Optimization under uncertainty and risk aversion

V. Leclère (ENPC)

June 15, 2022

Contents

Contents

An optimization problem

A generic optimization problem can be written

$$\begin{aligned} \min_{x} \quad & L(x) \\ s.t. \quad & g(x) \leq 0 \end{aligned}$$

where

- x is the decision variable
- L is the objective function
- g is the constraint function

An optimization problem with uncertainty

Adding uncertainty ξ in the mix

$$\begin{aligned} \min_x \quad & L(x, \tilde{\xi}) \\ s.t. \quad & g(x, \tilde{\xi}) \leq 0 \end{aligned}$$

Remarks:

- $\tilde{\xi}$ is unknown. Two main way of modelling it:
  - $\tilde{\xi} \in R$ with a known uncertainty set R, and a pessimistic approach. This is the robust optimization approach (RO).
  - $\tilde{\xi}$ is a random variable with known probability law. This is the Stochastic Programming approach (SP).
- Cost is not well defined.
  - RO : $\max_{\xi \in R} L(x, \xi)$.
  - SP : $\mathbb{E}\big[L(x, \boldsymbol{\xi})\big]$.
- Constraints are not well defined.
  - RO : $g(x, \xi) \leq 0, \qquad \forall \xi \in R$.
  - SP : $g(x, \boldsymbol{\xi}) \leq 0, \qquad \mathbb{P} - a.s.$.

An optimization problem with uncertainty

Adding uncertainty ξ in the mix

$$\begin{aligned} \min_x \quad & L(x, \tilde{\xi}) \\ s.t. \quad & g(x, \tilde{\xi}) \leq 0 \end{aligned}$$

Remarks:

- $\tilde{\xi}$ is unknown. Two main way of modelling it:
  - $\tilde{\xi} \in R$ with a known uncertainty set R, and a pessimistic approach. This is the robust optimization approach (RO).
  - $\tilde{\xi}$ is a random variable with known probability law. This is the Stochastic Programming approach (SP).
- Cost is not well defined.
  - RO : $\max_{\xi \in R} L(x, \xi)$.
  - SP : $\mathbb{E}\big[L(x, \boldsymbol{\xi})\big]$.
- Constraints are not well defined.
  - RO : $g(x, \xi) \leq 0, \qquad \forall \xi \in R$.
  - SP : $g(x, \boldsymbol{\xi}) \leq 0, \qquad \mathbb{P} - a.s.$.

An optimization problem with uncertainty

Adding uncertainty ξ in the mix

$$\begin{aligned} \min_{x} \quad & L(x, \tilde{\xi}) \\ s.t. \quad & g(x, \tilde{\xi}) \leq 0 \end{aligned}$$

Remarks:

- $\tilde{\xi}$ is unknown. Two main way of modelling it:
    - $\tilde{\xi} \in R$ with a known uncertainty set R, and a pessimistic approach. This is the robust optimization approach (RO).
    - $\tilde{\xi}$ is a random variable with known probability law. This is the Stochastic Programming approach (SP).
- Cost is not well defined.
    - RO : $\max_{\xi \in R} L(x, \xi)$.
    - SP : $\mathbb{E}\big[L(x, \boldsymbol{\xi})\big]$.
- Constraints are not well defined.
    - RO : $g(x, \xi) \leq 0, \qquad \forall \xi \in R$.
    - SP : $g(x, \boldsymbol{\xi}) \leq 0, \qquad \mathbb{P} - a.s.$.

Requirements and limits

- Stochastic optimization :
  - requires a law of the uncertainty ξ
  - can be hard to solve (generally require discretizing the support and blowing up the dimension of the problem)
  - there exists specific methods (like Bender's decomposition)
- Robust optimization :
  - requires an uncertainty set R
  - can be overly conservative, even for reasonable R
  - complexity strongly depend on the choice of R
- Distributionally robust optimization :
  - is a mix between robust and stochastic optimization
  - consists in solving a stochastic optimization problem where the law is chosen in a robust way
  - is a fast growing fields with multiple recent results
  - but is still hard to implement than other approaches

Requirements and limits

- Stochastic optimization :
  - requires a law of the uncertainty ξ
  - can be hard to solve (generally require discretizing the support and blowing up the dimension of the problem)
  - there exists specific methods (like Bender's decomposition)
- Robust optimization :
  - requires an uncertainty set R
  - can be overly conservative, even for reasonable R
  - complexity strongly depend on the choice of R
- Distributionally robust optimization :
  - is a mix between robust and stochastic optimization
  - consists in solving a stochastic optimization problem where the law is chosen in a robust way
  - is a fast growing fields with multiple recent results
  - but is still hard to implement than other approaches

Requirements and limits

- Stochastic optimization :
  - requires a law of the uncertainty ξ
  - can be hard to solve (generally require discretizing the support and blowing up the dimension of the problem)
  - there exists specific methods (like Bender's decomposition)
- Robust optimization :
  - requires an uncertainty set R
  - can be overly conservative, even for reasonable R
  - complexity strongly depend on the choice of R
- Distributionally robust optimization :
  - is a mix between robust and stochastic optimization
  - consists in solving a stochastic optimization problem where the law is chosen in a robust way
  - is a fast growing fields with multiple recent results
  - but is still hard to implement than other approaches

Contents

Contents

Definition

A risk measure ρ is a function associating a *determinist equivalent* to an *uncertain cost* $\boldsymbol{X}$, in the sense that it is the maximum amount of cash you are willing to pay to be rid of the uncertain cost.

Mathematically, consider a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Then the risk measure ρ is a function mapping the random variables $L^0(\Omega, \mathcal{F}, \mathbb{P}; \mathbb{R} \cup \{+\infty\})$ into $\mathbb{R} \cup \{+\infty\}$.

Warning : the definition and convention in risk measure litterature are not perfectly unified. For example, what I call $\rho(\boldsymbol{X})$ is sometimes called $-\rho(X)$ or $-\rho(-X)$... Or some assitional assumption (discussed later) are required of a risk measure by some author.

Definition

A risk measure ρ is a function associating a *determinist equivalent* to an *uncertain cost* $\boldsymbol{X}$, in the sense that it is the maximum amount of cash you are willing to pay to be rid of the uncertain cost.

Mathematically, consider a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Then the risk measure ρ is a function mapping the random variables $L^0(\Omega, \mathcal{F}, \mathbb{P}; \mathbb{R} \cup \{+\infty\})$ into $\mathbb{R} \cup \{+\infty\}$.

Warning : the definition and convention in risk measure litterature are not perfectly unified. For example, what I call $\rho(\boldsymbol{X})$ is sometimes called $-\rho(X)$ or $-\rho(-X)$... Or some assitional assumption (discussed later) are required of a risk measure by some author.

