

Jeff Linderoth* · Alexander Shapiro** · Stephen Wright***

The Empirical Behavior of Sampling Methods for Stochastic Programming

September 30, 2002

Abstract. We investigate the quality of solutions obtained from sample-average approximations to two-stage stochastic linear programs with recourse. We use a recently developed software tool executing on a computational grid to solve many large instances of these problems, allowing us to obtain high-quality solutions and to verify optimality and near-optimality of the computed solutions in various ways.

1. Introduction

Consider the following stochastic programming problem

$$\min_{x \in X} \left\{ f(x) := \mathbb{E}_P[F(x, \xi(\omega))] \right\}, \quad (1)$$

where $F(x, \xi)$ is a real valued function of two vector variables $x \in \mathbb{R}^n$ and $\xi \in \mathbb{R}^d$, X is a subset of \mathbb{R}^n , and $\xi(\omega)$ is viewed as a random vector having probability distribution P , which is assumed to be known. Often we omit the subscript P and simply write $\mathbb{E}[F(x, \xi(\omega))]$. In this paper, we discuss sampling techniques for obtaining and verifying approximate solutions of (1), and we present results for several standard problems from the literature. Some of the ideas presented

Jeff Linderoth: Industrial and Systems Engineering Department, Lehigh University, 200 West Packer Avenue, Bethlehem, PA 18015; jt13@lehigh.edu

Alexander Shapiro: School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA 30332; ashapiro@isye.gatech.edu

Stephen J. Wright: Computer Sciences Department, University of Wisconsin-Madison, 1210 West Dayton Street, Madison, WI 53706; swright@cs.wisc.edu

* Research supported by the Mathematical, Information, and Computational Sciences Division subprogram of the Office of Advanced Scientific Computing Research, U.S. Department of Energy, under Contract W-31-109-Eng-38, and by the National Science Foundation under Grant 9726385.

** Research supported by the Mathematical, Information, and Computational Sciences Division subprogram of the Office of Advanced Scientific Computing Research, U.S. Department of Energy, under Contract W-31-109-Eng-38, and by the National Science Foundation under Grant DMS-0073770.

*** Research supported by the Mathematical, Information, and Computational Sciences Division subprogram of the Office of Advanced Scientific Computing Research, U.S. Department of Energy, under Contract W-31-109-Eng-38, and by the National Science Foundation under Grants 9726385 and 0082065.

here are applicable to the general form (1), while others apply specifically to two-stage stochastic linear programs with recourse.

The two-stage stochastic linear program with recourse can be written in standard form as follows:

$$\min_x c^T x + Q(x) \text{ subject to } Ax = b, x \geq 0, \quad (2)$$

where $Q(x) := \mathbb{E}[Q(x, \xi(\omega))]$ and $Q(x, \xi)$ is the optimal value of the second-stage problem

$$\min_y q^T y \text{ subject to } Tx + Wy = h, y \geq 0. \quad (3)$$

Here some or all elements of the data vector $\xi = (q, h, T, W)$ can be random (that is, can depend on ω). In problems with *fixed recourse*, the matrix W is fixed (deterministic). If we define

$$F(x, \xi) := c^T x + Q(x, \xi), \quad X = \{x : Ax = b, x \geq 0\},$$

and assume that the optimal value of the second-stage problem (3) is finite for all $x \in X$ and almost every ξ , then problem (2) becomes a particular case of problem (1).

We report computational results in this paper for problems of the form (2), (3) with fixed recourse. We consider problems where the probability distribution of $\xi(\omega)$ has *finite support*; that is, $\xi(\omega)$ has a finite number of possible realizations, called *scenarios* $\{\xi_1, \xi_2, \dots, \xi_K\}$ with respective probabilities $p_k \in (0, 1)$, $k = 1, 2, \dots, K$. For such problems, the expected value function $Q(x)$ can be written as the finite sum

$$Q(x) = \sum_{k=1}^K p_k Q(x, \xi_k), \quad (4)$$

where $Q(x, \xi)$ is once again the optimal value of (3) and ξ_k denotes the data vector (q_k, h_k, T_k, W) , $k = 1, 2, \dots, K$. By defining y_k to be the instance of the second-stage variable vector corresponding to scenario k , we can combine (2), (3) into a single linear program as follows:

$$\min_{x, y_1, y_2, \dots, y_K} c^T x + \sum_{k=1}^K q_k^T y_k \text{ subject to} \quad (5a)$$

$$Ax = b, \quad x \geq 0, \quad (5b)$$

$$Wy_k = h_k - T_k x, \quad y_k \geq 0, \quad k = 1, 2, \dots, K. \quad (5c)$$

In many applications, K is an astronomically large number. This situation can arise when each component of the random data vector ξ has a number of possible values, independent of the other components. The total number of scenarios is then the product of the number of possible values for each component of ξ . For example, if ξ has 100 components, each of which has two possible values, then the total number of scenarios is 2^{100} , or approximately 10^{30} . In such situations, it is not practical to solve (2), (3), or equivalently (5), directly. It is not even possible to evaluate the function $Q(x)$, for a given x , by using

formula (4) in a reasonable amount of time, since the number of second-stage linear programs (3) to be solved is just too large.

We can however use sampling techniques, which consider only randomly selected subsets of the set $\{\xi_1, \xi_2, \dots, \xi_K\}$, to obtain approximate solutions. Sampling techniques can be applied in different ways. One approach uses sampling in an “interior” fashion. Such algorithms aim to solve the original problem (1), but resort to sampling whenever the algorithm requires an (approximate) value of $f(\cdot)$ or subgradient information for $f(\cdot)$ at some point x . Typically, a different sample is used each time function or subgradient information is required. We mention in particular the L-shaped method with embedded sampling of Dantzig and Infanger [18], the stochastic decomposition method of Hige and Sen [15, 17], and stochastic quasi-gradient methods (see Ermoliev [8]). A second fundamental approach to sampling is an “exterior” approach, in which a sample is selected from $\{\xi_1, \xi_2, \dots, \xi_K\}$ and a corresponding approximation to $f(\cdot)$ is defined from this sample. This approximate objective, known as a *sample-average approximation* of $f(\cdot)$, is then minimized using a deterministic optimization algorithm; no further sampling is performed. This approach is known variously as sample-path optimization [13, 29] or the stochastic counterpart method [31, 32].

In scenario-based approaches, each scenario can often (for portions of the algorithm) be considered independently from other scenarios. This makes such approaches well-suited to implementation on parallel computing platforms, where different processors execute the independent portions of the computation in parallel. The more independent the computational components of a solution algorithm, the more able it is to benefit from parallel computing.

Our computational experiments are performed on an unusual platform—a large federation of nondedicated, heterogeneous, loosely coupled processors, known as a *computational grid* or a *metacomputer* [9, 22]. The grid represents an extremely powerful computing resource that can provide large amounts of computing power over extended time periods, yet is also inexpensive, especially since we use only idle time on many of the workstations and PCs that make up the grid. We needed the computational power of the grid to determine optimal and near-optimal solutions to two-stage stochastic linear programs with recourse, and to test the convergence behavior of sample average approximations of $f(\cdot)$. Zakeri [39] has also used a computational grid to verify the quality of a given solution to a stochastic programming instance. Other recent use of computational grids to solve large, complex optimization problems have been described by Anstreicher et al. [1], Goux and Leyffer [11], and Chen, Ferris, and Linderoth [4].

The remainder of the paper is structured as follows. Section 2 defines the construction of sample-average approximations using different sampling techniques. In Section 3, we summarize results from Shapiro and Homem-de-Mello [35] concerning convergence of the minimizer of the sample-average approximation to the exact minimizer of the original problem under certain circumstances. In Section 4, we describe techniques for estimating lower and upper bounds on the optimal value of the problem (1), and for estimating the gap between these bounds. We describe several tests for optimality or near-optimality of a candi-

date solution in Section 5. In Section 6 we focus on the two-stage stochastic linear programming problem with recourse (2), outlining the software that we used to solve this problem. (A fuller description can be found in [21].) Section 7 briefly highlights the characteristics of the computational grid we used in this work. Section 8 describes the benchmark problems used in our computational tests, while Section 9 describes and analyzes the results obtained from our tests.

2. Sample-Average Approximations

Suppose that a random sample $\xi^1, \xi^2, \dots, \xi^N \sim P$ of N (not necessarily independent) realizations of the random vector $\xi(\omega)$ is generated. By replacing $f(\cdot)$ by an approximation based on this sample, we approximate the problem (1) by

$$\min_{x \in X} \left\{ \widehat{f}_N(x) := N^{-1} \sum_{i=1}^N F(x, \xi^i) \right\}. \quad (6)$$

The function \widehat{f}_N is a sample-average approximation (SAA) to the objective f of (1). We denote by \widehat{v}_N the optimal value and by \widehat{x}_N an optimal solution of (6).

Since the random realizations ξ^i have the same probability distribution P , it follows that $\widehat{f}_N(x)$ is an unbiased estimator of $f(x)$, for any x . If we perform Monte Carlo sampling, in which all ξ^i are independent of each other, then $\text{Var}[\widehat{f}_N(x)] = \sigma^2(x)/N$, where $\sigma^2(x) := \text{Var}[F(x, \xi(\omega))]$.

We also used Latin Hypercube sampling (LHS) to construct the sample-average function \widehat{f}_N . In this technique, initially proposed by McKay, Beckman, and Conover [25], a sample of size N is constructed by dividing the interval $(0, 1)$ into N subintervals of equal size and picking one number randomly from each subinterval. These N numbers are then shuffled, and the resulting sequence is used to generate random variates for a given distribution, possibly by performing an inverse transform. Assuming that the components of the random vector are independently distributed, the procedure is repeated for each component, yielding a stratified sample of size N for that vector. The sample-average function \widehat{f}_N so obtained is an unbiased estimator of f , but its variance $\text{Var}[\widehat{f}_N(x)]$ often is considerably less than that obtained from Monte Carlo sampling. Further discussion of LHS and its properties can be found in Avramidis and Wilson [2, Section 1.2.2], Stein [38], and Owen [28]. In stochastic programming, LHS methodology was used in Diwekar and Kalagnanam [7] and Bailey, Jensen and Morton [3]. Some other variance reduction techniques have been described in Hight [14]. Importance sampling variance reduction method has been used by Dantzig and Infanger [5].

From a computational point of view, one can regard the SAA problem (6) as a stochastic programming problem with a finite number of scenarios $\xi^1, \xi^2, \dots, \xi^N$ each with equal probability $p_i = N^{-1}$. Therefore, any numerical algorithm suitable for solving the corresponding stochastic problem with a discrete distribution can be applied to the SAA problem.

3. Convergence of the SAA Solution to the Exact Solution Set

A fundamental question central to any sampling-based solution approach is how large must the sample size be so that the solution to the sampled instance is a good (or optimal) solution for the true instance. In this section, we review a theory, developed recently in [35,20,37], concerning the rate of convergence of the solutions \hat{x}_N of the SAA problem (6) to the optimal solution set of the original problem (1). This theory suggests that for certain instances, the sample size necessary to obtain excellent solutions is small compared to the size of the whole sample space.

Unless stated otherwise, we make the following assumptions on the problem (1) throughout the remainder of the paper:

- (A1) The distribution of $\xi(\omega)$ has finite support, say $\Xi = \{\xi_1, \xi_2, \dots, \xi_K\}$;
- (A2) For every $\xi \in \Xi$, the function $F(\cdot, \xi) : X \rightarrow \mathbb{R}$ is convex and piecewise linear;
- (A3) The set X is polyhedral;
- (A4) The set S of optimal solutions of (1) is nonempty and bounded.

Assumption (A1) means that we deal with random data having a finite, although possibly very large, number of realizations (scenarios). In the case of two-stage linear programs of the form (2), (3) we have that $F(x, \xi) < +\infty$ if the corresponding second stage problem is feasible. We also have that $F(\cdot, \xi)$ is convex and piecewise linear, and hence $F(\cdot, \xi)$ is bounded from below if $F(x, \xi) > -\infty$ for at least one x . Therefore, assumption (A2) is satisfied by stochastic two-stage linear programs provided that the recourse is relatively complete, that is, for every $x \in X$ and $\xi \in \Xi$, the second stage problem (3) is feasible. For the problems used in our experiments, the set X is defined by linear constraints, and hence is polyhedral, and consequently assumption (A3) holds. For the considered problems, the final assumption (A4) is satisfied as well. Assumptions (A1) and (A2) imply that the expected value function $f(\cdot)$ is finite valued, piecewise linear and convex on X . These assumptions together with (A3) imply that the set S is convex and polyhedral.

For $\varepsilon \geq 0$ we denote by S_ε the set of all ε -optimal solutions of (1), that is, the set of points $\bar{x} \in X$ such that $f(\bar{x}) \leq v^* + \varepsilon$, where v^* is the optimal objective value in (1). In particular, for $\varepsilon = 0$ the set S_ε coincides with the set S of optimal solutions of (1).

Under assumptions (A1)–(A4), we have that for each $\varepsilon \geq 0$ there exist constants $C_\varepsilon > 0$ and $\beta_\varepsilon > 0$ such that the inequality

$$1 - P(\hat{x}_N \in S_\varepsilon) \leq C_\varepsilon e^{-\beta_\varepsilon N} \quad (7)$$

holds for all $N \in \mathbb{N}$. Note that the choice of the optimal solution \hat{x}_N of (6) here is arbitrary, so that the probability $P(\hat{x}_N \in S_\varepsilon)$ actually refers to the probability that the set of optimal solutions of the SAA problem (6) is included in the set S_ε . It follows from (7) that

$$\limsup_{N \rightarrow \infty} \frac{1}{N} \log [1 - P(\hat{x}_N \in S_\varepsilon)] \leq -\beta_\varepsilon. \quad (8)$$

In other words, the probability that \hat{x}_N is an ε -optimal solution of the true (expected value) problem (1) approaches 1 exponentially fast in N . The above result (7), and hence (8), is a consequence of the upper bound of Cramér's Large Deviation Theorem and the assumed polyhedral structure of the problem (cf., [35, Theorem 3.2] and [20, Section 2.2]). By using the lower bound of Cramér's Large Deviation Theorem, it is possible to show that a limit in the left hand side of (8) is attained, so that "lim sup" can be replaced by "lim".

