

A NEW FRAMEWORK OF DESIGNING SEQUENTIAL RANKING-AND-SELECTION PROCEDURES

Ying Zhong

Department of Management Sciences
City University of Hong Kong, Kowloon, HONG KONG

L. Jeff Hong

Department of Economics and Finance and Department of Management Sciences
City University of Hong Kong, Kowloon, HONG KONG

ABSTRACT

Many classical sequential procedures model the partial sum difference process between two competing alternatives as a Brownian motion process. In this paper, the marginal probability of eliminating the best alternative is considered while modeling the partial sum difference process. We adaptively allocate the total amount of the probability of incorrect selection α to every time point where the comparisons between alternatives are conducted and set the continuation regions to ensure the marginal probability of eliminating the best alternative does not exceed the probability assigned to each time point t . We show by examples that under our framework, the procedure can be easily developed for both indifference-zone (IZ) and IZ free formulations.

1 INTRODUCTION

In the simulation field, there is a large body of literature on the ranking-and-selection (R&S) problems which focus primarily on selecting the best alternative with the highest mean performance among a finite set of alternatives, see Bechhofer et al. (1995), Kim and Nelson (2006) and Chick (2006) for an overview. Based on the statistical inference used, the procedures currently available can be classified in two approaches: the Bayesian approach and the frequentist approach.

In the early development of frequentist procedures, procedures are designed under the situation where the observations come from physical experiments, e.g., agricultural experiments or clinical trials. These experiments generally require a quite long time period to get the sample results and observations collected are often in batches. Procedures being able to initially identify the total sample size and select the best at the end of the collection, i.e., stage-wise selection, are preferred. One major difficulty encountered, as to the procedure design, is that, for the case where the second best alternative is arbitrarily close to the best one, the sample size needed to select the best can be arbitrarily large which is hard to guarantee in practice. To address the issue, Bechhofer (1954) makes the first attempt by introducing an indifference-zone (IZ) formulation. In that work, the difference between the mean performance of the best alternative and that of the second best is assumed to be at least $\delta > 0$ which is called the IZ parameter. In the presence of equal and known variances for all the alternatives, the author shows by an example that, given a probability of correct selection (PCS) $1 - \alpha$, even a simple procedure which picks the alternative with the highest sample mean after sampling each alternative in a pre-determined number of times works well. Many procedures thereafter in the field adopt such formulation. Subset-selection(SS) formulation proposed in Gupta (1956) and Gupta (1965) offers another way to solve the problem. Instead of choosing the unique best alternative by introducing an IZ parameter, the procedure developed under this formulation returns a subset of the

original alternative set at the end of the experiment and can guarantee that the best alternative is in the subset with PCS $1 - \alpha$. Both ways of formulations make the original goal of the R&S problem practically achievable. In this paper, we mainly focus on the IZ formulation.

Associating with the rapid growth of computer technology, under the IZ formulation, literature in this realm starts to look into the problems where observations are generated by computer based simulations. As opposed to the traditional physical experiments, the time it takes to generate each observation reduces dramatically and observations arrive no longer in batches but one at a time. These new features naturally drive researchers to consider procedures which can make comparisons between current existing alternatives upon new arrivals and eliminate inferior ones if enough sample evidence is gathered. The earliest idea of sequential elimination can be traced back to Paulson (1964). Probably the most influential and frequently cited work in this line of research is that by Kim and Nelson (2001). The fully sequential procedure KN proposed in the paper perfectly fits the nature of the computer based simulations in a way that, as long as each current existing alternative collects one more observation, a new round of comparisons can be done. Since then, many fully sequential procedures have been developed (see, for instance, Jeff Hong (2006), Hong and Nelson (2007)). These fully sequential procedures show great advantages of reducing the total sample size and high efficiencies while dealing this type of problems. To implement such procedures, one common way is to model the partial sum difference process between two competing alternatives as a Brownian motion process and the process stops as the partial sum difference goes beyond the pre-defined continuation region at which we can make a decision.

