Stochastic Programming

A. Shapiro

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

Modeling and Basic Properties

Consider the optimization problem

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

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

Robust (worst case) approach:

$$\begin{array}{ll} \mathrm{Min}_{x\in\mathcal{X}} & \left\{f(x):=\mathrm{max}_{\xi\in\Xi}F(x,\xi)\right\} \\ \text{subject to} & c_i(x,\xi)\leq 0,\ i=1,...,q,\ \xi\in\Xi. \end{array} \tag{2}$$

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

$$F(x,d) = cx + b[d-x]_+ + h[x-d]_+. \tag{3}$$

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 $F(x,d)$.

*Recall that $[a]_+ = \max\{0, a\}$.

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

$$\operatorname*{Min}_{x \geq 0} \Big\{ f(x) := \mathbb{E}[F(x, D)] \Big\}. \tag{4}$$

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 (4) 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 (4) is equal to the quantile

$$\bar{x} = H^{-1}(\kappa), \quad \text{with} \quad \kappa = \frac{b-c}{b+h}. \tag{5}$$

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

$$\operatorname*{Min}_{x \geq 0} \left\{ f(x) := \max_{d \in [\ell, u]} F(x, d) \right\}. \tag{6}$$

Clearly we have to look at $x \in [\ell, u]$, in which case

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

and hence $x^* = \frac{h\ell + bu}{h+b}$ is the optimal solution of problem (6). This solution is quite different from solution "on average" of problem (4). In particular, if the holding cost $h = 0$, then $x^* = u$.

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,

$$W_0 = \sum_{i=1}^{n} x_i. \tag{7}$$

After one period of time our wealth becomes

$$W_1 = \sum_{i=1}^{n} \xi_i x_i, \tag{8}$$

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 (7) and "no borrowing" constraints $x_i \geq 0$. This leads to the optimization problem

$$\text{Max}_{x \geq 0} \ \mathbb{E}[W_1] \quad \text{subject to } \sum_{i=1}^n x_i = W_0. \tag{9}$$

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 (9) 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 $q(W_1 - \tau)$, and if W_1 exceeds τ we are rewarded by $r(W_1 - \tau)$, with $q > r$. This leads to the concept of utility function

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

and to the optimization problem

$$\text{Max}_{x \geq 0} \ \mathbb{E}\left[F(x,\xi)\right] \quad \text{subject to } \sum_{i=1}^{n} x_i = W_0, \tag{10}$$

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

Chance (probabilistic) constraints formulation

$$\begin{array}{ll} \mathsf{Max}_{x \geq 0} & \mu^{\mathsf{T}} x \\ \text{subject to} & \mathsf{Prob}\left(R^{\mathsf{T}} x < -b\right) \leq \alpha,\ \sum_{i=1}^{n} x_i = W_0, \end{array} \tag{11}$$

where $R^{\mathsf{T}} x = \sum_{i=1}^{n} R_i x_i$, $\mu = \mathbb{E}[R]$ 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 $N(\mu, \Sigma)$, then $R^{\mathsf{T}} x \sim N(\mu^{\mathsf{T}} x, x^{\mathsf{T}} \Sigma x)$ and the probabilistic constraint is equivalent to:

$$b + \mu^{\mathsf{T}} x - z_\alpha (x^{\mathsf{T}} \Sigma\, x)^{1/2} \geq 0, \tag{12}$$

where z_α is the $(1 - \alpha)$-quantile of the standard normal distribution. Note that $z_\alpha > 0$ and the left hand side of (12) is a concave function of x, provided that $\alpha \in (0, 1/2)$.

By convex duality, there exists $\lambda \geq 0$ such that problem (11) is equivalent to the problem

$$\begin{array}{ll} \mathrm{Max}_{x \geq 0} & \mu^{\mathsf{T}} x - \lambda (x^{\mathsf{T}} \Sigma\, x)^{1/2} \\ \text{subject to} & \sum_{i=1}^{n} x_i = W_0. \end{array} \tag{13}$$

The above problem can be viewed as a compromise between optimizing (maximizing) the expected return $\mu^{\mathsf{T}} x$ 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):

$$\begin{array}{ll} \mathrm{Max}_{x \geq 0} & \overbrace{\mathbb{E}[R^{\mathsf{T}} x]}^{\mu^{\mathsf{T}} x} - \lambda \overbrace{\mathrm{Var}[R^{\mathsf{T}} x]}^{x^{\mathsf{T}} \Sigma\, x} \\ \text{subject to} & \sum_{i=1}^{n} x_i = W_0. \end{array} \tag{14}$$

Equivalent formulations:

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

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

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

$$\begin{array}{ll} \text{Max}_{x\geq 0} & \sum_{i=1}^{n} \mu_i x_i \\ \text{s.t.} & \sum_{i=1}^{n} x_i = W_0,\ (x^{\mathsf{T}} \Sigma x)^{1/2} \leq \gamma^{1/2}. \end{array}$$

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

$$\begin{aligned} \underset{x_t \geq y_t}{\text{Min}} \ & \sum_{t=1}^{T} \mathbb{E}\Big\{c_t(x_t - y_t) + b_t[D_t - x_t]_+ + h_t[x_t - D_t]_+\Big\} \\ \text{s.t. } & y_{t+1} = x_t - D_t, \ t = 1, ..., T-1. \end{aligned} \tag{15}$$

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:

$$\begin{aligned} \underset{x_T \geq y_T}{\text{Min}} \; c_T(x_T - y_T) \; + \; & \mathbb{E}\Big\{ b_T[D_T - x_T]_+ \\ & + h_T[x_T - D_T]_+ \Big| D_{[T-1]} = d_{[T-1]} \Big\}. \end{aligned} \tag{16}$$

The expectation in (16) 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 (16) 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

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

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 (15) as the problem of minimization of the expectation in (15) 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

$$c_t x_t + \mathbb{E}\Big\{b_t[D_t - x_t]_+ + h_t[x_t - D_t]_+ + Q_{t+1}\left(x_t - D_t\right)\Big\}, \qquad (17)$$

$t = T, ..., 1$. By using convexity of the value functions $Q_t(\cdot)$ 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

$$\sum_{i=1}^{n} x_{i0} = W_0. \tag{18}$$

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

$$W_t = \sum_{i=1}^{n} \xi_{it} x_{i,t-1}, \tag{19}$$

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

$$\text{Max}\, \mathbb{E}\left[U(W_T)\right] \tag{20}$$

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

$$\begin{aligned} \underset{x_t \geq 0, W_{t+1}}{\text{Max}} \quad & \mathbb{E}\Big\{ Q_{t+1}(W_{t+1}, \xi_{[t+1]}) \Big| \xi_{[t]} \Big\} \\ \text{s.t. } & W_{t+1} = \sum_{i=1}^{n} \xi_{i,t+1} x_{i,t}, \ \sum_{i=1}^{n} x_{i,t} = W_t. \end{aligned} \tag{21}$$

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 $\xi_{[t]}$.

Consider the logarithmic utility function $U(W) := \ln W$. At stage $t = T - 1$ we solve the problem

$$\underset{x_{T-1} \geq 0}{\text{Max}} \ \mathbb{E}\left\{ \ln\left(\sum_{i=1}^{n} \xi_{i,T} x_{i,T-1} \right) \Big| \xi_{[T-1]} \right\} \ \text{s.t.} \ \sum_{i=1}^{n} x_{i,T-1} = W_{T-1}. \tag{22}$$

Its optimal value is

$$Q_{T-1}\left(W_{T-1}, \xi_{[T-1]}\right) = \nu_{T-1}\left(\xi_{[T-1]}\right) + \ln W_{T-1},$$

where $\nu_{T-1}\left(\xi_{[T-1]}\right)$ denotes the optimal value of (22) for $W_{T-1} = 1$. At stage $t = T - 2$ we solve problem

$$\begin{aligned} \underset{x_{T-2} \geq 0}{\text{Max}} \ & \mathbb{E}\left\{ \nu_{T-1}\left(\xi_{[T-1]}\right) + \ln\left(\sum_{i=1}^{n} \xi_{i,T-1} x_{i,T-2} \right) \Big| \xi_{[T-2]} \right\} \\ \text{s.t.} \ & \sum_{i=1}^{n} x_{i,T-2} = W_{T-2}. \end{aligned} \tag{23}$$

Of course, we have that

$$\mathbb{E}\left\{\nu_{T-1}\left(\xi_{[T-1]}\right)+\ln\left(\sum_{i=1}^{n}\xi_{i,T-1}x_{i,T-2}\right)\Big|\xi_{[T-2]}\right\}$$

$$=\mathbb{E}\left\{\nu_{T-1}\left(\xi_{[T-1]}\right)\Big|\xi_{[T-2]}\right\}+\mathbb{E}\left\{\ln\left(\sum_{i=1}^{n}\xi_{i,T-1}x_{i,T-2}\right)\Big|\xi_{[T-2]}\right\},$$

and hence the optimal value of (23) can be written as

$$\begin{aligned}Q_{T-2}\left(W_{T-2},\xi_{[T-2]}\right)= \;& \mathbb{E}\left\{\nu_{T-1}\left(\xi_{[T-1]}\right)\Big|\xi_{[T-2]}\right\}\\ &+\nu_{T-2}\left(\xi_{[T-2]}\right)+\ln W_{T-2},\end{aligned}$$

where $\nu_{T-2}\left(\xi_{[T-2]}\right)$ is the optimal value of the problem

$$\underset{x_{T-2}\geq 0}{\text{Max}}\;\mathbb{E}\left\{\ln\left(\sum_{i=1}^{n}\xi_{i,T-1}x_{i,T-2}\right)\Big|\xi_{[T-2]}\right\}\;\text{ s.t. }\;\sum_{i=1}^{n}x_{i,T-2}=1.$$

An identical argument applies at earlier stages. Therefore, it suffices to solve at each stage $t = T-1, ..., 1, 0$, the corresponding optimization problem

$$\underset{x_t \geq 0}{\text{Max}} \ \mathbb{E}\left\{ \ln\left(\sum_{i=1}^{n} \xi_{i,t+1} x_{i,t} \right) \Big| \xi_{[t]} \right\} \ \text{s.t.} \ \sum_{i=1}^{n} x_{i,t} = W_t, \qquad (24)$$

in a completely myopic fashion.

By definition, we set ξ_0 to be constant, so that for the first stage problem, at $t = 0$, the corresponding expectation is unconditional. An optimal solution $\bar{x}_t = \bar{x}_t(W_t, \xi_{[t]})$ of problem (24) gives an optimal policy.

If the random process ξ_t is stagewise independent, then conditional expectations in (24) are the same as the corresponding unconditional expectations, and hence optimal values $\nu_t(\xi_{[t]}) = \nu_t$ do not depend on $\xi_{[t]}$ and are given by the optimal value of the problem

$$\underset{x_t \geq 0}{\text{Max}} \ \mathbb{E}\left\{\ln\left(\sum_{i=1}^{n} \xi_{i,t+1} x_{i,t}\right)\right\} \ \text{s.t.} \ \sum_{i=1}^{n} x_{i,t} = 1. \tag{25}$$

Also in the stagewise independent case the optimal policy can be described as follows. Let $x_t^* = (x_{1t}^*, ..., x_{nt}^*)$ be the optimal solution of (25), $t = 0, ..., T-1$. Such optimal solution is unique by strict concavity of the logarithm function. Then

$$\bar{x}_t(W_t) := W_t x_t^*, \ \ t = 0, ..., T-1,$$

defines the optimal policy.

Two and multistage stochastic programming

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

$$\underset{x \in \mathcal{X}}{\text{Min}}\, c^{\mathsf{T}} x + \mathbb{E}[Q(x, \boldsymbol{\xi})], \tag{26}$$

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

$$\underset{y}{\text{Min}}\, q^{\mathsf{T}} y \ \text{ s.t. } \ Tx + Wy = h,\ y \geq 0, \tag{27}$$

with $\xi = (q, T, W, h)$. In general, the feasible set $\mathcal{X}$ 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 $p_i > 0$, $i = 1, ..., K$. 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 (26)-(27) as one large linear program:

$$\begin{array}{ll} \underset{x, y_1, ..., y_K}{\text{Min}} & c^{\mathsf{T}} x + \sum_{k=1}^{K} p_k q_k^{\mathsf{T}} y_k \\ \text{subject to} & Ax = b, \\ & T_k x + W_k y_k = h_k,\ k = 1, ..., K, \\ & x \geq 0,\ y_k \geq 0,\ k = 1, ..., K. \end{array} \tag{28}$$

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

$$Q(x, \xi) = \sup \left\{ \pi^{\mathsf{T}}(h - Tx) : W^{\mathsf{T}}\pi \leq q \right\}. \tag{29}$$

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 (27) is empty. In that case, by the definition, $Q(x, \xi) = +\infty$. It also can happen that problem (27) 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

$$\partial Q(x, \xi) = -T^{\mathsf{T}} \mathfrak{D}(x, \xi), \tag{30}$$

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

$$\partial \mathcal{Q}(x) = \sum_{k=1}^{K} p_k \partial Q(x, \xi_k). \tag{31}$$

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

$$\partial \mathcal{Q}(x) = \int_{\Omega} \partial Q(x, \omega) dP(\omega). \tag{32}$$

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 (32) 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 (26) as follows: a feasible point $\bar{x} \in \mathcal{X}$ is an optimal solution of (26) iff

$$0 \in \partial \mathcal{Q}(\bar{x}) + N_{\mathcal{X}}(\bar{x}), \tag{33}$$

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

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

General formulation of two-stage stochastic programming problems

$$\operatorname*{Min}_{x\in\mathcal{X}} \Big\{ f(x) := \mathbb{E}[F(x,\omega)] \Big\}, \tag{34}$$

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

$$\operatorname*{Min}_{y\in\mathfrak{X}(x,\omega)} g(x,y,\omega). \tag{35}$$

Here $(\Omega,\mathcal{F},P)$ is a probability space, $\mathcal{X}\subset\mathbb{R}^n$, $g:\mathbb{R}^n\times\mathbb{R}^m\times\Omega\to\mathbb{R}$ and $\mathfrak{X}:\mathbb{R}^n\times\Omega\rightrightarrows\mathbb{R}^m$ 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

$$\mathfrak{X}(x,\omega) := \{y : T(\omega)x + W(\omega)y = h(\omega),\ y\geq 0\}. \tag{36}$$

The second stage problem (35) can be also written in the following equivalent form

$$\operatorname*{Min}_{y\in\mathbb{R}^m} \bar{g}(x,y,\omega), \tag{37}$$

where

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

By the interchangeability principle we have

$$\underbrace{\mathbb{E}\left[\inf_{y\in\mathbb{R}^m} \bar{g}(x,y,\omega)\right]}_{F(x,\omega)} = \inf_{\boldsymbol{y}\in\mathfrak{Y}} \mathbb{E}\Big[\bar{g}(x,\boldsymbol{y}(\omega),\omega)\Big], \tag{38}$$

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 (34)–(35) as one large problem:

$$\begin{aligned} \underset{x\in\mathbb{R}^n,\boldsymbol{y}\in\mathfrak{Y}}{\text{Min}} \quad & \mathbb{E}\left[g(x,\boldsymbol{y}(\omega),\omega)\right] \\ \text{s.t. } & x\in\mathcal{X},\ \boldsymbol{y}(\omega)\in\mathfrak{X}(x,\omega) \text{ a.e. } \omega\in\Omega. \end{aligned} \tag{39}$$

In particular, if $\Omega = \{\omega_1,...,\omega_K\}$ is finite, then problem (39) becomes

$$\begin{aligned} \underset{x,y_1,...,y_k}{\text{Min}} \quad & \sum_{k=1}^{K} g(x,y_k,\omega_k) \\ \text{s.t. } & x\in\mathcal{X},\ y_k\in\mathfrak{X}(x,\omega_k)\ k=1,...,K. \end{aligned} \tag{40}$$

Decision rules

Recall that the two stage problem (26) can be formulated as an optimization problem over $x \in \mathcal{X}$ and second stage decision $y(\cdot)$ considered as a function of the data. Suppose now that the recourse is fixed, i.e., only $T = T(\omega)$ and $h = h(\omega)$ of the second stage problem

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

are random. For a given $x \in \mathcal{X}$ the second stage problem attains its minimal value at an extreme point of the set $\{y : Wy = h - Tx, \ y \geq 0\}$, assuming that this set is nonempty and bounded. By the well known result of linear programming we have the following characterization of the extreme points

A feasible point $\bar{y}$ is an extreme point (basic solution) of the feasible set if there exists an index set $I \subset \{1, ..., m\}$ such that the corresponding submatrix W_I is nonsingular and $\bar{y}_i = 0$ for $i \in \{1, ..., m\} \setminus I$. That is, $W_I \bar{y}_I = h - Tx$, where $\bar{y}_I$ is the corresponding subvector of $\bar{y}$. This can be written as $\bar{y} = R_I(h - Tx)$, where the I-th submatrix of R_I is W_I^{-1} and other entries of R_I are zeros. It follows (under appropriate nondegeneracy conditions) that an optimal policy (optimal decision rule) $\bar{y} = \bar{y}(h, T)$ is a piecewise linear function of the data.

