

Authors are encouraged to submit new papers to INFORMS journals by means of a style file template, which includes the journal title. However, use of a template does not certify that the paper has been accepted for publication in the named journal. INFORMS journal templates are for the exclusive purpose of submitting to an INFORMS journal and should not be used to distribute the papers in print or online or to submit the papers to another publication.

Practical Nonparametric Sampling Strategies for Quantile-based Ordinal Optimization

Dongwook Shin

School of Business and Management, Hong Kong University of Science and Technology,
Clear Water Bay, Kowloon, Hong Kong
{dwshin@ust.hk}

Mark Broadie, Assaf Zeevi

Graduate School of Business, Columbia University, New York, NY 10025
{mnb2@columbia.edu, assaf@gsb.columbia.edu}

Given a finite number of stochastic systems, the goal of our problem is to dynamically allocate a finite sampling budget to maximize the probability of selecting the “best” system. Systems are encoded with the probability distributions that govern sample observations, which are unknown and only assumed to belong to a broad family of distributions that need not admit any parametric representation. The “best” system is defined as the one with the highest quantile value. The objective of maximizing the probability of selecting this “best” system is not analytically tractable. In lieu of that we use the rate function for the probability of error relying on large deviations theory. Our point of departure is an algorithm that naively combines sequential estimation and myopic optimization. This algorithm is shown to be asymptotically optimal, however, it exhibits poor finite-time performance and does not lend itself to implementation in settings with a large number of systems. To address this we propose practically implementable variants that retain the asymptotic performance of the former, while dramatically improving its finite-time performance.

Key words: Quantile, ordinal optimization, tractable procedures, large deviations theory

History:

1. Introduction

Given a finite number of alternative distributions, henceforth referred to as *systems*, we are concerned with the problem of selecting the “best” one, where performance of the systems is initially unknown but it is possible to sequentially sample from each of them. A critical assumption made in most academic studies is that a decision maker is primarily interested in the mean performances of the systems, i.e., the means of underlying distribu-

tions. However, such mean-based approaches do not accommodate various risk preferences of the decision maker. In this paper, we explore cases where downside or upside risk is more crucial than the average performance. In particular, the systems are compared based on p th quantiles of their distributions, with the value of p chosen according to the decision maker's risk preference.

In most cases, the underlying probability distributions rarely admit a parametric representation. Therefore, we focus on nonparametric settings where the probability distributions belong to a broad family that cannot be characterized by one or more parameters. While nonparametric approaches have an extensive range of applications, the main difficulty is loss of tractability. Hence, we are motivated to develop nonparametric sampling strategies that are implementable and perform well in practice.

The problem described above arises in a variety of applications. A widely used example is value-at-risk (VaR) that measures the p th quantile of a portfolio's value, with p typically chosen to be 1% or 5%. For large portfolios of complex derivative securities,¹ value-at-risk can be evaluated through Monte Carlo methods, which typically entails substantial computational burden. The quantile-based approach is quite relevant in the case where a portfolio manager needs to identify the portfolio with largest value-at-risk out of many alternatives; in particular, a portfolio manager needs to sequence the simulation trials in an efficient manner given a certain amount of time to increase the likelihood of identifying the portfolio with largest VaR.

Another example is selecting the best design of a telephone call center, where a typical measure of interest is the quantile of waiting time. For instance, if a conservative decision maker uses the quantiles for $p = 0.95$ to compare different designs, then the best design will have the smallest value of the 95% quantile of waiting times. The waiting time is a complex function of design parameters, such as staffing and routing policies, and is often evaluated through simulation. The decision maker's job is to configure a finite number of designs and to spend a certain amount of time testing different designs before deciding which is best.

The main objective of this paper is to design a dynamic sampling algorithm that minimizes the probability of selecting suboptimal systems subject to a given sampling budget.

¹ In practice, there can be some correlation in performance between portfolios if they consist of similar assets. However, following most of papers in the area of ordinal optimization and ranking and selection, our model assumes independent systems. (Our model serves as a good approximation when the degree of correlation is not significant.)

Since the aforementioned objective is not analytically tractable, our departure point will be an asymptotic benchmark characterized by a large sampling budget. In this regime, the rate function (Dembo and Zeitouni 2009, p. 4) of the probability of false selection, hereafter simply referred to as the rate function, can be written in tractable form.

Exact evaluation of the rate function, however, is possible only when the decision maker knows the underlying probability distributions in advance, and the absence of perfect prior information introduces an important new component into the optimization problem. Namely, one needs to allocate samples to maximize the rate function (exploitation) but at the same time take sufficient samples from each system to learn the underlying distributions (exploration).

As a benchmark for further analysis, we propose a naive algorithm that estimates unknown values of the rate function from the history of sample observations, and then allocates a sample in each subsequent stage as if the unknown values are equal to the sample estimates; this is often referred to as a certainty-equivalence approach. As will be explained later, this algorithm maximizes the objective asymptotically, however, its finite-time performance is poor since it spends too much of the sampling budget on exploration, which leaves less budget to exploit that knowledge and optimize the objective. Another drawback of this approach is the heavy computational burden since it requires solving a difficult nested optimization problem repeatedly, which makes the algorithm not practically implementable in many applications with a large number of systems.

The main contribution of this paper is to introduce sampling strategies that are practical to implement in applications with large number of systems, and improve their finite-time performance relative to the aforementioned class of algorithms. In more detail, our contributions are summarized as follows.

- (i) We introduce an alternative performance metric and show that it is closely aligned with the rate function when the gap between the quantiles of the best and second-best systems is sufficiently small. The alternative performance metric is structured around the quantiles and local density behavior, rather than global distribution functions, which lends itself to some key structural insights about near-optimal allocations.
- (ii) Building on the structural properties of the alternative performance metric, we propose a family of nonparametric dynamic sampling algorithms, hereafter referred to as Quantile Divergence (QD), with the aim of improving the finite-time performance of

the aforementioned naive algorithm. We rigorously show that this class of algorithms is near-optimal in a precise mathematical sense.

- (iii) We also propose variants of QD, hereafter referred to as Adaptive QD (AQD), that are computationally efficient when the number of systems is large. We show via numerical testing that these variants preserve the asymptotic performance of the QD algorithms and improve the computation time dramatically.

Our proposed algorithms are *nonparametric* in the sense that they do not postulate any parametric structure on underlying probability distributions. Despite having almost no prior information, these algorithms allocate samples in a way that is close to the best ex ante allocation under full information. When the underlying distributions are known to belong to a parametric family, more efficient algorithms can be designed, albeit at the cost of misspeciation; in particular, if the assumed parametric structure is not “close enough” to that of the true underlying family of distributions, such parametric algorithms may fail to identify the best system.

The remainder of this paper is organized as follows: In §2 we survey related literature on ordinal optimization, and ranking and selection problems, and relevant studies from the multi-armed bandit literature. In §3 we derive a tractable objective function using large deviations theory and formulate a dynamic optimization problem. In §4 we suggest dynamic algorithms and provide theoretical analyses on their performances. In §5 we propose adaptive variants of the algorithms that are practically implementable in large problem instances. In §6 we test the suggested policies numerically and compare with several benchmark policies. This paper has an online supplement with three parts. Appendix A states additional theoretical results; Appendix B contains the proofs of main theoretical results and auxiliary lemmas; and Appendix C contains additional numerical results to provide practical guideline on the implementation of proposed algorithms.²

2. Literature Review

Mean-based R&S Procedures. An area closely related to ordinal optimization is that of ranking and selection (R&S). While the goal of this paper is to minimize the probability of

² A very preliminary version of this paper appeared as Shin et al. (2016). In particular, §4 contains much stronger theoretical results in §3 of the prior paper, and Appendix A provides further theoretical results. New algorithms to improve those in the prior paper are given in §5. Lastly, §6 exhibits numerical experiments in more extensive settings than those in §4 of the previous paper.

false selection given a *fixed* sampling budget, the goal of the R&S problem is to take as few samples as possible to satisfy a desired guarantee on the probability of correctly selecting the best system; most of this literature considers “best” to be the largest mean. The traditional R&S approach traces back to the work of Bechhofer (1954) who established the Indifference Zone (IZ) formulation. (See the survey paper by Kim and Nelson (2006).) Most IZ procedures rely on the assumption that sample observations are normally distributed, and hence raise the risk of misspecification alluded to earlier.

Quantile-based R&S Procedures. Despite the wide usage of quantiles as a performance metric, the topic of quantile-based R&S procedure has not received much attention in the literature. Bekki et al. (2007) modify the traditional two-stage IZ procedure by Rinott (1978) to suggest a heuristic technique that addresses the issue of non-normality of quantile estimates, by averaging them over batches. (Thus the batch behavior would be nearly normal by the central limit theorem.) Batur and Choobineh (2010) suggest a two-stage procedure based on Rinott (1978), where a set of quantile values is compared between two systems. Lastly, Lee and Nelson (2014) suggest an R&S procedure based on bootstrapping, which can be applied to general performance measures including quantile, albeit with a heavy computational load. In contrast to these papers, we show that the (asymptotically) optimal allocation is unaffected by the use of batches, so that the allocation based on the normal approximation may give substantially suboptimal allocations.

Mean-based Ordinal Optimization Procedures. Several studies consider the probability of selecting the best system when a finite sampling budget is given. Notable recent examples include the Optimal Computing Budget Allocation (OCBA) rule by Chen et al. (2000), the Knowledge Gradient (KG) Bayesian framework of Frazier et al. (2008). These formulations provide an attractive, stylized analysis but they rely heavily on normality assumptions. In contrast, the “frequentist” approach of Glynn and Juneja (2004) based on large deviations theory allows for much broader scope but is difficult to implement due to (i) the challenges in estimating the moment generating function, as noted by Glynn and Juneja (2015), and (ii) the computational burden in optimizing the large deviations rate function, as noted by Pasupathy et al. (2015). Regarding (i), Glynn and Juneja show that this issue can be mitigated in restricted settings where upper bounds on suitable moments of underlying distributions are known (such information is rarely available in most practical applications). Regarding (ii), Pasupathy et al. show that the optimal allocation for the rate

function becomes tractable in a certain asymptotic limit characterized by the large number of systems. In the context of our work, the large deviations rate function also entails issues (i) and (ii); see §4.1. However, our work is significantly different from previous papers in that we circumvent these issues by introducing an approximation to the rate function in a certain asymptotic regime characterized by the quantiles of the systems.

The recent work of Shin et al. (2018) also addresses the issues with the large deviations rate function approach of Glynn and Juneja (2004) by proposing a two moment approximation and leveraging it to design a well-performing dynamic procedure called Welch Divergence (WD). This paper shares an important common theme with Shin et al. (2018) in that we too use the large deviations rate function to construct competitive sampling procedures. On the other hand, the key conclusions from the two papers are fundamentally different: the two-moment approximation of Shin et al. (2018) suggests that when systems are compared based on means, the probability of false selection can be approximated by the rate function corresponding to Gaussian distribution if the probability distributions of the best and second-best systems are “close” to each other. In contrast, when systems are compared based on quantiles, we show that the allocation based on the two-moment approximation can be significantly sub-optimal, as the (asymptotically) near-optimal allocation rule requires density information rather than just two moments. This calls into further questions the Gaussian assumptions that are prevalent in the literature.

Quantile-based Ordinal Optimization Procedures. Pasupathy et al. (2010) characterize the rate function associated with the probability of false selection using large deviations theory. However to build on this one requires knowledge of the underlying distribution functions. This paper extends their work along three dimensions: we strengthen their theoretical results on the rate function characterization by relaxing conditions on the underlying distributions; we derive an approximation to the rate function using only quantiles and local density estimates at particular points, which is simpler than having to estimate an entire distribution function; and we propose sampling algorithms that judiciously manage the issue of the lack of prior information on the underlying distributions.

Peng et al. (2019) propose Bayesian sampling algorithms for selecting the optimal quantile based on the premise that underlying probability distributions belong to a known parametric family, including normal, exponential, and Pareto. The efficacy of their algorithms is tested in a variety of (non-)normal cases. Our focus is to design and analyze

algorithms that do not rely on any parametric distributional assumption, and provide rigorous justification for their efficacy.

Quantile-based Best-arm Identification. Our research is closely related to that of pure exploration in the multi-armed bandit (MAB) problem, often referred to as best-arm identification; see Bubeck and Cesa-Bianchi (2012) for a comprehensive overview. The best-arm identification procedures seek the same goal of selecting the best arm (i.e., system, in the language of our paper). In the context of quantile-based pure exploration, Yu and Nikolova (2013) consider the problem where arms are compared based on value-at-risk. Szorenyi et al. (2015) proposes quantile-based online learning algorithms when rewards from the arms are not necessarily continuous real-valued. Our work differs significantly from these papers in that we deal with the setting with a fixed sampling budget. Further, from an analytical perspective, the two antecedent papers primarily use the concentration property of quantile estimators to determine the number of samples taken to satisfy a desired guarantee. In contrast, our main concern is to minimize the probability of false selection given a (sufficiently large) number of samples, which gives rise to the question of what is the behavior of the probability distributions of quantile estimators in the extreme tails. This question is addressed in §3.2.

3. Formulation

3.1. Model Primitives

Consider k stochastic systems, each of which is characterized by a distribution function $F_j(\cdot)$, $j = 1, \dots, k$, with its support denoted as \mathcal{H}_j . We define the system configuration as $\mathbf{F} = \{F_1, \dots, F_k\}$; we denote $\mathbf{F} \in \mathcal{C}$ if each distribution is continuous on its domain; and $\mathbf{F} \in \mathcal{D}$ if each distribution is discrete and, without loss of generality, has a support in the set of nonnegative integers. Fix $p \in (0, 1)$ that represents the quantile of interest and define the p th quantile of $F_j(\cdot)$ as

$$\xi_j = \inf\{x : F_j(x) \geq p\}. \quad (1)$$

Denote $\boldsymbol{\xi} = (\xi_1, \dots, \xi_k)$ as the k -dimensional vector of the p th quantiles. Without loss of generality, we take $\xi_1 > \xi_2 \geq \dots \geq \xi_k$. We make the following assumptions throughout the paper.

(F1) ξ_j is the unique solution x of $F_j(x-) \leq p \leq F_j(x)$ for each system j

(F2) $[\xi_k, \xi_1] \subset \cap_{j=1}^k \mathcal{H}_j$

Note that (F1) is a mild assumption that ensures $F_j(\cdot)$ is not flat around the p th quantile. Assumption (F2) ensures that each quantile estimated from sample observations can take any value in the interval $[\xi_k, \xi_1]$, in order to avoid trivial cases where the probability of false selection is zero. In our model the decision maker does not know the system configuration \mathbf{F} , but is able to sequentially take T independent samples from the systems to infer quantiles of the distributions, where T is exogenously given. The goal is to correctly identify the system with the largest p th quantile, i.e., system 1.

Let π denote an algorithm, which is a sequence of random variables, π_1, π_2, \dots , taking values in the set $\{1, \dots, k\}$; the event $\{\pi_t = j\}$ means a sample from system j is taken at stage t . Define X_{jt} as a random sample from system j in stage t such that $\pi_t = j$. The set of non-anticipating policies is denoted as Π , in which the sampling decision in stage t is determined by all the sampling decisions and samples observed in previous stages.

Let N_{jt}^π be the cumulative number of samples up to stage t from system j induced by algorithm π , and define $\alpha_{jt}^\pi := N_{jt}^\pi/t$ as the sampling rate for system j at stage t . We let $\tau_{jn}^\pi = \inf\{t \geq 1 | N_{jt}^\pi = n\}$ denote the n th time that system j is sampled. The sample distribution function for system j is defined as

$$\hat{F}_{jt}^\pi(x) = \frac{1}{N_{jt}^\pi} \sum_{n=1}^{N_{jt}^\pi} \mathbf{I}\{X_{j\tau_{jn}^\pi} \leq x\}, \quad (2)$$

where $\mathbf{I}\{A\}$ is one if A is true and zero otherwise.³ Denote $\hat{\xi}_{jt}^\pi$ as the quantile of the sample distribution function, i.e.,

$$\hat{\xi}_{jt}^\pi = \inf\{x : \hat{F}_{jt}^\pi(x) \geq p\}. \quad (3)$$

It is useful to define a static allocation rule denoted by $\pi(\boldsymbol{\alpha}) \in \Pi$, where $\boldsymbol{\alpha} \in \Delta$ with

$$\Delta = \left\{ (\alpha_1, \dots, \alpha_k) \in \mathcal{R}^k : \sum_{j=1}^k \alpha_j = 1 \text{ and } \alpha_j \geq 0 \text{ for all } j \right\}. \quad (4)$$

Specifically, the allocation rule $\pi(\boldsymbol{\alpha})$ is static in the sense that it allocates $\alpha_j T$ samples to system j independently of sample observations. This rule can be easily implemented in several ways; for example, one may draw one sample from each system and thereafter set $\pi_{t+1}(\boldsymbol{\alpha}) = \arg \max_j \{\alpha_j - \alpha_{jt}^{\pi(\boldsymbol{\alpha})}\}$ for $t \geq k$, with ties broken arbitrarily. We let Δ^0 denote the interior of Δ .

³ In our model, samples from system j are indexed by τ_{jn}^π , $n = 1, \dots, N_{jt}^\pi$, in order to emphasize the fact that the random variable X_{jt} does not need to be defined for times t when system j is not sampled.

DEFINITION 1 (CONSISTENCY). An algorithm $\pi \in \Pi$ is consistent if $N_{jt}^\pi \rightarrow \infty$ in probability for each j as $t \rightarrow \infty$.⁴

In the optimization problem we consider, we further restrict attention to a set of consistent policies, denoted as $\bar{\Pi} \subset \Pi$. Under such policies the sample quantiles are consistent estimators of the population counterparts, as formalized in the following proposition.

PROPOSITION 1 (**Consistency of quantile estimators**). *Under (F1),*

$$\hat{\xi}_{jt}^\pi \rightarrow \xi_j \text{ as } t \rightarrow \infty \quad (5)$$

in probability for any consistent algorithm $\pi \in \bar{\Pi}$.

Note that consistent algorithms ensure that the probability of false selection eventually converges to zero as the sampling budget grows to infinity. We remark that the property of consistency is not straightforward to verify for a dynamic algorithm, although any static allocation rule $\pi(\alpha)$ with $\alpha \in \Delta^0$ is consistent since $N_{jt} = \alpha_j t \rightarrow \infty$ as $t \rightarrow \infty$. The following is a simple example of a dynamic algorithm that is not consistent.⁵

EXAMPLE 1 (A DYNAMIC ALGORITHM THAT IS NOT CONSISTENT). Suppose $\mathbf{F} \in \mathcal{C}$ and each distribution has infinite support. Define a dynamic algorithm π in which n_0 samples are initially taken from each system and $\pi_t = \arg \max_j \{\hat{\xi}_{jt}\}$ for $t \geq kn_0$. Suppose $k = 2$ and $(\xi_1, \xi_2) = (1, 0)$. At stage $2n_0$, it can be easily seen that the event $A = \{\hat{\xi}_{1,2n_0} \in (-\infty, -1), \hat{\xi}_{2,2n_0} \in (1, \infty)\}$ occurs with positive probability. Conditional on this event, system 1 would not be sampled in subsequent stages if $\hat{\xi}_{2t} \geq 0$ for all $t \geq 2n_0$. The latter event occurs with positive probability because

$$\mathbb{P} \left(\sup_{t \geq 2n_0} |\hat{\xi}_{2t} - \xi_2| \leq 1 \right) \geq \frac{2}{1 - \gamma} \gamma^{n_0}, \quad (6)$$

where $\gamma = \exp(-2 \min(F_2(\xi_2 + 1) - p, p - F_2(\xi_2 - 1)))$ (see Section 2.3.2 of Serfling 2009). Combined with the fact that $\mathbb{P}(A) > 0$, system 1 is not sampled infinitely often with positive probability, and hence, the algorithm is not consistent.

Notational conventions. Throughout the paper, we use \hat{x}_t to denote sample estimate of an unknown value x in stage t . For brevity, the superscript π may be dropped when it is clear from the context. We use boldface letters for k -dimensional vectors; e.g., $\alpha := (\alpha_1, \dots, \alpha_k)$ and $\alpha_t := (\alpha_{1t}, \dots, \alpha_{kt})$.

⁴ In this paper, we consider the weak consistency but note that some of the results can be strengthened to strong consistency, albeit with a stronger condition; the analysis of this case is available in Appendix A of the online supplement.

⁵ See also the MAP algorithm in Peng et al. (2019), which is dynamic but not consistent if $p < 1/2$.

3.2. Large Deviations Preliminaries

The probability of false selection, denoted $\mathbf{P}(\text{FS}_t^\pi)$ with $\text{FS}_t^\pi := \{\hat{\xi}_{1t}^\pi < \max_{j \neq 1} \hat{\xi}_{jt}^\pi\}$, is a widely used criterion for the efficiency of a sampling algorithm (see, e.g., the survey paper by Kim and Nelson (2006)). However, the exact evaluation of $\mathbf{P}(\text{FS}_t^\pi)$ under dynamic sampling policies is not analytically tractable. In this subsection, following Glynn and Juneja (2004) we build a tractable objective associated with $\mathbf{P}(\text{FS}_t^\pi)$ based on large deviations theory. In particular, we fix a static allocation rule $\pi = \pi(\boldsymbol{\alpha})$ for some $\boldsymbol{\alpha} \in \Delta$ and characterize how fast $\mathbf{P}(\text{FS}_t^{\pi(\boldsymbol{\alpha})})$ converges to 0 as a function of $\boldsymbol{\alpha} \in \Delta$. We eliminate $\pi(\boldsymbol{\alpha})$ in the superscripts in order to improve clarity.

To begin, observe that N_{jt} is deterministic under the static allocation rule $\pi(\boldsymbol{\alpha})$ and that

$$\mathbf{P}(\hat{\xi}_{jt} > x) = \mathbf{P}\left(\sum_{n=1}^{N_{jt}} \mathbf{I}\{X_{j\tau_{jn}} < x\} < \lfloor pN_{jt} \rfloor\right), \quad (7)$$

where $\lfloor y \rfloor$ is the greatest integer less than y . Note that τ_{jn} is deterministic under the static allocation rule so that $\mathbf{I}\{X_{j\tau_{jn}} < x\}$ are independent Bernoulli random variables for $n = 1, \dots, N_{jt}$. Applying Cramer's theorem (Dembo and Zeitouni 2009, p. 27), the large deviation probability for $\hat{\xi}_{jt}$, $j = 1, \dots, k$, can be characterized as follows (see Lemma B.1 in the online supplement):

$$\begin{aligned} \lim_{t \rightarrow \infty} \frac{1}{t} \log \mathbf{P}(\hat{\xi}_{jt} \geq x) &= -\alpha_j I_j(x) \quad \text{for } x > \xi_j \\ \lim_{t \rightarrow \infty} \frac{1}{t} \log \mathbf{P}(\hat{\xi}_{jt} \leq x') &= -\alpha_j I_j(x') \quad \text{for } x' < \xi_j, \end{aligned} \quad (8)$$

where

$$I_j(x) = p \log \left(\frac{p}{F_j(x)} \right) + (1-p) \log \left(\frac{1-p}{1-F_j(x)} \right). \quad (9)$$

PROPOSITION 2 (Rate function). *Suppose (F1) and (F2) hold. For a static allocation rule $\pi(\boldsymbol{\alpha})$ for some $\boldsymbol{\alpha} \in \Delta^0$,*

$$\lim_{t \rightarrow \infty} \frac{1}{t} \log \mathbf{P}(\text{FS}_t^{\pi(\boldsymbol{\alpha})}) = -\rho(\boldsymbol{\alpha}), \quad (10)$$

with $\rho(\boldsymbol{\alpha}) = \min_{j \neq 1} \{G_j(\boldsymbol{\alpha})\}$, where

$$G_j(\boldsymbol{\alpha}) = \inf_x \{\alpha_1 I_1(x) + \alpha_j I_j(x)\}. \quad (11)$$

We remark that the result of Proposition 2 generalizes the results in Pasupathy et al. (2010). Specifically, the analogous results in Pasupathy et al. (2010) require that $F_j(\cdot)$ is twice differentiable in the continuous case, since their proof relies on a Taylor expansion of the logarithmic moment generating function up to second order terms.

An important implication of Proposition 2 is that $\mathbb{P}(\text{FS}_t^{\pi(\alpha)})$ behaves roughly like $\exp(-\rho(\alpha)t)$ for large values of t . Hence, it follows that $\rho(\cdot)$ is an appropriate measure of asymptotic efficiency that is closely associated with $\mathbb{P}(\text{FS}_t^\pi)$. Note that for each $x \in \mathcal{R}$, $\alpha_1 I_1(x) + \alpha_j I_j(x)$ is a continuous, linear (hence, concave) function of (α_1, α_j) . Since $G_j(\alpha)$ in (11) is a point-wise infimum thereof, it is also concave. Therefore, $\rho(\alpha)$, being a minimum of $G_j(\alpha)$ for $j = 2, \dots, k$, is concave for $\alpha \in \Delta$. We define $\rho^* = \max_{\alpha \in \Delta} \{\rho(\alpha)\}$.

3.3. Problem Formulation

Based on the relationship between $\rho(\alpha)$ and $\mathbb{P}(\text{FS}_t^{\pi(\alpha)})$ provided in Proposition 2, we define the *relative efficiency* \mathcal{R}_t^π for any given algorithm $\pi \in \Pi$ in stage t to be

$$\mathcal{R}_t^\pi = \frac{\rho(\alpha_t^\pi)}{\rho^*}. \quad (12)$$

By definition, the value of \mathcal{R}_t^π lies in the interval $[0, 1]$.

We are interested in designing an algorithm that maximizes the expected relative efficiency with the sampling budget T :

$$\sup_{\pi \in \bar{\Pi}} \mathbb{E}(\mathcal{R}_T^\pi). \quad (13)$$

It is computationally intractable to find an exact solution to the dynamic programming problem in (13) due to the curse of dimensionality. In addition, the underlying distribution functions are not known a priori, hence neither is the function $\rho(\cdot)$ nor its maximizer α^* . This adds another layer of difficulty in searching for the optimal sampling algorithm for the dynamic programming problem with finite T . Alternatively, we focus on the following asymptotic criteria.

DEFINITION 2 (ASYMPTOTIC OPTIMALITY). An algorithm $\pi \in \bar{\Pi}$ is asymptotically optimal if

$$\mathbb{E}(\mathcal{R}_t^\pi) \rightarrow 1 \text{ as } t \rightarrow \infty. \quad (14)$$

It is not difficult to check that an asymptotically optimal algorithm is consistent; otherwise, there exists some system j for which $\alpha_{jt} \rightarrow 0$ as $t \rightarrow \infty$ with positive probability, which in turn implies that the limit of $\mathbb{E}(\mathcal{R}_t^\pi)$ is less than one as $t \rightarrow \infty$.

ALGORITHM 1: Naive(n_0, m, c)

For each j , take n_0 samples and let $t = kn_0$.

repeat

If $N_{jt}^\pi \leq c \log t$ for some j , let $\pi_{t+\ell} = \min_j \{N_{jt}^\pi\}$ for $\ell = 1, \dots, \hat{m}$, with ties broken arbitrarily. Otherwise, solve for $\hat{\alpha}_t \in \arg \max_{\alpha \in \Delta} \{\hat{\rho}_t(\alpha)\}$, with ties broken arbitrarily, and let

$$\pi_{t+\ell} = \arg \max_j \{\hat{\alpha}_{jt} - \alpha_{jt}^\pi\} \quad (17)$$

for $\ell = 1, \dots, m$. Let $t = t + m$.

until $t \leq T$;

Return: $\arg \max_j \{\hat{\xi}_{jT}^\pi\}$.

