Certified Characterization of Privacy, Participation, and Convergence in Over-the-Air Federated Learning

Anonymous Authors[1]

Abstract

Over-the-air federated learning (OTA-FL) exploits wireless superposition to aggregate client updates in one transmission, reducing communication overhead. In truncated channel-inversion OTA-FL, the receive scaling factor η jointly controls normalized channel noise, privacy sensitivity, and client activation thresholds. We characterize this coupling for smooth nonconvex objectives under Rayleigh fading via certified truncation envelopes. The resulting bounds yield a computable feasible set of receive amplitudes satisfying convergence, privacy, and participation constraints. We show that apparent participation fairness can be misleading at large η, where clients may be nearly uniformly excluded, and introduce a dropped-weight constraint to rule out this degenerate regime. We further propose CARES, an adaptive receive-scaling algorithm that selects η_t online via contextual bandits, tracks privacy with zero-concentrated differential privacy, and enforces participation-envelope constraints. Experiments on image classification show that CARES achieves competitive accuracy and strong worst-client accuracy while maintaining formal (ε, δ)-DP using only implicit channel noise among the OTA baselines considered.

1. Introduction

Federated learning (FL) enables clients to train a shared model without exchanging raw data (McMahan et al., 2017), making it attractive for wireless edge systems such as IoT, autonomous devices, and emerging 6G networks. However, as the client populations and model sizes grow, the repeated transmission of high-dimensional updates creates a severe communication bottleneck (Kairouz & McMahan, 2021).

[1]Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Preliminary work. Under review by the International Conference on Machine Learning (ICML). Do not distribute.

Over-the-air federated learning (OTA-FL) addresses this bottleneck by exploiting waveform superposition over the wireless multiple-access channel where clients transmit simultaneously and the server directly receives an analog aggregate (Amiri & Gündüz, 2020; Zhu et al., 2019; Yang et al., 2020). This reduces communication latency and bandwidth usage, but also ties learning performance to physical-layer effects. For instance, channel noise can degrade optimization, yet it can also provide implicit privacy by masking client updates (Liu & Simeone, 2020; Liang et al., 2025). Meanwhile, fading, power constraints, and truncated channel inversion determine which clients can/cannot participate (Amiri & Gündüz, 2020; Hamidi et al., 2025). Thus, accuracy, privacy, and participation are coupled through the wireless channel.

In OTA-FL truncated channel inversion, we evaluate how this coupling is governed by the receive scaling factor η. This parameter controls the effective aggregation noise, transcript sensitivity, and channel threshold for client activation. Thus, small η can amplify noise and slow learning, whereas large η can exclude many clients. Moreover, standard participation-fairness metrics can be misleading because clients may appear treated similarly while being nearly uniformly inactive. Therefore, our main question is *how can OTA-FL adapt receive scaling while preserving convergence, privacy, and non-degenerate participation?*

To this end, we propose a certified and adaptive receive-scaling framework for privacy-preserving OTA-FL. The key idea is to separate certification from adaptation. First, we derive conservative truncation envelopes under Rayleigh fading that define a computable feasible set of receive amplitudes that satisfy convergence, privacy, and participation constraints. We introduce a dropped-weight constraint to exclude regimes where clients are nearly uniformly inactive. Second, within this certified safe set, we adapt η_t online using contextual bandits, drawing on LinUCB (Li et al., 2010) and Bandits with Knapsacks (Badanidiyuru et al., 2018), with zero-concentrated differential privacy (zCDP) budget treated as the limited resource.

Prior work has studied aggregation distortion, device selection, power control, beamforming, and robustness (Yang et al., 2022; Cao et al., 2021; Liu et al., 2024). Privacy-

preserving OTA-FL has explored channel noise as privacy noise, fading-aware privacy, scheduling, alignment, and privacy-budget constraints (Liu & Simeone, 2020; Liang et al., 2025; Wei et al., 2026; Yan et al., 2023). AdaScale (Kalarde et al., 2025) adapts receive scaling to reduce privacy leakage under convergence constraints. Rather than replacing such methods, our framework can act as a certification layer where given a receive-scaling rule, scheduling policy, or alignment choice, the proposed envelopes check whether the induced participation pattern satisfies dropped-weight and asymmetry limits while also satisfying privacy and convergence constraints.

Our work also involves differentially private and fairness-aware FL. DP-SGD and privacy accounting provide standard tools for private learning (Abadi et al., 2016), while zCDP gives convenient composition for Gaussian mechanisms (Bun & Steinke, 2016). Digital-FL methods such as DP-SCAFFOLD and Robust-HDP address data heterogeneity and heterogeneous privacy noise (Noble et al., 2022; Malekmohammadi et al., 2024), but do not model OTA superposition or channel truncation. Fairness-aware FL methods such as agnostic FL and q-FFL reweight client objectives using separately observable updates (Mohri et al., 2019; Li et al., 2019), which are unavailable after OTA aggregation. Recent OTA fairness work considers multi-objective or minimax criteria (Hamidi et al., 2025; Öksüz et al., 2024). In contrast, we certify participation through analytic truncation envelopes tied to channel heterogeneity. Our contributions are:

- Characterization of how receive scaling couples channel noise, privacy sensitivity, and participation in truncated channel-inversion OTA-FL for smooth nonconvex objectives under Rayleigh fading.
- Derivations of certified truncation envelopes that yield a computable feasible set satisfying convergence, privacy, and useful participation constraints.
- Identification of a degeneracy in participation-fairness metrics at large η, together with a dropped-weight constraint that rules out nearly uniform exclusion.
- Proposing CARES, a certified adaptive receive-scaling algorithm that selects η_t online using contextual bandits, tracks privacy via zCDP, and enforces convergence and participation-envelope constraints.
- Evaluation showing that CARES improves over fixed OTA scaling by 2.7 percentage points on MNIST and 8.9 percentage points on FEMNIST, while maintaining formal DP using only implicit channel noise among the OTA baselines considered.

2. System Model and Certified η Characterization

We consider K clients and one parameter server (PS). Client k holds dataset $\mathcal{D}_k$ of size n_k, with $n = \sum_k n_k$ and aggregation weight $p_k = n_k/n$. The global objective is $\min_{\mathbf{w}\in\mathbb{R}^d} F(\mathbf{w}) \triangleq \sum_{k=1}^K p_k F_k(\mathbf{w})$, where $F_k(\mathbf{w}) = \mathbb{E}_{\xi\sim\mathcal{D}_k}[\ell(\mathbf{w};\xi)]$. This is the standard weighted objective used in FedAvg-style FL (McMahan et al., 2017; Kairouz & McMahan, 2021). At round t, client k computes an unbiased stochastic gradient $\mathbf{g}_k^t$ at $\mathbf{w}^t$. We use one local step per round ($E = 1$). Let $h_k^t > 0$ be the effective uplink channel magnitude after phase compensation, with

$$h_k^t \sim \text{Rayleigh}(\mu_k), \qquad \mathbb{E}[h_k^t] = \mu_k\sqrt{\pi/2}. \tag{1}$$

Long-term channel statistics $\{\mu_k\}$ are treated as public side information. The PS selects a receive scaling factor $\eta > 0$. Exact OTA realization of the FedAvg weight would use $b_k^{t,\star} = \eta p_k/h_k^t$ before normalization (Amiri & Gündüz, 2020; Zhu et al., 2019). This inversion is feasible only if the transmit-power constraint is satisfied:

$$h_k^t \geq \frac{\eta p_k \|\mathbf{g}_k^t\|}{\sqrt{P_{\max}}} \triangleq h_{\min}^t(k). \tag{2}$$

We therefore use strict truncation:

$$b_k^t = \begin{cases} \eta p_k/h_k^t, & h_k^t \geq h_{\min}^t(k), \\ 0, & h_k^t < h_{\min}^t(k). \end{cases} \tag{3}$$

Let $\mathcal{K}^t$ denote the active set, then PS receives

$$\mathbf{y}^t = \sum_{k\in\mathcal{K}^t} h_k^t b_k^t \mathbf{g}_k^t + \mathbf{z}^t, \qquad \mathbf{z}^t \sim \mathcal{N}(\mathbf{0}, \sigma_z^2 I_d), \tag{4}$$

forms $\tilde{\mathbf{g}}^t = \mathbf{y}^t/\eta$, and updates $\mathbf{w}^{t+1} = \mathbf{w}^t - \alpha\tilde{\mathbf{g}}^t$.

Defining truncation bias as:

$$\mathbf{e}^t = -\sum_{k\notin\mathcal{K}^t} p_k \mathbf{g}_k^t. \tag{5}$$

The OTA estimator has the exact decomposition of:

$$\tilde{\mathbf{g}}^t = \sum_{k=1}^K p_k \mathbf{g}_k^t + \mathbf{e}^t + \frac{1}{\eta}\mathbf{z}^t. \tag{6}$$

Thus; η simultaneously controls normalized channel noise, power-feasible participation, and, as shown later, privacy sensitivity. Detailed derivations and modeling assumptions are given in Appendix A. We assume L-smooth objectives, unbiased stochastic gradients with bounded variance σ_g^2, conditionally independent client noise, Rayleigh fading independent of gradient norms, and clipped gradients $\|\mathbf{g}_k^t\| \leq G$, with formal statements given in Appendix A.1. Also, our analysis combines standard smooth nonconvex SGD descent arguments, Gaussian/zCDP privacy accounting, and gradient clipping as used in DP-SGD (Ghadimi & Lan, 2013; Dwork & Roth, 2014; Bun & Steinke, 2016; Abadi et al., 2016).

