Stochastic Programming

A. Shapiro

School of Industrial and Systems Engineering,
Georgia Institute of Technology,
Atlanta, Georgia 30332-0205, USA

eVITA Winter School 2009, Norway

Modeling and Basic Properties

Consider optimization problem:

$$\begin{array}{ll} \mathrm{Min}_{x \in X} & F(x, \xi) \\ \text{subject to} & c_i(x, \xi) \leq 0, \ i = 1, ..., q. \end{array}$$

Here $X \subset \mathbb{R}^n$ and $\xi \in \Xi \subset \mathbb{R}^d$ is a parameter vector representing "uncertainty" of the problem.

Robust (worst case) approach:

$$\begin{array}{ll} \mathrm{Min}_{x \in X} & \left\{ f(x) := \mathrm{max}_{\xi \in \Xi} F(x, \xi) \right\} \\ \text{subject to} & c_i(x, \xi) \leq 0, \ i = 1, ..., q, \ \forall \xi \in \Xi. \end{array}$$

Here Ξ is viewed as the uncertainty set for parameter vector Ξ.

Stochastic optimization approach: view ξ as a random vector with a known (given) probability measure (distribution) on Ξ.

The Newsvendor Problem

Suppose that a company has to decide about order quantity x of a certain product to satisfy demand d. The cost of ordering is $c > 0$ per unit. If the demand d is larger than x, then the newsvendor makes an additional order for the unit price $b \geq 0$. The cost of this is equal to $b(d-x)$ if $d > x$, and is zero otherwise. On the other hand, if $d < x$, then holding cost of $h(x-d) \geq 0$ is incurred. The total cost is then equal to

$$F(x,d) = cx + b[d-x]_+ + h[x-d]_+. \tag{1}$$

We assume that $b > c$, i.e., the back order penalty cost is *larger* than the ordering cost. The objective is to minimize the total cost $F(x,d)$.

Consider the case when the ordering decision should be made *before* a realization of the demand becomes known. One possible way to proceed in such situation is to view the demand D as a *random variable*. By capital D we denote the demand when viewed as a random variable in order to distinguish it from its particular realization d. We assume, further, that the probability distribution of D is *known*. This makes sense in situations where the ordering procedure repeats itself and the distribution of D can be estimated from historical data. Then it makes sense to talk about the expected value, denoted $\mathbb{E}[F(x, D)]$, of the total cost viewed as a function of the order quantity x. Consequently, we can write the corresponding optimization problem

$$\operatorname*{Min}_{x \geq 0} \Big\{ f(x) := \mathbb{E}[F(x, D)] \Big\}. \tag{2}$$

The above problem gives a very simple example of a *recourse action*. At the first stage, before a realization of the demand D is known, one has to make a decision about ordering quantity x. At the second stage after demand D becomes known, it may happen that $d > x$. In that case the company takes the recourse action of ordering the required quantity $d - x$ at the higher cost of $b > c$.

In the present case problem (2) can be solved in a closed form. Consider the cumulative distribution function (cdf)

$$H(x) := \text{Prob}(D \leq x)$$

of the random variable D. Note that $H(x) = 0$ for all $x < 0$, because the demand cannot be negative. It is possible to show that an optimal solution of problem (2) is equal to the quantile

$$\bar{x} = H^{-1}(\kappa), \quad \text{with} \quad \kappa = \frac{b-c}{b+h}. \tag{3}$$

Worst case approach.

Suppose that we know upper and lower bounds on the demand d, i.e., $\ell \leq d \leq u$. Consider the problem

$$\operatorname*{Min}_{x \geq 0} \left\{ f(x) := \max_{d \in [\ell, u]} F(x, d) \right\}. \tag{4}$$

We have here

$$f(x) = \max \left\{ cx + b[\ell - x]_+ + h[x - \ell]_+, cx + b[u - x]_+ + h[u - d]_+ \right\},$$

and if $b > h$, then $\bar{x} = u$ is the optimal solution of problem (4).

Example of financial planning

Suppose that we want to invest an amount of W_0 in n assets, x_i, $i = 1, ..., n$, in each. That is,

$$W_0 = \sum_{i=1}^{n} x_i. \tag{5}$$

After one period of time our wealth becomes

$$W_1 = \sum_{i=1}^{n} \xi_i x_i, \tag{6}$$

where $\xi_i = 1 + R_i$ and R_i is the return of the i-th asset. We would like to maximize W_1 by making an "optimal" distribution of our initial wealth. Of course, we have to make a decision about x_i *before* a realization of the returns R_i (of ξ_i) becomes known.

Suppose that we have an idea, may be from historical data, about probability distribution of $\xi = (\xi_1, ..., \xi_n)$. Then we may think about maximizing W_1 on *average*. That is, we would like to maximize the expected value $\mathbb{E}[W_1]$ of our wealth subject to the budget constraint $\sum x_i = W_0$ and "no borrowing" constraints $x_i \geq 0$. This leads to the optimization problem

$$\text{Max}_{x \geq 0} \ \mathbb{E}[W_1] \ \text{ subject to } \sum_{i=1}^n x_i = W_0. \tag{7}$$

We have that

$$\mathbb{E}[W_1] = \mathbb{E}\left[\sum_{i=1}^n \xi_i x_i\right] = \sum_{i=1}^n \mu_i x_i,$$

where $\mu_i = \mathbb{E}[\xi_i]$. Consequently, problem (7) has the simple optimal solution of investing everything into the asset with the maximal expected return.

Suppose now that we have a target wealth of τ. If W_1 falls short of τ we are penalized by $r(W_1 - \tau)$, and if W_1 exceeds τ we are rewarded by $q(W_1 - \tau)$, with $r > q$. This leads to the concept of utility function

$$U(w) = \begin{cases} r(w-\tau), & \text{if} \quad w \leq \tau \\ q(w-\tau), & \text{if} \quad w \geq \tau, \end{cases}$$

and to the optimization problem

$$\text{Max}_{x \geq 0} \ \mathbb{E}\left[F(x,\xi)\right] \quad \text{subject to } \sum_{i=1}^{n} x_i = W_0, \tag{8}$$

where $F(x,\xi) = U\left(\sum_{i=1}^{n} \xi_i x_i\right)$.

Chance (probabilistic) constraints formulation

$$\begin{array}{ll} \mathsf{Max}_{x \geq 0} & \mu_R^{\mathsf{T}} x \\ \text{subject to} & \mathsf{Prob}\left(R^{\mathsf{T}} x < -b\right) \leq \alpha,\ \sum_{i=1}^n x_i = W_0, \end{array} \tag{9}$$

where $R^{\mathsf{T}} x = \sum_{i=1}^n R_i x_i$, $\mu_R = \mathbb{E}[R]$ and $\alpha \in (0,1)$ is a chosen significance level. The above probability constraint means that the probability of loosing more than a given amount $b > 0$ is no more than α, and is called the *Value at Risk* constraint. If R has a (multivariate) normal distribution $N(\mu_R, \Sigma)$, then $R^{\mathsf{T}} x \sim N(\mu_R^{\mathsf{T}} x, x^{\mathsf{T}} \Sigma x)$ and the probabilistic constraint is equivalent to:

$$b + \mu_R^{\mathsf{T}} x - z_\alpha (x^{\mathsf{T}} \Sigma\, x)^{1/2} \geq 0, \tag{10}$$

where z_α is the $(1-\alpha)$-quantile of the standard normal distribution. Note that $z_\alpha > 0$ and the left hand side of (10) is a concave function of x, provided that $\alpha \in (0, 1/2)$.

By convex duality, there exists $\lambda \geq 0$ such that problem (9) is equivalent to the problem

$$\begin{array}{ll} \text{Max}_{x \geq 0} & \mu_R^{\mathsf{T}} x - \lambda (x^{\mathsf{T}} \Sigma\, x)^{1/2} \\ \text{subject to} & \sum_{i=1}^n x_i = W_0. \end{array} \tag{11}$$

The above problem can be viewed as a compromise between optimizing (maximizing) the expected return $\mu_R^{\mathsf{T}} x$ and minimizing risk term $\lambda (x^{\mathsf{T}} \Sigma\, x)^{1/2}$. In general (for non-normal distributions or nonlinear return functions), it could be difficult to handle probabilistic constraints numerically.

