Risk-Averse Dynamic Optimization and Control

Andrzej Ruszczyński

RUTGERS

SESO Paris

June 1, 2016

Coherent Risk Measures

Space of uncertain outcomes $\mathcal{Z} = \mathcal{L}_p(\Omega, \mathcal{F}, P)$, $p \in [1, \infty]$

A functional $\rho : \mathcal{Z} \to \overline{\mathbb{R}}$ is a coherent risk measure if it satisfies the following axioms

- Convexity: $\rho(\lambda Z + (1-\lambda)W) \le \lambda\rho(Z) + (1-\lambda)\rho(W)$ $\forall\ \lambda \in (0,1),\ Z, W \in \mathcal{Z}$
- Monotonicity: If $Z \le W$ then $\rho(Z) \le \rho(W)$, $\quad \forall\ Z, W \in \mathcal{Z}$
- Translation Equivariance: $\rho(Z + a) = \rho(Z) + a$, $\quad \forall\ Z \in \mathcal{Z},\ a \in \mathbb{R}$
- Positive Homogeneity: $\rho(\tau Z) = \tau\rho(Z)$, $\quad \forall\ Z \in \mathcal{Z},\ \tau \ge 0$

Kijima-Ohnishi (1993) – no monotonicity
Artzner-Delbaen-Eber-Heath (1999–) - spaces $\mathbb{R}^n$, $\mathcal{L}_\infty$

R.-Shapiro (2005) – spaces $\mathcal{L}_p$, ...

Good news: $\mathbb{E}[Z]$ is coherent

Space of uncertain outcomes $\mathcal{Z} = \mathcal{L}_p(\Omega, \mathcal{F}, P)$, $p \in [1, \infty]$

A functional $\rho : \mathcal{Z} \to \overline{\mathbb{R}}$ is a coherent risk measure if it satisfies the following axioms

- Convexity: $\rho(\lambda Z + (1-\lambda)W) \le \lambda\rho(Z) + (1-\lambda)\rho(W)$ $\forall\ \lambda \in (0,1),\ Z, W \in \mathcal{Z}$
- Monotonicity: If $Z \le W$ then $\rho(Z) \le \rho(W)$, $\quad \forall\ Z, W \in \mathcal{Z}$
- Translation Equivariance: $\rho(Z + a) = \rho(Z) + a$, $\quad \forall\ Z \in \mathcal{Z},\ a \in \mathbb{R}$
- Positive Homogeneity: $\rho(\tau Z) = \tau\rho(Z)$, $\quad \forall\ Z \in \mathcal{Z},\ \tau \ge 0$

Kijima-Ohnishi (1993) – no monotonicity
Artzner-Delbaen-Eber-Heath (1999–) - spaces $\mathbb{R}^n$, $\mathcal{L}_\infty$

R.-Shapiro (2005) – spaces $\mathcal{L}_p$, ...

Good news: $\mathbb{E}[Z]$ is coherent

Space of uncertain outcomes $\mathcal{Z} = \mathcal{L}_p(\Omega, \mathcal{F}, P)$, $p \in [1, \infty]$

A functional $\rho : \mathcal{Z} \to \overline{\mathbb{R}}$ is a coherent risk measure if it satisfies the following axioms

- Convexity: $\rho(\lambda Z + (1-\lambda)W) \le \lambda\rho(Z) + (1-\lambda)\rho(W)$ $\forall\ \lambda \in (0,1),\ Z, W \in \mathcal{Z}$
- Monotonicity: If $Z \le W$ then $\rho(Z) \le \rho(W)$, $\quad \forall\ Z, W \in \mathcal{Z}$
- Translation Equivariance: $\rho(Z + a) = \rho(Z) + a$, $\quad \forall\ Z \in \mathcal{Z},\ a \in \mathbb{R}$
- Positive Homogeneity: $\rho(\tau Z) = \tau\rho(Z)$, $\quad \forall\ Z \in \mathcal{Z},\ \tau \ge 0$

Kijima-Ohnishi (1993) – no monotonicity
Artzner-Delbaen-Eber-Heath (1999–) - spaces $\mathbb{R}^n$, $\mathcal{L}_\infty$

R.-Shapiro (2005) – spaces $\mathcal{L}_p$, ...

Good news: $\mathbb{E}[Z]$ is coherent

Space of uncertain outcomes $\mathcal{Z} = \mathcal{L}_p(\Omega, \mathcal{F}, P)$, $p \in [1, \infty]$

A functional $\rho : \mathcal{Z} \to \overline{\mathbb{R}}$ is a coherent risk measure if it satisfies the following axioms

- Convexity: $\rho(\lambda Z + (1-\lambda)W) \leq \lambda\rho(Z) + (1-\lambda)\rho(W)$
  $\forall\ \lambda \in (0,1),\ Z, W \in \mathcal{Z}$
- Monotonicity: If $Z \leq W$ then $\rho(Z) \leq \rho(W)$, $\quad \forall\ Z, W \in \mathcal{Z}$
- Translation Equivariance: $\rho(Z + a) = \rho(Z) + a$, $\quad \forall\ Z \in \mathcal{Z},\ a \in \mathbb{R}$
- Positive Homogeneity: $\rho(\tau Z) = \tau\rho(Z)$, $\quad \forall\ Z \in \mathcal{Z},\ \tau \geq 0$

Kijima-Ohnishi (1993) – no monotonicity
Artzner-Delbaen-Eber-Heath (1999–) - spaces $\mathbb{R}^n$, $\mathcal{L}_\infty$

R.-Shapiro (2005) – spaces $\mathcal{L}_p$, ...

Good news: $\mathbb{E}[Z]$ is coherent

Coherent Risk Measures

Space of uncertain outcomes $\mathcal{Z} = \mathcal{L}_p(\Omega, \mathcal{F}, P)$, $p \in [1, \infty]$

A functional $\rho : \mathcal{Z} \to \overline{\mathbb{R}}$ is a coherent risk measure if it satisfies the following axioms

- Convexity: $\rho(\lambda Z + (1-\lambda)W) \le \lambda\rho(Z) + (1-\lambda)\rho(W)$
  $\forall\ \lambda \in (0,1),\ Z, W \in \mathcal{Z}$
- Monotonicity: If $Z \le W$ then $\rho(Z) \le \rho(W)$, $\quad \forall\ Z, W \in \mathcal{Z}$
- Translation Equivariance: $\rho(Z + a) = \rho(Z) + a$, $\quad \forall\ Z \in \mathcal{Z},\ a \in \mathbb{R}$
- Positive Homogeneity: $\rho(\tau Z) = \tau\rho(Z)$, $\quad \forall\ Z \in \mathcal{Z},\ \tau \ge 0$

Kijima-Ohnishi (1993) – no monotonicity
Artzner-Delbaen-Eber-Heath (1999–) - spaces $\mathbb{R}^n$, $\mathcal{L}_\infty$

R.-Shapiro (2005) – spaces $\mathcal{L}_p$, ...

Good news: $\mathbb{E}[Z]$ is coherent

Coherent Risk Measures

Space of uncertain outcomes $\mathcal{Z} = \mathcal{L}_p(\Omega, \mathcal{F}, P)$, $p \in [1, \infty]$

A functional $\rho : \mathcal{Z} \to \overline{\mathbb{R}}$ is a coherent risk measure if it satisfies the following axioms

- Convexity: $\rho(\lambda Z + (1-\lambda)W) \leq \lambda\rho(Z) + (1-\lambda)\rho(W)$ $\forall\ \lambda \in (0,1),\ Z, W \in \mathcal{Z}$
- Monotonicity: If $Z \leq W$ then $\rho(Z) \leq \rho(W), \quad \forall\ Z, W \in \mathcal{Z}$
- Translation Equivariance: $\rho(Z + a) = \rho(Z) + a, \quad \forall\ Z \in \mathcal{Z},\ a \in \mathbb{R}$
- Positive Homogeneity: $\rho(\tau Z) = \tau\rho(Z), \quad \forall\ Z \in \mathcal{Z},\ \tau \geq 0$

Kijima-Ohnishi (1993) – no monotonicity
Artzner-Delbaen-Eber-Heath (1999–) - spaces $\mathbb{R}^n$, $\mathcal{L}_\infty$
R.-Shapiro (2005) – spaces $\mathcal{L}_p$, ...

