Guarantees on the Probability of Good Selection

David J. Eckman

Cornell University
Operations Research & Info Eng
dje88@cornell.edu

Shane G. Henderson

Cornell University
Operations Research & Info Eng
sgh9@cornell.edu

Winter Simulation Conference
December 10, 2018

Problem Setting

- Optimize a scalar performance measure over a finite number of alternatives.
- An alternative's performance is observed with simulation noise.

Examples:

Alternatives	Performance Measure
hospital bed allocations	expected blocking costs
ambulance base locations	expected call response time
MDP policy	expected discounted total cost

Two alternatives $\longrightarrow$ A/B testing.

More than two alternatives $\longrightarrow$ ranking and selection and exploratory MAB.

Selection of the Best in Software

E.g., Simio.

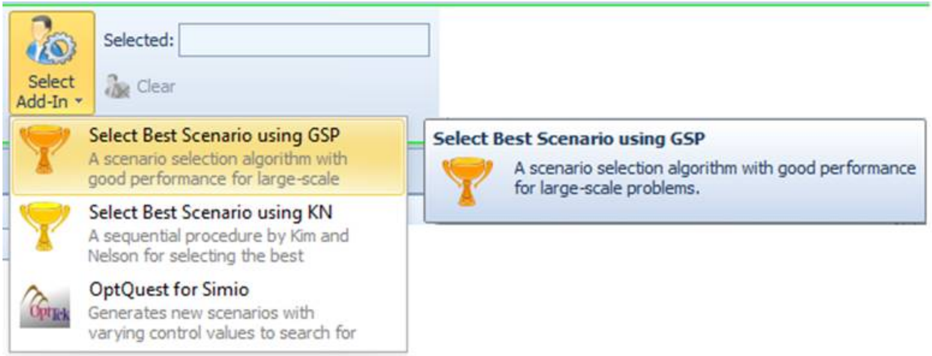

Model

Alternative 1	X_{11}	X_{12}	$\cdots$	i.i.d. $\sim F_1$ with mean θ_1
Alternative 2	X_{21}	X_{22}	$\cdots$	i.i.d. $\sim F_2$ with mean θ_2
$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$
Alternative k	X_{k1}	X_{k2}	$\cdots$	i.i.d. $\sim F_k$ with mean θ_k

Observations across alternatives are independent, unless CRN are used.

Marginal distributions F_i:

- Ranking and selection (R&S): normal (via batching + CLT)
- Multi-armed bandits: bounded support or sub-Gaussian with known variance bound

The vector $\boldsymbol{\theta} = (\theta_1, \theta_2, \ldots, \theta_k)$ represents the (unknown) problem instance.

- Assume that larger θ_i is better.

Selection Events

Let $\mathscr{D}$ be the index of the selected alternative.

- **Correct Selection:** *"Select one of the best alternatives."*

$$\text{CS} := \{\theta_{\mathscr{D}} = \theta_{[k]}\}.$$

- **Good Selection:** *"Select a δ-good alternative."*

$$\text{GS} := \{\theta_{\mathscr{D}} > \theta_{[k]} - \delta\}.$$

where $\theta_{[1]} \leq \theta_{[2]} \leq \cdots \leq \theta_{[k]}$ are the ordered mean performances.

Here δ represents the decision-maker's tolerance toward making a suboptimal selection.

"Close enough is good enough."

Fixed-Confidence Guarantees

Guarantee that a certain selection event occurs with high probability:

$$\mathbb{P}(\text{GS}) \text{ (or } \mathbb{P}(\text{CS})) \geq 1 - \alpha,$$

where $1 - \alpha$ is specified by the decision-maker.

Guarantee on PGS (PAC Selection)

$$\underbrace{\text{W.p. } 1-\alpha,}_{\textbf{Probably}} \text{ Alternative } \mathscr{D} \text{ is } \underbrace{\text{within } \delta}_{\textbf{Approximately}} \underbrace{\text{of the best}}_{\textbf{Correct}}.$$

Expected Opportunity Cost

Another popular criteria is the expected opportunity cost (EOC)—a.k.a. linear loss.

$$\mathbb{E}[\mathscr{L}_{OC}] = \mathbb{E}[\theta_{[k]} - \theta_{\mathscr{D}}].$$

EOC can give a loose upper bound on PGS via Markov's inequality:

$$\mathbb{P}(\text{GS}) = 1 - \mathbb{P}(\theta_{[k]} - \theta_{\mathscr{D}} \geq \delta) \geq 1 - \frac{\mathbb{E}[\theta_{[k]} - \theta_{\mathscr{D}}]}{\delta} = 1 - \frac{\mathbb{E}[\mathscr{L}_{OC}]}{\delta}.$$

- EOC can be harder for a decision-maker to interpret or quantify.
- EOC is commonly studied under a Bayesian framework.

