Introduction to modeling using stochastic programming

Andy Philpott
The University of Auckland

Tutorial presentation at SPX, Tuscon, October 9th, 2004

Summary

Introduction to basic concepts

Risk

Multi-stage stochastic programs

Non-standard models

Desiderata

Example: the newsvendor

A newsvendor must decide how many newspapers to order for the next day. Each newspaper costs the vendor α =\$1.00 and she sells them for $\alpha + \beta$ =\$1.50. Any papers left over can be sent back for recycling with a refund of γ =\$0.10 per newspaper. Her cost is

$$C(x, \xi) = \begin{cases} -0.5x, & \text{if } x < \xi \\ 0.9x - 1.4\xi, & \text{otherwise.} \end{cases}$$

What is the optimal choice of x if we know ξ? What is the optimal choice of x if ξ is random?

Some notation

Stochastic programmers use the notation ξ to denote all the random components in a problem. So ξ is a *random variable* (or vector of variables). A random variable is a function from a probability space $(\Omega, \mathcal{F}, P)$ to $\mathbb{R}^n$. The elementary events in Ω are denoted ω, and so we should write $\xi(\omega)$ instead of ξ. The expectation of a function $z(\xi)$ is then denoted

$$E[z(\xi)] = \int_{\Omega} z(\xi(\omega)) dP(\omega).$$

Solution to newsvendor problem

Suppose ξ has cumulative distribution function F. Then

$$E[C(x,\xi)] = \int_{-\infty}^{\infty} C(x,\xi(\omega))dF(\omega)$$

$$\begin{aligned} x^* &= F^{-1}\left(\frac{\beta}{\alpha+\beta-\gamma}\right) \\ &= F^{-1}\left(\frac{0.5}{1.4}\right) \\ &= F^{-1}(0.35714) \end{aligned}$$

Example

Suppose ξ is uniformly distributed on [0,100].

$$
\begin{aligned}
E[C(x,\xi)] &= \int_\Omega C(x,\xi(\omega))dP(\omega) \\
&= \int_0^x (0.9x - 1.4\omega)\frac{d\omega}{100} \\
&\quad + \int_x^{100} -0.5x\frac{d\omega}{100} \\
&= (\frac{1}{100})(0.7x^2 - 50x)
\end{aligned}
$$

$$x^* = F^{-1}(0.35714) = 35.714 \text{ with expected cost } -8.928$$

We rarely get a *closed form* like this for the expected value of a candidate solution for a stochastic optimization problem. In this case we can use this function to conceptually test out different approaches to minimizing this function. In practice we need a *simulation*.

Wait-and-See Problem

The optimal choice of x if we know ξ is to choose $x(\xi) = \xi$. This gives

$$C(x(\xi), \xi) = -0.5\xi$$

In the "simulation"

$$\begin{aligned} E[C(x(\xi), \xi)] &= E[-0.5\xi] \\ &= -25.0 \end{aligned}$$

Expected Value Problem

Take expectations of random variables, giving $E[\xi] = 50$.

$$C(x, E[\xi]) = \begin{cases} -0.5x, & \text{if } x < E[\xi] \\ 0.9x - 1.4E[\xi], & \text{otherwise.} \end{cases}$$
$$= \begin{cases} -0.5x, & \text{if } x < 50 \\ 0.9x - 70, & \text{otherwise.} \end{cases}$$

$C(x, E[\xi])$ has minimizer $\bar{x} = 50$ giving value -25.0, but in the "simulation"

$$E[C(\bar{x}, \xi)] = -7.5$$

This illustrates the famous "fallacy of averages".

Summary: WS, RP, EVPI, and VSS

V(WS) $=$ optimal value (cost) of problem with perfect information $= -25.0$

V(RP) $=$ optimal value (cost) of problem with imperfect information $= -8.928$

EVPI $=$ expected value of perfect information $=$ V(RP) - V(WS) $=$ 16.072

VSS $=$ value of stochastic solution $= E[C(\bar{x}, \xi)]$ - V(RP) $=$ 1.428

The newsvendor with scenarios

In general we cannot solve stochastic programs analytically. Approximate the randomness using *scenarios*, and solve a deterministic optimization problem.

