Some recent developments in stochastic programming

A. Shapiro

Georgia Institute of Technology,
Atlanta, Georgia 30332-0205, USA

Efficient Monte Carlo: From Variance Reduction to
Combinatorial Optimization, Denmark, 2008

Consider stochastic optimization problem:

$$\operatorname*{Min}_{x\in X} \Big\{ f(x) := \mathbb{E}_P[F(x,\boldsymbol{\xi})] \Big\}, \tag{1}$$

where $\boldsymbol{\xi}$ is a random vector having probability distribution P supported on set $\Xi \subset \mathbb{R}^d$, $F(x,\xi)$ is a real valued function and $X \subset \mathbb{R}^n$. For example, consider two-stage (linear) stochastic programming problem with recourse

$$\text{Min}_{x\in X}\, c^{\mathsf T}x + \mathbb{E}[Q(x,\boldsymbol{\xi})], \tag{2}$$

where $X = \{x : Ax = b,\ x \geq 0\}$ and $Q(x,\xi)$ is the optimal value of the second stage problem

$$\text{Min}_{y}\, q^{\mathsf T}y \ \text{ s.t. } \ Tx + Wy = h,\ y \geq 0, \tag{3}$$

with $\xi = (q, T, W, h)$. The feasible set X can be finite, i.e., integer first stage problem. Both stages can be integer (mixed integer) problems.

Suppose that the probability distribution P of $\boldsymbol{\xi}$ has a finite support, i.e., $\boldsymbol{\xi}$ can take values $\xi_1, ..., \xi_K$ (called *scenarios*) with respective probabilities $p_1, ..., p_K$. Then

$$\mathbb{E}_P[F(x, \boldsymbol{\xi})] = \sum_{k=1}^{K} p_k F(x, \xi_k).$$

In the case of two-stage (linear) stochastic programming problem with recourse we can write problem (2)-(3) as one large linear program:

$$\begin{array}{ll} \text{Min}_{x,y_1,...,y_K} & c^{\mathsf{T}}x + \sum_{k=1}^{K} p_k q_k^{\mathsf{T}} y_k \\ \text{subject to} & Ax = b, \\ & T_k x + W_k y_k = h_k,\ k = 1, ..., K, \\ & x \geq 0,\ y_k \geq 0,\ k = 1, ..., K. \end{array} \tag{4}$$

Even crude discretization of the distribution of $\boldsymbol{\xi}$ leads to an exponential growth of the number of scenarios.

Could such problems be solved numerically?

How do we know the probability distribution P?

Why do we optimize the expected value of the objective (cost) function?

Monte Carlo sampling approach

Let $\xi^1, ..., \xi^N$ be a generated (iid) random sample drawn from P and

$$\hat{f}_N(x) := N^{-1} \sum_{j=1}^{N} F(x, \xi^j)$$

be the corresponding sample average function. By the Law of Large Numbers, for a given $x \in X$, we have $\hat{f}_N(x) \to f(x) = \mathbb{E}_P[F(x, \boldsymbol{\xi})]$ w.p.1 as $N \to \infty$.

Notoriously slow convergence of order $O_p(N^{-1/2})$. By the Central Limit Theorem

$$N^{1/2}\left[\hat{f}_N(x) - f(x)\right] \Rightarrow N(0, \sigma^2(x)),$$

where $\sigma^2(x) := \mathsf{Var}[F(x, \boldsymbol{\xi})]$.

The sample average approximation (SAA) approach to Monte Carlo sampling optimization, the true problem is approximated by the sample average approximation problem:

$$\mathsf{Min}_{x\in X}\left\{\hat{f}_N(x) := N^{-1}\textstyle\sum_{j=1}^{N} F(x, \xi^j)\right\}.$$

Once the sample $\xi^1, ..., \xi^N \sim P$ is generated, the SAA problem becomes a deterministic optimization problem and can be considered as a stochastic programming problem with scenarios $\xi^1, ..., \xi^N \sim P$ each with probability $1/N$.

