This article has been accepted for publication in IEEE Transactions on Signal Processing. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TSP.2022.3187655

Active Sampling of Multiple Sources for Sequential Estimation

Arpan Mukherjee, *Student Member, IEEE*, Ali Tajer, *Senior Member, IEEE*, Pin-Yu Chen, Payel Das *Member, IEEE*

***Abstract*—Consider K processes, each generating a sequence of identical and independent random variables. The probability measures of these processes have random parameters that must be estimated. Specifically, they share a parameter θ common to all probability measures. Additionally, each process $i \in \{1, \ldots, K\}$ has a private parameter α_i. The objective is to design an active sampling algorithm for sequentially estimating these parameters in order to form reliable estimates for all shared and private parameters with the fewest number of samples. This sampling algorithm has three key components: (i) data-driven sampling decisions, which dynamically over time specifies which of the K processes should be selected for sampling; (ii) stopping time for the process, which specifies when the accumulated data is sufficient to form reliable estimates and terminate the sampling process; and (iii) estimators for all shared and private parameters. Owing to the sequential estimation being known to be analytically intractable, this paper adopts *conditional* estimation cost functions, leading to a sequential estimation approach that was recently shown to render tractable analysis. Asymptotically optimal decision rules (sampling, stopping, and estimation) are delineated, and numerical experiments are provided to compare the efficacy and quality of the proposed procedure with those of the relevant approaches.**

I. Introduction

A. Overview

Consider the canonical estimation problem, in which we have a collection of probability measures $\mathcal{P} \triangleq \{\mathbb{P}(\cdot \mid \theta) \ : \ \theta \in \Theta\}$ defined over a common measurable space. The nature picks θ, the statistician draws samples from $\mathbb{P}(\cdot \mid \theta)$, and the objective is to use these samples to estimate θ. Building up on this canonical model, assume that we have a collection of K probability measures $\mathcal{P}_i \triangleq \{\mathbb{P}_i(\cdot \mid \theta) \ : \ \theta \in \Theta\}$ for $i \in [K] \triangleq \{1, \cdots, K\}$. Similarly, the nature selects θ and the statistician is given the freedom to collect samples from any one of the K models, and the objective is to form a reliable estimate for θ. With the objective of estimating θ with the fewest number of samples, a fundamental question pertains to which model(s) are the most reliable for estimating θ. If we can determine, *a priori*, which model is expected to be the most informative about θ in its entire range, then the answer is clear: always sample from that model. For instance, when $K = 2$ and

$$\mathbb{P}_1(\cdot \mid \theta) \sim \mathcal{N}(\theta, 1) \quad \text{and} \quad \mathbb{P}_2(\cdot \mid \theta) \sim \mathcal{N}(\theta, 2) \ , \tag{1}$$

under the mean-squared error cost function, $\mathbb{P}_1$ is always a more reliable model to use. In general, however, different models can be selectively more descriptive about θ in different regimes. For instance, for $\Theta \in [0, 1]$ consider the following two models

$$\mathbb{P}_1(\cdot \mid \theta) \sim \mathcal{N}(0, \theta) \quad \text{and} \quad \mathbb{P}_2(\cdot \mid \theta) \sim \mathcal{N}(0, 1 - \theta) \ . \tag{2}$$

It can be readily verified that under the mean-squared cost function, when $\theta \in (0, 1/2)$, $\mathbb{P}_1$ is the more informative model, and when $\theta \in (1/2, 1)$, $\mathbb{P}_2$ is the more informative one. Thus, without knowing the value of θ, it is impossible to specify which model is more informative. In such scenarios, any sampling action that focuses on only one model is sub-optimal. An optimal strategy will involve alternating between models until one can be identified as most informative with sufficient confidence. The complexity of such decisions can be further compounded when there are more parameters involved. As an example, consider a generalization of (2) in which the mean values are also unknown, i.e.,

$$\mathbb{P}_1(\cdot \mid \theta, \mu_1) \sim \mathcal{N}(\mu_1, \theta) \ \text{ and } \ \mathbb{P}_2(\cdot \mid \theta, \mu_2) \sim \mathcal{N}(\mu_2, 1 - \theta) \ , \tag{3}$$

where $\theta \in (0, 1)$ and $\mu_1, \mu_2 \in \mathbb{R}$. In these models, even when the most informative model for θ is known, we still need to draw samples from both models in order to be able to estimate μ_1 and μ_2, since μ_i can be estimated exclusively from model $i \in \{1, 2\}$. In this paper, we consider a general setting of K models that can capture both aspects discussed (i.e., shared and private parameters). The objective is to design a sequential and data-adaptive sampling procedure such that we can form sufficiently reliable estimates for the parameters involved with the fewest number of samples.

The design of such a sequential procedure involves performing three intertwined tasks dynamically over time. The first decision pertains to forming and updating estimates about the parameters over time. Besides the ultimate interest in the estimates, they also guide the sampling process. The second decision specifies the next model to be selected for sampling. These decisions at time t are formed based on the sequence of models selected up to time t, the data collected, and the estimates formed. Finally, a stopping decision has to be specified, which terminates the sampling process when sufficient confidence about the fidelity of the estimates has been reached. Designing such data-acquisition and inference mechanisms is related to active (controlled) sampling for sequential design of experiments, the foundations of which was laid out by Chernoff for binary composite hypothesis testing [1] through incorporating a data-acquisition process that dynamically decides about taking a finite number of data-acquisition actions. Under the assumption of uniformly distinguishable hypothesis and statistically independent control actions, there exists a rich body of literature on greedy algorithms (making decisions with best immediate return) that achieve optimality in the asymptote of diminishing rates of erroneous decisions.

Unlike the extensive literature for detection and classification problems, active sampling is far less investigated for estimation problems. Sequential designs of experiments for detection and estimation have a wide range of real-world applications. In machine learning, we are generally provided with a large dataset, from which machine learning literature almost universally assumes selecting a subset of samples uniformly at random, with the objective of performing various learning tasks. However, a network-guided sequential sampling strategy is shown to provide significant improvements in terms of sample complexity in [2]–[4]. This framework also finds strong resemblance to that of multi-armed bandits [5], which have a broad set of applications such as clinical drug testing and trial, brain and behavior modelling and dynamic pricing [6], computerized adaptive testing (CAT) [7], detecting intrusions in computer networks, [8] and target detection in radar systems [9]. Next, we provide an overview of the related literature on active sampling for inference problems that are most relevant to the scope of this paper.

A. Mukherjee and A. Tajer are with the Electrical, Computer, and Systems Engineering Department, Rensselaer Polytechnic Institute, Troy, NY.

P.-Y. Chen and P. Das are IBM Thomas J. Watson Research Center, Yorktown, NY.

The results were presented in part at the 2021 IEEE International Conference on Acoustics, Speech, and Signal Processing.

This research was supported in part by the U. S. National Science Foundation under the CAREER Award ECCS-1554482 and the Rensselaer-IBM Artificial Intelligence Research Collaboration program.

B. Related Literature

Active Sampling for Detection. The notion of active sequential detection was first studied for binary composite hypothesis testing in [1], providing asymptotically optimal design rules for sequential hypothesis testing under uniformly distinguishable hypotheses. This approach and its extensions to broader settings in [2] and [10]–[24], sequentially take control actions that greedily maximize the immediate return. Specifically, the rules identify the most likely decision at every step, and selects the action that maximally reinforces this decision. The studies in [10] and [11] extend the results to the cases of an infinite number of available actions, and an infinite number of true states of nature, respectively. The studies in [12] and [25] provide alternate sampling strategies that are empirically shown to outperform Chernoff's rule in the non-asymptotic regime. In a more recent study, the assumptions of having uniformly distinguishable hypothesis is relaxed in [14], and a modified Chernoff's rule is designed for multi-hypothesis testing. This approach introduces a sequence of intervals during which actions are selected uniformly at random instead of selecting the action that maximizes the immediate return. An extension to a stationery Markov model is investigated in [15]. Detection of an anomalous process from a finite set of anomalous processes has been considered in [26], where it is shown that Chernoff's rule is asymptotically optimal even without the assumption of uniformly distinguishable hypotheses, or using uniform selections at certain instants as prescribed in [14]. More applications of Chernoff's rule can be found in [17], [27] and [28]. Active sequential hypothesis testing coupled with a switching cost for switching between actions has been studied in [17]. In this study, a modification of the Chernoff's procedure is shown to be asymptotically optimal in the regime of diminishing error probabilities. The study in [27] considers the problem of active state tracking for a dynamic stochastic system whose states are varying through a discrete time Markov chain, and a dynamic programming-based approach is proposed for estimating the underlying states. The problem of active classification of graphs based on connectivity through partially observable nodes is studied in [28]. The problem of active model selection in Markov random fields (MRFs) is studied in [2], where data acquisition is modelled as a sequential process. Induced by the underlying dependence structure of an MRF, different sampling decisions over time are not necessarily statistically independent, rendering Chernoff's rule sub-optimal in this setting, even in asymptotic regimes. This arises mainly due to the fact that Chernoffs rule focuses on maximizing the immediate return, without paying attention to the expected future return. The approaches in [2] incorporates the effect of sampling actions on the expected rewards from future sampling decisions. Chernoff's test has been extended to a distributed setup in [29], designing distributed and consensus-based Chernoff tests for decentralized detection.

Other related literature on the sequential design of experiments for inference include the studies presented in [30] and [31], which propose approaches that take an initial number of samples according to certain predesignated rule in order to guess the true state of nature, and then select actions that maximize proper information measures based on the predicted states. Furthermore, [32] studies the problem of sequential design of experiments in the finite-sample regime, and proposes a switching policy which achieves an average sample complexity that is within a constant factor of the optimal average sampling complexity, while being computationally simpler than the optimal solution. The problem of multi-hypothesis testing using multiple controls is studied in [18], which frames the problem as a dynamic program, whose optimal solution is intractable. Thus, two heuristic policies are proposed and analyzed in both asymptotic and non-asymptotic regimes.

Active Sampling for Estimation – Theoreitial Aspects. Compared to the extensive literature on active sampling for detection and classification problems, the theoretical aspects of parameter estimation in an active sampling setting are far less investigated. Sequential estimation, without the additional complexity of control actions, has a rich body of literature in statistics. Reviews of sequential point and interval estimation can be found in [33]–[37]. Furthermore, a representative list of active sampling for parameter estimation includes [7], [24], and [38]–[41].

Sequential estimation from one time series is investigated in [24], where a fixed sampling cost is charged every time a new observation is procured, and a unified cost function aggregating the estimation accuracy and the total cost of sampling the observations is adopted. The stochastic unified cost function captures a trade-off between estimation fidelity and agility of arriving at an estimate. This study formalizes the notion of asymptotic point-wise optimality (A.P.O.) for assessing the performance of any sequential design, and proposes sequential procedures that are asymptotically optimal.

This study was extended in [38] to active sampling for estimation by incorporating a set of control actions for collecting the data. In [38], the objective is to estimate a function of an unknown parameter underlying each control action, such that the parameters under different control actions are different. Hence, the data collected based on a control action does not provide any information about the parameters associated with other actions. In [40], this drawback is mitigated by considering the case of estimating a single shared unknown parameter, while allowing the distributions under different control actions to be non-identical. This study also considers estimating a linear function of a set of shared unknown parameters. In a different setting, an intertwined problem of model detection and model probability kernel estimation is considered in [22], where in addition to discerning the true model (detection), an unknown probability kernel should be estimated too (estimation). Thus, the generalized stopping rule has to cater to both the quality of estimation and the intertwined task of detection. In all these studies, a unified cost function involving a linear combination of the estimation cost and the sampling cost are considered. Such approaches, in general, are known to be analytically intractable. To address this issue, a different approach to sequential estimation is introduced in [41], where a single control is assumed, and the conventional average estimation cost function is replaced by an average *conditional* estimation cost function. In [41], the objective is to form a reliable estimate of an unknown parameter using the fewest number of samples such that a Bayes posterior risk falls below a pre-specified threshold. This study provides an exactly optimal solution for this setting. In this paper, we adopt the approach of [41] for the estimation cost function in the active sampling framework. In this framework, we propose asymptotically optimal sequential procedures, with the main emphasis being on the design of efficient active sampling strategies.

Active Sampling for Estimation – Applied Aspects. On the application side, one domain in which active sampling for estimation is used is in the context of computerized adaptive testing (CAT). The goal in this setting is to devise a sequence of sampling decisions (items) that match the trait level of an examinee, characterized by an unknown parameter. This general setting and objective of CAT, while relevant, has major differences in its settings, assumptions, objective, and performance metrics. First, CAT falls in the category of finite-horizon settings. Since it has a fixed number of items to choose from [7], [42]–[44] and each item can be used at most once. This enforces a hard constraint on its stopping time. Even though some investigations introduce an adaptive stopping time [45]–[48], the stopping time may not exceed the number of items in the pool. In contrast, we consider an infinite-horizon setting, in which each experiment can be sampled

This article has been accepted for publication in IEEE Transactions on Signal Processing. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TSP.2022.3187655

as many times as necessary. In sequential analyses, finite versus infinite-horizon have distinct guarantees and designs. Secondly, CAT considers binary observations from the examinees, using the item response theory with a parameterized logistic regression model. This is distinct from the setting in this paper, which is much more general; essentially, we consider any model that satisfies some mild regularity conditions. Third, the objective function optimized in CAT is different from the one investigated in this paper. Specifically, CAT optimizes measures associated with the sample covariance, such as the largest eigenvalue or the trace [48]. On the contrary, we optimize the average sample complexity with a constraint on the estimation cost. This cost function has no precedence in the active sequential literature and has only been considered in [41] for the case of single control action. Finally, the key performance metrics investigated in CAT are the properties of the estimator, such as the bias and mean squared error reduction. In sequential statistics, the canonical performance metric is the average sample complexity [1], [15], [49]–[52]. The results investigated in this paper are based on characterizing the average sample complexity for the proposed procedures and analyzing their optimality.

C. Contributions

Our framework for active sampling for sequential parameter estimation consists of three intertwined decision rules: (i) a sequence of estimators for dynamically updating the estimates of interest over time, (ii) a sampling rule that specifies how the sampling decisions are guided by and adapted to the data, and (iii) a stopping rule, at which the sampling process is terminated, and the available data is deemed sufficient for forming sufficiently reliable estimates. The overall theoretical contribution is designing a combination of rules that collectively admit a form of optimality, addressing a setting for which theoretical guarantees are hitherto unknown. Besides the theoretical guarantee, we also specify our decision rules' relevance (similarities and differences) to the existing ones in active sampling and sequential estimation literature.

- **Estimator.** Generally, sequential estimation problems are not as analytically tractable as their detection counterparts, often rendering a lack of optimality guarantees. In this paper, we adopt a new estimation cost function introduced in [41]. This cost function focuses on the *conditional* estimation cost, distinct from the conventional average estimation cost functions adopted in sequential estimation. This choice of the cost function facilitates tractable analysis and closed-form characterization of the decision rules that enjoy optimality guarantees. This cost function has no precedence in the active sampling frameworks. Besides the literature on the theoretical aspects, this cost function does not have precedence in the relevant applied domains (e.g., CAT), which generally focus on the conventional average estimation cost function, distinct from the average posterior cost function adopted in this paper.
- **Sampling rule.** Our setting and analysis generalize those of [41], which focuses on a single experiment. We consider a two-fold generalization. First, we consider multiple experiments, where one experiment can be chosen at each instant. Secondly, each experiment can have both shared and private parameters, each having different guarantees on the estimation cost. Thus, the sampling rule, at its core, has to balance between greedily sampling the most informative experiment and ensuring that the guarantees for the less informative experiments are satisfied. We adopt a sampling rule that ensures an optimal trade-off between exploiting the most informative experiment while providing sufficient exploration for the other experiments. This is a key distinction from the existing theoretical investigation of active sampling for sequential estimation [40] and their applications, e.g., in CAT. In CAT, specifically, the active selection of the experiments has two main differences. First, they do not have the notion of shared versus private parameters and consequently do not have the inherent estimation trade-off involved. Secondly, CAT by design is a finite-horizon problem, rendering the design of the sampling rules fundamentally a dynamic programming problem. This is distinct from our setting, which is infinite-horizon with a stochastic stopping time.
- **Stopping rule.** Finally, we propose a stopping rule that ensures that the respective guarantees on the estimation cost for both the shared and private parameters are satisfied while, in parallel, the average delay is minimized. Specifically, we provide universal lower bounds on the average delay required to achieve the guarantees on the estimation cost and then show that our proposed procedures match these bounds asymptotically. Furthermore, we show that the most relevant sampling rule in [40] is sub-optimal in the general setting. These stopping rules, being infinite-horizon, are quite distinct from those generally adopted in CAT.

