STOCHASTIC PROGRAMMING MODESLS: WAIT-AND-SEE versus HERE-AND-NOW [*]

Roger J-B Wets

Department of Mathematics
University of California, Davis
rjbwets@ucdavis.edu

Abstract. We introduce a number of stochastic programming models via examples and then proceed to derive one of the fundamental theorems in the field that brings to the fore the constrast between wait-and-see and here-and-now formulations.

Keywords: stochastic programming, wait-and-see, here-and-now,
measurable mappings, nonanticipativity

AMS Classification: 90C15

Date: September 17, 2000

---

[*] Research supported in part by a grant of the National Science Foundation

1 Introduction

For now, let's formulate optimization problems as:

$$\inf_{x\in X} F(x)$$

where

- X is a linear space, e.g. $I\!R^n$, $C([0,1];I\!R^n)$, etc.;
- $F : X \to \overline{I\!R}$ is an extended real-valued functions;
- $\operatorname{dom} F := \{x \in X \mid F(x) < \infty\}$, the *(effective) domain* of F correspond to the set of x's that are feasible, i.e., satisfy the constraints of the problem.

One refers to F as having a *minimum* if the infimum is attained for some $\bar{x}$, and this $\bar{x}$ is then said to be a *minimizer* of F.

A standard *linear* or *nonlinear programming problem*:

$$\begin{array}{ll} \min\limits_{x\in \boldsymbol{R}^n} & f_0(x) \\ \text{so that} & f_i(x) \le 0, \quad i = 1,\dots,s, \\ & f_i(x) = 0, \quad i = s+1,\dots,m, \\ & \qquad x \in C \subset I\!R^n \end{array}$$

can be cast in our initial formulation, simply define

$$F(x) = \begin{cases} f_0(x) & \text{if } f_i(x) \le 0,\ i = 1,\dots,s, \\ & \qquad f_i(x) = 0,\ i = s+1,\dots,m, \quad x \in C; \\ \infty & \text{otherwise.} \end{cases}$$

A standard *calculus of variation problem*:

$$\min_{x\in \mathrm{H}^1_0([0,1];\boldsymbol{R}^n)} F(x) = \int_0^1 f(t, x(t), \dot{x}(t))\, dt$$

is already in this form.

In this last instance, the decision is time dependent and requires the minimization of an integral functional. Stochastic optimization problems are of this type but involve a *nonanticipativity* constraint which restrict the choice of the decisions at any one time to the information that has been revealed so far. In our first model, the nonanticipativity restriction will only be implicit in the formulation of the problem. But we are going to see that bringing it to the fore will provide us with a deeper understanding of optimality conditions and, eventually, has lead to the development of a certain class of solution procedures.

2 Stochastic programming models

Our first model could be given the following interpretation. A decision $x \in I\!R^n$ must be selected before the value ξ of a random event, modeled by the random variable $\boldsymbol{\xi}$, is observed. The 'cost' of the decision is then $f(\xi, x)$ which possibly takes on the value ∞ if the pair (ξ, x) is not feasible (acceptable). The distribution μ of $\boldsymbol{\xi}$ is known. It's a probability measure defined on $(\Xi, \mathcal{S})$ with $\Xi \subset I\!R^N$ and $\mathcal{S}$ a σ-field. The stochastic optimization problem is

$$\inf_{x \in \boldsymbol{R}^n} Ef(x) = \int_\Xi f(\xi, x)\, \mu(d\xi) = E\{f(\boldsymbol{\xi}, x)\}$$

with $E\{\cdot\}$ denoting expectation. In order for the function Ef to be well-defined, one requires that for all x, the function $\xi \mapsto f(\xi, x)$ is measurable and the following definition of the integral is used to avoid possible ambiguities: For any extended real-valued measurable function $\psi : \Xi \to \overline{I\!R}$

$$\int_\Xi \psi(\xi)\, \mu(d\xi) = \int_\Xi \max\,[0, \psi(\xi)]\, \mu(d\xi) - \int_\Xi \max\,[0, -\psi(\xi)]\, \mu(d\xi)$$

with the convention that the value is ∞ if both integrals on the right are divergent.

2.1 Example (the newsboy problem)**.** *The newsboy problem is a one-period model in which a firm (newsboy) orders a nonnegative quantity $x \in I\!R_+$ of perishable items (newspapers) for resale. The firm purchases inventory at a fixed per-unit cost γ. The per-unit sale price is δ. The firm must purchase its inventory prior to the arrival of its customers who, in total, demand a nonnegative random amount $\boldsymbol{\xi}$ with distribution μ. As there is no initial inventory, the quantity ordered by the firm is the total amount available for sale; the firm's sales is the smaller of the demand and the inventory level. Either there is excess demand, which represents lost sales, or excess inventory, which perishes. The firm must select an inventory level to maximize its expected profit.*

Detail. Our cost function (negative profit) is

$$f(\xi, x) = \begin{cases} (\gamma - \delta)x & \text{if } 0 \le x \le \xi, \\ \gamma x - \delta\xi & \text{if } x > \xi \\ \infty & \text{if } x < 0. \end{cases}$$