2.1. Certified Truncation Envelopes

The true truncation probability of client k is $q_k^{\mathrm{true}}(\eta) = \Pr\Big(h_k^t < \frac{\eta p_k \|\mathbf{g}_k^t\|}{\sqrt{P_{\max}}}\Big)$. It depends on the random, training-dependent gradient norm and therefore is generally not available in closed form before training. Under Rayleigh fading, channel-gradient independence, and clipping $\|\mathbf{g}_k^t\| \le G$, we obtain the certified upper envelope:

$$\bar{q}_k(\eta) = 1 - \exp\left(-\frac{\eta^2 p_k^2 G^2}{2\mu_k^2 P_{\max}}\right) \ge q_k^{\mathrm{true}}(\eta). \tag{7}$$

Then, the dropped-weight envelope defined as:

$$\bar{\mathcal{E}}(\eta) = \sum_{k=1}^{K} p_k \bar{q}_k(\eta), \tag{8}$$

which satisfies $\mathbb{E}[S^t] \le \bar{\mathcal{E}}(\eta)$ for $S^t = \sum_{k\notin\mathcal{K}^t} p_k$. To measure envelope-level participation imbalance, let k_+ be $\arg\min_k \frac{p_k}{\mu_k}$, and k_-be $\arg\max_k \frac{p_k}{\mu_k}$, and define $a = \frac{p_{k_-}^2 G^2}{2\mu_{k_-}^2 P_{\max}}$, and $b = \frac{p_{k_+}^2 G^2}{2\mu_{k_+}^2 P_{\max}}$, with $a \ge b$. The envelope-based participation-asymmetry surrogate is then equals:

$$\bar{\mathcal{A}}(\eta) = \bar{q}_{k_-}(\eta) - \bar{q}_{k_+}(\eta) = e^{-b\eta^2} - e^{-a\eta^2} \ge 0. \tag{9}$$

The quantity $\bar{\mathcal{A}}$ compares the certified truncation risks of the effective hardest and easiest clients. Since it is a difference of upper envelopes, it is not claimed to upper-bound the realized participation gap pointwise. Moreover, $\bar{\mathcal{A}}(\eta)$ also vanishes in the degenerate regime where all clients are nearly always truncated. We therefore use $\bar{\mathcal{A}}$ together with the dropped-weight envelope $\bar{\mathcal{E}}$ and later impose $\bar{\mathcal{E}}(\eta) \le \tau$ to rule out uniform exclusion.

2.2. Certified η-Characterization

Theorem 2.1 (Certified OTA Receive-Scaling Characterization)**.** *Under Assumptions A.1–A.4 and the privacy scope in Assumption A.5, with fixed $\eta > 0$, $\alpha \le 1/(12L)$, and T rounds, the following certified statements hold. Sharper constants under a tighter step-size condition appear in Appendix C.1.*

Convergence.

$$\frac{1}{T}\sum_{t=0}^{T-1} \mathbb{E}\|\nabla F(\mathbf{w}^t)\|^2 \le \frac{C_0\left[F(\mathbf{w}^0) - F^\star\right]}{\alpha T} + C_1\alpha\sigma_g^2 + C_2\alpha\frac{\sigma_z^2 d}{\eta^2} + C_3 G^2 \bar{\mathcal{E}}(\eta). \tag{10}$$

The constants are specified in Appendix C.1. The channel-noise term decreases as $1/\eta^2$, while the truncation-bias envelope grows through $\bar{\mathcal{E}}(\eta)$.

Privacy. *For the aggregate OTA transcript, the mechanism satisfies record-level (ε_k, δ)-DP with respect to client k's local contribution, with*

$$\varepsilon_k(\eta) \le \frac{2\eta^2 p_k^2 G^2 T}{\sigma_z^2} + \sqrt{\frac{8\eta^2 p_k^2 G^2 T \ln(1/\delta)}{\sigma_z^2}}. \tag{11}$$

Thus the certified privacy loss is increasing in η. The transcript scope and conservative sensitivity accounting are given in Appendix C.4.

Participation envelope. *Let $\bar{\mathcal{A}}(\eta)$ be defined as in* (9)*. If $a = b$, then $\bar{\mathcal{A}}(\eta) \equiv 0$. In the nondegenerate case $a > b > 0$, $\bar{\mathcal{A}}(\eta)$ is increasing on $0 \le \eta \le \eta_{\mathrm{peak}}$, where*

$$\eta_{\mathrm{peak}} = \sqrt{\frac{\ln(a/b)}{a-b}}.$$

On this moderate-truncation branch, increasing η reduces normalized channel noise, increases privacy loss, and worsens the certified participation-asymmetry envelope.

Corollary 2.2 (Adaptive receive scaling)**.** *If $\{\eta_t\}_{t=0}^{T-1}$ is any predictable sequence chosen using only observations up to round $t-1$, with $\eta_t \in [\eta_{\min}, \eta_{\max}]$ and $\eta_{\min} > 0$, then, defining*

$$\bar{G}_T = \frac{1}{T}\sum_{t=0}^{T-1} \mathbb{E}\|\nabla F(\mathbf{w}^t)\|^2,$$

we have

$$\bar{G}_T \le \frac{C_0\big(F(\mathbf{w}^0) - F^\star\big)}{\alpha T} + C_1\alpha\sigma_g^2 + \frac{C_2\alpha\sigma_z^2 d}{T}\sum_{t=0}^{T-1}\eta_t^{-2} + \frac{C_3 G^2}{T}\sum_{t=0}^{T-1}\bar{\mathcal{E}}_t(\eta_t). \tag{12}$$

The cumulative zCDP cost of client k is

$$\rho_k^{1:T} = \sum_{t=0}^{T-1} \frac{2\eta_t^2 p_k^2 G^2}{\sigma_z^2}, \tag{13}$$

which converts to (ε_k, δ)-DP via

$$\varepsilon_k = \rho_k^{1:T} + 2\sqrt{\rho_k^{1:T}\ln(1/\delta)}.$$

2.3. Certified Feasibility

Let $\Gamma(\eta)$ denote the right-hand side of (10). For a compact design interval $[\eta_{\min}, \eta_{\max}]$ with $\eta_{\min} > 0$, defining:

$$\mathcal{C}_\rho = \{\eta \in [\eta_{\min}, \eta_{\max}] : \Gamma(\eta) \le \rho\}, \tag{14}$$

$$\mathcal{P}_\varepsilon = \{\eta \in [\eta_{\min}, \eta_{\max}] : \max_k \varepsilon_k(\eta) \le \varepsilon\}, \tag{15}$$

$$\mathcal{Q}_\psi = \{\eta \in [\eta_{\min}, \eta_{\max}] : \bar{\mathcal{A}}(\eta) \le \psi\}, \tag{16}$$

$$\mathcal{U}_\tau = \{\eta \in [\eta_{\min}, \eta_{\max}] : \bar{\mathcal{E}}(\eta) \le \tau\}. \tag{17}$$

The certified feasible set becomes:

$$\mathcal{F} = \mathcal{C}_\rho \cap \mathcal{P}_\varepsilon \cap \mathcal{Q}_\psi \cap \mathcal{U}_\tau. \tag{18}$$

The two participation-side constraints play different roles. $\mathcal{Q}_\psi$ controls envelope-level asymmetry between the effective hardest and easiest clients. $\mathcal{U}_\tau$ controls the total certified dropped weight. The latter is essential because $\bar{\mathcal{A}}(\eta) \to 0$ as $\eta \to \infty$ when all clients are uniformly excluded, not because participation is useful or equitable. Within the design interval, $\mathcal{P}_\varepsilon$ is an interval of the form $[\eta_{\min}, \eta_P]$ when nonempty, because $\varepsilon_k(\eta)$ is increasing. Similarly, $\mathcal{U}_\tau$ is an interval of the form $[\eta_{\min}, \eta_U]$ when nonempty, because $\bar{\mathcal{E}}(\eta)$ is nondecreasing. On the moderate-truncation branch, $\mathcal{Q}_\psi$ is also a lower interval. The convergence-safe set $\mathcal{C}_\rho$ is evaluated by one-dimensional search because Γ is not globally monotone in general.

Corollary 2.3 (Interval form under monotone convergence bound)**.** *If $\Gamma'(\eta) \le 0$ on $[\eta_{\min}, \eta_{\max}]$, i.e.,*

$$C_3 G^2 \bar{\mathcal{E}}'(\eta) \le \frac{2C_2 \alpha \sigma_z^2 d}{\eta^3} \qquad \forall \eta \in [\eta_{\min}, \eta_{\max}],$$

then $\mathcal{C}_\rho$ is either empty or an interval of the form $[\eta_C, \eta_{\max}]$. This condition can hold on the lower part of the low-truncation regime, before the truncation-bias term becomes comparable to the normalized channel-noise term.

Certified Infeasibility: If $\mathcal{F} = \emptyset$, no fixed η in the design range can be certified by Theorem 2.1 to satisfy all four constraints simultaneously. On the moderate-truncation branch, worsening the effective hard-client channel, while keeping the identity of k_- fixed, increases $\bar{\mathcal{A}}(\eta)$ and can shrink $\mathcal{Q}_\psi$, making certified infeasibility more likely. For any fixed target quadruple and design range, the disappearance of $\mathcal{F}$ can be checked numerically. Certified infeasibility is under analytic upper bounds. It is not an information-theoretic impossibility theorem.