Nonanticipativity

Consider the first stage problem (34). 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 (34):

$$\operatorname*{Min}_{x_1,\ldots,x_K} \sum_{k=1}^{K} p_k F(x_k, \omega_k) \text{ subject to } x_k \in \mathcal{X},\ k = 1, \ldots, K. \qquad (41)$$

Problem (41) is separable in the sense that it can be split into K smaller problems, one for each scenario:

$$\operatorname*{Min}_{x_k \in \mathcal{X}} F(x_k, \omega_k), \quad k = 1, \ldots, K, \qquad (42)$$

and that the optimal value of problem (41) is equal to the weighted sum, with weights p_k, of the optimal values of problems (42), $k = 1, \ldots, K$.

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 (34) (with a finite number of scenarios),

$$\underset{x_1,\ldots,x_K,z}{\mathsf{Min}} \sum_{k=1}^{K} p_k \bar{F}(x_k, \omega_k) \text{ s.t. } x_k = z,\ k = 1, ..., K, \tag{43}$$

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 (43) is:

$$\begin{array}{ll} \mathsf{Max}_{\lambda_1,\ldots,\lambda_K} & \left\{\mathsf{inf}_{x_1,\ldots,x_K} \sum_{k=1}^{K} p_k \left(\bar{F}(x_k, \omega_k) + \lambda_k^{\mathsf{T}} x_k\right)\right\}, \\ \text{subject to} & \sum_{k=1}^{K} p_k \lambda_k = 0. \end{array} \tag{44}$$

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 (43) and (44), 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 (43) and (44), respectively, then $\bar{x}_1 = ... = \bar{x}_K = \bar{z}$ and

$$\bar{x}_k \in \arg\min_{x_k} \left\{ \bar{F}(x_k, \omega_k) + \bar{\lambda}_k^{\mathsf{T}} x_k \right\}. \tag{45}$$

Multistage stochastic programming

Consider a multistage decision process of the from

$$\begin{array}{c}\text{decision } (x_1) \rightsquigarrow \text{observation } (\xi_2) \rightsquigarrow \text{decision } (x_2) \rightsquigarrow \\ \ldots\ldots \rightsquigarrow \text{observation } (\xi_T) \rightsquigarrow \text{decision } (x_T).\end{array} \tag{46}$$

Here $\xi_t \in \mathbb{R}^{d_t}$, $t = 1, ...,$ is a sequence of vectors with $\xi_{[t]} := (\xi_1, ..., \xi_t)$ representing history of the data process up to time t. At time period $t \in \{1, ..., T\}$ we observe the past, $\xi_{[t]}$, but future observations $\xi_{t+1}, ...,$ are uncertain. So our decision at time t should only depend on information available at that time, i.e., $x_t = x_t(\xi_{[t]})$ should be a function of $\xi_{[t]}$ and should not depend on future observations. This is the basic requirement of *nonanticipativity* of the decision process. A sequence $x_1, x_2(\xi_{[2]}), ...$ of such decisions is called a *policy* or a decision rule.

Risk neutral multistage stochastic programming

$$\operatorname*{Min}_{\pi\in\Pi} \mathbb{E}\Big[f_1(x_1) + f_2(x_2(\xi_{[2]}),\xi_2) + \cdots + f_T\left(x_T(\xi_{[T]}),\xi_T\right)\Big], \quad (47)$$

where Π is the set of policies $\pi = (x_1, x_2(\xi_{[2]}), ..., x_T(\xi_{[T]}))$ satisfying the feasibility constraints

$$x_1 \in \mathcal{X}_1,\ x_t(\xi_{[t]}) \in \mathcal{X}_t(x_{t-1}(\xi_{[t-1]}),\xi_t),\ t = 2,\ldots,T,$$

for almost every (a.e.) realization of the random process, $f_t : \mathbb{R}^{n_t} \times \mathbb{R}^{d_t} \to \mathbb{R}$ are real valued functions and $\mathcal{X}_t : \mathbb{R}^{n_{t-1}} \times \mathbb{R}^{d_t} \rightrightarrows \mathbb{R}^{n_t}$, $t = 2,\ldots,T$, are multifunctions. For example

$$f_t(x_t,\xi_t) := c_t^{\mathsf{T}} x_t,$$

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

$t = 2, ..., T$, $\mathcal{X}_1 := \{x_1 : A_1 x_1 \le b_1\}$, with $\xi_t = (c_t, A_t, B_t, b_t)$, corresponds to linear multistage stochastic programming.

Note that it is assumed here that the probability distribution of the random process ξ_t does not depend on our decisions. Note also that

$$\mathbb{E}[Z] = \mathbb{E}_{|\xi_1}\Big[\mathbb{E}_{|\xi_{[2]}}\big[\cdots \mathbb{E}_{|\xi_{[T-1]}}[Z]\big]\Big],$$

where $\mathbb{E}_{|\xi_{[t]}}$ denotes conditional expectation. This decomposition property of the expectation allows to write the multistage stochastic programming problem in the following nested form

$$\begin{aligned}\min_{x_1\in\mathcal{X}_1} \; & f_1(x_1) + \mathbb{E}_{|\xi_1}\Big[\inf_{x_2\in\mathcal{X}_2(x_1,\xi_2)} f_2(x_2,\xi_2) + \ldots \\ & + \mathbb{E}_{|\xi_{[T-2]}}\big[\inf_{x_{T-1}\in\mathcal{X}_T(x_{T-2},\xi_{T-1})} f_{T-1}(x_{T-1},\xi_{T-1}) \\ & + \mathbb{E}_{|\xi_{[T-1]}}[\inf_{x_T\in\mathcal{X}_T(x_{T-1},\xi_T)} f_T(x_T,\xi_T)]\big]\Big].\end{aligned}$$

This formulation assumes that: (i) the probability distribution of the data process is known (specified), (ii) the optimization is performed *on average* (both, with respect to different realizations of the random process, and with respect to time).

In the linear case the nested formulation can be written as (recall that ξ_1 is deterministic and hence $\mathbb{E}_{|\xi_1} = \mathbb{E}$)

$$\underset{\substack{A_1x_1=b_1\\ x_1\geq 0}}{\text{Min}} c_1^\top x_1 + \mathbb{E}\left[\underset{\substack{B_2x_1+A_2x_2=b_2\\ x_2\geq 0}}{\text{Min}} c_2^\top x_2 + \cdots + \mathbb{E}_{|\xi_{[T-1]}}\left[\underset{\substack{B_Tx_{T-1}+A_Tx_T=b_T\\ x_T\geq 0}}{\text{Min}} c_T^\top x_T\right]\right]$$

If the number of realizations (scenarios) of the process ξ_t is finite, then the above problem can be written as one large linear programming problem.

Numerical difficulties in solving multistage problems.

From a modeling point of view typically it is natural to assume that the random data process has a *continuous* distribution. This raises the question of how to compute the involved expectations (multivariate integrals). A standard approach is to discretize the random process by generating a finite number of possible realizations (called scenarios). These scenarios can be represented by the corresponding *scenario tree*.

How many scenarios are needed in order to approximate the "true" distribution of the random data process?

Note that solving the deterministic equivalent for the constructed scenario tree does not produce by itself an implementable policy for the "true" problem (with continuous distributions). This is because an actual realization of the data process could, and with probability one (w.p.1) will, be different from scenarios used in the constructed tree. In that case policy constructed for scenarios of the tree does not tell what decision to make. Of course, one can use only the first stage solution which is deterministic (does not depend on future observations) and update it as new observations become available - this is a rolling horizon approach. Such a rolling horizon approach requires resolving the corresponding multistage problem at every stage as new realization of the data becomes available.

Dynamic programming equations

Consider the last stage problem

$$\underset{x_T \in \mathcal{X}_T(x_{T-1}, \xi_T)}{\text{Min}} f_T(x_T, \xi_T). \tag{48}$$

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:

$$\begin{array}{ll} \underset{x_t}{\text{Min}} & f_t(x_t, \xi_t) + \mathbb{E}_{|\xi_{[t]}} \left[Q_{t+1}\left(x_t, \xi_{[t+1]}\right) \right] \\ \text{s.t.} & x_t \in \mathcal{X}_t(x_{t-1}, \xi_t). \end{array} \tag{49}$$

Its optimal value depends on the decision x_{t-1} at the previous stage and realization of the data process $\xi_{[t]} = (\xi_1, ..., \xi_t)$, 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:

$$\underset{x_1 \in \mathcal{X}}{\text{Min}} \, f_1(x_1) + \mathbb{E}\left[Q_2\left(x_t, \xi_2\right)\right]. \tag{50}$$

The dynamic programming equations:

$$Q_t\left(x_{t-1}, \xi_{[t]}\right) = \inf_{x_t \in \mathcal{X}_t(x_{t-1}, \xi_t)} \left\{ f_t(x_t, \xi_t) + \mathcal{Q}_{t+1}\left(x_t, \xi_{[t]}\right) \right\}, \tag{51}$$

where

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

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 feasible if it satisfies the feasibility constraints, i.e.,

$$x_t(\xi_{[t]}) \in \mathcal{X}_t(x_{t-1}(\xi_{[t-1]}), \xi_t), \ t = 2, ..., T, \text{ w.p.1.} \qquad (52)$$

A policy $\bar{x}_t(\xi_{[t]})$ is *optimal* if and only if it satisfies the dynamic programming equations, i.e.,

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

In the *stagewise independent* case $\mathcal{Q}_{t+1}(x_t)$ does not depend on $\xi_{[t]}$, and optimal solution $\bar{x}_t$ (satisfying the dynamic programming equations) is a function of $\bar{x}_{t-1}$ and ξ_t.

Nonanticipativity constraints

Consider the multistage problem (47). Suppose that there is a finite number of scenarios $\xi_1^k, ..., \xi_T^k$, $k = 1, ..., K$, with respective probabilities p_k. To each scenario $k \in \{1, ..., K\}$ corresponds a sequence $x_1^k, ..., x_T^k$ of decision vectors. The *nonanticipativity* principle requires that

$$x_t^k = x_t^\ell \text{ for all } k, \ell \text{ for which } \xi_{[t]}^k = \xi_{[t]}^\ell, \quad t = 1, \ldots, T. \tag{53}$$

Let $\mathfrak{X}$ be the space of all sequences $(x_1^k, ..., x_T^k)$, $k = 1, ..., K$ (this is a linear space of dimension $(n_1 + ... + n_T)K$) and $\mathfrak{L}$ be a linear subspace of $\mathfrak{X}$ defined by the nonanticipativity constraints (53). We can write the multistage problem in the form

$$\operatorname*{Min}_{\boldsymbol{x} \in \mathfrak{X}} \left\{ f(\boldsymbol{x}) := \sum_{k=1}^{K} \sum_{t=1}^{T} p_k f_t(x_t^k, \xi_t^k) \right\} \text{ s.t. } \boldsymbol{x} \in \mathfrak{L}. \tag{54}$$

In particular, the relaxed version of the linear multistage program

$$
\begin{array}{lccccccccc}
\text{Min} & \sum_{k=1}^{K} p_k \Big[c_1^{\mathsf{T}} x_1^k & + & (c_2^k)^{\mathsf{T}} x_2^k & + & (c_3^k)^{\mathsf{T}} x_3^k + & \dots & + & (c_T^k)^{\mathsf{T}} x_T^k \Big] & \\
\text{s.t.} & A_1 x_1^k & & & & & & & = & b_1, \\
& B_2^k x_1^k & + & A_2^k x_2^k & & & & & = & b_2^k, \\
& & & B_3^k x_2^k & + & A_3^k x_3^k & & & = & b_3^k, \\
& \dots & \dots & \dots & \dots & \dots & \dots & \dots & \dots & \dots \\
& & & & & & B_T^k x_{T-1}^k & + & A_T^k x_T^k = & b_T^k, \\
& x_1^k \geq 0, & & x_2^k \geq 0, & & x_3^k \geq 0, & \dots & & x_T^k \geq 0, & \\
& k = 1, \dots, K. & & & & & & & &
\end{array}
$$

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. Together with the nonanticipativity constraints (53) the considered problem becomes equivalent to the original formulation.

Monte Carlo sampling methods

Consider (two-stage) stochastic programming problem:

$$\operatorname*{Min}_{x \in \mathcal{X}} \Big\{ f(x) := \mathbb{E}[F(x, \xi)] \Big\}. \tag{55}$$

Even a crude discretization of the distribution of the random data vector $\xi \in \mathbb{R}^d$ typically results in an exponential growth of the number of scenarios with increase of the number d of random variables. For example, if components of random vector ξ are independently distributed and distribution of each component is discretized by r points, then the total number of scenarios is r^d. That is, although the input data grows linearly with increase of the dimension d, the number of scenarios grows exponentially. The standard approach to dealing with this issue is to generate a manageable number of scenarios in some "representative" 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

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

Consequently the true (expected value) problem is approximated by the so-called *sample average approximation* (SAA) problem:

$$\operatorname*{Min}_{x \in \mathcal{X}} \widehat{f}_N(x). \tag{56}$$

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 for a given $x \in \mathcal{X}$, by the Law of Large Numbers (LLN), $\widehat{f}_N(x)$ converges to $f(x)$ w.p.1 as N tends to infinity. It is possible to show that, under mild regularity conditions, this convergence is uniform on any compact subset of $\mathcal{X}$ (uniform LLN). It follows that the optimal value $\widehat{v}_N$ and an optimal solution $\widehat{x}_N$ of the SAA problem (56) converge w.p.1 to their counterparts of the true problem.

Central Limit Theorem type results. Notoriously slow convergence of order $O_p(N^{-1/2})$. By the CLT, for a given $x \in \mathcal{X}$,

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

where $\sigma^2(x) := \mathrm{Var}[F(x, \xi)]$ and " $\Rightarrow$ " denotes convergence in distribution.

