Indifference-Zone Ranking and Selection For More Than 10,000 Alternatives

Peter I. Frazier pf98@cornell.edu School of Operations Research & Information Engineering, Cornell University

Cornell University Operations Research and Information Engineering

Contribution

We construct a fully sequential elimination procedure for ranking and selection called the **Bayes-inspired IZ (BIZ)** procedure.

1. BIZ has the IZ guarantee:

$$\mathrm{PCS}(\pi,\theta) \geq P^*, \quad \text{for all } \theta \in \mathrm{PZ}(\delta).$$

2. The lower bound on the performance of BIZ is tight in continuous time:

$$\inf_{\theta\in\mathrm{PZ}(\delta)} \mathrm{PCS}(\pi,\theta) = P^*.$$

3. When the number of alternatives is large, BIZ samples much less than existing IZ procedures in numerical comparisons.

The Ranking & Selection (R&S) Problem

We have k alternatives or systems, e.g., different ways of operating a supply chain, that can be simulated. Each time we simulate alternative x, we observe an independent sample

$$y \sim \mathrm{Normal}(\theta_x, \sigma^2)$$

where θ_x is unknown. Here we assume σ^2 is known and constant, but the results can be generalized to heterogeneous unknown sampling variance.

Goal: use simulation efficiently to find the best alternative (the one with the largest θ_x).

Ranking and Selection is a fundamental form of **simulation optimization**.

Indifference Zone Guarantees on R&S Procedures

The Indifference Zone (IZ) guarantee is a **statistical guarantee on the performance** of a ranking and selection procedure.

The probability of correct selection $\mathrm{PCS}(\pi,\theta)$ is the probability that the procedure π selects the best alternative (the one with the largest θ_x).

The preference zone is the collection of system configurations where the best is better than the second best by at least $\delta > 0$:

$$\mathrm{PZ}(\delta) = \left\{\theta \in \mathbb{R}^k : \theta_{[1]} - \theta_{[2]} \geq \delta\right\}.$$

The indifference zone (IZ) are those system configurations outside the preference zone. A procedure π has an **IZ guarantee** with parameters δ and P^* if

$$\mathrm{PCS}(\pi,\theta) \geq P^* \quad \text{for all } \theta \in \mathrm{PZ}(\delta).$$

An Open Problem for 60 Years

Almost all IZ procedures created over the last 60 years share a common issue: **More samples are taken than required to meet the IZ guarantee**. This problem is especially pronounced when the number of alternatives is large. **Open Problem**: How can we design a procedure that takes no more samples than is required to meet the IZ guarantee?