II. Active Sampling Model

Consider a finite set of experiments $\{S_1, \cdots, S_K\}$. When experiment S_i for $i \in [K] \triangleq \{1, \cdots, K\}$ is selected, it generates a random variable whose probability distribution depends on two unknown parameters $\theta \in \Theta \subseteq \mathbb{R}$ and $\alpha_i \in \mathcal{A}_i \subseteq \mathbb{R}$. We call θ the *shared* unknown parameter of all experiments, and call α_i the *private* unknown parameter of experiment $i \in [K]$. As an example, consider a sensor network in which each sensor produces noisy observations of a shared parameter θ (such as the temperature and pressure), and the task is to estimate the shared parameter from the observations obtained from the sensors. Furthermore, the stochastic processes corresponding to each sensor $i \in [K]$ could also have unknown private parameters α_i governing the data generation, such as the noise variance of each sensor. Our objective is to have a sequential design of experiments such that we can estimate θ and $\boldsymbol{\alpha} \triangleq (\alpha_1, \cdots, \alpha_K)$ with the fewest number of experiments (samples). We denote the probability density function (pdf) of the samples generated by experiment S_i by $f_i(\cdot \mid \theta, \alpha_i)$, where for simplicity of notations we are assuming that the pdfs are well-defined without any zero over lower-dimensional manifolds. We also assume that the samples generated by an experiment are independent and identically distributed (i.i.d.). Our focus is on the Bayesian setting and we denote the prior pdfs of θ and α_i by π_θ and π_i, respectively.

A. Sampling Model

We consider a fully sequential data-acquisition mechanism, according to which we are allowed to gather one sample at-a-time. Our primary goal is to identify a sequence of experiments for forming reliable estimates for θ and $\boldsymbol{\alpha}$ using the minimum number of experiments on average. An optimal sampling process will require balancing a number of competing measures in an optimal way. The trade-offs we are facing include:

- **Delay versus fidelity:** Collecting more samples, when used judiciously, results in improved estimation fidelity at the expense of increased delay in forming estimates.
- **Model discrepancies:** Not all experiments are equally informative about the shared parameter. Nevertheless, *a priori*, it might not be known which experiments are the most informative about the realization of θ. Hence, this creates a natural exploration versus exploitation trade-off, according to which the experiments

This article has been accepted for publication in IEEE Transactions on Signal Processing. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TSP.2022.3187655

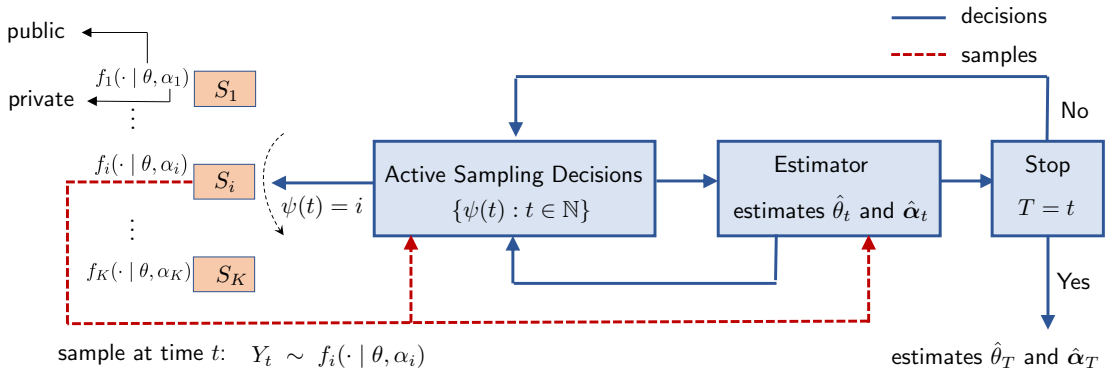

Fig. 1: Active sampling for sequential parameter estimation. $\hat{\theta}_t$ and $\hat{\boldsymbol{\alpha}}_t$ denote the estimates of the shared and private parameters respectively, $\{\psi(t) : t \in \mathbb{N}\}$ represents the sequence of selected experiments, and T represents the stopping time.

serve a two-fold purpose: identifying the most informative experiments and using the data from these experiments to estimate θ.

- **Shared versus private parameters:** While for estimating the shared parameter there is the tendency to identify the most informative experiment for θ, for estimating the private parameters each experiment should be selected sufficiently often to render a reliable estimate for its private parameter.

To formalize the sampling process, we need to specify three decision-making tasks. First, we specify a sampling rule $\psi : \mathbb{N} \to [K]$ where $\psi(t) \in [K]$ denotes the experiment to be selected at time t. We denote the sample collected from the experiment $S_{\psi(t)}$ by Y_t. Accordingly, we define the *ordered* sets

$$Y^t \triangleq \{Y_1, \cdots, Y_t\}, \quad \text{and} \quad \psi^t \triangleq \{\psi(1), \cdots, \psi(t)\}. \tag{4}$$

The observed filtration induced by a sampling rule ψ is denoted by $\mathcal{F}_t^\psi \triangleq \sigma(Y^t, \psi^t)$[1]. We assume that $\psi(t+1)$ is $\mathcal{F}_t^\psi$-measurable. The second decision-making task is a stopping rule that determines the instance at which we have accumulated sufficient evidence to form reliable estimates. Specifically, given a sampling rule ψ, we define an $\mathcal{F}_T^\psi$-stopping time that is a randomization T such that $\{T = t\} \in \mathcal{F}_t^\psi$, for all $t \in \mathbb{N}$. Finally, corresponding to a given sampling rule ψ, we define $\hat{\theta}(Y^t, \psi^t)$ and $\hat{\alpha}_i(Y^t, \psi^t)$ as the estimators for θ and α_i respectively[2]. We refer to the tuple

$$\Delta \triangleq (T, \psi, \Phi_T), \tag{5}$$

as the collection of rules involved in active sampling for sequential estimation, where we have defined

$$\Phi_t \triangleq (\hat{\theta}(Y^t, \psi^t), \hat{\boldsymbol{\alpha}}(Y^t, \psi^t)). \tag{6}$$

III. Problem Statements

In this section, we formalize the problem of active sampling for sequential estimation. There are two key figures of merit involved in characterizing the performance of the active sampling framework: the average delay (sample complexity) and the estimation costs incurred by the final estimates. There exists a tension between these two quantities, since improving one penalizes the other one. Specifically, improving the estimates necessitates collecting more samples, which in turn, penalizes the sample complexity. Capturing this trade-off, our formulation aims to maintain the estimation cost below a target threshold and, in parallel, minimize the average sample complexity.

[1]Unless otherwise stated, we use the shorthand $\mathcal{F}_t$ for $\mathcal{F}_t^\psi$, where the sampling rule ψ is clear from the context.

[2]Throughout the paper, sometimes we use the short-hands $\hat{\theta}_t$ and $\hat{\alpha}_{i,t}$ for $\hat{\theta}(Y^t, \psi^t)$ and $\hat{\alpha}_{i,t}(Y^t, \psi^t)$ respectively.

A. Estimating Shared Parameters

We start by focusing on the setting in which we have unknown shared parameters, and the private parameters are fully known. In this setting, we denote the pdf of the samples generated by the experiment S_i by $f_i(\cdot \mid \theta)$, for $i \in [K]$. Define $\ell(\theta, \hat{\theta})$ as a non-negative cost function that captures the fidelity of the estimate $\hat{\theta}$ with respect to the ground truth θ. Accordingly, $\mathbb{E}[\ell(\hat{\theta}, \theta)]$ is the associated average cost of the estimate $\hat{\theta}$, where the expectation is taken with respect to the data and the prior distribution of θ.

The posterior distribution of θ at time t given the set of samples and control actions taken up to time t is given by

$$\pi_\theta^t(\theta) = \frac{\pi_\theta(\theta) \prod_{i \in \psi^t} f_i(Y_i \mid \theta)}{\int_{v \in \Theta} \pi_\theta(v) \prod_{i \in \psi^t} f_i(Y_i \mid v) \, \mathrm{d}v}. \tag{7}$$

Based on the above posterior, we define the *conditional average* posterior cost for an estimate $\hat{\theta}_t$ formed at time t as

$$\mathsf{C}(\hat{\theta}_t \mid \mathcal{F}_t) \triangleq \mathbb{E}_t\big[\ell(\hat{\theta}_t, \theta) \mid \mathcal{F}_t\big], \tag{8}$$

where $\mathbb{E}_t$ denotes expectations with respect to π_θ^t. We denote the associated average posterior cost by

$$\mathsf{C}(\hat{\theta}_t) = \mathbb{E}_t[\ell(\theta, \hat{\theta}_t)]. \tag{9}$$

Note that the estimation cost $\mathsf{C}(\hat{\theta}_t)$ depends on the choice of the cost function $\ell(\cdot, \cdot)$. Thus, the algorithm design and its optimality properties depend on the choice of the cost function. In this paper, we use the quadratic cost function for our analyses. The sequence of conditional posterior cost functions $\{\mathsf{C}(\hat{\theta}_t \mid \mathcal{F}_t) : t \in \mathbb{N}\}$ and posterior cost functions $\{\mathsf{C}(\hat{\theta}_t) : t \in \mathbb{N}\}$ are $\mathcal{F}_t$-measurable and form increasing sequences of σ-fields. A natural approach to formulating sequential estimation is finding a solution to:

$$\tilde{\mathcal{P}}(\tilde{\beta}) \triangleq \begin{cases} \inf\limits_{\Delta} & \mathbb{E}[T] \\ \text{s.t.} & \mathsf{C}(\hat{\theta}_t) \leq \tilde{\beta} \end{cases}, \tag{10}$$

where $\tilde{\beta} \in \mathbb{R}_+$ controls the estimation quality. However, as discussed in [41] and its references [3-6], solving (10) even in simpler settings, e.g., $K = 1$, in which we do not have the action sampling decisions, is analytically intractable. In this paper, instead of (10), we adopt the approach of [41] and in (10) replace the average posterior cost $\mathsf{C}(\hat{\theta}_t)$ with the conditional average posterior cost $\mathsf{C}(\hat{\theta}_t \mid \mathcal{F}_t)$. Hence our objective is to minimize the average sample size such that the conditional average posterior cost falls below a prescribed threshold, formalized as

$$\mathcal{P}(\beta) \triangleq \begin{cases} \inf\limits_{\Delta} & \mathbb{E}[T] \\ \text{s.t.} & \mathsf{C}(\hat{\theta}_T \mid \mathcal{F}_T) \leq \beta \end{cases} , \tag{11}$$

where $\beta \in \mathbb{R}_+$ controls the quality of the estimate $\hat{\theta}_T$ given $\mathcal{F}_T$.

B. Estimating Shared and Private Parameters

Next, we consider a generalization of the framework, such that in addition to the shared parameter, experiment S_i also has a private statistically independent parameter α_i to estimate. To formalize this, we denote the posterior joint pdf of parameters θ and $\{\alpha_i : i \in [K]\}$ conditioned on Y^t and ψ^t by g_t. By defining the space $\mathcal{A} \triangleq \prod_{i=1}^{K} \mathcal{A}_i$ we have

$$g_t(\theta, \boldsymbol{\alpha}) \;=\; \frac{\pi_\theta(\theta) \prod\limits_{j=1}^{K} \pi_i(\alpha_j) \prod\limits_{i=1}^{t} f_i(Y_i \mid \theta, \boldsymbol{\alpha})}{\int_{\mathbf{u}\in\mathcal{A}} \int_{v\in\Theta} \pi_\theta(v) \prod\limits_{j=1}^{K} \pi_j(u_j) \prod\limits_{i=1}^{t} f_i(Y_i \mid v, u) \;\mathrm{d}v \,\mathrm{d}\mathbf{u}} . \tag{12}$$

Subsequently, the marginal posterior pdf of θ conditioned on Y^t and ψ^t is denoted by

$$\pi_\theta^t(\theta) \triangleq \int_{\mathbf{v}\in\mathcal{A}} g_t(\theta, \mathbf{v}) \,\mathrm{d}\mathbf{v} . \tag{13}$$

Furthermore, let $S_i^t \triangleq \{s \in \{1, \cdots, t\} \;:\; \psi(s) = i\}$ denote the set of instances up to time t at which experiment $i \in [K]$ is selected. Hence, the posterior pdf of α_i conditioned on θ, Y^t and ψ^t is

$$h_i^t(\alpha_i) \;=\; \frac{\pi_i(\alpha_i) \prod\limits_{s\in S_i^t} f_i(Y_s \mid \alpha_i, \theta)}{\int_{v\in\mathcal{A}_i} \pi_i(v) \prod\limits_{s\in S_i^t} f_i(Y_s \mid v, \theta) \,\mathrm{d}v} . \tag{14}$$

Based on the above definitions, we now define the *conditional average estimation cost* for estimating α_i as

$$\mathsf{D}(\hat{\alpha}_{i,t} \mid \mathcal{F}_t, \hat{\theta}_t) \;\triangleq\; \mathbb{E}_t^i\big[\ell(\hat{\alpha}_{i,t}, \alpha_i) \mid \mathcal{F}_t, \hat{\theta}_t\big] , \tag{15}$$

where $\mathbb{E}_t^i$ is the expectation with respect to $h_i^t(\alpha_i)$. This setting emphasizes a hierarchy of inference objectives in which the primary objective is estimating the shared parameter θ. The estimate of θ, subsequently, guides estimating the private parameters. By incorporating constants that capture the fidelity of the estimates for $\{\alpha_i : i \in [K]\}$, an optimal sequential estimation procedure can be found as the solution to:

$$\mathcal{P}(\boldsymbol{\beta}) \triangleq \begin{cases} \min\limits_{\Delta} & \mathbb{E}[T] \\ \text{s.t.} & \mathsf{C}(\hat{\theta}_T \mid \mathcal{F}_T) \leq \beta \\ \text{and} & \mathsf{D}(\hat{\alpha}_{i,T} \mid \mathcal{F}_T,\, \hat{\theta}_T) \leq \beta_i,\; \forall i \in [K] \end{cases} , \tag{16}$$

where we have defined $\boldsymbol{\beta} \triangleq [\beta, \beta_1, \cdots, \beta_K]$, and $\beta_i \in \mathbb{R}_+$ controls the estimation quality of α_i.

C. Technical Assumptions

In this section, we provide the assumptions under which the performance guarantees are established. The assumptions are mainly necessary for the existence and consistency of the maximum likelihood (ML) estimates and the existence of the Fisher information (FI) measures of the relevant parameters. To proceed, corresponding to the pdfs $f_i(x \mid \theta)$ (only shared parameters) and $f_i(x \mid \theta, \alpha_i)$ (shared and private parameters), we define the log-likelihood functions

$$\lambda_i(x \mid \theta) \;\triangleq\; \log f_i(x \mid \theta) , \quad \text{for } i \in [K] , \tag{17}$$

$$\lambda_i(x \mid \theta, \alpha_i) \;\triangleq\; \log f_i(x \mid \theta, \alpha_i) , \quad \text{for } i \in [K] . \tag{18}$$

- A_1: Parameter spaces Θ and $\{\mathcal{A}_i : i \in [K]\}$ are assumed to be non-empty and compact.
- A_2: We assume that the ML estimate of the parameter θ exists, and it is finite. Specifically, when we have only the shared parameter, $\mathbb{E}[|\lambda_i(x \mid \theta)|] < +\infty$ for all $\theta \in \Theta$. Similarly, for the setting with shared and private parameters, for each $i \in [K]$, we assume $\mathbb{E}[|\lambda_i(x \mid \theta, \alpha_i)|] < +\infty$ for all $\theta \in \Theta$ and $\alpha_i \in \mathcal{A}_i$.
- A_3: The log-likelihood functions $\{\lambda_i(x \mid \theta) : i \in [K]\}$ are assumed to be continuous and differentiable for all $\theta \in \Theta$. The first order derivative $\frac{\partial}{\partial\theta}\lambda_i(x \mid \theta)$ is assumed to be bounded, continuous, and differentiable everywhere, such that the second derivative $\frac{\partial^2}{\partial\theta^2}\lambda_i(x \mid \theta)$ exists and is bounded. Similarly, the second derivatives $\frac{\partial^2}{\partial\theta^2}\lambda_i(x \mid \theta, \alpha_i)$ and $\frac{\partial^2}{\partial\alpha_i^2}\lambda_i(x \mid \theta, \alpha_i)$ are assumed to exist and be bounded.
- A_4: The pdfs $\{f_i(x \mid \theta) : i \in [K]\}$ and $\{f_i(x \mid \theta, \alpha_i) : i \in [K]\}$ are assumed to have the same support.
- A_5: In the setting with only the shared parameter, let us denote the FI measures under the model $i \in [K]$ for the shared parameter by

$$\mathscr{I}_i(\theta) \;\triangleq\; -\mathbb{E}\left[\frac{\partial^2}{\partial\theta^2}\lambda_i(x \mid \theta)\right] . \tag{19}$$

Similarly, when we have shared and private parameters, let us denote the FI measures under the model $i \in [K]$ by

$$\mathscr{J}_i(\theta) \;\triangleq\; -\mathbb{E}\left[\frac{\partial^2}{\partial\theta^2}\lambda_i(x \mid \theta, \alpha_i)\right] , \tag{20}$$

$$\text{and} \quad \mathscr{J}_i(\alpha_i) \;\triangleq\; -\mathbb{E}\left[\frac{\partial^2}{\partial\alpha_i^2}\lambda_i(x \mid \theta, \alpha_i)\right] . \tag{21}$$

We assume that the FI measures are bounded and continuous functions of θ and α_i.