Delta method

Let $Y_N \in \mathbb{R}^d$ be a sequence of random vectors, converging in probability to a vector $\mu \in \mathbb{R}^d$. Suppose that there exists a sequence $\tau_N \to +\infty$ such that $\tau_N(Y_N - \mu) \Rightarrow Y$. Let $G : \mathbb{R}^d \to \mathbb{R}^m$ be a vector valued function, differentiable at μ. That is $G(y) - G(\mu) = M(y-\mu)+r(y)$, where $M := \nabla G(\mu)$ is the $m \times d$ Jacobian matrix of G at μ, and the remainder $r(y)$ is of order $o(\|y - \mu\|)$. It follows that $\tau_N [G(Y_N) - G(\mu)] \Rightarrow MY$. In particular, suppose that $N^{1/2}(Y_N - \mu)$ converges in distribution to a (multivariate) normal distribution with zero mean vector and covariance matrix Σ. Then it follows that

$$N^{1/2} [G(Y_N) - G(\mu)] \Rightarrow N(0, M\Sigma M^{\mathsf{T}}).$$

Infinite dimensional Delta Theorem. Let B_1 and B_2 be two Banach spaces, and $G : B_1 \to B_2$ be a mapping. It is said that G is directionally differentiable at a point $\mu \in B_1$ if the limit

$$G'_\mu(d) := \lim_{t \downarrow 0} \frac{G(\mu + td) - G(\mu)}{t} \tag{57}$$

exists for all $d \in B_1$. If, in addition, the directional derivative $G'_\mu : B_1 \to B_2$ is linear and continuous, then it is said that G is Gâteaux differentiable at μ. It is said that G is directionally differentiable at μ in the sense of Hadamard if the directional derivative $G'_\mu(d)$ exists for all $d \in B_1$ and, moreover,

$$G'_\mu(d) = \lim_{\substack{t \downarrow 0 \\ d' \to d}} \frac{G(\mu + td') - G(\mu)}{t}. \tag{58}$$

Theorem 1 (Delta Theorem) *Let B_1 and B_2 be Banach spaces, equipped with their Borel σ-algebras, Y_N be a sequence of random elements of B_1, $G : B_1 \to B_2$ be a mapping, and τ_N be a sequence of positive numbers tending to infinity as $N \to \infty$. Suppose that the space B_1 is separable, the mapping G is Hadamard directionally differentiable at a point $\mu \in B_1$, and the sequence $X_N := \tau_N(Y_N - \mu)$ converges in distribution to a random element Y of B_1. Then*

$$\tau_N \left[G(Y_N) - G(\mu)\right] \Rightarrow G'_\mu(Y), \tag{59}$$

and

$$\tau_N \left[G(Y_N) - G(\mu)\right] = G'_\mu(X_N) + o_p(1). \tag{60}$$

Let $\mathcal{X}$ be a compact subset of $\mathbb{R}^n$ and consider the space $B = C(\mathcal{X})$ of continuous functions $\phi : \mathcal{X} \to \mathbb{R}$. Assume that:
(A1) For some point $x \in \mathcal{X}$ the expectation $\mathbb{E}[F(x,\xi)^2]$ is finite.
(A2) There exists a measurable function $C : \Xi \to \mathbb{R}_+$ such that $\mathbb{E}[C(\xi)^2]$ is finite and

$$\left|F(x,\xi) - F(x',\xi)\right| \leq C(\xi)\|x - x'\|, \tag{61}$$

for all $x, x' \in \mathcal{X}$ and a.e. $\xi \in \Xi$.

We can view $Y_N := \widehat{f}_N$ as a random element of $C(\mathcal{X})$. Consider the min-function $V : B \to \mathbb{R}$ defined as $V(Y) := \inf_{x\in\mathcal{X}} Y(x)$. Clearly $\widehat{v}_N = V(Y_N)$. It is not difficult to show that for any $\mu \in C(\mathcal{X})$ and $\mathcal{X}^*(\mu) := \arg\min_{x\in\mathcal{X}} \mu(x)$,

$$V'_\mu(\delta) = \inf_{x\in\mathcal{X}^*(\mu)} \delta(x), \quad \forall \delta \in C(\mathcal{X}),$$

and the above directional derivative holds in the Hadamard sense.

By a functional CLT, under assumptions (A1) and (A2), $N^{1/2}(\widehat{f}_N - f)$ converges in distribution to a random element Y of $C(\mathcal{X})$. In particular, for any finite set $\{x_1, ..., x_m\} \subset \mathcal{X}$, the random vector $(Y(x_1), ..., Y(x_m))$ has a multivariate normal distribution with zero mean and the same covariance matrix as the covariance matrix of $(F(x_1, \xi), ..., F(x_m, \xi))$. In particular, for fixed $x \in \mathcal{X}$, $Y(x) \sim N(0, \sigma^2(x))$ with $\sigma^2(x) := \mathrm{Var}[F(x, \xi)]$.

Denote v^0 the optimal value and S^0 the set of optimal solutions of the true problem.

Theorem 2 *Suppose that the set $\mathcal{X}$ is compact, and assumptions (A1) and (A2) hold. Then*

$$\widehat{v}_N = \min_{x \in S^0} \widehat{f}_N(x) + o_p(N^{-1/2}),$$
$$N^{1/2}[\widehat{v}_N - v^0] \Rightarrow \inf_{x \in S^0} Y(x).$$

In particular, if the optimal set (of the true problem) $S^0 = \{x^0\}$ is a singleton, then

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

This result suggests that the optimal value of the SAA problem converges at a rate of $O_p(N^{-1/2})$. In particular, if $S^0 = \{x^0\}$, then $\widehat{v}_N$ converges to v^0 at the same rate as $\widehat{f}_N(x^0)$ converges to $f(x^0)$.

Validation analysis

How one can evaluate quality of a given (feasible) solution $\widehat{x} \in \mathcal{X}$? The SAA approach – statistical test based on estimation of $f(\widehat{x}) - v^0$, where v^0 is the optimal value of the true problem.
(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

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

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

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

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

where S^0 is the set of optimal solutions of the true problem, $(Y(x_1), ..., Y(x_k))$ has a multivariate normal distribution with zero mean vector and covariance matrix given by the covariance matrix of the random vector $(F(x_1, \xi), ..., F(x_k, \xi))$. 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 $\varepsilon \geq 0$.

Sample size estimates (by Large Deviations type bounds)

Consider an iid sequence $Y_1, \ldots, Y_N$ of replications of a real valued random variable Y, and let $Z_N := N^{-1} \sum_{i=1}^N Y_i$ be the corresponding sample average. Then for any real numbers a and $t > 0$ we have that $\mathrm{Prob}(Z_N \geq a) = \mathrm{Prob}(e^{tZ_N} \geq e^{ta})$, and hence, by Markov inequality

$$\mathrm{Prob}(Z_N \geq a) \leq \quad e^{-ta}\mathbb{E}[e^{tZ_N}] = e^{-ta}[M(t/N)]^N,$$

where $M(t) := \mathbb{E}[e^{tY}]$ is the *moment generating function* of Y. Suppose that Y has finite mean $\mu := \mathbb{E}[Y]$ and let $a \geq \mu$. By taking the logarithm of both sides of the above inequality, changing variables $t' = t/N$ and minimizing over $t' > 0$, we obtain

$$\frac{1}{N} \log \Big[\mathrm{Prob}(Z_N \geq a) \Big] \leq -I(a), \tag{62}$$

where $I(z) := \sup_{t \in \mathbb{R}} \{tz - \Lambda(t)\}$ is the conjugate of the logarithmic moment generating function $\Lambda(t) := \log M(t)$.

Suppose that $|\mathcal{X}| < \infty$, i.e., the set $\mathcal{X}$ is finite, and: (i) for every $x \in \mathcal{X}$ 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

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

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., $\text{Prob}(\widehat{S}_N^{\delta} \subset S^{\varepsilon}) \geq 1 - \alpha$.

Let $\mathcal{X} = \{x_1, x_2\}$ 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 $\widehat{f}_N(x_2) - \widehat{f}_N(x_1) > 0$. 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 $\widehat{f}_N(x_2) - \widehat{f}_N(x_1)$ is bigger than a specified critical value".

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

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

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

$$\Big|F(x',\xi) - F(x,\xi)\Big| \leq \kappa(\xi)\|x' - x\|, \ \ \forall \xi \in \Xi, \ \forall x', x \in \mathcal{X}.$$

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

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

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

*$O(1)$ denotes a generic constant independent of the data.

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

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

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

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

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

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

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

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

and $\widehat{x}_N = \|\bar{\xi}_N\|^{-\gamma}\bar{\xi}_N$, where $\gamma := \frac{2k-2}{2k-1}$ if $\|\bar{\xi}_N\| \leq 1$, 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 $\|\bar{\xi}_N\|^{\nu} \leq \varepsilon$, 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

$$\text{Prob}\left(\chi^2_n > N\varepsilon^{2/\nu}/\sigma^2\right).$$

Moreover, $\mathbb{E}[\chi^2_n] = n$ and the probability $\text{Prob}(\chi^2_n > n)$ 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

$$N > \frac{n\sigma^2}{\varepsilon^{2/\nu}} \tag{63}$$

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

An alternative approach is going back to Robbins and Monro (1951) and is known as the Stochastic Approximation (SA) method. Assume that the true problem is *convex*, i.e., the set $\mathcal{X} \subset \mathbb{R}^n$ is convex (and closed and bounded) and function $F(\cdot, \xi) : \mathcal{X} \to \mathbb{R}$ is convex for all $\xi \in \Xi$.

Also assume existence of the following *stochastic oracle*: given $x \in \mathcal{X}$ and a random realization $\xi \in \Xi$, the oracle returns the quantity $F(x, \xi)$ and a *stochastic subgradient* – a vector $G(x, \xi)$ such that $g(x) := \mathbb{E}[G(x, \xi)]$ is well defined and is a subgradient of $f(\cdot)$ at x, i.e., $g(x) \in \partial f(x)$. For example, use $G(x, \xi) \in \partial_x F(x, \xi)$.

Classical SA algorithm

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

where $\gamma_j = \theta/j$, $\theta > 0$, and $\Pi_X(x) = \arg\min_{z \in X} \|x - z\|_2$ is the orthogonal (Euclidean) projection onto $\mathcal{X}$. Theoretical bound (assuming $f(\cdot)$ is *strongly convex and differentiable*)

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

for an optimal choice of constant θ (here 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.

As a simple example consider $f(x) = \frac{1}{2}cx^2$, with $c = 0.2$ and $\mathcal{X} = [-1, 1] \subset \mathbb{R}$ and assume that there is no noise, i.e., $G(x, \xi) \equiv \nabla f(x)$. Suppose, further, that we take $\theta = 1$ (i.e., $\gamma_j = 1/j$), which will be the optimal choice for $c = 1$. Then the iteration process becomes

$$x_{j+1} = x_j - f'(x_j)/j = \left(1 - \tfrac{1}{5j}\right) x_j,$$

and hence starting with $x_1 = 1$,

$$\begin{aligned} x_j &= \textstyle\prod_{s=1}^{j-1} \left(1 - \frac{1}{5s}\right) = \exp\left\{-\sum_{s=1}^{j-1} \ln\left(1 + \frac{1}{5s-1}\right)\right\} \\ &> \exp\left\{-\textstyle\sum_{s=1}^{j-1} \frac{1}{5s-1}\right\} > \exp\left\{-\left(0.25 + \int_1^{j-1} \frac{1}{5t-1} dt\right)\right\} \\ &> \exp\left\{-0.25 + 0.2 \ln 1.25 - \tfrac{1}{5} \ln j\right\} > 0.8 j^{-1/5}. \end{aligned}$$

That is, the convergence is extremely slow. For example for $j = 10^9$ the error of the iterated solution is greater than 0.015. On the other hand for the optimal stepsize factor of $\theta = 1/c = 5$, the optimal solution $\bar{x} = 0$ is found in one iteration.

Robust SA approach (with averaging) (B. Polyak, 1990): consider

$$\tilde{x}_j = \sum_{t=1}^{j} \nu_t x_t, \quad \text{where } \nu_t = \frac{\gamma_t}{\sum_{\tau=1}^{j} \gamma_\tau}.$$

Let $D_{\mathcal{X}} = \max_{x \in \mathcal{X}} \|x - x_1\|_2$, and assume that

$$\mathbb{E}\left[\|G(x,\xi)\|_2^2\right] \leq M^2, \quad \forall x \in \mathcal{X},$$

for some constant $M > 0$. Then

$$\mathbb{E}\left[f(\tilde{x}_j) - v^0\right] \leq \frac{D_{\mathcal{X}}^2 + M^2 \sum_{t=1}^{j} \gamma_t^2}{2 \sum_{t=1}^{j} \gamma_t}.$$

For $\gamma_t = \frac{\theta D_{\mathcal{X}}}{M\sqrt{t}}$, after N iterations, we have

$$\mathbb{E}\left[f(\tilde{x}_N) - v^0\right] \leq \frac{\max\left\{\theta, \theta^{-1}\right\} M (D_{\mathcal{X}}^2 + \log N)}{2 D_{\mathcal{X}} \sqrt{N}}.$$

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

$$\mathbb{E}[f(\tilde{x}_N) - v^0] \leq \frac{D_{\mathcal{X}}^2}{2\gamma N} + \frac{\gamma M^2}{2}.$$

For optimal (up to factor θ) $\gamma = \frac{\theta D_{\mathcal{X}}}{M\sqrt{N}}$ we have

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

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

$$\text{Prob}\left\{f(\tilde{x}_N) - v^0 > \varepsilon\right\} \leq \frac{\kappa D_{\mathcal{X}} M}{\varepsilon\sqrt{N}}.$$

Hence the sample size $N \geq \frac{\kappa^2 D_{\mathcal{X}}^2 M^2}{\varepsilon^2 \alpha^2}$ guarantees that $\text{Prob}\left\{f(\tilde{x}_N) - v^0 > \varepsilon\right\} \leq \alpha$.

Mirror Decent SA method (A. Nemirovski)

Let $\|\cdot\|$ be a norm on $\mathbb{R}^n$ and $\omega : \mathcal{X} \to \mathbb{R}$ be continuously differentiable strongly convex on $\mathcal{X}$ with respect to $\|\cdot\|$ function, i.e., for $x, x' \in \mathcal{X}$:

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

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

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

For $\omega(x) = \frac{1}{2}x^T x$ we have that $P_x(y) = \Pi_{\mathcal{X}}(x - y)$. Set

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

$$\widetilde{x}_j = \sum_{t=1}^{j} \nu_t x_t, \quad \text{where } \nu_t = \frac{\gamma_t}{\sum_{\tau=1}^{j} \gamma_\tau}.$$

Then

$$\mathbb{E}\left[f(\tilde{x}_j) - v^0\right] \leq \frac{D^2_{\omega,\mathcal{X}} + (2c)^{-1} M_*^2 \sum_{t=1}^j \gamma_t^2}{2\sum_{t=1}^j \gamma_t},$$

where M_* is a positive constant such that

$$\mathbb{E}\left[\|G(x,\xi)\|_*^2\right] \leq M_*^2, \ \ \forall x \in \mathcal{X},$$

$\|x\|_* = \sup_{\|z\|\leq 1} x^T z$ is the dual norm of the norm $\|\cdot\|$, and

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

For constant step size $\gamma_j = \gamma$, $j = 1, ..., N$, with optimal (up to factor $\theta > 0$) stepsize $\gamma = \frac{\theta D_{\omega,\mathcal{X}}}{M_*}\sqrt{\frac{2c}{N}}$, we have

$$\mathbb{E}\left[f(\tilde{x}_N) - v^0\right] \leq \frac{\max\{\theta, \theta^{-1}\}\sqrt{2} D_{\omega,\mathcal{X}} M_*}{\sqrt{cN}}.$$

Large Deviations type bounds.

Suppose that

$$\mathbb{E}\left[\exp\left\{\|G(x,\xi)\|_*^2/M_*^2\right\}\right] \le \exp\{1\}, \ \ \forall x \in \mathcal{X}.$$

Then for the constant stepsize policy and any $\Omega \ge 1$:

$$\text{Prob}\left\{f(\tilde{x}_N) - v^0 > \frac{\sqrt{2}\max\{\theta,\theta^{-1}\}M_* D_{\omega,\mathcal{X}}(12+2\Omega)}{\sqrt{cN}}\right\} \le 2\exp\{-\Omega\}.$$

It follows that for $\varepsilon > 0$,$\alpha \in (0,1)$ and

$$N_{SA} = O(1)\varepsilon^{-2} D_{\omega,\mathcal{X}}^2 M_*^2 \log^2(1/\alpha),$$

we are guaranteed that $\text{Prob}\left\{f(\tilde{x}_N) - v^0 > \varepsilon\right\} \le \alpha$.

This can be compared with a corresponding estimate of the sample size for the SAA method:

$$N_{\text{SAA}} = O(1)\varepsilon^{-2} R^2 M_*^2 \Big[\log(1/\alpha) + n\log\left(RM_*/\varepsilon\right)\Big].$$

Example

Let $\mathcal{X} = \{x \in \mathbb{R}^n : \sum_{i=1}^n x_i = 1,\ x \geq 0\}$ be a standard simplex. Consider two setups for the Mirror Descent SA: *Euclidean setup*, where $\|\cdot\| = \|\cdot\|_2$ and $\omega(x) = \frac{1}{2}x^T x$, and ℓ_1*-setup*, where $\|\cdot\| = \|\cdot\|_1$ with $\|\cdot\|_* = \|\cdot\|_\infty$, and ω is the *entropy* function

$$\omega(x) = \sum_{i=1}^n x_i \ln x_i.$$

For the constant stepsize policy, the Euclidean setup leads to

$$\mathbb{E}\left[f(\tilde{x}_N) - v^0\right] \leq O(1) \max\left\{\theta, \theta^{-1}\right\} M N^{-1/2},$$

with $M^2 = \sup_{x\in\mathcal{X}} \mathbb{E}\left[\|G(x,\xi)\|_2^2\right]$ (note that the Euclidean diameter of $\mathcal{X}$ is $\sqrt{2}$).

