

SEQUENTIAL SAMPLING FOR SOLVING STOCHASTIC PROGRAMS

Güzin Bayraksan

Department of Systems and Industrial Engineering
The University of Arizona
Tucson, AZ 85721, USA

David P. Morton

Graduate Program in Operations Research
The University of Texas at Austin
Austin, TX 78712, USA

ABSTRACT

We develop a sequential sampling procedure for solving a class of stochastic programs. A sequence of feasible solutions, with at least one optimal limit point, is given as input to our procedure. Our procedure estimates the optimality gap of a candidate solution from this sequence, and if that point estimate is sufficiently small then we stop. Otherwise, we repeat with the next candidate solution from the sequence with a larger sample size. We provide conditions under which this procedure: (i) terminates with probability one and (ii) terminates with a solution which has a small optimality gap with a prespecified probability.

1 INTRODUCTION

Stochastic programming provides a framework for decision making under uncertainty, which extends deterministic mathematical programming to incorporate random variables and probabilistic statements. A major challenge in the analysis of stochastic programs of practical size is having to compute a difficult high-dimensional expectation with expensive function evaluations. Monte Carlo sampling-based methods can provide an attractive approximation. These methods replace the probabilistic statements in the model (e.g., expectation) with their sampling estimators (e.g., sample mean). They are usually justified asymptotically, by providing conditions under which the approximating solutions solve the stochastic program as the sample size grows large. However, practical implementations require a finite sample size. A question of interest is then, what this sample size should be to have a good approximate solution. One way to approach this is to sequentially increase the sample size until we reach a good solution. Such a procedure requires a reliable means of stopping and rules to increase the sample size. Moreover, some statement regarding the quality of the associated solution needs to be made. These issues are the focus of this paper.

We assume a sequence of candidate solutions, with at least one limit point that solves the original stochastic program, is given as an input. Such a sequence can be generated by solving a series of sampling problems with an increasing sample size. (See [Higle and Sen 1991b](#) for a computationally effective way to do so for a class of two-stage stochastic linear programs.) That said, we allow candidate solutions to be generated by any method. Given a candidate solution, we assess its quality with increasing sample size, measuring quality in terms of the candidate solution's optimality gap. At each iteration when assessing the candidate solution's quality, we have the option to: (i) reuse previously-generated sample observations and add additional, newly-generated samples, or (ii) generate an entire new set of observations. Our ability to use a warm start to more efficiently solve the next iteration's stochastic program favors the former choice. On the other hand, the risk of persisting with a "bad" set of samples suggests we should occasionally choose the latter option. We investigate this tradeoff. We terminate the procedure when a stopping criterion is satisfied, and we prove that asymptotically, this procedure yields a high quality solution with a desired probability.

In sequential sampling, the sample size is not fixed. Rather, it depends on the observations collected so far, and hence is a random variable. Two classic sequential problems in estimation involve forming fixed-width ([Chow and Robbins 1965](#)) and relative-width ([Nadas 1969](#)) confidence intervals for the mean by sequentially increasing the sample size. When simulating stochastic systems, the simulation run length can take the place of the sample size. For steady-state simulations, [Law and Kelton \(1982\)](#) and [Law, Kelton, and Koenig \(1981\)](#) survey sequential methods for constructing fixed-width and relative-width confidence intervals for the mean performance measure of a stochastic system. [Glynn and Whitt \(1992\)](#) provide conditions under which asymptotic validity of sequential stopping rules for simulations can be achieved. More recent work in the area of simulation involves selecting from a number of alterna-

tive system designs using sequential sampling; see Kim and Nelson (2001, 2006).

M-estimators (Huber 1981) generalize maximum likelihood estimators, and sequential sampling procedures have been developed for M-estimators (Hlávka 2000). These procedures differ from ours in that they assume differentiability and focus on estimation of an optimum solution. We mainly focus on stochastic programs with recourse, and our computational results are for two-stage stochastic linear programs with recourse (SLP-2). In SLP-2, the objective function is convex and typically also non-smooth. Moreover, we are indifferent to how close we are to the set of optimum solutions as long as the candidate solution's objective function value is close to the optimum value, i.e., the optimality gap is small.

There has been considerable work in stochastic programming on employing Monte Carlo methods. Ermoliev (1988) surveys stochastic quasi-gradient methods and Pflug (1988) surveys step-size and stopping rules. Dantzig and Glynn (1990) and Infanger (1992) use importance sampling to reduce estimation variance within a cutting-plane scheme, and Hige and Sen (1996b) update sampling-based cuts to ensure desirable asymptotics. For stochastic global optimization, Norkin, Pflug, and Ruszczyński (1998) sample within a branch-and-bound algorithm. Even though there has been work on how to stop the procedures, and how to assess the quality of the solutions resulting from these sampling-based methods (Dantzig and Infanger 1995, Hige and Sen 1991a, Hige and Sen 1996a), sequential issues that arise have received little attention.

Instead of embedding sampling within an optimization algorithm, another approach is to first sample a set of observations and then simply solve the resulting problem. There is a significant literature on large sample size properties of this approach; see, e.g., the survey of Shapiro (2003). For convex, piecewise linear stochastic programs that have a unique, sharp optimum, Shapiro, Kim, and Homem-de-Mello (2002) provide insight as to the sample size needed to find the optimal solution via large deviations theory. When independent samples are drawn at a sequence of iterations, Homem-de-Mello (2003) studies rates at which the sample sizes must grow to ensure consistency of the objective function estimator, and he derives associated error statements in the spirit of the law of the iterated logarithm. For stochastic nonlinear programs, Polak and Royset (2007) propose a procedure that approximately minimizes the computational effort required to reduce an initial optimality gap by a pre-specified fraction, in the context of so-called diagonalization schemes.

