PROGRESSIVE HEDGING IN STOCHASTIC OPTIMIZATION

Terry Rockafellar
University of Washington, Seattle

AMSO Optimise
Melbourne, Australia
18 June 2017

Decisions and Observations

- Decisions can take advantage of information known in advance
- Information can come from observations and evolve in time
- Observations may come from realizations of random variables
- Optimization of decisions most cope with uncertainty and risk
- Evolution of information leads to problems of high dimension

Scenario model for information

Suppose that information comes in stages $k = 1, 2, \ldots, N$.
Let $\xi_k \in \Xi_k$ stand for what is observed in stage k. Then

$$\xi = (\xi_1, \xi_2, \ldots, \xi_N) \text{ is a } \textbf{scenario}$$

Let $\Xi \subset \Xi_1 \times \cdots \times \Xi_N$ be the set of scenarios ξ to be considered, and give it the structure of a probability space

How might decisions and their optimization fit with this?

Uncertain Costs and Measures of Risk

Costs: not necessarily money, but low values preferred to high

Uncertainty: a cost can depend on the scenario: $X(\xi)$

$\Longrightarrow$ then X can be treated technically as a random variable

Risk measures: expressions $\mathcal{R}$ that assign to X a risk level $\mathcal{R}(X)$

Examples of risk measures

- $\mathcal{R}(X) = E[X] = \mu(X)$, expected value or mean value
- $\mathcal{R}(X) = \mu(X) + \lambda\sigma(X)$ with $\lambda > 0$ giving risk margin
- $\mathcal{R}(X) = \sup X$, essential supremum: worst-case risk
- $\mathcal{R}(X) = \mathrm{VaR}_\alpha(X)$, value-at-risk, α-quantile
- $\mathcal{R}(X) = \mathrm{CVaR}_\alpha(X)$, cond. value=at-risk, α-superquantile

Coherency: axiomatized properties for $\mathcal{R}$ to be "good"

above, only #1, #3, and #5 are coherent

The Basics for a Single Stage of Information

Decisions: $x \in R^n$ **Observations:** $\xi \in \Xi$

Cost objective: $g(x,\xi)$ **Feasible set:** $C(\xi) \subset R^n$

Observation first? $\Longrightarrow$ the decision can depend on ξ

get $x(\xi)$ for each ξ by minimizing $g(\cdot,\xi)$ over $C(\xi)$

Decision first? $\Longrightarrow$ the decision can't depend on ξ

choosing x just yields a random variable $g(x,\cdot) : \xi \to g(x,\xi)$

minimize $\mathcal{R}(\,g(x,\cdot)\,)$ in $x \in \cap_\xi C(\xi)$ for a risk measure $\mathcal{R}$

$g(x,\xi)$ convex in x, $\mathcal{R}$ coherent $\Rightarrow \mathcal{R}(\,g(x,\cdot)\,)$ convex in x

Possible mixture: the decision x could have two components

- a first part which can't depend on ξ
- a second part which can depend on ξ $\rightarrow$ "recourse action"

Example: Two-Stage Stochastic Linear Programming

First stage decision: $x_1 \in \boldsymbol{R}^{n_1}$

Second stage decision: $x_2(\xi) \in \boldsymbol{R}^{n_2}$

Constraints: $x_1 \geq 0,\ A_{11}x_1 \leq b_1$
$x_2 \geq 0,\ A_{22}(\xi)x_2(\xi) \leq b_1(\xi) - A_{21}(\xi)x_1$

Objective: minimize $c_1 \cdot x_1 + E_\xi[\, c_2(\xi) \cdot x_2(\xi)\,]$

Reduced problem:

minimize $c_1 \cdot x_1 + h(x_1)$ subject to $x_1 \geq 0,\ A_{11}x_1 \leq b_1,$ where

$$h(x_1) = E_\xi\Big[\min\Big\{ c_2(\xi) \cdot x_2 \,\Big|\, x_2 \geq 0,\ A_{22}(\xi)x_2 \leq b_2(\xi) - A_{21}(\xi)x_1 \Big\}\Big]$$