4. Proposed Algorithms and Main Theoretical Results

4.1. A Naive Algorithm

We first suggest a naive algorithm that iteratively estimates the maximizer of the rate function by utilizing the history of sample observations. Define $\hat{\rho}_t(\alpha) = \min_{j \neq b} \{\hat{G}_{jt}(\alpha)\}$ with $b = \arg \max_j \{\hat{\xi}_{jt}\}$,

$$\hat{G}_{jt}(\alpha) = \inf_x \left\{ \alpha_b \hat{I}_{bt}(x) + \alpha_j \hat{I}_{jt}(x) \right\} \text{ for } j \neq b, \quad (15)$$

and

$$\hat{I}_{jt}(x) = p \log \left(\frac{p}{\hat{F}_{jt}(x)} \right) + (1-p) \log \left(\frac{1-p}{1-\hat{F}_{jt}(x)} \right) \text{ for } j = 1, \dots, k. \quad (16)$$

Note that $\hat{\rho}_t(\alpha)$ is an estimator of $\rho(\alpha)$ in stage t with each $F_j(\cdot)$ replaced with its empirical counterpart, $\hat{F}_{jt}(\cdot)$. Define $\hat{\alpha}_t \in \arg \max_{\alpha \in \Delta} \{\hat{\rho}_t(\alpha)\}$ as a maximizer of the function $\hat{\rho}_t(\alpha)$, with ties broken arbitrarily. The naive algorithm is summarized in Algorithm 1, with n_0 , m , and c being tuning parameters; n_0 is the initial number of samples from each system, m is the batch size, and c controls the minimum sampling frequency. For ease of exposition, we assume T is a multiple of m .

The naive algorithm makes the current allocation (α_t^π) close to the target allocation $(\hat{\alpha}_t)$ in each stage; see equation (17). Essentially, it attempts to sample the system with greatest need of information; see, e.g., Chen and Lee (2011) for similar procedures of this kind. Note also that the naive algorithm (as well as those to be introduced in what follows) makes no use of the value of T , the problem horizon. A more forward looking approach typically requires dynamic programming, which may not be practically implementable and hence will not be discussed in this paper. Readers interested in dynamic programming approaches are referred to Peng et al. (2018a).

The following theorem shows that $\hat{\alpha}_t$ eventually approaches the optimal allocation α^* , and therefore α_t^π converges to α^* as $t \rightarrow \infty$, implying the algorithm is asymptotically optimal.

THEOREM 1 (Asymptotic performance of the naive algorithm). *Suppose $\mathbf{F} \in \{\mathcal{C}, \mathcal{D}\}$. Under (F1)-(F2), the naive algorithm is consistent and asymptotically optimal.*

REMARK 1 (POOR PERFORMANCE OF THE NAIVE ALGORITHM). From an implementation perspective, the naive algorithm requires a nested optimization in each stage. Specifically, the inner optimization loop involves solving for the infimum in (15) to find $\hat{G}_{jt}(\alpha)$ for given α , and the outer optimization task is to maximize $\hat{\rho}_t(\alpha) = \min_{j \neq b} \{\hat{G}_{jt}(\alpha)\}$ over $\alpha \in \Delta$. As we will see in Table 1 in §5.1, this becomes a significant computational burden as the number of systems increases. From a performance standpoint, despite the theoretical guarantee on its asymptotic performance, the naive algorithm exhibits poor finite-time performance in terms of the probability of false selection. In particular, in the naive algorithm we need to estimate the distribution functions in order to evaluate the rate function, which results in relatively poor performance; see numerical results in §6.

REMARK 2 (FORCED SAMPLING IN THE NAIVE ALGORITHM). An issue that can arise with the rate function estimator $\hat{\rho}_t(\cdot)$ based on the empirical distribution functions is a long period of no sampling from a particular system. In this case, the maximizer $\hat{\alpha}_t \in \arg \max_{\alpha \in \Delta} \{\hat{\rho}_t(\alpha)\}$ may lie at the boundary of Δ ; that is, we may have $\hat{\alpha}_{jt} = 0$ for some j . In this case, the naive algorithm does not sample system j in stage t . If this problem persists in subsequent stages, the empirical distribution may not converge to its theoretical counterpart “fast enough”. In order to prevent this issue, the naive algorithm ensures that each system is sampled at least $c \log t$ times in stage t . Note that the asymptotic result in Theorem 1 holds when $c \log t$ is replaced with any sublinear function of t , although such a change would affect the performance with finite t .

REMARK 3 (VARIANTS OF THE NAIVE ALGORITHM). In this paper we use the rate function estimator $\hat{\rho}_t(\cdot)$ based on empirical distribution functions. In cases with $\mathbf{F} \in \mathcal{C}$, one can also estimate the rate function based on “smooth” estimators; for example, see the kernel-type estimator in Reiss (1981). In cases with $\mathbf{F} \in \mathcal{D}$, one can also use discrete kernel estimators; for example, see Rajagopalan and Lall (1995). Smooth estimators can be useful when the sample observations are sparse, although the efficiency of these estimators depends on the choice of smoothing parameters whose optimal values are unknown

a priori. It is not our intention to compare different types of distribution estimators, but we remark that the performance of the naive algorithm may vary with different choices of these estimators.

4.2. Alternative Algorithm for Continuous Distributions

We discuss here the case in which the underlying distributions are continuous with the following condition.

(F3) $F_j(\cdot)$ is twice continuously differentiable and possesses a positive continuous density $f_j(\cdot)$ over the interval \mathcal{H}_j .

It is trivial to check that the smoothness assumption (F3) implies (F1) and (F2). Note that $F_j(\cdot)$ must be twice continuously differentiable in order to ensure that $I_j(x)$ can be closely approximated, using a Taylor expansion, by a quadratic function of x in a neighborhood of ξ_j , which will be key to the theoretical results in this subsection.

Let $\delta := \xi_1 - \xi_2$ be the gap between the best and the second best systems and define $\rho^\delta(\boldsymbol{\alpha}) := \min_{j \neq 1} \{G_j^\delta(\boldsymbol{\alpha})\}$ with

$$G_j^\delta(\boldsymbol{\alpha}) = \frac{(\xi_1 - \xi_j)^2}{2p(1-p) \left(1/(\alpha_1 f_1^2(\xi_1)) + 1/(\alpha_j f_j^2(\xi_j))\right)}. \quad (18)$$

We provide some intuition behind the definition of $\rho^\delta(\boldsymbol{\alpha})$. Under (F3), the sample quantile from $\alpha_j T$ independent observations from system j is asymptotically normal with mean ξ_j and variance $p(1-p)/(\alpha_j T f_j^2(\xi_j))$ (see, e.g., pp. 77-79 of Serfling 2009). Hence, one may consider $G_j^\delta(\boldsymbol{\alpha})$ as a measure of divergence between two quantiles, measured in units of standard errors. The greater the value of $G_j^\delta(\boldsymbol{\alpha})$ the better we can distinguish whether $\xi_1 > \xi_j$ or not with greater confidence. This in turn suggests that $\rho^\delta(\boldsymbol{\alpha})$ is closely aligned with $\rho(\boldsymbol{\alpha})$, the rate function of the probability of false selection. This intuitive observation is formalized in the following proposition.

PROPOSITION 3 (Characteristics of ρ^δ). *Consider a set of system configurations in \mathcal{C} , where each configuration \mathbf{F} satisfies (F1)-(F3) and $f_j(\xi_j)$ is in a compact set that does not include zero. Then, $\rho^\delta(\boldsymbol{\alpha})$ has a unique maximum $\boldsymbol{\alpha}^\delta \in \Delta^0$ and*

$$\frac{\rho(\boldsymbol{\alpha}^\delta)}{\rho^*} \rightarrow 1 \text{ as } \delta \rightarrow 0. \quad (19)$$

Proposition 3 states that one can achieve near-optimal performance with respect to $\rho(\boldsymbol{\alpha})$ by maximizing $\rho^\delta(\boldsymbol{\alpha})$ when δ is sufficiently close to 0. Note that the maximization of

$\rho^\delta(\boldsymbol{\alpha})$ is much simpler than that of $\rho(\boldsymbol{\alpha})$ because the former is strongly concave in $\boldsymbol{\alpha}$ (Appendix A.1) and does not involve a nested optimization structure. Further, using Theorem 2 of Peng et al. (2018b) one can show that the maximization problem $\max_{\boldsymbol{\alpha} \in \Delta} \{\rho^\delta(\boldsymbol{\alpha})\}$ can be transformed into a one-dimensional convex optimization problem, which can alleviate the computational complexity when the number of alternatives is large. In Proposition 3, the condition that $f_j(\xi_j)$ lies in a compact set is a relatively mild restriction. We provide three examples to illustrate the nature of this condition.

EXAMPLE 2 (EXPONENTIAL DISTRIBUTIONS). First, consider two exponential systems with means $(\mu + \delta, \mu)$ for $\mu > 0$, for which $(f_1(\xi_1), f_2(\xi_2)) \rightarrow ((1-p)/\mu, (1-p)/\mu)$ as $\delta \rightarrow 0$, satisfying the condition in Proposition 3. Second, consider two exponential systems with means $(\delta + 1/\delta, 1/\delta)$, for which $(\xi_1, \xi_2) = (-(\delta + 1/\delta) \log(1-p), -(1/\delta) \log(1-p))$. In this case, $(f_1(\xi_1), f_2(\xi_2)) \rightarrow (0, 0)$ as $\delta \rightarrow 0$, violating the condition in Proposition 3. Third, consider two exponential systems with means $(2\delta, \delta)$, for which $(\xi_1, \xi_2) = (-2\delta \log(1-p), -\delta \log(1-p))$. In this case, $(f_1(\xi_1), f_2(\xi_2)) \rightarrow (\infty, \infty)$ as $\delta \rightarrow 0$, violating the condition of Proposition 3.

Based on the approximation $\rho^\delta(\cdot)$, we now propose an algorithm called Quantile Divergence (QD) for continuous distributions that iteratively estimates $\boldsymbol{\alpha}^\delta = \arg \max_{\boldsymbol{\alpha} \in \Delta} \{\rho^\delta(\boldsymbol{\alpha})\}$ from the history of sample observations. Specifically, denote $\hat{\boldsymbol{\alpha}}_t^\delta$ the estimator of $\boldsymbol{\alpha}^\delta$ in stage t . Formally, $\hat{\boldsymbol{\alpha}}_t^\delta = \arg \max_{\boldsymbol{\alpha} \in \Delta} \{\hat{\rho}_t^\delta(\boldsymbol{\alpha})\}$, where

$$\hat{\rho}_t^\delta(\boldsymbol{\alpha}) = \min_{j \neq b} \frac{(\hat{\xi}_{bt} - \hat{\xi}_{jt})^2}{2p(1-p) \left(1/(\alpha_b \hat{f}_{bt}^2(\hat{\xi}_{bt})) + 1/(\alpha_j \hat{f}_{jt}^2(\hat{\xi}_{jt})) \right)}, \quad (20)$$

with $b = \arg \max_j \{\hat{\xi}_{jt}\}$ and

$$\hat{f}_{jt}(y) = \frac{1}{N_{jt}} \sum_{n=1}^{N_{jt}} K_{h(N_{jt})}(y - X_{j\tau_{jn}}) \quad (21)$$

is the kernel-based estimator with kernel $K(\cdot)$ and bandwidth $h(t) > 0$ for each t . A kernel with subscript h is called a scaled kernel and defined as $K_h(x) = h^{-1}K(x/h)$. The optimal choices of the kernel function and the bandwidth parameter depends on the true density functions that are a priori unknown (see, e.g., Silverman (1986)), but we impose standard regularity conditions on $K(\cdot)$ and $h(t)$, which are satisfied by almost any conceivable kernel such as normal, uniform, triangular, and others:

ALGORITHM 2: QD-C ($n_0, m, K(\cdot), h(\cdot)$)

For each j , take n_0 samples and let $t = kn_0$.

repeat

Let $b = \arg \max_j \{\hat{\xi}_{jt}^\pi\}$. If $\hat{\xi}_{bt}^\pi = \hat{\xi}_{jt}^\pi$ for system $j \neq b$, then take a sample from system

$$\pi_{t+\ell} = \arg \min_{i=j,b} \{\alpha_{it}^\pi\} \quad (22)$$

for $\ell = 1, \dots, m$, with ties broken arbitrarily. Otherwise, estimate the density estimators in (21) using the kernel $K(\cdot)$ and the bandwidth parameter $h(\cdot)$. Solve for $\hat{\alpha}_t^\delta = \arg \max_{\alpha \in \Delta} \{\hat{\rho}_t^\delta(\alpha)\}$ and let

$$\pi_{t+\ell} = \arg \max_j \{\hat{\alpha}_{jt}^\delta - \alpha_{jt}^\pi\} \quad (23)$$

for $\ell = 1, \dots, m$. Let $t = t + m$.

until $t \leq T$;

Return: $\arg \max_j \{\hat{\xi}_{jT}^\pi\}$.

$$(K1) \int |K(x)|dx < \infty \text{ and } \int K(x)dx = 1$$

$$(K2) |xK(x)| \rightarrow 0 \text{ as } |x| \rightarrow \infty$$

$$(K3) h(t) \rightarrow 0 \text{ and } th(t) \rightarrow \infty \text{ as } t \rightarrow \infty.$$

We propose an algorithm that matches α_t^π with $\hat{\alpha}_t^\delta$ in each stage, simultaneously ensuring that $\hat{\alpha}_t^\delta$ approaches α^δ as $t \rightarrow \infty$. The procedure is summarized in Algorithm 2, with n_0 and m being tuning parameters.

THEOREM 2 (Asymptotic performance of QD-C). *Suppose $\mathbf{F} \in \mathcal{C}$. Under (F1)-(F2), the QD-C algorithm is consistent, and if (F3) is further satisfied, then*

$$\mathbb{E}(\mathcal{R}_t^\pi) \rightarrow \frac{\rho(\alpha^\delta)}{\rho^*} \text{ as } t \rightarrow \infty. \quad (24)$$

To rephrase (24), the QD algorithm eventually allocates the sampling budget so that $\rho^\delta(\alpha)$ is maximized. Further, from Proposition 3 the loss in asymptotic efficiency due to maximizing $\rho^\delta(\cdot)$ instead of $\rho(\cdot)$ is negligible for a fixed but sufficiently small δ .

REMARK 4 (RELATION TO EXISTING ALGORITHMS). The QD-C algorithm shares the common theme with two sampling algorithms in the literature of mean-based ordinal optimization. First, QD-C is analogous to the WD algorithm proposed by Shin et al. (2018) in the sense that the former (respectively, the latter) repeatedly maximizes the divergence between sample quantiles (respectively, sample means). Second, QD-C can be related to

a family of sampling algorithms called OCBA (Chen and Lee 2011). Specifically, from the first order conditions it can be seen that $\hat{\alpha}_t^\delta$ of the QD-C algorithm satisfies

$$\frac{(\hat{\xi}_{bt} - \hat{\xi}_{it})^2}{(\hat{\alpha}_{bt}^\delta \hat{f}_{bt}^2(\hat{\xi}_{bt}))^{-1} + (\hat{\alpha}_{it}^\delta \hat{f}_{it}^2(\hat{\xi}_{it}))^{-1}} = \frac{(\hat{\xi}_{bt} - \hat{\xi}_{jt})^2}{(\hat{\alpha}_{bt}^\delta \hat{f}_{bt}^2(\hat{\xi}_{bt}))^{-1} + (\hat{\alpha}_{jt}^\delta \hat{f}_{jt}^2(\hat{\xi}_{jt}))^{-1}} \quad \text{for } i, j \neq b \quad (25)$$

$$\hat{\alpha}_{bt}^\delta = \sqrt{\sum_{j \neq b} (\hat{\alpha}_{jt}^\delta)^2 \frac{\hat{f}_{jt}^2(\hat{\xi}_{jt})}{\hat{f}_{bt}^2(\hat{\xi}_{bt})}}. \quad (26)$$

If we further assume that $\hat{\alpha}_{bt}^\delta \gg \hat{\alpha}_{jt}^\delta$, then the above formulas are equivalent to the OCBA formulas, except that the means and variances of sample means are replaced with those of sample quantiles. Both WD and OCBA are derived based on the premise that sample means are normal, and hence, QD-C can be considered as a comparable algorithm when sample quantiles are approximately normally distributed.

REMARK 5 (BIAS-VARIANCE TRADEOFF). The major advantage of QD-C over the naive algorithm is that the estimation of $\rho^\delta(\cdot)$ is more “localized” and does not require the estimation of the entire distribution functions; see Remark 1 for a discussion about the optimization error under the naive algorithm. In other words, exploration of the function $\rho^\delta(\cdot)$ requires less of the sampling budget than that of $\rho(\cdot)$, which allows us to exploit more of the budget to maximize $\rho^\delta(\cdot)$, with a sacrifice due to the gap between $\rho^\delta(\cdot)$ and $\rho(\cdot)$. We now rigorously analyze the bias and variance in this tradeoff. For ease of exposition, let $k = 2$ and, with a slight abuse of notation, denote $\rho(\alpha_{1t}^\pi) = \rho(\alpha_t^\pi)$ for $\alpha_t^\pi = (\alpha_{1t}^\pi, \alpha_{2t}^\pi)$. Using a second-order Taylor expansion of $\rho(\cdot)$ at α_1^* , observe that

$$\rho(\alpha_1^*) - \rho(\alpha_{1t}^\pi) = -\frac{\rho''(\alpha_1^*)}{2}(\alpha_{1t}^\pi - \alpha_1^*)^2 + o((\alpha_{1t}^\pi - \alpha_1^*)^2), \quad (27)$$

where $\rho''(\alpha_1^*)$ is the second derivative of $\rho(\cdot)$ at α_1^* , and the term $o((\alpha_{1t}^\pi - \alpha_1^*)^2)$ is sublinear in $(\alpha_{1t}^\pi - \alpha_1^*)^2$ as this squared difference tends to zero. Hence, the loss in the asymptotic efficiency, $\mathbb{E}(\rho(\alpha_1^*) - \rho(\alpha_{1t}^\pi))$, can be largely explained by the mean squared error (MSE), $\mathbb{E}(\alpha_{1t}^\pi - \alpha_1^*)^2$, which can be further decomposed as

$$\mathbb{E}(\alpha_{1t}^\pi - \alpha_1^*)^2 = \mathbb{E}(\alpha_{1t}^\pi - \mathbb{E}(\alpha_{1t}^\pi))^2 + (\mathbb{E}(\alpha_{1t}^\pi) - \alpha_1^*)^2. \quad (28)$$

The first term represents the variance component (due to noisy sample estimates) and the second represents the bias component (due to maximizing $\rho^\delta(\cdot)$ instead of $\rho(\cdot)$). As illustrated in Figure 1, most of the MSE contribution is due to variance under the naive

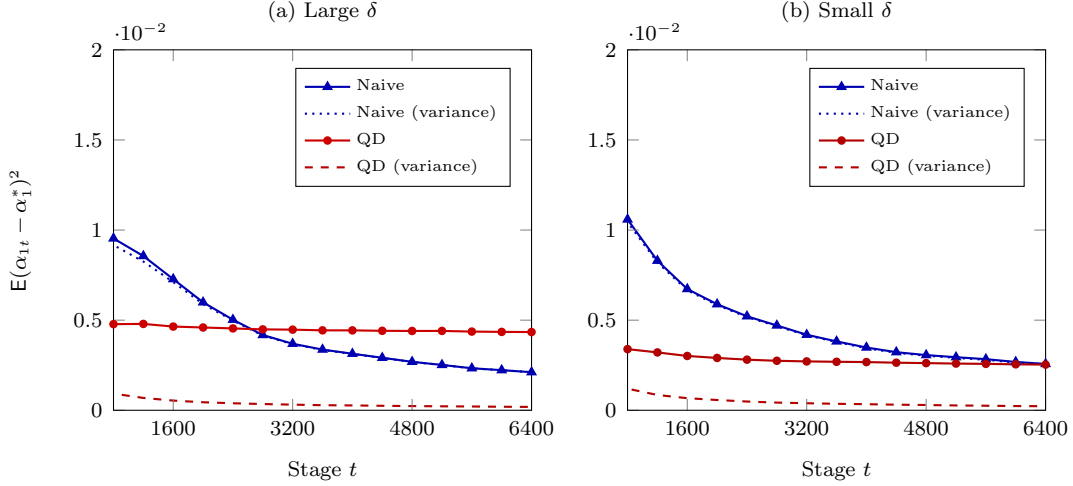


Figure 1 Bias-variance tradeoff. For each naive (blue) and QD-C (red) algorithm, the solid line is the mean squared error and the dotted (or dashed) line is its variance component. The gap between the solid and dotted (or dashed) lines corresponds to the bias component. The system configurations are (a) two normal systems with means $(0, 0)$ and standard deviations $(1, 3)$ and (b) two normal systems with the same means $(0, 0)$ and standard deviations $(1, 2)$. The 10% quantiles are $\xi = (0.13, 0.38)$ for (a) and $\xi = (0.13, 0.25)$ for (b). For both algorithms $(n_0, m) = (20, 10)$. Also, $c = 10$ for the naive algorithm and we use the normal kernel with the bandwidth parameter $h(n) = 1.06\sigma n^{-1/5}$ for QD-C, where σ is replaced with sample standard deviation.

algorithm. On the other hand, QD-C significantly reduces the variance, while introducing a small bias that decreases to zero as $\delta \rightarrow 0$. Overall, the QD-C algorithm has lower MSE's than the naive counterpart for small t , indicating it is more suitable in settings with small sampling budgets. Obviously, QD-C is even more attractive than the naive algorithm when δ is “small” because the loss due to the bias is “small” as well.

4.3. Alternative Algorithm for Discrete Distributions

We now develop an alternative algorithm for the case in which the underlying distributions are discrete, i.e., $\mathbf{F} \in \mathcal{D}$. To this end, define

$$\begin{aligned}
 h_{j1} &= \frac{F_j(\xi_1) - F_j(\xi_j)}{\xi_1 - \xi_j} \\
 h_{1j} &= \frac{F_1(\xi_1) - F_1(\xi_j)}{\xi_1 - \xi_j}
 \end{aligned} \tag{29}$$

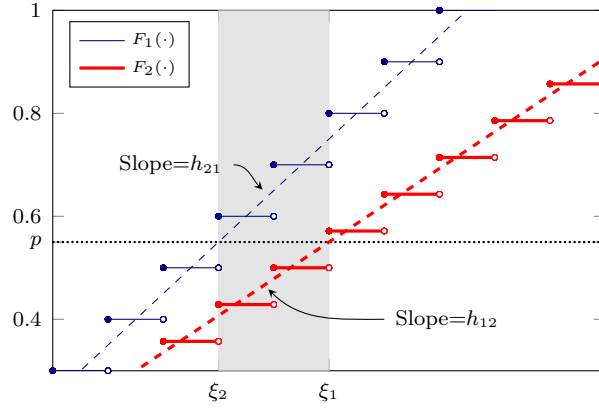


Figure 2 Illustration of h_{12} and h_{21} defined in (29) for $p = 0.55$. The step functions in thin and thick solid lines represent the distributions $F_1(\cdot)$ and $F_2(\cdot)$, respectively, and the dashed lines are the linear approximations.

for each $j \neq 1$ and let $\mathbf{h} = \{(h_{j1}, h_{1j}) \mid j \neq 1\}$. Figure 2 illustrates the definitions of h_{j1} and h_{1j} for $j = 2$. Also define $\epsilon \in (0, 1)$ as a constant such that, for each $j \neq 1$ and $x \in [\xi_j, \xi_1]$,

$$\begin{aligned}
 1 - \epsilon &\leq \frac{F_j(x)}{p + (x - \xi_j)h_{j1}}, \frac{1 - F_j(x)}{1 - p - (x - \xi_j)h_{j1}} \leq 1 + \epsilon \\
 1 - \epsilon &\leq \frac{F_1(x)}{p + (x - \xi_1)h_{1j}}, \frac{1 - F_1(x)}{1 - p - (x - \xi_1)h_{1j}} \leq 1 + \epsilon.
 \end{aligned} \tag{30}$$

The constant ϵ represents (multiplicative) errors when $F_j(x)$ is approximated by a linear function, $p + (x - \xi_j)h_{j1}$ or $p + (x - \xi_1)h_{1j}$.

We consider a set of discrete distributions that satisfies the following condition, which ensures that h_{j1} and h_{1j} defined in (29) are strictly positive for each $j \neq 1$.

(F3') The probability mass function, $f_j(x) = F_j(x) - F_j(x - 1)$, is positive for all integer $x \in \mathcal{H}_j$

Define $\rho^{\delta, \epsilon}(\boldsymbol{\alpha}) = \min_{j \neq 1} \{G_j^{\delta, \epsilon}(\boldsymbol{\alpha})\}$ with

$$G_j^{\delta, \epsilon}(\boldsymbol{\alpha}) = \frac{(\xi_1 - \xi_j)^2}{2p(1-p) \left(\frac{1}{\alpha_1 h_{1j}^2} + \frac{1}{\alpha_j h_{j1}^2} \right)}, \tag{31}$$

which has a similar structure as (18) in the continuous case, except that $f_j(\xi_j)$ is replaced with h_{j1} or h_{1j} . To provide some intuition behind the definition of $G_j^{\delta, \epsilon}(\boldsymbol{\alpha})$, recall that the asymptotic variance of $\hat{\xi}_{jt}$ is inversely proportional to $f_j(\xi_j)$ in the continuous case. In other words, when h_{1j} or h_{j1} is high (low), samples from system j are denser (sparser) around ξ_j so that the asymptotic variance of the sample quantile is low (high). Therefore, $G_j^{\delta, \epsilon}(\boldsymbol{\alpha})$ can be seen as an appropriate measure of divergence between two quantiles. The following proposition validates this intuition rigorously. We use the following notation:

$\bar{h} = \max_{j \neq 1} \{\max(h_{j1}, h_{1j})\}$, $r = \max_{j \neq 1} \{\xi_1 - \xi_j\} / \min_{j \neq 1} \{\xi_1 - \xi_j\}$, $\theta = r\bar{h}\delta$, and $u: \mathcal{R}_+ \rightarrow \mathcal{R}$ is a non-decreasing function of θ defined as

$$u(\theta) = \frac{2\theta p(1-p)(\theta^3 + 3\theta p(1-p) + p(1-p))}{3(p-\theta)^3(1-p-\theta)^3}. \quad (32)$$

PROPOSITION 4 (**Characteristic of $\rho^{\delta, \epsilon}(\cdot)$**). *Suppose $\mathbf{F} \in \mathcal{D}$. Under (F1)-(F2) and (F3'), for any $\alpha \in \Delta$*

$$\frac{\rho(\alpha^{\delta, \epsilon})}{\rho^*} \geq \frac{(1-u(\theta))(1-\epsilon)}{(1+u(\theta))(1+\epsilon)} \quad (33)$$

for $\theta = r\bar{h}\delta$ sufficiently small so that $u(\theta) < 1$.

