Rate-Optimal Budget Allocation for the Probability of Good Selection

David J. Eckman

Texas A&M University, ISEN

Winter Simulation Conference
December 17, 2024

Joint work with Taeho Kim.

Supported by NSF grant CMMI-2206972.

Ranking and Selection

Problem: Select from among k simulated alternatives (systems).

- Output of System i is assumed to be distributed as $N(\mu_i, \sigma_i^2)$ for $i = 1, 2, \ldots, k$.
- Means μ_i and variances σ_i^2 are unknown.
- Seek to minimize the expected (mean) output.

Approach:

1. Allocate a budget of n simulation replications across systems.
2. Calculate sample means $\mu_{i,n}$ for $i = 1, 2, \ldots, k$.
3. Select System $\mathcal{D}_n := \arg\min_{1 \leq i \leq k} \mu_{i,n}$.

Goal: Maximize the **probability of correct selection (PCS)**:

$$\mathrm{PCS}_n := \mathrm{P}\left(\mathcal{D}_n = \arg\min_{1 \leq i \leq k} \mu_i\right)$$

How should we allocate the n replications to achieve this goal?

Large Deviations Analysis for PCS

Assume $\mu_1 < \mu_2 \leq \ldots \leq \mu_k$ and that the means and variances are *known*.

Results from Glynn and Juneja (2004)

For any static allocation $\boldsymbol{\alpha} = (\alpha_1, \alpha_2, \ldots, \alpha_k) \in \Delta^k$,

$$\lim_{n\to\infty} \underbrace{-\frac{1}{n}\log(1 - \mathrm{PCS}_n)}_{\text{large deviations rate (LDR)}} = \min_{2\leq i\leq k} G_i(\boldsymbol{\alpha}), \text{ where}$$

$$G_i(\boldsymbol{\alpha}) := \lim_{n\to\infty} -\frac{1}{n}\log \mathrm{P}(\mu_{i,n} \leq \mu_{1,n}) = \frac{(\mu_i - \mu_1)^2}{2\left(\sigma_1^2/\alpha_1 + \sigma_i^2/\alpha_i\right)} \text{ for } i = 2, 3, \ldots, k.$$

The LDR is determined by the hardest system to distinguish from System 1.

The optimal allocation for maximizing the rate at which the PCS approaches 1 is

$$\boldsymbol{\alpha}^{\mathrm{pcs}} := \arg\max_{\boldsymbol{\alpha}\in\Delta^k} \min_{2\leq i\leq k} G_i(\boldsymbol{\alpha}).$$

Good Selection

Let $\delta > 0$ be a *user-specified tolerance* and assume that

$$\underbrace{\mu_1 \le \mu_2 \le \cdots \le \mu_\ell}_{\ell \text{ good systems}} \le \mu_1 + \delta < \underbrace{\mu_{\ell+1} \le \cdots \le \mu_k}_{k-\ell \text{ bad systems}} \text{ for some } 1 \le \ell < k.$$

Relaxed Goal: Maximize the **probability of good selection (PGS)**:

$$\mathrm{PGS}_n = \mathrm{P}(\mathcal{D}_n \in \{1, 2, \ldots, \ell\}).$$

In terms of the PGS goal, α^{pcs} allocates too much effort to the good systems.

Large Deviations Analysis for PGS

Results from Kim et al. (2022)

The optimal allocation for maximizing the rate at which the PGS approaches 1 is

$$\boldsymbol{\alpha}^{\mathrm{pgs}} := \arg\max_{\boldsymbol{\alpha}\in\Delta^k} \min_{\ell+1\le j\le k} \widetilde{G}_j(\boldsymbol{\alpha}), \text{ where}$$

$$\widetilde{G}_j(\boldsymbol{\alpha}) = \min_{x\in[\mu_1,\mu_j]}\left\{\frac{\alpha_j}{2\sigma_j^2}(x-\mu_j)^2 + \sum\nolimits_{1\le i<j}\frac{\alpha_i}{2\sigma_i^2}[(x-\mu_i)^+]^2\right\} \text{ for } j = 2,3,\ldots,k.$$

Can derive necessary and sufficient conditions for $\boldsymbol{\alpha}^{\mathrm{pgs}}$ based on the KKT conditions.

Zero-Sampling Phenomenon

Unlike $\boldsymbol{\alpha}^{\mathrm{pcs}}$, some components of $\boldsymbol{\alpha}^{\mathrm{pgs}}$ are **zero**!

- This only occurs for good systems.

What does zero sampling mean?

- *Asymptotically*, we allocate a vanishing fraction of the budget to that system.

