Proceedings of the 2017 Winter Simulation Conference
W. K. V. Chan, A. D'Ambrogio, G. Zacharewicz, N. Mustafee, G. Wainer, and E. Page, eds.

RATE-OPTIMALITY OF THE COMPLETE EXPECTED IMPROVEMENT CRITERION

Ye Chen

Department of Mathematics
University of Maryland
College Park, MD 20742, USA

Ilya O. Ryzhov

Robert H. Smith School of Business
University of Maryland
College Park, MD 20742, USA

ABSTRACT

Expected improvement (EI) is a leading algorithmic approach to simulation-based optimization. However, it was recently proved that, in the context of ranking and selection, some of the most well-known EI-type methods cause the probability of incorrect selection to converge at suboptimal rates. We investigate a more recent variant of EI (known as "complete EI") that was proposed by Salemi, Nelson, and Staum (2014), and summarize results showing that, with some minor modifications, complete EI can be made to achieve the optimal convergence rate in ranking and selection with independent Gaussian noise. This is the strongest theoretical guarantee available for any EI-type method.

1 INTRODUCTION

In simulation optimization (Fu 2015), the role of simulation is to provide information about a complex stochastic system, and the role of optimization is to use this information to improve decision-making. Often, simulations are expensive and time-consuming, which makes it necessary to pick and choose which information we want to acquire. For example, given a limited number of replications, we may wish to allocate most of them to a small number of "promising" decisions. This gives rise to several fundamental problems: we need to 1) characterize the impact of information on the quality of the ensuing decisions; 2) develop a metric for comparing different allocation strategies; 3) design optimal strategies (or at least good ones) with respect to this metric.

The ranking and selection (R&S) problem provides a stylized mathematical framework that allows us to precisely answer each of these questions. Numerous surveys and tutorials are available on simulation optimization in general and R&S in particular; some examples from previous WSC proceedings include Hong and Nelson (2009), Chau et al. (2014), and Jian and Henderson (2015). In R&S, there is a finite number of "alternatives," each of which has an unknown "value," and the optimization problem is to select the highest-valued alternative. A single simulation replication provides a noisy observation of the value of *one* alternative of our choosing. The relative simplicity of this problem allows for a much richer and more detailed theoretical characterization of optimal information collection than might be possible for more complex problems.

Consider a particularly simple version of R&S in which the simulation output is independent and normally distributed, and a sample of one alternative provides no information about any others. For this problem, Glynn and Juneja (2004) derived an explicit characterization of the *optimal* allocation strategy, expressed in terms of the proportions of the total budget assigned to each of the alternatives. If these proportions satisfy certain conditions, then the probability of incorrect selection (which occurs when a suboptimal alternative has a better sample mean than the optimal one) will converge to zero at an exponential rate, with the best possible exponent, as the total budget increases to infinity.

Although the optimal solution to this problem has been known for some time, it is difficult to implement in practice since the optimal proportions are themselves functions of the unknown values. For this reason, researchers have preferred to focus on heuristic methods that are easier to implement, such as optimal computing budget allocation (Chen and Lee 2010), Thompson sampling (Russo and Van Roy 2014), indifference-zone selection (Hong and Nelson 2005), or other approaches surveyed in Powell and Ryzhov (2012). In this paper, we focus specifically on the expected improvement (EI) class of methods (Jones, Schonlau, and Welch 1998, Chick, Branke, and Schmidt 2010, Powell and Ryzhov 2012). EI is a Bayesian approach to R&S (Chick 2006) that allocates simulations in a sequential manner: every new sample is used to update posterior distributions of the values, and the so-called "value of information" criterion is used to assign the next sample to the alternative that appears to have the most potential to improve the most recent estimate of the best value.

There is no single agreed-upon definition for the value of information. The literature has experimented with many variants, such as the classic EI criterion of Jones, Schonlau, and Welch (1998), the knowledge gradient criterion (Powell and Ryzhov 2012), the LL_1 criterion (Chick, Branke, and Schmidt 2010) and others. EI is a flexible concept that can be adapted to many problem classes in simulation optimization (Scott, Powell, and Simão 2010, Han, Ryzhov, and Defourny 2016) and often performs very well in practice (Branke, Chick, and Schmidt 2007). At the same time, Ryzhov (2016) showed that three of the most well-known variants of EI all produce suboptimal asymptotic allocations (i.e., none of them satisfies the optimality conditions of Glynn and Juneja 2004).

Recently, Salemi, Nelson, and Staum (2014) proposed a novel variant of EI, called "complete expected improvement" (CEI). Under the CEI criterion, the value of information for a seemingly-suboptimal alternative includes our uncertainty about the estimated best value, whereas classic EI uses only a point estimate of this quantity. Salemi, Nelson, and Staum (2014) derived the CEI criterion in the context of Gaussian Markov random fields, a powerful Bayesian learning model that allows a sample of one alternative to provide information about all others. This model is much more scalable than the simple version of R&S considered in the present paper, but it is also much more difficult to study (for example, the convergence rate analysis of Glynn and Juneja 2004 does not handle Gaussian Markov structure).