With $\mu(\eta) = \text{prob.}\,[\boldsymbol{\xi} \le \eta]$, the stochastic optimization problem is

$$\min_{x \in \boldsymbol{R}} Ef(x) = \int f(\xi, x)\, \mu(d\xi).$$

The optimal solution $\bar{x}$ must satisfy

$$\lim_{\eta \nearrow \bar{x}} \mu(\eta) := \mu(\bar{x}^-) \leq \frac{\delta - \gamma}{\delta} \leq \mu(\bar{x}),$$

which allows for the possibility of a (discontinuous) jump in μ at $\bar{x}$, as could happen when the random variable $\boldsymbol{\xi}$ has a discrete distribution. The figure below illustrates the possibilities. $\square$

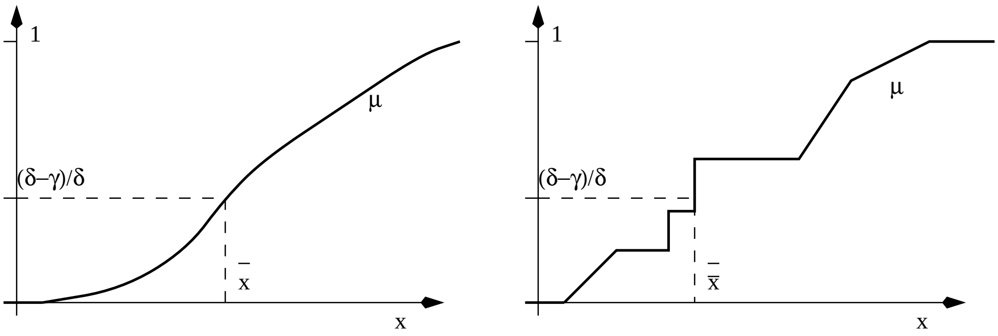

Figure 1: Optimal solution of the newsboy problem

2.2 Example (portfolio selection)**.** *An investor wants to allocate her wealth w among n assets (may include borrowing), i.e., choose*

$$x \in I\!R^n_+ \text{ with } \sum_{i=1}^{n} x_j \leq w,$$

so as to 'guarantee' a return r. The per-unit return on asset j is given by a random variable $\boldsymbol{\xi}_j$. The probability measure μ defines the joint distribution of $(\boldsymbol{\xi}_1, \ldots, \boldsymbol{\xi}_n) = \boldsymbol{\xi}$. Problems of this type are sometimes called tracking problems in that the investor is mainly interested in 'tracking' a pre-determined return.

Detail. The problem isn't well defined in that the return $\langle \boldsymbol{\xi}, x \rangle$ is a random variable and thus a return of r can usually not be guaranteed. Depending on the meaning one gives to the word 'guarantee', we are going to end up with different tracking problems. Let's begin with the following interpretation:

the investor's goal is to avoid falling short of the desired revenue r as much as possible. To this effect, let's introduce the *monitoring* function

$$\theta(\langle\xi,x\rangle - r) = \begin{cases} (\langle\xi,x\rangle - r)^2 & \text{if } \langle\xi,x\rangle \le r, \\ 0 & \text{otherwise.} \end{cases}$$

Our cost function reads,

$$f(\xi,x) = \begin{cases} \theta(\langle\xi,x\rangle - r) & \text{if } \sum_{i=1}^n x_j \le w,\ x \ge 0, \\ \infty & \text{otherwise,} \end{cases}$$

and the optimal portfolio is then

$$\bar{x} \in \operatorname{argmin}_{x\in \boldsymbol{R}^n} \int_{\boldsymbol{R}^n} f(\xi,x)\,\mu(d\xi).$$

This is a linear-quadratic programming problem whose structure can be exploited to design efficient solution techniques. However the integration of $(\langle\xi,x\rangle - r)^2$ over $\{\xi \in \Xi \,|\, \langle\xi,x\rangle \le r\}$ might require extensive computational effort. Maybe for this reason, variants of this problem have been proposed that are more amenable to simple implementations, but are less satisfactory from a modeling viewpoint.

Let's consider first the following variant with

$$\theta(\langle\xi,x\rangle - r) = (\langle\xi,x\rangle - r)^2.$$

This means that we now penalize deviations from r in both directions, on the upside as well as on the downside. The cost function is

$$f(\xi,x) = \begin{cases} (\langle\xi,x\rangle - r)^2 & \text{if } \sum_{i=1}^n x_j \le w,\ x \ge 0, \\ \infty & \text{otherwise.} \end{cases}$$

The stochastic optimization problem, $\inf_x Ef(x) = E\{f(\boldsymbol{\xi},x)\}$, becomes a quadratic programming problem

$$\begin{array}{ll} \inf_{x\in\boldsymbol{R}^n} & \langle x, Cx\rangle - 2r\langle\bar{\xi},x\rangle + r^2 \\ \text{so that} & \sum_{i=1}^n x_j \le w, \quad x \ge 0, \end{array}$$

where the elements of the matrix C are $C_{ij} = E\{\boldsymbol{\xi}_i\boldsymbol{\xi}_j\}$ and the vector $\bar{\xi} = E\{\boldsymbol{\xi}\}$.

