

Simulation-Based Confidence Bounds for Two-Stage Stochastic Programs*

Peter W. Glynn and Gerd Infanger

Version 3/21/2008

Dedicated to the memory of George B. Dantzig

Abstract

This paper provides a rigorous asymptotic analysis and justification of upper and lower confidence bounds proposed by Dantzig and Infanger (1995) for an iterative sampling-based decomposition algorithm, introduced by Dantzig and Glynn (1990) and Infanger (1992), for solving two-stage stochastic programs. Extensions of the theory to cover use of variance reduction, different iterative sampling sizes, and the dropping of cuts are also presented. An extensive empirical investigation of the performance of these bounds establishes that the bounds perform reasonably on realistic problems.

*Technical Report SOL 2008-1.

1 Introduction

Consider the two-stage stochastic linear program (with recourse) given by

$$\begin{aligned} \min \quad & cx + E z^*(x; \Lambda) \\ \text{s/t} \quad & Ax = b \\ & x \geq 0, \end{aligned} \tag{1}$$

where $z^*(x; \Lambda)$ is the minimum of the second-stage linear program

$$\begin{aligned} \min \quad & \nu y \\ & \delta y = \mu + \beta x \\ & y \geq 0, \end{aligned}$$

and $\Lambda \equiv (\nu, \beta, \delta, \mu)$ is random. When the sample space corresponding to Λ contains a large number of outcomes (or is countably infinite or continuous), use of sampling-based algorithms for solving (1) becomes attractive. Of course a key element in the reliable application of every numerical method is the ability to accurately assess its error. For a sampling-based procedure, such an assessment typically comes in the form of confidence statements. Such confidence statements have been previously proposed for the class of sampling-based procedures based on “sample-average approximation” (SAA). Mak, Morton, and Wood (1999) show, for example, how upper and lower one-sided confidence intervals for the minimal objective value can be calculated (that are asymptotically valid as the sample size tends to infinity). Their approach actually applies to much more general stochastic programs than (1). On the other hand, Shapiro and Homem-de-Mello (1998) propose a stopping criterion for SAA, as applied to (1), that is based on the asymptotic behavior of confidence regions related to the first-order optimality conditions for two-stage stochastic programs with recourse.

This paper’s main contribution is the rigorous development of upper and lower one-sided confidence intervals for the minimum objective value of (1), when the underlying solution algorithm is not SAA but is instead the iterative sampling-based decomposition approach introduced by Dantzig and Glynn (1990) and Infanger (1992). Iterative sampling-based decomposition is an appealing algorithm because it lends itself better than does SAA to applications that demand that one iterate until a reasonable error tolerance has been achieved. These confidence intervals provide a useful indicator of the quality of the final solution at algorithmic termination, and were first proposed in a technical report by Dantzig and Infanger (1995) that lacked a complete theoretical justification. The current paper provides the theory necessary to make these confidence bounds rigorous and covers both discrete and continuous distributions on Λ .

This paper is organized as follows. In Section 2, we carefully describe the iterative sampling-based decomposition algorithm and then prove the asymptotic validity (as the number of samples per iteration goes to infinity) of our upper and lower confidence bounds on the minimum of z^* of (1). Section 3 is concerned with extending our results to modified versions of our basic algorithm (to cover use of variance reduction, different iterative sample

sizes, and to allow the dropping of cuts). We conclude in Section 4 with an extensive empirical investigation of the accuracy of our confidence bounds in the context of a suite of test problems.

2 Confidence Bounds for Iterative Sampling-based Decomposition

We start by describing the iterative sampling-based decomposition algorithm that we shall study. A key feature is its reliance on the dual linear program (LP), denoted $L^D(x; \Lambda)$, given by

$$\begin{aligned} \max \pi(\mu + \beta x) \\ \pi \delta \leq \nu \end{aligned}$$

that is dual to the (primal) subproblem $L^P(x; \Lambda)$ associated with (1), namely

$$\begin{aligned} \min \nu y \\ \delta y &= \mu + \beta x \\ y &\geq 0. \end{aligned}$$

We assume throughout this paper that:

A1. For each $x \in H \equiv \{x : Ax = b, x \geq 0\}$, $L^P(x; \Lambda)$ and $L^D(x; \Lambda)$ are almost surely (a.s.) feasible LP's.

Let $y^*(x; \Lambda)$ be the minimizer of $L^P(x; \Lambda)$ and let $\pi^*(x; \Lambda)$ be the maximizer of $L^D(x; \Lambda)$ (both assumed to be measurable selections in Λ ; see Kall (1976) for details of the existence of such a measurable selection). An important parameter in the specification of the algorithm is the number n of primal subproblems that are sampled at every iteration of the decomposition algorithm. We denote the master solution associated with performing m iterations, each based on n subproblem samples, by $X_{m,n}$.

Algorithm

1. Set $X_{0,n} = x_0 \in H$ as our initial guess at some minimizer x^* of (1).
2. $m \leftarrow 1$.
3. Generate n independent realizations of $\Lambda = (\nu, \beta, \delta, \mu)$, call them $\Lambda_1 = (\nu_1, \beta_1, \delta_1, \mu_1), \dots, \Lambda_n = (\nu_n, \beta_n, \delta_n, \mu_n)$.
4. Solve the n LP's $L^D(X_{m-1,n}; \Lambda_1), \dots, L^D(X_{m-1,n}; \Lambda_n)$ for the corresponding maximizers $\pi^*(X_{m-1,n}; \Lambda_1), \dots, \pi^*(X_{m-1,n}; \Lambda_n)$.

5. Compute the sample means

$$G_{m,n} = \frac{1}{n} \sum_{j=1}^n \pi^*(X_{m-1,n}; \Lambda_j) \beta_j,$$

$$g_{m,n} = \frac{1}{n} \sum_{j=1}^n \pi^*(X_{m-1,n}; \Lambda_j) \mu_j.$$

6. Solve the LP

$$\begin{aligned} \min cx &+ \theta \\ Ax &= b \\ -G_{i,n}x + \theta &\geq g_{i,n}, \quad 1 \leq i \leq m \\ x &\geq 0, \end{aligned}$$

for the minimizer $X_{m,n}$. (Note that this LP is computing the minimizer of $cx + \max_{1 \leq i \leq m} [G_{i,n}x + g_{i,n}]$ over $x \in H$.)