Indifference-Zone Formulation

Bechhofer (1954) developed the idea of an indifference zone (IZ).

For an IZ parameter $\delta > 0$:

- **Preference Zone**: $\text{PZ}(\delta) = \{\boldsymbol{\theta} : \theta_{[k]} - \theta_{[k-1]} \geq \delta\}$
  "The best alternative is at least δ better than all the others."
- **Indifference Zone**: $\text{IZ}(\delta) = \{\boldsymbol{\theta} : \theta_{[k]} - \theta_{[k-1]} < \delta\}$
  "There are close competitors to the best alternative."

The parameter δ is described as the *smallest difference in performance worth detecting.*

- ...but that's not its role in the IZ formulation.

Space of Configurations

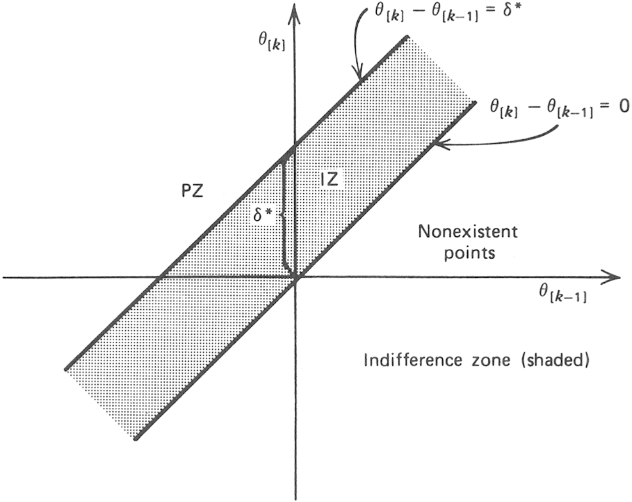

Goals of R&S Procedures

Two Frequentist Guarantees

Specify confidence level $1 - \alpha \in (1/k, 1)$ and $\delta > 0$ and guarantee

$$\mathbb{P}_{\boldsymbol{\theta}}(\text{CS}) \geq 1 - \alpha \quad \text{for all } \boldsymbol{\theta} \in \text{PZ}(\delta), \qquad \text{(Goal PCS-PZ)}$$

$$\mathbb{P}_{\boldsymbol{\theta}}(\text{GS}) \geq 1 - \alpha \quad \text{for all } \boldsymbol{\theta}. \qquad \text{(Goal PGS)}$$

$$\text{Goal PGS} \implies \text{Goal PCS-PZ}.$$

Goal PCS-PZ is the standard in the frequentist R&S community.

Goal PCS-PZ vs Goal PGS

Issues with Goal PCS-PZ

- Says nothing about a procedure's performance in IZ(δ).
- Configurations in PZ(δ) may be unlikely in practice:
  - when there are a large number of alternatives, or
  - when alternatives found by a search.
- Choice of δ restricts the problem.
- May require making Bayesian assumptions about $\boldsymbol{\theta}$.

Goal PGS has none of these issues!

Proving Goal PGS

Several ways to prove Goal PGS:

1. Lifting Goal PCS-PZ
2. Concentration inequalities
3. Multiple comparisons

Equivalence of Goals PCS-PZ and PGS

"When does Goal PCS-PZ $\implies$ Goal PGS?"

Intuition: More good alternatives $\implies$ more likely to pick a good alternative.

Scattered results since Fabian (1962), but none in the past 20 years.

Show that some R&S procedures delivering Goal PCS-PZ also deliver Goal PGS.

Equivalence Results: Condition 1

Condition 1 (Guiard 1996)

For all subsets $A \subset \{1, \ldots, k\}$, the joint distribution of the *estimators* of θ_i for $i \in A$ does not depend on θ_j for all $j \notin A$.

"Changing the mean of an alternative doesn't change the distribution of the estimators of other alternatives' means."

Limitation: Can only be applied to procedures without screening.

