VaR and CVaR: measuring and optimizing risk

Outline

Managerial motivation

Definitions and properties

Worked example by hand

Case study 1: mean-CVaR portfolio

Case study 2: two-stage supply chain

Sensitivity analysis

Exercises

Learning objectives (1/2)

At the end of this chapter you will be able to:

- define and compute VaR and CVaR on a discrete distribution;
- use the Rockafellar–Uryasev linear formulation;

Learning objectives (2/2)

At the end of this chapter you will be able to:

- build portfolios and logistics plans protected against the worst scenarios;
- recognize the statistical limits of tail estimates.

Prerequisites

- From this module, Chapter 7 «The Newsvendor and its variants»: the scenario-based LP, the positive-part trick and the mean-CVaR objective
- From this module, Chapter 2 «Background: linear programming and sensitivity analysis»: shadow prices and the sensitivity protocol

Managerial motivation

Introduction: Managerial motivation

- In this section we see **why** optimizing an average value is not enough
- We introduce the two risk measures of the chapter: **VaR** and **CVaR**
- We list the goals: definitions, LP model, case studies, statistical limits

The average is not enough

- Many models optimize an **average value**: a simple choice, but one that can hide **rare and very large** losses
- VaR and CVaR describe the **tail** of the loss distribution and turn risk aversion into quantitative decisions
- **VaR** answers: "what is a high loss threshold?"
- **CVaR** also answers: "how much do we lose *on average* when that threshold is exceeded?"

The language of tail risk

- After the financial crises, "how much do we lose in the worst scenarios?" has become a **regulatory** question as well as a managerial one
- VaR and CVaR are the standard language with which banks, insurers and, increasingly, supply chain managers speak about tail risk
- **The didactic surprise:** CVaR, an apparently sophisticated concept, is optimized with a simple **LP**

Summary

- Optimizing the average can hide rare and large losses
- VaR: a high loss **threshold**; CVaR: the **average loss beyond** the threshold
- CVaR is optimized with an LP: this is the core of the chapter

Definitions and properties

Introduction: Definitions and properties

- We formally define the **VaR** and the **CVaR** of a random loss
- We compare their properties: why CVaR is preferable for optimization
- We build the **Rockafellar–Uryasev LP model**, which computes and optimizes the CVaR over a set of scenarios

The problem in words

- *We decide* as before: portfolio weights, quantities, booked capacity
- *The objective* changes: not to minimize the **average** loss, but rather a measure of its **tail**, that is, how much is lost in the worst scenarios
- *The constraints* remain those of the underlying problem, plus the auxiliary linear constraints that define the CVaR

Random loss and confidence level

- Let L be the **random loss** (high values = worse outcomes)
- Let $\alpha \in (0, 1)$ be a **confidence level** (typically 0.90–0.99)
- As for the demand of the newsvendor, the random variable is in upper case and its realizations in lower case (ℓ_s)
- This is the exception, to follow the literature, to the convention that reserves capital letters for sets and matrices

VaR and CVaR

VaR and CVaR

$$\mathrm{VaR}_\alpha(L) = \inf\{\eta \in \mathbb{R} : \mathbb{P}(L \leq \eta) \geq \alpha\},$$

$$\mathrm{CVaR}_\alpha(L) = \text{average loss in the worst tail of mass } 1-\alpha.$$

Properties compared

Convexity	CVaR is convex in the decisions (if L is); VaR in general is not
Subadditivity	CVaR rewards diversification; VaR may violate it
Tail	VaR does not distinguish two tails with the same quantile but different severity
Optimization	CVaR admits a scenario-based LP formulation; VaR does not

Data of the scenario-based model

- Number of scenarios $k \in \mathbb{Z}_{\geq 1}$, indexed by $s \in \{1, 2, \ldots, k\}$
- For each scenario, the probability $\pi_s \in \mathbb{Q}_{\geq 0}$, with $\sum_{s=1}^{k} \pi_s = 1$
- For each scenario, the loss $\ell_s(\boldsymbol{x})$, a linear function of the decision vector $\boldsymbol{x}$, constrained to a polyhedron X
- The confidence level $\alpha \in (0, 1)$

