SHOULD I BOTHER TO SIMULATE? AN ECONOMIC APPROACH TO DISCRETE OPTIMIZATION VIA SIMULATION

Stephen E. Chick

INSEAD
Technology and Operations Management Area
Boulevard de Constance
77305 Fontainebleau FRANCE

Noah Gans

OPIM Department – Wharton School
University of Pennsylvania
3730 Walnut Street, Suite 500
Philadelphia, PA 19104-6340 U.S.A.

ABSTRACT

Simulations cost time and money to develop and run, but simulation output can inform decisions by reducing the uncertainty about each potential alternative. Managers that use stochastic simulation to decide which of a finite set of alternatives to implement are faced with the following questions: Should I invest time and money to build simulation tools in order to evaluate the alternatives by using stochastic simulations? If so, for how long, and in what order, should the simulations be run before selecting one of the alternatives to implement? This paper discusses recent work to answer those questions, assuming that the manager seeks to maximize the expected net present value of the total analysis costs and the potential economic value of the alternative that is eventually implemented.

1 INTRODUCTION

We summarize recent work from Chick and Gans (2007) and potential directions for future work. They propose a new approach to the problem of selecting the best simulated alternative, one that focuses on the expected net present value of decisions that are supported by simulation, as opposed to more typical approaches that focus on statistical aspects involved with ranking and selection approaches to discrete optimization with simulation.

The premise is that managers decide the operating characteristics of their firm's business processes, such as for manufacturing or for service delivery. Often the decision reflects the choice of one among a number of competing designs. To aid their decision-making managers may use stochastic or discrete event simulation. Simulation is a widely-used and relatively low cost 'insurance' mechanism to estimate the performance of alternative systems and to improve the chances that the best system is implemented.

If there is a fixed set of k alternative designs, one must decide how long to simulate each alternative and, given the simulation results, which design to implement.

One approach for selecting the best of a finite set of simulated systems is ranking and selection, and the last 10 years have seen exciting progress. Both Bayesian and frequentist approaches are possible (Chick 2006, Chen et al. 2000, Kim and Nelson 2006), and large-scale numerical comparisons have identified strengths and weaknesses of each approach (Branke et al. 2007).

Another approach is to examine the long-run convergence of solutions to an optimal value, in probability or almost surely. Andradóttir (2006) provides a thorough discussion, and Hong and Nelson (2006) provide another example of recent innovations.

A third approach is to combine heuristics that were developed for deterministic optimization over large solution spaces, with some adaptation to account for randomness. Some widely-used tools use simple policies, such as using the sample average of a handful of replications as the actual mean value of the output. More sophisticated algorithms have been developed, too (Boesel, Nelson, and Kim 2003). See also Fu, Glover, and April (2005).

These approaches provide the flexibility to assess a wide variety of operational and other measures of system performance. The assessment does not usually consider the financial costs of the analysis itself.

When system and simulation results are themselves financial measures, as when simulation is used at a tactical or strategic rather than operational level, a more direct economic approach may be appropriate. If the manager's goal is to maximize the expected net present value (NPV) of high-level system design choices, then the manager is faced with two countervailing costs. On the one hand, uncertainty about the expected NPV of each alternative compels the manager to simulate more to reduce the opportunity costs associated with an incorrect selection. On the other, a lengthy simulation analysis incurs direct costs and reduces the NPV of the system that is ultimately implemented, due to discounting.

Two classes of costs appear. If the time and cost to develop the simulation tools that are required to perform the analysis are sunk, as is implicitly assumed by the main approaches that are described above, then the costs include the marginal cost of additional simulations and the discounting due to the time required for the replications. While such discounting may be small in many applications, there are a number of applications where simulations can take a long time, or the number of alternatives may result in a long simulation analysis.

A second class of costs is whether to invest the time and money that is required to develop the simulation tools themselves. These "up-front" costs are typically ignored within the simulation optimization community. But those costs are crucial for a manager that must decide whether a system, or set of systems, warrants the time and money that is required to reduce the uncertainty about their performance by developing simulation tools. Analysis is not free.

Given these costs for analysis, and the potential revenue streams that result from implementing one of the system alternatives, should a manager even bother to develop simulation tools at all? If so, for how long should simulations run, and in what order should systems be simulated, before an alternative is selected for implementation. Managers are trained to make such decisions on the basis of the expected NPV of the alternatives?

