February 2018. The regression task is to predict the next-day closing price given the current day’s trading features. Each sample corresponds to a single trading day for a single stock and includes the trading date, the stock’s ticker name, and $d = 6$ daily price/volume features: open price, high price, low price, close price, adjusted close price, and trading volume. The dataset contains 619040 stock-day records in total across the ~ 505 S&P 500 constituents, corresponding to roughly 1260 trading days per stock over the 5-year span. We pre-process the data by selecting a representative subset of stocks and a window of approximately 14 months (~ 62 trading weeks) from the original 5-year coverage, and computing daily returns and normalized volume. We group together weekly trading measurements, resulting in stages of approximately 5 trading days. Each stage pools all available stock measurements for that week. The selected window yields $T = 62$ stages. For each stage, we randomly split 20% of the measurements as a test set $(\mathbf{X}_k^{\text{test}}, \mathbf{y}^{\text{test}})$, and the rest are used as the training set: their respective features are utilized as the candidate feature vectors $\mathcal{X}_k$. The number of candidate experiments is $p = 16$ at each stage.

Table 4: Dynamics parameter estimates (ρ, σ, σ_w) fitted to each model architecture’s residuals on each dataset (see App. L.3). For each real dataset, the underlined value in the ResNet/CNN/MLP columns marks the median across the three non-linear architectures per parameter; this median is the shared estimate used for all non-linear models on that dataset in our experiments.

Dataset	Linear			ResNet			CNN			MLP		
	ρ	σ	σ_w	ρ	σ	σ_w	ρ	σ	σ_w	ρ	σ	σ_w
Synthetic	0.90	0.90	0.90	0.90	0.90	0.90	0.90	0.90	0.90	0.90	0.90	0.90
Beijing PM2.5 [14]	0.92	80.29	0.98	0.95	<u>134.48</u>	0.03	<u>0.93</u>	131.68	<u>0.02</u>	0.92	135.93	<u>0.02</u>
S&P 500 [51]	0.66	0.78	0.22	0.96	0.75	0.01	<u>0.91</u>	<u>0.78</u>	<u>0.02</u>	<u>0.91</u>	0.89	<u>0.02</u>
Bike Sharing [21]	0.93	156.98	29.08	<u>0.95</u>	263.28	<u>0.03</u>	0.93	261.16	<u>0.03</u>	0.96	<u>261.75</u>	0.02
California Housing [52]	0.94	0.79	0.01	0.95	1.19	<u>0.02</u>	<u>0.91</u>	1.22	0.01	<u>0.91</u>	<u>1.21</u>	<u>0.02</u>
AT&T Downlink [13]	0.53	183.58	684.94	<u>0.94</u>	<u>363.75</u>	<u>0.04</u>	0.92	364.47	<u>0.04</u>	<u>0.94</u>	360.51	0.03
AT&T Uplink [13]	0.74	13.32	644.75	<u>0.91</u>	<u>20.04</u>	0.02	0.94	20.56	0.05	0.90	19.57	<u>0.03</u>

L.1.4 Bike Sharing Demand [21].

This dataset contains hourly bike rental counts from the Capital Bikeshare system in Washington, D.C., USA, spanning two full years from January 1, 2011 through December 31, 2012 (24 months). The regression task is to predict the total number of bike rentals in a given hour given weather and calendar features. Each sample corresponds to a single hourly measurement with $d = 8$ features, including season, holiday indicator, working day indicator, weather situation, normalized temperature, normalized feeling temperature, normalized humidity, and normalized wind speed. The dataset contains 17379 hourly records in total, corresponding to roughly 8690 hourly observations per year over the 2-year span.

We pre-process the data by applying ordinal encoding to the categorical features (season, weather situation). We group together daily measurements as feature vectors, resulting in 24-hour long stages with 24 samples each. After grouping the 17379 hourly records, we obtain $T = 724$ stages. We exclude the date-time information from the feature set. For each stage, we randomly split 20% of the measurements as a test set $(\mathbf{X}_k^{\text{test}}, \mathbf{y}^{\text{test}})$, and the rest are used as the training set: their respective features are utilized as the candidate feature vectors $\mathcal{X}_k$. The number of candidate experiments is $p = 19$ at each stage.

L.1.5 California Housing [52].

This dataset contains aggregated demographic and housing statistics from the 1990 US Decennial Census, covering 20640 census block groups distributed across the entire state of California (all 58 counties). The regression task is to predict the median house value for a block group given the remaining census-derived features. Each sample corresponds to a single block group with $d = 6$ features: median income, housing median age, average number of rooms, average number of bedrooms, block-group population, and average household size. The raw dataset additionally provides each block group’s latitude and longitude, which support geographic ordering. The dataset contains 20640 samples in total, all collected from the 1990 decennial census, yielding an average of ~ 356 block-group observations per California county.

We induce a natural ordering by sorting block groups geographically (by latitude, north-to-south) and partitioning them into stages of 50 consecutive block groups each. This geographic ordering introduces the drift, as the relationship between housing features and price varies across regions (e.g., the income–price mapping differs substantially between the San Francisco Bay Area and the Central Valley). The dataset yields $T = 412$ stages. For each stage, we randomly split 20% of the block groups as a test set $(\mathbf{X}_k^{\text{test}}, \mathbf{y}^{\text{test}})$, and the rest are used as the training set: their respective features are utilized as the candidate feature vectors $\mathcal{X}_k$. The number of candidate experiments is $p = 40$ at each stage.

L.1.6 AT&T Downlink and Uplink [13].

This multi-carrier cellular measurement dataset was collected during a cross-country driving campaign from Boston, MA to Los Angeles, CA in August 2022, traversing approximately 4,500 km across 13 US states (Massachusetts, New York, Pennsylvania, Ohio, Indiana, Illinois, Iowa, Nebraska, Wyoming, Colorado, Utah, Nevada, and California) over a multi-day trip and passing through

Table 5: Methods

Method	Description
STS-AE (ours)	Stationary optimal selection, Adaptive estimation
STS-AE-K (ours)	Stationary optimal selection, Adaptive estimation, Known β_*
STG-AE (ours)	Stage-wise optimal selection, Adaptive estimation
STG-AE-K (ours)	Stage-wise optimal selection, Adaptive estimation, Known β_*
STG-SE	Stage-wise optimal selection, Static estimation
STG-SGD	Stage-wise optimal selection, SGD estimation
ADM-AE	Adam optimized selection, Adaptive estimation
ADM-AE-K	Adam optimized selection, Adaptive estimation, Known β_*
ADM-SE	Adam optimized selection, Static estimation
ADM-SGD	Adam optimized selection, SGD estimation
RND-AE	Random selection, Adaptive estimation
RND-AE-K	Random selection, Adaptive estimation, Known β_*
RND-SE	Random selection, Static estimation
RND-SGD	Random selection, SGD estimation
GKS-AE	Greedy Kalman information-gain selection, Adaptive estimation
GKS-AE-K	Greedy Kalman information-gain selection, Adaptive estimation, Known β_*
GKS-SE	Greedy Kalman information-gain selection, Static estimation
GKS-SGD	Greedy Kalman information-gain selection, SGD estimation
GKS-RLS	Greedy Kalman information-gain selection, Forgetting-factor RLS estimation
RND-RLS	Random selection, Forgetting-factor RLS estimation
ADM-RLS	Adam optimized selection, Forgetting-factor RLS estimation
STG-RLS	Stage-wise optimal selection, Forgetting-factor RLS estimation

major cities including Buffalo, Cleveland, Columbus, Indianapolis, Chicago, Des Moines, Omaha, Denver, Salt Lake City, and Las Vegas. The corresponding measurement route is illustrated in Fig. 5. Throughout the trip, three smartphones each connected to one of the three major US mobile operators (AT&T, Verizon, and T-Mobile) concurrently logged downlink/uplink throughput at 500 ms intervals together with per-sample signal-quality and cellular-technology indicators (RSRP, RSRQ, SINR, serving frequency band, and 4G/5G technology type). We extract two separate prediction tasks from the AT&T traces: predicting *downlink* throughput and predicting *uplink* throughput, yielding the AT&T Downlink and AT&T Uplink datasets, with approximately 6,700 and 6,560 raw measurement samples, respectively, collected over the course of the cross-country trip (roughly 500 samples per traversed state on average).