Decision variables

Besides the decisions $\boldsymbol{x}$, we introduce one free variable and k non-negative variables:

$$\begin{cases} \eta = \text{candidate loss threshold (at the optimum: a VaR)} \\ \xi_s = \text{loss in excess of } \eta \text{ in scenario } s \end{cases}$$

$$\forall s \in \{1, 2, \ldots, k\}.$$

Rockafellar–Uryasev LP model

Rockafellar–Uryasev LP model for the CVaR

$$\min \quad \eta + \frac{1}{1-\alpha} \sum_{s=1}^{k} \pi_s \, \xi_s \tag{8.1a}$$

$$\text{subject to} \quad \xi_s \geq \ell_s(\boldsymbol{x}) - \eta, \quad \forall s \in \{1, 2, \ldots, k\}, \tag{8.1b}$$

$$\boldsymbol{x} \in X, \tag{8.1c}$$

$$\eta \gtreqless 0, \tag{8.1d}$$

$$\xi_s \geq 0, \qquad \forall s \in \{1, 2, \ldots, k\}. \tag{8.1e}$$

Description of the model (1/2)

- **Objective function:** the threshold η plus the weighted average of the tail excesses, amplified by $1/(1-\alpha)$
- *At the optimum η^* is a VaR and the objective value is the CVaR*: a single LP provides both measures
- **Tail constraints** $\xi_s \geq \ell_s(\boldsymbol{x}) - \eta$: they define the excesses; only the scenarios with a loss beyond the threshold contribute (k linear constraints)

Description of the model (2/2)

- **Constraint $\boldsymbol{x} \in X$**: collects the constraints of the underlying decision problem (polyhedral)
- **Definition of the variables:** η is *free* (the symbol $\gtreqless$), the ξ_s are non-negative

Why the model works

- The variable $\xi_s = (\ell_s - \eta)^+$ is the loss in excess of the candidate threshold η (the same positive-part trick as in the Newsvendor)
- Minimizing over η balances two forces: raising η reduces the excesses but pushes up the threshold
- The equilibrium point is exactly the α **quantile**
- If X is a polyhedron and ℓ_s is linear, the whole model is an **LP**

Summary

- VaR_α: α quantile of the loss; CVaR_α: average loss in the tail of mass $1-\alpha$
- CVaR is convex, subadditive, sensitive to the severity of the tail and optimizable with an LP; VaR is not
- Rockafellar–Uryasev model: free threshold η, excesses $\xi_s \geq 0$, k tail constraints
- At the optimum: η^* is a VaR, the optimal value is the CVaR

Worked example by hand

Introduction: Worked example by hand

- We compute VaR and CVaR by hand on an example with **six scenarios**
- We check that the Rockafellar–Uryasev LP returns the same values
- We compare mean, VaR and CVaR: three numbers, three stories

Example: six scenarios, $\alpha = 0.80$ (1/3)

- Equally likely losses: $\{2, 4, 5, 7, 12, 20\}$, each with probability $1/6$
- **VaR.** The cumulative distribution reaches 0.80 at the fifth ordered loss:

$$\mathrm{VaR}_{0.80} = 12$$

Example: six scenarios, $\alpha = 0.80$ (2/3)

- **CVaR with the formula.** The worst tail has mass 0.20
- The scenario with loss 20 weighs $1/6 = 0.167$, and the remaining mass 0.033 sits *on the point* 12:

$$\mathrm{CVaR}_{0.80} = 12 + \frac{1}{0.20} \cdot \frac{1}{6} (20 - 12) = 12 + 6.67 = 18.67$$

Example: six scenarios, $\alpha = 0.80$ (3/3)

- **Check with the LP** (script): $\eta^* = 12$, $\xi = (0, 0, 0, 0, 0, 8)$, value 18.6667
- The formula correctly handles the mass "straddling" the quantile without selecting scenarios with binary variables
- **Comparison.** Average loss 8.33; VaR 12; CVaR 18.67
- The mean reassures, the VaR gives a threshold, the CVaR says how much it hurts when things go badly

Summary

- The VaR is read on the cumulative distribution: it is the loss at which it reaches α
- The CVaR averages the tail of mass $1 - \alpha$, also counting the mass that sits on the quantile
- The LP returns the same values without binary variables

Case study 1: mean-CVaR portfolio

Introduction: mean-CVaR portfolio

- We apply the Rockafellar–Uryasev model to a **portfolio** of securities
- Risk is measured with the CVaR of the monthly loss, for a fixed expected return
- We will see that the CVaR looks **only at the tail** of the losses