However, there is a drawback to the procedure under the IZ formulation. As the IZ parameter δ is set by the experimenter's prior belief on the problem, the procedure's performance depends heavily on this parameter δ . On the one hand, if the real difference between the best and that of the second best is much larger than δ , the procedure tends to be very conservative. On the other hand, if the real difference is smaller than δ , the statistical validity of the procedure may no longer be guaranteed. To break down the IZ formulation, in Fan, Hong, and Nelson (2016)'s work, the authors propose a class of fully sequential procedures whose continuation regions are determined based on the Law of the Iterated Logarithm. Built under the IZ free formulation, these procedures are able to select the best alternative with a pre-specified PCS as long as it is unique. The numerical results show that, their procedures perform much better than KN procedure when the IZ parameter δ is set deviating from the true difference and the means of the alternatives are spread out over a wide range in a large-scale R&S problem. While determining the continuation regions, very complicated theories are involved in their procedures. In this paper, we propose a totally different view on how to construct the continuation regions and show that the procedure satisfying either IZ free or IZ formulation can be easily developed without employing many complicated theories.

The rest of this paper is organized as follows. In section 2, motivated by the literature on multi-armed bandit (MAB) problems, we develop a new framework to construct the continuation regions and one procedure developed under this framework is presented in the later part of the section. Section 3 shows some preliminary numerical results followed by conclusion and future possible research direction in Section 4.

2 A MARGINAL PROBABILITY AGGREGATION FRAMEWORK

2.1 Main Idea

To facilitate the illustration, we first introduce some standard notations in R&S field. Let $I = \{1, 2, \dots, k\}$ be the set of alternatives and X_{ir} denotes the r th output from alternative $i \in I$ which follows a normal distribution with unknown mean μ_i and variance $\sigma_i^2 < \infty$. All the outputs X_{ir} for $r = 1, 2, \dots$ are independent and identical distributed (i.i.d.) and X_{ir} is independent of X_{jr} for $i \neq j$. Without loss of generality, we assume $\mu_1 \leq \mu_2, \dots, \leq \mu_{k-1} < \mu_k$, i.e., k th alternative is the unique best one.

Our work is closely related to the literature on MAB, see Lai (2001) for a comprehensive review. Being different from R&S, MAB aims to maximize the total expected rewards as samples are sequentially drawn

from statistical populations (arms) with unknown and different mean rewards. To measure the performance of a procedure, the term *regret* is defined as the difference between the total expected rewards if the procedure is applied and that of an oracle, that repeatedly samples the best arm. Thus, the smaller the *regret* the better the procedure. Since the exact value of *regret* is generally hard to get, the smallest possible upper bound for *regret* is often pursued. One technique frequently used is to consider the upper bound of the marginal probability of making a mistake, i.e., fail to sample the arm with the highest mean reward, at every time point t . Aggregating all the upper bounds over the entire time horizon provides an upper bound to the total expected number of times the procedure fails to choose the optimal arm. Therefore, an upper bound for the *regret* can be gotten.

While designing our fully sequential procedure, we reverse the aggregation process. We adaptively allocate the total amount of incorrect selection probability α to every time point where the comparisons between alternatives are conducted. Since we only care about whether the best alternative is eliminated or not, the assigned probability, namely α_t , is served as the upper bound of the marginal probability of falsely eliminating alternative k at time point t . We then equally assign $\alpha_t/(k-1)$ to each pairwise comparison between the best and the others. If the decision boundary for each partial sum difference is chosen properly based on $\alpha_t/(k-1)$, the procedure's statistical validity can be provided. It worthwhile to note that, in MAB, because the observations are assumed to be sampled from distributions with bounded supports, to find the upper bound of the marginal probability of making a mistake, the *Hoeffding's Inequality* (Hoeffding 1963) is often used. For the case of R&S with unknown variances, since the observations are generated from normal distributions, we can actually transform the partial sum difference between two competing alternatives into a random variable which follows Student's t -distribution at time t . Consequently, finding the boundary of the partial sum difference can refer to that random variable.