Chick and Gans (2007) formulate and solve a simulation selection problem in which the manager seeks to maximize expected NPV. Their formulation is Bayesian: it assumes that a manager has prior beliefs concerning the distribution of the NPV of each of the alternatives and that she uses simulation output to update these beliefs. The system which the manager ultimately chooses to implement maximizes expected NPV with respect to the posterior distributions of her beliefs, as well as analysis costs and discounting costs.

Section 2 defines the problem, assuming that up-front simulation costs are sunk. Section 3 observes that the simulation selection problem can be expressed as a variant of the multi-armed bandit problem (Gittins 1979), a so-called stoppable bandit process (Glazebrook 1979). As a result, a variant of the well known Gittins-index policy is an optimal way to decide which system to simulate next and when to stop simulating in favor of implementing a system. At each step of the sequential procedure, the system with the largest Gittins index is identified. If the statistics of the system with the largest Gittins index fall into a continuation set, then that system is simulated one more time and its Gittins index is updated. If the statistics fall outside of the continuation set, then simulation stops and that system is implemented. The boundary of the continuation set therefore determines whether one should "learn" (by simulating) or "earn" (by implementing a system). Section 3 requires few distributional assumptions other than joint independence and bounded expectations. Simulation run times of the different systems are initially assumed to be equal.

Gittins indices are typically difficult to compute exactly in Bayesian problems. Section 4 overviews asymptotic approximations for the Gittins index associated with normally distributed simulation output. The approximations apply when simulation output is independent across systems with unknown means and known, potentially different, variances. The approximation is determined by the solution of a free boundary problem for a heat equation that shares characteristics with financial options, and that draws upon work on composite hypothesis tests initiated by (Chernoff 1961).

In summary, this paper suggests that an alternative approach to discrete optimization via simulation may be useful–that of linking financial measures (discount rate, simulation development costs, marginal cost of simulations) to the optimal control of simulation experiments that are designed to inform operational decisions. The area is ripe for further research. See Chick and Gans (2007) for a more description, proofs, and numerical examples of the ideas outlined below, along with some suggestions for a number of interesting research questions to address.

2 PROBLEM DESCRIPTION

A manager seeks to develop one of k projects, labelled $i = 1, \ldots, k$. The net present value (NPV) of each of the i projects is not known with certainty, however. The manager wishes to develop the project which maximizes her expected NPV, or to do nothing if the expected present value of all projects is negative. We represent the "do nothing" option as $i = 0$ with a sure NPV of zero.

2.1 Uncertain Project NPV's

Let X_i be the random variable representing the NPV of project i, where $X_0 \equiv 0$. If the manager is risk neutral and the distributions of all X_i's are known to her, then she will select the project with the largest expected NPV, $i^* = \arg\max_i\{\mathrm{E}[X_i]\}$.

Although we model NPVs as simple random variables, the systems that generate them may be quite complex. For example, a particular project's sequence of cash flows may involve the composition of several interrelated random processes describing the evolution of investments, $\mathcal{I}(t)$, revenues, $\mathcal{R}(t)$, and operating costs, $\mathcal{O}(t)$, over time. Nevertheless, given a continuous-time discount rate $\delta > 0$, each realization of these processes, ω_i, yields a sample $X(\omega_i) = \int_0^\infty [\mathcal{R}(t, \omega_i) - \mathcal{O}(t, \omega_i) - \mathcal{I}(t, \omega_i)]e^{-\delta t}dt$.

It may also be the case that the distributions of the X_i's are not known with certainty by the manager. Rather, she may believe that a given X_i may come from one of a family of probability distributions, $P_{X_i|\theta_i}$, indexed by parameter θ_i. We model her belief as taking the form of a probability distribution on θ_i, which we call P_{Θ_i}. For example, the manager may believe that X_i is normally distributed with a known variance, σ_i^2, but unknown mean. Then P_{Θ_i} represents a probability distribution for the mean. To ease notation, we will sometimes refer to the distribution as Θ_i. In this case, the expected NPV of project $i > 0$ is $\mathrm{E}[X_i] = \mathrm{E}[X(\Theta_i)] \triangleq \iint X(\theta_i) dP_{X_i|\theta_i} dP_{\Theta_i}$. We denote the vector of distributions for the projects by $\mathbf{\Theta} = (\Theta_1, \ldots, \Theta_k)$.

2.2 Simulation to Select the Best Project

If the distributions of the X_i's are not known, then the manager may be able to use simulation as a tool to reduce distributional uncertainty, before having to decide which project to develop. She may decide to simulate the outcome of project i a number of times, and she views the result of each run as a sample of X_i. She uses Bayes' rule to update her beliefs concerning Θ_i.