The case: eight securities with fat tails

- Eight sector stocks and $k = 220$ monthly scenarios generated with **fat tails** (Student t market factor, 4 degrees of freedom)
- It is the world in which the CVaR shows its value
- Scenario loss: $\ell_s(\boldsymbol{x}) = -\sum_{i=1}^{n} r_{is}\, x_i$, where r_{is} is the return of security i in scenario s
- We minimize $\mathrm{CVaR}_{0.90}$ with expected return $\geq 8\%$ per year

The model in Gurobi

```
eta = m.addVar(lb=-GRB.INFINITY, name="eta")  # threshold: free!
xi  = m.addVars(S, name="xi")
m.addConstr(x.sum() == 1)
m.addConstr(gp.quicksum(mu[i]*x[i] for i in range(n))
            >= r_min)
m.addConstrs((xi[s] >=
              -gp.quicksum(R[s, i]*x[i] for i in range(n))
              - eta for s in range(S)))
m.setObjective(eta + gp.quicksum(xi[s] for s in range(S))
               / (S * (1 - alpha)), GRB.MINIMIZE)
```

Results: risk measures

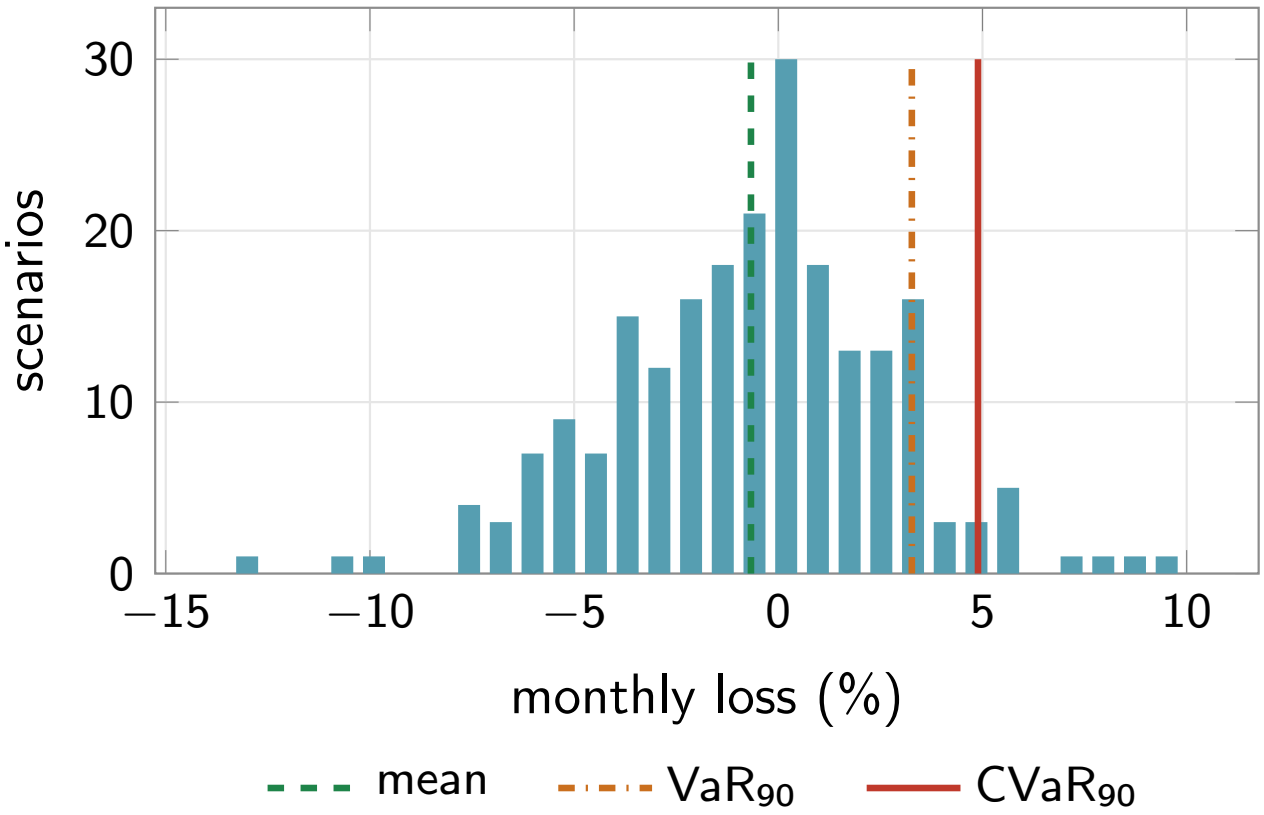
Monthly losses of the mean-CVaR portfolio:

	mean loss	VaR_{90}	CVaR_{90}
mean-CVaR	-0.0067	0.0327	0.0489

Results: compositions (%)