Corollary 2.4 (Certified Feasibility)**.** *If $\mathcal{F} \neq \emptyset$, any $\eta \in \mathcal{F}$ ensures that the certified convergence bound in* (10) *is at most ρ, the worst-client privacy loss satisfies $\max_k \varepsilon_k \le \varepsilon$, the envelope asymmetry satisfies $\bar{\mathcal{A}}(\eta) \le \psi$, and the expected dropped-weight envelope satisfies $\bar{\mathcal{E}}(\eta) \le \tau$.*

In the algorithm, we evaluate $\mathcal{F}$ on a finite arm grid and choose the certified arm that minimizes the convergence certificate at its privacy-feasible training horizon. A low-truncation closed-form approximation and additional feasibility details are given in Appendix D.

3. Algorithm

The certified characterization gives a static rule by computing $\mathcal{F}$ and selecting a fixed $\eta \in \mathcal{F}$. Since wireless channels and gradient statistics may vary over time, CARES adapts η_t online while keeping certification separate from adaptation. Certification restricts the available arms using zCDP and participation-envelope constraints; adaptation then chooses among these safe arms using a budget-aware contextual bandit. Hence the bandit never replaces the privacy or participation certificates, and convergence for the realized adaptive sequence follows from Corollary 2.2.

Let $\mathcal{H} = \{\eta_1, \dots, \eta_M\}$ be a finite receive-scaling grid covering $[\eta_{\min}, \eta_{\max}]$.

Certified initialization. The server estimates initial channel statistics $\{\hat{\mu}_k^0\}$ from data-independent pilots and converts the target (ε, δ) to the zCDP budget (Bun & Steinke, 2016)

$$\rho_{\max} = \left(\sqrt{\ln(1/\delta) + \varepsilon} - \sqrt{\ln(1/\delta)}\right)^2.$$

For each arm, it computes

$$T_{\text{eff}}^m = \left\lfloor \frac{\rho_{\max}}{\rho_{\text{inc}}(\eta_m)} \right\rfloor, \qquad \rho_{\text{inc}}(\eta_m) = \max_k \frac{2\eta_m^2 p_k^2 G^2}{\sigma_z^2},$$

and evaluates the convergence certificate at this arm-specific budget horizon. The certified feasible arm set is

$$\begin{aligned}\mathcal{F}_{\mathcal{H}} = \big\{\eta_m \in \mathcal{H} : \Gamma(\eta_m, T_{\text{eff}}^m) \le \rho_{\text{target}},\\ \rho_{\text{inc}}(\eta_m) \le \rho_{\max},\\ \bar{\mathcal{E}}_0(\eta_m) \le \tau, \quad \bar{\mathcal{A}}_0(\eta_m) \le \psi\big\}.\end{aligned}$$

In our experiments we set $\rho_{\text{target}} = 2 \min_m \Gamma(\eta_m, T_{\text{eff}}^m)$. If $\mathcal{F}_{\mathcal{H}} \neq \emptyset$, CARES warm-starts at

$$\eta_{\text{init}} \in \arg\min_{\eta_m \in \mathcal{F}_{\mathcal{H}}} \Gamma(\eta_m, T_{\text{eff}}^m).$$

If $\mathcal{F}_{\mathcal{H}} = \emptyset$, it reports certified infeasibility and may run only in diagnostic best-effort mode, for which no joint certificate is claimed.

Adaptive training. At round t, the server forms the safe arm set

$$\begin{aligned}\mathcal{H}^t = \big\{\eta_m \in \mathcal{H} : \rho_{\text{inc}}(\eta_m) \le \rho_{\text{rem}}^t,\\ \bar{\mathcal{A}}_t(\eta_m) \le \psi, \quad \bar{\mathcal{E}}_t(\eta_m) \le \tau\big\}.\end{aligned} \tag{19}$$

For each safe arm, the context is

$$\begin{aligned}\phi^t(\eta_m) = \big(\eta_m^{-2}, \eta_m^2, e^{-a_t \eta_m^2}, e^{-b_t \eta_m^2},\\ \bar{\mathcal{E}}_t(\eta_m), \rho_{\text{rem}}^t, d_\mu^t\big) \in \mathbb{R}^7,\end{aligned}$$

where

$$a_t = \max_k \frac{p_k^2 G^2}{2(\hat{\mu}_k^t)^2 P_{\max}}, \qquad b_t = \min_k \frac{p_k^2 G^2}{2(\hat{\mu}_k^t)^2 P_{\max}},$$

and $d_\mu^t = \|\hat{\mu}^t - \hat{\mu}^{t-1}\|_2$ captures channel drift. The UCB score is

$$U_t(\eta_m) = \hat{\theta}_m^\top \phi^t(\eta_m) + \alpha_{\text{UCB}} \sqrt{\phi^t(\eta_m)^\top A_m^{-1} \phi^t(\eta_m)}.$$

Rather than using the standard LinUCB score (Li et al., 2010), CARES follows a Bandits-with-Knapsacks ratio rule (Badanidiyuru et al., 2018):

$$\eta^t \in \arg\max_{\eta_m \in \mathcal{H}^t} \frac{U_t(\eta_m)}{\rho_{\mathrm{inc}}(\eta_m)}. \tag{20}$$

This favors arms with high estimated utility per unit privacy cost.

After the OTA update, the contextual model is updated with

$$r^t = \widetilde{\Delta \mathrm{acc}}^{\,t} - \lambda_2 \bar{\mathcal{A}}_t(\eta^t) - \lambda_3 \bar{\mathcal{E}}_t(\eta^t), \tag{21}$$

where $\widetilde{\Delta \mathrm{acc}}^{\,t}$ is a smoothed accuracy gain. The privacy cost is omitted from the reward because it is already accounted for in (20). The remaining budget is updated as

$$\rho_{\mathrm{rem}}^{t+1} = \rho_{\mathrm{rem}}^{t} - \rho_{\mathrm{inc}}(\eta^t).$$

Proposition 3.1 (Budget-aware decision-layer regret)**.** *Under a linear reward approximation and Assumption A.6, the BwK contextual decision layer satisfies*

$$\mathbb{E}[\mathrm{Reg}(T_{\mathrm{budget}})] \le \mathcal{O}\Big(R\sqrt{M d_\phi T_{\mathrm{budget}} \ln T_{\mathrm{budget}}}\Big),$$

where $d_\phi = 7$ and $T_{\mathrm{budget}} = \rho_{\max}/\min_m \rho_{\mathrm{inc}}(\eta_m)$. This is a surrogate decision-layer guarantee; the FL convergence guarantee is given separately by Corollary 2.2.

Algorithm 1 gives the complete pseudocode.

Algorithm 1 CARES

Require: targets $\varepsilon, \delta, \psi, \tau$; arm set $\mathcal{H}$; weights λ_2, λ_3
Estimate $\{\hat{\mu}_k^0\}$ from pilots; compute $\rho_{\max}$
Compute T_{eff}^m and $\Gamma(\eta_m, T_{\mathrm{eff}}^m)$ for all m
Form $\mathcal{F}_{\mathcal{H}}$
if $\mathcal{F}_{\mathcal{H}} \neq \emptyset$ **then**
  $\eta_{\mathrm{init}} \leftarrow \arg\min_{\eta_m \in \mathcal{F}_{\mathcal{H}}} \Gamma(\eta_m, T_{\mathrm{eff}}^m)$
else
  Report certified infeasibility
  $\eta_{\mathrm{init}} \leftarrow \mathrm{median}(\mathcal{H})$ ▷ best-effort mode only
end if
Warm-start contextual model at η_{init}; $\rho_{\mathrm{rem}}^0 \leftarrow \rho_{\max}$
for $t = 0, \ldots, T-1$ **do**
  Update channel estimates $\{\hat{\mu}_k^t\}$
  Form $\mathcal{H}^t$ via (19)
  if $\mathcal{H}^t = \emptyset$ **then**
    Terminate
  end if
  Select η^t via (20)
  Execute OTA-FL round with η^t; compute r^t via (21)
  Update contextual bandit model
  $\rho_{\mathrm{rem}}^{t+1} \leftarrow \rho_{\mathrm{rem}}^t - \rho_{\mathrm{inc}}(\eta^t)$
end for
return $\mathbf{w}^T$

4. Experiments

We evaluate CARES on MNIST digit classification with heterogeneous federated clients and Rayleigh fading channels. Images are compressed to 64 PCA dimensions, and the training set is split across $K = 20$ clients using a Dirichlet distribution with $\alpha_{\mathrm{dir}} = 0.5$. This creates non-IID client data with $p_{\max} \approx 0.09$. We train a two-layer network with 128 hidden units, learning rate 0.1, clipping norm $G = 1$, and channel noise $\sigma_z = 0.05$. The receive-scaling grid is $\mathcal{H} \subset [0.25, 0.40]$ with $M = 10$ arms. The target privacy budget is $(\varepsilon^\star, \delta) = (500, 10^{-5})$ because, in the considered physical regime with $\sigma_z = 0.05$ and $\eta \in [0.25, 0.40]$, each OTA round incurs a non-negligible zCDP cost, so much tighter budgets would terminate training before meaningful convergence. This value is larger than typical digital DP benchmarks because our goal is to compare OTA mechanisms under a common cumulative zCDP budget in a high-dimensional analog aggregation setting, not to claim deployment-level privacy at small ε. We therefore report full accuracy–privacy curves along with $\mathrm{acc@}\varepsilon^\star$. Results are averaged over ten random seeds and reported as mean $\pm$ one standard deviation. Channels include path loss, log-normal shadowing, and fast Rayleigh fading, producing client channel scales that vary by roughly a factor of five. The certified feasible set $\mathcal{F}_{\mathcal{H}}$ is nonempty for all seeds, and CARES static selects $\eta^\star = 0.40$ by minimizing the certified convergence bound over the arm-specific budget horizon.