Definition

A risk measure ρ is a function associating a *determinist equivalent* to an *uncertain cost* $\boldsymbol{X}$, in the sense that it is the maximum amount of cash you are willing to pay to be rid of the uncertain cost.

Mathematically, consider a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Then the risk measure ρ is a function mapping the random variables $L^0(\Omega, \mathcal{F}, \mathbb{P}; \mathbb{R} \cup \{+\infty\})$ into $\mathbb{R} \cup \{+\infty\}$.

Warning : the definition and convention in risk measure litterature are not perfectly unified. For example, what I call $\rho(\boldsymbol{X})$ is sometimes called $-\rho(X)$ or $-\rho(-X)$... Or some assitional assumption (discussed later) are required of a risk measure by some author.

Interpretation

Assume that you consider an uncertain cost $\boldsymbol{X}$.

- $\rho(\boldsymbol{X})$ can be seen as the maximum price you are ready to pay for an insurance covering this cost.
- The choice of ρ is difficult as it is highly subjective.
- We will discuss some "natural" properties one can ask of ρ and suggest some possible choice of ρ.

Risk neutral case: expectation

We started our discussion by assuming that we have a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. In particular we assume knowledge of a reference probability $\mathbb{P}$.

We say that we are in the risk-neutral setting when the chosen risk measure is simply the expectation with respect to $\mathbb{P}$:

$$\rho[\cdot] = \mathbb{E}_{\mathbb{P}}[\cdot]$$

This choice is justified if you do not have any aversion to risk (e.g : you are willing to pay 100 € for having 1/10 chance of getting 1000 €)

It can also be justified if you are repeating the same operation a large number of times by law of large number.

Risk neutral case: expectation

We started our discussion by assuming that we have a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. In particular we assume knowledge of a reference probability $\mathbb{P}$.

We say that we are in the risk-neutral setting when the chosen risk measure is simply the expectation with respect to $\mathbb{P}$:

$$\rho[\cdot] = \mathbb{E}_{\mathbb{P}}[\cdot]$$

This choice is justified if you do not have any aversion to risk (e.g : you are willing to pay 100 € for having 1/10 chance of getting 1000 €)

It can also be justified if you are repeating the same operation a large number of times by law of large number.

Some interesting properties

We are now giving some properties that have intuitive interest.

- Monotonicity: $\boldsymbol{X} \leq \boldsymbol{Y} \quad \Longrightarrow \quad \rho(\boldsymbol{X}) \leq \rho(\boldsymbol{Y})$
- Subadditivity: $\rho[\boldsymbol{X}+\boldsymbol{Y}] \leq \rho[\boldsymbol{X}]+\rho[\boldsymbol{Y}]$
- Translation equivariance: $\rho[\boldsymbol{X}+c] = \rho[\boldsymbol{X}]+c$
- Normalization: $\rho[0]=0$
- Positive homogeneity: $\rho[\lambda\boldsymbol{X}] = \lambda\rho[\boldsymbol{X}]$ for all $\lambda > 0$
- Convexity: $\rho[\lambda\boldsymbol{X}+(1-\lambda)\boldsymbol{Y}] \leq \lambda\rho[\boldsymbol{X}]+(1-\lambda)\rho[\boldsymbol{Y}]$
- Law invariance: if $\boldsymbol{X}$ and $\boldsymbol{Y}$ have the same law, then $\rho[\boldsymbol{X}]=\rho[\boldsymbol{Y}]$.

Remark : With positive homogenity, subbaditivity and convexity are equivalent.

Some interesting properties

We are now giving some properties that have intuitive interest.

- Monotonicity: $\boldsymbol{X} \leq \boldsymbol{Y} \implies \rho(\boldsymbol{X}) \leq \rho(\boldsymbol{Y})$
- Subadditivity: $\rho[\boldsymbol{X} + \boldsymbol{Y}] \leq \rho[\boldsymbol{X}] + \rho[\boldsymbol{Y}]$
- Translation equivariance: $\rho[\boldsymbol{X} + c] = \rho[\boldsymbol{X}] + c$
- Normalization: $\rho[0] = 0$
- Positive homogeneity: $\rho[\lambda \boldsymbol{X}] = \lambda \rho[\boldsymbol{X}]$ for all $\lambda > 0$
- Convexity: $\rho[\lambda \boldsymbol{X} + (1-\lambda)\boldsymbol{Y}] \leq \lambda \rho[\boldsymbol{X}] + (1-\lambda)\rho[\boldsymbol{Y}]$
- Law invariance: if $\boldsymbol{X}$ and $\boldsymbol{Y}$ have the same law, then $\rho[\boldsymbol{X}] = \rho[\boldsymbol{Y}]$.

Remark : With positive homogenity, subbaditivity and convexity are equivalent.

Coherent risk measures

- A risk measure ρ is said to be convex if it is monotone, convex and translation equivariant.
- A risk measure ρ is said to be coherent if it is monotone, subbaditive, translation equivariant and positive homogeneous.
- A convex risk measure is coherent if it is positive homogeneous.

Mathematicians mostly agree that the right modeling tool are law invariant coherent risk measures, in particular because if ρ is a coherent law invariant risk measure, then there exists a set of probability $\mathcal{Q}$ such that

$$\forall \boldsymbol{X}, \qquad \rho[\boldsymbol{X}] = \sup_{\mathbb{Q}\in\mathcal{Q}} \mathbb{E}_{\mathbb{Q}}[\boldsymbol{X}]$$

And reciprocally, every risk measure defined in such a way is coherent and law invariant.

Coherent risk measures

- A risk measure ρ is said to be convex if it is monotone, convex and translation equivariant.
- A risk measure ρ is said to be coherent if it is monotone, subbaditive, translation equivariant and positive homogeneous.
- A convex risk measure is coherent if it is positive homogeneous.

Mathematicians mostly agree that the right modeling tool are law invariant coherent risk measures, in particular because if ρ is a coherent law invariant risk measure, then there exists a set of probability $\mathcal{Q}$ such that

$$\forall \boldsymbol{X}, \qquad \rho[\boldsymbol{X}] = \sup_{\mathbb{Q} \in \mathcal{Q}} \mathbb{E}_{\mathbb{Q}}[\boldsymbol{X}]$$

And reciprocally, every risk measure defined in such a way is coherent and law invariant.

Contents

Extreme cases

We start with two example of coherent risk measures that are limit cases of some families of risk measures.

