

Most Correlated Arms Identification

Che-Yu Liu, Sébastien Bubeck

Department of Operations Research and Financial Engineering,
Princeton University

cheliu@princeton.edu, sbubeck@princeton.edu

April 23, 2014

Abstract

We study the problem of finding the most mutually correlated arms among many arms. We show that adaptive arms sampling strategies can have significant advantages over the non-adaptive uniform sampling strategy. Our proposed algorithms rely on a novel correlation estimator. The use of this accurate estimator allows us to get improved results for a wide range of problem instances.

1 Introduction

We define the *most correlated arms identification* problem as follows. Let K be an integer and $X^t = (X_{1,t}, \dots, X_{K,t})^\top$, $t = 1, 2, \dots$ be a sequence of independent, identically distributed Gaussian random vectors with zero mean and an *unknown* covariance matrix $\Sigma = (\sigma_{ij})_{i,j \in [K]}$. We use \mathbb{P}_Σ to denote the corresponding probability measure (on the natural probability space for this model) and make the following assumptions on the entries of Σ : $\forall i, j \in [K]$, $\sigma_{ii} = 1$ and $\sigma_{ij} \geq 0$. Under these assumptions, Σ can also be seen as the correlation matrix. Consider now an agent facing K arms. At each time step $t = 1, 2, \dots$, the agent selects a subset of arms $A_t \subset [K]$ and observes the corresponding values in X^t , denoted by $X_{A_t}^t = \{X_{i,t}, i \in A_t\}$. The agent is allowed to employ an *adaptive sampling strategy*, that is, the selection of A_t may depend on the past observations $H_{t-1} = \{A_s, X_{A_s}^s\}_{s \in [t-1]}$. This is in contrast to the non-adaptive *uniform sampling* strategy, where A_t is always chosen to be $[K]$. The task of the agent is to return as rapidly as possible a subset of h most mutually correlated arms. More precisely, the agent is interested in finding

$$S^* \in \operatorname{argmax}_{S \subset [K], |S|=h} \sum_{j, \ell \in S, j \neq \ell} \sigma_{j, \ell}.$$

Throughout this paper, we assume that there is a unique solution S^* to the above problem and we call this unique solution the *optimal subset*. Furthermore an arm $i \in [K]$ is called *optimal* if it belongs to S^* and is called *suboptimal* otherwise.

The goal of this paper is to devise algorithms that output reliably S^* while using as few *samples* as possible. Throughout this paper, the word *sample* is reserved for the realized value of a real-valued random variable. In comparison, the sampling outcome of a multi-dimensional random vector is called a *sample vector*. Following the tradition of the adaptive-exploration literature, we study our problem in two different settings.

Fixed-Budget: In the fixed-budget setting, we are given a fixed budget of n samples and are asked to output a subset of arms \hat{S} with size h as soon as the sampling budget has been used up. The task here is to devise a sampling and decision strategy that achieves a small probability of returning a wrong subset, that is, one wants to minimize $\mathbb{P}_\Sigma(\hat{S} \neq S^*)$.

Fixed-Confidence: In the fixed-confidence setting, we are given a fixed confidence level $\delta > 0$ and there is no explicit budget limitation. We look for a sampling, stopping and decision strategy that returns S^* with probability at least $1 - \delta$ when it stops, regardless of the underlying correlation matrix Σ . The performance of the strategy is evaluated by the number of samples obtained before it terminates, either in expectation or with high probability.

In Section 4 and 5, we present two algorithms SR-C and SE-C that are designed for the above two settings respectively. The main innovation in our algorithms is the use of an accurate correlation estimator, called difference-based correlation estimator. In Section 2, we define and analyze this new estimator and argue that it can substantiate in an optimal way the intuition that the estimation task becomes easier when the correlation is close to 1, a feature that the classical correlation estimator can not attain. Inspired by the statistical properties of the difference-based correlation estimator, we define the suboptimality ratio of B with respect to A (for the correlation matrix Σ) as

$$\mathcal{D}_\Sigma(A, B) = \frac{\sum_{(j,\ell) \in B^2 \setminus (A \cap B)^2, j \neq \ell} (1 - \sigma_{j,\ell})}{\sum_{(j,\ell) \in A^2 \setminus (A \cap B)^2, j \neq \ell} (1 - \sigma_{j,\ell})}$$

for two subsets of arms A and B with size h and with the convention that $\frac{0}{0} = 1$. This ratio is a measure of how uncorrelated the arms in B are compared to the arms in A . Note that $\mathcal{D}_\Sigma(A, B) \geq 1$ if only if $\sum_{j,\ell \in A, j \neq \ell} \sigma_{j,\ell} \geq \sum_{j,\ell \in B, j \neq \ell} \sigma_{j,\ell}$. Now for an arm $i \in [K]$, denote by $R_{i,\Sigma}$ the suboptimality ratio of the arm i with respect to the optimal set S^* , defined as

$$R_{i,\Sigma} = \min_{i \in B \subset [K], |B|=h} \mathcal{D}_\Sigma(S^*, B).$$

It is clear that $R_{i,\Sigma} = 1$ if arm i is optimal and $R_{i,\Sigma} > 1$ otherwise. Therefore, the order statistics of $R_{i,\Sigma}$ satisfies

$$1 = R_{(1),\Sigma} = \dots = R_{(h),\Sigma} < R_{(h+1),\Sigma} \leq \dots \leq R_{(K),\Sigma}.$$

These suboptimality ratios determine the number of times the suboptimal arms need to be drawn in our algorithms. More precisely, define the α function as $\alpha(\theta) = \frac{1}{2}(\log \theta - 1 + \frac{1}{\theta})$, for $\theta \geq 1$. We consider in Section 3 the non-adaptive sampling setting where we draw each arm an equal number of times and show that we need $\tilde{\Theta}\left(\frac{K}{\alpha(R_{(h+1),\Sigma})} + K\right)$ samples¹ to reliably identify S^* by providing a matching upper and lower bound. Furthermore, we prove in Section 4 and 5 that SR-C and SE-C

¹The notation $\tilde{\Theta}, \tilde{\mathcal{O}}$ hides a logarithmic factor.

need at most $\tilde{O}(H_C + K)$ samples to find S^* where H_C is defined as

$$H_C = \frac{h}{\alpha(R_{(h+1),\Sigma})} + \sum_{i=h+1}^K \frac{1}{\alpha(R_{(i),\Sigma})}.$$

Note that H_C can be much smaller than $\frac{K}{\alpha(R_{(h+1),\Sigma})}$ for a wide range of Σ , especially when the entries of Σ are inhomogeneous. This shows that adaptive sampling strategies could have significant advantages over non-adaptive uniform sampling. We close the paper with a discussion in Section 6 of possible improvements and challenges.

