



HAL
open science

Marginal and training-conditional guarantees in one-shot federated conformal prediction

Pierre Humbert, Batiste Le Bars, Aurélien Bellet, Sylvain Arlot

► **To cite this version:**

Pierre Humbert, Batiste Le Bars, Aurélien Bellet, Sylvain Arlot. Marginal and training-conditional guarantees in one-shot federated conformal prediction. 2024. hal-04579882

HAL Id: hal-04579882

<https://hal.science/hal-04579882>

Preprint submitted on 18 May 2024

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.



Distributed under a Creative Commons Attribution 4.0 International License

Marginal and training-conditional guarantees in one-shot federated conformal prediction

Pierre Humbert¹, Batiste Le Bars², Aurélien Bellet³ and Sylvain Arlot^{1,4}

¹ *Université Paris-Saclay, CNRS, Inria, Laboratoire de mathématiques d’Orsay, 91405, Orsay, France, e-mail: pierre.humbert@universite-paris-saclay.fr; sylvain.arlot@universite-paris-saclay.fr*

² *Inria Paris - Ecole Normale Supérieure, PSL Research University, e-mail: batiste.le-bars@inria.fr*

³ *PreMeDICaL project team, Inria, Univ Montpellier, France, e-mail: aurelien.bellet@inria.fr*

⁴ *Institut Universitaire de France (IUF)*

Abstract: We study conformal prediction in the one-shot federated learning setting. The main goal is to compute marginally and training-conditionally valid prediction sets, at the server-level, in only one round of communication between the agents and the server. Using the quantile-of-quantiles family of estimators and split conformal prediction, we introduce a collection of computationally-efficient and distribution-free algorithms that satisfy the aforementioned requirements. Our approaches come from theoretical results related to order statistics and the analysis of the Beta-Beta distribution. We also prove upper bounds on the coverage of all proposed algorithms when the nonconformity scores are almost surely distinct. For algorithms with training-conditional guarantees, these bounds are of the same order of magnitude as those of the centralized case. Remarkably, this implies that the one-shot federated learning setting entails no significant loss compared to the centralized case. Our experiments confirm that our algorithms return prediction sets with coverage and length similar to those obtained in a centralized setting.

Keywords and phrases: conformal prediction, one-shot federated learning, prediction set, tolerance region, uncertainty quantification.

1. Introduction

1.1. Problem statement and motivation

We consider the one-shot federated learning (FL) set-prediction problem, where a set of agents connected to a central server try to compute a valid prediction set in only one round of communication, and without sharing their raw data [25].

Formally, assume that a data set \mathcal{D} is distributed over $m \in \mathbb{N}^*$ agents connected to a central server. In addition, suppose that the server has an independent test point $(X, Y) \in \mathcal{X} \times \mathcal{Y}$ following the same distribution as the elements $(X_{i,j}, Y_{i,j})$ of \mathcal{D} —which are assumed to be identically distributed— but whose outcome Y is unobserved. The goal of the server is to compute a distribution-free *valid* prediction set for Y . One can aim at marginal validity, that is, for a given miscoverage level $\alpha \in (0, 1)$, constructing from \mathcal{D} a set $\hat{\mathcal{C}}(X)$ such that

$$\mathbb{P}\left(Y \in \hat{\mathcal{C}}(X)\right) \geq 1 - \alpha, \quad (1)$$

whatever the data distribution. In addition, the set $\hat{\mathcal{C}}(\cdot)$ must be computed in a single round of communication between the agents and the server to satisfy the one-shot constraint—a condition motivated by the fact that the number of communication rounds is often the main bottleneck in FL [27]. Note also that we want the set to be small

(for instance, with respect to the counting or the Lebesgue measure), a property called *efficiency*.

The probability in (1) is taken with respect to (X, Y) and the data set \mathcal{D} . However, in practice, we only have access to one particular data set. Another quantity of interest is therefore the training-conditional miscoverage rate defined by

$$\alpha(\mathcal{D}) := \mathbb{P}(Y \notin \hat{\mathcal{C}}(X) \mid \mathcal{D}), \quad (2)$$

where the probability is now only taken with respect to the test point (X, Y) . Notice that $\alpha(\mathcal{D})$ is a $\sigma(\mathcal{D})$ -measurable random variable and that the marginal guarantee (1) corresponds to a bound on the miscoverage rate (2) on average over all possible datasets. In other words, Eq. (1) is equivalent to a control on the expectation of the training-conditional miscoverage rate: $\mathbb{P}(Y \notin \hat{\mathcal{C}}(X)) = \mathbb{E}[\alpha(\mathcal{D})] \leq \alpha$. However, the random variable $\alpha(\mathcal{D})$ can have high variance, and it is important to also control its deviation from the desired upper-bound α with high probability. For any $\alpha, \beta \in (0, 1)$, we are therefore also interested in constructing prediction sets with training-conditional coverage guarantees, that is, of the form

$$\mathbb{P}(1 - \alpha(\mathcal{D}) \geq 1 - \alpha) \geq 1 - \beta. \quad (3)$$

Prediction sets satisfying (3) are also known in the statistical literature as (α, β) -tolerance regions or ‘‘Probably Approximately Correct’’ (PAC) predictive sets. This type of guarantee dates back to [63] and an overview can be found in [32]. More recent works on this subject include [8, 29, 47, 48, 61].

1.2. Related works

Our contribution takes place in the FL framework, a rather recent paradigm that allows training from decentralized data sets stored locally by multiple agents [27]. In this framework, the learning is made without exchanging raw data, making FL advantageous when data are highly sensitive and cannot be centralized for privacy or security reasons. So far, the design of FL algorithms has mainly focused on the learning step of statistical machine learning, the goal being to fit a (pointwise) predictor to decentralized data sets while minimizing the amount of communication [see e.g. 28, 37, 43]. However, quantifying the uncertainty in the prediction of these FL algorithms has not been widely studied yet.

Conformal prediction (CP) methods have become the state-of-the-art to construct marginally and conditionally valid distribution-free prediction sets [46, 54, 61, 62]. Unfortunately, one of the key steps of CP methods is the ordering of some computed scores, which is not possible in FL settings without sharing the full local data sets or performing many agent-server communication rounds. These methods are thus not well-suited to the constraints of FL in which agents process their data locally and only interact with a central server by sharing some aggregate statistics. Constructing a valid prediction set is even more difficult in the one-shot FL setting [14, 20, 36, 55, 65, 66] considered in this work, where only one round of communication between the agents and the server is allowed.

To our knowledge, [38] is the first paper considering the one-shot federated set-prediction problem with conformal prediction. Its idea is to locally calculate some quantiles of computed scores for each agent and to average them in the central server. Unfortunately, [38] does not prove that the corresponding prediction sets are valid and its method is non-robust, especially when the size of local data sets is small (see Appendix I.2 for more details). To address these issues, [25] has recently proposed a family of estimators called *quantile-of-quantiles*. The idea is that each agent sends to the server a local empirical quantile of its scores and the server aggregates them by computing a quantile of these

quantiles. Interestingly, these estimators can therefore be seen as a clever way of aggregating several quantiles calculated locally by m agents, who each locally use the (centralized) split CP method (see Section 2.1 for details). Last but not least, when $m = 1$ this approach exactly recovers split CP, emphasizing that methods based on quantile-of-quantiles estimators generalize centralized split CP to the FL setting. However, an important limitation of [25] is that in order to determine which order of quantile to select for marginal validity (1), the proposed methodology can be computationally intensive at the server level. Moreover, the training-conditional guarantee (3) given by [25] is obtained with a conservative procedure, leading to large prediction sets in practice. The present paper shows how to solve these two issues.

Outside the one-shot FL setting considered here, we can also mention [39] and [50], which focus on data-heterogeneous settings but require many communication rounds between the agents and the server. Finally, we can also mention recent works on federated evaluation of classifiers [12], federated quantile computation [3, 49], and on uncertainty quantification with Bayesian FL [15, 31] which, although related to our work, do not study CP and do not obtain formal coverage guarantees.

1.3. Contributions

In this work, we consider the quantile-of-quantiles family of estimators proposed in [25] and introduce several algorithms to find the appropriate order of quantiles so that the marginal condition (1) or the training-conditional condition (3) are satisfied. Each of these algorithms is computationally-efficient, distribution-free (depending only on the number of agents and the size of their local data sets) and specially tailored to satisfy the aforementioned conditions. Importantly, they come from novel theoretical results, which take their roots in the theory of order statistics. For clarity and simplicity of reading, all the contributions are first presented in the case where agents have the same number of data points $n \geq 1$ (Assumption 1 in Section 2.2).

For distribution-free marginal guarantees (1), in Section 3.1:

- We prove that when $\hat{\mathcal{C}}$ is obtained using our method, its probability of coverage $\mathbb{E}[1 - \alpha(\mathcal{D})] = \mathbb{P}(Y \in \hat{\mathcal{C}}(X))$ is lower bounded by the expectation of a random variable following a particular Beta-Beta distribution [11, 40] (when the scores are almost surely distinct, these two quantities are equal). We also derive a closed-form expression for this expectation (Theorem 5), improving the one obtained in [25] and leading to Algorithm 1.
- This closed-form expression remains difficult to compute for large values of n and m , leading to a quite computationally demanding algorithm. To tackle this problem, we show that the expectation of this Beta-Beta distribution is lower and upper bounded by the quantile function of a standard Beta distribution evaluated at particular values (Proposition 8). These bounds are tight and fast to compute, which makes them interesting for practical use: they lead to Algorithm 2, which is more computationally efficient.

In Section 3.2, we build the first (to the best of our knowledge) one-shot FL algorithms with distribution-free *training-conditional* guarantees (3). More precisely:

- We prove that $1 - \alpha(\mathcal{D})$ is stochastically larger (equal when the scores are a.s. distinct) than a Beta-Beta random variable (Theorem 10). We also show that the quantiles of the Beta-Beta are fast to compute, leading to an efficient algorithm that constructs training-conditionally valid prediction sets in one-shot FL (Algorithm 3).

- In order to obtain an even faster algorithm, we provide a tight bound on the cumulative distribution function (cdf) of the Beta-Beta distribution (Proposition 13) which allows the automatic selection of the empirical-quantiles order that the agents should send to the server (Algorithm 4).

Importantly, our results allow to trade-off the tightness of the bounds for computational efficiency. More generally, our contributions go beyond the setting of FL, in the sense that they investigate the open question of how several split CP estimators obtained over independent data sets should be aggregated to obtain valid prediction sets.

In Section 4, we give several upper-bounds on the probability of coverage of our prediction sets. We first derive an upper-bound of order $1 - \alpha + \mathcal{O}(m^{-1}n^{-1/2})$ for the marginal coverage of our methods from Section 3.1 (Algorithms 1–2). This result shows that our coverage is not too much above $1 - \alpha$ and strictly improves upon the work of [25], which did not provide such type of results. In the same vein, we also prove high-probability upper bounds on the training-conditional miscoverage rate obtained by the methods of Section 3.2 (Algorithms 3–4). Remarkably, these bounds are of order $1 - \alpha + \mathcal{O}((mn)^{-1/2})$, the same order of magnitude as those of the centralized case. Hence, up to constant factors, the one-shot federated learning setting does not incur any loss.

In Section 5, we extend our theoretical results and associated methods to the more complicated setting where the agents can have different data set sizes (Theorems 19–20 and Algorithms 5–6), improving again upon the results presented in [25].

Finally, in Section 6, we empirically evaluate the performance of our algorithms on standard CP benchmarks, validating that they produce prediction sets that are computationally efficient and close to those obtained when data are centralized.

2. Preliminaries

2.1. Split Conformal Prediction

Conformal Prediction (CP) is a framework to construct prediction sets satisfying (1) without relying on any distributional assumption on the data [62]. In a centralized setting, one of the most popular methods to perform CP is the split conformal method (split CP) [46]. Since the quantile-of-quantiles procedure studied in this paper generalizes split CP, let us detail here how it is defined and recall some of the key results of previous literature. Simple and full proofs of these results are provided in Appendix I.1.

Assume that we have access to a centralized dataset \mathcal{D} , that we split into a learning set \mathcal{D}^{lm} and a calibration set $\mathcal{D}^{cal} = (X_i, Y_i)_{1 \leq i \leq n_c}$, where $n_c \geq 1$ and the calibration data (X_i, Y_i) , $1 \leq i \leq n_c$, are i.i.d. and follow the same distribution as the independent test point (X, Y) .

First, a predictor \hat{f} is built from \mathcal{D}^{lm} only, and it is used to define a nonconformity score function $s = s_{\hat{f}} : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$, such that for every $(x, y) \in \mathcal{X} \times \mathcal{Y}$, $s_{\hat{f}}(x, y)$ measures how far the prediction $\hat{f}(x)$ is from the true output y . Whether we are in the regression or classification setting, many different score functions exist in the literature (see e.g. [4]). In regression, for instance, a common choice is the fitted absolute residual $s_{\hat{f}} : (x, y) \mapsto |y - \hat{f}(x)|$. In the sequel, we often write s instead of $s_{\hat{f}}$ for simplicity. Furthermore, note that split CP does not assume a particular choice of score function, so throughout the paper, we keep the function s abstract.

Second, we calculate the values of $s_{\hat{f}}$ taken on the calibration set \mathcal{D}^{cal} , called the nonconformity scores $S_i := s_{\hat{f}}(X_i, Y_i)$, $i = 1, \dots, n_c$.

Third, we compute the r -th smallest nonconformity score $S_{(r)} := \widehat{Q}_{(r)}(\mathcal{S}_{n_c}^{cal})$ for some $r \in \llbracket n_c \rrbracket := \{1, \dots, n_c\}$, where $\mathcal{S}_{n_c}^{cal} := (S_1, \dots, S_{n_c})$ and $\widehat{Q}_{(\cdot)}$ is the sample quantile function defined by

$$\forall r, N \geq 1, \forall \mathcal{S}' \in \mathbb{R}^N, \quad \widehat{Q}_{(r)}(\mathcal{S}') := \begin{cases} \mathcal{S}'_{(r)} & \text{if } r \leq N \\ +\infty & \text{otherwise,} \end{cases} \quad (4)$$

with $\mathcal{S}'_{(1)} \leq \dots \leq \mathcal{S}'_{(N)}$ the ordered values of $\mathcal{S}' = (\mathcal{S}'_1, \dots, \mathcal{S}'_N)$. Finally, the split CP prediction set is defined, for any $x \in \mathcal{X}$, by

$$\widehat{\mathcal{C}}_r(x) := \left\{ y \in \mathcal{Y} : s_{\widehat{f}}(x, y) \leq S_{(r)} \right\}. \quad (5)$$

Following Eq. (2) in Section 1, we define the training-conditional miscoverage rate of $\widehat{\mathcal{C}}_r$ by

$$\alpha_r(\mathcal{D}) := \mathbb{P}(Y \notin \widehat{\mathcal{C}}_r(X) \mid \mathcal{D}),$$

where (X, Y) is a test point, independent of \mathcal{D} and with the same distribution as the (X_i, Y_i) . We then have

$$1 - \alpha_r(\mathcal{D}) = \mathbb{P}(s_{\widehat{f}}(X, Y) \leq S_{(r)}) = F_S(S_{(r)})$$

where F_S is the common cdf of the scores S_i . It is well known that for every $r \in \llbracket n_c \rrbracket$,

$$\mathbb{P}(Y \in \widehat{\mathcal{C}}_r(X)) = \mathbb{E}[1 - \alpha_r(\mathcal{D})] \geq \frac{r}{n_c + 1},$$

with equality if the scores S_i are almost surely distinct [35, 62]. As a consequence, for any $\alpha \in (0, 1)$, if $(1 - \alpha)(n_c + 1) \leq n_c$, taking $r = \lceil (1 - \alpha)(n_c + 1) \rceil$ in Eq. (5) yields a prediction set satisfying Eq. (1), and if the S_i are almost surely distinct, its expected miscoverage rate is

$$\mathbb{E}[1 - \alpha_{\lceil (1 - \alpha)(n_c + 1) \rceil}(\mathcal{D})] = \frac{\lceil (1 - \alpha)(n_c + 1) \rceil}{n_c + 1} \leq 1 - \alpha + \frac{1}{n_c + 1}. \quad (6)$$

Regarding training-conditional guarantees, a straightforward consequence of [61, Proposition 2b] is that

$$\forall r \geq 1, \forall \beta \in (0, 1), \quad \mathbb{P}(1 - \alpha_r(\mathcal{D}) \geq F_{U(r; n_c)}^{-1}(\beta)) \geq 1 - \beta, \quad (7)$$

where $F_{U(r; n_c)}^{-1}$ denotes the quantile function of the $\text{Beta}(r, n_c - r + 1)$ distribution. In other words, $1 - \alpha_r(\mathcal{D})$ is *stochastically larger* than the $\text{Beta}(r, n_c - r + 1)$ distribution. Furthermore, Eq. (7) becomes an equality if the scores S_i are almost surely distinct — that is, in such a case, $\alpha_r(\mathcal{D})$ exactly follows a $\text{Beta}(r, n_c - r + 1)$ distribution. Finally, taking $r \in \llbracket n_c \rrbracket$ such that $F_{U(r; n_c)}^{-1}(\beta) \geq 1 - \alpha$, Eq. (7) implies that the split CP prediction set $\widehat{\mathcal{C}}_r$ satisfies Eq. (3), that is, $\widehat{\mathcal{C}}_r$ is a (α, β) -tolerance region —see also Eq. (99) in Appendix I.1.2.

Remark 1. When n_c tends to infinity, the optimal asymptotically training-conditionally valid r —given by Eq. (102) in Appendix I.1.2— yields a prediction set $\widehat{\mathcal{C}}_r$ with coverage between $1 - \alpha$ and $1 - \alpha + \mathcal{O}(1/\sqrt{n_c})$ with high probability —see Eq. (103) in Appendix I.1.2 for a precise statement.

The problem of split CP in a federated setting is that computing the quantile $S_{(r)}$ requires in general several (and often many) communications between the central server and the agents, hence it cannot be used in one-shot FL. Therefore, we consider in this paper another family of procedures (quantile-of-quantiles estimators), that we define in the next section.

Note that, although the first CP methods were the *split* and the related *full* methods [46, 62], many extensions based upon them and with similar guarantees have been proposed in the literature. Their principal novelty lies in a clever choice of the non-conformity score function s . In regression, [35] presents a method called locally weighted CP and provides theoretical insights for conformal inference. More recently, [54] has developed a variant of the split CP called Conformal Quantile Regression (CQR). Other recent alternatives have been proposed [19, 21, 22, 29, 45, 56]. We refer to [62], [4], and [17] for in-depth presentations of CP and to [41] for a curated list of papers related to CP.

2.2. Federated conformal prediction with the quantile-of-quantiles

We now present the quantile-of-quantiles family of estimators, first introduced in [25], and how it can be used to obtain valid prediction sets in a one-shot FL setting, that is, in a setting where only one round of communication between the agents and the server is allowed [20, 66] and where only aggregated statistics computed locally by the agents can be sent to the server. From now on, we assume that the decentralized data set \mathcal{D} is divided into a learning set \mathcal{D}^{lrn} and a calibration set $\mathcal{D}^{cal} = (X_{i,j}, Y_{i,j})_{1 \leq j \leq m, 1 \leq i \leq n_j}$ where $n_j \geq 1$ for every $j \in \llbracket m \rrbracket$. Among calibration data, agent $j \in \llbracket m \rrbracket$ has only access to $(X_{i,j}, Y_{i,j})_{1 \leq i \leq n_j}$. We assume that \mathcal{D}^{lrn} is independent from \mathcal{D}^{cal} , and that the calibration data $(X_{i,j}, Y_{i,j})$, $j \in \llbracket m \rrbracket$, $i \in \llbracket n_j \rrbracket$ are i.i.d. with the same distribution as the test point (X, Y) (which is independent from \mathcal{D}).

We assume that a (pointwise) predictor \hat{f} is learned on \mathcal{D}^{lrn} only, using for instance standard FL algorithms such as FedAvg [43]. Therefore, \hat{f} is independent from \mathcal{D}^{cal} . As in the centralized setting, \hat{f} is used to define a nonconformity score function $s = s_{\hat{f}} : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$ such that for every $(x, y) \in \mathcal{X} \times \mathcal{Y}$, $s_{\hat{f}}(x, y)$ measures how far the prediction $\hat{f}(x)$ is from the true output y . In the sequel, we only focus on the calibration of the prediction set and not on the learning part.

Remark 2. Note that \hat{f} does not have to be a point-wise predictor, that is, a function $\mathcal{X} \rightarrow \mathcal{Y}$. For instance, like in CQR [54], we can rely on the use of a score function s depending on $\hat{f} = (\hat{f}_-, \hat{f}_+)$ a pair of quantile functions $\mathcal{X} \rightarrow \mathcal{Y}$.

Remark 3. In the following, all probabilistic statements are valid conditionally to \mathcal{D}^{lrn} . This amounts to acting as if \hat{f} were deterministic, since \mathcal{D}^{lrn} appears only through \hat{f} and \mathcal{D}^{cal} is independent from \mathcal{D}^{lrn} .

For simplicity, from now on and until the end of Section 4, we also make the following assumption.

Assumption 1. Each agent $j \in \llbracket m \rrbracket$ has exactly $n_j = n \geq 1$ calibration data points.

Under Assumption 1, the calibration data set size is equal to nm . We refer to Section 5 for the more general case where agents have data sets of calibration of different sizes $(n_j)_{1 \leq j \leq m}$.

For calibration, the first step is to ask each agent $j \in \llbracket m \rrbracket$ to compute its n i.i.d. local calibration scores $\mathcal{S}_j := (S_{1,j}, \dots, S_{n,j})$, where $S_{i,j} = s_{\hat{f}}(X_{i,j}, Y_{i,j})$ is the score associated

to the i -th calibration data point of agent j . Then, the server and the agents jointly compute the quantile-of-quantiles (QQ) estimator, defined as follows.

Definition 4. (Quantile-of-quantiles [25]) For any $(\ell, k) \in \llbracket n \rrbracket \times \llbracket m \rrbracket$, the QQ estimator of order (ℓ, k) calculated on the sets of scores $(\mathcal{S}_j)_{1 \leq j \leq m}$ is

$$S_{(\ell, k)} := \widehat{Q}_{(k)} \left(\widehat{Q}_{(\ell)}(\mathcal{S}_1), \dots, \widehat{Q}_{(\ell)}(\mathcal{S}_m) \right), \quad (8)$$

where $\widehat{Q}_{(\cdot)}(\cdot)$ is the sample quantile function defined in Eq. (4).

In other words, each agent sends to the server its ℓ -th smallest local score, denoted by $\widehat{Q}_{(\ell)}(\mathcal{S}_j)$, and the server then computes the k -th smallest value of these scores, denoted by $S_{(\ell, k)}$. This strategy requires a single round of communication and thus fits the constraints of one-shot FL. Finally, for any $x \in \mathcal{X}$, we define the prediction set

$$\widehat{\mathcal{C}}_{\ell, k}(x) = \left\{ y \in \mathcal{Y} : s(x, y) \leq S_{(\ell, k)} \right\}, \quad (9)$$

similarly to the centralized split CP method —see Eq. (5)—, but with $S_{(r)}$ replaced by $S_{(\ell, k)}$.

The decentralized QQ prediction set can be seen as a generalization of (centralized) split CP since they coincide when $m = 1$. The crucial remaining component is a computationally efficient approach to identify a pair (ℓ, k) so that $\widehat{\mathcal{C}}_{\ell, k}$ defined in Eq. (9) satisfies the marginal condition (1) or the training-conditional condition (3) while being as small as possible. We investigate these points in the next sections.

2.3. Prediction set performance measure

A natural way to evaluate the performance of a valid prediction set $x \mapsto \widehat{\mathcal{C}}(x)$ is to measure its size $\mu(\widehat{\mathcal{C}}(x))$, where μ is some measure on \mathcal{Y} , for instance, the counting measure when \mathcal{Y} is finite, or the Lebesgue measure when $\mathcal{Y} \subset \mathbb{R}^p$ for some $p \geq 1$. This size should be minimized, either at a given $x \in \mathcal{X}$ or on average over $x = X$. For a prediction set of the form $\widehat{\mathcal{C}}_{\ell, k}$, as defined by Eq. (9), its size depends on μ , on the score function s , on the predictor \widehat{f} , and on the pair (ℓ, k) . In order to build general-purpose algorithms for choosing (ℓ, k) , a natural strategy is thus to select, among all the pairs such that $\widehat{\mathcal{C}}_{\ell, k}$ is marginally or conditionally valid, the one which also minimizes the size $\mu(\widehat{\mathcal{C}}_{\ell, k}(x))$. By Eq. (9), this size is a nondecreasing function of $S_{(\ell, k)}$. Furthermore, we know that

$$S_{(\ell, k)} \stackrel{a.s.}{=} F_S^{-1} \circ F_S(S_{(\ell, k)}) = F_S^{-1}(1 - \alpha_{\ell, k}(\mathcal{D})), \quad (10)$$

$$\text{where } 1 - \alpha_{\ell, k}(\mathcal{D}) := \mathbb{P}(Y \in \widehat{\mathcal{C}}_{\ell, k}(X) | \mathcal{D}) = F_S(S_{(\ell, k)})$$

is the coverage of $\widehat{\mathcal{C}}_{\ell, k}$, F_S is the cdf of the scores $S_{i, j}$, and F_S^{-1} its generalized inverse —see Eq. (35) in Appendix B.1. Hence, the size of $\widehat{\mathcal{C}}_{\ell, k}(x)$ is also a nondecreasing function of the coverage.

Therefore, the quantiles or the expectation of the coverage are good ways to measure the performance of prediction sets of the form $\widehat{\mathcal{C}}_{\ell, k}$. In the following, we will use these quantities as criteria (to be minimized) for choosing among pairs (ℓ, k) such that $\widehat{\mathcal{C}}_{\ell, k}$ is (marginally or training-conditionally) valid. More detailed arguments about this strategy can be found in Appendix A.

3. Choice of the quantiles for coverage guarantees

We now present theoretical results together with one-shot FL algorithms which return prediction sets with marginal (Section 3.1) or training-conditional (Section 3.2) guarantees.

3.1. Marginal guarantees

In this section, we present two strategies for choosing a pair (ℓ, k) ensuring that $\widehat{\mathcal{C}}_{\ell, k}(x)$, defined by Eq. (9), is a marginally valid prediction set, that is, satisfies (1) for a given $\alpha \in (0, 1)$. Following Section 2.3, in order to get the best possible prediction set, our strategy is to take (ℓ, k) such that the marginal coverage

$$\mathbb{P}(Y \in \widehat{\mathcal{C}}_{\ell, k}(X)) = \mathbb{P}(s_{\widehat{f}}(X, Y) \leq S_{(\ell, k)})$$

is above $1 - \alpha$ while being as small as possible.

Our first algorithm is based on the following theorem, which simplifies the formula given in [25, Theorem 3.2].

Theorem 5. *In the setting of Section 2.2, with Assumption 1, for any $(\ell, k) \in \llbracket n \rrbracket \times \llbracket m \rrbracket$, the set $\widehat{\mathcal{C}}_{\ell, k}$ defined by Eq. (9) satisfies*

$$\mathbb{P}(Y \in \widehat{\mathcal{C}}_{\ell, k}(X)) \geq M_{\ell, k} \tag{11}$$

where

$$M_{\ell, k} := \frac{k \binom{m}{k} \sum_{i_1=\ell}^n \cdots \sum_{i_{k-1}=\ell}^n \sum_{i_{k+1}=0}^n \cdots \sum_{i_m=0}^n \frac{\binom{n}{i_1} \cdots \binom{n}{i_{k-1}} \binom{n}{i_{k+1}} \cdots \binom{n}{i_m}}{\binom{mn}{i_1 + \cdots + i_{k-1} + \ell + i_{k+1} + \cdots + i_m}}}{(mn+1)B(\ell, n-\ell+1)}$$

and

$$B : (a, b) \in (0, +\infty)^2 \mapsto \int_0^1 t^{a-1} (1-t)^{b-1} dt$$

denotes the Beta function [59]. Moreover, when the associated scores $(S_{i, j})_{1 \leq j \leq m, 1 \leq i \leq n}$ and $S := s(X, Y)$ are almost surely distinct, Eq. (11) is an equality.

Theorem 5 is proved in Appendix C.1. It shows that we can lower bound the probability of coverage of the quantile-of-quantiles prediction set by a distribution-free quantity $M_{\ell, k}$, which depends only on m, n, ℓ and k . Furthermore, the lower bound is sharp as it becomes an equality when the scores have a continuous cdf. This is for instance the case with the fitted absolute residual if the noise distribution given X is almost surely atomless.

Theorem 5 suggests the following algorithm for the selection of ℓ and k , which we call QQM (QQ stands for *Quantile-of-Quantiles* and M for *Marginal*).¹

Algorithm 1 (QQM). *Given $\alpha \in (0, 1)$,*

$$\text{compute } (\ell^*, k^*) = \underset{(\ell, k) \in \llbracket n \rrbracket \times \llbracket m \rrbracket}{\operatorname{argmin}} \{M_{\ell, k} : M_{\ell, k} \geq 1 - \alpha\},$$

$$\text{and output } \widehat{\mathcal{C}}_{\ell^*, k^*}(x) = \{y \in \mathcal{Y} : s(x, y) \leq S_{(\ell^*, k^*)}\}.$$

¹QQM is a slight modification of the algorithm FedCP-QQ proposed in [25]. More specifically, Eq. (11) is simpler and easier to compute than the corresponding formula for $M_{\ell, k}$ in [25].

Remark 6. By convention, when the argmin defining (ℓ^*, k^*) in Algorithm 1 is empty, we define $\widehat{\mathcal{C}}_{\ell^*, k^*}(x) = \mathcal{Y}$. A similar convention is used in all of our algorithms.

