Risk-Averse Dynamic Optimization

Andrzej Ruszczyński

RUTGERS

Research supported by the NSF award CMMI-0965689

Outline

1. Risk-Averse Preferences
2. Coherent Risk Measures
3. Dynamic Risk Measurement
4. Markov Risk Measures
5. Risk-Averse Control Problems
6. Solution Methods

Uncertain Outcomes and Risk

Why Probabilistic Models?

- Wealth of results of probability theory
- Connection to real data via statistics
- Universal language (engineering, economics, medicine, . . .)

- Probability space $(\Omega, \mathcal{F}, P)$
- Decision space $\mathcal{X}$
- Random outcome (*e.g.*, cost) $Z_x(\omega)$, $\quad Z : \mathcal{X} \times \Omega \to \mathbb{R}$

Expected Value Model

$$\min_x \; \mathbb{E}[Z_x] = \int_\Omega Z_x(\omega)\, P(d\omega)$$

It optimizes the outcome on average (Law of Large Numbers?)

What is Risk?

Existence of unlikely and undesirable outcomes - high $Z_x(\omega)$ for some ω

Classical Utility Models

Expected Utility Models (von Neumann and Morgenstern, 1944)

$$\min_{x \in X} \mathbb{E}\left[u(Z_x)\right] \qquad \left(= \int_{\Omega} u\big(Z_x(\omega)\big)\, dP(\omega) \right)$$

$u : \mathbb{R} \to \mathbb{R}$ is a nondecreasing disutility function

Rank Dependent Utility (Distortion) Models (Quiggin, 1982; Yaari, 1987)

$$\min_{x \in X} \int_0^1 F_{Z_x}^{-1}(p)\, dw(p) \qquad F_{Z_x}^{-1}(\cdot) \text{ - quantile function}$$

$w : [0, 1] \to \mathbb{R}$ is a nondecreasing rank dependent utility function

Existence of utility functions is derived from systems of axioms, but in practice they are difficult to elicit

Mean–Risk Models

Two Objectives

- Minimize the expected outcome, the mean $\mathbb{E}[Z_x]$
- Minimize a scalar measure of uncertainty of Z_x, the risk $r[Z_x]$

$$r_1[Z] = \mathbb{V}\mathrm{ar}[Z] \qquad \text{(Markowitz' model)}$$

$$r_2[Z] = \left(\mathbb{E}[(Z - \mathbb{E}Z)_+^s]\right)^{1/s} \qquad \text{(semideviation)}$$

$$r_3[Z] = \min_\eta \mathbb{E}\Big[\max\Big(\eta - Z, \frac{p}{1-p}(Z-\eta)\Big)\Big] \qquad \text{(deviation from quantile)}$$

Mean–Risk Optimization

$$\min_{x\in X} \ \rho(Z_x) = \mathbb{E}[Z_x] + \kappa\, r[Z_x], \qquad 0 \le \kappa \le \kappa_{\max}$$

Interesting application of parametric optimization

$r[Z_x]$ is nonlinear w.r.t. probability and possibly nonconvex in x

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

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

⋮

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

⋮

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

$\vdots$

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

⋮

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

$\vdots$

Conjugate Duality of Risk Measures

Pairing of a linear topological space $\mathcal{Z}$ with a linear topological space $\mathcal{Y}$ of regular signed measures on Ω with the bilinear form

$$\langle \mu, Z \rangle = \mathbb{E}_{\mu}[Z] = \int_{\Omega} Z(\omega)\, \mu(d\omega)$$

We assume standard conditions on pairing and the polarity: $(\mathcal{Z}_+)^\circ = \mathcal{Y}_-$

Dual Representation Theorem

If $\rho : \mathcal{Z} \to \overline{\mathbb{R}}$ is a lower semicontinuous* coherent risk measure, then

$$\rho(Z) = \sup_{\mu \in \mathcal{A}} \langle \mu, Z \rangle, \quad \forall Z \in \mathcal{Z}$$

with a convex $\mathcal{A} \subset \mathcal{P}$ (set of probability measures in $\mathcal{Y}$).