Suppose we have scenarios ξ_i, $i = 1, 2, \ldots, N$, each with probability p_i. Let x be the number of newspapers bought, and let y_1 be the number sold, y_2 be the unsatisfied demand, and y_3 be the number recycled.

$$\min 1.0x + \sum_{i=1}^{N} p_i[Q(x, \xi_i)]$$

$$\begin{aligned} Q(x, \xi_i) = \ & \min -1.5y - 0.1v \\ & \text{s.t. } y_1 + y_2 = \xi_i, \\ & \quad y_1 + y_3 = x, \\ & \quad x, y_1, y_2, y_3 \geq 0 \end{aligned}$$

The two-stage stochastic program with (fixed) recourse

$$\begin{array}{ll} \text{RP: minimize} & c^\top x + E_\xi[Q(x,\xi)] \\ \text{subject to} & Ax = b, \\ & x \geq 0. \end{array}$$

where

$$Q(x, \xi(\omega)) = \min\{q^\top y \mid Wy = h(\omega) - T(\omega)x,\ y \geq 0\}.$$

In example we have

$$q^\top = \begin{bmatrix} -1.5 & 0 & -0.1 \end{bmatrix}, \quad W = \begin{bmatrix} 1 & 1 & 0 \\ 1 & 0 & 1 \end{bmatrix},$$

$$h(\omega) = \begin{bmatrix} \xi(\omega) \\ 0 \end{bmatrix}, \quad T(\omega) = \begin{bmatrix} 0 \\ -1 \end{bmatrix}.$$

Exterior sampling

Randomly sample from distribution of ξ, and solve *sample average approximation*. For newsvendor sample, say $\xi_i = 20, 40, 60, 80$, with equal probability and minimize

$$\begin{aligned} & \frac{1}{N}\sum_{i=1}^{N}[C(x,\xi_i)] \\ = & \frac{1}{4}(C(x,20)+C(x,40)+C(x,60)+C(x,80)) \end{aligned}$$

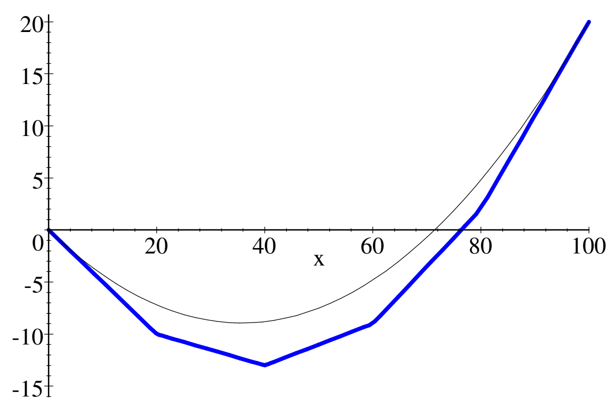

$\frac{1}{N}\sum_{i=1}^{N}[C(x,\xi_i)]$ has minimizer $\hat{x} = 40$ with flattering "optimal" value -13.0, but in the "simulation" $E[C(\hat{x},\xi)] = -8.8$, which is close to the value -8.928 from optimal $x^* = 35.714$.

Remarks

Approximating probability distributions affects the answer.

Differences obtained in optimal solution.

Differences obtained in optimal value.

Study these differences using *stability* theory.

Key statistic is the expected loss we incur by simulating the answer we get.

How do we build scenarios to make this expected loss small?

Moment matching

Our sampled distribution has the same mean 50, but its variance $= 500$

True variance of ξ is 833.

Construct a distribution with the same moments. Consider four equally-likely scenarios $\xi_i = 15, 45, 55, 85$. These give mean 50 and variance 833.

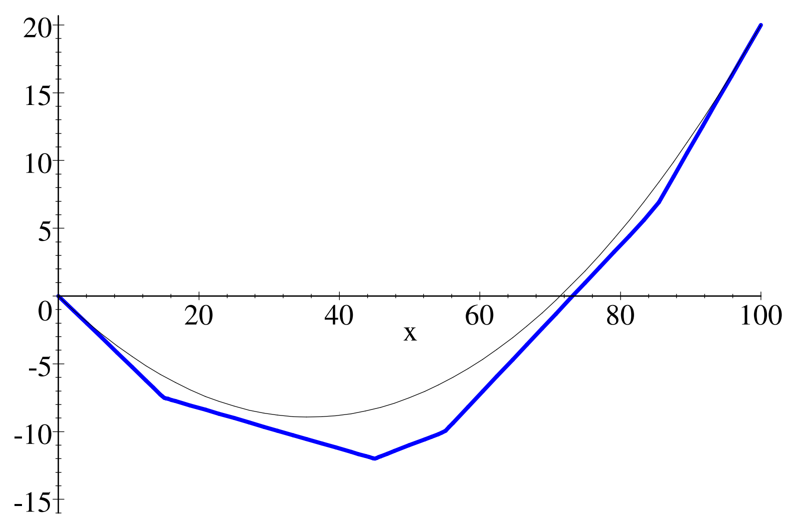

$\frac{1}{N}\sum_{i=1}^{N}[C(x,\xi_i)]$ has minimizer $\hat{x} = 45$ with "optimal" value -12.0, but in the "simulation" $E[C(\hat{x},\xi)] = -8.325$.

Jensen Bounding

Divide the interval [0,100] into subintervals and take conditional expectations to give scenario values. This guarantees that the approximate objective function lies below the true objective.

Example: choose five intervals giving $\xi_i = 10, 30, 50, 70, 90$, each with probability 0.2. (variance $=$ 800).