For the downlink task, each sample has $d = 12$ features comprising the cellular signal-quality and technology indicators listed above; for the uplink task, each sample has $d = 8$ features (the uplink direction does not expose a subset of the per-channel KPIs that are available on the downlink). We induce a stage structure by grouping measurements into road segments of approximately equal geographic length, with each segment constituting one stage; as the vehicle traverses regions with varying base station density, technology availability (4G LTE vs. 5G low/mid/mmWave bands), and terrain, the network conditions change substantially across stages, creating natural concept drift in the feature-throughput relationship. The downlink task spans $T = 335$ stages and the uplink task spans $T = 328$ stages. For each stage, we randomly split 20% of the measurements as a test set $(\mathbf{X}_k^{\text{test}}, \mathbf{y}^{\text{test}})$, and the rest are used as the training set: their respective features are utilized as the candidate feature vectors $\mathbf{X}_k$. The number of candidate experiments is $p = 16$ for both tasks at each stage.

L.2 Architectures

The three non-linear architectures are parameterized as follows. ResNet is a custom 1D residual network based on the ResNet architecture [27] with 5 convolutional layers, organized as an initial $1 \rightarrow 2$ -channel convolution of kernel size 3, followed by two 2-channel residual blocks each containing two convolutions of kernel sizes 3 and 1, an adaptive average pooling layer reducing the spatial dimension to 1, and a final linear scalar readout. The CNN consists of two 1D convolutional layers with output channels (3, 5) and kernel sizes (3, 5), each followed by a tanh activation, and uses the same adaptive-pool-plus-linear-head readout as the ResNet. Since both convolutional architectures expect a sequence input but our datasets provide d -dimensional feature vectors $\mathbf{x} \in \mathbb{R}^d$, we feed each vector to the network by reshaping it into a single-channel 1D signal of length d before the convolutions are applied. The MLP, in contrast, takes the feature vector directly and consists of three hidden fully-connected layers of width 2 each with tanh activations followed by a linear scalar output.

All three non-linear architectures use tanh activations and bias terms in every linear/convolutional layer; per-architecture parameter counts are reported in Table 1.

L.3 Parameter Estimation

Note that, in the real datasets, there is no ground truth dynamical model β_k . As we need ρ , σ^2 , and σ_w^2 in our algorithms, we measure these directly from data as follows. In the linear setting, we produce an estimate $\hat{\beta}_*$ of the mean fixed-point β_* by performing linear regression on the burn-in prefix only (the first $\lfloor 0.1 T \rfloor$ stages), and estimate a per-stage $\hat{\beta}_k$ through the linear regression using the p training samples. In the non-linear setting, we instead fit the non-linear model h on the same burn-in prefix via non-linear least squares:

$$\hat{\beta}_* = \arg \min_{\beta \in \mathbb{R}^d} \sum_{k=1}^K \|\mathbf{y}_k - h(\mathbf{X}_k, \beta)\|_2^2. \quad (134)$$

and estimate a per-stage $\hat{\beta}_k$ by fitting the non-linear model on each stage's p training samples:

$$\hat{\beta}_k = \arg \min_{\beta \in \mathbb{R}^d} \|\mathbf{y}_k - h(\mathbf{X}_k, \beta)\|_2^2, \quad k = 1, \dots, K. \quad (135)$$

Then, for both of the settings, considering our dynamics in Eq. (3), observe that residual errors satisfy

$$\mathbf{w}_k \equiv \beta_{k+1} - \rho \beta_k - (1 - \rho) \beta_* \quad (136)$$

Exploiting this relation, we first estimate value $\hat{\rho}$ by minimizing the average squared residual via:

$$\hat{\rho} = \arg \min_{\rho \in [0,1]} \left[\frac{1}{K-1} \sum_{k=1}^{K-1} \left\| \hat{\beta}_{k+1} - \rho \hat{\beta}_k - (1 - \rho) \hat{\beta}_* \right\|^2 \right], \quad (137)$$

where $K = \lfloor 0.1 T \rfloor$ is the number of burn-in stages (the first 10% of the T total stages), which are excluded from every reported metric; all estimates in this subsection are computed causally on this prefix, using no information from the evaluated stages. After obtaining $\hat{\rho}$, we estimate the drift noise variance σ_w^2 as the average of the mean square error of these estimated residuals:

$$\hat{\sigma}_w^2 = \frac{1}{(K-1)d\beta} \sum_{k=1}^{K-1} \|\hat{\mathbf{w}}_k\|^2, \quad \text{where} \quad \hat{\mathbf{w}}_k = \hat{\beta}_{k+1} - \hat{\rho} \hat{\beta}_k - (1 - \hat{\rho}) \hat{\beta}_*, \quad k \geq 1. \quad (138)$$

Finally, we estimate the measurement noise variance σ^2 as the average of the corresponding noise residuals:

$$\hat{\sigma}^2 = \frac{1}{(K-1)p} \sum_{k=1}^{K-1} \|\hat{\mathbf{e}}_k\|^2, \quad \text{where} \quad \mathbf{e}_k = \mathbf{y}_k - \mathcal{X}_k \hat{\beta}_k, \quad k \geq 1. \quad (139)$$

These parameters are then provided to the design selection algorithms we tested. Note that $(\hat{\rho}, \hat{\sigma}, \hat{\sigma}_w)$ characterize the pair (dataset, model class), and not the dataset alone. By Eq. (136), $\mathbf{w}_k$ is the drift of the *model parameters*, so $\hat{\sigma}_w$ lives in the parameter space of the architecture that is fitted. In addition, Eqs. (137)–(138) plug in the per-stage fits $\hat{\beta}_k$ in place of the unobserved β_k , so the estimation error of these fits is absorbed into $\hat{\mathbf{w}}_k$ along with the true drift. This is what produces the large $\hat{\sigma}_w$ that the linear model attains in Table 4 on the datasets where a linear fit is poor, namely 644.75 and 684.94 on the two AT&T tasks and 29.08 on Bike Sharing, against 0.02–0.05 for the non-linear architectures. Fitting a linear model to a strongly non-linear signal yields unstable per-stage coefficients, which inflate $\|\hat{\mathbf{w}}_k\|$. The same noisy $\hat{\beta}_k$ also act as regressors in Eq. (137), which in turn attenuates $\hat{\rho}$ towards zero, as observed on the AT&T tasks (0.74 and 0.53, against ≈ 0.9 for the non-linear architectures). These gaps thus reflect how well each model class fits the signal rather than a different underlying data drift.

L.4 Methods

We explored five design selection methods combined with four estimation methods, as described in Table 5. All of them were used on our datasets with the four model functions we discussed in the context of our synthetic dataset (see App. L.1.1): i.e., the linear, ResNet, CNN, and MLP models.

Table 6: Parameters for Sub-Synthetic Scenarios.

Param.	Sub-Syn. I	Sub-Syn. II	Description
p	40	15	# possible candidate $\mathbf{x}_i \in \mathbb{R}^d$
n	50	20	# selected $\mathbf{x}_i$ each stage
d	30	30	Dimension of $\mathbf{x}_i$
ρ	0.9	0.9	State transition constant
σ	0.9	0.9	Measurement noise variance
σ_w	0.9	0.9	Drift noise variance
T	2000	2000	Number of stages

L.4.1 Design Selection Methods

We consider the following design selection methods:

- **Stage-wise Optimal Selection (STG):** This selection method refers to our optimal stage-wise design policy in Sec. E.1.1, that is determined by solving the convex optimization Problem (64a). We solve this problem via the *cvxpy* convex optimization toolbox.
- **Stationary Optimal Selection (STS):** This selection method refers to our optimal steady-state design policies for linear and non-linear settings, described in Sections 3.3.2 and K.2.2. This is computed by solving Probs. (13) and (132). Contrary to the stage-wise optimization problem, computing the objective involves finding a solution to the DARE system of equations (Eqs. (12) and (133)). This prevents us from applying convex optimization solvers, like *cvxpy*, directly. Thus, we solve this problem with *Frank–Wolfe* algorithm by calculating the gradients of the objective functions in Eqs. (13a) and (132a), that involves finding a solution to the DARE systems of equations (Eqs. (12) and (133)), using the implicit function theorem [37]. Details of these computations are provided in App. J. Since the objective in Eq. (132a) depends on $\hat{\beta}_*$, in our experiments, we re-solve Prob. (132) at each stage to improve the selection as $\hat{\beta}_*$ is refined at each stage.
- **Greedy Kalman Information-Gain Selection (GKS):** This is an external design-selection baseline [61] that sequentially selects the n samples that maximize the one-step information gain. Starting from the *a priori* precision $\mathbf{M} = \Sigma_{\xi_{k|k-1}}^{-1}$, at each of the n picks it selects the candidate that maximizes the leverage score $\mathbf{x}_i^\top \mathbf{M}^{-1} \mathbf{x}_i$ (in the non-linear setting, the Jacobian row $\mathbf{g}(\mathbf{x}_i, \hat{\beta}_*)$ replaces $\mathbf{x}_i$), increments n_k at that index, and updates $\mathbf{M}^{-1}$ through a rank-one Sherman–Morrison update. Unlike our convex D-optimal designs (STG, STS), this greedy sensor-selection heuristic comes with no optimality guarantee, and thus isolates the benefit of the proposed design criterion.
- **Adam Selection (ADM):** This selection method performs projected gradient ascent to solve the Problem (9) using the Adam optimizer [34]. At each Adam iteration t , the method forms the gradient of the objective (9a) with respect to the i -th design variable as

$$g_k^{t,i} = \frac{n}{\sigma^2} \mathbf{x}_i^\top \mathbf{M}(\boldsymbol{\pi}_k^t)^{-1} \mathbf{x}_i, \quad (140)$$

where $\mathbf{M}(\boldsymbol{\pi}_k) = \frac{n_k}{\sigma^2} \mathbf{B}(\boldsymbol{\pi}_k) + \Sigma_{\xi_{k|k-1}}^{-1}$ from Eq. (9a). The iterate $\boldsymbol{\pi}_k^t$ is then updated using an Adam ascent step followed by projection onto the feasible simplex, and this procedure is repeated until convergence.

- **Random Selection (RND):** This selection method refers to random selection, that randomly selects n feature vectors from the candidates to construct the design matrix at each stage. The procedure is similar to the “balls-and-bins” method, by initializing the selection vector $\mathbf{n}_k \in \mathbb{N}^p$ with zero elements, selecting an index uniformly random, and incrementing the index by 1. After repeating this n times, the selection vector is determined.

L.4.2 Model Estimation

We also consider the following model estimation policies:

Table 7: Model error on synthetic data for Linear, ResNet, CNN, and MLP models. **Bold** indicates the best result per column and underline the second-best. Values are mean $\pm$ std over experiment seeds.

Method	Linear	ResNet	CNN	MLP
RND-SGD	12.43 $\pm$ 1.10	2.75 $\pm$ 0.45	3.06 $\pm$ 0.43	13.19 $\pm$ 2.01
RND-SE	14.20 $\pm$ 1.18	8.38 $\pm$ 1.87	9.91 $\pm$ 1.92	15.62 $\pm$ 2.61
RND-AE	8.18 $\pm$ 0.96	2.47 $\pm$ 0.42	2.47 $\pm$ 0.34	12.97 $\pm$ 2.22
RND-RLS	9.46 $\pm$ 0.98	3.36 $\pm$ 0.38	3.61 $\pm$ 0.38	13.25 $\pm$ 2.18
ADM-SGD	12.33 $\pm$ 1.07	2.58 $\pm$ 0.44	2.85 $\pm$ 0.36	12.41 $\pm$ 1.90
ADM-SE	11.45 $\pm$ 1.12	8.05 $\pm$ 1.72	9.75 $\pm$ 1.95	14.66 $\pm$ 2.38
ADM-AE	7.31 $\pm$ 0.87	2.43 $\pm$ 0.37	2.39 $\pm$ 0.33	12.11 $\pm$ 1.97
ADM-RLS	8.89 $\pm$ 0.81	2.98 $\pm$ 0.31	3.11 $\pm$ 0.35	12.38 $\pm$ 1.95
GKS-SGD	12.36 $\pm$ 1.17	2.51 $\pm$ 0.38	2.87 $\pm$ 0.37	12.94 $\pm$ 2.05
GKS-SE	11.18 $\pm$ 1.02	8.11 $\pm$ 1.68	9.58 $\pm$ 1.69	14.49 $\pm$ 2.45
GKS-AE	7.46 $\pm$ 0.81	2.41 $\pm$ 0.35	2.45 $\pm$ 0.34	12.08 $\pm$ 1.92
GKS-RLS	8.71 $\pm$ 0.91	3.01 $\pm$ 0.29	3.09 $\pm$ 0.31	12.55 $\pm$ 2.03
STG-SGD	11.36 $\pm$ 1.11	2.71 $\pm$ 0.41	2.88 $\pm$ 0.34	13.39 $\pm$ 2.03
STG-SE	10.47 $\pm$ 0.84	7.96 $\pm$ 1.67	9.62 $\pm$ 2.03	16.63 $\pm$ 2.90
STG-RLS	6.92 $\pm$ 0.61	2.92 $\pm$ 0.39	2.95 $\pm$ 0.34	11.96 $\pm$ 2.14
STG-AE (ours)	<u>6.32 $\pm$ 0.49</u>	<u>1.91 $\pm$ 0.37</u>	<u>2.03 $\pm$ 0.36</u>	<u>11.42 $\pm$ 2.31</u>
STS-AE (ours)	6.16 $\pm$ 0.47	1.82 $\pm$ 0.34	1.94 $\pm$ 0.38	11.13 $\pm$ 2.05

- **Adaptive Estimation (AE):** This estimation method refers to estimating full posterior via Kalman filter and EKF equations in the linear and non-linear settings in Eqs. (5)–(6) and (122a)–(122b) respectively, for both the model $\hat{\beta}_{k|k}$ and the covariance $\Sigma_{k|k}$ at each stage. In adaptive estimation, we assume by default that the mean fixed point β_* in Eq. (3) is unknown, and therefore use the estimator in Eq. (8); we also consider a sub-case in which the mean fixed point is assumed to be known, denoted by K.
- **Forgetting-factor Recursive Least Squares Estimation (RLS):** This is an external recursive-estimation baseline [43] that tracks the drifting model using a classic exponential forgetting factor $\gamma \in (0, 1]$. Its measurement update coincides with the (E)KF update in Eq. (5), but its time update models the parameters as a plain random walk with multiplicative covariance inflation,

$$\hat{\beta}_{k|k-1} = \hat{\beta}_{k-1|k-1}, \quad \Sigma_{\epsilon_{k|k-1}} = \frac{1}{\gamma} \Sigma_{\epsilon_{k-1|k-1}}, \quad (141)$$

rather than the mean-reverting dynamics of Eq. (3). It therefore discounts past observations geometrically and does not use the mean fixed-point β_* . We set $\gamma = 0.98$.

- **Static Estimation (SE):** This estimation method refers to estimating the model using priors [12] and the design matrix $\mathbf{X}_k$ along with the collected labels $\mathbf{y}_k$ at each stage, using Eq. (32).
- **Stochastic Gradient Descent Estimation (SGD):** This estimation method employs a standard stochastic gradient descent (SGD) update [59] for the model estimation. Given the design matrix $\mathbf{X}_k \in \mathbb{R}^{n \times d}$ and the observations $\mathbf{y}_k \in \mathbb{R}^n$ at stage k , the empirical squared loss is:

$$\mathcal{L}_k(\beta) = \frac{1}{2n} \|\mathbf{y}_k - \mathbf{X}_k \beta\|_2^2. \quad (142)$$

The method updates the model estimation at each stage via the SGD step:

$$\hat{\beta}_{k|k} = \hat{\beta}_{k-1|k-1} - \eta_k \nabla \mathcal{L}_k(\hat{\beta}_{k-1|k-1}), \quad (143)$$

until convergence, where the gradient is $\nabla \mathcal{L}_k(\hat{\beta}_{k-1|k-1}) = -\frac{1}{n} \mathbf{X}_k^\top (\mathbf{y}_k - \mathbf{X}_k \hat{\beta}_{k-1|k-1})$ and $\eta_k > 0$ is the learning rate. We run at most 200 epochs per stage, stopping early once the relative change in the loss (or the gradient norm) falls below 10^{-6} .




Figure 6: Estimation error between the estimation $\hat{\beta}_*$ and the true mean fixed-point β_* for the AE methods on the synthetic dataset, across four model architectures (rows, top to bottom: ResNet, CNN, MLP, Linear) and five AE methods (columns, left to right: STS-AE, STG-AE, ADM-AE, GKS-AE, RND-AE). Shaded regions represent ± 1 standard deviation from the mean.