Good news: $\mathbb{E}[Z]$ is coherent

Examples of Risk Measures

Mean–Semideviation (coherent): [Ogryczak–Ruszczyński, 1997]

$$\rho(Z) \triangleq \mathbb{E}[Z] + \kappa \left(\mathbb{E}\Big[\big((Z - \mathbb{E}[Z])_+\big)^p \Big] \right)^{1/p}, \qquad \kappa \in [0,1], \quad p \geq 1$$

Average Value at Risk (coherent): [Rockafellar–Uryasev, 2000-2002]

$$\begin{aligned} \text{AV@R}_\alpha^+(Z) &\triangleq \frac{1}{\alpha} \int_0^\alpha F_Z^{-1}(1-\beta)\, \mathrm{d}\beta \\ &= \min_{\eta \in \mathbb{R}} \left\{ \eta + \frac{1}{\alpha} \mathbb{E}\big[(Z-\eta)_+\big] \right\} \qquad \alpha \in (0,1] \end{aligned}$$

Entropic Risk Measure (convex, not coherent): [Howard-Matheson, 1972]

$$\rho(Z) \triangleq \frac{1}{\gamma} \ln \left(\mathbb{E}\big[e^{\gamma Z} \big] \right), \qquad \gamma > 0$$

Optimization of Risk Measures

"Minimize" over $x \in X$ a random outcome $Z_x(\omega) = f(x, \omega)$, $\omega \in \Omega$

Composite Optimization Problem

$$\min_{x \in X} \rho(Z_x) \qquad \text{(P)}$$

Dual Representation

A convex set $\mathcal{A}$ of probability measures exists, such that

$$\rho(Z_x) = \max_{\mu \in \mathcal{A}} \mathbb{E}_{\mu}[Z_x]$$

If $x \mapsto Z_x(\omega)$ is convex and $\hat{x} \in X$ is an optimal solution of (P), then a probability measure $\hat{\mu} \in \partial\rho(Z_{\hat{x}}) \subseteq \mathcal{A}$ exists such that

$$\min_{x \in X} \rho(Z_x) = \min_{x \in X} \mathbb{E}_{\hat{\mu}}[Z_x] = \min_{x \in X} \max_{\mu \in \mathcal{A}} \mathbb{E}_{\mu}[Z_x] = \max_{\mu \in \mathcal{A}} \inf_{x \in X} \mathbb{E}_{\mu}[Z_x]$$

Intriguing game against an opponent choosing the distribution $\mu \in \mathcal{A}$

How to Measure Risk of Sequences?

Probability space $(\Omega, \mathcal{F}, P)$ with filtration $\mathcal{F}_1 \subset \cdots \subset \mathcal{F}_T \subset \mathcal{F}$
Adapted sequence of random variables (costs) $Z_1, Z_2, \ldots, Z_T$
Spaces: $\mathcal{Z}_t = \mathcal{L}_{\bar{s}}(\Omega, \mathcal{F}_t, P)$, $\bar{s} \in [1, \infty)$, and $\mathcal{Z}_{t,T} = \mathcal{Z}_t \times \cdots \times \mathcal{Z}_T$

Conditional Risk Measure

A mapping $\rho_{t,T} : \mathcal{Z}_{t,T} \to \mathcal{Z}_t$ satisfying the monotonicity condition:

$$\rho_{t,T}(Z) \leq \rho_{t,T}(W) \text{ for all } Z, W \in \mathcal{Z}_{t,T} \text{ such that } Z \leq W$$

Dynamic Risk Measure

A sequence of conditional risk measures $\rho_{t,T} : \mathcal{Z}_{t,T} \to \mathcal{Z}_t$, $t = 1, \ldots, T$

$$\rho_{1,T}(Z_1, Z_2, Z_3, \ldots, Z_T) \in \mathcal{Z}_1 = \mathbb{R}$$
$$\rho_{2,T}(Z_2, Z_3, \ldots, Z_T) \in \mathcal{Z}_2$$
$$\rho_{3,T}(Z_3, \ldots, Z_T) \in \mathcal{Z}_3$$
$$\vdots$$

Evaluating Risk on a Scenario Tree

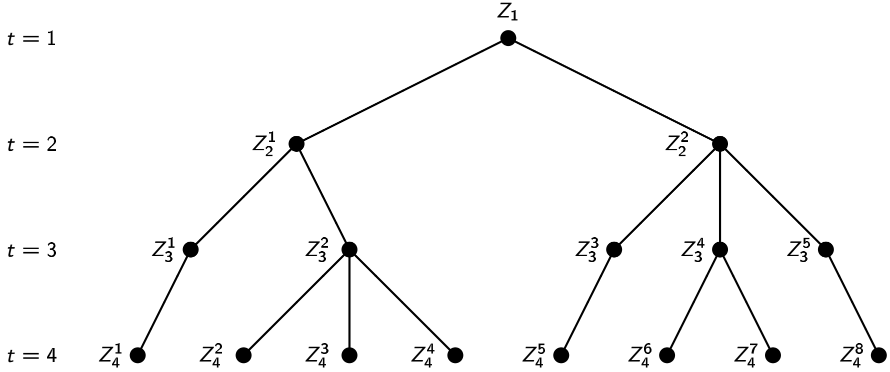

Evaluating Risk on a Scenario Tree

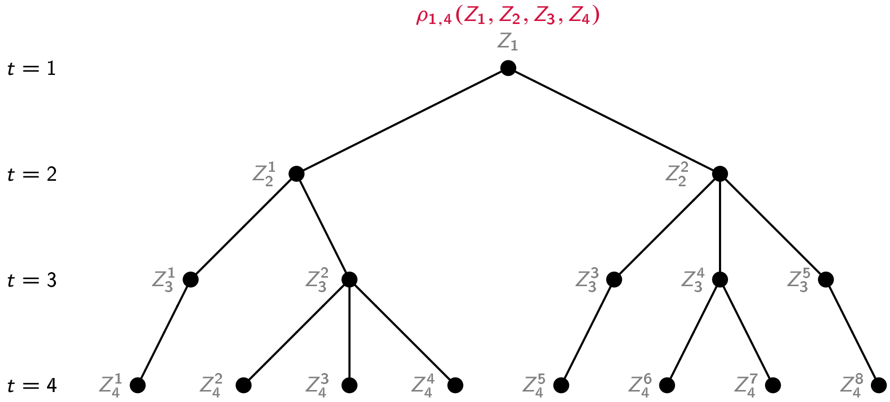

Evaluating Risk on a Scenario Tree

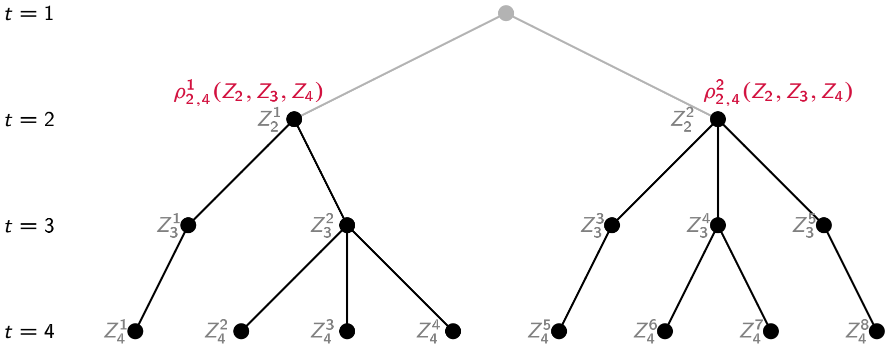

Time Consistency of Dynamic Risk Measures

A dynamic risk measure $\{\rho_{t,T}\}_{t=1}^{T}$ is time-consistent if for all $\tau < \theta$

$$Z_k = W_k,\ \ k = \tau, \ldots, \theta - 1 \quad \text{and} \quad \rho_{\theta,T}(Z_\theta, \ldots, Z_T) \le \rho_{\theta,T}(W_\theta, \ldots, W_T)$$

imply that $\quad \rho_{\tau,T}(Z_\tau, \ldots, Z_T) \le \rho_{\tau,T}(W_\tau, \ldots, W_T)$