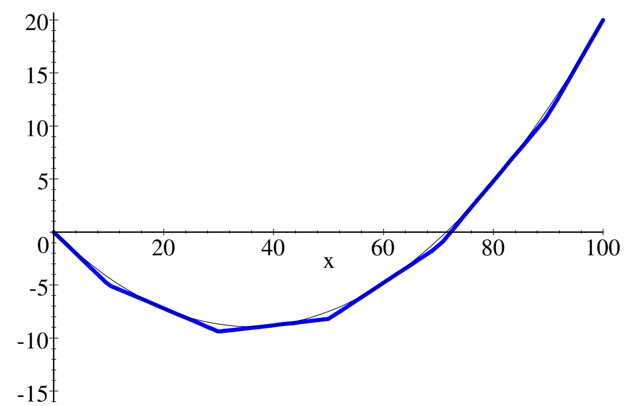

$\frac{1}{N}\sum_{i=1}^{N}[C(x,\xi_i)]$ has minimizer $\hat{x} = 30$ with "optimal" value -9.4, but in the "simulation" $E[C(\hat{x},\xi)] = -8.7$.

[Note: if uncertainty appears in objective more sophisticated bounding techniques are required: barycentric approximation]

Summary observations

To solve stochastic programming problems, optimization algorithms work with a finite number of scenarios that represent the uncertainty. Scenarios can be found by:

Exterior sampling procedures (see e.g. [Shapiro, 2003]).

Construction procedures to match moments can be used (see e.g. [Hoyland and Wallace,2001])

Scenario-reduction procedures (see e.g. [Dupačová, Gröwe-Kuska, Römisch, 2003])

Interior sampling procedures within an algorithm (see e.g. [Infanger, 1994], [Higle and Sen, 1996], [Pereira and Pinto, 1991])

Measure of scenario quality is the value of "solution" obtained when simulated using the true distribution.

What determines solution quality might not be means and variances. In the newsvendor problem the solution is determined by $F^{-1}(\cdot) \Rightarrow$ we should approximate F well in the vicinity of the solution. [See the tutorial by G. Pflug.]

Risk

Discussion above focused on expected cost. The newsvendor chooses

$$x^* = 35.714 \text{ with expected cost } -8.928$$

This gives the following plots of $C(x^*, \xi)$ and cdf $F(z)$ for total cost.

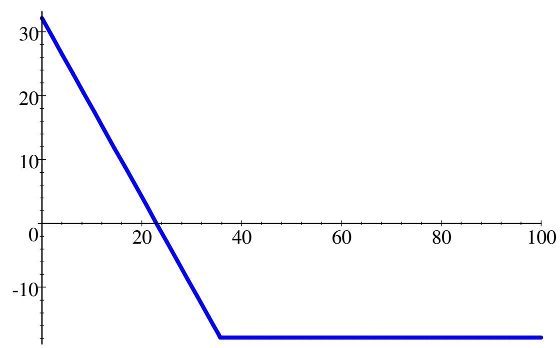

Cost as demand varies

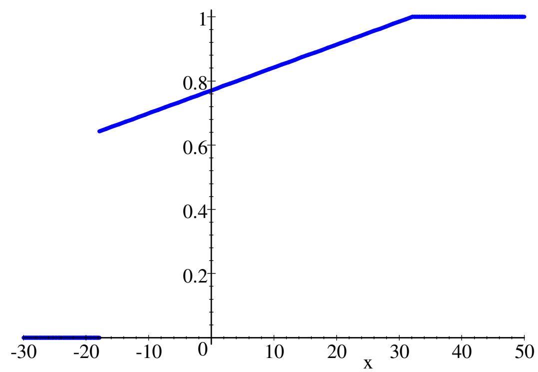

Cdf of cost

$$F(z) = \begin{cases} 0, & z < -17.857 \\ 0.6429 + 0.007142(z + 17.857), & -17.857 \leq z \leq 32.143 \\ 1, & \text{otherwise} \end{cases}$$

Risk measurement

Risk measures assign a numerical score to a probabilistic set of outcomes to allow risky values to be compared. Once a risk measure is available then this can be minimized or constrained. Common risk measures:

1. Probability of a (bad) set of outcomes - chance constraints
2. Utility functions
3. Variance
4. Semivariance
5. Mean absolute deviation
6. Value-at-Risk
7. Conditional Value at Risk

The newsvendor with chance constraints

Recall the demand for the newsvendor is uniform on $[0, 100]$. For any fixed policy x,

$$C(x,\xi) = \begin{cases} -0.5x, & \text{if } x < \xi \\ 0.9x - 1.4\xi, & \text{otherwise.} \end{cases}$$

A *chance constraint* might require that the probability of a cost above $20 be less than 0.05.

$$\begin{aligned} \Pr(C(x,\xi) &\geq 20)) = \Pr(0.9x - 1.4\xi \geq 20)) \\ &= \Pr(\xi \leq \frac{0.9x - 20}{1.4})) \\ &= \frac{0.9x - 20}{140} \end{aligned}$$

$$\begin{aligned} \Pr(C(x,\xi) \quad &\geq \quad 20)) \leq 0.05 \\ &\iff \frac{0.9x - 20}{140} \leq 0.05 \\ &\iff 0.9x \leq 27 \\ &\iff x \leq 30 \end{aligned}$$

[See the tutorial by R. Henrion.]

