

Variable-Sample Methods for Stochastic Optimization

TITO HOMEM-DE-MELLO
Ohio State University

In this article we discuss the application of a certain class of Monte Carlo methods to stochastic optimization problems. Particularly, we study *variable-sample* techniques, in which the objective function is replaced, *at each iteration*, by a sample average approximation. We first provide general results on the *schedule* of sample sizes, under which variable-sample methods yield consistent estimators as well as bounds on the estimation error. Because the convergence analysis is performed pathwisely, we are able to obtain our results in a flexible setting, which requires mild assumptions on the distributions and which includes the possibility of using different sampling distributions along the algorithm. We illustrate these ideas by studying a modification of the well-known *pure random search* method, adapting it to the variable-sample scheme, and show conditions for convergence of the algorithm. Implementation issues are discussed and numerical results are presented to illustrate the ideas.

Categories and Subject Descriptors: G.1.6 [Numerical Analysis]: Optimization—*global optimization*; G.3 [Probability and Statistics]: *Probabilistic algorithms (including Monte Carlo)*; I.6.1 [Simulation and Modeling]: Simulation Theory

General Terms: Algorithms, Design, Theory

Additional Key Words and Phrases: Monte Carlo methods, pathwise bounds, random search, stochastic optimization

1. INTRODUCTION

In the past few years, a great deal of attention has been devoted to theoretical and practical aspects of combining *simulation* and *optimization* techniques to solve practical problems. This is required in problems where the objective function cannot be evaluated exactly but rather must be estimated by simulation. Indeed, many simulation packages have now optimization procedures implemented in the software. Conversely, there has been considerable amount of research on optimization methods that incorporate sampling in order to deal with uncertainty factors.

This research was supported in part by a SEED grant from The Ohio State University and by National Science Foundation (NSF) Grant DMI-0115385.

Author's address: Department of Industrial, Welding and Systems Engineering, The Ohio State University, 1971 Neil Ave., Columbus, OH 43210-1271; email: homem-de-mello.1@osu.edu.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 1515 Broadway, New York, NY 10036 USA, fax: +1 (212) 869-0481, or permissions@acm.org.

© 2003 ACM 1049-3301/03/0400-0108 \$5.00

Perhaps the most common (and fairly general) way to obtain a model that captures the existing randomness is by defining a *random function* of the underlying parameters on a proper probability space and then optimizing the *expected value* of such function with respect to the decision variables. More formally, we have a probability space (Ω, \mathcal{F}, P) , a subset $X \subset \mathbb{R}^m$, a (measurable) function $G : X \times \Omega \rightarrow \mathbb{R}$ and we want to solve

$$\min_{x \in X} \left\{ g(x) := \mathbb{E}[G(x)] = \int_{\Omega} G(x, \omega) P(d\omega) \right\}. \quad (1)$$

Typically, the expected value in problem (1) cannot be computed exactly, so approximation methods are required. In the context of a simulation model, $\mathbb{E}[G(x)]$ may represent some performance measure. For example, x could denote the vector of mean service times in a queuing network, and $G(x)$ the average waiting time of the first 100 customers, or the maximum number of customers in the system up to some pre-specified time T . In such cases, one should keep in mind that a sample point $G(x, \omega_i)$ can be obtained by running one replication of the model. Such multiple-replication scheme renders our framework more suitable to *transient* simulation models; nevertheless, the ideas can in principle be applied to steady-state systems, see below.

One approach to approximate the function in (1) is to resort to *Monte Carlo methods*: in its basic form, the idea is to replace the expected value function with its corresponding Monte Carlo approximation

$$\hat{g}_N(x) = \frac{1}{N} \sum_{i=1}^N G(x, \omega_i),$$

(where $\omega_1, \dots, \omega_N$ form an independent and identically distributed sample) and then solve the resulting deterministic problem. This type of algorithm (sometimes called “sample path optimization”, or “sample average approximation”) has been well studied in the literature, see for instance [Rubinstein and Shapiro 1993] and references therein. One advantage of such method is its nice convergence properties; indeed, it is possible to show convergence of optimal solutions and optimal values under fairly general assumptions (see, e.g., Dupačová and Wets [1998], Robinson [1996], and Shapiro [1991; 1993]). In some cases, the solution of the approximating problem converges *exponentially fast* on the sample size N to a solution of the original problem. This type of phenomenon was observed by Shapiro and Homem-de-Mello [2000] in the context of piecewise linear convex stochastic programs, and studied by Kleywegt et al. [2001] in the context of discrete optimization.

The basic idea in the Monte Carlo method yields several possible variations. For example, suppose we have at hand an iterative method to solve the deterministic problem. Instead of fixing a sample from the beginning and then minimizing the resulting deterministic function, one may consider using *different samples* along the algorithm. That is, the idea is to use, at iteration k , the approximating function

$$\hat{g}_k(x) := \frac{G(x, \omega_1^k) + \dots + G(x, \omega_{N_k}^k)}{N_k},$$

where $\omega_1^k, \dots, \omega_{N_k}^k$ is a sample from some distribution P_k close to P . Notice that we use a new sample at every iteration (hence, the term “variable-sample method” used throughout the article) as opposed to fixing a large sample at the beginning and then optimizing the resulting deterministic function. The measure P_k is defined by the “user”, for example, by the underlying optimization algorithm being studied. It could represent, for instance, an approximation to an unknown steady-state distribution.

One apparent advantage of a variable-sample scheme over the basic Monte Carlo method is that, since we generate independent estimates of the objective function at different iterations, we avoid getting “trapped” in a single sample-path. For example, as the results in Kleywegt et al. [2001] and Shapiro and Homem-de-Mello [2000] show, some problems have the property that, for almost all ω , there exists a number $N_0 = N(\omega)$ such that the solution $x_{N_0}^*$ of the approximating problem $\min \hat{g}_{N_0}(x)$ coincides with the solution x^* of the original problem (1). Such N_0 , however, is difficult to determine in practice, so for a given sample size N there may be a positive probability that x_N^* is actually far away from x^* —which in turn implies the existence of “bad” sample-paths. This effect tends to be reduced once we generate independent estimates of the objective function, since resampling allows the iterates of the algorithm to get away from those “bad” sample paths.

Another advantage of a variable-sample scheme is that the sample sizes can increase along the algorithm, so that sampling effort is not wasted at the initial iterations of the algorithm. Also, because the estimates at different iterations are independent, one can perform *statistical tests* to compare those estimates, which in turn can lead to *stopping criteria* for the algorithm. Indeed, this type of approach has been successfully used in some gradient-based methods for continuous stochastic optimization; see, for instance, Cooper and Homem-de-Mello [2003], Homem-de-Mello et al. [1999] and Shapiro and Homem-de-Mello [1998].

The price to pay for the flexibility provided by a variable-sample scheme, of course, is that the function being optimized changes at every iteration. Therefore, the convergence results developed for the sample average approximation described above are no longer valid. For example, it is important to ensure that $\hat{g}_k(x) \rightarrow g(x)$ with probability one (w.p.1)—that is, it is desirable that $\hat{g}_k(x)$ be a *consistent* estimator of $g(x)$. Perhaps surprisingly, it turns out that, for such a property to hold, it is not enough that the sequence of sample sizes $\{N_k\}$ be increasing; as we show in Section 3, N_k must grow at a certain rate.

While, in some cases, consistency of estimators suffices to show convergence of a method adapted to the variable-sample scheme, in other situations we need more than that; in those cases, we must ensure that the error from the deterministic algorithm dominates the stochastic error $|\hat{g}_k(x) - g(x)|$, so that the convergence properties of the deterministic algorithm are preserved. While this assertion is quite intuitive, showing that such property holds for a given algorithm can be a difficult task. This can be made easier by imposing a proper *schedule* of sample sizes, so that we can bound the stochastic error $|\hat{g}_k(x) - g(x)|$.

In this article, we address these issues. We propose a framework to analyze methods that use the variable-sample scheme. In particular, we focus on

algorithms that use function evaluations only. We provide general results on consistency of estimators as well as bounds on $|\hat{g}_k(x) - g(x)|$ under the variable-sample scheme. These goals are accomplished by exploiting the fact that the estimates $\hat{g}_k(x)$ of the objective function $g(x)$ are obtained via averaging, which allows us to use some classical results from probability theory. We obtain generic bounds on the deviation $|\hat{g}_k(x) - g(x)|$, which can then be used to show convergence of a specific method. Such conditions translate directly to the choice of sample size N_k used to compute the average at each iteration. An important aspect of our study is that the analysis is completely performed in terms of *sample paths*. This is why it is necessary to resort to tools such as the law of the iterated logarithm instead of the more commonly used Central Limit Theorem. By doing so, we do not need to assume any distribution for the error $|\hat{g}_k(x) - g(x)|$.

Another feature of our analysis is that we allow the use of different sampling distributions P_k at each iteration to obtain the estimate \hat{g}_k . This feature can be exploited in several ways, for example by using sampling methods that yield *variance reduction* for the resulting estimators. Notice that the sampling distributions P_k can even yield *biased* estimators of $g(x)$, as long as the bias goes to zero at a specific rate. This underscores the generality of the results.

The above results provide a set of tools that can be used to show convergence of a variable-sample version of a specific deterministic optimization algorithm. Such convergence analysis can in principle be applied to discrete or continuous problems. However, because it deals solely with convergence of function values, the analysis cannot be applied directly to methods that use *gradient* information such as the steepest descent algorithm, for example. Indeed, in such cases one must impose conditions on the estimates of the gradients in order to have convergence; see Shapiro and Wardi [1996]. Nevertheless, the analysis can be conceivably used in methods where derivatives are approximated by *finite differences*, though we do not explore that road here.

A natural setting for methods that are based exclusively on function evaluations is that of *discrete stochastic optimization* problems. This class of models of the form (1) consists of problems where the feasibility set X is *finite* but typically very large, so that explicit enumeration is not feasible. Many techniques have been proposed for discrete problems in the deterministic setting; a particular class is that of *random search* procedures. In this type of algorithms, on every iteration a point is selected at random and compared with the current point based on the value of the objective function. A decision is made as to whether or not the new point should replace the current one, and the algorithm proceeds. Random search methods have proven to be quite useful, especially when the structure of the problem cannot be exploited.

