NEW DEVELOPMENTS IN RANKING AND SELECTION: An Empirical Comparison of Three Main Approaches

Jürgen Branke[2] Stephen E. Chick[1] Christian Schmidt[2]

[1](speaker)
Technology Management Area
INSEAD
Fontainebleau, France

[2]Institut AIFB
Universität Karlsruhe (TH)
D-76128 Karlsruhe, GERMANY

2005 Winter Simulation Conference

INSEAD

Selecting the Best of a Finite Set

1. There are a plethora of ranking and selection approaches
   - Indifference zone, VIP, OCBA, ETSS, ...
   - Each approach has variations, parameters, approximations leading to different allocation, stopping and selection rules
   - Optimizations more demanding of such procedures
2. Today: Which *sequential* selection procedure is "best" (given independent, Gaussian samples, unknown means/variances).
   - New procedures (stopping rules, allocations)
   - New measures and mechanisms to evaluate procedures
   - Summarize observations from what is believed to be the largest numerical experiment to date
   - Identify strengths/weaknesses of leading procedures

See also *Selecting a Selection Procedure* Branke, Chick, and Schmidt (2005), more allocations, experiments, ...

INSEAD

Outline

1. Overview for Ranking and Selection
   - What are Measures of a Good Procedure?
   - Problem Formulation
   - Evidence for Correct Selection and New Stopping Rules
   - Procedures Tested
2. Empirical Evaluation
   - Empirical Figures of Merit
   - Numerical Test Bed
   - Implementation
3. Summary of Qualitative Conclusions
   - Stopping Rules
   - Allocations
   - General Comments
4. General Summary
   - Which procedure to use?
   - Discussion (time permitting)

What are measures of a good procedure?

- Utopia: always find true best with zero effort.
  - Fact: Variability implies incorrect selections or infinite work.
- Theoretical properties:
  - Derivations are preferred to ad hoc approximations
  - Reasonable people may choose different assumptions
- Empirical properties:
  - Efficiency: Mean evidence for correct selection as function of mean number of samples
  - Controllability: Ease of setting parameters to achieve a targeted evidence level
  - Robustness: Dependency of procedure's effectiveness on underlying problem characteristics
  - Sensitivity: Effect of parameters on mean number of samples

INSEAD

Problem formulation

- Identify best of k systems (biggest mean).
- Let X_{ij} be output of jth replication of ith system:

$$\{X_{ij} : j = 1, 2, \ldots\} \overset{i.i.d.}{\sim} \texttt{Normal}\left(w_i, \sigma_i^2\right), \text{ system } i = 1, \ldots, k.$$

- True (unknown) order of means: $w_{[1]} \leq w_{[2]} \leq \ldots \leq w_{[k]}$
- Configuration:

$$\chi = (\mathbf{w}, \boldsymbol{\sigma}^2).$$

- Samples statistics: $\bar{x}_i$ and $\hat{\sigma}_i^2$ updated based on n_i observations seen so far.
- Order statistics: $\bar{x}_{(1)} \leq \bar{x}_{(2)} \leq \ldots \leq \bar{x}_{(k)}$
- If select (k), then $\{w_{(k)} = w_{[k]}\}$ is a correct selection event

Evidence for Correct Selection

- Loss function if system $\mathfrak{D}$ is chosen when means are $\mathbf{w}$:
  - Zero-one: $\mathcal{L}_{0-1}(\mathfrak{D}, \mathbf{w}) = \mathbb{1}\left\{w_{\mathfrak{D}} \neq w_{[k]}\right\}$
  - Expected opportunity cost (EOC): $\mathcal{L}_{oc}(\mathfrak{D}, \mathbf{w}) = w_{[k]} - w_{\mathfrak{D}}$
- Frequentist measures (distribution of $\mathfrak{D} = f(\mathbf{X})$)

$$\begin{aligned} \mathrm{PCS}_{\mathrm{iz}}(\chi) &\stackrel{\mathrm{def}}{=} 1 - \mathrm{E}\left[\mathcal{L}_{0-1}(\mathfrak{D}, \mathbf{w}) \mid \chi\right] \\ \mathrm{EOC}_{\mathrm{iz}}(\chi) &\stackrel{\mathrm{def}}{=} \mathrm{E}\left[\mathcal{L}_{oc}(\mathfrak{D}, \mathbf{w}) \mid \chi\right] \end{aligned}$$