The ℓ_1-setup corresponds to $D_{\omega,\mathcal{X}} = \sqrt{\ln n}$, $c = 1$ and $x_1 = \arg\min_X \omega = n^{-1}(1,...,1)^T$. The associated Mirror Descent SA is easily implementable: the prox mapping can be computed in $O(n)$ operations according to the explicit formula:

$$[P_x(y)]_i = \frac{x_i e^{-y_i}}{\sum_{k=1}^n x_k e^{-y_k}}, \quad i = 1,...,n.$$

The efficiency estimate guaranteed with the ℓ_1-setup is

$$\mathbb{E}\left[f(\widetilde{x}_N) - v^0\right] \leq O(1)\max\left\{\theta, \theta^{-1}\right\}\sqrt{\log n} M_* N^{-1/2},$$

with $M_*^2 = \sup_{x\in X}\mathbb{E}\left[\|G(x,\xi)\|_\infty^2\right]$, To compare the Euclidean and ℓ_1-setups, observe that $M_* \leq M$, and the ratio M_*/M can be as small as $n^{-1/2}$. When X is a standard simplex of large dimension, we have strong reasons to prefer the ℓ_1-setup to the usual Euclidean one.

Bounds by Mirror Decent SA method.

For 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

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

Also $v^0 \leq f(\tilde{x}_N)$ and by convexity of f,

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

That is, for any realization of the random sample $\xi^1, ..., \xi^N$,

$$f_*^N \leq v^0 \leq f^{*N}.$$

Computational (observable) counterparts of these bounds:

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

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

We have that $\mathbb{E}\left[f^{*N}\right] = \mathbb{E}\left[\overline{f}^N\right]$, and

$$\mathbb{E}\left[\underline{f}^N\right] \leq v^0 \leq \mathbb{E}\left[\overline{f}^N\right].$$

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 stagewise 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 stagewise independence 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 $1 - \alpha$. (Here $D_1, D_2, L_1, L_2, \sigma_1, \sigma_2$ are certain analogues of similar constants in the sample size estimate of two stage problem.)

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 $N = N_1^2$.

That is, 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. Another way of putting this is that the number of scenarios needed to solve T-stage problem would grow as $O\left(\varepsilon^{-2(T-1)}\right)$ with decrease of the error level $\varepsilon > 0$.

This indicates that from the point of view of the **number of scenarios**, complexity of multistage programming problems grows exponentially with increase of the number of stages. Furthermore, even if the SAA problem can be solved, its solution does not define a policy for the true problem and of use may be only the computed first stage solution.

Example (Multistage portfolio selection)
Consider the problem of multistage portfolio selection (18)–(20) with logarithmic utility function $U(W) = \ln W$ and stagewise independent data process $\xi_t = 1 + R_t$, $t = 1, ..., T$. Then the optimal value v^0 of the true problem is

$$v^0 = \ln W_0 + \sum_{t=0}^{T-1} \nu_t,$$

where ν_t is the optimal value of the problem*

$$\operatorname*{Max}_{x_t \geq 0} \mathbb{E}\left[\ln(\xi_{t+1}^{\mathsf{T}} x_t)\right] \text{ s.t. } e^{\mathsf{T}} x_t = 1.$$

Let the SAA method be applied with the respective sample ξ_t^j, $j = 1, ..., N_{t-1}$ of ξ_t, $t = 1, ..., T-1$.

*by e we denote vector of ones.

In that case the corresponding SAA problem is also stagewise independent and the optimal value of the SAA problem

$$\widehat{v}_{\mathcal{N}} = \ln W_0 + \sum_{t=0}^{T-1} \widehat{\nu}_{t,N_t}, \tag{64}$$

where $\widehat{\nu}_{t,N_t}$ is the optimal value of the problem

$$\operatorname*{Max}_{x_t \geq 0} \frac{1}{N_t} \sum_{j=1}^{N_t} \ln \left[(\xi_{t+1}^j)^{\mathsf{T}} x_t \right] \ \text{ s.t. } \ e^{\mathsf{T}} x_t = 1. \tag{65}$$

We can view $\widehat{\nu}_{t,N_t}$ as an SAA estimator of ν_t. Since here we solve a maximization rather than a minimization problem, $\widehat{\nu}_{t,N_t}$ is an upwards biased estimator of ν_t, i.e., $\mathbb{E}[\widehat{\nu}_{t,N_t}] \geq \nu_t$. We also have that $\mathbb{E}[\widehat{v}_{\mathcal{N}}] = \ln W_0 + \sum_{t=0}^{T-1} \mathbb{E}[\widehat{\nu}_{t,N_t}]$, and hence

$$\mathbb{E}[\widehat{v}_{\mathcal{N}}] - v^0 = \sum_{t=0}^{T-1} \Big(\mathbb{E}[\widehat{\nu}_{t,N_t}] - \nu_t \Big). \tag{66}$$

That is, for the logarithmic utility function, bias of the SAA estimator of the optimal value grows additively with increase of the number of stages. Also because the samples at different stages are independent of each other, we have that

$$\mathrm{Var}[\widehat{v}_{\mathcal{N}}] = \sum_{t=0}^{T-1} \mathrm{Var}[\widehat{\nu}_{t,N_t}]. \tag{67}$$

On the other hand, for the power utility function $U(W) := W^{\gamma}$, with $\gamma \in (0, 1]$, bias of the corresponding SAA estimator $\widehat{v}_{\mathcal{N}}$ grows with increase of the number of stages in a *multiplicative* way. In particular, if the respective relative biases are constant, then bias of $\widehat{v}_{\mathcal{N}}$ grows *exponentially* fast with increase of the number of stages.

Distributionally robust and risk averse approaches to stochastic programming

$$\operatorname*{Min}_{x\in\mathcal{X}} \left\{ g(x) := \sup_{Q\in\mathfrak{M}} \mathbb{E}_Q[F_x] \right\}, \tag{68}$$

where $\mathcal{X} \subset \mathbb{R}^n$, $F_x(\omega) = F(x, \xi(\omega))$, $F : \mathbb{R}^n \times \Xi \to \mathbb{R}$, $\xi : \Omega \to \Xi$ is a measurable mapping from Ω into $\Xi \subset \mathbb{R}^d$ and $\mathfrak{M}$ is a (nonempty) set of probability measures (distributions) on the sample space $(\Omega, \mathcal{F})$.

Let $\mathcal{Z}$ be a linear space of measurable functions $Z : \Omega \to \mathbb{R}$. We assume that $F_x \in \mathcal{Z}$ for all $x \in \mathcal{X}$. Consider

$$\rho(Z) := \sup_{Q\in\mathfrak{M}} \mathbb{E}_Q[Z] = \sup_{Q\in\mathfrak{M}} \int_\Omega Z(\omega) dQ(\omega), \quad Z \in \mathcal{Z}.$$

The functional $\rho : \mathcal{Z} \to \overline{\mathbb{R}}$ has the following properties:

(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:

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

Functionals $\rho : \mathcal{Z} \to \overline{\mathbb{R}}$ satisfying axioms (A1)–(A4) are called *coherent risk measures* (Artzner, Delbaen, Eber, Heath (1999)).

Examples of dual spaces

Space $\mathcal{Z} := L_p(\Omega, \mathcal{F}, P)$, where P is a (reference) probability measure on $(\Omega, \mathcal{F})$ and $p \in [1, \infty)$. That is, $\mathcal{Z}$ is the space of random variables $Z(\omega)$ having finite p-th order moment.

For $\zeta = dQ/dP$, space $\mathcal{Z}$ is paired with its dual space $\mathcal{Z}^* = L_q(\Omega, \mathcal{F}, P)$, where $1/p + 1/q = 1$, and the scalar product

$$\langle \zeta, Z \rangle := \int_\Omega \zeta(\omega) Z(\omega) dP(\omega), \ \ \zeta \in \mathcal{Z}^*, \ Z \in \mathcal{Z}.$$

We also consider space $\mathcal{Z} := L_\infty(\Omega, \mathcal{F}, P)$, of essentially bounded (measurable) functions $Z : \Omega \to \mathbb{R}$, paired with space $L_1(\Omega, \mathcal{F}, P)$.

Another example. Let Ω be a metric space equipped with its Borel sigma algebra $\mathcal{F}$, and $\mathcal{Z} := C(\Omega)$ be the space of continuous functions $Z : \Omega \to \mathbb{R}$ with the max-norm $\|Z\| = \sup_{\omega\in\Omega} |Z(\omega)|$. Its dual space $\mathcal{Z}^*$ is the space of finite signed measures on $(\Omega, \mathcal{F})$ with the scalar product

$$\langle \mu, Z \rangle := \int_\Omega Z(\omega) d\mu(\omega), \;\; \mu \in \mathcal{Z}^*, \; Z \in \mathcal{Z}.$$

This framework is suitable when the uncertainty set $\mathfrak{M}$ is defined by moment constraints.

We mainly consider the first example with the reference probability space $(\Omega, \mathcal{F}, P)$ and paired spaces $\mathcal{Z} = L_p(\Omega, \mathcal{F}, P)$ and $\mathcal{Z}^* = L_q(\Omega, \mathcal{F}, P)$.

Dual representation of risk functions

By Fenchel-Moreau theorem if $\rho : \mathcal{Z} \to \overline{\mathbb{R}}$ is convex (assumption (A1)) and lower semicontinuous, then

$$\rho(Z) = \sup_{\zeta \in \mathcal{Z}^*} \{\langle \zeta, Z \rangle - \rho^*(\zeta)\}, \ Z \in \mathcal{Z}, \tag{69}$$

where $\rho^*(\zeta) := \sup_{Z \in \mathcal{Z}} \{\langle \zeta, Z \rangle - \rho(Z)\}$. Note that maximization with respect to the dual space $\mathcal{Z}^*$ can be replaced by its subset

$$\mathfrak{A} := \mathrm{dom}(\rho^*) = \{\zeta \in \mathcal{Z}^* : \rho^*(\zeta) < +\infty\}.$$

It is possible to show that condition (A2) (monotonicity) holds iff $\zeta \succeq 0$ for every $\mu \in \mathfrak{A}$. Condition (A3) (translation equivariance) holds iff $\int_\Omega \zeta dP = 1$ for every $\zeta \in \mathfrak{A}$. If ρ is positively homogeneous, then $\rho^*(\zeta) = 0$ for every $\zeta \in \mathfrak{A}$.

If conditions (A1)–(A4) hold, then $\mathfrak{A}$ is a set of probability density functions and

$$\rho(Z) = \sup_{\zeta \in \mathfrak{A}} \int_{\Omega} \zeta(\omega) Z(\omega) dP(\omega). \tag{70}$$

That is

$$\rho(Z) = \sup_{Q \in \mathfrak{M}} \mathbb{E}_Q[Z],$$

where $\mathfrak{M}$ is the set of probability measures Q absolutely continuous with respect to the reference measure P and such that $dQ/dP \in \mathfrak{A}$.

Average Value-at-Risk (also called Conditional Value-at-Risk, Expected Shortfall, Expected Tail Loss, Expected Shortfall)

Chance (probabilistic) constraint:

$$\text{Prob}\Big\{C(x,\xi) > \tau\Big\} \leq \alpha. \tag{71}$$

Constraint (71) can be written as $\mathbb{E}\left[\mathbb{1}_{(0,\infty)}(Z_x)\right] \leq \alpha$, where $Z_x := C(x,\xi) - \tau$. Let $\psi : \mathbb{R} \to \mathbb{R}_+$ be nondecreasing, convex function such that $\psi(\cdot) \geq \mathbb{1}_{(0,\infty)}(\cdot)$. We have that

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

and hence

$$\inf_{t>0} \mathbb{E}[\psi(tZ_x)] \leq \alpha \tag{72}$$

is a conservative approximation of the chance constraint (71).

The choice $\psi^*(z) := [1+z]_+$ gives best conservative approximation. For this choice of ψ, (72) is equivalent to

$$\underbrace{\inf_{t\in\mathbb{R}} \left\{ t + \alpha^{-1}\mathbb{E}[Z_x - t]_+ \right\}}_{\mathsf{AV@R}_\alpha(Z_x)} \leq 0. \tag{73}$$

Note that the minimum in the left hand side of (73) is attained at $t^* = \mathsf{V@R}_\alpha(Z_x)$, where

$$\mathsf{V@R}_\alpha(Z) = H_Z^{-1}(1-\alpha) := \inf\{t : H_Z(t) \geq 1-\alpha\},$$

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

Constraint $\mathrm{Prob}(Z \leq 0) \geq 1 - \alpha$ is equivalent to $\mathrm{V@R}_\alpha(Z) \leq 0$. Therefore $\mathrm{AV@R}_\alpha(C(x,\xi)) \leq \tau$ is a conservative approximation of the chance constraint (71).

Note that $\rho(Z) = \mathrm{AV@R}_\alpha(Z)$ is a coherent risk measure. It is natural to take here $\mathcal{Z} = L_1(\Omega, \mathcal{F}, P)$ and $\mathcal{Z}^* = L_\infty(\Omega, \mathcal{F}, P)$. Dual representation

$$\mathrm{AV@R}_\alpha(Z) = \sup_{\zeta \in \mathfrak{A}} \int_\Omega Z(\omega)\zeta(\omega) dP(\omega),$$

where

$$\mathfrak{A} = \left\{ \zeta \in \mathcal{Z}^* : 0 \leq \zeta(\omega) \leq \alpha^{-1} \text{ a.e. } \omega \in \Omega, \int_\Omega \zeta(\omega) dP(\omega) = 1 \right\}.$$

Suppose that $Z^1, ..., Z^N$ is an iid sample of random variable Z. Then by replacing the true probability distribution P of Z by its empirical* estimate $P_N = \sum_{j=1}^N \delta(Z^j)$ we obtain sample estimate of $\theta := \mathsf{AV@R}_\alpha(Z)$:

$$\widehat{\theta}_N = \inf_{t \in \mathbb{R}} \left\{ t + \alpha^{-1} N^{-1} \sum_{j=1}^N [Z^j - t]_+ \right\}.$$

By the Delta theorem (Theorem 2) we have

$$\widehat{\theta}_N = \inf_{t \in [t^*, t^{**}]} \left\{ t + \alpha^{-1} N^{-1} \sum_{j=1}^N [Z^j - t]_+ \right\} + o_p(N^{-1/2}),$$

where $[t^*, t^{**}]$ is the set of $1 - \alpha$ quantiles of the distribution of the random vector Z. If the $(1 - \alpha)$-quantile $t^* = \mathsf{V@R}_\alpha(Z)$ is unique, then

$$\widehat{\theta}_N = \mathsf{V@R}_\alpha(Z) + \alpha^{-1} Y_N + o_p(N^{-1/2}),$$

where $Y_N := N^{-1} \sum_{j=1}^N [Z^j - t^*]_+$. This approximation can be reasonable when N is significantly bigger than $1/\alpha$.

*By $\delta(z)$ we denote measure of mass one at the point z.

Example Mean-variance risk function ($c > 0$):

$$\rho(Z) := \mathbb{E}[Z] + c\,\mathsf{Var}[Z], \;\; Z \in \mathcal{Z} := L_2(\Omega, \mathcal{F}, P).$$

Dual representation:

$$\rho(Z) = \sup_{\zeta \in \mathcal{Z},\, \mathbb{E}[\zeta]=1} \left\{ \langle \zeta, Z \rangle - (4c)^{-1} \mathsf{Var}[\zeta] \right\}.$$

Satisfies (A1) and (A3), does not satisfy (A2) and (A4).

