



Uncertainty quantification in presence of missing values

Margaux Zaffran, Aymeric Dieuleveut, Julie Josse, Yaniv Romano

► To cite this version:

Margaux Zaffran, Aymeric Dieuleveut, Julie Josse, Yaniv Romano. Uncertainty quantification in presence of missing values. ICSDS 2022 - IMS International Conference on Statistics and Data Science, Dec 2022, Florence, Italy. hal-03896384v1

HAL Id: hal-03896384

<https://hal.science/hal-03896384v1>

Submitted on 13 Dec 2022 (v1), last revised 9 Nov 2023 (v4)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Uncertainty quantification in presence of missing values

Margaux Zaffran^[1,2,3] Aymeric Dieuleveut^[3] Julie Josse^[2] Yaniv Romano^[4]

^[1]Electricité De France, Paris, France ^[2]INRIA, Montpellier, France ^[3]Ecole Polytechnique, Paris, France ^[4]Technion - Israel Institute of Technology, Haifa, Israel



Motivations and setting

Objectives

- Characterize the **impact of missing values** on **uncertainty** of the outcome.
- Propose a **methodology** outputting **predictive intervals** with **conditional coverage guarantees** with respect to **each pattern of missing values**.

- $(X, Y) \in \mathbb{R}^d \times \mathbb{R}$ random variables.

- Missing pattern** (mask) $M \in \{0, 1\}^d$: there are **2^d patterns**.

$$X = (1, \text{NA}, 2) \Rightarrow M = (0, 1, 0) \text{ and } X_{\text{obs}(M)} = (1, 2).$$

- Missing mechanism: Missing Completely At Random (**MCAR**)

for all $m \in \{0, 1\}^d$, $\mathbb{P}(M = m|X) = \mathbb{P}(M = m)$, i.e. $M \perp\!\!\!\perp X$.

- Framework: learn Y given $X_{\text{obs}(M)}$ and M .

- Most popular strategies to deal with missing values: **imputation**.

ϕ denotes an imputation function (e.g. replaces **NA** by a constant, the empirical mean, etc).

Exchangeability after imputation

Let $(X^{(k)}, M^{(k)}, Y^{(k)})_{k=1}^n$ be i.i.d.. Then, for any missing mechanism, for almost all imputation function ϕ : $(\phi(X_{\text{obs}(M^{(k)})}^{(k)}, M^{(k)}), M^{(k)}, Y^{(k)})_{k=1}^n$ is **exchangeable**.

Infinite data

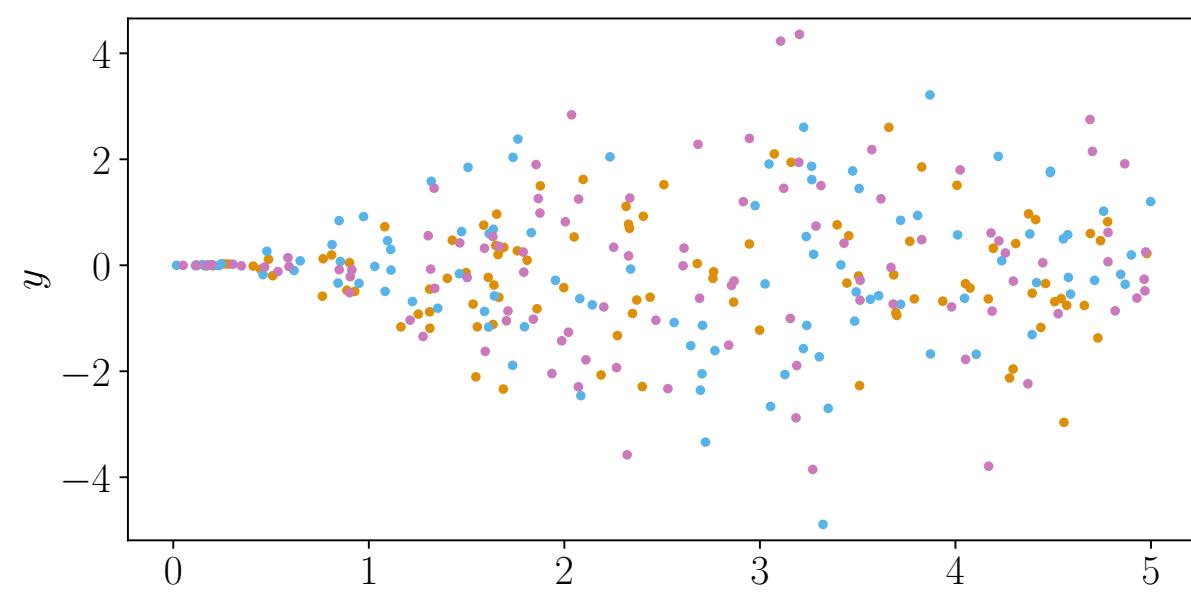
Consider **Impute-then-Regress** procedures, e.g. $g \circ \phi$. Define $g_{\delta, \phi}^* \in \operatorname{argmin}_{g: \mathbb{R}^d \rightarrow \mathbb{R}} \mathbb{E}[\rho_\delta(Y - g \circ \phi(X_{\text{obs}(M)}, M))]$, where ρ_δ is the **pinball loss** associated to the quantile of level δ .