- Bayesian measures (given all output $\mathcal{E}$, $\mathfrak{D}$ and posterior of $\mathbf{W}$)

$$\begin{aligned} \mathrm{PCS}_{Bayes} &\stackrel{\mathrm{def}}{=} 1 - E\left[\mathcal{L}_{0-1}(\mathfrak{D}, \mathbf{W}) \mid \mathcal{E}\right] \\ \mathrm{EOC}_{Bayes} &\stackrel{\mathrm{def}}{=} E\left[\mathcal{L}_{oc}(\mathfrak{D}, \mathbf{W}) \mid \mathcal{E}\right] \end{aligned}$$

- Similar for PGS_{δ^*}, for "good" selections (within δ^* of best)

INSEAD

Bayesian Evidence and Stopping Rules

- Bounds (approximate) for Bayesian measures
  - Normalized distance: $d^*_{jk} = d_{(j)(k)}\lambda_{jk}^{1/2}$, where $d_{(j)(k)} = (\bar{x}_{(k)} - \bar{x}_{(j)})$ and $\lambda_{jk}^{-1} = \left(\frac{\hat{\sigma}^2_{(j)}}{n_{(j)}} + \frac{\hat{\sigma}^2_{(k)}}{n_{(k)}}\right)$.

$$\begin{aligned} \mathrm{PCS}_{Bayes} &\geq \prod_{j:(j)\neq(k)} \Pr\left(W_{(k)} > W_{(j)} \,\middle|\, \mathcal{E}\right) \text{ (Slepian)} \\ &\approx \prod_{j:(j)\neq(k)} \Phi_{\nu_{(j)(k)}}(d^*_{jk}) \stackrel{\text{def}}{=} \mathrm{PCS}_{Slep} \text{ (Welch)} \end{aligned}$$

  - $\mathrm{EOC}_{Bonf} = \sum_{j:(j)\neq(k)} \lambda_{jk}^{-1/2}\Psi_{\nu_{(j)(k)}}\left[d^*_{jk}\right]$. ("newsvendor" loss)
  - $\mathrm{PGS}_{Slep,\delta^*} = \prod_{j:(j)\neq(k)} \Phi_{\nu_{(j)(k)}}(\lambda_{jk}^{1/2}(\delta^* + d_{(j)(k)}))$.
  - $\mathrm{PCS}_{Slep,\delta^*} = \prod_{j:(j)\neq(k)} \Phi_{\nu_{(j)(k)}}(\lambda_{jk}^{1/2}\max\{\delta^*, d_{(j)(k)}\})$ (Chen and Kelton 2005).

Bayesian Evidence and Stopping Rules

- New "adaptive" stopping rules provide flexibility
  1. Sequential ($\mathcal{S}$): Repeat sampling if $\sum_{i=1}^{k} n_i < B$ for a given total budget B. [Default for most previous VIP and all OCBA work]
  2. Repeat if $\mathrm{PCS}_{Slep,\delta^*} < 1 - \alpha^*$ for a given δ^*, α^*.
  3. Repeat if $\mathrm{PGS}_{Slep,\delta^*} < 1 - \alpha^*$ for a given δ^*, α^*.
  4. Repeat if $\mathrm{EOC}_{Bonf} > \beta^*$, for an EOC target β^*.
- We use PCS_{Slep} to denote $\mathrm{PCS}_{Slep,0}$.

