Multistage stochastic programs: Time-consistency, time-inconsistency and martingale bounds

Georg Ch. Pflug, joint work with Raimund Kovacevic and Alois Pichler

April 2016

The time-consistency problem

Many concrete multistage stochastic optimization problems we typically consider (hydrostorage and electricity production management, thermal plant optimization, asset-liability management, insurance management) contain risk (utility) functionals in their objective or constraints.

If we plan for a sequence of decision times, do we have to reconsider the optimal decisions at later stages, when more information comes in?

Conditional utility functionals

We consider a probability space $(\Omega, \mathcal{F}, P)$. Let $\mathcal{F}_1$ be a σ-field contained in $\mathcal{F}$. A mapping $\mathcal{U}(\cdot|\mathcal{F}_1) : L_p(\mathcal{F}) \to L_{p'}(\mathcal{F}_1)$ is called *conditional utility mapping* (with observable information $\mathcal{F}_1$) if the following conditions are satisfied for all Y, $\lambda \in [0,1]$:

(CA1) $\mathcal{U}(Y + Y_1|\mathcal{F}_1) = \mathcal{U}(Y|\mathcal{F}_1) + Y_1$, if $Y_1 \lhd \mathcal{F}_1$
(*predictable translation-equivariance*),

(CA2) $\mathcal{U}(\lambda Y + (1-\lambda)\tilde{Y}|\mathcal{F}_1) \geq \lambda \mathcal{U}(Y|\mathcal{F}_1) + (1-\lambda)\mathcal{U}(\tilde{Y}|\mathcal{F}_1)$
*(*concavity*),*

(CA3) $Y \leq \tilde{Y}$ *implies* $\mathcal{U}(Y|\mathcal{F}_1) \leq \mathcal{U}(\tilde{Y}|\mathcal{F})$ *(*monotonicity*).*

Let $\mathcal{F}_0 = (\Omega, \emptyset)$ be the trivial σ-algebra. Then $\mathcal{U}(\cdot|\mathcal{F}_0)$ is an unconditional utility functional.
Risk is just negative utility.

The extension of the classical Fenchel-Moreau Theorem to $L_{p'}$-valued concave functionals leads to a representation of the form

$$\mathcal{U}(Y|\mathcal{F}_1) = \mathbf{inf}\{\mathbb{E}(Y\,Z|\mathcal{F}_1) - \mathcal{U}^+_{\mathcal{F}_1}(Z) : Z \in \mathcal{S}_{\mathcal{F}_1}\},$$

If $\mathcal{U}$ is positively homogeneous

(CA4) $\mathcal{U}(\lambda Y|\mathcal{F}_1) = \lambda \mathcal{U}(Y|\mathcal{F}_1))$ for $\lambda > 0$

then

$$\mathcal{U}(Y|\mathcal{F}_1) = \mathbf{inf}\{\mathbb{E}(Y\,Z|\mathcal{F}_1) : Z \in \mathcal{S}_{\mathcal{F}_1}\}.$$

The relevant Z's are called *supergradients* and $\mathcal{S}_{\mathcal{F}_1}$ is called the *supergradient set*.

Conditional functionals have 3 "arguments"

- Concavity of $Y \mapsto \mathcal{U}_P(Y|\mathcal{F}_1)$
- Convexity of $P \mapsto \mathcal{U}_P(Y|\mathcal{F}_1)$
- Monotonicity of $\mathcal{F}_1 \mapsto \mathcal{U}_P(Y|\mathcal{F}_1)$

Example: The entropic functional

- *Primal form*

$$\mathcal{U}(Y) = -\frac{1}{\gamma} \log \mathbb{E}[\exp(-\gamma Y)].$$

- *Dual form*

$$\mathcal{U}(Y) = \inf\{\mathbb{E}(Y\, Z) + \frac{1}{\gamma}\mathbb{E}(Z \log Z) : \mathbb{E}(Z) = 1, Z \geq 0\}.$$

- *Conditional form*

$$\mathcal{U}(Y|\mathcal{F}_1) = -\frac{1}{\gamma} \log \mathbb{E}[\exp(-\gamma Y)|\mathcal{F}_1].$$

- *Dual conditional form*

$$\mathcal{U}(Y|\mathcal{F}_1) = \mathbf{inf}\{\mathbb{E}(Y\, Z|\mathcal{F}_1) + \frac{1}{\gamma}\mathbb{E}(Z \log Z|\mathcal{F}_1) : \mathbb{E}(Z|\mathcal{F}_1) = 1, Z \geq 0\}$$

Example: The average value-at-risk

- *Primal form.* $\mathbb{A}\mathrm{V@R}_\alpha(Y) = \frac{1}{\alpha}\int_0^\alpha G_Y^{-1}(p)\, dp$

$$\mathbb{A}\mathrm{V@R}_0(Y) = \mathrm{ess-}\inf(Y); \mathbb{A}\mathrm{V@R}_1(Y) = \mathbb{E}(Y).$$

- *Dual form*

$$\mathbb{A}\mathrm{V@R}_\alpha(Y) = \inf\{\mathbb{E}(Y\,Z) : \mathbb{E}(Z) = 1, 0 \le Z \le 1/\alpha\}.$$

- *Conditional form*

$$\mathbb{A}\mathrm{V@R}_\alpha(Y|\mathcal{F}_1) = \mathbf{sup}\{X - \frac{1}{\alpha}\mathbb{E}([Y - X]_-) : X \lhd \mathcal{F}_1\}.$$

- *Dual conditional form*

$$\mathbb{A}\mathrm{V@R}_\alpha(Y|\mathcal{F}_1) = \mathbf{inf}\{\mathbb{E}(Y\,Z|\mathcal{F}_1) : \mathbb{E}(Z|\mathcal{F}_1) = 1, 0 \le Z \le 1/\alpha\}.$$

Other names for this functional: *conditional value-at-risk* (Rockefellar and Uryasev (2002)), *expected shortfall* (Acerbi and Tasche (2002)) and *tail value-at-risk* (Artzner et al. (1999)). The name average value-at-risk is due to Föllmer and Schied (2004).