Proposition 4 validates the approximation of $\rho(\cdot)$ by $\rho^{\delta, \epsilon}(\cdot)$ for small δ (equivalently, small θ) and ϵ . This in turn implies α^* and $\alpha^{\delta, \epsilon}$, the maximizers of $\rho(\cdot)$ and $\rho^{\delta, \epsilon}(\cdot)$, respectively, are close when δ and ϵ are sufficiently small.

REMARK 6 (TIGHTNESS OF THE LOWER BOUND IN PROPOSITION 4). We note that $u(\theta)$ is a continuous, increasing function of $\theta \geq 0$ with $u(0) = 0$. Hence, when θ and ϵ are close to 0, the bound in (33) is tight. To see the tightness of the bound for different values of θ , we provide a numerical example in Figure 3 for the case of discrete uniform distributions. Note that the lower bound is tight for small values of θ , and while there is a noticeable gap for large values of θ , the actual values of $\rho(\alpha^{\delta, \epsilon})/\rho^*$ are still close to one.

We now present the QD-D algorithm, which iteratively maximizes $\rho^{\delta, \epsilon}(\cdot)$ from the history of sample observations. Let $\hat{\alpha}_t^{\delta, \epsilon} = \arg \max_{\alpha \in \Delta} \{\hat{\rho}_t^{\delta, \epsilon}(\alpha)\}$ be the estimator of $\alpha^{\delta, \epsilon*}$ in stage t , where

$$\hat{\rho}_t^{\delta, \epsilon}(\alpha) = \min_{j \neq b} \frac{(\hat{\xi}_{bt} - \hat{\xi}_{jt})^2}{2p(1-p) \left(1/(\alpha_b \hat{h}_{bjt}^2) + 1/(\alpha_j \hat{h}_{jbt}^2) \right)}, \quad (34)$$

with $b = \arg \max_j \{\hat{\xi}_{jt}\}$ and \hat{h}_{jbt} and \hat{h}_{bjt} defined as

$$\hat{h}_{jbt} = \begin{cases} (\hat{F}_{jt}(\hat{\xi}_{bt}) - \hat{F}_{jt}(\hat{\xi}_{jt})) / (\hat{\xi}_{bt} - \hat{\xi}_{jt}) & \text{if } \hat{F}_{jt}(\hat{\xi}_{bt}) - \hat{F}_{jt}(\hat{\xi}_{jt}) > 0 \\ (\hat{F}_{jt}(y_{jbt}) - \hat{F}_{jt}(\hat{\xi}_{jt})) / (y_{jbt} - \hat{\xi}_{jt}) & \text{otherwise} \end{cases} \quad (35)$$

$$\hat{h}_{bjt} = \begin{cases} (\hat{F}_{bt}(\hat{\xi}_{bt}) - \hat{F}_{bt}(\hat{\xi}_{jt})) / (\hat{\xi}_{bt} - \hat{\xi}_{jt}) & \text{if } \hat{F}_{bt}(\hat{\xi}_{bt}) - \hat{F}_{bt}(\hat{\xi}_{jt}) > 0 \\ (\hat{F}_{bt}(\hat{\xi}_{bt}) - \hat{F}_{bt}(y_{bjt})) / (\hat{\xi}_{bt} - y_{bjt}) & \text{otherwise,} \end{cases}$$

where $y_{jbt} := \inf\{y \geq \hat{\xi}_{bt} \mid \hat{F}_{jt}(y) - \hat{F}_{jt}(\hat{\xi}_{jt}) > 0\}$ is the smallest value of $y \geq \hat{\xi}_{bt}$ such that the empirical distribution $\hat{F}_{jt}(\cdot)$ has positive probability mass between $\hat{\xi}_{jt}$ and y . Similarly,

ALGORITHM 3: QD-D (n_0, m)

For each j , take n_0 samples and let $t = kn_0$

repeat

Let $b = \arg \max_j \{\hat{\xi}_{jt}^\pi\}$. If $\hat{\xi}_{bt}^\pi = \hat{\xi}_{jt}^\pi$ for system $j \neq b$, then take a sample from system

$\pi_{t+\ell} = \arg \min_{i=j,b} \{\alpha_{it}^\pi\}$ for $\ell = 1, \dots, m$. If there are multiple such j 's, choose the one with the smallest j . Otherwise, solve for $\hat{\alpha}_t^{\delta, \epsilon} = \arg \max_{\alpha \in \Delta} \{\hat{\rho}_t^{\delta, \epsilon}(\alpha)\}$ and let

$$\pi_{t+1} = \arg \max_j \{\hat{\alpha}_{jt}^{\delta, \epsilon} - \alpha_{jt}^\pi\} \quad (36)$$

for $\ell = 1, \dots, m$. Let $t = t + m$

until $t \leq T$;

Return: $\arg \max_j \{\hat{\xi}_{jT}^\pi\}$.

$y_{bjt} := \sup\{y \leq \hat{\xi}_{jt} | \hat{F}_{bt}(\hat{\xi}_{bt}) - \hat{F}_{bt}(y_{bjt}) > 0\}$. This ensures $\hat{h}_{jbt}, \hat{h}_{bjt} > 0$ for each $j \neq b$. The QD-D algorithm for the discrete case is summarized in Algorithm 3, with n_0 and m being parameters of the algorithm.

We note that the event $\{\hat{\xi}_{bt} = \hat{\xi}_{jt}\}$ for some $j \neq b$ can occur with positive probability, in which case $\hat{\alpha}_t^{\delta, \epsilon} = \arg \max_{\alpha \in \Delta} \{\hat{\rho}_t^{\delta, \epsilon}\}$ may not be well defined. When these cases occur, the QD-D algorithm takes a sample from system j or b , whichever was sampled less. When there are multiple systems satisfying such conditions, choose j arbitrarily.

THEOREM 3 (Asymptotic performance of QD-D). *Suppose $\mathbf{F} \in \mathcal{D}$. Under (F1)-(F2), the QD-D algorithm is consistent, and if (F3') is further satisfied, then*

$$\mathbb{E}(\mathcal{R}_t^\pi) \rightarrow \frac{\rho(\alpha^{\delta, \epsilon})}{\rho^*} \text{ as } t \rightarrow \infty. \quad (37)$$

Note that the preceding theorem, combined with Proposition 4 implies that the QD-D algorithm is near-optimal in an asymptotic regime as $t \rightarrow \infty$ and $\theta \rightarrow 0$.

5. Adaptive Heuristics for Fast Implementation

From an implementation standpoint, the QD algorithms (QD-C and QD-D) significantly improve the computation time of the naive algorithm by replacing the two-layer optimization problem with a single-layer one. However, there still exists a heavy computational burden for problems with a large number of alternatives and a large sampling budget since the convex optimization problem should be solved repeatedly in every stage. (See Table 1 later in this section for computation times of the proposed algorithms as the number of systems increases.) Hence, we suggest variants of the QD algorithms for continuous and discrete distribution that are practically implementable with a large number of systems.




Figure 3 Asymptotic performance of the QD-D algorithm, $\rho(\alpha^{\delta, \epsilon^*})/\rho^*$, and its lower bound defined in (37), as a function of $\theta = r\bar{h}\delta$. The system configuration consists of three systems with discrete uniform distributions with 10% quantiles $(0, -\delta, -2\delta)$, where the values of δ range from 10 to 100 and the length of supports is 10^4 for all three systems. The parameters for the QD-D algorithm are $m = 10$ and $n_0 = 20$.

ALGORITHM 4: AQD-C ($n_0, m, K(\cdot), h(\cdot)$)

For each j , take n_0 samples and let $t = kn_0$

repeat

Let $b = \arg \max_j \{\hat{\xi}_{jt}^\pi\}$. If $\hat{\xi}_{bt}^\pi = \hat{\xi}_{jt}^\pi$ for system $j \neq b$, then take a sample from system

$$\pi_{t+\ell} = \arg \min_{i=j,b} \{\alpha_{it}^\pi\} \quad (38)$$

for $\ell = 1, \dots, m$, with ties broken arbitrarily. If

$$\alpha_{bt}^\pi < \sqrt{\sum_{j \neq b} (\alpha_{jt}^\pi)^2 \frac{(\hat{f}_{jt}^\pi(\hat{\xi}_{jt}^\pi))^2}{(\hat{f}_{bt}^\pi(\hat{\xi}_{bt}^\pi))^2}}, \quad (39)$$

set $\pi_{t+\ell} = b$ for $\ell = 1, \dots, m$. Otherwise, estimate the density estimators in (21) using the kernel $K(\cdot)$ and the bandwidth parameter $h(\cdot)$. Set

$$\pi_{t+\ell} = \arg \min_{j \neq b} \left\{ \frac{(\hat{\xi}_{bt}^\pi - \hat{\xi}_{jt}^\pi)^2}{1/(\alpha_{bt}^\pi (\hat{f}_{bt}^\pi(\hat{\xi}_{bt}^\pi))^2) + 1/(\alpha_{jt}^\pi (\hat{f}_{jt}^\pi(\hat{\xi}_{jt}^\pi))^2)} \right\} \quad (40)$$

for $\ell = 1, \dots, m$. Let $t = t + m$

until $t \leq T$;

Return: $\arg \max_j \{\hat{\xi}_{jT}^\pi\}$.

In the case with continuous distributions, the adaptive variant of the QD-C algorithm, AQD-C, is summarized in Algorithm 4. To provide some intuition behind AQD-C, recall the first order conditions for the maximizer of $\hat{\rho}_t^\delta(\alpha)$ in (25)-(26). While QD-C solves for (25)-(26) in every stage, AQD-C attempts to balance the left and right sides of (25)-(26) so

ALGORITHM 5: AQD-D (n_0, m)

For each j , take n_0 samples and let $t = kn_0$

repeat

Let $b = \arg \max_j \{\hat{\xi}_{jt}^\pi\}$. If $\hat{\xi}_{bt}^\pi = \hat{\xi}_{jt}^\pi$ for some $j \neq b$, then take a sample from system

$\pi_{t+\ell} = \arg \min_{i=j,b} \{\alpha_{it}^\pi\}$ for $\ell = 1, \dots, m$. Otherwise, set $\pi_{t+\ell} = b$ for $\ell = 1, \dots, m$ if

$$\alpha_{bt}^\pi < \sqrt{\sum_{j \neq b} (\alpha_{jt}^\pi)^2 \frac{(\hat{h}_{jbt}^\pi)^2}{(\hat{h}_{bjt}^\pi)^2}}. \quad (41)$$

Otherwise, set

$$\pi_{t+\ell} = \arg \min_{j \neq b} \left\{ \frac{(\hat{\xi}_{bt}^\pi - \hat{\xi}_{jt}^\pi)^2}{1/(\alpha_{bt}^\pi (\hat{h}_{bjt}^\pi)^2) + 1/(\alpha_{jt}^\pi (\hat{h}_{jbt}^\pi)^2)} \right\} \quad (42)$$

for $\ell = 1, \dots, m$. Let $t = t + m$

until $t \leq T$;

Return: $\arg \max_j \{\hat{\xi}_{jT}^\pi\}$.

that they are satisfied asymptotically as $t \rightarrow \infty$. The adaptive variant of the QD algorithm for discrete case (AQD-D) is similarly defined in Algorithm 5, with $\hat{f}_{jt}(\cdot)$ in AQD-C replaced with \hat{h}_{jbt} or \hat{h}_{bjt} .

5.1. Discussion

Comparison to QD. In the case with $k = 2$ systems, the QD-C and AQD-C algorithms are identical. To see this, note that, when (39) is satisfied for some stage t , AQD-C takes a sample from system b . Further, (39) implies $\hat{\alpha}_t^\delta = \arg \max_{\alpha \in \Delta} \{\hat{\rho}_t^\delta\}$ satisfies $\hat{\alpha}_{bt}^\delta > \alpha_{bt}^\pi$ and $\hat{\alpha}_{jt}^\delta < \alpha_{jt}^\pi$ for $j \neq b$, and therefore, QD-C also samples system b . Since the preceding argument is true for each t , one can conclude that QD-C and AQD-C are identical when $k = 2$. The case with $k > 2$ is less obvious, but we show via numerical example that the allocations under QD-C and AQD-C coincide eventually as $t \rightarrow \infty$ (Figure 4). Observe that the distributions of α_{1t}^π under the QD-C and AQD-C algorithms are almost identical. Also, Figure 4 illustrates that $\alpha_{1t}^\pi \rightarrow \alpha_1^* = 0.33$ as $t \rightarrow \infty$ for all three algorithms, but the convergence rate under the naive algorithm is slower than those of the QD-C and AQD-C, which is consistent with Remark 1. A similar argument applies between the QD-D and AQD-D algorithms.

Scalability. Compared to the naive and two QD algorithms, the AQD algorithms exhibit orders of magnitude decrease in CPU time for larger k . Table 1 summarizes the computation time of the algorithms as a function of the number of systems with normal distributions with zero means and 10% percentiles $\xi_1 - \xi_j = 0.13$ for all $j \neq 1$. The relative CPU times




Figure 4 The sampling frequencies (α_{1t}) under the naive, QD-C, and AQD-C algorithms. The configuration is characterized by $k = 4$ normally distributed systems with $\mu = (0, \dots, 0)$ and $\sigma = (1, \dots, 1.2)$. In this case, the 10% quantiles are $\xi = (0.13, 0.38, 0.38, 0.38)$ and $\alpha_1^* = 0.33$. For all algorithms $(n_0, m) = (20, 10)$. Also, $c = 10$ for the naive algorithm and we use the normal kernel with the bandwidth parameter $h(n) = 1.06\sigma n^{-1/5}$ for QD-C and AQD-C, where σ is replaced with sample standard deviation.

Table 1 CPU times (in seconds) per stage. CPU times are estimated by taking averages over 100 simulation trials. (Operating system: Windows 7; processor: Intel core i7 2.7GHz; memory: 32GB RAM; software: Matlab.)

Algorithms	$k = 4$	$k = 40$	$k = 400$	$k = 4000$
Naive	0.057	0.064	9.287	2379
QD	0.012	0.043	4.567	206
AQD	0.003	0.003	0.010	0.135

illustrate the dramatic difference between the AQD and the other algorithms, in particular scalability of the AQD algorithms. Indeed, we observe such a dramatic improvement in CPU time with different system configurations.

Generality. In our main theoretical results in §4, certain conditions on the underlying distributions, (F1)-(F3) or (F3'), are imposed in order to characterize the performance in terms of the relative efficiency defined in (12). However, it is important to note that the naive, QD-C, and QD-D algorithms are consistent so that $P(\text{FS}_t^\pi) \rightarrow 0$ as $t \rightarrow \infty$ under the mild conditions (F1)-(F2); see Theorems 1-3. We remark that the AQD-C and AQD-D algorithms can be applied in general settings where (F3) or (F3') is not necessarily satisfied.

ALGORITHM 6: HH (n_0, m, β)

For each j , take n_0 samples and let $t = kn_0$.

repeat

Let $b = \arg \max_j \{\hat{\xi}_{jt}^\pi\}$ and $b' = \arg \max_{j \neq b} \{\hat{\xi}_{jt}^\pi\}$.

Fix $v_t = \beta \hat{\xi}_{bt}^\pi + (1 - \beta) \hat{\xi}_{b't}^\pi$.

Define $\{\zeta_{jt}\}_{j=1}^k$ as follows

$$\zeta_{jt} = \begin{cases} p - \hat{F}_{bt}^\pi(v_t) & \text{for } j = b \\ \hat{F}_{jt}^\pi(v_t) - p & \text{for } j \neq b. \end{cases} \quad (44)$$

Set $\pi_{t+\ell} = \arg \min_j \{N_{jt}^\pi \zeta_{jt}^2\}$ for $\ell = 1, \dots, m$ and let $t = t + m$.

until $t \leq T$;

Return: $\arg \max_j \{\hat{\xi}_{jT}^\pi\}$.

6. Comparison with Benchmark Algorithms

6.1. Benchmark Sampling Algorithms

We compare the proposed algorithms with two benchmark algorithms: the equal allocation (EA) and a heuristic algorithm based on Hoeffding's inequality (HH). The HH algorithm is introduced as a simple benchmark that uses confidence intervals of sample quantiles; see Algorithm 6. It takes as input three tuning parameters: the number of initial samples n_0 , the batch size m , and $\beta \in (0, 1)$.

To provide some intuition behind the HH algorithm, let us denote ξ^p as the p th quantile and $\hat{\xi}^p$ as the sample quantile from N independent samples. From Hoeffding's inequality, observe that

$$\mathbf{P}(\hat{\xi}^{p+\zeta} < \xi^p) = \mathbf{P}\left(\sum_{i=1}^N (\mathbf{I}\{X_i \leq \xi^p\} - p) \geq \zeta N\right) \leq e^{-2N\zeta^2}. \quad (43)$$

Therefore, $\hat{\xi}^{p+\zeta}$ is greater than ξ^p with probability greater than or equal to $1 - \exp(-2N\zeta^2)$. In other words, for a given value of ζ , the value of $1 - \exp(-2N\zeta^2)$, or equivalently, $N\zeta^2$, can be considered as a proxy for the confidence level. The HH algorithm is designed to take a sample from the system with the least value of such a measure, $\arg \min_j \{N_{jt}^\pi \zeta_{jt}^2\}$, in stage t . The value of ζ_{jt} for each j depends on the parameter β . The optimal value of β may depend on underlying probability distributions, and hence, it may not be known a priori. In the following numerical experiments, we test the HH algorithm for different values of $\beta = 0.25, 0.5, 0.75$ but only show the results for $\beta = 0.5$ which gives a better overall performance in terms of the probability of false selection.

6.2. Numerical Experiments

We test our procedures for continuous (normal and Student's t) and discrete (uniform and Poisson) distributions. $P(\text{FS}_t^\pi)$ is estimated by counting the number of false selections out of v simulation trials, which is chosen so that:

$$\sqrt{\frac{P_t(1-P_t)}{v}} \leq \frac{P_t}{10}, \quad (45)$$

where P_t is the order of magnitude of $P(\text{FS}_t^\pi)$. This implies standard errors for the estimates of $P(\text{FS}_t^\pi)$ are at least ten times smaller than the value of $P(\text{FS}_t^\pi)$ so that we have sufficiently high confidence that the results are not driven by simulation error.

For the naive algorithm, we set $c = 10$ in all cases. For the QD-C and AQD-C algorithms for continuous distributions, we use the normal kernel with the bandwidth parameter $h(n) = 1.06\sigma n^{-1/5}$, where σ is replaced with the usual sample standard deviation. The tuning parameters chosen in our numerical experiments are by no means optimal but we provide numerical examples in Appendix C of the online supplement to show the sensitivity of performance with respect to these parameters.

In the following experiments, we test the naive and QD algorithms for cases with small number of systems ($k = 4$) since they are not practically implementable with larger problem instances; see CPU times of these algorithms in Table 1. The algorithms are implemented using MATLAB and our codes are available in the MATLAB Central File Exchange.⁶

Normal distributions. We consider configurations where the quantiles are monotonically decreasing; in particular, we set $\mu_j = 0$ and $\sigma_j = 1 + \gamma(j - 1)$ for $j = 1, \dots, k$ for some $\gamma > 0$. We vary $\gamma = 0.1, 1$ which corresponds to $\delta = 0.128, 1.28$, the gap in quantile between the best and second best systems. Also, to show the performance of our procedure for the cases with extreme quantiles and large number of systems, we vary $p = 0.1, 0.01$ and $k = 4, 400$.

In all cases of Figure 5, the QD-C and AQD-C algorithms outperform the others. In light of Theorem 2, the QD-C and AQD-C policies are near-optimal asymptotically when the gap δ in quantile between the best and second-best systems is sufficiently small. This is consistent with the observations from panels (a),(c), and (e) of Figure 5, where δ is sufficiently small and the finite-time performance of the two algorithms is competitive

⁶ <https://www.mathworks.com/matlabcentral/fileexchange/74423-sampling-algorithms-for-quantile-based-ordinal-optimization>

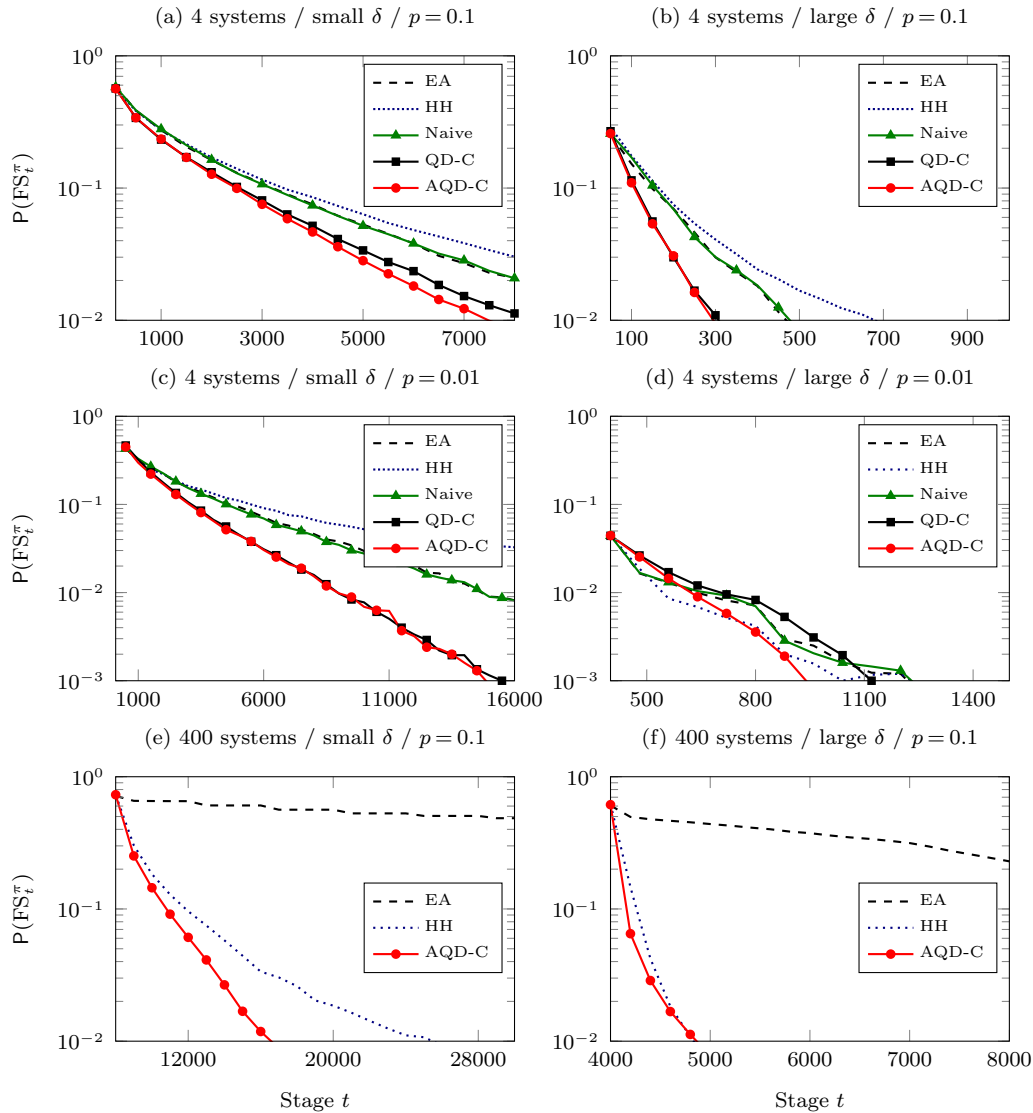


Figure 5 $P(\text{FS}_t^T)$ as a function of stage t on a log-linear scale in the normal cases. In all cases, means of systems are equally set to zero and standard deviations are $\sigma_j = 1 + \gamma(j - 1)$ for $j = 1, \dots, k$, where we vary $\gamma = 0.1, 1$ for the cases with small and large values of δ , respectively. We set $k = 4$ and $p = 0.1$ in panels (a)-(b), $k = 4$ and $p = 0.01$ in panels (c)-(d), and $k = 400$ and $p = 0.1$ in panels (e)-(f). We take $m = 10$ samples per batch for all cases and set $n_0 = 200, 20, 200, 100, 20, 10$ for (a)-(f), respectively. For the HH algorithm, the parameter β is tuned to be 0.5 for both configurations.

in terms of the probability of false selection. Somewhat surprisingly, the performance of QD-C and AQD-C are also preferable in cases with relatively large values of δ ; see panels (b),(d), and (f) in Figure 5. These results indicate that the finite-time performance of the two algorithms is competitive for both small and large values of δ , which complements the asymptotic results of Theorem 2 as $\delta \rightarrow 0$. Further, note that AQD-C slightly outperforms QD-C in some cases in Figure 5.

Notice that the poor performance of the naive algorithm is anticipated by Remark 1. The performance of the HH algorithm varies significantly for different configurations; it performs competitively with the AQD-C algorithm in panel (f) of Figure 5, but it is even worse than EA in panels (a)-(c). This shows that the performance of the HH algorithm highly depends on the choice of parameter β , which is not known a priori. Lastly, note that the EA algorithm performs very poorly in the case with $k = 400$ since too many samples are allocated to systems that are far from the best one, while the AQD-C algorithm safely discards seemingly non-best systems in early stages.

Heavy tailed distributions. In Figure 6, we consider two system configurations, each with four Student t -distributed systems. In both cases, we set means $(0, 0, 0, 0)$ and standard deviations $(1, 2, 3, 4)$ with $p = 0.1$ so that the p th quantile values monotonically decrease with system index.

It is noteworthy that the performance of quantile-based algorithms is robust against the presence of heavy tails in the sense that $\mathbb{P}(\text{FS}_t^\pi)$ converges to zero at an exponential rate (Proposition 2), as opposed to mean-based procedures under which the probability of false selection decreases to zero only at a polynomial rate (Broadie et al. 2007). This can be seen in Figure 6: On the left panel, the degrees of freedom is set to $\nu = 1$ so that the distributions are significantly heavy-tailed, while we set $\nu = 20$ for the less heavy-tailed distributions on the right panel of Figure 6. In both cases, the AQD-C and QD-C algorithms outperform the others, with AQD-C performing slightly better than QD-C on the left panel. Also, note that $\mathbb{P}(\text{FS}_t^\pi) \approx c_1 \exp(-c_2 t)$ implies that $\log \mathbb{P}(\text{FS}_t^\pi)$ is approximately linear in t . Hence, it can be seen that $\mathbb{P}(\text{FS}_t^\pi)$ converges to zero at an exponential rate under QD-C and AQD-C, indicating these algorithms can be attractive alternatives to mean-based procedures in many applications where heavy tails are prevalent.

Discrete cases: uniform and Poisson distributions. Recall from Theorem 3 that the asymptotic performance of the QD-D algorithm is characterized by $\rho(\boldsymbol{\alpha}^{\delta, \epsilon})/\rho^*$, which depends on the gap in means between the best and second-best systems (δ) and the error due to linear approximation of distribution functions (ϵ). To illustrate the effect of ϵ , we consider two families of discrete distributions: uniform distribution functions which can be closely approximated by a linear function (i.e., small ϵ) and Poisson distributions for which ϵ is relatively large. Also, to illustrate the effect of δ , we consider small and large values of δ for each of the two distributions. We fix $p = 0.1$ and $k = 4$ in all of these cases.