Define $\quad \rho_t(Z_{t+1}) = \rho_{t,T}(0, Z_{t+1}, 0, \ldots, 0)$

Nested Decomposition Theorem

Suppose a dynamic risk measure $\{\rho_{t,T}\}_{t=1}^{T}$ is time-consistent, and

$$\rho_{t,T}(0, \ldots, 0) = 0$$
$$\rho_{t,T}(Z_t, Z_{t+1}, \ldots, Z_T) = Z_t + \rho_{t,T}(0, Z_{t+1}, \ldots, Z_T)$$

Then for all t we have the representation

$$\rho_{t,T}(Z_t, \ldots, Z_T) = Z_t + \rho_t\Big(Z_{t+1} + \rho_{t+1}\big(Z_{t+2} + \cdots + \rho_{T-1}(Z_T)\big)\cdots\Big)$$

Coherent One-Step Conditional Risk Measures

Stronger assumptions about one-step measures $\rho_t : \mathcal{Z}_{t+1} \to \mathcal{Z}_t$:

- Convexity: $\rho_t(\lambda Z + (1-\lambda)W) \le \lambda\rho_t(Z) + (1-\lambda)\rho_t(W)$
  $\forall\ \lambda \in (0,1),\ Z, W \in \mathcal{Z}_{t+1}$
- Monotonicity: If $Z \le W$ then $\rho_t(Z) \le \rho_t(W)$, $\ \forall\ Z, W \in \mathcal{Z}_{t+1}$
- Predictable Translation Equivariance:
  $\rho_t(Z + W) = Z + \rho_t(W)$, $\forall\ Z \in \mathcal{Z}_t,\ W \in \mathcal{Z}_{t+1}$
- Positive Homogeneity: $\rho_t(\tau Z) = \tau\rho_t(Z)$, $\forall\ Z \in \mathcal{Z}_{t+1},\ \tau \ge 0$

Scandolo ('03), Riedel ('04), R.-Shapiro ('06), Cheridito-Delbaen-Kupper ('06), Föllmer-Penner ('06), Artzner-Delbaen-Eber-Heath-Ku ('07), Pflug-Römisch ('07)

Example: Conditional Mean–Semideviation

$$\rho_t(Z_{t+1}) = \mathbb{E}\left[Z_{t+1}|\mathcal{F}_t\right] + \kappa\,\mathbb{E}\Big[\big(Z_{t+1} - \mathbb{E}\left[Z_{t+1}|\mathcal{F}_t\right]\big)_+^p\Big|\mathcal{F}_t\Big]^{1/p}$$

Here $p \in [1, \bar{p}]$ and $\kappa \in [0,1]$ may be $\mathcal{F}_t$-measurable

Coherent One-Step Conditional Risk Measures

Stronger assumptions about one-step measures $\rho_t : \mathcal{Z}_{t+1} \to \mathcal{Z}_t$:

- Convexity: $\rho_t(\lambda Z + (1-\lambda)W) \le \lambda\rho_t(Z) + (1-\lambda)\rho_t(W)$
  $\forall\ \lambda \in (0,1),\ Z, W \in \mathcal{Z}_{t+1}$
- Monotonicity: If $Z \le W$ then $\rho_t(Z) \le \rho_t(W)$, $\ \forall\ Z, W \in \mathcal{Z}_{t+1}$
- Predictable Translation Equivariance:
  $\rho_t(Z + W) = Z + \rho_t(W),\ \forall\ Z \in \mathcal{Z}_t,\ \ W \in \mathcal{Z}_{t+1}$
- Positive Homogeneity: $\rho_t(\tau Z) = \tau\rho_t(Z),\ \forall\ Z \in \mathcal{Z}_{t+1},\ \tau \ge 0$

Scandolo ('03), Riedel ('04), R.-Shapiro ('06), Cheridito-Delbaen-Kupper ('06), Föllmer-Penner ('06), Artzner-Delbaen-Eber-Heath-Ku ('07), Pflug-Römisch ('07)

Example: Conditional Mean–Semideviation

$$\rho_t(Z_{t+1}) = \mathbb{E}[Z_{t+1}|\mathcal{F}_t] + \kappa\mathbb{E}\Big[\big(Z_{t+1} - \mathbb{E}[Z_{t+1}|\mathcal{F}_t]\big)_+^p\Big|\mathcal{F}_t\Big]^{1/p}$$

Here $p \in [1, \bar{p}]$ and $\kappa \in [0,1]$ may be $\mathcal{F}_t$-measurable

Coherent One-Step Conditional Risk Measures

Stronger assumptions about one-step measures $\rho_t : \mathcal{Z}_{t+1} \to \mathcal{Z}_t$:

- Convexity: $\rho_t(\lambda Z + (1-\lambda)W) \leq \lambda\rho_t(Z) + (1-\lambda)\rho_t(W)$
  $\forall\ \lambda \in (0,1),\ Z, W \in \mathcal{Z}_{t+1}$
- Monotonicity: If $Z \leq W$ then $\rho_t(Z) \leq \rho_t(W)$, $\ \forall\ Z, W \in \mathcal{Z}_{t+1}$
- Predictable Translation Equivariance:
  $\rho_t(Z + W) = Z + \rho_t(W),\ \forall\ Z \in \mathcal{Z}_t,\ \ W \in \mathcal{Z}_{t+1}$
- Positive Homogeneity: $\rho_t(\tau Z) = \tau\rho_t(Z),\ \forall\ Z \in \mathcal{Z}_{t+1},\ \tau \geq 0$

Scandolo ('03), Riedel ('04), R.-Shapiro ('06), Cheridito-Delbaen-Kupper ('06), Föllmer-Penner ('06), Artzner-Delbaen-Eber-Heath-Ku ('07), Pflug-Römisch ('07)

Example: Conditional Mean–Semideviation

$$\rho_t(Z_{t+1}) = \mathbb{E}\left[Z_{t+1}|\mathcal{F}_t\right] + \kappa\,\mathbb{E}\Big[\big(Z_{t+1} - \mathbb{E}\left[Z_{t+1}|\mathcal{F}_t\right]\big)_+^p\,\Big|\,\mathcal{F}_t\Big]^{1/p}$$

Here $p \in [1, \bar{p}]$ and $\kappa \in [0,1]$ may be $\mathcal{F}_t$-measurable

Coherent One-Step Conditional Risk Measures

Stronger assumptions about one-step measures $\rho_t : \mathcal{Z}_{t+1} \to \mathcal{Z}_t$:

- Convexity: $\rho_t(\lambda Z + (1-\lambda)W) \le \lambda\rho_t(Z) + (1-\lambda)\rho_t(W)$
  $\forall\ \lambda \in (0,1),\ Z, W \in \mathcal{Z}_{t+1}$
- Monotonicity: If $Z \le W$ then $\rho_t(Z) \le \rho_t(W)$, $\ \forall\ Z, W \in \mathcal{Z}_{t+1}$
- Predictable Translation Equivariance:
  $\rho_t(Z + W) = Z + \rho_t(W)$, $\ \forall\ Z \in \mathcal{Z}_t,\ \ W \in \mathcal{Z}_{t+1}$
- Positive Homogeneity: $\rho_t(\tau Z) = \tau\rho_t(Z)$, $\forall\ Z \in \mathcal{Z}_{t+1},\ \tau \ge 0$

Scandolo ('03), Riedel ('04), R.-Shapiro ('06), Cheridito-Delbaen-Kupper ('06), Föllmer-Penner ('06), Artzner-Delbaen-Eber-Heath-Ku ('07), Pflug-Römisch ('07)

Example: Conditional Mean–Semideviation

$$\rho_t(Z_{t+1}) = \mathbb{E}[Z_{t+1}|\mathcal{F}_t] + \kappa\,\mathbb{E}\Big[\big(Z_{t+1} - \mathbb{E}[Z_{t+1}|\mathcal{F}_t]\big)_+^p \Big| \mathcal{F}_t\Big]^{1/p}$$

Here $p \in [1, \bar{p}]$ and $\kappa \in [0,1]$ may be $\mathcal{F}_t$-measurable

Stronger assumptions about one-step measures $\rho_t : \mathcal{Z}_{t+1} \to \mathcal{Z}_t$:

- Convexity: $\rho_t(\lambda Z + (1-\lambda)W) \le \lambda \rho_t(Z) + (1-\lambda)\rho_t(W)$ $\forall\ \lambda \in (0,1),\ Z, W \in \mathcal{Z}_{t+1}$
- Monotonicity: If $Z \le W$ then $\rho_t(Z) \le \rho_t(W)$, $\forall\ Z, W \in \mathcal{Z}_{t+1}$
- Predictable Translation Equivariance: $\rho_t(Z + W) = Z + \rho_t(W),\ \forall\ Z \in \mathcal{Z}_t,\ \ W \in \mathcal{Z}_{t+1}$
- Positive Homogeneity: $\rho_t(\tau Z) = \tau \rho_t(Z),\ \forall\ Z \in \mathcal{Z}_{t+1},\ \tau \ge 0$

Scandolo ('03), Riedel ('04), R.-Shapiro ('06), Cheridito-Delbaen-Kupper ('06), Föllmer-Penner ('06), Artzner-Delbaen-Eber-Heath-Ku ('07), Pflug-Römisch ('07)

Example: Conditional Mean–Semideviation

$$\rho_t(Z_{t+1}) = \mathbb{E}[Z_{t+1}|\mathcal{F}_t] + \kappa \mathbb{E}\Big[\big(Z_{t+1} - \mathbb{E}[Z_{t+1}|\mathcal{F}_t]\big)_+^p \big|\mathcal{F}_t\Big]^{1/p}$$

Here $p \in [1, \bar{p}]$ and $\kappa \in [0,1]$ may be $\mathcal{F}_t$-measurable

Coherent One-Step Conditional Risk Measures

Stronger assumptions about one-step measures $\rho_t : \mathcal{Z}_{t+1} \to \mathcal{Z}_t$:

- **Convexity:** $\rho_t(\lambda Z + (1-\lambda)W) \le \lambda \rho_t(Z) + (1-\lambda)\rho_t(W)$
  $\forall\ \lambda \in (0,1),\ Z, W \in \mathcal{Z}_{t+1}$
- **Monotonicity:** If $Z \le W$ then $\rho_t(Z) \le \rho_t(W)$, $\ \forall\ Z, W \in \mathcal{Z}_{t+1}$
- **Predictable Translation Equivariance:**
  $\rho_t(Z + W) = Z + \rho_t(W),\ \forall\ Z \in \mathcal{Z}_t,\ \ W \in \mathcal{Z}_{t+1}$
- **Positive Homogeneity:** $\rho_t(\tau Z) = \tau \rho_t(Z),\ \forall\ Z \in \mathcal{Z}_{t+1},\ \tau \ge 0$

Scandolo ('03), Riedel ('04), R.-Shapiro ('06), Cheridito-Delbaen-Kupper ('06), Föllmer-Penner ('06), Artzner-Delbaen-Eber-Heath-Ku ('07), Pflug-Römisch ('07)

Example: Conditional Mean–Semideviation

$$\rho_t(Z_{t+1}) = \mathbb{E}\left[Z_{t+1} | \mathcal{F}_t\right] + \kappa \mathbb{E}\left[\left(Z_{t+1} - \mathbb{E}\left[Z_{t+1} | \mathcal{F}_t\right]\right)_+^p \Big| \mathcal{F}_t\right]^{1/p}$$

Here $p \in [1, \bar{p}]$ and $\kappa \in [0,1]$ may be $\mathcal{F}_t$-measurable

Local Property

A conditional risk measure $\rho_{t,T} : \mathcal{Z}_{t,T} \to \mathcal{Z}_t$ has the local property, if for every event $A \in \mathcal{F}_t$ we have the equation

$$\rho_{t,T}(\mathbb{1}_A Z_t, \mathbb{1}_A Z_{t+1}, \ldots, \mathbb{1}_A Z_T) = \mathbb{1}_A \, \rho_{t,T}(Z_t, Z_{t+1}, \ldots, Z_T)$$

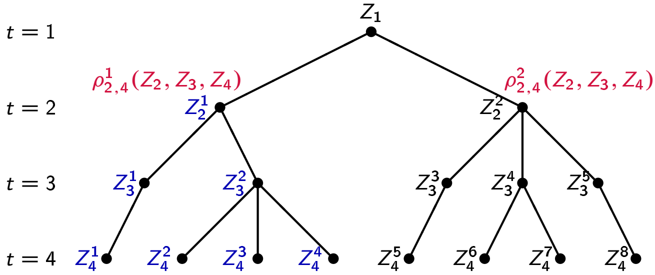

Automatic for coherent conditional risk measures

A conditional risk measure $\rho_{t,T} : \mathcal{Z}_{t,T} \to \mathcal{Z}_t$ has the local property, if for every event $A \in \mathcal{F}_t$ we have the equation

$$\rho_{t,T}(\mathbb{1}_A Z_t, \mathbb{1}_A Z_{t+1}, \dots, \mathbb{1}_A Z_T) = \mathbb{1}_A\, \rho_{t,T}(Z_t, Z_{t+1}, \dots, Z_T)$$

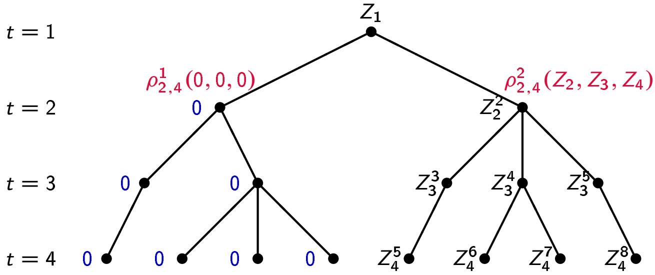

Automatic for coherent conditional risk measures

Law Invariance for Static Risk Measures

$\rho : \mathcal{Z} \to \mathbb{R}$ is law invariant if $Z \overset{\mathcal{D}}{\sim} W \;\Rightarrow\; \rho(Z) = \rho(W)$