Another variant has the following cost function:

$$f(\xi,x) = \begin{cases} (\langle\xi,x\rangle - \langle\bar{\xi},x\rangle)^2 & \text{if } \langle\bar{\xi},x\rangle \ge r,\ \sum_{i=1}^n x_j \le w,\ x \ge 0, \\ \infty & \text{otherwise.} \end{cases}$$

The corresponding stochastic optimization problem is as usual, $\inf Ef(x)$, which again takes the form of a quadratic programming problem:

$$\begin{aligned} \inf_{x \in \boldsymbol{R}^n} \quad & \langle x, \Sigma x\rangle \\ \text{so that} \quad & \langle \bar{\xi}, x\rangle \geq r,\ \sum_{i=1}^{n} x_j \leq w,\ x \geq 0, \end{aligned}$$

where Σ is the covariance matrix of $\boldsymbol{\xi}$. As now formulated, the goal of the investor is to minimize the variance of her return under the additional constraints that the expected return be at least r. This is the so-called *mean/variance portfolio optimization problem.* Markowitz (1959) original formulation of this problem has $\langle \bar{\xi}, x\rangle = r$ instead of the $\geq$ inequality. If one views r as a parameter and let for each r, $v(r)$ be the optimal value of the corresponding quadratic programming problem. Then v traces out the so-called *efficient frontier*, see Fig. 2; for each r, there is a solution that minimize variance for an expected return of at least r. □

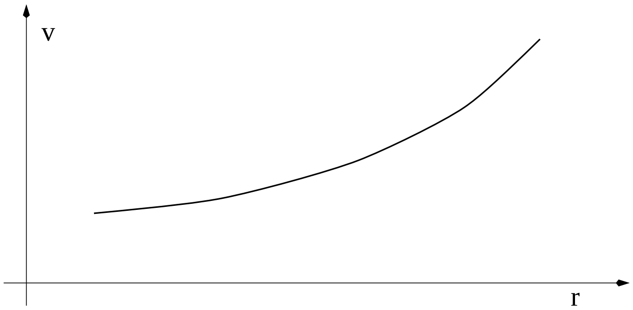

Figure 2: Efficient frontier

2.3 Exercise (mean/variance and tracking problems). *Consider again the mean/variance problem but insist on having $\langle \bar{\xi}, x\rangle = r$. Assuming this problem to be feasible, show that the mean/variance portfolio optimization problem is then equivalent to a tracking problem but with an adjusted value for the return.*

Guide. With $\langle \bar{\xi}, x\rangle = r$, $\langle x, \Sigma x\rangle = \langle x, Cx\rangle - 2r\langle \bar{\xi}, x\rangle + r^2$ and the quadratic program takes the same form as when $\theta(\langle \xi, x\rangle - r) = (\langle \xi, x\rangle - r)^2$, except for the additional constraint $\langle \bar{\xi}, x\rangle = r$. The resulting problem is a (convex) quadratic programming problem (with linear constraints), there exists a multiplier, call it λ_r, which can be associated to this extra constraint so that an optimal solution of the (modified) mean/variance problem is also a solution of a tracking problem with monitoring function $\theta(\langle \xi, x\rangle - r) = (\langle \xi, x\rangle - (r + \lambda_r/2))^2$. Note that λ_r could be negative as well as positive or 0. $\square$

The mean/variance portfolio optimization problem can also be derived from an expected utility maximization argument. Suppose the investor's utility function is a concave function $u : \mathbb{R}_+ \to \mathbb{R}_+$, at least twice differentiable. Taylor expansion around r yields

$$u(\tau) = \tilde{u}(\tau) + o(|\tau - r|^2) \quad \text{where} \quad \tilde{u}(\tau) = u(r) + u'(r)(\tau - r) + \tfrac{1}{2}u''(r)(\tau - r)^2.$$

Replacing u by $\tilde{u}$ in the expected utility maximization problem:

$$\max_{x \in \boldsymbol{R}^n} E\{u(\langle \boldsymbol{\xi}, x\rangle)\} \ \text{ so that } \ \langle \bar{\xi}, x\rangle = r, \ \sum_{i=1}^{n} x_j \le w, \ x \ge 0,$$

yields the (modified) mean/variance problem considered in 1.C.

2.4 Example (network capacity expansion)**.** *For a power transmission network, let e_i be the external flow at node i, i.e., the difference between demand and supply at node i. The internal flow y_j on arc j is limited by its capacity γ_j. Total supply exceeds total demand but the capacity of the transmission lines has to be expanded from γ_j to $\gamma_j + x_j$, with v_j an upper bound on x_j, to render the problem feasible.*
The cost of such an expansion is $\sum_{j=1}^{n} \psi_j(x_j)$. A deterministic version of the capacity expansion problem would be:

$$\min_{0 \le x_j \le v_j} \sum_{j=1}^{n} \psi_j(x_j)$$

so that

$$0 \le y_j \le \gamma_j + x_j, \; j = 1, \ldots, n, \; \sum_{\to i} y_j - \sum_{\leftarrow i} y_j \ge e_i \; i = 1, \ldots, m;$$

$\sum_{\to i} y_j$ is the (internal) flow into i whereas $\sum_{\leftarrow i} y_j$ is the flow from i to the other nodes.