Morton (1998) develops stopping rules for a class of algorithms that use asymptotically-normal optimality gap estimates, formed as a difference of upper- and lower-bound estimators. For minimization problems, upper-bound estimators can be formed for a fixed, feasible solution, and

asymptotic normality is easy to achieve. A natural lower-bound estimator arises from optimizing a sample-mean objective function (Mak, Morton, and Wood 1999, Norkin, Pflug, and Ruszczyński 1998). Unfortunately, this estimator is not, in general, asymptotically normal, and so the approach of Morton (1998) does not apply. In this paper, we overcome this difficulty, and allow use of the non-normal optimized sample-mean estimator. We improve on Morton (1998) in two other important ways: The confidence interval we form on the optimality gap uses the (observable) sample variance estimator instead of the (unknown) population variance, and we develop sampling procedures under weaker moment conditions. Importantly, our main results do not require independent and identically distributed (i.i.d.) sampling, and hence can be applied when using other sampling schemes that are designed, e.g., to reduce variance. Another attractive feature of our sequential approach is its flexibility in how observations can be generated at each iteration. One option is to use a single stream of observations in which at each iteration we simply augment the existing set of observations with a few new additional samples. Alternatively, the observations from the previous iteration can be discarded and we can generate an entirely new set of observations of increased size. Intermediate options also exist, and are permitted by the theory we develop.

2 FRAMEWORK

We consider

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

where the expectation is taken with respect to random vector $\tilde{\xi}$. We assume $X \neq \emptyset$, X is compact, $E \sup_{x \in X} |f(x, \tilde{\xi})| < \infty$, and $f(\cdot, \tilde{\xi})$ is lower semicontinuous (lsc) on X , with probability one (w.p.1). This ensures $E f(\cdot, \tilde{\xi})$ is lsc, and hence (SP) has a finite optimal solution, which is achieved on X .

Our procedure works as follows: At iteration k , we are given $\hat{x}_k \in X$. We select a sample size, n_k , and estimate \hat{x}_k 's optimality gap. The procedure stops when the optimality gap estimate falls below a certain level. Otherwise, we continue with $\hat{x}_{k+1} \in X$ and $n_{k+1} \geq n_k$. Let X^* be the set of optimal solutions to (SP). We assume:

- A1.** The sequence of feasible candidate solutions $\{\hat{x}_k\}$ has at least one limit point in X^* , with probability one (w.p.1).

Such a sequence can be found by solving a series of sampling problems

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

with optimal solutions x_n^* and with sample sizes $n = m_k$, such that $m_k \rightarrow \infty$ as $k \rightarrow \infty$. Other ways to find $\{\hat{x}_k\}$ are also possible. We do assume that the method that generates $\{\hat{x}_k\}$ does not depend on the sampled observations used in our evaluation procedures.

For any $x \in X$, let $\mu_x = Ef(x, \xi) - z^*$ and $\sigma^2(x) = \text{var}[f(x, \xi) - f(x_{\min}^*, \xi)]$, where $x_{\min}^* \in \arg \min_{y \in X^*} \text{var}[f(y, \xi) - f(x, \xi)]$. If (SP) has a unique optimum, i.e., $X^* = \{x^*\}$, then, of course, $x_{\min}^* = x^*$.

Let $\xi^1, \xi^2, \dots, \xi^n$ be a sample, which could be i.i.d. as ξ or could be drawn in another way, e.g., to reduce variance. Suppose we have at hand $G_n(x)$, that uses this sample of size n to estimate μ_x , and similarly we have an estimator $s_n^2(x) \geq 0$ of $\sigma^2(x)$. We define

$$D_n(x) = \frac{1}{n} \sum_{i=1}^n [f(x, \xi^i) - f(x_{\min}^*, \xi^i)], \quad (1)$$

where x_{\min}^* is defined as above for this x . The estimators $D_n(x)$, $G_n(x)$ and $s_n^2(x)$ all use the same n observations and we assume the following:

- A2.** Let $\{x_k\}$ be a feasible sequence (i.e., $x_k \in X$) with x as one of its limit points. Let sample size n_k satisfy $n_k \rightarrow \infty$ as $k \rightarrow \infty$. Then, $\liminf_{k \rightarrow \infty} P(|G_{n_k}(x_k) - \mu_x| > \delta) = 0$ for any $\delta > 0$.
- A3.** $G_n(x) \geq D_n(x)$, w.p.1 for all $x \in X$ and $n \geq 1$.
- A4.** $\liminf_{n \rightarrow \infty} s_n^2(x) \geq \sigma^2(x)$, w.p.1 for all $x \in X$.
- A5.** $\sqrt{n}(D_n(x) - \mu_x) \Rightarrow N(0, \sigma^2(x))$ as $n \rightarrow \infty$ for all $x \in X$, where $N(0, \sigma^2(x))$ is a normal random variable with mean zero and variance $\sigma^2(x)$. Here, " \Rightarrow " denotes convergence in distribution.

As an example of estimators that satisfy the above assumptions, let $\xi^1, \xi^2, \dots, \xi^n$ be i.i.d. as ξ , let x_n^* solve (SP_n) for this sample, let $\bar{f}_n(x) = \frac{1}{n} \sum_{i=1}^n f(x, \xi^i)$ and form:

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

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

With this choice of estimators, Assumption A3 is immediate since $\bar{f}_n(x_n^*) \leq \bar{f}_n(x_{\min}^*)$, w.p.1. $D_n(x)$ is simply a sample mean of i.i.d.r.v.s, and hence A5 holds by the standard central limit theorem (CLT), provided $\sigma^2(x) < \infty$. Sufficient conditions for A4 to hold under i.i.d. sampling are given in Bayraksan and Morton (2006). Note that when (SP) has multiple optimal solutions we cannot expect $\{x_n^*\}$ to have a single limit point and hence we cannot expect $s_n^2(x)$ to converge as $n \rightarrow \infty$. However, A4 is a form of consistency for $s_n^2(x)$ in the sense that it is bounded below by $\sigma^2(x)$. A sufficient condition under which A2 holds is that $\bar{f}_n(x)$ converges uniformly to continuous limit $Ef(x, \xi)$

on X , w.p.1, which holds under i.i.d. sampling and compact X provided $f(\cdot, \xi)$ is continuous on X , w.p.1, and $E \sup_{x \in X} |f(x, \xi)| < \infty$.

