

000 FEDCF: FAIR FEDERATED CONFORMAL PREDICTION

001
002
003 **Anonymous authors**

004 Paper under double-blind review

005 006 007 ABSTRACT

008
009
010 Conformal Prediction (CP) is a widely used technique for quantifying uncertainty
011 in machine learning models. In its standard form, CP offers probabilistic guaran-
012 tees on the coverage of the true label, but it is agnostic to sensitive attributes in the
013 dataset. Several recent works have sought to incorporate fairness into CP by ensur-
014 ing conditional coverage guarantees across different subgroups. One such method
015 is Conformal Fairness (CF). In this work, we extend the CF framework to the Fed-
016 erated Learning setting and discuss how we can audit a federated model for fair-
017 ness by analyzing the fairness-related gaps for different demographic groups. We
018 empirically validate our framework by conducting experiments on several datasets
019 spanning multiple domains, fully leveraging the exchangeability assumption.

020 021 1 INTRODUCTION

022
023 Ensuring model fairness is a critical thrust of trustworthy machine learning (ML). ML models, when
024 not calibrated for fairness, are prone to developing biases at each stage of an ML pipeline, as re-
025 flected by their predictions Mehrabi et al. (2021). We define bias as disparate performance (i.e.,
026 accuracy for classification) between different sub-populations. In the data collection phase, mea-
027 surement bias may occur due to disproportionate data collection on sub-populations, while repre-
028 sentation bias manifests from a lack of training data on specific strata. During training, these biases
029 are inductively learned by the model—leading to incorrect predictions in safety-critical tasks. These
030 models are also susceptible to algorithmic bias, resulting from regularization and optimization tech-
031 niques during model training, which incorrectly generalize for marginalized groups. To mitigate
032 these risks, many ML models must adhere to regulations placed by local governing bodies (Hirsch
033 et al., 2023). Towards model compliance, Komala et al. (2024); Agrawal et al. (2024); Jones et al.
034 (2025) have proposed approaches to enhance model fairness in varying tasks, including federated
035 graph learning and representation learning.

036 Developing robust ML frameworks with mathematically rigorous guarantees is also essential for
037 building actionable, trustworthy ML models for safety-critical tasks. In this frontier, researchers
038 have increasingly explored Conformal Prediction (CP)—an uncertainty quantification (UQ) technique
039 that only assumes statistical exchangeability—to develop trustworthy ML models (Vovk et al., 2005).
040 Unlike traditional point-wise prediction in ML, CP guarantees that the correct outcome will be in
041 a prediction set with a user-specified property. Practitioners have adopted CP due to its model-free
042 assumption and post-hoc application (Cherian & Bronner, 2020). Additionally, users can apply
043 CP to structured data (such as graphs), which cannot be used with traditional IID-based methods
044 (Maneriker et al., 2025). However, vanilla CP is not calibrated for fairness and can be inherently
045 unfair Cresswell et al. (2025).

046 Several approaches have been proposed at the intersection of fairness and CP—each catering to differ-
047 ent tasks and notions of fairness. Romano et al. (2020a) developed a CP approach for the regression
048 setting to ensure equalized coverage across protected groups. Lu et al. (2022) considers equalized
049 coverage in a classification task for medical imaging. Zhou & Sesia (2024) extends Romano et al.
050 (2020a) and provides an algorithm adaptive to sensitive groups to increase the predictive power
051 of the CP sets when several sensitive attributes are present, and focuses on the classification task.
052 Lastly, Vadlamani et al. (2025) provides a framework to ensure fair coverage of positive outcomes,
053 without requiring protected attributes at inference time, unlike prior work. Orthogonally, CP has
been used to enhance the fairness of other tasks. To mitigate bias in LLM-based recommender
systems, Fayyazi et al. (2025) explores iteratively using fairness-aware CP.

054 While there are several approaches to integrating fairness into CP, these methods are not considered
 055 when the training data is decentralized (i.e., available only to clients) and the ML model is stored
 056 on a centralized server. Extending these CP methods to the federated learning (FL) setting is es-
 057 sential because tasks that benefit from fair uncertainty quantification (such as those in healthcare and
 058 finance) also often have privacy considerations, making it infeasible to keep data on a centralized
 059 server. Thus, we extend the work of Vadlamani et al. (2025) using recent literature in Federated CP
 060 (Lu et al., 2023).

061 **Key Contributions:** We develop Federated Conformal Fairness (FedCF) and extend the Conformal
 062 Fairness (CF) framework (Vadlamani et al., 2025) to the FL setting while maintaining the theoretical
 063 fairness guarantees provided by CF.

064 **Extending CF Theory to FL.** We discuss how to bound conditional coverage according to a user-
 065 specified fairness notion when data is decentralized. To facilitate this, we provide a sufficient set
 066 of terms that a client can compute using local data, how the server should aggregate these terms to
 067 bound the conditional coverage, and theoretically prove the validity of our approach.

068 **Descent-Based CF Formulation.** We revise the original CF algorithm to reduce the number of
 069 communication rounds required to construct a fair conformal predictor. We do this by reformulating
 070 the original iterative approach presented in Vadlamani et al. (2025) into a descent-based optimization
 071 framework. This allows FedCF to be embedded directly into an FL-style algorithm and to integrate
 072 naturally with existing FL systems.

073 **Flexible Aggregation Protocols.** We consider the client-server communication overhead and its
 074 tradeoff with preserving data privacy. Specifically, we propose two approaches, with one having
 075 less communication overhead and the other being more privacy-preserving of client data.

076 **Real-World Empirical Validation.** We evaluate FedCF on several datasets with naturally induced
 077 *data heterogeneity*, including tabular, graph, and image datasets, and for multiple popular fairness
 078 metrics, and observe that FedCF can control for a particular coverage gap level while maintaining
 079 the original CP coverage guarantee.

082 2 BACKGROUND

084 2.1 CONFORMAL PREDICTION

086 Conformal Prediction (CP) (Vovk et al., 2005) is a widely used framework for quantifying predictive
 087 uncertainty in ML. CP provides rigorous statistical guarantees without imposing assumptions on the
 088 model, requiring only that the data are *exchangeable*—a broader condition than IID and compatible
 089 with non-IID or structured settings (e.g., graphs).

090 We focus on split (inductive) CP in the classification setting. Let $\mathbf{x}_i \in \mathcal{X} = \mathbb{R}^d$ and $y_i \in \mathcal{Y} =$
 091 $\{0, \dots, C-1\}$ denote features and labels. Given a calibration dataset, $\mathcal{D}_{\text{calib}} = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$, our
 092 goal is to construct a set-valued predictor \mathcal{C} such that, for an exchangeable test point $(\mathbf{x}_{\text{test}}, y_{\text{test}})$,

$$093 1 - \alpha \leq \Pr[y_{\text{test}} \in \mathcal{C}(\mathbf{x}_{\text{test}})] \leq 1 - \alpha + \frac{1}{|\mathcal{D}_{\text{calib}}|}, \quad (1)$$

095 where $1 - \alpha \in (0, 1)$ is the target *coverage level*. We refer to Equation 1 as the *coverage guarantee*.
 096 Concretely, given a non-conformity score $s : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$, define the *conformal quantile* as

$$097 \hat{q}(\alpha) = \text{Quantile}\left(\frac{\lceil (n+1)(1-\alpha) \rceil}{n}; \{s(\mathbf{x}_i, y_i)\}_{i=1}^n\right).$$

099 The resulting prediction set $\mathcal{C}_{\hat{q}(\alpha)}(\mathbf{x}_{\text{test}}) = \{y \in \mathcal{Y} : s(\mathbf{x}_{\text{test}}, y) \leq \hat{q}(\alpha)\}$ satisfies the guarantee in 1.

100 **Evaluating CP:** Two standard metrics are used: (1) *Coverage*, the estimated test-time probability,
 101 $\Pr[y_{\text{test}} \in \mathcal{C}_{\hat{q}(\alpha)}(\mathbf{x}_{\text{test}})]$; and (2) *Efficiency*, the average prediction set size, $|\mathcal{C}_{\hat{q}(\alpha)}(\mathbf{x}_{\text{test}})|$. These are
 102 typically in tension as achieving a higher desired coverage often necessitates larger sets.

104 2.2 FEDERATED LEARNING

106 A key contributor to developing strong deep learning models is providing a large amount of qual-
 107 ity training data (Kaplan et al., 2020). However, in certain domains, such as healthcare and fi-
 108 nance, collecting large amounts of data may be prohibitive due to privacy concerns. Federated

learning (FL) McMahan et al. (2017) is a framework for collaborative learning that keeps training data decentralized and private. Given K clients that will participate in training, each client has its own training data that it wants to keep private. The goal is to optimize a global loss function, L , that is the weighted average of local risk functions, ℓ_k . Formally, FL finds weights θ^* s.t. $\theta^* = \operatorname{argmin}_{\theta} \left\{ L(\theta) = \sum_{k=1}^K w_k \cdot \mathbb{E}_{(x^{(k)}, y^{(k)}) \sim P_k} [\ell_k(\theta; x^{(k)}, y^{(k)})] \right\}$, where P_k is client k 's local distribution and $w \in \Delta^k$ are weights (Lu et al., 2023).

2.3 FEDERATED CONFORMAL PREDICTION (FCP)

Setting: In FL, the development and calibration datasets are partitioned over K clients. Meaning, each client $k \in \{1, \dots, K\} = \mathcal{K}$ retains a private calibration set $\mathcal{D}_{\text{calib}}^{(k)} = \left\{ (x_i^{(k)}, y_i^{(k)}) \right\}_{i=1}^{n_k}$ drawn from an unknown local distribution P_k . The goal is to still construct a prediction set function \mathcal{C} such that for any test point $(x_{\text{test}}, y_{\text{test}}) \sim Q_{\text{test}}$, where $Q_{\text{test}} = \sum_{k=1}^K \gamma_k P_k$ is the mixutre distribution with weights $\gamma_k \propto (n_k + 1)$ (Lu et al., 2023), Equation 1 is satisfied. This is done while respecting the communication and privacy constraints of FL.

Partial exchangeability and the FCP algorithm. Lu et al. (2023) introduce partial exchangeability: within each client, the multiset $\left\{ s(x_1^{(k)}, y_1^{(k)}), \dots, s(x_{n_k}^{(k)}, y_{n_k}^{(k)}), s(x_{\text{test}}, y_{\text{test}}) \right\}$ is exchangeable with probability γ_k . Under this assumption, the FCP method aggregates all non-conformity scores, orders them, and selects the $(1 - \alpha)(N + K)$ -th statistic as follows

$$\hat{q}(\alpha) = \operatorname{Quantile} \left(\frac{\lceil (N + K)(1 - \alpha) \rceil}{N}; \left\{ (x_i^{(k)}, y_i^{(k)}) \right\}_{k,i} \right)$$

where $N = \sum_{k=1}^K n_k$. The prediction set $C_{\alpha}(\mathbf{x}) = \{y : S(\mathbf{x}, y) \leq \hat{q}_{\alpha}\}$ then satisfies

$$1 - \alpha \leq \Pr[y_{\text{test}} \in C_{\alpha}(\mathbf{x}_{\text{test}})] \leq 1 - \alpha + \frac{K}{N+K}. \quad (2)$$

Communication-efficient quantile sketches. To preserve privacy, instead of transmitting all N scores to the server, each client can send a mergeable sketch (e.g., using T-Digest (Dunning, 2021) or DDSketch (Masson et al., 2019)). Doing so will loosen the guarantee given in Equation 2.

2.4 CONFORMAL FAIRNESS

While CP provides marginal coverage guarantees, it is agnostic to sensitive attributes within the data. So different groups can receive systematically different coverages. The Conformal Fairness (CF) framework (Vadlamani et al., 2025) formalizes the notion of fairness for prediction sets by considering the disparity in conditional coverage between sensitive groups, all while retaining the validity based on the CP exchangeability. At a high level, CF adapts group-fairness notions (e.g., Demographic Parity and Equal Opportunity) to the set-valued outputs of CP and then tunes a score threshold to satisfy a user-specified “closeness” criterion, c , on inter-group disparities. Using the exchangeability assumption, CF can be applied to non-IID/structured data.

From point predictions to set-based fairness. Let $\mathcal{C}_{\lambda}(\mathbf{x}) = \{y \in \mathcal{Y} : s(\mathbf{x}, y) \leq \lambda\}$ denote the CP prediction set at score threshold λ . CF adapts classical group-fairness metrics by replacing point-prediction events (i.e., $\tilde{y} = \hat{Y}$) with set-membership events (i.e., $\tilde{y} \in \mathcal{C}_{\lambda}(X)$) and evaluating disparities across groups \mathcal{G} and, when appropriate, advantaged labels \mathcal{Y}^+ . For example, a set-based Demographic Parity-style constraint can be written as,

$$|\Pr[\tilde{y} \in \mathcal{C}_{\lambda}(X) | X \in g_a] - \Pr[\tilde{y} \in \mathcal{C}_{\lambda}(X) | X \in g_b]| \leq c \quad \forall g_a, g_b \in \mathcal{G}, \tilde{y} \in \mathcal{Y}^+,$$

with analogous set-based forms for other common group-fairness metrics.

Conditional coverage as the fairness control knob. To evaluate a chosen fairness notion, CF filters the calibration data to the relevant subpopulation (e.g., a group or a group-and-label slice) via a filter function, F_m , and uses conditional coverage estimates under that filter. It then searches a threshold space Λ to identify λ_{opt} that satisfies the closeness criterion across groups (and labels, if required) while maintaining CP validity. A key technical ingredient is that CP coverage holds when labels are fixed to a particular \tilde{y} , which underpins group- and class-conditional control in CF.

Guarantees and trade-offs. CF provides a theoretically grounded procedure to bound fairness disparities (as defined above) without sacrificing CP's finite-sample coverage guarantees, which is

Table 1: Important notation used for coverage gap calculation.

Notation	Definition	Notation	Definition
n_k	$\left \mathcal{D}_{\text{calib}}^{(k)} \right $	$\mathcal{D}_{\text{calib}}^{(k)}$	Client k 's calibration dataset.
$n_k^{(g, \tilde{y})}$	$\left \mathcal{S}_k^{(g, \tilde{y})} \right $	$\mathcal{S}_k^{(g, \tilde{y})}$	$\left\{ (\mathbf{x}_i, y_i) \in \mathcal{D}_{\text{calib}}^{(k)} \mid F_M(\mathbf{x}_i, y_i, g, \tilde{y}) = 1 \right\}$
γ_k	$\Pr[E_k]$	E_k	The event \mathbf{x}_{test} is exchangeable with $\mathcal{D}_{\text{calib}}^{(k)}$.
$\pi^{(g, \tilde{y})}$	Point estimate for Term (IV) in Equation 3.	$L^{(g, \tilde{y})}, U^{(g, \tilde{y})}$	Bounds for Term (IV) in Equation 3.
		$\alpha_k^{(g, \tilde{y}); \lambda}$	$\sum_{(\mathbf{x}_i, -) \in \mathcal{S}_k^{(g, \tilde{y})}} \mathbf{1}[s(\mathbf{x}_i, \tilde{y}) \leq \lambda]$

empirically backed by CF reducing fairness violations across several metrics and remains effective with multiple sensitive attributes (intersectional groups). In practice, satisfying stricter fairness closeness c increases the average prediction set size, reflecting the fairness–efficiency trade-off.

Practical advantages. Unlike many conditional-CP baselines that require group membership at inference time or are model-specific, CF's set-based metrics and thresholding procedure do not require protected attributes at test time and apply across different non-conformity scores and data modalities, making it compatible with downstream deployment constraints.

3 FEDCF THEORY AND METHODOLOGY

In this section, we begin by establishing the theoretical and methodological foundations. We first redefine the concept of the *coverage gap* (3.1) and introduce a descent-based reformulation of the CF Framework (3.2), both within the federated setting. Following this, we present the Federated Conformal Fairness (FedCF) Framework.