The newsvendor with a utility function

Utility functions allow one to take a distribution of outcomes and assign a score by taking the expectation of a nonlinear function of the random payoff z. Thus the utility of z (which seek to maximize) is

$$\mathcal{U} = \int u(z(\omega))dP(\omega)$$

For example consider the newsvendor's payoff

$$z(x,\xi) = \begin{cases} 0.5x, & \text{if } x < \xi \\ -0.9x + 1.4\xi, & \text{otherwise.} \end{cases}$$

If she has utility function $(40 + z)^{0.5}$ then the expected utility of policy x is

$$\mathcal{U} = \int_0^x (40 - 0.9x + 1.4\omega)^{0.5} \frac{1}{100} d\omega + \frac{100 - x}{100}(40 + 0.5x)^{0.5}$$

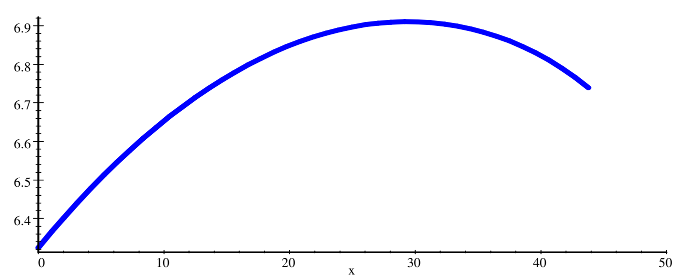

and optimal purchase is $x^* = 29.382$.

The newsvendor with variance

$$C(x,\xi) = \begin{cases} -0.5x, & \text{if } x < \xi \\ 0.9x - 1.4\xi, & \text{otherwise.} \end{cases}$$

$$E[C(x,\xi)] = (\frac{1}{100})(0.7x^2 - 50x)$$

$$E[C(x,\xi)^2] = \int_0^x (0.9x - 1.4\omega)^2 \frac{d\omega}{100} + \int_x^{100} (-0.5x)^2 \frac{d\omega}{100}$$

$$\begin{aligned} \text{var}[C(x,\xi)] &= E[C(x,\xi)^2] - (E[C(x,\xi)])^2 \\ &= 0.0065334x^3 - .000049x^4 \end{aligned}$$

$$\text{minimize } E[C(x,\xi)] + \lambda \text{var}[C(x,\xi)]$$

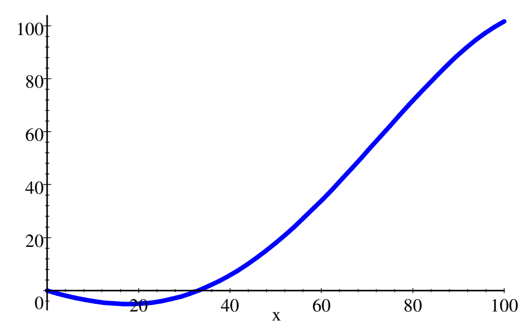

The newsvendor with Value at Risk

Recall the daily payoff for the newsvendor is

$$z(x,\xi) = \begin{cases} 0.5x, & \text{if } x < \xi \\ -0.9x + 1.4\xi, & \text{otherwise.} \end{cases}$$

If $x = 30$, the profit is plotted below as a function of demand as a solid line.

