

Fixed-Width Sequential Stopping Rules for a Class of Stochastic Programs

Güzin Bayraksan and Péguy Pierre-Louis
Systems and Industrial Engineering
The University of Arizona, Tucson, AZ 78721
guzinb@sie.arizona.edu, pierrelp@email.arizona.edu

Abstract

Monte Carlo sampling-based methods are frequently used in stochastic programming when exact solution is not possible. A critical component of Monte Carlo sampling-based methods is determining when to stop sampling to ensure the desired quality of the solutions. In this paper, we develop stopping rules for sequential sampling procedures that depend on the width of an optimality gap confidence interval estimator. The procedures solve a sequence of sampling approximations with increasing sample size to generate solutions and stop when the width of the confidence interval on the current solution's optimality gap plus an inflation factor falls below a pre-specified value, ϵ . We first present a method that takes the schedule of sample sizes as an input and provide guidelines on the growth of sample sizes. Then, we present a method that increases the sample sizes according to the current estimates of the optimality gap. The larger the estimates, the larger the increases in sample sizes. We provide conditions under which the procedures find ϵ -optimal solutions and terminate in a finite number of iterations with probability one and present empirical performance on test problems.

Key words. Stochastic programming; Monte Carlo sampling; stopping rules; confidence intervals

1 Introduction

We consider a stochastic optimization problem of the form

$$z^* = \min_{x \in X} E f(x, \tilde{\xi}), \quad (\text{SP})$$

where f is a real-valued function of decisions x and a realization of the random vector $\tilde{\xi}$, whose distribution is assumed known. $\Xi \subseteq \mathbb{R}^{d_{\tilde{\xi}}}$ denotes the support of $\tilde{\xi}$, $X \subseteq \mathbb{R}^{d_x}$ denotes the set of constraints that the decision vector x must obey and E is the expectation operator. (SP) includes a large class of problems found in statistics, simulation and operations research literature. It also embodies problems that are not typically thought of having objectives of minimizing an expected cost. For instance, when f is an indicator function of an event, the above minimizes the probability of that event. A fairly large class of stochastic problems with minimax objectives can also be cast as (SP) [43].

While other sampling schemes are possible, consider $\tilde{\xi}^1, \tilde{\xi}^2, \dots, \tilde{\xi}^n$ independent and identically distributed (i.i.d.) as $\tilde{\xi}$ and form a sampling problem as

$$z_n^* = \min_{x \in X} \frac{1}{n} \sum_{i=1}^n f(x, \tilde{\xi}^i). \quad (\text{SP}_n)$$

Let x^* denote an optimal solution to (SP) with optimal objective function value z^* . Similarly, let x_n^* denote an optimal solution to (SP_n) with optimal objective function value z_n^* . For a certain class of problems,

an optimal solution to (SP_n) provides an optimal solution to (SP) for large enough sample sizes, n , with probability one (w.p.1). In addition, the probability that the set of optimal solutions to (SP_n) is a subset of the set of optimal solutions to (SP) goes to one exponentially fast, see e.g., the survey by Shapiro [42]. The class of stochastic programs we consider fall into this category. The statement “an event happens w.p.1 for n large enough” means that for almost every $\omega = \{\tilde{\xi}^1, \tilde{\xi}^2, \dots\}$, there exists an integer $N(\omega)$ such that the event happens for all samples of size n from ω , $\{\tilde{\xi}^1, \tilde{\xi}^2, \dots, \tilde{\xi}^n\}$, with $n \geq N(\omega)$. Here, $N(\omega)$ is a random variable since it depends on the sequence of realizations ω . Let X^* denote the set of optimal solutions to (SP) and let X_n^* denote the set of optimal solutions to (SP_n) . The class of stochastic programs we consider satisfies the following:

- (A1) The event $A_n = \{X_n^* \subseteq X^*\}$ happens w.p.1 for n large enough,
- (A2) $1 - P(A_n) \leq Ce^{-\beta n}$, with $C, \beta > 0, \forall n$.

When X^* is a singleton, i.e., $X^* = \{x^*\}$, A_n corresponds to the event that (SP_n) has a unique optimal solution which is equal to x^* . Suppose $\tilde{\xi}^1, \tilde{\xi}^2, \dots, \tilde{\xi}^n$ are i.i.d. as $\tilde{\xi}$ and $E \sup_{x \in X} |f(x, \tilde{\xi})| < \infty$. Then, (A1) is satisfied, for instance, when $|\Xi| < \infty$, $f(\cdot, \tilde{\xi})$ is piecewise-linear convex for every $\tilde{\xi} \in \Xi$, X is a closed convex polyhedral and $X^* \neq \emptyset$ is bounded [44]. Alternatively, (A1) is satisfied when X is a finite set [26]. Slightly more stringent conditions are required for (A2). Note that the condition in (A2) is satisfied for all n and is typically referred to as the exponential rate of convergence. Such results are often obtained through large deviations theory. Other convergence results in stochastic programming based on large deviations theory include [7, 23]. The samples can also be generated in a non-i.i.d. fashion, for instance, to reduce variance and also to have a higher rate of convergence, i.e., to have a larger value of β in (A2). Homem-de-Mello [20] provides conditions under which (A1) and (A2) are satisfied under non-i.i.d. sampling. In this paper, we also explicitly assume that observations of $\tilde{\xi}$ can be generated and instances of (SP_n) can be solved for sufficiently large n .

While the class of stochastic optimization problems we consider constitutes a larger group, our main focus is on stochastic programs with recourse. The class of stochastic programs with recourse that satisfy the above conditions under appropriate sampling schemes include two-stage stochastic linear programs with finite Ξ and two-stage stochastic integer programs with integer first stage. Many important applications of stochastic programs in the literature have been successfully modeled by this class of programs. Some of these applications include strategic supply chain planning [1], interdiction of nuclear material smuggling [32], power capacity expansion planning [22], stochastic scheduling [24] and shortest path with random arc failures [45], among others.

Consistency results like (A1) suggest solving a sampling problem (SP_n) with a sufficiently large sample size n to obtain a high-quality solution. However, this requires a reliable way to determine the sample size so that the quality of the resulting solution is of a desired accuracy. When “accuracy” is measured as the probability of obtaining an optimal solution to (SP), (A2) can be used, provided reliable estimates of C and β are available, or these values are known. In fact, such estimates of sample sizes, n , have been suggested in the literature to find optimal solutions to (SP) with a desired probability by solving a sampling problem, see e.g., [42]. These sample size estimates provide valuable insight into the complexity of solving stochastic problems with a sampling approximation. Unfortunately, they can be too conservative for practical purposes. In contrast to these theoretical sample size estimates, empirical tests indicate that relatively small sample sizes with reasonable solution times of (SP_n) yield solutions within an estimated 1% of optimality [45].

In this paper, we approach the problem of finding a minimal sample size to reach a desired accuracy in a more algorithmic fashion. We consider sequential sampling procedures and mainly focus on *external sampling* where a sampling problem with increasing sample size is solved to generate candidate solutions. We provide stopping rules and guidelines on how to increase the sample sizes such that these procedures stop at the minimal sample size, finding solutions within a desired accuracy. Our measure of “accuracy” is defined as being sufficiently close to the optimal objective function value. We measure this through a confidence interval formed on the optimality gap of the current candidate solution. The procedures stop when the width of the confidence interval on the optimality gap plus an inflation factor falls below a pre-specified, fixed value, ϵ . We show that the procedures stop w.p.1 for a given $\epsilon > 0$ and asymptotically find ϵ -optimal solutions to (SP), w.p.1. Our preliminary computational results indicate that the sample sizes when the procedures stop are far smaller than the theoretical estimates, while still providing high-quality solutions

with high probability (see §6.3).

There are methods developed in the literature that assess the quality of a given approximate solution by forming confidence intervals on its optimality gap [2, 30, 35]. Our stopping rules can use these; in fact, we make use of the confidence interval procedures in [2]. However, we note that these confidence interval methods are *static* in nature. That is, the candidate solution and the sample size are *fixed* and are *not adaptive, hence not random*. When such confidence intervals are used in a truly algorithmic setting, as we do in this paper, the sequence of candidate solutions and the sample size when the procedure stops become random variables. Therefore, their analysis is quite different from their static counterparts. Sequential analyses have been studied extensively in statistics [6, 12, 33], and in simulation of stochastic systems [13, 21, 25, 28, 29]; see also the stopping rules for random search algorithms for deterministic global optimization [4, 5, 9, 15]. In the context of stochastic programming, however, even though stopping rules have been proposed and used for a number of sampling-based methods for stochastic programs [8, 16, 17, 27, 35], stopping criteria that *control* the quality of the obtained solutions and their sequential nature have not been fully investigated. An exception is the stochastic quasi-gradient (SQG) methods where gradients and subgradients are replaced by their sampling estimates within a steepest descent method. Pflug [37] surveys stopping rules and step sizes for SQG methods. SQG can handle larger classes of problems but does not make use of the underlying structure of the problems. Our stopping rules for sequential sampling procedures are rooted in sample average approximations which can readily use mature tools developed for the specific problem at hand to efficiently solve the sampling problems. Another attractive feature of the stopping rules developed in this paper is that they can handle *both independent and dependent* estimates, such as the ones obtained by augmenting the previous observations with new observations at each iteration. We investigate the empirical performance of independent vs. dependent estimates in §6.

A number of researchers studied conditions on the growth of sample sizes when solving sampling-based approximations of stochastic programs. Polak and Royset [38] estimate sample sizes to minimize computational effort when solving a class of stochastic nonlinear programs. Futschik and Pflug [10] propose an optimal sampling strategy aimed at reducing the computational costs incurred in solving stochastic discrete programs through Monte Carlo sampling-based approximations. Homem-de-Mello [19] establishes conditions on the sample sizes such that the objective function estimators are consistent for variable sample methods. Pasupathy [36] establishes growth rates for the sample sizes in relation to the sequence of error tolerances to ensure optimal convergence of the solutions of retrospective-approximation algorithms for stochastic root finding and simulation optimization. In [36], the focus is on efficient convergence of solutions (x_n^*) to an optimal solution (x^*) whereas we are indifferent to how close we are to the set of optimal solutions as long as the optimality gap is sufficiently small. Similar to our goal, Morton [31] develops stopping rules theory and studies the minimal required growth in the sample sizes for a class of sampling-based algorithms with asymptotically normal optimality gap estimators. Unfortunately, asymptotic normality is not typically satisfied by the optimality gap estimators for stochastic programs. Bayraksan and Morton [3] improve on [31] by relaxing the asymptotic normality assumption and by replacing the unknown variance by a sample variance estimator in the sample size growth conditions. This paper continues in this line of research and provides alternative stopping rules for a class of stochastic programs.

We note several differences between the stopping rules of this paper and those of [3]. First, the stopping rules in [3] use *point estimators* of the optimality gap and variance, whereas the stopping rules of this paper focus on the *interval estimator* of the optimality gap. In [3], the stopping criteria depend on the point estimator of the optimality gap falling below a pre-specified *fraction* of the sampling variance; hence are considered *relative-width*. In contrast, the stopping rules in this paper depend on the *width* of the interval estimator falling below a pre-specified level, ϵ , and are considered *fixed-width* (see §2 for details). The implication of this is that the relative-width stopping rules of [3] provide a quality statement with respect to the sample variance estimator—if the variability is large, a worse quality statement is made. In contrast, the fixed-width stopping rules of this paper always aim to find an ϵ -optimal solution. That is, they use a “fixed” quality statement. Note that the variability of the problem is typically captured in the sampling error part of an interval estimator, and this part shrinks to zero as the sample size increases even if the variability is large. Another major difference is that, in this paper, we mainly focus on finding the candidate solutions via solving a sampling problem with increasing sample size, whereas in [3], the candidate solutions can be found by any method that eventually generates optimal solutions, w.p.1. We also restrict our attention to problems that satisfy (A1) and (A2) and sometimes focus on a specific class of interval estimators. These

restrictions pay off in several ways. First, they allow us to obtain stronger results, e.g., asymptotic validity results with probability one instead of $1 - \alpha$ as in [3]. Second, compared to [3], we significantly increase our analysis on the growth of sample sizes. For instance, we analyze the growth needed in the sample sizes used for generating versus evaluating a candidate solution as well as the interplay between them such that the expected sample sizes used are finite. We examine the rate of increase in the sample sizes when the sequential procedures stop. Third, we present a way to increase the sample sizes in a stochastic way, by using the current estimates, in an effort to make the procedures more adaptive to the problem at hand.

The rest of the paper is organized as follows. In the next section, we briefly go over how to construct confidence intervals on the optimality gap of a given solution and list assumptions on the statistical estimators to be used within the stopping rules. Next, we review and extend relevant results for asymptotic behavior of random variables with random sample sizes. In §4, we develop a fully sequential procedure where the sample sizes are increased according to deterministic schedules. We show that the procedure stops in a finite number of iterations using finite number of sample sizes for $\epsilon > 0$ and asymptotically finds an ϵ -optimal solution with probability one. We provide conditions on the growth of sample sizes such that the expected number of samples used is finite. In §5, we develop a procedure where the sample sizes are increased according to the current statistical estimates. In §6, we present computational results on two-stage stochastic linear and integer programs and provide guidelines for implementation. We end in §7 with a summary and conclusions.

2 Confidence Interval Estimation of Optimality Gap

Given a feasible solution $x \in X$, we denote the optimality gap of x by $\mu_x = Ef(x, \tilde{\xi}) - z^*$. Given a feasible solution $x \in X$, a sample of size n , $\tilde{\xi}^1, \tilde{\xi}^2, \dots, \tilde{\xi}^n$, and a desired value for $\alpha \in (0, 1)$, suppose we have at hand a method to form a $(1 - \alpha)$ -level confidence interval on μ_x . The confidence interval consists of two parts: (i) a point estimate of μ_x plus (ii) a sampling error part. We denote the point estimator of the optimality gap by $G_n(x)$ and the sampling error by $v_n^\alpha(x)$. Then, a $(1 - \alpha)$ -level confidence interval (CI) on μ_x is formed by

$$[0, G_n(x) + v_n^\alpha(x)]. \quad (1)$$

Since $\mu_x \geq 0$ for all $x \in X$, we inherently assume that $G_n(x) \geq 0$ and $v_n^\alpha(x) \geq 0$, w.p.1. If, for instance, there exists a $G'_n(x)$ which can take negative values, a new estimator $G_n(x) = \max\{G'_n(x), 0\}$ can be used instead.

Our sequential stopping rules depend on the confidence interval formed in (1). Given a sequence of candidate solutions, we form confidence intervals on their optimality gaps and stop when the width of the confidence interval plus an inflation factor falls below a pre-specified value, ϵ . This inflation factor is a non-negative function of the sample size, n , and it shrinks to zero as $n \rightarrow \infty$. We denote it by $h(n)$. Examples include $h(n) = 1/\sqrt{n}$ or $h(n) = 1/\ln(n)$. Note that tightening the stopping criterion in this manner is quite common in classical sequential statistical procedures, see, e.g., [6]. The procedures we develop in this paper stop at candidate solution x using sample size n the first time

$$G_n(x) + v_n^\alpha(x) + h(n) \leq \epsilon$$

is satisfied. Note that, in contrast, the sequential sampling procedure developed in [3] stops the first time $G_n(x) \leq h' s_n(x) + \epsilon'$ is satisfied for some positive constants $h' > 0$ and $\epsilon' > 0$, where $s_n^2(x)$ is a sampling variance estimator typically used in forming the sampling error, $v_n^\alpha(x)$, e.g., see (2b) and (3) below.