Time consistency for final processes

Let

$$\mathcal{U}(\cdot|\mathcal{F}_1) : L_p(\Omega, \mathcal{F}, P) \to L_{p'}(\Omega, \mathcal{F}_1, P)$$

be a conditional utility-type mapping and let $\mathcal{U}(\cdot)$ be its unconditional counterpart.

Definition. (see also Artzner at al. 2007). The functional $\mathcal{U}(\cdot|\mathcal{F}_1)$ is called *time consistent*, if for all $Y, \tilde{Y} \in L_p(\Omega, \mathcal{F}, P)$ the implication

$$\mathcal{U}(Y|\mathcal{F}_1) \leq \mathcal{U}(\tilde{Y}|\mathcal{F}_1) \; a.s. \Longrightarrow \mathcal{U}(Y) \leq \mathcal{U}(\tilde{Y})$$

holds.

↯ $\mathbb{A}$V@R is not time-consistent.

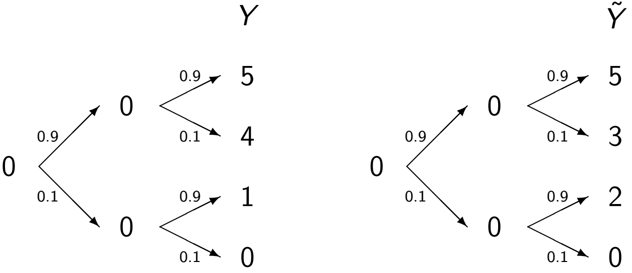

$$\mathbb{A}\text{V@R}_{0.1}(Y|\mathcal{F}_1) = (4;0) \geq (3;0) = \mathbb{A}\text{V@R}_{0.1}(\tilde{Y}|\mathcal{F}_1)$$

while

$$\mathbb{A}\text{V@R}_{0.1}(Y) = 0.9 < 1.8 = \mathbb{A}\text{V@R}_{0.1}(\tilde{Y}).$$

Theorem. (Kusuoka) Any law invariant, positively homogeneous acceptability functional $\mathcal{U}$ on L^∞ has the representation

$$\mathcal{U}(Y) = \inf_{\mu \in \mathcal{M}} \int_0^1 \mathbb{A}\mathrm{V@R}_\alpha(Y)\,\mu(\mathrm{d}\alpha), \tag{1}$$

where $\mathcal{M}$ is a set of probability measures on $[0,1]$.

Positively homogeneous utility functionals are typically not time consistent

Theorem. Suppose that the positively homogeneous functional $\mathcal{U}$ has a Kusuoka representation

$$\mathcal{U}(Y) = \inf\{\int_0^1 \mathbb{A}\text{V@R}_\alpha(Y)\, d\mu(\alpha) : \mu \in \mathcal{M}\}.$$

If

$$\inf\{\mu([\epsilon, 1-\epsilon]) : \mu \in \mathcal{M}\} > 0$$

for some $\epsilon > 0$ and

$$\sup\{\mu([0,\gamma]) : \mu \in \mathcal{M}\} \to 0$$

for $\gamma \to 0$, then $\mathcal{U}$ is not time-consistent as such, but has to be randomly decomposed for ensuring time-consistency (see later).
The only exceptions are

- the expectation
- the essential infimum
- the essential supremum

Definition. A functional $\mathcal{U}(\cdot|\mathcal{F}_1)$ is called *acceptance consistent*, if for all $Y \in L_p(\Omega, \mathcal{F}, \mu)$ the implication

$$\operatorname{ess\,inf} \mathcal{U}(Y|\mathcal{F}_1) \leq \mathcal{U}(Y)$$

holds. It is called *rejection consistent*, if

$$\operatorname{ess\,sup} \mathcal{U}(Y|\mathcal{F}_1) \geq \mathcal{U}(Y).$$

(adapted from Weber, 2006).

Definition. The functional $\mathcal{U}(\cdot|\mathcal{F}_1)$ is called

(i) *compound convex* , if for all $Y \in dom\mathcal{U}$

$$\mathcal{U}(Y) \leq \mathbb{E}\left(\mathcal{U}(Y|\mathcal{F}_1)\right).$$

(ii) *compound concave*, if for all $Y \in dom\mathcal{U}$

$$\mathcal{U}(Y) \geq \mathbb{E}\left(\mathcal{U}(Y|\mathcal{F}_1)\right).$$

Theorem.

(i) compound convexity implies rejection consistency.

(ii) compound concavity implies acceptance consistency.

Remark. Let $\mathcal{F} = (\mathcal{F}_0, \ldots, \mathcal{F}_T)$ be a filtration.

(i) $\mathcal{U}(\cdot|\cdot)$ is compound convex, iff $Y_t = \mathcal{U}(Y|\mathcal{F}_t)$ is a submartingale.