Notation

v^0 is the optimal value of the true problem
S^0 is the optimal solutions set of the true problem
S^ε is the set of ε-optimal solutions of the true problem
$\hat{v}_N$ is the optimal value of the SAA problem
$\hat{S}_N^\varepsilon$ is the set of ε-optimal solutions of the SAA problem
$\hat{x}_N$ is an optimal solution of the SAA problem

Convergence properties

Vast literature on statistical properties of the SAA estimators $\hat{v}_N$ and $\hat{x}_N$:

Consistency. By the Law of Large Numbers, $\hat{f}_N(x)$ converge (pointwise) to $f(x)$ w.p.1. Under mild additional conditions, this implies that $\hat{v}_N \to v^0$ and $\text{dist}(\hat{x}_N, S^0) \to 0$ w.p.1 as $N \to \infty$. In particular, $\hat{x}_N \to x^0$ w.p.1 if $S^0 = \{x^0\}$. (Consistency of Maximum Likelihood estimators, Wald (1949)).

Central Limit Theorem type results.

$$\hat{v}_N = \min_{x \in S^0} \hat{f}_N(x) + o_p(N^{-1/2}).$$

In particular, if $S^0 = \{x^0\}$, then

$$N^{1/2}[\hat{v}_N - v^0] \Rightarrow N(0, \sigma^2(x^0))$$

(Shapiro, 1991).

These results suggest that the optimal value of the SAA problem converges at a rate of $\sqrt{N}$. In particular, if $S^0 = \{x^0\}$, then $\hat{v}_N$ converges to v^0 at the same rate as $\hat{f}_N(x^0)$ converges to $f(x^0)$.

Sample size estimates (by Large Deviations type bounds)

Suppose that $|X| < \infty$, i.e., the set X is finite, and: (i) for every $x \in X$ the expected value $f(x) = \mathbb{E}[F(x, \boldsymbol{\xi})]$ is finite, (ii) there are constants $\sigma > 0$ and $a \in (0, +\infty]$ such that

$$M_x(t) \leq \exp\{\sigma^2 t^2/2\}, \ \ \forall t \in [-a, a], \ \forall x \in X \setminus S^{\varepsilon},$$

where $M_x(t)$ is the moment generating function of the random variable $F(u(x), \boldsymbol{\xi}) - F(x, \boldsymbol{\xi}) - \mathbb{E}[F(u(x), \boldsymbol{\xi}) - F(x, \boldsymbol{\xi})]$ and $u(x)$ is a point of the optimal set S^0. Choose $\varepsilon > 0$, $\delta \geq 0$ and $\alpha \in (0, 1)$ such that $0 < \varepsilon - \delta \leq a\sigma^2$. Then for sample size

$$N \geq \frac{2\sigma^2}{(\varepsilon - \delta)^2} \log \left(\frac{|X|}{\alpha} \right)$$

we are guaranteed, with probability at least $1 - \alpha$, that any δ-optimal solution of the SAA problem is an ε-optimal solution of the true problem, i.e., $\text{Prob}(\widehat{S}^{\delta}_N \subset S^{\varepsilon}) \geq 1 - \alpha$
(Kleywegt, Shapiro & Homem-de-Mello, 2001).

Let $X = \{x_1, x_2\}$ with $f(x_2) - f(x_1) > \varepsilon > 0$ and suppose that random variable $F(x_2, \boldsymbol{\xi}) - F(x_1, \boldsymbol{\xi})$ has normal distribution with mean $\mu = f(x_2) - f(x_1)$ and variance σ^2. By solving the corresponding SAA problem we make the correct decision (that x_1 is the minimizer) if $\hat{f}_N(x_2) - \hat{f}_N(x_1) > 0$. Probability of this event is $\Phi(\mu\sqrt{N}/\sigma)$. Therefore we need the sample size $N > z_\alpha^2 \sigma^2 / \varepsilon^2$ in order for our decision to be correct with probability at least $1 - \alpha$.

In order to solve the corresponding optimization problem we need to test $H_0 : \mu \leq 0$ versus $H_a : \mu > 0$. Assuming that σ^2 is known, by Neyman-Pearson Lemma, the uniformly most powerful test is: "reject H_0 if $\hat{f}_N(x_2) - \hat{f}_N(x_1)$ is bigger than a specified critical value".

