Posterior-Based Stopping Rules for Bayesian Ranking-and-Selection Procedures

David Eckman

Dept of Industrial Engineering and Management Sciences

IEMS Seminar
September 29, 2020

Joint work with Shane Henderson (Cornell)

Overview

Ranking-and-selection procedures deliver Bayesian guarantees by tracking a posterior quantity and checking a stopping rule.

Common Practice

Use a conservative bound as a surrogate for the posterior quantity.

- **Tacit assumption:** Posterior quantity is costly to compute.

We show how to compute posterior quantities at modest cost.

- Can lead to savings in a procedure's total sample size.
- Monte Carlo estimators help to achieve further gains.

Optimization via Simulation

Decision-making under uncertainty involving a complex system.

- Optimize a scalar performance measure.

Example: Locating ambulance bases in a metropolitan area.

- Minimize the *expected* call response time.

Ranking-and-Selection (R&S):

1. Restrict attention to a finite set of alternatives (i.e., solutions).
2. Use **stochastic simulation** to evaluate performance.

Bayesian R&S Framework

Treat the performances of the alternatives as *random variables*:

- W_i is the performance of Alternative $i = 1, \ldots, k$.
- $\boldsymbol{W} = (W_1, \ldots, W_k)$ denotes the unknown problem instance.
- The ordered performances are $W_{[1]} \leq W_{[2]} \leq \cdots \leq W_{[k]}$.

Assume that larger W_i is better.

- Alternative $[k]$ is (one of) the best.

Approach:

1. The decision-maker places a *prior distribution* on $\boldsymbol{W}$.
2. Obtain observations (replications) from the alternatives.
3. Use Bayes' rule to update the *posterior distribution* of $\boldsymbol{W}$.
4. Return to Step 2 and repeat until a **stopping rule** is satisfied.

Assumptions

Let X_{ij} denote the jth observation from Alternative i and define

$$\overrightarrow{X_i} = \{X_{i1}, X_{i2}, \ldots\}.$$

Assumptions:

1. For each $i = 1, \ldots, k$, the sequence $\overrightarrow{X_i}$ consists of i.i.d. observations $X_{ij} \sim \mathcal{N}(w_i, \sigma_i^2)$.
2. The sequences $\overrightarrow{X_1}, \overrightarrow{X_2}, \ldots, \overrightarrow{X_k}$ are independent (i.e., no CRN).
3. The decision-maker has independent beliefs about $W_1, \ldots, W_k$.

Given 2–3, the posterior distribution of $\boldsymbol{W}$ has a *product* form.

Marginal Posterior Distribution of W_i

Suppose that the conjugate reference prior distribution is used.

Given n_i observations from Alternative i, let $\bar{x}_i$ and s_i^2 be the sample mean and variance.

Known Variances

$$W_i \sim \mathcal{N}(\bar{x}_i, \sigma_i^2/n_i) \equiv \mathcal{N}(\mu_i, \rho_i^2).$$

Unknown Variances

$$W_i \sim t_{n_i-1}(\bar{x}_i, s_i^2/n_i) \equiv t_{\nu_i}(\mu_i, \rho_i^2).$$

Bayesian Selection Criteria

Let $\mathbb{P}(\,\cdot \mid \mathcal{E})$ denote the posterior probability, given the evidence $\mathcal{E}$.

Posterior PCS of Alternative *i*

$$\text{pPCS}_i = \mathbb{P}(W_i = W_{[k]} \mid \mathcal{E}).$$

Posterior PGS of Alternative *i*

$$\text{pPGS}_i = \mathbb{P}(W_i \geq W_{[k]} - \delta \mid \mathcal{E}) \quad \text{where } \delta > 0.$$

Posterior EOC of Alternative *i*

$$\text{pEOC}_i = \mathbb{E}[W_{[k]} - W_i \mid \mathcal{E}].$$

Bayesian Guarantees

For $1-\alpha \in (1/k, 1)$ or $\beta > 0$, guarantee that

$$\text{pPCS}_d \geq 1-\alpha, \qquad \text{(pPCS Guarantee)}$$

$$\text{pPGS}_d \geq 1-\alpha, \text{ or} \qquad \text{(pPGS Guarantee)}$$

$$\text{pEOC}_d \leq \beta, \qquad \text{(pEOC Guarantee)}$$

where d is the index of the selected alternative.

Run a R&S procedure until a Bayesian guarantee is delivered.