Baselines and metrics. We compare with FedAvg, FedAvg+DP, OTA-FFL budget, and two CARES variants. FedAvg is used only as a nonprivate upper reference. FedAvg+DP adds Gaussian noise calibrated to the same $(\varepsilon^\star, \delta)$ target using weighted sensitivity $\Delta = 2p_{\max}G$. OTA-FFL budget uses fixed $\eta = 0.293$ and stops before exhausting the zCDP budget. All private methods use the same cumulative privacy accounting. The primary metric is $\mathrm{acc@}\varepsilon^\star$, the interpolated accuracy when cumulative privacy loss reaches $\varepsilon^\star = 500$. We also report best global accuracy, worst-client accuracy at the target, and final cumulative ε.

4.1. MNIST Privacy–Utility Results

Table 1 and Figures 1a–1b summarize the results. FedAvg reaches $93.4 \pm 0.3\%$ accuracy without privacy and serves only as an upper reference. At $\varepsilon^\star = 500$, FedAvg+DP reaches $78.5 \pm 2.0\%$, budget-matched OTA-FFL reaches 77.9%, CARES static reaches 80.6%, and CARES adaptive reaches 78.3%. Thus, CARES static improves over the fixed OTA baseline by 2.7 percentage points at the same privacy budget.

This gain comes from certified receive-scaling selection.

Table 1. MNIST PCA-64 results under privacy target $\varepsilon^\star = 500$, $\delta = 10^{-5}$. Mean over ten random seeds; $\pm$ standard deviation shown where all seeds contribute to that measurement point. FedAvg is a nonprivate upper reference. †OTA-FFL without budget enforcement exceeds $\varepsilon^\star$ and is shown only as a diagnostic.

Method	Best acc.	Acc. at $\varepsilon^\star$	Final ε	Worst-client acc.
FedAvg	0.934 ± 0.003	–	∞	0.925 ± 0.003
FedAvg+DP	0.785 ± 0.020	0.785 ± 0.020	$\leq$500	0.707 ± 0.030
OTA-FFL†	0.819	0.793	593	0.739
OTA-FFL budget	0.819	0.779	$\leq$500	0.716
CARES static	0.824 ± 0.008	**0.806**	$\leq$500	**0.740**
CARES adaptive	0.815	0.783	$\leq$500	0.730

The fixed OTA baseline uses the heuristic value $\eta = 0.293$, while CARES static selects $\eta^\star = 0.40$. The larger scaling reduces effective channel noise and improves per-round learning, but consumes privacy budget faster. CARES therefore uses fewer but higher-quality rounds and achieves better budget-normalized accuracy.

On stationary MNIST channels, the adaptive variant achieves 78.3% at the budget, competitive with the baselines but below the static rule. This is expected as the certified feasible set is stable across rounds, so exploration provides limited gain over committing to $\eta^\star$. The adaptive layer's value appears under nonstationary channels in Appendix H, where it achieves the highest peak accuracy (80.4%) and strongest worst-client accuracy (73.8%).

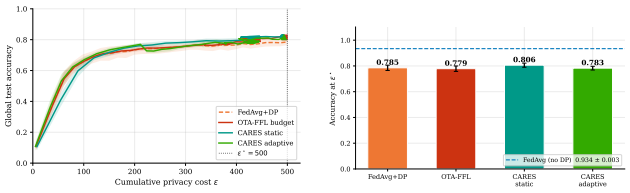

(a) *(b)*

Figure 1. MNIST results. Mean over ten random seeds; $\pm$ standard deviation shown where all seeds contribute to that measurement point. FedAvg is a nonprivate upper reference. †OTA-FFL without budget enforcement exceeds $\varepsilon^\star$.

4.2. Certified Trilemma

Figure 2 shows the certified tradeoff among convergence, privacy, and participation. It evaluates the convergence bound, privacy cost, dropped-weight envelope, and participation-asymmetry envelope as functions of η. The feasible region is nonempty, showing that MNIST admits receive scalings that satisfy all certified constraints.

The tradeoff is controlled by η. Increasing η lowers the normalized OTA noise σ_z/η and improves the convergence certificate, but it also increases the per-round privacy cost since $\rho_{\text{inc}} \propto \eta^2$ and may worsen participation imbalance. Feasible arms are exactly those that remain below all certified thresholds.

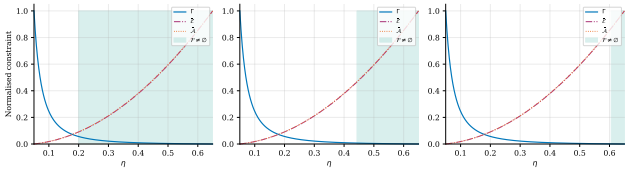

Figure 2. Certified receive-scaling feasibility on MNIST. The normalized convergence certificate Γ, privacy cost ε, dropped-weight envelope $\bar{\mathcal{E}}$, and asymmetry envelope $\bar{\mathcal{A}}$ are plotted as functions of η; the shaded region indicates the certified feasible set $\mathcal{F}$. Increasing σ_z shrinks the feasible set and can make the joint constraints infeasible.

Sweeping $\sigma_z \in \{0.05, 0.11, 0.15\}$ with the convergence target calibrated at $\sigma_z = 0.05$ shows the expected pattern. The feasible set is wide at low noise, narrows at medium noise, and disappears at high noise. In the infeasible case, CARES correctly reports that no certified operating point exists.

Figure 5a shows that the adaptive controller explores early and then concentrates on safe receive-scaling values. All selected arms remain certified feasible.

5. Conclusion

We showed that OTA receive scaling is a shared control knob for convergence, privacy, and participation under truncated channel inversion. The feasible set $\mathcal{F}$ certifies scalings that satisfy zCDP privacy, convergence, and participation-envelope constraints. It therefore serves both as a design rule and as a diagnostic tool. A nonempty $\mathcal{F}$ identifies a certified operating region, while an empty set indicates that the desired guarantees cannot be certified under the current channel and noise conditions.

CARES builds on this certificate by separating safety from adaptation. It selects only certified receive scalings, either statically for a target privacy horizon or adaptively under changing channels. Experiments show gains of 2.7 percentage points on MNIST and 8.9 percentage points on FEMNIST over fixed OTA scaling, with larger benefits on harder, more noise-sensitive tasks. The framework remains conservative by design, since privacy is certified only for the aggregate OTA transcript and participation is bounded through worst-case envelopes. Future work should include tighter certificates, per-client privacy accounting, stronger fairness guarantees, multiple local steps, and MIMO extensions.

References

Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., and Zhang, L. Deep learning with differential privacy. In *Proceedings of the 2016 ACM*

SIGSAC conference on computer and communications security, pp. 308–318, 2016.

Amiri, M. M. and Gündüz, D. Federated learning over wireless fading channels. *IEEE transactions on wireless communications*, 19(5):3546–3557, 2020.

Badanidiyuru, A., Kleinberg, R., and Slivkins, A. Bandits with knapsacks. *Journal of the ACM (JACM)*, 65(3):1–55, 2018.

Bun, M. and Steinke, T. Concentrated differential privacy: Simplifications, extensions, and lower bounds. In *Theory of cryptography conference*, pp. 635–658. Springer, 2016.

Cao, X., Zhu, G., Xu, J., Wang, Z., and Cui, S. Optimized power control design for over-the-air federated edge learning. *IEEE Journal on Selected Areas in Communications*, 40(1):342–358, 2021.

Dwork, C. and Roth, A. The algorithmic foundations of differential privacy. *Foundations and trends® in theoretical computer science*, 9(3-4):211–487, 2014.

Ghadimi, S. and Lan, G. Stochastic first-and zeroth-order methods for nonconvex stochastic programming. *SIAM journal on optimization*, 23(4):2341–2368, 2013.

Hamidi, S. M., Bereyhi, A., Asaad, S., and Poor, H. V. Over-the-air fair federated learning via multi-objective optimization. *IEEE Communications Letters*, 2025.

Kairouz, P. and McMahan, H. B. Advances and open problems in federated learning. *Foundations and trends in machine learning*, 14(1-2):1–210, 2021.

Kalarde, F. M., Liang, B., Dong, M., Ahmed, Y. A. E., and Cheng, H. T. Privacy enhancement in over-the-air federated learning via adaptive receive scaling. *arXiv preprint arXiv:2510.03860*, 2025.

Li, L., Chu, W., Langford, J., and Schapire, R. E. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th international conference on World wide web*, pp. 661–670, 2010.

Li, T., Sanjabi, M., Beirami, A., and Smith, V. Fair resource allocation in federated learning. *arXiv preprint arXiv:1905.10497*, 2019.

Liang, H., Wen, H., Wu, K., and Xing, H. Differential privacy as a perk: Federated learning over multiple-access fading channels with a multi-antenna base station. *arXiv preprint arXiv:2510.23463*, 2025.

Liu, D. and Simeone, O. Privacy for free: Wireless federated learning via uncoded transmission with adaptive power control. *IEEE Journal on Selected Areas in Communications*, 39(1):170–185, 2020.