Let the sample size n_k satisfy $n_k \rightarrow \infty$ as $k \rightarrow \infty$. We make the following assumptions with respect to the point estimator of the optimality gap and the sampling error part of the interval estimator:

(A3) $\lim_{k \rightarrow \infty} P(\sup_{x \in X} |G_{n_k}(x) - \mu_x| > \delta) = 0$ for any $\delta > 0$,

(A4) Let $\alpha \in (0, 1)$ be given. Then, $\lim_{k \rightarrow \infty} P(\sup_{x \in X} v_{n_k}^\alpha(x) > \delta) = 0$, for any $\delta > 0$.

In words, assumption (A3) requires the optimality gap estimator to uniformly converge in probability over X to the actual optimality gap. Similarly, assumption (A4) requires the sampling error to uniformly converge to zero in probability over X .

As an example of an interval estimator given in (1), consider the single replication procedure (SRP) [2]. Let $\bar{f}_n(x) = \frac{1}{n} \sum_{i=1}^n f(x, \tilde{\xi}^i)$. With this notation, when the sampling problem (SP_n) is solved, $z_n^* = \bar{f}_n(x_n^*)$. For any $x \in X$, the point estimator of the optimality gap, μ_x , and its associated sample variance for SRP are obtained by the following:

$$G_n(x) = \frac{1}{n} \sum_{i=1}^n f(x, \tilde{\xi}^i) - \min_{x \in X} \frac{1}{n} \sum_{i=1}^n f(x, \tilde{\xi}^i) = \bar{f}_n(x) - \bar{f}_n(x_n^*), \quad (2a)$$

$$s_n^2(x) = \frac{1}{n-1} \sum_{i=1}^n \left[(f(x, \tilde{\xi}^i) - f(x_n^*, \tilde{\xi}^i)) - (\bar{f}_n(x) - \bar{f}_n(x_n^*)) \right]^2, \quad (2b)$$

where the same set of observations, $\tilde{\xi}^1, \tilde{\xi}^2, \dots, \tilde{\xi}^n$, are used in both terms in (2a) and the corresponding terms in (2b). Note that $G_n(x) \geq 0$, w.p.1. Given the optimality gap and variance estimators in (2a) and (2b), the single replication procedure (SRP) forms a $(1 - \alpha)$ -level confidence interval on μ_x through

$$\left[0, G_n(x) + \frac{t_{n-1, \alpha} s_n(x)}{\sqrt{n}} \right], \quad (3)$$

where $t_{n, \alpha}$ is the $1 - \alpha$ quantile of the Student's t distribution with n degrees of freedom. Similarly, let z_α be the $1 - \alpha$ quantile of the standard normal. Under conditions provided in [2] the point estimators in (2a) and (2b) are consistent (e.g., $G_n(x) \rightarrow \mu_x$, w.p.1 as $n \rightarrow \infty$ for all $x \in X$), and the confidence interval in (3) is asymptotically valid, i.e.,

$$\liminf_{n \rightarrow \infty} P \left(\mu_x \leq G_n(x) + \frac{z_\alpha s_n(x)}{\sqrt{n}} \right) \geq 1 - \alpha. \quad (4)$$

It has been observed that, for some problems, SRP can perform poorly with small sample sizes [2]. Variants of this procedure enhance empirical performance by making more than one replication and hence lowering the probability of undercoverage. For example, the averaged two-replication procedure, A2RP, selects n even and divides the observations into two random partitions. Then, $G_{n/2, i}(x)$ and $s_{n/2, i}^2(x)$ are calculated as given in (2a) and (2b) for partition $i = 1, 2$. These gap and variance estimators are pooled to obtain A2RP estimators, $G_n(x) = \frac{1}{2}(G_{n/2, 1}(x) + G_{n/2, 2}(x))$ and $s_n^2 = \frac{1}{2}(s_{n/2, 1}^2(x) + s_{n/2, 2}^2(x))$. Then, the A2RP confidence interval on μ_x is calculated in the same way as in (3). The number of replications can be increased and confidence intervals can be estimated by ArRP by pooling r estimators. In the rest of the paper, we let A1RP denote SRP and use the two interchangeably.

Suppose ArRP, $r \geq 1$, is used to form the confidence interval in (1). Then, a sufficient condition for (A3) to be satisfied is that the sample means $\bar{f}_{n_k}(x)$ converge uniformly to $Ef(x, \tilde{\xi})$ on X , w.p.1. We note that uniform convergence in this setting is not a restrictive assumption, since for problems that satisfy (A1), this is usually already satisfied, for instance, when X is finite [20, 26]. A sufficient condition for (A4) to be satisfied is that for some $M > 0$, $\lim_{n \rightarrow \infty} P(\sup_{x \in X} s_n^2(x) < M) = 1$. One set of conditions that ensure this is given in [2].

Note that there are two main differences between, for instance, the single replication procedure (SRP) (or any other procedure) to form confidence intervals, for a given $x \in X$ and a given sample size n , and a sequential procedure that uses such a confidence interval as a stopping rule. SRP is static, that is, x is fixed and the resulting consistency and validity results such as (4) are shown as $n \rightarrow \infty$ for a fixed $x \in X$. In contrast, in the sequential procedures: (i) the candidate solution x changes and hence is a random variable and (ii) the sample size *when* the procedure stops is also a random variable. The analysis of random variables indexed by random sample sizes can significantly differ from their counterparts with deterministic sample sizes. Next, we review and extend relevant results for random variables that are indexed by random sample sizes.

3 Limit Theorems for Random Variables with Random Indices

Consider a sequence of random variables $\{Y_k, k = 1, 2, \dots\}$ and suppose the index k is replaced by a random index $\{K(\epsilon), \epsilon \geq 0\}$ such that $K(\epsilon) \rightarrow \infty$, w.p.1 as $\epsilon \downarrow 0$. Usually, these indices are defined as sample sizes but they can also represent other indices. For example, k could be the iteration number of a procedure and Y_k the random variable generated at iteration k , such as the candidate solution or a subgradient estimate generated at iteration k . Then, the random index $K(\epsilon)$ could be the *random* iteration number *when* the algorithm stops with respect to a stopping rule for a given ϵ and $Y_{K(\epsilon)}$ would then be the value of the random variable Y_k at iteration $K(\epsilon)$. For a definition of random indices in the context our of sequential procedures, see §4.1 and §5. The result we review below is concerned with convergence of the random variables as the random index tends to infinity.

Theorem 1 ([14, 39]) *Let $\{Y_k, k = 1, 2, \dots\}$ be a sequence of random variables and $\{K(\epsilon), \epsilon \geq 0\}$ be a family of positive, integer-valued random variables such that $K(\epsilon) \rightarrow \infty$, w.p.1 as $\epsilon \downarrow 0$. Suppose*

$$Y_k \rightarrow Y, \text{ w.p.1 as } k \rightarrow \infty. \quad (5)$$

Then, $Y_{K(\epsilon)} \rightarrow Y$, w.p.1 as $\epsilon \downarrow 0$.

In Theorem 1, there is no independence assumption between the random variables $\{Y_k, k = 1, 2, \dots\}$ and the random index $\{K(\epsilon), \epsilon \geq 0\}$. In fact, in many settings, including ours, $K(\epsilon)$ is defined by the random variables $\{Y_k, k = 1, 2, \dots\}$. Without independence, general limiting results for $Y_{K(\epsilon)}$ like above are harder to obtain under weaker conditions [14]. However, consider a special case of (5) where $Y_k = a$, w.p.1 for k large enough where $a \in \mathbb{R}$ is a constant. In this case, it is possible to obtain a slightly stronger result, which we state in the proposition below.

Proposition 2 *Let $\{Y_k, k = 1, 2, \dots\}$ be a sequence of random variables and $\{K(\epsilon), \epsilon \geq 0\}$ be a family of positive, integer-valued random variables such that $K(\epsilon) \rightarrow \infty$, w.p.1 as $\epsilon \downarrow 0$. Let $a \in \mathbb{R}$. Suppose*

$$Y_k = a, \text{ w.p.1 for } k \text{ large enough.} \quad (6)$$

Then, $Y_{K(\epsilon)} = a$, w.p.1 for ϵ small enough.

Proof. Our proof is similar to that of Theorem 1 in [39]. Let $A = \{\omega : Y_{K(\epsilon, \omega)}(\omega) \neq a \text{ for } \epsilon \text{ small enough and } K(\epsilon, \omega) \rightarrow \infty \text{ as } \epsilon \downarrow 0\}$. It is sufficient to show that $P(A) = 0$. To this end, select an element ω from A . Then, for every k , one can find an $\epsilon \equiv \epsilon(\omega)$ such that $K(\epsilon, \omega) \geq k$ and

$$|Y_{K(\epsilon, \omega)}(\omega) - a| > 0. \quad (7)$$

By increasing k , one can obtain a sequence, which we denote by $\{\epsilon_t\}$ such that $\epsilon_t \downarrow 0$ as $t \rightarrow \infty$. Now, set $k_t = K(\epsilon_t, \omega)$. Since $\omega \in A$, $k_t \rightarrow \infty$ as $t \rightarrow \infty$. Then, by (7),

$$|Y_{k_t}(\omega) - a| > 0,$$

for the sequence $Y_{k_t}(\omega)$. Hence, ω is contained in the set $B = \{\omega : Y_k \neq a \text{ for } k \text{ large enough and } K(\epsilon) \rightarrow \infty \text{ as } \epsilon \downarrow 0\}$, i.e., $A \subseteq B$. By hypothesis, $P(B) = 0$ and hence the desired result follows. ■

The above result is not surprising, as once the index reaches the point where (6) is satisfied, the index being random does not have much of an impact. We are now ready to present our sequential procedures.

4 Fully Sequential Procedure

4.1 Description

In our sequential procedures, we use two independent streams of observations, $\tilde{\xi}_1^1, \tilde{\xi}_1^2, \dots$ and $\tilde{\xi}_2^1, \tilde{\xi}_2^2, \dots$, one to generate a sequence of candidate solutions, which we denote as $\{\hat{x}_k\}$, and the other to form confidence

intervals on the optimality gap of the candidate solutions. At iteration k , we use m_k observations to solve (SP $_{m_k}$) and generate \hat{x}_k by setting $\hat{x}_k = x_{m_k}^*$. Independently, we use n_k observations to calculate $G_{n_k}(\hat{x}_k)$ and $v_{n_k}^\alpha(\hat{x}_k)$ to form a CI on $\mu_{\hat{x}_k}$ as in (1). We stop when this CI plus an inflation factor falls below a pre-specified value, ϵ . Otherwise, at iteration $k + 1$, we have the options to either reuse the previously generated n_k observations and generate $n_{k+1} - n_k$ new observations, or generate an entirely new set of n_{k+1} observations. Similarly, we can either augment m_k with $m_{k+1} - m_k$ new observations or generate an entirely new set of m_{k+1} observations. When warm-starting can be used, augmentation can reduce solution times. On the other hand, generating an entirely new set of observations prevents being stuck in a bad sample path [19] and hence, can increase the performance of the sequential procedures. To this end, we define two new parameters as input to the procedures, k_f^m and k_f^n , denoting the resampling frequency for generating candidate solutions and for assessing solution quality, respectively. For instance, when $k_f^n = 1$, independent observations for assessing solution quality are generated at every iteration. We set $k_f^n = \infty$ when samples are always augmented.

The first method we study, called the fully sequential procedure (FSP), increases the sample sizes according to deterministic schedules. Schedule $\{m_k\}$ dictates increases in sample sizes to generate candidate solutions and schedule $\{n_k\}$ dictates increases in sample sizes to assess the quality of the candidate solutions. We assume both $m_k \rightarrow \infty$ and $n_k \rightarrow \infty$ as $k \rightarrow \infty$. Later, we will provide other conditions on these sample sizes under which certain desired properties are satisfied. Below, we give a short statement of the procedure.

FSP:

Input: Sample size schedules $\{m_k\}$ and $\{n_k\}$, inflation factor $h(n)$, desired values of $\epsilon > 0$, $0 < \alpha < 1$, and resampling frequencies k_f^m and k_f^n .

Output: A candidate solution $\hat{x}_{K(\epsilon)}$ with an estimated optimality gap of ϵ .

0. (Initialization) Set $k = 1$, generate m_1 observations $\tilde{\xi}_1^1, \tilde{\xi}_1^2, \dots, \tilde{\xi}_1^{m_1}$ and independently generate n_1 observations $\tilde{\xi}_2^1, \tilde{\xi}_2^2, \dots, \tilde{\xi}_2^{n_1}$.

1. (Find candidate solution) Using the observations $\tilde{\xi}_1^1, \dots, \tilde{\xi}_1^{m_k}$, solve (SP $_{m_k}$) to obtain $x_{m_k}^*$. Set $\hat{x}_k = x_{m_k}^*$.

2. (Assess solution quality) Using the observations $\tilde{\xi}_2^1, \dots, \tilde{\xi}_2^{n_k}$, calculate $G_{n_k}(\hat{x}_k)$ and $v_{n_k}^\alpha(\hat{x}_k)$.

3. (Check stopping criterion) If

$$G_{n_k}(\hat{x}_k) + v_{n_k}^\alpha(\hat{x}_k) + h(n_k) \leq \epsilon \tag{8}$$

then set $K(\epsilon) = k$, $N(\epsilon) = n_k$, $M(\epsilon) = m_k$, and $\hat{x}_{K(\epsilon)} = \hat{x}_k$ and stop. Output candidate solution $\hat{x}_{K(\epsilon)}$ with confidence interval on $\mu_{\hat{x}_{K(\epsilon)}}$ as $[0, \epsilon]$.

4. (Increase sample size) If k_f^m divides $k+1$ then sample $\tilde{\xi}_1^1, \tilde{\xi}_1^2, \dots, \tilde{\xi}_1^{m_{k+1}}$ independently of samples generated in previous iterations. Else, sample $m_{k+1} - m_k$ additional observations, $\tilde{\xi}_1^{m_k+1}, \dots, \tilde{\xi}_1^{m_{k+1}}$. Similarly, if k_f^n divides $k + 1$ then sample $\tilde{\xi}_2^1, \tilde{\xi}_2^2, \dots, \tilde{\xi}_2^{n_{k+1}}$ independently of samples generated in previous iterations. Else, sample $n_{k+1} - n_k$ additional observations, $\tilde{\xi}_2^{n_k+1}, \dots, \tilde{\xi}_2^{n_{k+1}}$. Set $k = k + 1$ and go to 1.

If FSP stops according to (8), the iteration when the procedure stops, $K(\epsilon)$, is a random variable defined by

$$K(\epsilon) = \inf_{k \geq 1} \{k : G_{n_k}(\hat{x}_k) + v_{n_k}^\alpha(\hat{x}_k) + h(n_k) \leq \epsilon\}. \tag{9}$$

It is the first iteration when the inflated CI falls below ϵ . Later we provide conditions under which $K(\epsilon)$ is finite w.p.1. The sample sizes when the procedure stops, $N(\epsilon)$ and $M(\epsilon)$, used to form the CI and to generate the candidate solution respectively, are then defined by $N(\epsilon) = n_{K(\epsilon)}$ and $M(\epsilon) = m_{K(\epsilon)}$. We denote the solution found by FSP as $\hat{x}_{K(\epsilon)}$. We would like $\hat{x}_{K(\epsilon)}$ to have an optimality gap of at most ϵ with a desired confidence, i.e., $P(\mu_{\hat{x}_{K(\epsilon)}} \leq \epsilon) \geq 1 - \alpha$. In order to achieve such a result, first, we need to be able to find ϵ -optimal solutions. (A1) suggests that this can be achieved if $M(\epsilon)$ grows sufficiently large. Next, due to the stopping rule (8), we need the confidence interval to be able to detect an ϵ -optimal solution with a high probability. (8) implies that $P(\mu_{\hat{x}_{K(\epsilon)}} \leq \epsilon) \geq P(\mu_{\hat{x}_{K(\epsilon)}} \leq G_{N(\epsilon)}(\hat{x}_{K(\epsilon)}) + v_{N(\epsilon)}^\alpha(\hat{x}_{K(\epsilon)}))$ and when the probability on the right-hand side is $\geq 1 - \alpha$, we can achieve our desired result. However, note that even in the case where the sample size, n , and the candidate solution, \hat{x} , are not random (as opposed to random variables $N(\epsilon)$ and $\hat{x}_{K(\epsilon)}$), similar results stated in (4) are shown only asymptotically, as $n \rightarrow \infty$. Therefore, we also need the sample size $N(\epsilon)$ to grow.