- As opposed to *fixed-budget* settings.
- Our focus is on **stopping rules** instead of **allocation rules.**

Delivering Bayesian Guarantees

Consequence of the Stopping Rule Principle

It is valid to stop and select an alternative whenever its posterior quantity (e.g., pPCS/pPGS/pEOC) crosses some threshold.

Advantages:

1. Can *repeatedly compute* pPCS/pPGS/pEOC without sacrificing statistical validity.
2. Flexibility in allocating simulation replications.

Visualizing Stopping Rules for pPCS and pPGS

For $k = 2$, stop when $\left| n(\bar{X}_1 - \bar{X}_2) \right| \geq \sqrt{2n}\sigma\Phi^{-1}(1-\alpha) - \delta n$.

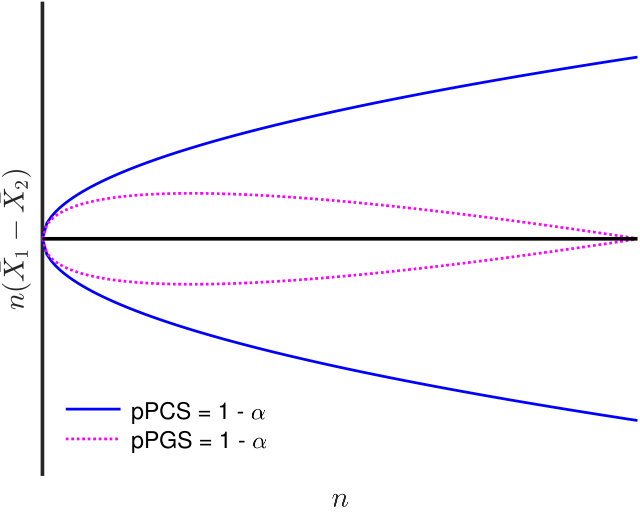

The pPCS Guarantee

Issues

- Small differences in performances $\implies$ long run times.
- Hard to justify the computational effort needed to detect differences in performance that are *not practically significant.*
- Some practical problems may have *multiple* optimal solutions.

The pPGS guarantee has none of these issues.

Computational Considerations

Motivation: pPGS and pEOC are k-dimensional integrals.

- Naive numerical integration becomes intractable for large k.

Common Practice

Use **bounds** on the posterior quantity of interest of the alternative with the *highest posterior mean*, denoted as Alternative (k).

- Results in conservative procedures, e.g., excessive sampling.

Tradeoff between simulation time and computational time.

- How precisely to check the stopping rule?

pPGS Bounds

Compute a *lower* bound on $\text{pPGS}_{(k)}$ – stop when it exceeds $1 - \alpha$.

Via Bonferroni's Inequality

$$\text{pPGS}_{(k)} \geq 1 - \sum_{j \neq (k)} \mathbb{P}(W_{(k)} < W_j - \delta \mid \mathcal{E}) =: \text{pPGS}_{(k)}^{\text{Bonf}}.$$

Via Slepian's Inequality

$$\text{pPGS}_{(k)} \geq \prod_{j \neq (k)} \mathbb{P}(W_{(k)} \geq W_j - \delta \mid \mathcal{E}) =: \text{pPGS}_{(k)}^{\text{Slep}}.$$

Tightness of pPGS Bounds

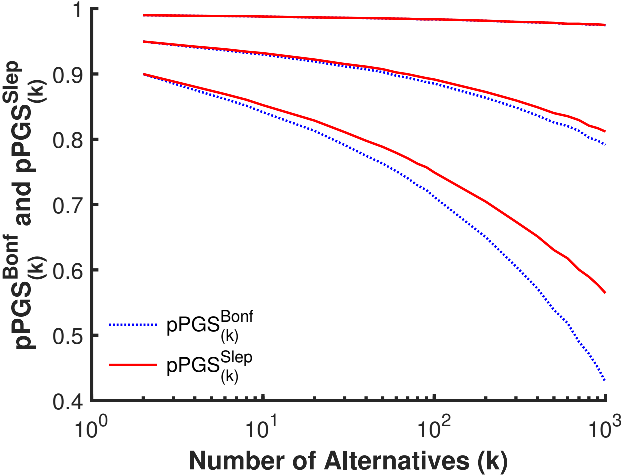

$\mathrm{pPGS}^{\mathrm{Bonf}}_{(k)}$ and $\mathrm{pPGS}^{\mathrm{Slep}}_{(k)}$ in a slippage configuration of posterior means in which $\mathrm{pPGS}_{(k)} = 1 - \alpha$ for $1 - \alpha = 0.90$, 0.95, and 0.99.