Motivated by the above discussion, we consider the application of the variable-sample framework to random search methods for discrete stochastic optimization problems in this article. Because the purpose of the article is to illustrate the use of some general techniques to prove convergence—rather than proposing a new method—we focus on the simplest kind of random search algorithms, namely the *pure* random search method. In that algorithm, on every

iteration a point is selected *uniformly* on X and independently of the previous choices; the new point replaces the current one if the value of the objective function is improved.

Although the pure random search method is not particularly efficient, we feel it constitutes an important example of application of the ideas developed in the article, for several reasons: first, we believe the framework presented here can be easily extended to other methods that include pure random search as a special case. Such is the case, for example, of the modifications proposed by Dorea and Gonçalves [1993], where the choice of the new point is made conditionally on previous selections. Second, some proofs developed for more sophisticated algorithms rely on convergence of pure random search to guarantee convergence of the algorithm. Such is the case of published proofs for genetic algorithms (see, e.g., Rudolph [1996]) and for the so-called *ant colony optimization* method [Gutjahr 2002], to name a few examples. Thus, our results can be readily adapted to show convergence of variable-sample versions of those methods. Third, pure random search allows for an additional degree of flexibility by letting the underlying distributions *depend on the decision variables* x ; in that case, we have $P_k = P_{x_k}$, where x_k is the point obtained in the k th iteration. Finally, in the context of deterministic optimization, *stopping rules* have been developed for pure random search algorithm (see e.g. Dorea [1990] and Hart [1999]). Those stopping rules can in principle be applied to the variable-sample setting—which is possible since the estimates obtained at different iterations are independent—though we do not pursue that in this article. In Homem-de-Mello [2001], we apply the methodology to a more complex algorithm, the *simulated annealing* method.

An important aspect of using a variable-sample scheme is its implementation. The results in Section 3 show that a necessary condition for convergence is that the sample sizes increase at a certain rate. A naive approach of increasing the sample size too often, however, will be too slow in practice to have any usefulness. We address this issue by proposing an *adaptive* version of a variable-sample scheme. The idea is to allow the algorithm to decide automatically what is a “good” sample size. This is accomplished by means of *t-tests*. Once the “good” sample size is achieved, it is increased only occasionally in order to fulfill the requirements for theoretical convergence.

The remainder of this article is organized as follows: we start by briefly reviewing the literature on discrete stochastic optimization. In Section 3, we formally introduce *variable-sample* methods. We provide general conditions under which those methods produce consistent estimates as well as pathwise bounds on the estimation error $|\hat{g}_k(x) - g(x)|$. The analysis in these sections is general, in that it does not depend on the particular algorithm being used for optimization. Then, in Section 4, we present a variable-sample modification of the pure random search (PRS) algorithm, and show that the modification preserves the convergence properties of PRS for deterministic problems, as long as the sample sizes N_k grow at a specific rate. Implementation aspects are discussed in Section 5, where the ideas are illustrated by some numerical results. Finally, in Section 6, we present some concluding remarks.

2. A BRIEF LITERATURE REVIEW

Several methods have proposed in the literature to handle problems of the form (1) when the feasible set X is finite. Here, we mention general *random search* procedures such as the ones discussed in Alrefaei and Andradóttir [2001], Andradóttir [1995, 1996], Gong et al. [1992] and Yan and Mukai [1992]. Adapted versions of the *simulated annealing* method have been studied by Alrefaei and Andradóttir [1999], Fox and Heine [1995], Gelfand and Mitter [1989], and Gutjahr and Pflug [1996]. Another approach is the *ordinal optimization*, proposed by Ho et al. [1992], where the order of the function values are estimated, rather than the function values themselves. Yakowitz et al. [2000] discuss a method where quasi-Monte Carlo techniques are used to select *low-dispersion* points in the feasibility set. An adaptation of the classical *branch-and-bound* method to the context of stochastic optimization is studied by Norkin et al. [1998]. Boesel and Nelson [1998] and Allen et al. [2002] present an alternative procedure based on the combination of *genetic algorithms* with *ranking and selection* techniques. In the particular case where the function $G(\cdot, \omega)$ is the optimal value of a linear programming problem and the set X is polyhedral, *stochastic integer programming* techniques can be applied; see, for instance, the bibliography on stochastic programming compiled by Van der Vlerk [2003].

The basic fixed-sample Monte Carlo approach described in Section 1 has also been applied to discrete stochastic optimization problems. As seen earlier, in that case the expected value function is replaced by its corresponding sample average approximation, and the resulting deterministic problem is solved by some discrete optimization method. Morton and Wood [1998] use this approach to derive upper and lower bounds to the optimal value, and show that the gap decreases with the sample size. Kleywegt et al. [2001] show some theoretical properties of the method. Besides showing convergence of optimal values, they resort to large deviations techniques to show that the solution of the approximating problem converges *exponentially fast* on the sample size N to a solution of the original problem. On the implementation side, they propose solving a sample average approximation of the problem a few times and then using ranking and selection procedures as a second step. We refer to Kleywegt et al. [2001] for details.

3. VARIABLE-SAMPLE METHODS

In this section, we establish a framework to analyze simulation-based methods that utilize different samples along the iterations. The framework presented is general in that we do not assume any particular structure for the problem, which could be either discrete or continuous.

We start with a few definitions. Let N_1, N_2, \dots be a sequence of integer numbers representing the size of the sample used at each iteration. We call $\{N_k\}$ the *schedule* of sample sizes associated with the algorithm under scrutiny. *We assume that the sample used at any given iteration is independent and identically distributed, and that this sample is independent of previous samples.* Notice that the independent and identically distributed assumption

regards only the sample at a given iteration—that is, samples at different iterations can be drawn from different distributions. With that setting, let Ω^{N_k} denote the N_k -fold Cartesian product of the sample space Ω , and let P_k be a probability measure on Ω^{N_k} . Also, let $\tilde{\Omega} = \Omega^{N_1} \times \Omega^{N_2} \times \dots$, and let \tilde{P} denote the corresponding probability distribution on $\tilde{\Omega}$ generated by the P_k s. Notice that such construction is valid because of the assumption of independence between samples at different iterations. In particular, it implies that $\tilde{P}(A_k) = P_k(A_k)$, where A_k is any event in the σ -algebra corresponding to iteration k .

Notice that a point $\omega = (\omega_1^1, \dots, \omega_{N_1}^1, \omega_1^2, \dots, \omega_{N_2}^2, \dots) \in \tilde{\Omega}$ represents a sample-path followed along the iterations of the algorithm. Define now the following random variables on $(\tilde{\Omega}, \tilde{P})$:

$$G_i^k(x, \omega) := G(x, \omega_i^k), \quad k = 1, 2, \dots, i = 1, \dots, N_k.$$

Now, for each $\omega \in \tilde{\Omega}$, define the approximating functions

$$\hat{g}_k(x) := \frac{G_1^k(x, \omega) + \dots + G_{N_k}^k(x, \omega)}{N_k}, \quad k = 1, 2, \dots \quad (2)$$

(we omit the dependence of $\hat{g}_k(x)$ on ω for brevity). The function $\hat{g}_k(x)$ is the approximation to the original function $g(x)$ that is used in the k th iteration. Notice that, conditionally on x , the estimates $\{\hat{g}_k(x)\}$, $k = 1, 2, \dots$, are all *independent* of each other.

3.1 Consistency of Estimators

One approach to show convergence of a variable-sample simulation-based method is to show that the algorithm converges for \tilde{P} -almost all $\omega \in \tilde{\Omega}$. It is natural to think that, in order to have convergence for almost all sample paths, we must have that the estimators used at each iteration are *consistent*, that is, for all $x \in X$ and \tilde{P} -almost all $\omega \in \tilde{\Omega}$, we must have that

$$\lim_{k \rightarrow \infty} \hat{g}_k(x) = g(x). \quad (3)$$

It is interesting to notice that, although condition (3) may seem a direct consequence of the law of large numbers, this is not the case. Two factors contribute to that: first, we do not impose that $\mathbb{E} \hat{g}_k(x) = g(x)$, that is, the estimator $\hat{g}_k(x)$ is allowed to be *biased*. Second, even when this is not the case—for example, when all measures P_k are identical—it could happen that, in principle, $\hat{g}_k(x, \omega)$ does not get close to $g(x)$ with a sample of size N_k . To illustrate the latter point, consider the function $G(x, z) = z$, where z is 0 or 1 with probability 1/2 each, and a sequence ω formed by 2ℓ ones followed by $2\ell + 1$ zeros, $\ell = 0, 1, 2, \dots$. That is, $\omega = (0, 1, 1, 0, 0, 0, 1, 1, 1, 1, \dots)$. Suppose that $N_k = k$, $k = 1, 2, \dots$. Then, after a little algebra, we see that

$$\frac{\ell(n)}{2\ell(n) + 1} \leq \frac{\sum_{j=1}^n \omega_j}{n} \leq \frac{\ell(n) + 1}{2\ell(n) + 1}, \quad (4)$$

where $\ell(n)$ is the smallest nonnegative integer ℓ such that $n \leq (\ell + 1)(2\ell + 1)$. The above inequalities imply that

$$\lim_{n \rightarrow \infty} \frac{\sum_{j=1}^n \omega_j}{n} = \frac{1}{2} = \mathbb{E}\omega,$$

so this sample path satisfies the law of large numbers for the overall sequence. However, we have that $\omega_i^k = u_k$, $i = 1, \dots, N_k$, where $u_k = 0$ if k is odd and $u_k = 1$ otherwise. Hence, we have that $(\omega_1^k + \dots + \omega_{N_k}^k)/N_k = u_k$ and thus the limit in (3) does not exist.

Our initial task is therefore to show that pathological cases such as the one described above happen only on a set of \tilde{P} -probability zero. Moreover, we must impose conditions on the bias $\mathbb{E}\hat{g}_k(x) - g(x)$. Then, we will be able conclude that (3) holds. We start with the following assumptions:

Assumption A1. For each $x \in X$, there exists a positive constant $M = M(x) > 0$ such that $\sup_{k,i} G_i^k(x) \leq M$ with probability one.