Law Invariance for Dynamic Risk Measures

$\{\rho_{t,T} : \mathcal{Z}_T \to \mathcal{Z}_t\}_{t=1}^{T}$ is law invariant if

$$Z_T|\mathcal{F}_t \overset{\mathcal{D}}{\sim} W_T|\mathcal{F}_t, \text{ a.s. } t = 1, \dots, T-1 \\ \Rightarrow \rho_{t,T}(Z_T) = \rho_{t,T}(W_T), \; t = 1, \dots, T$$

Caution: If we defined law invariance in a static way:

$$Z_T \overset{\mathcal{D}}{\sim} W_T \;\Rightarrow\; \rho_{t,T}(Z_T) = \rho_{t,T}(W_T), \; t = 1, \dots, T$$

we would restrict the set of law invariant nd time consistent risk measures to only the expected value, the "max" measure, or the entropic measure (in continuous time) (Kupper-Schachermayer (2009), Shapiro (2012))

Probability Space: $(\Omega, \mathcal{F}, P)$ with filtration $\mathcal{F}_1 \subset \cdots \subset \mathcal{F}_T \subset \mathcal{F}$
Decision Variables: $x_t(\omega)$, $\omega \in \Omega$, $t = 1, \ldots, T$
Nonanticipativity: Each x_t is $\mathcal{F}_t$-measurable
Cost per Stage: $Z_t(x_t)$ with realizations $Z_t(x_t(\omega), \omega)$, $\omega \in \Omega$
Objective Function: Time-consistent dynamic measure of risk

Interchangeability for Time-Consistent Measures

$$\min_{x_1, x_2(\cdot), \ldots, x_T(\cdot)} \Bigg\{ Z_1(x_1) + \rho_1\Bigg(Z_2(x_2) + \rho_2\Bigg(Z_3(x_3) + \ldots \\ \cdots + \rho_{T-2}\Big(Z_{T-1}(x_{T_1} + \rho_{T-1}(Z_T(x_T))\Big) \cdots \Bigg)\Bigg)\Bigg\}$$

$$= \min_{x_1} \Bigg\{ Z_1(x_1) + \rho_1\Bigg[\min_{x_2} \Bigg(Z_2(x_2) + \rho_2\Bigg[\min_{x_3} \Bigg(Z_3(x_3) + \ldots \\ \cdots + \rho_{T-2}\Big[\min_{x_{T-1}} \big(Z_{T-1}(x_{T_1}) + \rho_{T-1}(\min_{x_T} Z_T(x_T))\big)\Big] \cdots \Bigg)\Bigg]\Bigg)\Bigg]\Bigg\}$$

Interchangeability on a Scenario Tree

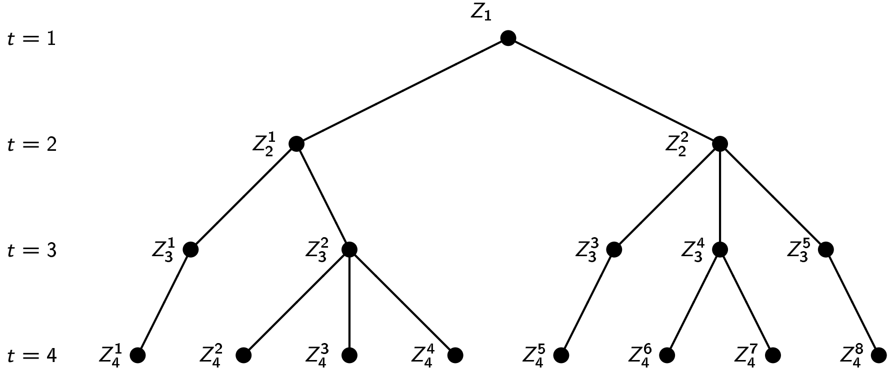

Interchangeability on a Scenario Tree

$$\min_{x_1, x_2(\cdot), x_3(\cdot), x_4(\cdot)} \rho_{1,4}(Z_1, Z_2, Z_3, Z_4)$$

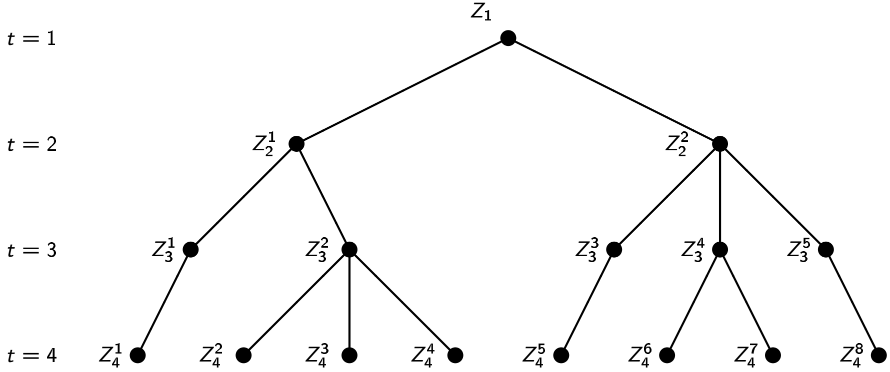

Interchangeability on a Scenario Tree

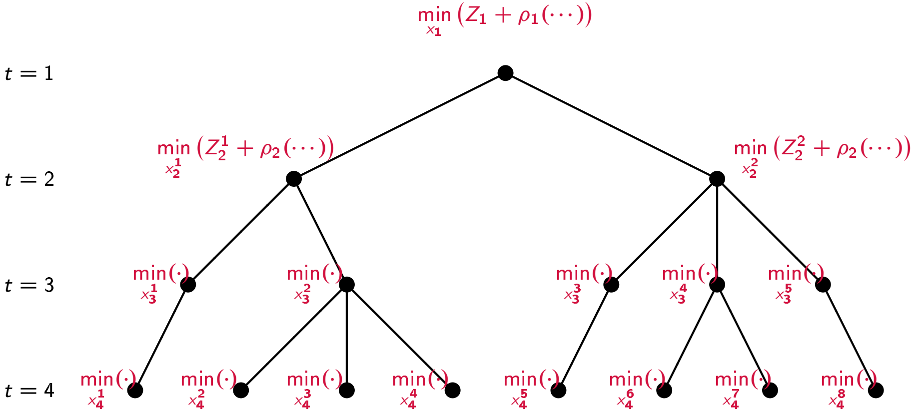

Controlled Markov Models

- State space $\mathcal{X}$ (Borel)
- Control space $\mathcal{U}$ (Borel)
- Feasible control set $U : \mathcal{X} \rightrightarrows \mathcal{U}$, $t = 1, 2, \dots$
- Controlled transition kernel $Q : \text{graph}(U) \to \mathcal{P}(\mathcal{X})$, $t = 1, 2, \dots$
  $\mathcal{P}(\mathcal{X})$ - set of probability measures on $\mathcal{X}$
- Cost functions $c : \mathcal{X} \times \mathcal{U} \to \mathbb{R}$, $t = 1, 2, \dots$
- State history $h_t = (x_1, \dots, x_t) \in \mathcal{X}^t$ (up to time $t = 1, 2, \dots$)
- Policy $\pi_t : \mathcal{X}^t \to \mathcal{U}$, $t = 1, 2, \dots$ (always supported in $U(x_t)$)
- Markov policy $\pi_t : \mathcal{X} \to \mathcal{U}$, $t = 1, 2, \dots$
  (stationary if $\pi_t = \pi_1$ for all t)

$$x_t \longrightarrow u_t = \pi_t(x_t)$$

$$(x_t, u_t) \longrightarrow x_{t+1} \sim Q(x_t, u_t)$$

Risk-Neutral Total Cost Problem

Infinite horizon expected cost problem:

$$\min_{\pi_1,\pi_2,\dots} \mathbb{E}^{\Pi}\left[\sum_{t=1}^{\infty} \alpha^{t-1} c_t(x_t,u_t)\right], \quad \alpha \in (0,1]$$

with controls $u_t = \pi_t(x_1,\dots,x_t)$

Two Cases:

Discounted models (with $\alpha < 1$) and transient models (with $\alpha = 1$)

Standard Results:

- A deterministic Markov policy is optimal
- Optimal policy can be found by dynamic programming equations

Our Intention

Introduce risk aversion to the problem by replacing
the expected value by dynamic risk measures

Using Dynamic Risk Measures for Markov Decision Processes

- Controlled Markov process x_t^{Π}, $t = 1, \ldots, T$
- Policy $\Pi = \{\pi_1, \pi_2, \ldots, \pi_T\}$ with $u_t = \pi_t(x_t)$ implies measure P^{Π}
- Cost sequence $Z_t^{\Pi} = c(x_t^{\Pi}, \pi_t(x_t^{\Pi}))$ (bounded), $t = 1, \ldots, T$,
- Dynamic time-consistent risk measure

$$J_T(\Pi) = Z_1^{\Pi} + \rho_1^{\Pi}\Big(Z_2^{\Pi} + \cdots + \rho_{T-1}^{\Pi}(Z_T^{\Pi}) \cdots \Big)$$

- Risk-averse optimal control problem: $\min_{\Pi} \lim_{T\to\infty} J_T(\Pi)$

Difficulties

- Probability measure P^{Π}, processes x_t^{Π} and Z_t^{Π} depend on policy Π
- The risk measures $\rho_t^{\Pi}(\cdot)$ depend on Π and may depend on history; no Markov policies

Idea

We only need to measure risk of random sequences that may occur

History $h_t = (x_1, \ldots, x_t)$. Process $Z_t^{\Pi}(h_t) = c(x_t, \pi_t(h_t))$, $t = 1, \ldots, T$

A family of conditional risk measures $\{\rho_{t,T}^{\Pi}\}_{t=1,\ldots,T}^{\Pi \in \varPi}$ is stochastically conditionally time-consistent if for all feasible policies Π, Π', all $1 \le t \le T-1$, and for all histories $h_t \in \mathcal{X}^t$, the relations

$$Z_t^{\Pi}(h_t) = Z_t^{\Pi'}(h_t)$$

$$\left(\rho_{t+1,T}^{\Pi}(Z_{t+1}^{\Pi}, \ldots, Z_T^{\Pi}) \middle| H_t^{\Pi} = h_t\right) \preceq_{\text{st}} \left(\rho_{t+1,T}^{\Pi'}(Z_{t+1}^{\Pi'}, \ldots, Z_T^{\Pi'}) \middle| H_t^{\Pi'} = h_t\right)$$

imply

$$\rho_{t,T}^{\Pi}(Z_t^{\Pi}, \ldots, Z_T^{\Pi})(h_t) \le \rho_{t,T}^{\Pi'}(Z_t^{\Pi'}, \ldots, Z_T^{\Pi'})(h_t)$$

The conditional stochastic order $\preceq_{\text{st}}$:

$$\begin{aligned} Q_t^{\Pi}(h_t)\Big(\{y : Z_t^{\Pi}(h_t) + \rho_{t+1,T}^{\Pi}(Z_{t+1}^{\Pi}, \ldots, Z_T^{\Pi})(h_t, y) > \eta\}\Big) \\ \le Q_t^{\Pi'}(h_t)\Big(\{y : Z_t^{\Pi'}(h_t) + \rho_{t+1,T}^{\Pi'}(Z_{t+1}^{\Pi'}, \ldots, Z_T^{\Pi'})(h_t, y) > \eta\}\Big) \end{aligned}$$

A family of process-based dynamic risk measures $\{\rho^{\Pi}_{t,T}\}^{\Pi\in\varPi}_{t=1,\dots,T}$ for a Markov decision problem is Markovian if for all Markov policies $\Pi \in \varPi$, for any measurable and bounded $c_1, \dots, c_T : \mathcal{X} \times \mathcal{U} \to \mathbb{R}$, and for all $h_t = (x_1, \dots, x_t)$ and $h'_t = (x'_1, \dots, x'_t)$ such that $x_t = x'_t$, we have

$$\rho^{\Pi}_{t,T}\big(c_t(X_t, \pi_t(X_t)), \dots, c_T(X_T, \pi_T(X_T))\big)(h_t) = \rho^{\Pi}_{t,T}\big(c_t(X_t, \pi_t(X_t)), \dots, c_T(X_T, \pi_T(X_T))\big)(h'_t).$$

If the current state x_t is the same, and the same Markov policy Π is used, then the risk is the same. The risk measure can be written as a function of the state:

$$\rho^{\Pi}_{t,T}\big(c_t(X_t, \pi_t(X_t)), \dots, c_T(X_T, \pi_T(X_T))\big)(x_t)$$

Structure of Markovian Risk Measures

For a fixed history-dependent policy Π and every $h_t \in \mathcal{X}^t$, we write

$$v_t^{c,\Pi}(h_t) = \rho_{t,T}^{\Pi}\big(c_t(X_t, \pi_t(H_t)), \dots, c_T(X_T, \pi_T(H_T))\big)(h_t)$$

If a family of process-based dynamic risk measures $\left\{\rho_{t,T}^{\Pi}\right\}_{t=1,\dots,T}^{\Pi\in\varPi}$ is Markovian, translation-invariant, and stochastically conditionally time-consistent, then there exist transition risk mappings

$$\sigma_t : \big\{\big(x, Q_t(x,u)\big) \,:\, u \in U(x),\ x \in \mathcal{X}\big\} \times \mathcal{V} \to \mathbb{R}, \quad t = 1, \dots, T-1$$

($\mathcal{V}$ - space of measurable bounded functions on $\mathcal{X}$)
such that for all $\Pi \in \varPi$, for all $t = 1, \dots, T-1$, and all $h_t \in \mathcal{X}^t$, the functional $\sigma_t\big(x_t, Q_t(x_t, \pi_t(h_t), \cdot)\big)$ is a law-invariant risk measure on $\big(\mathcal{X}, \mathcal{B}(\mathcal{X}), Q_t\big)$ and for any $c = \{c_t\}_{t=1\dots T}$, we have

$$v_t^{c,\Pi}(h_t) = c_t(x_t, \pi_t(h_t)) + \sigma_t\big(x_t, Q_t(x_t, \pi_t(h_t)), v_{t+1}^{c,\Pi}(h_t, \cdot)\big),\ t = 1 \dots T-1$$

For a fixed history-dependent policy Π and every $h_t \in \mathcal{X}^t$, we write

$$v_t^{c,\Pi}(h_t) = \rho_{t,T}^{\Pi}\big(c_t(X_t, \pi_t(H_t)), \dots, c_T(X_T, \pi_T(H_T))\big)(h_t)$$

If a family of process-based dynamic risk measures $\{\rho_{t,T}^{\Pi}\}_{t=1,\dots,T}^{\Pi\in\varPi}$ is Markovian, translation-invariant, and stochastically conditionally time-consistent, then there exist transition risk mappings

$$\sigma_t : \big\{\big(x, Q_t(x,u)\big) : u \in U(x),\ x \in \mathcal{X}\big\} \times \mathcal{V} \to \mathbb{R}, \quad t = 1, \dots, T-1$$

($\mathcal{V}$ - space of measurable bounded functions on $\mathcal{X}$)
such that for all $\Pi \in \varPi$, for all $t = 1, \dots, T-1$, and all $h_t \in \mathcal{X}^t$, the functional $\sigma_t\big(x_t, Q_t(x_t, \pi_t(h_t), \cdot)\big)$ is a law-invariant risk measure on $\big(\mathcal{X}, \mathcal{B}(\mathcal{X}), Q_t\big)$ and for any $c = \{c_t\}_{t=1\dots T}$, we have (for Markovian Π)

$$v_t^{c,\Pi}(x_t) = c_t(x_t, \pi_t(x_t)) + \sigma_t\big(x_t, Q_t(x_t, \pi_t(x_t)), v_{t+1}^{c,\Pi}(\cdot)\big), \ t = 1 \dots T-1$$

Finite Horizon Risk-Averse Control Problem

Consider a controlled Markov process $\{X_t\}$ with $u_t = \pi_t(X_1, \ldots, X_t)$.

