

NPS-OR-94-003

2

NAVAL POSTGRADUATE SCHOOL

Monterey, California

AD-A278 654



DTIC
ELECTE
APR 26 1994
S B D

STOPPING RULES FOR A CLASS OF SAMPLING-BASED STOCHASTIC PROGRAMMING ALGORITHMS

David P. Morton

January 1994

Approved for public release; distribution is unlimited.

Prepared for:
National Research Council
2101 Constitution Avenue
Washington, DC 20418

DTIC/NSA/CSS/STO

94-12691



94 4 25 103

NAVAL POSTGRADUATE SCHOOL
MONTEREY, CA 93943-5000


Rear Admiral T. A. Mercer
Superintendent

Harrison Shull
Provost

This report was prepared for the Naval Postgraduate School and funded by the National Research Council.


Reproduction of all or part of this report is authorized.

This report was prepared by:




DAVID P. MORTON
NRC Fellow

Reviewed by:



PETER PURDUE
Professor and Chairman
Department of Operations Research

Released by:



PAUL J. MARIO
Dean of Research

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE January 1994		3. REPORT TYPE AND DATES COVERED Technical
4. TITLE AND SUBTITLE Stopping Rules for a Class of Sampling-Based Stochastic Programming Algorithms			5. FUNDING NUMBERS	
6. AUTHOR(S) David P. Morton				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School Monterey, CA 93943			8. PERFORMING ORGANIZATION REPORT NUMBER NPS-OR-94-003	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) National Research Council 2101 Constitution Avenue Washington, DC 20418			10. SPONSORING / MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the author and do not necessarily reflect the views of the Department of Defense and the U. S. Government.				
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) Decomposition and Monte Carlo sampling-based algorithms hold much promise for solving stochastic programs with many scenarios. A critical component of such algorithms is a stopping criterion to ensure the quality of the solution. In this paper, we develop a stopping rule theory for a class of algorithms that estimate bounds on the optimal objective function value by sampling. We provide rules for selecting sample sizes and terminating the algorithm under which asymptotic validity of confidence intervals for the quality of the proposed solution can be verified. These rules are applied to a multistage stochastic linear programming algorithm due to Pereira and Pinto. DTIC (unclassified) Avail: GPO 5				
14. SUBJECT TERMS Stopping Rules, Monte Carlo Sampling, Stochastic Programming			15. NUMBER OF PAGES 28	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UL	

Stopping Rules for a Class of Sampling-Based Stochastic Programming Algorithms

by

David P. Morton

Abstract

Decomposition and Monte Carlo sampling-based algorithms hold much promise for solving stochastic programs with many scenarios. A critical component of such algorithms is a stopping criterion to ensure the quality of the solution. In this paper, we develop a stopping rule theory for a class of algorithms that estimate bounds on the optimal objective function value by sampling. We provide rules for selecting sample sizes and terminating the algorithm under which asymptotic validity of confidence intervals for the quality of the proposed solution can be verified. These rules are applied to a multistage stochastic linear programming algorithm due to Pereira and Pinto.

Keywords: Stopping Rules, Monte Carlo Sampling, Stochastic Programming

January 1994

Department of Operations Research
Naval Postgraduate School
Monterey, CA 93943-5000

Accession For	
NTIS GRA&I	☒
DTIC TAB	☐
Unannounced	☐
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	

This work was supported by the National Research Council under a Research Associateship at the Naval Postgraduate School, Monterey, California.

1 Introduction

Practical planning problems with deterministic forecasts of inherently uncertain parameters often yield unsatisfactory solutions. Stochastic programming formulations allow uncertain parameters to be modeled as random variables with known distribution, but the size of the resulting mathematical programs can be formidable. Dantzig [5] and Beale [1] introduced stochastic programming with recourse; some example applications from the literature include capacity expansion planning [7], forest harvest planning [12], hydroelectric scheduling [20,27], and portfolio management [22,26].

Stochastic programming algorithms are typically one of three types: (i) exact solution procedures, (ii) approximation and bounding schemes, and (iii) sampling-based methods. Exact methods include simplex-based algorithms that exploit special structure of bases [17], decomposition or L-shaped schemes [2,31], interior point methods [24], and the Progressive Hedging algorithm [29]. A classic approximation scheme involves calculating deterministic lower and upper bounds via the inequalities of Jensen and Edmundson-Madansky, respectively; see [3,4,11,21] for extensions and alternatives. Stochastic quasigradient (SQG) methods [9,13] are sampling-based. Another set of sampling-based algorithms are rooted in the L-shaped method [6,14,19,27]. In many models, as the number of random parameters and the number of scenarios grow, exact solution procedures and approximation and bounding schemes become more difficult to apply due to required computational effort. Sampling-based algorithms may provide an attractive alternative for such models. Stopping criteria for SQG methods are examined in [28]. In general, however, sampling-based approaches lack stopping rules that can control *a priori* the quality of the proposed solution. In this paper, we develop rules designed to rectify this shortcoming for a particular class of sampling-based algorithms.

A host of questions arise when one replaces deterministic upper and lower objective function bounds generated by a decomposition algorithm with estimates formed from sample means. How should an "optimal" solution be characterized and how should the sampling procedure proceed so as to ensure an appropriately defined notion of convergence? The primary purpose of this paper is to provide a framework in which these issues may be addressed. The stopping rules we develop are comprised of two components: (i) a criterion for terminating the algorithm and (ii) a rule for selecting the sample sizes. The main results, detailed in §2, provide stopping rules that guarantee

asymptotic validity of the desired confidence interval statements as the interval width shrinks and the sample sizes grow. Application of these methods to an algorithm for a class of multistage stochastic linear programs [27] is described in §3. Empirical coverage results for a simple example are given in §4, and the paper is summarized in §5.

2 Stopping Rules

This section begins by outlining our framework of assumptions on the underlying sampling-based algorithm. Consider the following general optimization problem:

$$\begin{aligned} z^* = & \text{minimize } z(x) \\ & \text{subject to } x \in \mathcal{X}. \end{aligned} \quad (1)$$

Suppose we have at hand an algorithm that at each iteration k selects a *sample size* n_k , produces a feasible decision x_k , and generates estimates for upper and lower bounds on the optimal objective function value denoted $U_k(n_k)$ and $L_k(n_k)$, respectively. It is assumed that at each iteration k the *difference random variable*, $D_k(n_k) = U_k(n_k) - L_k(n_k)$, satisfies the central limit theorem (CLT):

$$\sqrt{n_k}(D_k(n_k) - \mu_k) \Rightarrow N(0, \sigma_k^2) \text{ as } n_k \rightarrow \infty, \quad \text{where } \sigma_k > 0 \quad (2)$$

and where $\Rightarrow$ denotes convergence in distribution; $N(\mu, \sigma^2)$ denotes a normal random variable with mean μ and variance σ^2 . The sample size parameter will typically be suppressed when referring to the upper, lower, and difference random variables. The true upper and lower bounds at iteration k are denoted u_k and l_k and satisfy $U_k \Rightarrow u_k$ and $L_k \Rightarrow l_k$ as $n_k \rightarrow \infty$, where $l_k \leq z^* \leq u_k$, and (necessarily) $\mu_k = u_k - l_k$. If U_k and L_k are independent and satisfy respective CLT's then hypothesis (2) follows as a consequence.

The algorithm is terminated on the first iteration, T , in which the difference random variable drops below zero; i.e.,

$$T = \inf_{k \geq 1} \{k : D_k \leq 0\}. \quad (3)$$

The feasible decision x_T generated at the random stopping iteration satisfies $u_T \geq z(x_T)$. In addition, we assume that, at the stopping iteration, the algorithm permits re-evaluation of the difference random variable through independent resampling. The algorithm is said to *stop correctly*

if $z(x_T) \leq z^* + \epsilon$ where ϵ is a positive *confidence interval width*. Stopping correctly is ensured if $\mu_T \leq \epsilon$; we use this observation and the CLT hypothesis (2) to prescribe sample sizes, n_k , under which a statement can be made regarding the probability that the algorithm stops correctly. At the heart of the procedures we develop is the idea that the sample sizes must increase as the algorithm proceeds.