Remark: Note that the function $\alpha : [1, +\infty) \rightarrow \mathbb{R}$ defined above is a positive, strictly increasing function. Moreover, $\alpha(\theta) = \Theta((1 - \theta)^2)$ when $\theta \rightarrow 1$ and $\alpha(\theta) = \Theta(\log(\theta))$ when $\theta \rightarrow +\infty$. It follows from these facts that for any $q > 0$, there exists two constants $c_1(q), c_2(q) > 0$ that only depend on q such that $c_1(q)\alpha(\theta) \leq \alpha(\theta^q) \leq c_2(q)\alpha(\theta)$ for any $\theta \geq 1$.

Related Work

The problem we study in this paper is similar in spirit to the *best arm identification problem* in multi-armed bandits. In the latter problem, an agent repeatedly selects an arm and observes a sample reward drawn from the arm’s reward distribution, and then he is asked to return the arm with the highest mean reward. The Successive Elimination algorithm in [Even-Dar et al. \[2006\]](#) was shown to find the single best arm i^* with high probability with $\tilde{O}\left(\sum_{i \neq i^*} \frac{1}{\Delta_i}\right)$ samples where Δ_i is the gap between the mean reward of the best arm and that of a suboptimal arm i . [Bubeck et al. \[2009\]](#) and [Audibert et al. \[2010\]](#) study the same problem under the fixed-budget setting, in particular the Successive Rejects procedure of [Audibert et al. \[2010\]](#) was shown to require essentially as many samples as Successive Elimination. The algorithms that we propose in Sections 4 and 5 for most correlated arms identification are inspired by these algorithms for best arm identification. In particular they share the same high-level idea that the more suboptimal an arm is with respect to the best arms, the less samples we need to distinguish it from the best arms. In the fixed-budget setting, this is done by distributing the sample budget to the arms in an adaptive way based on their correlation estimates. For the fixed-confidence setting we build confidence intervals on correlations among arms and uniformly sample all the arms until we have enough confidence to identify and exclude some suboptimal arms.

Our work is different from [Arias-Castro et al. \[2012a\]](#), [Arias-Castro et al. \[2012b\]](#) and [Castro et al. \[2013\]](#) where the task is the detection of the presence of a sparse and correlated subset of components from samples of a high-dimensional Gaussian distribution. Their work focuses on determining whether there is a correlated subset while the task in this paper is to accurately identify the subset with largest mutual correlations. Perhaps more importantly, on the contrary to these work, our algorithms adapt to the potential heterogeneity in the correlation matrix Σ .

2 Correlation Estimators

In this section, we describe a new correlation estimator, called difference-based correlation estimator. It is intuitive that the task of correlation estimation becomes easier when the correlation is

close to 1. We are going to see that the classical correlation estimator fails to capture this intuition while the difference-based correlation estimator can quantify it in an optimal way. The use of this novel estimator will allow us to largely reduce the sample complexity of our algorithms for a wide range of correlation matrices Σ .

2.1 Classical Correlation Estimator

For $j \neq \ell \in [K]$, consider the classical correlation estimator defined as

$$\begin{aligned}\tilde{\sigma}_{j,\ell,t} &= \frac{1}{t} \sum_{s=1}^t X_{j,s} X_{\ell,s} = \frac{1}{t} \left(\sum_{s=1}^t \frac{(X_{j,s} + X_{\ell,s})^2}{4} - \sum_{s=1}^t \frac{(X_{j,s} - X_{\ell,s})^2}{4} \right) \\ &= \sigma_{j\ell} + \left(\frac{1 + \sigma_{j\ell}}{2} \left(\frac{1}{t} \sum_{s=1}^t \frac{(X_{j,s} + X_{\ell,s})^2}{2(1 + \sigma_{j\ell})} - 1 \right) - \frac{1 - \sigma_{j\ell}}{2} \left(\frac{1}{t} \sum_{s=1}^t \frac{(X_{j,s} - X_{\ell,s})^2}{2(1 - \sigma_{j\ell})} - 1 \right) \right).\end{aligned}$$

By observing that $\sum_{s=1}^t \frac{(X_{j,s} + X_{\ell,s})^2}{2(1 + \sigma_{j\ell})}$ and $\sum_{s=1}^t \frac{(X_{j,s} - X_{\ell,s})^2}{2(1 - \sigma_{j\ell})}$ are independent, chi-square distributed with t degrees of freedom, we see that the magnitude of the fluctuation of $\tilde{\sigma}_{j,\ell,t}$ around $\sigma_{j\ell}$ only has a very mild dependency on $\sigma_{j\ell}$. Its estimation accuracy improves only slightly even when $\sigma_{j\ell} = 1$.

2.2 Difference-Based Correlation Estimator

Our main idea for getting a good correlation estimator is to draw inspiration from the likelihood ratio tests. Assume that $K = 2$ and consider the following testing problem.

$$(T) \quad \begin{cases} H_0 : \Sigma = \Sigma_0 \\ H_1 : \Sigma = \Sigma_1 \end{cases}$$

where $\Sigma_0 = \begin{pmatrix} 1 & \rho_0 \\ \rho_0 & 1 \end{pmatrix}$ and $\Sigma_1 = \begin{pmatrix} 1 & \rho_1 \\ \rho_1 & 1 \end{pmatrix}$ with $1 > \rho_0 > \rho_1 \geq 0$. Since the likelihood ratio test is an optimal test, it must be able to distinguish H_0 and H_1 with high accuracy in the case when ρ_0 is close to 1 while ρ_1 is bounded away from 1. This suggests that the study of the likelihood ratio test may be helpful in constructing a good correlation estimator. Denote by f_Q the probability density function of a probability distribution Q . It is easy to verify that the likelihood ratio statistics of the testing problem (T) can be written as

$$\Lambda((X^s)_{s=1,\dots,t}) = \frac{f_{\mathcal{N}(0, (1+\rho_0)I_t)} \left(\left(\frac{X_{1,s} + X_{2,s}}{\sqrt{2}} \right)_{s=1,\dots,t} \right)}{f_{\mathcal{N}(0, (1+\rho_1)I_t)} \left(\left(\frac{X_{1,s} + X_{2,s}}{\sqrt{2}} \right)_{s=1,\dots,t} \right)} \cdot \frac{f_{\mathcal{N}(0, (1-\rho_0)I_t)} \left(\left(\frac{X_{1,s} - X_{2,s}}{\sqrt{2}} \right)_{s=1,\dots,t} \right)}{f_{\mathcal{N}(0, (1-\rho_1)I_t)} \left(\left(\frac{X_{1,s} - X_{2,s}}{\sqrt{2}} \right)_{s=1,\dots,t} \right)}.$$