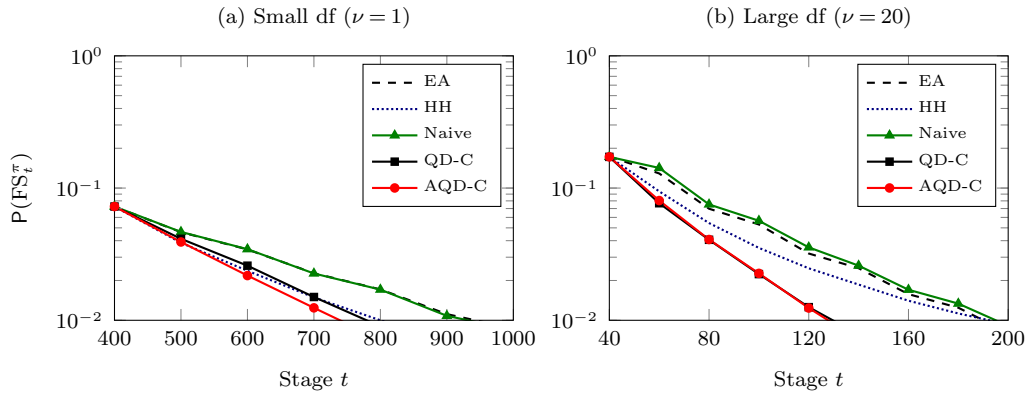


Figure 6 $P(\text{FS}_t^T)$ as a function of stage t on a log-linear scale in the t -distribution cases. For both cases, there are four systems with means $(0, 0, 0, 0)$ and standard deviations $(1, 2, 3, 4)$, while the degrees of freedom are $\nu = 1$ on the left panel and $\nu = 20$ on the right panel. We set $p = 0.1$ and $m = 10$ for both cases and let $n_0 = 100, 10$ for (a) and (b), respectively. For the HH algorithm, the parameter β is tuned to be 0.5.

In the top panels of Figure 7, two configurations with uniform distributions are considered. In each case, the support of the distribution function is $\{1, \dots, u_j\}$ for system j with $u_j = 1000 - \delta(j - 1)$; we let $\delta = 50$ in (a) and $\delta = 200$ in (b). The values of ϵ are small (in the order of 10^{-3}) in both cases, while $\theta = 0.0064, 0.15$ for (a) and (b), respectively. Further, the asymptotic efficiencies characterized in Theorem 3, $\rho(\boldsymbol{\alpha}^{\delta, \epsilon^*})/\rho^*$, are 0.89 and 0.61 for (a) and (b), respectively, because the gap in quantiles δ (or equivalently, θ) is smaller in the first configuration. It is interesting to observe that the performance of the AQD-D algorithm is still superior to the others in terms of the probability of false selection even with the relatively low asymptotic efficiency in Figure 7(b).

In the bottom panels of Figure 7, each system j is characterized by a Poisson distribution with parameter λ_j . In particular, we let $\lambda_j = 1000 - \delta(j - 1)$ with $\delta = 5, 20$ in panels (c) and (d), respectively. As opposed to the uniform case, the values of ϵ are relatively large due to the non-linearity of Poisson distribution functions; in particular, $\epsilon = 0.101, 1.814$ for (c) and (d), respectively. The large values of ϵ deteriorate the asymptotic efficiencies significantly; $\rho(\boldsymbol{\alpha}^{\delta, \epsilon^*})/\rho^* = 0.72, 0.56$ for (c) and (d). Nevertheless, observe that the finite-time performance of AQD-D in terms of the probability of false selection is still superior to the other algorithms, indicating its performance is robust against a wide spectrum of ϵ and δ .

7. Concluding Remarks

We have shown how the problem of minimizing the probability of false selection can be analyzed using large deviations theory and certain approximations thereof. By analyzing

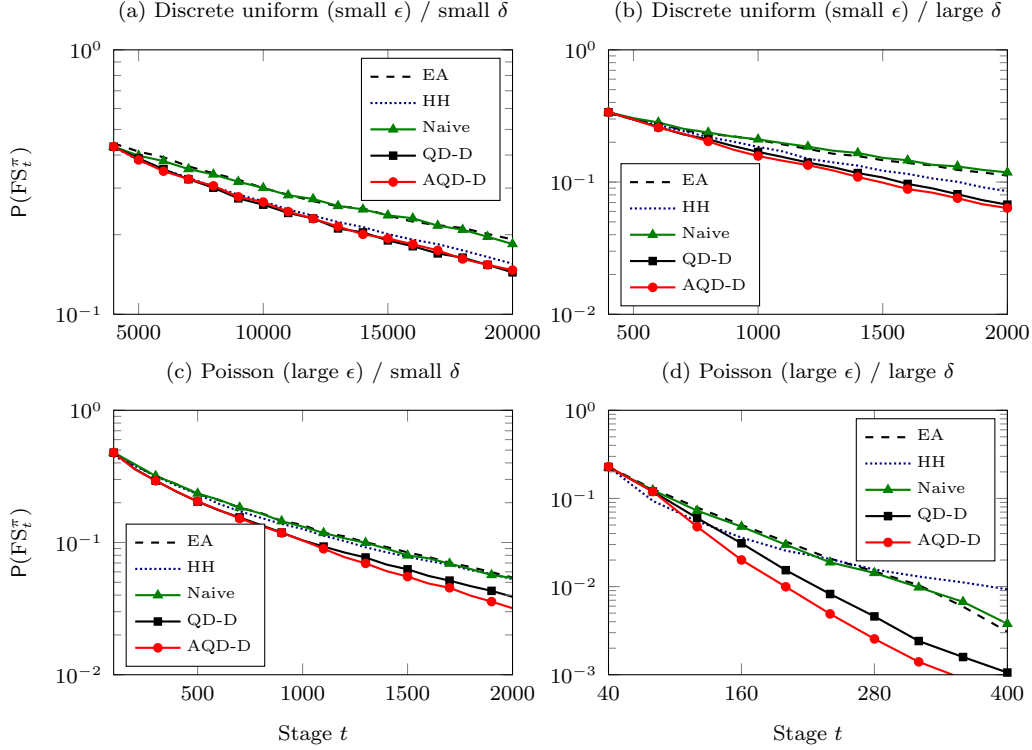


Figure 7 $P(\text{FS}_t^T)$ as a function of stage t on a log-linear scale in the discrete cases. In (a) and (b), the system configuration is characterized by $k = 4$ discrete uniform distributions, each on the range $[0, u_j]$ with $u_j = 1000 - \delta(j - 1)$ for $j = 1, \dots, k$. We set $\delta = 50, 200$ for (a) and (b), respectively. In (c) and (d), the system configuration is characterized by $k = 4$ Poisson distributions, each with parameter $\lambda_j = 1000 - \delta(j - 1)$. We set $\delta = 5, 20$ for (c) and (d), respectively. For all cases $p = 0.1$. We set $(n_0, m) = (1000, 10), (100, 10), (100, 10), (10, 1)$ for (a)-(d), respectively. For the HH algorithm, the parameter β is tuned to be 0.5.

the rate function of the probability of false selection in an asymptotic regime, we obtain a tractable objective function and structural insights that guide algorithm design. Although it is tractable, the rate function is computationally expensive to estimate from sample observations. A significant contribution of our work is that a nearly-optimal algorithm is obtained using a simple alternative to the rate function, which can be estimated using a surprisingly small number of sample observations with high accuracy.

In a wide range of applications, selecting multiple systems is likely to be a topic of interest. The theoretical results in our work can serve as a base for the case with multiple selections. On a methodological level, one can use different nonparametric estimators for quantiles, other than the sample quantile used in this paper, which can be leveraged to design even more efficient algorithms.

References

- Batur D, Choobineh F (2010) A quantile-based approach to system selection. *Eur. J. Oper. Res.* 202:764–772.
- Bechhofer RE (1954) A single-sample multiple decision procedure for ranking means of normal populations with known variances. *Annals of Math. Statist.* 25:16–39.
- Bekki JM, Fowler JW, Mackulak GT, Nelson BL (2007) Using quantiles in ranking and selection procedures. Henderson SG, Biller B, Hsieh MH, Tew JD, Barton RR, eds., *Proc. 2007 Winter Simulation Conf.*, 1722–1728 (Piscataway, NJ: IEEE).
- Broadie M, Han M, Zeevi A (2007) Implications of heavy tails on simulation-based ordinal optimization. Henderson SG, Biller B, Hsieh MH, Tew JD, Barton RR, eds., *Proc. 2007 Winter Simulation Conf.*, 439–447 (Piscataway, NJ: IEEE).
- Bubeck S, Cesa-Bianchi N (2012) Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends Machine Learn.* 5:1–122.
- Chen CH, Lee LH (2011) *Stochastic simulation optimization: an optimal computing budget allocation*, volume 1 (World scientific).
- Chen CH, Lin J, Yücesan E, Chick SE (2000) Simulation budget allocation for further enhancing the efficiency of ordinal optimization. *Discrete Event Dynamic Systems* 10:251–270.
- Dembo A, Zeitouni O (2009) *Large deviations techniques and applications* (New York: Springer Science & Business Media).
- Frazier P, Powell WB, Dayanik S (2008) A knowledge-gradient policy for sequential information collection. *SIAM J. Control Optim.* 47:2410–2439.
- Glynn P, Juneja S (2004) A large deviations perspective on ordinal optimization. Ingalls RG, Rossetti MD, Smith JS, Peters BA, eds., *Proc. 2004 Winter Simulation Conf.*, 577–585 ((IEEE, Piscataway, NJ)).
- Glynn P, Juneja S (2015) Ordinal optimization—empirical large deviations rate estimators, and stochastic multi-armed bandits. *arXiv:1507.04564* .
- Kim SH, Nelson BL (2006) Selecting the best system. Henderson SG, Nelson BL, eds., *Handbooks in Operations Research and Management Science: Simulation*, volume 13, chapter 17, 501–534 (Boston: Elsevier).
- Lee S, Nelson BL (2014) Bootstrap ranking & selection revisited. Tolk A, Diallo SY, Ryzhov IO, Yilmaz L, Buckley S, Miller JA, eds., *Proc. 2014 Winter Simulation Conf.*, 3857–3868 (Piscataway, NJ: IEEE).
- Pasupathy R, Hunter SR, Pujowidianto NA, Lee LH, Chen CH (2015) Stochastically constrained ranking and selection via score. *ACM Transactions on Modeling and Computer Simulation (TOMACS)* 25(1):1.
- Pasupathy R, Szechtman R, Yücesan E (2010) Selecting small quantiles. Johansson B, Jain S, Montoya-Torres J, Hagan J, Yücesan E, eds., *Proc. 2010 Winter Simulation Conf.*, 2762–2770 (Piscataway, NJ: IEEE).

- Peng Y, Chen C, Fu M, Hu J, Ryzhov I (2019) Efficient sampling allocation procedures for selecting the optimal quantile, *Forthcoming in INFORMS J. Comput.*
- Peng Y, Chong EKP, Chen CH, Fu MC (2018a) Ranking and selection as stochastic control. *IEEE Transactions on Automatic Control* 63(8):2359–2373.
- Peng Y, Xu J, Lee L, Hu J, Chen CH (2018b) Efficient simulation sampling allocation using multifidelity models. *IEEE Transactions on Automatic Control* 64(8):3156–3169.
- Rajagopalan B, Lall U (1995) A kernel estimator for discrete distributions. *Nonparametric Statist.* 4(4):409–426.
- Reiss R (1981) Nonparametric estimation of smooth distribution functions. *Scandinavian J. of Statist.* 116–119.
- Rinott Y (1978) On two-stage selection procedures and related probability-inequalities. *Commun. Statist.* 7:799–811.
- Serfling RJ (2009) *Approximation theorems of mathematical statistics* (New York: John Wiley & Sons).
- Shin D, Broadie M, Zeevi A (2016) Tractable sampling strategies for quantile-based ordinal optimization. *Proc. 2016 Winter Simulation Conf.*, 847–858 (IEEE Press).
- Shin D, Broadie M, Zeevi A (2018) Tractable sampling strategies for ordinal optimization. *Oper. Res.* 66(6):1693–1712.
- Silverman BW (1986) *Density estimation for statistics and data analysis* (New York: CRC press).
- Szorenyi B, Busa-Fekete R, Weng P, Hüllermeier E (2015) Qualitative multi-armed bandits: A quantile-based approach. *32nd Internat. Conf. on Machine Learn.*, 1660–1668.
- Yu JY, Nikolova E (2013) Sample complexity of risk-averse bandit-arm selection. *Proc. 23rd Internat. Joint Conf. on Artificial Intelligence*, 2576–2582 (AAAI Press).

Submitted to *INFORMS Journal on Computing*
manuscript (Please, provide the manuscript number!)

Authors are encouraged to submit new papers to INFORMS journals by means of a style file template, which includes the journal title. However, use of a template does not certify that the paper has been accepted for publication in the named journal. INFORMS journal templates are for the exclusive purpose of submitting to an INFORMS journal and should not be used to distribute the papers in print or online or to submit the papers to another publication.

Online Supplement for Practical Nonparametric Sampling Strategies for Quantile-based Ordinal Optimization

Dongwook Shin

School of Business and Management, Hong Kong University of Science and Technology,
Clear Water Bay, Kowloon, Hong Kong
{dwshin@ust.hk}

Mark Broadie, Assaf Zeevi

Graduate School of Business, Columbia University, New York, NY 10025
{mnb2@columbia.edu, assaf@gsb.columbia.edu}

Key words: Quantile, ordinal optimization, tractable procedures, large deviations theory

The online supplement comprises three parts: Appendix A states additional theoretical results; Appendix B contains the proofs of main theoretical results and auxiliary lemmas; and Appendix C contains additional numerical results to provide practical guideline on the implementation of proposed algorithms.

Appendix A: Additional Theoretical Results

A.1. Properties of the Function $\rho^\delta(\cdot)$

Recall the definition of $\rho^\delta(\boldsymbol{\alpha}) = \min_{j \neq 1} \{G_j^\delta(\boldsymbol{\alpha})\}$, where

$$G_j^\delta(\boldsymbol{\alpha}) = \frac{(\xi_1 - \xi_j)^2}{2p(1-p)(1/(\alpha_1 f_1^2(\xi_1)) + 1/(\alpha_j f_j^2(\xi_j)))}.$$

It is trivial to check that $G_j^\delta(\boldsymbol{\alpha})$ is a smooth function of $\boldsymbol{\alpha} \in \Delta^0$. However, $\rho^\delta(\boldsymbol{\alpha})$ is only piecewise smooth because its derivative does not exist at $\boldsymbol{\alpha}$ such that $G_i^\delta(\boldsymbol{\alpha}) = G_j^\delta(\boldsymbol{\alpha})$ for some $i \neq j$.

LEMMA A.1. *Assume (F1)-(F3). Then, $G_j^\delta(\boldsymbol{\alpha})$ is strongly concave in $\boldsymbol{\alpha}$.*

From the previous lemma, it is straightforward to see that $\rho^\delta(\boldsymbol{\alpha})$, being the minimum of $G_j^\delta(\boldsymbol{\alpha})$ over $j \neq 1$, is also strongly concave.

A.2. Strong Consistency of the Proposed Algorithms

In the main text we have shown the (weak) consistency of the naive, QD-C, and QD-D algorithms, respectively in Theorems 1, 2, and 3, respectively; that is, $N_{jt}^\pi \rightarrow \infty$ in probability as $t \rightarrow \infty$. In this section, we are interested in strengthening of these results to convergence with probability one. To this end, we need a stronger notion of consistency:

DEFINITION A.1 (STRONG CONSISTENCY). An algorithm $\pi \in \Pi$ is strongly consistent if $N_{jt}^\pi \rightarrow \infty$ almost surely for each j as $t \rightarrow \infty$.

Note that the naive and QD-D algorithms can be shown to be strongly consistent without additional conditions. However, for the strong consistency of the QD-C algorithm, which is structured around the kernel-based density estimators, we need the following additional conditions:

- (F4) $F_j(\cdot)$ possesses a positive, uniformly continuous density $f_j(\cdot)$ over the interval \mathcal{H}_j .
- (K4) K is of bounded variation.
- (K5) $K(x) \rightarrow 0$ as $|x| \rightarrow \infty$.
- (K6) $\sum_{n=1}^{\infty} \exp(-\gamma th(t)^2) < \infty$ for every $\gamma > 0$.

Note that (F4) is different than (F3) because it does not require twice differentiable $F_j(\cdot)$, but instead requires uniform continuity of the density function. Note also that (K5) does not follow from (K1); a sufficient condition for (K5) is that the kernel K is uniformly continuous. The last condition (K6) is stronger than (K3); for instance, the sequence of $h(t) = \log(t)/t$ satisfies (K3) but not (K6). The stronger consistency results are summarized in the following proposition.

PROPOSITION A.1 (Strong consistency of the proposed algorithms). *The naive algorithm is strongly consistent. For $\mathbf{F} \in \mathcal{D}$, QD-D algorithm is strongly consistent under (F1)-(F2). For $\mathbf{F} \in \mathcal{C}$, QD-C algorithm is strongly consistent if underlying distributions satisfy (F1),(F2), and (F4) and the kernel and bandwidth parameters satisfy (K1)-(K6).*

A.3. Strengthening of Proposition 3

An important implication of Proposition 3 is that, after taking $T \rightarrow \infty$, the probability of false selection (on a logarithmic scale) depends on the underlying distributions only through the densities at quantiles, $\{f_j(\xi_j)\}_{j=1}^k$, as the quantiles get closer to each other. To formalize this property more precisely, consider a stylized sequence of system distributions indexed by t , denoted as $\{F_{1,t}(\cdot), \dots, F_{k,t}(\cdot)\}_{t=1}^{\infty}$. Let $\xi_{j,t}$ denote the p -th quantile of the distribution function $F_{j,t}(\cdot)$ with $\xi_{1,t} \geq \dots \geq \xi_{k,t}$ and define $\delta_t = \xi_{1,t} - \xi_{2,t}$. For $t = 1$, we fix $F_{j,1}(\cdot) = F_j(\cdot)$ so that $\xi_{j,1} = \xi_j$ for each j and $\delta_1 = \delta$. For $t \geq 2$, we fix $F_{1,t}(\cdot) = F_{1,1}(\cdot)$ for system 1 and $F_{j,t}(\cdot)$ for system $j \neq 1$ is shifted from $F_j(\cdot)$ so that $F_{j,t}(x - \delta_1 + \delta_t) = F_j(x)$. Essentially, this makes each $\xi_{j,t}$ approaches $\xi_{1,t}$ as $t \rightarrow \infty$, while maintaining the ranking, $\xi_{1,t} > \xi_{2,t} \geq \dots \geq \xi_{k,t}$, for each t . For brevity, we let the system configuration, $(F_{1,t}(\cdot), \dots, F_{k,t}(\cdot))$, be characterized by δ_t and define $\mathbb{P}(\text{FS}_t; \delta_t)$ as the probability of false selection under the configuration, $(F_{1,t}(\cdot), \dots, F_{k,t}(\cdot))$.

PROPOSITION A.2 (Strengthening of Proposition 3). *Under assumptions (F1)-(F3), if $t\delta_t^2 \rightarrow \infty$ and $\delta_t \rightarrow 0$ as $t \rightarrow \infty$, then for any static algorithm $\pi(\boldsymbol{\alpha})$ that satisfies $N_{jt}/t \rightarrow \alpha_j$ as $t \rightarrow \infty$ for some $\boldsymbol{\alpha} \in \Delta^0$,*

$$\frac{1}{t\delta_t^2} \log \mathbb{P}(\text{FS}_t^{\pi(\boldsymbol{\alpha})}; \delta_t) \rightarrow -\rho^\delta(\boldsymbol{\alpha}) \text{ as } t \rightarrow \infty. \quad (\text{A.1})$$

The proof for the preceding proposition is provided in §B.1. Proposition A.2 suggests a relationship between the sampling budget T and the difference in quantiles δ . In particular, the asymptotic behavior of $\mathbb{P}(\text{FS}_T^\pi)$ is related to $T\delta^2$ for sufficiently small δ and large T ; if the difference in quantiles halves, one may need approximately four times more sampling budget to lower $\mathbb{P}(\text{FS}_T^\pi)$ to a certain level.

Appendix B: Proofs

This section is organized as follows. In §B.1 we present proofs for main results and those for auxiliary lemmas are collected in §B.2. We use the following notation for the purpose of asymptotic analysis: for real-valued functions $\ell_1(x)$ and $\ell_2(x)$, we write $\ell_1(x) = o(\ell_2(x))$ if $|\ell_1(x)|/|\ell_2(x)| \rightarrow 0$ as $\ell_2(x) \rightarrow 0$.

B.1. Proofs for Main Results

Proof of Proposition 1. Let \mathcal{F}_t be the σ -field generated by the samples and sampling decisions taken up to stage t (i.e., $\{(\pi_\tau, X_{\pi_\tau, \tau})\}_{\tau=1}^t$), with the convention that \mathcal{F}_0 is the nominal σ -algebra associated with underlying probability space. Fix j and define a random variable P_n so that $N_{jP_n} = n$. (Note that P_n depends on the index j .) We show that $\{X_{jP_n}\}_{n=1}^\infty$ is a sequence of independent and identically distributed random variables.

First, observe that the characteristic function of X_{jP_n} is

$$\begin{aligned} \mathbb{E}(e^{i\theta X_{jP_n}}) &= \sum_{t=1}^{\infty} \mathbb{E}(e^{i\theta X_{jt}} | P_n = t) \mathbb{P}(P_n = t) \\ &= \sum_{t=1}^{\infty} \mathbb{E}(e^{i\theta X_{jt}}) \mathbb{P}(P_n = t) \\ &= \mathbb{E}(e^{i\theta X_j}), \end{aligned} \tag{B.1}$$

where the second equality follows from the fact that the event $\{P_n = t\}$ is adapted to the filtration \mathcal{F}_{t-1} , and the last follows from the fact that the sequence $\{X_{jt}\}_{t=1}^\infty$ is identically distributed.

Next, we show that X_{jP_m} and X_{jP_n} are independent for $n > m$. Observe that, for any measurable sets A_1 and A_2 ,

$$\begin{aligned} &\mathbb{P}(X_{jP_m} \in A_1, X_{jP_n} \in A_2) \\ &= \sum_{t_1 < t_2} \mathbb{P}(X_{jt_1} \in A_1, X_{jt_2} \in A_2) \mathbb{P}(P_m = t_1, P_n = t_2) \\ &= \sum_{t_1 < t_2} \mathbb{E}(\mathbf{I}\{X_{jt_1} \in A_1\} \mathbf{I}\{X_{jt_2} \in A_2\}) \mathbb{P}(P_m = t_1, P_n = t_2) \\ &= \sum_{t_1 < t_2} \mathbb{E}(\mathbb{E}(\mathbf{I}\{X_{jt_1} \in A_1\} \mathbf{I}\{X_{jt_2} \in A_2\}) | \mathcal{F}_{t_2-1}) \mathbb{P}(P_m = t_1, P_n = t_2) \\ &\stackrel{(a)}{=} \sum_{t_1 < t_2} \mathbb{E}(\mathbf{I}\{X_{jt_1} \in A_1\}) \mathbb{E}(\mathbf{I}\{X_{jt_2} \in A_2\}) \mathbb{P}(P_m = t_1, P_n = t_2) \\ &= \sum_{t_1 < t_2} \mathbb{P}(X_{jt_1} \in A_1) \mathbb{P}(X_{jt_2} \in A_2) \mathbb{P}(P_m = t_1, P_n = t_2) \\ &\stackrel{(b)}{=} \mathbb{P}(X_j \in A_1) \mathbb{P}(X_j \in A_2) \sum_{t_1 < t_2} \mathbb{P}(P_m = t_1, P_n = t_2) \\ &= \mathbb{P}(X_j \in A_1) \mathbb{P}(X_j \in A_2) \\ &\stackrel{(c)}{=} \mathbb{P}(X_{jP_m} \in A_1) \mathbb{P}(X_{jP_n} \in A_2), \end{aligned} \tag{B.2}$$

where (a) follows from the fact that $X_{j t_1} \in \mathcal{F}_{t_2-1}$ for $t_1 < t_2$ and (b) and (c) follow from (B.1). Therefore, $\{X_{j P_n}\}_{n=1}^\infty$ is independent and identically distributed. The sample p th quantile from an independent and identically distributed sequence converges to the p th quantile of the distribution almost surely (see, e.g., p. 75 of Serfling 2009), and hence, also in probability. This completes the proof. \square

The proof for Proposition 2 requires the following lemma. Proofs for all auxiliary lemmas are provided in §B.2.

LEMMA B.1 (Large deviations for sample quantiles). *Under (F1)-(F2), Equation (8) holds for a static algorithm $\pi(\boldsymbol{\alpha})$ for any $\boldsymbol{\alpha} \in \Delta^0$.*

Proof of Proposition 2. In this proof, we fix a static algorithm $\pi = \pi(\boldsymbol{\alpha})$ for some $\boldsymbol{\alpha} \in \Delta^0$ and suppress π in the superscripts to improve clarity. Observe that

$$\max_{j=2, \dots, k} \mathbb{P}(\hat{\xi}_{1t} \leq \hat{\xi}_{jt}) \leq \mathbb{P}(\text{FS}_t) \leq (k-1) \max_{j=2, \dots, k} \mathbb{P}(\hat{\xi}_{1t} \leq \hat{\xi}_{jt}). \quad (\text{B.3})$$

Hence, if, for each $j = 2, \dots, k$,

$$\lim_{t \rightarrow \infty} \frac{1}{t} \log \mathbb{P}(\hat{\xi}_{1t} \leq \hat{\xi}_{jt}) = -G_j(\boldsymbol{\alpha}), \quad (\text{B.4})$$

then

$$\lim_{t \rightarrow \infty} \frac{1}{t} \log \mathbb{P}(\text{FS}_t) = - \min_{j=2, \dots, k} \{G_j(\boldsymbol{\alpha})\}. \quad (\text{B.5})$$

Therefore, the rate function of $\mathbb{P}(\text{FS}_t)$ can be immediately obtained once we prove (B.4) for some $j \neq 1$.