State-of-the-Art and New Procedures Tested

- Indifference-zone (IZ): $\mathcal{KN}++$ (Kim and Nelson 2001)
- OCBA Allocations with all stopping rules
  - Usual $\mathcal{OCBA}$ allocation (Chen 1996; PCS_{Slep} objective)
  - $\mathcal{OCBA}_{LL}$ for EOC_{Bonf} objective (He, Chick, and Chen 2005)
  - $\mathcal{OCBA}_{\delta*}$: Like $\mathcal{OCBA}$ but with $\mathrm{PGS}_{\delta*}$-allocation
  - $\mathcal{OCBA}_{max,\delta*}$: Like $\mathcal{OCBA}$, with max replacing $+$ in $\mathrm{PGS}_{\delta*}$-allocation (cf. Chen and Kelton 2005)
- VIP Allocations (Chick and Inoue 2001) with all stopping rules
  - Sequential $\mathcal{LL}$ allocation (for EOC_{Bonf} objective)
  - Sequential 0-1 allocation (for PCS_{Bonf} objective)
- Equal allocation with all stopping rules

- Names: Allocation(stop rule), e.g. $\mathcal{LL}(\mathrm{EOC}_{Bonf})$.

INSEAD

Comparing Procedures

- Theoretical evaluation:
  - Hard. Different objectives. Each makes approximations.
  - Can link large-sample EVI $\mathcal{LL}$ with small-sample $\mathcal{OCBA}_{LL}$
- Empirical measures of effectiveness:
  - Parameters of procedures implicitly define *efficiency curves*,

    $$(E[N], \log \mathrm{PICS}_{\mathrm{iz}}) \text{ or } (E[N], \log \mathrm{EOC}_{\mathrm{iz}})$$

    "More efficient" procedures have lower efficiency curves.
  - Efficiency ignores how to set parameter to achieve desired target $\mathrm{PICS}_{\mathrm{iz}}$ or $\mathrm{EOC}_{\mathrm{iz}}$
  - *Target curves* relate procedures parameter with desired target,

    $$(\log \alpha^*, \log \mathrm{PICS}_{\mathrm{iz}}) \text{ or } (\log \beta^*, \log \mathrm{EOC}_{\mathrm{iz}})$$

    "Conservative" procedures are below diagonal
    "Controllable": Can pick parameters to get desired target
  - Robust: Efficient and controllable over range of configs.

Configurations: Stylized

- Slippage configuration (SC): All worst systems tied for second.

$$\begin{aligned} X_{1j} &\sim \texttt{Normal}\left(0, 2\rho/(1+\rho)\right) \\ X_{ij} &\sim \texttt{Normal}\left(-\delta, 2/(1+\rho)\right) \text{ for } i = 2, \ldots, k \\ \delta^* &= \gamma\delta. \end{aligned}$$

  Best has largest variance if $\rho > 1$. $\mathrm{Var}[X_{1j} - X_{ij}]$ constant for all ρ. γ allows δ^* to differ from difference in means.

- Monotone decreasing means (MDM): Equally spaced means.

$$\begin{aligned} X_{ij} &\sim \texttt{Normal}\left(-(i-1)\delta, 2\rho^{2-i}/(1+\rho)\right) \\ \delta^* &= \gamma\delta. \end{aligned}$$

- Tested hundreds of combinations of $k \in \{2, 5, 10, 20, 50\}$; $\rho \in \{0.125, 0.177, 0.25, 0.354, 0.5, 0.707, 1, 1.414, 2, 2.828, 4\}$; $n_0 \in \{4, 6, 10\}$; $\delta \in \{0.25, 0.354, 0.5, 0.707, 1\}$; $\delta^* \in \{0.05, 0.1, \ldots, 0.6\}$.