- Normal (i.i.d.): Bechhofer (1954), Dudewicz and Dalal (1975), Rinott (1978)
- Normal (CRN): Clark and Yang (1986), Nelson and Matejcik (1995)
- Bernoulli: Sobel and Huyett (1957)
- Support $[a, b]$: Naive Algorithm of Even-Dar et al. (2006)

Equivalence Results: Condition 2

Condition 2 (Hayter 1994)

For all alternatives $i = 1, \ldots, k$,

$$\mathbb{P}_{\boldsymbol{\theta}}(\text{Select Alternative } i)$$

is nonincreasing in θ_j for every $j \neq i$.

"Improving the mean of an alternative doesn't help any other alternative get selected."

Limitation: Checking the monotonicity of $\mathbb{P}_{\boldsymbol{\theta}}(\text{Select Alternative } i)$ is hard.

Equivalence Results: Condition 2

Procedure not satisfying Condition 2

1. Take n_0 samples of each alternative.
2. Eliminate all but the two alternatives with the highest means.
3. Take n_1 additional samples for the two surviving alternatives.
4. Select the surviving alternative with the highest overall mean.

Consider the three-alternative case: $\theta_1 < \theta_2 < \theta_3$.

- Track $\mathbb{P}_{\boldsymbol{\theta}}(\text{Select Alternative } 2)$ as θ_1 increases up to θ_2.
- Fix $n_0 \geq 1$ and consider $n_1 = 0$ and $n_1 = \infty$ as extreme cases.

Equivalence Results: Condition 3

Condition 3

For all alternatives $i = 1, \ldots, k$,

$$\mathbb{P}_{\boldsymbol{\theta}}(\text{Select some alternative, } j \text{, for which } \theta_j < \theta_i)$$

is nonincreasing in θ_i.

"Improving the mean of an alternative doesn't help inferior alternatives get selected."

Condition 2 $\implies$ Condition 3.

Sampling Efficiency

Sequential selection procedures screen out (eliminate) inferior systems.

- They are among the most efficient at delivering Goal PCS-PZ.

"Do the procedures of Kim and Nelson (2001) and Frazier (2014) achieve Goal PGS?"

Even if they do, they may be inefficient for problem instances in the IZ.

There may be an opportunity to design more efficient procedures delivering Goal PGS.

Concentration Inequalities

Regularity conditions for multi-armed bandits enable the use of confidence inequalities.

- E.g., Hoeffding and Chernoff bounds.

General approach:

1. Bound the probability an estimator differs from its mean value by at least $\delta/2$.
2. Use Bonferroni's inequality to sum over all alternatives.

The Envelope Procedure of Ma and Henderson (2017) uses confidence bands that hold *throughout the entire procedure*.

- Tracks upper and lower confidence limits for each alternative's mean performance.

Multiple Comparisons

Let Y_i be the estimator of the mean performance θ_i.

Assume that the selected alternative is $\mathscr{D} = \arg\max_{i=1,\ldots,k} Y_i$.

Multiple Comparisons with the Best (MCB)

$$\mathscr{B} = \{Y_i - Y_{[k]} - (\theta_i - \theta_{[k]}) < \delta, \ \forall i \neq [k]\}$$

$$\begin{aligned} \mathbb{P}_{\boldsymbol{\theta}}(\mathscr{B}) \geq 1 - \alpha \implies & \mathbb{P}_{\boldsymbol{\theta}}\{Y_{\mathscr{D}} - Y_{[k]} - (\theta_{\mathscr{D}} - \theta_{[k]}) < \delta\} \geq 1 - \alpha, \\ \implies & \mathbb{P}_{\boldsymbol{\theta}}(\text{GS}) \geq 1 - \alpha. \end{aligned}$$

Deriving Goal PGS from MCB results in a *conservative* selection procedure.

Frequentist and Bayesian Frameworks

Different perspectives on what is random and what is fixed.

Frequentist

PGS = The probability that the random alternative chosen by the procedure is good for the fixed problem instance.

Bayesian

PGS = The ***posterior*** probability that—given the observed data—the random problem instance is one for which the fixed alternative chosen by the procedure is good.

"How do these guarantees differ on a practical level?"

Design for Frequentist PGS

Design the procedure to satisfy the PGS guarantee for the least favorable configuration (LFC), i.e., the hardest problem instance.

The LFC is often the so-called slippage configuration (SC).

- Fix a best alternative, j, and set $\theta_i = \theta_j - \delta$ for all $i \neq j$.

Frequentist procedures are conservative: they often overdeliver on PGS.

Frequentist PGS

Ex: Two alternatives with observations $X_{1j} \sim N(\theta_1, \sigma^2)$ and $X_{2j} \sim N(\theta_2, \sigma^2)$ for $j = 1, \ldots, n$ where σ^2 is known.