Example Mean-upper-semideviation risk function of order $p \in [1, +\infty)$, $Z \in \mathcal{Z} := L_p(\Omega, \mathcal{F}, P)$, $c \geq 0$ and

$$\rho(Z) := \mathbb{E}[Z] + c\Big(\mathbb{E}_P \Big\{ [Z - \mathbb{E}_P[Z]]_+^p \Big\} \Big)^{1/p}.$$

Here ρ satisfies (A1),(A3),(A4), and also (A2) (monotonicity) if $c \leq 1$. The max-representation

$$\rho(Z) = \sup_{\zeta \in \mathfrak{A}} \int_\Omega Z(\omega)\zeta(\omega) dP(\omega)$$

holds with $\mathfrak{A} = \{\zeta : \zeta = 1 + h - \int_\Omega h dP,\, \|h\|_q \leq c,\, h \geq 0\}$.

Example ϕ-divergence

Consider a convex continuous function $\phi : \mathbb{R}_+ \to \mathbb{R}_+$ such that $\phi(1) = 0$ and $\phi(x) > 0$ for $x > 1$. Let

$$\mathfrak{A} := \left\{ \zeta \geq 0 : \int_\Omega \phi(\zeta(\omega)) dP(\omega) \leq c, \int_\Omega \zeta(\omega) dP(\omega) = 1 \right\}$$

for some $c > 0$. For example we can take $\phi(x) := |x - 1|^p$, $p \in [1, \infty)$. In that case it will be natural to use the space $\mathcal{Z} = L_p(\Omega, \mathcal{F}, P)$ and

$$\mathfrak{A} = \left\{ \zeta \geq 0 : \|\zeta - 1\|_p \leq c^{1/p}, \int_\Omega \zeta dP = 1 \right\}.$$

For $\phi(x) := x \ln x - x + 1$ we have that $\int_\Omega \phi(\zeta(\omega)) dP(\omega)$ defines the Kullback-Leibler divergence, denoted $D_{KL}(\zeta \| P)$.

It is possible to show that in case of ϕ-divergence the corresponding functional can be written in the form

$$\rho(Z) = \inf_{\lambda \geq 0, \mu} \Big\{ \lambda c + \mu + \mathbb{E}_P[(\lambda\phi)^*(Z - \mu)] \Big\},$$

where $(\lambda\phi)^*$ is the conjugate function of $\lambda\phi$.

In particular for the Kullback-Leibler divergence,

$$\rho(Z) = \inf_{\lambda > 0} \Big\{ \lambda c + \lambda \mathbb{E}_P[e^{Z/\lambda}] \Big\}.$$

Law invariance

It is said random variables $Z, Z' : \Omega \to \mathbb{R}$ are distributionally equivalent if $H_Z = H_{Z'}$, where $H_Z(z) = P(Z \leq z)$ denotes the cumulative distribution function (cdf) of Z with respect to the reference distribution P.

It is said that a risk measure $\rho : \mathcal{Z} \to \mathbb{R}$ is *law invariant* if for any distributionally equivalent $Z, Z' \in \mathcal{Z}$ it follows that $\rho(Z) = \rho(Z')$. That is, law invariant risk measure $\rho(Z)$ is a function of the cdf H_Z of random variable Z.

Consider a nonempty set $\mathfrak{A} \subset \mathcal{Z}^*$ and the corresponding functional

$$\rho_{\mathfrak{A}}(Z) = \sup_{\zeta \in \mathfrak{A}} \int_{\Omega} \zeta(\omega) Z(\omega) dP(\omega), \; Z \in \mathcal{Z},$$

It is said that the uncertainty set $\mathfrak{A}$ is law invariant if $\zeta \in \mathfrak{A}$ and ζ' is distributionally equivalent to ζ, then $\zeta' \in \mathfrak{A}$.

(i) If the uncertainty set $\mathfrak{A}$ is law invariant, then the corresponding functional $\rho_{\mathfrak{A}}$ is law invariant.
(ii) Conversely, if the functional $\rho_{\mathfrak{A}}$ is law invariant and the set $\mathfrak{A}$ is convex and weakly* closed, then $\mathfrak{A}$ is law invariant.

It is said that $\sigma : [0,1) \to \mathbb{R}_+$ is a *spectral* function if $\sigma(\cdot)$ is right side continuous, monotonically nondecreasing and such that $\int_0^1 \sigma(t)dt = 1$.

Suppose that the reference probability measure P is *nonatomic*. Then the dual representation (70) of a law invariant coherent risk measure $\rho : \mathcal{Z} \to \mathbb{R}$ can be written as

$$\rho(Z) = \sup_{\zeta \in \mathfrak{A}} \int_0^1 H_\zeta^{-1}(t) H_Z^{-1}(t) dt,$$

where H_Z and H_ζ are cdfs of Z and ζ respectively. That is

$$\rho(Z) = \sup_{\sigma \in \Upsilon} \int_0^1 \sigma(t) H_Z^{-1}(t) dt, \tag{74}$$

where

$$\Upsilon = \{\sigma = H_\zeta^{-1} : \zeta \in \mathfrak{A}\}.$$

Kusuoka representation

For a probability measure (distribution) μ on the interval $[0,1)$ consider spectral function

$$\sigma(t) := \int_0^t (1-\alpha)^{-1} d\mu(\alpha),\ t \in [0,1).$$

This equation defines a mapping $\sigma = \mathbb{T}\mu$ from the set of probability measures on $[0,1)$ to the set of spectral functions. The mapping $\mathbb{T}$ is one-to-one and onto and its inverse is

$$(\mathbb{T}^{-1}\sigma)(\alpha) = (1-\alpha)\sigma(\alpha) + \int_0^\alpha \sigma(t)dt,\ \alpha \in [0,1). \qquad (75)$$

For $\sigma = \mathbb{T}\mu$ we have

$$\int_0^1 \sigma(t)H^{-1}(t)dt = \int_0^1 \int_0^t (1-\alpha)^{-1} H^{-1}(t) d\mu(\alpha)dt =$$

$$\int_0^1 (1-\alpha)^{-1} \int_\alpha^1 H^{-1}(t)dtd\mu(\alpha) = \int_0^1 \mathsf{AV@R}_{1-\alpha}(H)d\mu(\alpha).$$

Thus the dual representation (74) leads to the Kusuoka representation of the law invariant coherent risk measure ρ,

$$\rho(Z) = \sup_{\mu \in \mathfrak{M}} \int_0^1 \mathsf{AV@R}_{1-\alpha}(Z) d\mu(\alpha),$$

where $\mathfrak{M}$ is a set of probability measures on [0,1) given by $\mathfrak{M} = \mathbb{T}^{-1}(\Upsilon)$.

For example, consider the absolute semideviation risk measure

$$\rho(Z) := \mathbb{E}[Z] + c\mathbb{E}\left\{ \left[Z - \mathbb{E}[Z] \right]_+ \right\}, \ c \in [0,1].$$

Its (minimal) Kusuoka representation is

$$\rho(Z) = \sup_{\alpha \in (0,1)} \left\{ (1 - c\alpha)\mathsf{AV@R}_1(Z) + c\alpha \mathsf{AV@R}_\alpha(Z) \right\}.$$

Interchangeability principle

Let $\mathcal{R} : \mathcal{Z} \to \mathbb{R}$ be a *monotone* functional, i.e. if $Z, Z' \in \mathcal{Z}$ and $Z \succeq Z'$, then $\mathcal{R}(Z) \geq \mathcal{R}(Z')$. Consider

$$F(\omega) := \inf_{y \in Y} f(y, \omega), \tag{76}$$

where Y is an abstract set and $f : Y \times \Omega \to \mathbb{R} \cup \{+\infty\}$ is an extended real valued function. Let $\mathcal{Y}$ be the set of mappings $\eta : \Omega \to Y$ such that $f_\eta \in \mathcal{Z}$, where $f_\eta(\cdot) := f(\eta(\cdot), \cdot)$.

Suppose that the minimum in (76) is attained at $\bar{y}(\omega) \in Y$ for $\omega \in \Omega$, and hence $F(\omega) = f(\bar{y}(\omega), \omega)$. Then by monotonicity of $\mathcal{R}$, assuming that $F \in \mathcal{Z}$, we have that

$$\mathcal{R}(F) = \inf_{\eta \in \mathcal{Y}} \mathcal{R}(f_\eta).$$

That is the minimization operator and functional $\mathcal{R}$ can be interchanged. For monotone functionals this interchangeability holds in general (without assuming existence of minimizers). Moreover, the following implication holds

$$\bar{\eta}(\cdot) \in \arg\min_{y \in Y} f(y, \cdot) \Rightarrow \bar{\eta} \in \arg\min_{\eta \in \mathcal{Y}} \mathcal{R}(f_\eta).$$

It is possible to give simple examples showing that the converse implication

$$\bar{\eta} \in \arg\min_{\eta \in \mathcal{Y}} \mathcal{R}(f_\eta) \Rightarrow \bar{\eta}(\cdot) \in \arg\min_{y \in Y} f(y, \cdot)$$

may not hold, unless the functional $\mathcal{R}$ is *strictly* monotone.

Definition 1 *It is said that a functional* $\mathcal{R} : \mathcal{Z} \to \mathbb{R}$ *is strictly monotone, if* $Z, Z' \in \mathcal{Z}$, $Z \succeq Z'$ *and* $Z \neq Z'$ *imply that* $\mathcal{R}(Z) > \mathcal{R}(Z')$.

Conditions for (strict) monotonicity

In the setting of $\mathcal{Z} = L_p(\Omega, \mathcal{F}, \mathbb{P})$ and $\mathcal{Z}^* = L_q(\Omega, \mathcal{F}, \mathbb{P})$, a convex continuous functional $\mathcal{R} : \mathcal{Z} \to \mathbb{R}$ is (strictly) monotone iff for every $Z \in \mathcal{Z}$ and $\zeta \in \partial\mathcal{R}(Z)$ it follows that $\zeta(\omega) \geq 0$ ($\zeta(\omega) > 0$) for a.e. $\omega \in \Omega$.

In particular if $\mathcal{R}$ is a coherent risk measure and

$$\mathcal{R}(Z) = \sup_{\zeta \in \mathfrak{A}} \int_\Omega \zeta(\omega) Z(\omega) d\mathbb{P}$$

its dual representation with $\mathfrak{A}$ being a convex weakly* closed bounded subset of $\mathcal{Z}^*$, then $\partial\mathcal{R}(0) = \mathfrak{A}$.

For example AV@R$_\alpha(Z)$ is not strictly monotone for $\alpha \in (0, 1)$.

In the setting of $\mathcal{Z} = C(\Omega)$ and $\mathcal{R}(Z) = \sup_{P \in \mathfrak{M}} \mathbb{E}_P[Z]$, with $\mathfrak{M}$ being convex and weakly* closed, the functional $\mathcal{R} : \mathcal{Z} \to \mathbb{R}$ is strictly monotone iff for every $\mu \in \mathfrak{M}$ it follows that $\mu(A) > 0$ for every open set $A \subset \Omega$.

In particular, if the set $\mathfrak{M}$ is defined by q moment constraints, then by the Richter-Rogosinski Theorem the maximum is attained at a probability measure supported on a finite set of no more than $q + 1$ points. It follows that $\mathcal{R}$ is monotone, but is not strictly monotone if the set Ω has more than $q + 1$ points.

Dynamic equations and time consistency

Consider two stage risk averse stochastic program

$$\min_{x\in\mathcal{X}_1, y(\cdot)\in\mathcal{X}_2(x,\cdot)} \mathcal{R}\Big(g(x, y(\omega), \omega)\Big), \tag{77}$$

where $\mathcal{X}_1 \subset \mathbb{R}^n$, $g : \mathbb{R}^n \times \mathbb{R}^k \times \Omega \to \mathbb{R}$ and $\mathcal{X}_2 : \mathbb{R}^n \times \Omega \rightrightarrows \mathbb{R}^k$ is a multifunction. An alternative formulation is

$$\min_{x\in\mathcal{X}_1} \mathcal{R}\Big(\underbrace{\min_{y\in\mathcal{X}_2(x,\omega)} g(x, y, \omega)}_{f(x,\omega)}\Big), \tag{78}$$

where $f_x(\omega) = f(x, \omega)$ is the optimal value of the second stage problem.

An optimal solution $(\bar{x}, \bar{y}(\cdot))$ of problem (77) is time consistent if $\bar{y}(\cdot)$ is an optimal solution of the second stage program, i.e.,

$$\bar{y}(\omega) \in \arg\min_{y\in\mathcal{X}_2(\bar{x},\omega)} g(\bar{x}, y, \omega),\ \omega \in \Omega.$$

For the two stage programs the above means that the optimal values of problems (77) and (78) are the same and if $\bar{x}$ is an optimal solution of the first stage problem and $\bar{y}(\omega)$, $\omega \in \Omega$, is an optimal solution of the second stage problem

$$\min_{y \in \mathcal{X}_2(\bar{x},\omega)} g(\bar{x}, y, \omega),$$

then $(\bar{x}, \bar{y}(\cdot))$ is an optimal solution of problem (77). The converse of that is true if $\mathcal{R}$ is *strictly* monotone.

If $\mathcal{R}$ is not strictly monotone, then the first stage problem (77) could have an optimal solution $(\bar{x}, \bar{y}(\cdot))$ such that conditional on $x = \bar{x}$, the solution $\bar{y}(\omega)$ is not optimal for the second stage problem for some $\omega \in \Omega$. That is for some $\omega \in \Omega$ the corresponding value $g(\bar{x}, \bar{y}(\omega), \omega)$ is strictly bigger than the minimal value

$$\min_{y \in \mathcal{X}_2(\bar{x}, \omega)} g(\bar{x}, y, \omega).$$

Such solutions are not time consistent. That is, without strict monotonicity time inconsistent optimal policies could exist already for two stage problems and finite number of scenarios.

Tree representation of risk averse multistage setting.

Notation: Ω_t set of nodes at stage $t = 1, ..., T$, $K_t = |\Omega_t|$, $C_a \subset \Omega_{t+1}$ set of children nodes of $a \in \Omega_t$. With every node $a \in \Omega_t$ we can associate risk function

$$\rho^a : \mathbb{R}^{|C_a|} \to \mathbb{R},\ a \in \Omega_t,\ t = 1, ..., T-1.$$

For example, let $p^a \in \mathbb{R}^{|C_a|}$ be conditional probability vector of moving from node a to its children nodes, and

$$\rho^a(Z) := \mathbb{E}_{p^a}[Z],\ a \in \Omega_t,\ t = 1, ..., T-1,$$

or

$$\rho^a(Z) := \mathbb{E}_{p^a}[Z] + c_a \mathbb{E}_{p^a}\Big[Z - \mathbb{E}_{p^a}[Z]\Big]_+,\ c_a \in [0, 1].$$

Note that $\mathbb{R}^{K_{t+1}} = \mathbb{R}^{|C_{a_1}|} \times \cdots \times \mathbb{R}^{|C_{a_{K_t}}|}$, where $\{a_1, ..., a_{K_t}\} = \Omega_t$. Define

$$\rho_{t+1} := (\rho^{a_1}, ..., \rho^{a_{K_t}}) : \mathbb{R}^{K_{t+1}} \to \mathbb{R}^{K_t},\ \ t = 1, ..., T-1.$$

With the considered tree is associated sequence of (finite) sigma algebras $\mathcal{F}_1 \subset \cdots \subset \mathcal{F}_T$, with $\mathcal{F}_T = 2^{\Omega_T}$ and $\mathcal{F}_1 = \{\emptyset, \Omega_T\}$. Let $\mathcal{Z}_t$ be the space of all $\mathcal{F}_t$-measurable functions $Z : \Omega_T \to \mathbb{R}$, i.e., $Z(\cdot)$ is constant on every C_a, $a \in \Omega_t$. The space $\mathcal{Z}_t$ can be identified with $\mathbb{R}^{K_t}$. It is said that $\rho_{t+1} : \mathcal{Z}_{t+1} \to \mathcal{Z}_t$ is a conditional risk mapping if the following conditions hold
(A'1) Convexity:

$$\rho_{t+1}(\tau Z_1 + (1-\tau)Z_2) \preceq \tau\rho_{t+1}(Z_1) + (1-\tau)\rho_{t+1}(Z_2)$$

for all $Z_1, Z_2 \in \mathcal{Z}_{t+1}$ and $\tau \in [0,1]$.
(A'2) Monotonicity: If $Z_1, Z_2 \in \mathcal{Z}_{t+1}$ and $Z_2 \succeq Z_1$, then $\rho_{t+1}(Z_2) \succeq \rho_{t+1}(Z_1)$.
(A'3) Translation Equivariance: If $Y \in \mathcal{Z}_t$ and $Z \in \mathcal{Z}_{t+1}$, then $\rho_{t+1}(Z+Y) = \rho_{t+1}(Z) + Y$,
(A'4) Positive Homogeneity:

$$\rho_{t+1}(\tau Z) = \tau\rho_{t+1}(Z), \quad Z \in \mathcal{Z}_{t+1},\ \tau > 0.$$