In this paper, we translate the CEI criterion to our simpler R&S model, which allows us to study its asymptotic behaviour theoretically. Our main result, proved rigorously in Chen and Ryzhov (2017) and summarized in this paper, is that CEI achieves the *optimal* convergence rate, with a minor modification to the method as originally presented in Salemi, Nelson, and Staum (2014). This is one of the strongest optimality results for any R&S heuristic to date, with the caveat that we are considering a relatively simple version of R&S. For instance, Russo (2017) develops a new class of heuristics for the same problem that can provably achieve rate-optimality, but this requires tuning of a parameter, whereas CEI is entirely tuning-free. Peng and Fu (2017) shows that the optimal allocation may also be recovered by reverse-engineering the computational form of EI, but one first requires an approximate solution to the optimality conditions, whereas CEI requires no additional computational effort over classic EI, and also has a more intuitive interpretation in our opinion.

The contribution of this result to the simulation literature is twofold. First, our work complements Salemi, Nelson, and Staum (2014) and provides new theoretical evidence in support of the CEI concept, even though our results are not directly applicable to the Gaussian Markov model considered in that paper. Second, our work strengthens the body of theoretical support for EI-type methods in general, helping to close the gap between the purely theoretical analysis of Glynn and Juneja (2004) and the more practical issues that drove the development of EI-type methods in the first place.

2 PRELIMINARIES

Let there be M alternatives, and let μ_x be the unknown value of alternative $x \in \{1,...,M\}$. The best alternative $x_* = \arg\max_x \mu_x$ is assumed to be unique. Let $\{x^n\}_{n=0}^{\infty}$ be a sequence of alternatives chosen for sampling. For every x^n, we observe $W_{x^n}^{n+1} \sim \mathcal{N}\left(\mu_{x^n}, \lambda_{x^n}^2\right)$, where the variances $\lambda_x^2 > 0$ are assumed to be known for simplicity. Define $\mathscr{F}^n$ to be the sigma-algebra generated by $x^0, W_{x^0}^1, ..., x^{n-1}, W_{x^{n-1}}^n$. This definition allows x^n to be a random variable adapted to past information; thus, the sequence $\{x^n\}$ may be obtained from an adaptive allocation rule (in fact, under most such rules, each x^n is also $\mathscr{F}^n$-measurable).

Let $N_x^n = \sum_{m=0}^{n-1} 1_{\{x^m = x\}}$ be the number of replications assigned to alternative x up to time n, taking $N_x^0 = 0$ by convention. Denote by

$$\theta_x^n = \frac{1}{N_x^n} \sum_{m=0}^{n-1} 1_{\{x^m = x\}} W_{x^m}^{m+1}, \tag{1}$$

$$\sigma_x^n = \frac{\lambda_x}{\sqrt{N_x^n}} \tag{2}$$

the sample mean (for alternative x) and its standard deviation. By convention, θ_x^0 may be any arbitrary value (representing some prior guess of μ_x).

Let $x_*^n = \arg\max_x \theta_x^n$ be the alternative with the highest estimated mean after n replications have been expended. We assume that, if no additional simulations can be conducted, x_*^n will be the "selected" alternative, i.e., the one reported as being the best. We say that "correct selection" occurs at time n if $x_*^n = x_*$, i.e., the selected alternative is also the best one. The probability of correct selection (PCS), written as $P(x_*^n = x_*)$, depends on the allocation rule used to assign the simulation budget (i.e., on the law of the sequence $\{x^n\}$); we denote a generic allocation rule by π.

Suppose that, under allocation π, the limit $\alpha_x = \lim_{n\to\infty} \frac{N_x^n}{n}$ exists a.s. and is non-zero for every x. If the observations $W_{x^n}^{n+1}$ are mutually independent, Glynn and Juneja (2004) proves that the limit

$$\Gamma^\alpha = -\lim_{n\to\infty} \frac{1}{n} \log P(x_*^n \neq x_*)$$

exists and can be expressed in terms of the proportions $\alpha = (\alpha_1, ..., \alpha_M)$. In words, the probability of *incorrect* selection converges to zero at an exponential rate where the exponent contains a constant multiplier Γ^α that depends on the limiting allocation α. The theoretical optimal allocation is the vector

$$\alpha^* = \arg \max_{\alpha \in \mathbb{R}_{++}^M, \sum_x \alpha_x = 1} \Gamma^\alpha,$$

which makes the probability of incorrect selection vanish as quickly as possible.