- The risk neutral case is represented by

$$\rho[\boldsymbol{X}] = \mathbb{E}[\boldsymbol{X}].$$

- The robust case is represented by

$$\rho[\boldsymbol{X}] = \sup_{\omega} \boldsymbol{X}(\omega).$$

Polyhedral risk measure

A very practical and simple way of constructing a coherent risk measure, known as *polyhedral risk measures*, consists in considering a finite set of probability $\mathcal{Q} = \{\mathbb{Q}^k\}_{k\in[K]}$, for example each given by an expert, and define

$$\rho[\boldsymbol{X}] = \max_{k\in[K]} \ \mathbb{E}_{\mathbb{Q}^k}[\boldsymbol{X}].$$

Let's take an example with a coin flip : $\Omega = \{H, T\}$. We have two expert, one thinking that it is equilibrated (i.e $\mathbb{Q}^1 = (0.5, 0.5)$), and the other thinking that T will happen with probability 0.7 (i.e $\mathbb{Q}^2 = (0.3, 0.7)$)). Therefore we would have

$$\rho[\boldsymbol{X}] = \max\left\{0.5\boldsymbol{X}(H) + 0.5\boldsymbol{X}(T), 0.3\boldsymbol{X}(H) + 0.7\boldsymbol{X}(T)\right\}$$

Polyhedral risk measure

A very practical and simple way of constructing a coherent risk measure, known as *polyhedral risk measures*, consists in considering a finite set of probability $\mathcal{Q} = \{\mathbb{Q}^k\}_{k\in[K]}$, for example each given by an expert, and define

$$\rho[\boldsymbol{X}] = \max_{k\in[K]} \mathbb{E}_{\mathbb{Q}^k}[\boldsymbol{X}].$$

Let's take an example with a coin flip : $\Omega = \{H, T\}$. We have two expert, one thinking that it is equilibrated (i.e $\mathbb{Q}^1 = (0.5, 0.5)$), and the other thinking that T will happen with probability 0.7 (i.e $\mathbb{Q}^2 = (0.3, 0.7)$)). Therefore we would have

$$\rho[\boldsymbol{X}] = \max\Big\{0.5\boldsymbol{X}(H) + 0.5\boldsymbol{X}(T), 0.3\boldsymbol{X}(H) + 0.7\boldsymbol{X}(T)\Big\}$$

Mean-variance

- A very natural (but misleading) way of modeling risk aversion consists in considering the Markovitz expectation/variance trade-off, i.e.

  $$\rho(\boldsymbol{X}) = \mathbb{E}[\boldsymbol{X}] + \alpha var(\boldsymbol{X}), \qquad \alpha \geq 0$$

  Not satisfying as $\mathbb{E}[\boldsymbol{X}]$ is in €, and $var(\boldsymbol{X})$ is in €^{2}.
- We can adjust by considering the standard deviation

  $$\rho(\boldsymbol{X}) = \mathbb{E}[\boldsymbol{X}] + \alpha\sigma(\boldsymbol{X}), \qquad \alpha \geq 0$$

This is a wrong approach as this is non-monotonous. Compare $\boldsymbol{X}$ and $\boldsymbol{Y}$ given as follows: $\mathbb{P}(\boldsymbol{X} = 11) = 1$, $\mathbb{P}(\boldsymbol{Y} = 10) = 0.9$ and $\mathbb{P}(\boldsymbol{Y} = 0) = 0.1$.

Mean-variance

- A very natural (but misleading) way of modeling risk aversion consists in considering the Markovitz expectation/variance trade-off, i.e.

$$\rho(\boldsymbol{X}) = \mathbb{E}[\boldsymbol{X}] + \alpha var(\boldsymbol{X}), \qquad \alpha \geq 0$$

  Not satisfying as $\mathbb{E}[\boldsymbol{X}]$ is in €, and $var(\boldsymbol{X})$ is in €^{2}.
- We can adjust by considering the standard deviation

$$\rho(\boldsymbol{X}) = \mathbb{E}[\boldsymbol{X}] + \alpha\sigma(\boldsymbol{X}), \qquad \alpha \geq 0$$

This is a wrong approach as this is non-monotonous. Compare $\boldsymbol{X}$ and $\boldsymbol{Y}$ given as follows: $\mathbb{P}(\boldsymbol{X} = 11) = 1$, $\mathbb{P}(\boldsymbol{Y} = 10) = 0.9$ and $\mathbb{P}(\boldsymbol{Y} = 0) = 0.1$.

Mean-variance

- A very natural (but misleading) way of modeling risk aversion consists in considering the Markovitz expectation/variance trade-off, i.e.
$$\rho(\boldsymbol{X}) = \mathbb{E}[\boldsymbol{X}] + \alpha var(\boldsymbol{X}), \qquad \alpha \geq 0$$
Not satisfying as $\mathbb{E}[\boldsymbol{X}]$ is in €, and $var(\boldsymbol{X})$ is in €^{2}.
- We can adjust by considering the standard deviation
$$\rho(\boldsymbol{X}) = \mathbb{E}[\boldsymbol{X}] + \alpha\sigma(\boldsymbol{X}), \qquad \alpha \geq 0$$

This is a wrong approach as this is non-monotonous.
Compare $\boldsymbol{X}$ and $\boldsymbol{Y}$ given as follows: $\mathbb{P}(\boldsymbol{X} = 11) = 1$,
$\mathbb{P}(\boldsymbol{Y} = 10) = 0.9$ and $\mathbb{P}(\boldsymbol{Y} = 0) = 0.1$.

Mean-variance

- A very natural (but misleading) way of modeling risk aversion consists in considering the Markovitz expectation/variance trade-off, i.e.

$$\rho(\boldsymbol{X}) = \mathbb{E}[\boldsymbol{X}] + \alpha var(\boldsymbol{X}), \qquad \alpha \geq 0$$

  Not satisfying as $\mathbb{E}[\boldsymbol{X}]$ is in €, and $var(\boldsymbol{X})$ is in €^{2}.
- We can adjust by considering the standard deviation

$$\rho(\boldsymbol{X}) = \mathbb{E}[\boldsymbol{X}] + \alpha\sigma(\boldsymbol{X}), \qquad \alpha \geq 0$$

This is a wrong approach as this is non-monotonous.
Compare $\boldsymbol{X}$ and $\boldsymbol{Y}$ given as follows: $\mathbb{P}(\boldsymbol{X} = 11) = 1$, $\mathbb{P}(\boldsymbol{Y} = 10) = 0.9$ and $\mathbb{P}(\boldsymbol{Y} = 0) = 0.1$.
We have

$$\rho[\boldsymbol{X}] = 10, \qquad \rho[\boldsymbol{Y}] \approx 9 + 0.42\alpha$$

Semi deviation model

- We can adapt Markovitz's approach to make it coherent by considering semi-deviation instead of standard deviation.
- The semi-deviation of $\boldsymbol{X}$ is defined as

$$\sigma^+(\boldsymbol{X}) = \sqrt{\mathbb{E}[(X - \mathbb{E}[X])_+^2}$$

- We can consider

$$\rho[\boldsymbol{X}] = \mathbb{E}[X] + \alpha\sigma^+(\boldsymbol{X}),$$

which is coherent for $\alpha \geq 0$.