3.1 FEDCF: EXTENDING COVERAGE GAP TO THE FEDERATED SETTING

Let F_M be a filter function for some fairness metric. Then, we define the fairness-specific coverage level for positive label $\tilde{y} \in \mathcal{Y}^+$ in group $g \in \mathcal{G}$ as $\Pr[s(\mathbf{x}_{\text{test}}, \tilde{y}) \leq \lambda \mid F_M(\mathbf{x}_{\text{test}}, y_{\text{test}}, g, \tilde{y}) = 1]$. In the federated setting, since the data is decentralized, we cannot directly estimate this quantity. To address this, we rewrite the quantity—using notation in Table 1—as,

$$\Pr[s(\mathbf{x}_{\text{test}}, \tilde{y}) \leq \lambda \mid F_M(\mathbf{x}_{\text{test}}, y_{\text{test}}, g, \tilde{y}) = 1] = \sum_{k=1}^K \left(\underbrace{\Pr[s(\mathbf{x}_{\text{test}}, \tilde{y}) \leq \lambda \mid F_M(\mathbf{x}_{\text{test}}, y_{\text{test}}, g, \tilde{y}) = 1, E_k]}_{\text{I}} \right. \\ \left. \cdot \underbrace{\Pr[F_M(\mathbf{x}_{\text{test}}, y_{\text{test}}, g, \tilde{y}) = 1 \mid E_k]}_{\text{II}} \cdot \underbrace{\Pr[E_k]}_{\text{III}} \cdot \underbrace{\left(\Pr[F_M(\mathbf{x}_{\text{test}}, y_{\text{test}}, g, \tilde{y}) = 1] \right)^{-1}}_{\text{IV}} \right). \quad (3)$$

With this reformulation, we can estimate different terms individually—either locally on each client or globally on the server. The following theorem presents two types of fairness-specific conditional coverage estimates: (1) **interval bounds** and (2) **point estimates**. The estimates for the individual terms are provided in Lemmas B.1, B.2, and B.3 in Appendix B. For clarity, Table 1 summarizes the primary notation used in our main theorem. A full notations table can be found in Appendix A.

Theorem 3.1. *The fairness-specific coverage level (Equation 3) can be bounded as*

$$L_{\text{cov}}(\lambda, F_M, g, \tilde{y}) \leq \Pr[s(\mathbf{x}_{\text{test}}, \tilde{y}) \leq \lambda \mid F_M(\mathbf{x}_{\text{test}}, y_{\text{test}}, g, \tilde{y}) = 1] \leq U_{\text{cov}}(\lambda, F_M, g, \tilde{y}),$$

where

$$L_{\text{cov}}(\lambda, F_M, g, \tilde{y}) = \sum_{k=1}^K \frac{\gamma_k \alpha_k^{(g, \tilde{y}); \lambda} n_k^{(g, \tilde{y})}}{(n_k^{(g, \tilde{y})} + 1)(n_k + 1)U^{(g, \tilde{y})}} \text{ and } U_{\text{cov}}(\lambda, F_M, g, \tilde{y}) = \sum_{k=1}^K \frac{\gamma_k (\alpha_k^{(g, \tilde{y}); \lambda} + 1)}{(n_k + 1)L^{(g, \tilde{y})}}. \quad (4)$$

216 *If the data is IID, using MLE estimates for each term, we get the following estimate for the fairness-*
 217 *specific coverage level*

$$219 \quad \Pi_{cov} = \Pr[s(\mathbf{x}_{test}, \tilde{y}) \leq \lambda \mid F_M(\mathbf{x}_{test}, y_{test}, g, \tilde{y}) = 1] = \sum_{k=1}^K \frac{\gamma_k \alpha_k^{(g, \tilde{y}); \lambda}}{n_k \pi^{(g, \tilde{y})}}. \quad (5)$$

222 Theorem 3.1 gives us bounds for the coverage level, which we can convert into bounds for the
 223 coverage gap between groups for fairness evaluation. The interval bounds provide finite sample
 224 guarantees that are typically seen in the CP literature, at the cost of being a more conservative
 225 estimate. Conversely, point estimates provide tight coverage estimates, but assume IID data, and
 226 may violate those guarantees.

228 3.2 FEDCF: REVISITING THE CONFORMAL FAIRNESS ALGORITHM

231 **Algorithm 1** Descent-Based CF Optimization

```

232 1: procedure FAIR_OPT_DESCENT( $\tilde{y}$ ,  $\lambda_0$ ,  $\mathcal{G}$ ,  $c$ ,  

233    $F_M$ , num_rounds,  $\eta$ ,  $\mu$ )
234   2:  $\lambda_{opt} = 1$ 
235   3: for  $(t = 0; t++; t < \text{num\_rounds}) do
236   4:    $cg_t = \text{coverage\_gap}(\lambda_0, F_M, \tilde{y}, \mathcal{G})$ 
237   5:   if  $cg_t \leq c$  and  $t = 0$  then
238   6:     return  $\lambda_0$ 
239   7:   else if  $cg_t \leq c$  and  $\lambda_t < \lambda_{opt}$  then
240   8:      $\lambda_{opt} = \lambda_t$ 
241   9:   end if
242   10:   $b_{t+1} = \mu \cdot b_t + (cg_t - c)$ 
243   11:   $\Delta\lambda = (\lambda_{opt} - \lambda_t) \mathbf{1}_{[b_{t+1} \geq 0]}$ 
244   12:   $+ (\lambda_t - \lambda_0) \mathbf{1}_{[b_{t+1} < 0]}$ 
245   13:   $p_t = \max\{\lceil \log_2(\frac{\eta}{\Delta\lambda}) \rceil, 0\}$ 
246   14:   $\eta_t = \text{update\_lr}(\eta, p_t, b_{t+1})$ 
247   15:   $\lambda_{t+1} = \lambda_t + \eta_t b_{t+1}$ 
248   16:  end for
249   17:  return  $\lambda_{opt}$ 
250 end procedure$ 
```

250 *heres to our fairness constraint. FedCF solves, $\lambda_{opt} = \min_{\lambda \in \Lambda} \lambda$, subject to $cg_t - c \leq 0$. We*
 251 *solve this using a framework analogous to Gradient Descent (GD) with Momentum (Polyak, 1964).*
 252 *Let η and μ be the initial learning rate and momentum constant, respectively. The update rule for λ_t*
 253 *is,*

$$\lambda_{t+1} = \lambda_t + \eta_t \cdot b_{t+1} = \lambda_t + \eta \cdot (1/2)^{p_t} \cdot b_{t+1} \in [\lambda_0, \lambda_{opt}],$$

254 *where $b_{t+1} = \mu \cdot b_t + (cg_t - c)$ is the modified step size, $p_t = \max\{\lceil \log_2(\frac{\eta}{\Delta\lambda}) \rceil, 0\}$ determines*
 255 *the scaling factor to update η by and ensure $\lambda_{t+1} \in [\lambda_0, \lambda_{opt}]$ where $\Delta\lambda = (\lambda_{opt} - \lambda_t) \mathbf{1}_{[b_{t+1} \geq 0]} +$*
 256 *$(\lambda_t - \lambda_0) \mathbf{1}_{[b_{t+1} < 0]}$.*

257 We do not stop immediately once a satisfactory λ is found; instead, we continue exploring to check
 258 whether a smaller λ exists. This algorithm directly applies to the federated setting, with one impor-
 259 tant consideration: the computation of the coverage gap in Line 4.

262 3.3 FEDCF: THE END-TO-END FEDERATED CONFORMAL FAIRNESS FRAMEWORK

264 Having established the sufficient terms to compute the fairness-specific coverage gap, we now
 265 present the FedCF framework. We discuss FedCF in the context of the interval-bounds esti-
 266 mates from Theorem 3.1, noting the discussion also applies to the point-estimate case by setting
 267

One drawback of the original Conformal Fairness algorithm is that it creates a sampled, discretized search space and iterates to find the minimal λ satisfying the fairness specification. This process requires computing the coverage gap every iteration and for each positive label, which becomes even more inefficient in the federated setting, where coverage gap computation requires client-server communication.

Algorithm 1 describes the core descent-based CF algorithm.¹ The algorithm takes as input the calibration set for each client, $\{\mathcal{D}_{\text{calib}}^{(k)}\}_{k \in \mathcal{K}}$, the set of (positive) labels, the set of sensitive groups, \mathcal{G} , a closeness criterion, c , and a filtering function, F_M . Additionally, we initialize a threshold, λ_0 , as the \hat{q} value given by the FCP algorithm (Lu et al., 2023), to ensure $1 - \alpha$ coverage is still satisfied.

For a given threshold λ_t , we can compute the coverage gap cg_t and evaluate whether it ad-

¹We omit the iteration over the positive labels for brevity and present just the core optimization.

270 $L_{\text{cov}} = U_{\text{cov}} = \Pi_{\text{cov}}$. The fairness-specific coverage gap is given by,
 271
 272 $\text{cg}(\lambda, F_M, \tilde{y}, \mathcal{G}) := \max_{g_a \in \mathcal{G}} \{U_{\text{cov}}(\lambda, F_M, g_a, \tilde{y})\} - \min_{g_b \in \mathcal{G}} \{L_{\text{cov}}(\lambda, F_M, g_b, \tilde{y})\}$ (6)
 273
 274 $= \max_{g_a, g_b \in \mathcal{G}} \{U_{\text{cov}}(\lambda, F_M, g_a, \tilde{y}) - L_{\text{cov}}(\lambda, F_M, g_b, \tilde{y})\}.$ (7)

275 While equations 6 and 7 are mathematically equivalent, their formulations lead to two different communication
 276 and aggregation strategies demonstrating the tradeoff between **communication overhead** and **privacy**. We present the *communication efficient* protocol in the main paper and the
 277 *enhanced privacy* protocol in Appendix D. In Appendix D, we also present a hybrid protocol, where
 278 clients select whether to use the *communication efficient* or *enhanced privacy* protocol. We include
 279 FedCF extensions concerning differential privacy in Appendix F.
 280

281 Note that U_{cov} and L_{cov} depend on $L^{(g, \tilde{y})}$ and $U^{(g, \tilde{y})}$, respectively. Since these quantities are also
 282 computed in a federated manner on the server, we compute them *prior* to computing the coverage
 283 gap for any particular λ^2 . Given that these priors are available on the server, we can compute the
 284 fairness-specific coverage gap. From Theorem 3.1, each client computes and sends two values for
 285 each $(g, \tilde{y}) \in \mathcal{G} \times \mathcal{Y}^+$ pair: $\frac{\alpha_k^{(g, \tilde{y}); \lambda} \cdot n_k^{(g, \tilde{y})}}{(n_k^{(g, \tilde{y})} + 1) \cdot (n_k + 1)}$ and $\frac{\alpha_k^{(g, \tilde{y}); \lambda} + 1}{n_k + 1}$. Once the server receives these pairs
 286 from each client, it proceeds to aggregate these quantities to derive L_{cov} and U_{cov} for each (g, \tilde{y}) .
 287 U_{cov} is limited 1 to reconcile $\Pr[\cdot] \leq 1$ for any event. The final coverage gap is determined as per
 288 Equation 6. Algorithms 2 and 3 describe the federated coverage gap algorithm.
 289

290 **Communication Complexity and Privacy Implications.** Each client is responsible for sending
 291 messages of size totaling $\mathcal{O}(2 \cdot |\mathcal{G}| |\mathcal{Y}^+|)$ to the server per server round. While this is linear in terms
 292 of the number of (g, \tilde{y}) pairs, we note that with enough λ s, the server can learn the distribution of
 293 $\Pr[s(\mathbf{x}_{\text{test}}, \tilde{y}) \leq \lambda \mid F_M(\mathbf{x}_{\text{test}}, y_{\text{test}}, g, \tilde{y}) = 1, E_k]$.
 294

295 **Algorithm 2** Server-side Aggregation for
 296 Coverage Gap

297 1: **procedure** SERVERCG($\lambda, F_M, \tilde{y}, \mathcal{G}, \mathcal{K}$)
 2: $n_list = [0]_{\mathcal{K}}$
 3: $l_list = [0]_{\mathcal{K} \times \mathcal{G}}, u_list = [0]_{\mathcal{K} \times \mathcal{G}}$
 4: **for** client $k \in \mathcal{K}$ **in parallel do**
 5: $(l_list[k], u_list[k], n_list[k])$
 6: = CLIENTCG($k, \lambda, F_M, \tilde{y}, \mathcal{G}$)
 7: **end for**
 8: $N = \sum_{k \in \mathcal{K}} n_list[k], K = |\mathcal{K}|$
 9: $U_{\text{cov}} = [0]_{\mathcal{G}}, L_{\text{cov}} = [0]_{\mathcal{G}}$
 10: **for** client $k \in \mathcal{K}$ **do**
 11: $\gamma_k = ((n_list[k] + 1) / (N + K))$
 12: $U_{\text{cov}} += (\gamma_k / L^{(g, \tilde{y})}) \cdot u_list[k]$
 13: $L_{\text{cov}} += (\gamma_k / U^{(g, \tilde{y})}) \cdot l_list[k]$
 14: **end for**
 15: $U_{\text{cov}} = \text{element_wise_min}(U_{\text{cov}}, [1]_{\mathcal{G}})$
 16: cov_gap = $\max_{g \in \mathcal{G}} U_{\text{cov}}[g] - \min_{g \in \mathcal{G}} L_{\text{cov}}[g]$
 17: **return** cov_gap
 18: **end procedure**

315 4 EXPERIMENTS

316 4.1 SETUP.

317 **Datasets.** We evaluate the FedCF framework on four multi-class datasets in different domains: (1,
 318 2) ACSIncome and ACSEducation (Ding et al., 2021), (3) Pokec- $\{\mathbf{n}, \mathbf{z}\}$ (Takac & Zabovsky, 2012),
 319 (4) Fitzpatrick Groh et al. (2021). These datasets were not originally for FL, so we partitioned them
 320 to form our clients. For the ACS datasets, we use state and territory data to partition the information
 321

322 ²The algorithm for computing the prior is similar to that of the coverage gap, so we omit it here

324 into clients. We consider **six** different partitioning schemes, based on common regional definitions
 325 in the U.S., which result in 4 (small), 8 (large), and 51 (all) clients. We also consider equivalent
 326 schemes for just the continental U.S.. For Pokec- $\{n, z\}$, each graph is treated as a separate client,
 327 as they originate from distinct partitions of the larger *Pokec* social network. Finally, for Fitzpatrick,
 328 since there is no predetermined partitioning scheme, we use a Dirichlet partitioner (Yurochkin et al.,
 329 2019) with concentration parameter of 0.5 to split $\mathcal{D}_{\text{train}}/\mathcal{D}_{\text{valid}}/\mathcal{D}_{\text{calib}}$ for $K \in \{2, 4, 8\}$ clients.
 330 We use a 30%/20%/25%/25% stratified split for the full $\mathcal{D}_{\text{train}}/\mathcal{D}_{\text{valid}}/\mathcal{D}_{\text{calib}}/\mathcal{D}_{\text{test}}$. We elaborate
 331 more on the datasets and experimental setup in Appendix C.

332 **Base Models.** For the ACS datasets, we use XGBoost (Chen & Guestrin, 2016). For Pokec- $\{n, z\}$,
 333 we use GraphSAGE (Hamilton et al., 2017) with GCN aggregation. For Fitzpatrick, we use ResNet-
 334 18 (He et al., 2016). Each of these models is trained using FedAvg (McMahan et al., 2017).
 335

336 **Baseline.** We construct a federated fairness-agnostic conformal predictor targeting a coverage
 337 level of $1 - \alpha = 0.9$ using FCP Lu et al. (2023) with T-Digest Dunning (2021). For the non-
 338 conformity score, we adopt APS (Romano et al., 2020b) and RAPS (Angelopoulos et al., 2022) for
 339 all datasets, as well as DAPS (H. Zargarbashi et al., 2023), a graph-specific method, for Pokec- $\{n, z\}$.
 340 We then assess fairness using $\lambda = \hat{q}(\alpha)$ for three popular group-fairness metrics, reformulated in
 341 Table 2—Demographic Parity, Equal Opportunity, and Predictive Equality.