The minimization step of Algorithm 1 comes from the fact that all pairs (ℓ, k) such that $M_{\ell, k} \geq 1 - \alpha$ ensure the marginal coverage, by Theorem 5, but we need to select one specific pair. For the prediction set to be as small as possible, following Section 2.3, we should minimize the coverage. In Algorithm 1, we minimize $M_{\ell, k}$ since it is equal to the expected coverage when the scores cdf is continuous.

By Theorem 5, the set $\widehat{\mathcal{C}}_{\ell^*, k^*}$ is marginally valid. It is also nontrivial when mn is large enough, according to the following lemma.

Lemma 7. The argmin defining (ℓ^*, k^*) in Algorithm 1 is non-empty —hence $\widehat{\mathcal{C}}_{\ell^*, k^*}$ is nontrivial— if and only if $mn \geq \alpha^{-1} - 1$.

This lemma is proved in Appendix C.2. Remarkably, the necessary and sufficient condition in Lemma 7 depends on the calibration samples sizes only through the total number mn of calibration data points, hence it is the same for the centralized case (with one agent holding mn calibration data points) and for the one-shot FL case (with $m \geq 2$ agents, each having access to n calibration data points).

A critical limitation of QQM is that computing (ℓ^*, k^*) , and even a single $M_{\ell, k}$, is costly when $m \times n$ is large, preventing the approach to scale to a large number of data points or agents. For example, based on techniques from [34], [25] describes an algorithm that can be used to compute a single $M_{\ell, k}$ with a worst-case complexity of $\mathcal{O}(m^4 n \log(n))$. Although this complexity could be slightly improved using more advanced ideas from [34], the overall complexity of Algorithm 1 would still remain too high for very large data sets (e.g., when the number of agents m is large). Our second strategy enables us to find a valid pair (ℓ, k) much faster. It is based on sharp upper and lower bounds over $M_{\ell, k}$ that can be computed efficiently.

Proposition 8. Let $n, m \geq 1$ be two integers, $(\ell, k) \in \llbracket n \rrbracket \times \llbracket m \rrbracket$, and $M_{\ell, k}$ be defined in Theorem 5. Then,

$$F_{U(\ell; n)}^{-1} \left(\frac{k - 1/2}{m + 1/2} \right) < M_{\ell, k} < F_{U(\ell; n)}^{-1} \left(\frac{k}{m + 1/2} \right), \quad (12)$$

where $F_{U(\ell; n)}^{-1}$ is the quantile function of the Beta($\ell, n - \ell + 1$) distribution [59].

Proposition 8 is proved in Appendix C.3. By combining Proposition 8 and Theorem 5, if a pair (ℓ, k) is such that the left-hand side of Eq. (12) is greater or equal to $1 - \alpha$, then the associated prediction set $\widehat{\mathcal{C}}_{\ell, k}$ is marginally valid. As a consequence, for any $\ell \in \llbracket n \rrbracket$ and $\alpha \in (0, 1)$, if we set

$$k = \tilde{k}_{m, n}(\ell, \alpha) := \lceil (m + 1/2) \cdot F_{U(\ell; n)}(1 - \alpha) + 1/2 \rceil, \quad (13)$$

then $1 - \alpha \leq F_{U(\ell; n)}^{-1} \left(\frac{\tilde{k}_{m, n}(\ell, \alpha) - 1/2}{m + 1/2} \right) \leq M_{\ell, \tilde{k}_{m, n}(\ell, \alpha)} \leq \mathbb{P} \left(Y \in \widehat{\mathcal{C}}_{\ell, \tilde{k}_{m, n}(\ell, \alpha)}(X) \right)$,

provided that $\tilde{k}_{m, n}(\ell, \alpha) \in \llbracket m \rrbracket$. Therefore, choosing the associated pair $(\ell, \tilde{k}_{m, n}(\ell, \alpha))$ leads to a marginally-valid prediction set. Following our idea to minimize $M_{\ell, k}$ among marginally-valid pairs (ℓ, k) in Algorithm 1, it is here natural to choose ℓ by minimizing the upper-bound $F_{U(\ell; n)}^{-1} \left(\frac{\tilde{k}_{m, n}(\ell, \alpha)}{m + 1/2} \right)$ provided by Proposition 8, which leads to Algorithm 2 below.

Algorithm 2 (QQM-Fast). Given $\alpha \in (0, 1)$,

$$\begin{aligned} \text{compute } \tilde{\ell} &= \underset{\ell \in \llbracket n \rrbracket \text{ s.t. } \tilde{k}_{m,n}(\ell, \alpha) \in \llbracket m \rrbracket}{\operatorname{argmin}} \left\{ F_{U(\ell;n)}^{-1} \left(\frac{\tilde{k}_{m,n}(\ell, \alpha)}{m + 1/2} \right) \right\} \quad \text{and } \tilde{k}_{m,n}(\tilde{\ell}, \alpha), \\ &\text{where } \tilde{k}_{m,n}(\ell, \alpha) \text{ is defined by Eq. (13),} \\ \text{and output } \widehat{\mathcal{C}}_{\tilde{\ell}, \tilde{k}_{m,n}(\tilde{\ell}, \alpha)}(x) &= \left\{ y \in \mathcal{Y} : s(x, y) \leq S_{(\tilde{\ell}, \tilde{k}_{m,n}(\tilde{\ell}, \alpha))} \right\}. \end{aligned}$$

By construction, following the arguments detailed after Proposition 8, the set $\widehat{\mathcal{C}}_{\tilde{\ell}, \tilde{k}_{m,n}(\tilde{\ell}, \alpha)}$ is marginally valid. It is also nontrivial when mn is large enough, according to the next lemma.

Lemma 9. *The argmin defining $\tilde{\ell}$ in Algorithm 2 is non-empty —hence $\widehat{\mathcal{C}}_{\tilde{\ell}, \tilde{k}_{m,n}(\tilde{\ell}, \alpha)}$ is nontrivial— if and only if*

$$(1 - \alpha)^n \leq \frac{m - 1/2}{m + 1/2}. \quad (14)$$

In particular, condition (14) holds true if $n(m - 1/2) \geq \alpha^{-1} - 1$.

Lemma 9 is proved in Appendix C.4. In addition, QQM-Fast is computationally efficient. Indeed, the functions $F_{U(\ell;n)}$ and $F_{U(\ell;n)}^{-1}$ are fast to evaluate even when n is large (it takes a few milliseconds for $n = 10^6$). Furthermore, the minimization step in Algorithm 2 requires at most n evaluations, and the complexity of Algorithm 2 is almost not impacted by the value of m . For example, on a standard personal machine (Intel i5 with 4 CPU at 2.50GHz), with $m = n = 10^6$, QQM-Fast takes a few seconds to return the valid pair $(\tilde{\ell}, \tilde{k}_{m,n}(\tilde{\ell}, \alpha))$ using the SciPy implementation [60] of the cdf $F_{U(\ell;n)}$ and the quantile function $F_{U(\ell;n)}^{-1}$ of the Beta distribution, whereas QQM can take several hours.

While the two methods presented above ensure that the marginal coverage at level $1 - \alpha$ is satisfied whatever the data distribution, one may wonder how much above $1 - \alpha$ it can be. This question is answered in detail in Section 4.1.

3.2. Training-conditional guarantees

In this section, we present two algorithms for choosing (ℓ, k) such that the quantile-of-quantiles prediction set $\widehat{\mathcal{C}}_{\ell,k}(x)$ defined by Eq. (9) is a distribution-free training-conditionally valid prediction set, that is, satisfies Eq. (3). In other words, the goal is to select ℓ and k such that the miscoverage random variable

$$\alpha_{\ell,k}(\mathcal{D}) = \mathbb{P}(Y \notin \widehat{\mathcal{C}}_{\ell,k}(X) \mid \mathcal{D}) \quad (15)$$

is smaller than $\alpha \in (0, 1)$ with probability at least $1 - \beta \in (0, 1)$, that is,

$$\mathbb{P}\left(\mathbb{P}(Y \in \widehat{\mathcal{C}}_{\ell,k}(X) \mid \mathcal{D}) \geq 1 - \alpha\right) = \mathbb{P}(1 - \alpha_{\ell,k}(\mathcal{D}) \geq 1 - \alpha) \geq 1 - \beta. \quad (16)$$

Our first algorithm is based on the following theorem.

Theorem 10. *In the setting of Section 2.2, with Assumption 1, for any $(\ell, k) \in \llbracket n \rrbracket \times \llbracket m \rrbracket$ and any $\beta \in (0, 1)$, the miscoverage random variable $\alpha_{\ell,k}(\mathcal{D})$ defined by Eq. (15) satisfies*

$$\mathbb{P}(1 - \alpha_{\ell,k}(\mathcal{D}) \geq F_{U(\ell;n,k;m)}^{-1}(\beta)) \geq 1 - \beta, \quad (17)$$

where $F_{U(\ell;n,k;m)}^{-1} := F_{U(\ell;n)}^{-1} \circ F_{U(k;m)}^{-1}$ and for every $1 \leq r \leq N$, $F_{U(r;N)}^{-1}$ is the quantile function of the Beta($r, N - r + 1$) distribution.

Moreover, when the associated scores $(S_{i,j})_{1 \leq i \leq n, 1 \leq j \leq m}$ and $S := s(X, Y)$ are almost surely distinct, Eq. (17) is an equality and, for any $\beta' \in (0, 1)$ such that $\beta \leq 1 - \beta'$,

$$\mathbb{P}(F_{U_{(\ell:n, k:m)}}^{-1}(\beta) \leq 1 - \alpha_{\ell, k}(\mathcal{D}) \leq F_{U_{(\ell:n, k:m)}}^{-1}(1 - \beta')) = 1 - \beta - \beta'. \quad (18)$$

The interval $[F_{U_{(\ell:n, k:m)}}^{-1}(\beta), F_{U_{(\ell:n, k:m)}}^{-1}(1 - \beta')]$ is thus a two-sided fluctuation interval for the coverage random variable $1 - \alpha_{\ell, k}(\mathcal{D})$.

Theorem 10 is proved in Appendix D.1. Eq. (17) is sharp since it becomes an equality when the scores are a.s. distinct. It can be seen as a generalization of Eq. (7) to the case of the quantile-of-quantiles estimator. Theorem 10 is based on the fact that if $U_{(\ell:n, k:m)}$ is a random variable with cdf $F_{U_{(\ell:n, k:m)}}$, then $1 - \alpha_{\ell, k}(\mathcal{D})$ stochastically dominates $U_{(\ell:n, k:m)}$ in general, and they have the same distribution when the scores are a.s. distinct by Lemma 27 in Appendix B.2 (which also implies that Theorem 5 holds true with $M_{\ell, k} = \mathbb{E}[U_{(\ell:n, k:m)}]$).

Remark 11. The distribution associated to $F_{U_{(\ell:n, k:m)}}$ is a particular case of the Beta-Beta distribution [11, 40]. It can also be seen as the cdf of the k -th order statistics of a sample of m independent $\text{Beta}(\ell, n - \ell + 1)$ random variables [9, 26].

Theorem 10 suggests the following algorithm to select ℓ and k for training-conditional validity, which we call QQC (QQ stands for *Quantile-of-Quantiles* and C for *Conditional*).

Algorithm 3 (QQC). Given $\alpha \in (0, 1)$ and $\beta \in (0, 1)$,

$$\text{compute } (\ell_c^*, k_c^*) = \underset{(\ell, k) \in \llbracket n \rrbracket \times \llbracket m \rrbracket}{\text{argmin}} \left\{ F_{U_{(\ell:n, k:m)}}^{-1}(1 - \beta) : F_{U_{(\ell:n, k:m)}}^{-1}(\beta) \geq 1 - \alpha \right\} \quad (19)$$

$$\text{and output } \widehat{\mathcal{C}}_{\ell_c^*, k_c^*}(x) = \{y \in \mathcal{Y} : s(x, y) \leq S_{(\ell_c^*, k_c^*)}\}.$$

By construction and Theorem 10, the set $\widehat{\mathcal{C}}_{\ell_c^*, k_c^*}$ satisfies the training-conditional condition (16). It is also nontrivial when mn is large enough, according to the following lemma.

Lemma 12. The argmin defining (ℓ_c^*, k_c^*) in Algorithm 3 is non-empty —hence $\widehat{\mathcal{C}}_{\ell_c^*, k_c^*}$ is nontrivial— if and only if

$$mn \geq \frac{\log(\beta)}{\log(1 - \alpha)}. \quad (20)$$

Lemma 12 is proved in Appendix D.2. Furthermore, as long as m and n are not too large, Algorithm 3 is computationally efficient since $F_{U_{(\ell:n)}}^{-1}$ and $F_{U_{(k:m)}}^{-1}$ can be computed quickly (see the comments below Algorithm 2 in Section 3.1).

Note that any pair (ℓ, k) satisfying the condition $F_{U_{(\ell:n, k:m)}}^{-1}(\beta) \geq 1 - \alpha$ leads to a conditionally valid prediction set. As a selection criterion among these pairs (ℓ, k) , following Section 2.3, we minimize $F_{U_{(\ell:n, k:m)}}^{-1}(1 - \beta)$ which, according to Eq. (18) in Theorem 10, is equal to the $(1 - \beta)$ -quantile of the distribution of $1 - \alpha_{\ell, k}(\mathcal{D})$ when the scores are a.s. distinct.

QQC may require to evaluate $F_{U_{(\ell:n, k:m)}}^{-1}$ for many pairs $(\ell, k) \in \llbracket n \rrbracket \times \llbracket m \rrbracket$, which can be too costly when n and m are large. In the sequel, we propose a simpler sufficient condition for Eq. (16) to be satisfied, leading to a faster algorithm. It is based on the following lower bound on the quantile function $F_{U_{(\ell:n, k:m)}}^{-1}$.

Proposition 13. Let $n, m \geq 1$ be two integers, $k \in \llbracket m \rrbracket$, $\ell \in \llbracket n \rrbracket$, and $F_{U_{(\ell:n, k:m)}}^{-1}$ be defined as in Theorem 10. For any $\beta \in (0, 1)$, we have

$$F_{U_{(\ell:n, k:m)}}^{-1}(\beta) \geq F_{U_{(\ell:n)}}^{-1} \left(\frac{k}{m+1} - \sqrt{\frac{\log(1/\beta)}{2(m+2)}} \right). \quad (21)$$

Proposition 13 is proved in Appendix D.3. From this result, we see that for any $\ell \in \llbracket n \rrbracket$, taking (if possible)

$$k = \tilde{k}_{m,n}^{\text{cond}}(\ell, \alpha, \beta) := \left\lceil (m+1) \left(F_{U_{(\ell;n)}}^{-1}(1-\alpha) + \sqrt{\frac{\log(1/\beta)}{2(m+2)}} \right) \right\rceil \quad (22)$$

implies that $F_{U_{(\ell;n, \tilde{k}_{m,n}^{\text{cond}}(\ell, \alpha, \beta); m)}}^{-1}(\beta) \geq 1 - \alpha$ and by Theorem 10, $\widehat{\mathcal{C}}_{\ell, \tilde{k}_{m,n}^{\text{cond}}(\ell, \alpha, \beta)}$ satisfies the training-conditional condition (16). Choosing ℓ similarly to Algorithm 3 leads to the following algorithm.

Algorithm 4 (QQC-Fast). *Given $\alpha \in (0, 1)$ and $\beta \in (0, 1)$,*

$$\begin{aligned} \text{compute} \quad \tilde{\ell}^{\text{cond}} &= \underset{\ell \in \llbracket n \rrbracket \text{ s.t. } \tilde{k}_{m,n}^{\text{cond}}(\ell, \alpha, \beta) \leq m}{\text{argmin}} \left\{ F_{U_{(\ell;n, \tilde{k}_{m,n}^{\text{cond}}(\ell, \alpha, \beta); m)}}^{-1}(1-\beta) \right\} \\ \text{and} \quad \tilde{k}_{m,n}^{\text{cond}}(\tilde{\ell}^{\text{cond}}, \alpha, \beta), \quad &\text{where } \tilde{k}_{m,n}^{\text{cond}}(\ell, \alpha, \beta) \text{ is defined by Eq. (22)}. \\ \text{Output} \quad \widehat{\mathcal{C}}_{\tilde{\ell}^{\text{cond}}, \tilde{k}_{m,n}^{\text{cond}}(\tilde{\ell}^{\text{cond}}, \alpha, \beta)}(x) &= \left\{ y \in \mathcal{Y} : s(x, y) \leq S_{(\tilde{\ell}^{\text{cond}}, \tilde{k}_{m,n}^{\text{cond}}(\tilde{\ell}^{\text{cond}}, \alpha, \beta))} \right\}. \end{aligned}$$

By construction, following the arguments detailed after Proposition 13, the set $\widehat{\mathcal{C}}_{\tilde{\ell}^{\text{cond}}, \tilde{k}_{m,n}^{\text{cond}}(\tilde{\ell}^{\text{cond}}, \alpha, \beta)}$ satisfies the training-conditional condition (16). It is also nontrivial when mn is large enough, as shown by the following lemma.

Lemma 14. *The argmin defining $\tilde{\ell}^{\text{cond}}$ in Algorithm 4 is non-empty —hence the set $\widehat{\mathcal{C}}_{\tilde{\ell}^{\text{cond}}, \tilde{k}_{m,n}^{\text{cond}}(\tilde{\ell}^{\text{cond}}, \alpha, \beta)}$ is nontrivial— if and only if*

$$\frac{m}{m+1} - \sqrt{\frac{\log(1/\beta)}{2(m+2)}} \geq (1-\alpha)^n. \quad (23)$$

For instance, it holds true when

$$n \geq \frac{\log(1/3)}{\log(1-\alpha)} \quad \text{and} \quad m \geq \max \left\{ 2, \frac{9}{2} \log(1/\beta) - 2 \right\}. \quad (24)$$

Lemma 14 is proved in Appendix D.4. In addition, Algorithm 4 is computationally efficient. Indeed, $F_{U_{(\ell;n, k;m)}}^{-1} := F_{U_{(\ell;n)}}^{-1} \circ F_{U_{(k;m)}}^{-1}$ by Theorem 10, so Algorithm 4 requires at most n evaluations of the functions $F_{U_{(\ell;n)}}$, $F_{U_{(\ell;n)}}^{-1}$ and $F_{U_{(k;m)}}^{-1}$, which are fast to evaluate even when n, m are large as explained below Algorithm 2.

Remark 15. *Note that the computational complexities of all proposed algorithms have two parts: (i) the choice of a valid pair (ℓ, k) , and (ii) the computation of $S_{(\ell, k)}$ given the data and (ℓ, k) . The second part is not an issue since \widehat{f} is supposed to be given and the score functions usually can be computed fastly. So, all the complexities discussed throughout Section 3 are about the first part, which can be reused across multiple data sets, score functions, or predictors \widehat{f} , as long as m (the number of agents) and n (the size of local data sets) remain fixed.*

4. Upper bounds on the coverage

In the previous section, we focused on the problem of choosing ℓ and k such that $1 - \alpha_{\ell, k}(\mathcal{D})$ is above $1 - \alpha$, either in expectation (Section 3.1), or with high probability (Section 3.2). In the present section, we raise the natural and important question of upper-bounding

$1 - \alpha_{\ell,k}(\mathcal{D})$ in expectation or with high-probability, as a function of k, ℓ, m , and n . In particular, when the coverage is guaranteed to be larger than $1 - \alpha$, it is interesting to evaluate how far from $1 - \alpha$ it can be, in order to quantify the quality of the corresponding algorithms.

4.1. Marginal coverage upper bounds

In this section, we provide an upper bound on the expectation of the coverage $\mathbb{E}[1 - \alpha_{\ell,k}(\mathcal{D})]$ obtained by Algorithms 1–2 of Section 3.1. Recall that, by Theorem 5, when the scores are a.s. distinct, $\mathbb{E}[1 - \alpha_{\ell,k}(\mathcal{D})] = \mathbb{P}(Y \in \widehat{\mathcal{C}}_{\ell,k}(X)) = M_{\ell,k}$. Hence, in that case it is sufficient to upper-bound $M_{\ell,k}$, as done by the following result.

Theorem 16. *Given $\alpha \in (0, 1)$, let (ℓ^*, k^*) and $(\tilde{\ell}, \tilde{k}_{m,n}(\tilde{\ell}, \alpha))$ be the pairs respectively returned by Algorithms 1 and 2 of Section 3.1, where $\tilde{k}_{m,n}(\ell, \alpha)$ is defined by Eq. (13). Then, some $n_0(\alpha), m_0 \geq 1$ exist such that, if $n \geq n_0(\alpha)$ and $m \geq m_0$, we have*

$$1 - \alpha \leq M_{\ell^*, k^*} \leq M_{\tilde{\ell}, \tilde{k}_{m,n}(\tilde{\ell}, \alpha)} \leq 1 - \alpha + \frac{C}{(2m+1)\sqrt{n+2}}, \quad (25)$$

where C denotes a numerical constant (for instance, the result holds true with $m_0 = 18$ and $C = 27$).

Theorem 16 is proved in Appendix E.1. Combined with Theorem 5, it shows that, if the scores are a.s. distinct, and when m, n are sufficiently large, Algorithms 1 and 2 return prediction sets with marginal coverage between $1 - \alpha$ and $1 - \alpha + \mathcal{O}(n^{-1/2}m^{-1})$. In particular, their marginal coverages tend to $1 - \alpha$ when either n or m tends to infinity, a result that was not obtained previously for any one-shot FL marginally-valid prediction set [25].

Note that we do not exactly recover the upper bound $1 - \alpha + 1/(mn)$ obtained in the centralized case when there are mn calibration points —Eq. (6) in Section 2.1. For Algorithm 1, following [25] we conjecture the stronger result $M_{\ell^*, k^*} \leq 1 - \alpha + \mathcal{O}(1/mn)$. While our experiments in Section 6.1 corroborate this, proving it would require a different analysis that we leave for future work. For Algorithm 2, the results of Section 6.1.1 suggest that Eq. (25) is tight (up to the value of C).

4.2. Training-conditional coverage upper bounds

We now provide high-probability upper bounds on the coverage $1 - \alpha_{\ell,k}(\mathcal{D})$ obtained by Algorithms 3–4 of Section 3.2.

Theorem 17. *Let $\alpha \in (0, 1)$ and $\beta \in (0, 1)$ be given. Recall that the pairs (ℓ_c^*, k_c^*) and $(\tilde{\ell}^{\text{cond}}, \tilde{k}_{m,n}^{\text{cond}}(\tilde{\ell}^{\text{cond}}, \alpha, \beta))$ respectively denote the pairs (ℓ, k) chosen by Algorithms 3 and 4. In the setting of Section 2.2, with Assumption 1, assuming in addition that the scores $S_{i,j}, S$ are a.s. distinct, some $n'_0(\alpha), m'_0(\beta) \geq 1$ exist such that, for any $n \geq n'_0(\alpha)$ and $m \geq m'_0(\beta)$,*

$$\mathbb{P} \left(1 - \alpha_{\ell_c^*, k_c^*}(\mathcal{D}) \leq 1 - \alpha + \frac{\Delta(\beta)}{\sqrt{(m+2)(n+2)}} \right) \geq 1 - \beta \quad (26)$$

$$\text{and} \quad \mathbb{P} \left(1 - \alpha_{\tilde{\ell}^{\text{cond}}, \tilde{k}_{m,n}^{\text{cond}}(\tilde{\ell}^{\text{cond}}, \alpha, \beta)}(\mathcal{D}) \leq 1 - \alpha + \frac{\Delta(\beta)}{\sqrt{(m+2)(n+2)}} \right) \geq 1 - \beta, \quad (27)$$

with $\Delta(\beta) = 12 \max\{2\sqrt{\log(1/\beta)}, 1\}$.

Theorem 17 is proved in Appendix E.2. Remarkably, the bounds are of the form $1 - \alpha + \mathcal{O}(1/\sqrt{mn})$, exactly as for split CP in the centralized case (see Remark 1 in Section 2.1). Therefore, up to constants, using one-shot FL training-conditionally valid prediction sets —also known as (α, β) -tolerance regions— such as Algorithms 3 and 4 incurs no loss compared to a centralized algorithm such as split CP [61].

A consequence of Theorem 17 is that, under the same assumptions, for any $n \geq n'_0(\alpha)$ and $m \geq m'_0(\beta)$,

$$\mathbb{P} \left(1 - \alpha \leq 1 - \alpha_{\ell_c^*, k_c^*}(\mathcal{D}) \leq 1 - \alpha + \frac{\Delta(\beta)}{\sqrt{(m+2)(n+2)}} \right) \geq 1 - 2\beta \quad (28)$$

(by Theorem 10 and the union bound), and a similar results holds true for the output of Algorithm 4. This shows that the deviations of the coverage are (at most) of order $(mn)^{-1/2}$, as for split CP in the centralized case [61, Proposition 2a]. The proof of this theorem, and by extension Eq. (28), rely on the fact that

$$[F_{U(\ell:n, k:m)}^{-1}(\beta), F_{U(\ell:n, k:m)}^{-1}(1-\beta)] \subseteq \left[1 - \alpha, 1 - \alpha + \frac{\Delta(\beta)}{\sqrt{(m+2)(n+2)}} \right]$$

when the pair (ℓ, k) is chosen by our algorithms. Then, thanks to Eq. (18) in Theorem 10, we know that, with probability $1 - 2\beta$, the coverage random variable $1 - \alpha_{\ell, k}(\mathcal{D})$ belongs to the first interval, which concludes the proof. This means that, when (ℓ, k) is the pair chosen by either Algorithm 3 or Algorithm 4, we have that $[F_{U(\ell:n, k:m)}^{-1}(\beta), F_{U(\ell:n, k:m)}^{-1}(1-\beta)]$ contains $1 - \alpha_{\ell, k}(\mathcal{D})$ with probability $1 - 2\beta$, and that both $F_{U(\ell:n, k:m)}^{-1}(\beta)$ and $F_{U(\ell:n, k:m)}^{-1}(1-\beta)$ tend to $1 - \alpha$ at a rate of at least $(mn)^{-1/2}$. According to the numerical experiments of Section 6.1, this rate $(mn)^{-1/2}$ seems exact (that is, not improvable in worst-case) when (ℓ, k) is chosen by Algorithm 4.

Remark 18. *The proof of Theorem 17 also shows that*

$$\mathbb{P} \left(1 - \alpha_{\ell_n, \tilde{k}_{m,n}^{cond}(\ell_n, \alpha, \beta)}(\mathcal{D}) \leq 1 - \alpha + \frac{\Delta(\beta)}{\sqrt{(m+2)(n+2)}} \right) \geq 1 - \beta$$

for $\ell_n = \lceil n(1 - \alpha) \rceil$. This inequality could be used to build prediction sets with a smaller computational complexity than Algorithm 4, at the price of being slightly more conservative. By Theorem 31 in Appendix G.2, the proof can also be generalized to any sequence $(\ell_n)_{n \geq 1}$ such that $\ell_n \in \llbracket n \rrbracket$ for every n and $\ell_n = n(1 - \alpha) + o(\sqrt{n})$ as $n \rightarrow +\infty$.

5. Generalization to different n_j

In the previous sections, for simplicity, we assume that all agents have the same number n of (calibration) data points. Let us now consider the general case, that is, the setting described in Section 2.2 *without* Assumption 1. Since the m agents have data sets of possibly different sizes n_1, \dots, n_m , one needs to adapt the order of the local quantile to each agent. Given $(\ell_1, \dots, \ell_m) \in \llbracket n_1 \rrbracket \times \dots \times \llbracket n_m \rrbracket$, the QQ estimator (Definition 4) is now defined by

$$\mathcal{S}_{(\ell_1, \dots, \ell_m, k)} := \widehat{Q}_{(k:m)} \left(\widehat{Q}_{(\ell_1:n_1)}(\mathcal{S}_1), \dots, \widehat{Q}_{(\ell_m:n_m)}(\mathcal{S}_m) \right), \quad (29)$$

where $\mathcal{S}_j = (\mathcal{S}_{i,j})_{1 \leq i \leq n_j}$ denotes the set of local scores computed by the j -th agent. In words, each agent $j \in \llbracket m \rrbracket$ now sends to the server its ℓ_j -th smallest score and the server

computes the k -th smallest value of these values. The associated prediction set is, for any $x \in \mathcal{X}$,

$$\widehat{\mathcal{C}}_{\ell_1, \dots, \ell_m, k}(x) := \{y \in \mathcal{Y} : s(x, y) \leq S_{(\ell_1, \dots, \ell_m, k)}\}, \quad (30)$$

and its miscoverage rate is denoted by

$$\alpha_{\ell_1, \dots, \ell_m, k}(\mathcal{D}) := \mathbb{P}(Y \notin \widehat{\mathcal{C}}_{\ell_1, \dots, \ell_m, k}(X) \mid \mathcal{D}). \quad (31)$$

When the n_j are different, the difficulty is that the random variables $(\widehat{Q}_{(\ell_j; n_j)}(\mathcal{S}_j))_{1 \leq j \leq m}$ in the right-hand side of Eq. (29) are not identically distributed as they are computed on data sets of different sizes. Hence, in this setting, the counterpart of Theorem 5 of Section 3.1 (marginal guarantee) is as follows.

Theorem 19. *In the setting of Section 2.2, without Assumption 1, for any $(\ell_1, \dots, \ell_m, k) \in \llbracket n_1 \rrbracket \times \dots \times \llbracket n_m \rrbracket \times \llbracket m \rrbracket$, the set $\widehat{\mathcal{C}}_{\ell_1, \dots, \ell_m, k}$ defined by Eq. (30) satisfies*

$$\mathbb{P}(Y \in \widehat{\mathcal{C}}_{\ell_1, \dots, \ell_m, k}(X)) \geq M_{\ell_1, \dots, \ell_m, k} \quad (32)$$

with

$$M_{\ell_1, \dots, \ell_m, k} = 1 - \frac{1}{1 + \sum_{j=1}^m \frac{1}{n_j}} \sum_{j=k}^m \sum_{a \in \mathcal{P}_j} \sum_{i_1 = \ell_{a_1}}^{n_{a_1}} \dots \sum_{i_j = \ell_{a_j}}^{n_{a_j}} \sum_{i_{j+1}=0}^{\ell_{a_{j+1}}-1} \dots \sum_{i_m=0}^{\ell_{a_m}-1} \frac{\binom{n_{a_1}}{i_1} \dots \binom{n_{a_m}}{i_m}}{\binom{n_1 + \dots + n_m}{i_1 + \dots + i_m}},$$

where \mathcal{P}_j denotes the set of permutations $a = (a_1, \dots, a_m)$ of $\{1, \dots, m\}$ such that $a_1 < a_2 < \dots < a_j$ and $a_{j+1} < a_{j+2} < \dots < a_m$. Moreover, when the associated scores $(S_{i,j})_{1 \leq j \leq m, 1 \leq i \leq n_j}$ and $S := s(X, Y)$ are almost surely distinct, Eq. (32) is an equality.