L.5 Performance Metrics

We evaluate the quality of each method (i.e., a combined design selection and estimation process) using the following metrics. In order to display the per-stage metrics into a single number per experiment in Table 2, we take the average over the number of stages T in Table 1.

Test RMSE. In both synthetic and real data scenarios, over the test set, we measure the root mean square error (RMSE) between the collected test labels and predicted labels. In particular, at each stage, we collect labels for the test set as:

$$\mathbf{y}_k^{\text{test}} = f(\mathbf{X}_k^{\text{test}}, \beta_k) + \epsilon_k, \quad \epsilon_k \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}), \quad (144)$$

where $f(\mathbf{X}, \beta) = \mathbf{X}\beta$ in the linear setting and $f = h$ in the non-linear setting. Predictions are computed as $\mathbf{y}_k^{\text{pred}} = f(\mathbf{X}_k^{\text{test}}, \hat{\beta}_{k|k})$. Then, we measure the test RMSE as:

$$\begin{aligned} \text{Test RMSE}_k &= \sqrt{\frac{1}{n} \sum_{i=1}^n (y_{ik}^{\text{test}} - y_{ik}^{\text{pred}})^2} \\ &= \sqrt{\frac{1}{n} \|\mathbf{y}_k^{\text{test}} - \mathbf{y}_k^{\text{pred}}\|_2^2}. \end{aligned} \quad (145)$$

To report the average measurement per experiment, we discard the first τ stages as a warm-up (burn-in) period, where $\tau = \lfloor 0.1 T \rfloor$ corresponds to 10% of the total number of stages, and average the per-stage test RMSE over the remaining $T - \tau$ stages. Specifically, the mean test RMSE per experiment is:

$$\overline{\text{Test RMSE}} = \frac{1}{T - \tau} \sum_{k=\tau+1}^T \text{Test RMSE}_k. \quad (146)$$

Note that, as test labels are available in both synthetic and real data scenarios, we report this metric for all experiments.




Figure 7: Unknown vs. known β_* test RMSE comparisons across the five AE methods on the synthetic dataset; lower is better. Rows (top to bottom) correspond to the five AE methods: STS-AE, STG-AE, ADM-AE, GKS-AE, RND-AE. Columns (left to right) correspond to the four model architectures: Linear, ResNet, CNN, MLP. Shaded regions represent ± 1 standard deviation from the mean.

Model Error. In synthetic experiments, where the true model β_k is available, we measure the model error at each stage k with the Euclidean distance between the true model β_k and the estimator $\hat{\beta}_{k|k}$. In particular:

$$\text{Model Error}_k = \left\| \beta_k - \hat{\beta}_{k|k} \right\|_2. \quad (147)$$

In the non-linear setting, we instead evaluate the discrepancy through the model outputs over the test set:

$$\text{Model Error}_k = \left\| h(\mathbf{X}_k^{\text{test}}, \beta_k) - h(\mathbf{X}_k^{\text{test}}, \hat{\beta}_{k|k}) \right\|_2. \quad (148)$$

The mean model error per experiment is:

$$\overline{\text{Model Error}} = \frac{1}{T - \tau} \sum_{k=\tau+1}^T \text{Model Error}_k. \quad (149)$$

where T is the number of stages from Table 1. Note that we report this for synthetic experiments only, as we do not have such a true model β_k in the real data scenarios.




Figure 8: Unknown vs. known β_* model error comparisons across the five AE methods on the synthetic dataset; lower is better. Rows (top to bottom) correspond to the five AE methods: STS-AE, STG-AE, ADM-AE, GKS-AE, RND-AE. Columns (left to right) correspond to the four model architectures: Linear, ResNet, CNN, MLP. Shaded regions represent ± 1 standard deviation from the mean.

Objective Value. Lastly, for the linear model utilized in synthetic sub-scenarios in App. L.7, we use the objective value at each stage. From the maximization problem in Eq. (64a) the objective formula for the stage-wise optimal design is:

$$\text{Objective Value}_k = \log \det \left[\frac{n_k}{\sigma^2} \mathbf{B}(\boldsymbol{\pi}_k) + \boldsymbol{\Sigma}_{\boldsymbol{\xi}_{k|k-1}}^{-1} \right], \quad (150)$$

where $\boldsymbol{\pi}_k \in [0, 1]^{p_k}$ is the fractional policy at stage $k \in \mathbb{N}$, $\mathbf{B}(\boldsymbol{\pi}) = \sum_{i=1}^{p_k} \pi_i \mathbf{x}_i \mathbf{x}_i^\top$, and $\boldsymbol{\Sigma}_{\boldsymbol{\xi}_{k|k-1}}^{-1} \in \mathbb{R}^{d \times d}$ is the a priori covariance given by Eq. (61b). The mean objective value per experiment is:

$$\overline{\text{Objective Value}} = \frac{1}{T - \tau} \sum_{k=\tau+1}^T \text{Objective Value}_k. \quad (151)$$

Note that, for the optimal stationary policy (STS-AE), the discrete Riccati equation implies that, if we start with $\overline{\boldsymbol{\Sigma}}_{\boldsymbol{\xi}}$ (as we do), then $\boldsymbol{\Sigma}_{\boldsymbol{\xi}_{k|k-1}}^{-1} = \boldsymbol{\Sigma}_{\boldsymbol{\xi}_{k|k}}^{-1} = \overline{\boldsymbol{\Sigma}}_{\boldsymbol{\xi}}$. Hence, the objective remains constant in time under the stationary policy.

Table 8: Test RMSE on synthetic and real datasets, comparing each -AE method (unknown β_*) with its -K variant (known β_*) for the Linear, ResNet, CNN, and MLP models. **Bold** indicates the best result per model and underline the second-best. The -AE values are reproduced from Table 2. Mean $\pm$ std reported over 100 trials with different random seeds.