To obtain the desired growths in the random sample sizes w.p.1 as $\epsilon \downarrow 0$, we inflate the confidence interval by $h(n_k)$. Otherwise, when $G_{n_k}(\hat{x}_k) + v_{n_k}^\alpha(\hat{x}_k) = 0$, the procedure stops at iteration k regardless of the value of $\epsilon > 0$. For instance, when SRP is used, this can happen when $\hat{x}_k = x_{n_k}^*$ even when $\mu_{\hat{x}_k} \gg \epsilon$. To prevent this from happening, a positive inflation factor that shrinks to zero while $n_k \rightarrow \infty$ as $k \rightarrow \infty$ can be used. Examples of such inflation factors include $h(n) = 1/n$, $h(n) = 1/\log(n)$, etc. To see the effect of $h(n_k)$, let $m_k = n_k = n_0 + k - 1$ for some initial sample size $n_0 \geq 2$. That is, the sample sizes are increased one by one starting from n_0 . Let $h(n_k) = 1/\sqrt{n_k}$. Then, $P(K(\epsilon) \leq k) = 0$ for $k \leq \lfloor \frac{1}{\epsilon^2} \rfloor + 1 - n_0$, where $\lfloor \cdot \rfloor$ returns the largest integer less than or equal to its argument. With the inflation factor, by taking $\epsilon \downarrow 0$, the iteration $K(\epsilon)$, and hence $N(\epsilon)$ and $M(\epsilon)$, can be made arbitrarily large, w.p.1.

4.2 Asymptotic Validity, Finite Stopping and Finite Sample Sizes

In this section, we present three main properties of FSP. First, we show that FSP finds an ϵ -optimal solution with probability one as $\epsilon \downarrow 0$. Second, we guarantee, under certain assumptions, the finiteness of the number of iterations and the sample sizes used when the procedure stops. Finally, we provide conditions under which the expected number of iterations and the expected sample sizes when the procedure stops are finite.

Theorem 3 *Assume (A1) holds. Let $\{m_k\}$, $\{n_k\}$ be positive, nondecreasing sequences such that $m_k \rightarrow \infty$ and $n_k \rightarrow \infty$ as $k \rightarrow \infty$. Let $h(\cdot)$ be a nonnegative, nonincreasing function such that $h(n_k) \rightarrow 0$ as $n_k \rightarrow \infty$, and suppose $k < \infty$ implies that $m_k < \infty$, $n_k < \infty$, and $h(n_k) > 0$.*

(i) *Consider FSP that stops at iteration $K(\epsilon)$ according to stopping criterion (9), defined in terms of ϵ , using $N(\epsilon) = n_{K(\epsilon)}$ and $M(\epsilon) = m_{K(\epsilon)}$ samples. Then,*

$$\lim_{\epsilon \downarrow 0} P(\mu_{\hat{x}_{K(\epsilon)}} \leq \epsilon) = 1.$$

(ii) *In addition, assume (A3) and (A4) hold. Given $\epsilon > 0$,*

$$P(K(\epsilon) < \infty) = P(N(\epsilon) < \infty) = P(M(\epsilon) < \infty) = 1.$$

Proof. (i) Let $Y_k = \mu_{\hat{x}_k} = E[f(\hat{x}_k, \tilde{\xi}) | \hat{x}_k] - z^*$. Then, by (A1) and the fact that $m_k \rightarrow \infty$ as $k \rightarrow \infty$, $Y_k = 0$, w.p.1 for k large enough. The hypotheses of the theorem ensure $K(\epsilon) \rightarrow \infty$, w.p.1 as $\epsilon \downarrow 0$. Therefore, using Proposition 2, $Y_{K(\epsilon)} = \mu_{\hat{x}_{K(\epsilon)}} = 0$, w.p.1 for ϵ small enough; hence, the desired result follows.

(ii) We first show that FSP stops in a finite number of iterations, w.p.1. Note that

$$\begin{aligned} P(K(\epsilon) = \infty) &\leq \lim_{k \rightarrow \infty} P(G_{n_k}(\hat{x}_k) + v_{n_k}^\alpha(\hat{x}_{n_k}) + h(n_k) > \epsilon) \\ &\leq \lim_{k \rightarrow \infty} P(|G_{n_k}(\hat{x}_k) - \mu_{\hat{x}_k}| + v_{n_k}^\alpha(\hat{x}_{n_k}) > \epsilon - h(n_k) - \mu_{\hat{x}_k}) \\ &\leq \lim_{k \rightarrow \infty} P\left(\sup_{x \in X} |G_{n_k}(x) - \mu_x| + \sup_{x \in X} v_{n_k}^\alpha(x) > \epsilon - h(n_k) - \mu_{\hat{x}_k}\right) \\ &= 0, \end{aligned} \tag{10}$$

where (10) follows from (A1), (A3), (A4) and the fact that $h(n_k)$ is shrinking to 0 as k , hence n_k , tends to infinity. $P(N(\epsilon) < \infty) = 1$ and $P(M(\epsilon) < \infty) = 1$ follow from the hypothesis that $k < \infty$ implies $n_k < \infty$ and $m_k < \infty$. ■

Theorem 3 shows that FSP, asymptotically, finds an ϵ -optimal solution with probability one. While this result is asymptotic, in practice, FSP can be used for a given $\epsilon > 0$ and for ϵ small enough we can expect $P(\mu_{\hat{x}_{K(\epsilon)}} \leq \epsilon) \gtrsim 1 - \alpha$. Note that even in much simpler settings, such as forming confidence intervals on the mean, validity of sequential procedures are only shown asymptotically, see, e.g., [6]. Nevertheless, sequential procedures are successfully used in practice in different fields such as clinical statistical studies [11]. Note that part (i) of Theorem 3 is stronger than $\lim_{\epsilon \downarrow 0} P(\mu_{\hat{x}_{K(\epsilon)}} \leq \epsilon) \geq 1 - \alpha$. Such results have been shown for other sequential procedures [3]. For the class of problems we consider, i.e., problems that satisfy (A1), we are able to show that this probability, in the limit, is one.

By part (ii) of Theorem 3, for a given $\epsilon > 0$, FSP terminates in a finite number of iterations and uses a finite number of samples to generate the approximate solution and to assess its quality, w.p.1. Under more

restrictive conditions, it is also possible to show that the expected values of these random variables are finite. It is well known that $P(K(\epsilon) < \infty) = 1$ does not necessarily imply $EK(\epsilon) < \infty$. To show that the expected values of the sample sizes and expected number of iterations is finite, we confine ourselves to problems where X^* is a singleton, i.e., (SP) has a unique optimal solution, and ArRP is used to form the CI in (8) for some $r \geq 1$. FSP can be run with ArRP by changing the calculation of the gap and variance estimators in step 2 of the procedure. A minor additional requirement we impose on the sample size schedules $\{n_k\}$ when using ArRP for assessing solution quality is that n_k is divisible by r for all $k \geq 1$. The proposition below summarizes the conditions on the schedule of sample sizes $\{m_k\}$ and $\{n_k\}$ under which the expected sample sizes and the expected number of iterations of FSP are finite for a given $\epsilon > 0$.

Proposition 4 *Assume (A2) holds and (SP) has a unique optimal solution, $X^* = \{x^*\}$. Let $\{m_k\}, \{n_k\}$ be positive, nondecreasing sequences such that $m_k \rightarrow \infty$ and $n_k \rightarrow \infty$ as $k \rightarrow \infty$. Let $h(\cdot)$ be a nonnegative, nonincreasing function such that $h(n_k) \rightarrow 0$ as $n_k \rightarrow \infty$, and suppose $k < \infty$ implies that $m_k < \infty, n_k < \infty$, and $h(n_k) > 0$. Consider FSP that uses ArRP, $r \geq 1$, based on SRP given in (2)-(3), to calculate the confidence intervals and stops at iteration $K(\epsilon)$ according to stopping criterion (9), defined in terms of ϵ , using $N(\epsilon) = n_{K(\epsilon)}$ and $M(\epsilon) = m_{K(\epsilon)}$ samples. Given $\epsilon > 0$,*

(i) if

$$\sum_{k=1}^{\infty} e^{-\beta m_k} < \infty \text{ and } \sum_{k=1}^{\infty} e^{-\beta n_k} < \infty \text{ for all } \beta > 0, \quad (11)$$

then, $EK(\epsilon) < \infty$,

(ii) if

$$\sum_{k=1}^{\infty} (n_{k+1} - n_k) e^{-\beta m_k} < \infty \text{ and } \sum_{k=1}^{\infty} (n_{k+1} - n_k) e^{-\beta n_k} < \infty \text{ for all } \beta > 0, \quad (12)$$

then, $EN(\epsilon) < \infty$,

(iii) if

$$\sum_{k=1}^{\infty} (m_{k+1} - m_k) e^{-\beta m_k} < \infty \text{ and } \sum_{k=1}^{\infty} (m_{k+1} - m_k) e^{-\beta n_k} < \infty \text{ for all } \beta > 0,$$

then, $EM(\epsilon) < \infty$.

Proof. (i) Let B_k denote the event $\{G_{n_k}(\hat{x}_k) + v_{n_k}^\alpha(\hat{x}_k) + h(n_k) > \epsilon\}$, and let $\{A\}^c$ denote the complement of event A . Since $K(\epsilon)$ is a nonnegative random variable,

$$EK(\epsilon) = \sum_{k=0}^{\infty} P(K(\epsilon) > k) \leq 1 + \sum_{k=1}^{\infty} P(B_k). \quad (13)$$

Let k_0 be the smallest positive integer such that $h(n_k) \leq \epsilon$ for all $k \geq k_0$. Note that since ArRP is used to calculate the confidence intervals, for $k \geq k_0$,

$$B_k \subseteq \{\hat{x}_k = x_{n_k/r,i}^*, \forall i = 1, \dots, r\}^c \subseteq \left\{ \{x_{m_k}^* \neq x^*\} \cup \left(\bigcup_{i=1}^r \{x_{n_k/r,i}^* \neq x^*\} \right) \right\}. \quad (14)$$

Otherwise, when $\hat{x}_k = x_{n_k/r,i}^*$ for all $i = 1, 2, \dots, r$, $G_{n_k}(\hat{x}_k) = v_{n_k}^\alpha(\hat{x}_k) = 0$ and the condition for event B_k is not satisfied for $k \geq k_0$. Using the right-hand side of (14) and continuing from (13),

$$\begin{aligned} EK(\epsilon) &\leq k_0 + \sum_{k=k_0}^{\infty} P(x_{m_k}^* \neq x^*) + \sum_{i=1}^r \sum_{k=k_0}^{\infty} P(x_{n_k/r,i}^* \neq x^*) \\ &\leq k_0 + C \sum_{k=k_0}^{\infty} e^{-\beta m_k} + \sum_{i=1}^r \left(C \sum_{k=k_0}^{\infty} e^{-\beta n_k/r} \right) \end{aligned} \quad (15)$$

$$\begin{aligned} &= k_0 + C \sum_{k=k_0}^{\infty} e^{-\beta m_k} + rC \sum_{k=k_0}^{\infty} e^{-\beta n_k/r} \\ &< \infty, \end{aligned} \quad (16)$$

where (15) follows from (A2) and (16) follows from (11).

(ii) Let k_0 be defined as above. Then,

$$\begin{aligned} EN(\epsilon) &= \sum_{n=0}^{\infty} P(N(\epsilon) > n) \\ &= n_{k_0} + \sum_{k=k_0}^{\infty} (n_{k+1} - n_k) P(K(\epsilon) > k) \end{aligned} \quad (17)$$

$$\leq n_{k_0} + C \sum_{k=k_0}^{\infty} (n_{k+1} - n_k) e^{-\beta m_k} + rC \sum_{k=k_0}^{\infty} (n_{k+1} - n_k) e^{-\beta n_k/r}, \quad (18)$$

where (17) follows from the fact that $P(N(\epsilon) > n) = 1$ for $0 \leq n \leq n_{k_0} - 1$; $P(N(\epsilon) > n) = P(K(\epsilon) > k_0)$ for $n_{k_0} \leq n \leq n_{k_0+1} - 1$, etc. The right hand side of (18) is finite by (12). Proof of (iii) is analogous. ■

A few remarks about Proposition 4 are now in order. First, note that when m_k and n_k both grow of order $O(\ln^2 k)$, $O(k)$, or higher such as $O(k^2)$, etc., the conditions in part (i) of Proposition 4 are satisfied but not when they both grow of order $O(\ln k)$. To see this, suppose SRP is used in step 2 of FSP and consider the following schedules of sample sizes: $m_k = n_k = \lceil c_0 + c_1 \ln^2 k \rceil$, where $c_0 \geq 2$ and $c_1 > 0$ are two constants and $\lceil \cdot \rceil$ returns the smallest integer greater than or equal to its argument. These guarantee $EK(\epsilon) < \infty$. In contrast, when the schedules are changed to $m_k = n_k = \lceil c_0 + c_1 \ln k \rceil$ the conditions in (11) are satisfied if and only if $c_1 \beta > 1$. Since reliable estimation of β can be a difficult problem, we do not consider such schedules. Therefore, the conditions on the sample size schedules in Proposition 4 are stated for all $\beta > 0$.

As a second remark, conditions in parts (ii) and (iii) of Proposition 4 indicate that there needs to be a “balance” between the growths in sample sizes m_k and n_k . For instance, when m_k is increased too fast relative to n_k , even though FSP stops in a finite number of iterations w.p.1 and $EK(\epsilon) < \infty$, we cannot guarantee $EM(\epsilon) < \infty$. In this case, even though good solutions can be found quickly with a high probability it might take time for the procedure to detect that it is indeed a good solution. This might happen, for instance, when $m_k = k^k$ and $n_k = \lceil 2 + \ln^2 k \rceil$.

4.3 Rate of increase of $M(\epsilon)$ and $N(\epsilon)$

In this section, we establish the rate of increase of $M(\epsilon)$ and $N(\epsilon)$ as $\epsilon \downarrow 0$ for certain schedules of sample sizes. Recall that while $\{m_k\}$ and $\{n_k\}$ are deterministic sequences, $M(\epsilon)$ and $N(\epsilon)$ are the random variables that represent the sample sizes used when FSP stops for a given $\epsilon > 0$. For our analysis, we set the inflation factor $h(n) = 1/\sqrt{n}$ and for simplicity, set $m_k = n_k$, even though m_k could be a multiple of n_k . The class of schedules we consider are defined as a function of the iteration number, k , by:

$$m_k = n_k = \lceil c_0 + c_1 g(k) \rceil, \quad (19)$$

where $c_0, c_1 > 0$ are two constants. We assume $n_1 \geq 2$. The growth function, $g(k)$, is a nonnegative, nondecreasing function of k with the following two properties:

$$\lim_{k \rightarrow \infty} g(k) = \infty \quad \text{and} \quad \lim_{k \rightarrow \infty} (g(k) - g(k-1)) \leq U, \quad (20)$$

where $U \in [0, \infty)$. The first limit property in (20) guarantees that $m_k \rightarrow \infty$ and $n_k \rightarrow \infty$ as $k \rightarrow \infty$, while the second limit property ensures that the increase in sample sizes is finite as $k \rightarrow \infty$. Examples of functions that meet the requirements in (20) include linear and sub-linear functions such as $g(k) = k$ and $g(k) = \ln^2 k$. A quadratic increase, $g(k) = k^2$, in contrast, does not satisfy the second limit condition in (20). In Proposition 6 below, we show that both $N(\epsilon)$ and $M(\epsilon)$ grow at the same rate as $1/\epsilon^2$, w.p.1 for sample size schedules defined according to (19) and (20). Before presenting Proposition 6, we introduce the following lemma.