7. $m \leftarrow m + 1$ and return to 3.

The above algorithm produces a sequence of iterative approximations $X_{1,n} \dots, X_{m,n}$ to the solution of (1). The associated sequence of approximating optimal objective values is given by $((cX_{j,n} + \max_{1 \leq i \leq j} [G_{i,n}X_{j,n} + g_{i,n}]) : j \geq 1)$. Note that the family of n independent realizations of Λ that we generate at the m -th iteration are different from those generated at previous iterations. To make this difference notationally clear, we henceforth denote the n realizations of $\Lambda = (\nu, \beta, \delta, \mu)$ generated at iteration i by $\Lambda_{i,1} = (\nu_{i,1}, \beta_{i,1}, \delta_{i,1}, \mu_{i,1}), \dots, \Lambda_{i,n} = (\nu_{i,n}, \beta_{i,n}, \delta_{i,n}, \mu_{i,n})$. Relative to this notation, Step 3 of the algorithm is requiring that $(\Lambda_{i,j} : i \geq 1, j \geq 1)$ is a collection of independent and identically distributed (iid) replicates (copies) of Λ . Let $\mathcal{F}_{m,n} = \sigma(\Lambda_{j,i} : 1 \leq j \leq m, 1 \leq i \leq n)$ (with $\mathcal{F}_{0,n} = \{0, \Omega\}$) be the σ -algebra generated by the $\Lambda_{i,j}$'s that are needed by the first m iterations of the algorithm when sample size n is used at each iteration.

Our first result is an upper bound that is widely used in the setting of sampling-based optimization algorithms; see Asmussen and Glynn (2007). Let $z(x) = cx + E z^*(x; \Lambda)$ and note that because x^* is a minimizer of (1), it follows that $z(X_{m,n}) \geq z(x^*) \equiv z^*$. Hence $z(X_{m,n})$ is an upper bound on the minimum z^* of (1). Because $E z^*(\cdot; \Lambda)$ is not typically available in closed form (for otherwise one would not be applying sampling to solve (1)), $z(X_{m,n})$ must be computed by sampling. Suppose that we generate ℓ iid copies of Λ , independently of the previous realizations used to compute $X_{m,n}$. More precisely, suppose that we generate $\Lambda_{m+1,1}, \dots, \Lambda_{m+1,\ell}$; these random variates are clearly independent of $\mathcal{F}_{m,n}$. Then,

$$cX_{m,n} + \frac{1}{\ell} \sum_{j=1}^{\ell} z^*(X_{m,n}; \Lambda_{m+1,j}) \tag{2}$$

is an estimator for $z(X_{m,n})$. If we set $\sigma^2(x) \equiv \text{var } z^*(x; \Lambda)$ and

$$s_\ell^2(x) \equiv \frac{1}{\ell-1} \sum_{i=1}^{\ell} \left(z^*(x; \Lambda_{m+1,i}) - \frac{1}{\ell} \sum_{j=1}^{\ell} z^*(x; \Lambda_{m+1,k}) \right)^2,$$

we note that $\ell^{-1} s_\ell^2(X_{m,n})$ is an unbiased estimator for the conditional variance $\ell^{-1} \sigma_\ell^2(X_{m,n})$ of the estimator (2) (conditional on $\mathcal{F}_{m,n}$). The following upper (conditional) confidence bound on z^* thus follows easily from the central limit theorem (CLT).

Proposition 1. Assume A1. Suppose that $\text{var } z^*(x; \Lambda) < \infty$ for all $x \in H$. If η is chosen so that $P(N(0, 1) > \eta) = \alpha$, then

$$\liminf_{\ell \rightarrow \infty} P(z^* \leq cX_{m,n} + \frac{1}{\ell} \sum_{j=1}^{\ell} z^*(X_{m,n}; \Lambda_{m+1,j}) + \eta \sqrt{\frac{s_\ell^2(X_{m,n})}{\ell}} \mid \mathcal{F}_{m,n}) \geq 1 - \alpha \text{ a.s.}$$

(with equality when $z^*(x, \Lambda) > 0$).

The more subtle and difficult issue is that of a lower confidence bound on z^* . Our next result is the main theoretical contribution of this paper.

Theorem 1. Assume A1. If $\sigma^2(x^*) < \infty$, then

$$\liminf_{n \rightarrow \infty} P(z^* \geq cX_{m,n} + \max_{1 \leq i \leq m} (G_{i,n}X_{m,n} + g_{i,n}) - \eta \sqrt{\frac{\sigma^2(x^*)}{n}}) \geq 1 - \alpha,$$

provided that η is selected so that $P(N(0, 1) \leq \eta) = (1 - \alpha)^{1/m}$.

Proof. Because $X_{m,n}$ is a minimizer of

$$\min_{x \in H} [cx + \max_{1 \leq i \leq m} (G_{i,n}x + g_{i,n})],$$

it follows that

$$cX_{m,n} + \max_{1 \leq i \leq m} (G_{i,n}X_{m,n} + g_{i,n}) \leq cx^* + \max_{1 \leq i \leq m} (G_{i,n}x^* + g_{i,n}).$$

Hence,

$$\begin{aligned} P(cX_{m,n} + \max_{1 \leq i \leq m} (G_{i,n}X_{m,n} + g_{i,n}) \leq cx^* + z^*(x^*) + \frac{\eta\sigma(x^*)}{\sqrt{n}}) \\ \geq P(\max_{1 \leq i \leq m} (G_{i,n}x^* + g_{i,n}) \leq z^*(x^*) + \frac{\eta\sigma(x^*)}{\sqrt{n}}) \\ = P(\max_{1 \leq i \leq m} (G_{i,n}x^* + g_{i,n}) - z^*(x^*) \leq \frac{\eta\sigma(x^*)}{\sqrt{n}}). \end{aligned} \tag{3}$$