Comparing with the traditional fully sequential procedures which employ the Brownian process, our procedure is more conservative, in the sense that at every time point t , we do not separately consider the case where some alternatives may be eliminated prior to t but always assume that all the alternatives are in contention. Equally assigning $\alpha_t/(k-1)$ to each pairwise comparison results a loose boundary for the partial sum difference. However, due to the simplicity of constructing the continuation region, under our framework, the procedure for either IZ or IZ free formulation can be easily designed.

2.2 Procedure

In this part, we give a detail description of the procedure designed based on the main idea listed above. An IZ parameter $\delta \geq 0$, such that $\mu_1 \leq \mu_2, \dots, \mu_{k-1} < \mu_k - \delta$, is assumed. Notice that we allow $\delta = 0$, which implies an IZ free formulation.

Procedure 1.

1. **Setup:** Select PCS $1 - \alpha$ ($0 < \alpha \leq 1 - 1/k$). Set IZ parameter δ
2. **Initialization:** Let $I = \{1, 2, \dots, k\}$ be the set of alternatives in contention. For each $i \in I$, simulate one observations X_{i1} from alternative i . Set $t = 1$.
3. **Update:** Take one additional observation $X_{i,t+1}$ from each alternative $i \in I$ and calculate

$$\begin{aligned} \bar{X}_i(t+1) &= \frac{1}{t+1} \sum_{s=1}^{t+1} X_{is}, \forall i \in I \\ S_{ij}^2(t+1) &= \frac{1}{t} \sum_{s=1}^{t+1} [X_{is} - X_{js} - (\bar{X}_i(t+1) - \bar{X}_j(t+1))]^2, \forall j \in I \setminus \{i\} \end{aligned}$$

4. Screening: Let $I^{old} = I$, $t = t + 1$ and

$$W_{ij}(\alpha, k, t) = \sqrt{t} S_{ij}(t) \times \text{T.INV} \left(1 - \frac{\alpha_t}{(k-1)}, t-1 \right)$$

where $\alpha_t = \alpha/t^2$ and T.INV(p, df) returns the inverse of Student's t cumulative distribution function using the degrees of freedom (df) for the corresponding probability (p). Define

$$I = I^{old} \setminus \left\{ i \in I^{old} : t(\bar{X}_j(t) - \bar{X}_i(t)) \geq W_{ij}(\alpha, k, t) - \delta \times t \text{ for some other } j \in I^{old} \right\}$$

5. Stopping Rule: If $|I| = 1$, stop and select the alternative whose index is in I as the best. Otherwise, go to **Update**.

Remark 1. Notice that $W_{ij}(\alpha, k, t) - \delta \times t$ is the decision boundary we designed for the partial sum difference process between alternatives i and j . The larger δ we choose, the tighter decision boundary we have.

Remark 2. By setting the IZ parameter $\delta = 0$ or $\delta > 0$, Procedure 1 is able to switch between the IZ free and IZ formulations respectively.

Remark 3. Limited by Stein (1945)'s result, the first stage sampling in many traditional fully sequential procedures is served as the variance estimation of the alternatives and the estimated variances cannot be updated in the future. For these procedures, there exists a dilemma of choosing the first stage sample size. Since Procedure 1 does not involve Stein (1945)'s result, we allow the sample variance of each alternative to be updated every time we collect new observations and eliminations could start at $t = 2$.