We model the running of simulations as occurring at sequence of discrete stages $t = 0, 1, 2, \ldots$, and we represent Bayesian updating of prior beliefs and sample outcomes, $\{(\mathbf{\Theta}_t, \mathbf{X}_t) \,|\, t = 0, 1, \ldots\}$ as follows. If project $i > 0$ is simulated at stage t with sample outcome $x_{i,t}$, then $X_{i,t} = x_{i,t}$ and $X_{j,t} = 0$ for all $j \neq i$. In turn, Bayes' rule is used to determine $\mathbf{\Theta}_{t+1}$:

$$dP_{\Theta_{i,t+1}}(\theta_i \,|\, x_{i,t}, \Theta_{i,t}) \;=\; \frac{dP_{X_i\,|\,\theta_i}(x_{i,t} \,|\, \theta_i)\, dP_{\Theta_{i,t}}(\theta_i)}{\int_{\theta_i} dP_{X_i\,|\,\theta_i}(x_{i,t} \,|\, \theta_i)\, dP_{\Theta_{i,t}}(\theta_i)}$$

for all $\theta_i \in \Omega_{\Theta_i}$, while $\Theta_{j,t+1} = \Theta_{j,t}$ for all $j \neq i$. So the evolution of the manager's beliefs regarding the distribution of outcomes of each project is Markovian. We also assume that simulation results, hence the evolution of the manager's beliefs, are independent from one project to the next.

If, in theory, simulation runs could be performed at zero cost and in no time, then the manager might simulate each of the k systems infinitely, until all uncertainty regarding the θ_i's was resolved. At this point the problem would revert to the original case in which the distributions and means of the X_i are known.

But simulation runs do take time and cost money. We assume that each run of system i costs $\$c_i$ and takes η_i units of time to complete. Thus, given a continuous-time discount rate of $\delta > 0$, the decision to simulate system i costs the manager c_i plus a reduction of $\Delta_i = \int_0^{\eta_i} e^{-\delta s} ds < 1$ times the expected NPV of the (unknown) project that is eventually chosen.

There may also be associated up-front costs associated with the development of the simulation tool, itself. For example it may cost time and money to develop the underlying simulation platform, independent of which projects end up being evaluated. Additional costs may be required to be able to simulate particular projects. Furthermore, these project-specific costs may be inter-related.

In order to determine whether it is optimal to incur up-front costs, one needs to determine the value of having the simulation tools themselves. This is typical when solving stochastic dynamics programs – one solves for the optimal policy of decisions at later stages first, given the outcomes that result from earlier decisions, then recursively works backwards through the natural sequence of decisions.

To start, then, we make two simplifying assumptions regarding the costs of simulation. First, we ignore all up-front costs for the simulation tool, assuming that the necessary facilities exist to simulate all k projects. Second, we assume that $\eta_i \equiv \eta$ for all k projects. This allows us to define a common $\Delta \equiv \Delta_i$ for the projects as well. Section 6 argues that these assumptions may be relaxed.

Even with these simplifications, the availability of a simulation tool to sample project outcomes makes the manager's problem much more complex. Rather than simply choosing the project that maximizes expected NPV, she must choose a sequence of simulation runs and, ultimately select a project, so that the discounted stream of costs and terminal expected value, together, maximize expected NPV.

We define a number of indices in order to track the manager's choices as they proceed. Let $T \in \{t = 0, 1, 2, \ldots\}$ be the stage at which the manager selects a system to implement. For $t < T$, define $i(t) \in \{1, \ldots, k\}$ to be the index of the project simulated at time t, and define $I(T) \in \{0, \ldots, k\}$ to be the ultimate choice of project.

A *selection policy* is the choice of a sequence of simulation runs, a stopping time, and a final project. Define Π to be the set of all *non-anticipating* selection policies, whose choice at time $t = 0, 1, \ldots$ depends only on system history up to t: $\{\mathbf{\Theta}_0, \mathbf{X}_0, \ldots, \mathbf{\Theta}_{t-1}, \mathbf{X}_{t-1}, \mathbf{\Theta}_t\}$. Given prior distributions $\mathbf{\Theta} = (\Theta_1, \ldots, \Theta_k)$ and a policy $\pi \in \Pi$, the expected discounted value of the future stream of rewards is

$$V^{\pi}(\mathbf{\Theta}) \;=\; \mathrm{E}_{\pi}\left[\sum_{t=0}^{T-1} -\Delta^t c_{i(t)} \;+\; \Delta^T X_{I(T),T} \;\,|\, \mathbf{\Theta}_0 = \mathbf{\Theta}\right], \tag{1}$$

where $X_{I(T),T}$ is the unknown NPV of the selected system, $I(T)$, when a system is selected (at time T).