342
 343 Table 2: Formulations for Conformal Fairness Metrics.

Metric	Definition
Demographic (or Statistical) Parity	$ \Pr[\tilde{y} \in \mathcal{C}_\lambda(X) \mid X \in g_a] - \Pr[\tilde{y} \in \mathcal{C}_\lambda(X) \mid X \in g_b] < c, \forall g_a, g_b \in \mathcal{G}, \forall \tilde{y} \in \mathcal{Y}^+$
Equal Opportunity	$ \Pr[\tilde{y} \in \mathcal{C}_\lambda(X) \mid Y = \tilde{y}, X \in g_a] - \Pr[\tilde{y} \in \mathcal{C}_\lambda(X) \mid Y = \tilde{y}, X \in g_b] < c, \forall g_a, g_b \in \mathcal{G}, \forall \tilde{y} \in \mathcal{Y}^+$
Predictive Equality	$ \Pr[\tilde{y} \in \mathcal{C}_\lambda(X) \mid Y \neq \tilde{y}, X \in g_a] - \Pr[\tilde{y} \in \mathcal{C}_\lambda(X) \mid Y \neq \tilde{y}, X \in g_b] < c, \forall g_a, g_b \in \mathcal{G}, \forall \tilde{y} \in \mathcal{Y}^+$

350 **Evaluation Metrics:** We report two key metrics: (1) *efficiency*, and (2) *worst-case fairness disparity*.
 351 The latter captures the largest difference in conditional coverage across groups, under the
 352 chosen fairness metric. For example, under *Demographic Parity*, we report:

$$\max_{\tilde{y} \in \mathcal{Y}^+} \max_{g_a, g_b \in \mathcal{G}} |\Pr[\tilde{y} \in \mathcal{C}_\lambda(\mathbf{x}_{\text{test}}) \mid \mathbf{x}_{\text{test}} \in g_a] - \Pr[\tilde{y} \in \mathcal{C}_\lambda(\mathbf{x}_{\text{test}}) \mid \mathbf{x}_{\text{test}} \in g_b]|. \quad (8)$$

355 More details on the experimental setup can be found in Appendix C.

356 4.2 RESULTS

359 In each figure, we use a **solid** line to represent the *average* efficiency of the **base federated conformal predictors** across different thresholds and a **dashed** line to represent the corresponding *average* worst-case fairness disparity. The bar plot shows the efficiency and worst-case fairness disparity using FedCF, while the **dots** indicate the *desired* fairness disparity. We report the average base
 360 performance for clarity and readability. In all experiments, FedCF achieves an actual fairness disparity
 361 within the specified closeness criterion, c , which may not be the case with the base federated
 362 conformal predictor.

366 **Preserves Key Characteristics of CF.** Two important characteristics of the CF framework are
 367 that it is (1) agnostic to the specific non-conformity score function and (2) supports intersectional
 368 fairness. We demonstrate that our FedCF framework preserves these two characteristics via the
 369 Pokec- $\{n, z\}$ dataset. Pokec- $\{n, z\}$ each have two sensitive attributes: *region* and *gender*. In addition
 370 to considering each attribute individually, we can treat each *pair* of attributes as distinct and apply
 371 FedCF. Furthermore, Pokec- $\{n, z\}$ is a graph dataset. Recently, several developments have been
 372 made in graph CP research on non-conformity scores that utilize the graph structure. In addition
 373 to two standard CP methods—APS and RAPS—we also provide results using DAPS. Figure 1 shows
 374 how the FedCF framework can achieve the desired fairness criterion with minimal cost to efficiency
 375 for different non-conformity scores and when considering multiple groups.

376 **Robust Performance with Different Numbers of Clients.** An important trade-off in trustworthy
 377 FL is between predictive utility and maintaining fairness/privacy guarantees for each of its clients,

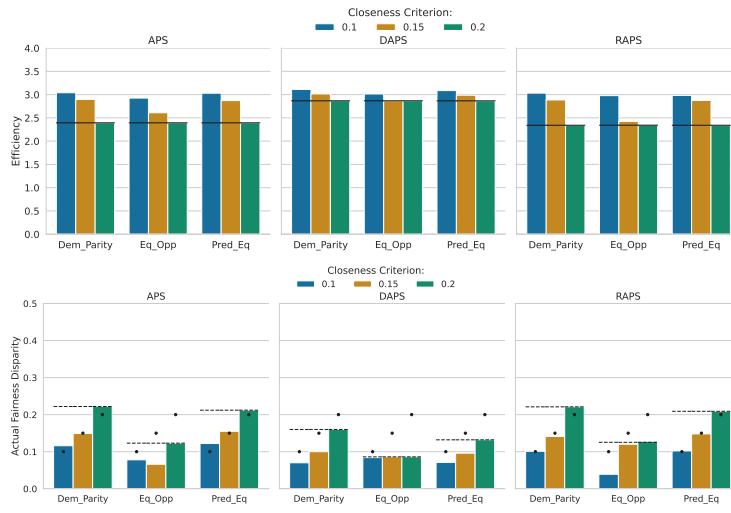


Figure 1: **Pokec- $\{n, z\}$** using **both** sensitive attributes. The top plots present the efficiency results, while the bottom plots are for the fairness disparities for (a) APS, (b) DAPS, and (c) RAPS. In all cases, FedCF achieves the desired closeness criteria better than the base federated conformal predictors.

which becomes increasingly challenging as the number of clients increases (Wen et al., 2023). To demonstrate how our framework can adapt to a varying number of clients, we use a Dirichlet partitioner with the Fitzpatrick dataset to evaluate performance with $K \in \{2, 4, 8\}$ clients in addition to a centralized setup with a single client. We see in Table 3 that as the number of clients increases, the baseline performance worsens, but the FedCF framework can still control for the necessary closeness criterion. We omit Equal Opportunity for Fitzpatrick as it is not meaningful in the context of this dataset, which aims to predict a skin condition, and the sensitive attribute is the skin type. People with certain skin types are known to be more likely to develop certain skin conditions, so the true positive rates of a classifier will typically not equalize, resulting in a degenerate (meaningless) solution. Finally, Fitzpatrick is a relatively small dataset with fewer than 17K points. Splitting the dataset into different splits and clients will result in a small number of points per (g, \tilde{y}) pair, making the interval bounds from Theorem 3.1 quite large.

Table 3: **Fitzpatrick using RAPS**. Each entry is of the form, **efficiency/fairness disparity**. We bold the lower fairness disparity value for each comparison. This table contains the results for $K \in \{1, 2, 4, 8\}$ clients. We observe that FedCF consistently matches or outperforms the base federated conformal predictor and is below the desired closeness criterion, c .

(a) $c = 0.1$

Metric	1 client		2 clients		4 clients		8 clients	
	Base	Ours	Base	Ours	Base	Ours	Base	Ours
Dem_Parity	2.356 / 0.151	3.647 / 0.103	2.565 / 0.163	4.149 / 0.153	2.844 / 0.293	4.684 / 0.099	3.502 / 0.166	5.214 / 0.089
Pred_Eq	2.356 / 0.111	3.647 / 0.109	2.543 / 0.177	4.140 / 0.177	2.837 / 0.287	4.564 / 0.102	3.502 / 0.171	5.293 / 0.083

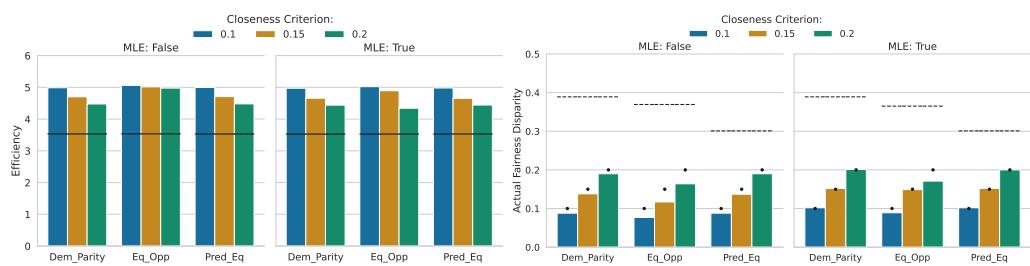
(b) $c = 0.15$

Metric	1 client		2 clients		4 clients		8 clients	
	Base	Ours	Base	Ours	Base	Ours	Base	Ours
Dem_Parity	2.356 / 0.151	2.356 / 0.151	2.565 / 0.163	2.793 / 0.163	2.837 / 0.291	3.347 / 0.114	3.498 / 0.166	3.859 / 0.088
Pred_Eq	2.356 / 0.111	2.356 / 0.111	2.543 / 0.177	2.778 / 0.177	2.844 / 0.289	3.296 / 0.144	3.496 / 0.170	3.867 / 0.087

(c) $c = 0.2$

Metric	1 client		2 clients		4 clients		8 clients	
	Base	Ours	Base	Ours	Base	Ours	Base	Ours
Dem_Parity	2.356 / 0.151	2.356 / 0.151	2.541 / 0.161	2.541 / 0.161	2.844 / 0.293	3.260 / 0.164	3.500 / 0.166	3.528 / 0.146
Pred_Eq	2.356 / 0.111	2.356 / 0.111	2.541 / 0.177	2.541 / 0.177	2.837 / 0.287	3.256 / 0.167	3.501 / 0.171	3.524 / 0.150

432 **Efficiency vs Fairness Trade-Off.** To make FedCF an actionable framework, it is essential to un-
 433 derstand the utility trade-off when imposing fairness constraints. Using the interval-based approach
 434 to estimate the fairness-specific coverage gap gives a finite-sample guarantee for controlling fairness
 435 gaps. However, sometimes imposing fairness may result in a severe cost to utility. For example, a de-
 436 generate conformal predictor (one with near-full efficiency) is “fair,” but completely impractical for
 437 use. By relaxing the theoretical guarantees, we can improve the efficiency by considering a tighter
 438 estimate for the coverage gap, using point estimates through MLE. Figure 2 compares the efficiency
 439 and fairness disparities when using the interval bounds vs the point estimate on the ACSEducation
 440 dataset. We observe that with the interval bounds, we always get within the closeness criterion, but
 441 the efficiencies are quite high. Alternatively, using point estimates may exceed the desired closeness
 442 criterion, but be more fair than the baseline and not sacrifice as much efficiency.



452 **Figure 2: ACSEducation using RAPS.** The left two plots are the efficiency plots for (a) using the
 453 interval bounds and (b) using the MLE estimate. Similarly, the two right (c) using the interval bounds
 454 and (d) using the MLE estimate. We observe that with the MLE estimates, FedCF achieves lower
 455 efficiency at the cost of a higher worst-case fairness disparity. Both the interval bounds and MLE
 456 estimates outperform the base federated conformal predictor in controlling for fairness disparity.

458 5 DISCUSSION

461 **On Data Heterogeneity.** In a practical federated setting, the data distribution will vary between
 462 clients—resulting in data heterogeneity across the FL system, which can affect performance at infer-
 463 ence time (Wen et al., 2023). To address these concerns, we evaluate FedCF on varying partitioning
 464 schemes. For Fitzpatrick, we use a probabilistic partitioning scheme to ensure the data is distributed
 465 in a particular manner. For the ACS and Pokec- $\{n, z\}$ datasets, the partitioning is naturally induced
 466 by state and region information in the datasets.

467 **On Data Requirements.** A major limitation of CP is that to achieve a desired coverage rate, practi-
 468 cally, you require a large enough calibration dataset such that the interval width for the CP guarantee
 469 is tight enough. This is exacerbated in the CF and FedCF framework as it requires sufficient calibra-
 470 tion data for each group-positive label pair (for each client). If we consider intersectional fairness,
 471 the multiplicative increase in the number of groups further increases the data requirements.

473 **On Interval Bounds.** We provide two ways of estimating the fairness-specific coverage level with
 474 intervals and point estimates, but we can choose different intervals by considering the following.
 475 Suppose the event $s(\mathbf{x}_{test}, y_{test}) \leq \lambda$ (conditioned on F_M) is a Bernoulli random trial of some
 476 unknown probability p . We want to estimate p as that is our fairness-specific coverage level. If we
 477 also treat the calibration scores as Bernoulli trials (by exchangeability), we can estimate p using a
 478 Binomial Proportion Confidence Interval. Several results provide tighter or looser bounds (Wallis,
 479 2013), which can be used to get better efficiency vs. fairness trade-offs.

480 6 CONCLUSION

483 In this work, we extended the Conformal Fairness framework to a federated setting, introducing
 484 the novel and comprehensive FedCF framework. We reformulated the CF framework to use a
 485 descent-based approach to make it more efficient for FL applications. Additionally, we developed
 486 theoretically grounded protocols to enable coverage gap calculations in a federated manner. The

486 FedCF framework offers clients a choice of participation protocols, including *communication efficient*
 487 and *enhanced privacy* options. We conducted experiments on various non-conformity scores
 488 and datasets– including graph data where we leverage the exchangeability assumption from CP.
 489

490 **Extensibility and Future Work** An important application of FedCF is that it can be used to audit
 491 federated conformal predictors for fairness (discussed in Appendix G). In the future, we will explore
 492 how FedCF can be extended to split learning (Gupta & Raskar, 2018). Unlike federated Learning,
 493 which trains full models locally and aggregates updates, split learning divides the model across
 494 clients and server, sharing only partial computations (enhanced privacy, reduced compute).
 495

496 REFERENCES

498 Nimesh Agrawal, Anuj Kumar Sirohi, Sandeep Kumar, et al. No prejudice! fair federated graph
 499 neural networks for personalized recommendation. In *Proceedings of the AAAI Conference on*
 500 *Artificial Intelligence*, volume 38, pp. 10775–10783, 2024.

501 Anastasios Angelopoulos, Stephen Bates, Jitendra Malik, and Michael I. Jordan. Uncertainty sets
 502 for image classifiers using conformal prediction, 2022. URL <https://arxiv.org/abs/2009.14193>.

504 Daniel J Beutel, Taner Topal, Akhil Mathur, Xinchi Qiu, Javier Fernandez-Marques, Yan Gao,
 505 Lorenzo Sani, Hei Li Kwing, Titouan Parcollet, Pedro PB de Gusmão, and Nicholas D Lane.
 506 Flower: A friendly federated learning research framework. *arXiv preprint arXiv:2007.14390*,
 507 2020.

509 Tianqi Chen and Carlos Guestrin. Xgboost: A scalable tree boosting system. In *Proceedings of the*
 510 *22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD*
 511 *'16*, pp. 785–794, New York, NY, USA, 2016. Association for Computing Machinery. ISBN
 512 9781450342322. doi: 10.1145/2939672.2939785. URL <https://doi.org/10.1145/2939672.2939785>.

514 John Cherian and Lenny Bronner. How the washington post estimates outstanding votes for the
 515 2020 presidential election. *Retrieved September*, 13:2023, 2020.

517 Jesse C Cresswell, Bhargava Kumar, Yi Sui, and Mouloud Belbahri. Conformal prediction sets can
 518 cause disparate impact. In *The Thirteenth International Conference on Learning Representations*,
 519 2025.

520 Frances Ding, Moritz Hardt, John Miller, and Ludwig Schmidt. Retiring adult: New datasets for fair
 521 machine learning. *Advances in neural information processing systems*, 34:6478–6490, 2021.

523 Jinshuo Dong, Aaron Roth, and Weijie J Su. Gaussian differential privacy. *Journal of the Royal*
 524 *Statistical Society Series B: Statistical Methodology*, 84(1):3–37, 2022.