The estimates (7) and (8) hold in particular for $\varepsilon = 0$, that is, for $S_\varepsilon = S$. Of course, the probability $P(\hat{x}_N \in S_\varepsilon)$ increases with ε , and hence for larger values of ε the exponential constant β_ε is larger. In the case in which the solution to (1) is a singleton (that is, $S = \{x_0\}$ for some $x_0 \in X$), there exists $\beta > 0$ such that

$$\lim_{N \rightarrow \infty} \frac{1}{N} \log [1 - P(\hat{x}_N = x_0)] = -\beta. \quad (9)$$

We can refine the expression (9) by quantifying the conditioning of the singleton solution. Denote by $T_X(x_0)$ the tangent cone to the feasible set X at x_0 and by $F'_\omega(x_0, d)$ and $f'(x_0, d)$ the directional derivatives of $F(\cdot, \omega)$ and $f(\cdot)$, respectively, at x in the direction d . Because of the assumed polyhedral structure of problem (1) and since the optimal solution x_0 is unique, we have by the theory of linear programming that it is *sharp*, that is

$$f'(x_0, d) > 0, \quad \forall d \in T_X(x_0) \setminus \{0\}. \quad (10)$$

Furthermore, from [35], we have the following property: There exists a *finite* set $\Delta \subset T_X(x_0) \setminus \{0\}$ of directions, independent of the sample, such that if $\hat{f}'_N(x_0, d) > 0$ for every $d \in \Delta$, then $\hat{x}_N = x_0$. We call

$$\kappa := \max_{d \in \Delta} \frac{\text{Var} [F'_\omega(x_0, d)]}{[f'(x_0, d)]^2} \quad (11)$$

the *condition number* of the true problem (1). The above definition is motivated by the result which states that the sample size N required to achieve a given probability of the event " $\hat{x}_N = x_0$ " is roughly proportional to the condition number κ . Specifically, for large N , we have from [37, Theorem 1] that the probability of this event can be approximated as follows:

$$P(\hat{x}_N = x_0) \approx 1 - \frac{\nu e^{-N/(2\kappa)}}{\sqrt{4\pi N/(2\kappa)}}, \quad (12)$$

where $\nu \geq 1$ is a small constant depending on the problem. It follows by comparing this expression with (9) that for large κ , the exponential constant β is small, and therefore that one needs a large sample to attain a given probability of the event " $\hat{x}_N = x_0$ ". Such problems can be viewed as ill conditioned. Note that for a given d , the expression in the right hand side of (11) is the ratio of the variance of $F'_\omega(x_0, d)$ to its squared expected value.

It is impractical to try to compute the numerical value of the condition number κ . (For one thing, it depends on the optimal solution x_0 which, of course,

is not known a priori.) Rather, the value of the above analysis is conceptual. It indicates that instances that possess a sharp optimal solution will be well-conditioned, and can be solved exactly with a small sample size. In other problems, the objective function is flat near the minimizer. These cases correspond to a very large value of κ from (11). The value of κ may even be infinite, since the directional derivative of $f(\cdot)$ may be zero in some directions. (This happens if the set of optimal solutions of the true problem is not a singleton.) For these instances, the sample size necessary to obtain a good solution to the true instance could be quite large. One of the goals of this paper is to test the convergence behavior of the sample average approximation on a number of real stochastic programming instances to empirically verify if this conceptual analysis extends to practice.

4. Estimating the Optimal Value

Although the results discussed in the previous section indicate a high probability of identifying an exact optimal solution of (1) under certain circumstances, they have little to say about identifying the optimal objective value v^* . Even when N is large and \hat{x}_N is identical to x^* for different samples of size N , the corresponding optimal objective values \hat{v}_N of (6) will in general be different for each sample. Further, the expected value of \hat{v}_N is an underestimate of v^* , as we discuss in Section 4.1.

Sampling methodology can be used to obtain estimates of upper and lower bounds on the optimal value v^* of the true problem (1). In this section, we give details on obtaining these estimates (with approximate confidence intervals).

4.1. Lower Bound Estimates

Denote by v^* and \hat{v}_N the optimal values of the true (1) and SAA (6) problems, respectively. It is well known, and can be easily shown, that

$$\mathbb{E}[\hat{v}_N] \leq v^*. \quad (13)$$

The expected value $\mathbb{E}[\hat{v}_N]$ can be estimated as follows. Generate M independent samples, $\xi^{1,j}, \dots, \xi^{N,j}$, $j = 1, \dots, M$, each of size N , and solve the corresponding SAA problems

$$\min_{x \in X} \left\{ \hat{f}_N^j(x) := N^{-1} \sum_{i=1}^N F(x, \xi^{i,j}) \right\}, \quad (14)$$

for each $j = 1, \dots, M$. Let \hat{v}_N^j be the optimal value of problem (14), and compute

$$L_{N,M} := \frac{1}{M} \sum_{j=1}^M \hat{v}_N^j. \quad (15)$$

The estimate $L_{N,M}$ is an unbiased estimate of $\mathbb{E}[\widehat{v}_N]$. Because of the property (13), it provides a statistical lower bound for the true optimal value v^* .

When the M batches $\xi^{1,j}, \xi^{2,j}, \dots, \xi^{N,j}$, $j = 1, \dots, M$, are i.i.d. (although the elements *within* each batch do not need to be i.i.d.; they can for instance be obtained by LHS), we have by the Central Limit Theorem that

$$\sqrt{M} [L_{N,M} - \mathbb{E}[\widehat{v}_N]] \Rightarrow N(0, \sigma_L^2), \text{ as } M \rightarrow \infty, \quad (16)$$

where $\sigma_L^2 := \text{Var}[\widehat{v}_N]$ and “ \Rightarrow ” denotes convergence in distribution. The sample variance estimator of σ_L^2 is

$$s_L^2(M) := \frac{1}{M-1} \sum_{j=1}^M \left(\widehat{v}_N^j - L_{N,M} \right)^2. \quad (17)$$

Defining z_α to satisfy $P\{N(0,1) \leq z_\alpha\} = 1 - \alpha$, and replacing σ_L by $s_L(M)$, we can obtain an approximate $(1 - \alpha)$ -confidence interval for $\mathbb{E}[\widehat{v}_N]$ to be

$$\left[L_{N,M} - \frac{z_{\alpha/2} s_L(M)}{\sqrt{M}}, L_{N,M} + \frac{z_{\alpha/2} s_L(M)}{\sqrt{M}} \right]. \quad (18)$$

Note that the above confidence interval is based on the asymptotic result (16), and therefore is approximate. For small values of M , one can use $t_{\alpha/2, M-1}$ critical values instead of $z_{\alpha/2}$, which will produce a slightly bigger confidence intervals.

The idea of the statistical lower bound (15), and statistical upper bounds discussed in the next section, was introduced by Norikin, Pflug, and Ruszczyński [27] and developed further by Mak, Morton, and Wood [24].

4.2. Upper Bound Estimates

An upper bound can be obtained by noting that for any $\hat{x} \in X$, we have immediately from (1) that $f(\hat{x}) \geq v^*$. Hence, by choosing \hat{x} to be a near-optimal solution (possibly obtained by solving an SAA problem (6)), and by using some unbiased estimator of $f(\hat{x})$, we can obtain an estimate of an upper bound for v^* . In particular, we can generate T independent batches of samples of size \bar{N} , denoted by $\xi^{1,j}, \xi^{2,j}, \dots, \xi^{\bar{N},j}$, $j = 1, 2, \dots, T$, where each batch has the unbiased property, namely

$$\mathbb{E} \left[\widehat{f}_N^j(x) := \bar{N}^{-1} \sum_{i=1}^{\bar{N}} F(x, \xi^{i,j}) \right] = f(x), \text{ for all } x \in X. \quad (19)$$

We can then use the average value defined by

$$U_{\bar{N}, T}(\hat{x}) := T^{-1} \sum_{j=1}^T \widehat{f}_N^j(\hat{x}), \quad (20)$$

as an estimate of $f(\hat{x})$. By applying the Central Limit Theorem again, we have that

$$\sqrt{T} [U_{\bar{N},T}(\hat{x}) - f(\hat{x})] \Rightarrow N(0, \sigma_U^2(\hat{x})), \text{ as } T \rightarrow \infty, \quad (21)$$

where $\sigma_U^2(\hat{x}) := \text{Var} [\widehat{f}_{\bar{N}}(\hat{x})]$. We can estimate $\sigma_U^2(\hat{x})$ by the sample variance estimator $s_U^2(\hat{x}; T)$ defined by

$$s_U^2(\hat{x}; T) := \frac{1}{T-1} \sum_{j=1}^T [\widehat{f}_{\bar{N}}^j(\hat{x}) - U_{\bar{N},T}(\hat{x})]^2. \quad (22)$$

By replacing $\sigma_U^2(\hat{x})$ with $s_U^2(\hat{x}; T)$, we can proceed as above to obtain a $(1 - \alpha)$ -confidence interval for $f(\hat{x})$:

$$\left[U_{\bar{N},T}(\hat{x}) - \frac{z_{\alpha/2} s_U(\hat{x}; T)}{\sqrt{T}}, U_{\bar{N},T}(\hat{x}) + \frac{z_{\alpha/2} s_U(\hat{x}; T)}{\sqrt{T}} \right]. \quad (23)$$

4.3. Estimating the Gap

For a given feasible solution $\hat{x} \in X$ we may wish to estimate the optimality gap $f(\hat{x}) - v^*$. Consider the difference

$$\text{Gap}_{N,M,\bar{N},T}(\hat{x}) := U_{\bar{N},T}(\hat{x}) - L_{N,M}, \quad (24)$$

between the upper and lower bound estimates, defined in (20) and (15), respectively. By the Law of Large Numbers we have that $U_{\bar{N},T}(\hat{x})$ tends to $f(\hat{x})$ with probability one as $\bar{N}T$ tends to ∞ , and $L_{N,M}$ tends to $\mathbb{E}[\widehat{v}_N]$ with probability one as M tends to ∞ . Moreover, under the assumptions (A1), (A2) and (A4), $\mathbb{E}[\widehat{v}_N]$ tends to v^* with probability one as N tends to ∞ . It follows that $\text{Gap}_{N,M,\bar{N},T}(\hat{x})$ tends to $f(\hat{x}) - v^*$ with probability one as N , M , \bar{N} , and T all tend to ∞ . Of course, if \hat{x} is not an optimal solution, then $f(\hat{x}) - v^*$ is strictly positive. Three factors contribute to the error in the statistical estimator $\text{Gap}_{N,M,\bar{N},T}(\hat{x})$ of the gap $f(\hat{x}) - v^*$. These factors are

- (i) variance of $U_{\bar{N},T}(\hat{x})$;
- (ii) variance of $L_{N,M}$;
- (iii) bias $v^* - \mathbb{E}[\widehat{v}_N]$.

Note that $U_{\bar{N},T}(\hat{x})$ and $L_{N,M}$ are unbiased estimators of $f(\hat{x})$ and $\mathbb{E}[\widehat{v}_N]$, respectively. Variances of these estimators can be estimated from the generated samples and may be reduced by increasing the respective sample sizes \bar{N} , M and T . We have here that $\text{Gap}_{N,M,\bar{N},T}(\hat{x})$ is an unbiased estimator of $f(\hat{x}) - \mathbb{E}[\widehat{v}_N]$, and that

$$f(\hat{x}) - \mathbb{E}[\widehat{v}_N] \geq f(\hat{x}) - v^*.$$

That is, $\text{Gap}_{N,M,\bar{N},T}(\hat{x})$ overestimates the true gap $f(\hat{x}) - v^*$, and has bias $v^* - \mathbb{E}[\widehat{v}_N]$.

The bias constitutes the most serious problem in estimating the $\text{Gap}_{N,M,\bar{N},T}(\hat{x})$. We have by the Central Limit Theorem that $N^{1/2}(\widehat{f}_N(x) - f(x))$ converges in

distribution to $Y(x)$, where each $Y(x)$ is a random variable having normal distribution with mean zero and the same covariance structure as $F(x, \xi(\omega))$. Furthermore, it is known that $N^{1/2}(\widehat{v}_N - v^*)$ converges in distribution to $\min_{x \in S} Y(x)$ (see Shapiro [34]). Under mild additional conditions, we have that the expected value of $N^{1/2}(\widehat{v}_N - v^*)$ converges to the expected value of $\min_{x \in S} Y(x)$. Although each $Y(x)$ has mean (expected value) zero, the mean of the minimum of these random variables could be negative and would tend to be smaller for problems with a larger set S of optimal solutions. For a finite sample size N , one may expect that problems having a large set of nearly optimal solutions will exhibit a similar behavior. That is, for ill conditioned problems, this bias may be relatively large and tends to zero at a rate of $O(N^{-1/2})$. The bias can be reduced by increasing the sample size N of the corresponding SAA problems or by sampling more intelligently (by using LHS, say). Of course, an increase in N leads to a larger problem instance to be solved, while increases in \bar{N} , M and T to reduce components (i) and (ii) of the error lead only to more instances of the same size to be solved. (See [20, Section 3.3] for a further discussion of the bias problem.)