We have that with each coherent risk measure ρ^a, $a \in \Omega_t$ is associated set $\mathfrak{A}_{t+1}(a) \subset \mathbb{R}^{K_{t+1}}$ of probability vectors such that

$$\rho^a(Z) = \max_{p \in \mathfrak{A}_{t+1}(a)} \mathbb{E}_p[Z].$$

Let $\nu = (\nu_a)_{a \in \Omega_t}$ be a probability distribution on Ω_t, and

$$\mathfrak{C}_{t+1} := \left\{ \mu = \textstyle\sum_{a \in \Omega_t} \nu_a p^a : p^a \in \mathfrak{A}_{t+1}(a) \right\}.$$

We have that

$$\mathbb{E}_\mu[Z|\mathcal{F}_t] = \mathbb{E}_{p^a}[Z],\ Z \in \mathcal{Z}_{t+1},$$

and it follows that

$$\rho_{t+1}(Z) = \max_{\mu \in \mathfrak{C}_{t+1}} \mathbb{E}_\mu[Z|\mathcal{F}_t].$$

Risk averse multistage programming (nested formulation):

$$\begin{aligned} \underset{x_1 \in \mathcal{X}_1}{\text{Min}} \quad & f_1(x_1) + \rho_2\Big[\inf_{x_2 \in \mathcal{X}_2(x_1,\omega)} f_2(x_2,\omega) + \ldots \\ & + \rho_{T-1}\Big[\inf_{x_{T-1} \in \mathcal{X}_T(x_{T-2},\omega)} f_{T-1}(x_{T-1},\omega) \\ & + \rho_T[\inf_{x_T \in \mathcal{X}_T(x_{T-1},\omega)} f_T(x_T,\omega)]\Big]\Big]. \end{aligned} \tag{79}$$

Here ω is an element of $\Omega := \Omega_T$, the objective functions $f_t : \mathbb{R}^{n_{t-1}} \times \Omega \to \mathbb{R}$ are real valued functions and $\mathcal{X}_t : \mathbb{R}^{n_{t-1}} \times \Omega \rightrightarrows \mathbb{R}^{n_t}$, $t = 2, \ldots, T$, are multifunctions such that $f_t(x_t, \cdot)$ and $\mathcal{X}_t(x_{t-1}, \cdot)$ are $\mathcal{F}_t$-measurable for all x_t and x_{t-1}. Note that if the corresponding risk measures ρ^a are defined as conditional expectations, then the above multistage problem becomes a risk neutral multistage problem.

Consider function $\varrho : \mathcal{Z}_1 \times \cdots \times \mathcal{Z}_T \to \mathbb{R}$ defined as

$$\varrho(Z_1, \ldots, Z_T) := Z_1 + \rho_2\Big[Z_2 + \cdots + \rho_{T-1}\Big[Z_{T-1} + \rho_T[Z_T]\Big]\Big]. \quad (80)$$

By condition (A'3) we have that

$$\rho_{T-1}\Big[Z_{T-1} + \rho_T[Z_T]\Big] = \rho_{T-1} \circ \rho_T\Big[Z_{T-1} + Z_T\Big].$$

By continuing this process we obtain that

$$\varrho(Z_1, \ldots, Z_T) = \bar{\rho}(Z_1 + \cdots + Z_T),$$

where $\bar{\rho} := \rho_2 \circ \cdots \circ \rho_T$. We refer to $\bar{\rho} : \mathcal{Z}_T \to \mathbb{R}$ as the composite risk measure.

That is,

$$\bar{\rho}(Z_1 + \cdots + Z_T) = Z_1 + \rho_2\Big[Z_2 + \cdots + \rho_{T-1}\Big[Z_{T-1} + \rho_T[Z_T]\Big]\Big],$$

defined for $Z_t \in \mathcal{Z}_t$, $t = 1, \ldots, T$. Conditions (A'1)–(A'4) imply that $\bar{\rho}$ is a coherent risk measure.

We can write the risk averse multistage programming problem

$$\begin{array}{ll} \underset{x_1, \boldsymbol{x}_2, \ldots, \boldsymbol{x}_T}{\text{Min}} & \bar{\rho}\Big[f_1(x_1) + f_2(\boldsymbol{x}_2(\omega), \omega) + \cdots + f_T\left(\boldsymbol{x}_T(\omega), \omega\right)\Big] \\ \text{s.t.} & x_1 \in \mathcal{X}_1,\ \boldsymbol{x}_t(\omega) \in \mathcal{X}_t(\boldsymbol{x}_{t-1}(\omega), \omega),\ t = 2, \ldots, T. \end{array}$$

Distributionally robust multistage stochastic programming.

We can write the risk neutral multistage problem as

$$\operatorname*{Min}_{\pi \in \Pi} \mathbb{E}_P[Z^\pi], \tag{81}$$

where P is the probability distribution of random vector $\xi_{[T]} = (\xi_1, ..., \xi_T)$, Π is a set of policies satisfying the feasibility constraints

$$x_1 \in \mathcal{X}_1,\ x_t(\xi_{[t]}) \in \mathcal{X}_t(x_{t-1}(\xi_{[t-1]}), \xi_t),\ t = 2, ..., T-1,$$

and $Z^\pi = Z^\pi(\xi_{[T]})$ is defined as

$$Z^\pi := f_1(x_1) + f_2(x_2(\xi_{[2]}), \xi_2) + \cdots + f_T(x_T(\xi_{[T]}), \xi_T).$$

It looks natural to formulate the following distributionally robust analogue of problem (81). Consider a set $\mathfrak{M}$ of probability distributions of $\xi_{[T]}$ supported on a set $\Xi \subset \mathbb{R}^{d_1} \times \cdots \times \mathbb{R}^{d_T}$ equipped with its Borel sigma algebra $\mathcal{B}$, and the min-max program

$$\operatorname*{Min}_{\pi \in \Pi} \sup_{Q \in \mathfrak{M}} \mathbb{E}_Q[Z^\pi]. \tag{82}$$

However, there is a problem with formulation (82).

The expectation operator has the following property (recall that ξ_1 is deterministic)

$$\mathbb{E}[Z] = \mathbb{E}_{|\xi_1}\left[\mathbb{E}_{|\xi_{[2]}}\left[\cdots \mathbb{E}_{|\xi_{[T-1]}}(Z)\right]\right].$$

For $Z = Z(\xi_{[T]}) \in \mathcal{Z}$ and $Q \in \mathfrak{M}$ we have that

$$\mathbb{E}_Q[Z] = \mathbb{E}_Q\left[\mathbb{E}_{Q|\xi_{[2]}}\left[\cdots \mathbb{E}_{Q|\xi_{[T-1]}}[Z]\right]\right],$$

and hence

$$\sup_{Q\in\mathfrak{M}} \mathbb{E}_Q[Z] \leq \sup_{Q\in\mathfrak{M}} \mathbb{E}_Q\left[\sup_{Q\in\mathfrak{M}} \mathbb{E}_{Q|\xi_{[2]}}\left[\cdots \sup_{Q\in\mathfrak{M}} \mathbb{E}_{Q|\xi_{[T-1]}}[Z]\right]\right]. \quad (83)$$

In a certain sense the the right hand side of (83) represents the tightest upper bound for the left hand side among all possible coherent and time consistent upper bounds.

Suppose that the right hand side of (83) is finite for all $Z \in \mathcal{Z}$. Then there exists a bounded set $\widehat{\mathfrak{M}} \subset \mathcal{Z}^*$ of probability measures such that for any $Z \in \mathcal{Z}$,

$$\sup_{Q \in \widehat{\mathfrak{M}}} \mathbb{E}_Q[Z] = \sup_{Q \in \mathfrak{M}} \mathbb{E}_Q \left[\sup_{Q \in \mathfrak{M}} \mathbb{E}_{Q|\xi_{[1]}} \left[\cdots \sup_{Q \in \mathfrak{M}} \mathbb{E}_{Q|\xi_{[T-1]}}[Z] \right] \right].$$

The analysis simplifies if we assume that the set $\mathfrak{M}$ consists of product measures (the rectangular case). That is,

$$\mathfrak{M} := \{Q = Q_1 \times \cdots \times Q_T : Q_t \in \mathfrak{M}_t,\ t = 1, ..., T\},$$

where $\Xi = \Xi_1 \times \cdots \times \Xi_T$ and $\mathfrak{M}_t$ is a set of probability measures on $(\Xi_t, \mathcal{B}_t)$, $t = 1, ..., T$.

If we view $\xi_1, ..., \xi_T$ as a random process having distribution $Q \in \mathfrak{M}$, then this means that random vectors ξ_t, $t = 1, ..., T$, are mutually independent with respective marginal distributions $Q_t \in \mathfrak{M}_t$. For $\mathfrak{M} \ni Q = Q_1 \times \cdots \times Q_T$ we can write

$$\mathbb{E}_{Q|\xi_{[T-1]}}[Z] = \int_{\Xi_T} Z(\xi_{[T-1]}, \xi_T) dQ_T(\xi_T) =: \mathbb{E}_{Q_T|\xi_{[T-1]}}[Z],$$

and hence

$$\sup_{Q \in \mathfrak{M}} \mathbb{E}_Q[Z] \leq \sup_{Q_1 \in \mathfrak{M}_1} \mathbb{E}_{Q_1} \left[\sup_{Q_2 \in \mathfrak{M}_2} \mathbb{E}_{Q_2|\xi_{[1]}} \left[\cdots \sup_{Q_T \in \mathfrak{M}_T} \mathbb{E}_{Q_T|\xi_{[T-1]}}[Z] \right] \right].$$

Dynamic Programming Equations.

For the last period T we have

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

$$\mathcal{Q}_T(x_{T-1}, \omega) := \rho_T[Q_T(x_{T-1}, \omega)],$$

and for $t = T - 1, \ldots, 2$,

$$\mathcal{Q}_t(x_{t-1}, \omega) := \rho_t \left[Q_t(x_{t-1}, \omega) \right],$$

where

$$Q_t(x_{t-1}, \omega) := \inf_{x_t \in \mathcal{X}_t(x_{t-1}, \omega)} \Big\{ f_t(x_t, \omega) + \mathcal{Q}_{t+1}(x_t, \omega) \Big\}.$$

Finally, at the first stage we solve the problem

$$\operatorname*{Min}_{x_1 \in \mathcal{X}_1} f_1(x_1) + \rho_2[Q_2(x_1, \omega)].$$

By using the conditional max-representation, we can write the dynamic programming equations in the form

$$Q_t(x_{t-1}, \omega) = \inf_{x_t \in \mathcal{X}_t(x_{t-1}, \omega)} \left\{ f_t(x_t, \omega) + \sup_{\mu \in \mathfrak{C}_{t+1}} \mathbb{E}_\mu \left[Q_{t+1}(x_t) \middle| \mathcal{F}_t \right](\omega) \right\}.$$

Time consistency of stochastic programming problems

Consider multistage stochastic optimization problem

$$\begin{array}{ll} \mathrm{Min}_{\pi\in\Pi} & \mathcal{R}\Big[f_1(x_1), \ldots, f_T(x_T,\omega)\Big], \\ \text{s.t.} & x_1 \in \mathcal{X}_1,\ x_t \in \mathcal{X}_t(x_{t-1},\omega),\ t = 2,...,T. \end{array} \tag{84}$$

Optimization in the above problem is performed over feasible policies adapted to filtration $\mathfrak{F} = (\mathcal{F}_1,...,\mathcal{F}_T)$. That is, $\mathcal{F}_1 \subset \cdots \subset \mathcal{F}_T$ is a sequence of sigma algebras defined on a probability space $(\Omega, \mathcal{F}, P)$, Π is the set of policies $\pi = (x_1, x_2(\omega),...,x_T(\omega))$ satisfying the feasibility constraints w.p.1, and $f_1 : \mathbb{R}^{n_1} \to \mathbb{R}$, $\mathcal{X}_1 \subset \mathbb{R}^{n_1}$, $f_t : \mathbb{R}^{n_t} \times \Omega \to \mathbb{R}$ and $\mathcal{X}_t : \mathbb{R}^{n_{t-1}} \times \Omega \rightrightarrows \mathbb{R}^{n_t}$, $t = 2,...,T$. It is assumed that $f_t(x_t,\cdot)$ and $\mathcal{X}_t(x_{t-1},\cdot)$ are $\mathcal{F}_t$-measurable. Filtration $\mathfrak{F}$ represents flow of information with progress of the time and $x_t : \Omega \to \mathbb{R}^{n_t}$ is an $\mathcal{F}_t$-measurable mapping defining our decision at time $t = 1,...,T$. It is assumed that $\mathcal{F}_1 = \{\emptyset, \Omega\}$, and hence decision x_1 is deterministic, made before observing any realization of the data process.

Denote $\mathcal{Z}_t := L_p(\Omega, \mathcal{F}_t, P)$ the space of $\mathcal{F}_t$-measurable variables $Z_t : \Omega \to \mathbb{R}$, $t = 1, ..., T$, having finite p-th order moments. The functional $\mathcal{R} : \mathcal{Z}_1 \times \cdots \times \mathcal{Z}_T \to \mathbb{R}$ represents a choice of the objective in solving the corresponding multistage problem. The standard choice of $\mathcal{R}$ is the expected value of the total sum of the costs, that is

$$\mathcal{R}(Z_1, ..., Z_T) := \mathbb{E}[Z_1 + ... + Z_T].$$

In that case problem (84) becomes the so-called risk neutral multistage stochastic program

$$\begin{array}{ll} \mathrm{Min}_{\pi \in \Pi} & \mathbb{E}\Big[f_1(x_1) + \cdots + f_T(x_T, \omega)\Big], \\ \text{s.t.} & x_1 \in \mathcal{X}_1,\ x_t \in \mathcal{X}_t(x_{t-1}, \omega),\ t = 2, ..., T. \end{array} \tag{85}$$

There are many other possible choices of the functional $\mathcal{R}$ leading to different formulations of risk averse multistage programs.

A well known quotation of Bellman: *"An optimal policy has the property that whatever the initial state and initial decision are, the remaining decisions must constitute an optimal policy with regard to the state resulting from the first decision."*

It is said that a policy solving the reference problem (84), before any realization of the uncertainty data became available, is time consistent if it remains optimal at the later stages.

This formulation is quite vague since it is not clearly specified what optimality at the later stages means. In order to make this precise we need to specify preferences

$$\mathcal{R}_{s,t} : \mathcal{Z}_s \times \cdots \times \mathcal{Z}_t \to \mathcal{Z}_s, \ 1 \leq s < t \leq T, \tag{86}$$

defining our objective from stage s to stage t. Note that since $\mathcal{F}_1$ is trivial, the space $\mathcal{Z}_1$ can be identified with $\mathbb{R}$ and hence the functional $\mathcal{R}_{1,t}$ is real valued. We assume that $\mathcal{R}_{1,T} = \mathcal{R}$.

Definition 2 *We say that an optimal policy* $\bar{\pi} = (\bar{x}_1, ..., \bar{x}_T)$ *of the reference problem* (84) *is time consistent if for all* $1 \leq s < t \leq T$, *the policy* $(\bar{x}_s, ..., \bar{x}_t)$ *is optimal for the problem*

$$\begin{array}{ll} \text{Min} & \mathcal{R}_{s,t}\Big[f_s(x_1), \ldots, f_t(x_t, \omega)\Big], \\ \text{s.t.} & x_\tau \in \mathcal{X}_\tau(x_{\tau-1}, \omega),\ \tau = s, ..., t, \end{array} \tag{87}$$

conditional on $\mathcal{F}_s$ *and* $\bar{x}_{s-1}$.