(ii) $\mathcal{U}(\cdot|\cdot)$ is compound concave, iff $Y_t = \mathcal{U}(Y|\mathcal{F}_t)$ is a supermartingale.

The $\mathbb{A}$V@R is compound convex, hence rejection consistent.

⚡ $\mathbb{A}$V@R is not acceptance consistent.

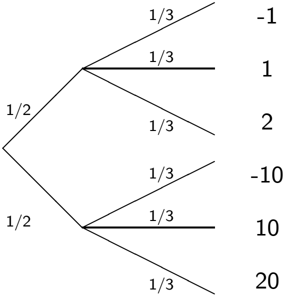

$$\mathbb{A}\text{V@R}_{2/3}(Y|\mathcal{F}_1) = 0, \qquad \mathbb{A}\text{V@R}_{2/3}(Y) = -2.$$

Definition. (Artzner (2008), Kupper (2008), Jobert (2000)) A pair of functionals $\mathcal{U}_0$ and $\mathcal{U}(\cdot|\mathcal{F}_1)$ is called *recursive*, if for all $Y \in L_p(\Omega, \mathcal{F}, \mu)$ the equation

$$\mathcal{U}_0(Y) = \mathcal{U}_0\left(\mathcal{U}_1(Y|\mathcal{F}_1)\right)$$

holds.

Of special interest are version-independent conditional functionals, which are auto-recursive (i.e. for which $\mathcal{U}_0(\cdot) = \mathcal{U}_1(\cdot|\mathcal{F}_0)$).

Examples.

- EC-functionals (i.e. functionals of the form $\mathbb{E}[\mathcal{U}(Y|\mathcal{F}_1)]$) are recursive.
- The entropic functional is auto-recursive.
- The $\mathbb{A}$V@R is not auto-recursive.

The relation between time consistency and recursivity

Theorem. (Artzner et. al., 2007) A pair $\mathcal{U}_0(\cdot)$, $\mathcal{U}_1(\cdot|\mathcal{F}_1)$ with translation equivariant $\mathcal{U}_1(\cdot|\mathcal{F}_1)$, the centered-at-zero property $\mathcal{U}_1(0|\mathcal{F}_1) = 0$ and monotonic $\mathcal{U}_0(\cdot)$ is time consistent if and only if it is recursive.

Enforcing time consistency (recursivity) by composition

Let, for each $t = 1, \dots, T$, conditional acceptability mappings $\mathcal{U}_{t-1} := \mathcal{U}(\cdot\,|\mathcal{F}_{t-1})$ be given. Introduce a multi-period probability functional $\mathcal{U}$ by compositions of the conditional acceptability mappings $\mathcal{U}_{t-1}$, $t = 1, \dots, T$, namely,

$$\begin{aligned}\mathcal{U}(Y;\mathcal{F}) &:= \mathcal{U}_0[Y_1 + \cdots + \mathcal{U}_{T-2}[Y_{T-1} + \mathcal{U}_{T-1}(Y_T)]\cdot] \\ &= \mathcal{U}_0 \circ \mathcal{U}_1 \circ \cdots \circ \mathcal{U}_{T-1}(\sum_{t=1}^{T} Y_t)\end{aligned}$$

for every $Y_t \in \mathcal{Y}_t$. (Ruszczynski and Shapiro, 2006). Notice that these functionals are recursive in a trivial way.

The nested $\mathbb{A}\mathsf{V@R}$

Example. Consider the conditional Average Value-at-Risk (of level $\alpha \in (0,1]$) as conditional acceptability mapping

$$\mathcal{U}_{t-1}(Y_t) := \mathbb{A}\mathsf{V@R}_\alpha(\cdot\,|\mathcal{F}_{t-1})$$

for every $t = 1, \ldots, T$. Then the multi-period probability functional

$$n\mathbb{A}\mathsf{V@R}_\alpha(Y;\mathcal{F}) = \mathbb{A}\mathsf{V@R}_\alpha(\cdot\,|\mathcal{F}_0) \circ \cdots \circ \mathbb{A}\mathsf{V@R}_\alpha(\cdot\,|\mathcal{F}_{T-1})(\sum_{t=1}^{T} Y_t)$$

is a concave multiperiod acceptability functional. It is called the *nested Average Value-at-Risk*.

The entropic functional

The nested entropic acceptability functional is $\mathcal{U}_0 \circ \mathcal{U}_1 \circ \cdots \circ \mathcal{U}_{T-1}(Y)$ with $\mathcal{U}_t(Y) = -\frac{1}{\gamma} \log \mathbb{E}[\exp(-\gamma Y)|\mathcal{F}_t]$, for Y nonnegative and nonvanishing.

The nested entropic functional collapses to the unconditional entropic functional.

The nested $\mathbb{A}$V@R

Example. The nested $\mathbb{A}$V@R has the following dual representation:

$$n\mathbb{A}\mathrm{V@R}_\alpha(Y;\mathcal{F}) = \inf\{\mathbb{E}[(Y_1+\cdots+Y_T)M_T] : 0 \le M_t \le \frac{1}{\alpha}M_{t-1},$$
$$\mathbb{E}(M_t|\mathcal{F}_{t-1}) = M_{t-1}, M_0 = 1, t = 1,\ldots,T\}.$$

The nested average value-at-risk $n\mathbb{A}$V@R is given by a linear stochastic optimization problem containing functional constraints.