First, observe that

$$\mathbb{P}(\hat{\xi}_{1t} \leq \hat{\xi}_{2t}) \geq \mathbb{P}(\hat{\xi}_{1t} \leq x) \mathbb{P}(\hat{\xi}_{2t} \geq x) \quad (\text{B.6})$$

for any $x \in \mathcal{R}$. Also, taking log on both sides and combining the result from Lemma B.1, it can be easily seen that

$$\liminf_{t \rightarrow \infty} \frac{1}{t} \log \mathbb{P}(\hat{\xi}_{1t} \leq \hat{\xi}_{2t}) \geq - \inf_{x \in \mathcal{R}_1 \cap \mathcal{R}_2} \{\alpha_1 I_1(x) + \alpha_2 I_2(x)\}, \quad (\text{B.7})$$

where the inf follows from the fact that (B.6) holds for any x . Also, in the continuous case, observe that $I_j(x)$ is non-increasing for $x < \xi_j$ and non-decreasing for $x > \xi_j$, and therefore, it suffices to search for the infimum for $\xi_2 \leq x \leq \xi_1$. Likewise, in the discrete case, observe that $I_j(x)$ is non-increasing for $x < \xi_j - 1$ and non-decreasing for $x > \xi_j$, so that the infimum is achieved in $[\xi_2 - 1, \xi_1]$. Therefore, (B.6) reduces to

$$\liminf_{t \rightarrow \infty} \frac{1}{t} \log \mathbb{P}(\hat{\xi}_{1t} \leq \hat{\xi}_{2t}) \geq \begin{cases} - \inf_{x \in [\xi_2, \xi_1]} \{\alpha_1 I_1(x) + \alpha_2 I_2(x)\} & \text{for } \mathbf{F} \in \mathcal{C} \\ - \inf_{x \in [\xi_2 - 1, \xi_1]} \{\alpha_1 I_1(x) + \alpha_2 I_2(x)\} & \text{for } \mathbf{F} \in \mathcal{D}. \end{cases} \quad (\text{B.8})$$

It remains to show the upper bound:

$$\limsup_{t \rightarrow \infty} \frac{1}{t} \log \mathbb{P}(\hat{\xi}_{1t} \leq \hat{\xi}_{2t}) \leq \begin{cases} - \inf_{x \in [\xi_2, \xi_1]} \{\alpha_1 I_1(x) + \alpha_2 I_2(x)\} & \text{for } \mathbf{F} \in \mathcal{C} \\ - \inf_{x \in [\xi_2 - 1, \xi_1]} \{\alpha_1 I_1(x) + \alpha_2 I_2(x)\} & \text{for } \mathbf{F} \in \mathcal{D}. \end{cases} \quad (\text{B.9})$$

(i) Continuous case. Let $\mathbf{x} = (x_1, x_2) \in \mathcal{R}^2$ with $x_1 \leq x_2$, fix $\eta > 0$, and consider a square centered at \mathbf{x} , $S_{\mathbf{x}}^\eta = \{\mathbf{x}' \in \mathcal{R}^2 \mid |x'_1 - x_1| \leq r_\eta, |x'_2 - x_2| \leq r_\eta\}$, with $r_\eta > 0$ chosen small enough so that

$$\max_{j=1,2} \{|I_j(x'_j) - I_j(x_j)|\} \leq \eta \quad \text{for any } \mathbf{x}' \in S_{\mathbf{x}}^\eta. \quad (\text{B.10})$$

Note that r_η always exists for each $\eta > 0$ because $I_j(\cdot)$ is a continuous function. Observe that

$$\begin{aligned} \mathbb{P}\left((\hat{\xi}_{1t}, \hat{\xi}_{2t}) \in S_{\mathbf{x}}^\eta\right) &= \mathbb{P}\left(\hat{\xi}_{1t} \in [x_1 - r_\eta, x_1 + r_\eta]\right) \mathbb{P}\left(\hat{\xi}_{2t} \in [x_2 - r_\eta, x_2 + r_\eta]\right) \\ &\leq \mathbb{P}\left(\hat{\xi}_{1t} \leq x_1 + r_\eta\right) \mathbb{P}\left(\hat{\xi}_{2t} \geq x_2 - r_\eta\right). \end{aligned} \quad (\text{B.11})$$

Applying Lemma B.1, it follows that

$$\begin{aligned} \limsup_{t \rightarrow \infty} \frac{1}{t} \log \mathbb{P}\left((\hat{\xi}_{1t}, \hat{\xi}_{2t}) \in S_{\mathbf{x}}^\eta\right) &\leq -(\alpha_1 I_1(x_1 + r_\eta) + \alpha_2 I_2(x_2 - r_\eta)) \\ &\leq \eta - (\alpha_1 I_1(x_1) + \alpha_2 I_2(x_2)), \end{aligned} \quad (\text{B.12})$$

where the second inequality follows from (B.10).

Now, fix $a > 0$ and consider a compact set $\Gamma_a = \{\mathbf{x} \in \mathcal{R}^2 \mid |x_1 - \xi_1| \leq a, |x_2 - \xi_2| \leq a\}$. Then the set, $\Gamma_a \cap \{\mathbf{x} \in \mathcal{R}^2 \mid x_1 \leq x_2\}$, can be covered by $M < \infty$ squares, each of which is centered at $\mathbf{x}^i = (x_1^i, x_2^i) \in \Gamma_a$ with $x_1^i \leq x_2^i$. Hence, it follows that

$$\begin{aligned} \mathbb{P}\left(\hat{\xi}_{1t} \leq \hat{\xi}_{2t}, (\hat{\xi}_{1t}, \hat{\xi}_{2t}) \in \Gamma_a\right) &\leq \sum_{i=1}^M \mathbb{P}\left((\hat{\xi}_{1t}, \hat{\xi}_{2t}) \in S_{\mathbf{x}^i}^\eta\right) \\ &\leq M \max_{i=1, \dots, M} \left\{ \mathbb{P}\left((\hat{\xi}_{1t}, \hat{\xi}_{2t}) \in S_{\mathbf{x}^i}^\eta\right) \right\}. \end{aligned} \quad (\text{B.13})$$

Therefore,

$$\begin{aligned} \limsup_{t \rightarrow \infty} \frac{1}{t} \log \mathbb{P}\left(\hat{\xi}_{1t} \leq \hat{\xi}_{2t}, (\hat{\xi}_{1t}, \hat{\xi}_{2t}) \in \Gamma_a\right) &\leq \max_{i=1, \dots, M} \left\{ \limsup_{t \rightarrow \infty} \frac{1}{t} \log \mathbb{P}\left(\hat{\xi}_{1t} \leq \hat{\xi}_{2t}, (\hat{\xi}_{1t}, \hat{\xi}_{2t}) \in S_{\mathbf{x}^i}^\eta\right) \right\} \\ &\leq \max_{i=1, \dots, M} \left\{ \eta - (\alpha_1 I_1(x_1^i) + \alpha_2 I_2(x_2^i)) \right\} \\ &\leq \eta - \inf_{x_1 \leq x_2, \mathbf{x} \in \Gamma_a} \left\{ \alpha_1 I_1(x_1) + \alpha_2 I_2(x_2) \right\}, \end{aligned} \quad (\text{B.14})$$

where the second inequality follows from (B.12). Since $I_j(x)$ is non-increasing for $x < \xi_j$ and non-decreasing for $x > \xi_j$, it suffices to search for the infimum with (x_1, x_2) that satisfies $\xi_2 \leq x_1 = x_2 \leq \xi_1$. By the arbitrariness of $\eta > 0$, we have that

$$\limsup_{t \rightarrow \infty} \frac{1}{t} \log \mathbb{P}\left(\hat{\xi}_{1t} \leq \hat{\xi}_{2t}, (\hat{\xi}_{1t}, \hat{\xi}_{2t}) \in \Gamma_a\right) \leq - \inf_{x \in [\xi_2, \xi_1]} \left\{ \alpha_1 I_1(x) + \alpha_2 I_2(x) \right\}. \quad (\text{B.15})$$

Finally, observe that

$$\mathbb{P}\left(\hat{\xi}_{1t} < \hat{\xi}_{2t}\right) \leq \mathbb{P}\left(\hat{\xi}_{1t} < \hat{\xi}_{2t} \text{ and } (\hat{\xi}_{1t}, \hat{\xi}_{2t}) \in \Gamma_a\right) + \mathbb{P}\left((\hat{\xi}_{1t}, \hat{\xi}_{2t}) \in \Gamma_a^c\right), \quad (\text{B.16})$$

and hence,

$$\begin{aligned} \limsup_{t \rightarrow \infty} \frac{1}{t} \log \mathbb{P}\left(\hat{\xi}_{1t} < \hat{\xi}_{2t}\right) &\leq \max \left\{ \limsup_{t \rightarrow \infty} \frac{1}{t} \log \mathbb{P}\left(\hat{\xi}_{1t} < \hat{\xi}_{2t} \text{ and } (\hat{\xi}_{1t}, \hat{\xi}_{2t}) \in \Gamma_a\right), \limsup_{t \rightarrow \infty} \frac{1}{t} \log \mathbb{P}\left((\hat{\xi}_{1t}, \hat{\xi}_{2t}) \in \Gamma_a^c\right) \right\} \\ &\leq \max \left\{ - \inf_{x \in [\xi_2, \xi_1]} \left\{ \alpha_1 I_1(x) + \alpha_2 I_2(x) \right\}, \limsup_{t \rightarrow \infty} \frac{1}{t} \log \mathbb{P}\left((\hat{\xi}_{1t}, \hat{\xi}_{2t}) \in \Gamma_a^c\right) \right\}, \end{aligned} \quad (\text{B.17})$$

where the last inequality follows from (B.15). One can show that the first term on the right-hand side of (B.17) is bounded below since $I_1(x)$ and $I_2(x)$ are bounded for x in the compact set $[\xi_2, \xi_1]$. Also, since

$$\begin{aligned} \mathbb{P}\left((\hat{\xi}_{1t}, \hat{\xi}_{2t}) \in \Gamma_a^c\right) &\leq 4 \max \left\{ \mathbb{P}\left(\hat{\xi}_{1t} < \xi_1 - a\right), \mathbb{P}\left(\hat{\xi}_{1t} > \xi_1 + a\right), \mathbb{P}\left(\hat{\xi}_{2t} < \xi_2 - a\right), \mathbb{P}\left(\hat{\xi}_{2t} < \xi_2 + a\right) \right\}, \end{aligned} \quad (\text{B.18})$$

combined with Lemma B.1, it can be seen that

$$\begin{aligned} & \lim_{a \rightarrow \infty} \limsup_{t \rightarrow \infty} \frac{1}{t} \log \mathbb{P} \left((\hat{\xi}_{1t}, \hat{\xi}_{2t}) \in \Gamma_a^c \right) \\ & \leq - \lim_{a \rightarrow \infty} \min \{ \alpha_1 I_1(\xi_1 + a), \alpha_1 I_1(\xi_1 - a), \alpha_2 I_2(\xi_2 + a), \alpha_2 I_2(\xi_2 - a) \} \\ & = -\infty, \end{aligned} \quad (\text{B.19})$$

where the equality follows from the fact that $I_j(x) \rightarrow \infty$ as $|x| \rightarrow \infty$ for each $j = 1, 2$. This implies that one can choose sufficiently large a so that the second term on the right-hand side of (B.17) is smaller than the first term. This completes the proof for the upper bound (B.9) in the continuous case.

(ii) Discrete case. Let $\mathbf{n} = (n_1, n_2) \in \mathcal{N}^2$ with $n_1 \leq n_2$, where \mathcal{N} is the set of non-negative integers, and observe that

$$\mathbb{P} \left((\hat{\xi}_{1t}, \hat{\xi}_{2t}) = \mathbf{n} \right) = \mathbb{P} \left(\hat{\xi}_{1t} = n_1 \right) \mathbb{P} \left(\hat{\xi}_{2t} = n_2 \right). \quad (\text{B.20})$$

Applying Lemma B.1, it follows that

$$\limsup_{t \rightarrow \infty} \log \mathbb{P} \left((\hat{\xi}_{1t}, \hat{\xi}_{2t}) = \mathbf{n} \right) = -(\alpha_1 I_1(n_1) + \alpha_2 I_2(n_2)). \quad (\text{B.21})$$

Now, fix $a > 0$ and consider a compact set $\Gamma_a = \{ \mathbf{n} \in \mathcal{N}^2 \mid |n_1 - \xi_1| \leq a, |n_2 - \xi_2| \leq a \}$. Denote $\{\mathbf{n}^i\}_{i=1}^M$ be the $M < \infty$ points such that $\mathbf{n}^i \in \Gamma_a$ and $n_1^i \leq n_2^i$. Observe that

$$\begin{aligned} \mathbb{P} \left(\hat{\xi}_{1t} \leq \hat{\xi}_{2t}, (\hat{\xi}_{1t}, \hat{\xi}_{2t}) \in \Gamma_a \right) & \leq \sum_{i=1}^M \mathbb{P} \left((\hat{\xi}_{1t}, \hat{\xi}_{2t}) = \mathbf{n}^i \right) \\ & \leq M \max_{i=1, \dots, M} \{ \mathbb{P} \left((\hat{\xi}_{1t}, \hat{\xi}_{2t}) = \mathbf{n}^i \right) \}, \end{aligned} \quad (\text{B.22})$$

from which we establish that

$$\begin{aligned} & \limsup_{t \rightarrow \infty} \frac{1}{t} \log \mathbb{P} \left(\hat{\xi}_{1t} \leq \hat{\xi}_{2t}, (\hat{\xi}_{1t}, \hat{\xi}_{2t}) \in \Gamma_a \right) \\ & \leq \max_{i=1, \dots, M} \left\{ \limsup_{t \rightarrow \infty} \frac{1}{t} \log \mathbb{P} \left(\hat{\xi}_{1t} \leq \hat{\xi}_{2t}, (\hat{\xi}_{1t}, \hat{\xi}_{2t}) = \mathbf{n}^i \right) \right\} \\ & \leq \max_{i=1, \dots, M} \{ -(\alpha_1 I_1(n_1^i) + \alpha_2 I_2(n_2^i)) \} \\ & \leq - \inf_{n_1 \leq n_2, \mathbf{n} \in \Gamma_a} \{ \alpha_1 I_1(n_1) + \alpha_2 I_2(n_2) \}, \end{aligned} \quad (\text{B.23})$$

where the second inequality follows from (B.22). It is not hard to verify that $I_j(n)$ is non-increasing for $n < \xi_j - 1$ and non-decreasing for $n > \xi_j$. Hence, it suffices to search for (n_1, n_2) such that $\xi_2 - 1 \leq n_1 = n_2 \leq \xi_1$. Therefore, we have that

$$\limsup_{t \rightarrow \infty} \frac{1}{t} \log \mathbb{P} \left(\hat{\xi}_{1t} \leq \hat{\xi}_{2t}, (\hat{\xi}_{1t}, \hat{\xi}_{2t}) \in \Gamma_a \right) \leq - \inf_{n \in [\xi_2 - 1, \xi_2]} \{ \alpha_1 I_1(n) + \alpha_2 I_2(n) \}. \quad (\text{B.24})$$

Finally, by the same argument as in the continuous case, the event $\{(\hat{\xi}_{1t}, \hat{\xi}_{2t}) \in \Gamma_a\}$ is negligible with a chosen sufficiently large. This gives us the desired result for the discrete case, which concludes the proof of the theorem. \square

Proof of Theorem 1. First, it is straightforward to see that the naive algorithm is consistent since $N_{jt} \geq c \log t$ for sufficiently large t . Hence, it follows that $\hat{\rho}_t(\boldsymbol{\alpha}) \rightarrow \rho(\boldsymbol{\alpha})$ as $t \rightarrow \infty$ for each $\boldsymbol{\alpha}$ because $\hat{I}_{jt}(x)$ converges to $I_j(x)$ pointwise for each j and each quantile estimator $\hat{\xi}_{jt}$ converges to ξ_j by Proposition 1. Let $\boldsymbol{\alpha}^* \in \arg \max_{\boldsymbol{\alpha} \in \Delta} \{ \rho(\boldsymbol{\alpha}) \}$ and observe that

$$|\hat{\rho}_t(\hat{\boldsymbol{\alpha}}_t) - \rho(\boldsymbol{\alpha}^*)| \leq \max(|\hat{\rho}_t(\hat{\boldsymbol{\alpha}}_t) - \rho(\hat{\boldsymbol{\alpha}}_t)|, |\hat{\rho}_t(\boldsymbol{\alpha}^*) - \rho(\boldsymbol{\alpha}^*)|), \quad (\text{B.25})$$

and therefore, $\hat{\rho}_t(\hat{\alpha}_t) \rightarrow \rho(\alpha^*)$ as $t \rightarrow \infty$. Combined with the fact that $\alpha_t - \hat{\alpha}_t \rightarrow 0$ as $t \rightarrow \infty$ almost surely by construction of Algorithm 1, it can be easily seen that $\mathcal{R}_t^\pi = \rho(\alpha_t)/\rho(\alpha^*) \rightarrow 1$ as $t \rightarrow \infty$ almost surely. Since \mathcal{R}_t^π is a bounded random variable, we establish that

$$\mathbb{E}(\mathcal{R}_t^\pi) = \mathbb{E}\left(\frac{\rho(\alpha_t)}{\rho(\alpha^*)}\right) \rightarrow 1 \text{ as } t \rightarrow \infty. \quad \square \quad (\text{B.26})$$

Proof of Proposition 3. As a slight abuse of notation, we consider a sequence of system configurations where each configuration is uniquely parametrized by $\delta > 0$. We assume without loss of generality that $\xi_1 > \xi_2 = \max_{j>2}\{\xi_j\}$ for every system configuration. We prove the proposition in four steps.

Step 1. We first show that $|G_j(\alpha) - G_j^\delta(\alpha)| = o(\delta^2)$ if $\xi_1 - \xi_j \rightarrow 0$ as $\delta \rightarrow 0$. Fix $j = 2, \dots, k$ such that $\xi_1 - \xi_j \rightarrow 0$ as $\delta \rightarrow 0$ and observe that

$$G_j(\alpha) = \inf_{x \in [\xi_j, \xi_1]} \{\alpha_1 I_1(x) + \alpha_j I_j(x)\} \quad (\text{B.27})$$

and observe that

$$\begin{aligned} I_j'(x) &= f_j(x) \left(\frac{p}{F_j(x)} - \frac{1-p}{1-F_j(x)} \right) \\ I_j''(x) &= F_j''(x) \left(\frac{p}{F_j(x)} - \frac{1-p}{1-F_j(x)} \right) + f_j^2(x) \left(-\frac{p}{F_j^2(x)} + \frac{1-p}{(1-F_j(x))^2} \right), \end{aligned} \quad (\text{B.28})$$

where $I_j'(x)$ and $I_j''(x)$ are the first and the second derivatives of $I_j(\cdot)$ at x , respectively. Hence, it can be easily seen that $I_j(\xi_j) = 0$, $I_j'(\xi_j) = 0$, and

$$I_j''(\xi_j) = \frac{f_j^2(\xi_j)}{p(1-p)}. \quad (\text{B.29})$$

Therefore, applying a second-order Taylor expansion at $x = \xi_j$, we obtain that

$$I_j(x) = \frac{(x - \xi_j)^2}{2} I_j''(\tilde{x}) \quad (\text{B.30})$$

for some \tilde{x} on the line segment between x and ξ_j . By (F3) we know that $I_j''(x)$ is continuous. Since $f_j^2(\xi_j) \leq f_{\max}$ and $(x - \xi_j)^2 \leq (\xi_1 - \xi_j)^2$, we obtain that

$$I_j(x) = \frac{(x - \xi_j)^2}{2p(1-p)} f_j^2(\xi_j) + o((\xi_1 - \xi_j)^2) \text{ as } \xi_1 - \xi_j \rightarrow 0. \quad (\text{B.31})$$

for $x \in [\xi_j, \xi_1]$. Note that the observations from (B.28)-(B.31) also hold for $j = 1$. Define $I_{1j}(x) = \alpha_1 I_1(x) + \alpha_j I_j(x)$ and let

$$\begin{aligned} I_{1j}^\delta(x) &= \alpha_1 I_1^\delta(x) + \alpha_j I_j^\delta(x) \\ &= \alpha_1 \frac{(x - \xi_1)^2}{2p(1-p)} f_1^2(\xi_1) + \alpha_j \frac{(x - \xi_j)^2}{2p(1-p)} f_j^2(\xi_j). \end{aligned} \quad (\text{B.32})$$

Now, let x^* and $x^{\delta*}$ be the minimizers of $I_{1j}(x)$ and $I_{1j}^\delta(x)$, respectively. Note that $I_{1j}^\delta(x)$ is a quadratic function of x . Using first order conditions, we establish that

$$\begin{aligned} x^{\delta*} &= \frac{\alpha_1 \xi_1 f_1^2(\xi_1) + \alpha_j \xi_j f_j^2(\xi_j)}{\alpha_1 f_1^2(\xi_1) + \alpha_j f_j^2(\xi_j)}, \\ I_{1j}^\delta(x^{\delta*}) &= \frac{(\xi_1 - \xi_j)^2}{2p(1-p) \left(1/(\alpha_1 f_1^2(\xi_1)) + 1/(\alpha_j f_j^2(\xi_j)) \right)}. \end{aligned} \quad (\text{B.33})$$

From the definitions of $I_{1j}(x)$ and $I_{1j}^\delta(x)$ and using (B.31), it can be seen that $|I_{1j}(x) - I_{1j}^\delta(x)| = o((\xi_1 - \xi_j)^2)$ for any $x \in [\xi_j, \xi_1]$. Also, since $I_j(x)$ is non-increasing for $x < \xi_j$ and non-decreasing for $x > \xi_j$ for any j , it

can be easily seen that $x^* \in [\xi_j, \xi_1]$. Likewise, $I_j^\delta(x)$ is decreasing for $x < \xi_j$ and increasing for $x > \xi_j$ for any j , and hence, $x^{\delta*} \in [\xi_j, \xi_1]$. Now, observe that $|I_{1j}(x^*) - I_{1j}^\delta(x^{\delta*})| \leq \max(|I(x^*) - I^\delta(x^*)|, |I^\delta(x^{\delta*}) - I(x^{\delta*})|)$, and each term of max is $o((\xi_1 - \xi_j)^2)$. Therefore, we obtain that $I_{1j}(x^*) = I_{1j}^\delta(x^{\delta*}) + o((\xi_1 - \xi_j)^2)$, and therefore, $|G_j(\boldsymbol{\alpha}) - G_j^\delta(\boldsymbol{\alpha})| = o(\delta^2)$ as $\delta \rightarrow 0$.

Step 2. Next, we show that if $\arg \min_{j \neq 1} \{G_j(\boldsymbol{\alpha})\} \rightarrow i$ as $\delta \rightarrow 0$ for some $i \neq 1$, then $\xi_1 - \xi_i \rightarrow 0$ as $\delta \rightarrow 0$. Observe that $G_i(\boldsymbol{\alpha}) \leq G_2(\boldsymbol{\alpha})$ for sufficiently small δ since $\rho(\boldsymbol{\alpha}) = G_i(\boldsymbol{\alpha})$ for $\delta \in (0, \delta_0)$. From the definition of $G_j(\cdot)$ in (10) and the fact that $I_j(x)$ is strictly concave with $I_j(\mu_j) = 0$, it can be easily seen that $G_2(\boldsymbol{\alpha}) \rightarrow 0$ as $\delta \rightarrow 0$, and therefore, we establish that $G_i(\boldsymbol{\alpha}) \rightarrow 0$ as $\delta \rightarrow 0$. Now, towards a contradiction, suppose that $\liminf_{\delta \rightarrow 0} (\xi_1 - \xi_i) \geq d$ for some constant $d > 0$. Recall that $G_i(\boldsymbol{\alpha}) = \inf_x \{\alpha_1 I_1(x) + \alpha_i I_i(x)\}$ and let x_i^* be the minimizer which lies in $[\mu_i, \mu_1]$. Note that $G_i(\boldsymbol{\alpha}) \rightarrow 0$ implies both $I_1(x_i^*)$ and $I_i(x_i^*)$ converge to zero. However, due to (9) and the assumption (F3), it can be seen that $I_1(x_i^*)$ and $I_i(x_i^*)$ can converge to zero only when $x_i^* - \xi_1 \rightarrow 0$ and $x_i^* - \xi_i \rightarrow 0$ as $\delta \rightarrow 0$, respectively. This is a contradiction due to the assumption that $\liminf_{\delta \rightarrow 0} (\mu_1 - \mu_i) \geq d$. Therefore, we have that $\xi_1 - \xi_i \rightarrow 0$ as $\delta \rightarrow 0$. Exactly the same arguments can be applied to show that if $\arg \min_{j \neq 1} \{G_j^\delta(\boldsymbol{\alpha})\} \rightarrow i$ as $\delta \rightarrow 0$ for some $i \neq 1$, then it must be that $\xi_1 - \xi_i \rightarrow 0$ as $\delta \rightarrow 0$, which will be omitted.

Step 3. Next, we show that $|\rho(\boldsymbol{\alpha}) - \rho^\delta(\boldsymbol{\alpha})| = o(\delta^2)$ as $\delta \rightarrow 0$. We consider three cases.

Case (a). Suppose that the limits of $\arg \min_{j \neq 1} \{G_j(\boldsymbol{\alpha})\}$ and $\arg \min_{j \neq 1} \{G_j^\delta(\boldsymbol{\alpha})\}$ exist as $\delta \rightarrow 0$

$$\lim_{\delta \rightarrow 0} \arg \min_{j \neq 1} \{G_j(\boldsymbol{\alpha})\} = \lim_{\delta \rightarrow 0} \arg \min_{j \neq 1} \{G_j^\delta(\boldsymbol{\alpha})\}. \quad (\text{B.34})$$

In this case, we have that $|\rho(\boldsymbol{\alpha}) - \rho^\delta(\boldsymbol{\alpha})| = |G_j(\boldsymbol{\alpha}) - G_j^\delta(\boldsymbol{\alpha})|$ for some $j \neq 1$ and sufficiently small δ . Hence we immediately establish from Step 1 that $|\rho(\boldsymbol{\alpha}) - \rho^\delta(\boldsymbol{\alpha})| = o(\delta^2)$ as $\delta \rightarrow 0$.

Case (b). Suppose that the limits of $\arg \min_{j \neq 1} \{G_j(\boldsymbol{\alpha})\}$ and $\arg \min_{j \neq 1} \{G_j^\delta(\boldsymbol{\alpha})\}$ exist as $\delta \rightarrow 0$ and

$$\lim_{\delta \rightarrow 0} \arg \min_{j \neq 1} \{G_j(\boldsymbol{\alpha})\} \neq \lim_{\delta \rightarrow 0} \arg \min_{j \neq 1} \{G_j^\delta(\boldsymbol{\alpha})\}. \quad (\text{B.35})$$

Without loss of generality, let $\arg \min_{j \neq 1} \{G_j(\boldsymbol{\alpha})\} \rightarrow 2$ and $\arg \min_{j \neq 1} \{G_j^\delta(\boldsymbol{\alpha})\} \rightarrow 3$ as $\delta \rightarrow 0$. By *Step 2*, we know that $|\xi_1 - \xi_3| \rightarrow 0$ as $\delta \rightarrow 0$. Now, consider a sufficiently small δ and observe that if $G_2(\boldsymbol{\alpha}) > G_3^\delta(\boldsymbol{\alpha})$, then $|G_2(\boldsymbol{\alpha}) - G_3^\delta(\boldsymbol{\alpha})| \leq |G_3(\boldsymbol{\alpha}) - G_3^\delta(\boldsymbol{\alpha})|$ since $G_3(\boldsymbol{\alpha}) \geq G_2(\boldsymbol{\alpha})$ for sufficiently small δ . Similarly, if $G_2(\boldsymbol{\alpha}) < G_3^\delta(\boldsymbol{\alpha})$, then $|G_2(\boldsymbol{\alpha}) - G_3^\delta(\boldsymbol{\alpha})| \leq |G_2(\boldsymbol{\alpha}) - G_2^\delta(\boldsymbol{\alpha})|$ since $G_3^\delta(\boldsymbol{\alpha}) \leq G_2^\delta(\boldsymbol{\alpha})$. Combining these observations, it can be seen that

$$\begin{aligned} |\rho(\boldsymbol{\alpha}) - \rho^\delta(\boldsymbol{\alpha})| &= |G_2(\boldsymbol{\alpha}) - G_3^\delta(\boldsymbol{\alpha})| \\ &\leq \max(|G_2(\boldsymbol{\alpha}) - G_2^\delta(\boldsymbol{\alpha})|, |G_3(\boldsymbol{\alpha}) - G_3^\delta(\boldsymbol{\alpha})|) \\ &= o(\delta^2) \end{aligned} \quad (\text{B.36})$$

as $\delta \rightarrow 0$, where the last equality follows from Step 1.