Assumption A2. For each $x \in X$, we have that $\lim_{k \rightarrow \infty} \mathbb{E}\hat{g}_k(x) = g(x)$.

A few words about the above assumptions. Assumption A1 says that all random variables are *uniformly bounded* with probability one. In a simulation model, this is the case for example when all input distributions (e.g., service times) have supports which are contained in a finite region that does not depend on k . As we shall see below, this assumption can be relaxed, at a certain expense. Assumption A2 says that the estimators $\hat{g}_k(x)$ are *asymptotically unbiased*, and obviously holds in case all $\hat{g}_k(x)$ are unbiased estimators of $g(x)$. We must also notice that the term “constant” in assumption A1 refers to ω rather than x , that is, constant means “nonrandom quantity”. This terminology is used throughout the article.

We consider now the following alternative assumptions A1' and A2'. Assumption A1' is clearly weaker than assumption A1—rather than requiring conditions on the support of the distributions, it only imposes a condition on the second moments. Most distributions used in simulation have finite variances, so the only practical requirement of assumption A1' is that there be a common bound for all k . Since the measures P_k are controlled by the user, this assumption might be easily enforceable. Assumption A2', on the other hand, deals with the special case when all probability measures P_k are identical.

Assumption A1'. For each $x \in X$, there exists a positive constant $M = M(x) > 0$ such that $\sup_k \text{Var}[G_1^k(x)] \leq M$.

Assumption A2'. All probability measures P_k are identical and the estimator $\hat{g}_k(x)$ is unbiased.

Before proceeding with the results, let us recall some basic facts from *large deviations* theory. Let Y_1, Y_2, \dots be a sequence of independent and identically distributed random variables with finite expectation μ , and for all $N \geq 1$ define

$\bar{Y}_N = \sum_{i=1}^N Y_i/N$; then, the weak law of large numbers says that, for any $\delta > 0$,

$$\lim_{N \rightarrow \infty} P(|\bar{Y}_N - \mu| \geq \delta) = 0, \quad (5)$$

and the large deviation theory asserts that the above probability converges to zero exponentially fast. Indeed, for any $N \geq 1$, Chernoff's bound yields

$$P(\bar{Y}_N - \mu \geq \delta) \leq \exp(-NI(\mu + \delta)). \quad (6)$$

Here, $I(\cdot)$ is the so-called *rate function* corresponding to the distribution of Y_1 , which is defined by $I(z) := \sup_{t \in \mathbb{R}} \{tz - \log M(t)\}$, where $M(t) := \mathbb{E}[e^{tY_1}]$ is the *moment generating function* of Y_1 (which is assumed to be finite in a neighborhood of zero). It is possible to show that $I(\cdot)$ is nonnegative, strictly convex and attains its minimum at μ (with $I(\mu) = 0$), so that the exponent on the right-hand side of (6) is strictly negative. Next, by applying inequality (6) to the process $\{-Y_i\}$, we have that

$$P(\bar{Y}_N - \mu \leq -\delta) \leq \exp(-NI(\mu - \delta))$$

and thus

$$P(|\bar{Y}_N - \mu| \geq \delta) \leq 2 \exp(-N\gamma(\delta)), \quad (7)$$

where $\gamma(\delta) := \min(I(\mu + \delta), I(\mu - \delta))$. This implies (5). It is possible to show that the exponential bound in the above inequality is asymptotically sharp, in the sense that

$$P(|\bar{Y}_N - \mu| > \delta) = \exp(-N\gamma(\delta) + o(N)). \quad (8)$$

The literature on large deviations theory is quite ample; we refer the reader to the books by Dembo and Zeitouni [1998] and Shwartz and Weiss [1995], for example, for comprehensive discussions. For our purposes, the results in (7) and (8) will suffice.

The above results, while very important from a *qualitative* point of view, are difficult to use directly since typically it is very hard to compute the rate function $I(\cdot)$. Thus, estimates for $I(\cdot)$ are needed. A useful one can be derived when the corresponding random variables are *uniformly bounded*, that is, $|Y_i| \leq M$ with probability one. In that case, we have that

$$I(z) \geq \frac{(z - \mu)^2}{2M^2} \quad \text{for all } z \in \mathbb{R} \quad (9)$$

(see, e.g., Shapiro and Homem-de-Mello [2000] for a proof). A similar result can be derived under the weaker assumption that the corresponding random variables have finite variance σ^2 . Then, there exists a neighborhood \mathcal{N} of μ such that

$$I(z) \geq \frac{(z - \mu)^2}{3\sigma^2} \quad \text{for all } z \in \mathcal{N}. \quad (10)$$

This follows directly from the Taylor expansion of the function I ; see Kleywegt et al. [2001] for a proof.

Another estimate of the deviation probability $P(|\bar{Y} - \mu| > \delta)$ can be obtained by a variant of the Central Limit Theorem. In that case, δ goes to zero with n ,

so that the deviations are never very large. In Chung [1974, Thm. 7.1.3], the following result is proved: suppose the independent and identically distributed sequence $\{Y_i\}$ has finite variance σ^2 and finite third central moment $\gamma^3 := \mathbb{E}(|Y_1 - \mu|^3)$. Let a_n be a sequence of real numbers increasing to infinity, and subject to the following growth condition:

$$\lim_{n \rightarrow \infty} \log \frac{n\gamma^3}{(n\sigma^2)^{3/2}} + \frac{a_n^2}{2}(1 + \varepsilon) = -\infty \quad \text{for some } \varepsilon > 0. \quad (11)$$

Then, for this ε , there exists N such that, for all $n \geq N$, we have

$$\exp^{-a_n^2(1+\varepsilon)/2} \leq P\left(\bar{Y} - \mu \geq \frac{a_n\sigma}{\sqrt{n}}\right) \leq \exp(-a_n^2(1-\varepsilon)/2). \quad (12)$$

We can now state the results:

PROPOSITION 3.1. *Suppose that Assumptions A1 and A2 hold. Suppose also that the schedule $\{N_k\}$ satisfies the following property:*

$$\sum_{k=1}^{\infty} \alpha^{N_k} < \infty \quad \text{for all } \alpha \in (0, 1). \quad (13)$$

Then, $\hat{g}_k(x) \rightarrow g(x)$ for \tilde{P} -almost all $\omega \in \tilde{\Omega}$.

PROOF. Fix $x \in X$, let $k \geq 1$ and $\delta > 0$. To simplify the notation, let $g_k := \mathbb{E}\hat{g}_k(x)$. Then, by inequality (7) above, we have that

$$P_k(|\hat{g}_k(x) - g_k| \geq \delta) \leq 2 \exp(-N_k \gamma_k(\delta)), \quad (14)$$

where $\gamma_k(\delta) := \min(I_k(g_k + \delta), I_k(g_k - \delta))$, and I_k is the rate function of $G_i^k(x)$. By Assumption A1, we have that $G_i^k(x) \leq M$ with probability one for all k and i and thus, from (9), we have

$$\gamma_k(\delta) \geq \frac{\delta^2}{2M^2}$$

and hence

$$\tilde{P}(|\hat{g}_k(x) - g_k| \geq \delta) = P_k(|\hat{g}_k(x) - g_k| \geq \delta) \leq 2 \exp(-N_k \delta^2 / (2M^2)). \quad (15)$$

It follows that

$$\sum_{k=1}^{\infty} \tilde{P}(|\hat{g}_k(x) - g_k| \geq \delta) \leq \sum_{k=1}^{\infty} 2 \exp(-\delta^2 / (2M^2))^{N_k}. \quad (16)$$

Notice that, if condition (13) holds, then the expression on the right-hand side of the above inequality is finite for all $\delta > 0$. By the Borel–Cantelli lemma (see, e.g., Chung [1974, Thm. 4.2.1]), we then have that

$$\tilde{P}(|\hat{g}_k(x) - g_k| \geq \delta \text{ infinitely often (i.o.)}) = 0 \quad \forall \delta > 0.$$

Finally, Assumption A2 implies that, given $\delta > 0$, $|g_k - g(x)| < \delta/2$ for k large enough. It follows that

$$\tilde{P}(|\hat{g}_k(x) - g(x)| \geq \delta/2 \text{ infinitely often}) = 0 \quad \forall \delta > 0 \quad (17)$$

and therefore $\hat{g}_k(x) \rightarrow g(x)$ for \tilde{P} -almost all $\omega \in \tilde{\Omega}$. \square

The above result can be strengthened in case the measures P_k are identical. As the proposition below shows, in that case, condition (13) is also *necessary* for convergence of $\hat{g}_k(x)$ to $g(x)$.

PROPOSITION 3.2. *Suppose that Assumptions A1' and A2' hold. Then, a sufficient condition to have $\hat{g}_k(x) \rightarrow g(x)$ for \tilde{P} -almost all $\omega \in \tilde{\Omega}$ is that the schedule $\{N_k\}$ satisfies condition (13). If, in addition, $\text{Var}[G_1^1(x)] > 0$, then condition (13) is also necessary.*

PROOF. Fix $x \in X$, let $k \geq 1$ and $\delta > 0$. Then, (7) and (8), together with Assumption A2', imply that there exists a sequence $\{c_k\}$ such that $c_k \rightarrow 0$ and

$$P_k(|\hat{g}_k(x) - g(x)| \geq \delta) \leq 2 \exp(-N_k \gamma(\delta)) \quad (18)$$

$$P_k(|\hat{g}_k(x) - g(x)| > \delta) \geq \exp(-N_k(\gamma(\delta) + c_k)), \quad (19)$$

where $\gamma(\delta) := \min(I(g(x) + \delta), I(g(x) - \delta))$, and I is the rate function of $G_i^k(x)$. Now, from (10) and Assumption A1', we have that there exists a neighborhood \mathcal{N} of zero such that

$$\gamma(\delta) \geq \frac{\delta^2}{3M^2} \quad \text{for all } \delta \in \mathcal{N}$$

and thus, for δ small enough and all k ,

$$\tilde{P}(|\hat{g}_k(x) - g(x)| \geq \delta) = P_k(|\hat{g}_k(x) - g(x)| \geq \delta) \leq 2 \exp(-N_k \delta^2 / (3M^2)). \quad (20)$$

On the other hand, since the random variables $G_i^k(x)$ are assumed to have positive variance, it follows that the rate function I is finite in a neighborhood of $g(x)$, that is, $\gamma(\delta) < \infty$ for δ small enough. In fact, $\gamma(\delta) \rightarrow 0$ as $\delta \rightarrow 0$.