3 SEQUENTIAL SAMPLING PROCEDURE

At iteration $k \geq 1$ of the sequential procedure, we select n_k observations to assess \hat{x}_k 's quality. We can choose to generate ξ^1, \dots, ξ^{n_k} independently of those generated in previous iterations. Or, we can augment the observations from the previous iteration by generating $n_k - n_{k-1}$ additional observations, $\xi^{n_{k-1}+1}, \dots, \xi^{n_k}$ (assume $n_0 = 0$). We form $G_{n_k}(\hat{x}_k)$, and $s_{n_k}^2(\hat{x}_k)$, and if the stopping criterion is not satisfied, we repeat with sample size, $n_{k+1} \geq n_k$ and \hat{x}_{k+1} . To simplify notation we suppress dependence on \hat{x}_k and n_k , and simply denote $\mu_k = \mu_{\hat{x}_k}$, $\sigma_k^2 = \sigma^2(\hat{x}_k)$, $D_k = D_{n_k}(\hat{x}_k)$, $G_k = G_{n_k}(\hat{x}_k)$ and $s_k = s_{n_k}(\hat{x}_k)$. We terminate at iteration

$$T = \inf_{k \geq 1} \{G_k \leq h's_k + \epsilon'\}, \quad (3)$$

i.e., the first time G_k 's width relative to s_k falls below $h' > 0$ plus a small positive number ϵ' , which ensures finite stopping (see next section).

The stopping criterion (3) is with respect to $h's_T + \epsilon'$, and the statement regarding the quality of the candidate solution when we stop is with respect to a larger relative term $hs_T + \epsilon$, where $h > h'$ and $\epsilon > \epsilon'$. (Typically, we choose the epsilon terms so they are small compared to h' .) For the procedure we propose we will show

$$\liminf_{h, h'} P(\mu_T \leq hs_T + \epsilon) \geq 1 - \alpha. \quad (4)$$

So, \hat{x}_T 's optimality gap is a fraction of the sample standard deviation plus ϵ , with a desired probability provided h is close enough to h' .

At iteration k , we choose the sample size according to

$$n_k \geq \left(\frac{1}{h-h'} \right)^2 (c_p + 2p \ln^2 k), \quad (5)$$

where $c_p = \max\{2 \ln(\sum_{k=1}^{\infty} k^{-p \ln k} / \sqrt{2\pi\alpha}), 1\}$. Here, $p > 0$ is a parameter we can choose, which affects the number of samples we generate. In the next section, we show if the sample size satisfies (5) then (4) holds under a finite moment generating function assumption. We also show the procedure stops in a finite number of steps. We summarize our procedure below.

Sequential Sampling Procedure:

Input: Values for $h > h' > 0$, $\epsilon > \epsilon' > 0$, $0 < \alpha < 1$, and $p > 0$. Method that generates $\{\hat{x}_k\}$ with at least one limit point in X^* . Resampling frequency k_f , a positive integer.

Output: Candidate solution, \hat{x}_T , and a $(1 - \alpha)$ -level confidence interval on its optimality gap, μ_T .

0. (Initialization) Set $k = 1$, calculate n_k as given in (5), and sample observations $\tilde{\xi}^1, \tilde{\xi}^2, \dots, \tilde{\xi}^{n_k}$.

1. Use $\tilde{\xi}^1, \tilde{\xi}^2, \dots, \tilde{\xi}^{n_k}$ to form G_k and s_k^2 .

2. If $\{G_k \leq h's_k + \varepsilon'\}$, then set $T = k$, and go to 4.

3. Set $k = k + 1$ and calculate n_k according to (5). If k_f divides k then sample observations $\tilde{\xi}^1, \tilde{\xi}^2, \dots, \tilde{\xi}^{n_k}$, independently of samples generated in previous iterations. Else, sample $n_k - n_{k-1}$ observations $\tilde{\xi}^{n_{k-1}+1}, \tilde{\xi}^{n_{k-1}+2}, \dots, \tilde{\xi}^{n_k}$ from the distribution of $\tilde{\xi}$. Go to 1.

4. Output candidate solution \hat{x}_T and a one-sided confidence interval on μ_T ,

$$[0, hs_T + \varepsilon]. \quad (6)$$

If $k_f = 1$ then at every iteration we sample n_k new observations. At the other extreme, if k_f is sufficiently large then we continually augment the existing set of observations throughout execution of the procedure.

4 ASYMPTOTIC VALIDITY & FINITE STOPPING

In this section, we first present a result that shows (4) holds under a finite moment generating function assumption. Then, we prove and discuss finite stopping of the algorithm. In the next section we relax the moment generating function assumption, replacing it with a finite r^{th} moment assumption. We state our results without proof and refer to Bayraksan and Morton (2007) for the proofs.

4.1 Finite Moment Generating Function

Assume for some $\gamma_0 > 0$

$$\sup_{n \geq 1} \sup_{x \in X} E \left[e^{\gamma \left(\frac{D_n(x) - \mu_x}{\sigma(x)/\sqrt{n}} \right)} \right] < \infty, \text{ for } |\gamma| \leq \gamma_0, \quad (7)$$

where $D_n(x)$ is defined in (1). When $\tilde{\xi}^1, \tilde{\xi}^2, \dots, \tilde{\xi}^n$ are i.i.d., a sufficient condition to ensure (7) is that the moment generating function (MGF) of the scaled random variables $[(f(x, \tilde{\xi}) - f(x_{\min}^*, \tilde{\xi})) - \mu_x] / \sigma(x)$ exists, i.e.,

$$\sup_{x \in X} E \left[e^{\gamma \left(\frac{(f(x, \tilde{\xi}) - f(x_{\min}^*, \tilde{\xi})) - \mu_x}{\sigma(x)} \right)} \right] < \infty \text{ for } |\gamma| \leq \gamma_0. \quad (8)$$

Below we state and prove the validity of the sequential sampling procedure under hypothesis (7). Our result, given in (4), is asymptotic, as $h \downarrow h'$, i.e., as the sample sizes grow. Theorem 1's proof rests on: (i) Fatou's Lemma, which provides inequalities when "lim inf" and an integral (or an infinite sum) are exchanged and (ii) a bound on the tail of a normal random variable.