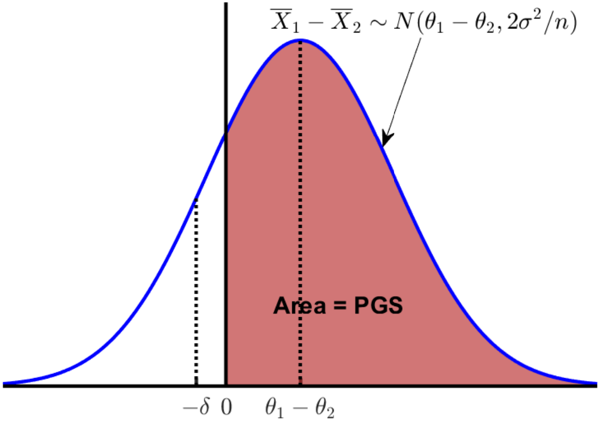

Design for Bayesian Guarantees

Stopping Rule Principle

It is valid to stop and select an alternative whenever its *posterior* PGS exceeds $1 - \alpha$.

Can use posterior PGS as a stopping rule for a variety of procedures:

- E.g., VIP, OCBA, and TTTS.

Advantages:

- Can *repeatedly compute* posterior PGS without sacrificing statistical validity.
- Complete flexibility in allocating simulation runs across alternatives.

Bayesian PGS

Ex: Two alternatives with observations $X_{1j} \sim N(\theta_1, \sigma^2)$ and $X_{2j} \sim N(\theta_2, \sigma^2)$ for $j = 1, \ldots, n$ where σ^2 is known, with a noninformative prior on $\theta_1 - \theta_2$.

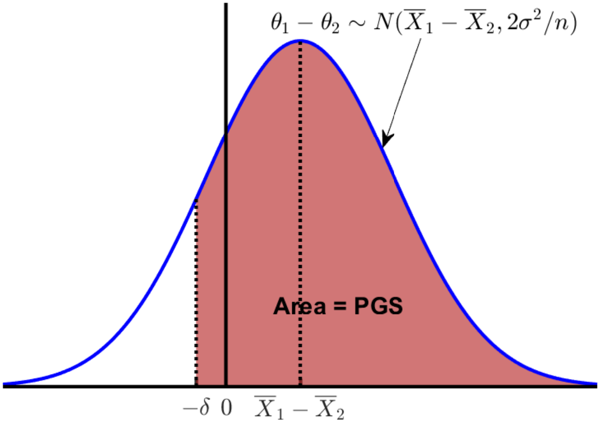

Continuation Regions

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

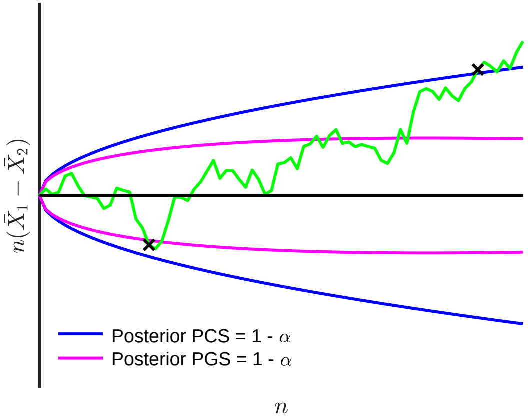

Interpreting Bayesian Guarantees

A Bayesian PGS guarantee will **NOT** deliver a frequentist guarantee that PGS exceeds $1-\alpha$ **for all** problem instances.

Its guarantee can still be interpreted in a frequentist sense.

1. Draw θ from the prior distribution.
2. Run the Bayesian procedure (with the stopping rule) on θ.

For repeated runs of Steps 1 and 2, the procedure will make a good selection w.p. $1-\alpha$.