Value at Risk I

- A very common risk measure is the Value at Risk of level β.
- It is the highest value that can take the cost if you forgo the $1-\beta$ worst case.
- $V@R_\beta(\boldsymbol{X}) := \inf_t \mathbb{P}(\boldsymbol{X} \geq t) \leq 1-\beta$.

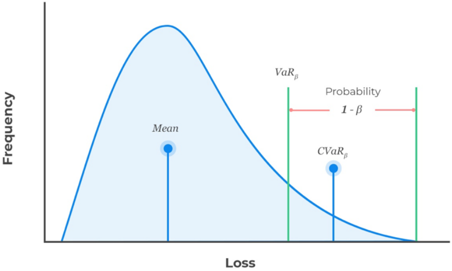

Value at Risk II

The Value at Risk is a pretty intuitive risk measure, with some heavy drawbacks:

- it can be hard to compute, and even more to use as a constraint (chance-constraints).
- it is positively homogene and translation equivariant
- but not coherent because it is not subadditive:
  - Consider X and Y two independznt random variable taking value 1 with probability 0.1, and 0 otherwise.

Value at Risk II

The Value at Risk is a pretty intuitive risk measure, with some heavy drawbacks:

- it can be hard to compute, and even more to use as a constraint (chance-constraints).
- it is positively homogene and translation equivariant
- but not coherent because it is not subadditive:
    - Consider X and Y two independznt random variable taking value 1 with probability 0.1, and 0 otherwise.

Value at Risk II

The Value at Risk is a pretty intuitive risk measure, with some heavy drawbacks:

- it can be hard to compute, and even more to use as a constraint (chance-constraints).
- it is positively homogene and translation equivariant
- but not coherent because it is not subadditive:
  - Consider X and Y two independznt random variable taking value 1 with probability 0.1, and 0 otherwise.
  - Then $X+Y$ take value 0 with probability 0.81, 1 with probability 0.18 and 2 with probability 0.01.

Value at Risk II

The Value at Risk is a pretty intuitive risk measure, with some heavy drawbacks:

- it can be hard to compute, and even more to use as a constraint (chance-constraints).
- it is positively homogene and translation equivariant
- but not coherent because it is not subadditive:
  - Consider X and Y two independznt random variable taking value 1 with probability 0.1, and 0 otherwise.
  - Then $X+Y$ take value 0 with probability 0.81, 1 with probability 0.18 and 2 with probability 0.01.
  - Thus $V@R_{0.85}(X+Y)=1$ while $V@R_{0.85}(X)=V@R_{0.85}(Y)=0$.

Tail Value at Risk

The Tail Value at Risk (TV@R) (a.k.a conditional value at risk, average value at risk, expected shortfall or superquantile) is a convexification of the V@R measure. It has various equivalent definitions (assuming $\boldsymbol{X}$ admits a density):

1. $TV@R_\beta[\boldsymbol{X}] = \mathbb{E}[\boldsymbol{X} | \boldsymbol{X} \geq V@R_\beta(\boldsymbol{X})]$
2. $TV@R_\beta[\boldsymbol{X}] = \frac{1}{1-\beta} \int_\beta^1 V@R_b(\boldsymbol{X}) db$
3. $TV@R_\beta[\boldsymbol{X}] = \min_{t \in \mathbb{R}} \left\{ t + \frac{1}{1-\beta} \mathbb{E}[\boldsymbol{X} - t]^+ \right\}$
4. $TV@R_\beta[\boldsymbol{X}] = \sup \left\{ \mathbb{E}_{\mathbb{Q}}[\boldsymbol{X}] \mid \frac{d\mathbb{Q}}{d\mathbb{P}} \leq \frac{1}{1-\beta} \right\}$

Tail Value at Risk

The Tail Value at Risk (TV@R) (a.k.a conditional value at risk, average value at risk, expected shortfall or superquantile) is a convexification of the V@R measure. It has various equivalent definitions (assuming $\boldsymbol{X}$ admits a density):

1. $TV@R_\beta[\boldsymbol{X}] = \mathbb{E}[\boldsymbol{X} | \boldsymbol{X} \geq V@R_\beta(\boldsymbol{X})]$
   $\rightsquigarrow$ intuitive interpretation
2. $TV@R_\beta[\boldsymbol{X}] = \frac{1}{1-\beta} \int_\beta^1 V@R_b(\boldsymbol{X}) db$
3. $TV@R_\beta[\boldsymbol{X}] = \min_{t \in \mathbb{R}} \left\{ t + \frac{1}{1-\beta} \mathbb{E}[\boldsymbol{X} - t]^+ \right\}$
4. $TV@R_\beta[\boldsymbol{X}] = \sup \left\{ \mathbb{E}_{\mathbb{Q}}[\boldsymbol{X}] \mid \frac{d\mathbb{Q}}{d\mathbb{P}} \leq \frac{1}{1-\beta} \right\}$

Tail Value at Risk

The Tail Value at Risk (TV@R) (a.k.a conditional value at risk, average value at risk, expected shortfall or superquantile) is a convexification of the V@R measure. It has various equivalent definitions (assuming $\boldsymbol{X}$ admits a density):

1. $TV@R_\beta[\boldsymbol{X}] = \mathbb{E}[\boldsymbol{X} | \boldsymbol{X} \geq V@R_\beta(\boldsymbol{X})]$
   $\leadsto$ intuitive interpretation
2. $TV@R_\beta[\boldsymbol{X}] = \frac{1}{1-\beta} \int_\beta^1 V@R_b(\boldsymbol{X}) db$
   $\leadsto$ upperbound of V@R
3. $TV@R_\beta[\boldsymbol{X}] = \min_{t \in \mathbb{R}} \left\{ t + \frac{1}{1-\beta} \mathbb{E}[\boldsymbol{X} - t]^+ \right\}$
4. $TV@R_\beta[\boldsymbol{X}] = \sup \left\{ \mathbb{E}_{\mathbb{Q}}[\boldsymbol{X}] \mid \frac{d\mathbb{Q}}{d\mathbb{P}} \leq \frac{1}{1-\beta} \right\}$

Tail Value at Risk

The Tail Value at Risk (TV@R) (a.k.a conditional value at risk, average value at risk, expected shortfall or superquantile) is a convexification of the V@R measure. It has various equivalent definitions (assuming $\boldsymbol{X}$ admits a density):