Theorem 1 Assume A3, A4 and A5 are satisfied and (7) holds for $\gamma_0 > 0$. Let $M^2 = \sup_{x \in X} \sigma^2(x)$, $\varepsilon > \varepsilon' > 0$, $p > 0$ and $0 < \alpha < 1$. Consider the sequential sampling procedure where the sample size is increased according to (5). If the procedure stops at iteration T according to (3) then,

$$\liminf_{h \downarrow h'} P(\mu_T \leq hs_T + \varepsilon) \geq 1 - \alpha. \quad (9)$$

Theorem 1 shows that for values of h close enough to h' , or, when the sample sizes n_k are large enough, we have the optimality gap of the solution when we stop within $[0, hs_T + \varepsilon]$ with at least the desired probability of $1 - \alpha$. We now turn our attention to finite stopping. As the following theorem indicates, the sequential procedure stops with probability one in a finite number of iterations.

Theorem 2 Assume A1 and A2 are satisfied. Let $\varepsilon' > 0$ and $h > h' > 0$ be fixed. Then, for the sequential sampling procedure where the sample size is increased according to (5), and the procedure stops at iteration T according to (3), we have $P(T < \infty) = 1$.

4.2 Weaker Moment Conditions

In this section, we prove a variant of Theorem 1 that assumes finite moments up to order r . Specifically, we relax the MGF assumption of the previous section to

$$\sup_{x \in X} E \left[f(x, \tilde{\xi}) \right]^r < \infty, \quad (10)$$

for some even integer $r \geq 2$. Note that (10) implies $\sup_{x \in X, y \in X^*} E[f(x, \tilde{\xi}) - f(y, \tilde{\xi})]^r < \infty$. We also simplify the analysis by only considering i.i.d. sampling. Under these assumptions, we select the sample size at each iteration k according to

$$n_k \geq \left(\frac{1}{h - h'} \right)^2 \left(c_{p,q} + 2pk^{2q/r} \right), \quad (11)$$

where $q > 1$, $p > 0$ and where $c_{p,q} = \max\{2 \ln(\sum_{k=1}^{\infty} \exp(-pk^{2q/r}) / \sqrt{2\pi\alpha}), 1\}$. The growth in the sample size is of order $O(k^{2q/r})$ and we must choose $q > 1$. So, if condition (10) holds for $r = 2$ we can choose q just larger than unity so we essentially have the sample size growing at a linear rate, and if condition (10) holds for $r = 4$ we can obtain a rate that essentially grows with $k^{1/2}$. In other words, a less restrictive assumption on the existence of moments implies a faster rate of growth for the sample sizes.

Under the moment generating function hypothesis, the sample size formula (5) contained a parameter, $p > 0$, which was at our disposal. Now, under the moment hypothesis (10), if we elect the slowest possible rate of growth of n_k

by choosing q just larger than unity then we can again view the sample size formula as being parameterized by a scalar term, p . We discuss how to choose these parameters in the next section.

We now state the validity and finite stopping of our procedure under the finite r^{th} moment assumption given in (10). Key to this proof is a lemma that establishes a bound on the central moments of a sample mean, which we state first.

Lemma 3 *Let X^1, X^2, \dots, X^n be i.i.d. random variables with mean μ and $\bar{X}_n = \frac{1}{n} \sum_{i=1}^n X^i$. If $E|X^1 - \mu|^r < \infty$ for some integer $r \geq 1$ then*

$$E(\bar{X}_n - \mu)^r \leq E|X^1 - \mu|^r \left(\frac{r}{n}\right)^{r/2}.$$

Theorem 4 *Assume A1-A4 are satisfied, (10) holds with $r \geq 2$ even, and that ξ^1, ξ^2, \dots are i.i.d. as ξ . Let $\varepsilon > \varepsilon' > 0$, $p > 0$, $q > 1$ and $0 < \alpha < 1$. Then, for the sequential sampling procedure where the sample size is increased according to (11), and the procedure stops at iteration T according to (3),*

$$P(T < \infty) = 1 \quad \text{and} \quad \liminf_{h \downarrow h'} P(\mu_T \leq hs_T + \varepsilon) \geq 1 - \alpha.$$

Under i.i.d. sampling, Theorem 4 differs from Theorem 1 in that the MGF assumption of (8) is replaced by the moment condition (10). Under this weaker assumption, the sample sizes are chosen according to (11) instead of (5), which require a larger number of sampled observations.

So far, we have characterized desirable theoretical properties of our sequential procedure. In the next section, we discuss issues that arise when implementing the procedure. In particular, we discuss how to choose p when using the sample size formula (5) and how to choose q and p when using (11). Then, we apply the procedure to two test problems that are two-stage stochastic linear programs with recourse to examine its performance.

5 CHOOSING PARAMETERS FOR n_k

At iteration k , we set n_k according to (5) under the MGF assumption and according to (11) under the r^{th} moment assumption. Then, using this many samples, we solve a sampling problem (SP_{n_k}) to estimate the optimality gap of the current candidate solution and its associated variance. Suppose the procedure terminates in T iterations. Then, we solve T sampling problems, $(\text{SP}_{n_1}), (\text{SP}_{n_2}), \dots, (\text{SP}_{n_T})$. Therefore, the computational effort exerted for the evaluation of the candidate solutions is approximately proportional to $\sum_{k=1}^T n_k$. (For some decomposition methods, empirical studies suggest the effort to solve a stochastic program grows linearly in the number of scenarios, see e.g., [Ruszczynski and Swietanowski 1997](#) and [Verweij, Ahmed, Kleywegt, Nemhauser, and Shapiro 2003](#)). In other words, the effort is proportional to $S_M(p)$ and $S_W(p, q)$, which are the following

two respective expressions

$$T \max \left\{ 2 \ln \left(\frac{\sum_{k=1}^{\infty} k^{-p \ln k}}{\sqrt{2\pi\alpha}} \right), 1 \right\} + 2p \sum_{k=1}^T \ln^2 k \quad (12a)$$

$$T \max \left\{ 2 \ln \left(\frac{\sum_{k=1}^{\infty} \exp(-pk^{2q/r})}{\sqrt{2\pi\alpha}} \right), 1 \right\} + 2p \sum_{k=1}^T k^{2q/r}. \quad (12b)$$

Assume, for the moment, that T is known. The parameter $p > 0$ for $S_M(p)$ and the parameters $p > 0$ and $q > 1$ for $S_W(p, q)$ are at our disposal. To reduce the computational effort, we would like to choose them to minimize $S_M(p)$ and $S_W(p, q)$. The following result helps to do so.