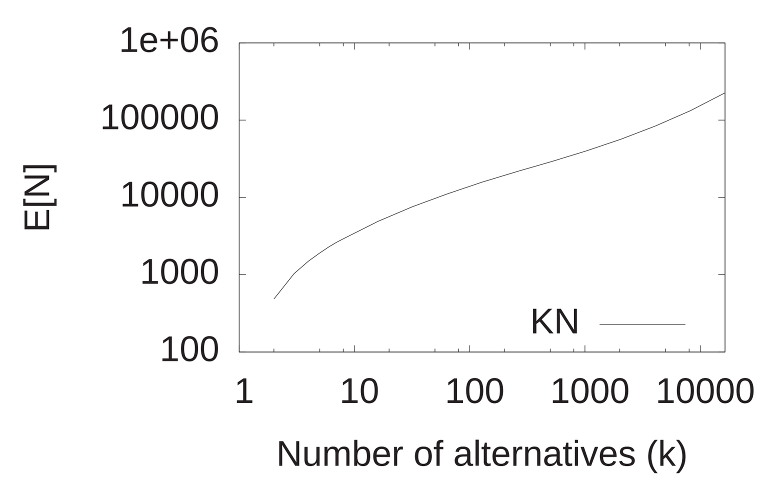

(a) Number of samples

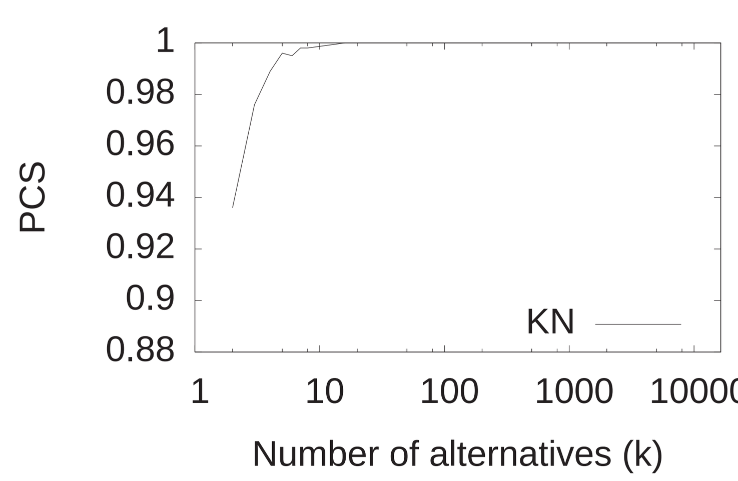

(b) Probability of correct selection

Figure: Monotone-decreasing-means configuration $\theta = [-\delta, -2\delta, \ldots, -k\delta]$ with $P^* = 0.9$, $\sigma^2 = 100$, $\delta = 1$, with the KN procedure of [Kim and Nelson, 2001] modified for known σ^2.

The BIZ Procedure

The Bayes-inspired IZ (BIZ) procedure is inspired by a Bayesian analysis, although the indifference-zone theoretical results are non-Bayesian.

Bayesian origin:
Let Q be a prior probability measure concentrated on least-favorable configurations.

$$X_* \sim \mathrm{Uniform}(1,\ldots,k), \qquad \theta_x = \begin{cases} \delta & \text{if } x = X_* \\ 0 & \text{if } x \neq X_* \end{cases}$$

The posterior probability under Q that x is best, given that $X_* \in A$, is

$$q_{tx}(A) = \exp\left(\frac{\delta}{\sigma^2} Y_{tx}\right) \Big/ \sum_{x' \in A} \exp\left(\frac{\delta}{\sigma^2} Y_{tx'}\right).$$

where Y_{tx} is the sum of all observations from alternative x by time t.

BIZ Summary: The BIZ procedure maintains a set of alternatives A_n that are "in contention". An alternative is eliminated from this set when its posterior probability $q_{tx}(A_n)$ drops below a lower threshold c. BIZ stops when the posterior probability of the best alternative exceeds an upper threshold P_n.

BIZ Procedure (Discrete-Time Version): Fix an elimination threshold $c \leq 1 - (P^*)^{1/(k-1)}$.

1. Let $A \leftarrow \{1,\ldots,k\}$, $t \leftarrow 0$, $P \leftarrow P^*$.
2. While $\max_{x\in A} q_{tx}(A) < P$
   - 2a. While $\min_{x\in A} q_{tx}(A) \leq c$
     - $*$ Let $z \in \arg\min_{x\in A} q_{tx}(A)$.
     - $*$ Let $P \leftarrow P/(1 - q_{tz}(A))$.
     - $*$ Eliminate z from A.
   - 2b. Sample from each $x \in A$ to obtain $Y_{t+1,x}$. Then increment t.
3. Select $\hat{x} \in \arg\max_{x\in A} q_{tx}(A)$ as our estimate of the best.

This discrete-time version of BIZ can be generalized to continuous time and heterogeneous unknown sampling variance.