Theorem 19 is proved in Appendix F.1. Furthermore, the miscoverage rate defined by Eq. (31) can be controlled with high probability thanks to the following result, which generalizes Theorem 10 in Section 3.2.

Theorem 20. *In the setting of Section 2.2, without Assumption 1, for any $(\ell_1, \dots, \ell_m, k) \in \llbracket n_1 \rrbracket \times \dots \times \llbracket n_m \rrbracket \times \llbracket m \rrbracket$, and any $\alpha \in (0, 1)$, the miscoverage random variable $\alpha_{\ell_1, \dots, \ell_m, k}(\mathcal{D})$ defined by Eq. (31) satisfies*

$$\mathbb{P}(1 - \alpha_{\ell_1, \dots, \ell_m, k}(\mathcal{D}) \geq 1 - \alpha) \geq \text{PB}\left(k - 1; (F_{U_{(\ell_j; n_j)}}(1 - \alpha))_{1 \leq j \leq m}\right), \quad (33)$$

where $F_{U_{(\ell_j; n_j)}}$ denotes the cdf of the Beta($\ell_j, n_j - \ell_j + 1$) distribution and for every $t \in \mathbb{R}$ and $u \in [0, 1]^m$, $\text{PB}(\cdot; u)$ denotes the cdf of a Poisson-Binomial (PB) random variable with parameter u ². Moreover, when the associated scores $(S_{i,j})_{1 \leq j \leq m, 1 \leq i \leq n_j}$ and $S := s(X, Y)$ are almost surely distinct, Eq. (33) is an equality.

Theorem 20 is proved in Appendix F.2. Theorem 19 and Theorem 20 imply that (i) if $M_{\ell_1, \dots, \ell_m, k} \geq 1 - \alpha$, then $\widehat{\mathcal{C}}_{\ell_1, \dots, \ell_m, k}$ is a marginally valid prediction set, and (ii) given $\beta \in (0, 1)$ if $\text{PB}\left(k - 1; (F_{U_{(\ell_j; n_j)}}(1 - \alpha))_{1 \leq j \leq m}\right) \geq 1 - \beta$, then $\widehat{\mathcal{C}}_{\ell_1, \dots, \ell_m, k}$ is an (α, β) -tolerance region, that is, a training-conditionally valid prediction set. It remains to select $(\ell_1, \dots, \ell_m, k)$, among those satisfying the desired condition. Applying directly the strategy of Algorithms 1 or 3, although theoretically possible, can be computationally untractable since the ℓ_j can be different, hence a parameter set of cardinality $m \cdot n^m$.

²Recall that the PB distribution with parameter $u \in [0, 1]^m$ is the distribution of $Z_1 + \dots + Z_m$ where the random variables Z_j are independent and respectively follow the Bernoulli(u_j) distribution. It therefore generalizes the Binomial(m, p) distribution, which corresponds to the case where $u_1 = \dots = u_m = p$. More details can be found in [58].

Therefore, we propose to fix $\ell_j = \lceil (1 - \alpha)(n_j + 1) \rceil$ for all $j \in \llbracket m \rrbracket$, similarly to the classical (centralized) split CP methodology. Then, we only need to select $k \in \llbracket m \rrbracket$, thereby reducing significantly the computational complexity.³ Following the principles described in Section 2.3 for choosing k , Theorems 19–20 lead to the two following algorithms.

Algorithm 5 (QQM- n_j [25]). *Given $\alpha \in (0, 1)$,*

$$\begin{aligned} & \text{compute } \ell_j^* = \lceil (1 - \alpha)(n_j + 1) \rceil \quad \text{for every } j \in \llbracket m \rrbracket \\ & \text{and } k^*(\ell_{1\dots m}^*) = \arg \min_{k \in \llbracket m \rrbracket} \{ M_{\ell_1^*, \dots, \ell_m^*, k} : M_{\ell_1^*, \dots, \ell_m^*, k} \geq 1 - \alpha \} . \\ & \text{Output } \widehat{\mathcal{C}}_{\ell_1^*, \dots, \ell_m^*, k^*(\ell_{1\dots m}^*)}(x) = \{ y \in \mathcal{Y} : s(x, y) \leq S_{(\ell_1^*, \dots, \ell_m^*, k^*(\ell_{1\dots m}^*))} \} . \end{aligned}$$

Algorithm 6 (QQC- n_j). *Given $\alpha \in (0, 1)$ and $\beta \in (0, 1)$,*

$$\begin{aligned} & \text{compute } \ell_j^* = \lceil (1 - \alpha)(n_j + 1) \rceil \quad \text{for every } j \in \llbracket m \rrbracket \\ & \text{and } k_c^*(\ell_{1\dots m}^*) = \operatorname{argmin}_{k \in \mathcal{K}} \left\{ \text{PB} \left(k - 1; (F_{U(\ell_j^*; n_j)}(1 - \alpha))_{1 \leq j \leq m} \right) \right\} , \\ & \text{where } \mathcal{K} := \left\{ k \in \llbracket m \rrbracket : \text{PB} \left(k - 1; (F_{U(\ell_j^*; n_j)}(1 - \alpha))_{1 \leq j \leq m} \right) \geq 1 - \beta \right\} . \\ & \text{Output } \widehat{\mathcal{C}}_{\ell_1^*, \dots, \ell_m^*, k_c^*(\ell_{1\dots m}^*)}(x) = \{ y \in \mathcal{Y} : s(x, y) \leq S_{(\ell_1^*, \dots, \ell_m^*, k_c^*(\ell_{1\dots m}^*))} \} . \end{aligned}$$

For reasons detailed below Theorem 20, Algorithm 5 yields a distribution-free marginally valid prediction set, and Algorithm 6 yields a distribution-free training-conditionally valid prediction set. The exact computation of $M_{\ell_1, \dots, \ell_m, k}$ can be done in the same way as that of $M_{\ell, k}$ (see details in [25, Appendix A.2.]). Furthermore, because $M_{\ell_1, \dots, \ell_m, k} = \mathbb{E}[Z]$ where Z has cdf $\text{PB}(k - 1; (F_{U(\ell_j; n_j)}(1 - \alpha))_{1 \leq j \leq m})$, another possibility is to numerically integrate this Poisson-Binomial cdf to approximate the expectation. There exist many algorithms to compute the Poisson-Binomial cdf of Eq. (33) efficiently, see for instance [23]. Moreover, precise approximations of it can be obtained when m is large. We refer to the recent review of [58] on the Poisson-Binomial distribution for more details.

Remark 21 (Other variants of Algorithms 5–6). *Following the comments below Theorem 20, Algorithm 5 (resp. Algorithm 6) remains marginally (resp. training-conditionally) valid with any other choice for the ℓ_j^* , leading to many possible variants. Let us detail three of them. First, in Algorithm 6, one can replace ℓ_j^* by the value of r_c defined by Eq. (34) when $|\mathcal{D}^{\text{cal}}| = n_j$, which corresponds to training-conditionally valid splitCP applied to the data of agent j alone. Second, when the argmin defining k^* or k_c^* is empty, it is possible to increase the ℓ_j^* (for instance, by adding 1 to each of them) until the argmin becomes nonempty (which happens for large enough ℓ_j^* if and only if Eq. (20) in Lemma 12) holds true. Third, when m is small, one can consider small sets \mathcal{L}_j^* of candidate values for ℓ_j^* , and replace the argmin in Algorithms 5–6 by an argmin over $(k, \ell_1^*, \dots, \ell_m^*) \in \mathcal{K} \times \mathcal{L}_1^* \times \dots \times \mathcal{L}_m^*$, which yields a tractable algorithm provided this set remains small. While our Theorems 19–20 prove coverage lower bounds for each of these variants, we leave their detailed study for future work.*

Remark 22. *If we set $n_1 = \dots = n_m = n$ and $\ell_1 = \dots = \ell_m = \ell \in \llbracket n \rrbracket$ in Eq. (32), respectively in Eq. (33), we exactly recover the results obtained in Theorem 5 of Section 3.1 (marginal guarantee), resp. in Theorem 10 of Section 3.2 (conditional*

³This strategy can also be used when $n_j = n$ (Section 3) at the cost of being less accurate than searching for all values of ℓ and k . See also Section 6.1.2 and Remark 18 in Appendix E.2.

guarantee). Indeed, the summation over permutations $a \in \mathcal{P}_j$ in Theorem 19 then becomes the $\binom{m}{k}$ of Theorem 5, after rearranging the terms. For Theorem 20, the parameters of the Poisson-Binomial $(F_{U_{(\ell_j:n_j)}}(1-\alpha))_{1 \leq j \leq m}$ are then equal, hence we get the Binomial($m; F_{U_{(\ell:n)}}(1-\alpha)$) distribution. This leads to Theorem 10 with $\beta = F_{U_{(\ell:n,k:m)}}(1-\alpha) = F_{U_{(k:m)}} \circ F_{U_{(\ell:n)}}(1-\alpha)$ since $F_{\text{Binomial}(m,p)}(k-1) = 1 - F_{U_{(k:m)}}(p)$, here used with $p = F_{U_{(\ell:n)}}(1-\alpha)$.

6. Numerical experiments

In this section, we first study the behavior of our algorithms in a generic setting and then evaluate them on real datasets. The code of our experiments is publicly available.⁴ In the experiments, we compare the performance of the following conformal-based prediction-set algorithms:

- Algorithms 1–4, respectively called QQM, QQM-Fast, QQC and QQC-Fast (when the n_j are equal);
- Algorithms 5–6, respectively called QQM- n_j and QQC- n_j (when the n_j may be different);
- CentralM: split conformal prediction $\widehat{\mathcal{C}}_r$ (as defined in Section 2.1) on the full (centralized) calibration data set \mathcal{D}^{cal} , with $r = \lceil (1-\alpha)(|\mathcal{D}^{\text{cal}}| + 1) \rceil$ so that it is marginally valid; in other words, CentralM is the centralized version of QQM and they coincide when $m = 1$;
- CentralC: split conformal prediction $\widehat{\mathcal{C}}_r$ on the full (centralized) calibration data set \mathcal{D}^{cal} with

$$r = r_c := \min \{ \tilde{r} \in \llbracket |\mathcal{D}^{\text{cal}}| \rrbracket : F_{U_{(\tilde{r}; |\mathcal{D}^{\text{cal}}|)}}^{-1}(\beta) \geq 1 - \alpha \}, \quad (34)$$

so that it is conditionally valid, by Eq. (7) in Section 2.1; in other words, CentralC is the centralized version of QQC and they coincide when $m = 1$;

- FedCP-Avg: federated approach proposed by [38], which averages the m quantiles of order $\lceil (n_j + 1)(1-\alpha) \rceil$ sent by the agents, hence the prediction set

$$\widehat{\mathcal{C}}_{\text{FedCP-Avg}}(x) := \left\{ y \in \mathcal{Y} : s(x, y) \leq \frac{1}{m} \sum_{j=1}^m \widehat{Q}_{\lceil (n_j+1)(1-\alpha) \rceil}(\mathcal{S}_j) \right\}.$$

All our experiments are done with $\alpha = 0.1$ and $\beta = 0.2$.

As explained in Section 2.3, we compare the performance of these methods either through the length of the prediction sets they produce (in Section 6.2, where these are intervals) or their coverage, which should both be as small as possible while keeping the associated set valid. Note that it is fair to compare split CP (CentralM and CentralC), FedCP-Avg and our quantile-of-quantiles prediction sets (Algorithms 1–6) through their coverages since they all are of the form $\{y \in \mathcal{Y} : s(x, y) \leq \overline{S}\}$ for some random variable \overline{S} .

6.1. Generic comparison.

In this section, we compare numerically the performance (measured by their coverages) of our algorithms and their centralized counterparts (CentralM and CentralC). This comparison is made under the mild assumption that F_S is continuous, which implies that the cdf of the training-conditional coverage distribution of each algorithm is known and,

⁴<https://github.com/pierreHmbt/One-shot-FCP>

more importantly, universal – see Eq. (7) and the discussion below for centralized algorithms, Theorem 10 for federated algorithms with equal n_j and Theorem 20 for federated algorithms with general n_j . We also have closed-form formulas for the expected coverage when F_S is continuous, by Eq. (6) for centralized algorithms, and by Theorems 5 and 19 for our federated algorithms. Under this *generic setting*, we thus consider simultaneously all learning problems, data distributions, predictors \hat{f} and score functions s such that the scores cdf F_S is continuous.

Throughout this subsection, for each algorithm considered, denoting by $1 - \alpha(\mathcal{D})$ its (random) coverage, we are interested in $F_{1-\alpha(\mathcal{D})}$ the coverage cdf, $\Delta\mathbb{E} := \mathbb{E}[1 - \alpha(\mathcal{D})] - (1 - \alpha)$ the difference between the expected coverage and the nominal coverage $(1 - \alpha)$, and, for several values of $\zeta \in (0, 1)$, $\Delta q_\zeta := F_{1-\alpha(\mathcal{D})}^{-1}(\zeta) - (1 - \alpha)$ the difference between the ζ -quantile of the coverage and the nominal coverage.

6.1.1. Equal n_j

We first compare QQM, QQM-Fast, QQC, QQC-Fast, CentralM and CentralC under Assumption 1, that is, when each agent $j \in \llbracket m \rrbracket$ has the same number $n_j = n$ of calibration data points. In Section 4, we prove that $\Delta\mathbb{E} = \mathcal{O}(m^{-1}n^{-1/2})$ for our marginally valid federated prediction sets (Theorem 16) and that $\Delta q_{1-\beta} = \mathcal{O}(m^{-1/2}n^{-1/2})$ for our training-conditionally valid federated prediction sets (Theorem 17). But these only are *upper bounds*, and we would like to know whether they match the true order of magnitude of their coverages —at least in worst case, since one can for instance get $\Delta\mathbb{E} = 0$ for QQM by choosing $\alpha = 1 - M_{\ell,k}$ for some $(\ell, k) \in \llbracket n \rrbracket \times \llbracket m \rrbracket$, or for CentralM by choosing $\alpha \in \{\frac{r}{n_c+1} : r \in \llbracket n_c + 1 \rrbracket\}$. Furthermore, knowing the true (worst-case) order of magnitude of the coverage of each algorithm is crucial to determine when one of the federated prediction sets we propose should be preferred to another, and for knowing precisely what may be lost by considering the one-shot FL setting instead of the centralized setting.

Methods. In order to answer these questions, for each algorithm considered, we evaluate below the rates of convergence towards zero of $\Delta\mathbb{E}$, Δq_β and $\Delta q_{1-\beta}$ as m and n increase. The choice of the quantile orders β and $(1 - \beta)$ comes from the fact that (i) the conditional validity is equivalent to $\Delta q_\beta \geq 0$, so for prediction sets satisfying this condition, knowing how tight is this inequality in worst case is meaningful, and (ii) since $\beta = 0.2$, the β and $(1 - \beta)$ -quantiles of the coverage are the bounds of an interval containing “typical values” of the coverage. In particular, by definition, the coverage $1 - \alpha(\mathcal{D})$ is below $(1 - \alpha) + \Delta q_{1-\beta}$ with probability $1 - \beta = 80\%$.

We consider values of (m, n) in the grid $\{\lfloor 10^{i/3} \rfloor : i \in \llbracket 9 \rrbracket\}^2$. For each pair (m, n) and each algorithm, we compute $\Delta\mathbb{E}$, Δq_β and $\Delta q_{1-\beta}$. Then, for each algorithm, using empirical risk minimization with the Huber loss [24], we robustly fit the log-linear regression model $\log y = \log(c) - \gamma \log(m) - \delta \log(n) + \varepsilon$, where ε is some residual term, $c > 0$ and $\gamma, \delta \in \mathbb{R}$ are the model parameters, and y is either $\Delta\mathbb{E}$, Δq_β or $\Delta q_{1-\beta}$. Note that by construction, $\Delta\mathbb{E} \geq 0$ for marginally-valid algorithms, and $\Delta q_{1-\beta} \geq \Delta q_\beta \geq 0$ for conditionally-valid algorithms. The estimated values of c, γ, δ for each algorithm and each quantity of interest are reported in Table 1. Plots showing the values of $\Delta\mathbb{E}$, Δq_β , $\Delta q_{1-\beta}$ (together with their log-linear model approximation) are provided in Figures 1 and 3, and in Appendix J.1.

In order to compare the performance of the algorithms more closely, we also plot in Figures 2 and 4 the full cdfs of the coverage distribution for two specific pairs (m, n) , $(200, 20)$ and $(20, 200)$, and report the corresponding expectations, standard-deviations (computed by numerically integrating the cdfs), and ζ -quantiles for $\zeta \in \{\beta, 1 - \beta\}$ in Table 2.

Method	$\Delta\mathbb{E} \approx c_1 m^{-\gamma_1} n^{-\delta_1}$			$\Delta q_\beta \approx c_2 m^{-\gamma_2} n^{-\delta_2}$			$\Delta q_{1-\beta} \approx c_3 m^{-\gamma_3} n^{-\delta_3}$		
	c_1	γ_1	δ_1	c_2	γ_2	δ_2	c_3	γ_3	δ_3
CentralM	0.502	0.999	0.999	0.235	0.498	0.498	0.269	0.501	0.501
QQM	0.924	1.026	1.019	0.463	0.492	0.487	0.522	0.501	0.495
QQM-Fast	0.336	1.001	0.582	0.227	0.483	0.499	0.400	0.512	0.504
CentralC	0.257	0.501	0.501	0.590	1.014	1.014	0.504	0.499	0.499
QQC	0.341	0.507	0.502	0.508	0.910	0.847	0.637	0.502	0.497
QQC-Fast	0.676	0.498	0.501	0.400	0.503	0.509	0.958	0.496	0.498

TABLE 1

Estimated parameters of the log-linear model $\log y = \log(c_i) - \gamma_i \log(m) - \delta_i \log(n)$ where y is either $\Delta\mathbb{E}$, Δq_β or $\Delta q_{1-\beta}$, for the six algorithms compared in Section 6.1.1; see text for details.




FIG 1. Marginally-valid algorithms: log-log plot of $\Delta\mathbb{E}$ as a function of m (left) or n (right). Lines show the approximation $\log \Delta\mathbb{E} \approx \log(c_1) - \gamma_1 \log(m) - \delta_1 \log(n)$ with c_1, γ_1, δ_1 given by Table 1. Plain lines and dots correspond to $n = 10$ (left) or $m = 10$ (right). Dashed lines and diamonds correspond to $n = 10^4$ (left) or $m = 10^4$ (right).

Marginally-valid algorithms. For QQM and CentralM, according to Table 1, $\Delta\mathbb{E}$ decrease to zero at the same rate, close to $m^{-1}n^{-1}$, with only a multiplicative constant ≈ 2 in favour of CentralM. Overall, the difference between expected coverages of QQM and CentralM is small, as shown by Figure 1. In comparison, QQM-Fast shows a worse performance with $\Delta\mathbb{E}$ decreasing to zero at a rate close to $m^{-1}n^{-1/2}$ (Table 1). This leads to a significant gap between the performances of QQM-Fast and the two other ones as shown by Figure 1. For $(m, n) \in \{(200, 20), (20, 200)\}$, the coverage cdfs of the three algorithms have similar shapes (Figure 2) apart from a shift corresponding to already mentioned differences between expectations, and a slightly larger standard-deviation for federated algorithms compared to CentralM (see also Table 2). The $(1 - \beta)$ -quantile of the coverage is also slightly larger for the federated algorithms compared to CentralM, even if they all are of the same order of magnitude $(1 - \alpha) + \mathcal{O}(m^{-1/2}n^{-1/2})$ at first order (see Tables 1 and 2).

Note also that Table 1 suggests that our theoretical expected coverage upper bound (Theorem 16) $(1 - \alpha) + \mathcal{O}(m^{-1}n^{-1/2})$ provides the correct worst-case order of magnitude for QQM-Fast (even if the estimated exponent for n is $\delta_1 = 0.58$ instead of exactly $1/2$), and supports our conjecture that it could be improved to $(1 - \alpha) + \mathcal{O}(m^{-1}n^{-1})$ for QQM.

Overall, these numerical results show that among one-shot FL prediction sets, QQM should be preferred to QQM-Fast as long as its computational complexity remains tractable (keeping in mind Remark 15 about the computational complexity of our federated algorithms). Then, using QQM yields a mild loss compared to CentralM, with a one-shot FL algorithm. Note however that the expected coverage of QQM-Fast also converges to $1 - \alpha$ when mn tends to infinity, as proved by Theorem 16 and illustrated in this section, so when m or n is too large so that one must use QQM-Fast, the loss in terms of




FIG 2. Marginally-valid algorithms: cdf of the coverage $1 - \alpha(\mathcal{D})$. The mean of each distribution is shown by a dot and the quantiles of order β and $1 - \beta$ by crosses. Left: $(m, n) = (200, 20)$. Right: $(m, n) = (20, 200)$. See Table 2 for additional information.

Method	$(m, n) = (200, 20)$				$(m, n) = (20, 200)$			
	$\mathbb{E}[\cdot]$	Std	q_β	$q_{1-\beta}$	$\mathbb{E}[\cdot]$	Std	q_β	$q_{1-\beta}$
CentralM	0.90002	0.00474	0.89605	0.90403	0.90002	0.00474	0.89605	0.90403
QQM	0.90004	0.00604	0.89500	0.90515	0.90012	0.00603	0.89510	0.90522
QQM-Fast	0.90084	0.00577	0.89601	0.90572	0.90070	0.00585	0.89580	0.90563
CentralC	0.90402	0.00466	0.90013	0.90796	0.90402	0.00466	0.90013	0.90796
QQC	0.90524	0.00569	0.90048	0.91004	0.90526	0.00589	0.90036	0.91025
QQC-Fast	0.91084	0.00558	0.90618	0.91556	0.91046	0.00560	0.90579	0.91519

TABLE 2

Expectation, standard-deviation, β -quantile and $(1 - \beta)$ -quantile of the coverage $1 - \alpha(\mathcal{D})$ of the algorithms compared in Section 6.1.1 (equal n_j).

expected coverage remain very small; for instance, when $m = 200$, the expected coverage of QQM-Fast is 0.90084 while the one of QQM is 0.90004 (Table 2).

Conditionally-valid algorithms. The $(1 - \beta)$ -quantiles of the coverages of the three conditionally-valid algorithms are of the same order of magnitude $1 - \alpha + \mathcal{O}(m^{-1/2}n^{-1/2})$, with constants within a factor 2 in front of the remainder term (Table 1, where the estimated γ_3, δ_3 all are very close to $1/2$; see also Figure 19 in Appendix J.1). This matches our theoretical upper bound (Theorem 10), suggesting its tightness. Bigger differences appear when considering the β -quantiles of the coverage, which seem to be of order $1 - \alpha + \mathcal{O}(m^{-1}n^{-1})$ for CentralC and QQC (with γ_2, δ_2 slightly smaller than 1 for QQC), while it is clearly larger, of order $1 - \alpha + \mathcal{O}(m^{-1/2}n^{-1/2})$, for QQC-Fast, according to Table 1 and Figure 3. A good summary of what happens here may be found in Figure 4: CentralC and QQC tightly adjust the β -quantile of the coverage (which is the one that must be above $1 - \alpha$), better than QQC-Fast.

Overall, as expected, the computationally more efficient method QQC-Fast is also a bit more conservative than the other two methods, which would suggest to use QQC (in a one-shot FL context) if one has enough computational power, leading to only a small loss compared to the centralized case. Note finally that all three training-conditionally valid algorithms also are marginally valid in the settings considered by Figure 4, since the expectations of their coverages are larger than their β -quantiles (which is not surprising since $\beta = 0.2 < 1/2$ and the medians here are close to the expectations).




FIG 3. Conditionally-valid algorithms: log-log plot of Δq_β as a function of m (left) or n (right). Lines show the approximation $\log \Delta q_\beta \approx \log(c_2) - \gamma_2 \log(m) - \delta_2 \log(n)$ with c_2, γ_2, δ_2 given by Table 1. Plain lines and dots correspond to $n = 10$ (left) or $m = 10$ (right). Dashed lines and diamonds correspond to $n = 10^4$ (left) or $m = 10^4$ (right).




FIG 4. Conditionally-valid algorithms: cdf of the coverage $1 - \alpha(\mathcal{D})$. The mean of each distribution is shown by a dot and the quantiles of order β and $1 - \beta$ by crosses. Left: $(m, n) = (200, 20)$. Right: $(m, n) = (20, 200)$. See Table 2 for additional information.

6.1.2. Different n_j

We now study the case where the n_j are different, so that only QQM- n_j and QQC- n_j are available in the one-shot FL case.

Methods. We consider several settings with a total of $N = 4000$ calibration points distributed across m agents, where m divides N . For each value of m considered, the values of (n_1, \dots, n_m) are chosen (once for all) randomly, according to a multinomial distribution with parameters N and $(1/m, \dots, 1/m)$; in other words, each of the N calibration points is uniformly assigned among the m agents, independently. Figure 20 in Appendix J.2.1 provides the exact values of the n_j for $m = 4$ and $m = 25$.

In addition to QQM- n_j and QQC- n_j , which are the only ones able to deal with agents having different number of data points in a one shot FL setting, we also consider several other algorithms for comparison:

- Centralized algorithms (CentralM and CentralC);
- Algorithms 5–6 with data split equally into m agents (hence having each $n_j = N/m$ points), that we call QQM- (N/m) and QQC- (N/m) , respectively;
- QQM and QQC, with data split equally into m agents, as in Section 6.1.1.

For each of these algorithms, we report the values of the expectation, standard-deviation

and ζ -quantile for $\zeta \in \{\beta, 1 - \beta\}$ of their coverages when $m \in \{4, 25\}$ in Table 3, and we plot the cdfs of their coverages when $m = 25$ in Figure 5. Since specific values of m can be misleading (as explained at the beginning of Section 6.1.1), we report $\Delta\mathbb{E}$ (respectively, Δq_β) as a function of m for marginally-valid (respectively, conditionally-valid) algorithms in Figure 6; for completeness, the values of $\Delta q_{1-\beta}$ for conditionally-valid algorithms are plotted in Figure 21 of Appendix J.2.1. Note that the values $m \in \{4, 25\}$ in Table 3 and $m = 25$ in Figure 5 have been chosen because they are typical of the worst-case performance of most algorithms, according to Figure 6.

Method	$(m, N) = (4, 4000)$				$(m, N) = (25, 4000)$			
	$\mathbb{E}[\cdot]$	Std	q_β	$q_{1-\beta}$	$\mathbb{E}[\cdot]$	Std	q_β	$q_{1-\beta}$
CentralM	0.90002	0.00474	0.89605	0.90403	0.90002	0.00474	0.89605	0.90403
QQM	0.90009	0.005659	0.89535	0.90485	0.90017	0.00706	0.89435	0.90617
QQM-(N/m)	0.90305	0.00558	0.89838	0.90775	0.90217	0.00581	0.89731	0.90709
QQM- n_j	0.90335	0.00558	0.89869	0.90804	0.90238	0.00583	0.89755	0.90733
CentralC	0.90402	0.00465	0.90012	0.90796	0.90402	0.00465	0.90012	0.90796
QQC	0.90503	0.00553	0.90040	0.90968	0.90549	0.00581	0.90062	0.91039
QQC-(N/m)	0.90969	0.00622	0.90444	0.91487	0.90679	0.00569	0.90203	0.91160
QQC- n_j	0.90997	0.00620	0.90474	0.91512	0.90698	0.00567	0.90222	0.91177

TABLE 3

Different n_j : Expectation, standard-deviation, β -quantile and $(1 - \beta)$ -quantile of the coverage $1 - \alpha(\mathcal{D})$ of the eight algorithms listed at the beginning of Section 6.1.2.




FIG 5. Different n_j : cdf of the coverage $1 - \alpha(\mathcal{D})$ of the algorithms considered in Section 6.1.2 with $N = \sum_{j=1}^m n_j = 4000$ data points distributed among $m = 25$ agents. The mean of each distribution is shown by a dot and the quantiles of order β and $1 - \beta$ by crosses. See Table 3 for additional information. Left: marginally-valid algorithms: CentralM, QQM with $n_j = N/m$ for all j , QQM-(N/m) (that is, Algorithm 5 with $n_j = N/m$ for all j) and QQM- n_j . Right: conditionally-valid algorithms: CentralC, QQC with $n_j = N/m$ for all j , QQC-(N/m) (that is, Algorithm 6 with $n_j = N/m$ for all j) and QQC- n_j .




FIG 6. Different n_j : log-log plot of the performance (left: $\Delta\mathbb{E}$; right: Δq_β) of the algorithms considered in Section 6.1.2 (left: marginally-valid algorithms; right: conditionally-valid algorithms) as a function of the number of agents m . The total number of data points $N = \sum_{j=1}^m n_j = 4000$ is fixed and m varies within the set of divisors of N . The values when $m = 1, 2$ for $QQC-(N/m)$ and $QQC-n_j$ are missing because there is no k such that the resulting sets are conditionally valid.

Comparison to centralized algorithms and to federated algorithms with $n_j = n$.