Theorem 5 *Let $S_M(p)$ and $S_W(p, q)$ be defined in (12). $S_M(\cdot)$ is convex on $\{p : p > 0\}$ and $S_W(\cdot, q)$ is convex on $\{p : p > 0\}$ for fixed $q > 1$. Furthermore, $S_M(p)$ and $S_W(p, q)$ are both bounded below by*

$$2T \ln \left(\frac{T}{\sqrt{2\pi\alpha}} \right). \quad (13)$$

When employing the sample size formula (5), the associated function $S_M(p)$ is convex, and associated minimizers p^* for various values of T are shown in the second and third columns of Table 1. When we use sample size formula (11) we seek to select p and q to solve $\min_{p>0, q>1} S_W(p, q)$. By Theorem 5, $S_W(\cdot, q)$ is convex and so we can minimize this function for a fixed value of q . The fourth and fifth columns of Table 1 show the results of doing so for $q = 1.5$. We know by the lower-bounding values shown in the final column of the table that these are suboptimal by no more than 2.5-5.5%.

The results of Table 1 can guide selection of p given rough estimates for T . Of course, the assumption that T is known is unrealistic. While we could view T as a random variable and attempt to minimize S_M or S_W in expectation, we will not do so. When the procedure terminates at a different iteration than that of the assumed T , the differences in sample sizes are quite modest. For instance, taking $h - h' = 0.5$ and $\alpha = 0.10$, and using $p^1 = 1.91 \times 10^{-1}$ (for $T = 50$), we have $n_T \geq 33, 56$, and 65 samples when $T = 1, 50$ and 100 respectively under the MGF assumption of (7). Similarly, we have $n_T \geq 37, 55$, and 63 when we instead use $p^2 = 1.53 \times 10^{-1}$ (for $T = 100$). For the weaker moment condition with $r = 2$ and $q = 1.5$, $n_T \geq 39, 52$, and 77 when $T = 1, 50$ and 100 respectively when we use $p^1 = 4.67 \times 10^{-3}$ and $n_T \geq 45, 50$, and 58 samples when we instead use $p^2 = 1.66 \times 10^{-3}$. Slightly smaller values of $S_W(p, q)$ (and n_T) can be obtained with even larger values of q but the results are more sensitive to having assumed the “wrong” value of T .

Table 1: Choices of p that minimize $S_M(p)$ and $S_W(p, q)$ for a given number of iterations, T . In $S_W(p, q)$, we fix $q = 1.5$ and $r = 2$. $LB = 2T \ln(T/(\sqrt{2\pi}\alpha))$ denotes the lower bound given in Theorem 5. The values for S_M and S_W (and LB) are reported are for $\alpha = 0.10$, but the same p^* values are optimal, e.g., for $\alpha = 0.05$.

T	p^*	$S_M(p^*)$	p^*	$S_W(p^*, 1.5)$	LB
10	4.07×10^{-1}	82	5.05×10^{-2}	78	74
50	1.91×10^{-1}	591	4.67×10^{-3}	552	530
100	1.53×10^{-1}	1,334	1.66×10^{-3}	1,243	1,198
500	1.04×10^{-1}	8,421	1.49×10^{-4}	7,822	7,598
1,000	9.08×10^{-2}	18,333	5.27×10^{-5}	17,031	16,583

6 APPLICATION TO SLP-2

We apply the sequential sampling procedure to two-stage stochastic linear programs with recourse (SLP-2s). We present computational results on two test problems from the literature, PGP2 and APL1P. PGP2 is an electric power generation model with 3 stochastic parameters and 576 scenarios (Higle and Sen 1996b). APL1P, another power generation model, has 5 independent stochastic parameters and 1280 scenarios (Infanger 1992). These small problems can be solved exactly, and so it is unnecessary to use simulation to solve them. Our purpose for considering them here is that they allow us to assess the performance of our sequential sampling procedures, e.g., we can test whether confidence intervals cover the optimality gap they are designed to estimate. In addition, we use these as test problems since we know they pose challenges for the optimality gap estimators we use.

6.1 Gap and Variance Estimators

We use four methods to form the gap and variance estimators. These are the single replication procedure (SRP), the averaged two-replication procedure (A2RP), and their δ -optimal versions (δ -optimal SRP, and δ -optimal A2RP), which we now detail. For the SRP estimators, step 1 of the procedure of §3 computes gap and variance estimators according to equations (2), i.e., step 1 becomes:

- 1.a. Solve (SP $_{n_k}$) using i.i.d. $\tilde{\xi}^1, \tilde{\xi}^2, \dots, \tilde{\xi}^{n_k}$. Obtain $x_{n_k}^*$,
- 1.b. Form $G_k = \frac{1}{n_k} \sum_{i=1}^{n_k} \left(f(\hat{x}_k, \tilde{\xi}^i) - f(x_{n_k}^*, \tilde{\xi}^i) \right)$, and $s_k^2 = \frac{1}{n_k - 1} \sum_{i=1}^{n_k} \left[\left(f(\hat{x}_k, \tilde{\xi}^i) - f(x_{n_k}^*, \tilde{\xi}^i) \right) - \left(\bar{f}_{n_k}(\hat{x}_k) - \bar{f}_{n_k}(x_{n_k}^*) \right) \right]^2$.