Model	Method	Beijing PM2.5	S&P 500	Bike Sharing	Cal. Housing	AT&T Downlink	AT&T Uplink	Synthetic
Linear	RND-AE	142.53 $\pm$ 0.50	2.71 $\pm$ 0.73	267.70 $\pm$ 1.00	3.72 $\pm$ 0.01	146.30 $\pm$ 1.69	20.77 $\pm$ 0.95	3.20 $\pm$ 0.05
	RND-AE-K	142.39 $\pm$ 0.52	2.68 $\pm$ 0.61	265.95 $\pm$ 1.07	3.66 $\pm$ 0.02	142.12 $\pm$ 1.59	20.64 $\pm$ 0.83	3.16 $\pm$ 0.06
	ADM-AE	141.46 $\pm$ 0.70	1.79 $\pm$ 0.57	265.40 $\pm$ 0.31	3.70 $\pm$ 0.02	142.34 $\pm$ 1.09	12.17 $\pm$ 0.58	2.74 $\pm$ 0.08
	ADM-AE-K	141.01 $\pm$ 0.46	1.74 $\pm$ 0.51	261.40 $\pm$ 0.38	3.48 $\pm$ 0.01	138.98 $\pm$ 0.92	11.19 $\pm$ 0.53	2.66 $\pm$ 0.06
	GKS-AE	141.65 $\pm$ 0.52	1.88 $\pm$ 0.55	263.01 $\pm$ 0.41	3.71 $\pm$ 0.02	143.51 $\pm$ 1.14	15.67 $\pm$ 0.93	2.79 $\pm$ 0.09
	GKS-AE-K	141.21 $\pm$ 0.53	1.85 $\pm$ 0.53	261.86 $\pm$ 0.40	3.62 $\pm$ 0.02	139.18 $\pm$ 1.08	14.88 $\pm$ 0.87	2.68 $\pm$ 0.08
	STG-AE (ours)	138.61 $\pm$ 0.60	1.68 $\pm$ 0.51	252.97 $\pm$ 0.39	3.66 $\pm$ 0.01	138.97 $\pm$ 0.85	11.10 $\pm$ 0.81	2.66 $\pm$ 0.08
	STG-AE-K (ours)	138.50 $\pm$ 0.41	1.65 $\pm$ 0.52	250.31 $\pm$ 0.34	3.32 $\pm$ 0.02	135.78 $\pm$ 0.90	10.88 $\pm$ 0.69	2.62 $\pm$ 0.07
	STS-AE (ours)	133.82 $\pm$ 0.63	1.75 $\pm$ 0.53	244.18 $\pm$ 0.31	3.68 $\pm$ 0.01	132.78 $\pm$ 0.90	10.51 $\pm$ 0.78	2.56 $\pm$ 0.06
	STS-AE-K (ours)	133.14 $\pm$ 0.55	1.71 $\pm$ 0.49	242.70 $\pm$ 0.25	<u>3.38 $\pm$ 0.01</u>	129.83 $\pm$ 0.87	10.39 $\pm$ 0.71	2.51 $\pm$ 0.04
ResNet	RND-AE	99.24 $\pm$ 0.39	0.76 $\pm$ 0.03	180.01 $\pm$ 0.35	1.32 $\pm$ 0.15	131.51 $\pm$ 0.33	7.66 $\pm$ 0.27	1.29 $\pm$ 0.10
	RND-AE-K	97.26 $\pm$ 0.46	0.71 $\pm$ 0.02	177.91 $\pm$ 0.32	1.26 $\pm$ 0.11	128.14 $\pm$ 0.35	7.51 $\pm$ 0.24	1.21 $\pm$ 0.09
	ADM-AE	96.31 $\pm$ 0.32	0.75 $\pm$ 0.02	177.88 $\pm$ 0.34	1.23 $\pm$ 0.12	127.44 $\pm$ 0.41	6.86 $\pm$ 0.23	1.29 $\pm$ 0.09
	ADM-AE-K	95.28 $\pm$ 0.36	0.69 $\pm$ 0.02	173.49 $\pm$ 0.31	1.21 $\pm$ 0.10	123.63 $\pm$ 0.35	6.68 $\pm$ 0.24	1.18 $\pm$ 0.09
	GKS-AE	94.89 $\pm$ 1.38	0.81 $\pm$ 0.05	179.93 $\pm$ 0.62	1.29 $\pm$ 0.19	129.45 $\pm$ 0.55	10.95 $\pm$ 0.86	1.11 $\pm$ 0.12
	GKS-AE-K	93.12 $\pm$ 1.31	0.78 $\pm$ 0.05	176.98 $\pm$ 0.59	1.21 $\pm$ 0.20	126.39 $\pm$ 0.51	10.24 $\pm$ 0.71	1.06 $\pm$ 0.11
	STG-AE (ours)	92.22 $\pm$ 0.38	0.67 $\pm$ 0.03	174.81 $\pm$ 0.27	1.24 $\pm$ 0.17	123.32 $\pm$ 0.29	6.74 $\pm$ 0.32	0.98 $\pm$ 0.07
	STG-AE-K (ours)	91.48 $\pm$ 0.20	0.63 $\pm$ 0.02	171.48 $\pm$ 0.28	<u>1.18 $\pm$ 0.18</u>	<u>121.61 $\pm$ 0.30</u>	6.63 $\pm$ 0.31	0.95 $\pm$ 0.06
	STS-AE (ours)	<u>87.28 $\pm$ 0.32</u>	<u>0.60 $\pm$ 0.04</u>	177.88 $\pm$ 0.33	1.19 $\pm$ 0.16	122.43 $\pm$ 0.35	<u>6.60 $\pm$ 0.34</u>	<u>0.89 $\pm$ 0.08</u>
	STS-AE-K (ours)	86.58 $\pm$ 0.31	0.57 $\pm$ 0.03	<u>172.63 $\pm$ 0.28</u>	1.15 $\pm$ 0.15	120.51 $\pm$ 0.32	6.56 $\pm$ 0.32	0.86 $\pm$ 0.05
CNN	RND-AE	96.32 $\pm$ 0.21	0.76 $\pm$ 0.02	180.92 $\pm$ 0.19	1.32 $\pm$ 0.10	131.46 $\pm$ 0.30	7.79 $\pm$ 0.27	1.39 $\pm$ 0.07
	RND-AE-K	94.36 $\pm$ 0.18	0.73 $\pm$ 0.01	178.96 $\pm$ 0.15	1.28 $\pm$ 0.09	128.89 $\pm$ 0.29	7.19 $\pm$ 0.29	1.32 $\pm$ 0.08
	ADM-AE	95.26 $\pm$ 0.24	0.74 $\pm$ 0.01	179.96 $\pm$ 0.16	1.25 $\pm$ 0.09	128.43 $\pm$ 0.25	6.76 $\pm$ 0.25	1.37 $\pm$ 0.07
	ADM-AE-K	93.88 $\pm$ 0.20	0.69 $\pm$ 0.01	175.91 $\pm$ 0.13	1.23 $\pm$ 0.08	126.94 $\pm$ 0.26	6.18 $\pm$ 0.23	1.32 $\pm$ 0.06
	GKS-AE	95.31 $\pm$ 0.22	0.73 $\pm$ 0.03	178.92 $\pm$ 0.18	1.26 $\pm$ 0.15	128.45 $\pm$ 0.24	6.89 $\pm$ 0.28	1.38 $\pm$ 0.08
	GKS-AE-K	94.11 $\pm$ 0.20	0.69 $\pm$ 0.02	175.58 $\pm$ 0.15	1.23 $\pm$ 0.11	125.84 $\pm$ 0.21	6.02 $\pm$ 0.25	1.31 $\pm$ 0.07
	STG-AE (ours)	93.23 $\pm$ 0.17	0.66 $\pm$ 0.04	177.93 $\pm$ 0.15	1.20 $\pm$ 0.12	123.47 $\pm$ 0.26	5.86 $\pm$ 0.22	1.39 $\pm$ 0.12
	STG-AE-K (ours)	91.42 $\pm$ 0.18	0.61 $\pm$ 0.02	174.85 $\pm$ 0.11	1.15 $\pm$ 0.08	<u>120.92 $\pm$ 0.32</u>	5.80 $\pm$ 0.21	1.31 $\pm$ 0.07
	STS-AE (ours)	88.23 $\pm$ 0.13	0.62 $\pm$ 0.04	173.93 $\pm$ 0.17	1.24 $\pm$ 0.10	121.39 $\pm$ 0.27	5.78 $\pm$ 0.26	1.35 $\pm$ 0.10
	STS-AE-K (ours)	86.48 $\pm$ 0.13	0.59 $\pm$ 0.01	170.74 $\pm$ 0.31	<u>1.19 $\pm$ 0.09</u>	119.47 $\pm$ 0.28	5.71 $\pm$ 0.24	1.30 $\pm$ 0.06
MLP	RND-AE	95.70 $\pm$ 0.65	0.78 $\pm$ 0.03	180.80 $\pm$ 0.26	1.35 $\pm$ 0.12	128.36 $\pm$ 0.35	9.74 $\pm$ 0.34	3.86 $\pm$ 0.62
	RND-AE-K	93.49 $\pm$ 0.66	0.75 $\pm$ 0.02	177.83 $\pm$ 0.29	1.31 $\pm$ 0.11	126.65 $\pm$ 0.34	9.18 $\pm$ 0.32	3.80 $\pm$ 0.60
	ADM-AE	93.65 $\pm$ 0.76	0.75 $\pm$ 0.02	179.99 $\pm$ 0.32	1.34 $\pm$ 0.11	127.46 $\pm$ 0.31	7.95 $\pm$ 0.41	3.61 $\pm$ 0.55
	ADM-AE-K	92.85 $\pm$ 0.73	0.71 $\pm$ 0.02	176.74 $\pm$ 0.30	1.28 $\pm$ 0.13	125.19 $\pm$ 0.28	7.91 $\pm$ 0.38	3.59 $\pm$ 0.53
	GKS-AE	93.88 $\pm$ 0.89	0.81 $\pm$ 0.08	179.85 $\pm$ 0.30	1.34 $\pm$ 0.11	128.10 $\pm$ 0.35	8.46 $\pm$ 0.35	3.59 $\pm$ 0.54
	GKS-AE-K	91.74 $\pm$ 0.85	0.77 $\pm$ 0.08	175.94 $\pm$ 0.26	1.29 $\pm$ 0.10	126.85 $\pm$ 0.31	8.31 $\pm$ 0.34	3.51 $\pm$ 0.52
	STG-AE (ours)	91.65 $\pm$ 0.86	0.69 $\pm$ 0.02	177.97 $\pm$ 0.32	1.28 $\pm$ 0.12	124.29 $\pm$ 0.34	8.21 $\pm$ 0.33	2.98 $\pm$ 0.65
	STG-AE-K (ours)	90.88 $\pm$ 0.81	0.65 $\pm$ 0.01	<u>174.48 $\pm$ 0.28</u>	1.24 $\pm$ 0.09	122.67 $\pm$ 0.32	8.09 $\pm$ 0.32	2.94 $\pm$ 0.63
	STS-AE (ours)	89.70 $\pm$ 0.68	0.72 $\pm$ 0.03	174.99 $\pm$ 0.33	1.31 $\pm$ 0.16	121.45 $\pm$ 0.29	7.90 $\pm$ 0.29	2.63 $\pm$ 0.57
	STS-AE-K (ours)	88.37 $\pm$ 0.69	<u>0.68 $\pm$ 0.01</u>	172.71 $\pm$ 0.28	<u>1.27 $\pm$ 0.11</u>	118.92 $\pm$ 0.28	7.84 $\pm$ 0.29	2.56 $\pm$ 0.53