5. Testing for (Approximate) Optimality

5.1. Testing for a Unique Solution with Repeated Independent Samples

Perhaps the simplest scheme for optimality verification is to ascertain whether a phenomenon predicted by the results discussed in Section 3 occurs; that is, to see if the solutions of (6) for a number of different samples of a given size N yield identical solutions \widehat{x}_N . For stochastic two-stage linear programs with unique optimal solution, we would expect to see the same minimizer \widehat{x}_N for almost all samples, for sufficiently large N . Although this test will not in general yield coincident values of \widehat{x}_N if the solution set S is not a singleton, or if the problem is poorly conditioned and does not allow identification of the true solution for any reasonable value of N , we can still sometimes use the minimizers \widehat{x}_N to learn about the dimensionality of S and its diameter.

Our implementation of this test is as follows. We choose N , and solve a single SAA problem (6) for this N , using a very tight convergence tolerance. We use the solution \widehat{x}_N so calculated as a starting point for solving a number (say, 10) of SAAs for the same N , with different samples, again using a tight convergence tolerance. If all these SAAs terminate without moving away from the starting point \widehat{x}_N , we conclude that it is likely that \widehat{x}_N coincides with x^* . If not, we repeat the entire process for a larger value of N .

5.2. Comparing Several Points with Repeated Independent Samples

Given two candidates $x^0, x^1 \in X$ for an optimal solution of the problem (1), we can validate whether one of these points has a significantly smaller value

of the objective function by using the statistical (paired) t -test. We generate T independent batches of samples of size \bar{N} , denoted by $\xi^{1,j}, \xi^{2,j}, \dots, \xi^{\bar{N},j}$, $j = 1, 2, \dots, T$, where each batch has the unbiased property (19). We have from this property that

$$\mathbb{E} \left[\widehat{f}_{\bar{N}}^j(x^1) - \widehat{f}_{\bar{N}}^j(x^0) \right] = \mathbb{E} \left[\bar{N}^{-1} \sum_{i=1}^{\bar{N}} \left(F(x^1, \xi^{i,j}) - F(x^0, \xi^{i,j}) \right) \right] = f(x^1) - f(x^0),$$

so the average value defined by

$$D_{\bar{N},T}(x^0, x^1) := T^{-1} \sum_{j=1}^T \left[\widehat{f}_{\bar{N}}^j(x^1) - \widehat{f}_{\bar{N}}^j(x^0) \right] \quad (25)$$

can be used as an estimate of $f(x^1) - f(x^0)$. By applying the Central Limit Theorem again, we have that

$$\sqrt{T} \left[D_{\bar{N},T}(x^0, x^1) - (f(x^1) - f(x^0)) \right] \Rightarrow N(0, \sigma_D^2(x^0, x^1)), \text{ as } T \rightarrow \infty, \quad (26)$$

where $\sigma_D^2(x^0, x^1) := \text{Var} \left[\widehat{f}_{\bar{N}}^j(x^1) - \widehat{f}_{\bar{N}}^j(x^0) \right]$. We can use the sample variance estimator $s_D^2(x^0, x^1; T)$ defined by

$$s_D^2(x^0, x^1; T) := \frac{1}{T-1} \sum_{j=1}^T \left[\widehat{f}_{\bar{N}}^j(x^1) - \widehat{f}_{\bar{N}}^j(x^0) - D_{\bar{N},T}(x^0, x^1) \right]^2$$

to construct a $(1 - \alpha)$ -confidence interval for $f(x^1) - f(x^0)$, as follows:

$$\left[D_{\bar{N},T}(x^0, x^1) - \frac{z_{\alpha/2} s_D(x^0, x^1; T)}{\sqrt{T}}, D_{\bar{N},T}(x^0, x^1) + \frac{z_{\alpha/2} s_D(x^0, x^1; T)}{\sqrt{T}} \right]. \quad (27)$$

Mak, Morton, and Wood [24] use a common-random numbers technique like that described above in connection with estimating the optimality gap defined in Section 4.3.

5.3. Testing KKT Conditions I

In discussing the tests of optimality in this section, we assume that the feasible set X for the problem (1) and the SAA problem (6) is defined by linear constraints, that is,

$$X := \{x \in \mathbb{R}^n : a_i^T x + b_i \leq 0, i \in \mathcal{I}\}, \quad (28)$$

where \mathcal{I} is a finite index set. It follows that X is polyhedral.

Another test for near-optimality of a candidate solution to (1) is based on statistical verification of the first-order Karush-Kuhn-Tucker (KKT) optimality conditions, using subgradients from the sample-average approximations (6) to build up an approximation to the subdifferential of $f(\cdot)$ in the vicinity of the

candidate solution \hat{x} . Hige and Sen [16, Section 3] discuss a test that is related to the one described below. It differs in that it does not attempt to calculate explicitly the distance between the subdifferential of f and the normal cone to X at \hat{x} , but rather uses a kind of sampling technique to estimate this distance.

We discuss first the situation in which $f(\cdot)$ is differentiable at \hat{x} ; that is, its subdifferential $\partial f(\hat{x})$, at \hat{x} , is a singleton. In the case of two-stage linear stochastic programming the (optimal value) function $Q(\cdot, \xi)$ is convex, and since it is assumed that the corresponding expected value function is real valued, the subdifferential of $f(x)$ can be taken inside the expectation operator (see, e.g., Ioffe and Tihomirov [19, Theorem 4, page 381]). It follows that $\partial f(\hat{x})$ is a singleton if $\partial Q(\hat{x}, \xi(\omega))$ is a singleton with probability one. This typically holds if the involved probability distributions are continuous. On the other hand if the number of scenarios is finite, then the expected value function is piecewise linear and cannot be differentiable everywhere except in the trivial case in which it is affine. Nevertheless, it may still be differentiable at the considered point \hat{x} .

Assuming that $F(\cdot, \xi(\omega))$ is differentiable at a given point x with probability one, we can (under mild regularity conditions and, in particular, if $F(\cdot, \xi(\omega))$ is convex and $f(\cdot)$ real valued) take the derivatives inside the expected value, that is,

$$\nabla f(x) = \mathbb{E}[\nabla_x F(x, \xi(\omega))]. \quad (29)$$

The following first-order optimality conditions hold at an optimal solution x^* of the true problem (1). There exist Lagrange multipliers $\lambda_i \geq 0$, $i \in \mathcal{I}(x^*)$, such that

$$0 \in \partial f(x^*) + \sum_{i \in \mathcal{I}(x^*)} \lambda_i a_i, \quad (30)$$

where

$$\mathcal{I}(x^*) := \{i \in \mathcal{I} : a_i^T x^* + b_i = 0\}$$

is the index set of constraints active at x^* . In particular, if $f(\cdot)$ is differentiable at x^* , that is, $\partial f(x^*) = \{\nabla f(x^*)\}$ is a singleton, then the conditions (30) can be written in the familiar KKT form:

$$\nabla f(x^*) + \sum_{i \in \mathcal{I}(x^*)} \lambda_i a_i = 0. \quad (31)$$

We can also write (31) in the form

$$-\nabla f(x^*) \in N_X(x^*), \quad (32)$$

where $N_X(x^*)$, the normal cone to X at x^* , can be written as follows:

$$N_X(x^*) = \left\{ \sum_{i \in \mathcal{I}(x^*)} \lambda_i a_i : \lambda_i \geq 0, i \in \mathcal{I}(x^*) \right\}. \quad (33)$$

Given a candidate solution \hat{x} , we propose in our KKT test to estimate the distance

$$\delta(\hat{x}) := \text{dist}(-\nabla f(\hat{x}), N_X(\hat{x})),$$

where $\text{dist}(a, C)$ denotes the distance from point a to set C . The KKT conditions hold at \hat{x} if and only if $\delta(\hat{x}) = 0$.

Assuming that equation (29) holds at the point \hat{x} , the gradient $\nabla \hat{f}_N(\hat{x})$ of the sample-average function provides an unbiased and consistent estimator of $\nabla f(\hat{x})$. It follows that the quantity $\hat{\delta}_N(\hat{x})$ defined by

$$\hat{\delta}_N(\hat{x}) := \text{dist}\left(-\nabla \hat{f}_N(\hat{x}), N_X(\hat{x})\right) \quad (34)$$

gives a consistent estimator of $\delta(\hat{x})$. Moreover, the covariance matrix of $\nabla \hat{f}_N(\hat{x})$ can be estimated, allowing a confidence region for $\nabla f(\hat{x})$ to be constructed, and therefore allowing a statistical validation of the KKT conditions. This test is discussed in detail in [36].

Consider now the case in which the function $f(\cdot)$ in (1) is convex but is not differentiable everywhere, as typically happens in the two-stage linear problem with recourse when the distribution is discrete. We can write the first-order conditions (30) as follows:

$$\exists \gamma \in \partial f(x^*) \text{ such that } -\gamma \in N_X(x^*). \quad (35)$$

Defining

$$\Delta(\hat{x}) := \inf_{\gamma \in \partial f(\hat{x})} \text{dist}(-\gamma, N_X(\hat{x})), \quad (36)$$

we have that $\Delta(\hat{x}) = 0$ if and only if \hat{x} satisfies the conditions (30). Similarly to (34), we estimate $\Delta(\hat{x})$ by $\hat{\Delta}_N(\hat{x})$ defined as follows:

$$\hat{\Delta}_N(\hat{x}) := \inf_{\gamma \in \partial \hat{f}_N(\hat{x})} \text{dist}(-\gamma, N_X(\hat{x})). \quad (37)$$

It is known that the Hausdorff distance between $\partial f(x)$ and its sample-average estimate $\partial \hat{f}_N(x)$ converges with probability 1 to zero, and moreover the convergence is uniform in x on any compact subset of \mathbb{R}^m if the distribution of $\xi(\omega)$ has a finite support [35, Lemma 2.4]. In this sense, $\partial \hat{f}_N(x)$ provides a consistent estimator of $\partial f(x)$, and it follows that $\hat{\Delta}_N(\hat{x})$ is a consistent estimator of $\Delta(\hat{x})$.

It is, of course, impractical to calculate the entire subdifferential $\partial \hat{f}_N(\hat{x})$. However, when f is convex and piecewise linear (as in the case of two-stage linear programs with recourse), then in the neighborhood of any point \hat{x} there exist a finite number of points x^1, x^2, \dots, x^ℓ such that \mathcal{Q} is differentiable at these points and $\partial f(\hat{x}) = \text{conv}\{\nabla f(x^1), \nabla f(x^2), \dots, \nabla f(x^\ell)\}$. This observation suggests the following test.

- (i) Generate M points x^j , $j = 1, 2, \dots, M$, randomly distributed in a small neighborhood of \hat{x} .
- (ii) Calculate an estimate $\hat{\gamma}_j$ of a subgradient of f at each point x^j . This can be done either by using a common sample for all subgradients, that is, $\hat{\gamma}_j \in \partial \hat{f}_N(x^j)$, where \hat{f}_N is defined by the same sample for each argument x^j ; or by using a different sample for each point x^j .

- (iii) Calculate the minimal distance of the convex hull of $\hat{\gamma}_j$, $j = 1, 2, \dots, M$, to $-N_X(\hat{x})$, given by the optimal value $\hat{\delta}_N^M(\hat{x})$ of the following problem:

$$\begin{aligned} \min_{\alpha} \text{dist} \left(-\sum_{i=1}^M \alpha_i \hat{\gamma}_i, N_X(\hat{x}) \right), \quad \text{subject to} \quad (38) \\ \sum_{j=1}^M \alpha_j = 1, \quad \alpha_j \geq 0, \quad j = 1, 2, \dots, M. \end{aligned}$$

The convex hull of $\hat{\gamma}_j$, $j = 1, 2, \dots, M$, converges (in the Hausdorff metric) to $\partial f(\hat{x})$ with probability 1 as N and M both tend to infinity. It follows that this test is a consistent KKT test.

Note that the convex hull of $\hat{\gamma}_j$, $j = 1, 2, \dots, M$, actually converges to the convex hull of the set $\cup_{\|x-\hat{x}\| \leq \epsilon} \partial f(x)$, where $\epsilon > 0$ defines a small neighborhood of \hat{x} . Although, since $f(\cdot)$ is piecewise linear, for ϵ sufficiently small this set will coincide with $\partial f(\hat{x})$, we should in practice choose ϵ to be a value corresponding to the error we are prepared to tolerate in \hat{x} . That is, if we wish only to verify that \hat{x} is within ϵ of an exact solution, it suffices to use this value of ϵ in this test.

5.4. Testing KKT Conditions II

Here we consider a variant of the test proposed in the previous section. As above, the aim is to estimate the quantity $\Delta(\hat{x})$ defined by (36). Rather than building up an estimate of the subgradient by taking randomly distributed points in a small neighborhood of \hat{x} , we use different samples of size N to generate different elements of the true subdifferential $\partial f(\hat{x})$. The approach is as follows.

- (i) For given N and M , choose M samples of size N , that is, $\{\xi^{1,j}, \xi^{2,j}, \dots, \xi^{N,j}\}$, $j = 1, 2, \dots, M$.
- (ii) For each $j = 1, 2, \dots, M$, and defining \hat{f}_N^j as in (14), calculate an element $\hat{\gamma}_j \in \partial \hat{f}_N^j(\hat{x})$.
- (iii) Calculate the minimal distance of the convex hull of $\hat{\gamma}_j$, $j = 1, 2, \dots, M$, to $-N_X(\hat{x})$ by solving (38).