Lemma 5 *Assume (A1) holds and (SP) has a unique optimal solution, $X^* = \{x^*\}$. Given n and $x \in X$, let $G_n(x)$ and $s_n^2(x)$ be calculated using ArRP, $r \geq 1$ based on SRP given in (2). Furthermore, assume $M(\epsilon)$*

and $N(\epsilon)$ are independent sequences of positive, integer-valued random variables such that $M(\epsilon) \rightarrow \infty$ and $N(\epsilon) \rightarrow \infty$, w.p.1 as $\epsilon \downarrow 0$. Let $\hat{x}_\epsilon \equiv x_{M(\epsilon)}^*$ be obtained by solving a sampling problem with sample size $M(\epsilon)$. Then, w.p.1, there exists $\epsilon_0 > 0$, dependent on the sample path, such that $G_{N(\epsilon)}(\hat{x}_\epsilon) = 0$ and $s_{N(\epsilon)}(\hat{x}_\epsilon) = 0$ for all $\epsilon < \epsilon_0$.

Proof. We will show the result for SRP. The same result for ArRP, $r \geq 2$, follows from the fact that ArRP estimators are formed by averaging r SRP estimators. Consider x_m^* and x_n^* obtained by solving sampling problems (SP $_m$) and (SP $_n$), respectively. As $m \rightarrow \infty$ and $n \rightarrow \infty$, by (A1) and unique optimality, $x_m^* = x^* = x_n^*$, w.p.1 for n and m large enough. The hypotheses of the lemma states $M(\epsilon) \rightarrow \infty$ and $N(\epsilon) \rightarrow \infty$, w.p.1 as $\epsilon \downarrow 0$. Therefore, using a vector version of Proposition 2, $\hat{x}_\epsilon = x_{M(\epsilon)}^* = x^* = x_{N(\epsilon)}^*$, w.p.1 for ϵ small enough. When this happens, $G_{N(\epsilon)}(\hat{x}_\epsilon) = \bar{f}_{N(\epsilon)}(\hat{x}_\epsilon) - \bar{f}_{N(\epsilon)}(x_{N(\epsilon)}^*) = 0$. Hence, $G_{N(\epsilon)}(\hat{x}_\epsilon) = 0$, w.p.1 for ϵ small enough. When $\hat{x}_\epsilon = x_{N(\epsilon)}^*$, $s_{N(\epsilon)}^2(\hat{x}_\epsilon) = 0$. ■

When the class of schedules that satisfy (19) and (20) is used for FSP, it can be shown that $M(\epsilon)$ and $N(\epsilon)$ asymptotically increase at the rate of $1/\epsilon^2$, w.p.1 under a unique optimality assumption. This result is formally expressed in the next proposition.

Proposition 6 Assume (A1) holds and (SP) has a unique optimal solution, $X^* = \{x^*\}$. Let $\{m_k\}$, $\{n_k\}$ be positive, nondecreasing sequences that satisfy (19)-(20) and let $h(n) = 1/\sqrt{n}$. Consider FSP that uses ArRP, $r \geq 1$, based on SRP given in (2)-(3), to calculate the confidence intervals and stops at iteration $K(\epsilon)$ according to stopping criterion (9), defined in terms of ϵ , using $N(\epsilon) = n_{K(\epsilon)}$ and $M(\epsilon) = m_{K(\epsilon)}$ samples. Then,

$$\lim_{\epsilon \downarrow 0} \frac{M(\epsilon)}{1/\epsilon^2} = 1, \quad \text{and} \quad \lim_{\epsilon \downarrow 0} \frac{N(\epsilon)}{1/\epsilon^2} = 1, \quad \text{w.p.1.}$$

Proof. Let $G_{n_k} = G_{n_k}(\hat{x}_k)$, $s_{n_k} = s_{n_k}(\hat{x}_k)$, $t_k = t_{n_k-1, \alpha}$. When FSP stops, (8) is satisfied. Rearranging (8), we have

$$\frac{\left(\sqrt{N(\epsilon)}G_{N(\epsilon)} + t_{K(\epsilon)}s_{N(\epsilon)} + 1\right)^2}{\epsilon^2} \leq N(\epsilon). \quad (21)$$

Note that (8) is not satisfied at iteration $K(\epsilon) - 1$. Rearranging the terms again, adding n_1 to the right-hand side and noting that $n_{K(\epsilon)-1} \geq n_{K(\epsilon)} - c_1 [g(K(\epsilon)) - g(K(\epsilon) - 1)] - 1$, we obtain

$$N(\epsilon) \leq n_1 + 1 + c_1 [g(K(\epsilon)) - g(K(\epsilon) - 1)] + \frac{\left(\sqrt{n_{K(\epsilon)-1}}G_{n_{K(\epsilon)-1}} + t_{K(\epsilon)-1}s_{n_{K(\epsilon)-1}} + 1\right)^2}{\epsilon^2}, \quad (22)$$

where we set $m_{K(\epsilon)-1} = m_1$ and $n_{K(\epsilon)-1} = n_1$ when $K(\epsilon) = 1$. The inequality (22) is satisfied since, if the procedure stops at iteration 1, i.e., $K(\epsilon) = 1$, then, $N(\epsilon) = n_1$, otherwise, it is less than or equal to the sum of the last three terms on the right-hand side of (22). Dividing each term in (21) and (22) by $1/\epsilon^2$ and taking limits as $\epsilon \downarrow 0$, we obtain,

$$1 \leq \lim_{\epsilon \downarrow 0} \frac{N(\epsilon)}{1/\epsilon^2} \leq 1, \quad \text{w.p.1.} \quad (23)$$

In (23), the left-hand side “1” is obtained as a consequence of Lemma 5 ($\lim_{\epsilon \downarrow 0} \sqrt{N(\epsilon)}G_{N(\epsilon)}(\hat{x}_\epsilon) = 0$, w.p.1) and the fact that $t_{K(\epsilon)} \rightarrow z_\alpha$, w.p.1 as $\epsilon \downarrow 0$. The right-hand side “1” is obtained by noting that as $\epsilon \downarrow 0$: (i) $\epsilon^2(n_1 + 1 + c_1 [g(K(\epsilon)) - g(K(\epsilon) - 1)]) \rightarrow 0$, w.p.1 since $n_1 + 1$ is a constant and $\lim_{\epsilon \downarrow 0} c_1 [g(K(\epsilon)) - g(K(\epsilon) - 1)] < \infty$, w.p.1, and (ii) $K(\epsilon) - 1 \rightarrow \infty$, $m_{K(\epsilon)-1} \rightarrow \infty$, and $n_{K(\epsilon)-1} \rightarrow \infty$; hence, the numerator of the last term in (22) tends to 1 by again invoking Lemma 5. By (19), the same result holds for $M(\epsilon)$. ■

Proposition 6 indicates that as we try to obtain higher-quality solutions using FSP, we need to use larger sample sizes with asymptotic increases of $1/\epsilon^2$. Sample size increases of order $O(\epsilon^{-2})$ are quite common in Monte Carlo sampling-based algorithms. For stochastic discrete optimization, for instance, the sample size estimates to obtain ϵ -optimal solutions with a desired probability via sampling-based approximations require

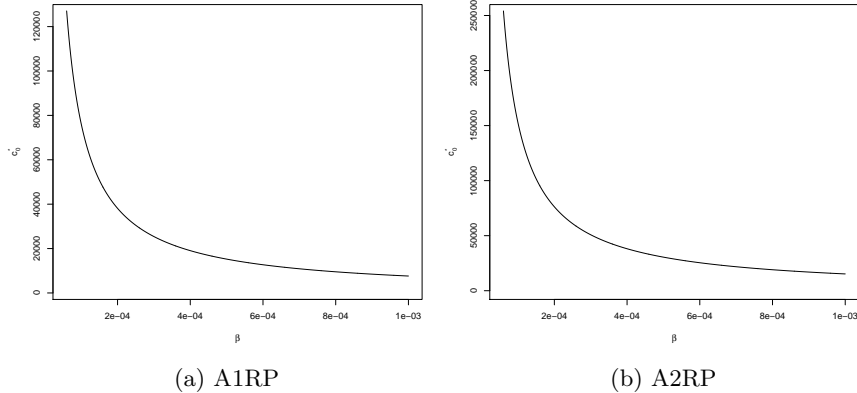


Figure 1: Values of c_0^* vs β for FSP with (a) A1RP and (b) A2RP

similar increases [26]; see also the sample size estimates for a robust stochastic approximation method [34]. Classical sequential procedures in statistics also exhibit similar behavior, see, e.g., [12].

4.4 Parameter Values of Linear Schedules

One class of sample size schedules that satisfy (19)-(20) is linear schedules with $g(k) = k - 1$. Here, we set $g(k) = k - 1$ so that the initial sample size is $n_1 = \lceil c_0 \rceil$. Linear schedules are quite popular in sequential statistics as this is a natural way to increase the number of observations. Such schedules have also been used in stochastic programming. The stochastic decomposition method of Hige and Sen [18], for instance, increases the sample sizes one by one due to computational efficiency. For this class of schedules, Proposition 6 shows, as $\epsilon \downarrow 0$, $N(\epsilon)$ grows of order $O(\epsilon^{-2})$, w.p.1, under certain conditions. However, we know from Theorem 3 that for $\epsilon > 0$, $P(N(\epsilon) < \infty) = 1$; and moreover, $EN(\epsilon) < \infty$ by Proposition 4. We might then try to determine the values of the constants c_0 and c_1 that minimize the expected sample sizes when the procedure stops. Since exact minimization of $EN(\epsilon) = EM(\epsilon)$ may not always be possible, we can instead minimize the upper bound provided in (18). For a given value of r , by setting k_0 in (18) to $k_0 = 1$ and ignoring the integrality requirement, we obtain:

$$EN(\epsilon) \leq u(c_0, c_1) = c_0 + C \frac{c_1}{1 - e^{-\beta c_1}} e^{-\beta c_0} + Cr \frac{c_1}{1 - e^{-\beta c_1/r}} e^{-\beta c_0/r}. \quad (24)$$

The upper bound $u(c_0, c_1)$ in (24) has the following properties:

- $u(c_0, \cdot)$ is an increasing function of c_1 on $(0, \infty)$,
- $u(\cdot, c_1)$ is a convex function on $c_0 > 0$.

When $ArRP$ is used, by the first property, we can set $c_1 = r$. That is, we increase the sample sizes by r observations at each iteration. Then, we can determine the value of c_0 that minimizes $u(\cdot, r)$ by $c_0^* = \{c_0 > 0 : \frac{\partial}{\partial c_0} u(c_0, r) = 0\}$. In the case of FSP with A1RP, $c_0^* = \frac{1}{\beta} \ln(\frac{2\beta C}{1 - e^{-\beta}})$; and in the case of FSP with A2RP, $c_0^* = \frac{2}{\beta} \ln\left(\frac{2A}{-B + \sqrt{B^2 + 4A}}\right)$ where $A = \frac{2\beta C}{1 - e^{-2\beta}}$ and $B = \frac{2\beta C}{1 - e^{-\beta}}$. We can then round up c_0^* to the nearest integer divisible by $r = 1, 2$ to find the initial sample sizes.

Figure 1 shows the values of c_0^* as a function of β for both A1RP and A2RP. Note that c_0^* is a decreasing function of β and tends to be large for small values of β . It is approximately twice as large for A2RP than for A1RP. To see a practical value of c_0^* , we used one of our test problems, denoted 10D (see §6 for details). 10D is a static stochastic knapsack problem with 10 decision variables and 10 independent normally distributed stochastic parameters [26]. Exploiting the properties of the normal distribution, we derived an estimate for β that is $\hat{\beta} = 9.9625 \times 10^{-8}$. We note that 10D is a test problem designed to have a slow rate of convergence, i.e., β for 10D is small. Using $\hat{\beta}$, we obtain $c_0^* = 7.65 \times 10^7$ and $c_0^* = 1.50 \times 10^8$ for A1RP and A2RP, respectively. Unfortunately, since $u(c_0, c_1)$ is based on worst-case bounds, these initial sample size values are too large to be of practical value. They exceed by far the sample sizes utilized by FSP in our computational experiments (see §6).

5 Sequential Procedure with Stochastic Schedules

In this section, we present a variant of FSP that determines the next values of the sample sizes according to the current estimates of the optimality gap and the sampling error. Since the sample sizes are determined according to statistical estimates, they now form stochastic schedules. We denote these stochastic schedules as $\{M_k\}$ and $\{N_k\}$ and denote the sequential procedure with stochastic schedules as SSP. In our subsequent analysis, we set $M_k = N_k$ and $h(n) = 1/\sqrt{n}$ and assume $v_n^\alpha(x) = \frac{t_{n-1,\alpha} s_n(x)}{\sqrt{n}}$. We note that M_k can be set to be a multiple of N_k and $v_n^\alpha(x)$ is in a form satisfied by ArRP, $r \geq 1$.

Just like the FSP, our aim is to stop at the first iteration when the inflated CI falls below ϵ . The stopping iteration of SSP is similarly denoted as

$$K'(\epsilon) = \inf_{k \geq 1} \left\{ k : G_{N_k}(\hat{x}_k) + \frac{t_{N_k-1,\alpha} s_{N_k}(\hat{x}_k) + 1}{\sqrt{N_k}} \leq \epsilon \right\}. \quad (25)$$

When the stopping criteria in (25) is not satisfied, we would like to estimate the next sample size. There could be a couple of reasons why the current confidence interval is large. First, the current candidate solution might have a large optimality gap. Second, the gap estimator might be biased. For instance, it is well-known that ArRP estimators are biased, that is, on average they tend to overestimate the optimality gap for a fixed sample size. Even though we do not exactly know the current optimality gap and the nature of bias for a particular problem, these two are approximately captured in the estimator $G_{N_k}(\hat{x}_k)$. A third reason why the CI might be too large to stop is that the sampling error is large. In this case, we need larger sample sizes to reduce the sampling error to stop at a fixed-width CI. Again, we do not know the variance, hence the exact sampling error, but we can use our current estimate $s_{N_k}^2(\hat{x}_k)$ instead.