1. $TV@R_\beta[\boldsymbol{X}] = \mathbb{E}[\boldsymbol{X} | \boldsymbol{X} \geq V@R_\beta(\boldsymbol{X})]$
   $\leadsto$ intuitive interpretation
2. $TV@R_\beta[\boldsymbol{X}] = \frac{1}{1-\beta} \int_\beta^1 V@R_b(\boldsymbol{X}) db$
   $\leadsto$ upperbound of V@R
3. $TV@R_\beta[\boldsymbol{X}] = \min_{t \in \mathbb{R}} \left\{ t + \frac{1}{1-\beta} \mathbb{E}[\boldsymbol{X} - t]^+ \right\}$
   $\leadsto$ linear formulation
4. $TV@R_\beta[\boldsymbol{X}] = \sup \left\{ \mathbb{E}_{\mathbb{Q}}[\boldsymbol{X}] \mid \frac{d\mathbb{Q}}{d\mathbb{P}} \leq \frac{1}{1-\beta} \right\}$

Tail Value at Risk

The Tail Value at Risk (TV@R) (a.k.a conditional value at risk, average value at risk, expected shortfall or superquantile) is a convexification of the V@R measure. It has various equivalent definitions (assuming $\boldsymbol{X}$ admits a density):

1. $TV@R_\beta[\boldsymbol{X}] = \mathbb{E}[\boldsymbol{X} | \boldsymbol{X} \geq V@R_\beta(\boldsymbol{X})]$
   $\rightsquigarrow$ intuitive interpretation
2. $TV@R_\beta[\boldsymbol{X}] = \frac{1}{1-\beta} \int_\beta^1 V@R_b(\boldsymbol{X}) db$
   $\rightsquigarrow$ upperbound of V@R
3. $TV@R_\beta[\boldsymbol{X}] = \min_{t \in \mathbb{R}} \left\{ t + \frac{1}{1-\beta} \mathbb{E}[\boldsymbol{X} - t]^+ \right\}$
   $\rightsquigarrow$ linear formulation
4. $TV@R_\beta[\boldsymbol{X}] = \sup \left\{ \mathbb{E}_{\mathbb{Q}}[\boldsymbol{X}] \mid \frac{d\mathbb{Q}}{d\mathbb{P}} \leq \frac{1}{1-\beta} \right\}$
   $\rightsquigarrow$ coherent risk measure

Tail Value at Risk

The Tail Value at Risk (TV@R) (a.k.a conditional value at risk, average value at risk, expected shortfall or superquantile) is a convexification of the V@R measure. It has various equivalent definitions (assuming $\boldsymbol{X}$ admits a density):

1. $TV@R_\beta[\boldsymbol{X}] = \mathbb{E}[\boldsymbol{X} | \boldsymbol{X} \geq V@R_\beta(\boldsymbol{X})]$
   $\rightsquigarrow$ intuitive interpretation
2. $TV@R_\beta[\boldsymbol{X}] = \frac{1}{1-\beta} \int_\beta^1 V@R_b(\boldsymbol{X}) db$
   $\rightsquigarrow$ upperbound of V@R
3. $TV@R_\beta[\boldsymbol{X}] = \min_{t \in \mathbb{R}} \left\{ t + \frac{1}{1-\beta} \mathbb{E}[\boldsymbol{X} - t]^+ \right\}$
   $\rightsquigarrow$ linear formulation
4. $TV@R_\beta[\boldsymbol{X}] = \sup \left\{ \mathbb{E}_{\mathbb{Q}}[\boldsymbol{X}] \mid \frac{d\mathbb{Q}}{d\mathbb{P}} \leq \frac{1}{1-\beta} \right\}$
   $\rightsquigarrow$ coherent risk measure

Remark : with β going from 0 to 1 we smoothly get from the expectation to the worst case.

Convex combination of expectation and TV@R

- TV@R has very good properties:
    - law invariant coherent risk measure
    - linear programming formulation
    - upper bound of V@R
    - decent interpretation
- It is, however, quite risk averse and is sensitive only to the $1-\beta$ worst cases.
- A very common practice consists in considering

$$\rho[\boldsymbol{X}] = \lambda\mathbb{E}[\boldsymbol{X}] + (1-\lambda)\,TV@R_\beta[\boldsymbol{X}]$$

(with $\lambda \in [0,1], \beta \in [0,1]$) which is coherent with more flexibility in representing risk aversion.

- Actually most coherent risk measures can be represented as convex combinations of TV@R.

Contents

1. Introduction and motivations
   - How to add uncertainty in an optimization problem
2. Risk Measures
   - Definitions
   - Examples
   - Distributionally Robust Optimization
3. Multistage Stochastic Optimization under uncertainty
   - Stochastic optimal control problem
   - Dynamic Programming principle
   - Risk averse multistage
4. Practical application: two-stage case

General idea

- In stochastic optimization we assume that the true underlying probability $\mathbb{P}$ is known.
- In practice we generally do not have the probability $\mathbb{P}$, but only some data and a-priori.
- Classicaly we approximate the true probability $\mathbb{P}$ by the empirical probability $\hat{\mathbb{P}}_N$.
- In, Distributionally Robust Optimization, we
  - choose a distance d on the probability distribution
  - consider all the probability $\mathbb{Q}$ that are close to the empirical probability, i.e. $d(\mathbb{Q}, \hat{\mathbb{P}_N} \leq \varepsilon)$
  - and take a robust approach on the probability distribution :

$$\rho[\boldsymbol{X}] = \sup_{\mathbb{Q}: d(\mathbb{Q}, \hat{\mathbb{P}}_N) \leq \varepsilon} \mathbb{E}_{\mathbb{Q}}[\boldsymbol{X}]$$

Contents

1. Introduction and motivations
   - How to add uncertainty in an optimization problem
2. Risk Measures
   - Definitions
   - Examples
   - Distributionally Robust Optimization
3. Multistage Stochastic Optimization under uncertainty
   - Stochastic optimal control problem
   - Dynamic Programming principle
   - Risk averse multistage
4. Practical application: two-stage case

Contents

Stochastic Controlled Dynamic System

A discrete time controlled stochastic dynamic system is defined by its *dynamic*

$$\boldsymbol{x}_{t+1} = f_t(\boldsymbol{x}_t, \boldsymbol{u}_t, \boldsymbol{\xi}_{t+1})$$

and initial state

$$\boldsymbol{x}_0 = \boldsymbol{\xi}_0$$

The variables

- $\boldsymbol{x}_t$ is the *state* of the system,
- $\boldsymbol{u}_t$ is the *control* applied to the system at time t,
- $\boldsymbol{\xi}_t$ is an exogeneous noise.

Usually, $\boldsymbol{x}_t \in \mathbb{X}_t$ and $\boldsymbol{u}_t$ beglongs to a set depending upon the state: $\boldsymbol{u}_t \in U_t(\boldsymbol{x}_t)$.