Recall that $\pi^*(x; \Lambda_{i,j})$ is a maximizer of $L^D(x; \Lambda_{i,j})$, so

$$\pi^*(X_{i-1,n}; \Lambda_{i,j})(\mu_{i,j} + \beta_{i,j}x) \leq \pi^*(x; \Lambda_{i,j})(\mu_{i,j} + \beta_{i,j}x)$$

because $\pi^*(X_{i-1,n}; \Lambda_{i,j})$ is feasible for $L^D(x; \Lambda_{i,j})$. By duality,

$$\pi^*(x; \Lambda_{i,j})(\mu_{i,j} + \beta_{i,j}x) = \nu_{i,j}y^*(x; \Lambda_{i,j}).$$

Hence,

$$\pi^*(X_{i-1,n}; \Lambda_{i,j})(\mu_{i,j} + \beta_{i,j}x) \leq \nu_{i,j}y^*(x; \Lambda_{i,j}) \text{ a.s.}$$

We conclude that

$$G_{i,n}x + g_{i,n} \leq \frac{1}{n} \sum_{j=1}^n z^*(x; \Lambda_{i,j}) \text{ as}$$

for $x \in H$, so that

$$G_{i,n}x^* + g_{i,n} - z^*(x^*) \leq \frac{1}{n} \sum_{j=1}^n z^*(x^*; \Lambda_{i,j}) - z^*(x^*) \text{ a.s.}$$

The probability (3) may therefore be lower bounded by

$$P\left(\max_{1 \leq i \leq m} n^{1/2} \left(\frac{1}{n} \sum_{j=1}^n z^*(x^*; \Lambda_{i,j}) - z^*(x^*)\right) \leq \sigma(x^*)\eta\right) \quad (4)$$

Note that the above probability is 1 if $\sigma^2(x^*) = 0$. On the other hand, if $\sigma^2(x^*) > 0$, then $\{z^*(x^*; \Lambda_{i,j}) - z^*(x^*) : 1 \leq i \leq m, 1 \leq j \leq n\}$ is a collection of iid mean zero random variables with finite variance $\sigma^2(x^*)$. Hence,

$$\begin{aligned} & \left(\frac{n^{1/2}}{\sigma(x^*)} \left(\frac{1}{n} \sum_{j=1}^n z^*(x^*; \Lambda_{i,j}) - z^*(x^*)\right), \dots, \left(\frac{n^{1/2}}{\sigma(x^*)} \left(\frac{1}{n} \sum_{j=1}^n z^*(x^*; \Lambda_{m,j}) - z^*(x^*)\right)\right) \right) \\ & \Rightarrow (N_1(0, 1), \dots, N_m(0, 1)) \end{aligned}$$

as $n \rightarrow \infty$ where $(N_1(0, 1), \dots, N_m(0, 1))$ are iid normal random variables with mean zero and unit variance. In view of (4), it follows that when $\sigma^2(x^*) > 0$,

$$\begin{aligned} \liminf_{n \rightarrow \infty} P(cX_{m,n} + \max_{1 \leq i \leq m} G_{i,n}X_{m,n} + g_{i,n}) & \leq cx^* + z^*(x^*) + \frac{\eta}{\sqrt{n}}\sigma(x^*) \\ & \geq P(\max_{1 \leq i \leq m} N_i(0, 1) \leq \eta) \\ & = P(N(0, 1) \leq \eta)^m = 1 - \alpha, \end{aligned}$$

proving the theorem. //

Theorem 1 provides the key lower bound. Note that the lower bound involves one quantity, namely $\sigma^2(x^*)$, that is not computed during the course of our algorithm. Theorem 1 effectively reduces the problem of obtaining a lower bound on z^* to computing $\sigma^2(x^*)$ or (at least) to computing an upper bound on $\sigma^2(x^*)$.

Since x^* is unknown, $\sigma^2(x^*)$ can not be estimated via the conditional sample variance $s_\ell^2(x^*)$. While the question of computing an upper bound on $\sigma^2(x^*)$ deserves further attention, we suggest here three different ways of bounding $\sigma^2(x^*)$.

Method 1: We bound $\sigma^2(x^*)$ by $\sigma_*^2 \equiv \sup\{\sigma^2(x) : x \in H\}$ and estimate σ_*^2 . To estimate σ_*^2 , we sample r points χ_1, \dots, χ_r at random from within H , compute $s_\ell^2(\chi_i)$ at each of the r points and bound σ_*^2 via the sample-based quantity

$$\hat{\sigma}_*^2(r, \ell) = \max_{1 \leq i \leq r} s_\ell^2(\chi_i). \quad (5)$$

An alternative is to generate independent Λ realizations of the second-stage subproblems at each of the r points, thereby calculating

$$\tilde{s}_\ell^2(\chi_i) = \frac{1}{\ell - 1} \sum_{i=1}^{\ell} \left(z^*(\chi_i; \Lambda_{m+i,j}) - \frac{1}{\ell} \sum_{k=1}^{\ell} z^*(\chi_i; \Lambda_{m+i,k}) \right)^2$$

for $1 \leq i \leq r$, and then compute

$$\tilde{\sigma}_*^2(r, \ell) = \max_{1 \leq i \leq r} \tilde{s}_\ell^2(\chi_i).$$

Note that because $s_\ell^2(\chi_i) \leq \max_{1 \leq k \leq r} s_\ell^2(\chi_k)$, it follows that

$$\sigma^2(\chi_i) \leq E [\hat{\sigma}_*^2(r, \ell) | \chi_1, \dots, \chi_r]$$

for $1 \leq i \leq r$. Furthermore if $\sigma(\cdot)$ is continuous over H and if the sequence χ_1, χ_2, \dots is a.s. asymptotically dense in H , then $\max_{1 \leq i \leq r} \sigma^2(\chi_i) \rightarrow \sigma_*^2$ a.s. as $r \rightarrow \infty$. It follows, under these conditions, that $\hat{\sigma}_*^2(r, \ell)$ is biased high for σ_*^2 . In addition, $P(\tilde{\sigma}_*^2(r, \ell) \leq \sigma_*^2 - \epsilon) \rightarrow 0$ as $r \rightarrow \infty$ for each $\epsilon > 0$, so that

$$\liminf_{n \rightarrow \infty, r \rightarrow \infty} P(z^* \geq cX_{m,n} + \max_{1 \leq i \leq m} (G_{i,n}X_{m,n} + g_{i,n}) - \eta \sqrt{\frac{\hat{\sigma}_*^2(r, \ell)}{n}} \geq 1 - \alpha, \quad (6)$$

provided that η is selected so that $P(N(0, 1) \leq \eta) = (1 - \alpha)^{1/m}$. An identical conclusion follows for $\tilde{\sigma}_*^2(r, \ell)$. An implementation of this method requires specifying r , ℓ , and the algorithm used to generate the χ 's.