In order to illustrate the underlying ideas in a simple setting, we assume in §2.1 that the difference random variables at each iteration are normally distributed. Under this assumption, we show it suffices to increase the sample size at a rate proportional to $\log k$ to guarantee that the probability of stopping correctly satisfies a prescribed confidence level. In §2.2, the normality assumption is replaced with the CLT hypothesis (2) and an $\mathcal{O}(\log^2 k)$ sample-size formula is provided under which the results of §2.1 can be recovered in the form of an asymptotic validity result. Moreover, we indicate that with respect to required computational effort, the $\mathcal{O}(\log^2 k)$ sample-size formula is preferable to the $\mathcal{O}(\log k)$ formula. In §2.3 we verify the asymptotic validity result under a weaker history-dependent CLT hypothesis in which the difference random variable may depend on the (potentially) random history of the algorithm in previous iterations. In §2.4 we address issues associated with finite stopping times and incorporation of sample variance estimators.

2.1 Stopping Rules: Normal Differences

In this subsection we replace the CLT hypothesis (2) with the following more restrictive assumption: At the k^{th} iteration of the algorithm, we choose a sample size n_k and then observe the random variable:

$$D_k \sim N(\mu_k, \sigma_k^2/n_k), \quad \text{where } \sigma_k > 0. \quad (4)$$

Example 1 This example indicates that a fixed sample size, $n_k = n$, can lead to unattractive results. Suppose $u_k = \mu_k = \mu > \epsilon$ for $k = 1, \dots, K$; $\mu_k = 0$ for $k \geq K + 1$; and $\sigma_k = \sigma > 0$ for all k . The algorithm will not stop correctly if and only if it terminates prior to iteration $K + 1$. By choosing K sufficiently large we can make the probability we stop correctly, $P\{D_1 > 0, \dots, D_K > 0\} = [P\{D_1 > 0\}]^K$, arbitrarily close to 0. ■

Theorem 2 provides a sample-size formula that overcomes the difficulty illustrated in Example 1; in particular, it ensures that the probability the algorithm stops correctly satisfies a minimum prescribed confidence level of $1 - \alpha$; e.g., 0.95. We require the following lemma regarding bounds on the tail of a normal distribution (see Feller Chapter VII §1 [10]).

Lemma 1 *If $z > 0$ then*

$$P\{N(0, 1) \geq z\} \leq \frac{1}{\sqrt{2\pi}} \frac{1}{z} e^{-z^2/2}.$$

Theorem 2 *Assume (3), (4), and define*

$$n_k \geq \left(\frac{\sigma_k}{\epsilon}\right)^2 (\beta + 2p \ln k), \quad (5)$$

where $\epsilon > 0$, $\beta = \max\{2 \ln [\zeta(p)/(\sqrt{2\pi} \alpha)], 1\}$, $0 < \alpha < 1$, and $\zeta(p) = \sum_{k=1}^{\infty} k^{-p}$, $p > 1$. Then $P\{\mu_T \leq \epsilon\} \geq 1 - \alpha$.

Proof

$D_T \leq 0$ implies

$$\begin{aligned} P\{\mu_T \geq \epsilon\} &\leq P\{\mu_T \geq D_T + \epsilon\} \\ &= \sum_{k=1}^{\infty} P\{D_1 > 0, \dots, D_{k-1} > 0, D_k \leq 0, D_k \leq \mu_k - \epsilon\} \\ &\leq \sum_{k=1}^{\infty} P\{D_k \leq \mu_k - \epsilon\}. \end{aligned}$$

Now $\sum_{k=1}^{\infty} [\alpha/\zeta(p)] k^{-p} = \alpha$; thus, it suffices to show

$$P\{D_k \leq \mu_k - \epsilon\} \leq \frac{\alpha}{\zeta(p)} k^{-p}.$$

To this end consider

$$\begin{aligned} P\{D_k \leq \mu_k - \epsilon\} &= P\left\{\frac{D_k - \mu_k}{\sigma_k/\sqrt{n_k}} \leq -\epsilon\sqrt{n_k}/\sigma_k\right\} \\ &= P\{N(0, 1) \geq \epsilon\sqrt{n_k}/\sigma_k\} \\ &\leq P\{N(0, 1) \geq (\beta + 2p \ln k)^{1/2}\} \end{aligned}$$

Since $\beta \geq 1$, the tail bound from Lemma 1 yields:

$$P\{D_k \leq \mu_k - \epsilon\} \leq \frac{1}{\sqrt{2\pi}} \exp\left[-\frac{1}{2}(\beta + 2p \ln k)\right].$$

Hence it suffices to show

$$\frac{1}{\sqrt{2\pi}} \exp \left[-\frac{1}{2} (\beta + 2p \ln k) \right] \leq \frac{\alpha}{\zeta(p)} k^{-p},$$

and this inequality follows from the definition of β . ■

Remark 1 The coverage result of Theorem 2 states that $[0, \epsilon]$ is a $(1 - \alpha) \cdot 100\%$ confidence interval for μ_T . In terms of the optimization problem (1) this implies:

$$P \{ z(x_T) \leq z^* + \epsilon \} \geq 1 - \alpha.$$

Remark 2 Tables of values of the Riemann-Zeta function, $\zeta(p)$, may be found in Dwight [8]; $\zeta(p)$ is also an intrinsic function in some mathematical software packages such as *Mathematica* [32]. We return to the topic of choosing values for the parameter p to minimize the total number of samples required in §2.2.

Remark 3 The term “ k^{th} iteration” should be interpreted liberally; a better phrase is “ k^{th} stopping cycle” with one possible definition as follows: A stopping cycle consists of a number of algorithm iterations in which a fixed sample size is used *plus* one resampling iteration in which the sample-size formula (5) is applied. The fixed sample-size phase of a stopping cycle is terminated when a heuristic pre-test is passed; e.g., a negative difference is observed. The sample-size formula and stopping criterion described above are then applied to a realization of the difference random variable generated from an independent sample. This idea, illustrated in the application of §3, helps to minimize computational burden.

Remark 4 The purpose of our analysis is to control the quality of the proposed solution x_T . Any number of heuristic stopping rules can generate a feasible solution, $x_{T'}$. By resampling the associated difference random variable, one can form a $(1 - \delta) \cdot 100\%$ confidence interval of the form $[z^*, z^* + D_{T'}^+ + w_\delta \sigma_{T'} / \sqrt{n}]$ for $z(x_{T'})$ where n denotes the resampling sample size, w_δ satisfies $P\{N(0, 1) \leq w_\delta\} = 1 - \delta$, and $D_{T'}^+ \equiv \max \{D_{T'}, 0\}$. (We use $D_{T'}^+$ because it is known that $\mu_{T'} \geq 0$.) The disadvantage of results based on heuristic stopping rules is that we have inadequate control of the interval width, $D_{T'}^+ + w_\delta \sigma_{T'} / \sqrt{n}$. The primary purpose of Theorem 2 (and the subsequent generalizations we present in this paper) is to provide *a priori* control on the confidence interval width.

Remark 5 We regard the specific values of z^* and $z(x_T)$ of secondary interest relative to controlling the quality of the proposed solution, x_T . However, if U_T and L_T are each normally distributed (or more generally satisfy respective CLT's) then one can also develop confidence intervals for these quantities through the resampling procedure.

Remark 6 The coverage result of Theorem 2 does not depend on any convergence structure of the optimization algorithm. The sample-size formula is designed so that if the algorithm does not converge then the stopping criterion will not be satisfied with probability, at least, $1 - \alpha$. This property contributes to the conservative nature of the coverage result since many algorithms have some underlying convergence structure. We provide conditions under which finite termination can be ensured (with probability one) in §2.4.