Experimental Results

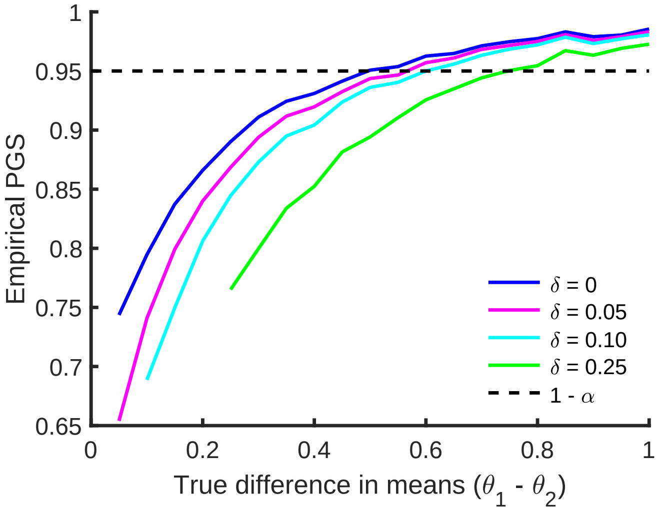

Observations

1. For hard problem instances, procedures with Bayesian PGS guarantees underdeliver on empirical PGS.
   - Gap becomes more pronounced for more tolerant good selection.
2. Hard problems look easier because of a "means-spreading" phenomenon.
   - Similar issues arise in predicting the runtime of a procedure.

Practical Implications

A decision-maker's preference may depend on the situation:

1. A one-time, critical decision.
2. Repeated problem instances (i.e., using R&S for control).
3. R&S after search, where the problem instance is random.

Computational Considerations

Bayesian procedures with fixed-confidence guarantees pose computational challenges.

1. Checking whether the posterior PGS stopping condition has been met.
2. Calculating or estimating posterior PGS for a given alternative.

Setup:

- Assume that observations are normally distributed and i.i.d.
- Assume a multivariate normal prior with independent beliefs.
- Let W_i denote the (random) mean performance of Alternative i.

The posterior distribution of $\boldsymbol{W} = (W_1, \ldots, W_k)$ is a multivariate normal *(if variances are known)* or a multivariate t *(if variances are unknown)* distribution.

Computing Posterior PGS

The posterior PGS of Alternative i is

$$\text{pPGS}_i = \mathbb{P}(W_i > W_j - \delta, \text{ for all } j \neq i \mid \mathcal{E}),$$

where $\mathbb{P}(\cdot \mid \mathcal{E})$ is the probability under the posterior of $\boldsymbol{W}$ given the evidence $\mathcal{E}$.

When there are k alternatives, this amounts to a k-dimensional integral.

- Becomes intractable for large k, unless we condition on W_i.

Conditioning on W_i leads to a one-dimensional integral:

$$\text{pPGS}_i = \mathbb{E}\left[\prod_{j \neq i} \mathbb{P}(W_i > W_j - \delta \mid W_i, \mathcal{E}) \mid \mathcal{E}\right].$$

Slepian's Bound on Posterior PGS

Slepian's inequality can be used to get a *cheap* lower bound on posterior PGS.

$$\text{pPGS}_i = \mathbb{P}(W_i > W_j - \delta, \text{ for all } j \neq i \mid \mathcal{E}) \geq \prod_{j \neq i} \mathbb{P}(W_i > W_j - \delta \mid \mathcal{E}) =: \text{pPGS}_i^{\text{Slep}}.$$

Terminate the first time any $\text{pPGS}_i^{\text{Slep}}$ exceeds $1 - \alpha$ and select that alternative.

As k increases, the tightness of Slepian's bound deteriorates.

- Appears to deteriorate slower for values of PGS close to 1.
- Using it as a stopping condition will lead to longer run-lengths than necessary.

Extensions: PGS for Continuous Optimization

Embed the R&S problem in a continuous domain $\mathcal{D}$ with objective function $\theta : \mathcal{D} \mapsto \mathbb{R}$.

Assume some structural property of θ, e.g., convex or Lipschitz continuous.

Goal PGS

Select a (random) solution $x_{\mathscr{D}} \in \mathcal{D}$ such that

$$\mathbb{P}(\theta(x_{\mathscr{D}}) > \theta(x^*) - \delta) \geq 1 - \alpha,$$

where $x^* \in \arg\max_{x \in \mathcal{D}} \theta(x)$.

See Nesterov and Vial (2008), for example.

Extensions: Good Subset Selection

Instead of selecting a single alternative, return a subset of alternatives, $\mathcal{S}$.

Two main purposes:

1. Make a final selection based on secondary performance measures.
2. Use the subset as input to a selection procedure.

Under the frequentist framework, good subset selection is defined as

$$\text{GSS} = \{\exists i \in \mathcal{S} \text{ s.t. } \theta_i \geq \theta_{[k]} - \delta\}.$$

Is this the right definition of a “good” subset?

Extensions: Good Subset Selection

Under the Bayesian framework, good subset selection is defined as

$$\text{GSS} = \{\exists i \in \mathcal{S} \text{ s.t. } W_i \geq W_{[k]} - \delta\}.$$

Bayesian subset selection can be done at any time.