Under the additional assumption of normally distributed observations, Glynn and Juneja (2004) gives a complete characterization of the optimal proportions α_x^* in the form of the conditions

$$\left(\frac{\alpha_{x_*}^*}{\lambda_{x_*}}\right)^2 = \sum_{x \neq x_*} \left(\frac{\alpha_x^*}{\lambda_x}\right)^2, \tag{3}$$

$$\frac{(\mu_x - \mu_{x_*})^2}{\frac{\lambda_x^2}{\alpha_x^*} + \frac{\lambda_{x_*}^2}{\alpha_{x_*}^*}} = \frac{(\mu_y - \mu_{x_*})^2}{\frac{\lambda_y^2}{\alpha_y^*} + \frac{\lambda_{x_*}^2}{\alpha_{x_*}^*}}, \qquad \text{for all } x, y \neq x_*. \tag{4}$$

Clearly, (3)-(4) depend on the unknown performance values and cannot be known in advance to the decision-maker. One possible approach, widely adopted in the optimal computing budget allocation (OCBA) literature (Chen and Lee 2010), is to periodically substitute sample means in place of μ_x and solve the resulting approximations of (3)-(4). Even then, most modern work on OCBA-type methods prefers to introduce additional approximations to make the optimality conditions easier to solve; for instance, Pasupathy et al. (2014) gives a principled approach for constructing such approximations.

Since we focus on the EI class of methods, the remainder of this section briefly introduces the EI concept. EI adopts a Bayesian view of R&S; thus, to derive the method, we assume that $\mu_x \sim \mathcal{N}\left(\theta_x^0, \left(\sigma_x^0\right)^2\right)$, and that μ_x and μ_y are independent for $x \neq y$. The parameters $\left(\theta_x^0, \sigma_x^0\right)$ characterize our prior beliefs about the unknown value μ_x. For simplicity, let us suppose that $\sigma_x^0 = \infty$, indicating a non-informative prior; in that case, it is well-known (DeGroot 1970) that the posterior distribution of μ_x given $\mathscr{F}^n$ is normal with parameters $\left(\theta_x^n, \sigma_x^n\right)$ given by (1)-(2) above. Thus, under the Bayesian setting, *estimation* occurs exactly as in the frequentist setting, but *allocation* will be based on the purely Bayesian idea of uncertain beliefs.

One of the first EI methods, and arguably the best-known, was proposed by Jones, Schonlau, and Welch (1998). In the context of our R&S problem with independent normal priors, this method assigns $x^n = \arg\max_x v_x^n$, where

$$v_x^n = \mathbb{E}\left(\max\left\{\mu_x - \theta_{x_*^n}^n, 0\right\} \mid \mathscr{F}^n\right) \tag{5}$$

$$= \sigma_x^n f\left(-\frac{\left|\theta_x^n - \theta_{x_*^n}^n\right|}{\sigma_x^n}\right) \tag{6}$$

and $f(z) = z\Phi(z) + \phi(z)$ with ϕ, Φ being the standard normal pdf and cdf, respectively. From (5), we see that the EI criterion v_x^n may be viewed as a measure of the potential of the unknown value μ_x to improve over the current point estimate of the best value. The criterion is adapted to $\mathscr{F}^n$ and may be efficiently recomputed after each new observation using the closed-form expression in (6).

Ryzhov (2016) provided the first explicit characterization of the limiting allocation under EI, given by

$$\lim_{n\to\infty} \frac{N_{x_*}^n}{n} = 1, \tag{7}$$

$$\lim_{n\to\infty} \frac{N_x^n}{N_y^n} = \frac{\lambda_x^2 \left(\mu_y - \mu_{x_*}\right)^2}{\lambda_y^2 \left(\mu_x - \mu_{x_*}\right)^2}, \qquad \text{for all } x, y \neq x_*. \tag{8}$$

It is easy to see that (7)-(8) do not match (3)-(4) except in the limiting case where $\alpha_{x_*}^* \to 1$. However, this limiting case does not arise for any finite M, meaning that, under EI, the probability of incorrect selection will not converge at an exponential rate. Ryzhov (2016) also derives the limiting allocations for two other EI-type methods; however, although one of these does achieve an exponential rate, that rate is not the best possible.

3 ALGORITHM AND MAIN RESULTS

Section 3.1 provides background on the complete expected improvement criterion and explains our modifications to the procedure. Section 3.2 states our main results and sketches out the argument (the full details are omitted due to space considerations, but can be found in Chen and Ryzhov 2017).