2.2 Stopping Rules: CLT Differences

In §2.1, under the assumption of normally distributed difference random variables, we derived a valid confidence interval for all positive interval widths, ϵ . In this subsection, we replace the normality assumption with an asymptotic normality hypothesis (2) and provide conditions under which the confidence interval of §2.1 is valid in the limit as the interval width shrinks and the sample sizes grow. In particular, we show

$$\lim_{\epsilon \downarrow 0} P \{ z(x_{T(\epsilon)}) \leq z^* + \epsilon \} \geq 1 - \alpha, \quad (6)$$

where the stopping time, $T = T(\epsilon)$ is once again the first iteration in which the observed difference drops below zero.

Example 1 was appropriate under the assumptions of §2.1 only because the confidence interval statement was made for all positive interval widths; for sufficiently small ϵ , adequate coverage results are achieved. In Example 2 we again use an identical sample size for all iterations and construct a problem in which the probability of stopping correctly is zero in the limit as the interval width shrinks to zero.

Example 2 Let $D_k \sim N(k^{-1}, n_k^{-1})$ be independent and define $m = \lceil \epsilon^{-1} \rceil - 1$. (The ceiling operator, $\lceil \cdot \rceil$, yields the smallest integer greater than or equal to its argument.) If the algorithm terminates

prior to iteration $m + 1$ it has stopped incorrectly. Thus,

$$\begin{aligned} P\{\mu_{T(\epsilon)} \leq \epsilon\} &\leq P\{D_1 > 0, \dots, D_m > 0\} \\ &= \prod_{k=1}^m P\{N(0, 1) < k^{-1}\sqrt{n_k}\}. \end{aligned}$$

Let $\epsilon = Cn_k^{-1/2} > 0$ so that the sample size has the same value for all iterations and observe:

$$\begin{aligned} P\{\mu_{T(\epsilon)} \leq \epsilon\} &\leq \prod_{k=1}^m P\{N(0, 1) < k^{-1}(m+1)C\} \\ &\leq \prod_{k=\lceil m/2 \rceil}^m P\{N(0, 1) < (m+1)C/\lceil m/2 \rceil\} \\ &\leq [P\{N(0, 1) < 3C\}]^{\lceil m/2 \rceil}. \end{aligned}$$

Hence $P\{\mu_{T(\epsilon)} \leq \epsilon\} \rightarrow 0$ as $\epsilon \downarrow 0$. ■

Example 2 demonstrates that a fixed sample size, no matter how large, may lead to "confidence" intervals with unsatisfactory coverage properties. The key idea, once again, is that we must increase the sample size as the algorithm proceeds; Theorem 3 provides stopping rules under which the desired confidence intervals are asymptotically valid.

Theorem 3 Assume (2), (3), and

$$\sup_{k \geq 1} E \exp \left[\gamma \left(\frac{D_k - \mu_k}{\sigma_k / \sqrt{n_k}} \right) \right] \text{ is bounded for } |\gamma| \leq \gamma_0. \quad (7)$$

Let $\beta' = \max \{2 \ln [\phi(p)/(\sqrt{2\pi} \alpha)], 1\}$, where $\phi(p) = \sum_{k=1}^{\infty} k^{-p \ln k}$, $p > (2\gamma_0^2)^{-1}$ and $0 < \alpha < 1$.

If

$$n_k \geq \left(\frac{\sigma_k}{\epsilon} \right)^2 (\beta' + 2p \ln^2 k) \quad (8)$$

then $\lim_{\epsilon \downarrow 0} P\{\mu_{T(\epsilon)} \leq \epsilon\} \geq 1 - \alpha$.

Proof

Let $Z_k = \sigma_k^{-1} \sqrt{n_k} (D_k - \mu_k)$. We begin as in the proof of Theorem 2 and see that it suffices to show

$$\lim_{\epsilon \downarrow 0} \sum_{k=1}^{\infty} P\{Z_k \leq -(\beta' + 2p \ln^2 k)^{1/2}\} \leq \alpha. \quad (9)$$

Next we show the order of the limit and sum on the left hand side of (9) may be exchanged with equality by employing the dominated convergence theorem. Applying Markov's inequality (see, e.g., Loève §9 [23]) to $e^{-\gamma Z_k}$, $\gamma > 0$ we have

$$P \left\{ Z_k \leq -(\beta' + 2p \ln^2 k)^{1/2} \right\} \leq \exp \left[-\gamma(\beta' + 2p \ln^2 k)^{1/2} \right] E e^{-\gamma Z_k}. \quad (10)$$

With $\gamma = \gamma_0$, the right hand side of (10) is bounded above by $\left\{ \sup_{k \geq 1} E e^{-\gamma_0 Z_k} \right\} k^{-\gamma_0 \sqrt{2p}}$. As $\gamma_0 \sqrt{2p} > 1$, the order of the limit and sum may be exchanged with equality and it remains to show

$$\sum_{k=1}^{\infty} \lim_{\epsilon \downarrow 0} P \left\{ Z_k \leq -(\beta' + 2p \ln^2 k)^{1/2} \right\} \leq \alpha.$$

By applying the CLT (2) we may complete the proof in a fashion analogous to that of Theorem 2. ■

Under a weaker hypothesis on the difference random variables, we have recovered the coverage result of Theorem 2 in an asymptotic sense. The asymptotic validity result of Theorem 3 may be interpreted as follows: For sufficiently small choices of ϵ , $[z^*, z^* + \epsilon]$ is an *approximate* $(1 - \alpha) \cdot 100\%$ confidence interval for $z(x_T)$.

We now address two issues with respect to the hypotheses of Theorem 3. First, we provide conditions under which the bounded expectation assumption (7) holds and then we examine the function $\phi(p)$. Verification of (7) is straightforward when each difference random variable may be expressed as a sample mean of independent and identically distributed random variables (i.i.d.r.v.'s): in particular, it suffices to verify that the underlying random variables have moment generating functions for $|\gamma| \leq \gamma_0$. Indeed, this observation is at the core of an elementary proof of the central limit theorem for i.i.d.r.v.'s (see, e.g., Hogg and Craig [16]). Proposition 4 summarizes this result; we use the following notation: For each k , $D_{k1}, D_{k2}, \dots$ denote i.i.d.r.v.'s, $D_k = \frac{1}{n_k} \sum_{i=1}^{n_k} D_{ki}$, $ED_{k1} = \mu_k$, and $E(D_{k1} - \mu_k)^2 = \sigma_k^2$ where $0 < \sigma_k^2 < \infty$.

Proposition 4 *If $\sup_{k \geq 1} E \exp \left[\gamma \left(\frac{D_{k1} - \mu_k}{\sigma_k} \right) \right] < \infty$ then $\sup_{k \geq 1} E \exp \left[\gamma \left(\frac{D_k - \mu_k}{\sigma_k / \sqrt{n_k}} \right) \right]$ is bounded.*

We now turn our attention to $\phi(p) = \sum_{k=1}^{\infty} k^{-p \ln k}$. By comparing terms of the $\phi(p)$ series and $\sum_{k=1}^{\infty} k^{-2}$ for sufficiently large k we conclude $\phi(p) < \infty$ if $p > 0$. Proposition 5 provides bounds on $\phi(p)$ that facilitate numerical function evaluations and further characterizes both $\phi(p)$ and $\zeta(p)$ so that we may subsequently address the issue of choosing good values for the parameter p .