- A_6: We assume that the log-likelihood functions under two sufficiently distinguishable parameters θ and $\bar{\theta}$ are also distinguishable, that is,

$$\mathbb{E}_\theta\left[\sup_{\bar{\theta}}\{\lambda_i(x \mid \theta) - \lambda_i(x \mid \bar{\theta}) \;:\; |\theta - \bar{\theta}| \;>\; \epsilon\}\right] \;<\; 0 . \tag{22}$$

Similarly, for the setting with shared and private parameters, under sufficiently distinguishable parameters θ and $\bar{\theta}$, the log-likelihoods are also distinguishable, i.e.,

$$\mathbb{E}_\theta\left[\sup_{\bar{\theta}}\{\lambda_i(x \mid \theta, \alpha) - \lambda_i(x \mid \bar{\theta}, \alpha) \;:\; |\theta - \bar{\theta}| \;>\; \epsilon\}\right] \;<\; 0 , \tag{23}$$

and, for $\alpha, \bar{\alpha} \in \mathcal{A}_i$ we have

$$\mathbb{E}_\alpha\left[\sup_{\bar{\alpha}}\{\lambda_i(x \mid \theta, \alpha) - \lambda_i(x \mid \theta, \bar{\alpha}) \;:\; |\alpha - \bar{\alpha}| \;>\; \epsilon\}\right] \;<\; 0 . \tag{24}$$

Assumptions A_1, A_2, A_3, and A_6 collectively establish the existence of the ML estimates of the shared and private parameters. Assumptions A_3 and A_5 prove the existence of the FI measures corresponding to the shared and private parameters. Assumption A_4 requires every experiment to have the same support for the different parameters. We note that these assumptions are satisfied by a wide range of distributions, including continuous distributions (e.g., Gaussian and exponential) and discrete distributions (e.g., Bernoulli).

IV. Active Sampling for Shared Parameters

In this section, we describe the active estimation procedure that solves the problem specified in (11) for the setting in which we only have the shared parameter.

A. Bayesian and Maximum Likelihood Estimators

The goal of the optimal estimators is to minimize the conditional average cost $\mathsf{C}(\hat{\theta}_T \mid \mathcal{F}_T)$ for given choices of T and ψ^T captured by $\mathcal{F}_T$. For $\mathsf{C}(\hat{\theta}_T \mid \mathcal{F}_T)$ we have

$$\mathsf{C}(\hat{\theta}_T \mid \mathcal{F}_T) = \mathbb{E}_t\left[\sum_{t=0}^{\infty} \ell(\hat{\theta}_t, \theta)\mathbb{1}_{\{t=T\}} \mid \mathcal{F}_t\right] \tag{25}$$

$$= \sum_{t=0}^{\infty} \mathbb{E}_t\Big[\ell(\hat{\theta}_t, \theta) \mid \mathcal{F}_t\Big]\, \mathbb{1}_{\{t=T\}} \tag{26}$$

$$\geq \sum_{t=0}^{\infty} \inf_u\ \mathbb{E}_t\left[\ell(u, \theta) \mid \mathcal{F}_t\right] \mathbb{1}_{\{t=T\}}\ . \tag{27}$$

The transition from (25) to (26) is due to the indicator function $\mathbb{1}_{\{t=T\}}$ being $\mathcal{F}_t$-measurable. We denote the Bayes optimal estimator by

$$\nu_t^{\mathsf{B}} \ \triangleq\ \arg\inf_u\ \mathbb{E}_t\left[\ell(u, \theta) \mid \mathcal{F}_t\right]\ . \tag{28}$$

Accordingly, we denote the conditional average cost associated with the Bayesian estimate by

$$\mathsf{C}_t^{\mathsf{B}} \ \triangleq\ \mathbb{E}_t\Big[\ell(\nu_t^{\mathsf{B}}, \theta) \mid \mathcal{F}_t\Big]\ . \tag{29}$$

Hence, from (25)-(29), we have

$$\mathsf{C}(\hat{\theta}_T \mid \mathcal{F}_T)\ \geq\ \sum_{t=0}^{\infty} \mathsf{C}_t^{\mathsf{B}} \cdot \mathbb{1}_{\{T=t\}}\ =\ \mathsf{C}_T^{\mathsf{B}}\ . \tag{30}$$

This indicates that for any stopping time T, using Bayes optimal estimator at stopping time minimizes the estimation cost. There are several possible choices of the cost function $\ell(\cdot,\cdot)$. We could choose $\ell(\cdot,\cdot)$ to be the maximum a-posteriori probability (MAP) cost function, i.e.,

$$\ell^{\mathsf{MAP}}(\hat{\theta}, \theta) \triangleq \begin{cases} 0\ , & \text{if } \|\hat{\theta} - \theta\| \leq c \\ 1\ , & \text{if } \|\hat{\theta} - \theta\| > c \end{cases}\ , \tag{31}$$

and, $\mathsf{C}_t^{\mathsf{MAP}}(\hat{\theta}_t \mid \mathcal{F}_t) \triangleq \mathbb{E}_t[\ell^{\mathsf{MAP}}(\hat{\theta}_t, \theta) \mid \mathcal{F}_t]$ is the posterior cost corresponding to the MAP cost function. As another example, $\ell(\cdot,\cdot)$ could be chosen to be the median estimation cost, i.e.,

$$\ell^{\mathsf{median}}(\hat{\theta}, \theta) \triangleq \|\hat{\theta} - \theta\|_1\ , \tag{32}$$

in which case, the posterior cost function is given by $\mathsf{C}_t^{\mathsf{median}}(\hat{\theta}_t \mid \mathcal{F}_t) \triangleq \mathbb{E}_t[\ell^{\mathsf{median}}(\hat{\theta}_t\theta) \mid \mathcal{F}_t]$. As mentioned, throughout this paper, we consider the minimum mean squared error (MMSE) cost function, i.e.,

$$\ell(\hat{\theta}, \theta) \triangleq (\hat{\theta} - \theta)^2\ . \tag{33}$$

Corresponding to this, the conditional average posterior cost $\mathsf{C}(\hat{\theta}_t \mid \mathcal{F}_t)$ at time t is given by $\mathsf{C}(\hat{\theta}_t \mid \mathcal{F}_t) \triangleq \mathbb{E}_t[(\hat{\theta}_t - \theta)^2 \mid \mathcal{F}_t]$. The Bayesian estimator under the MMSE cost function becomes the MMSE estimator, which we denote by

$$\nu_t^{\mathsf{MMSE}} \ \triangleq\ \mathbb{E}_t\big[\theta \mid \mathcal{F}_t\big]\ . \tag{34}$$

By specializing (29) to the quadratic (MMSE) cost function and denoting the conditional average MMSE by $\mathsf{C}_t^{\mathsf{MMSE}}$ we have

$$\mathsf{C}(\hat{\theta}_T \mid \mathcal{F}_t)\ \geq\ \sum_{t=0}^{\infty} \mathsf{C}_t^{\mathsf{MMSE}} \cdot \mathbb{1}_{\{T=t\}}\ =\ \mathsf{C}_T^{\mathsf{MMSE}}\ , \tag{35}$$

rendering the MMSE estimate as the optimal estimate under the quadratic cost function. Besides the MMSE estimator ν_t^{MMSE}, we also use the ML estimator for designing our algorithm. Specifically, the MMSE and the ML estimates each serve a specific purpose. We use the MMSE estimate for producing the final estimates for the parameters of interest, and use the ML estimate for guiding the sampling decisions. We denote the ML estimator of θ by

$$\nu_t^{\mathsf{ML}} \ \triangleq\ \arg\max_{\theta\in\Theta}\ \sum_{i\in\psi^t} \lambda_i(Y_i \mid \theta)\ . \tag{36}$$

B. Chernoff-like Sampling Rule

Our sampling rule follows the spirit of the sequential experimental design due to Chernoff [1], which addresses the problem of active sampling for binary composite hypothesis testing. Under the assumption of uniformly distinguishable hypothesis and a finite set of control actions, at each round, Chernoff's rule decides in favor of the design that maximizes the *immediate* return. Such return, in the context of hypothesis testing, is characterized by a function of the Kullback-Leibler (KL) divergence of the models under different hypotheses. Specifically, Chernoff's rule determines the maximum likelihood (ML) decision about the most likely hypothesis at each instant, and then chooses an action that maximally reinforces this decision.

In the context of sequential estimation, at each time step $t \in \mathbb{N}$, we wish to select the experiment that results in the *most informative* observation, that is the one which is likely to produce the largest reduction in the estimation cost. As a relevant measure for comparing the informativeness of various experiments in the sequential estimation framework, we adopt the FI measure. Specifically, for selecting the experiment $\psi(t) \in [K]$, we compute the ML estimate generated by ν_{t-1}^{ML} based on the sequence of samples accumulated up to time $t-1$. We then select the experiment that maximizes the FI measure computed at the ML estimate, i.e.,

$$\psi(t)\ =\ \arg\max_{i\in[K]}\ \mathscr{I}_i(\nu_{t-1}^{\mathsf{ML}})\ , \tag{37}$$

where a potential tie is broken by selecting one uniformly at random. This sampling rule is greedy in the sense that it only focuses on exploiting the most informative experiment. We will show that this rule is optimal only for estimating the shared parameter. It loses its optimality when we also have private parameters, for which we provide an alternate sampling rule in Section V.

C. Stopping Rule

Finally, we specify the stopping rule that characterizes the end of the sampling procedure. The rule is directly driven by the decision quality constraint specified in the formulation of problem $\mathcal{P}(\beta)$ in (11). Specifically, based on $\mathcal{P}(\beta)$, we are interested in minimizing the number of samples such that the average posterior estimation cost falls below the target reliability threshold β. Thus, we set the stopping time as the first time that the cost $\mathsf{C}(\nu_T^{\mathsf{MMSE}} \mid \mathcal{F}_T)$ falls below β, i.e.,

$$T\ \triangleq\ \inf\ \big\{t \in \mathbb{N}\ :\ \mathsf{C}(\nu_t^{\mathsf{MMSE}} \mid \mathcal{F}_t)\ \leq\ \beta\big\}\ . \tag{38}$$

The structure of (38) is similar to that of [41], with the key difference that the posterior variance $\mathsf{C}(\nu_T^{\mathsf{MMSE}} \mid \mathcal{F}_T)$ not only depends on the estimator, but also it depends on the sampling path ψ^T, which does not exist in [41].

D. Performance Guarantees

In this section, we evaluate the optimality of the active sampling procedure for sequential estimation. First, we provide a universal

lower bound on the average sample complexity of any procedure that solves (11). Next, we provide a high-probability upper bound on the sample complexity achieved by our proposed procedure for a range of the prescribed guarantee β. Subsequently, we remark that the average sample complexity of our proposed procedure achieves the lower bound in the asymptote of small values of β. To this end, corresponding to θ, we define $V_\beta(\theta)$, which is instrumental in characterizing the sample complexity of our sequential estimation problem:

$$V_\beta(\theta) \triangleq \inf_{\mathbf{p}\in\mathcal{Q}_K} \frac{1}{\beta}\left(\sum_{i\in[K]} p(i)\mathscr{I}_i(\theta)\right)^{-1}, \tag{39}$$

where $\mathcal{Q}_K$ denotes the K-dimensional probability simplex and we have defined $\mathbf{p} \triangleq [p_1, \cdots, p_K]$. It can be readily verified that under the assumptions A_1-A_6, (39) can be simplified to $V_\beta(\theta) = \min_{i\in[K]} \frac{1}{\beta.\mathscr{I}_i(\theta)}$. Based on this definition, we now provide a lower bound on the average sample complexity. Note that for Theorem 1 and all the other subsequent theorems, expectation $\mathbb{E}$ and probability $\mathbb{P}$ measures are with respect to the measures induced by the randomness in the observations, the control actions, and the stochastic stopping rule.

Theorem 1 (Converse)**.** *Under assumptions A_1-A_6, for any sequential procedure Δ, and for any $h > 0$, there exists a constant $C(h) > 0$ such that for any $\beta \in (0, C(h))$, we have*

$$\mathbb{E}[T] \geq \left(V_\beta(\theta) - \frac{h}{\beta}\right)(1-h) . \tag{40}$$

Proof. See appendix B. ■

Theorem 2 (Achievability)**.** *Under assumptions A_1-A_6, for any $h > 0$, there exists a constant $C'(h) > 0$ such that for any $\beta \in (0, C'(h))$, the proposed procedure achieves*

$$\mathbb{P}\left\{T \leq V_\beta(\theta) + \frac{h}{\beta} + 1\right\} = 1 . \tag{41}$$

Proof. See appendix C. ■

Theorem 3 (Achievability)**.** *In the asymptote of $\beta \to 0$, the proposed procedure satisfies*

$$\lim_{\beta\to 0} \frac{\mathbb{E}[T]}{V_\beta(\theta)} \leq 1 . \tag{42}$$

Proof. See Appendix D. ■

Note that the upper bound on the average delay provided in Theorem 3 matches the universal lower bound on the average delay obtained in Theorem 1. Specifically, since Theorem 1 holds for any $h > 0$, we can take the supremum over h, followed by the limit with respect to β. Thus, in the asymptote of $\beta \to 0$, the lower-bound on the average sample complexity specified in (40) becomes $\limsup_{\beta\to 0} \mathbb{E}[T]/V_\beta(\theta) \geq 1$, which is also the upper-bound on the average sample complexity of the proposed procedure, specified in (42). This establishes the optimality of the proposed procedure Δ, in the asymptote of a diminishing guarantee on the estimation cost.

Furthermore, note that under the presence of a single control, when $K = 1$, our procedure reduces to that of [41]. We note that [41] addresses a problem that is a special case of the problem we consider in two ways. First, it focuses only on one data stream (experiment). Secondly, which is also an artifact of the first point, in [41] there is no notion of active selection of the experiments/streams. By setting $K = 1$, our asymptotic bound on the average sample complexity provided in Theorem 3 reduces to the result of [41]. Specifically, it can be readily verified that for the case of $K = 1$, we have

$$\lim_{\beta\to 0} \frac{\mathbb{E}[T]}{(1/\beta)} \leq \mathscr{J}(\theta)^{-1} . \tag{43}$$

Although [41] does not provide any expression for the average delay, it proves the optimality of the proposed rules $(\nu_T^{\mathsf{MMSE}}, T)$, which are the same estimate and stopping rule that we specified in (34) and (38). Furthermore, for $K = 1$, it is shown in [41] that the proposed procedure achieves optimality in all regimes (both asymptotic and non-asymptotic, for all values of β). This is due to the fact that when $K = 1$, the control action has only one experiment to choose from.

V. Active Sampling for Shared and Private Parameters

In this section, we extend the active sampling procedure and the attendant performance guarantees to the settings in which the experiments have both shared and private parameters.