3.1 Complete Expected Improvement

Recently, Salemi, Nelson, and Staum (2014) proposed a new variant of EI, called "complete EI" or CEI, in which (5) is replaced by

$$v_x^n = \mathbb{E}\left(\max\left\{\mu_x - \mu_{x_*^n}, 0\right\} \mid \mathscr{F}^n\right), \tag{9}$$

and (6) can be replaced by

$$v_x^n = \sqrt{\left(\sigma_x^n\right)^2 + \left(\sigma_{x_*^n}^n\right)^2} f\left(-\frac{\left|\theta_x^n - \theta_{x_*^n}^n\right|}{\sqrt{\left(\sigma_x^n\right)^2 + \left(\sigma_{x_*^n}^n\right)^2}}\right) \tag{10}$$

for $x \neq x_*^n$. The advantage of CEI over classic EI is that it incorporates additional uncertainty about the best value. Salemi, Nelson, and Staum (2014) developed this criterion for a more general Gaussian Markov

---

Step 0: Initialize $n=0$ and $N_x^n=0$, $\theta_x^n=0$, $\sigma_x^n=\infty$ for all x.

Step 1: Check whether

$$\left(\frac{N^n_{x^n_*}}{\lambda_{x^n_*}}\right)^2 < \sum_{x\neq x^n_*}\left(\frac{N^n_x}{\lambda_x}\right)^2. \tag{11}$$

Step 2: If (11) holds, assign $x^n = x^n_*$. Otherwise, assign

$$x^n = \arg\max_{x\neq x^n_*}\sqrt{\left(\sigma^n_x\right)^2+\left(\sigma^n_{x^n_*}\right)^2}\, f\left(-\frac{\left|\theta^n_x-\theta^n_{x^n_*}\right|}{\sqrt{\left(\sigma^n_x\right)^2+\left(\sigma^n_{x^n_*}\right)^2}}\right).$$

Step 3: Collect new information $W^{n+1}_{x^n}$, update posterior parameters using (1)-(2). Increment n by 1 and return to step 1.

Figure 1: Description of modified CEI (mCEI) procedure.

model with prior correlations between values, so the original presentation of CEI also included the posterior covariance between μ_x and $\mu_{x^n_*}$ in the calculation. In (10) we give a simpler version of CEI adapted to our R&S model with independent priors and observations.

The definition (9) implies that $v^n_{x^n_*}=0$ for all n. For this reason, our analysis makes an additional modification to the method. Essentially, we introduce a new condition to determine whether the next replication should be assigned to x^n_*; if this condition is not met, we then assign $x^n = \arg\max_{x\neq x^n_*} v^n_x$ where v^n_x is computed using (10). As far as we can tell, Salemi, Nelson, and Staum (2014) does not explicitly comment on the issue of assigning replications to x^n_*. However, in the broader literature, this is acknowledged as a difficult problem. For example, Russo (2017) observes that the popular Thompson sampling algorithm (Russo and Van Roy 2014), if applied to R&S, will sample x^n_* too often. Russo (2017) then proposes a correction which essentially assigns a pre-specified proportion β of the budget to x^n_*, and uses Thompson sampling (or an equivalent criterion) to choose between the remaining alternatives in all other cases. It follows that the optimal convergence rate can be achieved, but only if β is tuned properly. In light of this issue, it is especially interesting that the modified version of CEI achieves the optimal rate without any tuning.

Figure 1 states the procedure; the modification consists of condition (11), which is trivial to implement. It is easy to see that mCEI does not involve any tunable parameters. We begin our analysis by proving a basic consistency result for the mCEI procedure.

Theorem 1 Under mCEI, $N^n_x\to\infty$ a.s. for all x.

Proof. Fix a sample path ω and define $A=\{x\,|\,N^n_x(\omega)\to\infty\}$. It is obvious that A is non-empty. Suppose that A^c is also non-empty and that $x\in A^c$; then, the time

$$K_1=\sup\{n\,|\,x^n(\omega)=x\}$$

is finite. There must also exist a finite time $K_2\geq K_1$ such that, for all $n\geq K_2$, we have $x^n_*(\omega)\neq x$; if this were not the case, we would be able to find some $n>K_1$ for which $x^n_*(\omega)=x$ and (11) would hold, implying that $x^n(\omega)=x$ and contradicting the definition of K_1.

Since this analysis applies to any $x\in A^c$ and there are only finitely many such x, it follows that there exists $K_3\geq K_2$ such that both $x^n(\omega)\in A$ and $x^n_*(\omega)\in A$ for all $n\geq K_3$. Consequently, for $n\geq K_3$, we can write

$$\sigma^n_{x^n_*(\omega)}(\omega) \leq \max_{x\in A} \sigma^n_x(\omega). \tag{12}$$

The RHS of (12) vanishes to zero as $n \to \infty$ due to (2) and the definition of A. Therefore, $\sigma^n_{x^n_*(\omega)}(\omega) \to 0$. Since the CEI criterion satisfies

$$v^n_x \leq \frac{1}{\sqrt{2\pi}} \cdot \sqrt{(\sigma^n_x)^2 + \left(\sigma^n_{x^n_*}\right)^2},$$

it follows that, for any $x \in A$, we have $v^n_x(\omega) \to 0$. On the other hand, for any $x \in A^c$, $\lim_{n\to\infty} \sigma^n_x(\omega) > 0$ and it straightforwardly follows that $\liminf_{n\to\infty} v^n_x(\omega) > 0$.