The implementation under A2RP involves small changes. We select n_k even and divide the observations into two random partitions and calculate G_k^i and s_k^{2i} as above (steps 1.a and 1.b) for each sample containing $n_k/2$ observations, $i = 1, 2$. We then pool these gap and variance estimators to obtain $\bar{G}_k = \frac{1}{2}(G_k^1 + G_k^2)$ and $\bar{s}_k^2 = \frac{1}{2}(s_k^{21} + s_k^{22})$. For the sequential procedure that uses A2RP, we use these

pooled estimators for the stopping criterion in step 3 of the procedure. The gap and variance estimators formed by SRP and A2RP under i.i.d. sampling satisfy assumptions A2-A5 of §2 (Bayraksan and Morton 2006).

We also use δ -optimal versions of these methods in which we solve the sampling problem(s) in step 1.a suboptimally, i.e., to within $100 \cdot \delta\%$ of optimal. Our motivation for doing so is two-fold. First, it is unnecessary (and computationally wasteful) to numerically optimize a sample mean orders of magnitude more precisely than the confidence interval width output by our procedure. Second, we found in Bayraksan and Morton (2006) that numerically precise optimization of the (SP $_n$) used in assessing solution quality can degrade coverage properties, even in the non-sequential setting due to an x_n^* coinciding with the candidate solution. This property also motivated development of the A2RP estimators described above. When implementing the δ -optimal procedures, we set the suboptimality level to $\delta = 10^{-3}$, which was found to yield good coverage results in the non-sequential setting (Bayraksan and Morton 2006).

6.2 Generating Candidate Solutions

The method that generates candidate solutions is an input to our procedure, and any method suffices provided its sequence has at least one limit point in X^* . We generate $\{\hat{x}_k\}$ by solving a separate sampling problem (SP $_{m_k}$) with increasing sample sizes, as described below.

- i. Set $m_k = m_1$. Sample i.i.d. observations (independent of those used in evaluation procedures) $\tilde{\xi}^1, \tilde{\xi}^2, \dots, \tilde{\xi}^{m_k}$,
- ii. Solve (SP $_{m_k}$) to obtain $x_{m_k}^*$,
- iii. Set $\hat{x}_k = x_{m_k}^*$. Sample $m_{k+1} - m_k$ i.i.d. observations $\tilde{\xi}^{m_k+1}, \tilde{\xi}^{m_k+2}, \dots, \tilde{\xi}^{m_{k+1}}$. Set $k = k + 1$ and go to ii.

In step i, we use a separate stream of i.i.d. observations from the distribution of $\tilde{\xi}$, independent from those generated in steps 0 and 2 of the sequential procedure. In our experiments, we set $m_k = 2n_k$, i.e., we are willing to spend more computational effort to find high-quality candidate solutions and use fewer sample sizes to evaluate them. All limit points of this sequence of candidate solutions is in the set of optimal solutions, X^* , under mild conditions satisfied by our test problems. So, assumption A1 is satisfied.

6.3 Parameters Used

Both of the test problems satisfy assumptions stated in §2 ($X \neq \emptyset$ and is compact, $f(\cdot, \tilde{\xi})$ is continuous on X , w.p.1., etc.) Moreover, the random vector $\tilde{\xi}$ of these test problems each has finite discrete distributions with independent components. Therefore, the MGF assumption (7) is satisfied for all γ_0 and hence we use the sample size formula (5). We set $\alpha = 0.10$ and design the procedure for $T = 50$, i.e., we set $p = 1.91 \times 10^{-1}$ (see §5). Table 2 lists the values of h' and

Table 2: Parameters and the corresponding initial sample sizes used in the tests. Other parameters are the same for both test problems: $\alpha = 0.10$, $\varepsilon = 2 \times 10^{-7}$, $\varepsilon' = 1 \times 10^{-7}$, and $p = 1.91 \times 10^{-1}$.

	PGP2	APL1P
h'	0.025	0.015
h	0.312	0.217
n_1	100	200

Table 3: Effect of resampling frequency, k_f . Results are from applying the sequential procedure with SRP to PGP2. We report total solution time (in seconds), total time spent in assessing solution quality (in seconds), total number of sampling problems, (SP_n s), solved for assessing solution quality, and an average time spent per sampling problem for assessing solution quality. We also report average value of T , the iteration the sequential procedure stopped along with a 90% confidence interval half-width for ET .

k_f	Sol Time	Assess Time	(SP_n)s Solved	Assess Time/ (SP_n)	T
100	307.97	102.91	2152	0.05	54.14 ± 16.44
75	279.56	93.51	1950	0.05	41.46 ± 12.78
50	243.99	82.28	1728	0.05	30.50 ± 8.96
25	196.42	67.08	1339	0.05	16.27 ± 2.71
12	162.62	61.27	1007	0.06	11.07 ± 1.72
6	151.95	66.55	808	0.08	8.65 ± 1.40
3	133.70	64.40	606	0.11	6.13 ± 0.79
1	135.10	77.35	417	0.19	4.19 ± 0.59

h used for the two test problems PGP2 and APL1P. With the given parameters, the sequential procedure uses $n_1 \geq 100$ for PGP2 and $n_1 \geq 200$ for APL1P. When implementing the procedure, we set $\varepsilon = 2 \times 10^{-7}$ and $\varepsilon' = 1 \times 10^{-7}$.

6.4 Computational Results

We use the regularized decomposition algorithm of Ruszczyński (1986). Specifically, we modified the implementation of Ruszczyński and Świetanowski (1997) to warm-start the algorithm when additional samples are added to the current problem. This helps us to quickly find a solution when we augment a problem with a few additional samples. However, as mentioned before, our method also allows for creating an entirely new set of samples. Augmenting is computationally attractive but we may get trapped, for an extended number of iterations, in a bad sample path. We first examine the trade-off between augmentation and resampling by changing the resampling frequency, k_f . We only examine the resampling frequency issue in the context of assessing solution quality, i.e., we always augment the sampling problem (SP_{m_k}) used for generating the candidate solutions.