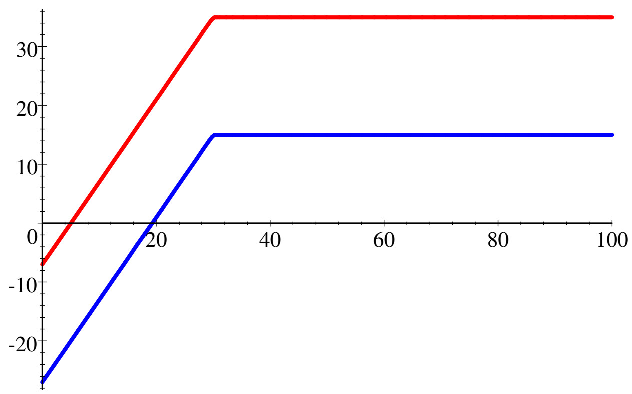

Payoff as demand varies

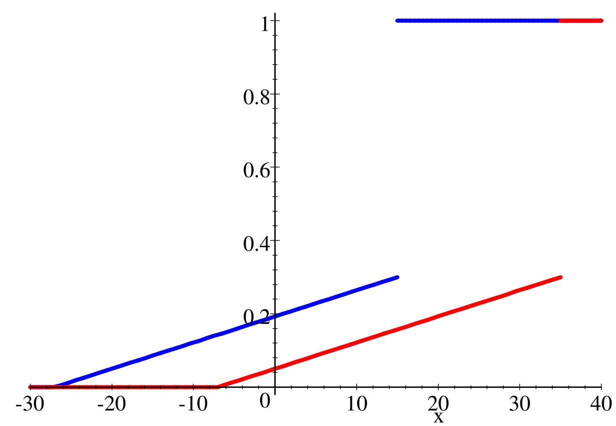

Cdf of payoff

$$\begin{aligned} VaR_a &= \inf\{r : \Pr(z(x,\xi) + r < 0) \leq \alpha\} \\ VaR_{0.05} &= \$20 \end{aligned}$$

Value at Risk

VaR was popular with banks and regulators in recent years, but leads to some paradoxes. For example, let $z_1(\omega)$ and $z_2(\omega)$ be two *independent* identically distributed payoffs with probability densities:

$$f(\omega) = \begin{cases} 0.05, & -2 \le z_i(\omega) < 0, \\ 0.9, & 0 \le z_i(\omega) \le 1, \\ 0, & \text{otherwise.} \end{cases}$$

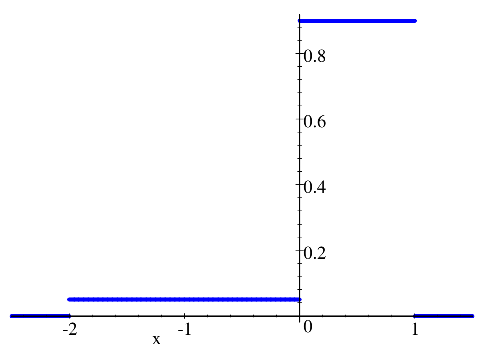

Density of payoff

$$VaR_{0.1}(z_1) = VaR_{0.1}(z_2) = 0$$

What is $VaR_{0.1}(z_1 + z_2)$?

Easy to show that $z_1 + z_2$ has density

$$f_{z_1+z_2}(\omega) = \begin{cases} 0, & z_1(\omega) + z_2(\omega) < -4 \\ 0.0025, & -4 \leq z_1(\omega) + z_2(\omega) < -2 \\ 0.04875, & -2 \leq z_1(\omega) + z_2(\omega) < -1 \\ 0.09125, & -1 \leq z_1(\omega) + z_2(\omega) < 0 \\ 0.45, & 0 \leq z_1(\omega) + z_2(\omega) < 1 \\ 0.405, & 1 \leq z_1(\omega) + z_2(\omega) < 2 \\ 0, & 2 \leq z_1(\omega) + z_2(\omega) \end{cases}$$

$$\begin{aligned} \Pr(z_1 + z_2 &< -0.49314) \\ &= (0.0025)(2) + 0.04875 + 0.09125(1 - 0.49314) \\ &= 0.1. \end{aligned}$$

$$VaR_{0.1}(z_1 + z_2) = 0.49314 > 0$$

Coherent risk measures

Example shows that $VaR_\alpha(z)$ is not a coherent risk measure as it is not *subadditive*.

A *coherent* risk measure ρ is:

subadditive $\rho(z_1 + z_2) \leq \rho(z_1) + \rho(z_2))$

positively homogeneous $\rho(\lambda z) = \lambda \rho(z)$, $\lambda \geq 0$

translation invariant $\rho(z + a) \leq \rho(z) - a$

monotonic $z_1 \geq z_2 \Rightarrow \rho(z_1) \leq \rho(z_2)$

[Artzner, Belbaen, Eber, Heath, 1997].

Conditional Value at Risk

$$CVaR_a = \frac{1}{\alpha}\int_0^{\alpha} VaR\gamma(\xi)d\gamma$$

If z has density

$$f(\omega) = \begin{cases} 0.05, & -2 \leq \omega < 0, \\ 0.9, & 0 \leq \omega \leq 1, \\ 0, & \text{otherwise.} \end{cases}$$

then

$$\begin{aligned} CVaR_{0.1}(z) &= \frac{1}{0.1}\int_0^{0.1} VaR\gamma(z)d\gamma \\ &= 1 \end{aligned}$$