Liu, Y., Liu, D., Jin, R., Zhu, G., and Shi, Q. Over-the-air federated edge learning with error-feedback one-bit quantization and power control. In *2024 10th International Conference on Computer and Communications (ICCC)*, pp. 2205–2210. IEEE, 2024.

Malekmohammadi, S., Yu, Y., and Cao, Y. Noise-aware algorithm for heterogeneous differentially private federated learning. *arXiv preprint arXiv:2406.03519*, 2024.

McMahan, B., Moore, E., Ramage, D., Hampson, S., and y Arcas, B. A. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, pp. 1273–1282. Pmlr, 2017.

Mohri, M., Sivek, G., and Suresh, A. T. Agnostic federated learning. In *International conference on machine learning*, pp. 4615–4625. PMLR, 2019.

Noble, M., Bellet, A., and Dieuleveut, A. Differentially private federated learning on heterogeneous data. In *International conference on artificial intelligence and statistics*, pp. 10110–10145. PMLR, 2022.

Öksüz, H. Y., Molinari, F., Sprekeler, H., and Raisch, J. Boosting fairness and robustness in over-the-air federated learning. *IEEE Control Systems Letters*, 8:682–687, 2024.

Wei, X., Wang, T., Huang, R., Shen, C., Yang, J., and Poor, H. V. Differentially private wireless federated learning using orthogonal sequences. *IEEE Transactions on Information Theory*, 2026.

Yan, N., Wang, K., Pan, C., Chai, K. K., Shu, F., and Wang, J. Over-the-air federated averaging with limited power and privacy budgets. *IEEE Transactions on Communications*, 72(4):1998–2013, 2023.

Yang, H., Qiu, P., Liu, J., and Yener, A. Over-the-air federated learning with joint adaptive computation and power control. In *2022 IEEE International Symposium on Information Theory (ISIT)*, pp. 1259–1264. IEEE, 2022.

Yang, K., Jiang, T., Shi, Y., and Ding, Z. Federated learning via over-the-air computation. *IEEE transactions on wireless communications*, 19(3):2022–2035, 2020.

Zhu, G., Wang, Y., and Huang, K. Broadband analog aggregation for low-latency federated edge learning. *IEEE transactions on wireless communications*, 19(1):491–506, 2019.

A. Model Details and Assumptions

This section collects the modeling assumptions and records the exact OTA estimator decomposition used in the analysis.

The ideal FedAvg stochastic aggregate is $\hat{\mathbf{g}}^t = \sum_{k=1}^K p_k \mathbf{g}_k^t$. Under strict truncated channel inversion,

$$b_k^t = \begin{cases} \eta p_k / h_k^t, & h_k^t \geq h_{\min}^t(k), \\ 0, & h_k^t < h_{\min}^t(k), \end{cases} \qquad h_{\min}^t(k) = \frac{\eta p_k \|\mathbf{g}_k^t\|}{\sqrt{P_{\max}}}.$$

Thus an active client contributes exactly $p_k \mathbf{g}_k^t$ after normalization, while an inactive client contributes zero. Therefore,

$$\tilde{\mathbf{g}}^t = \frac{\mathbf{y}^t}{\eta} = \sum_{k \in \mathcal{K}^t} p_k \mathbf{g}_k^t + \frac{1}{\eta}\mathbf{z}^t = \hat{\mathbf{g}}^t + \mathbf{e}^t + \frac{1}{\eta}\mathbf{z}^t,$$

where $\mathbf{e}^t = -\sum_{k \notin \mathcal{K}^t} p_k \mathbf{g}_k^t$. The truncation bias $\mathbf{e}^t$ is not generally zero mean because the activity event depends on $\|\mathbf{g}_k^t\|$ through the activation threshold.

A.1. Assumptions

Assumption A.1 (L-smoothness)**.** Each local objective F_k is L-smooth. Hence $F = \sum_k p_k F_k$ is also L-smooth.

Assumption A.2 (Stochastic gradients)**.** For all k, t,

$$\mathbb{E}[\mathbf{g}_k^t \mid \mathbf{w}^t] = \nabla F_k(\mathbf{w}^t), \qquad \mathbb{E}[\left\|\mathbf{g}_k^t - \nabla F_k(\mathbf{w}^t)\right\|^2 \mid \mathbf{w}^t] \leq \sigma_g^2.$$

Client stochastic-gradient noises are conditionally independent given $\mathbf{w}^t$.

Assumption A.3 (Clipping)**.** All transmitted gradients are clipped so that $\|\mathbf{g}_k^t\| \leq G$.

Assumption A.4 (Rayleigh fading)**.** The channel magnitudes satisfy $h_k^t \sim \text{Rayleigh}(\mu_k)$ and are independent of the current stochastic-gradient norm conditioned on $\mathbf{w}^t$. The long-term parameters $\{\mu_k\}$ are estimated from data-independent pilots and treated as public side information.

Assumption A.5 (Privacy transcript)**.** The DP guarantee is for the aggregate transcript

$$\mathcal{T} = \{(\mathbf{y}^t, \eta_t)\}_{t=0}^{T-1}$$

and public system parameters. It excludes data-dependent active-set indicators, per-client received powers, and per-round client success/failure indicators.

Assumption A.6 (Bandit reward model)**.** For the decision-layer regret statement only, the surrogate reward is assumed to admit a linear contextual approximation with conditionally zero-mean R-subgaussian noise.

Remark A.7 (Scope of privacy). If activity patterns or per-client received powers are observable, then participation itself may leak information because activity depends on the clipped gradient norm. Such full-transcript privacy would require additional mechanisms such as activity masking, dummy transmissions, or anonymous access.

B. Participation Envelopes

Lemma B.1 (Truncation and participation envelopes)**.** *Under Assumptions A.3 and A.4,*

$$q_k^{\text{true}}(\eta) = \Pr\left(h_k^t < \frac{\eta p_k \|\mathbf{g}_k^t\|}{\sqrt{P_{\max}}}\right) \leq \bar{q}_k(\eta) = 1 - \exp\left(-\frac{\eta^2 p_k^2 G^2}{2\mu_k^2 P_{\max}}\right).$$

Consequently, for

$$S^t = \sum_{k \notin \mathcal{K}^t} p_k, \qquad \bar{\mathcal{E}}(\eta) = \sum_{k=1}^K p_k \bar{q}_k(\eta),$$

we have $\mathbb{E}[S^t] \leq \bar{\mathcal{E}}(\eta)$. *Moreover, if*

$$c_k = \frac{p_k^2 G^2}{2\mu_k^2 P_{\max}}, \qquad a = \max_k c_k, \qquad b = \min_k c_k,$$

then $\bar{\mathcal{A}}(\eta) = e^{-b\eta^2} - e^{-a\eta^2}$ *is nonnegative and unimodal when* $a > b$*, with peak*

$$\eta_{\text{peak}} = \sqrt{\frac{\ln(a/b)}{a-b}}.$$

Finally, the truncation bias satisfies $\mathbb{E}\left\|\mathbf{e}^t\right\|^2 \le G^2\bar{\mathcal{E}}(\eta)$.

Proof. The Rayleigh CDF gives $\Pr(h_k^t < x) = 1 - \exp(-x^2/(2\mu_k^2))$. Substituting the activation threshold and using $\|\mathbf{g}_k^t\| \le G$ yields the bound on q_k^{true}. The dropped-weight envelope follows by linearity of expectation. For the bias term,

$$\left\|\mathbf{e}^t\right\| \le \sum_{k\notin\mathcal{K}^t} p_k \left\|\mathbf{g}_k^t\right\| \le GS^t,$$

and since $S^t \in [0,1]$, $\left\|\mathbf{e}^t\right\|^2 \le G^2 S^t$. Taking expectations gives the claim. The properties of $\bar{\mathcal{A}}$ follow by differentiating $e^{-b\eta^2} - e^{-a\eta^2}$ and solving $ae^{-a\eta^2} = be^{-b\eta^2}$. □

Remark B.2 (Why $\bar{\mathcal{E}}$ is needed). The asymmetry envelope $\bar{\mathcal{A}}(\eta)$ goes to zero as $\eta \to \infty$ because all clients become nearly inactive. A small $\bar{\mathcal{A}}$ alone therefore does not imply useful participation. The dropped-weight constraint $\bar{\mathcal{E}}(\eta) \le \tau$ rules out this degenerate regime.

C. Proofs of the Main Characterization

C.1. Convergence

The proof follows the standard smooth nonconvex SGD descent argument (Ghadimi & Lan, 2013). The only additional terms are the truncation bias $\mathbf{e}^t$ and normalized channel noise $\mathbf{z}^t/\eta$.