Formally, we define the manager's *simulation selection problem* to be the to choice of a selection policy $\pi^* \in \Pi$ that maximizes $V^{\pi^*}(\mathbf{\Theta}) = \sup_{\pi \in \Pi} V^{\pi}(\mathbf{\Theta})$.

The solution to the simulation selection problem determines the optimal expected value to running simulations. The expected value of the objective function can then be compared with the up-front costs, in order to determine whether or not a manager should invest in the development of simulation tools themselves.

3 SIMULATION SELECTION AND BANDITS

The simulation selection problem is closely related to a class of sequential decision problem known as the multi-armed bandit problem. Chick and Gans (2007) demonstrate that simulation selection problems can be reduced to multi-armed bandits. More precisely, they show that the simulation selection problem can be expressed as a stoppable family of alternative bandit processes, if up-front costs are sunk. They also show that the sufficient conditions of Glazebrook (1979) for index-based policies to be optimal are satisfied, given certain relatively weak technical conditions (e.g., bounded expected one-period rewards).

The result implies that a class of simple, index-based policies is optimal for simulation selection. They also prove that there is a Gittins index policy that is optimal, *and* that the Gittins index for each arm is proportional to the optimal value function in Equation (2).

$$
\begin{aligned}
&V_i^{\pi_i^*}(\Theta_{i,t}) \\
&\quad= \max\{-c_i + \Delta \mathrm{E}[V_i^{\pi_i}(\Theta_{i,t+1}) \,|\, \Theta_{i,t}, t \neq T_i], \\
&\qquad\quad (1-\Delta)\mathrm{E}[X(\Theta_{i,t})] + \Delta \mathrm{E}[V_i^{\pi_i^*}(\Theta_{i,t})]\Big\} \\
&\quad= \max\{-c_i + \Delta \mathrm{E}[V_i^{\pi_i}(\Theta_{i,t+1}) \,|\, \Theta_{i,t}, t \neq T_i], \\
&\qquad\quad \mathrm{E}[X(\Theta_{i,t})]\}\,. \qquad (2)
\end{aligned}
$$

The optimal value function, if it can be computed for each project individually, can therefore serve as a Gittins index, as is needed to optimally solve the simulation selection problem. The Gittins index results simplifies the problem of identifying the optimal policy for a k-dimensional problem into the solution of k selection problems, each of which has a single alternative.

4 GITTINS INDEX APPROXIMATION FOR NORMAL OUTPUT WITH KNOWN VARIANCE

At a high level, the optimal policy is straightforward. At each t it compares the expected discounted value of optimal stopping for each project and selects the one with the highest $V_i^{\pi_i^*}(\Theta_{i,t})$. If that project is $i=0$, then abandonment is most favorable and no project is further simulated or implemented. If the best project is some $i>0$ and $t=T_i^*$ then project i is implemented. Otherwise, project i is simulated, Bayes' rule is used to calculate $\Theta_{i,t+1}$, $V_i^{\pi_i^*}(\Theta_{i,t+1})$ is determined, and the comparison begins again.

The foundation of the optimal policy is the repeated determination of the various $V_i^{\pi_i^*}(\Theta_{i,t})$'s, as in Equation (2). The calculation of each $V_i^{\pi_i^*}(\Theta_{i,t})$, what we call the optimal expected discounted reward (OEDR), is itself a difficult task for which exact solutions are not available.

Chick and Gans (2007) develop diffusion approximations, $\mathcal{V}_i$, for the OEDRs in the spirit of (Chernoff 1961). The diffusion approximations are asymptotically appropriate when the discount rate over the duration of a simulation replication is small, as is usually the case in simulation. Repeated sampling leads to realizations of a scaled Brownian motion with drift.

Approximating the OEDR involves the solution of a so-called free boundary problem for a heat equation that is obtained from the diffusion approximation. The boundary is "free" since it is determined by equating the two maximands in the value function, rather than on a known, pre-specified boundary. A comparison of the maximands in the continuous-time analogue of Equation (2) determines the free boundary between a continuation set, $\mathcal{C}$, in which it is optimal to continue simulating a project, and a stopping set, in which it is optimal to stop simulating and implement the project. If the boundary is never reached, then the NPV of simulating forever is better than the expected NPV of implementing a poor system.