Computing pPGS

Conditioning on $W_{(k)}$ leads to a one-dimensional integral:

$$\begin{aligned} \text{pPGS}_{(k)} &= \mathbb{E}\left[\mathbb{P}(W_{(k)} \geq W_j - \delta \text{ for all } j \neq (k) \mid W_{(k)}, \mathcal{E}) \mid \mathcal{E}\right] \\ &= \int_{-\infty}^{\infty} \left[\prod_{j \neq (k)} F_{W_j|\mathcal{E}}(w+\delta)\right] f_{W_{(k)}|\mathcal{E}}(w) dw, \end{aligned}$$

Average times (seconds) for numerical integration of $\text{pPGS}_{(k)}$ for 1000 RPI.

k	10	100	1000	10,000
Average Time	0.005	0.009	0.029	0.138

pEOC Bounds

Compute an *upper* bound on $\text{pEOC}_{(k)}$ – stop when it is below β.

$$\text{pEOC}_{(k)} \leq \sum_{j \neq (k)} \mathbb{E}[(W_j - W_{(k)})^+ \mid \mathcal{E}] =: \text{pEOC}_{(k)}^{\text{Bonf}}.$$

Or integrate the tail of $W_{[k]} - W_{(k)}$ and apply the $\text{pPGS}_{(k)}^{\text{Slep}}$ bound:

$$\begin{aligned}
\text{pEOC}_{(k)} = \mathbb{E}[W_{[k]} - W_{(k)} \mid \mathcal{E}] &= \int_0^\infty \mathbb{P}(W_{[k]} - W_{(k)} > \delta \mid \mathcal{E})\, d\delta \\
&= \int_0^\infty \left[1 - \text{pPGS}_{(k)}\right] d\delta \\
&\leq \int_0^\infty \left[1 - \prod_{j \neq (k)} \mathbb{P}(W_{(k)} \geq W_j - \delta \mid \mathcal{E})\right] d\delta =: \text{pEOC}_{(k)}^{\text{Slep}}.
\end{aligned}$$

Tightness of pEOC Bounds

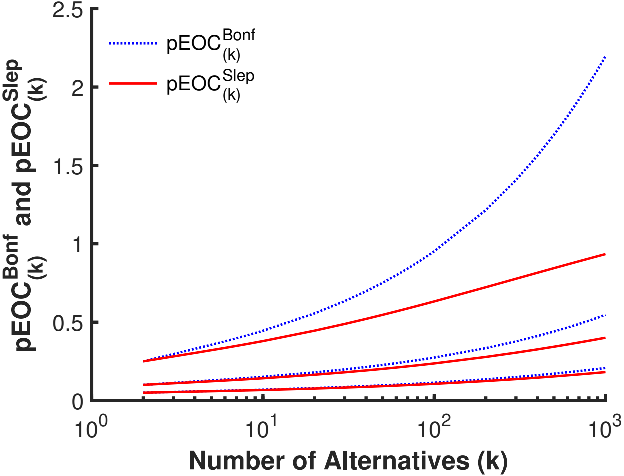

$\mathrm{pEOC}^{\mathrm{Bonf}}_{(k)}$ and $\mathrm{pEOC}^{\mathrm{Slep}}_{(k)}$ in a slippage configuration of posterior means in which $\mathrm{pEOC}_{(k)} = \beta$ for $\beta = 0.05$, 0.10, and 0.25.

Computing pEOC

$$\mathrm{pEOC}_{(k)} = \int_0^{\infty} \int_{-\infty}^{\infty} \left[1 - \prod_{j \neq (k)} F_{W_j|\mathcal{E}}(w+\delta) \right] f_{W_{(k)}|\mathcal{E}}(w)\, dw\, d\delta.$$

Average times (seconds) for numerical integration of $\mathrm{pEOC}_{(k)}$ for 1000 RPI.

k	10	100	1000
Average Time	1.20	1.66	4.54

Experimental Setup

Checked stopping rules by exactly computing pPGS and pEOC.

Compared to stopping rules using bounds:

- Stop when $\text{pPGS}_{(k)}^{\text{Bonf}} \geq 1 - \alpha$ or $\text{pPGS}_{(k)}^{\text{Slep}} \geq 1 - \alpha$.
- Stop when $\text{pEOC}_{(k)}^{\text{Bonf}} \leq \beta$ or $\text{pEOC}_{(k)}^{\text{Slep}} \leq \beta$.