Risk averse formulation (Markowitz, 1952):

$$\begin{array}{ll} \text{Max}_{x \geq 0} & \overbrace{\mathbb{E}[R^{\mathsf{T}} x]}^{\mu_R^{\mathsf{T}} x} - \lambda \overbrace{\text{Var}[R^{\mathsf{T}} x]}^{x^{\mathsf{T}} \Sigma\, x} \\ \text{subject to} & \sum_{i=1}^n x_i = W_0. \end{array} \tag{12}$$

Equivalent formulations:

$$\begin{array}{ll} \text{Max}_{x\geq 0} & \sum_{i=1}^{n} \mu_i x_i - \lambda\, x^T \Sigma x \\ \text{s.t.} & \sum_{i=1}^{n} x_i = W_0, \end{array}$$

$$\begin{array}{ll} \text{Min}_{x\geq 0} & x^T \Sigma x \\ \text{s.t.} & \sum_{i=1}^{n} \mu_i x_i \geq \tau,\ \sum_{i=1}^{n} x_i = W_0, \end{array}$$

$$\begin{array}{ll} \text{Max}_{x\geq 0} & \sum_{i=1}^{n} \mu_i x_i \\ \text{s.t.} & \sum_{i=1}^{n} x_i = W_0,\ x^T \Sigma x \leq \gamma. \end{array}$$

Supply Chain Network Design

A supply chain is a network of suppliers, manufacturing plants, warehouses, and distribution channels organized to acquire raw materials, convert these raw materials to finished products, and distribute these products to customers. Denote by $\mathcal{S}, \mathcal{P}$ and $\mathcal{C}$ the respective (finite) sets of suppliers, processing facilities and customers. The union $\mathcal{N} := \mathcal{S} \cup \mathcal{P} \cup \mathcal{C}$ of these sets is viewed as the set of nodes of a directed graph $(\mathcal{N}, \mathcal{A})$, where $\mathcal{A}$ is a set of arcs (directed links) connecting these nodes in a way representing flow of the products. The processing facilities include manufacturing centers $\mathcal{M}$, finishing facilities $\mathcal{F}$ and warehouses $\mathcal{W}$, i.e., $\mathcal{P} = \mathcal{M} \cup \mathcal{F} \cup \mathcal{W}$. Further, a manufacturing center $i \in \mathcal{M}$ or a finishing facility $i \in \mathcal{F}$ consists of a set of manufacturing or finishing machines $\mathcal{H}_i$. Thus the set $\mathcal{P}$ includes the processing centers as well as the machines in these centers. Let $\mathcal{K}$ be the set of products flowing through the supply chain.

The supply chain configuration decisions consist of deciding which of the processing centers to build (major configuration decisions), and which processing and finishing machines to procure (minor configuration decisions). Assign a binary variable $x_i = 1$, if a processing facility i is built or machine i is procured, and $x_i = 0$ otherwise. The operational decisions consist of routing the flow of product $k \in \mathcal{K}$ from the supplier to the customers. By y_{ij}^k we denote the flow of product k from a node i to a node j of the network where $(ij) \in \mathcal{A}$. A deterministic mathematical model for the supply chain design problem can be written as follows

$$
\begin{array}{ll}
\underset{x\in X,\, y\geq 0}{\text{Min}} & \sum_{i\in\mathcal{P}} c_i x_i + \sum_{k\in\mathcal{K}} \sum_{(ij)\in\mathcal{A}} q_{ij}^k y_{ij}^k \\
\text{s.t.} & \sum_{i\in\mathcal{N}} y_{ij}^k - \sum_{\ell\in\mathcal{N}} y_{j\ell}^k = 0, \;\; j\in\mathcal{P},\; k\in\mathcal{K}, \\
& \sum_{i\in\mathcal{N}} y_{ij}^k \geq d_j^k, \;\; j\in\mathcal{C},\; k\in\mathcal{K}, \\
& \sum_{i\in\mathcal{N}} y_{ij}^k \leq s_j^k, \;\; j\in\mathcal{S},\; k\in\mathcal{K}, \\
& \sum_{k\in\mathcal{K}} r_j^k \left(\sum_{i\in\mathcal{N}} y_{ij}^k\right) \leq m_j x_j,\; j\in\mathcal{P}.
\end{array}
\tag{13}
$$

Here c_i denotes the investment cost for building facility i or procuring machine i, q_{ij}^k denotes the per-unit cost of processing product k at facility i and/or transporting product k on arc $(ij) \in \mathcal{A}$, d_j^k denotes the demand of product k at node j, s_j^k denotes the supply of product k at node j, r_j^k denotes per-unit processing requirement for product k at node j, m_j denotes capacity of facility j, $X \subset \{0,1\}^{|\mathcal{P}|}$ is a set of binary variables and $y \in \mathbb{R}^{|\mathcal{A}|\times|\mathcal{K}|}$ is a vector with components y_{ij}^k. The set X represents logical dependencies and restrictions.

The objective function of (13) is aimed at minimizing total investment and operational costs. Assume that at time a decision about vector $x \in X$ should be made, i.e., which facilities to built and machines to procure, there is an uncertainty about parameters involved in operational decisions represented by vector $y \in \mathbb{R}^{|\mathcal{A}|\times|\mathcal{K}|}$. This naturally classifies decision variables x as first stage decision variables and y as second stage decision variables. Note that problem (13) can be written in the following equivalent form as a two stage program:

$$\underset{x\in X}{\text{Min}}\, c^{\mathsf{T}}x + Q(x,\xi), \tag{14}$$

where $Q(x,\xi)$ is the optimal value of the second stage problem

$$
\begin{array}{ll}
\underset{y\geq 0}{\text{Min}} & \sum_{i\in\mathcal{P}} c_i x_i + \sum_{k\in\mathcal{K}} \sum_{(ij)\in\mathcal{A}} q_{ij}^k y_{ij}^k \\
\text{s.t.} & \sum_{i\in\mathcal{N}} y_{ij}^k - \sum_{\ell\in\mathcal{N}} y_{j\ell}^k = 0, \;\; j \in \mathcal{P},\; k \in \mathcal{K}, \\
& \sum_{i\in\mathcal{N}} y_{ij}^k \geq d_j^k, \;\; j \in \mathcal{C},\; k \in \mathcal{K}, \\
& \sum_{i\in\mathcal{N}} y_{ij}^k \leq s_j^k, \;\; j \in \mathcal{S},\; k \in \mathcal{K}, \\
& \sum_{k\in\mathcal{K}} r_j^k \left(\sum_{i\in\mathcal{N}} y_{ij}^k\right) \leq m_j x_j,\; j \in \mathcal{P}.
\end{array}
\tag{15}
$$

with $\xi = (q, d, s, R, M)$ being vector of the involved parameters.The above optimization problem depends on the data vector ξ and decision vector x. If some of the data parameters are uncertain, then the deterministic problem (14) does not make much sense since it depends on unknown parameters.

Suppose that we can model uncertain components of the data vector ξ as random variables with a specified probability distribution. Then we can formulate the following stochastic programming problem

$$\operatorname*{Min}_{x \in X} c^{\mathsf{T}} x + \mathbb{E}[Q(x, \xi)], \tag{16}$$

where the expectation is taken with respect to the probability distribution of the random vector ξ. That is, the cost of the second stage problem enters the objective of the first stage problem *on average*. A distinctive feature of the above stochastic programming problem is that the first stage problem here is a combinatorial problem with binary decision variables and finite feasible set X. On the other hand, the second stage problem is a linear programming problem and its optimal value $Q(x, \xi)$ is convex in x (if x is viewed as a vector in $\mathbb{R}^{|\mathcal{P}|}$).

The concept of two-stage (linear) stochastic programming problem with recourse

$$\underset{x \in X}{\text{Min}}\, c^{\mathsf{T}} x + \mathbb{E}[Q(x, \boldsymbol{\xi})], \tag{17}$$

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

$$\underset{y}{\text{Min}}\, q^{\mathsf{T}} y \ \text{ s.t. } \ Tx + Wy = h,\ y \geq 0, \tag{18}$$

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 $p_1, ..., p_K$. In that case

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

where

$$Q(x, \xi_k) = \inf \left\{ q_k^{\mathsf{T}} y_k : T_k x + W_k y_k = h_k,\ y_k \geq 0 \right\}.$$

It follows that we can write problem (17)-(18) as one large linear program:

$$\begin{array}{ll} \underset{x, y_1, ..., y_K}{\text{Min}} & c^{\mathsf{T}} x + \sum_{k=1}^{K} p_k q_k^{\mathsf{T}} y_k \\ \text{subject to} & Ax = b, \\ & T_k x + W_k y_k = h_k,\ k = 1, ..., K, \\ & x \geq 0,\ y_k \geq 0,\ k = 1, ..., K. \end{array} \tag{19}$$

Even crude discretization of the distribution of the data vector $\boldsymbol{\xi}$ leads to an exponential growth of the number of scenarios with increase of its dimension d.