For the risk neutral problem it is natural to define preferences as conditional expectations $\mathcal{R}_{s,t} := \mathbb{E}_{|\mathcal{F}_s}$. In that case any optimal policy of the reference problem is time consistent. This follows from the decomposability property of the expectation operator

$$\mathbb{E}_{|\mathcal{F}_s}[\cdot] = \mathbb{E}_{|\mathcal{F}_s}\Big[\mathbb{E}_{|\mathcal{F}_{s+1}}\Big[\cdots \mathbb{E}_{|\mathcal{F}_t}[\cdot]\Big]\Big],\ 1 \leq s < t \leq T.$$

The additive case

One of the most common cases considered in applications is optimization of the total cost. In that case the reference problem can be written as

$$\begin{array}{ll} \text{Min} & \varrho\Big[f_1(x_1) + \cdots + f_T(x_T, \omega)\Big], \\ \text{s.t.} & x_1 \in \mathcal{X}_1,\ x_t \in \mathcal{X}_t(x_{t-1}, \omega),\ t = 2, ..., T, \end{array} \tag{88}$$

for some functional $\varrho : \mathcal{Z}_T \to \mathbb{R}$. In case of $\varrho(\cdot) := \mathbb{E}[\cdot]$ this becomes the risk neutral problem. For this type of problems the preference system can be defined by functionals $\rho_{t,T} : \mathcal{Z}_T \to \mathcal{Z}_t$, $t = 1, ..., T-1$, with $\rho_{1,T} = \varrho$ and the corresponding conditional problems

$$\begin{array}{ll} \text{Min} & \rho_{t,T}\Big[f_t(x_t, \omega) + \cdots + f_T(x_T, \omega)\Big], \\ \text{s.t.} & x_\tau \in \mathcal{X}_\tau(x_{\tau-1}, \omega),\ \tau = t, ..., T. \end{array} \tag{89}$$

Again in the risk neutral case $\rho_{t,T}(\cdot) := \mathbb{E}_{|\mathcal{F}_t}[\cdot]$. As another example consider the following max-type functional

$$\varrho(Z) := \operatorname{ess\,sup}(Z(\omega)). \tag{90}$$

This corresponds to the so-called robust approach to multistage optimization under uncertainty. It is natural to define here $\rho_{t,T}$ as the essential supremum conditional on $\mathcal{F}_t$. Similar to conditional expectations, such system of preferences has the following recursive property

$$\rho_{s,T}\left(\rho_{t,T}(\cdot)\right) = \rho_{s,T}(\cdot),\ 1 < s < t \leq T-1. \tag{91}$$

In case of the max-type functional ϱ there is no guarantee that every optimal solution of the reference problem is time consistent. Reason for this is that the conditional ess sup mappings $\rho_{s,T}$ are not *strictly* monotone.

The following definition in slightly different forms was used by several authors: Wang (1999), Riedel (2004), Cheridito et al. (2006), Artzner et al. (2007), Ruszczynski (2010).

Definition 3 *The preference system* $\{\rho_{t,T}\}$ *is said to be dynamically consistent if the implication*

$$Z, Z' \in \mathcal{Z}_T \; and \; \rho_{t,T}(Z) \succeq \rho_{t,T}(Z') \implies \rho_{s,T}(Z) \succeq \rho_{s,T}(Z').$$

holds for all $1 \leq s < t \leq T-1$.

It turns out that the above 'forward' property of dynamic consistency is not always sufficient to ensure that every optimal policy is time consistent.

Definition 4 *A dynamically consistent preference system* $\{\rho_{t,T}\}$ *is said to be strictly dynamically consistent if the implication*

$$Z, Z' \in \mathcal{Z}_T,\ \rho_{t,T}(Z) \succeq \rho_{t,T}(Z') \text{ and } \rho_{t,T}(Z) \neq \rho_{t,T}(Z') \text{ imply that } \rho_{s,T}(Z) \succeq \rho_{s,T}(Z') \text{ and } \rho_{s,T}(Z) \neq \rho_{s,T}(Z')$$

holds for all $1 \leq s < t \leq T-1$.

- Suppose that preference mappings $\rho_{t,T}$, $1 \leq t < T-1$, are monotone (strictly monotone) and the preference system is recursive. Then the preference system is (strictly) dynamically consistent. (ii) Conversely, if the preference system is dynamically consistent, translation equivariant and $\rho_{t,T}(0) = 0$, $t = 1, ..., T-1$, then the preference system is recursive.

- Suppose that the preference system is recursive and the preference mappings $\rho_{t,T}$, $t = 1, ..., T-1$, are monotone. Then the following holds. (i) If $\bar{\pi} \in \Pi$ is the unique optimal solution of the reference problem (84), then $\bar{\pi}$ is time consistent. (ii) If moreover the preference mappings $\rho_{t,T}$ are strictly monotone, then every optimal solution of the reference problem is time consistent.

The Average Value-at-Risk measure has conditional analogues $\mathsf{AV@R}_{\alpha|\mathcal{F}_t}$, which can be used to define a preference system. This preference system is not recursive, i.e.,

$$\mathsf{AV@R}_{\alpha}\left(\mathsf{AV@R}_{\alpha|\mathcal{F}_t}(\cdot)\right) \neq \mathsf{AV@R}_{\alpha}(\cdot)$$

for a subalgebra $\mathcal{F}_t \subset \mathcal{F}$, $\mathcal{F}_t \neq \mathcal{F}$. A respective recursive system is obtained by using nested mappings

$$\rho_{t,T}(\cdot) := \mathsf{AV@R}_{\alpha|\mathcal{F}_t}\left(\mathsf{AV@R}_{\alpha|\mathcal{F}_{t+1}}\left(\cdots \mathsf{AV@R}_{\alpha|\mathcal{F}_{T-1}}(\cdot)\right)\right).$$

These mappings are monotone, but are not strictly monotone.

It is said that preference mappings $\rho_{t,T}$ are decomposable via a family of one-step mappings $\rho_t : \mathcal{Z}_t \to \mathcal{Z}_{t-1}$, if they can be represented as compositions

$$\rho_{t,T}(\cdot) = \rho_{t+1}(\cdots \rho_T(\cdot)).$$

It is said that $\rho_{t,T}$ are translation equivariant if $Z \in \mathcal{Z}_T$ and $Y \in \mathcal{Z}_t$ imply that $\rho_{t,T}(Z + Y) = \rho_{t,T}(Z) + Y$.

If the preference system is recursive, then

$$\rho_{t,T}(Z) = \rho_{t,T}\Big(\rho_{t+1,T}(\cdots \rho_{T-1,T}(Z))\Big),$$

which gives a decomposition of $\rho_{t,T}$ with the associate one-step mappings $\rho_\tau := \rho_{\tau-1,T}$ considered as mappings from $\mathcal{Z}_\tau$ to $\mathcal{Z}_{\tau-1}$. Conversely, if mappings ρ_t are translation equivariant and $\rho_t(0) = 0$, then the corresponding decomposable preference system is recursive.

For decomposable systems it is possible to write dynamic programming equations. Suppose that the mappings $\rho_t : \mathcal{Z}_t \to \mathcal{Z}_{t-1}$ are monotone and translation equivariant. Then *dynamic programming equations* for the reference problem can be written going backwards in time. That is, the cost-to-go (value) function at stage $t = T, ..., 2$ is

$$Q_t(x_{t-1}, \omega) := \operatorname*{ess\,inf}_{x_t \in \mathcal{X}_t(x_{t-1}, \omega)} \Big\{ f_t(x_t, \omega) + \rho_{t+1}[Q_{t+1}(x_t, \omega)] \Big\}, \quad (92)$$

with the term $\rho_{T+1}[Q_{T+1}(x_T, \omega)]$, at stage $t = T$, omitted. At first stage the problem

$$\operatorname*{Min}_{x_1 \in \mathcal{X}_1} f_1(x_1) + \rho_2[Q_2(x_1, \omega)] \quad (93)$$

should be solved.

Consider policies $\bar{\pi} = (\bar{x}_1, ..., \bar{x}_T)$ with $\bar{x}_1$ being an optimal solution of problem (93) at the first stage and $\bar{x}_t = \bar{x}_t(\bar{x}_{t-1}, \omega)$, $t = 2, ..., T$, being an optimal solution of the problem

$$\underset{x_t \in \mathcal{X}_t(x_{t-1}, \omega)}{\text{Min}} f_t(x_t, \omega) + \rho_{t+1}[Q_{t+1}(x_t, \omega)]. \tag{94}$$

Provided that $\bar{\pi} \in \Pi$, we refer to such policy $\bar{\pi}$ as a *solution of dynamic programming equations* (92)–(93). If such minimizers $\bar{x}_t$ do exist, then the corresponding policy $\bar{\pi}$ is optimal for the reference problem, and moreover is time consistent.

If mappings ρ_t are not *strictly* monotone, there may exist optimal policies which are not time consistent and are not solutions of the dynamic programming equations. In particular, this can happen for max-type and nested AV@R systems.

Risk Averse Multistage Portfolio Selection

A nested formulation of multistage portfolio selection* can be written as (recall that e denotes vector of ones)

$$\begin{array}{ll} \text{Min} & \left\{\bar{\rho}(W_T) := \rho_1\Big[\cdots \rho_{T-1|W_{T-2}}\Big[\rho_{T|W_{T-1}}\left[W_T\right]\Big]\Big]\right\} \\ \text{s.t.} & W_{t+1} = \xi_{t+1}^{\mathsf{T}} x_t,\ e^{\mathsf{T}} x_t = W_t,\ x_t \geq 0,\ t = 0, \ldots, T-1. \end{array}$$

If we set $\rho_{t|W_{t-1}} := \mathbb{E}_{|W_{t-1}}$, $t = 1, ..., T$, then $\bar{\rho}(W_T) = \mathbb{E}[W_T]$. Now let, for example,

$$\rho_{t|W_{t-1}}(\cdot) := (1-\beta)\mathbb{E}_{|W_{t-1}}(\cdot) + \beta \mathsf{AV@R}_\alpha(\cdot|W_{t-1}),$$

$\alpha \in (0,1)$, $\beta \in (0,1)$. Suppose that the random process ξ_t is *stagewise independent*.

*Note that we formulated here the problem as a minimization rather than maximization problem.

Let us write dynamic programming equations. At the last stage we have to solve problem

$$\begin{array}{ll} \underset{x_{T-1}\geq 0, W_T}{\mathsf{Min}} & \rho_{T|W_{T-1}}[W_T] \\ \mathsf{s.t.} & W_T = \xi_T^{\mathsf{T}} x_{T-1},\ e^{\mathsf{T}} x_{T-1} = W_{T-1}. \end{array} \tag{95}$$

Since W_{T-1} is a function of $\xi_{[T-1]}$, by the stagewise independence we have that ξ_T, and hence W_T, are independent of W_{T-1}. Therefore we have in (95) that $\rho_{T|W_{T-1}}[W_T] = \rho[W_T]$, where

$$\rho(\cdot) = (1-\beta)\mathbb{E}(\cdot) + \beta \mathsf{AV@R}_\alpha(\cdot) \tag{96}$$

is the corresponding (unconditional) risk measure.

It follows by positive homogeneity of $\rho(\cdot)$ that the optimal value of (95) is $Q_{T-1}(W_{T-1}) = W_{T-1}\nu_{T-1}$, where ν_{T-1} is the optimal value of

$$\begin{array}{ll} \underset{x_{T-1}\geq 0, W_T}{\text{Min}} & \rho[W_T] \\ \text{s.t.} & W_T = \xi_T^{\mathsf{T}} x_{T-1},\ e^{\mathsf{T}} x_{T-1} = 1, \end{array} \tag{97}$$

and an optimal solution of (95) is $\bar{x}_{T-1}(W_{T-1}) = W_{T-1}x^*_{T-1}$, where x^*_{T-1} is an optimal solution of (97).

And so on we obtain that the optimal policy $\bar{x}_t(W_t)$ here is *myopic*. That is, $\bar{x}_t(W_t) = W_t x^*_t$, where x^*_t is an optimal solution of

$$\begin{array}{ll} \underset{x_t\geq 0, W_{t+1}}{\text{Min}} & \rho[W_{t+1}] \\ \text{s.t.} & W_{t+1} = \xi_{t+1}^{\mathsf{T}} x_t,\ e^{\mathsf{T}} x_t = 1. \end{array} \tag{98}$$

Note that the composite risk measure $\bar{\rho}$ is quite complicated here.

An alternative, multiperiod risk averse approach can be formulated as

$$\begin{array}{l} \text{Min } \Big\{\rho[W_T] = (1-\beta)\mathbb{E}[W_T] + \beta\Big(r + \alpha^{-1}\mathbb{E}[W_T - r]_+\Big)\Big\} \\ \text{s.t. } W_{t+1} = \xi_{t+1}^{\mathsf{T}} x_t,\ e^{\mathsf{T}} x_t = W_t,\ x_t \geq 0,\ t = 0, \ldots, T-1. \end{array} \tag{99}$$

Here $r \in \mathbb{R}$ is the (additional) first stage decision variable. After r is decided, at the first stage, the problem comes to minimizing $\mathbb{E}[U(W_T)]$ at the last stage, where

$$U(W) := (1-\beta)W + \beta\alpha^{-1}[W - r]_+$$

can be viewed as a disutility function.

It is possible to write dynamic programming equations for this problem. Note that here r is the *first stage* decision variable, the optimal policy is not myopic and the property of time consistency is not satisfied

Chance constrained problems.

Consider problem

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

where $\mathcal{X} \subset \mathbb{R}^n$ 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 $C(x, \xi) := \max_{1\leq i\leq q} C_i(x, \xi)$.

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 $x \in \mathcal{X}$. 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

$$\operatorname*{Min}_{x \in \mathcal{X}} f(x) \text{ subject to } C(x, \xi^j) \leq 0, \;\; j = 1, ..., N. \tag{100}$$

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

Theorem 3 (Calafiore, Campi, Garatti) *Suppose that the convexity condition holds and let $\bar{x}_N$ be an optimal solution of the above problem* (100)*. Then*

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

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

$$\operatorname*{Min}_{x\in\mathcal{X}} f(x) \text{ s.t. } \widehat{p}_N(x) \leq \gamma. \tag{101}$$

Note that we can use the significance level γ, of the SAA problem (101), 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 $\bar{x} \in \mathcal{X}$ 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 $k > Np$,

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

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 $\bar{x} \in \mathcal{X}$ 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 $\widehat{v}_N^m$, $m = 1,\ldots,M$. That is, solve M times the corresponding SAA problem at the significance level γ.

We can view $\widehat{v}_N^m$, $m = 1, \ldots, M$, as an iid sample of the random variable $\widehat{v}_N$, where $\widehat{v}_N$ is the optimal value of the respective SAA problem at significance level γ. Next we rearrange the calculated optimal values in the nondecreasing order $\widehat{v}_N^{(1)} \leq \cdots \leq \widehat{v}_N^{(M)}$. We use the random quantity $\widehat{v}_N^{(L)}$ as a lower bound of the true optimal value v^0. It is possible to show that with probability at least $1 - \beta$, the random quantity $\widehat{v}_N^{(L)}$ is below the true optimal value v^0, i.e., $\widehat{v}_N^{(L)}$ is indeed a lower bound of the true optimal value with confidence at least $1 - \beta$.

Convex approximations

Consider chance constraint:

$$\text{Prob}\Big\{C(x,\xi) > \tau\Big\} \leq \alpha. \tag{102}$$

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 (see (72))

$$\inf_{t>0} \mathbb{E}[\psi(tZ_x)] \leq \alpha \tag{103}$$

is a conservative approximation of the chance constraint (102). The "best" choice $\psi(z) := [1+z]_+$ leads to the AV@R approximation (see (73)).

Bernstein approximation

Take $\psi(t) := e^t$. Suppose that: (i) The components ξ_j, $j = 1, ..., d$, of random vector ξ are independent of each other random variables. (ii) The moment generating functions $M_j(t) = \mathbb{E}[e^{t\xi_j}]$, $j = 1, ..., d$, are finite valued for all $t \in \mathbb{R}$ and are efficiently computable. (iii) The constraint function $C(x, \xi)$ is affine in ξ:

$$C(x, \xi) = g_0(x) + \sum_{j=1}^{d} \xi_j g_j(x).$$

For $z = (z_0, z_1, ..., z_d)$ and $\Lambda_j(z_j) = \log M_j(z_j)$ consider function

$$\Phi(z) := \log\left(\mathbb{E}\Big[\exp\left\{z_0 + \textstyle\sum_{j=1}^{d} \xi_j z_j\right\}\Big]\right) = z_0 + \textstyle\sum_{j=1}^{d} \Lambda_j(z_j).$$

Note that function $\Lambda_j(\cdot)$ is convex and monotonically nondecreasing if $\xi_j \geq 0$ w.p.1.