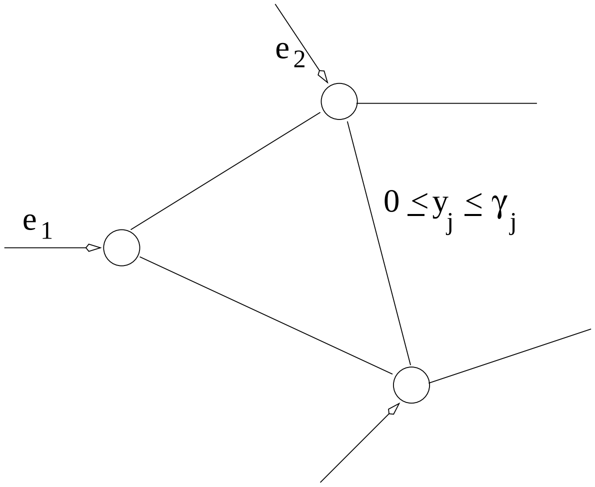

Figure 3: Power transmission network

But external flows, due to breakdowns in the supply generators and variations in the demand, are actually random, i.e., must be modeled as a random vector $\boldsymbol{\xi} = (\boldsymbol{\xi}_1, \ldots, \boldsymbol{\xi}_m)$ *with possible values in* $\Xi \subset \mathbb{R}^m$ *and (joint) distribution* μ.

Detail. Let's consider two different pairs of stochastic programming models for this capacity expansion problem. In the first pair, a monitoring function is associated with the level of demand satisfaction at each node i:

$$\theta_i(\tau) = \begin{cases} 0 & \text{if } \tau < 0, & \text{(supply exceeds demand)} \\ \kappa_i \tau^2 / 2\delta_i & \text{if } \tau \in [0, \delta_i], & \text{(low level excess demand)} \\ \kappa_i \tau - \kappa_i \delta_i / 2 & \text{if } \tau > \delta_i, & \text{(high level excess demand).} \end{cases}$$

In our first model, the budget available for the capacity expansion is fixed, say β; eventually, β can be varied and the dependence of the solution on β can be (parametrically) analyzed. The cost function of the stochastic optimization problem, $\min_x Ef(x)$, is

$$f(\xi, x) = \begin{cases} Q(\xi, x) & \text{if } 0 \le x_j \le v_j,\ j = 1, \ldots, n,\ \ \sum_{j=1}^n \psi_j(x_j) \le \beta, \\ \infty & \text{otherwise,} \end{cases}$$

where

$$Q(\xi, x) = \inf_{y \in \boldsymbol{R}^n} \{ \sum_{i=1}^m \theta_i(\xi_i - \sum_{\to i} y_j + \sum_{\leftarrow i} y_j) \,|\, 0 \le y_j \le \gamma_j + x_j,\ j = 1, \ldots, n. \}$$

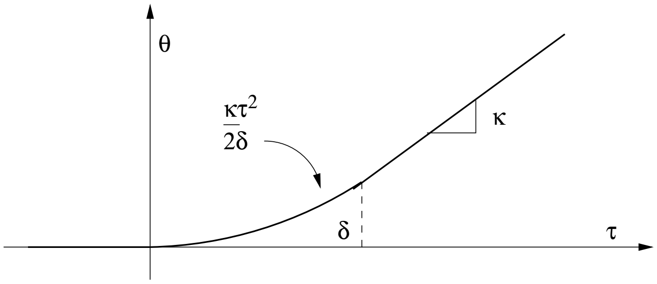

Figure 4: Monitoring function for capacity expansion problem

One refers to such a problem as a (two-stage) *stochastic program with recourse.* The first stage decision x has to be made 'now,' before the supply/demand situation can be observed. The second stage or recourse decisions, i.e., decision about flows y, are made after the supply/demand situations are known. The function Q is called the *recourse cost function.*

Our second model is almost identical except that the capacity expansion cost is included in the objective, not in the form of a budgetary constraints. The function f is adjusted as follows:

$$f(\xi, x) = \begin{cases} \sum_{j=1}^n \psi_j(x_j) + \lambda Q(\xi, x) & \text{if } 0 \le x_j \le \upsilon_j,\ j = 1, \ldots, n, \\ \infty & \text{otherwise,} \end{cases}$$

where $\lambda > 0$ is a coefficient chosen to weigh the relative importance of shortage costs with respect to expansion costs.

In the second pair of models, one insists on having the demand satisfied sufficiently often. One interprets this to mean that with sufficiently high probability there will exist feasible flows $0 \le y \le \gamma + x$ that will satisfy the demand. Let

$$A = \{(\xi, x) \in \Xi \times I\!R^n \,|\, \exists\, y \in [0, \gamma + x] \text{ so that } \sum_{\to i} y_j - \sum_{\leftarrow i} y_j \ge \xi_i,\ \forall\, i\}.$$

We want to restrict the choice of x to those for which $\mu(\{\xi \in \Xi \,|\, (\xi, x) \in A\}) \ge \alpha$ where α is the reliability level to be achieved. With

$$g(\xi, x) = 1_A(\xi, x), \quad \text{the indicator function of } A,$$

the problem can be formulated as follows:

$$\min_{x \in \boldsymbol{R}^n} F(x) \text{ where } F(x) = \begin{cases} \sum_{j=1}^n \psi_j(x_j) & \text{if } 0 \le x \le \upsilon,\ \ E\{g(\boldsymbol{\xi}, x)\} \ge \alpha, \\ \infty & \text{otherwise.} \end{cases}$$

This formulation of the capacity expansion problem doesn't fit in the class of problems of the type: $\inf_x E\{f(\xi,x)\}$, i.e., the minimization of an integral function. However, the constraints of the problem involve an integral functional. One refers to such a problem as a *stochastic programs with chance constraints.*