Consider the stopping criteria in (25). Multiplying each side by N_k and rearranging,

$$-N_k \epsilon + N_k G_{N_k}(\hat{x}_k) + \sqrt{N_k} (t_{N_k-1,\alpha} s_{N_k}(\hat{x}_k) + 1) \leq 0. \quad (26)$$

One possible way to find the next sample size is to replace N_k that appears in (26) by a generic variable n and solve for the smallest n , keeping the current point estimators $G_{N_k}(\hat{x}_k)$ and $s_{N_k}^2(\hat{x}_k)$ intact, such that (26) is satisfied. Notice that here, we are using the current estimates $G_{N_k}(\hat{x}_k)$ and $s_{N_k}^2(\hat{x}_k)$ to find the next sample size even though these are affected by the sample size as well. As mentioned above, these are our current best indicators of the optimality gap, bias and variance, therefore, we use them in calculating the next sample size. Estimating the next sample size by fixing the point estimators and solving for a value of n that satisfies the stopping criterion is quite common in sequential statistics [12]; this approach can be viewed analogous. The resulting sample size estimate is $N_{k+1} = \inf_{n \in \mathbb{Z}_+} \left\{ n \geq \frac{(t_{n-1,\alpha} s_n(\hat{x}_k) + 1)^2}{(\epsilon - G_{N_k}(\hat{x}_k))^2} \right\}$. There are several dangers in using this sample size estimate. First, when the current optimality gap estimate is much larger than ϵ , which can happen easily at early iterations, the next sample size estimate N_{k+1} can actually be less than the current sample size N_k . This is rather contradictory. When the gap estimate is large, this could be an indication of either a low-quality candidate solution or high bias. In either case, a large optimality gap estimate suggests increasing the sample size in an effort to obtain a higher quality solution and/or reduce bias. On the other hand, when the gap estimate is close to ϵ , it results in a very large sample size increase, which is also somewhat contradictory by the same reasoning.

Using similar ideas, we propose an alternative way to increase the sample sizes. Let $b_k = t_{N_k-1,\alpha} s_{N_k}(\hat{x}_k) + 1$ and $c_k = N_k G_{N_k}(\hat{x}_k)$. Starting from (26), fixing b_k and c_k and replacing N_k by n , we can rewrite the stopping criterion as a quadratic expression in the square root of the sample size n as follows

$$-\epsilon n + b_k \sqrt{n} + c_k \leq 0. \quad (27)$$

When this stopping rule is not satisfied at iteration k , the sample size N_{k+1} for the next iteration can be estimated as the smallest integer value for n such that (27) holds. To find such a value, consider the quadratic polynomial in \sqrt{n} given in (27). Note that $b_k > 0$ and $c_k \geq 0$. The corresponding discriminant to the quadratic polynomial can be computed as $\Delta_k = b_k^2 + 4\epsilon c_k \geq b_k^2 > 0$. Hence, the quadratic polynomial in (27) has two roots: $\nu_- = \frac{-b_k + \sqrt{\Delta_k}}{-2\epsilon} \leq 0$ and $\nu_+ = \frac{b_k + \sqrt{\Delta_k}}{2\epsilon} > 0$. Thus, we can select $N_{k+1} = \lceil \nu_+^2 \rceil$. This choice of N_{k+1} guarantees that $N_{k+1} > N_k$. Note that with $n = N_k$, the left-hand-side of (27) is

positive (assuming at iteration k we did not yet stop) and $-\epsilon < 0$, so, as a consequence of the properties of a quadratic polynomial with two real roots, $N_{k+1} > N_k$. When ArRP, $r \geq 1$ is used, the value of N_{k+1} can be set to the smallest integer greater than or equal to $\lceil \nu_+^2 \rceil$ that is divisible by r .

This is a somewhat heuristic way to increase the sample sizes but it has the right implication. For instance, all else equal, if the optimality gap estimate is larger, then the sample size increases will be larger. Similarly, if the variance estimate is larger, the sample size increases will be larger as well. We note that this might not be the only way to estimate the next sample sizes. For instance, more elaborate adaptive sequential procedures can be designed aiming to correct bias and/or variance but we do not pursue it in this paper and leave it for future work. Nevertheless, with these sample size increases we are able to show similar results as FSP but under more restrictive conditions. Below, we first provide a brief statement of the sequential procedure with stochastic schedules (SSP). Then, we prove and discuss several theoretical properties.

SSP:

Input: Initial sample sizes $m_0 = n_0$, values of $\epsilon > 0$, $0 < \alpha < 1$, and resampling frequencies k_f^n and k_f^m .

Output: A candidate solution $\hat{x}_{K'(\epsilon)}$ with an estimated optimality gap of ϵ .

0. (Initialization) Set $k = 1$, $M_1 = N_1 = n_0$, generate M_1 observations $\tilde{\xi}_1^1, \tilde{\xi}_1^2, \dots, \tilde{\xi}_1^{M_1}$ and independently generate N_1 observations $\tilde{\xi}_2^1, \tilde{\xi}_2^2, \dots, \tilde{\xi}_2^{N_1}$.

1. (Find candidate solution) Using the observations $\tilde{\xi}_1^1, \dots, \tilde{\xi}_1^{M_k}$, solve (SP_{M_k}) to obtain $x_{M_k}^*$. Set $\hat{x}_k = x_{M_k}^*$.

2. (Assess solution quality) Using the observations $\tilde{\xi}_2^1, \dots, \tilde{\xi}_2^{N_k}$, calculate $G_{N_k}(\hat{x}_k)$, $s_{N_k}^2(\hat{x}_k)$ and $v_{N_k}^\alpha(\hat{x}_k)$.

3. (Check stopping criterion) If $G_{N_k}(\hat{x}_k) + v_{N_k}^\alpha(\hat{x}_k) + 1/\sqrt{N_k} \leq \epsilon$, then set $K'(\epsilon) = k$, $N'(\epsilon) = N_k$, $M'(\epsilon) = M_k$, and $\hat{x}_{K'(\epsilon)} = \hat{x}_k$ and stop. Output candidate solution $\hat{x}_{K'(\epsilon)}$ with confidence interval on $\mu_{\hat{x}_{K'(\epsilon)}}$ as $[0, \epsilon]$.

4. (Increase sample size)

4.1. Solve the following quadratic equation in terms of \sqrt{n} :

$$\frac{1}{n} (N_k G_{N_k}(\hat{x}_k)) + \frac{1}{\sqrt{n}} (t_{N_k-1, \alpha} s_{N_k}(\hat{x}_k) + 1) = \epsilon.$$

4.2. The above quadratic equation has two roots, $\nu_+ > 0$ and $\nu_- \leq 0$. Thus, there is a unique solution of $\sqrt{n} = \nu_+$. Consequently, N_{k+1} is chosen as $N_{k+1} = \lceil \nu_+^2 \rceil$.

4.3. Set $M_{k+1} = N_{k+1}$. If k_f^m divides $k + 1$ then sample $\tilde{\xi}_1^1, \tilde{\xi}_1^2, \dots, \tilde{\xi}_1^{M_{k+1}}$ independently of samples generated in previous iterations. Else, sample $M_{k+1} - M_k$ additional observations, $\tilde{\xi}_1^{M_k+1}, \dots, \tilde{\xi}_1^{M_{k+1}}$. Similarly, if k_f^n divides $k + 1$ then sample $\tilde{\xi}_2^1, \tilde{\xi}_2^2, \dots, \tilde{\xi}_2^{N_{k+1}}$ independently of samples generated in previous iterations. Else, sample $N_{k+1} - N_k$ additional observations, $\tilde{\xi}_2^{N_k+1}, \dots, \tilde{\xi}_2^{N_{k+1}}$. Set $k = k + 1$ and go to 1.

Like FSP, the sample sizes when the procedure stops are denoted by $N'(\epsilon) = N_{K'(\epsilon)}$ and $M'(\epsilon) = m_{K'(\epsilon)}$. In SSP, sample sizes when the procedure stops can be larger than FSP, but when fewer optimization problems are solved, solution times can be accelerated. The overall solution time, of course, depends on the size of the optimization problems solved as well as the number. In early iterations, the statistical estimators may be poor and/or the candidate solution may be of low-quality, leading to large values of the next sample sizes. In this case, larger optimization problems need to be solved and benefits from numerical speedup techniques such as warm-starting may be lost. The deterministic schedules in FSP can be increased slowly, such as of order $O(\ln^2 k)$. This can be desirable if, for instance, obtaining samples is costly, the solution procedure allows for quick updates (e.g., via warm-starting), or solving a problem with a large sample size is computationally burdensome. In contrast, in SSP, the sample sizes are increased in random jumps given the current gap and variance estimates. The larger the estimates, the larger the increases in sample sizes.

We now provide equivalent results for SSP as the ones stated in Theorem 3 for FSP. Namely, SSP finds an ϵ -optimal solution with probability one as $\epsilon \downarrow 0$ and for $\epsilon > 0$ it stops in a finite number of iterations using finite number of samples, w.p.1.

Theorem 7 Assume (A1) holds.

(i) Consider SSP that stops at iteration $K'(\epsilon)$ according to stopping criterion (25), defined in terms of ϵ , using $N'(\epsilon) = N_{K'(\epsilon)}$ and $M'(\epsilon) = m_{K'(\epsilon)}$ samples. Then,

$$\lim_{\epsilon \downarrow 0} P\left(\mu_{\hat{x}_{K'(\epsilon)}} \leq \epsilon\right) = 1.$$

(ii) In addition, assume (SP) has a unique optimum solution, $X^* = \{x^*\}$ and that SSP uses ArRP, $r \geq 1$, based on SRP given in (2)-(3), to calculate the confidence intervals. Given $\epsilon > 0$,

$$P(K'(\epsilon) < \infty) = P(N'(\epsilon) < \infty) = P(M'(\epsilon) < \infty) = 1.$$

Proof. (i) Let $Y_m = \mu_{x_m^*} = E[f(x_m^*, \tilde{\xi}) | x_m^*] - z^*$ for some $m \rightarrow \infty$. By (A1), $Y_m = 0$, w.p.1 for m large enough. Similarly define $Y_{M'(\epsilon)} = \mu_{x_{M'(\epsilon)}^*} = E[f(x_{M'(\epsilon)}^*, \tilde{\xi}) | x_{M'(\epsilon)}^*] - z^*$ and note that $\hat{x}_{K'(\epsilon)} = x_{M'(\epsilon)}^*$. For SSP, $M'(\epsilon) \rightarrow \infty$, w.p.1 as $\epsilon \downarrow 0$. Therefore, using Proposition 2, we obtain that $Y_{M'(\epsilon)} = 0$, w.p.1 for ϵ small enough. The desired result follows.

(ii) Let $G_{N_k} = G_{N_k}(\hat{x}_k)$, $s_{N_k} = s_{N_k}(\hat{x}_k)$, $t_k = t_{N_k-1, \alpha}$ and $\Delta_k = (t_k s_{N_k} + 1)^2 + 4\epsilon N_k G_{N_k}$. Note that

$$P(K'(\epsilon) = \infty) \leq P\left(\bigcap_{k=1}^{\infty} \{N_{k+1} > N_k\}\right) \quad (28)$$

$$= P\left(\bigcap_{k=1}^{\infty} \left\{\left\lceil \left(\frac{(t_k s_{N_k} + 1) + \sqrt{\Delta_k}}{2\epsilon}\right)^2 \right\rceil > N_k\right\}\right). \quad (29)$$

The right-hand side of (29) is positive only if $N_k G_{N_k} \rightarrow \infty$ and $s_{N_k} \rightarrow \infty$ with positive probability as $k \rightarrow \infty$. Let $A = \{\omega : K'(\epsilon) = \infty\}$. Suppose $P(A) > 0$. Conditioned on A , $M_k = N_k \rightarrow \infty$, w.p.1 as $k \rightarrow \infty$. For every $k = 1, 2, 3, \dots$ construct a decreasing sequence of $\bar{\epsilon}$ such that $k = \lceil \frac{1}{\bar{\epsilon}} \rceil$. Set $M_k = M(\bar{\epsilon})$ and $N_k = N(\bar{\epsilon})$. Then, conditioned on A , $M(\bar{\epsilon}) = N(\bar{\epsilon}) \rightarrow \infty$, w.p.1 as $\bar{\epsilon} \downarrow 0$. Therefore, by Lemma 5, conditioned on A , $N_k G_{N_k} \rightarrow 0$ and $s_{N_k} \rightarrow 0$, w.p.1 as $k \rightarrow \infty$. This implies that right-hand side of (29) is 0, which is a contradiction. This shows $P(K'(\epsilon) < \infty) = 1$. Now, $P(N'(\epsilon) = \infty)$ is also less than or equal to the right-hand side of (28), hence is 0. By noting that $M'(\epsilon) = N'(\epsilon)$, we have $P(M'(\epsilon) < \infty) = 1$. ■

Conditions in part (ii) of Theorem 7 differ from those of Theorem 3. In order to ensure finiteness of the number of iterations and the sample sizes used when SSP stops, we assumed unique optimality and focused on using ArRP for assessing solution quality. Under these conditions, we can further show equivalent results for SSP like those of FSP. We do this in the next two propositions. First, we show that for $\epsilon > 0$, SSP stops in a finite expected number of iterations, i.e., $EK'(\epsilon) < \infty$. Then, by a slight modification in the initial sample sizes, we obtain asymptotic sample size increases for SSP similar to FSP. That is, asymptotically $N(\epsilon) = M(\epsilon)$ grows of order $O(\epsilon^{-2})$, w.p.1.

Proposition 8 Assume (A2) holds and (SP) has a unique optimal solution, $X^* = \{x^*\}$. Consider SSP that uses ArRP, $r \geq 1$, based on SRP given in (2)-(3), to calculate the confidence intervals and stops at iteration $K'(\epsilon)$ according to stopping criterion (25), defined in terms of ϵ , using $N'(\epsilon) = N_{K'(\epsilon)}$ and $M'(\epsilon) = m_{K'(\epsilon)}$ samples. Given $\epsilon > 0$, $EK'(\epsilon) < \infty$.

Proof. Let B_{N_k} denote the event $\{G_{N_k}(\hat{x}_k) + v_{N_k}^\alpha(\hat{x}_k) + h(N_k) > \epsilon\}$. Then,

$$\begin{aligned} EK'(\epsilon) &= 1 + \sum_{k=1}^{\infty} P(K'(\epsilon) > k) \\ &\leq 1 + P(B_{N_1}) + \sum_{k=2}^{\infty} P(B_{N_k}). \end{aligned} \quad (30)$$

Note that N_k is an integer-valued random variable that takes on values $N_k \geq n_0 + k - 1$ for $k \geq 2$ and \hat{x}_k is obtained by solving an independent sampling problem with sample size $M_k = N_k$. Using arguments similar

to (14) and continuing from (30),

$$\begin{aligned} EK'(\epsilon) &\leq 2 + \sum_{k=2}^{\infty} \sum_{N_k=n_0+k-1}^{\infty} P(B_{N_k}) \\ &\leq 2 + C(r+1)e^{-\beta/r(n_0-1)} \sum_{k=2}^{\infty} \sum_{j=0}^{\infty} e^{-\beta/r(k+j)}, \end{aligned} \quad (31)$$

where (31) follows from (A2). The right hand side of (31) is finite for $\beta > 0$. ■

We now modify SSP so that the initial sample sizes are selected according to

$$m_0 = n_0 = \left\lceil \max \left\{ \bar{n}_0, \ln \frac{1}{\epsilon} \right\} \right\rceil, \quad (32)$$

where $\bar{n}_0 \geq 2$ is a finite constant (for instance, one can set $\bar{n}_0 = 50$). The initial sample sizes (32) are equal to \bar{n}_0 for all practical purposes but tend to infinity as $\epsilon \downarrow 0$. For example, with $\bar{n}_0 = 50$, values of $\epsilon < e^{-50}$ will result in $n_0 > 50$. By selecting the sample sizes as in (32), one is willing to use more samples as more precision is demanded. With the initial sample sizes defined as above, it is possible to show that the sample sizes used by SSP, $N'(\epsilon) = M'(\epsilon)$, asymptotically increase at the same rate as $1/\epsilon^2$, w.p.1.