Examples

- Stock of water in a dam:
    - $\boldsymbol{x}_t$ is the amount of water in the dam at time t,
    - $\boldsymbol{u}_t$ is the amount of water turbined at time t,
    - $\boldsymbol{\xi}_{t+1}$ is the inflow of water in $[t, t+1[$.
- Boat in the ocean:
    - $\boldsymbol{x}_t$ is the position of the boat at time t,
    - $\boldsymbol{u}_t$ is the direction and speed chosen for $[t, t+1[$,
    - $\boldsymbol{\xi}_{t+1}$ is the wind and current for $[t, t+1[$.
- Subway network:
    - $\boldsymbol{x}_t$ is the position and speed of each train at time t,
    - $\boldsymbol{u}_t$ is the acceleration chosen at time t,
    - $\boldsymbol{\xi}_{t+1}$ is the delay due to passengers and incident on the network for $[t, t+1[$.

More considerations about the state

- Physical state: the physical value of the controlled system. e.g. amount of water in your dam, position of your boat...
- Information state: physical state and information you have over noises. e.g.: amount of water and weather forecast...
- Knowledge state: your current belief over the actual information state (in case of noisy observations). Represented as a distribution law over information states.

The state in the Dynamic Programming sense is the information required to define an optimal solution.

Optimization Problem

We want to solve the following optimization problem

$$\begin{aligned}
\min_{\boldsymbol{u}} \quad & \mathbb{E}\Big[\sum_{t=0}^{T-1} L_t\big(\boldsymbol{x}_t, \boldsymbol{u}_t, \boldsymbol{\xi}_{t+1}\big) + K\big(\boldsymbol{x}_T\big)\Big] \\
s.t. \quad & \boldsymbol{x}_{t+1} = f_t(\boldsymbol{x}_t, \boldsymbol{u}_t, \boldsymbol{\xi}_{t+1}), \qquad \boldsymbol{x}_0 = \boldsymbol{\xi}_0 \\
& \boldsymbol{u}_t \in \mathcal{U}_t(\boldsymbol{x}_t, \boldsymbol{\xi}_{t+1}) \\
& \sigma(\boldsymbol{u}_t) \subset \sigma\big(\boldsymbol{\xi}_0, \cdots, \boldsymbol{\xi}_{t+1}\big)
\end{aligned}$$

Optimization Problem

We want to solve the following optimization problem

$$\begin{aligned} \min_{\boldsymbol{u}} \quad & \mathbb{E}\Big[\sum_{t=0}^{T-1} L_t(\boldsymbol{x}_t, \boldsymbol{u}_t, \xi_{t+1}) + K(\boldsymbol{x}_T)\Big] \\ s.t. \quad & \boldsymbol{x}_{t+1} = f_t(\boldsymbol{x}_t, \boldsymbol{u}_t, \xi_{t+1}), \qquad \boldsymbol{x}_0 = \xi_0 \\ & \boldsymbol{u}_t \in \mathcal{U}_t(\boldsymbol{x}_t, \xi_{t+1}) \\ & \sigma(\boldsymbol{u}_t) \subset \sigma(\xi_0, \cdots, \xi_{t+1}) \end{aligned}$$

1. We want to minimize the expectation of the sum of costs.
2. The system follows a dynamic given by the function f_t.
3. There are constraints on the controls.
4. The controls are functions of the past noises (= non-anticipativity).

Optimization Problem

We want to solve the following optimization problem

$$\begin{aligned} \min_{\Phi} \quad & \mathbb{E}\Big[\sum_{t=0}^{T-1} L_t\big(\boldsymbol{x}_t, \boldsymbol{u}_t, \xi_{t+1}\big) + K\big(\boldsymbol{x}_T\big) \Big] \\ s.t. \quad & \boldsymbol{x}_{t+1} = f_t(\boldsymbol{x}_t, \boldsymbol{u}_t, \xi_{t+1}), \qquad \boldsymbol{x}_0 = \xi_0 \\ & \boldsymbol{u}_t \in \mathcal{U}_t(\boldsymbol{x}_t, \xi_{t+1}) \\ & \boldsymbol{u}_t = \Phi(\xi_0, \cdots, \xi_{t+1}) \end{aligned}$$

1. We want to minimize the expectation of the sum of costs.
2. The system follows a dynamic given by the function f_t.
3. There are constraints on the controls.
4. The controls are functions of the past noises (= non-anticipativity).

Optimization Problem with independence of noises

If noises at time independent, the optimization problem is equivalent to

$$\begin{aligned}
\min_{\pi} \quad & \mathbb{E}\Big[\sum_{t=0}^{T-1} L_t\big(\boldsymbol{x}_t, \boldsymbol{u}_t, \boldsymbol{\xi}_{t+1}\big) + K\big(\boldsymbol{x}_T\big)\Big] \\
s.t. \quad & \boldsymbol{x}_{t+1} = f_t(\boldsymbol{x}_t, \boldsymbol{u}_t, \boldsymbol{\xi}_{t+1}), \qquad \boldsymbol{x}_0 = \boldsymbol{\xi}_0 \\
& \boldsymbol{u}_t \in \mathcal{U}_t(\boldsymbol{x}_t, \boldsymbol{\xi}_{t+1}) \\
& \boldsymbol{u}_t = \pi_t(\boldsymbol{x}_t, \boldsymbol{\xi}_{t+1})
\end{aligned}$$

Keeping only the state

For notational ease, we want to formulate Problem (**??**) only with states. Let $\mathcal{X}_t(x_t, \xi_{t+1})$ be the reachable states, i.e.,

$$\mathcal{X}_t(x_t, \xi_{t+1}) := \Big\{ x_{t+1} \in \mathbb{X}_{t+1} \quad \Big| \quad \exists u_t \in \mathcal{U}_t(x_t, \xi_{t+1}), \quad x_{t+1} = f_t(x_t, u_t, \xi_{t+1}) \Big\}.$$

And $c_t(x_t, x_{t+1}, \xi_{t+1})$ the transition cost from x_t to x_{t+1}, i.e.,

$$c_t(x_t, x_{t+1}, \xi_{t+1}) := \min_{u_t \in U_t(x_t, \xi_{t+1})} \Big\{ L_t(x_t, u_t, \xi_{t+1}) \quad \Big| \quad x_{t+1} = f_t(x_t, u_t, \xi_{t+1}) \Big\}.$$

Then, under independance of noises, the optimization problem reads

$$\begin{aligned} \min_{\psi} \quad & \mathbb{E}\Big[\sum_{t=0}^{T-1} c_t(\boldsymbol{x}_t, \boldsymbol{x}_{t+1}, \boldsymbol{\xi}_{t+1}) + K(x_T) \Big] \\ s.t. \quad & \boldsymbol{x}_{t+1} \in \mathcal{X}_t(\boldsymbol{x}_t, \boldsymbol{\xi}_{t+1}), \qquad x_0 = \xi_0 \\ & \boldsymbol{x}_{t+1} = \psi_t(\boldsymbol{x}_t, \boldsymbol{\xi}_{t+1}) \end{aligned}$$