If rather than a common reliability level, we want to specify a reliability level, say α_i at each node i, let μ_i be the marginal distribution of $\boldsymbol{\xi}_i$ and E^{μ_i} the expectation with respect to this marginal distribution. The problem becomes

$$\min_{x\in \boldsymbol{R}^n} F(x) = \begin{cases} \sum_{j=1}^n \psi_j(x_j) & \text{if } 0 \le x \le v, \\ & \quad E^{\mu_i}\{g(\boldsymbol{\xi},x)\} \ge \alpha_i,\ i=1,\dots,m, \\ \infty & \text{otherwise.} \end{cases}$$

This problem involves a number of chance constraints but is actually a much simpler problem to solve than the previous one. Indeed, with $\hat{e}_i$ the α_i-quantile of the distribution μ_i, i.e., $\mu_i[\boldsymbol{\xi}_i \le \hat{e}_i] = \alpha_i$, it's easy to show that the solution of the preceding problem can be obtained by solving the deterministic optimization problem:

$$\begin{array}{lc} \min_{x\in \boldsymbol{R}^n} & \sum_{j=1}^n \psi_j(x_j) \\ \text{so that} & 0 \le x \le v,\ \ 0 \le y \le \gamma + x, \\ & \sum_{\to i} y_j - \sum_{\leftarrow i} y_j \ge \hat{e}_i,\ i=1,\dots,m. \end{array}$$

If the network is large, one can even rely on a decomposition technique to take advantage of the network structure of the y-part of the problem. □

2.5 Example (eutrophication management). *Eutrophication management deals with the pollution control of lakes, rivers, etc. Our lake has an industrial zone on its East end and is used for recreational purposes on the West end. The available controls are the installation of sewage treatment plants in various locations along the rivers feeding the lake as well as the building of reed basins, the reeds absorbing cl-phosphates that are the major cause of pollution.*

Purification installations should be chosen so that each section (I to IV) of the lake meets certain minimal water quality levels. Pollution is dependent on the atmospheric conditions about which there is enough statistical data to reliably estimate their distribution. In the lake, transport of the pollution occurs mostly from East to West as the major inlets are in the East part of the lake and the outlet is on the West end.

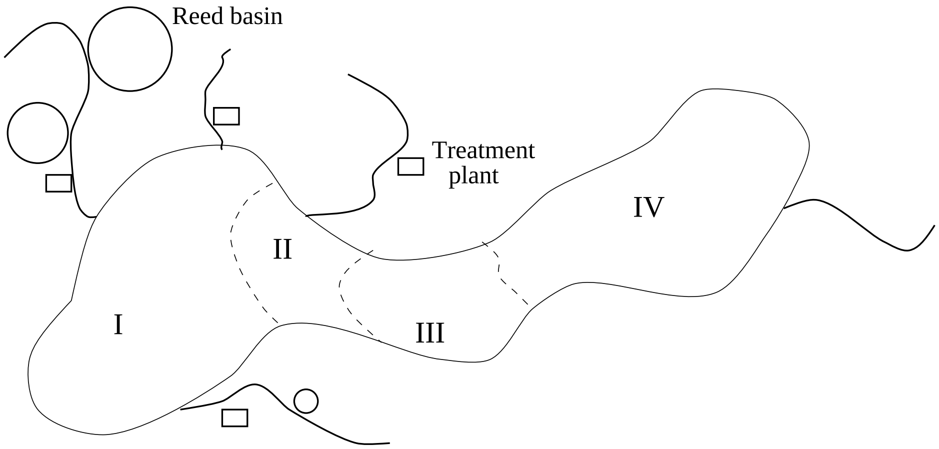

Figure 5: Lake Balaton (Hungary): Eutrophication management

Detail. For $j = 1, \ldots, n$, let x_j be the size of the treatment plant or reed basin in location j; v_j is an upper bound on x_j. The relation

$$y = T(\xi)x - d(\xi)$$

determines the water quality levels y_i, $i = I, \ldots, IV$ in the four lake sections. The weather pattern, designated by the parameter ξ, affects both the exogenous pollution $d(\xi)$ and the transport of the pollutants, given by $T(\xi)x$, from one section to another. This latter component is controlled to a large extent by the choice of x. The technical and budgetary limitations on the choice of x are included in the (linear) constraints $Ax \leq b$. Again, we use monitoring functions to evaluate deviations from the desired quality level in each particular section of the lake:

$$\theta_i(\tau) = \begin{cases} 0 & \text{if } \tau < \gamma_i^l, \qquad \text{(excellent)} \\ \rho_i(\tau - \gamma_i^l)^2/2(\gamma_i^u - \gamma_i^l) & \text{if } \tau \in [\gamma_i^l, \gamma_i^u], \quad \text{(acceptable)} \\ \rho_i(\tau - \gamma_i^u) - \rho_i(\gamma_i^l + \gamma_i^u)/2 & \text{if } \tau > \gamma_i^u, \qquad \text{(unacceptable).} \end{cases}$$