The two fractions above can be seen as the likelihood ratio statistics of the following two testing problems.

$$(T_1) \quad \begin{cases} H_0 : \frac{X_{1,s} + X_{2,s}}{\sqrt{2}} \sim \mathcal{N}(0, 1 + \rho_0) \\ H_1 : \frac{X_{1,s} + X_{2,s}}{\sqrt{2}} \sim \mathcal{N}(0, 1 + \rho_1) \end{cases}, \quad (T_2) \quad \begin{cases} H_0 : \frac{X_{1,s} - X_{2,s}}{\sqrt{2}} \sim \mathcal{N}(0, 1 - \rho_0) \\ H_1 : \frac{X_{1,s} - X_{2,s}}{\sqrt{2}} \sim \mathcal{N}(0, 1 - \rho_1) \end{cases}.$$

Now observe that the testing problem (T_2) is always easier than the problem (T_1) since we always have $\frac{1+\rho_0}{1+\rho_1} \leq \frac{1-\rho_1}{1-\rho_0}$. This suggests that the hardness of the problem stays roughly the same if we replace the original problem (T) by problem (T_2) . Moreover, for problem (T_2) , it is natural to consider tests based on the test statistics $\frac{1}{t} \sum_{s=1}^t \frac{(X_{1,s} - X_{2,s})^2}{2}$. This leads us to define the difference-based correlation estimator for $\sigma_{j\ell}$ with $j \neq \ell \in [K]$ as

$$\hat{\sigma}_{j,\ell,t} = 1 - \frac{1}{t} \sum_{s=1}^t \frac{(X_{j,s} - X_{\ell,s})^2}{2}.$$

To analyze its statistical properties, we further define the random variable $Y_{j,\ell,t}$ as

$$Y_{j,\ell,t} = \frac{1}{1 - \sigma_{j\ell}} \sum_{s=1}^t \frac{(X_{j,s} - X_{\ell,s})^2}{2}$$

with the convention that $\frac{0}{0} = 1$. Then $Y_{j,\ell,t}$ follows a chi-square distribution with t degrees of freedom and we have the following relations

$$1 - \hat{\sigma}_{j,\ell,t} = (1 - \sigma_{j\ell}) \frac{Y_{j,\ell,t}}{t} \text{ and } \hat{\sigma}_{j,\ell,t} = \sigma_{j\ell} - (1 - \sigma_{j\ell}) \left(\frac{Y_{j,\ell,t}}{t} - 1 \right).$$

For a fixed number of samples, the deviation of the difference-based correlation estimator from the true correlation is proportional to $1 - \sigma_{j\ell}$. In other words, the accuracy of estimation increases when $\sigma_{j\ell}$ approaches 1.

2.3 Optimality of the Difference-Based Correlation Estimator

To illustrate the strength of the difference-based correlation estimator, we return to the testing problem (T) . The testing accuracy of a test φ is measured by its maximal risk, defined as

$$\mathcal{R}_t(\varphi) = \max(\mathbb{P}_{\Sigma_0}(\varphi = 1), \mathbb{P}_{\Sigma_1}(\varphi = 0)).$$

Now, define $R = \frac{1-\rho_1}{1-\rho_0} > 1$ and consider the test φ^* which outputs 0 if $1 - \hat{\sigma}_{1,2,t} \leq (1 - \rho_0)R^{\frac{1}{2}}$ and outputs 1 otherwise. Using Lemma 5 in the Appendix, one has the following upper bound on the maximal risk of φ^* ,

$$\mathcal{R}_t(\varphi^*) \leq \exp\left(-t \cdot \alpha(R^{\frac{1}{2}})\right) \leq \exp(-c_1(1/2) \cdot t \cdot \alpha(R)).$$

On the other hand, using Lemma 6 and Lemma 7 in the Appendix, we know that the maximum risk for any test φ is lower bounded as follows.

$$\begin{aligned} \mathcal{R}_t(\varphi) &\geq \frac{1}{4} \exp(-\text{KL}(\mathcal{N}(0, \Sigma_0)^{\otimes t}, \mathcal{N}(0, \Sigma_1)^{\otimes t})) \\ &= \frac{1}{4} \exp(-t \cdot \text{KL}(\mathcal{N}(0, \Sigma_0), \mathcal{N}(0, \Sigma_1))) \geq \frac{1}{4} \exp(-ct \cdot \alpha(R)) \end{aligned}$$

where c is an universal constant. We conclude that the test φ^* using the difference-based correlation estimator is an optimal test in the sense that it needs $\mathcal{O}(\frac{1}{\alpha(R)})$ samples to reliably identify the correct Σ and no test can do better.

3 Non-Adaptive Sampling Setting

In this section, we study the most correlated arms identification problem under the classical non-adaptive setting where we have access to m full sample vectors $(X^t)_{t=1,\dots,m}$ (this corresponds to $n = Km$ samples of arms). We show that a naive decision policy based on the difference-based correlation estimator needs $n = \tilde{\mathcal{O}}\left(\frac{K}{\alpha(R_{(h+1),\Sigma})} + K\right)$ samples to find S^* . We also provide a matching lower bound (up to a logarithmic factor on K) on problem instances with a certain correlation matrix structure.

3.1 Upper Bound

Consider the decision policy that simply outputs the empirically optimal subset \hat{S} , as described in Figure 1.

Input: m full samples $(X^t)_{t=1,\dots,m}$.
Output: $\hat{S} = \arg \max_{S \subset [K], |S|=h} \sum_{j,\ell \in S, j \neq \ell} \hat{\sigma}_{j,\ell,m}$

Figure 1: A naive decision policy under the non-adaptive sampling setting.

Theorem 1 *The probability of error of the decision policy in Figure 1 is bounded as*

$$\mathbb{P}_{\Sigma}(\hat{S} \neq S^*) \leq K(K-1) \exp(-cm \cdot \alpha(R_{(h+1),\Sigma}))$$

where c is an universal constant.