Let us start with marginally-valid algorithms, that is, the ones appearing in the top half of Table 3 and on the left of Figures 5–6. The typical ordering of their performance (for instance measured by their expected coverages) is not surprising: first CentralM, then QQM (with a small loss compared to CentralM, as already studied in Section 6.1.1), and finally QQM- (N/m) and QQM- n_j , with similar performances (and a larger loss compared to CentralM). In addition to the increase of expected coverage, the quantiles also increases slightly, with the same ordering. Nevertheless, QQM- n_j has a reasonably good performance, with an expected coverage below $1 - \alpha + 0.005$.

One possible reason explaining the fact that QQM- n_j is more conservative than the federated algorithms we propose for the case of equal n_j (with the same overall sample size) can be the suboptimality of the choice $\ell_j^* = \lceil (1 - \alpha)(n_j + 1) \rceil$ made for computational reasons. Indeed, we see on Table 3 and Figures 5–6 that QQM- (N/m) and QQM- n_j have similar performances and are not as good as QQM or CentralM.

Similar comments can be made about conditionally-valid algorithms, that is, the ones appearing in the bottom half of Table 3 and on the right of Figures 5–6, whose performance should primarily be measured by the β and $(1 - \beta)$ -quantiles of their coverages.

Impact of the number of agents m . Let us now consider when the number of agents m vary while the total number of samples N is fixed. Once again, let us start with marginally-valid algorithms. First, notice that when $m = 1$ and $m = N$, the four algorithms considered choose the same pair (ℓ, k) by definition, hence they have the same performance. Then, putting aside the value $m = 1$, Figure 6 shows that the performance of both QQM- (N/m) and QQM- n_j generally improve when m increases, while the one of QQM is approximately constant. The fact that distributing the data across more agents can yield a better performance may seem surprising, but it can easily be explained. Indeed, as shown by Table 4 and Figure 22 in Appendix J.2.2, for QQM- (N/m) , when both m and $n := N/m$ vary, $\Delta\mathbb{E}$ is approximately proportional to $m^{-1}(N/m)^{-1/2} = m^{-1/2}N^{-1/2}$. So, in the experiment of Figure 6, where N is fixed, $\Delta\mathbb{E}$ is approximately proportional to $m^{-1/2}$, hence it decreases when m increases. In addition, we propose the following intuition for explaining the behaviors of both QQM- n_j and QQM- (N/m) : when m increases, the algorithm has more options for $k \in \llbracket m \rrbracket$ (while each ℓ_j remains fixed equal to ℓ_j^*), allowing to adjust more precisely the coverage (in worst case), hence a

better performance. The fact that QQM behaves differently (with $\Delta\mathbb{E}$ roughly constant when m varies) should not be surprising, since the number of candidate values for (ℓ, k) is equal to $N = nm$ (hence it remains constant), and Table 1 shows that for QQM, $\Delta\mathbb{E}$ is approximately proportional to $(mn)^{-1} = N^{-1}$.

For conditionally-valid algorithms, Figure 6 shows a similar behavior for Δq_β (decreasing function of m for QQC- n_j and QQC- (N/m) , roughly constant for QQC), for which we propose the same intuitive explanation (see also Table 4 and Figure 23 in Appendix J.2.2). Note however that QQC- n_j and QQC- (N/m) do not coincide with QQC and CentralC when $m = 1$ because of the choice made for ℓ_j in Algorithm 6, which prevents to find a value of k satisfying the training-conditional guarantee. Other choices would be possible (see Remark 21 in Section 5).

Conclusion. Our one-shot FL algorithms dealing with non-equal n_j provide prediction sets only slightly more conservative than their centralized counterparts, while being computationally tractable, hence they seem effective for performing CP in one-shot FL with different n_j . Note that the values of m for which QQM- n_j and QQC- n_j are the most conservative are the small values $m \geq 2$, which precisely are the ones for which QQM- n_j and QQC- n_j could be improved (while remaining tractable) by choosing among a few values of ℓ_j for every $j \in \llbracket m \rrbracket$ (see Remark 21).

6.2. Real data

In this section, we evaluate the performance of each algorithm (in terms of coverage and length of the returned prediction sets) on five public-domain regression data sets also considered in [54] and [56]: physicochemical properties of protein tertiary structure (bio) [51], bike sharing (bike) [16], communities and crimes (community) [52], Tennessee’s student teacher achievement ratio (star) [1], and concrete compressive strength (concrete) [64].

6.2.1. Setup

For each experiment, we split the full data set into three parts: a learning set (40%), a calibration set (40%), and a test set (20%). To simulate a FL scenario, we also split the calibration set in m disjoint subsets of equal size n . We consider scenarios where either $m > n$ or $m < n$, with $\max\{m/n, n/m\} \in \{4, 8\}$ depending on the data set; the exact values of (m, n) for each data set are given in Appendix J.3. All features are then standardized to have zero mean and unit variance. For each algorithm, we compute the empirical coverage obtained on the test set and the average length of the prediction set over the test set. These two metrics are collected over 50 different learning-calibration-test random splits.

Prediction sets are constructed using the Conformalized Quantile Regression method (CQR) [54], a popular variant of split CP directly compatible with our approach, following Remark 2 in Section 2.2. In CQR, \hat{f} is $(\hat{f}_{\alpha/2}, \hat{f}_{1-\alpha/2})$ where \hat{f}_δ is a quantile regressor of order δ [30] and $s(x, y) = \max(\hat{f}_{\alpha/2}(x) - y, y - \hat{f}_{1-\alpha/2}(x))$, so that the prediction set $\{y \in \mathbb{R} : s(x, y) \leq \hat{q}\} = [\hat{f}_{\alpha/2}(x) - \hat{q}, \hat{f}_{1-\alpha/2}(x) + \hat{q}]$ has a size adaptive to heteroscedasticity. In the following experiments $\hat{f}_{\alpha/2}$ and $\hat{f}_{1-\alpha/2}$ are quantile regression forests [44] built from the learning set only. The number of trees in the forest is set to 1000, the two parameters controlling the coverage rate on the learning data are tuned using cross-validation (within the learning set), and the remaining hyperparameters are set as done in [54].




FIG 7. Empirical coverages of prediction intervals ($\alpha = 0.1$) constructed by various methods, aggregated across all data sets. Our methods are shown in bold font. See the beginning of Section 6.2.2 for details. Left: $m > n$. Right: $m < n$.




FIG 8. Coverage (left) and average length (right) of prediction intervals for 50 random learning-calibration-test splits of the bike data set. The miscoverage is $\alpha = 0.1$, $\beta = 0.2$, and the calibration set is split into $m = 40$ disjoint subsets of equal size $n = 10$. See the beginning of Section 6.2.2 for details.

6.2.2. Results

Figure 7 displays the boxes of the empirical coverages obtained by each method over all the data sets and all the 50 different data splits (one point represents the empirical coverage obtained on one random split of one data set). Figures 8 and 9 show the empirical coverages as well as the lengths of the intervals obtained on the bike data set. All the results on other individual data sets are provided in Appendix J.3. For each box, the white circle indicates the mean, the left-end of the box corresponds to the empirical quantile of order $\beta = 0.2$ and the right-end to the empirical quantile of order $1 - \beta = 0.8$.

Marginally-valid algorithms. First, on average, CentralM, QQM and QQM-Fast all return intervals with coverage greater than $1 - \alpha = 0.9$ (the nominal coverage), without being too far from it. Importantly, our two one-shot FL methods return prediction sets with coverage and length close to those returned by split CP (CentralM). So, there is only a small loss when using our one-shot FL algorithms (whose coverage/length are slightly larger in terms of both expectation and dispersion). It is interesting to note that the results obtained with our two FL methods are quite similar, with slightly better results for QQM notably when $m > n$ (left panel of Figure 7). As QQM-Fast is much faster, it is a good alternative to QQM and can be preferred in real applications. Note finally that for marginal algorithms, the results for $m > n$ and $m < n$ are comparable.




FIG 9. Coverage (left) and average length (right) of prediction intervals for 50 random learning-calibration-test splits of the bike data set. The misscoverage is $\alpha = 0.1$, $\beta = 0.2$, and the calibration set is split into $m = 10$ disjoint subsets of equal size $n = 40$. See the beginning of Section 6.2.2 for details.

Conditionally-valid algorithms. The comparison between the three training conditionally valid algorithms on Figures 7–9 is similar to the comparison made above between the three marginally-valid algorithms. The main difference is the loss of using our fast one-shot FL algorithm QQC-Fast instead of QQC is larger than for our marginally-valid one-shot FL algorithms. Therefore, our advice here is to use QQC as long as it is computationally tractable, that is, when the total number mn of calibration data points is not huge—since its complexity might be proportional to mn .

Note that on Figures 7–9, the empirical β -quantile of the coverage of some training-conditionally-valid algorithms is slightly smaller than $1 - \alpha$. This apparent contradiction with the theoretical guarantees about these algorithms is due to the fact that the coverages reported in these experiments are only estimations obtained from finite test sets (200 points for the bike dataset, for instance).

Finally, as in Section 6.1, we can notice that the (empirical) marginal coverages of training-conditional algorithms are higher than those of the marginal algorithms.

FedCP-Avg. First, in our experiments, FedCP-Avg appears to be marginally valid for each data set, and training-conditionally valid at a confidence level β for all data sets except bike, community and star (each time when $m < n$). So, even without general theoretical guarantees (and even some counterexamples, see Appendix I.2), it here turns out to be valid in most cases (but not all, which is an issue). Second, FedCP-Avg is much more conservative than all competing algorithms in most data sets, and the few cases where it is less conservative than training-conditional algorithms precisely are the cases where FedCP-Avg is not training-conditionally valid. Therefore, FedCP-Avg appears to be clearly worse than our one-shot FL algorithms in all our experiments, a fact which is also supported by the few theoretical results presented in Appendix I.2.

Conclusion. Overall, these experiments support the fact that our one-shot FL methods are well-suited for building prediction sets in a federated setting, with only a mild loss compared to the centralized case.

7. Discussion

All prediction sets proposed in this paper are of the form $\widehat{\mathcal{C}}_{\bar{\ell}, \bar{k}}$ —as defined by Eq. (9)— with

$$(\bar{\ell}, \bar{k}) \in \operatorname{argmin}_{(\ell, k) \in \mathcal{E}} \{\operatorname{crit}(\ell, k)\},$$

where $\mathcal{E} \subset \llbracket n \rrbracket \times \llbracket m \rrbracket$ is a set of marginally (resp. training-conditionally) valid pairs, and $\operatorname{crit}(\ell, k)$ is an upper bound on the expectation (resp. some quantile) of the coverage $1 - \alpha_{\ell, k}(\mathcal{D})$ when the scores cdf is continuous. Note that we adopt here Assumption 1 to maintain notation simplicity, but the comments made in this paragraph also apply to Algorithms 5–6 in Section 5. The selection criteria (crit) used in Algorithms 1–6 come from arguments presented in Section 2.3 (and Appendix A), and they are validated by the theoretical and numerical results obtained in Sections 4 and 6. Yet, other choices are possible, leading to numerous variants of our algorithms. For instance, one could modify Algorithms 3–4 by taking

$$\operatorname{crit}(\ell, k) = F_{U_{(\ell, n, k, m)}}^{-1}(1 - \beta')$$

with any $\beta' \in (0, 1)$ —instead of fixing $\beta' = \beta$ —, leading to the algorithms respectively defined by Eq. (57) and (58) in Appendix E.2. Remarkably, theoretical coverage upper bounds similar to Theorem 17 can also be obtained for these variants, as proved by Eq. (59) and (60) in Appendix E.2. One could also use this criterion for selecting among marginally-valid pairs (ℓ, k) , or use $\operatorname{crit}(\ell, k) = M_{\ell, k}$ for selecting among training-conditionally-valid pairs (ℓ, k) . The function crit can also be chosen by the user —for instance, with a specific application in mind—, together with a set \mathcal{E} chosen among those used by our Algorithms 1–6, and the resulting prediction set would satisfy the corresponding distribution-free coverage lower bounds proved in this paper.

This work brings many possible future research directions. It would be interesting to investigate how our quantile-of-quantiles estimators can be adapted to the heterogeneous case, that is, when agents have data following different distributions. This could potentially extend the approaches of [39] and [50] to the one-shot setting. Another interesting line of research would be to propose and study differentially private versions of our algorithms. On a more theoretical aspect, it would be interesting to control the deviations of the coverage of marginal prediction sets built in Section 3.1, and conversely to control the expected coverage of the training-conditional prediction sets of Section 3.2. Finally, our paper focuses on the calibration step, making it particularly suited for split-based conformal methods. It would be interesting to study how our FL approach could be extended to full conformal prediction, and to the more general framework of nested conformal prediction [21].

References

- [1] ACHILLES, C. M., BAIN, H. P., BELLOTT, F., BOYD-ZAHARIAS, J., FINN, J., FOLGER, J., JOHNSTON, J. and WORD, E. (2008). Tennessee’s student teacher achievement ratio (STAR) project. *Harvard Dataverse* **1** 2008.
- [2] AL-HUSSAINI, E. K. (2012). Composition of cumulative distribution functions. *Journal of Statistical Theory and Applications ISSN* **1538** 7887.
- [3] ANDREW, G., THAKKAR, O., MCMAHAN, B. and RAMASWAMY, S. (2021). Differentially private learning with adaptive clipping. *Advances in Neural Information Processing Systems* **34** 17455–17466.

- [4] ANGELOPOULOS, A. N. and BATES, S. (2023). Conformal prediction: A gentle introduction. *Foundations and Trends® in Machine Learning* **16** 494–591.
- [5] BAGNOLI, M. and BERGSTROM, T. (2005). Log-concave probability and its applications. *Econom. Theory* **26** 445–469.
- [6] BAGNOLI, M. and BERGSTROM, T. C. (1989). Log-concave probability and its applications.
- [7] BALAKRISHNAN, N. (2007). Permanents, order statistics, outliers, and robustness. *Revista matemática complutense* **20** 7–107.
- [8] BIAN, M. and BARBER, R. F. (2023). Training-conditional coverage for distribution-free predictive inference. *Electronic Journal of Statistics* **17** 2044–2066.
- [9] CASTELLARES, F., LEMONTE, A. J. and SANTOS, M. A. (2021). A note on explicit expressions for moments of order statistics. *Spanish journal of statistics* **3** 45–53.
- [10] CONNOR, S. B. and FEWSTER, C. J. (2022). Integrals of incomplete beta functions, with applications to order statistics, random walks and string enumeration. *Brazilian Journal of Probability and Statistics* **36** 185–198.
- [11] CORDEIRO, G. M. (2013). A simple formula based on quantiles for the moments of beta generalized distributions. *Journal of Statistical Computation and Simulation* **83** 1932–1943.
- [12] CORMODE, G. and MARKOV, I. L. (2023). Federated Calibration and Evaluation of Binary Classifiers. *Proceedings of the VLDB Endowment* **16** 3253–3265.
- [13] DAVID, H. A. and NAGARAJA, H. N. (2004). *Order statistics*. John Wiley & Sons.
- [14] DENNIS, D. K., LI, T. and SMITH, V. (2021). Heterogeneity for the Win: One-Shot Federated Clustering. In *ICML*.
- [15] EL MEKKAOUI, K., MESQUITA, D., BLOMSTEDT, P. and KASKI, S. (2021). Federated stochastic gradient Langevin dynamics. In *Uncertainty in Artificial Intelligence* 1703–1712. PMLR.
- [16] FANAEE-T, H. and GAMA, J. (2013). Event labeling combining ensemble detectors and background knowledge. *Progress in Artificial Intelligence* 1–15.
- [17] FONTANA, M., ZENI, G. and VANTINI, S. (2023). Conformal prediction: a unified review of theory and new challenges. *Bernoulli* **29** 1–23.
- [18] GHOSH, J. K. (1971). A new proof of the Bahadur representation of quantiles and an application. *Ann. Math. Statist.* **42** 1957–1961.
- [19] GUAN, L. (2023). Localized conformal prediction: A generalized inference framework for conformal prediction. *Biometrika* **110** 33–50.
- [20] GUHA, N., TALWALKAR, A. and SMITH, V. (2019). One-shot federated learning. *arXiv preprint arXiv:1902.11175*.
- [21] GUPTA, C., KUCHIBHOTLA, A. K. and RAMDAS, A. (2022). Nested conformal prediction and quantile out-of-bag ensemble methods. *Pattern Recognition* **127** 108496.
- [22] HAN, X., TANG, Z., GHOSH, J. and LIU, Q. (2022). Split localized conformal prediction. *arXiv preprint arXiv:2206.13092*.
- [23] HONG, Y. (2013). On computing the distribution function for the Poisson binomial distribution. *Computational Statistics & Data Analysis* **59** 41–51.
- [24] HUBER, P. J. (1992). Robust estimation of a location parameter. In *Breakthroughs in statistics: Methodology and distribution* 492–518. Springer.
- [25] HUMBERT, P., LE BARS, B., BELLET, A. and ARLOT, S. (2023). One-Shot Federated Conformal Prediction. In *International Conference on Machine Learning* 14153–14177. PMLR.
- [26] JONES, M. C. (2004). Families of distributions arising from distributions of order statistics. *Test* **13** 1–43.
- [27] KAIROUZ, P., MCMAHAN, H. B., AVENT, B., BELLET, A., BENNIS, M., BHAGOJI, A. N., BONAWITZ, K., CHARLES, Z., CORMODE, G., CUMMINGS, R.

- et al. (2021). Advances and open problems in federated learning. *Foundations and Trends® in Machine Learning* **14** 1–210.
- [28] KARIMIREDDY, S. P., KALE, S., MOHRI, M., REDDI, S. J., STICH, S. U. and SURESH, A. T. (2020). SCAFFOLD: Stochastic Controlled Averaging for On-Device Federated Learning. In *ICML*.
- [29] KIVARANOVIC, D., JOHNSON, K. D. and LEEB, H. (2020). Adaptive, distribution-free prediction intervals for deep networks. In *International Conference on Artificial Intelligence and Statistics* 4346–4356. PMLR.
- [30] KOENKER, R. and BASSETT JR, G. (1978). Regression quantiles. *Econometrica: journal of the Econometric Society* 33–50.
- [31] KOTELEVSKII, N., VONO, M., DURMUS, A. and MOULINES, E. (2022). Fedpop: A bayesian approach for personalised federated learning. *Advances in Neural Information Processing Systems* **35** 8687–8701.
- [32] KRISHNAMOORTHY, K. and MATHEW, T. (2009). *Statistical tolerance regions: theory, applications, and computation*. John Wiley & Sons.
- [33] LAHIRI, S. N. (1992). On the Bahadur-Ghosh-Kiefer representation of sample quantiles. *Statist. Probab. Lett.* **15** 163–168.
- [34] LEBRUN, R. (2013). Efficient time/space algorithm to compute rectangular probabilities of multinomial, multivariate hypergeometric and multivariate Pólya distributions. *Statistics and Computing* **23** 615–623.
- [35] LEI, J., G’SSELL, M., RINALDO, A., TIBSHIRANI, R. J. and WASSERMAN, L. (2018). Distribution-free predictive inference for regression. *Journal of the American Statistical Association* **113** 1094–1111.
- [36] LI, Q., HE, B. and SONG, D. (2021). Practical One-Shot Federated Learning for Cross-Silo Setting. In *IJCAI*.
- [37] LI, T., SAHU, A. K., ZAHEER, M., SANJABI, M., TALWALKAR, A. and SMITH, V. (2020). Federated optimization in heterogeneous networks. *Proceedings of Machine Learning and Systems* **2** 429–450.
- [38] LU, C. and KALPATHY-CRAMER, J. (2021). Distribution-free federated learning with conformal predictions. *arXiv preprint arXiv:2110.07661*.
- [39] LU, C., YU, Y., KARIMIREDDY, S. P., JORDAN, M. and RASKAR, R. (2023). Federated Conformal Predictors for Distributed Uncertainty Quantification. In *International Conference on Machine Learning* 22942–22964. PMLR.
- [40] MAKGAI, S. L. et al. (2019). From Beta to Dirichlet Frontiers, PhD thesis, University of Pretoria.
- [41] MANOKHIN, V. (2022). Awesome Conformal Prediction.
- [42] MARCHAL, O. and ARBEL, J. (2017). On the sub-Gaussianity of the Beta and Dirichlet distributions.
- [43] MCMAHAN, B., MOORE, E., RAMAGE, D., HAMPSON, S. and Y ARCAS, B. A. (2017). Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics* 1273–1282. PMLR.
- [44] MEINSHAUSEN, N. and RIDGEWAY, G. (2006). Quantile regression forests. *Journal of machine learning research* **7**.
- [45] NDIAYE, E. (2022). Stable conformal prediction sets. In *International Conference on Machine Learning* 16462–16479. PMLR.
- [46] PAPADOPOULOS, H., PROEDROU, K., VOVK, V. and GAMMERMAN, A. (2002). Inductive confidence machines for regression. In *European Conference on Machine Learning* 345–356. Springer.
- [47] PARK, S., BASTANI, O., MATNI, N. and LEE, I. (2019). PAC Confidence Sets for Deep Neural Networks via Calibrated Prediction. In *International Conference on Learning Representations*.

- [48] PARK, S., DOBRIBAN, E., LEE, I. and BASTANI, O. (2021). PAC Prediction Sets Under Covariate Shift. In *International Conference on Learning Representations*.
- [49] PILLUTLA, K., LAGUEL, Y., MALICK, J. and HARCHAOU, Z. (2022). Differentially Private Federated Quantiles with the Distributed Discrete Gaussian Mechanism. In *International Workshop on Federated Learning: Recent Advances and New Challenges*.
- [50] PLASSIER, V., MAKNI, M., RUBASHEVSKII, A., MOULINES, E. and PANOV, M. (2023). Conformal Prediction for Federated Uncertainty Quantification Under Label Shift. In *Proceedings of the 40th International Conference on Machine Learning* (A. KRAUSE, E. BRUNSKILL, K. CHO, B. ENGELHARDT, S. SABATO and J. SCARLETT, eds.). *Proceedings of Machine Learning Research* **202** 27907–27947. PMLR.
- [51] RANA, P. (2013). Physicochemical properties of protein tertiary structure data set. *UCI Machine Learning Repository*.
- [52] REDMOND, M. (2011). Communities and crime unnormalized data set. *UCI Machine Learning Repository*.
- [53] RÉNYI, A. (1953). On the theory of order statistics. *Acta Math. Acad. Sci. Hung* **4** 48–89.
- [54] ROMANO, Y., PATTERSON, E. and CANDÈS, E. (2019). Conformalized quantile regression. *Advances in neural information processing systems* **32**.
- [55] SALEHKALEYBAR, S., SHARIF-NASSAB, A. and GOLESTANI, S. J. (2021). One-Shot Federated Learning: Theoretical Limits and Algorithms to Achieve Them. *Journal of Machine Learning Research* **22** 1–47.
- [56] SESIA, M. and ROMANO, Y. (2021). Conformal prediction using conditional histograms. *Advances in Neural Information Processing Systems* **34** 6304–6315.
- [57] SKORSKI, M. (2023). Bernstein-type bounds for beta distribution. *Modern Stochastics: Theory and Applications* **10** 211–228.
- [58] TANG, W. and TANG, F. (2023). The Poisson binomial distribution—Old & new. *Statistical Science* **38** 108–119.
- [59] TEMME, N. M. (1996). *Special functions: An introduction to the classical functions of mathematical physics*. John Wiley & Sons.
- [60] VIRTANEN, P., GOMMERS, R., OLIPHANT, T. E., HABERLAND, M., REDDY, T., COURNAPEAU, D., BUROVSKI, E., PETERSON, P., WECKESSER, W., BRIGHT, J., VAN DER WALT, S. J., BRETT, M., WILSON, J., MILLMAN, K. J., MAYOROV, N., NELSON, A. R. J., JONES, E., KERN, R., LARSON, E., CAREY, C. J., POLAT, İ., FENG, Y., MOORE, E. W., VANDERPLAS, J., LAXALDE, D., PERKTOLD, J., CIMRMAN, R., HENRIKSEN, I., QUINTERO, E. A., HARRIS, C. R., ARCHIBALD, A. M., RIBEIRO, A. H., PEDREGOSA, F., VAN MULBREGT, P. and SciPy 1.0 CONTRIBUTORS (2020). SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods* **17** 261–272.
- [61] VOVK, V. (2012). Conditional validity of inductive conformal predictors. In *Asian conference on machine learning* 475–490. PMLR.
- [62] VOVK, V., GAMMERMAN, A. and SHAFER, G. (2005). *Algorithmic learning in a random world*. Springer Science & Business Media.
- [63] WILKS, S. S. (1941). Determination of sample sizes for setting tolerance limits. *The Annals of Mathematical Statistics* **12** 91–96.
- [64] YEH, I.-C. (1998). Modeling of strength of high-performance concrete using artificial neural networks. *Cement and Concrete research* **28** 1797–1808.
- [65] YUROCHKIN, M., AGARWAL, M., GHOSH, S., GREENEWALD, K., HOANG, T. N. and KHAZAENI, Y. (2019). Bayesian Nonparametric Federated Learning of Neural Networks. In *ICML*.
- [66] ZHANG, Y., WAINWRIGHT, M. J. and DUCHI, J. C. (2012). Communication-efficient

algorithms for statistical optimization. *Advances in neural information processing systems* **25**.

Appendix A: Prediction set performance measure: complements

Recall that we prove in Section 2.3 that for every (ℓ, k) , the size of $\widehat{\mathcal{C}}_{\ell,k}(x)$ is a nondecreasing function of the coverage $1 - \alpha_{\ell,k}(\mathcal{D})$.

As a consequence, for any $\beta \in (0, 1)$, minimizing the $(1 - \beta)$ -quantile of the coverage $1 - \alpha_{\ell,k}(\mathcal{D})$ yields a minimizer of the $(1 - \beta)$ -quantile of the size of $\widehat{\mathcal{C}}_{\ell,k}(x)$, simultaneously for all $x \in \mathcal{X}$. We therefore use this strategy for building the two training-conditionally-valid algorithms described in Section 3.2.

The problem is a bit more difficult when one measures the performance of a predictor set by its size *on average over* \mathcal{D} (even at a single $x \in \mathcal{X}$). As an illustration, let us consider regression with the fitted absolute residual score function. Then,

$$\forall x \in \mathcal{X}, \quad \widehat{\mathcal{C}}_{\ell,k}(x) = [\widehat{f}(x) - S_{(\ell,k)}, \widehat{f}(x) + S_{(\ell,k)}]$$

is a prediction interval of length

$$2S_{(\ell,k)} = 2F_S^{-1}(1 - \alpha_{\ell,k}(\mathcal{D}))$$

almost surely. In this example, the size of $\widehat{\mathcal{C}}_{\ell,k}(x)$ (its length) is a nondecreasing function of the coverage $1 - \alpha_{\ell,k}(\mathcal{D})$, but the relationship between size and coverage is highly non-linear in general, hence the expected size $\mathbb{E}[\widehat{\mathcal{C}}_{\ell,k}(x)]$ cannot easily be linked to the expected coverage $\mathbb{E}[1 - \alpha_{\ell,k}(\mathcal{D})]$. Nevertheless, following what is usually done for split CP (see Section 2.1), choosing (ℓ, k) by minimizing $\mathbb{E}[1 - \alpha_{\ell,k}(\mathcal{D})]$ —or at least its value when F_S is continuous, see Sections 3.1 and 2.1— is a natural choice, which is efficient according to numerical experiments.

The latter argument might seem questionable, especially when F_S is not guaranteed to be continuous. Yet, the strategy we propose is reasonable for the following reasons. $S_{(\ell,k)}$ is a nondecreasing function of (ℓ, k) (for the lexicographic order on $\llbracket n \rrbracket \times \llbracket m \rrbracket$), hence both the expected size and the expected coverage are nondecreasing functions of (ℓ, k) . Therefore, minimizing the value $M_{\ell,k}$ of the expected coverage $\mathbb{E}[1 - \alpha_{\ell,k}(\mathcal{D})]$ when F_S is continuous (see Section 3.1) over the set of (marginally or conditionally) valid pairs (ℓ, k) yields some (ℓ, k) minimal for the lexicographic order in the set of valid pairs. So, even if this does not choose the optimal valid pair (ℓ, k) , it yields a performance close to its optimal value —at least numerically. We therefore follow this strategy when building marginally-valid algorithms in Section 3.1.

Appendix B: Key preliminary results on order statistics

In this section, we provide some known and new important results about order statistics that play a key role in our proofs. We refer to Appendix G for proofs and additional results, and to [13] for an in-depth presentation on this topic.

B.1. Order statistics

Let us first introduce some notation. Given a real-valued random variable Z with arbitrary cumulative distribution function (cdf) F_Z , its quantile function is defined as the generalized inverse of F_Z :

$$\forall p \in (0, 1), \quad F_Z^{-1}(p) := \inf\{x \in \mathbb{R} : F_Z(x) \geq p\}. \quad (35)$$

A key property of the generalized inverse —straightforward from Eq. (35)— is that

$$\forall p \in [0, 1], \quad F_Z \circ F_Z^{-1}(p) \geq p \quad \text{with equality if} \quad p \in \text{Im}(F_Z). \quad (36)$$

In particular, when F_Z is continuous, Eq. (36) is an equality for every $p \in [0, 1] = \text{Im}(F_Z)$. Given a sample Z_1, \dots, Z_N , we denote by $Z_{(1:N)} \leq \dots \leq Z_{(N:N)}$ the corresponding order statistics, so that $Z_{(i:N)} = \widehat{Q}_{(i)}(Z_1, \dots, Z_N)$ for every $i \in \llbracket N \rrbracket$, using the notation defined by Eq. (4) in Section 2.1. When the sample size is clear from context, we note $Z_{(i)} := Z_{(i:N)}$.

We first recall the well-known following link between general order statistics and uniform order statistics (see e.g. [13]).