Could stochastic programming problems be solved numerically?

What does it mean to solve a stochastic program?

How do we know the probability distribution of the random data vector?

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

Basic properties

For any realization ξ, the function $Q(\cdot, \xi)$ is convex piecewise linear. By the duality theory of linear programming we can write it in the following equivalent form

$$Q(x, \xi) = \sup \left\{ \pi^{\mathsf{T}}(h - Tx) : W^{\mathsf{T}}\pi \leq q \right\}. \tag{20}$$

It follows that the expectation function $\mathcal{Q}(x) = \mathbb{E}[Q(x, \boldsymbol{\xi})]$ is convex, and if P has a finite support (i.e., the number of scenarios is finite), then $\mathcal{Q}(x)$ is piecewise linear. Note that it can happen that, for some (x, ξ), the feasible set of problem (18) is empty. In that case, by the definition, $Q(x, \xi) = +\infty$. It also can happen that problem (18) is unbounded from below, and hence $Q(x, \xi) = -\infty$. That is, we can view $Q(x, \xi)$ as an extended real valued function.

Since $Q(\cdot, \xi)$ is a piecewise linear function, it can be differentiable everywhere only in the trivial case when it is linear. Nevertheless, if $Q(\cdot, \xi)$ is finite at a point $\bar{x}$, then it has a nonempty set of subgradients. The set of all subgradients is called *subdifferential* and denoted by $\partial Q(\bar{x}, \xi)$. Recall that $z \in \partial Q(\bar{x}, \xi)$ if

$$Q(x, \xi) \geq Q(\bar{x}, \xi) + z^{\mathsf{T}}(x - \bar{x}), \text{ for all } x.$$

The function $Q(\cdot, \xi)$ is differentiable at a point x iff $\partial Q(x, \xi) = \{z\}$ is a singleton, in which case $\nabla_x Q(x, \xi) = z$. The set $\partial Q(x, \xi)$ is convex, and since $Q(\cdot, \xi)$ is piecewise linear, is polyhedral. By duality theory we have that

$$\partial Q(x, \xi) = -T^{\mathsf{T}} \mathfrak{D}(x, \xi), \tag{21}$$

where $\mathfrak{D}(x, \xi) := \arg\max \left\{ \pi^{\mathsf{T}}(h - Tx) : W^{\mathsf{T}} \pi \leq q \right\}$.

If P has a finite support, then the subdifferential of the expectation function $\mathcal{Q}(\cdot)$ is given* by

$$\partial \mathcal{Q}(x) = \sum_{k=1}^{K} p_k \partial Q(x, \xi_k). \tag{22}$$

Therefore, $\mathcal{Q}(\cdot)$ is differentiable at x iff all functions $Q(\cdot, \xi_k)$, $k = 1, ..., K$, are differentiable at x. If the probability distribution P is continuous, then the situation is more subtle. It is possible to show that if $\mathcal{Q}(\cdot)$ is finite valued in a neighborhood of x, then

$$\partial \mathcal{Q}(x) = \int_{\Omega} \partial Q(x, \omega) dP(\omega). \tag{23}$$

For a given x, the above integral is understood as the set of all vectors of the form $\int_{\Omega} G(\omega) dP(\omega)$ such that $G(\omega) \in \partial Q(x, \omega)$ is an integrable selection of $\partial Q(x, \omega)$.

*The summation of the sets is understood here pointwise, i.e., the sum of two sets A and B is the set $\{a + b : a \in A,\ b \in B\}$.

It follows from (23) that $\partial \mathcal{Q}(x)$ is a singleton, and hence $\mathcal{Q}(\cdot)$ is differentiable at x, iff $\partial Q(x, \omega)$ is a singleton with probability one, i.e., for P-almost every $\omega \in \Omega$.

Loosely speaking we may say that, typically, for continuous distributions the expectation function $\mathbb{E}[Q(x, \boldsymbol{\xi})]$ *is differentiable, while in the case of discrete distributions it is not.*

We can formulate optimality conditions for the stochastic problem (17) as follows: a feasible point $\bar{x} \in X$ is an optimal solution of (17) iff

$$0 \in \partial \mathcal{Q}(\bar{x}) + N_X(\bar{x}), \tag{24}$$

where $N_X(\bar{x})$ is the normal cone to X at $\bar{x}$,

$$N_X(\bar{x}) = \left\{ z : z^{\mathsf{T}}(x - \bar{x}) \leq 0, \text{ for all } x \in X \right\}.$$

General formulation of two-stage stochastic programming problems

$$\operatorname*{Min}_{x \in X} \Big\{ f(x) := \mathbb{E}[F(x, \omega)] \Big\}, \tag{25}$$

where $F(x, \omega)$ is the optimal value of the second stage problem

$$\operatorname*{Min}_{y \in \mathcal{X}(x, \omega)} g(x, y, \omega). \tag{26}$$

Here $(\Omega, \mathcal{F}, P)$ is a probability space, $X \subset \mathbb{R}^n$, $g : \mathbb{R}^n \times \mathbb{R}^m \times \Omega \to \mathbb{R}$ and $\mathcal{X} : \mathbb{R}^n \times \Omega \rightrightarrows \mathbb{R}^m$ is a multifunction. In particular, the linear two-stage problem can be formulated in the above form with $g(x, y, \omega) := c^{\mathsf{T}} x + q(\omega)^{\mathsf{T}} y$ and

$$\mathcal{X}(x, \omega) := \{ y : T(\omega) x + W(\omega) y = h(\omega),\ y \geq 0 \}. \tag{27}$$

The second stage problem (26) can be also written in the following equivalent form

$$\operatorname*{Min}_{y\in\mathbb{R}^m} \bar{g}(x,y,\omega), \tag{28}$$

where

$$\bar{g}(x,y,\omega) := \begin{cases} g(x,y,\omega), & \text{if } y \in \mathcal{X}(x,\omega) \\ +\infty, & \text{otherwise.} \end{cases}$$

By the interchangeability principle we have

$$\underbrace{\mathbb{E}\left[\inf_{y\in\mathbb{R}^m} \bar{g}(x,y,\omega)\right]}_{F(x,\omega)} = \inf_{\boldsymbol{y}\in\mathfrak{Y}} \mathbb{E}\big[\bar{g}(x,\boldsymbol{y}(\omega),\omega)\big], \tag{29}$$

where $\mathfrak{Y}$ is a functional space, e.g., $\mathfrak{Y} := \mathcal{L}_p(\Omega,\mathcal{F},P;\mathbb{R}^m)$ with $p \in [1,+\infty]$.

Consequently, we can write two-stage problem (25)–(26) as one large problem:

$$\begin{aligned} \underset{x\in\mathbb{R}^n,\boldsymbol{y}\in\mathcal{Y}}{\text{Min}} \quad & \mathbb{E}\left[g(x,\boldsymbol{y}(\omega),\omega)\right] \\ \text{s.t.} \quad & x\in X,\ \boldsymbol{y}(\omega)\in\mathcal{X}(x,\omega) \text{ a.e. } \omega\in\Omega. \end{aligned} \tag{30}$$

Nonanticipativity

Consider the first stage problem (25). Assume that the number of scenarios is finite, i.e., $\Omega = \{\omega_1,\ldots,\omega_K\}$ with respective (positive) probabilities $p_1,\ldots,p_K$. Let us relax the first stage problem by replacing vector x with K vectors $x_1, x_2, \ldots, x_K$, one for each scenario. We obtain the following relaxation of problem (25):

$$\underset{x_1,\ldots,x_K}{\text{Min}} \sum_{k=1}^{K} p_k F(x_k,\omega_k) \text{ subject to } x_k\in X,\ k=1,\ldots,K. \tag{31}$$

Problem (31) is separable in the sense that it can be split into K smaller problems, one for each scenario:

$$\underset{x_k \in X}{\text{Min}}\, F(x_k, \omega_k), \quad k = 1, \ldots, K, \tag{32}$$

and that the optimal value of problem (31) is equal to the weighted sum, with weights p_k, of the optimal values of problems (32), $k = 1, \ldots, K$.