	ENE	FIN	TEC	SAN	UTL	MAT
mean-CVaR	9.9	1.9	10.5	8.6	46.5	22.5

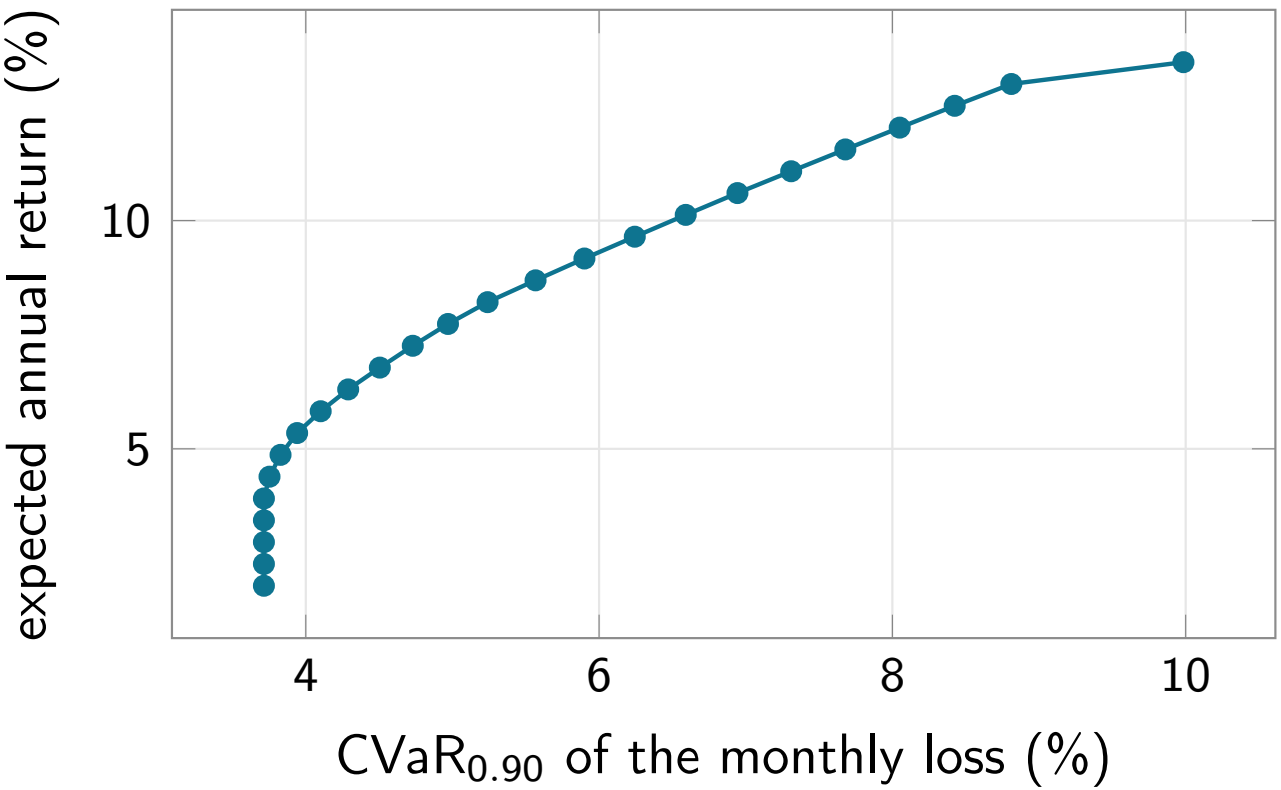
Loss distribution (mean-CVaR)



How to read the histogram

- Histogram of the monthly losses of the mean-CVaR portfolio over the 220 scenarios
- Vertical lines: mean, VaR_{90} and CVaR_{90}
- The CVaR is the **average of the tail** to the right of the VaR

Return–CVaR frontier



Obtained by varying the minimum required return.