Time consistent decisions

Let a stochastic multistage decision problem be given, for simplicity assume that it is defined on the basis of a stochastic tree. Its solution is called time-consistent, if the solutions of every conditional problem (that is when the process is conditioned to a given node at some stage t) can be chosen among the solutions of the original problem (when the decisions at times $1, \ldots, t-1$ are kept fixed).

Proposition. If the objective is a nested acceptability functional (and no other constraints are present), then the decision problem leads to time-consistent decisions.

Non-nested functionals make the trouble

Time inconsistency appears in a natural way in optimality problems. We want to find

$$\max\{\mathbb{E}(Y) : \mathbb{A}\text{V@R}_{0.05}(Y) \geq 2\} \qquad \text{or} \qquad \max \mathbb{E}(Y) + \mathbb{A}\text{V@R}_{0.05}(Y).$$

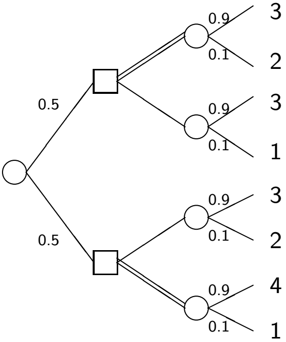

double line = optimal decision

The conditional problem given the first node:

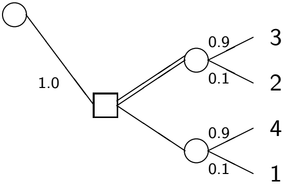

Nested functionals make the trouble

Time consistency contradicts information monotonicity.

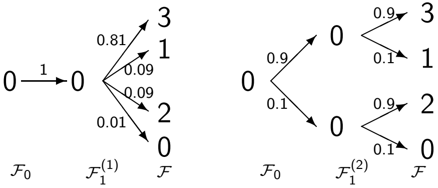

In both examples, the final income Y is the same, but in the right example, the filtration is finer. One calculates

$$\mathbb{A}\mathrm{V@R}_{0.1}[\mathbb{A}\mathrm{V@R}_{0.1}(Y|\mathcal{F}_1^{(1)})] = 0.9 > 0 = \mathbb{A}\mathrm{V@R}_{0.1}[\mathbb{A}\mathrm{V@R}_{0.1}(Y|\mathcal{F}_1^{(2)})].$$

Notice that

$$\mathbb{E}[\mathbb{A}\mathrm{V@R}_{0.1}(Y|\mathcal{F}_1^{(1)})] = \mathbb{E}[\mathbb{A}\mathrm{V@R}_{0.1}(Y|\mathcal{F}_1^{(2)})] = 0.9.$$

Information monotonicity of pos. homogeneous compositions

Proposition. Let $\mathcal{U}\left(Y;\mathcal{F}\right)$ be a composition of information monotone, positively homogeneous acceptability mappings $\mathcal{U}_t$ with supergradient sets $\mathcal{S}_t\left(\cdot\right)$. The composition is information monotone if and only if the following *nesting condition* holds:

$$\mathcal{S}_{t-1}(\mathcal{F}')\cdot\mathcal{S}_t(\mathcal{F})\subseteq S_t(\mathcal{F}')$$

if $\mathcal{F}'\subseteq\mathcal{F}$.

A composition of conditional positively homogeneous functionals

$$\mathcal{U}_0 \circ \cdots \circ \mathcal{U}_{t-1} \circ \mathcal{U}_t \circ \mathcal{U}_{t+1} \circ \cdots \circ \mathcal{U}_T$$

is information monotone only if

$$\mathcal{U}_0 = \cdots = \mathcal{U}_{t-1} = \mathbb{E} = \mathbb{A}\mathrm{V@R}_1$$

$$\mathcal{U}_{t+1} = \cdots \mathcal{U}_T = \text{essinf} \ = \mathbb{A}\mathrm{V@R}_0$$

A good example is

$$\mathcal{U}(Y_1, \ldots, Y_T) = \sum_{t=2}^{T} w_t \mathbb{E}[\mathbb{A}\mathrm{V@R}_{\alpha_t}(Y_t | \mathcal{F}_{t-1})]$$

$\mathbb{A}\mathrm{V@R}_0$ and $AVaR_1$ are also the only functionals, which are time-consistent.

Example: The nested entropic functional

Nested entropic functionals are defined as compositions $\mathcal{U}_0 \circ \mathcal{U}_1 \circ \cdots \circ \mathcal{U}_{T-1}$, where $\mathcal{U}_t(Y|\mathcal{F}_t) = -\frac{1}{\gamma_t} \log \mathbb{E}[\exp(-\gamma_t Y)|\mathcal{F}_t]$, and $\gamma_t \geq 0$ for $Y \in \mathcal{Y}$ (Now γ may depend on t !).
The correct dual pairing for the entropic functional is given by the Zygmund spaces L_{exp} and $L \log^+ L$.
Nested entropic functionals are information monotone, iff $\gamma_0 \leq \gamma_1 \leq \cdots \leq \gamma_{T-1}$.