Now let $X \subset \mathbb{R}^n$ be a set of finite diameter $D := \sup_{x',x\in X} \|x'-x\|$. Suppose that: (i) for every $x \in X$ the expected value $f(x) = \mathbb{E}[F(x,\boldsymbol{\xi})]$ is finite, (ii) there is a constant $\sigma > 0$ such that

$$M_{x',x}(t) \leq \exp\{\sigma^2 t^2/2\}, \;\; \forall t \in \mathbb{R}, \; \forall x', x \in X,$$

where $M_{x',x}(t)$ is the moment generating function of the random variable $F(x',\boldsymbol{\xi}) - F(x,\boldsymbol{\xi}) - \mathbb{E}[F(x',\boldsymbol{\xi}) - F(x,\boldsymbol{\xi})]$, (iii) there exists $\kappa : \Xi \to \mathbb{R}_+$ such that its moment generating function is finite valued in a neighborhood of zero and

$$\left| F(x',\xi) - F(x,\xi) \right| \leq \kappa(\xi)\|x'-x\|, \;\; \forall \xi \in \Xi, \; \forall x', x \in X.$$

Choose $\varepsilon > 0$, $\delta \in [0,\varepsilon)$ and $\alpha \in (0,1)$. Then for sample size

$$N \geq \frac{8\sigma^2}{(\varepsilon-\delta)^2}\left[n \log\left(\frac{O(1)DL}{(\varepsilon-\delta)^2}\right) + \log\left(\frac{2}{\alpha}\right)\right] \bigvee \left[\beta^{-1}\log\left(\frac{2}{\alpha}\right)\right],$$

we are guaranteed that $\text{Prob}\left(\widehat{S}_N^\delta \subset S^\varepsilon\right) \geq 1-\alpha$.

In particular, if $\kappa(\xi) \equiv L$, then the estimate takes the form

$$N \geq O(1) \left(\frac{LD}{\varepsilon - \delta}\right)^2 \left[n \log \left(\frac{O(1)DL}{\varepsilon - \delta} \right) + \log \left(\frac{1}{\alpha} \right) \right].$$

Suppose further that for some $c > 0$, $\gamma \geq 1$ and $\bar{\varepsilon} > \varepsilon$ the following growth condition holds

$$f(x) \geq v^0 + c[\mathrm{dist}(x, S^0)]^\gamma, \quad \forall x \in S^{\bar{\varepsilon}},$$

and that the problem is convex. Then, for $\delta \in [0, \varepsilon/2]$, we have the following estimate of the required sample size:

$$N \geq \left(\frac{O(1)LD}{c^{1/\gamma} \varepsilon^{(\gamma-1)\gamma}}\right)^2 \left[n \log \left(\frac{O(1)\bar{D}L}{\varepsilon} \right) + \log \left(\frac{1}{\alpha} \right) \right],$$

where $\bar{D}$ is the diameter of $S^{\bar{\varepsilon}}$. In particular, if $S^0 = \{x^0\}$ is a singleton and $\gamma = 1$, we have the estimate (independent of ε):

$$N \geq O(1) c^{-2} L^2 \left[n \log(O(1) c^{-1} L) + \log(\alpha^{-1}) \right].$$

Example Let $F(x,\xi) := \|x\|^{2k} - 2k\left(\xi^{\mathsf{T}} x\right)$, where $k \in \mathbb{N}$ and

$$X := \{x \in \mathbb{R}^n : \|x\| \le 1\}.$$

Suppose, that $\xi \sim N(0, \sigma^2 I_n)$. Then $f(x) = \|x\|^{2k}$, and for $\varepsilon \in [0,1]$, the set of ε-optimal solutions of the true problem is

$$\{x : \|x\|^{2k} \le \varepsilon\}.$$

Let $\bar{\xi}_N := (\xi^1 + ... + \xi^N)/N$. The corresponding sample average function is

$$\hat{f}_N(x) = \|x\|^{2k} - 2k\left(\bar{\xi}_N^{\mathsf{T}} x\right),$$

and $\hat{x}_N = \|\bar{\xi}_N\|^{-\gamma}\bar{\xi}_N$, where $\gamma := \frac{2k-2}{2k-1}$ if $\|\bar{\xi}_N\| \le 1$, and $\gamma = 1$ if $\|\bar{\xi}_N\| > 1$. Therefore, for $\varepsilon \in (0,1)$, the optimal solution of the SAA problem is an ε-optimal solution of the true problem iff $\|\bar{\xi}_N\|^{\nu} \le \varepsilon$, where $\nu := \frac{2k}{2k-1}$.