Theorem

For almost all functions $\phi \in \mathcal{F}_\infty^I$, $g_{\delta, \phi}^* \circ \phi$ is Bayes optimal for the pinball-risk of level δ .

A universally consistent learner trained on deterministically imputed data set will be Bayes optimal.
⇒ it will reach conditional coverage with respect to the missing data pattern.

Finite sample: Conformalized Quantile Regression (CQR, Romano et al., 2019)



Randomly split the data to obtain a **proper training set** and a **calibration set**. Keep the **test set**.

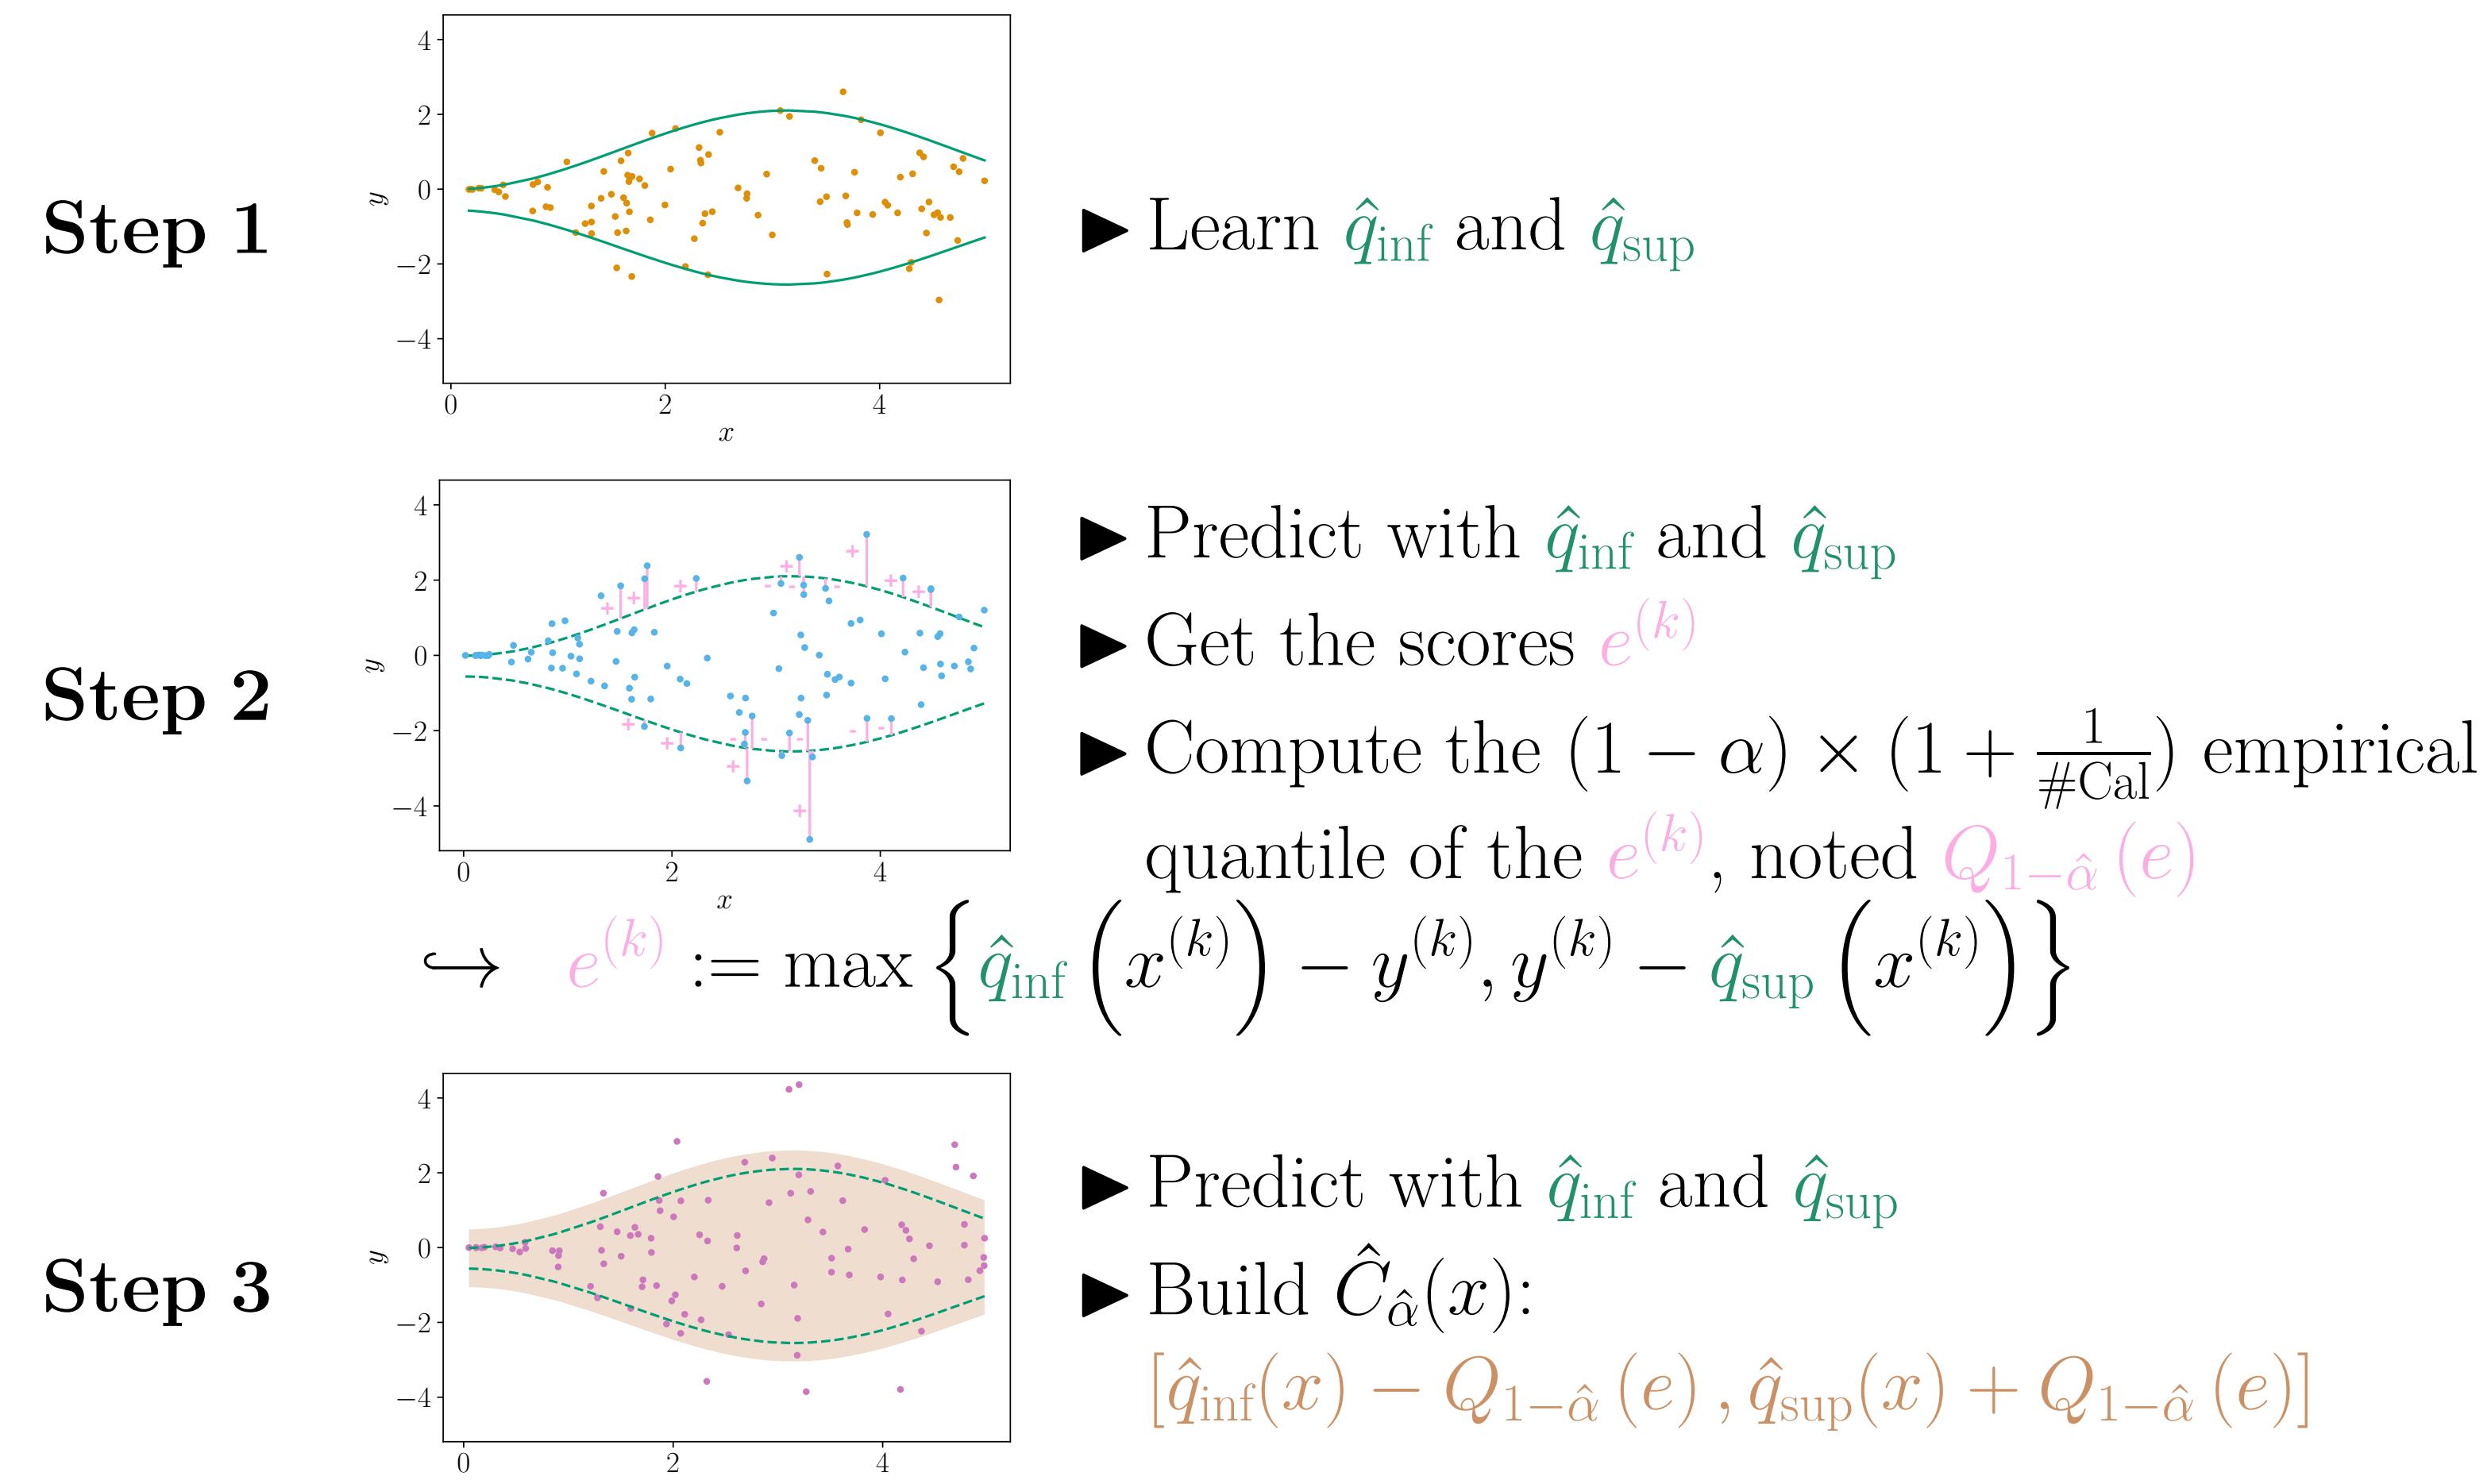
- Given any quantile regression functions \hat{q}_{inf} and \hat{q}_{sup}