Case (c). Lastly, suppose that the limit of $\arg \min_{j \neq 1} \{G_j(\boldsymbol{\alpha})\}$ or $\arg \min_{j \neq 1} \{G_j^\delta(\boldsymbol{\alpha})\}$ does not exist as $\delta \rightarrow 0$. In this case, fix $i \neq 1$ and $i' \neq 1$ and consider two sets of intervals such that

$$\begin{aligned} S_i &= \{\delta \geq 0 \mid \arg \min_{j \neq 1} \{G_j(\boldsymbol{\alpha})\} = i\} \\ S_{i'} &= \{\delta \geq 0 \mid \arg \min_{j \neq 1} \{G_j^\delta(\boldsymbol{\alpha})\} = i'\}. \end{aligned} \quad (\text{B.37})$$

Let $R_{ii'} = S_i \cup S_{i'}$ and observe that

$$\begin{aligned} \lim_{\substack{\delta \rightarrow 0 \\ \delta \in R_{ii'}}} \arg \min_{j \neq 1} \{G_j(\boldsymbol{\alpha})\} &= i \\ \lim_{\substack{\delta \rightarrow 0 \\ \delta \in R_{ii'}}} \arg \min_{j \neq 1} \{G_j^\delta(\boldsymbol{\alpha})\} &= i'. \end{aligned} \quad (\text{B.38})$$

Hence, we can apply the analysis of Case (a) (respectively, Case (b)) if $i = i'$ (respectively, if $i \neq i'$) to establish that $|\rho(\boldsymbol{\alpha}) - \rho^\delta(\boldsymbol{\alpha})| = o(\delta^2)$ as $\delta \rightarrow 0$. These arguments hold for arbitrary i and i' , which completes the proof of this step.

Step 4. Finally, we show that, for each $\boldsymbol{\alpha} \in \Delta^0$, $\rho(\boldsymbol{\alpha}) \geq c\delta^2$ for some positive constant c . From the observation in *Step 3*, it can be easily seen that

$$|\rho(\boldsymbol{\alpha}^*) - \rho(\boldsymbol{\alpha}^\delta)| \leq \max(|\rho(\boldsymbol{\alpha}^*) - \rho^\delta(\boldsymbol{\alpha}^*)|, |\rho^\delta(\boldsymbol{\alpha}^\delta) - \rho(\boldsymbol{\alpha}^\delta)|) = o(\delta^2) \text{ as } \delta \rightarrow 0. \quad (\text{B.39})$$

Further, from the assumption that $f_j(\xi_j) \geq f_{\min}$, we have that, for $\boldsymbol{\alpha}_{\text{eq}} = (1/k, \dots, 1/k)$,

$$\begin{aligned} \rho^\delta(\boldsymbol{\alpha}_{\text{eq}}) &\geq \min_{j \neq 1} \frac{(\xi_1 - \xi_j)^2}{2p(1-p)(2k/f_{\min})} \\ &\geq \min_{j \neq 1} \frac{\delta^2}{2p(1-p)(2k/f_{\min})} \\ &\geq c\delta^2 \end{aligned} \quad (\text{B.40})$$

for some positive constant c . Hence, from $\rho^\delta(\boldsymbol{\alpha}^\delta) \geq \rho^\delta(\boldsymbol{\alpha}_{\text{eq}})$ we establish that $\rho^\delta(\boldsymbol{\alpha}^\delta) \geq c\delta^2$. By Proposition 3 and (B.39), $\rho^\delta(\boldsymbol{\alpha}^\delta) \geq c\delta^2$ implies $\rho^* \geq c\delta^2 - o(\delta^2)$ as $\delta \rightarrow 0$, and therefore, we establish that

$$\frac{\rho(\boldsymbol{\alpha}^\delta)}{\rho^*} = 1 + \frac{o(\delta^2)}{\rho^*} \rightarrow 1 \text{ as } \delta \rightarrow 0, \quad (\text{B.41})$$

which completes the proof. \square

To prove Theorem 2, we define additional notation. Let $\mathbf{x} = (x_1, \dots, x_k) \in \mathcal{R}^k$ and $\mathbf{y} = (y_1, \dots, y_k) \in \mathcal{R}_+^k$. Define $B(\mathbf{x}) := \{i \in \{1, \dots, k\} \mid x_i = \max_j \{x_j\}\}$. Note that we allow the case with $|B(\mathbf{x})| > 1$, where $|B(\mathbf{x})|$ represents the cardinality of the set $B(\mathbf{x})$. Let Θ be the set of all $(\mathbf{x}, \mathbf{y}) \in \mathcal{R}^k \times \mathcal{R}_+^k$ with $|B(\mathbf{x})| < k$. For each $(\mathbf{x}, \mathbf{y}) \in \Theta$ and $\boldsymbol{\alpha} \in \Delta$, define

$$J(\boldsymbol{\alpha}; \mathbf{x}, \mathbf{y}) = \min_{i \in B(\mathbf{x}), j \notin B(\mathbf{x})} \frac{(x_i - x_j)^2}{(1/(\alpha_i y_i^2) + 1/(\alpha_j y_j^2))}. \quad (\text{B.42})$$

Define $\boldsymbol{\alpha}(\mathbf{x}, \mathbf{y})$ to be the maximizer of $J(\boldsymbol{\alpha}; \mathbf{x}, \mathbf{y})$, i.e.,

$$\boldsymbol{\alpha}(\mathbf{x}, \mathbf{y}) = \arg \max_{\boldsymbol{\alpha} \in \Delta} J(\boldsymbol{\alpha}; \mathbf{x}, \mathbf{y}). \quad (\text{B.43})$$

Note that $J(\boldsymbol{\alpha}; \mathbf{x}, \mathbf{y})$ is a strictly concave function of $\boldsymbol{\alpha}$ since it is the minimum of the strictly concave functions. Hence, $\boldsymbol{\alpha}(\mathbf{x}, \mathbf{y})$ in (B.43) is well defined. The following two lemmas are required for the proof of Theorem 2.

LEMMA B.2. *Fix $(\mathbf{x}, \mathbf{y}) \in \Theta$. Consider a sequence of parameters $(\mathbf{x}^t, \mathbf{y}^t) \in \Theta$ such that $x_i^t \rightarrow x_i$ and $y_j^t \rightarrow y_j$ as $t \rightarrow \infty$ for each j . Then, $\boldsymbol{\alpha}(\mathbf{x}^t, \mathbf{y}^t) \rightarrow \boldsymbol{\alpha}(\mathbf{x}, \mathbf{y})$ as $t \rightarrow \infty$ and the vector $\boldsymbol{\alpha}(\mathbf{x}, \mathbf{y})$ is strictly positive.*

LEMMA B.3. *Suppose the kernel function and the bandwidth parameter satisfy (K1)-(K3). Then, if $N_{jt} \rightarrow \infty$ in probability as $t \rightarrow \infty$, $\hat{f}_{jt}(\hat{\xi}_{jt}) \rightarrow f_j(\xi_j)$ in probability as $t \rightarrow \infty$. Suppose the kernel function and the bandwidth parameter further satisfy (K4)-(K7) and that $f_j(\cdot)$ is uniformly continuous. Then, if $N_{jt} \rightarrow \infty$ almost surely as $t \rightarrow \infty$, $\hat{f}_{jt}(\hat{\xi}_{jt}) \rightarrow f_j(\xi_j)$ almost surely as $t \rightarrow \infty$.*

Proof of Theorem 2. In this proof we fix $\pi = \text{QD-C}$ and omit π in superscripts for simple exposition.

(i) *Consistency.* We prove the consistency of the algorithm in three steps.

Step 1. We first show that $\hat{\alpha}_t^\delta$ converges to a random variable in probability as $t \rightarrow \infty$. Define $s_{jt} = \sup\{s < t \mid \pi_s = j\}$ as the last time that system j is sampled before t . Since s_{jt} is non-decreasing in t , its pathwise limit is well defined, hence let $s_{j\infty} = \lim_{t \rightarrow \infty} s_{jt}$. If $s_{j\infty}$ were infinite for all $j = 1, \dots, k$ with probability one, the proof of consistency would be complete. Hence, in what follows we assume that $s_{j\infty}$ can be finite for some j . Note that system j is not sampled between s_{jt} and $t - 1$, so we have that $\hat{\xi}_{jt} = \hat{\xi}_{js_{jt}} \rightarrow \hat{\xi}_{js_{j\infty}}$ almost surely, and hence also in probability, as $t \rightarrow \infty$. Likewise, it can be seen that $\hat{f}_{jt} \rightarrow \hat{f}_{js_{j\infty}}$ in probability as $t \rightarrow \infty$. Hence, for each $\alpha \in \Delta$, using the continuous mapping theorem, we have that

$$\hat{\rho}_t^\delta(\alpha) \rightarrow \hat{\rho}_\infty^\delta(\alpha) \text{ in probability as } t \rightarrow \infty, \quad (\text{B.44})$$

where $\hat{\rho}_\infty^\delta(\alpha)$ is a random variable such that $\hat{\rho}_\infty^\delta(\alpha) = \min_{j \neq b} \{G_{j\infty}^\delta(\alpha)\}$, $b = \arg \max_{j=1, \dots, k} \{\hat{\xi}_{js_{j\infty}}\}$, and

$$G_{j\infty}^\delta(\alpha) = \frac{(\hat{\xi}_{bs_{b\infty}} - \hat{\xi}_{js_{j\infty}})^2}{2p(1-p)(1/\alpha_b \hat{f}_{bs_{b\infty}}^2 + 1/\alpha_j \hat{f}_{js_{j\infty}}^2)}. \quad (\text{B.45})$$

for $j \neq b$. Note from Proposition 3 that $\hat{\rho}_t^\delta(\cdot)$ and $\hat{\rho}_\infty^\delta(\cdot)$ are strictly concave functions and have unique maximizers, so let $\hat{\alpha}_t^\delta = \arg \max_{\alpha \in \Delta} \{\hat{\rho}_t^\delta(\alpha)\}$ and $\hat{\alpha}_\infty^\delta = \arg \max_{\alpha \in \Delta} \{\hat{\rho}_\infty^\delta(\alpha)\}$. Applying Theorem 5.2 of Knight (1999), we have that $\hat{\alpha}_t^\delta \rightarrow \hat{\alpha}_\infty^\delta$ in probability as $t \rightarrow \infty$. Note also that $\hat{\rho}_\infty^\delta(\alpha)$ is zero if $\hat{\xi}_{bs_{b\infty}} = \hat{\xi}_{is_{i\infty}}$ for some $i \neq b$, which occurs with zero probability. Further, conditional that $\hat{\xi}_{bs_{b\infty}} \neq \hat{\xi}_{is_{i\infty}}$ for all $i \neq b$, we have that $\hat{\rho}_\infty^\delta(\hat{\alpha}_\infty^\delta) \geq \hat{\rho}_\infty^\delta(\alpha_{\text{eq}}) > 0$, where $\alpha_{\text{eq}} = (1/k, \dots, 1/k)$. Hence, we have that $\mathbb{P}(\hat{\rho}_\infty^\delta(\hat{\alpha}_\infty^\delta) > 0) = \mathbb{P}(\hat{\alpha}_\infty^\delta \in \Delta^0) = 1$.

Step 2. We next show that $u_{it} := \hat{\alpha}_{it}^\delta - \alpha_{it}$, $i = 1, \dots, k$, converges to zero in probability as $t \rightarrow \infty$. To improve the clarity of exposition, we fix i and use $j, l \in \{1, \dots, k\}$ to index systems. Fix $K < \infty$ and define $A_t = \{j \in \{1, \dots, k\} \mid N_{jt} < K\}$. Note that the random variable N_{jt} is monotone with t . Also, since $A_{t+1} \subset A_t$ for each t , the limit $A_\infty = \{j \in \{1, \dots, k\} \mid N_{j\infty} < K\}$ is well defined by the monotone convergence theorem. It is easily seen that $\hat{\alpha}_{js_{jt}}^\delta - \alpha_{js_{jt}} \geq 0$; otherwise, $\hat{\alpha}_{ls_{jt}}^\delta - \alpha_{ls_{jt}} < 0$ for all $l = 1, \dots, k$, leading to a contradiction that $1 = \sum_{l=1}^k \hat{\alpha}_{ls_{jt}}^\delta < \sum_{l=1}^k \alpha_{ls_{jt}} = 1$. By definition, $\alpha_{jt} \leq K/t$ for $j \in A_t$. For $j \notin A_t$, observe that

$$\alpha_{jt+1} = \begin{cases} \frac{t}{t+1} \alpha_{jt} & \text{if } \pi_t \neq j \\ \frac{t}{t+1} \alpha_{jt} + \frac{1}{t+1} & \text{if } \pi_t = j. \end{cases} \quad (\text{B.46})$$

Further, observe that for $j \notin A_t$,

$$\alpha_{jt} \leq \alpha_{jt-1} \leq \dots \leq \alpha_{js_{jt}+1} = \frac{s_{jt}}{s_{jt}+1} \alpha_{js_{jt}} + \frac{1}{s_{jt}+1} \leq \hat{\alpha}_{js_{jt}}^\delta + \frac{1}{s_{jt}+1}, \quad (\text{B.47})$$

where the last inequality follows from the previous observation that $\hat{\alpha}_{js_{jt}}^\delta - \alpha_{js_{jt}} \geq 0$ and the other relations follow from (B.46). Hence, we can write that

$$\begin{aligned} \hat{\alpha}_{it}^\delta - \alpha_{it} &= \sum_{j \neq i} (\alpha_{jt} - \hat{\alpha}_{jt}^\delta) \\ &\leq \sum_{j \notin A_t} \left(\hat{\alpha}_{js_{jt}}^\delta - \hat{\alpha}_{jt}^\delta + \frac{1}{s_{jt}+1} \right) + \sum_{j \in A_t} \left(\frac{K}{t} - \hat{\alpha}_{jt}^\delta \right) \\ &=: v_t \end{aligned} \quad (\text{B.48})$$

where the inequality follows from (B.47) and the definition of A_t . From *Step 1* recall that $\hat{\alpha}_{jt}^\delta \rightarrow \hat{\alpha}_{j\infty}^\delta$ in probability as $t \rightarrow \infty$. Also, by the definition of s_{jt} it is trivial to check that $\hat{\alpha}_{js_{jt}}^\delta \rightarrow \hat{\alpha}_{js_{j\infty}}^\delta$ almost surely, and

hence also in probability as $t \rightarrow \infty$. Hence, by the continuous mapping theorem and the monotonicity of the set A_t (i.e., $A_{t+1} \subset A_t$ for each t), we have that

$$v_t \rightarrow v_\infty := v_\infty^a - v_\infty^b \text{ in probability as } t \rightarrow \infty, \quad (\text{B.49})$$

where

$$\begin{aligned} v_\infty^a &:= \sum_{j \notin A_\infty} \left(\hat{\alpha}_{js_{j\infty}}^\delta - \hat{\alpha}_{j\infty}^\delta + \frac{1}{s_{j\infty} + 1} \right) \\ v_\infty^b &:= \sum_{j \in A_\infty} \hat{\alpha}_{j\infty}^\delta. \end{aligned} \quad (\text{B.50})$$

Now, fix any $\epsilon \in (0, 1)$ and $\epsilon_1 \in (0, \epsilon)$. Then, observe that

$$\begin{aligned} \mathbf{P}(u_{it} > \epsilon) &\leq \mathbf{P}(v_t > \epsilon) \\ &= \mathbf{P}(v_t - v_\infty > \epsilon - v_\infty) \\ &\leq \mathbf{P}(v_t - v_\infty > \epsilon - v_\infty^a), \end{aligned} \quad (\text{B.51})$$

where the first inequality follows from the fact that $u_{it} \leq v_t$ from (B.48) and the second follows from the definition of v_∞ in (B.49) and the fact that $v_\infty^b = \sum_{j \in A_\infty} \hat{\alpha}_{j\infty}^\delta > 0$. Let $B_\infty(\epsilon, \epsilon_1) := \{v_\infty^a < \epsilon - \epsilon_1\}$ and $B_\infty^c(\epsilon, \epsilon_1)$ be the complement of the set $B_\infty(\epsilon, \epsilon_1)$. Then we can write that

$$\begin{aligned} \mathbf{P}(v_t - v_\infty > \epsilon - v_\infty^a) &= \mathbf{P}(\{v_t - v_\infty > \epsilon - v_\infty^a\} \cap B_\infty(\epsilon, \epsilon_1)) + \mathbf{P}(\{v_t - v_\infty > \epsilon - v_\infty^a\} \cap B_\infty^c(\epsilon, \epsilon_1)) \\ &\leq \mathbf{P}(v_t - v_\infty > \epsilon_1) + \mathbf{P}(B_\infty^c(\epsilon, \epsilon_1)). \end{aligned} \quad (\text{B.52})$$

Note that $\mathbf{P}(v_t - v_\infty > \epsilon_1)$ in (B.52) converges to zero by (B.49). Further, note that $s_{j\infty} \rightarrow \infty$ almost surely (hence also in probability) as $K \rightarrow \infty$ for $j \notin A_\infty$. Therefore, $\mathbf{P}(B_\infty^c(\epsilon, \epsilon_1))$ in (B.52) can be made arbitrarily close to zero by letting K be sufficiently large. Combining these observations into (B.51), we deduce that $\mathbf{P}(u_{it} > \epsilon) \rightarrow 0$ as $t \rightarrow \infty$.

To show that $\mathbf{P}(u_{it} < -\epsilon) \rightarrow 0$ as $t \rightarrow \infty$, observe that, if $u_{it} < -\epsilon$, then there exists $j \neq i$ such that $u_{jt} \geq \epsilon/(k-1)$; otherwise, it would lead to a contradiction to the fact that $\sum_{j=1}^k u_{jt} = \sum_{j=1}^k (\hat{\alpha}_{jt}^\delta - \alpha_{jt}) = 0$. Therefore,

$$\mathbf{P}(u_{it} < -\epsilon) \leq \mathbf{P}(u_{jt} \geq \epsilon/(k-1)) \quad (\text{B.53})$$

for some $j \neq i$. The right-hand side of the preceding inequality converges to zero by our previous observations from (B.51) and (B.52). Hence, we conclude that $u_{it} = \hat{\alpha}_{jt}^\delta - \alpha_{jt} \rightarrow 0$ in probability as $t \rightarrow \infty$. Using the result from *Step 1*, this implies that $\alpha_t \rightarrow \hat{\alpha}_\infty^\delta$ in probability as $t \rightarrow \infty$.

Step 3. Lastly, fix any $K_1 < \infty$ and $\epsilon_2 \in (0, 1)$. Observe that

$$\begin{aligned} &\mathbf{P}(N_{jt} < K_1) \\ &= \mathbf{P}(\alpha_{jt} < K_1/t) \\ &= \mathbf{P}(\alpha_{jt} - \hat{\alpha}_{j\infty}^\delta - K_1/t < -\hat{\alpha}_{j\infty}^\delta) \\ &= \mathbf{P}(\alpha_{jt} - \hat{\alpha}_{j\infty}^\delta - K_1/t < -\hat{\alpha}_{j\infty}^\delta, \hat{\alpha}_{j\infty}^\delta \geq \epsilon_2) + \mathbf{P}(\alpha_{jt} - \hat{\alpha}_{j\infty}^\delta - K_1/t < -\hat{\alpha}_{j\infty}^\delta, \hat{\alpha}_{j\infty}^\delta < \epsilon_2) \\ &\leq \mathbf{P}(\alpha_{jt} - \hat{\alpha}_{j\infty}^\delta - K_1/t < -\epsilon_2) + \mathbf{P}(\hat{\alpha}_{j\infty}^\delta < \epsilon_2) \end{aligned} \quad (\text{B.54})$$

From *Step 2*, we have that $\alpha_{jt} - \hat{\alpha}_{j\infty}^\delta - K_1/t \rightarrow 0$ in probability as $t \rightarrow \infty$, so that $\mathbf{P}(\alpha_{jt} - \hat{\alpha}_{j\infty}^\delta - K_1/t < -\epsilon_2) \rightarrow 0$ as $t \rightarrow \infty$. Further, $\mathbf{P}(\hat{\alpha}_{j\infty}^\delta < \epsilon_2)$ in (B.54) can be made arbitrarily close to zero by letting $\epsilon_2 \rightarrow 0$

because $P(\hat{\alpha}_\infty^\delta \in \Delta^0) = 1$. Hence, we conclude that $P(N_{jt} < K_1) \rightarrow 0$ as $t \rightarrow 0$ for any $K_1 < \infty$, which implies that $N_{jt} \rightarrow \infty$ in probability as $t \rightarrow \infty$. This concludes the proof of consistency.

(ii) *Relative efficiency.* Given the consistency of the algorithm, it follows that $\hat{\xi}_{jt} \rightarrow \xi_j$ and $\hat{f}_{jt} \rightarrow f_j(\xi_j)$ as $t \rightarrow \infty$ in probability for $j = 1, \dots, k$ by Proposition 1 and Lemma B.3. Applying Lemma B.2 with $(\mathbf{x}^t, \mathbf{y}^t)$ replaced with $(\hat{\xi}_t, \hat{\mathbf{f}}_t)$ and (\mathbf{x}, \mathbf{y}) replaced with $(\boldsymbol{\xi}, \mathbf{f})$, where $\mathbf{f} = (f_1(\xi_1), \dots, f_k(\xi_k))$, it follows that $\hat{\alpha}_t^\delta \rightarrow \alpha^\delta$ as $t \rightarrow \infty$. Moreover, from part (i) of this proof we have that $\hat{\alpha}_t^\delta - \alpha_t \rightarrow 0$ in probability as $t \rightarrow \infty$. Consequently, we have that $\alpha_t \rightarrow \alpha^\delta$ as $t \rightarrow \infty$ and that $\rho(\alpha_t)/\rho(\alpha^\delta) \rightarrow 1$ in probability as $t \rightarrow \infty$. Since $\rho(\alpha_T)/\rho(\alpha^\delta)$ is bounded, we also have that

$$\mathbb{E} \left(\frac{\rho(\alpha_t)}{\rho(\alpha^\delta)} \right) \rightarrow 1 \text{ as } t \rightarrow \infty, \quad (\text{B.55})$$

from which (24) follows. From Proposition 3 we have that $\rho(\alpha^\delta)/\rho(\alpha^*) \rightarrow 1$ as $\delta \rightarrow 0$, which completes the proof of this theorem. \square

To show Proposition 4 we need the following lemma.

LEMMA B.4. *Consider $\mathbf{x} = (x_1, \dots, x_k) \in \mathbf{R}^k$ and $\mathbf{y} = (y_1, \dots, y_k) \in \mathbf{R}^k$ with $k \geq 1$. Suppose that there exists a constant $c > 0$ such that $1 - c \leq x_j/y_j \leq 1 + c$ for each $j = 1, \dots, k$, then*

$$1 - c \leq \frac{\min_{j=1, \dots, k} x_j}{\min_{j=1, \dots, k} y_j} \leq 1 + c. \quad (\text{B.56})$$

Proof of Proposition 4. For ease of notation we define $G_j(\boldsymbol{\alpha}; \mathbf{h}) = G_j^{\delta, \epsilon}(\boldsymbol{\alpha})$ for each $j \neq 1$ and let $\rho(\boldsymbol{\alpha}; \mathbf{h}) = \rho^{\delta, \epsilon}(\boldsymbol{\alpha})$. Further, we define

$$\begin{aligned} I_{1j}(x; \mathbf{h}) &= p \log \left(\frac{p}{p + (x - \xi_1)h_{1j}} \right) + (1 - p) \log \left(\frac{1 - p}{1 - p - (x - \xi_1)h_{1j}} \right) \\ I_{j1}(x; \mathbf{h}) &= p \log \left(\frac{p}{p + (x - \xi_j)h_{j1}} \right) + (1 - p) \log \left(\frac{1 - p}{1 - p - (x - \xi_j)h_{j1}} \right). \end{aligned} \quad (\text{B.57})$$

It is useful to observe that

$$\frac{\rho(\boldsymbol{\alpha})}{\rho^{\delta, \epsilon}(\boldsymbol{\alpha})} = \frac{\rho(\boldsymbol{\alpha}; \mathbf{h})}{\rho^{\delta, \epsilon}(\boldsymbol{\alpha})} \frac{\rho(\boldsymbol{\alpha})}{\rho(\boldsymbol{\alpha}; \mathbf{h})} = \frac{\rho(\boldsymbol{\alpha}; \mathbf{h})}{\rho^{\delta}(\boldsymbol{\alpha}; \mathbf{h})} \frac{\rho(\boldsymbol{\alpha})}{\rho(\boldsymbol{\alpha}; \mathbf{h})}. \quad (\text{B.58})$$

Further, it can be seen that if $\arg \min_{j \neq 1} \{G_j(\boldsymbol{\alpha})\} \rightarrow i$ as $\delta \rightarrow 0$ for some $i \neq 1$, then $\xi_1 - \xi_i \rightarrow 0$ as $\delta \rightarrow 0$. The proof of the preceding statement follows exactly the same arguments given in *Step 2* of the proof for Proposition 3, and hence skipped.

Now, we prove the proposition in three steps. In *Step 1*, we find the lower and upper bounds of the second term on the right-hand side of (B.58). In *Step 2*, we find the lower and upper bounds of the first term on the right-hand side of (B.58). In *Step 3*, we combine these observations to conclude the proof.

Step 1. We first show that $1 - \epsilon \leq \rho(\boldsymbol{\alpha})/\rho(\boldsymbol{\alpha}; \mathbf{h}) \leq 1 + \epsilon$. To this end, it suffices to show that $1 - \epsilon \leq G_j(\boldsymbol{\alpha})/G_j(\boldsymbol{\alpha}; \mathbf{h}) \leq 1 + \epsilon$ for each j , since the desired result follows immediately from Lemma B.4. Using the relationship $\log(1 \pm z) \leq \pm z$ for $0 < z < 1$ and assumption (F3), it can be seen that

$$1 - \epsilon \leq \frac{I_{1j}(x)}{I_{1j}(x; \mathbf{h})}, \frac{I_{j1}(x)}{I_{j1}(x; \mathbf{h})} \leq 1 + \epsilon \quad (\text{B.59})$$

for each $j \neq 1$ and $x \in [\xi_j, \xi_1]$. Now, let x^* and x^h be the minimizers of $\alpha_1 I_1(x) + \alpha_j I_j(x)$ and $\alpha_1 I_1(x; \mathbf{h}) + \alpha_j I_{j1}(x; \mathbf{h})$, respectively, so that $G_j(\boldsymbol{\alpha}) = \alpha_1 I_1(x^*) + \alpha_j I_j(x^*)$ and $G_j(\boldsymbol{\alpha}; \mathbf{h}) = \alpha_1 I_1(x^h; \mathbf{h}) + \alpha_j I_{j1}(x^h; \mathbf{h})$. We establish that

$$\frac{G_j(\boldsymbol{\alpha})}{G_j(\boldsymbol{\alpha}; \mathbf{h})} \leq \frac{\alpha_1 I_1(x^h) + \alpha_j I_j(x^h)}{\alpha_1 I_1(x^h; \mathbf{h}) + \alpha_j I_{j1}(x^h; \mathbf{h})} \leq (1 + \epsilon), \quad (\text{B.60})$$

where the first inequality follows by the definition of x^h and the second follows from (B.59). Likewise, we can establish the lower bound as

$$\frac{G_j(\boldsymbol{\alpha})}{G_j(\boldsymbol{\alpha}; \mathbf{h})} \geq \frac{\alpha_1 I_1(x^*) + \alpha_j I_j(x^*)}{\alpha_1 I_{1j}(x^*; \mathbf{h}) + \alpha_j I_{j1}(x^*; \mathbf{h})} \geq (1 - \epsilon). \quad (\text{B.61})$$

This completes the proof of *Step 1*.