It is obvious that there are infinitely many n for which condition (11) does not hold. From here, it is straightforward to show that we can find large enough $n \geq K_3$ for which (11) does not hold and

$$\max_{x\in A} v^n_x(\omega) < \min_{x\in A^c} v^n_x(\omega),$$

which would imply that $x^n(\omega) \in A^c$, contradicting the definition of K_3. We conclude that $A^c = \emptyset$. □

3.2 Main Results and Proof Sketch

We first introduce some additional notation. Let $\alpha^n_x = \frac{N^n_x}{n}$ be the proportion of the first n replications assigned to alternative x, and let

$$t^n_x = \frac{(\mu_x - \mu_{x_*})^2}{\frac{\lambda^2_x}{N^n_x} + \frac{\lambda^2_{x_*}}{N^n_{x_*}}}$$

for $x \neq x_*$. We now state our main results, which essentially show that conditions (3)-(4) are achieved in the limit by the mCEI procedure.

Theorem 2 (First optimality condition.) Under mCEI,

$$\lim_{n\to\infty} \left(\frac{\alpha^n_{x_*}}{\lambda_{x_*}}\right)^2 - \sum_{x\neq x_*} \left(\frac{\alpha^n_x}{\lambda_x}\right)^2 = 0,$$

almost surely.

Theorem 3 (Second optimality condition.) Under mCEI,

$$\lim_{n\to\infty} \frac{t^n_x}{t^n_y} = 1$$

almost surely for any $x, y \neq x_*$.

The proofs of these results are highly technical and outside the scope of the present paper; we refer readers to Chen and Ryzhov (2017) for the details. Below, we sketch the structure of the proof (particularly focusing on Theorem 3, which is more difficult).

First, our analysis of the procedure is frequentist, viewing μ_x as a fixed constant. To some degree, the frequentist and Bayesian versions of the R&S problem considered in this paper are interchangeable, since they both use (1)-(2) for estimation. Moreover, although Bayesian arguments are used to derive the CEI criterion, one could easily apply the closed-form expression in (10) in a frequentist setting. Since we prove almost sure convergence, and the value of μ_x is always fixed on a given sample path, this is not a major issue.

Second, by the arguments given in Section 4.1 of Ryzhov (2016), it is sufficient to prove Theorems 2-3 for a simplified version of mCEI with (11) replaced by

$$\left(\frac{N^n_{x_*}}{\lambda_{x_*}}\right)^2 < \sum_{x\neq x_*} \left(\frac{N^n_x}{\lambda_x}\right)^2,$$

and (10) replaced by

$$v_x^n = \sqrt{(\sigma_x^n)^2 + \left(\sigma_{x_*}^n\right)^2} f\left(-\frac{|\mu_x - \mu_{x_*}|}{\sqrt{(\sigma_x^n)^2 + \left(\sigma_{x_*}^n\right)^2}}\right). \tag{13}$$

That is, we replace x_*^n by its limit x_*, and similarly replace the posterior means by the true values. This version of mCEI cannot be implemented in practice, but it has the same asymptotic behaviour (on a fixed sample path) as the procedure in Figure 1; by the arguments in Ryzhov (2016), the error in the sample means does not substantially affect the tail behaviour of the function f.

From this point on, however, the analysis requires considerable new content over Ryzhov (2016): even after the simplification in (13), the CEI criterion does not decline independently across alternatives, as it would for classic EI. The difficulty comes from the fact that the ratio $\frac{v_x^n}{v_y^n}$, for $x, y \neq x_*$, can now change when x_* is sampled. To handle this problem, it is necessary to bound the number of samples of x_* in between two samples of x. The discussion below will summarize the bound while omitting the technical details.

By applying the Mills ratio (Ruben 1962), one can derive the expansion

$$\frac{\log v_x^n}{\log v_y^n} = \frac{t_x^n}{t_y^n} \cdot \frac{\left(1 + \mathscr{O}\left(\frac{\log t_x^n}{t_x^n}\right)\right)}{\left(1 + \mathscr{O}\left(\frac{\log t_y^n}{t_y^n}\right)\right)}. \tag{14}$$

Then, if the LHS of (14) can be shown to converge to 1, Theorem 3 will follow because $t_x^n \to \infty$ by Theorem 1. Since $v_x^n \to 0$ for all x, the LHS of (14) is equal to the ratio $\frac{|\log v_x^n|}{|\log v_y^n|}$ of the absolute values for large enough n. When either x or x_* is sampled, the absolute value $|\log v_x^n|$ will increase, which may affect the ratio of the absolute values. When y is sampled, the ratio will decrease, and when any other alternative z not equal to x, y or x_* is sampled, the LHS of (14) will be unaffected.