- For any (**finite**) sample size n

- If the $(X^{(k)}, Y^{(k)})$ are **exchangeable**

$$\mathbb{P}(Y \in \hat{C}_{\hat{\alpha}}(X)) \geq 1 - \alpha$$

⇒ CQR is **marginally valid** on imputed data sets.



Proposed algorithms

Idea: generate **additional missing values** in the calibration set.

Test point			
3	NA	NA	1
Initial calibration set			
$x^{(1)}$	-1	-10	6
$x^{(2)}$	4	NA	-2
$x^{(3)}$	5	1	1
$x^{(4)}$	0	NA	1

Calibration set used in CQR with masking

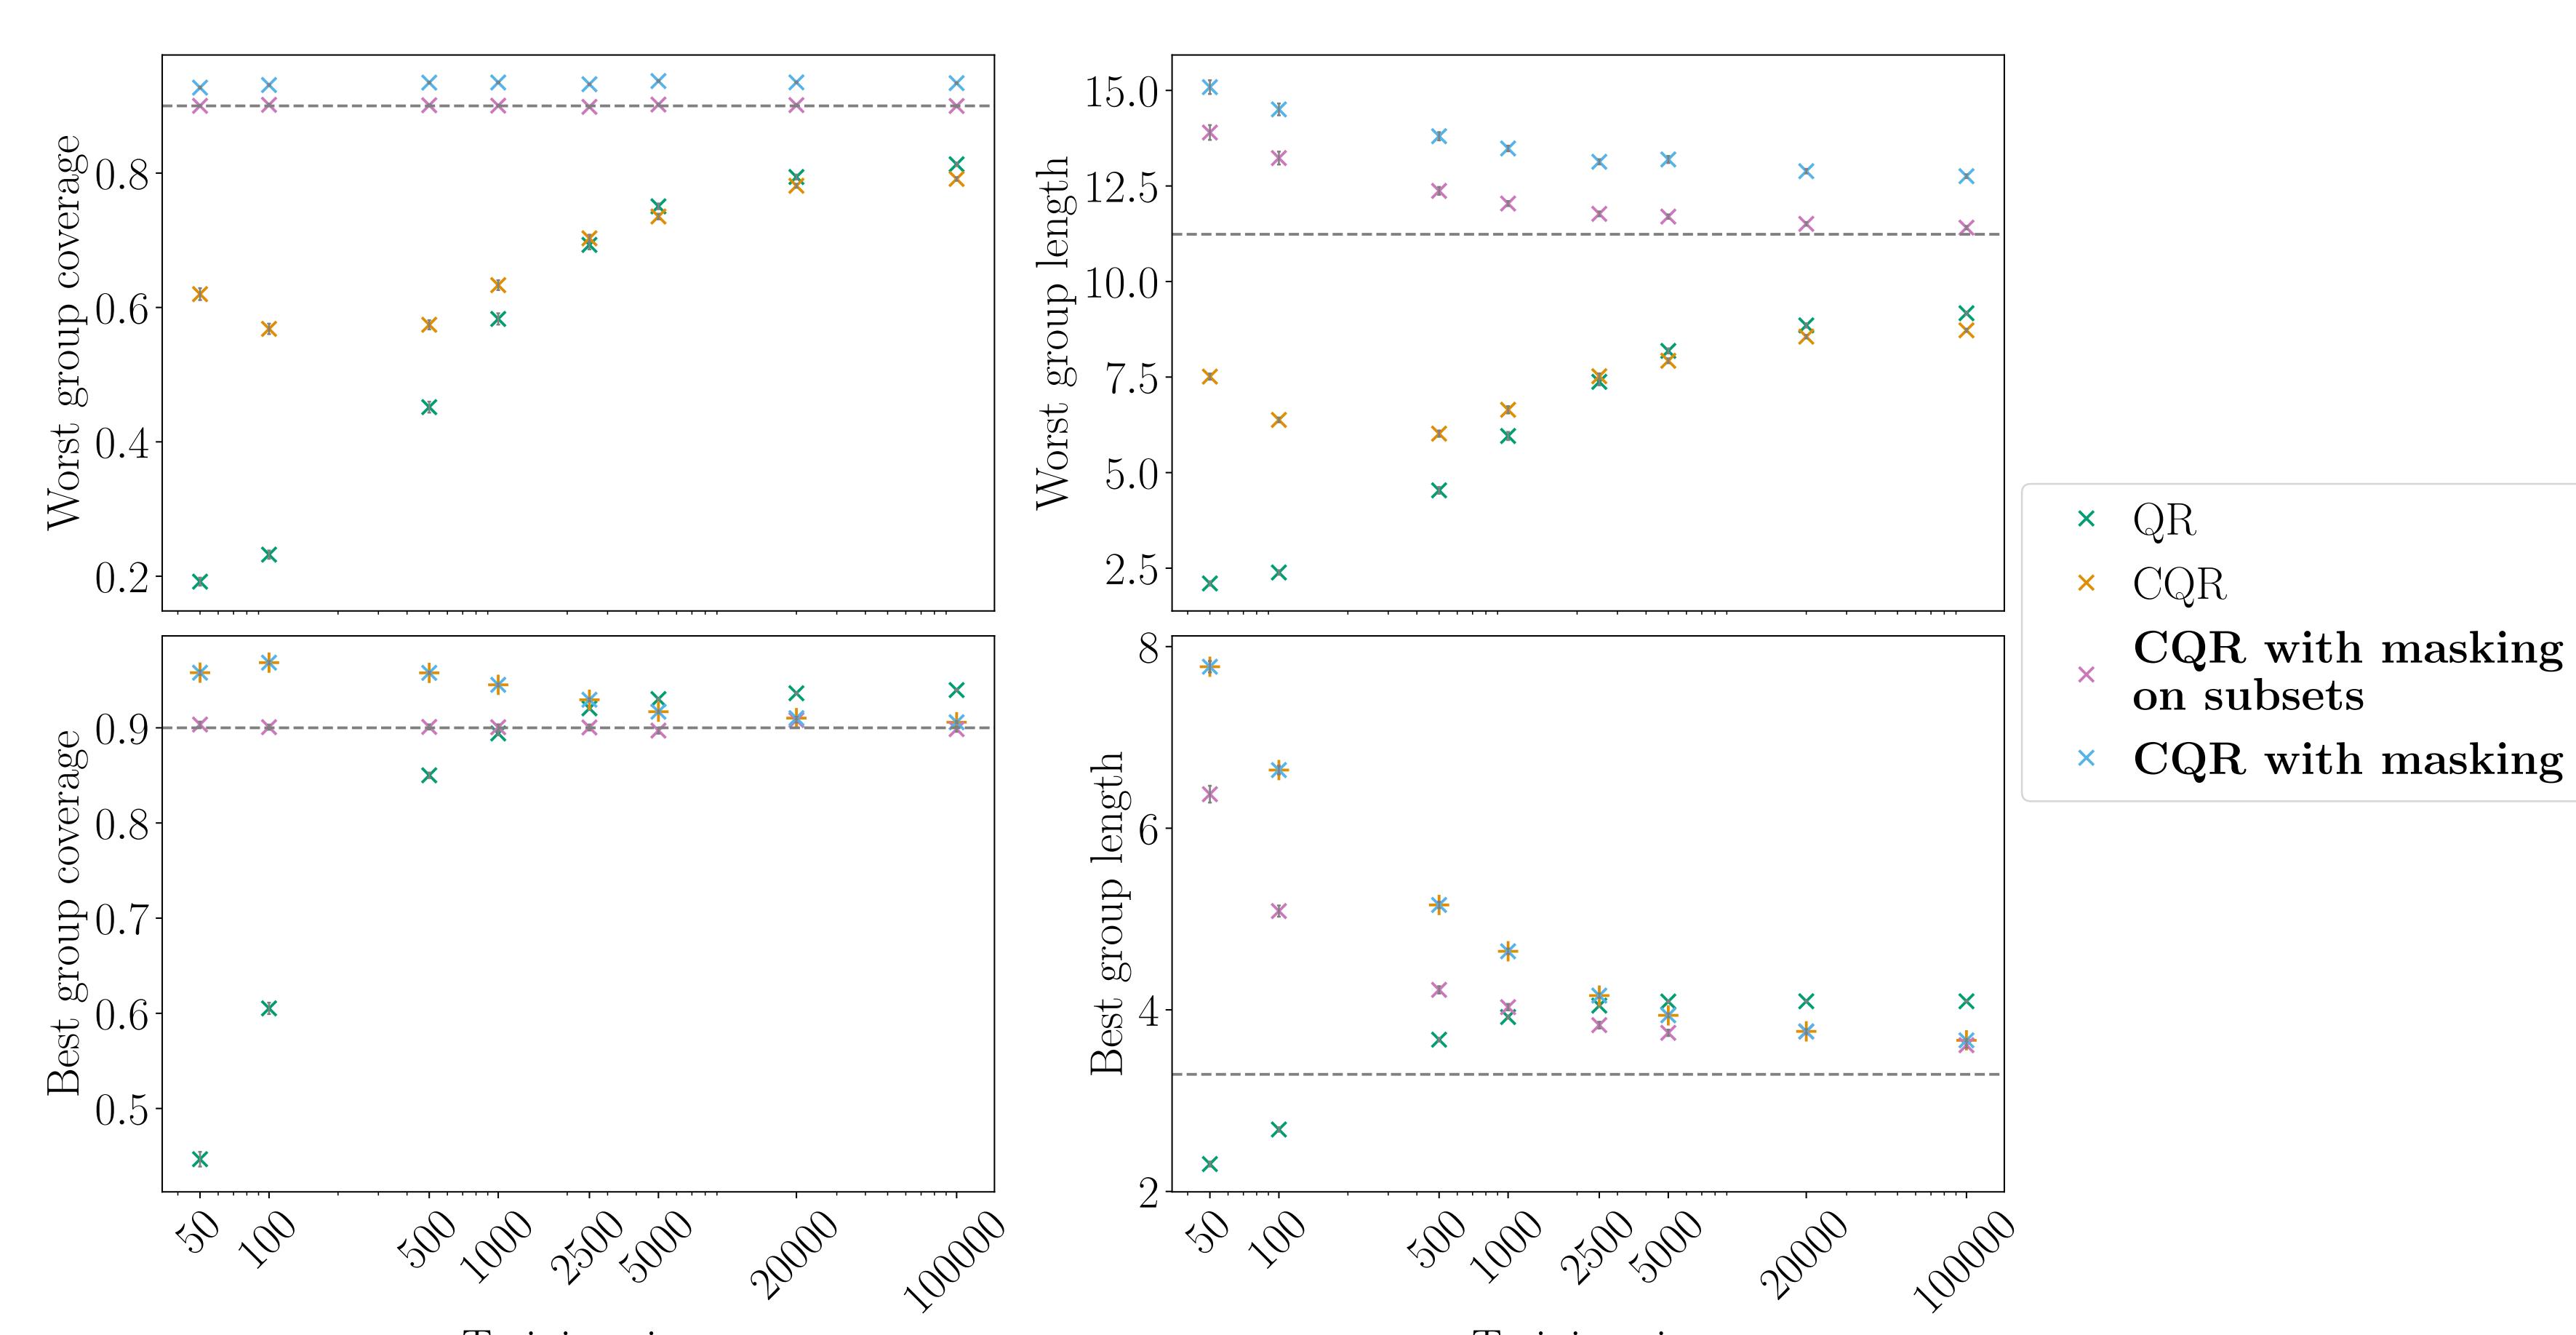
$\tilde{x}^{(1)}$			
$\tilde{x}^{(2)}$	-1	NA	NA
$\tilde{x}^{(3)}$	4	NA	NA
$\tilde{x}^{(4)}$	5	NA	NA