The *nonanticipativity* constraint: $(x_1, \ldots, x_K) \in \mathfrak{L}$, where

$$\mathfrak{L} := \{(x_1, \ldots, x_K) : x_1 = \ldots = x_K\} \subset \mathbb{R}^n \times \cdots \times \mathbb{R}^n.$$

Dualization of the nonanticipativity constraints

Consider the following formulation of stochastic programming problem (25) (with a finite number of scenarios),

$$\underset{x_1,...,x_K,z}{\text{Min}} \sum_{k=1}^{K} p_k \bar{F}(x_k, \omega_k) \text{ s.t. } x_k = z,\ k = 1, ..., K, \tag{33}$$

where $\bar{F}(x_k, \omega_k) = F(x_k, \omega_k)$ if $x_k \in X$ and $\bar{F}(x_k, \omega_k) = +\infty$ otherwise. The nonaticipativity constraints $x_k = z$, $k = 1, ..., K$, are written here with additional variable z. The (Lagrangian) dual of problem (33) is:

$$\begin{array}{ll} \text{Max}_{\lambda_1,...,\lambda_K} & \left\{\text{inf}_{x_1,...,x_K} \sum_{k=1}^{K} p_k \left(\bar{F}(x_k, \omega_k) + \lambda_k^{\mathsf{T}} x_k\right)\right\}, \\ \text{subject to} & \sum_{k=1}^{K} p_k \lambda_k = 0. \end{array} \tag{34}$$

Note the separable structure

$$\inf_{x_1,\ldots,x_K} \sum_{k=1}^{K} p_k \left(\bar{F}(x_k, \omega_k) + \lambda_k^{\mathsf{T}} x_k \right) = \sum_{k=1}^{K} p_k \left[\inf_{x_k} \left(\bar{F}(x_k, \omega_k) + \lambda_k^{\mathsf{T}} x_k \right) \right].$$

If the functions $F_k(\cdot, \omega_k)$ are piecewise linear (e.g., in the case of linear two-stage stochastic programming), then there is no duality gap between (33) and (34), and both problems have optimal solutions provided that their optimal value is finite. Moreover, if $(\bar{\lambda}_1, ..., \bar{\lambda}_K)$ and $(\bar{x}_1, ..., \bar{x}_K, \bar{z})$ are optimal solutions of (33) and (34), respectively, then $\bar{x}_1 = ... = \bar{x}_K = \bar{z}$ and

$$\bar{x}_k \in \arg\min_{x_k} \left\{ \bar{F}(x_k, \omega_k) + \bar{\lambda}_k^{\mathsf{T}} x_k \right\}. \tag{35}$$

Multistage models. Consider the newsvendor problem. Suppose now that the company has a planning horizon of T periods. We model the demand as a random process D_t indexed by the time $t = 1, ..., T$. At the beginning, at $t = 1$, there is (known) inventory level y_1. At each period $t = 1, ..., T$ the company first observes the current inventory level y_t and then places an order to replenish the inventory level to x_t. This results in order quantity $x_t - y_t$ which clearly should be nonnegative, i.e., $x_t \geq y_t$. After the inventory is replenished, demand d_t is realized and hence the next inventory level, at the beginning of period $t + 1$, becomes $y_{t+1} = x_t - d_t$. We allow backlogging and the inventory level y_t may become negative. The total cost incurred in period t is

$$c_t(x_t - y_t) + b_t[d_t - x_t]_+ + h_t[x_t - d_t]_+,$$

where c_t, b_t, h_t are the ordering cost, holding cost and backorder penalty cost per unit, respectively, at time t. We assume that $b_t > c_t > 0$ and $h_t \geq 0$, $t = 1, \ldots, T$.

The objective is to minimize the expected value of the total cost over the planning horizon. This can be written as the following optimization problem

$$\begin{aligned} \underset{x_t \geq y_t}{\text{Min}} \; & \sum_{t=1}^{T} \mathbb{E}\Big\{ c_t(x_t - y_t) + b_t[D_t - x_t]_+ + h_t[x_t - D_t]_+ \Big\} \\ \text{s.t. } & y_{t+1} = x_t - D_t, \; t = 1, ..., T-1. \end{aligned} \tag{36}$$

Consider the demand process D_t, $t = 1, ..., T$. We denote by $D_{[t]} := (D_1, ..., D_t)$ the history of the demand process up to time t, and by $d_{[t]} := (d_1, ..., d_t)$ its particular realization. At each period (stage) t, our decision about the inventory level x_t should depend only on information available at the time of the decision, i.e., on an observed realization $d_{[t-1]}$ of the demand process, and not on future observations. This principle is called the *nonanticipativity* constraint.

At the last stage $t = T$, for observed inventory level y_T, we need to solve the problem:

$$\underset{x_T \geq y_T}{\text{Min}} \; c_T(x_T - y_T) + \mathbb{E}\left\{ b_T[D_T - x_T]_+ + h_T[x_T - D_T]_+ \middle| D_{[T-1]} = d_{[T-1]} \right\}. \tag{37}$$

The expectation in (37) is conditional on the realization $d_{[T-1]}$ of the demand process prior to the considered time T. The optimal value (and the set of optimal solutions) of problem (37) depends on y_T and $d_{[T-1]}$, and is denoted $Q_T(y_T, d_{[T-1]})$. At stage $t = T - 1$ we solve the problem

$$\begin{aligned} \underset{x_{T-1} \geq y_{T-1}}{\text{Min}} \quad & c_{T-1}(x_{T-1} - y_{T-1}) \\ & + \mathbb{E}\Big\{ b_{T-1}[D_{T-1} - x_{T-1}]_+ + h_{T-1}[x_{T-1} - D_{T-1}]_+ \\ & \quad + Q_T\left(x_{T-1} - D_{T-1}, D_{[T-1]}\right) \Big| D_{[T-2]} = d_{[T-2]} \Big\}. \end{aligned}$$

Its optimal value is denoted $Q_{T-1}(y_{T-1}, d_{[T-2]})$. Proceeding in this way backwards in time we write the following *dynamic programming* equations

$$Q_t(y_t, d_{[t-1]}) = \min_{x_t \geq y_t} c_t(x_t - y_t) + \mathbb{E}\Big\{ b_t[D_t - x_t]_+ \\ + h_t[x_t - D_t]_+ + Q_{t+1}\left(x_t - D_t, D_{[t]}\right) \Big| D_{[t-1]} = d_{[t-1]} \Big\},$$

$t = T-1, ..., 2$. Finally, at the first stage we need to solve problem

$$\operatorname*{Min}_{x_1 \geq y_1} c_1(x_1 - y_1) + \mathbb{E}\Big\{ b_1[D_1 - x_1]_+ + h_1[x_1 - D_1]_+ + Q_2\left(x_1 - D_1, D_1\right) \Big\}.$$

Let $\bar{x}_t$, $t = T-1, ..., 1$, be an optimal solution of the corresponding dynamic programming equation. We see that $\bar{x}_t$ is a function of y_t and $d_{[t-1]}$, for $t = 2, ..., T$, while the first stage (optimal) decision $\bar{x}_1$ is independent of the data.

Under the assumption of the stagewise independence, $\bar{x}_t = \bar{x}_t(y_t)$ becomes a function of y_t alone. Note that y_t, in itself, is a function of $d_{[t-1]} = (d_1, ..., d_{t-1})$ and decisions $(x_1, ..., x_{t-1})$. Therefore we may think about a sequence of possible decisions $x_t = x_t(d_{[t-1]})$, $t = 1, ..., T$, as functions of realizations of the demand process available at the time of the decision (with the convention that x_1 is independent of the data). Such a sequence of decisions $x_t(d_{[t-1]})$ is called a *policy*. That is, a policy is a rule which specifies our decisions, based on information available at the current stage, for any possible realization of the demand process. By definition, a policy $x_t = x_t(d_{[t-1]})$ satisfies the nonanticipativity constraint. A policy is said to be *feasible* if it satisfies other constraints with probability one (w.p.1). In the present case a policy is feasible if $x_t \geq y_t$, $t = 1, ..., T$, for almost every realization of the demand process.

We can formulate optimization problem (36) as the problem of minimization of the expectation in (36) with respect to all feasible policies. An optimal solution of such problem will give us an optimal policy. We have that a policy $\bar{x}_t$ is optimal if it is given by optimal solutions of the respective dynamic programming equations. In the present case under the assumption of stagewise independence, an optimal policy $\bar{x}_t = \bar{x}_t(y_t)$ is a function of y_t alone. Moreover, in that case it is possible to give the following characterization of the optimal policy. Let x_t^* be an (unconstrained) minimizer of

$$c_t x_t + \mathbb{E}\Big\{ b_t[D_t - x_t]_+ + h_t[x_t - D_t]_+ + Q_{t+1}\left(x_t - D_t\right) \Big\}, \quad t = T, ..., 1. \tag{38}$$

By using convexity of the value functions it is not difficult to show that $\bar{x}_t = \max\{y_t, x_t^*\}$ is an optimal policy. Such policy is called the *basestock* policy.