525 Ted Dunning. The t-digest: Efficient estimates of distributions. *Software Impacts*, 7:100049, 2021.
 526 ISSN 2665-9638. doi: <https://doi.org/10.1016/j.simpa.2020.100049>. URL <https://www.sciencedirect.com/science/article/pii/S2665963820300403>.

528 Cynthia Dwork. Differential privacy. In *International colloquium on automata, languages, and*
 529 *programming*, pp. 1–12. Springer, 2006.

531 Cynthia Dwork, Aaron Roth, et al. The algorithmic foundations of differential privacy. *Foundations*
 532 *and trends® in theoretical computer science*, 9(3–4):211–407, 2014.

533 Úlfar Erlingsson, Vitaly Feldman, Ilya Mironov, Ananth Raghunathan, Kunal Talwar, and
 534 Abhradeep Thakurta. Amplification by shuffling: From local to central differential privacy
 535 via anonymity. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete*
 536 *Algorithms*, pp. 2468–2479. SIAM, 2019.

538 European Parliament and Council of the European Union. Regulation (EU) 2024/1689 of the European
 539 Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act) and amending Regulations (EC) No 300/2008, (EU) No

540 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Direc-
 541 tives 2014/90/EU, (EU) 2016/797 and (EU) 2019/882. Official Journal of the European Union,
 542 2024. URL <https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng>.

543 Arya Fayyazi, Mehdi Kamal, and Massoud Pedram. Facter: Fairness-aware conformal threshold-
 544 ing and prompt engineering for enabling fair llm-based recommender systems. In Forty-second
 545 International Conference on Machine Learning, 2025.

546 Matthew Groh, Caleb Harris, Luis Soenksen, Felix Lau, Rachel Han, Aerin Kim, Arash Koochek,
 547 and Omar Badri. Evaluating deep neural networks trained on clinical images in dermatology with
 548 the Fitzpatrick 17k dataset. In Proceedings of the IEEE/CVF Conference on Computer Vision and
 549 Pattern Recognition, pp. 1820–1828, 2021.

550 Otkrist Gupta and Ramesh Raskar. Distributed learning of deep neural network over multiple agents.
 551 Journal of Network and Computer Applications, 116:1–8, 2018. ISSN 1084-8045. doi: <https://doi.org/10.1016/j.jnca.2018.05.003>. URL <https://www.sciencedirect.com/science/article/pii/S1084804518301590>.

552 Soroush H. Zargarbashi, Simone Antonelli, and Aleksandar Bojchevski. Conformal prediction
 553 sets for graph neural networks. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara
 554 Engelhardt, Sivan Sabato, and Jonathan Scarlett (eds.), Proceedings of the 40th International
 555 Conference on Machine Learning, volume 202 of Proceedings of Machine Learning Research,
 556 pp. 12292–12318. PMLR, 23–29 Jul 2023. URL <https://proceedings.mlr.press/v202/h-zargarbashi23a.html>.

557 William L. Hamilton, Rex Ying, and Jure Leskovec. Inductive representation learning on large
 558 graphs. CoRR, abs/1706.02216, 2017. URL <http://arxiv.org/abs/1706.02216>.

559 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recogni-
 560 tion. In 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp.
 561 770–778, 2016. doi: 10.1109/CVPR.2016.90.

562 Dennis Hirsch, Timothy Bartley, Aravind Chandrasekaran, Davon Norris, Srinivasan Parthasarathy,
 563 and Piers Norris Turner. Business Data Ethics: Emerging Models for Governing AI and Advanced
 564 Analytics. Springer Nature, 2023.

565 Charles Jones, Fabio De Sousa Ribeiro, Mélanie Roschewitz, Daniel C Castro, and Ben Glocker.
 566 Rethinking fair representation learning for performance-sensitive tasks. In The Thirteenth
 567 International Conference on Learning Representations, 2025.

568 Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child,
 569 Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language
 570 models. CoRR, abs/2001.08361, 2020. URL <https://arxiv.org/abs/2001.08361>.

571 CR Komala, Ashok Kumar, N Hema, S Nagarani, Ajay Singh Yadav, M Rajendiran, R Srinivasan,
 572 and V Vijayan. Fair-automl: Enhancing fairness in machine learning predictions through auto-
 573 mated machine learning and bias mitigation techniques. In AIP Conference Proceedings, volume
 574 3193, pp. 020005. AIP Publishing LLC, 2024.

575 Charles Lu, Andréanne Lemay, Ken Chang, Katharina Höbel, and Jayashree Kalpathy-Cramer. Fair
 576 conformal predictors for applications in medical imaging. Proceedings of the AAAI Conference
 577 on Artificial Intelligence, 36(11):12008–12016, Jun. 2022. doi: 10.1609/aaai.v36i11.21459.
 578 URL <https://ojs.aaai.org/index.php/AAAI/article/view/21459>.

579 Charles Lu, Yaodong Yu, Sai Praneeth Karimireddy, Michael Jordan, and Ramesh Raskar. Federated
 580 conformal predictors for distributed uncertainty quantification. In International Conference on
 581 Machine Learning, pp. 22942–22964. PMLR, 2023.

582 Pranav Maneriker, Codi Burley, and Srinivasan Parthasarathy. Online fairness auditing through
 583 iterative refinement. In Proceedings of the 29th ACM SIGKDD Conference on Knowledge
 584 Discovery and Data Mining, KDD ’23, pp. 1665–1676, New York, NY, USA, 2023. Association
 585 for Computing Machinery. ISBN 9798400701030. doi: 10.1145/3580305.3599454. URL
 586 <https://doi.org/10.1145/3580305.3599454>.

594 Pranav Maneriker, Aditya T Vadlamani, Anutam Srinivasan, Yuntian He, Ali Payani, et al. Confor-
 595 mal prediction: A theoretical note and benchmarking transductive node classification in graphs.
 596 *Transactions on Machine Learning Research*, May 2025.

597

598 Charles Masson, Jee E. Rim, and Homin K. Lee. Ddsketch: a fast and fully-mergeable quantile
 599 sketch with relative-error guarantees. *Proc. VLDB Endow.*, 12(12):2195–2205, August 2019.
 600 ISSN 2150-8097. doi: 10.14778/3352063.3352135. URL <https://doi.org/10.14778/3352063.3352135>.

601

602 Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas.
 603 Communication-Efficient Learning of Deep Networks from Decentralized Data. In Aarti Singh
 604 and Jerry Zhu (eds.), *Proceedings of the 20th International Conference on Artificial Intelligence
 605 and Statistics*, volume 54 of *Proceedings of Machine Learning Research*, pp. 1273–1282. PMLR,
 606 20–22 Apr 2017. URL <https://proceedings.mlr.press/v54/mcmahan17a.html>.

607

608 Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. A survey
 609 on bias and fairness in machine learning. *ACM computing surveys (CSUR)*, 54(6):1–35, 2021.

610

611 New York City Council. Local law 4 of 2021: Fair chance act, 7 2021. URL <https://www.nyc.gov/site/cchr/law/fair-chance-act.page>.

612

613 New York City Council. Local law 144 of 2021: Prohibiting automated employment decision tools.
 614 *New York City Register*, 7 2023. URL <https://www.nyc.gov/site/dca/about/automated-employment-decision-tools.page>.

615

616 B.T. Polyak. Some methods of speeding up the convergence of iteration methods. *USSR
 617 Computational Mathematics and Mathematical Physics*, 4(5):1–17, 1964. ISSN 0041-5553. doi:
 618 [https://doi.org/10.1016/0041-5553\(64\)90137-5](https://doi.org/10.1016/0041-5553(64)90137-5). URL <https://www.sciencedirect.com/science/article/pii/0041555364901375>.

617

618 Yaniv Romano, Rina Foygel Barber, Chiara Sabatti, and Emmanuel Candès. With malice toward
 619 none: Assessing uncertainty via equalized coverage. *Harvard Data Science Review*, 2020a.

620

621 Yaniv Romano, Matteo Sesia, and Emmanuel Candes. Classification with valid and adaptive cover-
 622 age. *Advances in neural information processing systems*, 33:3581–3591, 2020b.

623

624 Lubos Takac and Michal Zabovsky. Data analysis in public social networks. In *International
 625 scientific conference and international workshop present day trends of innovations*, volume 1,
 626 2012.

627

628 U.S. Equal Employment Opportunity Commission. Questions and answers to clarify and provide
 629 a common interpretation of the uniform guidelines on employee selection procedures. *Federal
 630 Register*, 44(43), 1979. URL <https://www.eeoc.gov/laws/guidance/questions-and-answers-clarify-and-provide-common-interpretation-uniform-guidelines>.

631

632 Aditya T. Vadlamani, Anutam Srinivasan, Pranav Maneriker, Ali Payani, and Srinivasan
 633 Parthasarathy. A generic framework for conformal fairness. In *The Thirteenth International
 634 Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=xiQNFY133p>.

635

636 Vladimir Vovk, Alexander Gammerman, and Glenn Shafer. *Algorithmic learning in a random world*,
 637 volume 29. Springer, 2005.

638

639 Sean Wallis. Binomial confidence intervals and contingency tests: Mathematical fundamentals and
 640 the evaluation of alternative methods. *Journal of Quantitative Linguistics*, 20(3):178–208, 2013.
 641 doi: 10.1080/09296174.2013.799918. URL <https://doi.org/10.1080/09296174.2013.799918>.

642

643 Jie Wen, Zhixia Zhang, Yang Lan, Zhihua Cui, Jianghui Cai, and Wensheng Zhang. A survey on
 644 federated learning: challenges and applications. *International journal of machine learning and
 645 cybernetics*, 14(2):513–535, 2023.

648 Mikhail Yurochkin, Mayank Agarwal, Soumya Ghosh, Kristjan Greenewald, Trong Nghia Hoang,
649 and Yasaman Khazaeni. Bayesian nonparametric federated learning of neural networks, 2019.
650 URL <https://arxiv.org/abs/1905.12022>.

651
652 Yanfei Zhou and Matteo Sesia. Conformal classification with equalized coverage for adaptively
653 selected groups. *Advances in Neural Information Processing Systems*, 37:108760–108823, 2024.

654
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701

702 **A NOTATION TABLE**
703704 Table 4: Common notation used in FedCF.
705

706 Notation	707 Defintion
\mathcal{G}	The set of all demographic groups.
$\mathcal{Y}, \mathcal{Y}^+$	The set of labels and positive/advantaged labels, respectively.
g and \tilde{y}	The group $g \in \mathcal{G}$ and $\tilde{y} \in \mathcal{Y}^+$ under consideration.
F_M	Filter function for fairness metric M .
c	Closeness criterion for a fairness specification.
λ	Threshold used for constructing test prediction sets.
\mathcal{K}	The set of clients, $\{1, \dots, K\}$.
$\mathcal{D}_{\text{train}}^{(k)}/\mathcal{D}_{\text{valid}}^{(k)}/\mathcal{D}_{\text{calib}}^{(k)}$	Client k 's train/validation/calibration dataset.
n_k and N	$ \mathcal{D}_{\text{calib}}^{(k)} $ and $\sum_{k=1}^K n_k$, respectively.
$\mathcal{S}_k^{(g, \tilde{y})}$ and $n_k^{(g, \tilde{y})}$	$\left\{(\mathbf{x}_i, y_i) \in \mathcal{D}_{\text{calib}}^{(k)} \mid F_M(\mathbf{x}_i, y_i, g, \tilde{y}) = 1\right\}$ and $ \mathcal{S}_k^{(g, \tilde{y})} $
$\alpha_k^{(g, \tilde{y}); \lambda}$	$\sum_{(\mathbf{x}_i, -) \in \mathcal{S}_k^{(g, \tilde{y})}} \mathbf{1}[s(\mathbf{x}_i, \tilde{y}) \leq \lambda]$
E_k and γ_k	The event \mathbf{x}_{test} is exchangeable with data from client k and $\Pr[E_k]$.
$L^{(g, \tilde{y})}, U^{(g, \tilde{y})}$	Bounds for prior (term \textcircled{I}).
$\pi^{(g, \tilde{y})}$	Point estimate for prior (term \textcircled{I}).
$L_{\text{cov}}, U_{\text{cov}}$	Bounds for fairness-specific coverage level.
Π_{cov}	Point estimate for fairness-specific coverage level.

729 **B PROOFS**
730731 **B.1 PROOF OF THEOREM 3.1**
732

733 Recall, since the data is distributed across clients in the federated setting, we reformulated the
734 fairness-specific coverage level as Equation 3. In doing so, the computation of the coverage level is
735 split between the clients and the server. We present bounds and point estimates for each of the terms
736 in Equation 3 across Lemmas B.1, B.2, and B.3, leading to a proof of Theorem 3.1.

737 **B.1.1 CLIENT-SIDE ESTIMATES**
738

739 Since each client operates independently with its own dataset, we can derive interval bounds for
740 terms \textcircled{I} and \textcircled{II} . For the point estimates approach, we use maximum likelihood estimators (MLEs)
741 for each term, providing the tightest estimates.

742 **Lemma B.1.** *For each client k , group g , positive label \tilde{y} , and threshold λ , we get the following
743 interval bounds:*

$$744 \frac{\alpha_k^{(g, \tilde{y}); \lambda}}{n_k^{(g, \tilde{y})} + 1} \leq \Pr[s(\mathbf{x}_{\text{test}}, \tilde{y}) \leq \lambda \mid F_M(\mathbf{x}_{\text{test}}, y_{\text{test}}, g, \tilde{y}) = 1, E_k] \leq \frac{\alpha_k^{(g, \tilde{y}); \lambda} + 1}{n_k^{(g, \tilde{y})} + 1} \quad (9)$$

745 If the data are IID, then we can use an MLE point estimate, given by the following:

$$746 \Pr[s(\mathbf{x}_{\text{test}}, \tilde{y}) \leq \lambda \mid F_M(\mathbf{x}_{\text{test}}, y_{\text{test}}, g, \tilde{y}) = 1, E_k] = \frac{\alpha_k^{(g, \tilde{y}); \lambda}}{n_k^{(g, \tilde{y})}} \quad (10)$$

747 The proof of Lemma B.1 is as follows:

748 *Proof.* We first observe that,

$$749 \begin{aligned} 750 \Pr[s(\mathbf{x}_{\text{test}}, \tilde{y}) \leq \lambda \mid F_M(\mathbf{x}_{\text{test}}, y_{\text{test}}, g, \tilde{y}) = 1, E_k] \\ 751 = \Pr_{\mathbf{x}_{\text{test}} \sim P_k} [s(\mathbf{x}_{\text{test}}, \tilde{y}) \leq \lambda \mid F_M(\mathbf{x}_{\text{test}}, y_{\text{test}}, g, \tilde{y}) = 1], \end{aligned}$$

since exchangeability with the elements in k is true iff \mathbf{x}_{test} is sampled from k 's local distribution, P_k . The interval bounds follow from the conditional coverage guarantees given in CF (Vadlamani et al., 2025).

For the point estimate, we can model the event that the predicted score $s(\mathbf{x}_{\text{test}}, \tilde{y})$ falls below λ as a Bernoulli random variable with success probability p . We can treat the $n_k^{(g, \tilde{y})}$ calibration points as individual Bernoulli trials, to then construct a maximum likelihood estimate (MLE) for p , which will be $\hat{p} = \frac{\alpha_k^{(g, \tilde{y}); \lambda}}{n_k^{(g, \tilde{y})}}$. \square

Lemma B.1 bounds the fair-conditional coverage for a particular group-label pair for the test covariate $(\mathbf{x}_{\text{test}}, y_{\text{test}})$. We next bound the coverage of the test covariate satisfying the Fairness Metric (F_M), conditioned on the test point being exchangeable with data from client k using Lemma B.2, and provide the proof below.