Calibration set used in CQR with masking on subsets

$$e^{(k)} = \max \left\{ \hat{q}_{\text{inf}}(\tilde{x}^{(k)}) - y^{(k)}, y^{(k)} - \hat{q}_{\text{sup}}(\tilde{x}^{(k)}) \right\}$$

Appropriate coverage conditionally to the missing patterns

On Gaussian linear data with $d = 10$, focus on **2 extreme missing patterns**: largest and smallest number of missing values.



Particular case: $X \sim \mathcal{N}(\mu, \Sigma)$, and M is MCAR. Then, $\mu_m \equiv \mu$ and $\Sigma_m \equiv \Sigma$.

Oracle intervals

Under the Gaussian linear model, for any $m \in \{0, 1\}^d$, the oracle length is given by:

$$\mathcal{L}_\alpha^*(m) = 2 \times q_{1-\alpha/2}^{N(0,1)} \times \sqrt{\beta_{\text{mis}(m)} \Sigma_{\text{mis}(m)|\text{obs}(m)} \beta_{\text{mis}(m)}^T + \sigma_\varepsilon^2},$$

with $\Sigma_{\text{mis}(m)|\text{obs}(m)} = \Sigma_{\text{mis}(m), \text{mis}(m)} - \Sigma_{\text{mis}(m), \text{obs}(m)} \Sigma_{\text{obs}(m), \text{obs}(m)}^{-1} \Sigma_{\text{obs}(m), \text{mis}(m)}$.

- The oracle intervals depend on the regression coefficients.
- Additional **heteroskedasticity** is generated by the missing values.
- The oracle intervals depend on the **mask** in a **non-linear** fashion.
→ even under MCAR data, it is useful to add the mask as feature.

- As the training size increases, **QR** and **CQR** improve conditional coverage.
- CQR with masking on subsets** is not over-conservative on the easiest group, but requires more calibration data than **CQR with masking**.
- As the training size increases, **CQR with masking on subsets** → oracle length.