- the expectation in the objective signals **risk neutrality**
- replacing it by a risk measure $\mathcal{R}$ can capture **risk aversity**

Multistage Stochastic Programming More Genenarlly

Interplay between decisions and observations in N stages:

$$x_1,\ \xi_1,\ x_2,\ \xi_2, \ldots, x_N,\ \xi_N \quad \text{with} \quad x_k \in \boldsymbol{R}^{n_k},\ \xi_k \in \Xi_k$$

$$x = (x_1, \ldots, x_N) \in \boldsymbol{R}^n = \boldsymbol{R}^{n_1} \times \cdots \times \boldsymbol{R}^{n_N}$$

$$\xi = (\xi_1, \ldots, \xi_N) \in \Xi \ \subset \Xi_1 \times \cdots \Xi_N \quad \textbf{scenarios}$$

Nonanticipativity of decisions

x_k can respond to $\xi_1, \ldots, \xi_{k-1}$ but not to $\xi_k, \ldots, \xi_N$:

$$x(\xi) = (x_1, x_2(\xi_1), x_3(\xi_1, \xi_2), \ldots, x_N(\xi_1, \xi_2, \ldots, \xi_{N-1}))$$

Simplifying assumptions for this talk:

- the scenario space Ξ has only finitely many elements ξ
- each scenario $\xi \in \Xi$ has known probability $p(\xi) > 0$

$\longrightarrow$ Ξ is a discrete probability space

Problem Formulation in this Setting — With Convexity

Problem ingredients: for each scenario $\xi \in \Xi$, let

$$C(\xi) = \text{nonempty closed convex set in } R^n$$
$$g(x,\xi) = \text{continuous convex function of } x \in C(\xi)$$

Traditional stochastic programming — risk neutral

Over all nonanticipative decision choices

$$x(\cdot) : \xi \to (x_1, x_2(\xi_1), x_3(\xi_1,\xi_2), \ldots, x_N(\xi_1,\xi_2,\ldots,\xi_{N-1}))$$

such that $x(\xi) \in C(\xi)$ for all scenarios $\xi \in \Xi$ minimize

$$\mathcal{G}(x(\cdot)) = E_\xi[\, g(x(\xi),\xi)\,] = \sum_{\xi\in\Xi} p(\xi) g(x(\xi),\xi)$$

Risk-averse objective as an alternative:

$\mathcal{G}(x(\cdot)) = \mathrm{CVaR}_\alpha(G(x(\cdot))$ for the r.v. $G(x(\cdot)) : \xi \to g(x(\xi),\xi)$

$\mathrm{CVaR}_\alpha(X) =$ conditional expectation in upper α-tail of X

Nonanticipativity Posed as a Subspace Constraint

Goal: introducing "Lagrange multiplier" for nonanticipativity

$\mathcal{L}$ = **all** functions from scenario space Ξ to decision space R^n,

$$x(\cdot) : \xi = (\xi_1, \ldots, \xi_N) \mapsto x(\xi) = (x_1(\xi), \ldots, x_N(\xi))$$

$$\Xi \;\subset \Xi_1 \times \cdots \Xi_N, \qquad R^n = R^{n_1} \times \cdots \times R^{n_N}$$

Expectation inner product: for $x(\cdot),\ w(\cdot) \in \mathcal{L}$

$$\langle x(\cdot), w(\cdot)\rangle = E_\xi[\,x(\xi){\cdot}w(\xi)\,] = \sum_{\xi\in\Xi} p(\xi) \sum_{k=1}^N x_k(\xi){\cdot}w_k(\xi)$$

Nonanticipativity subspace: with $\xi = (\xi_1, \ldots, \xi_{k-1}, \xi_k, \ldots, \xi_N)$

$$\mathcal{N} = \{x(\cdot) \in \mathcal{L} \,|\, x_k(\xi) \text{ depends only on } \xi_1, \ldots, \xi_{k-1}\}$$

$$\longrightarrow \quad x(\cdot) \text{ is nonanticipative} \iff x(\cdot) \in \mathcal{N}$$