Proposition 5

(i) Let $v_N = \sqrt{2p} [\ln N - (2p)^{-1}]$ and $K_p = \sqrt{\pi/p} \exp [(4p)^{-1}]$, $p > 0$. Then

$$K_p P \{N(0, 1) \geq v_{N+1}\} \leq \phi(p) - \sum_{k=1}^N k^{-p \ln k} \leq K_p P \{N(0, 1) \geq v_N\}$$

and

$$K_p (P \{N(0, 1) \geq v_N\} - P \{N(0, 1) \geq v_{N+1}\}) \leq N^{-p \ln N}.$$

(ii) $\ln \zeta(p)$ and $\ln \phi(p)$ are convex functions on $(1, \infty)$ and $(0, \infty)$, respectively.

Proof

Based on the inequalities

$$(k+1)^{-p \ln(k+1)} \leq \int_k^{k+1} u^{-p \ln u} du \leq k^{-p \ln k}$$

we find

$$\int_{N+1}^{\infty} u^{-p \ln u} du \leq \phi(p) - \sum_{k=1}^N k^{-p \ln k} \leq \int_N^{\infty} u^{-p \ln u} du.$$

A change of variables yields

$$\int_y^{\infty} u^{-p \ln u} du = (2p)^{-1/2} \exp [(4p)^{-1}] \int_{v_y}^{\infty} e^{-u^2/2} du = K_p P \{N(0, 1) \geq v_y\},$$

where $v_y = \sqrt{2p} [\ln y - (2p)^{-1}]$. This establishes the first part of (i), and the error bound follows from

$$\int_N^{\infty} u^{-p \ln u} du - \int_{N+1}^{\infty} u^{-p \ln u} du \leq N^{-p \ln N}.$$

For result (ii) we require the infinite series version of Hölder's inequality:

$$\sum_{k=1}^{\infty} |a_k b_k| \leq \left[\sum_{k=1}^{\infty} |a_k|^r \right]^{1/r} \left[\sum_{k=1}^{\infty} |b_k|^s \right]^{1/s}$$

where $r, s > 1$ and $1/r + 1/s = 1$. For the Riemann-Zeta function we must show

$$\ln \zeta(\lambda p_1 + (1 - \lambda)p_2) \leq \lambda \ln \zeta(p_1) + (1 - \lambda) \ln \zeta(p_2).$$

Using the definition of $\zeta(p)$ we see it suffices to show:

$$\sum_{k=1}^{\infty} (k^{-p_1})^{\lambda} (k^{-p_2})^{1-\lambda} \leq \left[\sum_{k=1}^{\infty} k^{-p_1} \right]^{\lambda} \left[\sum_{k=1}^{\infty} k^{-p_2} \right]^{1-\lambda}.$$

This inequality follows from Hölder's inequality. The proof for $\ln \phi(p)$ is virtually identical. ■

The $\mathcal{O}(\log^2 k)$ sample-size formula of Theorem 3 is driven, in large part, by the moment generating function hypothesis (7). By making the stronger assumption,

$$\sup_{k \geq 1} E \exp \left[\gamma \left(\frac{D_k - \mu_k}{\sigma_k / \sqrt{n_k}} \right)^2 \right] \text{ is bounded for } |\gamma| \leq \gamma_0,$$

the $\mathcal{O}(\log k)$ sample-size formula of Theorem 2 can be recovered; the difficulty, however, lies in developing analogs of Proposition 4. The moment generating function hypothesis is attractive because it is easily verified *and* because with respect to required computational effort, the $\mathcal{O}(\log^2 k)$ sample-size formula is preferable to the $\mathcal{O}(\log k)$ formula for practical problems. For the purpose of verifying this statement we will assume σ_k/ϵ is constant as would be the case if ϵ is selected proportional to σ_k (see subsequent Remark 7). Minimizing the total number of samples or "work" over T iterations then corresponds to minimizing

$$W_\zeta(p) = T \cdot \max \left\{ 2 \ln \left[\zeta(p) / (\sqrt{2\pi} \alpha) \right], 1 \right\} + 2p \left(\sum_{k=1}^T \ln k \right)$$

for sample-size formula (5) or

$$W_\phi(p) = T \cdot \max \left\{ 2 \ln \left[\phi(p) / (\sqrt{2\pi} \alpha) \right], 1 \right\} + 2p \left(\sum_{k=1}^T \ln^2 k \right)$$

for sample-size formula (8). From Proposition 5, part (ii) it is clear that $W_\zeta(p)$ and $W_\phi(p)$ are convex functions on $(1, \infty)$ and $(0, \infty)$, respectively. Table 1 displays the results of minimizing the respective work functions for various choices of T with $\alpha = 0.05$. While T , of course, is unknown *a priori*, rough estimates for T (or ET) may be available for certain classes of problems.

T	formula (5) $\mathcal{O}(\log k)$			formula (8) $\mathcal{O}(\log^2 k)$		
	p^*	$\zeta(p^*)$	$W_\zeta(p^*)$	p^*	$\phi(p^*)$	$W_\phi(p^*)$
10	1.5	2.612 375	106	0.4	5.048 588	96
25	1.35	3.459 237	323	0.25	9.379 868	291
50	1.29	4.046 193	731	0.19	14.865 183	660
10^2	1.24	4.761 075	1 630	0.155	22.270 678	1 473
10^3	1.15	7.254 695	21 715	0.09	94.647 997	19 720
10^4	1.11	9.676 075	269 211	0.065	325.046 04	246 153
10^5	1.09	11.694 841	3 199 021	0.050	1 175.999 4	2 944 556

Table 1: Optimal Choices of p

For each value of T , Table 1 contains the minimized (within 0.1%) work function values, the corresponding minimizers, p^* , and $\zeta(p^*)$ and $\phi(p^*)$ for reference. Table 1 indicates that the $\mathcal{O}(\log^2 k)$ sample-size formula outperforms the $\mathcal{O}(\log k)$ formula for values of T dramatically larger than of practical interest (see Remark 3). The reason is as follows: Sample-size formula (5) requires $p > 1$, and this leads to a multiplier greater than 2 on the $\sum_{k=1}^T \ln k$ term of the work function while the corresponding term for $W_\phi(p)$ can have any positive coefficient. Theorem 2, of course, could be restated with the $\mathcal{O}(\log^2 k)$ formula; the bounding series $\zeta(p^*)$ was initially selected because it lead to a better rate of growth. When applying Theorem 3, one may not be able to use the recommended values of p^* due to the $p > (2\gamma_o^2)^{-1}$ condition. In many practical problems, however, the underlying random variables may be bounded, and under the i.i.d. hypotheses of Proposition 4 one can then choose any $p > 0$.

Well-designed stopping cycles (Remark 3) help minimize required computational effort by reducing the number of times that the sample size is increased. However, we wish to emphasize that even under poorly designed stopping cycle schemes the increase in computational effort as the algorithm proceeds is relatively modest. As an example, consider $\epsilon = 2\sigma_k/\sqrt{30}$ and $\alpha = 0.05$. The sample sizes required by the $\mathcal{O}(\log^2 k)$ sample-size formula with $p = 0.155$ at the 1st, 10th, 100th, and 1000th stopping cycles are 78, 91, 128, and 189.

2.3 Stopping Rules: History-Dependent CLT Differences

We can generalize the results of Theorem 3 with respect to the assumptions on the interaction of the algorithm with the upper and lower bound estimates. In the development above, we assume $\{\mu_k, \sigma_k\}_{k=1}^\infty$ is simply a sequence of constants, but in many applications these parameters may depend on the random history of the algorithm through iteration $k - 1$ which we denote $\mathcal{H}_{k-1}$. For example, in an L-shaped algorithm for two stage stochastic programs in which cuts are obtained by sampling (see Dantzig and Glynn [6]), the sequence of master program decisions depends on which scenarios were (randomly) selected to compute cuts in previous iterations. In this example, it is clear μ_k , σ_k , and D_k are random variables that are sample-path dependent. A realization of $\mathcal{H}_k$ may be thought of as the information necessary to reconstruct exactly the steps of the algorithm through iteration k . In this more general setting, we will assume that conditioned on the history

random variable, the mean and variance are constants and the difference random variables satisfy

$$P \left\{ \frac{D_k - \mu_k}{\sigma_k / \sqrt{n_k}} \leq y \mid \mathcal{H}_{k-1} \right\} \rightarrow \frac{1}{\sqrt{2\pi}} \int_{-\infty}^y e^{-u^2/2} du. \quad (11)$$

Theorem 6 Assume (9), (7), (11), and define n_k and the corresponding parameters as in Theorem 3. Then $\lim_{\epsilon \downarrow 0} P \{ \mu_{T(\epsilon)} \leq \epsilon \} \geq 1 - \alpha$.