B.1.2 SERVER-SIDE ESTIMATES

Terms III and IV require a global view of the clients' data, so they are handled on the server.

For term III , we follow the setup by Lu et al. (2023), where given $n_k = |\mathcal{D}_{\text{calib}}^{(k)}|$, $\gamma_k := \Pr[E_k] \propto n_k + 1$ and $\sum_{k=1}^K \gamma_k = 1$. Finally, for term IV , we have that

$$\Pr[F_M(\mathbf{x}_{\text{test}}, y_{\text{test}}, g, \tilde{y}) = 1] = \sum_{k=1}^K \Pr[F_M(\mathbf{x}_{\text{test}}, y_{\text{test}}, g, \tilde{y}) = 1 \mid E_k] \cdot \Pr[E_k].$$

Using Lemma B.2, we can get an interval-bound and point-estimate as shown in the following lemma.

Lemma B.2. *For each client k , group g , and positive label \tilde{y} , we get the following interval bounds:*

$$\frac{n_k^{(g, \tilde{y})}}{n_k + 1} \leq \Pr[F_M(\mathbf{x}_{\text{test}}, y_{\text{test}}, g, \tilde{y}) = 1 \mid E_k] \leq \frac{n_k^{(g, \tilde{y})} + 1}{n_k + 1} \quad (11)$$

If the data are IID, then we can use an MLE point estimate, given by the following:

$$\Pr[F_M(\mathbf{x}_{\text{test}}, y_{\text{test}}, g, \tilde{y}) = 1 \mid E_k] = \frac{n_k^{(g, \tilde{y})}}{n_k} \quad (12)$$

Proof. To demonstrate the finite sample guarantee, we note that $F_M(\mathbf{x}, y, g, \tilde{y}) = 1$ for all $(\mathbf{x}, y) \in \mathcal{D}_{\text{calib}}^{(k)}$ are all exchangeable Bernoulli trials. Observe that conditioning on E_k implies $\mathcal{D}_{\text{calib}}_+^{(k)} := \mathcal{D}_{\text{calib}}^{(k)} \cup \{\mathbf{x}_{\text{test}}\}$ is an exchangeable sequence of length $n_k + 1$. Treating this as a finite 'bag' of covariates, we have

$$\forall_{(\mathbf{x}, y) \in \mathcal{D}_{\text{calib}}_+^{(k)}}, \Pr[F_M(\mathbf{x}, y, g, \tilde{y}) = 1 \mid E_k] = \frac{\sum_{(\mathbf{x}_i, y_i) \in \mathcal{D}_{\text{calib}}_+^{(k)}} F_M(\mathbf{x}_i, y_i, g, \tilde{y})}{n_k + 1}.$$

In other words, we have defined the probability of randomly selecting a covariate with $F_M(\mathbf{x}, y, g, \tilde{y}) = 1$. Since this applies to all points we know,

$$\Pr[F_M(\mathbf{x}_{\text{test}}, y_{\text{test}}, g, \tilde{y}) = 1 \mid E_k] = \frac{\sum_{(\mathbf{x}_i, y_i) \in \mathcal{D}_{\text{calib}}_+^{(k)}} F_M(\mathbf{x}_i, y_i, g, \tilde{y})}{n_k + 1}. \quad (13)$$

Since we implicitly condition on $F_M(\mathbf{x}, y, g, \tilde{y})$, $\forall (\mathbf{x}, y) \in k$ (effectively making them deterministic), we can calculate the following bounds 13,

$$\begin{aligned} \frac{\sum_{(\mathbf{x}_i, y_i) \in \mathcal{D}_{\text{calib}}^{(k)}} F_M(\mathbf{x}_i, y_i, g, \tilde{y})}{n_k + 1} &\leq \Pr[F_M(\mathbf{x}_{\text{test}}, y_{\text{test}}, g, \tilde{y}) = 1 \mid E_k] \\ &\leq \frac{\sum_{(\mathbf{x}_i, y_i) \in \mathcal{D}_{\text{calib}}^{(k)}} F_M(\mathbf{x}_i, y_i, g, \tilde{y}) + 1}{n_k + 1}, \end{aligned} \quad (14)$$

810 where the $+1$ term comes from the unknown value of $F_M(\mathbf{x}_{\text{test}}, y_{\text{test}}, g, \tilde{y})$. Substituting the sums
 811 with $n_k^{(g, \tilde{y})}$, proves the interval bounds.
 812

813 For the point estimate, the event that $F_M = 1$ can be modeled as a Bernoulli random variable with
 814 success probability p . We can use the full n_k calibration points as n_k Bernoulli trials to construct an
 815 MLE for p , which will be $\hat{p} = \frac{n_k^{(g, \tilde{y})}}{n_k}$. \square
 816

817 Lastly, we use Lemma B.3, to bound $\Pr[F_M(\mathbf{x}_{\text{test}}, y_{\text{test}}, g, \tilde{y}) = 1]$. The proof of Lemma B.3 lever-
 818 ages the result of Lemma B.2.

819 **Lemma B.3.** *For each client k , group g , and positive label \tilde{y} , we get the following interval bounds:*

$$821 L^{(g, \tilde{y})} = \sum_{k=1}^K \gamma_k \frac{n_k^{(g, \tilde{y})}}{n_k + 1} \leq \Pr[F_M(\mathbf{x}_{\text{test}}, y_{\text{test}}, g, \tilde{y}) = 1] \leq \sum_{k=1}^K \gamma_k \frac{n_k^{(g, \tilde{y})} + 1}{n_k + 1} = U^{(g, \tilde{y})}. \quad (15)$$

823 If the data are IID, then we can use an MLE point estimate, given by the following:

$$824 \Pr[F_M(\mathbf{x}_{\text{test}}, y_{\text{test}}, g, \tilde{y}) = 1] = \sum_{k=1}^K \gamma_k \frac{n_k^{(g, \tilde{y})}}{n_k} = \pi^{(g, \tilde{y})}. \quad (16)$$

827 *Proof.* To achieve this result, we first use the law of total probability to separate
 828 $\Pr[F_M(\mathbf{x}_{\text{test}}, y_{\text{test}}, g, \tilde{y})]$ into terms known by the server and the client:
 829

$$830 \Pr[F_M(\mathbf{x}_{\text{test}}, y_{\text{test}}, g, \tilde{y})] = \sum_{k=1}^K \Pr[F_M(\mathbf{x}_{\text{test}}, y_{\text{test}}, g, \tilde{y}) = 1 \mid E_k] \cdot \Pr[E_k] \quad (17)$$

833 Then, substituting the bounds for term ② (see Lemma B.2), and $\gamma_k = \Pr[E_k]$ from term ④ into
 834 Equation 17, we complete the proof. \square
 835

836 Having proved the Lemmas, we can move on to proving,

837 **Theorem 3.1.** *The fairness-specific coverage level (Equation 3) can be bounded as*

$$838 L_{\text{cov}}(\lambda, F_M, g, \tilde{y}) \leq \Pr[s(\mathbf{x}_{\text{test}}, \tilde{y}) \leq \lambda \mid F_M(\mathbf{x}_{\text{test}}, y_{\text{test}}, g, \tilde{y}) = 1] \leq U_{\text{cov}}(\lambda, F_M, g, \tilde{y}),$$

839 where

$$840 L_{\text{cov}}(\lambda, F_M, g, \tilde{y}) = \sum_{k=1}^K \frac{\gamma_k \alpha_k^{(g, \tilde{y}); \lambda} n_k^{(g, \tilde{y})}}{(n_k^{(g, \tilde{y})} + 1)(n_k + 1) U^{(g, \tilde{y})}} \text{ and } U_{\text{cov}}(\lambda, F_M, g, \tilde{y}) = \sum_{k=1}^K \frac{\gamma_k (\alpha_k^{(g, \tilde{y}); \lambda} + 1)}{(n_k + 1) L^{(g, \tilde{y})}}. \quad (4)$$

843 If the data is IID, using MLE estimates for each term, we get the following estimate for the fairness-
 844 specific coverage level
 845

$$846 \Pi_{\text{cov}} = \Pr[s(\mathbf{x}_{\text{test}}, \tilde{y}) \leq \lambda \mid F_M(\mathbf{x}_{\text{test}}, y_{\text{test}}, g, \tilde{y}) = 1] = \sum_{k=1}^K \frac{\gamma_k \alpha_k^{(g, \tilde{y}); \lambda}}{n_k \pi^{(g, \tilde{y})}}. \quad (5)$$

849 *Proof.* Substituting the bounds for terms ①, ②, and ④ which were established via Lemmas B.1,
 850 B.2, B.3 respectively and the defintion of ③ into Equation 3 completes the proof. \square
 851

852 On closer inspection, we observe that Terms ① and ② can be combined and bound together.

854 **Lemma B.4.** *Using the defintions of $\alpha_k^{(g, \tilde{y}); \lambda}$ and $n_k + 1$ we have,*

$$855 \frac{\alpha_k^{(g, \tilde{y}); \lambda}}{n_k + 1} \leq \Pr[s(\mathbf{x}_{\text{test}}, \tilde{y}) \leq \lambda \mid F_M(\mathbf{x}_{\text{test}}, y_{\text{test}}, g, \tilde{y}) = 1, \mathbf{x}_{\text{test}} \stackrel{\text{exc.}}{\sim} k] \\ 856 \cdot \Pr[F_M(\mathbf{x}_{\text{test}}, y_{\text{test}}, g, \tilde{y}) = 1 \mid \mathbf{x}_{\text{test}} \stackrel{\text{exc.}}{\sim} k] \leq \frac{\alpha_k^{(g, \tilde{y}); \lambda} + 1}{n_k + 1}. \quad (18)$$

861 *Proof.* First, observe that,

$$862 \Pr[s(\mathbf{x}_{\text{test}}, \tilde{y}) \leq \lambda \mid F_M(\mathbf{x}_{\text{test}}, y_{\text{test}}, g, \tilde{y}) = 1, E_k] \cdot \Pr[F_M(\mathbf{x}_{\text{test}}, y_{\text{test}}, g, \tilde{y}) = 1 \mid E_k] \\ 863 = \Pr[s(\mathbf{x}_{\text{test}}, \tilde{y}) \leq \lambda, F_M(\mathbf{x}_{\text{test}}, y_{\text{test}}, g, \tilde{y}) = 1 \mid E_k]$$

864 Then consider the following Bernoulli random variables (R.V) $\mathbf{1}[s(\mathbf{x}, y) \leq \lambda] \cdot F_M(\mathbf{x}, y, g, \tilde{y})$
 865 for all $(\mathbf{x}, y) \in k \cup \{(\mathbf{x}_{\text{test}}, y_{\text{test}})\}$ which form an exchangeable sequence (using the assumption
 866 $\mathbf{x}_{\text{test}} \stackrel{\text{exc.}}{\sim} k$). Additionally, observe $\alpha_k^{(g, \tilde{y}); \lambda} = \sum_{(\mathbf{x}_i, y_i) \in k} \mathbf{1}[s(\mathbf{x}, y) \leq \lambda] \cdot F_M(\mathbf{x}, y, g, \tilde{y})$ is an
 867 equivalent definition of $\alpha_k^{(g, \tilde{y}); \lambda}$. The rest of the proof follows from the proof of Lemma B.2 by
 868 using $\mathbf{1}[s(\mathbf{x}, y) \leq \lambda] \cdot F_M(\mathbf{x}, y, g, \tilde{y})$ as the Bernoulli R.V instead of $F_M(\mathbf{x}, y, g, \tilde{y})$ and $\alpha_k^{(g, \tilde{y}); \lambda}$ in
 869 place of $n_k^{(g, \tilde{y})}$. \square
 870
 871

872 Using the above result, we can tighten the lower-bound of Theorem 3.1.
 873

874 **Corollary B.1.** *Swapping terms ① and ② with the combined term in Lemma B.4, the lower bound
 875 in Theorem 3.1 can be simplified and tightened to*

$$876 \quad L_{\text{cov}}(\lambda, F_M, g, \tilde{y}) = \sum_{k=1}^K \frac{\gamma_k \alpha_k^{(g, \tilde{y}); \lambda}}{(n_k + 1) U^{(g, \tilde{y})}} \quad (19)$$

880 *Proof.* Instead of substituting terms ① and ② into Equation 3 as in the proof of Theorem 3.1, we
 881 can instead use Lemma B.4 to update the bounds. Observe that the upper bound remains the same
 882 as Theorem 3.1 while the lower bound becomes tighter. \square
 883

884 C ADDITIONAL EXPERIMENT DETAILS

885 C.1 DATASETS

888 We present a summary of common dataset statistics in Table 5 and go into more details on each
 889 dataset in the following sections.
 890

891 Table 5: Dataset Statistics. T refers to Tabular, G refers to Graph, and V refers to vision.
 892 *ACS datasets have six (6) groups if using the *continental* split schemes (see Section C.2).
 893 ^ Number of inputs after removing those with unknown group information

Name	Type	Size	# Labeled	# Groups	# Classes
ACSIncome	T	1,664,500	ALL	race(9)*	4
ACSEducation	T	1,664,500	ALL	race(9)*	6
Fitzpatrick	V	16,012^	ALL	skin type(6)	9
Name	Type	(V , E)	# Labeled	# Groups	# Classes
Pokec-{n, z}	G	(133, 138, 1, 458, 258)	17,594	region(2), gender(2)	4

903 C.2 FOLKTABLES DATASETS

905 In the fairness space, the American Community Services (ACS) datasets from the `Folktables`
 906 library are a widely used set of tabular data (Ding et al., 2021). The data is taken across the 51 U.S.
 907 states and territories. For our federated setup and each dataset below, we consider the following 6
 908 partitioning schemes:
 909

- 910 (1.) **All:** We consider each U.S. state and territory to be its own client
- 911 (2.) **Large:** We follow the U.S. Census Bureau’s division of the U.S. into the Northeast, the
 912 Midwest, the South, and the West
- 913 (3.) **Small:** We follow the Bureau of Economic Analysis’s division of the U.S. into New Eng-
 914 land, the Mideast, the Great Lakes, the Plains, the Southeast, the Southwest, the Rocky
 915 Mountain, and the Far West.
- 916 (4-6.) **Continental All, Continental Large, Continental Small:** The same as 1 to 3, but we only
 917 consider the *continental* U.S.–removing Alaska, Hawaii, and Puerto Rico.