By Assumption A.2,

$$\mathbb{E}[\hat{\mathbf{g}}^t \mid \mathbf{w}^t] = \nabla F(\mathbf{w}^t), \qquad \mathbb{E}\left\|\hat{\mathbf{g}}^t - \nabla F(\mathbf{w}^t)\right\|^2 \le \sigma_g^2.$$

Using L-smoothness and $\tilde{\mathbf{g}}^t = \hat{\mathbf{g}}^t + \mathbf{e}^t + \mathbf{z}^t/\eta$,

$$\mathbb{E}[F(\mathbf{w}^{t+1})] \le \mathbb{E}[F(\mathbf{w}^t)] - \alpha\mathbb{E}\langle\nabla F(\mathbf{w}^t), \tilde{\mathbf{g}}^t\rangle + \frac{L\alpha^2}{2}\mathbb{E}\left\|\tilde{\mathbf{g}}^t\right\|^2.$$

The ideal aggregate gives the usual descent term, the channel noise is zero mean, and the truncation cross-term is controlled by Young's inequality:

$$-\alpha\langle\nabla F, \mathbf{e}^t\rangle \le \frac{\alpha}{2}\|\nabla F\|^2 + \frac{\alpha}{2}\left\|\mathbf{e}^t\right\|^2.$$

Moreover,

$$\mathbb{E}\left\|\tilde{\mathbf{g}}^t\right\|^2 \le 3\mathbb{E}\left\|\nabla F(\mathbf{w}^t)\right\|^2 + 3\sigma_g^2 + 3G^2\bar{\mathcal{E}}(\eta) + 3\frac{\sigma_z^2 d}{\eta^2},$$

where the truncation term follows from

$$\left\|\mathbf{e}^t\right\| \le G\sum_{k\notin\mathcal{K}^t} p_k = GS^t, \qquad (S^t)^2 \le S^t, \qquad \mathbb{E}[S^t] \le \bar{\mathcal{E}}(\eta).$$

Substituting these bounds, summing over $t = 0, \dots, T-1$, using $F(\mathbf{w}^T) \ge F^\star$, and applying $\alpha \le 1/(12L)$ gives

$$\frac{1}{T}\sum_{t=0}^{T-1}\mathbb{E}\left\|\nabla F(\mathbf{w}^t)\right\|^2 \le \frac{4[F(\mathbf{w}^0) - F^\star]}{\alpha T} + 6L\alpha\sigma_g^2 + 6L\alpha\frac{\sigma_z^2 d}{\eta^2} + (2 + 6L\alpha)G^2\bar{\mathcal{E}}(\eta).$$

Thus the proof gives the sharper constants

$$C_0 = 4, \qquad C_1 = C_2 = 6L, \qquad C_3 = 2 + 6\alpha L.$$

The main text uses the looser constants

$$C_1 = C_2 = 12L, \qquad C_3 = 2 + 12\alpha L,$$

which simplify presentation while preserving the same dependence on η, σ_z^2, d, and $\bar{\mathcal{E}}(\eta)$.

C.2. Adaptive Receive Scaling

For a predictable sequence $\{\eta_t\}_{t=0}^{T-1}$, the same one-step descent argument applies round by round with η replaced by η_t and with time-varying channel envelopes $\bar{\mathcal{E}}_t$. Summing gives

$$\frac{1}{T}\sum_{t=0}^{T-1}\mathbb{E}\left\|\nabla F(\mathbf{w}^t)\right\|^2 \le \frac{C_0[F(\mathbf{w}^0)-F^\star]}{\alpha T} + C_1\alpha\sigma_g^2 + \frac{C_2\alpha\sigma_z^2 d}{T}\sum_{t=0}^{T-1}\frac{1}{\eta_t^2} + \frac{C_3G^2}{T}\sum_{t=0}^{T-1}\bar{\mathcal{E}}_t(\eta_t).$$

This proves Corollary 2.2; the privacy part follows from adaptive zCDP composition in Appendix C.3.

C.3. Privacy

The aggregate observation in round t is a Gaussian mechanism with noise covariance $\sigma_z^2 I_d$. For client k, changing one sample can change the clipped transmitted update by at most $2G$, so the aggregate sensitivity is conservatively bounded by

$$\Delta_{k,t} \le 2\eta_t p_k G.$$

Because participation is data-dependent, we do not discount silent rounds in the formal accounting. The Gaussian mechanism is $\rho_{k,t}$-zCDP with

$$\rho_{k,t} = \frac{\Delta_{k,t}^2}{2\sigma_z^2} \le \frac{2\eta_t^2 p_k^2 G^2}{\sigma_z^2}.$$

By adaptive composition of zCDP (Bun & Steinke, 2016),

$$\rho_k^{1:T} = \sum_{t=0}^{T-1}\frac{2\eta_t^2 p_k^2 G^2}{\sigma_z^2}.$$

The zCDP-to-approximate-DP conversion gives

$$\varepsilon_k = \rho_k^{1:T} + 2\sqrt{\rho_k^{1:T}\ln(1/\delta)}.$$

For fixed $\eta_t = \eta$, this reduces to (11) in the main text. Solving $\varepsilon = \rho + 2\sqrt{\rho\ln(1/\delta)}$ gives the zCDP budget

$$\rho_{\max} = \left(\sqrt{\ln(1/\delta)+\varepsilon} - \sqrt{\ln(1/\delta)}\right)^2.$$

C.4. Privacy Scope and Conservative Accounting

The differential privacy guarantee in Theorem 2.1 is stated for the aggregate OTA transcript $\mathcal{T} = \{(\mathbf{y}^t, \eta_t)\}_{t=0}^{T-1}$ together with public system parameters. Long-term channel statistics $\{\mu_k\}$ are assumed to be estimated from data-independent pilots and are treated as public side information. The transcript does not include data-dependent activity indicators, per-round client success or failure indicators, or per-client received powers during gradient transmission. If such information is observable, participation itself may leak data-dependent signals because activation depends on $\|\mathbf{g}_k^t\|$. In that case, full-transcript privacy would require additional protections such as activity masking, anonymous access, or dummy transmissions.

Our accounting is conservative. Since participation is data-dependent, we do not discount silent rounds in the formal DP bound. At every round, client k is charged the worst-case clipped-gradient sensitivity

$$\Delta_{k,t} = 2\eta_t p_k G.$$

This bound holds independently of the mini-batch sampling rule because the transmitted clipped gradient lies in the radius-G ball. Hence two neighboring datasets can change the transmitted message by at most $2G$, and the corresponding received aggregate by at most $2\eta_t p_k G$. Tighter accounting based on subsampling amplification or participation-aware privacy analysis may reduce this cost, but is outside the scope of the present certificate.

D. Feasibility and Low-Truncation Approximation

Let

$$B(\eta) = C_2\alpha\frac{\sigma_z^2 d}{\eta^2} + C_3 G^2 \bar{\mathcal{E}}(\eta)$$

be the η-dependent part of the convergence certificate. The first term decreases with η, while the second increases because $\bar{\mathcal{E}}(\eta)$ is nondecreasing. Hence $B(\eta)$ is not assumed to be globally unimodal, and the convergence-safe set is evaluated by one-dimensional grid search.

In the low-truncation regime, $c_k\eta^2 \ll 1$, so

$$\bar{\mathcal{E}}(\eta) = \sum_k p_k(1 - e^{-c_k\eta^2}) \approx K_{\mathrm{LT}}\eta^2, \qquad K_{\mathrm{LT}} = \frac{G^2}{2P_{\max}}\sum_k \frac{p_k^3}{\mu_k^2}.$$

Then

$$B_{\mathrm{LT}}(\eta) = \frac{A}{\eta^2} + D\eta^2, \qquad A = C_2\alpha\sigma_z^2 d, \qquad D = C_3 G^2 K_{\mathrm{LT}},$$

and the closed-form minimizer is

$$\eta_{\mathrm{LT}}^\star = \left(\frac{A}{D}\right)^{1/4} = \left(\frac{C_2\alpha\sigma_z^2 d}{C_3 G^2 K_{\mathrm{LT}}}\right)^{1/4}.$$

This approximation is used only for intuition and diagnostics; the certified feasible set in the experiments is computed using the exact envelopes over the arm grid.

Proposition D.1 (Bound-Optimal Receive Scaling)**.** *If $\mathcal{F} \neq \emptyset$, then $\mathcal{F}$ is compact and Γ is continuous on $\mathcal{F}$. Therefore,*

$$\eta^\star \in \arg\min_{\eta\in\mathcal{F}} \Gamma(\eta)$$

exists. In practice, we compute this minimizer by one-dimensional search over the feasible arm grid.

In the low-truncation regime, $\bar{\mathcal{E}}(\eta) \approx K_{\mathrm{LT}}\eta^2$, where

$$K_{\mathrm{LT}} = \frac{G^2}{2P_{\max}}\sum_k \frac{p_k^3}{\mu_k^2}.$$

The η-dependent part of the convergence certificate is then

$$B_{\mathrm{LT}}(\eta) = C_2\alpha\frac{\sigma_z^2 d}{\eta^2} + C_3 G^2 K_{\mathrm{LT}}\eta^2,$$

with minimizer

$$\eta_{\mathrm{LT}}^\star = \left(\frac{C_2\alpha\sigma_z^2 d}{C_3 G^2 K_{\mathrm{LT}}}\right)^{1/4}.$$

When $\mathcal{F}$ is an interval, the feasible low-truncation design is the projection of $\eta_{\mathrm{LT}}^\star$ onto $\mathcal{F}$; otherwise it is the minimizer of $B_{\mathrm{LT}}(\eta)$ over $\mathcal{F}$.

E. Algorithmic Details

E.1. BwK Contextual Bandit Implementation

At round t, the current channel estimates define

$$a_t = \max_k \frac{p_k^2 G^2}{2\hat{\mu}_k^{t\,2} P_{\max}}, \qquad b_t = \min_k \frac{p_k^2 G^2}{2\hat{\mu}_k^{t\,2} P_{\max}}.$$

For each arm $\eta_m \in \mathcal{H}^t$ the contextual feature vector is

$$\phi^t(\eta_m) = \left(\eta_m^{-2},\, \eta_m^2,\, e^{-a_t\eta_m^2},\, e^{-b_t\eta_m^2},\, \bar{\mathcal{E}}_t(\eta_m),\, \rho_{\mathrm{rem}}^t,\, d_\mu^t\right) \in \mathbb{R}^7,$$

where $d_\mu^t = \|\hat{\mu}^t - \hat{\mu}^{t-1}\|_2$. The first two features capture channel-noise amplification and privacy-cost scaling; the exponential features encode the participation envelopes through the channel statistics; $\bar{\mathcal{E}}_t$ captures total dropped-weight risk; ρ_{rem}^t exposes the remaining budget; and d_μ^t captures channel drift.