Their approximations assume that the simulation output for a given alternative is independent and normally distributed with a known variance, σ^2, and unknown mean, θ (we drop subscripts that identify the system to simplify notation, as the Gittins index result allows each system to be handled indvidually). While this assumption may not satisfy the uniform boundedness condition, the analysis below results in a well-defined finite OEDR when the initial prior distributions for the unknown means are proper. We suppose that θ has a $\texttt{Normal}\left(\mu_0, \sigma_0^2\right)$ prior distribution. Some of these restrictive assumptions can be relaxed.

Define $n_0 = \sigma^2/\sigma_0^2$ (the so-called effective number of samples in the prior distribution), and redefine $t = n_0 + n$, where n is the number of simulation observations seen so far for the single system in question.

Chick and Gans (2007) show:

- The diffusion approximation to the Gittins index (via the free boundary problem) is reasonable under conditions that are typically valid in simulation.
- The solution involves solving what might be called the value and the optimal exercise time of a perpetual American call option on regular (not geometric) Brownian motion, with unknown drift that is inferred through time.
- The free boundary problem has three parameters, $\delta > 0$, $c \geq 0$, and $\sigma > 0$. One can re-parameterize the problem so that only *one standardized* problem must be solved to handle *any* values of those parameters.
- The OEDR of the standardized problem is a function of the mean of θ and the effective number of replications (initially μ_0 and t_0).
- Numerical techniques implement the necessary calculations in a tractable amount of time.

Table 1: Simulation Selection Procedure

1. Identify economic parameters (discount factor $\delta > 0$ and costs $c_i \geq 0$ per replication). Provide prior distributions for each of the k alternative systems and initialize y_i, t_i for each system (see below). Include system 0 as an option ('do nothing' option with a guaranteed NPV of $\mathcal{V}_0 = 0$) if appropriate.
2. Compute the OEDR $\mathcal{V}_{\ell,i}$ for each alternative $i = 0, 1, \ldots, k$, and set $t = \sum_{i=1}^{k} t_i$, where $\ell \in \{1, 2(\kappa)\}$, depending on the values of c_i, δ, σ_i.
3. While (a system is not yet implemented):
   (a) Increment $t \leftarrow t + 1$.
   (b) Identify the system with largest index, $i(t) = \arg\max_{i=0,1,\ldots,k} \mathcal{V}_{\ell,i}$ (break ties randomly).
   (c) If $(y_{i(t)}, t_{i(t)})$ is not in the continuation set, $\mathcal{C}_{\ell,i}$, then stop simulating and implement the appropriate system, otherwise run a simulation for that system to get output $x_{i,t_{i(t)}+1}$.
   (d) Update $y_{i(t)} \leftarrow y_{i(t)} + x_{i,t_{i(t)}+1}$; $t_{i(t)} \leftarrow t_{i(t)} + 1$ and $\mathcal{V}_{\ell,i(t)}$ for system $i(t)$.

5 SIMULATION SELECTION PROCEDURE

Table 1 is the simulation selection procedure that results from the above analysis, which assumes independent, normally distributed output with known variances.

Step 1 of the procedure requires prior distributions for the unknown means. There are at least two options for generating these initial priors. Expert judgment may provide a prior distribution, $\texttt{Normal}(\mu_{0i}, \sigma_{0i})$, for the unknown means θ_i, for $i = 1, 2, \ldots, k$. If that is done, initialize $t_i = \sigma_i^2/\sigma_{0i}^2$ and $y_i = t_i\mu_{0i}$. Alternatively, default assessments can be implemented by running n_0 replications for each system and setting $y_i = \sum_{j=1}^{n_0} x_{i,j}$ and $t_i = n_0$.

Chick and Gans (2007) describe how to compute the OEDR $\mathcal{V}_{\ell,i}$ and the boundary of the continuation set.

It is possible that the system being implemented is the 'do nothing' option, which has a known NPV of $\mathcal{V}_0 = 0$. Alternative systems with a known, positive expected NPV can be included by replacing the 0-arm with the option of implementing that better alternative if the maximum OEDR is $\mathcal{V}_0 =$ known expected NPV (e.g. for comparing mutually exclusive alternatives with an existing system).

6 EXTENSIONS AND DISCUSSION

Section 4 assumed jointly independent Gaussian output with known variances. Chick and Gans (2007) argue that those approximations can be generalized to the following scenarios, if a few additional technical conditions hold.