Method 2: The second method uses the fact that $\sigma_*^2 \leq \sup\{E z^*(x; \Lambda)^2 : x \in H\}$. But

$$E z^*(x; \Lambda)^2 \leq E \sup_{x \in H} z^*(x; \Lambda)^2 \equiv E w(\Lambda)^2.$$

Assuming that $z^*(x; \Lambda)$ is non-negative a.s., the random variable $w(\Lambda)$ can be computed as the (random) solution of the optimization problem

$$\begin{aligned} \max_{\pi, x} \quad & \pi(\mu + \beta x) \\ & \pi \delta \leq \nu \\ & Ax = b \\ & x \geq 0 \end{aligned}$$

It follows that

$$\frac{1}{\ell} \sum_{i=1}^{\ell} w(\Lambda_{m+1,i})^2 \rightarrow E w(\Lambda)^2 \geq \sigma_*^2 \quad \text{a.s.}$$

as $\ell \rightarrow \infty$, so that $\ell^{-1} \sum_{i=1}^{\ell} w(\Lambda_{m+1,i})^2$ is an estimator of an upper bound on $\sigma^2(x^*)$, which can be used to implement a lower bound on z^* (as in (2) of Method 1).

Method 3: Method 3 works with the variance $\sigma^2(x; \Lambda)$ directly, rather than the second moment bound $E z(x; \Lambda)^2$. We start by observing that

$$\sigma^2(x) = \frac{1}{2} E (z^*(x; \Lambda_{m+1,1}) - z^*(x; \Lambda_{m+1,2}))^2.$$

But $z^*(x; \Lambda_{m+1,i})$ is a maximizer of $L^D(x; \Lambda_{m+1,i})$.

Assume that $\delta_{m+1,i} = D$ is deterministic. The feasible region of both dual linear programs are then contained in $F = \{\pi : \pi D \leq (\nu_{m+1,1} \vee \nu_{m+1,2})\}$, where $(\nu_{m+1,1} \vee \nu_{m+1,2})$ is the component-wise maximum of $\nu_{m+1,1}$ and $\nu_{m+1,2}$. Note that

$$\begin{aligned} |z^*(x; \Lambda_{m+1,1}) - z^*(x; \Lambda_{m+1,2})| &\leq \sup\{|\pi(\nu_{m+1,1} + \beta_{m+1,1}x) - \pi(\nu_{m+1,2} + \beta_{m+1,2}x)| : \pi \in F\} \\ &= \sup\{|\pi(\nu_{m+1,1} - \nu_{m+1,2}) + \pi(\beta_{m+1,1} - \beta_{m+1,2})x| : \pi \in F\}. \end{aligned}$$

So,

$$\begin{aligned} \sigma^2(x^*) &\leq \frac{1}{2} E \sup_{x \in H} \{(\pi(\nu_{m+1,1} - \nu_{m+1,2}) + \pi(\beta_{m+1,1} - \beta_{m+1,2})x)^2 : \pi \in F\} \\ &\leq \frac{1}{2} E \sup\{(\pi(\nu_{m+1,1} - \nu_{m+1,2}) + \pi(\beta_{m+1,1} - \beta_{m+1,2})x)^2 : \pi \in F, x \in H\} \equiv W \end{aligned}$$

Note that W is the maximum of a quadratic objective over a linear feasible set (with the feasible set having a “size” equal to that of a master plus a single subproblem). Let W_1, \dots, W_ℓ be ℓ iid replicates of W . Then,

$$\frac{1}{\ell} \sum_{i=1}^{\ell} W_i \rightarrow E W \geq \sigma^2(x^*) \quad \text{a.s.},$$

yielding an upper bound on $\sigma^2(x^*)$. This method avoids sampling over H , as does method 2.

We conclude this section by noting that Proposition 1 provides an upper bound on z^* that involves a confidence statement that is conditional on $\mathcal{F}_{m,n}$, whereas Theorem 1 offers a lower bound on z^* relative to the unconditional distribution. A typical user of such confidence bounds prefers that either both statements be conditional or that both be unconditional.

The argument underlying Theorem 1 fails badly, when one conditions on $\mathcal{F}_{m,n}$ (because, in that case, the $\Lambda_{i,j}$'s with $i \leq m$ and $j \leq n$ are no longer iid). As a consequence, we now offer a version of the upper bound (Proposition 1) that is unconditional.

Proposition 2. Assume A1 and that $E z^*(x; \Lambda)^k$ is finite and continuous on H for $1 \leq k \leq 4$. If $\sigma^2(\cdot)$ is positive and H is compact, then

$$\liminf_{\ell \rightarrow \infty} P(z^* \leq cX_{m,n} + \frac{1}{\ell} \sum_{j=1}^{\ell} z^*(X_{m,n}; \Lambda_{m+1,j}) + \eta \sqrt{\frac{s_{\ell}^2(X_{m,n})}{\ell}}) \geq 1 - \alpha.$$

Proof. Because $E |z^*(x; \Lambda)|^3 / \sigma^3(x)$ is bounded over H , the Berry-Esséen theorem (Feller (1971)) implies that

$$P(-z \leq \frac{\ell^{1/2}}{\sigma(x)} (\sum_{j=1}^{\ell} z^*(x; \Lambda_{m+1,j}) - z^*(x))) \rightarrow 1 - \alpha \quad (7)$$

as $\ell \rightarrow \infty$ uniformly in $x \in H$. Also Chebyshev's inequality guarantees that for each $\epsilon > 0$,

$$P(|\frac{s_{1,\ell}^2(x)}{\sigma^2(x)} - 1| > \epsilon) \rightarrow 0 \quad (8)$$

as $\ell \rightarrow \infty$, uniformly in $x \in H$. Relations (7) and (8) clearly imply the result. //

3 Extensions

We present here several extensions to our key lower bound result, namely Theorem 1.