- Can calculate $\text{pPGSS}_{\mathcal{S}}$ for any subset $\mathcal{S}$, but it's computationally expensive.
- Selecting the smallest $\mathcal{S}$ such that $\text{pPGSS}_{\mathcal{S}} \geq 1 - \alpha$ is challenging.

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.

Equivalence of Goals PCS-PZ and PGS

Key Approach

Pair each $\boldsymbol{\theta} \in \text{IZ}(\delta)$ with a $\boldsymbol{\theta}^* \in \text{PZ}(\delta)$ and show that

$$\mathbb{P}_{\boldsymbol{\theta}}(\text{GS}) \geq \mathbb{P}_{\boldsymbol{\theta}^*}(\text{GS}) = \mathbb{P}_{\boldsymbol{\theta}^*}(\text{CS}) \geq 1 - \alpha.$$

Constructing $\boldsymbol{\theta}^*$

For an arbitrary configuration $\boldsymbol{\theta} \in \mathrm{IZ}(\delta)$, define subsets

$$\mathcal{G} = \{i : \theta_i > \theta_k - \delta\} \text{ "good" } \quad \text{and}$$
$$\mathcal{B} = \{i : \theta_i \leq \theta_k - \delta\} \text{ "bad." }$$

Define the configuration $\boldsymbol{\theta}^*$ by

$$\theta_i^* = \begin{cases} \theta_i & \text{for } i \in \mathcal{B} \cup \{k\}, \\ \theta_k - \delta & \text{for } i \in \mathcal{G} \setminus \{k\}. \end{cases}$$

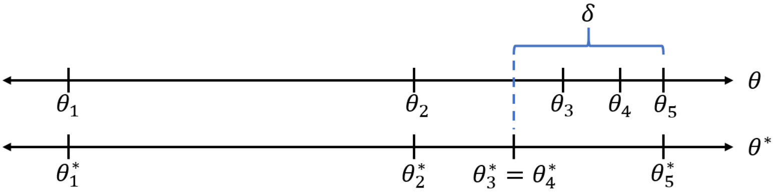

Sketch Proof of Condition 1

Assume ties in estimators Y_i occur with probability zero.

Fix an arbitrary configuration $\boldsymbol{\theta}$ and define $\mathcal{G}$, $\mathcal{B}$, and $\boldsymbol{\theta}^*$ accordingly.

$$\begin{aligned}
\mathbb{P}_{\boldsymbol{\theta}}(\text{GS}) &\geq \mathbb{P}_{\boldsymbol{\theta}}(Y_k > Y_i \text{ for all } i \in \mathcal{B}) \\
&= \mathbb{P}_{\boldsymbol{\theta}^*}(Y_k^* > Y_i^* \text{ for all } i \in \mathcal{B}) && (\star) \\
&\geq \mathbb{P}_{\boldsymbol{\theta}^*}(Y_k^* > Y_i^* \text{ for all } i \neq k) \\
&= \mathbb{P}_{\boldsymbol{\theta}^*}(\text{CS}) \\
&\geq 1 - \alpha.
\end{aligned}$$

($\star$) Condition 1 with $A = \mathcal{B} \cup \{k\}$. Note that $\theta_i^* = \theta_i$ for all $i \in \mathcal{B}$.

Sketch Proof of Condition 2

Fix an arbitrary configuration $\boldsymbol{\theta}$.

Repeatedly shift the mean performance of the *worst good* alternative down to $\theta_k - \delta$.

Each time, PGS is reduced.

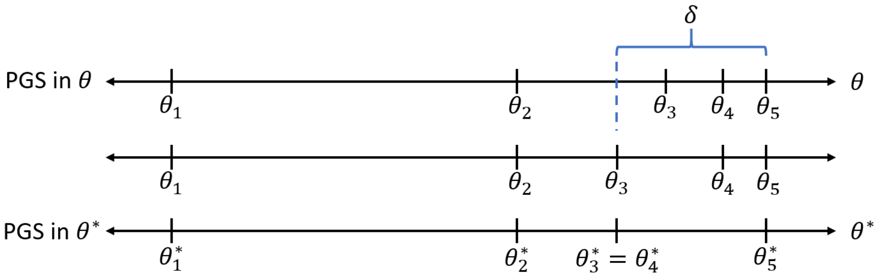

Final result: $\mathbb{P}_{\boldsymbol{\theta}}(\text{GS}) \geq \mathbb{P}_{\boldsymbol{\theta}^*}(\text{GS})$.