- Samples from a one-parameter member of the exponential family of distributions can be handled (exponential, Bernoulli, Poisson, ...)
- Autocorrelated output, if a "batch means" analysis is appropriate. Such autocorrelation is typical for the analysis of many queueing or inventory systems.
- Different runtime durations across systems.

We also have some results that show how up-front costs, in certain simplified simulation analysis problems, can be used to answer whether the manager should optimally invest in developing simulation tools or not. In particular, when the simulation tools for each alternative must be developed separately, or when there is only a single alternative, we can quantify the tradeoffs between the prior uncertainty about the unknown mean NPV of a simulated alternative (modeled by t_0, μ_0), and the economic parameters (up-front costs, time to implement tools, the firm's discount rate).

The Gittins index results apply when samples are independent and normally distributed with *unknown* variance, but the approximations of Section 4 do not apply when the variance is unknown. A nonoptimal *ad hoc* solution would be to plug in the sample variance for the true variance, or apply some fudge factor (e.g. by plugging in the variance of a Student random variable with the appropriate degrees of freedom, $\nu = t_i - 1$, for the known variance). A better approximation for the Gittins index of simulation selection problems when the variance is unknown would be useful. The ability to model the expected reward of a simulation selection problem when there are multiple alternatives (not just when there is one alternative) would help a manager assess whether or not it is economically optimal to invest in more complicated simulation tools, such as those when a single simulation platform allows a large number of simulations to be simulated (e.g., with each simulation differing in its definition by a different set of input/design parameters).

ACKNOWLEDGMENTS

The research of Noah Gans was supported by the Fishman-Davidson Center for Service and Operations Management. Both authors appreciate the support of the Wharton-INSEAD Alliance.

REFERENCES

Andradóttir, S. 2006. An overview of simulation optimization via random search. In *Handbook in Operations Research and Management Science: Simulation*, ed. S. Henderson and B. Nelson. Elsevier.

Boesel, J., B. L. Nelson, and S.-H. Kim. 2003. Using ranking and selection to 'clean up' after simulation optimization. *Operations Research* 51:814–825.

Branke, J., S. Chick, and C. Schmidt. 2007. Selecting a selection procedure. *Management Science* 53:to appear.

Chen, C.-H., J. Lin, E. Yücesan, and S. E. Chick. 2000, jul. Simulation budget allocation for further enhancing the efficiency of ordinal optimization. *Discrete Event Dynamic Systems: Theory and Applications* 10 (3): 251–270.

Chernoff, H. 1961. Sequential tests for the mean of a normal distribution. In *Proc. Fourth Berkeley Symposium on Mathematical Statistics and Probability*, Volume 1, 79–91. Univ. California Press.

Chick, S. E. 2006. Subjective probability and Bayesian methodology. In *Handbook in Operations Research and Management Science: Simulation*, ed. S. Henderson and B. Nelson. Elsevier.

Chick, S. E., and N. Gans. 2007. An economic analysis of simulation selection problems. *Management Science* 53:to appear.

Fu, M., F. W. Glover, and J. April. 2005. Simulation optimization: A review, new developments, and applications. In *Proc. 2005 Winter Simulation Conference*, ed. M. Kuhl, N. Steiger, F. Armstrong, and J. Joines, 83–95. Piscataway, NJ: IEEE, Inc.

Gittins, J. C. 1979. Bandit problems and dynamic allocation indices. *J. Royal Statistical Society, Series B* 41:148–177.

Glazebrook, K. D. 1979. Stoppable families of alternative bandit processes. *Journal of Applied Probability* 16:843–854.

Hong, L. J., and B. L. Nelson. 2006. Discrete optimization via simulation using COMPASS. *Operations Research* 54:115–129.

Kim, S.-H., and B. L. Nelson. 2006. Selecting the best system. In *Handbook in Operations Research and Management Science: Simulation*, ed. S. Henderson and B. Nelson. Elsevier.

AUTHOR BIOGRAPHIES

STEPHEN E. CHICK is a Professor of Technology and Operations Management at INSEAD. He has worked in the automotive and software sectors prior to joining academia, now teaches operations management, and is interested in the health sector. He enjoys Bayesian statistics, stochastic models, and simulation.

NOAH GANS is an Associate Professor in the OPIM Department at the Wharton School. He worked in business consulting before entering academia, and now teaches service operations management. He is interested in call center operations and enjoys stochastic models and applied probability.