A. Bayesian and Maximum Likelihood Estimators

Following the same line of arguments as in (25)-(35), we use the MMSE estimator for minimizing the average posterior conditional MMSE estimation costs $\mathsf{C}(\hat{\theta}_T \mid \mathcal{F}_T)$ and $\mathsf{D}(\hat{\alpha}_{i,T} \mid \mathcal{F}_T, \hat{\theta}_T)$ defined in (9) and (15), respectively. Accordingly, we denote the estimators for the parameters θ and $\{\alpha_i : i \in [K]\}$ by

$$\nu_t^{\mathsf{MMSE}} \triangleq \mathbb{E}_t\big[\theta \mid \mathcal{F}_t\big] , \text{ and } \zeta_{i,t}^{\mathsf{MMSE}} \triangleq \mathbb{E}_t^i\big[\alpha_i \mid \mathcal{F}_t,\ \nu_t^{\mathsf{MMSE}}\big] . \tag{44}$$

Furthermore, in our sampling rule, we leverage the ML estimates of the parameters θ and $\{\alpha_i : i \in [K]\}$. For this purpose, let $\phi_i \triangleq [\theta, \alpha_i]$ denote the vector containing the shared and the private parameters. Furthermore, we denote the ML estimators by

$$\nu_t^{\mathsf{ML}} \triangleq \arg\max_{\theta\in\Theta} \sum_{i\in\psi^t} \lambda_i(Y_i \mid \theta, \boldsymbol{\alpha}) , \tag{45}$$

$$\zeta_{i,t}^{\mathsf{ML}} \triangleq \arg\max_{\alpha_i\in\mathcal{A}_i} \sum_{s\in\psi_i^t} \lambda_i(Y_s \mid \theta, \alpha_i) , \tag{46}$$

$$\text{and} \quad \phi_{i,t}^{\mathsf{ML}} \triangleq \arg\max_{\phi\in[\theta\times\mathcal{A}_i]} \sum_{s\in\psi_i^t} \lambda_i(Y_s \mid \phi) , \tag{47}$$

where we have defined $\psi_i^t \triangleq \big\{t \in \{1, \cdots, t\} : \psi(t) = i\big\}$ as the ordered sequence of time instants during $\{1, \cdots, t\}$ at which experiment $i \in [K]$ is selected for sampling.

B. Sampling Rules

To accommodate the distinct levels of tolerance for the estimation costs associated with the shared and the private parameters, we need a sampling rule that is adaptive to the thresholds $\boldsymbol{\beta}$ imposed on the cost functions. Before we formally specify our sampling rule, we discuss an adaptation of the greedy sampling rule that we used in Section IV-B, and show that such a greedy approach becomes sub-optimal in this setting, caused by insufficient exploration.

1) Greedy Sampling Rule: The sampling rule described in Section IV-B aims at selecting the experiment that maximizes the FI measure at the current ML estimate. By generalizing this approach, a greedy sampling strategy aims at selecting the experiment that maximizes the FI measure for the parameter ϕ_i for all $i \in [K]$. This can be formalized as:

$$\psi^{\mathsf{c}}(t) \triangleq \arg\max_{i\in[K]} \mathsf{tr}\Big\{\mathscr{I}_i(\phi_{i,t}^{\mathsf{ML}})\Big\} , \tag{48}$$

where $\mathscr{I}_i(\phi_i)$ represents the Fisher information matrix (FIM) associated with ϕ_i, for $i \in [K]$. This sampling rule does not use

the information that the guarantees required for different estimation qualities associated with different experiments may not be the same. This becomes problematic when the experiment that maximizes the FI measure requires a considerably weaker guarantee (i.e., it has a large value of tolerance) on the estimation cost. However, due to the sampling rule being agnostic to the tolerance levels, the greedy rule in (48) continues sampling the same experiment, even after the target estimation quality is achieved. This renders the greedy sampling rule prone to insufficient exploration. This is stated more formally in the following theorem.

Theorem 4. *There exists $\beta_i > 0$ for any experiment $i \in [K]$, such that any sequence of sequential procedure Δ^{c} that involves a Chernoff-based control action $\psi^{\mathsf{c}}(t)$ defined in (48) is sub-optimal.*

Proof. See Appendix H. ∎

This theorem shows that a greedy sampling strategy based on selecting the most informative experiment, along with any choice of stopping rule that satisfies the constraints in (16) at stopping, renders an infinite average sample complexity for specific choices of the estimation guarantees.

2) Cost-aware Sampling Rule: Next, we propose a sampling rule that maintains a balance between exploiting the control actions that maximize the return and exploring actions that have not yet been sufficiently sampled. This ensures that the algorithm does not get stuck in using only the most informative experiments. This is critically needed to avoid insufficient exploration of the less informative experiments, since this leads to significant disparity among different estimation qualities. The more explored experiments will be oversampled, achieving estimation qualities stronger than the prescribed thresholds. This penalizes the overall sample complexity of the sampling process. To circumvent such oversampling, we propose a sampling rule that involves randomly sampling from the distribution defined as

$$\bar{q}_t \triangleq \arg\min_{q\in\mathcal{Q}_K} \left\{ \frac{1}{\beta}\Big(\sum_{i\in[K]} q(i)\mathscr{J}_i(\nu_t^{\mathsf{ML}})\Big)^{-1} + \sum_{i\in[K]} \frac{1}{\beta_i}\Big(q(i)\mathscr{J}_i(\zeta_{i,t}^{\mathsf{ML}})\Big)^{-1}\right\}. \quad (49)$$

Note that $\bar{q}_t$ does not necessarily place the entire mass on *one* of the control actions, thus facilitating exploration. Furthermore, the distribution $\bar{q}_t$ converges to a limiting distribution in the limit of $t \to \infty$. This is attributed to the fact that the FI measures are computed at the ML estimates, and, by the strong consistency of the ML estimates [53], $\nu_t^{\mathsf{ML}} \xrightarrow{a.s.} \theta$, and $\zeta_{i,t}^{\mathsf{ML}} \xrightarrow{a.s.} \alpha_i$ for every $i \in [K]$. Thus, as $t \to \infty$, the ML estimates converge to the respective ground truths θ and $\{\alpha_i : i \in [K]\}$, and $\bar{q}_t$ converge to its limiting distribution. We denote this limiting distribution by

$$q^* \triangleq \arg\min_{q\in\mathcal{Q}_K} \left\{ \frac{1}{\beta}\Big(\sum_{i\in[K]} q(i)\mathscr{J}_i(\theta)\Big)^{-1} \quad (50)$$

$$+ \sum_{i\in[K]} \frac{1}{\beta_i}\Big(q(i)\mathscr{J}_i(\alpha_i)\Big)^{-1}. \quad (51)$$

C. Stopping Rule

We design a stopping rule that takes into account the fidelity guarantees on the shared parameter and the private parameters. Specifically, at each time instant, we compute the conditional posterior MMSE cost of the estimated shared parameter $\mathsf{C}(\nu_t^{\mathsf{MMSE}} \mid \mathcal{F}_t)$ and the private parameters $\{\mathsf{D}(\zeta_{i,t}^{\mathsf{MMSE}} \mid \mathcal{F}_t,\ \nu_t^{\mathsf{MMSE}}) : i \in [K]\}$. Our proposed stopping rule T is given by

$$T \triangleq \inf\Big\{ t\in\mathbb{N} :\ \mathsf{C}(\nu_t^{\mathsf{MMSE}} \mid \mathcal{F}_t) \le \beta, \ \mathsf{D}(\zeta_{i,t}^{\mathsf{MMSE}} \mid \mathcal{F}_t,\ \nu_t^{\mathsf{MMSE}}) \le \beta_i,\ \forall\, i\in[K]\Big\}. \quad (52)$$

Based on this, the sampling process terminates at the first instant at which all estimation cost constraints are satisfied.

D. Performance Guarantees

In this section, we present the optimality guarantees of the proposed active sampling procedure. We begin by stating a lower bound on the average sample complexity for any algorithm that solves $\mathcal{P}(\boldsymbol{\beta})$, defined in (16). Next, we provide an upper bound on the average sample complexity of the proposed sequential procedures. Specifically, we provide a high probability upper bound on the sample complexity as well as an asymptotic upper bound on the average sample complexity. The average sample complexity is shown to achieve the universal lower bound asymptotically up to a constant term. To characterize the average sample complexity, we define

$$W_{\boldsymbol{\beta}}(\theta,\boldsymbol{\alpha}) \triangleq \inf_{q\in\mathcal{Q}_K} \frac{1}{\beta}\Big(\sum_{i\in[K]} q(i)\mathscr{J}_i(\theta)\Big)^{-1} + \sum_{i\in[K]} \frac{1}{\beta_i}\Big(q(i)\mathscr{J}_i(\alpha_i)\Big)^{-1}. \quad (53)$$

Next, we define

$$V(\theta) \triangleq \frac{1}{\beta}\left(\sum_{i\in[K]} q^*(i)\mathscr{J}_i(\theta)\right)^{-1},$$
$$\text{and}\quad V_i(\alpha_i) \triangleq \frac{1}{\beta_i}\left(q^*(i)\mathscr{J}_i(\alpha_i)\right)^{-1},\ \forall\, i\in[K]. \quad (54)$$

Accordingly, we define define

$$V_{\max}(\theta,\boldsymbol{\alpha}) \triangleq \max\Big(V(\theta), V_1(\alpha_1),\cdots,V_K(\alpha_K)\Big),$$
$$\text{and}\quad V_{\min}(\theta,\boldsymbol{\alpha}) \triangleq \min\Big(V(\theta), V_1(\alpha_1),\cdots,V_K(\alpha_K)\Big). \quad (55)$$

Furthermore, define $\beta_{\max}$ and $\beta_{\min}$ as the maximum and minimum tolerance levels on the estimation costs, i.e.,

$$\beta_{\max} \triangleq \max\{\beta,\beta_1,\cdots,\beta_K\},\ \text{ and }\ \beta_{\min} \triangleq \min\{\beta,\beta_1\cdots,\beta_K\}. \quad (56)$$

Based on these definitions, we provide a lower bound on the average sample complexity.

Theorem 5 (Converse)**.** *Under assumptions A_1-A_6, for any sequential procedure Δ, and for any $h > 0$ there exist constants $C(h) > 0$ and $\{D_i(h) > 0 : i \in [K]\}$, such that for any $\beta \in (0, C(h))$ and $\beta_i \in (0, D_i(h))$ for all $i \in [K]$, we have*

$$\mathbb{E}[T] \ge \frac{W_{\boldsymbol{\beta}}(\theta,\boldsymbol{\alpha}) - h}{K+1}\cdot\Big(1-(K+1)h\Big). \quad (57)$$

Proof. See Appendix E. ∎

Theorem 6 (Achievability)**.** *Under assumptions A_1-A_6, for any $h > 0$ there exist constants $L(h) > 0$ and $\{M_i(h) > 0 : i \in [K]\}$, such that for any $\beta \in (0, L(h))$ and $\beta_i \in (0, M_i(h))$ for all $i \in [K]$, the proposed procedure comprised of the sequence of estimates in (44), sampling rule in (49) and the stopping rule in (52) satisfies*

$$\mathbb{P}\left\{ T \le \frac{W_{\boldsymbol{\beta}}(\theta,\boldsymbol{\alpha})}{K+1} + \Big(V_{\max}(\theta,\boldsymbol{\alpha}) - V_{\min}(\theta,\boldsymbol{\alpha})\Big)\right.$$

This article has been accepted for publication in IEEE Transactions on Signal Processing. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TSP.2022.3187655

$$+\left(\frac{1}{\beta_{\min}}+\frac{1}{\beta_{\max}}\right)h+2\Bigg\} \;=\; 1\;. \quad (58)$$

Proof. See Appendix F. ■

Theorem 7 (Achievability)**.** *In the asymptote of $\boldsymbol{\beta}\to\mathbf{0}$, there exists $h\in(0,1)$ such that the proposed procedure comprised of the sequence of estimates in (44), sampling rule in (49), and the stopping rule in (52) satisfies*

$$\lim_{\boldsymbol{\beta}\to\mathbf{0}}\frac{\mathbb{E}[T]}{W_{\boldsymbol{\beta}}(\theta,\boldsymbol{\alpha})}\;\leq\;\frac{1}{K+1}+h\;. \quad (59)$$

Proof. See Appendix G. ■

VI. Numerical Experiments

A. Shared Parameter

Consider a network of K sensors, where K is an even number. Sensor $i\in[K]$ generates its samples according to the distribution $\mathcal{N}(0,\sigma_i^2(\theta))$, where

$$\sigma_i^2(\theta)\triangleq\begin{cases}\frac{(i-1)^2}{K(K-i+1)}+\frac{K-2i+2}{K-i+1}\cdot\theta\;, & \forall i\in\left\{1,\ldots,\frac{K}{2}\right\}\\ \frac{i}{K}+\frac{K-2i}{i}\cdot\theta\;, & \forall i\in\left\{\frac{K}{2}+1,\ldots K\right\}\end{cases}. \quad (60)$$

The choice of the variance values for the Gaussian distribution corresponding to each sensor makes them the *most informative* one under a specific regime of the underlying shared parameter θ. More specifically, note that the FI measure for the Gaussian distribution $\mathcal{N}(0,\sigma^2)$ with respect to σ is $\mathcal{I}\triangleq 1/2\sigma^4$. It can be readily verified that each sensor maximizes the FI under the following regimes of θ: for $\theta\in((i-1)/K\;,\;i/K)$ for $i\in[K]$, sensor i maximizes the FI. In our evaluations, we set $K=4$. θ is assumed to have a uniform prior distribution $\mathsf{Unif}[0.01,0.99]$. It can be readily verified that if the true value of θ is less than 0.25, Sensor 1 is the most informative sensor. Otherwise, if $0.25<\theta<0.5$, Sensor 2 becomes the most informative one. Similarly, for $0.5<\theta<0.75$ and $0.75<\theta<1$, Sensor 3 and Sensor 4, respectively, become the most informative ones. Thus, an effective sampling rule is characterized by its ability to identify and converge to the best sensor using as few samples as possible. For our experiment, we set $\theta=0.2$.

Figure 2 shows the average number of samples required $\mathbb{E}[T]$ versus various levels of tolerance β, and compares them against the following four approaches for sensor selection.

1. *Random selection:* Random selection forms a baseline for comparison. Essentially, it refers to sampling one of the sensors S_1, S_2, S_3, or S_4 uniformly at random. The same stopping rule specified in (38) is used for fair comparison.
2. *Genie-aided sampling:* In this setting, we consider a genie-aided scenario in which the sampling rule is informed what the most informative sensor for estimating the unknown θ is.
3. *Approach of [40]:* The algorithm prescribed in [40] proposes a different stopping rule based on a fixed cost of sampling c, while keeping the same sensor-selection policy. The approach trades off the estimation performance against the accumulated cost of sampling, and thus does not have an explicit performance guarantee on the estimation cost (or a counterpart of β in our setting). For comparison, for any given β, we find out a value of c that ensures that the estimation cost of [40] falls below β, and use that to generate the variations of $\mathbb{E}[T]$ versus β.

The performance shown in Figure 2 correspond to averaging over 100 Monte Carlo realizations. Note that in Figure 2, Chernoff-like greedy sampling corresponds to the proposed sampling rule, which is also a special case of the look-ahead active sampling rule proposed in (49), for the case that there is only the shared parameter. This figure shows that our sampling rule outperforms the random sampling strategy, and the procedure prescribed in [40]. Note that the genie-aided sampling strategy uses the fewest number of samples to meet the prescribed level β. This is because this sampling strategy knows the most informative sensor from $t=1$ and always samples from that sensor. On the other hand, our proposed Chernoff-like greedy sampling rule requires a few more samples to guess the most informative sensor before it starts drawing samples from that sensor. Thus, its sample complexity is worse than that of the genie-aided strategy but better than those of all other strategies. Furthermore, while the approach of [40] uses the same sampling strategy like ours, its sample complexity suffers due to the choice of stopping rule. Specifically, the stopping rule in [40] is designed to minimize a unified objective comprising the estimation cost and delay, which is different from the objective in (11). Finally, the random sampling strategy puts equal sampling effort on each sensor, thus, requiring a more significant average number of samples to reach the same guarantee on the estimation cost.

Furthermore, to gain more insight regarding the scaling behavior of the proposed algorithm with respect to the number of sensors K, Figure 3 plots the average number of samples against K. Clearly, as the number of sensors increases, the number of samples required by the proposed strategy in identifying the most informative sensor increases, thus, increasing the average sample complexity. For this experiment, we set $\beta=0.005$, and the other parameters remain the same.

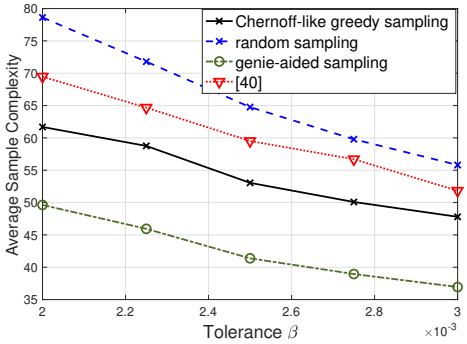

Fig. 2: Average sample complexity versus prescribed tolerance β.

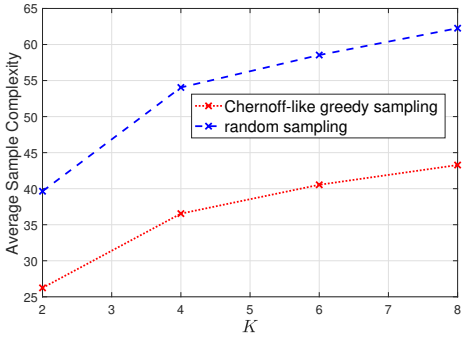

Fig. 3: Average sample complexity versus varying number of sensors K.