Lemma 23. *For any $N \geq 1$, let U_1, \dots, U_N be independent random variables with uniform distribution over $[0, 1]$, and Z_1, \dots, Z_N be independent and identically distributed real-valued random variables with common cdf F_Z . For any $r \in \llbracket N \rrbracket$, let $U_{(r)} = U_{(r:N)}$ and $Z_{(r)} = Z_{(r:N)}$ respectively denote the corresponding r -th order statistics. Then, we have*

$$F_Z(Z_{(r)}) \stackrel{d}{=} F_Z \circ F_Z^{-1}(U_{(r)}) \geq U_{(r)}, \quad \text{hence} \quad F_Z(Z_{(r)}) \succeq U_{(r)}, \quad (37)$$

that is, $F_Z(Z_{(r)})$ stochastically dominates $U_{(r)}$. In particular, Eq. (37) implies that

$$\forall a \in \mathbb{R}, \quad \mathbb{P}(a \leq F_Z(Z_{(r)})) \geq \mathbb{P}(a \leq U_{(r)}) \quad \text{and} \quad \mathbb{E}[F_Z(Z_{(r)})] \geq \mathbb{E}[U_{(r)}]. \quad (38)$$

Furthermore, when F_Z is continuous,

$$F_Z(Z_{(r)}) \stackrel{d}{=} U_{(r)} \quad (39)$$

both follow the $\text{Beta}(r, N - r + 1)$ distribution, and they have the same expectation and cdf.

Lemma 23 is proved in Appendix H.1 for completeness. It shows that concentration/deviation inequalities for order statistics can be obtained in general from results on uniform order statistics. Therefore, any result about the Beta distribution can be used to bound $F_Z(Z_{(r)})$ —hence the bounds on the coverage of (centralized) split CP in Section 2.1, for instance.

We now present non-asymptotic bounds on the quantile function of uniform order statistics, which rely on a concentration result for the Beta distribution [42, Theorem 1].

Proposition 24. *For any integer $N \geq 1$ and $r \in \llbracket N \rrbracket$, the quantile function $F_{U_{(r:N)}}^{-1}$ of the r -th order statistics of a sample of N independent standard uniform variables satisfies the following inequalities, for any $\delta > 0$:*

$$F_{U_{(r:N)}}^{-1}(\delta) \geq \frac{r}{N+1} - \sqrt{\frac{\log(1/\delta)}{2(N+2)}} \quad (40)$$

$$\text{and} \quad F_{U_{(r:N)}}^{-1}(1-\delta) \leq \frac{r}{N+1} + \sqrt{\frac{\log(1/\delta)}{2(N+2)}}. \quad (41)$$

Proposition 24 is proved in Appendix H.2.

Finally, our coverage upper bounds (Theorems 16 and 17) rely on the following non-asymptotic upper bound on the increments of the quantile function $F_{U_{(\ell:n)}}^{-1}$ of the $\text{Beta}(\ell, n - \ell + 1)$ distribution, showing that it is locally $\mathcal{O}(n^{-1/2})$ -Lipschitz.

Theorem 25. *For any $n \geq 1$, $\ell \in \llbracket n \rrbracket$, $\varepsilon \in (0, 1)$, and $x \in (0, 1 - \varepsilon)$, we have*

$$0 \leq F_{U_{(\ell:n)}}^{-1}(x + \varepsilon) - F_{U_{(\ell:n)}}^{-1}(x) \leq \frac{\varepsilon}{\sqrt{n+2}} \frac{\sqrt{2}}{g(\min\{x, 1 - x - \varepsilon\})}, \quad (42)$$

$$\text{where } \forall t \in \left(0, \frac{1}{2}\right], \quad g(t) := \sup_{\delta \in (0, \min\{t, 1-t\})} \frac{t - \delta}{\sqrt{\log(1/\delta)}} \geq \frac{t}{2\sqrt{\log(2/t)}}$$

$$\text{and } \forall t \in [1/3, 1/2], \quad \frac{\sqrt{2}}{g(t)} \leq 9.$$

Theorem 25 is proved in Appendix H.4. To the best of our knowledge, such an upper bound has never been proved previously. It is a key result of the paper, used in the proof of our coverage upper bounds.

Remark 26. The upper bound in Eq. (42) is always well defined since g is defined on $(0, 1/2]$ and $\min\{x, 1-x-\varepsilon\} \leq 1/2$ for every $\varepsilon \in (0, 1)$ and $x \in (0, 1-\varepsilon)$.

B.2. Order statistics of order statistics

The federated learning methods of Section 3 are all based on the quantile-of-quantiles family of estimators (Definition 4 in Section 2.2), which are order statistics of order statistics, so we focus on them in this subsection.

Given a collection $(Z_{i,j})_{1 \leq i \leq n, 1 \leq j \leq m}$ of m samples of size n , for every $j \in \llbracket m \rrbracket$, we denote the order statistics of the j -th sample $(Z_{i,j})_{1 \leq i \leq n}$ by

$$Z_{(1:n),j} \leq Z_{(2:n),j} \leq \cdots \leq Z_{(n:n),j}.$$

Then, for every $\ell \in \llbracket n \rrbracket$, we denote the order statistics of the sample of the ℓ -th order statistics $(Z_{(\ell:n),j})_{1 \leq j \leq m}$ by

$$Z_{(\ell:n),(1:m)} \leq Z_{(\ell:n),(2:m)} \leq \cdots \leq Z_{(\ell:n),(m:m)}.$$

For every $\ell \in \llbracket n \rrbracket$ and $k \in \llbracket m \rrbracket$, $Z_{(\ell:n),(k:m)} =: Z_{(\ell,n,k:m)}$ is the (ℓ, k) -th order statistics of order statistics of $(Z_{i,j})_{1 \leq i \leq n, 1 \leq j \leq m}$, which can also be written as follows with the notation defined by Eq. (4) in Section 2.1:

$$Z_{(\ell:n,k:m)} = \widehat{Q}_{(k)} \left(\left(\widehat{Q}_{(\ell)}(\mathcal{Z}_j) \right)_{j \in \llbracket m \rrbracket} \right)$$

$$\text{where } \forall j \in \llbracket m \rrbracket, \quad \mathcal{Z}_j := (Z_{i,j})_{i \in \llbracket n \rrbracket} \quad \text{so that} \quad Z_{(\ell:n),j} = \widehat{Q}_{(\ell)}(\mathcal{Z}_j).$$

In the following, we write $Z_{(\ell),j}$ (resp. $Z_{(\ell,k)}$) instead of $Z_{(\ell:n),j}$ (resp. $Z_{(\ell:n,k:m)}$) when the sample sizes n and m are clear from context.

Our first key result is a link with order statistics of uniform order statistics, similar to Lemma 23.

Lemma 27. For any $n, m \geq 1$, let $(U_{i,j})_{1 \leq i \leq n, 1 \leq j \leq m}$ be independent random variables with uniform distribution over $[0, 1]$, and $(Z_{i,j})_{1 \leq i \leq n, 1 \leq j \leq m}$ be independent and identically distributed real-valued random variables with common cdf F_Z . For every $\ell \in \llbracket n \rrbracket$ and $k \in \llbracket m \rrbracket$, let $U_{(\ell,k)}$ and $Z_{(\ell,k)}$ respectively denote the corresponding (ℓ, k) -th order statistics of order statistics, as defined at the beginning of Appendix B.2. Then, we have

$$F_Z(Z_{(\ell,k)}) \stackrel{d}{=} F_Z \circ F_Z^{-1}(U_{(\ell,k)}) \geq U_{(\ell,k)}, \quad \text{hence } F_Z(Z_{(\ell,k)}) \succeq U_{(\ell,k)}, \quad (43)$$

that is, $F_Z(Z_{(\ell,k)})$ stochastically dominates $U_{(\ell,k)}$. In particular, Eq. (43) implies that

$$\forall a \in \mathbb{R}, \quad \mathbb{P}(a \leq F_Z(Z_{(\ell,k)})) \geq \mathbb{P}(a \leq U_{(\ell,k)}) \quad \text{and} \quad \mathbb{E}[F_Z(Z_{(\ell,k)})] \geq \mathbb{E}[U_{(\ell,k)}]. \quad (44)$$

Furthermore, when F_Z is continuous,

$$F_Z(Z_{(\ell,k)}) \stackrel{d}{=} U_{(\ell,k)} \quad (45)$$

hence they have the same expectation and cdf.

Lemma 27 is proved in Appendix H.5. It shows that concentration/deviation inequalities for order statistics of order statistics can be obtained in general from results on order statistics of uniform order statistics. Therefore, we now focus on the distribution of $U_{(\ell:n,k:m)}$.

By definition, $U_{(\ell:n,k:m)}$ is the k -th order statistics of a sample of m independent and identically distributed random variables following a $\text{Beta}(\ell, n - \ell + 1)$ distribution. This is called a Beta order statistics, and it follows a Beta-Beta distribution [2, 10, 26, 40]. The Beta-Beta distribution is difficult to analyze directly [9, 11] —as one can see from the formula of its expectation $M_{\ell,k}$ provided by Theorem 5 in Section 3.1—, so it is more convenient to first relate it to the well-known Beta distribution, as done by the next result.

Lemma 28. *Let $n, m \geq 1$ be two integers, $k \in \llbracket m \rrbracket$, $\ell \in \llbracket n \rrbracket$. Let $U_{(\ell:n)}$ and $U_{(k:m)}$ be defined as in Lemma 23, and $U_{(\ell:n,k:m)} = U_{(\ell,k)}$ as in Lemma 27. Then, their cdfs satisfy*

$$F_{U_{(\ell:n,k:m)}} = F_{U_{(k:m)}} \circ F_{U_{(\ell:n)}}. \quad (46)$$

Lemma 28 is proved in Appendix H.6. Since the cdfs of $U_{(\ell:n)}$ and $U_{(k:m)}$ are known (see Lemma 23), it shows that the cdf of $U_{(\ell:n,k:m)}$ is the composition of the cdf of a $\text{Beta}(k, m - k + 1)$ with the cdf of a $\text{Beta}(\ell, n - \ell + 1)$.

Appendix C: Proofs of Section 3.1 (marginal guarantees)

C.1. Proof of Theorem 5

Throughout the proof, we reason conditionally to \mathcal{D}^{lrn} . Since $(X_{i,j}, Y_{i,j})_{1 \leq i \leq n, 1 \leq j \leq m}$, (X, Y) are i.i.d., the associated scores $(S_{i,j})_{1 \leq i \leq n, 1 \leq j \leq m}$, S are i.i.d. Let us denote by F_S their cdf (given \mathcal{D}^{lrn}).

By definition of $\widehat{C}_{\ell,k}(X)$, we have $\mathbb{P}(Y \in \widehat{C}_{\ell,k}(X)) = \mathbb{P}(S \leq S_{(\ell,k)}) = \mathbb{E}[F_S(S_{(\ell,k)})]$. Furthermore, by Eq. (44) in Lemma 27 with $Z_{i,j} = S_{i,j}$ (hence $F_Z = F_S$), we have

$$\mathbb{E}[F_S(S_{(\ell,k)})] \geq \mathbb{E}[U_{(\ell,k)}] =: M_{\ell,k} \quad (47)$$

in general, with equality when F_S is continuous —or equivalently when the scores are a.s. distinct— by Eq. (45) in Lemma 27. It remains to prove the formula for $M_{\ell,k}$ announced in Theorem 5.

Recall that by definition, $U_{(\ell,k)} = U_{(\ell:n,k:m)}$ is the k -th order statistics of a sample $(U_{(\ell:n),j})_{1 \leq j \leq m}$ of m independent random variables with common distribution $\text{Beta}(\ell, n - \ell + 1)$. Therefore, its pdf is given by Eq. (64) in Appendix G.1 with $r = k$, $N = m$, $f_Z = f_{U_{(\ell:n)}}$ and $F_Z = F_{U_{(\ell:n)}}$. Using Eq. (65) with $r = \ell$ and $N = n$, we get that for every $t \in \mathbb{R}$,

$$\begin{aligned} f_{U_{(\ell,k)}}(t) &= \frac{m!}{(k-1)!(m-k)!} F_{U_{(\ell)}}(t)^{k-1} [1 - F_{U_{(\ell)}}(t)]^{m-k} f_{U_{(\ell)}}(t) \\ &= k \binom{m}{k} F_{U_{(\ell)}}(t)^{k-1} [1 - F_{U_{(\ell)}}(t)]^{m-k} f_{U_{(\ell)}}(t) \\ &= k \binom{m}{k} \left[\sum_{i=\ell}^n \binom{n}{i} t^i (1-t)^{n-i} \right]^{k-1} \left[1 - \sum_{i=\ell}^n \binom{n}{i} t^i (1-t)^{n-i} \right]^{m-k} f_{U_{(\ell)}}(t) \\ &= k \binom{m}{k} \left[\sum_{i=\ell}^n \binom{n}{i} t^i (1-t)^{n-i} \right]^{k-1} \left[\sum_{i=0}^{\ell-1} \binom{n}{i} t^i (1-t)^{n-i} \right]^{m-k} f_{U_{(\ell)}}(t) \end{aligned}$$

$$\text{since } 1 = \sum_{i=0}^n \binom{n}{i} t^i (1-t)^{n-i} = \sum_{i=0}^{\ell-1} \binom{n}{i} t^i (1-t)^{n-i} + \sum_{i=\ell}^n \binom{n}{i} t^i (1-t)^{n-i}.$$

Then, by developing the powers over the summations, we get that for every $t \in [0, 1]$,

$$\begin{aligned} & f_{U_{(\ell,k)}}(t) \\ &= k \binom{m}{k} \sum_{i_1=\ell}^n \cdots \sum_{i_{k-1}=\ell}^n \sum_{i_{k+1}=0}^{\ell-1} \cdots \sum_{i_m=0}^{\ell-1} \binom{n}{i_1} \cdots \binom{n}{i_{k-1}} \binom{n}{i_{k+1}} \cdots \binom{n}{i_m} \\ & \quad \times t^{i_1+\cdots+i_{k-1}+i_{k+1}+\cdots+i_m} (1-t)^{n(m-1)-(i_1+\cdots+i_{k-1}+i_{k+1}+\cdots+i_m)} f_{U_{(\ell)}}(t) \\ &= \frac{k \binom{m}{k}}{\text{B}(\ell, n-\ell+1)} \cdot \sum_{i_1=\ell}^n \cdots \sum_{i_{k-1}=\ell}^n \sum_{i_{k+1}=0}^{\ell-1} \cdots \sum_{i_m=0}^{\ell-1} \binom{n}{i_1} \cdots \binom{n}{i_{k-1}} \binom{n}{i_{k+1}} \cdots \binom{n}{i_m} \\ & \quad \times t^{i_1+\cdots+i_{k-1}+i_{k+1}+\cdots+i_m+\ell-1} (1-t)^{n(m-1)-(i_1+\cdots+i_{k-1}+i_{k+1}+\cdots+i_m)+n-\ell}, \end{aligned}$$

where in the last equality we used the definition of $f_{U_{(\ell)}}$. We now note that

$$\begin{aligned} & \int_0^1 t \cdot t^{i_1+\cdots+i_{k-1}+i_{k+1}+\cdots+i_m+\ell-1} (1-t)^{n(m-1)-(i_1+\cdots+i_{k-1}+i_{k+1}+\cdots+i_m)+n-\ell} dt \\ &= \int_0^1 t^{i_1+\cdots+i_{k-1}+\ell+i_{k+1}+\cdots+i_m} (1-t)^{mn-(i_1+\cdots+i_{k-1}+\ell+i_{k+1}+\cdots+i_m)} dt \\ &= \text{B}(i_1+\cdots+i_{k-1}+\ell+i_{k+1}+\cdots+i_m+1, mn-(i_1+\cdots+i_{k-1}+\ell+i_{k+1}+\cdots+i_m)+1) \\ &= \left[(mn+1) \binom{mn}{i_1+\cdots+i_{k-1}+\ell+i_{k+1}+\cdots+i_m} \right]^{-1}, \end{aligned}$$

since $\text{B}(b+1, a-b+1) = [(a+1) \binom{a}{b}]^{-1}$ for every $a, b \in \mathbb{N}$ such that $a \geq b$, here used with $a = mn$ and $b = (i_1+\cdots+i_{k-1}+\ell+i_{k+1}+\cdots+i_m)$. We end the proof using the definition of the expectation (and the fact that $U_{(\ell,k)}$ is supported by $[0, 1]$):

$$\begin{aligned} \mathbb{E}[U_{(\ell,k)}] &= \int_0^1 t \cdot f_{U_{(\ell,k)}}(t) dt \\ &= \frac{k \binom{m}{k} \sum_{i_1=\ell}^n \cdots \sum_{i_{k-1}=\ell}^n \sum_{i_{k+1}=0}^{\ell-1} \cdots \sum_{i_m=0}^{\ell-1} \frac{\binom{n}{i_1} \cdots \binom{n}{i_{k-1}} \binom{n}{i_{k+1}} \cdots \binom{n}{i_m}}{\binom{mn}{i_1+\cdots+i_{k-1}+\ell+i_{k+1}+\cdots+i_m}}}{(mn+1) \text{B}(\ell, n-\ell+1)}. \end{aligned}$$

□

Remark 29. *The first part of the proof of Theorem 5 —proving Eq. (11) in general, and showing that it is an equality when F_S is continuous— is identical to the first part of the proof of [25, Theorem 3.2]. The second part of the proof is new: we here compute $\mathbb{E}[U_{(\ell,k)}]$ from its pdf —contrary to [25, Theorem 3.2] which uses its cdf—, which leads to a slightly simpler formula for $M_{\ell,k}$.*

C.2. Proof of Lemma 7

The argmin defining (ℓ^*, k^*) in Algorithm 1 is non-empty if and only if

$$\max_{(\ell,k) \in \llbracket n \rrbracket \times \llbracket m \rrbracket} M_{\ell,k} \geq 1 - \alpha.$$

Since $M_{\ell,k} = \mathbb{E}[U_{(\ell:n,k:m)}]$ —see Eq. (47) in Appendix C.1—, it is a nondecreasing function of (ℓ, k) (for the lexicographic order), hence

$$\max_{(\ell,k) \in \llbracket n \rrbracket \times \llbracket m \rrbracket} M_{\ell,k} = M_{n,m} = \mathbb{E}[U_{(n:n,m:m)}] = \mathbb{E}\left[\max_{(i,j) \in \llbracket n \rrbracket \times \llbracket m \rrbracket} U_{i,j}\right]$$

where the $(U_{i,j})_{(i,j) \in \llbracket n \rrbracket \times \llbracket m \rrbracket}$ are independent standard uniform random variables, hence

$$M_{n,m} = \mathbb{E}[U_{(mn:mn)}] = \frac{mn}{mn+1}$$

by Eq. (69) with $r = N = mn$. Finally, we have proved that the argmin defining (ℓ^*, k^*) in Algorithm 1 is non-empty if and only if

$$\frac{mn}{mn+1} = 1 - \frac{1}{mn+1} \geq 1 - \alpha,$$

which is equivalent to $mn+1 \geq \alpha^{-1}$. \square

C.3. Proof of Proposition 8

Recall that by Eq. (47) in Appendix C.1, $M_{\ell,k} = \mathbb{E}[U_{(\ell:n,k:m)}]$ where $U_{(\ell:n,k:m)}$ is the k -th order statistics of a sample $(U_{(\ell:n),j})_{1 \leq j \leq m}$ of m independent random variables with common distribution $\text{Beta}(\ell, n - \ell + 1)$ and cdf $F_{U_{(\ell:n)}}$. Since $\ell \geq 1$ and $n - \ell + 1 \geq 1$, the function $1 - F_{U_{(\ell:n)}}$ is log-concave [5, Section 6.3 and Theorem 3] and thus [13, Equation (4.5.7)] shows that

$$\begin{aligned} F_{U_{(\ell:n)}}(\mathbb{E}[U_{(\ell:n,k:m)}]) &= F_{U_{(\ell:n)}}\left(\mathbb{E}[\widehat{Q}_{(k)}(U_{(\ell:n),1}, \dots, U_{(\ell:n),m})]\right) \\ &\leq 1 - \exp\left(-\sum_{i=0}^{k-1} \frac{1}{m-i}\right) < \frac{k}{m+1/2}, \end{aligned} \quad (48)$$

hence the upper bound of Eq. (12).

For the lower bound, we first remark that $\mathbb{E}[U_{(\ell,k)}] = 1 - \mathbb{E}[U_{(n-\ell+1,m-k+1)}]$. Indeed, defining $\tilde{U}_{i,j} = 1 - U_{i,j}$ for every $i, j \in \llbracket n \rrbracket \times \llbracket m \rrbracket$, on the one hand, we have $\tilde{U}_{(\ell,k)} = 1 - U_{(n-\ell+1,m-k+1)}$, and on the other hand, the fact that $(\tilde{U}_{i,j})_{i,j \in \llbracket n \rrbracket \times \llbracket m \rrbracket} \stackrel{d}{=} (U_{i,j})_{i,j \in \llbracket n \rrbracket \times \llbracket m \rrbracket}$ implies that $U_{(\ell,k)} \stackrel{d}{=} \tilde{U}_{(\ell,k)}$. So, we have proved that

$$U_{(\ell:n,k:m)} \stackrel{d}{=} 1 - U_{((n-\ell+1):n,(m-k+1):m)}$$

hence the result by taking the expectation. Therefore, using the upper bound of Eq. (12), we get that

$$M_{\ell,k} = 1 - M_{n-\ell+1,m-k+1} > 1 - F_{U_{(n-\ell+1:n)}}^{-1}\left(\frac{m-k+1}{m+1/2}\right) = F_{U_{(\ell:n)}}^{-1}\left(\frac{k-1/2}{m+1/2}\right)$$

since $F_{U_{(n-\ell+1:n)}}^{-1}(x) = 1 - F_{U_{(\ell:n)}}^{-1}(1-x)$, again by a symmetry argument ($U_{(\ell:n)} \stackrel{d}{=} 1 - U_{(n-\ell+1:n)}$) [see also 59, about this property of the incomplete Beta function]. \square

C.4. Proof of Lemma 9

The argmin defining $\tilde{\ell}$ in Algorithm 2 is non-empty if and only if

$$\min_{\ell \in \llbracket n \rrbracket} \tilde{k}_{m,n}(\ell, \alpha) \leq m.$$

Since $F_{U_{(\ell,n)}}(1 - \alpha) = \mathbb{P}(U_{(\ell,n)} \leq 1 - \alpha)$, it is a nonincreasing function of ℓ , hence, by Eq. (13),

$$\begin{aligned} \min_{\ell \in \llbracket n \rrbracket} \tilde{k}_{m,n}(\ell, \alpha) &= \tilde{k}_{m,n}(n, \alpha) = \lceil (m + 1/2) \cdot F_{U_{(n,n)}}(1 - \alpha) + 1/2 \rceil \\ &= \lceil (m + 1/2)(1 - \alpha)^n + 1/2 \rceil, \end{aligned}$$

using Eq. (65) in Appendix G.1. This directly leads to the necessary and sufficient condition (14).

Note that Eq. (14) can be rewritten

$$\left(1 + \frac{1}{m - 1/2}\right)^{1/n} \leq \frac{1}{1 - \alpha}.$$

The function $x \in [0, +\infty) \mapsto (1 + x)^{1/n}$ being concave, it is below its first-order approximation at $x = 0$, so that

$$\left(1 + \frac{1}{m - 1/2}\right)^{1/n} \leq 1 + \frac{1}{n(m - 1/2)}$$

and, if $n(m - 1/2) \geq \alpha^{-1} - 1$, this upper bound is smaller than

$$1 + \frac{1}{\alpha^{-1} - 1} = \frac{\alpha^{-1}}{\alpha^{-1} - 1} = \frac{1}{1 - \alpha}.$$

□

Appendix D: Proofs of Section 3.2 (training-conditional guarantees)

D.1. Proof of Theorem 10

As previously, we make this proof conditionally to \mathcal{D}^{lrn} without writing it explicitly in the following. Since $(X_{i,j}, Y_{i,j})_{1 \leq i \leq n, 1 \leq j \leq m}$, (X, Y) are i.i.d., the associated scores $(S_{i,j})_{1 \leq i \leq n, 1 \leq j \leq m}$, S are i.i.d. We denote by F_S their cdf (given \mathcal{D}^{lrn}). Now, notice that

$$1 - \alpha_{\ell,k}(\mathcal{D}) := \mathbb{P}(Y \in \hat{\mathcal{C}}_{\ell,k}(X) \mid \mathcal{D}) = \mathbb{P}(S \leq S_{(\ell,k)} \mid (S_{i,j})_{1 \leq i \leq n, 1 \leq j \leq m}) = F_S(S_{(\ell,k)}).$$

Therefore, applying Lemma 27 with $F_Z = F_S$, Eq. (44) with $a = F_{U_{(\ell,k)}}^{-1}(\beta)$ shows that

$$\begin{aligned} \mathbb{P}(1 - \alpha_{\ell,k}(\mathcal{D}) \geq F_{U_{(\ell,k)}}^{-1}(\beta)) &= \mathbb{P}(F_S(S_{(\ell,k)}) \geq F_{U_{(\ell,k)}}^{-1}(\beta)) \\ &\geq \mathbb{P}(U_{(\ell,k)} \geq F_{U_{(\ell,k)}}^{-1}(\beta)) = 1 - \beta, \end{aligned} \quad (49)$$

that is, Eq. (17) holds true. The formula for $F_{U_{(\ell,k)}}^{-1}$ follows from Lemma 28. When the scores are a.s. distinct—or equivalently when F_S is continuous—, Eq. (45) in Lemma 27 shows that $1 - \alpha_{\ell,k}(\mathcal{D}) \stackrel{d}{=} U_{(\ell,k)}$ hence Eq. (49) is an equality and Eq. (18) holds true. □

D.2. Proof of Lemma 12

The argmin defining (ℓ^*, k_c^*) in Algorithm 3 is non-empty if and only if

$$\max_{(\ell, k) \in \llbracket n \rrbracket \times \llbracket m \rrbracket} F_{U_{(\ell:n, k:m)}}^{-1}(\beta) \geq 1 - \alpha.$$

Since $U_{(\ell:n, k:m)}$ is almost surely a nondecreasing function of (ℓ, k) for the lexicographic order, $F_{U_{(\ell:n, k:m)}}^{-1}(\beta)$ is a nondecreasing function of (ℓ, k) , hence

$$\max_{(\ell, k) \in \llbracket n \rrbracket \times \llbracket m \rrbracket} F_{U_{(\ell:n, k:m)}}^{-1}(\beta) = F_{U_{(n:n, m:m)}}^{-1}(\beta) = F_{U_{(nm:nm)}}^{-1}(\beta) = \beta^{1/(mn)}$$

by Eq. (66) in Appendix G.1. So, (ℓ_c^*, k_c^*) is well defined if and only if $\beta^{1/(mn)} \geq 1 - \alpha$. The conclusion follows. \square

D.3. Proof of Proposition 13

By Eq. (46) in Lemma 28, $F_{U_{(\ell:n, k:m)}}^{-1}(\beta) = F_{U_{(\ell:n)}}^{-1} \circ F_{U_{(k:m)}}^{-1}(\beta)$. Furthermore, Eq. (40) in Proposition 24 with $N = m$, $r = k$ and $\delta = \beta$ shows that

$$F_{U_{(k:m)}}^{-1}(\beta) \geq \frac{k}{m+1} - \sqrt{\frac{\log(1/\beta)}{2(m+2)}}.$$

This proves Eq. (21) since $F_{U_{(\ell:n)}}^{-1}$ is nondecreasing. \square

D.4. Proof of Lemma 14

The argmin defining $\tilde{\ell}^{\text{cond}}$ in Algorithm 4 is non-empty if and only if

$$\min_{\ell \in \llbracket n \rrbracket} \tilde{k}_{m,n}^{\text{cond}}(\ell, \alpha, \beta) \leq m.$$

Since $F_{U_{(\ell:n)}}(1 - \alpha)$ is a nonincreasing function of ℓ , using also Eq. (66) in Appendix G.1, we have

$$\min_{\ell \in \llbracket n \rrbracket} \tilde{k}_{m,n}^{\text{cond}}(\ell, \alpha, \beta) = \tilde{k}_{m,n}^{\text{cond}}(n, \alpha, \beta) = \left\lceil (m+1) \left((1-\alpha)^n + \sqrt{\frac{\log(1/\beta)}{2(m+2)}} \right) \right\rceil$$

hence the argmin defining $\tilde{\ell}^{\text{cond}}$ in Algorithm 4 is non-empty if and only if

$$(m+1) \left((1-\alpha)^n + \sqrt{\frac{\log(1/\beta)}{2(m+2)}} \right) \leq m,$$

which is equivalent to Eq. (23).