We have that $\bar{\xi}_N \sim N(0, \sigma^2 N^{-1} I_n)$, and hence $N\|\bar{\xi}_N\|^2/\sigma^2$ has the chi-square distribution with n degrees of freedom. Consequently, the probability that $\|\bar{\xi}_N\|^\nu > \varepsilon$ is equal to the probability

$$\mathbb{P}\left(\chi_n^2 > N\varepsilon^{2/\nu}/\sigma^2\right).$$

Moreover, $\mathbb{E}[\chi_n^2] = n$ and the probability $\mathbb{P}(\chi_n^2 > n)$ increases and tends to $1/2$ as n increases. Consequently, for $\alpha \in (0, 0.3)$ and $\varepsilon \in (0, 1)$, for example, the sample size N should satisfy

$$N > \frac{n\sigma^2}{\varepsilon^{2/\nu}} \tag{5}$$

in order to have the property: "with probability $1 - \alpha$ an (exact) optimal solution of the SAA problem is an ε-optimal solution of the true problem". Note that $\nu \to 1$ as $k \to \infty$.

Stochastic Approximation (SA) approach

Suppose that the problem is convex, i.e., the feasible set X is convex and $F(\cdot, \xi)$ is convex for all $\xi \in \Xi$. Classical SA algorithm

$$x_{j+1} = \Pi_X(x_j - \gamma_j G(x_j, \xi^j)),$$

where $G(x, \xi) \in \partial_x F(x, \xi)$ is a calculated gradient, Π_X is the orthogonal (Euclidean) projection onto X and $\gamma_j = \theta / j$. Theoretical bound (assuming $f(\cdot)$ is strongly convex and differentiable)

$$\mathbb{E}[f(x_j) - v^0] = O(j^{-1}),$$

for an optimal choice of constant θ (recall that v^0 is the optimal value of the true problem). This algorithm is very sensitive to choice of θ, does not work well in practice.

Robust SA approach (B. Polyak, 1990). Constant step size variant: fixed in advance sample size (number of iterations) N and step size $\gamma_j \equiv \gamma$, $j = 1, ..., N$: $\tilde{x}_N = \frac{1}{N}\sum_{j=1}^{N} x_j$. Theoretical bound

$$\mathbb{E}[f(\tilde{x}_N) - v^0] \le \frac{D_X^2}{2\gamma N} + \frac{\gamma M^2}{2},$$

where $D_X = \max_{x \in X} \|x - x_1\|_2$ and $M^2 = \max_{x \in X} \mathbb{E}\|G(x,\xi)\|_2^2$. For optimal (up to factor θ) $\gamma = \frac{\theta D_X}{M\sqrt{N}}$ we have

$$\mathbb{E}\left[f(\tilde{x}_N) - v^0\right] \le \frac{D_X M}{2\theta\sqrt{N}} + \frac{\theta D_X M}{2\sqrt{N}} \le \frac{\kappa D_X M}{\sqrt{N}},$$

where $\kappa = \max\{\theta, \theta^{-1}\}$. By Markov inequality it follows that

$$\text{Prob}\left\{f(\tilde{x}_N) - v^0 > \varepsilon\right\} \le \frac{\kappa D_X M}{\varepsilon\sqrt{N}},$$

and hence to the sample size estimate $N \ge \frac{\kappa^2 D_X^2 M^2}{\varepsilon^2 \alpha^2}$.

Mirror Decent SA method (Nemirovski)
Let $\|\cdot\|$ be a norm on $\mathbb{R}^n$ and $\omega(x)$ be a continuously differentiable strongly convex on X with respect to $\|\cdot\|$, i.e., for $x, x' \in X$:

$$\omega(x') \geq \omega(x) + (x' - x)^{\mathsf{T}} \nabla \omega(x) + \tfrac{1}{2} c \|x' - x\|^2.$$

Prox mapping $P_x : \mathbb{R}^n \to X$:

$$P_x(y) = \arg\min_{z \in X} \left\{ \omega(z) + (y - \nabla \omega(x))^{\mathsf{T}} z \right\}.$$

For $\omega(x) = \frac{1}{2}\|x\|^2$ we have that $P_x(y) = \Pi_X(x - y)$. Set

$$x_{j+1} = P_{x_j}(\gamma_j G(x_j, \xi^j)).$$

For constant step size $\gamma_j = \gamma$, $j = 1, ..., N$, with optimal

$$\gamma = \frac{D_{\omega,X}}{M_*}\sqrt{\frac{2c}{N}},$$

where $M_* = \max_{x \in X} \mathbb{E}\|G(x,\xi)\|_*^2$, with dual norm $\|\cdot\|_*$, and

$$\tilde{x}_N = N^{-1} \sum_{j=1}^{N} x_j$$

we have

$$\mathbb{E}\left[f(\tilde{x}_N) - v^0\right] \le D_{\omega,X}\sqrt{\frac{2M_*^2}{cN}},$$

where

$$D_{\omega,X} = \left[\max_{z \in X} \omega(z) - \min_{x \in X} \omega(x)\right]^{1/2}.$$

Validation analysis

How one can evaluate quality of a given (feasible) solution $\hat{x} \in X$? Two basic approaches: (1) Evaluate the gap $f(\hat{x}) - v^0$. (2) Verify the KKT optimality conditions at $\hat{x}$.