Delbaen (2001), Föllmer–Schied (2002), R.–Shapiro (2005), Rockafellar–Uryasev–Zabarankin (2006), . . .

* *Lower semicontinuity is automatic if ρ is finite and $\mathcal{Z}$ is a Banach lattice*

Optimization of Risk Measures

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

Composite Optimization Problem

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

Theorem (R.–Shapiro, 2005)

Let $x \mapsto Z_x(\omega)$ be convex and $\rho(\cdot)$ be coherent. Suppose that $\hat{x} \in X$ is an optimal solution of (P) and $\rho(\cdot)$ is continuous at $Z_{\hat{x}}$. Then there exists a probability measure $\hat{\mu} \in \partial\rho(Z_{\hat{x}}) \subset \mathcal{A}$ such that $\hat{x}$ solves

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

We also have the duality relation:

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

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) $Z_1, Z_2, \dots, Z_T$
Spaces: $\mathcal{Z}_t = \mathcal{L}_p(\Omega, \mathcal{F}_t, P)$, $p \in [1, \infty]$, 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:

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

Dynamic Risk Measure

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

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

Evaluating Risk on a Scenario Tree

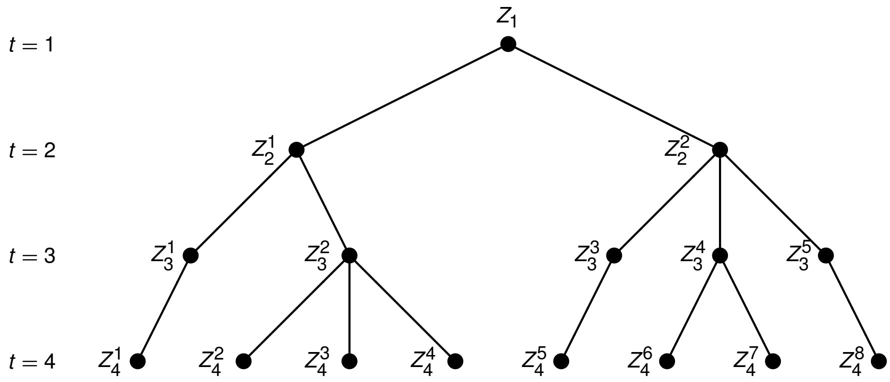

Evaluating Risk on a Scenario Tree

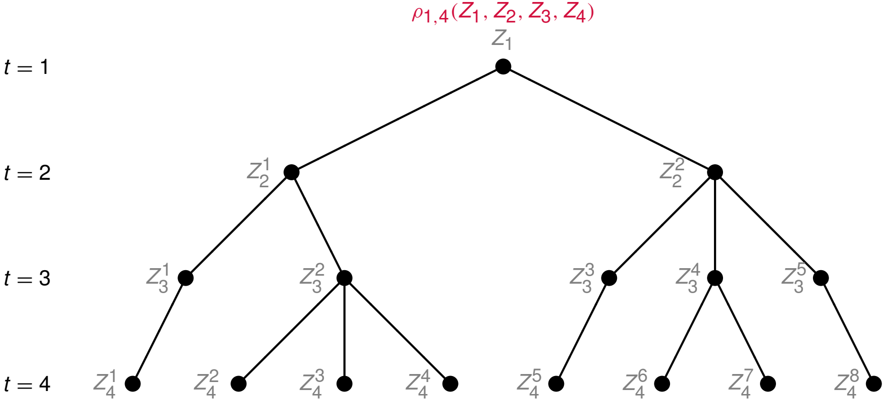

Evaluating Risk on a Scenario Tree

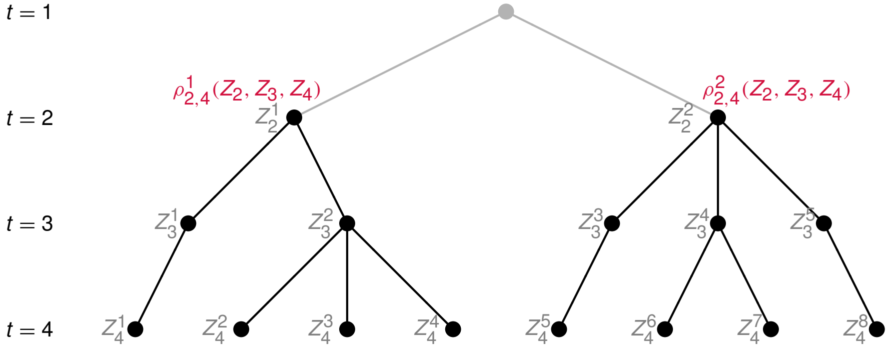

Time Consistency of Dynamic Risk Measures

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

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

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

Define $\rho_{\tau,\theta}(Z_\tau, \dots, Z_\theta) = \rho_{\tau,T}(Z_\tau, \dots, Z_\theta, 0, \dots, 0), \quad 1 \le \tau \le \theta \le T$

Risk-Averse Equivalence Theorem

Suppose $\{\rho_{t,T}\}_{t=1}^{T}$ satisfies the conditions:

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

Then it is time-consistent if and only if for all $\tau \le \theta$:

$$\rho_{\tau,T}\big(Z_\tau, \dots, Z_\theta, \dots, Z_T\big) = \rho_{\tau,\theta}\big(Z_\tau, \dots, Z_{\theta-1}, \rho_{\theta,T}(Z_\theta, \dots, Z_T)\big)$$

Collapsing Subtrees by Conditional Risk Measures

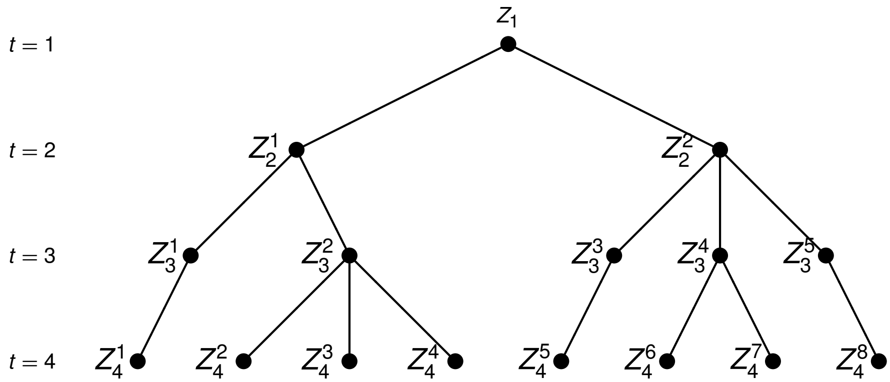

Collapsing Subtrees by Conditional Risk Measures

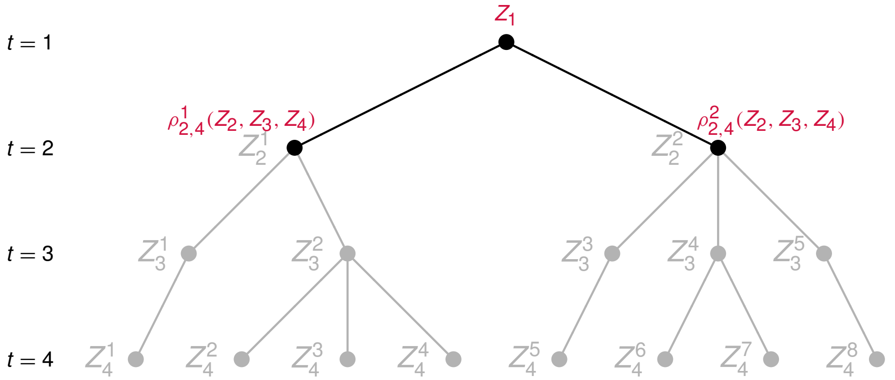

Recursive Structure of Dynamic Risk Measures

Define one-step conditional risk measures $\rho_t : \mathcal{Z}_{t+1} \to \mathcal{Z}_t$:

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

Nested Decomposition Theorem

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

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

Then for all t we have the representation

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

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

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

Here $s \in [1, p]$ 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

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

Here $s \in [1, p]$ 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

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

Here $s \in [1,p]$ 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

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

Here $s \in [1, p]$ 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

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

Here $s \in [1, p]$ 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

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

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

Multistage Risk-Averse Optimization Problems

Probability Space: $(\Omega, \mathcal{F}, P)$ with filtration $\mathcal{F}_1 \subset \cdots \subset \mathcal{F}_T \subset \mathcal{F}$
Decision Variables: $x_t(\omega)$, $\omega \in \Omega$, $t = 1, \dots, T$
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 Principle

$$\min_{x_1, x_2(\cdot), \dots, x_T(\cdot)} \Big\{ Z_1(x_1) + \rho_1\Big(Z_2(x_2) + \rho_2\Big(Z_3(x_3) + \dots$$

$$\cdots + \rho_{T-2}\big(Z_{T-1}(x_{T_1} + \rho_{T-1}(Z_T(x_T))\big) \cdots \Big)\Big)\Big\}$$

$$= \min_{x_1}\Big\{ Z_1(x_1) + \rho_1\Big[\min_{x_2}\Big(Z_2(x_2) + \rho_2\Big[\min_{x_3}\Big(Z_3(x_3) + \dots$$

$$\cdots + \rho_{T-2}\big[\min_{x_{T-1}}\big(Z_{T-1}(x_{T_1}) + \rho_{T-1}(\min_{x_T} Z_T(x_T))\big)\big] \cdots \Big)\Big]\Big)\Big]\Big\}$$