Note that this test is “stricter” than the test of the previous section, because it constructs an approximation to $\partial f(\hat{x})$ by sampling subgradients at the fixed point \hat{x} . Consequently, unless $\partial f(\hat{x})$ is a singleton, there is no guarantee that the convex hull of $\hat{\gamma}_j$, $j = 1, 2, \dots, M$, converges (in the Hausdorff metric) to $\partial f(\hat{x})$ with probability one as N and M tend to infinity. Hence, in general, this test is not a consistent KKT test.

6. An Algorithm for Two-Stage Stochastic Linear Programs with Recourse

In this section, we focus our attention on two-stage stochastic linear programs with recourse over a discrete scenario space, which are the subjects of our computational experiments. We discuss briefly the algorithm used in the experiments, referring the reader to Linderoth and Wright [21] for further details.

The problem we consider is (2), with second-stage problems defined by (3), with fixed recourse, and with a finite number K of scenarios. We thus have \mathcal{Q} defined by (4), where ξ_k defines the data vector (q_k, h_k, T_k, W) , $k = 1, 2, \dots, K$. The optimality conditions for each second-stage problem (3) with $\xi = \xi_k$ are that there exists a vector π_k such that

$$q_k - W^T \pi_k \geq 0 \perp y_k \geq 0, \quad W y_k = h_k - T_k \hat{x}.$$

(The symbol \perp represents a complementarity condition, which we can also state as requiring that the inner product of these two nonnegative vectors must be zero.) It can be shown that

$$-T_k^T \pi_k \in \partial Q(x; \xi_k).$$

From Rockafellar [30, Theorem 23.8], we have that

$$c + \sum_{k=1}^K p_k \partial Q(x; \xi_k) = c + \partial \mathcal{Q}(x),$$

and therefore

$$c - \sum_{k=1}^K p_k T_k^T \pi_k \in c + \partial \mathcal{Q}(x). \quad (39)$$

Since the evaluation of each $Q(x; \xi_k)$ requires solution of (3) and since π_k is simply the dual solution of this linear program, it follows that we can obtain an element of the subdifferential for the objective function in (2) for little more than the cost of evaluating the function alone.

Our tests require the solution of several SAA problems (6), which have the same form as described above except that we work with a subset of N (instead of K) scenarios, and the probabilities $p_k = 1/N$.

The approach described in [21] uses subgradient information to construct a lower bounding piecewise-linear approximation to the objective function in (2). Candidate iterations are generated by finding the minimizer of this model subject to the constraints $Ax = b$, $x \geq 0$ from (2) and a trust-region constraint of the form $\|x - x^I\|_\infty \leq \Delta$, where x^I is the ‘‘incumbent’’ (roughly speaking, the best point found so far), and Δ is a trust-region radius. This bundle-trust-region subproblem can be formulated as a linear program and therefore solved efficiently. Convergence to a solution (existence of a solution is specified in the assumption (A4)) is assured by means of rules for adjusting the value of Δ , for accepting candidate iterates as the new incumbent or rejecting them, and for managing the subgradient information.

7. Computational Grids

The algorithm described in Section 6 is implemented to run on a computational grid built using Condor [22]. The Condor system manages distributively owned collections of workstations known as *Condor pools*. A unique and powerful feature

of Condor is that each machine’s owner specifies the conditions under which jobs from other users are allowed to run on his machine. The default policy is to terminate a Condor job when a workstation’s owner begins using the machine, and migrate it to some other machine. Under this policy, Condor jobs use only compute cycles that would have otherwise been wasted. Because of the minimal intrusion of the Condor system, workstation owners often are quite willing to contribute their machines to Condor pools, and large pools can be built. Table 1 gives summary statistics about the resources that made up the computational grid used in many of the experiments of Section 9.

Number	Type of Machine	Location
414	Intel/Linux PC	Argonne
579	Intel/Linux PC	Wisconsin
31	Intel/Solaris PC	Wisconsin
133	Sun/Solaris Workstation	Wisconsin
512	Intel/Linux PC	New Mexico

Table 1. Computational Grid used for Experiments

The pool of workers typically grows and shrinks dynamically during the computation, as some workers become available to the job and others are reclaimed by their owners or by other Condor users. It is important for algorithms using this platform to be able to handle the dynamism of the environment, the possibility that the computations on some workers may never be completed, and the high latencies in transferring information between master and workers. To help implement our algorithm on the Condor platform, we used the runtime support library MW [12, 10], which supports parallel computations of the master-worker type on the computational grid.

A typical computation in this environment proceeds as follows: A master processor (typically the user’s workstation or a powerful server) maintains a queue of tasks which it wishes the worker processors to solve, and a list of currently available workers. MW manages these queues and takes responsibility for assigning tasks to available workers. When a worker completes its task, it returns the results to the master, which performs some computation in response, possibly generating more tasks for the task queue.

Parallelism in the algorithmic approach of Section 6 arises in two places. First, solution of the second-stage linear programs can be carried out independently for each scenario ξ_k . Typically, we group these scenarios into “chunks” and assign the solution of each chunk to a worker processor. The master processor manages the function and subgradient information returned by the workers. The second source of parallelism arises from considering more than one candidate iterate at a time. We maintain a “basket” of possible new incumbents; at any given time, tasks for evaluating the objective function for each iterate in the basket are being performed on the workers. When evaluation of the function for a point in the basket is completed, the master decides whether or not to make it the new incumbent, and solves a trust-region subproblem to generate a new candidate.

Name	Application	Scenarios	First-Stage Size	Second-Stage Size
20term	Vehicle Assignment	1.1×10^{12}	(3,64)	(124, 764)
gbd	Aircraft Allocation	6.5×10^9	(4, 17)	(5, 10)
LandS	Electricity Planning	10^6	(2, 4)	(7, 12)
ssn	Telecom Network Design	10^{10}	(1, 89)	(175, 706)
storm	Cargo Flight Scheduling	6×10^{81}	(185, 121)	(528, 1259)

Table 2. Test Problem Dimensions

Consideration of more than one candidate at a time makes the algorithm less “synchronous,” making it less sensitive to slow or unreliable workers.

For the statistical tests that do not require the solution of a master problem, the batches or scenarios can be processed completely independently of each other. No centralized control is needed, except to gather statistics and results of the runs. Such solution approaches lend themselves extremely well to computing platforms where the processors themselves are diverse and loosely coupled, such as the computational grid used in our experiments.

8. Test Problems

Here we briefly describe the test problems used in our experiments. All are two-stage stochastic linear programs with recourse obtained from the literature, though in some cases we refined the scenario distribution to increase the total number of scenarios. The input data sets we used, in SMPS format, can be obtained from the following site:

www.cs.wisc.edu/~swright/stochastic/sampling/

8.1. 20term

This problem is described by Mak, Morton, and Wood [24]. It is a model of a motor freight carrier’s operations, in which the first-stage variables are positions of a fleet of vehicles at the beginning of a day, and the second-stage variables move the fleet through a network to satisfy point-to-point demands for shipments (with penalties for unsatisfied demands) and to finish the day with the same fleet configuration as at the start.

8.2. gdb

This problem is derived from the aircraft allocation problem described by Dantzig [6, Chapter 28]. Aircraft of different types are to be allocated to routes in a way that maximizes profit under uncertain demand. Besides the cost of operating the aircraft, there are costs associated with bumping passengers when the demand for seats outstrips the capacity. In this model, there are four types of aircraft flying on five routes, and the first-stage variables are the number of aircraft of

each type allocated to each route. (Since three of the type-route pairs are infeasible, there are 17 variables in all.) The first-stage constraints are bounds on the available number of aircraft of each type. The second-stage variables indicate the number of carried passengers and the number of bumped passengers on each of the five routes, and the five second-stage constraints are demand balance equations for the five routes. The demands (the right-hand sides of the five demand balance equations) are random. In the original model, there are a total of only 750 scenarios. By refining the possible demand states (to allow between 13 and 17 states on each route), we obtain a model with 646,425 scenarios.

8.3. *LandS*

Our test problem LandS is a modification of a simple problem in electrical investment planning presented by Louveaux and Smeers [23]. The four first stage variables represent capacities of different new technologies, and the 12 second-stage variables represent production of each of three different “modes” of electricity from each of the four technologies. First-stage constraints represent minimum total capacity and budget restrictions. The second-stage constraints include capacity constraints for each of the four technologies, and three demand constraints, which have the form

$$\sum_{i=1}^4 y_{ij} \geq \xi_j, \quad j = 1, 2, 3,$$

where ξ_1 , ξ_2 , and ξ_3 are the (possibly uncertain) demands. All other data is deterministic.

In [23], ξ_2 and ξ_3 are fixed at 3 and 2, respectively, while ξ_1 takes on three possible values. The total number of scenarios is therefore 3. We modified the problem to allow 100 different values for each of these random parameters, namely,

$$\xi_j = .04(k - 1), \quad k = 1, 2, \dots, 100, \quad j = 1, 2, 3,$$

each with probability .01. Since we assume independence of these parameters, the total number of scenarios is 10^6 , each with equal probability 10^{-6} .

8.4. *SSN*

The SSN problem of Sen, Doverspike, and Cosares [33] arises in telecommunications network design. The owner of the network sells private-line services between pairs of nodes in the network. When a demand for service is received, a route between the two nodes with sufficient bandwidth must be identified for the time period in question. If no such route is available, the demand cannot be satisfied and the sale is lost. The optimization problem is to decide where to add capacity to the network to minimize the expected rate of unmet requests.

Since the demands are random (the scenarios are based on historical demand patterns), this is a stochastic optimization problem.

The formulation in [33] is as follows:

$$\min Q(x) \text{ subject to } e^T x \leq b, \quad x \geq 0, \quad (40)$$

where x is the vector of capacities to be added to the arcs of the network, e is the vector whose components are all 1, b (the budget) is the total amount of capacity to be added, and $Q(x)$ is the expected cost of adding capacity in the manner indicated by x . We have $Q(x) = \mathbb{E}_P Q(x, d(\omega))$, where

$$Q(x, d) := \min_{f,s} \{e^T s : Af \leq x + c, Ef + s = d(\omega), (f, s) \geq 0\}. \quad (41)$$

Here, $d = d(\omega)$ is the random vector of demands; e is the vector whose components are all 1; the columns of A are incidence vectors that describe the topology of the network; f is the vector whose components are the number of connections between each node pair that use a certain route; c is the vector of current capacities; s is the vector of unsatisfied demands for each request; and the rows of E have the effect of summing up the total flow between a given node pair, over all possible routes between the nodes in that pair.

In the data set for SSN, there are total of 89 arcs and 86 point-to-point pairs; that is, the dimension of x is 89 and of $d(\omega)$ is 86. Each component of $d(\omega)$ is an independent random variable with a known discrete distribution. Specifically, there are between three and seven possible values for each component of $d(\omega)$, giving a total of approximately 10^{70} possible complete demand scenarios.

8.5. Storm

This problem is based on a cargo flight scheduling application described by Mulvey and Ruszczyński [26, p.486 et seq.]. The aim is to plan cargo-carrying flights over a set of routes in a network, where the amounts of cargo are uncertain. Each first-stage variable represents the number of aircraft of a certain type assigned to a certain route, and the first-stage equations represent minimum frequencies for flights on each route, and balance equations for arrival and departure of aircraft of each type at the nodes in the network. Second-stage variables include the amounts of cargo to be shipped between nodes of the network, the amounts of undelivered cargo on each route, and the unused capacity on each leg of each route. Second-stage constraints include demand constraints and balance equations. Only the demands (the right-hand sides of the 118 demand constraints) are random. Mulvey and Ruszczyński [26] generate scenarios from a uniform distribution of $\pm 20\%$ around the basic demand forecast. In our data set, we used five scenarios for each demand, namely, the multiples .8, .9, 1.0, 1.1 and 1.2 of the basic demand. Each scenario was assigned a probability of 0.2. Since the demand scenarios were allowed to vary independently, our model has a total of 5^{118} scenarios, or approximately 6×10^{81} .

9. Computational Results

9.1. Optimal Value Estimates

Summaries of our experiments to determine upper and lower bounds on the optimal objective values are shown in Tables 3 and 4. Our experimental procedure is as described in Sections 4.1 and 4.2. For lower bound estimates, we calculated values of \hat{v}_N^j (see (14)) for $j = 1, 2, \dots, M$, with M between 7 and 10, and $N = 50, 100, 500, 1000, 5000$. The 95% confidence interval (18) for $\mathbb{E}[\hat{v}_N]$ is tabulated for each problem and for each sample size N in Tables 3 and 4. Monte Carlo (MC) was used to select the samples in Tables 3, while Latin Hypercube sampling (LHS) was used in Table 4.

These tables also show upper bound estimates. For the computed minimizer \hat{x}_N^j in each of the lower-bound trials, we estimated $f(\hat{x}_N^j)$ by sampling $T = 50$ batches of $\bar{N} = 20000$ scenarios, thereby obtaining a confidence interval of the form (23) for the value of the function f at each point \hat{x}_N^j . In Tables 3 and 4, we show on each line the 95% confidence interval obtained for the \hat{x}_N^j , $j = 1, \dots, M$ with the *lowest* value of $U_{\bar{N}, T}(\hat{x})$.

Further details of these experiments are given in Section A in the appendix. In particular, we show in Tables 13 through 22 the values of \hat{v}_N^j obtained in each lower-bound trial, and the corresponding confidence interval for $f(\hat{x}_N^j)$ for each computed solution \hat{x}_N^j . We also graph the results from Tables 3 and 4 in Figures 1, 2, 3, 4, and 5 in Section A. In these figures, the 95% confidence intervals for the lower and upper bounds are plotted as vertical bars, for each value of N , each problem, and each sampling technique. The graphs for MC and LHS are plotted on the same scale in each graph, to emphasize the improvements obtained with the latter technique.