Orthogonal complement:

$$\mathcal{N}^\perp = \{w(\cdot) \in \mathcal{L} \,|\, E_{\xi_k,\ldots,\xi_N}[\,w_k(\xi_1, \ldots, \xi_{k-1}, \xi_k \ldots, \xi_N)\,] = 0\,\}$$

Projection Tool for Achieving Nonanticipativity

Let $\mathcal{P}$ = projection onto $\mathcal{N}$, so that $\mathcal{I} - \mathcal{P}$ = projection onto $\mathcal{N}^{\perp}$

$x(\cdot)$ nonanticipative $\iff$ $[\mathcal{I} - \mathcal{P}](x(\cdot)) = 0$ (linear constraint)

Execution of projections relative to the information structure

- scenarios $\xi = (\xi_1, \ldots, \xi_N)$ and $\xi' = (\xi'_1, \ldots, \xi'_N)$ are information-equivalent at stage k if
$$(\xi_1, \ldots, \xi_{k-1}) = (\xi'_1, \ldots, \xi'_{k-1})$$
- let $A_k(\xi)$ = kth-stage equivalence class containing ξ
- then $x(\cdot) = \mathcal{P}(\hat{x}(\cdot))$ has its kth-stage component given by
$$x_k(\xi) = \frac{\sum_{\xi' \in A_k(\xi)} p(\xi')\hat{x}_k(\xi')}{\sum_{\xi' \in A_k(\xi)} p(\xi')} \quad \text{(conditional expectation)}$$

Stochastic Programming Problem Reformulated

minimize $\mathcal{G}(x(\cdot))$ over all functions $x(\cdot) \in \mathcal{C} \cap \mathcal{N} \subset \mathcal{L}$

$$\mathcal{N} = \text{the nonanticipativity subspace of } \mathcal{L}$$
$$\mathcal{C} = \left\{ x(\cdot) \in \mathcal{L} \,\middle|\, x(\xi) \in C(\xi) \subset \mathbb{R}^n \ \forall \xi \in \Xi \right\}$$
$$\mathcal{G}(x(\cdot)) = E_\xi[g(x(\xi), \xi)] = \textstyle\sum_{\xi \in \Xi} p(\xi) g(x(\xi), \xi)$$

$\mathcal{C} \subset \mathcal{L}$ is closed convex, $\mathcal{G} : \mathcal{C} \to \mathbb{R}$ is continuous convex

Computational strategy:

- introduce information cost "multipliers" $w(\cdot) \in \mathcal{N}^\perp$
- use these to relax the constraint of nonanticipativity
- thereby decompose the stochastic programming problem iteratively into solving "deterministic" subproblems

Progressive Hedging Approach (Rock.& Wets 1991)

- In iterations $\nu = 1, 2, \ldots$ solve "hindsight" problems for the separate scenarios ξ in which the cost $g(x,\xi)$ is modified to

$$g^\nu(x,\xi) + w^\nu(\xi){\cdot}x + \frac{r}{2}||x - x^\nu(\xi)||^2 \quad (r = \text{parameter} > 0)$$

with respect to the current $x^\nu(\cdot) \in \mathcal{N}$ and $w^\nu(\cdot) \in \mathcal{N}^\perp$

- This yields $\hat{x}^\nu(\xi)$ for each ξ, but the response function $\hat{x}^\nu(\cdot)$ won't be in $\mathcal{N}$; restore nonanticpativity by projecting $\hat{x}(\cdot)$ onto $\mathcal{N}$ and then generate an update for the information costs in $\mathcal{N}^\perp$

- the risk-neutral case of $E_\xi[g^\nu(x(\xi),\xi)] = \sum_{\xi\in\Xi} p(\xi) g(x(\xi),\xi)$ supports this decomposition into separate scenario subproblems
- the risk-averse case with CVaR_α has **no separability directly**, but **separability can be achieved,** as will be explained later