Remark 4. The statistical validity of Procedure 1 can be shown by simply summing up the marginal probabilities of falsely eliminating alternative k from $t = 2, 3, \dots$

$$\begin{aligned} & \Pr(\text{alternative } k \text{ is eliminated at } t | \mu_k - \mu_{k-1} > \delta) \\ \leq & \sum_{i=1}^{k-1} \Pr(\text{alternative } i \text{ eliminates alternative } k \text{ at } t | \mu_k - \mu_i > \delta) \\ = & \sum_{i=1}^{k-1} \Pr(t(\bar{X}_i(t) - \bar{X}_k(t)) \geq W_{ik}(\alpha, k, t) - \delta \times t | \mu_k - \mu_i > \delta) \\ \leq & \sum_{i=1}^{k-1} \Pr(t(\bar{X}_i(t) - \bar{X}_k(t) - (\mu_i - \mu_k)) \geq W_{ik}(\alpha, k, t) | \mu_k - \mu_i > \delta) \\ = & \sum_{i=1}^{k-1} \Pr\left(\frac{\sqrt{t}(\bar{X}_i(t) - \bar{X}_k(t) - (\mu_i - \mu_k))}{S_{ik}(t)} \geq \text{T.INV}\left(1 - \frac{\alpha_t}{(k-1)}, t-1\right) \mid \mu_k - \mu_i > \delta\right) \\ = & \sum_{i=1}^{k-1} \frac{\alpha_t}{(k-1)} = \frac{\alpha}{t^2} \end{aligned}$$

Since

$$\frac{\sqrt{t}(\bar{X}_i(t) - \bar{X}_k(t) - (\mu_i - \mu_k))}{S_{ik}(t)}$$

has a Student's t -distribution with $t - 1$ degrees of freedom, the first equality in the last line holds. The last inequality holds due to the assumption that $\mu_i < \mu_k - \delta \forall i \neq k$. Therefore, the PCS can be evaluated as,

$$\begin{aligned}
 PCS &= 1 - \Pr(\text{alternative } k \text{ is eliminated} | \mu_k - \mu_{k-1} > \delta) \\
 &\geq 1 - \sum_{t=2}^{\infty} \Pr(\text{alternative } k \text{ is eliminated at } t | \mu_k - \mu_{k-1} > \delta) \\
 &\geq 1 - \alpha \sum_{t=2}^{\infty} \frac{1}{t^2} \\
 &\geq 1 - \alpha \int_1^{\infty} \frac{1}{t^2} dt \geq 1 - \alpha
 \end{aligned}$$

So the Procedure 1 is statistically valid.

3 NUMERICAL RESULTS

For all the numerical experiments listed below, by default, the PCS $1 - \alpha$ is set to be 0.95. In order to take a glance over the growth rate of our decision boundary $W_{ij}(\alpha, k, t) - \delta \times t$ with respect to time t , we first consider a R&S problem with $k = 2$. We assume that observations generated from these two alternatives are normally distributed with mean $\mu_1 < \mu_2$ and variance $\sigma_1^2 = \sigma_2^2 = 1$. We simulate each alternative 2000 times. With these simulated observations, we accordingly calculate two decision boundaries for the partial sum difference between these two alternatives by letting $\delta = 0$ and $\delta = 0.2$ respectively. We plot the averaged decision boundaries for these two cases over 1,000 macro-replications in Figure 1.




Figure 1: The decision boundaries for $\delta = 0$ and $\delta = 0.2$.

In Figure 1, the dash line represents the averaged decision boundary for $\delta = 0$ and the other one is for the case $\delta = 0.2$. By observing the figure, we obtain several insights. First, the solid curve is dominated by the dash one. This is consistent with our analysis in Remark 1 as the decision boundary for $\delta > 0$ is always tighter than the one for $\delta = 0$. Second, while letting $\delta = 0$, the decision boundary has a sub-linear growth rate with respect to time t . This implies that the partial sum difference (grows linearly with t) between these two alternatives will go beyond the decision boundary and our procedure will stop in finite time. Third, the solid line indicates that, as δ is strictly greater than 0, the decision boundary would

intersect with the x-axis. This is quite similar to the ones in KN procedure except that, in our procedure, these curves go up first and then down.