Changing functionals as more information gets available may ensure time consistency

Decomposing the final $\mathbb{A}$V@R using random level $\mathbb{A}$V@R's

Let $\alpha \lhd \mathcal{F}_t$ be a random variable with values in [0,1]. Define the $\mathbb{A}$V@R with random level α as

$$\mathbb{A}\text{V@R}_\alpha(Y|\mathcal{F}_t) = \textbf{inf}\{\mathbb{E}(YZ|\mathcal{F}_t) : \mathbb{E}(Y|\mathcal{F}_t) = 1, 0 \leq Z; \alpha Z \leq 1\}.$$

It has an alternate characterization for $\alpha > 0$ by

$$\mathbb{A}\text{V@R}_\alpha(Y|\mathcal{F}_t) = \textbf{sup}\{Q - \frac{1}{\alpha}\mathbb{E}([Q - Y]_+|\mathcal{F}_t) : Q \lhd \mathcal{F}_t\}.$$

The $\mathbb{A}$V@R with random level obeys all properties like the usual $\mathbb{A}$V@R, i.e. translation-equivariance, concavity, monotonicity, and positive homogeneity. Moreover, $\alpha \mapsto \mathbb{A}\text{V@R}_\alpha$ is convex.

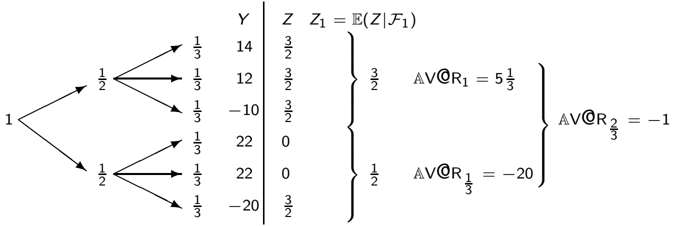

The total $\mathbb{A}\mathrm{V@R}_{\frac{2}{3}}$ is -1, while $\mathbb{A}\mathrm{V@R}_{\frac{2}{3}}(Y|\mathcal{F}_1) \equiv 1$.

Theorem. Nested decomposition of the $\mathbb{A}$V@R

Let $Y \in L^1(\mathcal{F}_T)$, $\mathcal{F}_t \subset \mathcal{F}_\tau \subset \mathcal{F}_T$.

1. For $\alpha \in [0,1]$ the Average Value-at-Risk obeys the decomposition

$$\mathbb{A}\text{V@R}_\alpha(Y) = \inf \mathbb{E}\left[Z_t \cdot \mathbb{A}\text{V@R}_{\alpha \cdot Z_t}(Y|\mathcal{F}_t)\right], \qquad (2)$$

   where the infimum is among all densities $Z_t \lhd \mathcal{F}_t$ with $0 \le Z_t$, $\alpha Z_t \le \mathbf{1}$ and $\mathbb{E}Z_t = 1$. For $\alpha > 0$ the infimum in (2) is attained.

2. Moreover if Z is the optimal dual density for the $\mathbb{A}$V@R, that is $\mathbb{A}\text{V@R}_\alpha(Y) = \mathbb{E}YZ$ with $Z \ge 0$, $\alpha Z \le \mathbf{1}$ and $\mathbb{E}Z = 1$, then $Z_t = \mathbb{E}[Z|\mathcal{F}_t]$ is the best choice in (2).

3. The *conditional* Average Value-at-Risk at random level $\alpha \lhd \mathcal{F}_t$ $(0 \le \alpha \le \mathbf{1})$ has the recursive (nested) representation

$$\mathbb{A}\text{V@R}_\alpha(Y|\mathcal{F}_t) = \mathbf{inf}\,\mathbb{E}\left[Z_\tau \cdot \mathbb{A}\text{V@R}_{\alpha \cdot Z_\tau}(Y|\mathcal{F}_\tau)\middle|\mathcal{F}_t\right], \qquad (3)$$

   where the infimum is among all densities $Z_\tau \lhd \mathcal{F}_\tau$ with $0 \le Z_\tau$, $\alpha Z_\tau \le \mathbf{1}$ and $\mathbb{E}[Z_\tau|\mathcal{F}_t] = \mathbf{1}$.

The simple special cases

- $\alpha = 0$

$$\text{essinf } Y = \text{essinf } (\text{essinf } (Y|\mathcal{F}_1))$$

- $\alpha = 1$

$$\mathbb{E}(Y) = \mathbb{E}[\mathbb{E}(Y|\mathcal{F}_1)]$$

Illustration

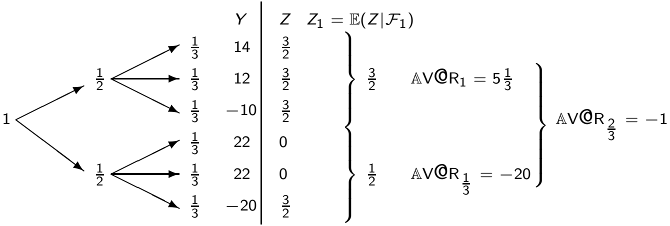

The total $\mathbb{A}$V@R is $\mathbb{A}\mathrm{V@R}_\alpha(Y) = \mathbb{E}[Z_1 \mathbb{A}\mathrm{V@R}_{\alpha Z_1}(Y|\mathcal{F}_1)] = -1$,
while $\mathbb{A}\mathrm{V@R}_{\frac{2}{3}}(Y|\mathcal{F}_1) \equiv 1$.