The cost function of the stochastic optimization problem, $\min_x Ef(x)$, is then

$$f(\xi, x) = \begin{cases} \sum_{i=I}^{IV} \theta_i(\langle T_i(\xi), x\rangle - d_i(\xi)) & \text{if } Ax \leq b,\ 0 \leq x \leq v, \\ \infty & \text{otherwise;} \end{cases}$$

here $T_i(\xi)$ is the i-th row of the matrix $T(\xi)$. If $\boldsymbol{\xi}$ has a discrete distribution, this is an extended linear-quadratic programming problem for which particularly efficient solution techniques have been designed, e.g., the Lagrangian finite-generation technique. $\square$

3 The price of nonanticipativity

All the examples that we have seen so far are characterized by the following sequence:

$$\boxed{\text{decision: } x \;\rightsquigarrow\; \text{observation: } \xi \;\rightsquigarrow\; \text{evaluation: } (\xi, x)}$$

In particular, this means that the decision can't depend on the value assumed by the random quantities $\boldsymbol{\xi}$. The problem is of the *here-and-now* type. In a problem of the *wait-and-see* type one would actually be allowed to wait making a decision until the random quantities can be observed. In other words, the decision may depend on ξ. Such a problem could be formulated as follows,

$$\inf_{x \in \mathcal{M}(\Xi, \mathcal{S}; \boldsymbol{R}^n)} I_f[x] := \int_\Xi f(\xi, x(\xi))\, \mu(d\xi),$$

where $\mathcal{M}(\Xi, \mathcal{S}; \mathbb{R}^n)$ is the space of all $\mathcal{S}$-measurable functions $x : \Xi \to \mathbb{R}^n$. Since the decisions are then made with perfect (complete) information, the quantity

$$\text{EVPI} = \inf_{x \in \boldsymbol{R}^n} Ef(x) \;\; - \inf_{x \in \mathcal{M}(\Xi, \mathcal{S}; \boldsymbol{R}^n)} I_f[x].$$

is called the *expected value of perfect information.* Obviously, EVPI is always nonnegative. It takes on the value 0 only if the cost associated the decisions (affected by the random parameters of the problem) is itself 0 or if the problem is basically nonstochastic, i.e., when there is $\bar{x} \in \operatorname{argmin}_{x \in \mathcal{M}} I_f[x]$ that is a constant function.

Although, an optimal solution of the wait-and-see problem is an element of a function space, the wait-and-see problem is in some way easier to solve than the here-and-now problem which usually requires the repeated evaluation of multidimensional integrals. The reason for this is that one can interchange integration and minimization in the wait-and-see problem! To deal with this technically, one needs to appeal to the theory of normal integrands and the associated integral functionals.

As far as our applications are concerned there is no loss of generality in working with the assumption that the σ-field $\mathcal{S}$ is μ-complete. This allows us to work with a definition of normal integrand that is easier to state. So, henceforth let's proceed with this assumption of μ-completeness of $\mathcal{S}$.

3.1 Definition (normal integrand). *Assuming that $\mathcal{S}$ is μ-complete, a normal integrand is an extended real-valued function $f : \Xi \times I\!R^n \to \overline{I\!R}$ with the following properties:*

- *(a) the function $(\xi, x) \mapsto f(\xi, x)$ is $\mathcal{S} \otimes \mathcal{B}$-measurable where $\mathcal{B}$ the Borel field on $I\!R^n$;*
- *(b) for every $\xi \in \Xi$, the function $x \mapsto f(\xi, x)$ is lsc.*

We are interested in the following properties of normal integrands:

3.2 Theorem (measurability of optimal values and solutions). *For any normal integrand $f : \Xi \times I\!R^n \to \overline{I\!R}$, let*

$$p(\xi) = \inf f(\xi, \cdot), \qquad P(\xi) = \operatorname{argmin} f(\xi, \cdot).$$

Then the function $p : \Xi \to \overline{I\!R}$ is measurable and the mapping $P : \Xi \rightrightarrows I\!R^n$ is closed-valued and measurable, i.e., the sets $P(\xi)$ are closed and $\{\xi \in \Xi \mid P(\xi) \cap O \neq \emptyset\} \in \mathcal{S}$ for all open subsets O of $I\!R^n$.

Therefore, the set $A = \{\xi \mid \operatorname{argmin}_x f(\xi, x) \neq \emptyset\} \subset \Xi$ is measurable, and it is possible for each $\xi \in A$ to select a minimizing point $x(\xi)$ in such a manner that the function $\xi \mapsto x(\xi)$ is measurable.

Proof. See [6, Theorem 14.37]. □

To an integrand f we associate the following integral functionals:

$$\begin{aligned} I_f[x] &:= \textstyle\int_\Xi f(\xi, x(\xi))\, \mu(d\xi) \quad \text{for } x \in \mathcal{M}(\Xi, \mathcal{S}; I\!R^n) \\ Ef(x) &:= \textstyle\int_\Xi f(\xi, x))\, \mu(d\xi) \quad \text{for } x \in I\!R^n \end{aligned}$$

3.3 Theorem (interchange criterion). *Let $f : \Xi \times I\!R^n \to \overline{I\!R}$ be a normal integrand. The minimization of I_f over $\mathcal{M}$ can be reduced to pointwise minimization in the sense that, as long as $I_f \not\equiv \infty$, one has*

$$\inf_{x \in \mathcal{M}} \int_\Xi f(\xi, x(\xi))\, \mu(d\xi) = \int_\Xi \Big[\inf_{x \in I\!R^n} f(\xi, x) \Big]\, \mu(d\xi).$$