Proof

Let $Z_k = \sigma_k^{-1} \sqrt{n_k} (D_k - \mu_k)$. We begin as in Theorem 2 and see it suffices to show

$$\lim_{\epsilon \downarrow 0} \sum_{k=1}^{\infty} P \left\{ Z_k \leq -(\beta' + 2p \ln^2 k)^{1/2} \right\} \leq \alpha.$$

From the proof of Theorem 3 we know hypothesis (7) permits exchanging the order of limit and summation with equality. Now observe

$$\begin{aligned} & \sum_{k=1}^{\infty} \lim_{\epsilon \downarrow 0} \int P \left\{ Z_k \leq -(\beta' + 2p \ln^2 k)^{1/2} \mid \mathcal{H}_{k-1} \right\} dP_{\mathcal{H}_{k-1}} \\ &= \sum_{k=1}^{\infty} \int \lim_{\epsilon \downarrow 0} P \left\{ Z_k \leq -(\beta' + 2p \ln^2 k)^{1/2} \mid \mathcal{H}_{k-1} \right\} dP_{\mathcal{H}_{k-1}} \\ &= \sum_{k=1}^{\infty} \int P \left\{ N(0, 1) \leq -(\beta' + 2p \ln^2 k)^{1/2} \right\} dP_{\mathcal{H}_{k-1}} \\ &= \sum_{k=1}^{\infty} P \left\{ N(0, 1) \leq -(\beta' + 2p \ln^2 k)^{1/2} \right\}. \end{aligned}$$

The remainder of the proof is analogous to that of Theorem 2. ■

Remark 7 The asymptotic validity results of Theorems 3 and 6 both still hold when we select ϵ proportional to σ_k ; i.e., a relative precision confidence interval. In particular, with $\epsilon = \bar{\epsilon} \sigma_k$ the only technical modification required is replacing $\lim_{\epsilon \downarrow 0}$ with $\lim_{\bar{\epsilon} \downarrow 0}$.

Remark 8 A sufficient condition for verifying (7) in the history-dependent setting is

$$E \exp \left[\gamma \left(\frac{D_k - \mu_k}{\sigma_k / \sqrt{n_k}} \right) \mid \mathcal{H}_{k-1} \right] < M \text{ for } |\gamma| \leq \gamma_0, \quad (12)$$

where M and γ_0 do not depend on the iteration k , the history, $\mathcal{H}_k$, or the confidence interval width, ϵ (which determines n_k). This test may often be more natural to apply than attempting to verify (7) directly.

2.4 Stopping Rules: Additional Considerations

In this subsection we address two additional issues. First, we describe how sample variance estimates may be incorporated in confidence interval construction, and second, we examine the issue of finite stopping times.

2.4.1 Sample Variance Estimators

If the confidence interval width, ϵ , is a sufficiently small, pre-selected, fixed value then we interpret Theorems 3 and 6 as providing an approximate absolute precision confidence interval. However, the population variance terms, σ_k^2 , are typically unknown; hence n_k is unknown and the procedure is not implementable. Alternatively, if ϵ is proportional to σ_k , we obtain a relative precision confidence interval (see Remark 7) with *unknown* width $\epsilon\sigma_T$. There are standard approaches to this difficulty based on well-known results from parameter estimation in statistics. For simplicity, we describe one possible approach in the setting of Theorem 3 for the relative precision case. We replace the CLT hypothesis (2) with the following assumption:

$$\frac{D_k - \mu_k}{s_k / \sqrt{n_k}} \Rightarrow N(0, 1) \text{ as } n_k \rightarrow \infty, \quad (13)$$

where s_k^2 is a sample variance estimator. We may then form a "sample-variance" analog of Theorem 3 by replacing σ_k^2 with s_k^2 in (7) and (8). However, this result is of limited value because the sample variance equivalent of the moment generating function hypothesis (7) may be difficult to verify; if D_k is a sample mean of i.i.d. normal random variables then $s_k^{-1} \sqrt{n_k} (D_k - \mu_k)$ is a Student's T random variable and does not have a moment generating function. (The situation for an absolute precision confidence interval is at best unimproved.)

A simple solution to this difficulty is as follows: At iteration T we can resample the difference random variable at the proposed solution x_T (see Remark 4). We denote this random variable $\tilde{D}_T$ and assume it satisfies the CLT (2) with mean μ_T and variance σ_T^2 . Coupled with a weakly consistent sample variance estimator, $\tilde{s}_T^2$, this ensures the sample variance version of the CLT hypothesis (13) holds for $\tilde{D}_T$. From this we infer

$$\left[z^*, z^* + \tilde{D}_T + \frac{w_s \tilde{s}_T}{\sqrt{n}} \right] \quad (14)$$

is an approximate $100(1-\delta)\%$ confidence interval for $z(x_T)$ where n denotes the (resampling) sample size and w_δ satisfies $P\{N(0,1) \leq w_\delta\} = 1 - \delta$. As Theorems 3 and 6 ensure μ_T is not too large, we have *a priori* control on the confidence interval width. Restated: For sufficiently small ϵ and for sufficiently large $n = n(\eta)$ we have

$$P\left\{\tilde{D}_T^+ + \frac{w_\delta \tilde{\sigma}_T}{\sqrt{n}} \geq \epsilon \sigma_T + \eta\right\} \approx \alpha \quad \text{for any } \eta > 0. \quad (15)$$

2.4.2 Finite Stopping Times

It would be undesirable if the stopping rules we have developed precluded finite termination for “well-behaved” algorithms. To this end, we introduce the notion of a *stopping tolerance* ϵ' satisfying $0 < \epsilon' < \epsilon$ and make the following assumptions. The stopping time is redefined as

$$T = \inf_{k \geq 1} \{k : D_k \leq \epsilon'\}. \quad (16)$$

Under this termination criterion $P\{T > m\} = P\{D_1 > \epsilon', D_2 > \epsilon', \dots, D_m > \epsilon'\}$, and we assume

$$P\{D_1 > \epsilon', D_2 > \epsilon', \dots, D_m > \epsilon'\} = \int \prod_{k=1}^m P\{D_k > \epsilon' | \mathcal{H}_{k-1}\} dP_{\mathcal{H}_{m-1}} \quad (17)$$

which is a natural generalization of inter-iteration independence of the difference random variables to the history-dependent setting. With regard to the convergence structure of the algorithm we assume there exists a subsequence of $\{\mu_k\}_{k=1}^\infty$ that converges to zero with probability one; i.e.,

$$P\{\omega : \mathcal{H}_\infty(\omega) \text{ implies } \exists \{\mu_{k_j}(\omega)\}_{j=1}^\infty \text{ such that } \mu_{k_j}(\omega) \rightarrow 0\} = 1. \quad (18)$$

In this framework the following modification of Theorem 6 incorporates a finite stopping result.

Theorem 7 Assume (7), (11), (16), $\epsilon' \geq 0$ and define

$$n_k \geq \left(\frac{\sigma_k}{\epsilon - \epsilon'}\right)^2 (\beta' + 2p \ln^2 k) \quad (19)$$

and the corresponding parameters as in Theorem 9.