Table 3 reports results of our tests when varying k_f from 100 down to 1 using the SRP estimators on PGP2. Each row of the table reports averages of the following, over 100 independent runs of the procedure: the total solution time, which includes both the time for generating the sequence of candidate solutions and the time for assessing their quality; the time spent assessing solution quality; the number of (SP_n)s solved in assessing solution quality; the time spent for assessing solution quality per sampling problem (SP_n); and, the average number of iterations of the sequential procedure (along with a 90% confidence interval halfwidth). Even though we report the total solution times, we focus on the computational effort for assessing solution quality since our approach allows for generating the sequence of candidate solutions in other ways. As k_f decreases, T and the total solution time both tend to decrease. However, the time spent in assessing solution quality first decreases, then, starts to increase. When we look at the time spent per sampling problem, the computational advantage of augmentation is in effect until $k_f = 25$, and as k_f further decreases, resampling increases the solution time per sampling problem. We did the same analysis for APL1P and found similar results for $k_f = 12$. For the rest of our computational experiments, we set $k_f = 25$ for PGP2 and $k_f = 12$ for APL1P.

We now turn to coverage properties of our sequential procedure, and examine these using the four different estimators described above. In Bayraksan and Morton (2006), we recommended the δ -optimal version of A2RP as a computationally-attractive alternative to the multiple-replication procedure of Mak, Morton, and Wood (1999). Both of those papers used non-sequential procedures and a single candidate solution, \hat{x} . Tables 4 and 5 summarize results for PGP2 and APL1P, respectively. The results are again based on 100 independent runs of the sequential procedures. We use the same stream of random numbers to feed each of the four estimation procedures. Also, since A2RP and its δ -optimal version use sample sizes that are even, we round up the sample sizes for SRP and its δ -optimal version so they are even, too. The tables report the number of iterations (T), the confidence interval width (CI), and the fraction of the 100 confidence intervals that contained the true optimality gap (\hat{p}). For each one of these values, we report the average of the 100 runs, along with their 90% confidence interval half-widths.

Tables 4 and 5 indicate that the A2RP method, on average, takes more iterations to terminate (T is larger), and the δ -optimal versions yield larger CIs and have more conservative coverage probabilities (\hat{p} is not below 0.90). The sequential procedure appears to work well for APL1P. The CI widths are within 0.5% of optimality for all methods and coverage probabilities are around or higher than 0.90. The results for PGP2 are more variable among the four methods. The variance of some of the frequently-obtained solutions is quite large for PGP2 and this results in larger

Table 4: Results for PGP2 ($k_f = 25$; $z^* = 447.324$ for PGP2).

Method	T	CI	\hat{p}
SRP	16.27 ± 2.71	10.49 ± 2.73	0.79 ± 0.07
δ -Opt SRP	14.97 ± 2.20	39.50 ± 17.37	0.96 ± 0.03
A2RP	107.28 ± 14.48	5.88 ± 2.13	0.76 ± 0.07
δ -Opt A2RP	103.29 ± 13.79	6.90 ± 2.02	0.98 ± 0.02

Table 5: Results for APL1P ($k_f = 12$; $z^* = 24,642.32$ for APL1P).

Method	T	CI	\hat{p}
SRP	15.40 ± 2.89	52.77 ± 8.54	0.88 ± 0.05
δ -Opt SRP	15.72 ± 2.90	54.98 ± 9.54	0.89 ± 0.05
A2RP	56.18 ± 13.05	66.10 ± 5.82	0.99 ± 0.02
δ -Opt A2RP	59.89 ± 13.30	73.24 ± 9.68	0.99 ± 0.02

CIs. A2RP and its δ -optimal version decrease the variance but the sequential procedure runs for a longer amount of time. For PGP2, the CIs obtained with δ -optimal A2RP, on average, are within 1.6% of optimality with high coverage probabilities. Overall, we view the results as consistent with our findings in the non-sequential setting, i.e., we view δ -optimal version of A2RP as preferable for these two test problems: Its coverage results are conservative, i.e., the procedure does not exhibit the risk of undercoverage that can arise from estimators rooted in the SRP and the numerically-precise A2RP.

7 CONCLUSIONS

In this paper, we develop a sequential sampling procedure to solve stochastic programs. We assume that a sequence of candidate solutions with at least one limit point that solves (SP) is given as an input to the procedure. Then, the sequential sampling procedure assesses the quality of these candidate solutions with increasing sample size and terminates according to a stopping criterion. The stopping criterion depends on the optimality gap estimate of the current solution and its associated variance. If the stopping criterion is satisfied, then the procedure outputs a confidence interval on the optimality gap of the current candidate solution. If not, the sample size is increased. The sequential procedure we develop allows for augmentation of additional observations to the ones that were previously generated, or, generation of an entirely new set of observations. Warm-starting techniques favor the first method and prevention of being trapped in a bad sample path favors the latter.

We provide rules to stop and to increase the sample sizes and a statement regarding the quality of the solution obtained when the procedure terminates. We show that this procedure asymptotically finds a high-quality solution with a desired probability, under certain assumptions. We also show that the procedure terminates in a finite number of

steps with probability 1. Then, we discuss several issues that arise in its implementation, such as determining what parameters to use. Our aim in choosing the parameters is to minimize the computational effort exerted while running the procedure. Finally, we apply the sequential sampling procedure to two two-stage stochastic linear programs with recourse from the literature. We test four different methods to assess a candidate solution’s quality that were previously developed. These are SRP, A2RP, and their δ -optimal versions. Our preliminary computational results indicate that the sequential sampling procedure with δ -optimal A2RP yields good coverage results with reasonable CI widths.

An area of future research is to develop more efficient sequential sampling procedures while maintaining the desired asymptotic properties. For instance, adaptive sequential methods can be designed where the sampling method takes into account the information obtained about the problem so far. Carefully designed adaptive methods can be more efficient while maintaining the desired asymptotic properties.