Interchangeability on a Scenario Tree

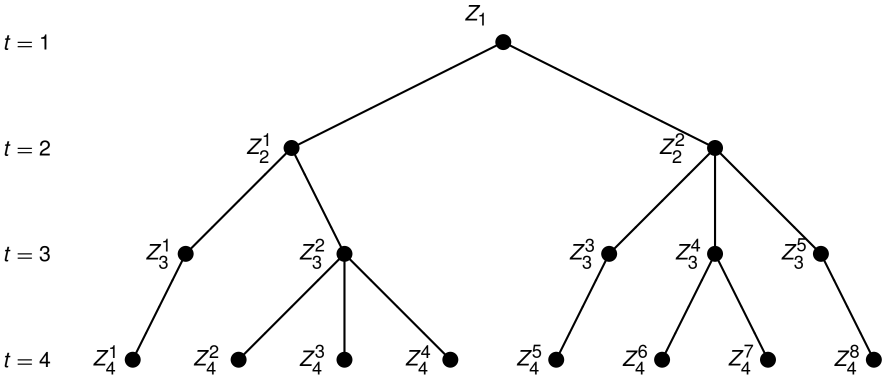

Interchangeability on a Scenario Tree

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

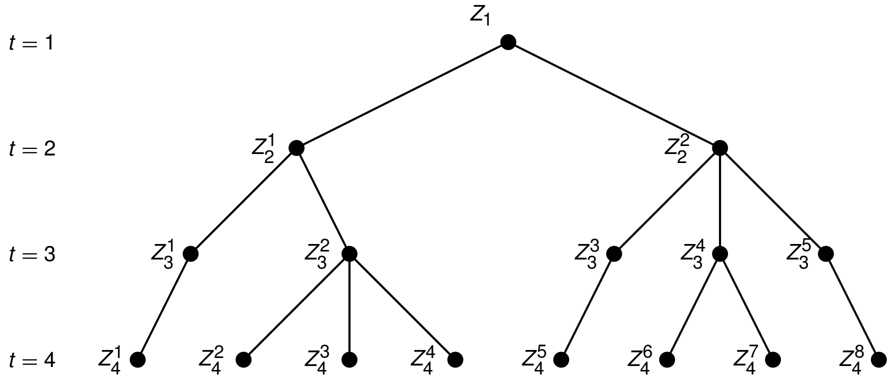

Interchangeability on a Scenario Tree

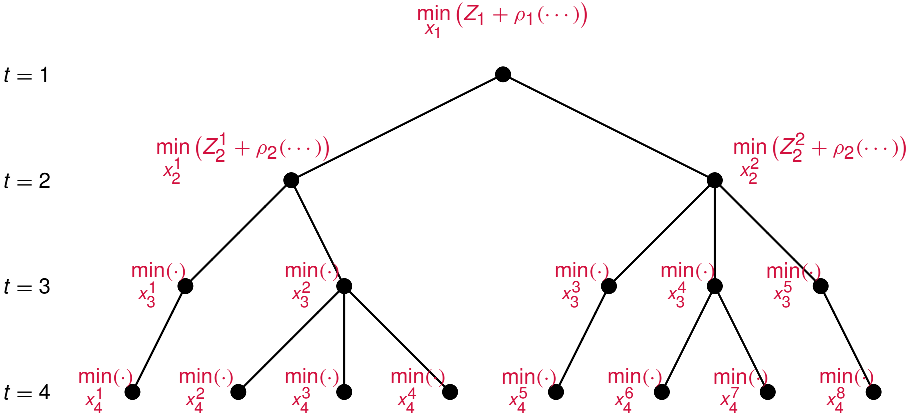

Controlled Markov Models

- State space $\mathcal{X}$ (Polish with Borel σ-algebra)
- Control space $\mathcal{U}$ (Polish with Borel σ-algebra)
- Feasible control sets $U_t : \mathcal{X} \rightrightarrows \mathcal{U}$, $t = 1, 2, \dots$
- Controlled transition kernels $Q_t : \text{graph}(U_t) \to \mathcal{P}$, $t = 1, 2, \dots$
  $\mathcal{P}$ - set of probability measures on $\mathcal{X}$
- Cost functions $c_t : \text{graph}(U_t) \to \mathbb{R}$, $t = 1, 2, \dots$
- State history $\mathcal{X}^t$ (up to time $t = 1, 2, \dots$)
- Policy $\pi_t : \mathcal{X}^t \to \mathcal{U}$, $t = 1, 2, \dots$ (always with values in $U_t(x_t)$)
- Markov policy $\pi_t : \mathcal{X} \to \mathcal{U}$, $t = 1, 2, \dots$
  (stationary if $\pi_t = \pi_1$ for all t)

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

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

Finite horizon expected cost problem:

$$\min_{\pi_1,\dots,\pi_T} \mathbb{E}\left[\sum_{t=1}^{T} c_t(x_t, u_t) + c_{T+1}(x_{T+1})\right]$$

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

Infinite horizon discounted expected cost problem:

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

- Both problems have optimal solutions in form of Markov policies
- Optimal policies can be found by dynamic programming equations

Our Intention

Introduce risk aversion to both problems by replacing the expected value by dynamic risk measures

Using Dynamic Risk Measures for Markov Decision Processes

- Controlled Markov process x_t, $t = 1, \dots, T, T+1$
- Policy $\Pi = \{\pi_1, \pi_2, \dots, \pi_T\}$ defines $u_t = \pi_t(x_t)$
- Cost sequence $c_t(x_t, u_t)$, $t = 1, \dots, T$, and $c_{T+1}(x_{T+1})$
- Dynamic time-consistent risk measure

$$J(\Pi) = c_1(x_1, u_1) + \rho_1\Big(c_2(x_2, u_2) + \rho_2\Big(c_3(x_3, u_3) + \\ \cdots + \rho_{T-1}\big(c_T(x_T, u_T) + \rho_T(c_{T+1}(x_{T+1}))\big)\cdots\Big)\Big)$$

- Risk-averse optimal control problem

$$\min_{\Pi} J(\Pi)$$

Difficulty

The value of $\rho_t(\cdot)$ is $\mathcal{F}_t$-measurable and is allowed to depend on the entire history of the process. We cannot expect a Markov optimal policy if our attitude to risk depends on the whole past

New Construction of a Conditional Risk Measure

- Consider functions of the state, in fixed space $\mathcal{V} = \mathcal{L}_p(\mathcal{X}, \mathcal{B}, P_0)$
- Additional argument: density on $(\mathcal{X}, \mathcal{B}, P_0))$ in the set

$$\mathcal{M} = \Big\{ m \in \mathcal{L}_q(\mathcal{X}, \mathcal{B}, P_0) : \int_{\mathcal{X}} m(x)\, P_0(dx) = 1,\ m \geq 0 \Big\}, \quad \frac{1}{p} + \frac{1}{q} = 1$$