Each arm m maintains a Gram matrix $A_m \in \mathbb{R}^{7\times 7}$ initialized to I_7 and a vector $\mathbf{b}_m \in \mathbb{R}^7$ initialized to $\mathbf{0}$. The ridge-regression estimate and UCB bound are

$$\hat{\theta}_m = A_m^{-1}\mathbf{b}_m, \qquad U_t(\eta_m) = \hat{\theta}_m^\top \phi^t + \alpha_{\text{UCB}}\sqrt{\phi^{t\top} A_m^{-1}\phi^t}.$$

After observing reward r^t for the selected arm $m^\star$, the model is updated by

$$A_{m^\star} \leftarrow A_{m^\star} + \phi^t\phi^{t\top}, \qquad \mathbf{b}_{m^\star} \leftarrow \mathbf{b}_{m^\star} + \tilde{r}^t\phi^t,$$

where $\tilde{r}^t = r^t - \bar{r}$ is the reward centered by a slow exponential moving average $\bar{r}$ for numerical stability.

For warm-starting, the arm m_{init} corresponding to η_{init} is given a synthetic prior observation:

$$A_{m_{\text{init}}} \leftarrow A_{m_{\text{init}}} + \phi^0\phi^{0\top}, \qquad \mathbf{b}_{m_{\text{init}}} \leftarrow \mathbf{b}_{m_{\text{init}}} + r_{\text{prior}}\phi^0,$$

biasing early selection toward the certified operating point without preventing exploration of other safe arms.

Reward smoothing. Raw per-round accuracy differences fluctuate by several percentage points due to channel noise, particularly in high-dimensional or many-class settings. The reward signal is smoothed using an exponential moving average of the accuracy baseline with coefficient $\alpha_{\text{ema}} = 0.3$. The smoothed delta is computed as $\widetilde{\Delta\text{acc}}^t = \text{acc}^t - \bar{\text{acc}}^t$, where $\bar{\text{acc}}^t = (1-\alpha_{\text{ema}})\bar{\text{acc}}^{t-1} + \alpha_{\text{ema}}\text{acc}^t$. This reduces reward variance by roughly a factor of $1/(2-\alpha_{\text{ema}}) \approx 0.59$ relative to the raw difference, meaningfully improving the bandit signal-to-noise ratio in noisy operating regimes.

E.2. Budget-Aware Surrogate Regret Analysis

The regret statement concerns only the adaptive decision layer and is not a convergence theorem for the FL objective. Convergence for the selected adaptive sequence is certified by Corollary 2.2. Assume for each arm m that the reward admits the linear approximation $\mathbb{E}[r^t \mid \eta_m, \phi^t] = \theta_m^{\star\top}\phi^t$ and that noise is conditionally zero-mean and R-subgaussian (Assumption A.6). Because each arm has a known deterministic cost $\rho_{\text{inc}}(\eta_m)$, the decision layer compares safe arms by reward per unit privacy budget. The BwK-style ratio policy satisfies the surrogate regret bound

$$\mathbb{E}[\text{Reg}(T_{\text{budget}})] \le \mathcal{O}\big(R\sqrt{M d_\phi T_{\text{budget}} \ln T_{\text{budget}}}\big), \quad d_\phi = 7.$$

The safe-arm restriction ensures all selected arms satisfy the current zCDP and participation-envelope constraints; global convergence and privacy guarantees are supplied separately by the certified analysis.

E.3. T_{eff}-Consistent Feasibility Computation

A subtlety arises in computing the certified feasible set when arms have different per-round costs. Each arm η_m can afford at most $T_{\text{eff}}^m = \lfloor \rho_{\max}/\rho_{\text{inc}}(\eta_m) \rfloor$ training rounds within the privacy budget. The convergence certificate $\Gamma(\eta_m, T_{\text{eff}}^m)$ therefore depends on the arm through both its noise amplification and its effective horizon. Using a common global horizon T for all arms in the feasibility check is incorrect: it underestimates Γ for arms that exhaust the budget quickly (large η, small T_{eff}), producing an overly optimistic feasibility assessment. The reference target ρ_{target} must therefore also be computed using the arm-wise T_{eff}^m, as done in Algorithm 1. This correction is critical for high-noise or highly heterogeneous settings where $T_{\text{eff}}^m \ll T$ for some arms.

E.4. Adaptive Convergence-Budget Filter

If a practitioner requires a hard adaptive convergence target maintained online, Corollary 2.2 can be converted into a convergence-budget filter. Define the per-round certified communication and truncation cost

$$B_t(\eta) = C_2\alpha\frac{\sigma_z^2 d}{\eta^2} + C_3 G^2 \bar{\mathcal{E}}_t(\eta).$$

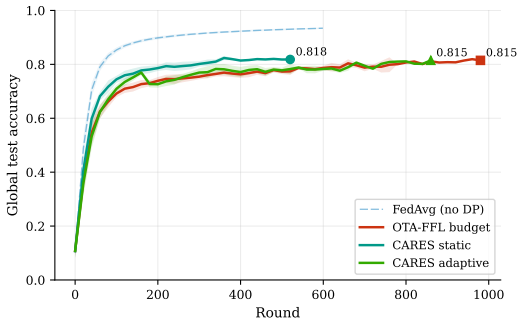

Figure 3. OTA training curves under the shared privacy budget as a function of communication round. Methods stop at different rounds because different η values consume different amounts of zCDP per round. The nonprivate FedAvg curve is an upper reference.

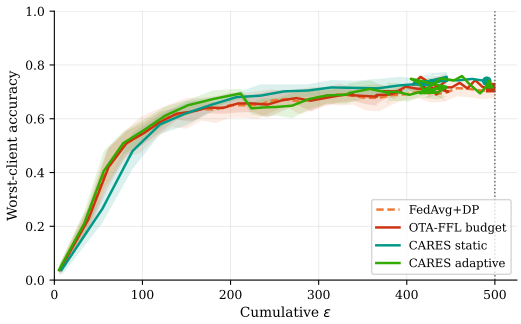

Figure 4. Worst-client accuracy versus cumulative privacy loss. CARES static achieves the highest worst-client accuracy at $\varepsilon^\star$ among private methods, indicating that certified arm selection also benefits participation fairness.

The remaining convergence budget evolves as $R^{t+1}_{\mathrm{conv}} = R^t_{\mathrm{conv}} - B_t(\eta^t)$, where $R^0_{\mathrm{conv}} = T(\rho - C_0F_0/(\alpha T) - C_1\alpha\sigma_g^2)$ for a target convergence level ρ. A look-ahead filter replaces $\mathcal{H}^t$ in (19) by

$$\mathcal{H}^t_{\mathrm{conv}} = \left\{\eta_m \in \mathcal{H}^t_{\mathrm{pp}} : B_t(\eta_m) + (T-t-1)B^t_{\min} \le R^t_{\mathrm{conv}}\right\},$$

where $B^t_{\min} = \min_{\eta\in\mathcal{H}^t_{\mathrm{pp}}} B_t(\eta)$. This ensures that after selecting η_m at round t, sufficient convergence budget remains to complete the remaining rounds at the minimum certified per-round cost. This filter is conservative and optional; the main algorithm certifies the realized sequence post-hoc via Corollary 2.2.

F. Experimental Protocol

All methods share the settings described in Section 4. FedAvg+DP uses weighted sensitivity $\Delta = 2p_{\max}G$ matching the worst-case OTA sensitivity; standard sensitivity $\Delta = 2G$ would inflate noise by $1/p_{\max} \approx 11$ and produce a misleadingly weak baseline. MNIST results are averaged over ten seeds; FEMNIST results over two seeds. The channel follows path-loss exponent two with 4 dB log-normal shadowing and fast Rayleigh fading, giving per-client scale parameters $\hat{\mu}_k$ with a roughly five-fold near-to-far range.

G. Additional Experimental Diagnostics

Figures 3–5b collect the main diagnostic results. All private OTA methods terminate before exceeding $\varepsilon^\star$, confirming that the budget-tracking mechanism operates correctly.

Motivation. Figure 6 compares FedAvg, FedAvg+DP, and OTA-FFL without budget enforcement across training rounds. OTA-FFL reaches 79.3% accuracy after 600 rounds but accumulates $\varepsilon \approx 525$, exceeding the target, while FedAvg+DP reaches only $\sim 78.5\%$ under the same budget. This motivates explicit budget enforcement and certified accounting.

H. Nonstationary Channel Stress Test

We evaluate CARES adaptive under nonstationary channels to identify the setting where adaptation provides the clearest benefit. Channel scale parameters $\hat{\mu}_k$ drift multiplicatively by up to $\pm 25\%$ every 40 rounds, simulating a scenario in which clients change their physical location during training.

Table 2 reports results over three random seeds. At $\varepsilon^\star = 500$, CARES static achieves 0.794 ± 0.012, outperforming budget-matched OTA-FFL at 0.785 ± 0.012. CARES adaptive achieves the highest peak accuracy at 0.804 ± 0.017 and the strongest worst-client accuracy at 0.738 ± 0.018, demonstrating that the adaptive controller responds to channel changes in a way that benefits fairness. However, its accuracy at the fixed privacy budget is 0.764 ± 0.009, below the static variant.