918 All Folktale datasets have a race attribute. When we partition the data using all the states and
 919 territories, we use the full version of race, which has 9 groups. However, when partitioning just with
 920 *continental* U.S., we combine some demographic groups—primarily those from Alaska, Hawaii, and
 921 Puerto Rico—into the appropriate ‘Other’ categories, resulting in a total of 6 groups.
 922

923 **ACSIIncome:** We used the standard ACSIIncome dataset from Folktale; however, we divided the
 924 targets into four classes by evenly splitting the income into 4 brackets. The sensitive attribute in this
 925 case is race, resulting in either 9 or 6 groups.
 926

927 **ACSEducation:** This is a custom dataset. We used the ACSTravelTime data and selected Edu-
 928 cation Level as the target. The education level was divided into 6 groups: {did not complete high
 929 school, has a high school diploma, has a GED, started an undergrad program, completed an under-
 930 grad program, and completed graduate or professional school}. ACSEducation also uses race as a
 931 sensitive attribute.
 932

933 C.3 NON-TABULAR DATASETS

934 **Pokec-{n,z}:** The Pokec-{n, z} dataset (Takac & Zabovsky, 2012) is a social network graph
 935 dataset collected from Pokec, a popular social network in Slovakia. Since several rows in the
 936 dataset are missing features, two commonly used subgraphs are the Pokec-z and Pokec-n datasets.
 937 The graphs have four labels corresponding to the fieldwork and two sensitive attributes: gender (2
 938 groups) and region (2 groups). Our experiments consider each attribute individually as well as in-
 939 tersectional fairness by creating an attribute with 4 groups. For our federated setup, we use each
 940 subgraph as a single client, resulting in 2 clients.
 941

942 **Fitzpatrick:** The Fitzpatrick dataset (Groh et al., 2021) contains clinical images classified based
 943 on the depicted skin condition. There are several levels of granularity regarding the skin condition la-
 944 bel. We use a version with 9 skin conditions: {inflammatory, malignant epidermal, genodermatoses,
 945 benign dermal, benign epidermal, malignant melanoma, benign melanocyte, malignant cutaneous
 946 lymphoma, malignant dermal}. There are 6 demographic groups based on the Fitzpatrick skin type.
 947 For our federated setup, we use a Dirichlet partitioner to split the data into $K \in \{2, 4, 8\}$ clients.
 948

949 C.4 HYPERPARAMETERS AND IMPLEMENTATION

950 To promote reproducibility, the source code for FedCF is provided in the supplementary material,
 951 along with the configuration files containing the hyperparameters used.
 952

953 The project was written using the Flower AI Federated Learning framework (Beutel et al., 2020) for
 954 both base model training and the FedCF framework.
 955

956 C.5 NON-CONFORMITY SCORES

957 **Adaptive Prediction Sets (APS)** The most popular CP method for classification problems is
 958 APS (Romano et al., 2020b). The scoring function first sorts the softmax logits in descending order
 959 and accumulates the class probabilities until the correct class is included. For tighter prediction sets,
 960 randomization is introduced through a uniform random variable.
 961

962 Formally, let $\hat{\pi}$ be a trained classification model with softmaxed output. If $\hat{\pi}(\mathbf{x})_{(1)} \geq \hat{\pi}(\mathbf{x})_{(2)} \geq$
 963 $\dots \geq \hat{\pi}(\mathbf{x})_{(K-1)}$, $u \sim U(0, 1)$, and r_y is the rank of the correct label, then
 964

$$965 s(\mathbf{x}, y) = \left[\sum_{i=1}^{r_y} \hat{\pi}(\mathbf{x})_{(i)} \right] - u \hat{\pi}(\mathbf{x})_y.$$

966 APS has two major drawbacks that have led to it being surpassed by other methods in recent CP
 967 literature. First, APS tends to produce large (less efficient) prediction sets. Second, it does not
 968

972 account for structure in its formulation. To address these issues, alternatives like RAPS and DAPS
 973 have emerged³.
 974

975 **Regularized Adaptive Prediction Sets (RAPS)** Angelopoulos et al. (2022) introduces a reg-
 976 ularization approach for APS. Given the same setup and notation as APS, define $o(\mathbf{x}, y) =$
 977 $|\{c \in \mathcal{Y} : \hat{\pi}(\mathbf{x})_y \geq \hat{\pi}(\mathbf{x})_c\}|$. Then,
 978

$$979 \quad s(\mathbf{x}, y) = \left[\sum_{i=1}^{r_y} \hat{\pi}(\mathbf{x})_{(i)} \right] - u\hat{\pi}(\mathbf{x})_y + \nu \cdot \max\{o(\mathbf{x}, y) - k_{reg}, 0\},$$

982 where ν and $k_{reg} \geq 0$ are regularization hyperparameters.
 983

984 **Diffusion Adaptive Prediction Sets (DAPS)** Graphs are rich with neighborhood information,
 985 with nodes often exhibiting homophily. This suggests that the non-conformity scores of connected
 986 nodes are likely to be related. To leverage this insight, DAPS H. Zargarbashi et al. (2023) incorpo-
 987 rates a one-step diffusion update on the non-conformity scores. Formally, if $s(\mathbf{x}, y)$ is a point-wise
 988 score function (e.g., APS), then the diffusion step yields a new score function

$$989 \quad \hat{s}(\mathbf{x}, y) = (1 - \delta)s(\mathbf{x}, y) + \frac{\delta}{|\mathcal{N}_{\mathbf{x}}|} \sum_{\mathbf{u} \in \mathcal{N}_{\mathbf{x}}} s(\mathbf{u}, y),$$

992 where $\delta \in [0, 1]$ is a diffusion hyperparameter and $\mathcal{N}_{\mathbf{x}}$ is the 1-hop neighborhood of \mathbf{x} .
 993

994 D FEDCF WITH ENHANCED PRIVACY

997 Preserving data privacy is a fundamental pillar of FL mechanisms, as they typically interact with
 998 sensitive client data. In this vein, we formulate an *enhanced privacy* version of FedCF.
 999

1000 D.1 ENHANCED PRIVACY

1002 To better preserve privacy (compared to the *communication efficient* approach), we can offload more
 1003 of the computation to the client-side, making it harder for the server-side to reverse-engineer or infer
 1004 distributional information from the sent quantities. Expanding Equation 7, we get

$$1005 \quad U_{\text{cov}}(\lambda, F_m, g_a, \tilde{y}) - L_{\text{cov}}(\lambda, F_m, g_b, \tilde{y})$$

$$1006 \quad = \sum_{k=1}^K \gamma_k \underbrace{\left\{ \frac{(\alpha_k^{(g_a, \tilde{y}); \lambda} + 1)}{(n_k + 1)L^{(g_a, \tilde{y})}} - \frac{\alpha_k^{(g_b, \tilde{y}); \lambda} n_k^{(g_b, \tilde{y})}}{(n_k^{(g_b, \tilde{y})} + 1)(n_k + 1)U^{(g_b, \tilde{y})}} \right\}}_{\text{Returned by the Client}}. \quad (20)$$

1011 In this formulation, the client sends back the summand for each group, positive label pair, making
 1012 the space complexity of the client’s message $\mathcal{O}(|\mathcal{G}|^2 |\mathcal{Y}^+|)$ —quadratic with respect to the number of
 1013 groups and linear with respect to positive labels.
 1014

1015 The data privacy improves with this approach compared to the *communication efficient* version,
 1016 since the data sent to the server is the difference of client-level summary statistics, which obfuscates
 1017 individual distribution information from the server. However, unlike the *communication efficient*
 1018 approach, the upper-coverage term (U_{cov}) is not separable from the aggregated sum, thus preventing
 1019 us from enforcing $U_{\text{cov}} < 1$. In limited data settings, this results in more conservative coverage gap
 1020 estimates, which increases the prediction set size when using the *enhanced privacy* approach.
 1021

1022 We provide a side-by-side comparison of the *communication efficient* and *enhanced privacy* version
 1023 of computing the federated coverage gap in Figure 3 in Appendix E.

1024 ³RAPS and DAPS have hyperparameters typically tuned on separate held-out data, but we fix them *a priori*
 1025 to preserve data for calibration and evaluation as well as to be consistent with what prior federated conformal
 1026 prediction works have done.

1026

D.2 HYBRID

1027

1028

1029

1030

1031

1032

1033

1034

1035

1036

1037

1038

In real-world scenarios, clients often have varying privacy and communication requirements. For example, clients in resource-constrained areas may not have the network bandwidth to send the necessary packets to the centralized server. In our proposed *hybrid* approach, a client may elect to be *communication efficient*, without preventing the remaining clients from using the *enhanced privacy* protocol. We present the full server-side algorithm, which combines the *communication efficient*, *enhanced privacy*, and *hybrid* protocols for the federated coverage gap, in Algorithm 6 in Appendix E.

1039

1040

1041

1042

1043

1044

1045

1046

1047

1048

1049

1050

D.3 EMPIRICAL COMPARISON

1051

1052

1053

1054

1055

1056

We conduct two experiments using the Fitzpatrick dataset and 8 clients, as well as the larger ACSIncome dataset with the *continental_all* partition scheme—48 clients—to test the *communication efficient*, *enhanced privacy*, and *hybrid* protocols. For the *hybrid* protocol, we randomly assign half the clients to each protocol. From Table 6, we observe that all configurations control the fairness disparity within the closeness criterion; however, if all clients agree upon the *communication efficient* protocol, FedCF achieves a better efficiency with a slightly worse fairness disparity, albeit still within the closeness criterion. Though with more data, we observe that the efficiency gaps are smaller as seen in Table 7.

1057

1058

1059

1060

1061

1062

1063

Table 6: **Fitzpatrick, 8 clients, APS.** Each entry is of the form, **efficiency/fairness disparity**. We bold the lower fairness disparity value for each comparison. We observe that the *communication efficient* approach produces the most efficient prediction sets, while having a similar or higher fairness disparity. The *enhanced privacy* approach and *hybrid* approach have similar performance (w.r.t efficiency and fairness disparity), with minor differences stemming from the stochasticity of FedCF, as they default to the same coverage-gap aggregation protocol (see Algorithm 6). All methods improve upon the baseline fairness disparity and control for the closeness criterion.

1064

1065

(a) Enhanced Privacy

Metric	$c = 0.1$		$c = 0.15$		$c = 0.2$	
	Base	Ours	Base	Ours	Base	Ours
Dem_Parity	3.671 / 0.136	7.041 / 0.047	3.671 / 0.136	4.978 / 0.101	3.672 / 0.136	3.940 / 0.111
Pred_Eq	3.676 / 0.134	7.042 / 0.047	3.675 / 0.134	4.765 / 0.094	3.672 / 0.134	3.8803 / 0.106

(b) Hybrid (50-50)

Metric	$c = 0.1$		$c = 0.15$		$c = 0.2$	
	Base	Ours	Base	Ours	Base	Ours
Dem_Parity	3.674 / 0.136	6.871 / 0.066	3.670 / 0.136	4.967 / 0.103	3.670 / 0.136	3.939 / 0.111
Pred_Eq	3.671 / 0.134	7.041 / 0.047	3.673 / 0.134	5.123 / 0.094	3.671 / 0.134	3.919 / 0.107

(c) Communication Efficient

Metric	$c = 0.1$		$c = 0.15$		$c = 0.2$	
	Base	Ours	Base	Ours	Base	Ours
Dem_Parity	3.674 / 0.137	6.053 / 0.104	3.674 / 0.136	4.890 / 0.103	3.672 / 0.136	3.935 / 0.111
Pred_Eq	3.671 / 0.134	6.308 / 0.109	3.674 / 0.134	4.931 / 0.094	3.674 / 0.134	3.876 / 0.106

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1080 Table 7: **ACSincome, Continental All, RAPS.** Each entry is of the form, **efficiency/fairness disparity**
 1081 We observe that with sufficient data, each protocol performs at a similar efficiency, and they
 1082 all decrease the baseline fairness disparity and control it within the closeness criterion. Our fairness
 1083 disparity values are bolded.

1084

1085

(a) Enhanced Privacy

Metric	$c = 0.1$		$c = 0.15$		$c = 0.2$	
	Base	Ours	Base	Ours	Base	Ours
Dem. Parity	2.609 / 0.148	3.037 / 0.086	2.610 / 0.148	2.634 / 0.138	2.613 / 0.148	2.613 / 0.148
Pred. Eq	2.607 / 0.160	3.294 / 0.063	2.610 / 0.161	2.661 / 0.138	2.609 / 0.161	2.609 / 0.161

1091

1092

(b) Hybrid (50-50)

Metric	$c = 0.1$		$c = 0.15$		$c = 0.2$	
	Base	Ours	Base	Ours	Base	Ours
Dem. Parity	2.608 / 0.148	3.039 / 0.085	2.609 / 0.148	2.633 / 0.138	2.596 / 0.148	2.596 / 0.148
Pred. Eq	2.606 / 0.160	3.277 / 0.079	2.606 / 0.160	2.657 / 0.138	2.595 / 0.161	2.595 / 0.161

1093

1094

(c) Communication Efficient

Metric	$c = 0.1$		$c = 0.15$		$c = 0.2$	
	Base	Ours	Base	Ours	Base	Ours
Dem. Parity	2.608 / 0.148	3.037 / 0.086	2.611 / 0.148	2.634 / 0.138	2.601 / 0.149	2.601 / 0.149
Pred. Eq	2.610 / 0.161	3.300 / 0.071	2.607 / 0.160	2.658 / 0.138	2.609 / 0.161	2.609 / 0.161

1095

1096

1097 With these empirical results, note that under the hybrid setting, clients that optimize for communication efficiency still benefit from the fact that they can operate over a limited bandwidth network connection. The required bandwidth for a particular client undergoes a factor of $\approx \frac{|\mathcal{G}|}{2}$ reduction—i.e. $\mathcal{O}(|\mathcal{G}|^2 ||\mathcal{Y}^+|) \rightarrow \mathcal{O}(2 \cdot |\mathcal{G}| ||\mathcal{Y}^+|)$, when a client selects the *communication efficient* protocol while ensuring the remaining clients benefit from the *enhanced privacy* protocol. For Fitzpatrick, this results in the communication overhead (in bytes) being reduced by a factor of three. ($|\mathcal{G}|/2 = 3$, for Fitzpatrick). For the ACS datasets using the small, large, or all client assignments, this reduction corresponds to $|\mathcal{G}|/2 = 4.5$

1098

1099

1100

1101

1102

1103

1104

1105

1106

1107

1108

1109

1110

1111

1112

1113

1114

1115

1116

1117

1118

1119

1120

1121

1122

1123

1124

1125

1126

1127

1128

1129

1130

1131

1132

1133

1134 E ALGORITHMS
1135
1136
1137
11381139 **Algorithm 4** More Communication Efficient Client-
1140 Side Computation for Coverage Gap

```

1: procedure CLIENTCG_COMM_EFFICIENT( $k, \lambda, F_M, \tilde{y}, \mathcal{G}$ )
2:    $l_k = [0]_{\mathcal{G}}$ 
3:    $u_k = [0]_{\mathcal{G}}$ 
4:   for  $g \in \mathcal{G}$  do
5:     if use_mle then
6:        $l_k[g] \leftarrow \frac{\alpha_k^{(g, \tilde{y}); \lambda}}{n_k}$ 
7:        $u_k[g] \leftarrow \frac{\alpha_k^{(g, \tilde{y}); \lambda}}{n_k}$ 
8:     else
9:        $l_k[g] \leftarrow \frac{\alpha_k^{(g, \tilde{y}); \lambda} \cdot n_k^{(g, \tilde{y})}}{((n_k^{(g, \tilde{y})} + 1) \cdot (n_k + 1))}$ 
10:       $u_k[g] \leftarrow \frac{\alpha_k^{(g, \tilde{y}); \lambda} + 1}{(n_k + 1)}$ 
11:    end if
12:   end for
13:   return  $l_k, u_k, n_k$ 
14: end procedure

```

Algorithm 5 Client-Side Computation for Coverage Gap with Enhanced Privacy

```

1: procedure CLIENTCG_PRIVATE( $k, \lambda, F_M, \tilde{y}, \mathcal{G}$ )
2:    $l_k = [0]_{\mathcal{G}}$ 
3:    $u_k = [0]_{\mathcal{G}}$ 
4:   for  $g \in \mathcal{G}$  do
5:     if use_mle then
6:        $l_k[g] \leftarrow \frac{\alpha_k^{(g, \tilde{y}); \lambda}}{(n_k \cdot \pi(g, \tilde{y}))}$ 
7:        $u_k[g] \leftarrow \frac{\alpha_k^{(g, \tilde{y}); \lambda}}{(n_k \cdot \pi(g, \tilde{y}))}$ 
8:     else
9:        $l_k[g] \leftarrow \frac{\alpha_k^{(g, \tilde{y}); \lambda} \cdot n_k^{(g, \tilde{y})}}{((n_k^{(g, \tilde{y})} + 1) \cdot (n_k + 1) \cdot U(g, \tilde{y}))}$ 
10:       $u_k[g] \leftarrow \frac{\alpha_k^{(g, \tilde{y}); \lambda} + 1}{((n_k + 1) \cdot L(g, \tilde{y}))}$ 
11:    end if
12:   end for
13:    $pw\_cg_k = [0]_{\mathcal{G} \times \mathcal{G}}$ 
14:   // Pairwise coverage gap
15:   for  $(g_a, g_b) \in \mathcal{G} \times \mathcal{G}$  do
16:      $pw\_cg_k[g_a, g_b] \leftarrow u_k[g_a] - l_k[g_b]$ 
17:   end for
18:   return  $pw\_cg_k, n_k$ 
19: end procedure

```

Figure 3: **Pseudocode for the two client-side protocols to compute the coverage gap.** The *enhanced privacy* version (on the right) includes the pairwise computation step, which results in a larger space complexity compared to the more *communication efficient* version (on the left).