Risk-averse optimal control problem:

$$\min_{\Pi} \; J_T(\Pi, x_1) = c_1(x_1, u_1) + \rho_1^{\Pi}\Big(c_2(X_2, u_2) + \cdots + \rho_{T-1}^{\Pi}\big(c_T(X_T, u_T) + \rho_T\big(c_{T+1}(X_{T+1})\big) \cdots \big)\Big)$$

Theorem

If the conditional measures ρ_t^{Π} are Markovian (+ general conditions), then the optimal solution is given by the dynamic programming equations:

$$v_{T+1}(x) = c_{T+1}(x), \quad x \in \mathcal{X}$$

$$v_t(x) = \min_{u \in U(x)} \Big\{ c_t(x, u) + \sigma_t\big(x, Q_t(x, u), v_{t+1}\big) \Big\}, \; x \in \mathcal{X}, \; t = T, \ldots, 1$$

Optimal Markov policy $\hat{\Pi} = \{\hat{\pi}_1, \ldots, \hat{\pi}_T\}$ - the minimizers above

Finite Horizon Risk-Averse Control Problem

Consider a controlled Markov process $\{X_t\}$ with $u_t = \pi_t(X_1, \dots, X_t)$.

Risk-averse optimal control problem:

$$\min_{\Pi} \; J_T(\Pi, x_1) = c_1(x_1, u_1) + \rho_1^{\Pi}\Big(c_2(X_2, u_2) + \cdots + \rho_{T-1}^{\Pi}\big(c_T(X_T, u_T) + \rho_T\big(c_{T+1}(X_{T+1})\big) \cdots \big)\Big)$$

Theorem

If the conditional measures ρ_t^{Π} are Markovian (+ general conditions), then the optimal solution is given by the dynamic programming equations:

$$v_{T+1}(x) = c_{T+1}(x), \quad x \in \mathcal{X}$$

$$v_t(x) = \min_{u \in U(x)} \Big\{ c_t(x,u) + \max_{\mu \in \mathcal{A}_t(x, Q_t(x,u))} \mathbb{E}_{\mu}\big[v_{t+1} \big] \Big\}, \; x \in \mathcal{X}, \; t = T, \dots, 1$$

Optimal Markov policy $\hat{\Pi} = \{\hat{\pi}_1, \dots, \hat{\pi}_T\}$ - the minimizers above

Infinite Horizon Risk (for stationary and coherent models)

Discounted risk measure $(0 < \alpha < 1)$

$$J_T^\alpha(\Pi, x) = Z_1^\Pi + \rho_1^\Pi\Big(\alpha Z_2^\Pi + \cdots + \rho_{T-1}^\Pi\big(\alpha^{T-1} Z_T^\Pi\big)\cdots\Big)$$

Optimal cost: $J^*(x) = \inf_{\Pi} \lim_{T\to\infty} J_T^\alpha(\Pi, x)$

Assume that the model is stationary, the conditional risk measures ρ_t, $t = 1, \ldots, T$, are Markovian (+ technical conditions). Then a bounded function $v : \mathcal{X} \to \mathbb{R}$ satisfies the dynamic programming equations

$$v(x) = \min_{u\in U(x)} \Big\{c(x,u) + \alpha\sigma\big(x, Q(x,u), v\big)\Big\}, \quad x \in \mathcal{X},$$

if and only if $v(\cdot) \equiv J^*(\cdot)$. Moreover, the minimizer $\pi^*(x)$, $x \in \mathcal{X}$, on the right hand side exists and defines an optimal Markov policy $\Pi^* = \{\pi^*, \pi^*, \ldots\}$.

If $\alpha = 1$ additional conditions of risk transient models

Infinite Horizon Risk (for stationary and coherent models)

Discounted risk measure $(0 < \alpha < 1)$

$$J_T^\alpha(\Pi, x) = Z_1^\Pi + \rho_1^\Pi\Big(\alpha Z_2^\Pi + \cdots + \rho_{T-1}^\Pi\big(\alpha^{T-1} Z_T^\Pi\big)\cdots\Big)$$

Optimal cost: $J^*(x) = \inf_{\Pi} \lim_{T\to\infty} J_T^\alpha(\Pi, x)$

Assume that the model is stationary, the conditional risk measures ρ_t, $t = 1, \ldots, T$, are Markovian (+ technical conditions). Then a bounded function $v : \mathcal{X} \to \mathbb{R}$ satisfies the dynamic programming equations

$$v(x) = \min_{u\in U(x)} \Big\{ c(x,u) + \alpha \max_{\mu\in\mathcal{A}(x,Q(x,u))} \mathbb{E}_\mu[v] \Big\}, \quad x \in \mathcal{X},$$

if and only if $v(\cdot) \equiv J^*(\cdot)$. Moreover, the minimizer $\pi^*(x)$, $x \in \mathcal{X}$, on the right hand side exists and defines an optimal Markov policy $\Pi^* = \{\pi^*, \pi^*, \ldots\}$.

If $\alpha = 1$ additional conditions of risk transient models

Solution Methods (for coherent mappings)

Value iteration

$$v^{k+1}(x) = \min_{u \in U(x)} \Big\{ c(x,u) + \alpha\sigma\big(x, Q(x,u), v^k\big) \Big\}, \quad x \in \mathcal{X}, \quad k = 1, 2, \dots$$

Policy iteration

- For $k = 0, 1, 2, \dots$, given a stationary Markov policy $\{\pi^k, \pi^k, \dots\}$, find the value function v^k by solving the nonsmooth equation

$$v(x) = c(x, \pi^k(x)) + \alpha\sigma\big(x, Q(x, \pi^k(x)), v\big), \quad x \in \mathcal{X},$$

  by specialized Newton's method or convex programming

- Find the next policy $\pi^{k+1}(\cdot)$ by one-step optimization

$$\pi^{k+1}(x) = \operatorname*{argmin}_{u \in U(x)} \Big\{ c(x,u) + \alpha\sigma\big(x, Q(x,u), v^k\big) \Big\}, \quad x \in \mathcal{X}$$

Policy Evaluation by Convex Optimization

How to solve $\quad v(x) = c(x, \pi^k(x)) + \alpha\, \sigma\big(x, Q(x, \pi^k(x)), v\big), \quad x \in \mathcal{X}$?

Newton Method

For $\ell = 1, 2, \dots$

- Find $M_\ell(x) \in \partial\sigma\big(x, Q(x, \pi^k(x)), v_\ell\big), \quad x \in \mathcal{X}$
- Solve the linear system $v(x) = c(x, \pi^k(x)) + \alpha\, \mathbb{E}_{M_\ell(x)}\big[v(\cdot)\big], \quad x \in \mathcal{X}$ to get $v_{\ell+1}$

Convex Programming (for finite state and control spaces)

$$\begin{aligned} \min_{v} \quad & \textstyle\sum_{x\in\mathcal{X}} v(x) \\ \text{s.t.} \quad & v(x) \geq c(x, \pi^k(x)) + \alpha\, \sigma\big(x, Q(x, \pi^k(x)), v(\cdot)\big), \quad x \in \mathcal{X} \end{aligned}$$

We use convex programming for policy evaluation only. Inclusion of minimization with respect to u would make the problem non-convex.