Extension 1: Different sample sizes at each iteration

In using an iterative sampling-based algorithm, it is natural to consider a modified version in which the sample size n_i describing the number of independent subproblems sampled at iteration i is permitted to depend on i . For example, as the number of iterations increases, we may wish to increase the sample size in order to improve the accuracy as the iterative scheme closes in on the optimizer.

Let the sequence $(n_i : i \geq 1)$ be a (deterministic) sequence of integers corresponding to the successive sample sizes used at the different iterations. In this setting, the cut generated at iteration i has an associated sample size n_i , so we write it as (G_{i,n_i}, g_{i,n_i}) . Assume that

for $i \geq 1$, $n_i = c_i n(1 + o(1))$, as $n \rightarrow \infty$, where $o(1)$ is a sequence that tends to zero as $n \rightarrow \infty$. Then, a slight adaptation of the proof in Theorem 1 shows that under the conditions stated there,

$$\liminf_{n \rightarrow \infty} P(z^* \geq cX_{m,n_m} + \max_{1 \leq i \leq m} (G_{i,n_i} X_{m,n_m} + g_{i,n_i}) - \eta \sqrt{\frac{\sigma^2(x^*)}{n_1}}) \geq 1 - \alpha,$$

where η is chosen so that $\prod_{i=1}^n P(N(0,1) \leq \sqrt{\frac{c_i}{c_1}} \eta) = (1 - \alpha)$. Such a value for η can easily be computed numerically, at least when n is not too large.

Extension 2: Importance sampling

Suppose that we use a sample distribution \tilde{P} (rather than the original P) from which to sample Λ , so as to concentrate more realizations of Λ in those regions that are most critical to determining the (expected) second-stage cost. We assume that \tilde{P} satisfies

$$P(\xi \in dx) = l(x)\tilde{P}(\xi \in dx),$$

where P is the nominal distribution at ξ associated with the original probability distribution. The use of \tilde{P} rather than P to generate outcomes of ξ is known as importance sampling. When ξ has densities p and \tilde{p} under P and \tilde{P} , respectively, $l(\cdot)$ can be expressed as $l(x) = p(x)/\tilde{p}(x)$ and is known as the likelihood ratio of P relative to \tilde{P} .

Because of the presence of importance sampling, we modify the i 'th cut $(G_{i,n}, g_{i,n})$ to

$$\begin{aligned} \tilde{G}_{i,n} &= \frac{1}{n} \sum_{j=1}^n \pi^*(X_{i-1,n}; \Lambda_{i,j}) \beta_{i,j} l(\Lambda_{i,j}), \\ \tilde{g}_{i,n} &= \frac{1}{n} \sum_{j=1}^n \pi^*(X_{i-1,n}; \Lambda_{i,j}) \mu_{i,j} l(\Lambda_{i,j}). \end{aligned}$$

and execute the algorithm with these modified cuts. The proof of Theorem 1 goes through with minor changes, namely $z^*(x; \Lambda_{i,j})$ must be replaced by $z^*(x; \Lambda_{i,j})l(\Lambda_{i,j})$, so that $\sigma^2(x^*)$ is replaced by $\text{var } z^*(x; \Lambda)l(\Lambda)$ (where $\text{var } (\cdot)$ denotes the variance operator associated with distribution \tilde{P}).

When \tilde{P} is chosen well, use of the sampling distribution \tilde{P} rather than P can produce significant variance reductions, particularly in settings where there are second-stage outcomes that occur rarely but have a big effect in the decisions to be made; see pp. 127–129 of Asmussen and Glynn (2007) for further detail.

Extension 3: Control Variates

In virtually all real-world applications of sampling-based methods for solving (1), a wide variety of expectations of the form $E g(\Lambda)$ for g \mathcal{R}^ℓ -valued) can be computed either analytically in closed form or numerically to high precision (e.g., the mean of the right-hand side

variables μ and β). Let $C = g(\lambda) - E g(\Lambda)$ and note that $E C = 0$. Such a random variable C is called a control variate and, when intelligently used, can be used to induce a variance reduction in our sampling-based procedure. The idea is to replace the cut estimators $(G_{i,n}, g_{i,n})$ generated at iteration i by

$$G_{i,n}^c = \frac{1}{n} \sum_{j=1}^n (\pi^*(X_{i-1,n}; \Lambda_{i,j}) \beta_{i,j} - \lambda_1 C_{i,j}),$$

$$g_{i,n}^c = \frac{1}{n} \sum_{j=1}^n (\pi^*(X_{i-1,n}; \Lambda_{i,j}) \mu_{i,j} - \lambda_2 C_{i,j}).$$

where $C_{i,j} = g(\Lambda_{i,j}) - E g(\Lambda)$ and λ_1, λ_2 are (deterministic) matrices (of row and column dimensions appropriate to make sense of the above expressions). One then executes the algorithm with these modified cuts.

As for importance sampling, the proof of Theorem 1 again goes through with minor changes. The key observation is to note that $z^*(x; \Lambda_{i,j})$ is replaced by

$$z^*(x, \Lambda_{i,j}) - \lambda_1 C_{i,j} x - \lambda_2 C_{i,j}$$

and $\sigma^2(x^*)$ is replaced by

$$\text{var}(z^*(x^*, \Lambda_{i,j}) - \lambda_1 C_{i,j} x^* - \lambda_2 C_{i,j}).$$

When λ_1 and λ_2 are chosen so that $\text{var}(z^*(x^*, \Lambda_{i,j}) - \lambda_1 C_{i,j} x^* - \lambda_2 C_{i,j})$ is smaller than $\sigma^2(x^*)$, the lower bound is tightened (in the sense that the length of the associated confidence interval for z^* is reduced). Note that a variance-minimizing choice of λ_1 and λ_2 (so as to minimize $\text{var}(z^*(x^*, \Lambda_{i,j}) - \lambda_1 C_{i,j} x^* - \lambda_2 C_{i,j})$) requires knowledge of x^* , complicating the statistical estimation of the optimal control coefficient matrices λ_1^* and λ_2^* . This represents an interesting opportunity for future research.