Statistical test based on estimation of $f(\hat{x}) - v^0$ (Norkin, Pflug & Ruszczynski 98, Mak, Morton & Wood 99):

(i) Estimate $f(\hat{x})$ by the sample average $\hat{f}_{N'}(\hat{x})$, using sample of a large size N'.

(ii) Solve the SAA problem M times using M independent samples each of size N. Let $\hat{v}_N^{(1)}, ..., \hat{v}_N^{(M)}$ be the optimal values of the corresponding SAA problems. Estimate $\mathbb{E}[\hat{v}_N]$ by the average $M^{-1} \sum_{j=1}^M \hat{v}_N^{(j)}$. Note that

$$\mathbb{E}\left[\hat{f}_{N'}(\hat{x}) - M^{-1} \sum_{j=1}^M \hat{v}_N^{(j)}\right] = \left(f(\hat{x}) - v^0\right) + \left(v^0 - \mathbb{E}[\hat{v}_N]\right),$$

and that $v^0 - \mathbb{E}[\hat{v}_N] > 0$.

The bias $v^0 - \mathbb{E}[\hat{v}_N]$ is positive and (under mild regularity conditions)

$$\lim_{N\to\infty} N^{1/2}\left(v^0 - \mathbb{E}[\hat{v}_N]\right) = \mathbb{E}\left[\max_{x\in S^0} Y(x)\right],$$

where $(Y(x_1), ..., Y(x_k))$ has a multivariate normal distribution with zero mean vector and covariance matrix given by the covariance matrix of the random vector $(F(x_1, \boldsymbol{\xi}), ..., F(x_k, \boldsymbol{\xi}))$. For ill-conditioned problems this bias is of order $O(N^{-1/2})$ and can be large if the ε-optimal solution set S^ε is large for some small $\varepsilon \geq 0$.

Common random numbers variant: generate a sample (of size N) and calculate the gap $\hat{f}_N(\hat{x}) - \inf_{x\in X} \hat{f}_N(x)$. Repeat this procedure M times (with independent samples), and calculate the average of the above gaps. This procedure works well for well conditioned problems, does not improve the bias problem.

KKT statistical test Let

$$X := \{x \in \mathbb{R}^n : c_i(x) = 0,\ i \in I,\ c_i(x) \le 0,\ i \in J\}.$$

Suppose that the probability distribution is continuous. Then $F(\cdot, \boldsymbol{\xi})$ is differentiable at $\hat{x}$ w.p.1 and $\nabla f(\hat{x}) = \mathbb{E}_P[\nabla_x F(\hat{x}, \boldsymbol{\xi})]$. KKT-optimality conditions at an optimal solution $x^0 \in S^0$ can be written as follows: $-\nabla f(x^0) \in C(x^0)$, where

$$C(x) := \left\{ y = \textstyle\sum_{i \in I \cup J(x)} \lambda_i \nabla c_i(x),\ \lambda_i \ge 0,\ i \in J(x) \right\},$$

and $J(x) := \{i : c_i(x) = 0,\ i \in J\}$. The idea of the KKT test is to estimate the distance $\delta(\hat{x}) := \text{dist}\,(-\nabla f(\hat{x}), C(\hat{x}))$, by using the sample estimator $\hat{\delta}_N(\hat{x}) := \text{dist}\left(-\nabla \hat{f}_N(\hat{x}), C(\hat{x})\right)$. The covariance matrix of $\nabla \hat{f}_N(\hat{x})$ can be estimated (from the same sample), and hence a confidence region for $\nabla f(\hat{x})$ can be constructed. This allows a statistical validation of the KKT conditions (Shapiro & Homem-de-Mello, 98).