B. Shared and Private Parameters

The second experiment showcases the performance of the active sequential estimation algorithm proposed in Section V. We consider

This article has been accepted for publication in IEEE Transactions on Signal Processing. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TSP.2022.3187655

a simple network of two sensors, each generating random samples of θ contaminated by noise. Each sensor is subject to a different level of noise variance. Sensors 1 and 2 have distributions $\mathcal{N}(\theta, \alpha_1)$ and $\mathcal{N}(\theta, \alpha_2)$, respectively. We assume a uniform prior for the mean, i.e., for given $a_0, b_0 \in \mathbb{R}$,

$$\pi_\theta(\theta) = \frac{1}{b_0 - a_0} \mathbb{1}_{\{a_0 \le \theta \le b_0\}} , \quad b_0 > a_0 > 0 . \tag{61}$$

Similarly, the variance of each sensor has a uniform prior, i.e.,

$$\pi_i(\alpha_i) = \frac{1}{b_i - a_i} \mathbb{1}_{\{a_i \le \alpha_i \le b_i\}} , \quad b_i > a_i > 0 , \ i \in \{1, 2\} . \tag{62}$$

The posterior distributions of the unknown mean and variance are analytically intractable and they are computed numerically. The performance of our proposed procedure for this setup is depicted in Fig. 4. For this experiment, we have set $\theta = 0.25$, $\alpha_1 = 0.25$, and $\alpha_2 = 4$. The corresponding priors are parameterized by $a_0 = 0.1$, $b_0 = 4$, $a_1 = 0.1$, $b_1 = 0.7$, $a_2 = 1$, and $b_2 = 5$. The confidence levels on the estimates of the variance α_1 and α_2 are fixed at $\beta_1 = 0.1$ and $\beta_2 = 0.05$. The results show that the proposed sampling strategy outperforms the random selection strategy, which we use as a baseline in this case, as well as the greedy sampling strategy described in Section V-B1. This matches our theoretical analysis, where we prove in Appendix H that the Chernoff-based greedy sampling rule is sub-optimal. Intuitively, it is clear from the setting that such a sampling strategy focuses on exploiting Sensor 1 in the long run. However, that could result in an insufficient exploration of Sensor 2, resulting in a bad estimate of the variance of the second sensor.

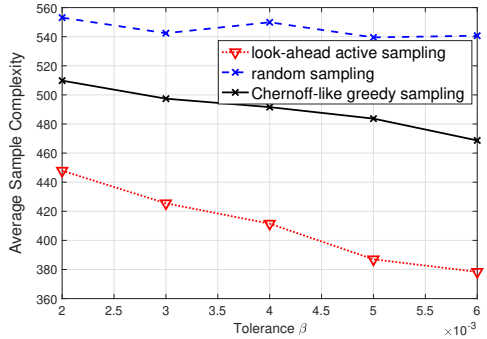

Fig. 4: Average sample complexity versus β, $\beta_1 = 0.05$, and $\beta_2 = 0.1$.

C. Real-world Data

In this section, we provide an experiment on a real-world dataset, comparing the performance of the proposed look-ahead sampling procedure against the random selection strategy. For this purpose, we use the Chicago beach weather station dataset, which consists of three different weather stations recording hourly measurements on various parameters, such as the air temperature, humidity, rain intensity, wind direction, and wind speed. We select two of these weather stations, namely the 63$^{\text{rd}}$ Street weather station and the Foster weather station. Our goal is to estimate the average air temperature in the month of September at 10 AM from noisy measurements of the temperature. We set the ground truth of the air temperature to the average temperature recorded by the two weather stations in September at 10 AM over the years 2019, 2020, and 2021. The true variance for each weather station is also set to the variance computed from the data over these years. It is noteworthy that the Foster weather station has a larger variance of 6.7975 in recorded temperatures, compared to the 63$^{\text{rd}}$ Street weather station, which has a variance of 4.1216. We consider the setting in which the variance of the weather

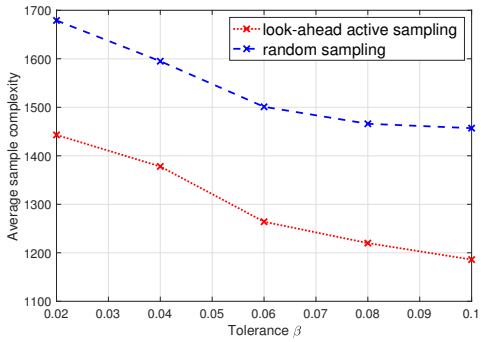

Fig. 5: Average sample complexity versus prescribed tolerance β.

station sensors is also unknown (in addition to the temperature that we intend to estimate). Each station is assumed to generate noisy measurements of the temperature, drawn from a Gaussian distribution with mean and variance set as specified. We set the levels required on the cost of estimating the variance for both the weather stations to $\beta_1 = \beta_2 = 0.1$. The prior distributions of the mean values are assumed to be $\mathsf{Unif}[10, 20]$, while those of the variances are assumed to be $\mathsf{Unif}[2, 8]$ for both the weather stations. Figure 5 demonstrates the average sample complexity against various levels of β. We observe that our proposed sequential procedure yields a far superior performance compared to random sampling, thus, clearly depicting the advantage of our proposed active sampling procedure.

VII. Conclusions

In this paper, we have investigated the problem of active sequential estimation. Specifically, we have considered a setup in which the problem instance may have both shared as well as private parameters, and the goal is to estimate each of the parameters while meeting a prescribed level of confidence on their estimation qualities. Sequential decision rules are proposed, where the sampling rule involves maximizing the Fisher Information measure when the setting comprises shared parameters only (greedy sampling rule), while in the setting of both private and shared parameters, the rule maintains a trade-off between exploiting the most informative action and exploring the scarcely sampled actions in order to meet the prescribed tolerances on the estimation quality of every parameter. The proposed decision rules have been shown to be asymptotically optimal, and we have provided numerical experiments to evaluate their advantage compared to the existing strategies.

Appendix A
Useful Lemmas

In this section, we provide a few lemmas which we will be using throughout our analysis.

Lemma 1 (Lemma 4.1, [40])**.** *When we have only a shared parameter (setting of Section IV), under assumptions A_1-A_6, for any $h > 0$, there exists $t(h) \in \mathbb{N}$ such that for all $t > t(h)$:*

$$\mathbb{P}\Big\{ t \cdot \inf_{\hat{\theta}_t,\, \psi^t} \mathsf{C}(\hat{\theta}_t \mid \mathcal{F}_t) \ \ge \ \beta V_\beta(\theta) - h \Big\} \ > \ 1 - h . \tag{63}$$

Lemma 2 (Theorem 4.3, [40])**.** *Under assumptions A_1-A_6, the sequences of estimates generated by ν_t^{MMSE}, which is characterized in Section IV-A, and sampling rules ψ^t specified in Section IV-B satisfy*

$$t \cdot \mathsf{C}(\nu_t^{\mathsf{MMSE}} \mid \mathcal{F}_t) \xrightarrow{t \to \infty} \beta V_\beta(\theta) . \tag{64}$$

Lemma 3. *When we have both shared and private parameters (the setting in Section V), under assumptions A_1-A_6, for any $h > 0$, there exists $t^*(h) < +\infty$ such that for any $t > t^*(h)$:*

$$\mathbb{P}\Bigg\{ \inf_{\Phi_t, \psi^t} \; t\Bigg[\frac{1}{\beta}\mathsf{C}(\hat{\theta}_t|\mathcal{F}_t) + \sum_{i\in[K]} \frac{1}{\beta_i}\mathsf{D}(\hat{\alpha}_{i,t}|\mathcal{F}_t)\Bigg] \;\geq\; W_{\boldsymbol{\beta}}(\theta, \boldsymbol{\alpha}) - h \Bigg\} \;>\; 1 - (K+1)h \; . \tag{65}$$

Proof. Let us denote the number of times experiment $i \in [K]$ is selected up to time $t \in \mathbb{N}$ by

$$T_i(t) \;\triangleq\; \sum_{s=1}^{t} \mathbb{1}_{\{\psi(s)=i\}} \; . \tag{66}$$

From the Cramér-Rao lower bound, we obtain the following inequalities.

$$\mathbb{P}\Big\{ \mathsf{C}(\hat{\theta}_t \mid \mathcal{F}_t) \;\geq\; \Big(\sum_{i\in[K]} T_i(t)\, \mathscr{J}_i(\theta) \Big)^{-1} \Big\} \;=\; 1 \; , \tag{67}$$

$$\mathbb{P}\Big\{ \mathsf{D}(\hat{\alpha}_{i,t} \mid \mathcal{F}_t) \;\geq\; \Big(T_i(t)\, \mathscr{J}_i(\alpha_i) \Big)^{-1} \Big\} \;=\; 1, \quad \forall i \in [K] \; . \tag{68}$$

Furthermore, let us denote the fraction of times any experiment $i \in [K]$ is selected up to time t by

$$p_i(t) \;\triangleq\; \frac{1}{t} \sum_{s=1}^{t} \mathbb{1}_{\{\psi(s)=i\}} \; . \tag{69}$$

Subsequently, from (67) and (68), we obtain that

$$\mathbb{P}\Big\{ t \cdot \mathsf{C}(\hat{\theta}_t \mid \mathcal{F}_t) \;\geq\; \Big(\sum_{i\in[K]} p_i(t)\, \mathscr{J}_i(\theta) \Big)^{-1} \Big\} \;=\; 1 \; , \tag{70}$$

$$\mathbb{P}\Big\{ t \cdot \mathsf{D}(\hat{\alpha}_{i,t} \mid \mathcal{F}_t) \;\geq\; \Big(p_i(t)\, \mathscr{J}_i(\alpha_i) \Big)^{-1} \Big\} = \; 1, \quad \forall i \in [K] \; . \tag{71}$$

Using the same argument as in [40, Lemma 4.1], any measurable control action will have a limiting distribution, or will be arbitrarily close to a limiting distribution such that for any arbitrarily small $h > 0$, there exists $t(h) < +\infty$ such that for every $t > t(h)$, from (70) we obtain

$$\mathbb{P}\Big\{ t \cdot \mathsf{C}(\hat{\theta}_t \mid \mathcal{F}_t) \;\geq\; \Big(\sum_{i\in[K]} p(i)\, \mathscr{J}_i(\theta) \Big)^{-1} - h \Big\} \;>\; 1 - h \; , \tag{72}$$

where $p(i)$ denotes the limiting distribution of $p_i(t)$. Using similar arguments, for every $i \in [K]$ and any $h > 0$, there exists $t_i'(h) < +\infty$ such that for all $t \geq t_i'(h)$:

$$\mathbb{P}\Big\{ t \cdot \mathsf{D}(\hat{\alpha}_{i,t} \mid \mathcal{F}_t) \;\geq\; \Big(p(i)\, \mathscr{J}_i(\alpha_i) \Big)^{-1} - h \Big\} \;>\; 1 - h \; . \tag{73}$$

By defining $t^*(h) \triangleq \max\{t(h), t_1'(h), \cdots, t_K'(h)\}$, combining (72) and (73) and taking the infimum with respect to all control actions and estimates, we obtain that

$$\mathbb{P}\Bigg\{ \inf_{\Phi_t, \psi^t} \; t\Bigg[\frac{1}{\beta}\mathsf{C}(\hat{\theta}_t|\mathcal{F}_t) + \sum_{i\in[K]} \frac{1}{\beta_i}\mathsf{D}(\hat{\alpha}_{i,t}|\mathcal{F}_t)\Bigg] \geq W_{\boldsymbol{\beta}}(\theta, \boldsymbol{\alpha}) - h \Bigg\} > 1 - (K+1)h \; , \tag{74}$$

for all $t > t^*(h)$, where (74) is obtained from the fact that for any two events A and B, $\mathbb{P}(A \cap B) \geq \mathbb{P}(A) + \mathbb{P}(B) - 1$. This completes our proof. ∎

Lemma 4. *When we have both shared and private parameters (the setting in Section V), under assumptions A_1-A_6, the sequences of estimates characterized in Section V-A and sampling rules specified in Section V-B satisfy*

$$t\Bigg[\frac{1}{\beta}\mathsf{C}(\nu_t^{\mathsf{MMSE}}|\mathcal{F}_t) + \sum_{i\in[K]} \frac{1}{\beta_i}\mathsf{D}(\zeta_{i,t}^{\mathsf{MMSE}}|\mathcal{F}_t, \nu_t^{\mathsf{MMSE}})\Bigg] \xrightarrow{t\to\infty} W_{\boldsymbol{\beta}}(\theta, \boldsymbol{\alpha}). \tag{75}$$

Proof. We denote the frequency of selecting experiment $i \in [K]$ under the sampling rule ψ^t by

$$p_i(t) \;\triangleq\; \frac{1}{t} \sum_{s=1}^{t} \mathbb{1}_{\{\psi(s)\,=\,i\}} \; . \tag{76}$$

Then, as $t \to \infty$, $p_i(t) \to \bar{q}_t(i)$ due to the law of large numbers. Furthermore, due to the strong consistency of ML estimates, ν_t^{ML} converges to θ, and $\zeta_{i,t}^{\mathsf{ML}}$ converges to α_i for all $i \in [K]$ [53].

Now, we will use a weak version of the Bernstein-Von-Mises theorem to establish the asymptotic normality of ν_t^{MMSE}. For this, let us consider a sequence of observations $\mathbf{Y}^t \triangleq \{Y_1, \cdots, Y_t\}$, which are not necessarily i.i.d. Let θ be an unknown parameter of the underlying pdfs of the observations $\mathbf{Y}^t$, and $\lambda(\cdot \mid \theta)$ denote the log-likelihood function corresponding to the observations. Next, we define the Fisher Information (FI) measure as

$$\mathcal{J}(\theta) \triangleq -\frac{1}{t}\mathbb{E}\left[\frac{\partial^2}{\partial\theta^2}\lambda(\mathbf{Y}^t \mid \theta)\right] \; . \tag{77}$$

Furthermore, define ν_t^{MMSE} as the MMSE estimator of θ, and define $z \triangleq \sqrt{t}(\nu_t^{\mathsf{MMSE}} - \theta)$. By a weak version of the Bernstein-Von-Mises theorem [54, Theorem 20.2], we have

$$\vartheta\left(z \mid \mathcal{F}_t\right) \xrightarrow{t\to\infty} \mathcal{N}(0, 1/\mathcal{J}(\theta)) \; , \tag{78}$$

where $\vartheta(z \mid \mathcal{F}_t)$ represents the posterior distribution of z. The key is to find $\mathcal{J}(\theta)$ in our case. We have,

$$\mathcal{J}(\theta) = -\frac{1}{t}\mathbb{E}\left[\frac{\partial^2}{\partial\theta^2}\lambda(\mathbf{Y}^t \mid \theta)\right] \tag{79}$$

$$= -\frac{1}{t}\mathbb{E}\left[\frac{\partial^2}{\partial\theta^2}\sum_{i=1}^{K}\lambda_i(\mathbf{Y}_i^t \mid \theta)\right] \; , \tag{80}$$

where $\mathbf{Y}_i^t \triangleq \{Y_s : \psi(s) = i, s \in [t]\}$ is the sequence of observations from the experiment $i \in [K]$. Simplifying (80), we obtain

$$\mathcal{J}(\theta) = \frac{1}{t}\sum_{i=1}^{K} t p_i(t)\, \mathscr{J}_i(\theta) = \sum_{i=1}^{K} p_i(t)\, \mathscr{J}_i(\theta) \; . \tag{81}$$

Thus, we have

$$\vartheta\big(z|\mathcal{F}_t\big) \xrightarrow{t\to\infty} \mathcal{N}\Bigg(0, \Big(\sum_{i\in[K]} p_i(t)\, \mathscr{J}_i(\theta) \Big)^{-1}\Bigg) \; . \tag{82}$$

Furthermore, since $p(t) \triangleq [p_1(t), \cdots, p_K(t)]$ converges to the limiting distribution q^* defined in (51), we have

$$\vartheta\big(z|\mathcal{F}_t\big) \xrightarrow{t\to\infty} \mathcal{N}\Bigg(0, \Big(\sum_{i\in[K]} q^*(i)\, \mathscr{J}_i(\theta) \Big)^{-1}\Bigg) \; . \tag{83}$$