Filtered probability space $(\Omega, \mathscr{F}, P, \mathbb{F})$
Filtration $\mathbb{F}$ is generated by n-dimensional Brownian motion $\{W_t\}_{t\in[0,T]}$
Controlled diffusion process with initial value $\zeta \in \mathcal{L}_2(\Omega, \mathscr{F}_t, \mathbb{P}; \mathbb{R}^n)$:

$$dX_s^{t,\zeta;u} = b(s, X_s^{t,\zeta;u}, u_s)\, ds + \sigma(s, X_s^{t,\zeta;u}, u_s)\, dW_s, \quad s \in [t, T],$$
$$X_t^{t,\zeta;u} = \zeta,$$

with functions $b : [0, T] \times \mathbb{R}^n \times U \to \mathbb{R}^n$ and $\sigma : [0, T] \times \mathbb{R}^n \times U \to \mathbb{R}^{n\times d}$.
Cost rate $c : [0, T] \times \mathbb{R}^n \times U \to \mathbb{R}$; Final cost $\Psi : \mathbb{R}^n \to \mathbb{R}$.
Cost accumulated in the interval $[t, T]$

$$\xi_{t,T}(u, \zeta) := \int_t^T c(s, X_s^{t,\zeta;u}, u_s)\, ds + \Psi(X_T^{t,\zeta;u}), \quad \text{a.s.}$$

All functions are assumed to be sufficiently regular (Lipschitz or bounded).

Risk-Averse Control Problem for Diffusion Processes

$$\min_{u(\cdot)\in\mathcal{U}} \varrho_{0,T}\left(\int_0^T c(s, X_s^{0,x_0;u}, u_s)\, ds + \Psi(X_T^{0,x_0;u})\right)$$

$$dX_s^{0,x_0;u} = b(s, X_s^{0,x_0;u}, u_s)\, ds + \sigma(s, X_s^{0,x_0;u}, u_s)\, dW_s, \quad s \in [0, T]$$

where $\{\varrho_{t,r}\}_{0\le t\le r\le T}$ is a dynamic risk measure on the space of square-integrable adapted processes on $[0, T] \times \Omega$

Time consistency: $\varrho_{t,r}(Y_r) = \varrho_{t,s}\big(\varrho_{s,r}(Y_r)\big)$, for all $t \le s \le r$
Local property: $\varrho_{t,r}(\mathbb{1}_A Y_r) = \mathbb{1}_A\, \varrho_{t,r}(Y_r)$, for all events $A \in \mathcal{F}_t$.

Structure of $\varrho_{t,r}(\cdot)$ [Coquet, Hu, Mémin, Peng (2002)]

Under minor conditions, a generator $g : [0, T] \times \mathbb{R} \times \mathbb{R}^n \to \mathbb{R}$ exists, such that $\varrho_{t,r}(\xi) = Y_t$, where (Y, Z) solve backward stochastic differential equation

$$-dY_s = g(s, Y_s, Z_s)\, ds - Z_s'\, dW_s, \quad s \in [t, r], \quad Y_r = \xi.$$

If g is convex, pos.-homogeneous, independent of Y, then ϱ is coherent

Value function

$$V(t,x) = \inf_{u(\cdot)\in\mathcal{U}} \varrho_{t,T}\left(\int_t^T c(s, X_s^{t,x;u}, u_s)\, ds + \Psi(X_T^{t,x;u})\right)$$

Dynamic Programming Equation

For any $(t,x) \in [0,T) \times \mathbb{R}^n$ and all $r \in [t,T]$, we have

$$V(t,x) = \inf_{u(\cdot)\in\mathcal{U}} \varrho_{t,r}\left[\int_t^r c\big(s, X_s^{t,x;u}, u_s\big)\, ds + V\big(r, X_r^{t,x;u}\big)\right].$$

Related decoupled forward–backward system:

$$\begin{aligned}
dX_s^{t,x;u} &= b(s, X_s^{t,x;u}, u_s)\, ds + \sigma(s, X_s^{t,x;u}, u_s)\, dW_s, \quad s \in [t,r] \\
X_t^{t,x;u} &= x \\
-dY_s^{t,x;u} &= \big[c(s, X_s^{t,x;u}, u_s) + g(s, Z_s^{t,x;u})\big]\, ds - Z_s^{t,x;u}\, dW_s, \quad s \in [t,r] \\
Y_r^{t,\zeta;u} &= V\big(r, X_r^{t,x;u}\big)
\end{aligned}$$

Laplacian operator:

$$\big[\mathcal{L}^{\alpha} w\big](t,x) = \partial_t w(t,x) +$$
$$+ \sum_{i,j=1}^{n} \frac{1}{2}\big(\sigma(t,x,\alpha)\sigma(t,x,\alpha)^{\top}\big)_{ij}\partial_{x_i x_j} w(t,x) + \sum_{i=1}^{n} b_i(t,x,\alpha)\partial_{x_i} w(t,x).$$

Risk-Averse HJB Equation

On the space $\mathcal{C}_{\mathrm{b}}^{1,2}([0,T]\times\mathbb{R}^n)$, we consider the following equation

$$\min_{\alpha\in U}\Big\{c(t,x,\alpha) + \big[\mathcal{L}^{\alpha} v\big](t,x) + g\Big(t,[\mathcal{D}_x v\cdot\sigma^{\alpha}](t,x)\Big)\Big\} = 0 \quad \forall (t,x)$$
$$v(T,x) = \Psi(T,x), \quad x\in\mathbb{R}^n.$$

If the functions b and σ are bounded, then the value function $V(t,x)$ is a viscosity solution of the risk-averse HJB equation.
Conversely, if the HJB equation has a solution, it is equal to $V(t,x)$.

Extensions and Current Research on Risk-Averse Control

- Partially Observable Markov Processes (with Jingnan Fan)
  - process-based risk measures
  - transition risk mappings on the observable part
  - dynamic programming equations
- Risk-Averse Control of Clinical Trials (with Darinka Dentcheva and Curtis McGinity)
  - new dynamic models of clinical trials
  - approximate dynamic programming methods
- Risk-Averse Control of Continuous-Time Markov Chains (with Darinka Dentcheva)
  - transition risk is infinesimal short intervals
  - set-valued analysis $\rightrightarrows$ risk multigenerators
  - risk-averse Kolmogorov-like equations
- Risk-Averse Control of Diffusion Processes (with Jianing Yao)
  - extension of the Hamilton-Jacobi-Bellman equation
  - approximation by risk-averse Markov chains

References

- A. Ruszczyński, Risk-averse dynamic programming for Markov decision processes, *Mathematical Programming*, Series B 125 (2010) 235–261
- Ö. Çavuş and A. Ruszczyński, Computational methods for risk-averse undiscounted transient Markov models, *Operations Research*, 62 (2), 2014, 401–417.
- Ö. Çavuş and A. Ruszczyński, Risk-averse control of undiscounted transient Markov models, *SIAM J. on Control and Optimization*, 52(6), 2014, 3935–3966.
- J. Fan and A. Ruszczyński, Process-based risk measures for observable and partially observable discrete-time controlled systems, *submitted for publication*.
- D. Dentcheva and A. Ruszczyński, Risk measures for continuous-time Markov chains, *submitted for publication*.
- A. Ruszczyński and J. Yao, Risk-averse control of diffusion processes, *submitted for publication*.