The following problem is the conservative approximation of the original chance constrained problem for $\psi(t) = e^t$:

$$\begin{array}{ll} \underset{x\in\mathcal{X},\, t>0}{\text{Min}} & f(x) \\ \text{s.t.} & \underbrace{g_0(x) + \sum_{j=1}^{d} t\Lambda_j\Big(t^{-1}g_j(x)\Big)}_{t\Phi\Big(t^{-1}(g_0(x),\ldots,g_d(x))\Big)} - t\log\alpha \leq 0. \end{array}$$

This problem is convex if the set $\mathcal{X}$ is convex, $f(\cdot)$ and $g_0(\cdot)$ are convex, and each $g_j(\cdot)$ for $j = 1, ..., d$, is either convex and $\xi_j \geq 0$ w.p.1, or $g_j(\cdot)$ is affine.

Approximate dynamic programming

Basic idea is to approximate the cost-to-go functions by a class of computationally manageable functions. Consider linear multistage problem and assume stagewise independence of the data process. Since cost-to-go functions $\mathcal{Q}_t(\cdot)$ are convex it is natural to approximate these functions by piecewise linear functions given by maximum of cutting hyperplanes.

Stochastic Dual Dynamic Programming (SDDP) method (Pereira and Pinto, 1991).

For trial decisions $\bar{x}_t$, $t = 1, ..., T-1$, at the backward step of the SDDP algorithm, piecewise linear approximations $\mathfrak{Q}_t(\cdot)$ of the cost-to-go functions $\mathcal{Q}_t(\cdot)$ are constructed by solving problems

$$\operatorname*{Min}_{x_t \in \mathbb{R}^{n_t}} (c_t^j)^{\mathsf{T}} x_t + \mathfrak{Q}_{t+1}(x_t) \text{ s.t. } B_t^j \bar{x}_{t-1} + A_t^j x_t = b_t^j, \ x_t \geq 0,$$

$j = 1, ..., N_t$, and their duals, going backward in time $t = T, ..., 1$.

Denote by v^0 and $\widehat{v}_N$ the respective optimal values of the true and SAA problems.

By construction

$$\mathcal{Q}_t(\cdot) \geq \mathfrak{Q}_t(\cdot),\ t = 2, ..., T.$$

Therefore the optimal value of

$$\operatorname*{Min}_{x_1 \in \mathbb{R}^{n_1}} c_1^{\mathsf{T}} x_1 + \mathfrak{Q}_2(x_1) \text{ s.t. } A_1 x_1 = b_1,\ x_1 \geq 0$$

gives a lower bound for the optimal value $\widehat{v}_N$ of the SAA problem.

We also have that

$$v^0 \geq \mathbb{E}[\widehat{v}_N].$$

Therefore *on average* $\widehat{v}_N$ is also a lower bound for the optimal value of the true problem.

The approximate cost-to-go functions $\mathfrak{Q}_2, ..., \mathfrak{Q}_T$ and a feasible first stage solution $\bar{x}_1$ define a feasible policy. That is for a realization (sample path) $\xi_1, ..., \xi_T$ of the data process, $\bar{x}_t = \bar{x}_t(\xi_{[t]})$ are computed recursively in $t = 2, ..., T$ as a solution of

$$\underset{x_t}{\text{Min}}\, c_t^\top x_t + \mathfrak{Q}_{t+1}(x_t) \text{ s.t. } B_t \bar{x}_{t-1} + A_t x_t \leq b_t.$$

In the *forward step* of the SDDP algorithm M sample paths (scenarios) are generated and the corresponding $\bar{x}_t$, $t = 2, ..., T$, are used as trial points in the next iteration of the backward step.

It is essential for convergence of this algorithm that at each iteration in the forward step the paths (scenarios) are *resampled*, i.e., generated independently of the previous iteration.

Note that the functions $\mathfrak{Q}_2, ..., \mathfrak{Q}_T$ and $\bar{x}_1$ define a feasible policy also for the *true* problem.

Convergence of the SDDP algorithm

It is possible to show that, under mild regularity conditions, the SDDP algorithm converges as the number of iterations go to infinity. That is, the computed optimal values and generated policies converge w.p.1 to their counterparts of the considered SAA problem. However, the convergence can be very slow and one should take such mathematical proofs very cautiously.

Moreover, it should be remembered that the SAA problem is just an approximation of the "true" problem. It is possible to show that, in a certain probabilistic sense, the SAA problem converges to the "true" problem as all sample sizes N_t, $t = 2, ..., T$, tend to infinity.

Stopping criteria

The policy value $\mathbb{E}\left[\sum_{t=1}^{T} c_t^{\mathsf{T}} \bar{x}_t(\xi_{[t]})\right]$ can be estimated in the forward step of the algorithm. That is, let $\xi_2^i, ..., \xi_T^i$, $i = 1, ..., M$, be sample paths (scenarios) generated at a current iteration of the forward step, and

$$\vartheta_i := \sum_{t=1}^{T} (c_t^i)^{\mathsf{T}} \bar{x}_t^i, \ i = 1, ..., M,$$

be the corresponding cost values. Then $\mathbb{E}[\vartheta_i] = \mathbb{E}\left[\sum_{t=1}^{T} c_t^{\mathsf{T}} \bar{x}_t(\xi_{[t]}^i)\right]$, and hence

$$\bar{\vartheta} = \frac{1}{M} \sum_{i=1}^{M} \vartheta_i$$

gives an unbiased estimate of the policy value.

Also

$$\widehat{\sigma}^2 = \frac{1}{M-1}\sum_{i=1}^{M}(\vartheta_i - \bar{\vartheta})^2$$

estimates variance of the sample $\vartheta_1, ..., \vartheta_M$. Hence

$$\bar{\vartheta} + z_\alpha \widehat{\sigma}/\sqrt{M}$$

gives an *upper* bound for the policy value with confidence of about $100(1-\alpha)\%$. Here z_α is the corresponding critical value.

At the same time this gives an upper bound for the optimal value of the corresponding multistage problem, SAA or the "true" problem depending from what data process the random scenarios were generated.

Typical example of behavior of the lower and upper bounds produced by the SDDP algorithm for an SAA problem

Risk averse approach

How to control risk, i.e., to reduce chances of extreme costs, at every stage of the time process.

Value-at-Risk of a random outcome (variable) Z at level $\alpha \in (0,1)$:

$$\mathsf{V@R}_\alpha(Z) = \inf\{t : F_Z(t) \geq 1-\alpha\},$$

where $F_Z(t) = \mathsf{Prob}(Z \leq t)$ is the cdf of Z. That is, $\mathsf{V@R}_\alpha(Z)$ is the $(1-\alpha)$-quantile of the distribution of Z.

Note that $\mathsf{V@R}_\alpha(Z) \leq c$ is equivalent to $\mathsf{Prob}(Z > c) \leq \alpha$. Therefore it could be a natural approach to impose constraints (chance constraints) of $\mathsf{V@R}_\alpha(Z) \leq c$ for $Z = cost$, chosen constant c and significance level α at every stage of the process.

There are two problems with such approach. It is difficult to handle chance constraints numerically and could lead to infeasibility problems.

Average Value-at-Risk (also called *Conditional Value-at-Risk*)

$$\mathsf{AV@R}_\alpha(Z) = \inf_{t\in\mathbb{R}} \left\{ t + \alpha^{-1}\mathbb{E}[Z-t]_+ \right\}$$

Note that the minimum in the above is attained at $t^* = \mathsf{V@R}_\alpha(Z)$. If the cdf $F_Z(z)$ is continuous, then

$$\mathsf{AV@R}_\alpha(Z) = \mathbb{E}\Big[Z | Z \geq \mathsf{V@R}_\alpha(Z) \Big].$$

It follows that $\mathsf{AV@R}_\alpha(Z) \geq \mathsf{V@R}_\alpha(Z)$. Therefore the constraint $\mathsf{AV@R}_\alpha(Z) \leq c$ is a conservative approximation of the chance constraint $\mathsf{V@R}_\alpha(Z) \leq c$.

In the problem of minimizing expected cost $\mathbb{E}[Z]$ subject to the constraint $\mathsf{AV@R}_\alpha(Z) \leq c$, we impose an infinite penalty for violating this constraint. This could result in infeasibility of the obtained problem. Instead we can impose a finite penalty and consider problem of minimization of $\mathbb{E}[Z] + \kappa \mathsf{AV@R}_\alpha(Z)$ for some constant $\kappa > 0$. Note that this is equivalent to minimization of $\rho(Z)$, where

$$\rho(Z) = (1 - \lambda)\mathbb{E}[Z] + \lambda \mathsf{AV@R}_\alpha(Z)$$

for $\lambda \in (0, 1)$ and $\kappa = \frac{\lambda}{1-\lambda}$.

This leads to the following (nested) formulation of risk averse multistage problem.

$$\begin{aligned}\operatorname*{Min}_{A_1x_1\leq b_1} \quad & c_1^{\mathsf{T}}x_1 + \rho_{2|\xi_1}\Big[\inf_{\substack{B_2x_1+A_2x_2=b_2\\ x_2\geq 0}} c_2^{\mathsf{T}}x_2 + \ldots \\ & +\rho_{T-1|\xi_{[T-2]}}\Big[\inf_{\substack{B_{T-1}x_{T-2}+A_{T-1}x_{T-1}=b_{T-1}\\ x_{T-1}\geq 0}} c_{T-1}^{\mathsf{T}}x_{T-1} \\ & +\rho_{T|\xi_{[T-1]}}\big[\inf_{\substack{B_Tx_{T-1}+A_Tx_T=b_T\\ x_T\geq 0}} c_T^{\mathsf{T}}x_T\big]\Big]\Big],\end{aligned}$$

with

$$\rho_{t|\xi_{[t]}}(\cdot) := (1-\lambda)\mathbb{E}_{|\xi_{[t]}}[\cdot] + \lambda \mathsf{AV@R}_{\alpha|\xi_{[t]}}(\cdot)$$

being conditional analogue of $\rho(\cdot)$.

We can write the risk averse multistage programming problem as

$$\begin{array}{cl} \underset{x_1, x_2(\cdot), \ldots, x_T(\cdot)}{\text{Min}} & \bar{\rho}\Big[F_1(x_1) + F_2(x_2(\xi_{[2]}), \xi_2) + \cdots + F_T\left(x_T(\xi_{[T]}), \xi_T\right)\Big] \\ \text{s.t.} & x_1 \in \mathcal{X}_1,\ x_t(\xi_{[t]}) \in \mathcal{X}_t(x_{t-1}(\xi_{[t-1]}), \xi_t),\ t = 2, \ldots, T, \end{array}$$

where $F_t(x_t, \xi_t) = c_t^{\mathsf{T}} x_t$ and

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

$$\begin{aligned} \bar{\rho}(Z_1 + \ldots + Z_T) &= \rho_{|\xi_1}\Big(\rho_{|\xi_{[2]}}\Big(\cdots \rho_{|\xi_{[T-1]}}(Z_1 + \ldots + Z_T)\Big)\Big) \\ &= Z_1 + \rho_{|\xi_1}\Big(Z_2 + \rho_{|\xi_{[2]}}\Big(+ \cdots \rho_{|\xi_{[T-1]}}(Z_T)\Big)\Big) \end{aligned}$$

is the corresponding composite risk measure. The optimization is performed over (nonanticipative) policies $x_1, x_2(\xi_{[2]}), \ldots, x_T(\xi_{[T]})$ satisfying the feasibility constraints.

With some modifications the SDDP algorithm can be applied to the above multistage problem. Assuming the stagewise independence, the dynamic programming equations for the adaptive risk averse problem take the form

$$Q_t\left(x_{t-1},\xi_t\right) = \inf_{x_t\in\mathbb{R}^{n_t}} \left\{c_t^{\mathsf{T}} x_t + \mathcal{Q}_{t+1}(x_t) : B_t x_{t-1} + A_t x_t = b_t,\ x_t \geq 0\right\},$$

$t = T, ..., 2$, where $\mathcal{Q}_{T+1}(\cdot) \equiv 0$ and

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

Since ξ_{t+1} is independent of $\xi_{[t]}$, the cost-to-go functions $\mathcal{Q}_{t+1}\left(x_t\right)$ do not depend on the data process. In order to apply the backward step of the SDDP algorithm we only need to know how to compute subgradients of the cost-to-go functions.

The value of this problem corresponds to the total objective

$$\begin{aligned}\bar{\rho}(Z_1 + ... + Z_T) &= \rho_{|\xi_{[1]}}\Big(\cdots \rho_{|\xi_{[T-1]}}(Z_1 + ... + Z_T)\Big) \\ &= Z_1 + \rho_{|\xi_{[1]}}\Big(Z_2 + \cdots + \rho_{|\xi_{[T-1]}}(Z_T)\Big)\end{aligned}$$

The dynamic programming equations of the risk averse formulation of the SAA program take the form

$$Q_t^j(x_{t-1}) = \inf_{x_t}\Big\{(c_t^j)^{\mathsf{T}}x_t + \mathcal{Q}_{t+1}(x_t) : B_t^j x_{t-1} + A_t^j x_t = b_t^j,\ x_t \geq 0\Big\},$$

$j = 1, ..., N_t$, $t = T, \ldots, 2$, and

$$\mathcal{Q}_{t+1}(x_t) = \rho\left(Q_{t+1}^1(x_t), ..., Q_{t+1}^{N_{t+1}}(x_t)\right),$$

with $\mathcal{Q}_{T+1}(\cdot) \equiv 0$ and the first stage problem

$$\operatorname*{Min}_{A_1 x_1 \leq b_1} c_1^{\mathsf{T}}x_1 + \rho\left(Q_2^1(x_1), ..., Q_2^{N_2}(x_1)\right).$$

For $\rho(\cdot) = (1-\lambda)\mathbb{E}[\cdot] + \lambda \mathsf{AV@R}_\alpha(\cdot)$, and $(Z_1, ..., Z_N) = \left(Q^1_{t+1}(x_t), ..., Q^N_{t+1}(x_t)\right)$ we have that

$$\mathcal{Q}_{t+1}(x_t) = \frac{1-\lambda}{N_{t+1}} \sum_{j=1}^{N_{t+1}} Z_j + \lambda \left(Z_\iota + \frac{1}{\alpha N_{t+1}} \sum_{j: Z_j > Z_\iota} \left[Z_j - Z_\iota \right] \right),$$

where Z_ι is the $(1-\alpha)$-quantile of $Z_1, ..., Z_{N_{t+1}}$. Note that if $N_{t+1} < (1-\alpha)^{-1}$, then $Z_\iota = \max\{Z_1, ..., Z_{N_{t+1}}\}$.

A subgradient of $\mathcal{Q}_{t+1}(x_t)$ is given by

$$\nabla \mathcal{Q}_{t+1}(x_t) = \frac{1-\lambda}{N} \sum_{j=1}^{N_{t+1}} \nabla Q^j_{t+1}(x_t) +$$

$$\lambda \left(\nabla Q^\iota_{t+1}(x_t) + \frac{1}{\alpha N_{t+1}} \sum_{j: Z_j > Z_\iota} \left[\nabla Q^j_{t+1}(x_t) - \nabla Q^\iota_{t+1}(x_t) \right] \right).$$

These formulas allow construction of cuts in the backward step of the SDDP algorithm. In the forward step trial points are generated in the same way as in the risk neutral case.

Remarks

Unfortunately there is no easy way for evaluating value of the risk objective of generated policies, and hence constructing a corresponding upper bound. Some suggestions were made in the recent literature. However, in larger problems the optimality gap (between the upper and lower bounds) never approaches zero in any realistic time. Therefore stopping criteria based on stabilization of the lower bound (and may be optimal solutions) could be reasonable. Also it should be remembered that there is no intuitive interpretation for the risk objective $\bar{\rho}(cost)$ of the total cost. Rather the goal is to control risk at every stage of the process.