Extension 4: Dropping Cuts

In large-scale problems, it is often the case that the number of iterations required to accurately compute z^* is large. Our iterative algorithm can be slowed significantly when one insists on keeping all m cuts. Instead, one may prefer on iteration m to (deterministically) retain the last j_m cuts (with $1 \leq j_m \leq m$), so that step 6 of our algorithm is implemented by solving

$$\begin{aligned} \min cx &+ \theta \\ Ax &= b \\ -G_{i,n}x + \theta &\geq g_{i,n}, \quad m - j_m + 1 \leq i \leq m \\ x &\geq 0, \end{aligned}$$

for the minimizer $X_{m,n}$. Our proof goes through without change to this setting. Note, however, that if the retained cuts are not deterministically specified (as would be the case if we were to choose to retain the “best” j_m cuts), our proof fails to be valid and a more sophisticated analysis would be required.

4 Numerical Results

We tested the confidence bounds methods described earlier in this paper on the following test problems: APL1P (Infanger (1992)) and PGP2 (Louveaux and Smeers (1988)) are electric power planning problems, CEP1 (Higle and Sen (1994)) is a machine capacity expansion planning problem, STORM (Mulvey and Ruszczyński (1992)) and STORMG2, a variant of STORM, are freight scheduling problems, SCTAP1 (Ho (1980)) is a traffic assignment problem, and SNL and SNS are variants of the problem SSN (Sen, Doverspike and Cosares (1994)), a large telecommunications planning problem. Table 1 summarizes the size characteristics of each of the above problems. The table also includes the optimal objective value z^* for each problem, as computed by solving the problems exactly using decomposition without any sampling. Because our lower bound confidence interval procedure involves bounding the second-stage standard deviation $\sigma^* \equiv \sigma(x^*)$, we also report the exact value of σ^* for each problem.

Problem	Size Master	Size Sub	RV	Scenarios	z^*	σ^*
APL1P	3, 3	6, 9	5	1280	24642.3	4808.8
PGP2	2, 4	4, 16	3	576	447.3	77.6
CEP1	9, 8	7, 15	3	216	355159.5	420458.4
STORM	1290, 122	527, 1259	1	40	15.569×10^6	0.2819×10^6
STORMG2	186, 122	529, 1259	1	1000	15.580×10^6	1.3295×10^6
SCTAP1	31, 49	61, 96		480	248.5	24.01
SNS	2, 8	18, 58	10	2560000	3.6121	1.7863
SNL	2, 8	2, 8	14	40960000	7.7727	9.9475

Table 1: Test data, problem dimensions (rows, columns), number of (independent) random variables, number of universe scenarios, optimal objective z^* and standard deviation of the optimal second-stage cost σ^*

We slightly modified the above problems by adding box constraints (where necessary) to our first-stage feasible region, in particular lower and upper bounds to each first stage variable. As we shall see below, the box bounds both make the first stage region compact (which is convenient theoretically) and permit us to sample points within the first-stage region as required by Method 1 in Section 3. The box bounds were set reasonably wide, in particular, for APL1P and PGP2 the lower/upper box bounds were 0/5000; for CEP1 the bounds were set at 0/1000; for STORM, STORMG2 and SCTAP1, the bounds were 0/100; and for SNS and SNL the lower/upper bounds were 0/10.

In implementing our confidence bounds, we chose estimator (5) to upper bound the second-stage variance $\sigma^2(x^*)$. As just noted, this estimator requires sampling a sequence of points within the first-stage feasible region. There are many such procedures that have been proposed for generating such a sequence; see, for example, Smith (1984).

Our approach starts by sampling a point uniformly within the box. If the sampled point falls within the feasible region $H = \{x : Ax = b, x \geq 0\}$, we add it to the sample of feasible points. If the sampled point turns out infeasible, we project the point onto the feasible

region H by solving the problem:

$$\min \|x - \hat{x}^k\|^2, \text{ s/t } x \in H$$

to obtain as its optimal solution a feasible boundary point $\hat{x}_b^k \in H$. In order to obtain a feasible point in the interior of H , we compute a random convex combination of consecutive feasible boundary points as

$$\hat{x}_{feas}^k = \alpha^k \hat{x}_b^k + (1 - \alpha^k) \hat{x}_b^{k-1},$$

where each α^k is an independently distributed uniform random parameter between 0 and 1 and $\hat{x}_{feas}^k, k = 1, \dots$ are feasible points with respect to H that are added to the sample of feasible points. We call this the *projection algorithm* and used this to generate $r = 30$ and $r = 50$ points within the feasible region. Of course, one obtains different lower bounds for these two different values of the parameter r . We used the projection algorithm rather than Smith (1984) because the projection algorithm could be very easily implemented within the framework of the stochastic programming software (Infanger (1997)) used for obtaining the numerical results of this paper (because solving the projection problem merely involves adding a quadratic term to the objective of the master problem).

The upper bound and lower bound results (for $r = 30$) are provided in Table 2, along with empirical coverages (based on running the sampling-based algorithm on each test problem 100 independent times) for the confidence intervals. The lower and upper bounds are reported as a (signed) percentage of the true objective value, so that a lower bound of -2% means that the computed lower bound was 2% lower than the true value, while an upper bound of 5% means that it was 5% above the true value. We note that the

Model	Iteration	n	Lower Bound		Upper Bound	
			Mean %	Coverage %	Mean %	Coverage %
APL1P	20	100	-3.97	100	3.96	97
PGP2	20	100	-16.57	100	3.50	90
CEP1	6	100	-21.04	99	25.38	92
STORM	20	20	-1.19	100	0.97	100
STORMG2	30	100	-1.63	100	1.92	100
SCTAP1	3	100	-1.60	95	1.66	94
SNS	20	500	-5.23	98	5.58	100
SNL	20	500	-6.65	100	5.67	99

Table 2: Summary Results, lower (for $r = 30$) and upper bound estimates and coverage; based on 100 replications each.

upper and lower confidence bound estimate are reasonably small for all problems, and the coverage results for the lower and upper bound estimates are good. The coverage of the lower bound estimate is nearly always close to 100%, indicating that the estimates of the optimal standard deviation at the optimal solution via the its maximum over the first-stage feasible region is a numerically valid approach.