Step 2. It remains to bound the first term of (B.58); in particular, $1 - u(\theta) \leq \rho(\boldsymbol{\alpha}; \mathbf{h})/\rho^\delta(\boldsymbol{\alpha}; \mathbf{h}) \leq 1 + u(\theta)$. Using Lemma B.4 it suffices to show that $1 - u(\theta) \leq G_j(\boldsymbol{\alpha}; \mathbf{h})/G_j^\delta(\boldsymbol{\alpha}; \mathbf{h}) \leq 1 + u(\theta)$ for each j . To this end, recall the definitions of $\rho(\boldsymbol{\alpha}; \mathbf{h}) = \min_{j \neq 1} \{G_j(\boldsymbol{\alpha}; \mathbf{h})\}$ and $\rho^\delta(\boldsymbol{\alpha}; \mathbf{h}) = \min_{j \neq 1} \{G_j^\delta(\boldsymbol{\alpha}; \mathbf{h})\}$, where

$$\begin{aligned} G_j(\boldsymbol{\alpha}; \mathbf{h}) &= \inf_x \{\alpha_1 I_{1j}(x; \mathbf{h}) + \alpha_j I_{j1}(x; \mathbf{h})\} \\ G_j^\delta(\boldsymbol{\alpha}; \mathbf{h}) &= \inf_x \{\alpha_1 I_{1j}^\delta(x; \mathbf{h}) + \alpha_j I_{j1}(x; \mathbf{h})\}, \end{aligned} \quad (\text{B.62})$$

and

$$\begin{aligned} I_{1j}^\delta(x; \mathbf{h}) &= \frac{(x - \xi_j)^2}{2p(1-p)/h_{1j}^2} \\ I_{j1}^\delta(x; \mathbf{h}) &= \frac{(x - \xi_j)^2}{2p(1-p)/h_{j1}^2}. \end{aligned} \quad (\text{B.63})$$

Applying a second-order Taylor expansion for $I_{j1}(x; \mathbf{h})$ at $x = \xi_j$, we obtain that

$$\begin{aligned} I_{j1}(x; \mathbf{h}) &= I_{j1}(\xi_j; \mathbf{h}) + (x - \xi_j)I'_{j1}(\xi_j; \mathbf{h}) + \frac{(x - \xi_j)^2}{2}I''_{j1}(\xi_j; \mathbf{h}) + \frac{(x - \xi_j)^3}{6}I'''_{j1}(\tilde{x}; \mathbf{h}) \\ &= \frac{(x - \xi_j)^2}{2p(1-p)}h_{j1}^2 + \frac{(x - \xi_j)^3}{6}I'''_{j1}(\tilde{x}; \mathbf{h}) \\ &= I_{j1}^\delta(x; \mathbf{h}) + \frac{(x - \xi_j)^3}{6}I'''_{j1}(\tilde{x}; \mathbf{h}) \end{aligned} \quad (\text{B.64})$$

for $x \in [\xi_j, \xi_1]$, where \tilde{x} lies between x and ξ_j and

$$I'''_{j1}(x; \mathbf{h}) = 2h_{j1}^3 \left(\frac{-p}{(p + (x - \xi_j)h_{j1})^3} + \frac{1-p}{(1-p - (x - \xi_j)h_{j1})^3} \right). \quad (\text{B.65})$$

Note that the first equality of (B.64) follows from the fact that $I_{j1}(\xi_j; \mathbf{h}) = 0$ and $I'_{j1}(\xi_j; \mathbf{h}) = 0$, and the last equality follows from the definition of $I_{j1}^\delta(\cdot)$ in (B.63). From (B.64), we have that

$$\frac{I_{j1}(x; \mathbf{h})}{I_{j1}^\delta(x; \mathbf{h})} \leq 1 + \left| \frac{(x - \xi_j)^3 I'''_{j1}(\tilde{x}; \mathbf{h})/6}{(x - \xi_j)^2 h_{j1}^2 / (2p(1-p))} \right| \leq 1 + u(\theta), \quad (\text{B.66})$$

which follows from the definition of $u(\theta)$ and straightforward algebra. Likewise, we can establish the lower bound as

$$\frac{I_{j1}(x; \mathbf{h})}{I_{j1}^\delta(x; \mathbf{h})} \geq 1 - u(\theta). \quad (\text{B.67})$$

To summarize our observations in Step 4 so far, we have that

$$1 - u(\theta) \leq \frac{I_{j1}(x; \mathbf{h})}{I_{j1}^\delta(x; \mathbf{h})} \leq 1 + u(\theta). \quad (\text{B.68})$$

Further, exactly the same steps from (B.63) to (B.67) can be applied to show that

$$1 - u(\theta) \leq \frac{I_{1j}(x; \mathbf{h})}{I_{1j}^\delta(x; \mathbf{h})} \leq 1 + u(\theta) \quad (\text{B.69})$$

(the detailed derivation is omitted).

Next, define x^h and $x^{h\delta}$ as the minimizers of $\alpha_1 I_{1j}(x; \mathbf{h}) + \alpha_j I_{j1}(x; \mathbf{h})$ and $\alpha_1 I_{1j}^\delta(x; \mathbf{h}) + \alpha_j I_{j1}^\delta(x; \mathbf{h})$, respectively. Then we can write

$$\frac{G_j(\boldsymbol{\alpha}; \mathbf{h})}{G_j^\delta(\boldsymbol{\alpha}; \mathbf{h})} \leq \frac{\alpha_1 I_{1j}(x^{h\delta}; \mathbf{h}) + \alpha_j I_{j1}(x^{h\delta}; \mathbf{h})}{\alpha_1 I_{1j}^\delta(x^{h\delta}; \mathbf{h}) + \alpha_j I_{j1}^\delta(x^{h\delta}; \mathbf{h})} \leq 1 + u(\theta), \quad (\text{B.70})$$

where the first inequality follows by the definition of $x^{h\delta}$ and the second follows from (B.68)-(B.69). Likewise, we can write

$$\frac{G_j(\boldsymbol{\alpha}; \mathbf{h})}{G_j^\delta(\boldsymbol{\alpha}; \mathbf{h})} \geq \frac{\alpha_1 I_{1j}(x^h; \mathbf{h}) + \alpha_j I_{j1}(x^h; \mathbf{h})}{\alpha_1 I_{1j}^\delta(x^h; \mathbf{h}) + \alpha_j I_{j1}^\delta(x^h; \mathbf{h})} \leq 1 - u(\theta). \quad (\text{B.71})$$

The proof of this step is complete from the last two inequalities.

Step 3. Finally, observe that

$$\begin{aligned} \frac{\rho(\boldsymbol{\alpha}^{\delta, \epsilon^*})}{\rho(\boldsymbol{\alpha}^*)} &= \frac{\rho(\boldsymbol{\alpha}^{\delta, \epsilon^*})}{\rho(\boldsymbol{\alpha}^*; \mathbf{h})} \frac{\rho(\boldsymbol{\alpha}^*; \mathbf{h})}{\rho(\boldsymbol{\alpha}^*)} \\ &\geq \frac{\rho(\boldsymbol{\alpha}^{\delta, \epsilon^*}; \mathbf{h})}{\rho(\boldsymbol{\alpha}^*; \mathbf{h})} \frac{1 - \epsilon}{1 + \epsilon} \\ &= \frac{\rho^\delta(\boldsymbol{\alpha}^{\delta, \epsilon^*}; \mathbf{h})}{\rho(\boldsymbol{\alpha}^*; \mathbf{h})} \frac{\rho(\boldsymbol{\alpha}^{\delta, \epsilon^*}; \mathbf{h})}{\rho^\delta(\boldsymbol{\alpha}^{\delta, \epsilon^*}; \mathbf{h})} \frac{1 - \epsilon}{1 + \epsilon} \\ &\geq \frac{\rho^\delta(\boldsymbol{\alpha}^*; \mathbf{h})}{\rho(\boldsymbol{\alpha}^*; \mathbf{h})} \frac{\rho(\boldsymbol{\alpha}^{\delta, \epsilon^*}; \mathbf{h})}{\rho^\delta(\boldsymbol{\alpha}^{\delta, \epsilon^*}; \mathbf{h})} \frac{1 - \epsilon}{1 + \epsilon}, \end{aligned} \quad (\text{B.72})$$

where the first inequality follows from the fact that $1 - \epsilon \leq \rho(\boldsymbol{\alpha})/\rho(\boldsymbol{\alpha}; \mathbf{h}) \leq 1 + \epsilon$ for any $\boldsymbol{\alpha} \in \Delta^0$ and the second inequality follows from the fact that $\boldsymbol{\alpha}^{\delta, \epsilon^*}$ is the maximizer of $\rho^\delta(\cdot; \mathbf{h}) = \rho^{\delta, \epsilon}(\cdot)$. Using the observations from *Step 2*, it can be seen that

$$\begin{aligned} \frac{\rho^\delta(\boldsymbol{\alpha}^*; \mathbf{h})}{\rho(\boldsymbol{\alpha}^*; \mathbf{h})} &\geq \frac{1}{1 + u(\theta)} \\ \frac{\rho(\boldsymbol{\alpha}^{\delta, \epsilon^*}; \mathbf{h})}{\rho^\delta(\boldsymbol{\alpha}^{\delta, \epsilon^*}; \mathbf{h})} &\geq 1 - u(\theta), \end{aligned} \quad (\text{B.73})$$

which establishes the desired result of the proposition. \square

The proof of Theorem 3 requires the following lemma.

LEMMA B.5. *If $N_{jt} \rightarrow \infty$ almost surely (respectively, in probability) as $t \rightarrow \infty$ for each j , then $\hat{h}_{jbt} \rightarrow h_{j1}$ and $\hat{h}_{bjt} \rightarrow h_{1j}$ almost surely (respectively, in probability) as $t \rightarrow \infty$.*

Proof of Theorem 3. We omit the proof for consistency because it follows from the identical steps as those in the proof of Theorem 2, with \hat{f}_{jt} being replaced by \hat{h}_{jbt} or \hat{h}_{bjt} and $f_j(\xi_j)$ by h_{j1} or h_{1j} , along with Lemma B.5.

Given the consistency of the algorithm, it follows that $\hat{\xi}_{jt} \rightarrow \xi_j$ and $\hat{h}_{jt} \rightarrow h_j(u)$ as $t \rightarrow \infty$ in probability for all j by Proposition 1 and Lemma B.5. Applying Lemma B.2 with $(\mathbf{x}^t, \mathbf{y}^t)$ replaced with $(\hat{\boldsymbol{\xi}}_t, \hat{\mathbf{h}}_t)$ and (\mathbf{x}, \mathbf{y}) replaced with $(\boldsymbol{\xi}, \mathbf{h})$, where $\hat{\mathbf{h}}_t = (\hat{h}_{1t}(u), \dots, \hat{h}_{kt}(u))$ and $\mathbf{h} = (h_1(u), \dots, h_k(u))$, it follows that $\hat{\boldsymbol{\alpha}}_t^{\delta, \epsilon} \rightarrow \boldsymbol{\alpha}^{\delta, \epsilon^*}$ as $t \rightarrow \infty$. Moreover, by construction of Algorithm 3 it is not difficult to see that the term $\boldsymbol{\alpha}_t - \hat{\boldsymbol{\alpha}}_t^{\delta, \epsilon} \rightarrow 0$. Consequently, we have that $\boldsymbol{\alpha}_t \rightarrow \boldsymbol{\alpha}^{\delta, \epsilon^*}$ as $t \rightarrow \infty$ and that $\rho(\boldsymbol{\alpha}_t)/\rho(\boldsymbol{\alpha}^{\delta, \epsilon^*}) \rightarrow 1$ in probability as $t \rightarrow \infty$. Since $\rho(\boldsymbol{\alpha}_t)/\rho(\boldsymbol{\alpha}^{\delta, \epsilon^*})$ is bounded, we establish that

$$\mathbb{E} \left(\frac{\rho(\boldsymbol{\alpha}_t)}{\rho(\boldsymbol{\alpha}^{\delta, \epsilon^*})} \right) \rightarrow 1 \text{ as } t \rightarrow \infty. \quad \square \quad (\text{B.74})$$

Proof of Lemma A.1. In the definition of $G_j^\delta(\boldsymbol{\alpha})$, p , ξ_j , and $f_j(\xi_j)$ are constants that are independent of $\boldsymbol{\alpha}$. Hence, without loss of generality, we can show that the function $\ell_j(\boldsymbol{\alpha})$, defined as

$$\ell_j(\boldsymbol{\alpha}) = \frac{1}{1/\alpha_1 + 1/\alpha_j} \quad (\text{B.75})$$

is strongly concave. Observe that the Hessian matrix for $\ell_j(\boldsymbol{\alpha})$ with respect to α_1 and α_j is

$$\nabla^2 \ell_j(\boldsymbol{\alpha}) = \frac{1}{(1/\alpha_1 + 1/\alpha_j)^2} \begin{bmatrix} \frac{2}{\alpha_1^3} \left(\frac{1/\alpha_1}{1/\alpha_1 + 1/\alpha_j} - 1 \right) & \frac{2/\alpha_1^2 \alpha_j^2}{1/\alpha_1 + 1/\alpha_j} \\ \frac{2/\alpha_1^2 \alpha_j^2}{1/\alpha_1 + 1/\alpha_j} & \frac{2}{\alpha_j^3} \left(\frac{1/\alpha_j}{1/\alpha_1 + 1/\alpha_j} - 1 \right) \end{bmatrix}. \quad (\text{B.76})$$

After some straightforward algebra, the determinant of the Hessian can be written as

$$|\nabla^2 \ell_j(\boldsymbol{\alpha})| = \frac{4((\alpha_1 + \alpha_j) - (\alpha_1^4 + \alpha_j^4))}{\alpha_1 \alpha_j (\alpha_1 + \alpha_j)^3}. \quad (\text{B.77})$$

In the remainder of the proof, we show that there exists a positive constant d such that

$$\min_{\boldsymbol{\alpha} \in \Delta} |\nabla^2 \ell_j(\boldsymbol{\alpha})| \geq 4d. \quad (\text{B.78})$$

First, it is trivial to check that $|\nabla^2 \ell_j(\boldsymbol{\alpha})| \rightarrow \infty$ as $\alpha_1 \rightarrow 0$ and/or $\alpha_j \rightarrow 0$. Hence, there exists a small $d \in (0, 1)$ such that

$$\min_{\boldsymbol{\alpha} \in \Delta} |\nabla^2 \ell_j(\boldsymbol{\alpha})| = \min_{\boldsymbol{\alpha} \in \Delta^d} |\nabla^2 \ell_j(\boldsymbol{\alpha})|, \quad (\text{B.79})$$

where $\Delta^d = \{\boldsymbol{\alpha} : d \leq \sum_{j=1}^k \alpha_j \leq 1 - kd\}$. Further, since $(x + y)/2 \geq \sqrt{xy}$, observe that for $\boldsymbol{\alpha} \in \Delta^d$,

$$\begin{aligned} |\nabla^2 \ell_j(\boldsymbol{\alpha})| &\geq \frac{(\alpha_1 + \alpha_j) - (\alpha_1^4 + \alpha_j^4)}{(\alpha_1 + \alpha_j)^5} \\ &\geq \frac{(\alpha_1 + \alpha_j) - (\alpha_1 + \alpha_j)^4}{(\alpha_1 + \alpha_j)^5} \\ &\geq (\alpha_1 + \alpha_j) - (\alpha_1 + \alpha_j)^4 \\ &\geq d, \end{aligned} \quad (\text{B.80})$$

which completes the proof of the lemma. \square

Proof for Proposition A.1. First, the strong consistency of the naive algorithm immediately follows, because $N_{jt}^\pi \geq c \log t \rightarrow \infty$ as $t \rightarrow \infty$ almost surely.

Next, we show that QD-C is strongly consistent under the additional assumptions in the statement of Proposition A.1. Fix a sequence of samples, $\{(X_{j1}, X_{j2}, \dots)\}_{j=1}^k$ and let $A \subset \{1, 2, \dots, k\}$ be the set of systems that are sampled only finitely many times and let $I := \{1, 2, \dots, k\} \setminus A$. Suppose, towards a contradiction, that F is non-empty and let $\tau < \infty$ be the last time that the systems in A are sampled. It is straightforward to see, using the proof of Proposition 1, that if QD-C is strongly consistent, then $\hat{\xi}_{jt} \rightarrow \xi_j$ almost surely as $t \rightarrow \infty$. Further, from Lemma B.3 we observe that \hat{f}_{jt} converges to $f_j(\xi_j)$ almost surely as $t \rightarrow \infty$ for each system $j \in I$. Also, for any system $j \in A$, $\hat{\xi}_{jt}$ and \hat{f}_{jt} are constant subsequent to the stage τ . Therefore, as $t \rightarrow \infty$,

$$\hat{\xi}_{jt} \rightarrow \hat{\xi}_{j\infty} = \begin{cases} \xi_j & \text{for } j \in I \\ \hat{\xi}_{j\tau} & \text{for } j \in A, \end{cases} \quad (\text{B.81})$$

and

$$\tilde{f}_{jt}(\hat{\xi}_{jt}) \rightarrow \tilde{f}_{j\infty} = \begin{cases} f_j(\xi_j) & \text{for } j \in I \\ \tilde{f}_{jt}(\hat{\xi}_{jt}) & \text{for } j \in A. \end{cases} \quad (\text{B.82})$$

Let $\hat{\boldsymbol{\xi}}_t = (\hat{\xi}_{1t}, \dots, \hat{\xi}_{kt})$ and $\hat{\mathbf{f}}_t = (\hat{f}_{1t}(\hat{\xi}_{jt}), \dots, f_{kt}(\hat{\xi}_{jt}))$. We may assume that $|B(\hat{\boldsymbol{\xi}}_\infty)| < k$, i.e., $\max_j \{\hat{\xi}_{j\infty}\} > \min_j \{\hat{\xi}_{j\infty}\}$, since the event, $\hat{\xi}_{1\infty} = \dots = \hat{\xi}_{k\infty}$, occurs with zero probability under the assumption (F3). Define $\hat{\boldsymbol{\alpha}}_t = \boldsymbol{\alpha}(\hat{\boldsymbol{\xi}}_t, \hat{\mathbf{f}}_t)$ for each $t \geq k$ and $\hat{\boldsymbol{\alpha}}_\infty = \boldsymbol{\alpha}(\hat{\boldsymbol{\xi}}_\infty, \hat{\mathbf{f}}_\infty)$. Applying Lemma B.2 with $(\mathbf{x}^t, \mathbf{y}^t)$ replaced with $(\hat{\boldsymbol{\xi}}_t, \hat{\mathbf{f}}_t)$ and (\mathbf{x}, \mathbf{y}) replaced with $(\hat{\boldsymbol{\xi}}_\infty, \hat{\mathbf{f}}_\infty)$, it follows that $\hat{\boldsymbol{\alpha}}_t \rightarrow \hat{\boldsymbol{\alpha}}_\infty$ as $t \rightarrow \infty$ with $\hat{\boldsymbol{\alpha}}_\infty > 0$. Further, by construction of the algorithm, $\boldsymbol{\alpha}_t - \hat{\boldsymbol{\alpha}}_t \rightarrow 0$, and therefore, $\boldsymbol{\alpha}_t \rightarrow \hat{\boldsymbol{\alpha}}_\infty$ as $t \rightarrow \infty$. However, this contradicts our assumption because $\alpha_{jt} \rightarrow 0$ for each $j \in A$ since such a system is sampled only finitely many times. Consequently, F is empty with probability 1 and QD-C is strongly consistent.

The proof for the strong consistency of QD-D follows using exactly the same logical steps as that for QD-C, with \hat{f}_{jt} being replaced by \hat{h}_{jbt} or \hat{h}_{bjt} , $f_j(\xi_j)$ by h_{j1} or h_{1j} , and Lemma B.3 by Lemma B.5. Hence, this will be omitted. \square

The proof for Proposition A.2 requires the following lemma.

LEMMA B.6 (Moderate deviations for sample quantiles). *Under assumptions (F1)-(F3), for any positive sequence $\{\delta_t\}$ such that $t\delta_t^2 \rightarrow \infty$ and $\delta_t \rightarrow 0$ as $t \rightarrow \infty$, a static algorithm $\pi(\boldsymbol{\alpha})$ for some $\boldsymbol{\alpha} \in \Delta^0$ satisfies*

$$\frac{1}{t\delta_t^2} \log \mathbb{P}(|\hat{\xi}_{jt} - \xi_j| > \delta_t) \rightarrow -\frac{\alpha_j f_j^2(\xi_j)}{2p(1-p)} \text{ as } t \rightarrow \infty \text{ for } j = 1, \dots, k \quad (\text{B.83})$$

Proof of Proposition A.2. In this proof, we fix a static algorithm $\pi = \pi(\boldsymbol{\alpha})$ for some $\boldsymbol{\alpha} \in \Delta^0$ and suppress π in the superscripts for clarity. Observe that

$$\max_{j=2, \dots, k} \mathbb{P}(\hat{\xi}_{1t} \leq \hat{\xi}_{jt}; \delta_t) \leq \mathbb{P}(\text{FS}_t; \delta_t) \leq (k-1) \max_{j=2, \dots, k} \mathbb{P}(\hat{\xi}_{1t} \leq \hat{\xi}_{jt}; \delta_t). \quad (\text{B.84})$$

Hence, if, for each $j = 2, \dots, k$,

$$\lim_{t \rightarrow \infty} \frac{1}{t\delta_t^2} \log \mathbb{P}(\hat{\xi}_{1t} \leq \hat{\xi}_{jt}; \delta_t) = -G_j(\boldsymbol{\alpha}), \quad (\text{B.85})$$

then

$$\lim_{t \rightarrow \infty} \frac{1}{t\delta_t^2} \log \mathbb{P}(\text{FS}_t; \delta_t) = -\min_{j=2, \dots, k} \{G_j(\boldsymbol{\alpha})\}. \quad (\text{B.86})$$

Therefore, the rate function of $\mathbb{P}(\text{FS}_t; \delta_t)$ can be immediately obtained once we prove (B.85) for some $j \neq 1$.

For simplicity, we consider a two-system case, but the proof can be extended in a straightforward manner to $k \geq 3$. Without loss of generality, assume that $\delta_1 = \xi_1 - \xi_2 = 1$. Observe that

$$\begin{aligned} \mathbb{P}(\text{FS}_t; \delta_t) &= \mathbb{P}(\hat{\xi}_{1t} < \hat{\xi}_{2t}; \delta_t) \\ &= \mathbb{P}(\hat{\xi}_{1t} - \xi_{1t} < \hat{\xi}_{2t} - \xi_{2t} - \delta_t; \delta_t) \\ &= \mathbb{P}(\hat{\xi}_{1t} - \xi_1 < \hat{\xi}_{2t} - \xi_2 - \delta_t) \\ &= \mathbb{P}(\hat{\xi}_{1t} - \xi_1 + \delta_t/2 < \hat{\xi}_{2t} - \xi_2 - \delta_t/2). \end{aligned} \quad (\text{B.87})$$

Define $\eta_{1t} = (\hat{\xi}_{1t} - \xi_1)/\delta_t + 1/2$ and $\eta_{2t} = (\hat{\xi}_{2t} - \xi_2)/\delta_t - 1/2$. Then, $\mathbb{P}(\text{FS}_t; \delta_t)$ can be written as

$$\mathbb{P}(\text{FS}_t; \delta_t) = \mathbb{P}(\eta_{1t} < \eta_{2t}). \quad (\text{B.88})$$

Note that for $x \leq 1/2$, $\mathbb{P}(\eta_{1t} < x) = \mathbb{P}(\hat{\xi}_{1t} < \delta_t(x - 1/2) + \xi_1)$ and therefore,

$$\lim_{t \rightarrow \infty} \frac{1}{t\delta_t^2} \log \mathbb{P}(\eta_{1t} < x) = - \left(x - \frac{1}{2}\right)^2 \frac{\alpha_1 f_1^2(\xi_1)}{2p(1-p)}. \quad (\text{B.89})$$

Likewise, we have that, for $x \geq -1/2$,

$$\lim_{t \rightarrow \infty} \frac{1}{t\delta_t^2} \log \mathbb{P}(\eta_{2t} > x) = - \left(x + \frac{1}{2}\right)^2 \frac{\alpha_2 f_2^2(\xi_2)}{2p(1-p)}. \quad (\text{B.90})$$

Using the same steps in the proof of Proposition 2, it can be easily seen that

$$\lim_{t \rightarrow \infty} \frac{1}{t\delta_t^2} \log \mathbb{P}(\eta_{1t} < \eta_{2t}) = - \inf_{x \in [-1/2, 1/2]} \{\alpha_1 \mathcal{F}_1(x) + \alpha_2 \mathcal{F}_2(x)\}, \quad (\text{B.91})$$

where

$$\begin{aligned} \mathcal{F}_1(x) &:= \left(x - \frac{1}{2}\right) \frac{f_1^2(\xi_1)}{2p(1-p)} \\ \mathcal{F}_2(x) &:= \left(x + \frac{1}{2}\right) \frac{f_2^2(\xi_2)}{2p(1-p)}. \end{aligned} \quad (\text{B.92})$$

Since $\alpha_1 \mathcal{F}_1(x) + \alpha_2 \mathcal{F}_2(x)$ is a quadratic function of x , using first order conditions, one can show that the infimum is

$$\frac{1}{2p(1-p)(1/\alpha_1 f_1^2(\xi_1) + 1/\alpha_2 f_2^2(\xi_2))}, \quad (\text{B.93})$$

which is attained at

$$x = \frac{\alpha_1 f_1^2(\xi_1) - \alpha_2 f_2^2(\xi_2)}{2(\alpha_1 f_1^2(\xi_2) + \alpha_2 f_2^2(\xi_2))}. \quad (\text{B.94})$$

This completes the proof. \square

B.2. Proofs for Auxiliary Lemmas

Proof of Lemma B.1. First, fix $x < \xi_j$ and observe that

$$\mathbb{P}\left(\hat{\xi}_{jt}^p \leq x\right) = \mathbb{P}\left(\sum_{n=1}^{N_{jt}} \mathbf{I}\{X_{j\tau_{jn}} \leq x\} > [pN_{jt}]\right). \quad (\text{B.95})$$

Note that $\mathbf{I}\{X_j \leq x\}$ is a Bernoulli random variable with success probability $F_j(x)$. Hence, applying Cramer's theorem, it can be seen that

$$\frac{1}{t} \log \mathbb{P}\left(\sum_{n=1}^{N_{jt}} \mathbf{I}\{X_{j\tau_{jn}} \leq x\} > [pN_{jt}]\right) \rightarrow -\alpha_j \sup_{\lambda} \{\lambda x - \Lambda_j(\lambda)\} \quad (\text{B.96})$$

as $t \rightarrow \infty$, where $\Lambda_j(\lambda) = \log \mathbb{E}[\exp(\lambda \mathbf{I}\{X_j \leq x\})]$. After some straightforward algebra one can show that

$$\sup_{\lambda} \{\lambda x - \Lambda_j(\lambda)\} = p \log \left(\frac{p}{F_j(x)}\right) + (1-p) \log \left(\frac{1-p}{1-F_j(x)}\right). \quad (\text{B.97})$$

From identical arguments, one can prove the case with $x' > \xi_j^p$ and this concludes the proof of the lemma.