For the sufficient condition, remark that when $m \geq 2$ we have

$$\frac{m}{m+1} \geq \frac{2}{3}.$$

If in addition

$$m \geq \frac{9}{2} \log(1/\beta) - 2,$$

then

$$\sqrt{\frac{\log(1/\beta)}{2(m+2)}} \leq \frac{1}{3}, \quad \text{hence} \quad \frac{m}{m+1} - \sqrt{\frac{\log(1/\beta)}{2(m+2)}} \geq \frac{1}{3}.$$

Finally, if we also have $n \geq \frac{\log(1/3)}{\log(1-\alpha)}$, then $(1-\alpha)^n \leq 1/3$ and we get that Eq. (23) holds true, which proves Eq. (24). \square

Appendix E: Proofs of Section 4 (upper bounds)

E.1. Proof of Theorem 16

Let $\alpha \in (0, 1)$ be fixed. Let $(\ell_1, k_1) := (\ell^*, k^*)$ as defined by Algorithm 1, and $(\ell_2, k_2) := (\tilde{\ell}, \tilde{k}_{m,n}(\ell, \alpha))$ as defined by Algorithm 2. First, by definition of Algorithms 1–2 and by Proposition 8, we have

$$1 - \alpha \leq M_{\ell_1, k_1} := \min_{(\ell, k)} \{M_{\ell, k} : M_{\ell, k} \geq 1 - \alpha\} \leq M_{\ell_2, k_2} \leq F_{U_{(\ell_2, n)}}^{-1} \left(\frac{\tilde{k}_{m,n}(\ell_2, \alpha)}{m + 1/2} \right). \quad (50)$$

Second, by definition of Algorithm 2, for every $\ell \in \llbracket n \rrbracket$ such that $\tilde{k}_{m,n}(\ell, \alpha) \leq m$, we have

$$\begin{aligned} F_{U_{(\ell_2, n)}}^{-1} \left(\frac{\tilde{k}_{m,n}(\ell_2, \alpha)}{m + 1/2} \right) &\leq F_{U_{(\ell, n)}}^{-1} \left(\frac{\tilde{k}_{m,n}(\ell, \alpha)}{m + 1/2} \right) \\ &= F_{U_{(\ell, n)}}^{-1} \left(\frac{\lceil (m + 1/2) \cdot F_{U_{(\ell, n)}}(1 - \alpha) + 1/2 \rceil}{m + 1/2} \right) \\ &\leq F_{U_{(\ell, n)}}^{-1} \left(F_{U_{(\ell, n)}}(1 - \alpha) + \frac{3}{2m + 1} \right). \end{aligned} \quad (51)$$

By Theorem 25 with $x = F_{U_{(\ell, n)}}(1 - \alpha)$, hence $F_{U_{(\ell, n)}}^{-1}(x) = 1 - \alpha$, and $\varepsilon = \frac{3}{2m+1}$, we get that if $0 < F_{U_{(\ell, n)}}(1 - \alpha) < 1 - \frac{3}{2m+1}$, then

$$\begin{aligned} &F_{U_{(\ell, n)}}^{-1} \left(F_{U_{(\ell, n)}}(1 - \alpha) + \frac{3}{2m + 1} \right) \\ &\leq 1 - \alpha + \frac{3}{(2m + 1)\sqrt{n + 2}} \frac{\sqrt{2}}{g(\min\{F_{U_{(\ell, n)}}(1 - \alpha), 1 - F_{U_{(\ell, n)}}(1 - \alpha) - \frac{3}{2m+1}\})}. \end{aligned} \quad (52)$$

Third, note that for every $\alpha \in (0, 1)$, by Eq. (71) in Appendix G.2 with $p = 1 - \alpha$,

$$\begin{aligned} &\sqrt{n} (U_{(\lceil (1-\alpha)n \rceil)} - (1 - \alpha)) \xrightarrow[n \rightarrow +\infty]{\mathcal{L}} \mathcal{N}(0, \alpha(1 - \alpha)) \\ \text{hence} \quad &F_{U_{(\lceil (1-\alpha)n \rceil; n)}}(1 - \alpha) \xrightarrow[n \rightarrow +\infty]{} \frac{1}{2}. \end{aligned} \quad (53)$$

As a consequence, some $n_0(\alpha)$ exists such that, for every $n \geq n_0(\alpha)$, we have $1/3 \leq F_{U_{(\lceil (1-\alpha)n \rceil; n)}}(1 - \alpha) \leq 7/12$. Therefore, for every $n \geq n_0(\alpha)$ and $m \geq 18$,

$$\min \left\{ F_{U_{(\lceil (1-\alpha)n \rceil; n)}}(1 - \alpha), 1 - F_{U_{(\lceil (1-\alpha)n \rceil; n)}}(1 - \alpha) - \frac{3}{2m + 1} \right\} \geq \frac{1}{3}$$

and by Eq. (50)–(52), we get that

$$1 - \alpha \leq M_{\ell_1, k_1} \leq M_{\ell_2, k_2} \leq 1 - \alpha + \frac{27}{(2m + 1)\sqrt{n + 2}}.$$

Note that the choice of the constants $1/3$ and $7/12$ in the proof is arbitrary. Other choices would lead to the same result with other values for C , m_0 and $n_0(\alpha)$. For instance, C can be made as close to $3\sqrt{2}/g(1/2) \leq 16.1$ as desired, at the price of enlarging m_0 and $n_0(\alpha)$. \square

E.2. Proof of Theorem 17

The proof relies on the following lemma.

Lemma 30. *Let $n, m \geq 1$, $\ell \in \llbracket n \rrbracket$, $\alpha, \beta, \beta' \in (0, 1)$ and let $\tilde{k}_{m,n}^{\text{cond}}(\ell, \alpha, \beta)$ be defined by Eq. (22). Then, some $n'_0(\alpha), m'_0(\beta, \beta') \geq 1$ exist such that, if $n \geq n'_0(\alpha)$ and $m \geq m'_0(\beta, \beta')$, we have $\lceil n(1-\alpha) \rceil \in \llbracket n \rrbracket$, $\tilde{k}_{m,n}^{\text{cond}}(\lceil n(1-\alpha) \rceil, \alpha, \beta) \in \llbracket m \rrbracket$ and*

$$F_{U_{(\lceil n(1-\alpha) \rceil: n, \tilde{k}_{m,n}^{\text{cond}}(\lceil n(1-\alpha) \rceil, \alpha, \beta): m)}}^{-1}(1-\beta') \leq 1-\alpha + \frac{12 \max\{\sqrt{\log(1/\beta)} + \sqrt{\log(1/\beta')}, 1\}}{\sqrt{(m+2)(n+2)}}. \quad (54)$$

Proof. The first condition holds true when $n \geq 1/\alpha$, hence it suffices to take $n'_0(\alpha) \geq n'_{0,a}(\alpha) = 1/\alpha$. For the second condition, which can be rewritten

$$\tilde{k}_{m,n}^{\text{cond}}(\lceil n(1-\alpha) \rceil, \alpha, \beta) = \left\lceil (m+1) \left(F_{U_{(\lceil n(1-\alpha) \rceil: n)}}(1-\alpha) + \sqrt{\frac{\log(1/\beta)}{2(m+2)}} \right) \right\rceil \leq m,$$

we first notice that by Eq. (53), some $n'_{0,b}(\alpha)$ exists such that, for every $n \geq n'_{0,b}(\alpha)$,

$$\frac{1}{3} \leq F_{U_{(\lceil (1-\alpha)n \rceil: n)}}(1-\alpha) \leq \frac{2}{3}. \quad (55)$$

Therefore, some $m'_{0,a}(\beta)$ exists such that, if $m \geq m'_{0,a}(\beta)$ and $n \geq n'_{0,b}(\alpha)$,

$$\tilde{k}_{m,n}^{\text{cond}}(\lceil n(1-\alpha) \rceil, \alpha, \beta) \leq \left\lceil (m+1) \left(\frac{2}{3} + \sqrt{\frac{\log(1/\beta)}{2(m+2)}} \right) \right\rceil \leq m,$$

hence the second condition holds true if $m'_0(\beta, \beta') \geq m'_{0,a}(\beta)$.

It remains to prove Eq. (54). For every $\ell \in \llbracket n \rrbracket$, if $\tilde{k}_{m,n}^{\text{cond}}(\ell, \alpha, \beta) \in \llbracket m \rrbracket$, we have

$$\begin{aligned} & F_{U_{(\ell: n, \tilde{k}_{m,n}^{\text{cond}}(\ell, \alpha, \beta): m)}}^{-1}(1-\beta') \\ &= F_{U_{(\ell: n)}}^{-1} \circ F_{U_{(\tilde{k}_{m,n}^{\text{cond}}(\ell, \alpha, \beta): m)}}^{-1}(1-\beta') \quad \text{by Eq. (46) in Lemma 28} \\ &\leq F_{U_{(\ell: n)}}^{-1} \left(\frac{\tilde{k}_{m,n}^{\text{cond}}(\ell, \alpha, \beta)}{m+1} + \sqrt{\frac{\log(1/\beta')}{2(m+2)}} \right) \quad \text{by Eq. (41) in Proposition 24} \\ &\leq F_{U_{(\ell: n)}}^{-1} \left(F_{U_{(\ell: n)}}(1-\alpha) + \frac{\sqrt{\log(1/\beta)} + \sqrt{\log(1/\beta')}}{\sqrt{2(m+2)}} + \frac{1}{m+1} \right) \quad \text{by Eq. (22)} \\ &\leq 1-\alpha + \frac{\varepsilon_m}{\sqrt{n+2}} \frac{\sqrt{2}}{g(\min\{x_n, 1-x_n-\varepsilon_m\})} \end{aligned} \quad (56)$$

by Theorem 25 with

$$x = x_n := F_{U_{(\ell: n)}}(1-\alpha) \quad \text{and} \quad \varepsilon = \varepsilon_m := \frac{\sqrt{\log(1/\beta)} + \sqrt{\log(1/\beta')}}{\sqrt{2(m+2)}} + \frac{1}{m+1},$$

assuming also that $x_n + \varepsilon_m < 1$. Now, taking $\ell = \lceil (1-\alpha)n \rceil$ in Eq. (56), when $m \geq m'_{0,a}(\beta)$ and $n \geq \max\{n'_{0,a}(\alpha), n'_{0,b}(\alpha)\}$, we have $\ell \in \llbracket n \rrbracket$, $\tilde{k}_{m,n}^{\text{cond}}(\lceil (1-\alpha)n \rceil, \alpha, \beta) \in \llbracket m \rrbracket$ and $x_n \in [1/3, 2/3]$ by Eq. (55). Since some $m'_{0,b}(\beta, \beta')$ exists such that $\varepsilon_m \leq 1/12$ as soon as $m \geq m'_{0,b}(\beta, \beta')$, we get that for $m \geq \max\{m'_{0,a}(\beta), m'_{0,b}(\beta, \beta')\}$ and $n \geq$

$\max\{n'_{0,a}(\alpha), n'_{0,b}(\alpha)\}$, we have $x_n + \varepsilon_m < 1$ and $\min\{x_n, 1 - x_n - \varepsilon_m\} \geq 1/4$, hence Eq. (56) leads to

$$\begin{aligned} & F_{U_{(\lceil(1-\alpha)n\rceil:n, \tilde{k}_{m,n}^{\text{cond}}(\lceil(1-\alpha)n\rceil, \alpha, \beta):m)}}^{-1} (1 - \beta') \\ & \leq 1 - \alpha + \left(\frac{\sqrt{\log(1/\beta)} + \sqrt{\log(1/\beta')}}{\sqrt{2(m+2)}} + \frac{1}{m+1} \right) \frac{1}{\sqrt{n+2}} \frac{\sqrt{2}}{g(1/4)} \\ & \stackrel{(\text{if } m \geq 12)}{\leq} 1 - \alpha + \frac{12 \max\{\sqrt{\log(1/\beta)} + \sqrt{\log(1/\beta')}, 1\}}{\sqrt{(m+2)(n+2)}}, \end{aligned}$$

where we assume at the last line that $m \geq 12$, hence $\sqrt{m+2}/(m+1) \leq 0.9949267 - 1/\sqrt{2}$, and we use that $g(1/4) \geq G(1/4, 0.0316) \geq 0.1175$ hence $0.9949267 \times \frac{\sqrt{2}}{g(1/4)} \leq 11.9749 \leq 12$. This ends the proof, by taking $m'_0(\beta, \beta') = \max\{m'_{0,a}(\beta), m'_{0,b}(\beta, \beta'), 12\}$ and $n'_0(\alpha) = \max\{n'_{0,a}(\alpha), n'_{0,b}(\alpha)\}$. \square

We can now prove Theorem 17. We actually prove a more general result, about the following generalizations of Algorithms 3 and 4, respectively, where $\beta' \in (0, 1)$ is given:

$$(\ell_A, k_A) := \underset{(\ell, k) \in \llbracket n \rrbracket \times \llbracket m \rrbracket}{\operatorname{argmin}} \left\{ F_{U_{(\ell:n, k:m)}}^{-1} (1 - \beta') : F_{U_{(\ell:n, k:m)}}^{-1} (\beta) \geq 1 - \alpha \right\} \quad (57)$$

$$\ell_B := \underset{\ell \in \llbracket n \rrbracket \text{ s.t. } \tilde{k}_{m,n}^{\text{cond}}(\ell, \alpha, \beta) \leq m}{\operatorname{argmin}} \left\{ F_{U_{(\ell:n, \tilde{k}_{m,n}^{\text{cond}}(\ell, \alpha, \beta):m)}}^{-1} (1 - \beta') \right\} \quad \text{and} \quad k_B := \tilde{k}_{m,n}^{\text{cond}}(\ell_B, \alpha, \beta). \quad (58)$$

Algorithm 3 – Proof of Eq. (26) Let us first consider the generalization of Algorithm 3 defined by Eq. (57). Let $\ell_n := \lceil n(1 - \alpha) \rceil$ and $k_n := \tilde{k}_{m,n}^{\text{cond}}(\ell_n, \alpha, \beta)$. By Lemma 30, when $n \geq n'_0(\alpha)$ and $m \geq m'_0(\beta, \beta')$, we have $\ell_n \in \llbracket n \rrbracket$ and $k_n \in \llbracket m \rrbracket$, hence, $F_{U_{(\ell_n:n, k_n:m)}}^{-1}(\beta) \geq 1 - \alpha$ by Proposition 13 and Eq. (22). Therefore, by definition (57) of (ℓ_A, k_A) ,

$$\begin{aligned} F_{U_{(\ell_A:n, k_A:m)}}^{-1} (1 - \beta') & \leq F_{U_{(\ell_n:n, k_n:m)}}^{-1} (1 - \beta') \\ & \leq 1 - \alpha + \frac{12 \max\{\sqrt{\log(1/\beta)} + \sqrt{\log(1/\beta')}, 1\}}{\sqrt{(m+2)(n+2)}} \end{aligned}$$

by Eq. (54) in Lemma 30. Finally, using Eq. (18) in Theorem 10 —with $\beta = 0$ and $(\ell, k) = (\ell_A, k_A)$ —, we get that if the scores $S_{i,j}, S$ are a.s. distinct, when $n \geq n'_0(\alpha)$ and $m \geq m'_0(\beta, \beta')$, the following inequality holds true with probability $1 - \beta'$:

$$1 - \alpha_{\ell_A, k_A}(\mathcal{D}) \leq 1 - \alpha + \frac{12 \max\{\sqrt{\log(1/\beta)} + \sqrt{\log(1/\beta')}, 1\}}{\sqrt{(m+2)(n+2)}}. \quad (59)$$

Eq. (26) follows by taking $\beta' = \beta$ and defining $m'_0(\beta) := m'_0(\beta, \beta)$.

Algorithm 4 – Proof of Eq. (27) Similarly, for the generalization of Algorithm 4 defined by Eq. (58), we get that when $n \geq n'_0(\alpha)$ and $m \geq m'_0(\beta, \beta')$, the following inequality holds true with probability $1 - \beta'$:

$$1 - \alpha_{\ell_B, k_B}(\mathcal{D}) \leq 1 - \alpha + \frac{12 \max\{\sqrt{\log(1/\beta)} + \sqrt{\log(1/\beta')}, 1\}}{\sqrt{(m+2)(n+2)}}. \quad (60)$$

Eq. (27) follows by taking $\beta' = \beta$, since Eq. (55) still holds true for such a sequence ℓ_n by [13, Section 10.2]. \square

Appendix F: Proofs of Section 5 (different n_j)

Before proving Theorems 19–20, notice that the proof of Lemma 27 straightforwardly generalizes to the case of the QQ estimator $S_{(\ell_1, \dots, \ell_m, k)}$ defined by Eq. (29). So, if $(U_{i,j})_{i \in [\ell_j], j \in [m]}$ denote independent standard uniform variables, we have

$$F_S(S_{(\ell_1, \dots, \ell_m, k)}) \stackrel{d}{=} F_S \circ F_S^{-1}(U_{(\ell_1, \dots, \ell_m, k)}) \geq U_{(\ell_1, \dots, \ell_m, k)}, \quad (61)$$

with equality when F_S is continuous.

F.1. Proof of Theorem 19

First, by Eq. (61), we have

$$\begin{aligned} \mathbb{P}\left(Y \in \widehat{\mathcal{C}}_{\ell_1, \dots, \ell_m, k}(X)\right) &= \mathbb{P}\left(S \leq S_{(\ell_1, \dots, \ell_m, k)}\right) \\ &= \mathbb{E}\left[F_S(S_{(\ell_1, \dots, \ell_m, k)})\right] \geq \mathbb{E}\left[U_{(\ell_1, \dots, \ell_m, k)}\right] =: M_{\ell_1, \dots, \ell_m, k}. \end{aligned}$$

It remains to prove the formula for $M_{\ell_1, \dots, \ell_m, k}$. The difference with the proof of Theorem 5 is that the variables $(\widehat{Q}_{(\ell_j: n_j)}(\mathcal{S}_j))_{1 \leq j \leq m}$ are not identically distributed, so $U_{(\ell_1, \dots, \ell_m, k)}$ is the k -th order statistics of a set of independent but *not* identically distributed (inid) random variables. Its cdf can still be computed as follows. By [7, Equation (16)], we have, for every $t \in [0, 1]$,

$$\begin{aligned} &F_{U_{(\ell_1, \dots, \ell_m, k)}}(t) \\ &= \sum_{j=k}^m \sum_{a \in \mathcal{P}_j} F_{U_{(\ell_{a_1}: n_1)}}(t) \cdots F_{U_{(\ell_{a_j}: n_j)}}(t) \cdot \left[1 - F_{U_{(\ell_{a_{j+1}}: n_{j+1})}}(t)\right] \cdots \left[1 - F_{U_{(\ell_{a_m}: n_m)}}(t)\right], \end{aligned}$$

where \mathcal{P}_j denotes the set of permutations (a_1, \dots, a_m) of $\{1, \dots, m\}$ for which $a_1 < a_2 < \dots < a_j$ and $a_{j+1} < a_{j+2} < \dots < a_m$. Then, using Eq. (65) and (67) in Appendix G.1 and rearranging the terms, we get that for every $t \in [0, 1]$,

$$\begin{aligned} &F_{U_{(\ell_1, \dots, \ell_m, k)}}(t) \\ &= \sum_{j=k}^m \sum_{\mathcal{P}_j} \sum_{i_1=\ell_{a_1}}^{n_{a_1}} \cdots \sum_{i_j=\ell_{a_j}}^{n_{a_j}} \sum_{i_{j+1}=0}^{\ell_{a_{j+1}}-1} \cdots \sum_{i_m=0}^{\ell_{a_m}-1} \binom{n_{a_1}}{i_1} \cdots \binom{n_{a_m}}{i_m} t^{i_1+\dots+i_m} (1-t)^{n_1+\dots+n_m-(i_1+\dots+i_m)}. \end{aligned}$$

We finally obtain the result using that

$$\begin{aligned} M_{\ell_1, \dots, \ell_m, k} &= \mathbb{E}[U_{(\ell_1, \dots, \ell_m, k)}] = \int_0^1 [1 - F_{U_{(\ell_1, \dots, \ell_m, k)}}(t)] dt \\ \text{and} \quad &\int_0^1 t^{i_1+\dots+i_m} (1-t)^{n_1+\dots+n_m-(i_1+\dots+i_m)} dt \\ &= B(i_1 + \dots + i_m + 1, n_1 + \dots + n_m - (i_1 + \dots + i_m) + 1) \\ &= (n_1 + \dots + n_m + 1)^{-1} \binom{n_1 + \dots + n_m}{i_1 + \dots + i_m}^{-1}, \end{aligned}$$

using a property of the Beta function already used in Appendix C.1. \square

F.2. Proof of Theorem 20

By the definition (30) of $\widehat{\mathcal{C}}_{\ell_1, \dots, \ell_m, k}$ and Eq. (61), we have

$$\begin{aligned} \mathbb{P}(1 - \alpha_{\ell_1, \dots, \ell_m, k}(\mathcal{D}) \geq 1 - \alpha) &= \mathbb{P}(F_S(S_{(\ell_1, \dots, \ell_m, k)}) \geq 1 - \alpha) \\ &\geq \mathbb{P}(U_{(\ell_1, \dots, \ell_m, k)} \geq 1 - \alpha), \end{aligned} \quad (62)$$

with equality when F_S is continuous. By definition of $U_{(\ell_1, \dots, \ell_m, k)} = \widehat{Q}_{(\ell_j: n_j)}((U_{(\ell_j: n_j)})_{1 \leq j \leq m})$, we have

$$\mathbb{P}(U_{(\ell_1, \dots, \ell_m, k)} \geq 1 - \alpha) = \mathbb{P}\left(\sum_{j=1}^m \mathbb{1}_{\{U_{(\ell_j: n_j)} \leq 1 - \alpha\}} \leq k - 1\right).$$

Furthermore, $\sum_{j=1}^m \mathbb{1}_{\{U_{(\ell_j: n_j)} \leq 1 - \alpha\}}$ is a sum of m independent Bernoulli variables with respective parameters $p_j = \mathbb{P}(U_{(\ell_j: n_j)} \leq 1 - \alpha) = F_{U_{(\ell_j: n_j)}}(1 - \alpha)$. Therefore, it follows a Poisson-Binomial distribution of parameters (p_1, \dots, p_m) . Using Eq. (62), we obtain Eq. (33). \square

Appendix G: Classical results about order statistics distribution

We recall in this section some useful well-known results about order statistics [13]. Throughout this section, U_1, \dots, U_N, \dots denote a sequence of independent standard uniform variables, and Z_1, \dots, Z_N, \dots a sequence of i.i.d. random variables with common cdf F_Z . For any integers $1 \leq r \leq N$, $U_{(r:N)} = \widehat{Q}_{(r)}(U_1, \dots, U_N)$ and $Z_{(r:N)} = \widehat{Q}_{(r)}(Z_1, \dots, Z_N)$ denote the corresponding r -th order statistics.

G.1. Exact distribution

For every $r \in \llbracket N \rrbracket$, the cdf of $Z_{(r:N)}$ is given by

$$\forall t \in \mathbb{R}, \quad F_{Z_{(r:N)}}(t) = \sum_{i=r}^N \binom{N}{i} F_Z(t)^i [1 - F_Z(t)]^{N-i}. \quad (63)$$

If we further assume that F_Z is continuous with corresponding probability density function (pdf) $f_Z = F'_Z$, then, for every $r \in \llbracket N \rrbracket$, the pdf of $Z_{(r:N)}$ is given by

$$\forall t \in \mathbb{R}, \quad f_{Z_{(r:N)}}(t) = \frac{N!}{(r-1)!(N-r)!} F_Z(t)^{r-1} [1 - F_Z(t)]^{N-r} f_Z(t). \quad (64)$$

In the case of uniform order statistics, $U_{(r:N)}$ follows a Beta($r, N - r + 1$) distribution and its cdf $F_{U_{(r:N)}}$ and pdf $f_{U_{(r:N)}}$ are respectively given by

$$\forall t \in [0, 1], \quad F_{U_{(r:N)}}(t) = \sum_{i=r}^N \binom{N}{i} t^i (1-t)^{N-i} \quad \text{and} \quad f_{U_{(r:N)}}(t) = \frac{t^{r-1} (1-t)^{N-r}}{B(r, N-r+1)} \quad (65)$$

$$\text{where } B : (a, b) \in (0, +\infty)^2 \mapsto \int_0^1 t^{a-1} (1-t)^{b-1} dt$$

denotes the Beta function [59]. In particular, for every $N \geq 1$, we have

$$\forall t \in [0, 1], \quad F_{U_{(N:N)}}(t) = t^N \quad \text{and} \quad F_{U_{(N:N)}}^{-1}(t) = t^{1/N}. \quad (66)$$

Since $\sum_{i=0}^N \binom{N}{i} t^i (1-t)^{N-i} = 1$, Eq. (65) also implies that

$$\forall t \in [0, 1], \quad 1 - F_{U_{(r:N)}}(t) = \sum_{i=0}^{r-1} \binom{N}{i} t^i (1-t)^{N-i}, \quad (67)$$

$$\text{hence} \quad F_{U_{(1:N)}}(t) = 1 - (1-t)^N. \quad (68)$$

In addition, closed-form formulas are available for the expectation and variance of uniform order statistics:

$$\mathbb{E}[U_{(r:N)}] = \frac{r}{N+1} \quad \text{and} \quad \text{Var}(U_{(r:N)}) = \frac{r(N-r+1)}{(N+1)^2(N+2)}. \quad (69)$$

G.2. Asymptotics

Central order statistics are known to be asymptotically normal. More precisely, following [18, 33], we have the following result.

Theorem 31 (Asymptotic normality of central uniform order statistics [18, 33]). *Let $p \in (0, 1)$, $\gamma \in \mathbb{R}$ and $(\ell_N)_{N \geq 1}$ be a sequence such that $\ell_N \in \llbracket N \rrbracket$ for every $N \geq 1$ and $\ell_N = pN + \gamma\sqrt{N} + o(\sqrt{N})$ when $N \rightarrow +\infty$. Then, we have*

$$\sqrt{N}(U_{(\ell_N:N)} - p) \xrightarrow[N \rightarrow +\infty]{\mathcal{L}} \mathcal{N}(\gamma, p(1-p)). \quad (70)$$

In particular, Theorem 31 implies that

$$\sqrt{N}(U_{(\lceil pN \rceil:N)} - p) \xrightarrow[N \rightarrow +\infty]{\mathcal{L}} \mathcal{N}(0, p(1-p)). \quad (71)$$

By Lemma 23, Eq. (70)–(71) imply that when F_Z is continuous, under the assumptions of Theorem 31,

$$\sqrt{N}(F_Z(Z_{(\ell_N:N)}) - p) \xrightarrow[N \rightarrow +\infty]{\mathcal{L}} \mathcal{N}(\gamma, p(1-p)) \quad (72)$$

$$\text{and} \quad \sqrt{N}(F_Z(Z_{(\lceil pN \rceil:N)}) - p) \xrightarrow[N \rightarrow +\infty]{\mathcal{L}} \mathcal{N}(0, p(1-p)). \quad (73)$$

Appendix H: Proofs of Appendix B

H.1. Proof of Lemma 23

First, for every $i \in \llbracket N \rrbracket$, let us define $\tilde{Z}_i := F_Z^{-1}(U_i)$, so that $\tilde{Z}_1, \dots, \tilde{Z}_N$ are independent and identically distributed, with cdf F_Z . Therefore, $(Z_1, \dots, Z_N) \stackrel{d}{=} (\tilde{Z}_1, \dots, \tilde{Z}_N)$, hence

$$Z_{(r)} \stackrel{d}{=} \tilde{Z}_{(r)}. \quad (74)$$

Second, remark that for any integers $1 \leq r \leq N$, any $\mathcal{S} \in \mathbb{R}^N$ and any nondecreasing function $h : \mathbb{R} \rightarrow \mathbb{R}$, we have

$$h(\hat{Q}_{(r)}(\mathcal{S})) = \hat{Q}_{(r)}(h(\mathcal{S})) \quad \text{where} \quad h(\mathcal{S}) := (h(\mathcal{S}_i))_{i \in \llbracket N \rrbracket} \quad (75)$$

and $\widehat{Q}_{(r)}$ is defined by Eq. (4) in Section 2.1.

Indeed, let $\sigma \in \llbracket N \rrbracket$ be some permutation of $\llbracket N \rrbracket$ such that $\mathcal{S}_{\sigma(1)} \leq \dots \leq \mathcal{S}_{\sigma(N)}$, hence $\widehat{Q}_{(r)}(\mathcal{S}) = \mathcal{S}_{\sigma(r)}$. Then, $h(\mathcal{S}_{\sigma(1)}) \leq \dots \leq h(\mathcal{S}_{\sigma(N)})$ since h is nondecreasing, and we get that

$$\widehat{Q}_{(r)}(h(\mathcal{S})) = h(\mathcal{S}_{\sigma(r)}) = h(\widehat{Q}_{(r)}(\mathcal{S})).$$

From Eq. (74) and (75), since F_Z^{-1} is nondecreasing, we obtain that

$$\begin{aligned} F_Z(Z_{(r)}) &\stackrel{d}{=} F_Z(\widetilde{Z}_{(r)}) = F_Z\left(\widehat{Q}_{(r)}(F_Z^{-1}(U_1, \dots, U_N))\right) = F_Z \circ F_Z^{-1}\left(\widehat{Q}_{(r)}(U_1, \dots, U_N)\right) \\ &= F_Z \circ F_Z^{-1}(U_{(r)}). \end{aligned}$$

Using Eq. (36), this proves Eq. (37) in the general case —Eq. (38) follows directly—, and Eq. (39) when F_Z is continuous. The distribution of the uniform order statistic $U_{(r)}$ is well known, see for instance [13] or Appendix G.1. \square

Remark 32. In the case of a general cdf F_Z , another consequence of Lemma 23 is that, for every $a, b \in \mathbb{R}$ such that $a \leq b$,

$$\mathbb{P}(U_{(r)} \leq b) \leq \mathbb{P}(F_Z(Z_{(r)}) \leq F_Z \circ F_Z^{-1}(b)) \quad (76)$$

$$\text{and} \quad \mathbb{P}(a \leq U_{(r)} \leq b) \leq \mathbb{P}(a \leq F_Z(Z_{(r)}) \leq F_Z \circ F_Z^{-1}(b)). \quad (77)$$

Proof. Since $F_Z \circ F_Z^{-1}$ is nondecreasing,

$$U_{(r)} \leq b \quad \text{implies that} \quad F_Z \circ F_Z^{-1}(U_{(r)}) \leq F_Z \circ F_Z^{-1}(b), \quad (78)$$

hence Eq. (76) thanks to Eq. (37). Since $F_Z \circ F_Z^{-1}(U_{(r)}) \geq U_{(r)}$ by Eq. (36), Eq. (78) also shows that

$$a \leq U_{(r)} \leq b \quad \text{implies that} \quad a \leq F_Z \circ F_Z^{-1}(U_{(r)}) \leq F_Z \circ F_Z^{-1}(b),$$

hence Eq. (77) thanks to Eq. (37). \square

H.2. Proof of Proposition 24

We start by recalling some concentration result about the Beta distribution [42, Theorem 1].