Thus, we consider a period of time in between two samples of x and bound the growth of $\frac{|\log v_x^n|}{|\log v_y^n|}$ during this period. Formally, suppose that $x^n = x$ and let $m = \min\left\{\ell > 0 \,|\, x^{n+\ell} = x\right\}$; then, the period of interest is between times n and $n+m$. During this period, x is measured exactly twice (at the endpoints), and we prove in Chen and Ryzhov (2017) that

$$\sum_{k=n}^{n+m} 1_{\{x^n = x_*\}} = \mathscr{O}(\log n). \tag{15}$$

The RHS of (15) is an upper bound; it is possible for the LHS to be zero for some periods between two samples of x.

It can also be proved that, for each sample of x_* occurring between times n and $n+m$, the increase in the ratio $\frac{|\log v_x^k|}{|\log v_y^k|}$ is at most of order $\mathscr{O}\left(\frac{1}{n}\right)$. Therefore, the total growth of the ratio during this period is of order $\mathscr{O}\left(\frac{\log n}{n}\right)$, which vanishes to zero as $n \to \infty$. Furthermore, since $\frac{|\log v_x^n|}{|\log v_y^n|} \leq 1$ at the start of the period (by the definition of n), and since $\limsup_{n\to\infty} \frac{|\log v_x^n|}{|\log v_y^n|} \geq 1$ because both x and y will be sampled infinitely often, we are able to conclude that $\limsup_{n\to\infty} \frac{|\log v_x^n|}{|\log v_y^n|} = 1$. Applying a symmetry argument to the reciprocal of the ratio, we find that the LHS of (14) converges to 1.

4 NUMERICAL EXAMPLE

We present a numerical illustration of the mCEI method on a small synthetic problem. Two additional benchmarks were implemented. The first of these is the classic EI method from Jones, Schonlau, and

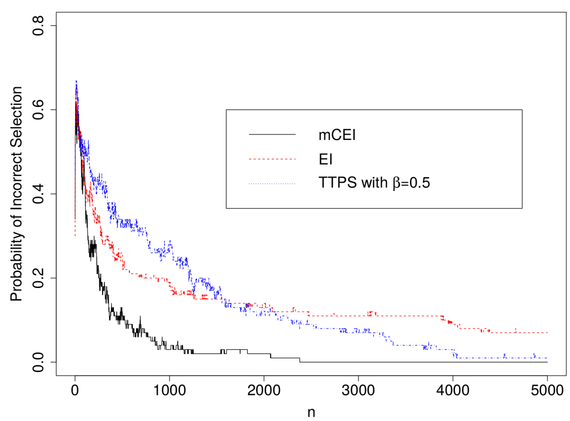

(a) Probability of incorrect selection.

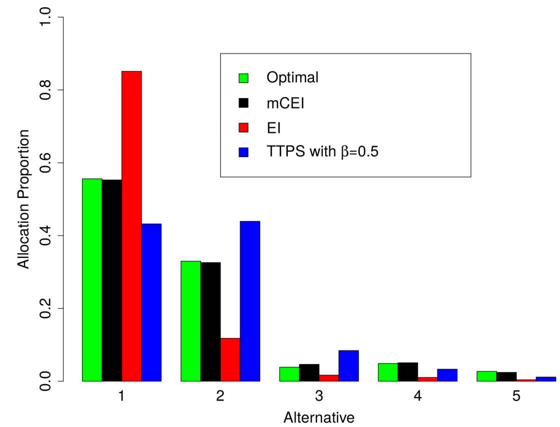

(b) Simulation allocations after 5000 samples.

Figure 2: Comparison between mCEI and benchmark methods on the example problem.

Welch (1998), given in (6). From (7)-(8), we do not expect this method to perform optimally in the long term; however, we include it because it is the fundamental procedure in the EI class of methods and thus a natural benchmark for mCEI. We also implemented the TTPS ("top-two probability sampling") method from Russo (2017). This method assigns a fixed proportion β of the sampling budget to alternative x_*^n and allocates the rest based on a Thompson sampling-like criterion. TTPS is an important benchmark since it can be made to achieve the optimal convergence rate if β is chosen correctly; however, since tuning β may be time-consuming in practice, Russo (2017) explicitly recommends setting $\beta = 0.5$ and derives a bound on the gap between the resulting convergence rate and the optimal one. We follow this recommendation in order to briefly comment on the tuning issue.