Proof First, observe that if $\hat{S} \neq S^*$, then

$$\left(\sum_{(j,\ell) \in \hat{S}^2 \setminus (S^* \cap \hat{S})^2, j \neq \ell} (1 - \sigma_{j,\ell}) \right) \geq R_{(h+1),\Sigma} \left(\sum_{(j,\ell) \in S^{*2} \setminus (S^* \cap \hat{S})^2, j \neq \ell} (1 - \sigma_{j,\ell}) \right)$$

However, by definition of \tilde{S} , one has

$$\left(\sum_{(j,\ell) \in \hat{S}^2 \setminus (S^* \cap \hat{S})^2, j \neq \ell} (1 - \hat{\sigma}_{j,\ell,m}) \right) \leq \left(\sum_{(j,\ell) \in S^{*2} \setminus (S^* \cap \hat{S})^2, j \neq \ell} (1 - \hat{\sigma}_{j,\ell,m}) \right)$$

Therefore, if $\hat{S} \neq S^*$, then one of the following two events must hold true

$$\begin{aligned} & \{\exists j \neq \ell \text{ such that } (1 - \hat{\sigma}_{j,\ell,m}) \geq R_{(h+1),\Sigma}^{\frac{1}{2}}(1 - \sigma_{j,\ell})\} \\ & \{\exists j \neq \ell \text{ such that } (1 - \hat{\sigma}_{j,\ell,m}) \leq R_{(h+1),\Sigma}^{-\frac{1}{2}}(1 - \sigma_{j,\ell})\}. \end{aligned}$$

Therefore, using Lemma 5, one gets

$$\mathbb{P}_{\Sigma}(\hat{S} \neq S^*) \leq K(K-1) \exp\left(-m \cdot \alpha\left(R_{(h+1),\Sigma}^{\frac{1}{2}}\right)\right) \leq K(K-1) \exp\left(-m \cdot c_1 \left(\frac{1}{2}\right) \alpha\left(R_{(h+1),\Sigma}\right)\right)$$

which completes the proof. ■

3.2 Lower Bound

Let $1 > \rho_h > \rho_{h+1} \geq \dots \geq \rho_K$ and consider the correlation matrix Σ defined by

$$\sigma_{j\ell} = \begin{cases} 1, & \text{if } j = \ell \text{ or } j, \ell < h \\ \rho_j \rho_\ell, & \text{if } j \neq \ell \text{ and } j, \ell \geq h \\ \rho_\ell, & \text{if } j < h \leq \ell \\ \rho_j, & \text{if } j \geq h > \ell \end{cases}$$

It is easy to verify that the optimal subset S^* is $[h]$ and for $i \geq h + 1$, $R_{(i),\Sigma} = R_{i,\Sigma} = \frac{1-\rho_i}{1-\rho_h}$. Now let $\rho'_{h+1} = 1 - \frac{(1-\rho_h)^2}{1-\rho_{h+1}}$ and define another correlation matrix Σ' which is constructed from Σ by replacing ρ_{h+1} by ρ'_{h+1} in the definition of Σ . By observing that $\rho'_{h+1} > \rho_h$, we see that the optimal subset for Σ' is $S_{\Sigma'}^* = [h-1] \cup \{h+1\}$ and

$$R_{(h+1),\Sigma'} = R_{h,\Sigma'} = \frac{1-\rho_h}{1-\rho'_{h+1}}, R_{(i),\Sigma} = R_{i,\Sigma} = \frac{1-\rho_i}{1-\rho'_{h+1}}, \text{ for } i \geq h+2.$$

Note that $R_{(i),\Sigma'} \geq R_{(i),\Sigma}$ for any $i \geq h+1$. So the identification problem with Σ' is always easier than that with Σ . The following theorem shows that for any decision policy with uniform sampling, its probability of error is at least $\Omega(\exp(-cm \cdot \alpha(R_{(h+1),\Sigma})))$ under one of the two problem instances with underlying correlation matrix Σ and Σ' .

Theorem 2 *Let φ be a decision policy that outputs the subset $\varphi((X^t)_{t=1,\dots,m})$ when $(X^t)_{t=1,\dots,m}$ is observed. Then*

$$\max(\mathbb{P}_\Sigma(\varphi \neq S^*), \mathbb{P}_{\Sigma'}(\varphi \neq S_{\Sigma'}^*)) \geq \frac{1}{4} \exp(-cm \cdot \alpha(R_{(h+1),\Sigma})) \geq \frac{1}{4} \exp(-cm \cdot \alpha(R_{(h+1),\Sigma'})).$$

where c is an universal constant.

Proof The first inequality can be easily proved by using Lemma 6 and Lemma 8 in Appendix. The second inequality is trivial. ■

4 Fixed-Budget Setting: Successive Rejects for Correlation

In this section, we present and analyze a new algorithm, SR-C (Successive Rejects for Correlation). See Figure 2 for its complete description. SR-C proceeds by rounds. At the end of each round, it computes an estimate for the suboptimality of each arm and rejects an arm that seems least likely to be an optimal arm. And during the next round, it samples the subset of existing arms a certain number of times. SR-C stops when there are only h arms left and outputs the set of remaining arms. It is easy to verify that SR-C does not exceed the sample budget n (the total number of samples it uses is $n_1 + \dots + n_{K-h-1} + (h+1)n_{K-h}$). The following theorem shows that SR-C needs at most $\tilde{\mathcal{O}}(H_C + K)$ samples to find the optimal subset. More precisely, when the sample budget n is of order $\mathcal{O}(H_C \log(K)^2)$, we have a non-trivial upper bound on the probability of error.

Input: sample budget n .

Let $S_1 = [K]$, $\overline{\log(\frac{K}{h})} = 1 + \sum_{i=h+1}^K \frac{1}{i}$, $n_0 = 0$ and for $k = 1, \dots, K - h$, define

$$n_k = \left\lceil \frac{n - K - 1}{\overline{\log(\frac{K}{h})(K + 1 - k)}} \right\rceil$$

For each round $k = 1, \dots, K - h$, do

(1) For $t = n_{k-1} + 1, \dots, n_k$, choose $A_t = S_k$ and sample A_t from X^t .