Table 3 gives the upper bound mean values and coverage results for various sample sizes n . The sample spaces for problems SNS and SNL are rather large, with a total number of universe scenarios of 2.56 million and 40.96 million outcomes, respectively. Both problems have a large standard deviation of the second-stage cost and therefore command a large sample size for obtaining reasonably small confidence bounds. We therefore include results for $n = 500$ as part of the analysis. It is important to note that even when the confidence intervals are wide, the coverage is still (as predicted) around 95%, confirming the theory.

Problem	n	95% upper bound	
		avg value %	coverage %
APL1P	30	6.97	97
	50	4.90	93
	100	3.96	97
PGP2	30	9.38	92
	50	6.52	92
	100	3.50	90
CEP1	30	33.54	92
	50	25.38	92
	100	19.18	96
STORM	20	0.97	100
STORMG2	30	3.32	96
	50	2.36	95
	100	1.92	100
SCTAP1	30	3.20	94
	50	2.71	100
	100	1.66	94
SNS	30	18.33	94
	50	17.24	100
	100	10.49	97
	500	5.58	100
SNL	30	19.37	99
	50	16.28	100
	100	11.94	100
	500	5.67	99

Table 3: Upper bound average value and coverage

Table 4 gives the results for the lower bound estimate for various sample sizes ℓ and for sample sizes $r = 30$ and $r = 50$. The table also presents the average percentage value (over 100 replications) of $(\hat{\sigma}_*(r, \ell) - \sigma^*)/\sigma^*$ in percent of the standard deviation at the optimal solution σ^* as well as the coverage (the fraction of instances where $\hat{\sigma}_*(r, \ell)$ was greater than σ^*).

As expected, the coverage of our upper bound estimate ($\hat{\sigma}_*(r, \ell)$) is almost always 100%. We also note that the lower confidence bound estimate are reasonably tight for all problems, and the coverage results for the lower bound estimates are always close to 100%, for all sample sizes n tested. Using sample size $r = 30$ resulted in good lower bounds for all

Problem	ℓ	$r = 30$				$r = 50$			
		$\hat{\sigma}_*(r, \ell)$		95% lower bound		$\hat{\sigma}_*(r, \ell)$		95% lower bound	
		avg value %	covg %	avg value %	covg %	avg value %	covg %	avg value %	covg %
APL1P	30	27.23	100	-8.02	100	32.73	100	-8.37	100
	50	21.03	100	-5.94	100	23.82	100	-5.98	100
	100	15.43	100	-3.97	100	17.12	100	-4.06	100
PGP2	30	299.14	98	-33.96	100	374.05	99	-40.77	100
	50	294.36	99	-25.66	100	344.40	100	-29.44	100
	100	265.75	100	-16.57	100	321.37	100	-19.46	100
CEP1	30	30.05	100	-41.65	100	32.42	100	-44.04	100
	50	23.65	100	-31.49	99	24.71	100	-32.12	98
	100	16.28	100	-21.04	99	17.98	100	-19.81	98
STORM	20	32.58	100	-1.19	100	33.17	100	-2.22	100
STORMG2	30	15.72	100	-3.08	100	17.68	100	-3.23	100
	50	11.99	100	-2.22	100	13.87	100	-2.46	100
	100	8.48	100	-1.63	100	9.56	100	-1.65	100
SCTAP1	30	23.64	100	-3.12	100	25.60	100	-3.04	98
	50	17.80	100	-2.31	98	19.83	100	-2.27	98
	100	12.87	100	-1.60	95	14.43	100	-1.44	94
SNS	30	28.78	100	-25.21	100	31.69	100	-26.70	100
	50	21.28	100	-18.20	99	23.24	100	-18.69	100
	100	15.12	100	-12.20	99	17.42	100	-13.21	100
	500	6.31	100	-5.23	98	7.56	100	- 5.14	98
SNL	30	25.52	100	-22.78	100	28.27	100	-23.40	100
	50	20.40	100	-17.45	100	23.54	100	-18.39	100
	100	15.18	100	-13.11	100	18.76	100	-13.45	100
	500	9.78	95	-6.65	100	10.59	100	- 6.73	99

Table 4: Average value and coverage of the standard deviation of second-stage cost, lower bound average value and coverage, each for sample sizes $r = 30$ and $r = 50$

problems. As expected, choosing $r = 50$ leads to a larger upper bound estimate for the standard deviation at the optimal solution and thus to a weaker (but more reliable) lower bound estimate.

For problem PGP2 the lower bound results reflect a large variation of the variance of the optimal second-stage cost over different feasible first-stage solutions. For example, using a sample size of $r = 30$, and a sample size of $n = 100$ the estimated maximum second-stage standard deviation is 265.75% larger than the true optimal variance of the second-stage cost, with a coverage of 100%. The corresponding average lower bound is 16.57% smaller than the true objective of the problem, with a coverage of also 100%. The standard deviation of the second-stage cost at the optimal first-stage solution is overestimated by using our procedure. Our computational experience indicates that typically the variance turns out to be large at feasible points having a large cost, at which a significant penalty cost for unserved demand needs to be paid. One could further restrict the feasible region for the first-stage problem by rejecting points that have an expected cost that are, with high probability, larger than

the best upper bound estimate. A statistical test could be used to help to decide which such points to reject. How to best reject based on objective estimates is subject to further research.

Tables 5 and 6 present statistics for the upper and lower bound estimates as a function of the number of iterations. Because this set of experiments was intended to offer insight into the question of iterative convergence of the bounds, we selected STORM2 (due to its large size) and SNL (because of its large sample space) as candidate problems. As one might expect, the results demonstrate that the bounds have excellent coverage characteristics. They also make clear that if the number of iterations is large, eventually the noise associated with the fixed sample size n used at every iteration dominates the accuracy and no further iterative improvement can be gained.