Progressive Hedging in Stochastic Programming

Algorithm statement in the risk-neutral case with parameter $r > 0$

Having $x^\nu(\cdot) \in \mathcal{N}$ and $w^\nu(\cdot) \in \mathcal{M}$, get $\hat{x}^\nu(\cdot) \in \mathcal{L}$ by

$$\hat{x}^\nu(\xi) = \operatorname*{argmin}_{x \in C(\xi)} \Big\{ g(x,\xi) + x{\cdot}w^\nu(\xi) + \frac{r}{2}||x - x^\nu(\xi)||^2 \Big\}$$

Then get $x^{\nu+1}(\cdot) \in \mathcal{N}$ and $w^{\nu+1}(\cdot) \in \mathcal{N}^\perp$ by:

$$x^{\nu+1}(\cdot) = \mathcal{P}(\hat{x}^\nu(\cdot)), \qquad w^{\nu+1}(\cdot) = w^\nu(\cdot) + r[\hat{x}^\nu(\cdot) - x^{\nu+1}(\cdot)]$$

Convergence theorem — when a solution pair $x(\cdot)$, $w(\cdot)$, exists

The sequence $\{(x^\nu(\cdot), w^\nu(\cdot))\}_{\nu=1}^\infty$ generated by the algorithm will always converge to a particular solution pair $(x^*(\cdot), w^*(\cdot))$, with

$$\begin{aligned} &||x^{\nu+1}(\cdot) - x^*(\cdot)||^2 + \tfrac{1}{r^2}||w^{\nu+1}(\cdot) - w^*(\cdot)||^2 \\ &\qquad < \; ||x^\nu(\cdot) - x^*(\cdot)||^2 + \tfrac{1}{r^2}||w^\nu(\cdot) - w^*(\cdot)||^2 \end{aligned}$$

Adaptation of Progressive Hedging to Risk-Averse Case

CVaR Minimization formula: Rock.& Uryasev (2000, 2002)

$$\mathrm{CVaR}_\alpha(X) = \min_{z\in R}\left\{z + \tfrac{1}{1-\alpha}E[\max\{0, X - z\}]\right\}$$

Consequence: for the random variable $G(x(\cdot)) : \xi \to g(x(\xi), \xi)$

$$\mathrm{CVaR}_\alpha(G(x(\cdot)) = \min_{z\in R}\left\{z + \tfrac{1}{1-\alpha}E_\xi[\max\{0, g(x(\xi), \xi) - z\}]\right\}$$

Risk-averse stochastic programming problem, expanded

minimizing $\mathrm{CVaR}_\alpha(G(x(\cdot))$ over $x(\cdot) \in \mathcal{C} \cup \mathcal{N}$ is equivalent to minimizing $E_\xi[\,h(z, x(\xi), \xi)\,]$ over $z \in R$, $x(\cdot) \in \mathcal{C} \cup \mathcal{N}$, where

$$h(z, x(\xi), \xi) = z + \tfrac{1}{1-\alpha}\max\{0, g(x(\xi), \xi) - z\}$$

- Incorporate z within $x(\cdot)$ as an extra first-stage variable
- Then just apply the risk-neutral version of the progressive hedging algorithm with h taking the place of g

Some References

[1] R.T. Rockafellar and R.J-B Wets (1991) "Scenarios and policy aggregation in optimization under uncertainty," *Mathematics of Operations Research* 16, 119–147.

[2] R.T. Rockafellar and S. Uryasev (2002) "Conditional value-at-risk for general loss distributions," *Journal of Banking and Finance* 26, 1443–1471.

[3] R.T. Rockafellar (2017) "Solving stochastic programming problems with risk measures by progressive hedging" *Mathematical Programming B*, accepted

[4] R.T. Rockafellar and Jie Sun (2017) "Solving monotone stochastic variational inequalities and complementarity problems by progressive hedging," *Mathematical Programming B*, accepted

website: www.math.washington.edu/∼rtr/mypage.html