(2) For each $i \in S_k$, compute $U_{i,k} = \max_{A \subset S_k, |A|=h} \min_{i \in B \subset S_k, |B|=h} \hat{\mathcal{D}}_k(A, B)$ where

$$\hat{\mathcal{D}}_k(A, B) = \frac{\sum_{(j,\ell) \in B^2 \setminus (A \cap B)^2, j \neq \ell} (1 - \hat{\sigma}_{j,\ell,n_k})}{\sum_{(j,\ell) \in A^2 \setminus (A \cap B)^2, j \neq \ell} (1 - \hat{\sigma}_{j,\ell,n_k})}$$

with the notation $\frac{0}{0} = 1$.

(3) Let $i_k = \arg \max_{i \in S_k} U_{i,k}$ and $S_{k+1} = S_k \setminus \{i_k\}$.

Output: $\hat{S} = S_{K-h+1}$.

Figure 2: SR-C (Successive Rejects for Correlation) algorithm

Theorem 3 *The probability of error of SR-C satisfies*

$$\mathbb{P}_\Sigma(\hat{S} \neq S^*) \leq K^3 \exp\left(-\frac{c(n - K - 1)}{\overline{\log(\frac{K}{h})} H_C}\right)$$

where c is an universal constant.

Proof Without loss of generality, we assume that $S^* = [h]$ and

$$1 = R_{1,\Sigma} = \dots = R_{h,\Sigma} < R_{h+1,\Sigma} \leq \dots \leq R_{K,\Sigma}.$$

For $k = 1, \dots, K - h$, define the event

$$G_k = \bigcap_{j,\ell \in [K], j \neq \ell} \left\{ R_{K+1-k,\Sigma}^{-\frac{1}{4}} (1 - \sigma_{j,\ell}) < 1 - \hat{\sigma}_{j,\ell,n_k} < R_{K+1-k,\Sigma}^{\frac{1}{4}} (1 - \sigma_{j,\ell}) \right\}$$

and $G = \bigcap_{k \in [K-h]} G_k$. G is the event when the correlation estimators are close to the true correlations. Observe that when G occurs, the following inequalities hold true

$$R_{K+1-k,\Sigma}^{-\frac{1}{2}} \mathcal{D}_\Sigma(A, B) < \hat{\mathcal{D}}_k(A, B) < R_{K+1-k,\Sigma}^{\frac{1}{2}} \mathcal{D}_\Sigma(A, B).$$

The rest of the proof is organized into two steps.

Step One: In this step, we prove that when G occurs, SR-C always finds the optimal subset, i.e. $S_{K-h+1} = S^*$. We do this by induction. Fix a $k \in [K - h]$ and assume that none of the optimal

arms have ever been rejected before the round k . We need to show that none of them are rejected at the end of the round k .

First, there must exist a suboptimal arm \underline{i} that belongs to $\{K + 1 - k, \dots, K\} \cap S_k$ since $|S_k| = K + 1 - k$. Then for any B such that $\underline{i} \in B \subset S_k, |B| = h$, one has

$$\hat{\mathcal{D}}_k(S^*, B) > R_{K+1-k, \Sigma}^{-\frac{1}{2}} \mathcal{D}_\Sigma(S^*, B) \geq R_{K+1-k, \Sigma}^{\frac{1}{2}}$$

where the last inequality follows from $\mathcal{D}_\Sigma(S^*, B) \geq R_{\underline{i}, \Sigma} \geq R_{K+1-k, \Sigma}$. Therefore, we have $U_{\underline{i}, k} > R_{K+1-k, \Sigma}^{\frac{1}{2}}$.

Next, for any A such that $A \subset S_k, |A| = h$, one has

$$\hat{\mathcal{D}}_k(A, S^*) < R_{K+1-k, \Sigma}^{\frac{1}{2}} \mathcal{D}_\Sigma(A, S^*) \leq R_{K+1-k, \Sigma}^{\frac{1}{2}}$$

This shows that for any optimal arm $\bar{i} \in S^*$, we have $U_{\bar{i}, k} < R_{K+1-k, \Sigma}^{\frac{1}{2}} < U_{\underline{i}, k}$. Thus, \underline{i} is not removed at the end of the round k .

Step Two: From step one, we know that the probability of error of SR-C is bounded by $\mathbb{P}(G^c)$. Using Lemma 5, one obtains

$$\begin{aligned} \mathbb{P}(G^c) &\leq K(K-1) \sum_{k=1}^{K-h} \exp\left(-n_k \cdot \alpha\left(R_{K+1-k, \Sigma}^{\frac{1}{4}}\right)\right) \\ &\leq K(K-1) \sum_{k=1}^{K-h} \exp\left(-\frac{c_1(\frac{1}{4})(n-K-1)\alpha(R_{K+1-k, \Sigma})}{\log(\frac{K}{h})(K+1-k)}\right) \\ &= K(K-1) \sum_{k=h+1}^K \exp\left(-\frac{c_1(\frac{1}{4})(n-K-1)\alpha(R_{k, \Sigma})}{\log(\frac{K}{h})k}\right) \\ &\leq K(K-1) \sum_{k=h+1}^K \exp\left(-\frac{c_1(\frac{1}{4})(n-K-1)}{\log(\frac{K}{h})H_C}\right) \\ &\leq K^3 \exp\left(-\frac{c_1(\frac{1}{4})(n-K-1)}{\log(\frac{K}{h})H_C}\right), \end{aligned}$$

which completes the proof. ■

5 Fixed-Confidence Setting: Successive Elimination for Correlation

In this section, we introduce and study a new algorithm, SE-C (Successive Elimination for Correlation). See Figure 3 for its complete description. SE-C is a more dynamic algorithm than SR-C. SE-C updates the suboptimality estimates for the existing arms at each time step rather than at the

end of each round. And while SR-C always rejects an arm at each of some prespecified time steps, SE-C eliminates arms only when those arms are judged suboptimal with enough confidence. The following theorem shows that SE-C needs at most $\tilde{\mathcal{O}}(H_C + K)$ samples to find the optimal subset.