Moreover, as long as this common value is not $-\infty$*, one has that for all* $\bar{x} \in \mathcal{M}(\Xi, \mathcal{S}; I\!R^n)$ *that*

$$\bar{x} \in \operatorname{argmin}_{x \in \mathcal{M}} I_f[x] \quad \Longleftrightarrow \\ \bar{x}(\xi) \in \operatorname{argmin}_{x \in \boldsymbol{R}^n} f(\xi, x) \text{ for } \mu\text{-almost every } \xi \in \Xi.$$

Proof. Let $p(\xi) = \inf_x f(\xi, x)$ and $P(\xi) = \operatorname{argmin}_x f(\xi, x)$, recalling from 1.G that these depend measurably on ξ. For any $x \in \mathcal{M}$ we have $f(\xi, x(\xi)) \geq p(\xi)$ for all ξ; hence $\inf_{\mathcal{M}} I_f \geq \int_\Xi p(\xi)\, \mu(d\xi)$.

To prove the inequality in the other direction, for $\varepsilon > 0$, let

$$A_\varepsilon(\xi) = \{x \in I\!R^n \,|\, f(\xi, x) \leq \max\,[\,p(\xi) + \varepsilon, -\varepsilon^{-1}\,]\}.$$

The mapping $\xi \mapsto A_\varepsilon(\xi)$ is closed-valued and measurable (*Variational Analysis*, Proposition 14.33) and nonempty-valued for μ-almost all ξ, otherwise one would have $I_f \equiv \infty$. Consequently, there is a measurable selection $x_\varepsilon : \Xi \to I\!R^n$ such that $x_\varepsilon(\xi) \in A_\varepsilon(\xi)$ for μ-almost all ξ, cf. [6, Corollary 14.6]. And for this selection, one has

$$\int_\Xi f(\xi, x_\varepsilon(\xi))\, \mu(d\xi) \leq \int_{\Xi_\varepsilon} p(\xi)\, \mu(d\xi) + \varepsilon\mu(\Xi_\varepsilon) + \varepsilon^{-1}\mu(\Xi \setminus \Xi_\varepsilon),$$

where $\Xi_\varepsilon = \{\xi \in \Xi \,|\, p(\xi) + \varepsilon > -\varepsilon^{-1}\}$. Next, let $\varepsilon \searrow 0$, to obtain

$$\inf_{x \in \mathcal{M}} I_f[x] \leq \lim_{\varepsilon \searrow 0} \int_\Xi f(\xi, x(\xi))\, \mu(d\xi) \leq \int_\Xi p(\xi)\, \mu(d\xi).$$

If a function $\bar{x} \in \mathcal{M}$ attains the minimum of I_f on $\mathcal{M}$, since $f(\xi, \bar{x}(\xi)) \geq \inf f(\xi, \cdot) = p(\xi)$ for all ξ, that's equivalent to having $\mu(\{\xi \mid f(\xi, \bar{x}(\xi)) > p(\xi)\}) = 0$ under our assumption that $\int_\Xi p(\xi)\, \mu(d\xi)$ is finite. This is identical to the stated criterion, because $f(\xi, \bar{x}(\xi)) > p(\xi)$ means that $\bar{x}(\xi) \notin \operatorname{argmin} f(\xi, \cdot)$. □

The fact that a wait-and-see problem doesn't require any computational effort that goes much beyond solving a number, possibly large, of deterministic optimization problems, leads us to consider the following variant of the here-and-now problem, $\inf_{x \in \boldsymbol{R}^n} Ef(x)$:

$$\inf_{x \in \mathcal{M}, \hat{x} \in \boldsymbol{R}^n} I_f[x] \text{ so that } \; x(\xi) = \hat{x} \text{ for } \mu\text{-almost all } \xi \in \Xi.$$

We now allow the decision to depend on ξ but then introduce a constraint that restrict the choice of x to the linear subspace of functions that are μ-almost surely constant. The constraints $x(\xi) = \hat{x}$ μ-a.s. are the *nonanticipativity constraints*. They render explicit the requirement that the choice of x can't really depend on ξ; the μ-a.s. condition is included for purely technical reasons.

Sufficient optimality conditions for this problem would be: $(x^*, \hat{x}^*)$ is an optimal pair if they are feasible,

$$\text{for } \mu\text{-almost all } \xi\text{:}\quad x^*(\xi) = \hat{x}^*,$$

and there exists $w \in \mathcal{M}(\Xi, \mathcal{S}; \mathbb{R}^n)$ such that

$$(x^*, \hat{x}^*) \in \operatorname{argmin}_{x \in \mathcal{M}, \hat{x} \in \boldsymbol{R}^n} \left[I_f[x] + \int_\Xi \langle w(\xi), x(\xi) \rangle \, \mu(d\xi) - \langle \int_\Xi w(\xi) \, \mu(d\xi), \hat{x} \rangle \right];$$

for each ξ, the multiplier $w(\xi)$ is attached to the constraint $x(\xi) - \hat{x} = 0$. In this last expression, we can split the minimization with respect to $x \in \mathcal{M}$ and that with respect to $\hat{x} \in \mathbb{R}^n$. With

$$f_w(\xi, x) = f(\xi, x) + \langle w(\xi), x \rangle,$$

these sufficient conditions become

- $x^*(\xi) = \hat{x}^*$, μ-a.s.,
- $x^* \in \operatorname{argmin}_{x \in \mathcal{M}} I_{f_w}[x]$,
- $\hat{x}^* \in \operatorname{argmin}_{\hat{x} \in \boldsymbol{R}^n} \langle \bar{w}, \hat{x} \rangle$ where $\bar{w} = E\{w(\boldsymbol{\xi})\}$.