Notice that for $t < \tau$

$$\mathbb{A}\mathrm{V@R}_\alpha(Y|\mathcal{F}_t) \leq \mathbb{E}[\mathbb{A}\mathrm{V@R}_\alpha(Y|\mathcal{F}_\tau)|\mathcal{F}_t)] \leq \mathbb{E}(Y|\mathcal{F}_t)$$

A typical multistage decision problem

Let $H(x_0, \xi_1, \ldots, x_{T-1}, \xi_T)$ be some (concave in x) profit function depending on the random scenario process $\xi = (\xi_1, \ldots, \xi_T)$ and the decisions $x = (x_0, \ldots, x_{T-1})$
The multistage decision problem is

$$\begin{array}{rl} \text{maximize} & \mathbb{E}H(x,\xi) + \gamma \cdot \mathbb{A}\text{V@R}\left[H(x,\xi)\right] \\ \text{s.t.} & x \lhd \mathcal{F} \\ & x \in \mathcal{X}, \end{array} \tag{4}$$

where $H(x, \xi)$ is a short notation for $H(x_0, \xi_1, \ldots, x_{T-1}, \xi_T)$.

As typical for Markov decision processes, we define the *value function*

$$\mathcal{V}_t\left(x_{0:t-1},\alpha,\gamma\right) := \text{esssup}_{\ x_{t:T}}\mathbb{E}\left[H\left(x_{0:T}\right)\middle|\,\mathcal{F}_t\right]+\gamma\cdot\mathbb{A}\mathrm{V@R}_\alpha\left(H\left(x_{0:T}\right)|\mathcal{F}_t\right)$$

The value function depends on

- the decisions up to time $t-1$, $x_{0:t-1}$, where $x_{t:T}$ is chosen such that $(x_{0:T})=(x_{0:t-1},\,x_{t,T})\in\mathcal{X}$,
- the random model parameters $\alpha\lhd\mathcal{F}_t$ and $\gamma\lhd\mathcal{F}_t$ and
- the current status of the system due to the filtration $\mathcal{F}_t$.

Evaluated at initial time $t=0$ and assuming the sigma-algebra $\mathcal{F}_0$ trivial the value function relates to the initial problem as

$$\begin{aligned}\sup_{x_{0:T}}\mathbb{E}H\left(x_{0:T}\right)+\gamma\cdot\mathbb{A}\mathrm{V@R}_\alpha\left(H\left(x_{0:T}\right)\right) &= \\ &= \text{esssup}_{\ x_{0:T}}\mathbb{E}\left[H\left(x_{0:T}\right)|\mathcal{F}_0\right]+\gamma\cdot\mathbb{A}\mathrm{V@R}_\alpha\left(H\left(x_{0:T}\right)|\mathcal{F}_0\right)\\ &= \mathcal{V}_0\left([],\alpha,\gamma\right).\end{aligned}$$

Theorem. Dynamic Programming Principle. Assume that H is upper semi-continuous with respect to x and ξ valued in some convex, compact subset of $\mathbb{R}^n$.

1. The value function evaluates to

$$\mathcal{V}_T\left(x_{0:T-1},\alpha,\gamma\right)=\left(1+\gamma\right)\text{esssup}_{\,x_T}H\left(x_{0:T}\right)$$

at terminal time T.

2. For any $t<\tau$, $(t,\tau\in\mathbf{T})$ the recursive relation

$$\begin{aligned}&\mathcal{V}_t\left(x_{0:t-1},\alpha,\gamma\right)\\=\;&\text{esssup}_{\,x_{t:\tau-1}}\text{essinf}_{\,Z_{t:\tau}}\mathbb{E}\left[\mathcal{V}_\tau\left(x_{0:\tau-1},\alpha\cdot Z_{t:\tau},\gamma\cdot Z_{t:\tau}\right)\middle|\,\mathcal{F}_t\right],\end{aligned}$$

where $Z_{t:\tau}\lhd\mathcal{F}_\tau$, $0\leq Z_{t:\tau}$, $\alpha Z_{t:\tau}\leq\mathbf{1}$ and $\mathbb{E}\left[Z_{t:\tau}|F_t\right]=\mathbf{1}$, holds true.

Definition. Let $\alpha \in [0, 1]$ be a fixed level.

1. $Z = (Z_t)_{t \in \mathbf{T}}$ is a feasible (for the nonanticipativity constraints) process of densities if
   1. 1.1 Z_t is a martingale with respect to the filtration $\mathcal{F}_t$ and
   2. 1.2 $0 \leq Z_t$, $\alpha Z_t \leq \mathbb{1}$ and $\mathbb{E} Z_t = 1$ $(t \in \mathbf{T})$.
2. The intermediate densities are $Z_{t:\tau} := \frac{Z_\tau}{Z_{t-1}}$ $(0 < t < \tau)$ and $Z_{0:\tau} := Z_\tau$.

Verification Theorem. Let x and Z be feasible for the multistage decision problem.

1. Suppose that $\mathcal{W}$ satisfies

$$\mathcal{W}_T\left(x_{0:T-1}, \alpha Z_{0:T}, \gamma Z_{0:T}\right) \geq \left(1+\gamma Z_{0:T}\right) H\left(x_{0:T}\left(\xi_{0:T}\right)\right) \quad \text{and}$$
$$\mathcal{W}_t\left(x_{0:t-1}, \alpha Z_{0:t}, \gamma Z_{0:t}\right)$$
$$\geq \operatorname{esssup}_{x_t} \mathbb{E}\left[\mathcal{W}_{t+1}\left(x_{0:t}, \alpha \cdot Z_{0:t+1}, \gamma \cdot Z_{0:t+1}\right) \middle| \mathcal{F}_t\right],$$

then the process $\mathcal{W}_t\left(x_{0:t-1}, \alpha Z_{0:t}, \gamma Z_{0:t}\right)$ $(t \in \mathbf{T})$ is a super-martingale dominating $\mathcal{V}\left(x_{0:t-1}, \alpha Z_{0:t}, \gamma Z_{0:t}\right)$, i.e. $\mathcal{V} \leq \mathcal{W}$.