Keeping only the state

For notational ease, we want to formulate Problem (**??**) only with states. Let $\mathcal{X}_t(x_t,\xi_{t+1})$ be the reachable states, i.e.,

$$\mathcal{X}_t(x_t,\xi_{t+1}) := \Big\{ x_{t+1} \in \mathbb{X}_{t+1} \quad \Big| \quad \exists u_t \in \mathcal{U}_t(x_t,\xi_{t+1}), \quad x_{t+1} = f_t(x_t,u_t,\xi_{t+1}) \Big\}.$$

And $c_t(x_t,x_{t+1},\xi_{t+1})$ the transition cost from x_t to x_{t+1}, i.e.,

$$c_t(x_t,x_{t+1},\xi_{t+1}) := \min_{u_t \in U_t(x_t,\xi_{t+1})} \Big\{ L_t(x_t,u_t,\xi_{t+1}) \quad \Big| \quad x_{t+1} = f_t(x_t,u_t,\xi_{t+1}) \Big\}.$$

Then, under independance of noises, the optimization problem reads

$$\begin{aligned} \min_{\psi} \quad & \mathbb{E}\Big[\sum_{t=0}^{T-1} c_t(\boldsymbol{x}_t,\boldsymbol{x}_{t+1},\boldsymbol{\xi}_{t+1}) + K(x_T) \Big] \\ s.t. \quad & \boldsymbol{x}_{t+1} \in \mathcal{X}_t(\boldsymbol{x}_t,\boldsymbol{\xi}_{t+1}), \qquad x_0 = \boldsymbol{\xi}_0 \\ & \boldsymbol{x}_{t+1} = \psi_t(\boldsymbol{x}_t,\boldsymbol{\xi}_{t+1}) \end{aligned}$$

Contents

Bellman's Principle of Optimality

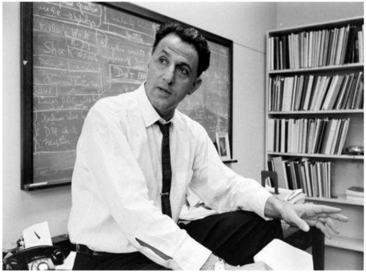

Richard Ernest Bellman (August 26, 1920 – March 19, 1984)

An optimal policy has the property that whatever the initial state and initial decision are, the remaining decisions must constitute an optimal policy with regard to the state resulting from the first decision (Richard Bellman)

The shortest path on a graph illustrates Bellman's Principle of Optimality

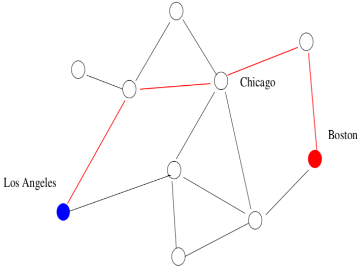

For an auto travel analogy, suppose that the fastest route from Los Angeles to Boston passes through **Chicago**. *The principle of optimality translates to obvious fact that the Chicago to Boston portion of the route is also the fastest route for a trip that starts from Chicago and ends in Boston. (Dimitri P. Bertsekas)*

Idea behind dynamic programming

If noises are time independent, then

1. The cost to go at time t depends only upon the current state.
2. We can compute recursively the cost to go for each position, starting from the terminal state and computing optimal trajectories backward.

Optimal cost-to-go of being in state x at time t is:
At time t, V_{t+1} gives the cost of the future. Dynamic

Programming is a time decomposition method.

Idea behind dynamic programming

If noises are time independent, then

1. The cost to go at time t depends only upon the current state.
2. We can compute recursively the cost to go for each position, starting from the terminal state and computing optimal trajectories backward.

Optimal cost-to-go of being in state x at time t is:
At time t, V_{t+1} gives the cost of the future. Dynamic

Programming is a time decomposition method.

Dynamic Programming Principle

Assume that the noises $\boldsymbol{\xi}_t$ are time-independent and exogeneous. The Bellman's equation writes

$$\begin{cases} V_T(x) & = K(x) \\ \hat{V}_t(x,\xi) & = \min\limits_{y \in \mathcal{X}_t(x,\xi)} c_t(x,y,\xi_{t+1}) + V_{t+1}(y) \\ V_t(x) & = \mathbb{E}\Big[\hat{V}_t(x,\boldsymbol{\xi}_{t+1})\Big] \end{cases}$$

An optimal state trajectory is obtained by $\boldsymbol{x}_{t+1} = \psi_t^V(\boldsymbol{x}_t)$, with

$$\psi_t^V(x,\xi) \in \mathop{\arg\min}_{y \in \mathcal{X}_t(x,\xi)} \underbrace{c_t(x,y,\xi)}_{\text{current cost}} + \underbrace{V_{t+1}(y)}_{\text{future costs}},$$

Dynamic Programming Principle

Assume that the noises $\boldsymbol{\xi}_t$ are time-independent and exogeneous. The Bellman's equation writes

$$\begin{cases} V_T(x) & = K(x) \\ \hat{V}_t(x, \xi) & = \min\limits_{y \in \mathcal{X}_t(x,\xi)} c_t(x, y, \xi_{t+1}) + V_{t+1}(y) \\ V_t(x) & = \mathbb{E}\Big[\hat{V}_t(x, \boldsymbol{\xi}_{t+1})\Big] \end{cases}$$

An optimal state trajectory is obtained by $\boldsymbol{x}_{t+1} = \psi_t^V(\boldsymbol{x}_t)$, with

$$\psi_t^V(x, \xi) \in \mathop{\arg\min}_{y \in \mathcal{X}_t(x,\xi)} \quad \underbrace{c_t(x, y, \xi)}_{\text{current cost}} + \underbrace{V_{t+1}(y)}_{\text{future costs}} ,$$