We then exam the performance of our procedure for the same problem as the previous example. We test different settings of δ where $\delta = 0$ and $(\mu_2 - \mu_1)/\delta = 1/4, 1/2, 1, 2, 4$. We report the estimated PCS and the average total sample size with 95% confidence interval over 2,000 macro-replications in Table 1.

Table 1: Performance test for Procedure 1 with PCS=0.95

$\mu_2 - \mu_1$	Procedure 1: $\delta = 0$		Procedure 1: $(\mu_2 - \mu_1)/\delta$				
			1/4	1/2	1	2	4
0.2	0.98		0.80	0.95	0.98	0.99	0.99
	2.74×10^3	0.60×10^2	2.06×10^2	5.46×10^2	1.08×10^3	1.60×10^3	
	$\pm 0.06 \times 10^3$	$\pm 0.02 \times 10^2$	$\pm 0.06 \times 10^2$	$\pm 0.14 \times 10^2$	$\pm 0.02 \times 10^3$	$\pm 0.04 \times 10^3$	
0.15	0.99		0.83	0.96	0.98	0.98	0.98
	5.23×10^3	1.07×10^2	4.02×10^2	1.06×10^3	2.01×10^3	3.17×10^3	
	$\pm 0.12 \times 10^3$	$\pm 0.03 \times 10^2$	$\pm 0.11 \times 10^2$	0.02×10^3	$\pm 0.04 \times 10^3$	$\pm 0.08 \times 10^3$	
0.1	0.98		0.84	0.97	0.98	0.98	0.98
	1.32×10^4	2.83×10^2	1.03×10^3	2.74×10^3	5.26×10^3	7.99×10^3	
	$\pm 0.02 \times 10^4$	$\pm 0.07 \times 10^2$	$\pm 0.03 \times 10^3$	$\pm 0.04 \times 10^3$	$\pm 0.12 \times 10^3$	$\pm 0.16 \times 10^3$	
0.05	0.98		0.87	0.98	0.98	0.98	0.98
	6.20×10^4	1.46×10^3	5.21×10^3	1.32×10^4	2.50×10^4	3.81×10^4	
	$\pm 0.11 \times 10^4$	$\pm 0.03 \times 10^3$	$\pm 0.12 \times 10^3$	$\pm 0.02 \times 10^4$	$\pm 0.04 \times 10^4$	$\pm 0.07 \times 10^4$	

From Table 1, it can be easily observed that, whatever value δ takes, the smaller the true difference $\mu_2 - \mu_1$, the larger total sample size Procedure 1 needs to tell the difference between these two alternatives. Notably, no matter how small the true difference $\mu_2 - \mu_1$ is, the procedure can deliver a PCS greater than the one desired by letting $\delta = 0$. This is the most conservative way to choose the best alternative as it always requires a larger total sample size than those with $\delta > 0$ (see each row). But the difference becomes smaller as $(\mu_2 - \mu_1)/\delta$ increases. Like many other traditional sequential procedures developed under IZ formulation, for the case $\delta > 0$, our procedure fails to correctly choose the best alternative with a PCS higher than the desired one as the true difference $\mu_2 - \mu_1$ is much smaller than δ (see the third column). Table 1 suggests that, for a particular problem, if the lower bound for the difference between the best alternative and the second best alternative is unknown, $\delta = 0$ is preferred over $\delta > 0$. Otherwise, δ should be set equal to the lower bound.