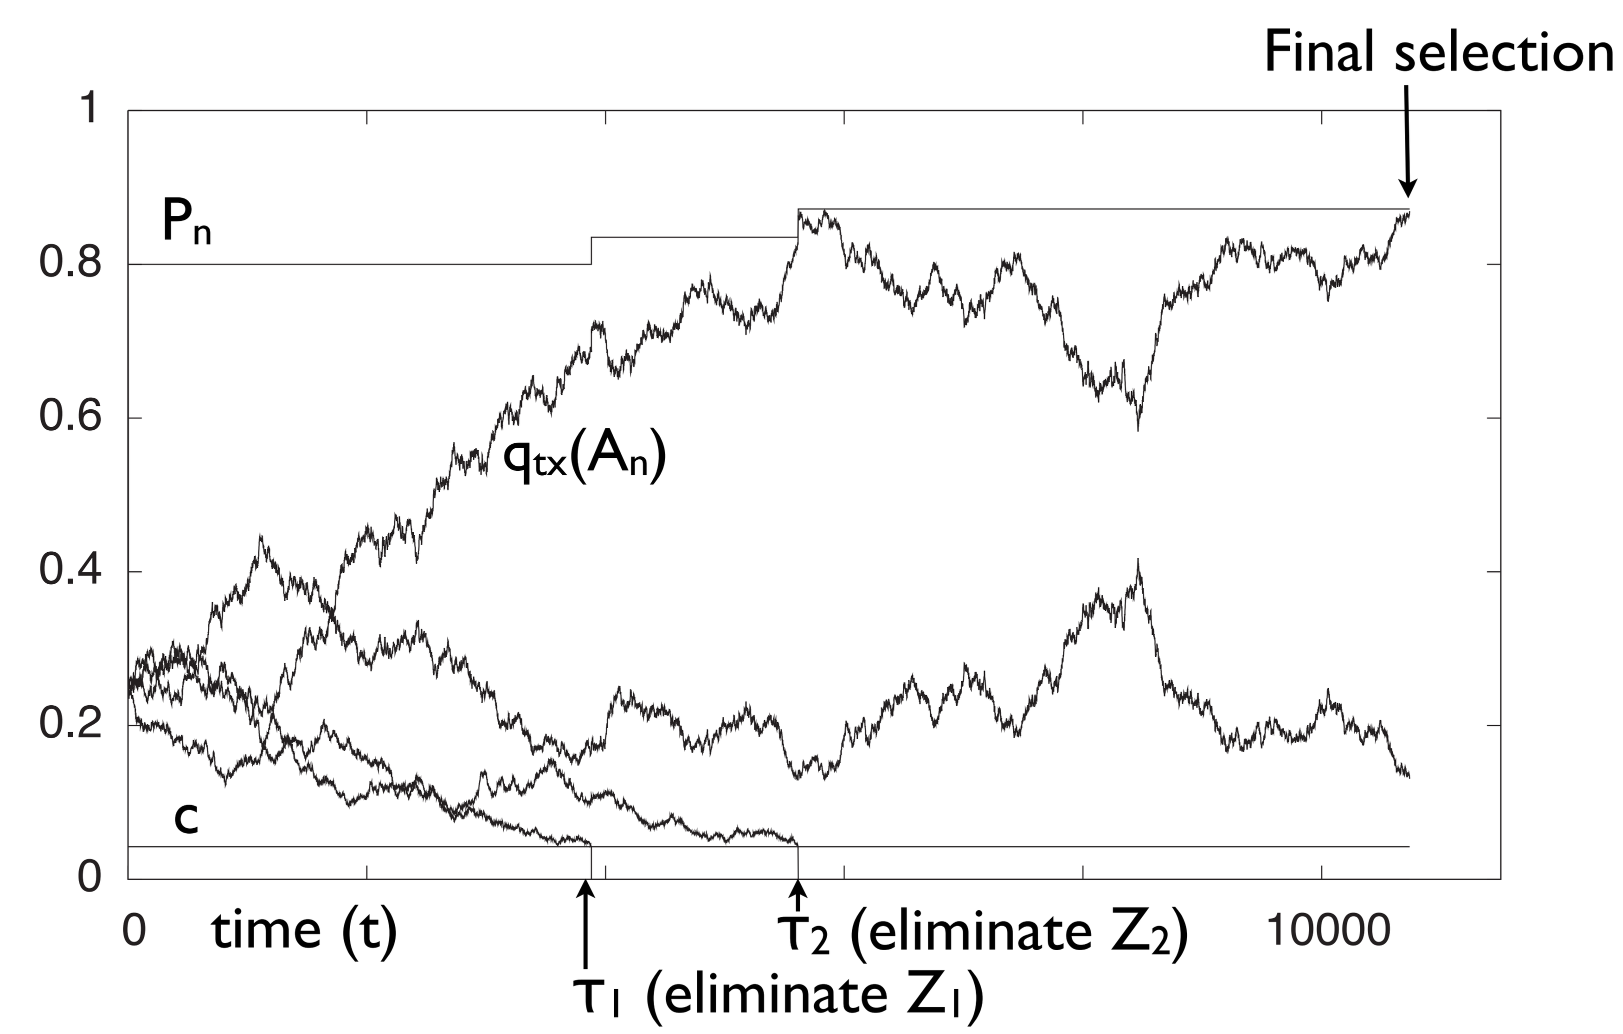

Example: In the example above, we start with all alternatives in contention, $A_0 = \{1,2,3,4\}$. At time τ_1, we eliminate the worst alternative because its posterior probability drops below c. The alternative eliminated is $Z_1 = 2$, and the new set of alternatives in contention is $A_1 = \{1,3,4\}$. At time τ_2 we eliminate $Z_2 = 3$, leaving $A_2 = \{1,4\}$. Eventually, the posterior probability of alternative 1 exceeds the upper threshold P_2 and we select it as the best.