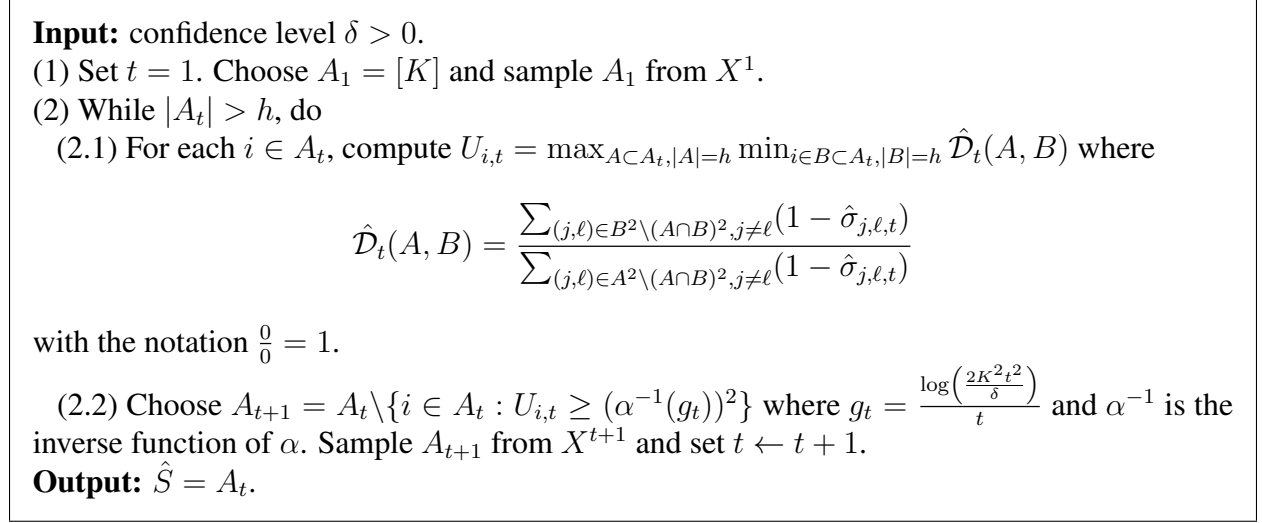


Figure 3: SE-C (Successive Elimination for Correlation) algorithm

Theorem 4 *With probability at least $1 - \delta$, SE-C returns the optimal subset S^* and the number of samples it uses is bounded by*

$$\mathcal{O} \left(h \max \left(1, \frac{\log \left(\frac{2K^2}{\delta \alpha(R_{(h+1), \Sigma})} \right)}{\alpha(R_{(h+1), \Sigma})} \right) + \sum_{i=h+1}^K \max \left(1, \frac{\log \left(\frac{2K^2}{\delta \alpha(R_{(i), \Sigma})} \right)}{\alpha(R_{(i), \Sigma})} \right) \right) = \tilde{\mathcal{O}}(H_C + K).$$

Proof Without loss of generality, we assume that $S^* = [h]$ and

$$1 = R_{1, \Sigma} = \dots = R_{h, \Sigma} < R_{h+1, \Sigma} \leq \dots \leq R_{K, \Sigma}.$$

Define the event

$$G = \bigcap_{j, \ell \in [K], j \neq \ell} \{ \forall t \geq 1, (\alpha^{-1}(g_t))^{-1} (1 - \sigma_{j,\ell}) < 1 - \hat{\sigma}_{j,\ell,t} < \alpha^{-1}(g_t) (1 - \sigma_{j,\ell}) \}.$$

Observe that when G occurs, the following inequalities hold true

$$(\alpha^{-1}(g_t))^{-2} \mathcal{D}_{\Sigma}(A, B) < \hat{\mathcal{D}}_t(A, B) < (\alpha^{-1}(g_t))^2 \mathcal{D}_{\Sigma}(A, B).$$

Once again, we use Lemma 5 to bound the probability of G^c as

$$\mathbb{P}(G^c) \leq K(K-1) \sum_{t=1}^{+\infty} \exp(-tg_t) \leq \sum_{t=1}^{+\infty} \frac{\delta}{2t^2} \leq \delta.$$

Now to prove the theorem, it is enough to show that under the event G , SE-C always terminates with the optimal subset S^* and the number of samples it uses is bounded as claimed. These two facts are proved separately in the following two steps.

Step One: Fix a $t \geq 1$. Assume that all of the optimal arms are still present at the beginning of time t , that is, $S^* \in A_t$. For any A such that $A \subset A_t, |A| = h$, one has

$$\hat{\mathcal{D}}_t(A, S^*) < (\alpha^{-1}(g_t))^2 \mathcal{D}_\Sigma(A, S^*) \leq (\alpha^{-1}(g_t))^2.$$

This implies that for any optimal arm $\bar{i} \in S^*$, we have $U_{\bar{i},t} < (\alpha^{-1}(g_t))^2$ and thus \bar{i} is not eliminated at the end of round t .

Step Two: Consider a suboptimal arm \underline{i} . If \underline{i} is not eliminated before or at the end of time t , then

$$\begin{aligned} (\alpha^{-1}(g_t))^2 &> U_{\underline{i},t} \\ &\geq \min_{\underline{i} \in B \subset A_t, |B|=h} \hat{\mathcal{D}}_t(S^*, B) \\ &\geq (\alpha^{-1}(g_t))^{-2} \min_{\underline{i} \in B \subset A_t, |B|=h} \mathcal{D}_\Sigma(S^*, B) \\ &= (\alpha^{-1}(g_t))^{-2} R_{\underline{i},\Sigma} \end{aligned}$$

Therefore, \underline{i} cannot be sampled more than t times as long as t is such that $g_t \leq \alpha \left(R_{\underline{i},\Sigma}^{\frac{1}{4}} \right)$. The latter inequality holds true as long as $g_t \leq c_2 \left(\frac{1}{4} \right) \alpha \left(R_{\underline{i},\Sigma} \right)$. Thus, any suboptimal arm \underline{i} cannot be sampled more than $\mathcal{O} \left(\max \left(1, \frac{\log \left(\frac{2K^2}{\delta \alpha \left(R_{\underline{i},\Sigma} \right)} \right)}{\alpha \left(R_{\underline{i},\Sigma} \right)} \right) \right)$ times. Moreover, the algorithm stops once all the suboptimal arms have been eliminated. Thus, an optimal arm cannot be sampled more than $\mathcal{O} \left(\max \left(1, \frac{\log \left(\frac{2K^2}{\delta \alpha \left(R_{h+1,\Sigma} \right)} \right)}{\alpha \left(R_{h+1,\Sigma} \right)} \right) \right)$ times, which completes the proof. \blacksquare