Conditional Value at Risk

Recall $z_1 + z_2$ has density

$$f_{z_1+z_2}(\omega) = \begin{cases} 0, & z_1(\omega) + z_2(\omega) < -4, \\ 0.0025, & -4 \leq z_1(\omega) + z_2(\omega) < -2, \\ 0.04875, & -2 \leq z_1(\omega) + z_2(\omega) < -1, \\ 0.09125, & -1 \leq z_1(\omega) + z_2(\omega) < 0, \\ 0.45, & 0 \leq z_1(\omega) + z_2(\omega) < 1, \\ 0.405, & 1 \leq z_1(\omega) + z_2(\omega) < 2, \\ 0, & 2 \leq z_1(\omega) + z_2(\omega), \end{cases}$$

then

$$\begin{aligned} & CVaR_{0.1}(z_1 + z_2) \\ = & \frac{1}{0.1} \int_0^{0.1} VaR\gamma(z) d\gamma \\ = & \frac{(3)(0.0025) + (1.5)(0.04875) + (0.5)\,(1 + 0.49314)\,(0.09125)\,(1 - 0.49314}{0.1} \\ = & 1.1515 \end{aligned}$$

The two-stage stochastic program with (fixed) recourse

$$
\begin{array}{ll}
\text{RP: minimize} & c^\top x + E_\xi[Q(x,\xi)] \\
\text{subject to} & Ax = b, \\
 & x \geq 0.
\end{array}
$$

where

$$
Q(x, \xi(\omega)) = \min\{q^\top y \mid Wy = h(\omega) - T(\omega)x,\ y \geq 0\}.
$$

Multi-stage models

Consider a reservoir into which we may pump water from a nearby river at the start of each week at a known cost c per unit. The reservoir is drawn down each week to meet a random demand, and is filled by a random stream inflow.

Let

$$\begin{aligned} x(t) &= \text{reservoir level at start of week } t \\ y(t) &= \text{pumped amount at start of week } t \\ D(t) &= \text{random demand net inflow during week } t \\ s(t) &= \text{spill from reservoir during week } t \end{aligned}$$

$$\begin{array}{lll} \text{minimize} & \sum_1^T c(t)y(t) & \\ \text{subject to} & x(t+1) = x(t) + y(t) - s(t) - D(t), & t = 1, \ldots, T, \\ & x(t) \leq a, & t = 1, \ldots, T, \\ & x(t), s(t), y(t) \geq 0, & t = 1, \ldots, T. \end{array}$$

(Observe that we would never pump and spill in the same week.)

Multi-stage models

For a stochastic model, we need to know the information structure.

WS Model: pumping/spill decisions are made with knowledge of random demand and inflow.

Inflow,demand $\rightarrow$ pump or spill $\rightarrow$inflow,demand$\rightarrow$ pump or spill $\rightarrow \cdots$

HN Model: pumping decisions are made before random demand and inflow are realized.

Pump$\rightarrow$inflow,demand,spill$\rightarrow$ pump$\rightarrow$inflow,demand,spill$\rightarrow$ pump $\rightarrow \cdots$

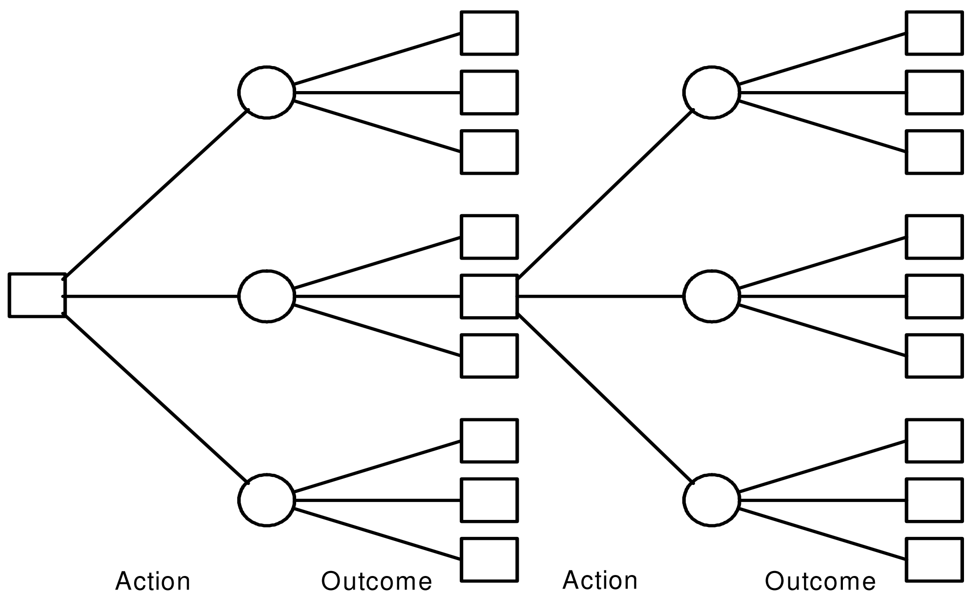

Scenario Tree

Work with a scenario tree of nodes $n \in \mathcal{N}$. Each node has a set of immediate successors, denoted n^+ and except for the root node 0, each node has a unique predecessor n^-. Each node has a probability $p(n)$, defined by the probability of the scenario for leaf nodes, and $p(n) = \sum_{i \in n^+} p(i)$ for all other nodes.