L.6 Computing Infrastructure

All computations were performed on a high-performance computing cluster managed by SLURM. The experiments were run on CPU nodes equipped with dual-socket Intel Xeon E5-2680 v2 processors (40 logical cores, 2.80 GHz) and 100 GB of allocated memory per job. Jobs were submitted via *sbatch* using single-core execution.

Table 9 reports the average per-stage wall-clock time (in seconds) of each design-selection method on this infrastructure, averaged over all stages, estimation variants, and seeds, and reported separately for each dataset and model architecture. The reported times are dominated by the design-selection step; the model-estimation overhead (AE, RLS, SE, SGD) is comparable across selection methods. These times are modest, and since each stage corresponds to a batch of newly collected data, they sit comfortably within a setting where retraining happens periodically.

L.7 Additional Synthetic Sub-Scenarios (I&II)

In addition to the synthetic experiment in the experimental setup in Sec. 4.1, we consider two synthetic sub-scenarios for the linear setting, (Sub-Scenarios I and II) summarized in Table 6, in which we simulate the same dynamics in Eq. (3). In these sub-scenarios, we simulate the linear model in Eq. (4) to analyze the impacts of different setups on our setting.

In particular, in the sub-scenario I, we set the parameters the same as the main synthetic dataset in Table 1 ($d = 30$, $p = 40$, $n = 50$), but the set of candidate experiments $\mathcal{X} = \{\mathbf{x}_i\}_{i=1}^p \subset \mathbb{R}^d$ is fixed through the stages. In sub-scenario II, we set $d = 30$, $p = 15$, and $n = 20$; note that, in this case, $p < d$.

Table 9: Average per-stage wall-clock time (seconds) of each design-selection method, reported per dataset and model architecture, averaged over estimation variants, stages, and 100 seeds.

Model	Method	Beijing PM2.5	S&P 500	Bike Sharing	Cal. Housing	AT&T Downlink	AT&T Uplink	Synthetic
Linear	STG	0.060 $\pm$ 0.006	0.042 $\pm$ 0.004	0.064 $\pm$ 0.005	0.092 $\pm$ 0.005	0.083 $\pm$ 0.006	0.069 $\pm$ 0.010	0.689 $\pm$ 0.021
	STS	0.002 $\pm$ 0.001	0.004 $\pm$ 0.001	0.002 $\pm$ 0.001	0.005 $\pm$ 0.001	0.013 $\pm$ 0.003	0.007 $\pm$ 0.002	0.023 $\pm$ 0.005
	GKS	0.003 $\pm$ 0.003	0.002 $\pm$ 0.001	0.003 $\pm$ 0.003	0.004 $\pm$ 0.003	0.002 $\pm$ 0.003	0.002 $\pm$ 0.003	0.009 $\pm$ 0.004
	ADM	0.080 $\pm$ 0.009	0.069 $\pm$ 0.009	0.078 $\pm$ 0.009	0.080 $\pm$ 0.009	0.059 $\pm$ 0.029	0.044 $\pm$ 0.037	0.117 $\pm$ 0.007
	RND	0.002 $\pm$ 0.003	0.001 $\pm$ 0.002	0.001 $\pm$ 0.002	0.002 $\pm$ 0.003	0.001 $\pm$ 0.003	0.001 $\pm$ 0.003	0.002 $\pm$ 0.001
ResNet	STG	1.280 $\pm$ 0.063	1.320 $\pm$ 0.095	1.290 $\pm$ 0.140	1.590 $\pm$ 0.143	1.290 $\pm$ 0.108	1.280 $\pm$ 0.117	4.660 $\pm$ 5.500
	STS	2.960 $\pm$ 3.160	1.670 $\pm$ 2.170	0.516 $\pm$ 0.349	8.700 $\pm$ 6.590	1.890 $\pm$ 1.500	2.290 $\pm$ 1.170	8.720 $\pm$ 7.800
	GKS	0.130 $\pm$ 0.057	0.119 $\pm$ 0.068	0.125 $\pm$ 0.057	0.154 $\pm$ 0.085	0.153 $\pm$ 0.092	0.130 $\pm$ 0.068	0.306 $\pm$ 0.252
	ADM	0.242 $\pm$ 0.065	0.130 $\pm$ 0.074	0.222 $\pm$ 0.074	0.257 $\pm$ 0.088	0.232 $\pm$ 0.069	0.180 $\pm$ 0.081	0.454 $\pm$ 0.254
	RND	0.119 $\pm$ 0.057	0.115 $\pm$ 0.062	0.115 $\pm$ 0.055	0.132 $\pm$ 0.080	0.120 $\pm$ 0.060	0.117 $\pm$ 0.059	0.280 $\pm$ 0.246
CNN	STG	0.771 $\pm$ 0.082	0.758 $\pm$ 0.087	0.778 $\pm$ 0.115	0.971 $\pm$ 0.117	0.772 $\pm$ 0.091	0.785 $\pm$ 0.119	3.090 $\pm$ 6.330
	STS	2.280 $\pm$ 2.420	1.820 $\pm$ 2.470	0.554 $\pm$ 0.730	6.880 $\pm$ 4.360	1.750 $\pm$ 1.430	2.200 $\pm$ 1.170	7.370 $\pm$ 4.800
	GKS	0.113 $\pm$ 0.069	0.104 $\pm$ 0.082	0.128 $\pm$ 0.098	0.126 $\pm$ 0.082	0.117 $\pm$ 0.079	0.112 $\pm$ 0.073	0.175 $\pm$ 0.132
	ADM	0.205 $\pm$ 0.051	0.109 $\pm$ 0.077	0.195 $\pm$ 0.051	0.222 $\pm$ 0.080	0.201 $\pm$ 0.059	0.173 $\pm$ 0.076	0.315 $\pm$ 0.129
	RND	0.093 $\pm$ 0.049	0.091 $\pm$ 0.068	0.089 $\pm$ 0.046	0.105 $\pm$ 0.072	0.090 $\pm$ 0.052	0.095 $\pm$ 0.065	0.162 $\pm$ 0.131
MLP	STG	0.563 $\pm$ 0.037	0.489 $\pm$ 0.050	0.702 $\pm$ 0.046	0.551 $\pm$ 0.050	1.340 $\pm$ 0.100	0.659 $\pm$ 0.049	3.620 $\pm$ 7.230
	STS	0.320 $\pm$ 0.152	0.240 $\pm$ 0.040	0.434 $\pm$ 0.460	1.060 $\pm$ 0.718	0.360 $\pm$ 0.110	0.312 $\pm$ 0.122	7.500 $\pm$ 8.713
	GKS	0.069 $\pm$ 0.034	0.067 $\pm$ 0.054	0.071 $\pm$ 0.034	0.080 $\pm$ 0.043	0.069 $\pm$ 0.036	0.069 $\pm$ 0.040	0.101 $\pm$ 0.063
	ADM	0.070 $\pm$ 0.034	0.065 $\pm$ 0.045	0.127 $\pm$ 0.063	0.090 $\pm$ 0.048	0.073 $\pm$ 0.036	0.073 $\pm$ 0.037	0.210 $\pm$ 0.070
	RND	0.064 $\pm$ 0.035	0.060 $\pm$ 0.043	0.062 $\pm$ 0.033	0.068 $\pm$ 0.039	0.061 $\pm$ 0.034	0.064 $\pm$ 0.036	0.085 $\pm$ 0.062