Thus, from (83), we have $\nu_t^{\mathsf{MMSE}} \xrightarrow{p} \theta$. Using this fact in conjunction with the Bernstein-Von Mises theorem, for all $i \in [K]$, we have

$$\omega_i\big(u_i \mid \nu_t^{\mathsf{MMSE}}, \mathcal{F}_t\big) \xrightarrow{t\to\infty} \mathcal{N}\Big(0, \big(q^*(i)\, \mathscr{J}_i(\alpha_i) \big)^{-1}\Big) \; , \tag{84}$$

where we have defined $u_i \triangleq \sqrt{t}(\zeta_{i,t}^{\mathsf{MMSE}} - \alpha_i)$ and $\omega_i(u_i \mid \nu_t^{\mathsf{MMSE}}, \mathcal{F}_t)$ represents the posterior distribution of u_i. The limit of the sequence of MMSE is then calculated by establishing the finiteness of the first and second order moments of $\vartheta(z \mid \mathcal{F}_t)$ and $\omega_i(u_i \mid \nu_t^{\mathsf{MMSE}}, \mathcal{F}_t)$. For $\vartheta(z \mid \mathcal{F}_t)$, this can be done following the same approach of [40,

This article has been accepted for publication in IEEE Transactions on Signal Processing. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TSP.2022.3187655

Lemma 7.7]. For $\omega_i(u_i \mid \mathcal{F}_t)$, this can be done following the same approach of [24, Lemma 3.1]. This yields

$$t \cdot \mathsf{C}(\nu_t^{\mathsf{MMSE}} \mid \mathcal{F}_t) \xrightarrow{t\to\infty} \Big(\sum_{i\in[K]} q^*(i) \mathscr{I}_i(\theta) \Big)^{-1}, \tag{85}$$

and $\quad t \cdot \mathsf{D}(\zeta_{i,t}^{\mathsf{MMSE}} \mid \mathcal{F}_t,\ \nu_t^{\mathsf{MMSE}}) \xrightarrow{t\to\infty} \Big(q^*(i) \mathscr{I}_i(\alpha_i) \Big)^{-1} \quad \forall i \in [K]. \tag{86}$

Finally, combining (85) and (86), we obtain

$$t\Big[\frac{1}{\beta} \mathsf{C}(\nu_t^{\mathsf{MMSE}} | \mathcal{F}_t) + \sum_{i\in[K]} \frac{1}{\beta_i} \mathsf{D}(\zeta_{i,t}^{\mathsf{MMSE}} | \mathcal{F}_t, \nu_t^{\mathsf{MMSE}}) \Big] \xrightarrow{t\to\infty} W_{\boldsymbol{\beta}}(\theta, \boldsymbol{\alpha}). \tag{87}$$

This concludes our proof. ∎

Appendix B
Proof of Theorem 1

Consider a constant $L \in \mathbb{N}$. Using the constraint in $\mathcal{P}(\beta)$ defined in (11), we have

$$\begin{aligned} 1 \ &= \ \mathbb{P}\Big\{ \mathsf{C}(\hat{\theta}_T \mid \mathcal{F}_T) \ \le \ \beta \Big\} && (88)\\ &= \ \mathbb{P}\Big\{ \mathsf{C}(\hat{\theta}_T \mid \mathcal{F}_T) \ \le \ \beta \ , \ T \ \le \ L \Big\} \\ &\quad + \ \mathbb{P}\Big\{ \mathsf{C}(\hat{\theta}_T \mid \mathcal{F}_T) \ \le \ \beta \ , \ T \ > \ L \Big\} . && (89) \end{aligned}$$

Define $T_L(i) \ \triangleq \ \sum_{s=1}^{L} \mathbb{1}_{\{\psi(s)=i\}}$ as the number of times that experiment $i \in [K]$ is chosen up to time L. Then, we have

$$\begin{aligned} \mathsf{C}(\hat{\theta}_T \mid \mathcal{F}_T) \ &\ge \ \Big(\sum_{i\in[K]} T_T(i) \mathscr{I}_i(\theta) \Big)^{-1} && (90)\\ &\ge \ \frac{1}{T} \Big(\sum_{i\in[K]} \mathscr{I}_i(\theta) \Big)^{-1} && (91)\\ &\ge \ \frac{1}{L} \Big(\sum_{i\in[K]} \mathscr{I}_i(\theta) \Big)^{-1} && (92)\\ &= \ C(L) , && (93) \end{aligned}$$

where the first inequality is a result of the Cramér-Rao lower bound [55], and $C(L)$ is a positive constant. Choosing $\beta \in (0, C(L))$ ensures that the first term in (89) becomes $\mathbb{P}\Big\{ \mathsf{C}(\hat{\theta}_T \mid \mathcal{F}_T) \ \le \ \beta \ , \ T \ \le \ L \Big\} = 0$. Using Lemma 1, there exists $t(h)$ such that for all $t > t(h)$, we have

$$\mathbb{P}\Big\{ t \cdot \mathsf{C}(\hat{\theta}_t \mid \mathcal{F}_t) \ \ge \ \beta V_\beta(\theta) - h \Big\} \ > \ 1 - h \, . \tag{94}$$

Leveraging (94), let us expand the the second term in (89). Choosing $L = t(h)$, we have

$$\begin{aligned} &\mathbb{P}\Big\{ \mathsf{C}(\hat{\theta}_T \mid \mathcal{F}_T) \ \le \ \beta \ , \ T \ > \ t(h) \Big\} \\ &= \ \mathbb{P}\Big\{ T\mathsf{C}(\hat{\theta}_T | \mathcal{F}_T) \le T\beta, T > t(h) | T\mathsf{C}(\hat{\theta}_T | \mathcal{F}_T) \ge \beta V_\beta(\theta) - h \Big\} \\ &\quad \times \ \mathbb{P}\Big\{ T \cdot \mathsf{C}(\hat{\theta}_T \mid \mathcal{F}_T) \ \ge \ \beta V_\beta(\theta) - h \Big\} \\ &+ \mathbb{P}\Big\{ T\mathsf{C}(\hat{\theta}_T, \theta | \mathcal{F}_T) \le T\beta, T > t(h) | T\mathsf{C}(\hat{\theta}_T | \mathcal{F}_T) < \beta V_\beta(\theta) - h \Big\} \\ &\quad \times \ \mathbb{P}\Big\{ T \cdot \mathsf{C}(\hat{\theta}_T \mid \mathcal{F}_T) \ < \ \beta V_\beta(\theta) - h \Big\} && (95)\\ &\le \ \mathbb{P}\Big\{ T\mathsf{C}(\hat{\theta}_T | \mathcal{F}_T) \le T\beta, T\mathsf{C}(\hat{\theta}_T | \mathcal{F}_T) \ge \beta V_\beta(\theta) - h \Big\} + h && (96)\\ &\le \ \mathbb{P}\Big\{ T \ \ge \ \frac{\beta V_\beta(\theta) - h}{\beta} \Big\} \ + \ h \, , && (97) \end{aligned}$$

where (95) follows from the law of total probability, and (97) is a result of the stopping rule in (38). Thus, combining (89) with (97), and due to the choice of $\beta < C(h)$, we obtain

$$\mathbb{P}\Big\{ T \ \ge \ \frac{\beta V_\beta(\theta) - h}{\beta} \Big\} \ \ge \ 1 \ - \ h \, , \tag{98}$$

where we have defined

$$C(h) \ \triangleq \ \frac{1}{t(h)} \Big(\sum_{i\in[K]} \mathscr{I}_i(\theta) \Big)^{-1} . \tag{99}$$

The result readily follows by applying Markov's inequality to (98).

Appendix C
Proof of Theorem 2

Using Lemma 2, we obtain that for any $h > 0$ there exists $t(h) < +\infty$ such that for any $t > t(h)$, we have

$$\mathbb{P}\big\{ t \cdot \mathsf{C}(\nu_t^{\mathsf{MMSE}} \mid \mathcal{F}_t) \ \le \ \beta V_\beta(\theta) \ + \ h \big\} \ = \ 1 \, . \tag{100}$$

Now, observe that if $T - 1 > t(h)$,

$$\begin{aligned} 1 \ &= \ \mathbb{P}\big\{ (T-1) \cdot \mathsf{C}(\nu_{T-1}^{\mathsf{MMSE}} \mid \mathcal{F}_{T-1}) \ \le \ \beta V_\beta(\theta) \ + \ h \big\} && (101)\\ &\le \ \mathbb{P}\big\{ (T-1)\beta \ \le \ \beta V_\beta(\theta) \ + \ h \big\} && (102)\\ &= \ \mathbb{P}\Big\{ T \ \le \ \frac{\beta V_\beta(\theta) + h}{\beta} + 1 \Big\} , && (103) \end{aligned}$$

where (102) is a result of the stopping rule in (38). Furthermore,

$$\begin{aligned} &\mathbb{P}\Big\{ T \ \le \ \frac{\beta V_\beta(\theta) + h}{\beta} + 1 \Big\} \\ &= \ \mathbb{P}\Big\{ T \ \le \ \frac{\beta V_\beta(\theta) + h}{\beta} + 1, T - 1 \ > \ t(h) \Big\} \\ &+ \ \mathbb{P}\Big\{ T \ \le \ \frac{\beta V_\beta(\theta) + h}{\beta} + 1 \ \Big| \ T - 1 \ \le \ t(h) \Big\} \mathbb{P}\big(T - 1 \ \le \ t(h) \big) \end{aligned} \tag{104}$$

Now, $\mathbb{P}(T - 1 \ \le \ t(h)) = \mathbb{P}\{ \mathsf{C}(\nu_T^{\mathsf{MMSE}} \mid \mathcal{F}_T) \le \beta, T \le t(h) + 1 \}$. Following the same steps as in (90)-(93), we can show that there exists a constant $C'(h) \ \triangleq \ (\sum_{i\in[K]} \mathscr{I}_i(\theta))^{-1} / (t(h) + 1)$, such that $\mathbb{P}(T - 1 \ \le \ t(h)) = 0$ for any $\beta \in (0, C'(h))$. Finally, combining (103) and (104), we obtain that for any $\beta \in (0, C'(h))$,

$$\mathbb{P}\Big\{ T \ \le \ \frac{\beta V_\beta(\theta) + h}{\beta} + 1 \Big\} \ = \ 1 \, . \tag{105}$$

This completes the proof.

Appendix D
Proof of Theorem 3

Using Lemma 2, for any $\epsilon > 0$, there exists $T_\epsilon < +\infty$, such that for all $t \ge T_\epsilon$,

$$t\mathsf{C}(\nu_t^{\mathsf{MMSE}} \mid \mathcal{F}_t) \ \in \ [\beta V_\beta(\theta) - \epsilon, \beta V_\beta(\theta) + \epsilon] \, . \tag{106}$$

Now, at the instant before stopping we have

$$\begin{aligned} T - 1 \ &= \ (T-1)\mathbb{1}_{\{T-1 \le T_\epsilon\}} + (T-1)\mathbb{1}_{\{T-1 > T_\epsilon\}} && (107)\\ &\le T_\epsilon + \frac{\beta V_\beta(\theta) + \epsilon}{\beta} + 1 \, , && (108) \end{aligned}$$

where (108) follows from the definition of the stopping rule and (106). Furthermore, note that $\sup_{Y^t, \psi^t} T_\epsilon < \ + \infty$ owing to Lemma 2. Thus, taking average throughout (108), dividing by $1/\beta$, and taking the limit of $\beta \to 0$, we obtain

$$\lim_{\beta\to 0} \frac{\mathbb{E}[T]}{V_\beta(\theta)} \ \le \ 1 \, . \tag{109}$$

This article has been accepted for publication in IEEE Transactions on Signal Processing. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TSP.2022.3187655

APPENDIX E
PROOF OF THEOREM 5

Define the event

$$\mathcal{S}_t \triangleq \Big\{ \mathsf{C}(\hat{\theta}_t \mid \mathcal{F}_t) \le \beta,\ \mathsf{D}(\hat{\alpha}_{i,t} \mid \mathcal{F}_t,\ \hat{\theta}_t) \le \beta_i\ \forall\, i \in [K] \Big\}. \tag{110}$$

From the constraint in (16), for any constant $L < +\infty$, we have

$$1 = \mathbb{P}\{\mathcal{S}_T\} = \mathbb{P}\{\mathcal{S}_T,\ T \le L\} + \mathbb{P}\{\mathcal{S}_T,\ T > L\}. \tag{111}$$

By the Cramér-Rao lower bound, if $T \le L$,

$$\mathsf{C}(\hat{\theta}_T \mid \mathcal{F}_T) \ge \frac{1}{L}\Big(\sum_{i\in[K]} \mathscr{J}_i(\theta) \Big)^{-1}. \tag{112}$$

Similarly, for any $i \in [K]$ and $T \le L$, we have

$$\mathsf{D}(\hat{\alpha}_{i,T} \mid \mathcal{F}_T,\ \hat{\theta}_T) \ge \frac{1}{L}\big(\mathscr{J}_i(\alpha_i \mid \hat{\theta}_T) \big)^{-1} \tag{113}$$

$$\ge \inf_{\theta\in\Theta} \frac{1}{L}\big(\mathscr{J}_i(\alpha_i \mid \theta) \big)^{-1}, \tag{114}$$

where we have defined:

$$\mathscr{J}_i(\alpha_i \mid \theta) \triangleq \mathbb{E}\left[\frac{\partial^2}{\partial \alpha_i^2} \lambda_i(x \mid \theta, \alpha_i) \right]. \tag{115}$$

Define the quantities

$$C(L) \triangleq \frac{1}{L}\left(\sum_{i\in[K]} \mathscr{J}_i(\theta) \right)^{-1},$$

$$\text{and} \quad D_i(L) \triangleq \inf_{\theta\in\Theta} \frac{1}{L}\left(\mathscr{J}_i(\alpha_i \mid \theta) \right)^{-1}. \tag{116}$$

Similarly to (90)-(93), choosing $\beta \in (0, C(L))$ and $\beta_i \in (0, D_i(L))$ for every $i \in [K]$,

$$\mathbb{P}\big(\mathcal{S}_T,\ T \le L\big) = 0. \tag{117}$$

Expanding the second term on the right hand side of (111), we obtain

$$\begin{aligned}
&\mathbb{P}\{\mathcal{S}_T,\ T > L\} \\
&= \mathbb{P}\{T \cdot \mathsf{C}(\hat{\theta}_T \mid \mathcal{F}_T) \le T\beta, \\
&\qquad T \cdot \mathsf{D}(\hat{\alpha}_{i,T} \mid \mathcal{F}_T,\ \hat{\theta}_T) \le T\beta_i\ \forall\, i \in [K],\ T > L\} \qquad (118) \\
&\le \mathbb{P}\{T \cdot \mathsf{C}(\hat{\theta}_T \mid \mathcal{F}_T) \le T\beta, \\
&\qquad T \cdot \mathsf{D}(\hat{\alpha}_{i,T} \mid \mathcal{F}_T) \le T\beta_i\ \forall\, i \in [K],\ T > L\} \qquad (119) \\
&\le \mathbb{P}\left\{ \left[\frac{T\mathsf{C}(\hat{\theta}_T \mid \mathcal{F}_T)}{\beta} + \sum_{i\in[K]} \frac{T\mathsf{D}(\hat{\alpha}_{i,T} \mid \mathcal{F}_T)}{\beta_i} \right] \le (K+1)T, \right. \\
&\qquad \left. T > L \right\}, \qquad (120)
\end{aligned}$$

where (119) follows from the fact that since $\mathsf{D}(\hat{\alpha}_{i,T} \mid \mathcal{F}_T,\ \hat{\theta}_T) \le \beta_i$ at stopping, we can take an expectation over $\mathsf{D}(\hat{\alpha}_{i,T} \mid \mathcal{F}_T,\ \theta)$ with respect to $\mathbb{E}_t$ and obtain $\mathsf{D}(\hat{\alpha}_{i,T} \mid \mathcal{F}_T) \le \beta_i$, i.e., $\{\mathsf{D}(\hat{\alpha}_{i,T} \mid \mathcal{F}_T,\ \hat{\theta}_T) \le \beta_i\} \subseteq \{\mathsf{D}(\hat{\alpha}_{i,T} \mid \mathcal{F}_T) \le \beta_i\}$. Furthermore, for any $h > 0$, define the event

$$\mathcal{E}_t \triangleq \left\{ t\left[\frac{\mathsf{C}(\hat{\theta}_t \mid \mathcal{F}_t)}{\beta} + \sum_{i\in[K]} \frac{\mathsf{D}(\hat{\alpha}_{i,t} \mid \mathcal{F}_t)}{\beta_i} \right] \ge W_{\boldsymbol{\beta}}(\theta, \boldsymbol{\alpha}) - h \right\} \tag{121}$$