The CVaR looks at the tail

- With an expected return of 8% per year the portfolio gains 0.67% per month on average, but in the worst 10% of the months it loses 4.89% on average
- The CVaR concentrates the portfolio on UTL and MAT and keeps a share of TEC, which although mediocre on average moves against the trend in the worst scenarios
- Measuring risk with the variance leads to the Markowitz portfolio, a quadratic model: it is in the module *NLP Modelling*, "The Markowitz portfolio"

Summary

- With an expected return of 8% the portfolio gains 0.67% per month on average, but in the worst 10% of the months it loses 4.89% on average
- The portfolio concentrates on UTL and MAT, with securities that protect in the worst scenarios
- Varying the minimum return traces the return–CVaR frontier

Case study 2: two-stage supply chain

Introduction: two-stage supply chain

- We use the CVaR in a **two-stage logistics** problem: first capacity is booked, then the recourse is decided
- A parameter λ balances average cost and tail risk
- We will measure the **cost of resilience**: how much average cost is needed to reduce the CVaR

The case: two suppliers

- Before observing the scenario we *book* capacity from the suppliers; then, having seen the scenario, we decide the **recourse**
- **F1** is cheap but fragile: in 12% of the scenarios its availability collapses to 30%
- **F2** is expensive but reliable
- Shortage penalty: 40 euro/unit; $k = 400$ scenarios

Data of the model

- Two suppliers $a \in \{1, 2\}$ (1 cheap but fragile, 2 expensive but reliable)
- Booking costs $c_a^{\text{pre}} \in \mathbb{Q}_{>0}$ and purchase costs $c_a^{\text{uso}} \in \mathbb{Q}_{>0}$
- k scenarios $s \in \{1, 2, \ldots, k\}$ with demand $d_s \in \mathbb{Q}_{\geq 0}$
- Availability $\delta_{as} \in (0, 1]$: the datum that models the failure
- Shortage penalty $g \in \mathbb{Q}_{>0}$
- Risk weight $\lambda \in [0, 1]$ and level $\alpha \in (0, 1)$

Decision variables

Non-negative variables:

- x_a = capacity booked from supplier a: **first-stage**, the same in all scenarios
- f_{as} = purchase from supplier a in scenario s
- u_s = unmet demand in scenario s

The f_{as} and the u_s are **recourse** variables.

Two-stage model with CVaR

Two-stage model with CVaR

$$\min \ (1-\lambda)\Big[\sum_{a=1}^{2} c_a^{\text{pre}} x_a + \sum_{s=1}^{k} \pi_s \, \phi_s\Big] + \lambda \, \mathrm{CVaR}_\alpha \tag{8.2a}$$

$$\text{subject to} \quad f_{as} \le \delta_{as} \, x_a, \quad \forall a \in \{1.2\}, \ \forall s \in \{1, 2, \ldots, k\}, \tag{8.2b}$$

$$f_{1s} + f_{2s} + u_s \ge d_s, \qquad \forall s \in \{1, 2, \ldots, k\}, \tag{8.2c}$$

$$x_a, \ f_{as}, \ u_s \ge 0, \qquad \forall a \in \{1.2\}, \ \forall s \in \{1, 2, \ldots, k\},$$

where $\phi_s = \sum_{a=1}^{2} c_a^{\text{uso}} f_{as} + g \, u_s$ is the recourse cost of scenario s and CVaR_α is the CVaR of the total cost, linearized with the variables η, ξ_s of (8.1).

Description of the model (1/2)

- **Objective function:** combines, with weight λ, the average cost (bookings plus the expected value of purchases and penalties) and the CVaR of the total cost
- **Capacity constraints** $f_{as} \leq \delta_{as} x_a$: in scenario s supplier a can be used only up to the booked capacity x_a, reduced by the availability δ_{as} ($2\,k$ linear constraints)

Description of the model (2/2)

- **Demand constraints:** demand is covered by the flows or, failing that, counted as a shortage u_s, which in the objective pays the penalty (k linear constraints)
- The model is never infeasible: it "buys" infeasibility at the price of the penalty
- **Last constraints:** they define the variables of the model

Results as λ varies

λ	suppl. 1	suppl. 2	mean cost	CVaR_{90}	service
0.0	161.5	134.2	1,724	2,902	98.7%
0.5	92.7	212.7	1,803	2,217	99.9%
1.0	68.5	236.9	1,906	2,139	99.0%

Columns suppl. 1 and suppl. 2 report the booked capacity.