Proposition 9 *Suppose (A1) holds and (SP) has a unique optimal solution, $X^* = \{x^*\}$. Let $m_0 = n_0$ be chosen according to (32). Consider SSP that uses ArRP, $r \geq 1$, based on SRP given in (2)-(3), to calculate the confidence intervals and stops at iteration $K'(\epsilon)$ according to stopping criterion (25), defined in terms of ϵ , using $N'(\epsilon) = N_{K'(\epsilon)}$ and $M'(\epsilon) = m_{K'(\epsilon)}$ samples. Then,*

$$\lim_{\epsilon \downarrow 0} \frac{M'(\epsilon)}{1/\epsilon^2} = 1, \quad \text{and} \quad \lim_{\epsilon \downarrow 0} \frac{N'(\epsilon)}{1/\epsilon^2} = 1, \quad \text{w.p.1.}$$

Proof. Let $G_{N_k} = G_{N_k}(\hat{x}_k)$, $s_{N_k} = s_{N_k}(\hat{x}_k)$, $t_k = t_{N_k-1, \alpha}$ and let $\Delta_k = (t_k s_{N_k} + 1)^2 + 4\epsilon N_k G_{N_k}$. When SSP stops, (25) is satisfied. Rearranging (25), we have

$$\frac{\left(\sqrt{N'(\epsilon)} G_{N'(\epsilon)} + t_{K'(\epsilon)} s_{N'(\epsilon)} + 1 \right)^2}{\epsilon^2} \leq N'(\epsilon). \quad (33)$$

By the way the sample sizes are increased, we also have

$$N'(\epsilon) \leq n_0 + r + \frac{\left(t_{K'(\epsilon)-1} s_{N_{K'(\epsilon)-1}} + 1 + \sqrt{\Delta_{K'(\epsilon)-1}} \right)^2}{4\epsilon^2}, \quad (34)$$

where we set $M_{K'(\epsilon)-1} = N_{K'(\epsilon)-1} = n_0$ when $K'(\epsilon) = 1$. Dividing each term in (33) and (34) by $1/\epsilon^2$, taking limits as $\epsilon \downarrow 0$ and following the same line of logic as in the proof of Proposition 6, we obtain

$$1 \leq \lim_{\epsilon \downarrow 0} \frac{N'(\epsilon)}{1/\epsilon^2} \leq 1, \quad \text{w.p.1.} \quad (35)$$

As $N'(\epsilon) = M'(\epsilon)$, the same result holds for $M'(\epsilon)$. ■

The technical reason for selecting n_0 according to (32) becomes apparent in the proof of Proposition 9. This hinges on the fact that unlike FSP, as $\epsilon \downarrow 0$, even though $M'(\epsilon) = N'(\epsilon) \rightarrow \infty$, w.p.1, the same is not true of $K'(\epsilon)$. With $m_0 = n_0$ fixed (not depending on ϵ), as $\epsilon \downarrow 0$, $K'(\epsilon)$ can be 2 with positive probability. In this case, $G_{N_{K'(\epsilon)-1}}(\hat{x}_{K'(\epsilon)-1}) = G_{n_0}(x_{m_0}^*)$ and $s_{N_{K'(\epsilon)-1}}^2(\hat{x}_{K'(\epsilon)-1}) = s_{n_0}^2(x_{m_0}^*)$ are random variables that do not depend on ϵ . Hence, the right-hand limit in (35) cannot be guaranteed to be 1, w.p.1. However, selecting the sample sizes according to (32) ensures that n_0 tends to infinity as $\epsilon \downarrow 0$.

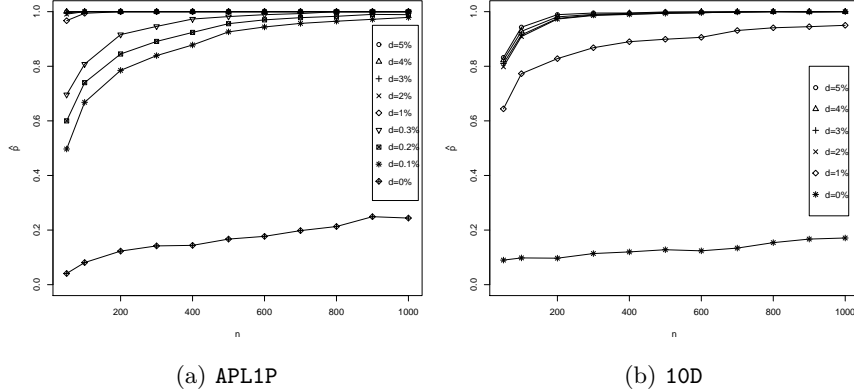


Figure 2: Probability of obtaining ϵ -optimal ($\epsilon = d \cdot z^*$) solutions for (a) APL1P and (b) 10D

6 Computational Results

In this section, we empirically examine the small sample behavior of the procedures. We use two test problems, APL1P and 10D, whose solutions are known, to study the full extent of the performance of our procedures and to examine different parameter settings. Then, we use a large scale problem, 20TERM, whose solution is unknown, to illustrate the application of the procedures in practice. APL1P, a two-stage stochastic linear program with recourse, is an electric power generation model with 5 independent stochastic parameters and 1280 scenarios [22]. 10D is a static stochastic knapsack problem with 10 binary decision variables and 10 independent stochastic parameters [26]. Since APL1P has a small number of realizations, and 10D has normally distributed random variables, we can calculate the actual optimality gaps to verify our procedures. Both problems are known to be challenging for solution via a sampling approximation. That is, their sampling approximations are known to have a low rate of convergence (β in (A2) is small). Both problems satisfy assumptions (A1) and (A2) under i.i.d. sampling and have unique optimal solutions. Our third test problem, 20TERM, is a motor freight scheduling problem modeled as a two-stage stochastic linear program with 40 stochastic parameters and 1.1×10^{12} scenarios [30]. Due to its large scale, its optimal solution is not known. Therefore, we use this problem to illustrate the application of our procedures in practice. Note that this problem also satisfies assumptions (A1) and (A2) under i.i.d. sampling.

We begin in §6.1 by investigating the characteristics of the two test problems APL1P and 10D. Then, in §6.2, we analyze the effect of resampling to determine the values of k_f^m and k_f^n . In §6.3, we apply FSP and SSP to APL1P and 10D and compare the two methods; and in §6.4, we examine the behavior of the procedures as $\epsilon \downarrow 0$ on these two test problems. Next, in §6.5, we report results on 20TERM, providing guidelines on how to choose ϵ . Finally, we summarize our findings and present implementation guidelines in §6.6.

6.1 APL1P and 10D

Before we present results of our computational experiments with FSP and SSP, we first examine the two test problems in more detail to gain insight. Figure 2 shows estimates of probability of obtaining ϵ -optimal solutions for APL1P and 10D. That is, it shows estimated values of $P(\mu_{x_n^*} \leq \epsilon)$, denoted \hat{p} , for various values of n . To obtain these estimates we solved 1,000 independent sampling problems (SP_n) for each value of n and calculated the frequency of the ϵ -optimal solutions obtained out of these 1,000 sampling problems. We set ϵ to different percentages of the optimal objective function value, $\epsilon = d \cdot z^*$, with $d = 0\%, \dots, 5\%$. The estimated probabilities in Figure 2 have 90% confidence interval half-widths of at most 0.01. It appears that for both problems, the probability of obtaining optimal solutions is indeed quite low, even for large sample sizes. However, as ϵ increases, the probability of obtaining ϵ -optimal solutions dramatically increases. This indicates that the fixed-width stopping rules are promising when used with moderate values of ϵ .

6.2 Effect of Resampling

We now examine the effect of augmenting the set of observations with a few additional ones versus generating an entire set of new observations through resampling. Note that augmentation leads to dependent estimates

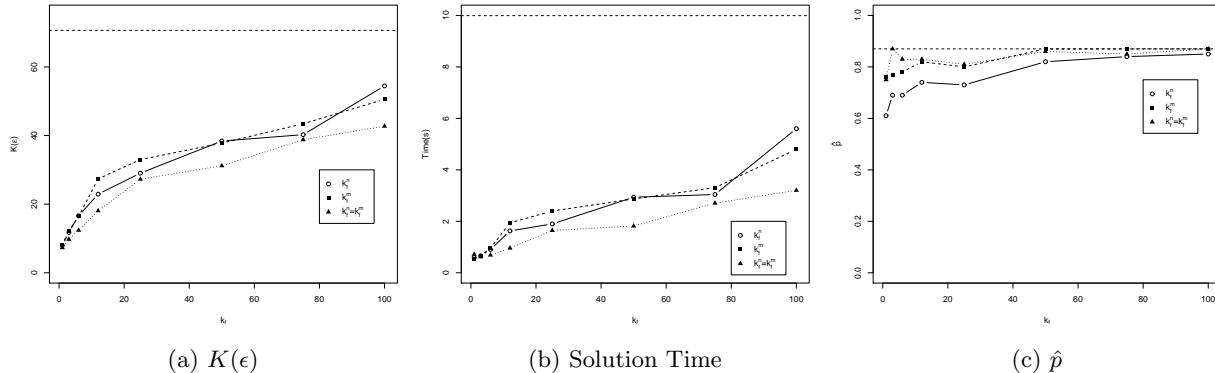


Figure 3: Effect of resampling

and can save computation time when warm-starting is used. Resampling, on the other hand, leads to independent estimates and can prevent getting stuck in a bad sample path. The frequency of resampling is controlled by the parameters k_f^m and k_f^n for the observations used for generating candidate solutions and for assessing solution quality, respectively. It has been observed for another stopping rule that as the resampling frequency is increased (i.e., k_f^n is decreased), the sequential sampling procedure stops on average in fewer iterations but this can lower the coverage probability below desired levels [3]. Our goal in this section is to find the values of k_f^m and k_f^n that result in minimal solution time and stopping iteration, and consequently sample sizes used, without sacrificing coverage probability.

To gain an understanding of the effect of resampling, we applied FSP using SRP≡A1RP with sample size schedules $m_k = n_k = 50 + 2(k - 1)$ on APL1P. We set $\epsilon = 0.2\%z^*$, $h(n) = 1/\sqrt{n}$, and $\alpha = 0.10$. We then tested three resampling schemes: (i) we set $k_f^m = \infty$ (always augment observations used for assessing solution quality) and only let k_f^n vary, (ii) we set $k_f^n = \infty$ (always augment observations used for generating candidate solutions) and only let k_f^m vary, and (iii) we set both resampling frequencies to be the same, $k_f^n = k_f^m$ and let them vary. In each of the three settings, we used resampling frequencies of 1 (resample an independent set of observations at every iteration), 3, 6, 12, 25, 50, 75, 100 (resample an independent set of observations every 100 iterations), and ∞ (always augment). We ran the procedures for each configuration 100 times and report the averages. We also conducted these experiments with 10D setting $\epsilon = 2\%z^*$ and for both problems using A2RP. The results were similar except that A2RP has higher coverage probabilities. For brevity, we only report our findings for APL1P using SRP.

Figure 3 shows the results of our experiments. We denote scheme (i) as k_f^n , scheme (ii) as k_f^m and scheme (iii) as $k_f^n = k_f^m$. From left to right, the first figure shows changes in the stopping iteration $K(\epsilon)$, the second figure shows the solution time of the procedure (in seconds) and the third shows changes in estimated coverage probability, $P(\mu_{\hat{x}_{K(\epsilon)}} \leq \epsilon)$, as k_f^n and k_f^m vary. The dotted horizontal line in each graph denotes the values for the configuration $k_f^n = k_f^m = \infty$. In all three resampling schemes, the more often we resample, the sooner the procedures stop. Performances of the three schemes with respect to $K(\epsilon)$ and running time are closest when values of k_f^n , k_f^m are small (≤ 6). As k_f^n , k_f^m increase, differences in this respect become more apparent. The $k_f^n = k_f^m$ scheme tends to stop earlier except for $k_f^n = k_f^m = \infty$. (This behavior is observed more consistently for APL1P than for 10D.) When $k_f^n = k_f^m$ scheme stops early, however, this does not seem to have an adverse effect on coverage probability.

Schemes k_f^m and $k_f^n = k_f^m$ result in similar coverage probabilities whereas coverage probability of k_f^n is consistently lower. We note that the coverage probabilities of the three schemes converge to the same value as k_f^n , k_f^m increase. We believe the differences in Figure 3c can be explained as follows. If the procedure starts with a *bad solution* (a solution that has an optimality gap higher than ϵ), when no resampling is employed, the procedure will tend to run longer until a *good solution* (an ϵ -optimal solution) is obtained due to the stopping rules. There is, of course, some noise in this process. When resampling scheme k_f^m is introduced, this allows the process to jump from a sample that generates bad solutions to a one that generates good solutions. This way, the procedure can stop earlier at a good solution. In contrast, under resampling scheme k_f^n , only the sample used to assess candidate solutions is regenerated. So, when the procedure starts with a bad solution, disrupting assessing solution quality causes early stopping with a bad solution.

	<i>G</i>	<i>B</i>	<i>Total</i>		<i>G</i>	<i>B</i>	<i>Total</i>		<i>G</i>	<i>B</i>	<i>Total</i>
<i>G</i>	55	6	61	<i>G</i>	41	35	76	<i>G</i>	45	30	75
<i>B</i>	1	38	39	<i>B</i>	15	9	24	<i>B</i>	11	14	25
<i>Total</i>	56	44	100	<i>Total</i>	56	44	100	<i>Total</i>	56	44	100

(a) $k_f^n = 1$ (b) $k_f^m = 1$ (c) $k_f^m = k_f^n = 1$

Table 1: Breakdown of solutions at the first iteration (columns) and when the procedure stops (rows) for APL1P under three resampling schemes

To verify our hypothesis and gain more insight, we examined the three schemes in more detail. In Table 1, we provide a breakdown of the quality of candidate solutions for all 3 resampling schemes at $k_f^n = 1$, $k_f^m = 1$ and $k_f^m = k_f^n = 1$ for APL1P. We denote ϵ -optimal solutions as *G* and all other solutions as *B*. In Table 1, the columns provide a breakdown of the quality of candidate solutions at the first iteration and the rows provide a breakdown of the quality of the candidate solutions when the procedure stops. Of the 100 runs made, 56 start with a good solution and 44 start with a bad solution. The solution breakdown at the start is the same for all the configurations since the same random number stream is used in all three cases. When the procedure stops, however, k_f^m and $k_f^m = k_f^n$ end up with 76 and 75 good solutions respectively, while k_f^n finished with only 61 good solutions. Let us focus on the case where the procedure starts with a bad solution (columns *B*). In this case, resampling the observations that generate the candidate solutions acts as a corrective action that drives the procedure away from bad solutions and towards good solutions. Such is the case for k_f^m where 35 of the 44 runs that started with a bad solution ended up with a good solution. For k_f^n , however, only 6 of the 44 runs starting with bad solutions ended up with good solutions. Results for $k_f^m = k_f^n$ are similar to k_f^m where 30 of 44 runs that started with a bad solution ended up with a good solution. When the procedure starts with a good solution (columns *G*), k_f^n scheme once again adjusts itself to the candidate solution and 55 out of 56 runs starting with good solutions stopped at a good solution. Resampling scheme k_f^m has a slightly adverse effect, 15 out of 56 runs starting with good solutions ended up at bad solutions. This adverse effect is, however, not significant enough to negatively affect coverage probability and is mitigated by the substantial gain in good solutions made from the runs starting with bad solutions. Scheme $k_f^m = k_f^n$ in this case is somewhere in between, and closer to k_f^m , with 45 out of 56 good starting solutions ending with good solutions.

We want to minimize computational effort while obtaining a coverage probability of at least $1 - \alpha$. Such conditions are satisfied for resampling scheme $k_f^m = k_f^n = 3$. In the rest of the paper, we use this value.