As illustrated in Fan, Hong, and Nelson (2016)'s work, the procedures developed under IZ free formulation show the advantage of dealing the large-scale R&S problem over the ones developed under IZ formulation. We also conduct a comparison between Procedure 1 with IZ parameter $\delta_1 = 0$ and the classical KN procedure in a similar setting. As in practice, the alternative means in the large-scale R&S problem usually spread out over a wide range, we consider a monotone increasing configuration of means in which $\mu_i = -5.5 + 0.5i$. An equal configuration of variances where $\sigma_i^2 = 5$ for all alternatives $i = 1, 2, \dots, k$ is assumed as well. We test the performance of these two procedures with the number of alternatives varying over $k = 20, 50, 100, 200$. For each k , different settings of IZ parameter δ_{KN} where $(\mu_k - \mu_{k-1})/\delta_{KN} = 1/4, 1/2, 1, 2, 4$, are applied to KN procedure. Since KN procedure requires a first stage sampling, in this experiment, we set the first stage sample size n_0 to be 10. The estimated PCS and average total sample size with 95% confidence interval based on 2,000 macro-replications are reported in Table 2.

From Table 2, we have several findings. First, from the fourth and fifth columns, we can conclude that if δ_{KN} is set near $\mu_k - \mu_{k-1}$, KN procedure tends to perform better than Procedure 1, i.e., KN procedure requires a smaller total sample size to deliver a PCS meeting the desire. However, if δ_{KN} is set much smaller than the true difference $\mu_k - \mu_{k-1}$ (less than half in this case), the KN procedure becomes quite conservative and demands a larger total sample size than Procedure 1. Second, the failure of the statistical

Table 2: Comparisons between Procedure 1 and KN with PCS=0.95

k	Procedure 1: $\delta_1 = 0$		KN: $(\mu_k - \mu_{k-1})/\delta_{KN}$				
			1/4	1/2	1	2	4
20	1.00		0.82	0.96	1.00	1.00	1.00
	3.46×10^3	2.95×10^2	5.70×10^2	1.27×10^3	2.80×10^3	5.99×10^3	
	$\pm 0.05 \times 10^3$	$\pm 0.02 \times 10^2$	$\pm 0.06 \times 10^2$	$\pm 0.02 \times 10^3$	$\pm 0.03 \times 10^3$	$\pm 0.08 \times 10^3$	
50	1.00		0.87	0.98	1.00	1.00	1.00
	3.80×10^3	6.72×10^2	1.06×10^3	2.04×10^3	4.35×10^3	9.04×10^3	
	0.05×10^3	$\pm 0.03 \times 10^2$	$\pm 0.01 \times 10^3$	$\pm 0.03 \times 10^3$	$\pm 0.05 \times 10^3$	$\pm 0.10 \times 10^3$	
100	1.00		0.88	0.99	1.00	1.00	1.00
	4.16×10^3	12.72×10^2	1.76×10^3	3.01×10^3	5.81×10^3	1.12×10^4	
	$\pm 0.05 \times 10^3$	$\pm 0.03 \times 10^2$	$\pm 0.01 \times 10^3$	$\pm 0.02 \times 10^3$	$\pm 0.06 \times 10^3$	$\pm 0.01 \times 10^4$	
200	1.00		0.92	0.99	1.00	1.00	1.00
	4.71×10^3	24.31×10^2	3.04×10^3	4.59×10^3	8.08×10^3	1.55×10^4	
	0.05×10^3	$\pm 0.04 \times 10^2$	$\pm 0.01 \times 10^3$	$\pm 0.03 \times 10^3$	$\pm 0.07 \times 10^3$	$\pm 0.01 \times 10^4$	

validity still exists for KN procedure when δ_{KN} is four times larger than the true difference in this case. Third, observing the average total sample sizes in each column, we find that an increasing k has more adverse impacts on KN procedure than on Procedure 1. As the number of alternatives increases, most of the newly added alternatives are obvious inferior candidates. Because the decision boundaries in Procedure 1 are tighter than those in KN procedure while t is small, these obvious inferior candidates can be quickly eliminated by Procedure 1 without incurring much sampling effort. Thus, Procedure 1 suffers less than KN procedure from an increasing k .