Next, let $\varepsilon > 0$ be arbitrary. Since $\gamma(\delta) \rightarrow 0$ and the sequence $\{c_k\}$ goes to zero, we have from (19) that, for δ sufficiently small and k sufficiently large,

$$\tilde{P}(|\hat{g}_k(x) - g(x)| > \delta) = P_k(|\hat{g}_k(x) - g(x)| > \delta) \geq \exp(-N_k \varepsilon) = [\exp(-\varepsilon)]^{N_k}. \quad (21)$$

Together, (20) and (21) imply that $\sum_{k=1}^{\infty} \tilde{P}(|\hat{g}_k(x) - g(x)| > \delta)$ is finite for all δ sufficiently small *if and only if* condition (13) holds. By applying the full statement of the Borel–Cantelli lemma (see, e.g., Chung [1974, Thms. 4.2.1 and 4.2.4]), we conclude that, under assumption A1', we have

$$(13) \text{ holds} \Rightarrow \exists \delta^0 > 0 : \tilde{P}(|\hat{g}_k(x) - g(x)| > \delta \text{ infinitely often}) = 0 \\ \forall \delta \in (0, \delta^0)$$

$$(13) \text{ does not hold} \Rightarrow \exists \delta^1 > 0 : \tilde{P}(|\hat{g}_k(x) - g(x)| > \delta \text{ infinitely often}) = 1 \\ \forall \delta \in (0, \delta^1).$$

It follows from the above implications that a necessary and sufficient condition to have $\hat{g}_k(x) \rightarrow g(x)$ for \tilde{P} -almost all $\omega \in \tilde{\Omega}$ is that condition (13) holds. \square

Some remarks about the above results are now in order. First, notice that condition (13) imposes a mild constraint on the schedule of sample sizes. Indeed, it is evident that such condition holds if $N_k \geq ck$, where c is any positive constant. Even a sublinear growth such as $N_k = \sqrt{k}$ satisfies (13). Notice however that not all increasing schedules satisfy (13): for example, with $N_k = c \log k$ we have, for any $\alpha > 0$,

$$\sum_{k=1}^{\infty} \alpha^{c \log k} = \sum_{k=1}^{\infty} e^{c \log k \log \alpha} = \sum_{k=1}^{\infty} k^{\log \alpha^c}$$

which converges if and only if $\alpha < \exp(-1/c)$. Therefore, condition (13) does not hold in that case. A somewhat surprising consequence of Proposition 3.2 is that, when the measures P_k are identical and $N_k = c \log k$, $\hat{g}_k(x)$ does *not* converge to $g(x)$ for \tilde{P} -almost all $\omega \in \tilde{\Omega}$. The proof of Proposition 3.2 shows why this happens—there exists some $\delta > 0$ such that the deviation $|\hat{g}_k(x) - g(x)|$ is larger than δ infinitely often with probability one. In that case, we can only conclude that there exists a *subsequence* of $\{\hat{g}_k(x)\}$ converging to $g(x)$. Nevertheless, as we shall see later, in some cases it suffices to ensure that $\tilde{P}(|\hat{g}_k(x) - g(x)| > \delta \text{ infinitely often}) = 0$ for some given $\delta > 0$ (rather than all $\delta > 0$) in order to have convergence of the algorithm; in those situations, the proof of Proposition 3.2 can be modified to show that even a schedule such as $N_k = c \log k$ is acceptable, provided c is sufficiently large.

Another remark concerns the necessity of condition (13) in Proposition 3.2. Observe the importance of the assumption of positive variance in that proposition, since otherwise $G_i^k(x) \equiv g(x)$ and thus (13) would not be necessary. Moreover, condition (13) is not necessary under the conditions of Proposition 3.1, particularly Assumption A2. Indeed, suppose, for example, that $G_i^k(x)$ has distribution with mean $g(x)$ and variance $\sigma_k^2 = 1/k$; clearly, $G_i^k(x)$ approaches the constant $g(x)$ with probability one as k grows and therefore any nondecreasing schedule $\{N_k\}$ (e.g., $N_k = 1$ for all k) guarantees that $\hat{g}_k(x)$ converges to $g(x)$ with probability one.

We conclude this section by proposing yet another alternative to Propositions 3.1 and 3.2. It requires a stronger assumption on the schedule $\{N_k\}$ but it requires weaker assumptions on the underlying random variables. A related result was derived by Cooper et al. [2003] in a different context.

PROPOSITION 3.3. *Suppose that Assumptions A1' and A2 hold. Suppose also that the schedule $\{N_k\}$ satisfies the following property:*

$$\sum_{k=1}^{\infty} \frac{1}{N_k} < \infty. \quad (22)$$

Then, $\hat{g}_k(x) \rightarrow g(x)$ for \tilde{P} -almost all $\omega \in \tilde{\Omega}$.

PROOF. Fix $x \in X$, let $k \geq 1$ and $\delta > 0$. We apply Chebyshev's inequality to obtain

$$P_k(|\hat{g}_k(x) - \mathbb{E}\hat{g}_k(x)| \geq \delta) \leq \frac{\text{Var}[\hat{g}_k(x)]}{\delta^2} = \frac{\text{Var}[G_1^k(x)]/\delta^2}{N_k} \leq \frac{M/\delta^2}{N_k}. \quad (23)$$

Thus, $\sum_{k=1}^{\infty} \tilde{P}(|\hat{g}_k(x) - \mathbb{E}\hat{g}_k(x)| \geq \delta)$ converges if and only if (22) holds. The remainder of the proof is identical to that of Proposition 3.1 and therefore is omitted. \square

3.2 Sample-Path Bounds

The results in the previous section ensure consistency of the estimators used at each iteration. Intuitively, this guarantees that, if k (and therefore N_k) is sufficiently large, then $\hat{g}_k(x)$ is close to $g(x)$ and so, in principle, when k is large a variable-sample method should not behave too differently from a hypothetical method that could solve the original problem (1). Notice however that the basic argument for using a variable-sample method is to update the sample sizes *as the algorithm progresses*; therefore, we may need stronger results than just consistency. In particular, it is desirable to derive *bounds* on the deviation $|\hat{g}_k(x) - g(x)|$.

The theorem below provides such bound. Notice that the result is not stated in terms of distributions, but rather for each sample path ω . In a sense, it corresponds to the law of the iterated logarithm in the standard independent and identically distributed case. Observe that conditions imposed on the schedule $\{N_k\}$ are stronger than before. We shall also impose the following assumption:

Assumption A3. For each $x \in X$, there exist a positive constant $M_1 = M_1(x)$ such that $\sup_k \mathbb{E}[|G_1^k(x) - \mathbb{E}\hat{g}_k(x)|^3] / (\text{Var}[G_1^k(x)])^{3/2} \leq M_1$.

Assumption A3 holds, for example, if the random variables $G_1^k(x)$ have uniformly bounded third moment (for all k) and their variances are uniformly bounded away from zero.

THEOREM 3.4. *Suppose that Assumption A3 holds. Suppose also that the schedule $\{N_k\}$ satisfies the following property:*

$$N_k \geq c_1 k^\rho \quad \text{for some } c_1 > 0 \text{ and some } \rho > 2. \quad (24)$$

Then, for \tilde{P} -almost all $\omega \in \tilde{\Omega}$, there exists $K = K(\omega) > 0$ such that

$$|\hat{g}_k(x) - g(x)| \leq \sigma_k(x) \sqrt{\frac{\log N_k}{N_k}} + |\mathbb{E}\hat{g}_k(x) - g(x)| \quad (25)$$

for all $k > K$, where $\sigma_k^2(x) := \text{Var}[G_1^k(x)]$.

If $N_k \leq c_2 k^\rho$ for some $c_2 > 0$ and some $\rho \leq 2$, then for \tilde{P} -almost all $\omega \in \tilde{\Omega}$ we have that

$$|\hat{g}_k(x) - g(x)| \geq \sigma_k(x) \sqrt{\frac{\log N_k}{N_k}} - |\mathbb{E}\hat{g}_k(x) - g(x)| \quad \text{infinitely often.} \quad (26)$$

PROOF. Fix $x \in X$ and let $k \geq 1$. Our goal is to apply inequalities (12) to estimate the deviation probabilities of $\hat{g}_k(x) - \mathbb{E}\hat{g}_k(x)$. Let $\sigma_k^2 := \text{Var}[G_1^k(x)]$ and $\gamma_k^3 := \mathbb{E}[|G_1^k(x) - \mathbb{E}\hat{g}_k(x)|^3]$. Let $\delta > 0$ be arbitrarily chosen, and define $a_k := \sqrt{1/(1+\delta) \log N_k}$. Clearly, $a_k \rightarrow \infty$. Moreover, the sequence $\{a_k\}$ satisfies

the growth condition (11), since

$$\begin{aligned} \limsup_{k \rightarrow \infty} \log \frac{N_k \gamma_k^3}{(N_k \sigma_k^2)^{3/2}} + \frac{\alpha_k^2}{2}(1 + \varepsilon) &= \limsup_{k \rightarrow \infty} \log \frac{\gamma_k^3}{\sigma_k^3} + \log N_k^{-1/2} + \frac{1 + \varepsilon}{2(1 + \delta)} \log N_k \\ &\leq \log M_1 + \limsup_{k \rightarrow \infty} \left(\frac{1 + \varepsilon}{2(1 + \delta)} - \frac{1}{2} \right) \log N_k \\ &= -\infty \quad \text{for } \varepsilon < \delta. \end{aligned}$$

Therefore, the conditions for (12) are satisfied when $\varepsilon < \delta$. Fix now an $\varepsilon < \delta/(1 + 2\delta)$. Thus, from (12), we have

$$\begin{aligned} 2 \exp \left(-\frac{1 + \varepsilon}{2(1 + \delta)} \log N_k \right) &\leq P_k \left(|\hat{g}_k(x) - \mathbb{E} \hat{g}_k(x)| \geq \frac{\sigma_k}{\sqrt{1 + \delta}} \sqrt{\frac{\log N_k}{N_k}} \right) \\ &\leq 2 \exp \left(-\frac{1 - \varepsilon}{2(1 + \delta)} \log N_k \right), \end{aligned}$$

that is,

$$2N_k^{-\frac{1+\varepsilon}{2(1+\delta)}} \leq \tilde{P} \left(|\hat{g}_k(x) - \mathbb{E} \hat{g}_k(x)| \geq \frac{\sigma_k}{\sqrt{1 + \delta}} \sqrt{\frac{\log N_k}{N_k}} \right) \leq 2N_k^{-\frac{1-\varepsilon}{2(1+\delta)}} \quad (27)$$

for k large enough, say $k \geq K$.