Multistage portfolio selection. Suppose that we can rebalance our portfolio at several, say T, periods of time. That is, at the beginning we choose values x_{i0} of our assets subject to the budget constraint

$$\sum_{i=1}^{n} x_{i0} = W_0. \tag{39}$$

At the period $t = 1, ..., T$, our wealth is

$$W_t = \sum_{i=1}^{n} \xi_{it} x_{i,t-1}, \tag{40}$$

where $\xi_{it} = (1 + R_{it})$ and R_{it} is the return of the i-th asset at the period t. Our objective is to maximize the expected utility

$$\text{Max}\, \mathbb{E}\left[U(W_T)\right] \tag{41}$$

at the end of the considered period, subject to the balance constraints $\sum_{i=1}^{n} x_{it} = W_t$ and $x_t \geq 0$, $t = 0, ..., T-1$.

We use notation $x_t = (x_{1t}, ..., x_{nt})$ and $\xi_t = (\xi_{1t}, ..., \xi_{nt})$, and $\xi_{[t]} = (\xi_1, .., \xi_t)$ for 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. *Dynamic programming equations:* the cost-to-go function $Q_t(W_t, \xi_{[t]})$ is given by the optimal value of

$$\begin{aligned} \underset{x_t \geq 0, W_{t+1}}{\text{Max}} \quad & \mathbb{E}\Big\{ Q_{t+1}(W_{t+1}, \xi_{[t+1]}) \Big| \xi_{[t]} \Big\} \\ \text{s.t. } & W_{t+1} = \sum_{i=1}^{n} \xi_{i,t+1} x_{i,t}, \ \sum_{i=1}^{n} x_{i,t} = W_t. \end{aligned} \tag{42}$$

If the process ξ_t is stagewise independent, i.e., ξ_t is (stochastically) independent of $\xi_1, ..., \xi_{t-1}$, for $t = 2, ..., T$, then the cost-to-go (value) function $Q_t(W_t)$, $t = 1, ..., T-1$, does not depend on $\xi_{[t]}$.

Multistage stochastic programming. We can write the two-stage problem using the following *nested formulation*:

$$\underset{Ax=b,\, x\geq 0}{\text{Min}} c^{\mathsf{T}}x + \mathbb{E}\left[\underset{Tx+Wy=h,\, y\geq 0}{\text{Min}} q^{\mathsf{T}}y\right]. \tag{43}$$

In the above, $y = y(\xi)$ is considered as a function of the random data $\xi = (q, T, W, h)$ and in that sense is random. If the number of scenarios is finite, we associate with every possible realization ξ_k of the data the corresponding second stage decision variable y_k. This can be extended to the following nested formulation of a multistage stochastic programming problem:

$$\begin{aligned}\underset{x_1\in X}{\text{Min}}\, f_1(x_1) + \mathbb{E}\Big[& \underset{x_2\in\mathcal{X}_2(x_1,\xi_2)}{\text{Min}} f_2(x_2,\xi_2) + \cdots \\ & + \mathbb{E}\Big[\underset{x_T\in\mathcal{X}_T(x_{T-1},\xi_T)}{\text{Min}} f_T(x_T,\xi_T)\Big]\Big],\end{aligned} \tag{44}$$

where $\xi_1, ..., \xi_T$ is a random process (ξ_1 is deterministic).

For example, in the linear case $f_t(x_t, \xi_t) := c_t^\top x_t$,

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

$\xi_t = (c_t, B_t, A_t, b_t)$, $t = 2, ..., T$, is considered as a random process, $\xi_1 = (c_1, A_1, b_1)$ is supposed to be known, and hence the nested formulation can be written as

$$\underset{\substack{A_1 x_1 = b_1 \\ x_1 \geq 0}}{\text{Min}} c_1^\top x_1 + \mathbb{E}\left[\underset{\substack{B_2 x_1 + A_2 x_2 = b_2 \\ x_2 \geq 0}}{\text{Min}} c_2^\top x_2 + \cdots + \mathbb{E}\left[\underset{\substack{B_T x_{T-1} + A_T x_T = b_T \\ x_T \geq 0}}{\text{Min}} c_T^\top x_T\right]\right]. \tag{45}$$

If the number of realizations (scenarios) of the process ξ_t is finite, then problem (45) can be written as one large linear programming problem. There are several possible formulations of the above multistage program.

Dynamic programming equations Consider the last stage problem

$$\underset{x_T \in \mathcal{X}_T(x_{T-1}, \xi_T)}{\text{Min}} f_T(x_T, \xi_T). \tag{46}$$

The optimal value of this problem, denoted $Q_T(x_{T-1}, \xi_T)$, depends on the decision vector x_{T-1} and data ξ_T. At stage $t = 2, ..., T-1$, we write the problem:

$$\begin{array}{ll} \underset{x_t}{\text{Min}} & f_t(x_t, \xi_t) + \mathbb{E}\left\{ Q_{t+1}\left(x_t, \xi_{[t+1]}\right) \middle| \xi_{[t]} \right\} \\ \text{s.t.} & x_t \in \mathcal{X}_t(x_{t-1}, \xi_t). \end{array} \tag{47}$$

Its optimal value depends on the decision x_{t-1} at the previous stage and realization of the data process $\xi_{[t]}$, and denoted $Q_t\left(x_{t-1}, \xi_{[t]}\right)$. The idea is to calculate the (so-called *cost-to-go* or *value*) functions $Q_t\left(x_{t-1}, \xi_{[t])}\right)$, recursively, going backward in time.

At the first stage we finally need to solve the problem:

$$\underset{x_1 \in X}{\text{Min}} f_1(x_1) + \mathbb{E}\left[Q_2\left(x_t, \xi_2\right)\right]. \tag{48}$$

The dynamic programming equations:

$$Q_t\left(x_{t-1}, \xi_{[t]}\right) = \inf_{x_t \in \mathcal{X}_t(x_{t-1}, \xi_t)} \left\{ f_t(x_t, \xi_t) + \mathcal{Q}_{t+1}\left(x_t, \xi_{[t]}\right) \right\}, \tag{49}$$

where

$$\mathcal{Q}_{t+1}\left(x_t, \xi_{[t]}\right) := \mathbb{E}\left\{ Q_{t+1}\left(x_t, \xi_{[t+1]}\right) \middle| \xi_{[t]} \right\}.$$

If the random process is *Markovian* (i.e., the conditional distribution of ξ_{t+1} given $\xi_{[t]} = (\xi_1, ..., \xi_t)$ is the same as the conditional distribution of ξ_{t+1} given ξ_t), then $Q_t\left(x_{t-1}, \xi_t\right)$ is a function of x_{t-1} and ξ_t, and if it is *stagewise independent* (i.e., ξ_{t+1} is independent of $\xi_{[t]}$), then $\mathbb{E}\left[Q_{t+1}\left(x_t, \xi_{t+1}\right) \middle| \xi_t \right] = \mathcal{Q}_{t+1}(x_t)$ does not depend on ξ_t.

A sequence of (measurable) mappings $x_t(\xi_{[t]})$, $t = 1, ..., T$, is called a *policy* (recall that ξ_1 is deterministic). A policy is said to be feasible if it satisfies the feasibility constraints, i.e.,

$$x_t(\xi_{[t]}) \in \mathcal{X}_t(x_{t-1}(\xi_{[t-1]}), \xi_t), \ t = 2, ..., T, \text{ w.p.1.} \tag{50}$$

We can formulate the multistage problem (44) in the form

$$\begin{array}{cl} \underset{x_1, x_2(\cdot), ..., x_T(\cdot)}{\text{Min}} & \mathbb{E}\Big[f_1(x_1) + f_2(x_2(\xi_2), \xi_2) + ... + f_T\left(x_T(\xi_{[T]}), \xi_T\right)\Big] \\ \text{s.t.} & x_1 \in X, \ x_t(\xi_{[t]}) \in \mathcal{X}_t(x_{t-1}(\xi_{[t-1]}), \xi_t), \ t = 2, ..., T. \end{array}$$

Note that the above optimization is performed over feasible policies. A policy $\bar{x}_t(\xi_{[t]})$ is *optimal* if it satisfies the dynamic programming equations, i.e.,

$$\bar{x}_t(\xi_{[t]}) \in \arg \min_{x_t \in \mathcal{X}_t\left(\bar{x}_{t-1}(\xi_{[t-1]}), \xi_t\right)} \Big\{f_t(x_t, \xi_t) + \mathcal{Q}_{t+1}\left(x_t, \xi_{[t]}\right)\Big\}, \ \text{w.p.1.}$$