Bounds by Mirror Decent SA method. (Lan, Nemirovski & Shapiro, 2008). Iterates

$$x_{j+1} = P_{x_j}(\gamma_j G(x_j, \xi^j)).$$

Consider

$$f^N(x) := \sum_{j=1}^{N} \nu_j [f(x_j) + g(x_j)^{\mathsf{T}}(x - x_j)],$$

where $f(x) = \mathbb{E}[F(x, \boldsymbol{\xi})]$, $g(x) = \mathbb{E}[G(x, \boldsymbol{\xi})]$ and $\nu_j := \gamma_j / (\sum_{j=1}^N \gamma_j)$. Since $g(x) \in \partial f(x)$, it follows that

$$f_*^N := \min_{x \in X} f^N(x) \le v^0.$$

Also by convexity of f,

$$v^0 \le f(\tilde{x}_N) \le f^{*,N} := \sum_{j=1}^{N} \nu_j f(x_j).$$

Computational counterparts:

$$\underline{f}^N := \min_{x\in X} \sum_{j=1}^{N} \nu_j [F(x_j,\xi_j) + G(x_j,\xi_j)^{\mathsf{T}}(x - x_j)],$$

$$\overline{f}^N := \sum_{j=1}^{N} \nu_j F(x_j,\xi_j).$$

Theorem. Assume that there are positive a constants M_*^2, Ω^2 such that for all $x \in X$:

$$\mathbb{E}\left[(F(x,\boldsymbol{\xi}) - f(x))^2\right] \le \Omega^2 M_*^2,$$
$$\mathbb{E}\left[\|G(x,\boldsymbol{\xi})\|_*^2\right] \le M_*^2.$$

Then

$$\begin{aligned}
\mathbb{E}\left[f^{*,N} - \overline{f}^N\right] &= 0,\\
\mathbb{E}\left\{\left[f^{*,N} - \overline{f}^N\right]^2\right\} &\le \Omega^2 M_*^2 \textstyle\sum_{t=1}^N \nu_t^2,\\
\mathbb{E}\left[\underline{f}^N\right] &\le v^0,\\
\mathbb{E}\left[|\underline{f}^N - f_*^N|\right] &\le 12\Omega M_* \sqrt{\textstyle\sum_{t=1}^N \nu_t^2},\\
\mathbb{E}\left[f^{*,N} - f_*^N\right] &\le 10\Omega M_* \sqrt{\textstyle\sum_{t=1}^N \nu_t^2} + \frac{D_{\omega,x}^2 + (2\alpha)^{-1} M_*^2 \sum_{t=1}^N \gamma_t^2}{\sum_{t=1}^N \gamma_t}.
\end{aligned}$$

Complexity of multistage stochastic programming

Multistage models Let ξ_t be a random process. Denote $\xi_{[t]} := (\xi_1, .., \xi_t)$ the history of the process ξ_t up to time t. The values of the decision vector x_t, chosen at stage t, may depend on the information $\xi_{[t]}$ available up to time t, but not on the future observations. The decision process has the form

$$\text{decision}(x_0) \rightsquigarrow \text{observation}(\xi_1) \rightsquigarrow \text{decision}(x_1) \rightsquigarrow \\ \ldots \rightsquigarrow \text{observation}(\xi_T) \rightsquigarrow \text{decision}(x_T).$$

There are several ways how this decision process can be made precise. Nested formulation of a T-stage stochastic programming problem:

$$\operatorname*{Min}_{x_1 \in \mathcal{X}_1} F_1(x_1) + \mathbb{E}\Big\{ \operatorname*{Min}_{x_1 \in \mathcal{X}_2(x_2, \xi_2)} F_2(x_2, \xi_2) + \cdots + \mathbb{E}\Big[\operatorname*{Min}_{x_T \in \mathcal{X}_T(x_{T-1}, \xi_T)} F_T(x_T, \xi_T) \Big]$$

In linear case, $F_t(x_t, \xi_t) := c_t^{\mathsf{T}} x_t$ and

$$\mathcal{X}_t(x_{t-1}, \xi_t) := \{x_t : B_t x_{t-1} + A_t x_t = b_t,\ x_t \geq 0\},\ \ t = 2, ..., T.$$

The decisions $x_t = x_t(\xi_{[t]})$, $t = 2, ..., T$, are supposed to be functions of the history of the process up to time t. Such decision process (called a policy) is feasible if

$$x_t(\xi_{[t]}) \in \mathcal{X}_t(x_{t-1}(\xi_{[t-1]}), \xi_t) \text{ w.p.1.}$$

If the number of realizations (scenarios) of the process ξ_t is finite, then the above (linear) problem can be written as one large (linear) programming problem. In that respect it is convenient to represent the random process in a form of scenario tree.