Suppose now that the schedule $\{N_k\}$ satisfies (24). By summing over k in (27), we obtain

$$\sum_{k=K}^{\infty} \tilde{P} \left(|\hat{g}_k(x) - \mathbb{E} \hat{g}_k(x)| \geq \frac{\sigma_k}{\sqrt{1 + \delta}} \sqrt{\frac{\log N_k}{N_k}} \right) \leq \sum_{k=K}^{\infty} C_1 k^{-\frac{(1-\varepsilon)\rho}{2(1+\delta)}} \quad (28)$$

where $C_1 = 2/c_1$. Let $\delta \leq (\rho - 2)/4$. Since $\rho > 2$, δ is positive. We have $\rho \geq 2(1 + 2\delta)$ and thus, in (28), we obtain

$$\sum_{k=K}^{\infty} \tilde{P} \left(|\hat{g}_k(x) - \mathbb{E} \hat{g}_k(x)| \geq \frac{\sigma_k}{\sqrt{1 + \delta}} \sqrt{\frac{\log N_k}{N_k}} \right) \leq \sum_{k=K}^{\infty} C_1 k^{-\frac{(1-\varepsilon)(1+2\delta)}{1+\delta}}. \quad (29)$$

The expression on the right-hand side of the above inequality is finite if and only if the exponent of k is less than -1 , that is, if and only if $\varepsilon < \delta/(1 + 2\delta)$. Since ε was fixed above to satisfy such condition, it follows that the expression on the left-hand side of (29) is finite when $\rho > 2$.

Suppose now that the schedule $\{N_k\}$ satisfies $N_k \leq c_2 k^\rho$ for $\rho \leq 2$. Next, recall that ε was fixed above to be smaller than δ . It follows from (27) that

$$\sum_{k=K}^{\infty} C_2 k^{-\rho/2} \leq \sum_{k=K}^{\infty} \tilde{P} \left(|\hat{g}_k(x) - \mathbb{E} \hat{g}_k(x)| \geq \frac{\sigma_k}{\sqrt{1 + \delta}} \sqrt{\frac{\log N_k}{N_k}} \right)$$

(where $C_2 = 2/c_2$) and so the series on the right-hand side of the above inequality diverges if $\rho \leq 2$.

The above conclusions, together with the Borel–Cantelli lemma, imply that

$$\begin{aligned} \tilde{P}\left(|\hat{g}_k(x) - \mathbb{E}\hat{g}_k(x)| \geq \frac{\sigma_k}{\sqrt{1+\delta}} \sqrt{\frac{\log N_k}{N_k}} \text{ infinitely often}\right) &= 0 \quad \text{if } \rho > 2 \\ \tilde{P}\left(|\hat{g}_k(x) - \mathbb{E}\hat{g}_k(x)| \geq \frac{\sigma_k}{\sqrt{1+\delta}} \sqrt{\frac{\log N_k}{N_k}} \text{ infinitely often}\right) &= 1 \quad \text{if } \rho \leq 2. \end{aligned}$$

The first equation holds for all $0 < \delta \leq (\rho - 2)/4$, whereas the second holds for any $\delta > 0$. Therefore, the assertion of the theorem follows. \square

Remark. The above theorem shows that, in a sense, $\rho > 2$ is the weakest requirement on ρ that yields a bound of order $\sqrt{\log N_k/N_k}$. This means that more restrictive bounds such as $O(N_k^{-1/2})$ require at least $\rho > 2$. Similarly, if $\rho \leq 2$, then we cannot expect a convergence rate faster than $\sqrt{\log N_k/N_k}$.

Notice that the right-hand side in (25) has a component due to the bias $\mathbb{E}[G_1^k(x)] - g(x)$. If this bias dominates the term $\sqrt{\log N_k/N_k}$, then, of course, the error $|\hat{g}_k(x) - g(x)|$ will be the order of the bias. Under Assumption A2'' below, Theorem 3.4 yields a direct consequence.

Assumption A2''. For each $x \in X$, there exists a positive constant $D = D(x)$ such that

$$|\mathbb{E}[G_i^k(x)] - g(x)| \leq D \sqrt{\frac{\log N_k}{N_k}} \quad \text{for all } k \geq 1. \quad (30)$$

COROLLARY 3.5. *Suppose that Assumptions A1', A2'' and A3 hold. Suppose also that the schedule $\{N_k\}$ satisfies the following property:*

$$N_k \geq ck^{2+\delta} \quad \text{for some } c > 0 \text{ and some } \delta > 0. \quad (31)$$

Then, there exists a constant $C = C(x) > 0$ such that, for \tilde{P} -almost all $\omega \in \tilde{\Omega}$, there exists $K = K(\omega) > 0$ such that

$$|\hat{g}_k(x) - g(x)| \leq C \sqrt{\frac{\log N_k}{N_k}} \quad (32)$$

for all $k > K$.

The above results provide the desired bound on the deviation $|\hat{g}_k(x) - g(x)|$. Note that no assumptions were made on the distribution of $G_i^k(x)$, other than some boundedness assumptions on the first three moments. This underscores the generality of the results. Another remark is that, clearly, (32) implies that $\hat{g}_k(x) \rightarrow g(x)$ with probability one, which was the conclusion of Propositions 3.1–3.3. Those propositions, however, use weaker assumptions on the schedule $\{N_k\}$.

3.3 Cumulative Samples

The results in the previous sections focus on a certain sampling structure—namely, it is assumed that samples used at different iterations are independent. It is natural to think of an alternative scheme, where at each iteration a new sample is *appended* to the previous one. In other words, using the notation defined earlier, if the sample used at the first iteration is $\omega_1^1, \dots, \omega_{N_1}^1$, then the sample used at the second iteration is $\omega_1^1, \dots, \omega_{N_1}^1, \omega_1^2, \dots, \omega_{N_2}^2$ and so on. Thus, the estimator $\bar{g}_k(x)$ used at iteration k is defined as

$$\bar{g}_k(x) := \frac{G_1^1(x, \omega) + \dots + G_{N_k}^k(x, \omega)}{N_1 + \dots + N_k}, \quad k = 1, 2, \dots \quad (33)$$

for each $\omega \in \tilde{\Omega}$.

It is clear that, under such scheme, consistency follows immediately from the strong law of large numbers when the measures P_k are identical. If the measures P_k are not identical, then we need extra conditions on the bias $|\mathbb{E}[G_i^k(x)] - g(x)|$, as the proposition below shows:

PROPOSITION 3.6. *Suppose that Assumptions A1' and A2'' hold. Then, for all $x \in X$,*

$$\lim_{k \rightarrow \infty} \bar{g}_k(x) = g(x) \quad (34)$$

for \tilde{P} -almost all $\omega \in \tilde{\Omega}$, provided that $\lim_{k \rightarrow \infty} N_1 + \dots + N_k = \infty$.

PROOF. Fix $x \in X$, and let $Y_i^k = G_i^k(x) - \mathbb{E}[G_i^k(x)]$, $k = 1, 2, \dots, i = 1, \dots, N_k$. Then, $\mathbb{E}[Y_i^k] = 0$ and, by assumption A1', $\mathbb{E}|Y_i^k|^2 \leq M$. It follows from a classical result in Probability that $\lim_{k \rightarrow \infty} (Y_1^1 + \dots + Y_{N_k}^k)/(N_1 + \dots + N_k) = 0$ with probability one (see, e.g., Chung [1974, p. 125]) and hence, by Assumption A2'', we have that

$$\begin{aligned} \lim_{k \rightarrow \infty} \frac{G_1^1(x) + \dots + G_{N_k}^k(x)}{N_1 + \dots + N_k} &= \lim_{k \rightarrow \infty} \frac{\mathbb{E}[G_1^1(x)] + \dots + \mathbb{E}[G_{N_k}^k(x)]}{N_1 + \dots + N_k} \\ &\leq \lim_{k \rightarrow \infty} \frac{\sum_{i=1}^k N_i \left(g(x) + D \sqrt{\frac{\log N_i}{N_i}} \right)}{\sum_{i=1}^k N_i} \\ &\leq g(x) + \lim_{k \rightarrow \infty} \frac{D \sum_{i=1}^k \sqrt{N_i \log N_i}}{\sum_{i=1}^k N_i} \\ &= g(x) \end{aligned}$$

for \tilde{P} -almost all $\omega \in \tilde{\Omega}$. Similarly, one can show that $\lim_{k \rightarrow \infty} (G_1^1(x) + \dots + G_{N_k}^k(x))/(N_1 + \dots + N_k) \geq g(x)$, so the assertion of the proposition follows. \square

A bound similar to the one given by Theorem 3.4 can also be derived in this case. It is a direct consequence of the *law of the iterated logarithm*—which can be seen from the “log log” term on the bound.

PROPOSITION 3.7. *Suppose that any of the conditions below is satisfied:*

- (i) *Assumption A2' holds;*
- (ii) *Assumption A1 holds;*
- (iii) *There exist constants $A > 0$ and $0 < \varepsilon < 1$ such that*

$$\frac{\Gamma_k}{\Sigma_k^3} \leq \frac{A}{(\log \Sigma_k)^{1+\varepsilon}},$$

where $\Sigma_k^2 := \sum_{i=1}^k N_i \sigma_i^2$, $\sigma_i^2 := \text{Var}[G_1^i(x)]$, $\Gamma_k = \sum_{i=1}^k N_i \gamma_i^3$, and also $\gamma_i^3 := \mathbb{E}[|G_1^i(x) - \mathbb{E}[G_1^i(x)]|^3]$.