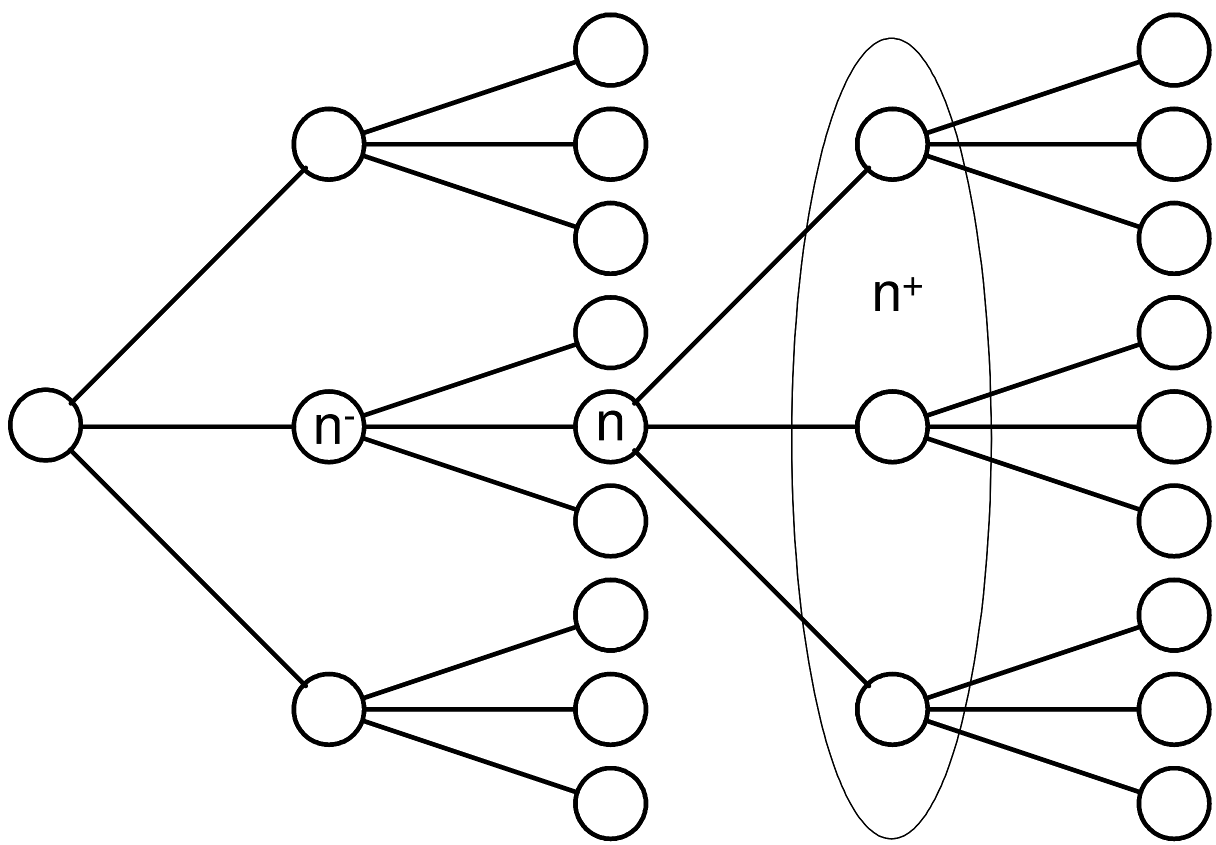

Formulation: WS

minimize $\sum_{n \in \mathcal{N}} p(n)c(n)y(n)$

subject to $x(n) = x(n^-) + y(n) - s(n) - D(n), \quad n \in \mathcal{N} \setminus \{0\},$

$x(n) \leq a, \quad n \in \mathcal{N},$

$x(n), s(n), y(n) \geq 0, \quad n \in \mathcal{N}.$

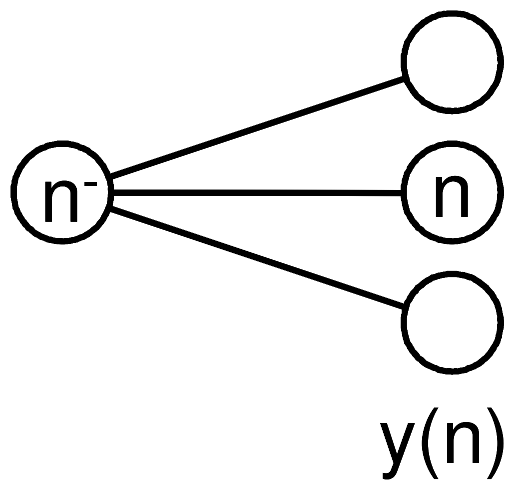

Formulation: HN

minimize $\sum_{n \in \mathcal{N}} p(n)c(n)y(n)$

subject to $x(n) = x(n^-) + y(n^-) - s(n) - D(n), \quad n \in \mathcal{N} \setminus \{0\},$