Risk Transition Mapping Associated with a Kernel $Q : \text{graph}(U) \to \mathcal{M}$

A measurable functional $\sigma : \mathcal{V} \times \mathcal{X} \times \mathcal{M} \to \mathbb{R}$ satisfying

(i) For every $x \in \mathcal{X}$ the functional $v \mapsto \sigma\big(v, x, Q(x, u(x))\big)$ is a coherent measure of risk on $\mathcal{V}$

(ii) For every $v \in \mathcal{V}$ the function $x \mapsto \sigma\big(v, x, Q(x, u(x))\big)$ is in $\mathcal{V}$

Dual Representation

If $\sigma(\cdot, x, m)$ is lsc, then there exist convex sets $\mathcal{A}(x, m)$ such that

$$\sigma(v, x, m) = \sup_{\mu \in \mathcal{A}(x,m)} \langle v, \mu \rangle$$

Assumption: The controlled kernels Q_t have values in the set $\mathcal{M}$ (with densities with respect to P_0)

A one-step conditional risk measure $\rho_t : \mathcal{Z}_{t+1} \to \mathcal{Z}_t$ is a Markov risk measure with respect to the controlled Markov process $\{x_t\}$, if there exists a risk transition mapping $\sigma_t : \mathcal{V} \times \mathcal{X} \times \mathcal{M} \to \mathbb{R}$ such that for all $v \in \mathcal{V}$ and for all measurable $u_t \in U_t(x_t)$ we have

$$\rho_t(v(x_{t+1})) = \sigma_t\big(v, x_t, Q_t(x_t, u_t)\big)$$

Duality: $$\rho_t(v(x_{t+1})) = \sup_{\mu \in \mathcal{A}_t(x_t, Q_t(x_t, u_t))} \langle v, \mu \rangle$$

$\mathcal{A}_t(x_t, Q_t(x_t, u_t))$ – controlled multikernel

In the risk neutral setting, when $\rho_t(v(x_{t+1})) = \mathbb{E}[v(x_{t+1}) | \mathcal{F}_t]$ we have a single-valued controlled kernel $\mathcal{A}_t(x_t, Q_t(x_t, u_t)) = \{Q_t(x_t, u_t)\}$.
Risk-averse preferences ⇆ Ambiguity in the transition kernel