Suppose also that Assumption A1' holds and that $\mathbb{E}[G_1^k(x)] = g(x)$ for all $x \in X$ and all $k = 1, 2, \dots$. Then, for \tilde{P} -almost all $\omega \in \tilde{\Omega}$, there exist positive constants C and $K = K(\omega)$ such that

$$|\bar{g}_k(x) - g(x)| \leq C \sqrt{\frac{\log \log(N_1 + \dots + N_k)}{N_1 + \dots + N_k}} \quad \forall k \geq K, \quad (35)$$

provided that $\lim_{k \rightarrow \infty} N_1 + \dots + N_k = \infty$.

PROOF. Fix $x \in X$. Then, conditions (i)–(iii) above, together with independence of the variables $G_i^k(x)$, imply that we can use the law of the iterated logarithm for the sequence $G_1^1(x), \dots, G_{N_k}^k(x)$, so that, with probability one,

$$\limsup_{k \rightarrow \infty} \frac{(G_1^1(x) - \mathbb{E}[G_1^1(x)]) + \dots + (G_{N_k}^k(x) - \mathbb{E}[G_{N_k}^k(x)])}{\sqrt{2\Sigma_k^2 \log \log \Sigma_k}} = 1 \quad (36)$$

$$\liminf_{k \rightarrow \infty} \frac{(G_1^1(x) - \mathbb{E}[G_1^1(x)]) + \dots + (G_{N_k}^k(x) - \mathbb{E}[G_{N_k}^k(x)])}{\sqrt{2\Sigma_k^2 \log \log \Sigma_k}} = -1 \quad (37)$$

(see, e.g., Chung [1974], and Rao [1984]). Since $\mathbb{E}[G_1^k(x)] = g(x)$ for all k and $\Sigma_k^2 = \sum_{i=1}^k N_i \sigma_i^2 \leq M \sum_{i=1}^k N_i$, it follows from the above inequalities that, given $\delta > 0$, there exists $K = K(\omega)$ such that

$$|(G_1^1(x) - g(x)) + \dots + (G_{N_k}^k(x) - g(x))| \leq (1 + \delta) \sqrt{2\Sigma_k^2 \log \log \Sigma_k} \quad \forall k \geq K$$

and thus

$$\begin{aligned} \left| \frac{G_1^1(x) + \dots + G_{N_k}^k(x)}{N_1 + \dots + N_k} - g(x) \right| &\leq (1 + \delta) \sqrt{\frac{2M \log \log \sqrt{M(N_1 + \dots + N_k)}}{N_1 + \dots + N_k}} \\ &\leq C \sqrt{\frac{\log \log(N_1 + \dots + N_k)}{N_1 + \dots + N_k}} \end{aligned}$$

for some $C > 0$ and k large enough. \square

The task of showing convergence of a variable-sample simulation-based method is facilitated by the above results. Assuming that a proof of convergence of the underlying deterministic algorithm is given (in the deterministic

setting), the goal becomes to prove that the convergence properties of such algorithm are kept when, at each iteration k , one replaces the original function $g(x)$ by an approximating function $\hat{g}_k(x)$ such that $\hat{g}_k(x) \rightarrow g(x)$ with probability one as k goes to infinity.

The main task is to show that, in some sense, the deterministic error *dominates* the stochastic error resulting from approximating $g(x)$ by $\hat{g}_k(x)$. In that sense, Theorem 3.4 and Proposition 3.7 are crucial, as they provide upper bounds on the error $|\hat{g}_k(x) - g(x)|$. An important aspect of those bounds is that they are *distribution-free*, which allows for applicability of those results in fairly general contexts. Moreover, because those bounds are derived for sample paths, one can analyze the underlying algorithm for each individual sample path, which typically leads to stronger “with probability one” results. Finally, we emphasize that the properties derived in the previous section hold both when new samples are drawn at each iteration as well as when samples are accumulated from one iteration to the next. Therefore, the convergence results will be valid under either sampling scheme.

In the next section we will discuss the application of the ideas developed above to a random search method.

4. RANDOM SEARCH METHODS

Random search methods have become quite popular in global optimization. The basic idea in those methods—roughly speaking, of course—is to search “at random” for points in the feasible set, and then compare the value of the objective function at those points. Many known algorithms fall into this class of methods. The simplest form of a random search method is the *pure* random search. The algorithm can be described very briefly:

Algorithm PRS

```

Select a point  $x_1$  in  $X$  at random;
 $k := 1$ ;
Repeat
  Choose a point  $y$  in  $X$  at random;
  If  $g(y) < g(x_k)$ 
    then  $x_{k+1} := y$ 
    else  $x_{k+1} := x_k$ ;
   $k := k + 1$ ;
Until {stopping criterion is satisfied}.

```

Of course, it is well known that the above algorithm is not particularly efficient, and may take considerable time to escape local minima. Nevertheless, under appropriate conditions on X and g the algorithm is *convergent* with probability one. One such case occurs when X is finite, since at every iteration there is probability at least $1/|X|$ of sampling a minimizer x^* of g —and after that we have $x_k = x^*$ for all k . More generally, when X is compact and g is continuous the algorithm converges to the set of ε -optimal solutions (i.e., the set S_ε such that such that $|g(x) - g(x^*)| < \varepsilon$ for all $x \in S_\varepsilon$) with probability one. One advantage of the simplicity of the method is that it allows for the development of easily implementable *stopping rules* for the algorithm (see Dorea [1990] and Hart [1999]).

Algorithm PRS can be easily modified to incorporate other sampling schemes instead of the uniform sampling used in it. For example, suppose we define a metric on X and let $B(x, \delta)$ denote the set of points that lie within a distance δ from x . We can modify the selection of y in X in the algorithm to use the following sampling scheme:

—Select a point y in X according to a conditional distribution that depends on $B_\delta := B(x_1, \delta) \cup \dots \cup B(x_k, \delta)$.

This conditional distribution can simply assign zero likelihood to the set B_δ [Devroye 1978] or, more generally, it can sample from B_δ with a certain pre-specified probability p [Dorea and Gonçalves 1993]. Such scheme is also convergent, and it has been observed that the expected number of steps until convergence is smaller than for the pure random search algorithm [Dorea and Gonçalves 1993].

Consider now the pure random search method adapted to random functions. Using the variable-sample scheme described in Section 3, Algorithm PRS becomes the following.

Algorithm SPRS

Select a point x_1 in X at random;
 $N_1 :=$ initial sample size;
 $k := 1$;
Repeat
 Choose a point y in X at random;
 Generate a sample $\omega_1^k, \dots, \omega_{N_k}^k$ from a distribution P_k ;
 Compute $\hat{g}_k(x_k), \hat{g}_k(y)$ according to (2) or (33);
 If $\hat{g}_k(y) < \hat{g}_k(x_k)$
 then $x_{k+1} := y$
 else $x_{k+1} := x_k$;
 Update N_k to obtain N_{k+1} ;
 $k := k + 1$;
Until {stopping criterion is satisfied}.

The theorem below shows that Algorithm SPRS does converge to the set of optimal solutions. For simplicity, we shall assume that X is finite, but the result is easily extendable to the continuous case with ε -optimal solutions mentioned earlier. Notice that, besides the probability space $(\tilde{\Omega}, \tilde{P})$ that drive the sampling points ω_i^k , in the above scheme there is another underlying probability space—call it (Ξ, P) —to govern the choice of states in X .

THEOREM 4.1. *Suppose that the feasibility set X is finite. Assume that either Assumptions A1 and A2 hold, or Assumptions A1' and A2' hold. Consider Algorithm SPRS, with the sample size schedule satisfying (13). Then, for \tilde{P} -almost all $\omega \in \tilde{\Omega}$ and P -almost all $\xi \in \Xi$, there exists $k_0 = k_0(\omega, \xi)$ such that*

$$x_k \in S^* \text{ for all } k \geq k_0, \quad (38)$$

where S^* is the set of optimal solutions of (1).

PROOF. As shown in the proofs of Propositions 3.1 and 3.2, if $\{N_k\}$ satisfies (13), then we have that, for each $x \in X$ and each $\delta > 0$,

$$\tilde{P}(|\hat{g}_k(x) - \mathbb{E}\hat{g}_k(x)| \geq \delta \text{ infinitely often}) = 0.$$

This means that, given $\delta > 0$ and $\omega \in \tilde{\Omega}$, there exists $K_x^1 = K_x^1(\delta, \omega)$ such that $|\hat{g}_k(x) - \mathbb{E}\hat{g}_k(x)| < \delta/2$ for all $k \geq K_x^1$. Moreover, by Assumption A2 (or A2') there exists $K_x^2 = K_x^2(\delta)$ such that $|g(x) - \mathbb{E}\hat{g}_k(x)| < \delta/2$ for all $k \geq K_x^2$. By taking $K := \max_{x \in X} \max(K_x^1, K_x^2)$, we have that

$$|g(x) - \hat{g}_k(x)| < \delta \text{ for all } k \geq K \text{ and all } x \in X. \quad (39)$$

Now, let $\delta := (1/2) \min_{x, y \in X} |g(x) - g(y)|$. Let x, y be two arbitrary points in X with $g(x) < g(y)$. When $k \geq K$, we have

$$\begin{aligned} \hat{g}_k(x) - \hat{g}_k(y) &= \hat{g}_k(x) - g(x) + g(x) - g(y) + g(y) - \hat{g}_k(y) \\ &< \min_{x, y \in X} |g(x) - g(y)| - (g(y) - g(x)) \leq 0. \end{aligned}$$

It follows that, when a point $x^* \in S^*$ is selected at some iteration $k_0 \geq K$ —which happens with P -probability one—we have that $\hat{g}_k(x_k) < \hat{g}_k(y)$ for all $y \notin S^*$ and all $k \geq k_0$. From that iteration on, $x_k \in S^*$. \square

Theorem 4.1 shows that any *accumulation point* of a sequence $\{x_k\}$ produced by Algorithm SPRS will belong to the optimal set S^* . Of course, in case S^* is a singleton the whole sequence converges to the element in S^* .