INSEAD

Configurations: Randomized

- SC and MDM are unlikely to be found in practice
- Randomized problem may be more representative

- Randomized problem instances (RPI1):
  - Sample χ randomly (conjugate prior)

$$\begin{aligned} p(\sigma_i^2) &\sim \texttt{InvGamma}\,(\alpha, \beta) \\ p(W_i \,|\, \sigma_i^2) &\sim \texttt{Normal}\,(\mu_0, \sigma_i^2/\eta). \end{aligned}$$

  - We set $\beta = \alpha - 1 > 0$: standardize mean of variances to be 1. Increase η: means more similar (OCBA, VIP and $\eta \to 0$); Increase α: reduce variability in the variances.
  - Tested all combinations of $k \in \{2, 5, 10\}$; $\eta \in \{.707, 1, 1.414, 2\}$; $\alpha \in \{2.5, 100\}$.
- Also tested other RPI experiments

INSEAD

Summary: Numerics

- 20,000 combinations of allocation-stopping rule-configuration. Each generates an efficiency and target curve
- Each curve estimated with at least 100,000 macro-replications of each allocation/stopping rule combination

- CRN across configurations
- C++, Gnu Scientific Libary for cdfs and Mersenne twister RNG (Matsumoto and Nishimura 1998, 2002 revised seeding)
- FILIB++ (Lerch et al. 2001) for interval arithmetic (stability for $\mathcal{LL}_1$, 0-1_1, and sometimes $\mathcal{OCBA}$)
- Mixed cluster of up to 120 nodes: Linux 2.4 and Windows XP; Intel P4 and AMD Athlon; 2 to 3 GHz.
- Distributed via JOSCHKA-System (Bonn et al. 2005).

Flexible Stopping Rules Help

General observations for efficiency

- Flexible stopping beats $\mathcal{S}$ for VIP, OCBA, and Equal; all configs; PICS_{iz} and EOC_{iz}.
- For SC, MDM: EOC_{Bonf} beats PCS_{Slep} beats $\mathcal{S}$

Example in Figure

- Equal allocation, $\mathcal{KN}$++
- SC; $k = 2$; $\delta^* = 0.5$; $\rho = 1$
- NB: Equal and $\mathcal{KN}$++ are optimal if $k = 2$, $\rho = 1$, difference is stopping rule.

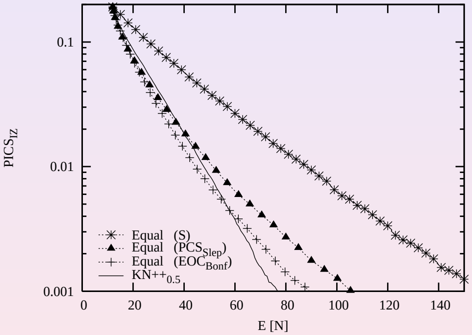

Efficiency of Allocations for SC, MDM

Observations for SC and MDM

- Equal performs poorly if $k \neq 2$, or unequal variances.
- *NO* procedure is controllable (robustly).
- $\mathcal{OCBA}$, $\mathcal{OCBA}_{LL}$, $\mathcal{LL}$ with EOC_{Bonf} typically most efficient.
- Often, $\exists\delta^*$ so that $\mathcal{KN}{++}$ is most efficient, but $\mathcal{KN}{++}$ extremely conservative at that δ^*

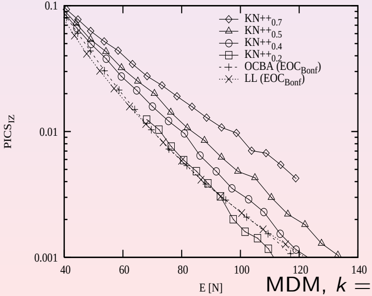

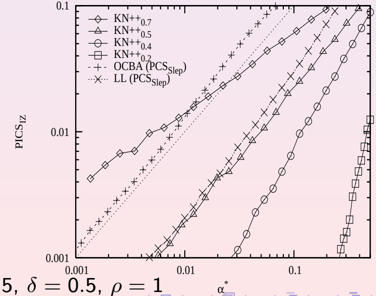

MDM, $k = 5$, $\delta = 0.5$, $\rho = 1$

INSEAD

Efficiency of Allocations for RPI1

RPI brings possibility of very close means

- Important to use $\text{PBS}_{\delta^*} = 1 - \text{PGS}_{\delta^*}$, not $\text{PICS} = 1 - \text{PCS}$.
- $\exists \delta^*$ such that $\text{PGS}_{Slep,\delta^*}$ more efficient than EOC_{Bonf}, *even* for EOC_{iz}, but only EOC_{Bonf} is controllable for EOC_{iz}
- Only $\text{PGS}_{Slep,\delta^*}$ is controllable for $\text{PGS}_{\text{iz},\delta^*}$