Several comments about these results are in order. The use of LHS produced significant improvements over MC sampling on all problems; in fact, the improvements were dramatic for all problems except SSN. The bias $v^* - \mathbb{E}[\hat{v}_N]$, which can be estimated by comparing the corresponding entries in the last two columns of Tables 3 and 4, is significantly better with the LHS scheme of Table 4. For the two smaller problems gbd and LandS, the results from Table 4 strongly suggest that the optimal solution value is identified exactly; we discuss this issue further in Section 9.2. For 20term and storm, the relative gap appears small, with storm showing a remarkable improvement in the variance of the lower bound estimate \hat{v}_N with the use of LHS. For 20term, gbd, LandS, and storm, the upper bound estimate is remarkably unaffected by the value of N , suggesting that the approximate solutions \hat{x}_N^j obtained even for sample sizes $N = 100$ are of similar quality to this obtained for $N = 5000$. This fact is particularly surprising in the case of storm, a problem with a huge number of scenarios.

The conditioning of SSN is significantly worse than than of the other problems. A sample size of at least $N = 500$ was required to identify approximate solutions of reasonable quality, and the quality continues to improve as N is increased to 1000 and 5000. LHS produces improvements in both the optimality

gap and the variance of the upper- and lower-bound estimates, and our results indicate strongly that the true optimal value is about 9.90.

Problem	N	$E\hat{v}_N$ (95% conf. int.)	Best $f(\hat{x}_N^j)$ (95% conf. int.)
20term	50	253361.33 \pm 944.06	254317.96 \pm 19.89
20term	100	254024.89 \pm 800.88	254304.54 \pm 21.20
20term	500	254324.33 \pm 194.51	254320.39 \pm 27.36
20term	1000	254307.22 \pm 234.04	254333.83 \pm 18.85
20term	5000	254340.78 \pm 85.99	254341.36 \pm 20.32
gbd	50	1678.62 \pm 66.73	1655.86 \pm 1.34
gbd	100	1595.24 \pm 42.41	1656.35 \pm 1.19
gbd	500	1649.66 \pm 13.60	1654.90 \pm 1.46
gbd	1000	1653.50 \pm 12.32	1655.70 \pm 1.49
gbd	5000	1653.13 \pm 4.37	1656.40 \pm 1.31
LandS	50	227.19 \pm 4.03	225.71 \pm 0.12
LandS	100	226.39 \pm 3.99	225.55 \pm 0.12
LandS	500	226.02 \pm 1.43	225.61 \pm 0.12
LandS	1000	225.96 \pm 0.76	225.70 \pm 0.13
LandS	5000	225.72 \pm 0.52	225.70 \pm 0.12
ssn	50	4.11 \pm 1.23	12.68 \pm 0.05
ssn	100	7.66 \pm 1.31	11.20 \pm 0.05
ssn	500	8.54 \pm 0.34	10.28 \pm 0.04
ssn	1000	9.31 \pm 0.23	10.09 \pm 0.03
ssn	5000	9.98 \pm 0.21	9.86 \pm 0.05
storm	50	15506271.7 \pm 22043.4	15499092.17 \pm 845.26
storm	100	15482549.9 \pm 19213.8	15499056.00 \pm 623.30
storm	500	15498139.8 \pm 4152.8	15498468.02 \pm 684.65
storm	1000	15500879.4 \pm 4847.1	15498893.02 \pm 695.90
storm	5000	15498121.3 \pm 1879.0	15498646.89 \pm 696.08

Table 3. Lower and Upper Bound Estimates for v^* , Monte Carlo Sampling.

9.2. Convergence of Approximate Solution

Taking the approximate solutions \hat{x}_N^j , $j = 1, 2, \dots, M$ in the Latin Hypercube experiments with $N = 5000$, we calculated pairwise distances between six of these solutions, for the problems 20term, gbd, LandS, SSN, and storm. The resulting “distance matrices” are shown in Tables 5—9. Table 6 confirms that we appear to have identified an exact solution for gbd; that is, $\hat{x}_N^j \equiv x^*$, where $\{x^*\} = S$ is the unique solution of (2), (3), (4). For 20term and SSN, Tables 5 and 8 suggest an “ill conditioned” problem—one in which a small change to the data results in a large change to the solution set S . For these problems, it is quite possible that S is a set with a fairly large diameter. It is even more likely that S lies in a shallow basin; that is, there is a large, feasible neighborhood of S within which the value of f is not much different from the optimal value v^* . For the storm problem, on the other hand, Table 9 indicates that the approximate solutions \hat{x}_N^j , $i = j, 2, \dots, M$ are quite close together for $N = 5000$, suggesting that they may be converging to a unique minimizer x^* that will be identified at a larger sample size N . We report below on further investigations of this

Problem	N	$E\hat{v}_N$ (95% conf. int.)	Best $f(\hat{x}_N^j)$ (95% conf. int.)
20term	50	254307.57 \pm 371.80	254328.69 \pm 4.73
20term	100	254387.00 \pm 252.13	254312.50 \pm 5.77
20term	500	254296.43 \pm 117.95	254315.82 \pm 4.59
20term	1000	254294.00 \pm 95.22	254310.33 \pm 5.23
20term	5000	254298.57 \pm 38.74	254311.55 \pm 5.56
gbd	50	1644.21 \pm 10.71	1655.628 \pm 0.00
gbd	100	1655.62 \pm 0.00	1655.628 \pm 0.00
gbd	500	1655.62 \pm 0.00	1655.628 \pm 0.00
gbd	1000	1655.62 \pm 0.00	1655.628 \pm 0.00
gbd	5000	1655.62 \pm 0.00	1655.628 \pm 0.00
LandS	50	222.59 \pm 2.75	225.647 \pm 0.004
LandS	100	225.57 \pm 0.16	225.630 \pm 0.004
LandS	500	225.65 \pm 0.05	225.628 \pm 0.004
LandS	1000	225.64 \pm 0.03	225.633 \pm 0.005
LandS	5000	225.62 \pm 0.02	225.624 \pm 0.005
ssn	50	10.10 \pm 0.81	11.380 \pm 0.023
ssn	100	8.90 \pm 0.36	10.542 \pm 0.021
ssn	500	9.87 \pm 0.22	10.069 \pm 0.026
ssn	1000	9.83 \pm 0.29	9.996 \pm 0.025
ssn	5000	9.84 \pm 0.10	9.913 \pm 0.022
storm	50	15497683.7 \pm 1078.8	15498722.14 \pm 17.97
storm	100	15499255.3 \pm 1011.7	15498739.40 \pm 18.34
storm	500	15498661.4 \pm 280.8	15498733.67 \pm 17.71
storm	1000	15498598.1 \pm 148.5	15498735.73 \pm 19.93
storm	5000	15498657.8 \pm 73.9	15498739.41 \pm 19.11

Table 4. Lower and Upper Bound Estimates for v^* , Latin Hypercube Sampling.

issue, but note here that the result is somewhat surprising; given the very large number of scenarios for this problem we might have expected it to display the same symptoms of ill conditioning as 20term or SSN.

0.000	16.1	26.5	296	278	263
16.1	0.000	20.3	297	279	264
26.5	20.3	0.000	291	273	259
296	297	291	0.000	49.2	67.0
278	279	273	49.2	0.000	31.9
263	264	259	67.0	31.9	0.000

Table 5. Distance Matrix for Six SAAs of 20term with $N = 5000$, Latin Hypercube Sampling.

0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00

Table 6. Distance Matrix for Six SAAs of gbd with $N = 5000$, Latin Hypercube Sampling.

0.000	0.0566	0.0980	0.0566	0.0566	0.0800
0.0566	0.000	0.0980	0.000	0.000	0.0566
0.0980	0.0980	0.000	0.0980	0.0980	0.1497
0.0566	0.000	0.0980	0.000	0.000	0.0566
0.0566	0.000	0.0980	0.000	0.000	0.0566
0.0800	0.0566	0.1497	0.0566	0.0566	0.000

Table 7. Distance Matrix for Six SAAs of LandS with $N = 5000$, Latin Hypercube Sampling.

0.00	19.9	13.3	22.0	16.9	21.8
19.9	0.00	21.6	27.3	23.0	18.1
13.3	21.6	0.00	22.4	19.5	22.3
22.0	27.3	22.4	0.00	26.4	30.6
16.9	23.0	19.5	26.4	0.00	26.7
21.8	18.1	22.3	30.6	26.7	0.00

Table 8. Distance Matrix for Six SAAs of SSN with $N = 5000$, Latin Hypercube Sampling.

0.000	0.0187	0.0037	0.0105	0.0147	0.0069
0.0187	0.000	0.0153	0.0095	0.0078	0.0124
0.0037	0.0153	0.000	0.0069	0.0122	0.0033
0.0105	0.0095	0.0069	0.000	0.0097	0.0036
0.0147	0.0078	0.0122	0.0097	0.000	0.0105
0.0069	0.0124	0.0033	0.0036	0.0105	0.000

Table 9. Distance Matrix for Six SAAs of storm with $N = 5000$, Latin Hypercube Sampling.

We performed the repeated sampling test of Section 5.1 on the gbd, LandS, and storm problems. For gbd and LandS, we solved an SAA with $N = 5000$ (using LHS), then solved 10 more different instances with different samples of the same size, with a very tight convergence tolerance of 10^{-12} . The results are as expected from the observations above; all 10 additional trials terminated without stepping away from their starting point, effectively confirming that all instances have the same solution \hat{x}_N .

For storm, we tried repeated sampling with sample sizes $N = 20000$ and $N = 50000$ (with LHS), but did not obtain convergence of the solution of the SAA to an identical value. Close examination of the solutions, however, suggested a solution set S with a very specific structure. In the 10 trials for $N = 20000$, the 10 solutions were essentially identical in all but one of their 121 components. The exception was the 112th component, which varied between 1.46003333 and 1.47593333. In the 10 trials for $N = 50000$, the 10 solutions were again identical in all components except the 112th, which took values between 1.46123333 and 1.47210000. It seems likely that the condition number κ as defined by (11) is very large (possibly infinite), yet this problem is quite well behaved, in that for small ε , the ε -optimal solution set S_ε is close to being a short line segment, rather than a large multidimensional basin.

For SSN, we tested with two SAAs at each sample size $N = 10000$, 20000, and 40000, all using LHS, but found that the two solutions obtained for each value of N differed. We conclude that either the solution set S for SSN is not a singleton, or else that the values of N we tried were not large enough to identify the unique solution x^* .

Quantity	Estimate (95% conf. int.)
$f(x^0)$	$9.9016607 \pm .0035396$
$f(x^1)$	$9.9016731 \pm .0035396$
$f(x^2)$	$9.9043870 \pm .0035302$
$f(x^1) - f(x^0)$	$1.24118 \times 10^{-5} \pm 2.910 \times 10^{-7}$
$f(x^2) - f(x^0)$	$2.72648 \times 10^{-3} \pm 4.96724 \times 10^{-5}$

Table 10. Confidence Intervals for Function Values and Gaps for SSN at Three Points

9.3. Comparison of Several Points

We performed a comparison of several candidate solution points for the SSN problem, using the technique described in Section 5.2. We compared three points. The first point x^0 was obtained by solving an SAA with sample size $N = 250000$, generated using LHS. To obtain the second point x^1 , we generated an SAA with a *different* sample of size $N = 250000$, and applied our trust-region algorithm starting from the first point x^0 , and initial trust-region radius 10^{-3} . The point x^1 was the first point accepted by this algorithm as having a lower value of \hat{f}_N for this sample. For the point x^2 , we used the same SAA as was used to obtain x^1 , but this time applied our algorithm with an initial trust-region radius of 1. The point x^2 was the first point accepted by the algorithm. (At the time it was generated, the trust-region radius had been reduced to .4 by the algorithm.)

We applied the test of Section 5.2 $\bar{N} = 25000$ (samples generated by LHS) and $T = 2000$ to obtain confidence intervals for the function values at x^0 , x^1 , and x^2 , and the gaps $f(x^1) - f(x^0)$ and $f(x^2) - f(x^0)$. Results are shown in Table 10. They show that x^0 almost certainly has the lowest function value of the three. In this sense, the second SAA (which was used to obtain x^1 and x^2) appears to give a less accurate approximation to the true function f in the vicinity of these three points than does the first SAA (for which x^0 was the minimizer). It also confirms the impression gained from Table 4 that the optimal value x^* for SSN is very close to 9.90. We can also conclude that the solution set for SSN is flat in the neighborhood of the optimal solution.

9.4. KKT Tests for Near-Optimality

We now report on the KKT tests described in Sections 5.3 and 5.4, applied to the problems 20term and SSN.

For 20term, we set \hat{x} to be a computed solution of an SAA with $N = 100000$. The two first-stage equality constraints were active at this point, of course, but the single inequality constraint was inactive. We calculated approximations to the distances $\hat{\delta}_N^M(\hat{x})$ defined in (38) using an ℓ_1 norm, since this allowed us to formulate the distance calculation as a linear program, which we solved with a simple AMPL model. For the test of Section 5.3, we generated each of the M subgradients $\hat{\gamma}_i$, $i = 1, 2, \dots, M$, at points x in a neighborhood of radius 10^{-2} around \hat{x} . We used MC sampling with $N = 100000$ to generate the subgradients, using a different sample for each point.