In essence, the proof of Theorem 4.1 is based on the fact that the approximation $\hat{g}_k(x)$ approaches $g(x)$ uniformly on X . Although the theorem is proved for the pure random search method, it is clear that the same argument can be used in other settings where only *selection* of feasible points and *comparison* of function values are involved. An important requirement, however, is that every point in X be visited *infinitely often* with probability one. While this is trivially true for the pure random search method, not all algorithms satisfy the requirement. For example, consider the modified version of the pure random search method, mentioned earlier, where points are selected according to conditional distributions. If those conditional distributions block any point that had been previously selected, then a variable-sample version of such algorithm may not be convergent. Nevertheless, if one allows selection of previously visited points with *some* probability—as in Dorea and Gonçalves [1993]—then the requirement is satisfied and so the proof of Theorem 4.1 will be applicable to such algorithm.

An important feature of Algorithm SPRS results from the fact that the points are selected independently of previous choices. This implies that we can allow the underlying distributions to *depend on the decision variables* x , that is, we can have

$$g(x) = \int_{\Omega} G(x, \omega) P_x(d\omega).$$

This, in turn, yields considerable flexibility to the model, since in that case we can have $P_k = P_{x_k}$, that is, P_k may depend on the point selected at iteration k .

It is also worthwhile mentioning that, although Theorem 4.1 covers a number of random-search type methods, some algorithms may require more elaborate

proofs. Such is the case, for example, of a variable-sample version of the *simulated annealing* (SA) method. The reason is that the SA algorithm involves, in addition to selection of feasible points and comparison of function values, probabilities of moves that depend directly on function values. Nevertheless, the tools derived in Section 3—particularly Theorem 3.4—allow for derivation of a proper sample size schedule that ensures convergence of a variable-sample version of SA for stochastic problems. We refer to Homem-de-Mello [2001] for details.

5. A PRACTICAL IMPLEMENTATION OF SAMPLE SIZE UPDATES

The results in the previous section ensure that, as long as the sample sizes grow at a certain rate, a variable-sample version of random-search type methods will converge to the correct set of solutions. A simplistic implementation of that concept, however, can be impractical in terms of computational effort. The main issue is that, since the sample size is increased automatically along the algorithm, as the number of iterations grows the sampling procedure will take longer and longer.

A key observation to remedy the problem comes from the proof of Theorem 4.1. As discussed earlier, the basic property used in that proof is the fact that, once the approximations $\hat{g}_k(x)$ get “sufficiently close” to $g(x)$, the variable-sample scheme starts behaving exactly like the deterministic method. This suggests that *sample sizes need to increase only up to a certain point*.

One way to detect whether it is necessary to increase the sample size further is to perform a *t-test* to check if the values of $\hat{g}_k(\cdot)$ at the current point x_k and at the candidate point y are *statistically different*. If they are, then the sample size is not increased. In order to ensure theoretical convergence, we can increase the sample size (regardless of the t-test) only at some specific points in time.

As a concrete example, consider Algorithm SPRS described in Section 4, and suppose we want to implement a linear schedule. The step “update N_k to obtain N_{k+1} ” can be specified as

- (1) Perform a *paired* t-test between $\hat{g}_k(x_k)$ and $\hat{g}_k(y)$ to test the hypothesis $H_0 : g(x_k) = g(y)$;
- (2) If the *p*-value of the test is large (say, at least 0.2), then $N_{k+1} := N_k$; otherwise, $N_{k+1} := N_k + C$, where C is some constant.
- (3) Every K iterations (K specified a priori), do $N_{k+1} := N_k + C$.

Notice that we must perform a paired rather than an independent t-test, since the estimates $\hat{g}_k(x_k)$ and $\hat{g}_k(y)$ use the same random numbers. Also, if $N_1 \geq C$, then clearly we have that $N_k \geq kC/K$, so the linear schedule applies.

5.1 Numerical Results

We present now some numerical results to illustrate the adaptive approach discussed above. For that, we consider a stochastic version of the *traveling salesman problem* (TSP) defined in the following way. The nodes in the graph are numbered $0, 1, \dots, n$. Node zero is fixed as the starting (and ending) point; therefore, there are $n!$ possible tours. The cost of traveling to or from node zero to

any other node is constant, whereas the other arcs in the graph have a *random* cost associated with them. The graph is directed, so arcs (a, b) and (b, a) can have different costs. The goal is to find the tour with least *expected* cost.

Because the cost of each tour is simply the sum of the costs of the arcs that compose the tour, the objective function is *linear* with respect to the random variables. Therefore, it is clear that the problem can be solved with a deterministic algorithm, with the cost of each arc being replaced by its expected value. That is, there is no need to use a stochastic optimization method to solve the problem; nevertheless, we discuss this example because it allows us to verify whether the stochastic optimization method converges to the correct solution.

In order to test the adaptive approach discussed above, we incorporated those ideas into Algorithm SPRS—a variable-sample implementation of pure random search—described in Section 4. Let us call the resulting algorithm SPRS-AVS (the suffix “AVS” stands for “adaptive variable-sample”). For the sake of comparison, we also consider two alternative versions of Algorithm SPRS:

- SPRS-FVS: A variable-sample scheme with *fixed sample size*;
- SPRS-FFS: A *fixed sample* scheme with fixed sample size.

Of course, given the results derived in Section 3, Algorithm SPRS-FVS theoretically does not converge. Algorithm SPRS-FFS, on the other hand, falls into the framework of the *sample average approximation* (SAA) approach, studied by Kleywegt et al. [2001] in the context of discrete stochastic optimization. In Kleywegt et al. [2001], the following important properties are proved: for almost all sample paths, there exists a number \bar{N} such that, if the sample size used for the SAA is larger than \bar{N} , then the solution given by the SAA algorithm coincides with the solution of the original problem. Moreover, the probability that such event occurs (i.e., the two solutions coincide) goes to one *exponentially fast* with the sample size. Thus, for a fixed sample size N , Algorithm SPRS-FFS will converge to the optimal solution with some probability.

We used the following methodology in our tests. First, we randomly generated a cost matrix for a 6-node TSP (i.e., $n = 6$); the cost of each arc was generated independently from a discrete uniform distribution on $[4, 20]$. The resulting matrix was

$$C = \begin{pmatrix} 14 & 7 & 4 & 10 & 7 & 17 \\ 8 & 4 & 14 & 18 & 6 & 12 \\ 17 & 4 & 8 & 17 & 7 & 8 \\ 11 & 14 & 18 & 13 & 11 & 15 \\ 15 & 7 & 18 & 17 & 15 & 11 \\ 9 & 11 & 12 & 14 & 7 & 9. \end{pmatrix}$$

Each element C_{ij} corresponds to the *mean* cost of arc (i, j) . The actual cost of arc (i, j) was assumed to have uniform distribution on $(C_{ij} - 4, C_{ij} + 4)$. As pointed out above, we can solve this problem by solving a deterministic TSP. The optimal solution is the sequence 4, 1, 3, 2, 5, 6 corresponding to a cost equal to 36. We must notice here we ran several test cases, all with comparable results; thus, we chose one of them—the one corresponding to the above matrix C —as a representative.

Table I. Evolution of the Algorithms Under Study for the Stochastic TSP

CPU	AVS		FVS			FFS		
	10	50	10	100	500	10	100	500
0	0 (10)	0 (50)	0	0	0	0	0	0
30	0 (46)	0 (72)	0	0	0	8	0	0
60	0 (56)	0 (78)	10	0	0	15	0	0
90	0 (60)	0 (84)	18	0	0	15	0	0
120	0 (61)	0 (84)	11	0	0	15	0	0
150	1 (68)	0 (84)	2	0	0	15	0	0
180	14 (76)	0 (89)	16	20	0	15	20	0
210	18 (78)	14 (91)	15	20	0	15	20	0
240	18 (79)	20 (92)	9	20	0	15	20	0
270	18 (79)	20 (92)	4	20	0	15	20	0
300	18 (79)	20 (92)	18	20	0	15	20	0

We then tested the following configurations:

- Algorithm SPRS-AVS with $N_0 = 10$ and 50 and the constants $C = 10$, $K = 100$;
- Algorithm SPRS-FVS with $N = 10$, 100 and 500;
- Algorithm SPRS-FFS with $N = 10$, 100 and 500.

The rationale for these choices was that Algorithm SPRS-AVS, being adaptive, can start with a low sample size and increase it as needed. Algorithms SPRS-FVS and SPRS-FFS, on the other hand, require a fixed sample size to be established a priori. Thus, we used small, medium and large sample sizes to compare the behavior of the algorithms.

We did twenty independent runs for each configuration. To ensure fairness in the comparison, we used a different *seed* to generate the stream of random numbers for each run; the same seeds were used in all configurations. Moreover, the sequence of permutations used by the random search procedure was generated from an independent random stream and fixed across all runs; in other words, all runs see the same sequence of permutations. We made this choice because the purpose of the test is to compare different sampling schemes rather than the random search method itself; thus, by using the same sequence of permutations, we can compare how fast each algorithm reaches the optimal solution, if ever.

Table I displays the results. In order to compare the progress of the algorithms, we keep track of the best solution obtained after a certain amount of CPU time. Each row in Table I corresponds to a certain number of units of CPU time, and the columns indicate how often (out of twenty runs) the optimal solution was found by that time. The numbers in parentheses correspond to the average sample size used in the AVS method at the selected times. Since all the algorithms were coded in the same environment and have the same structure, the comparison really addresses the efficiency of each method. For brevity, we dropped the prefix SPRS from the names of the algorithms in the table.