4 CONCLUSION

In this paper, we provided a new framework to construct the continuation region for the partial sum difference between two competing alternatives without employing Brownian motion process. We adaptively allocate the tolerable probability of incorrect selection α to every time point where the comparisons between alternatives a conducted. The decision boundaries are designed to ensure that at every time point t , the probability of eliminating the best alternative is less than the probability assigned to that time. Under this framework, procedures can be easily designed and keep the advantage of solving the large-scale R&S problem over the traditional fully sequential procedures developed under the IZ formulation.

We are working on incorporating the first stage sampling into our procedure. Being different from the traditional fully sequential procedure, we would look at the first stage sampling from a new angle. As our procedures are likely to make mistakes while t is small, restricting procedures from making eliminations in a larger first stage sampling would allow us to assign more amount of α_t to every time point thereafter and lead to a tighter bound. A proper defined first stage sample size may significantly boost our procedures' efficiencies.

ACKNOWLEDGMENTS

We gratefully acknowledge the support from the Hong Kong Research Grants Council under grants GRF 16203214 and GRF 11270116.

REFERENCES

Bechhofer, R. E. 1954. "A single-sample multiple decision procedure for ranking means of normal populations with known variances". *The Annals of Mathematical Statistics*:16–39.

- Bechhofer, R. E., D. M. Goldsman, and T. J. Santner. 1995. *Design and analysis of experiments for statistical selection, screening, and multiple comparisons*. Wiley.
- Chick, S. E. 2006. "Subjective Probability and Bayesian Methodology". In *Elsevier Handbooks in Operations Research and Management Science: Simulation*, edited by S. G. Henderson and B. L. Nelson, 225–257. Elsevier.
- Fan, W., L. J. Hong, and B. L. Nelson. 2016. "Indifference-Zone-Free Selection of the Best". *Operations Research* 64 (6): 1499–1514.
- Gupta, S. S. 1956. *On a decision rule for a problem in ranking means*. Ph. D. thesis, University of North Carolina at Chapel Hill.
- Gupta, S. S. 1965. "On some multiple decision (selection and ranking) rules". *Technometrics* 7 (2): 225–245.
- Hoeffding, W. 1963. "Probability inequalities for sums of bounded random variables". *Journal of the American statistical association* 58 (301): 13–30.
- Hong, L. J., and B. L. Nelson. 2007. "Selecting the best system when systems are revealed sequentially". *Iie Transactions* 39 (7): 723–734.
- Jeff Hong, L. 2006. "Fully sequential indifference-zone selection procedures with variance-dependent sampling". *Naval Research Logistics (NRL)* 53 (5): 464–476.
- Kim, S.-H., and B. L. Nelson. 2001. "A fully sequential procedure for indifference-zone selection in simulation". *ACM Transactions on Modeling and Computer Simulation (TOMACS)* 11 (3): 251–273.
- Kim, S.-H., and B. L. Nelson. 2006. "Selecting the best system". *Handbooks in operations research and management science* 13:501–534.
- Lai, T. L. 2001. *Sequential analysis*. Wiley Online Library.
- Paulson, E. 1964. "A sequential procedure for selecting the population with the largest mean from k normal populations". *The Annals of Mathematical Statistics*:174–180.
- Stein, C. 1945. "A two-sample test for a linear hypothesis whose power is independent of the variance". *The Annals of Mathematical Statistics* 16 (3): 243–258.

AUTHOR BIOGRAPHIES

YING ZHONG is a Ph.D. candidate with Hong Kong Ph.D. Fellowship in Department of Management Science in the College of Business at City University of Hong Kong. His research interests include ranking & selection, chance constrained selection of the best and multi-armed bandit problems with covariates. His email address is yzhong4-c@my.cityu.edu.hk.

L. JEFF HONG is a chair professor in the Department of Economics and Finance, and the Department of Management Sciences, at the City University of Hong Kong. His research interests include Monte Carlo methods, financial engineering, and stochastic optimization. He is currently an associate editor for *Operations Research*, *Naval Research Logistics* and *ACM Transactions on Modeling and Computer Simulation*. His email address is jeffhong@cityu.edu.hk.