We sample $\mathbf{x}_{ik} \in \mathbb{R}^d$ in $\mathcal{X}_k$ from a normal distribution with zero mean and unit variance, i.e., $\mathbf{x}_{ik} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$; that are sampled independently at each stage. Beyond the aforementioned set of candidate feature vectors, we create another synthetic set of feature vectors $\mathbf{X}_k^{\text{test}} = \{\mathbf{x}_{ik}\}_{i=1}^{n/4} \subset \mathbb{R}^d$, coupled with labels $\mathbf{y}^{\text{test}} \in \mathbb{R}^{n/4}$ generated via Eq. (2) with the ground truth model β_k . We use this test set to measure the predictive performance of our estimates (see also the performance metrics in Sec. 4.1). In both synthetic scenarios, we set the parameters in the model dynamics (Eq. (3)) as $\rho = 0.9$, and the noise parameters to $\sigma = 0.9$ and $\sigma_w = 0.9$ (see Table 6). The prior mean is randomly sampled as $\beta_0 \sim \mathcal{N}(\mathbf{0}, \mathbf{1})$ and the prior covariance is constructed as $\Sigma_0 = c\mathbf{I}$, where $c = 50$. The first stage model is sampled as $\beta_1 \sim \mathcal{N}(\beta_0, \Sigma_0)$. Moreover, the mean fixed-point is randomly sampled as $\beta_* \sim \mathcal{N}(\mathbf{0}, \mathbf{1})$. We run each scenario for a total of $T = 2000$ stages; we repeat this 100 times with different random seeds to compute confidence intervals.

We note that we *purposefully design the sub-scenario II* to pose a significant challenge for inference, as the number of candidate features is smaller than the dimensions (i.e. $p < d$). This makes estimation and tracking of changes due to drift even more challenging. Moreover, conditions of Thm. 2 are violated; to address this, we regularize the linear estimates in Eq. (8) by adding a small quadratic penalty (i.e., performing ridge regression); this modified estimator still approximates β_* quite well in practice, even if the estimator does not converge (see App. M.2.4), learning missing directions as they are revealed. In contrast, in Sub-Scenario I, the conditions of Thm. 2 are always satisfied, as the design matrix is always rank d .

M Additional Experimental Results

M.1 Model Error Across Architectures for the Synthetic Scenarios

Recall that, in the synthetic scenarios, we have access to the ground truth models. Table 7 reports the model error on synthetic data across the four model architectures (Linear, ResNet, CNN, MLP) over the synthetic dataset. Our proposed methods STS-AE and STG-AE achieve the lowest error under every architecture, with STS-AE yielding the best overall result on CNN (1.94) and ResNet(1.82). The gap between adaptive (-AE) and the (-SE), (-SGD), and (-RLS) estimation baselines is substantial regardless of the selection strategy, underscoring the importance of Kalman-based parameter tracking when the model drifts. On the selection side, the benefit of optimal design over greedy, random or Adam-based selection is present across all architectures but is especially pronounced for the linear model, where STS-AE (6.16) halves the error of RND-SGD (12.43). We note that it is the margin of the external greedy and the Adam-based baselines over random selection that narrows on the non-linear architectures, rather than ours. On ResNet and CNN, GKS-AE (2.41, 2.45) and ADM-AE (2.43, 2.39) improve upon RND-AE (2.47, 2.47) by at most a few percent, while on the linear model they improve by 9% and 11%, respectively; STS-AE, by contrast, stays 19–26% below all three on




Figure 9: Test RMSE (Eq. (145)) and model error (Eqs. (147)–(148)) comparisons with respect to different values of ρ , σ , and σ_w on the synthetic dataset, for the CNN (rows 1–2), MLP (rows 3–4), and linear model (rows 5–6). Within each pair of rows, the upper row reports test RMSE and the lower row reports model error. Other parameters as in Table 4. Lower is better. Shaded regions represent ± 1 standard deviation from the mean.

both architectures. This is consistent with the fact that neither baseline solves Prob. (9) to optimality. The greedy rule commits the budget to the few candidates of highest leverage, and projected Adam ascent converges to a point that need not be optimal, and the resulting loss is architecture-dependent, whereas our convex solvers attain the optimum by Thms. 3 and 5. Among nonlinear architectures, CNN and ResNet consistently track the true parameters more accurately than MLP, likely due to their structural inductive biases (convolutions and skip connections) producing a more favorable loss landscape for the EKF updates.




Figure 10: Performance w.r.t. model error, test RMSE, and objective value for synthetic sub-scenario I as a function of the stage k . Shaded regions represent ± 1 standard deviation from the mean. We observe that our methods STG-AE and STS-AE achieve lower model error and test RMSE, as well as higher objective value compared to other methods.

M.2 Additional Synthetic Scenarios

M.2.1 Convergence of $\hat{\beta}_*$ Across Architectures and Impact of Estimation

We plot the convergence of the mean fixed-point estimator evaluated in (8) and (126) for linear and non-linear model functions, respectively, in Fig. 6 for the synthetic scenario. We report the estimation difference $\|\hat{\beta}_* - \beta_*\|_2$ between the estimation $\hat{\beta}_*$ and the true mean fixed-point β_* for STG-AE, STS-AE, ADM-AE, GKS-AE, and RND-AE described in Table 5 in Fig. 6: note that estimates are somewhat affected by the selection process, as the latter determines X_k , this is why we report these across different methods, even though (8) and (126) stay the same. Similarly, for the non-linear setting, we report $\|h(X_k^{\text{test}}, \beta_*) - h(X_k^{\text{test}}, \hat{\beta}_*)\|_2$ for the MLP, CNN, and ResNet models in Fig. 6. For all methods, we observe that as k increases, the estimation difference decreases; in particular, our estimator $\hat{\beta}_*$ converges to the true mean-fixed point β_* .

To further assess the estimation quality of β_* , in Table 8 and Figs. 7–8, we compare the AE methods with their K variants, in which the mean fixed-point β_* is known and used during the model estimation process. For all methods, we observe that the AE methods exhibit very similar performance as their K variants, for all metrics. Across the four model architectures and seven datasets, the gap between each AE method and its K variant is uniformly small, typically within a few percent of the test RMSE, indicating that our estimator $\hat{\beta}_*$ is sufficiently close to the true β_* that the downstream design and estimation policies behave nearly identically in both cases. In most cells, the best result per model is a K variant (typically STS-AE-K, with STG-AE-K winning the remaining cases), and the corresponding AE counterpart is the second-best, indicating that our methods sit immediately below their K variants without any baseline able to bridge the gap. The relative ordering across methods is also preserved between the two settings: STS-based methods outperform STG-based ones, which in turn outperform GKS, ADM and RND, irrespective of whether β_* is known or estimated. Since the K variants act as an oracle upper bound, the closeness of the AE methods to this bound corroborates the empirical convergence of $\hat{\beta}_*$ shown in Fig. 6, and indicates that our estimation procedure does not meaningfully limit performance in practice.

M.2.2 Impact of System Parameters

In Fig. 9, for CNN, MLP, and linear models, we report the test RMSE (Eq. (145)) and the model error (Eqs. (147)–(148)) under the impact of varying system parameters ρ , σ , and σ_w for each of them while keeping the others fixed at their default values (see Table 4). Similar to the results for ResNet model in Fig. 2, our proposed methods achieve lower model error compared to the baselines.




Figure 11: Estimation error between the estimation $\hat{\beta}_*$ and the true mean fixed-point β_* for the AE methods on the synthetic dataset. Rows (top to bottom) correspond to the two synthetic sub-scenarios: Sub-scenario I and Sub-scenario II. Columns (left to right) correspond to the five AE methods: STS-AE, STG-AE, ADM-AE, GKS-AE, RND-AE. Shaded regions represent ± 1 standard deviation from the mean.