Next, we will use Lemma 3 and set $L = t^*(h)$. Expanding (120) we have

$$\begin{aligned}
&\mathbb{P}\left\{ \left[\frac{T\mathsf{C}(\hat{\theta}_T \mid \mathcal{F}_T)}{\beta} + \sum_{i\in[K]} \frac{T\mathsf{D}(\hat{\alpha}_{i,T} \mid \mathcal{F}_T)}{\beta_i} \right] \le (K+1)T, \right. \\
&\qquad \left. T > t^*(h) \right\} \\
&= \mathbb{P}\left\{ \left[\frac{T\mathsf{C}(\hat{\theta}_T \mid \mathcal{F}_T)}{\beta} + \sum_{i\in[K]} \frac{T\mathsf{D}(\hat{\alpha}_{i,T} \mid \mathcal{F}_T)}{\beta_i} \right] \le (K+1)T, \right. \\
&\qquad \left. T > t^*(h),\ \mathcal{E}_T \right\} \\
&+ \mathbb{P}\left\{ \left[\frac{T\mathsf{C}(\hat{\theta}_T \mid \mathcal{F}_T)}{\beta} + \sum_{i\in[K]} \frac{T\mathsf{D}(\hat{\alpha}_{i,T} \mid \mathcal{F}_T)}{\beta_i} \right] \le (K+1)T, \right. \\
&\qquad \left. T > t^*(h) \,\middle|\, \overline{\mathcal{E}_T} \right\} \cdot \mathbb{P}\big(\overline{\mathcal{E}_T}\big) \qquad (122) \\
&\le \mathbb{P}\Big\{ W_{\boldsymbol{\beta}}(\theta, \boldsymbol{\alpha}) - h \le (K+1)T \Big\} + (K+1)h, \qquad (123)
\end{aligned}$$

where (122) is a result of applying the law of total probability, and (123) is a result of Lemma 3. Finally, (123) can be rewritten using (111) and (117), as follows.

$$\mathbb{P}\left\{ T \ge \frac{W_{\boldsymbol{\beta}}(\theta, \boldsymbol{\alpha}) - h}{K+1} \right\} > 1 - h. \tag{124}$$

Subsequently, applying the Markov's inequality yields

$$\mathbb{E}[T] \ge \frac{W_{\boldsymbol{\beta}}(\theta, \boldsymbol{\alpha}) - h}{K+1} \cdot \Big(1 - (K+1)h \Big). \tag{125}$$

APPENDIX F
PROOF OF THEOREM 6

Using Lemma 4, we have the following two convergence properties.

$$t \cdot \mathsf{C}\Big(\nu_t^{\mathsf{MMSE}} \mid \mathcal{F}_t \Big) \xrightarrow{t\to\infty} \beta V(\theta),$$

$$\text{and} \quad t \cdot \mathsf{D}(\zeta_{i,t}^{\mathsf{MMSE}} \mid \mathcal{F}_t,\ \nu_t^{\mathsf{MMSE}}) \xrightarrow{t\to\infty} \beta_i V_i(\alpha_i). \tag{126}$$

Thus, for any $h > 0$, there exists $t(h) < +\infty$ such that for all $t > t(h)$,

$$\mathbb{P}\left\{ t \cdot \mathsf{C}(\nu_t^{\mathsf{MMSE}} \mid \mathcal{F}_t) \le \beta V(\theta) + h \right\} = 1, \tag{127}$$

$$\mathbb{P}\left\{ t \cdot \mathsf{D}(\zeta_{i,t}^{\mathsf{MMSE}} \mid \mathcal{F}_t,\ \nu_t^{\mathsf{MMSE}}) \le \beta_i V_i(\alpha_i) + h \right\} = 1. \tag{128}$$

Define $(T - \tau_i)$ as the i^{th} time instant at which the guarantee on one of the parameters is met. Furthermore, define $S(i) \in [K]$ as the experiment to which the parameter belongs, in case that it is a private parameter. Thus, $\tau_{K+1} = 0$. Furthermore, using (127) and (128), we obtain that for any $h > 0$ and $T - \tau_i - 1 \ge t(h)$,

$$\mathbb{P}\left\{ T - \tau_i \le \frac{V_{S(i)}(\alpha_{S(i)}) + h}{\beta_{S(i)}} + 1 \right\} = 1, \tag{129}$$

$$\text{or,} \quad \mathbb{P}\left\{ T - \tau_i \le \frac{\beta V(\theta) + h}{\beta} + 1 \right\} = 1, \tag{130}$$

depending on whether the parameter on which the tolerance guarantee is achieved is shared or private. Combining (129) and (130), we obtain:

$$\mathbb{P}\left\{ (K+1)T \le \frac{\beta V(\theta) + h}{\beta} + \sum_{i\in[K]} \frac{\beta_i V_i(\alpha_i) + h}{\beta_i} \right.$$

This article has been accepted for publication in IEEE Transactions on Signal Processing. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TSP.2022.3187655

$$+\sum_{i=1}^{K}\tau_i+(K+1)\Bigg\} \;=\; 1\;. \tag{131}$$

Furthermore, define $\beta_{\min} \triangleq \min\{\beta,\beta_1,\cdots,\beta_K\}$ and $\beta_{\max} \triangleq \max\{\beta,\beta_1,\cdots,\beta_K\}$. Note that using (127) and (128), along with the definition of $(t-\tau_i)$, for $t-\tau_i-1>t(h)$, we have

$$\mathbb{P}\Bigg\{T-\tau_i\le V_{\max}(\theta,\boldsymbol{\alpha})+\frac{h}{\beta_{\min}}+1\Bigg\} \;=\; 1\;. \tag{132}$$

Similarly, (126) combined with the definition of $(T-\tau_j)$ for any experiment $S(j)\in[K]$ and for $t-\tau_j>t(h)$, we have

$$\mathbb{P}\Bigg\{T-\tau_j\ge\frac{\beta_{S(j)}V_{S(j)}(\alpha_{S(j)})-h}{\beta_{S(j)}}\Bigg\} \;=\; 1\;, \tag{133}$$

$$\text{or,}\quad \mathbb{P}\Bigg\{T-\tau_j\ge\frac{\beta V(\theta)-h}{\beta}\Bigg\} \;=\; 1\;, \tag{134}$$

resulting in

$$\mathbb{P}\Bigg\{T-\tau_j\ge V_{\min}(\theta,\boldsymbol{\alpha})+\frac{h}{\beta_{\max}}\Bigg\} \;=\; 1\;. \tag{135}$$

Combining (132) and (135), we obtain that the following relationship holds with probability 1:

$$\tau_j-\tau_i\le V_{\max}(\theta,\boldsymbol{\alpha})-V_{\min}(\theta,\boldsymbol{\alpha})+\left(\frac{1}{\beta_{\min}}+\frac{1}{\beta_{\max}}\right)h+1\;.$$

Combining (131) and (136), we obtain that with probability 1

$$\begin{aligned}T\le 1&+\frac{W_{\boldsymbol{\beta}}(\theta,\boldsymbol{\alpha})}{K+1}\\&+\frac{K}{K+1}\left(V_{\max}(\theta,\boldsymbol{\alpha})-V_{\min}(\theta,\boldsymbol{\alpha})+\left(\frac{1}{\beta_{\min}}+\frac{1}{\beta_{\max}}\right)h+1\right)\;.\end{aligned} \tag{136}$$

Finally, to ensure that $T-\tau_i-1>t(h)$ for every parameter, we follow the same steps as (112), (113) and (116) in order to obtain constants $L(h)$ and $M_i(h)$ for every $i\in[K]$, such that $\beta\in(0,L(h))$ and $\beta_i\in(0,M_i(h))$ ensure the required conditions. The corresponding choices of $L(h)$ and $M_i(h)$ are given by:

$$\begin{aligned}L(h) \;&\triangleq\; \frac{1}{t(h)}\left(\sum_{i\in[K]}\mathscr{J}_i(\theta)\right)^{-1},\\ \text{and}\quad M_i(h) \;&\triangleq\; \inf_{\theta\in\Theta}\frac{1}{t(h)}\left(\mathscr{J}_i(\alpha_i\mid\theta)\right)^{-1}.\end{aligned} \tag{137}$$

Appendix G
Proof of Theorem 7

From Lemma 4, there exists $T_\epsilon<+\infty$ such that for all $t\ge T_\epsilon$,

$$\begin{aligned}t\Bigg\{\frac{\mathsf{C}(\nu_t^{\mathsf{MMSE}}\mid\mathcal{F}_t)}{\beta}+\sum_{i\in[K]}\frac{\mathsf{D}(\zeta_{i,t}^{\mathsf{MMSE}}\mid\mathcal{F}_t)}{\beta_i}\Bigg\}\\ \in[W_{\boldsymbol{\beta}}(\theta,\boldsymbol{\alpha})-\epsilon,W_{\boldsymbol{\beta}}(\theta,\boldsymbol{\alpha})+\epsilon]\;.\end{aligned} \tag{138}$$

Let $a\in[K]$ be the first experiment for which the estimation guarantee is achieved, and let $(T-\tau_a)$ denote the time instant at which this guarantee is achieved. Then, using (138), along with the stopping criterion in (52) we have

$$\begin{aligned}T-\tau_a-1 \;&=\; (T-\tau_a-1)\mathbb{1}_{\{T-\tau_a-1\le T_\epsilon\}}\\&\quad+(T-\tau_a-1)\mathbb{1}_{\{T-\tau_a-1>T_\epsilon\}} &(139)\\&\overset{(138)}{\le}\; T_\epsilon+\frac{W_{\boldsymbol{\beta}}(\theta,\boldsymbol{\alpha})+\epsilon}{K+1}+1\;. &(140)\end{aligned}$$

Hence,

$$T\;\le\;T_\epsilon+\frac{W_{\boldsymbol{\beta}}(\theta,\boldsymbol{\alpha})+\epsilon}{K+1}+\tau_a+2\;. \tag{141}$$

Furthermore, using (85), there exists $T_0(\epsilon)<+\infty$ such that for all $t\ge T_0(\epsilon)$

$$t\mathsf{C}(\nu_t^{\mathsf{MMSE}}\mid\mathcal{F}_t)\in[\beta V(\theta)-\epsilon,\beta V(\theta)+\epsilon]\;, \tag{142}$$

where $V(\theta)$ is defined in (54). Similarly, using (86) there exists $T_i(\epsilon)<+\infty$ for $i\in[K]$, such that for all $t\ge T_i(\epsilon)$

$$t\mathsf{D}(\zeta_{i,t}^{\mathsf{MMSE}}\mid\mathcal{F}_t)\in[\beta_i V_i(\alpha_i)-\epsilon,\beta_i V_i(\alpha_i)+\epsilon]\;, \tag{143}$$

where $V_i(\alpha_i)$ is defined in (54). Let $T_\epsilon^*\triangleq\max\{T_0(\epsilon),\cdots,T_K(\epsilon)\}$. Furthermore, let $b\in[K]$ be the experiment selected at the stopping time. Thus, $\tau_b=0$. Subsequently, for any $\epsilon>0$, (143) yields

$$\begin{aligned}T\;&\le\;T^*(\epsilon)+\frac{V_b(\alpha_b)+\epsilon}{\beta_b}+2 &(144)\\&\le\;T^*(\epsilon)+V_{\max}(\theta,\boldsymbol{\alpha})+\frac{\epsilon}{\beta_b}+2\;, &(145)\end{aligned}$$

where (144) is a result of the fact that at stopping, $\mathsf{D}(\zeta_{i,t}^{\mathsf{MMSE}}\mid\nu_t^{\mathsf{MMSE}},\mathcal{F}_t)\le\beta_i$, which implies that $\mathsf{D}(\zeta_{i,t}^{\mathsf{MMSE}}\mid\mathcal{F}_t)\le\beta_i$, which is obtained by taking an expectation over $\mathsf{D}(\zeta_{i,t}^{\mathsf{MMSE}}\mid\nu_t^{\mathsf{MMSE}},\mathcal{F}_t)$ with respect to the measure $\mathbb{E}_t^i$. Thus, noting that $\tau_a\le T$, we have

$$\tau_a\;\le\;T_\epsilon^*+V_{\max}(\theta,\boldsymbol{\alpha})+\frac{\epsilon}{\beta_b}+2\;. \tag{146}$$

Finally, combining (141) and (146), we obtain:

$$T\;\le\;T_\epsilon+T_\epsilon^*+\frac{W_{\boldsymbol{\beta}}(\theta,\boldsymbol{\alpha})}{K+1}+V_{\max}(\theta,\boldsymbol{\alpha})+\frac{\epsilon}{\beta_b}+4\;. \tag{147}$$

Note that $\sup_{Y_t,\psi^t}T_\epsilon<+\infty$, and $\sup_{Y_t,\psi^t}T_\epsilon^*<+\infty$ due to Lemma 4. The proof is completed by taking the expectation on both sides of (147), dividing by $W_{\boldsymbol{\beta}}(\theta,\boldsymbol{\alpha})$, and taking the limit.

Appendix H
Proof of Theorem 4

By contradiction, we prove that a purely Chernoff-based sampling strategy is sub-optimal. We begin by assuming that the Chernoff-based sampling strategy described in section V-B1 combined with an almost surely finite stopping time τ^{c} is asymptotically optimal. This implies that every constraint given in the problem (16) is satisfied at stopping. First, let us define the following quantities that are instrumental for our argument:

$$a^*\;\triangleq\;\underset{i\in[K]}{\arg\max}\;\Big\{\mathscr{J}_i(\theta)+\mathscr{J}_i(\alpha_i)\Big\}\;, \tag{148}$$

$$\mathcal{U}\;\triangleq\;\big\{\{\theta,\alpha_i\}:\theta\in\Theta,\alpha_i\in\mathcal{A}_i\;\forall i\in[K]\big\}\;, \tag{149}$$

$$\begin{aligned}\mathcal{U}_i\;\triangleq\;\Bigg\{&\theta\in\Theta,\alpha_i\in\mathcal{A}_i\;:\;\mathscr{J}_i(\theta)+\mathscr{J}_i(\alpha_i)\\&>\;\mathscr{J}_j(\theta)+\mathscr{J}_j(\alpha_i)\;\forall\,j\in[K]\setminus i\Bigg\}\;,\end{aligned} \tag{150}$$

where a^* represents the most informative experiment that maximizes the overall FI, $\mathcal{U}$ is a set containing all possible pairs of the parameters $\{\theta,\alpha_i\}$ for every $i\in[K]$, and $\mathcal{U}_i$ represents a subset of $\mathcal{U}$ that maximizes the FI computed for the pair $\{\theta,\alpha_i\}$ for the experiment $i\in[K]$. It can be readily verified that $(\theta,\alpha_{a^*})\in\mathcal{U}_{a^*}$. Now, by the strong consistency of the ML estimate [53], we have,

$$\nu_t^{\mathsf{ML}}\to\theta,\qquad\zeta_{a^*,t}^{\mathsf{ML}}\to\alpha_{a^*}\quad\text{a.s.} \tag{151}$$

This implies that there exists a finite $m(\epsilon_1,\epsilon_2)\in\mathbb{N}$ such that, almost surely, $|\nu_t^{\mathsf{ML}}-\theta|\;<\;\epsilon_1$, and $|\zeta_{a^*,t}^{\mathsf{ML}}-\alpha_{a^*}|\;<\;\epsilon_2$ for every $t\;\ge$