The cost of resilience

- The risk-neutral decision maker ($\lambda = 0$) buys from the cheap supplier and accepts that in 12% of the scenarios the shortage makes costs explode (CVaR 2,902)
- As λ grows, capacity migrates towards the reliable supplier
- $+79$ euro of average cost buy -685 euro of CVaR
- This is the **cost of resilience**: the number needed in a discussion about reshoring and dual sourcing

Summary

- Two stages: capacity booked before the scenario, purchases and shortage after it
- The objective combines average cost and CVaR with weight λ; the shortage makes the model always feasible
- As λ grows, capacity moves to the reliable supplier: the CVaR drops much more than the average cost rises

Sensitivity analysis

Introduction: Sensitivity analysis

- We see which parameters to vary and with which expected effect
- We warn about the **statistical limits** of tail estimates

Parameters to vary

α	higher $\Rightarrow$ a more extreme tail estimated from fewer scenarios: more scenarios are needed for stable estimates
λ	controls the trade-off between average performance and protection
π_s	reweighted scenarios or stress tests (adding extreme scenarios)
$\|S\|$	statistical stability: repeat with different samples, as in the newsvendor chapter

Beware of tail estimates

Observation 1: Tails estimated from few scenarios

With $\alpha = 0.99$ and $k = 220$ scenarios the tail contains only 2–3 scenarios: the estimated CVaR is almost noise.

Rule of thumb: at least a few dozen scenarios are needed beyond the quantile.

With discrete distributions the VaR may not be unique (η^* is a VaR): always compare η^* with the empirical quantile.

Summary

- α, λ, π_s and the number of scenarios are the levers of the analysis
- Extreme tails with few scenarios give noisy estimates: dozens of scenarios beyond the quantile are needed
- η^* is *a* VaR: compare it with the empirical quantile

Exercises

Introduction: Exercises

- Five exercises to consolidate the concepts of the chapter
- They range from computation by hand to modifying the models of the case studies, to be solved with the lab script

Exercise: check

- Recompute by hand the VaR and CVaR of the 6-scenario example with $\alpha = 0.90$
- Check with the LP
- Careful: the tail has mass 0.10 and the scenario with loss 20 weighs $0.167 > 0.10$

Exercise: subadditivity

- Build a scenario-based example with two activities in which

$$\mathrm{VaR}(L_1 + L_2) > \mathrm{VaR}(L_1) + \mathrm{VaR}(L_2)$$

- Check that the CVaR does not violate subadditivity

Exercise: frontier in α

- For the portfolio, plot the optimal CVaR for

$$\alpha \in \{0.80;\ 0.90;\ 0.95;\ 0.99\}$$

  at equal return
- How does the composition change?
- When do the numerical instabilities begin?

Exercise: CVaR-constrained

- Reformulate the portfolio as

$$\max \sum_{i=1}^{n} \mu_i x_i \quad \text{subject to} \quad \mathrm{CVaR}_{0.90} \leq \bar{c}$$

- Check that by varying $\bar{c}$ one traces the same frontier
- Read the dual of the constraint

Exercise: stress test

- Add to the supply chain an extreme scenario (demand +50% and F1 at 10%) with probability 1%
- How much does the optimal booking change for $\lambda = 0.5$?
- And the CVaR?

Final summary (1/2)

Risk measures:

- VaR: loss threshold (α quantile); CVaR: average loss in the tail of mass $1 - \alpha$
- CVaR is convex, subadditive and optimized with an LP

Rockafellar–Uryasev model:

- Free threshold η and excesses $\xi_s \geq \ell_s(\boldsymbol{x}) - \eta$
- At the optimum η^* is a VaR and the optimal value is the CVaR

Final summary (2/2)

Case studies:

- Portfolio: the mean-CVaR portfolio controls the tail of the losses
- Supply chain: λ measures the cost of resilience

Limits:

- Extreme tails with few scenarios give noisy estimates
- η^* is *a* VaR: compare it with the empirical quantile