1163
1164
1165
1166
1167
1168
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1180
1181
1182
1183
1184
1185
1186
1187

1188

1189

1190

1191

1192

1193

1194

1195

1196

1197

Algorithm 6 Full Server-side Aggregation for Coverage Gap

```

1: procedure SERVERCG( $\lambda_0, F_M, \tilde{y}, \mathcal{G}, \text{formulations}$ )
2:    $n\_list = [0]_{\mathcal{K}}$ 
3:    $l\_list = [0]_{\mathcal{K} \times \mathcal{G}}, u\_list = [0]_{\mathcal{K} \times \mathcal{G}}$                                  $\triangleright$  Used for comm. efficient formulations
4:    $pw\_cg\_list = [0]_{\mathcal{K} \times \mathcal{G} \times \mathcal{G}}$                                  $\triangleright$  Used for private formulations
5:   for client  $k \in \mathcal{K}$  in parallel do
6:     if  $\text{formulations} == \text{COMM\_EFFICIENT}$  then
7:       Receive  $(l_k, u_k, n_k) = \text{CLIENTCG\_COMM\_EFFICIENT}(k, \lambda_0, F_M, \tilde{y}, \mathcal{G})$ 
8:        $l\_list[k] \leftarrow l_k, u\_list[k] \leftarrow u_k$ 
9:     else
10:      Receive  $(pw\_cg_k, n_k) = \text{CLIENTCG\_PRIVATE}(k, \lambda_0, F_M, \tilde{y}, \mathcal{G})$ 
11:       $pw\_cg\_list[k] \leftarrow pw\_cg_k$ 
12:    end if
13:     $n\_list[k] \leftarrow n_k$ 
14:  end for

15:  // Initialize final coverage variables
16:   $N = \sum_{k \in \mathcal{K}} n\_list[k], K = |\mathcal{K}|, U_{\text{cov}} = [0]_{\mathcal{G}}, L_{\text{cov}} = [0]_{\mathcal{G}}, PW_{\text{cov}} = [0]_{\mathcal{G} \times \mathcal{G}}$ 
17:  all_comm_efficient = all( $\text{formulations}[k] == \text{COMM\_EFFICIENT}$ )
18:  for client  $k \in \mathcal{K}$  do
19:     $\gamma_k = ((n\_list[k] + 1)/(N + K))$ 
20:    if all_comm_efficient then
21:       $U_{\text{cov}} += \left(\gamma_k / L^{(g, \tilde{y})}\right) \cdot u\_list[k]$                                  $\triangleright$  Standard operations are element-wise
22:       $L_{\text{cov}} += \left(\gamma_k / U^{(g, \tilde{y})}\right) \cdot l\_list[k]$ 
23:    else
24:      if  $\text{formulations}[k] == \text{COMM\_EFFICIENT}$  then
25:         $PW_{\text{cov}} += \gamma_k \cdot (u\_list[k] \ominus l\_list[k]^{\top})$   $\ominus$  is pairwise differences between two vectors.
26:      else
27:         $PW_{\text{cov}} += \gamma_k \cdot pw\_cg\_list[k]$ 
28:      end if
29:    end if
30:  end for

31:  if all_comm_efficient then
32:     $U_{\text{cov}} = \text{element\_wise\_min}(U_{\text{cov}}, [1]_{\mathcal{G}})$                                  $\triangleright$  Limit upper coverage prior to coverage gap calculation
33:    cov_gap =  $\max_{g \in \mathcal{G}} U_{\text{cov}}[g] - \min_{g \in \mathcal{G}} L_{\text{cov}}[g]$ 
34:  else
35:    cov_gap =  $\min \left\{ \max_{g_a, g_b \in \mathcal{G}} PW_{\text{cov}}[g_a, g_b], 1 \right\}$                                  $\triangleright$  Limit Coverage Gap to 1
36:  end if
37:  return cov_gap
38: end procedure

```

1234

1235

1236

1237

1238

1239

1240

1241

1242 **F DIFFERENTIAL PRIVACY IN FEDCF**
 1243
 1244
 1245

1246 FedCF can also be extended to formally consider (ϵ, δ) -differential privacy (DP), a mathematically
 1247 rigorous framework for data privacy (Dwork, 2006), where δ is the probability that ϵ -DP is violated.
 1248 We can embed DP within our framework via client shuffling and additive noise approaches. Client
 1249 shuffling is a global DP approach that is performed after the client sends data. Before the server
 1250 receives the data, it goes through a trusted, centralized shuffler to anonymize which client has sent
 1251 what data (Erlingsson et al., 2019). Our framework can accommodate client shuffling due to its
 1252 parallelism with client-side computation and its additive aggregation approach.

1253 For additive noise, we propose augmenting the values each client sends back with Gaussian
 1254 noise (Dwork et al., 2014; Dong et al., 2022), such that a client returns,

$$1255 \quad h = \frac{(\alpha_k^{(g_a, \tilde{y}); \lambda} + 1)}{(n_k + 1)L^{(g_a, \tilde{y})}} - \frac{\alpha_k^{(g_b, \tilde{y}); \lambda} n_k^{(g_b, \tilde{y})}}{(n_k^{(g_b, \tilde{y})} + 1)(n_k + 1)U^{(g_b, \tilde{y})}} + X, \quad (21)$$

1258 where X is a Gaussian random variable (R.V). For the *communication efficient* approach, one would
 1259 add a Gaussian R.V. to the upper coverage and lower coverage terms returned by the client. To ensure
 1260 (ϵ, δ) -DP, we make $X \sim \mathcal{N}\left(0, \frac{2 \ln(1.25/\delta)(\Delta g)^2}{\epsilon^2}\right)$, where Δh is the sensitivity of h —or how much h
 1261 can change if one of the points in the client’s dataset changes. For FedCF, h can be affected by data
 1262 changes in the covariates (or non-conformity scores), labels, and group memberships.

1264
 1265
 1266
 1267 **F.1 EXAMPLE: DIFFERENTIAL PRIVACY BOUNDS FOR ENHANCED PRIVACY PROTOCOL**
 1268
 1269

1270 Observe using the *enhanced privacy* approach, $\Delta h \leq \frac{1}{n_k} \left(\frac{1}{L^{(g_a, \tilde{y})}} + \frac{1}{U^{(g_b, \tilde{y})}} \right)$. For the *communication*
 1271 *efficient* approach $\Delta h \leq \frac{1}{n_k L^{(g_a, \tilde{y})}}$ for the upper coverage term and $\Delta h \leq \frac{1}{n_k U^{(g_b, \tilde{y})}}$ for the lower
 1272 coverage term. The server will know the sensitivity used by each client and their choice of ϵ and δ .

1273 To demonstrate how the server can estimate the coverage gap, we will consider an example using
 1274 the *enhanced privacy* approach. The result from server aggregation is,

$$1277 \quad \text{cov_gap_est}(\lambda, F_m, g_a, g_b, \tilde{y}) \\ 1278 \quad = \sum_{k=1}^K \gamma_k \underbrace{\left\{ \frac{(\alpha_k^{(g_a, \tilde{y}); \lambda} + 1)}{(n_k + 1)L^{(g_a, \tilde{y})}} - \frac{\alpha_k^{(g_b, \tilde{y}); \lambda} n_k^{(g_b, \tilde{y})}}{(n_k^{(g_b, \tilde{y})} + 1)(n_k + 1)U^{(g_b, \tilde{y})}} + X_k^{(g_a, g_b, \tilde{y})} \right\}}_{\text{Returned by the Client}}, \quad (22)$$

1281 where $X_k^{(g_a, g_b, \tilde{y})} \sim \mathcal{N}(0, \sigma_{k; (g_a, g_b, \tilde{y})}^2)$ such that $\sigma_{k; (g_a, g_b, \tilde{y})}^2$ provides (ϵ_k, δ_k) -DP for the client.
 1282 Then observe,

$$1285 \quad \text{cov_gap_est}(\lambda, F_m, g_a, g_b, \tilde{y}) \\ 1286 \quad = \sum_{k=1}^K \gamma_k \underbrace{\left\{ \frac{(\alpha_k^{(g_a, \tilde{y}); \lambda} + 1)}{(n_k + 1)L^{(g_a, \tilde{y})}} - \frac{\alpha_k^{(g_b, \tilde{y}); \lambda} n_k^{(g_b, \tilde{y})}}{(n_k^{(g_b, \tilde{y})} + 1)(n_k + 1)U^{(g_b, \tilde{y})}} \right\}}_{\text{true coverage gap}} + \sum_{k=1}^K \underbrace{\gamma_k X_k^{(g_a, g_b, \tilde{y})}}_{\text{Guassian R.V.}} \quad (23)$$

$$1291 \quad = \text{cov_gap}(\lambda, F_m, g_a, g_b, \tilde{y}) + X, \quad X \sim \mathcal{N}\left(0, \sum_{k=1}^K \gamma_k^2 \sigma_{k; (g_a, g_b, \tilde{y})}^2\right) \quad (24)$$

1294 Using a prespecified probability β we can accept or reject the statement
 1295 $\text{cov_gap_est}(\lambda, F_m, g_a, g_b, \tilde{y}) \leq c$. In other words, we can check whether,

1296

1297

1298

1299

1300

1301

1302

1303

1304

1305

1306

1307

1308

1309

1310

1311

1312

1313

1314

1315

1316

1317

1318

$$\text{cov_gap}(\lambda, F_m, g_a, g_b, \tilde{y}) + X \leq c \implies X \leq c - \text{cov_gap}(\lambda, F_m, g_a, g_b, \tilde{y})$$

$$\implies \frac{X}{\underbrace{\sqrt{\sum_{k=1}^K \gamma_k X_k^{(g_a, g_b, \tilde{y})}}}_{\text{Standard Normal RV}}} \leq \frac{c - \text{cov_gap}(\lambda, F_m, g_a, g_b, \tilde{y})}{\sqrt{\sum_{k=1}^K \gamma_k X_k^{(g_a, g_b, \tilde{y})}}}.$$

Then, if $\Phi\left(\frac{c - \text{cov_gap}(\lambda, F_m, g_a, g_b, \tilde{y})}{\sqrt{\sum_{k=1}^K \gamma_k X_k^{(g_a, g_b, \tilde{y})}}}\right) > \beta$, where Φ is the CDF of the standard normal distribution,

we can accept the coverage gap as being less than c . In other words, with probability β , the closeness criterion is satisfied with λ .

While using Gaussian noise results in a PAC-style guarantee, one could instead add strictly positive noise via an exponential mechanism Dwork et al. (2014), where the noise $X \sim \exp(\frac{\epsilon}{2\Delta h})$ is selected to satisfy ϵ -DP, i.e., $(\epsilon, 0)$ -DP. This would result in an overestimate of the actual coverage gap. If the overestimate satisfies the closeness criterion, then the server would assert that the exact coverage gap also satisfies the closeness criterion—thus restoring the strict (non-PAC) guarantee in FedCF.

G FEDCF FOR AUDITING

Auditing tools are vital for regulatory bodies to ensure ML models comply with fairness and safety standards (Maneriker et al., 2023). In this regard, we present how FedCF can be used to determine if a federated conformal predictor is *fair* according to the regulator’s specification of fairness and closeness criterion, c (U.S. Equal Employment Opportunity Commission, 1979; New York City Council, 2021; 2023; European Parliament and Council of the European Union, 2024).

To assess compliance, FedCF can use the global threshold (λ) values used by the previously trained conformal-predictor and provide it to each client. Then, the client should send the sufficient values calculated via Algorithm 4 (or Algorithm 5) to compute the federated coverage gap. The server would aggregate these values using Algorithm 6. If the calculated coverage gap is below c , then the server can assert that the conformal predictor is fair.

Our auditing approach does not require all clients to provide data for auditing. As discussed in Section 3.1, our guarantees hold assuming that the test-point, $(\mathbf{x}_{\text{test}}, y_{\text{test}}) \sim \sum_{k=1}^K \gamma_k P_k$, is sampled from a mixture of client distributions where γ_k is the probability the test point is sampled from P_k , or equivalently is exchangeable with data from client k . Thus, if a subset of clients used to train the original federated conformal predictor provides auditing data, then the audit guarantees will hold assuming that $(\mathbf{x}_{\text{test}}, y_{\text{test}})$ are sampled from a mixture consisting of the subset of clients used for auditing. This result allows clients to *independently* decide if they would like to submit data for auditing.

The auditing tool provided by FedCF can also be used to ascertain the *marginal* fairness with respect to each client. Using the auditing procedure described above with data from one client, FedCF can determine if the global, federated conformal predictor maintains fairness with respect to data from a single client. If the computed coverage gap is less than c , then the fairness guarantees hold with regard to $(\mathbf{x}_{\text{test}}, y_{\text{test}}) \sim P_k$, i.e., the client’s marginal distribution.

1342

1343

1344

1345

1346

1347

1348

1349

1350 H MORE RESULTS

1352 Here, we provide additional results for the ACS and Pokec- $\{n,z\}$ datasets. Recall, in each figure,
 1353 we use a **solid** line to represent the *average* efficiency of the **base federated conformal predictors**
 1354 across different thresholds and a **dashed** line to represent the corresponding *average* worst-case
 1355 fairness disparity. The bar plot shows the efficiency and worst-case fairness disparity using FedCF,
 1356 while the **dots** indicate the *desired* fairness disparity. We report the average base performance for
 1357 clarity and readability

1359 H.1 IMPACT OF DATA HETEROGENEITY ON ACSEDUCATION: US vs CONTINENTAL US

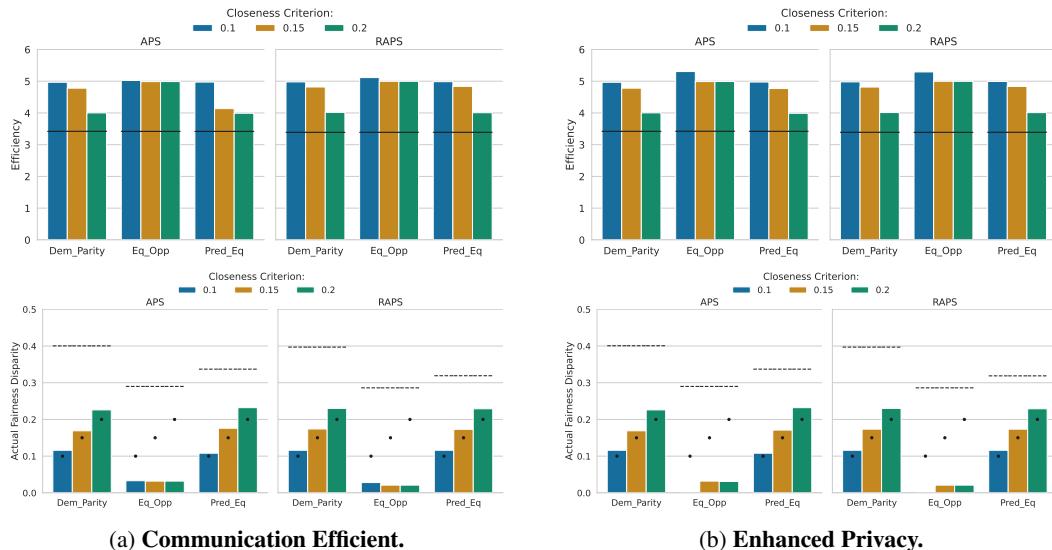


Figure 4: **ACSEDUCATION, Small, Interval Bounds.** The plots in the top row indicate the efficiency with the corresponding fairness disparity plots in the bottom row. We observe that when all US states are included (and Puerto Rico), the closeness criterion is satisfied. However, the efficiency for Equal Opportunity is high for all closeness criterion values, especially compared to the continental US version of ACSEDUCATION in Figure 5. This result stems from a conservative coverage gap estimate during calibration due to limited covariate representation for some groups.