Theorem 33 (taken from [42]). *For any $a, b > 0$, the Beta distribution $\text{Beta}(a, b)$ is ω -sub-Gaussian, with $\omega = \frac{1}{4(a+b+1)}$. This implies that, if $Z \sim \text{Beta}(a, b)$, for any $\varepsilon > 0$,*

$$\max \left\{ \mathbb{P} \left(Z \geq \frac{a}{a+b} + \varepsilon \right), \mathbb{P} \left(Z \leq \frac{a}{a+b} - \varepsilon \right) \right\} \leq e^{-2(a+b+1)\varepsilon^2}.$$

Note that Theorem 33 does not provide the “best” concentration inequality for the Beta distribution. Indeed, although the bound is sharp for the symmetric case $a = b$, when the Beta distribution is skewed, Theorem 33 has been improved with the Bernstein-type bound of [57]. Nevertheless, it is sufficient for our needs and yields simpler formulas.

We can now prove Proposition 24. Since $U_{(r:N)}$ follows a $\text{Beta}(r, N-r+1)$ distribution, Theorem 33 with $a = r$, $b = N-r+1$ and $\varepsilon = \sqrt{\frac{\log(1/\delta)}{2(N+2)}}$ shows that

$$F_{U_{(r:N)}} \left(\frac{r}{N+1} - \sqrt{\frac{\log(1/\delta)}{2(N+2)}} \right) = \mathbb{P} \left(U_{(r:N)} \leq \frac{r}{N+1} - \sqrt{\frac{\log(1/\delta)}{2(N+2)}} \right) \leq \delta, \quad (79)$$

hence Eq. (40). Similarly,

$$F_{U_{(r:N)}} \left(\frac{r}{N+1} + \sqrt{\frac{\log(1/\delta)}{2(N+2)}} \right) = 1 - \mathbb{P} \left(U_{(r:N)} \geq \frac{r}{N+1} + \sqrt{\frac{\log(1/\delta)}{2(N+2)}} \right) \geq 1 - \delta \quad (80)$$

by Theorem 33, hence Eq. (41). \square

H.3. Deviation inequalities for general order statistics

The combination of Theorem 33 with Lemma 23 yields the following deviation inequalities, which are not directly useful in our study of one-shot FL CP algorithms, but can be of interest in other contexts such as the non-asymptotic analysis of (centralized) split CP —see Appendix I.1.

Corollary 34. *For any $N \geq 1$, let Z_1, \dots, Z_N be independent and identically distributed real-valued random variables with common cdf F_Z . Then, for any $\delta > 0$ and $r \in \llbracket N \rrbracket$, the following inequality holds true:*

$$\mathbb{P} \left(F_Z(Z_{(r)}) \geq \frac{r}{N+1} - \sqrt{\frac{\log(1/\delta)}{2(N+2)}} \right) \geq 1 - \delta. \quad (81)$$

Moreover, if F_Z is continuous, we also have

$$\mathbb{P} \left(F_Z(Z_{(r)}) \leq \frac{r}{N+1} + \sqrt{\frac{\log(1/\delta)}{2(N+2)}} \right) \geq 1 - \delta \quad (82)$$

$$\text{and} \quad \mathbb{P} \left(\left| F_Z(Z_{(r)}) - \frac{r}{N+1} \right| \leq \sqrt{\frac{\log(2/\delta)}{2(N+2)}} \right) \geq 1 - \delta. \quad (83)$$

Proof. By Eq. (38) in Lemma 23 with $a = \frac{r}{N+1} - \sqrt{\frac{\log(1/\delta)}{2(N+2)}}$ (and using its notation), we have

$$\begin{aligned} & \mathbb{P} \left(F_Z(Z_{(r)}) \geq \frac{r}{N+1} - \sqrt{\frac{\log(1/\delta)}{2(N+2)}} \right) \\ & \geq \mathbb{P} \left(U_{(r:N)} \geq \frac{r}{N+1} - \sqrt{\frac{\log(1/\delta)}{2(N+2)}} \right) \geq 1 - \delta \end{aligned} \quad (84)$$

by Eq. (79) in the proof of Proposition 24, hence Eq. (81). If we also assume that F_Z is continuous, combining Eq. (39) in Lemma 23 and Eq. (80) in the proof of Proposition 24 shows that

$$\begin{aligned} \mathbb{P} \left(F_Z(Z_{(r:N)}) \leq \frac{r}{N+1} + \sqrt{\frac{\log(1/\delta)}{2(N+2)}} \right) &= \mathbb{P} \left(U_{(r:N)} \leq \frac{r}{N+1} + \sqrt{\frac{\log(1/\delta)}{2(N+2)}} \right) \\ &= F_{U_{(r:N)}} \left(\frac{r}{N+1} + \sqrt{\frac{\log(1/\delta)}{2(N+2)}} \right) \geq 1 - \delta, \end{aligned}$$

hence Eq. (82). Finally, replacing δ by $\delta/2$ in Eq. (81)–(82) implies Eq. (83) by the union bound. \square

Remark 35. For a general cdf F_Z , instead of Eq. (82), we have

$$\mathbb{P}\left(F_Z(Z_{(r)}) \leq F_Z \circ F_Z^{-1}\left(\frac{r}{N+1} + \sqrt{\frac{\log(1/\delta)}{2(N+2)}}\right)\right) \geq 1 - \delta. \quad (85)$$

Proof. Combine Eq. (76) with $b = \frac{r}{N+1} + \sqrt{\frac{\log(1/\delta)}{2(N+2)}}$ and Eq. (80). \square

Remark 36. Eq. (73) in Section G.2 shows that for any $p \in (0, 1)$, $F_Z(Z_{(\lceil pN \rceil)})$ converges towards p at the rate \sqrt{N} . Therefore, the order of magnitude $1/\sqrt{N}$ of the deviations in Corollary 34 is optimal (at least, asymptotically).

H.4. Proof of Theorem 25

We proceed in four steps.

Step 1: mean value theorem. Since $F_{U(\ell;n)}^{-1}$ is continuous on $[x, x+\varepsilon]$ and differentiable on $(x, x+\varepsilon)$ —it is even \mathcal{C}^∞ on $(0, 1)$ —, by the mean-value theorem, some $c \in (x, x+\varepsilon)$ exists such that

$$\begin{aligned} F_{U(\ell;n)}^{-1}(x+\varepsilon) - F_{U(\ell;n)}^{-1}(x) &= (F_{U(\ell;n)}^{-1})'(c) \times \varepsilon = \frac{\varepsilon}{f_{U(\ell;n)}(F_{U(\ell;n)}^{-1}(c))} \\ &\leq \frac{\varepsilon}{\inf_{[F_{U(\ell;n)}^{-1}(x), F_{U(\ell;n)}^{-1}(x+\varepsilon)]} f_{U(\ell;n)}}. \end{aligned} \quad (86)$$

Therefore, it remains to get a uniform lower bound on the density $f_{U(\ell;n)}$.

Step 2: density lower bound with the inverse cdf $F_{U(\ell;n)}^{-1}$. Since $f_{U(\ell;n)}$ is unimodal [6], a straightforward proof by exhaustion (depending on the relative position of $\operatorname{argmax} f_{U(\ell;n)}$ and $[F_{U(\ell;n)}^{-1}(x), F_{U(\ell;n)}^{-1}(x+\varepsilon)]$) shows that

$$\inf_{[F_{U(\ell;n)}^{-1}(x), F_{U(\ell;n)}^{-1}(x+\varepsilon)]} f_{U(\ell;n)} \geq \min\left\{\sup_{[0, F_{U(\ell;n)}^{-1}(x)]} f_{U(\ell;n)}, \sup_{[F_{U(\ell;n)}^{-1}(x+\varepsilon), 1]} f_{U(\ell;n)}\right\}. \quad (87)$$

Let us consider separately the two terms of the right-hand side of Eq. (87). On the one hand, for every $\delta \in (0, x)$,

$$\begin{aligned} \sup_{[0, F_{U(\ell;n)}^{-1}(x)]} f_{U(\ell;n)} &\geq \frac{1}{F_{U(\ell;n)}^{-1}(x) - F_{U(\ell;n)}^{-1}(\delta)} \int_{F_{U(\ell;n)}^{-1}(\delta)}^{F_{U(\ell;n)}^{-1}(x)} f_{U(\ell;n)}(t) dt \\ &= \frac{F_{U(\ell;n)} \circ F_{U(\ell;n)}^{-1}(x) - F_{U(\ell;n)} \circ F_{U(\ell;n)}^{-1}(\delta)}{F_{U(\ell;n)}^{-1}(x) - F_{U(\ell;n)}^{-1}(\delta)} \\ &= \frac{x - \delta}{F_{U(\ell;n)}^{-1}(x) - F_{U(\ell;n)}^{-1}(\delta)}. \end{aligned} \quad (88)$$

On the other hand, similarly, for every $\delta' \in (0, 1 - x - \varepsilon)$,

$$\begin{aligned} \sup_{[F_{U(\ell;n)}^{-1}(x+\varepsilon), 1]} f_{U(\ell;n)} &\geq \frac{1}{F_{U(\ell;n)}^{-1}(1 - \delta') - F_{U(\ell;n)}^{-1}(x+\varepsilon)} \int_{F_{U(\ell;n)}^{-1}(x+\varepsilon)}^{F_{U(\ell;n)}^{-1}(1 - \delta')} f_{U(\ell;n)}(t) dt \\ &= \frac{(1 - \delta') - (x + \varepsilon)}{F_{U(\ell;n)}^{-1}(1 - \delta') - F_{U(\ell;n)}^{-1}(x + \varepsilon)}. \end{aligned} \quad (89)$$

Step 3: use of concentration results. Given Eq. (88) and (89), it remains to show upper bounds on $F_{U(\ell;n)}^{-1}(x) - F_{U(\ell;n)}^{-1}(\delta)$ and $F_{U(\ell;n)}^{-1}(1 - \delta') - F_{U(\ell;n)}^{-1}(x + \varepsilon)$. On the one hand, by Eq. (40)–(41) in Proposition 24 with $r = \ell$ and $N = n$,

$$\forall \delta \in (0, 1 - x), \quad F_{U(\ell;n)}^{-1}(x) - F_{U(\ell;n)}^{-1}(\delta) \leq F_{U(\ell;n)}^{-1}(1 - \delta) - F_{U(\ell;n)}^{-1}(\delta) \leq \sqrt{\frac{2 \log(1/\delta)}{n + 2}}.$$

Therefore, using Eq. (88), we obtain that

$$\sup_{[0, F_{U(\ell;n)}^{-1}(x)]} f_{U(\ell;n)} \geq \sup_{\delta \in (0, \min\{x, 1-x\})} \left\{ \frac{x - \delta}{F_{U(\ell;n)}^{-1}(x) - F_{U(\ell;n)}^{-1}(\delta)} \right\} \geq g(x) \sqrt{\frac{n+2}{2}} \quad (90)$$

$$\text{with } g(x) := \sup_{\delta \in (0, \min\{x, 1-x\})} \left\{ \frac{x - \delta}{\sqrt{\log(1/\delta)}} \right\}.$$

On the other hand, similarly, for every $\delta' \in (0, x + \varepsilon)$,

$$F_{U(\ell;n)}^{-1}(1 - \delta') - F_{U(\ell;n)}^{-1}(x + \varepsilon) \leq F_{U(\ell;n)}^{-1}(1 - \delta') - F_{U(\ell;n)}^{-1}(\delta') \leq \sqrt{\frac{2 \log(1/\delta')}{n + 2}}.$$

Therefore, using Eq. (89), we obtain that

$$\begin{aligned} \sup_{[F_{U(\ell;n)}^{-1}(x+\varepsilon), 1]} f_{U(\ell;n)} &\geq \sup_{\delta' \in (0, \min\{x+\varepsilon, 1-(x+\varepsilon)\})} \left\{ \frac{(1 - \delta') - (x + \varepsilon)}{F_{U(\ell;n)}^{-1}(1 - \delta') - F_{U(\ell;n)}^{-1}(x + \varepsilon)} \right\} \\ &\geq \sup_{\delta' \in (0, \min\{x+\varepsilon, 1-(x+\varepsilon)\})} \left\{ \frac{1 - (x + \varepsilon) - \delta'}{\sqrt{\log(1/\delta')}} \right\} \times \sqrt{\frac{n+2}{2}} \quad (91) \end{aligned}$$

$$= g(1 - x - \varepsilon) \sqrt{\frac{n+2}{2}}. \quad (92)$$

Step 4: conclusion. Combining Eq. (86), (87), (90) and (92), we obtain that

$$F_{U(\ell;n)}^{-1}(x + \varepsilon) - F_{U(\ell;n)}^{-1}(x) \leq \frac{\sqrt{2}}{\sqrt{n+2}} \frac{\varepsilon}{\min\{g(x), g(1 - x - \varepsilon)\}}. \quad (93)$$

Now, let us rewrite the definition of g as

$$\forall t \in (0, 1), \quad g(t) := \sup_{\delta \in (0, \min\{t, 1-t\})} G(t, \delta) \quad \text{where } \forall \delta > 0, \quad G(t, \delta) := \frac{t - \delta}{\sqrt{\log(1/\delta)}}.$$

Note that $t \mapsto G(t, \delta)$ is increasing whatever $\delta > 0$, so that g is nondecreasing on $(0, 1/2]$. In addition, for every $t \in (0, 1/2]$, using that $t \leq 1 - t$, we have $G(t, \delta) \leq G(1 - t, \delta)$ for every $\delta > 0$, hence $g(t) \leq g(1 - t)$. This proves that

$$\forall t \in (0, 1), \quad g(\min\{t, 1 - t\}) = \min\{g(t), g(1 - t)\},$$

hence

$$\begin{aligned} \min\{g(x), g(1 - x - \varepsilon)\} &\geq \min\{g(x), g(1 - x), g(x + \varepsilon), g(1 - x - \varepsilon)\} \\ &= \min\left\{g(\min\{x, 1 - x\}), g(\min\{x + \varepsilon, 1 - x - \varepsilon\})\right\} \\ &= g(\min\{x, 1 - x, x + \varepsilon, 1 - x - \varepsilon\}) \quad \text{since } g \text{ is nondecreasing on } (0, 1/2] \end{aligned}$$

$$= g(\min\{x, 1 - x - \varepsilon\}),$$

By Eq. (93), we obtain that Eq. (42) holds true.

Let us finally prove the inequalities stated about the function g . First, for every $t \in (0, 1/2]$, we have $t/2 \leq \min\{t, 1 - t\}$, hence

$$g(t) \geq G(t, t/2) = \frac{t}{2\sqrt{\log(2/t)}}.$$

Second, since g is nondecreasing over $(0, 1/2]$, for every $t \in [1/3, 1/2]$,

$$g(t) \geq g\left(\frac{1}{3}\right) \geq G\left(\frac{1}{3}, 0.047\right) \geq 0.163 \geq \frac{\sqrt{2}}{9}.$$

□

H.5. Proof of Lemma 27

We proceed similarly to the proof of Lemma 23. First, let us define $\tilde{Z}_{i,j} = F_Z^{-1}(U_{i,j})$ for every $i \in \llbracket n \rrbracket$ and $j \in \llbracket m \rrbracket$, so that $(\tilde{Z}_{i,j})_{1 \leq i \leq n, 1 \leq j \leq m}$ has the same distribution as $(Z_{i,j})_{1 \leq i \leq n, 1 \leq j \leq m}$. Therefore, the corresponding order statistics of order statistics have the same distributions:

$$Z_{(\ell,k)} \stackrel{d}{=} \tilde{Z}_{(\ell,k)}. \quad (94)$$

Second, using twice Eq. (75) in the proof of Lemma 23, since F_Z^{-1} is nondecreasing, we have

$$\begin{aligned} \tilde{Z}_{(\ell,k)} &= \widehat{Q}_{(k)} \left(\left(\widehat{Q}_{(\ell)}(F_Z^{-1}(U_{1,j}), \dots, F_Z^{-1}(U_{n,j})) \right)_{j \in \llbracket m \rrbracket} \right) \\ &= \widehat{Q}_{(k)} \left(F_Z^{-1} \left(\widehat{Q}_{(\ell)}(U_{1,j}, \dots, U_{n,j}) \right)_{j \in \llbracket m \rrbracket} \right) \\ &= F_Z^{-1} \left(\widehat{Q}_{(k)} \left(\widehat{Q}_{(\ell)}(U_{1,j}, \dots, U_{n,j}) \right)_{j \in \llbracket m \rrbracket} \right) = F_Z^{-1}(U_{(\ell,k)}), \end{aligned}$$

hence Eq. (94) implies that

$$F_Z(Z_{(\ell,k)}) \stackrel{d}{=} F_Z(\tilde{Z}_{(\ell,k)}) = F_Z \circ F_Z^{-1}(U_{(\ell,k)}).$$

Using Eq. (36), this proves Eq. (43) in the general case, and Eq. (45) when F_Z is continuous. □

Remark 37. In the case of a general cdf F_Z , another consequence of Lemma 27 is that, for every $a, b \in \mathbb{R}$ such that $a \leq b$,

$$\mathbb{P}(U_{(\ell,k)} \leq b) \leq \mathbb{P}(F_Z(Z_{(\ell,k)}) \leq F_Z \circ F_Z^{-1}(b)) \quad (95)$$

$$\text{and } \mathbb{P}(a \leq U_{(\ell,k)} \leq b) \leq \mathbb{P}(a \leq F_Z(Z_{(\ell,k)}) \leq F_Z \circ F_Z^{-1}(b)). \quad (96)$$

Proof. Since $F_Z \circ F_Z^{-1}$ is nondecreasing,

$$U_{(\ell,k)} \leq b \quad \text{implies that} \quad F_Z \circ F_Z^{-1}(U_{(\ell,k)}) \leq F_Z \circ F_Z^{-1}(b), \quad (97)$$

hence Eq. (95) thanks to Eq. (43). Since $F_Z \circ F_Z^{-1}(U_{(\ell,k)}) \geq U_{(\ell,k)}$ by Eq. (36), Eq. (97) also shows that

$$a \leq U_{(\ell,k)} \leq b \quad \text{implies that} \quad a \leq F_Z \circ F_Z^{-1}(U_{(\ell,k)}) \leq F_Z \circ F_Z^{-1}(b),$$

hence Eq. (96) thanks to Eq. (43). □

H.6. Proof of Lemma 28

By Eq. (39) in Lemma 23 with $F_Z = F_{U_{(\ell:n)}}$ (which is continuous), $r = k$ and $N = m$, we have

$$F_{U_{(\ell:n)}}(Z_{(k:m)}) \stackrel{d}{=} U_{(k:m)}$$

where $Z_{(k:m)}$ is the k -th order statistics of an i.i.d. sample of m random variables distributed as $U_{(\ell:n)}$. In other words, with the notation of Lemma 27, we have $Z_{(k:m)} \stackrel{d}{=} U_{(\ell:n,k:m)}$, hence

$$F_{U_{(\ell:n)}}(U_{(\ell:n,k:m)}) \stackrel{d}{=} U_{(k:m)} \stackrel{d}{=} F_{U_{(k:m)}}^{-1}(U)$$

since $F_{U_{(k:m)}}$ is continuous (where U follows the uniform distribution over $[0, 1]$). So,

$$U_{(\ell:n,k:m)} \stackrel{d}{=} F_{U_{(\ell:n)}}^{-1} \circ F_{U_{(k:m)}}^{-1}(U) = (F_{U_{(k:m)}} \circ F_{U_{(\ell:n)}})^{-1}(U)$$

has cdf $F_{U_{(k:m)}} \circ F_{U_{(\ell:n)}}$. □

Appendix I: Previous and additional results on related works

I.1. Proofs of Section 2.1 (theoretical analysis of split CP)

In this section, we prove the results stated about split conformal prediction (split CP) in Section 2.1. Although these results are mostly classical, we think useful to prove them here for completeness, since these proofs are particularly simple with the notation introduced in this paper and the results of Appendix B.1.

I.1.1. Marginal guarantees

By Lemma 23, the expected coverage of $\widehat{\mathcal{C}}_r$ is

$$\mathbb{P}(Y \in \widehat{\mathcal{C}}_r(X)) = \mathbb{E}[1 - \alpha_r(\mathcal{D})] = \mathbb{E}[F_S(S_{(r)})] \geq \mathbb{E}[U_{(r:n_c)}]$$

with equality when F_S is continuous —or, equivalently, when the scores are almost surely distinct. Using Eq. (69) in Appendix G.1, we get that

$$\mathbb{P}(Y \in \widehat{\mathcal{C}}_r(X)) \geq \frac{r}{n_c + 1}$$

with equality when F_S is continuous. This proves that $\widehat{\mathcal{C}}_r$ is marginally valid when $r = \lceil (1 - \alpha)(n_c + 1) \rceil$, and Eq. (6) when F_S is continuous.

I.1.2. Training-conditional guarantees

By Lemma 23,

$$\mathbb{P}(Y \in \widehat{\mathcal{C}}_r(X) \mid \mathcal{D}) = 1 - \alpha_r(\mathcal{D}) = F_S(S_{(r)}) \succeq U_{(r:n_c)} \quad (98)$$

which follows the Beta($r, n_c - r + 1$) distribution, hence

$$\mathbb{P}(1 - \alpha_r(\mathcal{D}) \geq F_{U_{(r:n_c)}}^{-1}(\beta)) \geq \mathbb{P}(U_{(r:n_c)} \geq F_{U_{(r:n_c)}}^{-1}(\beta)) = 1 - F_{U_{(r:n_c)}} \circ F_{U_{(r:n_c)}}^{-1}(\beta) = 1 - \beta,$$

which proves Eq. (7). Note that Eq. (7) is similar to [61, Proposition 2b], even if its statement and proof are quite different. Lemma 23 also proves that Eq. (7) is an equality when F_S is continuous, that is, $1 - \alpha_r(\mathcal{D}) \stackrel{d}{=} \text{Beta}(r, n_c - r + 1)$.

Similarly to Algorithm 3 and its generalization given by Eq. (57) in Appendix E.2, in order to have the best possible training-conditionally valid prediction set $\widehat{\mathcal{C}}_r$ with split CP, Eq. (7) suggests the choice

$$\begin{aligned} r_c^* &:= \operatorname{argmin}_{r \in \llbracket n_c \rrbracket} \{ F_{U(r:n_c)}^{-1}(1 - \beta') : F_{U(r:n_c)}^{-1}(\beta) \geq 1 - \alpha \} \\ &= \min \{ r \in \llbracket n_c \rrbracket : F_{U(r:n_c)}^{-1}(\beta) \geq 1 - \alpha \}. \end{aligned} \quad (99)$$

Equality (99) above comes from the fact that $F_{U(r:n_c)}^{-1}(1 - \beta')$ is an increasing function of $r \in \llbracket n_c \rrbracket$ for every $\beta' \in (0, 1)$, hence minimizing $F_{U(r:n_c)}^{-1}(1 - \beta')$ is equivalent to minimizing r , whatever $\beta' \in (0, 1)$.

By Theorem 31 in Appendix G.2, for any $\gamma \in \mathbb{R}$, when $r = r_{n_c} = (1 - \alpha)n_c + \gamma\sqrt{n_c} + o(n_c)$ as $n_c \rightarrow +\infty$, we have

$$\mathbb{P}(1 - \alpha_{r_{n_c}}(\mathcal{D}) \geq 1 - \alpha) \geq \mathbb{P}(U_{(r:n_c)} \geq 1 - \alpha) \quad \text{by Eq. (98)} \quad (100)$$

$$\begin{aligned} &= \mathbb{P}\left(\sqrt{n_c}[U_{(r:n_c)} - (1 - \alpha)] \geq 0\right) \\ &\xrightarrow{n_c \rightarrow +\infty} \Phi\left(\frac{\gamma}{\sqrt{\alpha(1 - \alpha)}}\right) \end{aligned} \quad (101)$$

where Φ denotes the standard normal cdf. Taking $\gamma = \Phi^{-1}(1 - \beta)\sqrt{\alpha(1 - \alpha)}$, this proves that any

$$r = r_{n_c} = (1 - \alpha)n_c + \Phi^{-1}(1 - \beta)\sqrt{\alpha(1 - \alpha)}\sqrt{n_c} + o(\sqrt{n_c}) \in \llbracket n_c \rrbracket \quad (102)$$

is (asymptotically) training-conditionally valid, that is, satisfies Eq. (3). Furthermore, the choice (102) is not improvable in general since Eq. (100) is an equality when F_S is continuous. Then, for any r satisfying Eq. (102) and any $\beta' \in (0, 1)$, when F_S is continuous, by Theorem 31, the $(1 - \beta')$ -quantile of the distribution $\text{Beta}(r_{n_c}, n_c - r_{n_c} + 1)$ of the coverage $1 - \alpha_{r_{n_c}}(\mathcal{D})$ of $\widehat{\mathcal{C}}_{r_{n_c}}$ satisfies

$$F_{U(r_{n_c}:n_c)}^{-1}(1 - \beta') = 1 - \alpha + \frac{\sqrt{\alpha(1 - \alpha)}[\Phi^{-1}(1 - \beta') + \Phi^{-1}(1 - \beta)]}{\sqrt{n_c}} + o\left(\frac{1}{\sqrt{n_c}}\right). \quad (103)$$

In addition, combining Lemma 23 and Eq. (81) shows that

$$r = \left\lceil (n_c + 1) \left(1 - \alpha + \sqrt{\frac{\log(1/\beta)}{2(n_c + 2)}} \right) \right\rceil \quad (104)$$

yields a training-conditionally valid prediction set $\widehat{\mathcal{C}}_r$ whatever $n_c \geq 1$, provided $r \in \llbracket n_c \rrbracket$. For instance, when $n_c \geq \max\{2(\alpha^{-1} - 1), (8/3)\alpha^{-2} \log(1/\beta)\}$, we have $(n_c + 1)(1 - \alpha) \leq (1 - \alpha/2)n_c$ and $n_c\sqrt{n_c + 2}/(n_c + 1) \geq \sqrt{n_c}\sqrt{3/4} \geq \sqrt{2 \log(1/\beta)}/\alpha$, hence $(n_c + 1)\sqrt{\frac{\log(1/\beta)}{2(n_c + 2)}} \leq n_c\alpha/2$, so that $r \in \llbracket n_c \rrbracket$. Note that the choice (104) is similar to what can be deduced from [61, Proposition 2a], up to minor differences, and completely different notation and proof; see also [8, Theorem 1] for a rewriting of this result with notation closer to ours. If r satisfies Eq. (104), when F_S is continuous, Eq. (39) and (41) show

that for every $\beta' \in (0, 1)$, with probability at least $1 - \beta'$, the coverage $1 - \alpha_r(\mathcal{D})$ of $\widehat{\mathcal{C}}_r$ is smaller than

$$F_{U(r;n_c)}^{-1}(1 - \beta') \leq 1 - \alpha + \frac{\sqrt{\log(1/\beta)} + \sqrt{\log(1/\beta')}}{\sqrt{2(n_c + 2)}} + \frac{1}{n_c + 1}. \quad (105)$$

Notice that this upper bound is of order $1 - \alpha + \mathcal{O}(1/\sqrt{n_c})$ as Eq. (103), with only a slightly worse dependence on α and β, β' .

I.2. Additional results: average of quantiles

In the one-shot FL setting, an alternative approach to quantile-of-quantiles is to use an average of quantiles, as proposed by [38]. The idea is that each agent returns its $\lceil (n_j + 1)(1 - \alpha) \rceil$ -th smallest score, and then the central server averages these scores. Using the notation of Section 2.2, the associated prediction set is

$$\widehat{\mathcal{C}}_{\ell_1, \dots, \ell_m}^{\text{Avg}}(x) := \left\{ y \in \mathcal{Y} : s(x, y) \leq \frac{1}{m} \sum_{j=1}^m \widehat{Q}_{(\ell_j)}(\mathcal{S}_j) \right\} \quad (106)$$

with $\ell_j = \lceil (n_j + 1)(1 - \alpha) \rceil$ for every $j \in \llbracket m \rrbracket$. To the best of our knowledge, no coverage guarantee has ever been proved for $\widehat{\mathcal{C}}_{\ell_1, \dots, \ell_m}^{\text{Avg}}$. We study this prediction set in this section, showing that its marginal coverage strongly depends on the scores distribution, so that one cannot always find some $\ell_j \in \llbracket n \rrbracket$ such that the marginal coverage is larger than a given $1 - \alpha$.

I.2.1. First remark

A first simple remark is that when each agent $j \in \llbracket m \rrbracket$ has $n_j = 1$ calibration point, then one must take $\ell_j = 1$ and the marginal coverage of $\widehat{\mathcal{C}}^{\text{Avg}}$ cannot be close to 1 in general since it is equal to

$$\mathbb{E} \left[F_S \left(\frac{1}{m} \sum_{j=1}^m \widehat{Q}_{(1)}(\mathcal{S}_j) \right) \right] = \mathbb{E} \left[F_S \left(\frac{1}{m} \sum_{j=1}^m S_{1,j} \right) \right] \approx F_S(\mathbb{E}[S])$$

(at least when m is large) by the law of large numbers. Therefore, $\widehat{\mathcal{C}}^{\text{Avg}}$ cannot be marginally valid for all distributions without assuming that $n_j \geq n_0(\alpha)$ for some large enough $n_0(\alpha)$, at least for a large part of the agents $j \in \llbracket m \rrbracket$. The next result shows that such an assumption is not sufficient.

I.2.2. Negative result for Bernoulli scores

We now prove that when $n_j = n$ for every $j \in \llbracket m \rrbracket$, for every $\alpha \in (0, 1)$, some distribution of the scores exists such that the marginal coverage of $\widehat{\mathcal{C}}_{\ell_1, \dots, \ell_m}^{\text{Avg}}$ is smaller than $1 - \alpha$ whatever $\ell_j \in \llbracket m \rrbracket$, for arbitrary large m and n . Therefore, contrary to our quantile-of-quantiles prediction sets of Section 3, the ℓ_j cannot be chosen in such a way that $\widehat{\mathcal{C}}_{\ell_1, \dots, \ell_m}^{\text{Avg}}$ is a marginally-valid distribution-free prediction set, even when the data set is large enough.