M.2.3 Results for Sub-Scenario I

In Fig. 10, we plot the three metrics (model error, test RMSE, and objective value) for the STS, STG, GKS, ADM, and RND design selection methods coupled with the different model estimation policies, as described in Table 5. To avoid clutter, we skip the K variants of these algorithms and report their performance below (see Fig. 12). We observe that our methods STS-AE and STG-AE outperform other methods both in model error and test RMSE. Overall, as the total number of possible candidate features p is higher than the feature dimension d in this setting (see Tab. 6), selection methods play a more important role. In particular, we observe that STG-SE and ADM-SE outperform RND-AE in all three metrics. Overall, adaptive methods (AE) have better results compared to static (SE), SGD, and forgetting-factor RLS methods, highlighting the importance of estimating full posterior via Kalman filter equations in Eqs. (5)–(6) for both the model $\hat{\beta}_{k|k}$ and the covariance $\Sigma_{k|k}$ at each stage. Comparing the design selection methods within the same model estimation policies, conducting an optimal design policy (STS and STG) yields a lower error compared to GKS, ADM and RND selections. Moreover, in Fig. 10c, stage-wise and steady state optimal design selection methods tend to have a higher objective value in comparison due to the same reason above, as they find the maximum value of the function (150) optimally. Additionally, the adaptive methods yield a higher objective value at each stage compared to the static and SGD estimation methods. Figs. 10b and 10c also let us verify empirically the link between our design objective and the metric we report, discussed in Sec. 4.1. At every stage, the ranking of the methods by objective value in Fig. 10c is the reverse of their ranking by test RMSE in Fig. 10b, and we observe no stage at which one method attains both a higher objective value and a higher test RMSE than another. We stress that the objective value of a method is computed from its own posterior covariance estimate, so this agreement is expected only among methods that estimate the covariance consistently. A method whose covariance is badly mis-estimated can report a high objective while predicting poorly, which is why we report the test RMSE rather than the objective as our headline metric in Table 2.

We provide additional results comparing the -K variants of the above methods; we also provide numerical evidence of the convergence of estimator $\hat{\beta}_*$. In Fig. 11, we report the estimation difference $\|\hat{\beta}_* - \beta_*\|_2$ between the estimation $\hat{\beta}_*$ and the true mean fixed-point β_* , when designs are selected via STS-AE, STG-AE, ADM-AE, GKS-AE, and RND-AE for Sub-Scenario I. For all methods, we observe that as k increases, the estimation difference decreases rapidly; in particular, our estimator $\hat{\beta}_*$ converges to the true mean-fixed point β_* .

In Fig. 12, we compare AE methods with their K variants, where the real value of the mean fixed-point β_* in Eq. (3) is assumed to be known. We report the model error (Eq. (147)), the test RMSE (Eq. (145)), and the objective value (Eq. (150)) in Figs. 12a, 12b, and 12c for the STS-AE method, 12d, 12e, and 12f for the STG-AE method, 12g, 12h, and 12i for the ADM-AE method, 12j, 12k, and




Figure 12: Unknown vs. known β_* comparisons across the five AE methods in synthetic sub-scenario I. Rows (top to bottom) correspond to the five AE methods: STS-AE, STG-AE, ADM-AE, GKS-AE, RND-AE. Columns (left to right) correspond to: model error (Eq. (147)) between the estimation $\hat{\beta}_{k|k}$ and the true model β_k , lower is better; test RMSE (Eq. (145)) between predicted labels y_k^{pred} and test labels y_k^{test} , lower is better; objective value (Eq. (150)), higher is better (maximization problem). Shaded regions represent ± 1 standard deviation from the mean.

12l for the GKS-AE method, and 12m, 12n, and 12o for the RND-AE method respectively. For all methods, we observe that the AE methods exhibit very similar results as their K variants for all metrics, and that the estimator indeed approximates β_* quite well; thus the Euclidean difference between the estimation and the parameter itself converges to 0. Therefore, we justify that for the adaptive methods, we do not have the restriction of the assumption of knowing the true mean fixed-point β_* .




Figure 13: Unknown vs. known β_* comparisons across the five AE methods in synthetic sub-scenario II. Rows (top to bottom) correspond to the five AE methods: STS-AE, STG-AE, ADM-AE, GKS-AE, RND-AE. Columns (left to right) correspond to: model error (Eq. (147)) between the estimation $\hat{\beta}_{k|k}$ and the true model β_k , lower is better; test RMSE (Eq. (145)) between predicted labels $\mathbf{y}_k^{\text{pred}}$ and test labels $\mathbf{y}_k^{\text{test}}$, lower is better; objective value (Eq. (150)), higher is better (maximization problem). Shaded regions represent ± 1 standard deviation from the mean.

M.2.4 Results for Sub-Scenario II

In Fig. 14 we plot the three metrics (model error, test RMSE, and objective value) for the STS, STG, GKS, ADM, and RND design selection methods coupled with the different model estimation policies, as described in Table 5. To avoid clutter, we again skip the K variants of these algorithms and report their performance below (see Fig. 13). We observe that our methods STS-AE and STG-AE outperform other methods both in model error and test RMSE. Similar behavior as in Scenario I in terms of the superiority of adaptive (-AE) methods is observed in this scenario as well. Our selection methods STS and STG yield a lower error compared to GKS, ADM and RND. Moreover, in Fig.




Figure 14: Performance w.r.t. model error, test RMSE, and objective value for synthetic sub-scenario II as a function of the stage k . Shaded regions represent ± 1 standard deviation from the mean.

14c, adaptive methods tend to have a higher objective value in comparison due to the Kalman filter adaptations of the covariance $\Sigma_{\xi_{k|k}}$ at each stage, where the STS and STG design selections find the maximum value of the function (150) optimally, yielding a higher objective value at each stage compared to GKS, ADM and RND design selections. As observed in Sub-Scenario I, the objective value and the test RMSE remain consistent with each other in this scenario as well. The ranking of the methods by objective value in Fig. 14c is again the reverse of their ranking by test RMSE in Fig. 14b at every stage, so a design that attains a higher objective also yields a lower test error.

We again provide additional results comparing the -K variants, as well as illustrating the convergence of $\hat{\beta}_*$. Similar to the Sub-Scenario I, we plot the convergence of the mean fixed-point estimator evaluated in (8) for AE methods in Fig. 11 for Sub-Scenario II. We report the estimation difference $\|\hat{\beta}_* - \beta_*\|_2$ between the estimation $\hat{\beta}_*$ and the true mean fixed-point β_* for STS-AE, STG-AE, ADM-AE, GKS-AE, and RND-AE described in Table 5 in Figs. 11f, 11g, 11h, 11i, and 11j respectively. For all methods, we observe that as k increases, the estimation difference decreases; in particular, our estimator $\hat{\beta}_*$ converges to the true mean-fixed point β_* .

In Fig. 13, we compare AE methods with their K variants. We again observe that for all metrics, AE methods closely track their K variants, and that the estimator approximates β_* .

N Limitations

Our framework rests on a specific dynamics model in which the parameter β_k undergoes Gaussian oscillations around a mean fixed point β_* (Eq. (3)). Real-world parameter drift can exhibit richer behavior (level shifts, regime changes, or non-Gaussian innovations) that this model does not capture. The choice is motivated by two considerations: Gaussian oscillations are a standard default when the underlying dynamics are unknown; moreover, deployed model parameters in practice neither diverge nor collapse but rather fluctuate around an operating point, which this model captures compactly. Our modeled dynamics also come with important advantages. To begin with, the resulting DARE characterizing the steady-state covariance (Eq. (10)) admits a closed-form solution, which is what makes Thm. 4 and the steady-state policies analytically tractable. Furthermore, our ability to extend our analysis to EKF also exploits (Eq. (3)), as we can linearize the model (i.e., compute the Jacobian) around the mean fixed-point. Extending the framework to richer dynamics, such as regime-switching or heavy-tailed innovations, is left for future work.

Another limitation concerns the non-linear setting. The EKF linearization in Eq. (16) renders the underlying least-squares estimation non-convex, and consequently the mean fixed point β_* is not guaranteed to be identifiable in closed form. We approximate it via a “status-quo” estimate that minimizes the square-error loss while ignoring dynamics; this is an approximation rather than an exact recovery of β_* . Empirically, the approximation is sufficient: the convergence plots in Fig. 6 show that $\hat{\beta}_*$ tracks the true mean fixed point closely across all four architectures, the Table 8 and Figs. 7–8 show that the methods with unknown mean fixed point β_* exhibit very similar performance as their variants in which β_* is assumed to be known, and the test RMSE results in Table 2 indicate

that this is enough to obtain consistently strong downstream performance. Nevertheless, extending guarantees, e.g., when the mean fixed-point is identifiable, is an interesting open direction.