ACKNOWLEDGMENTS

The authors thank Andrzej Ruszczyński and Artur Świetanowski for access to their regularized decomposition code. This research was partially supported by the National Science Foundation under Grant DMI-0217927.

REFERENCES

- Bayraksan, G., and D. P. Morton. 2006. Assessing solution quality in stochastic programs. *Mathematical Programming* 108:495–514.
- Bayraksan, G., and D. P. Morton. 2007. A sequential sampling procedure for stochastic programming. Technical Report, University of Arizona, Available at www.sie.arizona.edu/faculty/guzinb.
- Chow, Y. S., and H. Robbins. 1965. On the asymptotic theory of fixed-width sequential confidence intervals for the mean. *Annals of Mathematical Statistics* 36:457–462.
- Dantzig, G. B., and P. W. Glynn. 1990. Parallel processors for planning under uncertainty. *Annals of Operations Research* 22:1–21.
- Dantzig, G. B., and G. Infanger. 1995, November. A probabilistic lower bound for two-stage stochastic programs. Technical Report SOL 95-6, Department of Operations Research, Stanford University.
- Ermoliev, Y. 1988. Stochastic quasigradient methods. In *Numerical Techniques for Stochastic Optimization*, ed. Y. Ermoliev and R. Wets, 141–185. Springer-Verlag, Berlin.
- Glynn, P. W., and W. Whitt. 1992. The asymptotic validity of sequential stopping rules for stochastic simulations. *The Annals of Applied Probability* 2:180–198.

- Higle, J. L., and S. Sen. 1991a. Statistical verification of optimality conditions for stochastic programs with recourse. *Annals of Operations Research* 30:215–240.
- Higle, J. L., and S. Sen. 1991b. Stochastic decomposition: An algorithm for two-stage linear programs with recourse. *Mathematics of Operations Research* 16:650–669.
- Higle, J. L., and S. Sen. 1996a. Duality and statistical tests of optimality for two stage stochastic programs. *Mathematical Programming* 75:257–275.
- Higle, J. L., and S. Sen. 1996b. *Stochastic decomposition: A statistical method for large scale stochastic linear programming*. Kluwer Academic Publishers, Dordrecht.
- Hlávka, Z. 2000. *Robust sequential methods*. Ph. D. thesis, Charles University, Prague.
- Homem-de-Mello, T. 2003. Variable-sample methods for stochastic optimization. *ACM Transactions on Modeling and Computer Simulation* 13:108–133.
- Huber, P. 1981. *Robust statistics*. Wiley, New York.
- Infanger, G. 1992. Monte Carlo (importance) sampling within a benders decomposition algorithm for stochastic linear programs. *Annals of Operations Research* 39:69–95.
- Kim, S. H., and B. L. Nelson. 2001. A fully sequential procedure for indifference-zone selection in simulation. *ACM TOMACS* 11:251–273.
- Kim, S. H., and B. L. Nelson. 2006. On the asymptotic validity of fully sequential selection procedures for steady-state simulation. *Operations Research* 54:475–488.
- Law, A. M., and W. D. Kelton. 1982. Confidence intervals for steady-state simulations ii: a survey of sequential procedures. *Management Science* 28:550–562.
- Law, A. M., W. D. Kelton, and L. W. Koenig. 1981. Relative width sequential confidence intervals for the mean. *Communications in Statistics* B10:29–39.
- Mak, W. K., D. P. Morton, and R. K. Wood. 1999. Monte Carlo bounding techniques for determining solution quality in stochastic programs. *Operations Research Letters* 24:47–56.
- Morton, D. P. 1998. Stopping rules for a class of sampling-based stochastic programming algorithms. *Operations Research* 46:710–718.
- Nadas, A. 1969. An extension of a theorem of chow and robbins on sequential confidence intervals for the mean. *Annals of Mathematical Statistics* 40 (2): 667–671.
- Norkin, V. I., G. Pflug, and A. Ruszczyński. 1998. A branch and bound method for stochastic global optimization. *Mathematical Programming* 83:425–450.
- Pflug, G. 1988. Stepsize rules, stopping times and their implementations in stochastic quasigradient algorithms. In *Numerical Techniques for Stochastic Optimization*. Springer Verlag, Berlin.
- Polak, E., and J. Royset. 2007. Efficient sample sizes in stochastic nonlinear programming. *Journal of Computational and Applied Mathematics*. Forthcoming.
- Ruszczyński, A. 1986. A regularized decomposition method for minimizing a sum of polyhedral functions. *Mathematical Programming* 35:309–333.
- Ruszczyński, A., and A. Świetanowski. 1997. Accelerating the regularized decomposition method for two stage stochastic linear problems. *European Journal of Operational Research* 101:328–342.
- Shapiro, A. 2003. Monte Carlo sampling methods. In *Stochastic Programming, Handbooks in Operations Research and Management Science*, ed. A. Ruszczyński and A. Shapiro. Elsevier, Amsterdam.
- Shapiro, A., J. Kim, and T. Homem-de-Mello. 2002. Conditioning of convex piecewise linear stochastic programs. *Mathematical Programming* 94:1–19.
- Verweij, B., S. Ahmed, A. Kleywegt, G. Nemhauser, and A. Shapiro. 2003. The sample average approximation method applied to stochastic vehicle routing problems: A computational study. *Computational and Applied Optimization* 24:289–333.

AUTHOR BIOGRAPHIES

DAVID P. MORTON is an Associate Professor in the Graduate Program in Operations Research in the Mechanical Engineering Department at The University of Texas at Austin. His research interests include computational stochastic programming, including simulation-based approximations in stochastic programming. His email address is morton@mail.utexas.edu and his web page is www.me.utexas.edu/~orie/Morton.html.

GUZIN BAYRAKSAN is an Assistant Professor in the Department of Systems and Industrial Engineering at the University of Arizona. Her research interests include stochastic optimization, particularly sampling-based methods for stochastic programming. Her email address is guzinb@sie.arizona.edu and her web page is www.sie.arizona.edu/faculty/guzinb.