Why does zero sampling occur?

- *Easier* to determine that some good systems are good than other good systems.

Numerical Example

Example with $k = 10$ systems (2 good, 8 bad) and where $\mu_1 = 1$ and $\delta = 1$.

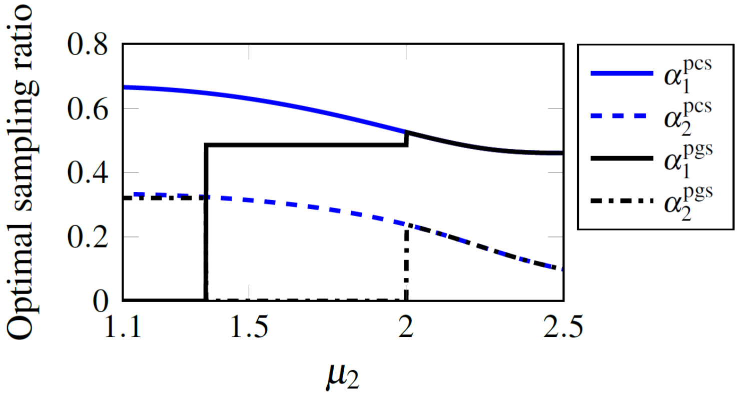

Systems 1 and 2 take turns receiving zero sampling.

When Does Zero Sampling Occur?

Sufficient conditions

For any System $j \le \ell$, if $\frac{\mu_{\ell+1}-\mu_i}{\sigma_i} > \frac{\mu_{\ell+1}-\mu_j}{\sigma_j}$ and $\mu_i \le \mu_j$ for some $i \ne j$, then $\alpha_j^{\mathrm{pgs}} = 0$.

Stronger sufficient conditions

For any System $j \le \ell$, if $\mu_i \le \mu_j$ and $\sigma_i^2 < \sigma_j^2$ for any $i \ne j$, then $\alpha_j^{\mathrm{pgs}} = 0$.

Necessary and sufficient conditions (for $k = 3$)

Assume $\mu_1 < \mu_2 \le \mu_1 + \delta < \mu_3$.

- $\alpha_1^{\mathrm{pgs}} > 0$ and $\alpha_2^{\mathrm{pgs}} > 0$ if and only if $\frac{\mu_3-\mu_1}{\sigma_3+\sigma_1} = \frac{\mu_3-\mu_2}{\sigma_3+\sigma_2}$;
- $\alpha_1^{\mathrm{pgs}} > 0$ and $\alpha_2^{\mathrm{pgs}} = 0$ if and only if $\frac{\mu_3-\mu_1}{\sigma_3+\sigma_1} > \frac{\mu_3-\mu_2}{\sigma_3+\sigma_2}$; and
- $\alpha_1^{\mathrm{pgs}} = 0$ and $\alpha_2^{\mathrm{pgs}} > 0$ if and only if $\frac{\mu_3-\mu_1}{\sigma_3+\sigma_1} < \frac{\mu_3-\mu_2}{\sigma_3+\sigma_2}$.

Plug-in Algorithm with Batch Allocation

1. Run n_0 replications of each system and update estimates $\mu_{i,n}$ and $\sigma^2_{i,n}$.
2. Define $\mathcal{I}_n := \{i : \mu_{i,n} \le \min_{1\le j\le k} \mu_{j,n} + \delta\}$, the set of all systems that *look* good.
3. Solve the following optimization problem:

$$\boldsymbol{\alpha}_n = \arg\max_{\boldsymbol{\alpha}\in\Delta^k} \min_{j\notin\mathcal{I}_n} \widetilde{G}_{j,n}(\boldsymbol{\alpha}),$$

   where $\widetilde{G}_{j,n}(\boldsymbol{\alpha})$ is a **plug-in version** of $\widetilde{G}_j(\boldsymbol{\alpha})$.
4. Draw a sample of size B from a *multinomial* distribution with probability $\boldsymbol{\alpha}_n$.
5. Take prescribed additional replications and update $\mu_{i,n}$ and $\sigma^2_{i,n}$ accordingly.
6. If budget not exceeded, return to Step 2.