Dynamic programming equations. Going recursively backwards in time. At stage T consider

$$Q_T(x_{T-1}, \xi_T) := \inf_{x_T \in \mathcal{X}_T(x_{T-1}, \xi_T)} F_T(x_T, \xi_T).$$

At stages $t = T-1, ..., 2$, consider

$$Q_t(x_{t-1}, \xi_{[t]}) := \inf_{x_t \in \mathcal{X}_t(x_{t-1}, \xi_t)} F_t(x_t, \xi_t) + \underbrace{\mathbb{E}\left[Q_{t+1}(x_t, \xi_{[t+1]}) \middle| \xi_{[t]}\right]}_{\mathcal{Q}_{t+1}(x_t, \xi_{[t]})}.$$

At the first stage solve:

$$\operatorname*{Min}_{x_1 \in \mathcal{X}_1} F_1(x_1) + \mathbb{E}[Q_2(x_1, \xi_1)].$$

A policy $\bar{x}_t = \bar{x}_t(\xi_{[t]})$ is optimal iff

$$\bar{x}_t \in \arg \min_{x_t \in \mathcal{X}_t(\bar{x}_{t-1}, \xi_t)} \left\{ F_t(x_t, \xi_t) + \mathbb{E}\left[Q_{t+1}(x_t, \xi_{[t+1]}) \middle| \xi_{[t]}\right] \right\}.$$

If the random process is between stages independent, i.e., ξ_{t+1} is independent of $\xi_{[t]}$, then $\mathcal{Q}_{t+1}(x_t) = \mathbb{E}[Q_{t+1}(x_t, \xi_{[t+1]}) | \xi_{[t]}]$ does not depend on $\xi_{[t]}$.

Conditional sampling. Let ξ_2^i, $i = 1, ..., N_1$, be an iid random sample of $\boldsymbol{\xi}_2$. Conditional on $\boldsymbol{\xi}_2 = \xi_2^i$, a random sample ξ_3^{ij}, $j = 1, ..., N_2$, is generated and etc. The obtained scenario tree is considered as a sample approximation of the true problem. Note that the total number of scenarios $N = \prod_{t=1}^{T-1} N_t$ and each scenario in the generated tree is considered with the same probability $1/N$. Note also that in the case of between stages independence of the corresponding random process, we have two possible strategies. We can generate a different (independent) sample ξ_3^{ij}, $j = 1, ..., N_2$, for every generated node ξ_2^i, or we can use the same sample ξ_3^j, $j = 1, ..., N_2$, for every ξ_2^i. In the second case we preserve the between stages condition for the generated scenario tree.

For $T = 3$, under certain regularity conditions, for $\varepsilon > 0$ and $\alpha \in (0,1)$, and the sample sizes N_1 and N_2 satisfying

$$O(1)\left[\left(\frac{D_1 L_1}{\varepsilon}\right)^{n_1} \exp\left\{-\frac{O(1)N_1\varepsilon^2}{\sigma_1^2}\right\} + \left(\frac{D_2 L_2}{\varepsilon}\right)^{n_2} \exp\left\{-\frac{O(1)N_2\varepsilon^2}{\sigma_2^2}\right\}\right] \le \alpha,$$

we have that any first-stage $\varepsilon/2$-optimal solution of the SAA problem is an ε-optimal first-stage solution of the true problem with probability at least $1 - \alpha$.

In particular, suppose that $N_1 = N_2$ and take $L := \max\{L_1, L_2\}$, $D := \max\{D_1, D_2\}$, $\sigma^2 := \max\{\sigma_1^2, \sigma_2^2\}$ and $n := \max\{n_1, n_2\}$. Then the required sample size $N_1 = N_2$:

$$\mathbf{N_1} \ge \frac{\mathbf{O(1)}\sigma^{\mathbf{2}}}{\varepsilon^{\mathbf{2}}}\left[\mathbf{n} \log\left(\frac{\mathbf{O(1)DL}}{\varepsilon}\right) + \log\left(\frac{\mathbf{1}}{\alpha}\right)\right],$$

with total number of scenarios $N = N_1^2$ (Shapiro, 2006).