The synthetic example has five alternatives with true values $\mu = (0.5, 0.4, 0.3, 0.2, 0.1)$, standard deviations $\lambda = (1, 0.6, 0.6, 1, 1)$, the initial prior means $\theta_x^0 \equiv 0$, and a budget of 5000 samples. Figure 2(a) shows the trajectory of the probability of incorrect selection, averaged over 100 macro-replications. Thus, the best alternative is $x_* = 1$, but the noise is greater for alternative 1 than for alternatives 2 and 3, which makes correct selection a bit more difficult.

By (7), we know that EI will not be able to achieve an exponential convergence rate, so it is unsurprising that it is eventually outperformed by TTPS; however, EI performs relatively well in the early stages. On the other hand, mCEI lags slightly behind EI during the first 200 replications, but subsequently discovers the best alternative very quickly. After 2500 samples, the empirical probability of incorrect selection is virtually zero under mCEI.

Figure 2(b) compares the allocations made by each method (also averaged over 100 macro-replications) to the optimal allocation, obtained by solving (3)-(4). As expected from (7), the EI allocation is far from optimal since it assigns most of the budget to the best alternative. The optimal proportion to assign to alternative 1 is slightly larger than 0.5; as a result, TTPS is not tuned optimally and thus consistently makes errors in all of the proportions. The allocation made by mCEI is very close to optimal.

Note that, even in this small problem, alternatives 3, 4 and 5 receive only about 10% of the budget under the optimal allocation. This suggests that, in some situations, the size of the problem may not necessarily determine its difficulty (aside from increasing the computational effort required to run a procedure), as many or even most of the alternatives may be similarly "irrelevant." Identifying characteristics that make problems more "difficult" may be an interesting subject for future work. At present, however, we only wish to illustrate the potential of mCEI to produce very close approximations of the optimal allocation, without any tuning, in a relatively small number of samples.

5 CONCLUSION

We have presented the first theoretical guarantee of rate-optimality for a method belonging to the EI class, in the context of ranking and selection with independent normal samples. The method in question is a modified version of the CEI criterion, first proposed by Salemi, Nelson, and Staum (2014); we show that it achieves the rate-optimality conditions of Glynn and Juneja (2004) asymptotically without the need for any tuning. This provides new theoretical support for EI-type methods in general and CEI in particular.

Our results raise several questions for further study. First, one wonders whether analogous results may be obtained for the more powerful Gaussian Markov framework of Salemi, Nelson, and Staum (2014), for which CEI was originally proposed. Unfortunately, such models do not fit the framework of Glynn and Juneja (2004), and so more fundamental questions about convergence rates under these models should first be addressed. A second, related question is whether CEI can retain its optimality properties under non-Gaussian sampling distributions; Ding and Ryzhov (2016) showed that some EI-type methods may not even be consistent in these settings, so some additional modifications or considerations may be necessary.

Finally, while this paper was under review, a new preprint by Qin, Klabjan, and Russo (2017) appeared on arXiv. This work also provides rate-optimality guarantees for CEI using a different theoretical technique. However, these authors use a tunable parameter to control the assignment of replications to x_*^n, exactly as in the TTPS algorithm. The mCEI procedure described in our paper addresses this problem in a different way that avoids the need for tuning.

REFERENCES

Branke, J., S. E. Chick, and C. Schmidt. 2007. "Selecting A Selection Procedure". *Management Science* 53 (12): 1916–1932.

Chau, M., M. C. Fu, H. Qu, and I. O. Ryzhov. 2014. "Simulation Optimization: A Tutorial Overview And Recent Developments In Gradient-Based Methods". In *Proceedings of the 2014 Winter Simulation Conference*, edited by A. Tolk, S. Y. Diallo, I. O. Ryzhov, L. Yilmaz, S. Buckley, and J. A. Miller, 21–35. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Chen, C.-H., and L. H. Lee. 2010. *Stochastic Simulation Optimization: An Optimal Computing Budget Allocation*. World Scientific.

Chen, Y., and I. O. Ryzhov. 2017. "Rate-Optimality Of Complete Expected Improvement". *Working paper, University of Maryland*. Available online at: http://scholar.rhsmith.umd.edu/sites/default/files/iryzhov/files/2017_chry_cei.pdf.

Chick, S. E. 2006. "Subjective Probability And Bayesian Methodology". In *Handbooks of Operations Research and Management Science, vol. 13: Simulation*, edited by S. G. Henderson and B. L. Nelson, 225–258. North-Holland Publishing, Amsterdam.

Chick, S. E., J. Branke, and C. Schmidt. 2010. "Sequential Sampling To Myopically Maximize The Expected Value Of Information". *INFORMS Journal on Computing* 22 (1): 71–80.

DeGroot, M. H. 1970. *Optimal Statistical Decisions*. John Wiley and Sons.