6.3 Results on APL1P and 10D

In this section, we compare the performance of the sequential procedures on APL1P and 10D. The results on 20TERM are presented in §6.5. We applied FSP using SRP with sample size schedules $n_k = m_k = n_0 + 2(k - 1)$, and using A2RP with sample size schedules $n_k = m_k = n_0 + 2(k - 1)$ and $n_k = m_k = n_0 + 100(k - 1)$ for various values of the initial sample size n_0 . We applied SSP using A2RP. Other parameter settings are as described in §6.2. In the subsequent graphs and tables, the fully sequential procedure with SRP and A2RP is denoted by F:SRP(2), F:A2RP(2), and F:A2RP(100), where the number in parentheses distinguishes between the cases where the sample size is increased two by two and one hundred by one hundred. The procedure with stochastic schedule of sample sizes using A2RP is denoted by S:A2RP. We ran each procedure 100 times for each initial sample size value n_0 .

Figure 4 shows the empirical coverage probabilities obtained as a result of the computational experiments for APL1P and 10D. The “0.2%” in Figure 4a and “2%” in Figure 4b show the probability of obtaining ϵ -optimal solutions with sample size n_0 for APL1P and 10D, respectively. These are the values from Figure 2. All procedures seem to have good coverage probabilities, except for F:SRP(2) for small initial sample sizes. SSP consistently finds ϵ -optimal solutions with high probability.

Tables 2 and 3 list the sample sizes when the procedures terminate along with the stopping iteration and the solution time (in seconds) for APL1P and 10D, respectively. All values are listed along with 90% confidence interval half-widths around their estimated means. The fully sequential procedure with SRP uses fewer sample sizes than all the other procedures while the sequential procedure with stochastic schedules uses the most. There is a considerable difference between the sample sizes used by FSP and SSP, especially at small initial sample sizes, n_0 . For APL1P, as the initial sample size increases, this difference decreases. This




Figure 4: Coverage probabilities of the sequential procedures for (a) APL1P and (b) 10D

Performance Measure	n_0	F:SRP(2)	F:A2RP(2)	F:A2RP(100)	S:A2RP
Sample sizes used: $M(\epsilon) = N(\epsilon)$ for FSP $M'(\epsilon) = N'(\epsilon)$ for SSP	50	67.48 \pm 2.58	106.26 \pm 6.90	409.00 \pm 34.37	3704.68 \pm 1018.42
	100	111.26 \pm 2.19	138.76 \pm 5.64	383.00 \pm 34.21	1904.78 \pm 521.12
	200	206.22 \pm 1.56	214.90 \pm 2.65	459.00 \pm 35.70	1434.06 \pm 264.63
	400	402.48 \pm 0.63	406.20 \pm 1.41	535.00 \pm 26.59	1108.10 \pm 187.94
	800	800.50 \pm 0.25	801.86 \pm 0.54	855.00 \pm 14.48	1045.44 \pm 82.43
	1000	1000.30 \pm 0.22	1000.38 \pm 0.23	1015.00 \pm 8.55	1082.52 \pm 55.61
Stopping iterations: $K(\epsilon)$ for FSP $K'(\epsilon)$ for SSP	50	9.74 \pm 1.29	29.13 \pm 3.45	4.59 \pm 0.34	2.27 \pm 0.11
	100	6.63 \pm 1.09	20.38 \pm 2.82	3.83 \pm 0.34	2.19 \pm 0.10
	200	4.11 \pm 0.78	8.45 \pm 1.32	3.59 \pm 0.36	2.04 \pm 0.12
	400	2.24 \pm 0.31	4.10 \pm 0.71	2.35 \pm 0.27	1.73 \pm 0.12
	800	1.25 \pm 0.13	1.93 \pm 0.27	1.55 \pm 0.14	1.44 \pm 0.12
	1000	1.15 \pm 0.11	1.19 \pm 0.12	1.15 \pm 0.09	1.10 \pm 0.06
Solution time (seconds)	50	0.62 \pm 0.09	2.45 \pm 0.40	1.69 \pm 0.27	12.00 \pm 4.36
	100	0.73 \pm 0.12	2.33 \pm 0.37	1.45 \pm 0.24	4.34 \pm 1.22
	200	0.94 \pm 0.15	1.74 \pm 0.26	1.83 \pm 0.26	3.20 \pm 0.67
	400	1.22 \pm 0.13	1.85 \pm 0.27	1.72 \pm 0.24	2.45 \pm 0.47
	800	1.65 \pm 0.14	2.14 \pm 0.22	2.07 \pm 0.20	2.28 \pm 0.27
	1000	1.85 \pm 0.15	1.92 \pm 0.14	1.92 \pm 0.15	1.96 \pm 0.16

Table 2: Summary of results for APL1P for select initial sample sizes n_0

is because at higher sample sizes, SSP stops quicker (see results for $K(\epsilon)$ and $K'(\epsilon)$ in Table 2), and also the sample size increases are not that large since the estimators are more reliable. For 10D, the estimators have a large bias; even at larger initial sample sizes, bias remains. This results in larger jumps in SSP and the sample sizes used by this procedure remain larger compared to the fully sequential procedures at all levels of initial sample sizes, n_0 .

To solve APL1P, we used the regularized decomposition code of Ruszczyński and Świanowski [40, 41] and modified this code to allow for warm-starting when new samples are augmented. We note that warm-starting is more effective when sample size increases are small. This, coupled with the fact that F:SRP(2) stops in fewer iterations, results in fast solution times for F:SRP(2) for APL1P. F:A2RP(2) also utilizes warm-starting effectively, but for small values of n_0 it tends to run for a long time. This is because A2RP estimators are more conservative (and hence more reliable for problems with undercoverage) than SRP estimators [2, 3]. F:A2RP(100), on the other hand, runs for a fewer number of iterations, especially for small values of n_0 . Even though larger sampling problems are solved, since a fewer number of problems are solved, solution times for F:A2RP(100) are smaller than those of F:A2RP(2). In SSP, the increases in sample sizes are much higher and the effect of warm-starting vanishes. Therefore, for APL1P, solution times of S:A2RP are generally higher than those of F:A2RP(100). As the initial sample size increases, solution times of all procedures become similar.

We solved 10D using a dynamic programming approach with no warm-starting. Once again, the differences in solution times of F:SRP(2) and F:A2RP(2) are explained by differences in $K(\epsilon)$. F:A2RP(100) improves on F:A2RP(2) by substantially decreasing $K(\epsilon)$. Even though larger problems are solved per it-

Performance Measure	n_0	F:SRP(2)	F:A2RP(2)	F:A2RP(100)	S:A2RP
Sample sizes used: $M(\epsilon) = N(\epsilon)$ for FSP $M'(\epsilon) = N'(\epsilon)$ for SSP	50	150.50 \pm 14.74	861.78 \pm 57.90	2668.00 \pm 149.02	6320.86 \pm 512.06
	100	220.42 \pm 20.02	837.42 \pm 57.06	2707.00 \pm 174.80	6277.70 \pm 434.33
	200	317.16 \pm 19.59	930.10 \pm 55.09	2737.00 \pm 177.92	6675.04 \pm 423.62
	400	495.60 \pm 14.09	917.06 \pm 42.79	2660.00 \pm 159.81	6862.56 \pm 478.47
	800	900.90 \pm 13.62	1141.40 \pm 35.11	2732.00 \pm 144.54	6461.80 \pm 418.29
	1000	1076.60 \pm 11.72	1236.08 \pm 31.46	2741.00 \pm 167.38	6297.38 \pm 409.92
Stopping iterations: $K(\epsilon)$ for FSP $K'(\epsilon)$ for SSP	50	51.25 \pm 7.37	406.89 \pm 28.95	27.18 \pm 1.49	3.77 \pm 0.21
	100	61.21 \pm 10.01	369.71 \pm 28.53	27.07 \pm 1.75	3.72 \pm 0.19
	200	59.58 \pm 9.79	366.05 \pm 27.55	26.37 \pm 1.78	3.99 \pm 0.27
	400	48.80 \pm 7.05	259.53 \pm 21.40	23.60 \pm 1.60	3.63 \pm 0.21
	800	51.45 \pm 6.81	171.70 \pm 17.55	20.32 \pm 1.45	3.43 \pm 0.23
	1000	39.30 \pm 5.86	119.04 \pm 15.73	18.41 \pm 1.67	3.08 \pm 0.21
Solution time (seconds)	50	0.95 \pm 0.08	11.58 \pm 1.62	4.69 \pm 0.57	2.39 \pm 0.25
	100	1.20 \pm 0.15	10.34 \pm 1.31	5.08 \pm 0.73	2.27 \pm 0.19
	200	1.41 \pm 0.18	12.19 \pm 1.49	5.21 \pm 0.62	2.65 \pm 0.27
	400	1.57 \pm 0.18	9.71 \pm 1.11	4.74 \pm 0.62	2.57 \pm 0.23
	800	3.12 \pm 0.40	10.78 \pm 1.43	4.73 \pm 0.60	2.48 \pm 0.25
	1000	3.59 \pm 0.52	9.80 \pm 1.48	4.91 \pm 0.68	2.25 \pm 0.20

Table 3: Summary of results for 10D for select initial sample sizes n_0

eration, the decrease in the number of iterations reduces the overall solution time. This trend continues with S:A2RP, resulting in the lowest solution times despite the largest sample sizes. Note that 10D does not have the computational speedup from warm-starting. As a final remark, we compared the theoretical sample size estimates to the sample sizes used by the sequential procedures. Using properties of 10D, we derived a sample size estimate $N \geq 481,600$ to guarantee solutions x_N^* that are within 2% optimality with a 90% probability [26]. It is well-known that such theoretical bounds are overly conservative; our results agree.

6.4 Effect of ϵ

In this section, we examine the performance of the procedures as ϵ shrinks. To do this, we ran 100 replications of the procedures with varying values of $\epsilon = d \cdot z^*$ by decreasing d by an amount of 0.01 from 0.05 (5% of optimality) down to 0.01 (1% of optimality), and then further to $d = 0.005$, $d = 0.002$ and $d = 0.001$ (0.1% of optimality). We set $n_0 = 100$ and compared the performance of the procedures with respect to $N(\epsilon)$ and $N'(\epsilon)$, $K(\epsilon)$ and $K'(\epsilon)$, and solution time. For brevity, we only report our results for APL1P as we obtained similar results for 10D. We also note that the values corresponding to $d = 0.002$ are listed in Table 2.

Table 4 shows the results of our experiments as a percentage change in the stopping iteration, sample sizes used, and the solution time compared to the base case of $d = 0.05$ for FSP and SSP with A2RP. There are substantial increases in the stopping iteration, sample sizes used, and solution time once we fall below $d = 0.01$ (within 1% of optimality). The increases in $N'(\epsilon)$ and solution time are most substantial for SSP. As ϵ shrinks, F:A2RP(100) performs better relative to F:A2RP(2) in terms of solution time as fewer optimization problems are solved (see changes in $K(\epsilon)$). Propositions 6 and 9 indicate asymptotic increases in $N(\epsilon)$ and $N'(\epsilon)$ proportional to $1/\epsilon^2$, w.p.1. The smallest value of ϵ used in the experiments is $\epsilon = 24.6423$ when $d = 0.001$. This value is still far from 0 for the asymptotic results to hold. Observe that the reason why we are able to obtain the results in Propositions 6 and 9 is that the optimality gap and variance estimate of the current candidate solution will be 0 for all sample sizes greater than a random N and M , based on the sample path, w.p.1 (Lemma 5). This will happen when all sampling problems return the same optimal solution. Figure 2 indicates that this probability is quite low even at higher sample sizes for APL1P. Therefore, a much smaller ϵ is needed to invoke these propositions for this problem.

We also examined how coverage probability changes as ϵ shrinks, depicted in Figure 5. F:A2RP(100) and S:A2RP have consistently high coverage. For FSP with SRP, coverage probability is high for large values of ϵ and as ϵ shrinks, it first decreases and then starts to increase again. We explain this as follows: When ϵ is large, as the results of §6.1 show, the probability of obtaining an ϵ -optimal solution is high. This results in good coverage with quick solution times. Then, as ϵ is decreased, this effect diminishes and the undercoverage problem of SRP shows. However, as ϵ is further decreased, the effect of Theorem 3 starts to be seen and the coverage improves. This trend can also be seen slightly in F:A2RP(2), although the coverage remains above 90%.

% Change in	$d =$		0.001	0.002	0.005	0.01	0.02	0.03	0.04	0.05
	Procedure									
$N(\epsilon), N'(\epsilon)$	F:A2RP(2)		102.7	38.8	8.7	2.5	0.2	0.0	0.0	0.0
	F:A2RP(100)		572.0	283.0	112.0	43.0	7.0	1.0	0.0	0.0
	S:A2RP		9938.9	1804.8	277.7	54.4	3.1	0.2	0.0	0.0
$K(\epsilon), K'(\epsilon)$	F:A2RP(2)		5135.0	1938.0	434.0	124.0	12.0	2.0	0.0	0.0
	F:A2RP(100)		572.0	283.0	112.0	43.0	7.0	1.0	0.0	0.0
	S:A2RP		132.0	119.0	98.0	49.0	7.0	1.0	0.0	0.0
Sol. Time (s)	F:A2RP(2)		4557.2	1347.2	258.8	71.7	4.5	2.9	0.0	0.0
	F:A2RP(100)		2390.9	826.2	223.4	67.1	12.4	0.7	0.2	0.0
	S:A2RP		19069.8	2210.5	362.1	68.1	1.7	1.3	3.0	0.0

Table 4: Percentage change in sample size, stopping iteration, and solution time for various values of d ($\epsilon = d \cdot z^*$) relative to $d = 0.05$ for APL1P




Figure 5: Coverage probability as ϵ shrinks

6.5 Selection of ϵ and Results on 20TERM

In this section, we illustrate how to apply the procedures in practice on a large-scale test problem, 20TERM, whose solution is unknown. Recall that 20TERM has 40 independent stochastic parameters and 1.1×10^{12} scenarios [30]. Based on our earlier findings, we set $k_f^m = k_f^n = 3$ and again used $\alpha = 0.10$ and $h(n) = 1/\sqrt{n}$. Next, we need to determine a value of ϵ to run the procedures. We proceed as follows.

Suppose the user is willing to use sample sizes of at most N_{\max} to solve the sampling approximation problems within the sequential procedures. We can use the stopping criterion (8) or (25) to estimate an appropriate value for ϵ that corresponds to N_{\max} . Rewriting the stopping rule yields the following lower bound on ϵ :

$$\epsilon \geq G_{N(\epsilon)} + \frac{t_{K(\epsilon)} s_{N(\epsilon)} + 1}{\sqrt{N(\epsilon)}} = \epsilon_{LB}. \quad (36)$$

An estimate $\widehat{\epsilon}_{LB}$ for ϵ_{LB} can be obtained by replacing $N(\epsilon)$ with N_{\max} , and estimating $G_{N(\epsilon)}$ and $s_{N(\epsilon)}^2$ through sampling as follows. For a moderate sample size $N_{\text{trial}} < N_{\max}$, we can run independent replications to collect values of $G_{N_{\text{trial}}}$ and $s_{N_{\text{trial}}}^2$. We can then use the averages of these values collected, denoted $\bar{G}_{N_{\text{trial}}}$ and $\bar{s}_{N_{\text{trial}}}^2$, as estimates for $G_{N(\epsilon)}$ and $s_{N(\epsilon)}^2$, respectively in (36). We select $\epsilon = \widehat{\epsilon}_{LB}$, where

$$\widehat{\epsilon}_{LB} = \bar{G}_{N_{\text{trial}}} + \frac{z_{\alpha} \bar{s}_{N_{\text{trial}}} + 1}{\sqrt{N_{\max}}}$$

For 20TERM, we used $N_{\max} = 750$ and calculated $\bar{G}_{N_{\text{trial}}}$ and $\bar{s}_{N_{\text{trial}}}$ by taking the averages of 25 replications with $n_0 = N_{\text{trial}} = 500$. Under these experimental settings, we obtained $\epsilon = \widehat{\epsilon}_{LB} = 50.61$. Using this value for ϵ , we ran the procedures on 20TERM and report the results in Table 5.