Consider the relaxed version of the linear multistage program (with finite number of scenarios):

$$
\begin{array}{lrcccccccl}
\text{Min} & \sum_{k=1}^{K} p_k \Big[c_1^{\mathsf{T}} x_1^k & + & (c_2^k)^{\mathsf{T}} x_2^k & + & (c_3^k)^{\mathsf{T}} x_3^k & + & \ldots & + & (c_T^k)^{\mathsf{T}} x_T^k \Big] \\
\text{s.t.} & A_1 x_1^k & & & & & & & & = b \\
& B_2^k x_1^k & + & A_2^k x_2^k & & & & & & = b \\
& & & B_3^k x_2^k & + & A_3^k x_3^k & & & & = b \\
& \ldots & \ldots & \ldots & \ldots & \ldots & \ldots & \ldots & \ldots & \ldots \\
& & & & & & & B_T^k x_{T-1}^k & + & A_T^k x_T^k = b \\
& x_1^k \geq 0, & & x_2^k \geq 0, & & x_3^k \geq 0, & & \ldots & & x_T^k \geq 0, \\
& k = 1, \ldots, K. & & & & & & & &
\end{array}
$$

In this problem all parts of the decision vector are allowed to depend on *all* parts of the random data, while each part x_t should be allowed to depend only on the data known up to stage t.

In order to correct this problem we should remember that at stage $t = 1, \ldots, T$, the scenarios that have the same history $\xi_{[t]}$ cannot be distinguished, so we need to enforce the *nonanticipativity constraints:*

$$x_t^k = x_t^\ell \text{ for all } k, \ell \text{ for which } \xi_{[t]}^k = \xi_{[t]}^\ell, \quad t = 1, \ldots, T. \qquad (51)$$

Together with the nonanticipativity constraints (51) the considered problem becomes equivalent to the original formulation.

Complexity of stochastic programs

Consider a two-stage stochastic programming problem. Even a crude discretization of the distribution of the random data vector ξ typically results in an exponential growth of the number of scenarios with increase of the number of random variables (dimension of ξ). The standard approach to dealing with this issue is to generate a manageable number of scenarios in some "representitative" way. For example, we can generate a random sample $\xi^1, ..., \xi^N$ of N realizations of the random vector ξ by using Monte Carlo sampling techniques. Then the expected value function $f(x) := \mathbb{E}[F(x, \xi)]$ can be approximated by the sample average function

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

where $p_j := 1/N$, $j = 1, ..., N$.

Consequently the true (expected value) problem stochastic problem can be approximated by the so-called sample average approximating (SAA) problem:

$$\operatorname*{Min}_{x \in X} \widehat{f}_N(x). \tag{52}$$

Note that once the sample is generated, the above SAA problem can be viewed as a two-stage problem with the corresponding set of scenarios $\left\{\xi^1, ..., \xi^N\right\}$ each scenario with equal probability $1/N$. A (naive) justification of the SAA method is that by the Law of Large Numbers, $\widehat{f}_N(x)$ converges to $f(x)$ w.p.1 as N tends to infinity. It is possible to show that, under mild regularity conditions, it follows that the optimal value $\widehat{v}_N$ and an optimal solution $\widehat{x}_N$ of the SAA problem (52) converge w.p.1 to their counterparts of the true problem.

It is known, however, that the convergence of Monte Carlo sampling based estimators is notoriously slow. By the Central Limit Theorem, for a fixed $x \in X$, the estimate $\widehat{f}_N(x)$ converges to $f(x)$ at a stochastic rate of $O_p(N^{-1/2})$. This indicates that by using Monte Carlo sampling techniques it is not possible to evaluate the expected value $f(x)$ very accurately.

By using Large Devitions exponential bounds it is possible to show that for given $\alpha \in (0,1)$, $\varepsilon > 0$ and $\delta \in [0, \alpha)$, the following estimate of the sample size N:

$$N \geq \frac{O(1)\sigma^2}{(\varepsilon - \delta)^2} \left[n \log \left(\frac{O(1)DL}{(\varepsilon - \delta)^2} \right) + \log \left(\frac{1}{\alpha} \right) \right] \tag{53}$$

guarantees that with probability at least $1 - \alpha$ any δ-optimal solution of the SAA problem is an ε-optimal solution of the true problem.

Here $O(1)$ is a generic constant, D is the diameter of the feasible set X (assumed to be finite), L is a Lipschitz constant of $F(\cdot, \xi)$ on X (holding for all $\xi \in \Xi$), and σ^2 is a constant representing a certain measure of variability of the objective function $F(x, \xi)$). That is, it is assumed that the moment generating function $M(t) := \mathbb{E}[e^{tZ}]$ of the random variable $Z := F(x', \xi) - F(x, \xi) - [f(x) - f(x')]$ satisfies

$$M(t) \leq \exp\left(\sigma^2 t^2/2\right), \ \ \forall t \in \mathbb{R}.$$

It is also assumed that the expected value function $f(x)$ is finite valued for all $x \in X$. Note that this implies that the recourse is relatively complete. In a sense, the sample size bound (53) gives an estimate of complexity of solving the true problem. It is proportional to the dimension n of the decision vector x and to $(\sigma/\varepsilon)^2$, while the other parameters appear on the logarithmic scale.

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)

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

for an optimal choice of constant θ (recall that v^* 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

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

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

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

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

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

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

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

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

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 solution $\widehat{x} \in X$? Two basic approaches: (1) Evaluate the gap $f(\widehat{x}) - v^*$, where v^* is the optimal value of the true problem. (2) Verify the KKT optimality conditions at $\widehat{x}$.

Statistical test based on estimation of $f(\widehat{x}) - v^*$: (i) Estimate $f(\widehat{x})$ by the sample average $\widehat{f}_{N'}(\widehat{x})$, using sample of a large size N'. (ii) Solve the SAA problem M times using M independent samples each of size N. Let $\widehat{v}_N^{(1)}, ..., \widehat{v}_N^{(M)}$ be the optimal values of the corresponding SAA problems. Estimate $\mathbb{E}[\widehat{v}_N]$ by the average $M^{-1} \sum_{j=1}^{M} \widehat{v}_N^{(j)}$. Note that $v^* - \mathbb{E}[\widehat{v}_N] \geq 0$ and

$$\mathbb{E}\left[\widehat{f}_{N'}(\widehat{x}) - M^{-1} \textstyle\sum_{j=1}^{M} \widehat{v}_N^{(j)}\right] = (f(\widehat{x}) - v^*) + (v^* - \mathbb{E}[\widehat{v}_N]).$$

Also $\mathbb{E}[\widehat{v}_N] = \mathbb{E}\left[M^{-1} \sum_{j=1}^{M} \widehat{v}_N^{(j)}\right]$ and hence $v^* - \mathbb{E}[\widehat{v}_N]$ represents the bias of the above average viewed as an estimator of v^*.

KKT statistical test Let

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

Suppose that the probability distribution is continuous. Then $F(\cdot, \boldsymbol{\xi})$ is differentiable at $\widehat{x}$ w.p.1 and

$$\nabla f(\widehat{x}) = \mathbb{E}_P \left[\nabla_x F(\widehat{x}, \boldsymbol{\xi})\right].$$

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

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

and $J(x) := \{i : c_i(x) = 0,\ i \in J\}$.

The idea of the KKT test is to estimate the distance

$$\delta(\widehat{x}) := \mathsf{dist}\left(-\nabla f(\widehat{x}), C(\widehat{x})\right),$$

by using the sample estimator

$$\widehat{\delta}_N(\widehat{x}) := \mathsf{dist}\left(-\nabla \widehat{f}_N(\widehat{x}), C(\widehat{x})\right).$$

The covariance matrix of $\nabla \widehat{f}_N(\widehat{x})$ can be estimated (from the same sample), and hence a confidence region for $\nabla f(\widehat{x})$ can be constructed. This allows a statistical validation of the KKT conditions.

Complexity of multistage stochastic programming

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

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

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 $N_1 = N_2$:

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

with total number of scenarios $N = N_1^2$. This indicates that the total number of scenarios needed to solve a T-stage stochastic program with a reasonable accuracy by the SAA method grows exponentially with increase of the number of stages T.