Tested three allocation rules:

- Equal Allocation (EA)
- Optimal Computing Budget Allocation (OCBA)
- Thompson Sampling (TS)

Experimental Setup

Tested random problem instances with $k = 100$ alternatives.

- $w_i \sim -\text{Weibull}(\text{scale} = 1.5,\ \text{shape} = 2)$ and $\sigma_i^2 \sim \chi^2(4)$.
- $\approx 36\%$ of alternatives are "good" for $\delta = 1$.

Used **splitting** to speed up the experiments.

- Ran 100 macroreplications with 50 splits each.

Plotted the empirical cdf of the **fraction savings** (replications):

$$S = \frac{N_b - N_e}{N_b},$$

where N_b and N_e are the sample sizes for *bounds* and *exact*.

Fraction Savings for pPGS

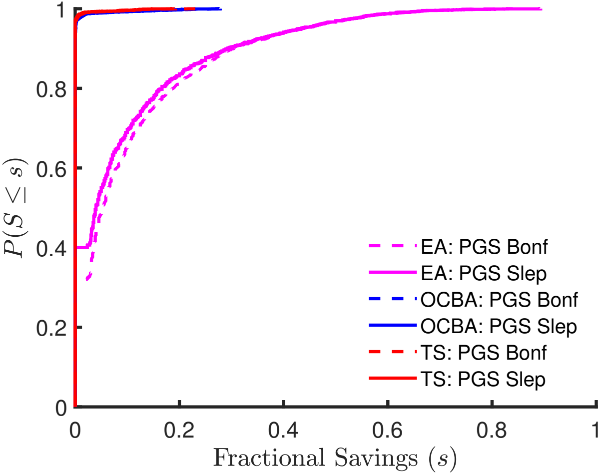

ECDF of fraction savings for the pPGS stopping rule with $1 - \alpha = 0.90$.

Fraction Savings for pEOC

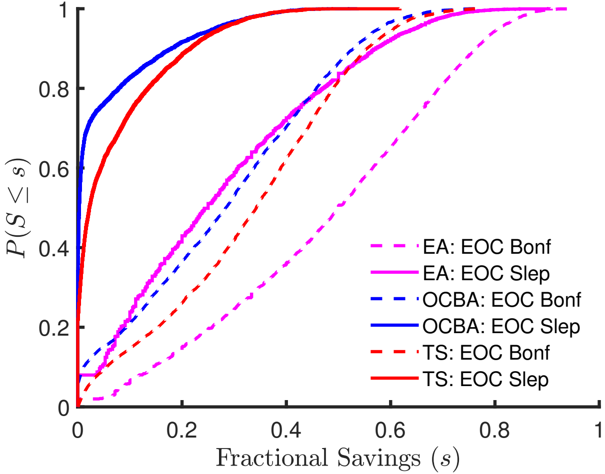

ECDF of fraction savings for the pEOC stopping rule with $\beta = 0.50$.

Average Fraction Savings

Bound	$\text{pPGS}_{(k)}^{\text{Bonf}}$	$\text{pPGS}_{(k)}^{\text{Slep}}$	$\text{pEOC}_{(k)}^{\text{Bonf}}$	$\text{pEOC}_{(k)}^{\text{Slep}}$
EA	0.11 ± 0.02	0.10 ± 0.02	0.48 ± 0.02	0.28 ± 0.02
TS	0.00 ± 0.00	0.00 ± 0.00	$\mathbf{0.28} \pm 0.02$	$\mathbf{0.05} \pm 0.01$
OCBA	0.00 ± 0.00	0.00 ± 0.00	$\mathbf{0.33} \pm 0.01$	$\mathbf{0.07} \pm 0.01$

Monte Carlo Precheck

Computing $\mathrm{pEOC}_{(k)}$ on every iteration can be expensive.

- Little use in computing $\mathrm{pEOC}_{(k)}$ when it's well above β.

Idea: Use Monte Carlo to **"precheck"** the pEOC stopping rule.

At each iteration:

1. Generate r random instances of $\boldsymbol{W}$ and compute

$$\widehat{\mathrm{pEOC}}_{(k)} = \frac{1}{r} \sum_{\ell=1}^{r} \left[W_{[k]}^{(\ell)} - W_{(k)}^{(\ell)} \right].$$

2. If $\widehat{\mathrm{pEOC}}_{(k)} < \beta$, compute $\mathrm{pEOC}_{(k)}$. Otherwise proceed to the next iteration.

This approach will not invalidate the Bayesian statistical guarantee.