2. Let $\mathcal{Y}$ satisfy

$$\mathcal{Y}_T\left(x_{0:T-1}, \alpha Z_{0:T}, \gamma Z_{0:T}\right) \leq \left(1+\gamma Z_{0:T}\right) H\left(x_{0:T}\left(\xi_{0:T}\right)\right) \quad \text{and}$$
$$\mathcal{Y}_t\left(x_{0:t-1}, \alpha Z_{0:t}, \gamma Z_{0:t}\right)$$
$$\leq \operatorname{essinf}_{Z_{t+1}} \mathbb{E}\left[\mathcal{Y}_{t+1}\left(x_{0:t}, \alpha \cdot Z_{0:t+1}, \gamma \cdot Z_{0:t+1}\right) \middle| \mathcal{F}_t\right],$$

then the process $\mathcal{Y}_t\left(x_{0:t-1}, \alpha Z_{0:t}, \gamma Z_{0:t}\right)$ is a sub-martingale dominated by $\mathcal{V}\left(x_{0:t-1}, \alpha Z_{0:t}, \gamma Z_{0:t}\right)$, i.e. $\mathcal{Y} \leq \mathcal{V}$.

The Algorithm

Step 0 Let $x^0_{0:T}$ be any feasible, initial solution of the problem (4). Set $k \leftarrow 0$. Set

$$\mathcal{Y}(x^0_{0:T}) = \mathbb{E}H(x^0_{0:T}) + \gamma \mathbb{A}\text{V@R}_\alpha(H(x^0_{0:T}))$$

Step 1 Find Z^k, such that $0 \leq Z^k \leq \frac{1}{\alpha}$, $\mathbb{E}Z^k = 1$ and define

$$Z^k_t := \mathbb{E}\left(Z^k | \mathcal{F}_t\right). \tag{5}$$

A good initial choice is often Z^k satisfying

$$\mathbb{E}Z^k H\left(x^k_{0:T}\right) = \mathbb{A}\text{V@R}_\alpha\left(H\left(x^k_{0:T}\right)\right). \tag{6}$$

Step 2 (check for local improvement). Choose

$$x^{k+1}_t \in \text{argmax}_{\, x_t \lhd \mathcal{F}_t} \mathbb{E}\left[H\left(x^k_{0:T}\right)\Big| \mathcal{F}_t\right] \tag{7}$$

$$+ \gamma Z^k_t \mathbb{A}\text{V@R}_{\alpha Z^k_t}\left(H\left(x^k_{0:T}\right)\Big| \mathcal{F}_t\right) \tag{8}$$

at any arbitrary stage t and a node specified by $\mathcal{F}_t$.

Step 3 (Verification). Accept $x_{0:t}^{k+1}$ if

$$\mathcal{Y}\left(x_{0:T}^{k}\right) \leq \mathbb{E}H\left(x_{0:T}^{k+1}\right) + \gamma \mathbb{A}\text{V@R}_{\alpha}\left(H\left(x_{0:T}^{k+1}\right)\right),$$

else try another feasible Z^k (for example $Z^k \leftarrow \frac{1}{2}\left(\mathbf{1} + Z^k\right)$, $Z^k \leftarrow (1+\alpha)\mathbf{1} - \alpha Z^k$ or $Z^k = \mathbf{1}_B$ $(P(B) \geq \alpha)$) and repeat Step 2. If no direction Z^k can be found providing an improvement, then $x_{0:T}$ is already optimal. Set

$$\mathcal{Y}\left(x_{0:T}^{k+1}\right) := \mathbb{E}H\left(x_{0:T}^{k+1}\right) + \gamma \mathbb{A}\text{V@R}_{\alpha}\left(H\left(x_{0:T}^{k+1}\right)\right), \qquad (9)$$

increase $k \leftarrow k+1$ and continue with Step 1 until

$$\mathcal{Y}\left(x_{0:T}^{k+1}\right) - \mathcal{Y}\left(x_{0:T}^{k}\right) < \varepsilon,$$

where $\varepsilon > 0$ is the desired improvement in each cycle k.

The martingale citerion

Let $\mathcal{Z}$ be the famlily of feasible processes (Z_t), i.e. they are martingales w.r.t. $\mathcal{F}_t$ and $0 \le \alpha Z_t \le \mathbf{1}$, $\mathbb{E}(Z_t) = 1$ for all t.
Theorem. If (x) and (Z) are optimal, then

$$M_t(x, Z) := \mathcal{V}_t(x_{0:t-1}, \alpha Z_{o,t}, \gamma Z_{0,t})$$

is a martingale w.r.t. $\mathcal{F}_t$. Conversely, if $M_t(x, Z)$ is a martingale and the argmax sets (for x) and the argmin sets (for Z) are nonempty, then x and Z are optimal.

A numerical example: The flowergirl problem

The flower girl problem is the multistage extension of the single stage newsboy problem and an example for multistage inventory control

The flowergirl may buy in the morning of each day t a certain amount x_t of flowers. If the random demand ξ_t is not met, a penalty has to be paid; if there are left-overs, the flowers may be sold the next day, but some proportion fades and cannot be sold any more. The problem is to find the optimal ordering policy, even for non-Markovian demand processes.