Theoretical Results: IZ Guarantees with Tight Bounds

Theorem 1: The BIZ procedure **satisfies the IZ guarantee**, i.e.,

$$\mathrm{PCS}(\mathrm{BIZ},\theta) \geq P^* \quad \forall \theta \in \mathrm{PZ}(\delta)$$

Theorem 2: If $\mathbb{T} = \mathbb{R}_+$, i.e., sampling occurs in continuous time, then **the lower bound on probability of correct selection is tight**, i.e.,

$$\inf_{\theta\in\mathrm{PZ}(\delta)} \mathrm{PCS}(\mathrm{BIZ},\theta) = P^*$$

These theoretical results assume common known sampling variance.

Numerical Experiments: Efficient Sampling

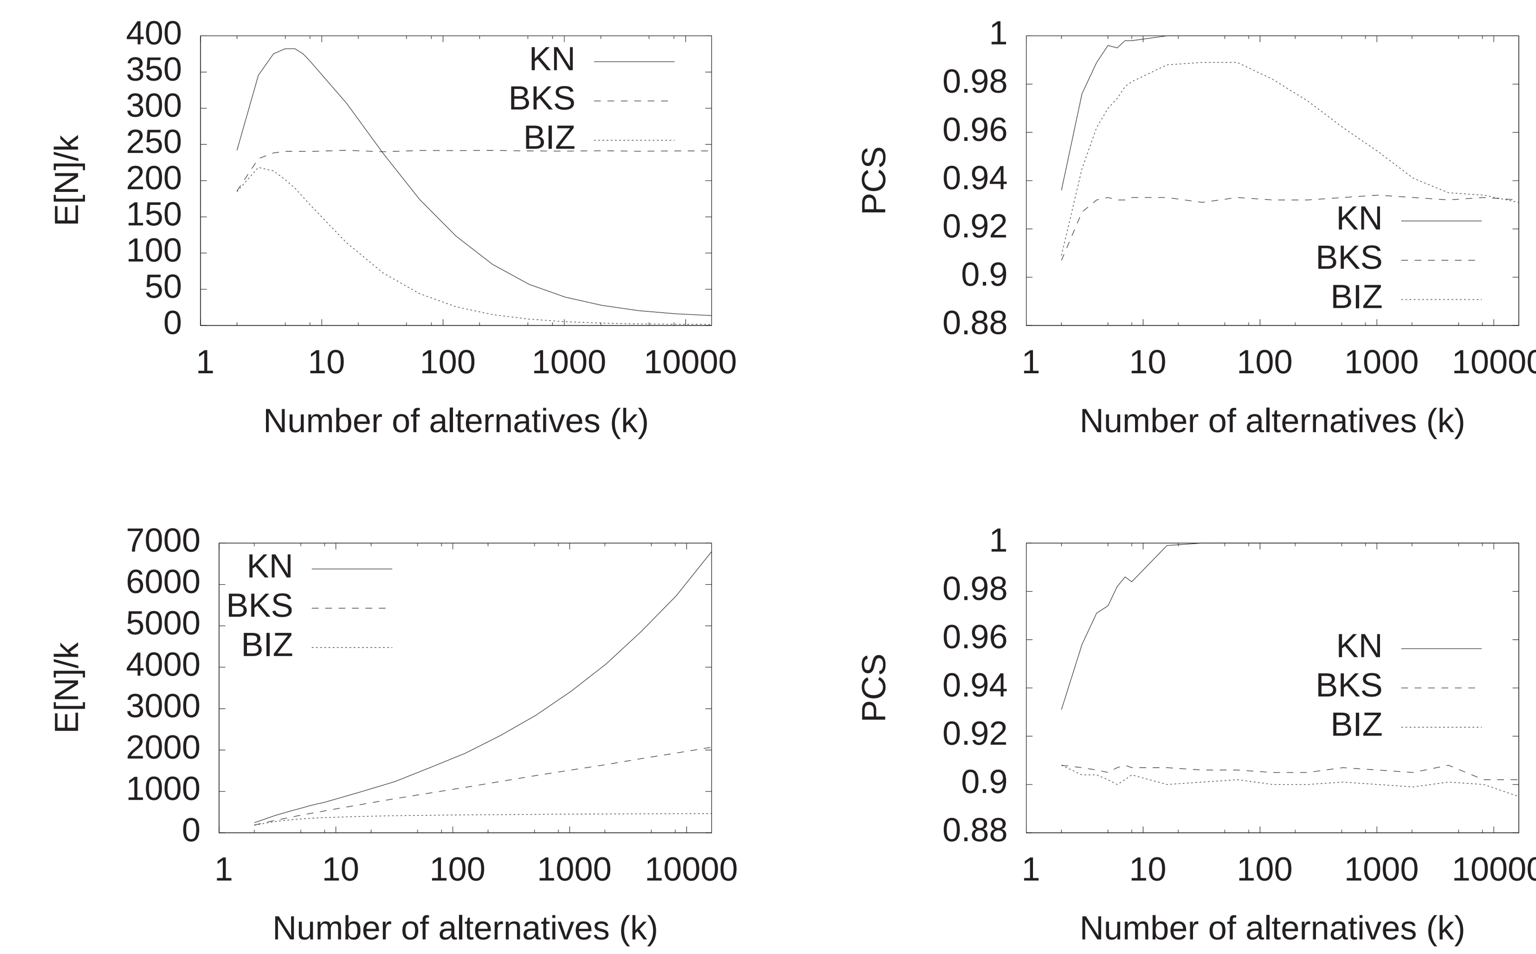

(top) **Monotone decreasing means configuration**, $\theta = [-\delta, -2\delta, \ldots, -k\delta]$. (bottom) **Slippage configuration**, $\theta = [0,\ldots,0,\delta]$. In both configurations, $P^* = 0.9$, $\sigma = 10$, $\delta = 1$, estimated with $\geq 10{,}000$ independent replications. BIZ uses $c = 1 - (P^*)^{1/(k-1)}$ (eliminate aggressively). BKS is $\mathcal{P}_B^*$ ($c = 0$) from [Bechhofer et al., 1968], which is a special case of BIZ with $c = 0$ and does not eliminate.