M	$\delta_N^M(\hat{x})$	λ
50	1065	
500	53.08	$(-100, -7 \times 10^{-12}, 0)$
1000	10.89	$(-100, -1 \times 10^{-11}, 0)$
2000	3.53	$(-100, -8 \times 10^{-14}, 0)$

Table 11. Discrepancy in KKT Conditions for 20term Using Test of Section 5.3

M	$\delta_N^M(\hat{x})$	λ
5	.254847	.050000
20	.250329	.050000
250	.238522	.049774
1000	.237175	.049733
2000	.236541	.049878

Table 12. Discrepancy in KKT Conditions for 20term Using Test of Section 5.3

Results are shown for different values of M in Table 11. The final column contains the values of λ , the multiplier vector for the active constraints at the nearest point in $N_X(\hat{x})$ to the approximate subdifferential of f . The third component of λ is uniformly zero since the third constraint is inactive, while the first two components are clearly converging to -100 and 0 , respectively. The distance to the subdifferential approximation appears to converge toward zero as more gradients are gathered, as we would expect, since the convex hull covers the subdifferential more fully as the number of gradients increases.

Although useful in verifying near-optimality, it is clear that this test is not a practical termination test for an algorithm, because the time required to compute enough gradients to verify optimality is too great.

For SSN, we performed the test described in Section 5.4 to the point x^0 discussed in Section 9.3. A subgradient was generated at each of the $T = 2000$ SAAs used in the experiment of Section 9.3, where each SAA has $\bar{N} = 25000$ and is obtained by LHS. The single inequality constraint in (40) is active at this solution, as are a number of the lower bounds of zero on the components of x . We performed this test using different values M for the size of the subgradient bundle. In each case we chose simply the first M of the 2000 subgradients in the list.

Results are shown in Table 12. The second column indicates the distance between the convex hull of the subgradient bundle and the normal cone of the constraints, and the third column indicates the Lagrange multiplier for the inequality constraint $e^T x \leq b$ in (40). We see that the distance does not shrink to zero as M increases, indicating either that the point x^0 is some distance from optimality, or that the convex hull of the subgradient bundle is not converging to $\partial f(x^0)$. We conjecture that the latter is the case, since we have good evidence from the other tests that x^0 is close to optimal, and we know in general that Hausdorff convergence of the subgradient bundle to $\partial f(x^0)$ need not occur. Our experience with this test indicates that it is not valuable.

10. Conclusions

We have described experiments with sample-average approximations to five two-stage stochastic linear programs with recourse. We first obtained confidence intervals for upper and lower bounds on the optimal value of these five problems, by solving many SAA instances for different sample sizes, using two different sampling techniques. We then investigated the optimality of various candidate solutions using a number of tests motivated by recently developed theory, including tests that are applicable to problems with a singleton solution set and tests based on satisfaction of the Karush-Kuhn-Tucker conditions for optimality. We found that the five test problem varied widely in difficulty, in that good approximations to the true solution could be obtained with small sample sizes in some cases, while large sample sizes were needed in others.

All these experiments were made possible by the availability of the ubiquitous computing power of the computational grid, and by software tools such as Condor and MW that enable us to run practical computations on the grid. All these ingredients, along with an algorithm that runs effectively on the grid, were necessary in building the software tool ATR (described comprehensively in [21]), which was needed to solve the multitude of instances required by this study.

References

1. K. Anstreicher, N. Brixius, J.-P. Goux, and J. T. Linderoth. Solving large quadratic assignment problems on computational grids. To appear in *Mathematical Programming*, available at http://www.optimization-online.org/DB_HTML/2000/10/233.html, 2001.
2. A. N. Avramidis and J. R. Wilson. Integrated variance reduction strategies for simulation. *Operations Research*, 44:327–346, 1996.
3. T. G. Bailey, P. Jensen, and D. P. Morton. Response surface analysis of two-stage stochastic linear programming with recourse. *Naval Research Logistics*, 46:753–778, 1999.
4. Q. Chen, M. C. Ferris, and J. T. Linderoth. Fatcop 2.0: Advanced features in an opportunistic mixed integer programming solver. *Annals of Operations Research*, 2000. To appear.
5. G. Dantzig and G. Infanger. Large-scale stochastic linear programs—Importance sampling and Bender’s decomposition. In C. Brezinski and U. Kulisch, editors, *Computational and Applied Mathematics I (Dublin, 1991)*, pages 111–120. North-Holland, Amsterdam, 1991.
6. G. B. Dantzig. *Linear Programming and Extensions*. Princeton University Press, Princeton, New Jersey, 1963.
7. U. M. Diwekar and J. R. Kalagnanam. An efficient sampling technique for optimization under uncertainty. *American Institute of Chemical Engineers Journal*, 43, 1997.
8. Y. Ermoliev. Stochastic quasigradient methods. In Y. Ermoliev and R. J.-B. Wets, editors, *Numerical techniques for stochastic optimization problems*. Springer-Verlag, Berlin, 1988.
9. I. Foster and C. Kesselman. *The Grid: Blueprint for a New Computing Infrastructure*. Morgan-Kaufmann, 1999. Chapter 1 “Grids in Context” by Larry Smarr and chapter 13 “High-Throughput Resource Management” by Miron Livny and Rajesh Raman.
10. J.-P. Goux, S. Kulkarni, J. T. Linderoth, and M. Yoder. Master-Worker : An enabling framework for master-worker applications on the computational grid. *Cluster Computing*, 4:63–70, 2001.
11. J.-P. Goux and S. Leyffer. Solving large MINLPs on computational grids. Numerical Analysis Report NA/200, Mathematics Department, University of Dundee, 2001.
12. J.-P. Goux, J. T. Linderoth, and M. Yoder. Metacomputing and the master-worker paradigm. Preprint MCS/ANL-P792-0200, Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, Ill., February 2000.

13. G. Gürkan, A. Y. Özge, and S. M. Robinson. Sample-path optimization in simulation. In *Proceedings of the Winter Simulation Conference*, pages 247–254, 1994.
14. J. L. Higle. Variance reduction and objective function evaluation in stochastic linear programs. *INFORMS Journal on Computing*, 10(2):236–247, 1998.
15. J. L. Higle and S. Sen. Stochastic decomposition: An algorithm for two-stage linear programs with recourse. *Mathematics of Operations Research*, 16:650–669, 1991.
16. J. L. Higle and S. Sen. Duality and statistical tests of optimality for two-stage stochastic programs. *Mathematical Programming*, 75:257–275, 1996.
17. J. L. Higle and S. Sen. *Stochastic Decomposition: A Statistical Method for Large Scale Stochastic Linear Programming*. Kluwer Academic Publishers, Boston, MA, 1996.
18. G. Infanger. *Planning Under Uncertainty: Solving Large Scale Stochastic Linear Programs*. Boyd and Fraser Publishing Company, 1994.
19. A. D. Ioffe and V. M. Tihomirov. *theory of Extremal Problems*. North-Holland, Amsterdam, 1979.
20. A. Kleywegt, A. Shapiro, and T. Homem-de Mello. The sample average approximation method for stochastic discrete optimization. *SIAM Journal on Optimization*, 2001. To appear.
21. J. T. Linderoth and S. J. Wright. Decomposition algorithms for stochastic programming on a computational grid. Preprint ANL/MCS-P875-0401, Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, Ill., April 2001. Available at http://www.optimization-online.org/DB_HTML/2001/04/315.html.
22. M. Livny, J. Basney, R. Raman, and T. Tannenbaum. Mechanisms for high throughput computing. *SPEEDUP*, 11, 1997. Available from http://www.cs.wisc.edu/condor/doc/htc_mech.ps.
23. F. V. Louveaux and Y. Smeers. Optimal investments for electricity generation: A stochastic model and a test problem. In Y. Ermoliev and R. J.-B. Wets, editors, *Numerical techniques for stochastic optimization problems*, pages 445–452. Springer-Verlag, Berlin, 1988.
24. 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.
25. M. D. McKay, R. J. Beckman, and W. J. Conover. A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics*, 21:239–245, 1979.
26. J. M. Mulvey and A. Ruszczyński. A new scenario decomposition method for large scale stochastic optimization. *Operations Research*, 43:477–490, 1995.
27. V. I. Norkin, G. Ch. Pflug, and A. Ruszczyński. A branch and bound method for stochastic global optimization. *Mathematical Programming*, 83:425–450, 1998.
28. A. B. Owen. Monte Carlo variance of scrambled equidistribution quadrature. *SIAM Journal on Numerical Analysis*, 34:1884–1910, 1997.
29. E. L. Plambeck, B. R. Fu, S. M. Robinson, and R. Suri. Sample-path optimization of convex stochastic performance functions. *Mathematical Programming, Series B*, 1996.
30. R. T. Rockafellar. *Convex Analysis*. Princeton University Press, Princeton, N.J., 1970.
31. R. Y. Rubinstein and A. Shapiro. Optimization of static simulation models by the score function method. *Mathematics and Computers in Simulation*, 32:373–392, 1990.
32. R. Y. Rubinstein and A. Shapiro. *Discrete Event Systems: Sensitivity Analysis and Stochastic Optimization by the Score Function Method*. Wiley, New York, 1993.
33. S. Sen, R. D. Doverspike, and S. Cosares. Network planning with random demand. *Telecommunications Systems*, 3:11–30, 1994.
34. A. Shapiro. Asymptotic analysis of stochastic programs. *Annals of Operations Research*, 30:169–186, 1991.
35. A. Shapiro and T. Homem de Mello. On the rate of convergence of optimal solutions of Monte Carlo approximations of stochastic programs. *SIAM Journal on Optimization*, 11:70–86, 2000.
36. A. Shapiro and T. Homem-de-Mello. A simulation-based approach to stochastic programming with recourse. *Mathematical Programming*, 81:301–325, 1998.
37. A. Shapiro, T. Homem-de Mello, and J. Kim. Conditioning of convex piecewise linear stochastic programs. Technical report, School of ISyE, Georgia Institute of Technology, Atlanta, GA 3032-0205, 2000.
38. M. Stein. Large sample properties of simulations using Latin Hypercube Sampling. *Technometrics*, 29:143–151, 1987.
39. Golbon Zakeri. Verifying a stochastic solution using meta-computing, 2000. Presentation at INFORMS Winter Meeting, Salt Lake City, USA.

A. Details of Bound Estimates

Here we tabulate the details of our experiments for determining upper and lower bounds, and graph the results from Tables 3 and 4. A description of the experimental methodology, and the values of the parameters N , \bar{N} , T , and so on, are given in Section 9.1.

N	i	$\hat{v}_N^{(i)}$	$\mathbb{E}(\hat{v}_N)$ (95% conf. int.)	$f(\hat{x}_N^{(i)})$ (95% conf. int.)	Best
50	1	252084	253361 ± 944.06	254397.80 ± 21.56	*
	2	253009		254333.82 ± 19.11	
	3	253060		254392.79 ± 19.71	
	4	251786		254351.03 ± 18.78	
	5	254045		254380.54 ± 24.75	
	6	253124		254425.12 ± 20.02	
	7	256794		254349.73 ± 20.70	
	8	253174		254383.53 ± 24.44	
	9	253176		254322.24 ± 22.90	
100	1	255012	254025 ± 800.88	254358.48 ± 21.85	*
	2	255350		254315.56 ± 20.19	
	3	253600		254300.62 ± 23.71	
	4	255024		254346.31 ± 20.59	
	5	253152		254321.04 ± 21.72	
	6	252012		254371.29 ± 31.39	
	7	253971		254321.12 ± 19.48	
	8	252768		254328.15 ± 21.92	
	9	255335		254314.30 ± 20.93	
500	1	254192	254324 ± 194.51	254377.86 ± 21.57	*
	2	254084		254316.99 ± 20.14	
	3	254601		254322.96 ± 21.55	
	4	253976		254334.15 ± 20.51	
	5	254368		254322.19 ± 21.06	
	6	254935		254353.95 ± 17.06	
	7	254294		254294.35 ± 25.35	
	8	254383		254303.03 ± 21.29	
	9	254086		254316.25 ± 17.01	
1000	1	254306	254307 ± 234.04	255168.89 ± 19.59	*
	2	254174		254343.38 ± 17.37	
	3	254626		254271.17 ± 27.10	
	4	254619		254338.87 ± 19.22	
	5	254541		254364.64 ± 21.14	
	6	253693		254343.03 ± 22.74	
	7	253852		254297.04 ± 19.35	
	8	254715		254303.15 ± 22.21	
	9	254239		254323.23 ± 23.83	
5000	1	254335	254341 ± 85.99	254314.01 ± 38.75	*
	2	254361		254332.26 ± 17.84	
	3	254484		254297.79 ± 27.06	
	4	254189		254316.99 ± 33.44	
	5	254483		254291.32 ± 28.22	
	6	254345		254319.74 ± 25.88	
	7	254471		254333.57 ± 28.14	
	8	254293		254309.17 ± 36.31	
	9	254106		254274.97 ± 33.27	

Table 13. 20term, Monte Carlo sampling.