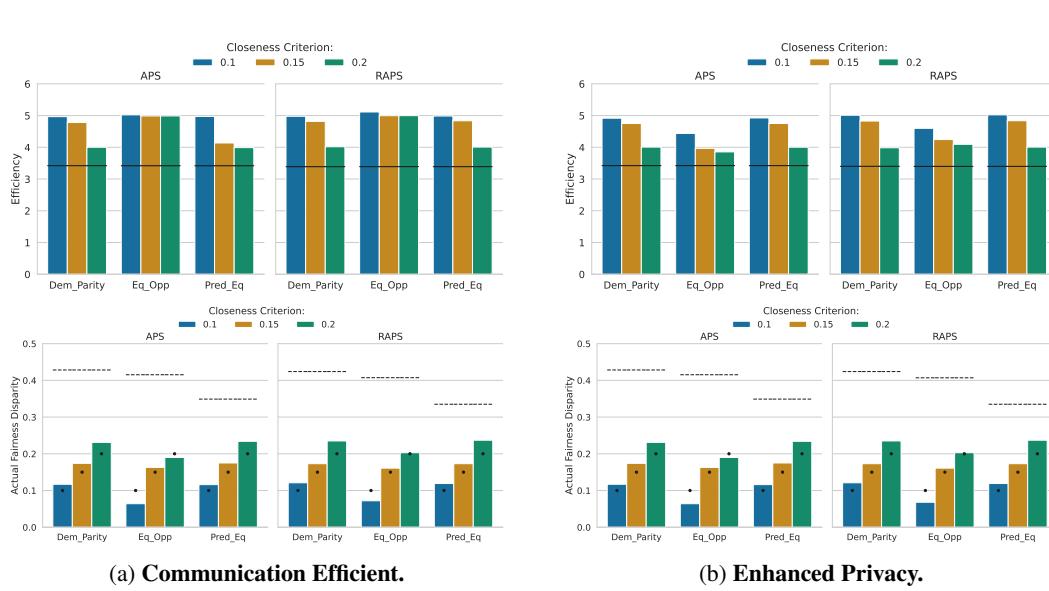


Figure 5: **ACSEducation, Continental Small, Interval Bounds.** The top row demonstrates the efficiency of FedCF when using the continental version of ACSEducation, and its fairness disparity on the bottom row. Compared to Figure 4, the efficiencies improved (particularly for Equal Opportunity using RAPS), due to increased covariate representation for all sensitive groups.

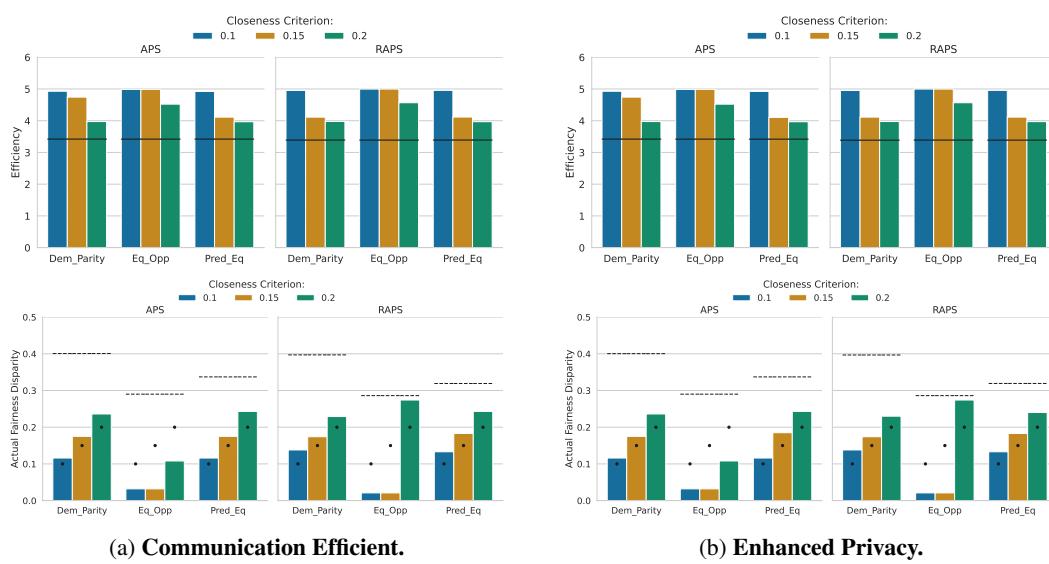
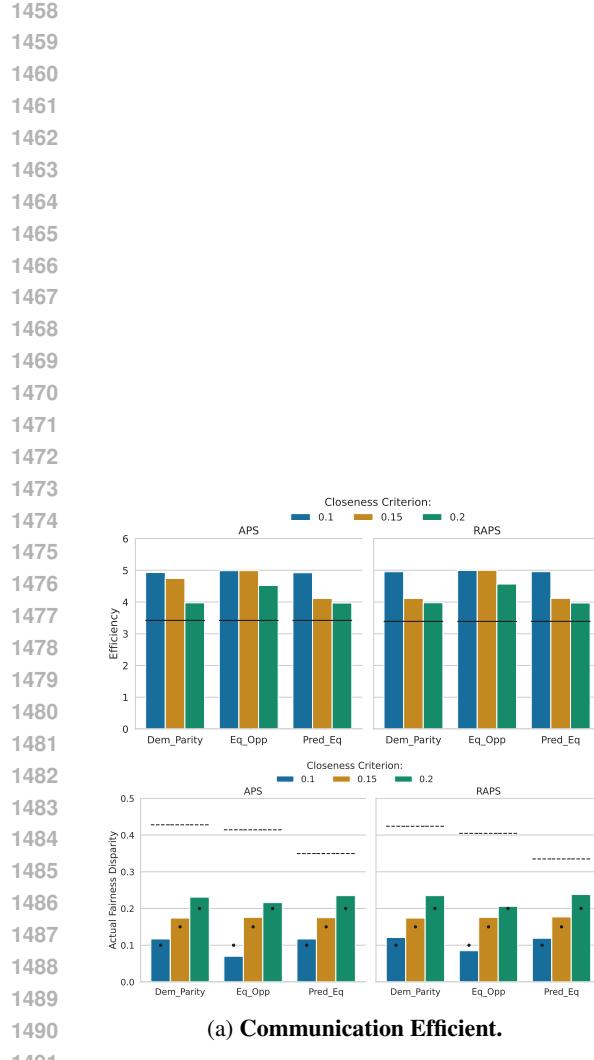
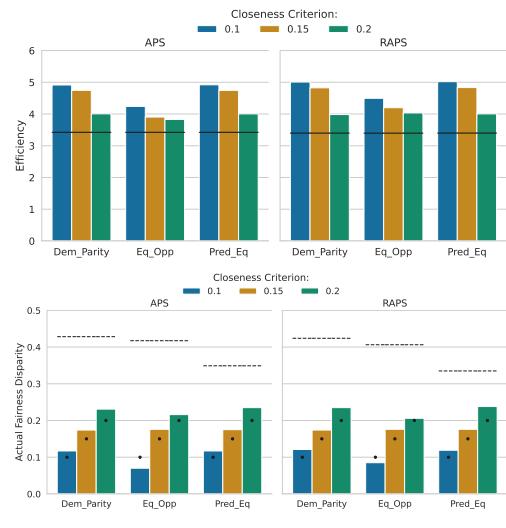


Figure 6: ACSEducation, Small, Point Estimates The plots in the top row indicate the efficiency with the corresponding fairness disparity plots in the bottom row. We observe that using point estimates will result in a similar or lower efficiency than using the interval bounds approach in Figure 4, at the cost of a similar or higher fairness violation. Because the MLE does not provide a finite sample guarantee, the violation can exceed the desired closeness criterion, but will be lower than the baseline federated conformal predictor.



(a) Communication Efficient.



(b) Enhanced Privacy.

Figure 7: **ACSEducation, Continental Small, Point Estimates.** The plots in the top row indicate the efficiency with the corresponding fairness disparity plots in the bottom row. We observe that using point estimates will result in a similar or lower efficiency than using the interval bounds approach in Figure 5, at the cost of a similar or higher fairness violation. Because the MLE does not provide a finite sample guarantee, the violation can exceed the desired closeness criterion, but will be lower than the baseline federated conformal predictor.

1512 H.2 IMPACT OF DIFFERENT SENSITIVE ATTRIBUTES FOR POKEC- $\{N, Z\}$
1513

1514

1515

1516

1517

1518

1519

1520

1521

1522

1523

1524

1525

1526

1527

1528

1529

1530

1531

1532

1533

1534

1535

1536

1537

1538

1539

1540

1541

1542

1543

1544

1545

1546

1547

1548

1549

1550

1551

1552

1553

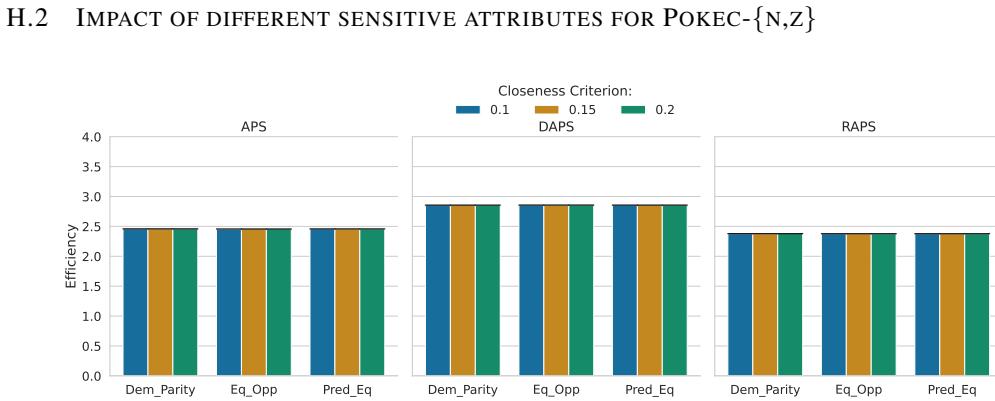
1554

1555

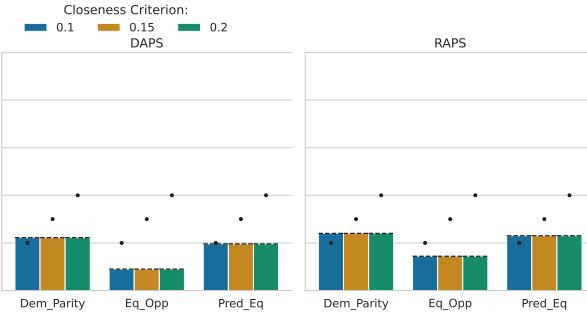
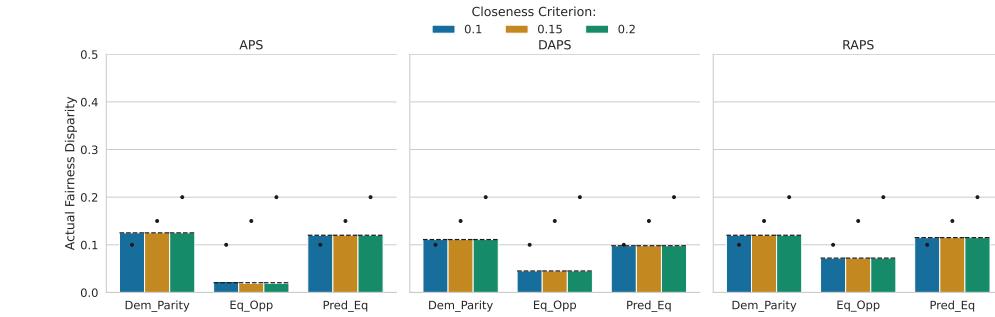
1556

1557

1558



(a) Communication Efficient.



(b) Enhanced Privacy.

Figure 8: **Pokec- $\{n,z\}$, gender.** For each plot (a) and (b), the top plots are for the efficiency, and the bottom plots are for the fairness disparity. The baseline disparity is within the closeness criterion, so we see no changes in efficiency when using FedCF. This is the case when using either the *communication efficient* and *enhanced privacy* protocols.

1563

1564

1565



Figure 9: **Pokec-{n,z}, region.** For each plot (a) and (b), the top plots are for the efficiency, and the bottom plots are for the fairness disparity. Note that while the baseline disparity is within the closeness criterion for the test set, the finite-sample guarantee from using the interval bounds ensures FedCF looks for a better threshold, resulting in a smaller violation with a small cost to efficiency. This is the case when using either the *communication efficient* and *enhanced privacy* protocols.



Figure 10: **Pokec- $\{n,z\}$, region and gender.** For each plot (a) and (b), the top plots are for the efficiency, and the bottom plots are for the fairness disparity. In the case of intersectional fairness, since there are more groups, the violation will be worse than considering a single sensitive attribute. We observe that in all cases, FedCF produces a threshold that satisfies the closeness criterion, at a slight cost to efficiency. This is the case when using either the *communication efficient* and *enhanced privacy* protocols.

1674 I ADDITIONAL RESULTS FROM REBUTTAL PHASE

1675 I.1 JUSTIFICATION FOR DESCENT-BASED APPROACH TO CF

1676 In this experiment, we define the discrete search space needed in (Vadlamani et al., 2025) as
 1677 $\Lambda = \text{linspace}(\hat{q}(\alpha), 2, \text{num_rounds})$, where $\hat{q}(\alpha)$ is the minimal lambda to satisfy $1 - \alpha$ cov-
 1678 erage of standard federated CP, and 2 represents the upper bound of the RAPS score (Angelopoulos
 1679 et al., 2022). We observe that more communication rounds are required for the iterative approach
 1680 to achieve performance comparable to our descent-based approach. This is because linspace deter-
 1681 mines the precision of the λ values, so by using fewer rounds, we will find a less precise, and in turn
 1682 more conservative, λ . This is not a limitation for our descent-based approach, which will continue
 1683 to converge to a constraint-satisfying λ . The results in Table 1 below demonstrate that it took 1000
 1684 rounds of the Iterative approach to achieve sufficient granularity and match the performance of the
 1685 Descent-Based approach, which only used 100 rounds. Thus, we demonstrate approximately a $10 \times$
 1686 speedup to achieve comparable performance, justifying the descent-based approach as a necessary
 1687 adaptation for the federated learning setting.

1688
 1689 **Table 8: ACSIncome (small) using RAPS.** Each entry is of the form, **efficiency/fairness disparity**.
 1690 We compare the Descent-Based (with 100 communication rounds) and the Iterative method (with
 1691 100 and 1000 communication rounds) across different closeness criteria (c).
 1692

1693 1694 1695 1696 1697 1698 1699 1700 1701 1702 1703 1704 1705 1706 1707 1708 1709 1710 1711 1712 1713 1714 1715 1716 1717 1718 1719 1720 1721 1722 1723 1724 1725 1726 1727	Method	$c = 0.1$	$c = 0.15$	$c = 0.2$
		Eff / Disp	Eff / Disp	Eff / Disp
	Descent-Based (rounds=100)	3.127 / 0.096	2.977 / 0.179	2.816 / 0.257
	Iterative (rounds=100)	3.239 / 0.119	3.137 / 0.139	2.890 / 0.213
	Iterative (rounds=1000)	3.120 / 0.137	2.989 / 0.170	2.821 / 0.251