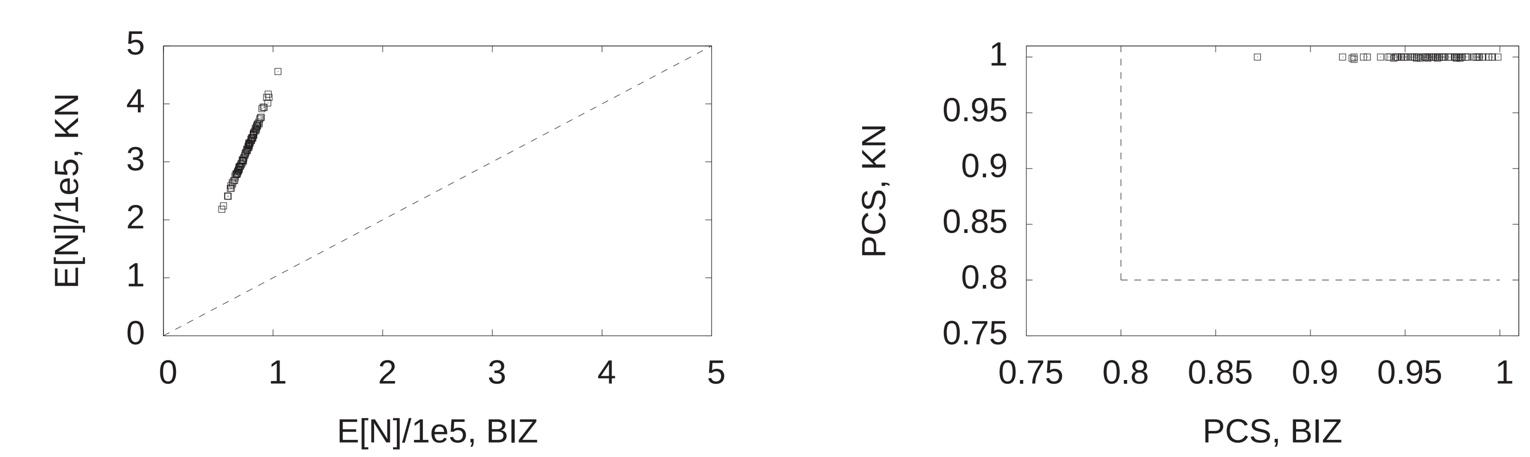

Random Problem Instances: $k = 100$, $P^* = 0.8$, θ was generated randomly from an independent normal prior and then adjusted to lie in $\mathrm{PZ}(\delta)$.

Funding

This work was supported in part by the Air Force Office of Scientific Research Young Investigator Program, #FA9550-11-1-0083.

References

[Frazier, 2011] Frazier, P.I. Indifference-Zone Ranking and Selection: Many Alternatives and Tight Bounds on Performance. In preparation

[Kim and Nelson, 2001] Kim, S.-H. and Nelson, B. L. (2001). A fully sequential procedure for indifference-zone selection in simulation. *ACM Trans. Model. Comput. Simul.*, 11(3):251–273.

[Bechhofer et al., 1968] Bechhofer, R., Kiefer, J., and Sobel, M. (1968). *Sequential Identification and Ranking Procedures*. University of Chicago Press, Chicago.