Ding, Z., and I. O. Ryzhov. 2016. "Optimal Learning With Non-Gaussian Rewards". *Advances in Applied Probability* 48 (1): 112–136.

Fu, M. C. 2015. *Handbook Of Simulation Optimization*. Springer.

Glynn, P., and S. Juneja. 2004. "A Large Deviations Perspective On Ordinal Optimization". In *Proceedings of the 2004 Winter Simulation Conference*, edited by R. G. Ingalls, M. D. Rossetti, J. S. Smith, and B. A. Peters, 577–585. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Han, B., I. O. Ryzhov, and B. Defourny. 2016. "Optimal Learning In Linear Regression With Combinatorial Feature Selection". *INFORMS Journal on Computing* 28 (4): 721–735.

Hong, L. J., and B. L. Nelson. 2005. "The Tradeoff Between Sampling And Switching: New Sequential Procedures For Indifference-Zone Selection". *IIE Transactions* 37 (7): 623–634.

Hong, L. J., and B. L. Nelson. 2009. "A Brief Introduction To Optimization Via Simulation". In *Proceedings of the 2009 Winter Simulation Conference*, edited by M. Rosetti, R. Hill, B. Johansson, A. Dunkin, and R. Ingalls, 75–85. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Jian, N., and S. G. Henderson. 2015. "An Introduction To Simulation Optimization". In *Proceedings of the 2015 Winter Simulation Conference*, edited by L. Yilmaz, W. K. V. Chan, I. Moon, T. M. K. Roeder, C. Macal, and M. D. Rossetti, 1780–1794. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Jones, D. R., M. Schonlau, and W. J. Welch. 1998. "Efficient Global Optimization Of Expensive Black-Box Functions". *Journal of Global Optimization* 13 (4): 455–492.

Pasupathy, R., S. R. Hunter, N. A. Pujowidianto, L. H. Lee, and C.-H. Chen. 2014. "Stochastically Constrained Ranking And Selection Via SCORE". *ACM Transactions on Modeling and Computer Simulation* 25 (1): 1:1–1:26.

Peng, Y., and M. C. Fu. 2017. "Myopic Allocation Policy With Asymptotically Optimal Sampling Rate". *IEEE Transactions on Automatic Control* 62 (4): 2041–2047.

Powell, W. B., and I. O. Ryzhov. 2012. *Optimal Learning*. John Wiley and Sons.

Qin, C., D. Klabjan, and D. Russo. 2017. "Improving The Expected Improvement Algorithm". *arXiv preprint arXiv:1705.10033*.

Ruben, H. 1962. "A New Asymptotic Expansion For The Normal Probability Integral And Mill's Ratio". *Journal of the Royal Statistical Society* B24 (1): 177–179.

Russo, D. 2017. "Simple Bayesian Algorithms For Best Arm Identification". *arXiv preprint arXiv:1602.08448*.

Russo, D., and B. Van Roy. 2014. "Learning To Optimize Via Posterior Sampling". *Mathematics of Operations Research* 39 (4): 1221–1243.

Ryzhov, I. O. 2016. "On The Convergence Rates Of Expected Improvement Methods". *Operations Research* 64 (6): 1515–1528.

Salemi, P., B. L. Nelson, and J. Staum. 2014. "Discrete Optimization Via Simulation Using Gaussian Markov Random Fields". In *Proceedings of the 2014 Winter Simulation Conference*, edited by A. Tolk, S. Y. Diallo, I. O. Ryzhov, L. Yilmaz, S. Buckley, and J. A. Miller, 3809–3820. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Scott, W. R., W. B. Powell, and H. P. Simão. 2010. "Calibrating Simulation Models Using The Knowledge Gradient With Continuous Parameters". In *Proceedings of the 2010 Winter Simulation Conference*, edited by B. Johansson, S. Jain, J. Montoya-Torres, J. Hugan, and E. Yücesan, 1099–1109. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

AUTHOR BIOGRAPHIES

YE CHEN is a Ph.D. student of Statistics in the Department of Mathematics, University of Maryland. His research interests include applied probability, statistical learning, and stochastic optimization. He was a finalist in WSC's Best Theoretical Paper Award competition in 2016. His e-mail address is yechen@math.umd.edu.

ILYA O. RYZHOV is an Associate Professor of Operations Management and Management Science in the Decision, Operations and Information Technologies department of the Robert H. School of Business at the University of Maryland. His research mostly focuses on simulation optimization and optimal learning, with applications in business analytics; he is a coauthor of the book *Optimal Learning* (Wiley, 2012). He was recognized in WSC's Best Theoretical Paper Award competition on three separate occasions (winner in 2012, finalist in 2009 and 2016), and was also a finalist in the 2014 INFORMS Junior Faculty Forum paper competition. His email address is iryzhov@rhsmith.umd.edu.