Interpretation of Bellman Value Function

The Bellman's value function $V_{t_0}(x)$ can be interpreted as the value of the problem starting at time t_0 from the state x.
More precisely we have

$$\begin{aligned} V_{t_0}(x) = \min \quad & \mathbb{E}\Big[\sum_{t=t_0}^{T-1} L_t\big(\boldsymbol{x}_t, \boldsymbol{u}_t, \boldsymbol{\xi}_{t+1}\big) + K\big(\boldsymbol{x}_T\big)\Big] \\ s.t. \quad & \boldsymbol{x}_{t+1} = f_t\big(\boldsymbol{x}_t, \boldsymbol{u}_t, \boldsymbol{\xi}_{t+1}\big), \qquad \boldsymbol{x}_{t_0} = x \\ & \boldsymbol{u}_t \in \mathcal{U}_t\big(\boldsymbol{x}_t, \boldsymbol{\xi}_{t+1}\big) \\ & \sigma\big(\boldsymbol{u}_t\big) \subset \sigma\big(\boldsymbol{\xi}_0, \cdots, \boldsymbol{\xi}_{t+1}\big) \end{aligned}$$

or

$$\begin{aligned} \min_{\psi} \quad & \mathbb{E}\Big[\sum_{t=t_0}^{T-1} c_t(x_t, x_{t+1}, \boldsymbol{\xi}_{t+1}) + K(x_T)\Big] \\ s.t. \quad & \boldsymbol{x}_{t+1} \in \mathcal{X}_t\big(\boldsymbol{x}_t, \boldsymbol{\xi}_{t+1}\big), \qquad x_{t_0} = x \\ & \boldsymbol{x}_{t+1} = \psi_t\big(\boldsymbol{x}_t\big) \end{aligned}$$

Interpretation of Bellman Value Function

The Bellman's value function $V_{t_0}(x)$ can be interpreted as the value of the problem starting at time t_0 from the state x.
More precisely we have

$$\begin{aligned} V_{t_0}(x) = \min \quad & \mathbb{E}\Big[\sum_{t=t_0}^{T-1} L_t\big(\boldsymbol{x}_t, \boldsymbol{u}_t, \boldsymbol{\xi}_{t+1}\big) + K\big(\boldsymbol{x}_T\big)\Big] \\ s.t. \quad & \boldsymbol{x}_{t+1} = f_t(\boldsymbol{x}_t, \boldsymbol{u}_t, \boldsymbol{\xi}_{t+1}), \qquad \boldsymbol{x}_{t_0} = x \\ & \boldsymbol{u}_t \in \mathcal{U}_t(\boldsymbol{x}_t, \boldsymbol{\xi}_{t+1}) \\ & \sigma(\boldsymbol{u}_t) \subset \sigma\big(\boldsymbol{\xi}_0, \cdots, \boldsymbol{\xi}_{t+1}\big) \end{aligned}$$

or

$$\begin{aligned} \min_{\psi} \quad & \mathbb{E}\Big[\sum_{t=t_0}^{T-1} c_t(x_t, x_{t+1}, \boldsymbol{\xi}_{t+1}) + K(x_T)\Big] \\ s.t. \quad & \boldsymbol{x}_{t+1} \in \mathcal{X}_t(\boldsymbol{x}_t, \boldsymbol{\xi}_{t+1}), \qquad x_{t_0} = x \\ & \boldsymbol{x}_{t+1} = \psi_t(\boldsymbol{x}_t) \end{aligned}$$

Contents

The problem of time consistency

The most natural way of writing a risk-averse multistage problem is to consider the following

$$\begin{aligned} \min \quad & \rho\Big[\sum_{t=1}^{T} L_t(\boldsymbol{X}_t, \boldsymbol{U}_t, \boldsymbol{\xi}_t)\Big] \\ s.t. \quad & \boldsymbol{X}_{t+1} = f_t(\boldsymbol{X}_t, \boldsymbol{U}_t, \boldsymbol{\xi}_t) \end{aligned}$$

Unfortunately there is no reason for the solution of the above problem to be also a solution of the same problem starting at time t_0

$$\begin{aligned} \min \quad & \rho\Big[\sum_{t=t_0}^{T} L_t(\boldsymbol{X}_t, \boldsymbol{U}_t, \boldsymbol{\xi}_t)\Big] \\ s.t. \quad & \boldsymbol{X}_{t+1} = f_t(\boldsymbol{X}_t, \boldsymbol{U}_t, \boldsymbol{\xi}_t) \end{aligned}$$

Composed risk measures

The solution to the above conundrum is to consider an alternative problem with nested risk measures

$$\min_{u_0} \rho\Big[L_0(x_0, u_0, \xi_0) + \min_{\boldsymbol{u}_1} \rho\big[L_1(\boldsymbol{X}_1, \boldsymbol{u}_1, \boldsymbol{\xi}_1) + \min_{\boldsymbol{u}_2} \rho[\dots]\big]\Big]$$

In which case Dynamic Programming principle easily apply by replacing expectation by ρ with

$$\begin{cases} V_T = K \\ V_t(x) = \min_u \rho\Big[L_t(x, u, \boldsymbol{\xi}_t) + V_{t+1}(f_t(x, u, \boldsymbol{\xi}_t))\Big] \end{cases}$$

Main downside : interpretation of what we are doing is not easy.

Composed risk measures

The solution to the above conundrum is to consider an alternative problem with nested risk measures

$$\min_{u_0} \rho\Big[L_0(x_0, u_0, \xi_0) + \min_{\boldsymbol{u}_1} \rho\big[L_1(\boldsymbol{X}_1, \boldsymbol{u}_1, \boldsymbol{\xi}_1) + \min_{\boldsymbol{u}_2} \rho[\dots]\big]\Big]$$

In which case Dynamic Programming principle easily apply by replacing expectation by ρ with

$$\begin{cases} V_T = K \\ V_t(x) = \min_u \rho\Big[L_t(x, u, \boldsymbol{\xi}_t) + V_{t+1}(f_t(x, u, \boldsymbol{\xi}_t))\Big] \end{cases}$$

Main downside : interpretation of what we are doing is not easy.

Contents

A newsvendor problem

We consider the following one-stage problem

$$\begin{aligned} \min_{x\in\mathbb{R}} \quad & \rho\Big[c^\top x + q_+(x-\boldsymbol{d})^+ + q_-(x-\boldsymbol{d})^-\Big] \\ \text{s.t.} \quad & x \geq 0 \end{aligned}$$

where

- x is a quantity of product bought
- $\boldsymbol{d}$ is a random demand
- c is the cost of buying a product
- q_+ is the destruction cost
- q_- is the shortage cost

Exercise

We assume that $\boldsymbol{d}$ is uniformly distributed on $\{d_1, \dots, d_n\}$
Write the following problem as an LP

$$\begin{aligned} \min_{x \in \mathbb{R}} \quad & \rho\Big[c^\top x + q_+(x - \boldsymbol{d})^+ + q_-(x - \boldsymbol{d})^-\Big] \\ \text{s.t.} \quad & x \geq 0 \end{aligned}$$

when

1. $\rho = \mathbb{E}$
2. $\rho = TV@R_\beta \qquad (\beta \in (0,1))$
3. $\rho = \alpha TV@R_\beta + (1-\alpha) TV@R_\beta \qquad (\beta \in (0,1), \quad \alpha \in [0,1])$