N	i	$\hat{v}_N^{(i)}$	$\mathbb{E}(\hat{v}_N)$ (95% conf. int.)	$f(\hat{x}_N^{(i)})$ (95% conf. int.)	Best
50	1	254239	254307.57 \pm 371.80	254336 \pm 5.56	*
	2	254091		254325 \pm 4.58	
	3	254697		254383 \pm 4.52	
	4	253949		254411 \pm 6.65	
	5	253705		254457 \pm 8.42	
	6	254262		254338 \pm 5.04	
	7	255210		254328 \pm 5.26	
100	1	254433	254387.00 \pm 252.13	254313 \pm 5.43	*
	2	254255		254346 \pm 4.38	
	3	254524		254324 \pm 5.48	
	4	253886		254371 \pm 7.40	
	5	254869		254390 \pm 4.57	
	6	254077		254337 \pm 5.03	
	7	254665		254327 \pm 4.12	
500	1	254148	254296.43 \pm 117.95	254312 \pm 5.58	*
	2	254156		254322 \pm 4.52	
	3	254294		254317 \pm 4.63	
	4	254537		254329 \pm 4.32	
	5	254170		254312 \pm 4.12	
	6	254487		254319 \pm 4.65	
	7	254283		254315 \pm 5.13	
1000	1	254395	254294.00 \pm 95.22	254312 \pm 5.14	*
	2	254205		254313 \pm 6.18	
	3	254180		254312 \pm 5.65	
	4	254519		254317 \pm 3.72	
	5	254316		254326 \pm 4.83	
	6	254165		254328 \pm 5.23	
	7	254278		254317 \pm 5.44	
5000	1	254245	254298.57 \pm 38.74	254315 \pm 4.89	*
	2	254290		254316 \pm 5.41	
	3	254326		254316 \pm 5.21	
	4	254390		254312 \pm 4.92	
	5	254306		254312 \pm 5.51	
	6	254233		254310 \pm 5.80	
	7	254300		254313 \pm 5.08	

Table 14. 20term, Latin Hypercube sampling.

N	i	$\hat{v}_N^{(i)}$	$\mathbb{E}(\hat{v}_N)$ (95% conf. int.)	$f(\hat{x}_N^{(i)})$ (95% conf. int.)	Best
50	1	1788.68	1678.62 ± 66.73	1657.15 ± 1.41	*
	2	1607.03		1665.08 ± 1.29	
	3	1647.35		1657.24 ± 1.30	
	4	1998.96		1664.08 ± 1.47	
	5	1592.97		1658.74 ± 1.59	
	6	1559.19		1667.80 ± 1.37	
	7	1693.98		1659.83 ± 1.33	
	8	1664.59		1653.89 ± 1.10	
	9	1654.02		1657.08 ± 1.28	
	10	1640.95		1662.60 ± 1.76	
	11	1703.55		1659.35 ± 1.31	
	12	1592.13		1655.95 ± 1.26	
100	1	1586.44	1595.24 ± 42.41	1659.33 ± 1.45	*
	2	1560.36		1656.47 ± 1.42	
	3	1500.98		1676.39 ± 1.56	
	4	1480.81		1655.86 ± 1.30	
	5	1685.07		1655.29 ± 1.50	
	6	1642.51		1656.95 ± 1.25	
	7	1612.27		1656.61 ± 1.16	
	8	1735.93		1656.60 ± 1.21	
	9	1645.40		1659.32 ± 1.44	
	10	1526.68		1656.53 ± 1.13	
	11	1565.53		1655.09 ± 1.31	
	12	1600.89		1662.06 ± 1.35	
500	1	1676.01	1649.66 ± 13.60	1655.72 ± 1.30	*
	2	1620.61		1655.45 ± 1.78	
	3	1643.91		1653.97 ± 1.15	
	4	1691.30		1657.15 ± 1.29	
	5	1651.25		1655.58 ± 1.36	
	6	1633.11		1655.90 ± 1.44	
	7	1666.47		1656.06 ± 1.29	
	8	1637.12		1656.17 ± 1.15	
	9	1640.52		1654.62 ± 1.39	
	10	1633.17		1655.65 ± 1.45	
	11	1619.83		1657.84 ± 1.12	
	12	1682.61		1654.91 ± 1.25	
1000	1	1659.00	1653.50 ± 12.32	1656.62 ± 1.44	*
	2	1642.25		1655.83 ± 1.95	
	3	1631.26		1656.26 ± 1.37	
	4	1658.69		1655.74 ± 1.16	
	5	1623.98		1656.38 ± 1.44	
	6	1654.58		1655.17 ± 1.21	
	7	1642.15		1655.36 ± 1.18	
	8	1709.83		1656.09 ± 1.45	
	9	1666.58		1654.83 ± 1.36	
	10	1656.75		1655.30 ± 1.34	
	11	1656.62		1655.28 ± 1.16	
	12	1640.28		1656.07 ± 1.44	
5000	1	1664.86	1653.13 ± 4.37	1656.85 ± 1.88	*
	2	1652.38		1656.22 ± 1.27	
	3	1640.53		1655.39 ± 1.35	
	4	1659.53		1654.62 ± 1.32	
	5	1651.32		1655.19 ± 1.46	
	6	1643.72		1655.35 ± 1.70	
	7	1655.29		1655.47 ± 1.58	
	8	1661.05		1656.36 ± 1.60	
	9	1652.66		1655.75 ± 1.65	
	10	1653.17		1655.16 ± 1.98	
	11	1660.30		1653.98 ± 1.45	
	12	1642.78		1655.09 ± 1.63	

Table 15. gbd, Monte Carlo sampling.

N	i	$\hat{v}_N^{(i)}$	$\mathbb{E}(\hat{v}_N)$ (95% conf. int.)	$f(\hat{x}_N^{(i)})$ (95% conf. int.)	Best
50	1	1617.97	1644.21 \pm 10.71	1656.77 \pm 0.00	*
	2	1661.83		1655.64 \pm 0.00	
	3	1650.97		1655.64 \pm 0.00	
	4	1626.21		1655.64 \pm 0.00	
	5	1659.45		1655.64 \pm 0.00	
	6	1642.91		1655.64 \pm 0.00	
	7	1660.43		1655.62 \pm 0.00	
	8	1618.77		1655.64 \pm 0.00	
	9	1657.15		1655.64 \pm 0.00	
	10	1646.41		1656.77 \pm 0.00	
100	1	1655.62	1655.62 \pm 0.00	1655.62 \pm 0.00	*
	2	1655.62		1655.62 \pm 0.00	
	3	1655.62		1655.62 \pm 0.00	
	4	1655.62		1655.62 \pm 0.00	
	5	1655.62		1655.62 \pm 0.00	
	6	1655.62		1655.62 \pm 0.00	
	7	1655.62		1655.62 \pm 0.00	
	8	1655.62		1655.62 \pm 0.00	
	9	1655.62		1655.62 \pm 0.00	
	10	1655.62		1655.62 \pm 0.00	
500	1	1655.62	1655.62 \pm 0.00	1655.62 \pm 0.00	*
	2	1655.62		1655.62 \pm 0.00	
	3	1655.62		1655.62 \pm 0.00	
	4	1655.62		1655.62 \pm 0.00	
	5	1655.62		1655.62 \pm 0.00	
	6	1655.62		1655.62 \pm 0.00	
	7	1655.62		1655.62 \pm 0.00	
	8	1655.62		1655.62 \pm 0.00	
	9	1655.62		1655.62 \pm 0.00	
	10	1655.62		1655.62 \pm 0.00	
1000	1	1655.62	1655.62 \pm 0.00	1655.62 \pm 0.00	*
	2	1655.62		1655.62 \pm 0.00	
	3	1655.62		1655.62 \pm 0.00	
	4	1655.62		1655.62 \pm 0.00	
	5	1655.62		1655.62 \pm 0.00	
	6	1655.62		1655.62 \pm 0.00	
	7	1655.62		1655.62 \pm 0.00	
	8	1655.62		1655.62 \pm 0.00	
	9	1655.62		1655.62 \pm 0.00	
	10	1655.62		1655.62 \pm 0.00	
5000	1	1655.62	1655.62 \pm 0.00	1655.62 \pm 0.00	*
	2	1655.62		1655.62 \pm 0.00	
	3	1655.62		1655.62 \pm 0.00	
	4	1655.62		1655.62 \pm 0.00	
	5	1655.62		1655.62 \pm 0.00	
	6	1655.62		1655.62 \pm 0.00	
	7	1655.62		1655.62 \pm 0.00	
	8	1655.62		1655.62 \pm 0.00	
	9	1655.62		1655.62 \pm 0.00	
	10	1655.62		1655.62 \pm 0.00	

Table 16. gbd, Latin Hypercube sampling.

N	i	$\hat{v}_N^{(i)}$	$\mathbb{E}(\hat{v}_N)$ (95% conf. int.)	$f(\hat{x}_N^{(i)})$ (95% conf. int.)	Best
50	1	232.93	227.19 ± 4.03	225.87 ± 0.12	*
	2	236.05		225.78 ± 0.10	
	3	238.39		225.86 ± 0.12	
	4	223.36		225.71 ± 0.14	
	5	221.39		225.77 ± 0.13	
	6	224.89		225.79 ± 0.12	
	7	222.25		225.81 ± 0.14	
	8	227.89		225.60 ± 0.10	
	9	228.55		225.63 ± 0.11	
	10	215.11		225.74 ± 0.11	
	11	228.30		225.67 ± 0.12	
100	1	230.03	226.39 ± 3.99	225.67 ± 0.12	*
	2	232.06		225.85 ± 0.12	
	3	230.17		225.68 ± 0.12	
	4	217.40		225.75 ± 0.13	
	5	218.49		225.74 ± 0.12	
	6	230.43		225.65 ± 0.12	
	7	229.71		225.71 ± 0.11	
	8	217.07		225.78 ± 0.12	
	9	221.84		225.70 ± 0.12	
	10	225.83		225.71 ± 0.14	
	11	237.29		225.93 ± 0.10	
500	1	227.30	226.02 ± 1.43	225.61 ± 0.13	*
	2	227.88		225.60 ± 0.11	
	3	227.60		225.64 ± 0.13	
	4	224.53		225.56 ± 0.12	
	5	223.53		225.60 ± 0.12	
	6	228.60		225.58 ± 0.10	
	7	222.73		225.69 ± 0.13	
	8	224.72		225.49 ± 0.11	
	9	227.69		225.61 ± 0.12	
	10	222.64		225.56 ± 0.13	
	11	229.01		225.61 ± 0.10	
1000	1	225.99	225.96 ± 0.76	225.62 ± 0.14	*
	2	226.33		225.70 ± 0.12	
	3	225.66		225.56 ± 0.10	
	4	226.38		225.64 ± 0.11	
	5	225.66		225.63 ± 0.10	
	6	226.22		225.66 ± 0.12	
	7	225.43		225.53 ± 0.11	
	8	223.47		225.61 ± 0.11	
	9	228.58		225.67 ± 0.12	
	10	227.07		225.69 ± 0.11	
	11	224.78		225.66 ± 0.10	
5000	1	224.17	225.72 ± 0.52	225.69 ± 0.12	*
	2	224.95		225.60 ± 0.14	
	3	225.83		225.50 ± 0.11	
	4	225.66		225.70 ± 0.14	
	5	226.49		225.63 ± 0.10	
	6	225.25		225.53 ± 0.12	
	7	226.79		225.70 ± 0.12	
	8	226.45		225.67 ± 0.14	
	9	225.64		225.64 ± 0.13	
	10	224.86		225.67 ± 0.14	
	11	226.88		225.66 ± 0.13	

Table 17. LandS, Monte Carlo sampling.

N	i	$\hat{v}_N^{(i)}$	$\mathbb{E}(\hat{v}_N)$ (95% conf. int.)	$f(\hat{x}_N^{(i)})$ (95% conf. int.)	Best
50	1	215.517	222.59 ± 2.75	225.726 ± 0.01	*
	2	222.507		225.682 ± 0.00	
	3	220.662		225.646 ± 0.00	
	4	223.54		225.651 ± 0.00	
	5	225.779		225.646 ± 0.00	
	6	226.508		225.769 ± 0.01	
	7	218.88		225.694 ± 0.00	
	8	230.732		225.65 ± 0.00	
	9	223.285		225.69 ± 0.00	
	10	218.493		225.669 ± 0.00	
100	1	225.269	225.57 ± 0.16	225.654 ± 0.01	*
	2	225.443		225.645 ± 0.00	
	3	225.755		225.629 ± 0.00	
	4	225.678		225.632 ± 0.00	
	5	225.596		225.635 ± 0.00	
	6	225.665		225.635 ± 0.00	
	7	225.493		225.642 ± 0.00	
	8	225.164		225.639 ± 0.00	
	9	225.61		225.632 ± 0.00	
	10	226.057		225.657 ± 0.00	
500	1	225.492	225.65 ± 0.05	225.631 ± 0.00	*
	2	225.674		225.634 ± 0.00	
	3	225.628		225.632 ± 0.00	
	4	225.543		225.63 ± 0.00	
	5	225.656		225.631 ± 0.00	
	6	225.699		225.632 ± 0.00	
	7	225.702		225.63 ± 0.01	
	8	225.671		225.627 ± 0.00	
	9	225.691		225.633 ± 0.00	
	10	225.735		225.63 ± 0.00	
1000	1	225.653	225.64 ± 0.03	225.634 ± 0.01	*
	2	225.649		225.633 ± 0.01	
	3	225.654		225.629 ± 0.00	
	4	225.655		225.633 ± 0.00	
	5	225.674		225.633 ± 0.00	
	6	225.536		225.632 ± 0.01	
	7	225.587		225.633 ± 0.00	
	8	225.663		225.628 ± 0.01	
	9	225.614		225.627 ± 0.00	
	10	225.696		225.633 ± 0.00	
5000	1	225.667	225.62 ± 0.02	225.632 ± 0.00	*
	2	225.571		225.629 ± 0.01	
	3	225.583		225.633 ± 0.01	
	4	225.627		225.625 ± 0.00	
	5	225.624		225.626 ± 0.00	
	6	225.678		225.631 ± 0.00	
	7	225.639		225.633 ± 0.00	
	8	225.576		225.631 ± 0.00	
	9	225.617		225.629 ± 0.00	
	10	225.649		225.629 ± 0.00	

Table 18. LandS, Latin Hypercube sampling.