Simulation Times (Replications)

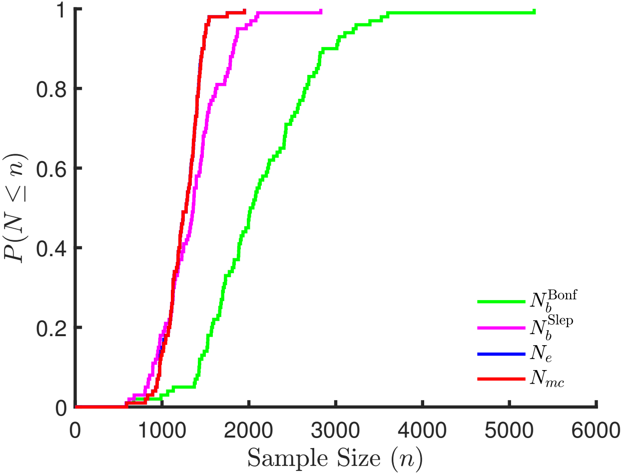

ECDF of the total sample size for the pEOC stopping rule with $\beta = 0.50$ and $n_0 = 5$ with the TS allocation.

Computation Times

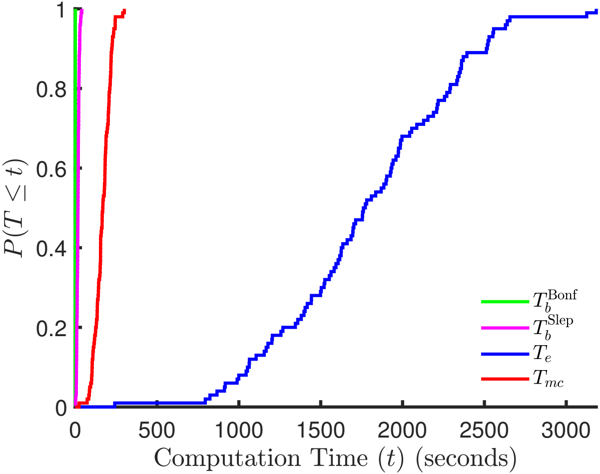

ECDF of the time spent checking the pEOC stopping rule with $\beta = 0.50$ and $n_0 = 5$ with the TS allocation.

Pure Monte Carlo

Idea: Directly check stopping rules using a Monte Carlo estimator.

Pros

- Sample sizes comparable to exactly checking the stopping rule.
- Computational times comparable to using bounds.

Cons

- Invalidates the Bayesian guarantee.

Practical jusification:

- Welch approximation.
- When assumptions of normally distributed outputs, independent sampling, and independent beliefs do not hold.

Pure Monte Carlo

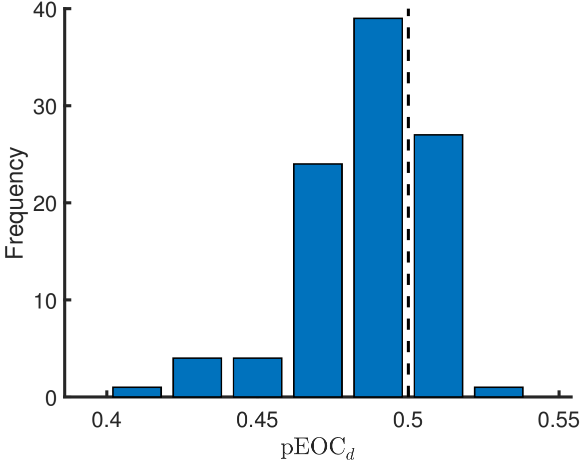

Histogram of pEOC_d upon termination when using the Monte Carlo estimator $\widehat{\mathsf{pEOC}}_{(k)}$ to directly check the stopping rule.

Takeaways

1. New cheaply computable Slepian bound on pEOC.
2. Savings vs. bounds on pPGS tend to be *minimal*.
3. Savings vs. bounds on pEOC can be *substantial*.
4. Using Monte Carlo precheck gives the best of both methods.
   - Smaller sample sizes with little extra computational time.
5. Pure Monte Carlo approach has practical appeal.

Acknowledgments

This material is based upon work supported by the Army Research Office under grant W911NF-17-1-0094 and by the National Science Foundation under grants DGE-1650441 and CMMI-1537394.

Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.