Empirical PGS

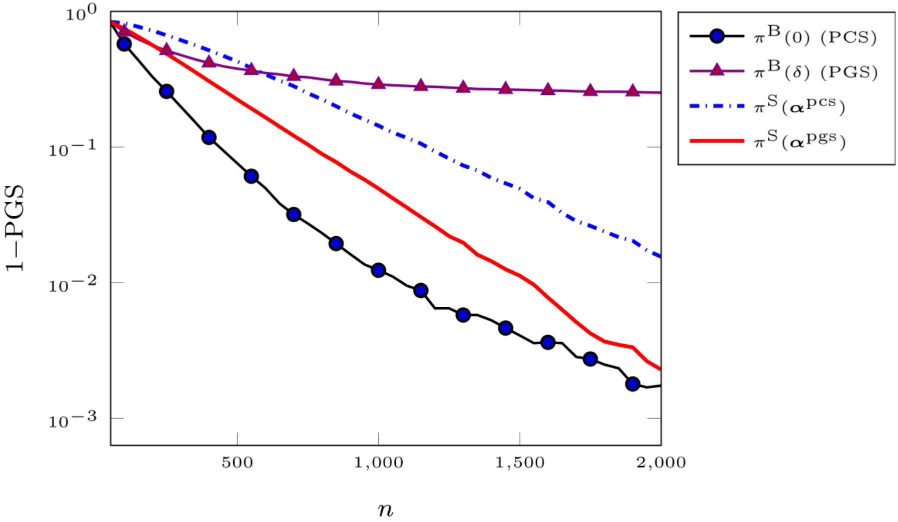

S = Static B = Batch

Two Bad Scenarios

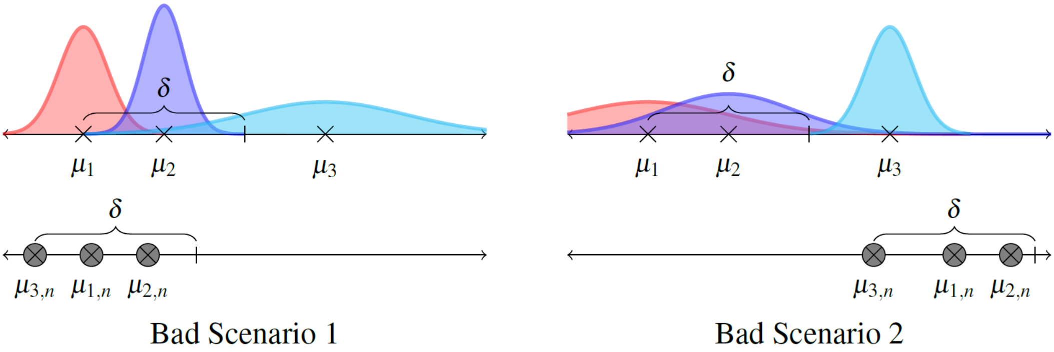

Sample Path of Sample Means

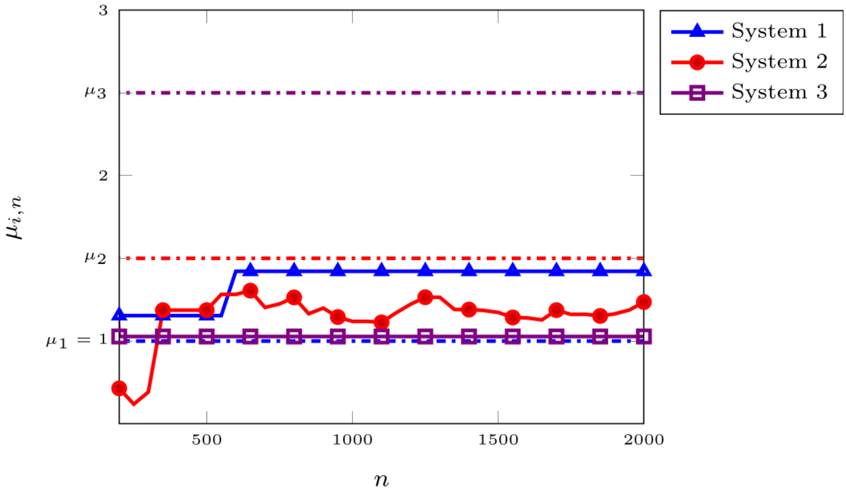

Sample path under Bad Scenario 1.

Samples paths like this can be used to prove that Algorithm $\pi^{\mathrm{B}}(\delta)$ is **inconsistent**.

A Mixture-Based Approach

We propose sampling according to a convex combination of $\boldsymbol{\alpha}^{\mathrm{pgs}}$ and $\boldsymbol{\alpha}^{\mathrm{pcs}}$:

$$\boldsymbol{\alpha}(\varepsilon) := (1-\varepsilon)\boldsymbol{\alpha}^{\mathrm{pgs}} + \varepsilon\boldsymbol{\alpha}^{\mathrm{pcs}}$$

for some $\varepsilon \in (0, 1)$.