Proposition 38. *In the setting of Section 2.2, with Assumption 1, for any $\ell_1, \dots, \ell_m \in \llbracket m \rrbracket$, the prediction set $\widehat{\mathcal{C}}_{\ell_1, \dots, \ell_m}^{\text{Avg}}$ defined by Eq. (106) satisfies*

$$\mathbb{P}(Y \in \widehat{\mathcal{C}}_{\ell_1, \dots, \ell_m}^{\text{Avg}}(X)) \leq \mathbb{P}(Y \in \widehat{\mathcal{C}}_{n, \dots, n}^{\text{Avg}}(X)). \quad (107)$$

Assuming in addition that the scores $S_{i,j}, S$ are i.i.d. Bernoulli(p) random variables for some $p \in [0, 1]$, for every $\ell \in \llbracket n \rrbracket$, we have

$$\mathbb{P}(Y \in \widehat{\mathcal{C}}_{\ell, \dots, \ell}^{\text{Avg}}(X)) = (1-p) + p \cdot F_{U_{(n-\ell+1:n)}}(p)^m \quad (108)$$

where $F_{U_{(n-\ell+1:n)}}$ is the cdf of the Beta($n-\ell+1, \ell$) distribution. Therefore, for every $c \in (0, 1)$, if $m = m_n$ and $p = p_n = 1 - \lceil \log(1/c)/m_n \rceil^{1/n}$ are such that $\log(m_n)/n \rightarrow +\infty$ as $n \rightarrow +\infty$, we have

$$\mathbb{P}(Y \in \widehat{\mathcal{C}}_{n, \dots, n}^{\text{Avg}}(X)) = (1-p_n) + p_n [1 - (1-p_n)^n]^{m_n} \xrightarrow[n \rightarrow +\infty]{} c. \quad (109)$$

Eq. (107)–(109) show that for any $\alpha \in (0, 1)$, when the scores follow a Bernoulli(p_n) distribution with

$$p_n = 1 - \left\lceil \frac{\log[2/(1-\alpha)]}{m_n} \right\rceil^{1/n} \quad \text{and} \quad \log(m_n) \gg n \rightarrow +\infty$$

—for instance, when $m_n \geq n^n$ —, the marginal coverage of $\widehat{\mathcal{C}}_{\ell_1, \dots, \ell_m}^{\text{Avg}}$ is strictly smaller than $1 - \alpha$ when n is large enough, whatever the $\ell_j, j \in \llbracket m \rrbracket$. So, $\widehat{\mathcal{C}}^{\text{Avg}}$ cannot be used as a distribution-free marginally-valid prediction set.

Note that the counterexample provided by Proposition 38 relies on the non-robustness of the empirical mean of the $\widehat{Q}_{(\ell_j)}(\mathcal{S}_j)$ defining $\widehat{\mathcal{C}}_{\ell_1, \dots, \ell_m}^{\text{Avg}}$: when the scores follow a Bernoulli distribution, it suffices to have one $\widehat{Q}_{(\ell_j)}(\mathcal{S}_j) \neq 1$ to make the expected coverage strictly smaller than $1 - p$, and such an “outlier” event occurs with a large probability if m is large enough, even when p is close to 1. This phenomenon enlightens the interest of our quantile-of-quantiles approach, which does not have such a drawback thanks to the robustness of the k -th empirical quantile.

Proof of Proposition 38. First, for every $j \in \llbracket m \rrbracket$, $\widehat{Q}_{(\ell_j)}(\mathcal{S}_j)$ is a nondecreasing function of ℓ_j , hence $\widehat{\mathcal{C}}_{\ell_1, \dots, \ell_m}^{\text{Avg}}(X)$ is a nondecreasing function of ℓ_j , which implies Eq. (107).

Second, when the scores are i.i.d. and follow a Bernoulli(p) distribution, for any $i \in \llbracket m \rrbracket$, $S_{(\ell:n),i}$ follows a Bernoulli distribution with parameter

$$\mathbb{P}(S_{(\ell:n),i} = 1) = \sum_{k=n-\ell+1}^n p^k (1-p)^{n-k} \binom{n}{k} = F_{U_{(n-\ell+1:n)}}(p)$$

by Eq. (65). Therefore, we have

$$\begin{aligned} \mathbb{P}(Y \in \widehat{\mathcal{C}}_{\ell, \dots, \ell}^{\text{Avg}}(X)) &= \mathbb{P}\left(S \leq \frac{1}{m} \sum_{j=1}^m S_{(\ell:n),j}\right) \\ &= \mathbb{P}(S = 0) + \mathbb{P}(S = 1 \text{ and } S_{(\ell:n),j} = 1 \text{ for every } j \in \llbracket m \rrbracket) \\ &= (1-p) + p \cdot \mathbb{P}(S_{(\ell:n),1} = 1)^m = (1-p) + p \cdot F_{U_{(n-\ell+1:n)}}(p)^m, \end{aligned}$$

which proves Eq. (108).

Third, when $m = m_n$ is such that $\log(m_n)/n \rightarrow +\infty$ as $n \rightarrow +\infty$, we have $p_n = 1 - [\log(1/c)/m_n]^{1/n} \in (0, 1)$ for n large enough, and $p_n \rightarrow 1$ and $m_n \rightarrow +\infty$ when $n \rightarrow +\infty$. So, Eq. (108) together with Eq. (68) shows that

$$\begin{aligned} \mathbb{P}(Y \in \widehat{\mathcal{C}}_{n, \dots, n}^{\text{Avg}}(X)) &= (1 - p_n) + p_n [1 - (1 - p_n)^n]^{m_n} \\ &= o(1) + [1 + o(1)] \left(1 - \frac{\log(1/c)}{m_n}\right)^{m_n} \\ &\xrightarrow{n \rightarrow +\infty} \exp(-\log(1/c)) = c. \end{aligned}$$

□

I.2.3. Uniform scores

We now study the case of scores following a uniform distribution.

Proposition 39. *Suppose that Assumption 1 holds true. When the scores $S_{i,j}, S$ are independent with a uniform distribution over $[a, b]$ for some $a < b$, for any $\ell \in \llbracket n \rrbracket$, we have*

$$\mathbb{P}(Y \in \widehat{\mathcal{C}}_{\ell, \dots, \ell}^{\text{Avg}}(X)) = \frac{\ell}{n+1}. \quad (110)$$

Proposition 39 shows that when the scores are uniform (and using this knowledge), $\widehat{\mathcal{C}}_{\ell, \dots, \ell}^{\text{Avg}}$ is marginally valid as soon as $\ell \geq (n+1)(1-\alpha)$, which suggests to take $\ell = \lceil (n+1)(1-\alpha) \rceil$. Nevertheless, the difference between the expected coverage and the nominal coverage is then equal to

$$\frac{\lceil (n+1)(1-\alpha) \rceil}{n+1} - (1-\alpha) \in \left[0, \frac{1}{n+1}\right),$$

so its worst-case value $1/(n+1)$ is independent from m , which is not a desirable property.

Proof. By definition of $\widehat{\mathcal{C}}^{\text{Avg}}$, we have

$$\begin{aligned} \mathbb{P}(Y \in \widehat{\mathcal{C}}_{\ell, \dots, \ell}^{\text{Avg}}(X)) &= \mathbb{P}\left(S \leq \frac{1}{m} \sum_{j=1}^m S_{(\ell:n), j}\right) = \mathbb{E}\left[F_S\left(\frac{1}{m} \sum_{j=1}^m S_{(\ell:n), j}\right)\right] \\ &= \mathbb{E}\left[\frac{1}{m} \sum_{j=1}^m F_S(S_{(\ell:n), j})\right] \end{aligned}$$

since $F_S(t) = (t-a)/(b-a)$ for every $t \in [a, b]$. Therefore, we get that

$$\mathbb{P}(Y \in \widehat{\mathcal{C}}_{\ell, \dots, \ell}^{\text{Avg}}(X)) = \mathbb{E}[F_S(S_{(\ell:n), 1})] = \mathbb{E}[U_{(\ell:n)}] = \frac{\ell}{n+1}$$

by Eq. (39) in Appendix B.1 and Eq. (69) in Appendix G.1. □

I.2.4. Exponential scores

We now study the case of scores following an exponential distribution.

Proposition 40. *Suppose that Assumption 1 holds true. When the scores $S_{i,j}, S$ are independent with an exponential distribution with parameter $\lambda > 0$, for any $\ell \in \llbracket n \rrbracket$, we have*

$$\mathbb{P}(Y \in \widehat{\mathcal{C}}_{\ell, \dots, \ell}^{\text{Avg}}(X)) = 1 - \prod_{j=1}^{\ell} \left(1 + \frac{1}{m(n-j+1)}\right)^{-m}. \quad (111)$$

Proposition 40 shows that when the scores are exponential (and using this knowledge), it is possible to choose ℓ such that the expected coverage of $\widehat{\mathcal{C}}_{\ell, \dots, \ell}^{\text{Avg}}$ is larger than $1 - \alpha$, provided that m and n are large enough. Indeed, by Eq. (111),

$$\begin{aligned} & \max_{\ell \in \llbracket n \rrbracket} \mathbb{P}(Y \in \widehat{\mathcal{C}}_{\ell, \dots, \ell}^{\text{Avg}}(X)) \\ &= \mathbb{P}(Y \in \widehat{\mathcal{C}}_{n, \dots, n}^{\text{Avg}}(X)) \\ &= 1 - \prod_{j=1}^n \left(1 + \frac{1}{m(n-j+1)}\right)^{-m} \xrightarrow{m \rightarrow +\infty} 1 - \exp\left(-\sum_{j=1}^n \frac{1}{j}\right) \xrightarrow{n \rightarrow +\infty} 1, \end{aligned}$$

where the first limit is taken for any fixed $n \geq 1$.

Proof of Proposition 40. A classical result about the order statistics of a sample of exponential random variables [53] shows that for any $i \in \llbracket m \rrbracket$,

$$S_{(\ell:n), i} \stackrel{d}{=} \frac{1}{\lambda} \sum_{j=1}^{\ell} \frac{E_{j,i}}{n-j+1}$$

where the $(E_{j,i})_{1 \leq j \leq n}$ are independent standard exponential variables. Therefore, we have

$$\begin{aligned} \mathbb{P}(Y \in \widehat{\mathcal{C}}_{\ell, \dots, \ell}^{\text{Avg}}(X)) &= \mathbb{E} \left[F_S \left(\frac{1}{m} \sum_{j=1}^m S_{(\ell:n), j} \right) \right] \\ &= 1 - \mathbb{E} \left[\exp \left(-\frac{\lambda}{m} \sum_{i=1}^m S_{(\ell:n), i} \right) \right] \\ &= 1 - \mathbb{E} \left[\exp \left(-\frac{1}{m} \sum_{i=1}^m \sum_{j=1}^{\ell} \frac{E_{j,i}}{n-j+1} \right) \right] \\ &= 1 - \mathbb{E} \left[\exp \left(-\sum_{j=1}^{\ell} \sum_{i=1}^m \frac{E_{j,i}}{m(n-j+1)} \right) \right] \\ &= 1 - \prod_{j=1}^{\ell} \mathbb{E} \left[\exp \left(-\sum_{i=1}^m \frac{E_{j,i}}{m(n-j+1)} \right) \right] = 1 - \prod_{j=1}^{\ell} \left(1 + \frac{1}{m(n-j+1)}\right)^{-m}. \end{aligned}$$

For the last equality, since $1/(m(n-j+1))$ is positive, $E_{j,i}/[m(n-j+1)]$ follows an exponential distribution with parameter $m(n-j+1)$. Therefore, $\sum_{i=1}^m \frac{E_{j,i}}{m(n-j+1)}$ follows a Gamma distribution with parameter $(m, m(n-j+1))$ and the value at -1 of its moment generating function is equal to $(1 + \frac{1}{m(n-j+1)})^{-m}$ since $-1 < m(n-j+1)$. \square

Appendix J: Additional experimental results

J.1. Generic comparison, equal n_j

In this section, we provide additional results in the generic setting of Section 6.1.

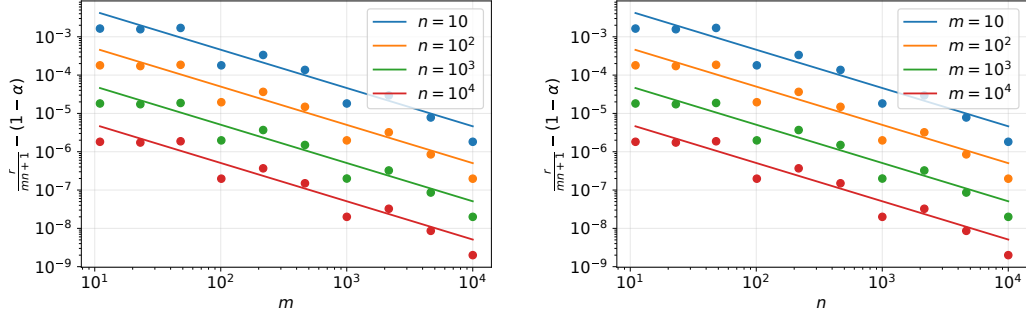


FIG 10. *CentralM*: log-log plot of $\Delta\mathbb{E}$ as a function of m (left) or n (right). Lines show the approximation $\log \Delta\mathbb{E} \approx \log(c_1) - \gamma_1 \log(m) - \delta_1 \log(n)$ with c_1, γ_1, δ_1 given by Table 1.

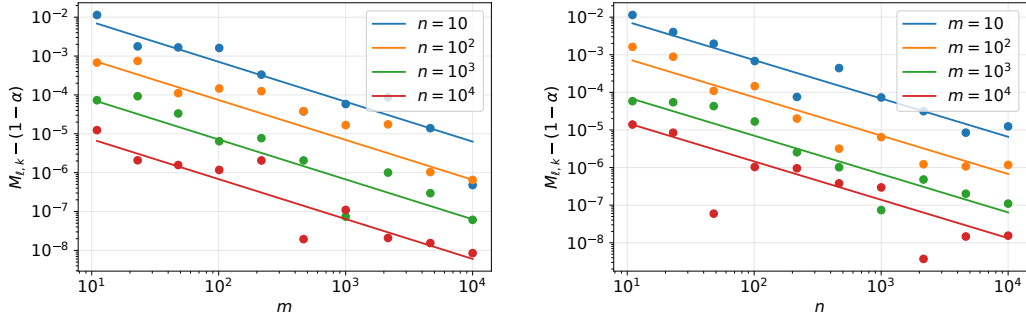


FIG 11. *QQM*: log-log plot of $\Delta\mathbb{E}$ as a function of m (left) or n (right). Lines show the approximation $\log \Delta\mathbb{E} \approx \log(c_1) - \gamma_1 \log(m) - \delta_1 \log(n)$ with c_1, γ_1, δ_1 given by Table 1.

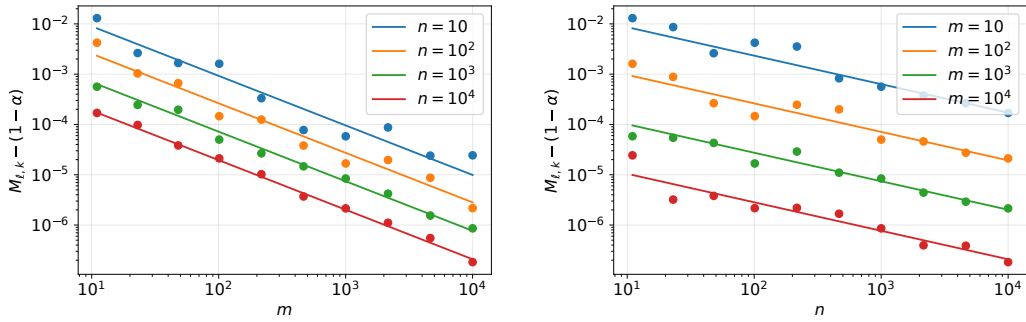


FIG 12. *QQM-Fast*: log-log plot of $\Delta\mathbb{E}$ as a function of m (left) or n (right). Lines show the approximation $\log \Delta\mathbb{E} \approx \log(c_1) - \gamma_1 \log(m) - \delta_1 \log(n)$ with c_1, γ_1, δ_1 given by Table 1.

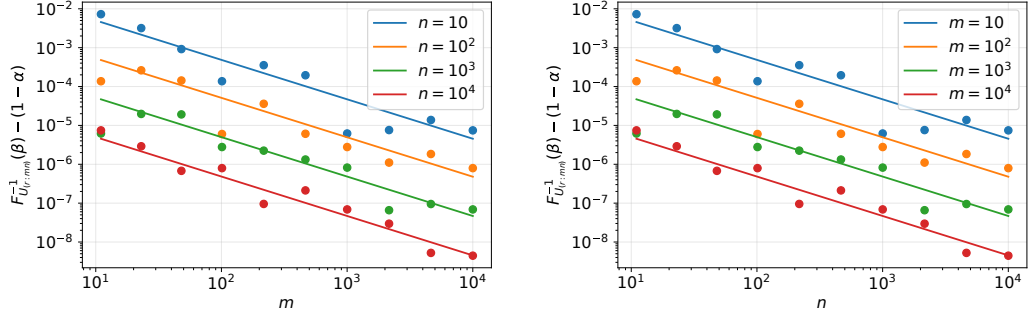


FIG 13. CentralC: log-log plot of Δq_β as a function of m (left) or n (right). Lines show the approximation $\log \Delta q_\beta \approx \log(c_2) - \gamma_2 \log(m) - \delta_2 \log(n)$ with c_2, γ_2, δ_2 given by Table 1.

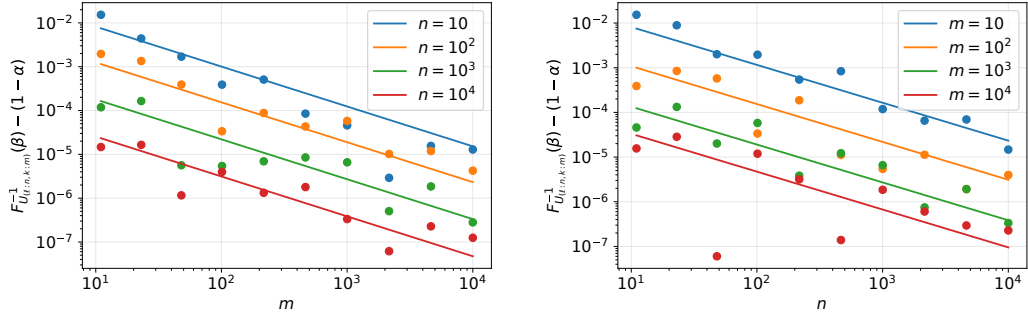


FIG 14. QQC: log-log plot of Δq_β as a function of m (left) or n (right). Lines show the approximation $\log \Delta q_\beta \approx \log(c_2) - \gamma_2 \log(m) - \delta_2 \log(n)$ with c_2, γ_2, δ_2 given by Table 1.

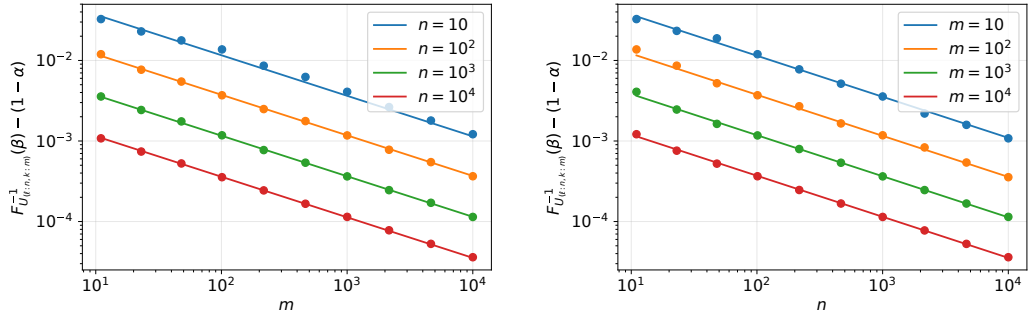


FIG 15. QQC-Fast: log-log plot of Δq_β as a function of m (left) or n (right). Lines show the approximation $\log \Delta q_\beta \approx \log(c_2) - \gamma_2 \log(m) - \delta_2 \log(n)$ with c_2, γ_2, δ_2 given by Table 1.




FIG 16. CentralC: log-log plot of $\Delta q_{1-\beta}$ as a function of m (left) or n (right). Lines show the approximation $\log \Delta q_{1-\beta} \approx \log(c_3) - \gamma_3 \log(m) - \delta_3 \log(n)$ with c_3, γ_3, δ_3 given by Table 1.

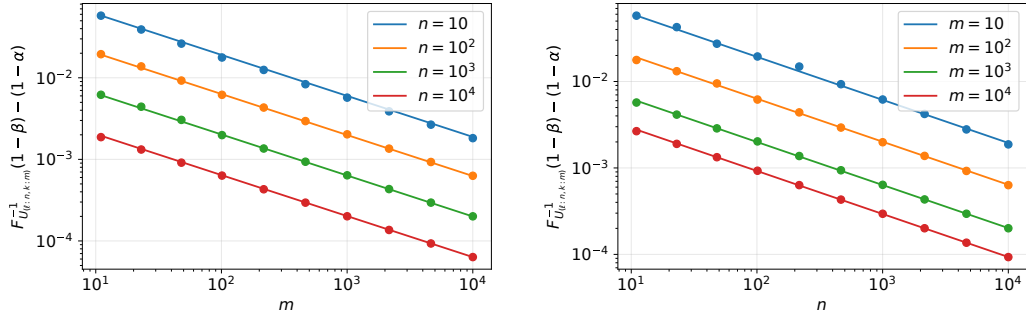


FIG 17. QQC: log-log plot of $\Delta q_{1-\beta}$ as a function of m (left) or n (right). Lines show the approximation $\log \Delta q_{1-\beta} \approx \log(c_3) - \gamma_3 \log(m) - \delta_3 \log(n)$ with c_3, γ_3, δ_3 given by Table 1.




FIG 18. QQC-Fast: log-log plot of $\Delta q_{1-\beta}$ as a function of m (left) or n (right). Lines show the approximation $\log \Delta q_{1-\beta} \approx \log(c_3) - \gamma_3 \log(m) - \delta_3 \log(n)$ with c_3, γ_3, δ_3 given by Table 1.




FIG 19. Conditionally-valid algorithms: log-log plot of $\Delta q_{1-\beta}$ as a function of m (left) or n (right). Lines show the approximation $\log \Delta q_{1-\beta} \approx \log(c_3) - \gamma_3 \log(m) - \delta_3 \log(n)$ with c_3, γ_3, δ_3 given by Table 1. Plain lines and dots correspond to $n = 10$ (left) or $m = 10$ (right). Dashed lines and diamonds correspond to $n = 10^4$ (left) or $m = 10^4$ (right).

J.2. Generic comparison, different n_j

This subsection provides experimental results that complement the ones of Section 6.1.2.

J.2.1. With random n_j

We start with experiments using n_j (random) that are not equal. Figure 20 gives the values of the n_j for $m \in \{4, 25\}$ in the experiments of Section 6.1.2. Figure 21 provides an equivalent of Figure 6 right (performance of conditionally-valid algorithms as a function of m) with $\Delta q_{1-\beta}$ instead of Δq_β .




FIG 20. Experiment of Section 6.1.2: values of $(n_j)_{j \in [m]}$ in increasing order. Left: $m = 4$. Right: $m = 25$.




FIG 21. Different n_j : log-log plot of the performance ($\Delta q_{1-\beta}$) of the conditionally-valid algorithms considered in Section 6.1.2 as a function of the number of agents m . The total number of data points $N = \sum_{j=1}^m n_j = 4000$ is fixed and m describes the set of divisors of N . The values when $m = 1, 2$ for $QQC-(N/m)$ and $QQC-n_j$ are missing because there is no k such that the resulting sets are conditionally valid.

J.2.2. Algorithms 5 and 6 with $n_j = N/m$

As in Section 6.1.2, we now consider Algorithms 5 and 6 with $n_j = N/m$ for every $h \in \llbracket m \rrbracket$, that we call $QQM-(N/m)$ and $QQC-(N/m)$. Table 4 gives the values of the coefficients of the robust log-linear regression obtained for $QQM-(N/m)$ and $QQC-(N/m)$ (following the exact same procedure as for QQM , QQM -Fast, QQC , and QQC -Fast, as described in Section 6.1.1). In Figure 22, we compare the performances (measured by $\Delta \mathbb{E}$) of the marginally-valid algorithms QQM , QQM -Fast, and $QQM-(N/m)$, as a function of either m or $n := N/m$. In Figure 23, we compare the performances (measured by Δq_β) of the conditionally valid algorithms QQC , QQC -Fast, and $QQC-(N/m)$, as a function of either m or $n = N/m$.

Method	$\Delta \mathbb{E} \approx c_1 m^{-\gamma_1} n^{-\delta_1}$			$\Delta q_\beta \approx c_2 m^{-\gamma_2} n^{-\delta_2}$			$\Delta q_{1-\beta} \approx c_3 m^{-\gamma_3} n^{-\delta_3}$		
	c_1	γ_1	δ_1	c_2	γ_2	δ_2	c_3	γ_3	δ_3
$QQM-(N/m)$	0.248	0.952	0.521	0.259	0.467	0.509	0.417	0.520	0.511
$QQC-(N/m)$	0.408	0.521	0.507	0.462	0.992	0.552	0.713	0.507	0.505

TABLE 4

Estimated parameters of the log-linear model $\log y = \log(c_i) - \gamma_i \log(m) - \delta_i \log(n)$ where y is either $\Delta \mathbb{E}$, Δq_β or $\Delta q_{1-\beta}$, for $QQM-(N/m)$ and $QQC-(N/m)$, introduced in Section 6.1.2; See also text of Section 6.1.1 for details.

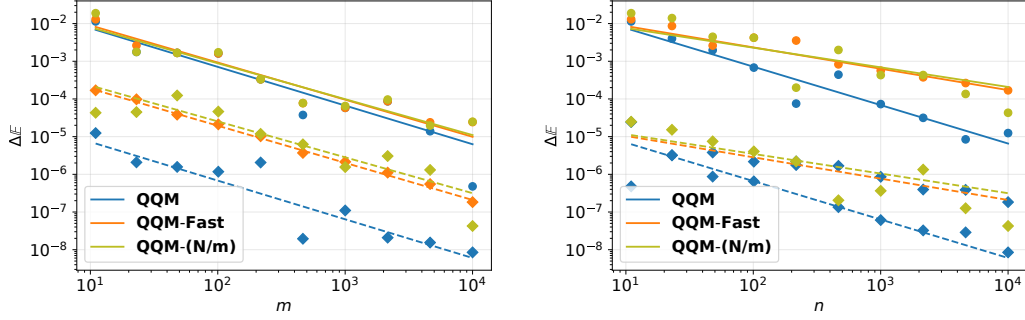


FIG 22. Marginally-valid algorithms: log-log plot of ΔE as a function of m (left) or n (right). Lines show the approximation $\log \Delta E \approx \log(c) - \gamma \log(m) - \delta \log(n)$. Plain lines and dots correspond to $n = 10$ (left) or $m = 10$ (right). Dashed lines and diamonds correspond to $n = 10^4$ (left) or $m = 10^4$ (right).

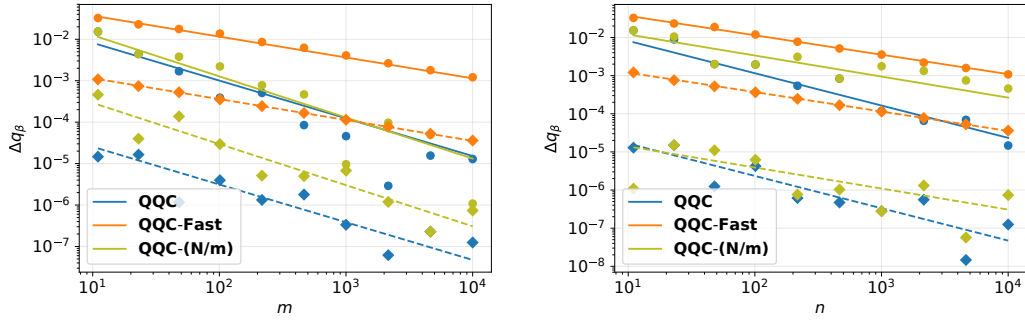


FIG 23. Conditionally-valid algorithms: log-log plot of Δq_β as a function of m (left) or n (right). Lines show the approximation $\log \Delta q_\beta \approx \log(c) - \gamma \log(m) - \delta \log(n)$. Plain lines and dots correspond to $n = 10$ (left) or $m = 10$ (right). Dashed lines and diamonds correspond to $n = 10^4$ (left) or $m = 10^4$ (right).

J.3. Real data: additional results on individual data sets

In this section, we present in Figures 24 to 31 the results of the experiments of Section 6.2 on individual data sets.




FIG 24. Coverage (left) and average length (right) of prediction intervals for 50 random learning-calibration-test splits. The misscoverage is $\alpha = 0.1$, $\beta = 0.2$, and the calibration set is split into $m = 80$ disjoint subsets of equal size $n = 10$. The white circle represents the mean and the name of the data set is located at the top of each plot.




FIG 25. Coverage (left) and average length (right) of prediction intervals for 50 random learning-calibration-test splits. The misscoverage is $\alpha = 0.1$, $\beta = 0.2$, and the calibration set is split into $m = 10$ disjoint subsets of equal size $n = 80$. The white circle represents the mean and the name of the data set is located at the top of each plot.




FIG 26. Same as Figure 24 (see its caption) with $m = 80$ and $n = 10$.




FIG 27. Same as Figure 25 (see its caption) with $m = 10$ and $n = 80$.




FIG 28. Same as Figure 24 (see its caption) with $m = 80$ and $n = 10$.




FIG 29. Same as Figure 25 (see its caption) with $m = 10$ and $n = 80$.




FIG 30. Same as Figure 24 (see its caption) with $m = 40$ and $n = 10$.




FIG 31. Same as Figure 25 (see its caption) with $m = 10$ and $n = 40$.