$m(\epsilon_1, \epsilon_2)$. We also note that the ground truth (θ, α_{a^*}) is an interior point of the set $\mathcal{U}_{a^*}$, since $\mathcal{U}_i$ is an open set by definition. Thus, if the FI measures $\mathscr{I}(\theta)$ and $\mathscr{J}(\alpha_i)$ for every $i \in [K]$ are continuous functions of the respective parameters θ and α_i, there exists an ϵ-ball around the point (θ, α_{a^*}), $\epsilon > 0$, such that any set of parameters $\{\theta', \alpha'_{a^*}\}$ within the ϵ-ball also maximizes the FI. Selecting $m(\epsilon_1, \epsilon_2)$ to be sufficiently large, such that $\epsilon < \min(\epsilon_1, \epsilon_2)$, we have that $(\nu_t^{\mathsf{ML}}, \zeta_{a^*,t}^{\mathsf{ML}}) \in \mathcal{U}_{a^*}$ for every $t \geq m(\epsilon_1, \epsilon_2)$. Next, we will show that for a range of the confidence interval β, the optimal stopping rule uses at least $m(\epsilon_1, \epsilon_2)$ samples with a high probability. Combining the fact that $\bar{p}_t \to p^*$, where $\bar{p}_t$ represents the probability mass function over the experiments $i \in [K]$ due to the Chernoff-based sampling strategy, and p^* denotes the distribution which selects the experiment a^* with probability 1, and the Bernstein-Von Mises Theorem ([54]), we obtain

$$t \cdot \mathsf{C}(\nu_t^{\mathsf{MMSE}} \mid \mathcal{F}_t^{\mathsf{c}}) \xrightarrow{t \to \infty} 1/\mathscr{J}_{a^*}(\theta) \,, \tag{152}$$

where we have defined $\mathcal{F}_t^{\mathsf{c}} \triangleq \{Y^t, \psi_{\mathsf{c}}^t\}$ and $\psi_{\mathsf{c}}^t \triangleq \{\psi^{\mathsf{c}}(1), \cdots, \psi^{\mathsf{c}}(t)\}$. Let τ^{c} denote the optimal stopping rule that minimizes (16) under ψ_{c}^t. Using (152), there exists $\tau_\epsilon < +\infty$, such that for all $t \geq \tau_\epsilon$, we have

$$1 \;=\; \mathbb{P}\Big\{ t \cdot \mathsf{C}(\nu_t^{\mathsf{MMSE}} \mid \mathcal{F}_t^{\mathsf{c}}) \;\geq\; 1/\mathscr{J}_{a^*}(\theta) - \epsilon \Big\} \,. \tag{153}$$

Furthermore, choosing $\tau^* = \max\{\tau_\epsilon, m(\epsilon_1, \epsilon_2)\}$, we have that for all $t \geq \tau^*$,

$$\begin{aligned} 1 \;=\; & \mathbb{P}\Big\{ \tau^{\mathsf{c}} \cdot \mathsf{C}(\nu_{\tau^{\mathsf{c}}}^{\mathsf{MMSE}} \mid \mathcal{F}_{\tau^{\mathsf{c}}}^{\mathsf{c}}) \;\geq\; 1/\mathscr{J}_{a^*}(\theta) - \epsilon, \tau^{\mathsf{c}} \;\leq\; \tau^* \Big\} \\ & + \; \mathbb{P}\Big\{ \tau^{\mathsf{c}} \cdot \mathsf{C}(\nu_{\tau^{\mathsf{c}}}^{\mathsf{MMSE}} \mid \mathcal{F}_{\tau^{\mathsf{c}}}^{\mathsf{c}}) \;\geq\; 1/\mathscr{J}_{a^*}(\theta) - \epsilon, \tau^{\mathsf{c}} \;>\; \tau^* \Big\} \\ \leq \; & \mathbb{P}\Big\{ \tau^{\mathsf{c}} \beta \;\geq\; 1/\mathscr{J}_{a^*}(\theta) - \epsilon, \tau^{\mathsf{c}} \;\leq\; \tau^* \Big\} \\ & + \; \mathbb{P}\Big\{ \tau^{\mathsf{c}} \cdot \mathsf{C}(\nu_{\tau^{\mathsf{c}}}^{\mathsf{MMSE}} \mid \mathcal{F}_{\tau^{\mathsf{c}}}^{\mathsf{c}}) \;\geq\; 1/\mathscr{J}_{a^*}(\theta) - \epsilon, \tau^{\mathsf{c}} \;>\; \tau^* \Big\} \,, \end{aligned} \tag{154}$$

where (154) holds since at the stopping time, $\mathsf{C}(\nu_{\tau^{\mathsf{c}}}^{\mathsf{MMSE}} \mid \mathcal{F}_{\tau^{\mathsf{c}}}^{\mathsf{c}}) \leq \beta$. Selecting $\beta < \frac{1}{\tau^*}\Big(\frac{1}{\mathscr{I}(\theta)} - \epsilon\Big)$, we observe that $\mathbb{P}\Big\{\tau^{\mathsf{c}}\beta \;\geq\; 1/\mathscr{J}_{a^*}(\theta) - \epsilon, \tau^{\mathsf{c}} \;\leq\; \tau^*\Big\} = 0$. Thus, we have

$$\mathbb{P}\Big\{ \tau^{\mathsf{c}} \;>\; \tau^* \Big\} \;=\; 1 \,. \tag{155}$$

Finally, note that for all $t > \tau^*$ and for each experiment $i \in [K] \setminus a^*$, $\mathsf{D}(\zeta_{i,\tau^{\mathsf{c}}}^{\mathsf{MMSE}} | \mathcal{F}_{\tau^{\mathsf{c}}}^{\mathsf{c}},\; \nu_{\tau^{\mathsf{c}}}^{\mathsf{MMSE}}) = \mathsf{D}(\zeta_{i,\tau^*}^{\mathsf{MMSE}} | \mathcal{F}_{\tau^*}^{\mathsf{c}},\; \nu_{\tau^*}^{\mathsf{MMSE}})$, since for all $t \geq \tau^*$, the Chernoff-based control action selects experiment a^*. Let $\gamma \triangleq \min_{i \in [K] \setminus a^*} \mathsf{D}(\zeta_{i,\tau^{\mathsf{c}}}^{\mathsf{MMSE}} | \mathcal{F}_{\tau^{\mathsf{c}}}^{\mathsf{c}}, \nu_{\tau^{\mathsf{c}}}^{\mathsf{MMSE}})$, $\gamma > 0$. Thus, if we choose $\beta_i < \gamma$ for any $i \in [K] \setminus a^*$, the sampling strategy forces the estimation cost $\mathsf{C}(\zeta_{i,\tau^{\mathsf{c}}}^{\mathsf{MMSE}} | \mathcal{F}_{\tau^{\mathsf{c}}}^{\mathsf{c}},\; \nu_{\tau^{\mathsf{c}}}^{\mathsf{MMSE}}) > \beta_i$, hence, failing to satisfy the constraints in 16, which is a contradiction to our initial assumption.

REFERENCES

[1] H. Chernoff, "Sequential design of experiments," *Annals of Mathematical Statistics*, vol. 30, no. 3, pp. 755–770, September 1959.

[2] A. Tajer, J. Heydari, and H. V. Poor, "Active sampling for the quickest detection of markov networks," *IEEE Transactions on Information Theory*, vol. 68, no. 4, pp. 2479–2508, 2022.

[3] C. Dimitrakakis and C. Savu-Krohn, "Cost-minimising strategies for data labelling: Optimal stopping and active learning," in *Proc. International Conference on Foundations of Information and Knowledge Systems*, Pisa, Italy, February 2008.

[4] S. Tong and D. Koller, "Active learning for parameter estimation in Bayesian networks," in *Proc. International Conference on Neural Information Processing Systems*, Denver, CO, November 2000.

[5] V. Anantharam, P. Varaiya, and J. Walrand, "Asymptotically efficient allocation rules for the multiarmed bandit problem with multiple plays—Part I: I.I.D. rewards," *IEEE Transactions on Automatic Control*, vol. 32, pp. 968 – 976, December 1987.

[6] D. Bouneffouf, I. Rish, and C. Aggarwal, "Survey on applications of multi-armed and contextual bandits," in *Proc. IEEE Congress on Evolutionary Computation*, Glasgow, UK, July 2020.

[7] H.-H. Chang and Z. Ying, "A global information approach to computerized adaptive testing," *Applied Psychological Measurement*, vol. 20, no. 3, pp. 213–229, September 1996.

[8] A. G. Tartakovsky, B. L. Rozovskii, R. B. Blazek, and H. Kim, "A novel approach to detection of intrusions in computer networks via adaptive sequential and batch-sequential change-point detection methods," *IEEE Transactions on Signal Processing*, vol. 54, no. 9, pp. 3372–3382, September 2006.

[9] M. Marcus and P. Swerling, "Sequential detection in radar with multiple resolution elements," *IRE Transactions on Information Theory*, vol. 8, no. 3, pp. 237–245, April 1962.

[10] S. Bessler, "Theory and applications of the sequential design of experiments, k-actions and infinitely many experiments: Part I. Theory," Ph.D. dissertation, Department of Statistics, Stanford University, March 1960.

[11] A. E. Albert, "The sequential design of experiments for infinitely many states of nature," *The Annals of Mathematical Statistics*, vol. 32, no. 3, pp. 774–799, September 1961.

[12] G. E. P. Box and W. J. Hill, "Discrimination among mechanistic models," *Technometrics*, vol. 9, no. 1, pp. 57–71, February 1967.

[13] W. J. Blot and D. A. Meeter, "Sequential experimental design procedures," *Journal of the American Statistical Association*, vol. 68, no. 343, pp. 586–593, Sep. 1973.

[14] S. Nitinawarat, G. K. Atia, and V. V. Veeravalli, "Controlled sensing for multihypothesis testing," *IEEE Transactions on Automatic Control*, vol. 58, no. 10, pp. 2451–2464, May 2013.

[15] S. Nitinawarat and V. V. Veeravalli, "Controlled sensing for sequential multihypothesis testing with controlled Markovian observations and non-uniform control cost," *Sequential Analysis*, vol. 34, no. 1, pp. 1–24, August 2015.

[16] K. Cohen and Q. Zhao, "Active hypothesis testing for anomaly detection," *IEEE Transactions on Information Theory*, vol. 61, no. 3, pp. 1432–1450, March 2015.

[17] N. K. Vaidhiyan and R. Sundaresan, "Active search with a cost for switching actions," in *Proc. Information Theory and Applications Workshop*, San Diego, CA, February 2015.

[18] M. Naghshvar and T. Javidi, "Active sequential hypothesis testing," *Annals of Statistics*, vol. 41, no. 6, pp. 2703–2738, December 2013.

[19] J. Wang, "Optimal sequential multi-class diagnosis," *Operations Research*, November 2020.

[20] K. Cohen, Q. Zhao, and A. Swami, "Optimal index policies for anomaly localization in resource-constrained cyber systems," *IEEE Transactions on Signal Processing*, vol. 62, no. 16, pp. 4224–4236, August 2014.

[21] K. Cohen and Q. Zhao, "Asymptotically optimal anomaly detection via sequential testing," *IEEE Transactions on Signal Processing*, vol. 63, no. 11, pp. 2929–2941, June 2015.

[22] J. Heydari and A. Tajer, "Quickest search and learning over correlated sequences: Theory and application," *IEEE Transactions on Signal Processing*, vol. 67, no. 3, pp. 638–651, February 2019.

[23] J. Heydari, A. Tajer, and H. V. Poor, "Quickest Linear Search over Correlated Sequences," *IEEE Transactions on Information Theory*, vol. 62, no. 10, pp. 5786–5808, October 2016.

[24] P. J. Bickel and J. A. Yahav, "Asymptotically pointwise optimal procedures in sequential analysis," in *Proc. Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Statistics*, Berkeley, CA, 1967, pp. 401–413.

[25] W. J. Blot and D. A. Meeter, "Sequential experimental design procedures," *Journal of the American Statistical Association*, vol. 68, no. 343, pp. 586–593, April 1973.

[26] K. Cohen and Q. Zhao, "Active hypothesis testing for anomaly detection," *IEEE Transactions on Information Theory*, vol. 61, no. 3, pp. 1432–1450, March 2015.

[27] D. Zois, M. Levorato, and U. Mitra, "Active classification for pomdps: A Kalman-like state estimator," *IEEE Transactions on Signal Processing*, vol. 62, no. 23, pp. 6209–6224, December 2014.

[28] J. G. Ligo, G. K. Atia, and V. V. Veeravalli, "A controlled sensing approach to graph classification," *IEEE Transactions on Signal Processing*, vol. 62, no. 24, pp. 6468–6480, December 2014.

[29] A. Rangi, M. Franceschetti, and S. Marano, "Distributed Chernoff test: Optimal decision systems over networks," *IEEE Transactions on Information Theory*, vol. 67, no. 4, pp. 2399–2425, April 2021.

[30] G. Schwarz, "Asymptotic shapes of bayes sequential testing regions," *Annals of Mathematical Statistics*, vol. 33, no. 1, pp. 224–236, March 1962.

[31] J. Kiefer and J. Sacks, "Asymptotically optimum sequential inference and design," *Annals of Mathematical Statistics*, vol. 34, no. 3, pp. 705–750, September 1963.

[32] S. P. Lalley and G. Lorden, "A control problem arising in the sequential design of experiments," *Annals of Probability*, vol. 14, no. 1, pp. 136–172, January 1986.

[33] T. L. Lai, "Sequential analysis: Some classical problems and new challenges," *Statistica Sinica*, vol. 11, no. 2, pp. 303 – 351, 2001.

[34] B. K. Ghosh, "Sequential analysis: Tests and confidence intervals (david siegmund)," *SIAM Review*, vol. 29, no. 2, pp. 315–318, 1987.

[35] M. Ghosh, N. Mukhopadhyay, and P. K. Sen, *Sequential Estimation*. John Wiley & Sons, New York, NY, 2011.

[36] M. Ghosh and N. Mukhopadhyay, "On two fundamental problems of sequential estimation," *Sankhyā: The Indian Journal of Statistics, Series B*, vol. 38, no. 3, pp. 203–218, 1976.

[37] F. J. Anscombe, "Large-sample theory of sequential estimation," *Mathematical Proceedings of the Cambridge Philosophical Society*, vol. 48, no. 4, pp. 600–607, 1952.

[38] V. J. Yohai, "Asymptotically optimal Bayes sequential design of experiments for estimation," *Annals of Statistics*, vol. 1, no. 5, pp. 822–837, September 1973.

[39] C.-D. Fuh and I. Hu, "Asymptotically efficient strategies for a stochastic scheduling problem with order constraints," *The Annals of Statistics*, vol. 28, no. 6, pp. 1670 – 1695, December 2000.

[40] G. Atia and S. Aeron, "Asymptotic optimality results for controlled sequential estimation," in *Proc. Annual Allerton Conference on Communication, Control, and Computing*, Monticello, IL, October 2013.

[41] G. V. Moustakides and T. Yaacoub and Y. Mei, "Sequential estimation based on conditional cost," in *Proc. IEEE International Symposium on Information Theory*, Aachen, Germany, July 2017.

[42] D. O. Segall, "Multidimensional adaptive testing," *Psychometrika*, vol. 61, no. 2, pp. 331–354, 1996.

[43] J. Mulder, V. D. Linden, and J. Wim, "Multidimensional adaptive testing with optimal design criteria for item selection," *Psychometrika*, vol. 74, no. 2, pp. 273–296, 2009.

[44] B. P. Veldkamp, V. D. Linden, and J. Wim, "Multidimensional adaptive testing with constraints on test content," *Psychometrika*, vol. 67, no. 4, pp. 575–588, 2002.

[45] S. W. Choi, M. W. Grady, and B. G. Dodd, "A new stopping rule for computerized adaptive testing," *Educational and Psychological Measurement*, vol. 70, no. 6, pp. 1–17, December 2010.

[46] S. B. Morris, M. Bass, and R. E. Neapolitan, "Stopping rules for computer adaptive testing when item banks have nonuniform information." *International Journal of Testing*, vol. 20, no. 2, pp. 146–168, July 2019.

[47] C. Wang, D. J. Weiss, and Z. Shang, "Variable-length stopping rules for multidimensional computerized adaptive testing," *Psychometrika*, vol. 84, no. 3, pp. 749–771, 2019.

[48] C. Wang, H.-H. Chang, and K. A. Boughton, "Deriving stopping rules for multidimensional computerized adaptive testing," *Applied Psychological Measurement*, vol. 37, no. 2, pp. 99–122, 2013.

[49] S. Nitinawarat, G. K. Atia, and V. V. Veeravalli, "Controlled sensing for multihypothesis testing," *IEEE Transactions on Automatic Control*, vol. 58, no. 10, pp. 2451–2464, October 2013.

[50] C. Qin, D. Klabjan, and D. Russo, "Improving the expected improvement algorithm," in *Proc. Advances in Neural Information Processing Systems*, Long Beach, CA, December 2017.

[51] A. Garivier and E. Kaufmann, "Optimal best arm identification with fixed confidence," in *Proc. Conference on Learning Theory*, New York, NY, June 2016.

[52] X. Shang, R. de Heide, P. Menard, E. Kaufmann, and M. Valko, "Fixed-confidence guarantees for Bayesian best-arm identification," in *Proc. International Conference on Artificial Intelligence and Statistics*, Sicily, Italy, August 2020, pp. 1823–1832.

[53] W. Newey and D. McFadden, "Large sample estimation and hypothesis testing," *Handbook of Econometrics*, vol. 4, pp. 2111–2245, 1986.

[54] A. DasGupta, *Asymptotic Theory of Statistics and Probability*, ser. Springer Texts in Statistics. Springer, New York, NY, 2008.

[55] S. M. Kay, *Fundamentals of Statistical Signal Processing: Estimation Theory*. Prentice Hall, NJ, 1997.