6 Discussion

This work is a first step towards understanding the hardness of finding the most correlated arms with an adaptive sampling scheme. We proposed two algorithms SR-C and SE-C, and we show that both succeed with at most $\tilde{\mathcal{O}} \left(h \cdot \alpha \left(R_{(h+1),\Sigma} \right)^{-1} + \sum_{i=h+1}^K \alpha \left(R_{(i),\Sigma} \right)^{-1} + K \right)$ samples. The result of Section 3.2 together with known arguments from the best arm identification literature strongly indicates that the term $\sum_{i=h+1}^K \alpha \left(R_{(i),\Sigma} \right)^{-1}$ is unavoidable. On the other hand it is clear that the term $h \cdot \alpha \left(R_{(h+1),\Sigma} \right)^{-1}$ is suboptimal in some cases. For example, consider the problem instances that are equivalent, up to a permutation of the arms, to the one with the correlation matrix Σ described in Section 3.2. In these cases there is always a block of $h - 1$ perfectly correlated arms, and thus one vector sample will allow to identify this block. Consequently, for these problem instances a trivial modification of SR-C and SE-C will result in a reduced upper bound $\tilde{\mathcal{O}} \left(\alpha \left(R_{(h+1),\Sigma} \right)^{-1} + \sum_{i=h+1}^K \alpha \left(R_{(i),\Sigma} \right)^{-1} + K \right)$. It is an interesting challenge to design algorithms

that can significantly reduce the term $\tilde{\mathcal{O}}(h \cdot \alpha(R_{(h+1),\Sigma})^{-1})$ for general problem instances. It is worth noting that a similar term $\tilde{\mathcal{O}}(h \cdot \Delta_{(h+1)}^{-2})$ was observed when the best arm identification algorithm Successive Rejects was used for the task of identifying the h best arms. This term can be significantly reduced by allowing the algorithm to accept seemingly optimal arms early, in the same way as seemingly suboptimal arms are rejected early, see [Kalyanakrishnan et al. \[2012\]](#) and [Bubeck et al. \[2013\]](#) for details. Unfortunately, the same trick can not be applied easily to the problem of most correlated arms identification. The main difficulty is that when an optimal arm is accepted early, the remaining optimal arms are not necessarily the most mutually correlated arms among all the remaining arms, thus making the identification of the remaining optimal arms difficult (if not impossible). To summarize, a novel algorithmic idea is needed to improve our upper bound. Another interesting direction of further work is to prove a lower bound on the number of samples that any adaptive strategy must use.

References

- E. Arias-Castro, S. Bubeck, and G. Lugosi. Detection of correlations. *The Annals of Statistics*, 40(1):412–435, 2012a.
- E. Arias-Castro, S. Bubeck, and G. Lugosi. Detecting positive correlations in a multivariate sample. *arXiv preprint arXiv:1202.5536*, 2012b.
- J.-Y. Audibert, S. Bubeck, and R. Munos. Best arm identification in multi-armed bandits. In *Proceedings of the 23rd Annual Conference on Learning Theory (COLT)*, 2010.
- S. Bubeck, R. Munos, and G. Stoltz. Pure exploration in multi-armed bandits problems. In *Proceedings of the 20th International Conference on Algorithmic Learning Theory (ALT)*, 2009.
- S. Bubeck, T. Wang, and N. Viswanathan. Multiple identifications in multi-armed bandits. In *Proceedings of the 30th International Conference on Machine Learning (ICML)*, 2013.
- R.M. Castro, G. Lugosi, and P.-A. Savalle. Detection of correlations with adaptive sensing. *arXiv preprint arXiv:1311.5366*, 2013.
- E. Even-Dar, S. Mannor, and Y. Mansour. Action elimination and stopping conditions for the multi-armed bandit and reinforcement learning problems. *Journal of Machine Learning Research*, 7: 1079–1105, 2006.
- S. Kalyanakrishnan, A. Tewari, P. Auer, and P. Stone. Pac subset selection in stochastic multi-armed bandits. In *Proceedings of the 29th International Conference on Machine Learning (ICML)*, 2012.
- A.B. Tsybakov. Introduction to nonparametric estimation. *Springer*, 2009.

A Technical Lemmas

Lemma 1 (*Concentration Inequalities for Chi-Square Distributions*) Let Y be a random variable following the chi-square distribution with degree of freedom $t \in \mathbb{N}$. Then for any $\theta \geq 1$, the following concentration inequalities hold:

$$\mathbb{P}\left(\frac{Y}{t} \leq \frac{1}{\theta}\right) \leq \exp(-t \cdot \alpha(\theta))$$

$$\mathbb{P}\left(\frac{Y}{t} \geq \theta\right) \leq \exp\left(-\frac{t}{2}(\theta - 1 - \log \theta)\right) \leq \exp(-t \cdot \alpha(\theta)).$$

Proof The proof is a simple application of Chernoff's bounding technique. ■

Lemma 2 Let μ_0 and μ_1 be two probability distributions on some set \mathcal{X} , with μ_0 absolutely continuous with respect to μ_1 . Let X be a random variable taking values in \mathcal{X} . Then for any measurable function $\varphi : \mathcal{X} \rightarrow \{0, 1\}$, one has

$$\max(\mathbb{P}_{X \sim \mu_0}(\varphi(X) = 1), \mathbb{P}_{X \sim \mu_1}(\varphi(X) = 0)) \geq \frac{1}{4} \exp(-\text{KL}(\mu_0, \mu_1)).$$

Proof See Chapter 2 in [Tsybakov \[2009\]](#) for a proof. ■

Lemma 3 Let $1 > \rho_0 > \rho_1 \geq 0$ and $R = \frac{1-\rho_1}{1-\rho_0}$. Let $\Sigma_0 = \begin{pmatrix} 1 & \rho_0 \\ \rho_0 & 1 \end{pmatrix}$ and $\Sigma_1 = \begin{pmatrix} 1 & \rho_1 \\ \rho_1 & 1 \end{pmatrix}$. Then

$$\alpha(R) \leq \text{KL}(\mathcal{N}(0, \Sigma_0), \mathcal{N}(0, \Sigma_1)) = \text{KL}(\mathcal{N}(\rho_0, 1 - \rho_0^2), \mathcal{N}(\rho_1, 1 - \rho_1^2)) \leq c \cdot \alpha(R).$$

where $c > 1$ is an universal constant.