Risk-averse approach Consider the following formulation of the financial planning problem:

$$\operatorname*{Max}_{x \geq 0} \ \mathbb{E}[W_1] - \lambda \mathbb{D}[W_1] \quad \text{s.t.} \ \sum_{i=1}^{n} x_i = W_0, \tag{54}$$

where $\mathbb{D}[W_1]$ is a measure of dispersion (variability) of W_1 and $\lambda \geq 0$ represents a compromise weight between maximizing returns and minimizing risk of investment.

Markowitz (1952) approach: take $\mathbb{D}[W_1] := \mathrm{Var}[W_1]$. In the present (linear) case

$$\mathrm{Var}[W_1] = \mathrm{Var}\left[\sum_{i=1}^{n} \xi_i x_i\right] = x^T \Sigma x,$$

where Σ is the covariance matrix of ξ.

This gives the following (equivalent) formulations of the corresponding risk-averse problem:

$$\begin{array}{ll} \text{Max}_{x \geq 0} & \sum_{i=1}^{n} \mu_i x_i - \lambda\, x^T \Sigma x \\ \text{s.t.} & \sum_{i=1}^{n} x_i = W_0, \end{array}$$

$$\begin{array}{ll} \text{Min}_{x \geq 0} & x^T \Sigma x \\ \text{s.t.} & \sum_{i=1}^{n} \mu_i x_i \geq \tau,\ \sum_{i=1}^{n} x_i = W_0, \end{array}$$

$$\begin{array}{ll} \text{Max}_{x \geq 0} & \sum_{i=1}^{n} \mu_i x_i \\ \text{s.t.} & \sum_{i=1}^{n} x_i = W_0,\ x^T \Sigma x \leq \gamma. \end{array}$$

Min-max approach to stochastic programming:

$$\operatorname*{Min}_{x \in X} \left\{ f(x) := \sup_{\mu \in \mathcal{A}} \mathbb{E}_\mu[F(x, \omega)] \right\},$$

where $F : \mathbb{R}^n \times \Omega \to \mathbb{R}$ and $\mathcal{A}$ is a set of probability measures (distributions) on the sample space $(\Omega, \mathcal{F})$.

Optimization of mean-risk models:

$$\operatorname*{Min}_{x \in X} \rho[F_x(\omega)],$$

where $\rho : \mathcal{Z} \to \mathbb{R} \cup \{+\infty\}$ is a mean-risk function, $\mathcal{Z}$ is a (linear) space of "allowable" functions $Z(\omega)$ and $F_x(\cdot) = F(x, \cdot) \in \mathcal{Z}$ for all $x \in X$.

Markowitz's approach: $\rho(Z) := \mathbb{E}[Z] + c\text{Var}[Z], \quad Z \in \mathcal{Z}$, where $c > 0$ is a weight constant (note that here we deal with minimization, rather than maximization, problem).

Axiomatic approach (coherent measures of risk), by Artzner, Delbaen, Eber, Heath (1999):
(A1) Convexity:

$$\rho(\alpha Z_1 + (1-\alpha)Z_2) \leq \alpha\rho(Z_1) + (1-\alpha)\rho(Z_2)$$

for all $Z_1, Z_2 \in \mathcal{Z}$ and $\alpha \in [0,1]$.
(A2) Monotonicity: If $Z_1, Z_2 \in \mathcal{Z}$ and $Z_2 \geq Z_1$, then $\rho(Z_2) \geq \rho(Z_1)$.
(A3) Translation Equivariance: If $a \in \mathbb{R}$ and $Z \in \mathcal{Z}$, then $\rho(Z + a) = \rho(Z) + a$.
(A4) Positive Homogeneity:

$$\rho(\alpha Z) = \alpha\rho(Z), \quad Z \in \mathcal{Z},\ \alpha > 0.$$

Space $\mathcal{Z}$ is paired with a linear space $\mathcal{Y}$ of finite signed measures on $(\Omega, \mathcal{F})$ such that the scalar product (bilinear form)

$$\langle \mu, Z \rangle := \int_{\Omega} Z(\omega) d\mu(\omega)$$

is well defined for all $Z \in \mathcal{Z}$ and $\mu \in \mathcal{Y}$. Typical examples $\mathcal{Z} := L_p(\Omega, \mathcal{F}, P)$ and $\mathcal{Y} := L_q(\Omega, \mathcal{F}, P)$, where $p, q \in [1, +\infty]$ such that $1/p + 1/q = 1$, and P is a probability (reference) measure on $(\Omega, \mathcal{F})$.

Dual representation of risk functions

By Fenchel-Moreau theorem if ρ is convex (assumption (A1)) and lower semicontinuous, then

$$\rho(Z) = \sup_{\mu \in \mathcal{A}} \{\langle \mu, Z \rangle - \rho^*(\mu)\} ,$$

where

$$\rho^*(\mu) := \sup_{Z \in \mathcal{Z}} \{\langle \mu, Z \rangle - \rho(Z)\},$$
$$\mathcal{A} := \mathrm{dom}(\rho^*) = \{\mu \in \mathcal{Y} : \rho^*(\mu) < +\infty\}.$$

It is possible to show that condition (A2) (monotonicity) holds iff $\mu \succeq 0$ for every $\mu \in \mathcal{A}$. Condition (A3) (translation equivariance) holds iff $\mu(\Omega) = 1$ for every $\mu \in \mathcal{A}$. If ρ is positively homogeneous, then $\rho^*(\mu) = 0$ for every $\mu \in \mathcal{A}$. If conditions (A1)–(A4) hold, then $\mathcal{A}$ is a set of probability measures and

$$\rho(Z) = \sup_{\mu \in \mathcal{A}} \mathbb{E}_\mu[Z].$$

Consequently, problem $\mathrm{Min}_{x \in X} \rho[F(x, \omega)]$ is equivalent to the min-max problem.

Chance constrained problem:

$$\operatorname{Min}_{x\in X} f(x) \text{ subject to } p(x) \leq \alpha,$$

where $X \subset \mathbb{R}^n$ is a closed set, $f : \mathbb{R}^n \to \mathbb{R}$ is a continuous function, $\alpha \in (0,1)$ is a given significance level, ξ is a random vector, whose probability distribution P is supported on set $\Xi \subset \mathbb{R}^d$, $C : \mathbb{R}^n \times \Xi \to \mathbb{R}$ and

$$p(x) := \text{Prob}\Big\{C(x,\xi) > 0\Big\}$$

is the probability that constraint is violated at point $x \in X$. Several chance constraints

$$\text{Prob}\Big\{C_i(x,\xi) \leq 0,\ i = 1,...,q\Big\} \geq 1-\alpha,$$

can be reduced to one chance constraint by employing the max-function $C(x,\xi) := \max_{1\leq i\leq q} C_i(x,\xi)$.

Example of portfolio selection. Consider the problem

$$\operatorname*{Max}_{x\geq 0} \sum_{i=1}^{n} \mu_i x_i \text{ s.t. } \sum_{i=1}^{n} x_i = W_0, \text{ Prob}\left\{\sum_{i=1}^{n} \xi_i x_i \geq b\right\} \geq 1-\alpha,$$

where $\xi = (\xi_1, ..., \xi_n)$ is vector of random returns and $\mu_i = \mathbb{E}[\xi_i]$. If random vector ξ has a multivariate normal distribution, $\xi \sim N(\mu, \Sigma)$, then $W_1 = \sum_{i=1}^n \xi_i x_i \sim N(\sum_{i=1}^n \mu_i x_i, x^{\mathsf{T}}\Sigma x)$ and

$$\text{Prob}\{W_1 \geq b\} = \text{Prob}\left\{Z \geq \frac{b - \sum_{i=1}^n \mu_i x_i}{\sqrt{x^{\mathsf{T}}\Sigma x}}\right\} = \Phi\left(\frac{\sum_{i=1}^n \mu_i x_i - b}{\sqrt{x^{\mathsf{T}}\Sigma x}}\right),$$

where $Z \sim N(0,1)$ and $\Phi(z) = \text{Prob}(Z \leq z)$. Therefore we can write the chance constraint in the form

$$b - \sum_{i=1}^{n} \mu_i x_i + z_\alpha \sqrt{x^{\mathsf{T}}\Sigma x} \leq 0,$$

where $z_\alpha := \Phi^{-1}(1-\alpha)$ (if $0 < \alpha \leq 1/2$).

There is a serious numerical problem with chance constraints. First, it is usually difficult even to check whether or not a given chance constraint is satisfied at a given point $x \in X$. Second, the feasible set of a chance constraint is convex only in very special cases. For example, the set

$$\text{Prob}\Big\{C_i(x,\xi) \leq 0,\ i = 1,...,q\Big\} \geq 1-\alpha,$$

is convex if $C_i(x,\xi)$ are convex (jointly in x and ξ) and ξ has an α-concave distribution (Prékopa).