\square

Proof of Lemma B.2. First, we show that $\boldsymbol{\alpha}(\mathbf{x}, \mathbf{y})$ is strictly positive. Note that the objective function of the optimization problem (B.43) is strictly positive at its optimal solution; to see this, notice that $\boldsymbol{\alpha}_e = (1/k, \dots, 1/k)$ is a feasible solution and the objective function is positive at $\boldsymbol{\alpha}_e$. Suppose that $\alpha_j(\mathbf{x}, \mathbf{y}) = 0$ for some j . Then the objective function of the optimization problem (B.43) is 0, which contradicts the optimality of $\boldsymbol{\alpha}(\mathbf{x}, \mathbf{y})$.

Without loss of generality, assume that $x_1 \geq \dots \geq x_k$ and let $b = |B(\mathbf{x})| < k$. Define $\epsilon \in (0, (x_1 - x_{b+1})/2)$ and let $t_0 < \infty$ such that $\max_j \{|x_j^t - x_j|\} \leq \epsilon$ and $\max_j \{|y_j^t - y_j|\} \leq \epsilon$ for all $t \geq t_0$. That is, for each $t \geq t_0$, $(\mathbf{x}^t, \mathbf{y}^t)$ is in a compact set defined by (\mathbf{x}, \mathbf{y}) and ϵ . Also, x_j^t is sufficiently close to x_j for each $j = 1, \dots, k$ so that $B(\mathbf{x}^t) = B(\mathbf{x})$ for $t \geq t_0$. In the rest of the proof, it suffices to consider $t \geq t_0$.

Suppose, towards a contradiction, that $\boldsymbol{\alpha}(\mathbf{x}^t, \mathbf{y}^t)$ does not converge to $\boldsymbol{\alpha}(\mathbf{x}, \mathbf{y})$. Then there exists a convergent subsequence, $\{t_1, t_2, \dots\}$, such that

$$\boldsymbol{\alpha}(\mathbf{x}^{t_n}, \mathbf{y}^{t_n}) \rightarrow \tilde{\boldsymbol{\alpha}} \neq \boldsymbol{\alpha}(\mathbf{x}, \mathbf{y}), \quad (\text{B.98})$$

as $n \rightarrow \infty$. Since $\boldsymbol{\alpha}(\mathbf{x}, \mathbf{y})$ is the unique maximizer of $H(\boldsymbol{\alpha}; \mathbf{x}, \mathbf{y})$ and $\tilde{\boldsymbol{\alpha}} \neq \boldsymbol{\alpha}(\mathbf{x}, \mathbf{y})$, it can be seen that

$$H(\tilde{\boldsymbol{\alpha}}; \mathbf{x}, \mathbf{y}) < H(\boldsymbol{\alpha}(\mathbf{x}, \mathbf{y}); \mathbf{x}, \mathbf{y}). \quad (\text{B.99})$$

On the other hand, since $\boldsymbol{\alpha}(\mathbf{x}^{t_n}, \mathbf{y}^{t_n})$ is the unique maximizer of $H(\boldsymbol{\alpha}; \mathbf{x}^{t_n}, \mathbf{y}^{t_n})$,

$$H(\boldsymbol{\alpha}(\mathbf{x}^{t_n}, \mathbf{y}^{t_n}); \mathbf{x}^{t_n}, \mathbf{y}^{t_n}) \geq H(\boldsymbol{\alpha}(\mathbf{x}, \mathbf{y}); \mathbf{x}^{t_n}, \mathbf{y}^{t_n}). \quad (\text{B.100})$$

Note that $H(\boldsymbol{\alpha}; \mathbf{x}, \mathbf{y})$ is continuous in $\boldsymbol{\alpha}$ and (\mathbf{x}, \mathbf{y}) because it is minimum of the continuous functions, $\hat{d}_{ij}(\boldsymbol{\alpha}; \boldsymbol{\mu}, \boldsymbol{\sigma})$. Since $\boldsymbol{\alpha}(\mathbf{x}^{t_n}, \mathbf{y}^{t_n}) \rightarrow \tilde{\boldsymbol{\alpha}}$, $\mathbf{x}^{t_n} \rightarrow \mathbf{x}$, and $\mathbf{y}^{t_n} \rightarrow \mathbf{y}$, taking $n \rightarrow \infty$ on both sides of (B.100), we obtain that

$$H(\tilde{\boldsymbol{\alpha}}; \mathbf{x}, \mathbf{y}) \geq H(\boldsymbol{\alpha}(\mathbf{x}, \mathbf{y}); \mathbf{x}, \mathbf{y}), \quad (\text{B.101})$$

which contradicts (B.99). Therefore, $\boldsymbol{\alpha}(\mathbf{x}^t, \mathbf{y}^t) \rightarrow \boldsymbol{\alpha}(\mathbf{x}, \mathbf{y})$ as $t \rightarrow \infty$ and the proof is complete. \square

Proof of Lemma B.3. We first show weak convergence. From Proposition 1 we have that $\hat{\xi}_{jt} \rightarrow \xi_j$ in probability as $t \rightarrow \infty$. Further, observe that $\{X_{jP_n}\}$ is a sequence of independent random variables, where a random variable P_n is define as $N_{jP_n} = n$; see (B.1) in the proof of Proposition 1. Since \hat{f}_{jt} is the density estimator of the independent sequence of random variables $\{X_{jP_n}\}$, we may apply Parzen (1962) and obtain that $\hat{f}_{jt}(x) \rightarrow f(x)$ almost surely as $t \rightarrow \infty$ for any $x \in \mathcal{H}_j$. Using the continuous mapping theorem, we establish the weak convergence of $\hat{f}_{jt}(\hat{\xi}_{jt})$ to $f_j(\xi_j)$.

To show strong convergence, it is easy to extend Proposition 1 to see that $\hat{\xi}_{jt} \rightarrow \xi_j$ almost surely if $N_{jt} \rightarrow \infty$ almost surely as $t \rightarrow \infty$, which follows from the independence argument above and p. 75 of Silverman (1986). Further, from Theorem 1 of Nadaraya (1965) we have that $\hat{f}_{jt}(x) \rightarrow f_j(x)$ almost surely as $t \rightarrow \infty$ under the additional conditions on the bandwidth parameter and the kernel. Combined with the continuous mapping theorem, the proof of the lemma is complete. \square

Proof of Lemma B.4. Without loss of generality, let $1 = \arg \min_j \{x_j\}$ and $2 = \arg \min_j \{y_j\}$. Then,

$$\frac{\min_j \{x_j\}}{\min_j \{y_j\}} = \frac{x_1}{y_2}. \quad (\text{B.102})$$

By definition, observe that

$$\frac{x_1}{y_2} \leq \frac{x_1}{y_1} \leq 1 + e \text{ and } \frac{x_1}{y_2} \geq \frac{x_2}{y_2} \geq 1 - e, \quad (\text{B.103})$$

which completes the proof of the lemma. \square

Proof of Lemma B.5. First, suppose $N_{jt} \rightarrow \infty$ in probability as $t \rightarrow \infty$. Then, we have that $\hat{\xi}_{jt} \rightarrow \xi_j$ in probability from Proposition 1. Also, from the same argument as in the proof for Proposition 1, it can be seen that $\hat{F}_{jt}(x)$ is an average of indicator random variables that are independent and identically distributed. Therefore, by Glivenko-Cantelli theorem, it follows that $\sup_{x \in \mathcal{X}_j} |\hat{F}_{jt}(x) - F_j(x)| \rightarrow 0$ in probability as $t \rightarrow \infty$. Consequently, from continuous mapping theorem, we have that $\hat{F}_{jt}(\hat{\xi}_{jt}) \rightarrow F_j(\xi_j)$ in probability as $t \rightarrow \infty$ for each j , from which the desired conclusion follows. The other case where $N_{jt} \rightarrow \infty$ almost surely as $t \rightarrow \infty$ follows immediately from the same steps as above, which will be omitted. \square

Proof of Lemma B.6. Note that N_{jt} under the *static* policy $\pi(\alpha)$ is a deterministic value for each j and t , satisfying $N_{jt}/t \rightarrow \alpha_j$ as $t \rightarrow \infty$. We first show the upper bound:

$$\limsup_{t \rightarrow \infty} \frac{1}{t\delta_t^2} \log \mathbb{P}(|\hat{\xi}_{jt} - \xi_j| > \delta_t) \leq -\frac{\alpha_j f_j^2(\xi_j)}{2p(1-p)}. \quad (\text{B.104})$$

Observe that, for $\theta < 0$,

$$\begin{aligned} \mathbb{P}(\hat{\xi}_{jt} > \xi_j + \delta_t) &= \mathbb{P}\left(\sum_{n=1}^{N_{jt}} \mathbf{I}\{X_{j\tau_{jn}} \leq \xi_j + \delta_t\} < pN_{jt}\right) \\ &\leq \exp(-\theta p N_{jt}) \mathbb{E}\left(\exp\left(\theta \sum_{n=1}^{N_{jt}} \mathbf{I}\{X_{j\tau_{jn}} \leq \xi_j + \delta_t\}\right)\right) \\ &= (\mathbb{E}(\exp(\theta(\mathbf{I}\{X_j \leq \xi_j + \delta_t\} - p))))^{N_{jt}}, \end{aligned} \quad (\text{B.105})$$

where the inequality follows from Chernoff's inequality. Further observe that

$$\begin{aligned} \frac{1}{t\delta_t^2} \log \mathbb{P}(\hat{\xi}_{jt} > \xi_j + \delta_t) &\leq -\frac{\alpha_j}{\delta_t^2} \sup_{\theta < 0} \{\theta p - \log \mathbb{E}(\exp(\theta \mathbf{I}\{X_j \leq \xi_j + \delta_t\}))\} \\ &= -\frac{\alpha_j}{\delta_t^2} \left(p \log\left(\frac{p}{F_j(\xi_j + \delta_t)}\right) + (1-p) \log\left(\frac{1-p}{1-F_j(\xi_j + \delta_t)}\right) \right), \end{aligned} \quad (\text{B.106})$$

where the inequality follows from (B.105) and the fact that the inequality holds for every $\theta < 0$; the equality follows from the fact that $\mathbf{I}\{X_j \leq \xi_j + \delta_t\}$ is a Bernoulli random variable with success probability $F_j(\xi_j + \delta_t)$. Now, from the second order Taylor expansion, we have that

$$\begin{aligned} \log(F_j(\xi_j + \delta_t)) &= \log(p) - \delta_t \frac{f_j(\xi_j)}{p} + \frac{\delta_t^2}{2} \left(\frac{f_j^2(\xi_j)}{p^2} - \frac{f_j'(\xi_j)}{p} \right) + o(\delta_t^2) \\ \log(1 - F_j(\xi_j + \delta_t)) &= \log(1-p) + \delta_t \frac{f_j(\xi_j)}{1-p} + \frac{\delta_t^2}{2} \left(\frac{f_j^2(\xi_j)}{(1-p)^2} - \frac{f_j'(\xi_j)}{(1-p)} \right) + o(\delta_t^2). \end{aligned} \quad (\text{B.107})$$

Substituting (B.107) into (B.106) and taking \limsup on both sides of the inequality, we establish that

$$\limsup_{t \rightarrow \infty} \frac{1}{t\delta_t^2} \log \mathbb{P}(\hat{\xi}_{jt} > \xi_j + \delta_t) \leq -\frac{\alpha_j f_j^2(\xi_j)}{2p(1-p)}. \quad (\text{B.108})$$

Similarly, one can easily show that

$$\limsup_{t \rightarrow \infty} \frac{1}{t\delta_t^2} \log \mathbb{P}(\hat{\xi}_{jt} < \xi_j - \delta_t) \leq -\frac{\alpha_j f_j^2(\xi_j)}{2p(1-p)}, \quad (\text{B.109})$$

and we complete the proof for the upper bound, (B.104).

Next, we show the lower bound:

$$\liminf_{t \rightarrow \infty} \frac{1}{t\delta_t^2} \log \mathbb{P}(|\hat{\xi}_{jt} - \xi_j| > \delta_t) \geq -\frac{\alpha_j f_j^2(\xi_j)}{2p(1-p)}. \quad (\text{B.110})$$

The lower bound can be shown using a standard argument using change of measure, which we spell out for completeness. To this end, define $Y_{j\tau}^t = \mathbf{I}\{X_{j\tau} \leq \xi_j + \delta_t\}$ for $\tau \leq t$ and let $Y_j^t = \mathbf{I}\{X_j \leq \xi_j + \delta_t\}$ be identically distributed with $Y_{j\tau}^t$, $\tau \leq t$. Also define $M_{jt}(\theta) = \mathbf{E}(\exp(\theta Y_j^t))$, $q_{jt} = F_j(\xi_j + \delta_t)$, and $q_{jt}(\theta) = F_j(\xi_j + \delta_t) \exp(\theta) / M_{jt}(\theta)$. If we let $\theta_t = \arg \max_{\theta \in \mathcal{R}} \{\theta p - \log M_{jt}(\theta)\}$, then it can be easily seen that $\theta_t > 0$ is well defined and $q_{jt}(\theta_t) = p$. For any function $g: \mathcal{R}^t \rightarrow \mathcal{R}$ with $\mathbf{E}|g(Y_{j1}^t, \dots, Y_{jt}^t)| < \infty$, $\mathbf{E}_\theta(g(Y_{j1}^t, \dots, Y_{jt}^t))$ denotes the expectation with respect to the measure $q_{jt}(\theta)$ for which $Y_{jt}^t = 1$ with probability $q_{jt}(\theta)$ and 0 with probability $1 - q_{jt}(\theta)$. Observe that

$$\begin{aligned} \mathbf{P}(\hat{\xi}_{jt} > \xi_j + \delta_t) &= \mathbf{P}\left(\sum_{n=1}^{N_{jt}} Y_{j\tau_{jn}} > pN_{jt}\right) \\ &= \mathbf{E}\left(\mathbf{I}\left\{\sum_{n=1}^{N_{jt}} Y_{j\tau_{jn}} > pN_{jt}\right\}\right) \\ &= \mathbf{E}_{\theta_t}\left(\mathbf{I}\left\{\sum_{n=1}^{N_{jt}} Y_{j\tau_{jn}} > pN_{jt}\right\} \prod_{n=1}^{N_{jt}} \left(\frac{q_{jt}}{p}\right)^{Y_{j\tau_{jn}}^t} \left(\frac{1-q_{jt}}{1-p}\right)^{(1-Y_{j\tau_{jn}}^t)}\right) \end{aligned} \quad (\text{B.111})$$

by the change of measure and the fact that $q_{jt}(\theta_t) = p$ by definition of θ_t . After some straightforward algebra, we have that

$$\begin{aligned} &\prod_{n=1}^{N_{jt}} \left(\frac{q_{jt}}{p}\right)^{Y_{j\tau_{jn}}^t} \left(\frac{1-q_{jt}}{1-p}\right)^{(1-Y_{j\tau_{jn}}^t)} \\ &= \exp\left(-\log\left(\frac{p}{q_{jt}} \sum_{n=1}^{N_{jt}} Y_{j\tau_{jn}}^t\right) - \log\left(\frac{1-p}{1-q_{jt}} \sum_{n=1}^{N_{jt}} (1-Y_{j\tau_{jn}}^t)\right)\right) \\ &= \exp\left(-N_{jt} \left(p \log\left(\frac{p}{q_{jt}}\right) + (1-p) \log\left(\frac{1-p}{1-q_{jt}}\right)\right)\right) \exp\left(-\beta_1 \sum_{n=1}^{N_{jt}} (Y_{j\tau_{jn}}^t - p)\right), \end{aligned} \quad (\text{B.112})$$

where $\beta_1 = \log(p/q_{jt}) + \log((1-p)/(1-q_{jt}))$. Therefore, rewriting (B.111),

$$\begin{aligned} &\mathbf{P}(\hat{\xi}_{jt} > \xi_j + \delta_t) \\ &= \exp\left(-N_{jt} \left(p \log\left(\frac{p}{q_{jt}}\right) + (1-p) \log\left(\frac{1-p}{1-q_{jt}}\right)\right)\right) \times \\ &\quad \mathbf{E}_{\theta_t}\left(\mathbf{I}\left\{\sum_{n=1}^{N_{jt}} Y_{j\tau_{jn}} > pN_{jt}\right\} \exp\left(-\beta_1 \sum_{n=1}^{N_{jt}} (Y_{j\tau_{jn}}^t - p)\right)\right), \end{aligned} \quad (\text{B.113})$$

Further, for some $\beta_2 > 0$, the expectation on the right-hand side of (B.113) can be bounded as

$$\begin{aligned} &\mathbf{E}_{\theta_t}\left(\mathbf{I}\left\{\sum_{n=1}^{N_{jt}} Y_{j\tau_{jn}} > pN_{jt}\right\} \exp\left(-\beta_1 \sum_{n=1}^{N_{jt}} (Y_{j\tau_{jn}}^t - p)\right)\right) \\ &\geq \mathbf{E}_{\theta_t}\left(\mathbf{I}\left\{pN_{jt} < \sum_{n=1}^{N_{jt}} Y_{j\tau_{jn}} < pN_{jt} + \beta_2 \sqrt{N_{jt}}\right\} \exp\left(-\beta_1 \beta_2 \sqrt{N_{jt}}\right)\right) \\ &= \exp\left(-\beta_1 \beta_2 \sqrt{N_{jt}}\right) \mathbf{P}_{\theta_t}\left(pN_{jt} < \sum_{n=1}^{N_{jt}} Y_{j\tau_{jn}} < pN_{jt} + \beta_1 \beta_2 \sqrt{N_{jt}}\right), \end{aligned} \quad (\text{B.114})$$

where the inequality follows from the fact that

$$\mathbf{I}\left\{\sum_{n=1}^{N_{jt}} Y_{j\tau_{jn}} > pN_{jt}\right\} \geq \mathbf{I}\left\{pN_{jt} < \sum_{n=1}^{N_{jt}} Y_{j\tau_{jn}} < pN_{jt} + \beta_2 \sqrt{N_{jt}}\right\}. \quad (\text{B.115})$$

Note that

$$\begin{aligned}
 & \mathbb{P}_{\theta_t} \left(pN_{jt} < \sum_{n=1}^{N_{jt}} Y_{j\tau_{jn}} < pN_{jt} + \beta_1\beta_2\sqrt{N_{jt}} \right) \\
 &= \mathbb{P}_{\theta_t} \left(0 < \frac{1}{\sqrt{N_{jt}}} \sum_{n=1}^{N_{jt}} (Y_{j\tau_{jn}} - p) < \beta_1\beta_2 \right) \\
 &\rightarrow c \in (0, 1),
 \end{aligned} \tag{B.116}$$

where the limit follows since $Y_{j\tau}$ with $\tau \leq t$ is a Bernoulli random variable with success probability p under $\mathbb{P}_{\theta_t}(\cdot)$, and hence, the distribution of the normalized sum $\sum_{n=1}^{N_{jt}} (Y_{j\tau_{jn}} - p)/\sqrt{N_{jt}}$ tends toward a standard normal distribution by the central limit theorem. Combining (B.113), (B.114), (B.116), and (B.107), we establish that

$$\liminf_{t \rightarrow \infty} \frac{1}{t\delta_t^2} \log \mathbb{P}(\hat{\xi}_{jt} > \xi_j + \delta_t) \geq -\frac{\alpha_j f_j^2(\xi_j)}{2p(1-p)}. \tag{B.117}$$

Likewise, one can easily obtain that

$$\liminf_{t \rightarrow \infty} \frac{1}{t\delta_t^2} \log \mathbb{P}(\hat{\xi}_{jt} < \xi_j - \delta_t) \geq -\frac{\alpha_j f_j^2(\xi_j)}{2p(1-p)}, \tag{B.118}$$

which completes the proof for the lower bound (B.110). Combining (B.104) and (B.110), the proof for the lemma is complete. \square

Appendix C: Parameter Sensitivity of the Algorithms

The algorithms proposed in this paper commonly take the number of initial samples n and the batch size m as parameters. These slightly affect the performance of the proposed algorithms, but we remark that the qualitative conclusions from §6.2 are robust relative to the choice of these parameters. In this section, we examine the sensitivity of performance with respect to parameters that are specific to each algorithm.

First, the naive algorithm takes c as an additional parameter, which ensures that each system is sampled at least at a logarithmic rate. To see how the performance of the naive algorithm is affected by the choice of c , we consider a two-system configuration with binomial distributions; the number of trials is 10 for both systems and the success probabilities are $\boldsymbol{\mu} = (0.62, 0.55)$, respectively. We set $p = 0.1$. In this case, the vector $\boldsymbol{\alpha}^* = (0.76, 0.24)$ maximizes the rate function. We vary $c = 0, 20$.

Figure 1(a) shows a sample path of α_{1t} under the naive algorithm with $c = 0$, in which system 1 is not sampled after $n_0 = 10$ initial samples. The relative frequency of α_{1t} for $t = 4000$ over many paths is illustrated in Figure 1(b) using 10^4 simulation replications. As noted in Remark 2, the naive algorithm may not sample a particular system for a long period without forced sampling, potentially increasing the likelihood of selecting a non-best system. Figure 1(c) presents a typical sample path α_{1t} under the naive algorithm with forced sampling and Figure 1(d) gives the relative frequency chart of α_{1t} for $t = 4000$.

Next, the QD-C and AQD-C algorithms take the kernel $K(\cdot)$ and $h(\cdot)$ as parameters. The optimal choice of the bandwidth parameter is $h(n) = dn^{-1/5}$, where n is the number of observations and the constant d depends on the true density function that is a priori unknown (Silverman 1986). In order to show the sensitivity of these algorithms with respect to the kernel and the bandwidth parameter, we use normal and triangular kernels and four bandwidth functions; $h(n) = 1.06\sigma n^{-1/5}$, where σ is replaced with sample




Figure 1 The frequencies of α_{1t} under the naive algorithms for different values of parameter c . In all cases, we consider two systems with binomial distributions parameterized by the number of trials $n = (10, 10)$ and the success probability $\mu = (0.62, 0.55)$. We set $p = 0.1$, in which case $\alpha_1^* = 0.76$. We use $n_0 = 10$ and $m = 10$. Panel (a) shows a sample path of the naive algorithm with $c = 0$, where the system 1 is not sampled over the long horizon. Panel (c) shows a sample path of the naive algorithm with $c = 20$ using the same sequence of random samples. Panels (b) and (d) give the relative frequency of α_{1t} at $t = 4000$ using 10^4 trials of the naive algorithm.

standard deviation, and $h(n) = h = 0.1, 1, 10$. We use the configuration with $k = 4$ normally distributed systems with $\mu = (0, 0, 0, 0)$ and $\sigma = (1, 1.1, 1.2, 1.3)$. We let $p = 0.1$.

Figures 2(a)-(b) illustrate the performance of the QD-C and AQD-C algorithms in terms of the probability of false selection when we use the normal kernel. It can be seen that the performance of the two algorithms is degraded when the bandwidth parameter is $h = 0.1$ (too little smoothing) or $h = 10$ (too much smoothing). The performance is robust when we use the adaptive bandwidth $h(n) = 1.06\sigma n^{-1/5}$. Figures 2(c)-(d) show the same experiments with the triangular kernel, in which similar conclusions hold.

References

- Knight K (1999) *Mathematical Statistics* (Florida: Chapman and Hall/CRC).
- Nadaraya EA (1965) On non-parametric estimates of density functions and regression curves. *Theory of Prob. & Its Applications* 10(1):186–190.

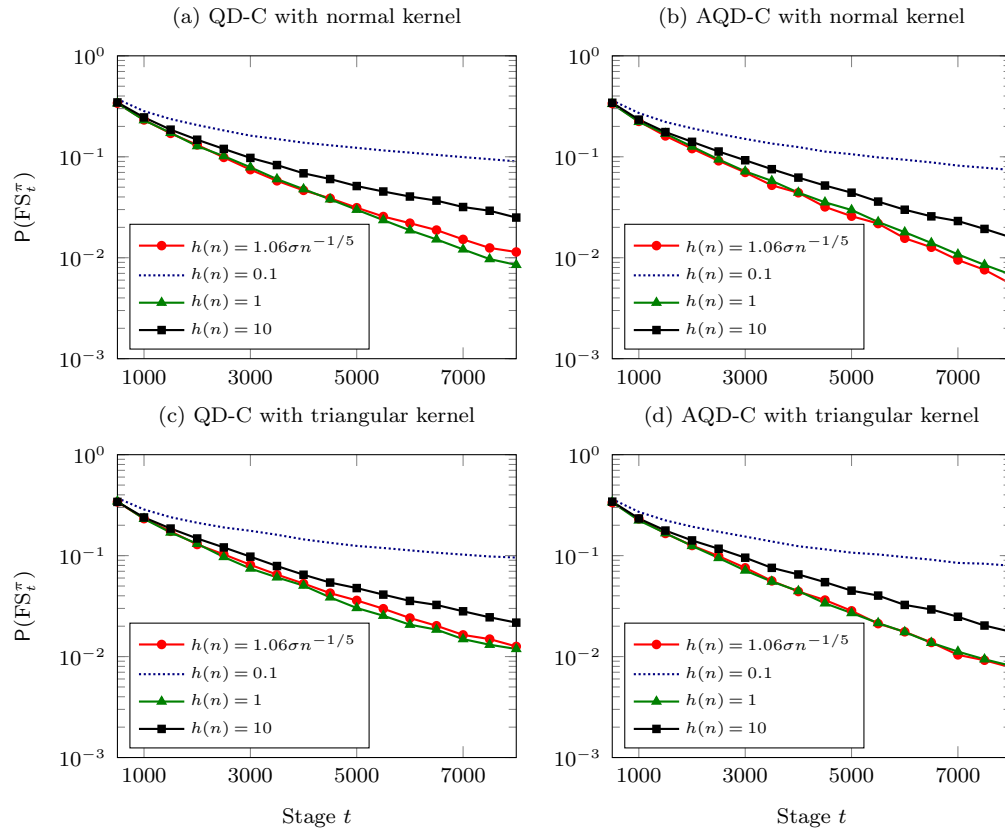


Figure 2 $P(\text{FS}_t^r)$ as a function of stage t on a log-linear scale for different kernels and bandwidth parameters. We test the QD-C and AQD-C algorithms in the case with four normally distributed systems; $\mu = (0, 0, 0, 0)$ and $\sigma = (1, 1.1, 1.2, 1.3)$. We let $p = 0.1$ and $(n_0, m) = (20, 10)$ for both algorithms. The performance of the two algorithms is not sensitive with respect to the choice of the kernel, but it degrades when the bandwidth parameter $h(\cdot)$ is too small or too large.

Parzen E (1962) On estimation of a probability density function and mode. *Annals of Math. Statist.* 33:1065–1076.

Serfling RJ (2009) *Approximation theorems of mathematical statistics* (New York: John Wiley & Sons).

Silverman BW (1986) *Density estimation for statistics and data analysis* (New York: CRC press).