N	i	$\hat{v}_N^{(i)}$	$\mathbb{E}(\hat{v}_N)$ (95% conf. int.)	$f(\hat{x}_N^{(i)})$ (95% conf. int.)	Best
50	1	5.75	4.11 ± 1.23	13.63 ± 0.05	*
	2	2.54		13.63 ± 0.06	
	3	7.16		15.76 ± 0.09	
	4	3.03		12.95 ± 0.07	
	5	5.74		13.39 ± 0.05	
	6	5.66		14.60 ± 0.05	
	7	2.99		17.97 ± 0.07	
	8	6.44		12.63 ± 0.06	
	9	2.56		15.25 ± 0.05	
	10	2.47		14.79 ± 0.06	
	11	0.85		15.63 ± 0.07	
100	1	6.71	7.66 ± 1.31	11.58 ± 0.04	*
	2	7.56		11.14 ± 0.04	
	3	8.82		11.92 ± 0.04	
	4	9.40		12.76 ± 0.04	
	5	6.37		11.44 ± 0.05	
	6	5.22		14.83 ± 0.06	
	7	6.57		13.71 ± 0.05	
	8	10.79		11.91 ± 0.05	
	9	6.84		11.91 ± 0.05	
	10	4.51		13.28 ± 0.05	
	11	11.44		12.17 ± 0.05	
500	1	8.60	8.54 ± 0.34	10.38 ± 0.05	*
	2	8.08		10.26 ± 0.04	
	3	7.85		10.44 ± 0.05	
	4	9.09		10.45 ± 0.05	
	5	8.50		10.64 ± 0.04	
	6	7.91		10.77 ± 0.05	
	7	9.11		10.86 ± 0.04	
	8	9.32		10.47 ± 0.04	
	9	9.30		10.59 ± 0.05	
	10	8.11		10.80 ± 0.05	
	11	8.11		10.51 ± 0.04	
1000	1	9.07	9.31 ± 0.23	10.04 ± 0.05	*
	2	9.02		10.08 ± 0.04	
	3	9.19		10.19 ± 0.05	
	4	9.81		10.27 ± 0.04	
	5	9.39		10.10 ± 0.04	
	6	9.65		10.12 ± 0.04	
	7	8.68		10.05 ± 0.05	
	8	9.43		10.17 ± 0.03	
	9	9.58		10.07 ± 0.04	
	10	8.83		10.19 ± 0.04	
	11	9.77		10.15 ± 0.04	
5000	1	10.28	9.98 ± 0.21	9.87 ± 0.06	*
	2	9.67		9.89 ± 0.06	
	3	9.68		9.96 ± 0.05	
	4	10.44		9.98 ± 0.04	
	5	9.69		9.94 ± 0.05	
	6	10.38		9.94 ± 0.04	
	7	10.37		9.95 ± 0.04	
	8	9.53		10.02 ± 0.05	
	9	9.98		9.97 ± 0.04	
	10	10.13		9.98 ± 0.04	
	11	9.65		9.92 ± 0.05	

Table 19. ssn, Monte Carlo sampling.

N	i	$\hat{v}_N^{(i)}$	$\mathbb{E}(\hat{v}_N)$ (95% conf. int.)	$f(\hat{x}_N^{(i)})$ (95% conf. int.)	Best
50	1	11.069	10.10 ± 0.81	12.629 ± 0.03	*
	2	9.785		11.467 ± 0.02	
	3	8.236		12.097 ± 0.03	
	4	10.417		11.385 ± 0.02	
	5	12.132		13.065 ± 0.02	
	6	9.816		12.323 ± 0.02	
	7	11.934		11.687 ± 0.02	
	8	9.7		11.562 ± 0.02	
	9	8.565		12.182 ± 0.02	
	10	9.347		12.061 ± 0.02	
100	1	8.491	8.90 ± 0.36	12.051 ± 0.03	*
	2	9.418		10.783 ± 0.03	
	3	9.887		10.813 ± 0.02	
	4	8.627		11.116 ± 0.03	
	5	8.614		10.517 ± 0.02	
	6	8.704		11.134 ± 0.02	
	7	8.741		11.375 ± 0.03	
	8	8.922		11.196 ± 0.02	
	9	9.638		10.918 ± 0.03	
	10	8.002		11.355 ± 0.03	
500	1	9.356	9.87 ± 0.22	10.233 ± 0.03	*
	2	9.886		10.048 ± 0.02	
	3	10.261		10.152 ± 0.02	
	4	10.168		10.326 ± 0.03	
	5	9.957		10.075 ± 0.03	
	6	10.21		10.189 ± 0.02	
	7	9.926		10.19 ± 0.02	
	8	9.687		10.175 ± 0.02	
	9	9.993		10.121 ± 0.02	
	10	9.218		10.236 ± 0.02	
1000	1	9.229	9.83 ± 0.29	10.046 ± 0.03	*
	2	9.907		9.99 ± 0.03	
	3	9.654		10.028 ± 0.03	
	4	9.922		9.975 ± 0.02	
	5	9.427		10.034 ± 0.02	
	6	10.925		10.082 ± 0.03	
	7	9.575		9.988 ± 0.03	
	8	9.694		9.974 ± 0.03	
	9	10.005		10.084 ± 0.02	
	10	10.003		10.099 ± 0.02	
5000	1	9.661	9.84 ± 0.10	9.941 ± 0.02	*
	2	9.974		9.934 ± 0.02	
	3	9.926		9.949 ± 0.02	
	4	10.146		9.93 ± 0.02	
	5	9.832		9.92 ± 0.02	
	6	9.652		9.918 ± 0.03	
	7	9.726		9.903 ± 0.02	
	8	9.82		9.927 ± 0.02	
	9	9.946		9.928 ± 0.03	
	10	9.736		9.901 ± 0.03	

Table 20. ssn, Latin Hypercube sampling.

N	i	$\hat{v}_N^{(i)}$	$\mathbb{E}(\hat{v}_N)$ (95% conf. int.)	$f(\hat{x}_N^{(i)})$ (95% conf. int.)	Best
50	1	15531965.9	15506271.7 \pm 22043.3	15499394.49 \pm 717.29	*
	2	15444631.9		15499858.44 \pm 809.75	
	3	15527925.1		15499137.46 \pm 649.78	
	4	15492684.7		15500436.31 \pm 705.40	
	5	15543334.8		15499410.71 \pm 831.47	
	6	15579111.7		15499016.18 \pm 640.24	
	7	15483314.5		15499775.67 \pm 650.70	
	8	15457776.9		15498142.30 \pm 792.63	
	9	15465732.8		15500048.44 \pm 843.03	
	10	15523480.9		15499029.75 \pm 719.52	
	11	15511112.3		15498506.54 \pm 722.76	
	12	15514189.1		15499466.80 \pm 646.17	
100	1	15510777.4	15482549.9 \pm 19213.8	15499228.78 \pm 553.94	*
	2	15528673.6		15498230.30 \pm 602.99	
	3	15498170.6		15499008.77 \pm 668.22	
	4	15459602.5		15498747.30 \pm 615.50	
	5	15477455.6		15498218.92 \pm 748.83	
	6	15440971.6		15499110.72 \pm 601.46	
	7	15416044.1		15498243.28 \pm 719.41	
	8	15518284.9		15498503.24 \pm 646.60	
	9	15473847.4		15499336.21 \pm 743.68	
	10	15518557.1		15499292.91 \pm 695.26	
	11	15477360.8		15498659.27 \pm 745.56	
	12	15470853.2		15499270.33 \pm 690.87	
500	1	15496050.6	15498139.8 \pm 4152.8	15498528.45 \pm 793.82	*
	2	15489507.6		15499189.25 \pm 704.01	
	3	15499018.5		15499121.32 \pm 688.68	
	4	15494267.3		15499165.45 \pm 775.66	
	5	15491105.1		15497602.67 \pm 1103.38	
	6	15503260.3		15499399.22 \pm 841.89	
	7	15508506.2		15498145.55 \pm 733.54	
	8	15509944.2		15499592.75 \pm 733.37	
	9	15486782.2		15498308.67 \pm 1002.77	
	10	15498430.0		15497527.86 \pm 1016.05	
	11	15504652.2		15498343.72 \pm 1024.06	
	12	15496153.1		15499285.83 \pm 1057.83	
1000	1	15498653.2	15500879.4 \pm 4847.1	15498803.79 \pm 686.55	*
	2	15520761.5		15499145.13 \pm 795.08	
	3	15504128.1		15498277.13 \pm 590.80	
	4	15497767.9		15498242.11 \pm 804.38	
	5	15504635.1		15498768.54 \pm 997.55	
	6	15491787.7		15498710.92 \pm 1001.64	
	7	15509053.8		15498785.86 \pm 902.27	
	8	15493874.9		15498893.94 \pm 1055.03	
	9	15492759.0		15498632.15 \pm 944.34	
	10	15491643.7		15498949.08 \pm 1036.71	
	11	15499917.0		15499356.79 \pm 861.47	
	12	15505570.6		15499461.20 \pm 871.86	
5000	1	15501515.1	15498121.3 \pm 1879.0	15498451.74 \pm 553.28	*
	2	15499211.4		15498824.11 \pm 634.32	
	3	15503271.1		15498747.77 \pm 720.78	
	4	15496814.0		15498836.80 \pm 711.78	
	5	15500288.4		15498459.42 \pm 780.61	
	6	15499461.5		15498279.30 \pm 774.23	
	7	15497473.5		15499035.93 \pm 835.64	
	8	15497455.0		15499447.97 \pm 779.99	
	9	15500551.3		15498035.30 \pm 628.43	
	10	15495827.0		15498760.45 \pm 612.69	
	11	15491199.8		15499014.30 \pm 623.51	
	12	15494386.9		15499530.54 \pm 681.05	

Table 21. storm, Monte Carlo sampling.

N	i	$\hat{v}_N^{(i)}$	$\mathbb{E}(\hat{v}_N)$ (95% conf. int.)	$f(\hat{x}_N^{(i)})$ (95% conf. int.)	Best
50	1	15496526.2	15497683.7 ± 1078.8	15498895.8 ± 18.76	*
	2	15500293.0		15498819.1 ± 21.67	
	3	15497849.4		15498882.3 ± 20.39	
	4	15497800.4		15498773.4 ± 16.51	
	5	15495986.3		15498741.4 ± 16.95	
	6	15497023.0		15498739.4 ± 19.10	
	7	15494804.3		15498721.2 ± 17.05	
	8	15500032.9		15498786.7 ± 17.48	
	9	15499062.3		15498733.7 ± 18.51	
	10	15497459.0		15498862.7 ± 17.85	
100	1	15500047.4	15499255.3 ± 1011.7	15498713.9 ± 20.10	*
	2	15498546.3		15498809.2 ± 19.41	
	3	15500509.7		15498729.2 ± 23.49	
	4	15500133.5		15498733.7 ± 18.36	
	5	15499785.9		15498724.8 ± 15.71	
	6	15498865.8		15498733.1 ± 17.13	
	7	15497825.8		15498783.3 ± 16.26	
	8	15497601.0		15498760.8 ± 21.35	
	9	15502359.7		15498755.8 ± 17.96	
	10	15496877.4		15498846.8 ± 19.48	
500	1	15498139.5	15498661.4 ± 280.8	15498745.9 ± 19.70	*
	2	15498694.2		15498744.0 ± 18.19	
	3	15499566.1		15498732.1 ± 18.98	
	4	15497929.2		15498725.3 ± 18.45	
	5	15498479.6		15498736.4 ± 24.17	
	6	15499010.3		15498746.0 ± 16.81	
	7	15498502.5		15498807.0 ± 23.88	
	8	15498701.5		15498724.7 ± 15.99	
	9	15498763.5		15498719.9 ± 22.71	
	10	15498827.9		15498735.5 ± 15.69	
1000	1	15498761.3	15498598.1 ± 148.5	15498752.0 ± 18.87	*
	2	15498188.8		15498740.8 ± 19.52	
	3	15498737.8		15498727.7 ± 20.36	
	4	15498497.0		15498773.9 ± 18.92	
	5	15498394.0		15498745.0 ± 21.36	
	6	15498858.9		15498733.1 ± 16.58	
	7	15498416.5		15498713.3 ± 17.30	
	8	15498567.1		15498740.4 ± 18.12	
	9	15498575.3		15498728.4 ± 16.98	
	10	15498983.9		15498737.3 ± 18.66	
5000	1	15498558.9	15498657.8 ± 73.9	15498728.4 ± 19.84	*
	2	15498709.6		15498733.5 ± 18.94	
	3	15498725.4		15498732.9 ± 20.16	
	4	15498486.7		15498732.7 ± 19.00	
	5	15498820.6		15498728.4 ± 15.04	
	6	15498737.1		15498728.8 ± 19.08	
	7	15498718.5		15498722.1 ± 18.16	
	8	15498728.3		15498743.8 ± 18.06	
	9	15498633.5		15498722.4 ± 22.42	
	10	15498459.0		15498730.7 ± 17.17	

Table 22. storm, Latin Hypercube sampling.




Fig. 1. 20term: Means and 95% Confidence Intervals for Optimal Value Bound Estimates




Fig. 2. gbd: Means and 95% Confidence Intervals for Optimal Value Bound Estimates




Fig. 3. LandS: Means and 95% Confidence Intervals for Optimal Value Bound Estimates




Fig. 4. SSN: Means and 95% Confidence Intervals for Optimal Value Bound Estimates




Fig. 5. storm: Means and 95% Confidence Intervals for Optimal Value Bound Estimates