The overall profit is given by a function

$$H(\xi, x(\xi)) = H(x_0, \xi_1, x_1(\xi_1), \xi_2, x_2(\xi_1, \xi_2), \ldots, \xi_T)$$

The flower-girl uses the objective

$$\mathbb{E}H(\xi, x(\xi)) + \gamma \mathbb{A}\text{V@R}_\alpha (H(\xi, x(\xi)))$$

for a suitable choice of α and γ and considers the problem

$$\begin{array}{ll} \text{maximize} & \mathbb{E}H(x) + \gamma \cdot \mathbb{A}\text{V@R}_\alpha (H(x)) \\ \text{subject to} & x \lhd \mathcal{F}, \\ & x \geq 0. \end{array} \tag{10}$$

The time complexity of the decomposed algorithm

stages	3	4	5	6	6	7	7	8	9	10
leaves	30	150	81	32	273	64	216	128	256	512
straight forward	14	639	262	58	1,970	264	3,390	1,203	5,231	26,157
Our algorithm	20	214	69	21	637	124	725	425	1,320	4,697

Run-times (CPU seconds) to solve the flowergirl problem for exemplary tree processes ξ with stages and leaves as indicated. (Our algorithm in its sequential, non-parallel implementation)

Extension for general distortion functionals

Decomposition Theorem. Let $\mathcal{U}_h(Y) = \int_0^1 G_Y^{-1}(u)h(u)\,du$ be a distortion functional.

1. $\mathcal{U}_h$ obeys the decomposition

$$\mathcal{U}_h(Y) = \inf \mathbb{E}\left[Z \cdot \mathcal{U}_Z(Y|\mathcal{F}_t)\right], \tag{11}$$

   where the infimum is among all feasible, positive random variables $Z \lhd \mathcal{F}_t$ satisfying $\mathbb{E}Z = 1$ and $h(U) \prec_{SSD} Z$ for $U \sim Uniform[0,1]$.

2. Let $\mathcal{F}_t \subset \mathcal{F}_\tau$. The utility functional obeys the nested decomposition

$$\mathcal{U}(Y|\mathcal{F}_t) = \operatorname{essinf} \mathbb{E}\left[Z_\tau \cdot \mathcal{U}_{Z_\tau}(Y|\mathcal{F}_\tau) \middle| \mathcal{F}_t\right],$$

   the essential infimum being among all feasible random variables $Z_\tau \lhd \mathcal{F}_\tau$.

			Y	Z	$\mathbb{E}(Z\|\mathcal{F})$	$Z'_{0.7}\cdot Z_{0.7}$		$Z'_{0.4}\cdot Z_{0.4}$		$\mathcal{R}_Z(Y\|\mathcal{F})$
		0.7	1	$\frac{2}{7}$		0		$\frac{5}{7}$		
	0.3	0.2	7	2.66	1	3.33	1	1.67	1	$6.07 = \frac{3}{5}\mathbb{A}\text{V@R}_{0.7} + \frac{2}{5}\mathbb{A}\text{V@R}_{0.4}$
		0.1	8	2.66		3.33		1.67		
1	0.4	0.7	0	0		0		0		
		0.3	3	$\frac{2}{3}$	0.2	0	0	1.67	0.5	$3 = 0\cdot\mathbb{A}\text{V@R}_1 + 1\cdot\mathbb{A}\text{V@R}_{0.7}$
	0.3	0.7	4	$\frac{38}{21}$		$\frac{40}{21}$		1.67		
		0.3	9	2.66	2.06	3.33	2.33	1.67	1.67	$5.94 = \frac{21}{31}\mathbb{A}\text{V@R}_{0.3} + \frac{10}{31}\mathbb{A}\text{V@R}_0$

Nested decomposition of $\mathcal{R} = \frac{3}{5}\text{U}\mathbb{A}\text{V@R}_{0.7}(Y) + \frac{2}{5}\text{U}\mathbb{A}\text{V@R}_{0.4}(Y)$.
We get

$$\mathcal{R}(Y) = \mathbb{E}[Z|\mathcal{R}_Z(Y|\mathcal{F}_t)] = 6.07\cdot 1\cdot 0.3 + 3\cdot 0.2\cdot 0.4 + 5.94\cdot 2.06\cdot 0.3 = 5.74$$

Conclusions

- Compositions of risk functionals are time consistent (but not interpretable) and not information monotone
- Final risk functionals are typically information monotone but not time consistent
- Exceptions are only the expectation and the (essential) infimum resp. supremum
- When using time-inconsistent functionals one one has to decide:
  - either to accept time-inconsistent decisions in a rolling horizon setup
  - or to accept path dependent utility functionals, i.e. decision criteria which depend on the actual path the scenario process takes.

References

- Georg Ch. Pflug, Alois Pichler. Multistage Stocahstic Optimization. Springer, 2014
- Raimund M. Kovacevic, Georg Ch. Pflug. Are time consistent valuations information-monotone? International Journal of Theoretical and Applied Finance, 17(1), 2014.
- Georg Ch. Pflug, Alois Pichler. Time-inconsistent multistage stochastic programs: Martingale bounds. European J of OR 249(1), 155-163, 2016.
- Georg Ch. Pflug, Alois Pichler. Time-consistent decisions and temporal decomposition of coherent risk functionals. to appear in: Mathematics of OR 2016.