Then

$$\lim_{\epsilon \downarrow \epsilon'} P\{\mu_{T(\epsilon)} \leq \epsilon\} \geq 1 - \alpha. \quad (\text{asymptotic validity})$$

If, in addition, (12), (17), (18) hold and $0 < \epsilon' < \epsilon$ then

$$P\{T(\epsilon) < \infty\} = 1. \quad (\text{finite stopping time})$$

Proof

By hypothesis (17)

$$P\{T < \infty\} \geq 1 - \int \prod_{k=1}^m P\{D_k > \epsilon' | \mathcal{H}_{k-1}\} dP_{\mathcal{H}_{m-1}}.$$

We know from the bounded convergence theorem it suffices to show

$$\lim_{m \rightarrow \infty} \prod_{k=1}^m P\{D_k > \epsilon' | \mathcal{H}_{k-1}\} = 0 \text{ w.p.1.}$$

For any $\epsilon' > 0$, condition (12) and $n_k/\sigma_k^2 \rightarrow \infty$ as $k \rightarrow \infty$ ensures there exists K such that $k \geq K$ implies $E\{|D_k - \mu_k| | \mathcal{H}_{k-1}\} < \epsilon'/4$. Applying (a conditional version of) Markov's inequality we find

$$P\{|D_k - \mu_k| \geq \epsilon'/2 | \mathcal{H}_{k-1}\} \leq E\{|D_k - \mu_k| | \mathcal{H}_{k-1}\} / (\epsilon'/2).$$

Define $\mathcal{K}_m = \{k : k \leq m, k \geq K, \mu_k \leq \epsilon'/2\}$, and observe

$$\prod_{k=1}^m P\{D_k > \epsilon' | \mathcal{H}_{k-1}\} \leq \prod_{k \in \mathcal{K}_m} P\{|D_k - \mu_k| \geq \epsilon'/2 | \mathcal{H}_{k-1}\} \leq (1/2)^{|\mathcal{K}_m|}.$$

The finite stopping result follows as (18) implies $|\mathcal{K}_m| \rightarrow \infty$ as $m \rightarrow \infty$ w.p.1.

The proof of the asymptotic validity result is virtually identical to that of Theorem 6. ■

3 Application

A T -stage stochastic linear program with recourse (SLP- T) may be formulated as follows:

$$\begin{aligned} & \text{minimize} && \sum_{t=1}^T \sum_{\omega_t \in \Omega_t} p_t^{\omega_t} c_t^{\omega_t} x_t^{\omega_t} \\ & \text{subject to} && -B_{t-1}^{\omega_t} x_{t-1}^{a(\omega_t)} + A_t^{\omega_t} x_t^{\omega_t} = b_t^{\omega_t}, \quad x_t^{\omega_t} \geq 0, \quad \omega_t \in \Omega_t \\ & \text{for} && t = 1, \dots, T \\ & \text{where} && B_0^{\omega_1} \equiv 0. \end{aligned}$$

SLP- T

A sample point (scenario) in the stage t sample space, Ω_t , is denoted ω_t . A stage $t \geq 2$ scenario, ω_t , has a unique stage $t-1$ ancestor denoted $a(\omega_t)$, and a stage $t < T$ scenario has a set of stage $t+1$ descendant scenarios denoted $\Delta(\omega_t)$. A stage t realization, $\xi_t^{\omega_t} = \text{vec}(c_t^{\omega_t}, b_t^{\omega_t}, B_{t-1}^{\omega_t}, A_t^{\omega_t})$, is to be read column-wise as a vector in $\mathbb{R}^{N_t}$, where $N_t = n_t + m_t + m_t \cdot n_{t-1} + m_t \cdot n_t$; $A_t^{\omega_t}$ is an $m_t \times n_t$ matrix and the remaining matrices and vectors are dimensioned to conform. We assume

ξ_t has finite support and a probability mass function given by $P\{\xi_t = \xi_t^{\omega_t}\} = p_t^{\omega_t}$. For notational convenience, we assume a first stage sample space, Ω_1 that is a singleton set where $\xi_1^{\omega_1}$ represents the known state at the time decisions are made in the first stage; clearly, $p_1^{\omega_1}$ has value one. At the time decisions are made in stage t , the observation $\xi_t^{\omega_t}$ and the previous stage's decision $x_{t-1}^{a(\omega_t)}$ are known to the decision maker; the goal is to find a first stage decision, x_1 , that minimizes the expected cost of operating the system over T stages.

Pereira and Pinto [27] have proposed a sampling and decomposition-based algorithm for SLP- T models with interstage independence of the stochastic parameters and a “manageable” number of descendants, $|\Delta(\omega_t)|$, for each node in the scenario tree. Note that if T is moderately large, application of exact methods or bounding and approximation schemes (see §1) may be computationally impractical. The basic idea behind the algorithm is to compute upper bound estimates by sampling paths through the scenario tree on a forward pass, and to compute deterministic outer-linearization cuts on a backward pass. Figures 1 and 2 are designed to illustrate this concept.




Figure 1: Forward pass (left to right)




Figure 2: Backward pass (right to left)

A sample path, $\omega = (\omega_1, \omega_2, \dots, \omega_T)$, specifies exactly one node per stage in the scenario tree and has the property that the nodes identify a path from the stage 1 node to a stage T leaf node. Nonanticipativity constraints are satisfied on a forward pass along a sample path: The first stage problem is solved with only previously generated cut information regarding the future. The first stage decision is then passed to the right-hand-side of a randomly selected second stage subproblem, $\text{sub}(\omega_2)$, and the process is then repeated from stage 2 to 3. A forward pass along a sample path

simulates operation of the system given the cuts collected in previous iterations. A new cut is then appended to the set of cuts at each stage on a backward pass along each sample path. Under the assumption that the stochastic parameters are interstage independent, the cuts computed for any stage t subproblem are valid for all other subproblems on that stage. The algorithm for SLP- $\mathcal{T}$ is summarized in Figure 3.

```

step 0  let  $k = 0$ ;
         append lower bounding cuts  $\theta_t \geq -M_t, t = 1, \dots, T - 1$ ;

step 1  solve the stage 1 master program and obtain  $(x_1^k, \theta_1^k)$ ;
         let  $\underline{z}^k = c_1 x_1^k + \theta_1^k$ ;

step 2  select a set of random sample paths  $\mathcal{S}^k$  according to  $p_T^{\omega_T}$ ;
         do  $\omega \in \mathcal{S}^k$ 
           do  $t = 2$  to  $T$ 
             form RHS of sub( $\omega_t$ ):  $B_{t-1}^{\omega_t} [x_{t-1}^{a(\omega_t)}]^k + b_t^{\omega_t}$ ;
             solve and obtain  $[x_t^{\omega_t}]^k$ ;
           enddo
         enddo
         let  $\bar{z}^k = c_1 x_1^k + \frac{1}{|\mathcal{S}^k|} \sum_{\omega \in \mathcal{S}^k} \sum_{t=2}^T c_t^{\omega_t} [x_t^{\omega_t}]^k$ ;

step 3  if  $\bar{z}^k - \underline{z}^k \leq 0$  then stop:  $x_1^k$  is the proposed solution;

step 4  do  $t = T - 1$  downto 1
         do  $\omega \in \mathcal{S}^k$ 
           do  $\omega_{t+1} \in \Delta(\omega_t)$ 
             form RHS of sub( $\omega_{t+1}$ ):  $B_t^{\omega_{t+1}} [x_t^{a(\omega_{t+1})}]^k + b_{t+1}^{\omega_{t+1}}$ ;
             solve and obtain dual variables;
           enddo
           use dual variables to compute cut;
           append the set of stage  $t$  cuts with  $\theta_t - G_t x_t \geq g_t$ ;
         enddo
       enddo

step 5  let  $k = k + 1$ ; goto step 1;