$x(n) \leq a, \quad n \in \mathcal{N},$

$x(n), s(n), y(n) \geq 0, \quad n \in \mathcal{N}.$

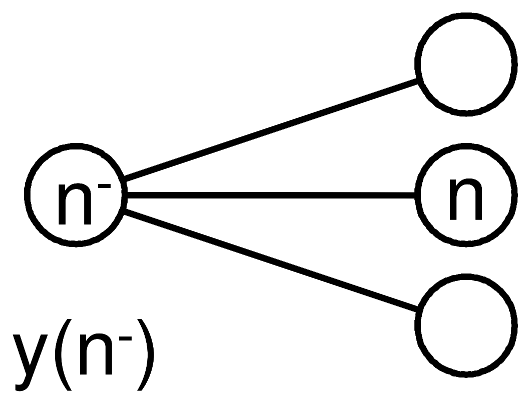

Split-variable formulation: HN Model

Let Ω denote the set of leaf nodes of the scenario tree, and for each $\omega \in \Omega$, let $n(t,\omega)$ be the unique node in the scenario tree corresponding to ω at time t.

$$
\begin{array}{lll}
\text{minimize} & \sum_{\omega\in\Omega} p(\omega) \sum_{t=1}^{t=T} c(t)y(t,\omega) & \\
\text{subject to} & x(t+1,\omega) = & \\
 & x(t,\omega) + y(t,\omega) - s(t+1,\omega) - D(t+1,\omega), & t = 1,\dots,T, \\
 & x(t,\omega) \le a, & t = 1,\dots,T, \\
 & x(t,\omega), s(t,\omega), y(t,\omega) \ge 0, & t = 1,\dots,T, \\
 & x(t,\omega_1) = x(t,\omega_2), & n(t,\omega_1) = n(t,\omega_2), \\
 & s(t,\omega_1) = s(t,\omega_2), & n(t,\omega_1) = n(t,\omega_2), \\
 & y(t,\omega_1) = y(t,\omega_2), & n(t,\omega_1) = n(t,\omega_2). \\
 & \text{(Nonanticipativity constraints)} &
\end{array}
$$

Decomposition

Multi-stage models are very hard to solve for large instances. For the pumping example over say 10 weeks with 20 demand outcomes per week, the number of nodes on the scenario tree is of the order of 10^{20}. This gives a very large linear program. Decomposition techniques must be used.

Decomposition in these problems is easiest when

1. all the randomness appears in the right-hand side of the constraints;

2. random variables are serially independent.

   For large problems sampling-based methods (e.g. [Pereira and Pinto, 1991], [Chen and Powell, 1999]) work best. These combine advantages of convexity of multi-stage LP with power of dynamic programming.

Dynamic programming

Assume serially independent $D(t)$. Let $C_t(x)$ be the minimum expected pumping cost incurred by pumping optimally from the start of week t to the start of week $T+1$, given that the reservoir currently holds x. Then $C_{T+1}(x) = 0$, and

$$C_t(x) = \min_{y \geq 0} \{cy + E_D[C_{t+1}(x + y - s - D(t))].$$

For discrete D the expectation is polyhedral and convex, and so $C_t(x)$ can be computed by solving linear programs.

Key observations:

1. For each week t, $C_t(x)$ must be determined for each possible reservoir level. Stochastic programming only computes this looking forward from the currently observed level, so there is less computation required.

2. Simulation of policies is easy for dynamic programming - hard for stochastic programming.

Non-standard models: optimization of probabilities

Sources of uncertainty affect the model

Uncertain events can be...

changes in demand for some resource (e.g. more demand for newspapers)

changes in environment (e.g. giving more inflows to a reservoir)

changes in prices for some commodity or financial instrument (e.g. increase in value of a stock)

changes in actions of competitors (e.g. what another agent is offering to a market)

Limitations of classical stochastic programming

Stochastic programming relies heavily on *convexity* to allow calculations. This requires uncertain events to be *exogenous*:

actions taken cannot affect outcomes
(e.g. increasing demand by changing prices).

actions taken cannot affect probability distributions
(e.g. learning about demand by experimenting).

actions taken cannot affect actions of competitors
(e.g. as in an oligopoly).

Desiderata for "applied" stochastic programming

All planning problems are stochastic, but some are more stochastic than others. Make sure the true problem is really stochastic before recommending stochastic programming.

Beware of the "fallacy of averages".

Be careful with correlation.

Simulation is a key tool of applied stochastic programming.

Beware of overselling stochastic programming policies: what decision-makers want is insights, tradeoffs and consistency in decision making, not just a model's "optimal" policies.

Be receptive to modeling key features of real problems that might not fit the paradigm.