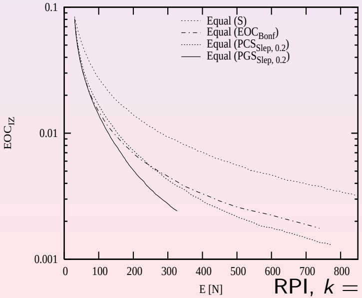

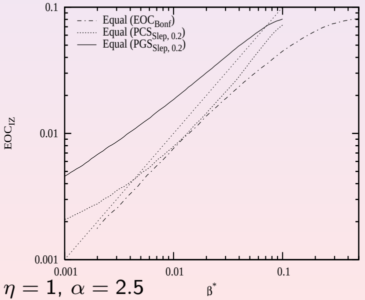

RPI, $k = 5$, $\eta = 1$, $\alpha = 2.5$

INSEAD

General Comments

Typically . . .

- $\mathcal{KN}$++ more efficient than original $\mathcal{OCBA}(\mathcal{S})$ and $\mathcal{LL}(\mathcal{S})$
- $\mathcal{LL}$, $\mathcal{OCBA}$, $\mathcal{OCBA}_{LL}$ with PGS_{Slep} or EOC_{Bonf} more efficient than $\mathcal{KN}$++
- $\mathcal{LL}$ beats 0-1 (even for PICS_{iz})
- OCBA and $\mathcal{LL}$ are greedy, but don't have that problem

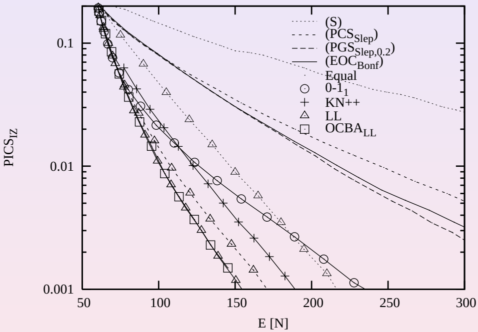

MDM, $k = 10$, $\delta = 0.5$, $\rho = 1$

INSEAD

Variations on the theme: Typically . . .

- OCBA: t (unknown σ^2) vs. normal ($\hat{\sigma}^2$) distribution approx.
  - same efficiency in allocation; but t better in stopping rule
- Student d.o.f. approximation for OCBA and VIP
  - Welch *slightly* beats Wilson and Pritsker (1984)
- Can use '+' or 'max' to include δ^* in allocation or stop rule (+ matches OCBA, 'max' like ETSS of Chen and Kelton).
  - '+' is more efficient than 'max'
  - Efficiency loss greater in stopping rule than in allocation.

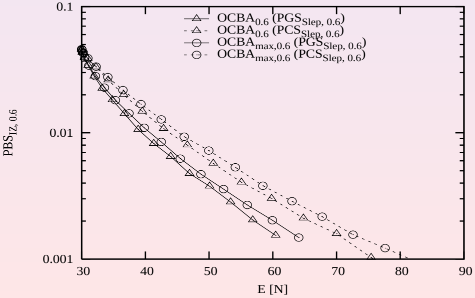

RPI, $k = 5$, $\eta = 1$, $\alpha = 100$

INSEAD

Which procedure to use (1)

- If budget constraint, use $\mathcal{OCBA}(\mathcal{S})$, $\mathcal{OCBA}_{LL}(\mathcal{S})$ or $\mathcal{LL}(\mathcal{S})$.

- No procedure controllable for SC and MDM.
- Only $\text{PGS}_{Slep,\delta^*}$ controllable for $\text{PGS}_{\text{iz},\delta^*}$; only EOC_{Bonf} controllable for EOC_{iz} in RPI.