In order for $\hat{x}^*$ to satisfy both the first and the last conditions, one must have

$$E\{w(\boldsymbol{\xi})\} = \bar{w} = 0.$$

In view of the interchange criterion in 1.H, the sufficient conditions for optimality can thus be restated as follows:

- $x^* : \Xi \to \mathbb{R}^n$ is μ-almost surely a constant function,

and there exists $w \in \mathcal{M}(\Xi, \mathcal{S}; \mathbb{R}^n)$ with $E\{w(\boldsymbol{\xi})\} = 0$ such that

- for μ-almost all ξ: $x^*(\xi) \in \operatorname{argmin}_{x \in \boldsymbol{R}^n} [\, f_w(\xi, x) = f(\xi, x) + \langle w(\xi), x \rangle \,]$.

Moreover, $\hat{x}^* \in \operatorname{argmin} Ef$.

The multipliers w can be viewed as an 'information' price system in the following sense: if the decision maker was allowed to adjust his decision, $\hat{x}^* \in \operatorname{argmin}_{x \in \boldsymbol{R}^n} Ef(x)$, after $\boldsymbol{\xi}$ is observed and the price system used for computing the charge for these adjustments would be $w(\xi)$, there would be no gain in making these adjustments. Indeed, in view of the optimality conditions, for μ-almost all $\xi \in \Xi$,

$$f(\xi, \hat{x}^*) + \langle w(\xi), \hat{x}^* \rangle \leq f(\xi, x) + \langle w(\xi), x \rangle, \ \forall x \in I\!\!R^n,$$

in particular, for all $x \in \operatorname{argmin} f(\xi, \cdot)$. It follows that for μ-almost all $\xi \in \Xi$:

$$f(\xi, \hat{x}^*) \leq f(\xi, x) + \langle w(\xi), x - \hat{x}^* \rangle \ \forall x \in I\!\!R^n.$$

The return from being able to wait to make a decision, say $x_\xi \in I\!\!R^n$, until the random elements are observed would be negated if the additional price charged would be

$$\langle w(\xi), x_\xi - \hat{x}^* \rangle.$$

The *nonanticipativity multipliers* w can thus also be viewed as an *equilibrium* price system.

At least in the convex case, when f is a *convex normal integrand*, i.e., $x \mapsto f(\xi, x)$ is convex for all $\xi \in \Xi$, and $\bar{x} \in \operatorname{argmin}_{\mathcal{M}} I_f$, one has

$$I_f[\bar{x}] + \int_\Xi \langle w(\xi), \bar{x}(\xi) \rangle \, \mu(d\xi) = Ef(\hat{x}^*)$$

which means that $E\{\langle w(\boldsymbol{\xi}), (\bar{x}(\boldsymbol{\xi}) - \hat{x}^* \rangle\} = \text{EVPI}$, as could have been expected; note that $E\{\langle w(\boldsymbol{\xi}), \hat{x}^* \rangle\} = 0$.

We didn't touch here on the question of the existence of such multipliers. This is not that simple an issue! Even when f is a convex normal integrand. Conditions that will guarantee their existence involve a 'classical' constraint qualification but also a 'nonanticipativity' constraint qualification that is germane to stochastic optimization problem. At this stage, we are not ready to deal with this issue. Let's just note that for all examples of the type '$\inf Ef$' mentioned earlier, except for the pair of stochastic programs with recourse in 1.D, both constraint qualifications are satisfied. And the preceding optimality conditions are both sufficient and necessary.

These optimality conditions also suggest a solution procedure for our stochastic optimization problem $\inf_x Ef(x)$. Indeed, any procedure that would sequentially (or progressively) generate improving estimates w^ν of the nonanticipativity multipliers could be used to find solutions $x^\nu \in \operatorname{argmin}_{x \in \mathcal{M}} I_{f^\nu}[x]$ with $f^\nu = f + \langle w^\nu, \cdot \rangle$ that would progressively satisfy the nonanticipativity constraint: the solution must be a μ-a.s. constant function.

Actually, such a procedure has already been proposed and is known as the 'Progressive Hedging Algorithm' [5] and a number of variants of this algorithm have been implemented [1], [3], [4] and [2].

References

[1] T. Helgason and S.W. Wallace. Approximate scenario solution in the progressive hedging algorithm. RH-08-89, Raumvísindastofnun Háskólans, 1989.

[2] J.M. Mulvey and A. Ruszczynski. A new scenario decomposition method for large-scale stochastic optimization. *Operations Research*, 43:477–490, 1995.

[3] J.M. Mulvey and H. Vladimirou. Evaluation of a distributed hedging algorithm for stochastic network programming. Statistics and Operations Research SOR 88-14, Princeton University, 1988.

[4] J.M. Mulvey and H. Vladimirou. Solving multistage investment problems: an application of scenario aggregation. Statistics and Operations Research SOR 88-1, Princeton University, 1988.

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

[6] R.T. Rockafellar and R.J-B Wets. *Variational Analysis*. Springer, 1998.