Two approaches to deal with chance constraints: sampling and convex approximations.

Generate a random sample $\xi^1, \ldots, \xi^N$ of N realizations of random vector ξ (by Monte Carlo sampling techniques) and consider problem

$$\operatorname*{Min}_{x \in X} f(x) \text{ subject to } C(x, \xi^j) \leq 0, \;\; j = 1, ..., N. \tag{55}$$

If the set X and functions $f(\cdot)$, $C(\cdot, \xi)$, $\xi \in \Xi$, are convex, then this is a convex problem.

Theorem 1 (Calafiore, Campi, Garatti) *Suppose that the convexity condition holds and let $\bar{x}_N$ be an optimal solution of the above problem* (55)*. Then*

$$\text{Prob}\,\{p(\bar{x}_N) > \alpha\} \leq B(n-1; \alpha, N),$$

where

$$B(k; \alpha, N) := \sum_{i=0}^{k} \binom{N}{i} \alpha^i (1-\alpha)^{N-i}, \;\; k = 0, ..., N.$$

By Chernoff inequality

$$B(n-1;\alpha,N) \leq \exp\left\{-\frac{(N\alpha - n + 1)^2}{2\alpha N}\right\}.$$

It follows that for $\beta \in (0,1)$ and

$$N \geq 2\alpha^{-1}\log(1/\beta),$$

we are guaranteed with probability at least $1-\beta$ that $\bar{x}_N$ is a *feasible* point of the true problem. This result only ensures feasibility of $\bar{x}_N$, doesn't say anything about optimality.

Note that $p(x) = \mathbb{E}_P[\mathbb{1}_{(0,\infty)}(C(x,\xi))]$. The corresponding sample average approximation:

$$\widehat{p}_N(x) = \frac{1}{N}\sum_{j=1}^{N} \mathbb{1}_{(0,\infty)}\left(C(x,\xi^j)\right)$$

is equal to the proportion of times that $C(x,\xi^j) > 0$. The SAA problem

$$\operatorname*{Min}_{x\in X} f(x) \text{ s.t. } \widehat{p}_N(x) \leq \gamma. \tag{56}$$

Note that we can use the significance level γ, of the SAA problem (56), different from α.

If $\gamma = \alpha$, then under mild regularity conditions, an optimal solution $\widehat{x}_N$ of the SAA problem converges w.p.1 to the set of optimal solutions of the true problem.

For a point $\bar{x} \in X$ we have that $\widehat{p}_N(\bar{x}) \leq \gamma$, i.e., $\bar{x}$ is a feasible point of the SAA problem, iff no more than γN times the event "$C(\bar{x}, \xi^j) > 0$" happens in N trials. Since probability of the event "$C(\bar{x}, \xi^j) > 0$" is $p(\bar{x})$, it follows that

$$\text{Prob}\Big\{\widehat{p}_N(\bar{x}) \leq \gamma\Big\} = B\Big(\lfloor \gamma N \rfloor; p(\bar{x}), N\Big).$$

By Chernoff inequality for $k > Np$,

$$B(k; p, N) \geq 1 - \exp\Big\{-N(k/N - p)^2/(2p)\Big\}.$$

It follows that if $p(\bar{x}) \leq \alpha$ and $\gamma > \alpha$, then $1 - \text{Prob}\Big\{\widehat{p}_N(\bar{x}) \leq \gamma\Big\}$ approaches zero at a rate of $\exp(-\kappa N)$, where $\kappa := (\gamma - \alpha)^2/(2\alpha)$.

Similarly, if $p(\bar{x}) = \alpha$ and $\gamma < \alpha$, then probability that $\bar{x}$ is a feasible point of the corresponding SAA problem approaches zero exponentially fast.

Optimality bounds

Given a point candidate solution $\bar{x} \in X$ how to verify its optimality. In order to verify feasibility of $\bar{x}$ we need to estimate the probability $p(\bar{x})$. By Monte Carlo sampling techniques, generate an iid sample $\xi^1, ..., \xi^N$ and estimate $p(\bar{x})$ by $\widehat{p}_N(\bar{x})$. Approximate $(1-\beta)$-confidence upper bound on $p(\bar{x})$:

$$U_{\beta,N}(\bar{x}) := \widehat{p}_N(\bar{x}) + z_\beta \sqrt{\widehat{p}_N(\bar{x})(1-\widehat{p}_N(\bar{x}))/N}.$$

A more accurate $(1-\beta)$-confidence upper bound is given by

$$U^*_{\beta,N}(\bar{x}) := \sup_{\rho \in [0,1]} \{\rho : B(k; \rho, N) \geq \beta\},$$

where $k := N\widehat{p}_N(\bar{x}) = \sum_{j=1}^N \mathbb{1}_{(0,\infty)}\left(G(\bar{x}, \xi^j)\right)$.

Lower bound for the optimal value

Choose two positive integers M and N, and let L be the *largest* integer such that

$$B(L-1;\theta_N,M) \leq \beta,$$

where $\theta_N := B\Big(\lfloor \gamma N \rfloor; \alpha, N\Big)$. Note that $\theta_N = (1-\alpha)^N$ for $\gamma = 0$. Next generate M independent samples $\xi^{1,m},\ldots,\xi^{N,m}$, $m = 1,\ldots,M$, each of size N, of random vector ξ. For each sample solve the associated optimization problem

$$\operatorname*{Min}_{x\in X} f(x) \;\text{ subject to }\; \sum_{j=1}^{N} \mathbb{1}_{(0,\infty)}\Big(C(x,\xi^{j,m})\Big) \leq \gamma N,$$

and hence calculate its optimal value $\widehat{\vartheta}_N^m$, $m = 1,\ldots,M$. That is, solve M times the corresponding SAA problem at the significance level γ.

We can view $\widehat{\vartheta}_N^m$, $m = 1, \ldots, M$, as an iid sample of the random variable $\widehat{\vartheta}_N$, where $\widehat{\vartheta}_N$ is the optimal value of the respective SAA problem at significance level γ. Next we rearrange the calculated optimal values in the nondecreasing order $\widehat{\vartheta}_N^{(1)} \leq \cdots \leq \widehat{\vartheta}_N^{(M)}$. We use the random quantity $\widehat{\vartheta}_N^{(L)}$ as a lower bound of the true optimal value ϑ^*. It is possible to show that with probability at least $1 - \beta$, the random quantity $\widehat{\vartheta}_N^{(L)}$ is below the true optimal value ϑ^*, i.e., $\widehat{\vartheta}_N^{(L)}$ is indeed a lower bound of the true optimal value with confidence at least $1 - \beta$.

Convex approximations

Consider chance constraint:

$$\text{Prob}\Big\{C(x,\xi) > \tau\Big\} \leq \alpha. \tag{57}$$

Let $Z_x = C(x,\xi) - \tau$ and $\psi : \mathbb{R} \to \mathbb{R}_+$ be nondecreasing, convex function such that $\psi(\cdot) \geq \mathbb{1}_{(0,\infty)}(\cdot)$. We have that

$$\text{Prob}\Big\{C(x,\xi) > \tau\Big\} = \mathbb{E}\left[\mathbb{1}_{(0,\infty)}(Z_x)\right]$$

and

$$\inf_{t>0} \mathbb{E}[\psi(tZ_x)] \geq \mathbb{E}\left[\mathbb{1}_{(0,\infty)}(Z_x)\right],$$

and hence

$$\inf_{t>0} \mathbb{E}[\psi(tZ_x)] \leq \alpha \tag{58}$$

is a conservative approximation of the chance constraint (57).

The choice $\psi(z) := [1+z]_+$ gives best conservative approximation. For this choice of ψ, (58) is equivalent to

$$\underbrace{\inf_{t\in\mathbb{R}} \left\{ t + \alpha^{-1}\mathbb{E}[Z_x - t]_+ \right\}}_{\mathsf{CVaR}_\alpha(Z_x)} \leq 0. \tag{59}$$

Note that the minimum in the left hand side of (59) is attained at $t^* = \mathsf{VaR}_{1-\alpha}(Z_x)$, where

$$\mathsf{VaR}_{1-\alpha}(Z) = F_Z^{-1}(1-\alpha) := \inf\{t : F_Z(t) \geq 1-\alpha\},$$

with $F_Z(t) := \mathsf{Prob}(Z \leq t)$ being the cdf of Z.