Markov Risk Evaluation

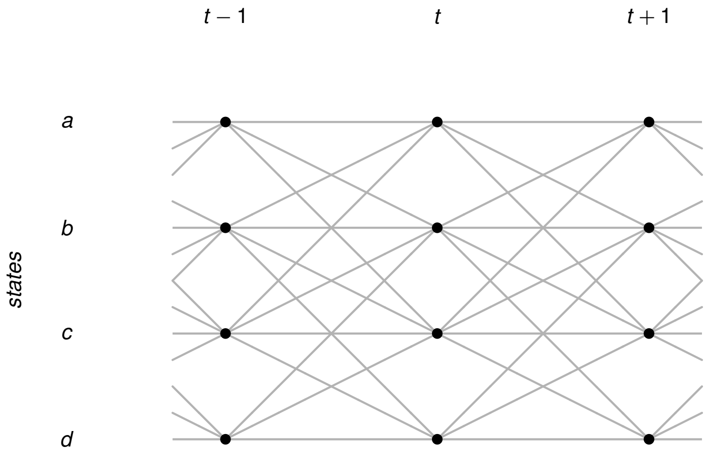

Markov Risk Evaluation

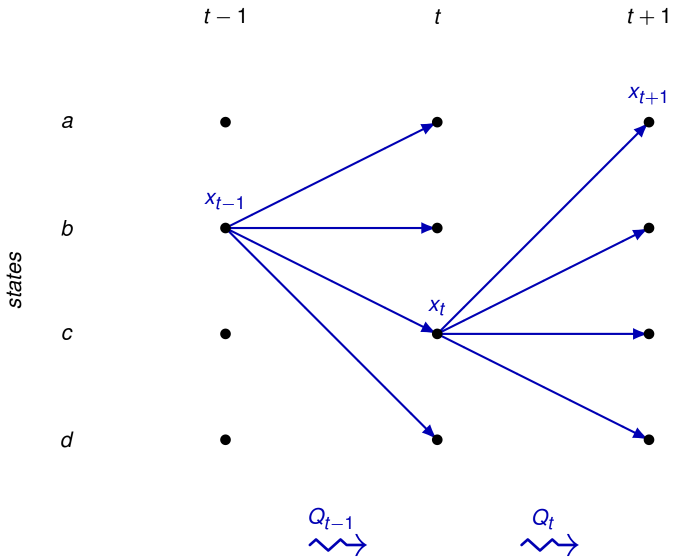

Markov Risk Evaluation

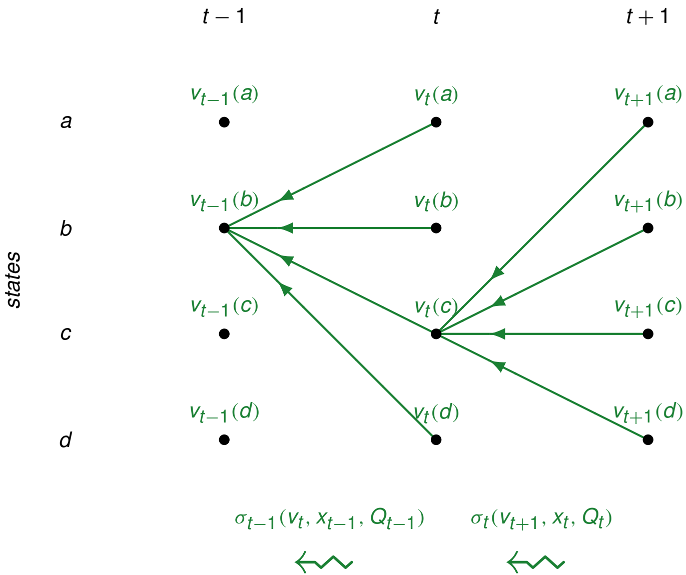

Finite Horizon Risk-Averse Control Problem

Consider a controlled Markov process $\{x_t\}$ with $u_t = \pi_t(x_1, \ldots, x_t)$.

Risk-averse optimal control problem:

$$\min_{\Pi} \; c_1(x_1, u_1) + \rho_1\Big(c_2(x_2, u_2) + \rho_2\Big(c_3(x_3, u_3) + \cdots$$
$$\cdots + \rho_{T-1}\big(c_T(x_T, u_T) + \rho_T(c_{T+1}(x_{T+1}))\big)\cdots\Big)\Big)$$

Theorem

If the conditional measures ρ_t are Markov (+ technical conditions), then the optimal solution is given by the dynamic programming equations:

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

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:

$$\min_{\Pi}\ c_1(x_1, u_1) + \rho_1\Big(c_2(x_2, u_2) + \rho_2\Big(c_3(x_3, u_3) + \cdots$$
$$\cdots + \rho_{T-1}\big(c_T(x_T, u_T) + \rho_T(c_{T+1}(x_{T+1}))\big) \cdots\Big)\Big)$$

Theorem

If the conditional measures ρ_t are Markov (+ technical conditions), then the optimal solution is given by the dynamic programming equations:

$$v_{T+1}(x) = c_{T+1}(x), \quad x \in \mathcal{X}$$
$$v_t(x) = \min_{u \in U_t(x)} \Big\{ c_t(x, u) + \sup_{\mu \in \mathcal{A}_t(x, Q_t(x,u))} \langle v_{t+1}, \mu \rangle \Big\}, \quad t = T, \dots, 1$$

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

Discounted Risk Measures for Infinite Sequences

- $\{\mathcal{F}_t\}$ - filtration on $(\Omega, \mathcal{F})$
- Z_t, $t = 1, 2 \ldots$ - adapted sequence of random variables
- $\mathcal{Z}_t = \mathcal{L}_p(\Omega, \mathcal{F}_t, P)$, $\mathcal{Z} = \mathcal{Z}_1 \times \mathcal{Z}_2 \times \cdots$
- $\rho_t : \mathcal{Z}_{t+1} \to \mathcal{Z}_t$ - conditional risk mappings