In Table 5, we report 90% confidence intervals on the sample sizes used, the stopping iterations, and the solution time computed out of 100 independent runs. The procedures stop on average in a small number of iterations using moderate sample sizes, resulting in relatively tight confidence intervals. In particular, SSP stops in slightly fewer number of iterations and faster overall solution time. The average solution times illustrate the increased difficulty in solving sampling instances of 20TERM compared to the two smaller test problems APL1P and 10D considered earlier.

Performance Measure	n_0	F:A2RP(2)	F:A2RP(100)	S:A2RP
Sample sizes used:	500	502.48 \pm 0.64	593.00 \pm 21.21	683.90 \pm 46.76
Stopping iterations:	500	2.24 \pm 0.32	1.93 \pm 0.21	1.54 \pm 0.11
Coverage probability estimate:	500	1.00 \pm 0.00	1.00 \pm 0.00	1.00 \pm 0.00
Solution time (seconds):	500	1906.83 \pm 208.30	2104.78 \pm 312.32	1794.30 \pm 214.36

Table 5: Summary of results for 20TERM

To test the quality of the solutions obtained, we conducted a separate study to estimate their optimality gaps. We used the *multiple replications procedure* (MRP) of Mak, Morton and Wood [30], which has been shown to be more conservative compared to A2RP in empirical studies [2]. That is, MRP typically results in larger CI widths than A2RP with more than the desired coverage probability. We generated 30 gap estimators, each using a sample size of 500, calculated their variance and formed a 90% confidence interval on the optimality gap of the solutions obtained. Note that these 30 samples of size 500 are independent of the ones used for estimating ϵ above. We checked the frequency of the MRP interval widths that were below ϵ . The row ‘‘Coverage probability estimate’’ in Table 5 shows the results. Even though MRP CIs can be more conservative, we found that all MRP CI widths were within ϵ . This suggests that we are 90% confident that our solutions are indeed ϵ -optimal solutions. To further see how these solutions compare to an actual optimum solution, we calculated a lower bound on z^* by solving a single scenario problem where all the stochastic parameters are set to their expected values, called the EV-problem. Since the second-stage decisions are continuous, the optimum value of the EV-problem results in a valid lower bound on z^* . We found $z_{\text{EV}}^* = 239, 272.8 \leq z^*$. Comparing z_{EV}^* to the ϵ used, we see that the ϵ -optimal solutions are indeed within at most 0.0211% of optimality. The results indicate that we have obtained high-quality ϵ -optimal solutions with a high probability (approximately 90%).

To further verify this method of selecting ϵ , we also tested it on APL1P and obtained favorable results. For instance, with $n_0 = N_{\text{trial}} = 100$ and $N_{\text{max}} = 1000$, we obtain an ϵ that is 0.52% of z^* . If the user wishes to solve smaller problems, e.g., $N_{\text{max}} = 500$, ϵ is slightly increased to 0.59%. For brevity, we do not report on the details of these experiments. However, Tables 2 and 4 can be used to see that the procedures performed well with this ϵ selection for APL1P.

Finally, we tested the effect of the choice of the inflation function $h(n)$. The results (sample sizes used, stopping iteration, coverage probability) remain virtually the same for the following choices of $h(n)$: $1/n$, $1/\sqrt{n}$, $1/\log n$, and $1/\log \log n$. Notice that for these test problems, ϵ typically takes on values that are much greater than the largest value of $h(n)$, and the procedures stop when both gap and variance estimators are relatively quite small; therefore, the choice of $h(n)$ does not have a significant effect.

6.6 Observations from Computational Experiments

In this section, we summarize our observations and provide preliminary guidelines for implementation based on our computational experiments.

- While increasing the sample sizes, we can choose to augment previously generated observations or generate an entirely new set of independent observations (resampling). Generating completely new independent observations results in quicker solution times but can lower coverage probability, as has been observed earlier for a different stopping rule in [3].
- The computational experiments indicate that when only the samples used to evaluate the candidate solutions are resampled (and the samples for generating candidate solutions are always augmented), the process somewhat adapts itself to the stream of candidate solutions generated. That is, if a bad solution stream is obtained, then there is a high likelihood that the procedure stops at a bad solution, the same is true for a good solution stream. This can result in a less than desired coverage probability at small initial sample sizes. On the other hand, if both sample streams—the one to generate candidate solutions and the one to evaluate them—are resampled, this effect diminishes. In this case, the procedures stop quicker and interestingly, the coverage probability remains high. Therefore, we recommend to resample an entirely new set of observations for both sample streams. We resample at a frequency of every 3 iterations, as this resulted in the quickest solution times with desired coverage levels.

- SRP stops faster than A2RP but can have low coverage for small initial sample sizes. This observation matches with our earlier observations in the nonsequential and sequential settings [2, 3]. It is well known that a valid interval estimator in the nonsequential setting can result in a lower coverage probability when used in a sequential setting [13]. Therefore, we recommend the use of more conservative but more reliable A2RP.
- Computationally, larger increases in sample sizes (e.g., 100 by 100 as opposed to 2 by 2) can be preferable. Larger increases in sample sizes result in fewer iterations; fewer optimization problems are solved in this case. Even though larger problems are solved, when considerably fewer of them are solved, overall solution times decrease.
- When the initial sample size is small to have reliable estimators and/or when there is a large bias or variance, SSP can yield large sample size estimates. This can be computationally taxing. Therefore, when the initial sample size is small, we recommend the use of FSP with larger increases in sample sizes. When the initial sample size is large enough (typically ≥ 500), SSP can result in moderate jumps with faster overall solution time.
- Theorems 3 and 7 suggest selecting ϵ as small as possible. However, as the experiments in §6.4 as well as the results of Propositions 6 and 9 show, the sample sizes when the procedures stop with a small ϵ can be quite large; hence, this can be computationally demanding. Our experiments indicate that with the use of *reliable* confidence interval estimators, high-quality solutions can be obtained with a high probability when moderate values of ϵ are used. In §6, we provide an approach to estimate ϵ by inverting the stopping rule and calculating an estimate for ϵ as a function of the targeted sample size. Empirical results indicate that this approach performs well in practice.

7 Conclusions

In this paper, we developed sequential stopping rules for Monte Carlo sampling-based procedures for a class of stochastic programs. The procedures solve sampling approximations with increasing sample sizes and construct confidence intervals on the optimality gap of the obtained solutions, and stop when the confidence interval's width plus an inflation factor falls below a pre-specified value, $\epsilon > 0$. We provided conditions under which the procedures find ϵ -optimal solutions and terminate in a finite number of iterations using finite number of sample sizes with probability one. We provided guidelines to increase the sample sizes when the sample sizes are given as input schedules, and presented a method that calculates the next sample size according to the current statistical estimates.

More reliable estimators will increase the success of the sequential procedures. Our ongoing work includes reducing the bias and variance of these estimators. We are also empirically comparing the relative-width stopping rules in [3] to the fixed-width rules of this paper. Other future work includes extending the procedures to allow for approximate solutions of the sampling approximations and designing adaptive sequential procedures.

Acknowledgements

The authors thank two anonymous referees and the associate editor for valuable suggestions that improved the paper, Anton Kleywegt and Tito Homem-de-Mello for providing the 10D instance, and Andrzej Ruszczyński and Artur Świanowski for access to their regularized decomposition code. This research is partially supported by the National Science Foundation under Grant EFRI-0835930.

References

- [1] A. Alonso-Ayuso, L.F. Escudero, A. Garín, and M.T. Ortuño. An approach for strategic supply chain planning under uncertainty based on stochastic 0-1 programming. *Journal of Global Optimization*, 26:97–124, 2003.
- [2] G. Bayraksan and D.P. Morton. Assessing solution quality in stochastic programs. *Mathematical Programming*, 108:495–514, 2006.

- [3] G. Bayraksan and D.P. Morton. A sequential sampling procedure for stochastic programming. *Operations Research*, 59(4):898–913, 2011.
- [4] B. Betro and F. Schoen. Optimal and suboptimal stopping rules for the multistart algorithm in global optimisation. *Mathematical Programming*, 57:445–458, 1992.
- [5] C. Boender and A. Rinnooy Kan. On when to stop sampling for the maximum. *Journal of Global Optimization*, 1:331–340, 1991.
- [6] Y. S. Chow and H. Robbins. On the asymptotic theory of fixed-width sequential confidence intervals for the mean. *Annals of Mathematical Statistics*, 36(2):457–462, 1965.
- [7] L. Dai, C.H. Chen, and J.R. Birge. Convergence properties of two-stage stochastic programming. *Journal of Optimization Theory and Applications*, 106:489–509, 2000.
- [8] G.B. Dantzig and G. Infanger. A probabilistic lower bound for two-stage stochastic programs. Technical Report SOL 95-6, Department of Operations Research, Stanford University, November 1995.
- [9] C. C. Y. Dorea. Stopping rules for a random optimization method. *SIAM Journal on Control and Optimization*, 28(4):841–850, 1990.
- [10] A. Futschik and G.Ch. Pflug. Optimal allocation of simulation experiments in discrete stochastic optimization and approximative algorithms. *European Journal of Operational Research*, 101:245–260, 1997.
- [11] B.K. Ghosh and P.K. Sen, editors. *Handbook of Sequential Analysis*. Marcel Dekker, Inc., New York, 1991.
- [12] M. Ghosh, N. Mukhopadhyay, and P. K. Sen. *Sequential Estimation*. Wiley, New York, 1997.
- [13] P. W. Glynn and W. Whitt. The asymptotic validity of sequential stopping rules for stochastic simulations. *The Annals of Applied Probability*, 2(1):180–198, 1992.
- [14] A. Gut. *Stopped Random Walks: Limit Theorems and Applications*. Springer-Verlag, New York, 1988.
- [15] W. E. Hart. Sequential stopping rules for random optimization methods with applications to multistart local search. *SIAM Journal on Optimization*, 9(1):270–290, 1998.
- [16] J.L. Higle and S. Sen. Statistical verification of optimality conditions for stochastic programs with recourse. *Annals of Operations Research*, 30:215–240, 1991.
- [17] J.L. Higle and S. Sen. Duality and statistical tests of optimality for two stage stochastic programs. *Mathematical Programming*, 75:257–275, 1996.
- [18] J.L. Higle and S. Sen. *Stochastic Decomposition: A Statistical Method for Large Scale Stochastic Linear Programming*. Kluwer Academic Publishers, Dordrecht, 1996.
- [19] T. Homem-de-Mello. Variable-sample methods for stochastic optimization. *ACM Transactions on Modeling and Computer Simulation*, 13(2):108–133, 2003.
- [20] T. Homem-de-Mello. On rates of convergence for stochastic optimization problems under non-i.i.d. sampling. *SIAM Journal on Optimization*, 19:524–551, 2008.
- [21] L.J. Hong and B.L. Nelson. The tradeoff between sampling and switching: New sequential procedures for indifference-zone selection. *IIE Transactions*, 37:623–634, 2005.
- [22] G. Infanger. Monte Carlo (importance) sampling within a Benders decomposition algorithm for stochastic linear programs. *Annals of Operations Research*, 39:69–95, 1992.
- [23] Y.M. Kaniovski, A.J. King, and R.J.-B. Wets. Probabilistic bounds (via large deviations) for the solutions of stochastic programming problems. *Annals of Operations Research*, 56:189–208, 1995.
- [24] B.D. Keller and G. Bayraksan. Scheduling jobs sharing multiple resources under uncertainty: A stochastic programming approach. *IIE Transactions*, 42:16–30, 2010.

- [25] S.-H. Kim and B.L. Nelson. On the asymptotic validity of fully sequential selection procedures for steady-state simulation. *Operations Research*, 54(3):475–488, 2006.
- [26] A.J. Kleywegt, A. Shapiro, and T. Homem-de-Mello. The sample average approximation method for stochastic discrete optimization. *SIAM Journal on Optimization*, 12(3):479–502, 2001.
- [27] G. Lan, A. Nemirovski, and A. Shapiro. Validation analysis of mirror descent stochastic approximation method. *Mathematical Programming*, 2011. Forthcoming. Online first DOI 10.1007/s10107-011-0442-6.
- [28] A.M. Law and W.D. Kelton. Confidence intervals for steady-state simulations II: A survey of sequential procedures. *Management Science*, 28:550–562, 1982.
- [29] A.M. Law, W.D. Kelton, and L.W. Koenig. Relative width sequential confidence intervals for the mean. *Communications in Statistics*, B10:29–39, 1981.
- [30] W.K. Mak, D.P. Morton, and R.K. Wood. Monte Carlo bounding techniques for determining solution quality in stochastic programs. *Operations Research Letters*, 24:47–56, 1999.
- [31] D.P. Morton. Stopping rules for a class of sampling-based stochastic programming algorithms. *Operations Research*, 46:710–718, 1998.
- [32] D.P. Morton, F. Pan, and K.J. Saeger. Models for nuclear smuggling interdiction. *IIE Transactions*, 39:3–14, 2007.
- [33] A. Nadas. An extension of a theorem of Chow and Robbins on sequential confidence intervals for the mean. *Annals of Mathematical Statistics*, 40(2):667–671, 1969.
- [34] A. Nemirovski, A. Juditsky, G. Lan, and A. Shapiro. Robust stochastic approximation approach to stochastic programming. *SIAM Journal on Optimization*, 19:1574–1609, 2009.
- [35] V.I. Norkin, G.Ch. Pflug, and A. Ruszczyński. A branch and bound method for stochastic global optimization. *Mathematical Programming*, 83:425–450, 1998.
- [36] R. Pasupathy. On choosing parameters in retrospective-approximation algorithms for stochastic root finding and simulation optimization. *Operations Research*, 58(4):889–901, 2010.
- [37] G. C. Pflug. Stepsize rules, stopping times and their implementations in stochastic quasigradient algorithms. In Y. Ermoliev and R. J.-B. Wets, editors, *Numerical Techniques for Stochastic Optimization*. Springer Verlag, Berlin, 1988.
- [38] E. Polak and J.O. Royset. Efficient sample sizes in stochastic nonlinear programming. *Journal of Computational and Applied Mathematics*, 217:301–310, 2008.
- [39] W. Richter. Limit theorems for sequences of random variables with sequences of random indices. *Theory of Probability and its Applications*, 10:74–84, 1965.
- [40] A. Ruszczyński. A regularized decomposition method for minimizing a sum of polyhedral functions. *Mathematical Programming*, 35:309–333, 1986.
- [41] A. Ruszczyński and A. Świetanowski. Accelerating the regularized decomposition method for two stage stochastic linear problems. *European Journal of Operational Research*, 101:328–342, 1997.
- [42] A. Shapiro. Monte Carlo sampling methods. In A. Ruszczyński and A. Shapiro, editors, *Handbooks in Operations Research and Management Science, Volume 10: Stochastic Programming*, pages 353–425. Elsevier, 2003.
- [43] A. Shapiro and S. Ahmed. On a class of minimax stochastic programs. *SIAM Journal on Optimization*, 14:1237–1249, 2004.
- [44] A. Shapiro and T. Homem-de-Mello. On rate of convergence of Monte Carlo approximations of stochastic programs. *SIAM Journal on Optimization*, 11:70–86, 2000.
- [45] B. Verweij, S. Ahmed, A. Kleywegt, G. Nemhauser, and A. Shapiro. The sample average approximation method applied to stochastic vehicle routing problems: a computational study. *Computational and Applied Optimization*, 24:289–333, 2003.