Let us comment on the results shown in Table I. First, we can see that Algorithm SPRS-AVS is fairly consistent—even when the initial sample size is small ($N_0 = 10$), it still converges to the optimal solution 90% of the time. With a larger initial sample size, convergence occurs for all runs. Algorithm

SPRS-FVS, on the other hand, behaves erratically for the small sample size $N = 10$ —the numbers show that the algorithm jumps in and out the optimal solution. This is expected, given the large error of the estimates. For $N = 100$, however, the algorithm behaves much better and converges perfectly—which is interesting to observe considering that the algorithm is not convergent in theory. The phenomenon is explained by (20): when δ in that inequality is less than $(1/2) \min_{x \neq x^*} g(x) - g(x^*)$ (where x^* denotes the optimal solution), the probability that the algorithm jumps out of x^* is small when N_k is large. Thus, in theory the algorithm will eventually jump out of x^* , but that did not happen before the time limit we imposed to the runs. Similarly, the zeros in the column corresponding to $N = 500$ result from the large amount of time spent per iteration, which prevented the algorithm to reach the optimal solution even once before time expired. Finally, the results for Algorithm SPRS-FFS corroborate the theory developed in Kleywegt et al. [2001]: the probability of convergence (which, from the table, can be estimated as 75% for $N = 10$) grows exponentially fast with N , so that it is very close to one when $N = 100$. When $N = 500$ we see the same problem as with Algorithm SPRS-FVS: the large amount of time spent per iteration prevented the algorithm to reach the optimal solution even once.

In summary, Table I shows that, if “optimal” sample sizes are known, then algorithms SPRS-FVS and SPRS-FFS are a better choice than the adaptive version SPRS-AVS. Under- or over-estimation of those sample sizes, however, will cause those algorithms to behave poorly. In those cases the adaptive version can be very useful.

6. CONCLUSIONS

We have presented a general framework to show convergence of a certain class of methods to solve stochastic optimization problems, which we called *variable-sample* methods. Such procedures essentially consist of incorporating sampling into deterministic algorithms that use function evaluations only. Although a complete proof of convergence will depend on the method under scrutiny, we have provided general results to aid in that task. In particular, we have given conditions on the schedule of sample sizes $\{N_k\}$ that ensure consistency of the estimators and also guarantee some bounds on the deviation from true values. The results provided are general, in that no particular distribution is assumed.

To illustrate the type of analysis made possible by this framework, we have presented a modification of a random search algorithm to adapt it to the variable-sample scheme. Our results provide a schedule of sample sizes that guarantees convergence of the algorithm, without making strong assumptions on the underlying distributions. The ideas can be extended to similar methods. We have also proposed an adaptive version of the algorithm that automatically controls the sample sizes, and illustrated the ideas with numerical results. The numbers show that the variable-sample scheme is a viable alternative to fixed-sample methods, particularly if information about appropriate sample sizes is difficult to obtain.

Some questions, of course, remain open: on the theoretical side, the study of *rates of convergence* is very important to provide some intuition on the behavior of the algorithm. Unfortunately, however, the very definition of rate of convergence is not standard in the stochastic optimization literature, so some further study is required. On the practical side, the implementation of an algorithm based on the techniques described in this article in a realistic setting would require the development of appropriate stopping criteria (such as the ones suggested in Dorea [1990] and Hart [1999]). One possibility to implement these criteria in the stochastic context would be to perform multiple t -tests or use analysis of variance techniques in order to compare estimates of the objective function. This would be possible under the variable-sample structure, since the estimates obtained at different iterations are independent.

ACKNOWLEDGMENT

The author is grateful to Michael Fu and two anonymous referees for comments and suggestions that led to significant improvement in the presentation of the results.

REFERENCES

- ALLEN, T., ITTIWATTANA, W., AND BERNSHTEYN, M. 2002. An elitist genetic algorithm incorporating sequential subset selection. Manuscript, Ohio State University.
- ALREFAEI, M. H. AND ANDRADÓTTIR, S. 1999. A simulated annealing algorithm with constant temperature for discrete stochastic optimization. *Manage. Sci.* 45, 748–764.
- ALREFAEI, M. H. AND ANDRADÓTTIR, S. 2001. A modification of the stochastic ruler method for discrete stochastic optimization. *Europ. J. Oper. Res.* 133, 1, 160–182.
- ANDRADÓTTIR, S. 1995. A method for discrete stochastic optimization. *Manage. Sci.* 41, 1946–1961.
- ANDRADÓTTIR, S. 1996. A global search method for discrete stochastic optimization. *SIAM J. Optim.* 6, 513–530.
- BOESEL, J. AND NELSON, B. L. 1998. Accounting for randomness in heuristic simulation optimization. In *Simulation: Past, Present and Future. 12th European Simulation Multiconference 1998*. SCS, San Diego, Calif.
- CHUNG, K. L. 1974. *A Course in Probability Theory*, 2nd. ed. Academic Press, New York, NY.
- COOPER, W. L., HENDERSON, S. G., AND LEWIS, M. E. 2003. Convergence of simulation-based policy iteration. *Probability in the Engineering and Informational Sciences*. Forthcoming issue.
- COOPER, W. L. AND HOMEM-DE-MELLO, T. 2003. Hybrid approaches to revenue management using stochastic optimization and Markov decision processes. Manuscript, available on *Optimization Online* (www.optimization-online.org).
- DEMBO, A. AND ZEITOUNI, O. 1998. *Large Deviations Techniques and Applications*. Springer-Verlag, New York.
- DEVROYE, L. P. 1978. Progressive global random search of continuous functions. *Math. Program.* 15, 3, 330–342.
- DOREA, C. C. Y. 1990. Stopping rules for a random optimization method. *SIAM J. Cont. Optim.* 28, 4, 841–850.
- DOREA, C. C. Y. AND GONÇALVES, C. R. 1993. Alternative sampling strategy for a random optimization algorithm. *J. Optim. Theory Appl.* 78, 2, 401–407.
- DUPAČOVÁ, J. AND WETS, R. J.-B. 1988. Asymptotic behavior of statistical estimators and of optimal solutions of stochastic optimization problems. *Ann. Stat.* 16, 1517–1549.
- FOX, B. L. AND HEINE, G. W. 1995. Probabilistic search with overrides. *Ann. Appl. Prob.* 5, 4, 1087–1094.
- GELFAND, S. B. AND MITTER, S. K. 1989. Simulated annealing with noisy or imprecise energy measurements. *J. Optim. Theory Appl.* 62, 49–62.

- GONG, W. B., HO, Y. C., AND ZHAI, W. 1992. Stochastic comparison algorithm for discrete optimization with estimation. In *Proceedings of the 31st. Conference Decision Control*. 795–800.
- GUTJAHR, W. AND PFLUG, G. C. 1996. Simulated annealing for noisy cost functions. *J. Global Optim.* 8, 1–13.
- GUTJAHR, W. J. 2002. ACO algorithms with guaranteed convergence to the optimal solution. *Info. Process. Lett.* 82, 3, 145–153.
- HART, W. E. 1999. Sequential stopping rules for random optimization methods with applications to multistart local search. *SIAM J. Optim.* 9, 1, 270–290 (electronic).
- HO, Y. C., SREENIVAS, R. S., AND VAKILI, P. 1992. Ordinal optimization of DEDS. *Discr. Event Dyn. Syst.* 2, 61–88.
- HOMEM-DE-MELLO, T. 2001. On convergence of simulated annealing for discrete stochastic optimization. Manuscript, Ohio State University.
- HOMEM-DE-MELLO, T., SHAPIRO, A., AND SPEARMAN, M. L. 1999. Finding optimal material release times using simulation based optimization. *Manage. Sci.* 45, 86–102.
- KLEYWEGT, A., SHAPIRO, A., AND HOMEM-DE-MELLO, T. 2001. The sample average approximation method for stochastic discrete optimization. *SIAM J. Optim.* 12, 2, 479–502.
- MORTON, D. P. AND WOOD, R. K. 1998. On a stochastic knapsack problem and generalizations. In *Advances in Computational and Stochastic Optimization, Logic Programming, and Heuristic Search: Interfaces in Computer Science and Operations Research*, D. L. Woodruff, Ed. Kluwer Academic Publishers, Dordrecht, Netherlands, Chap. 5, 149–168.
- NORKIN, V. I., PFLUG, G. C., AND RUSZCZYŃSKI, A. 1998. A branch and bound method for stochastic global optimization. *Math. Prog.* 83, 425–450.
- RAO, M. M. 1984. *Probability Theory with Applications*. Academic Press, New York.
- ROBINSON, S. M. 1996. Analysis of sample-path optimization. *Math. Oper. Res.* 21, 513–528.
- RUBINSTEIN, R. Y. AND SHAPIRO, A. 1993. *Discrete Event Systems: Sensitivity Analysis and Stochastic Optimization by the Score Function Method*. Wiley, Chichester, England.
- RUDOLPH, G. 1996. Convergence of evolutionary algorithms in general search spaces. In *Proceedings of the 3rd IEEE Conference on Evolutionary Computation*. IEEE Computer Society Press, Los Alamitos, Calif. 50–54.
- SHAPIRO, A. 1991. Asymptotic analysis of stochastic programs. *Ann. Oper. Res.* 30, 169–186.
- SHAPIRO, A. 1993. Asymptotic behavior of optimal solutions in stochastic programming. *Math. Oper. Res.* 18, 829–845.
- SHAPIRO, A. AND HOMEM-DE-MELLO, T. 1998. A simulation-based approach to two-stage stochastic programming with recourse. *Math. Prog.* 81, 301–325.
- SHAPIRO, A. AND HOMEM-DE-MELLO, T. 2000. On rate of convergence of Monte Carlo approximations of stochastic programs. *SIAM J. Optim.* 11, 70–86.
- SHAPIRO, A. AND WARDI, Y. 1996. Convergence analysis of stochastic algorithms. *Math. Oper. Res.* 21, 615–628.
- SHWARTZ, A. AND WEISS, A. 1995. *Large Deviations for Performance Analysis*. Chapman & Hall, London.
- VAN DER VLERK, M. 2003. *Stochastic programming bibliography*. Available at <http://129.125.9.214/BIBLIO/STOPROG.HTML>.
- YAKOWITZ, S., L'ECUYER, P., AND VÁSQUEZ-ABAD, F. 2000. Global stochastic optimization with low-dispersion point sets. *Oper. Res.* 6, 939–950.
- YAN, D. AND MUKAI, H. 1992. Stochastic discrete optimization. *SIAM J. Control and Optim.* 30, 594–612.

Received October 2001; revised August 2002 and February 2003; accepted February 2003