Figure 7 illustrates the dynamics. The left panel shows that all three methods track closely in accuracy versus privacy cost, with CARES adaptive slightly below the static rule at the budget boundary. The right panel shows that the adaptive

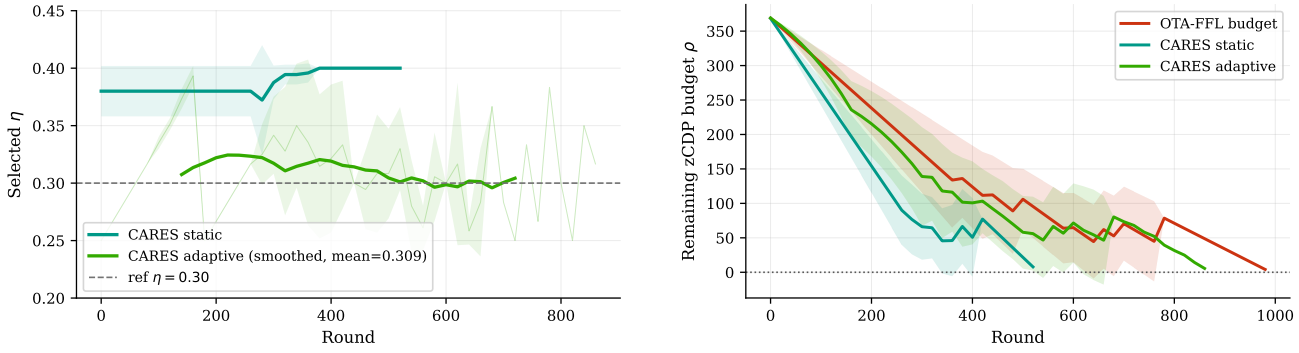

Figure 5. a) Receive-scaling trajectory. The static rule holds at the certified $\eta^\star$ until budget exhaustion. The adaptive rule explores within the certified safe set, with a smoothed mean near $\eta = 0.31$ across all seeds. b) Remaining zCDP budget over training. CARES static exhausts the budget faster because $\eta^\star = 0.40$ consumes more zCDP per round than the OTA-FFL baseline at $\eta = 0.293$, but the higher per-round quality more than compensates.

Table 2. Nonstationary channel results on MNIST PCA-64 over three seeds, $\varepsilon^\star = 500$, $\delta = 10^{-5}$.

Method	Best acc.	Acc. at $\varepsilon^\star$	Worst-client	Final ε
OTA-FFL budget	0.793 ± 0.005	0.785 ± 0.012	0.736 ± 0.003	495.4 ± 4.7
CARES static	0.798 ± 0.009	0.794 ± 0.012	0.726 ± 0.006	486.1 ± 2.8
CARES adaptive	$\mathbf{0.804 \pm 0.017}$	0.764 ± 0.009	$\mathbf{0.738 \pm 0.018}$	491.0 ± 6.7

controller varies η_t over time in response to changing channel statistics, while the static variant remains fixed. The variation in η_t is meaningful because when channels degrade, the controller reduces η_t to lower the per-round privacy cost and extend the training horizon, which benefits worst-client accuracy.

The results should not be interpreted as evidence that adaptation is universally inferior at fixed privacy budgets. Rather, the stationary and mildly nonstationary MNIST setting lacks the strong channel heterogeneity needed for the adaptive advantage to dominate over exploration cost. In settings with larger and more abrupt channel changes, such as mobile devices with significant position changes between training rounds, the adaptive controller is expected to provide a clearer improvement.

I. FEMNIST Supplementary Results

We also evaluate on FEMNIST, a harder 62-class benchmark with digits and lowercase and uppercase characters. We use the EMNIST byclass split compressed to 128 PCA dimensions, with $K = 100$ clients and Dirichlet concentration $\alpha_{\text{dir}} = 0.3$, producing a highly non-IID partition. We set $\sigma_z = 0.03$ because the weaker 62-class gradient signal makes $\sigma_z = 0.05$

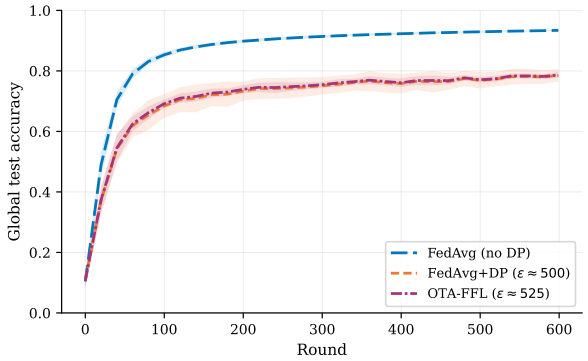

Figure 6. Motivation: implicit OTA noise versus explicit DP noise. OTA aggregation achieves better utility than FedAvg+DP but requires explicit zCDP accounting to respect the privacy budget.

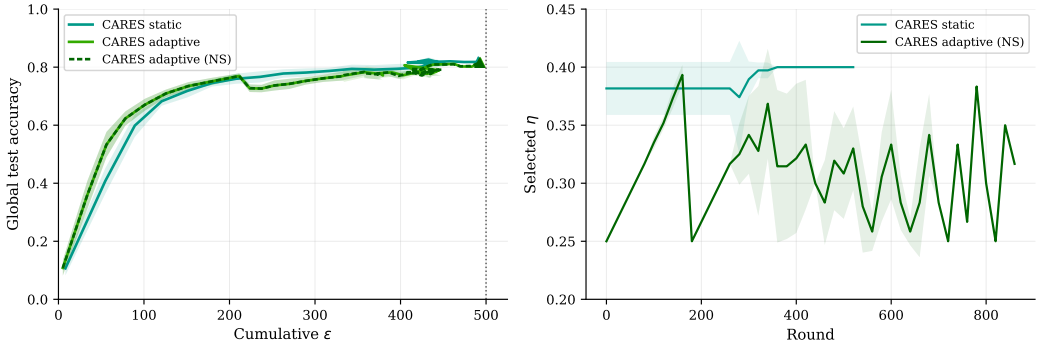

Figure 7. Nonstationary channel experiment. Left: accuracy versus cumulative privacy loss (mean ± one std. over three seeds). Right: receive-scaling trajectory, showing that the adaptive controller changes η_t in response to channel drift while the static variant remains at the certified initialization.

Table 3. FEMNIST results under $\varepsilon^\star = 500$, $\delta = 10^{-5}$, averaged over two seeds. OTA-FFL† runs without budget enforcement and terminates after T_{ref} rounds with $\varepsilon_{\mathrm{final}} \approx 54$, never reaching the target budget; it is shown as a diagnostic of the uncontrolled noise.

Method	Best acc.	Acc. at $\varepsilon^\star$	Worst-client	Final ε
FedAvg	0.586	–	–	∞
FedAvg+DP	0.552	0.552	0.362	≤ 500
OTA-FFL†	0.325	–	187	54
OTA-FFL budget	0.400	0.393	0.268	≤ 500
CARES static	**0.488**	**0.482**	**0.347**	≤ 500
CARES adaptive	0.392	0.390	0.263	≤ 500

overly noisy for OTA training. This value is held fixed across all FEMNIST methods, so the comparison remains fair; the change reflects the task's higher noise sensitivity rather than algorithm-specific tuning. All other hyperparameters follow the MNIST protocol.

Table 3 reports. FedAvg reaches 58.6% without privacy, while FedAvg+DP reaches 55.2% at $\varepsilon^\star = 500$ due to the small weighted sensitivity from $p_{\max} \approx 0.022$. OTA-FFL without budget enforcement reaches 32.5% because it consumes the privacy budget quickly and stops early, while budget-matched OTA-FFL reaches 39.3%. CARES static achieves $\mathbf{48.2}\%$, improving over budget-matched OTA-FFL by 8.9 percentage points and over unbudgeted OTA-FFL by 15.7 percentage points. CARES adaptive reaches 39.0%, again below the static rule on stationary channels.

The larger FEMNIST gain shows that receive-scaling selection matters more when classification is more noise-sensitive. On MNIST, certified arms have similar utility per privacy cost. On FEMNIST, minimizing $\Gamma(\eta, T_{\mathrm{eff}})$ selects a scaling with a much better convergence rate per unit privacy, and the benefit accumulates over training. These results support the claims of the main paper. The certified arm selection in CARES consistently outperforms a heuristic fixed scaling across both MNIST and FEMNIST, and the improvement is larger when the task is harder and the noise sensitivity is higher.

J. Limitations and Scope

The framework is deliberately conservative. The privacy guarantee applies only to the aggregate OTA transcript, not to richer physical-layer observations such as activity indicators or per-client received powers. If these signals are exposed, additional privacy protection is required. The truncation envelopes are worst-case bounds based on gradient clipping. They require no distributional assumptions on the gradients, but may overestimate realized client dropping; in our MNIST experiments, actual participation loss is often much smaller. The participation asymmetry envelope is a diagnostic proxy, not a direct fairness guarantee. It captures unequal truncation risk across clients but does not bound worst-client loss or accuracy, so empirical fairness metrics should also be reported. Finally, the convergence analysis assumes one local update per round. Multiple local steps would add client-drift terms. The adaptive component should therefore be viewed as a decision layer within the certified safe set.