- Some advantages and disadvantages of $\mathcal{KN}{++}$
  - Plus: Beats old $\mathcal{LL}(\mathcal{S})$, $\mathcal{OCBA}(\mathcal{S})$; robust to n_0; $1-\alpha^* \leq \text{PCS}_{\text{iz}}$; CRN
  - Minus: Not controllable; conservative (if want $1-\alpha^* = \text{PICS}_{\text{iz}}$ rather than $1-\alpha^* \leq \text{PCS}_{\text{iz}}$), e.g. large k, heterogeneous σ_i^2, δ^* too small.

INSEAD

Which procedure to use (2)

- We recommend $\mathcal{LL}$, $\mathcal{OCBA}_{LL}$ or $\mathcal{OCBA}$ allocation with $\text{PGS}_{Slep,\delta^*}$ or EOC_{Bonf} stopping rule (depending on goal)
  - Plus: Most efficient; controllable for RPI; robust; ability to incorporate sampling costs; how about PCS_{Bayes} and EOC_{Bayes} guarantees; prior information ok; ...
  - Minus: Sensitive to n_0 for extreme levels of evidence; slight degredation if many good systems; independence (although two-stage for VIP; recent work for OCBA).

- Do not use: 0-1; 'max' instead of '$+$' to bring in δ^* into allocation; normal distribution in stopping rule if variance unknown; small n_0 if extreme evidence levels desired; new 'small sample' EVI allocations.
- Caveats: Empirical observations limited to our testbed; assumed normality; no autocorrelation; no CRN; did not examine combinatorially large k

INSEAD

Discussion

- Link top procedures in large search spaces, assess with companion tools (DOvS; evolutionary algorithms; etc.)
- $\mathcal{KN}$++-like procedure with different number of reps/system.

- Standardized testbed. Performance evaluation criteria.
  - Within class: strengths and weaknesses
  - Across classes: broader testbed
- Economic basis for simulation projects. Why stop simulating? Statistical versus economic significance? e.g. mean # reps. versus simulation project costs and net revenues accrued from decision. (Chick and Gans 2005 suggest DP/bandit/real options approach.)

Bonn, M., F. Toussaint, and H. Schmeck. 2005.
JOSCHKA: Job-Scheduling in heterogenen Systemen mit Hilfe von Webservices.
In *PARS Workshop Lübeck*, ed. E. Maehle, in press.
Gesellschaft für Informatik.

Branke, J., S. E. Chick, and C. Schmidt. 2005.
Selecting a selection procedure.
working paper.

Chen, C.-H. 1996.
A lower bound for the correct subset-selection probability and its application to discrete event simulations.
IEEE Transactions on Automatic Control 41 (8): 1227–1231.

Chen, E. J., and W. D. Kelton. 2005.
Sequential selection procedures: Using sample means to improve efficiency.
European Journal of Operational Research 166:133–153.

Chick, S. E., and K. Inoue. 2001.
New two-stage and sequential procedures for selecting the best simulated system.
Operations Research 49 (5): 732–743.

He, D., S. E. Chick, and C.-H. Chen. 2005.
The opportunity cost and OCBA selection procedures in ordinal optimization.
submitted.

Kim, S.-H., and B. L. Nelson. 2001.

A fully sequential procedure for indifference-zone selection in simulation.
ACM TOMACS 11:251–273.

Lerch, M., G. Tischler, J. W. von Gudenberg, W. Hofschuster, and W. Kraemer. 2001.
The interval library filib++ 2.0 - design, features and sample programs.
Preprint 2001/4, University of Wuppertal.

Matsumoto, M., and T. Nishimura. 1998.
Mersenne twister: A 623-dimensionally equidistributed uniform pseudorandom number generator.
ACM TOMACS 8 (1): 3–30.

Wilson, J. R., and A. A. B. Pritsker. 1984.
Experimental evaluation of variance reduction techniques for queueing simulation using generalized concomitant variables.
Management Science 30:1459–1472.