Proof Recall the following general formula for the KL-divergence between two Gaussian distributions on \mathbb{R}^k ,

$$\text{KL}(\mathcal{N}(\mu_0, \Sigma_0), \mathcal{N}(\mu_1, \Sigma_1)) = \frac{1}{2} \left(\log \left(\frac{\det \Sigma_1}{\det \Sigma_0} \right) + \text{Tr}(\Sigma_1^{-1} \Sigma_0) - k + (\mu_1 - \mu_0)^T \Sigma_1^{-1} (\mu_1 - \mu_0) \right).$$

By straightforward application of the above formula, one has

$$\begin{aligned} \text{KL}(\mathcal{N}(0, \Sigma_0), \mathcal{N}(0, \Sigma_1)) &= \text{KL}(\mathcal{N}(\rho_0, 1 - \rho_0^2), \mathcal{N}(\rho_1, 1 - \rho_1^2)) \\ &= \frac{1}{2} \left(\log \frac{1 - \rho_1}{1 - \rho_0} + \frac{1 - \rho_0}{1 - \rho_1} - 1 + \log \frac{1 + \rho_1}{1 + \rho_0} + \frac{1 + \rho_0}{1 + \rho_1} - 1 \right) \\ &= \alpha(R) + \beta \left(\frac{1 + \rho_0}{1 + \rho_1} \right) \end{aligned}$$

where the function β is defined as $\beta(\theta) = \theta - 1 - \log \theta$, for $\theta \geq 1$. It is easy to see that β is a positive, strictly increasing function on $[1, +\infty[$. Since $\frac{1+\rho_0}{1+\rho_1} \leq \min(2, R)$, one has

$$\beta\left(\frac{1+\rho_0}{1+\rho_1}\right) \leq \beta(\min(2, R)).$$

Moreover, because $\beta(R) = \Theta((R-1)^2)$ when $R \rightarrow 1$, there exists a universal constant $c > 0$ such that $\beta(\min(2, R)) \leq c \cdot \alpha(R)$ for any $R \geq 1$, which completes the proof. \blacksquare

Lemma 4 Consider the two correlation matrices Σ, Σ' described in section 3.2. One has

$$\text{KL}\left(\mathcal{N}(0, \Sigma'), \mathcal{N}(0, \Sigma)\right) \leq c \cdot \alpha(R_{h+1, \Sigma})$$

where c is an universal constant.

Proof Let $Z_0, Z_h, Z_{h+1}, \dots, Z_K$ be i.i.d. standard Gaussian random variables. Define two random vectors $W = (W_1, \dots, W_K)$ and $W' = (W'_1, \dots, W'_K)$ as

$$W_i = \begin{cases} Z_0, & \text{if } i \leq h-1 \\ \rho_i Z_0 + \sqrt{1-\rho_i^2} Z_i, & \text{if } i \geq h \end{cases}$$

and

$$W'_i = \begin{cases} Z_0, & \text{if } i \leq h-1 \\ \rho_i Z_0 + \sqrt{1-\rho_i^2} Z_i, & \text{if } i \geq h \text{ and } i \neq h+1 \\ \rho'_{h+1} Z_0 + \sqrt{1-\rho'^2_{h+1}} Z_{h+1}, & \text{if } i = h+1 \end{cases}$$

Then it is easy to see that W and W' are two centered Gaussian vectors with covariance matrices Σ and Σ' . Moreover, conditioning on $W_1 = z$ for some $z \in \mathbb{R}$, W_2, \dots, W_K are independent and Gaussian distributed as $W_i \sim \mathcal{N}(z, 0)$ for $2 \leq i \leq h-1$ and $W_i \sim \mathcal{N}(\rho_i z, 1 - \rho_i^2)$ for $i \geq h$. Similarly, conditioning on $W'_1 = z$, W'_2, \dots, W'_K are independent and Gaussian distributed as $W'_i \sim \mathcal{N}(z, 0)$ for $2 \leq i \leq h-1$, $W'_{h+1} \sim \mathcal{N}(\rho'_{h+1} z, 1 - \rho'^2_{h+1})$ and $W'_i \sim \mathcal{N}(\rho_i z, 1 - \rho_i^2)$ for $i \geq h$ and $i \neq h+1$. In what follows, we use $\mathcal{L}(Y)$ to represent the distribution of a random

variable Y . Using the conditioning formula for KL-divergence, one has

$$\begin{aligned}
& \text{KL} \left(\mathcal{N}(0, \Sigma'), \mathcal{N}(0, \Sigma) \right) \\
&= \text{KL} \left(\mathcal{L}(W'_1), \mathcal{L}(W_1) \right) + \mathbb{E}_{z \sim W'_1} \left[\text{KL} \left(\mathcal{L}((W'_2, \dots, W'_K) | W'_1 = z), \mathcal{L}((W_2, \dots, W_K) | W_1 = z) \right) \right] \\
&= \mathbb{E}_{z \sim \mathcal{N}(0,1)} \left[\text{KL} \left(\mathcal{N}(\rho'_{h+1}z, 1 - \rho'^2_{h+1}), \mathcal{N}(\rho_{h+1}z, 1 - \rho^2_{h+1}) \right) \right] \\
&= \mathbb{E}_{z \sim \mathcal{N}(0,1)} \left[\frac{1}{2} \left(\log \left(\frac{1 - \rho^2_{h+1}}{1 - \rho'^2_{h+1}} \right) + \frac{1 - \rho'^2_{h+1}}{1 - \rho^2_{h+1}} - 1 + \frac{(\rho_{h+1} - \rho'_{h+1})^2 z^2}{1 - \rho^2_{h+1}} \right) \right] \\
&= \frac{1}{2} \left(\log \left(\frac{1 - \rho^2_{h+1}}{1 - \rho'^2_{h+1}} \right) + \frac{1 - \rho'^2_{h+1}}{1 - \rho^2_{h+1}} - 1 + \frac{(\rho_{h+1} - \rho'_{h+1})^2}{1 - \rho^2_{h+1}} \right) \\
&= \text{KL} \left(\mathcal{N}(\rho'_{h+1}, 1 - \rho'^2_{h+1}), \mathcal{N}(\rho_{h+1}, 1 - \rho^2_{h+1}) \right) \\
&\leq c \cdot \alpha \left(\frac{1 - \rho_{h+1}}{1 - \rho'_{h+1}} \right)
\end{aligned}$$

where c is a universal constant and the last step follows from Lemma 7. To conclude, it is enough to observe that $\alpha \left(\frac{1 - \rho_{h+1}}{1 - \rho'_{h+1}} \right) = \alpha \left(\left(\frac{1 - \rho_{h+1}}{1 - \rho_h} \right)^2 \right) = \alpha(R^2_{h+1, \Sigma}) \leq c_2(2)\alpha(R_{h+1, \Sigma})$. ■