Why mix with $\boldsymbol{\alpha}^{\mathrm{pcs}}$?

- $\boldsymbol{\alpha}^{\mathrm{pcs}}$ is aligned with making a correct (and therefore good) selection.
- Expected to perform better than ε-random sampling.

Performance of Mixing

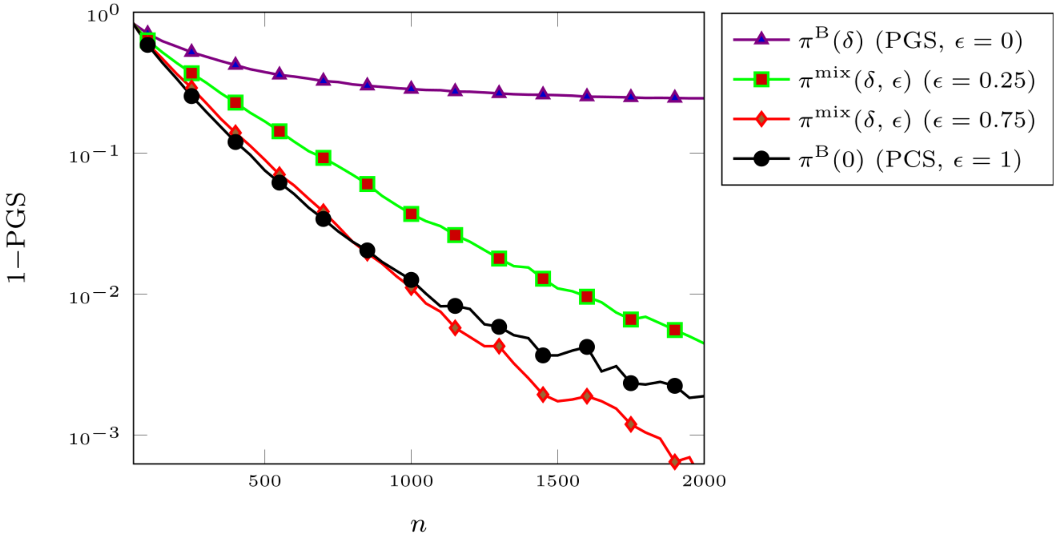

Performance of Mixing

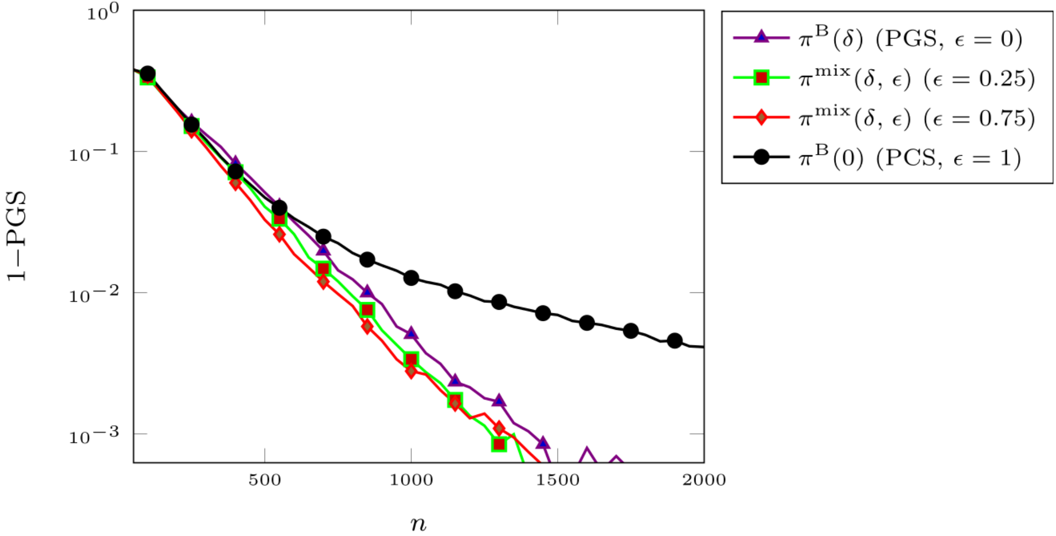

Conclusion

Takeaway

R&S problems having multiple good systems can have optimal allocations for PGS having *zero sampling ratios*, leading to performance issues with *adaptive* algorithms.

Future Research

- Necessary conditions for systems to have zero sampling ratios.
- Other algorithm-design remedies for issues caused by zero sampling ratios.
- Exploring other selection rules for cases with multiple good systems.