```

Figure 3: Decomposition Algorithm for SLP- $\mathcal{T}$

The algorithm of Figure 3 generates valid lower bounds (i.e., they contain no error due to sampling) and in the history-dependent sense, sample mean upper bounds on the optimal objective function value; applying Theorem 6 in this setting is straightforward. The first stage decision at the k^{th} iteration, x_1^k , is a random variable; it depends on the set of first stage cuts, and these cuts, in turn, depend on which sample paths were selected in previous iterations. However, given the history of the algorithm through the first $k - 1$ iterations, $\mathcal{H}_{k-1}$, the decision x_1^k is known, and the upper bound estimate $\bar{z}^k$ is the sum of a deterministic constant, $c_1 x_1^k$, and a sample mean of i.i.d.r.v.'s with realizations of the form $[z_2^\omega]^k = \sum_{i=2}^T c_i^{\omega_i} [x_i^{\omega_i}]^k$. Thus, the upper bound estimate satisfies a history-dependent CLT with mean $u_k = c_1 x_1^k + E z_2^k$ and variance $\sigma_{u_k}^2 = E (z_2^k - E z_2^k)^2$. The lower bound, $\underline{z}^k = l_k$, is deterministic when conditioned on $\mathcal{H}_{k-1}$. Thus the difference random variable $D_k = \bar{z}^k - \underline{z}^k$ satisfies the history-dependent CLT hypothesis (11) of Theorem 6 where $\mu_k = u_k - l_k$ and $\sigma_k^2 = \sigma_{u_k}^2$. The random vectors $\xi_t = \text{vec}(c_t, b_t, B_{t-1}, A_t)$, $t = 2, \dots, T$ are bounded and if we assume, for example, that the subproblems at each stage have bounded primal feasible regions then the random variable z_2^k is bounded. Thus the conditional moment generating function hypothesis (12) is satisfied (see Proposition 4), and hypothesis (7) then follows.

We now describe the recommended algorithm for SLP- $\mathcal{T}$ based on the stopping rule theory of §2. The modified algorithm utilizes the idea of stopping cycles (see Remark 3) which help to minimize computational burden. The existing steps of the algorithm in Figure 3 are modified as follows. To step 0 we append "define $\epsilon > 0$; let $\nu = 0$, $n_\nu > 0$, $p > 0$, $0 < \alpha < 1$, and let $\beta' = \max \{2 \ln [\phi(p)/(\sqrt{2\pi} \alpha)], 1\}$." See Section 2.2 for recommended values of p and the corresponding values of $\phi(p)$. In step 2 we select the random sample $\mathcal{S}^k$ to be of size n_ν . In step 3, if the heuristic pre-test is satisfied then a stopping cycle is complete and we go to step 6. The additional steps are detailed in Figure 4. (Note that we have redefined ϵ to be the $\hat{\epsilon}$ of Remark 7 so that we are employing a relative precision confidence interval.)

The sample-size formula (8) of Theorem 6 is satisfied by the modified algorithm, and hence the asymptotic validity result holds. Note $P\{\mu_T \leq \epsilon \sigma_T\} \geq 1 - \alpha$ implies $P\{l_T \leq z(x_T) \leq l_T + \epsilon \sigma_T\} \geq 1 - \alpha$; moreover, in this setting, l_T is a known lower bound when the algorithm terminates (see Remark 5). While σ_T is unknown, as a practical matter one may be satisfied that the solution

```

step 6  let  $\nu = \nu + 1$ ; and  $n_\nu = \lceil \epsilon^{-2} (\beta' + 2p \ln^2 \nu) \rceil$ 

step 7  independently select a set of new sample paths  $S^k$  according to  $p_T^{\omega_T}$  of size  $n_\nu$ ;
        do  $\omega \in S^k$ 
            do  $t = 2$  to  $T$ 
                form RHS of sub( $\omega_t$ ):  $B_{t-1}^{\omega_t} [x_{t-1}^{a(\omega_t)}]^k + b_t^{\omega_t}$ ;
                solve and obtain  $[x_t^{\omega_t}]^k$ ;
            enddo
        enddo
        let  $\bar{z}^k = c_1 x_1^k + \frac{1}{|S^k|} \sum_{\omega \in S^k} \sum_{t=2}^T c_t^{\omega_t} [x_t^{\omega_t}]^k$ ;

step 8  if  $\bar{z}^k - \underline{z}^k \leq 0$  then stop:  $x_1^k$  is the proposed solution; otherwise goto step 4

```

Figure 4: Additional “Stopping Rule” Steps to the Algorithm for SLP-T

x_T is of sufficiently high quality based on sample variance estimates observed during the course of the algorithm. Alternatively, the resampling procedure described in Section 2.4 may be applied to obtain a confidence interval of the form (14); we can also replace z^* by the known lower bound l_T in (14).

4 Empirical Coverage Results

In this section we present preliminary empirical coverage results that illustrate dangers associated with naive stopping rules and show empirical performance of the recommended stopping rules developed in §2. The simplistic “test problems” are motivated by Example 1, and we use a pseudo-random number generator [30] to directly form the difference random variables.

In the first group of these problems, the true gap is $\mu_k = 2/3$ for $k = 1, \dots, K$ and $\mu_k = 0$ for $k \geq K + 1$, and the pre-specified confidence interval width is $\epsilon = 1/3$. We terminate on the first iteration in which the difference random variable drops below zero; if the algorithm stops prior to iteration $K + 1$ it will have stopped incorrectly. The difference random variables are sample means of i.i.d.r.v.’s of the form $U(-\sqrt{3}, \sqrt{3}) + 2/3$, where $U(a, b)$ is a uniform random variable on the interval (a, b) . The choice of $(a, b) = (-\sqrt{3}, \sqrt{3})$ yields $\sigma_k = 1$. Table 2 depicts the empirical coverage results for two separate stopping rules on four test problems with different values of K . The first

stopping rule ignores the sequential nature of the problem and uses the fixed sample size that would generate an asymptotically valid 95% one-sided confidence interval of width $1/3$ for the true gap, if the algorithm consisted of only one iteration. We choose $\alpha = 0.05$ which yields $z_\alpha = 1.645$ and thus $n = \lceil (z_\alpha \sigma_k / \epsilon)^2 \rceil = 25$. In the second stopping rule, we increase the sample size at each iteration according to sample-size formula (8). The parameter p is selected from Table 1 for each value of K . Table 2 is based on 1000 replications and shows that ignoring the sequential nature of the problem does not lead to undercoverage results in this setting until K grows large and formula (8) leads to 100% coverage in each case.

K	$n = 25$	$n_k = \mathcal{O}(\log^2 k)$
	coverage	
10	0.997	1.000
100	0.963	1.000
1000	0.692	1.000
10000	0.029	1.000

Table 2: Coverage Results: $\epsilon = 1/3$ and $\mu = 2/3$

Next we consider the case in which $\mu_k = \epsilon = 1/3$ for $k = 1, \dots, K$. Thus the only modification to the previous group of test problems is that the difference random variables are sample means of i.i.d.r.v.'s of the form $U(-\sqrt{3}, \sqrt{3}) + 1/3$. This is a more demanding test of the stopping rule theory because the true gap and the confidence width are identical. The coverage results are summarized in Table 3. The coverage results obtained via sample-size formula (8) significantly exceed the 95% confidence level in each case. As one might expect, the naive stopping rule has very poor coverage results in this setting.

K	$n = 25$	$n_k = \mathcal{O}(\log^2 k)$
	coverage	
10	0.613	0.994
100	0.007	0.995
1000	0.000	0.993
10000	0.000	0.989

Table 3: Coverage Results: $\epsilon = 1/3$ and $\mu = 1/3$

5 Summary

In this paper we developed a stopping rule theory for a class of optimization algorithms that estimate upper and lower bounds on the optimal objective function value via sampling. While our immediate motivation lies in developing stopping rules for a class of Monte Carlo sampling-based stochastic programming algorithms, the theory may also be applicable to other optimization algorithms that use simulation techniques. In the main result, we assume that the difference random variables satisfy history-dependent central limit theorems and provide appropriate conditions and a sample-size formula under which the desired confidence interval for the objective function value of the proposed solution is asymptotically valid. We regard the recommended procedure as conservative because: (i) underlying convergence properties of the optimization algorithm are ignored in developing the methodology; and (ii) the normal tail bound used to derive the sample-size formula is not sharp, particularly in the early iterations. Moreover, through well-designed stopping cycles and because of the slow growth of the sample size formula, the recommended stopping rules are practical from an implementation standpoint.