For both problems the upper and lower bounds exhibit excellent coverage at every iteration, with the lower bound coverage almost always at 100%. We note that choosing a reasonable number of iterations is important. For example, the initial choices of 30 and 20 observations, respectively, is not quite sufficient for obtaining the best solution possible.

Iteration	Lower Bound			Upper Bound			Lower Bound	Upper Bound
	95% %	Mean %	Median %	Median %	Mean %	95% %	Coverage %	Coverage %
5	-3.80	-3.00	-2.98	3.10	3.22	4.83	100	100
10	-2.79	-2.15	-2.14	2.04	2.01	3.50	100	98
15	-2.39	-1.88	-1.86	1.96	2.03	3.51	100	100
20	-2.30	-1.73	-1.77	1.95	1.90	3.72	98	98
25	-2.32	-1.70	-1.70	2.06	1.98	3.31	100	99
30	-2.19	-1.63	-1.66	1.98	1.96	3.38	100	100
35	-2.19	-1.62	-1.65	1.72	1.80	3.64	100	99
40	-2.05	-1.60	-1.59	1.84	1.90	3.50	100	100
45	-2.02	-1.57	-1.59	1.79	1.86	3.43	100	98
50	-1.95	-1.52	-1.57	1.58	1.72	3.54	100	98
55	-1.93	-1.47	-1.53	1.96	1.84	3.37	100	98
60	-1.97	-1.51	-1.55	1.77	1.78	3.24	100	100
65	-1.84	-1.46	-1.50	1.77	1.71	3.02	100	98
70	-1.83	-1.43	-1.43	1.88	1.80	3.35	100	98
75	-1.92	-1.40	-1.45	1.74	1.69	2.97	100	96
80	-1.92	-1.40	-1.45	1.74	1.69	2.97	100	96

Table 5: Model STORMG2, statistics for upper bound and lower bound estimates as a function of iterations; sample size $n = 100$, $r = 30$, $\ell = 100$, 100 replications each.

Iteration	Lower Bound			Upper Bound			Lower Bound	Upper Bound
	95% %	Mean %	Median %	Median %	Mean %	95% %	Coverage %	Coverage %
5	-25.84	-23.42	-23.67	43.80	43.56	47.84	100	100
10	-15.09	-11.94	-12.16	5.41	5.90	11.83	100	99
15	-9.93	-8.32	-8.55	5.72	5.55	9.68	100	99
20	-8.15	-6.63	-6.69	5.84	5.98	10.55	100	99
25	-6.67	-5.58	-5.60	5.30	5.28	9.74	100	100
30	-5.86	-5.03	-4.99	4.78	5.07	8.88	100	100
35	-5.52	-4.63	-4.76	5.16	4.98	8.26	100	99
40	-5.39	-4.34	-4.33	4.82	5.00	8.56	100	100
45	-4.98	-4.25	-4.34	4.46	4.47	7.23	100	99
50	-5.07	-4.07	-4.09	4.36	4.48	8.79	100	100
55	-4.76	-3.92	-3.89	4.44	4.49	7.68	100	98
60	-4.72	-3.85	-3.87	4.14	4.41	8.29	100	100
65	-4.54	-3.77	-3.83	4.48	4.72	8.42	100	98
70	-4.51	-3.74	-3.79	4.45	4.44	8.22	100	98

Table 6: Model SNL, statistics for upper bound and lower bound estimates as a function of iterations; sample size $n = 500$, $r = 30$, $\ell = 500$, 100 replications each.

References

- [1] Asmussen, S. and P.W. Glynn (2007): *Stochastic Simulation: Algorithms and Analysis*, Springer Verlag, New York.
- [2] Dantzig, G.B. and P.W. Glynn (1990): Parallel processors for planning under uncertainty, *Annals of Operations Research* 22, 1–21.
- [3] Dantzig, G.B. and G. Infanger (1995): A probabilistic lower bound for two-stage stochastic programs, Report SOL 95-6, Dept of Operations Research, Stanford University, CA.
- [4] Feller, W. (1971): *An Introduction to Probability Theory and Its Applications*, Volume 2, John Wiley and Sons, Inc.
- [5] Hagle, J. L. and S. Sen (1994). Finite master programs in stochastic decomposition, *Mathematical Programming*, 67, pp. 143–168.
- [6] HO, J.K. (1980): A successive linear optimization approach to the dynamic traffic assignment problem, *Transportation Science*, 14(4), pp. 295–305.
- [7] Infanger, G. (1992): Monte Carlo (importance) sampling within a Benders decomposition algorithm for stochastic linear programs, *Annals of Operations Research*, 39, pp. 41–67.
- [8] Infanger, G. (1997): DECIS User’s Guide, Infanger Investment Technology, LLC, Mountain View, CA.
- [9] Infanger, G. and D. Morton (1996): Cut sharing for multistage stochastic linear programs with inter-stage dependency, *Mathematical Programming*, 75, pp. 241–256.
- [10] Kall, P. (1976): *Stochastic Linear Programming*, Springer Verlag, Berlin.
- [11] Louveaux, F.V. and Y. Smeers (1988): Optimal investment for electricity generation: A stochastic model and a test problem, in Ermoliev, Y. and Wets, R.J.-B. (eds.): *Numerical Techniques for Stochastic Optimization*, Springer Verlag, Berlin.
- [12] 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, pp. 47–56.
- [13] Mulvey, J. and A. Ruszczyński (1995) A new scenario decomposition method for large scale stochastic optimization, *Operations Research*, 43, pp. 477–490.
- [14] Shapiro, A. and T. Homem-de-Mello (1998): A simulation-based approach to two-stage stochastic programs with recourse, *Mathematical Programming*, 81, pp. 301–325.
- [15] Sen S., R. D. Doverspike, and S. Cosares (1994): Network planning with random demand, *Telecommunication Systems*, 3 pp. 11–30.
- [16] Smith, R.L. (1984): Efficient Monte Carlo procedures for generating points uniformly distributed over bounded regions, *Operations Research* 32, 1296–1308.