Fix the discount factor $\alpha \in (0, 1)$. For $T = 1, 2, \ldots$ define

$$\rho^{\alpha}_{1,T}(Z_1, Z_2, \ldots, Z_T) = \rho_{1,T}(Z_1, \alpha Z_2, \ldots, \alpha^{T-1} Z_T)$$
$$= Z_1 + \rho_1\Big(\alpha Z_2 + \rho_2\big(\alpha^2 Z_3 + \cdots + \rho_{T-1}(\alpha^{T-1} Z_T) \cdots\big)\Big)$$

Discounted Risk Measure

$$\varrho^{\alpha}(Z) = \lim_{T\to\infty} \rho^{\alpha}_{1,T}(Z_1, Z_2, \ldots, Z_T)$$

It is well defined, convex, monotone, and positively homogeneous, whenever $\max_t \operatorname{ess\,sup} |Z_t(\omega)| < \infty$

Discounted Infinite Horizon Problem

We consider a controlled stationary Markov process $\{x_t\}$, $t = 1, 2, \ldots$ with a discounted measure of risk ($0 < \alpha < 1$):

$$\min_{\Pi} \; J(\Pi, x_1) = \varrho^{\alpha}\big(c(x_1, u_1), c(x_2, u_2), \cdots\big)$$
$$= c(x_1, u_1) + \rho_1\Big(\alpha c(x_2, u_2) + \rho_2\big(\alpha^2 c(x_3, u_3) + \cdots\big)\Big)$$

Conditional Markov risk measures ρ_t are stationary, if they share the same risk transition mapping $\sigma : \mathcal{X} \times \mathcal{V} \times \mathcal{M} \to \mathbb{R}$

Theorem

If the conditional measures ρ_t are Markov and stationary, then the optimal value function $\hat{v}(x)$ satisfies the dynamic programming equation:

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

Optimal stationary Markov policy $\hat{\Pi} = \{\hat{\pi}, \hat{\pi}, \ldots\}$ - the minimizer above

Discounted Infinite Horizon Problem

We consider a controlled stationary Markov process $\{x_t\}$, $t = 1, 2, \dots$ with a discounted measure of risk ($0 < \alpha < 1$):

$$\min_{\Pi} J(\Pi, x_1) = \varrho^{\alpha}\big(c(x_1, u_1), c(x_2, u_2), \cdots\big)$$

$$= c(x_1, u_1) + \rho_1\Big(\alpha c(x_2, u_2) + \rho_2\big(\alpha^2 c(x_3, u_3) + \cdots\big)\Big)$$

Conditional Markov risk measures ρ_t are stationary, if they share the same risk transition mapping $\sigma : \mathcal{X} \times \mathcal{V} \times \mathcal{M} \to \mathbb{R}$

Theorem

If the conditional measures ρ_t are Markov and stationary, then the optimal value function $\hat{v}(x)$ satisfies the dynamic programming equation:

$$v(x) = \min_{u \in U(x)} \Big\{c(x, u) + \alpha \sup_{\mu \in \mathcal{A}(x, Q(x,u))} \langle v, \mu \rangle\Big\}, \quad x \in \mathcal{X}$$

Optimal stationary Markov policy $\hat{\Pi} = \{\hat{\pi}, \hat{\pi}, \dots\}$ - the minimizer above

Solution Methods

Value iteration

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

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 (by a specialized Newton method) the nonsmooth equation

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

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

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

Example: Service and Insurance

5 states - “new”, “good”, “bad”, “broken uninsured”, “broken insured”
5 controls - “nothing”, “insure”, “service”, “service and insure”, “replace”

As the cost of insurance is greater than the expected benefit from it,
no risk-neutral model can suggest insurance

Results for the mean-semideviation model
Value Iteration - 212 steps
Policy Iteration - 2 steps (with 3 and 2 Newton steps)

Optimal Policy

“new” - “service and insure”
“good” - “service”
“bad” - “do nothing”
“broken uninsured” and“broken insured” - “replace”