The applicability of the stopping rule theory was illustrated on an algorithm for a class of multistage stochastic linear programs (Pereira and Pinto [27]) that generate sample mean upper bound estimates and deterministically valid lower bounds. In other sampling-based stochastic programming algorithms for two stage programs (Dantzig and Glynn [6], Infanger [19], Higle and Sen [14]), sample mean upper bounds are readily available, but lower bound estimates have proved more difficult to analyze. Development and analysis of lower bound estimators for these sampling-based algorithms remains an active area of research [15,18,25]; the procedures we have developed here should be useful for asymptotically normal estimators.

References

- [1] Beale, E.M.L. (1955): On Minimizing a Convex Function Subject to Linear Inequalities, *Journal of the Royal Statistical Society* 17B, 173-184.
- [2] Birge, J.R. (1985): Decomposition and Partitioning Methods for Multistage Stochastic Linear Programs, *Operations Research* 33, 989-1007.
- [3] Birge, J.R. and R. J.-B. Wets (1989): Sublinear Upper Bounds for Stochastic Programs with Recourse, *Mathematical Programming* 43, 131-149.
- [4] Birge, J.R. and R. J.-B. Wets (1986): Designing Approximation Schemes for Stochastic Optimization Problems, in particular, for Stochastic Programs with Recourse, *Mathematical Programming Study* 27, 54-102.
- [5] Dantzig, G.B. (1955): Linear Programming Under Uncertainty, *Management Science* 1, 197-206.
- [6] Dantzig, G.B., P.W. Glynn (1990): Parallel Processors for Planning Under Uncertainty, *Annals of Operations Research* 22, 1-21.
- [7] Dantzig, G.B., P.W. Glynn, M. Avriel, J.C. Stone, R. Entriken, and M. Nakayama (1989): Decomposition Techniques for Multi-Area Generation and Transmission Planning Under Uncertainty, EPRI Report 2940-1.
- [8] Dwight, H.B. (1958): *Mathematical Tables of Elementary and Some Higher Mathematical Functions*, 2nd edition Dover Publications Inc., New York, New York.
- [9] Ermoliev, Y. (1983): Stochastic Quasigradient Methods and their Applications to Systems Optimization, *Stochastics* 9, 1-36.
- [10] Feller, W. (1957): *An Introduction to Probability Theory and Its Applications, Volume I*, John Wiley & Sons Inc., third edition.
- [11] Frauendorfer, K. (1988): Solving SLP Recourse Problems with Arbitrary Multivariate Distributions - The Dependent Case, *Mathematics of Operations Research* 13, 377-394.
- [12] Gassmann, H.I. (1989): Optimal Harvest of a Forest in the Presence of Uncertainty, *Canadian Journal of Forest Research* 19, 1267-1274.
- [13] Gaivoronski, A. (1988): Stochastic Quasigradient Methods and their Implementation, in Y. Ermoliev and R.J-B Wets (eds.) *Numerical Techniques for Stochastic Optimization*, Springer Verlag.
- [14] Higle, J.L. and S. Sen (1991): Stochastic Decomposition: An Algorithm for Two-Stage Linear Programs with Recourse, *Mathematics of Operations Research* 16, 650-669.
- [15] Higle, J.L. and S. Sen (1993): Duality and Statistical Tests of Optimality for Two Stage Stochastic Programs, Technical Report Systems and Industrial Engineering Department, University of Arizona.
- [16] Hogg, R.V. and A.T. Craig (1970): *Introduction to Mathematical Statistics*, Macmillan, third edition.

- [17] Kall, P. (1979): Computational Methods for Solving Two-Stage Stochastic Linear Programming Problems, *Journal of Applied Mathematics and Physics* 30, 261-271.
- [18] Infanger, G. (1992): Planning Under Uncertainty: Solving Large-Scale Stochastic Linear Programs, Technical Report SOL 92-8, Systems Optimization Laboratory, Department of Operations Research, Stanford University.
- [19] Infanger, G. (1992): Monte Carlo (Importance) Sampling within a Benders Decomposition Algorithm for Stochastic Linear Programs, *Annals of Operations Research* 39, 69-95.
- [20] Jacobs, J., G. Freeman, J. Grygier, D. Morton, G. Schultz, K. Staschus, J. Stedinger, and B. Zhang (1993): Stochastic Optimal Coordination of River-basin and Thermal Electric Systems (SOCRATES): A System for Scheduling Hydroelectric Generation Under Uncertainty, submitted to *Annals of Operations Research*.
- [21] Kall, P., A. Ruszczyński, and K. Frauendorfer (1988): Approximation Techniques in Stochastic Programming, in Y. Ermoliev and R.J-B Wets (eds.) *Numerical Techniques for Stochastic Optimization*, Springer Verlag.
- [22] Kusu, M.I., W.T. Ziemba (1986): A Bank Asset and Liability Management Model, *Operations Research*, 34, 356-376.
- [23] Loève, M. (1963): *Probability Theory*, D. Van Nostrand Company, Inc., Princeton, New Jersey, third edition.
- [24] Lustig, I.J., J.M. Mulvey, and T.J. Carpenter (1991): Formulating Two-Stage Stochastic Programs for Interior Point Methods, *Operations Research* 39, 757-770.
- [25] Morton, D.P. (1993): Algorithmic Advances in Stochastic Programming, Technical Report SOL 93-6, Systems Optimization Laboratory, Department of Operations Research, Stanford University.
- [26] Mulvey, J.M., H. Vladimirou (1989): Stochastic Network Optimization Models for Investment Planning, *Annals of Operations Research: Network Optimization and Applications* 20, 187-217.
- [27] Pereira, M.V.F. and L.M.V.G. Pinto (1991): Multi-stage Stochastic Optimization Applied to Energy Planning, *Mathematical Programming* 52, 359-375.
- [28] Pflug, G. C. (1988): Stepsize Rules, Stopping Times, and their Implementations in Stochastic Quasigradient Algorithms, in Y. Ermoliev and R.J-B Wets (eds.) *Numerical Techniques for Stochastic Optimization*, Springer Verlag.
- [29] Rockafellar, R.T. and R.J.-B. Wets (1991): Scenarios and Policy Aggregation in Optimization under Uncertainty, *Mathematics of Operations Research* 16, 119-147.
- [30] Schrage, L. (1979): A More Portable Fortran Random Number Generator, *ACM Trans. Math. Software* 5, 132-138.
- [31] Van Slyke, R.M., R.J-B Wets (1969): L-Shaped Linear Programs with Applications to Optimal Control and Stochastic Programming, *SIAM Journal of Applied Mathematics* 17, 638-663.
- [32] Wolfram, S. (1988): *Mathematica: A System for Doing Mathematics by Computer*, Addison-Wesley Publishing Company, Redwood City, California.

INITIAL DISTRIBUTION LIST

1. Research Office (Code 08) 1
Naval Postgraduate School
Monterey, CA 93943-5000
2. Dudley Knox Library (Code 52)..... 2
Naval Postgraduate School
Monterey, CA 93943-5002
3. Defense Technical Information Center 2
Cameron Station
Alexandria, VA 22314
4. Department of Operations Research (Code OR) 1
Naval Postgraduate School
Monterey, CA 93943-5000
5. David P. Morton (Code OR/Mo)..... 40
Naval Postgraduate School
Monterey, CA 93943-5000